LLR Computing Resources

A. Sartirana

Computing @ LLR

- The IT team at LLR **started managing computing services** to allow the LLR **CMS group** to be **present and active in WLCG**
  - in 2005 LLR **joined** the newborn **GRIF** T2 WLCG site
    - along the years LLR has become the **«CMS-pole» of GRIF;**
  - the local grid site turned out being a **very precious resource;**

- … along the years, as the **landscape evolved**, we've tried to adapt our computing offer and **keep up with new needs and new technologies**
  - other groups expressed computing needs (not necessarily "grid-compliant");
  - **HPC, manycores,** Cloud,...

General Guidelines

- We provide **"local" computing solutions** when/if they are needed
  - we are **complementary to** the activity of **computing centers;**
  - we provide users with tailored solutions to their specific needs.

- We try, as much as possible, to **mutualize resources and competences**
  - more **effective usage of resources;**
  - possibility to **absorb peaks** of requests;
  - **ease of administration.**

The Infrastructure

- Polytechnique provides a **fully equipped machine room**
  - **600kW** cooling and electric power, **400kW under USB**
    - electric energy/room maintenance/monitoring… all provided by the Ecole;
  - **~200m^2** shared with other(~12) labs
    - we currently have 10 racks (over 50. max possible 72);
  - **renewed Fall 2014.**

- Very good **WAN NW** by **IN2P3/Renater/Saphir**
  - **GrifOPN**
  - LHCOpn/**LHCOne**

The Grid Site

- LLR is one of the 6 **members of** the **GRIF Tier-2**
  - **~2.5k cores** (HTCondor + CREAM), **1PB** storage (DPM);
  - the **CMS "pole"** of GRIF
    - contacts with central CMS computing;
    - GRIF is one of the most reliable and active CMS T2's;
    - some general **CMS-related services/activities**: Xrootd EU redirector, CMS Middleware readiness testbed;
  - other active "local" VO's: ILC, T2K, Hess, Harpo.
- Local **Tier-3**
  - **~500 cores** (depends on needs);
  - shares T2 grid storage + NFS (~50TB)
    - ~50% CMS grid storage is user data;
  - **easy access**: batch, interactive access;
  - **important resource for our CMS group.**

The HPC Clusters

- Around 2012 we started managing some **HPC resources**
  - **driven by Galop** which needs a **local resource for developing** its PIC code (SMILEI)
    - **replicating PRACE/GENCI conditions** (high. tech. profile);
    - CMS as well: Higgs analysis in the VBF and ttH channels;
  - clusters are **mutualized with** the **LSI** lab
    - sharing usage, contributing to purchases and management;
    - MoU (to be signed) ruling collaboration.

- **JollyJumper:** 672 cores, FDR (56 Gb/s) IB connection
  - 2 x E5-2697 (12cores), 2.7GHz, 2GB/core ram;
- **Rantanplan:** 128 cores, QDR (40Gb/s) IB connection
  - 2 x E5620 (4cores), 2.4GHz, 6GB/core ram;
- **SPS NFS storage** (15TB) 10Gb/s and FDR IB NW.

The GridCL Platform

- **P2IO** proj: **GPGPU/manycore R&D with OpenCL** (2012)
  - **heterogeneous** manycore/GPU devices for testing the code
    - **portability** being one of the key features of OpenCL;

- ✓ 2 nodes with 2 Nvidia K20 GPU cards each;
- ✓ 2 nodes with 2 Intel Xeon Phi cards each;
- ✓ 1 node with 6 Nvidia Titan GPU cards;
- ✓ 1 node with 2 AMD FirePro GPU cards;

- **dev. activity** by members and externals (e.g. CERN);
- **training**: JDEV '13, IN2P3 School '16;
- "small" productions;
- **playground** for getting in touch with **new tech's**
  - e.g. CMS users currently testing **deep learning** algorithms (via python) on **GPU**-aware **dockers**.

What about Cloud?

- **Many Cloud and Cloud-related projects around us**
  - LAL/Orsay OpenSTACK Cloud. Indigo DataCloud/Fed. clouds projects.
- Some **attempts** to build a computing cloud **at LLR**
  - started with a stratuslab cloud project. Now looking to Openstack;
  - partners in the proposed P2IO CEPH project;
  - **not a clear need by users**. Thus we have **never really raised the priority** to a level which could push us to production mode.
- Should we raise the priority of such projects?
  - not clear if there is a **general strategy by in2p3;**
  - at equal money/manpower this is a **resources trade-in with other computing platforms** (Grid, HPC, …).

Summing up

- Longstanding experience in **managing and delivering computing services**
  - **driven by users** requirements. **Mutualizing** resources when possible;
  - begun with **CMS Grid.** Following the evolution of needs and technologies we slowly drifted to **include HPC;**
    - the same approach hasn't led us to Cloud yet;
- **Know-how** acquired
  - manage **heterogeneous resources** and computing services: HTC, HPC, GPUs, satellite services…;
  - "customer care": quickly provide **flexible solutions to users needs;**
  - deal with **mutualized/federated "environments"**: GRIF, LLR/LSI,….

Backup Slides

CMS Grid site