POWERING THE NEW IT GENERATION

June 2015 | Vol. 8 No. 2

VIRTUALIZATION REVIEW

VirtualizationReview.com

SPECIAL DEEP DIVE EDITION

KEEPING UP with the CHANGING DATACENTER

It's the era of cloud, virtualization and BYOD—falling behind is not an option.

THIS SPECIAL DEEP DIVE EDITION IS SPONSORED BY

WITH ADDITIONAL CONTENT PROVIDED BY VEEAM

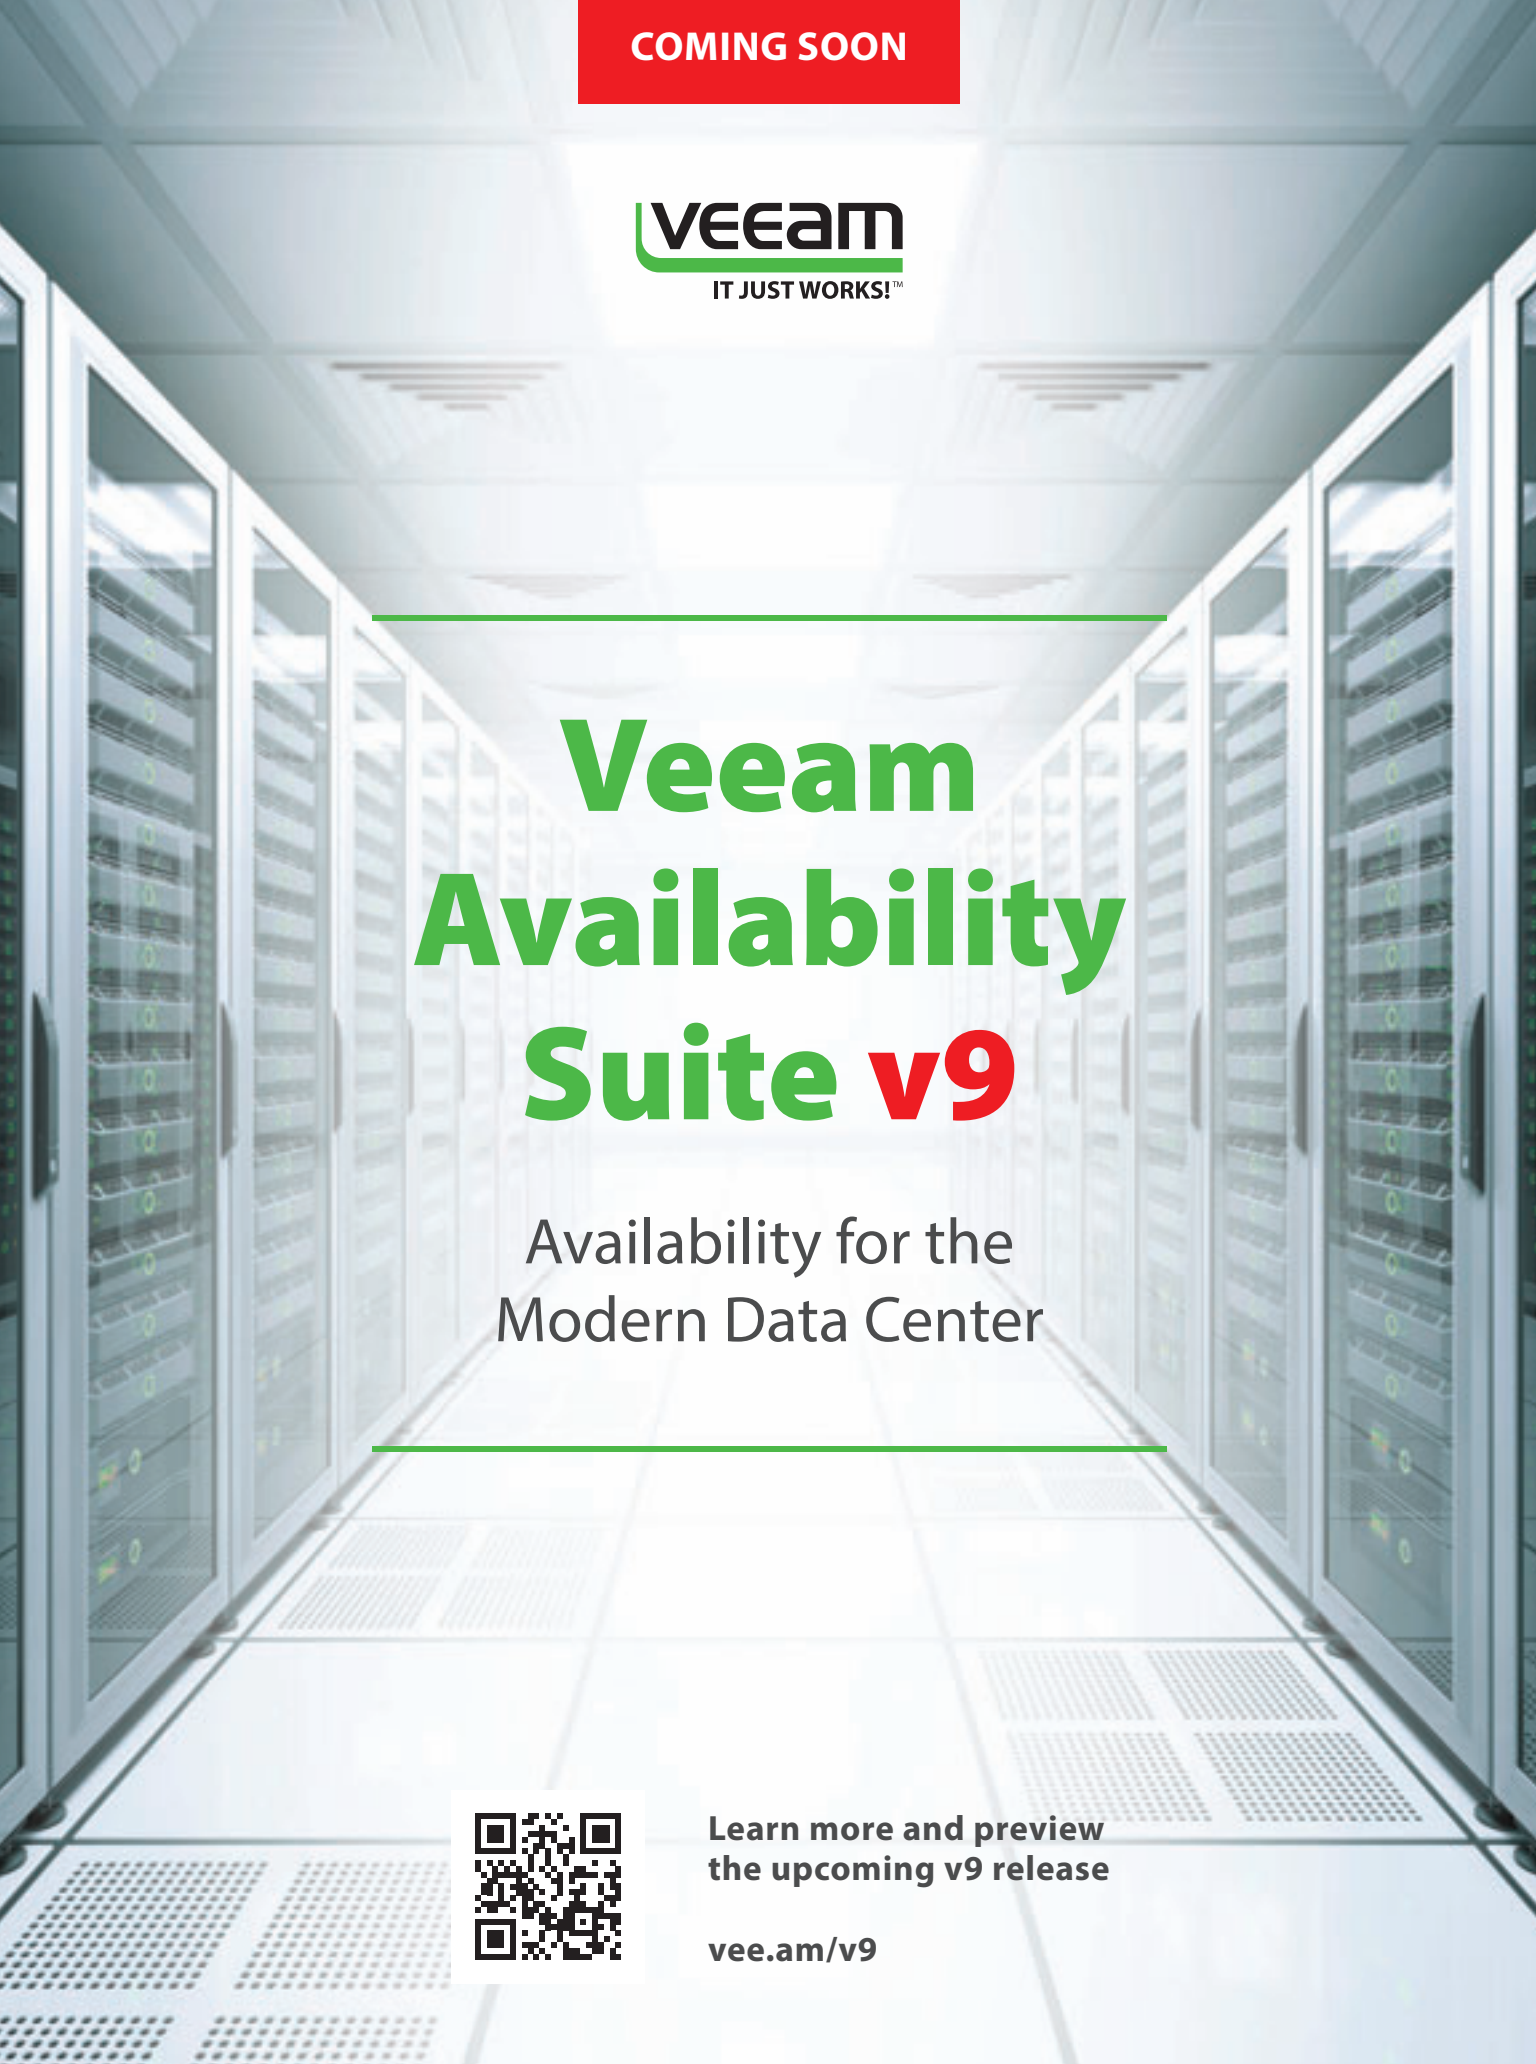
COMING SOON
VEEAM
IT JUST WORKS!™
Veeam Availability Suite v9
Availability for the Modern Data Center
Learn more and preview the upcoming v9 release
vee.am/v9

contents

June 2015 | VIRTUALIZATION REVIEW | VOL. 8, NO. 2

VIRTUALIZATION REVIEW

ED NOTE

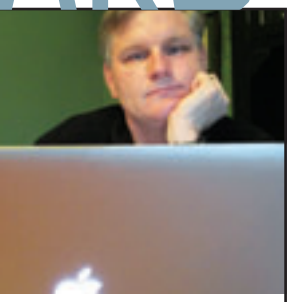

KEITH WARD

Are You Keeping up with Your Changing Datacenter?

Chances are near 100 percent that your datacenter doesn't look the same today as it did five years ago. And chances are just as high that your datacenter won't look the same five years from now as it does today. That saying, about the only constant being change? It's true; and nowhere is it more true than in your datacenter.

The change started with virtualization. Suddenly, one server no longer had to run just one app, wasting most of its compute capacity.

That led to the next phase, possibly the most disruptive of all: cloud computing. The cloud combines virtualization of all sorts—storage, compute, networking—with the ability to break through the traditional barrier of your own datacenter.

The cloud has upended the industry. Most enterprises have moved to a hybrid cloud system, with a combination of public and private clouds for things like provisioning and scalability. Apps now exist on desktops, laptops, tablets, phones, cars, thermostats … the list goes on from here to eternity.

Some, though, have stubbornly refused to adapt. For them, things work well, if slow. Dev is dev and ops is ops, and never the twain shall meet. Silos for data and management are good things. If this is you, then this issue has your name on it. In these pages, you'll find out about how things are changing, and how that change is mostly good, if scary. You'll learn about the different environments and how to get the most out of them.

If, on the other hand, you're gradually moving in the cloud/virtualization direction, this magazine will help you get there more quickly and smoothly.

And if you're a gung-ho early adopter, you'll be excited by some of the new ideas you'll find in these pages. You may even get some helpful advice.

The one thing you cannot do anymore is sit still, ignoring the changes going on in the datacenter. Closing your eyes and sticking your fingers in your ears is no longer an option; you're costing your company time and money if you do. Start here, and get moving. VR

POWERING THE NEW IT GENERATION

VIRTUALIZATION REVIEW

ID STATEMENT *Virtualization Review* magazine (ISSN 1941-2843) is published four times a year (July, August/September, October and December/January) by 1105 Media, Inc., 9201 Oakdale Avenue, Ste. 101, Chatsworth, CA 91311. Periodicals postage paid at Chatsworth, CA 91311-9998, and at additional mailing offices. Complimentary subscriptions are sent to qualifying subscribers. Annual subscription rates payable in U.S. funds for non-qualified subscribers are: U.S. $39.95, International $64.95.

Subscription inquiries, back issue requests, and address changes: Mail to: *Virtualization Review*, 9201 Oakdale Avenue, Suite 101, Chatsworth, CA 91311, E-mail privacy@1105media.com or call 818-814-5223, fax number 818-936-0267.

POSTMASTER: Send address changes to *Virtualization Review*, 9201 Oakdale Avenue, Suite 101, Chatsworth, CA 91311. Canada Publications Mail Agreement No: 40612608. Return Undeliverable Canadian Addresses to 9121 Oakdale Avenue, Suite 101, Chatsworth, CA 91311.

CORPORATE ADDRESS 1105 Media, 9201 Oakdale Ave. Ste 101, Chatsworth, CA 91311
www.1105media.com

MEDIA KITS Direct your Media Kit requests to Chief Revenue Officer Dan LaBianca, 972-687-6702 (phone), 972-687-6799 (fax), dlabianca@1105media.com

REPRINTS For single article reprints (in minimum quantities of 250-500), e-prints, plaques and posters contact: PARS International Phone: 212-221-9595. E-mail: 1105reprints@parsintl.com. www.magreprints.com/QuickQuote.asp

LIST RENTAL This publication's subscriber list, as well as other lists from 1105 Media, Inc., is available for rental. For more information, please contact our list manager, Jane Long, Merit Direct. Phone: 913-685-1301; E-mail: jlong@meritdirect.com; Web: www.meritdirect.com/1105

Reaching the Staff

Staff may be reached via e-mail, telephone, fax, or mail. A list of editors and contact information is also available online at VirtualizationReview.com.

E-mail: To e-mail any member of the staff, please use the following form: FirstinitialLastname@1105media.com

Irvine Office (weekdays, 9:00 a.m. – 5:00 p.m. PT)
Telephone 949-265-1520; Fax 949-265-1528
4 Venture, Suite 150, Irvine, CA 92618

Corporate Office (weekdays, 8:30 a.m. – 5:30 p.m. PT)
Telephone 818-814-5200; Fax 818-734-1522
9201 Oakdale Avenue, Suite 101, Chatsworth, CA 91311

The opinions expressed within the articles and other contentsherein do not necessarily express those of the publisher.

NEW

VEEAM
IT JUST WORKS!™

Veeam Management Pack v8 for System Center

Get app-to-metal visibility
of your critical infrastructure!

Learn more and preview
the v8 release

vee.am/mp

By James Green

AS MUCH AS RUNNING AN enterprise-class datacenter for the fun of it would be a dream job for some folks, that position doesn't exist, and funding a datacenter with Monopoly money hasn't happened yet. As it turns out, a datacenter is either the support system or the lifeblood of a business. In some cases, the services provided by the datacenter are internal-only and provide messaging services, file storage and host applications. Other times, the datacenter hosts customer-facing workloads like a Software-as-a-Service (SaaS) platform or an E-commerce site. In either example, it's vital to the organization that the datacenter systems stay online and available.

As one of the largest shifts in datacenter architecture in the past decade, virtualization has become the de facto method of provisioning server infrastructure in most enterprises. With the ability to drive efficiency and scale by packing more workloads into the same physical footprint, server virtualization is an economical advantage. With this advantage, however, comes one big risk: By the nature of this architecture, all the proverbial eggs are in one basket.

Bill Laing, corporate vice president of the Server and Cloud Division at Microsoft, famously wrote in a blog post, "The three truths of cloud computing are: hardware fails, software has bugs and people make mistakes." As such, designing a virtualized server infrastructure must take into account these risks to uptime. The system (vSphere, in the case of this article) must be designed with failure in mind.

Because failure is a guarantee, what can be done to minimize the impact on production workloads?

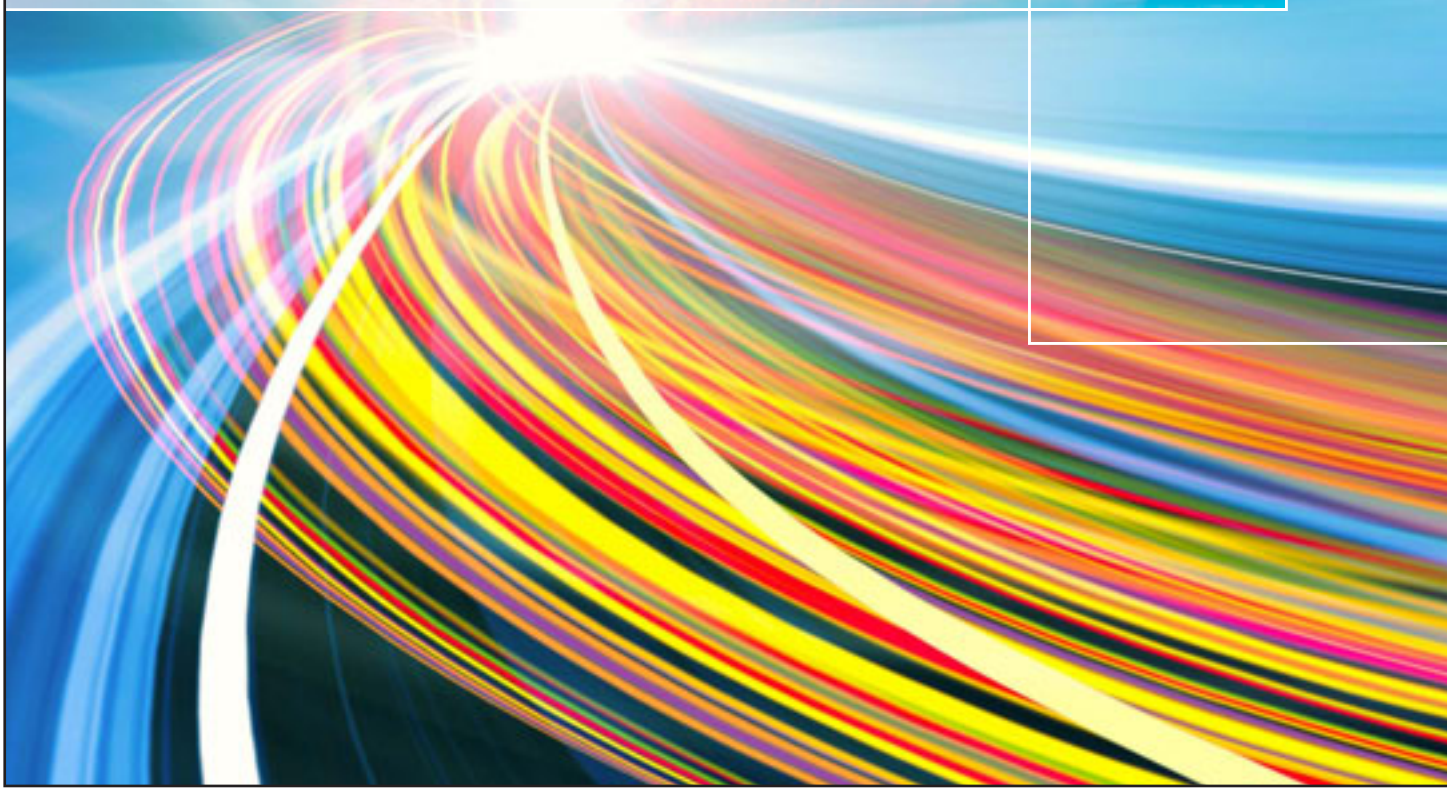

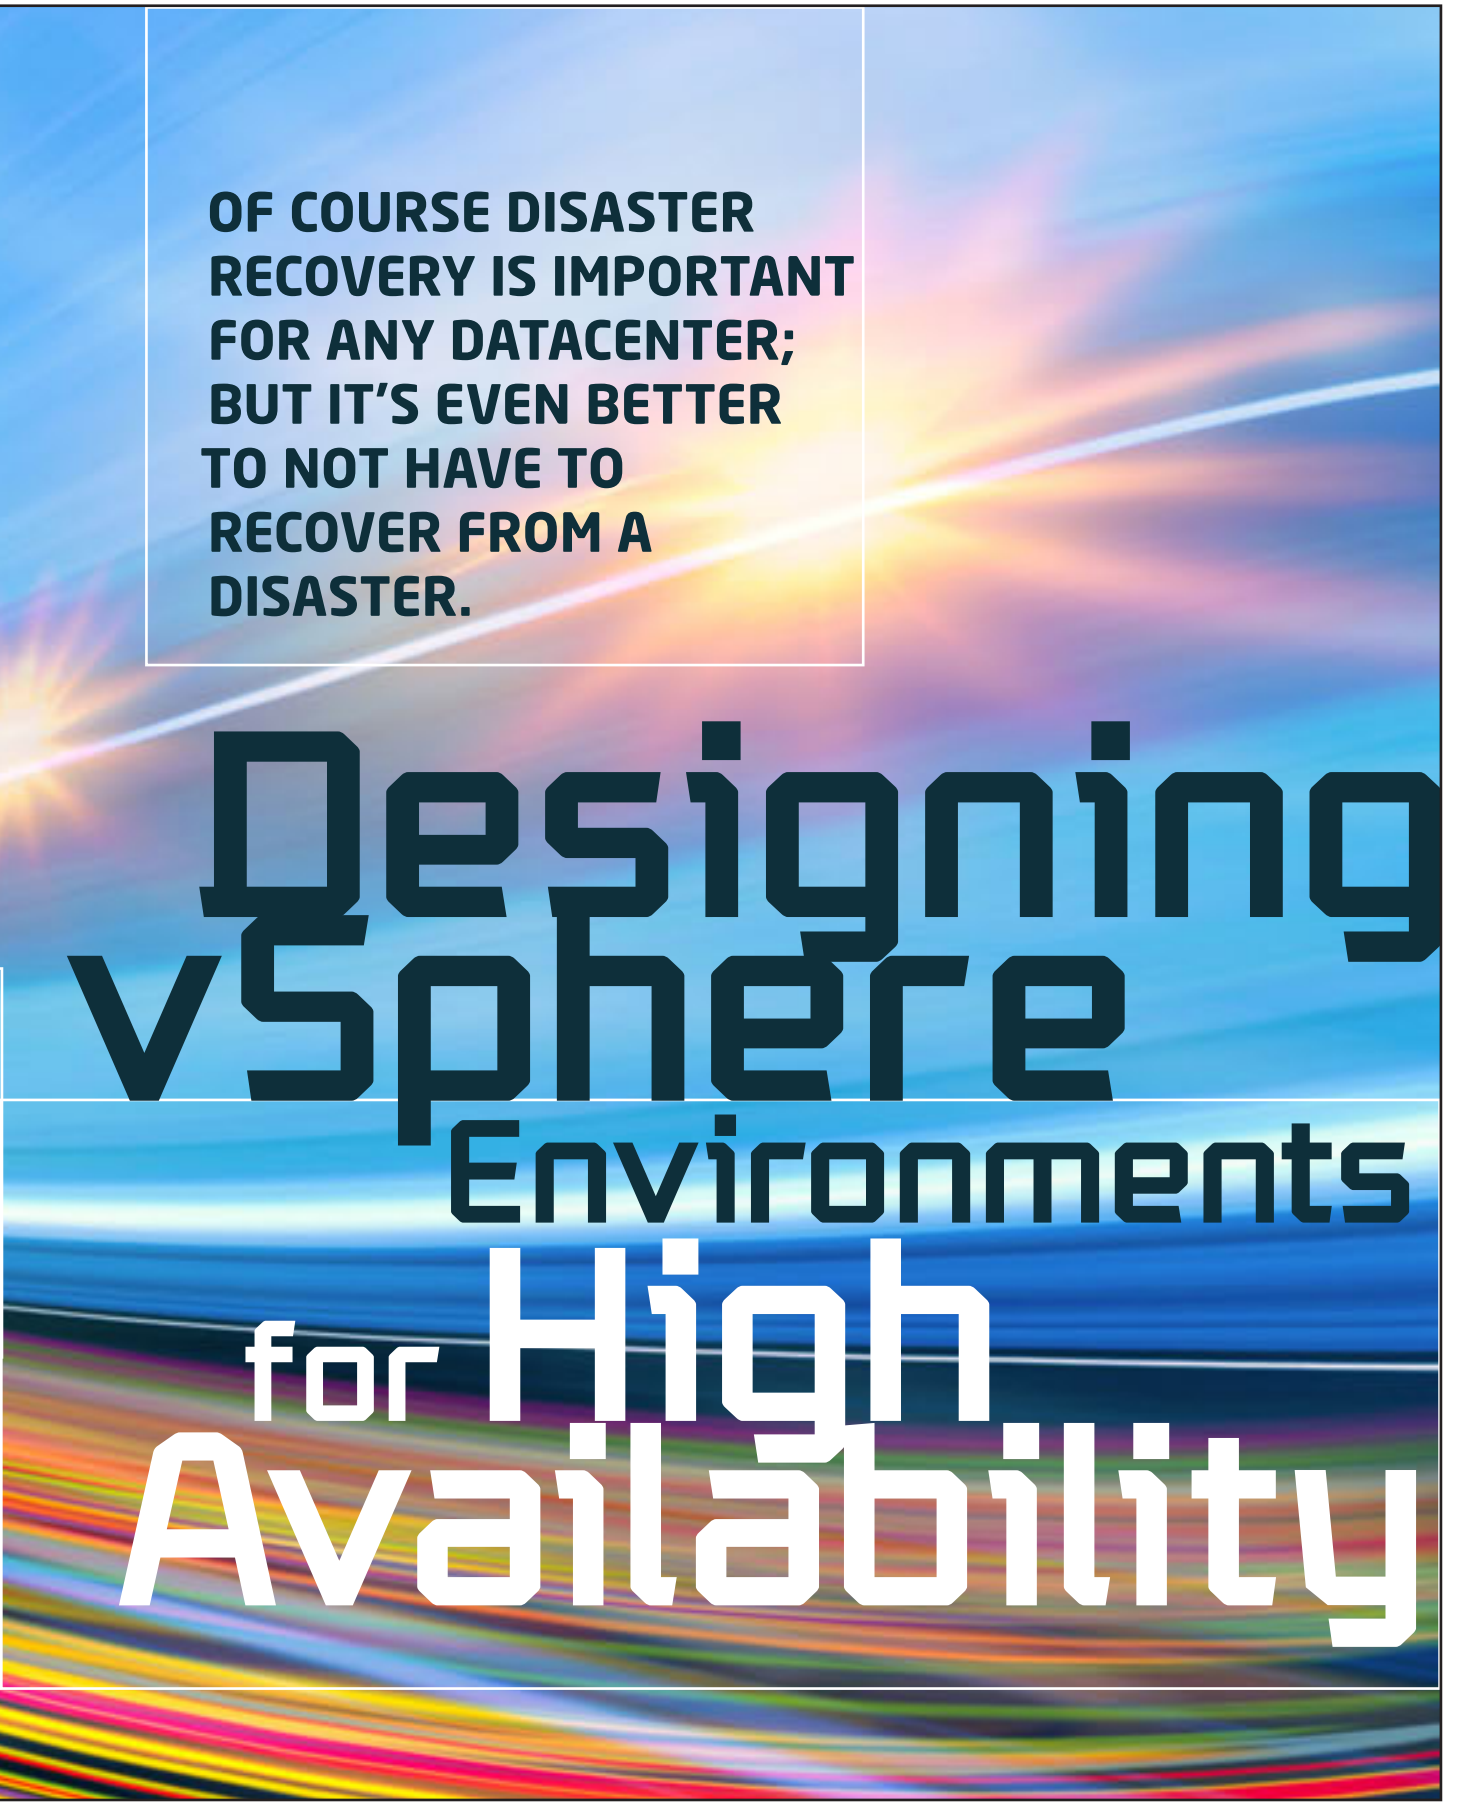

OF COURSE DISASTER RECOVERY IS IMPORTANT FOR ANY DATACENTER; BUT IT'S EVEN BETTER TO NOT HAVE TO RECOVER FROM A DISASTER.

Designing vSphere Environments for High Availability

Use a vSphere HA Cluster

Using a vSphere HA cluster (shown in **Figure 1**) is paramount to running a production environment on vSphere. Possibly the most important reason to cluster ESXi hosts, High Availability (HA) allows the platform to intelligently restart workloads when a failure has occurred. The impact of this design decision—HA vs. no HA—is enormous; HA can be the difference between five minutes of downtime and five hours.

Host Monitoring

The primary function of HA is to monitor for ESXi host failures in the cluster. If a host fails, either by losing power, experiencing a hardware malfunction, or losing access to network or storage resources, HA will automatically (within about five minutes) begin to re-launch those unavailable virtual machines (VMs) on another host.

Because rebooting VMs due to a small network blip would cause more harm than good, HA has a few mechanisms to verify whether failure has actually occurred.

Isolation Response Address: The first mechanism HA uses for failure verification is called the Isolation Response Address. When a host is participating in an HA cluster, it constantly attempts to ping this address to determine whether it still has network connectivity.

Out of the box, HA uses the default gateway of the management vmkernel interface as the address to test. A VMware Inc. recommended best practice that isn't configured automatically is to set up an alternate isolation response address. In the event that the default gateway of vmko is unreachable, HA will continue down the list of up to 10 other addresses to test. The Advanced Option to specify this alternate address is das.isolationaddressX, where X is a number from zero to nine.

Datastore Heartbeating: In conjunction with testing the availability of the Isolation Response Address, HA also uses a mechanism called Datastore Heartbeat to verify whether a failure has occurred. This is less useful in environments where many types of traffic share the same physical networking; but if the storage network is segregated, datastore heartbeats could help show that a host is up despite the isolation address being unreachable. An accidental HA failover would be avoided in this case.

Datastore Heartbeat selects two datastores based on the cluster configuration, and creates a folder on each called .vSphere-HA. The heartbeat files in this folder are updated on a regular basis by each host in the cluster. Capacity utilization isn't a concern, as this folder only consumes around 3MB on VMFS-5.

Admission Control

Because the point of HA is to be able to make VMs available again in the event of a failure, it would make little sense if precautions weren't taken to ensure that resources would be available to start those VMs once a failure occurs.

This is the purpose of Admission Control. In plain terms, Admission Control (when enabled and configured properly) is HA's guarantee that any running VMs in the cluster can be powered on on a surviving host. Without Admission Control, it's the Wild West when it comes to an HA failover, and only the strong survive the chaos.

Admission Control should be enabled and subsequently configured in every production environment. Not doing so is a major risk. Because different levels of resilience are required in different environments, the way Admission Control reserves spare capacity is configurable. The following three Admission Control Policy options are available:

1. Host Failures Cluster Tolerates: This is the default option. Interestingly, it's also probably the most misunderstood and incorrectly configured. First, while this setting is the default, it's not good enough as configured out of the box. This policy uses a calculation called "slot size" to determine how many VMs can run on a given host. When tuned properly, this is fairly accurate, but without tuning it's woefully unhelpful. The assumption when an administrator hasn't

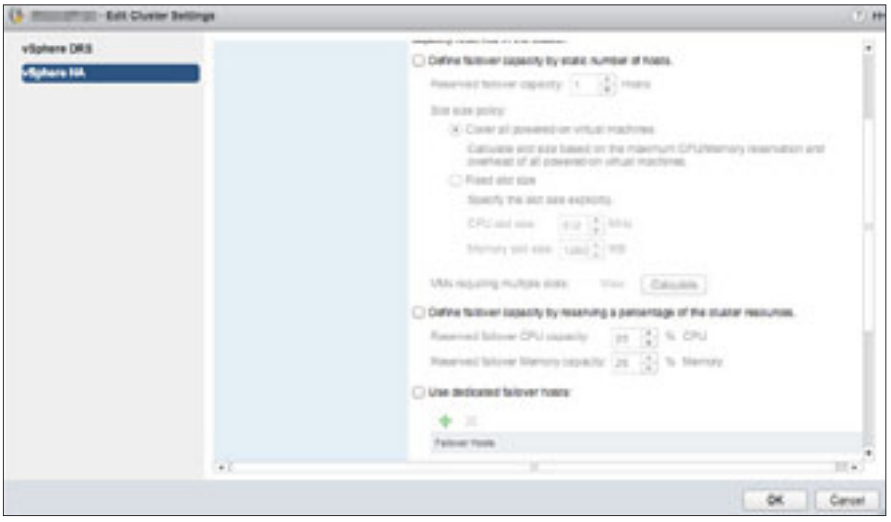

Figure 1. Creating a vSphere HA cluster.

configured otherwise is that each VM will use 32MHz CPU. In most cases, this is too small a value, and the calculations are skewed to allow more slots than should actually be available. From the Web client, an administrator should edit the cluster settings and manually specify slot size attributes that would accurately reflect the environment.

2. Percentage of Cluster Resources: This is a bit easier to understand and calculate than the default policy. Essentially, the percentage specified to be reserved (CPU and Memory can be specified independently) will be withheld from the aggregate of resources in the cluster. The proper setting for this would be the percentage of cluster resources that one host contributes (or a multiple of that). As an example, a cluster containing six hosts should be set to reserve 17 percent, which is just slightly more than one-sixth. If the cluster needed to tolerate two failures, then 34 percent should be reserved (two-sixths).

The benefit of this policy is that it's less complex than the slot size configurations. The drawback, however, is that it isn't dynamic. If a host is added to the cluster and this setting isn't updated, cluster resources won't be properly reserved by Admission Control.

3. Specify Failover Hosts: The final policy reserving cluster resources is similar to the idea of a hot spare in a RAID array. A specified host is online and idle, waiting to take over for a host with a failure. While this policy does have use cases, it typically isn't used due to the fact that it's more wasteful. Unless all other hosts will be run up to 100 percent utilization, this policy will actually cost more than a host's worth of resources.

VM Monitoring

VM Monitoring is an underutilized feature of vSphere HA, which is known primarily for handling host failures. VM Monitoring uses VMware Tools heartbeats, and observes network and storage IO to determine whether the guest OS in a VM is available. If a VM doesn't have any network or storage IO for a given period of time, and no heartbeats are being received from VMware Tools, HA can restart the VM. A typical example of this would be the Windows "Blue Screen of Death," in which VM Monitoring would properly diagnose this condition and restart the VM without needing administrator intervention. Although this won't likely fix the underlying issue, it does contribute to application uptime by getting the machine back up and running as quickly as possible.

Avoid Single Points of Failure

All of this HA failover business is assuming that a failure has actually occurred. When designing for vSphere environments, substantial consideration should be given to how failure of individual components can be sustained.

While catastrophic failure of an ESXi host is a possibility, much more likely is that a single network interface controller (NIC) will fail, or an intern will configure the wrong switchport, or a switch in the storage fabric will fail. While HA allows for minimal downtime (VMs begin restarting within five minutes), good resiliency in the physical and logical design will allow for no downtime when a single, redundant component fails.

vSwitch Uplinks

Because many VMs live on one piece of physical equipment, they all depend on the same few physical uplinks to get network traffic to the outside world. If 50 virtual servers all rely on four physical interfaces, a failure here could disrupt many workloads. Avoiding single points of failure in networking components is the key to staying online.

Physical NICs: Having multiple physical interfaces is critical to a successful vSphere implementation; not only for performance reasons, but for ensuring that VM traffic can still flow in the event of a failure on a physical link. An often-overlooked consideration is that although a vSwitch may have two uplinks, if both NICs are on the same PCI card, there's still a single point of failure. It's best when designing the vSwitch configuration to ensure it has uplinks from two different PCI cards, as shown in **Figure 2**.

Physical Switches: In the same way, if a vSwitch has two uplinks but they both run to the same physical switch, there is a single point of failure. A resilient vSwitch will contain uplinks that run to separate physical switches in a stack, or to separate line cards in a chassis-based switch.

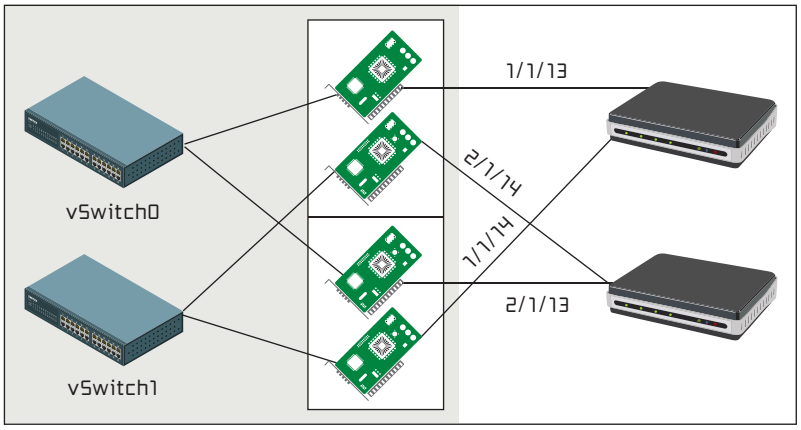

Figure 2. High availability is more likely when a vSwitch connects to two separate network interface controllers.

Storage Connectivity

On the storage side, considerations are exactly the same. It does little good to provide what looks like multiple paths to a LUN if they actually all terminate in one place. When designing vSphere storage connectivity, be sure to run two completely independent storage fabrics whenever possible.

Physical HBAs: It's quite common for Fibre Channel host bust adapters (FC HBAs) to contain two ports. As mentioned previously, the intention is not to provide extra throughput or to provide redundancy at the port level; it's to provide HA by having one run to each fabric and provide a storage path completely independent of the other.

Develop an Availability/Recovery Strategy

Despite lots of planning and hard work, disasters still happen. Components still fail and people make big enough mistakes that redundancy or HA isn't enough to protect against a failure. For the organization to keep functioning, a plan must be in place to provide capability to roll back changes, recover lost data, and resume business operations from a secondary location. This protection scheme should be multiple layers deep so that the response can be appropriate for the magnitude of the disaster.

Snapshots

A relatively simple but effective way to ensure destructive changes can be rolled back is to regularly perform snapshots. Any modern storage array will have the ability to take a point-in-time copy of metadata and store it for later access.

From a vSphere perspective, this means if machines are modified or corrupted, an admin can recover the working VM from an earlier snapshot in very short order. The ease with which snapshots can quickly correct a derailed situation makes this technology a must.

Backups

Sometimes, being able to recover from a snapshot is either not granular enough or snapshot data doesn't go far enough back in time. For this reason, a backup solution must also be in place. Many an organization has found out the hard way that snapshots are not backups. To be fully protected, the two technologies should be used in tandem.

Replication

Although it's a rare occurrence, site-level disasters also need to be accounted for when the datacenter being online is critical to the business. HA from a site perspective means having a recent (relative to recovery point objective) copy of data off-site at another facility. Depending on the pre-defined recovery time objective, this data may need to be available immediately, or in very short order. Once replicated data can be accessed, systems can be restored at the alternate location and business can resume.

Asynchronous: There are two different modes of replication, and which one is a good fit depends entirely on the criticality of the workload being replicated.

The most common replication mode is called asynchronous, which means that local site operations proceed as normal, and then in the background the changes are replicated to the remote site. This is most common because it's the most attainable from a cost perspective.

The downside is that an asynchronous replication strategy has a "lag time." Depending on the rate of change and the speed of the connection between sites, this lag time can mean seconds' to hours' worth of data lost in the event of a disaster at the local site.

Synchronous: The preferred mode of replication is much more expensive, due to the components needed to make it possible. Synchronous replication actively writes changes to both the local and the remote systems. This means there's effectively no loss of data in the event of a disaster at the local site.

As ideal as this sounds, it's quite complicated to deploy and manage, and is only worth the cost to larger organizations. In many cases, implementing a system that includes synchronous array replication would cost much more than the cost of an outage.

The 3-Step HA Plan

Designing a vSphere environment for HA is no easy task, but these three steps will make a huge difference:

1. Make use of vSphere HA. This is one of the primary purposes for clustering ESXi in the first place. Take care to tune the HA settings to be appropriate for the situation; remember the default settings aren't good enough.
2. When developing a vSphere design, carefully avoid single points of failure. Be sure to consider every component, and whether that component is actually made up of smaller components. Try to add redundancy at the most granular level possible.
3. Last, develop an availability and recovery plan that makes handling a failure possible, and make sure the strategy has multiple layers. Different recovery layers should be able to address different magnitudes of failure. Remember Bill Laing's truths: hardware fails, software has bugs and people make mistakes. Plan for them! VR

James Green is an independent blogger at virtadmin.com, a two-time vExpert, a serial Tech Field Day delegate, and works as a virtualization consultant/architect. Follow him on Twitter @jdgreen.

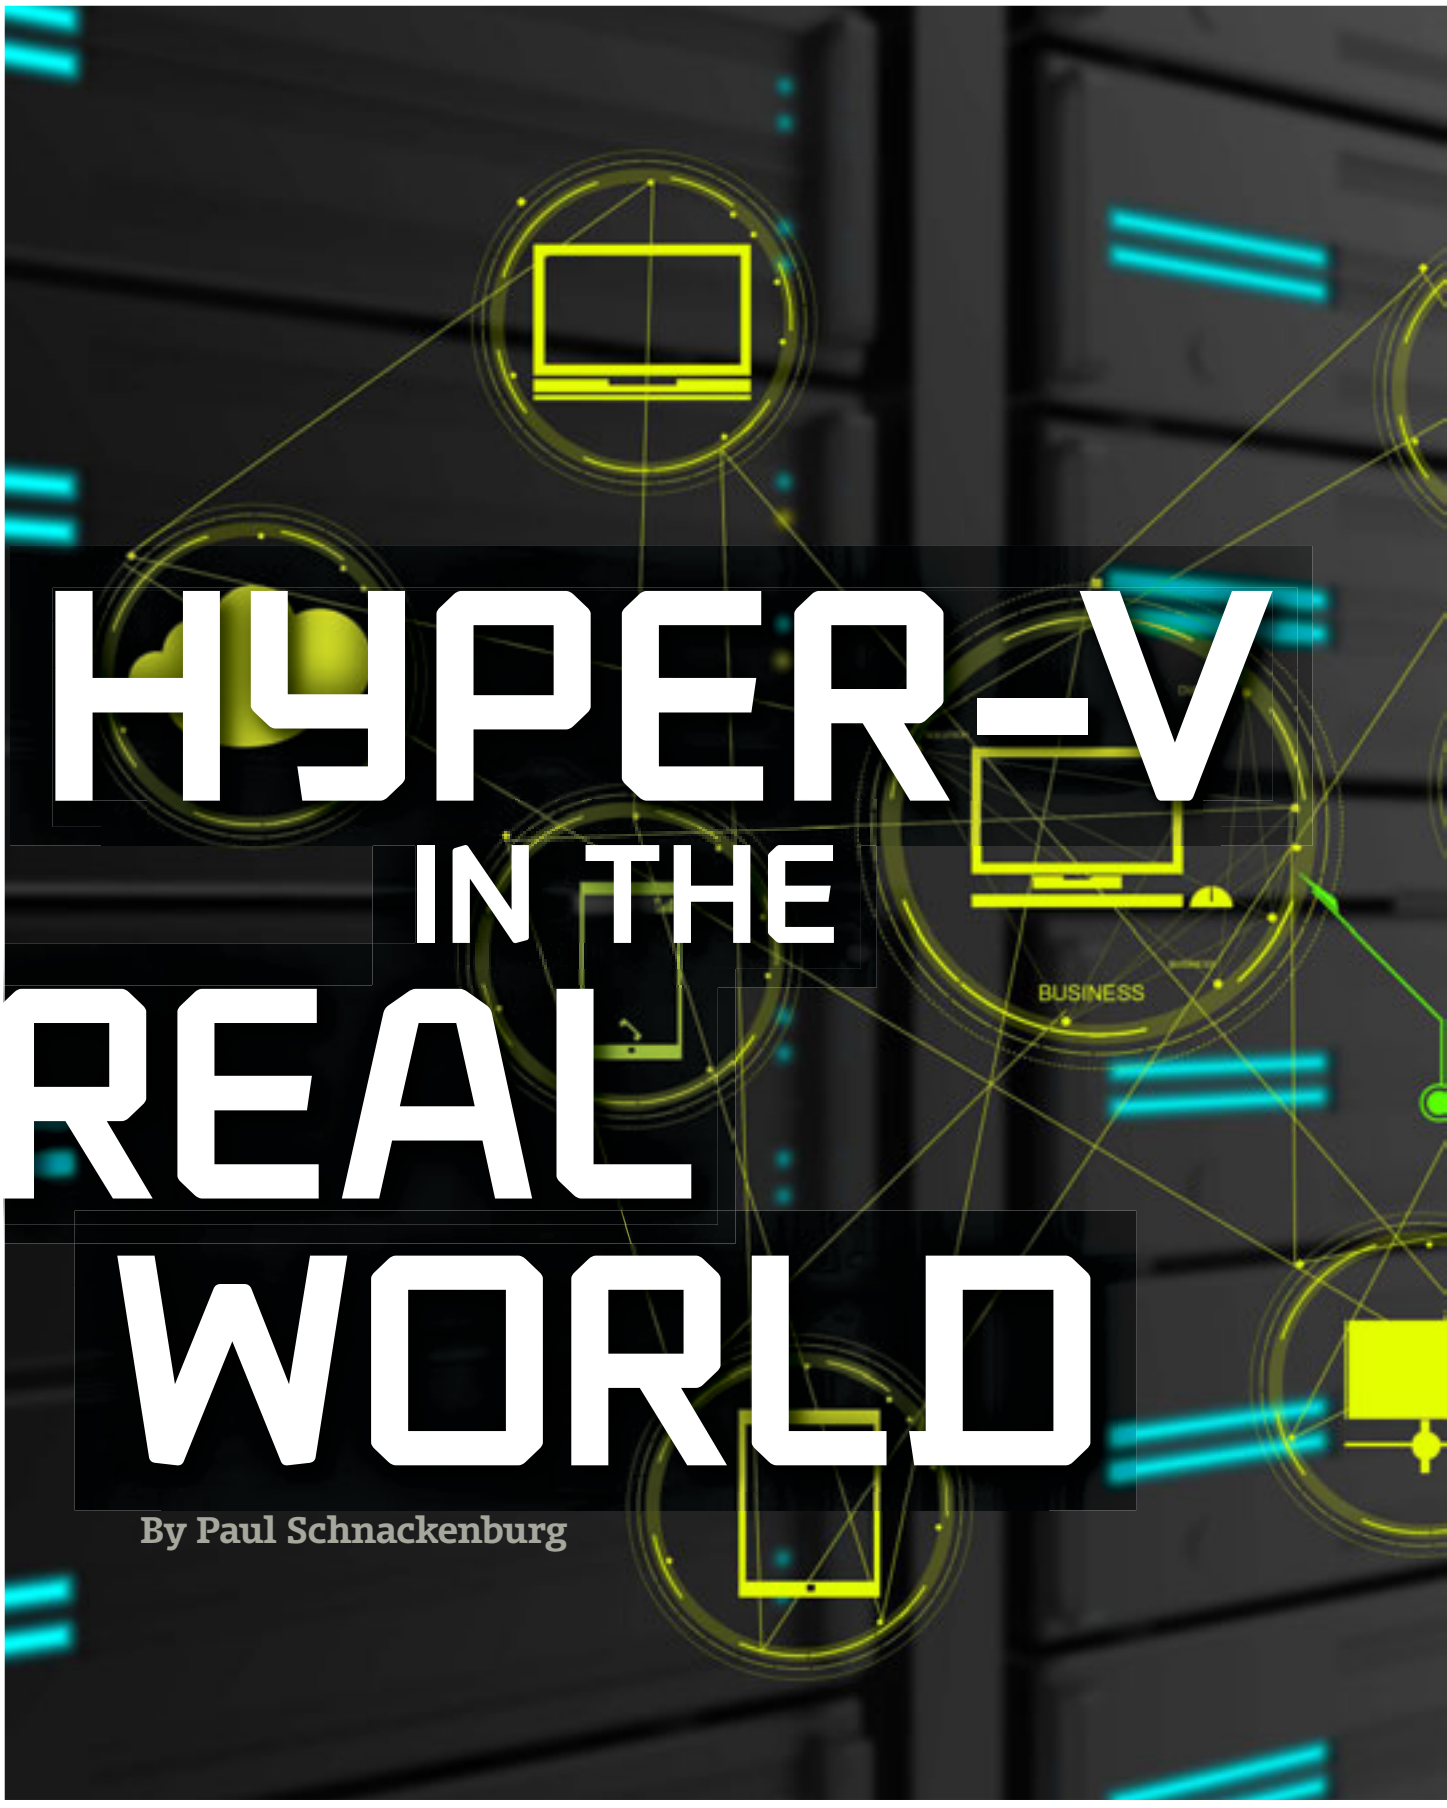

HYPER-V IN THE REAL WORLD

By Paul Schnackenburg

SHUTTERSTOCK.COM

Microsoft's virtualization platform has matured into a solid, dependable workhorse. Here's a primer for getting the most out of it.

THERE'S A GRADUAL SHIFT in the business world away from the market leader VMware Inc. for virtualization; partly because Microsoft Hyper-V is more cost effective in most scenarios, and partly because of the Microsoft Azure cloud platform. Enterprises looking at hybrid cloud functionality find Azure a more logical fit to expand their IT infrastructure into, compared to the limited public cloud footprint and functionality of VMware.

In this article I'll look at how to plan a Hyper-V implementation for your private cloud, including hosts, networking, storage, virtual machines (VMs) and management. I'll also look at how to configure it, along with some tips for ongoing maintenance and troubleshooting.

Planning for Success

As with any IT infrastructure implementation, it pays to do your upfront planning; a well-planned fabric will serve you better than some servers that your server vendor thought would work for your scenario.

The first step is to gather as much information as you can about the expected workloads: What kind of processor, memory, storage and networking resources will they require? What's the projected number of VMs, and how is that number going to change over time? Armed with this information, work through the following areas to assemble a plan of what kind of technologies and how much of them you'll need to achieve a highly performant fabric.

Hosts

While a single host running Hyper-V and a few VMs might work for a small branch office scenario, in most cases you'll want to cluster several hosts together. The industry trend is to scale out from a small number of very powerful (and expensive) hosts toward a larger number of commodity hardware, cost-effective hosts.

This is partly due to the "cluster overhead" effect. If you only have two hosts in your cluster and one host is down, the remaining host has to have enough capacity to host all the VMs. This means you can really only use 50 percent of the cluster's overall capacity.

With four hosts, on the other hand, when one is down you only lose 25 percent of the overall capacity. As you scale up toward the maximum of 64 hosts in a Hyper-V cluster, you gain better efficiency, although in eight-plus-node clusters there's often the need to survive two hosts being down. Just make sure your hosts are big enough in terms of processor cores, storage IO and memory to accommodate your biggest expected workloads (SQL, anyone?).

There are three versions of Hyper-V to select from: the free Hyper-V Server, Windows Server Standard (which comes with two Windows Server VM licenses for that host) and Windows Server Datacenter (which comes with unlimited Windows Server VM licenses for that host).

If you're going to run Linux or Windows Clients as VMs (for VDI), the free Hyper-V Server is appropriate. With the two Windows versions you can install hosts either with a GUI or in the Server Core (command-line only) mode. The latter has less overhead and less attack surface from a security point of view, but make sure your team is comfortable with command-line troubleshooting and Windows PowerShell.

The hosts should have processors from the same vendor (AMD or Intel), but they don't need to be the same model. Select hosts from the Windows Server Catalog to ensure support if you have to call Microsoft.

The first VM that runs on top of the hypervisor is called the parent partition. You should only run management and backup agents here; no other software, and in particular nothing resource heavy, as all resources should be available to the VMs.

Storage

Shared storage is a requirement for clustered Hyper-V hosts. This can be in the form of "local shared storage," which is common in Cluster in a Box (CiB) preconfigured servers where each host is connected to a Serially Attached SCSI (SAS) enclosure. This is appropriate for a small number of hosts (typically two to four).

Larger clusters will need SAN connectivity, either Fibre Channel or iSCSI, or alternatively Server Message Block (SMB) shares on a Scale-Out File Server (SOFS) cluster.

SOFS is a very attractive option to traditional SANs for Hyper-V and SQL Server workloads, both because it's easier to manage and it's more cost effective. A SOFS cluster consists of commodity hardware servers running Windows Server 2012 R2; behind them can be either a SAN or SAS enclosure with HDD and SSDs. The HDDs provide the capacity, while the SSDs provide the IO performance through storage tiering (a new feature in Windows Server 2012 R2).

Data is protected through two- or three-way mirroring (don't use Parity, as it's only for archive workloads), which can span enclosures to protect against a whole enclosure failing.

The point here is to not dismiss SOFS, because on a feature-by-feature basis it matches traditional SANs for both speed and performance, as well as being generally easier to manage.

These options aren't mutually exclusive, though. Say, for instance, you have an existing SAN and need to add a 16-node Hyper-V cluster. If you have spare ports in your Fibre Channel or iSCSI switch and don't mind purchasing 32 new Host Bus Adapters (HBAs), you could simply wire up the new hosts to LUNs on the SAN. But you could also set up two SOFS servers (with only four HBAs and four connections on the switch needed) in front of the SAN, then connect your Hyper-V nodes to the SMB shares on the SOFS.

Virtual Hard Drives

Use VHDX files rather than the older VHD format for virtual hard disks because they support larger disks (up to 64TB) and are more resilient to corruption. The old recommendation to use fixed-size disks (300GB virtual disk takes up 300GB on the underlying storage, no matter how much data is actually stored on the disk) for performance no longer applies.

Dynamic disks that only consume the actual amount of data stored in them on the underlying storage are very close

VEEAM
IT JUST WORKS!™

AVAILABILITY™
for the Modern Data Center

Take Availability for Your Modern Data Center Higher with Veeam Solutions for Microsoft

When your users have problems, recovery can't wait. Veeam has you covered with flexible options from instant VM recovery to file-level and object-level recovery of your critical Microsoft applications.

Microsoft Partner
Gold Application Development
Gold Datacenter

Microsoft Azure
Certified

Service Providers!

See the new Veeam Cloud Connect for Service Providers in Microsoft Azure Marketplace and take your Veeam Availability service to the cloud!

Learn more, visit hyperv.veeam.com

to the same performance. Just be mindful to not oversubscribe the underlying storage when the virtual disks keep growing, so you don't run out of disk space.

To control IO performance, you can use Storage Quality of Service (QoS) to set a maximum IOPS limit (in normalized 8KB chunks) on a per-virtual-disk basis (see **Figure 1**). You can also set a minimum IOPS value, but be aware that this is set on the Hyper-V host, so if the back-end storage can't deliver enough IOPS to satisfy all the VMs across all the hosts, the minimum IOPS aren't going to be delivered.

> In both the Linux and Windows worlds, the latest OSes are better at being virtualized.

Networking

Networking on Hyper-V hosts is used for many different types of traffic: there's the intra-cluster communication (heartbeat), Live Migration (VMs moving from one host to the other while still running), backup traffic, client access to the servers, and if you're using iSCSI or SOFS, there's storage traffic. In previous times these were segregated on separate 1 Gbps interfaces, and you needed a lot of these. Today, most servers come with a couple of 10 Gbps interfaces. To achieve the same separation of different types of traffic, you should use Network QoS (shown in **Figure 2**) to limit the bandwidth for each type of traffic (see "Resources" for more on traffic and other resources).

Windows Server has had NIC teaming since 2012, allowing teaming (either switch independently or with switch awareness through Link Aggregation Control Protocol) of several interfaces (max 32) for bandwidth and redundancy. For storage traffic (and perhaps Live Migration), you can use Remote Direct Memory Access (RDMA) networking cards, which enable SMB Direct. There are three different RDMA technologies: iWarp, Infiniband and ROCE. In essence they all do the same thing, bypassing the software stack for extremely fast traffic (10 Gbps, 40 Gbps and 56 Gbps) with zero CPU overhead. This is how SOFS can match even the fastest Fibre Channel SAN.

In Windows Server 2012 R2, Live Migration can be set to use one of three modes:

* **Compression.** The default, in which both the sending and receiving host's processor load is watched. If there's spare capacity, the data stream is compressed and decompressed on the fly, resulting in a 2x (or better, according to Microsoft) improvement for most scenarios.
* **RDMA.** If you have RDMA network cards, you can use the SMB setting. This will result in fast Live Migrations. Many hosts today accommodate 20 to 30 VMs; when the time comes to patch the cluster or do some other type of maintenance, moving all VMs off one host and doing the maintenance/reboot before moving on to the next host will be much faster with RDMA.
* **Single Root IO Virtualization network cards (SR-IOV).** If you have specific VMs that need low latency network access (to other VMs or clients, not to storage), consider using Single Root IO Virtualization network cards (your server motherboards need to support this, not just the NICs). SR-IOV bypasses the Hyper-V virtualized networking stack and projects a virtual NIC directly into the VMs. Note that you can still Live Migrate a VM using one or more SR-IOV interfaces; if the destination host also has SR-IOV NICs, they will be used (if they have spare SR-IOV virtual NICs available—the maximum number depends on the model of the NIC). Otherwise, the network traffic will fall back to using ordinary Hyper-V networking.

There are other hardware network-enhancing technologies to look at in your planning, like Receive Side Scaling (RSS) and Virtual RSS for

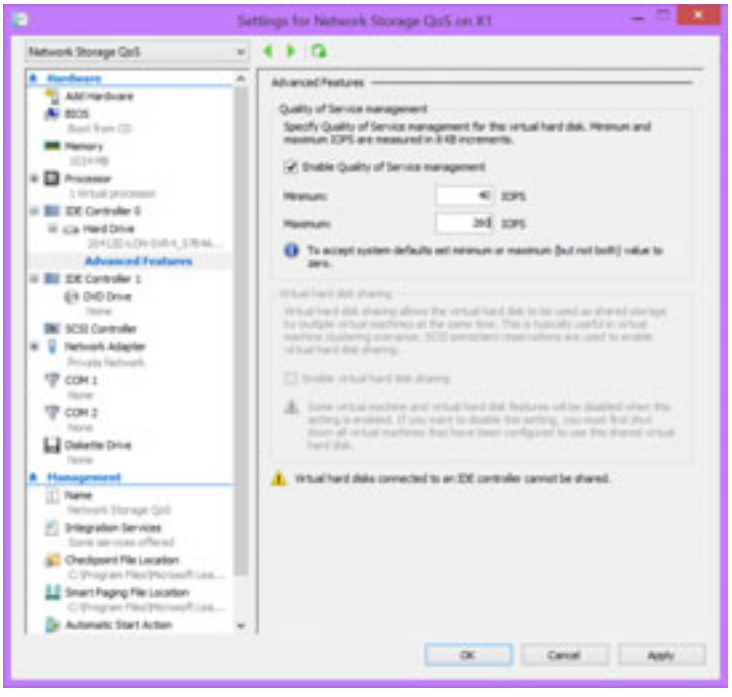

Figure 1. Storage Quality of Service management.

VEEAM
IT JUST WORKS!™

FAST, SECURE BACKUP TO THE CLOUD

NEW Veeam Cloud Connect

Simply choose a Veeam Service Provider and point your backups to their hosted backup repository

Learn more at vee.am/cc

spreading the CPU load of multiple data streams across multiple cores; IPSec Task Offload (TO) for encrypting and decrypting traffic; and Data Center Bridging (DCB), an alternative to network QoS.

Virtual Machine Considerations

Use the latest OS for your VMs if you can; in both the Linux and Windows worlds, the latest OSes are better at being virtualized. If the applications in your VMs can take advantage of multiple processors, you can safely assign as many (up to 64) vCPUs to them as you think they'll need. If the VM doesn't use them, the host will assign CPU resources to other VMs as needed.

If you have large hosts, make sure you're aware of the size of the Non Uniform Memory Access (NUMA) nodes across your Hyper-V hosts. A NUMA node is a combination of a number of Logical Processors (CPU cores) and a set size of memory. If you have really big VMs with larger amounts of memory and vCPU assigned to them than the size of a NUMA node, be aware that Hyper-V will project the NUMA topology into the VM so that the application can take advantage of spreading the load intelligently across the NUMA nodes.

Management

If you have just a few hosts, you can certainly use the built-in Hyper-V manager, along with Failover Cluster Manager, to operate a small cluster. For a fabric larger than, say, four or five hosts, System Center 2012 R2 Virtual Machine Manager is the best management option.

Virtual Machine Manager (see **Figure 3**) brings so much more than just managing VMs to the table:

* It can deploy the OS and configuration to a set of new, bare-metal servers you've just installed, either as SOFS or Hyper-V hosts.
* It can manage your Top of Rack switches, and automatically provision storage for VMs on any SAN or SOFS using SMI-S.
* It has a library for storing all VM components, and it can deploy single or combinations of multiple VMs as services/distributed applications.
* It manages Hyper-V and VMware VSphere environments.
* It can deploy clouds across both platforms, to abstract the details of the underlying resources away from users of your private clouds.

Hyper-V Configuration

Once your planning is complete, you've worked out which technology choices are needed for your environment, and you've purchased the boxes, it's time to configure it all.

As mentioned earlier, you can use Virtual Machine Manager to deploy the OS to both Hyper-V and SOFS hosts. Whether you should go through the necessary testing and fine-tuning to make this work reliably depends on how many new hosts you have. If it's less than about 20, it'll probably be faster to install Windows Server using your standard method (Configuration Manager, Windows Deployment Services or third party).

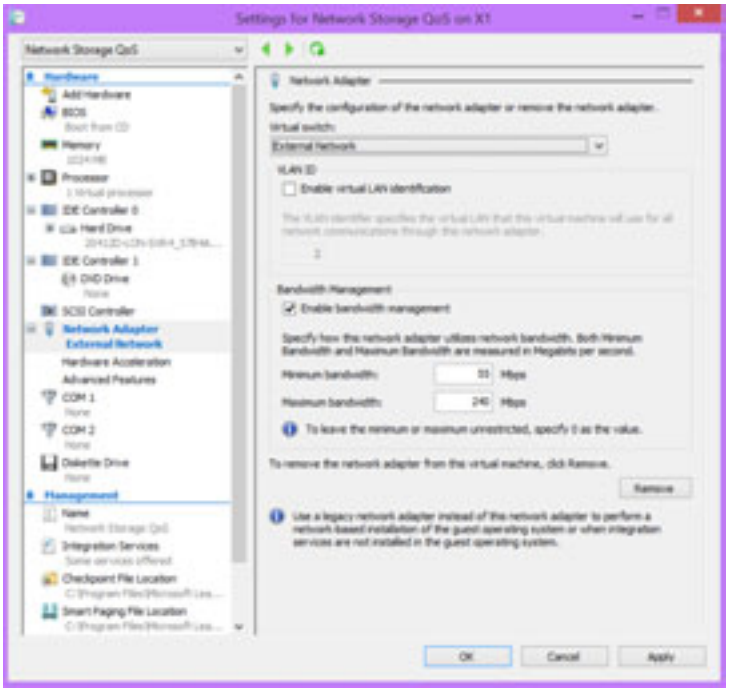

Figure 2. Network Quality of Service management.

The next step is to update all drivers and all firmware; even new servers are often delivered with out-of-date software. After this, you want to apply all Hyper-V and Cluster updates to all hosts.

If you're going to run anti-malware software on the Hyper-V hosts, make sure to exclude the relevant files, folders and processes.

Use Dynamic Memory for all guests, unless the workload inside a VM doesn't support it (Exchange doesn't, for example). If your security policy allows it, consider enabling Enhanced Session mode, which facilitates easy copy and paste from host to guest, along with RDP access, even if the VM has no network connection (or no OS).

Define which network interfaces to use for Live Migration traffic using cluster manager. Right-click on Networks and select Live Migration Settings; then uncheck the networks you don't want to use, and order the ones you do by using the priority values.

If you have VMs running as domain controllers, disable the time synchronization service from the

host and set up the PDC emulator for the root domain in the forest, to keep accurate time using an NTP server on the Internet.

If you're using SANs as cluster shared storage and your SAN supports Offloaded Data Transfer (ODX), make sure you enable this support on each host, as well as in any VM connecting directly to your SAN.

Ensure your VMs are running the latest Integration Components. These used to be installed through the host, but are now distributed through Windows Update (and, thus, through Windows Server Update Services or Configuration Manager), which should make this task easier. If you're using the Datacenter version for your hosts, enable Automatic VM activation and configure each VM with the right key; the days of tracking license keys across your fleet of VMs will be over.

Backing up your VMs is, of course, paramount. You can use Microsoft Data Protection Manager, Veeam Backup Free Edition or recently released Availability Suite v8. The built-in Hyper-V replica can also be used for backup and disaster recovery by asynchronously (every 30 seconds, or 5 to 15 minutes). It replicates VM disk writes to another Hyper-V host, first in the same datacenter and then (if you use extended replication) to another datacenter or Azure. These replica VMs can be failed over to in case of a disaster; if that happens, they can be injected with a different IP configuration to match the subnets in the failover datacenter or Azure virtual network.

Maintaining Hyper-V

Virtual Machine Manager works hand-in-hand with Operations Manager to monitor for issues in your infrastructure, as well as provide forecasting for your capacity utilization. Veeam offers comprehensive Management Packs for Operations Manager for both Hyper-V and VMware environments.

If you have specific performance issues or need to baseline your Hyper-V or storage nodes, you can use Performance Monitor; don't use Task Manager, as it's not Hyper-V-aware. Download and learn how to use Performance Analysis of Logs (PAL), which will help you manage Performance Monitor logging across several hosts.

Hyper-V is a very capable platform to build your business's private cloud on; and, along with System Center, it can be managed and monitored effectively. Technical insight and planning is required to navigate all the different options available. Hopefully, this article has provided some guidance toward a successful implementation. VR

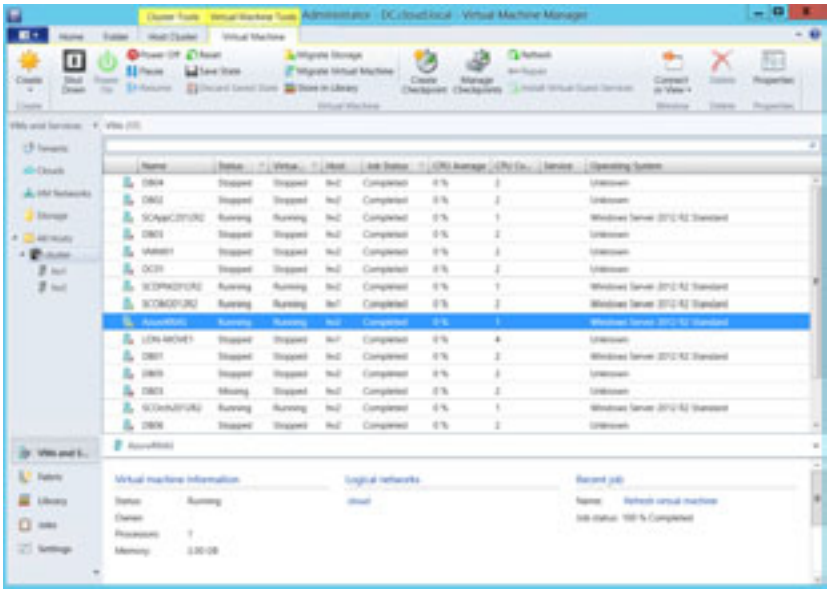

Figure 3. Microsoft System Center Virtual Machine Manager at work.

Hyper-V Resources

FREE HYPER-V SERVER:
bit.ly/1wvv3QY
NETWORK ARCHITECTURE SERIES (QOS & NIC TEAMING):
bit.ly/1pVaZIm
HYPER-V AND CLUSTER UPDATES FOR WINDOWS SERVER 2012:
bit.ly/1lbv9sk and bit.ly/1K9DM3G
HYPER-V AND CLUSTER UPDATES FOR WINDOWS SERVER 2012 R2:
bit.ly/1kxOyYS and bit.ly/1z5BfCw
ANTI-VIRUS EXCLUSIONS FOR HYPER-V HOSTS:
bit.ly/1v7MtHi
BARE METAL OS DEPLOYMENT IN MICROSOFT SYSTEM CENTER 2012 R2 VIRTUAL MACHINE MANAGER:
bit.ly/1zlBQzs
HYPER-V OFFLOADED DATA TRANSFER OVERVIEW:
bit.ly/1Qg6xzD
DEPLOY WINDOWS OFFLOADED DATA TRANSFERS:
bit.ly/1EaLZjy
AUTOMATIC VIRTUAL MACHINE ACTIVATION:
bit.ly/1KHAHsv
PERFORMANCE ANALYSIS OF LOGS (PAL):
bit.ly/1l9D2Nl

Paul Schnackenburg, MCSE, MCT, MCTS and MCITP, started in IT in the days of DOS and 286 computers. A regular contributor to Virtualization Review, *he runs Expert IT Solutions, which is focused on Windows, Hyper-V and Exchange Server solutions.*


10001
10110
011001

There's a dizzying selection of storage solutions out there. Which one is right for you? The answer will often be "more than one."

The Storage Explosion Is Here

By Dan Kusnetzky

Suppliers of storage hardware and software are presenting what appears to be a huge list of options. Which ones are best isn't always clear. Furthermore, it's not clear if there's a single option that addresses an enterprise's needs. All of the suppliers, however, promise that their solution is the best, the most cost-effective and makes the best possible use of ‹insert the name of your favorite storage technology here›.

SHUTTERSTOCK.COM

Marketecture Abounds

Like just about every other area of the IT market, suppliers of storage are always doing their level best to one-up their competitors, seeking ways to out-gun the others in the areas of storage performance, reliability, scalability and overall cost.

Although there's a great deal of noise in the market, a few things are clear. For one thing, each of the suppliers believes it and it alone is uniquely qualified to be the only source of storage technology. Also clear is the fact that there are a number of different types of technology from which enterprises may choose, and an overwhelming set of combinations for how this technology may be used together. Finding the right fit largely depends on what question the enterprise is asking.

Increasingly, storage services are being offered by suppliers of managed services, colocation or cloud services.

Why So Confusing?

There are a number of different approaches to storing applications and data, and each is useful in the right place and at the right time. Some approaches require the storage media be directly connected to clients and servers, while others attach storage devices to a storage appliance or a storage server.

If the storage device is directly connected to the client or server, there are many different storage interconnects from which to choose. Each offers a different mix of price and performance, and can control what media options are available.

As with direct-connect approaches, there are several different storage interconnects in use in today's datacenters when the storage devices are attached to a storage server or appliance. These servers may be connected to computing systems using a general-purpose LAN or special-purpose SAN.

The industry is also seeing the increasing use of system memory being used as a special form of storage for computationally intensive, extremely high-performance applications. Sometimes suppliers call this "distributed cache" or "in-memory database."

To add to the confusion, cloud services providers have begun to offer an array of new Storage-as-a-Service products. They're trying to convince enterprises that it's better, less complex and less costly to use those services rather than purchase, install and operate their own storage.

Different Technology for Different Needs

The industry has used various types of technology over the years, including:

* **Tape.** Different suppliers have offered paper tape, cassette tape, and reel-to-reel tape products. Several suppliers have offered direct access tape devices that could replace rotating media for large-scale storage applications.
* **Rotating media.** Different suppliers have offered rotating drums and a whole herd of different types of disk storage. While most of these were based on magnetic recording, some were based on optical recording technology.
* **Solid state.** Although solid-state storage has been available for decades, and its access times and throughput made it extremely desirable, the cost was prohibitive for most applications. Recently, however, the introduction of new technologies has resulted in the rapid adoption of flash memory.

As suppliers seek ways to offer flexible and inexpensive storage options, the market is seeing the emergence of distributed cache solutions using the system memory of low-cost, industry-standard servers, blades or distributed NoSQL database solutions using server clusters.

You've Got Options

Like most areas of IT, there are many different types of storage technology, and each has the ability to serve a different set of needs. If the enterprise carefully reviews its application portfolio, it will soon become clear that each application has a different storage profile.

Some applications require the storage and retrieval of huge amounts of data, and longer access times are acceptable. Other applications access huge amounts of data, but the access time must be kept to a minimum. Still, others require immediate access to data and any delay is unacceptable. Finding the right solution necessitates understanding your environment's unique requirements. Most fall into one of these categories:

* **Long-term storage.** The requirements for this type of storage usually include massive capacity and low cost per megabyte or gigabyte. Applications using this data typically are batch or analytical jobs.
* **Medium-term storage.** The requirements here lean more toward finding a good balance between performance and cost. This often means storing applications and data for remote or VDI desktops, servers or even handheld appli-

cations. Enterprises are often willing to compromise on storage performance to reduce overall cost.

* **Short-term storage.** The requirements for transactional or business intelligence applications often include the need for very short access time and medium levels of throughput. Enterprises are often willing to compromise on cost to obtain performance.

> Cloud services providers have begun to offer an array of new Storage-as-a-Service products.

* **Storage for high performance or technical processing.** The requirements for this type of workload often include extreme storage performance (seek performance and access time or latency), extreme needs for data throughput, and huge volumes of data. Shared cache, clustered NoSQL databases and in-memory databases are often used to address these requirements.

* **Flexible storage.** When the enterprise faces ever-changing, dynamic requirements, it will often turn to some from of distributed or hybrid storage. A local cache made of high-performance storage is deployed to improve the performance of off-site or cloud storage. The off-site storage may be at another enterprise-owned site, at a site managed by a managed services supplier or in the data-center of a cloud services provider.

How Are Suppliers Addressing These Requirements?

Although each supplier is addressing these storage requirements differently, there are some common threads:

* **Slow, but reliable storage for huge amounts of data.** A form of tape or optical technology often satisfies these requirements. Some cloud services providers are suggesting their Storage-as-a-Service offerings might be a replacement for this type of storage. It's not clear what type of storage they're actually using to address this need. Low-cost, low-performance rotating disk storage is very likely part of the cloud services provider's offering.

* **Fairly fast and inexpensive disks.** These can be used for client-side applications or server-side applications for small to midsize businesses.

* **High-speed, expensive disks.** These target server-side applications that need both a larger amount of storage and high levels of performance.

* **Flash and other forms of solid-state storage.** They're packaged as storage devices that address the needs of applications requiring very low access times or high levels of throughput. Typically these devices are much faster and more expensive than traditional disks, and offer less capacity.

* **Internal solid-state memory.** It's packaged by storage virtualization technology so that it appears to be a storage device.

What's the Best for Me?

The enterprise must take the time to survey its portfolio of workloads to learn the answers to the following questions:

* How much is the enterprise willing to pay for storage? High-performance storage typically is expensive.

* Does the enterprise really need massive storage capacity? There are many ways to address this type of need. The best answer usually is a compromise between cost, performance and storage capacity.

* Is the enterprise willing to use off-site storage? Enterprises in regulated environments may only be able to use on-site, locally controlled storage for regulated applications. Collaborative applications, e-mail and other non-regulated applications might be candidates for off-site cloud storage.

The Golden IT Rule

Most enterprises, by necessity, rely upon many different types of storage. This is partially due to the different needs of each application, and also due to the application's age. Older applications, for instance, are likely to be using older types of technology. Enterprises often follow the golden rule of IT, "If it's not broken, don't fix it," when dealing with these applications.

> The market is seeing the emergence of distributed cache solutions.

The more enterprise decision makers know about their applications and their requirements, the easier it is for them to select the right storage technology, storage location and determine if cloud storage is even a reasonable option. VR

Daniel Kusnetzky, a reformed software engineer and product manager, founded Kusnetzky Group LLC in 2006. He's literally written the book on virtualization and often comments on cloud computing, mobility and systems software. He's been a business unit manager at a hardware company and head of corporate marketing and strategy at a software company.

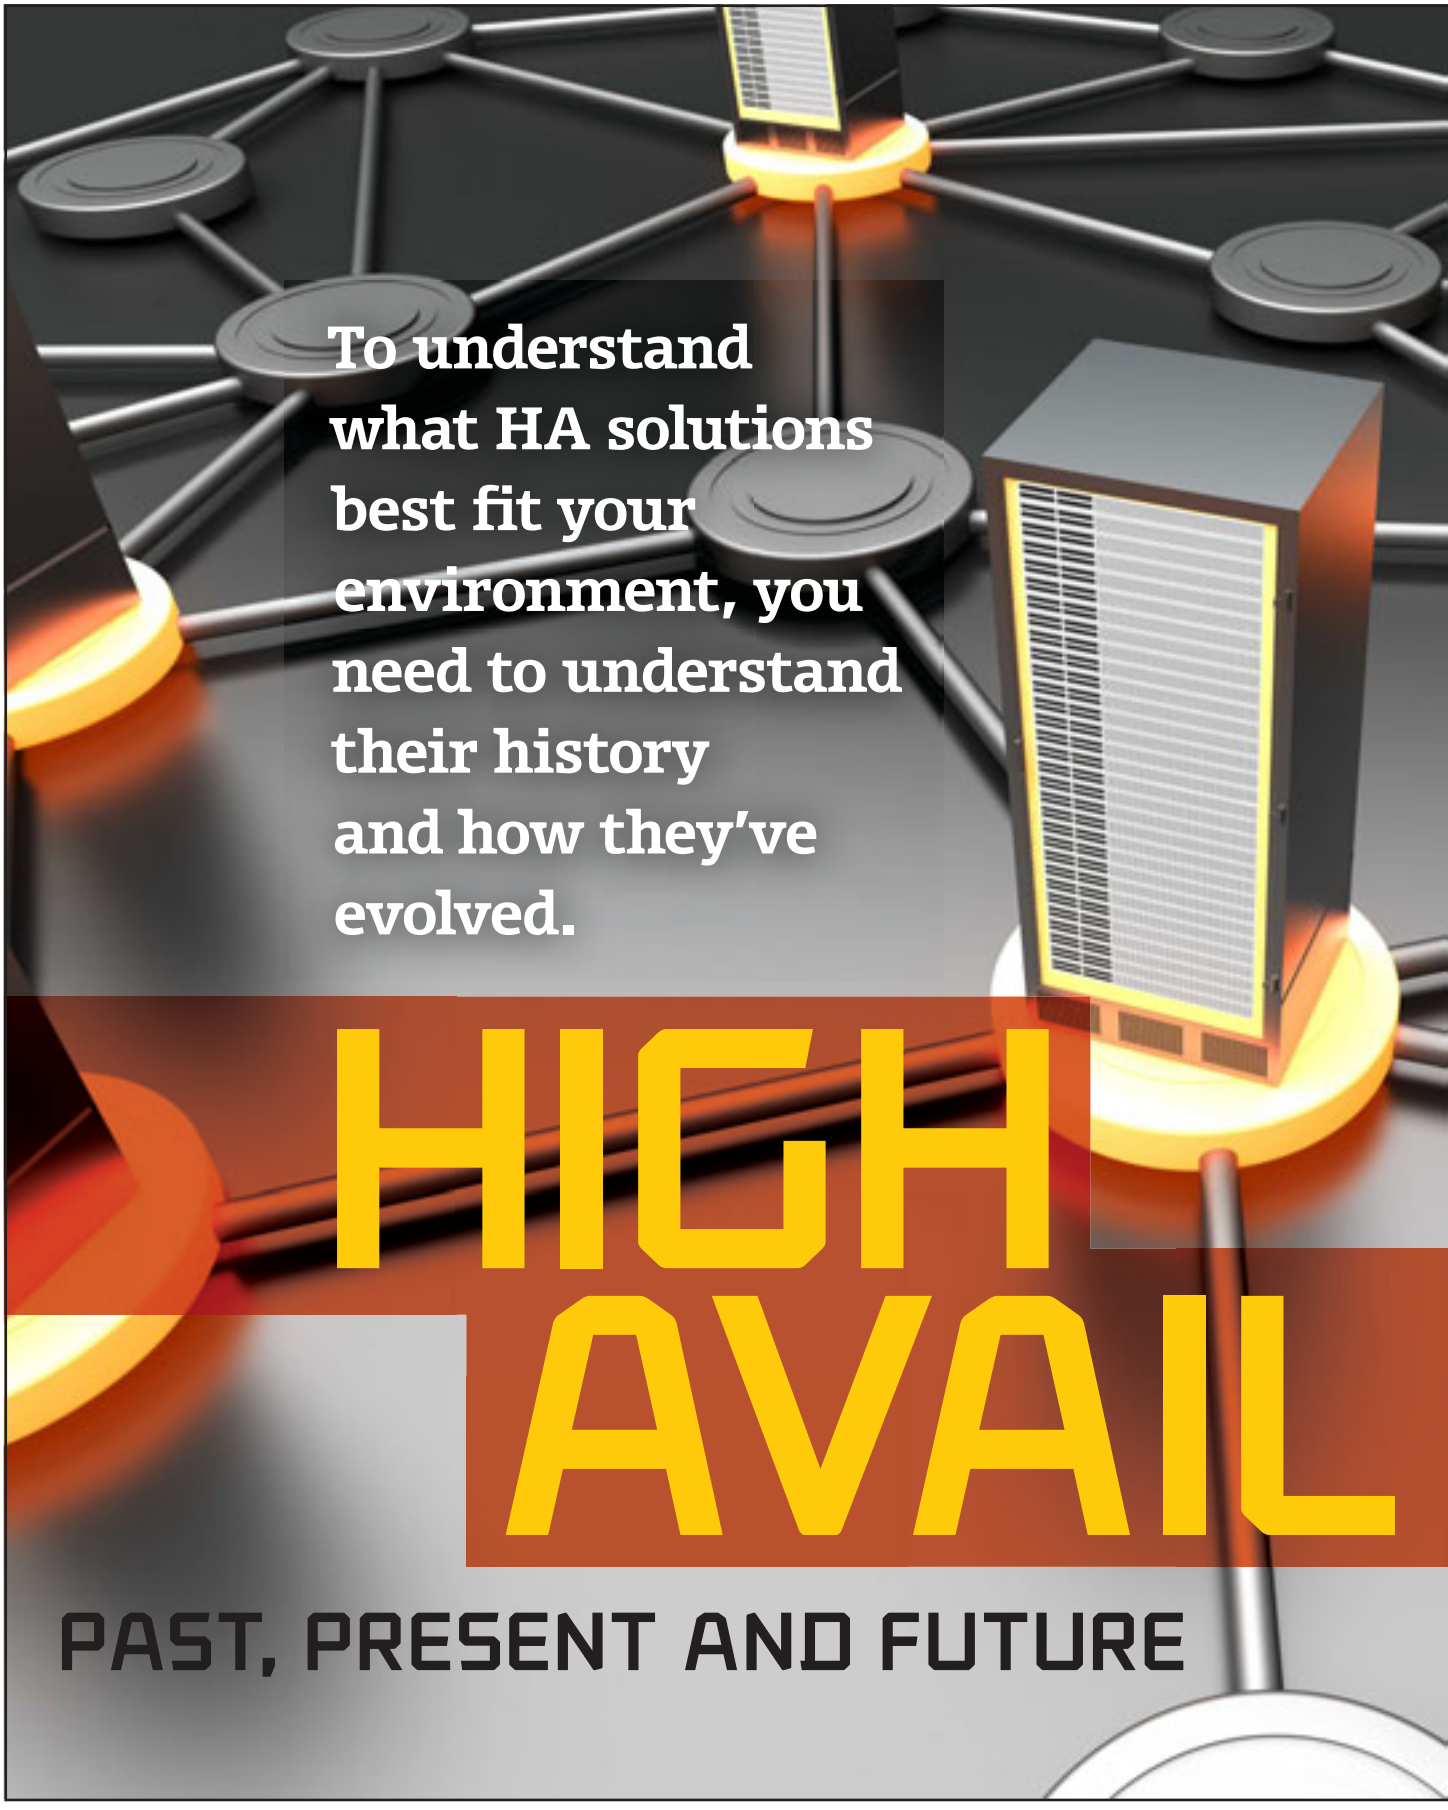

To understand what HA solutions best fit your environment, you need to understand their history and how they've evolved.

HIGH AVAIL

PAST, PRESENT AND FUTURE

SHUTTERSTOCK.COM

ABILITY:

By Dan Kusnetzky

HIGH AVAILABILITY (HA) is a topic with a great deal of history. Different approaches have been used over time to make sure applications, services, databases, networks, and storage remain available and reliable to support enterprises. As enterprises have grown increasingly reliant on information technology-based solutions, the need for these solutions to always be available has increased as well.

Most HA solutions rely on redundant hardware and special-purpose software designed to make the best use of that hardware. Virtualization and cloud computing are upending earlier approaches to HA. Organizations have learned that the use of virtualized access, applications, processing, network and storage makes the creation of HA solutions easier. They've also learned that virtualization makes it easier to use off-site cloud hosting as part of an HA solution.

Most HA solutions rely on redundant hardware and special-purpose software designed to make the best use of that hardware.

HA solutions can be expensive, though, and an enterprise's portfolio of IT solutions might not need the same level of availability. Business-critical functions are likely to need the highest levels of availability, while the requirements for business support functions are not likely to be as high.

Enterprises would be wise to understand all of the following approaches to HA, and make the proper choice for each of their workloads.

A Brief History of HA

When applications were more monolithic back in the 1960s through the 1990s, the UI, application logic, storage management, data management and networking functions were all hosted together on a single system. Back then, the industry focus was on making the systems themselves "fault tolerant."

This was accomplished by designing mainframe systems that used multiple processors, stacks of memory, storage adapters and network adapters; they included system firmware that monitored the health of individual components and moved workloads to surviving components in case a component failed or became unresponsive. IBM Corp. used "Parallel Sysplex," a special marketing catchphrase to describe these systems.

Parallel Sysplex failover took only a few microseconds or milliseconds. People using these workloads were usually unaware that a failure took place at all. These systems were extremely expensive when compared to standard off-the-shelf configurations, and were only used to host the most critical workloads.

IBM continues to make continuous processing mainframe configurations available today.

Suppliers such as DEC (now part of Hewlett-Packard Co.), Stratus Technologies and Tandem Computers (both part of HP now) developed similar technology in a smaller form factor—the minicomputer. IBM resold Stratus computers using the System/88 name.

As with the mainframe continuous processing systems, these systems were composed of redundant components and special-purpose firmware that detected failures and rapidly moved workloads so they could continue processing.

Failovers typically would only require milliseconds, and the users of these workloads were left unaware that a failure happened.

Because these systems were also quite expensive when compared to the off-the-shelf minicomputer competitors, they were only adopted to support the most critical workloads.

HP Integrity and Stratus ftServer systems are available today to address these business requirements.

Clustering

Suppliers hoping to address requirements for performance, reliability and availability worked to create more software-oriented solutions.

Rather than focusing on special-purpose hardware and firmware, these companies focused on special clustering and workload management software. The software orchestrated the use of either off-the-shelf networking solutions or special-purpose clustering networks.

Although clustered systems are likely to have been created by researchers as early as the 1960s, the first commercial offerings were the Datapoint ARCnet in 1977, which wasn't a commercial success, and the DEC VAXcluster in 1984, which was an overwhelming success and is still in use in many enterprises today.

These hardware configurations were used in a number of different ways. Each had a different goal and could be considered the earliest use of access, application, processing, networking and storage virtualization.

Customers deploy clusters, like those in **Figure 1**, to address the requirements for raw processing power, access availability, application availability, database availability, processing availability and even storage availability.

Different layers of virtualization technology are deployed, depending on the goals of the enterprise. Kusnetzky Group LLC has divided this virtual cake into seven layers, which you can read about at **VirtualizationReview.com/7LayerModel**.

Access Clusters

In access clusters, the basic cluster hardware configuration is used to make entire application systems available by using what is now thought of as "access virtualization" technology. Applications are installed on several cluster nodes, and if the node supporting the work of one group of users begins to

fail, workload access is shifted from the failing system to one of the surviving cluster nodes.

While this appears similar to an application cluster, the failover and workload management is being done at the access level rather than the application level. Applications aren't aware of this technology and don't need special APIs or to be specially architected for this failover to occur.

This type of cluster relies upon data being housed on a separate part of the cluster devoted to storage access, on the storage services of another cluster, or on a storage-area network (SAN) so that data remains available even if the systems hosting the applications themselves failed.

Because access virtualization is the main virtualization technology in this type of cluster, application and storage hosts might be housed in the same or different datacenters.

Suppliers such as Citrix Systems Inc., Microsoft and VMware Inc. supply this type of technology.

Application Clusters

In application clusters, the basic cluster hardware configuration is used to make applications or application components available by using what is now thought of as "application virtualization" technology.

Application virtualization technology is used to encapsulate applications or their underlying components. The application virtualization technology controls access to these virtualized components. As users request the use of these applications, the workload management portion of this technology reviews the available processing capacity of the systems it's monitoring, selects a system to execute the application based on policies and the availability of processing capacity, and then starts up the application or sends the user's requests to an already-running application instance.

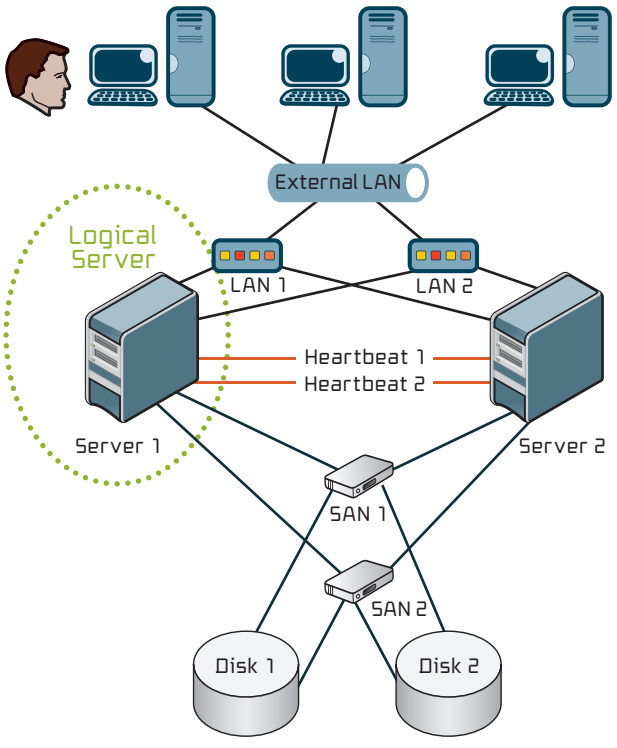

Figure 1. A typical two-node High-Availability cluster.

> **Less-critical applications might be happy executing on a cluster.**

If an underlying system is failing, the user's workloads are automatically moved to another system in the cluster, or connected to workloads already running on another system.

While this appears similar to an access cluster, the failover and workload management is done at the application-component level. Applications must be architected to work with the application virtualization's workload management tool to enable workload monitoring, management and migration. So, unlike access clusters, the applications are extremely aware of this technology and must use special APIs or be specially architected for failover to occur.

This type of cluster relies on data being housed on a separate part of the cluster devoted to storage access, or on a SAN so that it remains available even if the systems hosting the applications themselves fail. Access virtualization technology is often utilized, as well, so user access can be easily and automatically migrated from the failing systems to the new systems.

Application virtualization is the main virtualization technology in this type of cluster; storage hosts could be housed in the same or different datacenters.

Suppliers such as AppZero, Citrix, Microsoft, Novell Inc. and VMware offer application virtualization products today.

Processing Clusters

In processing clusters, the basic cluster hardware configuration is used to make entire system images available by using clustering managers, a form of "processing virtualization" technology.

Applications or their components are architected to access a cluster manager, and the cluster manager monitors the application and either restarts the application on another system or moves the working application to another system, depending on the type of failure. Workload management and migration are managed at a low level inside the OS.

As with the other types of clusters, this approach relies on data being housed on a separate part of the cluster devoted to storage access or on a SAN, so that it remains available even if the systems hosting the applications themselves fail. Also, access virtualization technology is utilized so user access can be easily migrated from the failing systems to the new systems automatically.

> Applications must be architected to work with the application virtualization's workload management tool.

In this case, a form of processing virtualization—cluster and workload management—is the main virtualization technology in this type of cluster. Storage hosts can be housed in the same or different datacenters.

Suppliers such as Citrix, Microsoft and VMware offer this type of processing virtualization today.

Database and Storage Clusters

Another use of the traditional cluster configuration is to support parallel- or grid-oriented databases or storage. The cluster manager's ability to support specially architected database technology, such as Oracle RAC or IBM PureScale DB/2, are typically database offerings designed for this type of configuration. While it does enhance database availability, the primary goals are database performance or scalability. New NoSQL databases, such as those offered by Couchbase, FoundationDB and MongoDB, are also designed to support large-scale clusters.

Special-purpose SANs are also built using this type of technology. Often, general-purpose systems access data stored in this system over a special-purpose, high-speed SAN.

Virtual Machine Software Emerges

A couple of processing virtualization technologies, virtual machine (VM) software and OS virtualization and partitioning, have emerged as the focus of today's HA strategies. Entire systems are encapsulated and workload monitoring and management combined with system image migration technology are replacing previous forms of clusters.

Applications running in these system images don't need to be written to use cluster APIs. If a virtual system appears to be in trouble due to a hardware failure, the entire virtual system can be moved to another host. This is a significantly simpler approach to HA. Failover can be managed in seconds or minutes.

Continuous processing systems, however, are better hosts for critical functions. Failover in that type of environment can take place in milliseconds or microseconds.

The Design Center Has Changed

The industry is in the final stages of a significant design center migration. In the past, the design center was keeping systems available and reliable through the use of special-purpose hardware and firmware. Now, the design center is using virtualization technology to assure that applications and their underlying components are available.

The new assumption is that hardware, regardless of whether the hardware is a system, network component or storage component, is going to fail; and properly designed software can provide a low-cost, simple-to-use strategy to address that failure.

Once a system image is encapsulated, it can be hosted on a local system, a system in another datacenter or on a system in a cloud services provider's datacenter.

How Much Availability Do You Really Need?

We're now in a world in which enterprises increasingly need their systems to be constantly available, *and* in a world in which these same enterprises need to do the most with a reduced IT budget and staff.

> Continuous processing systems are better hosts for critical functions.

Enterprises would be well advised to review their portfolio of applications to determine how much availability is necessary for each application, rather than how much is available. Some applications cannot be seen to fail, while it may be OK for other applications to become unavailable from time to time.

Business-critical applications are be best hosted on continuous processing systems. Less-critical applications might be happy executing on a cluster or even out in the cloud somewhere.

My advice is select the HA strategy right for each application, rather than using a "one-size-fits-all" approach. VR

Daniel Kusnetzky, a reformed software engineer and product manager, founded Kusnetzky Group LLC in 2006. He's literally written the book on virtualization and often comments on cloud computing, mobility and systems software. He has been a business unit manager at a hardware company and head

VEEAM LIVE | shows, whiteboards, specials, webcasts etc.

Join the **Veeam LIVE** team for an interactive session around **Virtualization topics!**

Microsoft Azure and Veeam:
Innovative solutions for Data Availability

What's new
in Veeam Availabilty Suite v8

Veeam R&D Inside Out
with Anton Gostev

Veeam Cloud Connect Deep Dive: Design
principles & deployment best practice

Learn and
Win Prizes
including
a trip to
VeeamON 2015

Watch live
for more info!

Learn more
about Veeam Live
vee.am/veeamlive

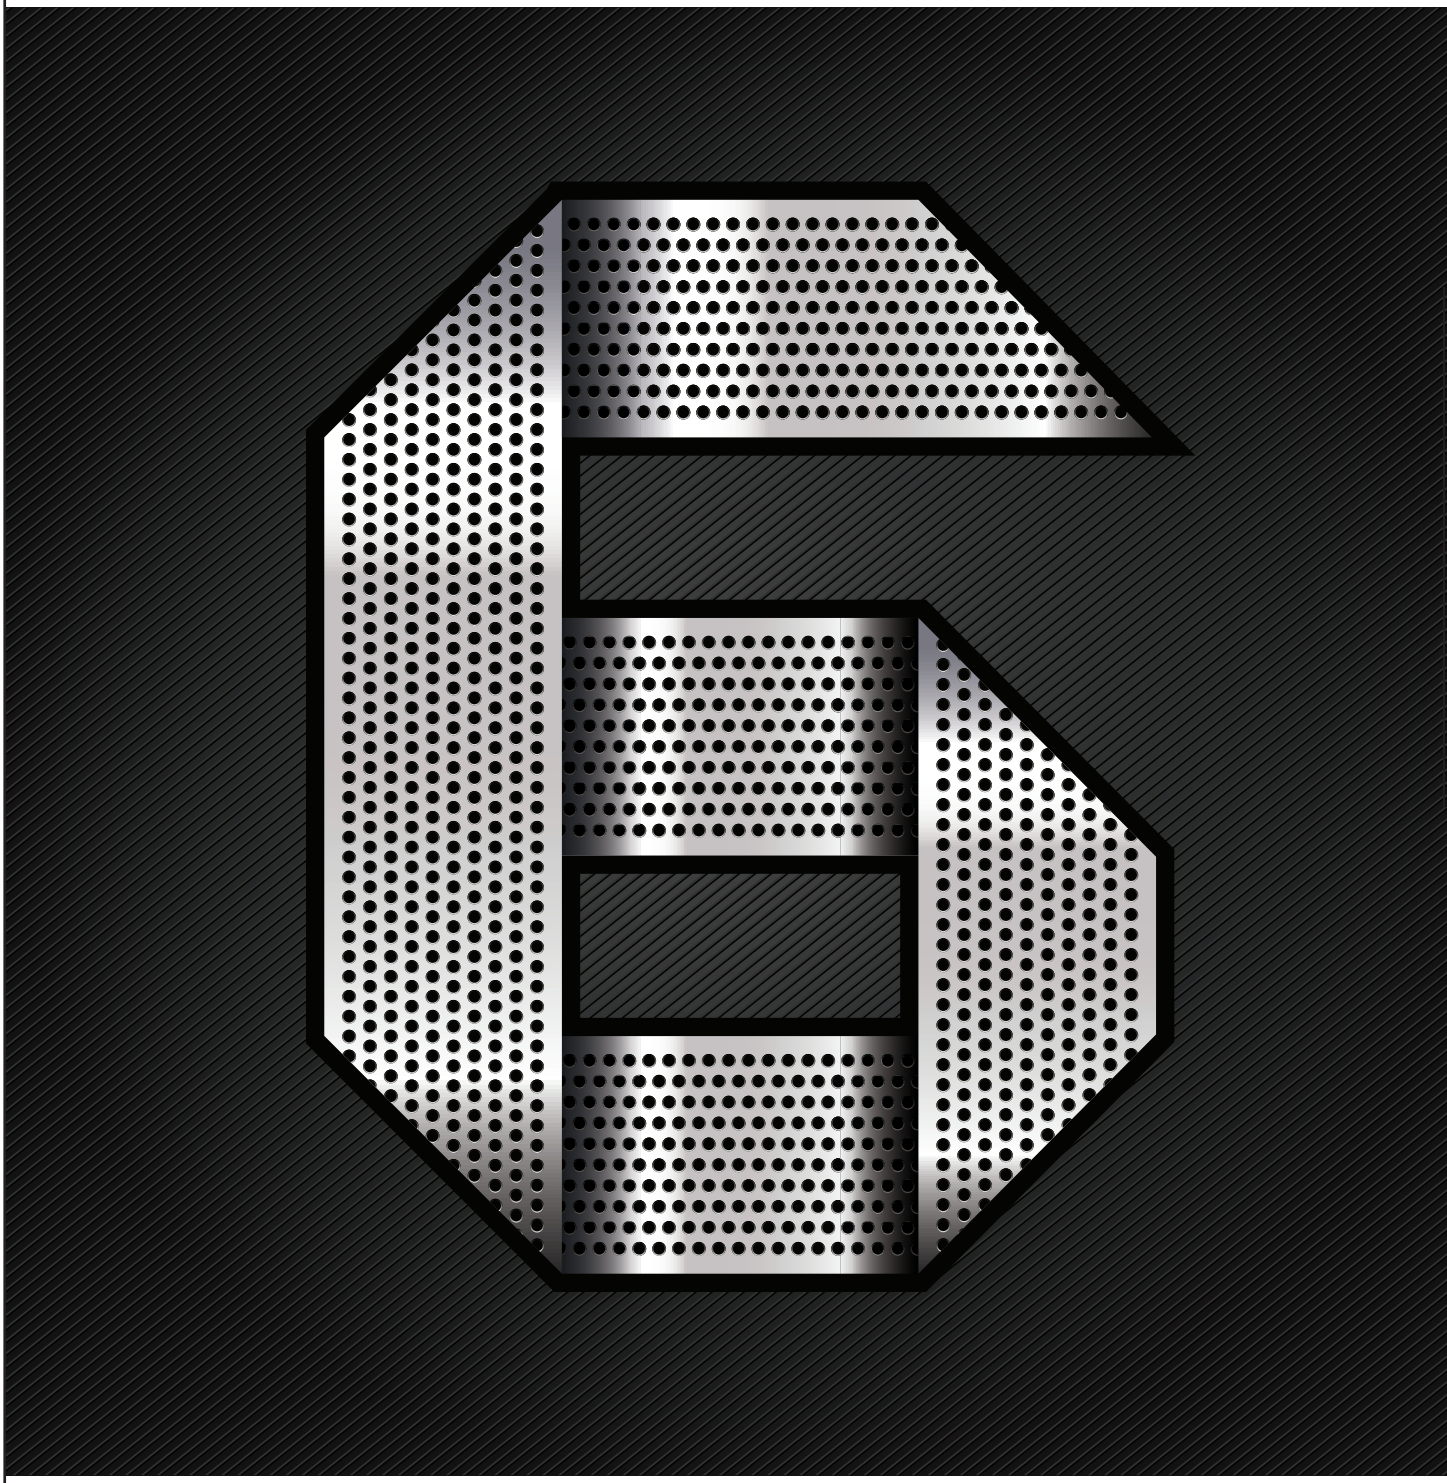

> If you try to use your legacy physical backup tools in your virtual environment, you're asking for a lot of pain and suffering.

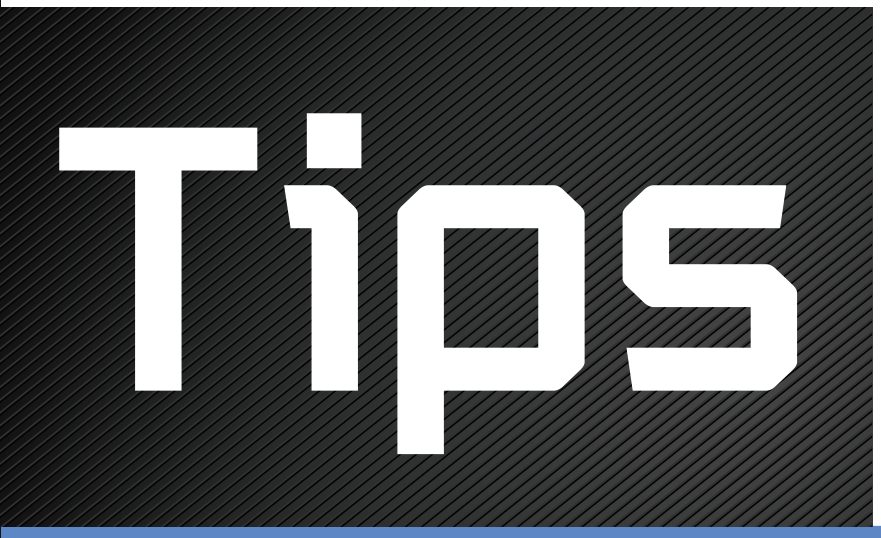

Almost no datacenters are homogeneous these days, and Hyper-V is becoming increasingly popular. Before creating your first virtual machine, however, you need to determine your backup strategy.

By Chris Henley

for Backing Up Hyper-V

Hyper-V is now a common hypervisor choice for both large and small businesses. Hyper-V offers great virtual machine (VM) capabilities, supports hardware and software innovations, and can save money. VMs running on Microsoft Hyper-V host the OSes and applications central to your day-to-day business operations. Your users interact with those VMs and their associated OSes and applications to perform their daily work and generate your business's critical data. To protect your VMs, the underlying OSes, the applications and the data contained in each VM, it's a good idea to adhere to the following best practices for backing up Hyper-V.

SHUTTERSTOCK.COM

DATE:

October
26-29th, 2015

LOCATION:

Aria Resort,
Las Vegas, NV

ATTENDEES:

2,500+
IT pros

Over 2,500 IT pros are expected to engage in over 80 break-out sessions and various contests at VeeamON 2015, a three-day mega-conference in Las Vegas. What's the caveat? If Vegas isn't enough, get ready for the data industry's premier party to wash all of that knowledge down.

Regarded by many as THE data availability event, VeeamON 2015 is set to eclipse expectations from last year. VeeamON 2014 saw over 1,300 IT professionals taking part in breakout sessions, contests, technical trainings and a slew of networking opportunities. This year, VeeamON 2015 is reaching even further to provide increased technical-training opportunities, education and fun for over 2,500 expected registrants.

Breakout sessions
Over 80 breakout sessions will allow attendees to learn about industry-leading products and innovations. Sessions will span into technical deep-dives on Veeam's v9 release, partner sessions and other special areas of focus. Throughout the 80+ sessions, attendees will be able to join discussions led by Veeam's partners. An even greater amount of technical presentations are slated for this year's event, offering greater insight into technology and best practices. Simply put, attendees are urged to come and learn how to enable the Always-On Business™.

Of course, there's THE party!
Widely considered as THE IT party of the year, the VeeamON 2015 party will be the perfect ending to three days of deep-dives, user groups, networking and trainings. Work hard and play hard: that's the Veeam mantra — and we do not like to disappoint! Let's face it, with over 2,500 people inside Vegas's hottest new club, this isn't a party you'll want to miss!

Additional exciting opportunities

VMCE (Veeam Certified Engineer) training
Offered before the main slate of programs, the VMCE training provides system engineers, backup and virtualization admin, solution architects, consultants and any other IT professionals using Veeam solutions an opportunity to become a Veeam Certified Engineer. What does this mean? In addition to being considered an expert on Veeam Software solutions, benefits include the ability to publicly use the Veeam Certified Engineer logo and have direct access to support for critical onsite issues.

Hands-on labs
New to VeeamON 2015, hands-on labs will allow attendees to interact with Veeam engineers and learn more about the mechanics behind why we say: **Veeam – It Just Works!**

Lab Warz
Lab Warz is back for VeeamON 2015! Arguably last year's most technologically interactive program, this program will give attendees the chance to extrapolate common software challenges in real time. Each station will be staffed with a Veeam systems engineer walking groups through live virtual scenarios and challenges. This is where the real techies will get to shine and showcase their skills.

Awards during last year's Lab Warz included a $10,000 first-place prize. Be prepared to win big!

Social Opportunities
VeeamON 2015 is the ideal venue for IT pro's looking to network. With everything from one-on-one time with Veeam executives and partners and top IT experts, to invaluable takeaways from numerous group discussions, VeeamON 2015 will allow attendees to connect on a global scale.

Join over 2,500 IT pros
VeeamON2015 provides a unique opportunity to meet and learn from the very best in data availability. Get the green light from your boss, TODAY! Register, book it and spread the word. VeeamON 2015 will be the event IT pros will be talking about!

RPOs AND RTOs: THE REASON FOR BACKUP

No business can sustain extended periods of downtime or significant data loss—the cost associated with these issues is simply unacceptable. A responsible IT organization works hard to find ways to minimize downtime and data loss for its organization.

A simple measure of how quickly you can recover from downtime and data loss is through recovery point objectives (RPOs) or recovery time objectives (RTOs). RTOs and RPOs are measured in increments of time. The ideal is to have your RPOs and RTOs measured in the smallest units of time possible, and as close to zero as possible.

For example, if it takes you four hours to recover a downed Microsoft Exchange Server VM, you would have an RTO of four hours for Exchange. If you could recover all the data for that Exchange server except for the 12 hours since the most recent backup occurred, then you would have an RTO of 12 hours for Exchange. The challenge is to find backup solutions that help you reduce your RPOs and RTOs for Hyper-V.

The 3-2-1 Rule of backup relates to how many copies of data should be available:

* Make at least **3** copies of your data
* Store those copies on at least **2** different types of physical media
* Store least **1** copy of that data at an offsite location

Following the rule should result in reduced RPOs and RTOs. At the very least, you won't suffer total data loss, and will always be able to recover. The 3-2-1 Rule should be the minimum standard for Hyper-V backup.

It's very likely you'll extend beyond the 3-2-1 Rule to meet the needs of your organization, given the different backup strategies, tools and techniques used now. Today, you don't want just copies of your data or VMs; you want to be able to recover VMs quickly.

Due to the need to recover individual pieces of data or applications, you want to be able to move applications to a specific point in time. This includes having granular recovery down to the individual file level.

All these needs and expectations call for high availability for your virtual environment. Backup, replication, monitoring, and management all play key roles in your data availability strategies, and backup in Hyper-V is all about your ability to recover. The following recommendations will have a big impact on your ability to provide high availability.

The 3-2-1 Rule should be the minimum standard for Hyper-V backup.

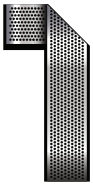

Get a Backup of the Full VM

VMs in Hyper-V are made up of several components, including the .vhdx file, the VM configurations and settings, binaries, snapshots, and more. When you make a backup of a VM, you have to back up everything associated with the VM. At its most basic level, the principles of backup demand that you be able to recover the VM in its entirety. How do you back up the entire VM? There are a couple of methods that work:

Export the VM: In the current version of Hyper-V on Windows Server 2012 R2, Microsoft includes VM Export, a tool that exports and packages a running VM. Once exported, the VM is completely portable.

VM Export gives you a basic ability to copy an entire VM, without interrupting the operation of the VM. It's not recommended to use VM Export as a backup strategy, however, because the only option for recovery is whole VM recovery, and regular exports of entire VMs on any kind of schedule are just not practical. There's also a significant amount of processor, memory and disk workload on the Hyper-V host machine as it creates the export.

Backup through software: Use a good backup software designed for Hyper-V (and VMware). There's a lot of good backup software on the market today that does an excellent job backing up entire VMs and providing granular recovery as well. Choose a solution that meets your needs and learn how to use it.

Every modern software package designed for Hyper-V should be able to give you a complete backup of your VM. Depending on the provider, some solutions may also offer the ability to get incremental, differential, synthetic, reversed incremental with restore points, forever incremental or other backup types, in addition to complete VMs. Advanced scheduling tools, replication functions, compression, deduplication, key storage integrations, WAN acceleration, proxies and other functions should also come with your purchase. The best method is to begin with full VM backups and build your high availability solutions from there.

Don't Use Physical Backup Tools for VMs

Although it may seem obvious, it's important to remember that a VM is not physical. Hyper-V is designed to abstract the hardware from the OS. If you try to use your legacy physical backup tools in your virtual environment—even if the manufacturer says it's supported—you're asking for a lot of pain and suffering.

Instead, look for an agentless virtualization backup tool. Both Microsoft and VMware recommend backing up their VMs without the use of agents on the VM, and both provide the mechanisms for agent-less backup.

Verify Your Backups

In the old backup days, there was a checkbox that said, "Verify backup on completion." The problem is that verification consisted only of a checksum comparison. The verification only told you that you had what was there. If what you backed up didn't work, you had a verified backup that also didn't work; but you wouldn't know it didn't work until you tried to get it to work, leading to frustration.

When choosing backup and recovery software, use a solution that offers verification through starting the VM in an isolated test environment, which will provide a true test of functionality.

Use Replication

Your goal with Hyper-V is to keep your workloads up and running. Replication offers near-continuous data protection, with very low RPO and RTO times. In addition, it allows failover to a replica VM in the event of foreseen or unforeseen events.

With replication, you can build a VM backup and, instead of making incremental updates every four hours, insert the updates every couple of minutes into the VM. This almost-constant synchronization is made to a backup VM ready to take over for the running VM on very little notice.

Just like full VM backup, there are a couple of ways to approach the buildout of replication in your environment:

1. Microsoft replication: Windows Server 2012 R2 has built-in replication capabilities for Hyper-V. It's easy to use and provides failover capabilities that will meet the needs of most IT environments. Replication can be configured onsite, offsite and to the cloud.

2. Replication software: Microsoft wasn't the first to use replication, but this is its first implementation of replication. There are also non-Microsoft solutions that provide advanced replication features like advanced failover and failback scenarios, and integrated backup and replication capabilities.

You will not likely use replication for all of your VMs. Instead, you might pick and choose the VM workloads that carry significant weight in terms of accessibility and access, then configure replication for those workloads. For example, you might configure replication for key authorizations, information processing, database operations and other mission-critical workloads. The point is to have another layer of protection for those key workloads, which will extend high availability beyond traditional backup and recovery.

Application-item-level recovery has become an important capability for today's environment.

Measure Recovery in Minutes—or Seconds

If you're measuring your RTOs in hours, you're doing it wrong. Modern data backup and recovery software should provide VM recovery measured in minutes.

Make sure you can get granular recovery down to the individual file level in your backups. The vast majority of recovery scenarios are file level, and not whole VM recoveries. When a user calls the help desk looking to recover a file, your backup software should provide the capability to recover a file without interrupting a running VM. The recovery time should also be measured in minutes—maybe even seconds.

Application item-level recovery has become an important capability for today's environment, to recover items from within the database files and within your backups for Microsoft Exchange, SharePoint, SQL and Active Directory. These recovery options should also allow for recovery to a running VM without interruption to the application. If a user needs to recover items to an Exchange mailbox, recovery should be just a few clicks and a few seconds away.

Backup in Hyper-V is all about your ability to recover.

The Cloud As Long-Term Backup Archive

Long-term backup archiving is a common regulatory and auditing requirement. In the past, this requirement was often fulfilled through the use of tape-based backup strategies. While tape can still be used in conjunction with VM backup strategies, it's become an outdated way to archive. Instead, use a reputable public cloud provider to lease storage, and move your data off-site to the cloud for long-term backup archiving. Many software backup solutions have built-in options for moving data to well-known cloud providers. Here are five good reasons to move your backup archives to the cloud:

1. Potentially large cost savings compared to tape.
2. Complete backup verification before archiving.
3. Most cloud providers have their own duplication strategy for data to ensure data availability, which will enhance your own on-site strategies.
4. SLAs for cloud archives are as good, or better than, tape.
5. No more tapes to worry about.

Hyper-V is a solid and improving product. If using it to host VMs in your own environment, make sure you know how to make its VMs highly available. It's been said that backup is dead; that may or may not be true. What we do know, though, is that in today's datacenter you need to do much more than just make copies of your VMs.

Gain Visibility into Your VM Environment

Hyper-V is a virtualization platform that runs on physical resources and allocates those shared resources to various VMs. Change is the rule for virtual environments. There are going to be new machines coming online, machines going through various use cases and lifecycles, and machines that will be abandoned. Remember, virtualization abstracts the hardware from the OS. This means you have very little insight from the VM about its performance compared to the resources it has allocated to it. When asking if a VM has enough processor allocated to it, the answer can be difficult to determine without external visibility into the Hyper-V environment.

Management and monitoring frameworks provide essential insight into exactly what's happening with each VM in your environment. Use a set of tools that makes sense based on the size and scope of your Hyper-V implementation. Use these tools to change your troubleshooting style from reactive to proactive. Start to identify issues before they happen and fix problems before they occur. Limit your downtime and data loss with prevention. ■

Chris Henley is senior manager, Microsoft Global Alliance for Veeam Software. He's focused on Windows Server, Hyper-V and Azure solutions from Microsoft. He's spent more than two decades working in technology, including an eight-year stint with Microsoft.

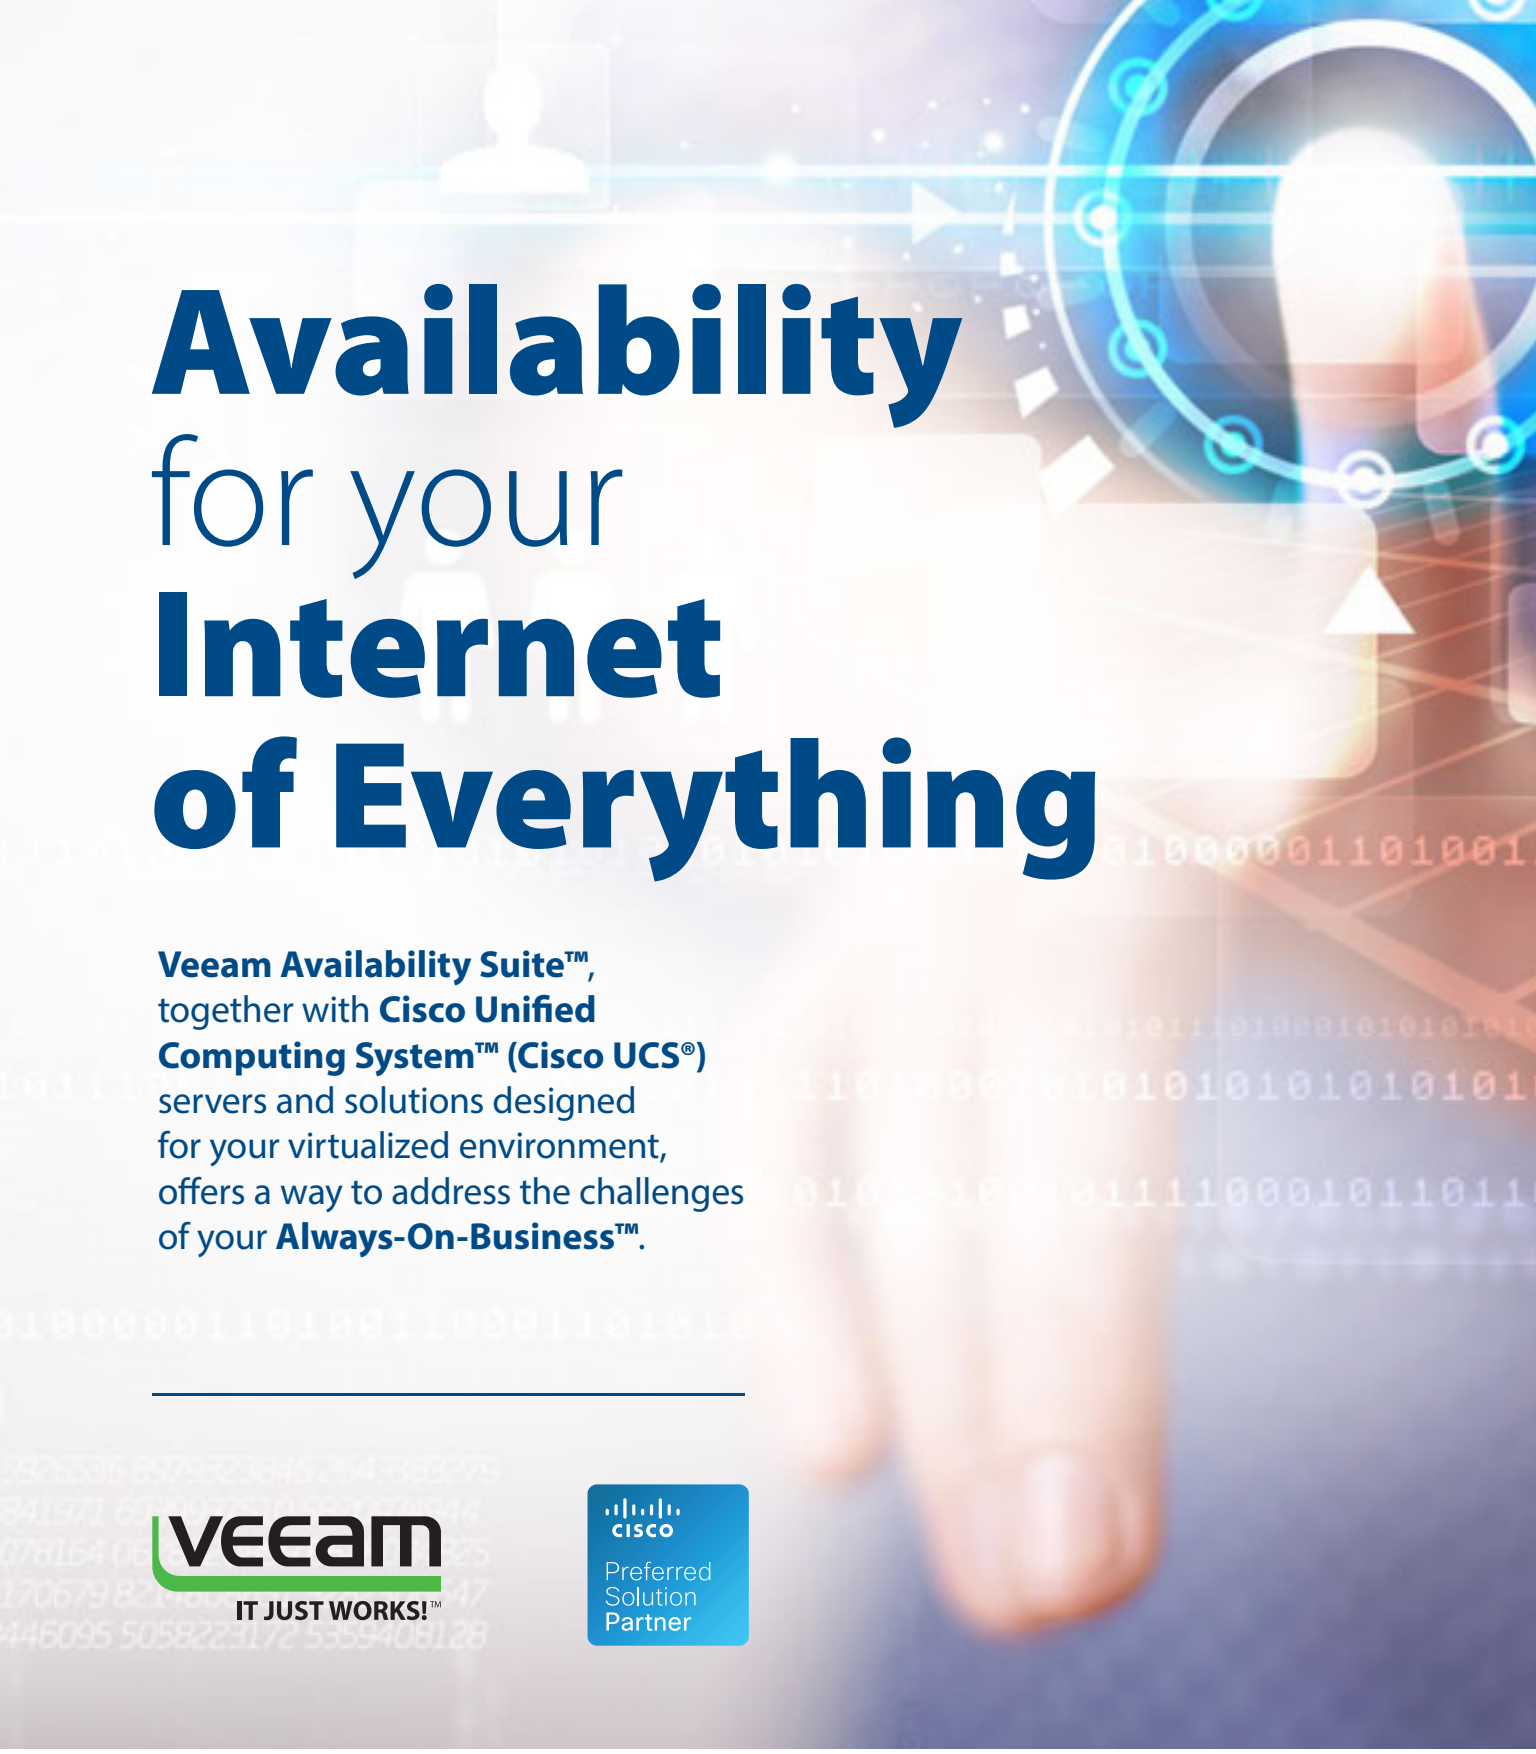

Availability for your Internet of Everything

Veeam Availability Suite™, together with **Cisco Unified Computing System™ (Cisco UCS®)** servers and solutions designed for your virtualized environment, offers a way to address the challenges of your **Always-On-Business™**.

VEEAM
IT JUST WORKS!™

cisco
Preferred Solution Partner

See how at:
veeam.com/veeam-cisco-ucs.html

3,2,1, Backup

The rules for backup are changing, and the cloud is a major reason why. It has significant advantages over the tried-and-true methods, and few downsides.

Luca Dell'Oca

A **common phrase** when talking about data protection and availability best practices is the "3-2-1 Rule." Although it's not new, it's still as valuable as ever.

The 3-2-1 Rule is a simple reminder of how to do backups and ensure proper protection for your data. It means:

- **3** 3 copies of any data
- **2** Saved onto at least 2 different media
- **1** With 1 copy saved off-site

The beauty of this rule is it's technology agnostic—it doesn't tell you which specific technology to use, as long as the final result is accomplished.

Having the first two local copies has become quite easy. Virtualization and modern storage technologies provide plenty of options to design highly available datacenters, where both production data and their local backup copies are protected and quickly usable in the event of a problem.

In the last few years, the shift to disk-based solutions for data protection is becoming more and more common. And now, both

SHUTTERSTOCK.COM

Tape still has limits, however, as a preferred off-site solution.

backup and restore activities have significantly reduced recovery point objective (RPO) and recovery time objective (RTO) values.

With the requirements for local data protection satisfied, it's time to turn to the "1" in the 3-2-1 Rule. Are you currently planning to send copies off-site, so they're protected from major problems in your datacenter?

If you haven't discussed this topic yet, or you're evaluating possible solutions right now, you've probably found there are many options. And all of those options come with pros and cons.

Tape Backup

Tape, for example, is no longer the preferred solution for the first local copy, which is the copy you use for quick restores. Tape backup performance, in terms of RPOs and RTOs, cannot satisfy current business demands, in which restores need to be completed in minutes—not hours. Although tape is certainly not dead, it has changed its role from a backup media to an archive media.

Tape still has limits, however, as a preferred off-site solution. To send tapes off-site, someone has to pull them from a local library, put them in a box and drive somewhere to store them in a secure location. This requires a lot of effort, with numerous manual activities that could go wrong. Tape is far from the ideal solution for modern datacenters that are ruled by automation and cloud-like technologies.

A more modern solution would involve building a second datacenter and deploying hardware and software technologies to create your second backup copy there. This may work, as long as it's financially feasible; but it incurs high capital and operational costs. In addition, your IT team now needs to manage two sites instead of one, with the same number of staff.

Some may suggest that there's a type of datacenter without these issues: a co-location facility. This solution means you would only be removing the capital costs of the second building, because that function would be shifted to a services provider. And all the other costs of owning a secondary location would still exist. That's why it's time to consider the cloud for off-site data.

Cloud Benefits

A cloud storage solution means no capital costs, unlimited storage space and pay as you go. Plus, there's no datacenter to build; an environment that's managed by the cloud services provider; and economical options to pay only for the amount of space used. Among the many cloud services out there, this particular solution is probably the most successful.

One requirement should be paramount when considering the plethora of solutions: It needs to fulfill your company's service-level agreements (SLAs). The fact that a cloud service is inexpensive just isn't enough—you need to know information such as the speed at which a solution can restore your data, and the potential cost of an entire day of non-production as you retrieve the data over the Internet.

Backup-as-a-Service (BaaS) solutions are so easy to set up and operate that services providers sometimes only compete on price, neglecting an emphasis on quality and service levels. Most of the available solutions are nothing more than simple sync-and-share file copies, with versioning of every single file that a company needs to protect.

When it comes to virtualized environments and enterprise-grade solutions, these solutions will start to show their limitations. The biggest one is that backups are managed just like huge files from a services provider's point of view. Imagine this: A full copy of your photo folder is really simple to create and you can do it by yourself by simply copying and pasting photos into a remote storage. That's what the simple solutions do.

File Format Issues

When it comes to enterprises, data is often in different formats; not just files, but also e-mail, shared projects, Active Directory objects and database tables. Most of the time, this data is all inside a virtual machine (VM). Image-based backups save entire VMs, but the items that need to be restored are inside those VMs. A simple backup done with a file copy isn't effective in these situations, because once the backup file is in the cloud, starting a restore means first locally retrieving the entire backup file itself.

In the meantime, your data creation continues to skyrocket—so all data protection activity involves an increasing amount of data. The problem is that increasing the bandwidth to deal with the the data deluge is a challenge.

It Does Not Compute

To overcome these limits, compute capabilites are needed. Compute capabilities applied to data availability, both on-premises and in the cloud, means a software component

Virtualization and modern storage technologies provide plenty of options to design highly available datacenters.

is running at the other end at the services provider. This component is capable of understanding the content of your backup sets and helping the on-premise components optimize both the uploads and retrieval operations.

On the way out of your infrastructure, you want compute to run optimization techniques on outgoing data. Because data is inside a VM most of the time, you don't want to waste time having agents inside a server. Modern technologies are able to work at the hypervisor layer and create image-based copies of those VMs.

Because of its increasing size, you want to leverage deduplication and compression as much as possible, so that the bandwidth for your chosen cloud solution is as optimized as possible.

A solution using local compute resources is only part of the answer; you also need compute capabilities at the services provider level. Again, in a sync-and-share solution, the services provider is only giving you a large storage space where you drop your data and indexing, and searches are done remotely by the customer component.

But what if you need to restore a single e-mail from the remote copy of your entire mail server? Thanks to modern solutions, this is an easy task when done using a local backup set. Why should moving to the cloud completely remove this capability and leave you with a "dump" storage solution? Do you really want to download the entire backup of your e-mail server just to open it and extract a single e-mail?

That's what happens most of the time. It's even worse when you add encryption to the picture; without compute capabilities, the services provider can't open encrypted backups. They have to send the entire backup to you, and your compute systems have to decrypt them, browse the content and, finally, restore the needed items.

A Two-Way Compute Street

The result is that RTO performance will be abysmal, which is why compute at both ends is necessary. With a solution capable of opening a remote copy directly at the services provider side, regardless of whether it's encrypted, compressed or deduplicated, you can browse its content without actually moving any data over the wire. This is because the interaction between the compute components at the two ends of communication do it for you.

Compute activities happen locally at the backup file's location at the services provider—and only when the needed piece of information is identified. The local compute resource receives only the minimum amount of data needed to complete a restore; if it's an application item instead of a file, the solution at the services provider needs to be able to understand the content of those remote copies, whether it's a Microsoft Exchange database, SQL table, Active Directory object and so on.

RPOs and RTOs

The services provider also needs to deploy some smart technology to leverage remote compute capabilities, which will save bandwidth and time. It's not wise to delegate data availability needs to a services provider if the SLAs end up being so much worse that savings alone cannot justify the increased RPOs and RTOs.

A solution with the ability to leverage compute with cloud storage allows you to effectively improve RTO and RPO values. Plus, the right solution will allow you to add additional capabilities to your solution, like application item restores. Consider all these factors when deciding on a backup solution in the cloud. ■

Luca Dell'Oca (vExpert, VCAP-DCD, CISSP) is an EMEA evangelist for Veeam Software based in Italy. A popular blogger and an active member of the virtualization community, Dell'Oca's career started in information security before focusing on virtualization. His main areas of expertise are VMware and storage design, with a deep focus on cloud services providers and large enterprises. You can follow him on Twitter @dellock6 or @Veeam.

How vSphere 6.0 Will

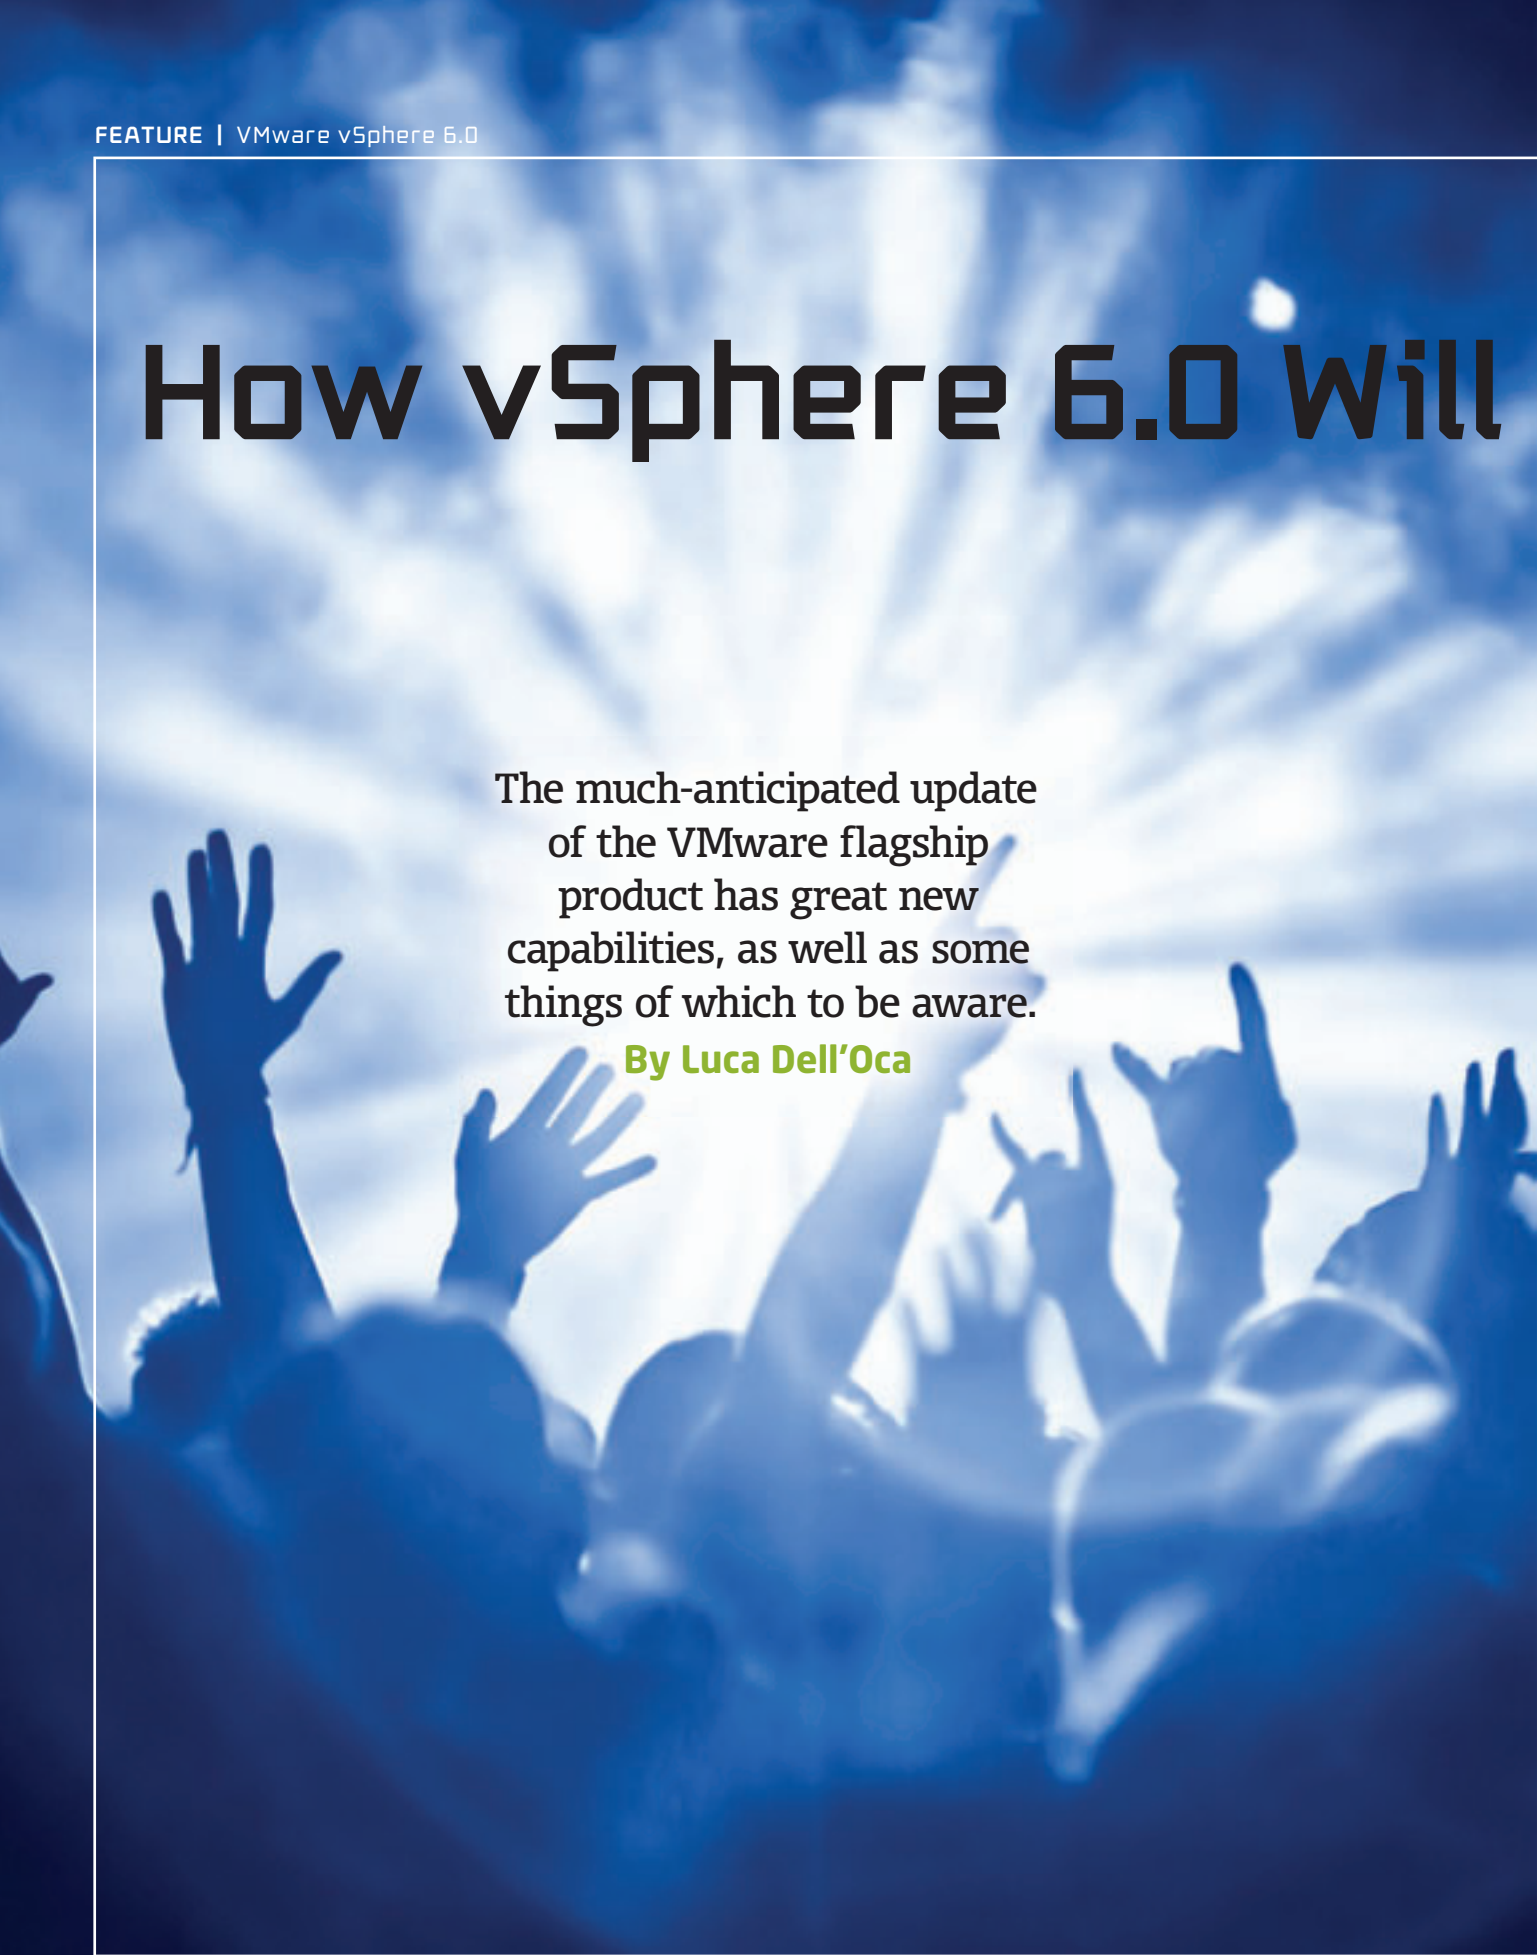

The much-anticipated update of the VMware flagship product has great new capabilities, as well as some things of which to be aware.

By Luca Dell'Oca

ROCK YOUR WORLD

A few months ago, VMware Inc. released vSphere 6.0. Whenever a new version of a solution comes out, the upgrade questions start to be asked: "Do I really need some of these new features?" "Is it best to go with in-place upgrades?" "What about upgrading using the 'rip-and-reinstall' technique?"

Today, vSphere is an entire ecosystem, rather than a single software solution. As such, there are several components contributing to the overall datacenter environment. All these components should be carefully checked before upgrading.

Hardware Check

Starting from the lowest level, there's the hardware. Each new vSphere release brings a new and updated hardware compatibility list (HCL)—a long list of certified hardware components guaranteed to work with the new version. Before even thinking about an upgrade, you should carefully check all your components and verify them with the HCL. Otherwise, an upgrade might fail or hardware might not have a driver after the upgrade, putting your infrastructure in an unsupported situation.

The ecosystem is also made of additional software solutions running side-by-side with vSphere, such as data protection. Before upgrading, it's vital to verify your chosen solution supports vSphere 6.0. Otherwise, you'll find yourself with a shiny new vSphere 6.0 environment and no way to protect the virtual machines (VMs). Plus, your data protection solution needs to fully support the new features of vSphere 6.0.

SHUTTERSTOCK.COM

With a maximum of 10,000 VMs and 1,000 hosts, many users are likely to migrate to vCenter appliance.

Keep in mind that "simple support" of existing features isn't the definition of a "supported solution." You must make sure that virtual volumes (VVOLs), for example—a new feature in vSphere 6.0—are supported, and that you can protect fault-tolerant VMs.

Once an entire environment has been checked and you've confirmed it can be successfully upgraded, the next step is to evaluate the new vSphere features that interest you, and find out what you need to know and understand about these features. Some of the enhancements are no-brainers, because they immediately become available when you upgrade. Other features require proper planning before deciding if you want to introduce them into your environment. Here's a look at some of them.

Increased Scalability

As is typical with previous versions of other software, vSphere 6.0 has higher limits for almost every parameter, including:

- More memory and CPU per host
- More memory and CPU per VM
- More VMs per host and cluster

Keep in mind that version 5.5 also has really high limits, so unless you have specific needs like a 128 CPU VM, this is just a "nice-to-have" improvement that comes with the upgrade. The support for 12TB of RAM, however, is available only if you buy one of the new servers with this capacity to replace one of the servers you bought for vSphere 5.5.

Long Distance vMotion

A perfect example of a new, free feature is Long Distance vMotion. As soon as you've upgraded to vSphere 6.0 and have the proper license, you'll have this new capability with no need to change anything else. And as long as the connection between your source and destination is below 150 ms latency, you can move a running VM with no downtime, even if there's no common component between the two sites or a single vCenter server. This is a feature I really like because it opens up great opportunities for datacenter migrations.

Fault Tolerance

Another new, no-brainer feature is Fault Tolerance (FT). FT has two capabilities that make it extremely useful:

1. **Support for VMs with up to 4 CPUs.** The previous limit of 1 CPU made FT difficult to use. If you have a critical VM that you want to protect by using FT to guarantee zero downtime, chances are this VM doesn't have just 1 CPU. Now, with 4 CPU support, many more workloads can be protected.

   However, FT doesn't solve all availability problems. For instance, when you replicate any VM CPU instructions to its twin, any corruption, program error or crash is also replicated.
2. **Snapshots for FT VMs.** Previously, there was no way to snapshot a VM protected by FT. Because data protection solutions for vSphere leverage snapshots, this meant you couldn't back up the VM from the hypervisor. vSphere 6.0 enables FT VMs to be snapshotted and properly backed up. By mixing FT for downtime avoidance with an availability solution to create protected copies of the same VM, the VM will be protected better and high service-level agreeements (SLAs) can be guaranteed.

vCenter Appliance

Another interesting improvement in vSphere 6.0 is the vCenter appliance, which was available in previous vSphere versions as an alternative to the full vCenter installation. In this latest release, the appliance has the same exact features and maximums of the full installation, and is completely supported by VMware in production environments.

With a maximum of 10,000 VMs and 1,000 hosts, many users are likely to migrate to the appliance. With the appliance, you can even remove any dependency on external and additional licenses for OS and the back-end database, for a ready-to-go solution that can be quickly deployed and used.

Obviously, for existing situations where the full vCenter installation is used, a potential drawback might be the need to migrate to the appliance. While this may be an easy task for small environments, the number of configurations that need to be recreated in the appliance for a large datacenter are too much.

In addition, some issues can't be solved, like a VM receiving a new ID and being backed up from scratch by any data protection solution. VMware has released a solution in the form of the VCS to VCVA Converter Appliance. The Converter, which started life as a "Fling," or unsupported product, has

FlexPod®

Is Your Application Data Protected?

FlexPod and Veeam

Veeam delivers application availability for the resilient converged FlexPod® infrastructure platform

Veeam adds:

- High-speed recovery
- Application-consistent recovery points
- Data loss avoidance
- Verified protection
- Cloud readiness

Learn more at
vee.am/netapp

cisco NetApp®

VEEAM IT JUST WORKS!™ Data protection software provider

VVOLs is, by far, the biggest innovation of vSphere 6.

been officially released. It allows you to migrate a vCenter installation to the appliance.

Virtual Volumes

There are additional features and new capabilities in vSphere 6.0, like the new VSAN 2.0 and the promising Instant Clones. Without a doubt, however, the most awaited feature of the latest release is Virtual Volumes (VVOLs).

So far, a VM has always been represented in the storage as a bunch of files stored in a common datastore. Some files are for the configuration parameters, and the large binary files represent the virtual disks. Because of the limits of storage protocols (mainly, the 256 maximum volumes per host), it's been impossible to assign a single LUN carved out of storage to a single VM.

One LUN—once prepared and exposed to vSphere to become a datastore—is the location for multiple VMs, all stored and executed from the same position. This has always been a necessary compromise, which leads to a lack of granularity in storage management.

When applying a policy to a single VM, you're actually applying it to all the VMs stored in the same datastore, because the storage itself has no visibility into the content of a LUN (with a few exceptions leveraging proprietary solutions). This limitation has always impacted day-to-day management. Think about storage snapshots, which forced you to snapshot an entire LUN, even if you were only interested in one VM.

With VVOLs, the new level of granularity is a single virtual disk. Finally, both the storage system and vSphere are able to see the same objects and, ultimately, talk the same language. Together, with this level of detail, come the storage capabilities and features available at the storage layer and exposed by the storage itself to vSphere to create per-VM policies and execute point actions—without impacting any other VM. For example, if you take a storage snapshot, only the chosen virtual disk will be cloned. This opens up a whole new world for storage management in virtualized environments; it is, by far, the biggest innovation of vSphere 6.0.

By applying different policies to specific VMs or even single disks, VVOLs paves the way toward a policy-based datacenter from a storage point of view. It's been designed to be open to third-party vendors via APIs, for data protection at the virtual disk level. Think about the possibility of protecting VMs directly at the storage layer, with minimum-to-no impact at the hypervisor layer.

VVOLs is a fascinating technology that's sure to suceed in the long run. For now, you just need to remember that because it's the first iteration of a completely new storage technology, it has some limits.

For instance, your storage array needs to specifically support VVOLs and be certified by VMware. There's no workaround for this, and even though the beta period of VVOLs was really long, now that it's generally available, not all storage vendors are ready to support it.

Also, if you own a storage array, it can't be updated to support VVOLs because it won't have enough compute power to manage the multitude of objects created by VVOLs. You need to have the right storage. If you don't have the right storage, you'll either need to update your storage firmware (if your vendor has already released the update) or buy a new model.

In addition, even with the right array, some vSphere features aren't supported yet. For example, storage replication isn't available with policies you can apply, so you won't be able to define a protection policy for critical workloads directly on VVOLs. FT (both the old single CPU or the new SMP one) is also unsupported, as are other features such as Storage I/O.

These limits will undoubtedly be removed in future versions of VVOLs, so vSphere can use the "old" datastore using VMFS or NFS side-by-side with VVOLs-enabled volumes. For users without storage vMotion licenses, VVOLs can be safely evaluated and adopted in any vSphere environment.

As with any major upgrade of an existing software solution, there are pros, cons, limits and notes that every user should carefully evaluate before proceding to the upgrade. But even with the limits (like any new technology), vSphere 6.0 has some interesting and useful features that will greatly improve the performance of a VMware environment. ■

Luca Dell'Oca (vExpert, VCAP-DCD, CISSP) is an EMEA evangelist for Veeam Software based in Italy. A popular blogger and an active member of the virtualization community, Dell'Oca's career started in information security before focusing on virtualization. His main areas of expertise are VMware and storage design, with a deep focus on cloud services providers and large enterprises. You can follow him on Twitter @dellock6 or @Veeam.

Be an IT Hero in just 15 minutes!

- **Jump** into a new infrastructure
- **Protect** your data
- **Dash** to recover files
- **Stretch** your virtualized capabilities
- **Brew** a mean pot of java

hp | veeam
IT JUST WORKS!™

In less than 15 minutes, this hero deployed the **HP ConvergedSystem 200-HC StoreVirtual** and protected his data using integrated snapshot intelligence with **Veeam**

Learn more at: vee.am/hp

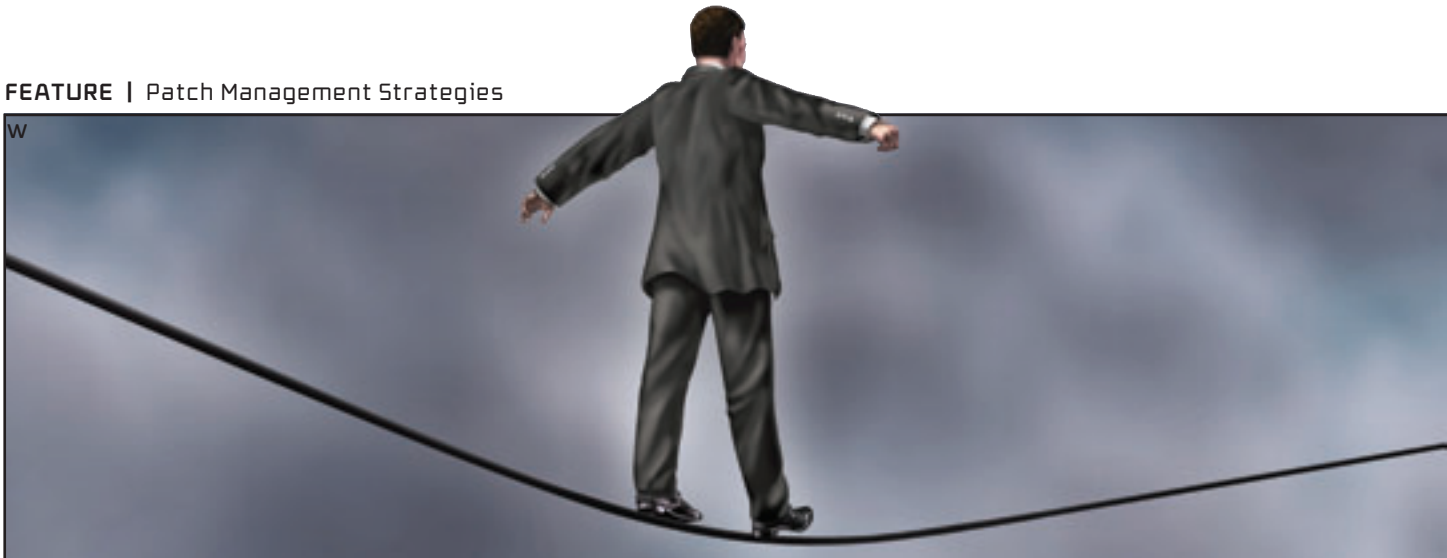

Taking the FEAR Out of Patches and Updates

When it comes to this crucial aspect of your datacenter, failing to plan could invite disaster.

By Mike Resseler

Every IT administrator knows the fear of patching or upgrading OSes and applications. Most have at least one story of something going wrong. But something going wrong isn't even the biggest issue.

The biggest issue is that the availability of your workload isn't present, forcing you to work quickly and under stress to get things fixed as soon as possible. Plus, you have to decide if you're going to try to fix the issue with the upgrade, patch it or revert to the previous version of the software. That brings up another worry—if you revert, will things be back to normal?

How do you find the root cause of the issue? What exactly happened? How do you make sure it won't happen again? Why wasn't this foreseen by the change management plan? These are questions an admin asks before starting an upgrade or patching. The ideal, of course, is to not have to ask these questions at all, through proper preparation.

Upgrades, patches and hotfixes are designed to resolve problems, add additional functionality and fix security issues. And although many admins still live by the "If it works, don't break it" motto, things are different today when it comes to keeping modern environments healthy and up-to-date.

Upgrading and Patching Tips

- Use a change plan: Sure, it's a lot more fun and a lot less time-consuming to simply start patching or upgrading. A change plan—workflows, steps to take, risks involved, affected systems and so on—is to assess minor or major changes in an IT environment. It can be in place for everything from large upgrades to small configuration changes. Having a change plan with the necessary control mechanisms can prevent many potential issues. To make sure you don't miss anything crucial, it's also important to understand which other workloads are dependent (both upstream and downstream) on a workload.
- Read the release notes: Every upgrade, patch and hotfix comes with release notes. Read them thoroughly to see what changes will be made and any known issues. Also, make sure you understand the reason for applying this change.
- Test, test, then test some more: Testing is crucial for knowing if a process will work, but it's often overlooked; many people simply install something and it if looks like it's running, they assume their job's done. Nope; it's also important to determine if the process breaks after a few days. You also need testing to determine if other process-dependent services break. Finally, test other services that seem unrelated, since they may actually be more related than you realize.
- Have a failback plan: Every upgrade, patch or hotfix should be able to uninstall itself, but this isn't always the case. Unfortunately, the world of upgrading and patching is far from perfect, and uninstalling and reverting aren't always possible. Worse, if things go very badly, failback won't work, either. What are you going to do when one of this situations appear? Will you reinstall the workload to the previous situation? Are you working temporarily with snapshots? Will you use your backups to do a disaster recovery? Whatever your solution is, make sure it's tested up front.
- Inform the stakeholders and the help desk: Something very simple, yet so often forgotten, is the importance of informing the parties who will be affected by an issue. Workloads getting upgraded or patched, and the dependent workloads (see your change plan), all have stakeholders you'll need to inform. You will, of course, want to make sure that your help desk is informed, as well, so they can know you're working on the issue. This will prevent the help desk from trying to fix something because they got an alarm or were informed by an end user.
- If possible, tackle non-critical servers first: This isn't a law set in stone, but if possible, try to work on your non-critical servers first and document the upgrade, assuming you haven't already done so through other means. Also, make sure you determine what non-critical servers or non-production servers are. What seems non-critical may, in fact, put an entire business unit out of business if it becomes unavailable. Make sure all your services and workloads are well documented, so that you know what really matters and what can be tolerable if it's missing for a longer period. For example, if you have five domain controllers, you might be able to run a few days with just four. Or a business Web site that runs on multiple servers might be able to handle the load with one fewer server. It's crucial to understand what's possible and what is not.

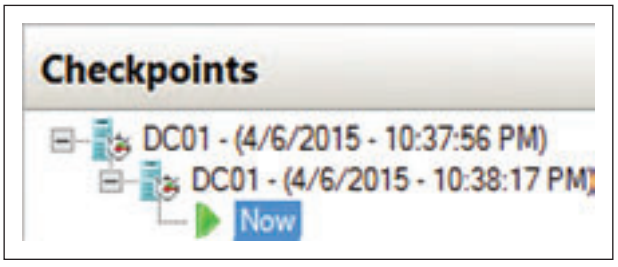

Figure 1. Example of a Checkpoints tree.

SHUTTERSTOCK.COM

Let's look at four patching methods, and their strengths and weaknesses.

1. Wait until others have updated

Waiting until others have updated is probably one of the most common upgrading or patching tips from IT pros. Wait a few weeks, at least, and also do some Internet research to find out if other people are encountering issues. Luckily, not everyone does this, or we'd be stuck with the chicken-and-egg syndrome.

In any event, waiting isn't a completely failsafe system. A patch or upgrade that succeeds in one environment can fail in another environment, for various reasons. And what about those applications specific to your industry? Will you know what's going on with them if the upgrade fails? When it concerns a security fix, waiting is certainly not advised, because your infrastructure remains exposed to the threat while waiting.

The waiting solution (although calling it a "solution" might be a bit overkill) may remove upgrade and patch issues, but it certainly won't protect you against specific issues in your environment. In addition, waiting may only be possible for upgrades and patches from large vendors. Finding this type of information for specific workloads might prove difficult. Finally, waiting can create a false sense of safety because you haven't tested against your own, unique environment.

2. Have a separate test environment

Having a separate test environment used to be—and probably for many out there still is—the dream of any IT pro. This environment is an exact copy of your production environment, so you can try out upgrades and patches.

Unfortunately, not many IT pros have this luxury. For most, it's very difficult to have a test environment; and while others may have a small test environment, it likely isn't even close to the real environment.

Even if your test environment isn't a perfect duplicate, though, it's still wise to do some tests and see what happens. A test environment won't protect you against every possible issue, but it will allow you to see major problems ahead of time.

3. Snapshots

Using snapshots (or checkpoints, as shown in **Figure 1**) on your virtual machines (VMs) is another way to try out a patch, upgrade or update. If something goes wrong, a snapshot will get you back in operation because it allows you to simply go back to the time of the snapshot. This functionality is embedded in every hypervisor, and can work for your testing needs.

However, this method comes with a warning: It's really not advisable to run snapshots for a long time on your production servers, because it can lead to potential commit problems (when you apply the checkpoint or snapshot) and will slow down your environment, so you'll need additional resources (such as more IOPs) to handle this. You shouldn't test for longer than a couple of days.

If using snapshots, remember to get the VM in some sort of offline state, which will cause downtime. The reason is that you might need to revert if the update didn't install as foreseen, which would result in losing all data that went to that server—starting from the moment of the snapshot until you revert. That might be difficult to explain to the application owners. And while it's entirely possible to use a snapshot with no downtime, you risk losing valuable company data.

There are certainly some use cases for working with snapshots, but they won't work all the time, and each situation will depend on which workload you're updating or patching. It's also important to make sure that snapshots don't live too long (or worse, get forgotten). Whenever you use these snapshots, have a solution in place that delivers full visibility in your environment so you can detect potential issues.

4. Backups

The closest simulation to your true production environment is the most recent recovery points of your backups. Instead of letting those backups just sit there and hoping you never need to use them, you can put backups to work and use them for testing upgrades, patches, updates and more. They can even be used to train new administrators.

Does this mean you need to recover all of your VMs and put them in an isolated environment? No. By leveraging technology that allows you to start your backups from the backup location (fully quarantined) you can test every backup automatically, including adding scripts to test if certain workloads are up, running and responding.

By leveraging this type of technology, you can also test upgrades, updates, patches and so on. Simply start such a job and leave it running for as long as you want to test. The VMs that run are real production VMs with production data and fully isolated from your environment, so you have a real-life situation to test all the work you need to do. And because you can be sure that the actual backup files aren't harmed when you close the environment, you can be assured that you still have the actual backup files in case of a disaster.

If something goes wrong, a snapshot will get you back in operation because it allows you to simply go back to the time of the snapshot.

Using this technology has some advantages:

- You can test everything out, record it or take screenshots and create the upgrade/update/patch manual safely.
- After the upgrades, you can leave the lab running and test (in combination with other needed workloads) if everything is still running as expected.
- You will have the possibility to examine an uninstall plan.
- You can prove testing and more for your change plan activities.
- If something goes wrong, stopping the virtual lab is enough. Then you can power it up again and start all over.

There are, of course, some small downsides. You're starting these VMs from the backup files and you probably won't have the same performance as in your production VMs. You'll also be using some of your compute resources for these tests. But these are the only downsides to leveraging a technology such as virtual labs.

Patching, updating and upgrading have always been challenging. Even for small patches, it's a good idea to create best practices and testing mechanisms to ensure availability of your workloads at all times. I've looked at some best practices and described four possibilities or strategies around testing patches. Depending on your situation, you can use one or combine different testing techniques. But no matter what your situation is, testing should always be included in the plan whenever applying a patch, update or upgrade. ■

Mike Resseler is a product strategy specialist for Veeam Software. Focused on technologies around Microsoft Hyper-V and System Center and with years of experience in the field, Resseler presents on many occasions at large events such as MMS, TechEd and TechDays. He has been awarded the MVP for System Center Cloud and Datacenter Management since 2010. His major hobby is discussing and developing solid disaster recovery scenarios. Additionally, he has enterprise-class experience in private cloud architecture and deployment, with marked focus on protection from the bottom to the top, and he holds certifications in many Microsoft technologies such as MCITP.

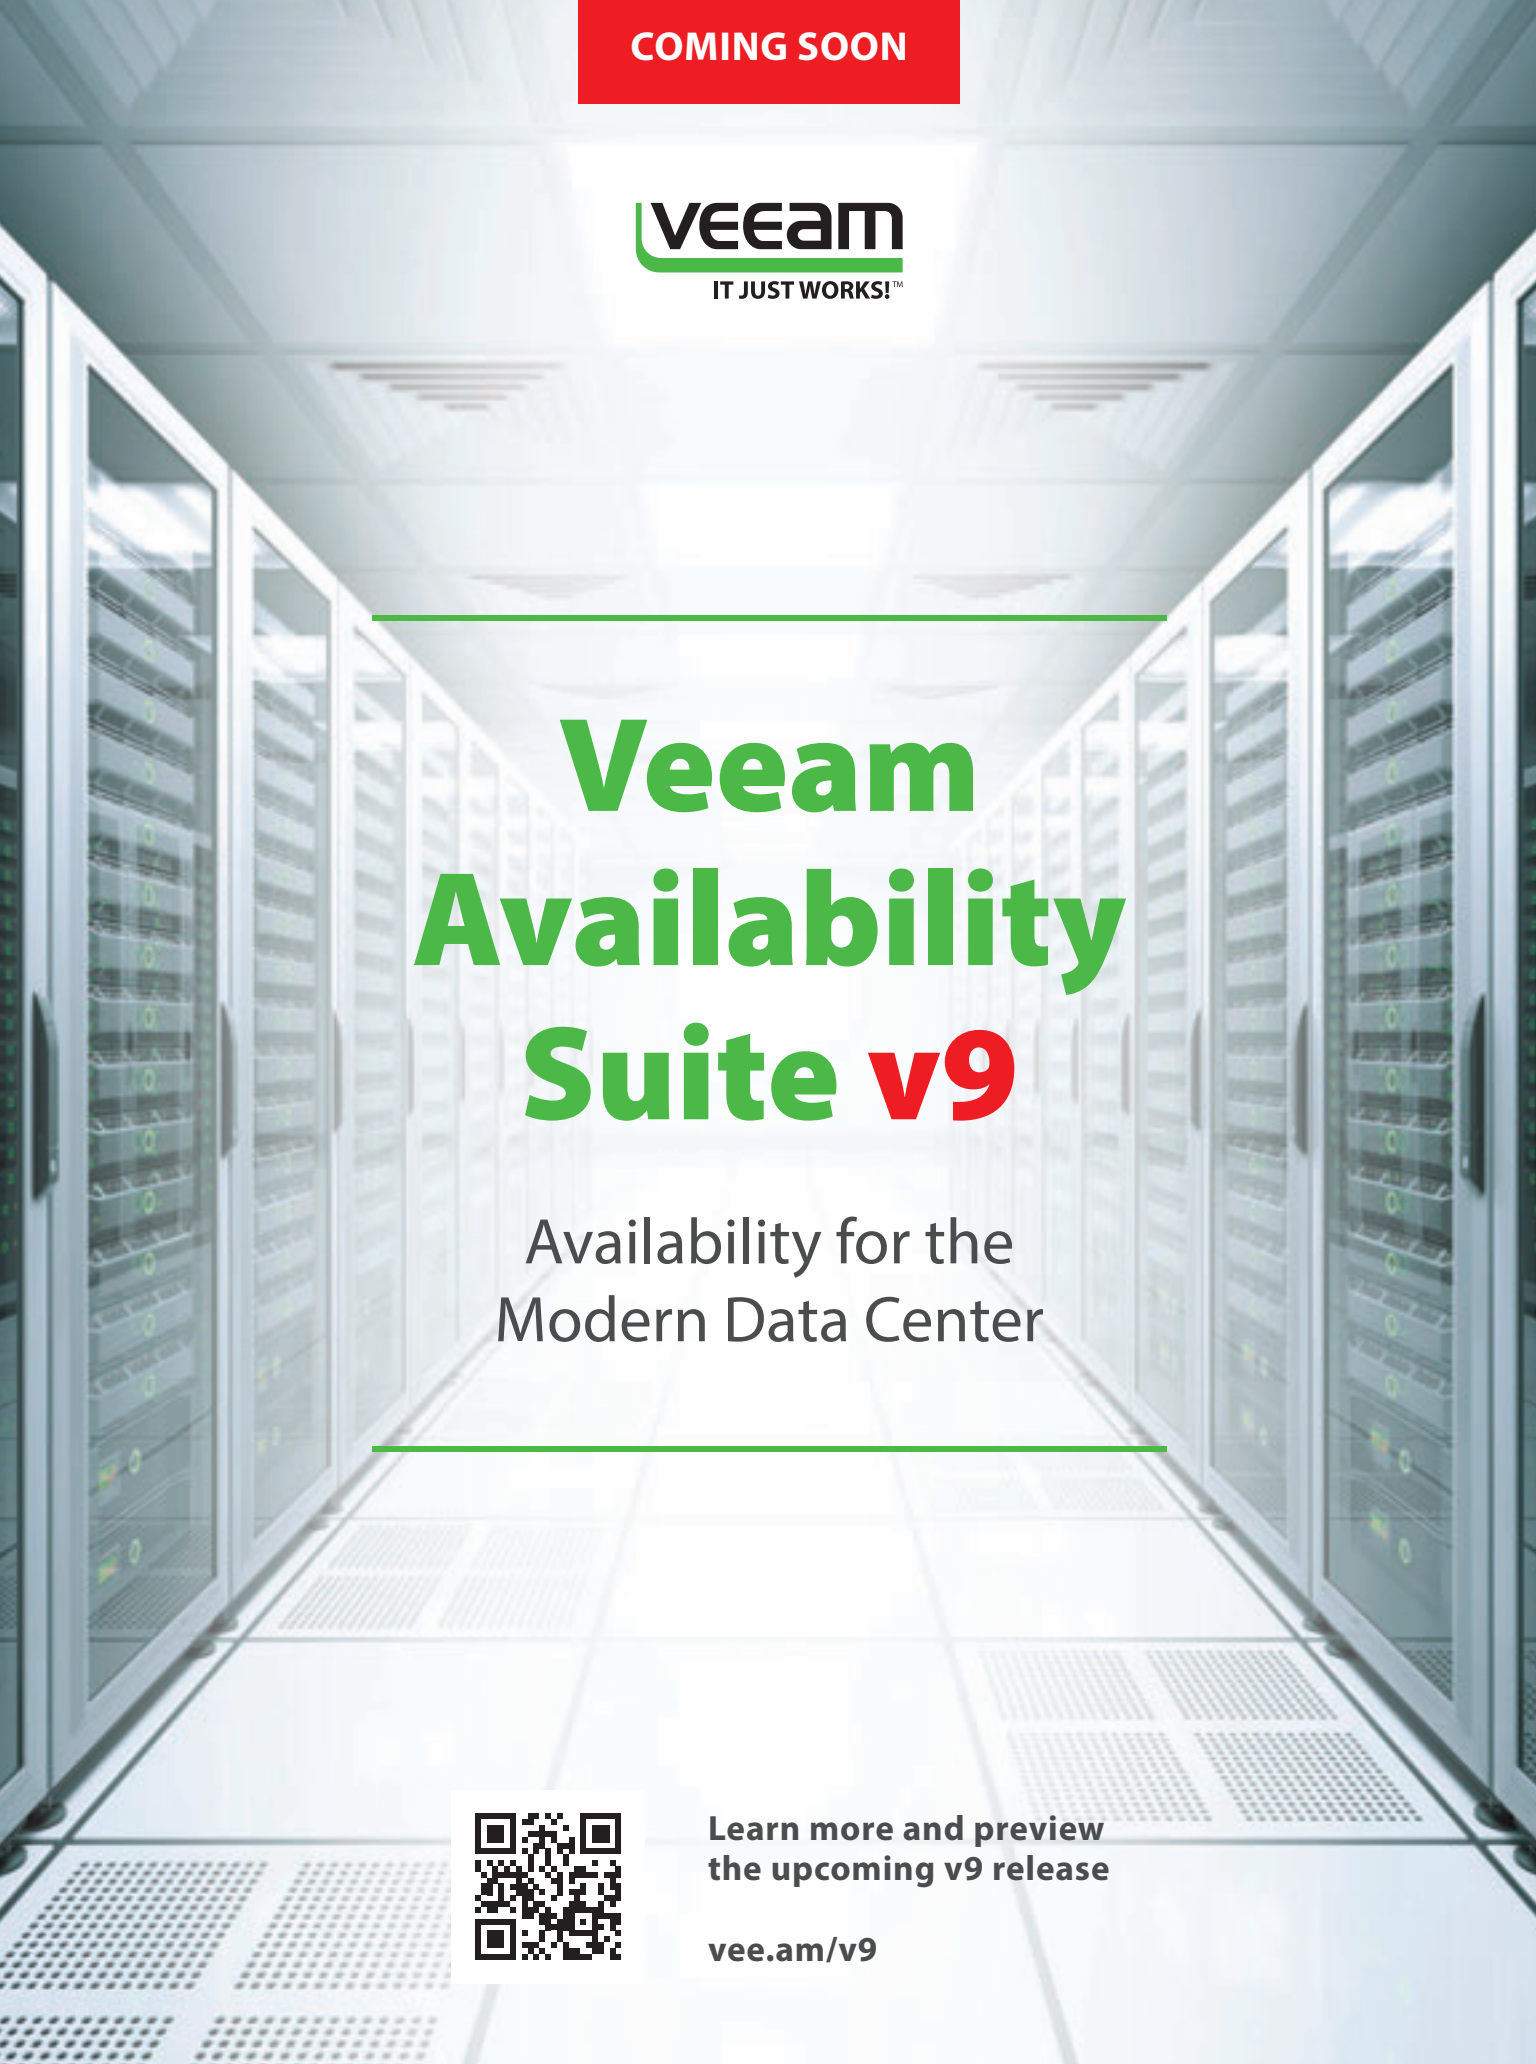
COMING SOON
VEEAM
IT JUST WORKS!™
Veeam Availability Suite v9
Availability for the Modern Data Center
Learn more and preview the upcoming v9 release
vee.am/v9

Windows Server 2003 Migrations Made Easy

Time is quickly running out to upgrade your Windows Server 2003 machines. When you upgrade, think "Cloud First" to smooth out the migration.

By Chris Henley

Windows Server 2003 may have been the quintessential client/server OS. The innovations and capabilities of Windows Server 2003 were, and in some cases still are, amazing: Among other things, Microsoft introduced the Volume Shadow Copy Service (VSS). VSS is still used with Hyper-V for things like image-based backups. The best indicator of the success of Windows Server 2003 is that Microsoft estimates there are tens of millions of instances, both physical and virtual, still running around the world.

But on July 15, 2015, those Windows Server 2003 machines will no longer have any support. If you're still running them, your business won't pass a security or compliance audit. The time has come for change.

The Challenge

Windows Server 2003 is 12 years old, and Windows Server 2012 R2 is the current Windows server OS. This means you'll be migrating from a 32-bit, physical, client/server OS to a 64-bit, virtual cloud OS. There's no direct upgrade path from Windows Server 2003 to Windows Server 2012 R2, and the changes couldn't be more significant. It's possible that you may have a 64-bit version of Windows Server 2003, or you could be running Windows Server 2003 as a virtual machine (VM), which will certainly aid in the migration process. That said, more than half of existing implementations are still the old physical configurations.

Think Cloud-First

Windows Server 2003 end of support (EOS) is a big opportunity for IT; the workloads and infrastructure designs of yesterday aren't typically the best way to provide the same services now. When assessing potential destinations for existing workloads, think cloud-first.

Public cloud solutions, such as Microsoft Office 365, often offer the best solution for Exchange e-mail services and are seemingly a no-brainer when compared to the old on-site Exchange Server running on physical hardware. Services like SQL are often more desirable on Microsoft Azure than traditional on-site implementations. Software-as-a-Service (SaaS) offerings from Microsoft and other online providers can provide significant cost savings and broad capabilities for your business. Infrastructure-as-a-Service (IaaS) offerings enable shops to run their own servers as VMs hosted on public cloud hardware, allowing better cost control and providing elasticity to your workloads.

If you're still on the fence about public cloud, please keep in mind that Windows Server 2012 R2 is what Microsoft says is "at the heart of Cloud OS," and you'll be using a private cloud design as you build your own networks—in reality, you can't escape the cloud. Finally, there will likely be some crossover between your leased public cloud and your own private on-premises cloud. Plan first, then move into to the method phase.

The Method

Because there's no direct upgrade from Windows Server 2003 to Windows Server 2012 R2, the method that most will follow for workloads that stay on-site are as follows:

1. Convert Windows Server 2003 from a physical to a VM. This can be done with the free Microsoft Disk2vhd tool, or with other third-party tools. Converting to a VM will also take care of the challenges of the 32-bit to 64-bit platform conversion, due to the hardware abstraction from Hyper-V or VMware.
2. VM backup and workload testing comes next. It's critically important to use the virtual platform as a basis to get a solid backup with granular recovery capabilities.
3. Install Windows Server 2012 R2 and configure the necessary roles and features to host the workloads that will be migrated from Windows Server 2003.
4. Perform the side-by-side migration from Windows Server 2003 to Windows Server 2012 R2.
   NOTE: Always use the migration guides located at **technet.microsoft.com** for the specific workload you're migrating.
5. Test, test, test!
6. Back up the newly configured servers.
7. Decommission the Windows Server 2003 boxes.

It's also possible that you'll be migrating not only OSes, but application versions.

There's no direct upgrade path from Windows Server 2003 to Windows Server 2012 R2, and the changes couldn't be more significant.

You've Got a Friend

Remember that at this stage of the Windows Server 2003 EOS lifecycle, you're not in this migration process alone. Millions of Windows Server 2003 boxes have been migrated, with almost all possible configurations and workloads. Make sure you find resources with experience that can help you with your migrations. Use Microsoft partners and third-party software to make the process easier.

If you're getting a late start on your Windows Server 2003 migration, resist the temptation to try and skip steps in order to finish faster. The migration process will take time and result in some downtime. Taking shortcuts is risky. Your business availability will be impacted during this process, so try to limit the impact as much as possible by following the recommendations.

Recognize the opportunities to adopt new technologies to provide competitive advantages to your business. Advances in virtualization, data protection and the cloud should be a part of every Windows Server 2003 migration.

Finally, as you migrate from Windows Server 2003 to Windows Server 2012 R2, don't forget that you're now running your own private cloud. Be sure to take advantage of the free training opportunities at Microsoft Virtual Academy, so you and your team can understand all the benefits of private cloud design. Migrating from Windows Server 2003 to Windows Server 2012 R2 is one of the biggest opportunities to hit your network in a long time. Enjoy it. ■

Chris Henley is senior manager, Microsoft Global Alliance for Veeam Software. He's focused on Windows Server, Hyper-V and Azure solutions from Microsoft. He's spent more than two decades working in technology, including an eight-year stint with Microsoft.

There Is No One-Size Recovery

Different applications need different tools. Do you have the right ones for your apps?

By Rick Vanover

When it comes to the applications in the datacenter, there are key problems around having the right type of recovery for specific problems that might occur. As an example, take one of the more common applications: Microsoft Exchange. When it comes to protecting Exchange, the story is quite simple, as there are scores of best-practice resources to not only virtualize Exchange, but protect it, as well.

But when the time comes to perform a recovery, what's the best course of action? Restoring an entire Exchange Server isn't a good idea in many situations, and not everyone has Exchange experts on staff who can handle entire Exchange databases, much less logs.

Fits-All

Recovering Exchange

There's an easier way, and in some cases you can do it for free. Veeam Explorer for Exchange (**bit.ly/1DuZVUX**) is the first one to check out. Veeam Explorer for Exchange provides five key recovery and e-discovery abilities:

1. Open content from an Exchange database (such as an e-mail or calendar appointment)
2. Export items to a PST file
3. Forward items as an e-mail attachment
4. Save items as MSG files
5. Search items through the database

These abilities can help answer the important question of who sent something with sensitive data out of the organization, should that type of e-discovery situation arise. These features, included in both the free and paid editions (Enterprise and higher) of Veeam Availability Suite, allow items to be restored directly back to the mailbox (or a different mailbox). This is in addition to whole virtual machine (VM) recovery, disk recovery and file recovery (to mount an Exchange database, for example). It's good to have options when things don't go as planned.

Recovering SQL Server

The story doesn't stop with Exchange. Another common datacenter application is Microsoft SQL Server. SQL databases are much like other applications, in which a whole system restore may not be the best course of action; and like Exchange, not everyone has access to a SQL database admin. Additionally, there are many SQL database administrators who utilize SQL Server Maintenance Plans or SQL Server Agent jobs that perform daily database full backups and truncate logs every two hours or so (BAK and TRN files); but rarely have those scripts been used for recovery.

One alternative to this situation is Veeam Explorer for SQL Server (**bit.ly/1GBRFr1**). Even with Veeam Backup Free Edition, SQL can benefit from enhanced availability in ways that might not have been possible before.

From a Veeam Backup image, an admin can restore a SQL database back to its original location or direct it to another SQL system. Additionally, the database can be exported as a script from the restore point taken with Veeam Backup Free Edition.

Paid editions of Veeam Availability Suite can also restore to a specific point in time, and even restore the database to a specific transaction. The important part here is that the restore process is driven from a wizard within the UI of Veeam Backup & Replication (or via the UI in Veeam Backup Free Edition). This takes the uncertainty out of the restore process, with the added protection of the option to send the restore task to a different SQL system.

This is, of course, in addition to the option to restore an entire system, as there may be other databases running that haven't encountered an issue. It's simply not the right thing to do if one application (and its corresponding SQL database) has an issue to restore other databases supporting other applications. That's additional data loss that can't be tolerated today.

In addition to SQL and Exchange, it's possible to use these same recovery techniques for Microsoft Active Directory and Microsoft SharePoint. In particular, AD recovery is an area where significant risk exists when

doing the wrong type of restore, including unnecessary hassles like role seizures when things don't go as planned.

A Happy Active Directory Ending

One story with a happy ending (which you can read about at **vee.am/ActiveDirectorySaved**) comes from a user who used a beta version of Veeam Backup & Replication v8 to restore more than 1,400 computer accounts. This could be a true disaster for a company to potentially have to go to every workstation (including remote systems) and re-add them to the AD domain.

Whether it's an entire organizational unit of computer accounts or a simple computer account or group in AD, the complexity involved in restoring these types of objects is extremely high. Like the other application explorers in their free mode, Veeam Backup Free Edition also provides support to restore an AD item.

Veeam Explorer for SharePoint can provide object recovery from Veeam backups. In particular, the ability to search across all SharePoint sites in a database makes finding what needs to be restored easy. SharePoint is also a special application composed of multiple components, including a SQL database. Like the other applications discussed, having to support all types of applications can be difficult, and not every organization will have expertise in every type of application.

Not everyone has access to a SQL database admin.

When it comes to addressing what can possibly go wrong and ensuring that datacenter availability is maintained, having application options is a must in today's IT practice. The important thing to remember is that deep application expertise isn't required to achieve this enhanced availability, because the right tool can make easy work of complicated application tasks, especially when administrators don't do such tasks on a daily basis. What's better, there are some powerful options available, even for free.

Remote/Branch Office Challenges

While I've painted a great picture for recovery, there still can be challenges brought about by the infrastructure. For instance, designing for remote or branch office (ROBO) sites has many requirements that might shift from site to site:

- The ROBO site might need to be fully self-sufficient
- Limited site-to-site bandwidth
- Small (or no) staff at remote sites
- A remote site might have a full application footprint

Too often, there's an ill-conceived notion that the ROBO site can't benefit from the same features as the main datacenter in terms of availability when things don't go as planned. When it comes to designing a backup architecture to unlock today's availability expectations, there are plenty of options to ensure no trade-off in capabilities at the ROBO.

Successful Veeam designs in the ROBO consistently put at least a backup repository, a proxy and a Veeam server (console) in each location. If data is to be sent to the central office, leveraging the Veeam Backup Copy Job with optional Built-In WAN Acceleration is the next natural step to move this data to a central datacenter. Additional retention can be kept in the central datacenter and provide e-discovery capabilities, should the need arise.

Are You Ready for E-Discovery?

Have you ever had an E-discovery situation in your professional work experience? E-discovery is common in lawsuits or any internal investigation where relevant information needs to be extracted or otherwise shared on what has happened on the data center systems. As such, e-discovery means different things to different people; but because Veeam Backup Free Edition includes this capability, it should be considered a significant step forward. What happens when an e-discovery request is made, but you don't have any tools to provide the data? This is where the built-in capabilities can save endless hours traversing live data and doing restores, to then repeat the process of traversing the data yet again. The search capabilities in the four Veeam Explorers for key applications (SharePoint, Exchange, SQL and Active Directory) can make easy work of finding that piece of information should a data-discovery event arise.

The infrastructure doesn't stop there in its ability to enhance the availability of the datacenter, however. Another free offering, Veeam Explorer for Storage Snapshots, provides another recovery technique. Veeam Explorer for Storage Snapshots reads storage snapshots for supported arrays (currently NetApp FAS and V-Series, HP StoreVirtual [LeftHand] and HP StoreServ [3Par]), and provides up to seven recovery scenarios. This includes file recovery from the VMs (both Windows and Linux VMs), whole VM recovery and the ability to run each of the Veeam Explorers for key applications.

What's better is that no prior preparation is required except a snapshot schedule for the array. Like the complicated application recovery, the comfort of the wizard-driven restore is used to easily restore from the storage snapshot.

It's clear there are options for the datacenter to keep availability high, whether the datacenter is a ROBO, has a mix of critical applications in play, or is leveraging modern storage systems. Veeam can give datacenter availability the options needed by businesses today. ■

Rick Vanover (vExpert, MCITP, VCP, Cisco Champion) is a product strategy specialist for Veeam Software. Vanover is a popular blogger, podcaster and active member of the virtualization community. His IT experience includes system administration and IT management. Follow him on Twitter @RickVanover or @Veeam.

VEEAM
IT JUST WORKS!™

Veeam Explorers

High-Speed Recovery of what you need when you need it

Backup is all about recovery and application-item level recovery is faster and easier with Veeam Explorers™. Zero in and recover exactly what you need:

Microsoft Active Directory

Microsoft SQL Server

Microsoft Exchange

Microsoft SharePoint

Storage Snapshots

> "The most common restore at our company is for items that users accidently delete in Microsoft Exchange. With Veeam, we search for the users' mailboxes, find the items (usually emails), and restore them back to their mailboxes in minutes."
>
> —
> **Raymond Ciscon**
> Senior Systems Engineer, MacLean-Fogg Company

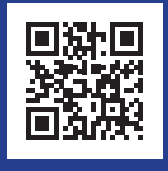

Learn more at vee.am/explorers

CEO Viewpoint:

Any Downtime Is Too Much Downtime

When Veeam CEO Ratmir Timashev sat down to talk technology, he ended up providing a stronger perspective on the man, his business values, and his strategy for building awareness around his not-so-secret plan to bring data availability to every business in the world.

Veeam: In response to World Backup Day, Veeam launched World Availability Day. Why was it so important to make that distinction?

Ratmir Timashev: We are not just technology innovators, but feel a responsibility to educate and support all types of businesses as they grow. The global enterprises that trust Veeam already know the critical role that availability plays in their business; the role that 24x7 access to all applications and data plays in daily operations.

Our goal with World Availability Day was to drive visibility across all businesses at every level, and really build awareness around the importance of going a step beyond backup. Let's be honest, even a small amount of unplanned downtime can affect a company's profitability and reputation. In some cases, a catastrophic event can even put the very

existence of their business in jeopardy. As both an entrepreneur and business leader, I see it as Veeam's corporate responsibility to deliver not just solutions, but education around this critical part of business development. If we can do that by simply harnessing even small opportunities like this to educate and empower businesses for the future, that's what we're going to do.

Q: Veeam has experienced tremendously fast growth, gaining more than 100,000 customers globally since 2006. How do you, as a global business leader, continue to stay so closely connected to your customers while achieving such remarkable growth?

Timashev: At this point in my career, I can say the idea of "practice makes perfect" certainly applies. I have always been passionate about creating great products and staying a step ahead of technology. That being said, when you have a great product, you need to be ready to meet the demand and do so very quickly. When I co-founded Aelita Software in 1997, within seven years we grew to $33 million. In 2004, we sold that business to one of our prime competitors, Quest Software, for $115 million. Quest was sold to Dell for $2.36 billion. Our products probably create close to $200 million in revenue for Dell. Admittedly, my companies have always had fast growth, so I can't imagine Veeam doing it any other way.

I am lucky in the fact that, as CEO, I get to travel the world meeting customers and partners and they tell me how they are making significant investments in a variety of business technologies, all in an effort to modernize their datacenters to provision IT services faster, strengthen security and control, and lower operational costs. I always ask myself, how will our solutions benefit today's user—whether it's a business professional, like me, who travels and needs access to information and applications while on the road, or consumers wishing to stream movies from Netflix or shop online—they demand access to data and applications 24x7. When you know the destination, the path is easy, and we can get there very quickly … or with "Veeam speed" as we like to say.

Q: Let's dig a little deeper into the impact that Veeam has made globally in helping organizations transition from legacy backup systems to embracing availability for the modern datacenter. No doubt the nature of the modern 24x7 business has directly affected your company's explosive growth. What has been your strategy?

Timashev: That's a good question, but again, it has a simple answer. We make it our business to know what is important to our customers. We listen to our customers. I know it sounds cliché, but to me it is a very serious matter. You cannot be a trusted advisor without knowing the pain points, the specific availability gaps [the gaps between what IT can provision/deliver and what users demand] when building a modern datacenter, without having that dialog and being informed.

I pride myself in being a technology innovator, so our specialty at Veeam is our laser-focused commitment to R&D, but you cannot build solid solutions without knowing what the goals are for your customers. As a matter of fact, one of our core values at Veeam is "Innovate and Iterate," so it's not just me, but all of the members of our Veeam team that support these forward-thinking ideas. This strategy translates perfectly to our success meeting the needs of our large enterprise customers like CBRE, Welch's, and American Standard, but it works just as well when accommodating small and midsize organizations.

Q: As the VeeamCEO, what motivates you professionally and how does that translate to keeping your growing global workforce productive and engaged?

Timashev: I think what motivates me to be successful in business are the same characteristics that motivate me in my personal life. My enthusiasm comes from finding solutions to challenges and helping to empower others to achieve whatever their goals may be. Whether it be professional business goals, educational goals or technology/innovation goals, bringing out the best in people is at the core of being an effective leader and creating tangible results. I apply those values at work, and it's important for me to reinforce those ideas to our employees.

My philosophy has always been that in order to have the best products, you need the best people, and I have been very strategic in ensuring the people that join our Veeam team are energetic self-starters who are empowered to take the initiative. That's why creating a place where they can grow, advance their careers, and contribute their talents to our business while also providing opportunities for them to serve and support causes within their own community is always top of mind. It's not just good business, it's the right thing to do.

Q: Speaking of big events and being engaged, last year's VeeamON conference was a landmark event for your company. Can you give us the inside scoop as to what we can expect at VeeamOn 2015?

Timashev: I was so pleased with the inaugural VeeamON event in 2014, and I can honestly say that this is now my favorite event of the year. The feedback we received from delegates, many of whom traveled from the farthest corners of the globe to attend, was outstanding. Proudly, it is on track to be even bigger and better this year. The opportunity for our global customers, development partners, and our Veeam Team to engage, network, share, and learn about the latest and greatest innovations is priceless. It truly is the world's premier datacenter availability event. Naturally, I cannot give away too much information at this time, but be prepared for something special. It's an exciting event and I would like to personally encourage readers to register and attend. ■

Welcome to Veeam

- Veeam is a global company with headquarters in Baar, Switzerland
- Veeam was founded in 2006
- Exponential revenue and customer growth

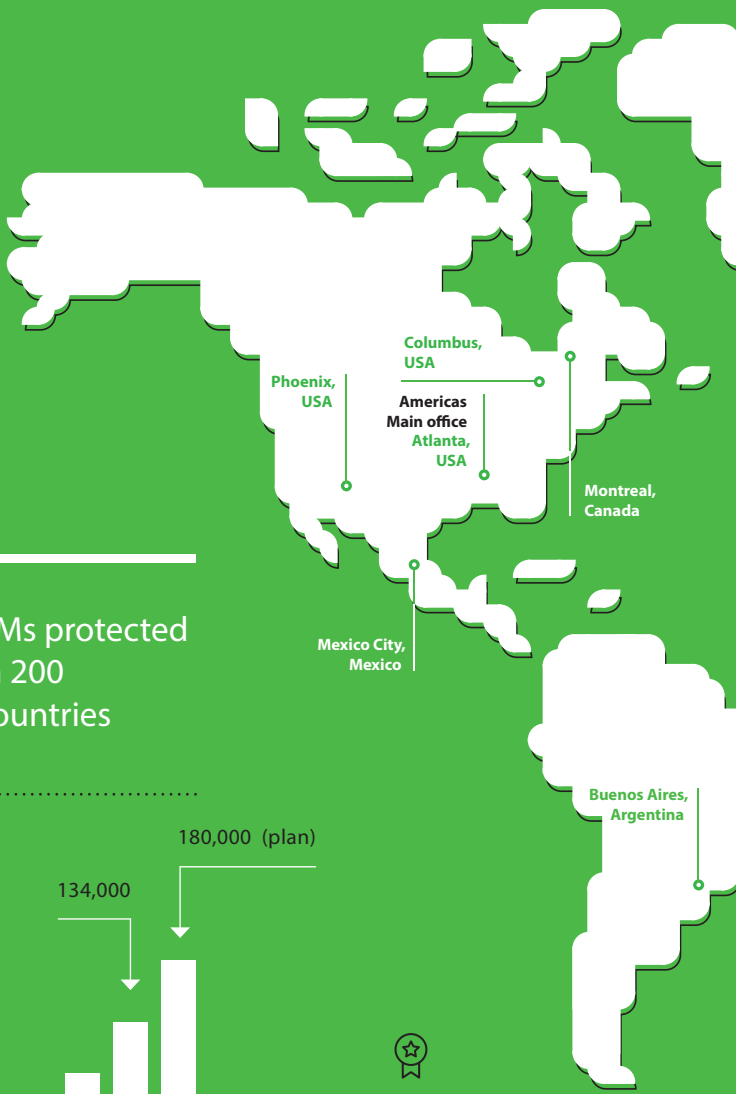

$1B Company's revenue target by 2018

8,4M VMs protected in 200 countries

Revenue bookings (M, USD)

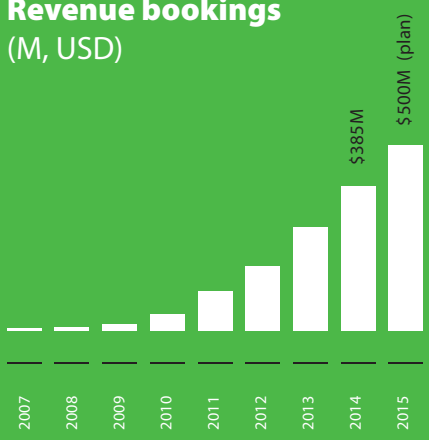

Number of customers

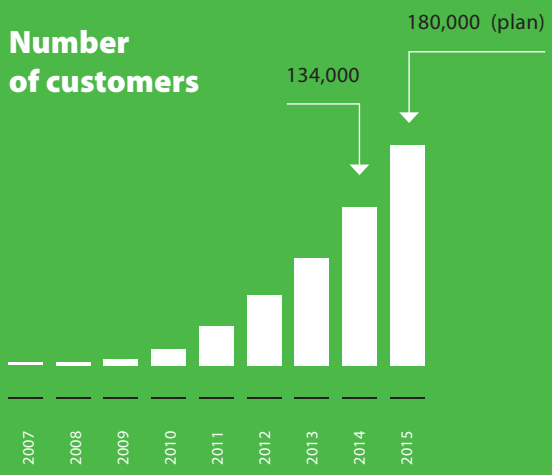

80+ top industry awards

Seven years of innovation

2008 Q1 — V1
- **2-in-1: backupand replication**
- **Instant File-Level Recovery**
- **Built-in deduplication**
- Synthetic full backups
- Replication multiple restore points

2008 Q3 — V2
- ESXi supportwith VCB
- Fastest VCBperformance
- Advanced VSSsupport

2009 Q1 — V3
- Instant File-LevelRecovery for Linux
- ESXi supportwithout VCB

2009 Q4 — V4
- Support forvStorage APIs
- **Support for CBT**
- Support for thin-provisioned disks
- Near-CDP
- Replication to ESXi

2010 Q4 — V5
- vPower
- Instant VM Recovery
- SureBackup
- U-AIR
- On-DemandSandbox
- Instant indexing

145,500

customers worldwide

- **13,000+** in Q4'14
- **4,000+** each month

30,500+

ProPartners

48% of **Global 2000**

70% of **Fortune 500**

are Veeam customers

2011 Q4 — **V6**
- Enterprise scalability
- Advanced replication
- Hyper-V support
- **1-Click File Restore**

2012 Q2 — **V6.1**
- **VeeamZIP**
- vPower for Hyper-V
- Veeam BackupFree Edition

2012 Q4 — **V6.5**
- **Veeam Explorer for Microsoft Exchange**
- **Veeam Explorer for Storage Snapshots**
- Support for ReFS
- Advanced backup monitoring, reporting and capacity planning

2013 Q3 — **V7**
- **Built-in WAN Acceleration**
- **Backup directly from Storage Snapshots**
- Veeam Explorer for Microsoft SharePoint
- Virtual Lab for Replicas
- Virtual Lab for Hyper-V

2014 Q4 — **V8**
- **Advanced storage integrations**
- **Veeam Cloud Connect**
- Veeam Explorer for SQL Server
- Veeam Explorer for Active Directory
- Encryption with password loss protection
- Backup I/O control

V9

VEEAM LIVE | shows, whiteboards, specials, webcasts etc.

Join the **Veeam LIVE** team for an interactive session around **Virtualization topics!**

Microsoft Azure and Veeam:
Innovative solutions for Data Availability

What's new
in Veeam Availabilty Suite v8

Veeam R&D Inside Out
with Anton Gostev

Veeam Cloud Connect Deep Dive: Design
principles & deployment best practice

Learn and
Win Prizes
including
a trip to
VeeamON 2015

Watch live
for more info!

Learn more
about Veeam Live
vee.am/veeamlive