Nominee: Joyent

Nomination title: Manta Storage Service: The convergence of storage and compute

Joyent Manta Storage Service is a cloud storage service that has a novel feature: a massively parallel compute system as a first-class citizen. By putting a high-performance compute cluster in the storage chassis, Manta brings the application to the data instead of the other way around. And by using a simple, popular compute model, a Unix virtual machine, Manta can run a wide range of analytic algorithms on terabyte-class and larger datasets without having to move it across the network. This convergence is transforming a process that is traditionally expensive and time-consuming into something that's fast, easy, and more cost effective. Whereas before only companies with large IT budgets could afford to expend the resources needed to analyze their data, Manta allows companies of all sizes to access the value of their data. Compute-on-storage is a fundamental paradigm shift, indicative of the changing economics and utility of object storage and high-performance data analysis.

Use cases include log analysis, cohort analysis, index generation, financial modeling, and any other data-intensive tasks. Manta allows you to perform modeling across very large data sets with minimal data latency. The object store service also provides a multi-datacenter replication with per-object replication controls, no object size limits, consistent writes and highly available reads, and a filesystem-like hierarchical namespace, including directory.

Why nominee should win

1. A high-performance, massively scaled compute cluster place directly inside an object store service. Large data sets can be accessed almost immediately by analytic algorithms, instead of waiting for data transport latency.

2. An easy-to-use, flexible, and tiered pipeline model. Data scientists and developers can be productive very quickly. Map-Reduce is one common pattern.

3. An unusually powerful remote VM capability that allows "logging into" a storage object for ad hoc analysis.

4. Filesystem: A copy-on-write soft link (a "snaplink") that makes archiving, duplication, and sharing of data efficient; a very familiar hierarchical namespace.