

UNAM T1 for ALICE

ICN-DGTIC / UNAM

Luciano Díaz Guy Paic Lukas Nellen Eduardo Murrieta
Fabián Romo J.Luis Gordillo Arión Pérez
november 2015

Antecedent


- 2007 First ALICE site in Latin America

@ICN-UNAM

32 Xeon 2.4GHz -32bits

1.5 GB memory/core

30 Mbps Internet2



Antecedent

- 2010 cluster upgrade

@ICN-UNAM

96 Xeon 2.4GHz -64bits

2 GB memory/core

60 Mbps Internet2



@ICN the resources are shared:

- ALICE
- AUGER
- EELA
- BIOMED
- Local users
 - HAWC (since 2012)
 - Several applications: mathematica, maple, MPI code, gaussian, ...

Proposal

2011

Grid of the Americas workshop concluded with:

Proposal to install a Tier-1 for ALICE in México

A collaboration project between:

ICN-UNAM and UNAM - DGSCA (now DGTIC) / Supercomputing department

with ALICE-CERN support

Location:

Clusters room @ DGTIC-UNAM

- | Electricity
- | Network connectivity
- | UPS
- | Cooling
- | Technical support

Requirements

- 1K cores or HEPSPROC06 equivalent (~26500 HS06)
- 1 PB disk storage
- 1 PB tape storage (maybe 10PB for Run3)
- 1Gbps link (now 10 Gbps)

2 steps process:

1. Start a Tier-2

- Accumulate good metrics: availability & reliability
- Cover the requirements

2. Become Tier-1

July 2014

We joined to ALICE production as T2:
SUPERCOMPUTO-UNAM

November 2014 MoU signed:
CERN-UNAM

SUPERCOMPUTO-UNAM

Cluster canek

- 32 servers each one:
 - 16 cores with HT enabled = 32 cores -1
- Total:
- 992 cores Intel Xeon E5-2670v2 2.5 GHz
 - 4 GB memory / core
 - 2 HDD 1TB in stripe (RAID 0)
 - 1Gbps interfaces
-
- ~11456 HS06 (based on <http://w3.hepix.org>)


SUPERCOMPUTO-UNAM


Storage

570 TB disk storage for ALICE with EOS

 1 redirector

 5 servers with 120TB c/u

 File system: XFS / RAID 6

 10Gbps interfaces

 One of the firsts storages with EOS

47. SPbSU - SE	ALICE::SPbSU::SE	38.14 TB	22.38 TB	15.76 TB	58.68%	547,753	FILE	-	-	-	-			Oct ...	Last...	23.10.2015 08:29	2	10	65.48%	
48. Strasbourg_IRES - SE	ALICE::Strasbourg_IRES::SE	116.9 TB	105.7 TB	11.22 TB	90.4%	2,075,392	FILE	81.07 TB	74.29 TB	6.784 TB	91.63%	v3.2.6				24.10.2015 04:29	12	0	0	
49. Subatech - EOS	ALICE::Subatech::EOS	388 TB	246.5 TB	141.5 TB	63.54%	9,407,732	FILE	388 TB	278 TB	110 TB	71.65%	3.3.6	0.3.118			24.10.2015 04:29	12	0	0	
50. SUT - SE	ALICE::SUT::SE	99 TB	50.26 TB	48.74 TB	50.77%	61,988	FILE	99.1 TB	45.11 TB	53.99 TB	45.52%	v3.1.0				24.10.2015 04:30	12	0	0	
51. Torino - SE	ALICE::Torino::SE	589 TB	457.6 TB	131.4 TB	77.68%	12,111,660	FILE	1.148 PB	929.6 TB	246.5 TB	79.04%	v3.3.4				24.10.2015 04:30	12	0	0.298%	
52. Trieste - SE	ALICE::Trieste::SE	342 TB	25.79 TB	316.2 TB	7.539%	528,499	FILE	342 TB	251.2 TB	90.82 TB	73.44%	v3.3.4				24.10.2015 04:31	12	0	0	
53. Troitsk - SE	ALICE::Troitsk::SE	54.57 TB	71.33 TB	-	130.7%	1,501,206	FILE	112.7 TB	68.82 TB	43.86 TB	61.07%	v4.2.3				24.10.2015 04:32	8	4	27.38%	
54. Trujillo - SE	ALICE::Trujillo::SE	30.92 TB	19 TB	11.92 TB	61.43%	339,955	FILE	185.5 TB	151.5 TB	33.99 TB	81.68%	v4.1.3				24.10.2015 04:32	12	0	0	
55. UNAM_T1 - EOS	ALICE::UNAM_T1::EOS	456 TB	131 TB	325 TB	28.73%	8,019,925	FILE	456 TB	185 TB	271 TB	40.57%	3.3.6	0.3.35			24.10.2015 04:32	12	0	0	
56. ZA_CHPC - SE	ALICE::ZA_CHPC::SE	107.7 TB	11.38 TB	96.32 TB	10.57%	406,944	FILE	-	-	-	-			Oct ...	1	12.04.2015 14:34	0	12	100%	
Total		26.03 PB	23.04 PB	5.592 PB		674,306,100		36.38 PB	27.43 PB	8.953 PB										

What is this about?

Tape storage elements

SUPERCOMPUTO-UNAM

Additional servers to support T2 operation:

ALICE VoBox

CREAM - CE with Torque / PBS

Cache WEB (squid) - Software ALICE / CVMFS

Virtualization server:

- SE-DPM - Just for EGI monitoring
- SiteBDII
- APEL - Accounting

For internal services:

Server with NTP, Gateway, DHCP, DNSmasq

SUPERCOMPUTO-UNAM

Internal network

- Cluster switch with 40Gbps capability
 - Nodes @ 1Gbps
 - Storage @ 10Gbps

Campus network

- UNAM core network @ 10Gbps

Internet2 connectivity

- 1 Gbps link



MonALISA Repository for ALICE

ALICE Repository

- ALICE Repository
- Google Map
- Shifter's dashboard
- Run Condition Table
- Production Overview
- Production info
- Job Information
- SE Information
- Services
- Network Traffic
- FTD Transfers
- CAF Monitoring
- SHUTTLE
- Build system
- HepSpec
- Dynamic charts

close all

This page: bookmark, URL

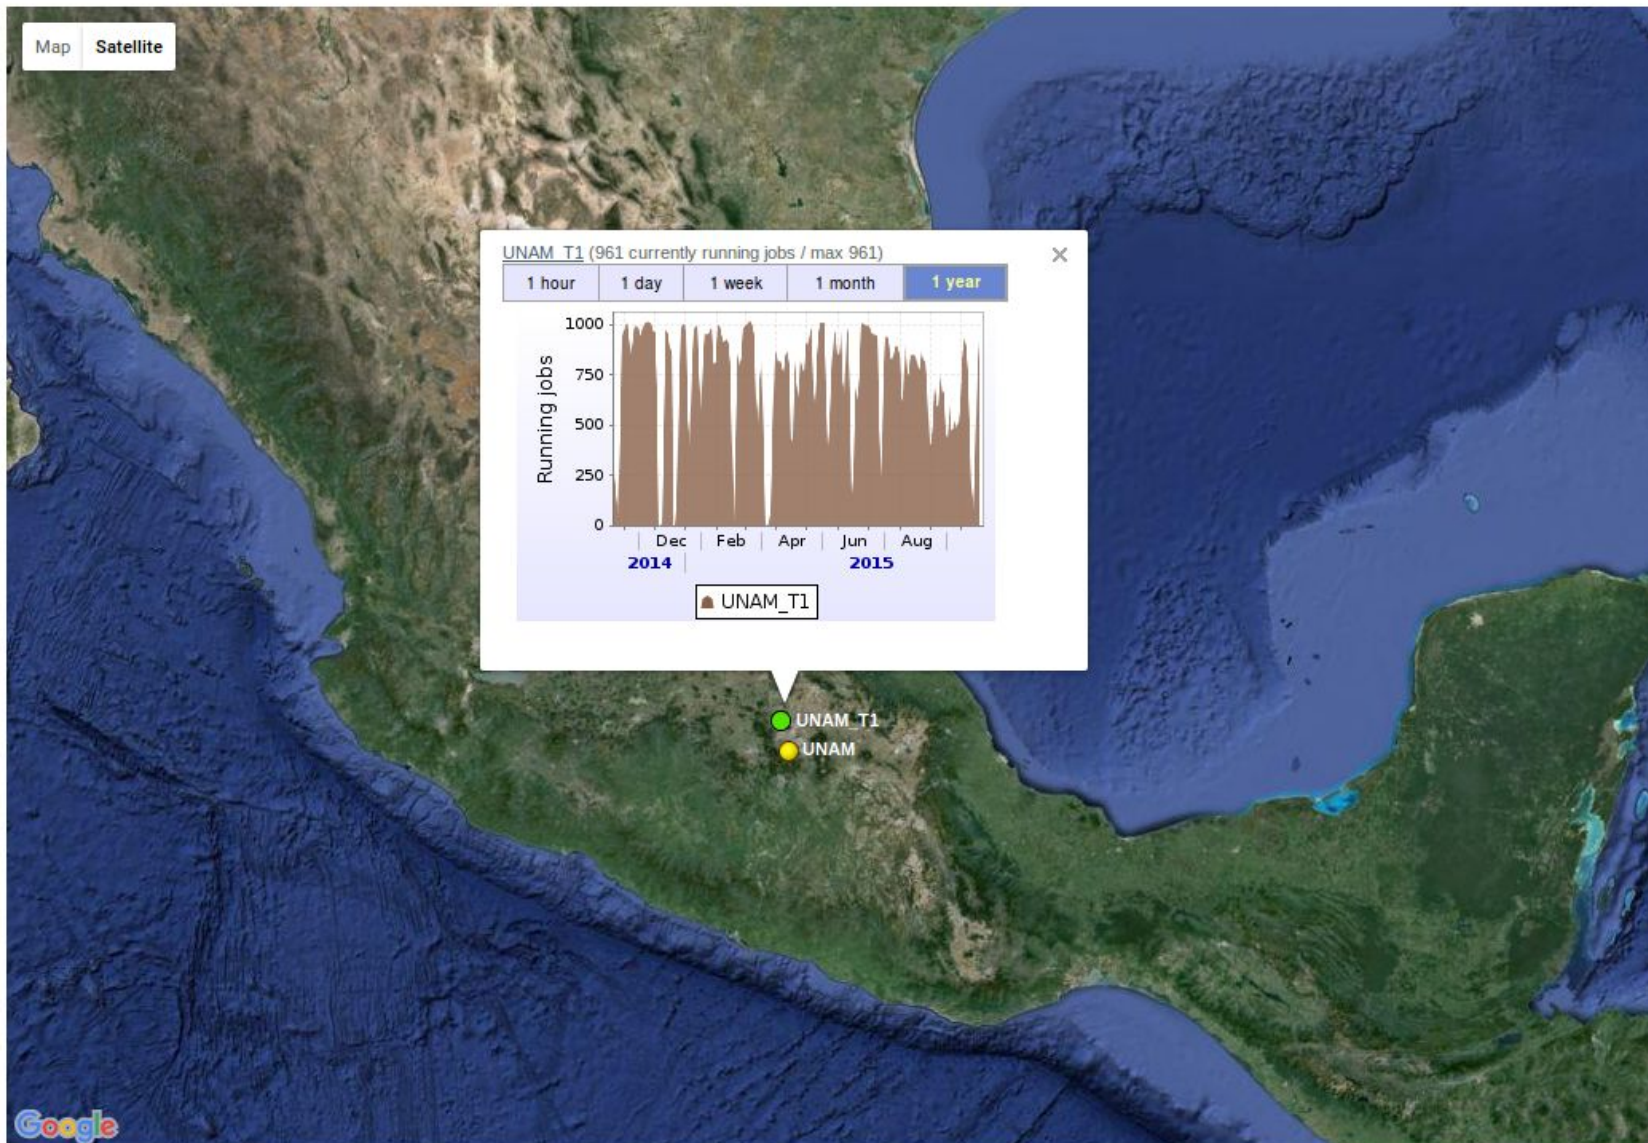
Running jobs trend




Running jobs trend

- 24h
- 12h
- 6h
- 1h

(click arrows for detailed view)



EGI in LA for LHC



- Site views
 - Summary plots
 - Job states
 - Jobs per site
 - Jobs per site table
 - Resource usage
- User views
- Task queue
- Job timings
- Memory profiles
- Federation views
 - Resources
 - Jobs
- JobAgents
- SE Information
- Services
- Network Traffic
- FTD Transfers
- CAF Monitoring
- SHUTTLE
- Build system
- HepSpec
- Dynamic charts

Armenia - Brasil - CERN - China - Czech Republic - Germany - Greece - HLT - Hungary - IN2P3 - India - Indonesia - INFN - Japan - Mexico

Nordic Countries Other Pakistan Poland RDIG Republic of Korea Romania Slovakia South Africa Spain Thailand The Netherlands UK


Ukraine US

Function :

Interval selection: - choose - or « 2014-06-01 23:00 - 2015-10-24 23:00 » Plot

What is this about?

DONE jobs



This page: bookmark, URL

Running jobs trend



Plans next year

CPU

- ∅ Add CPU capacity > 26500 HS06
 - | Proposal to buy 32 servers with CPU E5-2680v3

Disc storage

- | Add 500 TB disk storage to current EOS
 - | Servers already at DGTIC we are waiting to be delivered for ALICE use

Network

- | 10 Gbps link already requested to authorities:
 - | Work in progress from UNAM<->National network backbone (CUDI)
 - | Looking for 1 to 10 Gbps link upgrade to Internet2

Tape storage

- | Testing 8TB archive disks to propose an alternative to tape storage system based on ZFS and EOS
- | Waiting for the authorization of new personnel.