

DataCore acquires MayaData for Container-Native Storage

Storage vendors realize they must provide storage for containerized applications, and are moving to do that through internal development, extending their traditional products or by acquiring new technology.

According to Evaluator Group's **Spring 2021 Hybrid Cloud Study**, 84% of enterprises have already adopted or are planning to adopt containers. One-third will keep a mix of traditional and container-based workloads, 29% said it is a goal to move most or all workloads to containers, and 18% have already done so.¹

With the rise in the use of containers for production applications, these apps need persistent storage. Developers are taking responsibility for storage and data services when building applications. Kubernetes allows them to automate the process of allocating storage along with compute and network for containerized apps.

That has led to the emergence of **container-native storage (CNS)**, also known as cloud-native storage, designed for DevOps' storage and data services requirements. Besides allowing developers to deploy storage in Kubernetes, CNS must also provide enterprise services that IT ops require. Vendors are taking several approaches to providing storage for containers, either through CNS or **container storage interface (CSI)** drivers that expose capacity to Kubernetes

Startups have emerged specifically with CNS designed for the DevOps community. These include Diamanti, Ionir, Ondat, Portworx (now part of Pure Storage) and Robin.io. Established storage vendors have extended capabilities around their CSI drivers to bring data services to containerized applications. These include Dell EMC Container Storage Modules (CSM) and NetApp Trident and **Astra**. Some have adapted traditional storage products to run natively in containers, such as **IBM Spectrum Fusion** and Veritas InfoScale. And others have developed CNS as part of their container management platforms, such as Red Hat OpenShift Data Foundation and SUSE Rancher Longhorn.

These products enable provisioning persistent volumes through Kubernetes, but are in various stages of maturity and enterprise-readiness. The goal is to provide tools that gives developers self-service when deploying storage while also enabling IT Ops to maintain control over enterprise data.

Highlights

- DataCore acquires container-native storage startup MayaData
- MayaData's work on OpenEBS adds to DataCore's unified storage vision
- DataCore will develop MayaStor enterprise CNS

¹ Evaluator Group "Hybrid Cloud Matures" July 2021 Research Study

DataCore Unified Storage Vision

Software-defined storage vendor DataCore accelerated its CNS move by acquiring MayaData, 21 months after investing in the open-source startup that developed OpenEBS and MayaStor container-native storage. DataCore, a 23-year-old privately held company based in Fort Lauderdale, Florida, sees MayaData as the next piece of its DataCore One Unified Storage strategy. DataCore currently sells **SANsymphony** block storage, vFilo file storage and Swarm object storage.

Traditionally a block storage vendor, Data Core has built up its product portfolio through acquisition and partnerships over the past two years. MayaData is DataCore’s second acquisition of 2021. It bought Caringo in January for its Swarm object storage, 14 months after launching vFilo through an OEM deal with Hammerspace. DataCore internally developed SANsymphony, which has been available since 2000.

DataCore’s vision is to create a unified management plane to let users manage SANsymphony, Swarm, vFilo – and eventually MayaStor – from one interface.

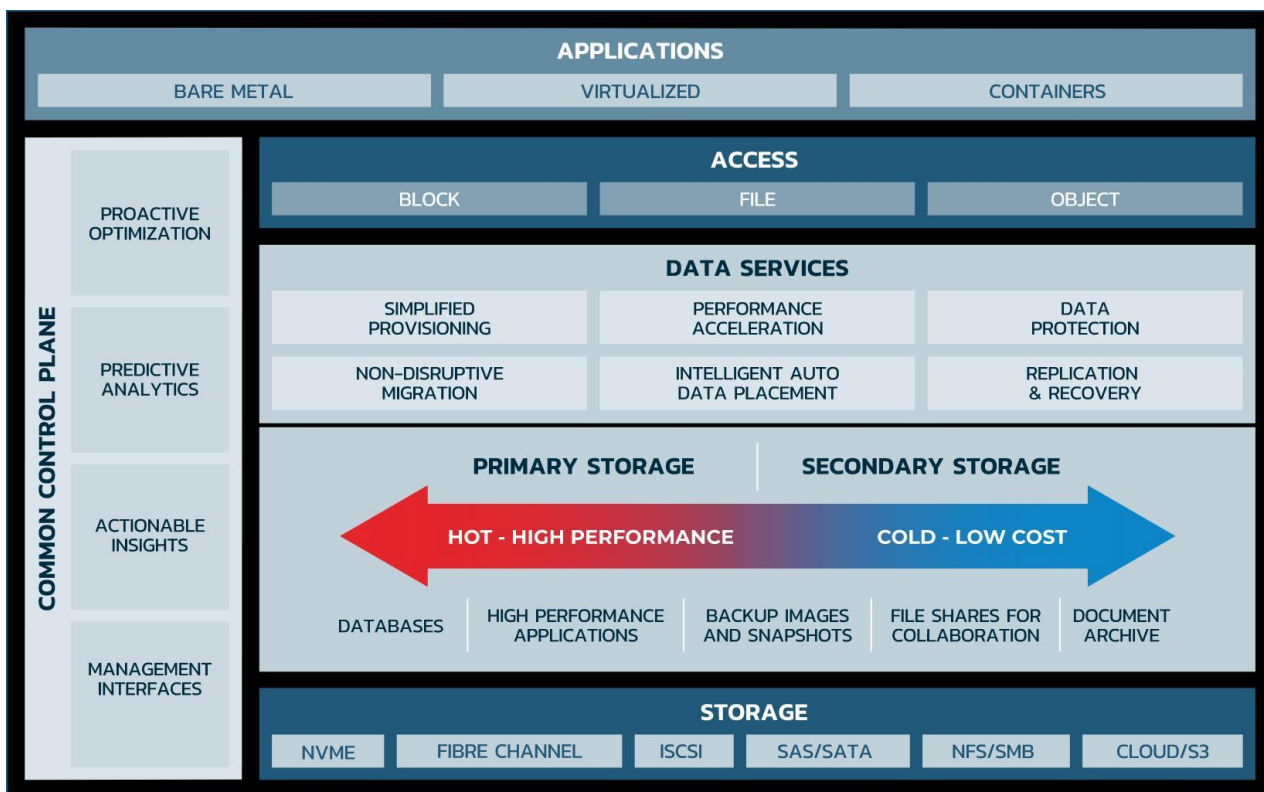


Figure 1: DataCore ONE storage vision includes block, file and object.

Where MayaStor fits

The MayaData acquisition is no surprise, considering its relationship with DataCore. DataCore was part of a group that invested \$26 million in MayaData in Feb. 2020 to help the startup enhance OpenEBS (Open Elastic Block Storage) open-source storage software for Kubernetes.

The investment gave Data Core an ownership stake in MayaData, and DataCore transferred engineers with container storage experience to the startup. The deal also included a cross-licensing arrangement giving MayaData access to DataCore's continuous data protection, mirroring and performance enhancing technology while granting DataCore license for anything MayaData produced.

As part of DataCore, MayaData will keep its branding and maintain distinct engineering and sales teams. DataCore, which claims 2020 was its 12th straight year of profitability, pledges an increased investment in MayaData's R&D and go-to-market teams, and to enhance OpenEBS community support. MayaData currently generates revenue by offering OpenEBS support, and lists Bloomberg, Comcast, Arista and Orange as customers. But it will not generate significant revenue until it releases its MayaStor enterprise CNS product – likely in 2022.

Most of MayaData's employees – largely made up of engineers – are expected to join DataCore. MayaData CEO Don Williams and founder and former CEO Evan Powell will take up advisory roles with DataCore to help the transition.

MayaStor Roadmap

OpenEBS is a CNS platform that provides persistent and containerized block storage for DevOps. Like other CNS software, OpenEBS dynamically provisions volumes and provides data services through containers. OpenEBS uses Kubernetes to orchestrate volume services. It is designed to run on any storage on-premises, at the edge, or in the cloud. It support bare metal and virtual environments, and solid-state drive (SSD) and hard-disk drive (HDD) media.

OpenEBS MayaStor, currently in beta, is designed to add enterprise capabilities to OpenEBS that DataCore will need to appeal to IT operations for production workloads. MayaStor is expected to be generally available in early 2022. MayaStor will be MayaData's choice for apps that require low latency and near disk throughput and replication, and with nodes that include high CPU and RAM resources.

DataCore is looking to accelerate development of MayaStor by adding R&D resources to MayaData. The target use case for MayaStor is applications requiring high IOPs.

MayaStor incorporates Intel's Storage Performance Development Kit for NVMe storage, and will support NVMe over Fabrics (NVMe-of) for shared storage. MayaData is also working on providing availability features such as multipathing to protect against node failures. A new control plane for MayaStor is also in development to better handle scale and resiliency.



MayaStor includes full backup/restore (through Velero open-source software) and synchronous replication. Other features that are planned include RAID support, on-demand capacity expansion, snapshots, clones, and incremental backup. These features are requirements for MayaStor to be fully enterprise ready.

Evaluator Group Comments

The MayaData acquisition is the next phase in DataCore's move to cloud-native storage. DataCore has been going down this path since its 2020 investment with MayaData. DataCore has treated MayaData similarly to a joint venture since making that investment, which included a transfer of some DataCore engineers to MayaData and bilateral licensing agreements.

By going the acquisition route, DataCore is following the path of Pure Storage, which bought Portworx in late 2020. But while MayaData OpenEBS has a following in the development community, it is not as mature as Portworx Enterprise software at the time of Pure's acquisition.

DataCore recognizes the need for enterprise-grade CNS, and is aware that MayaData is still a work in progress towards that goal. We will be watching closely over the next year or so to see DataCore's progress in bringing MayaStor from beta to a hardened enterprise product.

Copyright 2021 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, consequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.