

# Table of Contents

<b>Data Transfer Tools for Theory QCD Application.....</b>	<b>1</b>
Grid Tools.....	1
xrootd.....	1
Install and setup xrootd client tools "from scratch" (latest development version).....	1
Extra steps in case of problems (what I did to get it running on ubuntu 9.10).....	1
Transfer your files.....	2
Some transfer tests done by me.....	2
Setting up proxy server at CERN with public access.....	2
<b>Grid tools at CERN.....</b>	<b>4</b>
.....	<b>5</b>
<b>srmcp from dcache.....</b>	<b>6</b>
tests with big files from CERN intranet.....	6
initial problems.....	8
srmcp on my ubuntu box.....	8
<b>lcg-cp.....</b>	<b>9</b>

# Data Transfer Tools for Theory QCD Application

## Grid Tools

We investigate xrootd and FTS for this purpose.

### xrootd

**lxplus.cern.ch** : the tools are already installed in: /usr/bin

**SLC4,SLC5**: `yam install xrootd-client`

**SL4,SL5**: find and install rpm in linuxsoft.cern.ch extras repository

Alternatively see below to install "from scratch".

### Install and setup xrootd client tools "from scratch" (latest development version)

Minimal steps (should be fine for SLC4,SLC5,RHEL):

Get the installer from the xrootd homepage

- look for **xrd-installer** at <http://project-arda-dev.web.cern.ch/project-arda-dev/xrootd/site/download.html>

I installed the latest CVS development version which should in few weeks become the production version.

```
=bash xrd-installer --install
```

By default the client tools get installed in `~/xrdserver`. You need this:

```
export LD_LIBRARY_PATH=~/xrdserver/lib:$LD_LIBRARY_PATH
export PATH=~/xrdserver/bin:$PATH
```

You are ready to go.

### Extra steps in case of problems (what I did to get it running on ubuntu 9.10)

If installer fails to compile xrootd packages then check if you have all needed packages on the system (including the dev versions):

- ♦ <http://alien.cern.ch/twiki/bin/view/AliEn/HowToInstallXrootdNew>

To authenticate via Kerberos 5 make sure that you have krb5 package installed and configured to include CERN.CH:

If needed, add CERN.CH to kerberos realms in `[realms]` section in the configuration file `/etc/krb5.conf`:

```
[realms]
  CERN.CH = {
    default_domain = cern.ch
    kpasswd_server = afskrb5m.cern.ch
    admin_server = afskrb5m.cern.ch
    kdc = afsdb2.cern.ch
    kdc = afsdb3.cern.ch
```

```
kdc = afsdb1.cern.ch

v4_name_convert = {
  host = {
    rcmd = host
  }
}
```

I also made it default:

```
[libdefaults]
    default_realm = CERN.CH
```

### Transfer your files

```
kinit user@CERN.CH
```

You are ready to play with Castor transfer:

- **upload:** `xrdcp /etc/hosts root://castorpublic.cern.ch//castor/cern.ch/user/m/moscicki/tmp.test`
- **download and dump on screen:** `xrdcp root://castorpublic.cern.ch//castor/cern.ch/user/m/moscicki/tmp.test -`
- **download in verbose mode and pipe to /dev/null:** `xrdcp -d 3 -f root://castorpublic//castor/cern.ch/theory/pcqcd/L12T12_b5.8458_id1/cond/rome/L12T12_b5.84 /dev/null`
- for WAN transfers use `-s15` option to use up to 15 parallel streams to speedup transfers

Interactive command line client:

- `xrd castorpublic.cern.ch`
- browse the tree with `dirlist` and `cd`

## Some transfer tests done by me

Create a big random file on a local disk:

```
pcarda75: dd if=/dev/urandom of=20GB.RANDOM.TEST bs=20M count=1000
1000+0 records in
1000+0 records out
20971520000 bytes (21 GB) copied, 5908.11 s, 3.5 MB/s
```

Copy it to Castor using `xrootd` (intranet):

```
time xrdcp 20GB.RANDOM.TEST root://castorpublic.cern.ch//castor/cern.ch/user/m/moscicki/20GB.RAND
Disabling apmon monitoring since env variable APMON_CONFIG was not found
[xrootd] Total 20000.00 MB      |=====| 100.00 % [10.9 MB/s]

real    32m6.181s
user    0m9.925s
sys     2m37.810s
```

## Setting up proxy server at CERN with public access

Instructions are here:

[/afs/cern.ch/sw/arda/install/theory/xrootd](https://afs.cern.ch/sw/arda/install/theory/xrootd)

# Grid tools at CERN

Enable debug output for SOAP clients (srmcp,lcg-cp):

```
export CGSI_TRACE=1
```

Setup environment:

```
source /afs/cern.ch/project/gd/LCG-share/current/etc/profile.d/grid_env.sh
```

### Example:

```
grid-proxy-init  
globus-url-copy gsiftp://lxfsrk5801.cern.ch:2811///castor/cern.ch/user/m/moscicki/tmp.test file:/
```

# srmcp from dcache

Here is the trick:

```
srmcp -srm_protocol_version 2 srm://srm-public.cern.ch:8443/srm/managerV2?SFN=/castor/cern.ch/use
```

Warning: the destination path is relative!

## tests with big files from CERN intranet

Expiry of grid proxy:

```
[lxplus250] /afs/cern.ch/user/m/moscicki > date
Thu Feb 25 15:39:01 CET 2010
[lxplus250] /afs/cern.ch/user/m/moscicki > time srmcp -srm_protocol_version 2 srm://srm-public.cern.ch:8443/srm/managerV2?SFN=/castor/cern.ch/use
WARNING: SRM_PATH is defined, which might cause a wrong version of srm client to be executed
WARNING: SRM_PATH=/afs/cern.ch/project/gd/LCG-share/3.1.38-0/d-cache/srm
[main] ERROR gsi.CertificateRevocationLists - CRL /afs/cern.ch/project/gd/LCG-share2/certificate
java.security.GeneralSecurityException: [JGLOBUS-16] CRL data not found.
    at org.globus.gsi.CertUtil.loadCrl(CertUtil.java:526)
    at org.globus.gsi.CertificateRevocationLists.loadCrl(CertificateRevocationLists.java:174)
    at org.globus.gsi.CertificateRevocationLists.reload(CertificateRevocationLists.java:129)
    at org.globus.gsi.CertificateRevocationLists$DefaultCertificateRevocationLists.refresh(CertificateRevocationLists.java:200)
    at org.globus.gsi.CertificateRevocationLists.getDefault(CertificateRevocationLists.java:200)
    at org.globus.gsi.CertificateRevocationLists.getDefaultCertificateRevocationLists(CertificateRevocationLists.java:200)
    at org.globus.gsi.gssapi.GlobusGSSContextImpl.verifyChain(GlobusGSSContextImpl.java:717)
    at org.globus.gsi.gssapi.GlobusGSSContextImpl.initSecContext(GlobusGSSContextImpl.java:517)
    at org.globus.gsi.gssapi.net.GssSocket.authenticateClient(GssSocket.java:107)
    at org.globus.gsi.gssapi.net.GssSocket.startHandshake(GssSocket.java:145)
    at org.globus.gsi.gssapi.net.GssSocket.getOutputStream(GssSocket.java:166)
    at org.apache.axis.transport.http.HTTPSender.writeToSocket(HTTPSender.java:440)
    at org.apache.axis.transport.http.HTTPSender.invoke(HTTPSender.java:138)
    at org.apache.axis.strategies.InvocationStrategy.visit(InvocationStrategy.java:32)
    at org.apache.axis.SimpleChain.doVisiting(SimpleChain.java:118)
    at org.apache.axis.SimpleChain.invoke(SimpleChain.java:83)
    at org.apache.axis.client.AxisClient.invoke(AxisClient.java:165)
    at org.apache.axis.client.Call.invokeEngine(Call.java:2784)
    at org.apache.axis.client.Call.invoke(Call.java:2767)
    at org.apache.axis.client.Call.invoke(Call.java:2443)
    at org.apache.axis.client.Call.invoke(Call.java:2366)
    at org.apache.axis.client.Call.invoke(Call.java:1812)
    at org.dcache.srm.v2_2.SrmSoapBindingStub.srmStatusOfGetRequest(SrmSoapBindingStub.java:200)
    at sun.reflect.GeneratedMethodAccessor17.invoke(Unknown Source)
    at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25)
    at java.lang.reflect.Method.invoke(Method.java:585)
    at org.dcache.srm.client.SRMClientV2.handleClientCall(SRMClientV2.java:178)
    at org.dcache.srm.client.SRMClientV2.srmStatusOfGetRequest(SRMClientV2.java:449)
    at gov.fnal.srm.util.SRMGetClientV2.start(SRMGetClientV2.java:324)
    at gov.fnal.srm.util.SRMDispatcher.work(SRMDispatcher.java:817)
    at gov.fnal.srm.util.SRMDispatcher.main(SRMDispatcher.java:368)
[main] ERROR gsi.CertificateRevocationLists - CRL /afs/cern.ch/project/gd/LCG-share2/certificate
java.security.GeneralSecurityException: [JGLOBUS-16] CRL data not found.
    at org.globus.gsi.CertUtil.loadCrl(CertUtil.java:526)
    at org.globus.gsi.CertificateRevocationLists.loadCrl(CertificateRevocationLists.java:174)
    at org.globus.gsi.CertificateRevocationLists.reload(CertificateRevocationLists.java:129)
    at org.globus.gsi.CertificateRevocationLists$DefaultCertificateRevocationLists.refresh(CertificateRevocationLists.java:200)
    at org.globus.gsi.CertificateRevocationLists.getDefault(CertificateRevocationLists.java:200)
    at org.globus.gsi.CertificateRevocationLists.getDefaultCertificateRevocationLists(CertificateRevocationLists.java:200)
    at org.globus.gsi.gssapi.GlobusGSSContextImpl.verifyChain(GlobusGSSContextImpl.java:717)
    at org.globus.gsi.gssapi.GlobusGSSContextImpl.initSecContext(GlobusGSSContextImpl.java:517)
    at org.globus.gsi.gssapi.net.GssSocket.authenticateClient(GssSocket.java:107)
    at org.globus.gsi.gssapi.net.GssSocket.startHandshake(GssSocket.java:145)
    at org.globus.gsi.gssapi.net.GssSocket.getOutputStream(GssSocket.java:166)
```

## DataTransferForTheoryQCD < ArdaGrid < TWiki

```
at org.apache.axis.transport.http.HTTPSender.writeToSocket (HTTPSender.java:440)
at org.apache.axis.transport.http.HTTPSender.invoke (HTTPSender.java:138)
at org.apache.axis.strategies.InvocationStrategy.visit (InvocationStrategy.java:32)
at org.apache.axis.SimpleChain.doVisiting (SimpleChain.java:118)
at org.apache.axis.SimpleChain.invoke (SimpleChain.java:83)
at org.apache.axis.client.AxisClient.invoke (AxisClient.java:165)
at org.apache.axis.client.Call.invokeEngine (Call.java:2784)
at org.apache.axis.client.Call.invoke (Call.java:2767)
at org.apache.axis.client.Call.invoke (Call.java:2443)
at org.apache.axis.client.Call.invoke (Call.java:2366)
at org.apache.axis.client.Call.invoke (Call.java:1812)
at org.dcache.srm.v2_2.SrmSoapBindingStub.srmStatusOfGetRequest (SrmSoapBindingStub.java:2)
at sun.reflect.GeneratedMethodAccessor17.invoke (Unknown Source)
at sun.reflect.DelegatingMethodAccessorImpl.invoke (DelegatingMethodAccessorImpl.java:25)
at java.lang.reflect.Method.invoke (Method.java:585)
at org.dcache.srm.client.SRMClientV2.handleClientCall (SRMClientV2.java:178)
at org.dcache.srm.client.SRMClientV2.srmStatusOfGetRequest (SRMClientV2.java:449)
at gov.fnal.srm.util.SRMGetClientV2.start (SRMGetClientV2.java:324)
at gov.fnal.srm.util.SRMDispatcher.work (SRMDispatcher.java:817)
at gov.fnal.srm.util.SRMDispatcher.main (SRMDispatcher.java:368)
java.lang.RuntimeException: credential remaining lifetime is less than one minute
at org.dcache.srm.client.SRMClientV2.handleClientCall (SRMClientV2.java:167)
at org.dcache.srm.client.SRMClientV2.srmAbortFiles (SRMClientV2.java:347)
at gov.fnal.srm.util.SRMGetClientV2.abortAllPendingFiles (SRMGetClientV2.java:432)
at gov.fnal.srm.util.SRMGetClientV2.start (SRMGetClientV2.java:386)
at gov.fnal.srm.util.SRMDispatcher.work (SRMDispatcher.java:817)
at gov.fnal.srm.util.SRMDispatcher.main (SRMDispatcher.java:368)
srm client error:
java.lang.Exception: stopped
java.lang.RuntimeException: credential remaining lifetime is less than one minute
at org.dcache.srm.client.SRMClientV2.handleClientCall (SRMClientV2.java:167)
at org.dcache.srm.client.SRMClientV2.srmAbortFiles (SRMClientV2.java:347)
at gov.fnal.srm.util.SRMGetClientV2.abortAllPendingFiles (SRMGetClientV2.java:432)
at gov.fnal.srm.util.SRMGetClientV2.run (SRMGetClientV2.java:409)
at java.lang.Thread.run (Thread.java:595)

real    691m35.696s
user    6m6.926s
sys     2m18.016s
```

Now let's go to /tmp before to avoid disk quota exceeded:

```
[lxplus250] /tmp > date
Fri Feb 26 16:30:46 CET 2010
[lxplus250] /tmp > time srmcp -srm_protocol_version 2 srm://srm-public.cern.ch:8443/srm/managerv2
WARNING: SRM_PATH is defined, which might cause a wrong version of srm client to be executed
WARNING: SRM_PATH=/afs/cern.ch/project/gd/LCG-share/3.1.38-0/d-cache/srm
SRMClientV2 : srmReleaseFiles: try # 0 failed with error
SRMClientV2 : ; nested exception is:
    java.io.EOFException
SRMClientV2 : srmReleaseFiles: try again

real    5m18.570s
user    0m46.486s
sys     2m10.577s
```

It fails again (it maybe due to a previous attempt where quota was exceeded in a local directory /relative path!).

We try again:

```
[lxplus250] /tmp > time srmcp -srm_protocol_version 2 srm://srm-public.cern.ch:8443/srm/managerv2
WARNING: SRM_PATH is defined, which might cause a wrong version of srm client to be executed
WARNING: SRM_PATH=/afs/cern.ch/project/gd/LCG-share/3.1.38-0/d-cache/srm
```

tests with big files from CERN intranet



```
real    6m18.579s
user    0m45.207s
sys     1m29.777s
```

## initial problems

Version mismatch if v1 (default) is used:

Possible reason for this error is an outdate client SOAP protocol which is not understood by the server deployed at CERN:

```
> srmcp srm://srm-public.cern.ch/castor/cern.ch/user/m/moscicki/tmp.test file:///tmp/tmp.txt
WARNING: SRM_PATH is defined, which might cause a wrong version of srm client to be executed
WARNING: SRM_PATH=/afs/cern.ch/project/gd/LCG-share/3.1.38-0/d-cache/srm
SRMClientV1 : Method 'nsl:get' not implemented: method name or namespace not recognized
SRMClientV1 : get : try # 0 failed with error
SRMClientV1 : Method 'nsl:get' not implemented: method name or namespace not recognized
srm copy of at least one file failed or not completed
```

## smrcp on my ubuntu box

I copied the entire d-cache directory and LCG certificates directory

```
cp -a /afs/cern.ch/project/gd/LCG-share/3.1.38-0/d-cache .
scp -r lxplus:/afs/cern.ch/project/gd/LCG-share2/certificates .
```

A fix: `srmcp` is implemented as 2 bash scripts which finally call `java`. The magic line of the shell scripts is pointing to use `sh` which is wrong if on the system `sh!=bash` (which is often the case).

Workdir environment:

```
export X509_CERT_DIR=/home/moscicki/srmcp_standalone_client/certificates
moscicki@pcarda75 ~/srmcp_standalone_client
```

Getting these errors:

```
copy failed with the error
org.globus.ftp.exception.ServerException: Server refused performing the request. Custom message:
500-globus_xio: System error in connect: Connection refused
500-globus_xio: A system call failed: Connection refused
500 End.]. Nested exception is org.globus.ftp.exception.UnexpectedReplyCodeException: Custom me
```

Solution: the problem is caused by a wrongly configured `/etc/hosts` as described at:

<https://computing.llnl.gov/linux/slurm/faq.html#ubuntu>

Some systems by default will put your host in the `/etc/hosts` file as something like

```
127.0.1.1      snowflake.llnl.gov      snowflake
This will cause srun and other things to grab 127.0.1.1 as it's address instead of the correct address
```

A test of host configuration in python: `socket.gethostbyname(socket.gethostname())`

Anyway, the even if the client is wrongly configured, it is probably a `gsiftp` protocol or server implementation flaw that relies on a IP address sent by the client (and not using the client IP address from the connection itself).

# lcg-cp

lcg-cp srm://srm-public.cern.ch/castor/cern.ch/user/m/moscicki/tmp.test file:///tmp/tmp.txt

-- JakubMoscicki - 12-Jan-2010

- srmcp\_standalone\_client.tgz: java-based srmcp standalone client from dcache and certification authorities certificates

---

This topic: ArdaGrid > DataTransferForTheoryQCD

Topic revision: r10 - 2010-03-03 - JakubMoscicki



Copyright &© 2008-2022 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.  
or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)