Object Drives: DCM Lightning Talk

Mark Carlson/Toshiba

What are Object Drives?

- Key/Value semantics (Object store) among others
- Hosted software in some cases
- Interface changed from SCSI based to IP based (TCP/IP, HTTP)
- Channel (FC/SAS/SATA) interconnect moves to Ethernet network

This work is ongoing in the SNIA Object Drive TWG. Please join us at:
https://members.snia.org/apps/org/workgroup/objecttwg/

What is driving the market?

- A number of scale out storage solutions expand by adding identical storage nodes incrementally
  - Typically use an Ethernet interface and may be connected directly to the Internet
- Open source examples include:
  - Scale out file systems
    - Hadoop's HDFS
    - Lustre
  - Ceph
  - Swift (OpenStack object storage)
- Commercial examples also exist

Types of Object Drives

- Key Value Protocol (Object Drive)
  - Minimal incremental CPU/Memory requirements
  - Simple mapping to underlying storage
- In-Storage Compute (Object Drive)*
  - Enough CPU/Memory for Object Node Software to be embedded on the drive
  - General purpose download or factory installed
  - May have additional requirements such as solid state media and more/higher bandwidth networking connections
- In both cases, the interface abstracts the recording technology

*Jim Gray memorial

Key Value Protocol Object Drives

- Eliminates existing parts of the usual storage stack
  - Block drivers, logical volume manager, file system
  - And the *associated* bugs and maintenance costs
- Existing applications need to be re-written, or adapted
  - Mainly used by green field developed applications
  - Firmware is upgraded as an entire image
  - Hyperscale customers are already doing this
    - Using open source software and creating their own “apps” – Facebook, Google, etc.
  - Key Value organization of data is growing in popularity
    - Examples: Cassandra, NoSQL

In-Storage Compute Object Drives

- Same advantages as Key Value protocol plus
  - No need for a separate server to run Object Node service (other services still need a server but scale separately)
    - scaling is smoother – only adding drives
  - Additional features of the Object Node software can be deployed independently
  - Fewer hardware types that need to be maintained (for selected use cases)
  - Failure domains are more fine grained, thus overall data availability is enhanced

In-Storage Compute Future

- As data on the drive becomes colder, CPU/Memory becomes less utilized
  - Possible to host software that then uses this spare resource and works against the cold data
    - Extracting metadata
    - Performing preservation tasks
    - Other data services: advanced data protection, archiving, retention, deduplication, etc.
    - Data Analysis off-load