

## REPORT REPRINT

# There's nothing 'mini' about how MinIO approaches object storage

**NOVEMBER 27 2019**

**By Steven Hill, Liam Rogers**

Thanks to vendors like MinIO, object storage is starting to be recognized as the flexible and powerful storage platform of the future as it combines metadata-rich visibility and governance with the performance needed for analytics and other data-intensive workloads.

---

THIS REPORT, LICENSED TO MINIO, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, WAS PUBLISHED AS PART OF OUR SYNDICATED MARKET INSIGHT SUBSCRIPTION SERVICE. IT SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



### Introduction

With its open source, distributed high-performance object storage platform, MinIO is directly targeting large-scale, unstructured data workloads, referencing high-end use cases such as big data and machine learning where the volumes of data involved can be very large. The company claims to provide native performance and efficiency that other object storage offerings aren't geared for.

### 451 TAKE

We believe that object storage is a technology still waiting to be used to its full potential, and that future enterprises will need to adopt a model that efficiently combines file and object to deal with the overwhelming challenges caused by unstructured data growth. MinIO's open source object platform offers a highly efficient, lightweight and cloud-native approach to offering high-performance object storage, along with a substantial list of enterprise-class data protection and management capabilities. The company's recently announced partnership with Qumulo foreshadows a powerful combination of cloud-based, integrated file and object services that can serve the SMB/NFS needs of existing application environments with the advanced availability, governance and multi-cloud management capabilities provided by metadata-rich object storage.

### Context

MinIO was cofounded in 2014 by Anand Babu Periasamy (creator of Gluster), Garima Kapoor, Harshavardhana and Frederick Kautz. Today, most of the company's 42 US employees are based in Palo Alto, California, along with a smaller engineering team operating in India. In terms of user base, MinIO's software has been downloaded via GitHub more than 281 million times. The vendor is building its enterprise customer base via subscriptions to its SUBNET support service, and it claims that its annual recurring revenue is in the millions. MinIO maintains a small direct sales team as well as a partner ecosystem that it leverages as a path to market. The company started with a \$3.3m seed round, and in 2017 it raised a series A valued at \$20m.

In 451 Research's Voice of the Enterprise: Storage, Organizational Dynamics 2019 survey, 37% of organizations have object storage in use and another 7% have it in proof of concept. We anticipate that this number will increase moving forward, especially as customers and vendors alike learn to harness the metadata capabilities of object storage for better management of rapidly growing unstructured data. (We detailed the benefits of object storage and its metadata-rich framework as well as its ongoing role in data management in a past Technology & Business Insight report.)

### Products

MinIO offers an open source and hardware-agnostic object platform that can run on any heterogeneous combination of commodity hardware, but it also provides reference architectures for NVMe and disk-based appliances from multiple major system providers such as Cisco, Dell EMC, HPE and Lenovo. The company's software is designed as a high-performance, highly scalable distributed object storage platform for private cloud applications and features fully S3-compliant capabilities as well as native connectors for both Google Cloud Platform (GCP) and Microsoft Azure. The three components of MinIO's platform are the MinIO core object server, MinIO client and MinIO SDK. Performance is a key element of the platform's design, which leverages a combination of x64 SSE acceleration, streamlined coding and atomic metadata management on top of a single-layer, native object core – this translates to extremely efficient, highly scalable capabilities on commodity hardware.

## REPORT REPRINT

Although it's designed to be very lightweight (the core server code is about 40mb), MinIO's S3/S3-Select-compliant object platform offers a full range of enterprise-class features such as single-sign-on access control and encryption for data security, erasure coding and bit-rot monitoring for data protection, WORM-level immutability, and a unified federation model that supports single-pane-of-glass management for an unlimited number of MinIO instances within a single namespace. Another interesting feature is the vendor's ability to call AWS's serverless Lambda capabilities, which provides a useful model for incorporating highly flexible, demand-based custom services as part of a data ingest/management workflow.

On the container front, MinIO does not employ a container storage interface (CSI) driver to interact with Kubernetes like other storage firms as CSI is for block and file only. MinIO instead leverages the S3 API just as it would otherwise to offer existing Docker volumes with storage. MinIO itself can be deployed as a container via the Kubernetes command line or through a Helm Chart (Helm being an open source application package manager for Kubernetes-based apps) provided by the company.

### Competition

The main forces in object storage continue to be the hyperscalers – AWS (S3), Google (GCP Cloud Storage) and Microsoft (Azure Blob) – so the challenge for providers of on-premises and hybrid object storage lies in being easily extensible to these platforms, with the S3 API now serving as the de facto object interface for the industry. The ever-present incumbents in the storage market, Data Direct Networks, Dell EMC (ECS), Hitachi Vantara (HCP), HPE (HPE partners with vendor Scality), IBM (Cloud Object Storage) and NetApp (StorageGRID), all have object storage platforms among their portfolios. There are also several respected object storage specialists such as Caringo, Cloudian, Cohesity, Scality and SwiftStack that have been cultivating commercial hybrid object platforms for years, and although a commercial version of CEPH has been sold via vendors like Red Hat and SUSE as part of their Linux platforms, it also remains an open source rival to MinIO.

### SWOT Analysis

#### STRENGTHS

MinIO is looking to shatter the outdated impression that object storage is only for slow, cheap and deep use cases, and its published benchmarks for analytics and other high-performance workloads speak for themselves.

#### WEAKNESSES

As is the case with any open source project, the challenge lies in establishing a consumption and support model that appeals to enterprise customers that usually aren't interested in building their own storage infrastructure.

#### OPPORTUNITIES

There is a growing trend toward better integrating the capabilities of file and object to support existing application environments as well as leveraging the immense scalability and metadata of object. Partnerships like the one with Qumulo offer the best of both worlds.

#### THREATS

Object storage is rapidly becoming table stakes in the growing scale-out SDS sector, so there will be a lot of competition from traditional storage providers that are waking up to the flexibility and metadata-based management capabilities of object storage.