



Contribution ID: 75

Type: oral presentation

Architectures and methodologies for future deployment of multi-site Zettabyte-Exascale data handling platforms

Thursday, April 16, 2015 9:00 AM (15 minutes)

Several scientific fields, including Astrophysics, Astroparticle Physics, Cosmology, Nuclear and Particle Physics, and Research with Photons, are estimating that by the 2020 decade they will require data handling systems with data volumes approaching the Zettabyte distributed amongst as many as 10^{18} individually addressable data objects (Zettabyte-Exascale systems). It may be convenient or necessary to deploy such systems using multiple physical sites. This paper describes the findings of a working group composed of experts from several large European scientific data centres on architectures and methodologies that should be studied by building proof-of-concept systems, in order to prepare the way for building reliable and economic Zettabyte-Exascale systems. Key ideas emerging from the study are: the introduction of a global Storage Virtualization Layer which is logically separated from the individual storage sites; the need for maximal simplification and automation in the deployment of the physical sites; the need to present the user with an integrated view of their custom metadata and technical metadata (such as the last time an object was accessed, etc.); the need to apply modern efficient techniques to handle the large metadata volumes (e.g. Petabytes) that will be involved; and the challenges generated by the very large rate of technical metadata updates. It also addresses the challenges associated with the need to preserve scientific data for many decades. The paper is presented in the spirit of sharing the findings with both the user communities and data centre experts, in order to receive feedback and generate interest in starting prototyping work on the Zettabyte-Exascale challenges.

Primary author: DELFINO REZNICEK, Manuel (Universitat Autònoma de Barcelona (ES))

Co-authors: HEISS, Andreas (KIT); SANSUM, Andrew (RAL); DELAUNAY, Benoit (IN2P3); MATTHEWS, Brian (RAL); NEISSNER, Christian (IFAE); CORNEY, David (RAL); CANCIO, German (CERN); LAMANNA, Giovanni (CNRS); BIRD, Ian (CERN); COLLIER, Ian Peter (RAL); SHIERS, Jamie (CERN); NIEF, Jean-Yves (IN2P3); FLIX MOLINA, Jose (CIEMAT); DELL'AGNELLO, Luca (INFN); MAGGI, Marcello (INFN); PORTO FERNANDEZ, Maria Del Carmen (CIEMAT); SCHULZ, Markus (CERN); GASTHUBER, Martin (DESY); FUHRMANN, Patrick (DESY); MACCHI, Pierre-Etienne (IN2P3); BOCCALI, Tommaso (INFN); ACIN PORTELLA, Vanessa (IFAE); GUELZOW, Volker (DESY)

Presenter: DELFINO REZNICEK, Manuel (Universitat Autònoma de Barcelona (ES))

Session Classification: Track 3 Session

Track Classification: Track3: Data store and access