Should Storage Devices Stay Dumb or Become Smart?

A HotStorage'18 Breakout Session

Carlos Maltzahn
UC Santa Cruz

Tens of PCI-e flash storage devices connected to a fabric have a combined throughput that is two orders of magnitude higher than a typical 16-lane PCI-e bus of a storage server [2]. Current storage systems are designed to hide latency of storage devices instead of hiding latency of the local PCI-e bus and its local memory hierarchy. These systems are unable to make the full capability of flash storage devices available to the application without a significant expenditure of additional server CPUs and memory on data management. With faster fabrics and storage devices and the hyper-exponential growth of data, the gap between what fabric-connected storage devices could deliver and what host-based storage systems allow will only get wider. This is a hardware mismatch induced by storage system architectures as opposed to a problem of performance tuning existing storage systems. At the same time, computational resources within storage devices are reaching a level that could take over some of the data management overhead or even computational tasks from the host, i.e. make the storage device *smart*.

The idea of teaching traditionally *dumb* storage devices new tricks is not new: Active Disks [5] proposed in the late 1990s allow downloading of software into disk drives and securely executing it to assist with data management functionality. More recent explorations include Willow [6] which allows downloading and running custom software on small CPUs within flash devices and BlueDBM [3] which offloads application functionality to FPGAs within flash devices.

This recent interest is motivated by the disruptive trend of greater processing power within storage devices and the hope that this greater processing power will eventually reach a level that enables well-known and mature software frameworks such as the Linux kernel to run within these devices [2]. However, in spite of a number of technology announcements by component makers, none of them have been mass-produced so far.

In fact there is also the opposite trend such as Microsoft Azure's Denali Project [1] that aims to make flash storage devices *dumber* than they are today by moving the flash translation layer into the host to enable closer control of flash behavior and take more quickly advantage of raw flash device innovations. Storage systems vendors also tend to favor *dumb* storage devices since it increases the value of storage systems products and helps commoditize storage devices.

On the spectrum between dumb and smart devices are devices that are *collectively* smart while still assuming frugal CPU and DRAM resources within individual devices [4].

In this breakout session we will aim to elucidate the technical rationales for/against dumb and smart storage devices (and in between) and consider sensitivities to different operating environments such as the cloud datacenter and the edge.

References

[1] CAULFIELD, L. Project denali to define flexible ssds for cloud-scale applications. Web Page. azure.microsoft.com/en-us/blog/project-denali-to-define-flexible-ssds-for-cloud-scale-applications/, March 20 2018.

[2] DO, J. Softflash: Programmable storage in future data centers. In *SNIA SDC* (Santa Clara, CA, 2017).

[3] JUN, S.-W., ET AL. Bluedbm: An appliance for big data analytics. In *ISCA'15* (Portland, OR, June 13-17 2015).

[4] KUFELDT, P., MALTZAHN, C., FELDMAN, T., GREEN, C., MACKEY, G., AND TANAKA, S. Eusocial storage devices - offloading data management to storage devices that can act collectively. *;login: The USENIX Magazine 43*, 2 (2018), 16–22.

[5] RIEDEL, E., GIBSON, G. A., AND FALOUTSOS, C. Active storage for large-scale data mining and multimedia. In *24th international Conference on Very Large Databases (VLDB '98)* (New York, NY, 1998).

[6] SESHADRI, S., GAHAGAN, M., BHASKARAN, S., BUNKER, T., DE, A., JIN, Y., AND SWANSON, S. Willow: A user-programmable ssd. In *OSDI'14* (Broomfield, CO, October 6-8 2014).