Grid Security (2)

Grid Computing (M)

Richard Sinnott

UNIVERSITY
of
GLASGOW

Grids in a nutshell… and the security consequences

- Could be argued that Grid all boils down to "*dynamically*" establishing and managing Virtual Organisations (VO)
  - Definition of VO:
    - *dynamic* collection of distributed resources shared by *dynamic* collection of users from one or more organizations

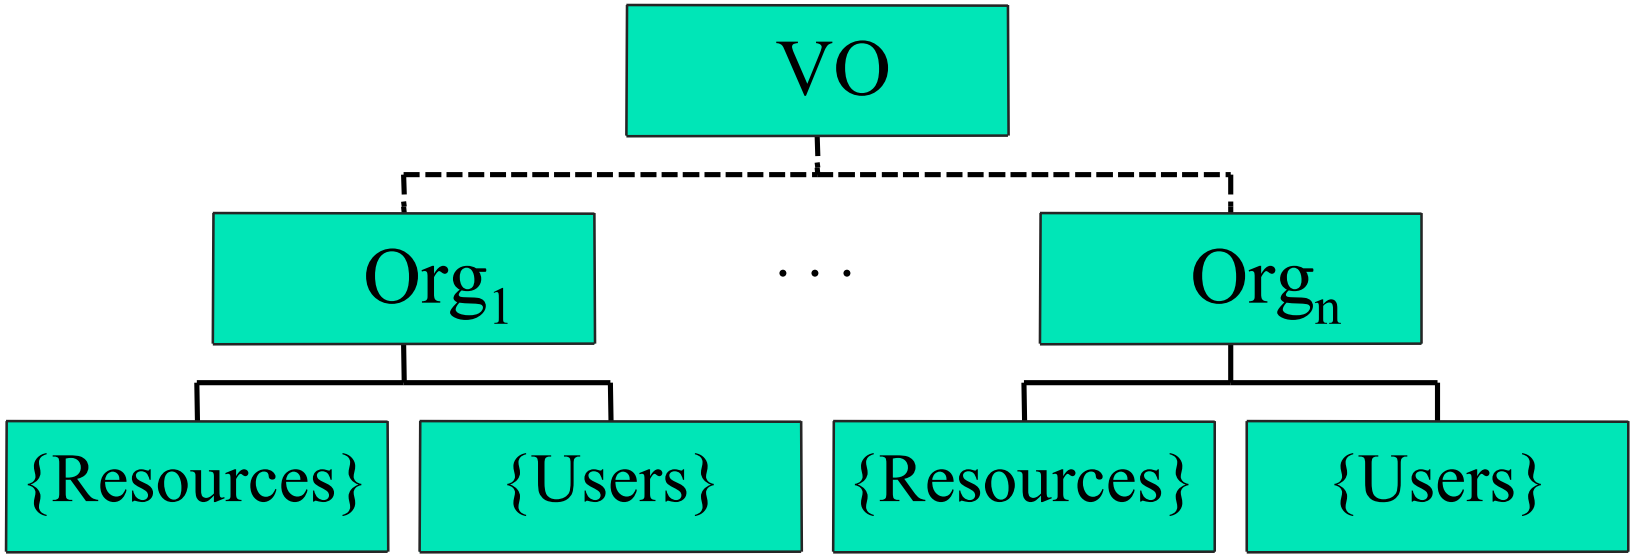

- VO technologies must scale
  - Dealing with potentially huge number of users, resources
  - Broad array of requirements from applications
    - Security, data management, high throughput computing…

Why are VOs important?

- Ability to **securely offer and access** dynamically changing distributed resources **in controlled manner** to dynamically changing groups of users
  - fundamental to way e-Science/e-Research undertaken

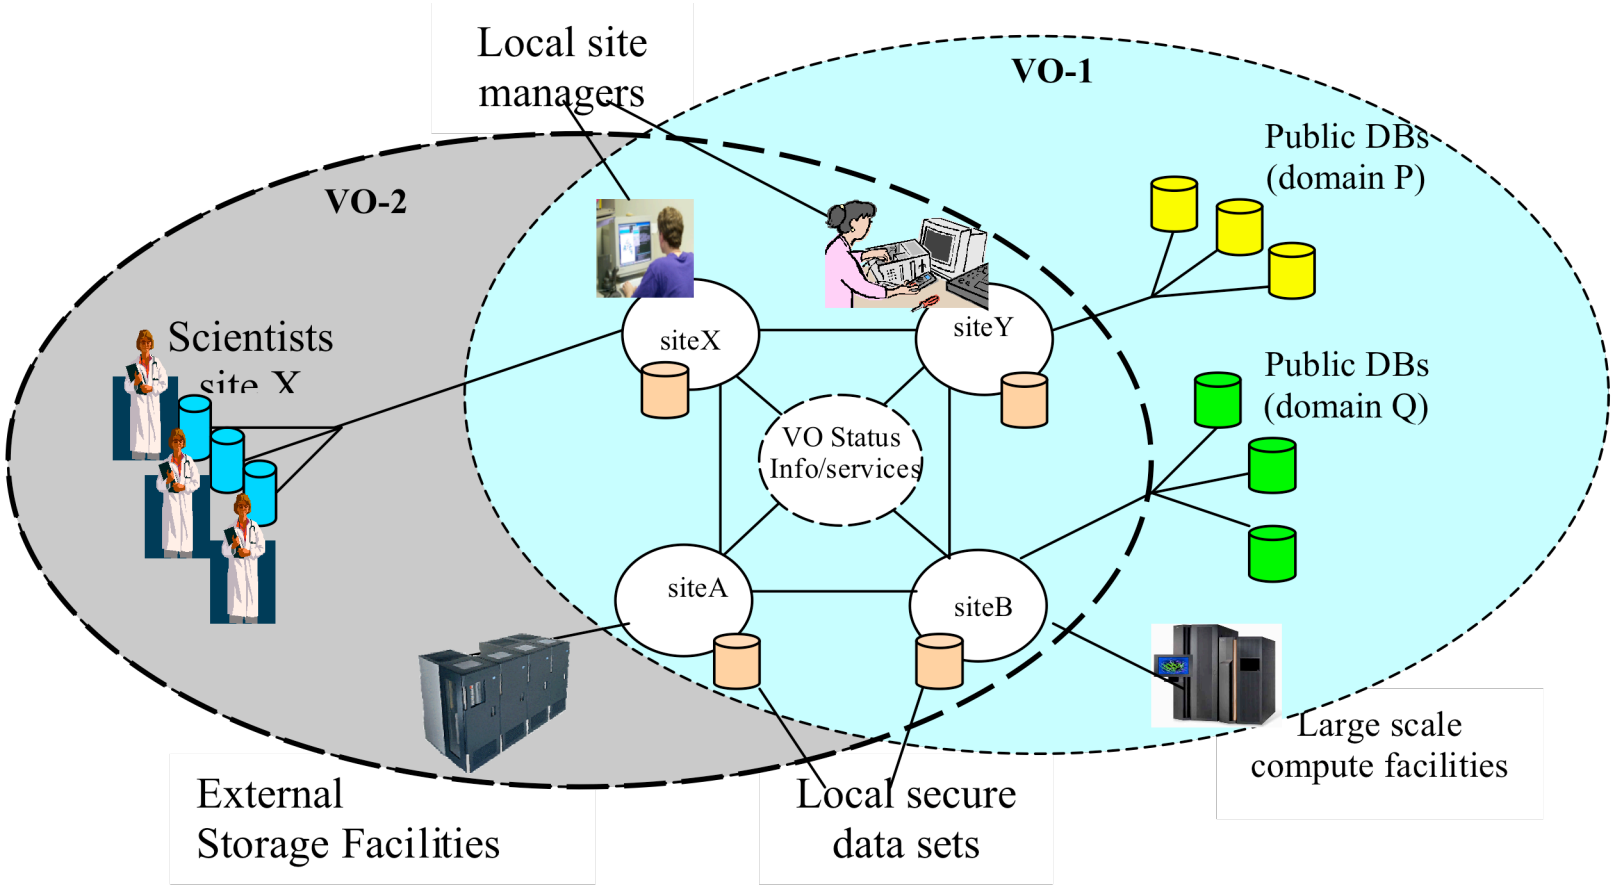

VO Practicalities

- VOs need rules/contracts (policies)
  - Who can do what, on what, in what context, …
- Policies can be direct assertions/obligations/prohibitions on specific resources/users
  - Policies can be local to VO members/resources
    - e.g. user X from site A can have access to P% resource B on site C
      - (site C responsible for local policy – autonomy!!!)
  - Policies can be on remote resources
    - users from site A can access / download data Y from site B provided they do not make it available outside of site A
      - …site B trusts site A to ensure this is the case
        - » and possibly to ensure that the security is comparable with site B
        - » … trust!!!

VO Global Policy Options

- Policies can be global across the VO
  - Compute load across VO should be balanced between all resources
    - Implies
      - scheduling
      - job management
      - accounting
      - … agreed by all VO members
    - Policy aims to try to keep steady state of resource usage
      - May include actions to be taken to maintain desired state
        - » e.g. if any site is performing less than 25% of the work of other sites, new jobs will be scheduled on that site until the workload is balanced
  - Any user using more than 25% of total VO resources have their future jobs not accepted until below this limit
    - Difficult - distributed job management
    - What if nobody else using resources and user has large job?
    - What if policies not explicitly defined, implicit, not implementable, …?
    - Promise you won't make this data public?

VO Global Policy Options

- VO members agree to share resources
  - "Give what you can when you can…" type policy
    - Good will and trust!
      - Easiest to achieve
    - Are we happy that others use our large resource and we get access to their smaller resource?
      - What if we are always busy? They are always free?
  - "Resource usage divided equally among VO members/organisations"
    - How do we measure resource use across VO?
      - Centralised interface (broker) through which all requests flow?
        - » Performance?
      - Job monitoring?
        - » Number of jobs completed? Time processing? Disks used?
        - » Monitoring all jobs, some jobs, jobs per user/per project/per site/per VO…
  - "Get what you give …" type policy
      - Each VO member/organisation receives credit equivalent to the resource utilisation they provide to other users
        - » What is unit of accounting?

VO Policy Issues

- Type and quality of resources vary
  - How do we compare different processors?
    - A 2 day job on a PC with PIII processor and 2GB RAM might complete in 5 minutes on a IBM P690 Regatta Server with 2TB RAM
  - How do we compare processors to disks to IO characteristics to available network at that resource site to …?
    - A 1 day job mining data in flat text files could be done in seconds if the data was indexed and in a DB
  - Often cannot be decided until know exact nature of jobs themselves
    - Some jobs lot more IO intensive
    - Some jobs require inter-process communication
    - Some jobs designed for specific hardware infrastructure, others more generic
    - Some jobs need to move lots of data to/from resource

Policy Considerations

- Do we always want to make such detailed agreements
  - Do we know before setting up VO exactly what policies will be/should be?
- Can we adapt to changing conditions?
- When should the VO take action to enforce it's policy?
  - Always for everything
    - Performance?
  - First violation (trust broken)
  - Sometimes based on statistical averaging of resource usage
- What action should the VO take?
  - Warn/cut-off
    - Demand more access to resources?
    - Restrict access to resources?
    - Remove user/resource from VO
      - Trust broken
  - Redirection
- What if policy violation beyond control of VO partner?
    - network failure, snooper accessing data in transit between sites

VO Consequences

- Members/organizations need to know what will be expected of them before they join VO and what it means to allow someone/some site to join their VO
  - …and consequences of what happens if they don't meet the agreements
- Individual sites trusted to implement the agreed policy
  - If some sites do not conform to policy (or violate) policy?
    - Security ramifications…?
      - Weakest link can affect all others!
        - » Totally secure supercomputing facility allowing access to scientist with own PC in remote location
        - » How do we know they are taking adequate security precautions?
      - Legal impact,
        - » e.g. Data protection act
      - Loss of trust
      - …
    - Increased load on other resources

Technologies for VO

- How does Grid technology meet these challenges?
  - Key that we need way to describe, implement and check/enforce policies

- Should be done at many levels
  - Abstract level to capture overall agreements
    - How best to describe resources, actions, people, …?
  - Design level to ensure that specific points where decisions needed are identified
    - Is there a generic way to achieve this…?
  - Implementation level to ensure that agreements/policies enforced in right places
    - Need to implement collections of rules that can be easily enforced across a variety of end systems

Technologies for VO

- Historically expression of policies not fine grained with Grid toolkits, e.g. Globus
  - For example policies on security based on PKI (previous lecture) and GSI (explored in lab)
    - Globus uses grid mapfile
      - "/C=UK/O=eScience/OU=Glasgow/L=Compserv/CN=john watt" jwatt
      - "/C=UK/O=eScience/OU=Glasgow/L=Compserv/CN=richard sinnott" ros
      - …
    - Users have X.509 certificates which are used to support PKI (single sign on)
    - Applications can check that invoker has appropriate credentials to invoke service
      - i.e. I know that the person with this certificate is registered in my grid mapfile
        - » provides for authentication but need finer grain security (rules/policies)
        - » i.e. authorisation

Authorization Technologies for VO

- Various technologies for authorization including
  - PERMIS
    - PrivilEge and Role Management Infrastructure Standards Validation
      - http://www.permis.org
  - Community Authorisation Service
    - http://www.globus.org/security/CAS/
  - AKENTI
    - http://www-itg.lbl.giv/security/akenti
  - CARDEA
    - http://www.nas.nasa.gov/Research/Reports/Techreports/2003/nas-03-020-abstract.html
  - VOMS
    - http://hep-project-grid-scg.web.cern.ch/hep-project-grid-scg/voms.html
  - All of them predominantly work at the local policy level

Standards for Generic Authorisation

Generic way to achieve authorisation defined in X.812|
ISO 10181-3 Access Control Framework

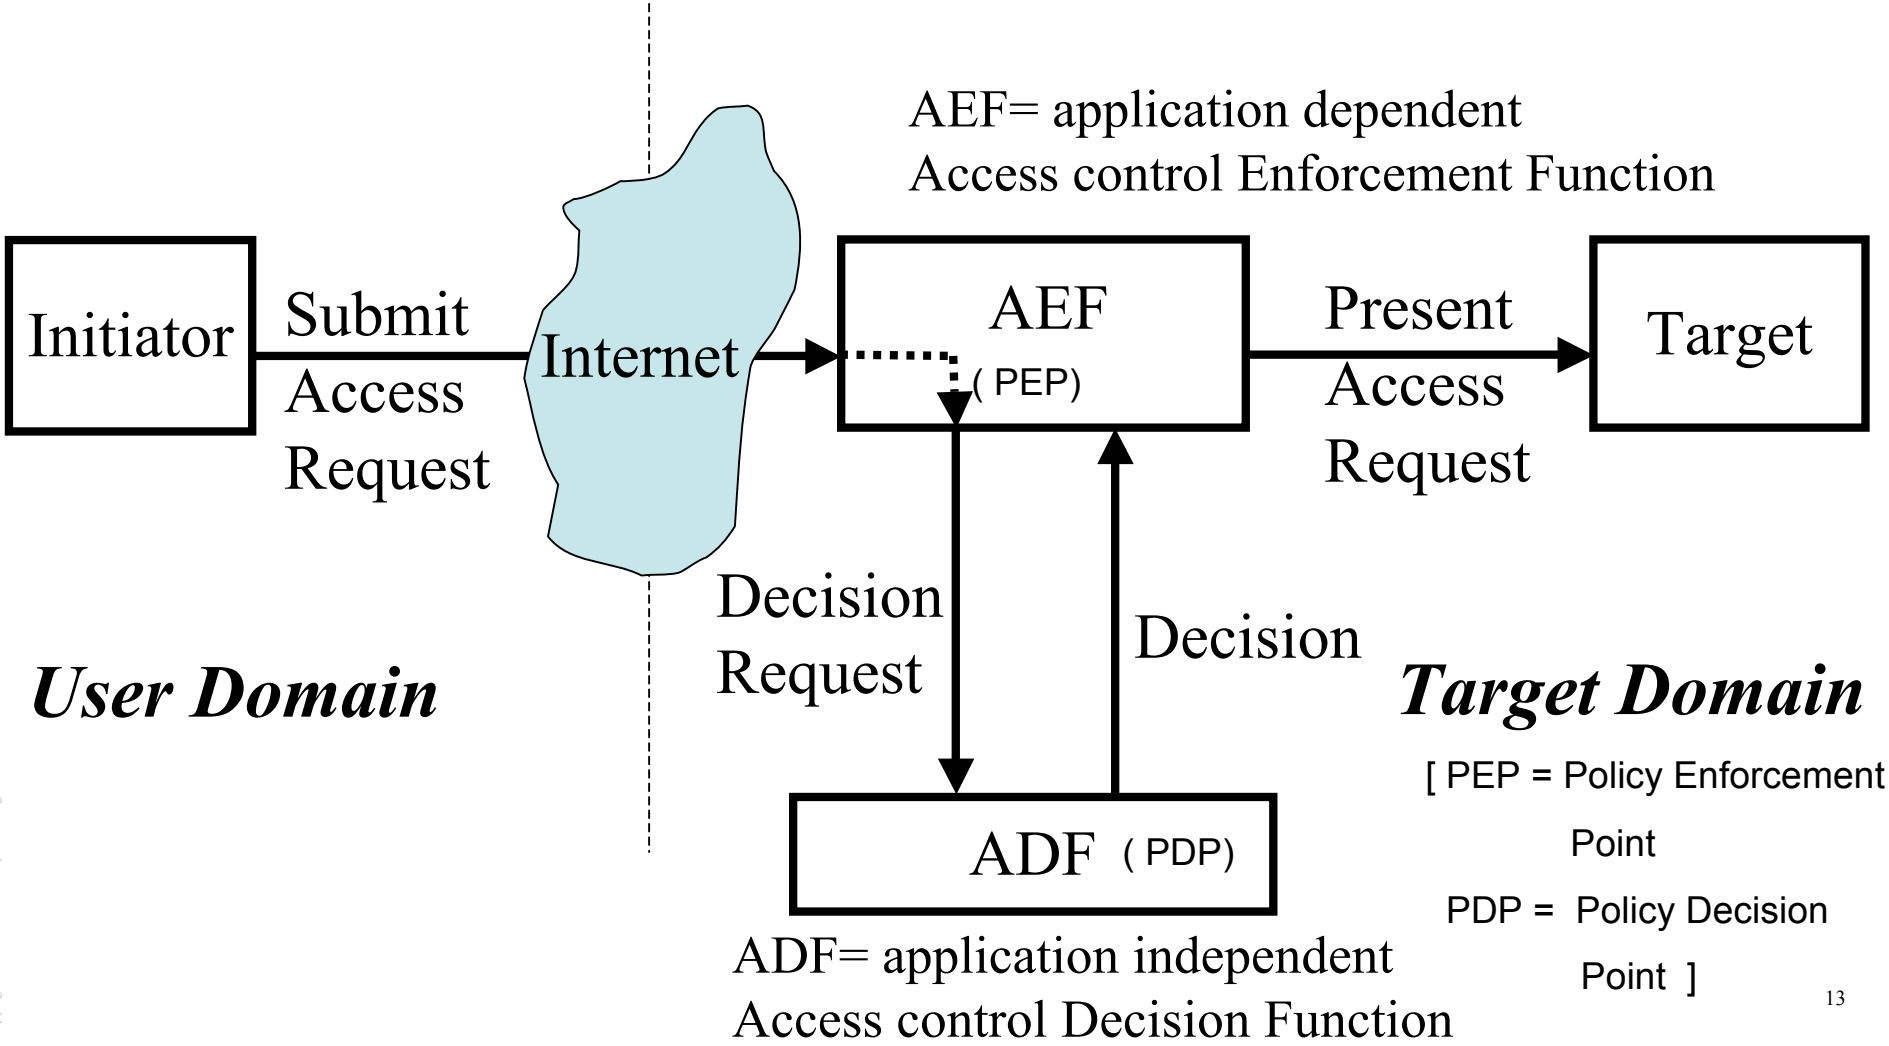

Grid APIs for Generic Authorisation

- Global Grid Forum (GGF) SAML AuthZ specification provides generic AEF approach for ALL Grid services
  - … or at least all GT3.3+ based services

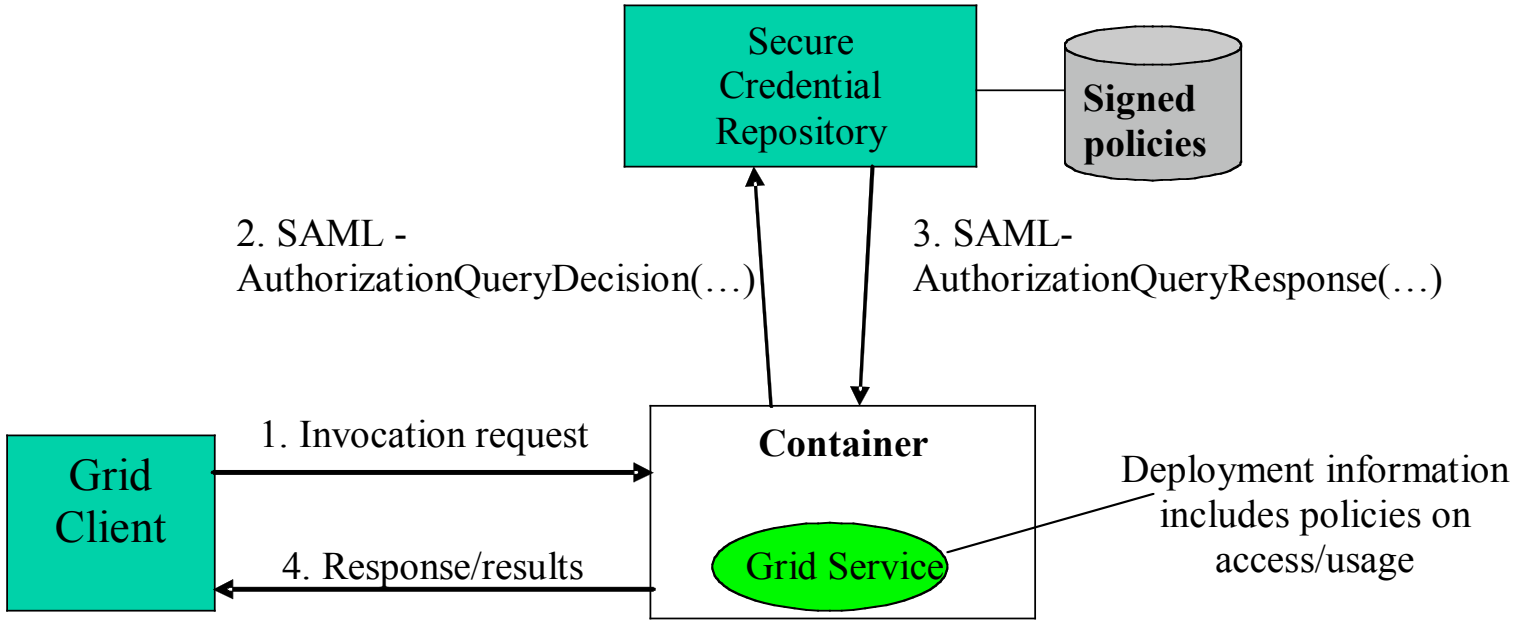

  - PDP application specific
    - Previous assignments have looked at PERMIS in detail (not this time!)
      - Default behaviour is if not explicitly granted by policy, then rejected

Role Based Access Controls

- Need to be able to express and enforce policies
  - Common approach is role based authorisation infrastructures
    - PERMIS, CAS, …
- Basic idea is to define:
  - *roles* applicable to specific VO
    - roles often hierarchical
      - Role X ≥ Role Y ≥ Role Z
      - Manager can do everything (and more) than an employee can do who can do everything (and more) than a trainee can do
  - *actions* allowed/not allowed for VO members
  - *resources* comprising VO infrastructure (computers, data resources etc)
- A policy then consists of sets of these rules
    - *{ Role x Action x Target }*
      - Can user with VO role X invoke service Y on resource Z?
    - Policy itself can be represented in many ways,
      - e.g. XML document, SAML, XACML, …

RBAC Policy Components

- Subject Policy
  - Specifies subject domains, e.g. dcs.gla.ac.uk
- Role Hierarchy Policy
  - Specifies hierarchy of role values, e.g. VO scientist, sys-admin
- SOA Policy
  - Specifies who is trusted to issue ACs (typically local sys-admin)
- Role Assignment Policy
  - Says which roles can be given to which subjects by which SOAs, with which validity times and whether delegation is allowed (depends on VO)
- Target Policy
  - Specifies the target domains covered by this policy (e.g. Grid services)
- Action Policy
  - Specifies the actions (methods/operations on Grid services) supported by the targets
- Target Access Policy
  - Specifies which roles are needed to access which targets for which actions, and under what conditions

PERMIS Based Authorisation

- PERMIS Policies created with PERMIS PolicyEditor (output is XML based policy)

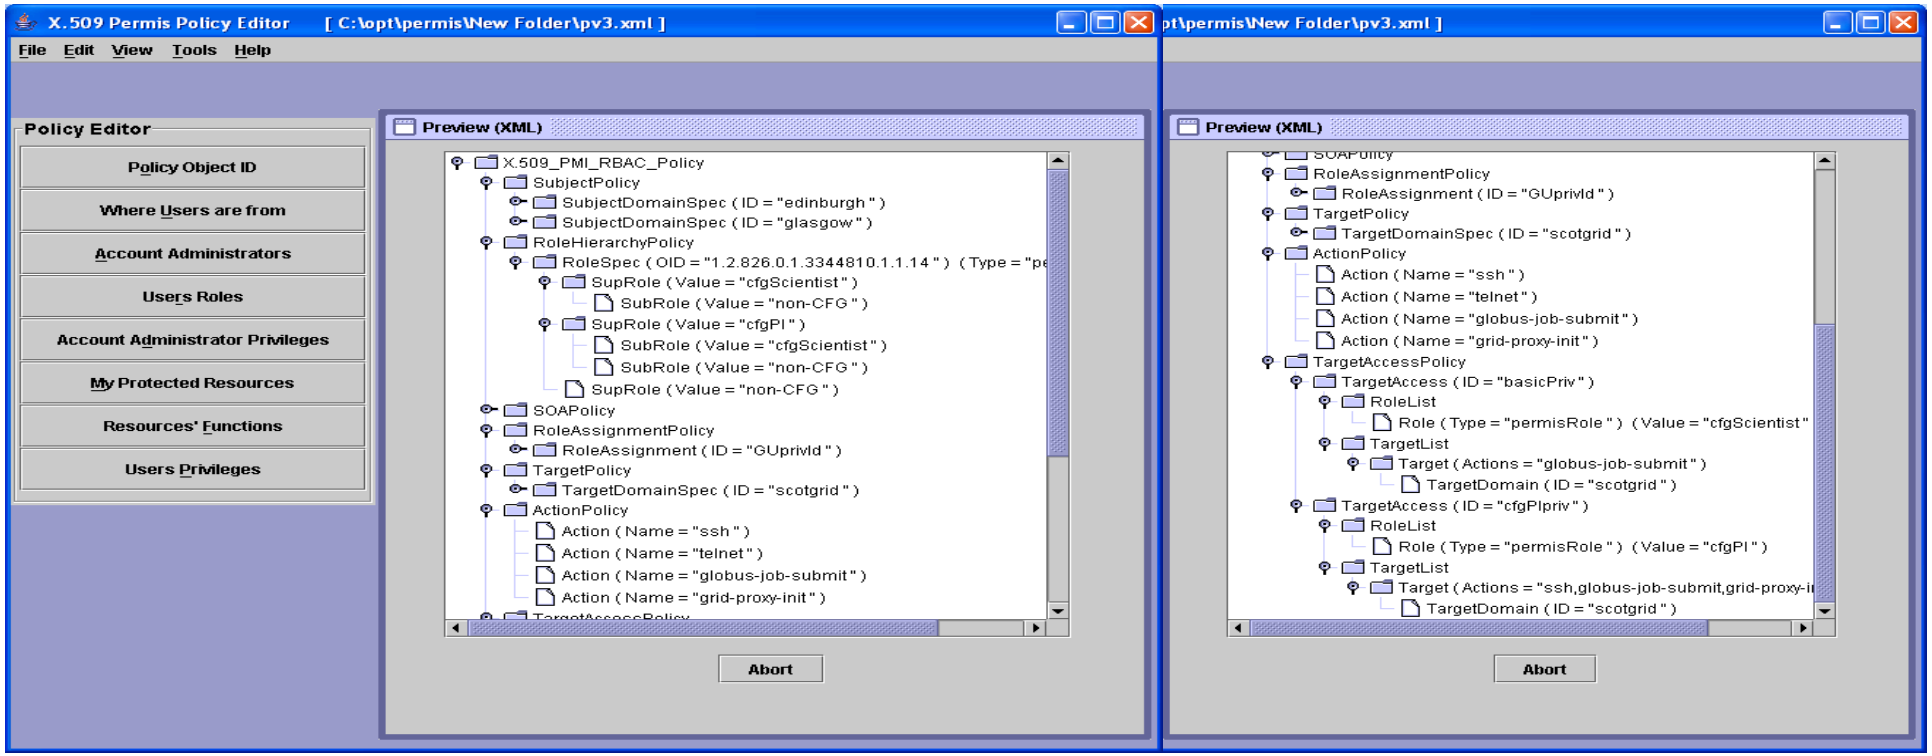

- Other PERMIS tools then used to sign policies
  - Associates roles with specific users
    - Policies stored as attribute certificates in LDAP server

Things for the future…

- Major WS-* efforts on many fronts
  - These range from proposal ideas, partial/full specifications, actual implementations (WS-Security)
    - Whole area driven by commercial players/politics
      - » WS-Notifications vs WS-Eventing
      - » WS-ReliableMessaging vs WS-Reliability
      - » WS-Orchestration vs WS-Co-ordination vs WS-Choreography
      - » WSFL vs BPEL vs …
      - » "WS-make me a cup of tea" vs "WS-make me a cup of coffee"
  - XACML (eXtensible Access Control MarkUp Language)
    - Richer possibilities for policy expression
      - Tools, complexity, …

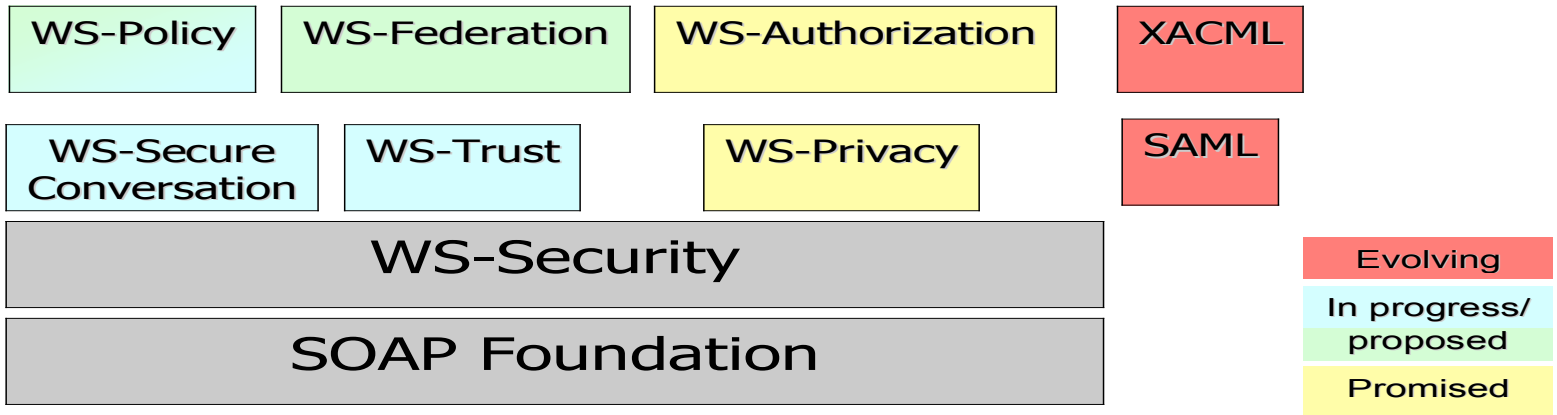

Introducing Shibboleth

- Shibboleth (http://shibboleth.internet2.edu)

Definition

Shibboleth [Hebrew for an ear of corn, or a stream or flood]

1. A word which was made the criterion by which to distinguish the *Ephraimites* from the *Gileadites*. The *Ephraimites*, not being able to pronounce sh, called the word sibboleth. See --Judges xii.
2. Hence, the criterion, test, or watchword of a party; a party cry or pet phrase.]

- Shibboleth will replace Athens as access mgt system across UK academia
  - UK federation went live on 30th November 2006

– Federations based on trust
  - or more accurately trust but verify
  - numerous international federations exist MAMS, SWITCH, HAKA, …

Typical Shibboleth Scenario

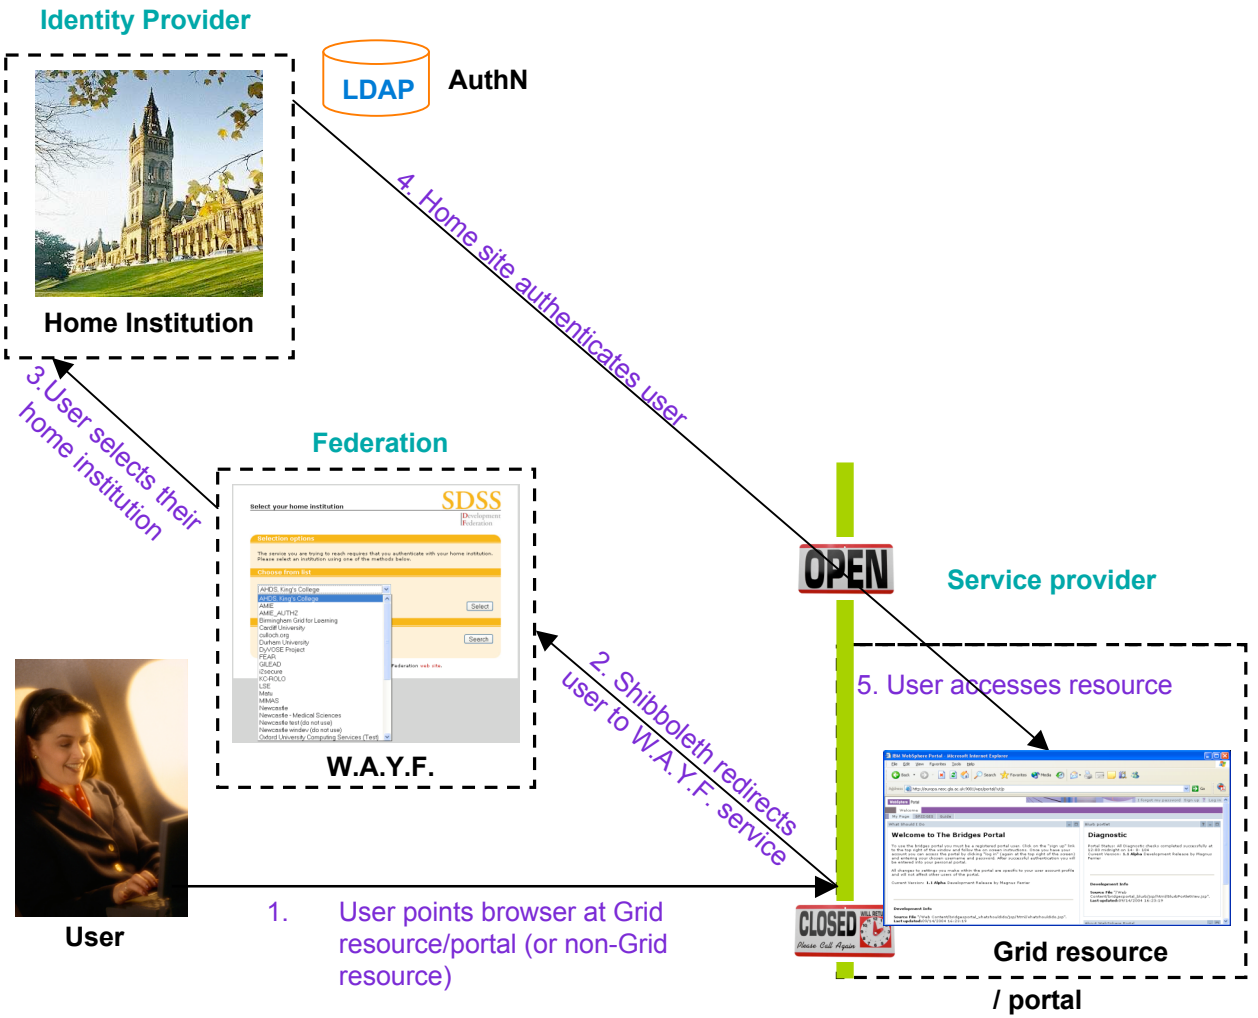

It's a start, but…

- Benefit from local authentication but really want finer grained control…
  - I know you have authenticated, but I need to know that you have sufficient/correct privileges to access my VO resources
  - can also return various other information needed to support ***authorisation*** decisions

Finer Grained Shibboleth Scenario

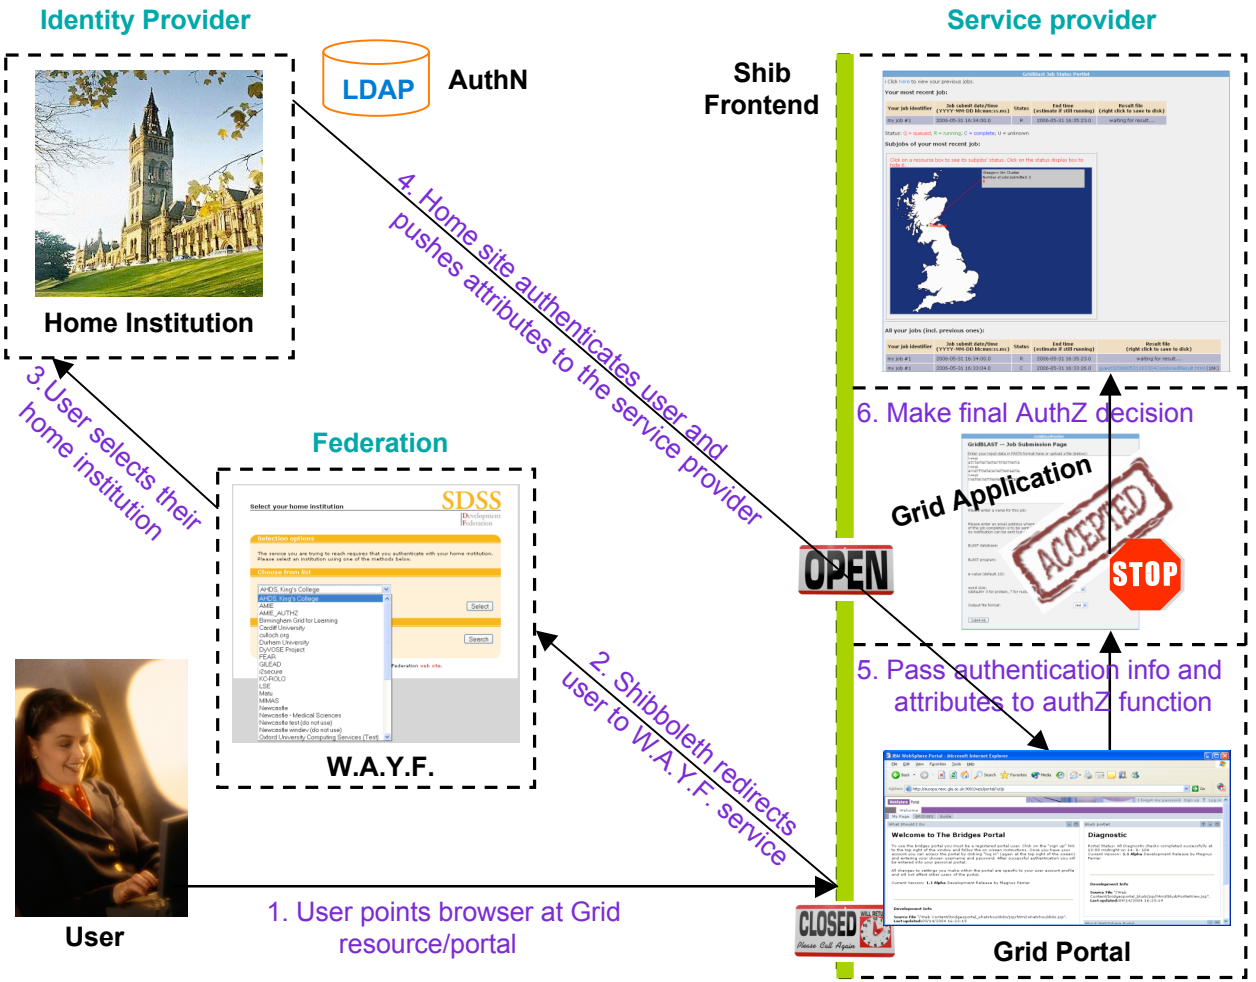

Ok, but…

- I can do authorisation but I want single-sign on to lots of distributed resources across my VO (or VOs)
  - Browser allows to keep session information so can access other resources without signing in again
    - Provided authorisation information valid for different service providers
      - Each service provider completely autonomous
    - Can configure attribute release/attribute acceptance policies per identity provider/service provider

Shibboleth issues…

- Federations are quite rigid/static and not in true dynamic Grid vision
  - Ok for some domains (e.g. clinical) where we don't really want truly dynamic Grids and you will hopefully never find new data/resources "on-the-fly"
- Federations and users must all pre-agree on security attributes (and their values)
  - Can enforce things like you can only use this service if you have a license for the software at your home site
  - Only users with role of "Glasgow Royal Infirmary consultant" can access this service/data set
  - eduPerson attributes being explored and various others on larger scale
  - Policies on attribute release, attribute negotiation etc all being worked on

Putting the "Dy" in DyVOSE

- **Dynamic PMI Case Study**

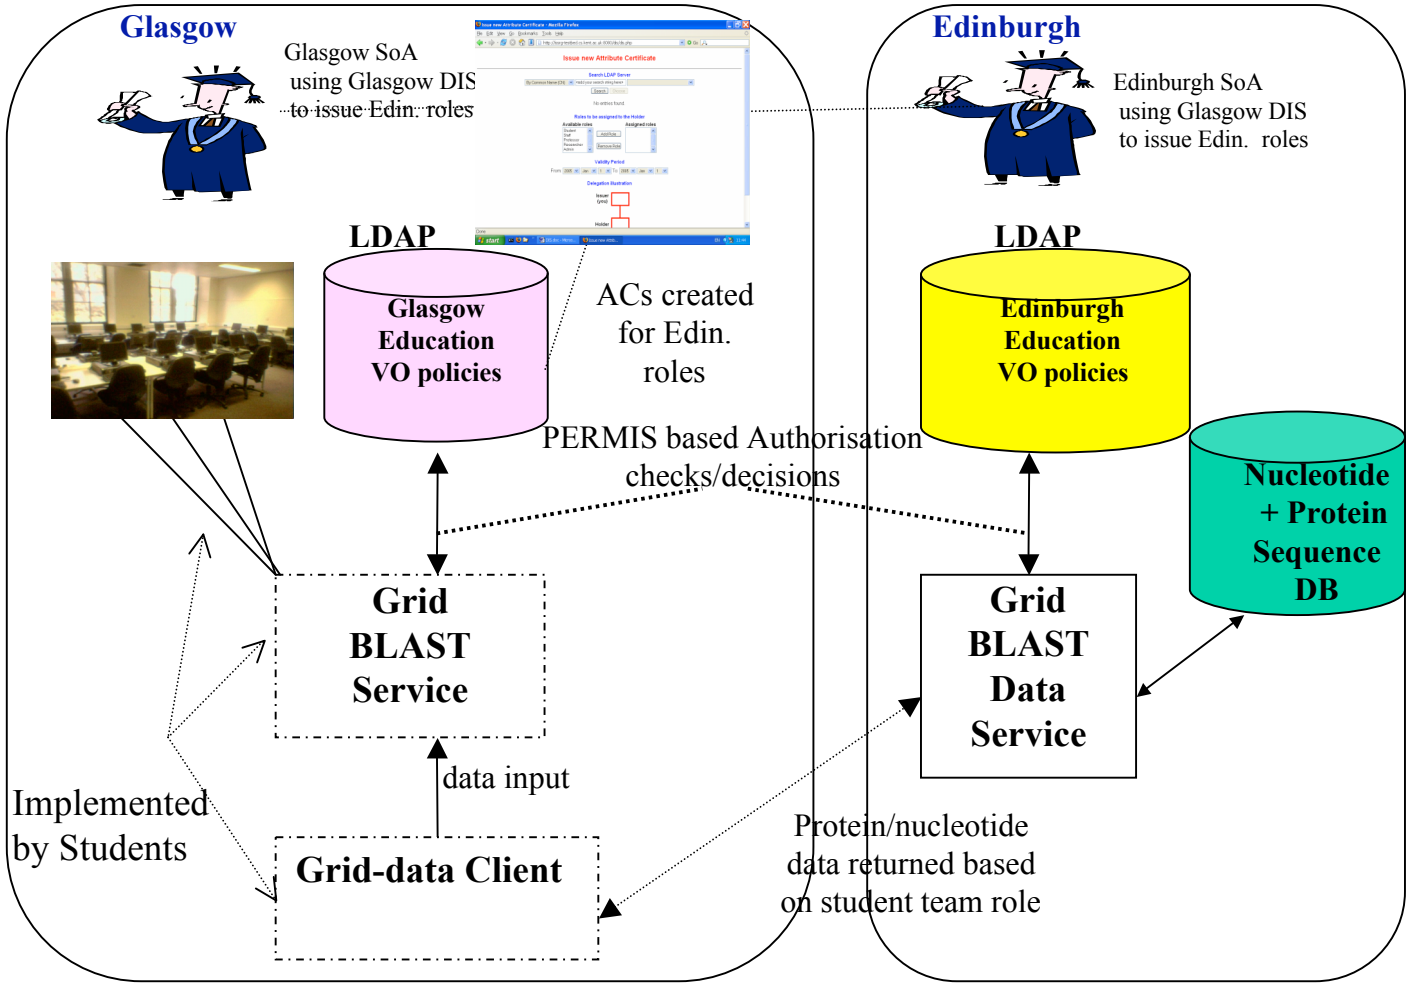

Shibboleth issues …ctd

- Trust underpins Shibboleth/Grids
  - What if remote site does not treat authentication as seriously as it should?

- University of Glasgow used to have
  - Multiple usernames/passwords for staff students
  - Now moved to single unified account management system based on Novell nSure active directory technology
    - Identity management based on
      - Human Resources information for staff
      - Registry for students
    - Based on this have Shib-enabled numerous non-Grid resources
      - **WebSurf**
        - » Student/staff service, e.g. courses registered, credits earned etc
      - **Moodle**
        - » Glasgow virtual e-learning environment
      - Various others

Conclusions

- VOs crucial to Grids
  - Must overcome limitations of PKI scalability, security
- Need way to express rules/policies
  - How detailed?
  - How dynamic?
  - What about performance…?
- Standards and specifications/implementations being put together
  - GGF AuthZ works
    - (but requires authZ/Grid technologies to implement it)
- Clear need for more experiences applying technologies
- Shibboleth is definitely coming and will influence how we interact with Grids and VOs in the future
  - At NeSC Glasgow we were first showing how to access Grids via Shibboleth technologies… hurrah!
  - We'll see these things in action later in Example Applications Lecture