Introduction to Programming and Computing for Scientists

Oxana Smirnova

Lund University

Tutorial 4b: Grid certificates and jobs

Step 1: Install public Certificate Authority certificates

- Before doing anything on the Grid, you will need to obtain **IGTF Certificate Authorities certificates**
  - Packages are available from IGTF and some Grid repositories
    - The packages include Certificate Revocation Lists (CRLs)
  - Regular updates for CRL and IGTF packages must be in place
    - Usually happens automagically
  - **Our virtual machines and cluster have them already:**
    - Start the course virtual machine
    - Inspect `/etc/grid-security/certificates`
      - Hint: use `ls -al`

/etc/grid-security/certificates in the course Virtual Machine

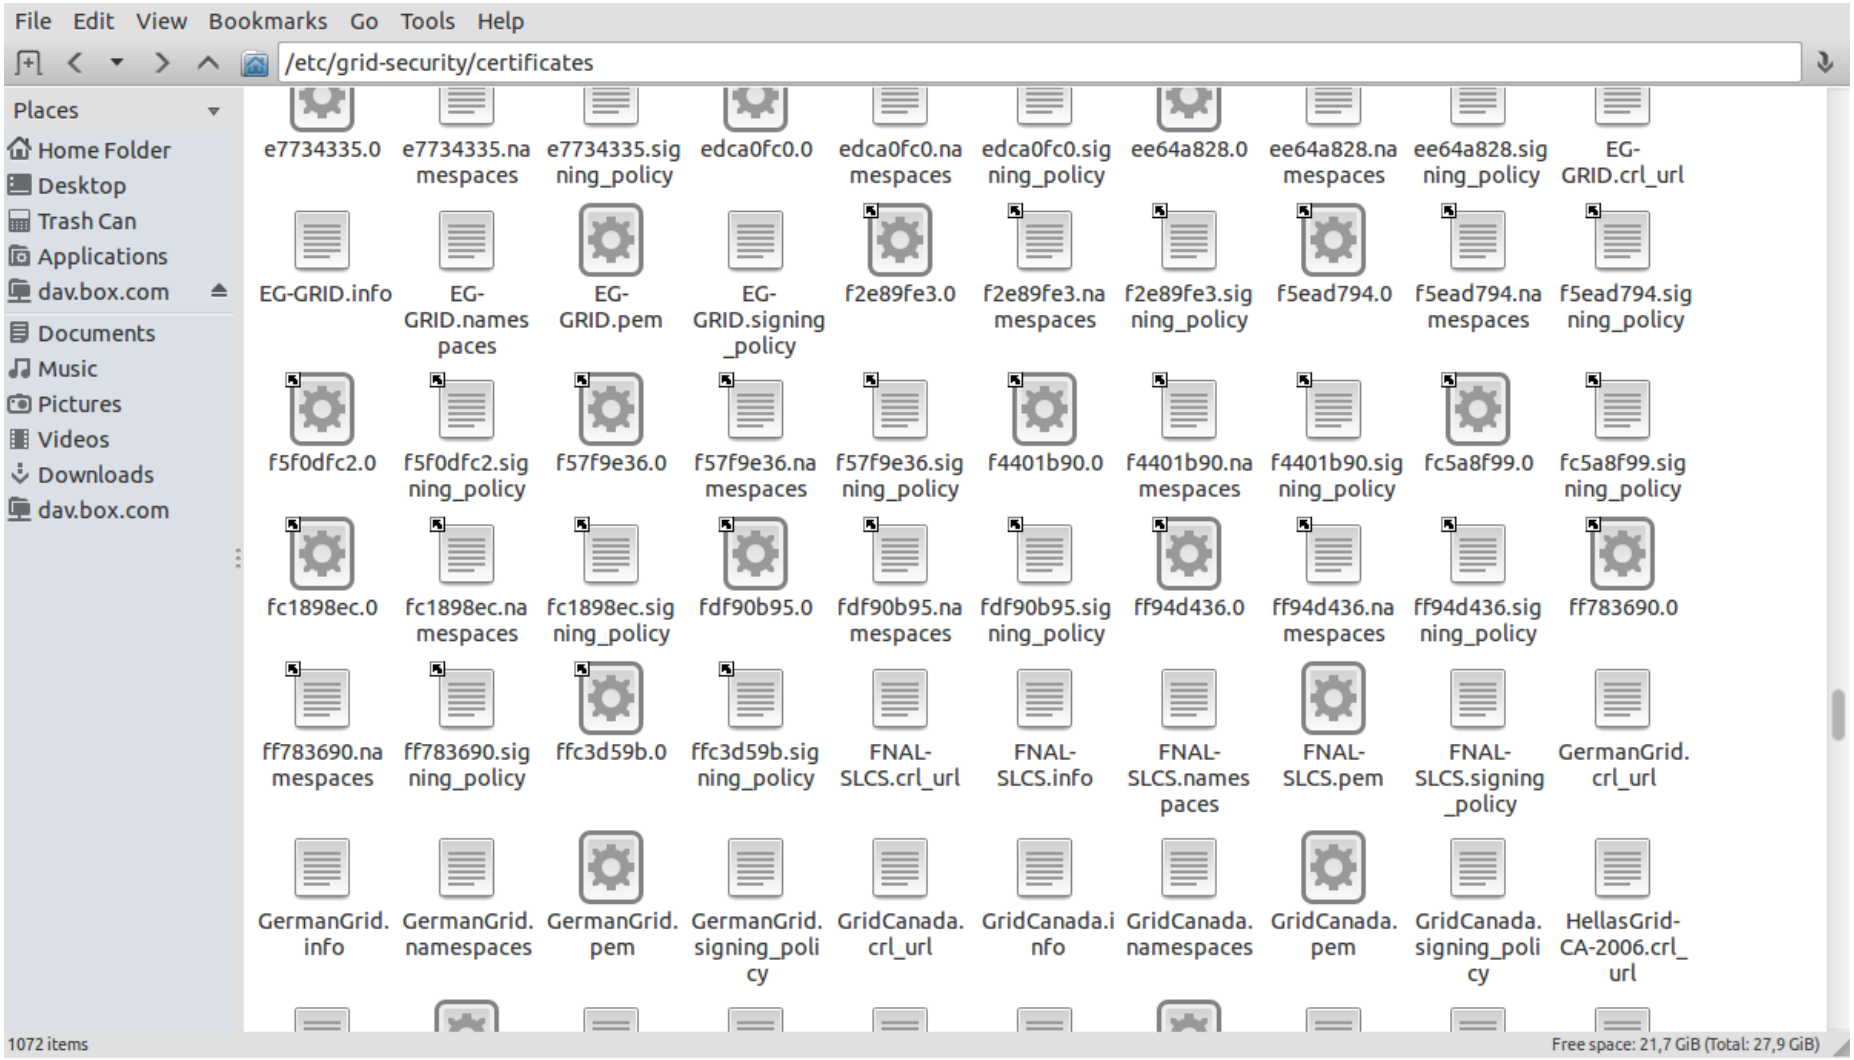

Step 2: get your own keys and certificates

There are two main ways of storing personal certificates:

- Two files: private key and public certificate
  - Grid uses PEM encoding for keys and certificates (ASCII)
  - Standard file names: **`userkey.pem`** and **`usercert.pem`**
    - Note: public key is inside the CA-signed certificate **`usercert.pem`**
- Single file: PKCS#12 formatted certificate, containing private and public keys, as well as CA signature and CRL info
  - PKCS#12 certificate (**`.p12`**) is used mostly by browsers, but can also replace PEM files in some Grid tools
  - One can convert PKCS#12 file to PEM files and vice versa

- Private keys must not be copied over the network!
- Private keys can not be stored in public machines like ours in this class!
- Because of this, we will log in to the Iridium cluster and create the keys there

Log in to Iridium and launch a Web browser there:

- From any virtual machine or a Linux machine log in to Iridium:

```
ssh -X yourlogin@pptest-iridium.lunarc.lu.se
```

- Launch a Web browser:

```
firefox &
```

- Some useful commands – a reminder:

```
ls -al
mkdir something
cd something
cp ~/dir/file1 file2
geany &
```

How to get a certificate

The easy way:

Google for “Grid certifikat Lund” and find a link to “Certifikat | LDC” - “inloggningsportalen”

Log in using your university credentials

Follow the instructions

The certificate in PKCS#12 format will be stored in your browser certificate store

You can export the certificate into a file (.p12) and extract PEM files, if necessary

How does it work:

Browser creates private and public keys for you

Public key is sent to TERENA CA, along with information that University provides about you

Signed public certificate is returned back to your browser, and merged with the private key to create a PKCS#12 certificate

Google is your friend

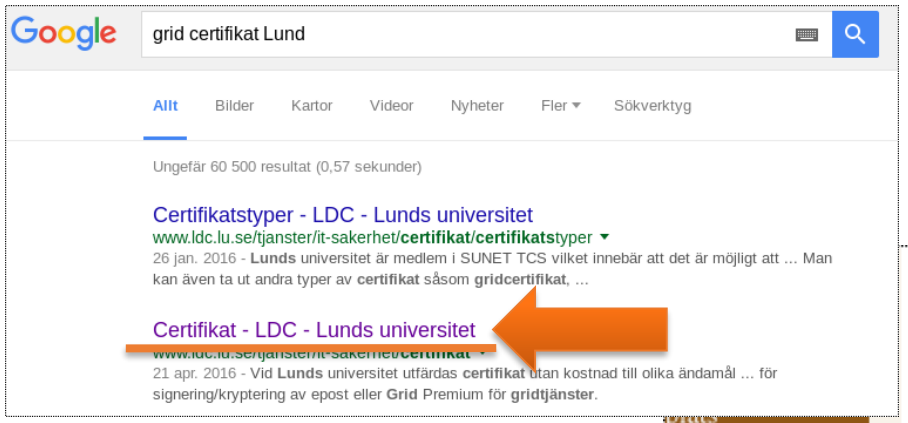

- Hint: Swedish SNIC Wiki also has good instructions!

Login to the CertCentral Portal

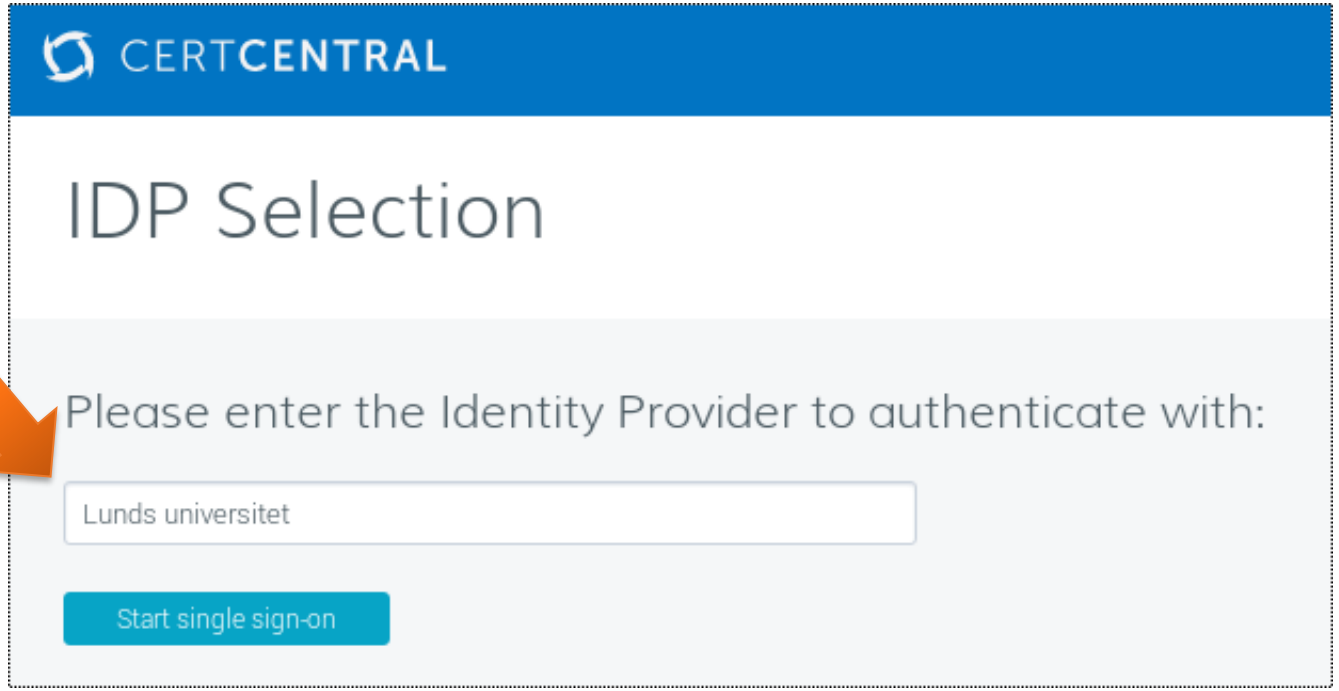

- Type Lunds universitet in the window
- If asked whether to remember Lunds universitet as Identity Provider, feel free to answer “yes”

Enter your LU credentials

Request “Grid Premium” certificate

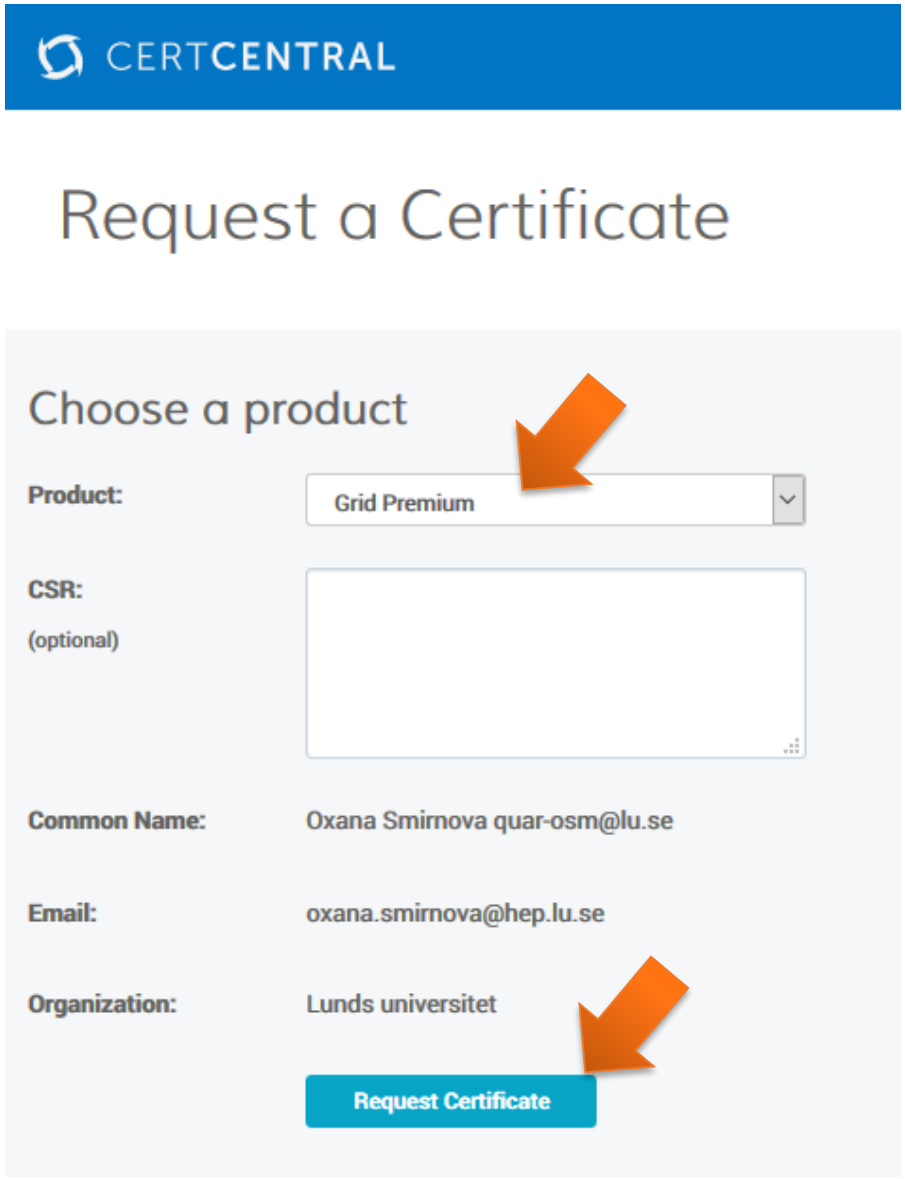

- Select “Grid Premium” Product
- If asked whether to trust the CA, feel free to tick “yes” everywhere

Backup (save) the certificate as a file

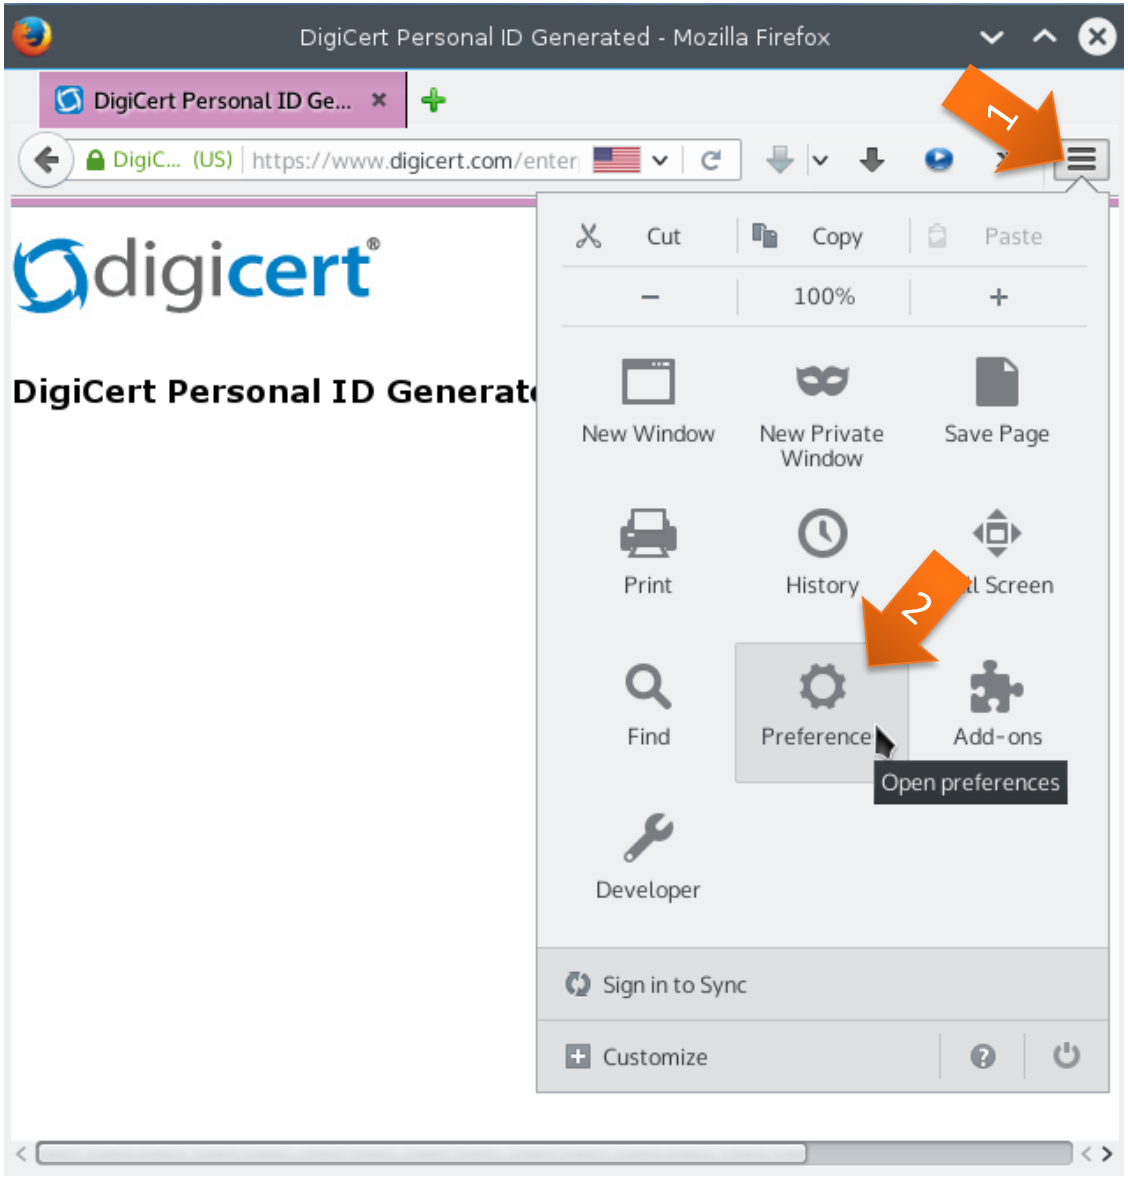

Go to Advanced – Certificates – View Certificates

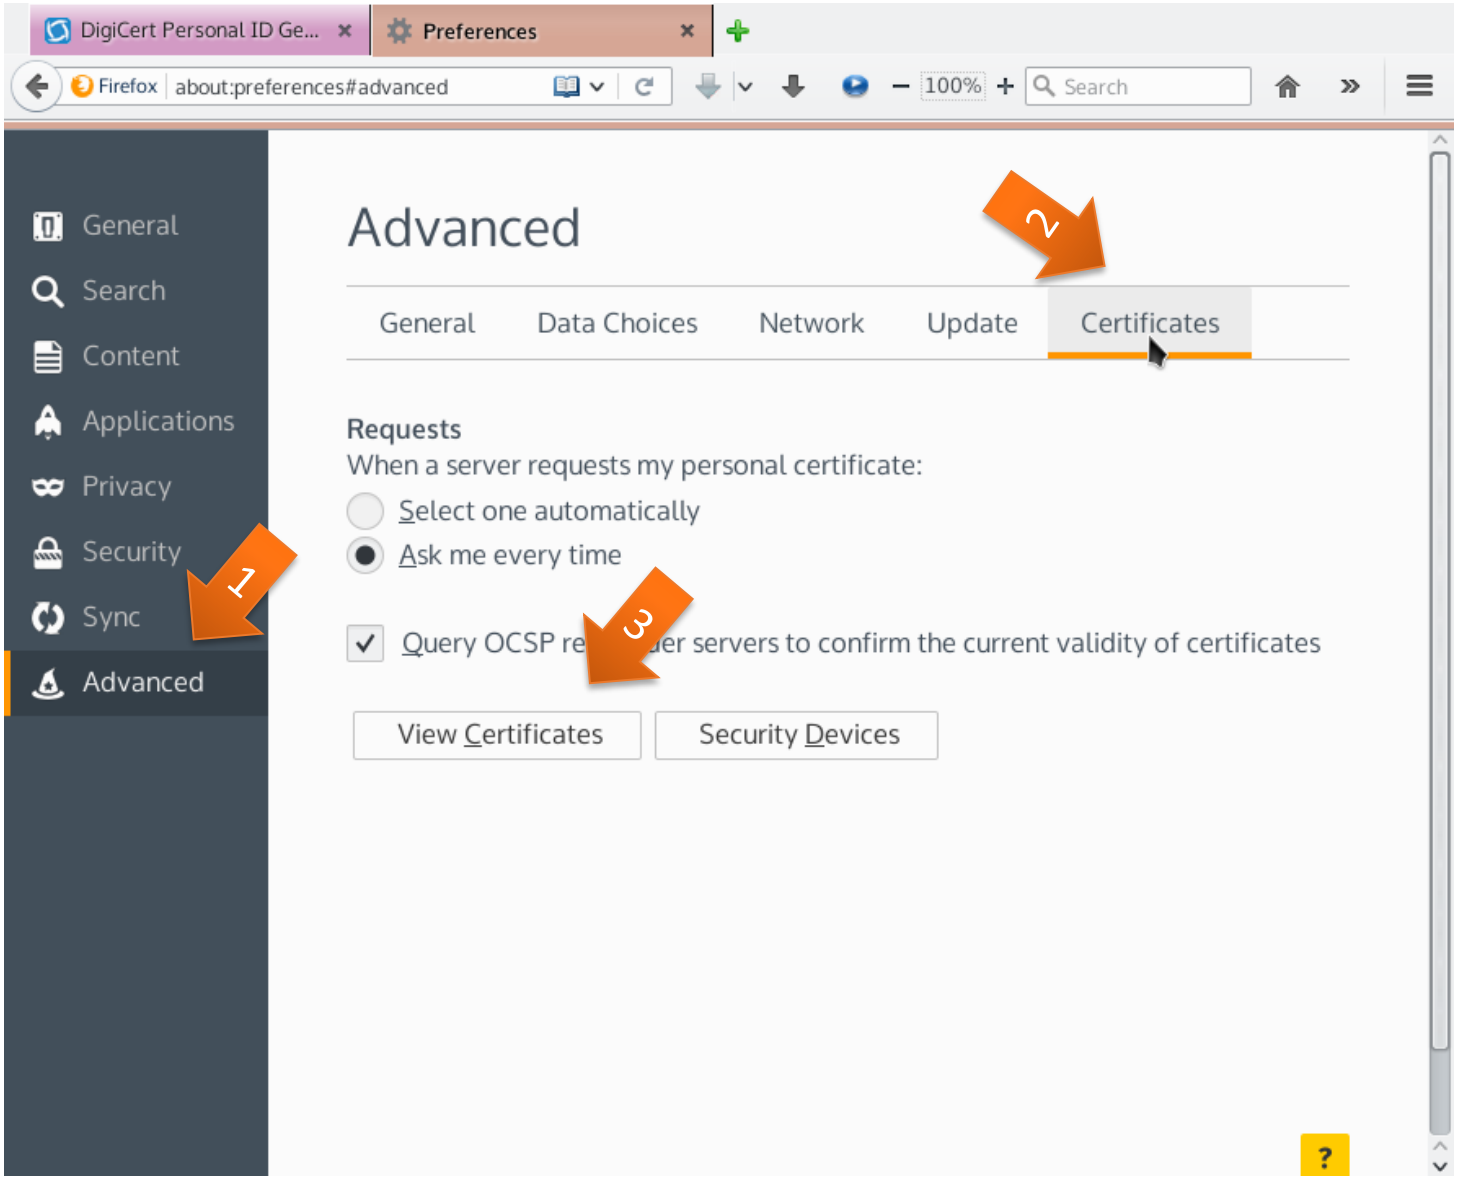

In Your Certificates, find yours, and use Backup to save it

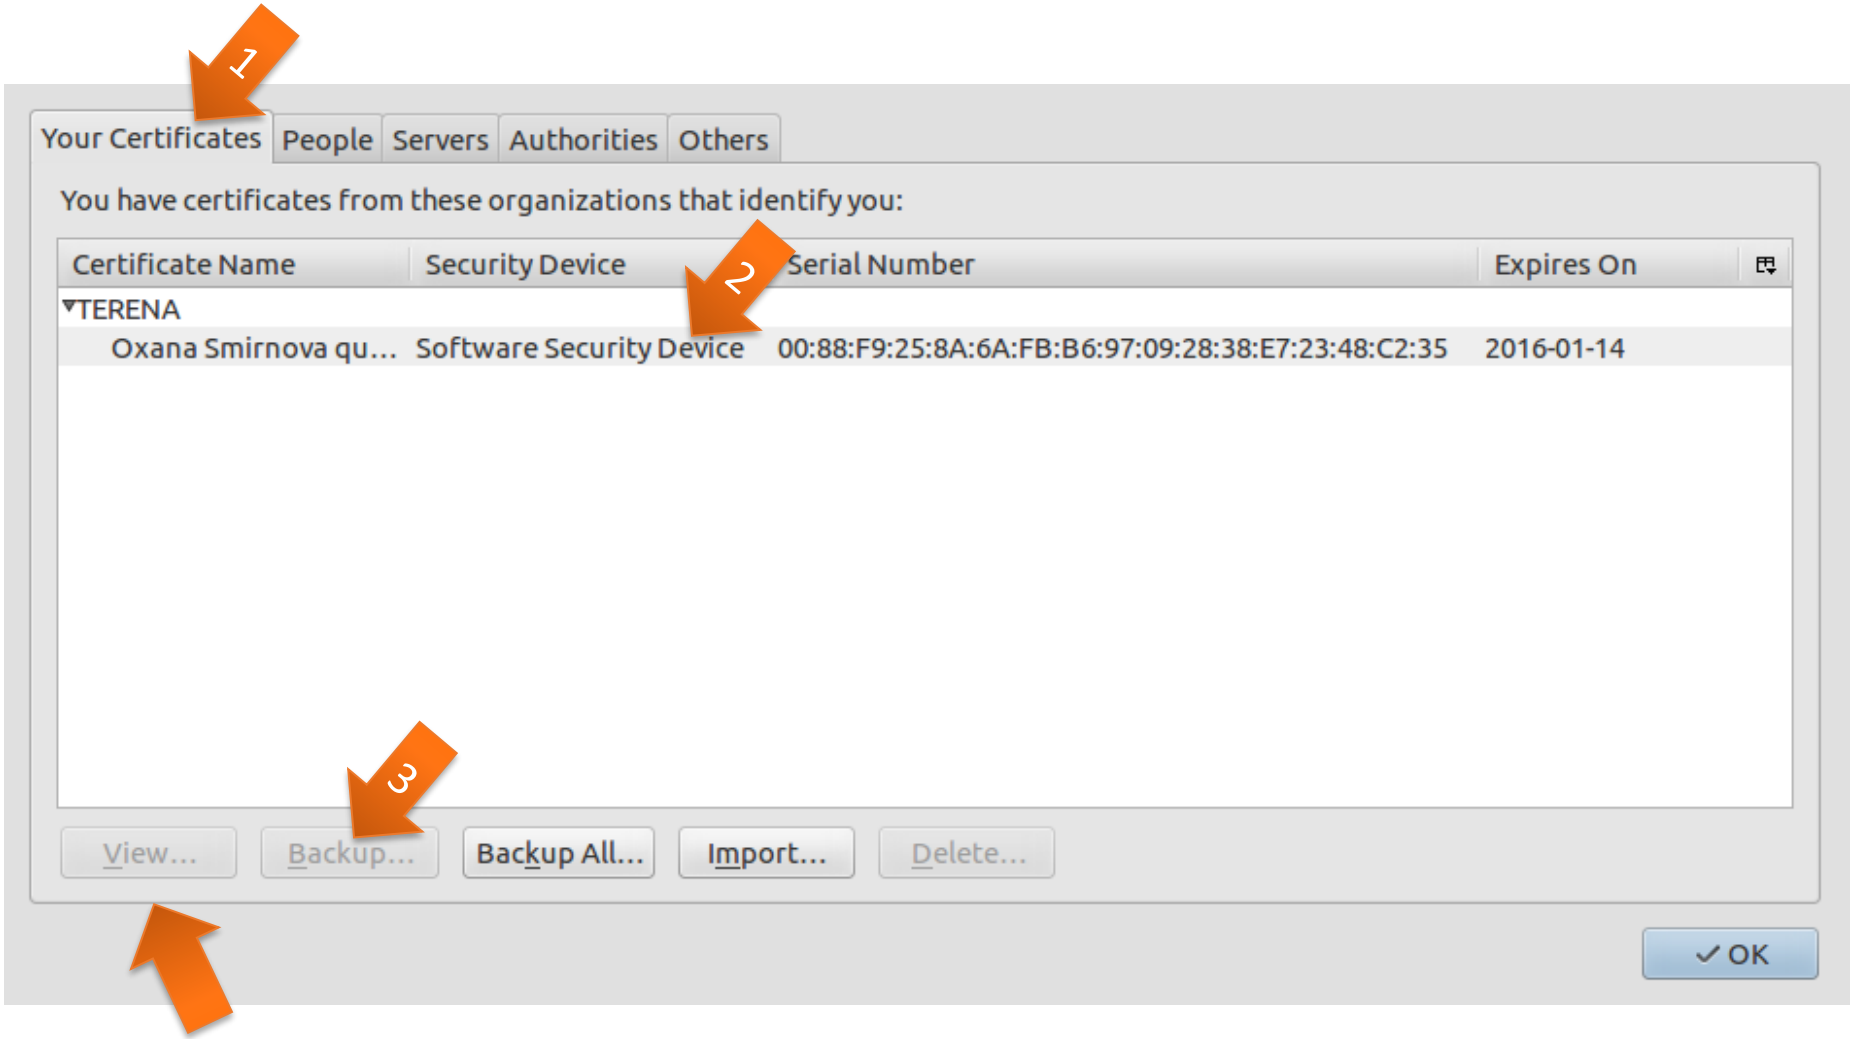

Take some time to view the certificate content

Backup the certificate as a file with extension .p12

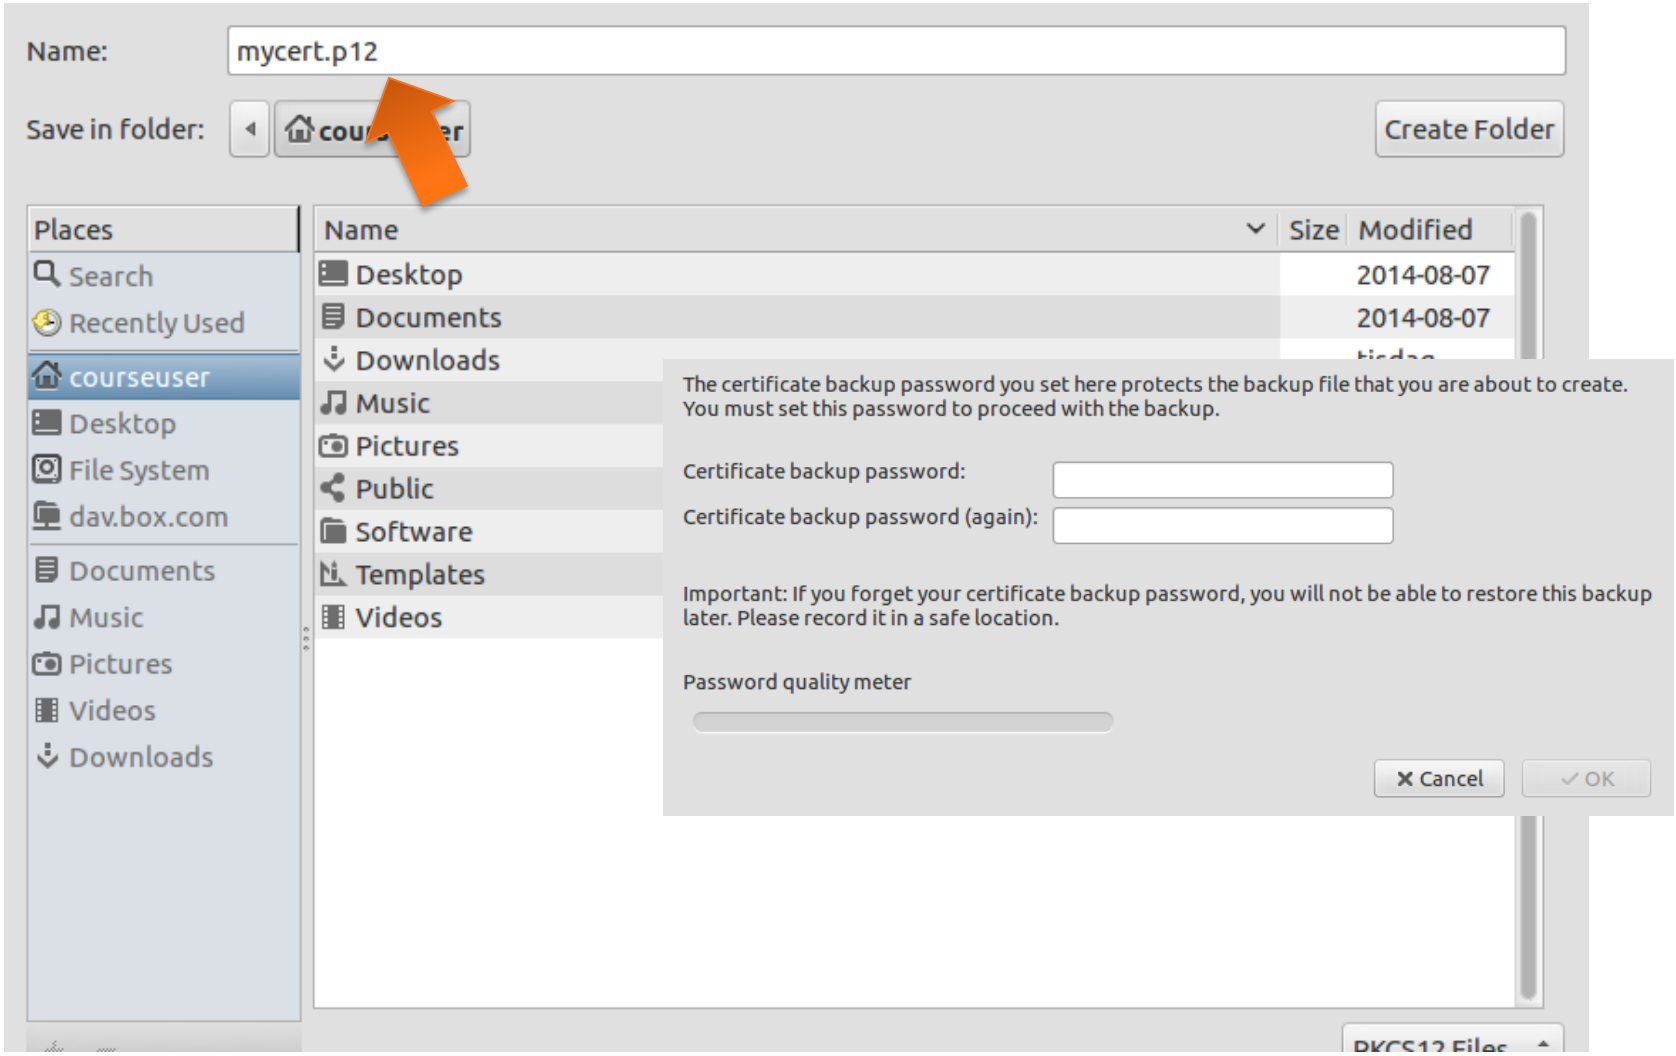

- Choose any password you like

Extract private and public keys

- Create a hidden directory **`~/.globus`**
  - This is the default location for Grid certificates
- Use **`openssl`** command to extract the keys inside **`~/.globus`** :
  - Private key:

```
openssl pkcs12 -nocerts -in mycert.p12 -out userkey.pem
```

  - Public key:

```
openssl pkcs12 -clcerts -nokeys -in mycert.p12 -out usercert.pem
```

  - Hint: Google for "grid certificate howto" to find where to copy-and-paste from
  - "*Import Password*" is the one you used to backup **`mycert.p12`** from the browser
  - "*pass phrase*" for PEM key is your own choice
    - You can use the same password in both cases
- Copy the files to **`userkey.pem`** and **`usercert.pem`** into **`~/.globus`**
  - Make sure that **`userkey.pem`** is readable only by you!
    - Hint: use **`ls -al`** and **`chmod`**

Summary of the steps:

```
oxana@bornholm:~/.globus >
oxana@bornholm:~/.globus >
oxana@bornholm:~/.globus > openssl pkcs12 -nocerts -in terena-15.p12 -out userkey-terena15.pem
Enter Import Password:
MAC verified OK
Enter PEM pass phrase:
Verifying - Enter PEM pass phrase:
oxana@bornholm:~/.globus > openssl pkcs12 -clcerts -nokeys -in terena-15.p12 -out usercert-terena15.pem
Enter Import Password:
MAC verified OK
oxana@bornholm:~/.globus > ls -al *terena15*
-rw-rw-r-- 1 oxana oxana 2279 нояб. 25 01:33 usercert-terena15.pem
-rw-rw-r-- 1 oxana oxana 2018 нояб. 25 01:32 userkey-terena15.pem
oxana@bornholm:~/.globus > chmod 400 userkey-terena15.pem
oxana@bornholm:~/.globus > ls -al *terena15*
-rw-rw-r-- 1 oxana oxana 2279 нояб. 25 01:33 usercert-terena15.pem
-r-------- 1 oxana oxana 2018 нояб. 25 01:32 userkey-terena15.pem
oxana@bornholm:~/.globus >
```

Step 3: Join a Virtual Organisation (Google for NorduGrid)

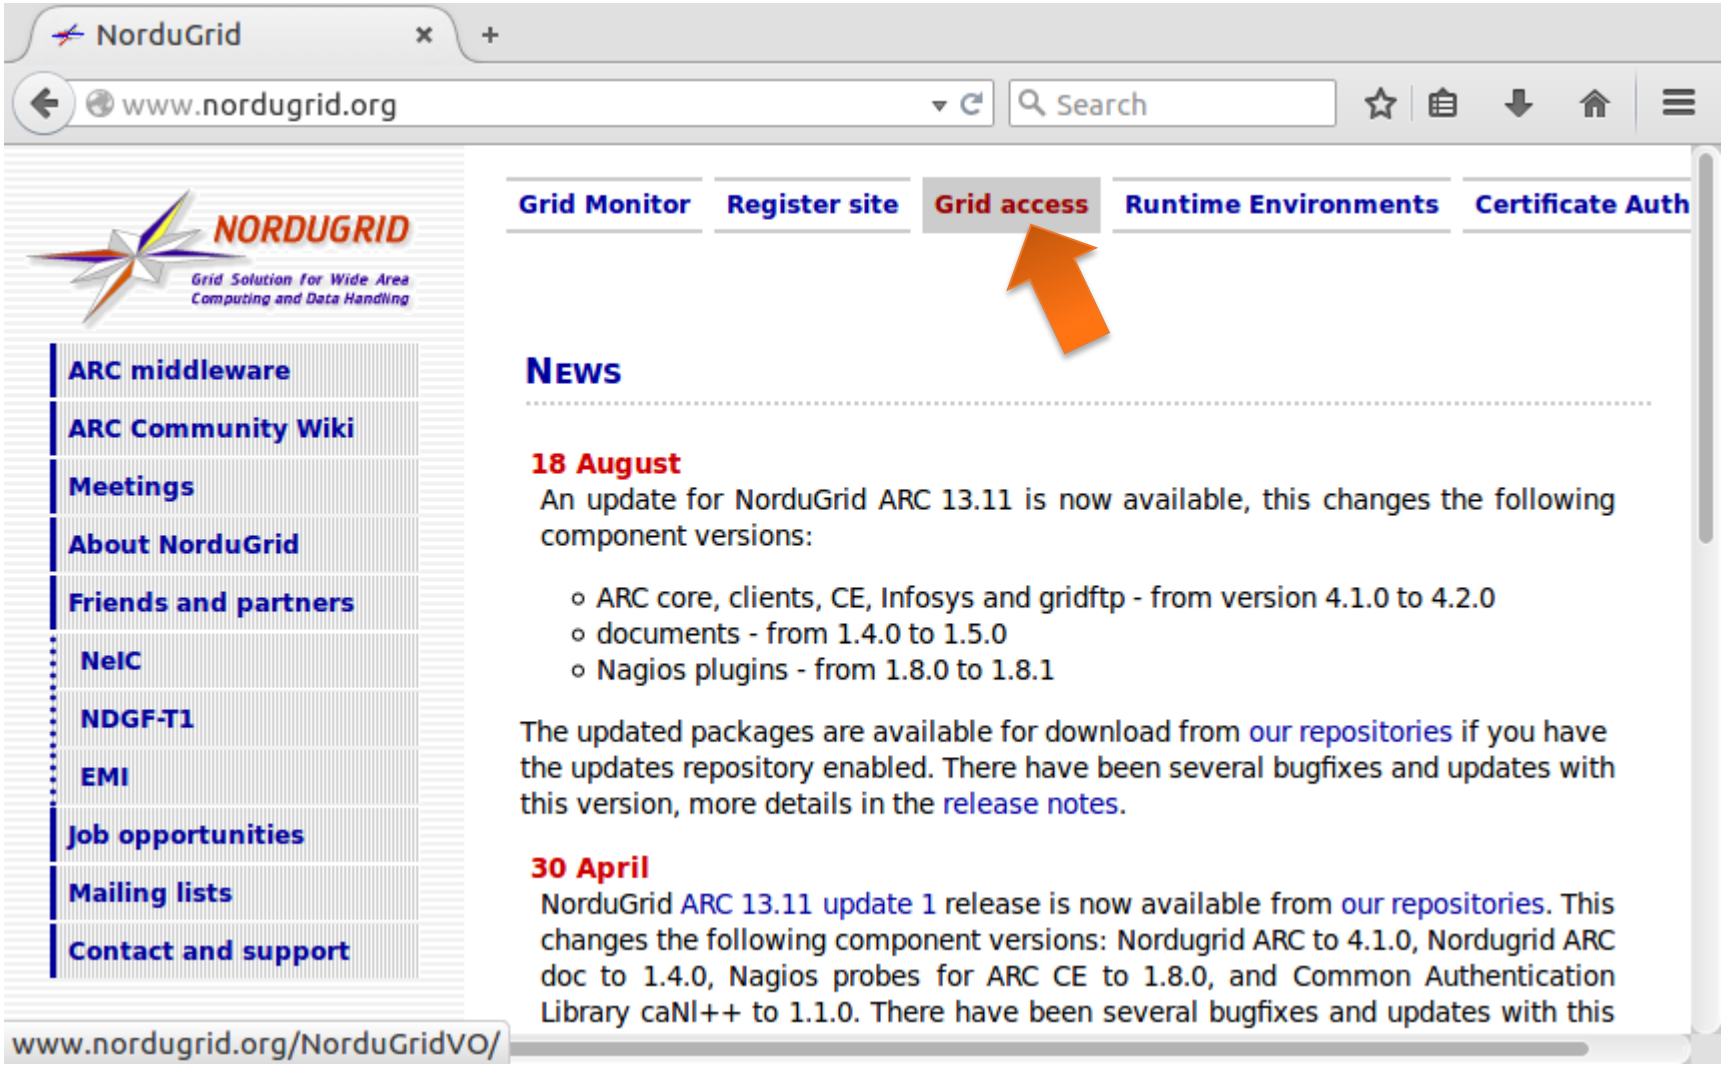

Find nordugrid.org and click Details

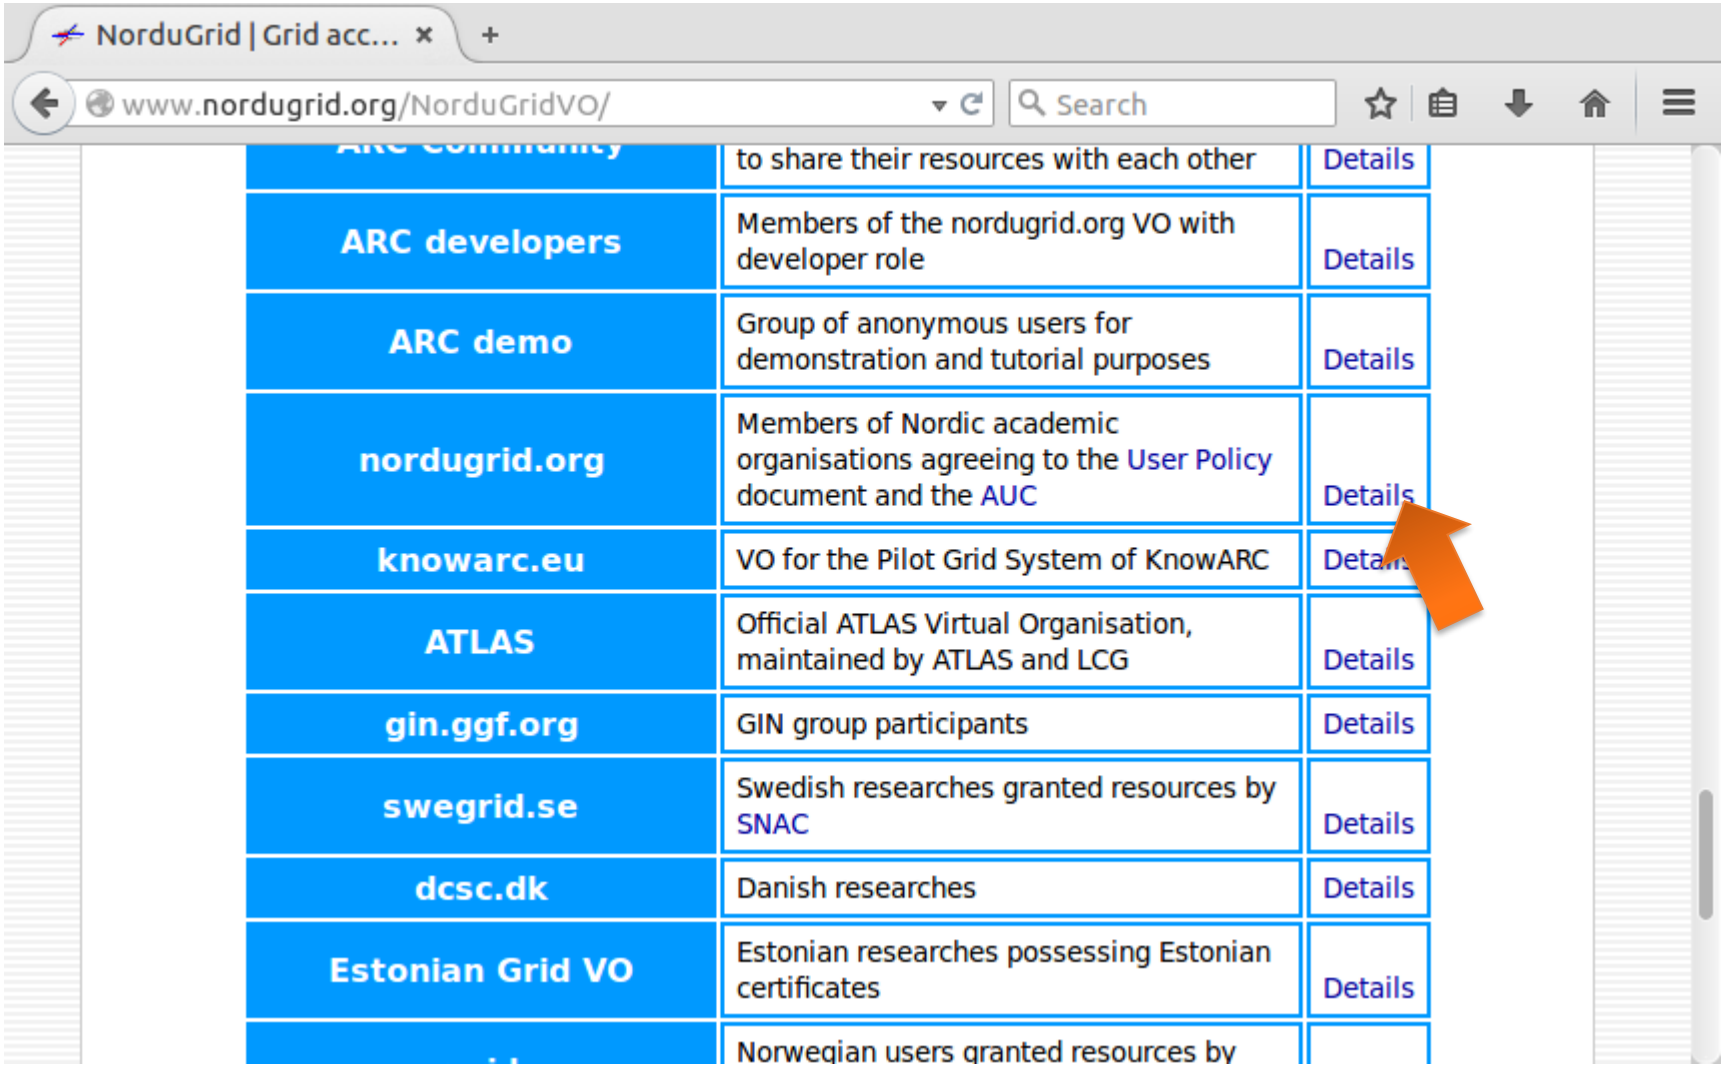

	to share their resources with each other	Details
ARC developers	Members of the nordugrid.org VO with developer role	Details
ARC demo	Group of anonymous users for demonstration and tutorial purposes	Details
nordugrid.org	Members of Nordic academic organisations agreeing to the User Policy document and the AUC	Details
knowarc.eu	VO for the Pilot Grid System of KnowARC	Detail
ATLAS	Official ATLAS Virtual Organisation, maintained by ATLAS and LCG	Details
gin.ggf.org	GIN group participants	Details
swegrid.se	Swedish researches granted resources by SNAC	Details
dcsc.dk	Danish researches	Details
Estonian Grid VO	Estonian researches possessing Estonian certificates	Details
	Norwegian users granted resources by	

Click “VOMS request interface”

The server requires your certificate:

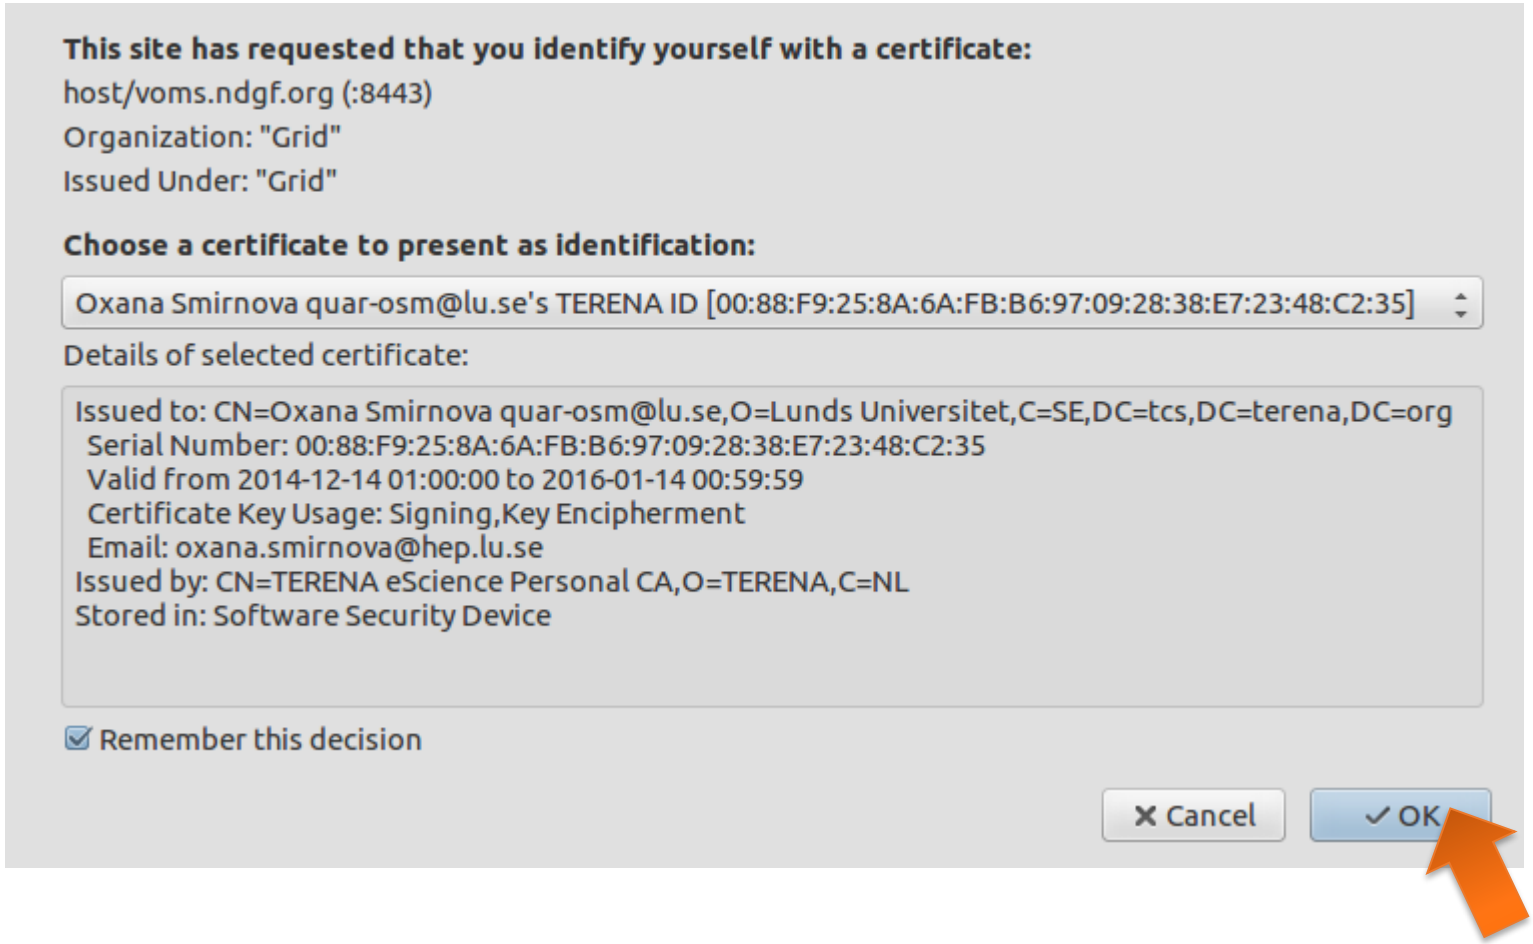

You have to establish trust with the server

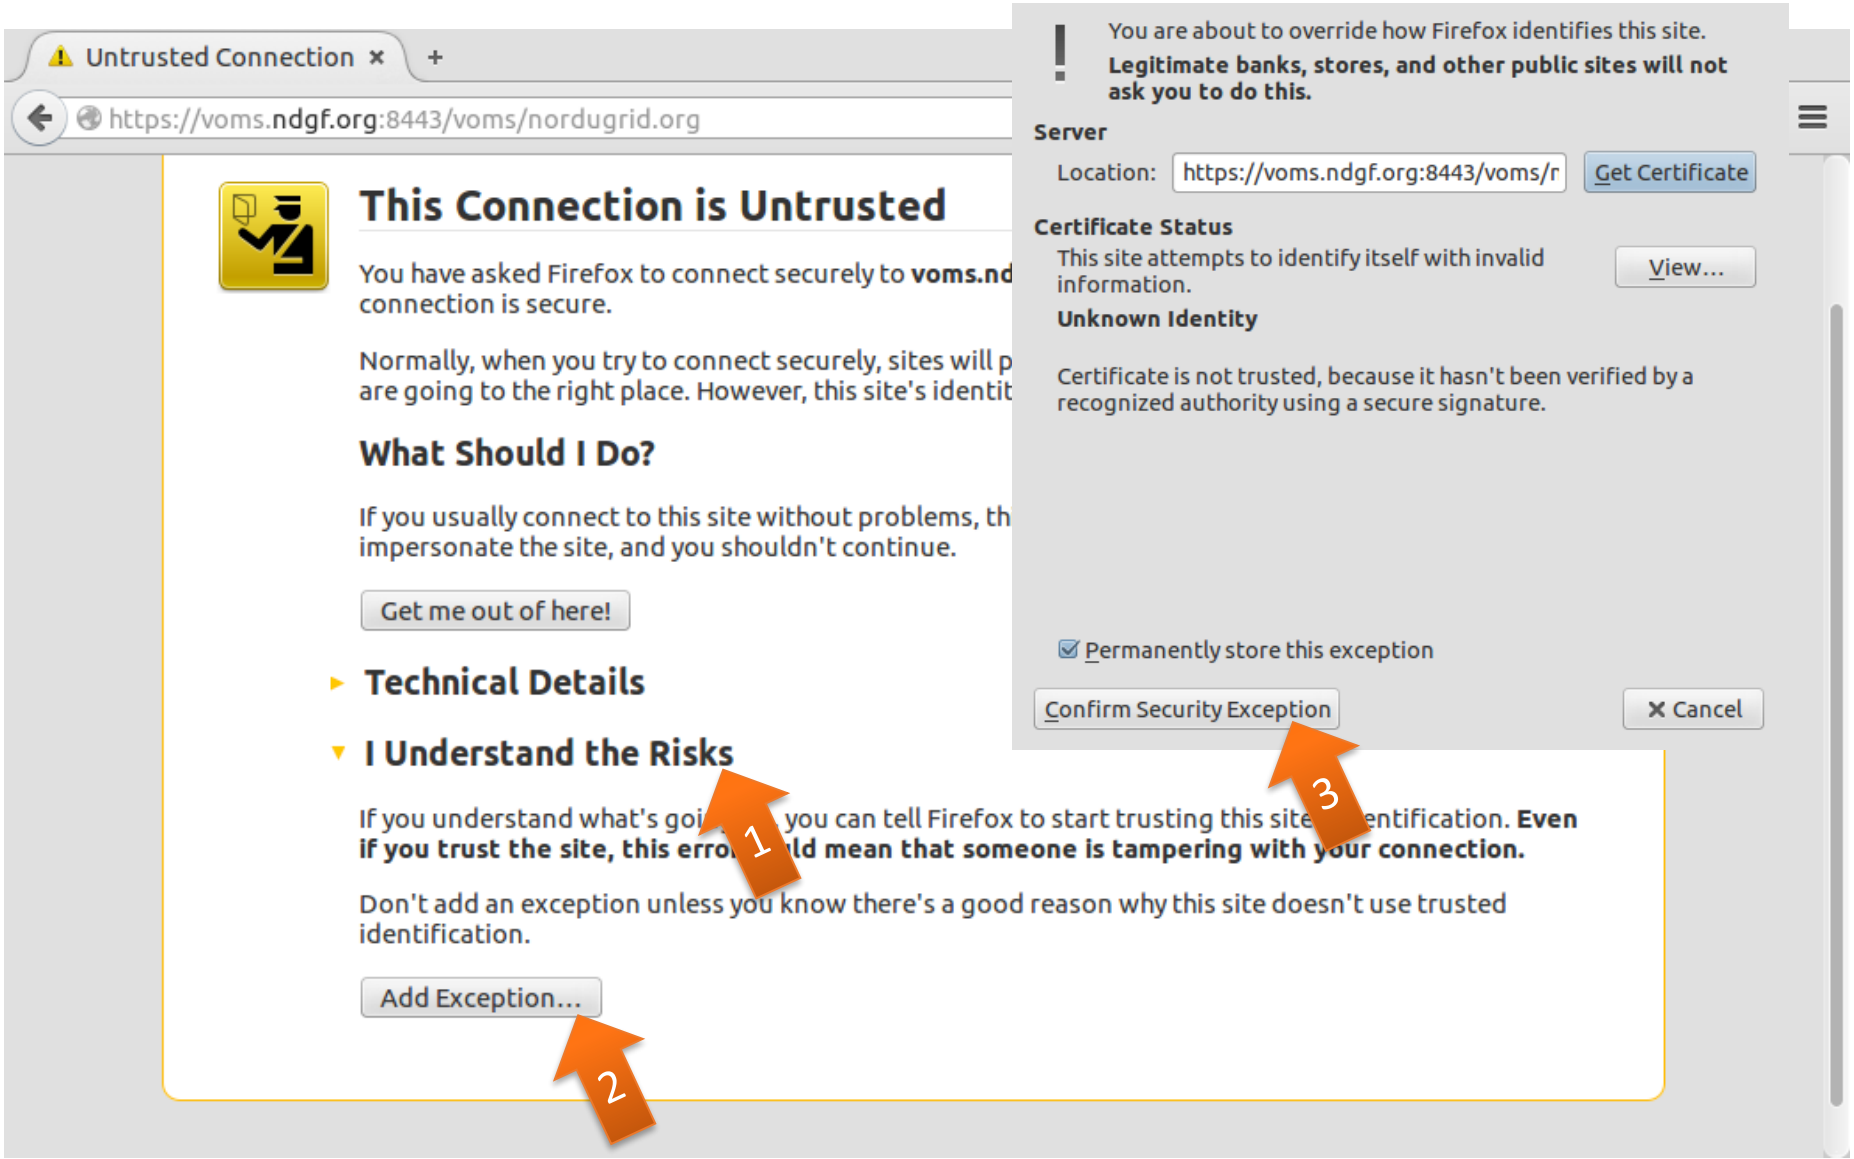

Fill in your details and request VO membership

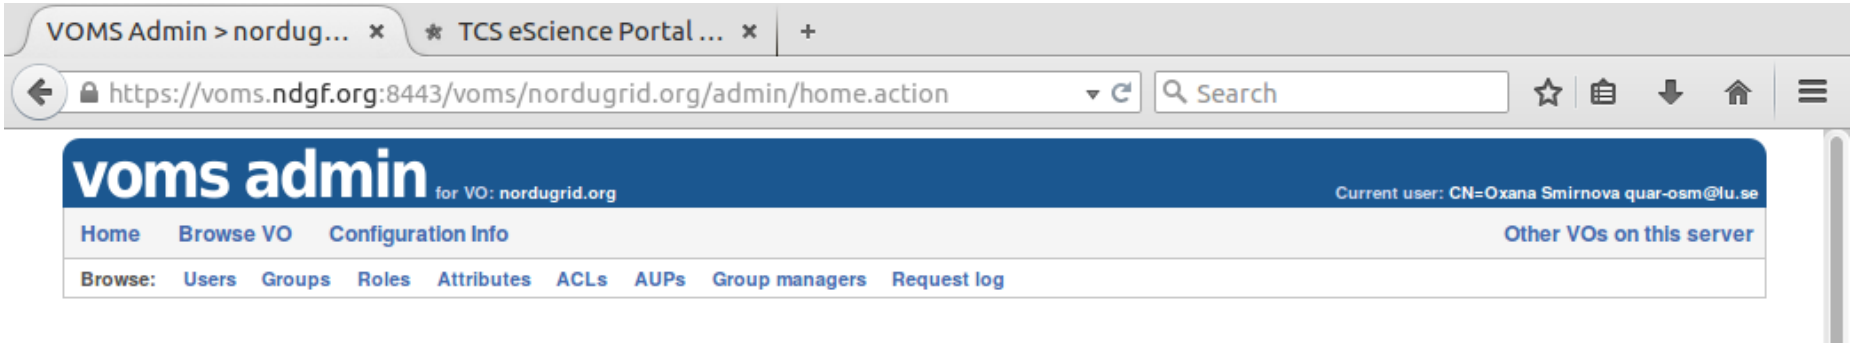

- Check your e-mail: VOMS will ask to confirm the request
  - Give teacher a moment to approve the request
- Ask the teacher if something is unclear
- What happens if your certificate is issued by a non-trusted CA?

On to Grid: create a proxy!

```
File  Edit  Tabs  Help
courseuser@Lubuntu-VirtualBox:~$ arcproxy
Enter pass phrase for private key:
Your identity: /DC=org/DC=terena/DC=tcs/C=SE/O=Lunds Universitet/CN=Oxana Smirnova quar-osm@lu.se
Proxy generation succeeded
Your proxy is valid until: 2014-12-15 12:32:44
courseuser@Lubuntu-VirtualBox:~$
```

- Simply type **`arcproxy`** and enter your Grid password (*PEM pass phrase* for the private key)
- Type **`arcproxy -I`** to check information about this proxy:
  - Issuer (must be you, of course)
  - Your Distinguished Name (DN), a.k.a Identity
  - Validity period

What actually `arcproxy` does?

- A **new** private/public key pair is created for each proxy
  - When a proxy expires, a new one must be created to continue working
    - Default expiration time is 24 hours
- A proxy is then constructed of:
  1. Public certificate (with public key embedded)
     - Certificate contains modified owner's Distinguished Name (has "*proxy*" appended to the name)
       - Owner's DN: **`/C=UK/O=Grid/OU=CenterA/L=LabX/CN=john doe`**
       - Proxy DN: **`/C=UK/O=Grid/OU=CenterA/L=LabX/CN=john doe/CN=proxy`**
     - Certificate is signed by the proxy owner's **real** private key
     - Certificate contains validity period
  2. Private key
  3. Optionally, Attribute Certificates – extensions containing additional information

The tale of two proxies

- A user always has to create a proxy certificate **P1**
  - Technically, it can be sent to the server, but it is a security breach
- Any Grid server (e.g. a Computing Element) creates itself a **delegated proxy P2** for each user request:
  1. Server generates a **new** private/public key pair (yes, that's a 3rd one...)
  2. Server returns the generated public key as a certificate request to the user
  3. User's tool signs that public key and inserts user information (DN etc), thus generating a public certificate. It uses the private key of proxy **P1** for performing signing operation.
     - It can also use the actual private key, but that will require entering password every time!
  4. User's tool sends the signed public certificate back to the server
  5. Server adds generated private key to that certificate and creates a delegated proxy **P2**

What's the use of VOMS

- A Grid user must become a member of a Virtual Organisation (VO)
  - *VOMS* is the most common VO management system
- A Grid cluster administrator gets the list of authorised users from the VOMS database
- VOMS can add extra VO information to your proxy, if necessary
  - For example, your VO role, group etc
  - You should use **`arcproxy`** with special command-line options to request such extra information to be added
    - We won't try it today

Summary of the proxies

- Luckily, all authentication and delegation procedures are a part of the protocol, you only need to create a proxy
- You have to create a proxy before every Grid activity
- Proxies expire quickly!
  - Resist temptation to create long-living proxy: this will undermine your security
- Proxies may have special extensions, specific to Virtual Organisations
- If you forget your Grid password (PEM pass phrase), and even the browser Import Password, you will have to request a new certificate

Workflow: Grid vs PC/cluster

PC/cluster		Grid	
Log in via SSH	• Different logins on different machines	Create proxy	• One for all machines
Familiarize with the environment	• OS, installed software, storage space, batch system, sysadmin etc	Create a Grid job description document	• Generalization of batch scripts, plus input/output data location etc
Customize the environment	• Pathes, environment variables, own software, scripts, data files etc	Test a couple of jobs, fix job description	
Prepare for batch submission	• Interactive execution of short jobs, optimization of batch scripts	Submit jobs to the Grid, check their status	• Same commands for all machines
Submit jobs to the batch system, check their status	• Different batch systems (or none) on different machines	Watch output appearing in the desired location	• Or fetch it manually
Log out			
Log in later to fetch the output			

Simplest Grid job submission

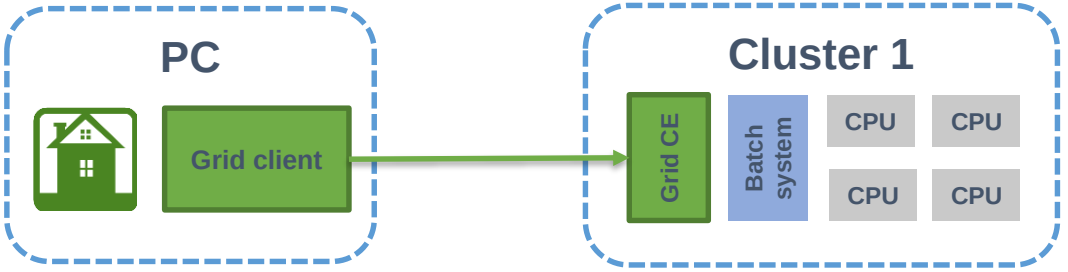

- Your Grid client should:
  - Create a proxy:
    - **`arcproxy`**
  - Submit the job description document to the cluster:
    - **`arcsub -c arc-iridium.lunarc.lu.se hello_grid.xrsl`**
      - **`arcsub`** will refuse submission if the cluster does not meet job requirements

See the ARC Clients manual for info about all ARC client commands:
http://www.nordugrid.org/documents/arc-ui.pdf

- The CE on the cluster should:
  - Check whether you are authorised
  - Fetch input file (if requested)
  - Convert job description to a batch script and start a batch job
  - Upload output file (if requested)

Simplest Grid job description: hello_grid.xrsl

```
&( executable = "/bin/echo" )
 ( arguments = "hello grid" )
 ( stdout = "stdout_file" )
 ( stderr = "error_file" )
 ( cputime = "13" )
 ( gmlog = "grid_log" )
 ( jobname = "hello_grid" )
```

- Yes, this is yet another language:
  XRSL – eXtended Resource Specification Language
  - File extension is **`.xrsl`**
- XRSL is not a standard language, but no standard exists
  - There are many other Grid languages and meta-languages
  - XRSL is an ARC extension of the original Grid language by Globus
    - It was actually modelled on the LDAP database query language
    - Is a list of attribute-value pairs

Main attributes of job description

Job attribute description	Attribute name (XRSL)	Example value
Main executable (binary or script)	**executable**	MyAnalysis.py
Arguments of the executable	**arguments**	-i input.dat -o output.dat
Input files	**inputfiles**	https://store.lu.se/physlab/2012/file1.dat
Output files	**outputfiles**	https://store.lu.se/physlab/2014/file1.dat
Standard input file	**stdin**	stdin.txt
Standard output file	**stdout**	stdout.txt
Standard error file	**stderr**	stderr.txt
Time (used by CPU)	**cputime**	1 hour
Memory (maximum needed, Mbytes)	**memory**	1000
Disk space (maximum needed, Mbytes)	**disk**	1000
Job name	**jobname**	My data analysis
Number of slots (cores) for the job	**count**	36
and many others: ARC job description language XRSL has 37 attributes, see http://www.nordugrid.org/documents/xrsl.pdf		

Create and submit your hello_grid.xrsl

- Prepare job description for the “Hello Grid” task:
  - Use Geany (or Vim, or Emacs) to create a file **`hello_grid.xrsl`**
  - Use at least the following XRSL attributes: **`executable`**, **`arguments`**, **`jobname`**
    - Hint: copy the example from the previous slide
- Submit your first Grid job to our Iridium cluster:
  - First, make sure you have a valid proxy:

    ```
    arcproxy -I
    ```

  - Use the **`arcsub`** command with explicit cluster selection:

    ```
    arcsub -c arc-iridium.lunarc.lu.se hello_grid.xrsl
    ```

  - Find the returned **job ID** (a long string that looks like a URL)
  - Check the job’s status:
    `arcstat <jobid>`
  - Check what the job “session directory” looks like on the cluster:
    `arcls <jobid>`
  - Check what does the job print out:
    `arccat <jobid>`

Manipulate the jobs: kill, retrieve

- Submit a couple more jobs
  - You may want to change the job names in **`hello_grid.xrsl`**
  - Or you may even want to change what do the jobs produce
- Check the status of all your jobs:

```
arcstat -a
```

- Terminate some of them and check the status afterwards:

```
arckill -k <jobid>
arcstat <jobid>
```

  - **`-k`** here means “keep the job files”, otherwise they will be wiped out
- Retrieve job results (download job output):

```
arcget -k <jobid>
```

  - **`-k`** here has the same meaning as for **`arckill`**
- Find where the downloaded files are, and look what is there
  - Inspect the content of the **`gmlog`** sub-directory: it has files useful for error diagnostics and debugging

If you have some time left

- Find the hidden directory **`~/.arc`** and file **`client.conf`** therein
- Open **`client.conf`** in Geany (or any other editor)
- Find blocks **`[registry/index1]`** , **`[registry/index2]`** etc and uncomment them and their content
  - Save **`client.conf`** and quit the editor
- Try to submit **`hello_grid.xrsl`** to the entire Grid

```
arcsub hello_grid.xrsl
```