

THETA

The Higher Education Technology Agenda

Really Big Data – Building a HPC-Ready Storage Platform for Research Datasets

The National Computational Infrastructure is Australia's national research computing facility which provides world-class services to Australian researchers, industry and government.

NCI is home to Raijinm one of the Southern Hemisphere's fastest supercomputers, ranked at #24 in the world-wide HPC Top 500 list on debut. NCI is also home to Australia's highest performance filesystems, and one of the nation's largest data catalogues.

To meet the growing demand for additional capacity for multi-petabyte data collections, NCI has recently commissioned an additional large-scale global persistent file storage platform known as gdata3.

The gdata3 storage system based on Intel's Enterprise Edition Lustre, with backend disk provided by NetApp's E-series arrays and EF-series all flash technologies.

Building a persistent HPC-ready data storage platform with a capacity of over 8-Petebytes and throughput of over 120 Gigabytes per second presents many challenges, both from a design and operational perspective.


This presentation will discuss the design solution used to deliver the gdata3 storage platform that will enable researchers using NCI's facilities to access and collaborate on collections of 1-3PB in size from both the High Performance 1.2 Petaflop supercomputer and Cloud systems.

Daniel Rodwell
Data Storage Services, NCI

SHARE THIS:



Loading...

 Follow