

IDC MarketScape: Worldwide Container Data Management 2023 Vendor Assessment

Johnny Yu

Phil Goodwin

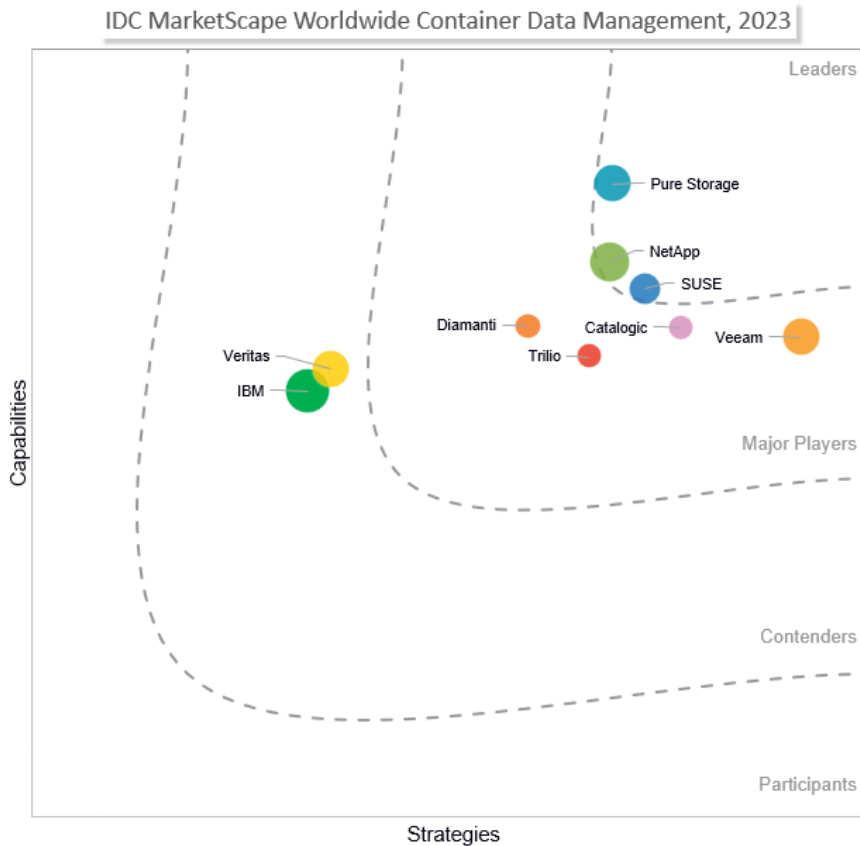
Greg Macatee

THIS IDC MARKETSCOPE EXCERPT FEATURES PURE STORAGE

IDC MARKETSCOPE FIGURE

FIGURE 1

IDC MarketScape Worldwide Container Data Management Vendor Assessment



Source: IDC, 2023

Please see the Appendix for detailed methodology, market definition, and scoring criteria.

IN THIS EXCERPT

The content for this excerpt was taken directly from IDC MarketScape: Worldwide Container Data Management 2023 Vendor Assessment (Doc # US51367723). All or parts of the following sections are included in this excerpt: IDC Opinion, IDC MarketScape Vendor Inclusion Criteria, Essential Guidance, Vendor Summary Profile, Appendix and Learn More. Also included is Figure 1.

IDC OPINION

As organizations increasingly adopt Kubernetes to enable faster application delivery, application portability, and other benefits, they must also adopt tools that can help them manage and maintain the infrastructure that supports containerized applications.

The emergence of cloud-native storage and cloud-native data protection products made it patently clear that storage administrators must treat applications running in container and Kubernetes environments differently from those running in physical and virtual ones. Tasks such as viewing storage resource usage, provisioning storage resources for applications to run and store their data, backing up and restoring these data, and quickly failing over the application into a new environment, should disaster strike, all require new tools specifically tuned for container architecture.

Traditional storage categorizes storage management and data management tasks separately. Although most physical and virtual storage platforms have native data management features such as data migration and backup and recovery, there exists an entirely separate data replication and protection software market designed to work with a range of storage platforms to offer more robust data management capabilities. IDC believes the container market is not yet at a point where this separation is necessary, nor are most customers actively seeking it.

Instead, physical storage and virtual machine (VM) experts are finding themselves responsible for supporting container and Kubernetes environments. They're not container experts, nor should they need to understand the technology to the same level as application developers, but they need tools that let them perform the same support tasks for applications running in containers as the ones they're using for applications running in their organizations' traditional environments. The better the product or platform is at making the tasks between traditional and container environments analogous, the greater the value-add for organizations that are running containerized applications in production.

IDC MARKETSCAPE VENDOR INCLUSION CRITERIA

This IDC MarketScape centers around products that provide data and storage management capabilities for containers (shortened to just "container data management" for brevity). This is not to be confused with products such as container schedulers and orchestrators, which are responsible for the creation and dissolution of the container environments themselves. Container data management products are instead responsible for what are generally categorized as data management capabilities and storage management capabilities.

Participating products in this IDC MarketScape all meet the following criteria:

- Product must be focused on containers and cannot be an existing data management product for traditional environments that requires a plug-in or add-on to extend its capabilities into container environments.

- Product provides either natively or through integrations with its partner ecosystem data management capabilities for container environments. Such capabilities include point-in-time data backup and recovery, backup of cluster components and metadata, data and application migration, and disaster recovery (DR).
- Product provides either natively or through integrations with its partner ecosystem storage management capabilities for container environments. Such capabilities include storage provisioning for containerized applications and persistent data that those applications may depend on, monitoring of storage resource consumption, storage scaling and optimization, and security.

Notably, this IDC MarketScape takes into consideration not only the individual participating product but also what possible container data management solution can be built using the product and its partners. This expands participants' areas of differentiation and allows each to showcase their unique value.

ADVICE FOR TECHNOLOGY BUYERS

IDC does not expect any one product to be the perfect solution for every organization. Container data management vendors differentiate themselves in a variety of ways, and a solution that is perfect for one organization could be suboptimal for another.

Here are some key points to consider when evaluating the right vendor for your organization:

- **Consider your own organization's level of container and Kubernetes expertise.** Some vendors have extremely robust and highly configurable products, but organizations need enough know-how to take advantage of these advanced features. Some vendors provide container and Kubernetes help beyond product support, such as by offering consulting services or fostering a user community via forums and official Discord and Slack channels.
- **Know your starting point, but plan to scale.** Find products that match where you are on your container and Kubernetes journey, and pay attention to how well those products scale. Some products are complete, out-of-the-box solutions, while others are tuned for organizations that have already made investments in container technology.
- **Select products that align with your container adoption plan.** Many organizations are containerizing their legacy applications, but many also have critical applications that must stay running on VMs. Similarly, some organizations use open source Velero to protect their container environments and want to continue to use it as they run more containerized applications.
- **Pay attention to what container management platforms a product supports.** While Red Hat OpenShift is widely adopted, managed Kubernetes services from the three major public clouds and other Kubernetes distributions also see use.

This IDC MarketScape highlights what we believe are the most important criteria when evaluating a container data management product, but it is not intended to be a buyer's guide. Instead, readers should use this document to identify strengths, weaknesses, and key points of differentiation between products in this market to narrow down the list of vendors that work best for them.

VENDOR SUMMARY PROFILES

This section briefly explains IDC's key observations resulting in a vendor's position in the IDC MarketScape. While every vendor is evaluated against each of the criteria outlined in the Appendix, the description here provides a summary of each vendor's strengths and challenges.

Pure Storage

IDC has identified Pure Storage as a Leader in this 2023 IDC MarketScape for container data management.

Portworx by Pure Storage is a Kubernetes data platform consisting of three parts: PX-Enterprise Kubernetes storage platform, backup and data protection tool PX-Backup, and a database-as-a-service platform called Portworx Data Services. The first two provide the core container storage and data management capabilities of the overall Portworx product, and the last component provides a data life-cycle management service for several popular database applications, including MySQL, MongoDB, Postgres, Kafka, and Cassandra.

PX-Backup provides backup and recovery and migration. The backups are application aware, capturing all the underlying resources, data, configurations, and objects needed to recover an application. PX-Backup allows recovery into a different environment, even non-Portworx volumes, so it provides a basic application migration function as well. Multicloud or multicluster application migrations can be handled with PX-Migrate, a tool within PX-Enterprise. PX-DR, an add-on to PX-Enterprise, provides multisite synchronous replication for zero-RPO DR within a metro area and asynchronous replication over WAN.

PX-Enterprise provides storage features such as dynamic provisioning, multicluster management through both a UI and CLI, automatic volume resizing for storage optimization, clusterwide encryption, and RBAC. Storage performance can be tuned through application I/O controls, and the PX-Fast feature provides an optimized I/O path for PVs using underlying NVMe storage.

Portworx's main strength is the sheer range of options the platform offers, both in terms of how it can be deployed and the functionality within the product itself. Multiple deployment options and wide support for container orchestrators and managed Kubernetes services mean Portworx can fit within nearly every organization's infrastructure. The product contains fine-level controls for every practical aspect of container storage and data management, whether it's granularity of backups, DR failover options, or application I/O controls. Portworx Data Services is also a unique, compelling offering, allowing quick deployment of what is the most popular and critical type of application organizations run on Kubernetes today.

Strengths

- Unique database-as-a-service capability that automates database deployment, day 2 operations, and data protection, with support for the most popular database applications
- Storage performance features such as application I/O tuning, data locality, and an optimized I/O path for NVMe (PX-Fast)
- Strong data availability capabilities: synchronous DR, multiple DR failover options, and ransomware protection through object lock
- PX-Security that enables multitenant authorization, authentication, and RBAC for PVs

- Wide range of container orchestrator support, including Mirantis, Nomad by HashiCorp, IBM, and Oracle's Kubernetes engines

Challenges

The sheer amount of control Portworx offers makes it a difficult recommendation for organizations without a high level of container or Kubernetes expertise. While its defaults are functional, organizations without enough understanding of their Kubernetes infrastructure are unlikely to use Portworx to its full potential.

Consider Pure Storage When

Portworx is generally well suited for most organizations looking for a comprehensive container data and storage management solution, but it aligns particularly well with those that are well underway on their container journey. Organizations that are running or looking to run mission-critical applications, databases, CI/CD tools, or AI/ML workloads in containers will want to consider Portworx by Pure Storage.

APPENDIX

Reading an IDC MarketScape Graph

For the purposes of this analysis, IDC divided potential key measures for success into two primary categories: capabilities and strategies.

Positioning on the y-axis reflects the vendor's current capabilities and menu of services and how well aligned the vendor is to customer needs. The capabilities category focuses on the capabilities of the company and product today, here and now. Under this category, IDC analysts will look at how well a vendor is building/delivering capabilities that enable it to execute its chosen strategy in the market.

Positioning on the x-axis, or strategies axis, indicates how well the vendor's future strategy aligns with what customers will require in three to five years. The strategies category focuses on high-level decisions and underlying assumptions about offerings, customer segments, and business and go-to-market plans for the next three to five years.

The size of the individual vendor markers in the IDC MarketScape represents the company size of each individual vendor within the specific market segment being assessed. Due to the nascency of the container data management market and the open source nature of some of the vendors within it, revenue isn't a useful metric for determining a participant's market share. Instead, IDC has opted to use employee count as a neutral representation for participants while presenting relevant information about each of them.

IDC MarketScape Methodology

IDC MarketScape criteria selection, weightings, and vendor scores represent well-researched IDC judgment about the market and specific vendors. IDC analysts tailor the range of standard characteristics by which vendors are measured through structured discussions, surveys, and interviews with market leaders, participants, and end users. Market weightings are based on user interviews, buyer surveys, and the input of IDC experts in each market. IDC analysts base individual vendor scores, and ultimately vendor positions on the IDC MarketScape, on detailed surveys and interviews with the vendors, publicly available information, and end-user experiences in an effort to

provide an accurate and consistent assessment of each vendor’s characteristics, behavior, and capability.

Market Definition

Cloud data management capabilities are defined by products that provide storage and data protection, security, and compliance for cloud-native infrastructure. Our definitions of pertinent market components follow:

- **Cloud-native storage:** The cloud-native storage submarket primarily includes highly portable, container-based, and persistent storage technology that stores the state of a production application natively using only the OS libraries and dependencies required to run the code.
- **Cloud-native protection/security/compliance:** The protection, security, and compliance submarket includes a broad range of providers and suppliers that offer functionality for infrastructure that is considered cloud native — that is, microservices based and automated (scalability, redundancy, provisioning, upgrades, etc.), using policy-driven resource allocation and life-cycle management through DevOps processes. This market continues to expand as these three interrelated topics gain prominence. For this reason, this subsegment draws from the following vendors/suppliers:
 - Global market-leading infrastructure product and/or solution suppliers whose offerings are used in non-native as well as cloud-native deployments, often as part of a generic or infrastructurewide deployment
 - Vendors that focus on protection/security/compliance and that provide some products or major functionality for deploying on cloud-native infrastructure but that are not considered market-leading vendors in the protection/security/compliance sector
 - Specialist vendors that focus almost entirely on such solutions for cloud-native infrastructure at a technology or infrastructure component level (e.g., for protecting Kubernetes and/or microservices environments)

LEARN MORE

Related Research

- *Container Infrastructure Software Market Assessment: Summary of In-Depth Interviews* (IDC #US49940521, December 2022)
- *IDC Market Glance: Cloud-Native Infrastructure Platforms, 3Q22* (IDC #US49642322, September 2022)

Synopsis

This IDC study evaluates and highlights the most important criteria that differentiate prominent vendors providing container data management capabilities. Storage administrators need to be able to perform the same back-end tasks for container environments as for traditional ones. Tools for monitoring and optimizing storage resource consumption, providing storage resources for applications, and performing data protection tasks such as backup and recovery and disaster recovery (DR) are necessary for running containers and Kubernetes at scale.

There’s a wide selection of products that can provide these capabilities, but the best product for each organization is difficult to pinpoint. This study evaluates and highlights the key differentiators of the most prominent container data management vendors, so buyers can pick products that align with their

stage of Kubernetes adoption, level of Kubernetes expertise, and long-term goals of their organization's entire application estate.

"Storage administrators are not always container or Kubernetes experts, yet they are still responsible for supporting containerized applications to the same extent as any other business application," said Johnny Yu, research manager, Storage and Computing Infrastructure Software Platforms at IDC. "Finding the right tools for this is challenging, and this IDC study aims to help organizations narrow down container data management products to the ones that best suit their needs."

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology, IT benchmarking and sourcing, and industry opportunities and trends in over 110 countries. IDC's analysis and insight helps IT professionals, business executives, and the investment community to make fact-based technology decisions and to achieve their key business objectives. Founded in 1964, IDC is a wholly owned subsidiary of International Data Group (IDG, Inc.).

Global Headquarters

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.com
www.idc.com

Copyright and Trademark Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, and web conference and conference event proceedings. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/about/worldwideoffices. Please contact IDC report sales at +1.508.988.7988 or www.idc.com/?modal=contact_repsales for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or web rights.

Copyright 2023 IDC. Reproduction is forbidden unless authorized. All rights reserved.

