# Grid site as a tool for data processing and data analysis. RO-11-NIPNE computing performance evolution

TEODOR IVANOAICA

*HORIA HULUBEI NATIONAL INSTITUTE FOR R&D IN PHYSICS AND NUCLEAR ENGINEERING (IFIN-HH)*

*POLITEHNICA UNIVERSITY OF BUCHAREST, 313 SPLAIUL INDEPENDENTEI, BUCURESTI, ROMANIA*

# Grid site as a tool for data processing and data analysis

Summary :

- What is GRID Computing.
- Needs and performance evolution.
- RO-11-NIPNE evolution.
- Upgrades and future implementations.

# Grid site as a tool for data processing and data analysis

## What is GRID Computing.

- **Grid computing** is the collection of computer resources from multiple locations to reach a common goal. – according to: https://en.wikipedia.org/wiki/Grid_computing
- **Grid computing** is a distributed architecture of large numbers of computers connected to solve a complex problem. – according to: http://searchdatacenter.techtarget.com/definition/grid-computing
- A scientist studying proteins logs into a computer and uses an entire network of computers to analyze data. A businessman accesses his company's network through a PDA in order to forecast the future of a particular stock. An Army official accesses and coordinates computer resources on three different military networks to formulate a battle strategy. All of these scenarios have one thing in common: They rely on a concept called **Grid computing**. – according to: http://computer.howstuffworks.com/grid-computing.htm

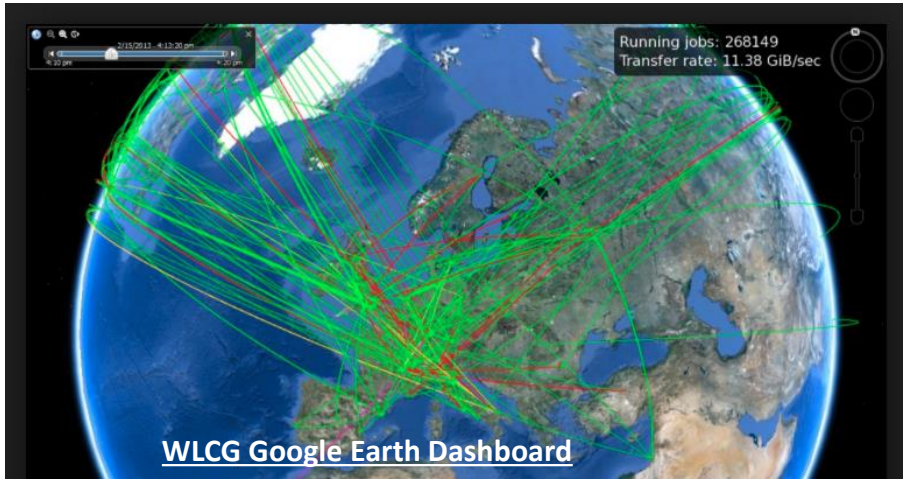

WLCG Google Earth Dashboard

The Worldwide LHC Computing Grid (WLCG) is a global collaboration of computer centres. It was launched in 2002 to provide a resource to store, distribute and analyse the 15 petabytes (15 million gigabytes) of data generated every year by the Large Hadron Collider (LHC).

# Grid site as a tool for data processing and data analysis

❑**The need**: **Approximately 600 million times per second, particles collide within the** Large Hadron Collider (LHC).

Each collision generates particles that often decay in complex ways into even more particles**.**

**Electronic circuits record the passage of each particle through a detector as a series of electronic signals, and send the data to the CERN Data Centre (DC) for digital reconstruction**.

**The digitized summary is recorded as a "collision event". Physicists must shift through the 30 petabytes or so of data produced annually to determine if the collisions have thrown up any interesting physics**

❑**The Solution should offer :**

- ❑ **A very big number of CPU's, storage …. (not feasible all in one place, redundancy back-ups ) – even more problems**
- ❑ **BIG IT Group – working around the clock for the Physicists**

**In 2002 everything turned into grid computing, every computing task is shared between multiple computing centers, structured in tiers.**

**How it works:**

- ❖ **Tier 0 ( CERN, and Wigner in Hungary) is storing the data sets emerging from the Data Acquisition, those centers are doing the initial processing of the data and Tape Storage. The data is than distributed to Tier 1's.**
- ❖ **Tier 1 – Data centers that have Data Storage pledged to each experiment, are offering a big amount of Computing Power ( $10^{th}$ of thousands of cores) that require access to large raw data sets for the data analysis and simulations.**
- ❖ **Tier 2 – Smaller data centers that will have da Data made available by the T-1's, mainly smaller sites that are organized and maintained by collaborating Research Institutes or Universities, supporting users tasks, reconstructions (in some cases), Monte Carlo Simulations etc.**
- ❖ **Tier 3 – Small local clusters , dedicated to an experiment, deployed by a group of users or an individual PC**

**During last years the requirements focus mainly on the storage, as the sites are evolving and add storage for the data, but also all the sites have increased the number of CPU's.**

# Grid site as a tool for data processing and data analysis

The Romanian participation in the LHCb started in 1996, our team, LHCb-Ro, being involved in: design, construction, commissioning, cosmic ray tests and calibration of the calorimeter modules.

During 2006 EGEE certified RO-11-NIPNE (LHCb) grid site with the following resources:

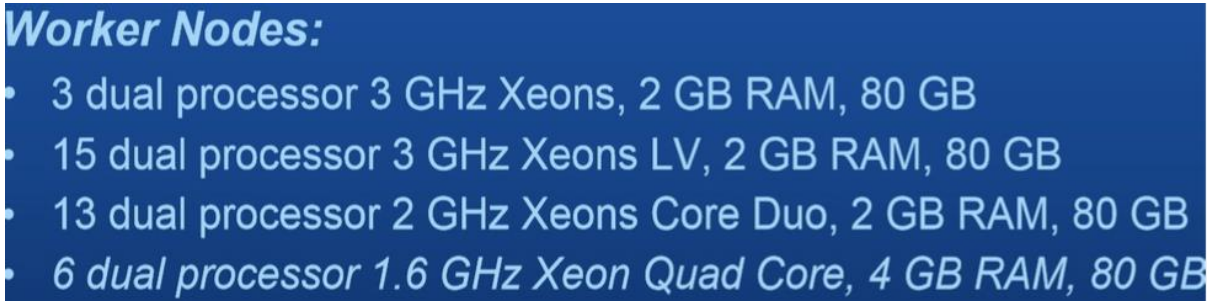


- In 2012 the old hardware 21 servers installed in 2007 were completely replaced by 10 servers Each server 8 cores 16GB ram.
- During 2013 we added 9 worker nodes 16 cores 32 GB ram.
- In february and march **2014** after updating the middleware and installing glexec and other collaboration requirements we added 5 new nodes with 24 cores and 48 GB ram

# Grid site as a tool for data processing and data analysis

Past Tense:

Our group started with :

- Only 28 cores installed
- Only 500GB storage on the SE
- 1 Gbps ethernet switch

Offering a performance of only ~ 233 k lSI2K in 2011

**RO-11-NIPNE Normalised CPU time (kSI2K) by SITE and VO.**

CUSTOM VOs. January 2011 - November 2011.

The following table shows the distribution of Normalised CPU time (kSI2K) grouped by SITE and VO (only information about the **selected VOs** is returned).

| Normalised CPU time [units 1K.SI2K.Hours] by SITE and VO | | | |
|---|---|---|---|
| SITE | lhcb | Total | % |
| RO-11-NIPNE | 233,562 | 233,562 | 100.00% |
| Total | 233,562 | 233,562 | |
| Percentage | 100.00% | | |

Click here for a csv dump of this table

Click here for a EXTENDED csv dump

The information in the previous table is also shown in the following graph.

RO-11-NIPNE Normalised CPU time (kSI2K) by SITE and VO

At that time the resources were moving to a new Data Center (**DFCTI/IFIN**-HH Department of Computational Physics and Information Technologies) solving all the problems that we previously had:

- Cooling issues
- UPS
- Stable Network Connectity (using a direct link to the core router)

# Grid site as a tool for data processing and data analysis

Present:

Offering a performance of only ~ 233 k lSI2K in 2011

We have at this time available:

- 23 Worker nodes, 8 cores , 16 cores
- 2GB ram/core , 16 GB , 32 GB
- Disk space 500 GB /node (even less on the 16 core nodes)
- Storage Element
- Middleware – EMI3.x

Normalised CPU time [units 1K.SI2K.Hours] by SITE and VO

| SITE | lhcb | Total | % |
|---|---|---|---|
| RO-11-NIPNE | 2,159,373 | 2,159,373 | 100.00% |
| Total | 2,159,373 | 2,159,373 | |
| Percentage | 100.00% | | |

Click here for a CSV dump of this table

Click here for a Extended CSV dump of this table

Click here for XML encoded data

First time over 2 000 000 K SI2K

At that time we moved the resources to a new Data Center (**DFCTI/IFIN**-HH Department of Computational Physics and Information Technologies) we solved all the problems that we previously had:

- Cooling issues
- UPS
- Stable Network Connectivty (using a direct link to the core router)

# Grid site as a tool for data processing and data analysis

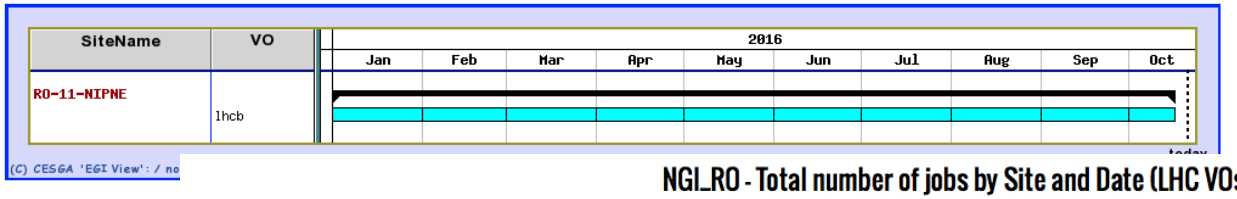


## NGI_RO - Total number of jobs by Site and Date (LHC VOs)

| | Site | Jan 2016 | Feb 2016 | Mar 2016 | Apr 2016 | May 2016 | Jun 2016 | Jul 2016 | Aug 2016 | Sep 2016 | Oct 2016 | Total | Percent |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| NIHAM | | 3,576 | 5,332 | 4,535 | 2,881 | 2,441 | 2,935 | 2,972 | 2,976 | 2,824 | 1,529 | **32,001** | **0.33%** |
| RO-02-NIPNE | | 166,350 | 164,288 | 200,218 | 155,475 | 147,999 | 144,033 | 172,449 | 118,248 | 147,203 | 65,547 | **1,481,810** | **15.44%** |
| RO-07-NIPNE | | 898,792 | 821,963 | 862,033 | 482,826 | 470,675 | 429,444 | 466,506 | 443,322 | 280,003 | 188,948 | **5,344,512** | **55.69%** |
| RO-11-NIPNE | | 7,258 | 10,972 | 13,781 | 10,433 | 7,234 | 6,267 | 7,352 | 10,685 | 8,832 | 8,025 | **90,839** | **0.95%** |
| RO-13-ISS | | 21,226 | 21,679 | 22,474 | 17,389 | 20,409 | 47,516 | 16,148 | 9,260 | 11,984 | 0 | **188,085** | **1.96%** |
| RO-14-ITIM | | 16,430 | 25,825 | 89,509 | 25,277 | 316,078 | 237,883 | 176,520 | 219,089 | 348,269 | 300,757 | **1,755,637** | **18.29%** |
| RO-15-NIPNE | | 3,389 | 4,940 | 4,288 | 5,467 | 5,572 | 2,519 | 3,898 | 3,720 | 0 | 0 | **33,793** | **0.35%** |
| RO-16-UAIC | | 33,093 | 27,443 | 83,182 | 60,838 | 58,927 | 95,914 | 43,990 | 123,969 | 119,883 | 22,351 | **669,590** | **6.98%** |
| **Total** | | **1,150,114** | **1,082,442** | **1,280,020** | **760,586** | **1,029,335** | **966,511** | **889,835** | **931,269** | **918,998** | **587,157** | **9,596,267** | |
| **Percent** | | **11.99%** | **11.28%** | **13.34%** | **7.93%** | **10.73%** | **10.07%** | **9.27%** | **9.70%** | **9.58%** | **6.12%** | | |

1 - 8 of 8 results ‹ 1 › Number of rows per page 30

## NGI_RO - Total number of jobs by Site and Date (Custom VOs)

| Site | Jan 2016 | Feb 2016 | Mar 2016 | Apr 2016 | May 2016 | Jun 2016 | Jul 2016 | Aug 2016 | Sep 2016 | Oct 2016 | Total | Percent |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| RO-07-NIPNE | 4,589 | 9,361 | 13,110 | 21,680 | 8,642 | 7,198 | 6,848 | 9,679 | 5,650 | 1,708 | **88,465** | **41.51%** |
| RO-11-NIPNE | 7,258 | 10,972 | 13,781 | 10,433 | 7,234 | 6,267 | 7,352 | 10,685 | 8,832 | 8,025 | **90,839** | **42.63%** |
| RO-15-NIPNE | 3,389 | 4,940 | 4,288 | 5,467 | 5,572 | 2,519 | 3,898 | 3,720 | 0 | 0 | **33,793** | **15.86%** |
| **Total** | **15,236** | **25,273** | **31,179** | **37,580** | **21,448** | **15,984** | **18,098** | **24,084** | **14,482** | **9,733** | **213,097** | |
| **Percent** | **7.15%** | **11.86%** | **14.63%** | **17.64%** | **10.06%** | **7.50%** | **8.49%** | **11.30%** | **6.80%** | **4.57%** | | |

1 - 3 of 3 results ‹ 1 › Number of rows per page 30

# Grid site as a tool for data processing and data analysis

The site is offering very good performances:

- 42,63% of the entire LHCb jobs that run in Romania are running on RO-11-NIPNE
- RO-07-NIPNE will have a bigger percentage as we managed to discuss and provide T2D (Tier 2 with Disk) for the Collaboration.
- During the year, according to monitoring portal we had ->
- Good availability and reliability of the resources, over 97%.
- At this time we run:

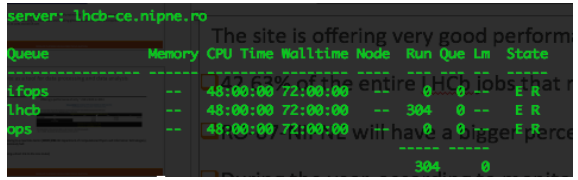



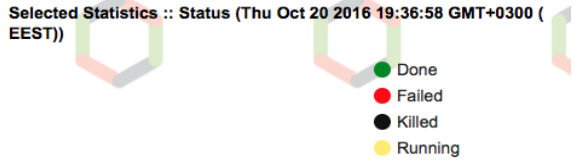


| | |
|---|---|
| Done | 12492 |
| Failed | 244 |
| Killed | 430 |
| Running | 19 |

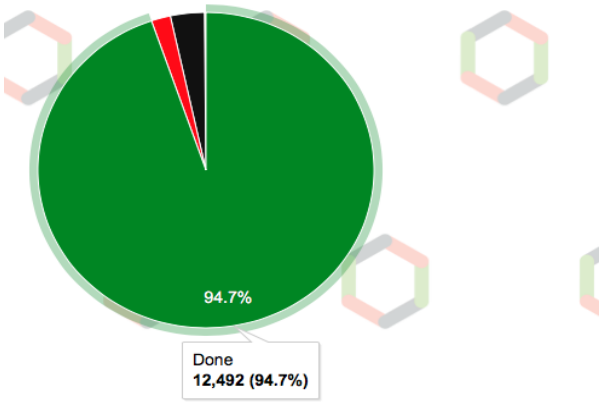

# Grid site as a tool for data processing and data analysis

Future Plans:

- ❑ Major upgrade is going to be needed as the resources have already reached 4 years old
- ❑ Pledges could be changed as the computing changed for some of the experiments (probably storage will be needed)
- ❑ Disk capacity not yet required
- ❑ Migration of the middleware as EMI is running out of support
- ❑ Adapt to new Computationa Physics requirements!

Thank you for your attention!