vl·e virtual laboratory for e·science

BiG Grid
the dutch e·science grid

Grid Security
on a global scale

Oscar Koeroo @ eth-0

NIKHEF pdp

Graphics: Real Time Monito
Gidon Moont, Imperial College London, see http://gridportal.hep.ph.ic.ac.uk/rtm

index

- Intro (wie, wat, waar?)
- Grid Computing
  - Wat? Hoe?
- Hoe zit de security in elkaar
  - End-to-End security
  - De uitdagingen

Wie ben ik?

nerdtest-score: 89

- Security middleware developer
  - Van afstudeerproject in 2003...
  - European Data Grid (tot 2003)
  - EGEE, EGEE-II en EGEE-III (tot april 2010)
  - ... en verder
- Werkzaam in diverse middleware groepen
  - Algemene security implementaties
  - Standaardisatie groepen
  - Actief in interoperability tussen Grid infra

Overal te vinden

Overal te vinden

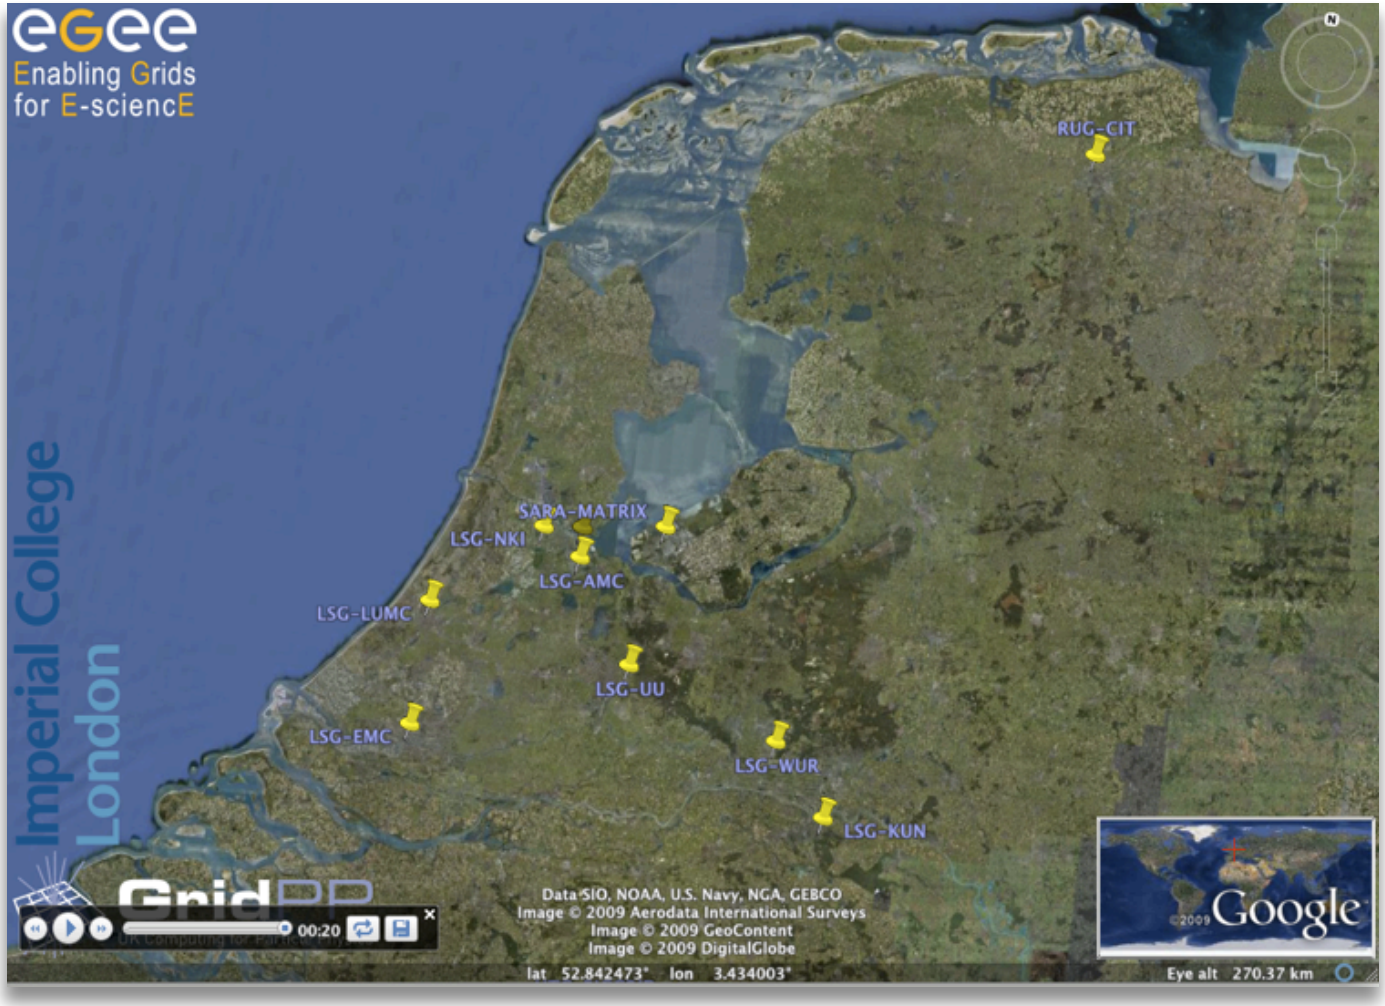
eGee
Enabling Grids for E-sciencE
Imperial College London
GridPP
UK Computing for Particle Physics
00:20
RUG-CIT
SARA-MATRIX
LSG-NKI
LSG-AMC
LSG-LUMC
LSG-UU
LSG-EMC
LSG-WUR
LSG-KUN
Data SIO, NOAA, U.S. Navy, NGA, GEBCO
Image © 2009 Aerodata International Surveys
Image © 2009 GeoContent
Image © 2009 DigitalGlobe
©2009 Google
lat 52.842473° lon 3.434003°
Eye alt 270.37 km
N

Overal te vinden

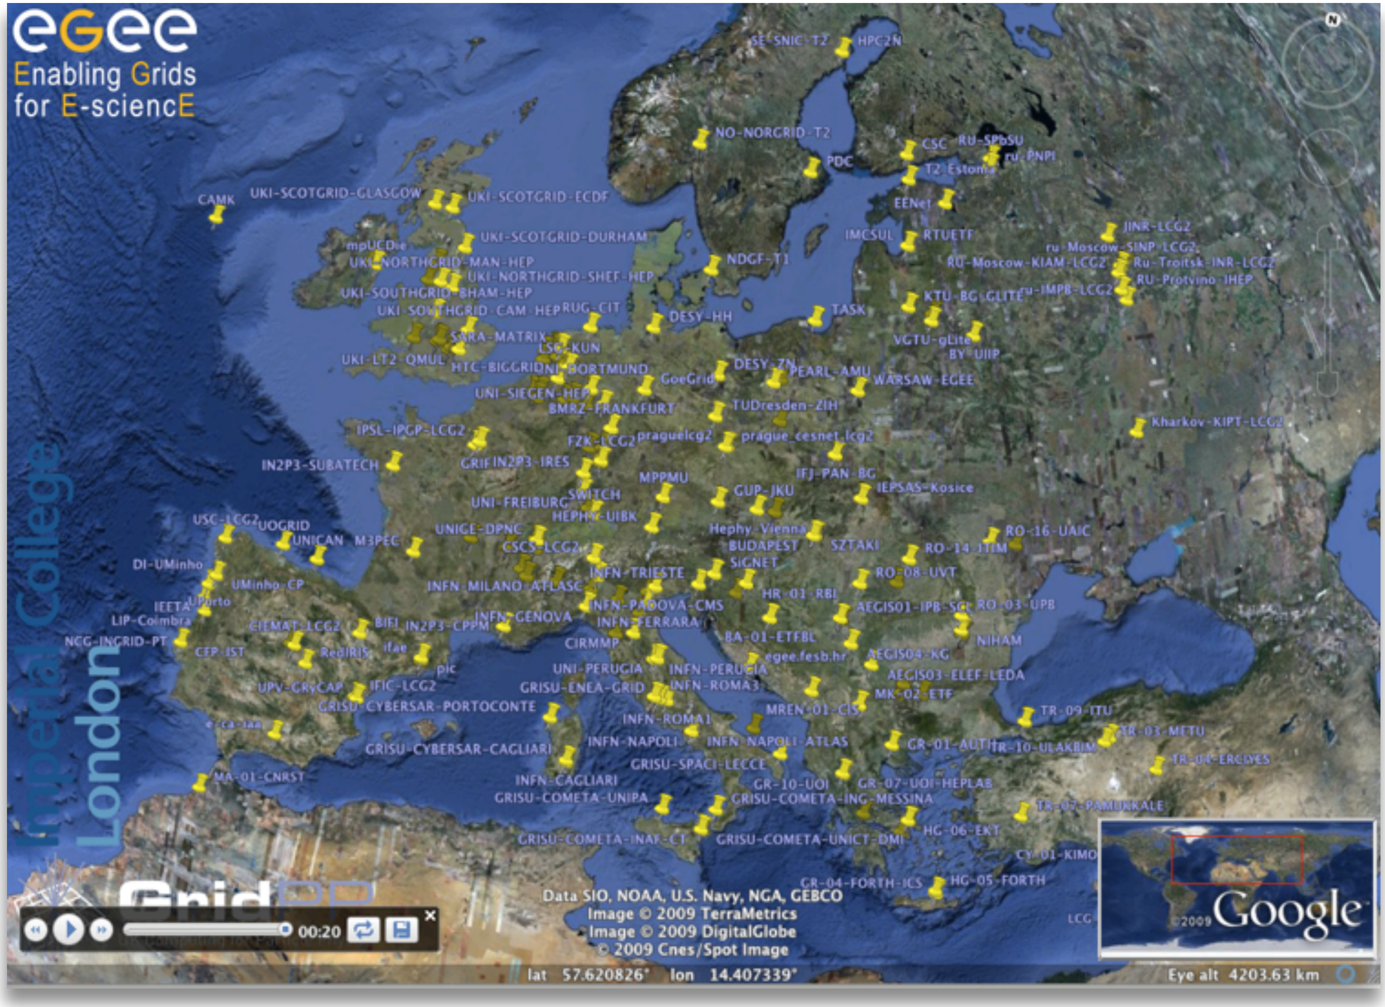

Data SIO, NOAA, U.S. Navy, NGA, GEBCO
Image © 2009 TerraMetrics
Image © 2009 DigitalGlobe
© 2009 Cnes/Spot Image

Overal te vinden

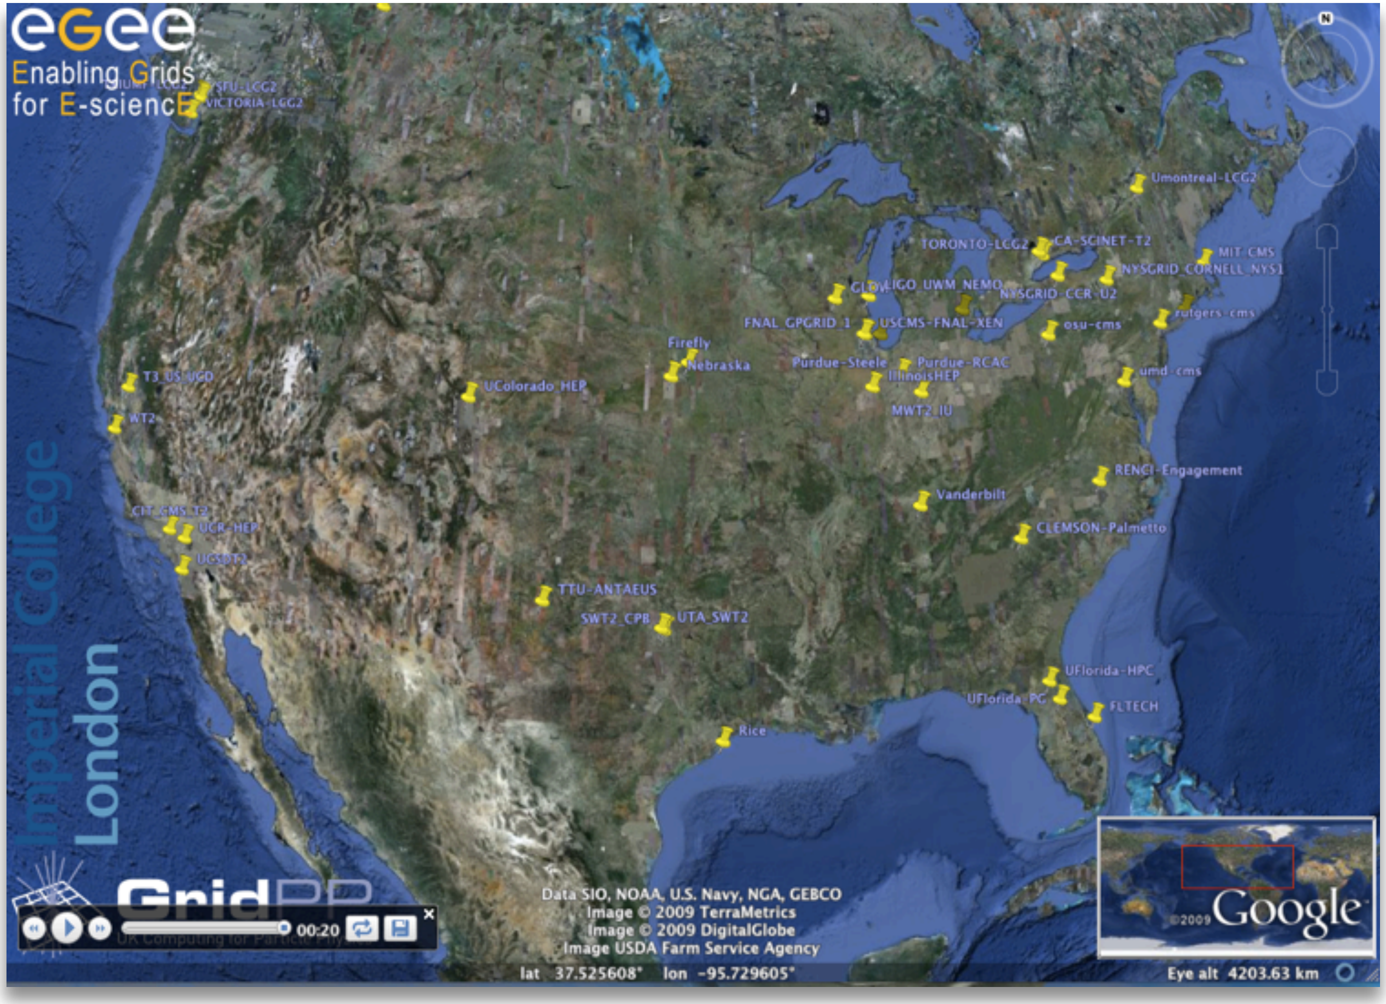

FOM institute for subatomic physics

Nikhef

NIKHEF

FOM

Nikhef samenwerkingsverband

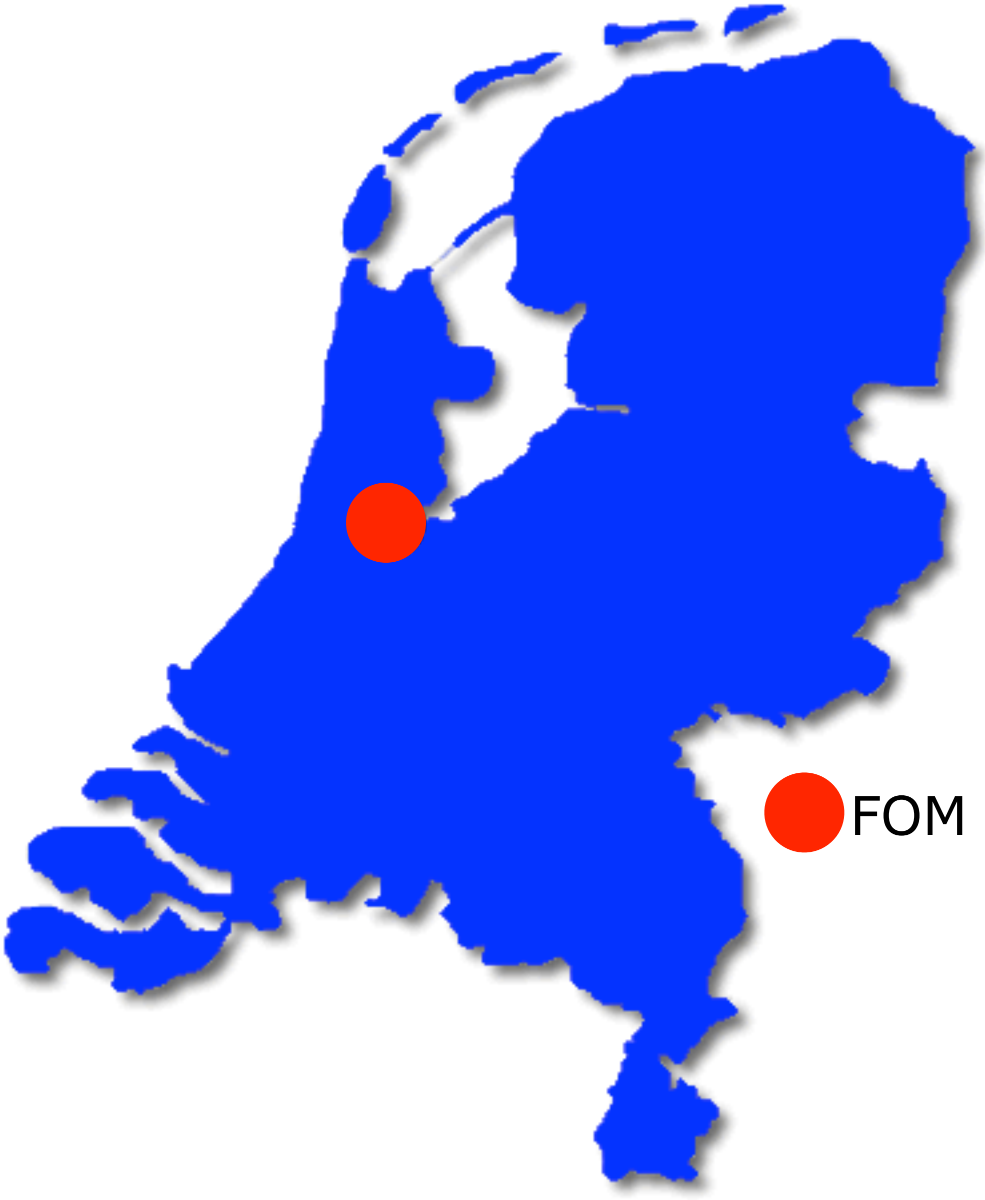

● FOM instituut voor subatomaire fysica

Nikhef samenwerkingsverband

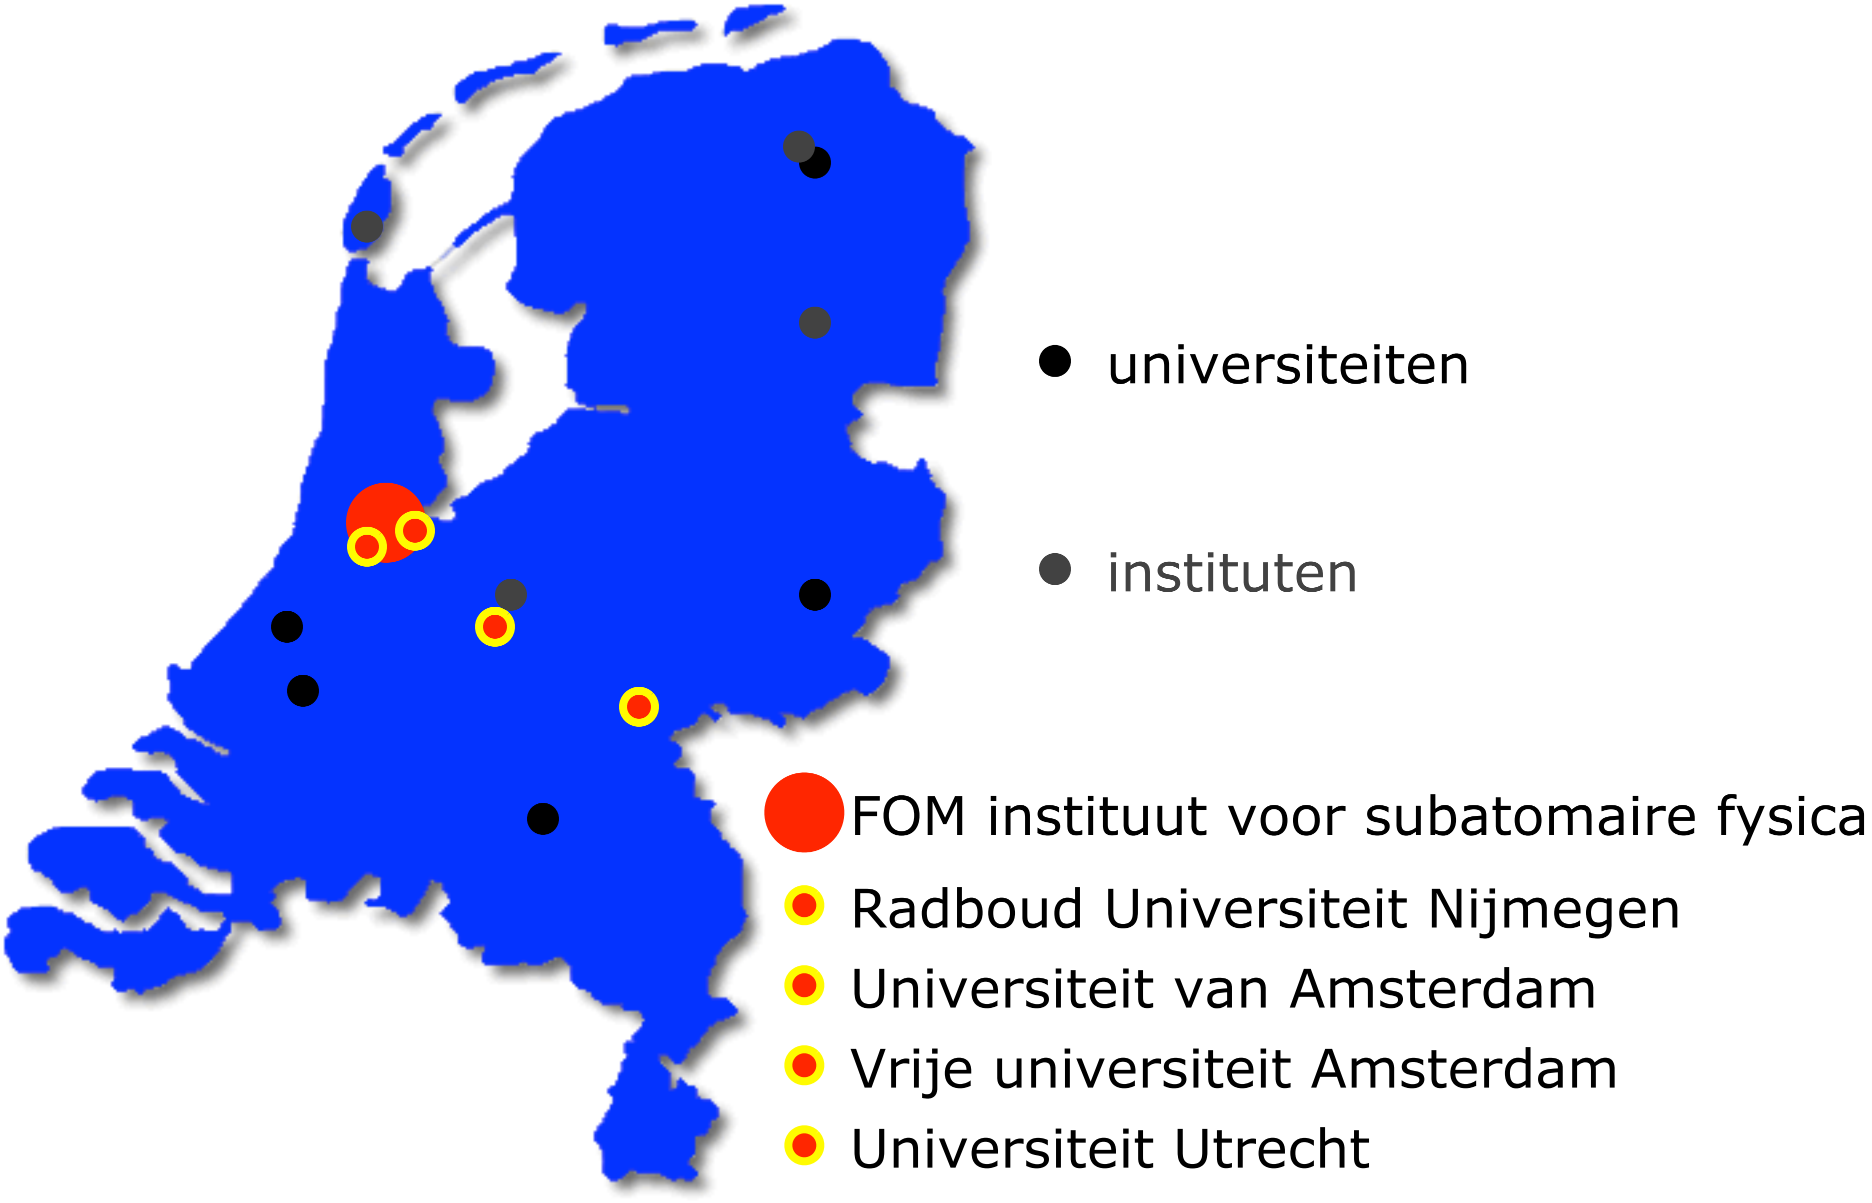

- ● universiteiten
- ● instituten
- ● FOM instituut voor subatomaire fysica
- ● Radboud Universiteit Nijmegen
- ● Universiteit van Amsterdam
- ● Vrije universiteit Amsterdam
- ● Universiteit Utrecht

deeltjesfysica

(effectief: LHC)

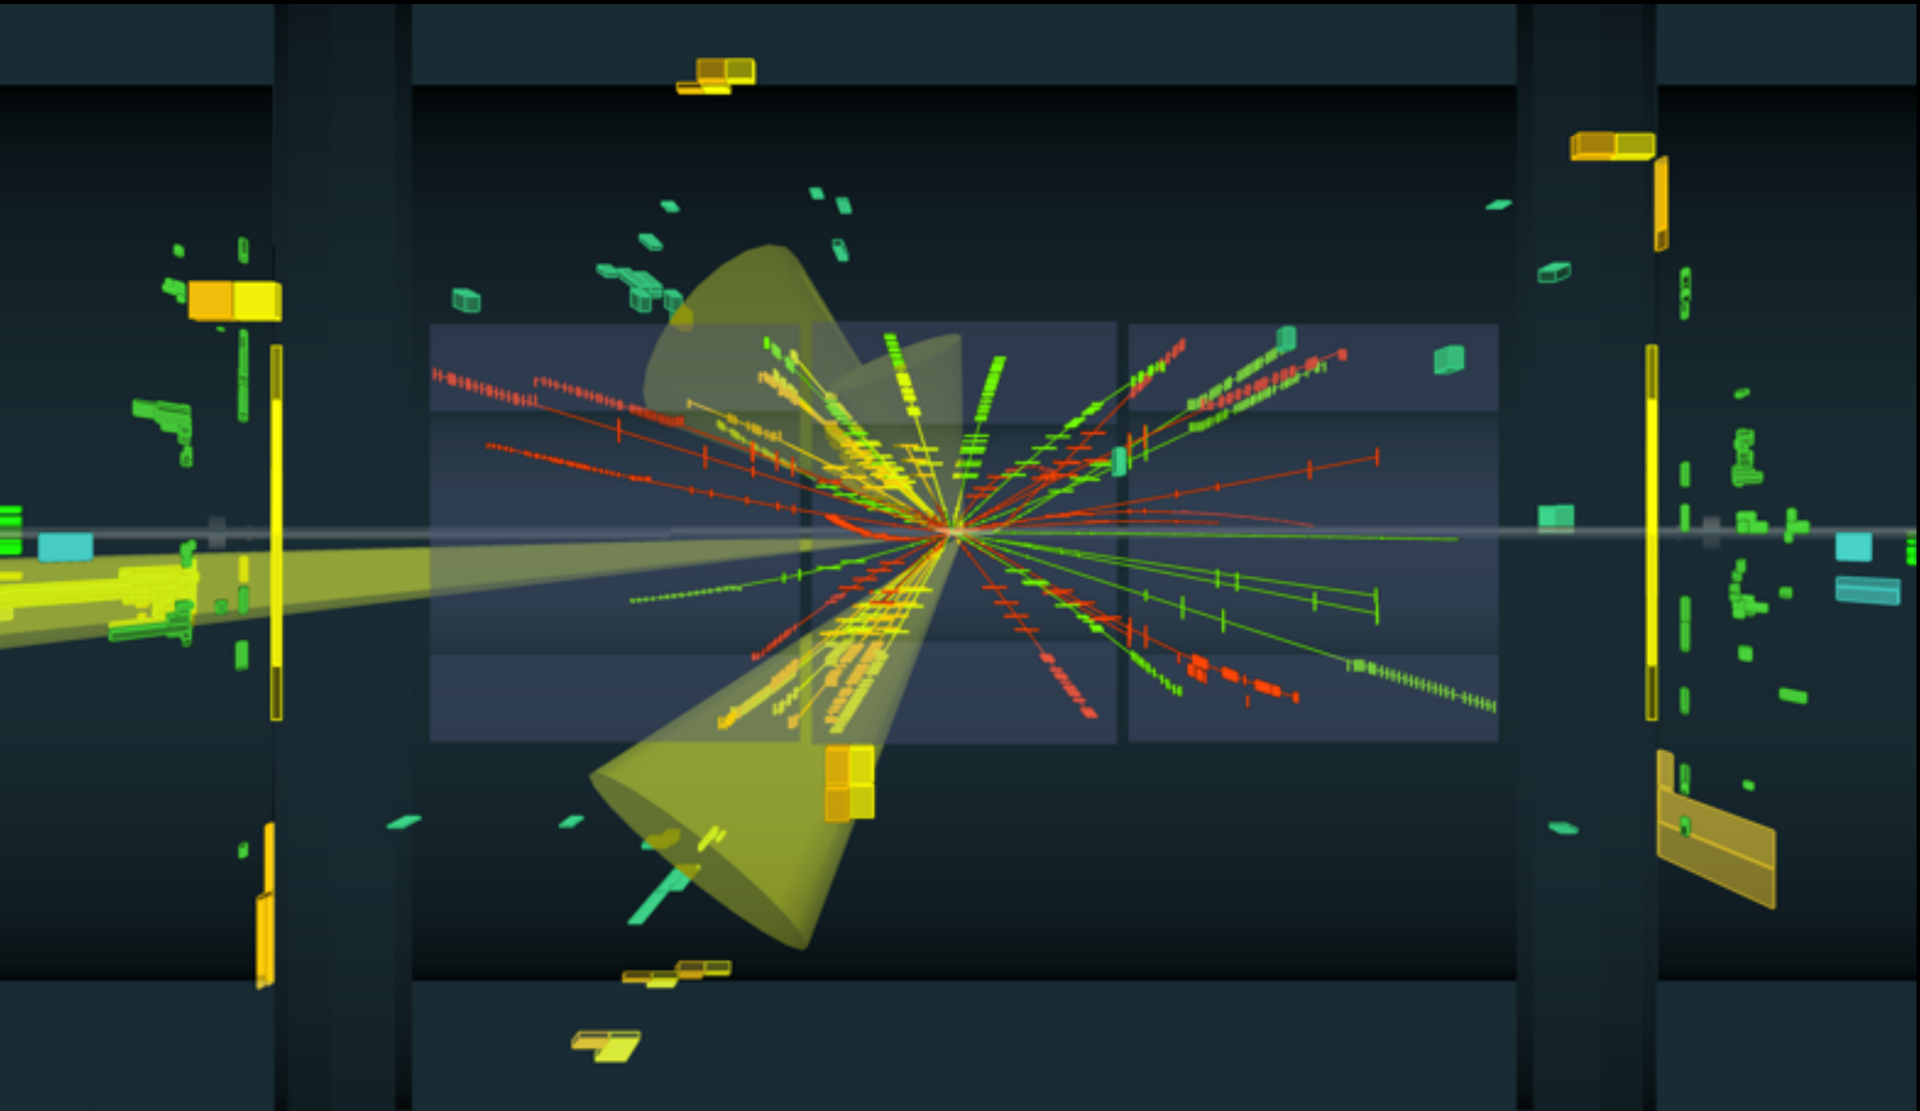

Jet Event at 2.36 TeV Collision Energy

2009-12-14, 04:30 CET, Run 142308, Event 482137

http://atlas.web.cern.ch/Atlas/public/EVTDISPLAY/events.html

astrodeeltjesfysica

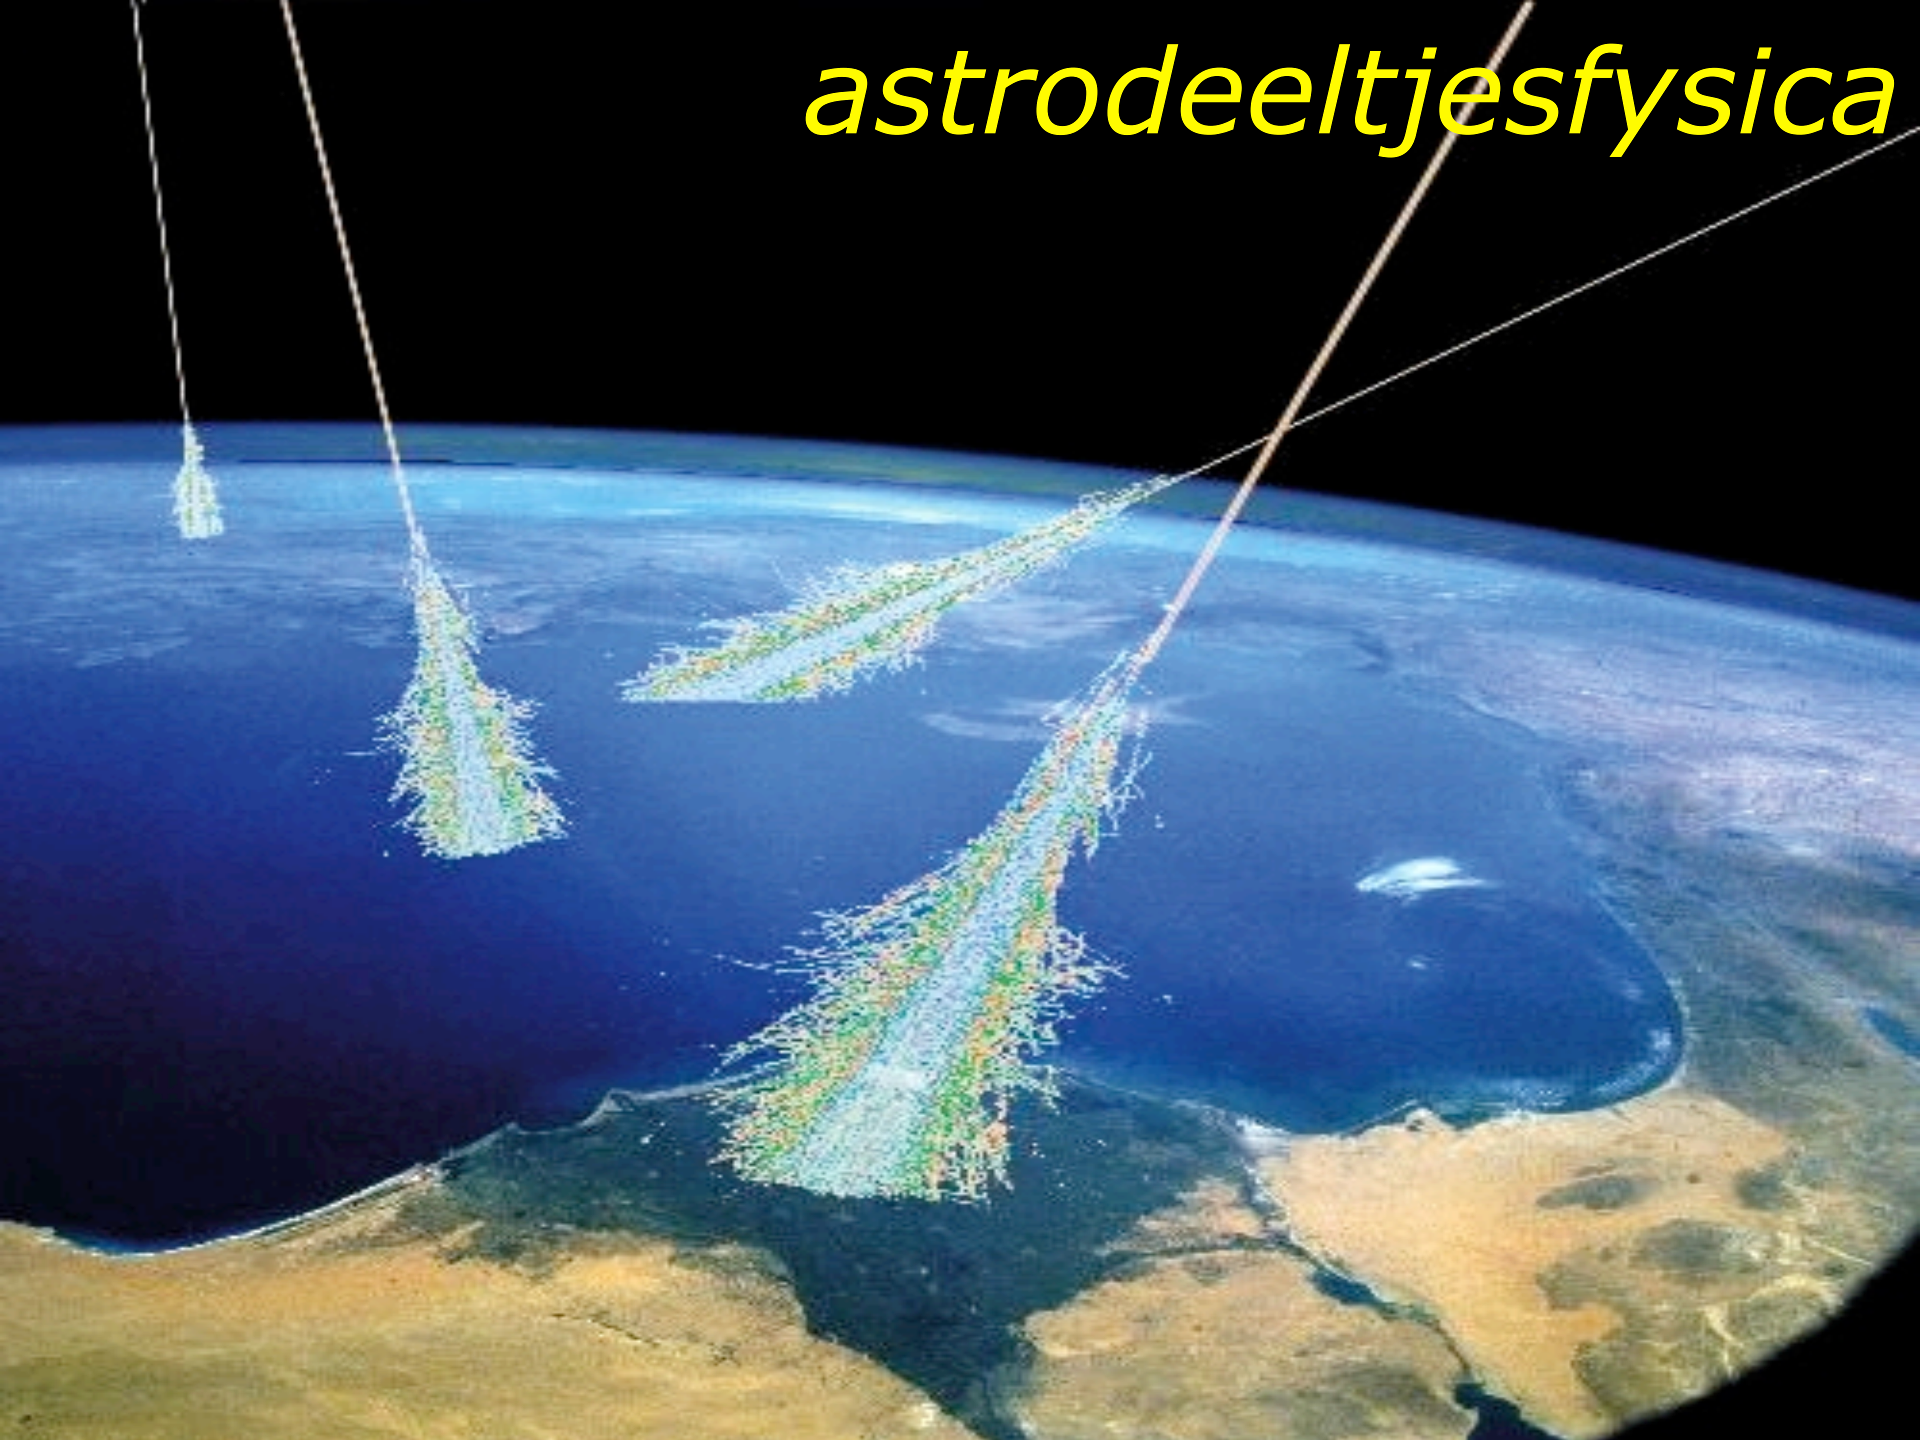

astrodeeltjesfysica

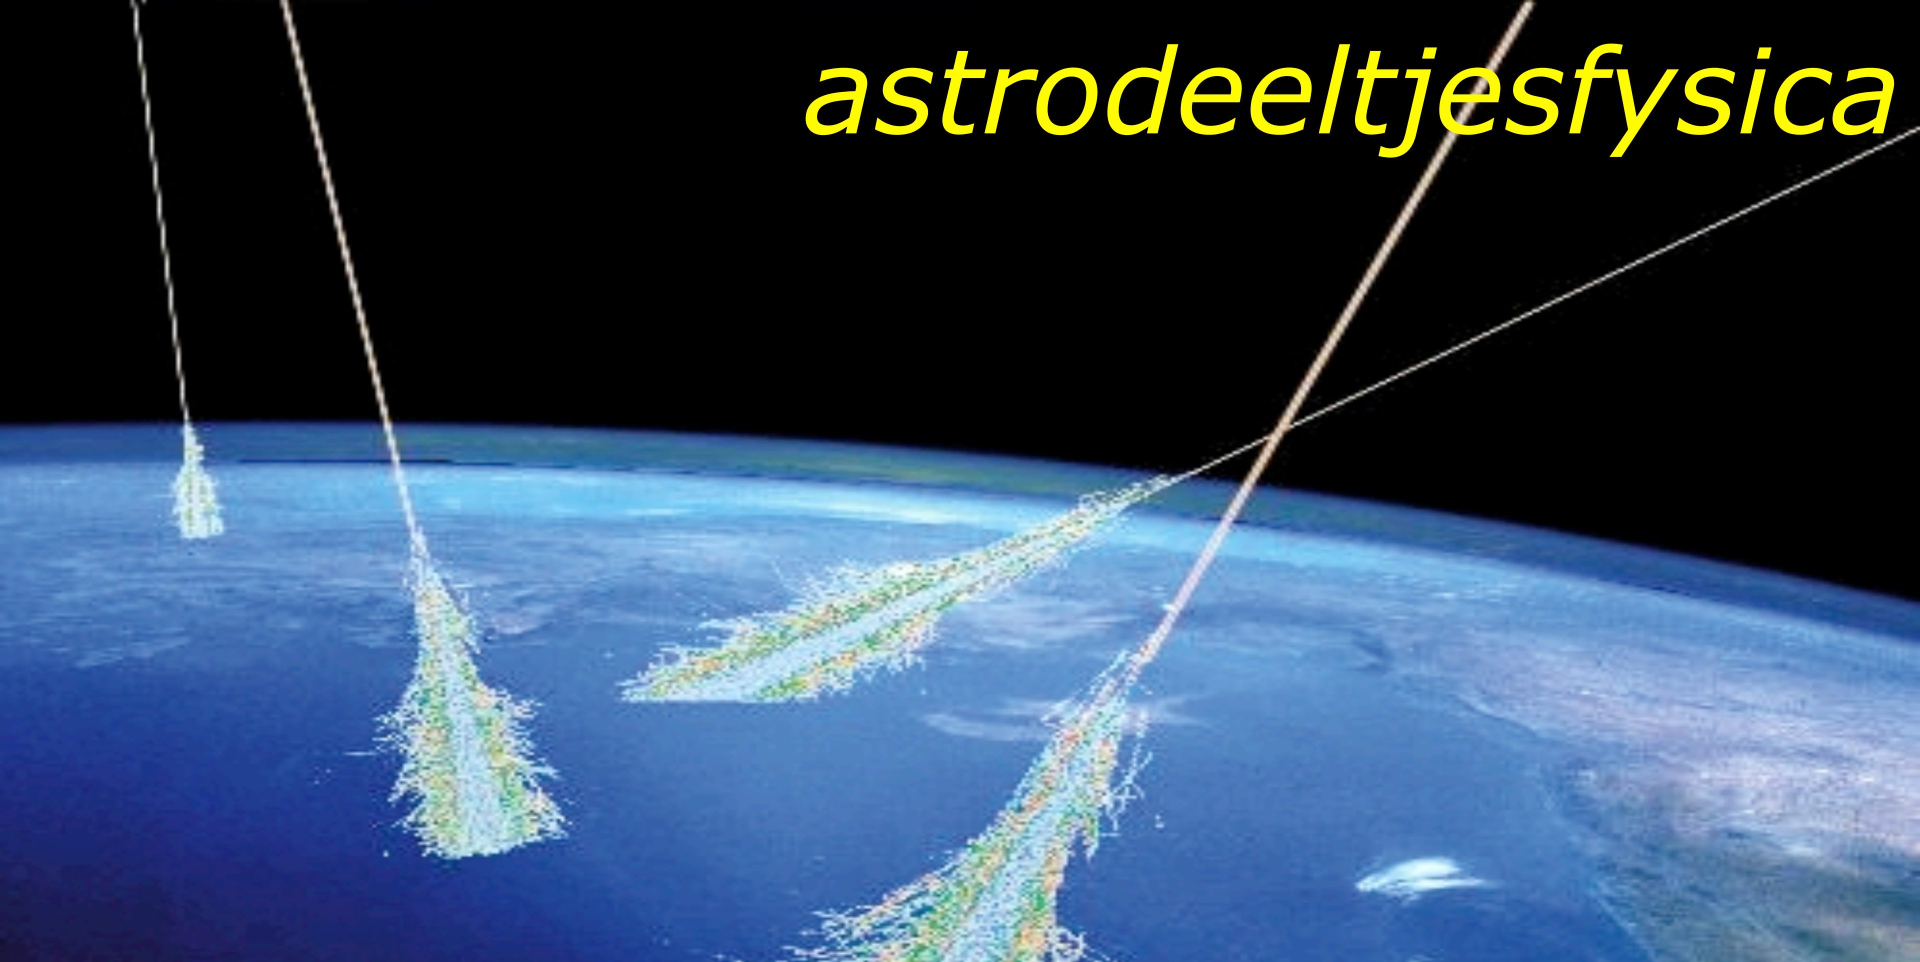

ANTARES
deep-sea
water-Cerenkov
neutrino telescope

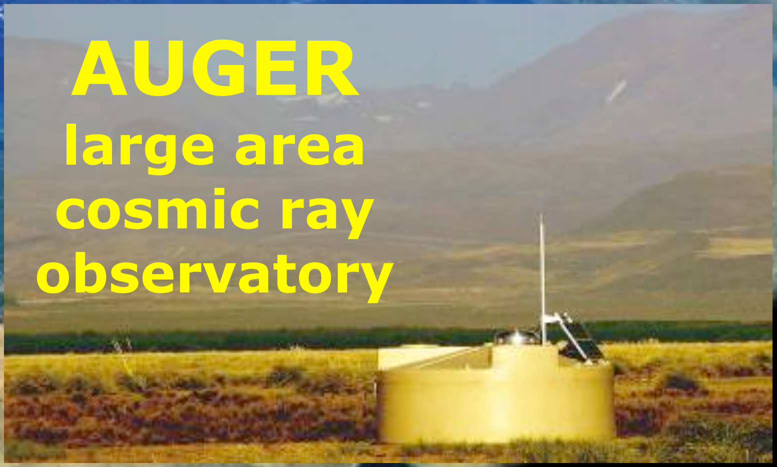

AUGER
large area
cosmic ray
observatory

gravitatie golven

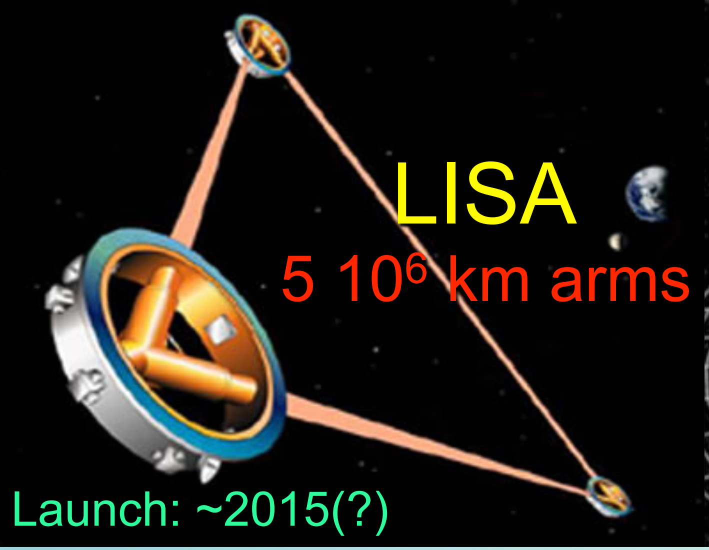

LISA

5 10^6 km arms

Launch: ~2015(?)

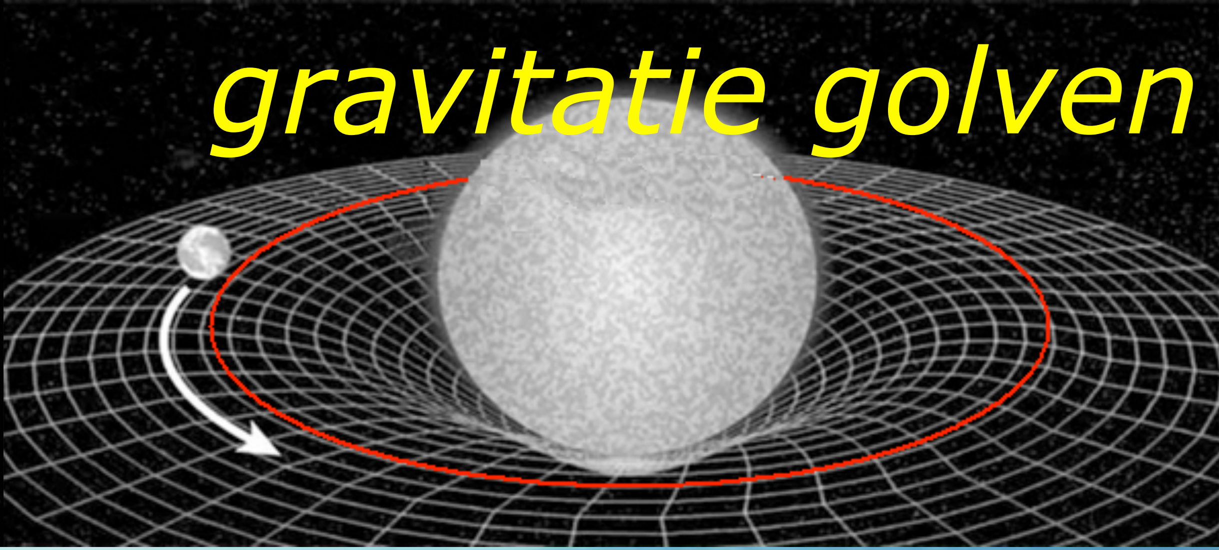

VIRGO

3 km arms

Theory

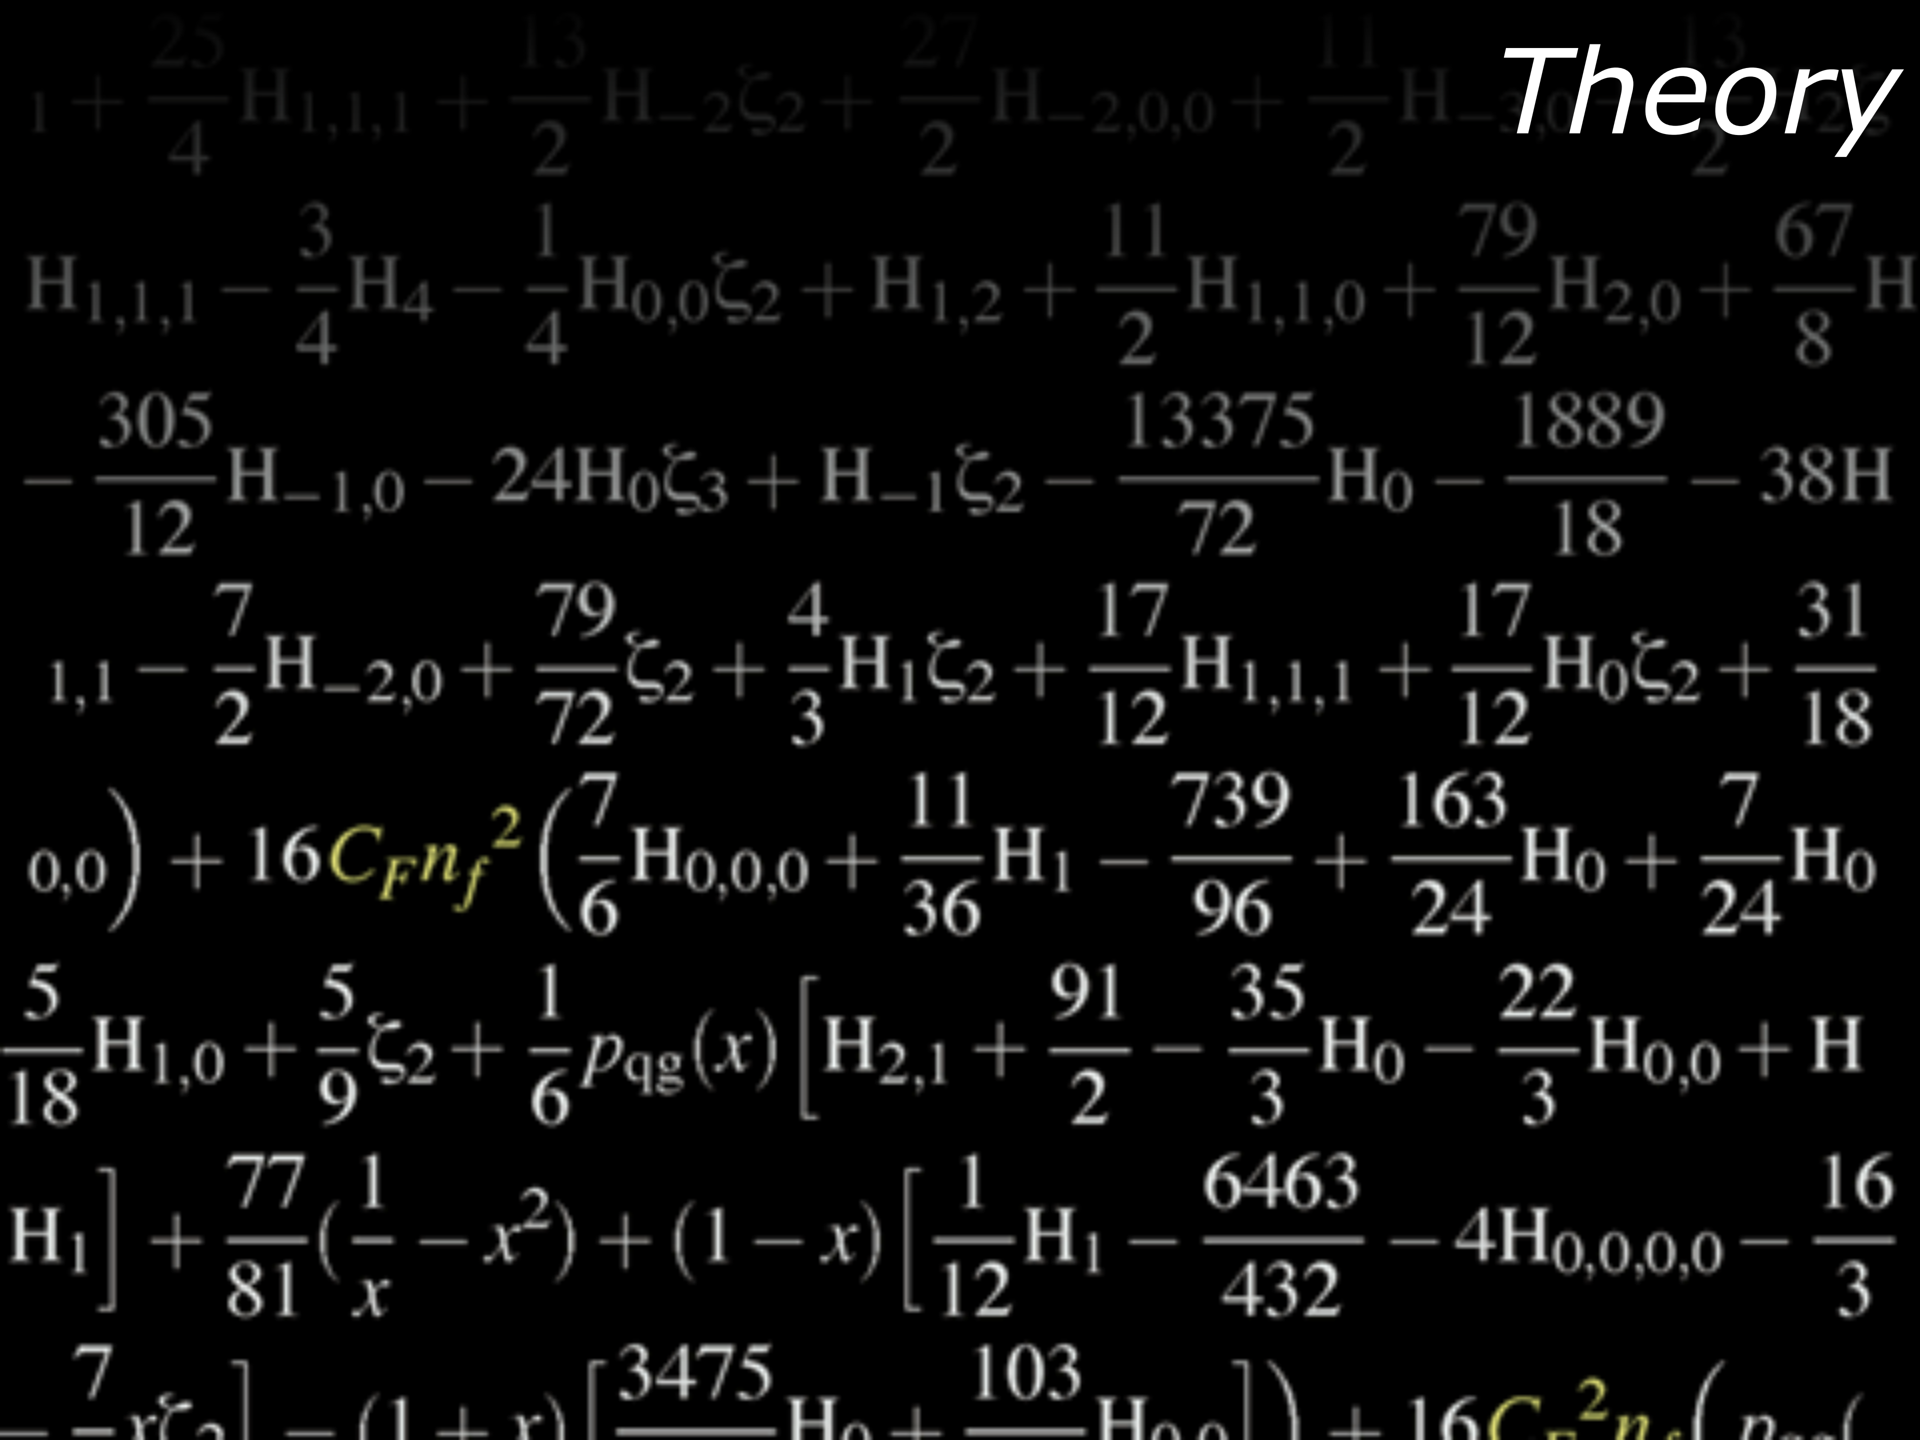

Grid computing

EGI.org at
Science Park Amsterdam

Proposal for the location of EGI.org

Stichting Nationale Computerfaciliteiten

BiG Grid

the dutch e-science grid

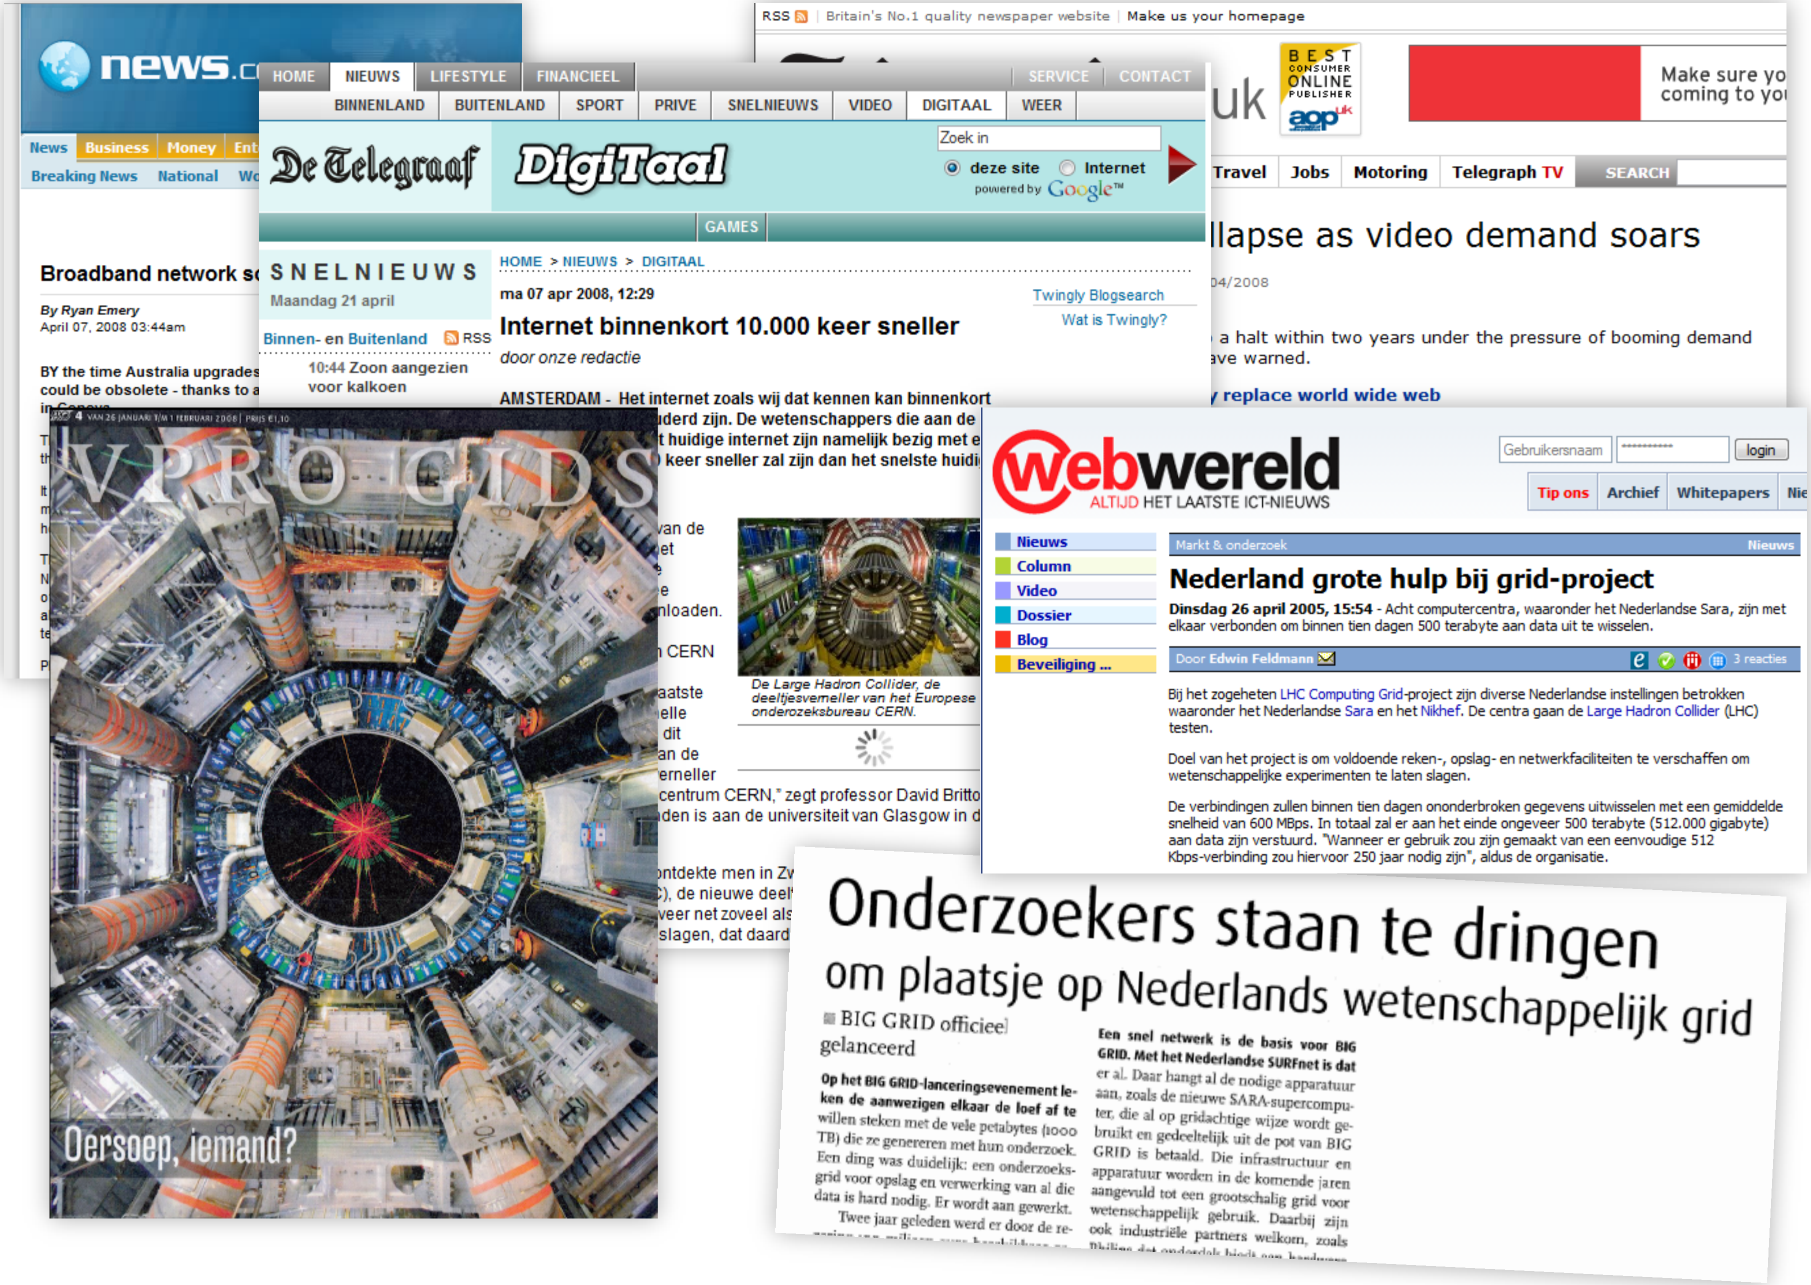

news.c
News Business Money Ent
Breaking News National Wo

Broadband network s

By Ryan Emery
April 07, 2008 03:44am

BY the time Australia upgrades
could be obsolete - thanks to a
in Geneva.

RSS | Britain's No.1 quality newspaper website | Make us your homepage

uk

BEST CONSUMER ONLINE PUBLISHER aop

Make sure yo
coming to yo

Travel Jobs Motoring Telegraph TV SEARCH

llapse as video demand soars

04/2008

a halt within two years under the pressure of booming demand
ave warned.

y replace world wide web

HOME NIEUWS LIFESTYLE FINANCIEEL SERVICE CONTACT
BINNENLAND BUITENLAND SPORT PRIVE SNELNIEUWS VIDEO DIGITAAL WEER

De Telegraaf DigiTaal

Zoek in
● deze site ○ Internet
powered by Google™

GAMES

SNELNIEUWS
Maandag 21 april

Binnen- en Buitenland RSS
10:44 Zoon aangezien voor kalkoen

HOME > NIEUWS > DIGITAAL

ma 07 apr 2008, 12:29

Twingly Blogsearch
Wat is Twingly?

Internet binnenkort 10.000 keer sneller

door onze redactie

AMSTERDAM - Het internet zoals wij dat kennen kan binnenkort
uderd zijn. De wetenschappers die aan de
t huidige internet zijn namelijk bezig met e
keer sneller zal zijn dan het snelste huidi

van de
et
e
nloaden.
n CERN
aatste
elle
dit
an de
erneller
centrum CERN," zegt professor David Britto
den is aan de universiteit van Glasgow in d

ontdekte men in Zw
), de nieuwe deel
veer net zoveel als
slagen, dat daard

De Large Hadron Collider, de deeltjesversneller van het Europese onderzoeksbureau CERN.

4 VAN 26 JANUARI T/M 1 FEBRUARI 2008 | PRIJS €1,10

VPRO GIDS

Oersoep, iemand?

webwereld
ALTIJD HET LAATSTE ICT-NIEUWS

Gebruikersnaam login

Tip ons Archief Whitepapers Nie

Nieuws
Column
Video
Dossier
Blog
Beveiliging ...

Markt & onderzoek Nieuws

Nederland grote hulp bij grid-project

Dinsdag 26 april 2005, 15:54 - Acht computercentra, waaronder het Nederlandse Sara, zijn met elkaar verbonden om binnen tien dagen 500 terabyte aan data uit te wisselen.

Door Edwin Feldmann 3 reacties

Bij het zogeheten LHC Computing Grid-project zijn diverse Nederlandse instellingen betrokken waaronder het Nederlandse Sara en het Nikhef. De centra gaan de Large Hadron Collider (LHC) testen.

Doel van het project is om voldoende reken-, opslag- en netwerkfaciliteiten te verschaffen om wetenschappelijke experimenten te laten slagen.

De verbindingen zullen binnen tien dagen ononderbroken gegevens uitwisselen met een gemiddelde snelheid van 600 MBps. In totaal zal er aan het einde ongeveer 500 terabyte (512.000 gigabyte) aan data zijn verstuurd. "Wanneer er gebruik zou zijn gemaakt van een eenvoudige 512 Kbps-verbinding zou hiervoor 250 jaar nodig zijn", aldus de organisatie.

Onderzoekers staan te dringen om plaatsje op Nederlands wetenschappelijk grid

BIG GRID officieel gelanceerd

Op het BIG GRID-lanceringsevenement leken de aanwezigen elkaar de loef af te willen steken met de vele petabytes (1000 TB) die ze genereren met hun onderzoek. Een ding was duidelijk: een onderzoeksgrid voor opslag en verwerking van al die data is hard nodig. Er wordt aan gewerkt.

Twee jaar geleden werd er door de re-

Een snel netwerk is de basis voor BIG GRID. Met het Nederlandse SURFnet is dat er al. Daar hangt al de nodige apparatuur aan, zoals de nieuwe SARA-supercomputer, die al op gridachtige wijze wordt gebruikt en gedeeltelijk uit de pot van BIG GRID is betaald. Die infrastructuur en apparatuur worden in de komende jaren aangevuld tot een grootschalig grid voor wetenschappelijk gebruik. Daarbij zijn ook industriële partners welkom, zoals

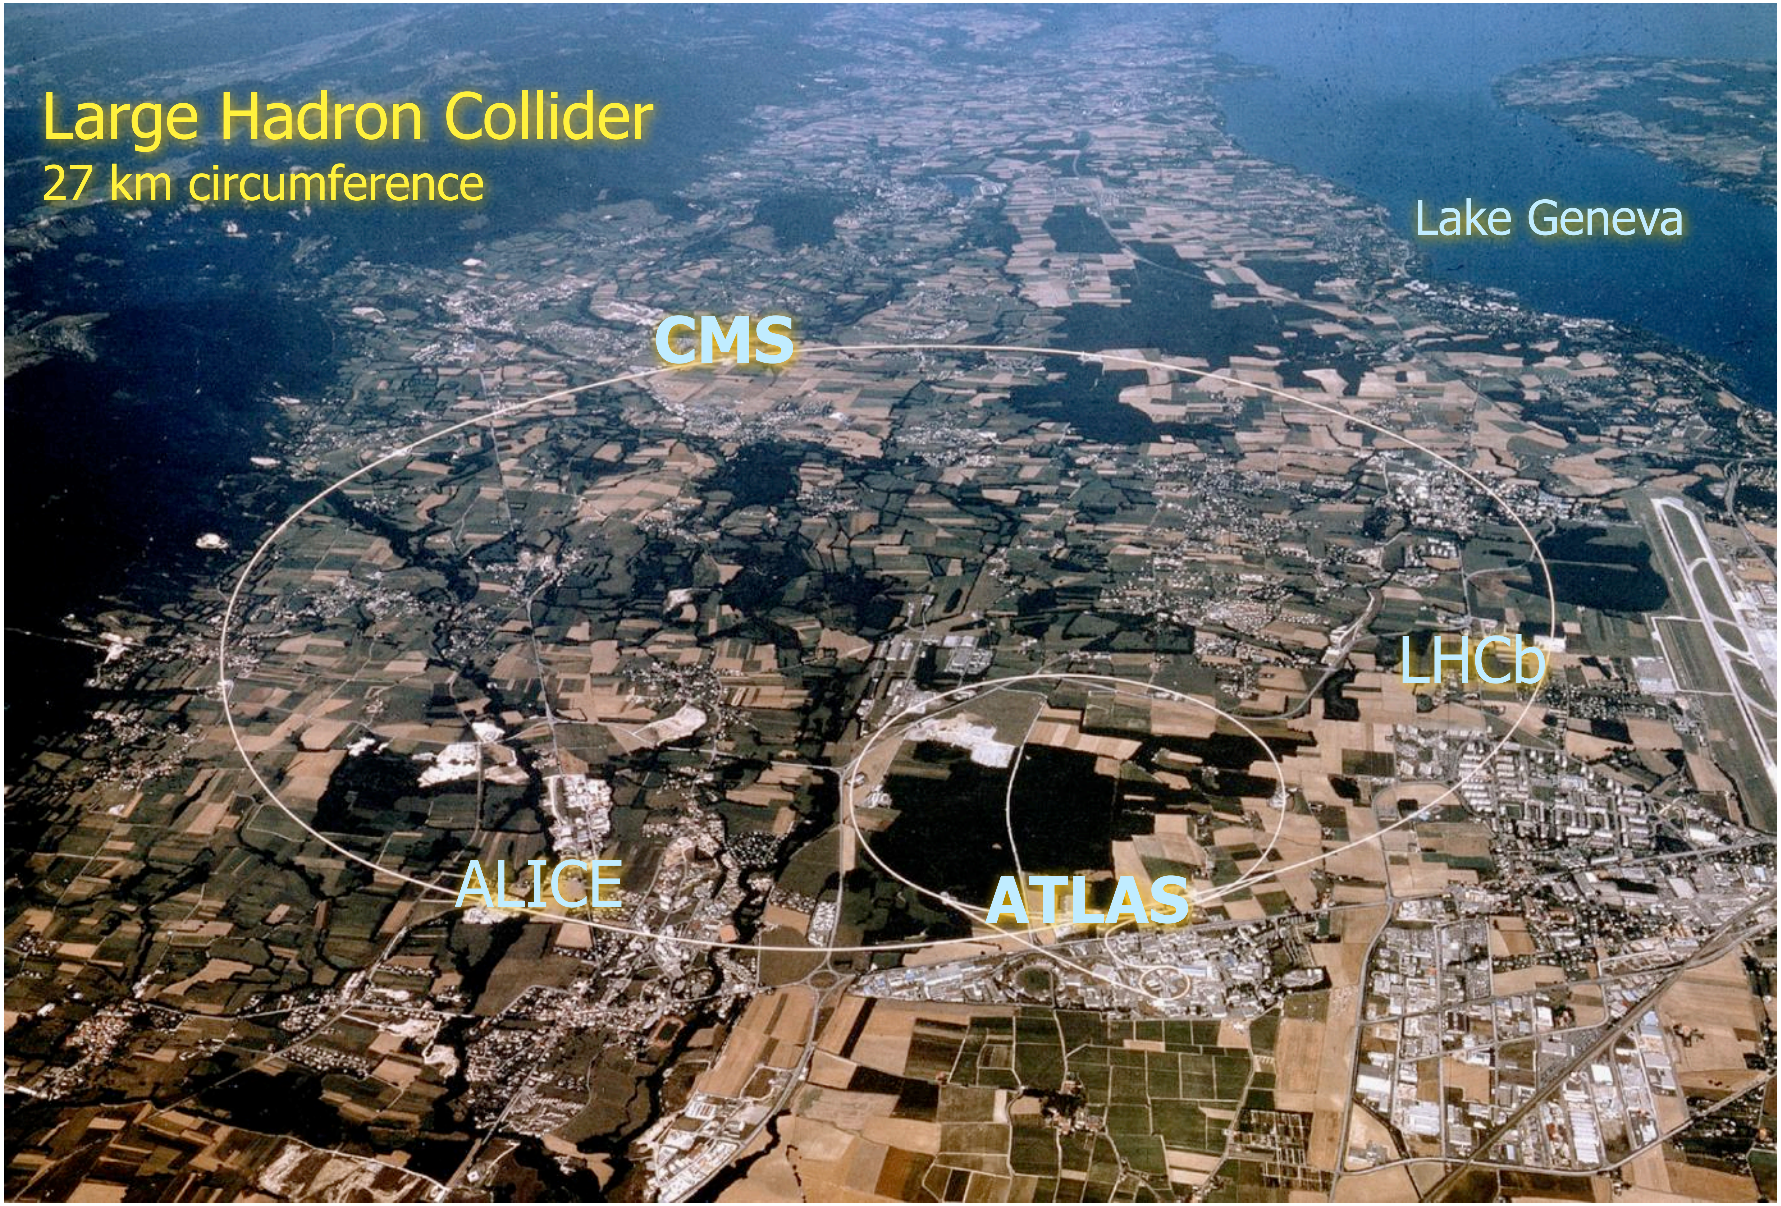
Large Hadron Collider
27 km circumference
Lake Geneva
CMS
LHCb
ALICE
ATLAS

Large Hadron Collider

27 km circumference

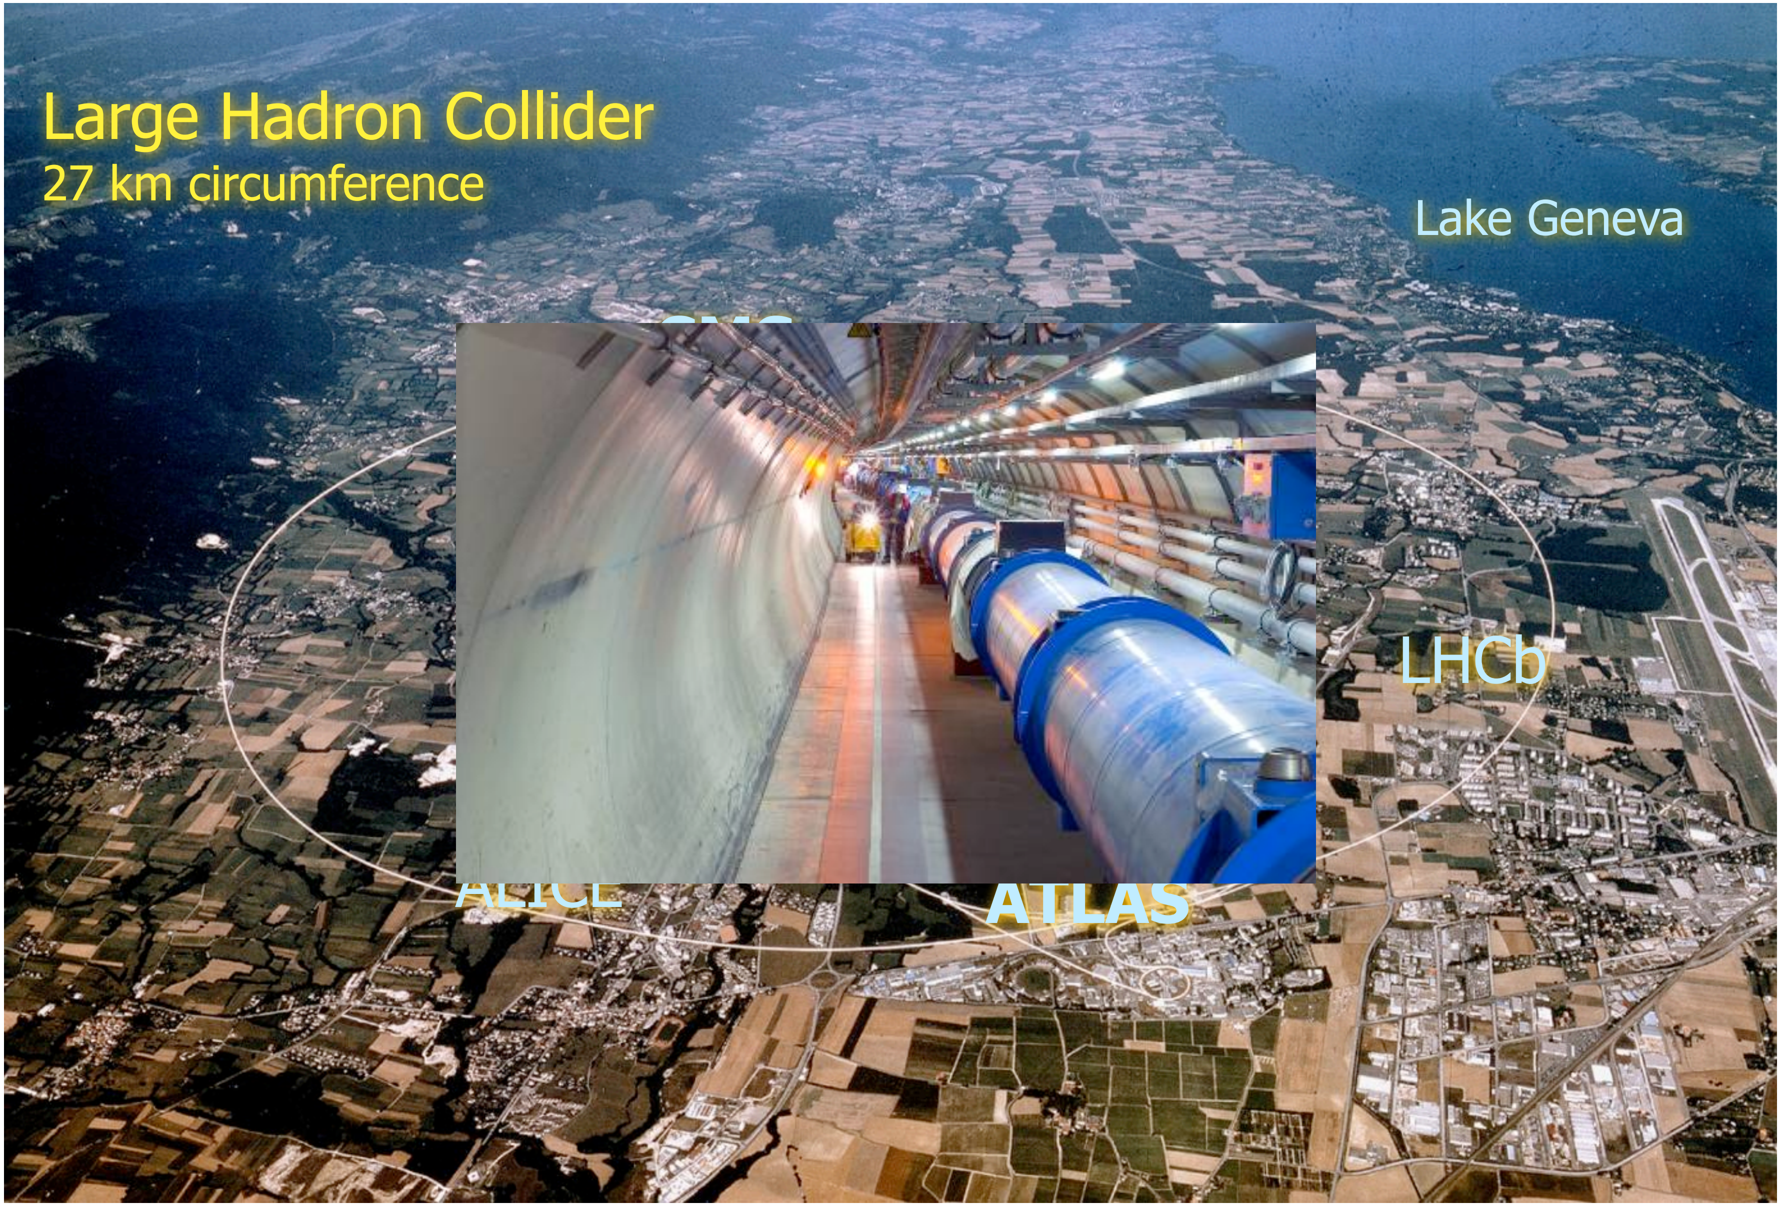

Large Hadron Collider

27 km circumference

The LHC Computing Challenge

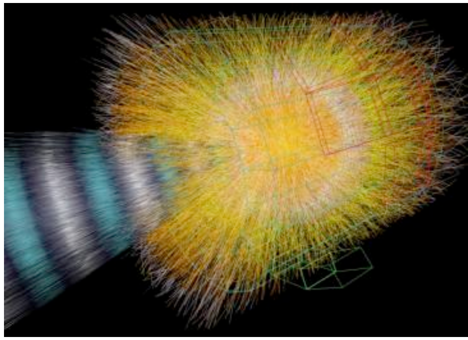

- **The scale and complexity of the data**
  ➔ **15 PetaBytes of new data each year**

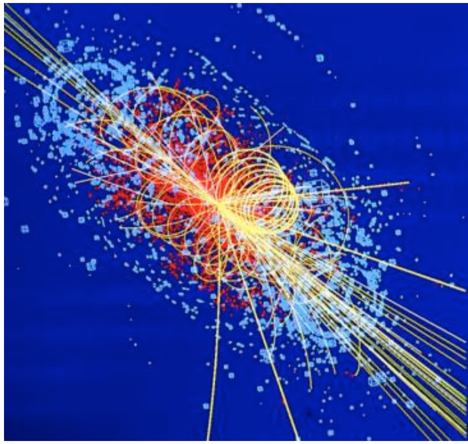

- **The computing capacity to support 7,000 researchers all actively analysing the data**
  ➔ **60’000 of (today's) fastest CPUs**

- **The way in which the data is accessed will depend on the physics that emerges**

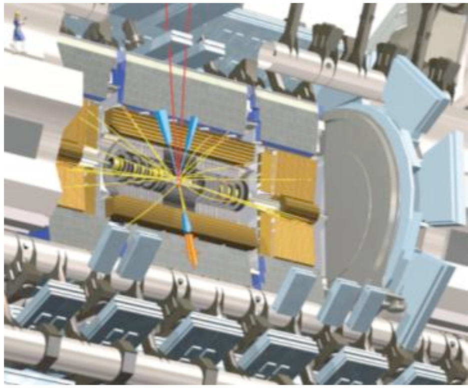

Astronomy & Astrophysics

Astronomy & Astrophysics

Astronomy & Astrophysics

LOFAR large distributed radio telescope

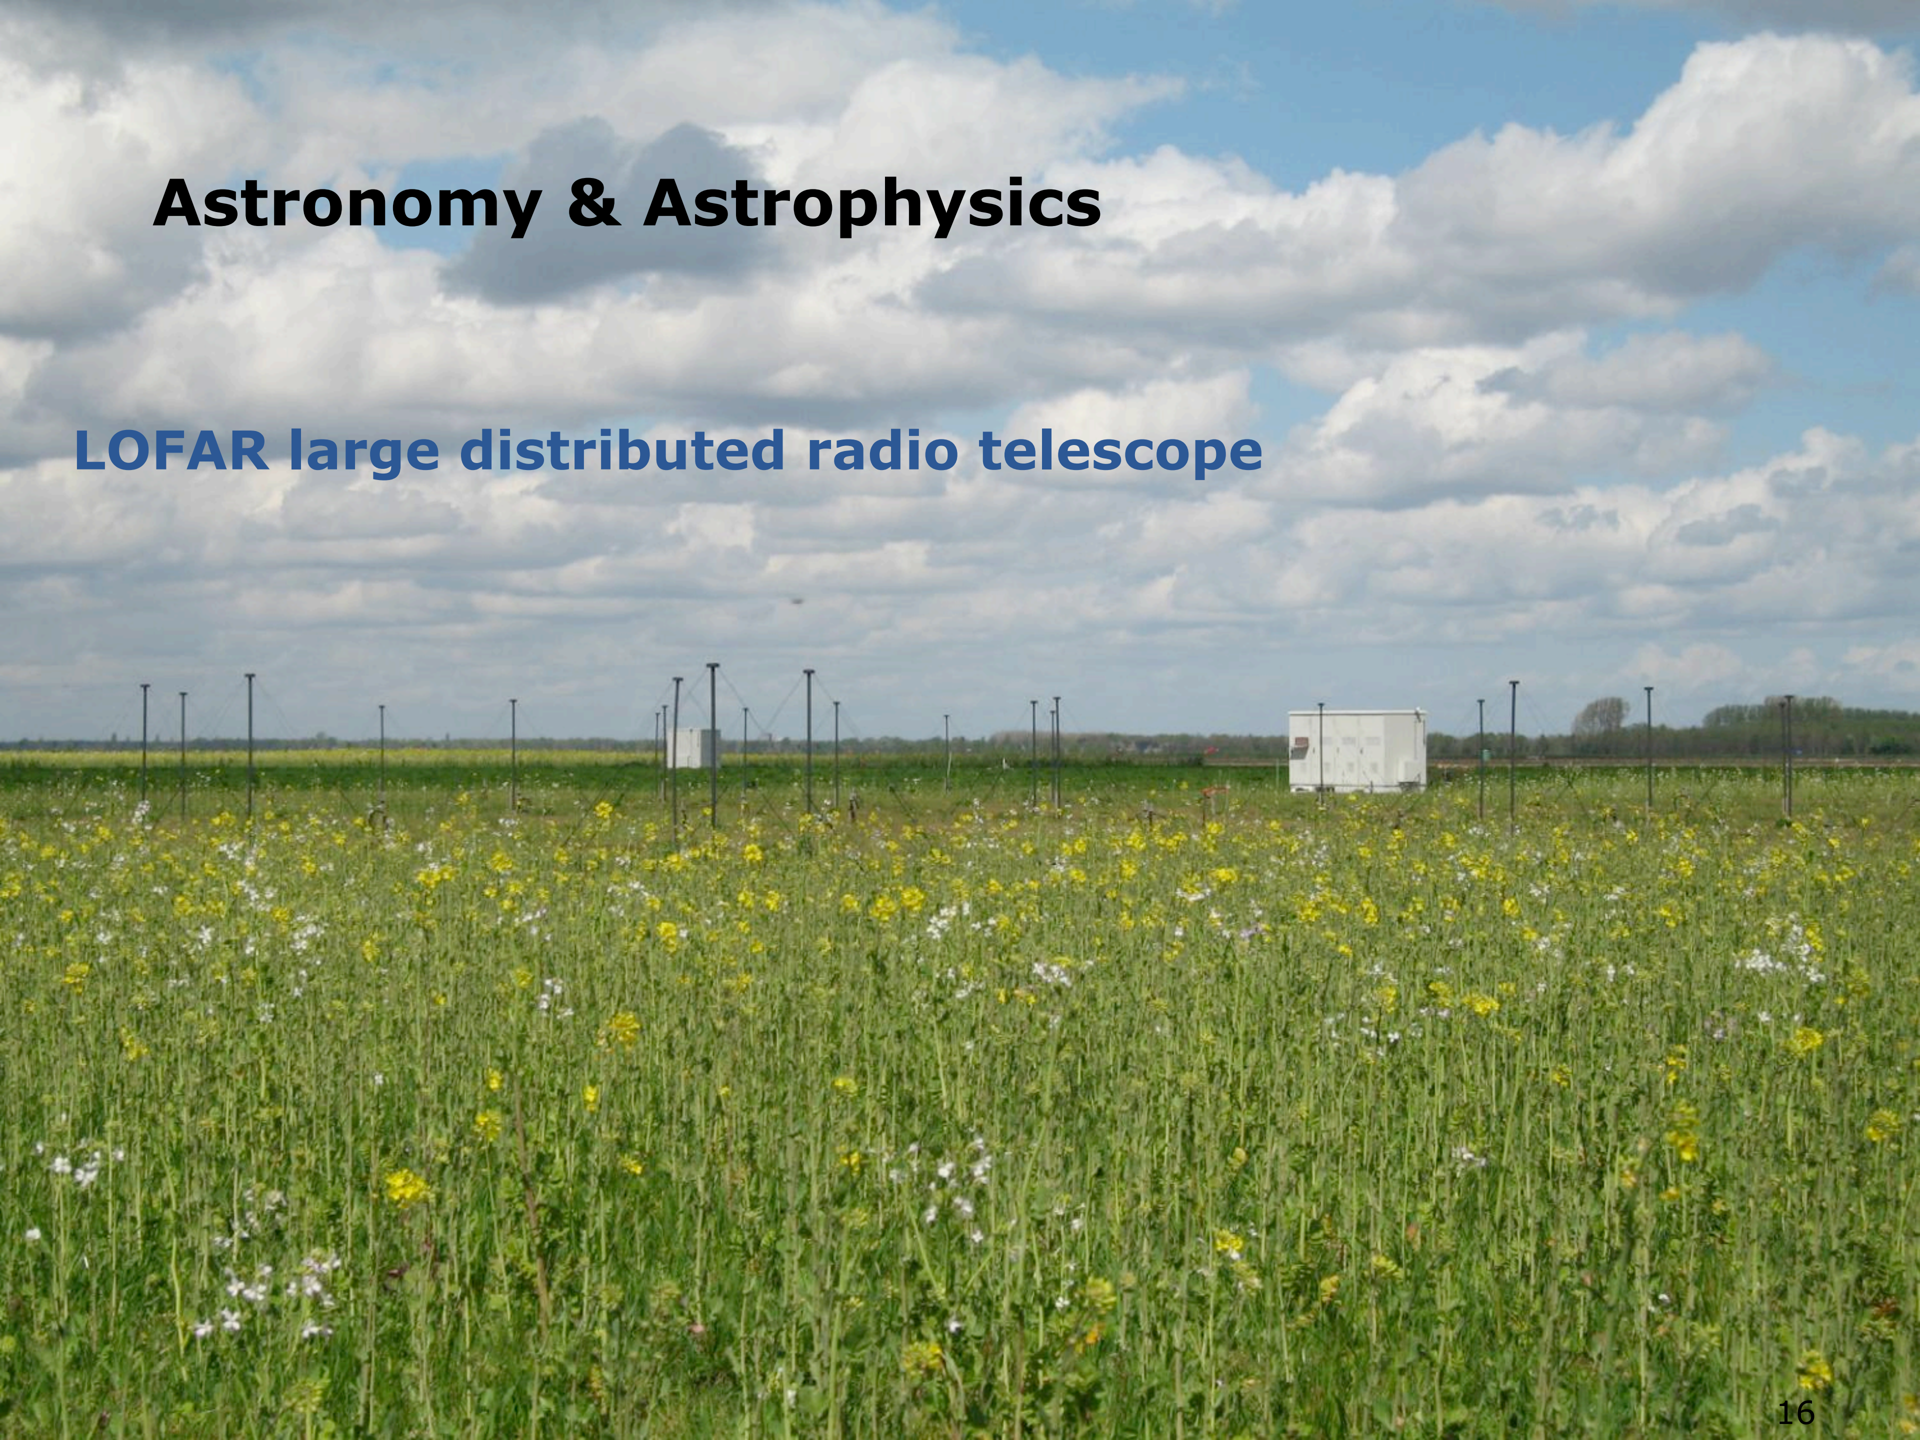

Astronomy & Astrophysics

LOFAR large distributed radio telescope

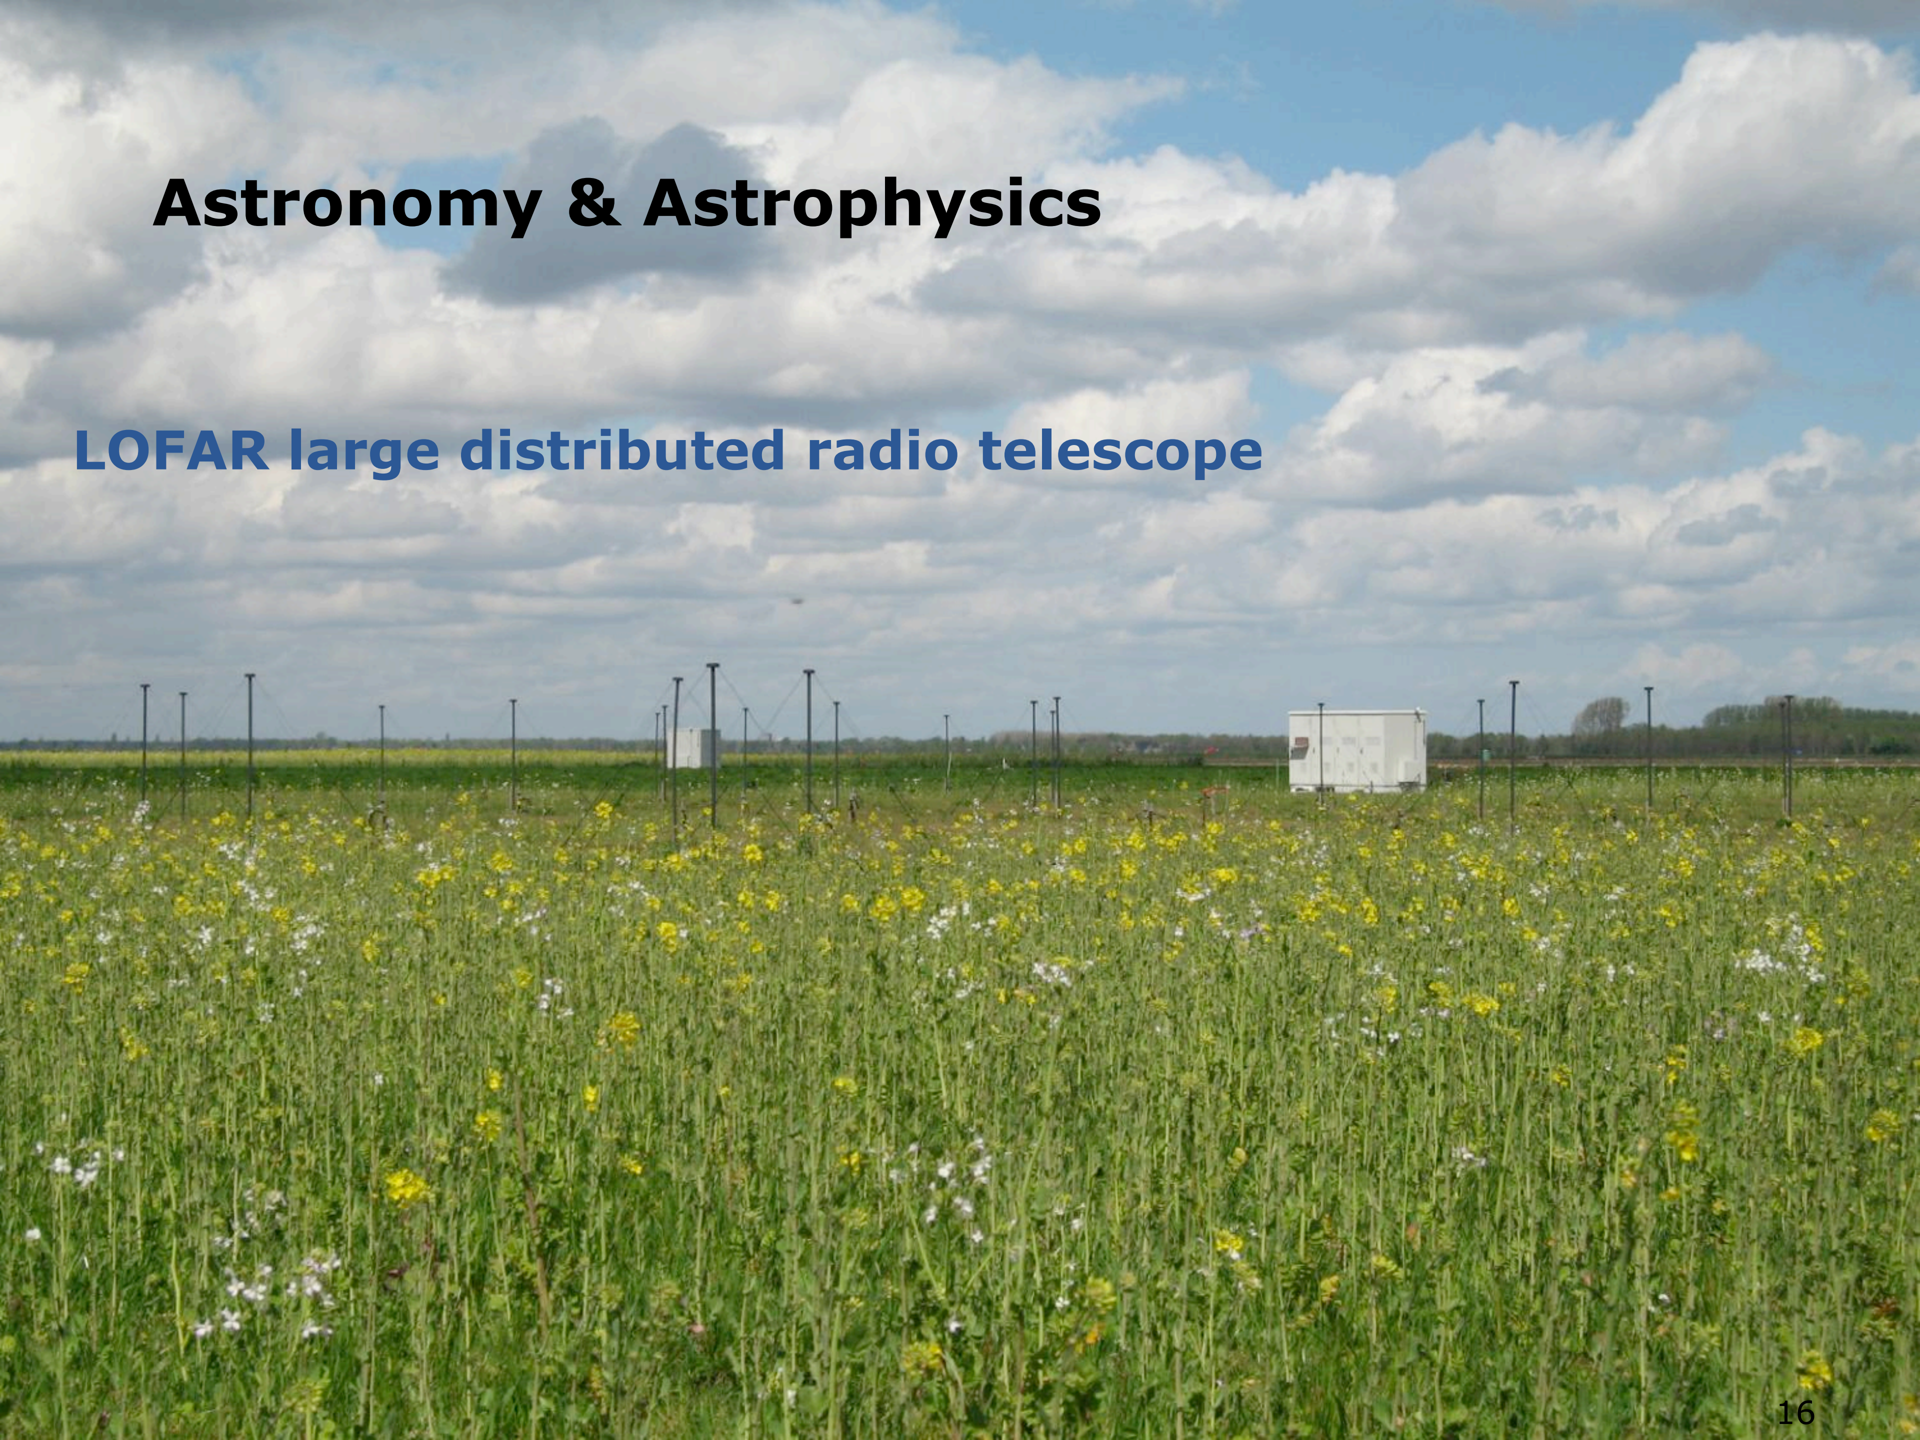

Astronomy & Astrophysics

LOFAR large distributed radio telescop

AUGER & ARGO Cosmic Ray Observator

Functional MRI analysis

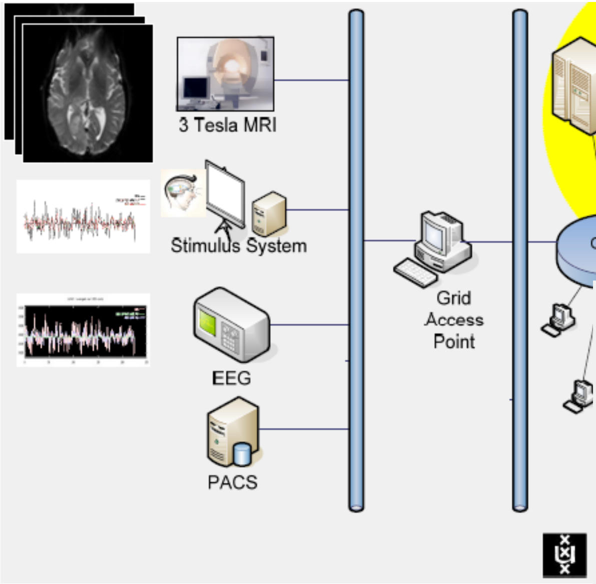

Storage of fMRI research data for sharing between groups and processing of image alignments

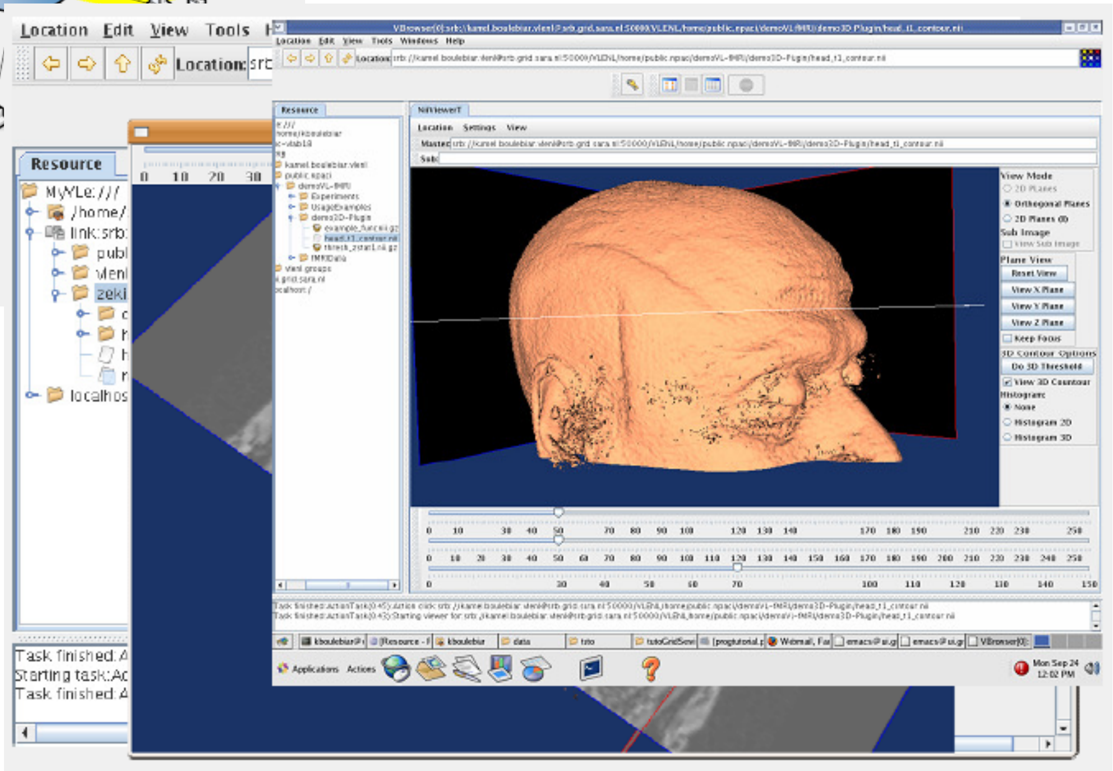

Task finished:ActionTask(0:45):Action click: srb://kamel.boulebiar.vlenl@srb.grid.sara.nl:50000/VLENL/home/public.npaci/demoVL-fMRI/demo3D-Plugin/head_t1_contour.nii
Task finished:ActionTask(0:43):Starting viewer for: srb://kamel.boulebiar.vlenl@srb.grid.sara.nl:50000/VLENL/home/public.npaci/demoVL-fMRI/demo3D-Plugin/head_t1_contour.nii

Research work by:
Silvia Olabarriaga (AMC, UvA)
Tristan Glatard (IvI,UvA)
Abdullah Ozsoy (IvI,UvA)

In silico drug discovery

- Diseases such as HIV/AIDS, SRAS, Bird Flu, Malaria etc. are a threat to public health due to world wide exchanges and circulation of persons
- Grids open new perspectives to *in silico* drug discovery
  - Reduced cost and adding an accelerating factor in the search for new drugs

International collaboration is required for:

- Early detection
- Epidemiological watch
- Prevention
- Search for new drugs
- Search for vaccines

WISDOM
Initiative for grid-enabled drug discovery against neglected and emergent diseases

Areas reporting confirmed occurence of H5N1* avian influenza in poultry and wild birds since 2003

Status as of 07 April 2006

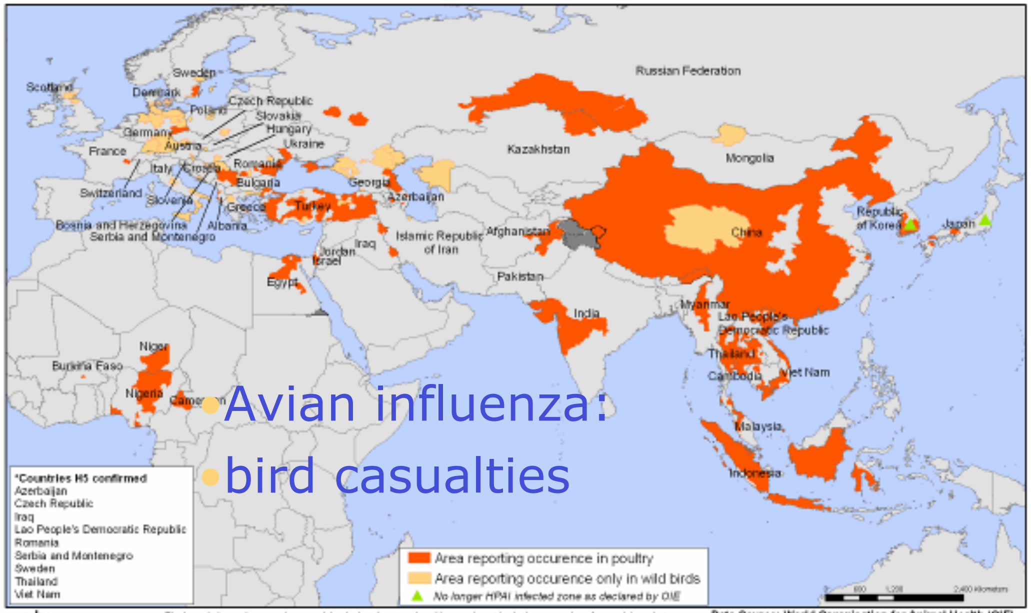

- Avian influenza:
- bird casualties

World Health Organization
© WHO 2006. All rights reserved

The boundaries and names shown and the designations used on this map do not imply the expression of any opinion whatsoever on the part of the World Health Organization concerning the legal status of any country, territory, city or area or of its authorities, or concerning the delimitation of its frontiers or boundaries. Dotted lines on maps represent approximate border lines for which there may not yet be full agreement.

Data Source: World Organisation for Animal Health (OIE) and national governments
Map Production: Public Health Mapping and GIS
Communicable Diseases (CDS) World Health Organization

Fusion

Commercial exploitation of fusion energy still needs to solve several outstanding problems requiring exceptional computing facilities including supercomputers and cluster-based grids

- Ion Kinetic Transport
- Massive Ray Tracing
- Stellarator Optimization

Interworking course-grained clusters and MPP systems across both the EGEE and DEISA grids

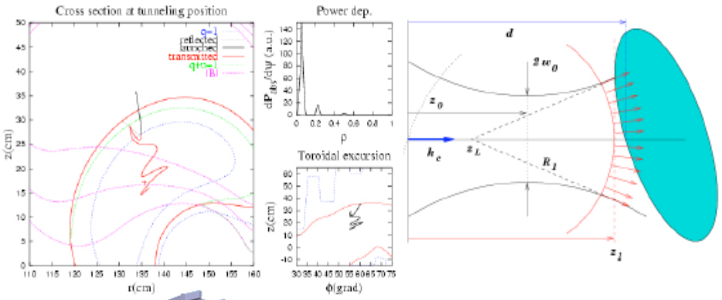

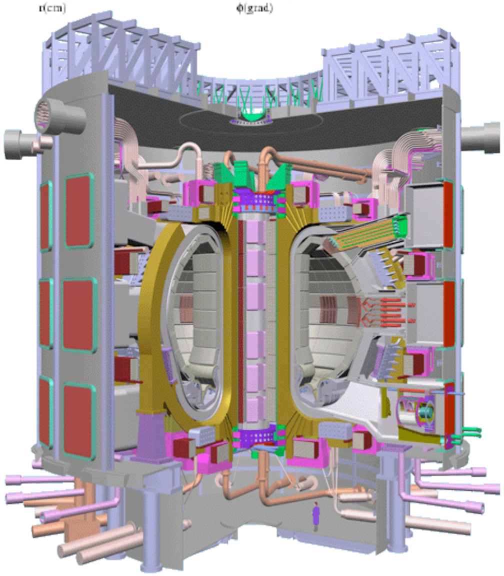

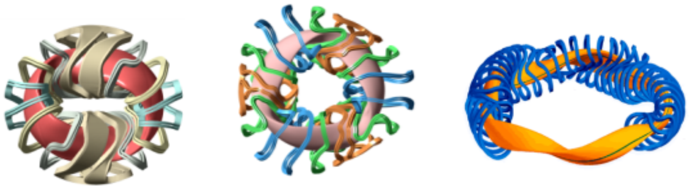

Grids in Science

The Grid is 'more of everything' as science struggles to deal with ever increasing complexity

more than one place on earth

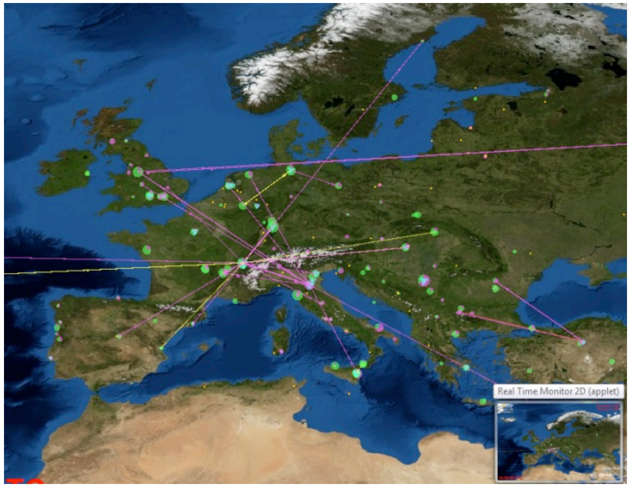

more than one computer

more than one science!

more than ...

Three essential ingredients for Grid

‘Access computing like the electrical power grid’

A grid combines resources that

- Are not managed by a single organization
- Use a common, open protocol … that is general purpose
- Provide additional qualities of service, *i.e.*, are usable as a collective and transparent resource

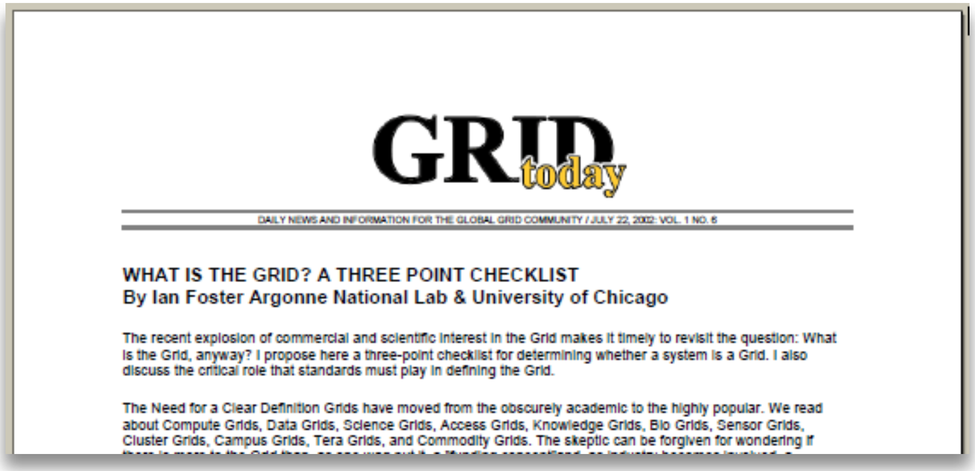

GRID today

DAILY NEWS AND INFORMATION FOR THE GLOBAL GRID COMMUNITY / JULY 22, 2002: VOL. 1 NO. 6

WHAT IS THE GRID? A THREE POINT CHECKLIST
By Ian Foster Argonne National Lab & University of Chicago

The recent explosion of commercial and scientific interest in the Grid makes it timely to revisit the question: What is the Grid, anyway? I propose here a three-point checklist for determining whether a system is a Grid. I also discuss the critical role that standards must play in defining the Grid.

The Need for a Clear Definition Grids have moved from the obscurely academic to the highly popular. We read about Compute Grids, Data Grids, Science Grids, Access Grids, Knowledge Grids, Bio Grids, Sensor Grids, Cluster Grids, Campus Grids, Tera Grids, and Commodity Grids. The skeptic can be forgiven for wondering if

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

What is Grid?

Cycle scavenging

- harvest idle compute power
- improve RoI on desktops

Cluster computing and storage

- What-if scenarios
- Physics event analysis
- Improve Data Centre Utilization

Cross-domain resource sharing

- more than one organization
- more than one application
- more than one …

- open protocols
- collective service

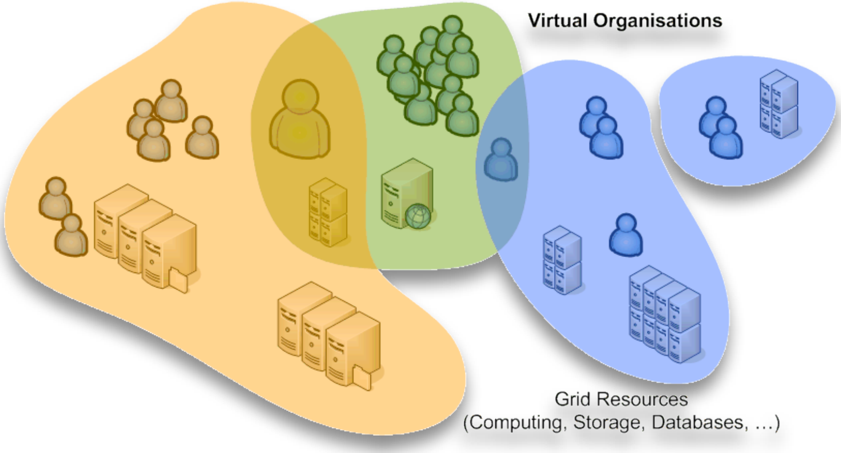

What is Grid?

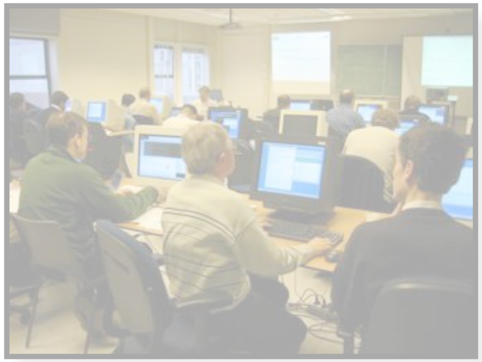

Cycle scavenging

- harvest idle compute power
- improve RoI on desktops

Cluster computing and storage

- What-if scenarios
- Physics event analysis
- Improve Data Centre Utilization

Cross-domain resource sharing

- more than one organization
- more than one application
- more than one …

- open protocols
- collective service

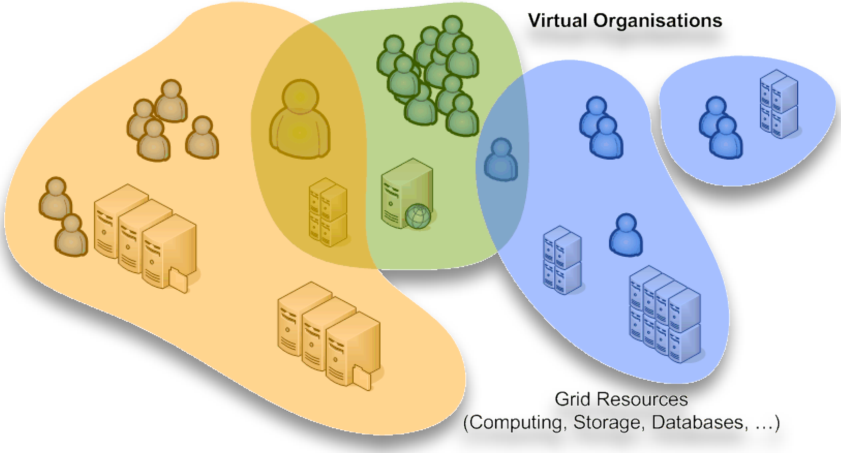

e-Infrastructure for Research

World Wide Web (1990) – sharing information

Grid (1997) – sharing computers and storage

Clouds (2007) – commoditizing the Grid

more than one place on earth

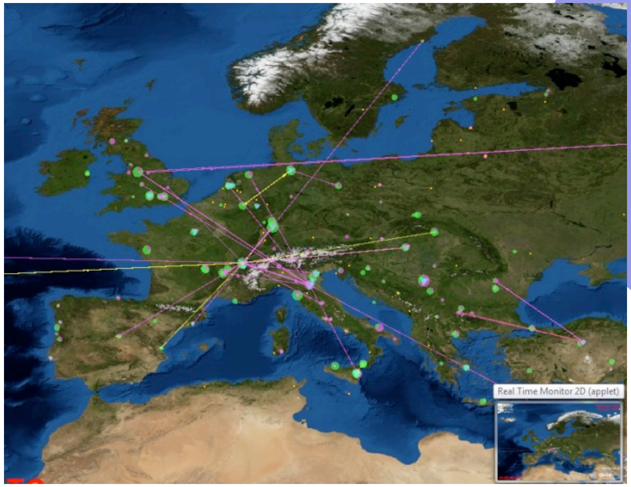

What Makes e-Research Happen ...

more than one computer

more than one science!

more than ...

Building Grid Infrastructures

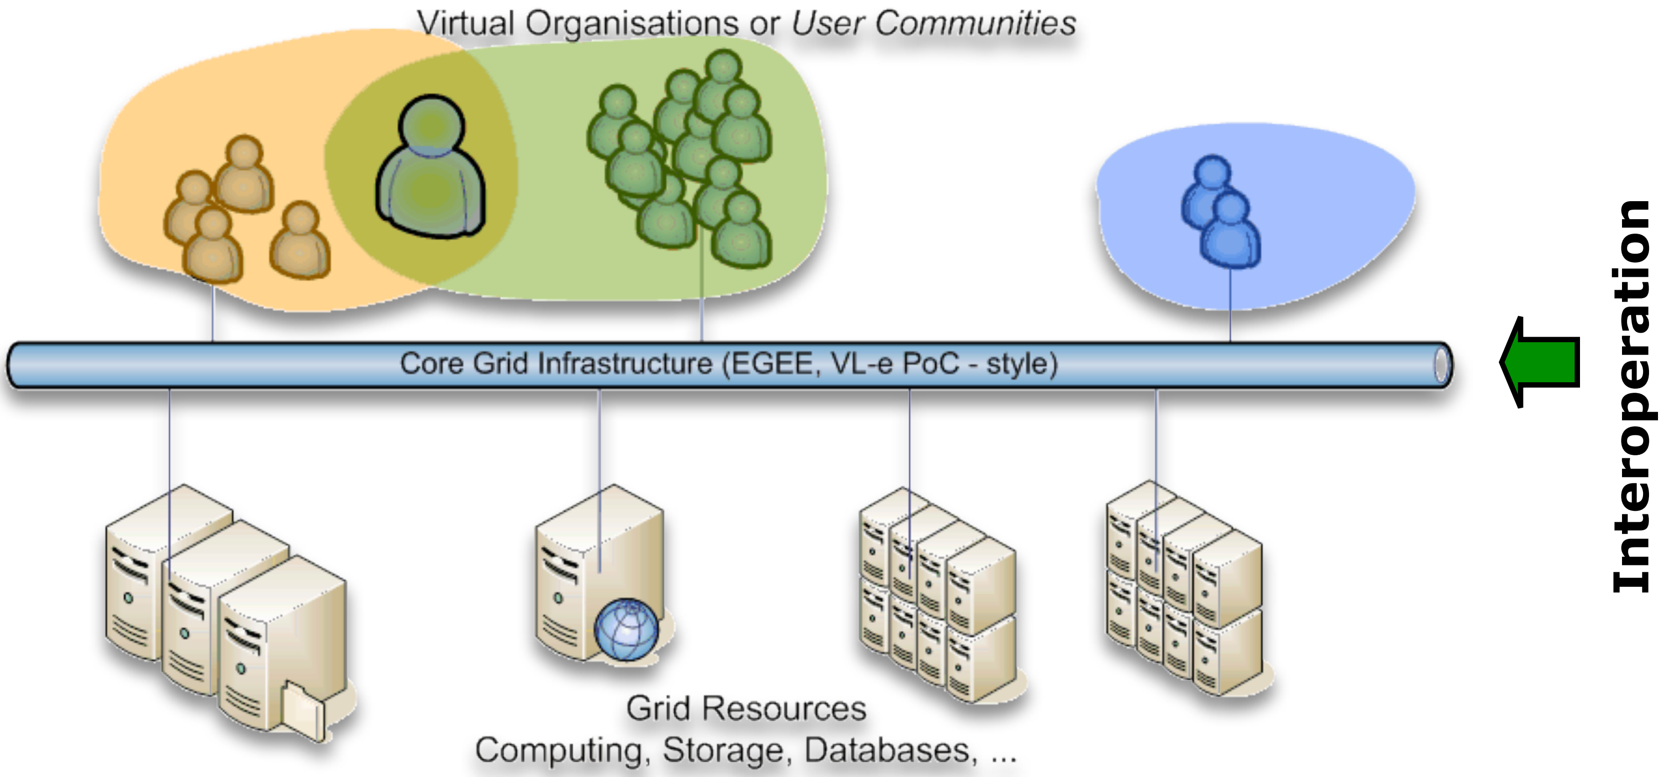
Virtual Organisations or *User Communities*
Core Grid Infrastructure (EGEE, VL-e PoC - style)
Interoperation
Grid Resources
Computing, Storage, Databases, ...

- Protocols: common syntax and sematics for grid operations
- APIs: making grid concepts accessible from the applications
- Portals and workflows: bridging the end-user gap

Standards

- Standards, such as those by IETF, OASIS, OGF, &c aid interoperability and reduce vendor lock-in

- as you go higher up the stack, you get less synergy
  - Transport: IP/TCP, HTTP, TLS/SSL, &c well agreed
  - Web services: SOAP used to be the solution for all …
    … but 'Web 2.0' shows alternatives tailored to
    specific applications gaining popularity
  - Grid standards:
    low-level job submission (BES, JSDL), management (DRMAA), basic security (OGSA-BSP Core, SC), high-level application toolkits (SAGA, GAT)

see also http://www.ogf.org/

Working at scale

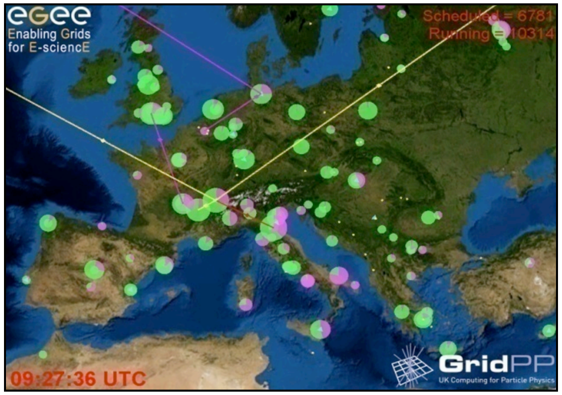
egee
Enabling Grids for E-sciencE
Scheduled = 6781
Running = 10314
09:27:36 UTC
GridPP
UK Computing for Particle Physics

Grid is an error amplifier …

'passive' controls are needed
to push work away
from failing resources

Failure-ping-pong – or *creeper and reaper* revisited

Resource information systems are the backbone of any real-life grid

Grid is much like the 'Wild West'

- almost unlimited possibilities – but as a community plan for scaling issues, and a novel environment
- users and providers *need to interact* and articulate needs

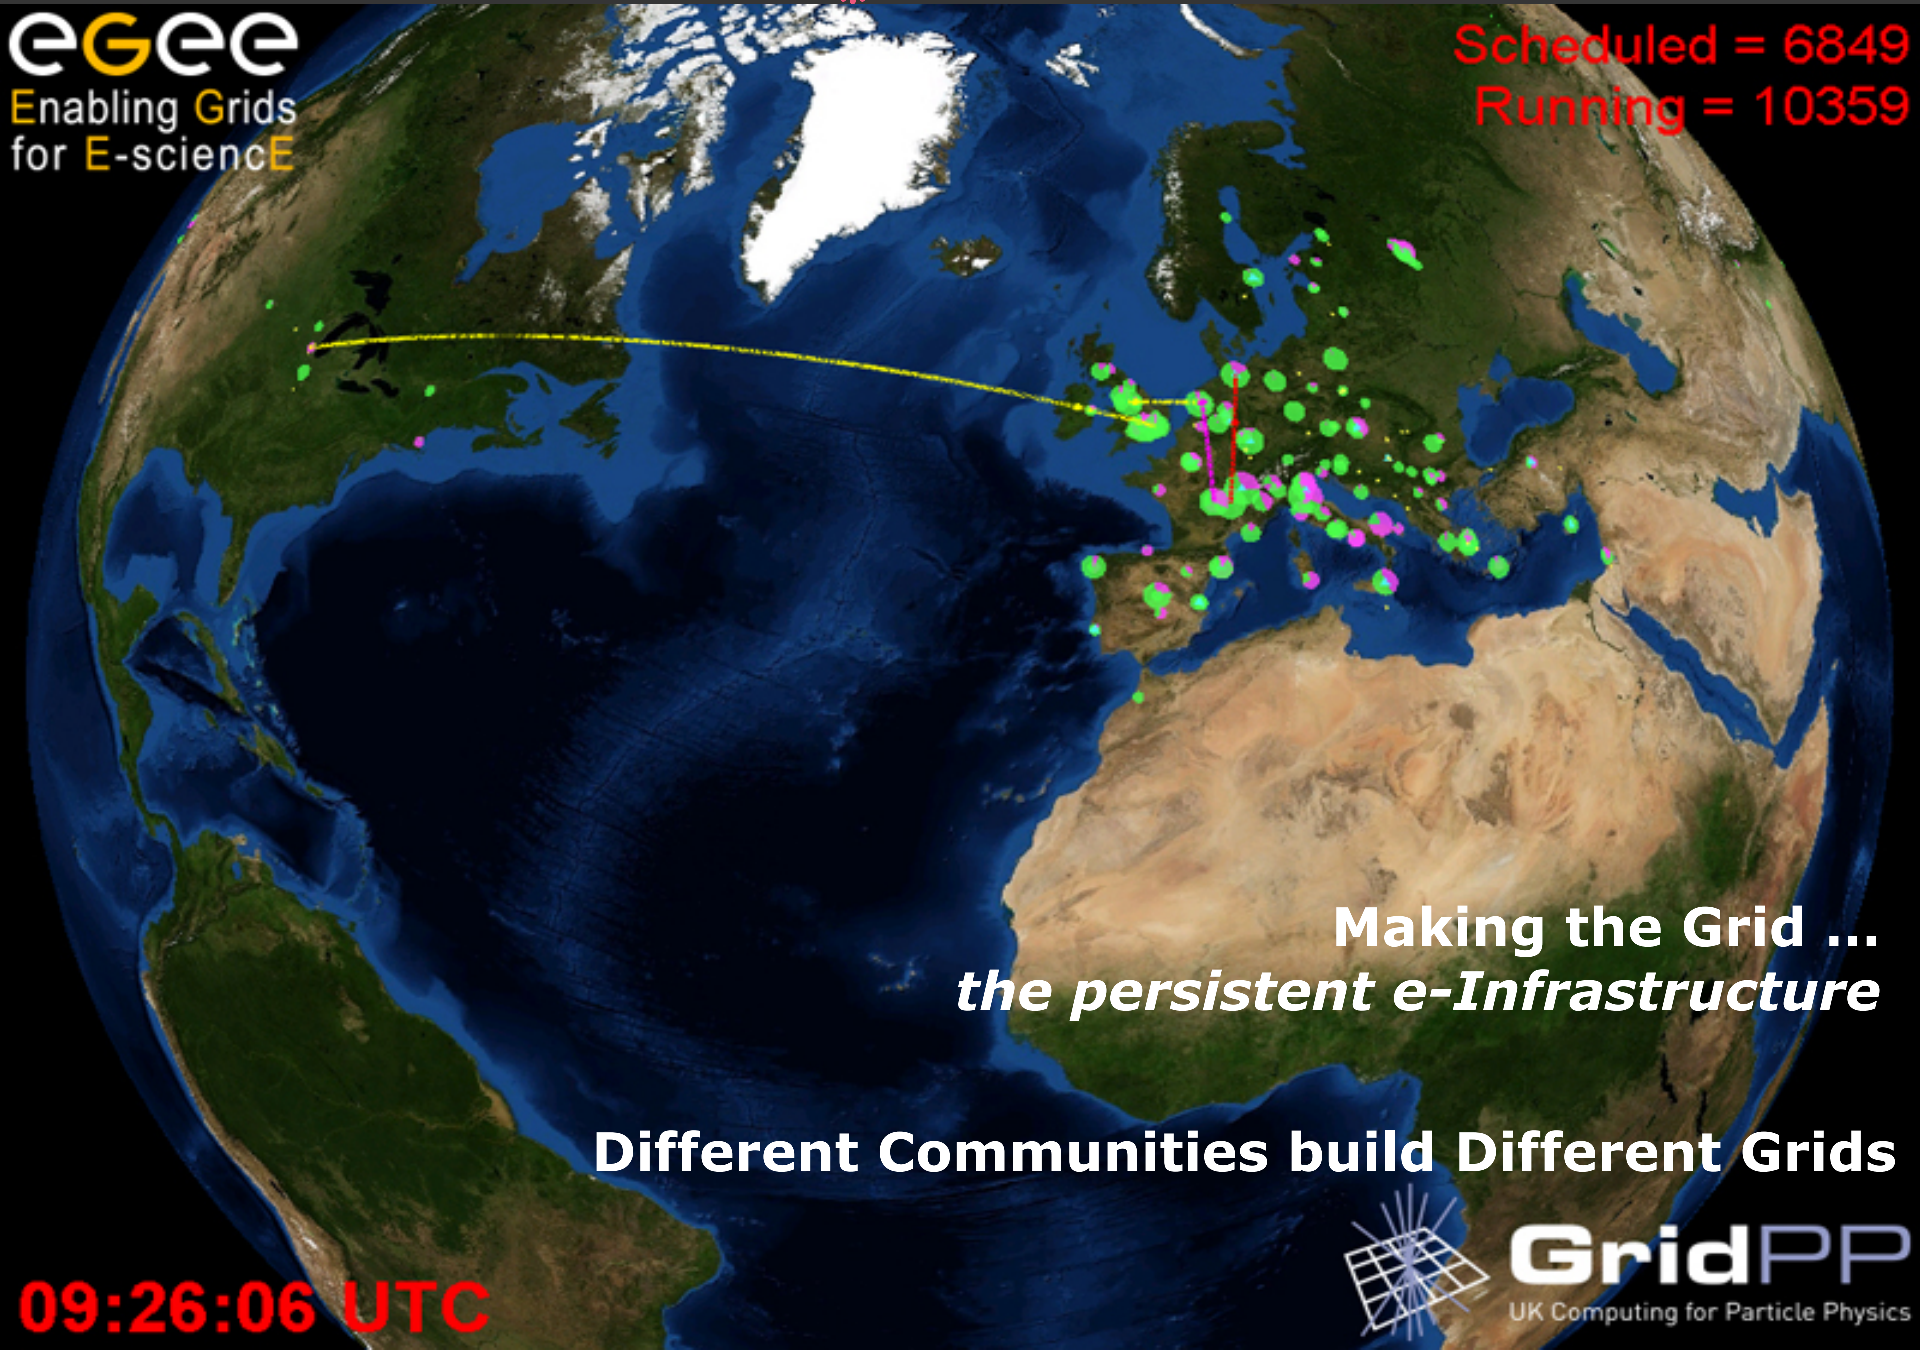
eGee
Enabling Grids for E-sciencE
Scheduled = 6849
Running = 10359
09:26:06 UTC
GridPP
UK Computing for Particle Physics

Making the Grid …
the persistent e-Infrastructure

Different Communities build Different Grids

Enabling the Grid – the Network

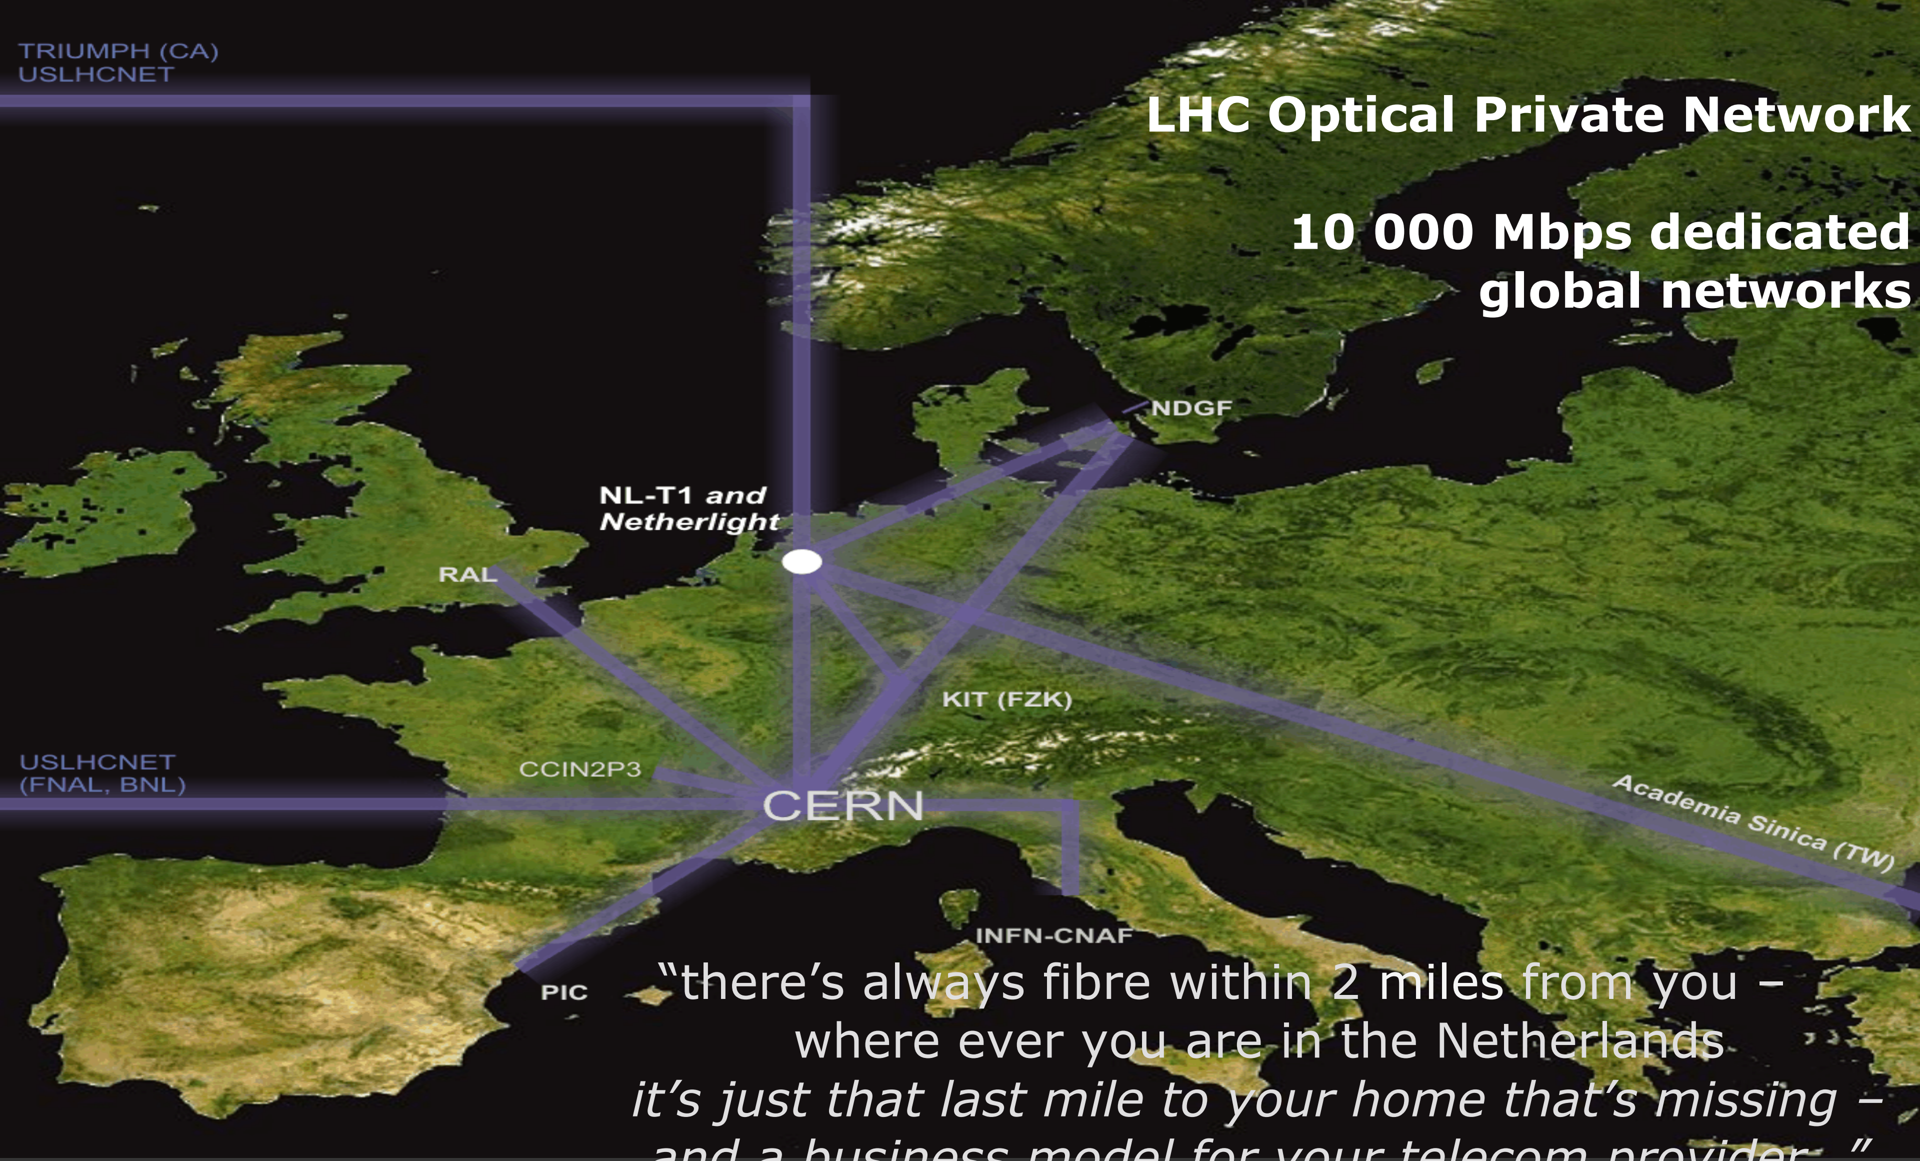

LHC Optical Private Network

10 000 Mbps dedicated global networks

"there's always fibre within 2 miles from you – where ever you are in the Netherlands
it's just that last mile to your home that's missing – and a business model for your telecom provider..."

There's always a network close to you

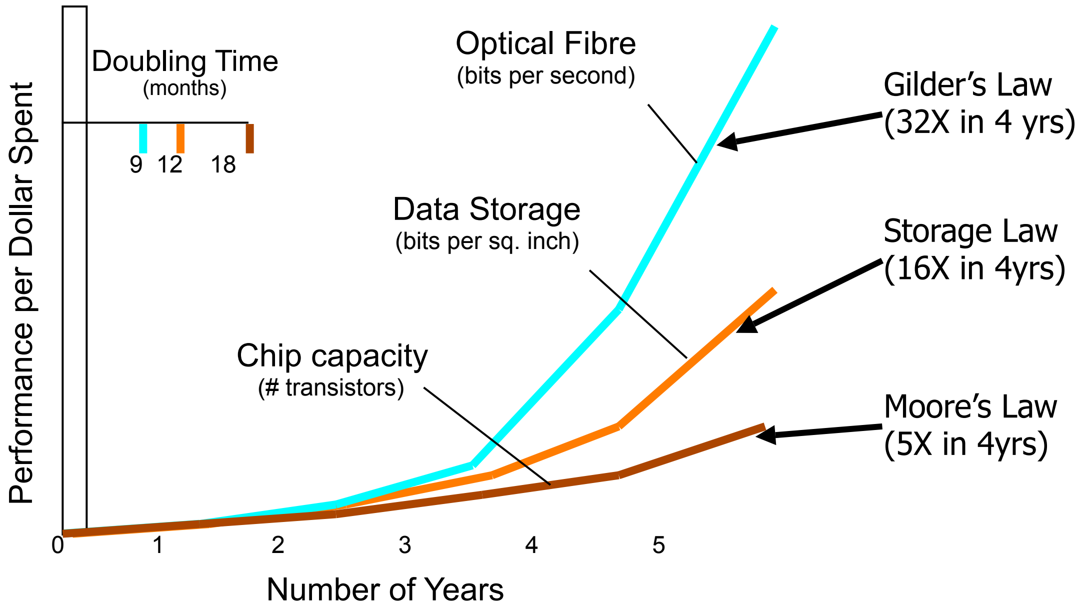

NL Light

SURF NET

SURFnet pioneered 'lambda' and hybrid networks in the world

- and likely contributed to the creation of a market for 'dark fibre' in the Netherlands

There's always fibre within 2 miles from you – where ever you are!

(it's just that last mile to your home that's missing – and the business model of your telecom provider…)

LHCOPN – current status

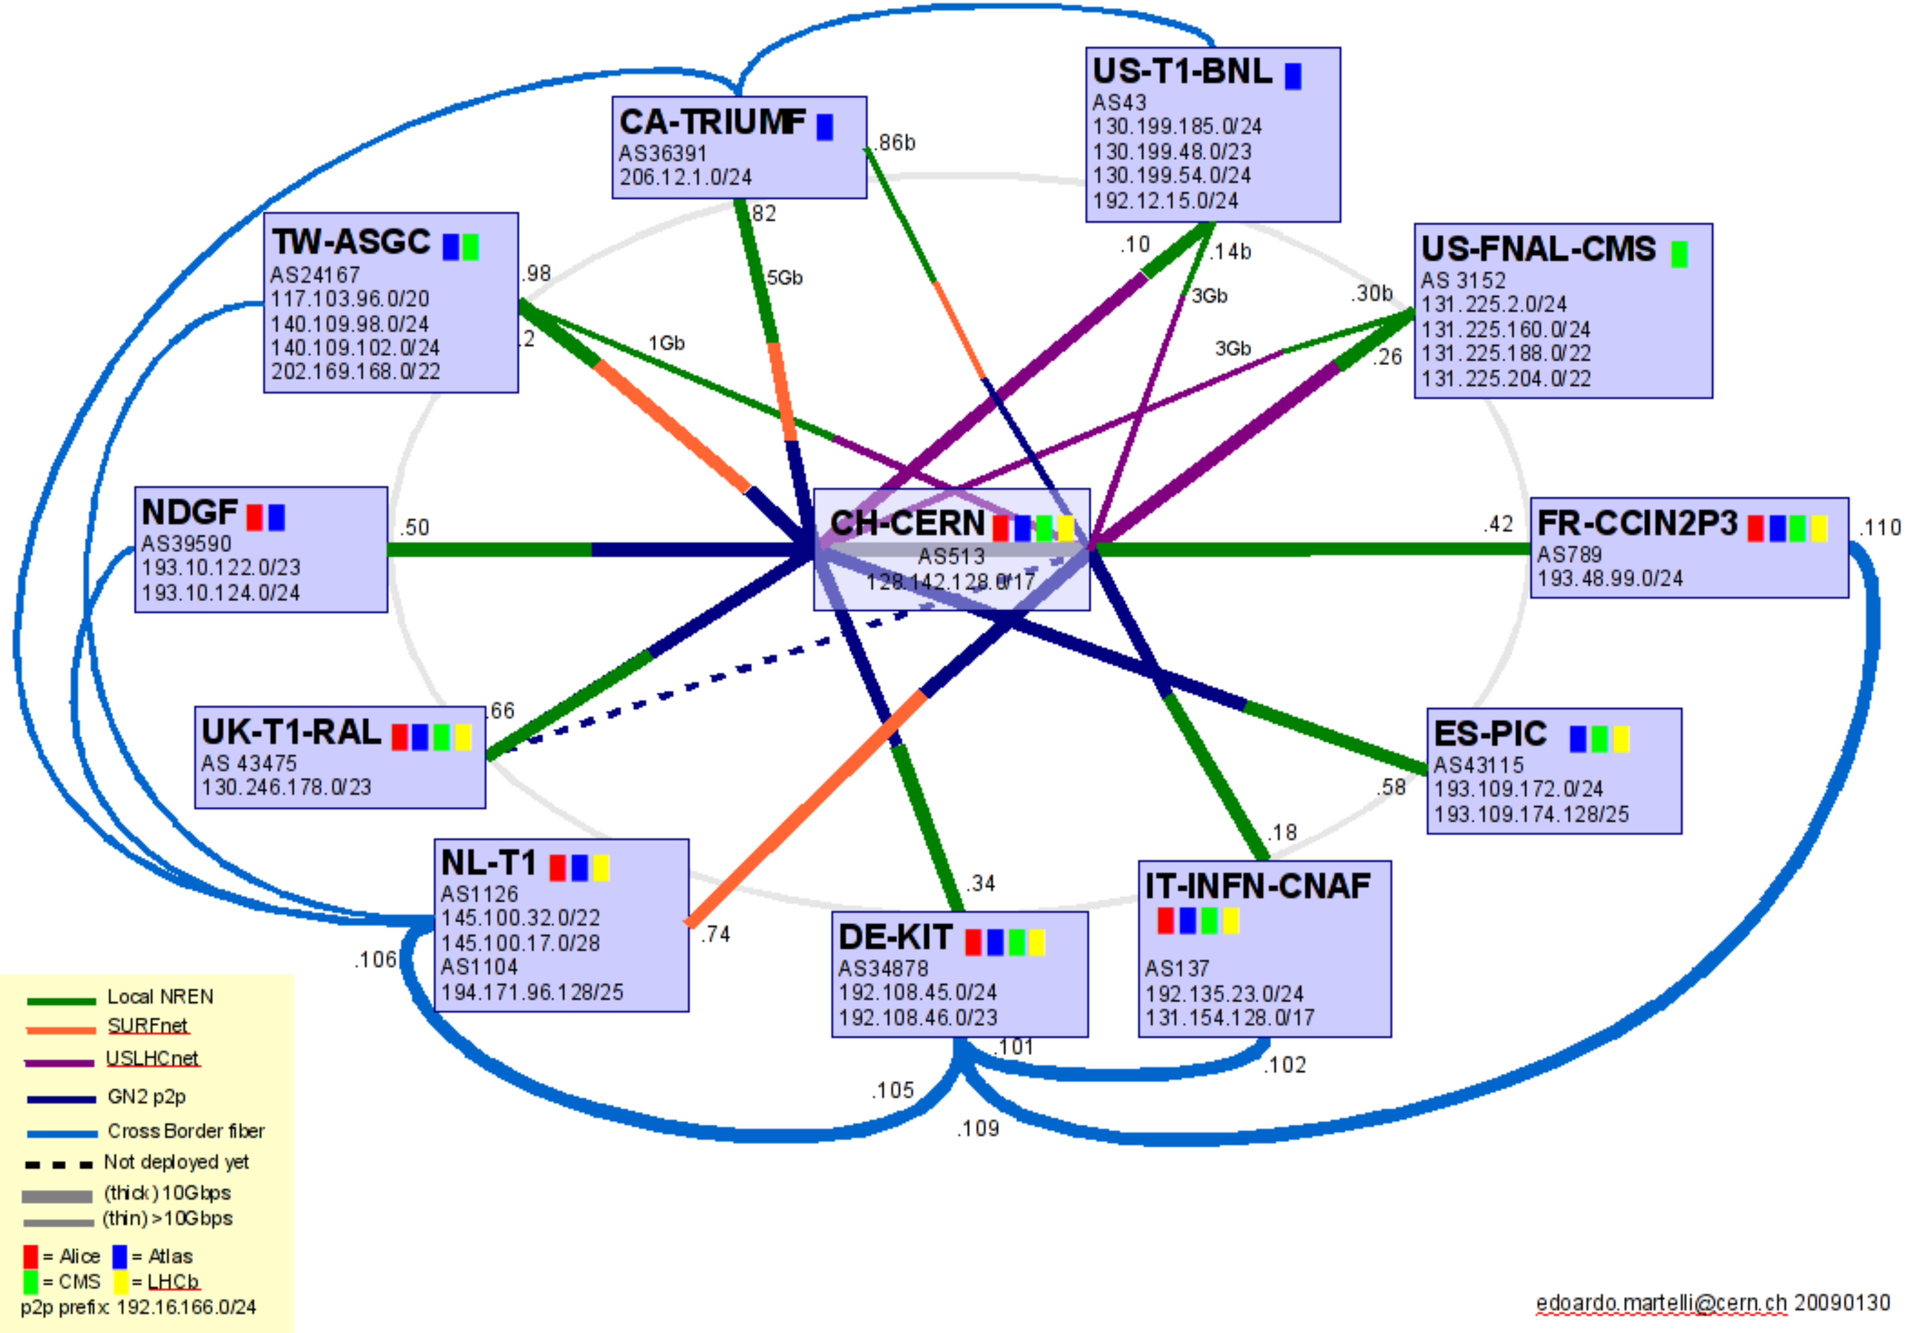

US-T1-BNL 🟦
AS43
130.199.185.0/24
130.199.48.0/23
130.199.54.0/24
192.12.15.0/24

CA-TRIUMF 🟦
AS36391
206.12.1.0/24

TW-ASGC 🟦🟩
AS24167
117.103.96.0/20
140.109.98.0/24
140.109.102.0/24
202.169.168.0/22

US-FNAL-CMS 🟩
AS 3152
131.225.2.0/24
131.225.160.0/24
131.225.188.0/22
131.225.204.0/22

NDGF 🟥🟦
AS39590
193.10.122.0/23
193.10.124.0/24

CH-CERN 🟥🟦🟩🟨
AS513
128.142.128.0/17

FR-CCIN2P3 🟥🟦🟩🟨
AS789
193.48.99.0/24

UK-T1-RAL 🟥🟦🟩🟨
AS 43475
130.246.178.0/23

ES-PIC 🟦🟩🟨
AS43115
193.109.172.0/24
193.109.174.128/25

NL-T1 🟥🟦🟨
AS1126
145.100.32.0/22
145.100.17.0/28
AS1104
194.171.96.128/25

DE-KIT 🟥🟦🟩🟨
AS34878
192.108.45.0/24
192.108.46.0/23

IT-INFN-CNAF
🟥🟦🟩🟨
AS137
192.135.23.0/24
131.154.128.0/17

Local NREN
SURFnet
USLHCnet
GN2 p2p
Cross Border fiber
Not deployed yet
(thick) 10Gbps
(thin) >10Gbps

🟥 = Alice 🟦 = Atlas
🟩 = CMS 🟨 = LHCb
p2p prefix: 192.16.166.0/24

edoardo.martelli@cern.ch 20090130

Firewall

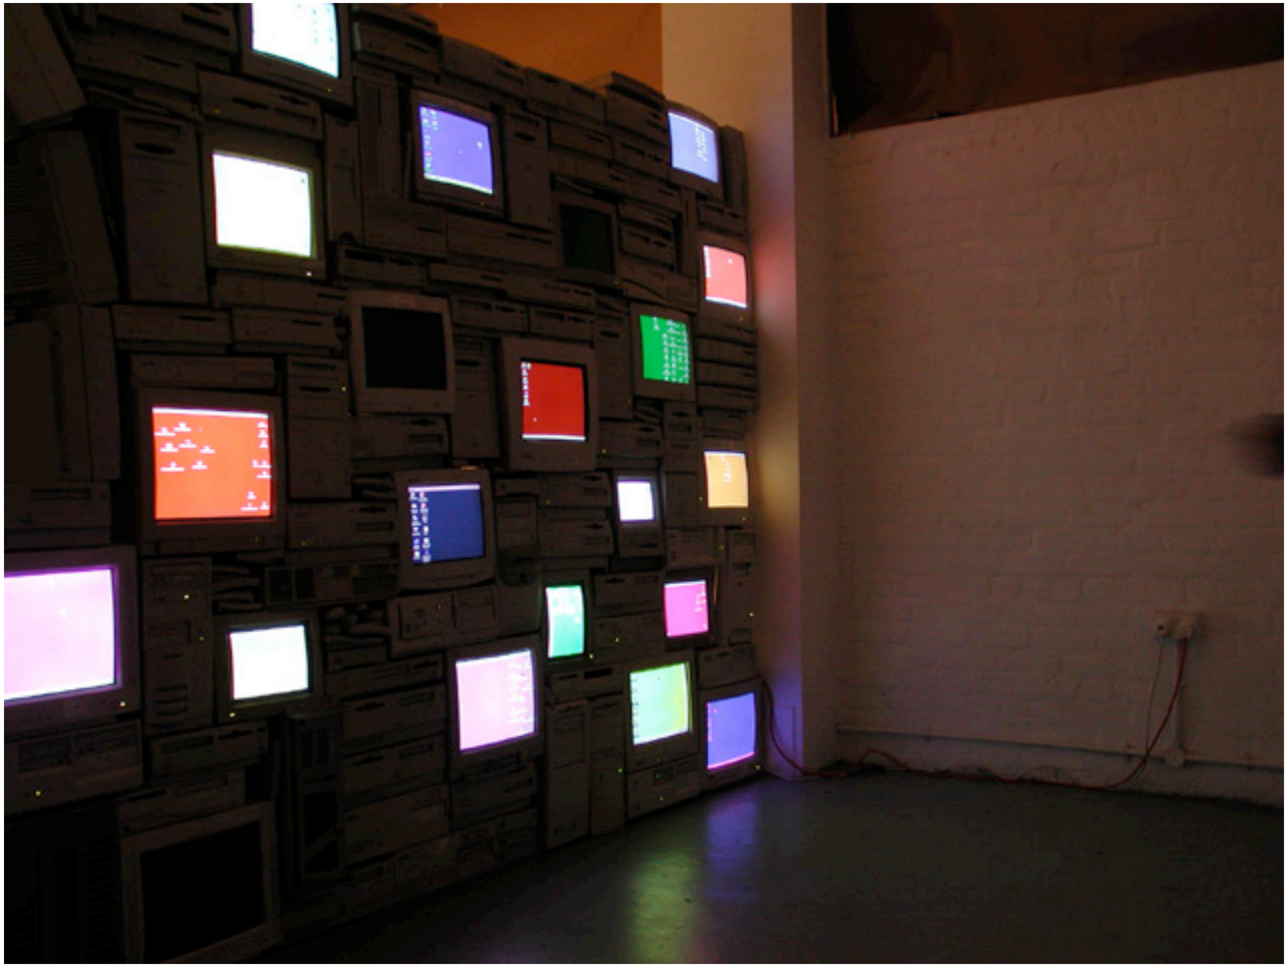

"*Firewall*" by Sandy Smith,
www.computersforart.org

Streams and Firewalls

- Data transfer target:
  300 MByte/s out of CERN to **each** of the ~10 T1s
  - 24 GBit/s aggregate bandwidth
  - you cannot traverse firewalls at that speed
  - For those of you who still believe in firewalls

- OPN – an Optical Private Network for the LHC
  - internal routing only (BGP)
  - all participants sign up to a common policy
  - exclusively for data transfers
  - no direct connections to 'The Internet'

"*Firewall*" by Sandy Smith,
www.computersforart.org

BiG Grid
the dutch e-science grid

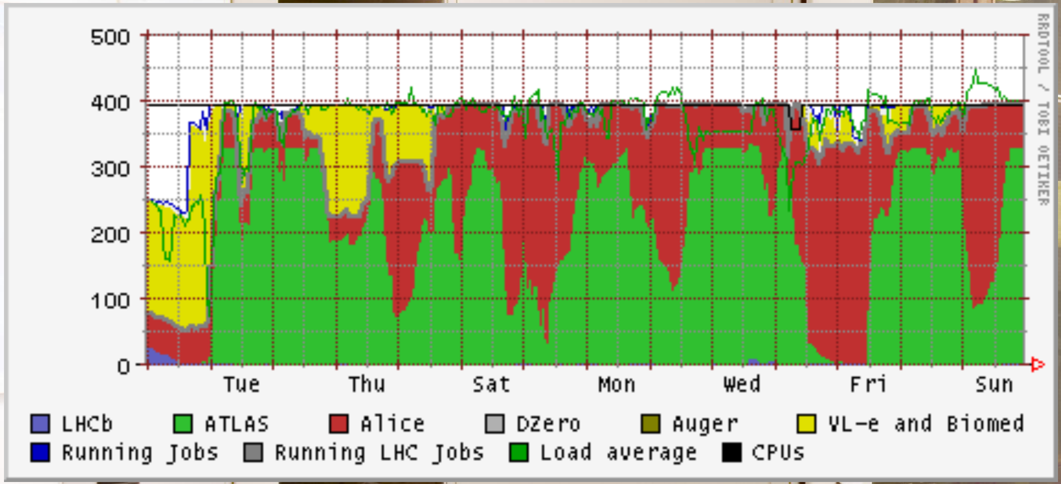

Nikhef (NDPF)

2550	processor cores
1 200	TByte disk
3x10	Gbps networks

SARA (GINA+LISA)

~2900	processor cores
850	TByte disk
1 500	TByte tape
4x 10	Gbps networks

RUG-CIT (Grid)

> 200	processor cores
34	TByte disk
10	Gbps networks

Philips Research Ehv

416	processor cores
126	TByte disk
1	Gbps networks

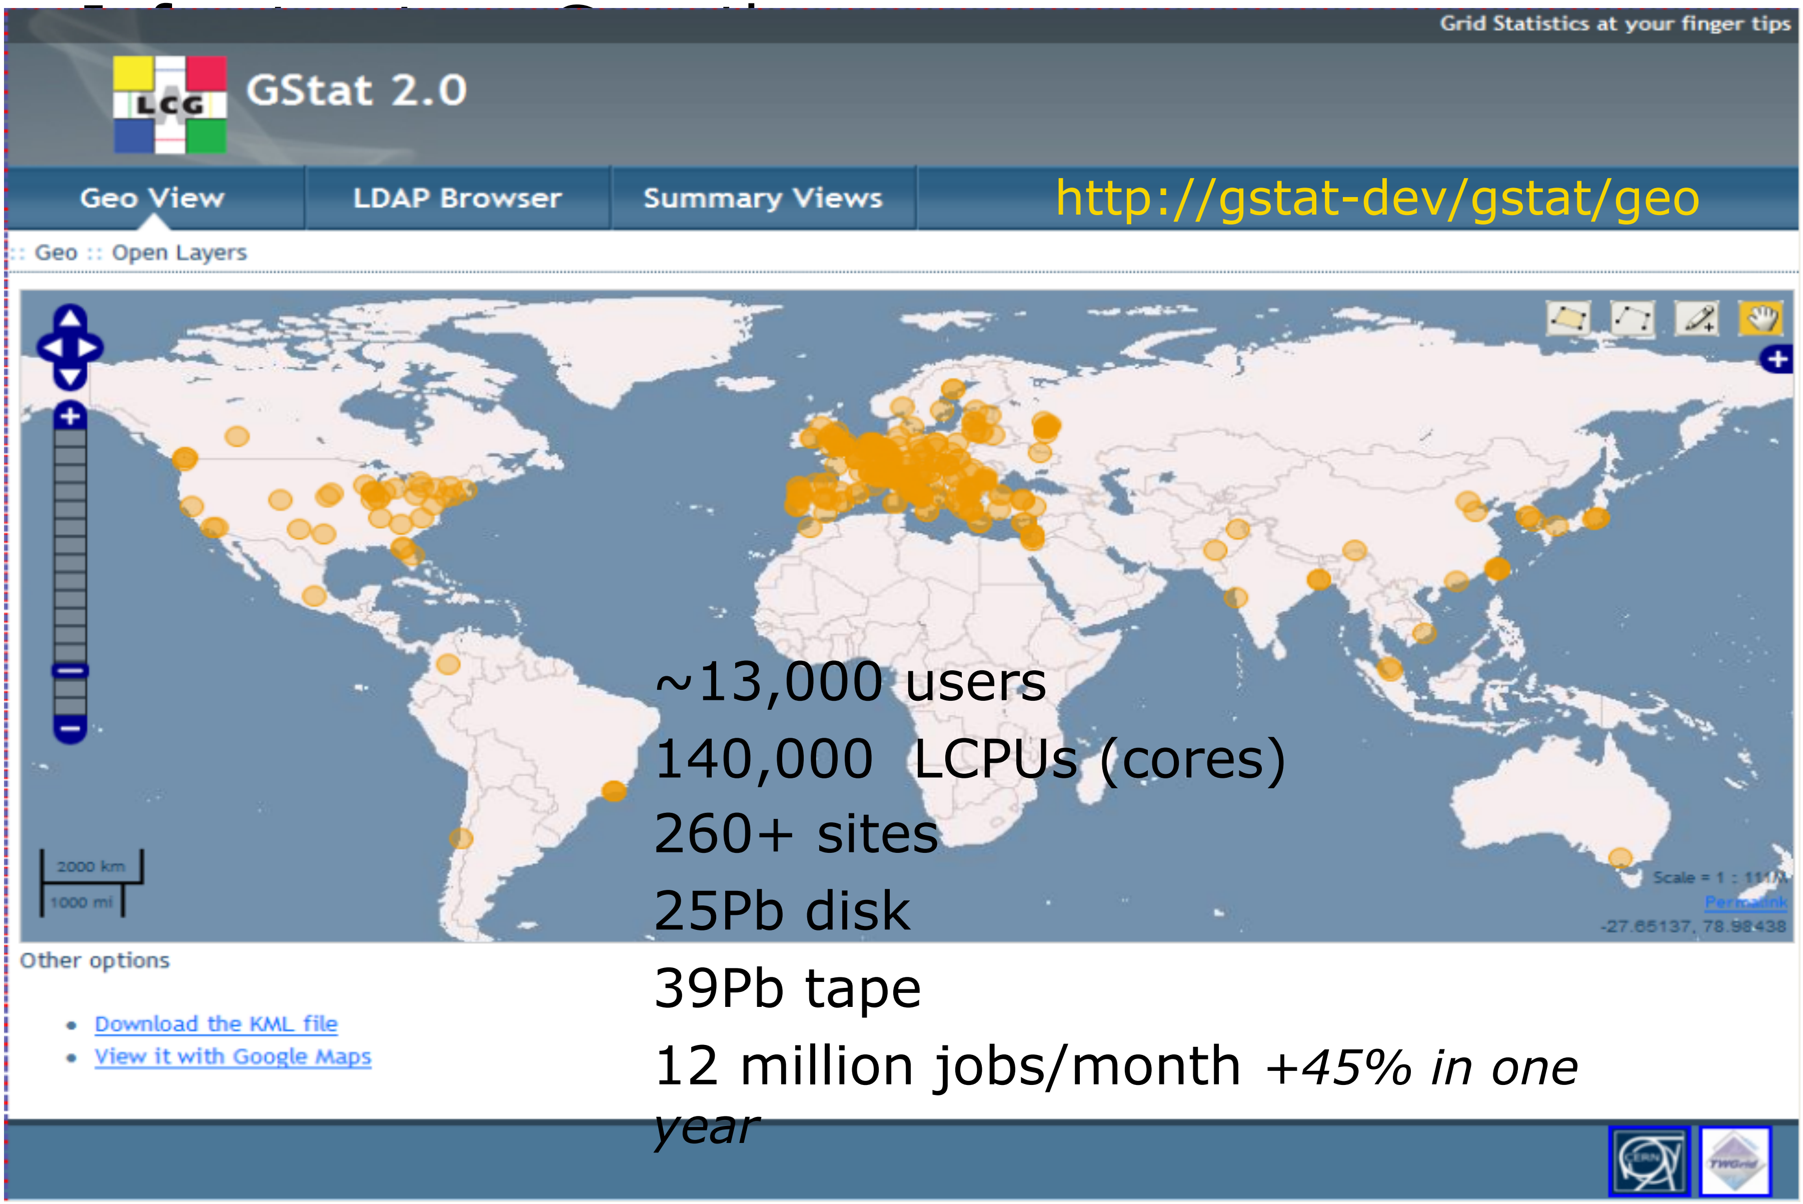

Grid Statistics at your finger tips

GStat 2.0

http://gstat-dev/gstat/geo

:: Geo :: Open Layers

Other options

- Download the KML file
- View it with Google Maps

~13,000 users

140,000 LCPUs (cores)

260+ sites

25Pb disk

39Pb tape

12 million jobs/month *+45% in one year*

The physical upgrade of the Nikhef DC from…

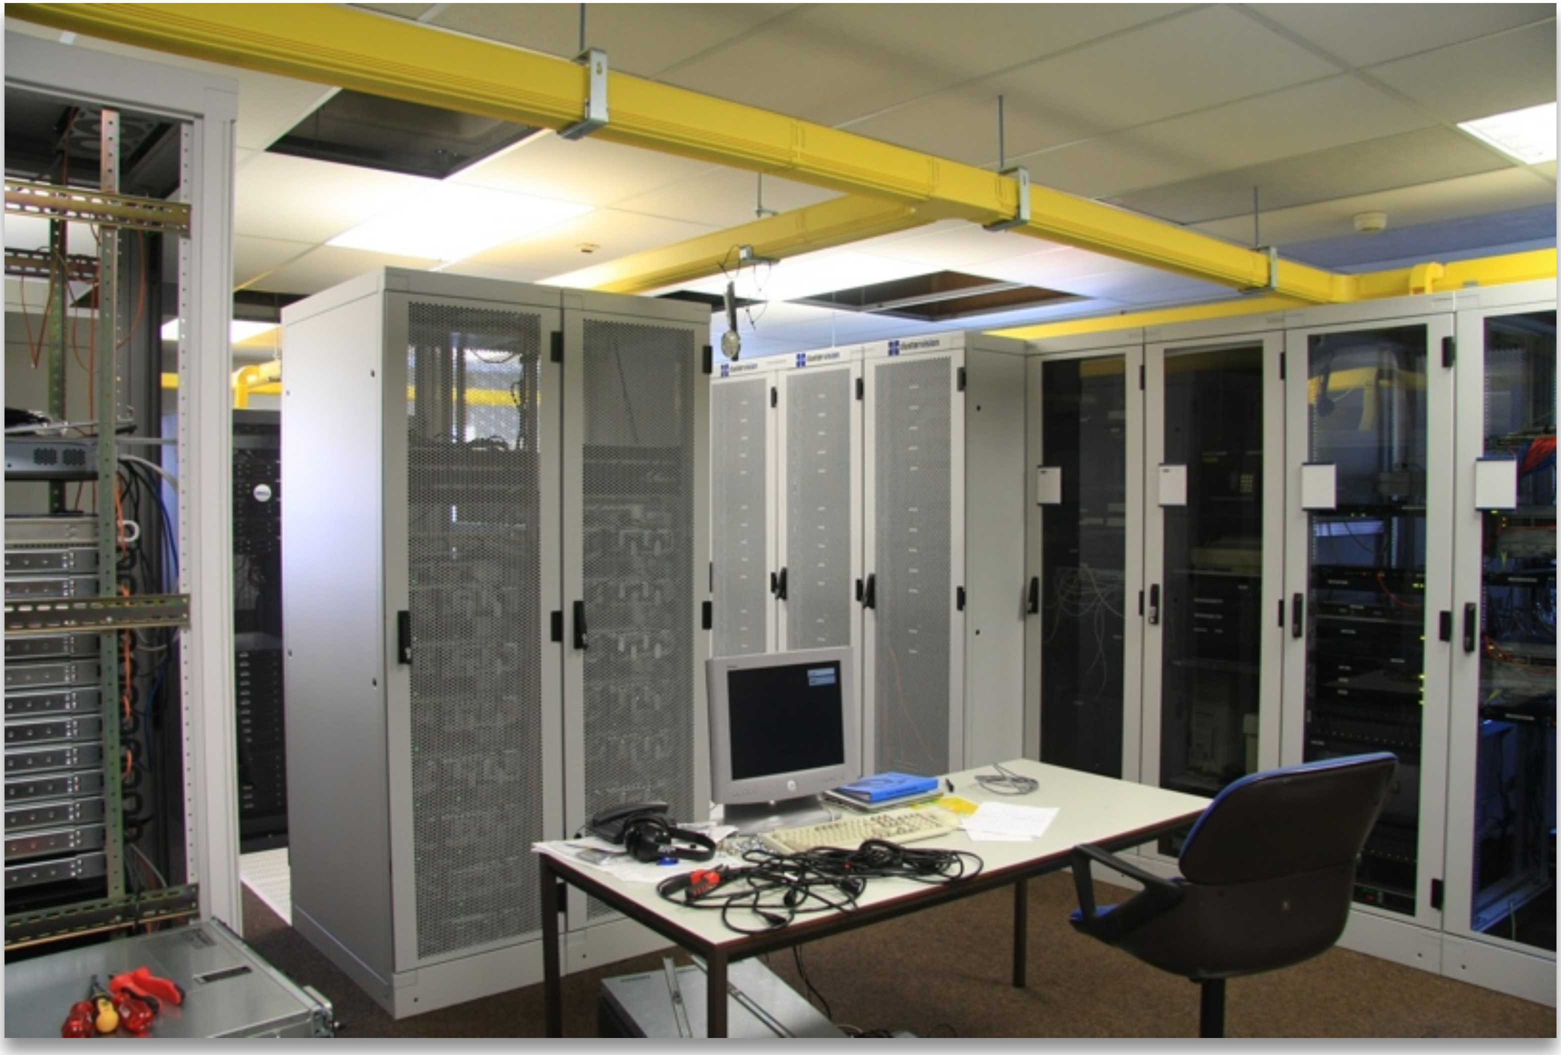

Too...

Internal network (until mid 2009)

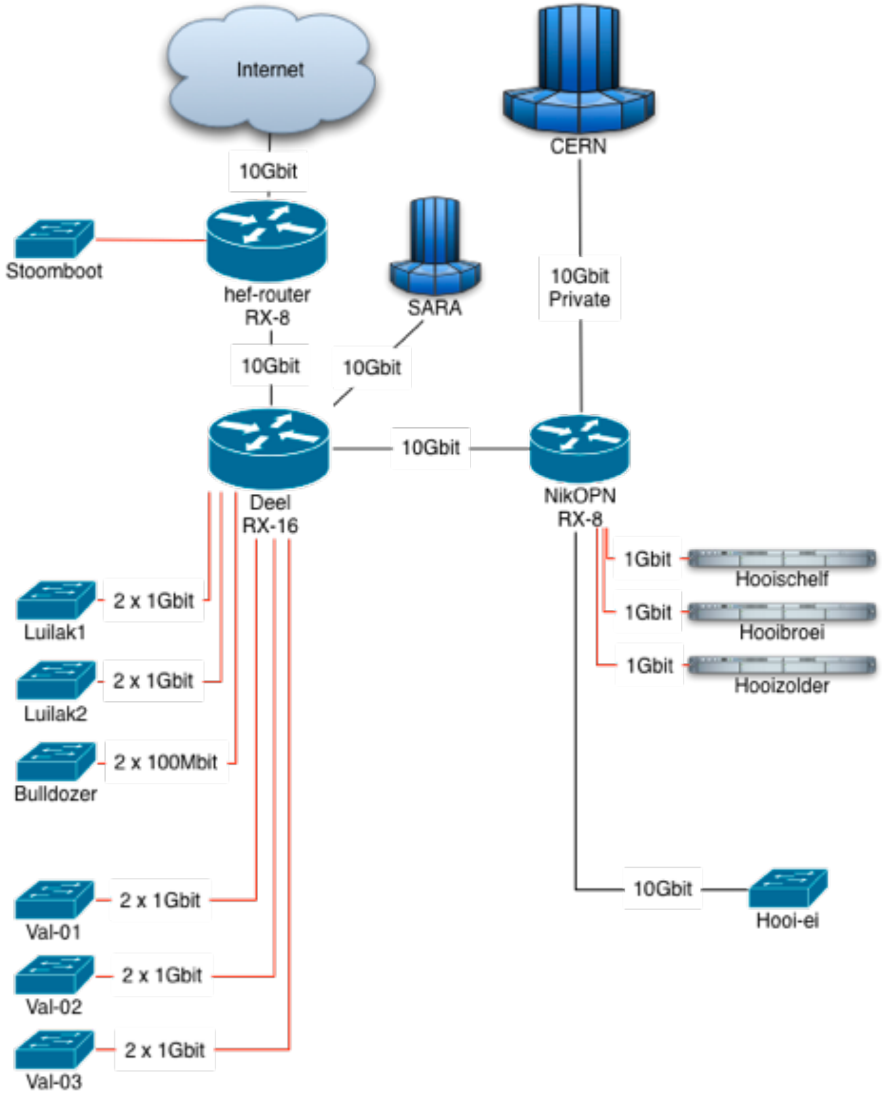

Internal network (late 2009)

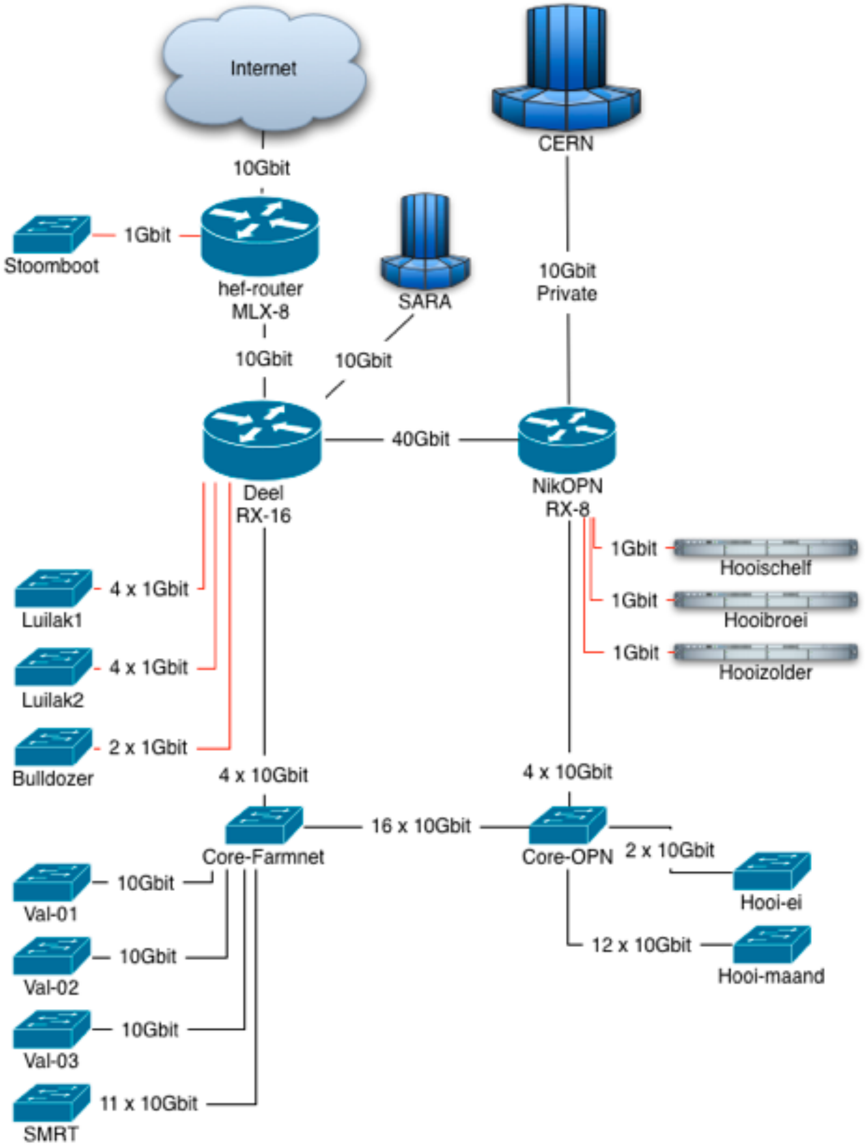

Think BIG

Examples: CERN Computer Centre

- not only systems management
- but also asset mngt and facilities
- *and you are not even allowed to look inside Google's data centers!*

And Why Do We Need It?

Enhanced Science needs more and more computations and

Collected data in science and industry grows exponentially

The Bible	**5 MByte**
Your own digital photographs	5 MByte/image
Bio-informatics databases	500 GByte each
Refereed journal papers	1 TByte/yr
Satellite world imagery	5 TByte/yr
Large Synoptic Survey Telescope	30 Tbyte/day
Internet Archive 1996-2002	100 Tbyte
Web downloads for Google indexing	4 PByte/yr
Large Hadron Collider physics	**20 PByte/yr**
Astronomy tomorrow: SKA	365 PByte/yr

1 Petabyte = 1 000 000 000 Megabyte

Why Grid computing – today?

- New applications need larger amounts of data or computation
- Larger, and growing, distributed user community
- Network grows faster than compute power/storage

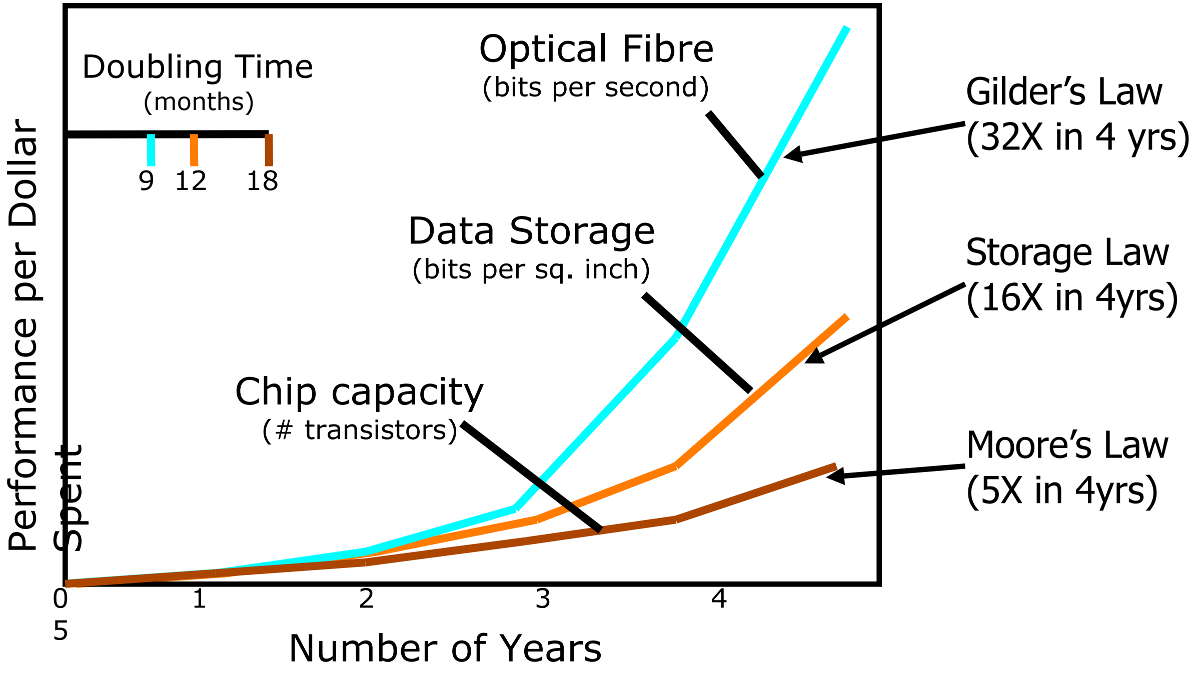

Graphic: "The Triumph of Light", Scientific American, January 2001

Typical grid topology for computational jobs

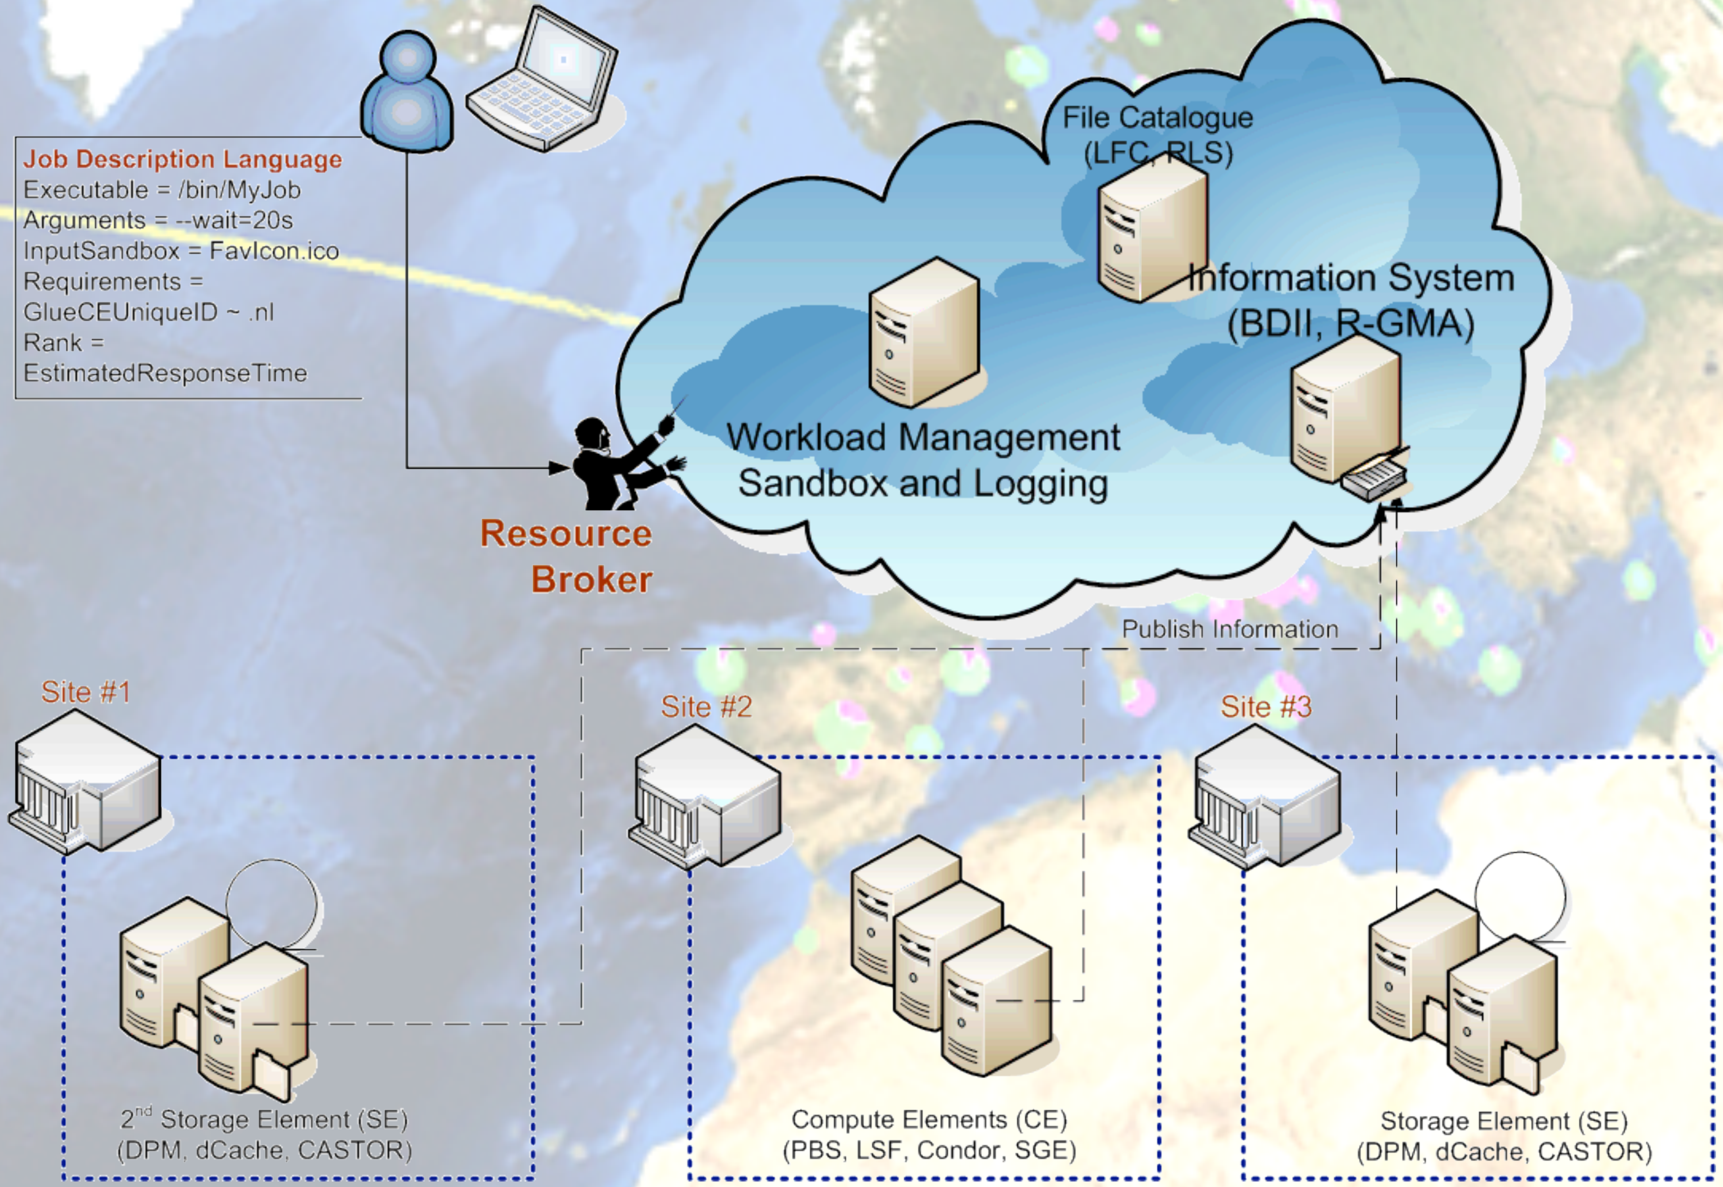

Job Description Language

```
Executable = /bin/MyJob
Arguments = --wait=20s
InputSandbox = FavIcon.ico
Requirements =
GlueCEUniqueID ~ .nl
Rank =
EstimatedResponseTime
```

Resource Broker
File Catalogue (LFC, RLS)
Information System (BDII, R-GMA)
Workload Management Sandbox and Logging
Publish Information
Site #1
Site #2
Site #3
2nd Storage Element (SE) (DPM, dCache, CASTOR)
Compute Elements (CE) (PBS, LSF, Condor, SGE)
Storage Element (SE) (DPM, dCache, CASTOR)

Virtual Organisations

The communities that make up the grid:

- **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**.

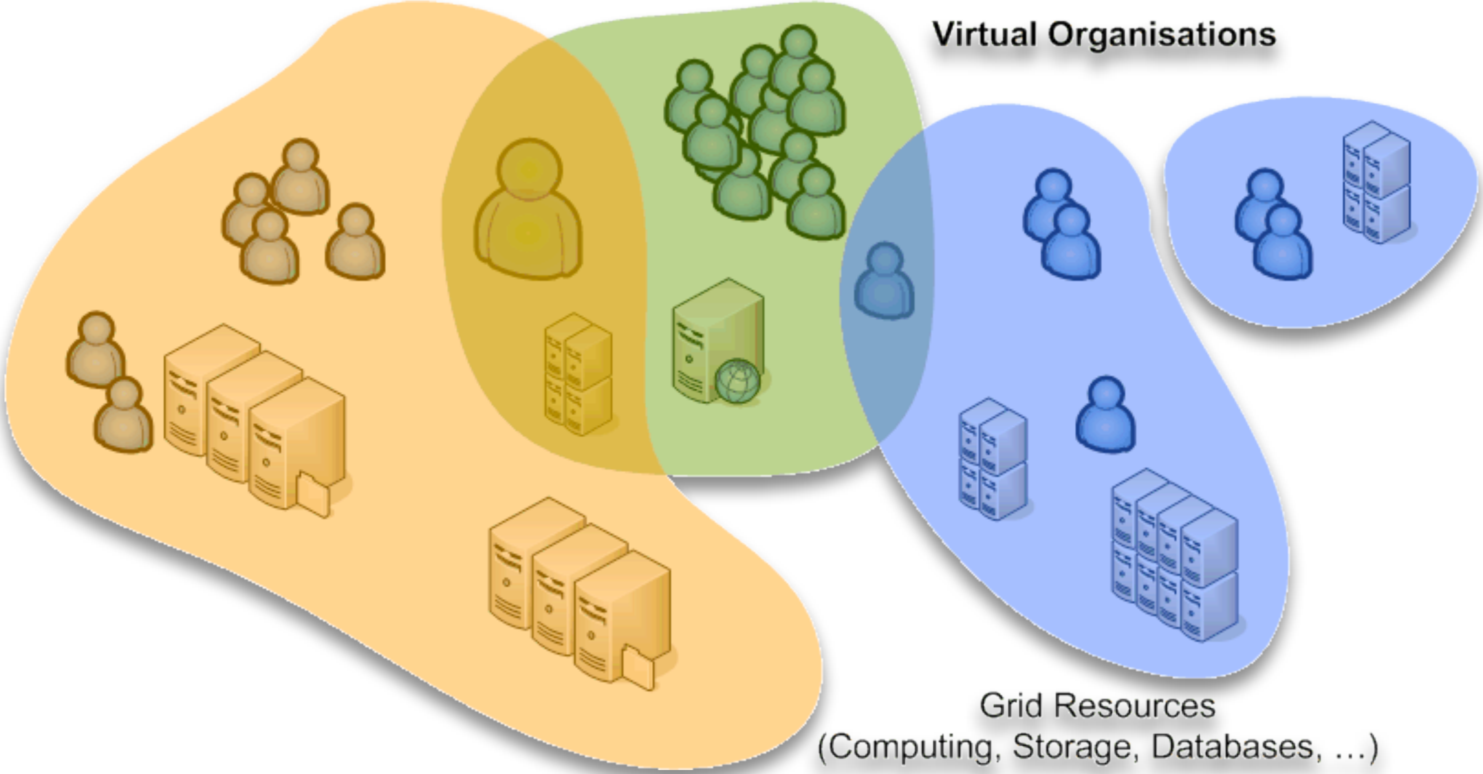

VOs and the infrastructure

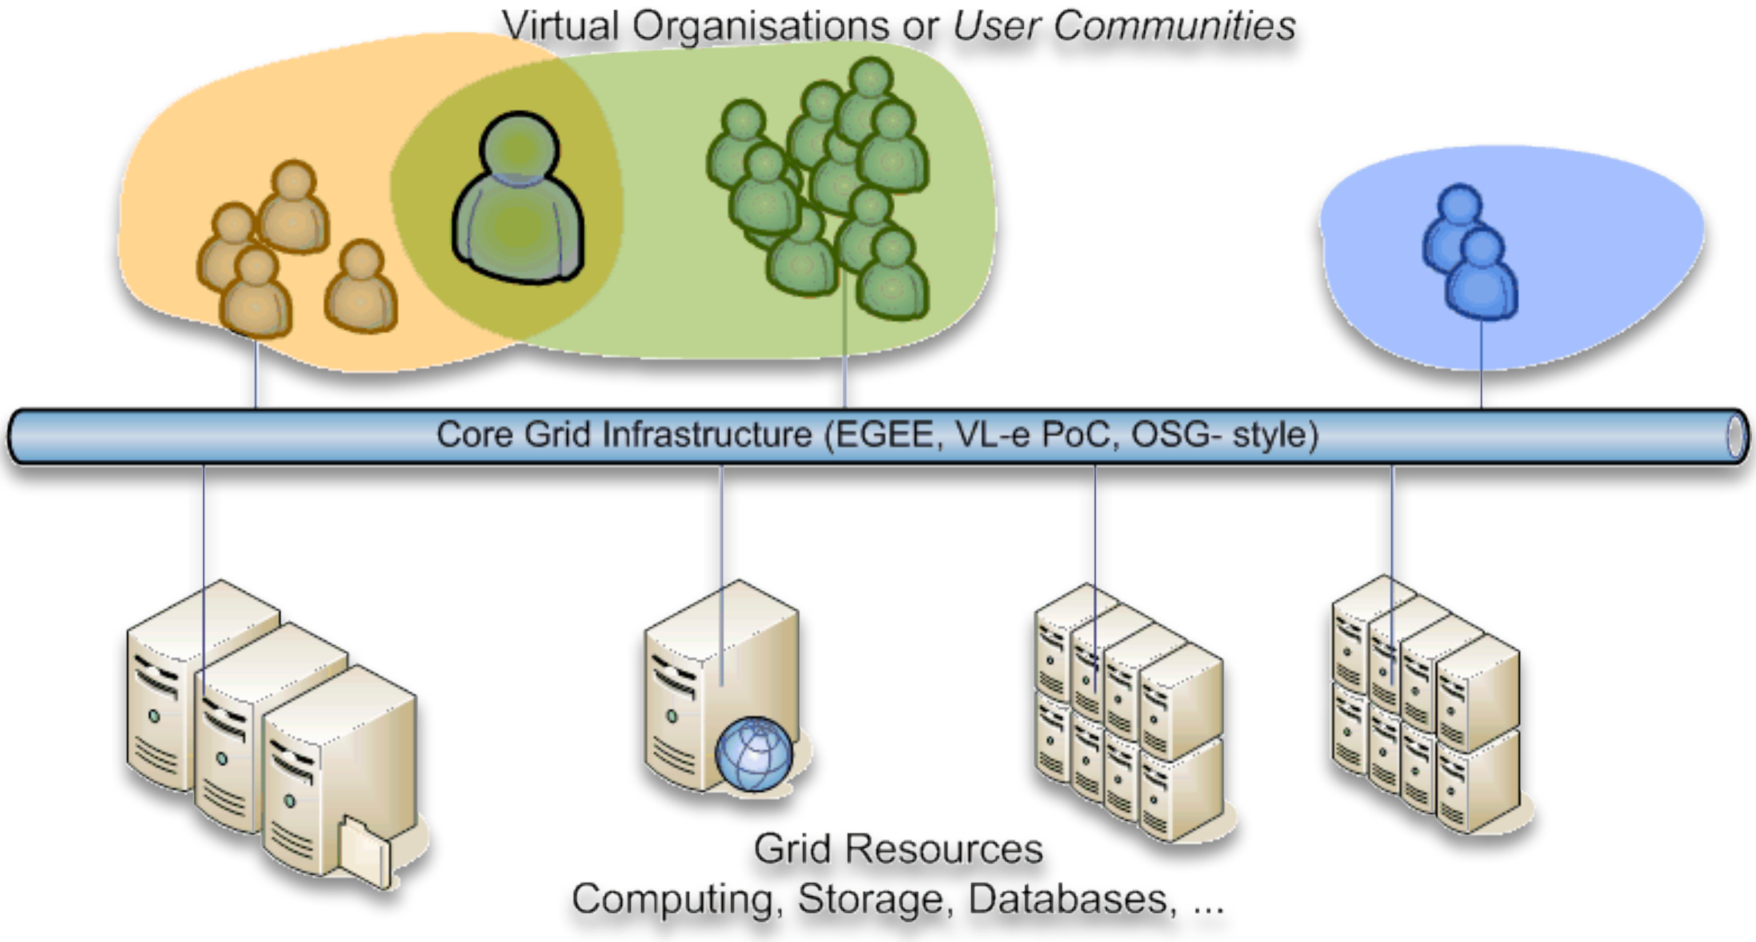
Virtual Organisations or *User Communities*
Core Grid Infrastructure (EGEE, VL-e PoC, OSG- style)
Grid Resources
Computing, Storage, Databases, ...

VO federation needs

- Trust establishment within the VO is separated in:
  - user identity (the user's *passport*)
  - group and roles within the VO (*visa)*
    as these are different from a persons organizational role

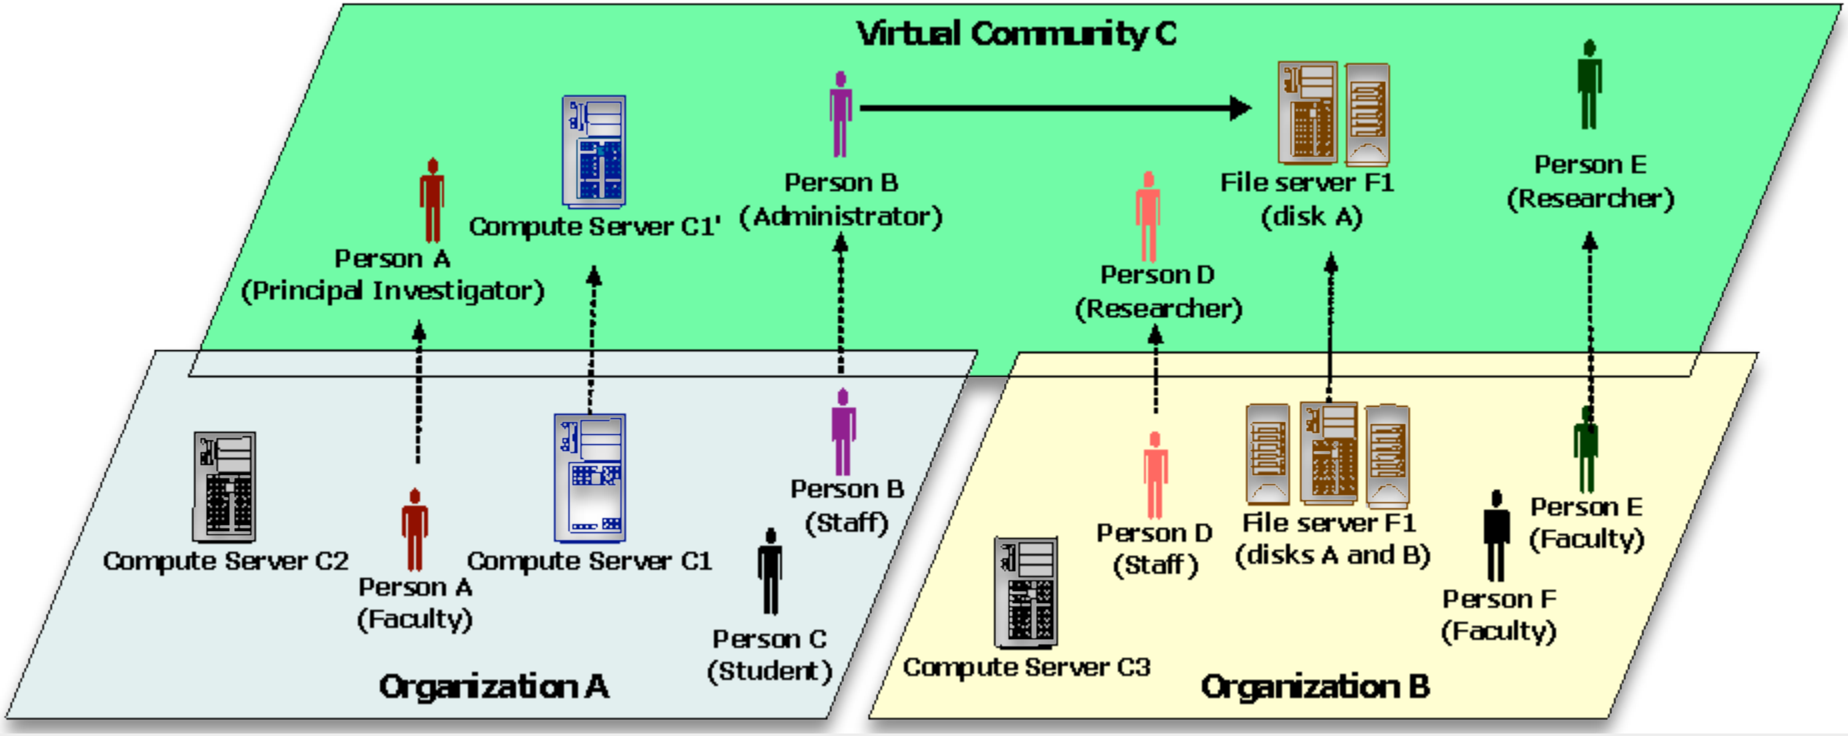

graphic: OGSA Architecture 1.0, OGF GFD-I.030

Trust relationships

- For the VO model to work, parties need to (minimally) trust each other in their VO interactions
  - the alternative would be that every user would have to register at and **every** resource provider…

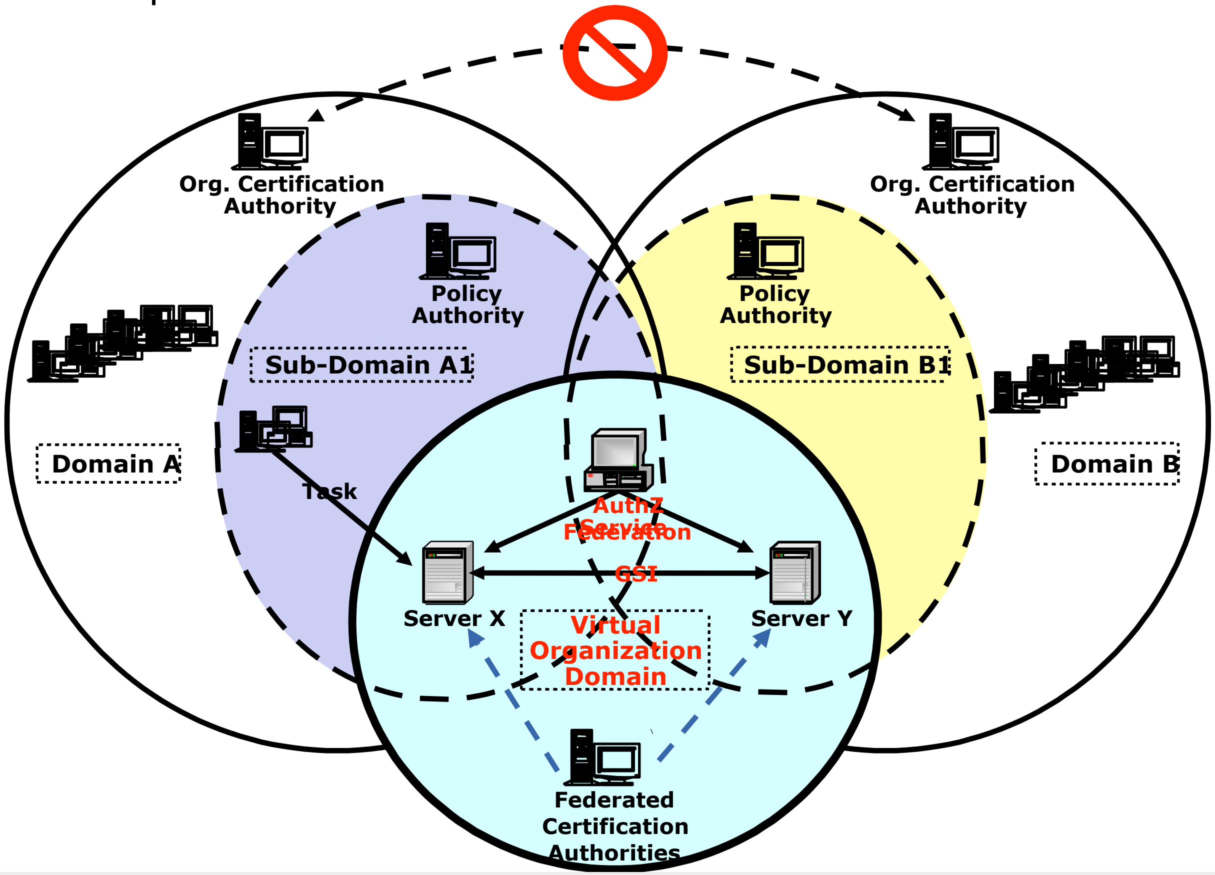

graphic from: Frank Siebenlist, Argonne Natl. Lab, *Globus Alliance*

Authentication models

> Direct user-to-site
  > passwords, enterprise PKI, Kerbe

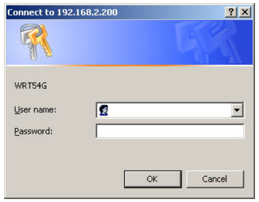

Connect to 192.168.2.200

WRT54G

User name:

Password:

OK Cancel

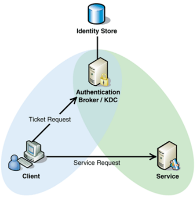

> PKI with trusted third parties

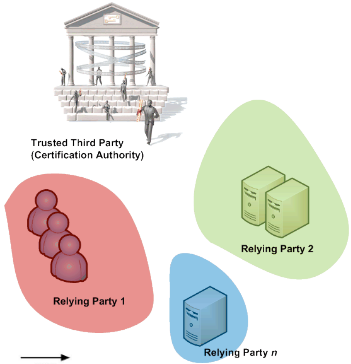

> Federated access
  > Controlled & policy based
  > Free-for-all, e.g., OpenID

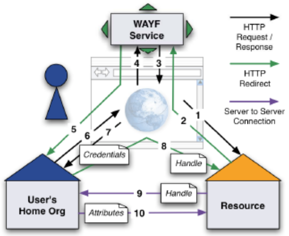

> Identity meta-system
  > Infocard type systems

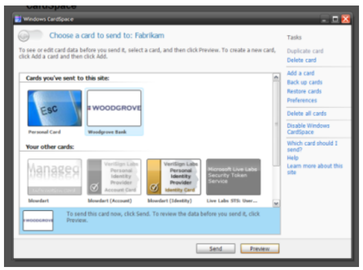

Windows CardSpace

Choose a card to send to: Fabrikam

To see or edit card data before you send it, select a card, and then click Preview. To create a new card, click Add a card and then click Add.

Cards you've sent to this site:

Personal Card

Woodgrove Bank

Your other cards:

blowdart

blowdart (Account)

blowdart (Identity)

Live Labs STS: User...

To send this card now, click Send. To review the data before you send it, click Preview.

Send Preview

Authentication models

> Direct user-to-site

  > passwords, enterprise PKI, Kerbe

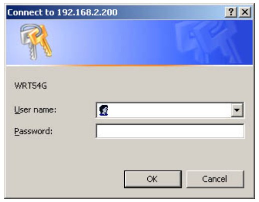

Connect to 192.168.2.200

WRT54G

User name:

Password:

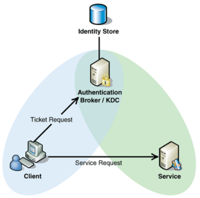

> PKI with trusted third parties

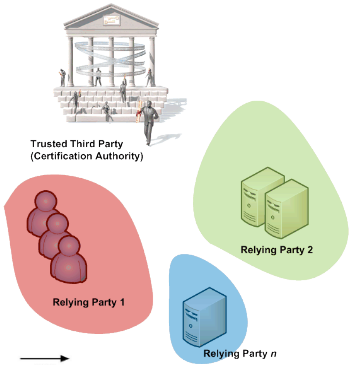

> Federated access

  > Controlled & policy based

  > Free-for-all, e.g., OpenID

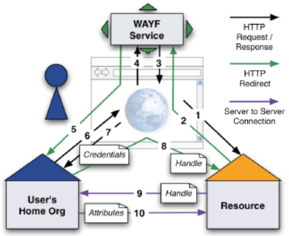

> Identity meta-system

  > Infocard type systems

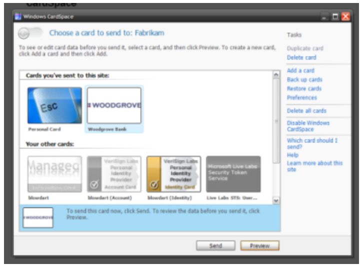

Windows CardSpace

Choose a card to send to: Fabrikam

To see or edit card data before you send it, select a card, and then click Preview. To create a new card, click Add a card and then click Add.

Cards you've sent to this site:

Personal Card

Woodgrove Bank

Your other cards:

blowdart

blowdart (Account)

blowdart (Identity)

Live Labs STS: User...

To send this card now, click Send. To review the data before you send it, click Preview.

Esc
WOODGROVE
Managed
VeriSign Labs Personal Identity Provider Account Card
VeriSign Labs Personal Identity Provider Identity Card
Microsoft Live Labs Security Token Service
Tasks
Duplicate card
Delete card
Add a card
Back up cards
Restore cards
Preferences
Delete all cards
Disable Windows CardSpace
Which card should I send?
Help
Learn more about this site
Send
Preview

User Identity

- Users and resources are typically part of more than one VO, but don't want many passwords

- Users and resource get a *single authentication token* (identity certificate)
  - that works across virtual organizations
  - issued by a party trusted by all ("CA"),
  - recognized by many resource providers, users, and VOs
  - satisfy traceability and persistency requirement
  - in itself does not grant any access, but provides a unique binding between an identifier and the subject

This is called your *(identity) certificate*

It is a *cryptographically protected statement* by the CA

- that you can use to prove your identity
  in combination with a ***private key*** and its ***passphrase***

Trusting the signature

- Paul’s digital signature is safe if:
  1. Paul’s private key is not compromised
  2. John knows Paul’s public key
- How can John be sure that Paul’s public key is really
  Paul’s public key and not someone else’s?
  - A *third party* guarantees the correspondence between public key and owner’s identity.
  - Both A and B must trust this third party

slide from EGEE NA3 Tutorial repository

Contacting the CA

- Each *CA* has different policies and practices

- Generate a cryptographic key pair
  - using a script like grid-cert-request
  - with your web browser
  - using jGridstart (Java Grid Start)
- Appear in-person to the Registration Authority (*RA*)
  - with a valid personal ID-card
- *RA* approves your request
- *CA* signs the approved request and sends you the cert

  - via mail: copy to your home directory
  - via the web: download into your browser and export to disk
  - via jGridstart: next -> next -> finish
- All use a network of *RAs* close to you

Your certificate (RFC 3280 / RFC 5280)

```
VisionMaster:~ okoeroo$ openssl x509 -text -noout -in ~/.globus/usercert.pem
Certificate:
    Data:
        Serial Number: 2812 (0xafc)
        Signature Algorithm: sha1WithRSAEncryption
        Issuer: C=NL, O=NIKHEF, CN=NIKHEF medium-security certification auth
        Validity
            Not Before: Dec 10 00:00:00 2009 GMT
            Not After : Dec 10 14:32:49 2010 GMT
        Subject: O=dutchgrid, O=users, O=nikhef, CN=Oscar Koeroo
        X509v3 extensions:
            X509v3 Subject Alternative Name:
                email:okoeroo@nikhef.nl
    Signature Algorithm: sha1WithRSAEncryption
        75:ef:19:f7:41:43:78:6b:32: ...
-----BEGIN CERTIFICATE-----
MIIEhTCCA22gAwIS0sAK/qZIPIt0GA8iWQo ...
-----END CERTIFICATE-----
```

How do the sites know me (and I them)?

IGTF International Grid Trust Federation AP | EU | TAG

International Grid Trust Federation

- All research grid infrastructures share the same base set of trusted third parties ('CAs')
- There is typically one in each country
- The credentials they issue are comparable in quality

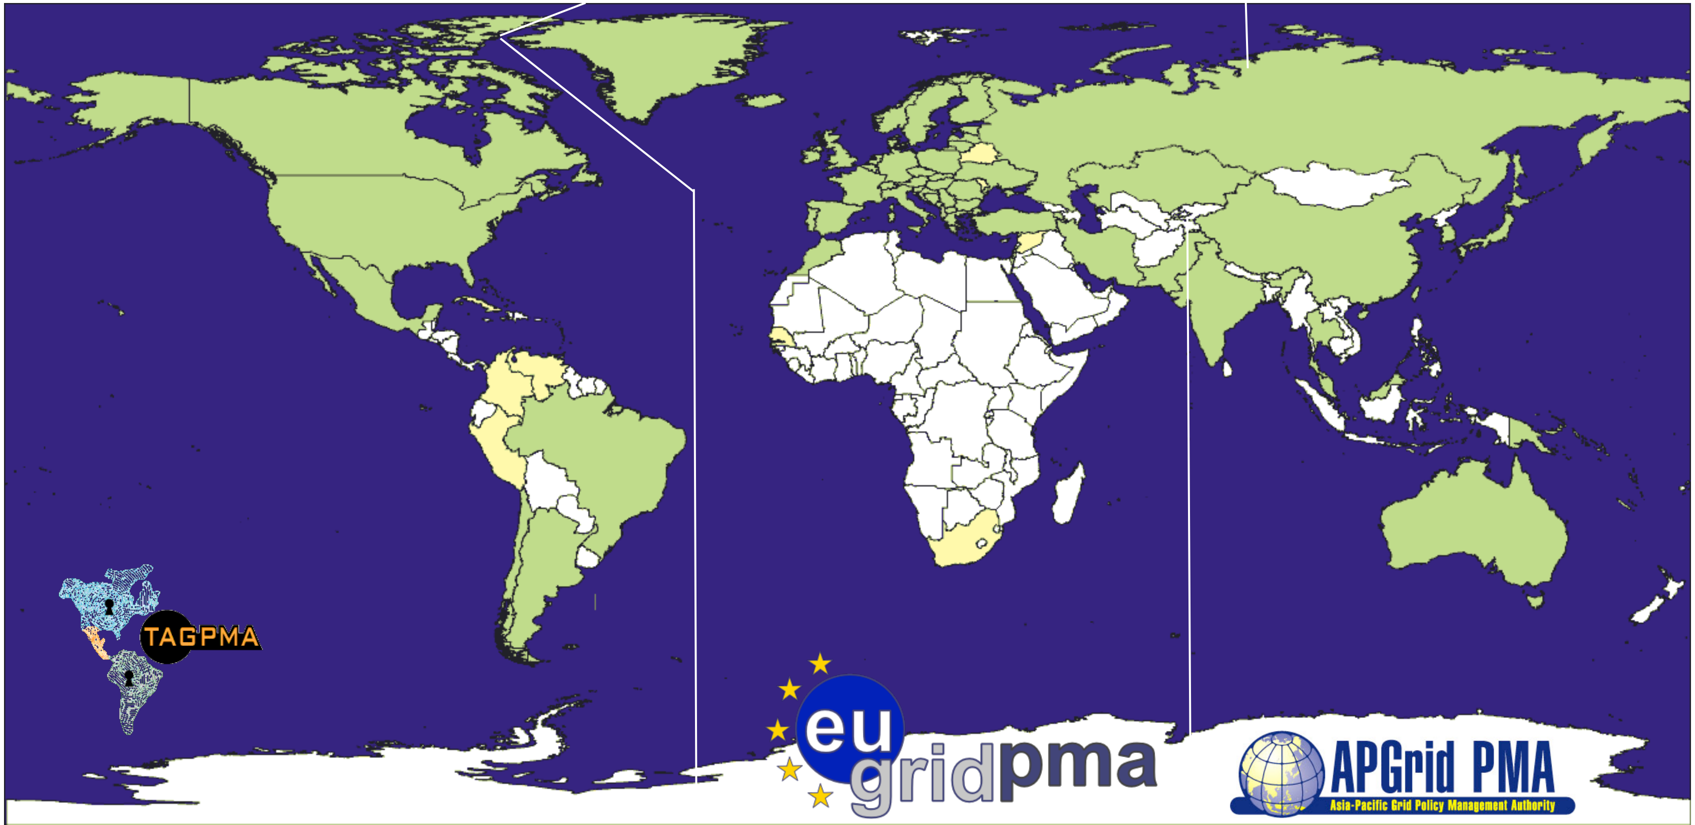

Requirements for (inter)national trust

- Identity vetting procedures
  - Based on (national) photo ID's
  - Face-to-face verification of applicants via a network of Registration Authorities
  - possible to trace the user in case of unlawful misconduct
  - Secure binding between the request and the identity vetting
  - Periodic renewal (once every year)
- Secure operation
  - off-line signing key or HSM-backed on-line secured systems
- Response to incidents
  - Timely revocation of compromised certificates

Single sign-on and delegation

- To authenticate with your certificate directly you would have to type a passphrase every time
- Also you need a way to send you *VOMS credentials* across

- In the Grid Security Infrastructure today, this is solved by *'proxy certificates'*
  - *a temporary key pair*
  - *in a temporary certificate signed by your 'long term' private key*
  - *valid for a limited time (default: 12 hours)*
  - *and itself not protected by a passphrase*

The grid-proxy-init tool

- User enters pass phrase, which is used to decrypt private key.
- Private key is used to sign a proxy certificate with its own, new public/private key pair.
  - User’s private key not exposed after proxy has been signed

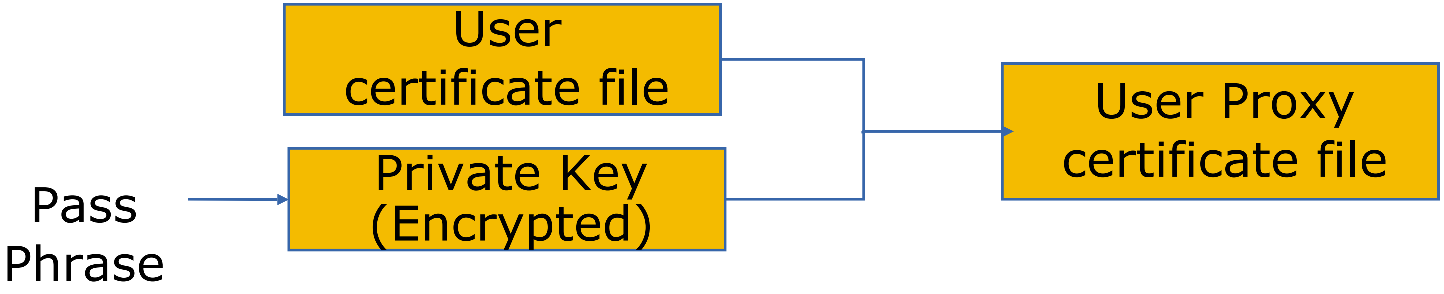

- Proxy placed in /tmp
  - the private key of the Proxy is *not* encrypted:
  - stored in local file: must be readable **only** by the owner;
  - proxy lifetime is short (typically 12 h) to minimize security risks.
- NOTE: *No* network traffic!

Daisy-chaining proxy delegation

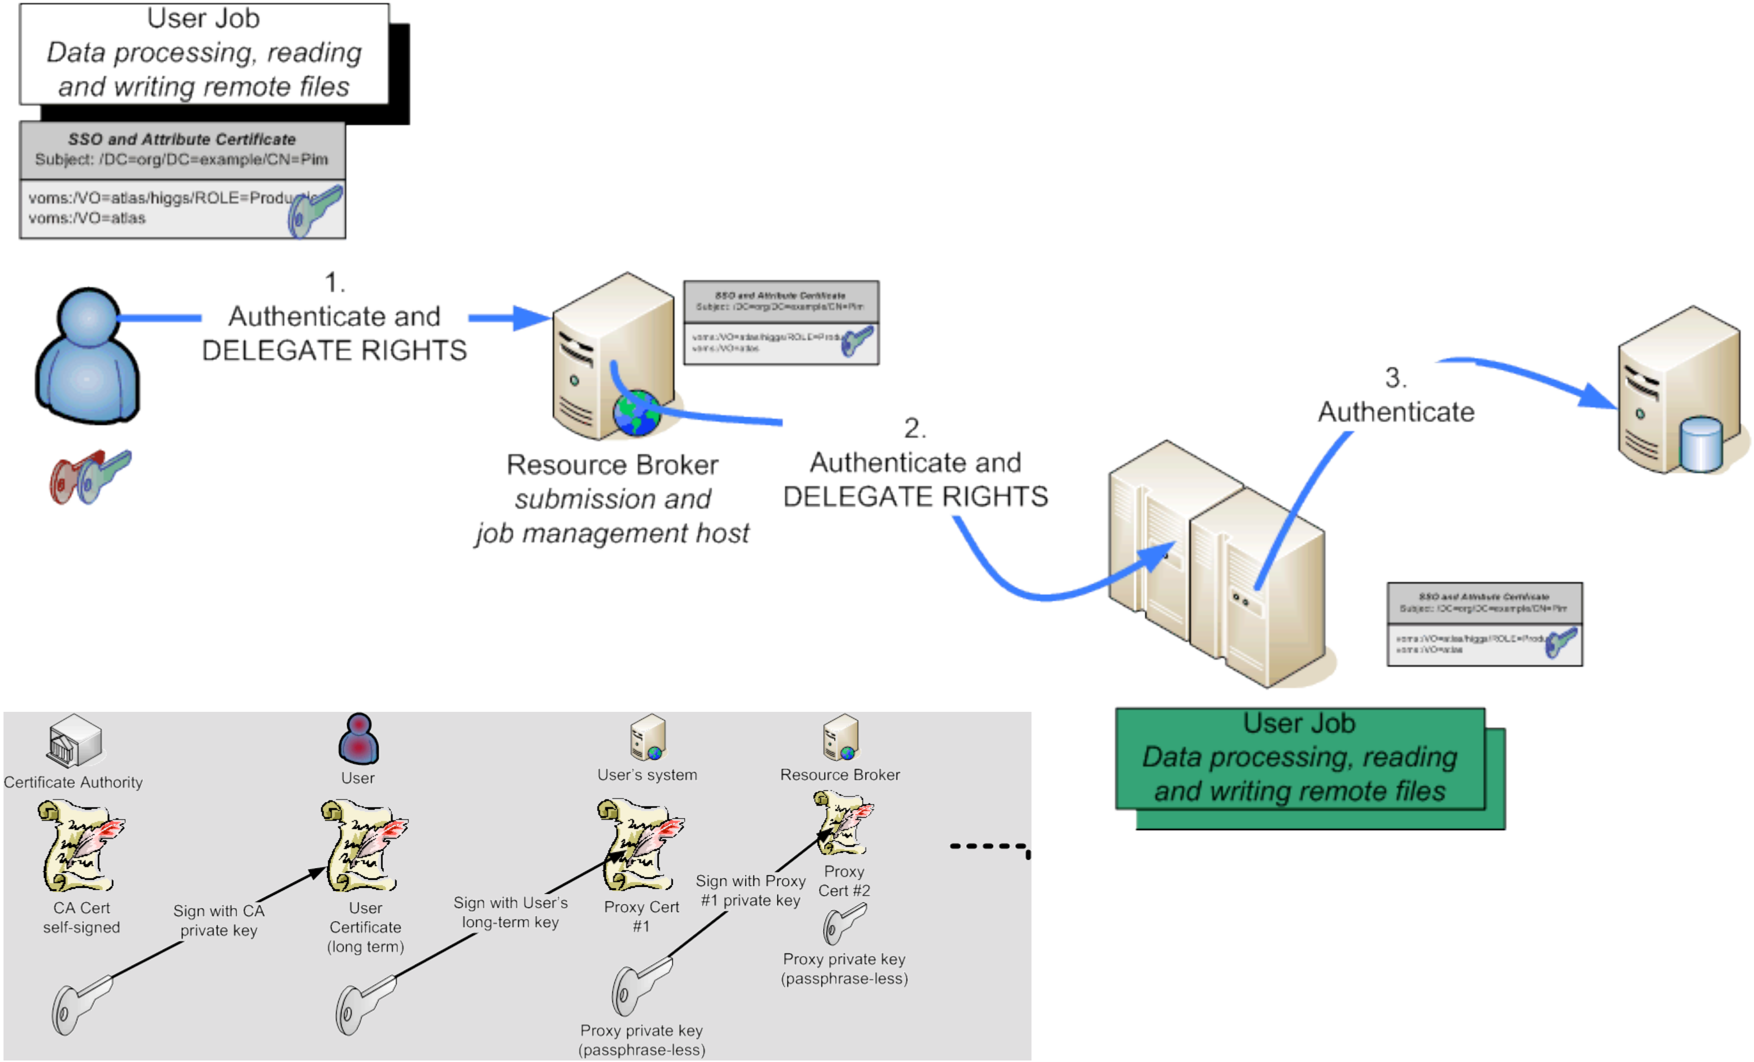

User Job
Data processing, reading and writing remote files

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

1. Authenticate and DELEGATE RIGHTS

Resource Broker
submission and job management host

2. Authenticate and DELEGATE RIGHTS

3. Authenticate

User Job
Data processing, reading and writing remote files

Certificate Authority
User
User's system
Resource Broker
CA Cert self-signed
Sign with CA private key
User Certificate (long term)
Sign with User's long-term key
Proxy Cert #1
Sign with Proxy #1 private key
Proxy Cert #2
Proxy private key (passphrase-less)
Proxy private key (passphrase-less)

Daisy-chaining proxy delegation

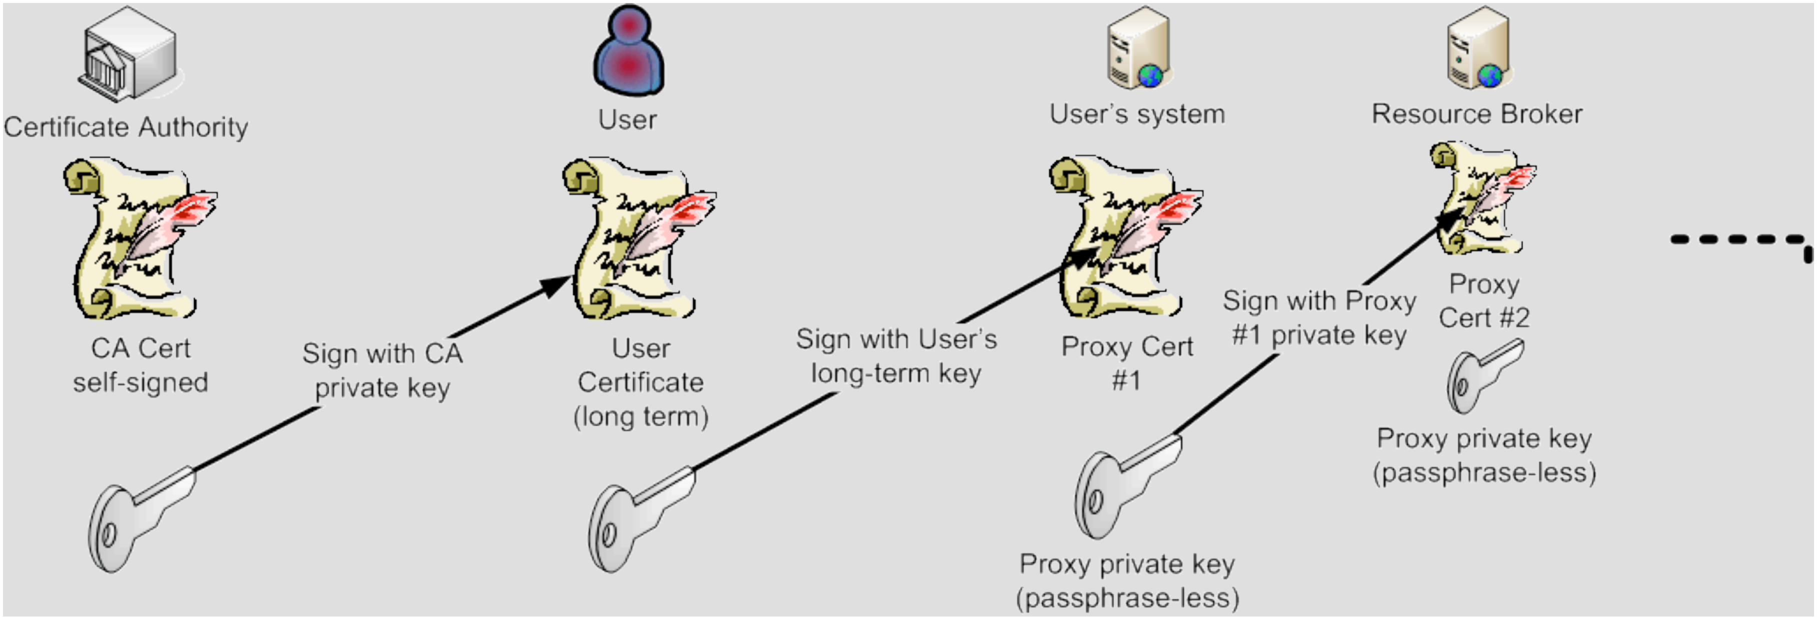
Certificate Authority
User
User's system
Resource Broker
CA Cert self-signed
Sign with CA private key
User Certificate (long term)
Sign with User's long-term key
Proxy Cert #1
Sign with Proxy #1 private key
Proxy Cert #2
Proxy private key (passphrase-less)
Proxy private key (passphrase-less)

Registering with your VO

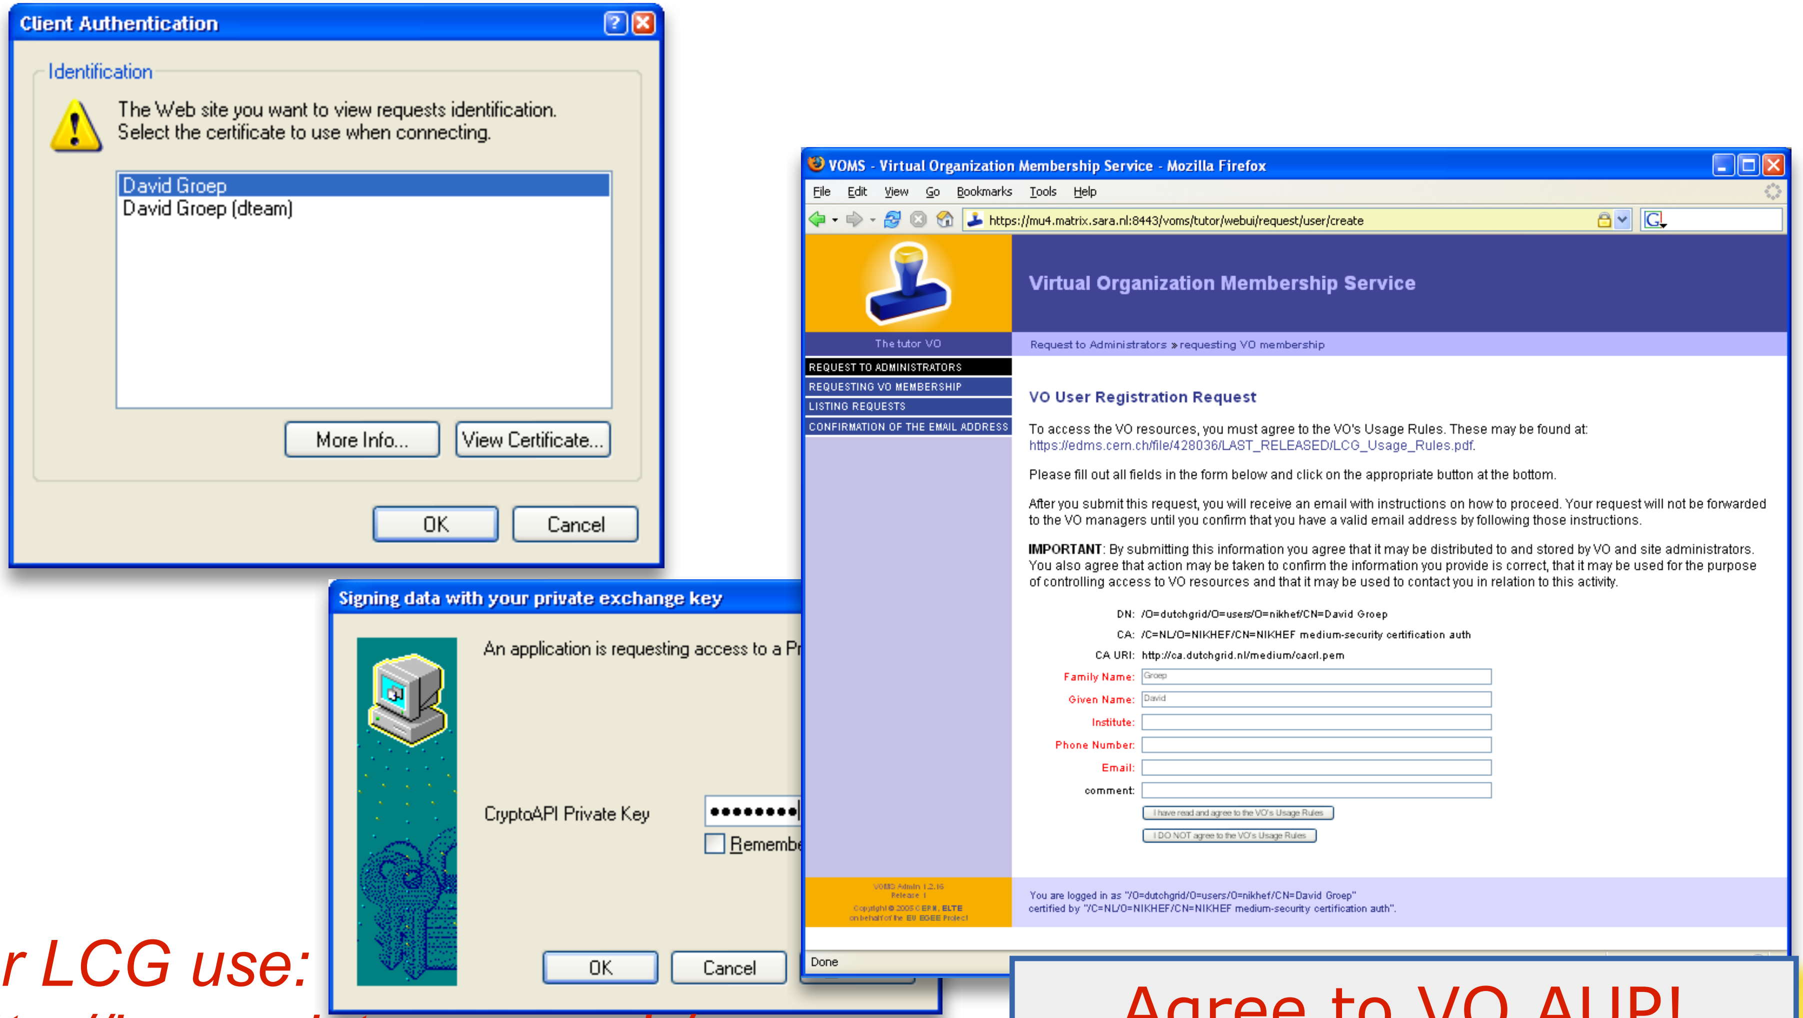

Client Authentication

Identification

The Web site you want to view requests identification. Select the certificate to use when connecting.

David Groep
David Groep (dteam)

Signing data with your private exchange key

An application is requesting access to a Pr

CryptoAPI Private Key ••••••••

□ Rememb

VOMS - Virtual Organization Membership Service - Mozilla Firefox

https://mu4.matrix.sara.nl:8443/voms/tutor/webui/request/user/create

Virtual Organization Membership Service

The tutor VO

Request to Administrators » requesting VO membership

VO User Registration Request

To access the VO resources, you must agree to the VO's Usage Rules. These may be found at: https://edms.cern.ch/file/428036/LAST_RELEASED/LCG_Usage_Rules.pdf.

Please fill out all fields in the form below and click on the appropriate button at the bottom.

After you submit this request, you will receive an email with instructions on how to proceed. Your request will not be forwarded to the VO managers until you confirm that you have a valid email address by following those instructions.

IMPORTANT: By submitting this information you agree that it may be distributed to and stored by VO and site administrators. You also agree that action may be taken to confirm the information you provide is correct, that it may be used for the purpose of controlling access to VO resources and that it may be used to contact you in relation to this activity.

DN: /O=dutchgrid/O=users/O=nikhef/CN=David Groep
CA: /C=NL/O=NIKHEF/CN=NIKHEF medium-security certification auth
CA URI: http://ca.dutchgrid.nl/medium/cacrl.pem
Family Name: Groep
Given Name: David
Institute:
Phone Number:
Email:
comment:

I have read and agree to the VO's Usage Rules
I DO NOT agree to the VO's Usage Rules

VOMS Admin 1.2.16
Release 1
Copyright © 2005 CERN, ELTE
on behalf of the EU EGEE Project

You are logged in as "/O=dutchgrid/O=users/O=nikhef/CN=David Groep" certified by "/C=NL/O=NIKHEF/CN=NIKHEF medium-security certification auth".

Done

for LCG use:
http://lcg-registrar.cern.ch/

Agree to VO AUP!

VO affiliation

- Per-VO Authorisations ("visa")
  - granted to a person or service by a virtual organisation
  - based on the 'passport' name
  - acknowledged by the resource owners
  - providers can still ban individual users,
and decide which privileges are granted to which VO attributes

- In your case, these 'visa' are called ***VOMS credentials***
- It is a cryptographically protected statement **by the VO**
- which is bound (by the VO) to your subject name

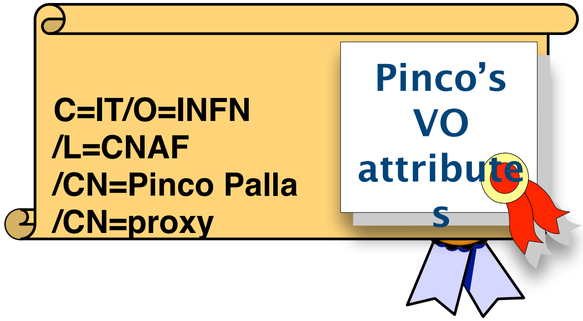

Embedding your VO affiliation

- The proxy can also be used as a *container* for other stuff
  - a 'plain' grid proxy does not indicate which VO you belong to
  - the VOMS credential is embedded as an *extension* in the proxy

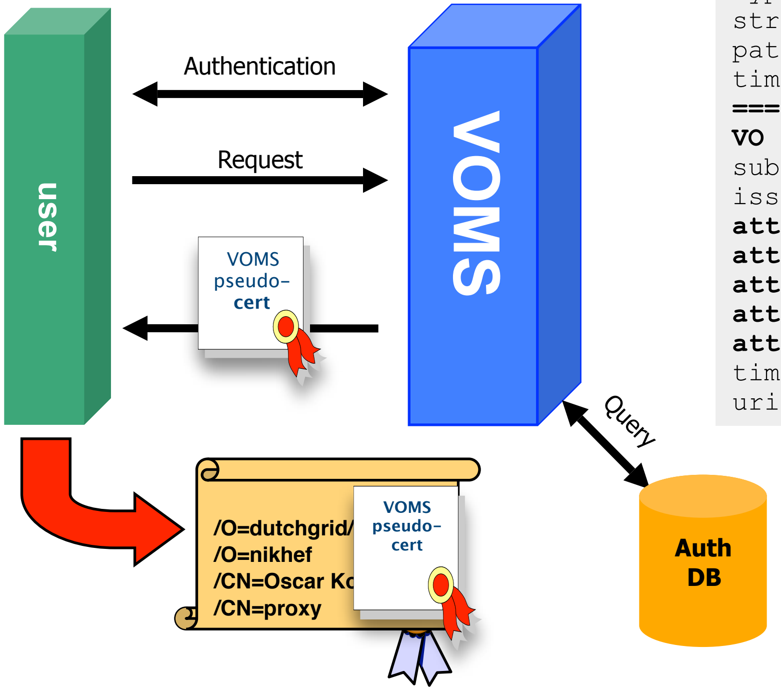

```
VisionMaster:~ okoeroo$ voms-proxy-info -all
subject   : /O=dutchgrid/O=users/O=nikhef/CN=Oscar Koeroo/CN=proxy
issuer    : /O=dutchgrid/O=users/O=nikhef/CN=Oscar Koeroo
identity  : /O=dutchgrid/O=users/O=nikhef/CN=Oscar Koeroo
type      : proxy
strength  : 1024 bits
path      : /tmp/x509up_u501
timeleft  : 11:59:30
=== VO dteam extension information ===
VO        : dteam
subject   : /O=dutchgrid/O=users/O=nikhef/CN=Oscar Koeroo
issuer    : /DC=ch/DC=cern/OU=computers/CN=voms.cern.ch
attribute : /dteam/Role=NULL/Capability=NULL
attribute : /dteam/ne/Role=NULL/Capability=NULL
attribute : /dteam/ne/SE/Role=NULL/Capability=NULL
attribute : /dteam/ne/SE/PDC/Role=NULL/Capability=NULL
attribute : /dteam/ne/pdc/Role=NULL/Capability=NULL
timeleft  : 11:59:40
uri       : voms.cern.ch:15004
```

Embedding your VO affiliation

- The proxy can also be used as a *container* for other stuff
  - a ‘plain’ grid proxy does not indicate which VO you belong to
  - the VOMS credential is embedded as an *extension* in the proxy

```
VisionMaster:~ okoeroo$ voms-proxy-info -all
subject   : /O=dutchgrid/O=users/O=nikhef/CN=Oscar Koeroo/CN=proxy
issuer    : /O=dutchgrid/O=users/O=nikhef/CN=Oscar Koeroo
identity  : /O=dutchgrid/O=users/O=nikhef/CN=Oscar Koeroo
type      : proxy
strength  : 1024 bits
path      : /tmp/x509up_u501
timeleft  : 11:59:30
=== VO dteam extension information ===
VO        : dteam
subject   : /O=dutchgrid/O=users/O=nikhef/CN=Oscar Koeroo
issuer    : /DC=ch/DC=cern/OU=computers/CN=voms.cern.ch
attribute : /dteam/Role=NULL/Capability=NULL
attribute : /dteam/ne/Role=NULL/Capability=NULL
attribute : /dteam/ne/SE/Role=NULL/Capability=NULL
attribute : /dteam/ne/SE/PDC/Role=NULL/Capability=NULL
attribute : /dteam/ne/pdc/Role=NULL/Capability=NULL
timeleft  : 11:59:40
uri       : voms.cern.ch:15004
```

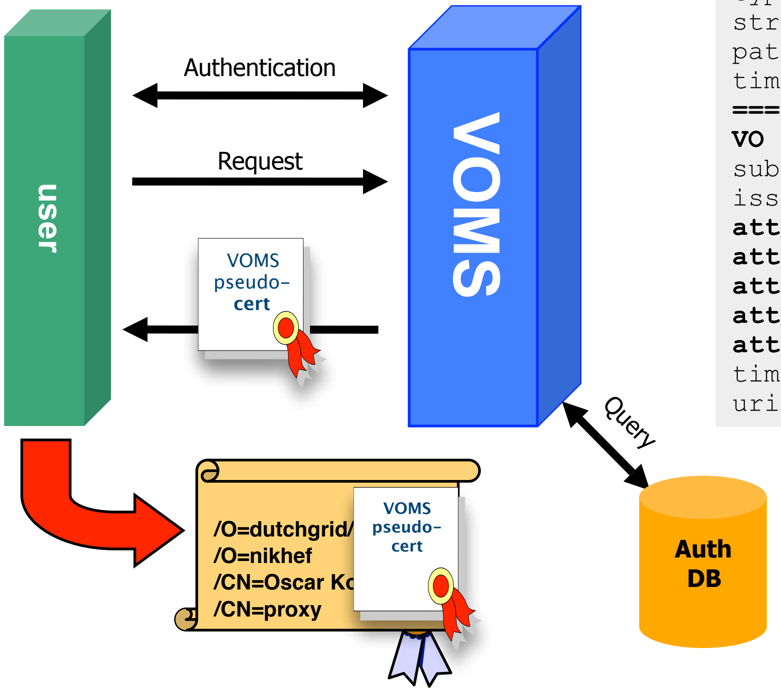

FQAN:
Fully Qualified Attribute Name

But what do users need to do?

- A Grid Security walk-through...

A walk-through

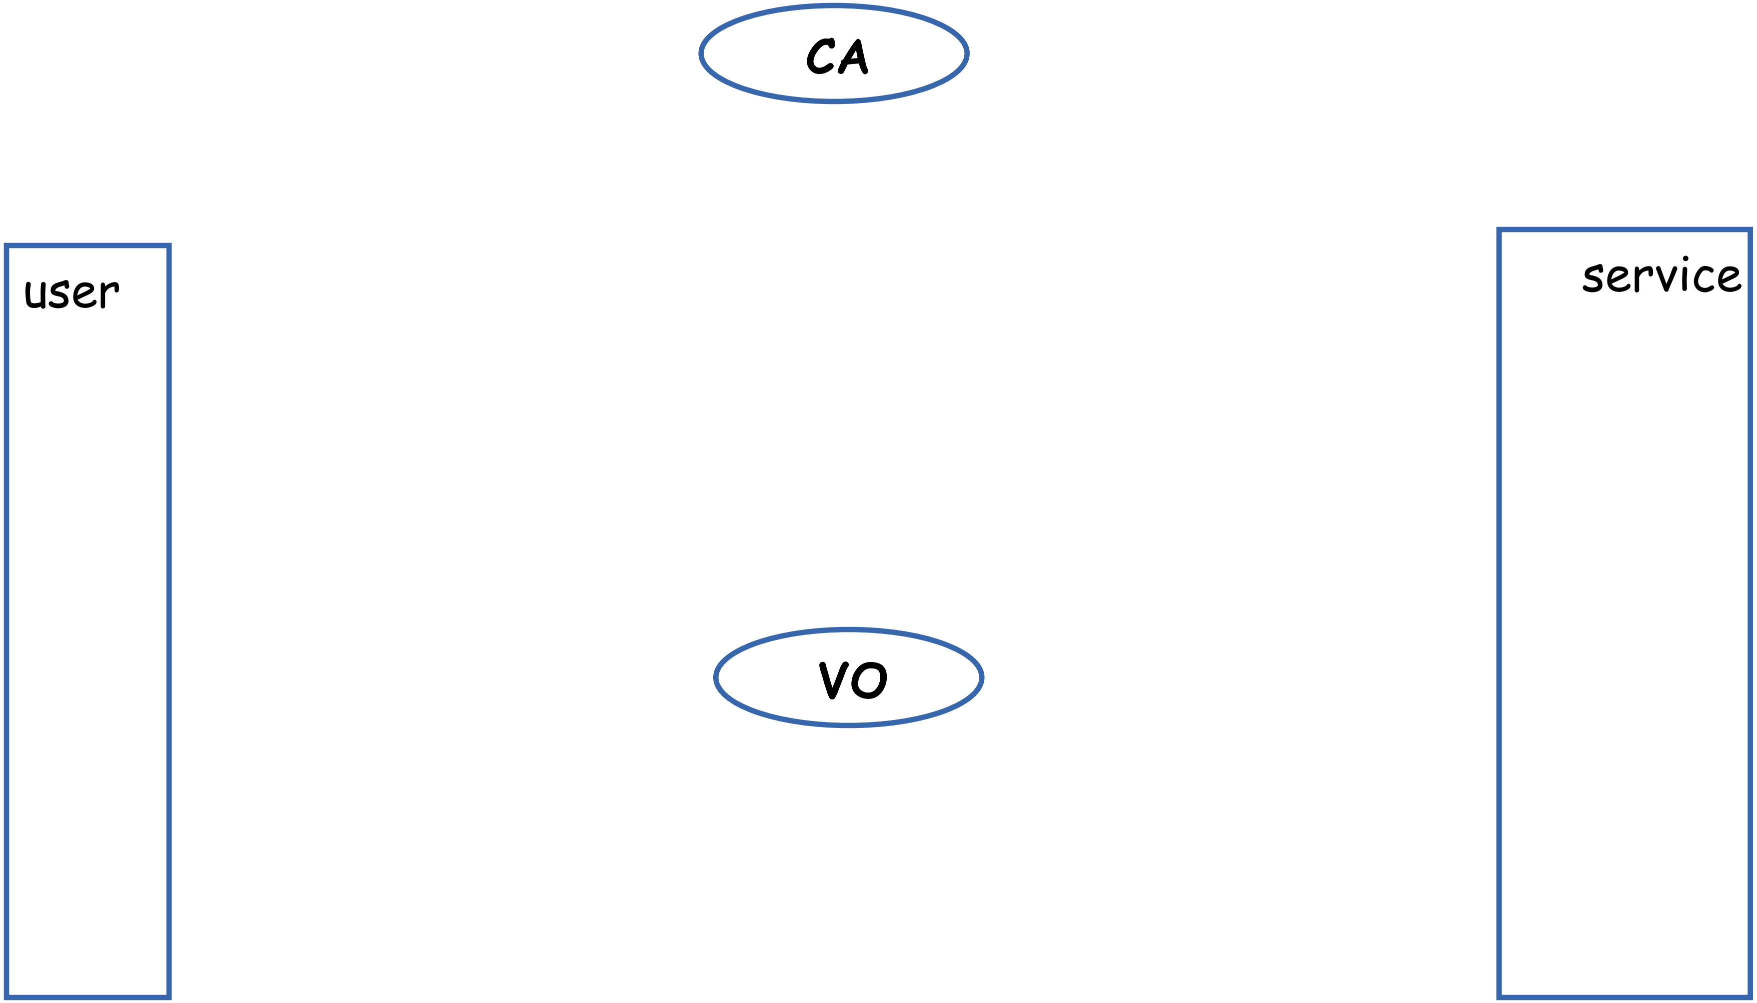

Certificate request

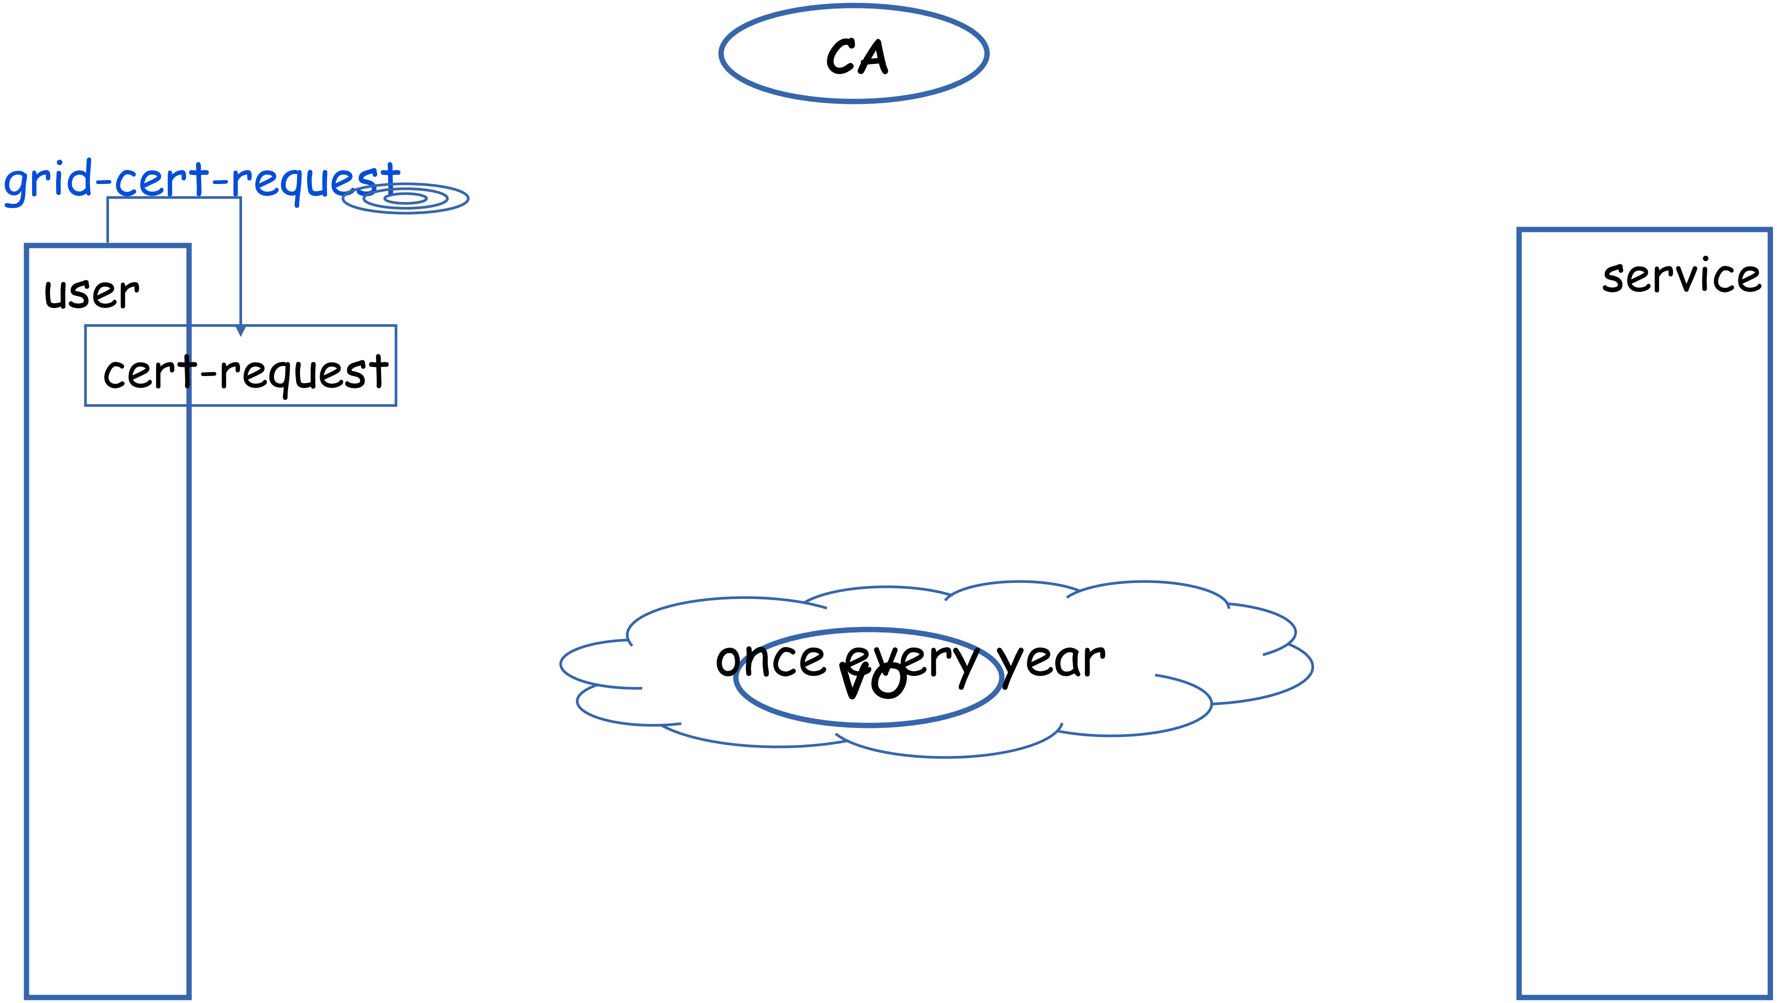

Certificate signing

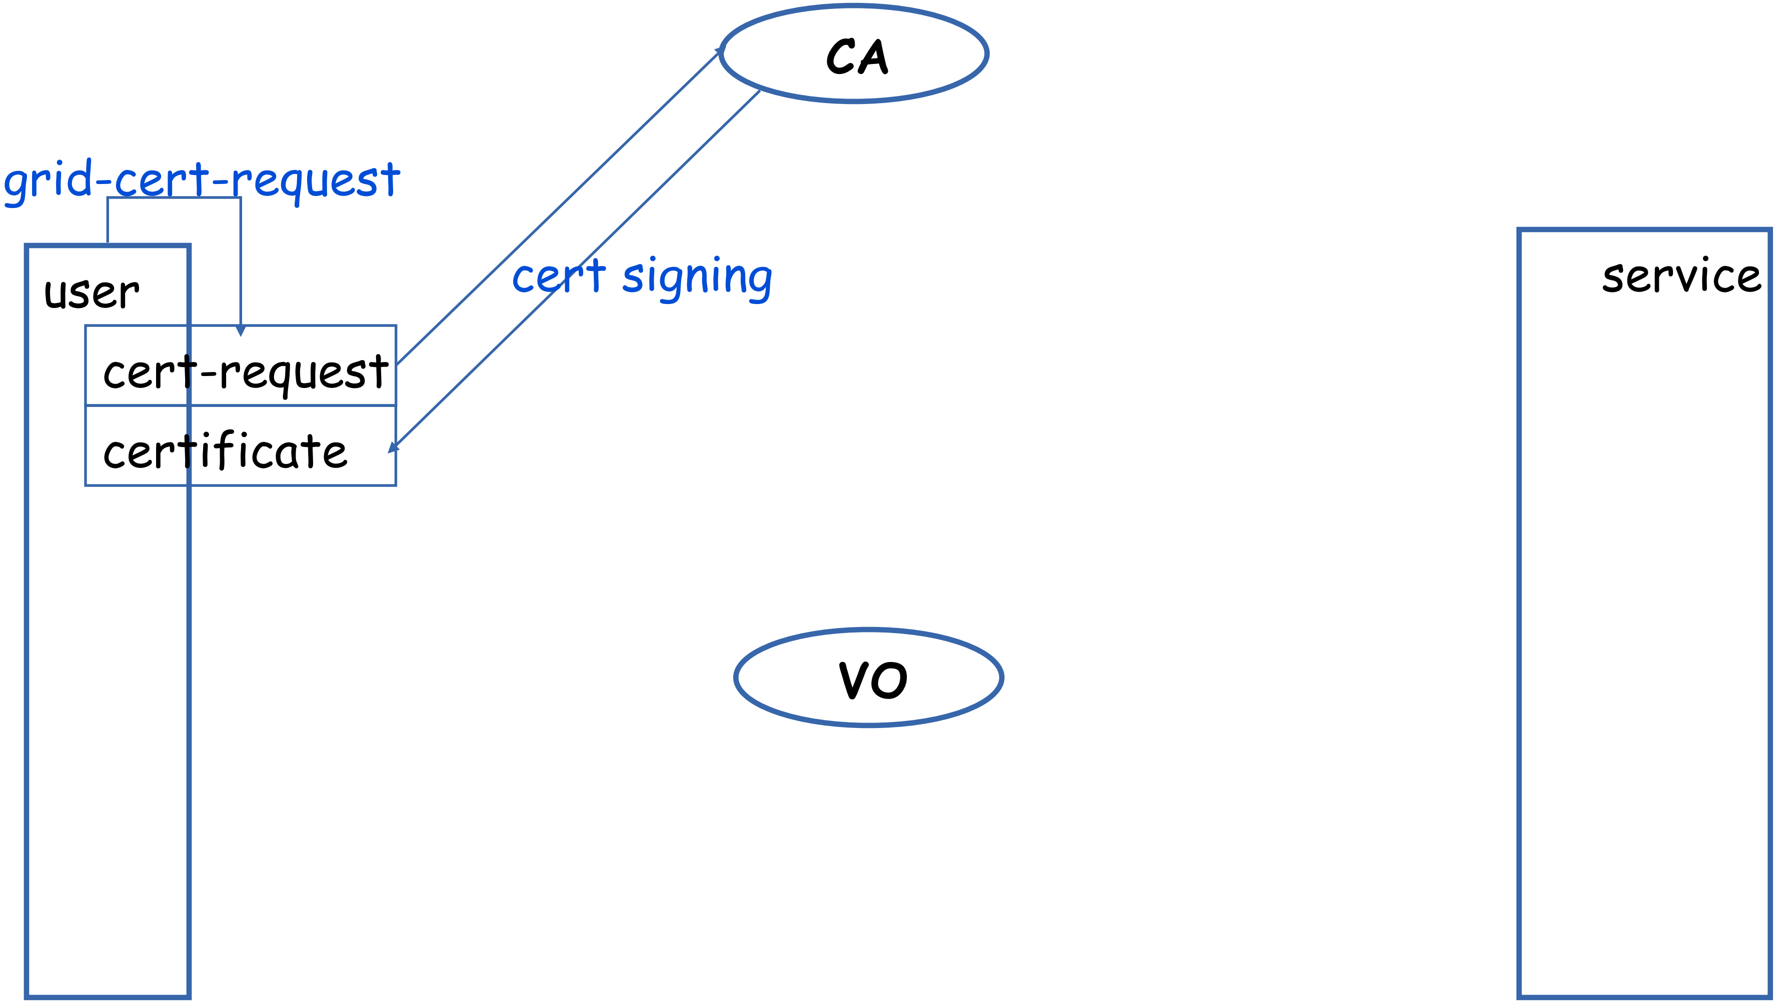

Importing your certificate in the browser

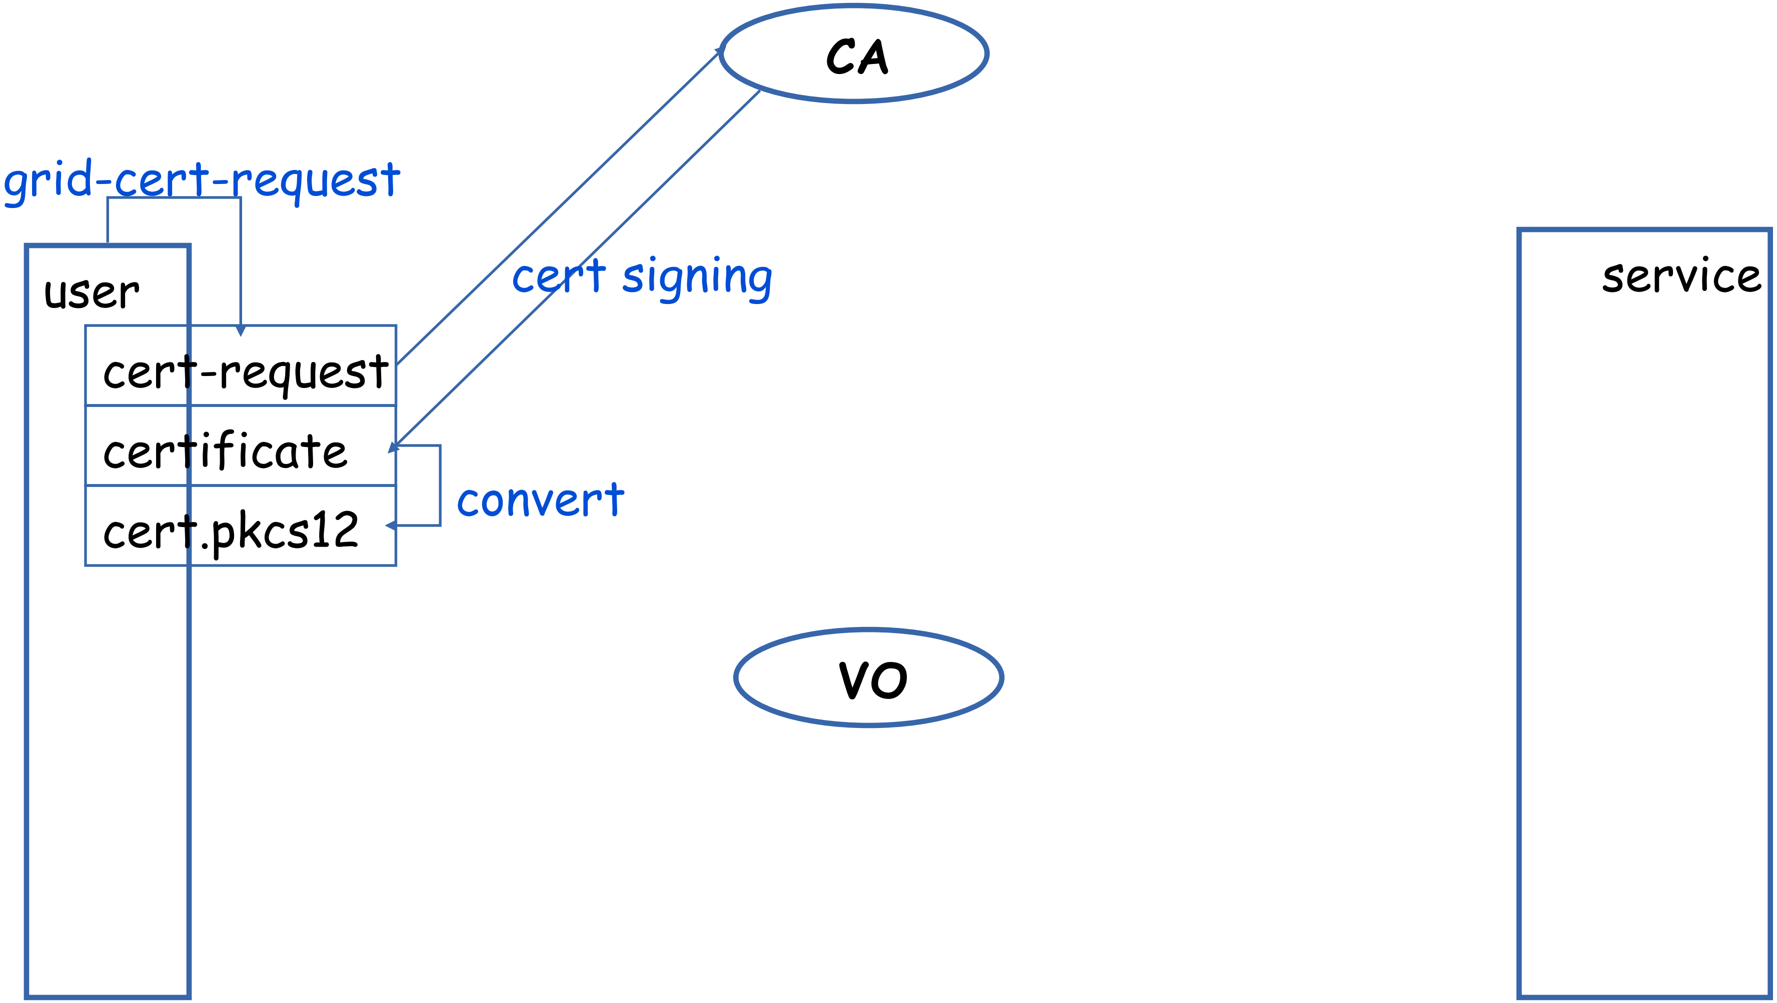

Configuration on the Server

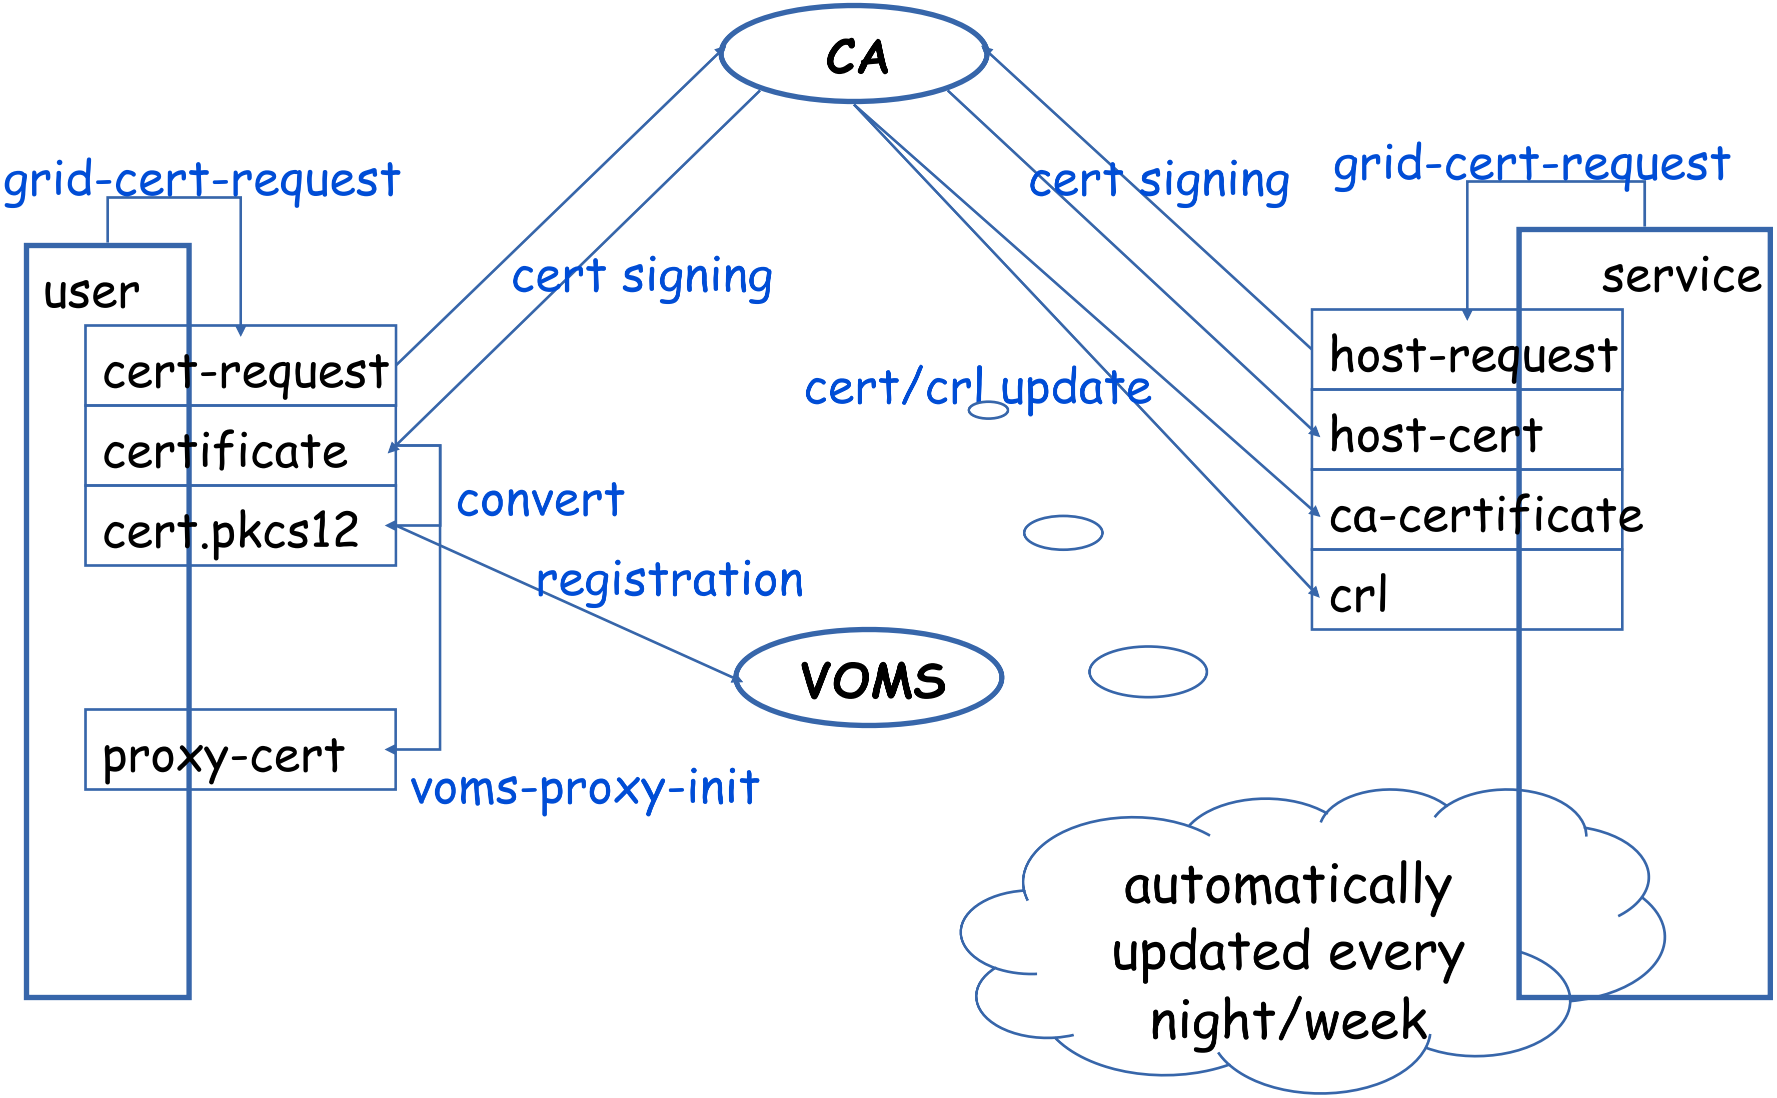

Using a Service

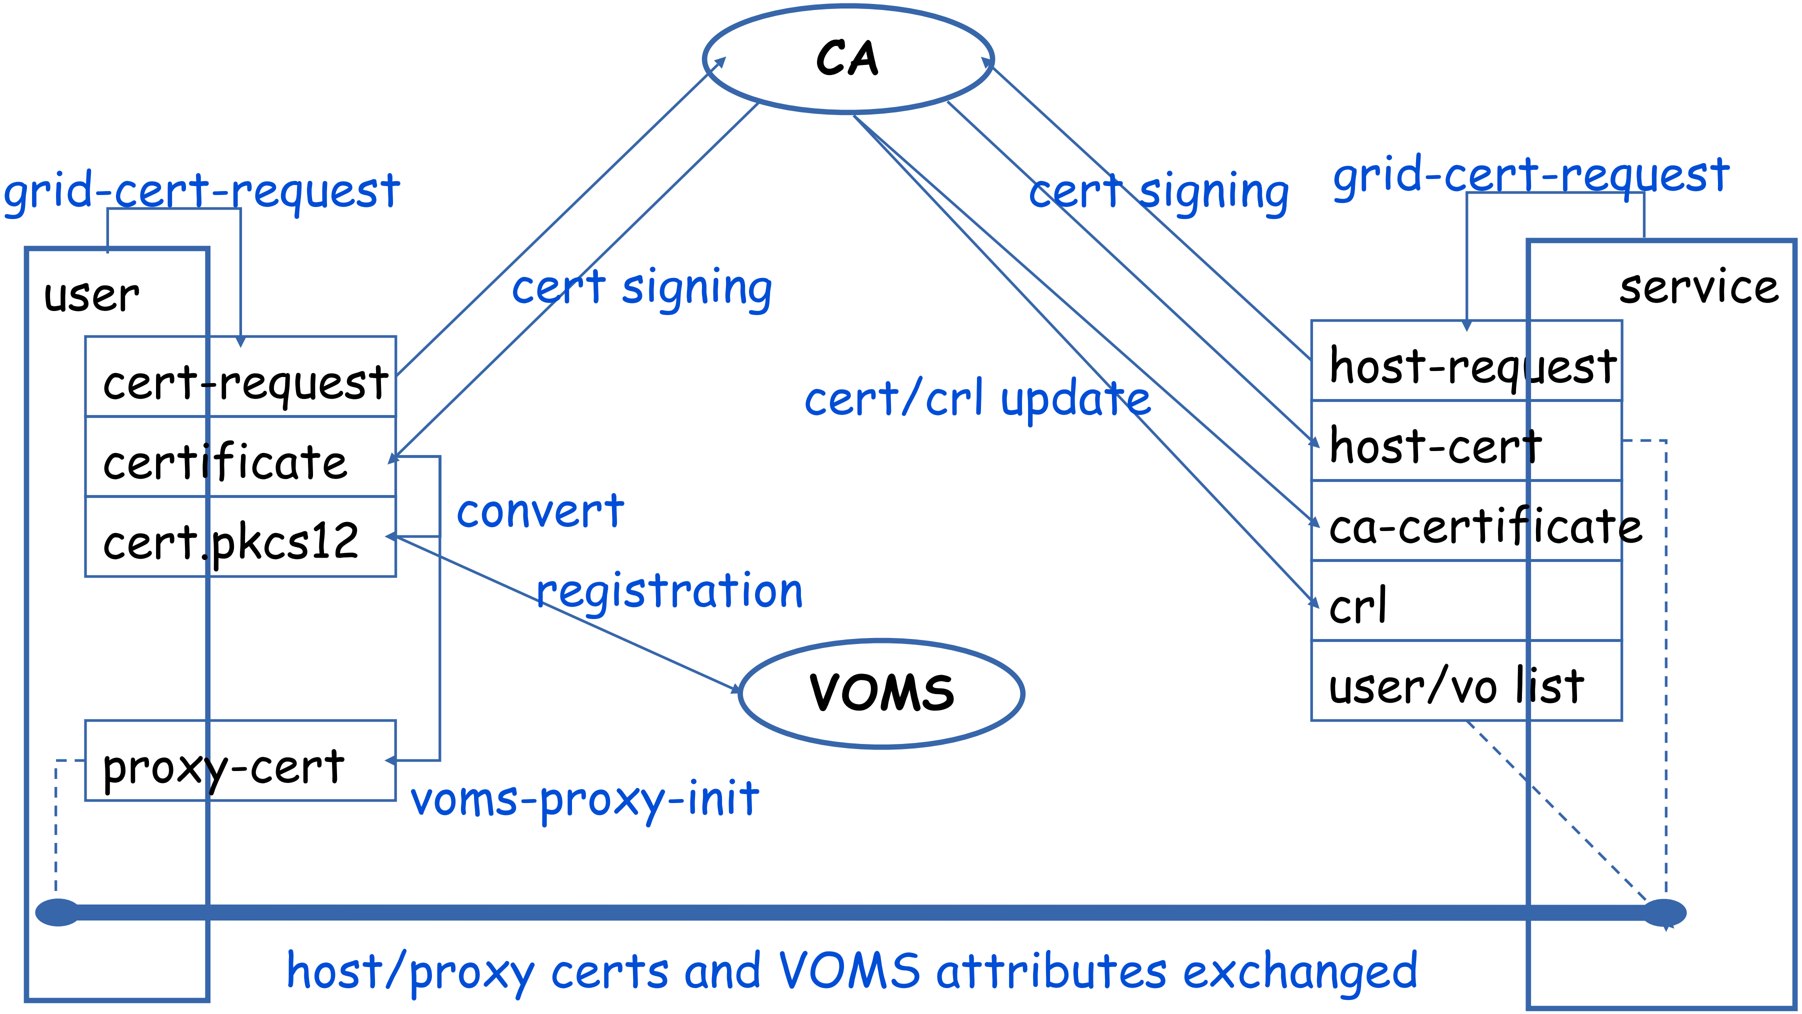

Intermediate summary

CA: authentication

VO: AUP, authorization and access

- new certificate: follow the web page instructions
- send to the appropriate CA (e.g. ca@dutchgrid.nl)
- save the answer
  - ~/.globus/usercert.pem
- import in web browser (.p12) and register with VO

only once

- new proxy certificate:
  - **voms-proxy-init –voms dteam**
- *use the Grid* :-)

~daily

List of credentials as input for the Grid (other info skipped due to time constraint)

- Subject ID of your certificate

  O=dutchgrid, O=users, O=nikhef, CN=Oscar Koeroo

- VO credentials

  Those FQANs which specify project, (sub)group and role affiliation

  /dteam/Role=VO-Admin

  /dteam/Role=software-manager

1. All credentials are cryptographically tied to each user
2. VO credentials (VOMS) sits inside a proxy certificate

AuthZ & Mapping

The tools

Job Submission Today

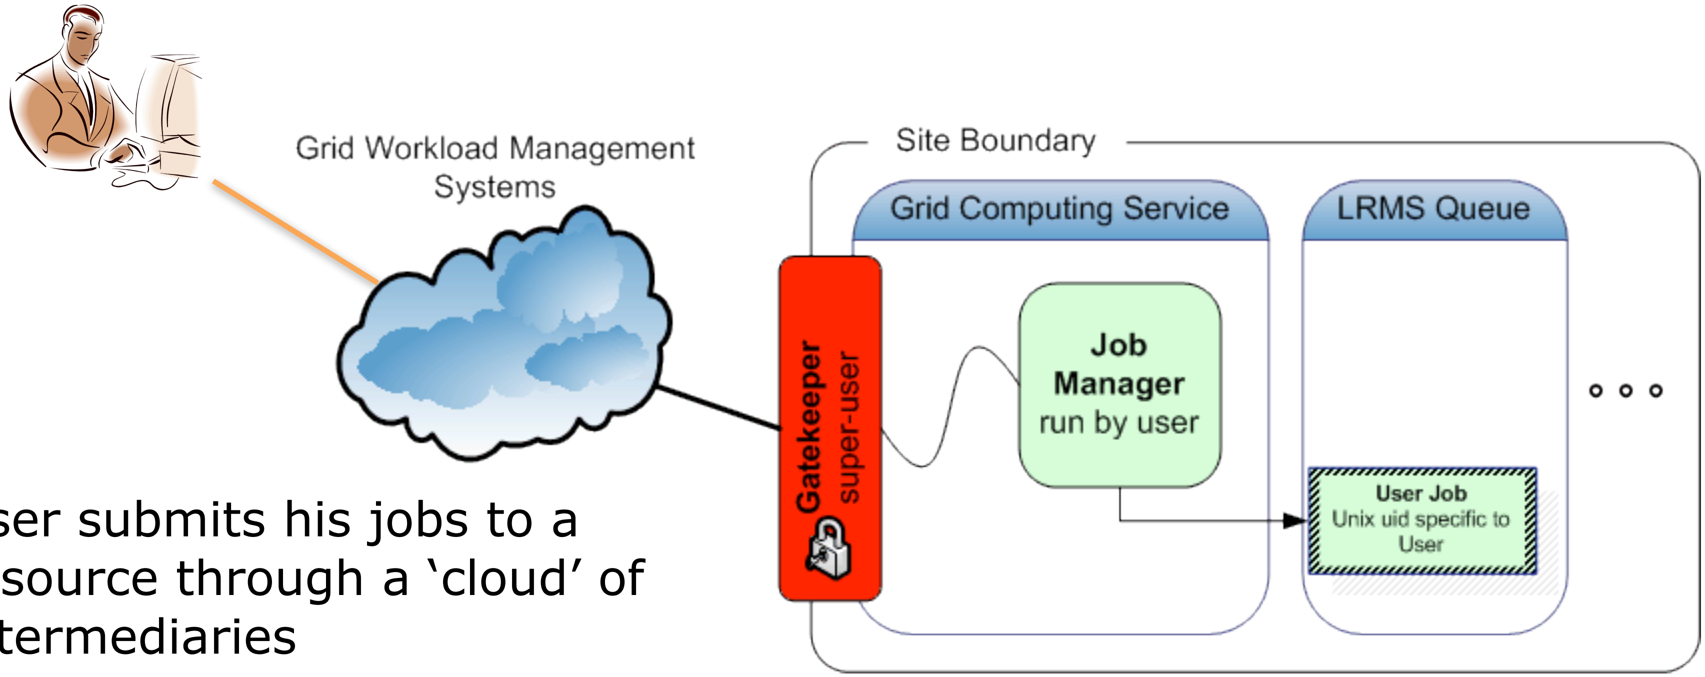

User submits his jobs to a resource through a 'cloud' of intermediaries

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

Example CE workflow

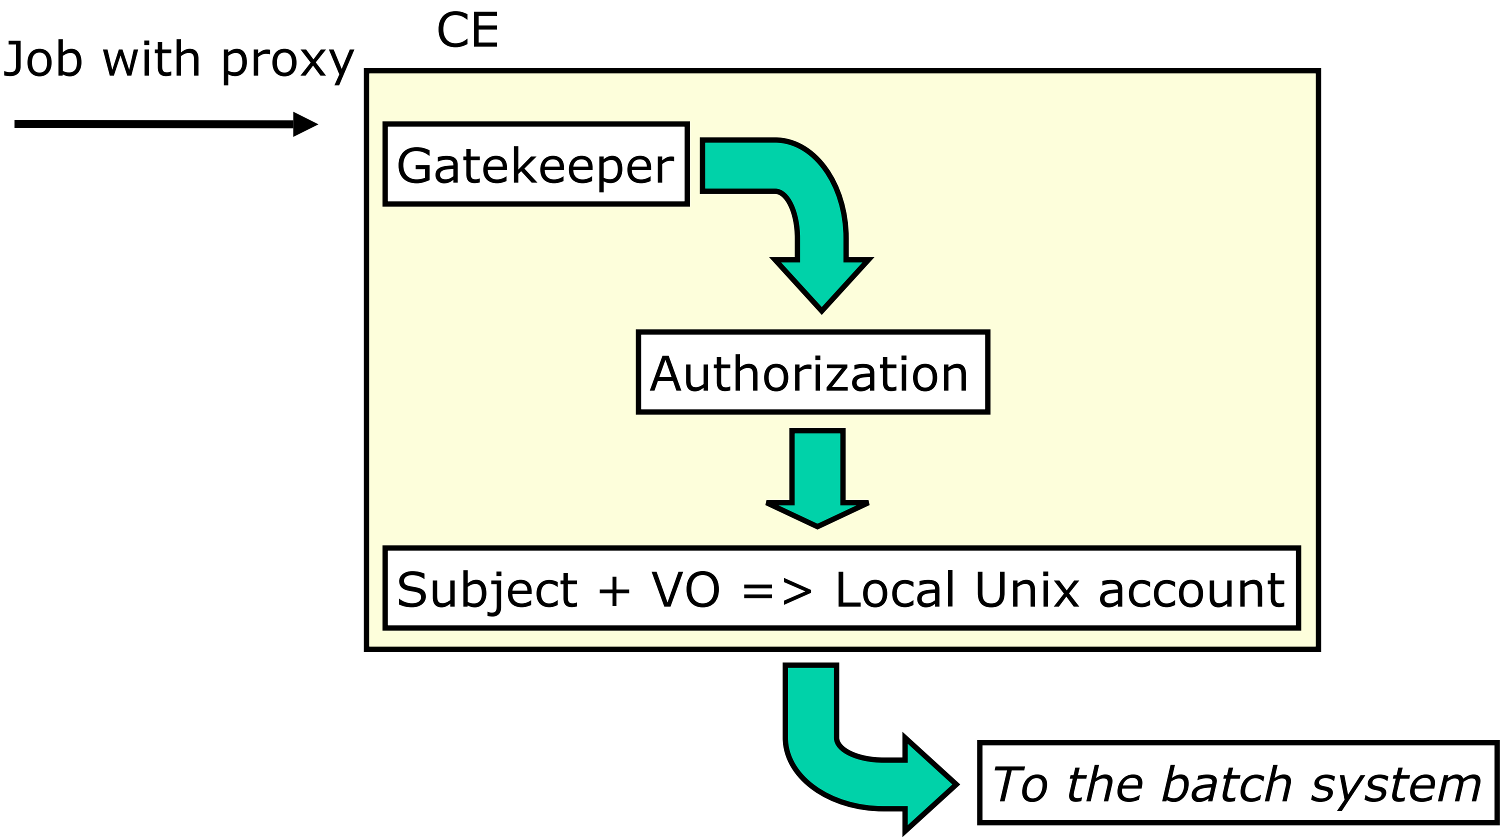

A multi-authority world

> Authorization elements

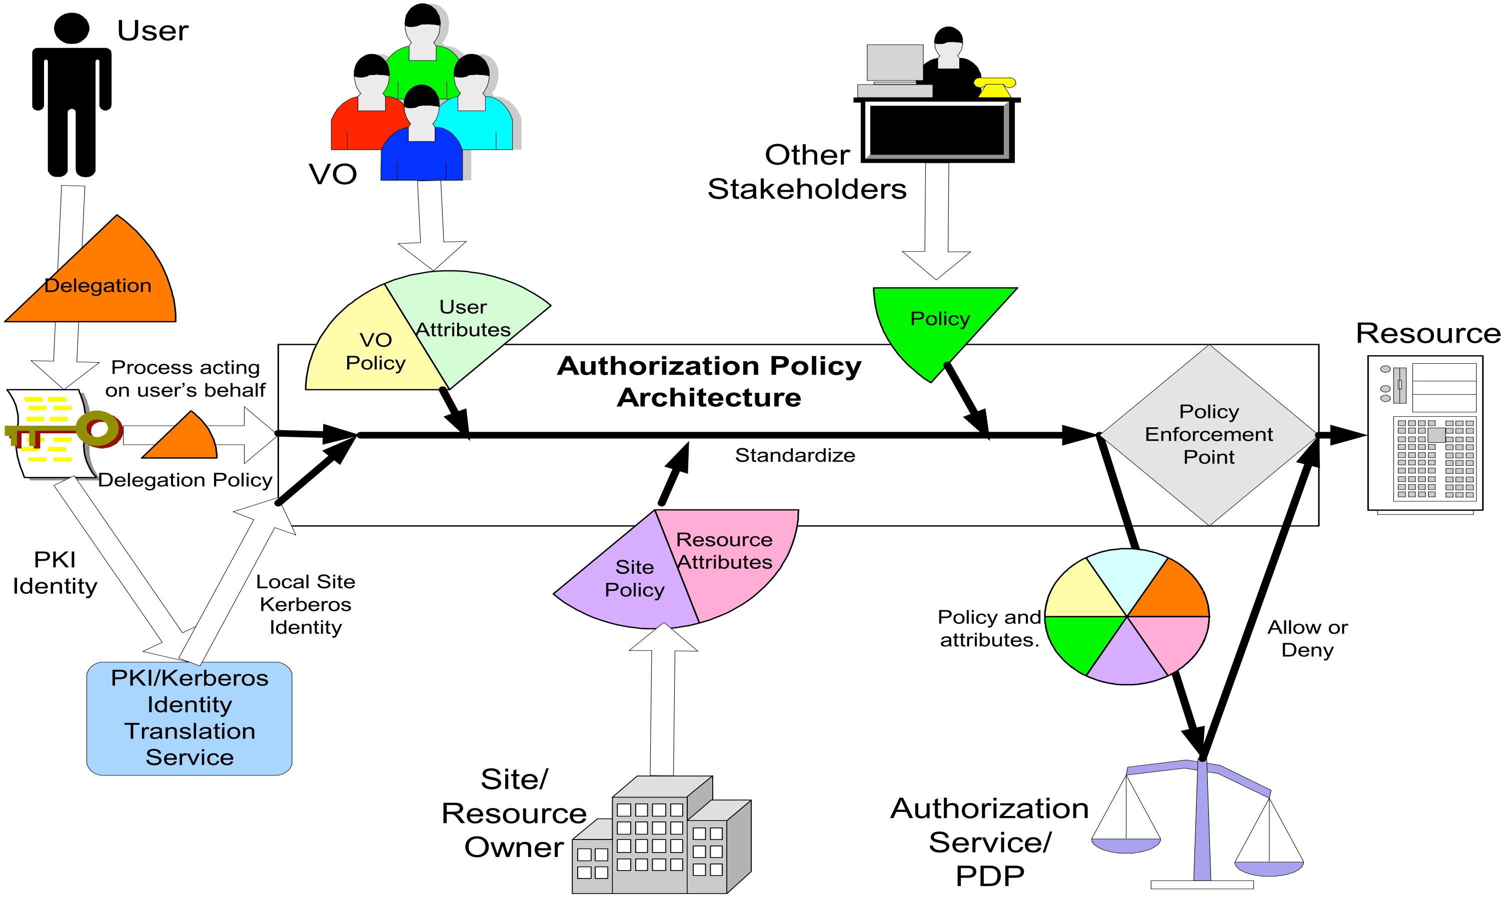

Graphic: OGSA Working Group

Authorization based on Subject-ID & VOMS

- Lots of different tools, libs and frameworks
  - All read a special file called ‘grid-mapfile’ or something similar (like from a database)
  - All give a binary ‘yes allowed’ or ‘no not allowed’ before giving access to any type of resource
- You could be banned from a site or globally from the Grid

```
[…]
"/O=dutchgrid/O=users/O=nikhef/CN=Jeffrey Templon”
"/O=dutchgrid/O=users/O=nikhef/CN=Martijn Steenbakkers"
"/O=dutchgrid/O=users/O=nikhef/CN=Oscar Koeroo"
"/O=gridtutorial/O=users/O=grid-tutorial/CN=Grid pupil 20"
"/cms/muon"
"/atlas/Production"
"/atlas/*"
"/lhcb/*"
```

To the Unix world...

grid identity

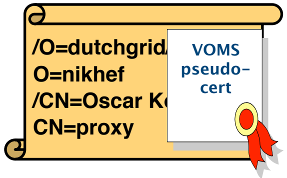

"/O=dutch[...]/CN=Oscar Koeroo"

of group "/pvier"

translate

```
pvier001:x:43401:2029:PoolAccount VL-e P4 no.1:/home/pvier001:/bin/sh
```

1. Unix does not talk Grid, so
   translation is needed between grid and local identity

1. this translation has to happen somewhere
2. something needs to do that

Account mapping libraries and services

- Clusters are typically Unix systems
  - Unix systems have accounts and groups
    uid=1001(okoeroo) gid=100(users) groups=16(cron)

- Mapping tools take the Subject-id and VOMS of the user and translate them to local Unix accounts
  - Per (sub)group affiliation
  - Per individual (bound to a group)

- People get mapped to poolaccounts
  - Poolaccounts are accounts unbound to a user at setup
  - Poolaccounts are assigned to users when they arrive at a site
  - Result:
    - Possible to track users on a cluster without ever meeting them in person

Stakkato

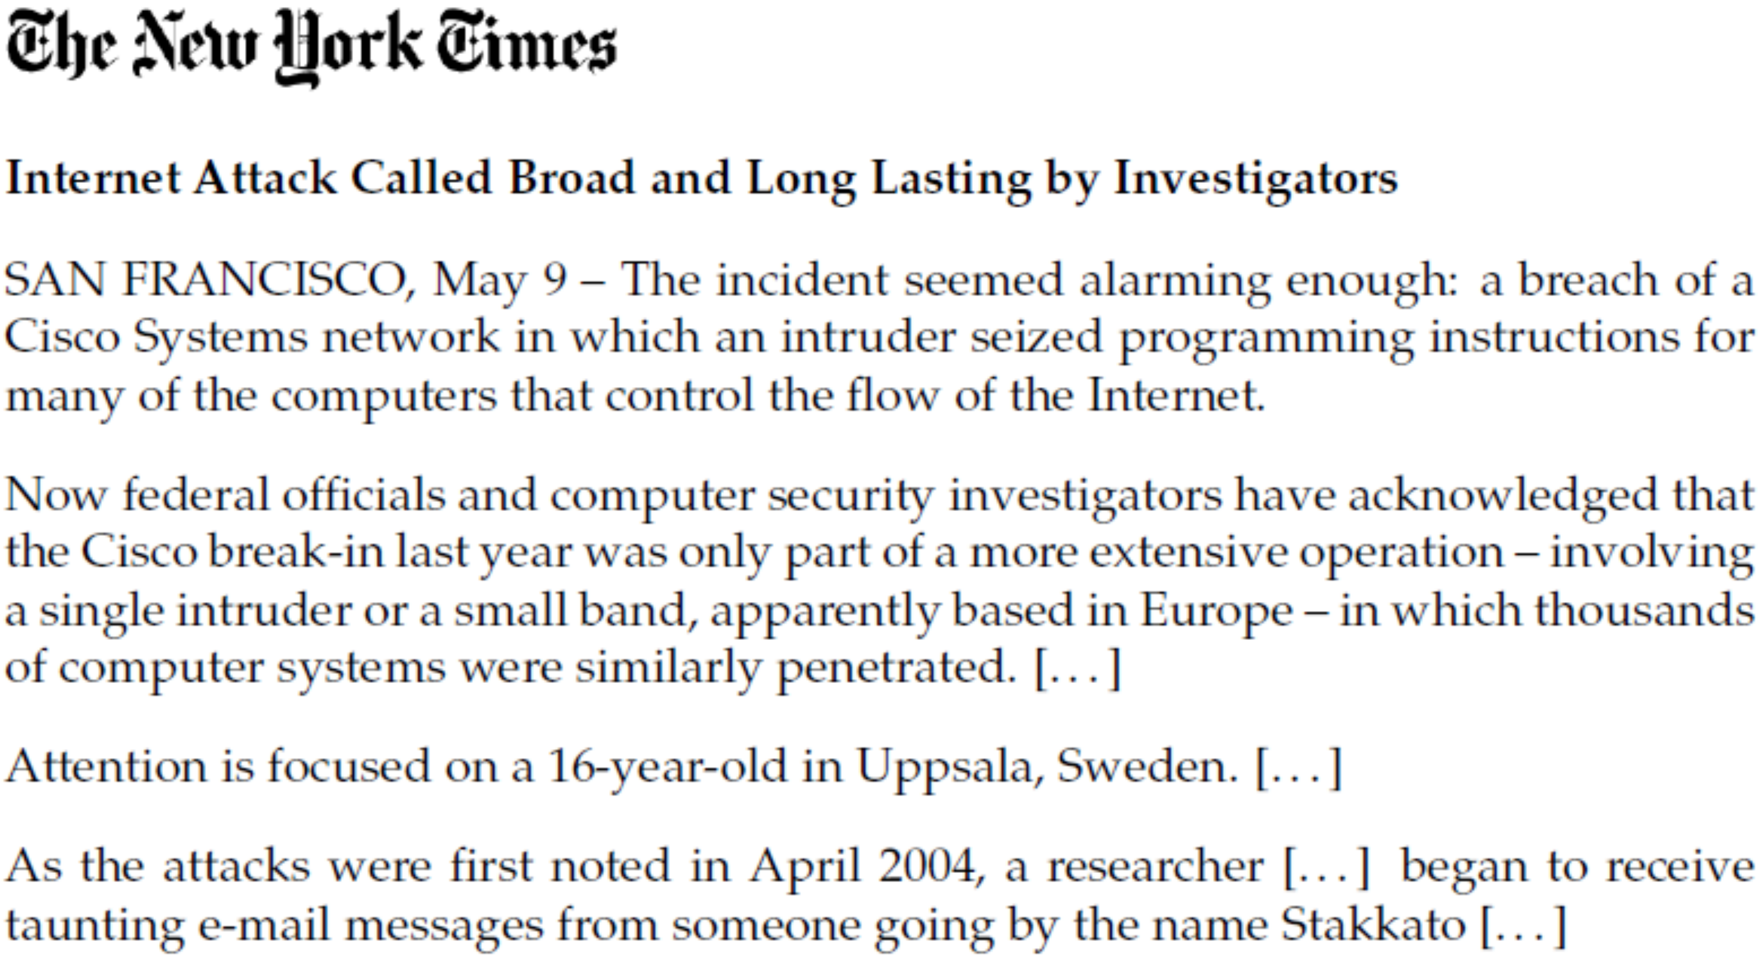

The New York Times

Internet Attack Called Broad and Long Lasting by Investigators

SAN FRANCISCO, May 9 – The incident seemed alarming enough: a breach of a Cisco Systems network in which an intruder seized programming instructions for many of the computers that control the flow of the Internet.

Now federal officials and computer security investigators have acknowledged that the Cisco break-in last year was only part of a more extensive operation – involving a single intruder or a small band, apparently based in Europe – in which thousands of computer systems were similarly penetrated. [. . .]

Attention is focused on a 16-year-old in Uppsala, Sweden. [. . .]

As the attacks were first noted in April 2004, a researcher [. . .] began to receive taunting e-mail messages from someone going by the name Stakkato [. . .]

Slide and info: Leif Nixon, NSC, Linköping – CCGrid06 key note “The Stakkato Intrusions”

Then, Nov 2007 and February 2008 …

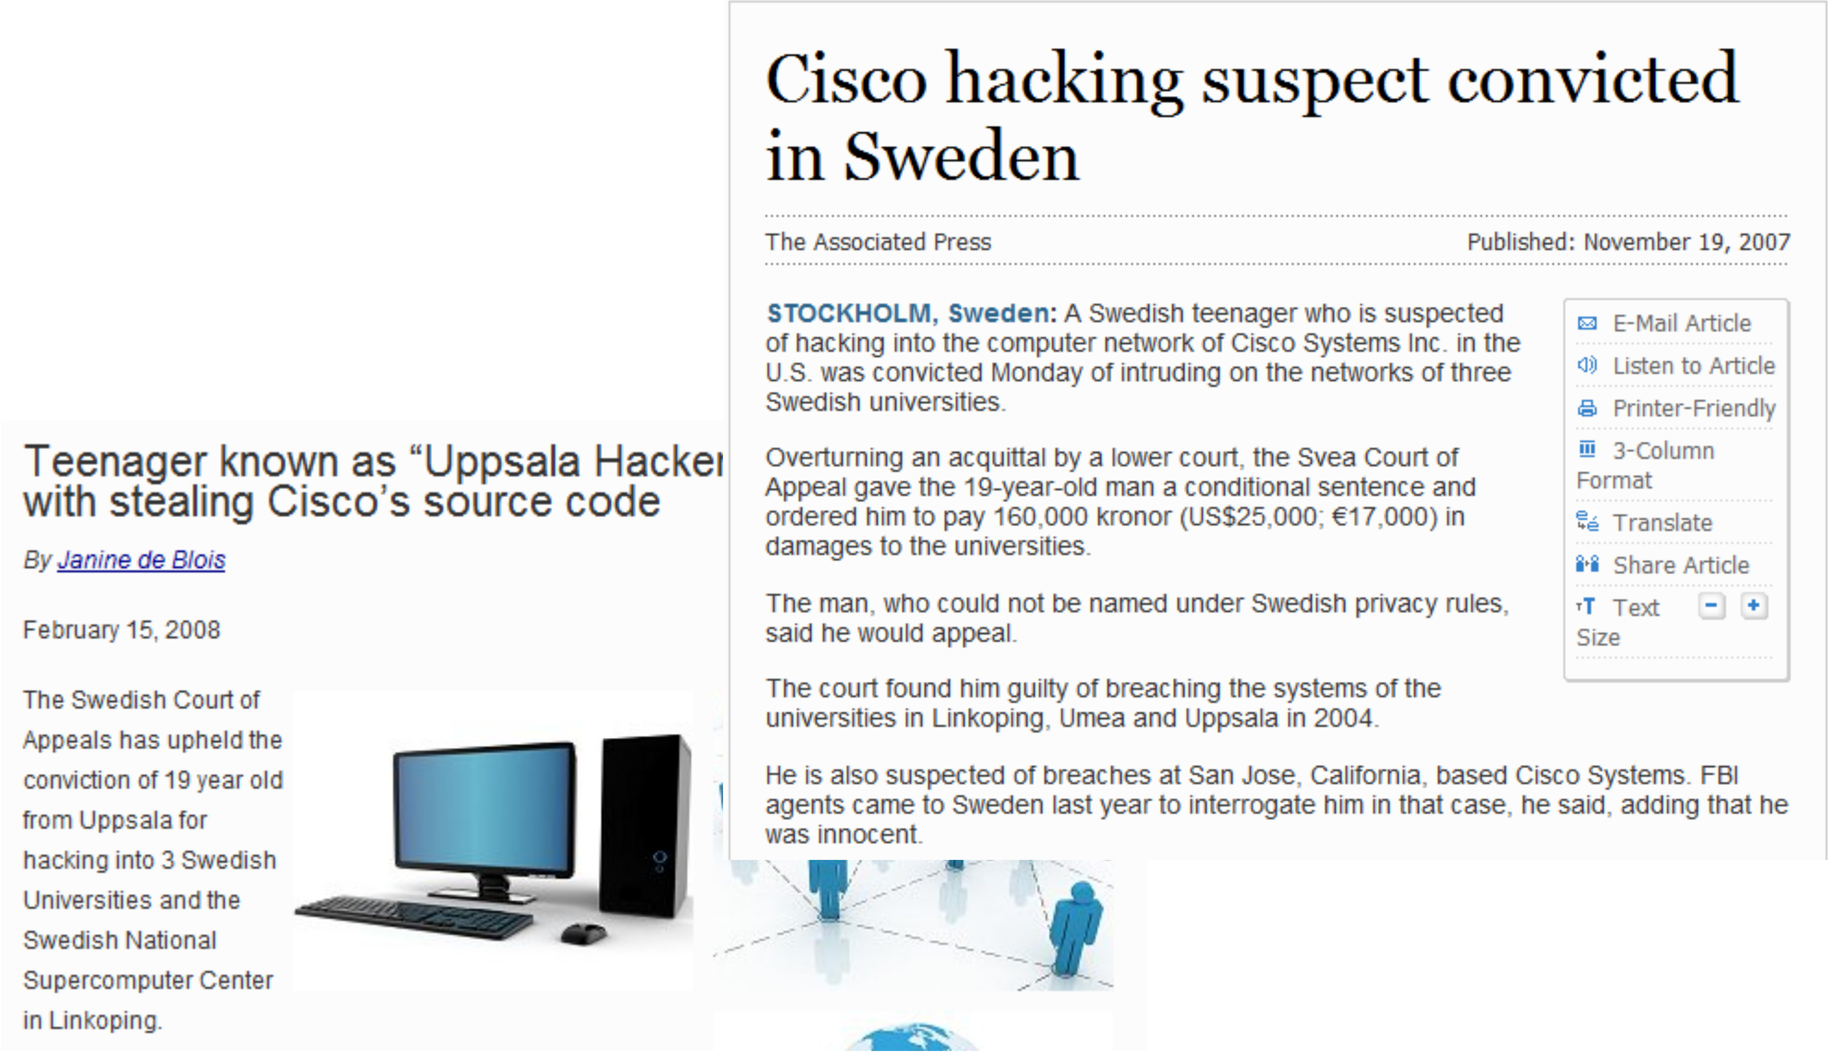

Cisco hacking suspect convicted in Sweden

The Associated Press   Published: November 19, 2007

STOCKHOLM, Sweden: A Swedish teenager who is suspected of hacking into the computer network of Cisco Systems Inc. in the U.S. was convicted Monday of intruding on the networks of three Swedish universities.

Overturning an acquittal by a lower court, the Svea Court of Appeal gave the 19-year-old man a conditional sentence and ordered him to pay 160,000 kronor (US$25,000; €17,000) in damages to the universities.

The man, who could not be named under Swedish privacy rules, said he would appeal.

The court found him guilty of breaching the systems of the universities in Linkoping, Umea and Uppsala in 2004.

He is also suspected of breaches at San Jose, California, based Cisco Systems. FBI agents came to Sweden last year to interrogate him in that case, he said, adding that he was innocent.

Teenager known as “Uppsala Hacker with stealing Cisco’s source code

By *Janine de Blois*

February 15, 2008

The Swedish Court of Appeals has upheld the conviction of 19 year old from Uppsala for hacking into 3 Swedish Universities and the Swedish National Supercomputer Center in Linkoping.

Is It Random: 3, 3, 3, 3, 3, 3, 3, 3, 3, 3, 3

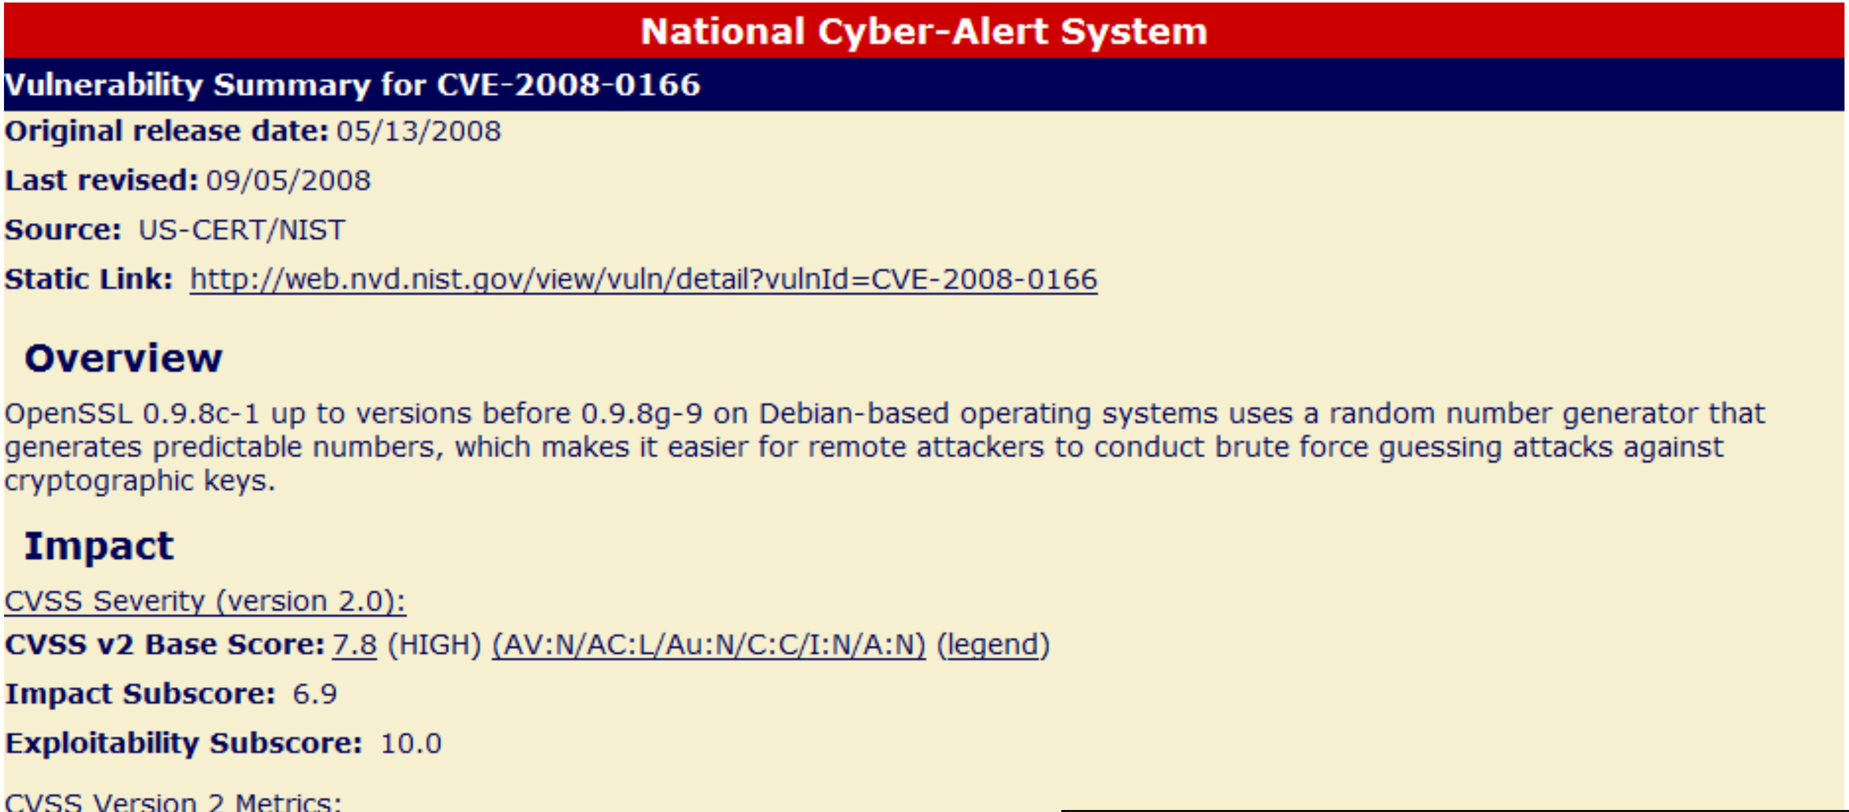

National Cyber-Alert System

Vulnerability Summary for CVE-2008-0166

Original release date: 05/13/2008

Last revised: 09/05/2008

Source: US-CERT/NIST

Static Link: http://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2008-0166

Overview

OpenSSL 0.9.8c-1 up to versions before 0.9.8g-9 on Debian-based operating systems uses a random number generator that generates predictable numbers, which makes it easier for remote attackers to conduct brute force guessing attacks against cryptographic keys.

Impact

CVSS Severity (version 2.0):

CVSS v2 Base Score: 7.8 (HIGH) (AV:N/AC:L/Au:N/C:C/I:N/A:N) (legend)

Impact Subscore: 6.9

Exploitability Subscore: 10.0

CVSS Version 2 Metrics:

Access Vector: Network exploitable

Access Complexity: Low

Authentication: Not required to exploit

Impact Type: Allows unauthorized disclosure of information

Only 163840 possible ssh keys!

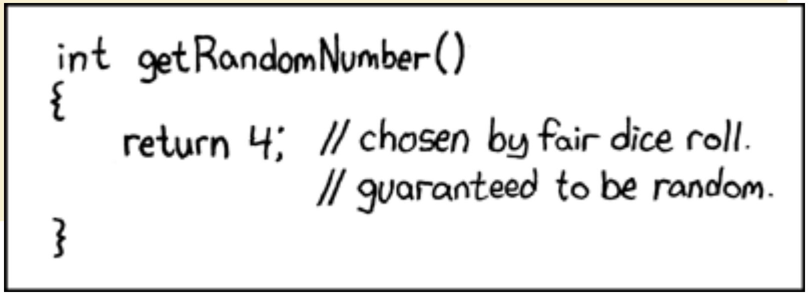

```
int getRandomNumber()
{
    return 4;  // chosen by fair dice roll.
               // guaranteed to be random.
}
```

http://xkcd.com/221/

More ssh

'XXXX-CERT-20080805'

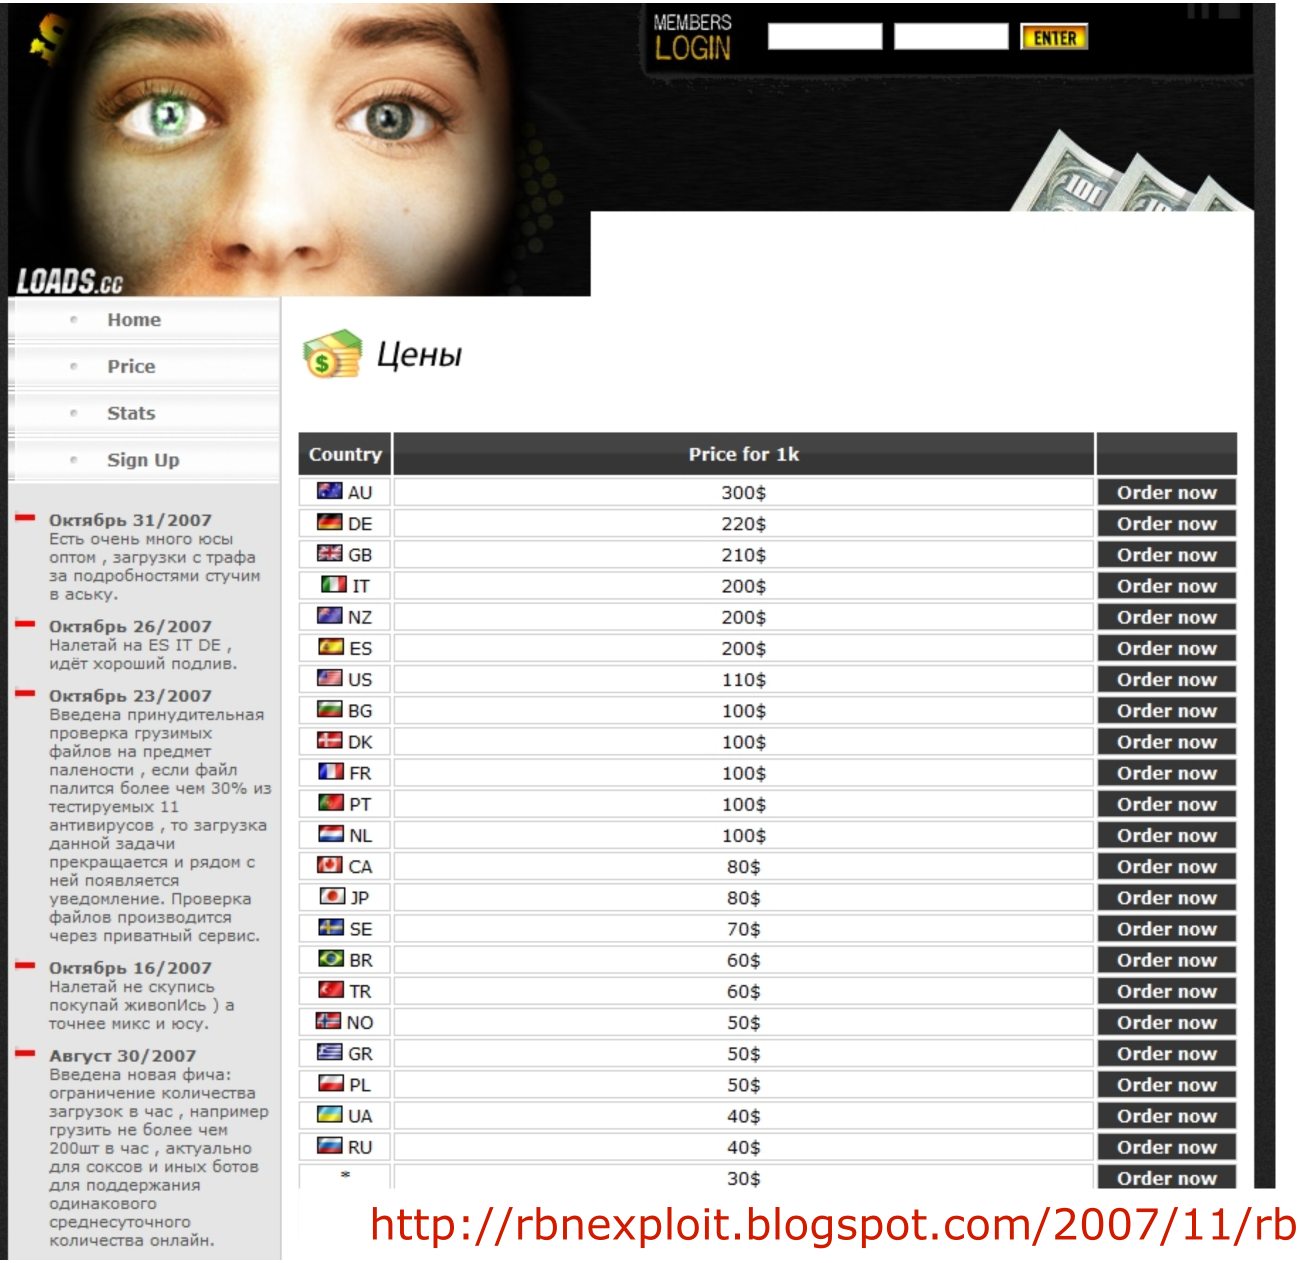

MEMBERS LOGIN

ENTER

LOADS.cc

- Home
- Price
- Stats
- Sign Up

Октябрь 31/2007
Есть очень много юсы оптом , загрузки с трафа за подробностями стучим в аську.

Октябрь 26/2007
Налетай на ES IT DE , идёт хороший подлив.

Октябрь 23/2007
Введена принудительная проверка грузимых файлов на предмет палености , если файл палится более чем 30% из тестируемых 11 антивирусов , то загрузка данной задачи прекращается и рядом с ней появляется уведомление. Проверка файлов производится через приватный сервис.

Октябрь 16/2007
Налетай не скупись покупай живопИсь) а точнее микс и юсу.

Август 30/2007
Введена новая фича: ограничение количества загрузок в час , например грузить не более чем 200шт в час , актуально для соксов и иных ботов для поддержания одинакового среднесуточного количества онлайн.

Цены

Country	Price for 1k	
AU	300$	Order now
DE	220$	Order now
GB	210$	Order now
IT	200$	Order now
NZ	200$	Order now
ES	200$	Order now
US	110$	Order now
BG	100$	Order now
DK	100$	Order now
FR	100$	Order now
PT	100$	Order now
NL	100$	Order now
CA	80$	Order now
JP	80$	Order now
SE	70$	Order now
BR	60$	Order now
TR	60$	Order now
NO	50$	Order now
GR	50$	Order now
PL	50$	Order now
UA	40$	Order now
RU	40$	Order now
*	30$	Order now

Price for 1000 infected consumer computers:

AU:	**US$ 300**
US:	**US$ 110**
NL:	**US$ 100**

And grid systems are better connected than xDSL systems, so …

http://rbnexploit.blogspot.com/2007/11/rbn-76-service-team-loads-cc-and-their.html

Incidents involving EGEE sites

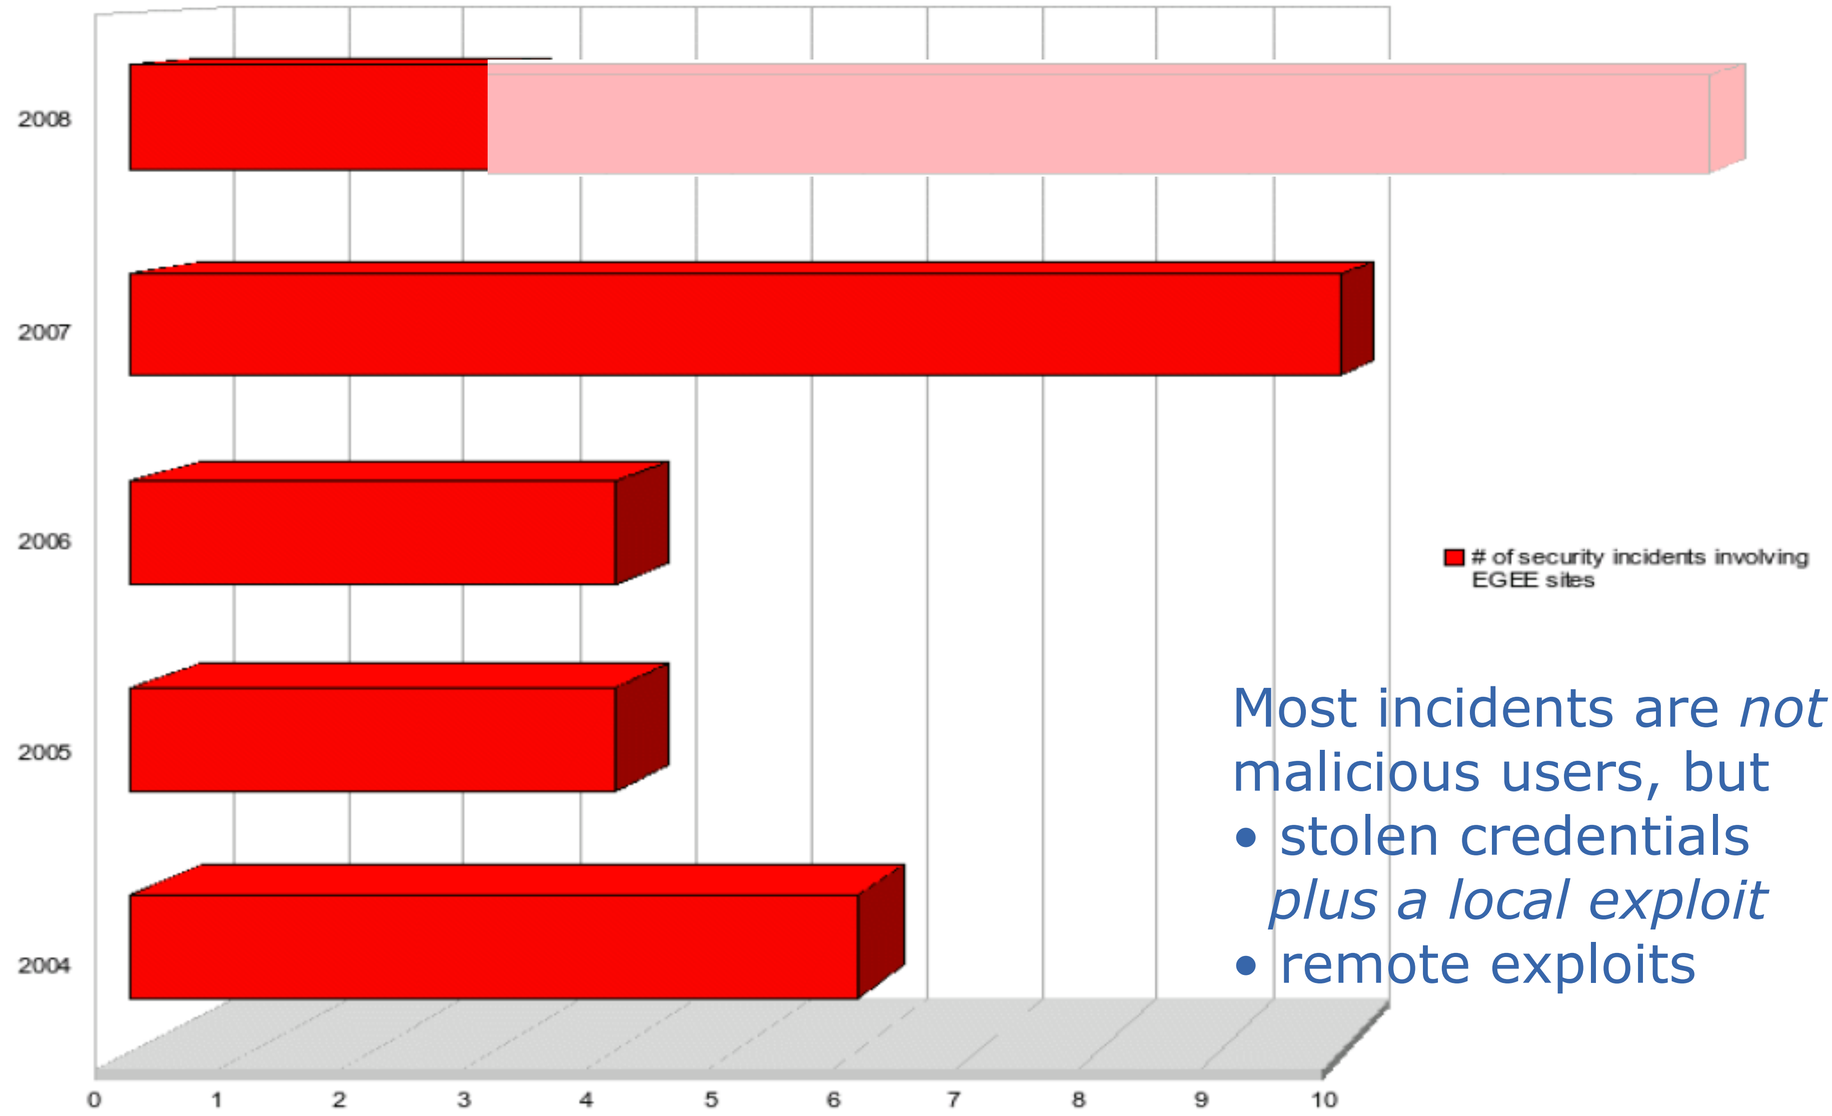

Most incidents are *not* malicious users, but

- stolen credentials *plus a local exploit*
- remote exploits

Romain Wartel, CERN and OSCT; http://romain.wartel.net/talks/20080409Wartel-short.pdf

But What About Containment?

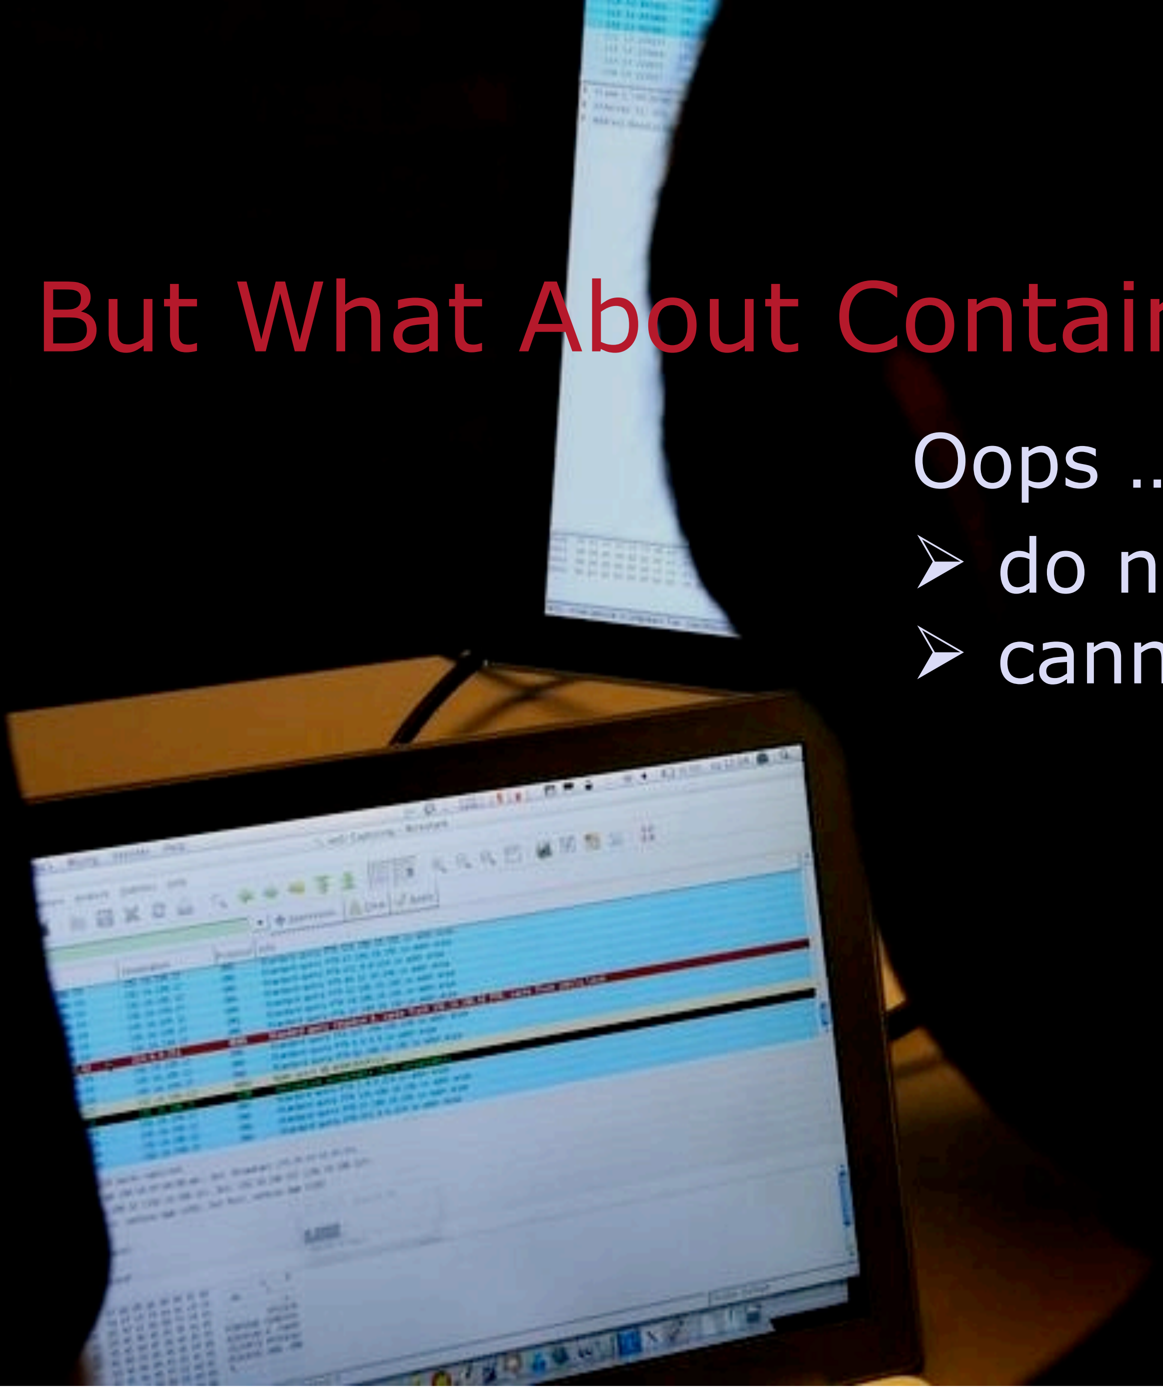

Oops … *ssh* keys

- do not expire
- cannot be revoked

Security And Availability For All Involved

Who are playing in the Grid Space, and thus: who get attacked?

- **Virtual Organisations or Communities: you and your colleagues**
- **Resource Centres and Grid Services:**
  CPU, Storage, Data base and service providers
  central services and coordination

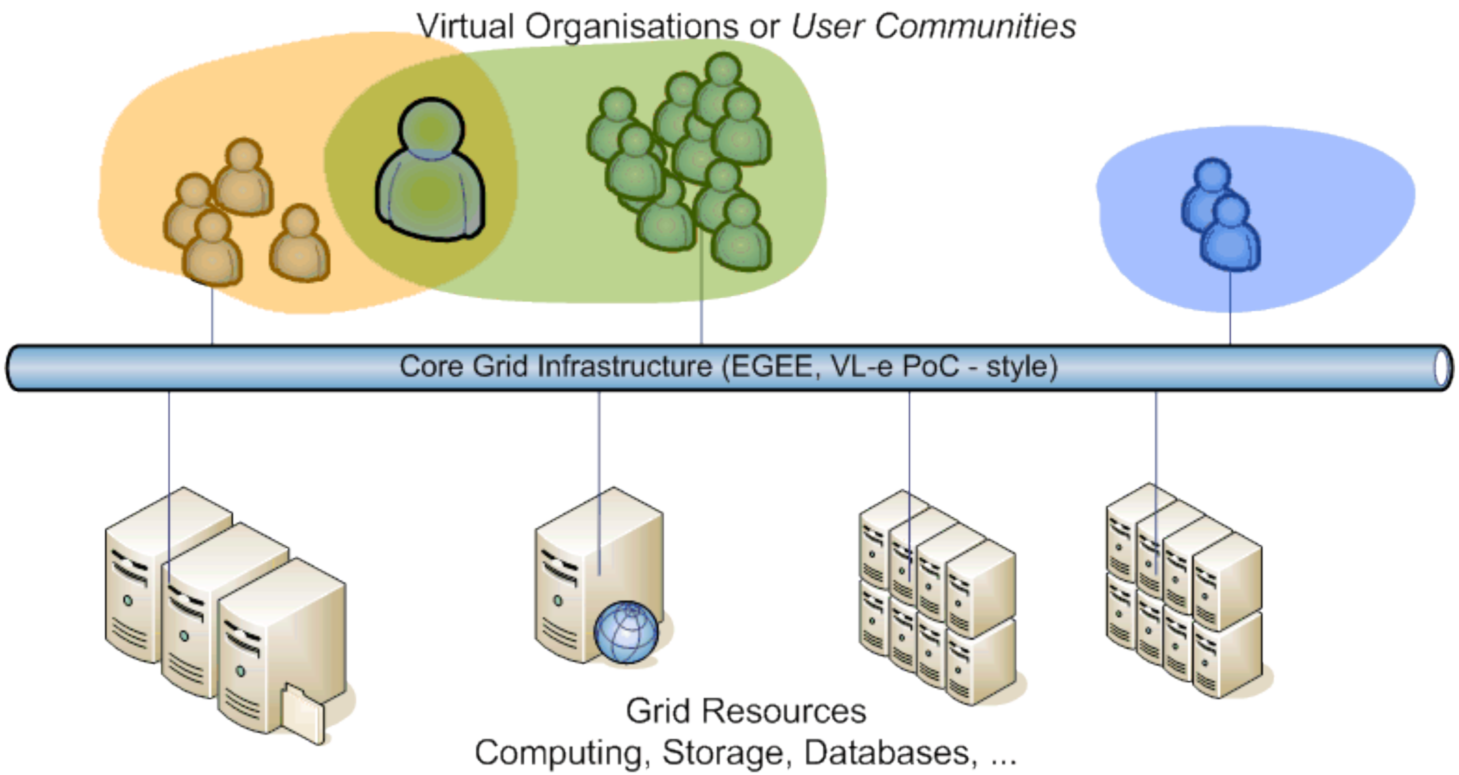
Virtual Organisations or *User Communities*
Core Grid Infrastructure (EGEE, VL-e PoC - style)
Grid Resources
Computing, Storage, Databases, ...

So, where does that leave us?

Is the grid safe? You never know …

- Strong authentication of users and resources by certificates
- Exposure is time-limited and revocable
- Community membership via secured 'visa'
- Encrypted and integrity-protected communications

- Grid and sites subject to policies, with data protection taken seriously,
  commensurate with the open, scientific nature of the infrastructure

- A vulnerability and risk assessment process to work on the software
- Auditing and incident response teams across Europe and the Grids

And you now know more-or-less how this works

But, as always, it remains a matter of *trust* …

Grid Security Middleware mechanisms for protecting the e-Infrastructure

Questions...?

Bonus slides

Hydra key store theory, and SSSS

- Keys are split for security and reliability reasons using Shamir's Secret Sharing Scheme (org.glite.security.ssss)
  - standalone library and CLI
  - modified Hydra service and Hydra client library/CLI
  - the client contacts all services for key registration, retrieval and to change permissions
    - there is no synchronization or transaction coordinator service

```
$ glite-ssss-split-passwd -q 5 3 secret
137c9547aba101ef 6ee7adbbaacac1ef 1256bcc160eda592
   fdabc259cdfbacc9 3113be83f203d794
$ glite-ssss-join-passwd -q 137c9547aba101ef NULL \
  1256bcc160eda592 NULL 3113be83f203d794
secret
```

Slides based on Akos Frohner, EGEE and CERN

Storage layering and interfaces

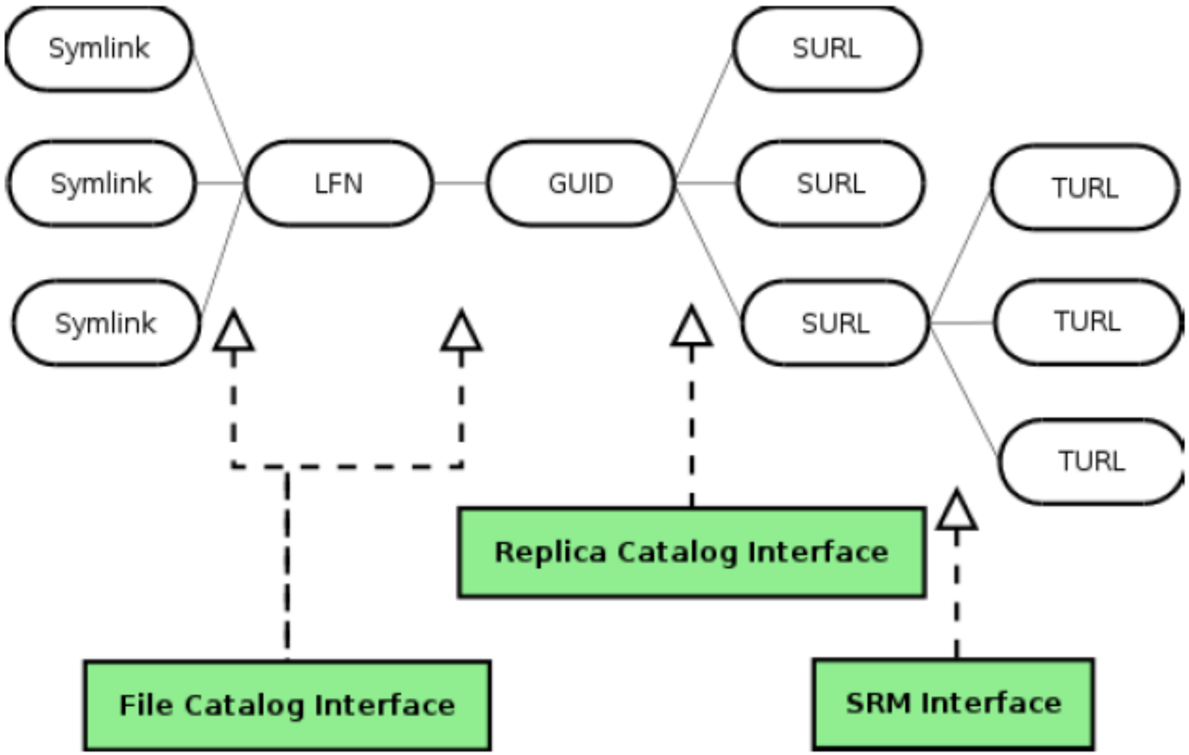

graphic: Peter Kunszt, EGEE DJRA1.4 gLite Architecture

DPM Architecture

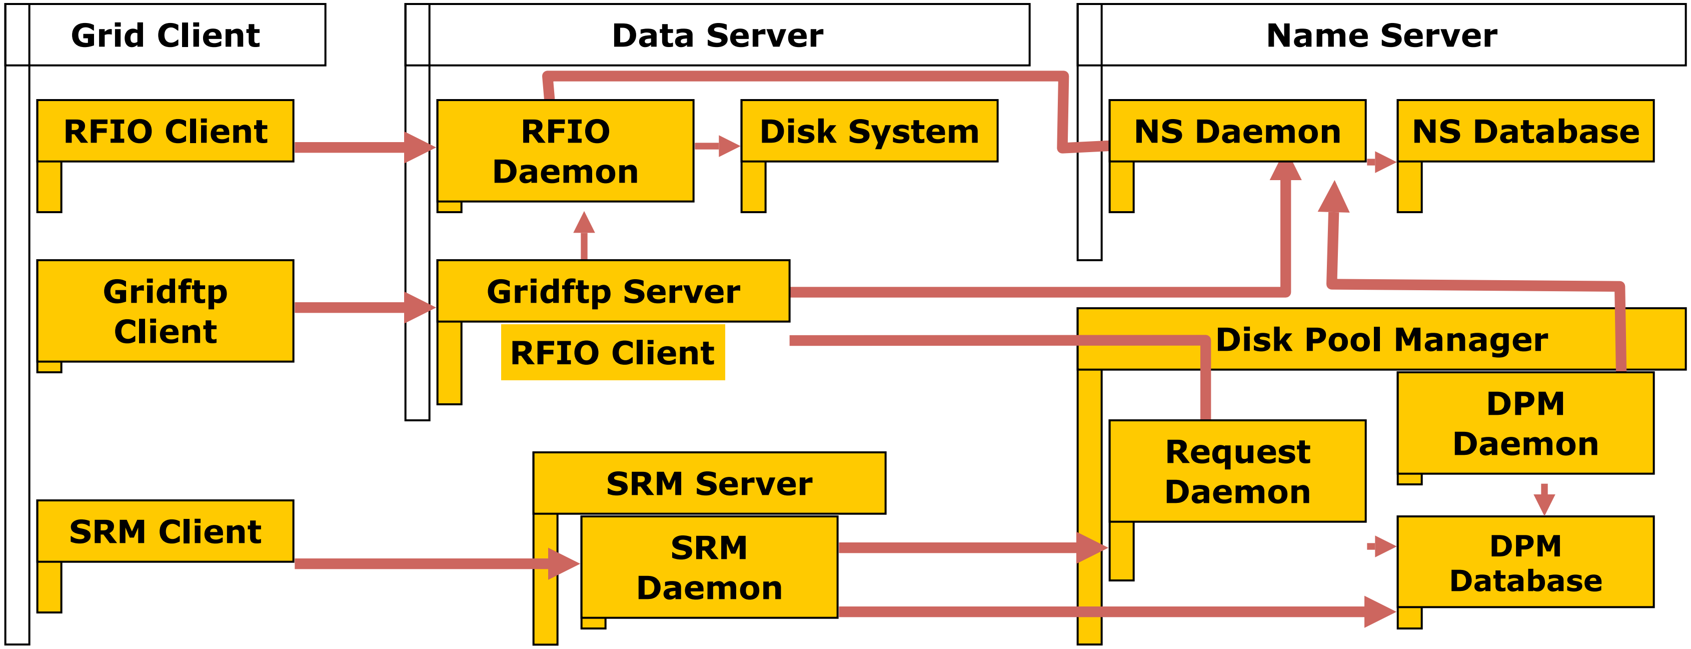

Slides and graphics: ‘ACLs in Light Weight Disk Pool Manager’ MWSG 2006, Jean Philippe Baud, CERN