

**A Light for Science**



**European Synchrotron Radiation Facility**

# The XRAY Grid

**ESRFUP-WP11**

# Outline

- Our Project: [ESRFUP-WP11](#)
- Our Collaboration: [XRAY VO](#)
- Our Resources: [Hardware and grid components](#)
- Our Deployment : [Grid services @ ESRF and partner sites](#)
- Our Issues: [Network](#)
- Our Figures: [Initial GridFTP transfer rates](#)

# ESRFUP-WP11

- ESRF is getting a Major Upgrade
  - Ambitious Program of 177M € over next 7 years
  - Budget decision had been taken only two weeks ago
- In preparation of this Upgrade Program
  - 13 Work packages already funded and ongoing
- Work Package 11: **Grid Study and Test Bed**
  - Funding for 50PMs over 18 months
  - Up to 150k€ for hardware

# The biggest incentive for the grid... ...is probably the data problem

- How to handle **TeraBytes** and possibly 10s of TBs per experiment?
- How will users access those data sets from home **institutions distributed all over Europe**?
- How to **transfer, replicate and archive** the data?
- How to guarantee **security** while allowing to **share** more widely between groups
- How best to manage the data analysis?


## Current Grid Technology ...

- is able to leverage data analysis and curation of these large data sets
  - It provides highly scalable tools to **store, archive, organize**, as well as **search, access and transfer data in a reliable way**
- can make better use of existing compute cycles
  - offering potentially large amounts of resources for **compute intensive applications**
- comes with a proven **security** Infrastructure (GSI)
- provides the framework to **foster collaboration** within large virtual organizations **by sharing** the compute, storage and software resources and ultimately knowledge

# The deliverables of ESRFUP-WP11

- Deliverable: Collaboration Agreement between 3(-4) partner labs for the **creation of a Synchrotron Radiation Virtual Organization** (Dec '08)
- Purchase compute and storage system, install and test gLite to form Grid Sites at the ESRF and 2(-3) other partner sites
  - Deliverable: **ESRF Grid Site operational** (Dec '08)
  - Deliverable: **Partner Grid Sites operational** (Jun '09)
- Organize a **workshop** with possible future partners
- Gridify one or two resource intensive applications on the test bed, write wrapper software, make added value analysis
  - Deliverable: **Test case software operational** (Oct '09)
- **Final report** on operational experiences with the international test bed installation including **future orientations for photon science grid activities** (Dec '09)

# XRAY Grid Testbed of ESRFUP-WP11





# Creation of a Synchrotron Radiation Virtual Organization



- The **XRAY VO** has been created
  - .. and registered with EGEE
    - <https://cic.gridops.org/index.php?section=vo&vo=xray.vo.eu-egee.org>
  
- **Virtual Organization Management Service** was set up
  - Enrolment URL:
    - <https://grid-voms.esrf.eu:8443/voms/xray.vo.eu-egee.org>
  
- Registered users are granted access to the resources of a VO according to their group membership and assigned role

# Hardware specs

- Storage Components

- Sun Fire X4500 Server (Thumper)
- 2 dual core AMD Opteron, 16 GB
- OpenSolaris based OS (SunOS 5.10)
- 24 TB internal storage



- Compute Nodes

- Sun Fire X2200 M2 Server
- 2 quad core AMD Opteron
- 16 GB main memory, 250 GB HDD
- Scientific Linux 4.7



# Hardware Specs

- Middleware servers

- Sun Fire X4150
- 2 x quad core Intel Xeon E5440
- 16GB main memory, 2x73 GB HDD
- Citrix XenServer 5.0 Enterprise Edition
- Virtual Machines on Scientific Linux 4.7 (i386 or x86\_64)

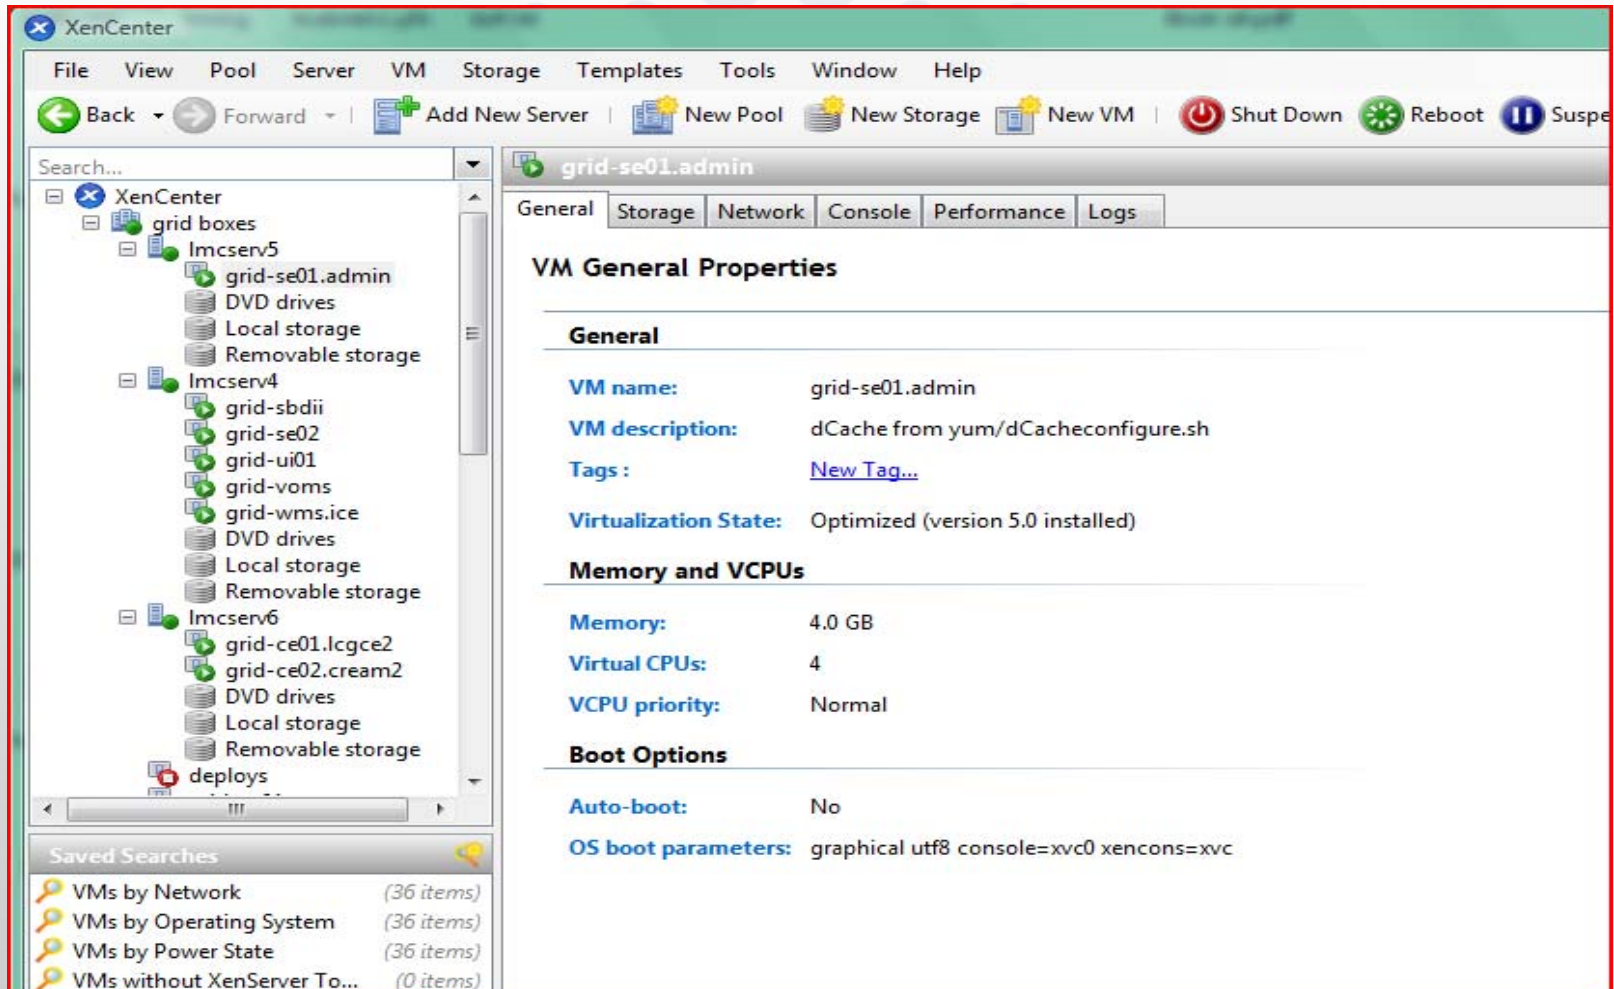


- Rack

- Sun Rack 900-38
- Network Switch, Extreme x450e-48p



# Virtualization using Citrix XenServer 5.0 EE




The screenshot displays the XenCenter web interface. The left-hand navigation pane shows a tree view of the XenCenter environment, including a 'grid boxes' folder containing several VMs like 'grid-se01.admin', 'grid-sbdii', and 'grid-wms.ice'. The main pane shows the configuration for the selected VM 'grid-se01.admin'. The 'General' tab is active, displaying the following properties:

VM General Properties	
<b>General</b>	
<b>VM name:</b>	grid-se01.admin
<b>VM description:</b>	dCache from yum/dCacheconfigure.sh
<b>Tags :</b>	<a href="#">New Tag...</a>
<b>Virtualization State:</b>	Optimized (version 5.0 installed)
<b>Memory and VCPUs</b>	
<b>Memory:</b>	4.0 GB
<b>Virtual CPUs:</b>	4
<b>VCPU priority:</b>	Normal
<b>Boot Options</b>	
<b>Auto-boot:</b>	No
<b>OS boot parameters:</b>	graphical utf8 console=xvc0 xencons=xvc

At the bottom of the interface, there is a 'Saved Searches' section with the following entries:

- VMs by Network (36 items)
- VMs by Operating System (36 items)
- VMs by Power State (36 items)
- VMs without XenServer To... (0 items)



# Site and VO Services

# The different choices

Computing Elements	
LCG-CE	Cream CE
Most widely deployed in EGEE	The next generation
Based on Globus gatekeeper	Based on web services

Storage Resource Manager	
Disk Pool Manager (DPM)	dCache
Most popular in EGEE	Popular with the largest Sites
Medium to high scalability	High scalability and versatility
Disk only storage	Disk storage + backend
Lightweight	High system complexity

## XRAY setup @ ESRF



- **14 Worker Nodes** with altogether **80 Cores**
- **12 TB of disk space (RAID-Z2)**
- **2 Computing Elements** in Test Bed
  - **Lcg-CE** (3.1.21, 32bit)
  - **Cream CE** (1.8.4, 32bit)
- **2 Storage Elements**
  - **dCache** (1.8.0-15p12, 64bit, chimera)
  - **DPM** (1.6.11, 64bit, mysql)
- **1 Site BDII**
- **1 User interface 32bit**
- **1 VOMS** (v1.8.8, 32bit, mysql)
- *Soon 1 WMS*
- *Pending transition to new network setup*

## XRAY setup @ DESY

- DESY enabled XRAY on their existing grid site
- Shipped one thumper (~**20TB**) for their dCache pool

## XRAY setup @ PSI

- 6 Worker Nodes with **48 Cores**
- Up to **20TB** of disk space
- 1 Computing Element: **Icg-CE**
- 1 Storage Element: **dCache**
- 1 User interface, Site-BDII and Monbox
- Cluster is currently on its way to PSI



## XRAY setup @ Soleil

- 6 Worker Nodes with **48 Cores**
- Up to **20TB** of disk space
- 1 Computing Element: Icg-CE
- 1 Storage Element: **DPM**
- 1 User interface, Site-BDII, and Monbox
- Cluster not yet installed
- Pending delivery of Thumper



# Handling the site administration

- Created a couple of **deployment and configuration scripts**
  - Similar to the current practice at the ESRF
- Facilitates the deployment and update of services
- Enables us to keep the configuration homogeneous

```

root@deploys:glite_repos/scripts
File Edit View Terminal Tabs Help
#!/bin/bash
#####
# File : change_network_params.sh
# Project : GRID
# Description : Change the network parameters
# Author(s) : Fernando Calvelo (fernando.calvelo@esrf.fr)
#
# Status : production
# Updated : 25/11/2008
#
# Copyright (c) 2008 by European Synchrotron Radiation Facility
# European Synchrotron Radiation Facility
# ALL RIGHTS RESERVED.
#####
# Usage:
# ./change_network_params-1jpp-1586.rpm
# - site = | esrf |
#####
REPOSITORY="/glite_repos/rx"

# Functions
print_command_syntax ()
{
    echo "#####"
    echo "Usage:"
    echo "#####"
}

rpm -q java-1.5.0-sun-1.5.0-1jpp java-1.5.0-sun-devel-1.5.0-1jpp-1586
if [ $? -ne 0 ]; then
    yum -y install /glite_repos/repos/jdk-1.5.0/jdk-1.5.0-15/java-1.5.0-sun-1.5.0-15-1586.rpm
    yum -y install /glite_repos/repos/jdk-1.5.0/jdk-1.5.0-15/java-1.5.0-sun-devel-1.5.0-1586.rpm
fi

rpm -q openssl-0.9.7a-43.17.el4_6.1
if [ $? -ne 0 ]; then
    yum -y --disablerepo=slc-* install openssl
fi

rpm -q bouncycastle-1.37-1jpp bouncycastle-jdk1.5-1.37-1jpp
if [ $? -ne 0 ]; then
    yum -y install /glite_repos/repos/rpms/bouncycastle-1.37-1jpp.noarch.rpm /glite_repos/rpms/bouncycastle-jdk1.5-1.37-1jpp.noarch.rpm
fi

rpm -q glite-WN-version-3.1.10-0
if [ $? -ne 0 ]; then
    yum -y groupinstall glite-WN
fi

rpm -q lcg-CA-1.25-1 glite-TORQUE_client-3.1.2-0
if [ $? -ne 0 ]; then
    yum -y install lcg-CA glite-TORQUE_client
fi

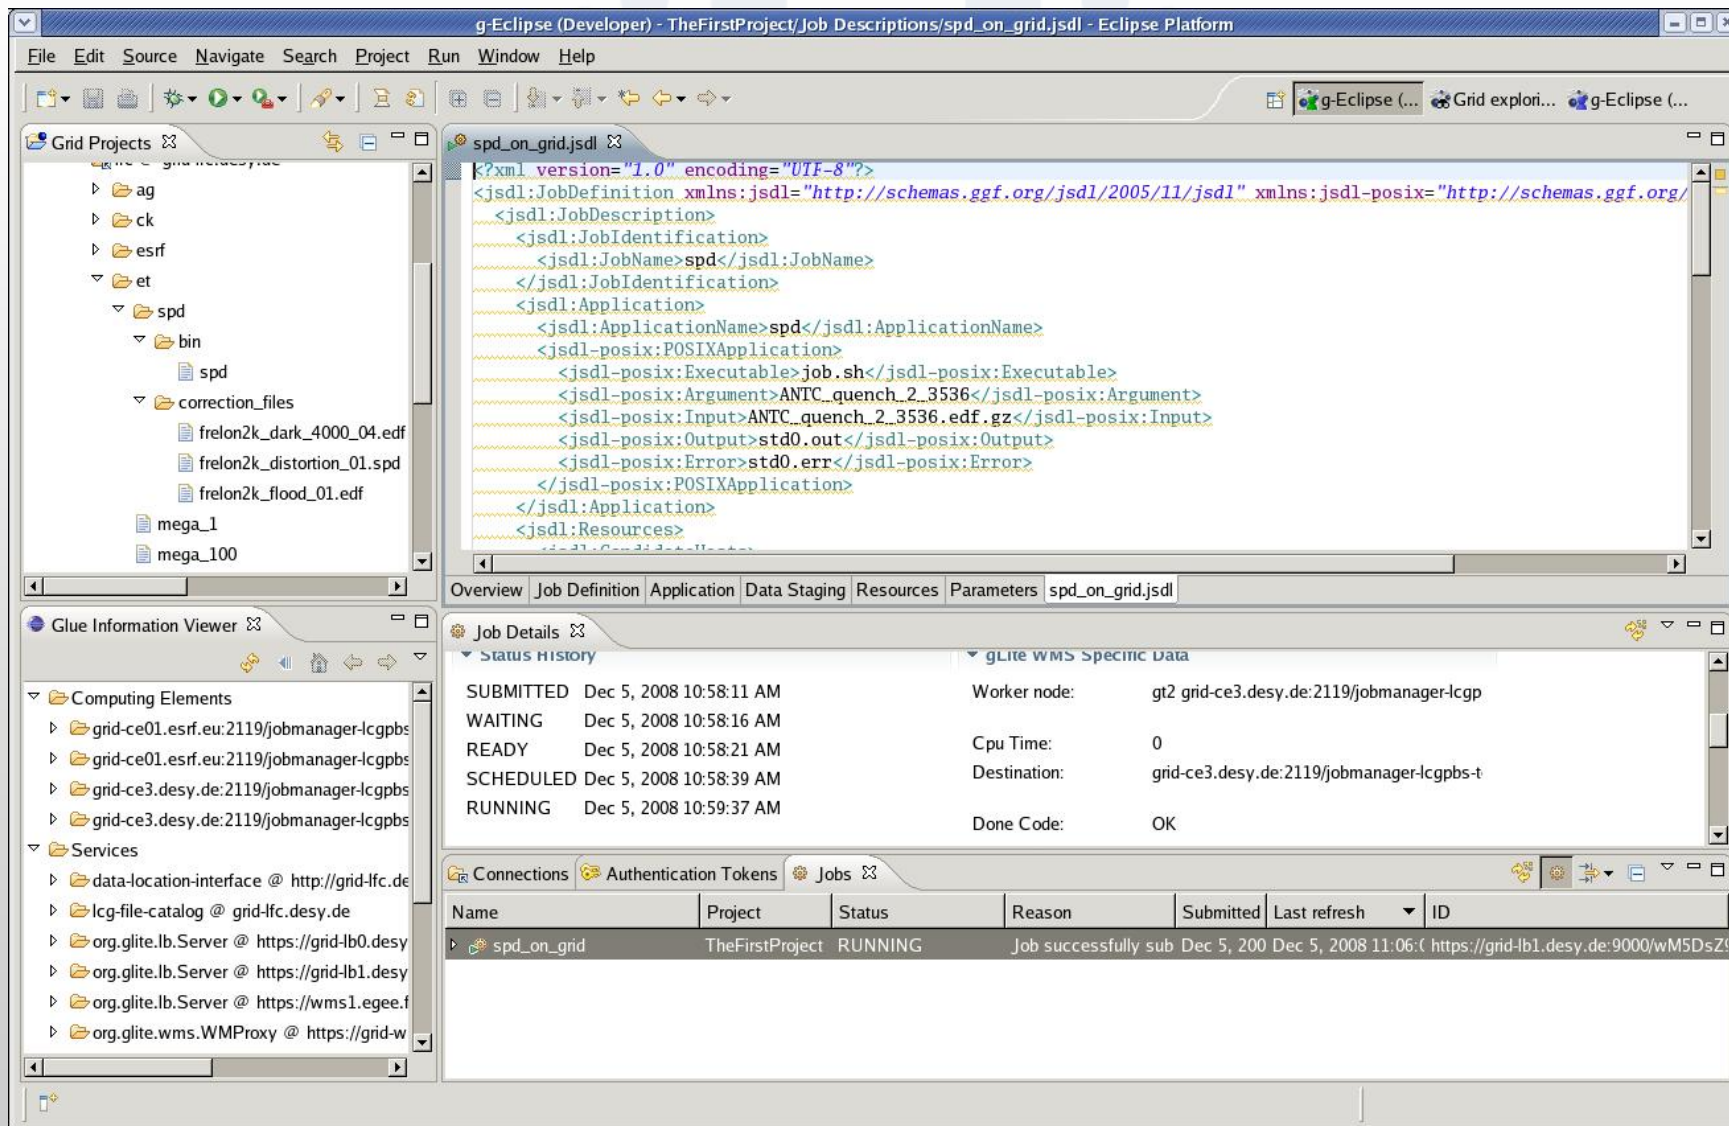
/opt/glite/yaim/bin/yaim -c -s /root/yaim/site-info.def -n glite-WN -n glite-TORQUE_client

yum -y --exclude=bouncycastle update

:
#@voas"
echo "Put here steps for install $NODE_TYPE node"

```

# gEclipse



The screenshot displays the gEclipse (Developer) IDE interface. The main editor shows the XML content of a JSDL file named 'spd\_on\_grid.jsdl'. The left sidebar contains a 'Grid Projects' tree view and a 'Glue Information Viewer' showing a list of services. The bottom section features a 'Job Details' table and a 'Connections' table.

```

<?xml version="1.0" encoding="UTF-8"?>
<jsdsl:JobDefinition xmlns:jsdsl="http://schemas.ggf.org/jsdsl/2005/11/jsdsl" xmlns:jsdsl-posix="http://schemas.ggf.org/...
<jsdsl:JobDescription>
  <jsdsl:JobIdentification>
    <jsdsl:JobName>spd</jsdsl:JobName>
  </jsdsl:JobIdentification>
  <jsdsl:Application>
    <jsdsl:ApplicationName>spd</jsdsl:ApplicationName>
    <jsdsl-posix:POSIXApplication>
      <jsdsl-posix:Executable>job.sh</jsdsl-posix:Executable>
      <jsdsl-posix:Argument>ANTC_quench_2_3536</jsdsl-posix:Argument>
      <jsdsl-posix:Input>ANTC_quench_2_3536.edf.gz</jsdsl-posix:Input>
      <jsdsl-posix:Output>std0.out</jsdsl-posix:Output>
      <jsdsl-posix:Error>std0.err</jsdsl-posix:Error>
    </jsdsl-posix:POSIXApplication>
  </jsdsl:Application>
</jsdsl:JobDescription>
</jsdsl:JobDefinition>
  
```

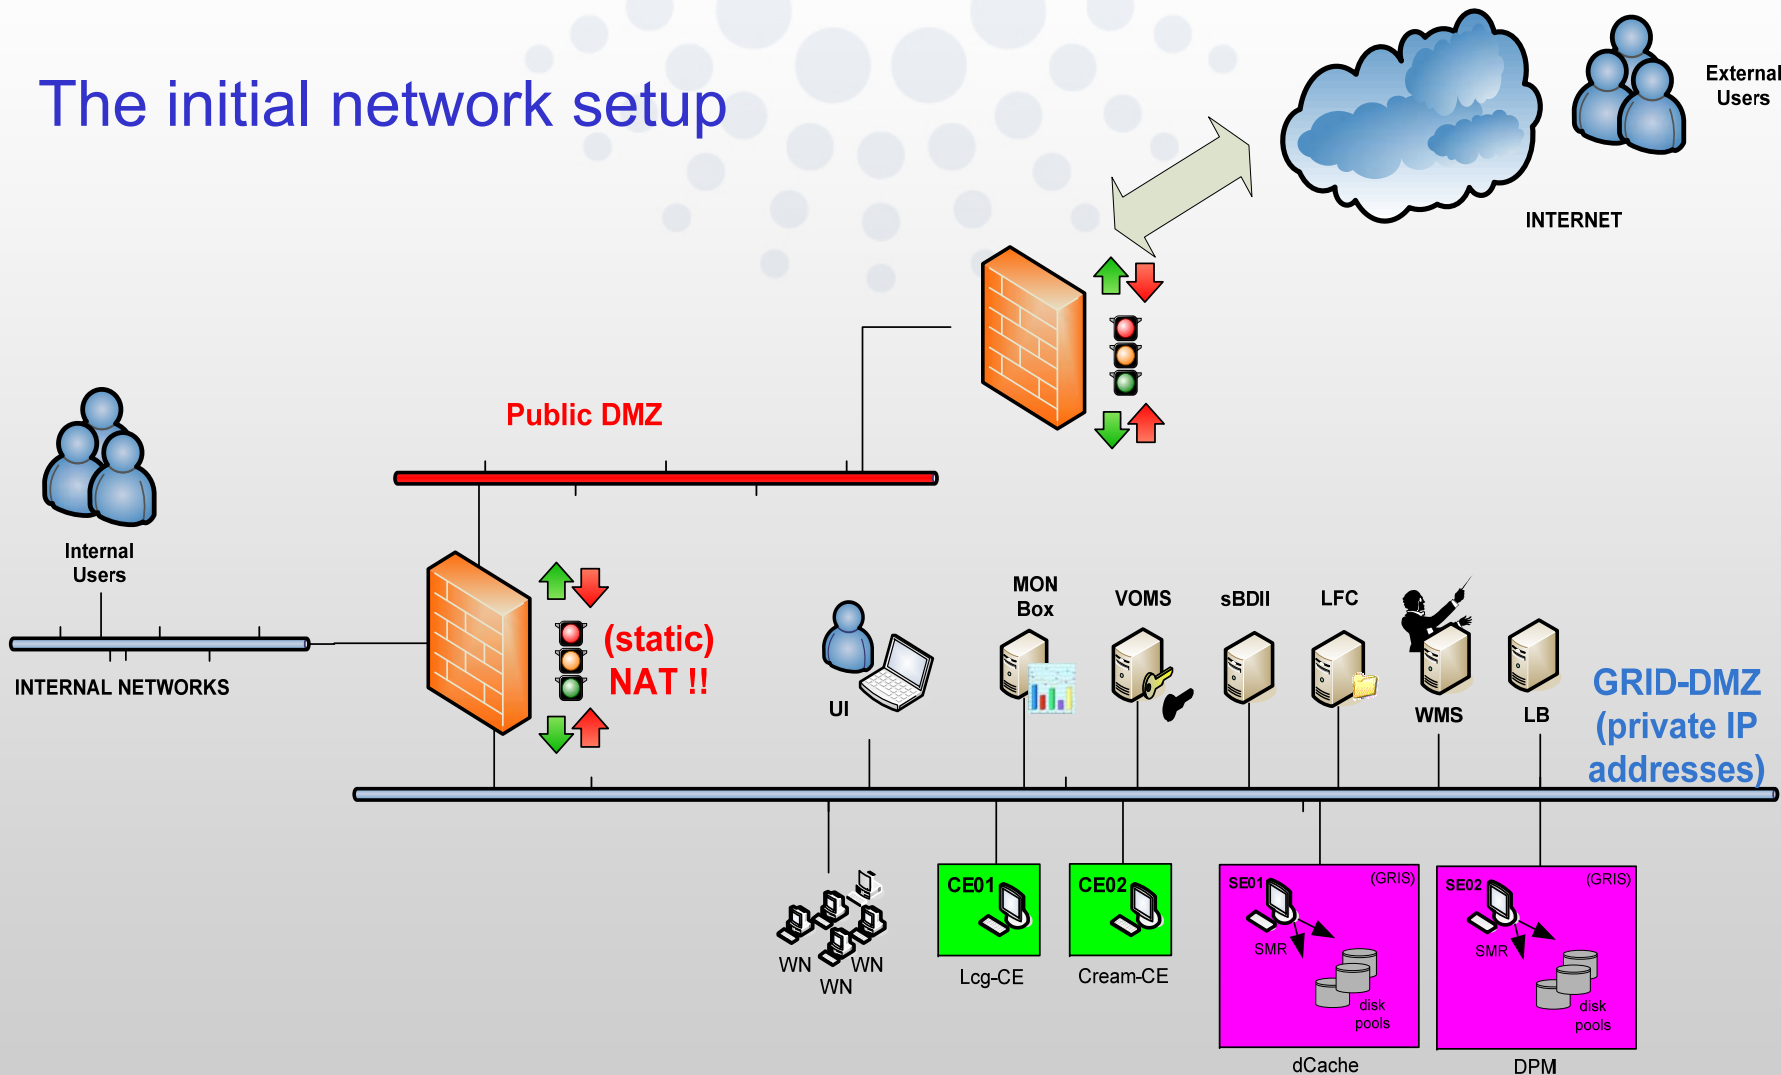
**Job Details Table:**

Status	Time	Worker node:	Cpu Time:	Destination:	Done Code:
SUBMITTED	Dec 5, 2008 10:58:11 AM	gt2 grid-ce3.desy.de:2119/jobmanager-lcgp			
WAITING	Dec 5, 2008 10:58:16 AM		0		
READY	Dec 5, 2008 10:58:21 AM			grid-ce3.desy.de:2119/jobmanager-lcgpbs-t	
SCHEDULED	Dec 5, 2008 10:58:39 AM				
RUNNING	Dec 5, 2008 10:59:37 AM				OK

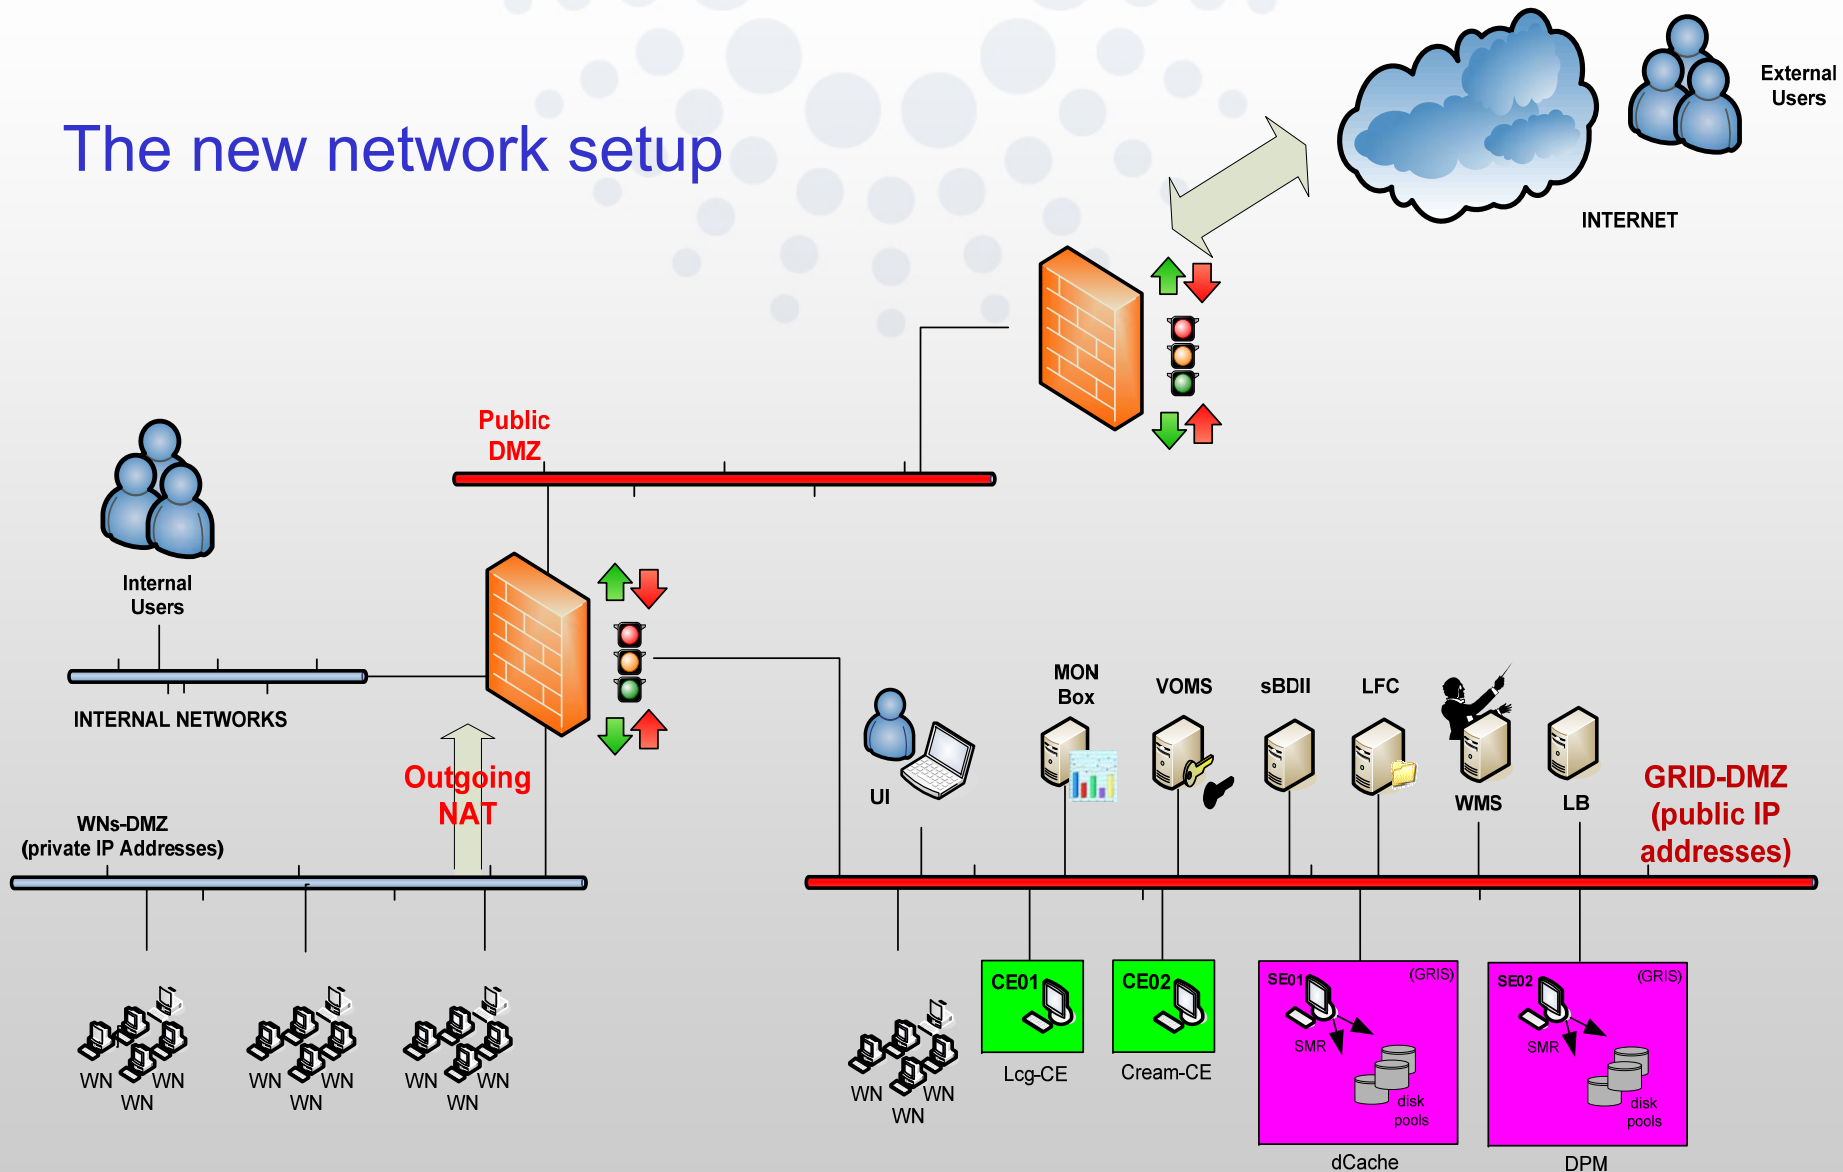
**Connections Table:**

Name	Project	Status	Reason	Submitted	Last refresh	ID
spd_on_grid	TheFirstProject	RUNNING	Job successfully sub	Dec 5, 2008 11:06:00	Dec 5, 2008 11:06:00	https://grid-lb1.desy.de:9000/wM5DsZ...

## The initial network setup



## The new network setup



# Some GridFTP transfer numbers

- GridFTP File Transfers ESRF → DESY SE
  - Single stream: ~ **1.5MB/s** (50 MB, 100 MB, or 500 MB files)
  - 8 parallel streams: **8 MB/s** for a 100 MB files
    - **Translates in 35 h for a 1 TB dataset**
- GridFTP File Transfer DESY SE → ESRF
  - ~ **1.5MB/s** for various files sizes
  - parallel streams blocked
- GridFTP File Transfer:
 

CERN → ESRF	SE	ESRF → CERN
• (45MB file)	<b>~3-5 MB/s</b>	<b>~3-5 MB/s</b>
• 8 parallel streams:	<b>13 MB/s</b>	(1TB/d)
- Uplink to Renater is **1Gb/s (125MB/s)**, but shared between ESRF, ILL, EMBL
  - **Do we need a dedicated >1Gb/s uplink for the grid ?!**

# Test the grid yourself

- Want to test the grid yourself?

→ start here:

<http://www.esrf.fr/Infrastructure/Computing/Grid/how-to-become-an-xray-user>

- Want to discuss grid?

→ contact [grid-admin@esrf.fr](mailto:grid-admin@esrf.fr)

# Tomorrow's XRAY Grid!?!

