



Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

XRootD for Neutrino and Precision Muon Experiments

Robert Illingworth

XRootD Workshop @ UCSD

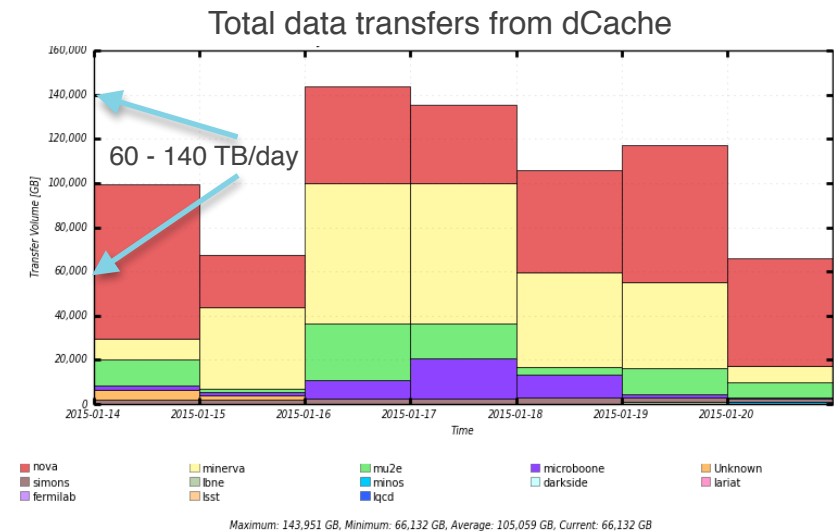
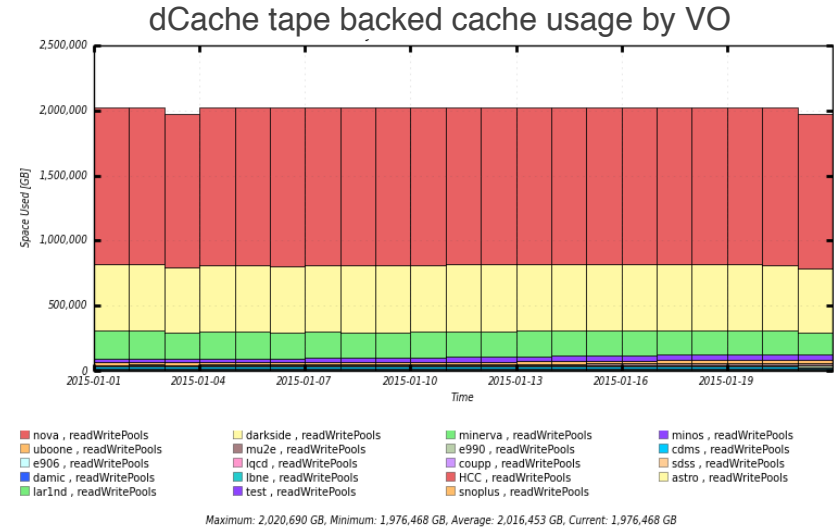
29 January 2015

Neutrino and precision muon experiments

- Fermilab has multiple running and upcoming neutrino, precision muon (and a few other categories) experiments
- By LHC standards these aren't large, but the cumulative data volume is significant
- The experiments don't have much effort to spare for computing infrastructure work, so they are encouraged to use common tools as much as possible

Computing models

- Currently data is largely centralized at Fermilab
 - 3 PB dCache (29 PB on tape)
 - Mostly accessed by gridftp
- Majority of the reconstruction and analysis at Fermilab
 - Some Monte-Carlo production at OSG sites or AWS
- As time goes on, expect more data to be replicated at other sites and more use of grid and cloud resources



Current use of XRootD

- Small scale use of XRootD from dCache
 - Generally people who want to directly open files from within root, for example: event displays; extracting metadata for cataloguing
 - It just works...

Potential other uses

- “Auxiliary” input files for Monte-Carlo
- Analysis at remote sites via federation
- Interface to storage elements

Auxiliary files

- Some tasks like Monte-Carlo generation need to read a limited subset of files from a dataset
 - For example, NOvA jobs require $\sim 10\%$ of $\sim 200\text{GB}$
 - Each file gets read multiple times
- The entire dataset gets used during large scale production
- Job efficiency is much higher when files can be read from a relatively local source
- Have been using cvmfs for this, but it's not ideal
 - tends to cause cache thrashing
- Alien cache cvmfs works better, but there are concerns about how widely it will be deployed

Auxiliary files

- Can XRootD help with this?
- Use federation
 - Only a real benefit if files can be accessed from a reasonably local source, so needs caching to be useful

Distributed analysis

- This use case is a familiar one - jobs read input data files locally/across the network as needed
- Most experiments can't provide much effort for managing datasets
 - Need a (semi-)automated system for doing data placement and cleanup
 - Again, what is wanted is some form of caching
- Experiment I/O is not optimized for high-latency operations
 - Not something anyone from the experiments has really thought about

Interface to Storage Elements

- Currently access to site storage elements often isn't convenient.
- Smaller experiments can't provide much in the way of personnel for management, so largely reliant on generic OSG service
- SRM is not the easiest interface to use
- Would OSG standard XRootD services provide more convenient access to storage?

The future - ELBNF

- In the process of becoming a fully international collaboration
- Fermilab centric model not an option here
- Needs fully distributed processing along the lines of the LHC experiments
- XRootD federation a good fit

Conclusions

- There are multiple neutrino and muon experiments with significant amounts of data but little manpower to manage it
- Currently data is highly centralized at Fermilab, but we'd like to change that
- Any solution has to leverage widely deployed (OSG) services and require minimal operations effort
- Can we benefit from XRootD to achieve this?