

Cloud-optimized HDD API Standardization Process in OCP

Proposed by Lawrence Ying (Google), Aaron Ogus & Michael McGrath (Microsoft), Jason Adrian (Facebook)

Last Updated: July 21, 2017

Background

Cloud companies store the majority of its online data on a relatively large distributed file system that runs on top of a large number of HDDs (typically >10,000 HDDs). While the exact implementation differs between each Cloud company, many of the desired properties and features at the HDD level are quite similar. Large Cloud companies already customize their HDD FW today in order to realize these desired features in production.

Process Proposal

Goal is to create a new OCP standardization process to facilitate consensus for Cloud Storage around a set of use cases and associated interfaces, in order to accelerate technology development and augment existing standards bodies (T10, T13, SATA-IO, TCG, etc).

Detailed Proposal

1. Any OCP member (the proposer) can initiate a new HDD standardization discussion for Cloud (the proposal). The soft guidance is that the feature must benefit a relatively large distributed file system (typically >10,000 HDDs).
2. The proposal should have (1) a well defined scope, and (2) a proposed high-level schedule.
3. The proposer is responsible to explain the problem statement (the pain point and the intent), as well as the value proposition of the proposal. Providing some concrete data here can help.
4. The proposer will work with others (explained below) to draft the OCP specification together.
 - a. All participants (including the proposer) needs to sign an OCP Joint Development Agreement on a per-company, per-project basis (see [OCP Storage Wiki](#)).
 - b. All communications and discussions will happen in a participants-only forum.
 - c. An OCP-owned online editing document (with trackable electronic revision history) will be used to draft the specification. Examples are Google Docs or Microsoft Word Online.
 - d. Proposal-specific OCP meetings can be created if it makes sense.
 - e. The final spec should describe both (1) the exact ATA/SCSI APIs and the detailed use cases, and (2) how these APIs should be tested to ensure it follows the spec.
5. The proposer will work with the OCP Storage Lead to bring the final spec to the OCP IC (Incubation Committee) for approval.
6. Once the spec is approved within OCP, the HDD companies and Controller companies can start offering these features in their HDD products.
7. OCP participants, who are also members of other storage standards bodies (such as T10, T13, SATA-IO, TCG, etc), can take the relevant interface-specific subset of the OCP approved specification to those standards bodies to be standardized (if applicable). Note that this is beyond the scope of OCP, but we believe this would greatly benefit the whole industry.

Additional Notes

1. As usual, this process results in specs that follow the OCP licensing terms. For more information on OCP Licensing and Branding, please see [this Wiki](#) and [this page](#).
2. OCP has plans to have an extensive certification process (still work-in-progress today).
3. In order to participate in the members-only forum, an existing OCP Member must sign the associated OCP Joint Development Agreement (one time). Please see [this page](#).
4. Kernel implementation requirements are beyond the scope of OCP, but note that multiple OCP members do work on and contribute to kernel development.