

## Grid Computing

### Security – Public Key Infrastructure, X.509 and Certificate Authorities (CA)

Paul A. Farrell  
Fall 2006

The Grid: Core Technologies  
Maozhen Li, Mark Baker  
John Wiley & Sons; 2005, ISBN 0-470-09417-6

Material from Dr John Watt, Grid Security Engineer  
National e-Science Centre, University of Glasgow  
Lecture slides by Amy Apon, Lawrie Brown,  
University of Aksansas

International Summer School  
on Grid Computing 2006  
Ischia (Naples) July 9-21, 2006

Paul A. Farrell 2006 KENT STATE Grid Computing 1

## Public Key Infrastructures (PKIs)

- PKIs provide a mechanism for privacy, integrity and authentication using public keys
  - Implemented with DIGITAL CERTIFICATES
    - Your UNIQUE virtualised identity
  - Issued by a CERTIFICATE AUTHORITY
    - Entity which administers certificates and issues them correctly
  - X509 (1988) is the standard for PKI certificates
    - Binds a globally unique X500 distinguished name to a public key
      - In reality, CAs tend to choose any name they want
    - Web browser compatible


Paul A. Farrell 2006 KENT STATE Grid Computing 2

## Why certificates?

- Certificates required for the large-scale use of public-key cryptography
- Anybody can create a public-private key pair
  - a malicious user can fool the originator into using their public key, and so get access to the information
- Solution: originator only trusts public keys that have been signed ("certified") by an authority
- Problem: user may not be familiar with another's certificate authority
  - certificate may also include a CA's public key signed by a "higher level" CA, which is more widely recognized
  - Gives rise to hierarchy of CAs

Paul A. Farrell 2006 KENT STATE Grid Computing 3

## Recall: Using a password to authenticate a client to a server



Paul A. Farrell 2006 KENT STATE Grid Computing 4

## Certificate-based Authentication

This is TLS (SSL) again. Recall:

Step 2: The server sends the client the SSL version number, random number Y, and its public key (packaged into a certificate)

Step 3: The client verifies that the server is who it says it is by examining the certificate. (Remember we said we would say more?)

## Certificate-based authentication details

Assume that the client has a private key and a certificate that contains the associated public key.

- The client generates random data
- It creates a digital signature of the data using the private key
- Client sends the data, digital signature, and its certificate across the network
- The server retrieves the packet containing the data, digital signature, and certificate
- Server extracts the client's public key from the certificate
- Server decrypts the digital signature using the client's public key
- Server compares the data with the decrypted signature to authorize the client

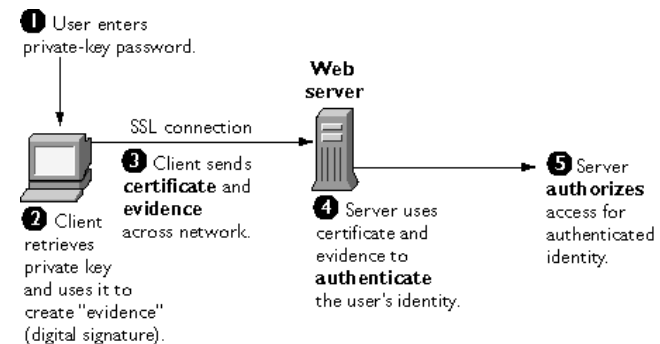
## Same process is used with user authentication

In the case of user authentication,

- The user enters a password that unlocks a local database and gives access to a private key.
- The client software retrieves the private key along with the associated public key certificate from the local database
- Continue with remaining client steps to authenticate a user to a server

No user password is sent across the network!

## Using a certificate to authenticate a client to a server



### Five types of certificates (used by Netscape)

---

1. Client SSL certificates
  - Used to identify clients to servers via SSL (client authentication).
  - Typically, the identity of the client is assumed to be the same as the identity of a human being, such as an employee in an enterprise.
2. Server SSL certificates
  - Used to identify servers to clients via SSL (server authentication).
  - Server authentication may be used with or without client authentication.
  - Server authentication is a requirement for an encrypted SSL session.

---

Paul A. Farrell 2006 KENT STATE Grid Computing 9

### Five types of certificates

---

3. S/MIME certificates
  - Used for signed and encrypted email.
  - As with client SSL certificates, the identity of the client is typically assumed to be the same as the identity of a human being, such as an employee in an enterprise.
  - A single certificate may be used as both an S/MIME certificate and an SSL certificate.
4. Object-signing certificates
  - Used to identify signers of Java code, JavaScript scripts, or other signed files.
5. Certificate Authority (CA) Certificates
  - Used to identify CAs.
  - Client and server software use CA certificates to determine what other certificates can be trusted.

---

Paul A. Farrell 2006 KENT STATE Grid Computing 10

### X.509 PKI

---

- Hierarchy of CAs top-down tree
- Top level (root) CA should be universally known
  
- Revocation of certificates
  - If private key has been exposed
  - Need to check Certificate Revocation List to check
  
- Format of certificate:
  - Public key being certified
  - Name
  - Validity period
  - Location (URL) of revocation list

---

Paul A. Farrell 2006 KENT STATE Grid Computing 11

### Contents of X.509 Certificates

---

- An X.509 v3 certificate binds a **distinguished name (DN)** to a public key
  
- A DN is a series of name-value pairs,
  - such as uid=doe
  - identify an entity--that is, the certificate **subject**.

---

Paul A. Farrell 2006 KENT STATE Grid Computing 12

## Example of X.509 Distinguished Name (DN)

uid=doe, e=doe@netscape.com, cn=John Doe,  
o=Netscape Communications Corp.,c=US where

- uid: user ID
- e: email address
- cn: the user's common name
- o: organization
- c: country

Paul A. Farrell 2006 KENT STATE Grid Computing 13

## X.509 Data Section

- The version number of the X.509 standard
- The certificate's serial number
  - Every certificate issued by a CA has a serial number that is unique among the certificates issued by that CA.
- Information about the user's public key
  - including the algorithm used and a representation of the key itself.
- The DN of the CA that issued the certificate.
- The period during which the certificate is valid
  - for example, between 1:00 p.m. on November 15, 1996 and 1:00 p.m. November 15, 1997
- The DN of the certificate subject
  - for example, in a client SSL certificate this would be the user's DN
- Optional **certificate extensions**

Paul A. Farrell 2006 KENT STATE Grid Computing 14

## X.509 Signature Section

- The cryptographic algorithm, or cipher, used by the issuing CA to create its own digital signature.
- The CA's digital signature, obtained by hashing all of the data in the certificate together and encrypting it with the CA's private key

Paul A. Farrell 2006 KENT STATE Grid Computing 15

## Full example in readable format

```
Certificate:
Data:
  Version: v3 (0x2)
  Serial Number: 3 (0x3)
  Signature Algorithm: PKCS #1 MD5 With RSA Encryption
  Issuer: OU=Ace Certificate Authority, O=Ace Industry, C=US
  Validity:
    Not Before: Fri Oct 17 18:36:25 1997
    Not After: Sun Oct 17 18:36:25 1999
  Subject: CN=Jane Doe, OU=Finance, O=Ace Industry, C=US
  Subject Public Key Info:
    Algorithm: PKCS #1 RSA Encryption
    Public Key:
      Modulus:
        00:ca:fa:79:98:8f:19:f8:d7:de:e4:49:80:48:e6:2a:2a:86:
        ed:27:40:4d:86:b3:05:c0:01:bb:50:15:c9:de:dc:85:19:22:
        43:7d:45:6d:71:4e:17:3d:f0:39:4b:5b:71:a8:51:a3:a1:00:
        98:ce:7f:47:50:2c:93:36:7c:01:6e:cb:89:06:41:72:b5:e9:
        73:49:38:76:ef:b6:8f:ac:49:bb:63:0f:9b:ff:16:2a:e3:0e:
        9d:3b:af:ce:9a:3e:48:65:de:96:61:d5:0a:11:2a:a2:80:b0:
        7d:c8:99:cb:0c:99:34:c9:ab:25:06:a8:31:ad:8c:4b:aa:54:
        91:f4:15
      Public Exponent: 65537 (0x10001)
  Extensions:
    Identifier: Certificate Type
    Critical: no
    Certified Usage:
      SSL Client
    Identifier: Authority Key Identifier
    Critical: no
    Key Identifier:
      12:f2:06:59:90:18:47:51:f5:89:33:5a:31:7a:e6:5c:fb:36:
      26:c9
  Signature:
    Algorithm: PKCS #1 MD5 With RSA E
```

Paul A. Farrell 2006 KENT STATE Grid Computing 16



## Types of Certificate

---

- Personal Certificates
  - Identify users
  - Allow digital signatures
- Host (server) certificates
  - To identify machines
- Subordinate CA certificates
  
- Certificate Policy statement (CPS)
  - Document on conditions under which certs are issued and level of assurance

---

Paul A. Farrell 2006 KENT STATE Grid Computing 21

## Certificate Authorities

---

- A CA also is in charge of *revoking* certificates
  - CA publishes a Certificate Revocation List
    - Download to your browser
    - Shows all invalid certificates in the organisation
  
- A CA **MUST** be explicitly trusted by the system
  - Trusted Root CAs list in Windows
  - In local SimpleCA for grid files are in
    - /etc/grid-security/certificates
  - Certificate cannot be used until the CA's root certificate has been accepted as trusted
    - Accepted very much like Software Licences i.e. nearly always!

---

Paul A. Farrell 2006 KENT STATE Grid Computing 22

## Digital Signatures

---

- CAs confirm the certificate's authenticity by digitally signing it
  - CA computes a hash of the certificate using an agreed (non-secret) algorithm
  - CA encrypts this hash with their private key and appends to bottom of certificate
  - Recipient computes their own hash of the info
  - Recipient decrypts the hash the CA sent (with the CA's public key) and compares with their own
    - Proves the CA signed the info and the info hasn't been tampered with
      - Encryption of the info is optional (for privacy)

---

Paul A. Farrell 2006 KENT STATE Grid Computing 23

## How CA Certificates are used to establish trust

---

- Certificate authorities (CAs) are entities that validate identities and issue certificates.
- They can be either independent third parties or organizations running their own certificate-issuing server software (such as the Netscape Certificate Server).
- A list of third-party certificate authorities is available at <https://certs.netscape.com/client.html>
  - Verisign most common used in US

---

Paul A. Farrell 2006 KENT STATE Grid Computing 24

## How CA Certificates are used to establish trust

- Any client or server software that supports certificates maintains a collection of **trusted CA certificates**.
- These CA certificates determine which other certificates the software can validate--in other words, which issuers of certificates the software can trust.
- In the simplest case, the software can validate only certificates issued by one of the CAs for which it has a certificate.
- It's also possible for a trusted CA certificate to be part of a chain of CA certificates, each issued by the CA above it in a certificate hierarchy.

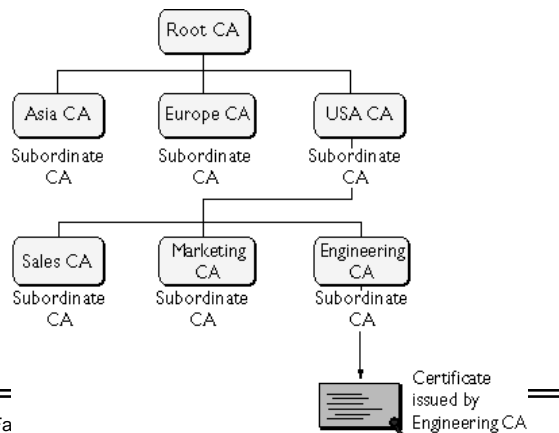
Paul A. Farrell 2006 KENT STATE Grid Computing 25

## CA Hierarchies

- In large organizations, it may be appropriate to delegate the responsibility for issuing certificates to several different certificate authorities. For example,
  - the number of certificates required may be too large for a single CA to maintain
  - different organizational units may have different policy requirements
  - or it may be important for a CA to be physically located in the same geographic area as the people to whom it is issuing certificates.
- It's possible to delegate certificate-issuing responsibilities to subordinate CAs.
- The X.509 standard includes a model for setting up a hierarchy of CAs

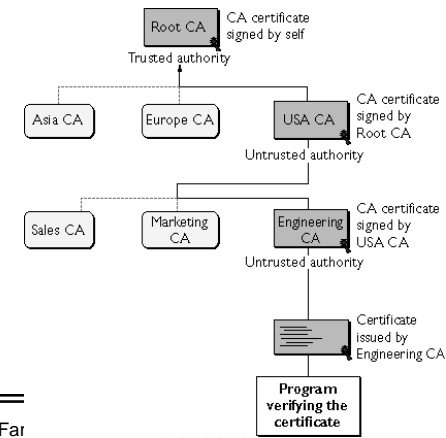
Paul A. Farrell 2006 KENT STATE Grid Computing 26

## X.509 CA Hierarchy Example



Paul A. Fa

## Certificate Chain Example



Paul A. Far

uting 28

## What happens in a certificate chain

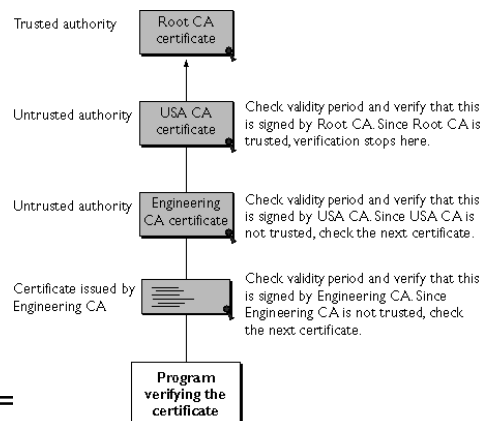
- Each certificate is followed by the certificate of its issuer.
- Each certificate contains the name (DN) of that certificate's issuer,
  - The same as the subject name of the next certificate in the chain.
  - the Engineering CA certificate contains the DN of the CA (that is, USA CA), that issued that certificate.
  - USA CA's DN is also the subject name of the next certificate in the chain.
- Each certificate is signed with the private key of its issuer.
- The signature can be verified with the public key in the issuer's certificate, which is the next certificate in the chain.
  - The public key in the certificate for the USA CA can be used to verify the USA CA's digital signature on the certificate for the Engineering CA.

## Verifying a Certificate Chain

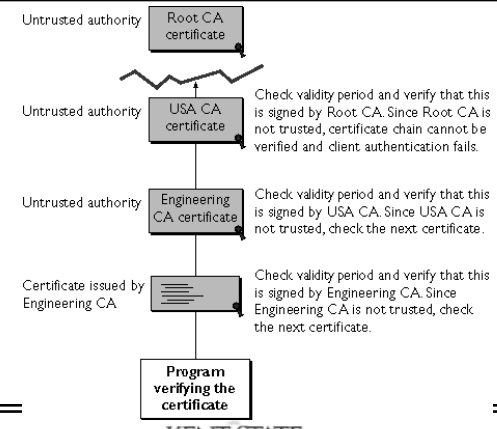
1. The certificate validity period is checked against the current time provided by the verifier's system clock.
2. The issuer's certificate is located. The source can be
  - either the verifier's local certificate database (on that client or server)
  - the certificate chain provided by the subject (for example, over an SSL connection).
3. The certificate signature is verified using the public key in the issuer's certificate.
4. If the issuer's certificate is trusted by the verifier in the verifier's certificate database, verification stops successfully here. Otherwise, the issuer's certificate is checked to make sure it contains the appropriate subordinate CA indication in the Netscape certificate type extension

Chain verification returns to step 1 to start again, but with this new certificate.

## A Valid Certificate Chain



## An Invalid Certificate Chain





## Proxies

---

- Proxies propagate through the grid by signing new versions of themselves
  - Using the proxy certificate and its private key, any number of new proxies may be spawned across each required resource.
- Achieves two goals
  - DELEGATION
    - Proxy acts on your behalf
  - SINGLE SIGN-ON
    - You only need to logon to the Grid once

## Firewalls

---

- Hardware or software component added to a network to prevent some communications forbidden by an organisation's administrative policy
- Two types
  - traditional firewall : a dedicated network device or computer on the boundary of a network
    - filters all traffic entering the networks
    - Works at some level of network stack
      - Network level : IP packet filtering
      - Transport (application) level : intercepting specific TCP socket traffic e.g. rlogin, telnet, www
  - personal firewall : a software application that filters traffic entering or leaving a single computer

## Proxy/Address Translation

---

- A proxy device, in hardware or software, can act as a firewall by responding to input packets) in the manner of an application, whilst blocking other packets
- It can be configured to pass certain packets to specific servers
- This reduces the risk of having internal systems compromised if there is misuse or misconfiguration
- Using private network addresses internally also helps