

From Dave Anderson

T10/00-295r0

Gene,

Here are some observations on the proposed OSD changes. I really wish I could have been at the meeting, but it just was not possible.

Comments on the differences:

- a. This was included in the original proposals because we had strong comments from customers that a migration path would be required. It seems to me that it could be done whether it is in the standard or not, and will be if there is strong enough demand among users.
- b. I really think that OSD products will implement this based on LUN's that subdivide a given storage space, and as is stated would more likely be a vendor-specific feature.
- c. If sessions are completely separate from Objects, is Jim suggesting that that there is no {OPEN/CLOSE} signal as to whether an object is in use or not? This was one of the things that some, including IBM (and me), thought was quite important for managing the objects?
- d. This change should have been made a long time ago. I had thought of it but never got around to writing it up.
- e. Good.
- f. I am not sure what this means. Sessions have always been optional.
- g. I certainly have no hang-up on changing the syntax of the commands. Is that what is being proposed?
- h. The LIST is a good idea. It was not proposed as it will require application level changes to incorporate, but if that is acceptable, there is no doubt that it makes the protocol much more powerful.

- i. 4.1 metadata attributes. I completely agree that more work is needed.
- 4.2 I agree that what is written is confusing. (You sort of had to be in the discussions to understand what was sought after in the fields in the proposal. A pretty weak rationalization, I admit.)
- 4.3 Good questions raised.
- 4.4 ?? The only fields returned other than Object ID are group remaining capacity, which is returned by several commands and for which there was a lot of disagreement – and session ID. If session ID is to be handled separately, then the only question is the remaining capacity which, as I said, could be dropped from all the responses with very little urging.
- 4.5 The LIST is new and worthwhile, as I have said.
- 4.6 OK by me.
- 4.7 Ditto
- 4.8 I vote for sessions not being persistent, but am willing to be convinced otherwise. (It's helpful in IO fencing.)
- 4.9 Here is the OPEN/CLOSE issue. I think that something is needed so that maintenance functions can have some confidence that an object is not in use, but perhaps a LOCK function is good enough.
- 4.10 I agree that this needs to be worked on. A complete locking mechanism is needed, not a full Distributed Lock Manager necessarily, but a mechanism that file systems can employ.
- 4.11
- 4.12 I fully support access controls for fencing client access to storage.
- 4.13 A lot of work has gone into security research, and I believe the approach suggested in the document appendix is a pretty good start. Still needs work though, I suspect.

Now, the question of changing SAM and whether commands can/should have data going in both directions. Our experience is that both customers and we – as target developers – prefer this. We have some customer unique commands that already do this. Customers helped choose this implementation over a pair of

commands. They are quite happy with it, and our firmware people also liked it as it made things easier to keep straight..

Slide comments:

1. Why cannot an OSD be a controller as easily as a drive? The document attempted to accommodate both.
2. See above
3. –
4. –
5. See earlier comments on sessions versus OPEN/CLOSE. It was thought that clients should be able to share sessions. If you think of a session as allocating a resource, this makes sense.
6. If Create cannot take data, then it becomes a separate, OPEN-like operation. This seemed wasteful in an environment where there are a lot of small file creates. As I mentioned before a more complex command is something we have experience with and think it is not at all a bad idea if a command can have data going in both directions. (I will leave to you to comment on how difficult it would be to change SAM.) I fully support the use of a single CREATE or DELETE(if you like that better than REMOVE) to cover both Objects and Groups. In fact that is how the first draft was written.
7. –
8. I am not sure what is meant/desired by the version of an object. Could this not be an attribute?
9. As I stated, I vote yes for the first two points. As to whether SCSI is the right protocol, my response is: that is where the world is today so that is where we started. It is not intended to be protocol specific, and I would be happy to see an instantiation on another protocol if that made sense.