

# GIGAOM

MARKET RADAR

## GigaOm Radar for Data Storage for Kubernetes v1.1

ENRICO SIGNORETTI | MAR 5, 2020 - 11:40 AM CST

TOPIC: **DATA STORAGE**



# GigaOm Radar for Data Storage for Kubernetes

## TABLE OF CONTENTS

- 1** Summary
- 2** About the GigaOm Radar
- 3** Market Categories and Deployment Types
- 4** Key Criteria Comparison
- 5** GigaOm Radar
- 6** Vendor Roundup/Overview
- 7** Conclusion
- 8** About Enrico Signoretti
- 9** About GigaOm
- 10** Copyright

## 1. Summary

Kubernetes adoption is quickly accelerating, and enterprises are now in a transition phase. In the last few years, we have seen an increasing interest in container-based application development. As a result, IT organizations started to implement a proof of concepts and laboratories, which moved later to development and test platforms. In this period, the entire industry has matured, both in terms of the core technology (container formats and development tools) and orchestrators, with several companies trying to push their solutions (i.e., Docker Swarm, Mesos DC/OS, Google Kubernetes, and others).

Now that Kubernetes is the clear winner, the number of organizations moving to the production phase is finally growing as well. In most cases, Kubernetes infrastructures are still relatively small, and applications running on them are fairly simple, with limited data storage needs. On the other hand, more and more stateful applications are migrating to these platforms, requiring additional resources and performance. At the same time, enterprises of all sizes are embracing hybrid cloud strategies that are becoming more complex and structured. We are quickly moving from a first adoption phase where data and applications are distributed manually and statically in different on-premises and cloud environments to a new paradigm in which data and application mobility is the key for flexibility and agility.

Now, organizations want the freedom to choose where applications and data should run dynamically, depending on several business, technical, and financial factors. They choose the public cloud for its flexibility and agility, while on-premises infrastructures are still a better option from efficiency, cost, and reliability perspectives. In this scenario, it is highly likely that development and testing are made on the public cloud while production could be on-premises, in the cloud, or both, depending on the business, regulatory, economic, and technical needs of the particular enterprise.

Kubernetes is instrumental in executing this vision, but it needs the right integration with infrastructure layers, such as storage, to make it happen. Persistent and reliable data storage, alongside data management and security, are vital factors to consider when evaluating Kubernetes deployments in enterprise environments today. These factors expand the scope of the orchestrator to a broader set of applications and use cases across different types of on-premises and cloud infrastructures. The goal is to provide a common data storage layer that is abstracted from physical and cloud resources, with a standard set of functionalities, services, protection, security, and management.

## 2. About the GigaOm Radar

### HOW TO READ THIS REPORT

The Market Roundup/Landscape/radar profile is part of a series of documents aimed at giving the reader the tools to better understand a technology, evaluate it, and explore the market to find the best solutions for his organization.

In this context, and in order to get a complete view of the state of the solutions available in the market, the reader should consider the following documents:

**Key Criteria to Evaluate Data Storage for Kubernetes**, is an introduction to the technology, defines the necessary evaluation metrics, the key criteria to evaluate new solutions, and the impact of the latter on the former. It is dedicated to those end users that are approaching a new technology for the first time or want an update on the latest evolution.

**GigaOm Radar for Data Storage for Kubernetes**, offers a brief 360° view of the market including market and technical positioning of most notable vendors, a short introduction of their solutions and differentiations, including a high level graphic comparison of the vendors.

**Vendor Profiles for Key Criteria**, are easy-to-read deep dive documents that cover a single vendor regarding the solutions described in the other reports. They provide more details on the solution, how the vendor approached the key criteria, and the impact that its solution has on the evaluation metrics. This document helps end-users to get a quick but complete evaluation of a single vendor.

### 3. Market Categories and Deployment Types

For a better understanding of the market and vendor positioning (**Table 1**), we assess how well solutions for Kubernetes data storage are positioned to serve specific market segments. These segments are:

- **Small-medium enterprise:** In this category we assess solutions on their ability to meet the needs of organizations ranging from small businesses to medium-sized organizations. Also assessed are departmental use cases in large enterprises, where ease of use and deployment are more important than extensive management functionality, data mobility, and feature set.
- **Large enterprise:** Here offerings are assessed on their ability to support large and business-critical projects, with emphasis on strong data management and mobility features, as well as on robust security and data protection. Scalability is another big differentiator, including the ability to manage multiple clusters from the same user interface.
- **ISP/MSP:** The number of solutions in the Internet service provider/managed service provider (ISP/MSP) category is still pretty limited. Products that do well here usually have the same characteristics as those in the large enterprise category, with an added focus on multi-tenancy and manageability.

In addition, we recognize three deployment models for solutions in this report—enterprise storage systems with a Container Storage Interface (CSI) plug-in, software-defined solutions with optimizations for containers, and cloud-native solutions.

- **Traditional storage arrays with CSI plug-ins:** This type of solution is the most common at the moment and, usually, is the first to be adopted by users. They are easy to deploy and allow for the reuse of storage resources already in place, with a minimal initial investment. On the other hand, most of these solutions are immature, and limited in operational performance, data management, and scalability. Some of the systems in this category are not designed to cope with the number of back-end operations necessary to run a Kubernetes cluster, and therefore risk creating bottlenecks that can impact service level agreements (SLAs) of critical production environments.
- **Software-defined storage with optimizations for containers:** The flexibility of this type of solution limits some of the deficiencies of traditional arrays while keeping the storage infrastructure compatible with traditional workloads and applications. The optimizations allow users to adopt Kubernetes gradually while granting good SLAs both to traditional and next-generation environments.
- **Cloud-native solutions:** These solutions are specifically designed to work with containers and Kubernetes across on-premises and public cloud environments, in both a hybrid and multi-cloud fashion. These solutions are usually architected around a set of core features focused on data management and mobility, with specific data services tailored for containers. They take advantage of storage resources local to each single node in the cluster, cloud storage, or traditional enterprise shared storage such as network-attached storage (NAS) and storage area network (SAN) infrastructures.

Table 1: Vendor Positioning

	MARKET SEGMENT			DEPLOYMENT MODEL		
	SMALL/MED ENTERPRISE	LARGE ENTERPRISE	ISP/MSP	TRADITIONAL STORAGE + CSI	SDS + OPTIMIZATIONS	CLOUD NATIVE
ARRIKTO	+++	+++	+++	-	-	+++
DATACORE	+++	++	+	-	+++	-
DATERA	+++	++	++	-	+++	-
DELL EMC	+	++	-	+++	-	-
DIAMANTI	-	+++	+	++	+++	+
HITACHI VANTARA	+	++	-	+++	-	-
HPE	+++	++	++	+++	+	-
IBM	+	++	-	+++	-	-
INFINIDAT	+++	+	-	+	+++	+
KAMINARIO	+++	++	+	-	+++	-
LINBIT	+++	++	++	-	++	++
MAYA DATA	+++	++	++	-	-	+++
NETAPP	+++	+++	++	+++	-	-
PORTWORX	+	+++	+++	-	-	+++
PURE STORAGE	+++	+++	++	+++	-	-
RANCHER LABS	+++	+++	+++	-	-	+++
RED HAT	+++	++	++	-	+++	-
ROBIN.IO	++	+++	+	-	-	+++
STORAGEOS	+++	++	++	-	-	+++
VMWARE	+++	+++	-	-	+++	-

**Legend:**

+++ Strong focus and perfect fit of the solution

++ The solution is good in this area, but there is still room for improvement

+ The solution has limitations and a narrow set of use cases

- Not applicable or absent.

Source: GigaOm 2020

## 4. Key Criteria Comparison

Following the general indications introduced in the report Key Criteria for Evaluating Data Storage for Kubernetes, **Table 2** quickly summarizes how each vendor included in this research performs in the areas that we consider differentiating and critical in this space. The objective is to give the reader a snapshot of the technical capabilities of different solutions and define the perimeter of the market landscape.

By combining the information provided in **Table 1** and **Table 2** the reader should be able to get a clear idea of the market and the technical solutions available in the field.

**Table 2: Key Criteria and Evaluation Metrics Comparison**

	KEY CRITERIA					EVALUATION METRICS					
	DATA SERVICES	PERFORMANCE	MULTI-TENDACY	SECURITY	MONITORING & ALERTING	ARCHITECTURE	SCALABILITY	FLEXIBILITY	EFFICIENCY	MANAGEABILITY EASE OF USE	PARTNER ECOSYSTEM
ARRIKTO	+++	++	+++	+++	++	+++	++	+++	++	+++	+
DATA CORE	+	+++	+++	++	++	++	++	+++	++	+++	+
DATERA	+	+++	++	+++	++	++	+++	++	+++	++	++
DELL EMC	+	++	++	++	++	+	+	+	++	+	+++
DIAMANTI	+++	+++	++	+++	+++	++	+++	++	+++	+++	+
HITACHI VANTARA	+	++	++	++	++	+	+++	+	++	+	++
HPE	+	++	++	++	+++	+	+	+	++	+	+++
IBM	+	++	+	++	+	+	+	+	++	+	++
INFINIDAT	++	+++	+++	++	++	++	+++	++	+++	++	++
KAMINARIO	+	++	++	++	+	+	+	+	++	++	+
LINBIT	+	++	+++	++	++	++	++	+++	++	+	+
MAYA DATA	+++	++	+	++	+++	++	++	+++	++	+++	++
NETAPP	+	+++	++	+++	+++	++	+++	++	+++	+++	+++
PORTWORX	+++	+++	++	+++	+++	+++	+++	+++	++	+++	+++
PURE STORAGE	+	+++	+++	+++	+	+++	+++	++	+++	++	+++
RANCHER LABS	++	+++	++	++	++	++	+++	++	++	+++	++
RED HAT	+	+	+	++	+++	+	+	+	+	+++	+++
ROBIN.IO	+++	++	+++	+++	+++	+++	++	+++	+++	+++	++
STORAGEOS	+	+++	+++	+++	++	+++	+++	++	++	++	+
VMWARE	+	++	+++	++	+++	++	+++	++	+++	+++	+++

**Legend:**

+++ Strong focus and perfect fit of the solution

++ The solution is good in this area, but there is still room for improvement

+ The solution has limitations and a narrow set of use cases

- Not applicable or absent.

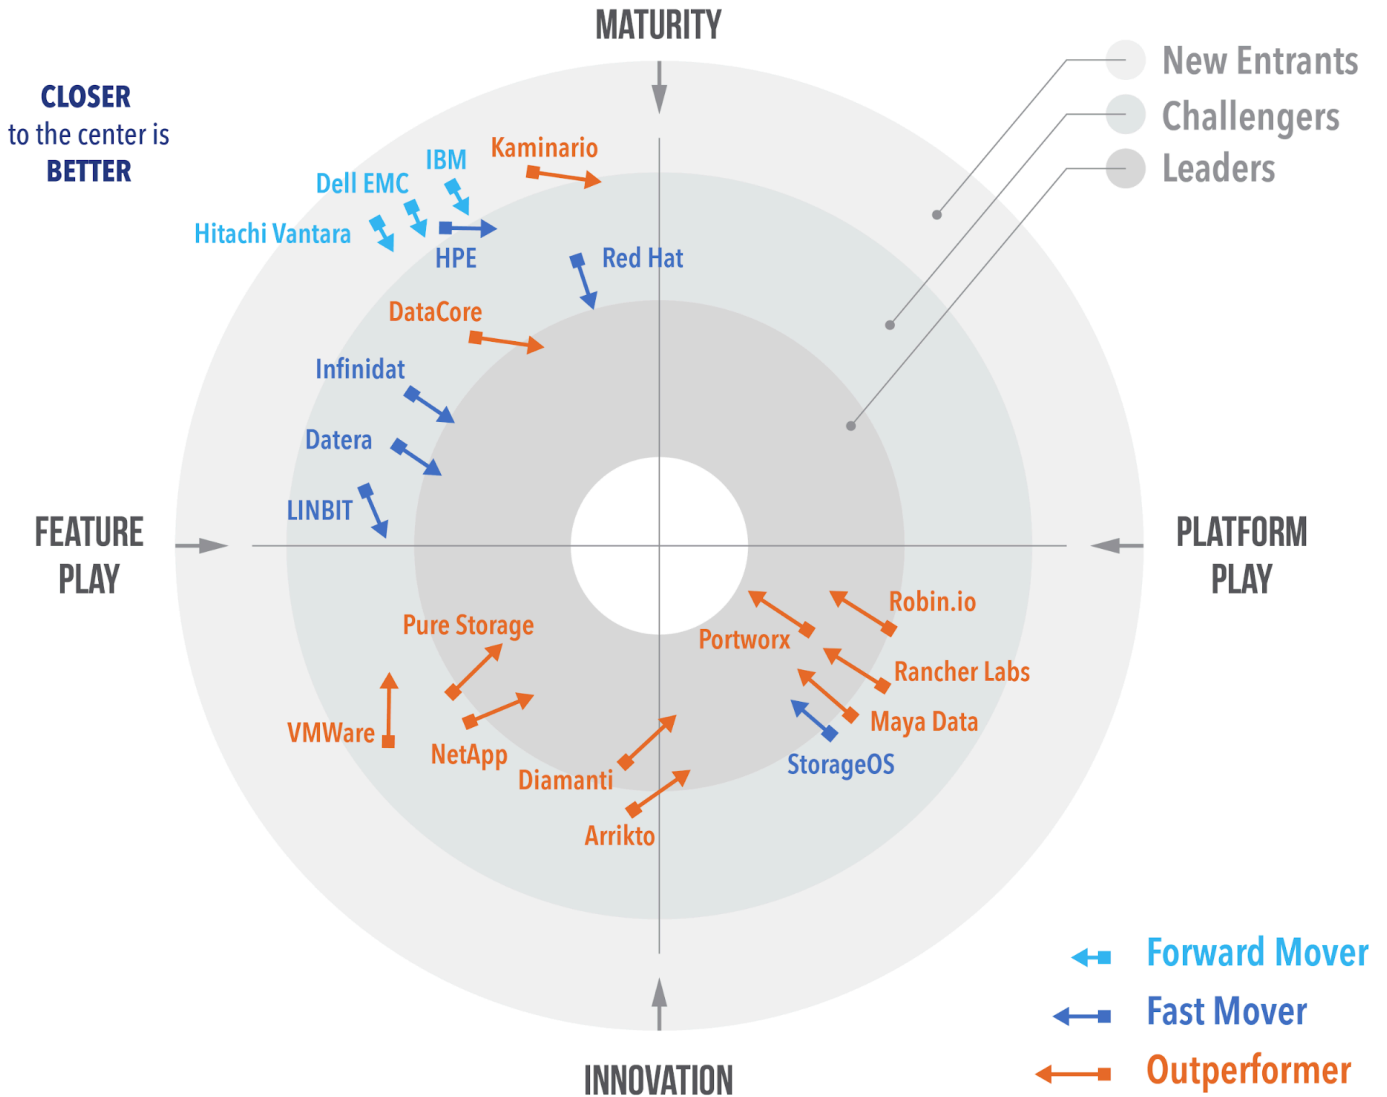
Source: GigaOm 2020



## 5. GigaOm Radar

This report synthesizes the analysis of key criteria and their impact on evaluation metrics to inform the GigaOm Radar graphic in **Figure 1**. The resulting chart is a forward-looking perspective on all the vendors in this report, based on their products' technical capabilities and feature sets.

**Figure 1: GigaOm Radar for Data Storage for Kubernetes**



Source: GigaOm 2020

©GigaOm Radar

## INSIDE THE GIGAOM RADAR

The GigaOm Radar weighs each vendor's execution, roadmap, and ability to innovate to plot solutions along two axes, each set as opposing pairs. On the Y axis, **Maturity** recognizes solution stability, strength of ecosystem, and a conservative stance, while **Innovation** highlights technical innovation and a more aggressive approach. On the X axis, **Feature Play** connotes a narrow focus on niche or cutting-edge functionality, while **Platform Play** displays a broader platform focus and commitment to a comprehensive feature set.

The closer to center a solution sits, the better its execution and value, with top performers occupying the inner Leaders circle. The centermost circle is almost always empty, reserved for highly mature and consolidated markets that lack space for further innovation.

The GigaOm Radar offers a forward-looking assessment, plotting the current and projected position of each solution over a 12- to 18-month window. Arrows indicate travel based on strategy and pace of innovation, with vendors designated as Forward Movers, Fast Movers, or Outperformers based on their rate of progression.

Note that the Radar excludes vendor market share as a metric. The focus is on forward-looking analysis that emphasizes the value of innovation and differentiation over incumbent market position.

## 6. Vendor Roundup/Overview

### Arrikto

Arrikto Rok and Rok Registry are an innovative data management automation and collaboration solution for Kubernetes that leverages a hybrid cloud architecture to enable instance-based NVMe drives to become persistent storage on the worker nodes coupled with a transparent object storage back-end for snapshots, versioning, and a globally federated distributed collaboration platform. This architecture design enables the platform to manage a sophisticated mechanism for snapshots, which is at the base of its copy data management, migration, disaster recovery, and data sharing functionality.

**Strengths:** An alternative approach to data storage and management for Kubernetes that combines the performance of local disk with the advanced data services enterprises expect without inserting a new storage layer. It can be deployed on-premises and in the public cloud.

**Challenges:** The approach is very innovative and the solution still has to prove its efficacy in traditional enterprise environments.

### DataCore

DataCore has a two-fold strategy around persistent storage for Kubernetes. First, the company is developing a CSI plug-in for its block and file storage products aimed at supporting its current customers. And second, it is investing in MayaData, is the sponsor behind OpenEBS and an open source solution designed for Kubernetes storage. This approach will help DataCore stay competitive while developing new solutions going forward.

**Strengths:** The CSI plug-in has a simple and efficient design and aligns with the latest CSI specifications. Easy to adopt and without additional license fees.

**Challenges:** CSI specifications are still immature and can not take advantage of all the potential offered by the DataCore platform.

### Datera

Datera Data Services Platform is a flexible, scale-out, software-defined storage solution that bridges the gap between legacy and modern infrastructures. It is integrated with Kubernetes CSI and aligns well with the latest API specifications.

**Strengths:** Storage Classes are implemented very well, enabling the DevOps team to take full advantage of the backend resources and flexibility of the solution.

**Challenges:** Advanced data services (data protection, remote replication, disaster recovery (DR)) are not yet implemented due to CSI limitations.

## Dell EMC

CSI-Plug-ins are now available for all Dell EMC storage systems, and the company is actively working to improve them. The implementation is still basic, and there are not features aimed at simplifying the deployment and management of Kubernetes storage in large infrastructures.

**Strengths:** Simplicity, free plug-ins, and ease of adoption for Dell EMC customers allow users to take advantage of existing infrastructures and start quickly without further investments.

**Challenges:** The solution is still immature and does not allow planning for hybrid cloud infrastructures, no disaster recovery options, and no official integration yet with Prometheus or other Kubernetes monitoring platforms.

## Diamanti

Diamanti offers an end-to-end solution that resembles hyper-converged infrastructures in the virtualization space. It comes with a fully supported standard Kubernetes distribution pre-installed, or the customer can choose Red Hat OpenShift instead.

**Strengths:** High resilience, performance, and ease of deployment and management are the most important characteristics of this platform, which also offers several features to overcome the limitations of CSI and Kubernetes networking. Thanks to the support of custom resource definitions (CRDs), Diamanti can seamlessly replicate data to the public cloud for migrations or DR.

**Challenges:** The solution is relatively expensive and focused on large enterprise deployments. Even though its basis is on commodity hardware and very efficient, Diamanti offloads many storage and network operations to the hardware, limiting hardware choice.

## Hitachi Vantara

Hitachi Vantara offers a CSI plug-in for all its storage systems. The plug-in is available for all Hitachi Vantara customers and provides a quick path to adopt Kubernetes for Hitachi Vantara customers.

**Strengths:** Hitachi has forward a looking-vision around containers, and Kubernetes in particular, which is aligned with its current product portfolio regarding data analytics and IIoT while protecting existing infrastructure investments.

**Challenges:** Current CSI plug-in is still immature and lacks some features such as remote replication for DR. This is due to the company strategy, focused on adhering to CSI specifications and APIs without

adding non-standard functionalities.

## HPE

HPE is working to offer in 2020 a CSI plug-in with dedicated CSPs (container storage providers) for all its platforms, including 3Par, Nimble Storage, and Primera alongside HPE Cloud Volumes. The goal is to create a uniform CSI user experience across all HPE platforms. Currently the company is in a transition phase, providing strong compliance with CSI specifications while missing some advanced features around data management, remote replication, disaster recovery orchestration, and advanced data placement in complex scenarios.

**Strengths:** Good roadmap, with several features and enhancements to be expected in 2020. HPE recently launched its HPE container platform, with the goal of providing end-to-end solutions for container management.

**Challenges:** Most of the advanced integration and features are not available yet.

## IBM

IBM offers CSI plug-ins for all its major storage platforms, both block and file. The block storage plug-in is based on an open source core and free to download and install. It complies with the latest CSI specifications while being supported with Kubernetes (up to 1.17) and Red Hat OpenShift. Even though the CSI plug-in supports the latest specifications, this solution doesn't offer advanced and automated data placement across systems, data copy management, integrated data protection, remote replication or other features to simplify disaster recovery orchestrations.

**Strengths:** IBM customers can take advantage of the installed base to start quickly with their Kubernetes deployments.

**Challenges:** The current implementation of the CSI plug-in does not enable a simplification of hybrid cloud infrastructure deployments.

## Infinidat

One of the main advantages offered by Infinidat is the possibility to consolidate a large number of applications, workloads, and data in very few systems. In this context, Infinidat provides a compelling solution for customers that want to consolidate Kubernetes applications alongside others.

**Strengths:** The new CSI plug-in aligns with the latest specifications. Additionally, Infinidat simplifies disaster recovery and data migrations across on-premises and major cloud providers thanks to its Neutrix Cloud service.

**Challenges:** Even though data volumes for Kubernetes can group for monitoring in the UI, there is not a specific integration with Prometheus yet.

## Kaminario

Kaminario is transitioning from being an on-premises storage vendor to a multi-cloud data storage platform. In this context, the roadmap shows a compelling potential for data mobility across clouds and platforms, including Kubernetes.

**Strengths:** The current CSI plug-in implementation covers on-premises installations and gives existing users the possibility to adopt this solution for Kubernetes alongside virtualized and physical systems.

**Challenges:** The solution does not offer data migration and replication, and does not support multi-cloud scenarios yet, features that will be released later this year.

## LINBIT

Building on its experience developing the open-source DRBD data replication solution for Linux systems, LINBIT's LINSTOR is a robust and complete software-defined storage solution that provides block storage resources to Kubernetes and other platforms. It supports DRBD, iSCSI and NVMe-oF, as well as common enterprise Linux distributions.

**Strengths:** Open source solution that can be deployed as a standalone storage system or within the Kubernetes cluster. Advanced volume placement and QoS functionalities allow for better performance alongside network traffic optimization.

**Challenges:** LINSTOR already supports snapshots, but a native replication mechanism for disaster recovery will be available only later this year.

## Maya Data

Maya Data OpenEBS Enterprise bundles enterprise functionality and support to the OpenEBS open source project, including a series of tools to improve data migration, mobility, visibility, and infrastructure hardening. Users and developers can adopt OpenEBS at no cost, while Maya Data offers a very flexible licensing.

**Strengths:** Maya Data offers a no lock-in approach for users that want to deploy an open source storage solution for Kubernetes in a multi-cloud environment.

**Challenges:** The solution offers tools to control and migrate data between Kubernetes cluster but it has some limitations with its current implementation, limiting its potential for high demanding enterprise use cases. Future releases of the product will address this issue.

## NetApp

NetApp is building Trident, a complete storage orchestration platform to address Kubernetes challenges. This allows simplifying data storage provisioning and management for DevOps teams while offering a consistent user experience across private and public cloud deployments.

**Strengths:** NetApp's overall strategy is solid, and Trident is a software component that fits very well in NetApp's Data Fabric vision. The solution shows a good feature set and an interesting roadmap, contributing to a good ROI on the NetApp solution.

**Challenges:** Trident is fully CSI compliant but still misses some advanced features that would simplify data migrations, remote replication, disaster recovery orchestration, and backup.

## Portworx

Portworx provides one of the most compelling solutions for enterprise storage dedicated to Kubernetes infrastructures. A cloud-native architecture combined with unique data services that enable enterprise organizations to deploy business and mission-critical applications without the limitations imposed by traditional solutions.

**Strengths:** Portworx Enterprise allows simplifying most of the operations that are currently limiting enterprise Kubernetes deployments for stateful applications, improving data protection and management processes with a positive impact on overall infrastructure TCO. It can work with local storage installed on cluster nodes or traditional enterprise shared storage.

**Challenges:** Most storage vendors offer free CSI plug-ins and backend resource orchestration, allowing enterprises to start with minimal or no investment.

## Pure Storage

Pure Storage has demonstrated again its ability to develop solutions with the right combination of performance and usability across a broad range of use cases. From this point of view, Pure Storage Orchestrator (PSO) has the right characteristics to allow smooth adoption of Kubernetes by leveraging existing storage resources and minimizing initial investment.

**Strengths:** PSO is a good solution and has the potential to become a key differentiator for Pure Storage when the end user is evaluating storage solutions ready to support containers alongside virtualized and physical systems.

**Challenges:** The solution is not ready for all use cases, especially when backup and remote replication for disaster recovery are involved. Pure Storage has best practices in place to mitigate this limitation while working on an implementation following the evolution of CSI specifications.

## Rancher Labs

Longhorn is an open source, lightweight, and efficient distributed block storage solution for Kubernetes that can be deployed on-premises and in the public cloud. Easy to deploy and manage, Longhorn offers a GUI dashboard, advanced snapshot management, and the ability to replicate snapshots remotely for disaster recovery. Rancher Labs offers subscription services for Longhorn including 24/7/365 technical support and consulting. Although originally developed by Rancher Labs in 2014, the Longhorn project is now managed by the Cloud Native Computing Foundation at the sandbox level and the roadmap looks solid. Rancher Labs' dedicated Longhorn engineering team is the main contributor to the project.

**Strengths:** High-performance, open source solution with a growing community of users and developers. Rancher Labs has built an end-to-end ecosystem for enterprise Kubernetes.

**Challenges:** Longhorn 1.0 was released in May 2020. While it supports volume level deduplication and compression in the backup store (which can store backups off-site in AWS S3), it doesn't yet support data reduction in the Kubernetes cluster.

## Red Hat

Red Hat Openshift Container Storage (OCS), based on Ceph, has been designed for simplicity and ease of use. Integrated with Openshift, all management operations are automated using Rook (a CNCF orchestration tool for storage) and automated, enabling DevOps teams with little or no storage knowledge to get a simplified hyperconvergence-like experience.

**Strengths:** Simplicity and ease of use of the solution, completely integrated with the Red Hat OpenShift Container Platform. The solution can be deployed on-premises as well as in the cloud. Simple support subscription model aligned with OCP support and licensing.

**Challenges:** Rigidity of the configuration and limited scalability can pose some risks for large deployments, and the customer may be forced to replace OCS with a standard Ceph installation.

## Robin.io

Robin.io takes an innovative and application-focused approach to Kubernetes data storage and management that enables the user to automate and simplify several aspects of application deployment, data management, and protection while providing better quality of service (QoS) and disaster recovery options. The solution can be deployed both in the public cloud with Google Cloud GKE and in on-premises Kubernetes clusters with solutions like Red Hat OpenShift. Robin.io also includes mechanisms to improve data placement for better performance. The user interface is well organized and easy to use, always giving the user important information on application status in a glance.



**Strengths:** Integrated stack for application and data management, which integrates with the storage layer and provides granular control over application deployment and performance across different environments.

**Challenges:** The Robin.io solution differentiates on its application management capability. However, the user must commit to the full solution stack to gain some of the advanced application management benefits offered by it.

## StorageOS

StorageOS is a cloud-native solution designed explicitly for Kubernetes. As such, the product is more flexible, easier to manage, and integrate with the Kubernetes platform than traditional storage systems. It can deploy both on-premises and in the cloud, providing the same functionalities in different environments.

**Strengths:** Good, lightweight, and efficient architectural design. StorageOS offers development licenses (up to 500GB of storage)

**Challenges:** StorageOS still misses data services features (e.g., snapshots and remote replication) and DR orchestration, limiting the possibility to use it in complex environments that need advanced data protection, data migrations, and fast cloning.

## VMware

Cloud Native Storage (CNS) in VMware vSphere is VMware's solution to provide storage to Kubernetes clusters installed on top of vSphere. It includes a Container Storage Interface (CSI) plugin for Kubernetes and the CNS control plane within vCenter. It is supported with vSAN, vVols, VMFS and NFS datastores. Perfectly integrated in the vSphere environment and very easy to use, it enables VMware customers to manage traditional and modern workloads seamlessly with their existing vSphere knowledge. Additionally, they can take advantage of the storage resources available to the VMware infrastructure for Kubernetes deployments efficiently and without additional costs.

**Strengths:** Well integrated with vSphere and its vSphere client, CNS in vSphere gives system administrators full visibility into storage resources allocated to K8s clusters installed on top of vSphere. When used with vSAN, it offers granular capacity and health monitoring for K8s workloads with better troubleshooting and traceability. It supports vSphere with Kubernetes, Tanzu Kubernetes Grid, and third-party Kubernetes distributions.

**Challenges:** The solution is still limited in the data protection and disaster recovery areas for modern workloads. Further integration with the backup ecosystem is on the current roadmap.

## 7. Conclusion

CSI specifications are still immature, and this is reflected in the products proposed by most vendors. Many storage vendors have chosen a conservative approach that follows the development of CSI specifications. This approach means that the integration between the storage platform and Kubernetes is limited, offering scarce support for all those features that are usually considered mandatory in an enterprise environment.

We identified three groups of vendors in this space, characterized by the level of sophistication of their approach and the features available on their platform. The most conservative ones provide a basic CSI plug-in and expose limited functionality from the array. In this group we find Dell EMC, Hitachi Vantara, and IBM. This severely limits the possibility of implementing Kubernetes for mission-critical environments, especially in hybrid cloud scenarios. HPE presents a similar approach, but its roadmap clearly shows an intention to align its solution with those of the second group.

The second group of vendors consists of Red Hat, Infinidat, Datera, Pure Storage, NetApp, DataCore, StorageOS, LINBIT, and VMware. These companies opted for a more sophisticated approach that gives additional options to the end-user. Some enterprise features are still missing, but the overall strategy of these vendors is much more aggressive, with clear roadmaps and best practices to overcome the limitations imposed by current CSI specifications. Of particular note in this group are NetApp and Pure Storage—their approaches are more holistic and offer greater flexibility in large infrastructure and hybrid cloud environments. Maya Data and Rancher Labs are worth a mention as well because of their open source cores (OpenEBS and Longhorn, respectively) and the innovative tools and services included in their enterprise support subscriptions and product lines.

The leading group is composed of three startups—Diamanti, Portworx, and Robin.io—with products designed specifically for Kubernetes and able to overcome the limitations imposed by CSI. In this group, it is worth noting the diversity of approaches. Diamanti is a hyper-converged solution aimed at building Kubernetes cloud-like experience for large enterprises, with high performance and ease of use as its primary characteristics. By contrast, Portworx emphasizes flexibility and consistent user experience across different environments. It offers an extensive feature set that extends data storage, with data management functionality aimed at building consistent data services across different clouds. Finally, Robin.io has opted for a third approach, with a compelling application-focused solution that hides and automates most storage operations behind the scenes.

## 8. About Enrico Signoretti



Enrico has 25+ years of industry experience in technical product strategy and management roles. He has advised mid-market and large enterprises across numerous industries and software companies ranging from small ISVs to large providers.

Enrico is an internationally renowned visionary author, blogger, and speaker on the topic of data storage. He has tracked the changes in the storage industry as a Gigaom Research Analyst, Independent Analyst and contributor to the Register.

## 9. About GigaOm

GigaOm provides technical, operational, and business advice for IT's strategic digital enterprise and business initiatives. Enterprise business leaders, CIOs, and technology organizations partner with GigaOm for practical, actionable, strategic, and visionary advice for modernizing and transforming their business. GigaOm's advice empowers enterprises to successfully compete in an increasingly complicated business atmosphere that requires a solid understanding of constantly changing customer demands.

GigaOm works directly with enterprises both inside and outside of the IT organization to apply proven research and methodologies designed to avoid pitfalls and roadblocks while balancing risk and innovation. Research methodologies include but are not limited to adoption and benchmarking surveys, use cases, interviews, ROI/TCO, market landscapes, strategic trends, and technical benchmarks. Our analysts possess 20+ years of experience advising a spectrum of clients from early adopters to mainstream enterprises.

GigaOm's perspective is that of the unbiased enterprise practitioner. Through this perspective, GigaOm connects with engaged and loyal subscribers on a deep and meaningful level.

## 10. Copyright

© [Knowingly, Inc.](#) 2020 "*GigaOm Radar for Data Storage for Kubernetes*" is a trademark of [Knowingly, Inc.](#). For permission to reproduce this report, please contact [sales@gigaom.com](mailto:sales@gigaom.com).