

CloudCRV - Cluster Deployment and Configuration Automation on the Cloud

Wednesday, 3 November 2010 15:00 (20 minutes)

With the development of virtualization technology and IaaS cloud, it is much easier for users to obtain large number of computing (virtual) resource. However, customizing these resources as a computing clusters remains a difficult task that requires in-depth IT knowledge and complex user-specific customizations. We develop a tool called CloudCRV (Cloud-Cluster-Role-VMs) to help users design, distribute and deploy a secure, functional cluster on the allocated resources.

We believe a predefined cluster in a whole can be distributed as a product to perform certain task (e.g. an ATLAS Tier3 cluster), we call this kind of product a Virtual Cluster Appliance (As extension to Virtual Appliance). The purpose of CloudCRV is to help the Cluster Designer to design such a product and to help the Cluster Managers to deploy it.

Most clusters can be abstracted to a set of Roles (e.g. a NFS server, or a Condor Head) and their relations (Condor Head depends on NFS server). Cluster Designer's work is to define the Roles and their relations. The Roles are defined with the help of configuration management systems such as Puppet or Cfengine. Once designed, the Virtual Cluster Appliance can be deployed at multiple sites by local Cluster Managers onto physical or virtual resources. CloudCRV provide interfaces to both Cloud Providers (such as EC2 and Nimbus), and to physical computers and libvirt based clusters via gPXE remote booting and image deployment.

In this contribution we demonstrates the process of designing and deploying such a Virtual Cluster Appliance with the help of CloudCRV.

Summary

Comments to the Reviewers:

We define the following three concepts:

VM:

A running virtual or physical computer that has a supported OS.

Role:

a set of requirements on a VM that once met, will let the VM provide certain functionality. E.g. NFS server, Condor Head, Condor Worker are all Roles. One Role can depend on another Role, e.g. Condor Worker depend on Condor Head.

Roles are defined by Configuration Management System Languages (e.g. Puppet/CFengine).

Cluster:

A collection of Roles.

CloudCRV contains two parts, the CloudCRV Designer will help a Cluster Designer (e.g. the Tier3 Working group of ATLAS) to design the cluster, test its functionalities, and package it. As a Cluster Designer, one has to have certain level of system administration knowledge.

The CloudCRV Manager, on the other hand, will help a Cluster Manager to customize and deploy a predefined cluster. The Cluster Manager need to have access to the hardware resources (e.g. and EC2 account, or a physical cluster), and little system administration knowledge is needed.

The key is design once, deploy multiple times. This gives a normal cloud user the power of running a cluster with least effort.

Primary author: YAO, Yushu (Lawrence Berkeley National Lab. (LBNL))

Presenter: YAO, Yushu (Lawrence Berkeley National Lab. (LBNL))

Session Classification: Grids and Clouds