

Top Five Use Cases for Hyperconverged Infrastructure in 2022



Pathfinder

December 2021

Commissioned by

vmware[®]

451 Research

S&P Global
Market Intelligence

©Copyright 2021 S&P Global Market Intelligence. All Rights Reserved.

About this paper

A Pathfinder paper navigates decision-makers through the issues surrounding a specific technology or business case, explores the business value of adoption, and recommends the range of considerations and concrete next steps in the decision-making process.

About the Author



Christian Perry

Senior Research Analyst, Infrastructure

Christian Perry is a Senior Research Analyst covering IT infrastructure at 451 Research, a part of S&P Global Market Intelligence. In this role, Christian covers emerging and disruptive infrastructure technologies, including hyperconverged infrastructure and composable infrastructure. He also manages the Voice of the Enterprise products – built on 451 Research’s proprietary global network of senior IT decision-makers – covering hyperconverged infrastructure. With more than 20 years of experience tracking and analyzing the IT infrastructure market, Christian brings broad knowledge to research around software-defined and cloud-native technologies that increasingly shape today’s market.

Prior to joining 451 Research, Christian was Practice Manager and Principal Analyst at Technology Business Research, where he directed the firm’s infrastructure research. Christian also held roles at Comdex, Treehouse Software, and Sandhills Publishing, and nearly 1,000 of his articles have appeared in technology publications, including Processor, Smart Computing, PC Today, Baseline and others. Christian holds a master’s degree in journalism with high honors from The University of Memphis.

Executive Summary

For many organizations, infrastructure modernization is a grueling, years-long process of trial and error to develop the type of agile environment required in this data-centric age of IT and business operations. Our research shows that a steadily rising number of IT teams are turning to hyperconverged infrastructure (HCI) to short-circuit this process, reducing time and frustration while accelerating the shift to cloud-focused, highly efficient IT operations across all environments. More often than not, enterprises are now modernizing 'in place' to support applications and processes rather than migrating everything to cloud, since hybrid environments ultimately provide a far more dynamic and flexible IT experience than either private or public clouds alone.

In 2022, we expect that the majority of HCI deployments across all major industries will be driven by five use cases: mission-critical applications, modern applications, edge deployments, hybrid cloud and disaster recovery. Our research shows that HCI is up for the task, as steady innovation across all of these areas now positions HCI as highly effective infrastructure for transitioning to modern IT operations. This encompasses support for vital pieces of the modernization puzzle, including containers, disaggregated resources, public cloud integration and others. On the surface, HCI may appear to be simply a one-stop shop for the latest and greatest infrastructure technologies, but the reality is far deeper, as current HCI users are leaning on it heavily as a holistic vehicle for IT transformation, both inside and outside of their datacenters.

Use Case: Mission-Critical Applications

HCI has received substantial attention in recent years for innovations that support emerging use cases, but from its inception, HCI has successfully supported mission-critical applications across all major industries. Nonetheless, HCI vendors continue to enhance their platforms with functions that further improve the performance, scalability and management of these applications. Whether running Microsoft SQL Server, Oracle databases, Microsoft Exchange, SAP HANA or other critical applications, HCI is now proven as a go-to deployment choice for organizations looking to simplify their environments and optimize their workload experience.

According to 451 Research's Voice of the Enterprise (VotE): Hyperconverged Infrastructure, Strategy & Workloads 2021 study, organizations are deploying a wide range of mission-critical workloads on HCI, including databases, data analytics, security, CRM, disaster recovery and ERP (see Figure 1). Even in limited, single-workload HCI deployments, we find that organizations often migrate more critical workloads to HCI over time to take advantage of its high performance (particularly when outfitted with all-flash storage, including NVMe), policy-based management, RAID tiering and right-sized scaling for varied application requirements.

Figure 1: Mission-Critical Workloads on HCI



Q. Which of the following workloads/applications are currently deployed on your organization's HCI? Please select all that apply.

Base: HCI is in use or POC (n=223)

Source: 451 Research's Voice of the Enterprise: Hyperconverged Infrastructure, Strategy and Workloads 2021

In fact, significant HCI innovation in recent years has extended scaling flexibility, which is critical in fast-growing businesses that demand agility. For example, a new line of business might require higher levels of compute (e.g., for analytics) than storage, or a new strategic initiative could carry higher storage than compute requirements. Some HCI platforms now allow IT teams to right-size resources as they scale, rather than adding both compute and storage, and potentially stranding resources. Called disaggregation, this capability essentially decouples compute and storage so that, for example, a cluster with excess compute can mount excess storage from another cluster.

Other recent innovations further bolster HCI's case for hosting mission-critical applications, including the ability to apply a compression-only value to certain applications, such as Microsoft SQL Server or Oracle. In these and similar applications, deduplication can slow backup and restoration by requiring a full, uncompressed backup to be written over the network, but with compression-only backups, IT teams can increase the performance of mission-critical applications.

On the security side, our research has revealed a steady migration of security workloads to HCI platforms. The consistent, highly predictable operations enabled by HCI are crucial for security applications, which can perform at less-than-optimal levels when deployed on more complex, stand-alone infrastructure – especially without the full complement of (often expensive) IT personnel required to keep the many disparate infrastructure pieces running without interruption. Essential to these consistent operations is HCI's support for automation and policy-based encryption, which simplify time-consuming tasks.

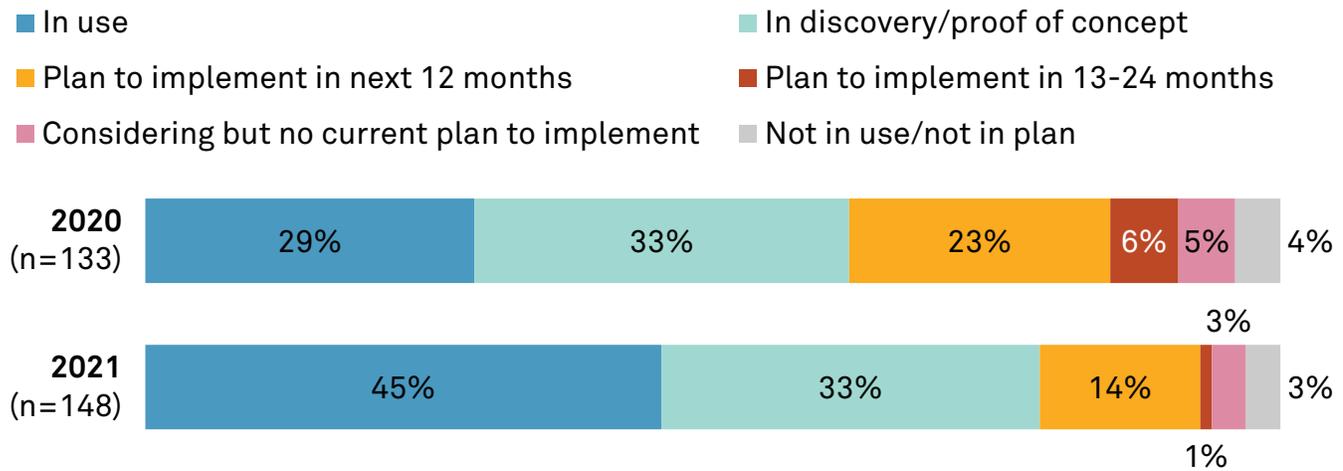
Use Case: Modern Applications

Alongside the ongoing support for mission-critical workloads, HCI has steadily evolved to support modern applications that are inherently integrated across private, hybrid and public clouds. Often termed ‘cloud native,’ these applications are considered essential to digital transformation and any modernization initiatives, whether designed to accelerate IT or business – or both. Nearly any application is considered a modern or cloud-native application if it has attributes generally associated with the concept, including the ability to be containerized, automated, divided into microservices or managed through DevOps processes. Modern applications are typically built from the ground up to support these capabilities, but legacy applications can also be rearchitected to bring them in line with modern functionality.

However, common challenges arise on the road to application modernization. Chief among these is the problem of effectively provisioning, managing and monitoring resources in a highly dynamic environment that inevitably contains a broad mix of VMs and containers spread across a wide assortment of deployment locations. Enterprises that gradually modernize their IT environments often find themselves burdened with multiple infrastructure platforms to host their applications. These increasingly complex ecosystems demand multiple skill sets that can be difficult to find or expensive to acquire, and even heavily staffed IT teams with deep expertise can encounter problems such as siloed operations.

More organizations are turning to HCI to consolidate application deployment on a platform engineered for the nuances of modern, cloud-native requirements. In 451 Research’s *VotE: Hyperconverged Infrastructure, Strategy & Workloads 2021* study, 45% of organizations using HCI said they deploy Kubernetes on those platforms, up from 29% in 2020 (see Figure 2). Further, another 33% of organizations have Kubernetes on HCI in the discovery or proof-of-concept stage.

Figure 2: Adoption of Kubernetes on HCI



Q. What is your organization's current implementation status for Kubernetes on HCI?

Base: HCI is in use/POC, abbreviated fielding

Source: 451 Research’s *Voice of the Enterprise: Hyperconverged Infrastructure, Strategy & Workloads 2021*

The rapid rise of containers and Kubernetes deployments on HCI continues to help organizations with the broader modernization of their infrastructure. For example, customers we speak with are now using HCI – along with Kubernetes – to accelerate software updates, enable microservices, ease deployment of stateless and stateful applications, shift on-premises applications to public cloud environments and more. On HCI, this is all accomplished under the umbrella of unified management for both VMs and containers.

As development teams work to build container-based applications and modernize legacy applications, platform integration and automation is critical to avoid complexity and prolonged development cycles. By integrating enterprise-grade Kubernetes platforms with hypervisors, modern HCI platforms can streamline the development process with built-in libraries, registries and APIs required for developers to optimize their tasks and stay on track with deployment schedules.

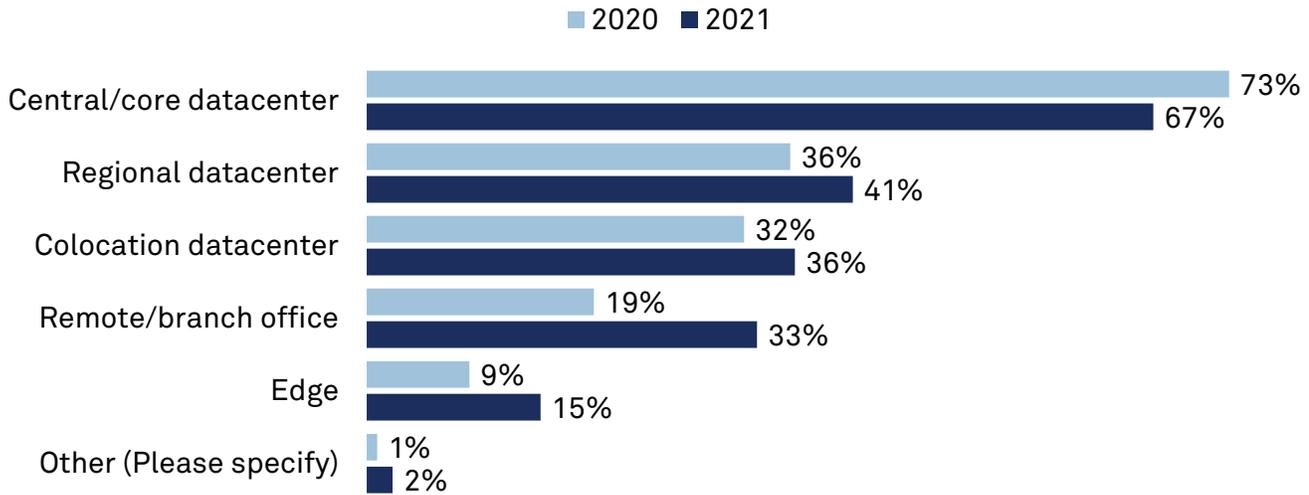
Use Case: Edge/ROBO

Edge IT – and its many variations, such as near-edge and far-edge – is now established as a central piece of business strategy as organizations of all sizes seek to capitalize on opportunities beyond the walls of core datacenters and major public cloud regions. In turn, there is a growing selection of edge-centric approaches and infrastructure designs, each with a unique set of capabilities designed to supply the IT resources needed for remote locations, such as capturing, analyzing and storing data.

In a perfect world, a core IT infrastructure environment deployed and perfected over many years within datacenter walls would be extended seamlessly to ROBO (remote office/branch office) environments or even more remote edge locations. But the reality is that deploying traditional IT infrastructure, such as stand-alone servers and storage, in these locations is often fraught with challenges, including complex deployment and management, lack of floor space and scaling limitations.

HCI, on the other hand, has historically been considered a strong fit for edge environments due to intrinsic design qualities that mitigate or eliminate those challenges. Early HCI edge deployments were typically limited to specific use cases such as video surveillance or VDI (virtual desktop infrastructure), but in recent years we have seen rapid expansion of HCI at the edge and ROBO for practically all use cases in those locations. HCI is now entrenched in core, regional and colocation datacenters, but our VotE: Hyperconverged Infrastructure, Strategy & Workloads 2021 study identified a significant increase in edge and ROBO deployments in 2021 compared with 2020 (see Figure 3).

Figure 3: HCI Deployment Locations: 2021 vs. 2020



Q. In which of the following physical locations has your organization deployed HCI? Please select all that apply.

Base: HCI is in-use/POC (n=224)

Source: 451 Research's Voice of the Enterprise: Hyperconverged Infrastructure, Strategy & Workloads 2021

Of considerable note here is that the ROBO space saw a rapid rise in HCI deployments in 2021, with 33% of respondents deploying HCI in those locations – up from 19% in 2020. The COVID-19 pandemic left a large number of IT teams scrambling to better outfit ROBO locations with more capable infrastructure that would not require specialized skills, in part because quarantine practices limited regular access to remote locations for core IT teams. Our *VotE: Hyperconverged Infrastructure, Technology & Platform Innovation 2020* study found that 55% of organizations with HCI already in place or in the POC stage were spending more on HCI as a result of the pandemic, compared with just 35% that were spending more on stand-alone servers and 37% that were spending more on stand-alone storage hardware.

Part of this growth is associated with HCI's 'right size' deployment options for almost any environment, regardless of space limitations, as well as its strength in remote management – a critical trait for ROBO and edge locations that rarely have full-time on-premises IT staff. HCI typically delivers a consistent, predictable experience, which eases the extension of IT resources to locations outside of core datacenters. After all, the last thing IT managers want is a remote location that requires a different set of operations and management processes.

This consistent operating experience is also a primary reason why customers are also increasingly deploying HCI for security purposes. At the edge, attackers constantly look to exploit weaknesses inherent in infrastructure that lacks enterprise-grade security. One workaround for insufficient security in these locations is to air-gap infrastructure, which removes network interfaces and prevents connections to outside networks. While this approach may be effective for preventing network intrusions, it effectively limits the value of remote deployments, which are most valuable when data generated in those locations can be accessed and analyzed for business advantages. With HCI, the enterprise operating experience extends seamlessly to ROBO and edge locations, ensuring a generally consistent, secure experience.

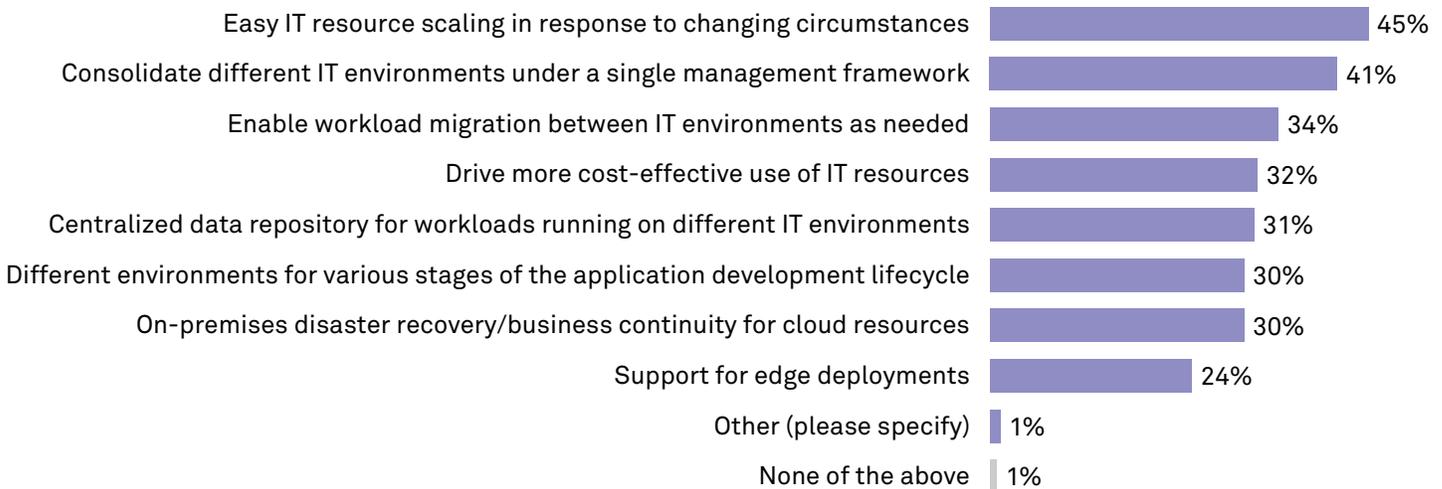
Use Case: Hybrid Cloud

Hybrid cloud is arguably the most challenging element of today’s IT strategies, due primarily to difficulty in managing all of the various components in a typical deployment. In many organizations, these environments are built incrementally over many years, which can easily lead to siloed technologies that require specialized skills and experience. This occurs for a multitude of reasons, beginning with infrastructure (and associated management platforms) that was designed for on-premises environments.

Another driver behind this splintered infrastructure reality is the lack of consistent internal buy-in for hybrid cloud due to different requirements across teams. For example, DevOps might need regular access to public cloud to spin up testing environments, while finance teams only need sporadic access for backups. The result is often an assortment of specialized infrastructure and tools that are complex and time-consuming to manage, and not easily assembled into a cohesive, unified data ecosystem.

Amid the ongoing struggle to build hybrid clouds, HCI has emerged as a logical choice for these deployments. Working with HCI’s inherently simplified, straightforward design, leading HCI vendors have successfully created powerful – yet easily managed – hybrid cloud platforms that customers can quickly deploy without the need for specialized skills. When integrated natively with broader virtualization, container and cloud platforms, HCI not only becomes a one-stop shop for application deployment across the enterprise, but also for flexible resource deployment in any location, whether on- or off-premises. Our research continues to show that HCI is well established as an effective choice for hybrid cloud infrastructure, with customers turning to HCI for a wide assortment of use cases (see Figure 4).

Figure 4: Hybrid Cloud Use Cases on HCI



Q. Which of the following use cases are driving your organization’s implementation (or planned implementation) of hybrid IT on on-premises HCI? Please select all that apply.

Base: Hybrid IT is in place or being implemented (n=242)

Source: 451 Research’s Voice of the Enterprise: Hyperconverged Infrastructure, Strategy & Workloads 2021

In our VotE: Hyperconverged Infrastructure, Strategy & Workloads 2021 study, 45% of surveyed organizations with HCI are using it to ease resource scaling as circumstances change across their environments. Scalability is one of the hallmarks of HCI, giving customers the ability to quickly and nondisruptively add resources to nodes or add nodes to clusters. HCI long ago proved it can meet the challenge of scaling resources on-premises, and now it is solving the same problem across clouds. Our research shows that a rapidly rising number of organizations with HCI are deploying HCI nodes and clusters on public cloud services. Scaling flexibility is the most common driver behind these deployments, supported by tight integration with popular public cloud destinations such as AWS, Azure and Google Cloud Platform.

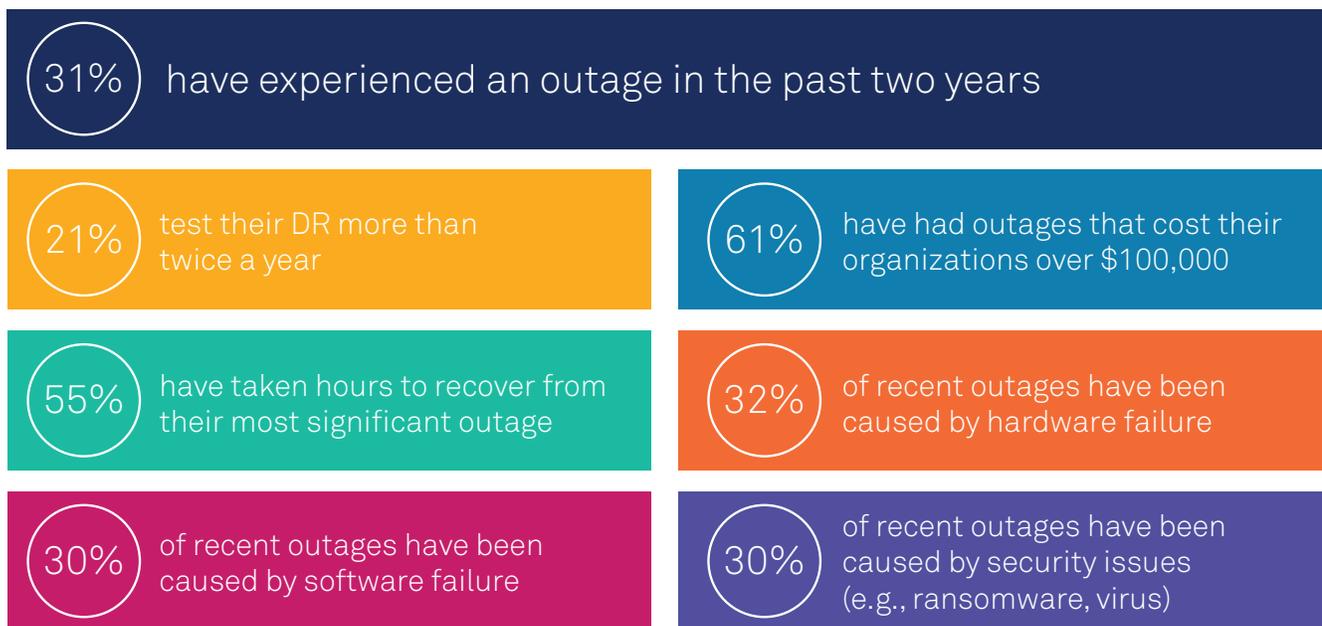
IT teams using HCI are spreading resources across multiple environments – in fact our research shows that 91% of organizations that deploy HCI resources on public cloud are doing so (or plan to) on at least two cloud providers. While this would likely generate excessive management complexity in a traditional IT environment, HCI eases the process. Our study found that 41% of HCI customers are using the platforms to consolidate different IT environments under a single management framework, and 34% are using HCI to boost workload mobility between environments.

The common denominator in all of these deployments is the single management plane across on- and off-premises environments, giving teams the freedom to deploy applications and resources where it makes the most sense to do so – without the headache of switching platforms. For IT teams, this approach removes the burden of managing multiple panes of glass and boosts their understanding and oversight of resource usage across clouds. The results are compelling. According to our study, 97% of HCI customers say that HCI is meeting their expectations for hybrid use cases, with 41% of those respondents indicating that HCI is exceeding their expectations. Further, 97% of these customers agree that HCI eases the process of deploying hybrid IT, with 52% indicating that they strongly agree with that statement.

Use Case: Disaster Recovery

Today's laser-sharp focus on data is delivering a wealth of strategic advantages to businesses, which now have a far deeper understanding of their customers and the markets in which they compete. However, this focus means that most data is now stored, whether for current projects, as part of initiatives for potential future use or for simple archival purposes. This drastically raises the possibility of data loss – through hardware failure, malware, natural disasters or various other threats. According to our *VotE: Storage, Data Management and Data Recovery 2021* study, nearly a third (31%) of surveyed organizations have experienced an outage in the past two years (see Figure 5).

Figure 5: Disaster Recovery Incidents



Source: 451 Research, *Voice of the Enterprise: Storage, Data Management and Disaster Recovery 2021*

A deeper examination of the current state of disaster recovery reveals an alarming picture. Our study found that 61% of surveyed organizations had outages that cost them over \$100,000, while 55% required hours to recover from their most recent outage, and only 21% test their DR implementations more than twice per year. Heavier data consumption in typical modern IT environments challenges traditional models for DR, which were not always designed to support hybrid cloud ecosystems and the unique requirements of cloud-native applications. But as with hybrid cloud, HCI continues to fill the gaps with native DR integration designed to flex with modern, cloud-heavy requirements.

Virtualization has eliminated the need to have identical hardware at production and DR sites, but storage continues to pose problems, since capacity does need to be the same at the DR site to ensure effective protection. Performance also remains a concern, particularly if a DR site is required to run mission-critical applications for an extended period during an outage at the production site. Again, intrinsic capabilities of HCI shine in the DR arena – HCI uses inexpensive industry-standard x86 servers and does not require complex external shared storage, so the same infrastructure used to run critical applications at the core can be used at DR sites without breaking budgets. Further, DR integrated into HCI platforms requires no specialized management tools or connections, which benefits DR and other remote sites that have limited or no local IT staff.

Modern HCI platforms cater to increasingly stringent DR requirements that demand fast recovery – for example, with asynchronous VM replication and recovery point objectives (RPOs) that can be as low as five minutes. As with other elements of the HCI experience, DR benefits from seamless, built-in automation and testing that lets IT managers quickly assign IP address changes during failover and consistently test their DR plans without disruption. Our research also shows that HCI users are increasingly taking advantage of stretched clusters, which are single clusters with nodes distributed across different physical locations. This innovation is of significant value in DR implementations, as an enterprise could share one half of a cluster at each of two production sites, with a third site hosting a second cluster – the DR cluster – to deliver resources for recovered VMs or even run applications in the event of an outage at the production sites. These stretched clusters can achieve an RPO of zero due to synchronous replication between the production sites.

Conclusions

Deciding on infrastructure that can successfully meet the highly dynamic requirements of modern hybrid IT environments is a seemingly impossible challenge for some IT decision-makers, including those who have yet to explore HCI. But enterprises that do move forward with HCI often discover this technology is the most sensible option for present and future use cases – including nearly all application requirements, edge deployments, hybrid cloud support and disaster recovery. While technically there is no ‘easy button’ for infrastructure modernization, our research shows that HCI represents an ‘easier button’ – one that streamlines the transition to modern IT operations with integrated support for developer-friendly processes and business-imperative agility.

Thanks to extensive innovation in key modernization areas, we expect that HCI will thrive as an infrastructure and operations base for essential hybrid IT requirements in 2022 – and in particular the five use cases outlined in this paper. Whether HCI is used for one or all of these use cases, the result will likely mirror what our research has consistently shown for years – that HCI accelerates IT transformation without the need for specialized skills or expertise. This shift to true software-defined IT can have a dramatic effect on IT’s ability to serve the business, especially as the infrastructure is relied upon to accommodate a wider set of use cases.



content provided by: **vmware**[®]

VMware hyperconverged infrastructure (HCI), powered by VMware vSAN, is the market leader in high-performance HCI solutions. VMware vSAN is the only storage software fully integrated with VMware vSphere, enabling customers to manage compute and storage with a single, integrated platform. VMware HCI reduces total storage cost of ownership (TCO) compared to traditional three-tier architecture by eliminating infrastructure silos. More than 30,000 customers and over 80% of Global 2000 have adopted VMware vSAN for flexible, resilient, and future-ready infrastructure.

VMware HCI provides the easiest path to a consistent operating model so companies can manage any application at the edge, in the public cloud or core data center with familiar tools. Its capabilities make VMware HCI the ideal platform for managing traditional virtual machines (VMs) and next-generation applications. It ensures consistent application performance, flexible scaling to adapt to changing requirements quickly, and centralized management.

VMware HCI accelerates business continuity and disaster recovery operations. Businesses need the agility to act decisively and respond appropriately to keep running in the face of many potential disruptions. VMware HCI offers capabilities that help IT prepare for, respond to, recover from, and return to normal following unplanned outages.



Contact a Connection Account Manager for more information.
1.800.800.0014 ■ www.connection.com/VMware

©2022 PC Connection, Inc. All rights reserved. Connection® and we solve IT® are trademarks of PC Connection, Inc. All other copyrights and trademarks remain the property of their respective owners. C1942814-1022

Copyright © 2021 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively, S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers, (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global Market Intelligence does not endorse companies, technologies, products, services, or solutions.

S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain non-public information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its websites, www.standardandpoors.com (free of charge) and www.ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.