



Enabling Grids for E-science

Authentication, Authorisation and Security

*Mike Mineter,
National e-Science Centre
mjm@nesc.ac.uk*


www.eu-egee.org




- This presentation can be re-used for academic purposes.
- However if you do so then please let training-support@nesc.ac.uk know. We need to gather statistics of re-use: no. of events, number of people trained. Thank you!!

- **Providers of resources (computers, databases,...) need risks to be controlled: they are asked to trust users they do not know**
 - They trust a VO
 - The VO trusts its members
- **User's need**
 - single sign-on: to be able to logon to a machine that can pass the user's identity to other resources
 - To trust owners of the resources they are using
- **Build middleware on layer providing:**
 - *Authentication*: know who wants to use resource
 - *Authorisation*: know what the user is allowed to do
 - *Security*: reduce vulnerability, e.g. from outside the firewall
 - *Non-repudiation*: knowing who did what
- **The “Grid Security Infrastructure” middleware is the basis of (most) production grids**

- **Achieved by Certification:**
 - User’s identity has to be certified by one of the national *Certification Authorities (CAs)*
 - mutually recognized <http://www.gridpma.org/>
 - In UK go to <http://www.grid-support.ac.uk/ca/ralist.htm> to find CA’s local “Registration Authorities”
 - Resources are also certified by CAs
- **User**
 - User joins a VO
 - Digital certificate is basis of AA
 - Identity passed to resources you use, where it is mapped to a local account
- **Virtual Organization negotiates rights to use resources**





User

Resource


- How does a user securely access the Resource without having an account on the machines in between or even on the Resource?
- How does the Resource know who a user is?
- How are rights and that they are allowed access?

Authentication: how is identity of user/site communicated?

Authorisation: what can a user do?

- **Launch attacks to other sites**
 - Large distributed farms of machines, perfect for launching a Distributed Denial of Service attack.
- **Illegal or inappropriate data distribution and access sensitive information**
 - Massive distributed storage capacity ideal for example, for swapping movies.
 - Growing number of users have data that must be private – biomedical imaging for example
- **Damage caused by viruses, worms etc.**
 - Highly connected infrastructure means worms could spread faster than on the internet in general.

- **Asymmetric encryption...**




- **.... and Digital signatures ...**

- A hash derived from the message and encrypted with the signer's private key
- Signature is checked by decrypting with the signer's public key


- **Are used to build trust**

- That a user / site is who they say they are
- And can be trusted to act in accord with agreed policies



- Every user has two keys: one *private* and one *public*:
 - it is *impossible* to derive the private key from the public one;
 - a message encrypted by one key can be decrypted **only** by the other one.



- Public keys are exchanged
- The sender encrypts using his private key
- The receiver decrypts using senders public key;
- The number of keys is $O(n)$

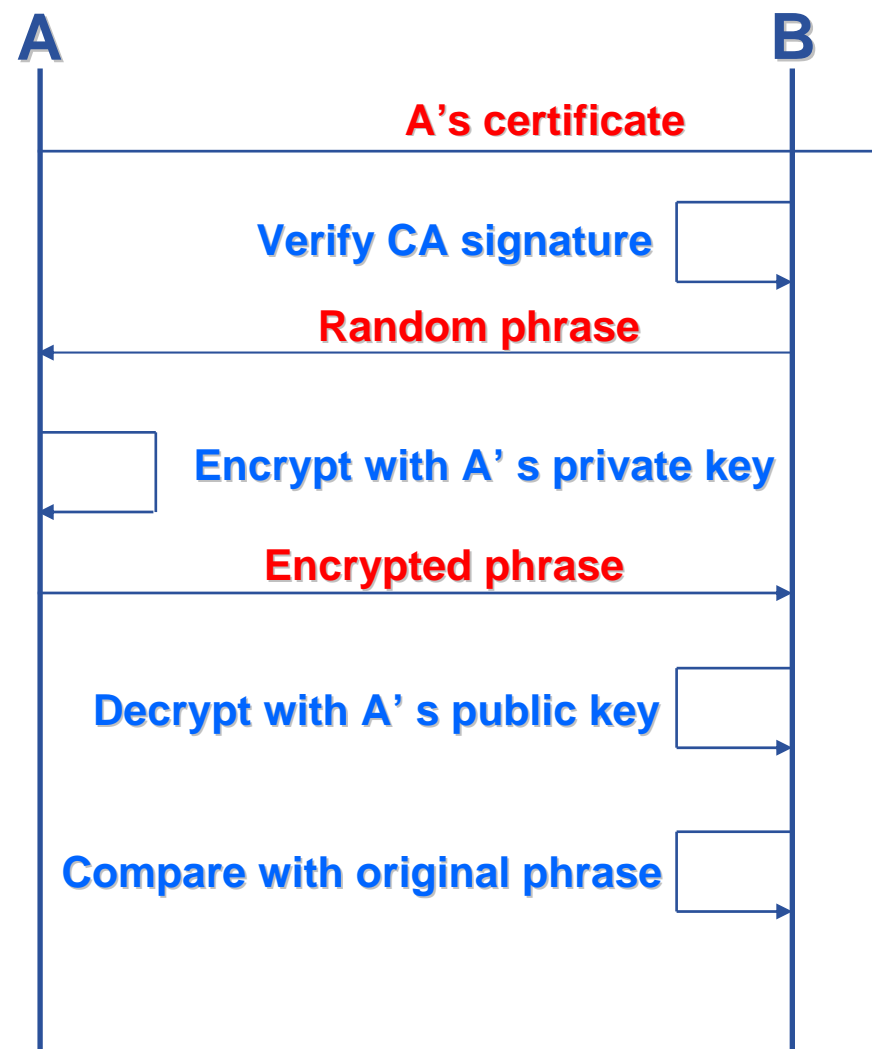


- Paul calculates the *hash* of the message
- Paul encrypts the hash using his *private* key: the encrypted hash is the digital signature.
- Paul sends the signed message to John.
- John calculates the hash of the message
- Decrypts signature, to get A, using Paul's *public* key.
- If hashes equal:
 1. message wasn't modified;
 2. hash A is from Paul's private key



Based on X.509 PKI:

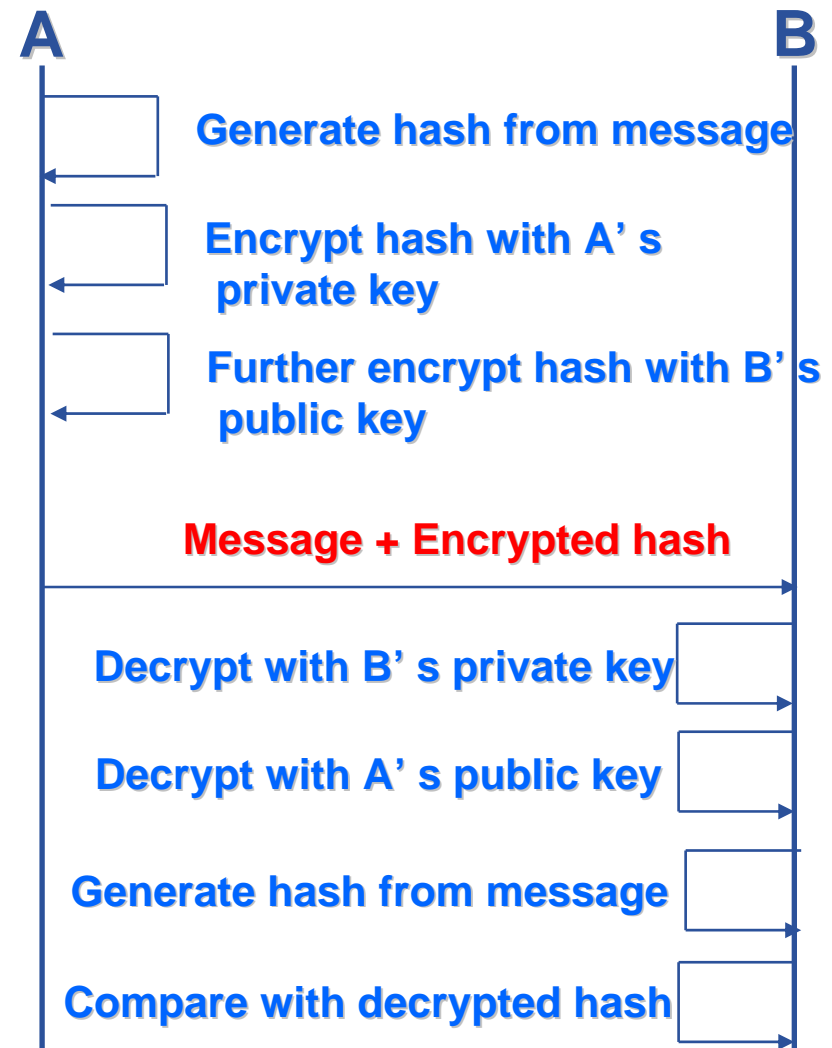
- every Grid transaction is mutually authenticated:
 1. A sends his certificate;
 2. B verifies signature in A's certificate using CA public certificate;
 3. B sends to A a challenge string;
 4. A encrypts the challenge string with his private key;
 5. A sends encrypted challenge to B
 6. B uses A's public key to decrypt the challenge.
 7. B compares the decrypted string with the original challenge
 8. If they match, B verified A's identity and A can not repudiate it.
 9. Repeat for A to verify B's identity



After A and B authenticated each other, for A to send a message to B:

- **Default: message integrity checking**
 - Not private – a test for tampering

- **For private communication:**
 - Encrypt all the message (not just hash) - Slower




- How can John be sure that Paul's public key is really Paul's public key and not someone else's?
 - A *third party* certifies correspondence between the public key and Paul's identity.
 - Both John and Paul trust this third party

The “third party” is called a Certification Authority (CA).

- An X.509 Certificate contains:

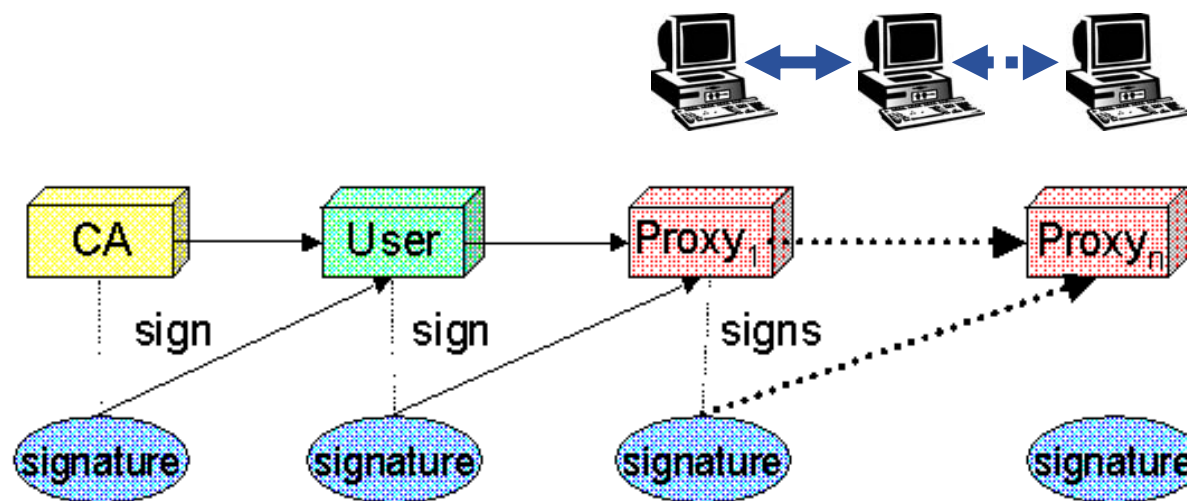
- owner's public key;
- identity of the owner;
- info on the CA;
- time of validity;
- Serial number;
- Optional extensions

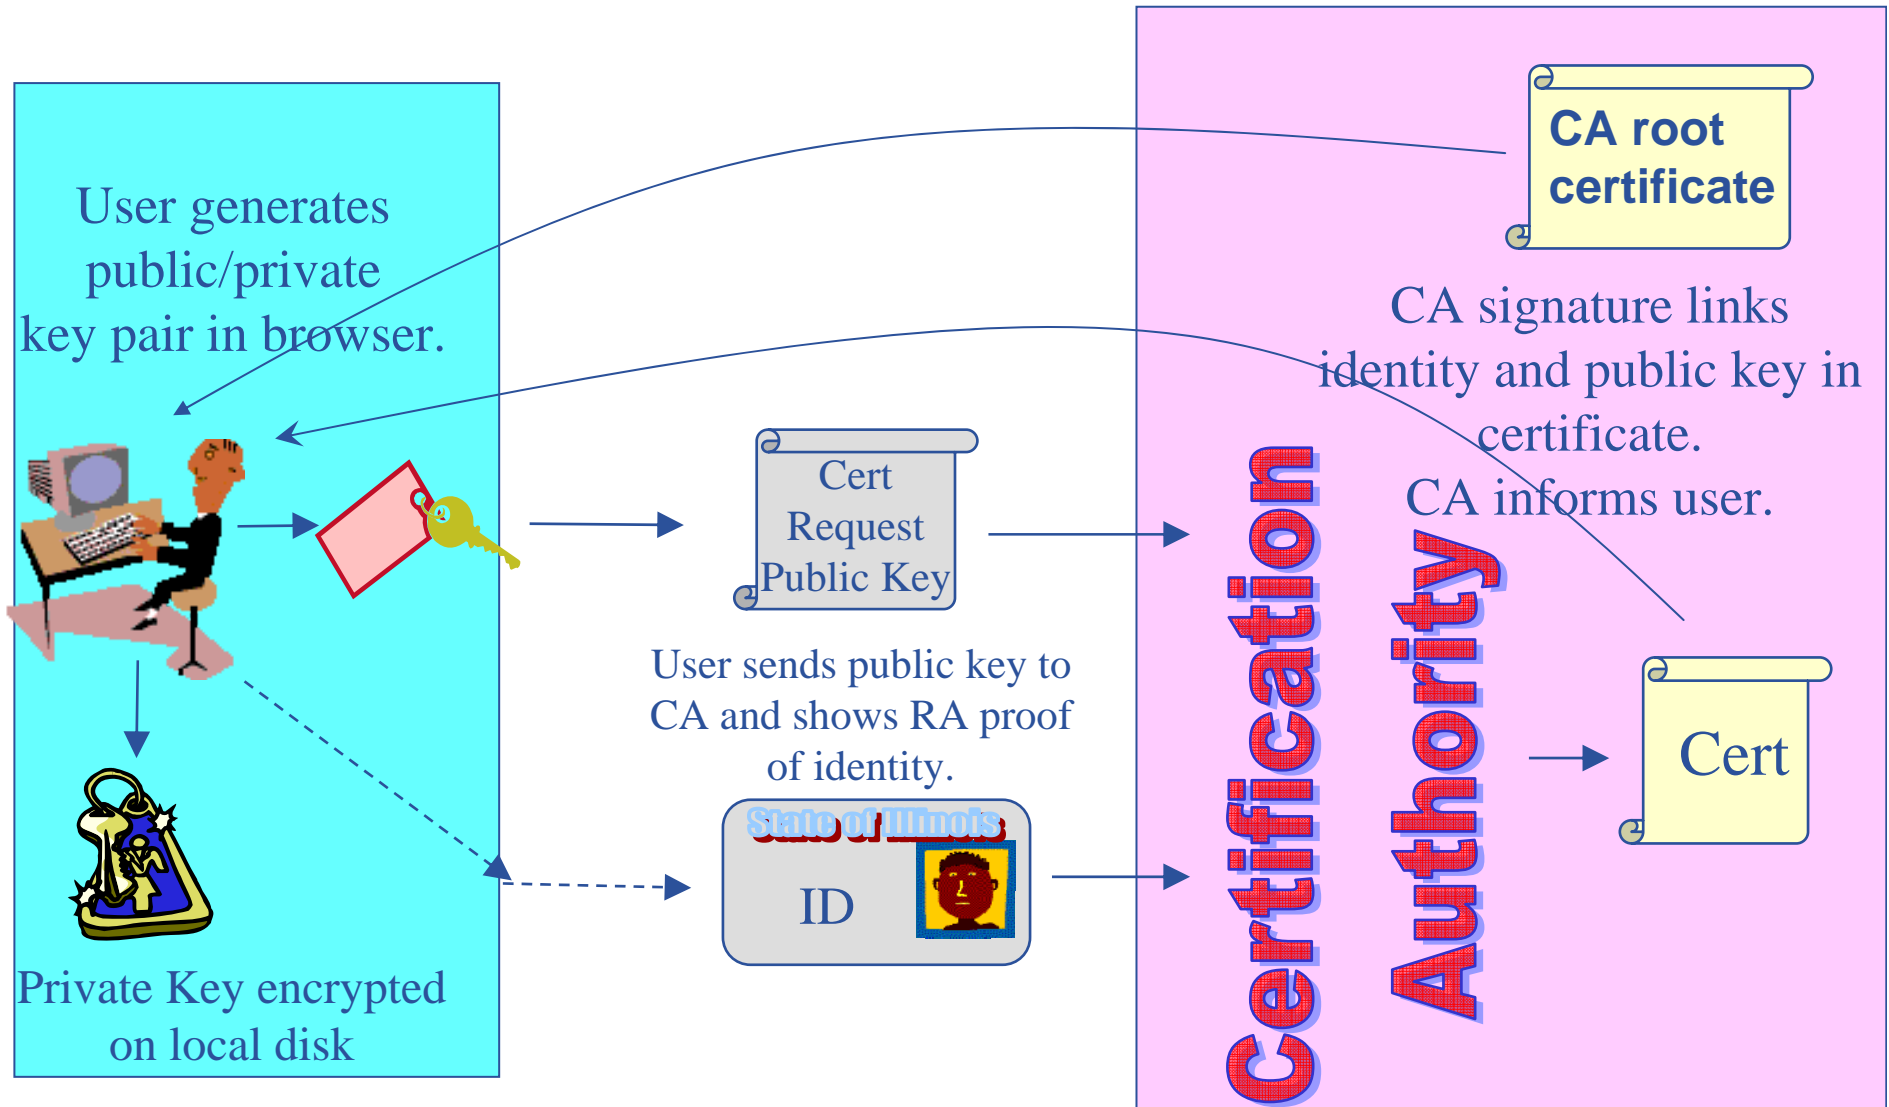


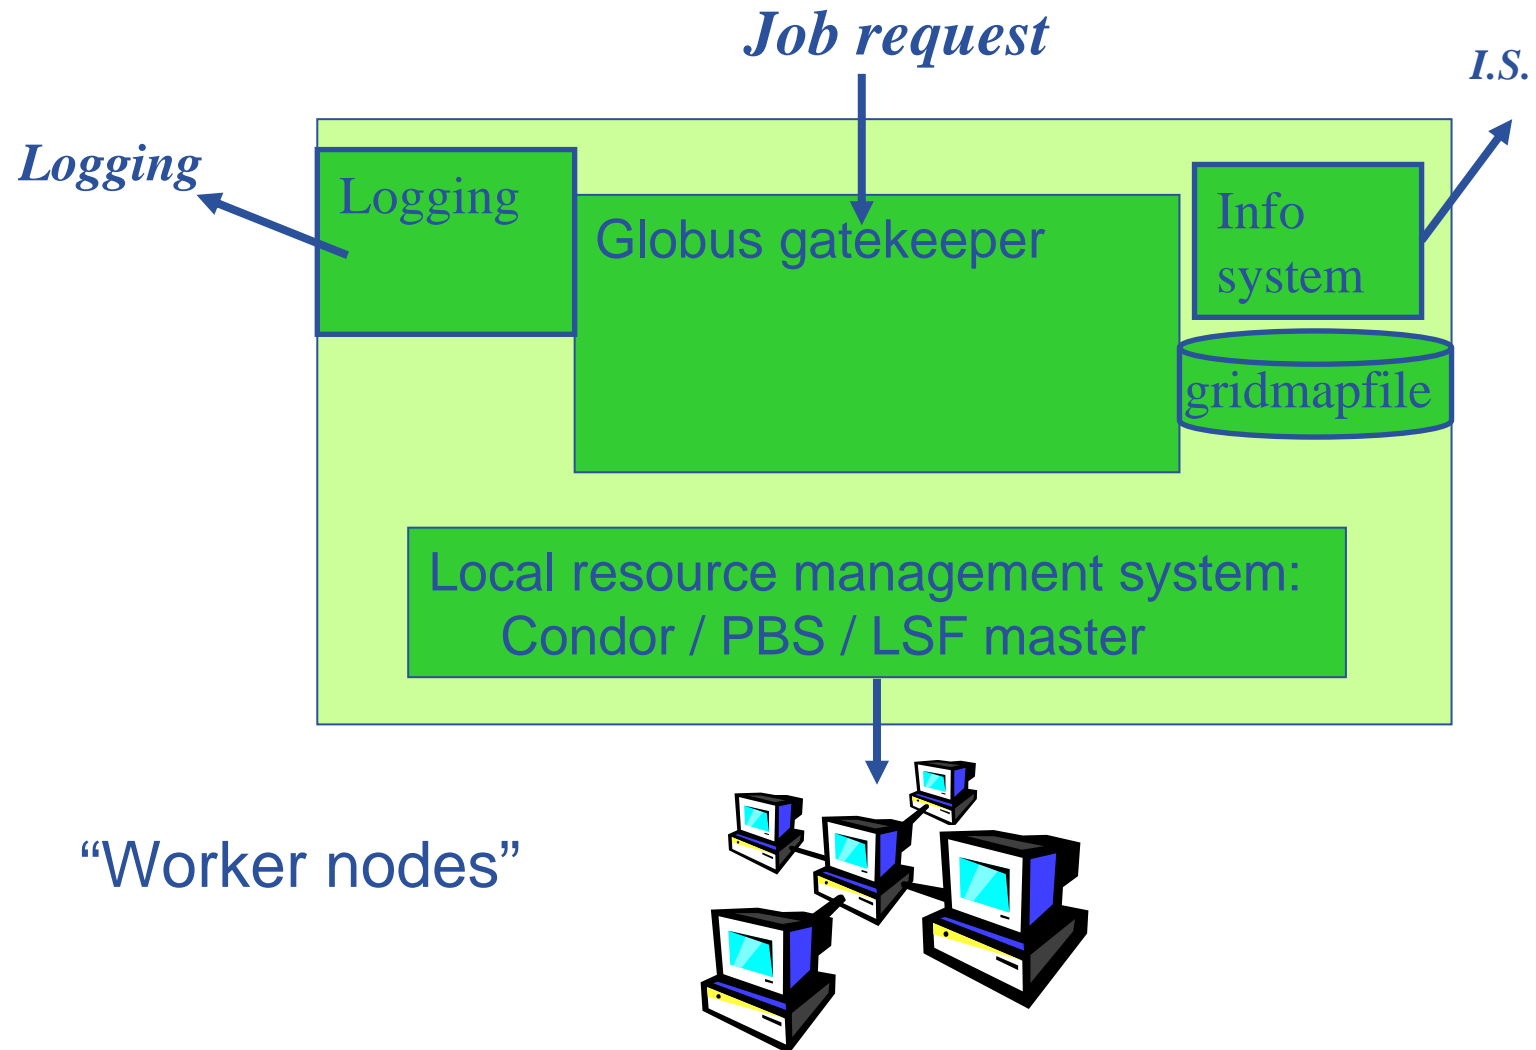
- digital signature of the CA

- User's identity has to be certified by one of the national *Certification Authorities (CAs)*
- Resources are also certified by CAs
- CAs are mutually recognized
<http://www.gridpma.org/>
- CAs each establish a number of people “registration authorities” RAs

- To support delegation: A delegates to B the right to act on behalf of A
- **proxy certificates extend X.509 certificates**
 - Short-lived certificates signed by the user's certificate or a proxy
 - Reduces security risk, enables delegation







- **Keep your private key secure – *on USB drive only***
- **Do not loan your certificate to anyone.**
- **Report to your local/regional contact if your certificate has been compromised.**
- **Do not launch a delegation service for longer than your current task needs.**

If your certificate or delegated service is used by someone other than you, it cannot be proven that it was not you.

Before VOMS

- User is authorised as a member of a single VO
- All VO members have same rights
- Gridmapfiles are updated by VO management software: map the user's DN to a local account
- **grid-proxy-init**

VOMS

- User can be in multiple VOs
 - Aggregate rights
- VO can have groups
 - Different rights for each
 - Different groups of experimentalists
 - ...
 - Nested groups
- VO has roles
 - Assigned to specific purposes
 - E.g. system admin
 - When assume this role
- Proxy certificate carries the additional attributes
- **voms-proxy-init**

VOMS – now in use on EGEE grid

- **Authentication based on X.509 PKI infrastructure**
 - Trust between **Certificate Authorities** (CA) and sites, CAs and users is established (offline)
 - CAs issue (long lived) **certificates** identifying sites and individuals (much like a passport)
 - Commonly used in web browsers to authenticate to sites
 - In order to reduce vulnerability, on the Grid user identification is done by using (short lived) **proxies** of their certificates
- **Proxies can**
 - Be **delegated** to a service such that it can act on the user's behalf
 - Include **additional attributes** (like VO information via the VO Membership Service VOMS)
 - Be stored in an **external proxy store** (myProxy)
 - Be **renewed** (in case they are about to expire)

- **Authentication**

- User obtains certificate from Certificate Authority
- Connects to UI by ssh (UI is the user's interface to Grid)
- Uploads certificate to UI
- Single logon – to UI - create proxy
- **Grid Security Infrastructure**

- **Authorisation**

- User joins Virtual Organisation
- VO negotiates access to Grid nodes and resources
- Authorisation tested by resource:

Credentials in proxy determine user's rights

