

A House Divided: Why Don't Cloud Storage and HPC Storage Share More Technology?

Panelists: Brent Welch (Google), Raghu Raja (Amazon) and Evan Burness (Microsoft)

PDSW 2019 panel discussion notes

- Audience: What is the most important characteristic of cloud storage vs. HPC storage, and how does that affect how you approach solutions in each space?
 - Panelist: cloud environments have a bigger focus on reliability and security, while HPC has emphasis on performance and ability to recompute batch jobs
 - Panelist: the economic model: pay as you go or buy a file system to run for 5 years or longer (in other words, the biggest differences are not technical). It requires customer education to get used to this model. One side effect of the pay as you go model is that it forces users to get more familiar with the efficiency of their data use.
 - Panelist: it's about the semantics. If you want big scale, then you have to be willing to surrender some POSIX semantics.
- Audience: Back to basics: how do you define "cloud storage" exactly? This could mean different things in different domains, and some interpretations are closer to HPC than others.
 - Panelist: It depends on user's perspective; some scientists would think of NERSC as a cloud, but in this context we are probably talking about public cloud providers that have layers (tiers) of storage, some of which could look similar to HPC solutions.
 - Panelist: is a cloud just "someone else's computer?". Agrees that it depends on the user's perspective. It's not actually important to the user unless it feels materially different for what they are trying to do.
 - Panelist: there are other kinds of storage systems that are familiar in the cloud environment, like scale-out key/value databases, not just file

systems. Flexibility and access to other data models are part of the difference.

- Audience: Will it change things for cloud providers when computing power becomes expensive for them? Right now cycles are cheap, so things like aggressive checksumming are viable; you just scale up more nodes.
 - Panelist: Reliability is always a worthwhile tradeoff, you should not optimize that out in the name of efficiency.
 - Panelist: efficiency is just as important as scale for cloud providers already. Things are often optimized within particular contexts. A generic cloud environment might not be perfectly optimal for every kind of workload, but it meets the needs of many users. However, there are specialized offerings that address key workloads. Efficiency is a big focus in those more specialized environments.
 - Panelist: you can already ask for an instance with SLURM, certain number of cores, etc. A cloud environment can really be very similar to an HPC one, except that it scales up and down.
- Audience: In a cloud environment you are charged for IOPs, bytes, etc. This is not how HPC works, though. In which direction will convergence happen? Will HPC become more cloud-like, or will cloud become more HPC-like?
 - Panelist: There is a big distinction between the capital investment model and the pay for what you use model. In a few years, cloud environments will be just as efficient as HPC environments for traditional HPC workloads, though. Look at the work going into custom solutions (hardware, network abstractions, etc.) to improve efficiency in cloud environments.
 - Panelist: Making HPC more cloud-like would require users in HPC to have more awareness of how they use resources. This would probably be a good thing!
 - Panelist: who gets charged for what part of a workflow? This is not a technical question. There are many ancillary things like cooling, support, etc. that aren't considered by most users. Storage is unique compared to other resources because it is often an under-utilized resource. The file system is not always at capacity or fully-subscribed. This is part of why it is harder to deal with.

- Audience: Is it a fair assumption that the diversity of cloud workloads are increasing?
 - Panel: yes.
- Audience: The above assumption requires more flexibility. Will you be able to have programmatically configured systems that can handle HPC workloads effectively? Maybe dynamically defined storage devices instead of direct attach?
 - Panelist: Right now you can already define custom VM shapes (number of cores etc.) for compute, though there isn't that much differentiation in storage yet.
 - Panelist: cloud providers are already doing this really: you can define the kind of remote attached storage that you want.