

Chirp for ATLAS User Output

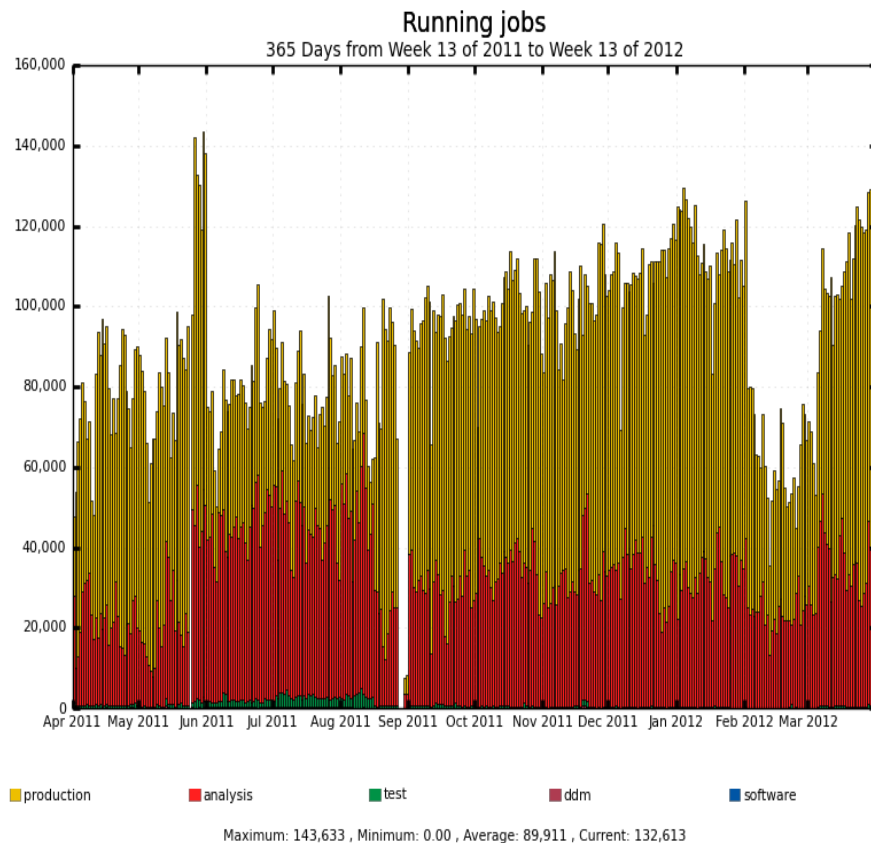
- Rod Walker, LMU, Munich

- Introduction
 - Chirp and ATLAS use-case
- Status & experience
- Conclusions

Chirp

- Part of the cctools package
- Distributed file system
 - can run non-root both for client and server
 - supports many auth methods, inc. X509
 - fuse mountable
 - full acl support

Use-case: ATLAS Analysis



- April 2011-April 2012:**
- more than 1500 ATLAS grid users
 - 20-40k running jobs
 - 300-400k finished jobs per day
 - 2 output files/job – 10Hz store rate

User Data flow

- User output(s) and log written from job to local SRM
 - e.g. Dcache, DPM, xrootd, storm
- Registered to ATLAS data handling system(DDM)
- User uses DDM client to pull files to local institute – typically ntuple to run Root over
- System can automatically move files between SRMs, but not to desktop

User Output File size

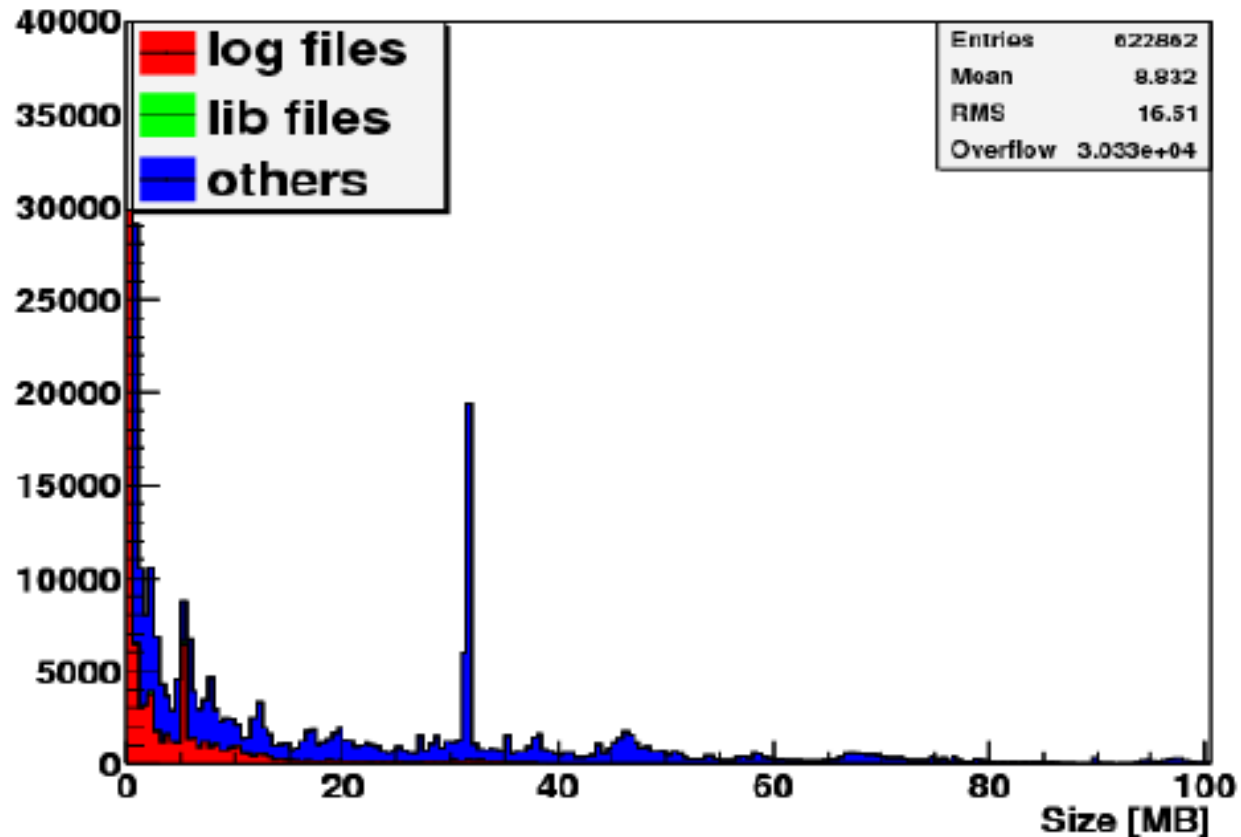


Fig 1: Snapshot of user output files at FZK from June 2010

Typical remote file copy $\sim 1\text{MB/s}$ – negligible time spent copying 8MB over WAN

Problem with that?

- Manual user step to get data
 - many jobs = many dq2-get and all SRMs need to work
 - site downtimes (planned or not) prevent access
- dq2-get is chaotic access to SRM
 - possible overload and interference with controlled DDM transfers
- Moving files with DDM better for SRMs but ..
 - adding many small files to the system slows it down
 - per file overhead typically 30-60s
 - often file size is irrelevant to transfer time

New data flow

- Store output files directly to remote Chirp server
- Also store to SRM for now
- User accesses outputs from Chirp server
 - directly using fuse
 - perhaps running hadd to merge them
 - copy then with cp, rather than dq2-get – big difference!
- Very positive feedback, but still small user base $O(10)$
 - users like posix access

Status

- Catch-all Chirp server at CERN
- ATLAS analysis client tools have user switches to enable storing output to Chirp
- Wiki linked to FAQ (low-profile)
 - <https://twiki.cern.ch/twiki/bin/view/Atlas/ChirpForUserOutput>
- Few regular users – added robot user
 - functional test jobs use Chirp(to create load)
 - 840k files on server, tmpwatch cleaning active
 - 280k sub-directories (someone asked!)
- All CERN login nodes allow FUSE
 - Chirp tools in ATLAS afs area


```
source /afs/cern.ch/atlas/offline/external/GRID/DA/cctools/setup.sh
```

```
[lxplus442] ~ $ mkdir chirp
```

```
[lxplus442] ~ $ fusermount -u chirp
```

```
[lxplus442] ~ $ ls chirp/voatlas92.cern.ch/RodWalker/
```

```
.....
```

Authorization

- must use X509 certificates for Authentication
 - server run under long-lived X509 proxy
- created user directories on Chirp server
 - created `.___acl` files by script
 - rather than chirp client auth manipulation
 - writable from the user jobs
 - full control for user, e.g. to share in group

Experience

- Chirp server froze a few times under (not so) heavy load
 - cctools support help track this to a trivial cause
 - ssl uses \$HOME/.rnd for every authentication process
 - this was on AFS and slowed it down a lot
 - simple fix: export RANDFILE=/tmp/chirp.server.rnd
- Operated as dev service, e.g.
 - no init.d script to start on reboot
 - user and acl's added via manual scripts, rather than interface to ATLAS user Db

Wish List

- VOMS integration
 - acl's applicable to e.g. /atlas/Role=pilot, /atlas/de, /atlas/phys-higgs
- Redundancy/Failover options
 - client tries list of servers?
 - maybe some redirect to pool of Chirp servers
- Web interface for user browse/download

Conclusions

- Chirp is easy to install and operate
 - very robust and good support when needed
- Users like it(posix access to outputs)
 - not actively advertised to users
 - lack redundancy/production deployment
 - lack regional/institute servers (instructions+support)
- Did not reach scale limit of single server
- Increase robot usage and perhaps enable by default in future