EGEE,06 Conference
Geneva, Switzerland September 25-29 2006

AFS pool account users and GRID interoperability

G. Bracco, P. D'Angelo, L. Giammarino*,
S.Migliori, A. Quintiliani, C. Scio**, F. Simoni

ENEA INFO, C.R. ENEA Frascati
V. E. Fermi 45 Frascati ROMA (Italy)
bracco@frascati.enea.it

(*) CASPUR, Roma Italy, (**) Esse3Esse, Roma, Italy

Outline

- AFS Users as Pool Account Users:

  modified software components for compatibility with AFS

  - Token required for locally mapped users: lcmaps
  - Token granting for VO users: gssklog/gssklogd

- ENEA-GRID and EGEE/EGEEII:

  gateway implementation for access to AIX resources

- Status and Discussion

lcmaps (1)

LCAS [Local Centre Authorization Service] and LCMAPS [Local Credential Mapping Service] are the EGEE services that make the connection between GRID users certificates and local UNIX userids on the local site. The systems support:

- The standard globus X509 certificate
  - gridmap file: [*distinguished_name local_userid*]
    "/O=dutchgrid/O=users/O=sara/CN=Dutch User" griduser
- The extension of X509 certificate holding VO information [VOMS service]
  - gridmap file: [*VO_data voms_pool_account_prefix*]
    "/VO=enea/GROUP=/enea" .enea

"voms pool account": predefined local users which are dynamically assigned to the members of the VO when running job on the site.
[e.g. for VO enea pool users are enea001, enea002, ..., enea020]

pool accounts can be used also with standard X509 certificates:
"/O=dutchgrid/O=users/O=sara/CN=Dutch User" .dteam

lcmaps (2)

LCMAPS [Local Credential Mapping Service] is a modular service.

A lcmaps_afs module exists with the purpose to obtain an AFS tokens for the local mapped user using gssklog but the standard module is not compatible with VOMS "pool account"

A modified module has been prepared to achieve compatibility:
gssklog is invoked always with -principal parameter

Note: due to the software architecture of the CE the relevant user processes on the CE can not be started in the same PAG shell; in fact they are managed by two independent services:

- authorization and user mapping are managed by lcas/lcmaps
- job submission is managed by the edg-gatekeeper

The reliability of a "user based" token instead that a "PAG based" token is delegated to the CE middle-ware architecture.

gssklogd/gssklog (1)

Gssklog package has been developed by Doug Engert (ANL) to obtain an AFS token using a GSS [Generic Security Service] implementation.

It has been used together with Globus Toolkit GSSAPI to obtain an AFS token from a X509 certificate.ftp://achilles.ctd.anl.gov/pub/DEE

The package include a daemon: [gssklogd] and a client [gssklog].

gssklogd runs on a server where the AFS cell keyfile is available and requires a service certificate and key. A map file is also required /etc/grid-security/afsgrid-mapfile containing the X509 distinguished name and the AFS userid(s):

"/O=Grid/O=Globus/OU=anl.gov/CN=John Doe" jdoe,altuser

The client gssklog requires the proxy file generated by the globus command grid-proxy-init; an AFS token is obtained by the command:

gssklog -server *server_address* -principal jdoe

gssklogd/gssklog (2)

EGEE uses of an extension of X509 certificate to hold VO information [VOMS] and the proxy is generated using command:

voms-proxy-init -voms *VO_name*

gssklogd has been modified to provide tokens for pool account users on the basis of the content of the VO information contained in the user certificate. The daemon requires now also the VO certificate.

The afsgrid-mapfile syntax has been extended, e.g. for ENEA VO:

"/VO=enea" .enea/020/

The client gssklog provides always the principal name:

gssklog -server *server_address* -principal enea004

Some users requires pool account without having VOMS extension

"/O=dutchgrid/O=users/O=sara/CN=Dutch User" .dteam

Status

- An EGEE Technical Note has been written and submitted to JRA1: Title “AFS Pool Account users”
- The patched source of lcmaps_afs module [lcmaps_afs.c] is available
- A patch to the standard gssklog-0.11 package is available.
- The patch to gssklog package has been presented to the AFS Best Practices Workshop 2006 [http://pmw.org/afsbpw06] and submitted to Doug Engert.
- Both components have been used in implementing the ENEA-INFO EGEE site, providing access to AIX resources from ENEA GRID infrastructure

ENEA

[Italian National Agency for New Technologies, Energy and Environment]

12 Research sites and a Central Computer and Network Service (ENEA-INFO) with 6 computer centres managing multi-platform resources for serial & parallel computation and graphical post processing.

ENEA GRID

INFO-ENEA computational resources:

- Hardware: ~100 hosts and ~650 cpu : IBM SP; SGI Altix & Onyx; Linux clusters 32/ia64/x86_64; Apple cluster; Windows servers. Most relevant resources: IBM SP5 192 cpu; 3 frames of IBM SP4 96 cpu
- software: commercial codes (fluent, ansys, abaqus..); elaboration environments (Matlab, IDL, SAS..)

ENEA GRID mission [started 1999]:

- provide a unified user environment and an homogeneous access method for all ENEA researchers, irrespective of their location.
- optimize the utilization of the available resources

ENEA GRID architecture

GRID functionalities (unique authentication, authorization, resource access and resource discovery) are provided using “mature”, multi-platform components:

Distributed File System: OpenAFS

Resource Manager: LSF Multicluster [www.platform.com]

Unified user interface: Java & Citrix Technologies

These components constitute the ENEA-GRID Middleware.

OpenAFS

- user homes, software and data distribution
- integration with LSF
- user authentication/authorization [still kerberos 4]

ENEA EGEE site implementation (1)

ENEA implementation requirements:

- provide access to the main ENEA-GRID resourses
- EGEE users => mapped to some standard ENEA-GRID AFS users
- Access to AIX platform => NO Middleware on Worker Nodes
- LSF resources must be used whenever possible
  - Batch submission: bsub
  - Information: lsload, bhosts, lshosts, bqueues,...
  - Prompt job execution: lsrun
  - Batch queues dedicated to EGEE GRID users

ENEA EGEE site implementation (2)

Implementation architecture:

- UI & SE: Linux, Standard EGEE configuration
- CE: Linux, modifications to implement a “gateway”
- A “Transmitter” linux element with WN packages installed
- WN: any platform/OS with support for AFS/LSF
  - share user homes with CE/Transmitter using AFS
  - delegate all grid commands concerning file transfer to/from GRID to the Transmitter by means of the LSF lsrun command.
  - AFS security and quota management guarantee reliability

ENEA EGEE site implementation (3)

How:

- Any WN middleware command concerning file transfer is wrapped with a lsrun script.

- The wrapped WN middleware is located in AFS.
- The relevant PATHs in the environment of the EGEE job on the WN are modified so that the wrapped middleware is used
  - LSF job-starter script for the EGEE dedicated queues

- EGEE GRID users require an AFS token:
  - gssklogd modified to be compatible with EGEE middleware
  - EGEE middleware on the CE must properly call gssklog.

EGEE Standard site layout

User Interface

Storage Element

Computing Element: accepts jobs sent by the RB and submits them to the local batch system (LSF)

Worker Nodes: they perform the computation and the WN middleware sends back job results to RB

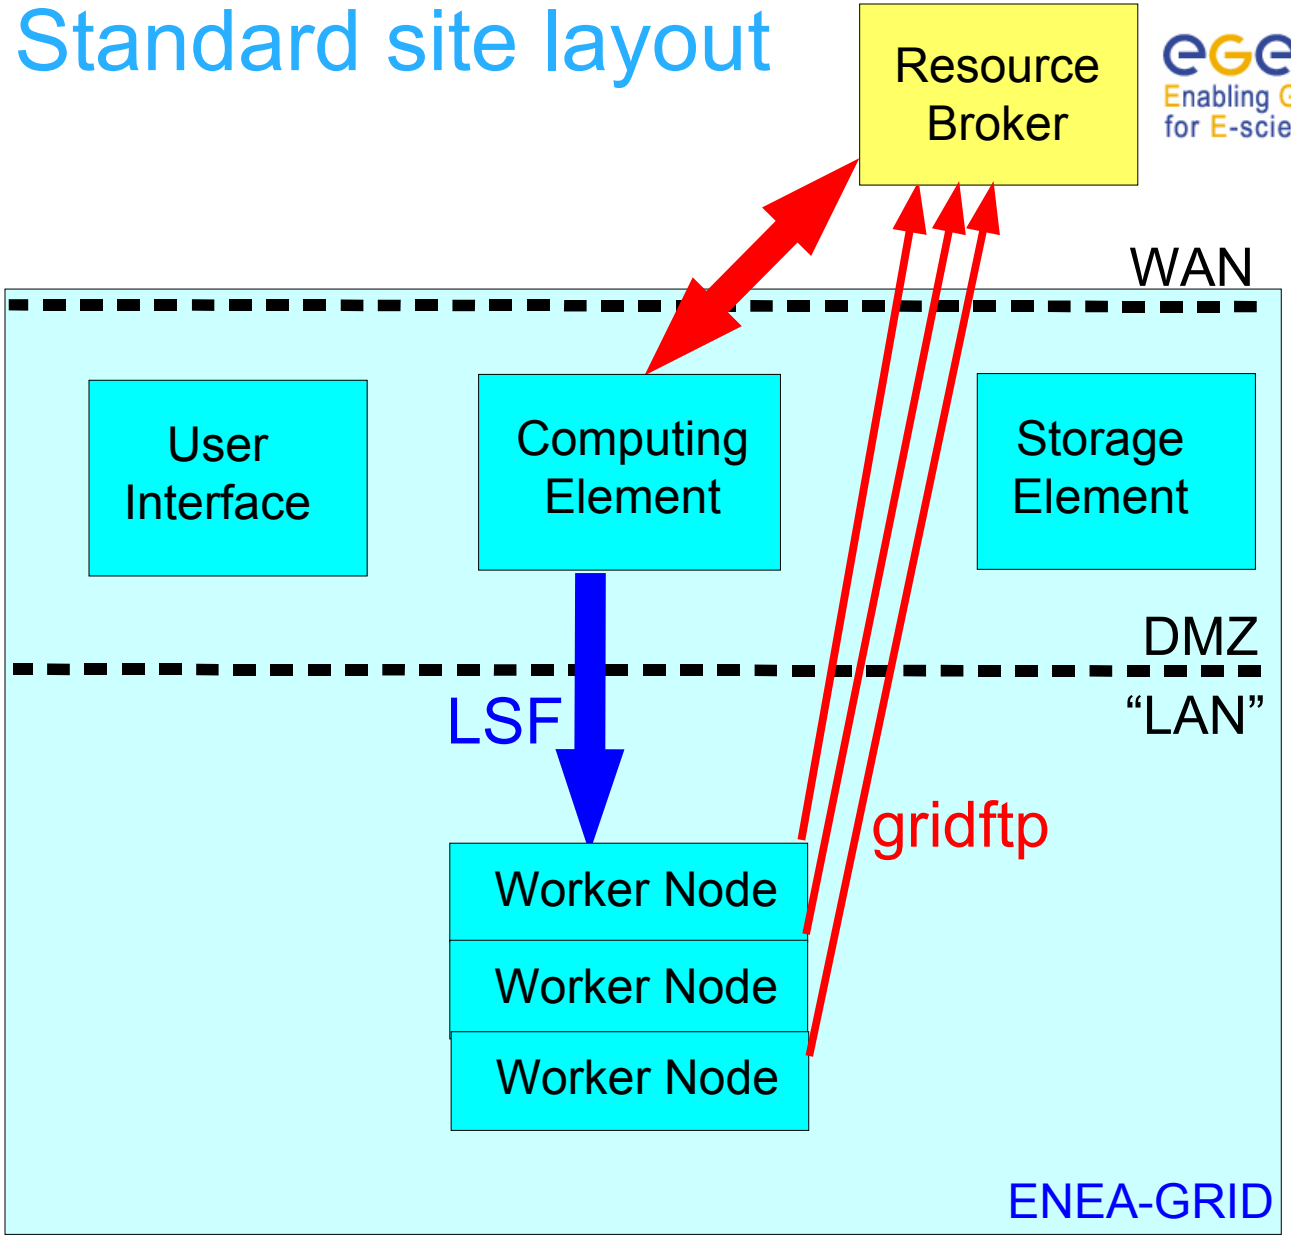

EGEE ENEA site implementation (1)

The Computing Element: a gateway system where all the grid commands are executed.

Worker Nodes: perform the computation and ask the CE to sends back job results to RB.

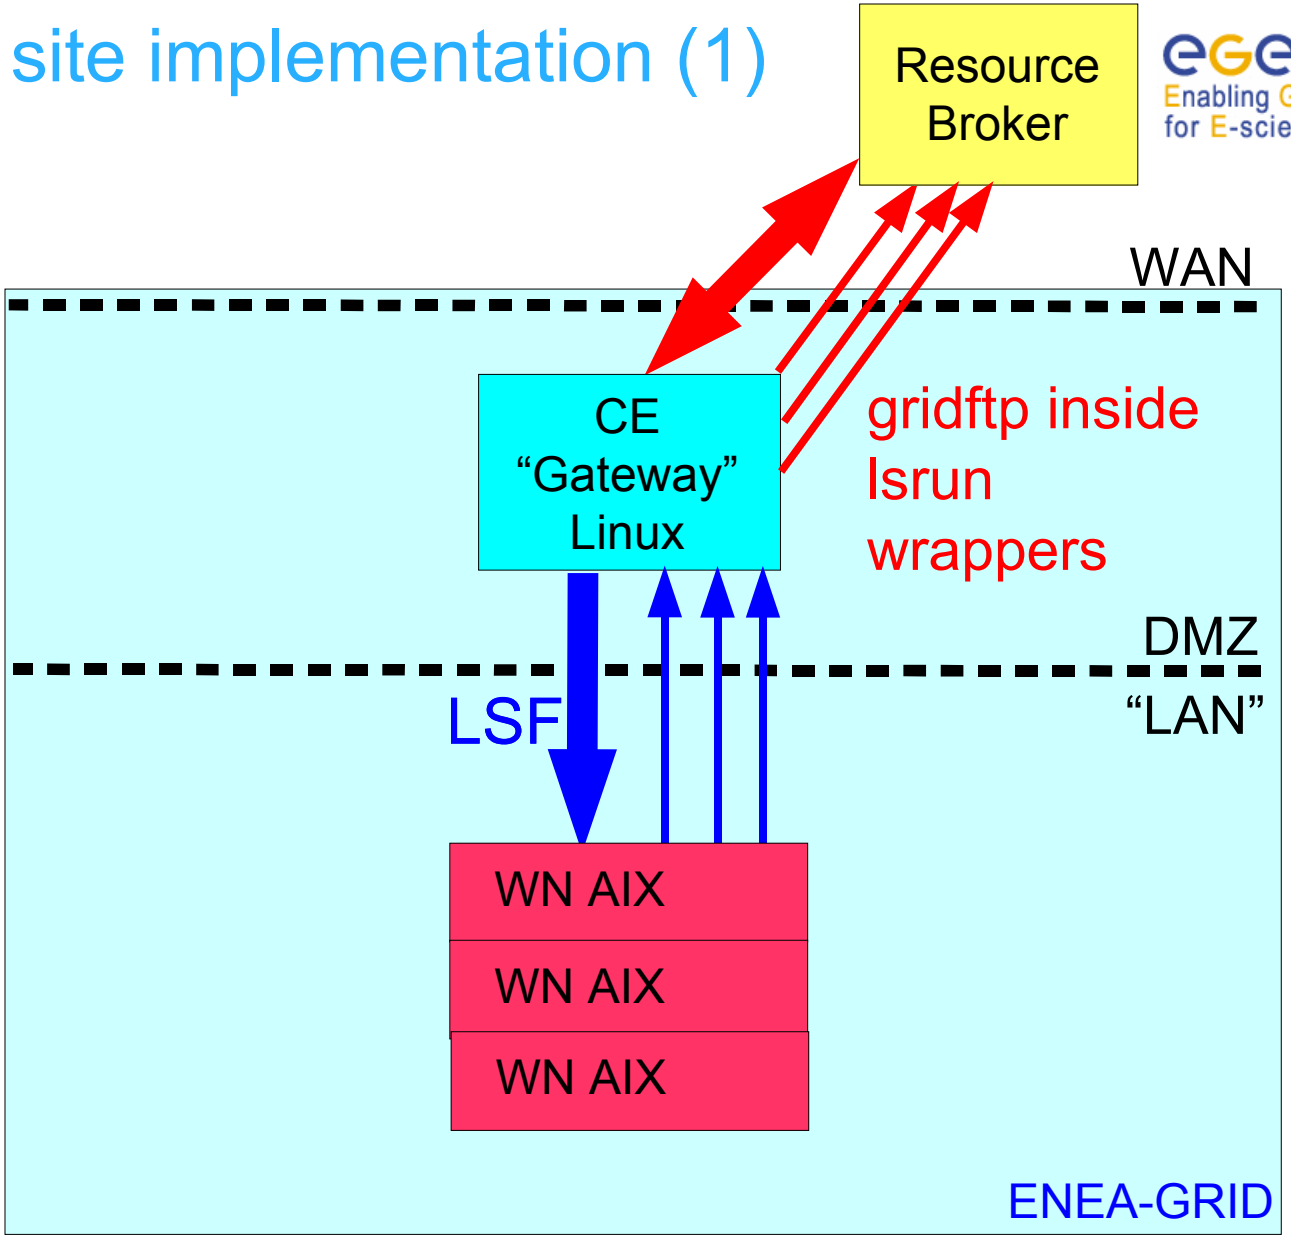

EGEE ENEA site implementation (3)

To simplify the installation a "Transmitter" Linux WN has been added

Grid data transfer commands are:

1) taken from AFS
2) wrapped with lsrun

e.g. globus-url-copy-> wrapper
wrapper: lsrun -m "transmitter" globus-url-copy $*

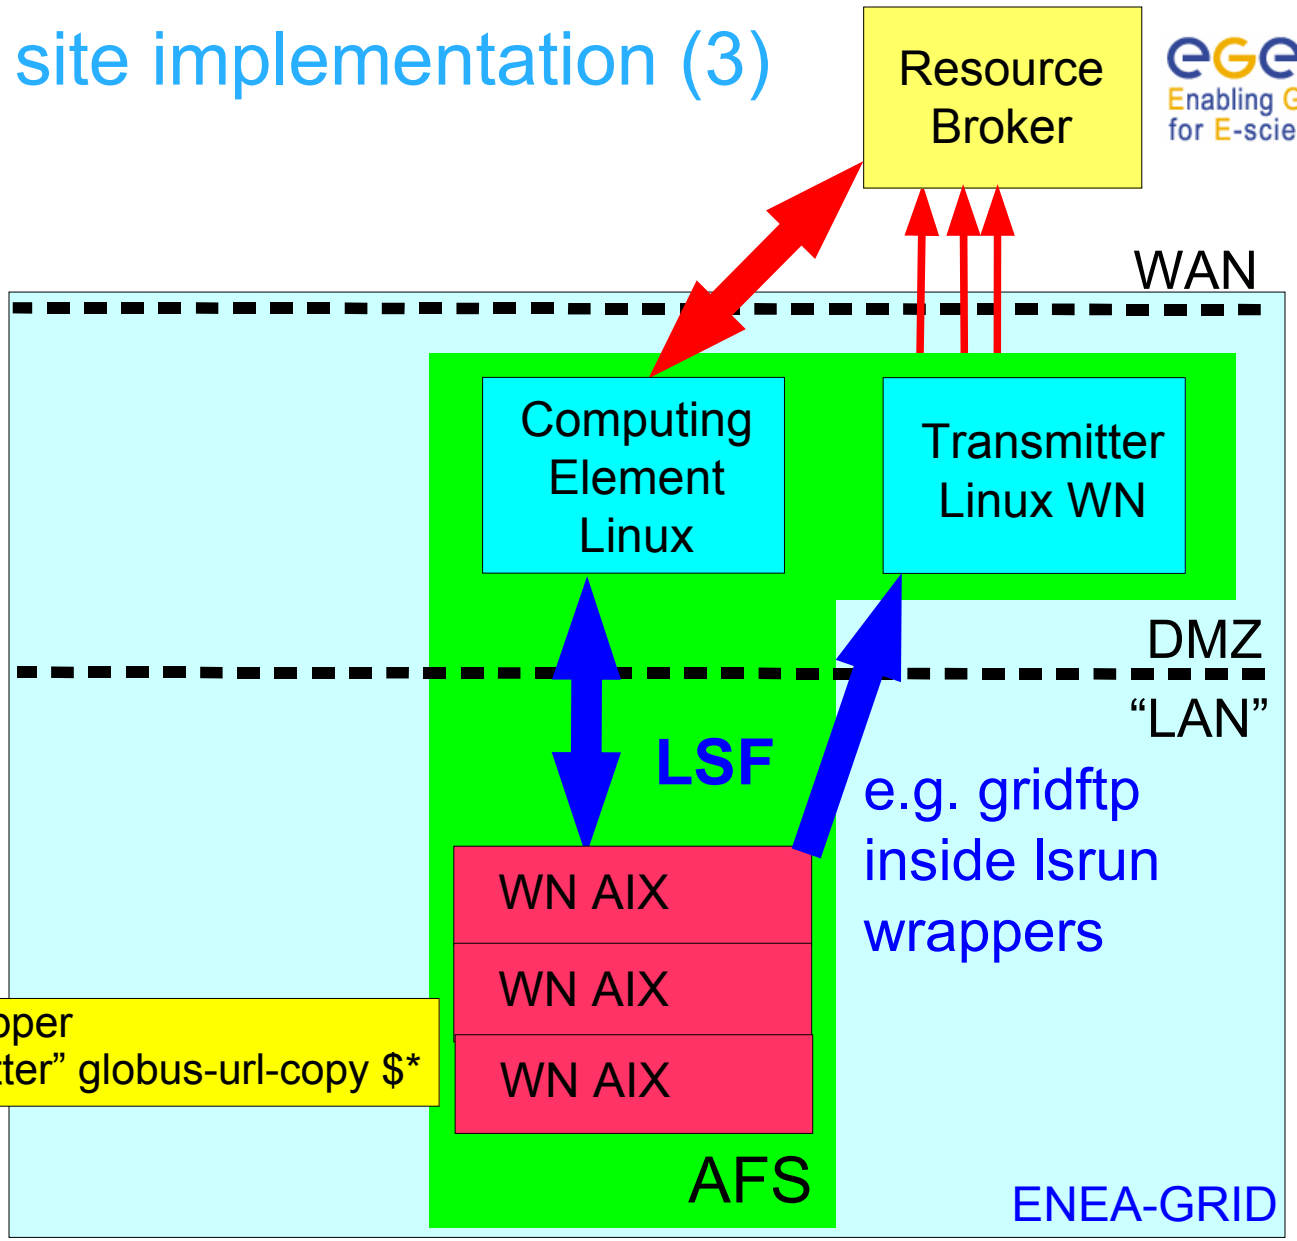

EGEE ENEA site implementation (4)

GRID users are AFS users so a modified gssklogd server is required to obtain an AFS token from the X509 certificate.

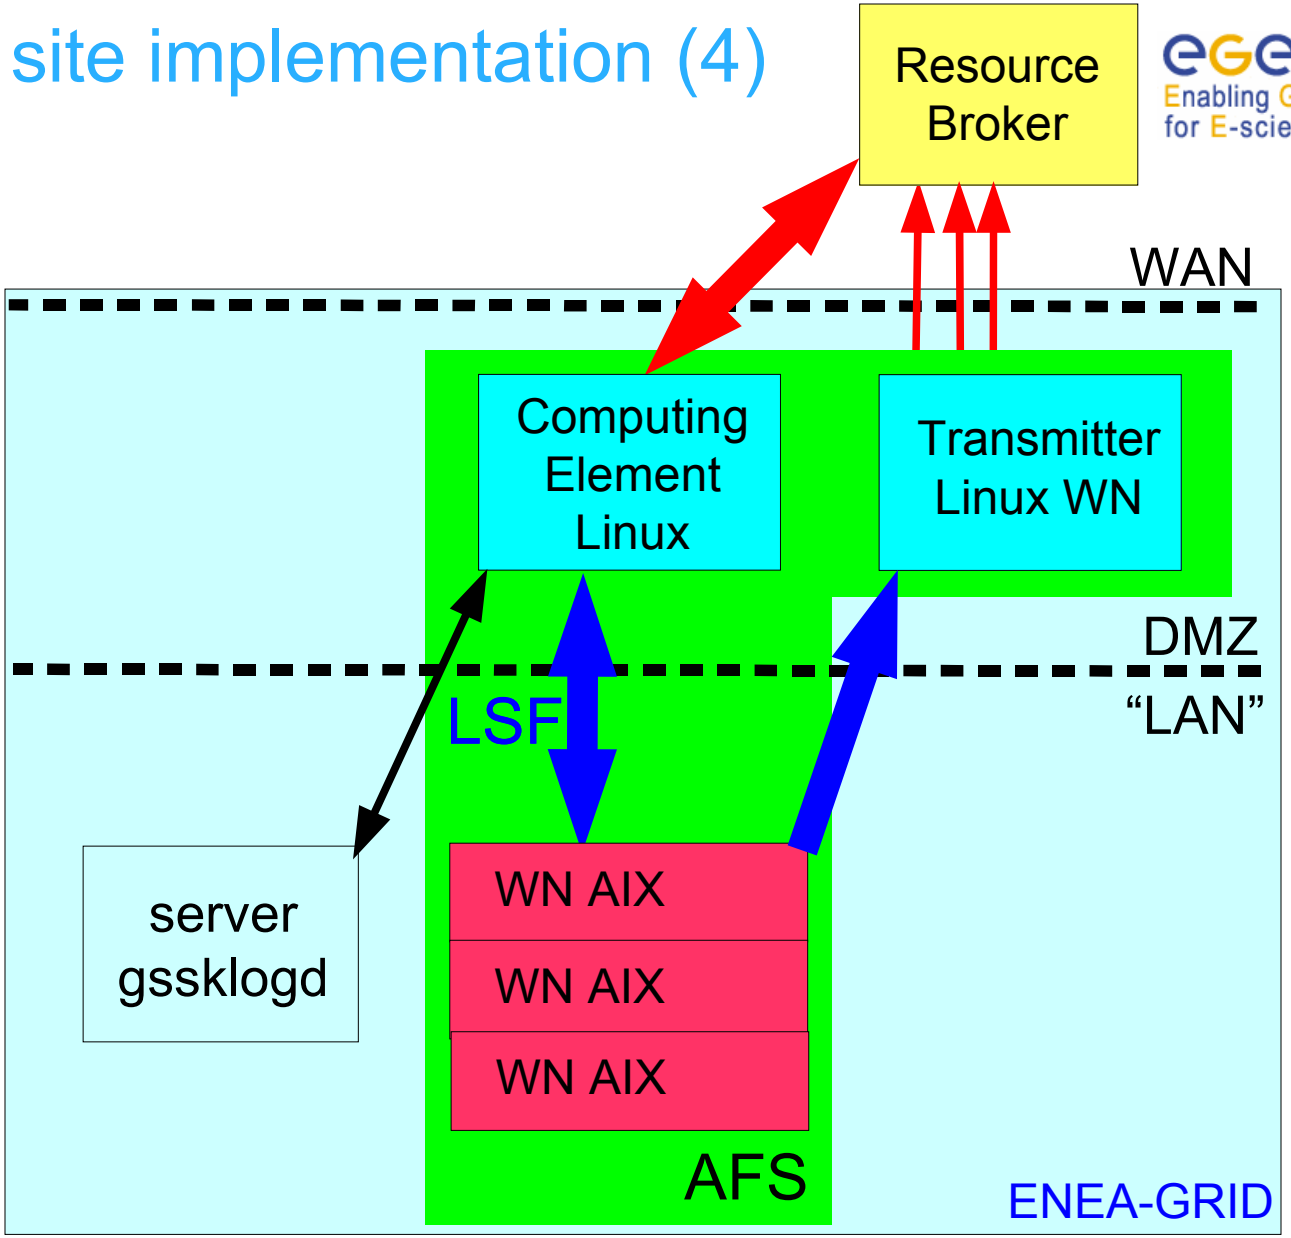

EGEE ENEA site implementation (5)

GRID users are AFS users so a modified gssklogd server is required to obtain an AFS token from the X509 certificate.

AFS can be used to provide to the gssklogd daemon the information about VO and CA

These information are automatically updated by the CE

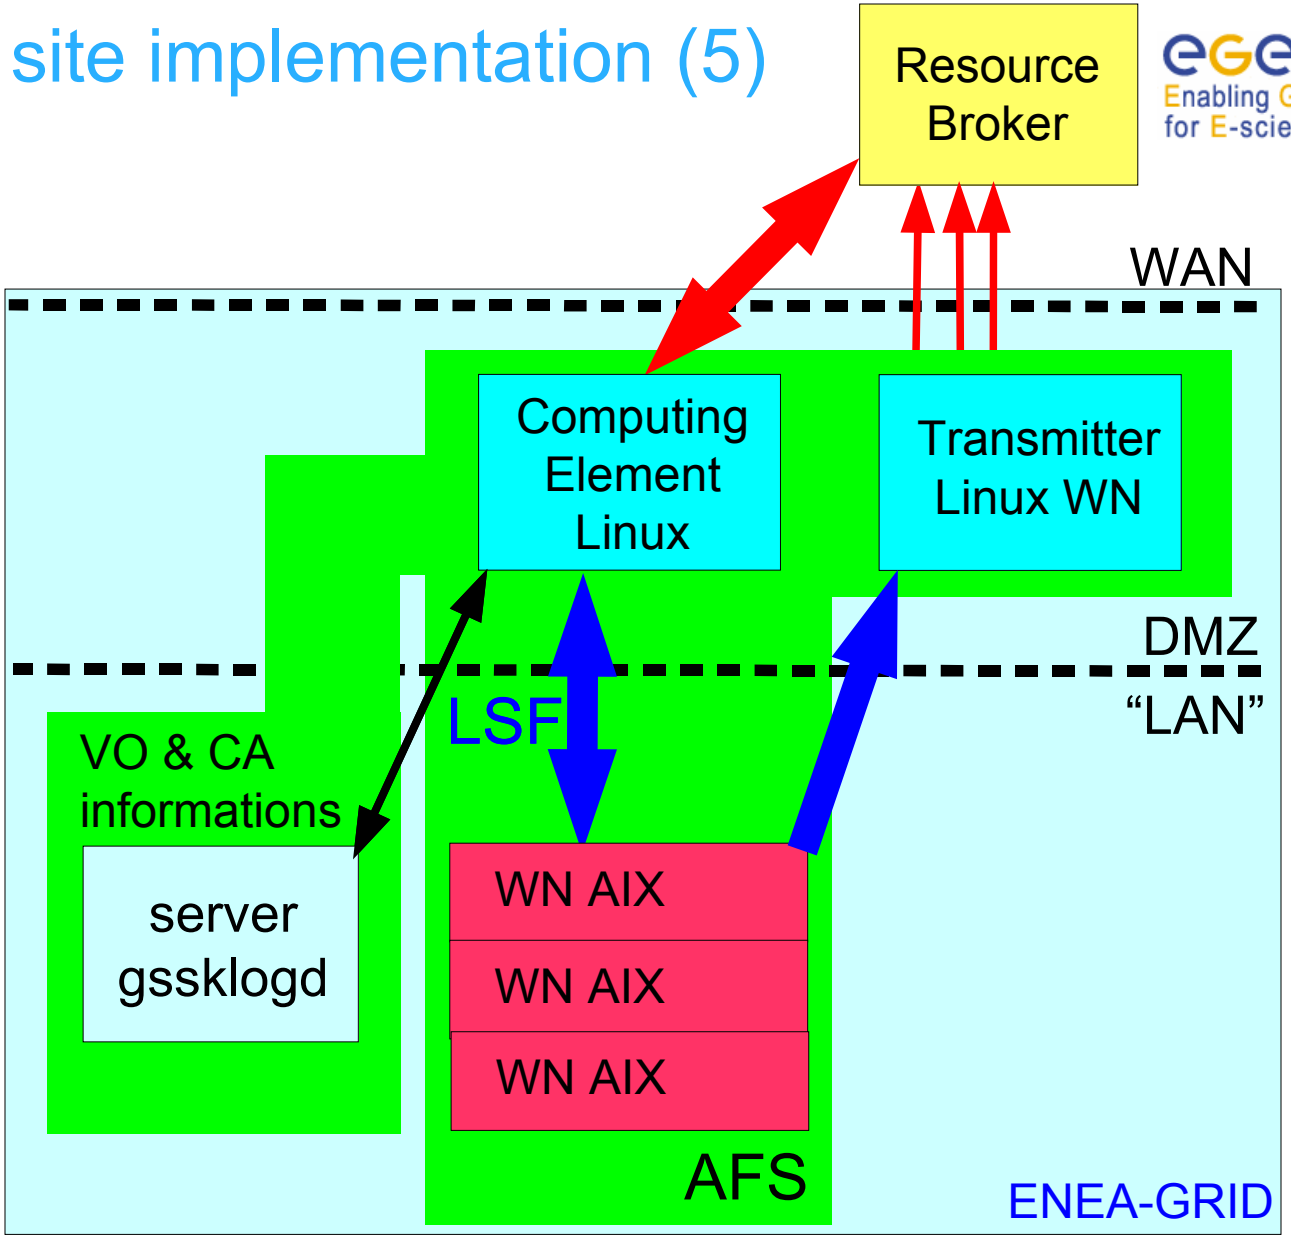

Status and discussion

ENEA Linux site: Glite 3.0 certified

CE/SE: egce.frascati.enea.it / egse.frascati.enea.it
WN: 14 Linux P4, 1.8 Ghz, 1 GB

ENEA AIX site: Glite 3.0 certification in progress

CE/SE: egceaix.frascati.enea.it / egseaix.frascati.enea.it
WN: 96 cpu SP4, 1.9 Ghz

BENEFITS of the gateway approach: Any Platform/OS with AFS/LSF

LIMITS of the gateway approach

- Not a completely standard site [but EGEE Certification job runs well]
  - GRID API are not available [How to "Flag" it ?]
  - Some WN monitoring components are unavailable
- Scalability : the number of "transmitter" can be increased

Conclusions and acknowledgements

Conclusions

- Patched versions of gssklog and lcmaps are available to enable AFS user as Pool Account local user
- A “gateway” approach provides access to unsupported platforms/OS

What next

- Prepare a Technical Report about the “gateway” implementation
- Add other ENEA-GRID platforms: ALTIX, Opteron Cluster, MacOSX (successful tests already performed).
- Application demonstration [FUSION VO]; better MPI support

ENEA GRID is operated with the support of many people in various ENEA sites: S. Taglienti, R. Guadagni, A. Perozziello, A. De Gaetano, S. Pecoraro, D. Giammattei, G. Mencuccini., M. De Rosa, M. Caiazzo, A. Palumbo, G. Elmo, S. Pierattini, M. Impara, G. Furini, C. Zini...

Many thanks to Vincenzo Ciaschini [INFN, CNAF, Bologna] for helping us with VOMS libraries utilization.