CIO × Microsoft × intel × snp technologies inc.

Innovate with Hyperconverged Infrastructure

8 reasons to consider hyperconverged infrastructure for your data center

Hyperconvergence offers speed and simplicity compared to the old ways of server configuration.

| By Andy Patrizio, Network World

Demand for on-premises data center equipment is shrinking as organizations move workloads to the cloud. But on-prem is far from dead, and one segment that's thriving is hyperconverged infrastructure (HCI).

HCI is a form of scale-out, software-integrated infrastructure that applies a modular approach to compute, network and storage capacity. Rather than silos with specialized hardware, HCI leverages distributed, horizontal blocks of commodity hardware and delivers a single-pane dashboard for reporting and management. Form factors vary: Enterprises can choose to deploy hardware-agnostic hyperconvergence software from vendors such as Nutanix and VMware, or an integrated HCI appliance from vendors such as HP Enterprise, Dell, Cisco, and Lenovo.

The market is growing fast. By 2023, Gartner projects 70% of enterprises will be running some form of hyperconverged infrastructure, up from less than 30% in 2019. And as HCI grows in popularity, cloud providers such as Amazon, Google and Microsoft are providing connections to on-prem HCI products for hybrid deployment and management.

So why is it so popular? Here are some of the top reasons.

1. Simplified design

A traditional data center design is comprised of separate storage silos with individual tiers of servers and specialized networking spanning the compute and storage silos. This worked in the pre-cloud era, but it's too rigid for the cloud era. "It's untenable for IT teams to take weeks or months to provision new infrastructure so the dev team can produce new apps and get to market quickly," says Greg Smith, vice president of product marketing at Nutanix.

"HCI radically simplifies data center architectures and operations, reducing the time and expense of managing data and delivering apps," he says.

2. Cloud integration

HCI software, such as from Nutanix or VMware, is deployed the same way in both a customer's data center and cloud instances; it runs on bare metal instances in the cloud exactly the same as it does in a data center. HCI "is the best foundation for companies that want to build a hybrid cloud. They can deploy apps in their data center and meld it with a public cloud," Smith says.

"Because it's the same on both ends, I can have one team manage an end-to-end hybrid cloud and with confidence that whatever apps run in my private cloud will also run in that public cloud environment," he adds.

3. Ability to start small, grow large

"HCI allows you to consolidate compute, network, and storage into one box, and grow

this solution quickly and easily without a lot of downtime," says Tom Lockhart, IT systems manager with Hastings Prince Edward Public Health in Bellville, Ontario, Canada.

In a legacy approach, multiple pieces of hardware – a server, Fiber Channel switch, host-based adapters, and a hypervisor – have to be installed and configured separately. With hyperconvergence, everything is software-defined. HCI uses the storage in the server, and the software almost entirely auto-configures and detects the hardware, setting up the connections between compute, storage, and networking.

"Once we get in on a workload, [customers] typically have a pretty good experience. A few months later, they try another workload, then another, and they start to extend it out of their data center to remote sites," says Chad Dunn, vice president of product management for HCI at Dell.

"They can start small and grow incrementally larger but also have a consistent operating model experience, whether they have 1,000 nodes or three nodes per site across 1,000 sites, whether they have 40 terabytes of data or 40 petabytes. They have consistent software updates where they don't have to retrain their people because it's the same toolset," Dunn added.

4. Reduced footprint

By starting small, customers find they can reduce their hardware stack to just what they need, rather than overprovision excessive capacity. Moving away from the siloed approach also allows users to eliminate certain hardware.

Josh Goodall, automation engineer with steel fabricator USS-POSCO Industries, says his firm deployed HCI primarily for its ability to do stretched clusters, where the hardware cluster is in two physical locations but linked together. This is primarily for use as a backup, so if one site went down, the other can take over the workload. In the process, though, USS-POSCO got rid of a lot of expensive hardware and software. "We eliminated several CPU [software] licenses, we eliminated the SAN from other site, we didn't need SRM [site recovery management] software, and we didn't need Commvault licensing. We saved between $25,000 and $30,000 on annual license renewals," Goodall says.

5. No special skills needed

To run a traditional three-tiered environment, companies need specialists in compute, storage, and networking. With HCI, a company can manage its environment with general technology consultants and staff rather than the more expensive specialists.

"HCI has empowered the storage generalist," Smith says. "You don't have to hire a storage expert, a network expert. Everyone has to have infrastructure, but they made the actual maintenance of infrastructure a lot easier than under a typical scenario, where a deep level of expertise is needed to manage under those three skill sets."

Lockhart of Hastings Prince Edward Public Health says adding new compute/storage/networking is also much faster when compared to traditional infrastructure. “An upgrade to our server cluster was 20 minutes with no down time, versus hours of downtime with an interruption in service using the traditional method,” he says.

“Instead of concentrating on infrastructure, you can expand the amount of time and resources you spend on workloads, which adds value to your business. When you don’t have to worry about infrastructure, you can spend more time on things that add value to your clients,” Lockhart adds.

6. Faster disaster recovery

Key elements of hyperconvergence products are their backup, recovery, data protection, and data deduplication capabilities, plus analytics to examine it all. Disaster recovery components are managed from a single dashboard, and HCI monitors not only the on-premises storage but also cloud storage resources. With deduplication, compression rates as high as 55:1, and backups can be done in minutes.

USS-POSCO Industries is an HP Enterprise shop and uses HPE’s SimpliVity HCI software, which includes dedupe, backup, and recovery. Goodall says he gets about 12-15:1 compression on mixed workloads, and that has eliminated the need for third-party backup software.

More importantly, recovery timeframes have dropped. “The best recent example is a Windows update messed up a manufacturing line, and the error wasn’t realized for a few weeks. In about 30 minutes, I rolled through four weeks of backups, updated the system, rebooted and tested a 350GB system. Restoring just one backup would have been a multi-hour process,” Goodall says.

7. Hyperconvergence analytics

HCI products come with a considerable amount of analytics software to monitor workloads and find resource constraints. The monitoring software is consolidated into a single dashboard view of system performance, including negatively impacted performance.

Hastings recently had a problem with a Windows 7 migration, but the HCI model made it easy to get performance info. “It showed that workloads, depending on time of day, were running out of memory, and there was excessive CPU queuing and paging,” Lockhart says. “We had the entire [issue] written up in an hour. It was easy to determine where problems lie. It can take a lot longer without that single-pane-of-glass view.”

8. Less time managing network, storage resources

Goodall says he used to spend up to 50% of his time dealing with storage issues and backup matrixes. Now he spends maybe 20% of his time dealing with it and most of his

time tackling and addressing legacy systems. And his apps are better performing under HCI. “We’ve had no issues with our SQL databases; if anything, we’ve seen huge performance gain due to the move to full SSDs [instead of hard disks] and the data dedupe, reducing reads and writes in the environment.”

Across every industry and geography, 93% of enterprises are committed to a hybrid multicloud strategy.

The top reasons and motivators for a hybrid cloud and edge come from a recent Harris Poll:

- A desire for the latest innovation from cloud providers
- Regulations to comply with
- The need to support remote work
- Or simply a need to keep up with competition

Whichever the reason, it’s clear that hybrid deployments are here to stay. Organizations are going to run workloads in the cloud, and outside of public cloud datacenters. But they’d still want access to all the benefits the cloud provides. Azure is a resilient, scalable cloud platform for infrastructure workloads and new services to build apps – all topped with a set of management and operations solutions and Microsoft is bringing the value of the cloud to any infrastructure with Azure Arc. Azure Arc is a bridge that extends the Azure platform so organizations can build applications and services with the flexibility to run across datacenters, edge, and multicloud environments on both new and existing hardware, virtualization, and Kubernetes platforms, IoT devices, and integrated systems.

Learn More

8 enterprise cloud strategy trends for 2022

Cloud adopters are shifting tactics to take advantage of fresh concepts and approaches. Here's what you can expect to see this year.

| By John Edwards, CIO

The sky's the limit when it comes to cloud innovation. As cloud technology advances and matures, enterprises are finding fresh ways to leverage what has become a technology linchpin.

Cloud computing will continue to shape the technology landscape for many years, says Bernie Hoecker, partner and enterprise cloud transformation leader with technology research and advisory firm ISG, noting that cloud adoption is rising rapidly. According to ISG research, cloud services spending increased 49% to over $450 billion during the pandemic and is expected to grow another 20% in 2022. "Although cloud computing is still in its infancy, cloud technology and business models are now the foundation for most digital services," Hoecker says.

To help you keep pace with the rapidly growing and evolving cloud market, here's a quick look at eight trends that promise to shape enterprise cloud strategy this year.

1. Cloud migration is going enterprise-wide

Cloud migration will continue evolving from a "lift-and-shift" approach to a carefully considered, methodical, and strategic modernization, predicts Ken Englund, a technology leader at professional services firm EY Americas. "Enterprise-wide cloud migration is the key to moving away from siloed tools and platforms to a more holistic business strategy," he notes. "More companies are seeing the value in scale and agility."

To fully optimize cloud migrations, business and engineering groups must work as a single team, Englund suggests. "When cloud migration is approached as an IT-only effort, the focus is often cost-oriented, which can hinder innovation and opportunities for bringing new capabilities to customers," he explains.

2. Cloud goals are shifting from efficiency to intimacy

Enterprises globally are going "all-in" on the cloud, relying on the technology's promise to increase efficiency and profitability, says Jenny Koehler, US cloud and digital leader at professional services firm PricewaterhouseCoopers. "In 2022, cloud transformations will take on a new role in building connectedness and intimacy between customers, employees, and business leaders," she predicts. "Organizations will look to the cloud to not only solve complex challenges, but also to build trust and identify new opportunities for stakeholders."

The real winners, Koehler notes, will be enterprises that leverage the cloud for intimacy by encouraging high degrees of C-suite engagement, creating implementation roadmaps, and establishing thoughtful plans.

3. The cloud is going vertical

As cloud services proliferate, industry-

specific — even sector-specific — business processes will increasingly be sourced from cloud providers and cloud-native third parties, predicts Mike Bechtel, chief futurist with business advisory firm Deloitte Consulting. “Retail product recommendations, hotel reservations, insurance claims administration … are industry-specific capabilities, and they’re well on their way to being source-able from cloud-native providers via API,” he explains.

The idea behind such “industry clouds” is simple, Bechtel notes, but the business implications are anything but simplistic. “It comes down to market competition and the scarcity of engineering talent,” he says. “Even leading organizations spend enormous amounts of time and resources on maintaining, administrating, and operating legacy industry systems that, on balance, don’t move their needle.”

An industry cloud isn’t a once-and-done initiative. “It essentially establishes a blueprint for how to think about the future, and the possible building blocks that need to come together,” Bechtel says. “Software engineering is the way to achieve that blueprint, one building block or use case at a time, but continuously evolving with new technology choices and improving these capabilities.”

4. Cloud management is becoming more challenging

As cloud adoption continues to scale, the operational complexity of managing and securing cloud resources is following in lockstep. “For many organizations, this complexity is further compounded by the need to scale within existing private cloud environments as well as being on at least one public cloud provider — often across multiple public cloud providers,” says Eric Drobisewski, senior architect at Liberty Mutual Insurance.

A modern platform-centric approach to scaling hybrid and multicloud ecosystems can help IT teams optimize the amount of time spent maintaining and securing cloud resources, freeing up time to focus on business differentiation through new digital capabilities, Drobisewski notes. “Organizations cannot afford to lose the agility they have created through cloud technology, and to sustain it they need to optimize their operating approach across a hybrid multicloud ecosystem.”

5. Cloud-native is taking center stage

Cloud-native has become ubiquitous among organizations that want to boost their agility, responsiveness, and time-to-market, says Manish Mathuria, CTO at digital engineering services firm Infostretch.

Cloud-native applications are designed according to the principles of cloud architecture, making them faster and easier to develop, deploy, update, and maintain. According to research firm IDC, by 2023 more than 500 million digital apps and services

will be developed and deployed using a cloud-native approach.

"That's hardly surprising, given the post-pandemic landscape many businesses now find themselves in," Mathuria says. As well as accelerating cloud migration in terms of productivity and daily operations, He believes that the trend toward cloud-native is also prompting many organizations with only a foot in the cloud-native door to fully embrace the approach and run with it.

Cloud-native is more than a technology, it's a philosophy, Mathuria states. "Cloud-native applications are often referred to as 'citizens of the cloud,' which not only means they reside there, but that they were built there from the ground up using every advantage that cloud architecture has at its disposal," he says. "That's the common mark of success.

6. Demand for portability and connectivity will increase

In 2022 there will be a greater emphasis on portability and connectivity, as well as on reining in the cost of these elastic services and rebalancing cloud versus on-premises workloads. "What makes this trend so important is that there will be a focus on mobile application development from both cloud-native and cloud-enabled in-house DevOps teams," says Neil Warren, managing director and North American head of the cloud center of excellence at technology services and consulting firm Capgemini.

Warren also predicts that applications will grow more platform-agnostic this year, enabling enterprises to concentrate on productivity and security issues. "The other lever to this would be greater utilization of hyperscale virtual desktops in a non-persistent approach," he adds. "These advancements can further reduce an enterprises' infrastructure complexity, reducing the burden on IT and improving data security."

7. AI is converging with data and the cloud

AI technology alone can't do very much to resolve important problems. Data and scalable computing power are also needed to complete the job. That's why leading enterprises are increasingly administering data, AI, and cloud (DAC) as a unified whole, says Anand Rao, global AI lead at PricewaterhouseCoopers.

"We'll see an influx of companies in 2022 take a lifecycle approach to managing these three interconnected operations when developing AI governance," he states. "When data, AI, and the cloud work together smoothly, end-to-end, the result is a supple and powerful system that realizes more value from data and solves more problems faster."

8. Expect a move toward application modernization

Enterprises have migrated an untold number of applications into the cloud over the past

several years. Yet in many cases, applications were hastily migrated without incorporating any of the re-architecting or re-engineering steps necessary to allow the applications to perform optimally within a cloud environment.

“Instead, many applications were ‘lifted and shifted,’ and therefore do not take advantage of key elements of cloud applications, such as rapid scalability, self-service, reserve pooling, or consumption-based costs,” warns John Rostern, senior vice president and global lead of cloud and Infrastructure security services at security consulting firm NCC Group.

To realize the promised benefits of cloud computing, organizations should begin a comprehensive and considered process where the architecture and deployment of applications in the cloud are examined against an accepted reference model, such as the AWS or Azure Well Architected Framework (WAF). “This process will serve as the basis for the evolution of these applications to cloud native,” Rostern explains.

While re-architecting and re-deploying applications and services to leverage the inherent characteristics of cloud computing will address this problem over the long term, such a transformation is typically a multiyear process, Rostern cautions. Meanwhile, there are short-term actions available to mitigate the impact of existing inefficiencies. Organizations, he says, should examine the problem from a business process standpoint to determine how well applications and services will be able to support operations. “This will reduce the friction in the system and thereby reduce the cost impact in the near to middle term until re-architecture and re-deployment can occur.”

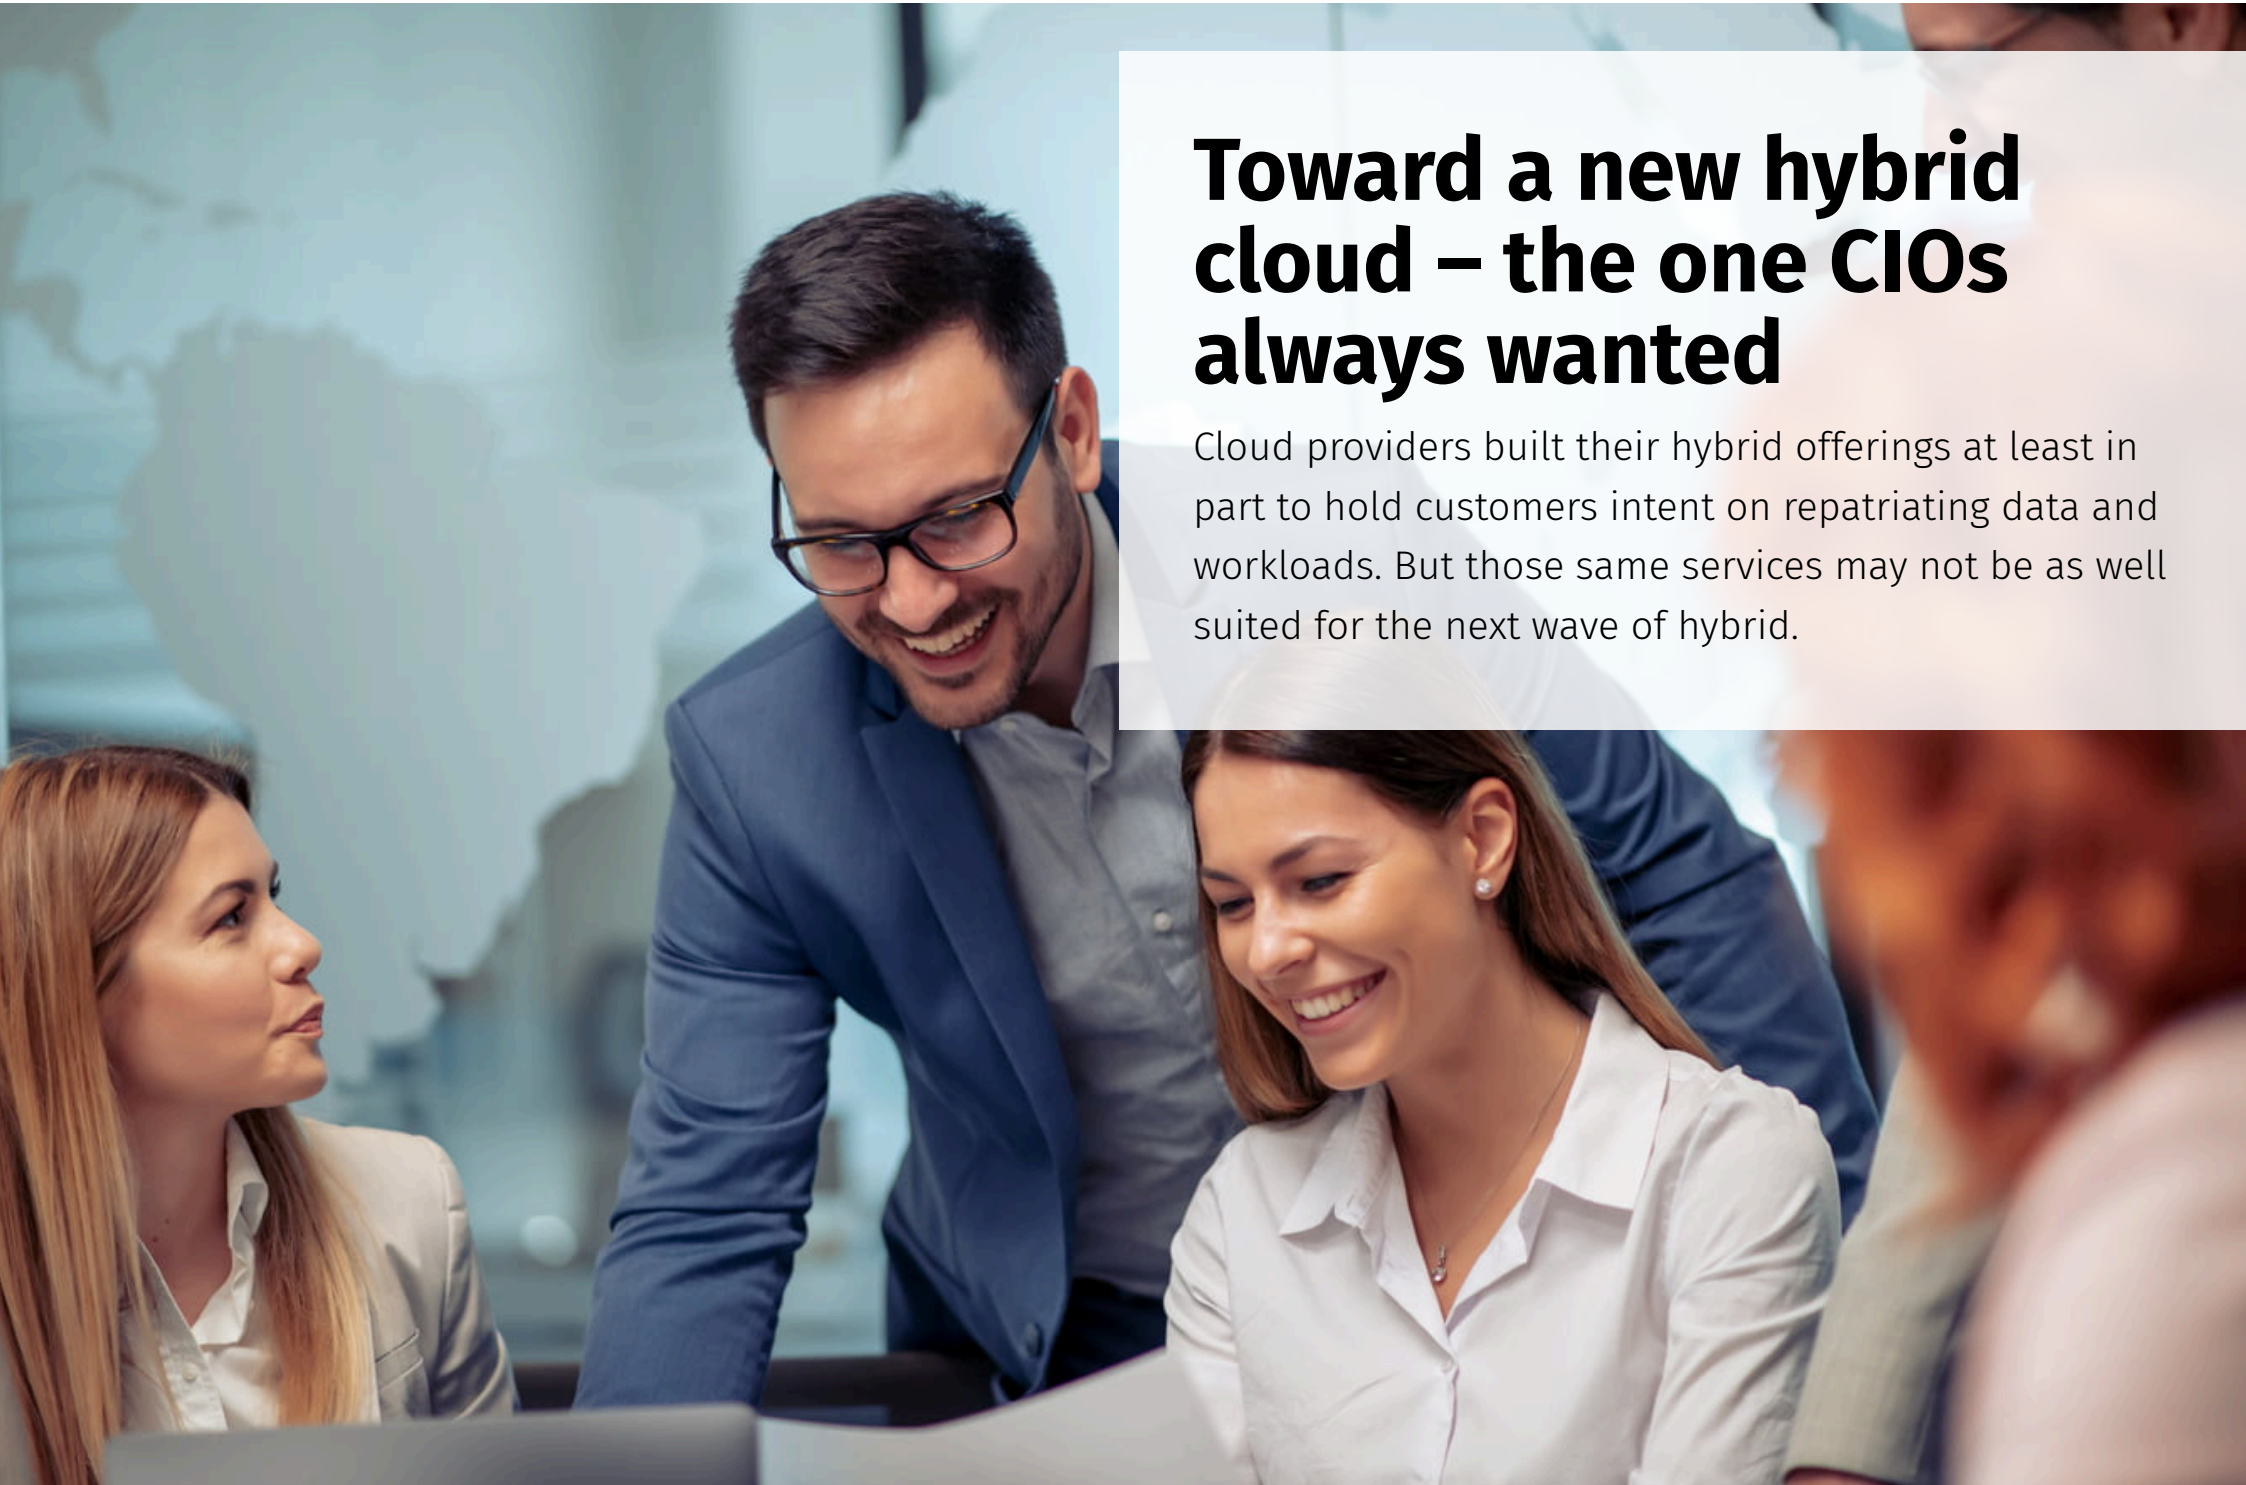

Toward a new hybrid cloud – the one CIOs always wanted

Cloud providers built their hybrid offerings at least in part to hold customers intent on repatriating data and workloads. But those same services may not be as well suited for the next wave of hybrid.

| By Mike Feibus, CIO

One giant leap … is no way to migrate to the cloud. That's what many companies discovered in their headlong rush to embrace the digital transformation.

The hangover from that overreach has been well documented. Like, for example, the inflexible, multiyear contracts that can sap potential savings as compute and storage needs stray from plan. Or divergent governance models, metadata retention policies and access costs that can create unwanted new data silos.

Which is why many of those companies have been looking for ways to move some data and workloads back into owned and co-located datacenters. The big three cloud providers all responded with new releases: AWS Outposts, Google Anthos and Microsoft Azure Arc, which are hybrid cloud offerings that more closely resemble two-way thoroughfares. And ever since, they've been packing more services into those hybrid packages.

Meanwhile, the next wave of modernization is taking hold – and cloud service providers and other XaaS vendors in the space should understand that it looks very different than the last one. By and large, these organizations watched the first wave from the sidelines and are hip to the pitfalls that tripped up early adopters. Some view early hybrid offerings more as escape hatches from the cloud than viable routes into the cloud.

In other words, companies now looking to modernize have no interest in taking any giant leaps. They want bread-and-butter flexibility, efficiency, governance and security. They want to take baby steps, each with a quick path to returns on their limited investments. In short, they want a new hybrid cloud, one that works from the data center up, not from the cloud down.

Doing more with … the same

Many of the CIOs in the second wave of modernization invariably sat out the first wave. But the call to modernize has gotten so deafening that it has become impossible for them to ignore any longer. The volume of data created in 2020, for example, mushroomed 56.6% to 64.2 zettabytes, a much faster clip than the 24.2% growth during 2019, according to Statista. The data management headaches that's creating are exacerbated by the jump in the different types of data. On top of that, the number of users who want access to the data is aggravating contention issues on aging assets and architectures.

As well, these CIOs can see that there's a revolution underway in discovery through data – and all the innovation is happening in the cloud-native space. Even if CIOs haven't been keeping up independently, frustrated data scientists that want to do what their peers are doing have made it abundantly clear to them what insights they are forgoing due to their aging infrastructure.

So decision-makers know they need to migrate to frameworks that can scale with fast-evolving needs. At the same time, though, budget constraints as well as an ever-tightening supply of skilled, experienced talent are also conspiring to shape the scope and pace of modernization plans. In HashiCorp's State of the Cloud Survey, in fact, respondents said that cost and the dearth of in-house skills were two of the top three most common inhibitors to IT cloud programs.

Baby steps

Regardless of whether they want to or need to, IT in this wave intends to take baby steps toward modernization. That way, they can keep a lid on costs and build the organizational pool of skills and experience organically. In the comfort of their own emerging private clouds.

But make no mistake: they are keenly motivated to improve efficiency and ROI. And to do that, they know they need to migrate to a software-defined architecture – as much as possible, on their existing hardware. They want to begin adopting container-based services. And they want a someday path to the public cloud that is push-button simple.

If that sounds like hybrid cloud to you, well, no arguments here. But you'd be better off calling it private cloud. Because for this first round, at least, decision-makers in this wave don't intend to stray from their own data centers. Besides, hybrid cloud to them sounds more like the ongoing repatriation efforts from the first wave than anything they're trying to accomplish.

And the trend is gaining steam. Private-cloud deployments are now responsible for 24.3% of workloads, up markedly from 16.6% in 2020, according to the Global Cloud Survey 2021 by data virtualization provider Denodo. In contrast, workloads in the other categories – hybrid cloud, public cloud and multi-cloud – each declined as a percent of the total. Hybrid cloud – the standard definition – is still home to 35.6% of all workloads, more than any other category, according to the survey.

Now that CIOs have a better understanding of what it is they want from hybrid cloud, traditional offerings delivered from on high will have more limited appeal. Of course, there's plenty of growth ahead for XaaS providers that stick to the current playbook. But for the long term, they might try delivering new hybrid cloud products that actually give CIOs the features and flexibility they need to modernize from the ground up. It will make their customers happy. And eventually, their investors will thank them as well.

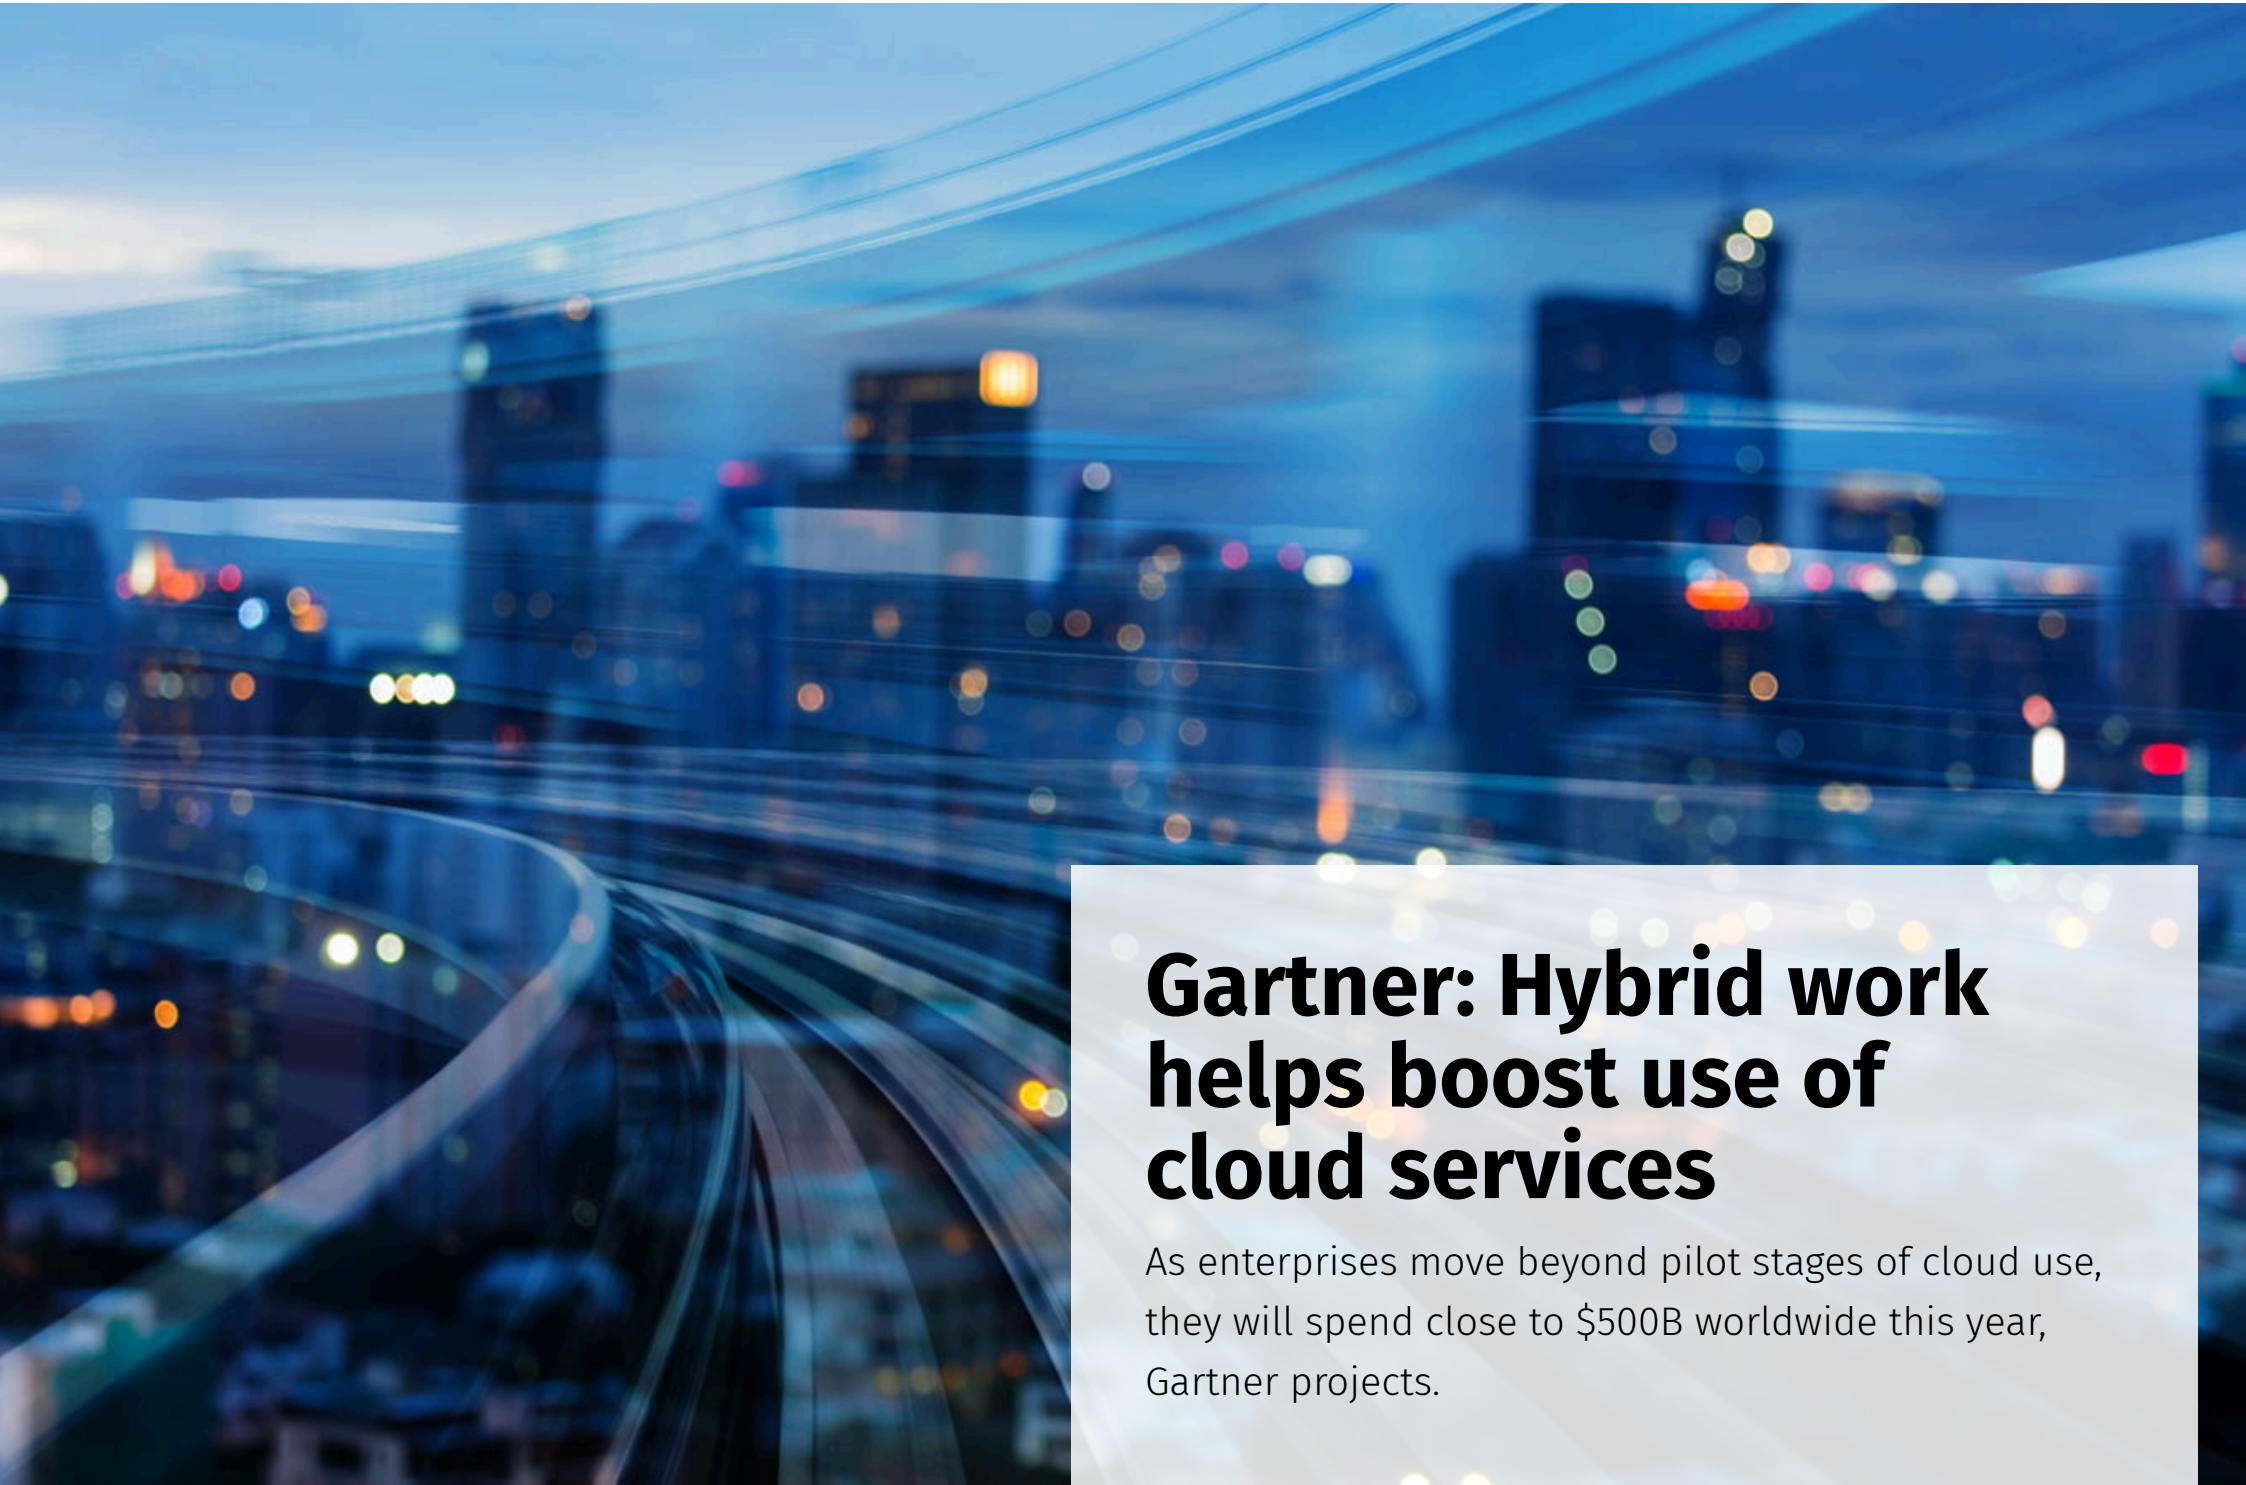

Gartner: Hybrid work helps boost use of cloud services

As enterprises move beyond pilot stages of cloud use, they will spend close to $500B worldwide this year, Gartner projects.

| By Andy Patrizio, Network World

Gartner projects that global spending on public cloud services will come in at $494.6 billion this year due to both growth in cloud-native infrastructure services and the trend towards hybrid work scenarios driven by the pandemic as well.

That's a 20.4% increase over the $410.9 billion in sales in 2021, just shy of the 21.2% growth to $599.8 billion that Gartner projects for 2023.

Because of the maturation of core cloud services, companies are past the tire-kicking stage and shifting their focus to capabilities that can affect digital business and enterprise operations.

"CIOs are beyond the era of irrational exuberance of procuring cloud services and are being thoughtful in their choice of public-cloud providers to drive specific, desired business and technology outcomes in their digital transformation journey," said Sid Nag, research vice president at Gartner in a statement. "Public-cloud services have become so integral that providers are now forced to address social and political challenges, such as sustainability and data sovereignty."

The greatest growth in spending is projected to be for infrastructure as a service (IaaS)—30.6% this year and another 30.5% next year. That is followed by desktop-as-a-service (DaaS), a consequence of the Covid lockdown and more people working from home. It is projected to grow 26.6% this year and another 23.7% next year.

However, DaaS is also by far the smallest category, just $2 billion this year compared to $176.6 billion for software as a service (SaaS) and $91.6 billion for IaaS. And with increasing numbers of companies ordering their staff to return to the office, don't be surprised if the DaaS projections are subject to change.

The third largest category in 2022 is projected to be application infrastructure services, or platform-as-a-service (PaaS), coming in at $109.6 billion. All told, SaaS, IaaS, and PaaS will account for 82% of total cloud sales.

Nag said that the dollar increase is more attributable to the cost of the services than to a massive expansion of purchases. "Cloud-native capabilities such as containerization, database PaaS, and artificial intelligence/machine learning contain richer features than commoditized compute such as IaaS or network-as-a-service," said Nag. "As a result, they are generally more expensive which is fueling spending growth."

Gartner also pointed to new product categories such as hyperscale edge computing and secure access service edge (SASE) as disruptors to adjacent markets and are forming new product categories, creating new revenue streams for cloud providers.

Hybrid cloud hurdles — and how to address them

Mixing public cloud services with on-premises clouds is fast becoming an architecture of choice for many enterprises, which all too often fall into these common traps.

| By Bob Violino, CIO

Many enterprises are moving into a hybrid world — whether it's the emerging workplace model or cloud environment. With the latter, a mix of private and public cloud services offers the flexibility organizations are looking for in a business environment that seems to be constantly shifting.

An August 2021 report by research firm Gartner notes that hybrid, multicloud, and edge environments "are growing and setting the stage for new distributed cloud models." The firm forecasts end-user spending on public cloud services will reach $396 billion in 2021 and grow 22% to reach $482 billion in 2022. By 2026, it predicts, public cloud spending will exceed 45% of all enterprise IT spending, up from less than 17% in 2021.

Meanwhile, adoption of private clouds is also on the rise. Research firm Global Industry Analysts predicts the global market for private cloud services will rise from $4.9 billion in 2020 to $13.2 billion by 2026, expanding at a compound annual growth rate of 18% over the period.

While some companies might opt to use only private or only public clouds, a mix seems to be particularly attractive to many.

"Hybrid cloud is increasingly one of the most predominant architectures we are seeing across enterprises today," says Nicholas Merizzi, a principal at consulting firm Deloitte Consulting. "It provides organizations with flexibility to carve their own transformation path to cloud that aligns with their priorities."

It's clear that organizations see the potential benefits of the hybrid cloud approach. But they'd better be prepared for some hurdles as well. Here are some of the challenges they might face, and how they can address them.

Operational complexities

Hybrid cloud essentially requires maintaining and managing two tightly integrated but separate ecosystems, which can create a complex operating environment, Merizzi says.

"Running parallel environments with multiple different technology platforms introduces ongoing operational complexities across areas such as monitoring, security, and production support," Merizzi says. Ensuring that operational processes and tools are applicable across both public cloud and on-premises private clouds becomes increasingly challenging.
"For instance, there may be a different suite of tools leveraged for cloud security versus that used for [on-premises] systems," Merizzi says.

Similarly, the underlying infrastructure ranging from storage to computing capacity has many differences when going from one environment to another. In addition, he says, upskilling staff to work across these different tools and environments introduces risks and efficiency challenges.

To minimize the complexity hurdle, organizations should look to drive commonality across the technologies for private and public cloud platforms, Merizzi says. This includes having tools that can extend from on-premises environments to the public cloud while maintaining the same operational experience.

Deloitte clients are "seeking software solutions to enable observability and end-to-end application tracing that can function across a hybrid cloud environment," Merizzi says. "Organizations need to drive greater operational visibility and measurement of the internal states of their systems through strong instrumentation."

The City of Wichita, Kan., is implementing software that allows it to better monitor its various IT environments, including cloud services.

"While the complexity of hybrid networks has increased dramatically, the software and applications available to support these have become more sophisticated as well," says Michael Mayta, CIO for the city. "When implemented from a global perspective and [by] applying various technologies such as automation, managing these networks" is possible.

Rising costs

Maintaining both on-premises private cloud and public cloud environments can result in significant cost overruns, if an organization is not disciplined in managing the underlying assets.

"Many organizations develop two sets of books when living in hybrid," including onsite private cloud and off-premises public cloud financials, Merizzi says. "In some cases, the teams even use different tools to assess their financials," he says. "This is leading to CIOs struggling to get an integrated view and projection of their finances."

This exposure is being amplified by the inability to turn off or decommission on-premises assets after having modernized and operationalized systems to the public cloud, Merizzi says. "We see these cost overruns surface at multiple junctions, including network circuits and software licensing," he says.

To minimize this problem, Deloitte recommends establishing an enterprise-wide cloud modernization office that focuses on multiple areas, including financial integration and asset decommissioning. Integrating the financial books for multiple environments provides an enterprise-wide financial picture, Merizzi says. Likewise, establishing a decommissioning process can help free up unused assets to reduce overall spending, he says.

Resource management tools can help keep costs under control. "Without a unified, centralized multicloud management platform, groups did not feel sure of their cost analyses and related decisions," says Arthur Hu, senior vice president and CIO at computer hardware provide Lenovo.

"To remedy this, we provided users with cost-planning tools that offered clear visibility into resource usage, billing, and expenditure forecasts," Hu says. "We supplemented these tools with intelligent resource optimization and rightsizing suggestions to help them achieve optimal cost efficiency."

Lack of coherent strategy

The cloud model is enticing, and many organizations might be tempted to deploy public and private clouds without really giving enough thought to what they hope to achieve and exactly how they plan on meeting their goals. This can lead to confusion, frustration, and cost overruns.

"A good cloud strategy that clearly outlines the value/business case of an expensive, complicated hybrid cloud solution is definitely a start," says Jeremy Roberts, analyst and research director for cloud and core infrastructure at IT research and advisory firm Info-Tech Research Group.

"Review your expected gain," Roberts says. "Why hybrid? Look at your environment. What workloads do you expect will benefit from the hybrid architecture? How? Don't just do it for the sake of doing it."

One organization Roberts dealt with seriously considered a hybrid deployment. "When we spoke more about their goals and drivers, we came to the conclusion that they were not especially interested in managing a private cloud and would instead like to focus on a SaaS [software as a service]-first, multicloud deployment, leveraging multiple public clouds instead of public/private," he says.

The reason was that the overhead and effort required to manage both public and private clouds and interoperability between them was not worth it, given the organization's overall goals, Roberts says.

"The case might be different for organizations with regulatory requirements that force them to keep some data/workloads on-premises, or those who need the performance they can only get from local services, but want to take advantage of the inherent characteristics shared by the public and private clouds," Roberts says.

Change in mindset

As with most major technology shifts, one of the biggest hurdles to overcome when moving to a hybrid cloud environment is the need for a mindset shift, Hu says.

"We had to ensure the businesses bought into the value of moving to a hybrid cloud environment," Hu says. Even the IT department had to be sold on the idea, he says. "We had to start with ourselves; IT wasn't totally convinced at the outset that this was the right direction," he says. "That quickly changed."

Lenovo started by piloting some "lighthouse" applications via a hybrid infrastructure and was immediately able to demonstrate clear value, Hu says. "Engineering productivity improved by 65%, thanks to the agile

practices enabled by the cloud platform and tools," he says. "Once we saw the significant value, we were able to successfully 'evangelize' to the rest of the company, and we continue to showcase impressive results."

Lack of standards or established practices

Although cloud services have been in use for years, a wholesale move to a hybrid cloud infrastructure is still uncharted waters for lots of organizations. "Because this was new territory for Lenovo, we knew we needed to establish clear standards for moving to the cloud and provide best practices," Hu says.

To address this, the company published clear guidance on how to choose the right cloud hosting environment, for private cloud, public cloud, or hybrid. "We made this selection part of every application's enterprise architecture review," Hu says.

The company also established a "5-R Migration Approach" that helped application owners determine their best path based on the application's current architecture state. This includes rehost (directly redeploy to cloud); refactor (modify architecture to adapt to the cloud); revise (modify or extend existing code beyond architectural modification); rebuild (discard the existing technology and develop from scratch); or replace (sunset the existing application and adopt an alternative solution).

Lenovo also developed application-level best practices for groups to reference, Hu says.

Market confusion

"Vendor selection and management is one of the thorniest topics for enterprises moving into hybrid cloud today," says Chris Kanaracus, research director for dedicated and hybrid cloud infrastructure at research firm IDC.

"That's because there's a lot of choice, which is a good thing, but the signal-to-noise ratio is getting higher all the time," Kanaracus says. "You have the well-established leading hyperscalers with AWS, Microsoft, and Google. But up-and-comers such as Oracle are making aggressive moves around cost and capabilities to grow share."

Other choices include Cisco, VMware, HPE, and IBM, "who have made pivots toward hybrid cloud," Kanaracus says. "They see an opportunity to provide the on-premises and edge components of hybrid cloud working in conjunction with the hyperscalers. But the hyperscalers want that business too, evidenced by AWS Outposts, Azure Stack, and Google Anthos."

Also, telecommunications companies and colocation providers are getting more into the mix, Kanaracus says. "It's just a lot to conceive of and manage from a CIO's perspective," he says. "Cost monitoring and management is already critical in hybrid, and this will only be a bigger deal going forward. Lots of vendors, large and small, are eager to be in this game. [There] are lots of choices for CIOs and thus the potential for confusion and

missteps. It's important to take this area of investment very seriously."

Additional resources

Additional Resources

1. Arc Solution Sheet

2.SNP's Hybrid Cloud Solutions - 4 Week Implementation

3.Accelerate Innovation Across Hybrid & Multicloud Environments with Azure Arc

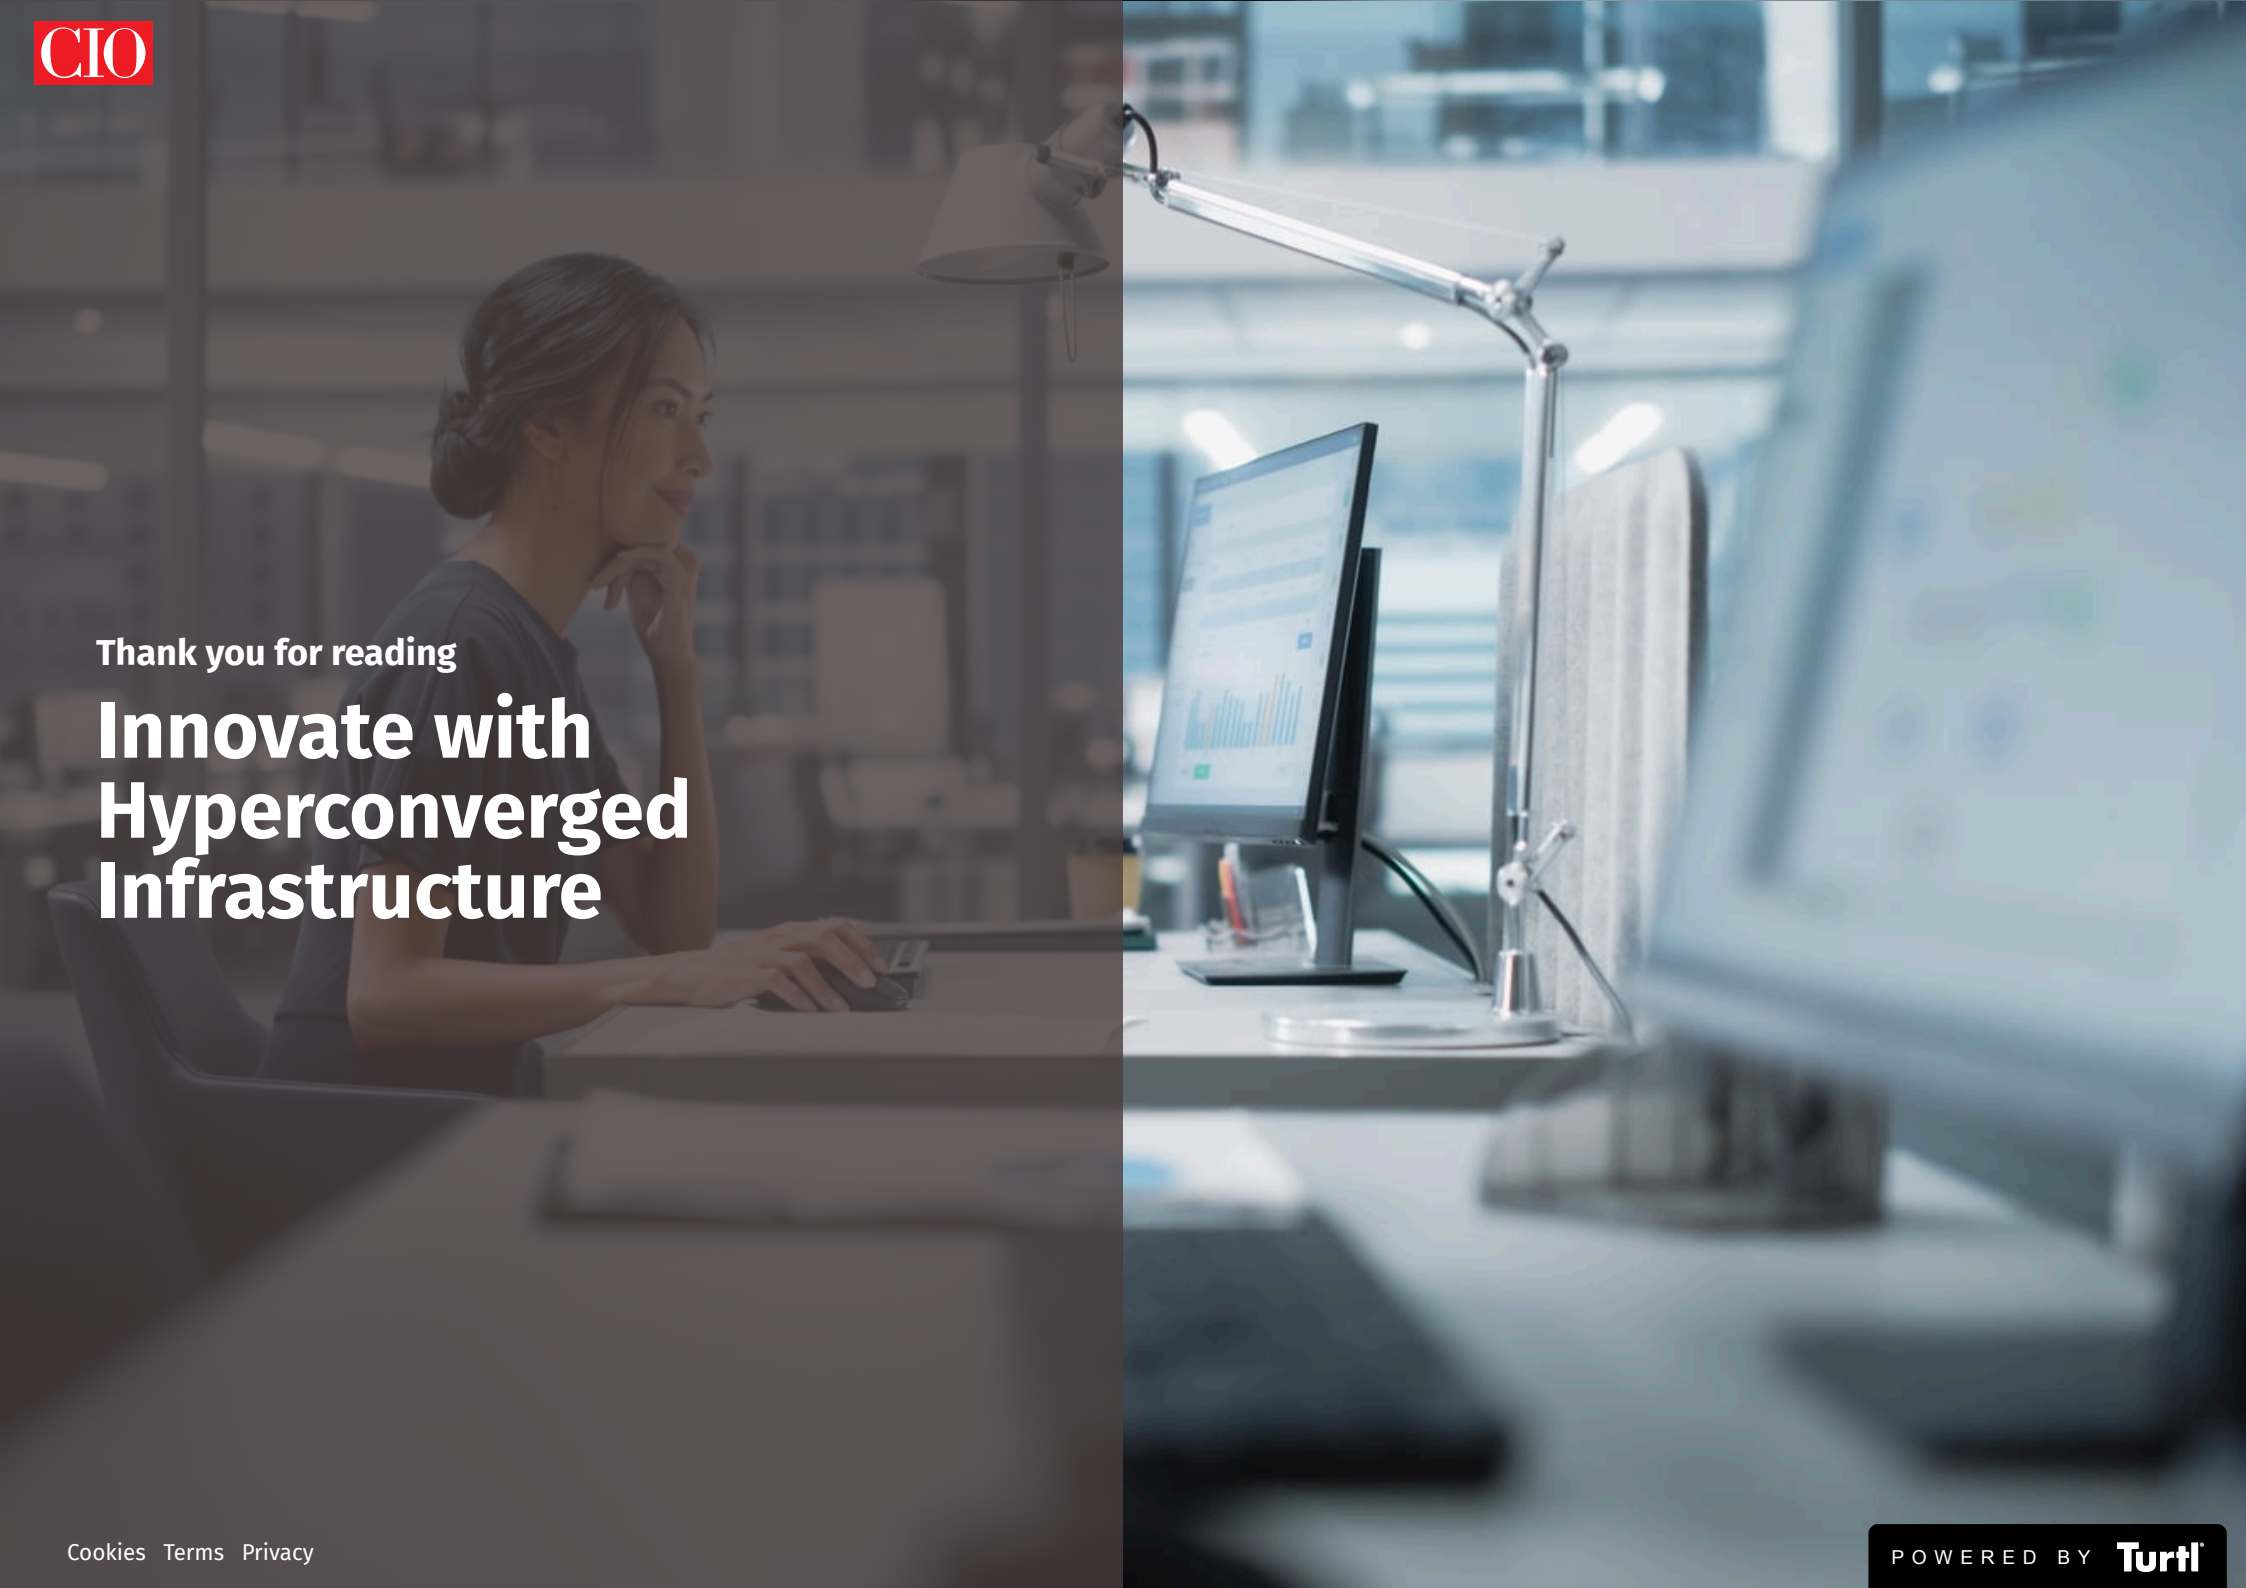

Thank you for reading

Innovate with Hyperconverged Infrastructure

Cookies Terms Privacy

POWERED BY Turtl