

## DTN performance using Lustre and CEPH filesystems

Florent Parent, Lixin Liu, Rolf Seuster, Thomas Tam, Pierre-Antoine Villeneuve, Jocelyn Picard, Simon Guilbault, Thomas Tam, Ryan Enge

### Abstract

Compute Canada and regional partners are deploying and managing advanced research computing (ARC) resources consisting of multi-petabytes storage systems and over 200,000 compute cores. These resources, located across five sites in Canada, are used by Canadian researchers from all regions in the country.

Each Compute Canada site is equipped with high speed networking equipment, allowing 100 Gbps connectivity to the CANARIE IP network. The multi-petabytes storage systems installed at most sites are using either Lustre parallel filesystem or IBM Spectrum Scale filesystem. CEPH storage is also installed in several sites with OpenStack cloud service. Data transfer nodes are used to allow transfers between local and remote sites. Globus service is used across Compute Canada sites to offer high performance file transfer service.

### Goals

1. Measure the single DTN node performance of a Lustre client (max. Lustre I/O performance), and the single Globus I/O node performance;
2. Measure the performance scaling using 2 Globus I/O nodes endpoints;
3. Using the same equipment and environment, measure the mdmFTP performance;
4. Measure the single DTN node performance of a CEPH client.

### Resources

The DTN and Lustre filesystems that will be used are already installed at Université Laval and Compute Canada. We are planning to install a DTN node connecting the Arbutus cluster CEPH storage at University of Victoria. We will run data transfer test between the University Laval Lustre storage and the CEPH storage at Arbutus.

Some configuration are required to provide adequate data measurements and monitoring. The data transfer tools are available online.

The 100G network in CANADA will be provided by the local NREN and CANARIE. Connectivity to SC19 will be

provided by iCAIR/StarLight, according to an agreed time schedule during the conference.

### Involved Parties

This list is partial, and will include more participants as needed.

- Florent Parent, Université Laval, [florent.parent@computecanada.ca](mailto:florent.parent@computecanada.ca)
- Lixin Liu, Simon Fraser University, [liu@sfu.ca](mailto:liu@sfu.ca)
- Rolf Seuster, University of Victoria, [seuster@uvic.ca](mailto:seuster@uvic.ca)
- Pierre-Antoine Villeneuve, Université Laval, [pierre-antoine.villeneuve@calculquebec.ca](mailto:pierre-antoine.villeneuve@calculquebec.ca)
- Jocelyn Picard, Université Laval, [jocelyn.picard@calculquebec.ca](mailto:jocelyn.picard@calculquebec.ca)
- Simon Guilbault, Université Laval, [simon.guilbault@calculquebec.ca](mailto:simon.guilbault@calculquebec.ca)
- Thomas Tam, CANARIE, [Thomas.Tam@canarie.ca](mailto:Thomas.Tam@canarie.ca)
- Ryan Enge, Arbutus site lead, [renge@uvic.ca](mailto:renge@uvic.ca)