



Enabling Grids for E-science

WMS+LB: installation, configuration & testing

Giuseppe La Rocca

INFN – Catania

giuseppe.larocca@ct.infn.it

EMBRACE-EGEE Tutorial

www.eu-egee.org

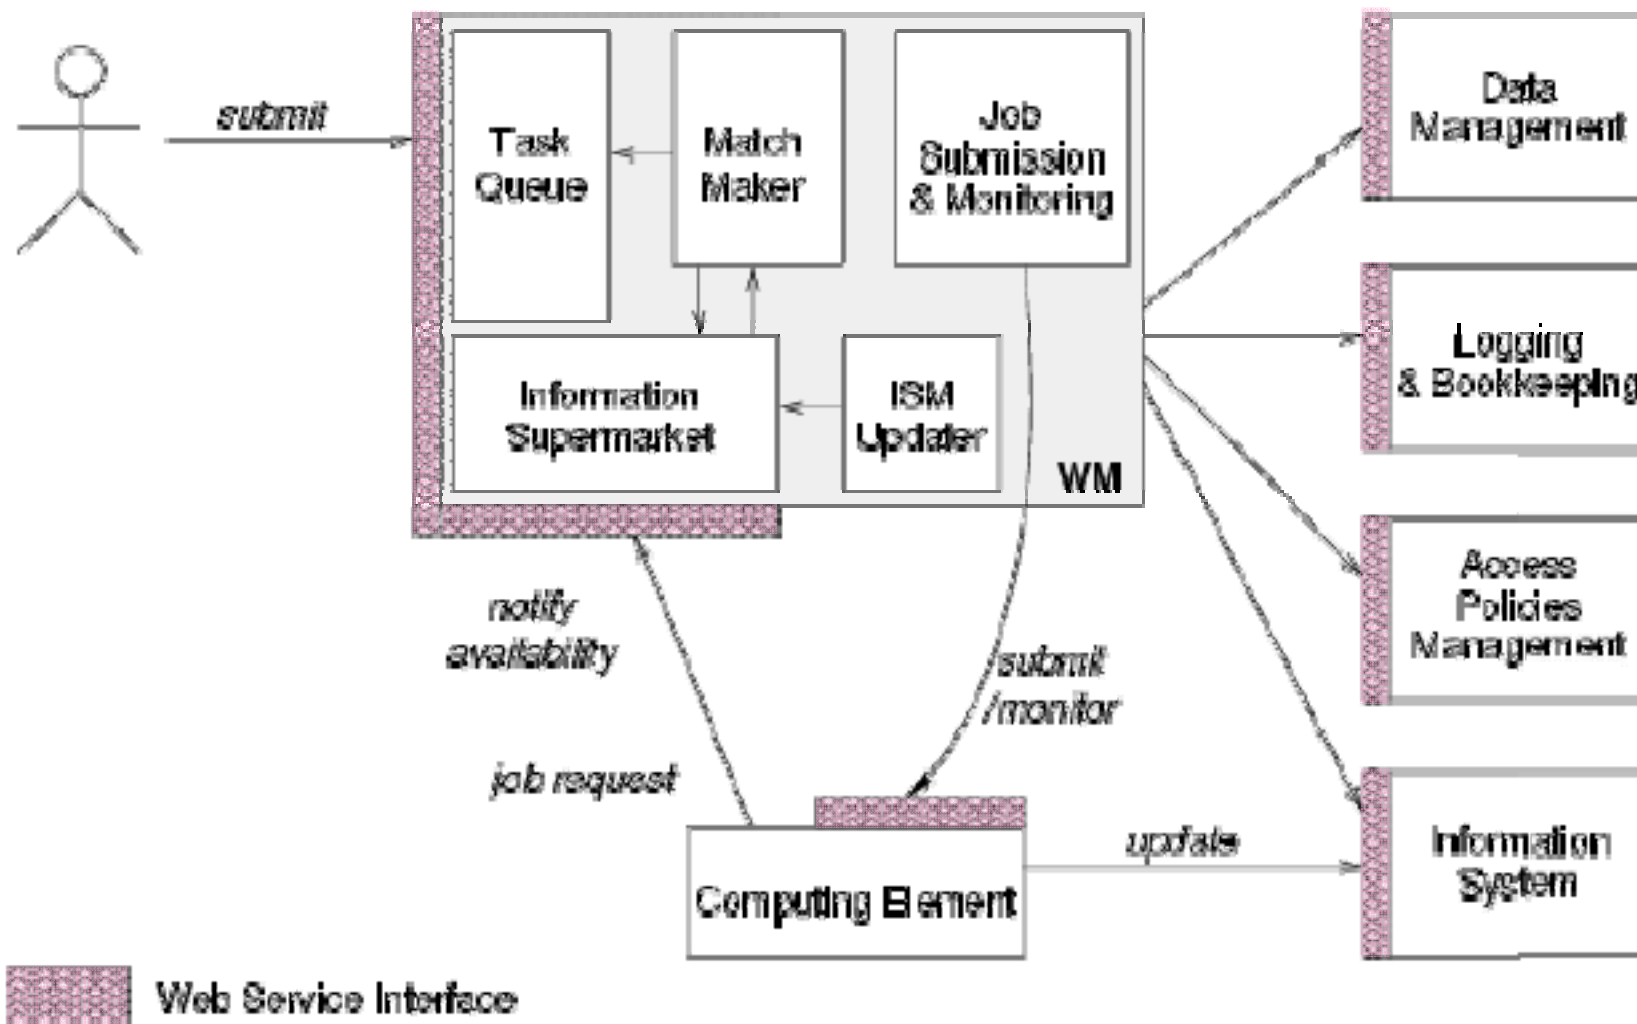


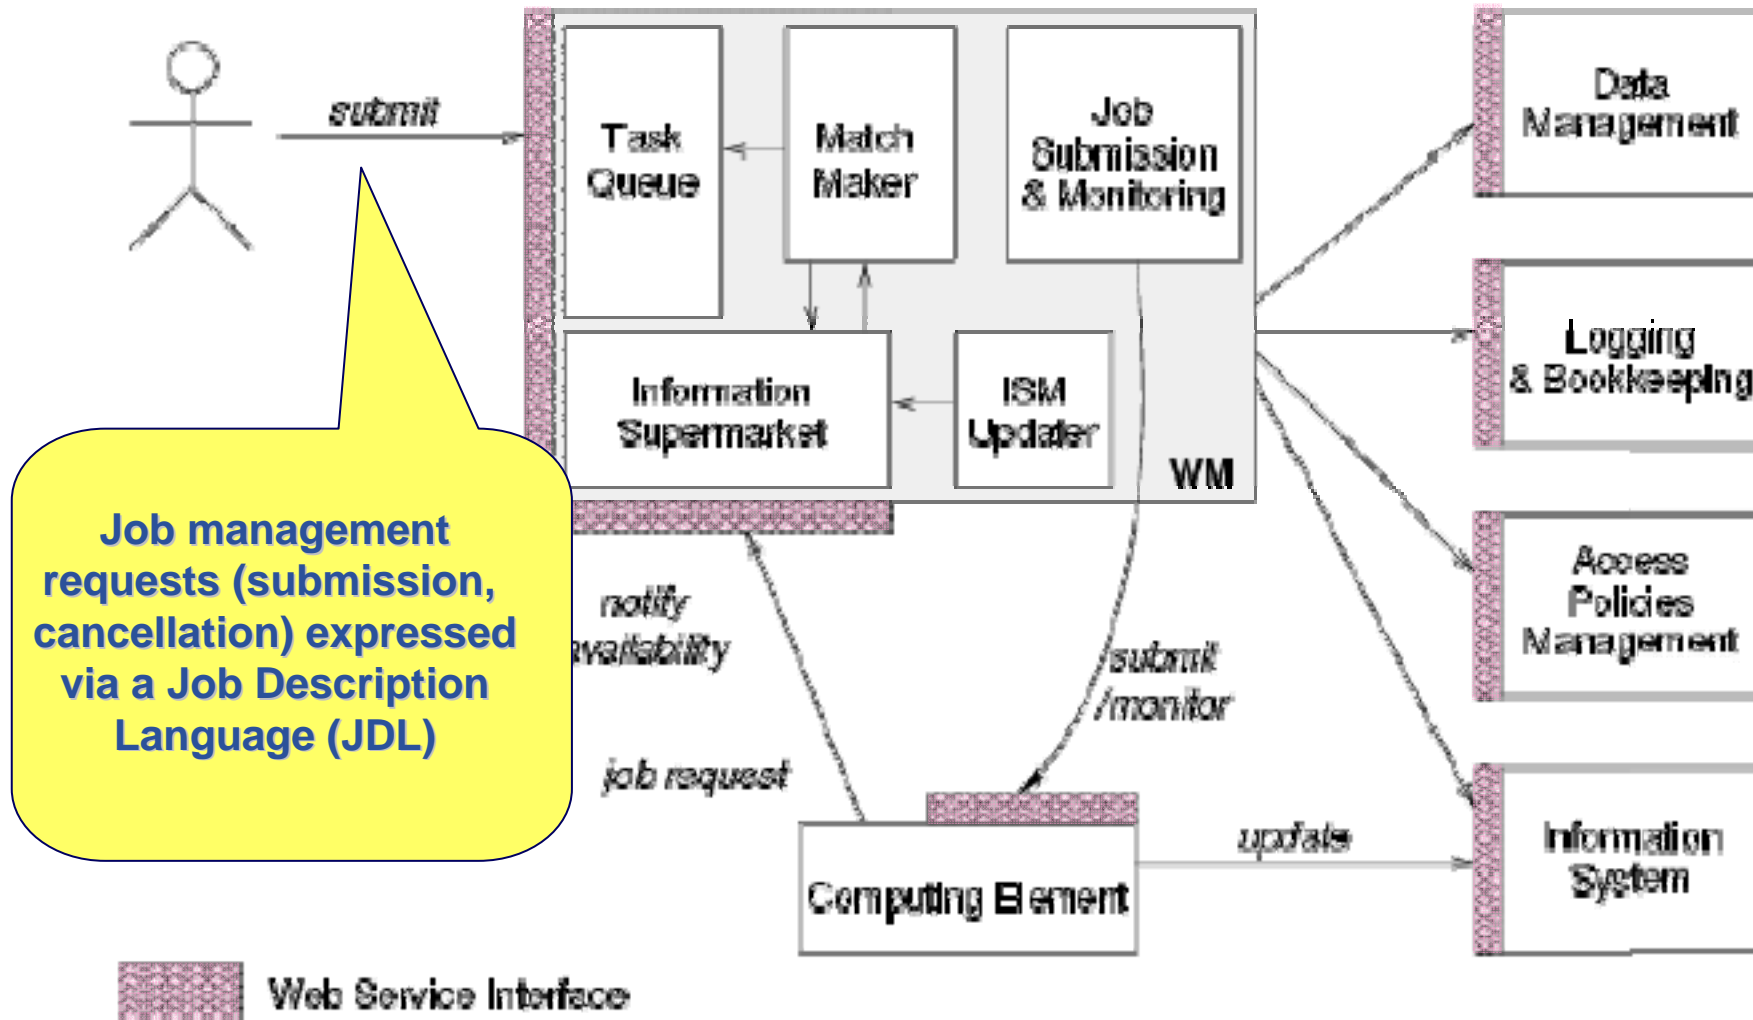
Information Society
and Media

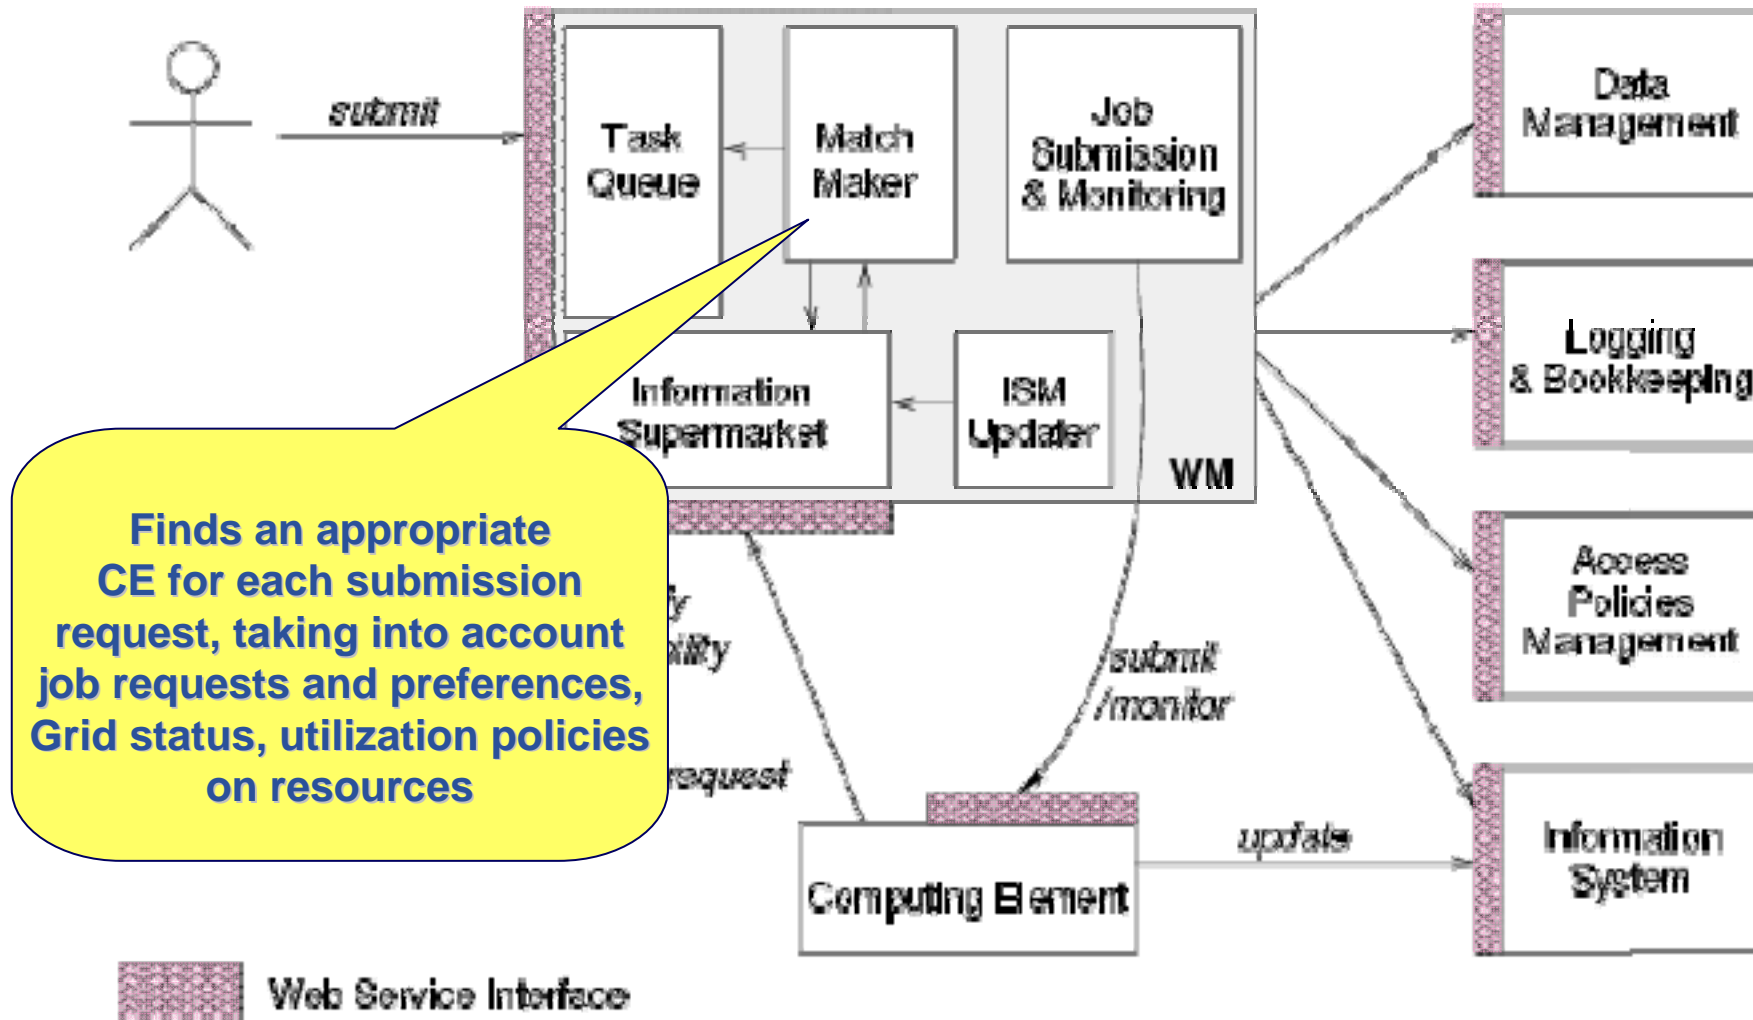


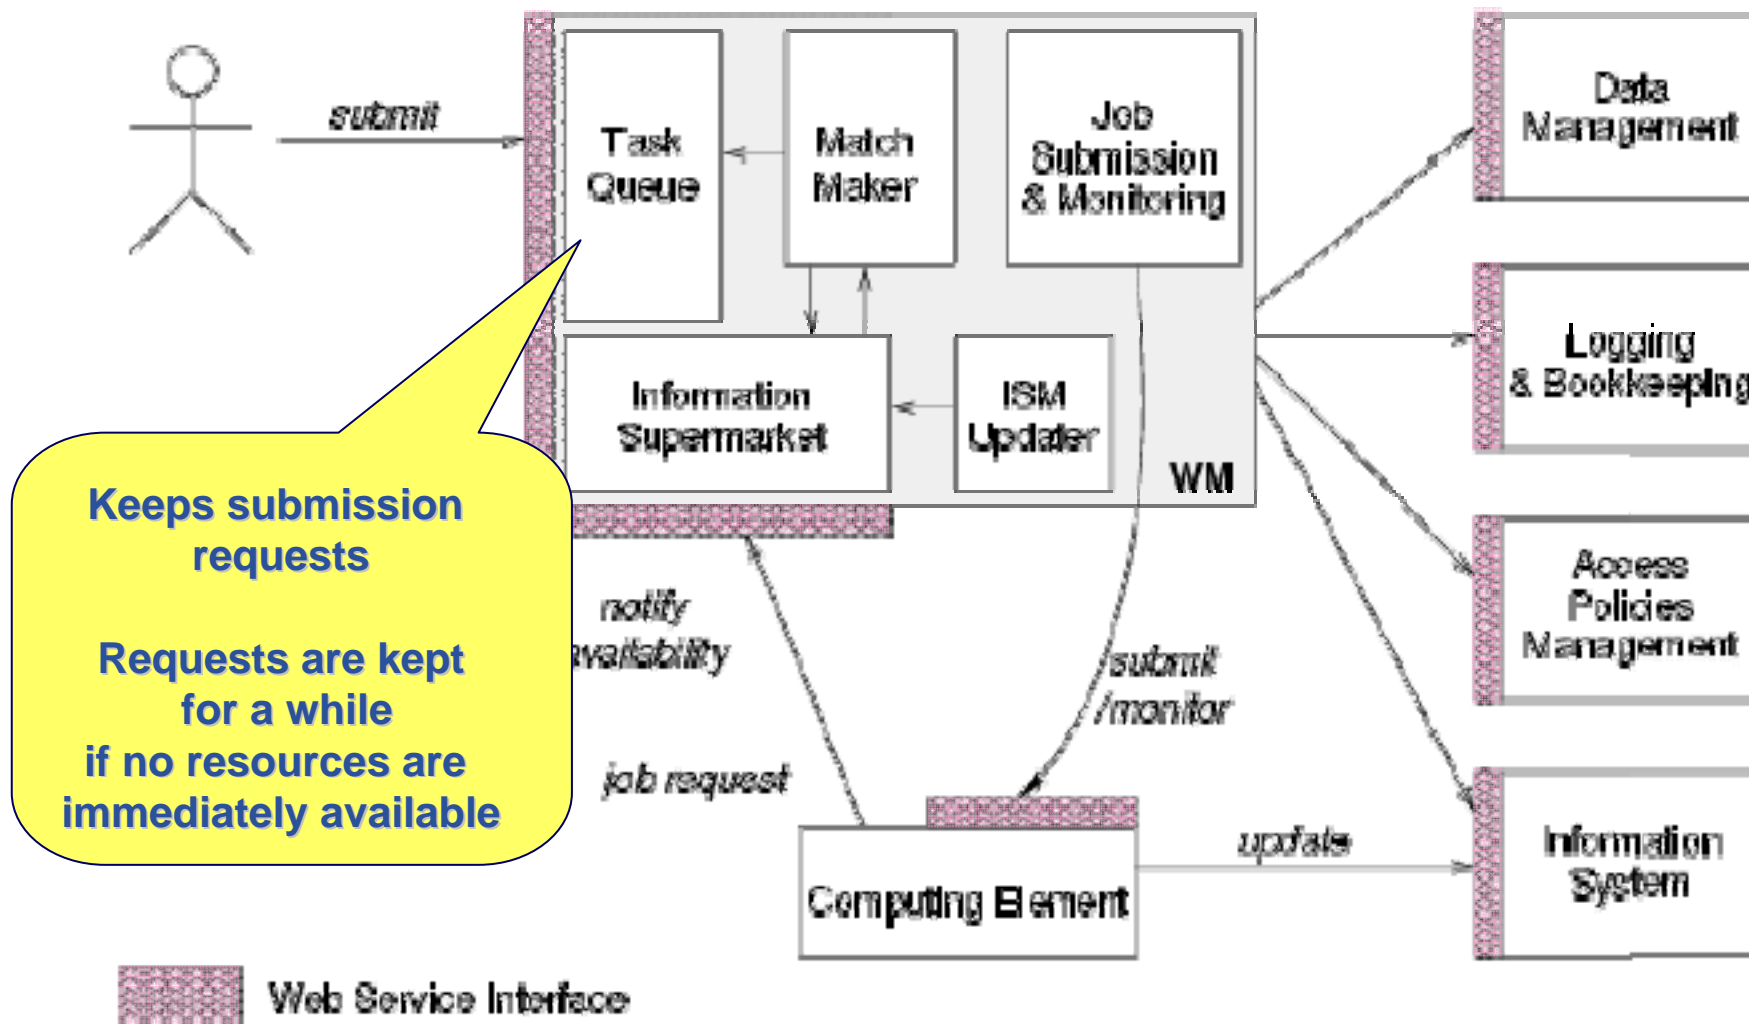
- **What is a Workload Management System (WMS) ?**
- **What is a Logging & Bookkeeping Service (LB) ?**
- **How to install a WMS+LB**
- **WMS+LB Testing**
- **Logs location and references**

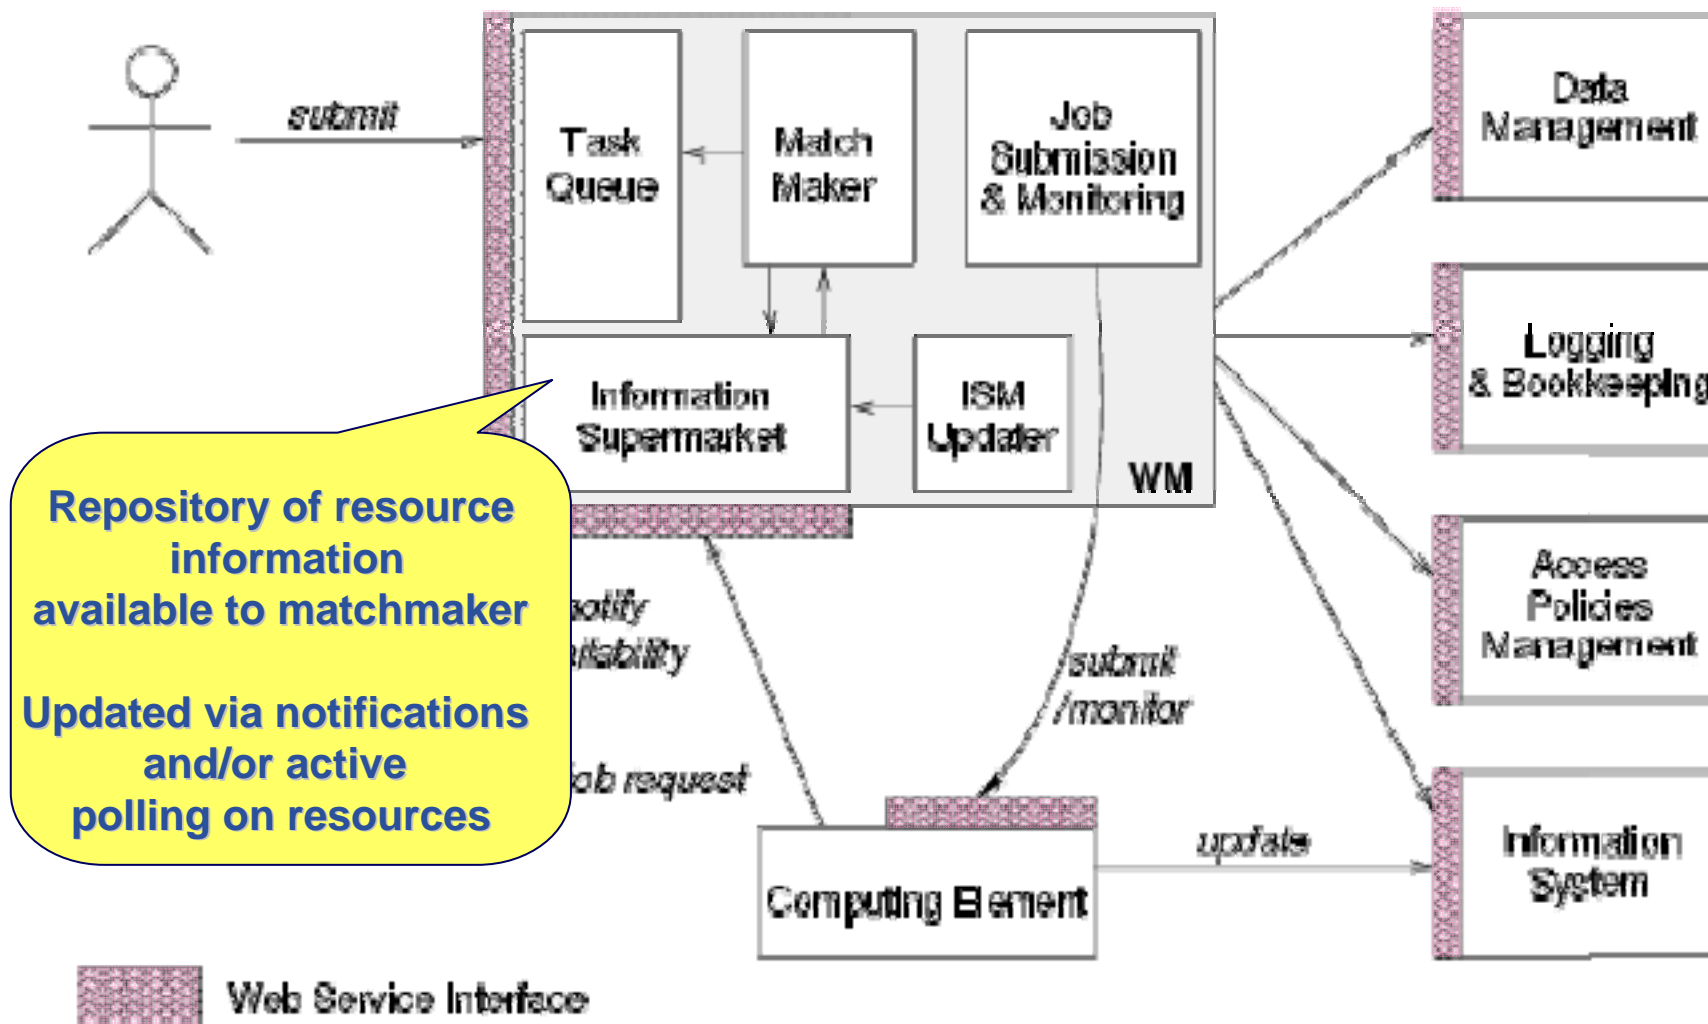
- **Resource Broker is an alternate for Workload Management System + Logging and Bookkeeping**
- **Accepts and satisfies the requests for job management coming from its clients (UI's)**
- **Requests are specified through JDL files using ClassAd**
- **WMS, taken a valid request, chooses the most appropriate action to satisfy it.**
- **Its main task is individuating the best suitable resources (CE, SE...)**

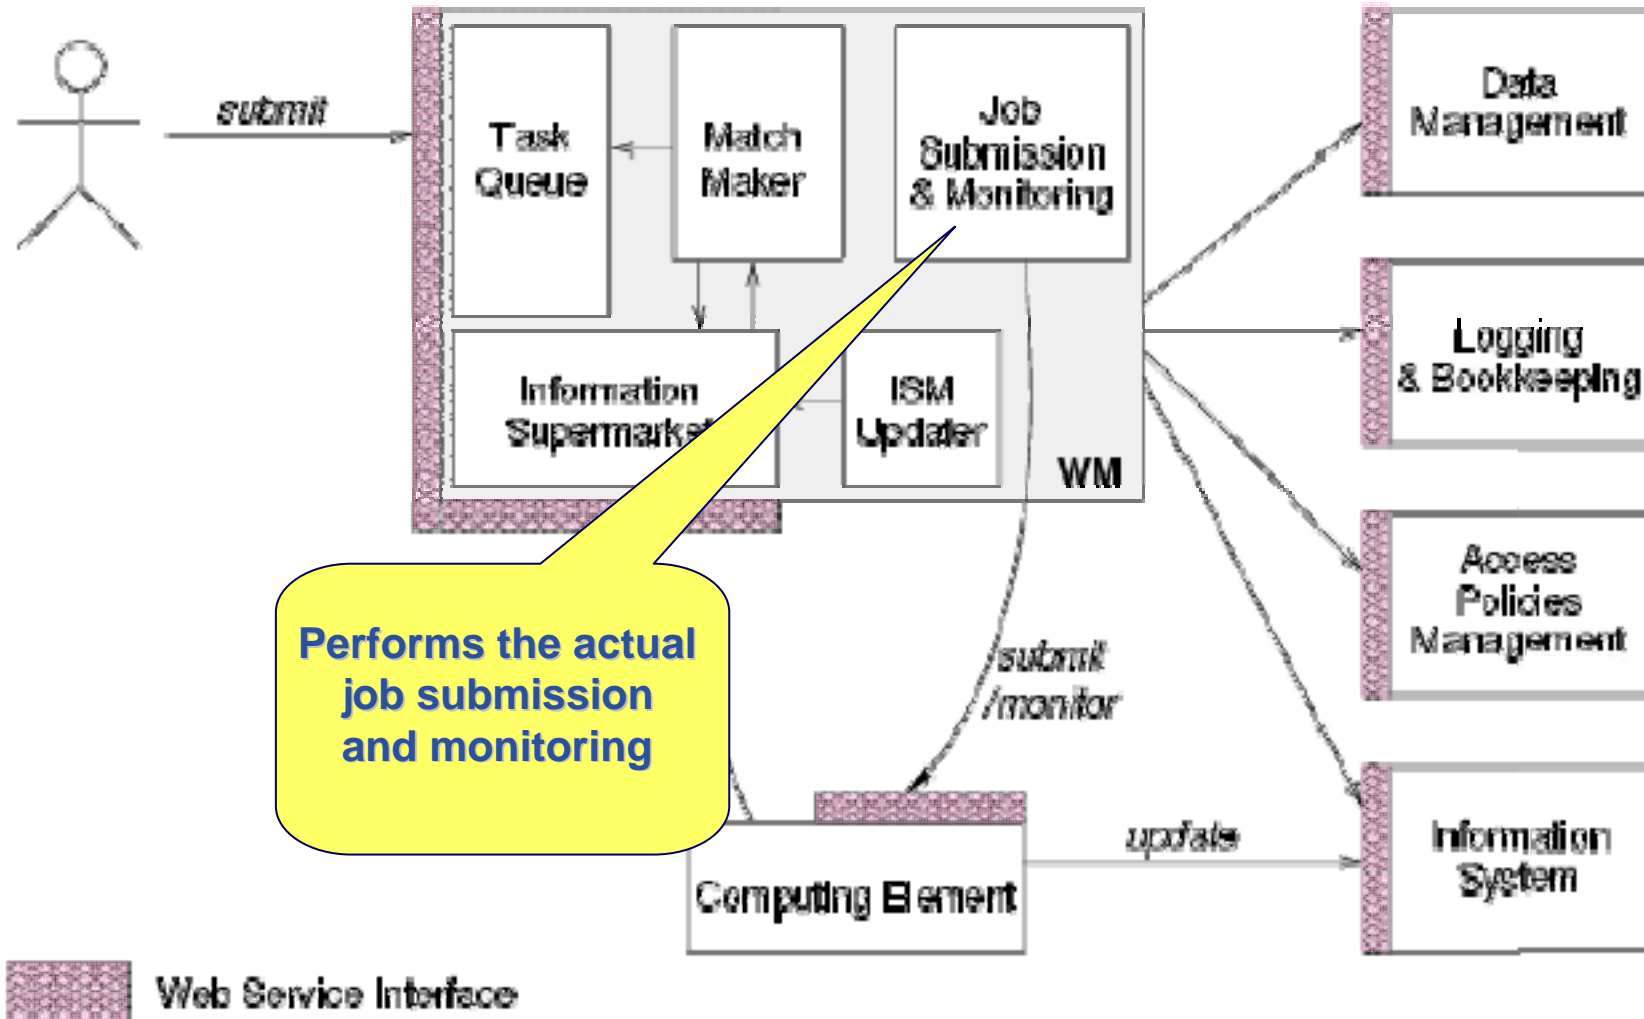


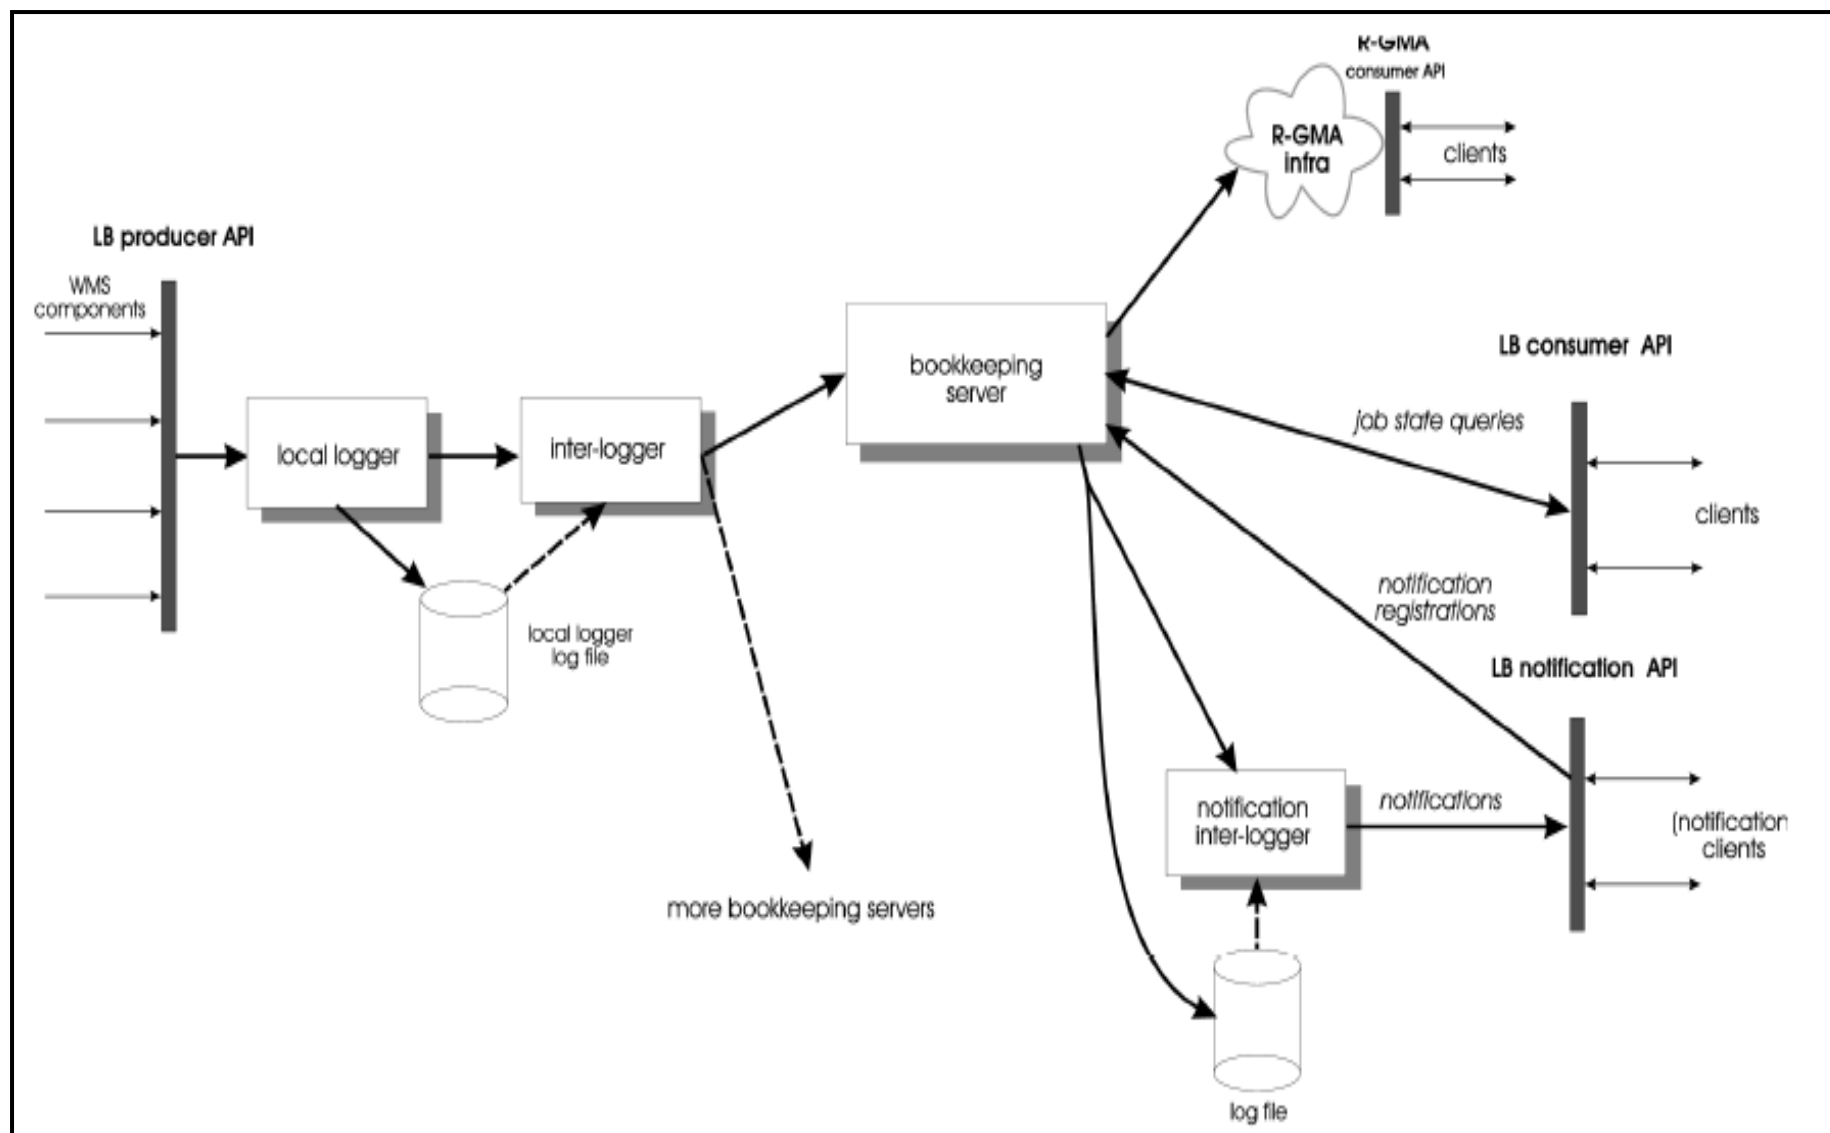












Logging and Bookkeeping (L&B) do:

- Gather events about jobs from various WMS, CE components, process them and provide a high level view of the state of jobs
- Provide interfaces for quering L&B, register for notifications

Basic Components are:

- **LocalLogger**
- **InterLogger**
- **Bookkeeping Server**



Installing the Workload Management System (WMS) + Logging and Bookkeeping System

- **Start from the Virtual Machine Base that you can download from :**

`https://gilda.ct.infn.it/GILDAVM/GILDAVM_Base.tar.bz2`

- **Verify that these packages are installed and properly configured:**
 - **Java SDK 1.4.2 (or greater)**
 - **edg_VO_Gilda rpm** (`https://gilda.ct.infn.it/RPMS/`)
 - **glite-yaim-3.0.0**
(`http://glitesoft.cern.ch/EGEE/gLite/APT/R3.0/rhel30/RPMS.Release3.0/`)
 - **gilda_ig-yaim-3.0.0**
(`http://grid018.ct.infn.it/apt/gilda_app-i386/utils/`)

- **Request host certificates for the WMS to a CA**
 - <https://gilda.ct.infn.it/CA/mgt/restricted/srvreq.php>
- **Copy host certificate (hostcert.pem and hostkey.pem) in /etc/grid-security**
- **Change the permissions**
 - `chmod 644 hostcert.pem`
 - `chmod 400 hostkey.pem`

- Because of SUN licence used for Java SDK, it is not possible to redistribute it with the middleware.
- You have to download Java SDK 1.4.2 from Sun web site:

<http://java.sun.com/j2se/1.4.2/download.html>

- Select **``Download J2SE SDK``**, and download the **``RPM in self-extracting file``**. Follow the instructions on the pages to extract the rpm.

- A general requirement for the gLite nodes is that they are synchronized.
- Configure the file `/etc/ntp.conf` by adding the lines dealing with your time server configuration such as, for instance:

```
# Prohibit general access to this service.
restrict default ignore
restrict 193.206.144.10 mask 255.255.255.255
nomodify notrap noquery
```

```
server 127.127.1.0 # local clock
fudge 127.127.1.0 stratum 10
server ntp-1.infn.it
```


- **Edit the file `/etc/ntp/step-tickers` adding a list of your time server(s) hostname(s)**

```
cat /etc/ntp/step-tickers
193.206.144.10
```

- **If you are running a kernel firewall, you will have to allow inbound communication on the NTP port.**
- **If you are using iptables, you can add the following to `/etc/sysconfig/iptables`**

```
-A INPUT -s <NTP-serverIP-1> -p udp --dport 123 -j
ACCEPT
```

- **You can then reload the firewall : `/etc/init.d/iptables restart`**

- **Activate the ntpd service with the following commands:**

```
# ntpdate <your ntp server name>
```

```
# service ntpd start
```

```
# chkconfig ntpd on
```

- **You can check ntpd's status by running the following command :**

```
# ntpq -p
```

Install glite-yaim and gilda_ig-yaim

- Download and install latest version of glite-yaim-3.0.0 -* on your machine

<http://glitesoft.cern.ch/EGEE/gLite/APT/R3.0/rhel30/RPMS.Release3.0/glite-yaim-3.0.0-11.noarch.rpm>

- Download and install the latest version of gilda_ig-yaim-3.0.0 -* on your machine

http://grid018.ct.infn.it/apt/gilda_app-i386/utils/gilda_ig-yaim-latest

- glite-yaim and gilda_ig-yaim provide a set of bash ``mini-scripts''. Each ``mini-script'' implements one bash function and it is stored in file with the same name of the function. Each function configures a specific middleware module. The functions are stored in the two directories:

`/opt/glite/yaim/functions` (glite-yaim functions)

`/opt/glite/yaim/functions/local` (gilda_ig-yaim functions)

- Copy the yaim configuration template file into the root dir:

```
cp /opt/glite/yaim/examples/gilda_ig-site-info.def
/root/my-site-info.def
```

- Open /root/my-site-info.def file using a text editor and set the following values according to your grid environment:

MY_DOMAIN=<your DOMAIN>

CE_HOST=grid010.ct.infn.it

RB_HOST=glite-rb.ct.infn.it

WMS_HOST=glite-rb.ct.infn.it

PX_HOST=grid001.ct.infn.it

BDII_HOST=grid004.ct.infn.it

MON_HOST=rgmasrv.ct.infn.it

FTS_HOST=gilda-fts.ct.infn.it

REG_HOST=rgmasrv.ct.infn.it

NTP_HOSTS="193.206.144.10"



For this tutorial substitute
grid018.ct.infn.it/rep with 192.168.0.50

```
OS_REPOSITORY="rpm  
http://grid018.ct.infn.it/rep slc306-i386 os  
updates extras"
```

```
LCG_REPOSITORY="rpm  
http://grid018.ct.infn.it/rep glite_sl3-i386  
3_0_0 3_0_0 externals 3_0_0 updates"
```

```
IG_REPOSITORY="rpm  
http://grid018.ct.infn.it/rep ig_sl3-i386  
3_0_0 utils"
```

```
GILDA_REPOSITORY="rpm  
http://grid018.ct.infn.it/rep gilda_app-i386  
app 3_0_0"
```

```
CA_REPOSITORY="rpm  
http://grid018.ct.infn.it/rep glite_sl3-i386  
security"
```

- **Add LFC settings:**

```
LFC_HOST=lfc-gilda.ct.infn.it
```

- **Configure DPM host:**

```
DPM_HOST=aliserv6.ct.infn.it
```

- **Check/Modify VO and VOMS specific configurations:**

```
VOS="gilda .."
```

```
ALL_VOMS="gilda .."
```

- **Choose a good password! And be sure that this file cannot ...**

```
MYSQL_PASSWORD=set_this_to_a_good_password
```

`SITE_EMAIL=<customize here>`



`#Specific VO settings`

`VO_GILDA_SW_DIR=$VO_SW_DIR/gilda`

`VO_GILDA_DEFAULT_SE=$DPM_HOST`

`VO_GILDA_STORAGE_DIR=$CLASSIC_STORAGE_DIR//gilda`

`VO_GILDA_QUEUES="short long infinite"`

`VO_GILDA_VOMS_SERVERS="vomss://voms.ct.infn.it:8443/voms
/gilda?/gilda"VO_GILDA_VOMSES="'gilda voms.ct.infn.it
15001 /C=IT/O=GILDA/OU=Host/L=INFN
Catania/CN=voms.ct.infn.it/Email=emidio.giorgio@ct.infn.
it gilda'"`

`JAVA_LOCATION="/usr/java/j2sdk1.4.2_12/"`



```
# These variables tell YAIM where to find additional  
configuration files.
```

```
[..]
```

```
# INFN-GRID: customized pool accounts list
```

```
USERS_CONF=/opt/glite/yaim/examples/gilda_ig-users.conf
```

```
GROUPS_CONF=/opt/glite/yaim/examples/gilda_ig-groups.conf
```

```
FUNCTIONS_DIR=/opt/glite/yaim/functions
```

```
YAIM_VERSION=3.0.0-1
```

```
[..]
```


- **We are ready to install the WMS:**

```
/opt/glite/bin/gilda_ig_install_node
/root/my-site-info.def GILDA_ig_RB
```

- This command will download and install all the needed packages

- **Now we can configure the node:**

```
/opt/glite/bin/gilda_ig_configure_node
/root/my-site-info.def GILDA_ig_RB
```



WMS firewall configuration

```
# Firewall configuration wirtten by redhat-config-
securitylevel
# Manual customization of this file is not recommeded.
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [0:0]
:RH-Firewall-1-INPUT - [0:0]
-A INPUT -j RH-Firewall-1-INPUT
-A FORWARD -j RH-Firewall-1-INPUT
-A RH-Firewall-1-INPUT -i lo -j ACCEPT
-A RH-Firewall-1-INPUT -s 193.206.144.10 -p udp --dport
123 -j ACCEPT
-A RH-Firewall-1-INPUT -p icmp --icmp-type any -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state
ESTABLISHED,RELATED -j ACCEPT
```

```
# Enable incoming SSH Connection.
```

```
-A RH-Firewall-1-INPUT -m state --state NEW -m
  tcp -p tcp -s XXX.XXX.XXX.XXX --dport 22 -j
  ACCEPT
```

```
# UI ports.
```

```
-A RH-Firewall-1-INPUT -m state --state NEW -m
  tcp -p tcp --dport 9000:9002 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m
  tcp -p tcp --dport 7772 -j ACCEPT
```

```
# UI / CE / SE / WN ports.
```

```
-A RH-Firewall-1-INPUT -m state --state NEW -m
  tcp -p tcp --dport 7771 -j ACCEPT
```

```
-A RH-Firewall-1-INPUT -m state --state NEW -m
  tcp -p tcp --dport 9002 -j ACCEPT
```

```
-A RH-Firewall-1-INPUT -m state --state NEW -m
  tcp -p tcp --dport 2170 -j ACCEPT
```

```
-A RH-Firewall-1-INPUT -m state --state NEW -m
  tcp -p tcp --dport 2811 -j ACCEPT
```

```
-A RH-Firewall-1-INPUT -m state --state NEW -m
  tcp -p tcp --dport 20000:25000 -j ACCEPT
```

```
# REJECT all traffic coming from the privileged ports.
-A RH-Firewall-1-INPUT -p tcp -m state --state
  NEW -j REJECT
-A RH-Firewall-1-INPUT -p udp -m state --state
  NEW -j REJECT
-A RH-Firewall-1-INPUT -p icmp --icmp-type any -
  j REJECT
-A RH-Firewall-1-INPUT -j REJECT --reject-with
  icmp-host-prohibited
COMMIT
```



WMS Testing

- Login to the User Interface and create, on your account, the following two configuration files:

```
$ cat WMS-test.txt
[
  VirtualOrganisation = "gilda";
  NSAddresses = "<your_WMS>:7772";
  LBAddresses = "<your_WMS>:9000";
]
```



```
$ cat edg_wl_ui_cmd_var.conf
```

```
[  
rank = - other.GlueCEStateEstimatedResponseTime;  
requirements = other.GlueCEStateStatus == "Production";  
RetryCount = 3;  
ErrorStorage = "/tmp";  
OutputStorage = "/tmp/jobOutput";  
ListenerPort = 44000;  
ListenerStorage = "/tmp";  
LoggingTimeout = 30;  
LoggingSyncTimeout = 30;  
LoggingDestination = "<your_WMS>:9002";  
# Default NS logger level is set to 0 (null)  
# max value is 6 (very ugly)  
NSLoggerLevel = 0;  
DefaultLogInfoLevel = 0;  
DefaultStatusLevel = 0;  
DefaultVo = "gilda";  
]
```

```
$ cat hostname.jdl
[
    Executable="/bin/hostame";
    Arguments="-f";
    StdOutput="std.out";
    StdError="std.err";
    OutputSandbox={"std.out","std.err"};
]
```

```
$ voms-proxy-init --voms gilda
```

```
Your identity: /C=IT/O=GILDA/OU=Personal
Certificate/L=INFN Catania/CN=Giuseppe La
Rocca/Email=giuseppe.larocca@ct.infn.it
Enter GRID pass phrase:
Creating temporary proxy
..... Done
Contacting voms.ct.infn.it:15001
[/C=IT/O=GILDA/OU=Host/L=INFN
Catania/CN=voms.ct.infn.it/Email=emidio.giorgio@ct.infn.
it] "gilda" Done
Creating proxy
.....
Done
Your proxy is valid until Thu Sep 14 21:01:57 2006
```

Tests your grid node as follow:

```
$ edg-job-submit --config-vo WMS-text.conf --config
  edg_wl_ui_cmd_var.conf hostname.jdl
```

```
Selected Virtual Organisation name (from --config-vo option):
  gilda
```

```
Connecting to host giular.trigrid.it, port 7772
```

```
Logging to host giular.trigrid.it, port 9002
```

```
*****
```

JOB SUBMIT OUTCOME

```
The job has been successfully submitted to the Network Server.
Use edg-job-status command to check job current status. Your
job identifier (edg_jobId) is:
```

```
- https://giular.trigrid.it:9000/WYJViI71z3MphZI6P09RKQ
```

```
*****
```

```

$ edg-job-status
  https://giular.trigrid.it:9000/WYJViI71z3MphZI6P09RKQ

*****
BOOKKEEPING INFORMATION:
Status info for the Job :
  https://giular.trigrid.it:9000/WYJViI71z3MphZI6P09RKQ
Current Status:          Scheduled
Status Reason:           Job successfully submitted to Globus
Destination:             gildace.oact.inaf.it:2119/jobmanager---
reached on:              Mon Sep 25 16:41:31 2006
*****

```

```

$ edg-job-status
  https://giular.trigrid.it:9000/WYJViI71z3MphZI6P09RKQ

*****
BOOKKEEPING INFORMATION:
Status info for the Job :
  https://giular.trigrid.it:9000/WYJViI71z3MphZI6P09RKQ
Current Status:      Done (Success)
Exit code:           0
Status Reason:       Job terminated successfully
Destination:         gildace.oact.inaf.it:2119/jobmanager---
reached on:          Mon Sep 25 16:45:14 2006
*****

```

```

$ edg-job-get-output --dir .
  https://giular.trigrid.it:9000/WYJViI71z3MphZI6P09RKQ

Retrieving files from host: giular.trigrid.it ( for
  https://giular.trigrid.it:9000/WYJViI71z3MphZI6P09RKQ )

*****

                JOB GET OUTPUT OUTCOME

Output sandbox files for the job:
- https://giular.trigrid.it:9000/WYJViI71z3MphZI6P09RKQ
have been successfully retrieved and stored in the
  directory:
/home/larocca/larocca_WYJViI71z3MphZI6P09RKQ

*****

```



Log files & init scripts

- Log files can be found in `/var/edgw1/`

| | |
|-------------------------|----------------------|
| logging | proxyrenewal |
| logmonitor | SandboxDir |
| jobcontrol | networkserver |
| workload_manager | |

- They are the best starting point to deal with eventual problems. You can monitor and follow the jobs
- Init scripts can be found in `/etc/init.d/`

- **Add the resource to black-list to the file**
`/opt/edg/etc/edg_wl.conf`

```
WorkloadManager = [
PipeDepth = 1;
NumberOfWorkerThreads = 1;
DispatcherType = "filelist";
Input = "${EDG_WL_TMP}/workload_manager/input.fl";
LogLevel = 5;
LogFile =
    "${EDG_WL_TMP}/workload_manager/log/events.log";
MaxRetryCount = 10;
IsmBlackList={"grid010.ct.infn.it(.*)"};
];
```

And restart the service `/etc/init.d/edg-wl-wm restart`



Troubleshooting

If the `glite-job-submit/glite-job-list-match` commands returns the following error message:

```
**** Error: API_NATIVE_ERROR ****
```

```
Error while calling the "NSClient::multi" native api
```

```
AuthenticationException: Failed to establish security context...
```

```
**** Error: UI_NO_NS_CONTACT ****
```

```
Unable to contact any Network Server
```

it means that there are authentication problems between the UI and the *Network Server*

- **Check your Proxy.**
 - **Maybe you have not a valid proxy. Remember to initialized the proxy with the VOMS extensions.**

```
$_ voms-proxy-info --all
```

```
subject    : /C=IT/O=GILDA/OU=Personal Certificate/L=INFN
Catania/CN=Giuseppe La
Rocca/Email=giuseppe.larocca@ct.infn.it/CN=proxy
issuer     : /C=IT/O=GILDA/OU=Personal Certificate/L=INFN
Catania/CN=Giuseppe La Rocca/Email=giuseppe.larocca@ct.infn.it
identity   : /C=IT/O=GILDA/OU=Personal Certificate/L=INFN
Catania/CN=Giuseppe La Rocca/Email=giuseppe.larocca@ct.infn.it
type       : proxy
strength   : 512 bits
path       : /tmp/x509up_u512
timeleft   : 11:59:55
```



**No VOMS
extensions!**



Verify the synchronization between the UI and the WMS.

```
Check if ntpd is running
/etc/init.d/ntpd status
ntpd (pid 1742) is running...
```

and if the date is correctly !

- Inspect the log file `/var/edgwl/networkserver/log/events.log`

```
05 Sep, 16:44:31 -F- " NS::main ": -----
05 Sep, 16:44:31 -F- " NS::main ": Starting Network Server...
```

```
05 Sep, 16:44:31 -F- " NSR::drop ": Already running in an unprivileged
account.
```

```
05 Sep, 16:44:31 -F- " NSR::daemon ": daemon started.
```

```
05 Sep, 16:44:31 -F- " NSR::init ": Listening on port: 7772
```

```
05 Sep, 16:44:31 -M- " NSR::init ": Authentication Handshaking Enabling:
Ok.
```

```
05 Sep, 16:44:31 -F- " NS2WM::init ": Initializing NS2WM Proxy...
```

```
05 Sep, 16:44:31 -F- " NS2WM::init ": File queue
```

```
is:/var/glite/workload_manager/input.fl
```

```
05 Sep, 16:44:31 -F- " NS2WM::init ": NS2WMPProxy Initialization Done.
```

```
05 Sep, 16:44:31 -F- " NS::main ": Initialization Done. Running Server...
```

```
05 Sep, 16:44:31 -F- " NS::main ": -----
```

```
05 Sep, 16:50:06 -F- "Manager::run": Exception Caught during
AgentAuthentication.Failed to acquire credentials...
```

**No CRL Installed or
No CA supported !**

- Inspect the log file `/var/edgwl/networkserver/log/events.log`

05 Sep, 17:00:43 -F- "NS2WM::convertProtocol": Converted String: [

```
arguments =
[ ad = [ requirements = ( other.GlueCEStateStatus == "Production" ) && (
other.GlueCatStatus == "Production" )
RetryCount = 3; Arguments = "-f" JobType = "normal
"; Executable = "/bin/hostname"; CertificateSubject
="/C=IT/O=GRID/OU=error/Certificate/INFN Catania/CN=Giuseppe La
Rocca/Email=giuseppe.larocca@ct.infn.it" StdOut = "hostname.out";
X509UserProxy = "/tmp/user.proxy.0xb74f6768.20050905170043677437";
OutputSandbox = { "hostname.err", "hostname.out" };
VirtualObjectSetOfGrids = "grid" run = -
other.l... E... response time, type: 'job'; to: 'c:
"hostname.err";
```

No subject found inside the grid-mapfile of the WMS!

05 Sep, 17:01:49 -F- "Manager::run": **Exception Caught during Client Authentication.**

- **gLite v3.0 Advanced Installation and Configuration Guide**
 - http://glite.web.cern.ch/glite/packages/R3.0/R20060502/doc/installation_guide_3.0-2.html
- **GILDA gLite 3.0 installation notes wiki**
 - <https://grid.ct.infn.it/twiki/bin/view/GILDA/GliteElementsInstallation>
- **GILDA gLite-3.0 installation instructions**
 - https://gilda.ct.infn.it/docs/GILDAsiteinstall-3_0_0.html
- **Glite3.0 Installation experience by SEE (look for WMSLB):**
 - <http://wiki.egee-see.org/index.php/GLite30>

