

In the dark about data storage?

Expert vendors such as Micron say there is an ongoing challenge in optimising solution architectures due to the fact that more options are now available for compute, storage and networking. PHOTO: MICRON

Flash. On-premise. Cloud. Hybrid. With a variety of different storage options now available, JAMES HAYES helps enterprise and data centre managers see the light when it comes to making the right choice.

Data centre veterans will recall the days when the specification of data storage provision was a relatively straightforward task. This was the era when storage repositories simply had to soak-up (or back-up) enterprise data on a relatively limited range of device formats. Back then, RAID promised to meet all their data storage needs for the foreseeable future, and ‘silo’ had yet to become a dirty word in the IT lexicon.

Today, beyond that foreseeable future and with a huge range of products/formats available and strategic imperatives for storage strategies to be built around, the optimisation of enterprise data storage proves as much of a challenge as the build and management of any other aspect of critical data centre infrastructure.

“The enterprise storage solutions market is more complex than ever,” says Paul Timms, MD at Buckinghamshire-based IT services specialist MCSA Group. “With ever-increasing options and choices available it’s become a minefield – and hence organisations sometimes delay making hard decisions about their data storage. The upside is that with hypervisors now being widely used, plus the maturity

around software-defined storage (SDS), the hardware vendor becomes less fixed. It’s easier to match-and-mix storage, server and network vendors.”

But with flexibility comes an inevitable degree of complexity as data centres endeavour to achieve performance parity between servers, networks and storage devices, says Doug Rollins, senior technical marketing engineer at US headquartered memory firm Micron. He reckons there is an ongoing challenge

“We have seen vendors launch strategies where we can put pay-per-use models on the customer site, with options to buy and therefore ‘sweat the asset’ at the end of the arrangement.”

Paul Timms,
MD,
MCSA Group

in optimising solution architectures due to the fact that more options are now available for compute, storage and networking. “We see increased interest in reliance on public cloud options, and pre-validated private/hybrid cloud solutions using reference architectures – those from Micron, local partners, for example, or the likes of VMware *Ready Nodes* and Microsoft *Azure Stack*. When faced with

many options, these approved designs can provide a starting point to optimise processing performance across node and network.”

Given the complexities of determining an optimal enterprise storage mix, data centre managers are well-advised to beware of the common pitfalls that the storage decision-making process can fall into. Increasingly, storage strategy is defined by specific needs of critical applications: alignment of storage solution options with those needs constitutes serious study for data centre techies – especially if their facilities are considering a transition to HPC platforms or hyper-converged infrastructures (HCIs) which would result in the adoption of SDS.

“‘Price-per-gigabyte’ used to be how storage was bought,” says Rollins. “Now we must think differently and not use old metrics for new challenges. Start by looking at what basic storage architectures are recommended by the most important applications. Oracle RAC, for example, is most proven on traditional arrays, while open source databases increasingly focus on hyper-converged or scale-out x86 architectures.”



From there, he says IT decision-makers can fine-tune the types of storage they need such as, for example, NVMe (non-volatile memory express) or SATA based on reference architectures and other tools available.

According to MCSA, another common pitfall for organisations is that they misjudge the amount of storage they will actually need, largely because it is hard to plan what the business will require over the next four-to-seven years – a typical storage investment cycle span.

Timms says: “We have seen vendors launch strategies where we can put pay-per-use models on the customer site, with options to buy and therefore ‘sweat the asset’ at the end of the arrangement. We’ve also seen an increase in lease models where customers are more comfortable in spreading the cost of storage solutions over the period of their active life, but [while] ensuring data sovereignty as well as having a choice of hardware and management software.”

He adds that the advent of cloud storage has brought about another change in the way customers can choose to pay for next-generation storage needs.

However, here Mark Scaife, head of cloud practice at Daisy Group, warns managers not to directly compare enterprise IT storage platforms to public cloud storage: “Tried-and-tested SAN or NAS devices still have a place in an on-premise or data centre infrastructure solution. Although their perceived cost-per-GB is higher than most cloud alternatives, there’s still a requirement for blistering IOPS and storage efficiencies that only local devices bring, along with storage protocol choice.”

Meanwhile, Dell EMC believes that the decision to use a public, private or hybrid cloud storage option depends on what an

organisation needs from cloud. Rob Lamb, the company’s CTO and cloud business director, reckons it generally comes down to application/workload requirements, as well as a degree of risk appetite. “If an application is latency-sensitive, then tiering some of its storage into a cloud service may cause operational and user-experience challenges. One benefit of cloud storage is that it’s easily consumed on an as-needed basis. That same ease of expansion can, however, lead to data sprawl and costs can mount quickly.”

Flash the cash?

Despite multiple arguments in flash’s favour, an ongoing thorny area for many data centre leaders is demonstrating quantifiable return on investment on hybrid- or all-flash storage adoption. C-suite executives want proof that this stuff delivers value for money, and that their IT function is not being lured into buying kit that, once installed, cannot be shown to have effected improvements.

US-based Excero, which has created what it describes as a “Software-Defined Block Storage” solution, says IT decision-makers must factor in the impact such technologies will have on the data centre infrastructure as a whole and not just evaluate data storage in isolation. According to Kirill Shoikhet, the firm’s chief architect, widespread adoption of flash in the data centre was seen as a revolution, while the transition in the way flash is accessed – from NVMe to SATA/SAS – is seen as an evolution.

“For example, while an improvement in storage media should, in theory, have a higher impact than an improvement in access protocol, the transition to NVMe is in fact significant: it shifts

system bottleneck locations, and provides separation between architectures created for flash and older architectures adapted for flash usage [usually as a cache extension].”

Shoikhet predicts that this separation will grow when persistent memory becomes mainstream due to changes in the way the storage performance is consumed, i.e., a significant skew toward higher write bandwidth.

For older deployments, Daisy Group’s Scaife points out that it’s a challenge to convince the C-suite on ROI because the ongoing maintenance of traditional SAN/NAS devices gets more expensive year-on-year. Plus, with the power and cooling costs associated with HDD versus more power-efficient and more performant SSD-based/hyper-converged solutions, he reckons it’s a losing battle. “Typically, in an average solution, we are almost at parity in the overall cost of an SSD solution to HDD although the latter does still have the edge in overall capacity density. But this comes with a footprint overhead and enterprises look to cloud to also offset [physical] placement of ‘cold’ storage.”

Shoikhet explains that all-flash arrays (AFAs) provide TCO improvements due to their ability to deliver a much higher read performance, especially random reads, which allows for the consolidation of workloads with a much lower footprint. But he goes on to point out that AFAs’ write-performance is still determined, mostly, by a write cache layer and this allows older storage architectures to compete. “The introduction of persistent memory as a major layer in storage hierarchies will require the flash layer to provide much improved write bandwidth which will change, at the same time, the storage service mix required from this layer.”



“Converged infrastructures bring an inherent rigidity in storage-to-compute and storage-to-network ratios which may – in some cases – impede the maximisation of flash utilisation.”

*Kirill Shoikhet,
Chief architect,
Excero*

Shoikhet is not altogether convinced that convergence is key for maximising the utilisation of flash/SSD storage in all cases: “I agree that the traditional siloed storage arrays cannot provide the agility and flexibility in storage allocation and access which are required

KVM CHOICE
Total Control in Computing

Specialist suppliers of Datacentre equipment call for a quote today!

HASSLE FREE PROCUREMENT OF: IT / NETWORK / POWER / INFRASTRUCTURE EQUIPMENT

0845 899 5010

sales@kvmchoice.com | sales@pduchoice.com
www.kvmchoice.com

myUTN-800

myUTN-80

SPECIAL PRICE

Dongleserver + Service^{plus} Bundle

Manage and Control Dongles in Network
+ Serviceplus Warranty Extension and
Advance Hardware Replacement

SEH Technology UK Ltd.
Phone: +44 (0) 1273-2346-81
E-Mail: info@seh-technology.co.uk
www.seh-technology.co.uk

to provide maximum utilisation of flash media, and that SDS infrastructure is the [prime] component in modernisation of data centre IT infrastructures. However, converged infrastructures bring an inherent rigidity in storage-to-compute and storage-to-network ratios which may – in some cases – impede the maximisation of flash utilisation.”

ROI equations around whether to transition to flash for enterprise storage arrays therefore continue to exercise decision-making processes. In this context, Timms emphasises the importance of not grasping at projections, and basing storage returns on line-of-business benefits. “Individual businesses have their own measurement around ROI on data storage. That said, we see flash become a must-have where competitive advantage is imperative, and speed of data transactions are critical to success.

“It’s difficult to say that spending £x will generate £y; however, it’s easier to say, ‘if we don’t do this then our competitors will gain significant advantage – therefore we can’t afford not to do it’. This is particularly the case in finance and Big Data-heavy sectors like retail and research/academia. Plus, flash is much more affordable now than it was just a few years ago.”

Micron’s Rollins is likely to agree here: “The value of flash goes up as costs go down and applications become focused on SSD performance optimisation. As applications include more measurement tools, there are more options for IT to model and validate the right architecture, and prove ROI for projects.”

But Thomas LaRock, head geek at SolarWinds, suspects this might not quite be the case. He believes it has become difficult for the IT function to measure and quantify storage ROIs, and adds that more data centres are adopting SSD/flash storage sometimes as part of an expensive shift towards a HCI model. “Often, monitoring tools have not evolved at the same pace – making it difficult to demonstrate ROI and to do so quickly. Add to this the fact that Microsoft and Amazon have made storage an affordable option for businesses [that arguably make it] easier to evaluate the costs, benefits and risks.”



IT decision-makers are advised to factor-in the impact that technologies such as hybrid- or all-flash arrays will have on the data centre infrastructure as a whole and not just evaluate storage in isolation. PHOTO COURTESY OF EXCELERO

The ‘third way’

While the pro-flash arguments might be gaining ground in terms of capex justifications, the advocacy of transition to software-defined storage models could prove more exacting.

SDS is sometimes cited as a ‘third way’ for data centre storage strategies even though, according to some industry opinion, confusion exists in the minds of many data centre managers about what it actually is and does.

“The key confusion is the definition,” says Timms. “Where some refer to SDS as a completely new type of storage, others say that it is the software which is key as it manages that underlying storage – and therefore the hardware functionality doesn’t matter.

“In our opinion, SDS is a new type of agile storage infrastructure; an enterprise platform layer that allows the use of multiple types of underlying storage

hardware, providing a storage platform which can easily be upgraded. The storage is provisioned using policies allowing users to become more agile in how they take advantage of virtualisation without requiring the purchase of new hardware.”

Scaife says that while some IT managers get SDS and are using it, others have no idea at all. He reckons this could be because SDS and cloud storage caused a market divergence. “If cloud had not happened, SDS/HCI would have had a much bigger impact in the data centre. Trouble is that cloud got in the way between the virtual- and converged eras, and now [some] IT managers are confused as to the best path to take.”

He continues by saying that some vendors are making a push for SDS/SDN, with niche players that are then gobbled-up by big fish. “However, the divergence is that most hardware/software vendors are spending more time developing cloud-based overlay solutions or hybrid software/migration

tools and have taken their eye off of the SDS/converged market. They are therefore being squeezed-out on price and a range of innovative, easily on-boarded alternatives.”

Comparing the deployment of traditional storage arrays connected via a separate SAN to that of a SDS infrastructure, which uses the same network infrastructure as business applications, again shows how much more challenging it has become to quantify and justify ROI. Shoikhet says: “A converged infrastructure might, at first glance, look like an easy case to measure, until you realise that you also need to take into account the effect of storage access requirements to the network infrastructure – how is the network configured to allow low-latency storage access to flash drives, for instance? And how is it over-provisioned to support bursts?”

He goes on to say that the expected performance of flash drives-per-storage nodes needs to be balanced with the number of PCIe lanes on one side, and the



“If an application is latency-sensitive, then tiering some of its storage into a cloud service may cause operational and user-experience challenges.”

Rob Lamb,
CTO,
Dell EMC

network interface bandwidth on the other. In addition, the choice between converged and disaggregated SDS, or a mix of both, becomes part of the analysis: “The storage-to-compute ratio implied by a choice of a converged building block may cause an imbalance in predominately storage- or compute-oriented environments. So, instead of being a separate item to analyse and quantify, storage infrastructure has become intertwined into other parts of the data centre infrastructure.”

While SDS is not necessarily synonymous with transition to HCI, hyper-convergence itself is increasingly associated with compute-intense HPC-driven applications like in-memory data analytics (think SAP HANA), high-frequency trading, and emergent AI-enabled applications. As a result, any capex entailed probably shouldn’t be considered

in isolation from the overall investment in infrastructural upgrade geared toward attainment of competitive goals.

It’s not just about raw compute-intense applications shaping the direction of storage provision: IoT is another change force phenomenon that will impact data storage deployment, predicts Dell EMC’s Lamb. “Gartner reckons that by 2020 there will be more than 20bn internet-connected ‘things’: most of these will generate data. We need to manage that generated data, and analyse it to make meaningful business decisions.”

The problem here is that pulling all these datasets back to a core enterprise data centre can be expensive in terms of bandwidth, storage and management costs. Lamb therefore proposes an approach that has edge gateways aggregating and analysing the data at, or close to, the point of origin. These gateways should only send on meaningful data to cloud or control centre. In these central locations, data lakes and HPC-optimised applications will then receive the data from the edge.

If Lamb is correct then the data that is not relayed back to base will still have to be stored, which suggests a need for decentralised ‘non-critical’ storage resources – possibly based on older platforms that perform well enough to be repurposed for the task.

Thus an SDS adoption in line with HCI transition need not necessarily signal a migration from *in situ* storage assets. As MCSA’s Timms points out SDS allows mixing of hardware and media – i.e., mixing flash and spinning disks – to provide a heterogeneous pool of storage resource. “Overall, this provides benefits in efficiency and cost, and the possibility of utilising existing hardware within data centres before considering full replacements.”

Guy England, director at Lenovo Data Centre Group, is likely to support this view and says HCI can co-exist with standard storage. But he also warns that HCI does not always address today’s underlying data expansion challenges. “Look at the progression of IT: every innovation produces vast volumes of data. Classic three-tier architectures have been doing their best to handle this scale and need for performance through AFAs and improved

Time to get flash?

Some of the latest flash-based storage units to look out for.



AccelStor’s *NeoSapphire H710* is an all-flash array in a 4U rack mount appliance. It’s claimed to deliver more than 1.2 million IOPS for 4KB random access, and comes with 10GbE or 16Gb Fibre channel connectivity, AccelStor’s *FlexiRemap* technology, and is claimed to offer “true high availability with no single point of failure” or downtime.

The *K2.N* from **Kaminario** comprises controller nodes (c.nodes) and media nodes (m.nodes) that are connected via a shared NVMe over fabric (NVMeF). When customers purchase two or more c.nodes and any number of m.nodes, Kaminario says they can create a truly active-active, scale-out storage cluster with a shared data reduction space.

Toshiba Memory Corporation claims its *PM5* 12Gbps SAS series and *CM5* NVMe series are the first enterprise-class SSDs with 64-layer 3D flash memory. Offering up to 30.72TB in a 2.5-inch form factor, the *PM5*s are also said to be the industry’s first MultiLink SAS architecture.

Toshiba says they are able to deliver up to 3,350MBps of sequential read and 2,720MBps of sequential write in MultiLink mode, and up to 400,000 random read IOPS in narrow or MultiLink mode. Available in capacities from 400GB to 30.72TB

Meanwhile, the dual-port PCIe Gen3 x4 *CM5* series is NVMeF-ready with

controller techniques, for example, but it hasn’t been enough. Instead, the most effective solution to manage this increased data demand is through the hybrid of SDS for its performance and cloud for its cost efficiencies.”

“To get back to pitfall-avoidance, my first maxim is this: know your data,” counsels LaRock at SolarWinds. “Not all data is created equal. Some storage

scatter-gather list and controller memory buffer features. The SSDs offer up to 800,000 random read and 240,000 random write IOPS for the five drive writes per day (DWPD) model, and up to 220,000 random write IOPS for the three DWPD model, both with a maximum power draw of 18W. Capacities range from 800GB to 15.36TB.

Tegile Systems’ dual-controller *IntelliFlash N Series* features 24 NVMe NAND flash SSDs in a 2U footprint, and is said to be capable of delivering up to three million IOPS with consistent 200 microsecond latency. The line up includes the *N5200* which is available in drive capacities from 960GB to 7.68TB, and the *N5800* which offers 800GB to 6.4TB.



FlexDrive enables users of **Tintri** arrays to expand capacity by adding as little as one drive at a time. The firm says because its products operate at the VM-level, they eliminate the need for traditional storage constructs such as RAID groups and shelves. As a result, customers can purchase a partially populated Tintri *EC6000* all-flash array, and add capacity by inserting drive(s) into any available disk slot.



solutions are built around the idea that data is just made of 0s and 1s – that’s not true. Some data needs fast access, some slow. Some data needs quick recovery, other data is less urgent. The common mistake here is applying a one-size-fits-all strategy which could see the IT function overstretch already shrinking resources and, ultimately, lead to department losses.” ■

Rittal – The System.

Faster – better – everywhere.

Data Centre In A Box

All the key data centre capabilities.....

.....just on a smaller scale.



Challenging the Edge:

The “Data Centre in a Box” concept enables equipment to be deployed in non-traditional Data Centre environments.

- TS-IT rack platform
- Demand-orientated climate control
- System monitoring
- Intelligent power rails

ENCLOSURES

POWER DISTRIBUTION

CLIMATE CONTROL

IT INFRASTRUCTURE

SOFTWARE & SERVICES

FRIEDHELM LOH GROUP

www.rittal.co.uk

