

Title: Exploiting Virtualization and Cloud Computing in ATLAS

The ATLAS Computing Model was designed around the concepts of grid computing; since the start of data-taking, this model has proven very successful in the federated operation of more than one hundred Worldwide LHC Computing Grid (WLCG) sites for offline data distribution, storage, processing and analysis. However, new paradigms in computing, namely virtualization and cloud computing, present improved strategies for managing and provisioning IT resources that could allow ATLAS to more flexibly adapt and scale its storage and processing workloads on varied underlying resources. In particular, ATLAS is developing a "grid-of-clouds" infrastructure in order to utilize WLCG sites that make resources available via a cloud API.

This work will present the current status of the Virtualization and Cloud Computing R&D project in ATLAS Distributed Computing. First, strategies for deploying PanDA queues on cloud sites will be discussed, including the introduction of a "cloud factory" for managing cloud VM instances. Next, performance results when running on virtualized/cloud resources at CERN LxCloud, StratusLab, and elsewhere will be presented. Finally, we will present the ATLAS strategies for exploiting cloud-based storage, including remote XROOTD access to input data, management of EC2-based files, and the deployment of cloud-resident LCG storage elements.

Author: Dan van der Ster

Co-Authors: Fernando Barreiro, Kaushik De, Rod Walker

Type: parallel

-- DanielVanDerSter - 20-Sep-2011

---

This topic: LCG > CHEP12ATLASCloudComputing

Topic revision: r2 - 2011-10-11 - unknown



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback