HPC for High Energy Physics

Humberto Salazar I

XXX Reunión Anulal de la DPyC, SMF. May, 23-25, 2016

LNS
Laboratorio Nacional de Supercómputo
del Sureste de México

Outline of the talk

- National Supercomputing Laboratory (LNS) design &Infrastructure
- Services
- High Energy Physics Aplications
- Installation of a Tier 2 node of the Worldwide LHC Computing Grid in LNS.
- Summary

LNS
Laboratorio Nacional de Supercómputo del Sureste de México

BUAP

Infrasestructure of LNS

- Data Center and Communications (Certification on International Computer Rooms Experts Association, **ICREA level III**)
- HPC Equipment
- Expertise Team

Data Center's

Availability of Electrical and Communications services
Nearly 100%

Communications Room

Electrical Room

Supercomputing Equipment

- FUJITSU Company
- 272 computer nodes
- Interconnection Infiniband FDR 56 Gb/s
- 200 TB /home
- 1 PB Storage (For all community: Academic & Industrial)
- ~ 200 TFLOPS on Performance

Supercomputing Equipment

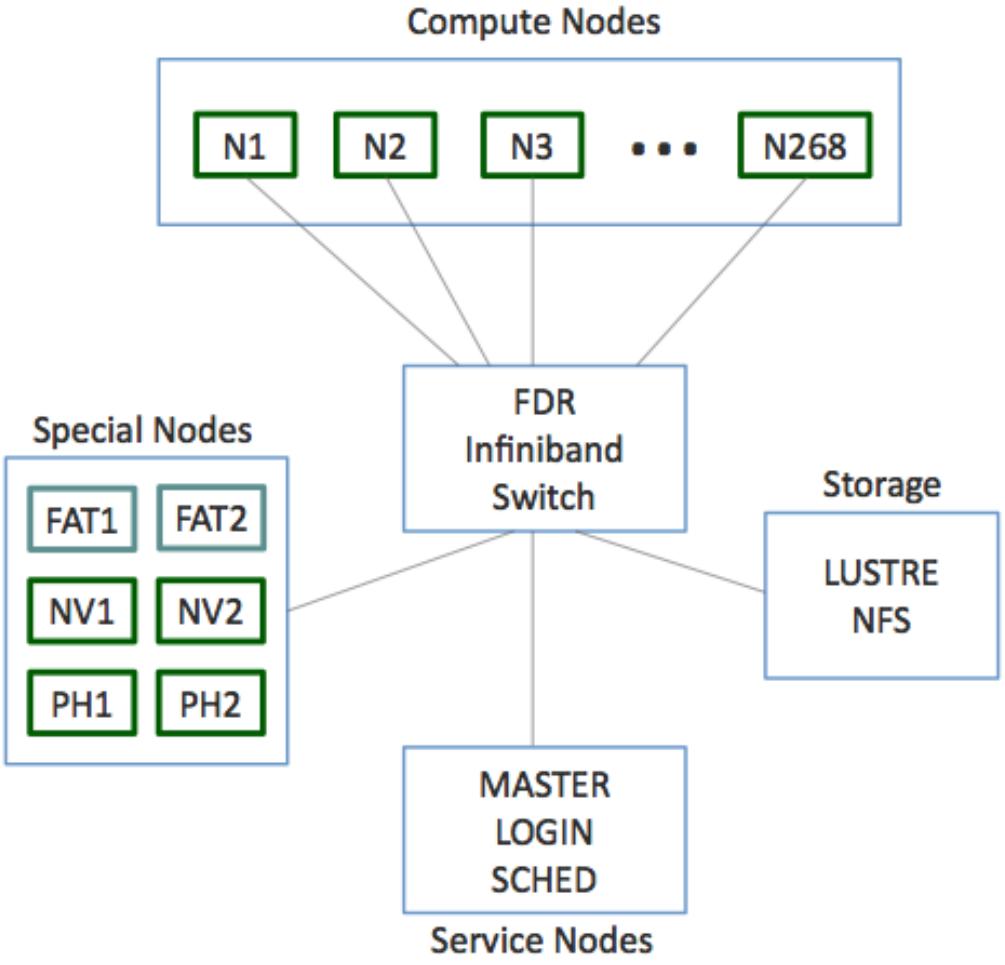

HPC Equipment

268 compute nodes:

- 2 x Intel Xeon E5-2680 v3 at 2.5 GHz
- 2 x 12 cores
- 128 GB DDR4 RAM
- AVX 2.0 (16 double precision floating point operations per clock cycle per core)

All compute nodes run CentOS Linux 6.6

Supercomputing Equipment

4 special compute nodes with GPUs:

- Same CPU as normal compute nodes
- 2 nodes with 2 NVIDIA K40 GPUs:
  - 2880 CUDA cores
  - 12 GB of memory
  - 1.43 TFLOPS DP peak performance
- 2 nodes with 2 Intel Xeon Phi coprocessors
  - 61 cores
  - 16 GB of memory
  - 1.208 TFLOPS DP peak performance

Supercomputing Equipment

3 service nodes:

- Master node
  Cluster monitoring and software deployment
- Login node
  User tools for code compiling, job execution and monitoring, etc.
- Job management node
  SLURM resource management

All servers run RedHat Linux 6.6

Fast data transfer network (computation and parallel filesystem):

Mellanox FDR Infiniband SX6518 director switch

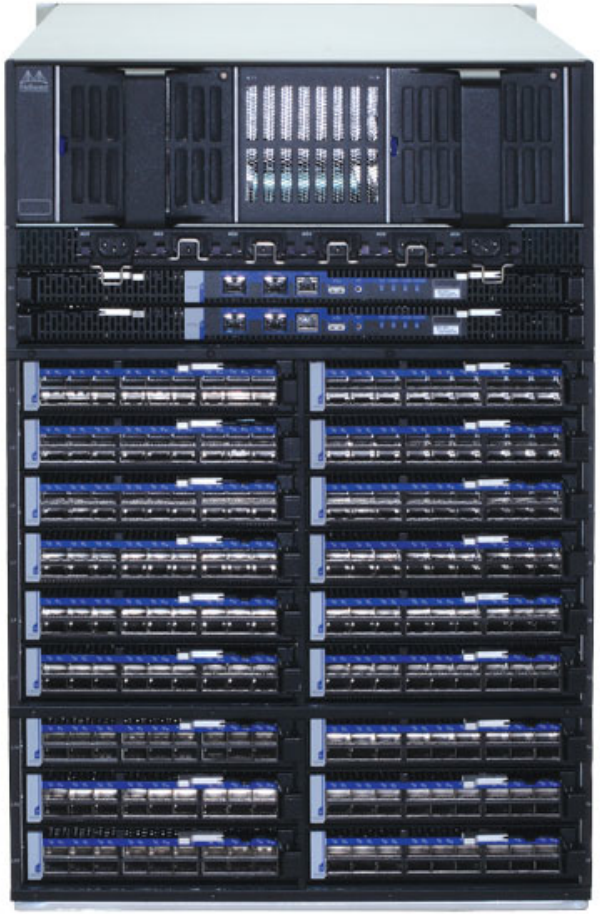

- Up to 324 FDR IB ports: 56 Gb/s full bidirectional bandwidth with sub 1 μs port latency.

Storage servers:

LUSTRE parallel distributed filesystem:

- 6 object storage servers (OSS)
  2 OSS share a 352 TB hardware
  RAID 6 object storage target (OST)

  1056 TB raw storage capacity

- 2 metadata servers (MDS) sharing
  a 32 TB hardware RAID 6 metadata
  target (MDT).

Storage servers /home:

NFS:

- 200 TB hardware RAID 6 cabinet
- XFS filesystem

Communications Infrastructure

Buap Network
Link @ 1 Gb/s (CUDI Network)
Link @ 5 Gb/s and 10 Gb/s (working on progress)

Installation and Configuration of LNS, from January to February 2015 at Data Center's of BUAP

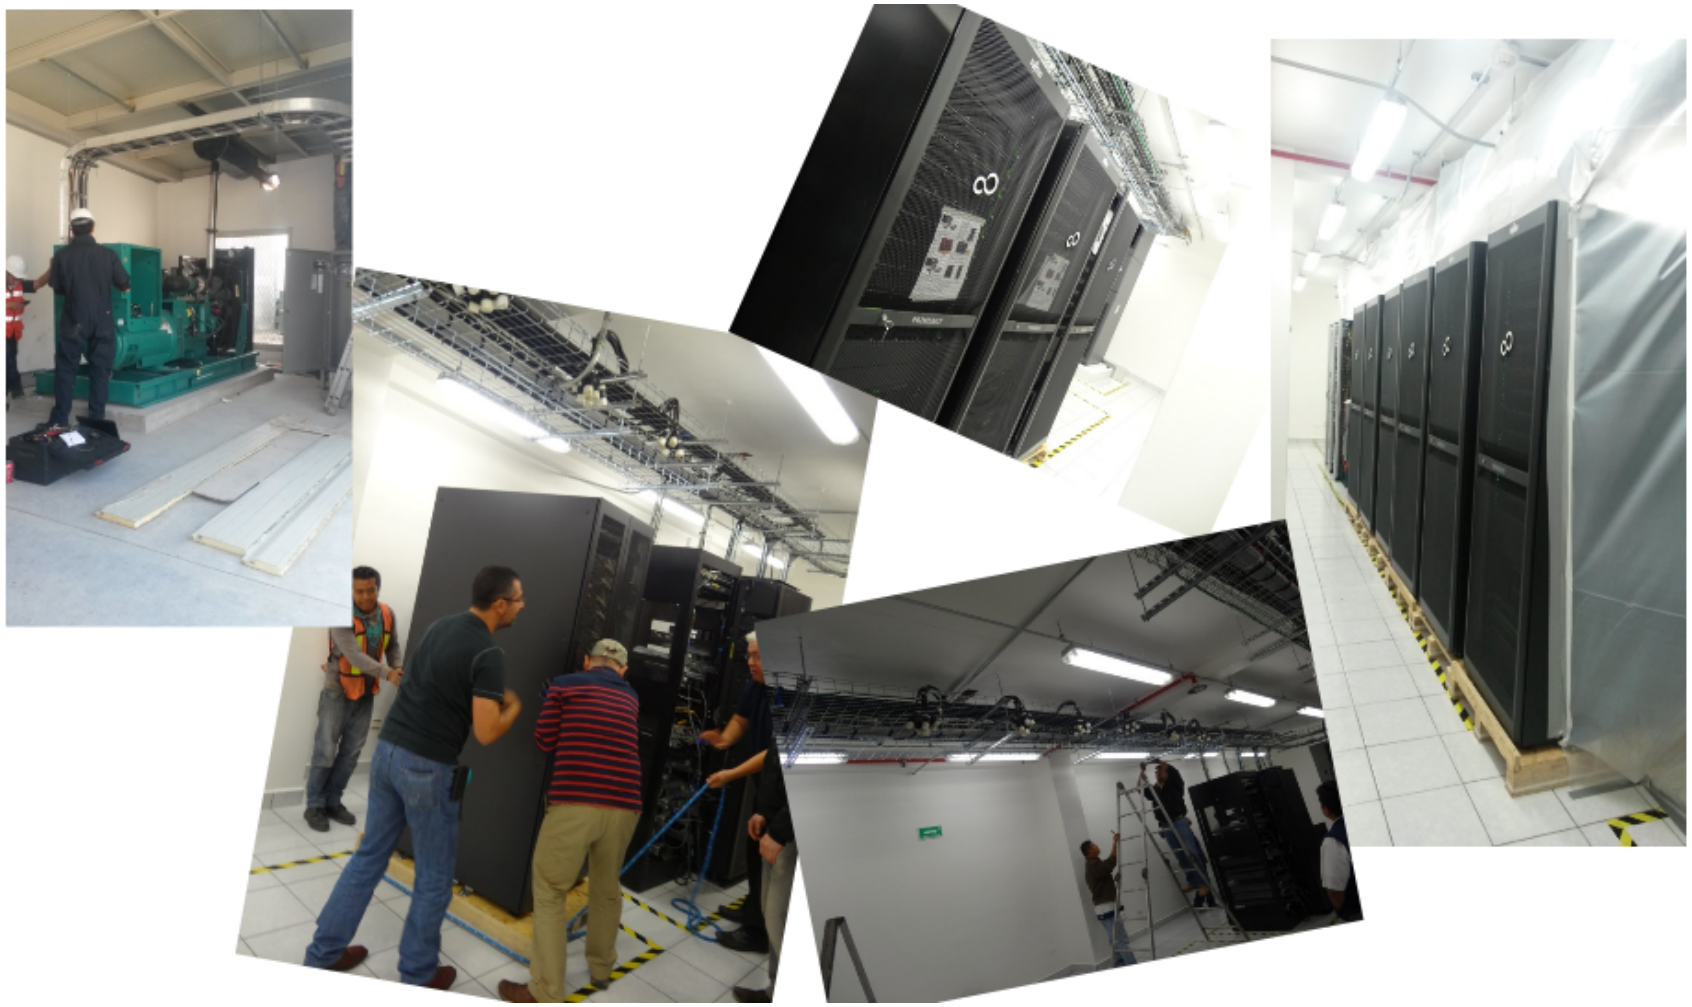

Data Center at LNS (AUGUST- 2015…)

Services of LNS

Academic and Scientific

BUAP-INAOE-UDLA

National Academic community

Industrial/Commercial

Services of LNS

First Stage (March-August 2015)
- 204 compute nodes
- 153.4 TFLOPS performance (Top500)
- Green 500
- 27 Projects running (BUAP-INAOE-UDLA consortium)
- 60 Researcher accounts

Services of LNS

Number of research projects by scientific field:

Condensed Matter Physics and Chemistry: 15
Biology and Physiology: 3
Mathematical Physics: 1
High Energy Physics: 6
Computational Science: 1
Plastic and Visual Arts: 1

Services of LNS

An important effort was made to provide a balanced set of commercial and free HPC applications:

Number of research groups using HPC applications in condensed matter physics and chemistry:

Gaussian: 7
CRYSTAL: 2
VASP: 3
TeraChem: 3
Molpro: 2

Abinit: 4
NWChem: 2
SIESTA: 1
ORCA: 2
Quantum Espresso: 3

Services of LNS

Biophysics and Physiology:

Sybyl: 2
NAMD: 2
Gromacs: 1
GULP: 2

Plastic and Visual Arts:

BLENDER: 1

Services of LNS

Second Stage (August 2015) 1st **Consolidation**

- 64 compute nodes upgrade

This upgrade increases the computing capacity, of 204+64= 268 (6432 cores) and position the cluster as one of the 500 most powerful supercomputing clusters in the world.

- Test of High Performance LINPACK in 200 Tflops

Services of LNS

RESULTS

- **7 National projects**
- **More of 20 consortium projects**

Use of LNS from December 2015

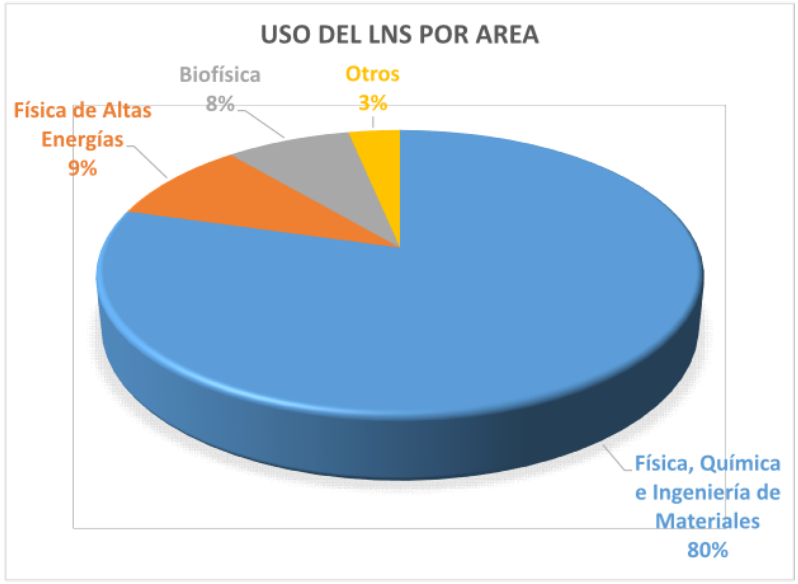

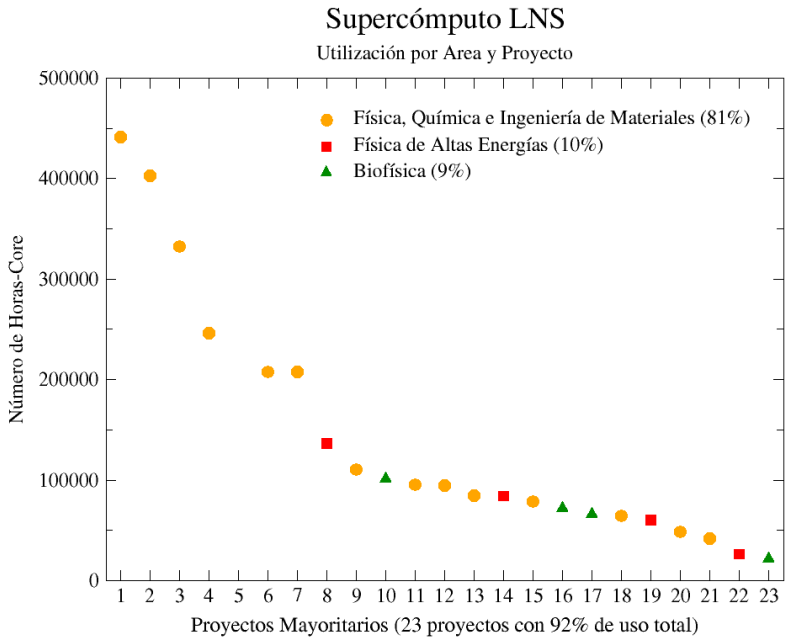

HPC on Medical Physics

Geant4

Applications

High energy and nuclear physics detectors

- ATLAS, CMS, HARP and LHCb at CERN and BaBar at SLAC

Accelerator and shielding

- Linacs for medical use

Medicine

- Radiotherapy
  - photon, proton and light ion beams
  - brachytherapy
  - boron and gadolinium neutron capture therapy
- Simulation of scanners
  - PET & SPECT with GATE (Geant4 Application for Tomographic Emission)

Space

- Satellites
  - effect of space environment on components (especially electronics)
  - shielding of instruments
  - charging effects
- Space environment
  - cosmic ray cut-offs
- Astronauts
  - dose estimates

Simulation of s
PET scanner using G
(Courtesy o
OpenGATE collabora

HPC on Medical Physics

GEANT4: Multithreading, run over cpus (E5-2680v3, 12 cores, 24 cores for node).
Posibility run over coprocessors, Intel Xeon Phi 7120p

GATE (Geant4 Application for Tomography Emission)

Created for Positron Emission Tomography (PET) and Single Photon Emission Computed Tomography (SPECT)
Open Source

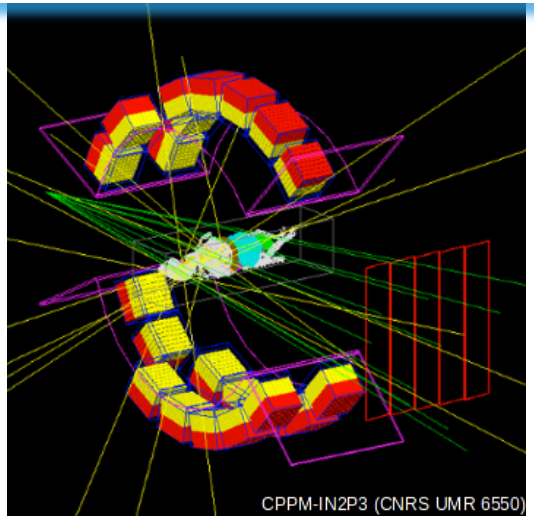

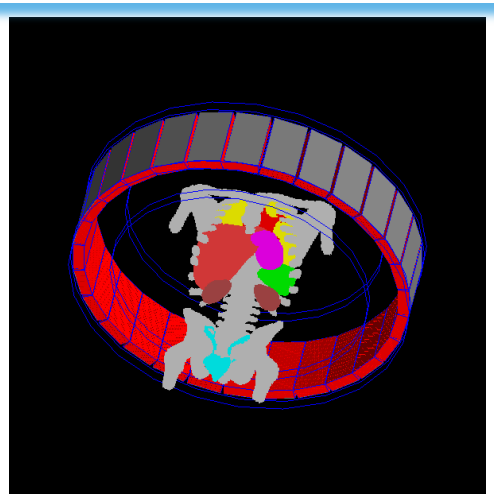

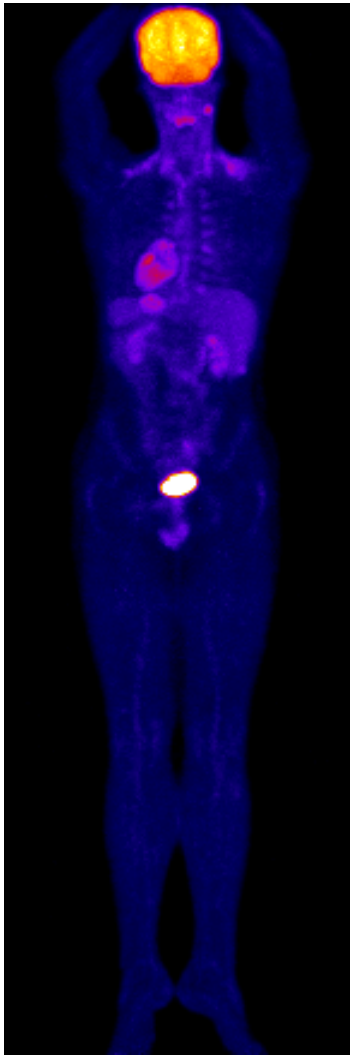

GATE (Geant4 Application for Tomography Emission)

Soon version for GPUs

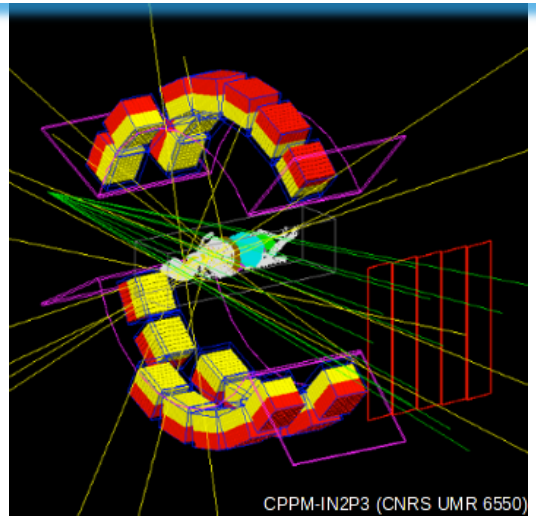

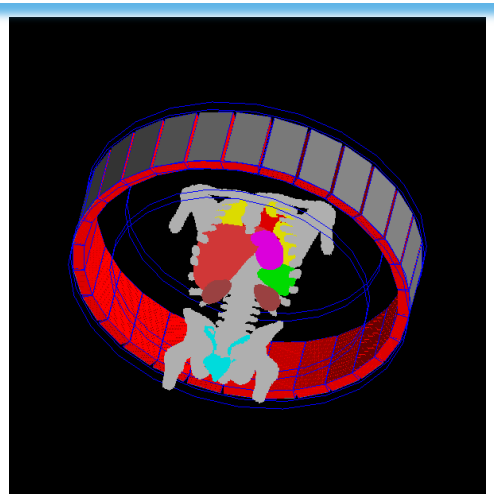

TOPAS (TOol for PArticle Simulation)

TOPAS wraps and extends the Geant4 Simulation Toolkit to make advanced Monte Carlo simulation of all forms of radiotherapy easier to use for medical physicists.

Not open source application!!!

HPC for High Energy Physics

CMSSW at LNS

HPC for High Energy Physics

CMS Software

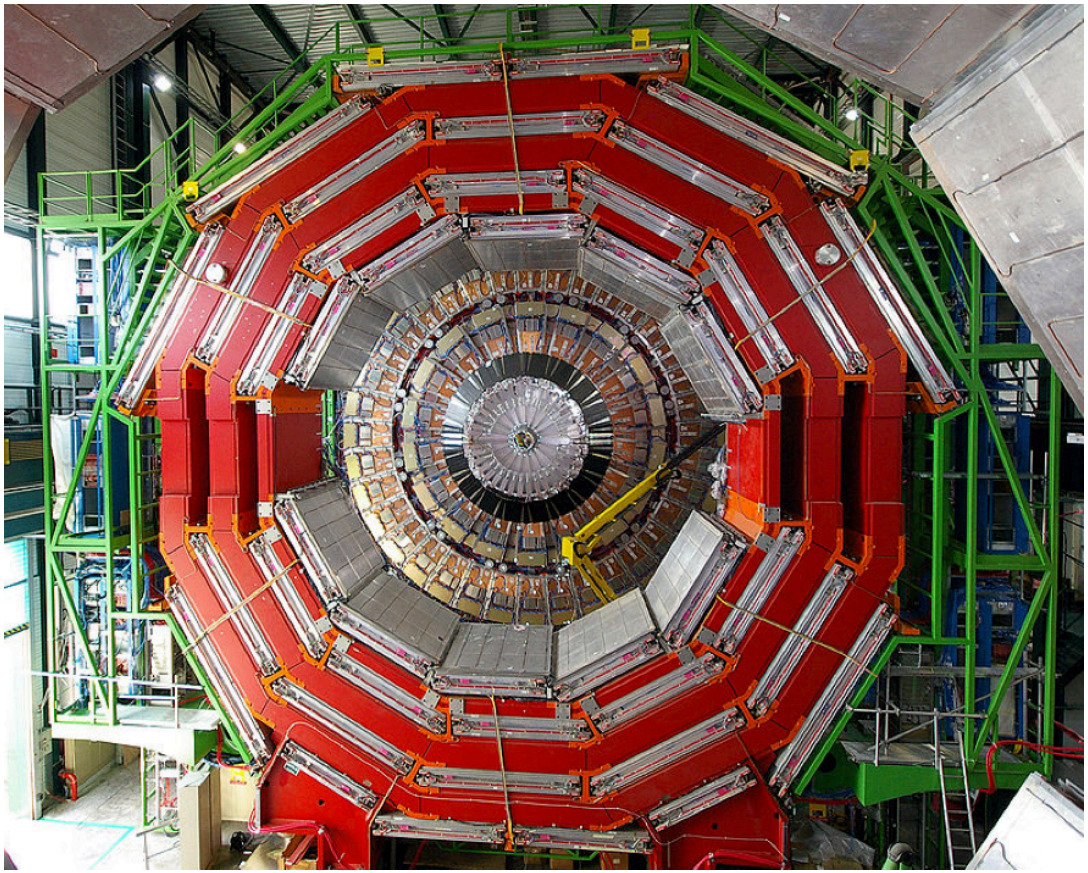

- CMSSW is the main software used by the CMS detector.
- There are many different releases of CMSSW
- On the LNS are 3 different version of CMSSW installed.

HPC for High Energy Physics

Versions of CMSSW at LNS

- 7_5_0_pre5, for the upgrade of the RPC´s on the CMS.
- 7_6_3_patch2 , for Charged Higgs.
- 8_0_7 for monitoring the RPC´s.

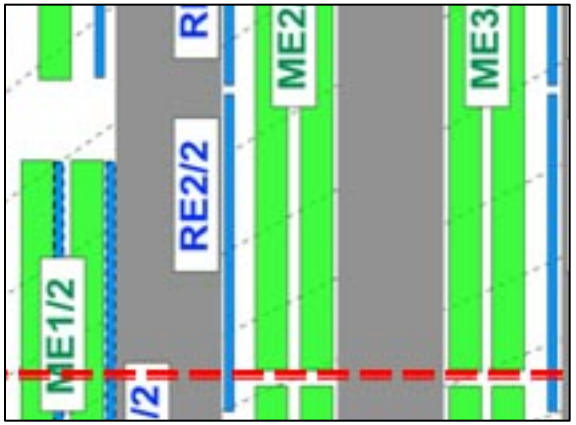

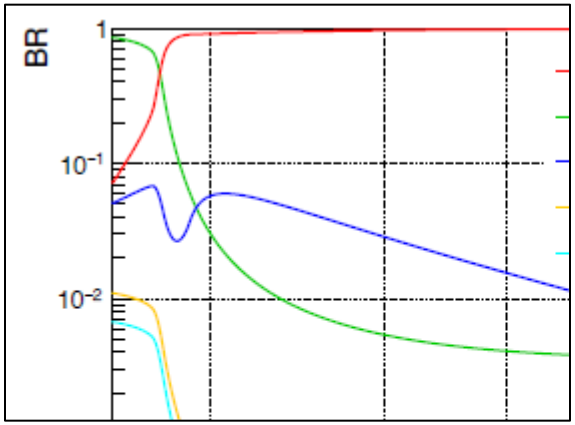

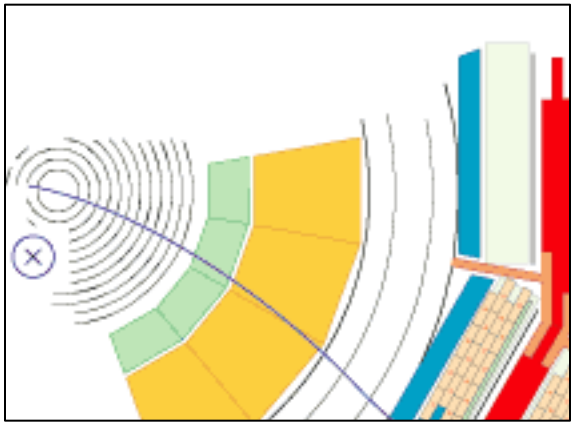

HPC for High Energy Physics

Thesis in progress

3 Ph. D. Thesis in progress…

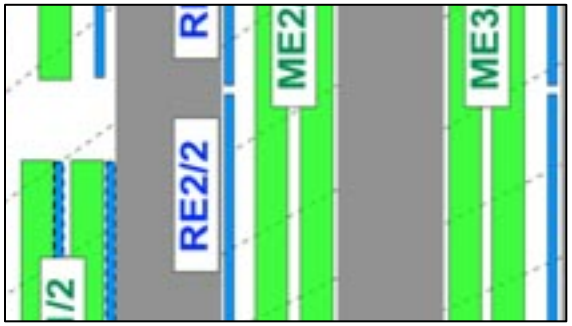

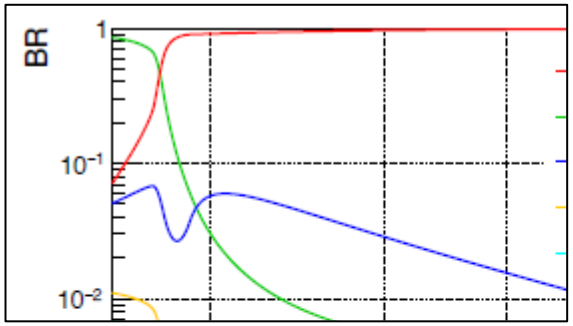

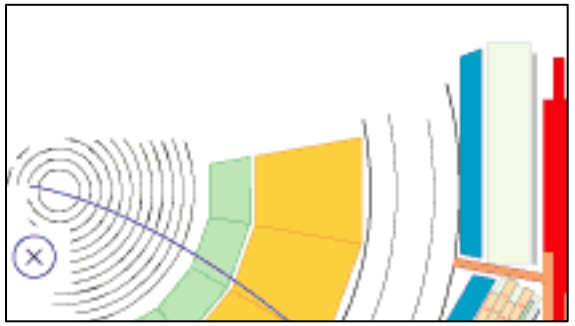

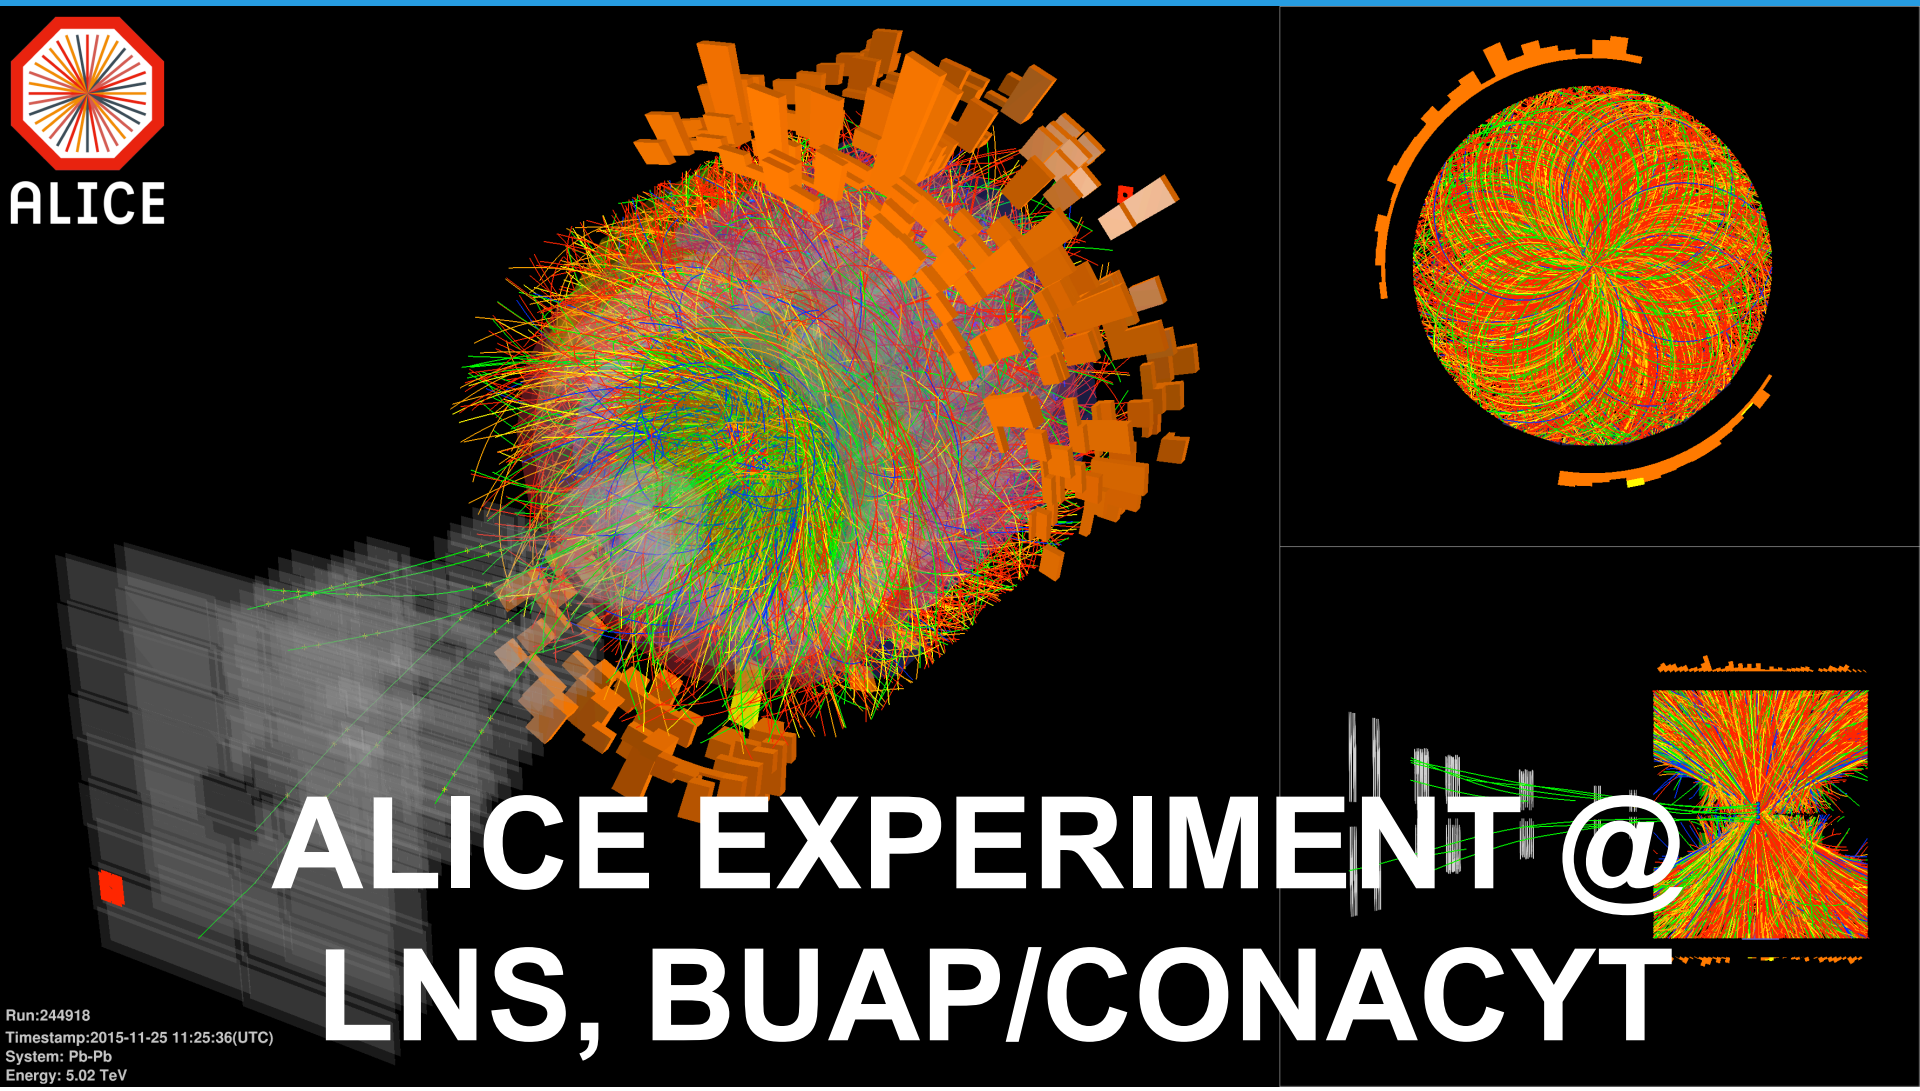

ALICE EXPERIMENT @ LNS, BUAP/CONACYT

A Large Ion Collider Experiment

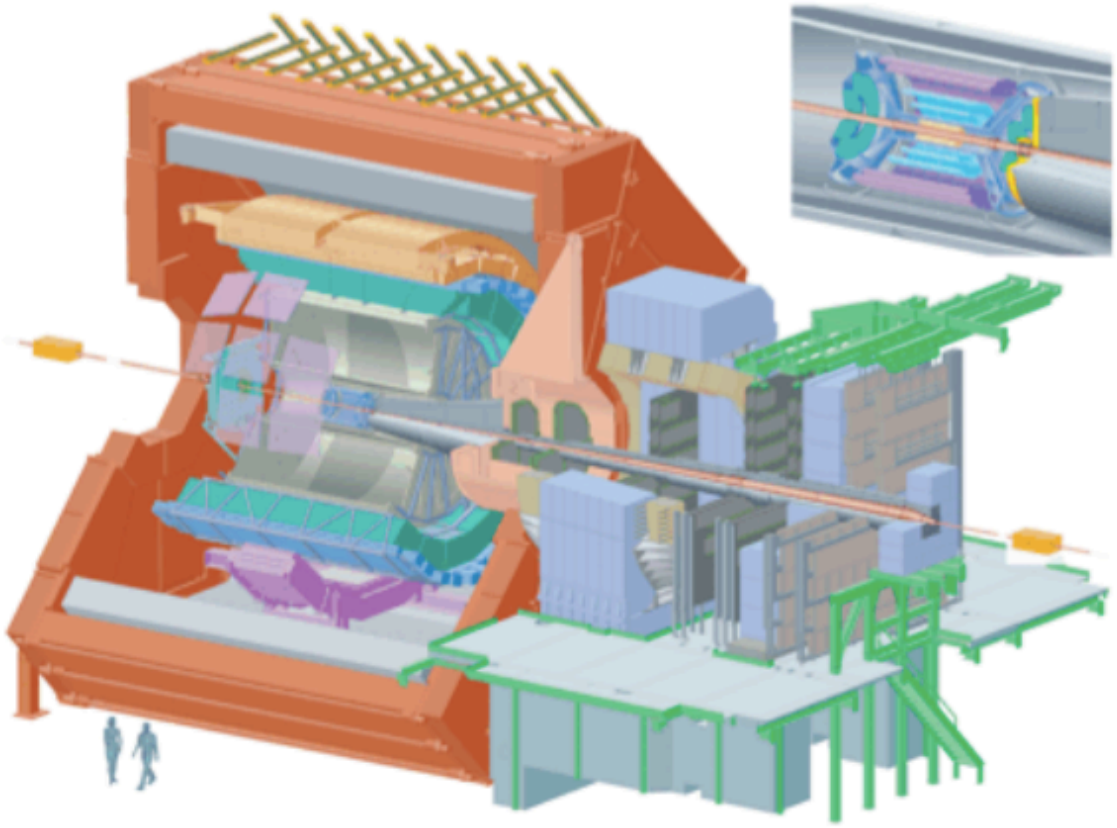

The ALICE Collaboration has built a **dedicated heavy-ion detector** to exploit the unique physics potential of nucleus-nucleus interactions at LHC energies. Our **aim is to study the physics of strongly interacting matter at extreme energy densities, where the formation of a new phase of matter, the quark-gluon plasma**, is expected. The existence of such a phase and its properties are key issues in QCD for the **understanding of confinement and of chiral-symmetry restoration.** For this purpose, we are carrying out a comprehensive study of the hadrons, electrons, muons and photons produced in the collision of heavy nuclei. **ALICE is also studying proton-proton collisions both as a comparison with lead-lead collisions and in physics areas where ALICE is competitive** with other LHC experiments.

INSTALLED SOFTWARE

Since April 2015, ALICE's software is running @ LNS

```
curl build-essential gfortran subversion cmake libmysqlclient-dev xorg-
dev libglu1-mesa-dev libfftw3-dev libssl-dev libxml2-dev libtool
automake git unzip libcgal-dev

clang-3.4

git-new-workdir
```

ROOT, GEANT 3, GEANT 4, ALIROOT, FASTJET, STARLIGHT, PYTHIA, CORSIKA, EPOS, FLUKA

MAIN ACTIVITIES OF ALICE @ LNS

- PHYSICS PERFORMANCE OF FAST INTERACTION TRIGGER
- PHYSICS PERFORMANCE OF COSMIC RAY TRIGGER
- MONTE CARLO STUDIES OF ATMOSPHERIC MUONS WITH UNDERGROUND DETECTORS
- MONTE CARLO STUDIES OF DIFFRACTIVE EVENTS IN PP COLLISIONS

PRODUCTS

- **Study of cosmic ray events with high muon multiplicity using the ALICE detector at the CERN Large Hadron Collider**, JCAP 1601 (2016) no.01, 032 arXiv:1507.07577
- **Study of high muon multiplicity cosmic ray events with ALICE at the CERN Large Hadron Collider, ICRC 2015 @ The Hague, The Netherlands**
- **Performance of Fast Interaction Trigger for ALICE upgrade,** Quark Matter 2015 - XXV International Conference on Ultrarelativistic Nucleus-Nucleus Collisions

Thesis in progress

- mu+/mu- cosmic charge ratio measurement for near vertical muons, Master thesis (Emma Gonzalez)
- Selection of diffractive events in p-p collisions, PhD. Thesis (Abraham Villatoro)
- Study of photo-production of dimuons , Bach. Thesis (Tania Martinez)
- Characterization of muon bundles, Bach. Thesis (Hebert Rodrigo Mojica)
- The LNS as ALICE-GRID site, Bach. Thesis (Y. Bañuelos)

  - The next step is to study the LNS features to be incorporated as an ALICE-Grid site.

HPC for High Energy Physics: Auger

- Offline framework (**Offline v2r9p1)**
- CDAS
- Corsika v75000 (For shower simulations for energies from 10^{17}-10^{18}eV)
- CONEX v2r4.37(For FD shower simulations for energies from 10^{17}-10^{18}eV)
- GEANT4 (Simulation on Detector Response)
- ROOT
- ADST (For analysis)

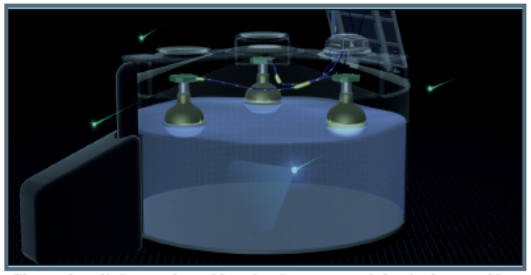

Cherenkov light produced by air-shower particles is detected by three photomultiplier tubes, which view the water volume

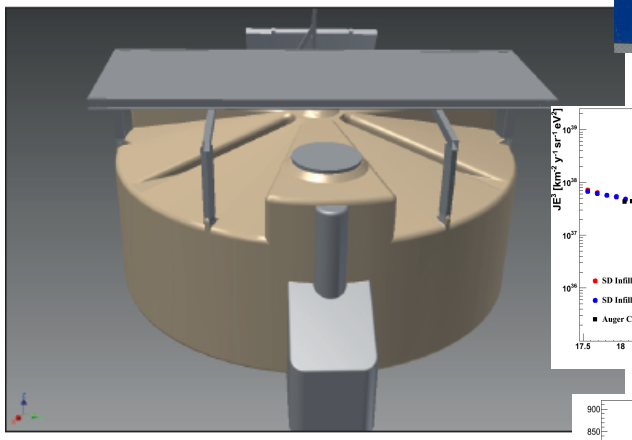

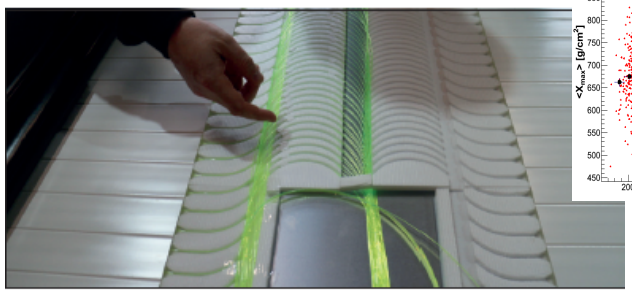

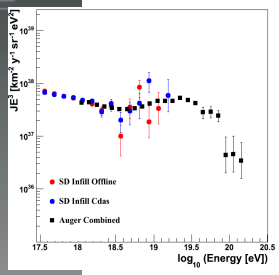

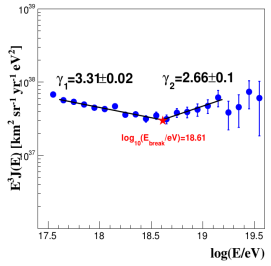

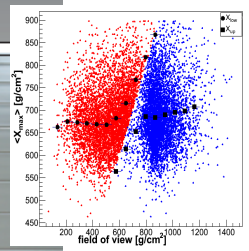

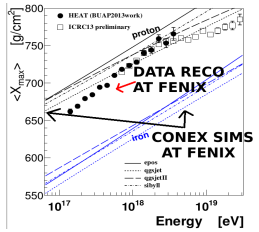

HPC for High Energy Physics: Auger

Thesis in progress

- M. C. Elsa Alejandra Parra Flores, estudiante de doctorado en física de la FCFM-BUAP
- Irving Gabriel Ocampo, estudiante de Ingeniería en Computación de FCC-BUAP

HPC for High Energy Physics: HAWC

- AERI (**Last version)**
- FLUKA
- Corsika v74100 (For shower simulations for energies around 100 TEV)
- GEANT4 (Simulation on Detector Response)
- ROOT, XCDF
- HAWCSIM
- SEDEB
- MAPS GENERATION
- OFFLINE RECO
- AROUND 100 TB OF HAWC DATA

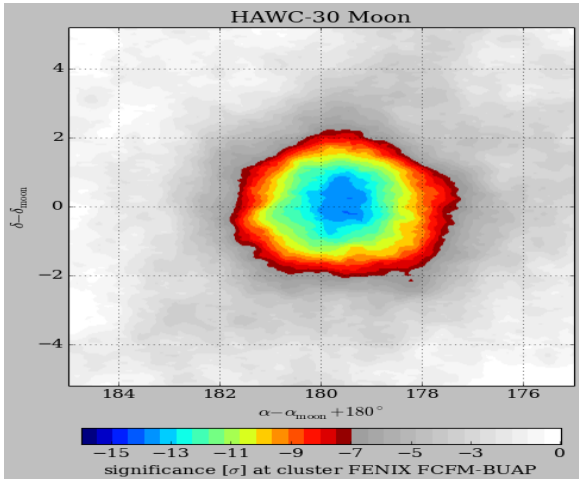

HPC for High Energy Physics: HAWC

Thesis in progress

- M. C. Céderik León de León Acuña, estudiante de doctorado en física de la FCFM-M. C. Alma Roberta Morales, estudiante doctorado en la FCFM.
- M.C. Alan Gilberto Chávez Meza estudiante de doctorado en física del IFM-UMSNH.
- Francisco Tapia Vázquez, estudiante de la licenciatura de la FCFM-UMSNH
- Giovanni Rangel Cortés, estudiante de la licenciatura de la FCFM-UMSNH.

CMS Tier 2 node required capabilities

LNS - T2 : CMS

REQUIREMENTS CMS

At least 788 Intel Xeon E5-2680v2 of processing power

At least 810 TB of storage:

- 30 TB of Stage-Out Space
- 250 TB of Group Space (125 TB per group) 200 TB of Central Space
- 170 TB of Local Space
- 160 TB of User Space (~40 Users of 4 TB each)

Availability 24x7

LNS CONFIG FOR CMS

- ❑ **20 nodes, 480 cores (960 Threads) Intel Xeon E5-2680 v3,** ✓
- ❑ **Storage 300 TB (Shared with ALICE)**✓
- ❑ **NIBA Networking 10 Gb/s** ✓
- ❑ **Infiniband FDR 56 Gb/s** ✓
- ❑ **5.3/2.6(HT) per core, DDR4, NEW 5.3 GB (HT)**✓
- ❑ **A**vailability 24x7 ✓

ALICE Tier 2 node required capabilities

LNS - T2 : ALICE

REQUIREMENTS ALICE

- **~ 1000 cores Intel Xeon E5-2680 v2 in HT**
- **Storage 1TB/4 cores (250 TB / 1000 cores)**
- **Networking (I/0) 10 Gb/s**
- **High-speed Infiniband networking close to 40 Gb/s**
- **Worker nodes min 4 GB RAM (DDR3) per core**
- **Jobs require max 10 GB of local scratch.**

LNS CONFIG. FOR ALICE

- ❑ **20 nodes, 480 cores Intel Xeon E5-2680 v3, 960 cores in HT** ✓
- ❑ **Storage 300 TB (Shared with ALICE) new can update to Tape robot storage** ✓
- ❑ **NIBA Networking 10 Gb/s** ✓
- ❑ **Infiniband FDR 56 Gb/s** ✓
- ❑ **5.3/2.6(HT) per core, DDR4** ✓
- ❑ **1XHDD 256 GB 6Gbps** ✓

SUMMARY

- LNS is already a Consolidated Computer Center
- LNS can offer Certificate Services for Industry and Scientific Research.
- LNS has began to configure GRID services for CMS, ALICE, with 20 nodes for each project and 300 TB of dedicated storage.

Thank you for your attention