

Cloud-Native Approach to Data Protection and Disaster Recovery



Pathfinder Report

July 2022

Commissioned by



451 Research

S&P Global
Market Intelligence

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 Authors



Henry Baltazar

Research Director, Storage

Henry Baltazar is a Research Director for the storage practice at 451 Research, a part of S&P Global Market Intelligence. Henry returned to 451 Research after spending nearly three years at Forrester Research as a senior analyst serving Infrastructure & Operations Professionals and advising Forrester clients on datacenter infrastructure technologies. Henry has evaluated and tested storage hardware and software offerings for more than 15 years as an industry analyst and as a journalist.



Jay Lyman

Senior Research Analyst, DevOps, cloud native, open-source software

Jay Lyman is a Senior Research Analyst with the Cloud Native and Applied Infrastructure & DevOps Channels at 451 Research, a part of S&P Global Market Intelligence. He covers infrastructure software, primarily hybrid and multi-cloud environments, management and orchestration, and enterprise use cases that center on the confluence of software development and IT operations known as DevOps. Jay's analysis encompasses evolving IT operations and software release models, as well as the technology used to create, deploy and support infrastructure and applications in today's enterprise and service-provider markets. This includes running the semi-annual Voice of the Enterprise: DevOps survey of both IT decision-makers and practitioners. Key areas of research also include cloud native, open source software and enterprise end users.

Executive Summary

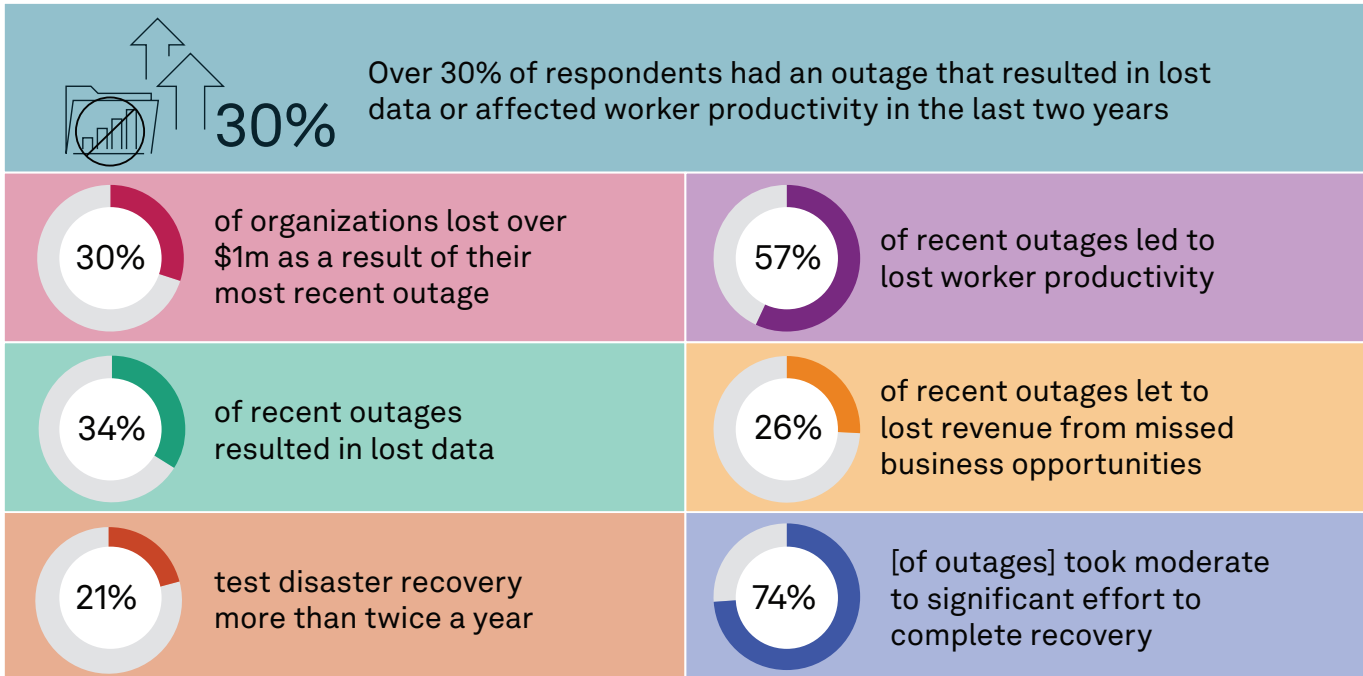
Cloud-native environments and workloads were created to overcome several issues with legacy environments while providing enhanced scalability, provisioning speed and workload portability, among other benefits. With outages continuing to be both common and costly incidents, organizations must focus on how they should implement or revamp their data protection and data management strategies with a cloud-native approach. To meet these needs, next-generation data protection tools must have flexibility to work across on-premises, public cloud and service provider environments while providing consistent protection that not only preserves data but can also rapidly restore entire environments quickly to reduce the negative consequences of an outage. To achieve this, modern data protection offerings must not only understand and cover all of the various aspects of Kubernetes, but they must also be built to efficiently take advantage of Kubernetes orchestration. Enterprises must also leverage DevOps and cloud-native approaches for benefits beyond speed and efficiency, such as business objectives and outcomes that include user experience and customer satisfaction.

Key Findings

- Meeting disaster recovery and business continuity requirements continues to be a challenge for enterprise organizations, particularly with the increasing threat of ransomware.
- While business continuity and resilience have become more critical in the market, business objectives and outcomes are a top priority for enterprise DevOps teams that are improving their deployments and broadening benefits beyond speed and efficiency.
- To achieve cloud-native advantages, including IT operations efficiency and security and developer speed and productivity, organizations need deeper and richer integration with cloud-native technology than simple plug-ins, such as container storage interface for containers and Kubernetes.
- The majority of the market favors backup/DR tools that are optimized for container platforms (58%) over sticking with legacy tools (29%).

Resiliency and Data Protection Requirements

Figure 1: The State of Disaster Recovery



Q. When was the last time your organization experienced an outage that resulted in lost data or affected worker productivity?
Base: All respondents (n=372)

Q. Which of the following effects did your organization experience as a result of your previous outages? Please select all that apply.
Base: Organizations that had negative effects from previous outages (n=253)

Q. Please estimate the total cost to your organization of its most recent cloud outage or downtime (from outage to full recovery, including direct costs, opportunity costs, etc.)
Base: Organizations with recent service outages/incidents and have estimated costs (n=192)

Q. How frequently does your organization test your disaster recovery plan? Base: All respondents (n=372)

Q. How much effort is required to resume normal operations after a failure (i.e., a fallback)? Base: All respondents (n=367)

Source: 451 Research's Voice of the Enterprise: Storage, Data Management and Disaster Recovery 2022

Meeting disaster recovery and business continuity requirements has been a top pain point for organizations for many years now, and it will only become more challenging going forward. Ransomware has become a key threat for almost all organizations surveyed in 451 Research's Voice of the Enterprise (VoE): Storage, Data Management and Disaster Recovery 2022 study; 59% of respondents said they are spending more on backup services and storage because of the increasing threat of ransomware. Nearly a third of respondents suffered through an outage in the past two years, and these outages were costly in many cases, with 30% of respondents claiming the outage led to losses over \$1m.

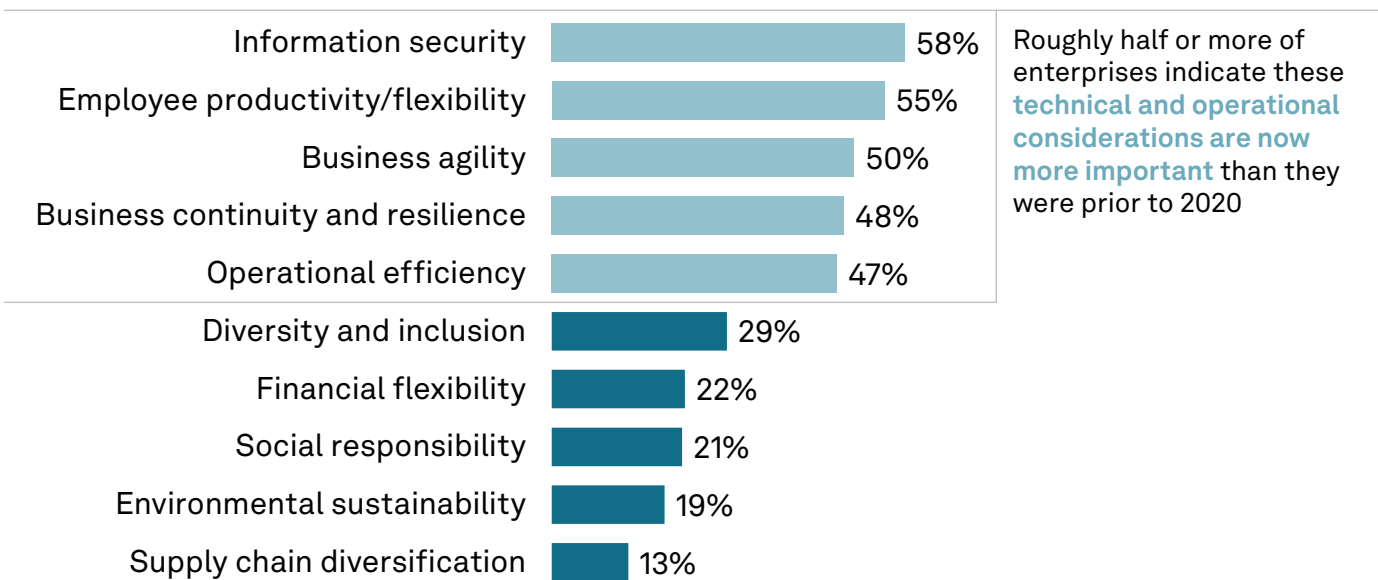
An outage can lead to many negative outcomes for organizations, with lost worker productivity being a key consequence for 57% of respondents in the study, and an additional 34% reported that the outage resulted in data loss. Other notable business consequences were lost revenue from missed business opportunities (26%), damaged reputation (18%) and lost customer loyalty (13%), which show that the impact of an outage goes well beyond an organization's infrastructure. Compliance-sensitive verticals such as healthcare (9%) and financial services (7%) reported that penalties related to compliance were a consequence of their most recent outage, and in some cases, employee termination was a consequence (healthcare, 9%; finance, 7%).

Disaster recovery and business continuity efforts are under a tremendous amount of time pressure, given that every hour of downtime increases the damage caused by an outage. In the study, 74% of respondents reported that it took a moderate to significant amount of effort to restore operations after their most recent outage, which should be alarming for companies. Only 21% of organizations currently test their disaster recovery plans more than twice a year, which can be a major issue in rapidly evolving cloud-native infrastructures, where frequent application and infrastructure changes could impact the recoverability of a workload. One potential solution for increasing the frequency of testing is automation, which can ensure testing is consistent with a reduced hands-on management burden.

Business and Cloud-Native Benefits

Global events such as the COVID-19 pandemic have highlighted the need for digital transformation and business agility, which were prioritized among top business objectives in our VotE: Digital Pulse, Business Reinvention & Transformation 2021 survey (see Figure 2). Enterprises that were unable to adapt to drastic changes in the market across verticals were left behind as those that were able to adjust – both organizationally and with new applications, services and features – found a way to survive and even prosper. The trials and tribulations of 2020 also drove organizations to prioritize business continuity and resilience, which enabled some to successfully navigate the difficult and disruptive period.

Figure 2: Business Objectives Prioritized



Q. Which of the following business objectives – if any – are now more important to your organization’s decision-making process than they were before 2020? Please select all that apply.

Base: All respondents, excluding don’t know (n=420)

Source: 451 Research’s Voice of the Enterprise: Digital Pulse, Business Reinvention & Transformation 2021

The criticality of digital transformation has also helped drive broader and deeper adoption of DevOps, whereby developers and IT operators collaborate for faster releases, more efficient IT management and other benefits. DevOps is typically a starting point or critical component of digital transformation. Our survey research also highlights how more mature DevOps deployments are focused on benefits beyond speed and efficiency, particularly business goals. Business objectives and outcomes are the leading priority (51%) for enterprises as they refine, improve and expand their DevOps deployments, ahead of a DevSecOps approach (46%) and end-to-end automation (also 46%), according to our VotE: DevOps, Workloads and Key Projects 2022 survey.

Further evidence of the growing enterprise focus on business benefits from DevOps is the ranking in our survey of improved user experience as the top desired outcome. Improving the experience for users of applications and services was the top outcome at 58% compared to organizational agility/flexibility (45%) and speeding time to market for new services and improvements (42%). Our survey also highlighted a correlation between digital transformation and prioritization of user experiences, with digital transformation leaders – those executing on their plans – much more likely to identify improved user experience as the top-sought outcome (65%) compared to 36% of digital transformation laggards that have no formal transformation plans.

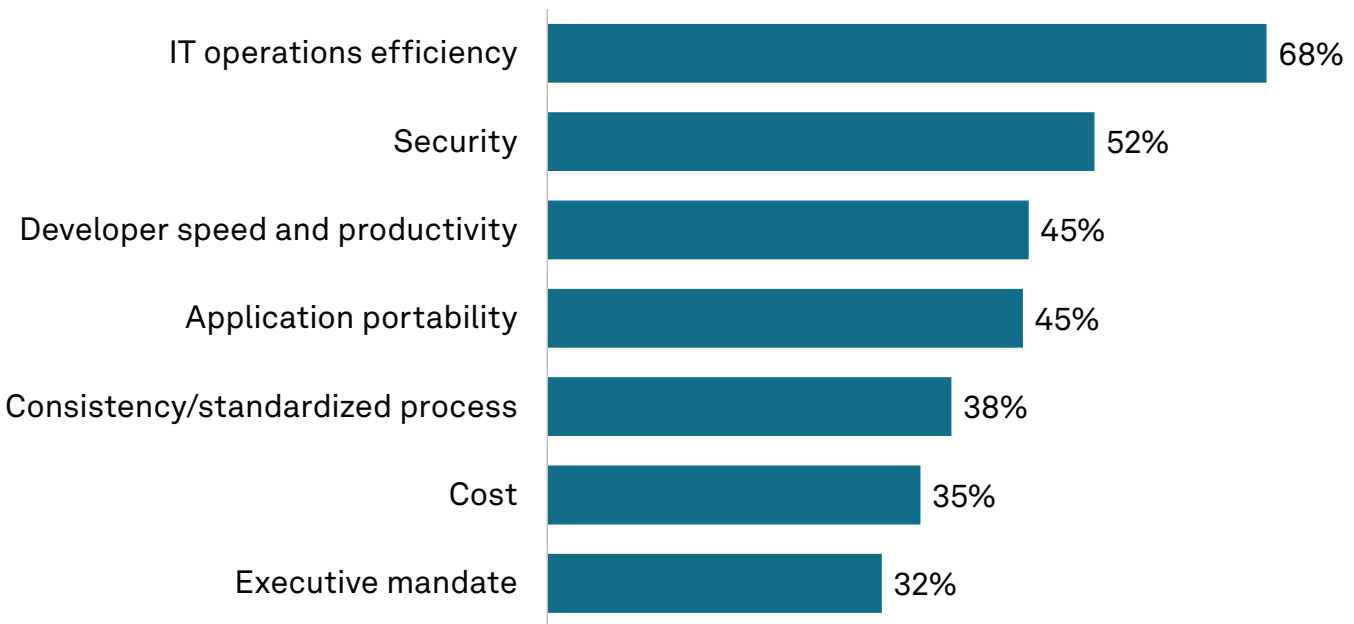
Enterprises are measuring DevOps success using business metrics (54%) as well as technical metrics such as quality (64%) and application performance (57%). Our analysis reveals that teams that have deployed DevOps for a longer period (three to four years) are much more likely to use business metrics to gauge their success (59%) compared to those with shorter DevOps implementations of one to two years (41%). More mature DevOps deployments are typically spread more broadly across the organization and thus are more likely to include business considerations.

Our research highlights that as DevOps is applied more broadly across the organization, it pulls in other critical stakeholders. This includes traditional IT administrators such as storage and networking teams, security teams, data analytics teams, and line-of-business and product managers. It is important that enterprises have the right platforms and tools to serve cross-discipline teams, and that they are aware of persisting silos and the need to promote collaboration among teams not necessarily accustomed to it.

Beyond Plug-ins – a Truly Cloud-Native Approach Is Needed for Full Realization of Benefits

The cloud-native trend – whereby applications are built from the ground up to take advantage of cloud architectures and operational capabilities such as auto-scaling and API provisioning – has grown in the enterprise, driven by benefits including IT operations efficiency, security and developer speed and productivity (see Figure 3). Cloud-native technologies such as containers and Kubernetes are increasingly the tools of choice for developers, IT operators and combined DevOps teams.

Figure 3: Cloud-Native Benefits



Q. What are the primary benefits of cloud-native technology, such as containers, Kubernetes or serverless, for your organization? Please select all that apply.

Base: All respondents (n=548)

Source: 451 Research's Voice of the Enterprise: DevOps, Workloads and Key Projects 2021

Many enterprises have, nonetheless, waded into cloud native gradually by grafting containers onto their existing VM-based infrastructure and application management. This is where plug-ins have enabled less friction with cloud-native deployment. However, there are limitations as organizations increasingly leverage containers without hypervisors or VMs (58% of containers run on cloud services, 46% on a standard OS, 33% on a container-specific OS and 31% on VMs, according to our VotE: DevOps, Workloads and Key Projects 2022 survey). This maturing adoption of cloud native, which includes serverless and service mesh technology as well as containers and Kubernetes, means that organizations require deeper and richer support for cloud-native technology beyond Docker or Kubernetes plug-ins.

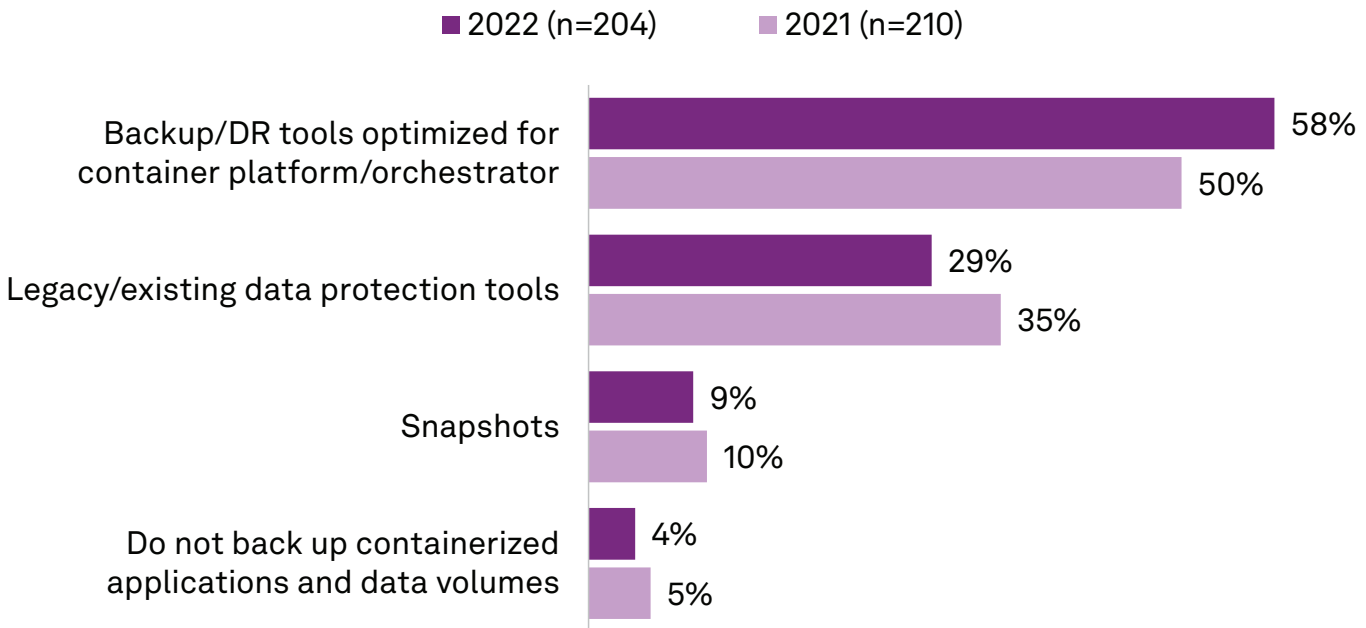
One example of this is the move to more stateful container applications driven by both end users and vendors. On the demand side, enterprises wanted to cast a wider cloud-native net across more of their portfolios beyond new and web applications. On the supply side, the evolution of Kubernetes included support for persistent data volumes and thus stateful applications – mostly operational workloads on databases (62%), data science workloads (62%) and analytic workloads (58%), according to our research. Cloud-native designs represent a dramatic shift from traditional, monolithic applications. Thus, the initial cloud-native applications were limited mostly to web and stateless applications that did not require data persistence. With the evolution of cloud native, we've seen a growing number of stateful applications containerized.

As cloud-native adoption matures, deeper and richer support can enable key cloud-native benefits. Such benefits include IT operations efficiency, such as the ability to slice up VMs to increase the utilization of physical servers and the management of larger-scale infrastructure with smaller teams. Another key cloud-native advantage centers on security. Containers offer more lightweight packaging that reduces the attack surface. At the same time, updates for distributed applications are simplified. Cloud native can also fuel developer speed and productivity by automating and abstracting manual tasks so software engineers can focus on new features, applications and innovation.

Things To Look For

Data protection for cloud-native environments and workloads is challenging for organizations and has forced them to rethink which tools they will rely on to protect cloud-native workloads. In our *VotE: Storage, Data Management and Data Protection 2022* study, the percentage of respondents who said they prefer backup and disaster recovery tools that were designed and optimized for cloud-native environments grew from 50% to 58% in the 2022 study, while those that prefer to rely on legacy tools dropped from 35% to 29% (see Figure 4).

Figure 4: Cloud-Native Data Protection Tools Are Preferred Over Legacy Tools



Q. What is your organization's primary data protection strategy for containerized applications and data volumes?
 Base: Organizations that use containers (n=204)
 Sources: 451 Research's Voice of the Enterprise: Storage, Data Management and Disaster Recovery 2021 and 2022

Cloud-native technologies were created to get beyond several limitations with legacy infrastructure such as limited workload mobility, inadequate scalability and slow provisioning of resources. To match the demanding requirements of cloud-native workloads, next-generation data protection tools and services must have the following capabilities:

- **Broad and consistent execution venue support.** The preferred venue for storing persistent data for containers is a split fairly evenly at 51% favoring public cloud and 49% preferring to store on legacy storage systems or newer software-defined storage options. Cloud-native backup and disaster recovery tools must be able to protect workloads regardless of whether they are using on-premises hardware or running on cloud services.
- **Flexible recovery options.** The ability to move workloads between various execution venues is a key attribute of cloud-native environments, and data protection tools must be able to facilitate recovery operations on-premises or in a public cloud based on the requirements of business stakeholders.

- **Comprehensive data protection.** Data protection and disaster recovery are not complete if the steps required to recreate an environment and complete a failover are not accounted for. These include complete protection for Kubernetes covering application configuration and Kubernetes objects, in addition to the application data.
- **Namespace awareness.** Organizations that would like to back up all the applications running in a namespace simultaneously can run into issues doing this manually, especially if there are too many pods to back up. Tools that understand Kubernetes APIs and can back up at a namespace level can potentially spare organizations from doing time-consuming extract, transform and load operations and can reduce the number of commands per volume that are issued to complete the backups.
- **Application consistency for distributed databases.** Next-generation data protection tools must understand the various snapshot procedures for protecting applications such as Cassandra, Kafka and Elasticsearch since they vary among data services. Using the wrong procedures for freezing workloads, taking snapshots and unfreezing can lead to data corruption and prevent organizations from completing a successful restoration when needed.
- **Enhanced security and resiliency.** Ransomware has had a major impact on data protection in recent years, and 59% of respondents reported that they are increasing their spending on backup and backup storage as a result of the growing ransomware threat. Modern platforms must have immutability and air-gap capabilities to protect backup repositories from attacks. These tools must also work well with role-based access control to restrict system access and simplify security profile management to minimize the potential impact of internal threats.

content provided by:



When protecting Kubernetes applications, it is essential to have a solution that can back up the entirety of a Kubernetes application, which consists of associated data, configuration, and Kubernetes objects. Otherwise, if only the data is protected but not the application configuration, then recovering apps can become a very slow and manual process. Portworx Backup provides application-aware and container-granular backup for full protection of Kubernetes applications. So, entire applications are quickly restored on demand via a few simple clicks. Cross-cloud mobility also gives organizations the flexibility to migrate applications across cloud, hybrid, and on-prem environments based on business needs.

With any data protection solution, security is of utmost importance. Ransomware attacks are growing in frequency, and organizations need to be prepared to protect their applications and data against potential attacks. Portworx Backup guarantees protection against ransomware attacks by leveraging object locked buckets. Easily recover all data from an attack with immutability and delete protection within defined retention periods. Also, granular role-based access controls restrict access to backups for heightened security, helping prevent malicious access.

Disaster recovery and ensuring business continuity is a top concern for most organizations. Service outages can have grave consequences, like losing data, potential revenue, and customer loyalty. To ensure business continuity and zero loss for mission-critical apps, synchronous and asynchronous disaster recovery is needed. With PX-DR, Portworx provides zero RPO and low RTO disaster recovery, instead of merely providing recovery from backups.

Portworx is the #1 Kubernetes data platform that enables organizations to operate, scale, and secure Kubernetes applications anywhere. Portworx Backup provides fast, easy, and secure data protection for all your cloud native applications.

Try Portworx Backup for free with a 30-day trial, or visit www.portworx.com for more information.

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 © 2022 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.