



Beating the tyranny of scale with a private cloud configured for Big Data

Bryan Lawrence (1,2,3), Victoria Bennett (2,4), Jonathan Churchill (5), Martin Juckes (1,2), Philip Kershaw (2,4), Sam Pepler (2), Matt Pritchard (2), Ag Stephens (1,2)

(1) National Centre for Atmospheric Science, UK, (2) Centre for Environmental Data Archival, RALSpace, STFC, Didcot, UK, (3) Department of Meteorology, University of Reading, Reading, UK, (4) National Centre for Earth Observation, (5) Scientific Computing Department, STFC, Didcot, UK

The Joint Analysis System, JASMIN, consists of a five significant hardware components: a batch computing cluster, a hypervisor cluster, bulk disk storage, high performance disk storage, and access to a tape robot. Each of the computing clusters consists of a heterogeneous set of servers, supporting a range of possible data analysis tasks - and a unique network environment makes it relatively trivial to migrate servers between the two clusters. The high performance disk storage will include the world's largest (publicly visible) deployment of the Panasas parallel disk system. Initially deployed in April 2012, JASMIN has already undergone two major upgrades, culminating in a system which by April 2015, will have in excess of 16 PB of disk and 4000 cores. Layered on the basic hardware are a range of services, ranging from managed services, such as the curated archives of the Centre for Environmental Data Archival or the data analysis environment for the National Centres for Atmospheric Science and Earth Observation, to a generic Infrastructure as a Service (IaaS) offering for the UK environmental science community.

Here we present examples of some of the big data workloads being supported in this environment - ranging from data management tasks, such as checksumming 3 PB of data held in over one hundred million files, to science tasks, such as re-processing satellite observations with new algorithms, or calculating new diagnostics on petascale climate simulation outputs. We will demonstrate how the provision of a cloud environment closely coupled to a batch computing environment, all sharing the same high performance disk system allows massively parallel processing without the necessity to shuffle data excessively - even as it supports many different virtual communities, each with guaranteed performance. We will discuss the advantages of having a heterogeneous range of servers with available memory from tens of GB at the low end to (currently) two TB at the high end.

There are some limitations of the JASMIN environment, the high performance disk environment is not fully available in the IaaS environment, and a planned ability to burst compute heavy jobs into the public cloud is not yet fully available. There are load balancing and performance issues that need to be understood.

We will conclude with projections for future usage, and our plans to meet those requirements.