

Extending Rocks Clusters into Amazon EC2 Using Condor

MURPA Seminar Friday 23 April at 10am

Speaker: Dr Phil Papadopolous, Program Director, UC Computing Systems, UCSD and Program Director SDSC

Venue: Seminar Room 135/26, Monash University Clayton

Abstract

The open-source Rocks clustering toolkit (www.rocksclusters.org) enables users to easily build both simple and complex clusters that are made up of both heterogeneous function and hardware. As part of our National Science Foundation award (#OCI-0721623), we have enabled user to build clusters that have some are all of their components as Xen-based virtual machines. By treating, Xen-based hardware as "just another brand of server", nearly all of the Rocks tools used on real hardware carry over into the virtual space. In particular, a node that has a particular logical function (e.g. compute, frontend, web, database, file server, and others), which we define as an "appliance", can easily be built as a VM or on real hardware. Recently, we developed the EC2 Roll (rolls are our fundamental, programmatic mechanism for defining the software and configuration of appliances) that allows us to easily build an EC2 version of any already-defined appliance. Finally, when coupled with the Condor Roll, we are able to automatically extend local clusters into the Amazon cloud using Condor as the scheduling and load manager. These three mechanisms together allow users to rigorously define the contents of a virtual machine in EC2 while keeping job management and submission on a local head node. Essentially, the local cluster is temporarily extended into the cloud, but with the full set of software required by the cluster owner.

In this talk, we describe the underlying mechanisms as well as some of the caveats to performing this kind of cluster extension.

Bio

Phil Papadopoulos received his PhD in 1993 from UC Santa Barbara in Electrical Engineering. He spent five years at Oak Ridge National Laboratory as part of the Parallel Virtual Machine (PVM) development team. Currently, he is program director of Grid and Cluster Computing at the San Diego Supercomputer Center. Papadopoulos is deeply involved in key research projects including the Biomedical Informatics Research Network (BIRN), OptIPuter, the Geosciences Network (GEON), and the Pacific Rim Applications and Grid Middleware Assembly (PRAGMA). He is also well known for development of the open-source Rocks Cluster toolkit, which has an installed base of 100s of clusters. His research interests focus on distributed and clustered systems and how they can be used more effectively in an expanding bandwidth-rich environment. Papadopoulos is co-investigator for the OptIPuter and CAMERA projects.