

# THE COMPLETE GUIDE TO: **Hyperconverged Infrastructure**





Credit: iStock

# Believe the hype



In the space of just a few years, hyperconverged infrastructure has evolved from a niche product category to multi-billion dollar market and a key area of focus for the industry's largest tech vendors.

As well as spurring the growth of start-ups such as Nutanix – valued at \$2 billion for its 2016 IPO – it has attracted heavyweights Dell EMC, VMware and HPE, which acquired another early market entrant, SimpliVity for \$650 million last year.

And this swift growth is unlikely to slow anytime soon. IDC stats showed that sales of hyperconverged systems expanded by 65 percent year on year during the first quarter of 2017, generating \$665 million.

Driving this interest is the promise of reduced infrastructure costs and removal of some of the headaches of running on premise infrastructure – both key factors when public cloud becomes ever-more accessible to all types of businesses.

More and more businesses are starting to see the benefits. A survey by analyst firm 451 Research late last year showed that hyperconverged infrastructure is being used by 40 percent of organizations surveyed.

As the enterprise customer case studies in this e-magazine show, it offers the ability to quickly and easily set up infrastructure for use cases such as virtual desktop infrastructure (VDI). But the hyperconverged market is quickly maturing, supporting a wider set of use cases, including more mission-critical applications.

It is also just one part of a wider move towards increased automation of data centre systems, leading to greater control of resources through software, and paving the way for further innovations around infrastructure architecture models.

Read on to find out more about the basics of hyperconverged infrastructure, who the main players are in the market and how the technology could be set to evolve in future. [Matthew Finnegan](#)

---

# Contents

- 4** Hyperconverged infrastructure explained
- 11** LCG overhauls storage infrastructure
- 14** The main players
- 20** Theme park operator Merlin selects SimpliVity
- 24** Composable infrastructure



Credit: iStock

# Hyperconverged infrastructure explained

How software is transforming the data centre

**H**yperconvergence is an approach to data centre infrastructure that combines server, storage and network functions in one box, all managed through a software layer. It promises greater flexibility, scalability and ease of management. “The aim is to simplify IT,” explains Jesse St Laurent, formerly VP of product strategy at SimpliVity, now HPE SimpliVity. “The amount of complexity is extremely high in IT

infrastructure and our goal is, both from a capital and operating perspective, to take those costs down by simplifying the infrastructure.”

The trend builds on developments around converged infrastructure, a similar approach which enabled customers to purchase pre-integrated and tested hardware and virtualization products from a vendor and its partners. Hyperconverged systems – which tend to rely on cheaper, commodity hardware – takes a software-centric architectural approach, allowing ‘software-defined’ compute and storage resources to be pooled and more easily managed.

St Laurent says that whereas traditional infrastructure has “ten to 12 different things to manage – appliances and software packages and stuff like that”, hyperconverged infrastructure relies on a “single unified interface embedded into something you already use – your hypervisor management toolkit”.

### **Use cases: From VDI to Hadoop**

Early uptake has mostly been among smaller and mid-sized businesses with tighter budgets and smaller IT teams, using hyperconverged appliances as their core infrastructure. But vendors claim larger enterprises are making substantial investments too, where operational costs and the need for agility are more of an issue.

Adoption has frequently centred around virtual desktop infrastructure (VDI) in particular, though disaster recovery and remote office or branch infrastructure deployments are common too.

One of the reasons is the ability to use locally attached storage rather than storage area networks (SANs). This is due to the inclusion of ‘software defined

storage’ tools in hyperconverged infrastructure products. “VDI is a good application because it is difficult to work out what resources you need to run it,” reveals 451 Research’s Abbott. “It has been hard to do that especially with back-end SANs.”

Increasingly, a wider range of virtualized workloads, including mission-critical applications, are being run on hyperconverged infrastructure, as well as distributed NoSQL and Hadoop-based applications.

The ease with which new modules can be added means hyperconverged systems can be scaled to meet the demands of big-data analytics.

“There are an emerging set of applications – typically cloud native in nature or generated using highly defined software structures to provide resiliency – you start to see those as applications in the hyperconverged world,” says VCE’s EMEA CTO, Nigel Moulton.

This includes “apps that would run in a Hadoop distribution or that would use containerized Linux or Cassandra-style object-oriented databases as a data store – those sorts of application environments lend themselves to a hyperconverged infrastructure system.”

## An alternative to public cloud?

The growing interest in hyperconverged infrastructure highlights a wider point: not all customers are keen to push workloads out to the public cloud, with security and performance concerns persisting. At the same time, many businesses with on-premise infrastructure are keen to access some of the efficiency benefits of cloud technologies and new application models.

“[IT teams] are starting to see if they can modernize their in-house architectures so that they are at least

getting somewhere towards a private cloud,” Abbott explains. “The hyperconverged vendors are saying that this is an easier way of doing it – it is modular, supports new application models, is cheaper than traditional big systems and easier to scale.”

St Laurent claims the total cost of ownership of hyperconverged systems can actually be lower than using public cloud providers. However, he adds that cost should not be the main consideration for where to place workloads. “The message is ‘don’t assume that to get the agility you are seeking in your business, you need to rush out and move to the public cloud’,” argues St Laurent. “Often IT organizations would prefer to be internal – but it is this need for agility.”

Hyperconverged infrastructure is also seen as a way of bridging the gap to the public cloud.

The University of Wolverhampton – which has 21,000 students and 2,400 staff – is using HPE’s HC 250 hyperconverged appliance, which supports the Helion OpenStack platform. This will make it easier to move workloads to public cloud providers where necessary.

“This gives us the ability to move services from our own private cloud to public cloud and between public cloud vendors,” says Dean Harris, assistant director for ICT Infrastructure at the university.

## Automation

Harris believes that investing in hyperconverged appliances allows the university’s IT department to be far more agile in reacting to business demands, acting as a ‘broker’ to internal resources or external cloud services.

“That puts us in a very strong position as an internal service provider,” he explains.

Automation is a key aspect of this. Harris adds that easier management means the university's first-line support and service desk can maintain its infrastructure. This frees up systems engineers to "deliver projects of value, and that makes a difference to the business".

"It is not just easier to manage, it is the skill level required," he says. "Currently I have my systems engineer managing my storage environment because they need to know about firewalls, networking and LUN structures. With hyperconverged systems it is just lumps of storage, it is logical, and it is not all about the technical detail, so your level of skill can be reduced and allow your higher level of expertise staff to really focus on the most value-added things."

## Challenges and drawbacks

There are a number of areas where hyperconverged systems might not be suitable, though. Large, mission-critical applications are less well-suited to the hyperconverged model – one of the reasons why adoption has been lower among larger enterprises so far. Applications such as finance or ERP software would need to be rewritten, and currently are more suited to converged infrastructure.

"When you get to things like Oracle or SAP or the big enterprise apps they don't really run very well on hyperconverged yet," argues 451 Research's Abbott.

"It is not such an obvious fit – you need to do a lot more optimization and administration work with all of those non-distributed traditional enterprise apps with high transaction processing layers and things like that."

According to VCE's Moulton: "If you think of an application like an ERP system or something like

Oracle Financials, if either of those applications can't see the attached database, they stop working. And in most enterprise organizations that is a heart-stopping moment if it happens.”

Another challenge for businesses is a change to the procurement model. “Some people can't get their head around buying their storage and servers in one,” says Abbott. “Storage and compute does not necessarily evolve at the same pace, so if some new disk drives come out, you might want to buy those. It depends what your life cycle of purchases is as well.”

He adds that while one of the benefits of the hyperconverged model is the “modularity and simplicity”, by offering more models and configurations to customers there is a “risk of introducing complexity”.

Lock-in remains an issue too. “It depends how you view it,” explains Abbott. “Nutanix and others say that you are just buying a standard storage and server and that in some ways that's open because it's basically standard – it is not proprietary as such. But you have to buy it all from them and you have to rely on whether all vendors are going to be around over the next few years – so that is a problem. The ideal would be a separate software layer that you could run any hardware underneath and turn it into a hyperconverged system.”

### What does the future hold?

Nevertheless, it is clear that hyperconverged vendors are targeting all manner of applications.

St Laurent says that while VDI has been a good entry point for early adoption, the SimpliVity has always been focused on the “core of the data centre”. He

says customers are already running a wide variety of software, including SQL Server, Exchange, SharePoint and Oracle, as well as VDI.

It is not quite there yet, though. “Hyperconvergence needs to continue to evolve in terms of the area of IT it can cover,” he says. “Can you hold your seven-year retention deep archive data on these environments? Can you run your absolute mission-critical, sub-millisecond response time on these environments? In terms of the bell curve of the market, hyperconvergence covers some portion of that today, but what is happening is we will see it continue to push out in both directions and we will see it a natural part of the infrastructure.”

451 Research’s John Abbott believes that hyperconvergence is another step towards software controlling the entire data centre.

“I think hyperconverged will be part of that,” he says, “but I don’t think it will be all of it.”

This is because some converged infrastructure architectures also include “a lot networking stuff which the hyperconverged people don’t really address”.

“The virtual networking side is where Cisco and VCE and Oracle puts some efforts in that the hyperconverged people haven’t done yet.

“The hyperconverged people like to say that they are the next generation after converged infrastructure but I am not sure that is the case. They have done a very good of merging storage and compute, but that is only part of the story.” [Matthew Finnegan](#)



Credit: iStock

# LCG overhauls storage infrastructure

London Capital Group gutted its entire storage infrastructure and replaced it with Nutanix to cut down on costs and make life simpler

London Capital Group (LCG) provides financial services such as spread betting and employs about 100 staff, with 35 in its UK IT team. It also has offices in Cyprus, Israel and Poland.

When CIO Blair Wright joined the company, LCG's entire infrastructure was at the end-of-life stage, and was a mishmash of traditional architecture – leading to

an unnecessary level of complexity. So Wright instigated a project to replace everything from the ground up, so that the firm was able to maintain its existing setup while building fresh systems with Nutanix at the same time.

The company decided that they'd commit fully to Nutanix's hyperconverged Xtreme Computing Platform for the enterprise, which is designed to consolidate compute, virtualization and storage into one product, with the ability to scale easily.

In the 12 months since LCG decided to make the switch, the business has been able to place its entire architecture on the Nutanix platform. It also has new trading platforms running on Nutanix, and even did a building move in the middle of the deployment.

This, plus a reduced data centre footprint of 60 percent, has been "phenomenal", according to Wright.

"Originally we were looking at Cisco equipment – Cisco servers backed with NAS or SAN behind it," Wright says. "Then we met the guys from Nutanix and had a chat with them, which for me was a different view on what we wanted to do."

"It's not something we could have considered doing had we gone for a traditional architecture," he explains. "Even half of that time would have probably been lost in the original installation and setup phase."

"I needed one system to manage rather than having multiple systems," Wright says. Now, in just over a year, the business' infrastructure is all running Nutanix and fully operational. "As of last week they cleared the last cage of old equipment."

LCG knew its systems needed rehailing and so did discuss other vendors – but the simplicity of both installation and management that Nutanix offered won

out. “What made a massive difference for me was that I didn’t have to have a specialized storage team just to run our storage solution,” Wright says. “The core thing for us has been not having to deal with a separate storage solution – having it all in one product, with support from Nutanix, has made our lives a lot easier.”

“The fact that as a provider they understand virtualization and understand VMware very well – they look at the virtualized world and how it interacts with everything – for us that support is invaluable.”

LCG is also running VMware NSX, which enables the microsegmentation of each system, providing a firewall around every server. “I haven’t worked anywhere else where we have the security we have today,” Wright adds.

He says that initially, because Nutanix is a new provider, he needed convincing. But when LCG looked at the traditional infrastructure, it all would have taken more time to deploy – plus the time to hire skilled staff.

“We were in a hurry to get the project done so we could get on with building the rest of the business,” he says. “The more we looked at it, the more we thought about it, the more it became a clear decision – if we do go with the technology, it will save us six months over going with something older.” [Tamlin Magee](#)



Credit: iStock

# The main players

Seven of the best hyperconverged vendors and products

**H**yperconverged infrastructure is increasingly being adopted by organizations of all sizes as they seek to simplify their back-end systems and cut down on power usage and rack space.

Analysts are expecting hyperconverged to continue to take off in a big way. Gartner, for instance, expects that the market will be worth nearly \$5 billion by 2019

when uses for the technology extend out into more mission-critical enterprise applications.

Here's a round-up of some of the main hyperconverged infrastructure vendors that your business needs to consider.

### Atlantis Computing

Atlantis was founded in 2006 in Mountain View, California, and counts more than 800 medium to large enterprises among its customers.

The company's USX allows for local servers with RAM, SAS, Flash or memory-channel storage to become a scalable and hyperconverged platform with existing servers – so they can build on existing infrastructure to get the benefits of hyperconverged.

According to the company, it's possible to get as much as 300TB of effective storage capacity in a three-server configuration. The company claims its offering can manage storage performance of as many as 3.3 million IOPS in a 3-blade cluster, and can boost the performance of VSAN clusters by five times. And the firm claims its Atlantis HyperScale product is the first on the market to offer all-flash hyperconvergence that integrates servers, storage and virtualization in an "economically viable" way.

Atlantis also supports multiple hypervisors, plus it supports Cisco, HP, Lenovo and SuperMicro server platforms.

### Dell EMC

EMC subsidiary VCE introduced its hyperconverged rackscale system VxRack in early 2015. According to the company, VxRack was specifically designed for the

simplification of deploying mobile, cloud and distributed Tier 2 applications – with the chief USP being the ease at which the user can scale out to thousands of servers.

The benefit of the business being so integrated with EMC means VxRack systems work with EMC ScaleIO, and the system itself is available with a choice of hypervisor, including VMware vSphere, KVM, or bare metal. VxRack is also optimized for VMware environments, the company says.

More recently it launched its VxRail appliance, targeted at smaller and mid-market customers and jointly engineered with VMware.

Also available is the Dell EMC XC Series of hyperconverged appliances, best suited for hypervisor support.

## Hitachi Data Systems

Japanese technology giant Hitachi launched its Hyper Scale-Out Platform in 2015 – hyperconverged architecture designed specifically with crunching big data in mind, across a distributed clustered architecture.

It's complemented by open source management and virtualization software, and according to Hitachi Data Systems, it's perfect for Hadoop environments – killing the need to move large data sets around before performing analysis.

HDS announced some updates at the start of 2017 with the HSP 400 series, including native integration with the Pentaho Enterprise Platform for big data. But HDP wants the HSP series to be the go-to option for a variety of workloads in the future. This year HSP is available in two configurations, with Serial Attached SCSI disk drives available now and an all-flash flavour due out soon.

## HPE SimpliVity

Founded in 2009, SimpliVity from Boston entered the hyperconverged market in 2012 with its OmniCube product, built on the OmniStack data virtualization platform. Again, the idea is to consolidate IT infrastructure below the hypervisor – and the company boasts a careful deployment of the plug-and-play gear which it claims could save up to three times in total cost of ownership (TCO) savings.

VP of product strategy Jesse St Laurent told us that the TCO of running on SimpliVity can in theory be lower than picking public cloud providers, but crucially it is a decent bit of agile infrastructure for internal on-premises operations because the cubes can be added or removed as needed. SimpliVity also offers OmniCube products that are integrated with the x86 Lenovo System, x3650 M5 servers and Cisco's UCS systems.

HPE picked up the company in January 2017 for \$650 million, and HPE SimpliVity now offers the HPE SimpliVity 380, which is described as dramatically simplifying IT by “combining infrastructure and advanced data services for virtualized workloads onto the best-selling server platform in the market.”

Although HPE bought SimpliVity, it maintains a wider portfolio of HC offerings.

HPE announced its Hyper Converged 380 infrastructure, specifically designed with the mid-market in mind and inspired by consumer technology – in other words, it's all supposed to be easy to get to know and you can move around it in a few clicks. It's built on top of the HPE ProLiant DL380 server, and packs other HPE features such as HPE StoreVirtual VSA for moving data

around. HPE promises it's a good option for disaster recovery and scales out easily, plus offers dedicated customer support for HPE hardware, software, private cloud, and VMware environments.

## Nutanix

Founded in 2009, San Jose-based Nutanix brought its first hyperconverged product to market in 2011. Although it carries a certain startup heritage, it has made a considerable effort to move into the Global 2000 enterprise space and now caters to that market, as well as SMB customers.

Nutanix offers two product families, Prism and Acropolis. The Prism line is designed to simplify the managing of virtual environments using a mixture of automation, plus remote management software and end-to-end analytics. The company pitches Acropolis as a feature-heavy software defined storage offering with built-in virtualization.

Together Nutanix claims its products are suitable for powering VDI, server virtualization, big data analytics, enterprise applications and running cloud environments. Because of its heritage, Nutanix has a certain weight behind its brand, but is known to be rather expensive out of the box, though ultimately you could save on power, staff and training. Nutanix powers other offerings too including the Dell XC Series products.

## Pivot3

Austin-based hyperconverged infrastructure business Pivot3 was founded in 2003 and introduced its first HCI offerings in 2007, making its mark in the industry with storage offerings designed to help surveillance

companies manage large amounts of data. Its vSTAC software platform for virtual storage and compute later became the company's flagship hyperconverged product. It built on this with the new Acuity software platform, announced in April 2017, which the company boasts is "priority-aware" and policy-based, and is "architected to overcome the performance, automation and consolidation limitations of conventional HCI".

### Scale Computing

Scale's tag line is 'virtualization made easy' and its HC3 product stripes and mirrors data across all of the drives in the cluster, meaning there's no single point of failure, according to the company. And if there is a failed node, it will figure out placement with available resources and shift VMs over to those nodes.

Scale says the clustered architecture makes its product reliable in the event of disaster recovery and also provides the scalability that's needed in highly available environments – all underlined with that common thread for hyperconverged, simplified management. Recently, the company introduced hybrid storage plus HyperCore Enhanced Automated Tiering to optimize data across the storage tiers and within the cluster. [Tamlin Magee](#)



Credit: Merlin Entertainment Group

# Theme park operator Merlin selects SimpliVity

Merlin picks SimpliVity to simplify its operations and support expansion

One of the things that struck me about Merlin was just how quickly they move when it comes to opening new sites and new attractions,” explains Sean Channon, global infrastructure architect at the firm.

“It’s just a really fast-paced business from that point of view.” However, he says that building attractions is the priority and that “IT systems are usually the last thing to go in”.

So to be able to prepare and deploy IT rapidly on an as-and-when-it's-needed basis was appealing.

“We can essentially build all the virtual servers and infrastructure for any attraction now in the data centre,” reveals Channon.

“We can do all our testing, the pre-live tasks, on those virtual servers – as soon as a site is handed over to us from a general contractor we can drop a SimpliVity OmniCube on-site, deploy it, and I think that takes all of about 30 minutes.”

“We can just take the virtual machines and systems down to the site,” he says. “It’s made meaningful savings for us, in terms of the amount of time it takes for us to stand up an attraction.”

Rather than picking any of the big vendors who were testing the waters in hyperconverged infrastructure or rival Nutanix, Merlin chose SimpliVity because of its deep integration with VMware and for the reliability it offered over alternative cloud solutions. Merlin has already made considerable space and cost savings, according to Channon.

“We started looking at SimpliVity in January 2015,” he says. “We ran quite an extensive proof of concept for a couple of months until about May 2015, and then our first live deployments were in June 2015, to Legoland California.”

### **Infrastructure set up in ‘under an hour’**

According to Channon, Merlin has 125 attractions at the moment, varying in size from the smaller Madame Tussauds-style attractions all the way through to the larger theme parks. Of these, across Europe, the UK, the US and Asia, Merlin is running 32 SimpliVity

OmniCubes across 24 sites. So although the business is in the process of rolling out SimpliVity to replace a hodgepodge of older legacy systems, an advantage is that the back end work can be performed and then optimized on virtual machines, so when a site's ready, it's just a case of bringing a box into the location, plugging and playing. This can take under an hour.

"It's the one thing that surprised us – scale – just how quickly from June last year we've managed to get so many boxes out into production," Channon says.

"I think that's testament to the fact it's very easy to deploy, and being VMware as it is, migrating onto it has been fairly easy as well."

From an operational point of view, it has "massively reduced" the complexity of Merlin's environments.

"If you looked at how our sites were composed previously, some of them were running purely on physical hardware, others were running HyperV or VMware on traditional servers," he says.

"Some attractions were a mix of both. So being able to take the complexity out and no longer having to worry about storage arrays and all the other intricacies required has reduced our day-to-day admin."

An example where the deployment has also helped in cost is the first attraction running SimpliVity – Legoland California.

"We essentially took them down from around three racks of IT equipment to just under half a rack when we were finished," Channon reveals.

And with a "fairly small" operations team trying to run the estates from the UK, having trimmed down the daily operations using this platform frees up IT staff to focus on providing better services.

## Selecting SimpliVity over Nutanix

Why SimpliVity over the big vendors, all of which are now taking HCI as a more serious prospect? Or Nutanix, the other top player in the HCI space?

According to Channon, when Merlin was looking into the competition at the start of 2015, the more established vendors just were not offering anything as appealing. VMware's EVO RAIL was kicking around, and there were some other propositions, but Channon claims these were all "fairly piecemeal" and there were add-ons that needed to be bolted on.

"They weren't as tightly integrated and that for us was a defining decision," he says.

Merlin did consider Nutanix, however.

"We had a very brief look at Nutanix," Channon explains. "With them versus SimpliVity, it was still going to be separate siloed infrastructure in terms of management, whereas with SimpliVity we can connect to one virtual central server and manage not only our SimpliVity estate but our more traditional IT as well."

Channon says he was not concerned about picking SimpliVity over more established IT vendors.

"You're always going to have concerns about working with a new and young product," he adds.

"But we definitely put them through their paces when we did the proof of concept. We went through as many scenarios as we possibly could to try and find a fault in the product – and we couldn't find anything that would prevent us from rolling out to production." **Tamlin Magee**



Credit: HPE

# Composable infrastructure

Data centre infrastructure has evolved quickly in recent years, with hyperconverged systems offering greater control

According to HPE the latest step in this evolution towards a fully ‘software-defined’ data centre is composable infrastructure, namely its Synergy platform, announced in 2015.

The idea behind composable systems is that hardware components – compute, storage and networking – are managed entirely via software commands. Pools of resources are created and assigned automatically to meet the demands of individual applications in near real-time.

“We are now at the point with composable infrastructure where you have infrastructure as

code: literally a single command line or line of code to deploy a piece of resource to allow you to run your infrastructure,” says David Chalmers, HPE’s vice president and chief technologist, EG EMEA.

Analysts believe that there will be a lot more action from the vendor community in future.

A research note from Moor Insights & Strategy last year states: “Over the next 12 to 24 months, the market is expected to ramp with additional new products, tighter integrations across vendors, and usability enhancements to make deploying and managing composable infrastructure easier for mainstream IT organizations to adopt more widely.”

### What are the benefits?

There are a number of potential advantages, according to Chalmers. Greater utilization of resources is likely to result in cost savings with regards to capital expenditure.

It can also speed up the deployment of applications, helping to meet demands from the wider business to create new products. In addition, composable systems allow IT teams to bridge the gap between older systems and newer cloud-native applications.

“We see Synergy as being the way to give the on-premise customer the flexibility of use that they need,” says HPE’s Chalmers. “It is not just buying one style of system for ERP and a different set of systems to run CRM, and then something else for different application modes. When you do that in the converged or even hyperconverged world, you end up with a very complicated mixture of systems. Composable is about saying ‘don’t do that, put them in a physical infrastructure that you can change very quickly’.

“So when ERP systems need [resources] because it is the month end, that is fine, we move the resource around. If it doesn’t need more because it is the first week or the first quarter, we divert that to the marketing application and so on and so forth.”

All of this, argues Chalmers, provides a compelling alternative to moving workloads out to the public cloud. Despite the claims of public cloud vendors, many businesses will want to keep at least some of their systems on-premise, particularly for regulatory purposes.

“People are recognizing that being focused on any one version of the approach – it is all public or all private – is not the way to go. Composable infrastructure gives people a dramatically more effective option on-prem to get that flexibility and control they need.”

### ‘Hyperconvergence’ rebranded?

The ‘composable’ concept bears a resemblance to hyperconverged systems that have brought improved software management to hardware systems.

However, there are some key differences. In contrast, says Chalmers, hyperconverged systems are more suited to single workloads that require scalability.

“We see a lot of people talk about ‘is this just a fancy way of describing hyperconverged’, and we would say no,” explains Chalmers.

“Hyperconverged is a step on the way. It tends to be where you have the same style of application running on all of the nodes. Because the nodes are relatively fixed in size, the whole point about hyperconverged is that it is incredibly simple, but it is also not designed to be flexible. It is designed to allow you to run a similar workload in a very easy way.”

Moor Insights & Strategy's note suggests that organizations already using hyperconverged systems are "a good target for a transition to composable infrastructure".

It states: "These customers are already familiar with the benefits of a single platform with server, storage, and network resources – and the efficiencies from an application-centric approach to resource deployment. Organizations that decide to take the next step to composable infrastructure from a converged or hyperconverged platform could experience additional benefits by pooling these resources and dynamically provisioning/reprovisioning their infrastructure as workload needs evolve."

### **A truly software-defined data centre**

Looking at the bigger picture, Chalmers sees composable infrastructure as another stepping-stone in the evolution of server architectures to the memory-centric – as opposed to processor-driven – architecture exhibited in HPE's 'The Machine' project.

"We had mainframes and we had minis and we had client servers, then we had server-client, Internet and converged, now we have composable. So what comes next?" Chalmers asks. "The answer is 'mesh' and memory-centric computing that is built around fundamental architectures turned round to be memory-centric, with different pools of resource. But it is built upon the composable concept of fluid pools of resource: instead of having processors tied tightly to memory, you separate memory from it." **Matthew Finnegan**



©2017 International Data Group.