Centre for High Performance Computing 2019 National Conference



Contribution ID: 177

Type: Talk

Updating the Outdated Storage Paradigm to Handle Complex Computational Workloads

Monday, 2 December 2019 12:00 (30 minutes)

Accelerating discovery in computational science and high performance computing environments requires compute, network and storage to keep pace with technological innovations. Within a single organization, interdepartmental and multi-site sharing of assets has become more and more crucial to success. Furthermore, as the growth of data is constantly expanding, storage workflows are exceeding the capabilities of the traditional filesystem. For most organizations, facing the challenge of managing terabytes, petabytes and even exabytes of archive data for the first time can force the redesign of their entire storage strategy and infrastructure. Increasing scale, level of collaboration and diversity of workflows are driving users toward a new model for data storage.

In the past, data storage usage was defined by the technology leveraged to protect data using a pyramid structure, with the top of the pyramid designated for SSD to store 'hot' data,' SATA HDDs used to store 'warm' data and tape used for the bottom of the pyramid to archive 'cold' data. Today, modern data centers have moved to a new two-tier storage architecture that replaces the aging pyramid model. The new two-tier paradigm focuses on the actual usage of data, rather than the technology on which it resides. The new two-tier paradigm combines a project tier that is file-based and a second or perpetual tier which is object based. The object based perpetual tier includes multiple storage media types, multi-site replication (sharing), cloud, and data management workflows. Data moves seamlessly between the two tiers as data is manipulated, analyzed, shared and protected – essentially creating yin and yang between the two storage tiers. Solutions designed to natively use the Perpetual Tier empower organizations to fully leverage their primary storage investments by reducing the overall strain on the Primary Tier, while at the same time, enabling data centers to realize numerous benefits of the Perpetual Tier that only increase as the amount of storage to manage increases.

The next logical question is how to manage data between the two tiers while maintaining user access and lowering overall administration burdens. Join us for a deeper look into the nuances of the two-tier system and data management between them. We will cover storage management software options; cloud vs. on-premise decisions; and using object storage to expand data access and create a highly effective storage architecture to break through data lifecycle management barriers

Supported Student

Presenter: BUNGAY, Craig (Spectra Logic) **Session Classification:** HPC Technology

Track Classification: Storage and IO