ESP-GRID
JISC

oerc

A case for Shibboleth and grid security: are we paranoid about identity?

UK e-Science All Hands Meeting, 2006

Mark Norman 19 Sept 2006

This talk

- The ESP-GRID project
- What is Shibboleth?
- Our requirements for grid security
  - How we usually express those requirements
  - What are the *real* requirements?
- Shibboleth vs (usual GSI-style) PKI
  - Assertion of permanent identifier versus
  - Assertion of current membership and suitability

The ESP-GRID Project

- The Evaluation of Shibboleth and PKI for Grids (ESP-GRID) Project
  - Ran July 04 to June 06
- Aim was to investigate whether and how Shibboleth offers solutions to issues of grid authentication, authorisation and security
- The first questions we hit were:
  - Do we *really* need explicit IDs asserted everywhere?
  - What would make things scalable?

Unforgeable or unrevocable?

- What's the best situation?
  - Using unforgeable paper that lasts for 10 years
  - or…
  - Checking with a trusted colleague to vouch for you every time ??
- The answer is (obviously) "it depends"
- It's all about authorisation…

What is Shibboleth?

- "Shibboleth is a system designed to exchange attributes across realms for the primary purpose of authorisation"
  - It's not strictly an authentication mechanism
  - Nor an authorisation mechanism
    - It enables both

- But in plainer speaking…

What is Shibboleth?

- It's all about how to transmit the authorisation and role information from your home institution to outside service providers
- And how those service providers can ask for that information
- Access management and the communication of authorisation credentials
- Aims: to separate authentication from authorisation
  - Devolve authentication to the 'home' organisation
  - Devolve the management of authorisation information as well

Can we use it on grids?

- It's not quite that easy!
  - Grids tend to use X.509 digital certificates
    - (Centrally/Nationally issued)
    - A bit hard to use (but that's a different matter)
  - Shibboleth is (so far) based in the web world
    - HTTP only
  - Some grid people think that
    - Certificates = secure
    - University libraries/SSO = insecure
      - (This is probably wrong, but grids *do* need higher security)

Grid security: what are we trying to secure?

- Current grid users have quite a deep level of control over the host machines
  - If you compare them with web browser/readers, for example
- Is this scaleable to high numbers of users?
  - Surely these people are Power Users (and we just haven't got many Regular Users yet)?
  - Are personal digital certificates suitable for Power Users and something else (e.g. Shib) good for ordinary Users?

Should we devolve ID management?

(Shibboleth is an ideal model for this)

- Currently we use 12 month credentials, checked once (possibly years ago)
  - And trust that the user hasn't gone bad, or if s/he does, we'll somehow get to hear about it
- When ID management is devolved:
  - Less central control of user policies (bad)
  - When someone steals something, they get their accounts immediately suspended (good)

Should we devolve ID management? (2)

(Shibboleth is an ideal model for this)

- Maybe we can't trust the local ID managers
  - They might be really sloppy with their ID management!
- But we trusted the University Card that they issued
  - So we were trusting them all along!?!

GSI-style PKI: it's identity, not AuthZ

- We all know that you should separate authentication and authorisation
- But having a certificate tends to mean that 'you're a member of the community for a year'
  - That's mixing authN and authZ
- It would be OK if someone had a long term certificate (e.g. 10 years) and we had sophisticated authZ

Emotional security

- "We must know who everyone is in case they do something bad"
  - Do we really mean that?
- How about: "If someone/something does something bad (on purpose or otherwise), I need to suspend their jobs and their activity until they are identified and dealt with"
  - So maybe you don't need to know their ID up front?

Real requirements

- Rapid suspension
  - Get the bad thing to stop ASAP
    - and you could go for slower (human mediated) identification/revocation? (It's not outside the requirement)
- Identification of wrong-doers
- Revocation of credentials for deliberate wrong-doers
- *These are the* real *requirements. The supposed requirement of needing to know an ID 'up front' is* one solution *for these requirements.*

A case for Shibboleth

- If we agree that…
  - numbers of users will grow
  - most current users are Power Users
  - most (future) regular users will be less technical
  - most regular users will be interested in (predictable) applications
  - (and maybe) we will have methods to deal with rogue users and jobs based upon suspend and then investigate
- then Shibboleth is attractive

Grids use GSI

- But grids are based on Grid Security Infrastructure
  - Tied inextricably to X.509 digital certificates
- Grids *do* have greater need for high security
- What about people using (e.g.) portals with applications in portlets?
  - The portal could authenticate them using username/password – or even via Shibboleth
  - The portal talks to the grid using GSI
  - (See other paper for the Customer-Service Provider model)

The C-SP model

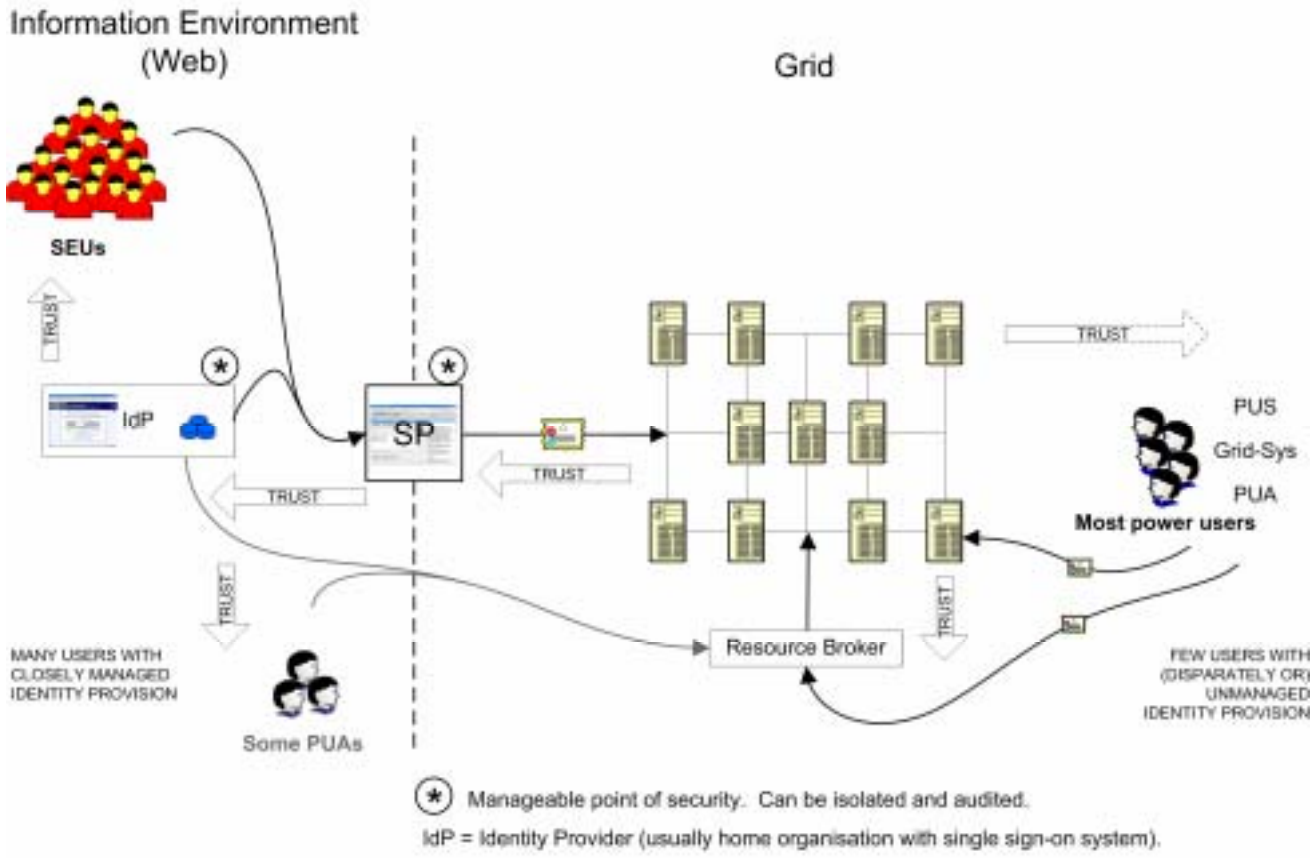

A last word about “identity”

- Is your identity a unique (scoped) identifier?
  - e.g. citizen 0001234567.uk (or mark.norman@oucs.ox.ac.uk)
- Or is your identity a whole list of things?
  - e.g. member of oucs,
  - member of oerc,
  - member of Oxford University,
  - part of ESP-GRID project,
  - staff,
  - “Mark Norman”,
  - male?

ESP-GRID
JISC

oerc

A case for Shibboleth and grid security: are we paranoid about identity?

UK e-Science All Hands Meeting, 2006

Mark Norman 19 Sept 2006