



Contribution ID: 32

Type: Oral presentation to parallel session

Forming an ad-hoc nearby storage framework, based on the IKAROS architecture and social networking services

Monday, 14 October 2013 14:39 (20 minutes)

Given the current state of I/O and storage systems in petascale systems, incremental solutions in most aspects are unlikely to provide the required capabilities in exascale systems. Traditionally I/O has been considered as a separate activity that is performed before or after the main simulation or analysis computation, or periodically for activities such as check-pointing, but still as separate overhead. I/O architectures, when designed in this way, have already shown not to be scalable as needed. At the same time, Grid computing implementations are not considered to be user friendly and the Virtual Organization (VO) approach is not very efficient for individual users and small groups.

We present an ad-hoc "nearby" storage framework, based on the IKAROS architecture and social networking services (such as Facebook), in order to address the above problems. IKAROS is, by design, able to add or remove nodes to or from the I/O system instance on the fly, without bringing everything down or losing data. Furthermore, it is capable to decide the file partition distribution schema, by taking into account user and application requests, as well as domain and VO policies. We are using existing social networking services, such as Facebook, in order to dynamically manage, share and publish meta-data. In this way, we do not need to build our own utilities for searching, sharing and publishing and at the same time we enable users to dynamically use the infrastructure, by creating ad-hoc storage formations. This results in a model which can scale both up and down and so can provide more cost effective infrastructures for both large scale and smaller size groups and consortiums. This approach gives us the opportunity to build a more efficient ad-hoc nearby storage: multiple instances of smaller capacity, higher bandwidth storage closer to the compute nodes.

To achieve this, IKAROS is using the JSON (JavaScript Object Notation) format for populating meta-data, which is very common within Web 2.0 technologies. Individual users or groups can dynamically change their meta-data management schema based on their needs. We use Facebook as a use case to show how easily IKAROS can connect with existing, successful infrastructures which already scales to millions of users. Although we are currently using Facebook, it should be emphasized that we are not constrained by this specific choice. IKAROS is responsible only for the core services and Facebook for the meta-data utilities. In this way, the responsibilities are kept separated and the two infrastructures can scale independently from each other.

Summary

Primary author: FILIPPIDIS, Christos (Nat. Cent. for Sci. Res. Demokritos (GR))

Co-authors: MARKOU, Christos (Nat. Cent. for Sci. Res. Demokritos (GR)); COTRONIS, Yiannis (University of Athens)

Presenter: FILIPPIDIS, Christos (Nat. Cent. for Sci. Res. Demokritos (GR))

Session Classification: Data Stores, Data Bases, and Storage Systems

Track Classification: Data Stores, Data Bases, and Storage Systems