

Disaster Recovery in the Age of Hybrid Cloud



**Special
Report**

October 2021

Commissioned by

NUTANIX[™]

451 Research

S&P Global
Market Intelligence

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

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.

Introduction

As infrastructure modernization accelerates the shift to hybrid cloud, organizations increasingly seek flexibility with their disaster recovery deployments. Whether data is currently used for a strategic initiative or simply saved for its potential, it is essential to protect that data amid a growing list of threats that include malware, natural disasters and hardware failure. The ability to set protection targets across multiple environments is crucial and now more feasible than ever, due to DR technologies that ease the management and orchestration process across clouds. When deployed in combination with transformative technologies such as hyperconverged infrastructure (HCI), modern DR can deliver assurance to enterprises, regardless of their data requirements and business SLAs.


Despite Innovation, Storage and Workload Challenges Persist

Enterprises that embark on infrastructure modernization initiatives are typically rewarded with significant improvements across IT, ranging from increased workload performance to better staff efficiency. For some, those improvements are tied closely to a substantial rise in automation. In 451 Research's Voice of the Enterprise (VotE): Hyperconverged Infrastructure, Technology and Platform Innovation 2020 study, we found that HCI users are automating a wide range of tasks on those platforms, including storage provisioning, VM sizing and provisioning, application deployment, disaster recovery and problem remediation. This is not merely a coincidence: all of the technology stacks required for automation (consoles, documentation, troubleshooting, skills, etc.) are built into HCI's management layer, making them inherently simpler, whereas in a legacy architecture environment, they must be individually deployed and configured to work together.

Automating previously manual (or semi-manual) tasks is nearly guaranteed to boost overall IT efficiency and is a necessary step for bridging the efficiency gap between on-premises and public cloud environments. However, it is just one step in the larger modernization picture, and not all infrastructure-related difficulties are inherently addressed by automation. This is particularly true for storage, where capacity, retention, protection and recovery all present ongoing challenges for many organizations.

While data has always been valued in any IT environment, today it holds a more prominent position than ever due to analytics and similar technologies that help organizations derive increased value from it. In fact, enterprises now tend to retain data if it holds even a glimpse of potential value, leading to significant storage requirements across all deployment locations throughout an organization's IT ecosystem. Meeting those capacity requirements is difficult and expensive in the face of rapidly expanding data stores fed by endpoints living everywhere from core datacenters to edge locations. Compounding this challenge is the need to ensure the data can be recovered in the event of an outage, whether caused by hardware failure, malware, natural disasters or other issues.

Figure 1: Top Storage Pain Points



Q: What are your organization's top pain points from a storage perspective? Please select top three choices that apply.
 Base: All respondents
 Source: 451 Research's Voice of the Enterprise: Storage, Data Management and Disaster Recovery 2021

Effective business leaders now understand the transformative, strategic value of data within their organizations; in turn, they expect it all to be available at any given time for new initiatives and projects. This places an immense architectural and administrative burden on IT teams because traditional storage implementations can buckle under the weight of these requirements. 451 Research's VotE: Storage, Data Management and Disaster Recovery 2021 study found that the rapid rise of data is contributing to a wide range of storage-centric challenges.

More than a quarter (28%) of respondents identified meeting disaster recovery requirements as a challenge, more than every other challenge aside from data/capacity growth. Even if disaster recovery technologies and plans are in place, there is an intricate labyrinth of associated factors that can impact the effectiveness of the DR initiative. For example, enterprises encounter difficulties with DR if the plan does not cover all the necessary bases (for example, all applications across all sites), if they fail to frequently test the plans due to complexity in most enterprise environments, or if data is not sufficiently protected.

As the specter of GDPR, Sarbanes-Oxley, PCI DSS, HIPAA and other regulatory laws continue to loom over today's enterprises, storage administrators face an increasingly stringent set of data retention requirements. Complicating their efforts is the need to ensure synchronization with compliance processes, including e-discovery. Our study found that 22% of organizations have encountered challenges with compliance/regulatory/governance requirements. Disaster recovery is of particular concern amid compliance requirements because the loss of data can lead to financial and other penalties.


These challenges – and others – are often amplified in the presence of hybrid cloud initiatives, which can lead to inconsistency across clouds and deployment locations when improperly deployed or managed. Simply deploying modernized infrastructure is not always enough; to the contrary, the modernization can introduce new problems altogether. Of the myriad challenges facing organizations in the hybrid cloud age, disaster recovery integration is arguably one of the most troublesome due to stringent requirements, complexity and the ever-looming threat of viruses and ransomware.

Enterprise Outages Run Rampant

In the VotE: Storage study, 30% of organizations said that security issues such as viruses and ransomware played a part in their most recent outages. Further, 73% of organizations said they are increasing their data protection spending as a result of the potential threat posed by ransomware. In some cases, this investment is designed to implement (or better implement) the 3-2-1 rule, which calls for three copies of data (one production and two backup copies) stored on two different media types (e.g., disk, tape, cloud storage), with one copy stored in an off-site location. Ransomware is driving the requirement for organizations to maintain an air-gapped or immutable copy of backups for added protection.

Figure 2: Disaster Recovery Trends

Disaster Recovery Incidents Continue to Plague Organizations



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

The disaster recovery landscape is littered with setbacks and hindrances, as evidenced by the 31% of organizations that have suffered through an outage in the last two years. Although the percentage in this year's study is similar to what we found a year ago, it is notable that there was an increase in the percentage of organizations with recent outages costing over \$100,000: 61% in this study, compared to 49% a year ago. Directly associated with those high outage costs is the length of outage time; 55% of organizations said they were able to recover from their most serious outage in hours, while 27% recovered in days, and 3% had outages that lasted months.

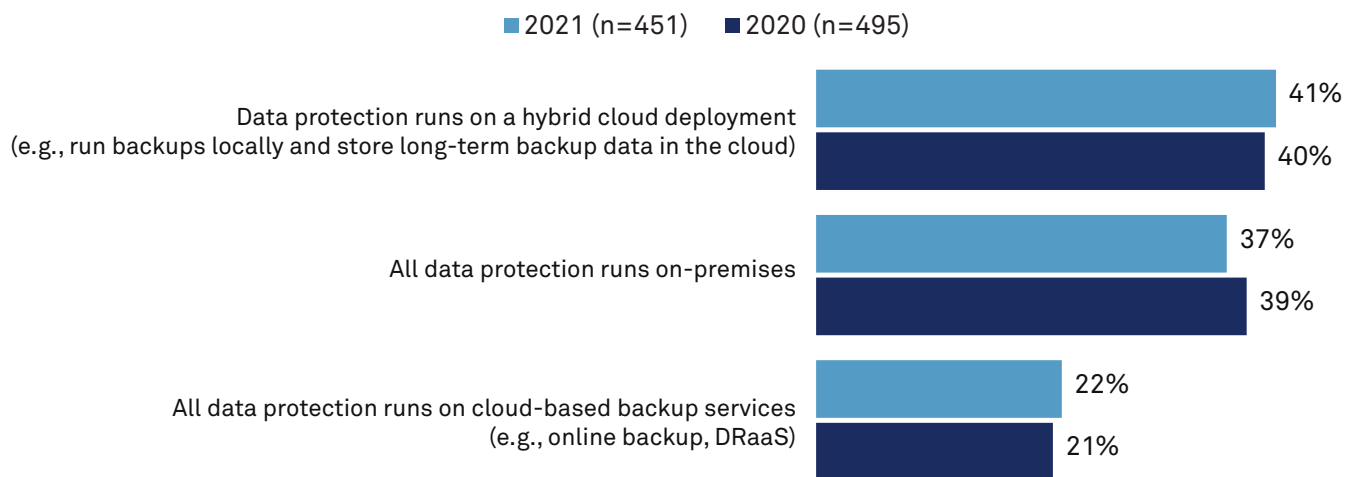
Only 21% of enterprises test their DR more than twice a year, which is problematic because a wide range of contributing factors can lead to outages. Most common, as reported in this study, are hardware and software failures. HCI can help to minimize the outages related to hardware and software failures through its standardized architecture and consistent, predictable operating experience. Further, HCI's inherent automation capabilities can help to minimize human-related errors.

Outages can impact a wide array of areas within a business. For example, our study found that organizations that experience outages suffer from lost worker productivity (51% of respondents), lost data (37%) and lost revenue from missed business opportunities (26%). Respondents also cited damaged reputation (22%), lost customer loyalty (17%) and penalties related to compliance (14%).

Modern HCI to the DR Rescue

Enterprises increasingly view hybrid cloud as an approach for solving outage-related challenges. Our *VotE: Storage* study found that many organizations (63%) are using cloud services for data protection, with 22% of those exclusively using cloud services such as online backup and disaster recovery as a service (DRaaS), in which a third-party provider delivers off-site DR capabilities (see Figure 3). However, 37% of organizations continue to run data protection exclusively on-premises, with some more likely than others to do so. For example, 47% of organizations within the government sector have all data protection on-premises.

Figure 3: Hybrid Data Protection



Q: Which of the following best describes your organization's current use of data protection (e.g., backup, disaster recovery)?

Base: All respondents

Source: 451 Research's Voice of the Enterprise: Storage, Data Management & Disaster Recovery 2021

Meanwhile, most HCI adopters reported success with deploying hybrid clouds on HCI platforms. Our *VotE: Hyperconverged Infrastructure, Strategy and Workloads 2021* study found that 41% of HCI adopters believe that HCI is exceeding expectations for hybrid IT (another 56% said it is meeting expectations), while 97% of HCI adopters agreed that HCI eases the process of deploying hybrid IT in their organization.

Hybrid cloud completely changes the playing field for disaster recovery because it not only helps enterprises work around the cost and time resources of traditional DR setups, but it also provides deeper resiliency. Because most enterprises already are using infrastructure both on- and off-premises, the flexibility to retain data in any location – or preferably multiple locations for maximum protection – not only makes sense but is now a critical requirement.

Although the number of enterprises moving data protection exclusively to the cloud is rising slightly, on-premises environments remain crucial to the DR strategies of most organizations. For example, many firms have significant investments in their IT environment across both infrastructure and skills. Compliance also plays a part in on-premises DR strategies because some regulations drive enterprises to keep data in-house either due to control capabilities or requirements to keep DR sites in specific regions. Another significant driver behind the use of on-premises DR is control over application and data recovery.

Meanwhile, cloud-based data protection can provide levels of security, availability and immutability that some organizations cannot duplicate on-premises. Part of this is due to the significant staff resources required to manage complex storage implementations in a typical datacenter, as well as staff resources needed to secure the infrastructure around the clock. In a typical cloud-based DR scenario, redundancy is built in through the distribution of data across datacenters. Cloud implementations also benefit from efficient patching of the underlying infrastructure, which is often a tremendous burden for on-premises IT staff, particularly when working to ensure the security and availability of critical data.

The importance of leveraging cloud capabilities such as infrastructure elasticity cannot be overstated. Innovative cloud DR offerings that enable use of 'pilot light' recovery sites are sized for only critical infrastructure services and workloads that require the highest level of protection and fastest recovery time objective (RTO). Elastic DR keeps cloud infrastructure for DR to a minimum until it is required, reducing costs. For example, an enterprise does not need an entire primary-site 10-node cluster replicated to a 10-node cluster in a public cloud; instead, it can simply deploy the minimum-required cluster size (e.g., three or four nodes) and quickly scale it to the size when required.

When enterprises are not using the cloud cluster(s) for scale-out DR, they are finding other use cases for that infrastructure, including dev/test, data analytics and capacity bursting. If their recovery point object (RPO) and RTO requirements are less demanding, they can even hibernate the DR cluster to long-term storage on the public cloud. These elastic configurations allow IT teams to use tiered DR more easily; for example, an enterprise can replicate Tier 0 (or 'hot') data on-premises and replicate Tier 1 or 2 data in the scale-out cloud environment as needed, on demand.

When using failback, in which production is restored to the on-premises environment (compared with failover, which switches production to another site), the public cloud cluster can be scaled back to the minimum size after failback occurs. Because all targets are using a common management plane on HCI across the hybrid cloud, there is no need to rebuild or change workload images, a process that can drain staff resources within a traditional DR implementation. Further, common software also tends to lead to efficient data transfer, in turn allowing lower egress costs when an enterprise enables a failback to its on-premises environment. Effective DR platforms integrate one-click failover and failback, along with isolated DR testing that ensures recovery will be ready when required.

Increasingly, enterprises seek to minimize their RPO and RTO times, where the former measures the interval between backups and the latter measures the time between an outage and data recovery. IT environments seeking simplicity often deploy asynchronous DR, which has a time lag between the initial write to the source system and the copying of that data to the target. Although some organizations might find acceptable levels of data loss tied to the frequency of replications, others might deem that delay – which might be an hour or more – unacceptable. However, some HCI platforms use lightweight snapshots to deliver more acceptable times (for example, an RTO of minutes and an RPO of 20 seconds or less).

Effective elastic DR platforms also provide licensing flexibility, whereby enterprises can decommission HCI clusters on-premises or anywhere else and reapply the same platform software licenses to DR nodes operating in their public cloud environment. If enterprises require capacity beyond the licenses they own, certain platforms allow purchase of consumption at a metered rate, which can provide more granular management of costs. Hybrid cloud platforms that allow enterprises to use existing virtual private clouds (VPCs) rather than require separate VPC configurations offer similar levels of flexibility. These capabilities lean on the integration of software-defined networking capabilities, which can simplify the connection of VPCs in DR deployments.

As previously noted, outages continue to plague enterprises, in turn forcing increased emphasis on disaster recovery. That emphasis has become more urgent in recent years with the increase in ransomware, but HCI platforms continue to evolve to help enterprises prevent ransomware attacks and optimize recovery in the event of an attack. For example, some modern platforms integrate a range of technologies designed to counter ransomware and its effects, including network segmentation (both physical and virtual), separate cluster credentials, immutable snapshots, network isolated restores and testing.

As our research shows, most enterprises perceive a hybrid approach that uses multiple environments as ideal. While hybrid DR can technically be accomplished by simply setting targets on- and off-premises, innovative implementations provide more power and flexibility for the overall DR process. Further, modern services allow enterprises to orchestrate their DR across multiple sites, relieving staff from the difficult, resource-intensive management process inherent in legacy hybrid DR approaches.

Finally, an effective hybrid DR implementation also tracks closely with the evolution of enterprise workloads and their unique requirements. For example, pandemic preparedness continues to drive an upsurge in virtual desktop infrastructure (VDI) deployments to accommodate the needs of remote employees. Modern DR technologies use dynamic data mobility to maximize use in VDI deployments and provide geographical distancing capabilities without data residency challenges.

Conclusions

Because data is now a highly valued component of every enterprise estate as workload proliferation continues, efforts to protect that component have risen in kind. Alongside this trend is the rise of hybrid cloud, which now presents compelling opportunities to protect data more easily, effectively and efficiently. And amid these developments is yet another innovative technology that is rising in enterprise deployments – hyperconverged infrastructure. The combination of HCI and hybrid cloud is quickly becoming the path of least resistance for enterprises seeking to optimize critical processes such as disaster recovery.

“The way we’ve selected and looked at all the [HCI] solutions, we wanted something that’s going to work with the cloud, either in part now or a long-stated road map of integrated cloud. So the type of things that we’re looking to leverage is backup and DR to the cloud, being able to restore and go live in the cloud if we have an emergency at a local site, like a hurricane or a winter storm or whatever.”

IT/engineering manager/staff

10,000-49,999 employees, \$10bn+, Energy

“[Having HCI compatible with public cloud] gives us the capability to scale easily. It also gives us disaster recovery and business continuity capabilities. And that’s what we’re actively looking for. That’s why a hybrid approach is something that is critical for us.”

IT/engineering manager/staff

5,000-9,999 employees, \$5bn-\$9.99bn, Manufacturing

As indicated in these quotes from IT professionals, hybrid cloud is now essential for business operations, particularly when used for DR. When they combine it with capabilities to easily set targets across destinations within the hybrid cloud ecosystem and tier data according to IT service priorities, enterprises can transform their ability to protect data and focus on other strategic initiatives.



Nutanix solutions unify operations across your cloud environment, bringing multicloud operability to your enterprise workloads—and making hybrid cloud architectures a reality.

A single management plane to manage infrastructure, apps and data in both your private cloud and public cloud environments dramatically reduces the operational complexity of migrating, extending or bursting your applications and data between clouds, and enables you to leverage your public cloud investments faster!

Visit www.nutanix.com/clusters for a free trial, or contact us directly at clusters@nutanix.com to get started today.

CONTACTS

The Americas

+1 877 863 1306

market.intelligence@spglobal.com

Europe, Middle East & Africa

+44 20 7176 1234

market.intelligence@spglobal.com

Asia-Pacific

+852 2533 3565

market.intelligence@spglobal.com

www.spglobal.com/marketintelligence

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.