

HYP_Link

DAOS Foundation Established to Drive Industry Ecosystem

Mark Nossokoff and Melissa Riddle
April 2024

RECENT DEVELOPMENT

Established in November 2023, [the DAOS Foundation](#) was formed under the auspices of the Linux Foundation to advance the governance and development of the Distributed Asynchronous Object Storage (DAOS) project. The broad range of founding members includes Argonne National Laboratory (ANL), Enakta Labs, Google Cloud Platform (GCP), Hewlett Packard Enterprise (HPE), and Intel. Accordingly, the DAOS Foundation aims to accelerate next-generation technical computing for a broad range of workloads, including both traditional HPC modeling and simulation (mod/sim) as well as modern Artificial Intelligence and Machine Learning (AI/ML).

ANALYST COMMENT

Parallel file systems are critical in achieving maximum system performance at scale for both traditional mod/sim and modern AI workloads. Currently dominated by Lustre and Spectrum Scale, most parallel file systems were developed and optimized for mod/sim workloads and retrofitted, as it were, to incorporate support for AI workloads. Advancements have also evolved on NFS-based NAS storage solutions to provide equivalent, and in some cases better, performance to parallel file systems at scale.

DAOS was conceived as an internal project by Intel in 2015 as a greenfield effort to optimally address large-scale system performance requirements for both mod/sim and AI workloads. DAOS was open-sourced in 2018, making it broadly available while also inviting corresponding community contributions and support. Further traction for DAOS occurred in 2019 when ANL began integrating it into its HPC and AI supercomputer environments, culminating in its integration into the flagship exascale system Aurora, which is anticipated to be fully accepted later this year. GCP provided further support in late 2023 by leveraging DAOS as the foundation for its Parallelstore storage service which targets large-scale generative AI and HPC simulation use cases.

New file systems typically take multiple generations to achieve stability and garner the necessary trust for use with critical applications. Creating the DAOS Foundation within the Linux Foundation is a good start in developing an ecosystem around DAOS. Broad industry adoption, however, will likely take some time as users pay close attention to the integration efforts on Aurora and outcomes delivered on Parallelstore. Should either effort not produce improvements over existing parallel file system and NAS-based solutions, DAOS may not mature further than internal ANL research and GCP-specific services.

Copyright Notice

Copyright 2024 Hyperion Research LLC. Reproduction is forbidden unless authorized. All rights reserved. Visit www.HyperionResearch.com to learn more. Please contact 612.812.5798 and/or email info@hyperionres.com for information on reprints, additional copies, web rights, or quoting permission.