

Xrootd

Back To The Future

Andrew Hanushevsky
SLAC National Accelerator Laboratory

WLCG Data/Storage TEG Meeting
January 24-25, 2011
NIKHEF, Amsterdam, Netherlands

The Project Vision

- Provide data access in a manner which is
 - Affordable
 - Minimize hardware, software and human requirements
 - Highly performant
 - Minimize use of CPU, disk, memory and network resources
 - Robust
 - Minimize failure modes and bottlenecks
 - Secure
 - Support widely used authentication & authorization models
 - Scalable
 - Open ended scale up and scale out technology
 - Supportable
 - Collaborative model with a straight-forward architecture

Vision & Experiments

- The vision dovetails with experimental concerns
 - This is should be expected when one considers that...
 - Xrootd was developed from the start with
 - Experimental involvement *during* the development process
- It continues to this day
 - The project is driven by experimental feedback
 - E.G. recent additions to the monitoring feature set driven by CMS
- The project is well aware of the bottom line
 - Experiments don't care how analysis is done as long as
 - The analysis gets done in a timely and efficient manner
 - I.E. it's all about high energy physics not the computer science

Vision & The Future

- Xrootd is an extremely well hardened system
 - In production use for over 10 years in many environments
- The number of new bugs is relatively small
 - Maintenance effort is very manageable
- Our primary goal is to evolve what we have
 - To meet evolving experimental operating modes
- The architecture is well positioned to do that
 - The plug-in model has proven to be a god-send
- The vision and platform form a sound foundation
 - This serves us well to meet future requirements

Some Things We See Comming

- Solidifying the federated storage model
 - Make it even easier with a far more predictable outcome
- More monitoring
 - The right information at the right time to the right places
- More storage management features
 - Simple multi-tiered storage
 - We already do it but the documentation is sparse
 - Automatic storage rebalancing
 - Very sensitive to what underlying file system is used
- Others as they bubble up in importance

Conclusion

- The next 3 years will be exciting
 - Nothing is really cast in concrete
 - The project is positioned to be highly adaptable
 - In a way that does not jeopardize existing production
 - Cluster IP is uniquely applicable elsewhere
 - To new and evolving technology and protocols
 - The architecture and software can work in many areas
 - The collaboration is expanding
 - Currently CERN, Duke, JINR, SLAC, and UCSD