



Integrating Federations in the International Grid Trust Fabric

> David Groep



Nikhef

Dutch national institute for sub-atomic physics

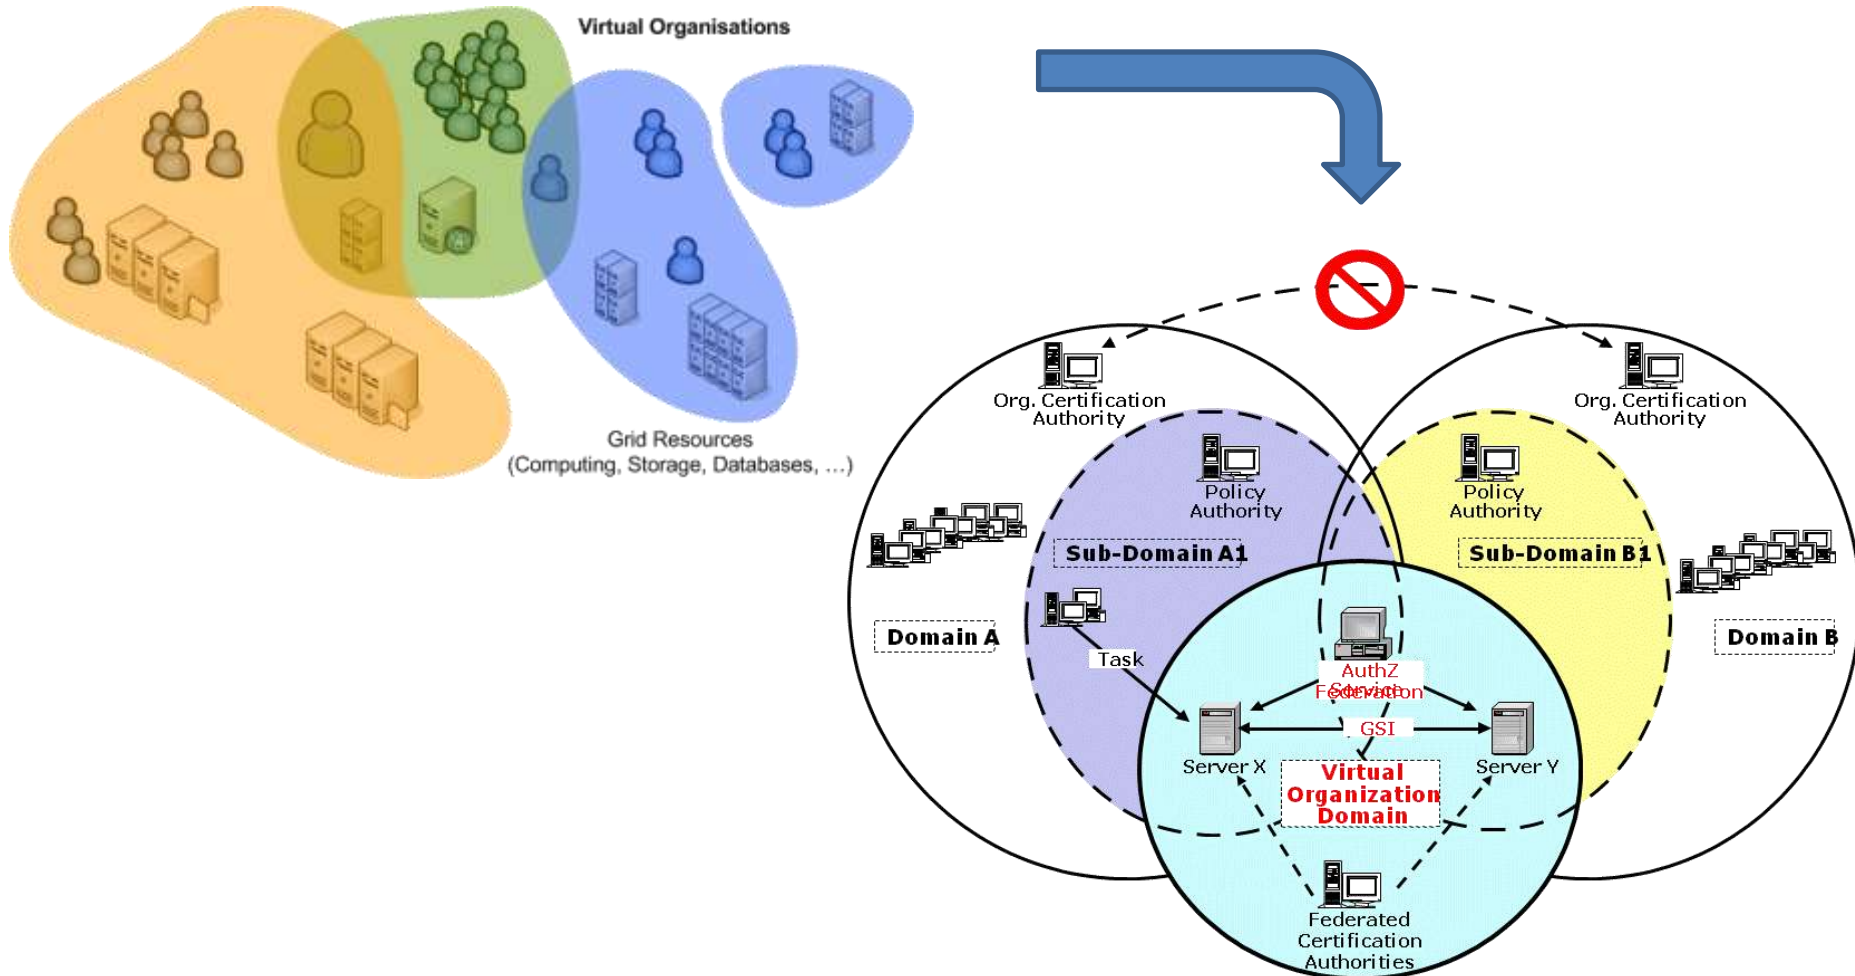
V Grids, Eduroam, Federations

- > Different terms, same issues
 - > How to provide access only the 'proper' people?
 - > Most of whom you've never met and will never meet
 - > Across organisations and countries
 - > Traceable and consistent over many years

- > Core issue is *trust*



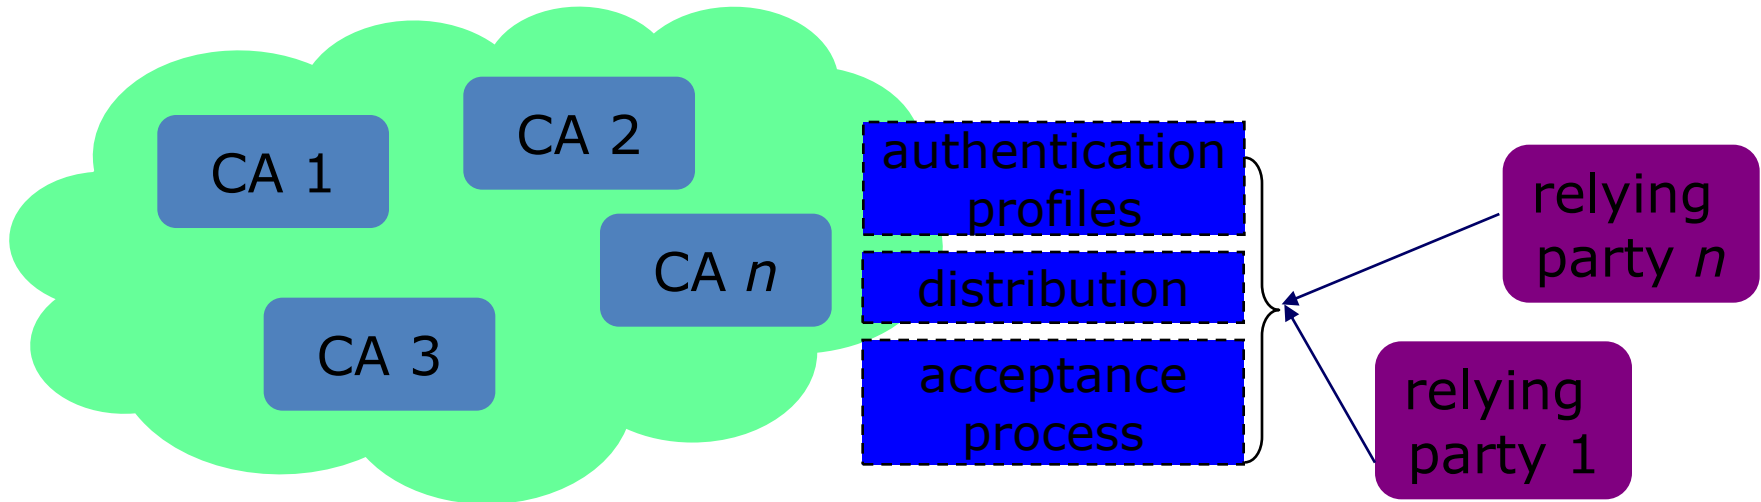
V Grid: Trust in 'Virtual Communities'



V Grid: Where to Root Trust?

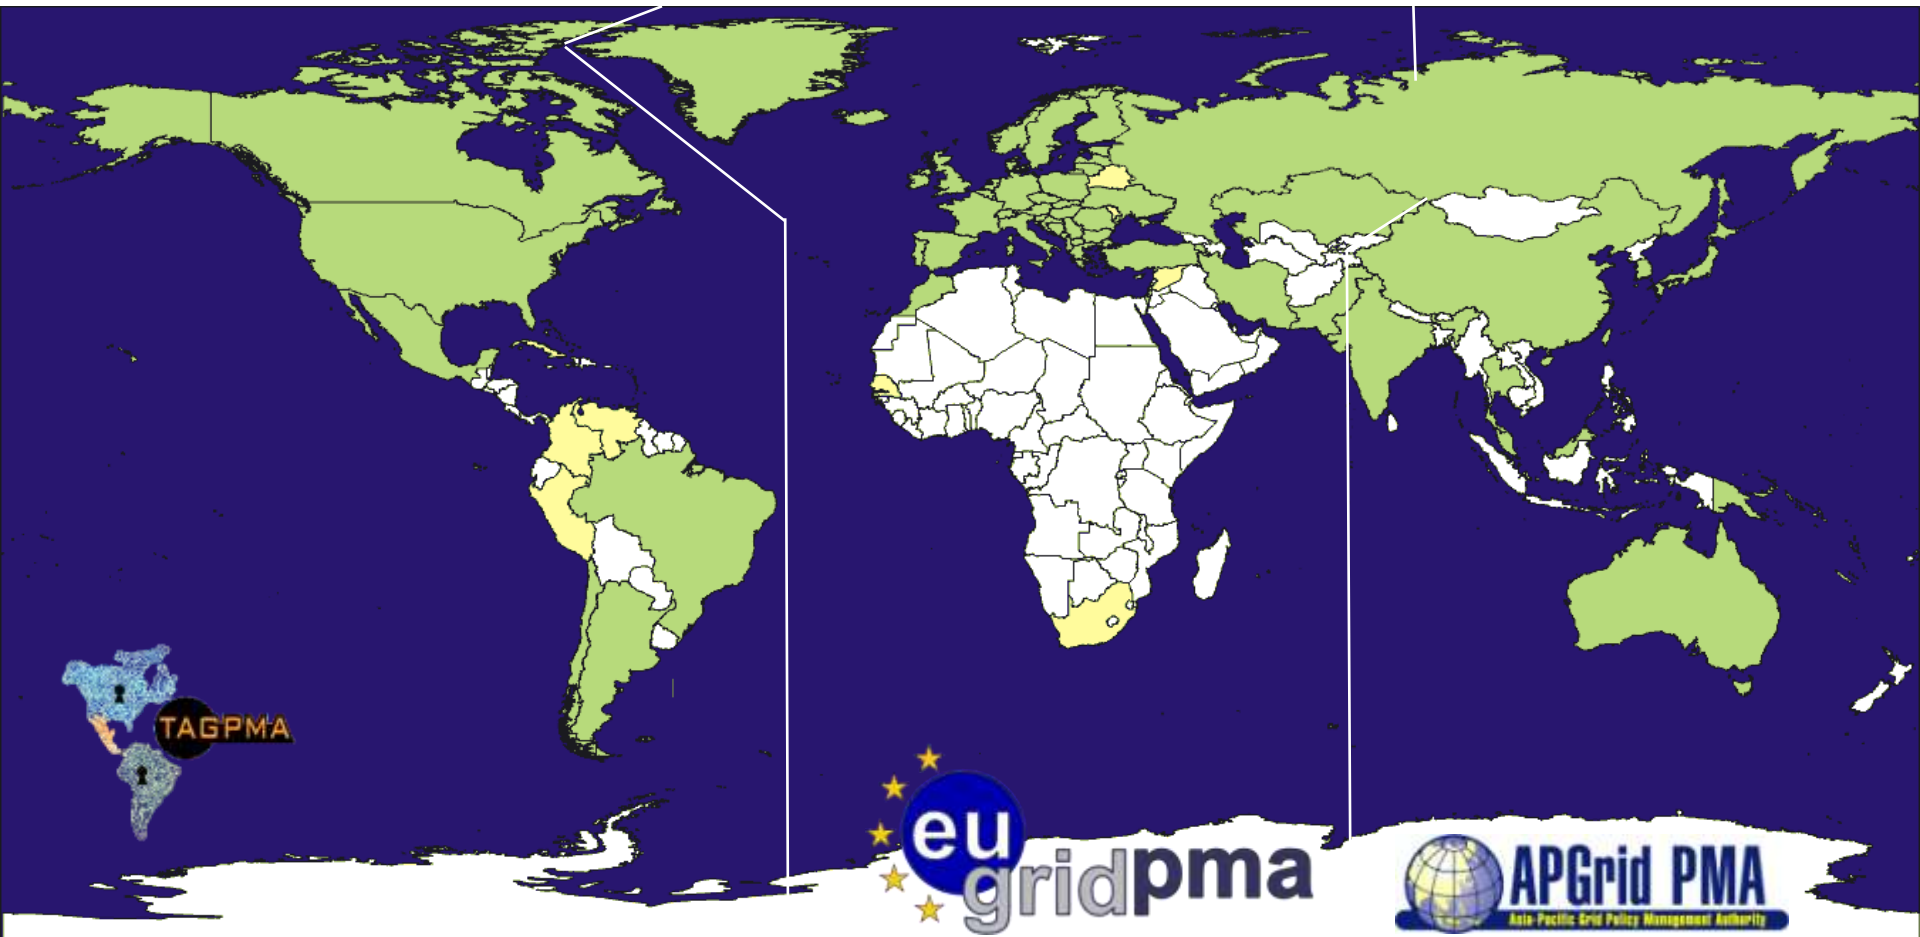
- > There is no *a priori* trust relationship between grid members or member organisations
 - > Community (VO) 'life time' can vary from days to decades
 - > people and resources are members of many VOs
 - > VOs can accept some responsibility, but are a 'bad source' for identity as VO loyalty may be the 'wrong way' (as far as resource providers are concerned)
- > ... but relationship user-VO-resource is required
 - > for authorising access, traceability and liability, incident handling, and accounting
- > ... so the grid needs trusted third parties

V Federation Model for Authentication



- > A Federation of many independent CAs
 - > common **minimum requirements** (in various flavours)
 - > trust domain as inspired by users and relying parties *where relying party is (an assembly of) resource providers*
 - > Using a peer-reviewed acceptance and audit process
- > No strict hierarchy with a single top
 - > spread of reliability, and failure containment (resilience)
 - > maximum leverage of national efforts and complementarities

V International Grid Trust Federation



Green: regular national CA accredited; Yellow: accreditation in progress

V IGTF Status Today

- > Already 77 accredited authorities
 - > EUGridPMA: 41 countries + 1 treaty organisation
 - > TAGPMA: 6 countries (but several authorities in US and BR)
 - > APGridPMA: 7 countries and economic regions (with several authorities in CN, JP, TW)

- > About 10 Major Relying Party members
 - > 6 in Europe: DEISA, EGEE, LCG, OSG, SEE-GRID, TERENA
 - > PRAGMA, TeraGrid, OSG , LCG in other PMAs as well

V Different technologies, different profiles

- > The IGTF established a single trust fabric, incorporating authorities using different techniques

Common Elements

- Unique Subject Naming
- Identifier Association
- Publication & IPR
- Contact and incident response
- Auditability

Profiles

- Classic
 - Real-time vetting (F2F or TTP)
 - 13 months life time
- SLCS
 - Existing IdM databases
 - 100k - 1Ms life time
- MICS
 - IdM Federation
 - 13 months max

<https://www.eugridpma.org/guidelines/>

V Guidelines: common elements in the IGTF

> Coordinated namespace

- > Subject names refer to a unique entity (person, host)
'Any single subject distinguished name (DN) in a certificate must be linked with one and only one entity for the whole lifetime of the service.'
- > Used as a basis (directly or indirectly) for authorization

> Common Naming

- > Single distribution for all PMAs

> Concerns and 'incident' handling

- > Guaranteed point of contact, to raise issues and concerns

> Requirement for documentation of processes

- > PMA-disclosed, detailed policy and practice statement
- > Open to auditing by federation peers

V Guidelines: secured X.509 CAs

- > Aimed at long-lived identity assertions
- > Identity vetting procedures
 - > Based on (national) photo ID's
 - > Face-to-face verification of applicants via a network of Registration Authorities or 'Trusted' Registrars
 - > Periodic renewal (once every year)
 - > Reasonable representation of the person's real name
- > Secured operation
 - > off-line signing key or HSM-backed on-line systems (140.2-3+)
- > Response to incidents
 - > Timely revocation of compromised certificates
 - > CRL issuance required (downloaded up to 400 times/minute!)

V Guidelines: short-lived credential service

- > Issue short-lived credentials based on another authentication system
 - > e.g. Kerberos CA based, or existing (federated) administration
 - > on-line issuing (with 140.2 level 2+ HSM or equivalent)
 - > Based on crypto-data held by the applicant
 - > Maybe only subset of entities in the database
 - > Reasonable representation of the person's real name
- > Same EE cert format, but no new user-held secrets
 - > Can act like a 'proxy' (RFC3820), and has similar risks
 - > The applicant can (and will) use it like a proxy
 - > Risk of EE key exposure is mitigated by its shorter life time
- > Same *common guidelines* apply
 - > reliable identity vetting to ensure uniqueness over CA life time

V Specifics of a SLCS

- > Sufficient information must be recorded and archived such that the association of entity and subject DN can be confirmed at a later date.
- > Qualifying IdMs must suspend or revoke authorization to use the service if the traceability to the person is lost.
- > The CP/CPS must describe:
 1. ...
 2. How it provides DN accountability, showing how they can verify enough identity information to trace back to the physical person for at least one year from the date of certification, and in keeping with audit retention requirements.
- > In the event documented traceability is lost, DN must never be reissued.

V MICS vs SLCS

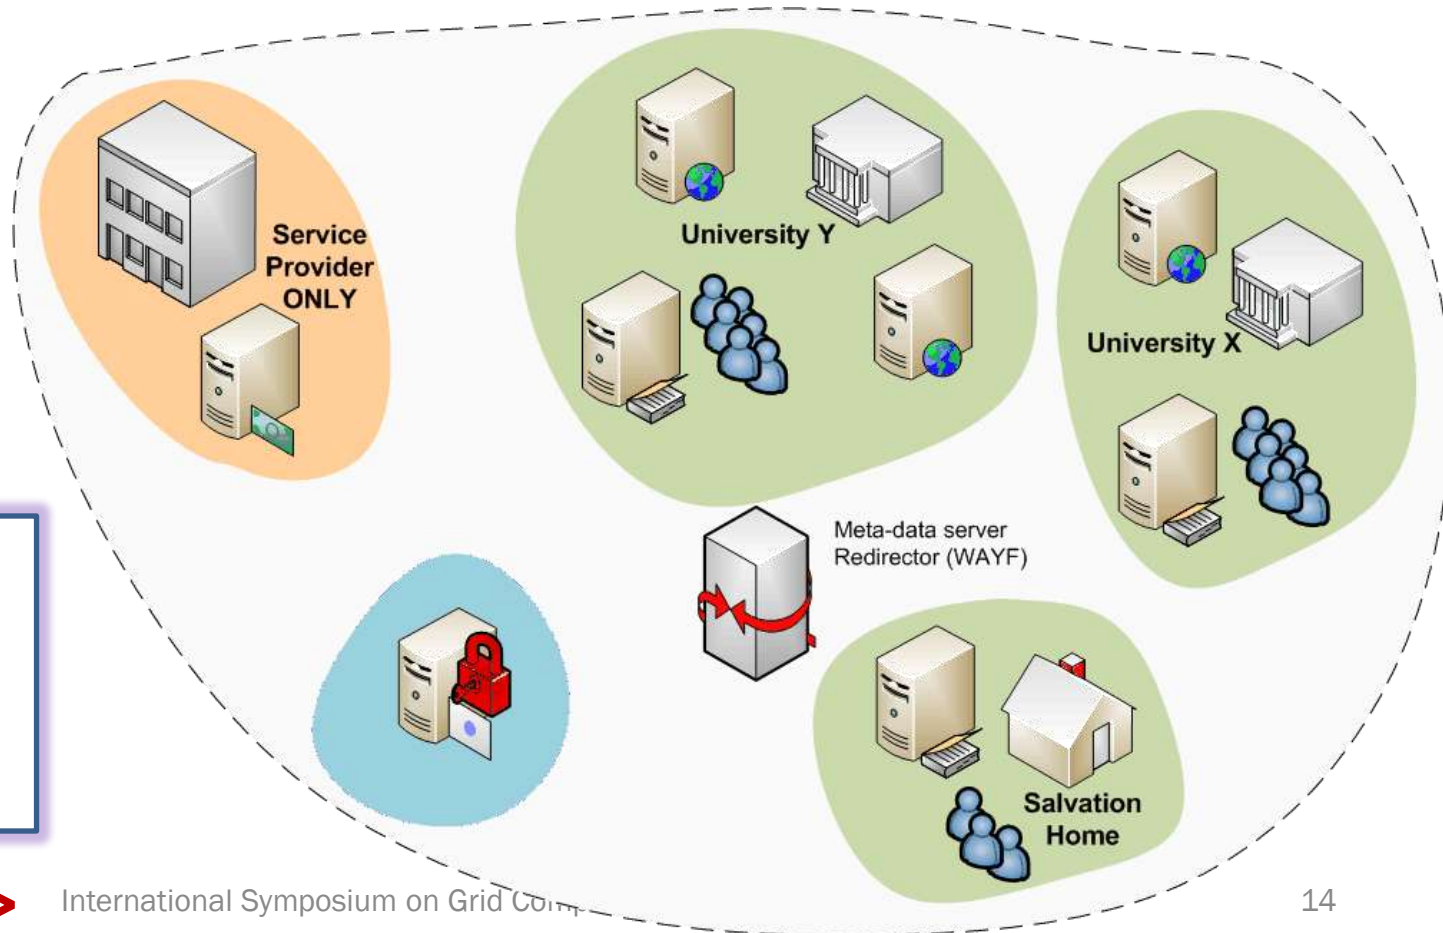
- > 'MICS' is essentially a Classic CA, but with the identity vetting time-shifted & off-loaded to federation or IdM ... which makes it look an awful lot like SLCS, but with better vetting and controls

But then:

- > The life time of the cert can be 13 months
- > To off-set this risk:
 - > The requirements on IdM access and -use are stronger
 - > Extra verification data is requested at issuance time to protect against weak or re-used IdM passwords

V Federations

- > A common Authentication and Authorization Infrastructure
- > Allow access to common resources with a single credential



grid structure was not too much different!

The small diagram shows a central point with arrows pointing to various icons representing Virtual Organisations and Grid Resources (Computing, Storage, Databases, ...).

V Federation basics

> Why Identity Federation?

- > To enable sharing of educational resources
- > Network (Wireless and/or not)
- > Applications
- > Online learning systems (e.g. for the Bologna Process)

> What is needed to set-up an Identity Federation

- > **Agreement on legal framework, policies, and trust** (this is where OpenID &c fail :)
- > Technology
- > Security
- > Common Language and Interoperability

> Identity Federations key element

- > *authentication performed by user home institution*
- > *authZ performed by the service provider***

With thanks to Licia Florio, TERENA

V Birth of Federation & Confederation

- > 2002: a problem to solve
 - > Provide wireless access only to authenticated users
 - > On-line anywhere, anytime
 - > eduroam: the operational, production quality, international confederation since ~2003



- > Solution?
 - > eduroam = education roaming
 - > federated network access for participating institutions
 - > Started in a very simple way ...
eduroam technology: 802.1X + RADIUS

With thanks to Licia Florio, TERENA



Eduroam Participating Countries

A highly successful confederation!

- › Eduroam in EU and APAN
 - › 500+ institutions connected



■ Countries that have joined
■ Countries in the process of joining
■ European Root (TERENA hosts the European Root, which is operated by UNI-C and SURFnet)



■ Countries that have joined
■ Countries in the process of joining
■ Asian Root
>>>EUROPE MAP

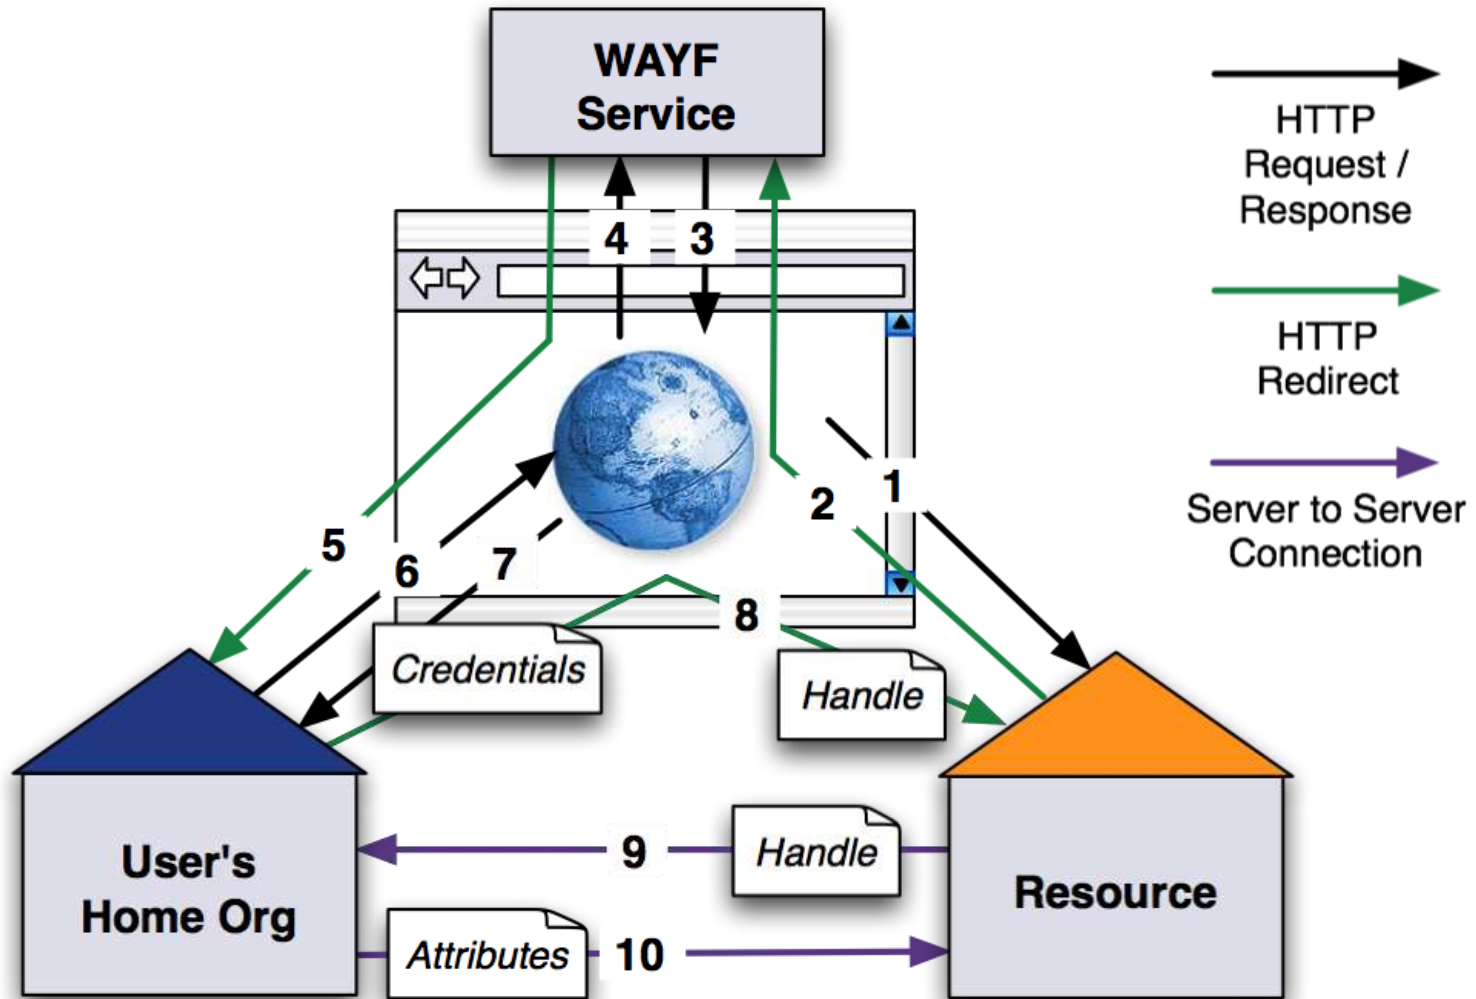
V Beyond the network

- > Research community requirements go beyond network access
 - > Increasing dynamics in the education system
 - > Students can access courses in other faculties
 - > Students take some course units abroad (Bologna)
 - > On-line courses are more common
 - > Users want to access the same services no matter where they are
 - > Grid: example of access to distributed resources

- > More institutions dealing with the same users means:
 - > Multiple registration of users
 - > Overhead to manage guest users
 - > Increased possibility of error in managing the users' records


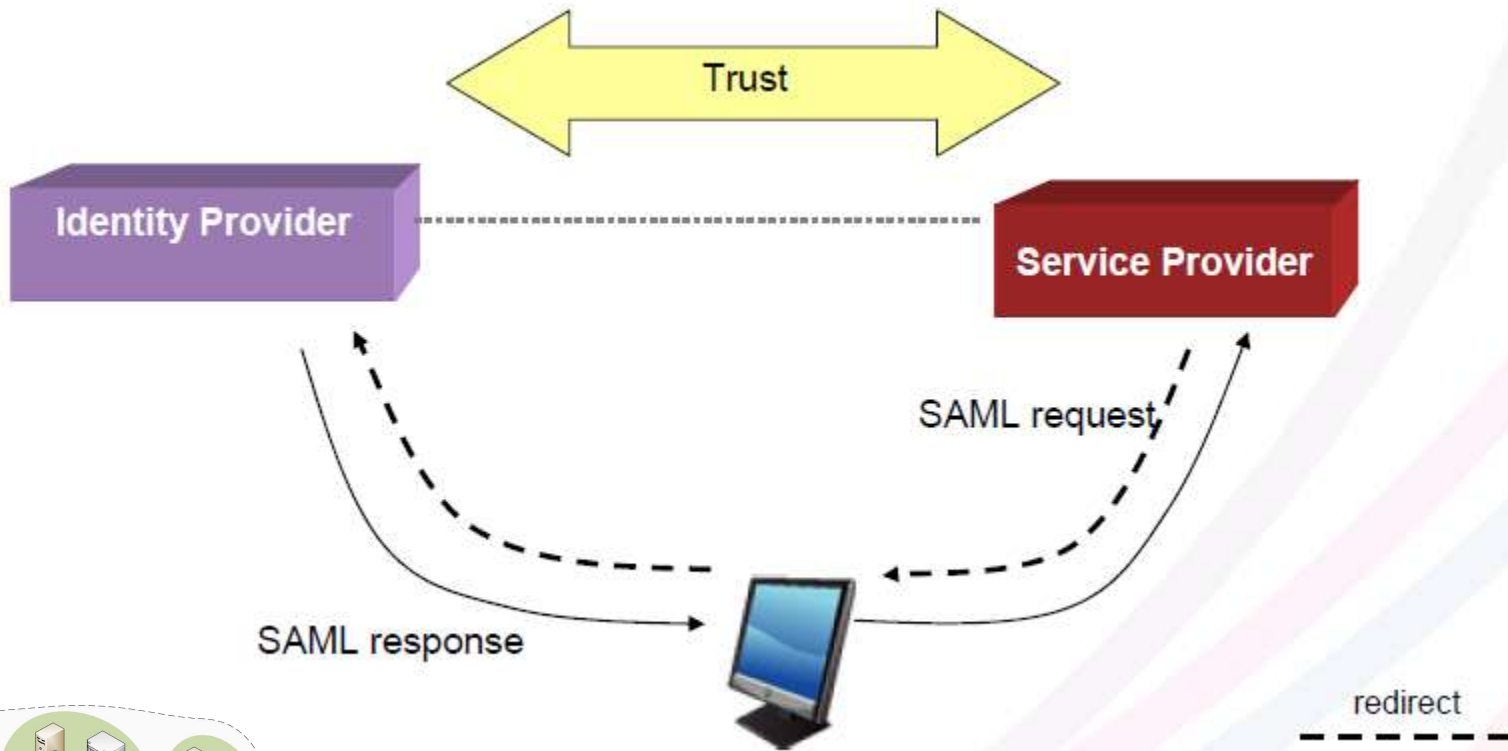
With thanks to Licia Florio, TERENA

Using the federation (example)



Graphic from: Christoph Witzig, SWITCH and EGEE

V Simplified Federation liaisons



With thanks to Licia Florio, TERENA

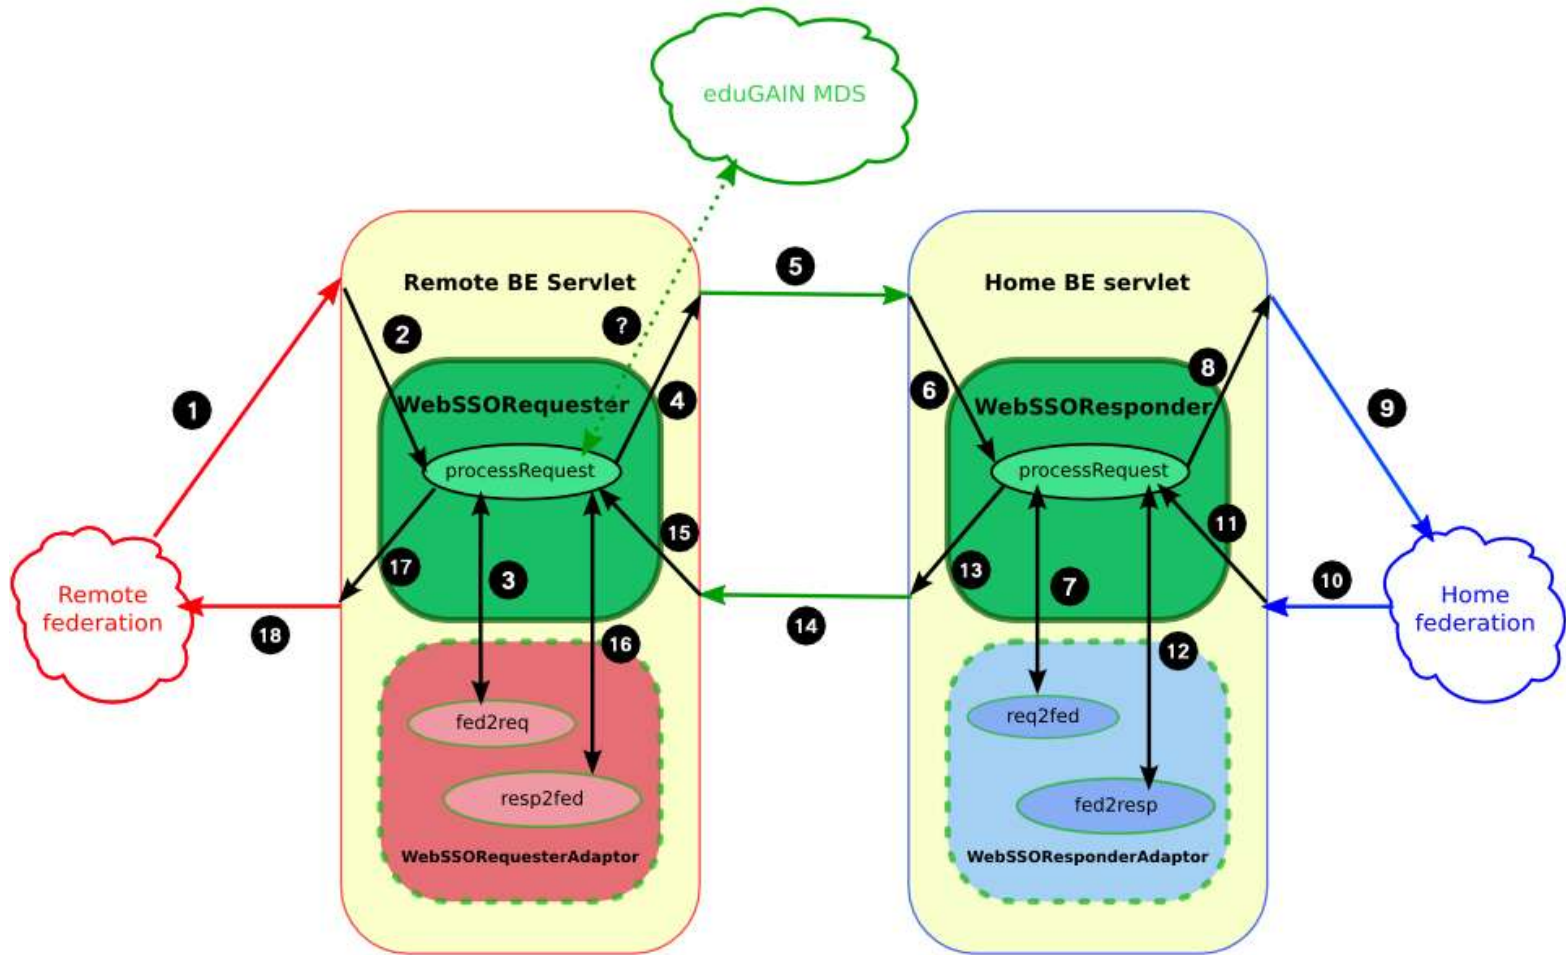
V Key elements in national federations

- > Scalability (to 1-10 M people per country)
 - > Various groups of people: students, walk-ins, staff
 - > Each group has different identity profile and LoA
- > Minimize personnel involvement
 - > Each helpdesk visit is costly
- > Compliance and regulations are an integral part
 - > Universities are far more visible than grids
- > Then, there are easy gains from the federation
 - > Large user base: many people already 'well known'
 - > Pretty good quality ID: paychecks, student IDs, exams

V European Landscape

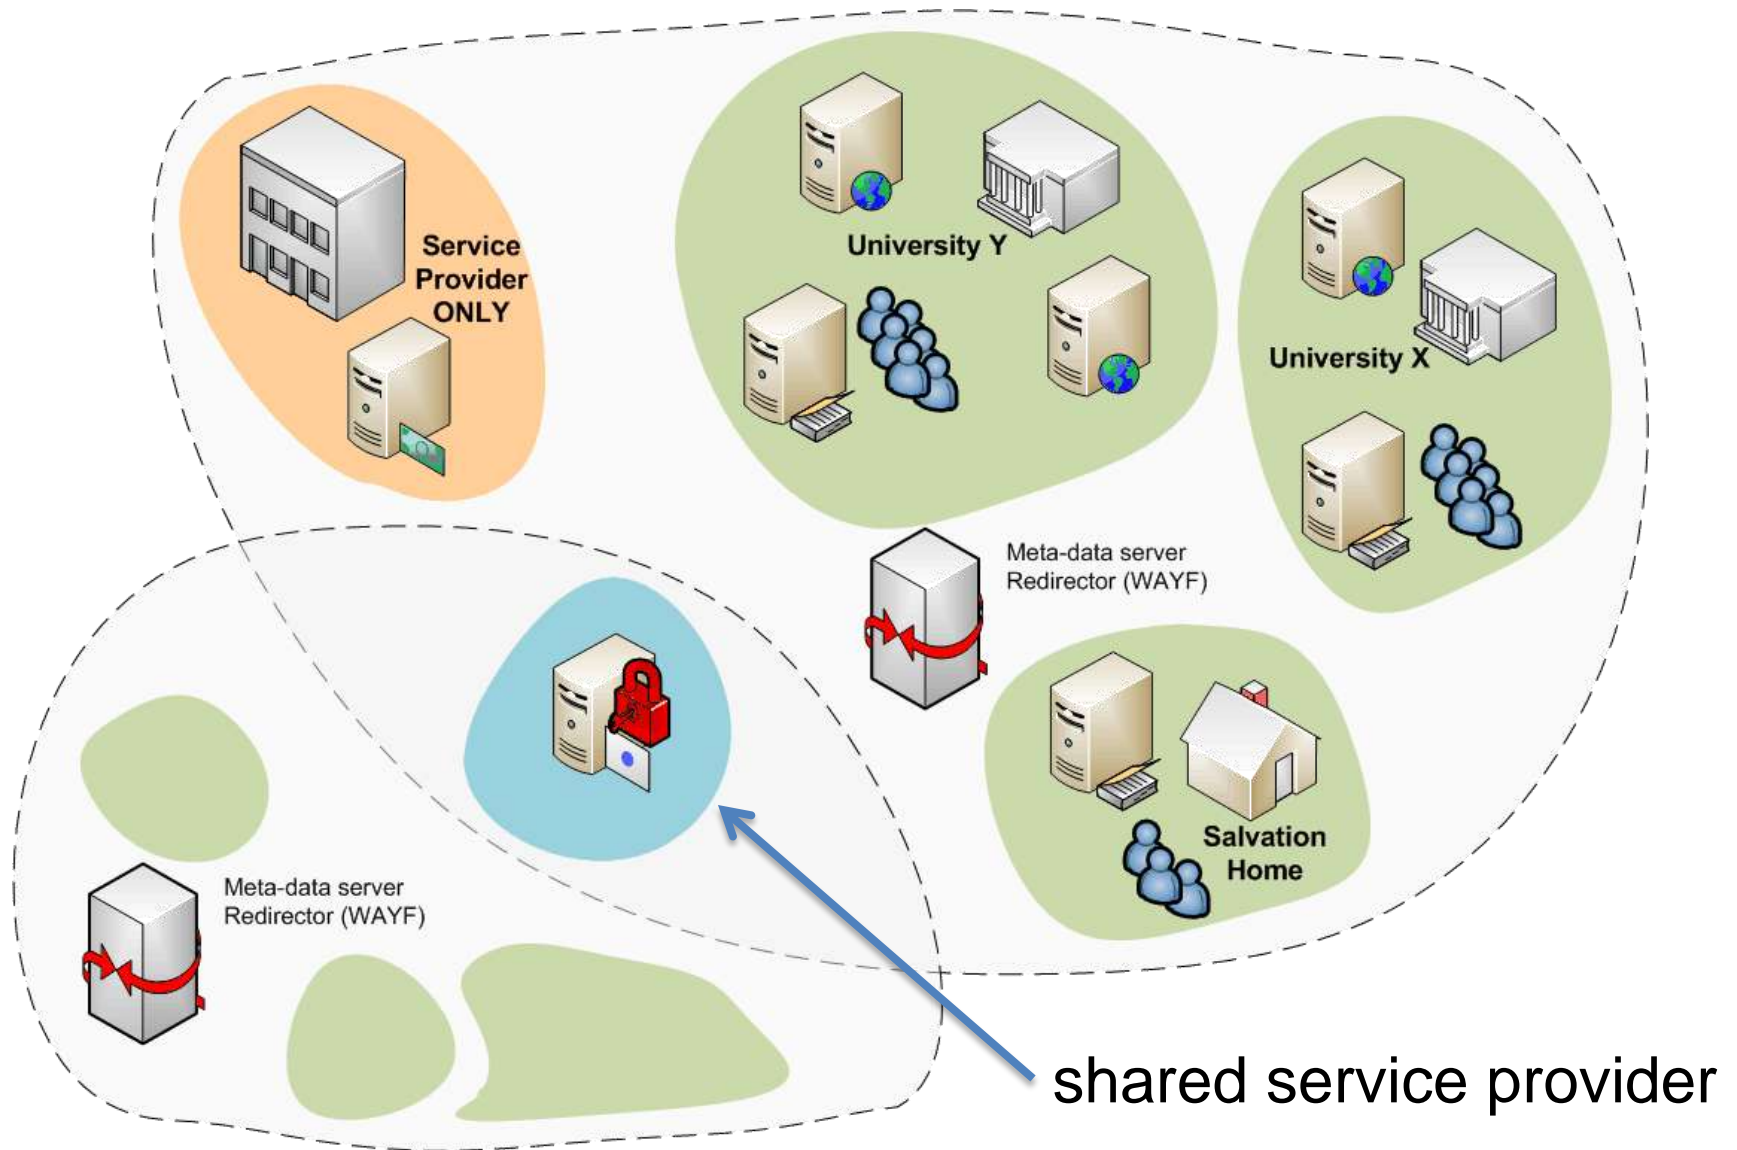
- > Identity Federations (or simply *federations*) are being developed at national level by the NRENs:
 - > Most European countries have one in production, such as SURFederation, FEIDE, WAYF, Swamid, UKAF, PAPI, HAKA, DFN, SWITCHaai, etc. – some as NREN function, others as a separate org.
- > Different (open source) technologies are used
 - > PAPI: Spain
 - > A-Select (Shib, PKI, ADFS, sSPHP,...): the Netherlands
 - > Shibboleth: UK, Finland, Switzerland, Germany
Well used technology, but certainly not the only one!
 - > Sun Federation Manager based upon Liberty Alliance spec: Norway
- > Interoperation through use of SAML (2.0)

V Confederation, the eduGAIN model



Graphic from: Diego Lopez, RedIRIS and GEANT2

V SP 'multi-federation'

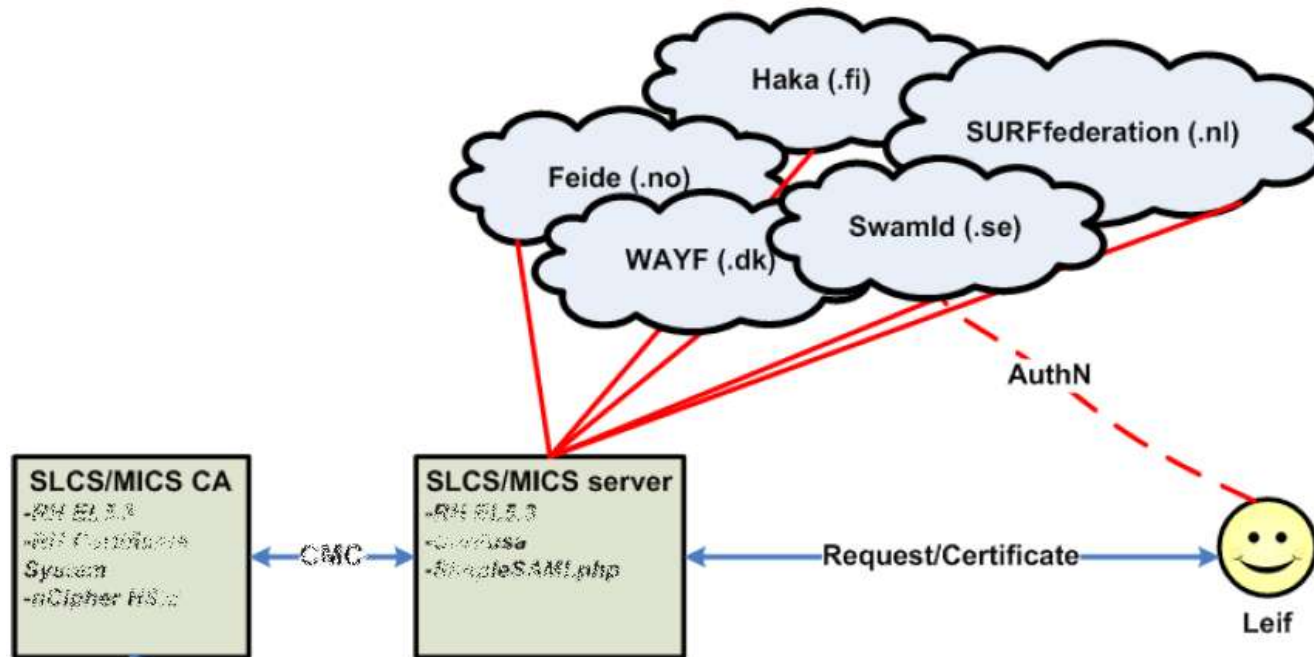


V Grid & Federations - a logical combination

- > Many of the e-Infrastructure users have a federated account anyway
 - > Could give users grid certificates within 5 minutes without any further hassle, if the trust level matches
 - > Allows users to 'try' the grid
 - > Allows for scaling the number of grid users significantly
- > Comparable issues and trust levels
 - > Federation assurance levels are going up, as more diverse services participate in the federation
 - > User account management of IdPs is improving rapidly ... far more than grid credential management!

✓ A Federated Grid CA

- > Use your federation ID
- > ... to authenticate to a service
- > ... that issues a certificate
- > ... recognised by the Grid today



Graphic from:
Jan Meijer, UNINETT

V Constraints on a Federated Grid CA

- > We are the first service to ask for a specific, higher, LoA
 - > LoA was slow in catching on
 - > various services in the federation have diverse value
 - > but waiting for it can take years – need specific agreements
- > IdPs and services in a federation loosely coupled
 - > How should loss of affiliation or expiry affect the CA and the issued certs
- > Pushing requirements on IdPs for just one SP requires effort
 - > needed to overcome LoA issue, but should be minimal
 - > Contract types and policy convergence required
 - > Certainly possible for ‘interesting’ services, and done already
- > Any human interaction (helpdesk) is expensive

✓ Matching the ‘easy’ requirements

- Persistent and unique naming
 - > IdPs historically tended to recycle login names
 - > even eduPersonPrincipalName is often recycled
 - > only eduPersonTargetedID is immune to thus, but not supported everywhere (and is usually opaque)
 - > *this adds a requirement to the federation or to the IdPs*
- Reasonable representation of names
 - > Given name, surname and nickname are usually considered privacy sensitive
 - > user-approved release of these appears doable
 - > requires evaluation of legal framework

✓ Matching the ‘harder’ requirements

- > How to protect against re-use of federation login and password (on e.g. Wiki’s, web sites, ISP mail)?
 - > The ‘second attribute’ is either privacy sensitive, or not standard across all IdPs, or difficult to collect
 - > IdPs improving credential management, as compensatory measure
- > Handling revocation
 - > The CA is ‘just a service provider’, and the IdP does not know about the issued certs per se
 - > Even single sign-out is already difficult
- > Assurance level at the IdP
 - > The CA SP is usually the first ‘high-LoA’ application

▼ Where to put the onus of the service?

Homogeneous federation

- > No service should impose requirements on IdPs, only on the federation
- > Federation should do everything internal to it
- > Don't ever bother IdPs
- > In the short term, only minimal assurance given
- > good in many cases, but maybe not here

Service based

- > Service can have contracts with specific IdPs (even for a subset of their users)
- > Federation *mediates* authN and attributes
- > Federation monitors IdPs for compliance in attribute exchange and authN
- > Today used for content licensed to specific IdPs

✓ Federated CAs: there already!

> SWITCH

- > Accredited 2007 under SLCS
- > Shib-only federation, dedicated software development
- > Leverages tightly-controlled SWITCHaai federation

> TERENA Grid CA (formerly NetherNordic SLCS/MICS)

- > Under accreditation today
- > Multi-technology federation, SAML2 based
- > Heterogeneous federations, with web access being the only common element

> CESNET

- > Same time line as TERENA or earlier

✓ SWITCH SLCS CA, since May 2007

- All IdP use comparable registration process
- SLCS eligibility added by explicit human interaction



enabling Grids for e-science

CP/CPS (3)

Initial Identify Validation


Because registration processes vary in details, we intend not to formulate a procedure, but requirements that the IdP must fulfill to enable access to the SLCS

1. Requester must contact RA
How: TBD by RA
2. RA enables access if
 1. User has valid AAI account
 2. a set of conditions are fulfilled (see next slide)
3. Requester can access SLCS
(after successful AAI log-on)

EGEE-II INFISO-RI-031688

SWITCH@elec: EUGRIDPMA Oct 5, 2006

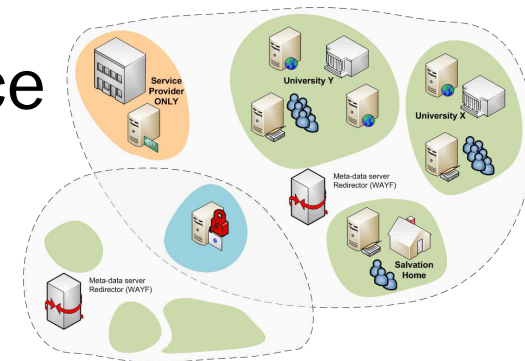
24



✓ New: TERENA Grid CA Service

- Initial partners:
FEIDE, SURFfederatie, HAKA, WAYF, Swamid, TERENA

- Trans-national, cross-federation service
 - > But not (yet) confederated



- How many SLCS/MICS CAs does Europe need ?
 - > Consolidate operational PKI skills in one place
 - > Better sustainability, in line with the European trend

V TERENA Grid CA Pilot principles

- > Leverage federations as much as possible
 - > Define criteria, not process
 - > IdPs with good ID vetting can enable eligibility *en-block*
- > Use only open interconnect (SAML2)
 - > Not a shib-only solution
- > Scalable contract model
 - > *Would like federations to join, but this runs into trouble for LoA and vetting criteria*
 - > IdPs can sign up with the service, through TERENA, if and only if they are members of their national federation
 - > Compliance ensured through the federation contracts

V Implementation

SWITCH-SLCS

- > Integrated client tool
- > Mimics shib exchange and relies on web-redirects,
- > works in pure-shib federation
- > Single command for users

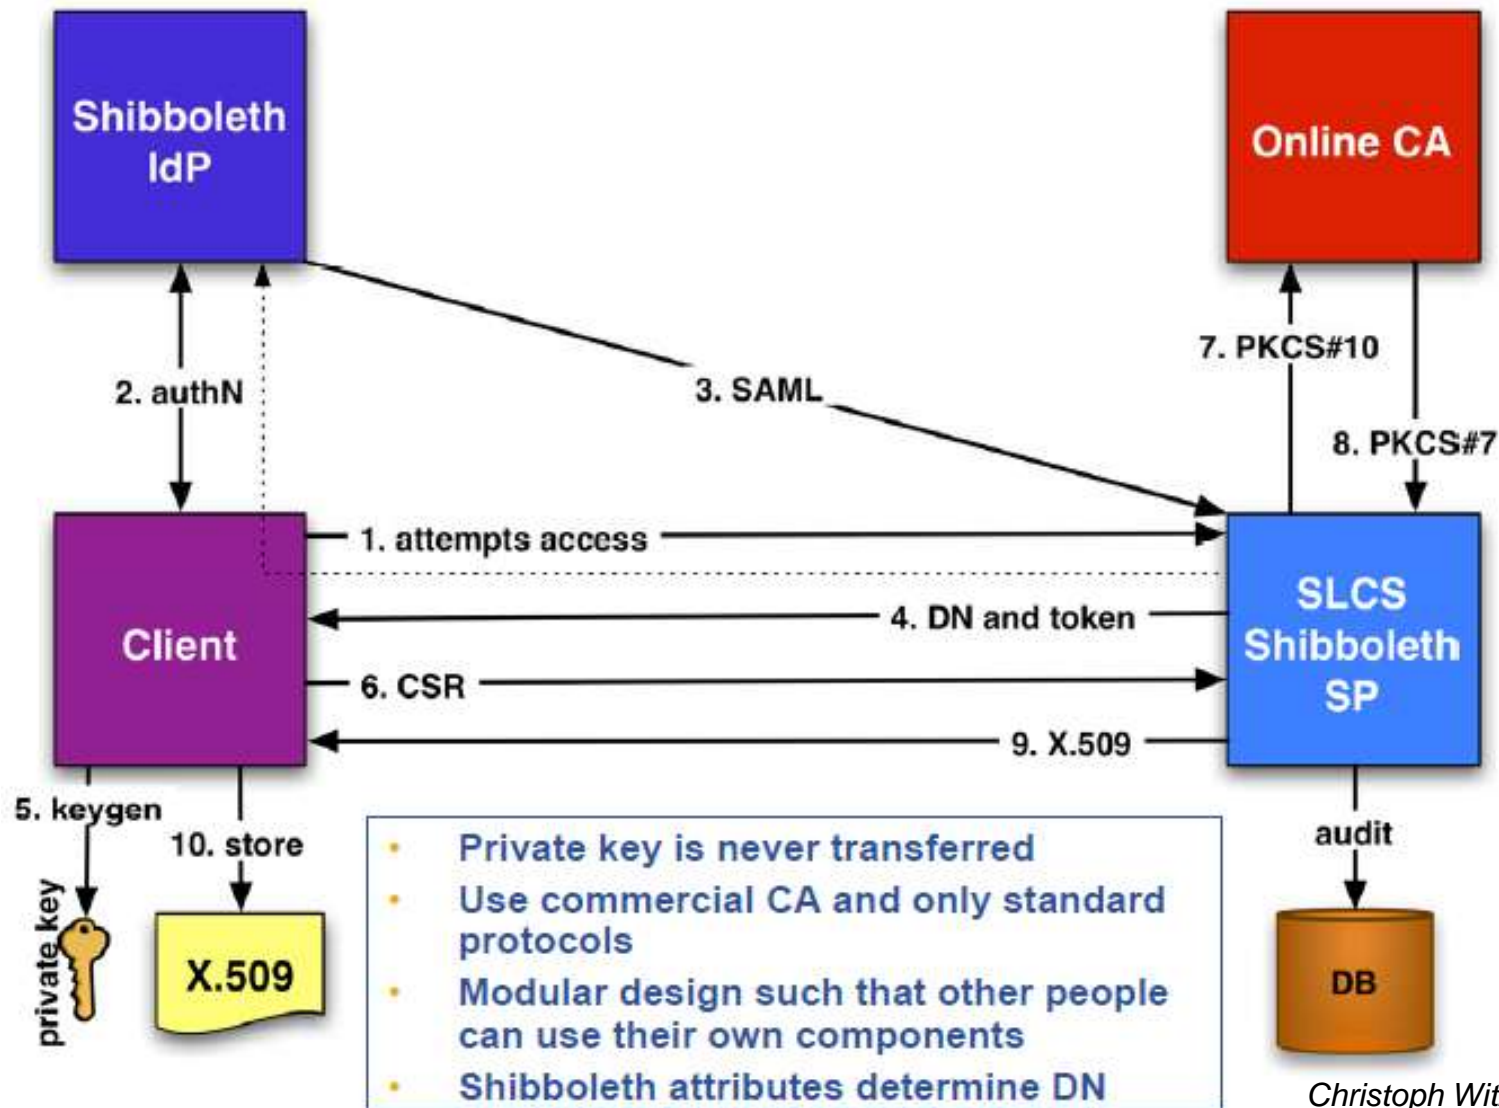
SWITCH 'glite-slcs'

TERENA Grid CA

- > mini-client supported by web browsers
- > actual authN happens through the web only – the only portable way ☹
- > compliant with all federations
- > two-step process for users
- > works with any credential (incl. e-banking, SMS, PKI, ...)

confusa.org

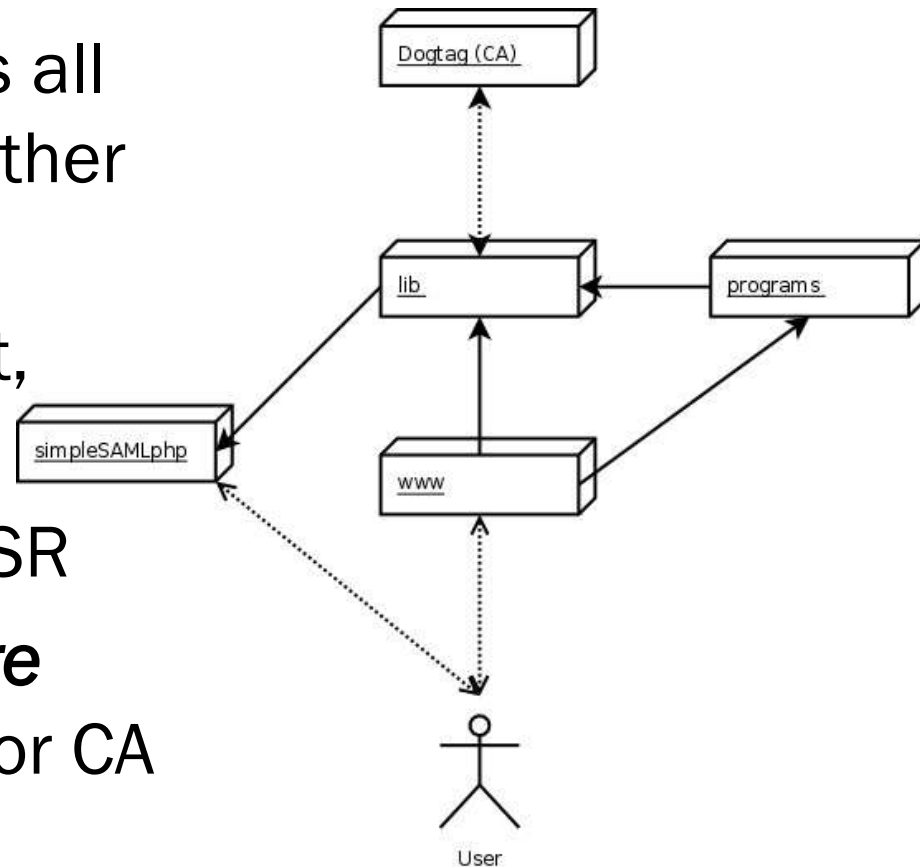
SWITCHaai SLCS CA model



Graphic from:
Christoph Witzig, SWITCH and EGEE

V Confusa model

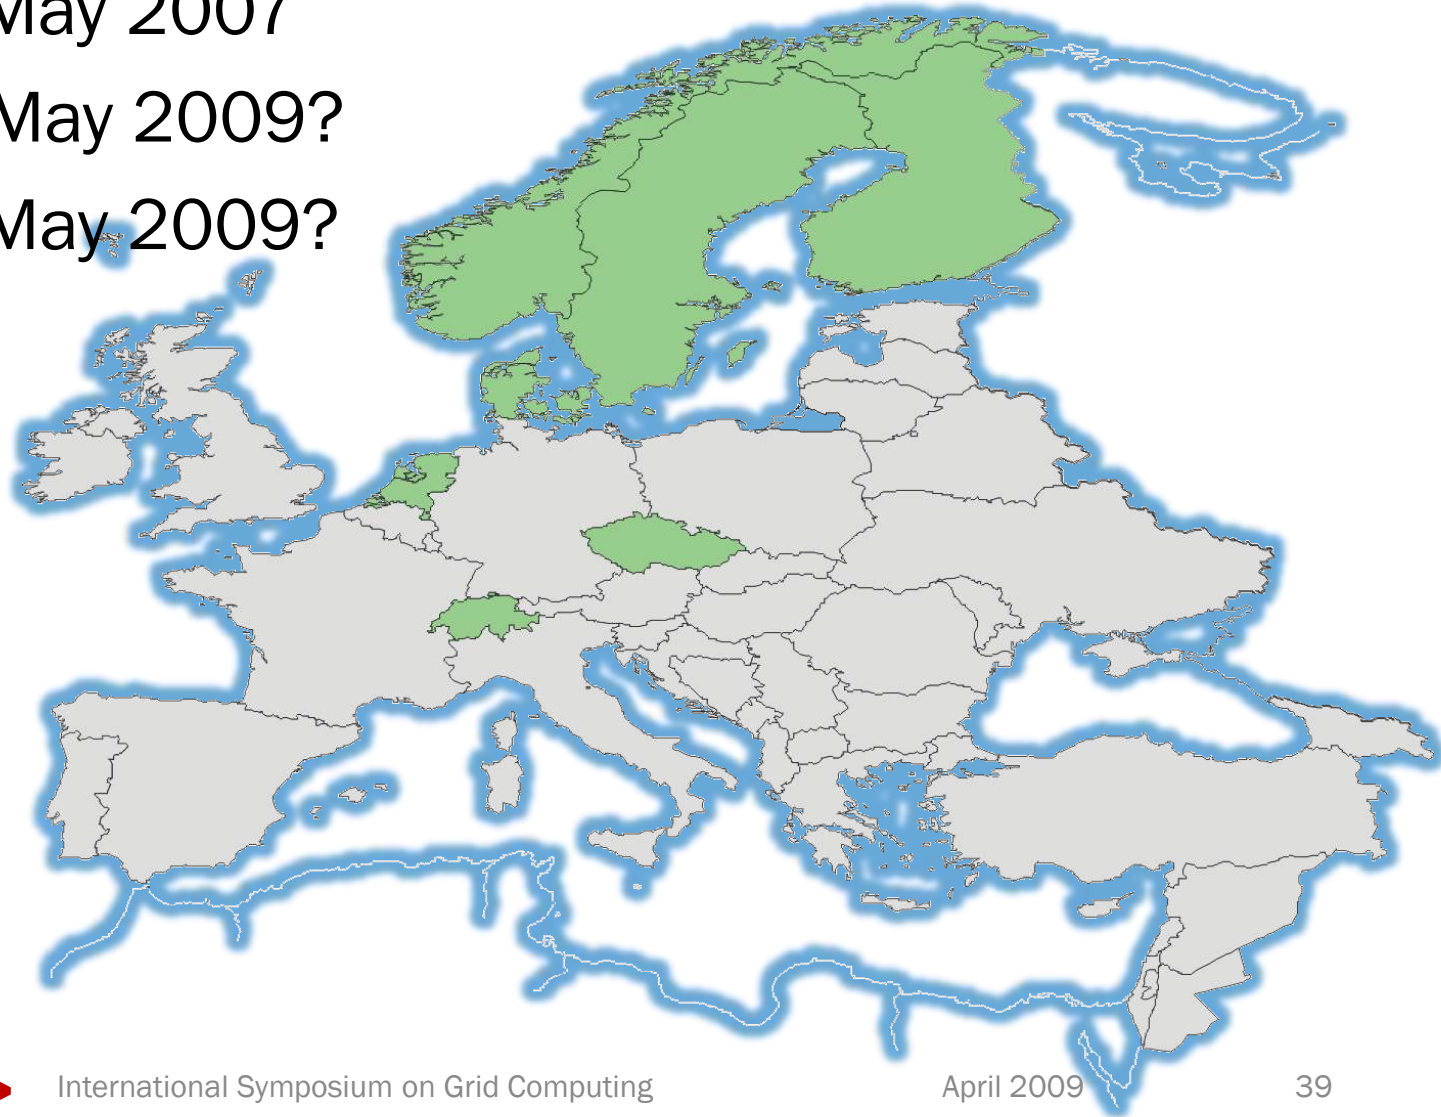
- > Create request out of band
- > SimpleSAMLphp handles all authN, integration with other Federations
- > Browser handles redirect, SAML-header magic
- > Client creates key and CSR
- > **We verify AuthN, compare attribs and create CMC for CA to sign**



Graphic from:
Hendrik Austad, UNINETT Sigma

✓ Federated CAs in Europe

- ✓ SWITCH: May 2007
- > TERENA: May 2009?
- > CESNET: May 2009?



V Joint vision and wide support

2004 policy by EUGridPMA and TERENA, approved by e-IRG:

- > The e-IRG notes the timely operation of the EUGridPMA in conjunction with the TACAR CA Repository and it expresses its satisfaction for a European initiative that serves e-Science Grid projects. [...] The e-IRG strongly encourages the EUGridPMA / TACAR to continue their valuable work [...]
- > *The e-IRG encourages work towards a common federation for academia and research institutes that ensures mutual recognition of the strength and validity of their authorization assertions.*

e-IRG Recommendations, 2004 and 2005, resp.