

NetJobs: Network Monitoring Using Grid Jobs

Etienne Dublé - CNRS/UREC
Alfredo Pagano - GARR



- Network Monitoring...
 - In the context of grids
 - In the context of EGI
- The idea
- System architecture
 - Global view
 - The Server, the Jobs and the Grid
 - User Interface
- Next steps

Network Monitoring...

- In the context of grids
- In the context of EGI

- GRIDs are *big users* and they will exercise the network
 - The LHC generating ~15 PetaBytes of raw data/year for sure is a big user
- Grid middleware can benefit from monitoring:
 - Example: Network aware job and data transfer scheduling
- When a problem occurs, a grid operator / user would like to check quickly **if the network is involved in the problem:**
 - This is especially important for grids because in such a complex environment the network is one of many layers

- e2emonit (pingER, UDPmon, IPERF)
- NPM (Network Performance Monitor)
 - PCP (Probe Control Protocol)
- Diagnostic Tool
- PerfSONAR_Lite-TSS
- PerfSONAR-MDM

- The EGEE/EGI project did not recommend any specific solution for network monitoring
 - A part of the grid is already monitored (LHCOPN, specific national initiatives, ...), and there are plans to monitor more links
 - Monitor all Tier-1 <-> Tier-2 links using PerfSONAR?
- PerfSONAR Lite TSS is dedicated to troubleshooting
 - *In this project we are trying to address the needs which are **not yet addressed***

- Our approach had to take into account:
 - High scalability
 - Security
 - Reliability
 - Cost-effectiveness
- And preferably:
 - A lightweight deployment

The idea:

“Instead of installing a probe at each site, run a grid job”

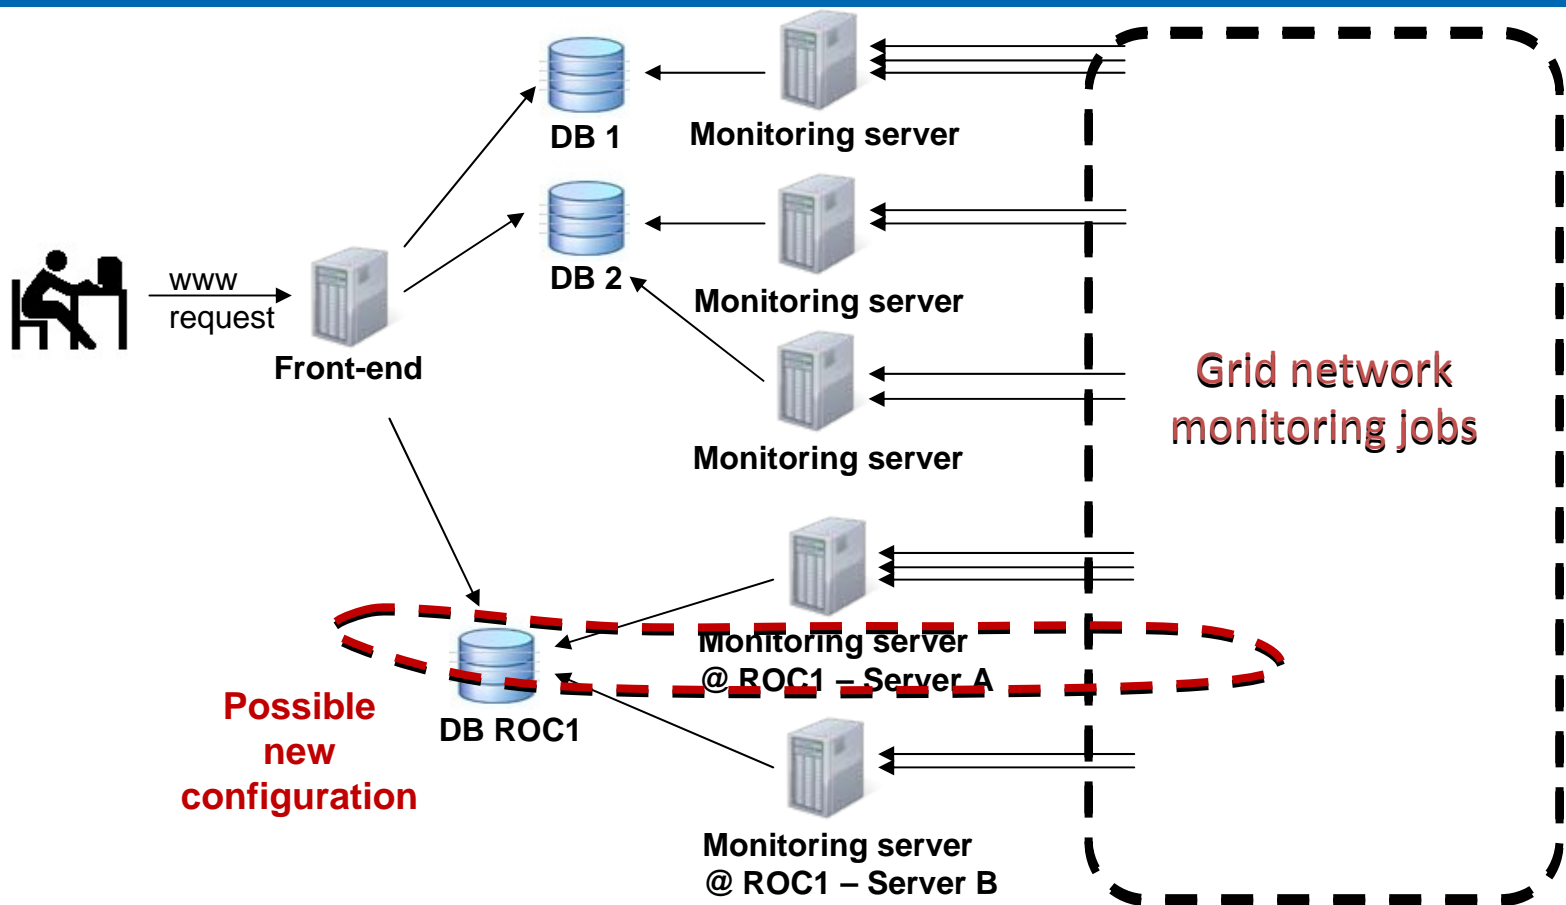
- Added value:
 - No installation/deployment needed in the sites
 - Monitoring 10 or 300 sites is just a matter of configuration
 - A monitoring system running on a proven architecture (the grid)
 - Possibility to use grid services (ex: AuthN and AuthZ)

- Limits:
 - Some low-level metrics can't be implemented in the job
 - Because we have no control of the “Worker Node” environment (hardware, software) where the job is running
 - Some sites will have to slightly update their middleware configuration
 - The maximum lifetime of jobs should be increased if it is too low (at least for the DN of the certificate that the system uses)

System architecture: Global view

System Architecture

the components

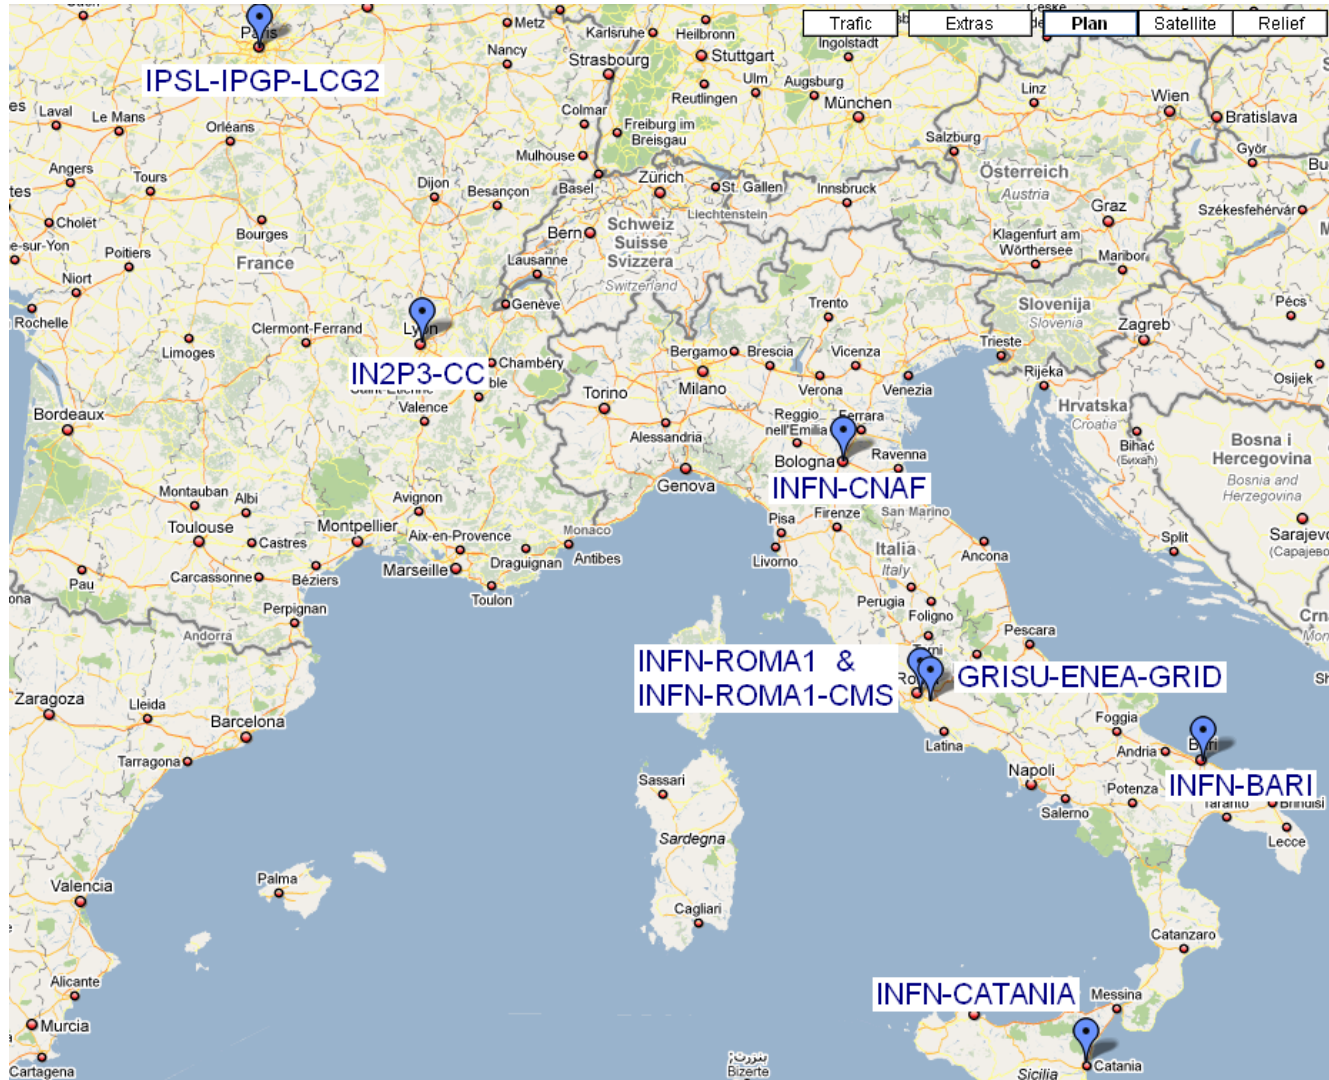


Frontend: Apache Tomcat, Ajax, Google Web Toolkit (GWT)

Monitoring server & Jobs: Python, bash script (portability is a major aspect for jobs)

Database: PostgreSQL

Current prototype: 8 Sites



- To Monitor all possible site-to-site paths will be too much: $N \times (N-1)$ paths
and $N \sim 300$ sites for a whole grid coverage
- We must restrict the number of these paths
 - To a specific VO, to an experiment, to the most used paths, etc.
 - We have studied this at <https://edms.cern.ch/document/1001777>

- The system is completely configurable about these paths and the scheduling of measurements
 - The admin specifies a list of scheduled tests, giving for each one
 - » The source and the remote site
 - » The type of test
 - » The time and frequency of the test
 - Users can contact and request the administrator to have a given path monitored (form available on the UI)
This request is then validated by the administrator.
- If you still have many paths, you can start several server instances (in order to achieve the needed performance)

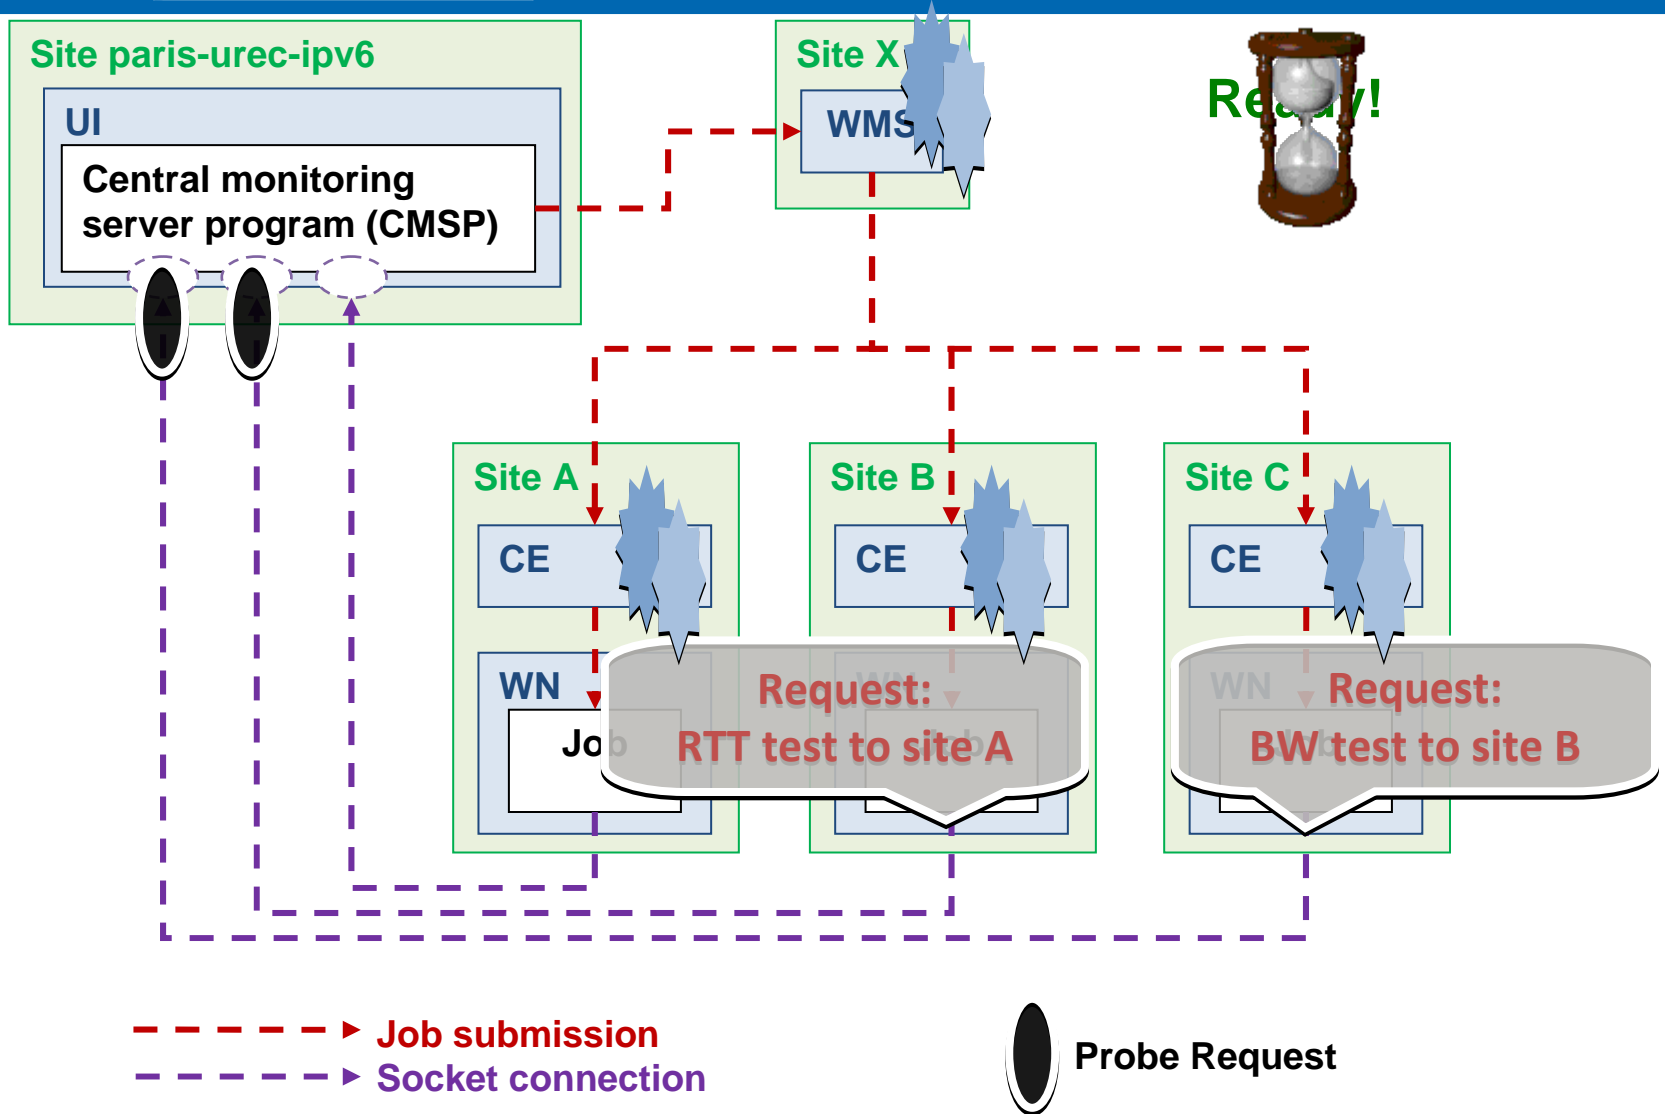
- Latency test
 - TCP RTT
 - Every 10 minutes
- Hop count
 - Iterative connect() test
 - Every 10 minutes
- MTU size
 - Socket (IP_MTU socket option)
 - Every 10 minutes
- Achievable Bandwidth
 - TCP throughput transfer via GridFTP transfer between 2 Storage Elements
 - Every 8h

In order to avoid too many connections these three measurements are done in the same test

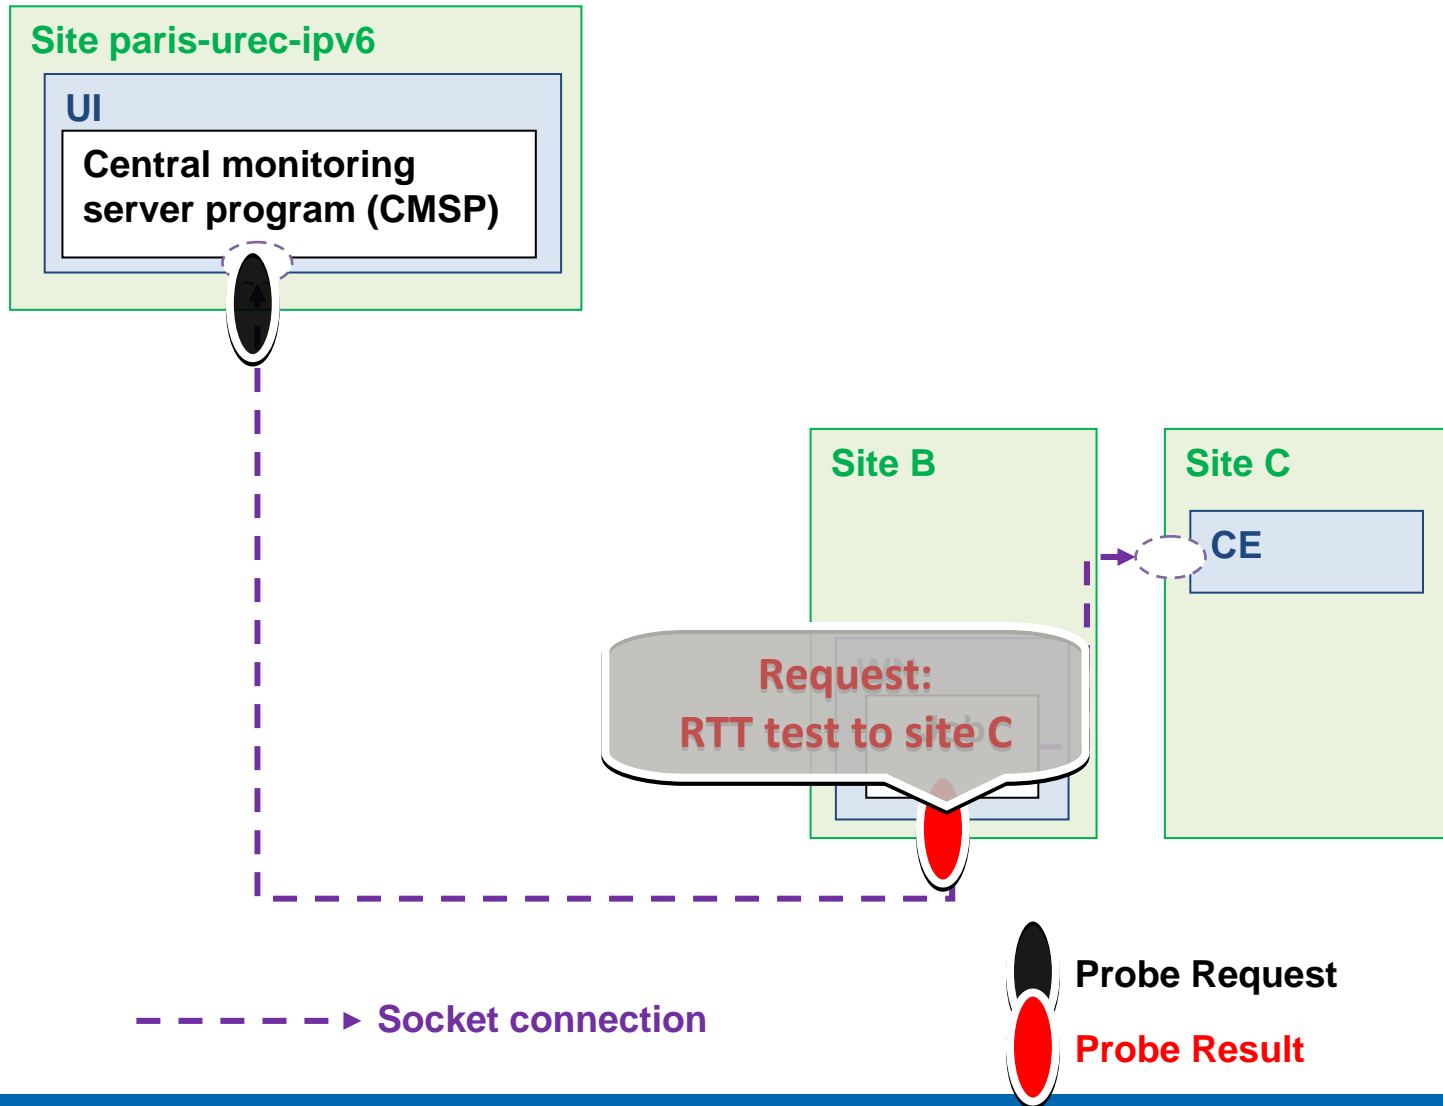
System architecture: The Server, the Jobs, and the Grid

- When running a job, the grid user is mapped to a Linux user of the Worker Node (WN):
 - This means the job is **not running as root** on the WN
 - Some low level operations are not possible (for example **opening an ICMP listening socket** is **not allowed**)
- **Heterogeneity** of the WN environments (various OS, 32/64 bits...)
 - Ex: making the job download and run an external tool may be tricky (except if it is written in an OS independent programming language)
- The system has to deal with the **grid mechanism overhead** (delays, job lifetime limit...)

Initialization of grid jobs



- Chosen design (**1 job <-> many probes**) is much more efficient than starting a job for each probe
 - Considering (grid-related) delays
 - Considering the handling of middleware failures (nearly 100% of failures occur at job submission, not once the job is running)
- TCP connection is initiated by the job
 - **No open port needed on the WN** → better for security of sites
- An **authentication mechanism** is implemented between the job and the server
- A job cannot last forever (GlueCEPolicyMaxWallClockTime), so actually there are *2 jobs running at each site*
 - A 'main' one, and
 - A 'redundant' one which is waiting and will become 'main' when the other one ends



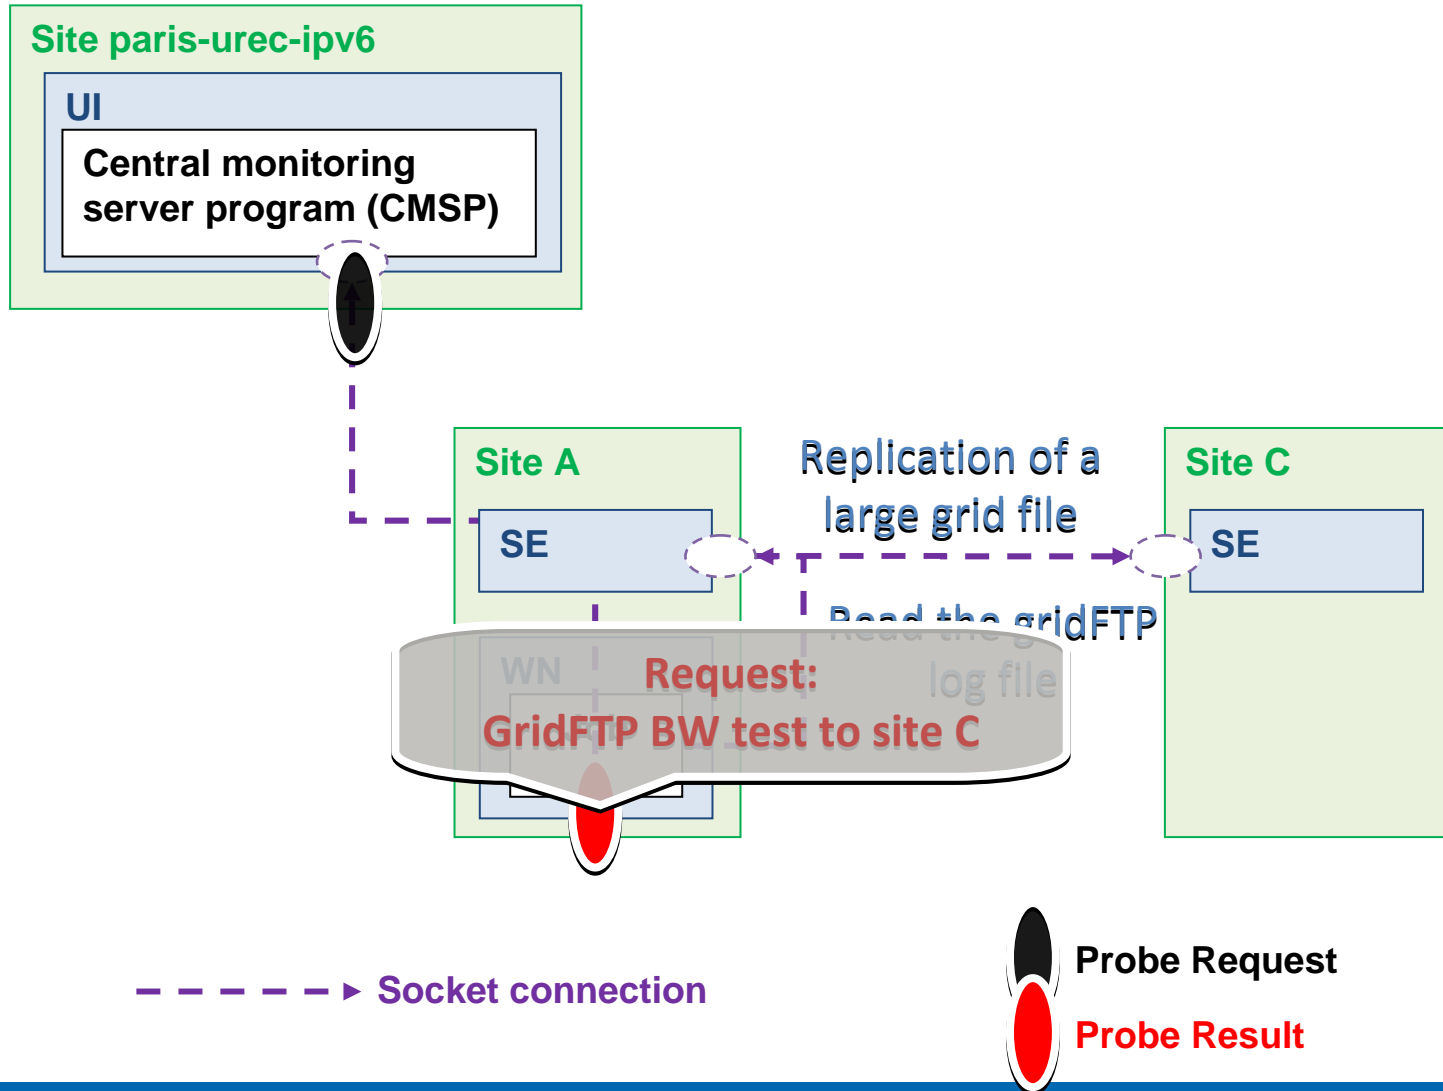
- The **RTT** measure is the time a **TCP 'connect()'** call takes:

- Because a connect() call involves a round-trip of packets:



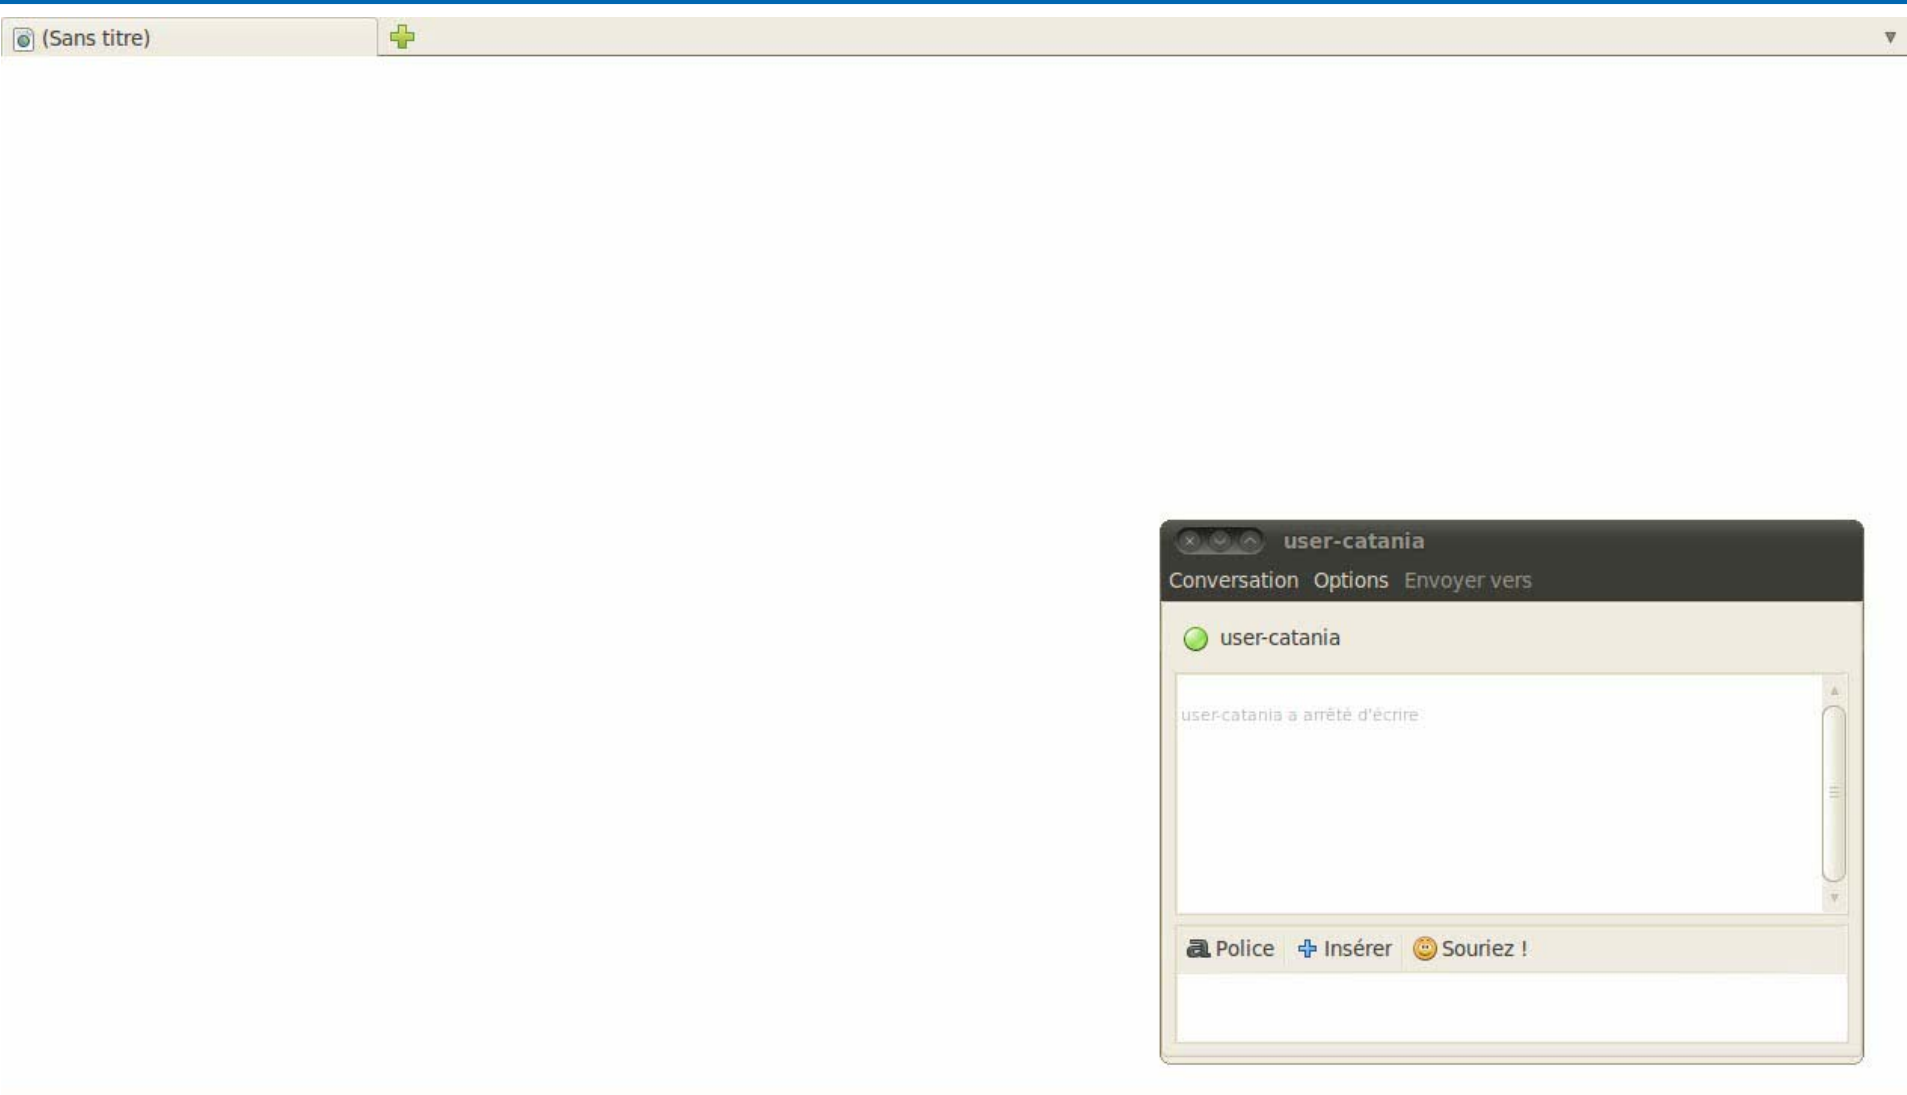
- Results very similar to the ones of 'ping'

- The **MTU** is given by the **IP_MTU socket option**
 - The **number of hops** is calculated in an iterative way
 - These measures require:
 - To connect to an accessible port (1) on a machine of the remote site
 - To close the connection (no data is sent)
 - *Note: This (connect/disconnect) is detected in the application log*
- (1): We use the port of the gatekeeper of the CE since it is known to be accessible (it is used by the grid middleware gLite)



- If the GridFTP log file is not accessible (cf. dCache?)
 - In this case we just do the transfer via globus-url-copy in a verbose mode in order to get the transfer rate.
- A passive version of this BW test is being developed
 - The job just reads the gridftp log file periodically (the system does not request additional transfers)
 - This is only possible if the log file is available on the Storage Element (i.e. it is a DPM)

System architecture: User Interface



NetJobs - A GRID Network Monitor based on GRID Jobs

Global DB

NetJobs A Grid Network Monitor based on Grid Jobs

Guest (user) << Measurements -

- Last 30 days
- Current Week
- Today

Jobs gLite ID + Logout +

Network Measurements (related to Global DB DB)

Search

Source site: infn-cnaf Destination site: infn-catania Probe Type: rtt From: 01/08/2010 00:00 To: 10/09/2010 10:41

Search Chart search results

Measurements

Date	Probe Type	Source Site	Dest Site	Result.
[-] Probe Type: rtt (100 Measurements)				
Fri Sep 10 2010 10:40:03 GMT+0200 (C	rtt	infn-cnaf	infn-catania	0.0249998569489
Fri Sep 10 2010 10:30:03 GMT+0200 (C	rtt	infn-cnaf	infn-catania	0.0250000953674
Fri Sep 10 2010 10:20:07 GMT+0200 (C	rtt	infn-cnaf	infn-catania	0.0249998569489
Fri Sep 10 2010 10:10:03 GMT+0200 (C	rtt	infn-cnaf	infn-catania	0.0250000953674
Fri Sep 10 2010 10:00:03 GMT+0200 (C	rtt	infn-cnaf	infn-catania	0.0249998569489
Fri Sep 10 2010 09:50:03 GMT+0200 (C	rtt	infn-cnaf	infn-catania	0.0250000953674
Fri Sep 10 2010 09:40:03 GMT+0200 (C	rtt	infn-cnaf	infn-catania	0.0249998569489
Fri Sep 10 2010 09:30:03 GMT+0200 (C	rtt	infn-cnaf	infn-catania	0.0250000953674
Fri Sep 10 2010 09:20:03 GMT+0200 (C	rtt	infn-cnaf	infn-catania	0.0249998569489

Page 1 of 52 Showing measurements 1 - 100 of 5176

Next steps

1. Near future:

- ✓ GridFTP passive BW test
- ✓ Email alerts



2. Other possible enhancements:

- ✓ Refresh measurements on-demand
(don't wait several hours for the next bw test...)
- ✓ Add more types of measurements?
- ✓ Consider adding a dedicated box (VObox?)
 - If some of the metrics needed are not available with the job-based approach
Ex: low level measurements requiring root privileges
 - The job would interact with this box and transport the results
 - This might be done in a restricted set of major sites
- ✓ Consider interaction with other systems (some probes may be already installed at some sites, we could benefit from them)



Thank You

Feedback, discussion, requests...

<http://netjobs.dir.garr.it/>

Wiki: <https://twiki.cern.ch/twiki/bin/view/EGI/GridNetworkMonitoring>

Contacts: etienne.duble@urec.cnrs.fr
alfredo.pagano@garr.it