Portals and Authentication

Issues and Solution Directions
from a CA and IGTF Perspective

David Groep
NIKHEF

Issues

- **Authentication**
  - a federated CA structure
  - Identity vetting, policies, requirements, and relying parties
  - Certificate ‘classes’ and their assurance
- **Authentication and Portals**
  - automated clients
  - user credential caches
  - AAI-backed Short-Lived Credential Service CAs

Federation Model for Grid Authentication

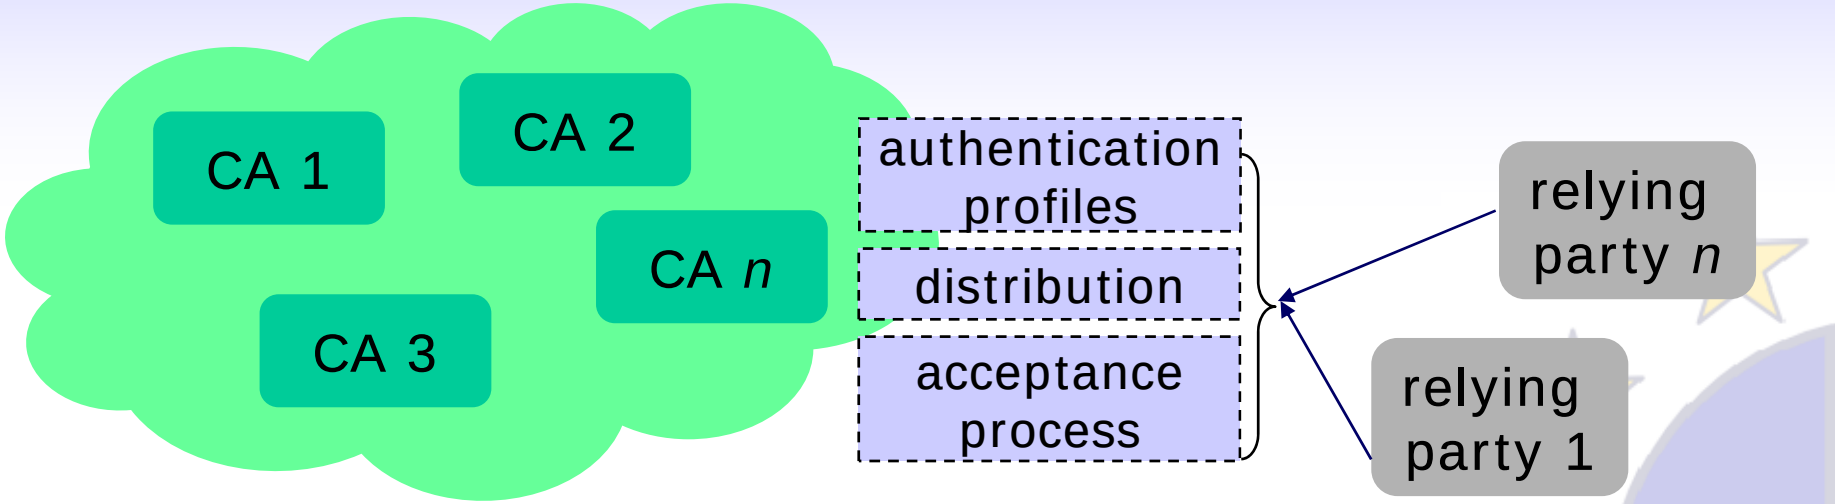

- A Federation of many independent CAs
  - common **minimum requirements** (in various flavours)
  - trust domain as required by users and relying parties *where relying party is (an assembly of) resource providers*
  - defined and peer-reviewed acceptance process
- No strict hierarchy with a single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

International Grid Trust Federation

IGTF
International Grid Trust Federation
AP | EU | TAG

Federation of 3 Regional "PMAs", that define common guidelines and accredit credential-issuing authorities

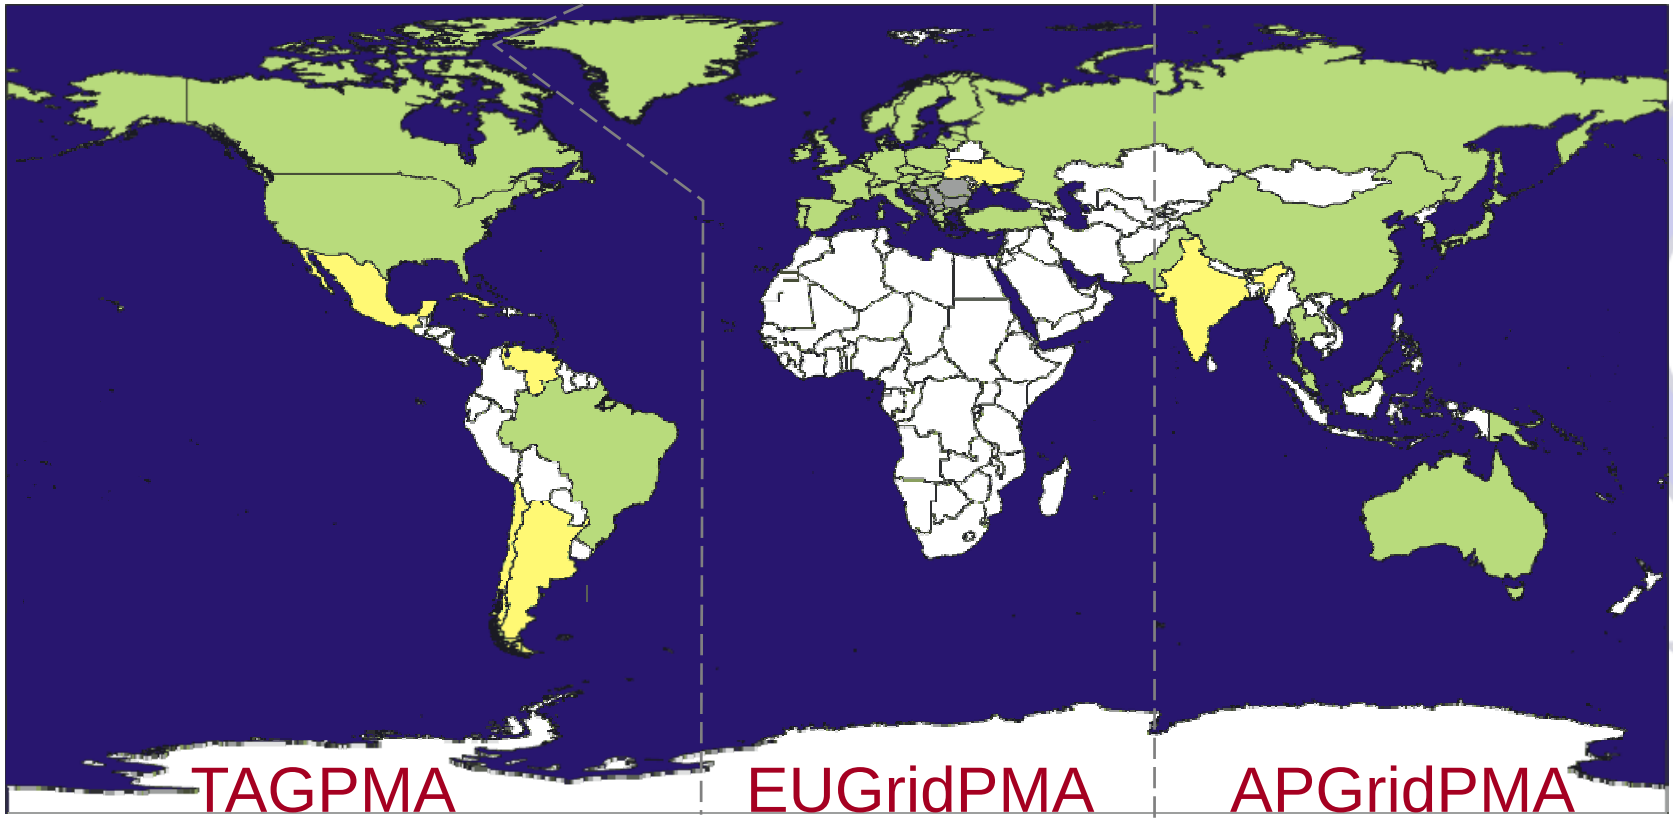

Relying Party issues to be addressed

Common Relying Party requests on the Authorities

1. standard accreditation profiles sufficient to assure **approximate parity**

   effectively, a single level of assurance sufficed then for relying parties – is changing today, as more diverse resources are being incorporated

2. monitor [] signing namespaces for **name overlaps**
3. a **forum** [to] participate and raise issues
4. [operation of] a **secure collection point** for information about CAs which you accredit
5. **common practices** where possible

6. **reasonable likeness** for a subject's name*
7. a subject's name should be forever **persistent***

list courtesy of the Open Science Grid (and wLCG and EGEE draft policy)*

Grid Relying Parties & resource providers

- In Europe
  - Enabling Grid for E-sciencE (EGEE) (~ 200 sites)
  - Distr. Eur. Infrastructure for Supercomputer Apps (DEISA) (~15 sites)
  - South Eastern Europe: SEE-GRID (10 countries)
  - *many national projects (NL BIG-GRID, VL-e, UK e-Science, Grid.IT, …)*
- In the Americas
  - EELA: E-infrastructure Europe and Latin America (24 partners)
  - WestGrid (6 sites), GridCanada, …
  - Open Science Grid (OSG) (~ 60 sites)
  - TeraGrid (~ 10 sites + many users)
- In the Asia-Pacific
  - AP Grid (~10 countries and regions participating, and growing)
  - Pacific Rim Applications and Grid Middleware Assembly (~15 sites)

data as per mid 2006

Building the federation

- Trust providers ('CAs') and relying parties ('sites') together shape the common requirements
  - Several *profiles* for different identity management models
  - Authorities demonstrate compliance with profile guidelines
  - Peer-review process within the federation to (re-) evaluate members on entry & periodically
  - reduces effort on the relying parties
    - single document to review and assess for all CAs under a profile
  - reduces cost for the authorities
    - but participation does come at a cost of involved participation …
- Ultimate trust decision *always* remains with the RP
- An authority is not necessarily limited to just 'grid' use

Evolution and revolution

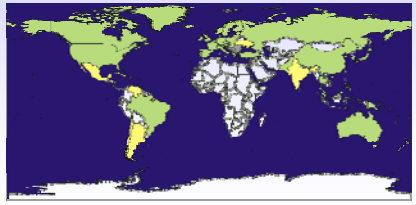

- Each CA is independent
  - constraints of manpower, local funding, national legislation &c
  - compliance is to minimum requirements

- Introduction of new features
  - through demand from within the subscriber base (per CA)
    most effective, especially if you bring along effort
  - through cross-fertilisation by peer CAs
    also effective, but can take a lot of time if effort is lacking
  - by raising the minimum requirements
    does not work well for innovations …

Certificate Assurance and LoA

Guidelines: secured X.509 CAs

Aimed at long-lived identity assertions, the 'traditional PKI' world

- Identity vetting procedures
  - Based on (national) photo ID's
  - Face-to-face verification of applicants via a network of distributed Registration Authorities
  - Periodic renewal (once every year)
  - revocation and CRL issuing required *and we have all RPs actually downloading the CRLs several times a day*
  - subject naming must be a reasonable representation of the entity name
- Secure operations
  - off-line signing key *or* HSM-backed on-line secured systems
  - data retention and audit trail requirements, traceability of certified entities
- Technical implementation
  - need to limit the number of issuing authorities for technical reasons *(most software and browsers cannot support $\mathcal{O}(1000)$ issuers)*
  - certificate profile and interoperability requirements

Short-lived or member integrated services

Aimed at short-lived ‘translations’, that are organisation/federation bound

- Identity vetting procedures
  - based on an existing ID Management system of sufficient quality**
  - Original identity vetting must be of sufficient quality to trace the individual for as long as name is in active use
  - If *documented* traceability is lost, the subject name can never be re-used
  - revocation and CRL issuing not required for assertion lifetimes << 1 Ms
  - subject naming must be a reasonable representation of the entity name
- Secure operations
  - HSM-backed on-line secured systems
  - data retention and audit trail requirements, traceability of certified entities
- Technical implementation
  - scaling of this model still needs to be demonstrated, and needs coordination
  - *most software and browsers cannot support O(1000) issuers*
  - *and a peer-review based trust fabric cannot do that either …*

Target audience for current profiles

- For users (personal certificates)
  - directly authenticating to end-systems
  - granting unrestricted access to (high-end) resources
- For hosts and services (networked connections)
  - proving identity of a network end-point to a user
    'did I get to the right system?'
  - 'abused' for other services, such as VOMS

Certificates today reflect these two audiences

Users versus Hosts and Services

- Users
  - high-quality identity vetting, so that the same subject name is quite surely bound to the person
  - ‘all’ CAs under the classic profile meet this bar

- Hosts (or ‘service’, e.g. ‘`CN=gatekeeper/ce.example.org`’)
  - the concept of ‘ownership’ of the (DNS) name is vague
  - can be a group of system admins, where the local RA will ensure (‘somehow’, ‘vaguely’) that the requestor is authorized
  - for some CAs, ‘service’ certificates can be requested by ‘service owners’, and no thorough checking is done with the system administrators
  - assurance level for host and service certs is really bound to the use of the DNS name only
  - when used outside securing TLS network-endpoint, the assurance level is ill-defined and varies widely across the IGTF

Hosts versus Robots

If hosts/service assurance level is so ill-defined, what then?

- Raise the assurance level
  - leads to intricate problems when used for the current purpose of securing network endpoints
- Introducing an ‘automated client’ class of certificates
  - identities for programs and services that act in an automated way towards the grid infrastructure
  - concept introduced by Mike Helm in 2002
  - initial criteria developed by Jens Jensen
  - not yet supported by all CAs, but interest is growing (actually, today only UK and NL do, with CZ coming up)

Profile matrix: towards multiple LoAs?

Today, a one-size-fits-all model has helped enormously

- prevent weird-interaction bugs in middleware
- wide interoperability
- based on common RP requirements

- Identity vetting requirements for Classic and MICS
  - only in-person allowed
  - *VOs need to collect this information and more anyway for incident response*

- Is a both stringent and looser LoA needed for other resource classes?

  - when risk profile changes, what about changing the RP requirements?
  - and if so, how do they change?

Current options in AuthN

Current authentication profile options

- **Service certs**
  - the CA *may* allow its use as an automated client
  - but the infrastructures should be wary of accepting them!
  - check of the policy may be needed
    - i.e. in NL, the 'hosts' class identifies network endpoints, as the verification is limited to finding the appropriate system admin; in DoEGrids they are quite weakly linked
- **User certs**
  - generate a proxy from the personal proxy of the portal owner
  - needs the owner to regularly provide the passphrase
  - but works in virtually all scenarios
- **Robots certs (see Jens' talk)**
  - where available (UK, NL, *soon CZ*) these are the preferred choice
  - protects private key from abuse outside the portal system

and, of course, these options can be mixed

downside: requires new Grid AUP/Policies (but no new CA requirements)

‘Alternatives’

- **Do nothing: just use a MyProxy solution**
  - all jobs are traceable to the requesting user
  - portal MyProxy server becomes a valuable target
  - traditional MyProxy is entirely within the current policy space
  - **serious issue:** ‘real’ users cannot handle any kind of credentials

In a pervasive AAI federation environment

- **Federation backed SLCS integrated with the portal**
  - SWITCHaai-like solutions
  - excellent for those countries that have a working AAI *that actually reaches all your researchers* (i.e. CH)
  - Authorize to portal based on AAI account, then generate a cert on the fly form the SLCS service
  - also entirely within current policy space
  - not too many countries have something pervasive …