New Project with Huawei - Project Goals and Initial Plan

Dirk Duellmann, CERN IT

Openlab Major Review Meeting
13 October 2011
CERN, Geneva

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

HUAWEI

CERN openlab

Why Cloud Storage & S3 Protocol?

- Cloud computing and storage gain rapidly in popularity
  - both as private infrastructure and as commercial service
  - Several investigations are taking place also in the HEP and broader science community
- Price point of commercial offerings may not (yet?) be comparable with services we run at CERN or WLCG sites, but
  - Changes in protocols, deployment model and offered consistency offer increased scalability with reduced deployment complexity (TCO)
  - Market is growing rapidly and we need to understand if promises can be confirmed with HEP work loads
  - We need to understand how cloud storage will integrate with (or change) current HEP computing models

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CERN Interest

- Main Interest - Scalability and TCO
  - can we run cloud storage systems to complement or consolidate existing storage services
- Focus: file storage for physics and service data
  - analysis disk pools, home directories, grid wide distributed user space, virtual machine image storage
    - very different in volume, number of concurrent clients, sequential vs random access,
- Which areas of the phase space can be covered by cloud storage with confirmed advantages in scalability and TCO?
  - immediate targets:
    - can we reliably setup and run a cloud storage service of medium scale (eg close to PB scale) ?
    - can we confirm scalability gain and deployment gains at CERN ?

WLCG Interest

- S3 Protocol is seen as potential standard interface to access (and federate) site storage
- Allowing to provide (buy or rent) storage services without change to user application
  - Eg large sites could run private clouds on acquired hardware
  - Smaller sites may run other implementations of S3 or rent capacity on demand
- Requirements
  - demonstrated success at one site (eg CERN)
  - demonstrate distribution across sites according to experiment computing models

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Interest for Huawei?

- Huawei is working on new products to implement cloud storage (eg S3)
  - new hardware architecture wrt to traditional storage systems
  - new software layering to implement a variety of storage services on top of a common highly-scalable clustered store
  - new deployment concepts with respect to hardware and software failures
- First installations exist now in the lab (visit and demo). Perfect point in time to start:
  - compatibility and performance tests against standard and open source solutions
  - get experience with continued operation of a larger setup under quasi production conditions
  - perform ramp-up scalability tests with a very large number of concurrent clients (~1000s)
- Openlab - you make it, we break it
  - We have shown several times that openlab can have a significant positive impact on product quality via a close interaction between experts at CERN and at the vendor
  - Openlab is a very attractive shop window once successful deployment has been reached

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

The Huawei Hardware Setup

- Huawei will provide initially a two rack system with some 768TB capacity in a single storage cluster along with internal networking and support infrastructure
- CERN will contribute the necessary client resources to run large scale tests according to the joint test plan

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Resources & Coordination

- The project will be hosted at CERN in the Data & Storage Services group
  - close to experts for all major storage services and investigations in IT
- A fellow is already starting preparatory work
  - a short-term student will join next week to construct a S3 storage test suite
- Will setup phone conferences and in-person visits as required
  - regular project review as part of the standard Openlab review meetings

Project Plan

- Huawei has joined the CERN Openlab initially for one year
  - work plan addressing the CERN and Huawei goals
- High level
  - CERN ready to receive hardware (Jan 2012)
    - start of commissioning phase
  - Huawei storage commissioned (Feb 2012)
    - passing initial functional tests
  - Low-scale Amazon S3 compatible style benchmark is performed (Jun 2012)
  - Ramping up in storage and client nodes to full capacity scale test (Oct 2012)
    - scheduled to not interfere with 2012 LHC running
  - Review / Dissemination according to mutual agreement

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Potential Future Extensions

- Areas of mutual interest beyond this initial program include:
  - Monitoring and integration in the overall computer centre management system
  - Metrics on the reliability rates of the data storage system and its components
  - Comparison of relaxed-POSIX filesystems
- Benchmarking of further HEP applications that could make use of the compute power embedded in the data storage system

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it