

On the challenges of deploying an unusual high performance hybrid object/file parallel storage system in JASMIN

Cristina del Cano Novales
Scientific Computing Department
STFC, RAL
Didcot, UK
cristina.del-cano-novales@stfc.ac.uk

Jonathan Churchill
Scientific Computing Department
STFC, RAL
Didcot, UK
jonathan.churchill@stfc.ac.uk

Athanasios Kanaris
Scientific Computing Department
STFC, RAL
Didcot, UK
athanasios.kanaris@stfc.ac.uk

Robert Döbbelin
Quobyte GmbH
Berlin, AG Charlottenburg, Germany
robert@quobyte.com

Felix Hupfeld
Quobyte GmbH
Berlin, AG Charlottenburg, Germany
felix@quobyte.com

Aleksander Trofimowicz
Quobyte GmbH
Berlin, AG Charlottenburg, Germany
alek@quobyte.com

Abstract— The demand for increasing data volume and velocity in environmental sciences continues to increase with higher data-rate Earth Observation platforms coming on line, such as ESA’s Sentinel satellites [1], and high-resolution climate modelling, such as the CMIP6 programmes [2], being just two examples. JASMIN [3], the UK’s Super-data-cluster hybrid HPC/cloud platform for environmental data analysis, is at the forefront of this deluge, with increasing demands from its over-10,000 users and its supported science communities. Since its inception in 2012, JASMIN has supported successfully the volume and velocity with just two tiers of storage: a single high-performance parallel file system (reaching ~20Pbytes capacity to date) and tape (of around 80Pbytes capacity). However, for the JASMIN phase-4 upgrade installed in the first half of 2018, we have needed to increase the number of tiers to straddle the capacity/performance/cost space more easily. For this, we have extended out from our well-known HPC parallel file system storage to add innovative software-defined hybrid object/file storage platforms of unusual choice and features sets for HPC systems.

In this WIP session, we briefly discuss our 42Pbyte Quobyte [4] software-defined storage system, which is the largest capacity deployment in the world and the first of its kind installed in an HPC context. We will discuss specifically some of the challenges in constructing a novel 5-tier CLOS non-blocking low latency routed Ethernet network of sufficient bandwidth (20’s Tb/s) and scale (1,000 100/50/40 Gb ports and 1,000 10Gb ports) to support JASMIN’s growth.

Preliminary results will be presented, showing our first WIP benchmarking output reaching >200Gbyte/s, and we hope to be able to report a first ever IO500 [5] benchmark result for Quobyte at SC18.

Keywords—

REFERENCES

- [1] ESA Sentinel UK Data Hub on JASMIN: <https://www.gov.uk/government/news/europes-earth-observation-programme-maximised-by-uk-data-hub>
- [2] Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., Taylor, K. E. 2016. Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization. *Geosci. Model Dev.*, 9, 1937-1958, doi:10.5194/gmd-9-1937-2016.
- [3] Lawrence, B.N., Bennet, V.L., Churchill, J., Jukes, M., Kershaw, P., Pascoe, S., Pepler, S., Pritchard, M., Stephens, A. 2013. Storing and manipulating environmental big data with JASMIN. *IEEE International Conference on Big Data*, San Jose, CA, USA, 6-9 Nov 2013: 68-75.
- [4] Quobyte, <http://www.quobyte.com>.
- [5] Kunkel, J.M., Bent, J., Lofstead, J., Markomanolis, G.S. 2017. Establishing the IO-500 Benchmark. WIP at PDSW-DISCS, SC17, Denver, Co, USA, 13 Nov 2017.