#### GRID related activities at the INRNE



Elena Puncheva (elly@inrne.bas.bg) Preslav Konstantinov (pkonst@inrne.bas.bg)

**IT Department** 

### Cooperation with CERN

- Cooperation agreement with CERN (June 2001) on "Development of Compute Fabrics for High-Energy Physics and other Sciences".
  - CERN will support INRNE in setting up a compute fabric by donating machines being phased out from it's compute fabrics and by providing training for scientists from INRNE to become familiar with the techniques required to deploy compute fabrics.
  - INRNE will make available the infrastructure necessary to install the equipment (networking, power, etc.) and send staff to CERN to be trained on fabric technologies and to ensure a close collaboration between the partners.



### Motivation

To make available adequate computing facilities such that the Bulgarian HEP scientists can contribute fully to the physics programme of LHC and to other experiments.

Applications from other sciences will also benefit from the available resources.



### First steps



- A local facility is deployed at INRNE in 2002.
- The compute fabric consists of 12 dual processor work nodes, 1 interactive node and 2 server nodes.
- Experience is gained in cluster deployment and management:
  - automatic, script-driven installation and configuration.
    - PBS and Maui scheduler used for workload management.
    - using LDAP to centrally store
       user and networked file system
       information.

## First steps (cont.)

- The computing resources are available for use by research groups from the INRNE and other institutes, as well.
- The facility has been most heavily used for cosmic ray simulations by scientists from the Mussala Environmental Observatory.
- Also available to students for course work and other exercises.

## Training

- Visits by INRNE staff to CERN IT to become familiar with the techniques used for deployment of compute fabrics.
- Involvement in projects of the Fabric and Infrastructure Operations group.
  - developed a prototype web application for displaying data collected by LEMON monitoring agents in the Measurement Repository.
  - performed an installation of Quattor from scratch and provided feedback to the developers in order to improve the installation manual.
  - Preparation (conditioning) and testing of hardware to be shipped to INRNE.

#### The LCG cluster BG-INRNE

- A second cluster, running LCG2 middleware deployed in late 2004 with the assistance of colleagues from CERN GRID Deployment group.
- Joined the LHC Computing Grid in February 2005. Site name BG-INRNE.



### LCG map



## Gstat map

ALBERTA-LCG2	<u>ok</u>	BEgrid-KULeuven	СТ	BEgrid-UGent	<u>OK</u>	BEgrid-ULB-VUB	<u>ok</u>	BEIJING-CNIC-LCG2-IA64
BelGrid-UCL	<u>ok</u>	BG-INRNE	<u>ok</u>	BG01-IPP	<u>SD</u>	BG02-IM	<u>SD</u>	BG04-ACAD SI
BHAM-LCG2	<u>ok</u>	BIFI	<u>ok</u>	BITLab-LCG	<u>OK</u> OK	BNL-LCG2	<u>SD</u>	BRISTOL-PP-LCG
CARLETONU-LCG2	<u>SD</u>	CAVENDISH-LCG2	<u>ok</u>	CERN-CIC	<u>OK</u>	CERN-PROD	<u>ok</u>	CESGA-EGEE 0
CIEMAT-LCG2	<u>ok</u>	CNB-LCG2	<u>SD</u>	<u>cpDIASie</u>	<u>SD</u>	CSCS-LCG2	<u>ct</u>	csQUBuk SI
<u>csUCCie</u>	<u>SD</u>	CY01-LCG2	<u>SD</u>	CYFRONET-IA64	<u>OK</u>	CYFRONET-LCG2	<u>ok</u>	DESY-HH OK OK O
egee.man.poznan.pl	<u>ok</u>	ekplcg2	<u>ok</u>	ESA-ESRIN	<u>SD</u>	FMPhI-UNIBA	<u>SD</u>	FZK-LCG2
giDCUie	<u>ok</u>	giDITie	<u>ok</u>	giITCie	<u>OK</u>	giITTAie	<u>JS</u>	gilTTRie 🛛
giRCSIie	<u>ok</u>	giULie	<u>ok</u>	giWITie	<u>SD</u>	GOG-Singapore	<u>SD</u>	GR-01-AUTH
GR-03-HEPNTUA	<u>ok</u>	GR-04-FORTH-ICS	<u>ok</u>	GR-05-DEMOKRITOS	<u>ok</u> <u>ok</u>	<u>GRIF</u> <u>SD</u>	<u>SD</u>	GSI-LCG2
Hephy-Vienna	<u>ok</u>	HG-01-GRNET	<u>ok</u>	HG-04-CTI-CEID	<u>OK</u>	HG-05-FORTH	<u>ok</u>	HPC2N 0
HR-01-RBI	<u>JS</u>	IEPSAS-Kosice	<u>SD</u>	ifae	<u>OK</u>	IFCA-LCG2	<u>ok</u>	IFIC-LCG2
IN2P3-CC	<u>ok</u>	IN2P3-CPPM	<u>ok</u>	IN2P3-LAPP		IN2P3-LPC OK OK	<u>ok</u>	IN2P3-SUBATECH
INFN-BARI	<u>ok</u>	INFN-BOLOGNA	<u>JS</u>	INFN-BOLOGNA-CMS	<u>OK</u>	INFN-CAGLIARI	<u>ok</u>	INFN-CATANIA SI
INFN-FERRARA	<u>ok</u>	INFN-FIRENZE	<u>ok</u>	INFN-FRASCATI	<u>OK</u>	INFN-LNL-2	<u>ok</u>	INFN-MILANO
INFN-NAPOLI-ATLAS	<u>sok</u>	INFN-PADOVA	<u>ok</u>	INFN-PERUGIA	<u>OK</u>	INFN-PISA2	<u>JS</u>	INFN-ROMA1
INFN-ROMA1-VIRGO	<u>SD</u>	INFN-ROMA2	<u>SD</u>	INFN-T1	<u>OK</u> OK	INFN-TORINO	<u>SD</u>	INTA-CAB

### Currently available resources

- one Computing Element (CE).
- one Storage Element (SE).
- one monitoring server (MON).
- ten dual processor (600 800 MHz) Work Nodes (WN),
   i. e. 20 job slots.
- 2 disk servers with about 1.5 TB space each holding CMS data to be processed by the high energy physics group at the Institute.



#### Features

- Network design with WNs on a private network.
  - benefits are conservation of public IP address space and no packet filtering on the WNs.
  - outbound connectivity for WNs provided by a separate NAT box.
  - Grid service nodes connected to both public and private networks.
  - network segments implemented as 802.1q VLANs using support in Linux kernel.
    - f traffic local to the site (job submission, monitoring, etc.) is confined to the private network.



#### NATed WNs



# Supported VOs

- Resources are allocated for the Virtual organisations:
  - CMS (cms).
  - Geant (geant4).
  - ✓ SixTrack (sixt).
  - Deployment Team (dteam).
- VO's have installed their
  specific software and
  attached the
  corresponding labels to
  the site.

#### Site General Information: ok

```
SiteUniqueID:
siteName:
                    BG-INRNE
GridVersion:
                    LCG-2 6 0
UserSupportContact: mailto: lcg-bg-inrne@inrne.bas.bg
SvsAdminContact:
                    mailto: lcg-bg-inrne@inrne.bas.bg
GlueSiteLatitude:
                   42.4
GlueSiteLongitude: 23.2
GlueCEUniqueID:
               cel.inrne.bas.bg:2119/jobmanager-lcgpbs-cms
                   GlueCEStateStatus:
                                                        Production
                   GlueCEPolicyMaxRunningJobs:
                                                        0
                   GlueCEPolicyMaxWallClockTime:
                                                        17280
                cel.inrne.bas.bg:2119/jobmanager-lcgpbs-dteam
                   GlueCEStateStatus:
                                                        Production
                   GlueCEPolicyMaxRunningJobs:
                                                        0
                   GlueCEPolicyMaxWallClockTime:
                                                        4320
                cel.inrne.bas.bg:2119/jobmanager-lcgpbs-geant4
                   GlueCEStateStatus:
                                                        Production
                   GlueCEPolicyMaxRunningJobs:
                                                        0
                   GlueCEPolicyMaxWallClockTime:
                                                        8640
                ce1.inrne.bas.bg:2119/jobmanager-lcgpbs-sixt
                   GlueCEStateStatus:
                                                        Production
                   GlueCEPolicyMaxRunningJobs:
                                                        0
                   GlueCEPolicyMaxWallClockTime:
                                                        8640
GlueSEUniqueID:
                sel.inrne.bas.bg
                   GlueSEName: BG-INRNE:disk
                   GlueSEPort: 2811
SubClusters:
                ce1.inrne.bas.bg
                   GlueHostOperatingSystemName: SL
                   GlueHostOperatingSystemRelease:
                                                        3.0.4
                   GlueHostOperatingSystemVersion:
                                                        3
                   GlueSubClusterPhysicalCPUs: 0
                   GlueSubClusterLogicalCPUs: 0
                   GlueHostApplicationSoftwareRunTimeEnvironment:
                        LCG-2
                                                      LCG-2_1_0
                       LCG-2 1 1
                                                      LCG-2_2_0
                       LCG-2_3_0
                                                      LCG-2_3_1
                        LCG-2_4_0
                                                      LCG-2_5_0
                        LCG-2_6_0
                                                      R-GMA
                        VO-cms-CMKIN_4_2_0_dar
                                                      VO-cms-CMKIN_4_4_0_dar
                        VO-cms-ORCA_8_7_1_SLC3_dar
                                                      VO-cms-slc3_ia32_gcc323
                        VO-dteam-testingthesite1
                                                      VO-dteam-geant42ndProd
                        VO-dteam-geant42ndProd_p01
                                                      VO-dteam-LCG-2_6_0
                        VO-dteam-geant42ndProd_p02
                                                      VO-dteam-ITUtest1
                        VO-dteam-december05prod-v22NovVO-dteam-december05prod-final
                        VO-dteam-december05prod-final-cand2VO-geant4-december05prod-final-cand03
```

GlueServiceUniqueID:



#### Resource usage

An example of 11 CMS jobs using about 30 hours CPU time and 7 dteam jobs.

Eile         Edit         View         Terminal         Tabs         Help           root@ce1:/var/spool/pbs/server_pr         root@wn21:/poolhome/dteam001/         root@ aether:~/LCG/bdiitest	
root@ce1:/var/spool/pbs/server_pr root@wn21:/poolhome/dteam001/ root@aether:~/LCG/bdiitest	
Job ID Username Queue Jobname SessID NDS TSK Memory Time S Ti	ne 🔺
10002.cel.farms cms001 cms STDIN 28714 1 192:0 R 30	:04
10003.cel.farms cms001 cms STDIN 29483 1 192:0 R 30	:01
wn30 10004.cel.farms cms001 cms STDIN 10233 1 192:0 R 29	: 59
wn29	-
10005.cel.farms cms001 cms S1D1N 9985 1 192:0 R 29 wn29	: 58
10006.cel.farms cms001 cms STDIN 6147 1 192:0 R 30	:01
10007.cel.farms cms001 cms STDIN 6395 1 192:0 R 30	:02
10008.cel.farms cms001 cms STDIN 24076 1 192:0 R 29	: 58
wn27 10009.cel.farms cms001 cms STDIN 24884 1 192:0 R 29	: 59
wn27 10010 celfarms cms001 cms STDIN 30418 1 19200 R 29	. 55
wn26	
10011.cel.farms cms001 cms S1D1N 30664 1 192:0 R 29 wn26	: 55
10012.cel.farms cms001 cms STDIN 12260 1 192:0 R 29	: 58
10026.cel.farms dteam001 dteam STDIN 7077 1 48:00 R 00 wn 25	:00
10027.cel.farms dteam001 dteam STDIN 19565 1 48:00 R 00	:00
10028.ce1.farms dteam001 dteam STDIN 20423 1 48:00 R 00	:00
10029.cel.farms dteam001 dteam STDIN 16230 1 48:00 R 00	:00
wn23 10030.cel.farms dteam001 dteam STDIN 16912 1 48:00 R 00	:00
wn23 10031.cel.farms dteam001 dteam STDIN 6252 1 48:00 R 00	:00
wn22	.00
wn21	.00
[[root@cel root]# qstat -Q	
dteam 0 7 yes yes 0 7 0 0 0 0 Execution	
cms 0 11 yes yes 0 11 0 0 0 0 Execution	
sixt 0 0 ves ves 0 0 0 0 0 0 Execution	1
geant4 0 0 yes yes 0 0 0 0 0 0 0 Execution	•

### Resource usage (cont.)

6386 CPU hours used by 301 CMS jobs in January 2006.

source: http://goc.grid-support.ac.uk/ gridsite/accounting/tree/country\_view. php?BASE=NormSumCPU&startYear =2006&startMonth=1&endYear= 2006&endMonth=2&ExecutingSite =BG-INRNE&Submit=Refresh



Histogram showing the Normalised CPU Time [units 1K.SI2K.Hours] per VO for each of the LHC Tier-1s **Caution:** The selected dataset may not be complete



### Cluster maintenance and support

- Deployment of new middleware releases with minimal downtime.
- Replacement of the hosts with newer ones.
- Monitoring results of the Site Functional Tests and taking corrective action when needed.
- Alerts sent to site admins by SMS.
- Successfully completed the Security Service Challenge exercise.

#### User support and induction

- The system was reconfigured to allow direct access for submitting jobs by users from the CMS research group.
- Provided guidance and assistance to them for obtaining personal user certificates and joining the CMS VO, so they can have access to all CMS GRID resources.
- Developed scripts for data transfer from CERN to INRNE, partially overcoming the restrictions on network use.

### Plans for the immediate future

- Deploying a 1 TB Storage Element with Storage Resource Manager (SRM) interface. Doubling the capacity of the server.
- Upgrading to the latest software release.
- Installing and maintaining a couple of User Interface (UI) machines for local users.
- Expanding the fabric by adding more work nodes.
- Identifying and helping potential users from the Institute, who collaborate with international projects, that have an established VO and contribute and use GRID resources.
- Exploring the possibility to support these VOs at our site.

# Participation in national GRID initiatives

- The Institute for Nuclear Research and Nuclear Energy is cofounder together with the Institute for Parallel Processing of the Bulgarian GRID Consortium.
- BGGC aims and goals



### Conclusion

- The Institute for Nuclear Research is a well established centre for GRID technologies in Bulgaria. It is a host to both technology experts and users.
- This is a result of close collaboration between CERN and INRNE.

# Please, come to see the demonstration!

# Thank you!