IT Project:
Hyper-converged Infrastructure

Your expert guide to implementing HCI

In this e-guide

In this e-guide:

The term hyper-convergence means different things to different people. But at its core, hyper-convergence brings together servers and storage network resources that are managed as a single unit using software. In some ways, it signals a return to a mainframe-like architecture before the rise of client-server, distributed computing model. In this e-guide, learn more about the benefits of hyper-convergence and how you can derive the greatest value from your investment in hyper-converged infrastructure. Plus, read about the movers and shakers in the hyper-converged market and why NVMe is a rapidly rising storage star.

Aaron Tan
Senior Editor, APAC

To reap the benefits of hyper-convergence, manage it right

Alastair Cooke

To gain the greatest value from your investment in hyper-converged infrastructure, it's critical to embrace the simplification it brings to data center operations. That simplified management of the data center, after all, ranks among the top benefits of hyper-convergence, and it would be unwise to replicate the constraints of your previous system's architecture.

If your organization is interested in hyper-converged infrastructure (HCI), it's smart to give some thought to how you'll manage it once it's implemented. You'll want to be sure this powerful new technology lives up to expectations. A well-integrated infrastructure product can make a data center simpler to operate, yes, but that doesn't mean your IT team won't have to manage the HCI to be sure it does all that it is capable of doing.

The operational model for HCI is based on VM-centered and policy-based management. Large pools of resources are controlled through these policies. Only policy exceptions need attention.

It is also important to remember that HCI is not a silver bullet for all of your virtualization management. No matter what the sales rep told you, you will still manage your hosts, VMs and networks.

Policies and pools

A policy is a set of controls you want to apply to a group of VMs. It might be a performance policy that guarantees certain resource levels, or it could be a data protection policy that mandates local backups and remote replication. You define the policy once and then apply it to a set of VMs. Each VM might have multiple policies attached to govern things such as performance, availability and disaster recovery.

Once the policy is applied, you do not need to check each VM to see that it is compliant; you'll be alerted when there is a policy breach. As an example, if you have a policy that requires a virtual machine to reside on your fastest tier of storage, you will be notified if someone migrates that VM to a slower tier.

Creating these policies requires some time and effort at the outset -- you will want to be sure you lock in the right policies for your particular needs. Once these standard policies are created, your VMs should be allocated a policy. You do not want to have a new policy created for every new VM. Policies should be standardized in the same way that cloud services are standardized, as a menu from which an option is selected.

Among the benefits of hyper-convergence is how an HCI product frees you from having to do certain things. Staff simply won't need to commit the same amount of time to data center tasks, and operations will be less complicated.

With most HCI installations, you'll create fewer clusters and fewer data stores. This results in fewer choices to make about where a new VM or application should reside. Data protection and storage performance policies are applied to VMs, not inherited from a data store or cluster. There will seldom be a need for more than one data store. With fewer, and larger, pools, we can share resource headroom and run at higher total utilization.

So, what's managed?

The virtualization team manages the HCI platform. There is no requirement to open tickets for storage administrators to make new logical unit numbers (LUNs) available for added VMs. The virtualization admins are the ones responsible for making sure storage resources perform for the applications in the VMs. Most often, this is a simple matter of assigning a storage policy and configuring the virtual machine.

Your IT team will be responsible for capacity planning on the storage. Often, all of the hyper-converged infrastructure's data stores will share the same underlying physical storage. You'll need to use the HCI platform to manage the capacity, as the hypervisor seldom understands the HCI capacity efficiency.

HCI enables simplification. One important example of this is the elimination of the dedicated storage network. All of the storage traffic goes over Ethernet, usually 10 gigabit Ethernet (GbE). In many HCI deployments, the 10 GbE will also carry management traffic and VM network traffic.

Managing the configuration and load on this network can be an important part, ensuring the delivery of applications on the HCI platform. You will still work with the network team to get virtual LANs provisioned and to troubleshoot any networking issues. In many HCI products, the Ethernet cards are the only expansion cards; there is no option to add any other networking.

Most hyper-converged infrastructure products plug into your existing hypervisor management, and all of the usual VM management remains. The skills you learned about configuring and managing virtual machines remain almost entirely unchanged with HCI. You still configure VMs with the right resources and configurations for the workloads. Some HCI platforms have their own versions of the hypervisor cloning and templating functions. These tend to be more efficient than the hypervisor tools, so using them can save time.

A few HCI offerings bring their own hypervisor. That means you won't plug into your existing hypervisor, and you will use the HCI vendor's interface for all management. There will be a learning curve to build familiarity with this new interface.

With this approach, the HCI vendor can simplify and integrate more deeply. That's one of the benefits of hyper-convergence.

What to avoid

Options in the hyper-converged market are plentiful, so you'll have a lot to consider when selecting a product. Just be sure to think about how you'll actually implement that system.

When adopting HCI, it's important to leave behind the baggage from your old infrastructure. Fresh design decisions should be made to match the capabilities of your chosen HCI product.

As an example, there's no reason to duplicate the data store design from your old storage environment onto HCI. On older storage, a LUN was a boundary for performance and for replication. There were high-performing LUNs and lower-cost LUNs. Some were snapshotted and replicated for data protection. VMs would be placed on these LUNs to satisfy the availability and performance requirements of their applications.

With HCI, availability and performance are policies applied to the VM. In many cases, a single data store will fulfil the requirements for a diverse set of application needs.

//

⤵ **Next article**

▪ Nutanix, Simplivity and Pivot3 lead hyper-converged pack: Forrester

Antony Adshead, Storage Editor

Startup specialists Nutanix, Simplivity and Pivot3 lead the hyper-converged infrastructure market, with giants such as EMC, Cisco, HPE and VMware tailing in the chasing pack.

That's according to analyst house Forrester, which evaluated 12 suppliers of hyper-converged infrastructure in terms of capabilities, company strategic positioning and market presence in Forrester Wave research included in its *Infrastructure Transformation Playbook For 2017*.

The 12 suppliers evaluated were: Nutanix, Simplivity and Pivot3, designated "leaders"; Atlantis Computing, Gridstore, EMC, Stratoscale, VMware, Cisco, HPE and Huawei, rated as "strong performers"; and Scale Computing, dubbed a "contender".

Hyper-converged infrastructure has been a rising star in the datacentre in the past couple of years. It combines server and storage hardware in one box, often with a software layer that allows for native hypervisor operations and the ability to discover and run in a grid-like fashion with other hyper-converged nodes from the same maker.

Hyper-converged infrastructure products are available as a hardware nodes as well as software products.

Out of the suppliers evaluated, Nutanix has the most customers (3,100-plus) and is rated by Forrester as "the largest supplier" of hyper-converged infrastructure products, with 1,800 employees and operations in 80 countries and customers that run more than 100 nodes. Its annual revenue, estimated by Forrester, is around $200m.

Second in Forrester's rating, Simplivity can only scale to 24 nodes but can scale to 1.3PB with its Omnistack product offering "a rich set of global multisite data management functions" to enterprise customers.

Third among the leaders is Pivot3, which has been operating longer than Nutanix and Simplivity and which has more than 2,000 customers for products. Historically, it focused on surveillance and media storage and has petabyte-scale capabilities.

Looking at the strong contenders, Forrester said Atlantis Computing, with its mid-size flash-based nodes, "has a small market footprint, but its overall capabilities give it a strong position".

Gridstore, which has rebranded to Hypergrid, with its flash-based hardware and Microsoft-only hypervisor compatibility is noted for its quality of service capabilities that give minimum performance guarantees.

EMC, with its VMware VSAN-based VxRail appliances is rated as, "in a good position to participate in the industry-wide shift toward software-defined infrastructure environments".

Stratoscale, founded in 2013, proclaims a desire to bring "Amazon Web services (AWS)-like capabilities to the datacentre" with its KVM Hypervisor and ZFS file system. It has few customers so far but Forrester believes it has, "one of the more advanced workload and data migration capabilities in our supplier group, as well as integrated high availability".

VMware, with its VSAN (Virtual SAN) claims more than 4,000 customers (unverified by Forrester) and has a rich feature set with tight integration with the VMware ecosystem.

Cisco recently entered the hyper-converged market with Hyperflex, and Forrester thinks that its "market ramp should be rapid", because of the existing underlying acceptance of its UCS server hardware.

Similar to Cisco, Forrester sees HPE's Proliant-based HC380 hyper-converged product as in its early stages in the market but with potential owing to its existing market presence.

Forrester notes the multi-hypervisor support of Huawei's FashionCube products and technical maturity, but it also has weak data services.

Alone in the third tier as a "contender" is Scale Computing, which is acknowledged as a small to medium-sized-targeted hyper-converged product with, "a large installed base in relation to its revenues".

↳ **Next article**

HPE's Simplivity purchase keeps hyper-convergence in the spotlight

Simon Robinson, Research Vice President

In January 2017, Hewlett Packard Enterprise (HPE) announced it would pay $650m cash for seven-year-old startup Simplivity.

The move has dispelled any lingering doubts over the potential for hyper-converged infrastructure to radically reshape the way organisations buy and deploy the platforms that run their IT operations.

The acquisition is the latest development in a rapidly evolving market that is becoming very real, in which large numbers of organisations have begun to deploy a range of software-defined and hyper-converged infrastructure capabilities to run an expanding range of traditional and emerging IT workloads.

Hyper-converged is also now becoming a core weapon in the armoury of suppliers looking to demonstrate to customers that they can meaningfully change the fundamental economics of running core and edge IT infrastructure

This is not only a linchpin of any IT transformation strategy, but absolutely essential for any supplier that wants to demonstrate ongoing relevance in the era of public cloud.

HPE's acquisition of Simplivity should further validate a market that until recently was mostly defined by a single startup specialist: Nutanix.

Let's not forget that the now-public company was only founded in 2009 – three years after Amazon unveiled its first cloud computing service – and continues to be a major force in the hyper-converged infrastructure market. Nutanix now has over 4,500 customers, almost four times the approximate 1,300 customers Simplivity has so far accrued.

One crucial test of whether a market is becoming mainstream is the extent to which it is embraced by large incumbent suppliers.

As far as hyper-converged infrastructure is concerned, we are now approaching that point.

VMware recently said it has exceeded 7,000 customers for VSAN, the hyper-converged infrastructure software that is optimised for VMware environments.
Meanwhile, the combined Dell-EMC has highlighted hyper-converged infrastructure as a key growth market. Although to date this has mostly been built on the Nutanix-based Dell XC Series appliance, the merged company is

focusing on its own stable of capabilities, which include the VxRail and VxRack hardware, that run VMware VSAN and the ScaleIO storage software stack.

Other major players are still hoping to make a run at hyper-converged infrastructure, including Cisco with HyperFlex, and there are still startups with momentum, including Pivot3 and Scale Computing.

Indeed, this broad range of participants highlights the diverse nature – and opportunity – that exists for hyper-converged infrastructure across a number of market segments.

Hyper-converged infrastructure initially had a reputation as primarily a remote office or branch office and departmental technology due to its promise of simplicity.

However, when asked where organisations had deployed hyper-converged infrastructure, 74.4% of respondents to 451's *Voice of the Enterprise (VoTE) server and converged infrastructure study* said core or central datacentres.

This compares with only 26% deployed in departmental or regional datacentres, 26% deployed with third-party colocation or outsourced service providers and 19.4% deployed in remote or branch offices.

Signs of hyper-converged evolution

We believe this signals the rapid maturity of hyper-converged from supportive, edge infrastructure to being a primary infrastructure solution.

Also, organisations indicate plans to migrate a wide range of workloads to hyper-converged infrastructure over the next two years.

This includes workloads typically associated with hyper-converged deployments, such as virtual desktop infrastructure (VDI), but also business-critical workloads such as ERP, email and industry-specific applications.

This highlights that there is still an awful lot of traditional on-premise infrastructure out there, and it is still worth fighting for. HPE's acquisition of SimpliVity is yet another development on that front, and signals a fascinating journey ahead.

↘ **Next article**

▪ # Storage predictions 2017: Hyper-converged to march on, and get NVMe

Antony Adshead, Storage Editor

Hyper-converged infrastructure has been a rapidly rising star of the storage and datacentre scene in the past year or two.

And it is predicted by Gartner that this rise will continue, with hyper-converged product sales set to more than double by 2019 to around $5 billion.

As that happens, according to the analyst house, hyper-converged infrastructure will increasingly break out from hitherto siloed applications and become a mainstream platform for application delivery.

Part of that evolution surely has to be the inclusion of NVMe connectivity, which should be a shoo-in for hyper-converged.

NVMe is another rapidly rising storage star. It's a standard based on PCIe, and offers vastly improved bandwidth and performance to drives than existing SAS and SATA connections.

In other words it's a super-rapid way of connecting drives via PCIe slots.

NVMe has so far been most notably been the basis for EMC's DSSD and startup E8's D24 but has not apparently been widely taken up by hyper-converged market players.

Surely though, widespread adoption is just a matter of time. With its PCIe connectivity, literally slotting in, NVMe offers the ability to push hyper-converged utility and scalability to wider sets of use cases than currently.

There are some vendors that focus on their NVMe/hyper-converged products, such X-IO (Axellio), Scalable Informatics, and DataON, but NVMe as standard in hyper-converged is almost certainly a trend waiting to happen.

↘ **Next article**

⚑ Q&A: From backup to data management

Aaron Tan, Senior Editor, APAC

In a world where data is the new gold, data management has become more critical than before. After all, if companies do not know what lies within and even outside their walls, there is no point talking about rolling out newfangled tools to draw insights from vast volumes of data.

The data protection laws passed by governments around the world does not make data management any easier, often limiting the retention of some data that might not be useful today, but could well be decades later.

Computer Weekly sat down with Scott Anderson, global senior vice-president for information protection solutions at Veritas, to discuss these challenges, and what Veritas is doing to help organisations better manage their data.

Q: Veritas is often seen as a backup company. How is it evolving to address the needs of enterprises today?

A: We've always been known as a backup and recovery company with products like NetBackup and Backup Exec as well as our appliances. Historically, these are storage products, founded on volume management,

file system availability and clustering. What we're trying to do is to transform into an enterprise data management company with a broader remit around managing data and information for enterprises in a hybrid world where everything sits on-premise and in the cloud.

In September 2016, we launched this concept called 360 data management. One of the things we started working on was this idea that NetBackup, in particular, has visibility into a tremendous amount of data.

We protect so much of the world's data with our backup and recovery products. And we saw a lot of information about the systems that we back up, and we thought it'll be a great way to leverage that data by providing global visibility to a customer. We also restore a lot of copies of data in point-in-time images, so we thought we're going to extend the value of those images so that they are not just used for recovery.

360 data management has four elements. The first is unified data protection, which comprises backup and recovery products that are integrated with our appliances. We launched that with NetBackup 8.0, and we have had physical appliances for nearly seven years now.

It has been a very successful business for us – we've shipped well over 20,000 appliances to 5,000 customers. In the purpose-built backup appliances market, we're the market leader in integrated systems, overtaking Dell-EMC, according to IDC. We've also focused on making

ourselves relevant in the cloud, by connecting to the cloud for long-term retention. We've recently announced partnerships with Amazon Web Services (AWS) and Microsoft Azure to protect the workloads in those clouds.

The next thing is what we call global visibility. This is where we bring the information map, where we put an agent on the NetBackup server that connects to Veritas' cloud. The agent takes the metadata that NetBackup is collecting and puts it literally on a map that shows you where your unstructured data sits physically.

It also shows a bunch of attributes, like the owner of that data, when it was created and modified, as well as file type. This data allows people to do better data management. We launched that at the end of September 2016 and have about 100 customers, and we have analysed over 40 petabytes of primary data against over 50 billion files and objects.

That solution allows customers to do a number of things. One is to look at data and in our data genomics report, we found that 41% of file-based data has been modified on primary storage and is three years or older. So, you can imagine from a backup point of view, you're doing deduplication but you're still looking at a file and trying to back it up, do segment matching and say don't put it on storage. But in some cases, you're still sending data without modifications to tape every week.

Wouldn't it make more sense to archive that data, which could include orphan data of employees who have left the company? We're able identify such data, such as Outlook PST files that some companies don't allow for compliance reasons. If we find those files, they either get deleted or archived for legal discovery. These are some of the things that we do to help companies free up storage, delay future storage purchases, better govern data and optimise backup policies. I believe these are unique capabilities that Veeam, Dell-EMC and Commvault do not have.

Another part of 360 data management is integrated resiliency, enabled through a disaster recovery solution called the Veritas Resiliency Platform. It allows customers to do disaster recovery (DR) through the cloud as well as on-premise. What we found from our customers was that one of three use cases always exists around DR: they are consolidating their data centres and want to ability to move applications in bulk.

They also want to move workloads to the cloud and with our latest release, we can migrate on-premise workloads to AWS. It also allows customers to achieve DR compliance in business continuity planning, which requires them to test and rehearse DR scenarios. The integration with Netbackup also enables them to meet different recovery time objectives (RTO) for different applications. For applications that have RTO requirements that fit the frequency of backups, organisations can now use the backup copy of not just the data, but also the business service.

The last thing is Veritas Velocity, which is a copy data management solution that integrates with Netbackup. If you have a copy of an Oracle database protected by Netbackup, you can use copy data management to provision that to other business users, and create virtual copies of the data so you're not creating more storage sprawl. This is useful in big data analytics as well as testing and development.

Q: How will 360 data management enable Veritas to stand out from the competition, especially from the likes of Veeam, as well as hyper converged suppliers?

A: My belief is that in backup and recovery, we're the king of scale. We've created a dedicated business unit that includes sales and product teams to target the SMB and commercial space, and go aggressively after Veeam. We're also very successful in the cloud with our support for Microsoft Azure and AWS.

We're also supporting new requirements for backup and recovery – you have to use data more effectively beyond backup, such as orchestrating higher requirements like bringing up business services. The hyper converged suppliers are doing data compression on their side, and that's effective. We're able to make a backup copy and do deduplication (which is a commodity technology) on the client side or in the media server on the storage.

In the hyper converged space, we have a product called Hyperscale for OpenStack, which offers a software-defined view of what some of those hyper converged solutions do, which is essentially bringing together compute and SSDs, and using an orchestration layer like OpenStack to move workloads across efficiently. What we've added are data nodes below the compute nodes. We can move data down every 15 minutes – almost creating a point-in-time copy of data – thus flushing data out to enable better utilisation of expensive tier-1 storage.

Q: With more enterprises deploying internet of things devices such as sensors, how can Veritas help to manage the vast volume and high velocity of data that's being created today?
A: We have a product called Veritas Access, a NAS based device that helps to manage a large number of small files. It supports multiple protocols, as well as the ingestion and export of files via Amazon's S3 service. It gives you the ability to do information management, but I think what you're talking about from a strategic standpoint is that whether data is at-rest or transient, extracting more metadata to enable better business processes, as well as intelligence on how long data should be stored and the different tiers of storage to use, will become more important.

Q: What are the some of the biggest challenges that customers in APAC are facing from a data management perspective?

A: There are some industries like financial services and healthcare that are very global in nature. They have similar challenges such as being heavily regulated and the need to manage large-scale data. But I do see differences in the maturity scale.

When I was in Bangkok, people were talking about Thailand 4.0, which aims to modernise the country into a digital economy. Countries like Thailand will skip a few steps during the digitisation process, but they need to think about whether they are going to classify data and develop information management and retention strategies upfront.

It's important because as you're creating data, you need to have good understanding of that data to manage it. In the US and other more developed economies, there are already massive amounts of data collected, but classifying all that data takes a large amount of effort, which means it's never going to get done.

Also, organisations in countries that are spearheading digitisation efforts will also have to take into account data protection laws not only at home, but also in countries where they operate. Organisations are custodians of customer and employee data that has to be managed from both compliance and cost standpoints.

Q: How should organisations decide whether or not to keep a piece of information or data? Something that might not be useful today might well be useful in future.

A: It's an interesting question - I have a young daughter and the health information about her today might be pertinent to her when she's 35 years old. But I also think there's a lot of data that isn't medical, or is trivial and irrelevant. At Veritas, we only retain data for two years, and we have employees who have had data for 15 years and were concerned about that policy.

There's some resistance, because people fear the need to go back to e-mails 14 years ago, which nobody does or is so infrequent. To overcome that, organisations need to understand what data they have, manage risk, use that data to generate revenue for the business, and make decisions about whether to keep a piece of data. At Veritas, we help customers do a dark data assessment, so they get a view of their unstructured data to help them make policy decisions.

⤵ Next article

Getting more CW+ exclusive content

As a CW+ member, you have access to TechTarget's entire portfolio of 120+ websites. CW+ access directs you to "platinum members-only resources" that are guaranteed to save you the time and effort of having to find such premium content down on your own, ultimately helping you to solve your toughest IT challenges more effectively—and faster—than ever before.

Take full advantage of your membership by visiting www.computerweekly.com/eproducts

Images; Fotalia