

# MIPP Software/Analysis meeting

Holger Meyer

Fermilab

9/27/07

## Outline

- Accessing Data from Mass Storage
  - Copying files
  - Using DSTs in ROOT without copying files

# Accessing Data from Mass Storage

- ~1 TB of local disk space on e907ana cluster is not enough to keep all useful data
  - Clean up unused files, please!
- All data from Batch processing (dsts for pass3, pass4a, pass4aii,... and output and dsts for some MC runs) reside in Fermilab Mass Storage and are accessible through [dCache](#).
  - Large Disk Cache in front of EnStore Tape storage
  - Use dccp to read/write files to dCache/Enstore. Don't use encp to read directly from tape.
    - dCache will serve file from tape the first time it is used, then serve from disk  
→ Faster than using tape copy every time.
    - Writes to dCache go to disk first, then get copied to tape by dCache  
→ Faster for the user, copy process finishes when dCache disk has the file.
  - Batch jobs copy job results to Mass Storage through gridFTP which goes through dCache
    - DSTs will already be on dCache disks if you use them 'soon' after jobs on the OSG finish.
      - 'soon' is days to weeks and depends on overall dCache usage

# Accessing Data from Mass Storage

- Didn't we use dccp already?
  - Well, some of us did. But it was never working on the newer e907ana systems and it broke on the older systems due to software and security upgrades.
- I had Vladimir Podstavkov make an rpm of the latest version and installed that on the SLF4.4 systems. Older systems will get dccp when they get upgraded.

- dccp not working on the new systems was one of the reasons to not upgrade other systems

- Firewall needs to have holes for the port range and nodes used by dCache
- Access files through dcap URIs, e.g. (see next slide for how to make this simpler)

```
dccp dcap://fndcal.fnal.gov:24725/pnfs/fnal.gov/usr/e907/mc/dst/lh2/mippHist10012857.01.mcdigireco.root .
```

- Access with valid Kerberos ticket...
  - Ports [24725](#) or [24736](#)
- ... or proxy certificate generated from DOEgrid x509 certificate (the one you also use for MIPP-DocDB access)
  - Ports [24525](#) or [24536](#)

# Accessing Data from Mass Storage

- dcap URIs are long, shell variable are short:

```
dccp
```

```
dcap://fndca1.fnal.gov:24725/pnfs/fnal.gov/usr/e907/reco3/FNAL/pass4aai/15/mippHist15930.pass4aai.root .
```

... is the same as ...

```
dccp ${DCAP_HOME}/reco3/FNAL/pass4aai/15/mippHist15930.pass4aai.root .
```

- How to find a file:
  - The pnfs namespace is mounted on all e907ana systems
    - `${DCAP_HOME}/reco3/FNAL/pass4aai/15/mippHist15930.pass4aai.root`  
... shows as ...  
`/pnfs/e907/reco3/FNAL/pass4aai/15/mippHist15930.pass4aai.root`
    - `/pnfs/e907/` looks just like an nfs mounted file system
      - It can only be mounted from within the .fnal.gov domain, not available off-site
      - `dccp /pnfs/e907/...` does not work.
- Accessing dCache from Universities
  - Should work if you have valid Kerberos ticket or DOEgrid certificate
  - Let me know if you have problems

# Using DSTs in ROOT

- ROOT supports the dcap protocol. You can open ROOT files from dCache directly without copying first:

```
[hmeyer@e907ana8 ~]$ root dcap://fndcal.fnal.gov:24725/pnfs/fnal.gov/usr/e907/reco3/FNAL/pass4aai/15/mippHist15930.pass4aai.root
*****
*                                     *
*           W E L C O M E  to  R O O T           *
*                                     *
*   Version   5.14/00e           29 March 2007   *
*                                     *
*   You are welcome to visit our Web site   *
*           http://root.cern.ch             *
*                                     *
*****
```

```
FreeType Engine v2.1.9 used to render TrueType fonts.
Compiled on 29 March 2007 for linux with thread support.
```

```
CINT/ROOT C/C++ Interpreter version 5.16.16, November 24, 2006
Type ? for help. Commands must be C++ statements.
Enclose multiple statements between { }.
root [0]
Attaching file dcap://fndcal.fnal.gov:24725/pnfs/fnal.gov/usr/e907/reco3/FNAL/pass4aai/15/mippHist15930.pass4aai.root as _file0...
root [1]
```

- Anamipp does not currently support this.
  - May be simple to upgrade IO module
  - Simple scripts to dccp file, run anamipp on it, delete it, work all right

# Accessing Data from Mass Storage

- Firewall issues: iptables may block the transfer
  - Use `dccp -A` for active mode, i.e. Client (e907ana) connects to the dCache pool node rather than the pool node connecting to your client
  - Open incoming connections from dCache pool nodes, currently subnet 131.225.13.0/255.255.255.0
    - Needed for ROOT to open DSTs/Histogram files directly
    - Port range used: ???
- Volatile dCache area:
  - `dcap://fndca1.fnal.gov:24725/pnfs/fnal.gov/usr/fermigrid/volatile/mipp`
  - Good for Grid job output that you want to check before archiving/deleting
- Outlook:
  - SRM – StorageResourceManagement
    - Middleware layer that negotiates door, port, authentication method,...
    - `srmcp`, `srmls`, etc. E.g. To list files in mipp-volatile dCache space:
      - `srmls srm://fndca1.fnal.gov:8443//pnfs/fnal.gov/usr/fermigrid/volatile/mipp`