



Grid Security, an introduction

- Introduction to security issues in grid infrastructures and cross-domain user collaborations



David Groep, Nikhef



Grid Security dealing with user-centric collaborations



V Grid: following research collaborations

Some things that may make a grid a bit 'special' compared to other distributed computing efforts

- > collaboration of individuals from different organisations
 - > most of the scientific grid communities today consist of people literally 'scattered' over many home organisations ... internationally
- > delegation – programs and services acting on your behalf – are an integral part of the architecture
 - > unattended operation
 - > resource brokering
 - > integrating compute, data access, and databases in the same task

V But ... what is Grid?

The word '*grid*' has been used in many ways

- > cluster computing
- > cycle scavenging
- > **cross-domain resource sharing**
- > ...

A clear definition for the grid?


- > Coordinates resources not subject to centralised control
- > Using standard, open and generic protocols & interfaces
- > Provides non-trivial qualities of collective service

www.ogf.org

Definition from Ian Foster in *Grid Today*, July 22, 2002; Vol. 1 No. 6,
see <http://www-fp.mcs.anl.gov/~foster/Articles/WhatIsTheGrid.pdf>

V Example: a biomedical imaging project

- > On functional MRI studies run from a 'standardized' workflow
- > People and systems involved (the 'vlemed' VO)
 - > medical doctors and the fMRI apparatus: AMC hospital
 - > data storage service: SARA Compute and Network services
 - > Compute services: Nikhef, Philips Research, SARA
 - > algorithm developers: University of Amsterdam
 - > Medical doctors and analysts (MD): AMC



SP1.3 Medical Imaging
simplified user scenario



Typical use case: WISDOM

Wide-area In-Silico Docking On Malaria

- > people and organisations
 - > Bio-informaticians and grid development: IN2P3 (FR)
 - > Service systems (brokers) provided by: RAL (UK), NIKHEF (NL)
 - > algorithms, and results analysed by: SCAI (DE)
 - > Compute resources: provided by over 45 independent organisations in ~15 countries, whose primary mission is usually HE Physics!
 - > VO management hosted by CERN (CERN), and the VO itself is managed by Vincent Breton (FR)

wLCG: implementing LHC computing


20 years est. life span
24/7 global operations
~ 4000 person-years of
science software investment

~ 5 000 physicists
~ 150 institutes
53 countries/economic regions



V Virtual Organisation

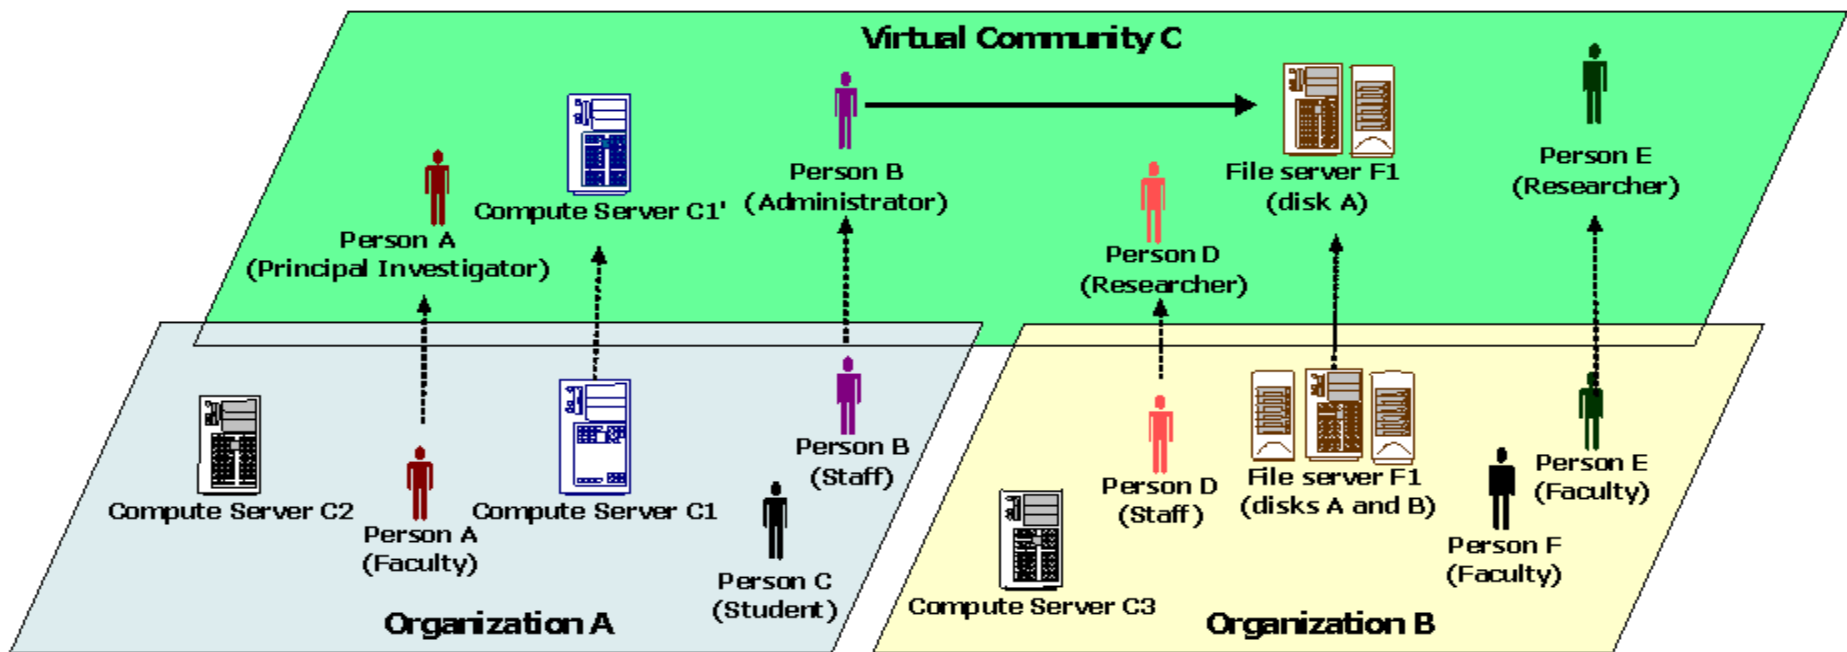
A set of individuals or organisations, **not under single hierarchical control**, (temporarily) **joining forces** to solve a particular problem at hand, bringing to the collaboration a subset of their resources, sharing those **at their discretion** and each **under their own conditions**.



- User driven
- Users are usually a member of more than one community
- Any “large” VO will have an internal structure, with groups, subgroups, and various roles

V Virtual vs. Organic structure


- > Virtual communities (“virtual organisations”) are many
- > An person will typically be part of many communities
 - > has different roles in different VO’s (distinct from organisational role)
 - > all at the same time, at the same set of resources, with SSO



graphic: OGSA Architecture 1.0, OGF GFD-I.030


V Before and parallel to the Grid ...

- > Each user in a collaboration gets *individual* access to many or most of the ICT resources of all participating groups
 - > Shared group accounts
 - > Individual accounts with the same name (and password)
 - > Permissive password sharing
- > Characteristics
 - > Gets more access than needed
 - > No centralized management
 - > Easy 'hopping' between sites, also for attackers ... !



Grid 'VOs': structuring communities on a sustainable infrastructure

- > Virtual Organisations as groupings of users
 - > E-infrastructures (EGI, BiG Grid) provide persistent infrastructure with a "bus-like" view for VOs: essentially *user communities*



Communities can exist without their 'own' resources and resource centres can do without local users



Granting Access

Policy framework

Authentication

Authorization and Virtual Organisation membership

GRID SECURITY MECHANISMS

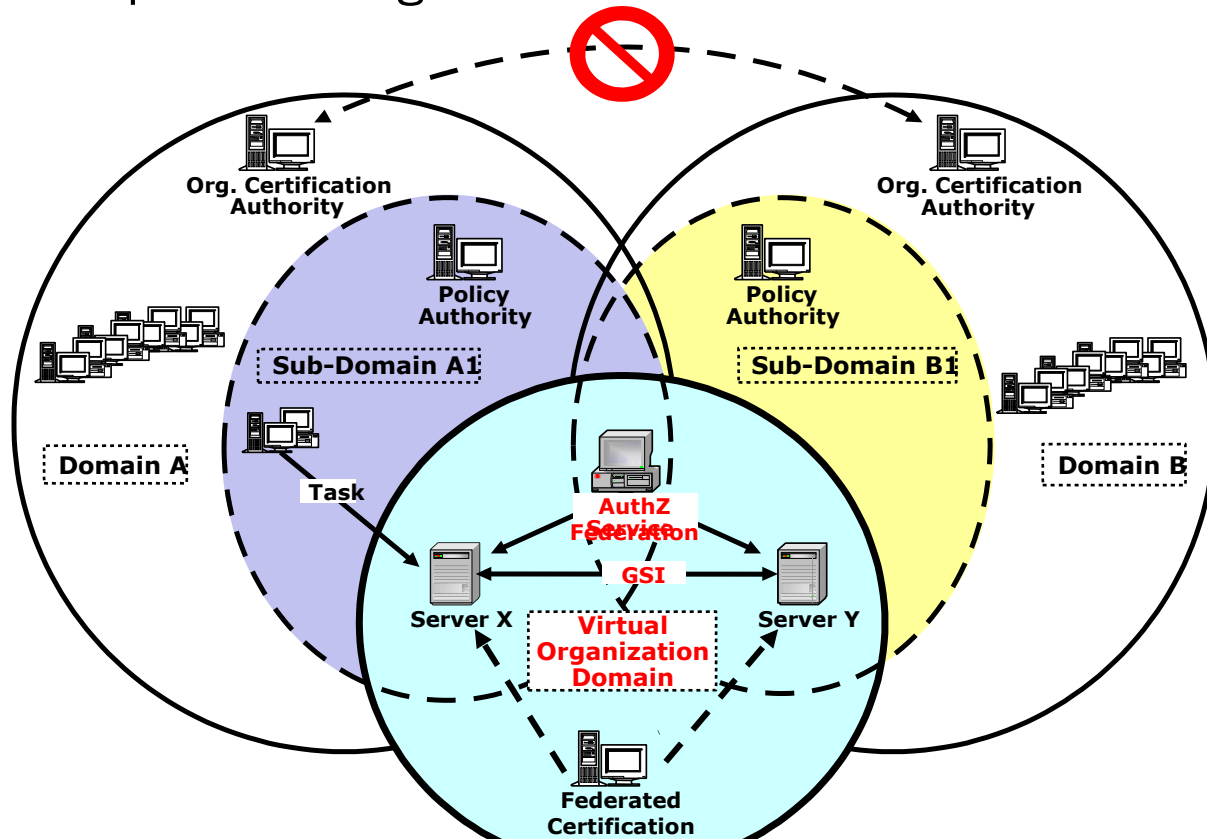
V Access and Allocation

- > But: why grant access to a user or community?
 - > Joint research programme
 - > Joint funding in projects
 - > Economic models,
either virtual 'pot money' or proper billing & settlement

- > Not too different from 'conventional' models
 - > 'Get an account because we work together'
 - > Allocations on supercomputers or large clusters
 - > Pay-per-use infrastructure (AWS EC2 & S3, etc...)

V Trust relationships

- > For the VO model to work, parties need a trust relationship
 - > **the alternative: every user would need to register at every resource!**
 - > need to provide a 'sign-on' for users that works across VOs



graphic from: Frank Siebenlist, Argonne Natl. Lab, *Globus Alliance*

V Elements of Trust


- > Authentication
 - > Who are you?
 - > Who says so?
-

- > Authorization
 - > Why should I let you in? What are you allowed to do?
 - > By whom? *Who* said you could do that?
 - > Community management and registration

- > Accounting (billing and settlement)
- > Incident Response
- > Compliance

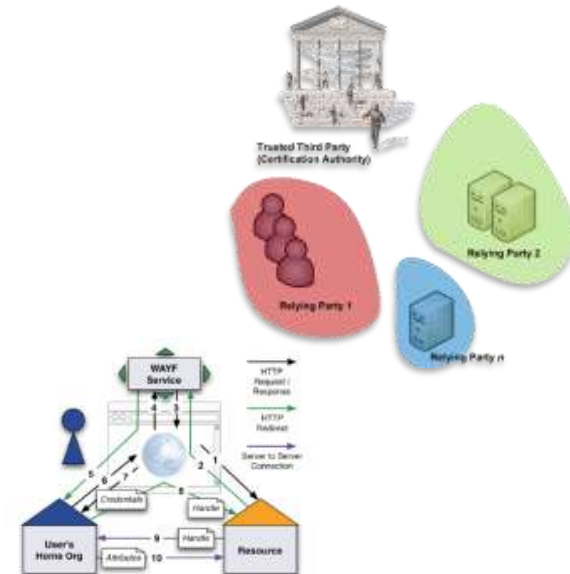
V Grid Security Policy ecosystem

- > A User and VO directed policy implementation



V Authentication models

- > Direct user-to-site
 - > passwords, enterprise PKI, Kerberos
 - > Usually with implicit authZ
- > PKI with trusted third parties
- > Federated access
 - > Controlled & policy based
 - > 'Free-for-all', e.g., OpenID
- > Identity meta-systems
 - > Infocard type systems




▼ Typical application domains

- Web access and direct user interactions
 - > Moving towards WebSSO & federations
 - > Or use client PKI where users already have certificates



- Task delegation (compute, data management)
 - > PKI 'Trusted Third Party' based
 - > Augmented with 'proxy' (RFC3820) delegation



V Grid authentication

With emergence of production grids: need for providing *cross-national* trust

Driven by resource owner – ‘relying party’ – needs

- > independent of users and Vos, who have a conflict of interest

- > National PKI?
 - > in general uptake of 1999/93/EC and e-Identification is (too) slow
- > Various commercial providers?
 - > Main commercial drive: secure web servers based on PKI
 - > Comodo, Verisign, Global Sign, Thawte, Verisign, SwissPost, ...
 - > primary market is server authentication, not end-user identities
 - > use of commercial CAs solves the ‘pop-up’ problem
 - ... so for (web) servers a pop-up free service is actually needed
- > Grass-roots CAs?
 - > usually project specific, and without documented policies
 - > unsuitable for the ‘production’ infrastructure

✓ Building a grid authentication infrastructures


> Grid research/academic PKIs

- > started off with pre-existing CAs, and some new ones
- > ‘reasonable’ assurance based on documented procedures
- > single assurance level inspired by grid-relying party** requirements
- > using a threshold model: *minimum requirements*

> Grid CA coordination driven by 2000 need to solve cross-national authentication issues *right now*

- > separation of AuthN and AuthZ allowed progress in the area
- > the policies convinced enough resource providers to ‘trust’ the AuthN assertions
- > there were and are individuals all over Europe (and the world) that need access to these resource providers

V Federation Model for Grid Authentication



> Federation of independent CAs

- > common **minimum requirements** (in various flavours)
- > trust domain as required by users and relying parties
where relying party is (an assembly of) resource providers
- > defined and peer-reviewed acceptance process

> No single top

- > leverage of national efforts and complementarities
- > Allow paced regional development, organisation and customisation

Guidelines: common elements in the IGTF

- > **Coordinated namespace**
 - > Subject names refer to a unique entity (person, host)
 - > Usable as a basis for authorization decisions
 - > This name uniqueness is essential for *all authentication profiles!*

- > **Trust anchor repository**
 - > Coordinated distribution for all trust anchors in the federation
 - > Trusted, redundant, sources for download, verifiable via TACAR

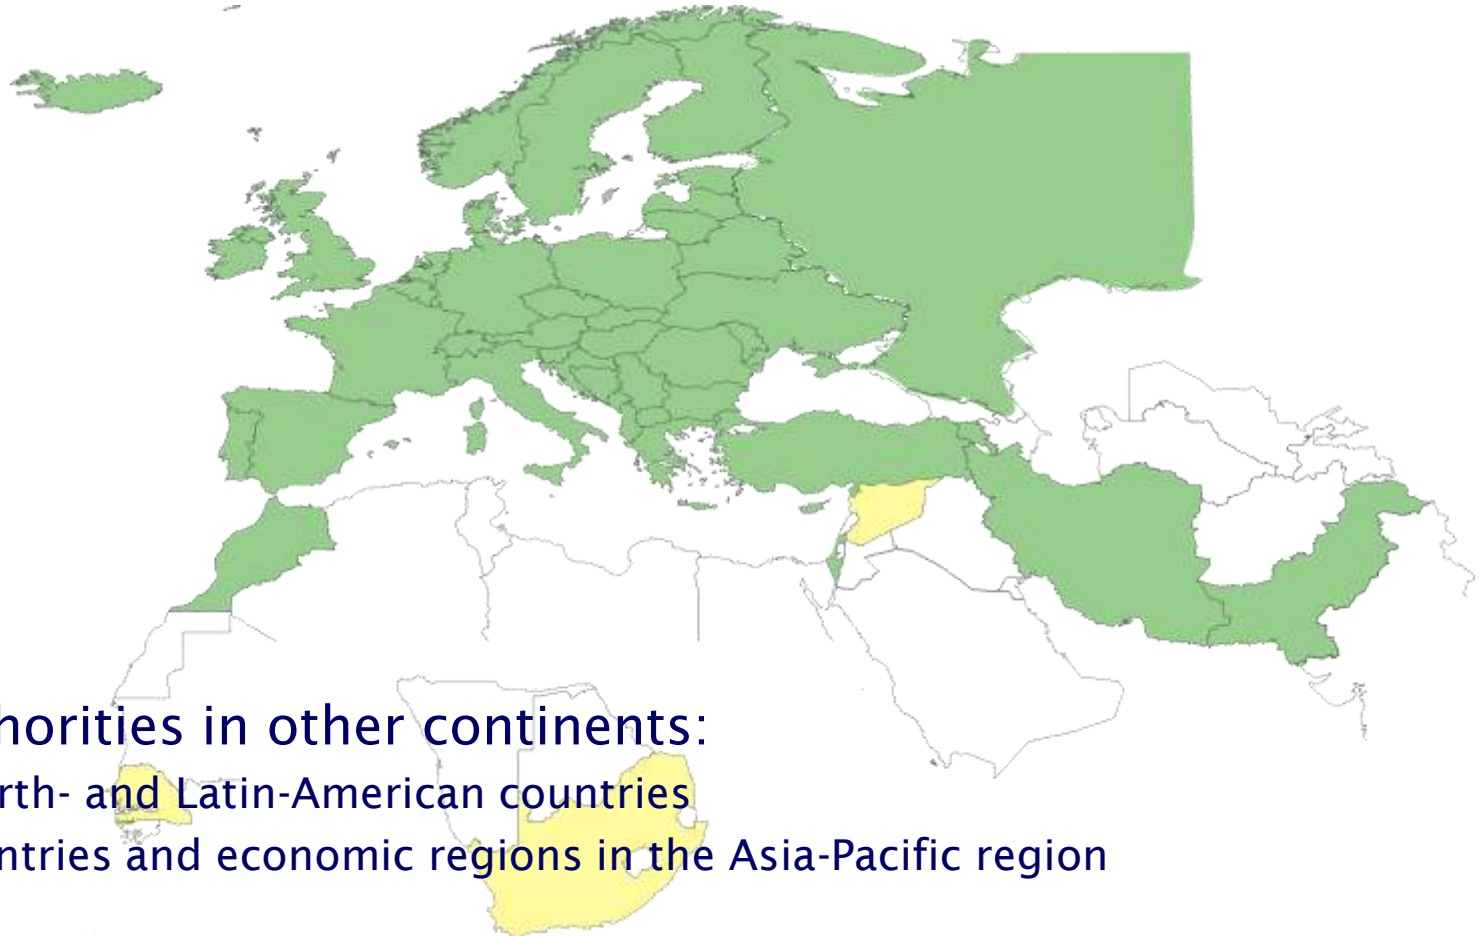
- > **Concerns, risk assessment, and incident handling**
 - > Guaranteed point of contact
 - > Forum to raise issues and concerns

- > **Documented processes of federation and authorities**
 - > Detailed policy and practice statement
 - > Auditing by federation peers

V Geographical coverage

Green: EMEA countries with an Accredited Authority

- 23 of 25 EU member states (all except LU, MT)
- + AM, BY, CH, HR, IL, IR, IS, MA, MD, MK, NO, PK, RS, RU, TR, ...



More Authorities in other continents:

- Most North- and Latin-American countries
- 13+ countries and economic regions in the Asia-Pacific region



Grouping users

VO management technologies

Delegation and access scenarios

AUTHORIZATION AND VIRTUAL ORGANISATIONS

✓ Authorization: VO representations

- > VO is a directory (database) with members, groups, roles
- > Based on identifiers issues at the authentication stage
- > Membership information is then to be conveyed to the resource providers
 - > configured statically, out of band
 - > in advance, by periodically pulling membership lists
LDAP directories, replicated databases (GUMS)
 - > in VO-signed assertions pushed with the request:
VOMS, Community AuthZ Service
- > Except for the **CA provided DN**, the VO is all the site will see
 - > Since VO is user-centric, it has a potential conflict of interest for identity

V VOMS: VO attributes in a X.509 container

Virtual Organisation Management System (VOMS)


- > developed by INFN for EU DataTAG and EGEE
- > used by VOs in EGEE, Open Science Grid, NAREGI, ...
- > push-model signed VO membership tokens
 - > using the traditional X.509 'proxy' certificate for trans-shipment
 - > fully backward-compatible with only-identity-based mechanisms

VOMS proxy with embedded VO assertion
Serial Number: 26423 (0x6737)
Issuer: O=dutchgrid, O=users, O=nikhef, CN=David Groep
Not Before: Oct 16 12:46:28 2006 GMT
Not After : Oct 17 00:51:28 2006 GMT
Subject: O=dutchgrid, O=users, O=nikhef, CN=David Groep, CN=proxy
Subject Public Key Info:
Public Key Algorithm: rsaEncryption
RSA Public Key: (512 bit)
X509v3 extensions:
1.3.6.1.4.1.8005.100.100.5:
0...0...0...0.....0W.U0O.M0K1.0...U./dteam/ne/ROLE=null/0...0...0
X509v3 Key Usage:
Digital Signature, Key Encipherment, Data Encipherment
Signature Algorithm: md5WithRSAEncryption

Attribute Certificate	
INTEGER	1
SUBJECT	/O=dutchgrid/O=users/O=nikhef/CN=David Groep
SERIAL	0396
ISSUER	/C=CH/O=CERN/CN=lcg-voms.cern.ch
OCTET STRING	/dteam/Role=NULL/Capability=NULL
OCTET STRING	/dteam/ne/Role=NULL/Capability=NULL
OBJECT	No revocation available
AuthorityKeyIdentifier	0...H...0.....<3...#..
SignatureAlgorithm	md5WithRSAEncryption



V VOMS model



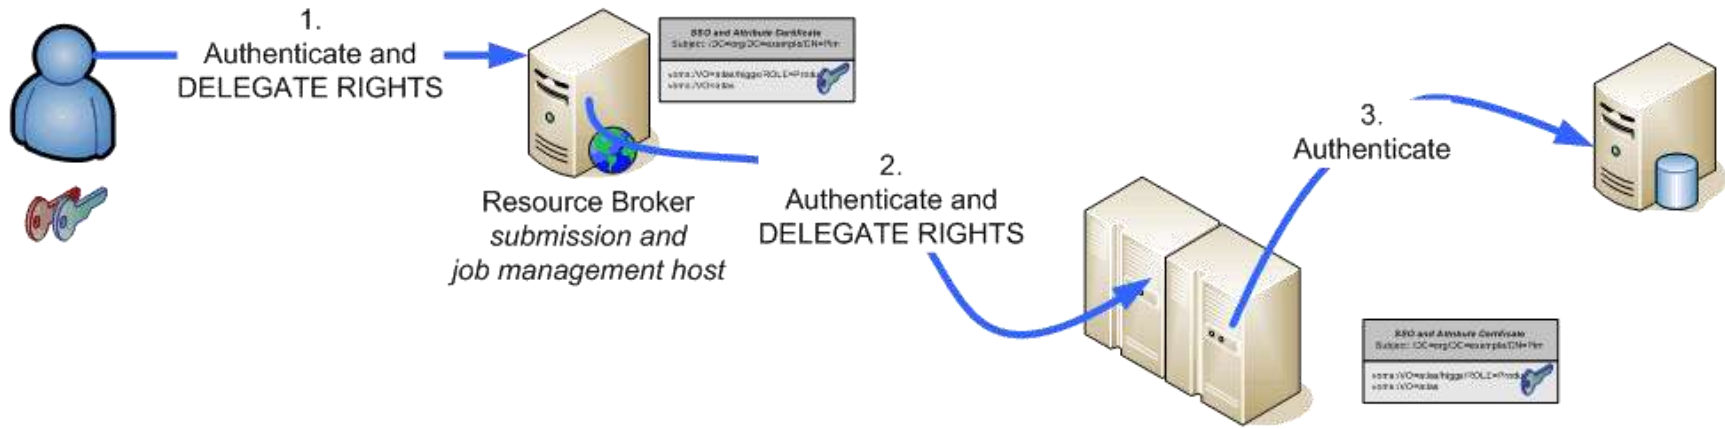
V Delegation

- > Mechanism to have someone, or some-thing – a program – act on your behalf
 - > as yourself
 - > with a (sub)set of your rights
- > Matches model of brokering and non-interactive (automated) operations
- > GSI (PKI) and recent SAML drafts define this
 - > GSI (PKI) through ‘proxy’ certificates (see RFC3820)
 - > SAML through *Subject Confirmation*, (linking to at least one key or name)

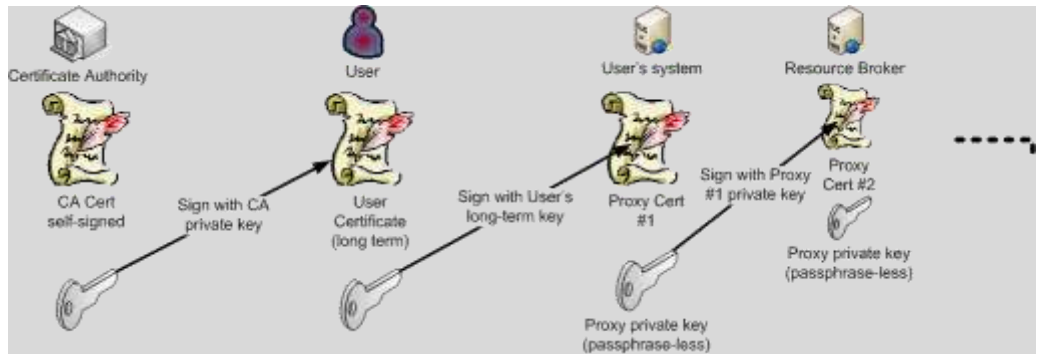
V Daisy-chaining proxy delegation

User Job
Data processing, reading
and writing remote files

SSO and Attribute Certificate
Subject: /DC=org/DC=example/CN=Pim
voms:/VO=atlas/higgs/ROLE=Production
voms:/VO=atlas



User Job
Data processing, reading
and writing remote files



V Acceptable Credentials on the Grid

'Let's not make the SSH mistake again'

'All Credentials Have A Life Time'

- > Long lived credentials must be revocable
- > Short lived (< 100ks) credentials may be left to expire

So we get

- > X.509 identity certificates: **<= 1 year**
- > Proxy credentials: **between 12 and ~24 hours**
- > VOMS attributes: **~ 24 hours**
- > Proxies in a managed credential store: **1Ms, ~11 days**
- > 'limited delegation' proxies prevents *creeper-reaper*-type exploits



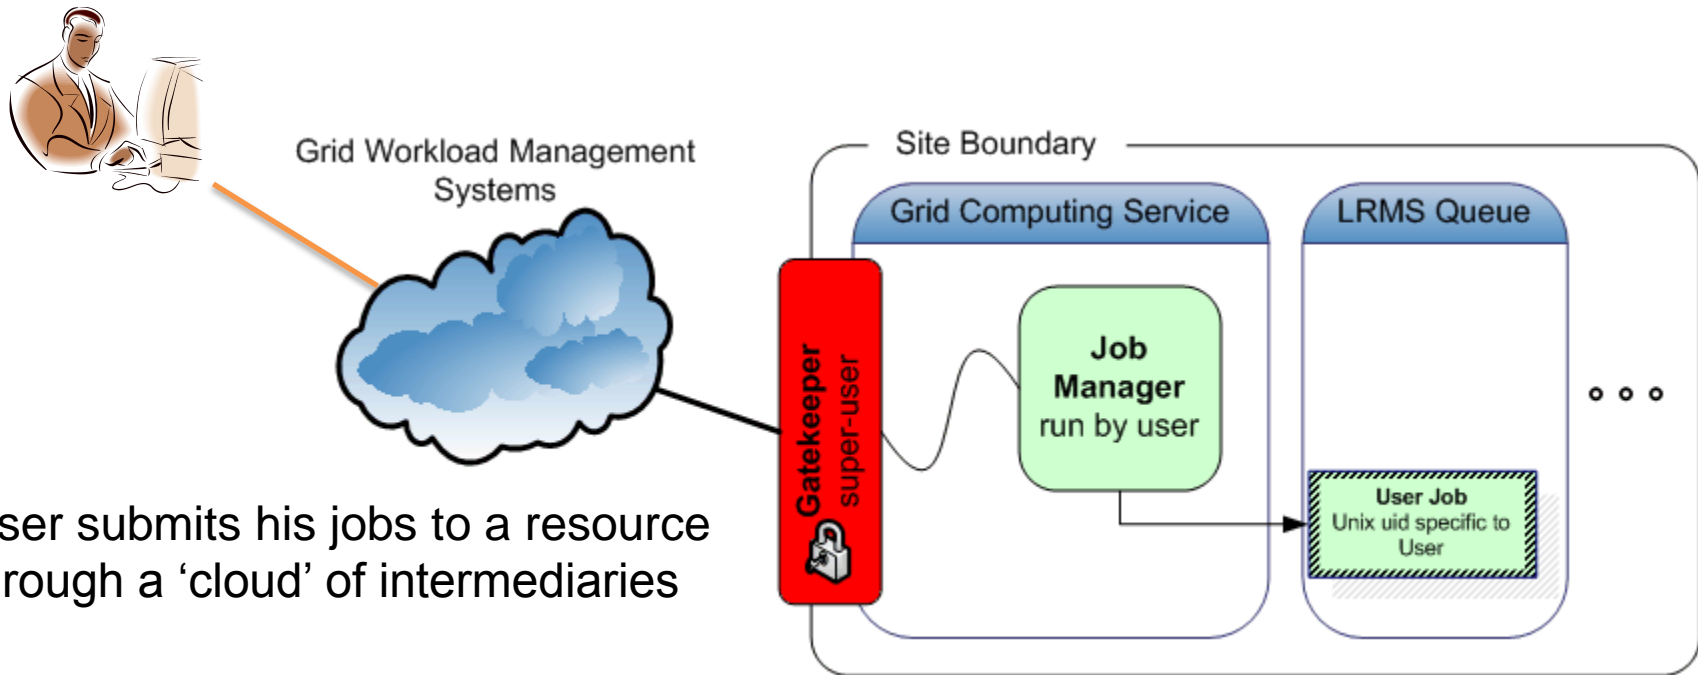
Example: running compute jobs

Tracing users and actions

Storage

ACCESS CONTROL AT THE SITE

V Accessing (compute) resources




User submits his jobs to a resource through a 'cloud' of intermediaries

Direct binding of payload and submitted grid job

- job contains all the user's business
- access control is done at the site's edge
- inside the site, the user job has a specific, site-local, system identity

V To the Unix world

grid identity



(X509, VOMS)

/dc=org/dc=example/CN=John Doe


translate

```
enmr001:x:43401:2029:PoolAccount eNMR 001:/home/enmr001:/bin/sh
```

- > Unix does not talk Grid, so translation is needed between grid and local identity
- > this translation has to happen somewhere
 - > On entry at the Gatekeeper
 - > When running tasks or accessing files

V Access Control on the CE

- > System access (authorization: LCAS, mapping: LCMAPS)
- > Embedded or through 'call-out hooks' in Grid middleware



V Access Control

Granting access: grid-mapfile

```
...
"/O=dutchgrid/O=users/O=nikhef/CN=David Groep" .dans
"/O=dutchgrid/O=users/O=nikhef/CN=Sven Gabriel" .dteam
"/O=dutchgrid/O=users/O=wageningen-universiteit/CN=Anonymised User" .lsg
"/O=dutchgrid/O=users/O=wageningen-universiteit/CN=Anonymised User" .lsg
"/alice/Role=lcgadmin" .alisgm
"/alice" .alice
"/atlas/Role=lcgadmin" .atlsm
"/atlas/Role=production" .atlb
"/atlas/Role=pilot" .atlpi
"/atlas/nl" .atlnl
"/atlas" .atlas
```

Denying access: ban_users.db

```
# This file contains the user subject DNs that are BANNED from this fabric
#
"/C=UK/O=eScience/OU=Cambridge/L=UCS/CN=anonymised user"
# from [UPDATE 5] Security incident - XXXXCERT-20080805, 04.Sept. 11:52
"/O=Grid/O=NorduGrid/OU=somesite.se/CN=Olof Palme"
"/O=Grid/O=NorduGrid/OU=somesite.se/CN=Alfred Nobel"
# 16-Jan-2009 banned compromised DN
"/C=CN/O=HEP/O=PKU/OU=PHYS/CN=Mao Zhedong"
# 23-Feb-2009 Security Service Challenge
# SG let Arjen in again after 18.Mar 2009
"/O=dutchgrid/O=users/O=nikhef/CN=Arnold Johan van Rijn"
```

V What does the site owner see?

Batch system

```
stro.nikhef.nl:
```

Job ID	Username	Queue	Jobname	SessID	NDS	TSK	Req'd Memory	Req'd Time	Elap S	Time	
3223967.stro.nikhef.	atlb021	atlas	STDIN	32473	1	--	--	66:00	R	--	wn-val-046
3227086.stro.nikhef.	atlb021	atlas	STDIN	22038	1	--	--	66:00	R	--	wn-val-004
3227691.stro.nikhef.	atlb019	atlas	STDIN	11290	1	--	--	66:00	R	--	wn-lui1-028
3228887.stro.nikhef.	atlb021	atlas	STDIN	1562	1	--	--	66:00	R	--	wn-val-091
3235888.stro.nikhef.	lhcbpi01	lhcb	STDIN	23903	1	--	--	33:00	R	32:11	wn-lui2-014
3236232.stro.nikhef.	atlb019	atlas	STDIN	26115	1	--	--	66:00	R	32:10	wn-bull-011

Gatekeeper audit log

```
PID: 13507 -- Requested service: jobmanager-pbs
PID: 13507 -- Authorized as local user: atlb019
PID: 13507 -- Authorized as local uid: 70019
PID: 13507 -- and local gid: 2036
PID: 13507 -- "/C=CA/O=Grid/OU=westgrid.ca/CN=Anony Mous" mapped to atlb019 (70019/2036)
PID: 13507 -- GATEKEEPER_JM_ID 2009-11-16.12:51:40.0000013507.0000000000 for
           /C=CA/O=Grid/OU=westgrid.ca/CN=Anony Mous on 142.90.256.257
PID: 13507 -- Child 13576 started
```

V Tracing the job

JobManager log

```
gmtime=20091116115140Z;uniqid=19095.1258344261;ug=70019:2036 2036;  
jobid=3243289.stro.nikhef.nl; tag=https://gazon.nikhef.nl:20082/19095/1258344261/;  
dry=no; jobtype=single; count=1;  
exec=https://condorg.triumf.ca:20014/home/atlasprod/Panda/pyfactory/20091105/runpilot3-  
wrapper.sh;  
args=;  
dir=/home/atlb019//gram_scratch_pCHATpWQJY;log=/home/atlb019/gram_job_mgr_13576.log;
```

Batch system syslog entry

```
Nov 16 11:51:40 gazon jobmanager-pbs[19374]: qsub success (atlb019:atlb)  
/home/atlb019/.globus/job/gazon.nikhef.nl/19095.1258344261/scheduler_pbs_job_script:  
3243289.stro.nikhef.nl
```

As well as regular entries created by the batch system(s) and any auditing data

V Storage: Virtual Ids or Unix domain?

- > Mapping to Unix credentials
 - > Lacks expression of VO attributes and rights
 - > Allows joint native and grid use of storage systems
- > Grid storage systems with grid meta-layer access control
 - > No need to allocate Unix-level resources or mappings
 - > Expresses both VO and site-level policies and ACLs
 - > Access *must* be via grid-aware mechanisms

Example: Disk Pool Manager DPM:

- > mapped to 'virtual UIDs': created on the fly first time system sees DN
- > VOMS roles are mapped to virtual GIDs
- > User can have one DN and several roles, so may be mapped to one UID and several GIDs

V Example Access Control Lists

- > LFC and DPM support Posix ACLs based on Virtual Ids
 - > Access Control Lists on files and directories
 - > Default Access Control Lists on directories: they are inherited by the sub-directories and files under the directory
- > Example
 - > `dpns-mkdir /dpm/cern.ch/home/dteam/jpb`
 - > `dpns-setacl -m d:u::7,d:g::7,d:o:5 /dpm/cern.ch/home/dteam/jpb`
 - > `dpns-getacl /dpm/cern.ch/home/dteam/jpb`
 - # file: /dpm/cern.ch/home/dteam/jpb
 - # owner: **/C=CH/O=CERN/OU=GRID/CN=Jean-Philippe Baud 7183**
 - # group: dteam
 - user::rwx
 - group::r-x #effective:r-x
 - other::r-x
 - default:user::rwx
 - default:group::rwx
 - default:other::r-x

V Handling E2E incidents in this system

- > Detection and coordination
 - > Globally unique identifiers (subject DNs, VO names)
 - > Policy ecosystem guidelines for auditing, log retention, and information exchange between participants
 - > Periodically tested through SSCs
- > Revocation

which, e.g., ssh keys don't have, but federated access does

 - > At the identity level, the Grid implements working revocation and CRL support for the PKI
 - > At the authorization level: VO-level banning, site bans
- > Recovery
 - > De-facto, the only transparent recovery is by revocation of identity
 - > Subject name (DN) is persistent for the user across incidents, so no re-registration needed



SUMMARY



V Summary

- > Grid and the VO make collaboration explicit at systems level
 - > Structure of researchers themselves drives VO structure
 - > This discloses the ‘interconnected vulnerabilities’ & incidents issue
- > Threats in distributed computing exist irrespective of Grid
 - > Multiple accounts across organisations, usually ill-managed
 - > Shared or semi-public group accounts or shared storage
 - > Grid middleware gives some additional handles ...
 - > ... but also exposes new risk surfaces
- > We have yet to see a grid-specific incident
 - > Many ‘traditional’ incidents propagate along research collaborations
 - > Using non-grid attack vectors, and without ‘grid’ controls to help