CONTENT PROVIDED BY

POWERING THE NEW IT GENERATION
VIRTUALIZATION REVIEW

CONTENT SPONSORED BY

UNITRENDS
Play IT Safe

DO YOU KNOW WHERE YOUR DATA IS?

Your data is your company's life. But more and more, data is spread far and wide: in virtual machines, in the cloud, in offsite storage. Learn how to protect your precious data, and recover it should the worst happen.

1105 MEDIA

#1 All-in-One Solution
Backup, Archiving, Replication & Instant Recovery

WORLD'S ONLY BACKUP APPLIANCE BUILT FOR VIRTUALIZATION-FIRST ENVIRONMENTS*

*protect physical servers at no additional cost
(if you're into that sort of thing)

DOWNLOAD FREE TRIAL OR REQUEST YOUR EVAL APPLIANCE TODAY!

HTTP://GO.UNITRENDS.COM/UNITRENDS-VIRTUAL-FIRST-BACKUP-APPLIANCE

UNITRENDS

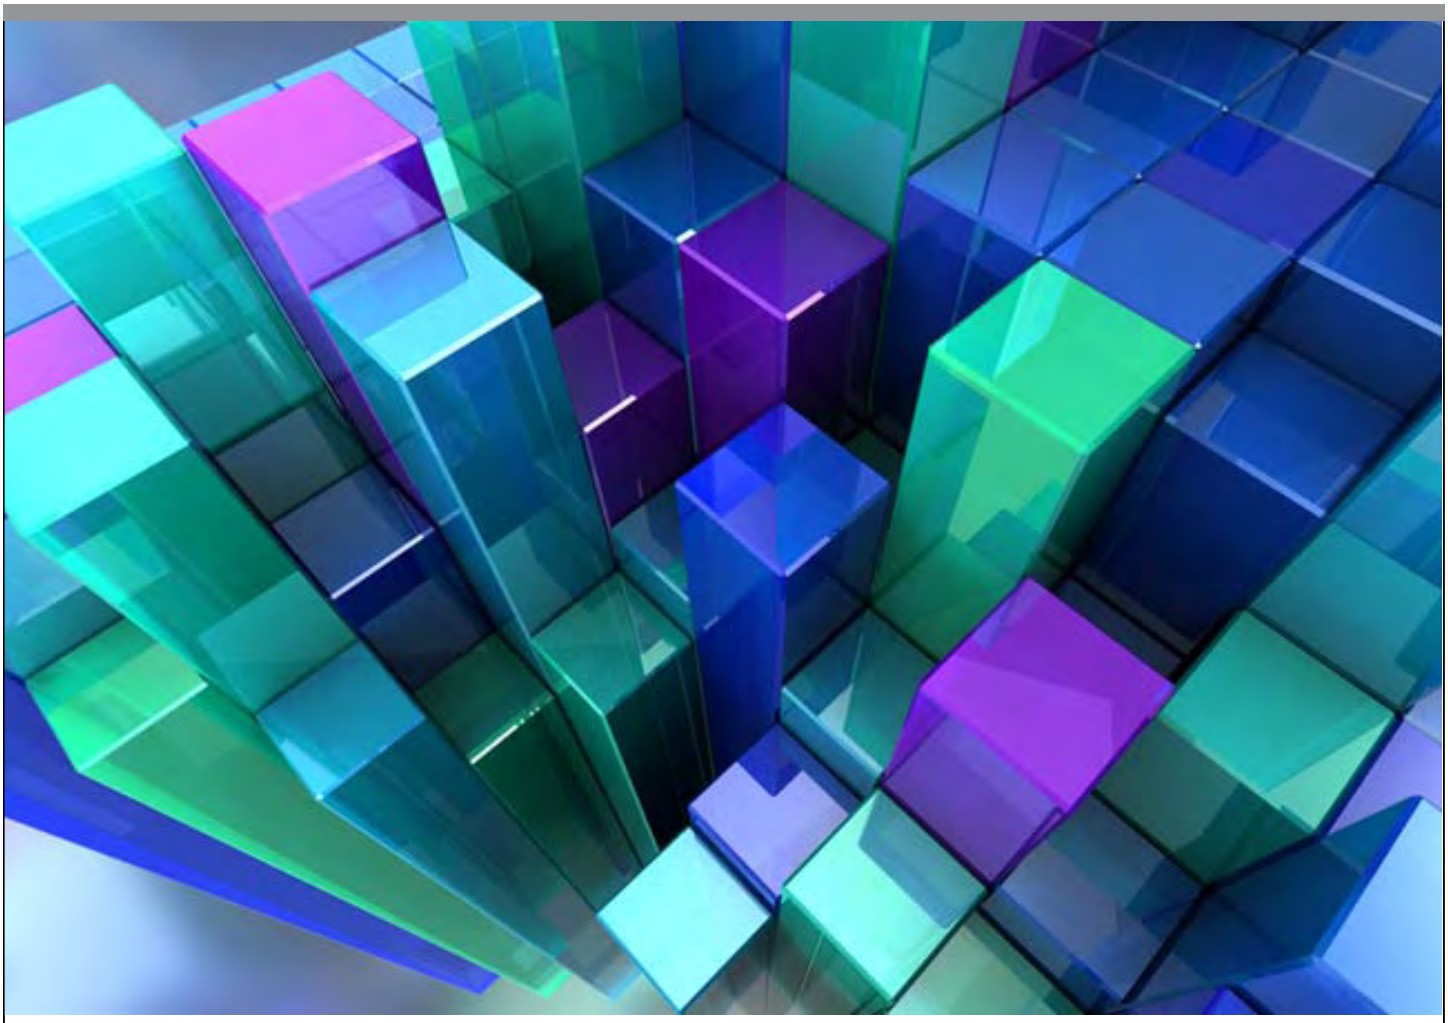

VMware, Microsoft Still Lead Gartner 'Magic Quadrant'

The latest x86 server virtualization report also shows an uptick for Red Hat. By Keith Ward

Gartner's latest Magic Quadrant report for x86 server virtualization infrastructure shows that the two leading companies in the space—VMware and Microsoft—continue to outpace the competition, but challengers are making the race more interesting.

As of July 2014, reports the analyst firm, VMware holds its commanding place in front. Microsoft is the only other company in the "Leaders" quadrant. Gartner said VMware "Continues to have dominant market share, and customers remain very satisfied with product capabilities and vendor support."

VMware "Continues to have dominant market share, and customers remain very satisfied with product capabilities and vendor support."

There are challenges for the leader, though. "... concern over price and vendor lock-in remains." Gartner also says that although VMware is still growing, so are the challenges of staying on top, mainly from market saturation and "competitive pricing pressure."

Much of that pressure is coming from Microsoft. Redmond is also maintaining its strong positioning, and Gartner says that on the server virtualization infrastructure front, it's getting closer to parity with the leader: "Microsoft has effectively closed most of the functionality gap with VMware in terms of the x86 server virtualization infrastructure." There are still gaps, according to the report, but Redmond "... can now meet the needs of most enterprises with respect to server virtualization."

The main challenge now for the company, Gartner says, is VMware's domination of the space. Microsoft is trying to compete on cost: "The most important factor in Microsoft's favor is price."

One company that made a jump in the Magic Quadrant was Red Hat. "Red Hat has made further progress in this year's Magic Quadrant, primarily due to a relatively strong tie between KVM adoption and OpenStack (roughly half of all OpenStack projects use KVM), and an increase in OpenStack adoption," the report states.

Gartner believes it has an insight into the company's virtualization goals: "Red Hat's strategy is to become the No. 3

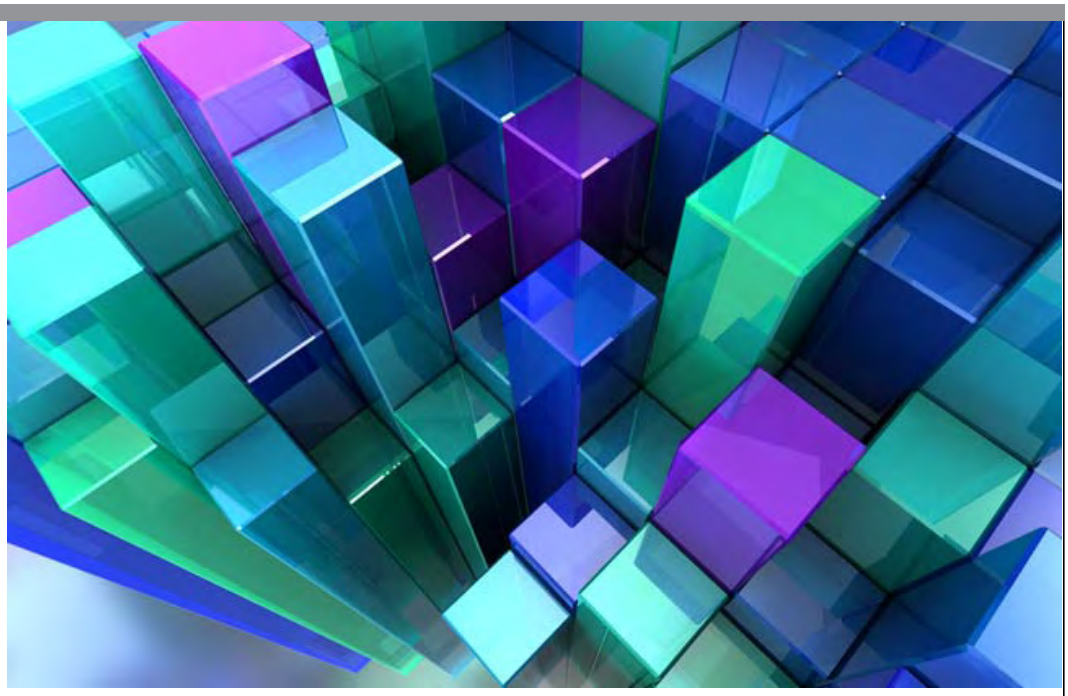

"Microsoft has effectively closed most of the functionality gap with VMware in terms of the x86 server virtualization infrastructure."

virtualization vendor behind VMware and Microsoft, especially in the RHEL [Red Hat Enterprise Linux] installed base and where there is concern over VMware lock-in."

The only vendor in the "Challengers" quadrant was Oracle. Gartner listed among its strengths that it "... has a large overall software installed base and financial strength, allowing Oracle to test and tune the hypervisor for optimal application performance."

However, Gartner added that Oracle's reluctance to embrace any non-Oracle virtualization components "limits its full growth potential."

Huawei, a self-described "global information and communication technology (ICT) solutions provider," is the only company added to the Magic Quadrant for server virtualization in the latest report. Gartner says "Huawei's FusionCloud suite—FusionSphere, FusionCube and FusionAccess—offers a good mix of physical, virtual and private cloud hardware and software combinations, with strength in its integrated capabilities." It entered under the "Niche Players" quadrant.

No companies were dropped from the Magic Quadrant list. **VR**

Keith Ward is the editor in chief of Virtualization Review.

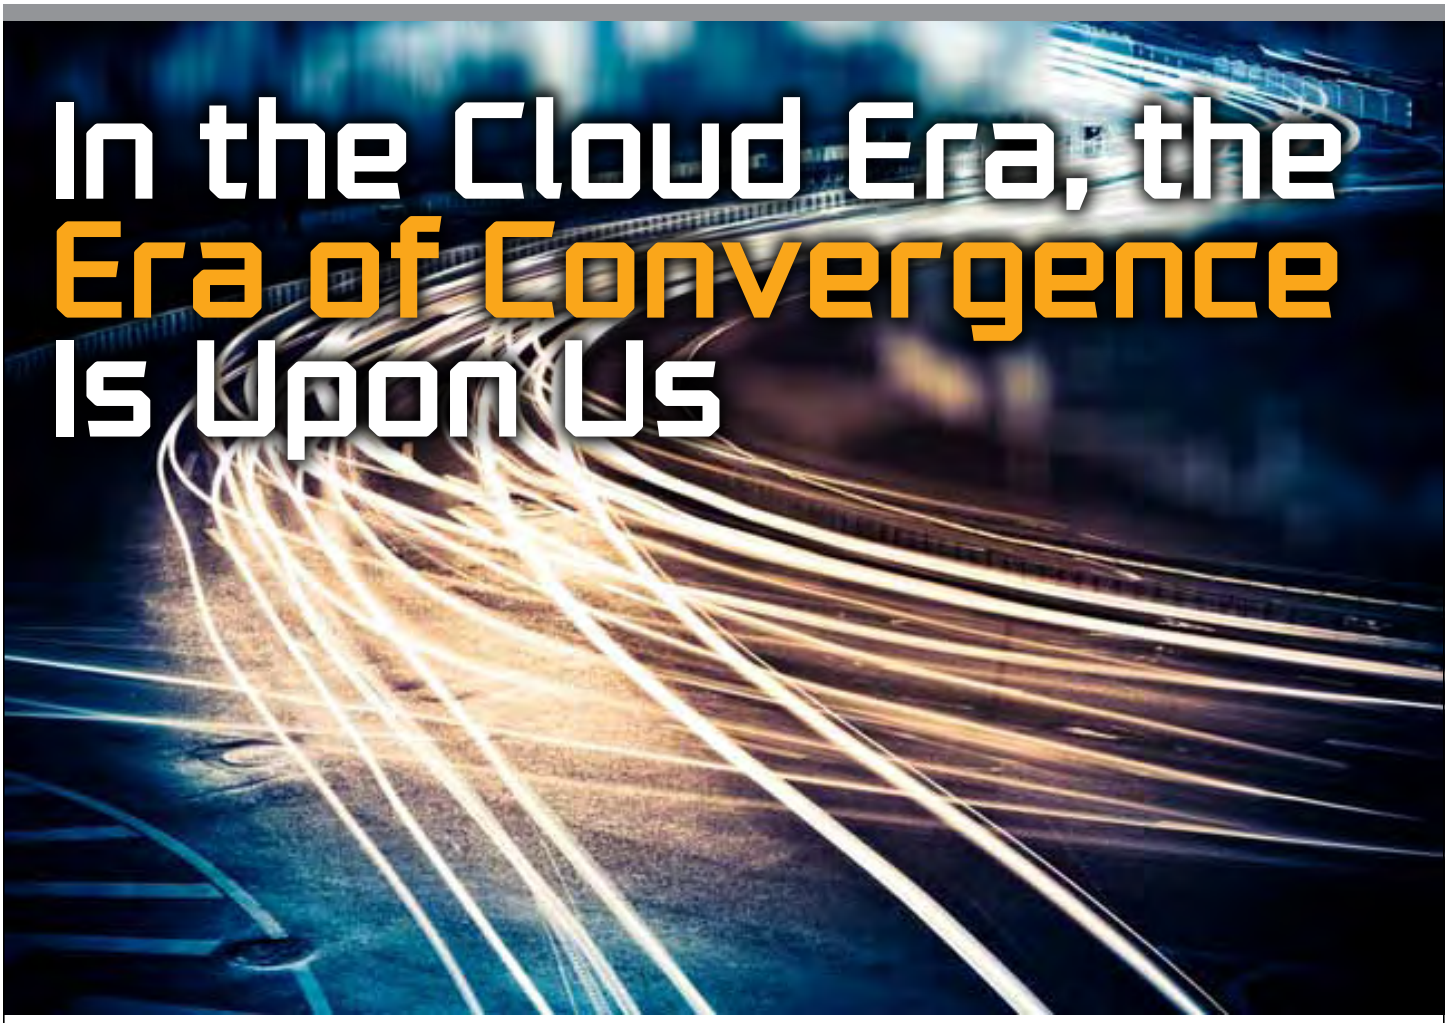

In the Cloud Era, the Era of Convergence Is Upon Us

What exactly is convergence and what is making vendors scramble to get included in this category?

By Arun Taneja

The era of IT infrastructure convergence is upon us. Every major vendor has some type of offering under this category. Startups and smaller players are also "talking" convergence. But what exactly is convergence and why are all the vendors so interested in getting included in this category? We will explain below the history of convergence, what it is, what it is not, what benefits accrue from such systems, who the players are, and who is leading the pack in true convergence.

Legacy Architectures Gone Wild

The fact is the fundamental makeup of IT infrastructure has remained the same for thirty years, maybe more. We have

the compute layer, the networking layer and the storage layer. Each layer has innovated in its own way and at its own pace but the layers have stayed intact. One could argue that the compute layer has followed Moore's law, the networking layer a quasi-Moore's law and the storage (or at least the HDD performance portion of it) a "flat line" law.

The issues created by the so-called I/O gap are well known and I will assume the reader is up to speed on its impact. As the amount of data kept growing at astronomical rates and the variety of data went from almost pure text to text, audio and video in a variety of formats, we kept throwing more and more hardware at the problem. We did this by simply adding more servers, each with more cores and higher speeds; networks with bigger switches and more bandwidth; and storage with more HDDs and more powerful controllers.

The overall IT stack today looks like a mishmash of technologies.

By the middle of the past decade, however, these infrastructures were at a breaking point, in spite of many "surrounding" technologies that surfaced in the 2000-2003 timeframe that kept them for utter collapse. And given the tsunami of data coming at us today it is only a question of time before the classic infrastructure will simply collapse on itself. But before we look at the "revolutionary" alternatives, let's look at some of these surrounding technologies, as they play an important role in the future of computing, regardless.

Technologies That Mitigate Infrastructure Issues

Technologies that have made the largest positive impact and allowed the current three-layer infrastructure to stay put, albeit in a wobbling state, include server and storage virtualization, data deduplication, compression, WAN Optimization, and flash in variety of implementations, including hybrid arrays and disk-based backup appliances. Of course, a list such as this would necessarily have to include cloud computing, cloud storage, and Hadoop (along with all its associated products) even if we would be hard pressed to call them "surrounding" technologies. I would also put scale-out architectures on that list.

Each of these technologies, in its own unique way, has helped keep the balance. For instance, server virtualization brought

consolidation and agility; storage virtualization delivered improvements in provisioning speed, capacity utilization, and management; data deduplication and compression enabled HDDs to be used economically for backup and restore and brought associated improvements in RTO's and RPO's and DR; and WAN Optimization made sure remote office employees didn't feel like second class citizens of the enterprise anymore. Of course, flash is on its way to revolutionizing application performance; cloud computing is fundamentally changing how we consume compute and storage resources and Hadoop is helping extract information out of mounds of collected data so better business decisions can be made.

The fundamental makeup of IT infrastructure has remained the same for 30 years, maybe more.

With the exception of cloud computing and Hadoop, however, all these technologies have been bolted on to the traditional three-layer architecture of the 1970's. As a result, the overall IT stack today looks like a mishmash of technologies, essentially with compute, networking and storage layers surrounded by the plethora of new technologies mentioned above. This raises the question: Is this the best way to run the railroad? The answer, as you guessed, is a definite no. Two essential approaches have appeared on the horizon: convergence, and what we call hyperconvergence. We will look at each in more detail.

Convergence Defined

In a bid to simplify the IT infrastructure, a number of vendors, especially the legacy players, started bundling specific configurations of compute, networking, storage and server virtualization and pre-testing them for interoperability and performance for targeted workloads. The first one to market was VCE, a joint collaboration of VMware, Cisco and EMC, with Cisco providing compute and networking components. Specific configurations were pre-tested for strict interoperability and performance for workloads, such as SAP or MS Exchange or VDI, etc. Management was simplified by adding software that viewed the unit as a whole; however, if a layer was not performing adequately, regular tools that came with that layer were used to diagnose and change configurations.

I think of this type of configuration as taking three atoms and combining them to create a molecule. You buy, deploy, run,

and manage the unit as a molecule. If you buy the right model number for the task, the probability that it will deliver the right SLA for the applications is higher than if you bought these layers separately from three different vendors and put them together yourself. The burden of deciding which models were appropriate to mix together to do a specific job was taken off the buyer. This simplification is far from trivial. With customers looking to deploy cloud-scale infrastructures, one could drop these molecules into place, knowing they work at a specific level of performance. Management became easier and deployment time went down from weeks or months to days. Just as importantly, the TCO was impacted favorably.

As the amount of data kept growing at astronomical rates, we kept throwing more and more hardware at the problem.

Over the past three years, all major players have jumped into this fray that the market calls Convergence. HP offers CloudSystem Matrix; NetApp worked with Cisco and VMware to offer FlexPod; Dell combined PowerEdge servers, EqualLogic arrays and Force10 network switches to deliver their converged solution as Active Infrastructure; and IBM offers PureFlex Systems, which combine IBM POWER and x86 server blades and Storwize V7000 storage with networking, server virtualization and management components.

While these converged systems provided ease of purchase, deployment and use along with significantly improved sharing of resources, a fundamentally different phenomenon was taking place in the market. At Taneja Group we call this hyperconvergence and consider it to be distinct from convergence. Alternatively, one could think of hyperconvergence as a continuation and maturation of convergence but we prefer to keep the categories separate on the fundamental assumption that players along the convergence axis cannot simply improve their products and become hyperconverged, without serious architectural changes. In other words, we believe that true hyperconvergence can only be achieved by starting with a clean slate and not by mixing existing pieces.

So what is hyperconvergence and how is it different from convergence?

Hyperconvergence Defined

We believe hyperconvergence occurs when you fundamental-

ly architect a new product with the following requirements:

1. A genuine integration of compute, networking, storage, server virtualization, primary storage data deduplication, compression, WAN optimization, storage virtualization and data protection. No need for separate appliances for disk-based data protection, WAN optimization or backup software. Full pooling and sharing of all resources. A true datacenter building block. Just stack the blocks, and they reform into a larger pool of complete datacenter infrastructure.

Given the tsunami of data coming at us today, it is only a question of time before the classic infrastructure will simply collapse on itself.

2. No need for separate acceleration or optimization solutions to be layered on or patched in. Performance (auto-tiering, caching and capacity optimizations all built in). As such, no need for separate flash arrays or flash caching software.

3. Scale-out to web scale, locally and globally, with the system presenting one image. Manageable from one or more locations. Radical improvement in deployment and management time due to automation.

4. VM centricity, i.e. full visibility and manageability at VM level. No LUNS or volumes or other low level storage constructs.

5. Policy-based data protection and resource allocation at a VM level.

6. Built-in cloud gateway, allowing the cloud to become a genuine, integrated tier for storage or compute, or both.

With today's converged systems one would have to add separate backup appliances, backup software, WAN optimization appliances, flash arrays, flash cashing software, cloud gateways, and more, to get to the conceptual equivalent of the above. But one would still not achieve VM-centricity, or web scale or space and power savings, or the ability to manage all these pieces as a whole. One could get close but no cigar. This is why we believe hyperconvergence is a separate and distinct category. According to our definition,

we believe three systems in the market fall into this category: Nutanix, SimpliVity and Scale Computing.

Nutanix came to market first with a "built from scratch" hyperconverged system that met most of the requirements from day 1. Missing initially was data deduplication and global management capability, which Nutanix added recently in rev 4.0. The Nutanix Virtual Computing Platform is VMware-based but Hyper-V was added as an option in rev 4.0. One can now build a cluster using VMware-based nodes and a separate cluster, using Hyper-V nodes, and manage the whole, globally, as one instance.

Management, even at the global level, becomes trivial, compared to managing traditional architectures or, for that matter, Converged architectures.

SimpliVity took a slightly different tack. Given that many of its developers came from Diligent Technologies (now IBM), the purveyor of in-line data deduplication appliances, SimpliVity started with the premise that data should be reduced to its smallest size at inception and kept that way for the entire lifecycle, whether it is being moved, stored or analyzed, except when it needs to be viewed by a user. In order to ensure that this capability stood out, SimpliVity developed a special PCIe card to handle the number crunching required by the deduplication algorithms, without impacting the ingest performance.

Nutanix, on the other hand, wanted to stay true to a 100 percent commodity hardware strategy, so they chose post-processing data deduplication for HDD to ensure zero impact on performance. For main memory and flash storage, Nutanix chose in-line data deduplication, which makes these small capacities effectively much larger.

Regardless of the differences, both products meet the essential premise of hyperconvergence and the differences between them are architectural and can only be evaluated in a hands-on evaluation.

Scale Computing, on the other hand, is targeted at the lower end of the market and uses KVM as the hypervisor. Given the open source nature of KVM, Scale was able to more tightly integrate KVM into the architecture (more so than one could with VMware or Hyper-V). At least at this point in time Scale

does not offer data deduplication but most other requirements of HyperConvergence are met in full. For smaller IT shops where IT specialists are rare, the ability to buy the whole infrastructure as a unit and manage it as such can be a gift from the heavens.

Benefits of hyperconvergence

The best way to think of hyperconverged systems is to think of them as "infrastructure in a box." You can start with the minimum number of nodes the vendor requires—two for Nutanix and SimpliVity and three for Scale Computing. All functionality we've mentioned as requirements is included in each node. Installation and deployment times are trivial. You decide on the importance of each VM you wish to run and assign each a priority, which will determine how much resource is to be made available to that VM, in terms of IOPs, throughput, latency, degree of protection, RTO/RPO, etc. The system does the rest. All relevant data is available on a VM by VM basis.

The best way to think of hyperconverged systems is to think of them as "infrastructure in a box."

If more resource is needed, given the mix and resource requirements of all VMs, the system alerts the operator that one or more additional nodes are needed. Adding the nodes is simple. The cluster recognizes the additional nodes automatically and storage and compute resources become available instantly. If remote clusters are installed, the clusters can recognize each other and present a single image to the IT administrator. All data is presented at the VM level and there are no low level storage tasks (provisioning, LUN creation, virtualization, expansion/contraction of volumes, balancing workloads, etc. etc.) to be performed. As such, the server virtualization administrator can easily manage the entire cluster, without requiring strict storage experience.

Management, even at the global level, becomes trivial, compared to managing traditional architectures or, for that matter, Converged architectures. All data exchange across the WAN happens efficiently, using WAN optimization methods, and only unique data is sent across and even that in a compressed fashion.

We believe this level of functionality and integration can only happen if one starts with a clean slate. It is hard, if not

impossible, to make this happen with equipment from ten or more vendors, each with its own idiosyncrasy. This is why we believe hyperconvergence may conceptually be viewed as an evolution of convergence, but in reality it is more "revolutionary" than not, even if most components of the technology it uses are well defined and mature.

Where Does Virtual SAN Fit In?

VMware announced Virtual SAN as a product that essentially allows a number of compute nodes with local HDD and flash storage (DAS) to pool their storage resources and make the pool available to all applications, running as VMs. vCenter becomes the central place to manage the entire cluster (no separate storage console). Configurations as large as 32 nodes were announced and the product is being targeted at midsize and large enterprises for all but tier-1 workloads. All vSphere services are available, including vSphere Data Protection Advanced, vCenter Site Recovery Manager, vMotion, Storage vMotion, DRS, etc.

The pain of managing large infrastructures has become so acute that hyperconvergence presents almost a panacea.

The Future of Hyperconverged Solutions

If the current reception to hyperconverged solutions by midsize and large IT is any indication, hyperconverged solutions will cut deeply into traditional architecture-based solutions. And they will do so very quickly. The pain of managing large infrastructures has become so acute that hyperconvergence presents almost a panacea. The combination of workload unpredictability, the pace at which new data is coming into the enterprise, and the requirement to deliver results instantly, all point to a new architecture that adapts and adjusts automatically, with little or no human intervention.

While traditional architectures keep improving in all these dimensions, incremental improvements are just not enough. hyperconverged solutions could not have arrived at a more opportune time. VR

Arun Taneja is founder, president and consulting analyst of the Taneja Group, an analyst firm specializing in storage and storage-centric server technologies.

5 Tips for Mastering Virtualization Backup

By David Davis

At too many companies today, getting reliable and 100 percent recoverable virtualization backup for the entire infrastructure is more art than science. You might get a daily message saying that backups have only partially completed, but you've come to accept it. Or maybe you get a message saying that backups have successfully completed, but you don't trust them. Backup and recovery should be a core function your IT department offers—and one you can perform entirely. Unfortunately, there are too many companies where this isn't the case. Now is the time to master your virtualization backups.

To help you, I've compiled a list of my top five tips for mastering virtualization backup.

1. **Select the Right Tool**
   So much of your success in backup and recovery is dependent on selecting the right tool for the job. If you have physical (non-virtualized) servers in your infrastructure, you need a tool that can back them up. If you're 100 percent virtualized, a virtualization-only backup tool might be the best answer. However, if you're like most companies out there—which are, let's say, 90 percent virtual and 10 percent physical—then you need a tool that's able to back up both your physical servers and your VMs. For the sake of simplicity (and licensing costs), I strongly believe that companies shouldn't have multiple backup tools. It just doesn't make sense.

Backup and recovery should be a core function your IT department offers—and one you can perform entirely. Unfortunately, there are too many companies where this isn't the case.

2. **Plan to Get Your Data Off-Site**
   You can't just stop at successfully backing up your data. You need to get that data off-site. Too many virtualization-specific backup tools stop at the point of backing up your data and offer few options to get data off-site. Not every company has the bandwidth to perform replication. You need a controlled and automated way to move data from your backup repository to a portable device (tape or disk) for off-site storage in the event of a disaster.

3. **Utilize Advanced Features**
   The backup tool you select should offer advanced features. Examples include:

   - **De-duplication and compression:** Ensure that backup data doesn't consume excessive storage.
   - **Application-consistent backups:** Utilize the Volume Shadow Copy Service (VSS), and backup tools to back up your data so that application data is quiesced and has integrity.

- **Verification and automated recovery testing:** Ensure that VM OSes and applications will truly function, once restored, in an automated fashion.

Many of the backup tools today offer one or more of these features, but they're executed differently. My recommendation is to ensure that your backup tool offers all the features you can afford, and that they're as efficient as possible.

4. **Make Sure It's Fast**
Companies make claims about how fast and reliable their backup products are, but have you ever tested your backup product? How fast can it get your largest server running again if it's lost? How fast can it restore 20 servers at once? After all, speed is always relative to your servers, storage and data.

Have you ever tested your backup product? How fast can it get your largest server running again if it's lost?

5. **Future-Proof It**
What if you want to move vSphere VMs to Hyper-V, recover your physical servers to vSphere or recover your vSphere VMs to a public infrastructure cloud? Will your backup tool support cross-hypervisor backup and recovery? What about off-site backup to a cloud? Ensure that your backup tool is innovative and has a history of launching new features that give you the greatest flexibility possible.

Bottom line: Look for a tool that covers all the bases for both physical server and virtual infrastructure backup, while offering numerous advanced features. VR

David Davis is a well-known virtualization and cloud computing expert, author, speaker, and analyst. David's library of popular video training courses can be found at Pluralsight.com. To contact David about his speaking schedule and his latest project, go to VirtualizationSoftware.com.

The Disaster Recovery Assurance Paradigm Shift

The cloud has shifted the frequency with which we need to test our systems' capabilities to come back from disaster, from months to right now and often. **By Steve Kahan**

When Gartner Group recently reported that each disaster recovery test a company performs costs $30,000 to $40,000, there is a big problem—it's a huge waste of money!

Disaster recovery assurance represents a paradigm shift for cloud operations and management. It unlocks huge value to companies regardless of size.

Paradigm shifts occur when an invention changes the behavior of people permanently. For instance, the automatic transmission in cars made them easy for everyone to use and led to two-car families, suburbs, two-salary households, and so on.

We see three major parts to this paradigm shift:

First, with today's tools, disaster recovery plans are tested once or twice per year. Maybe. It's hard to do. It's so hard in fact that tests need to be planned and prepared well in advance, usually two or three months in advance.

Once people adopt disaster recovery assurance, going back to testing twice per year will be unthinkable in the same way going back to cars without automatic transmissions is unthinkable.

DR assurance is about testing every few hours. Let's stop to think about this. From twice per year to several thousand times per year is three orders of magnitude—that's a paradigm shift. Once people adopt DR assurance, going back to testing twice per year will be unthinkable in the same way going back to cars without automatic transmissions is unthinkable.

Second, DR assurance unleashes huge value. DR tests are not only very disruptive but also very expensive because customers have to rent data center space, bring in consultants and so on. It's easily $30k per test for a small company, multiply that by ten for a large company. We have customers that need between 1500 and 4000 man-days per year to conduct the DR tests they need. This is why, for a Fortune 500 company, DR testing can easily cost several million dollars per exercise.

And what's worse is that within hours of a DR test, the results become obsolete. Clouds reconfigure virtualized resources all the time; CPUs, cores, servers, memory, all of the storage, and increasingly the networks—they are constantly reconfigured and reallocated. That's what clouds do.

And that is why testing frequently is the only way to make sure that recovery procedures are current, and to ensure compliance with business continuity requirements. Now

imagine the value of driving the cost of testing down from tens of thousands of dollars to a few dollars per DR test—several orders of magnitude less.

With this all of this potential savings and new recovery testing and procedures to think about, how can you make sure to choose the right partner to utilize and guide you through the shift? Choose a disaster recovery assurance provider that has all the core technologies (backup, archive and replication) that also hooks into the cloud and automates recovery testing with recovery time objectives and recovery time actuals.

For a Fortune 500 company, disaster recovery testing can easily cost several million dollars per exercise.

Disaster recovery assurance is going to change the way people protect their IT. They'll need to assuring that corporate business processes are recoverable at all times with 100 percent confidence and for a fraction of the cost today. VR

Steve Kahan is PHD Virtual's chief marketing officer. Prior to PHD Virtual, Steve was senior vice president of global marketing for Quest Software, and vice president of marketing at The Planet, Bindview, Postini, PentaSafe and was president and CEO of eSecurity. Steve holds a B.S. in Communications from Illinois State University.