

# Measurements from NOVA GENIE workflow over CVMFS

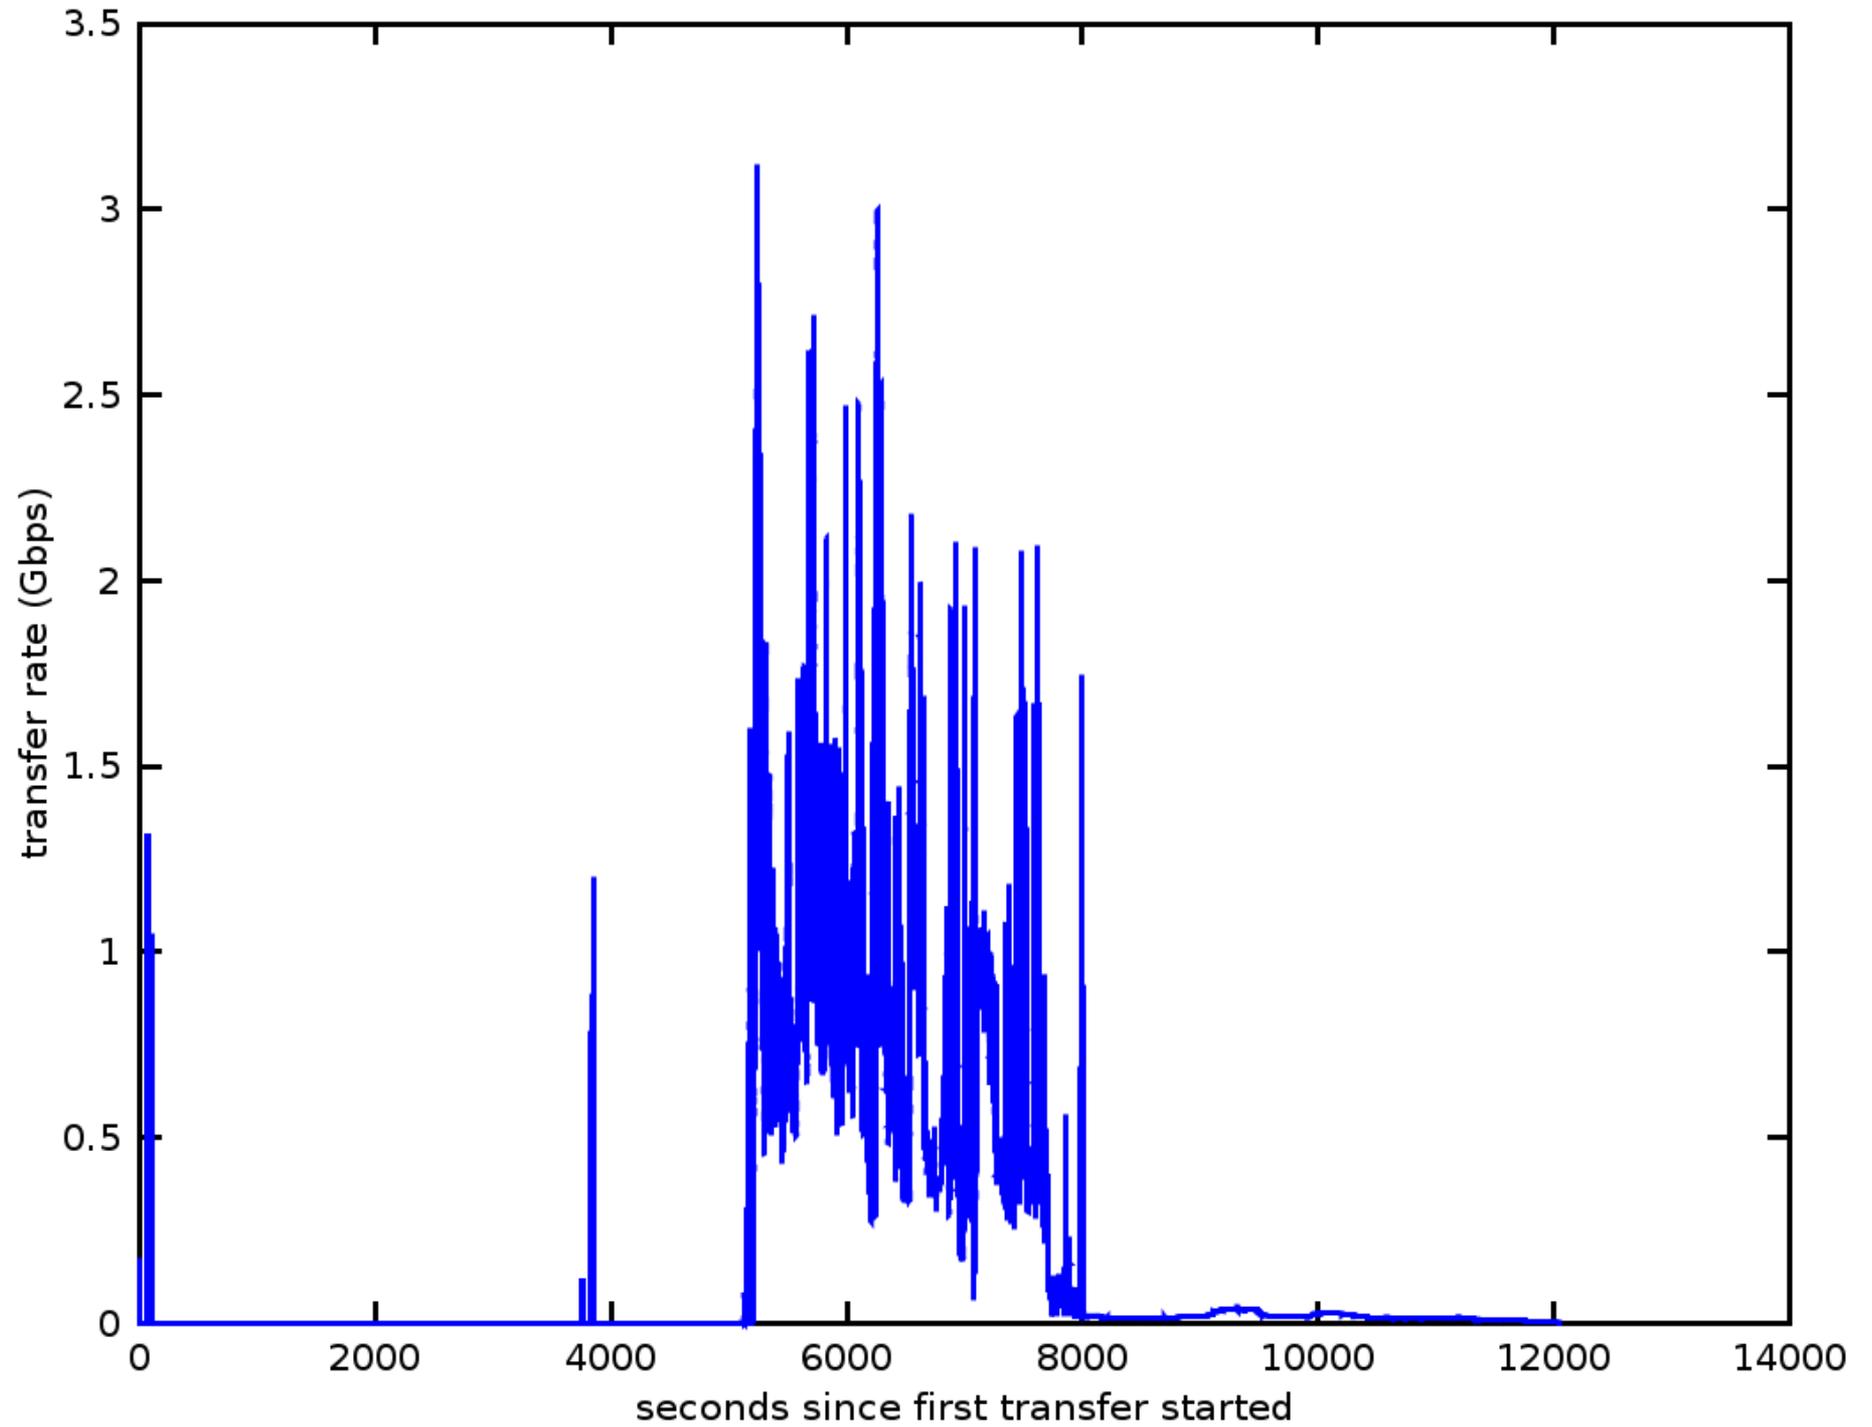
27 March 2014

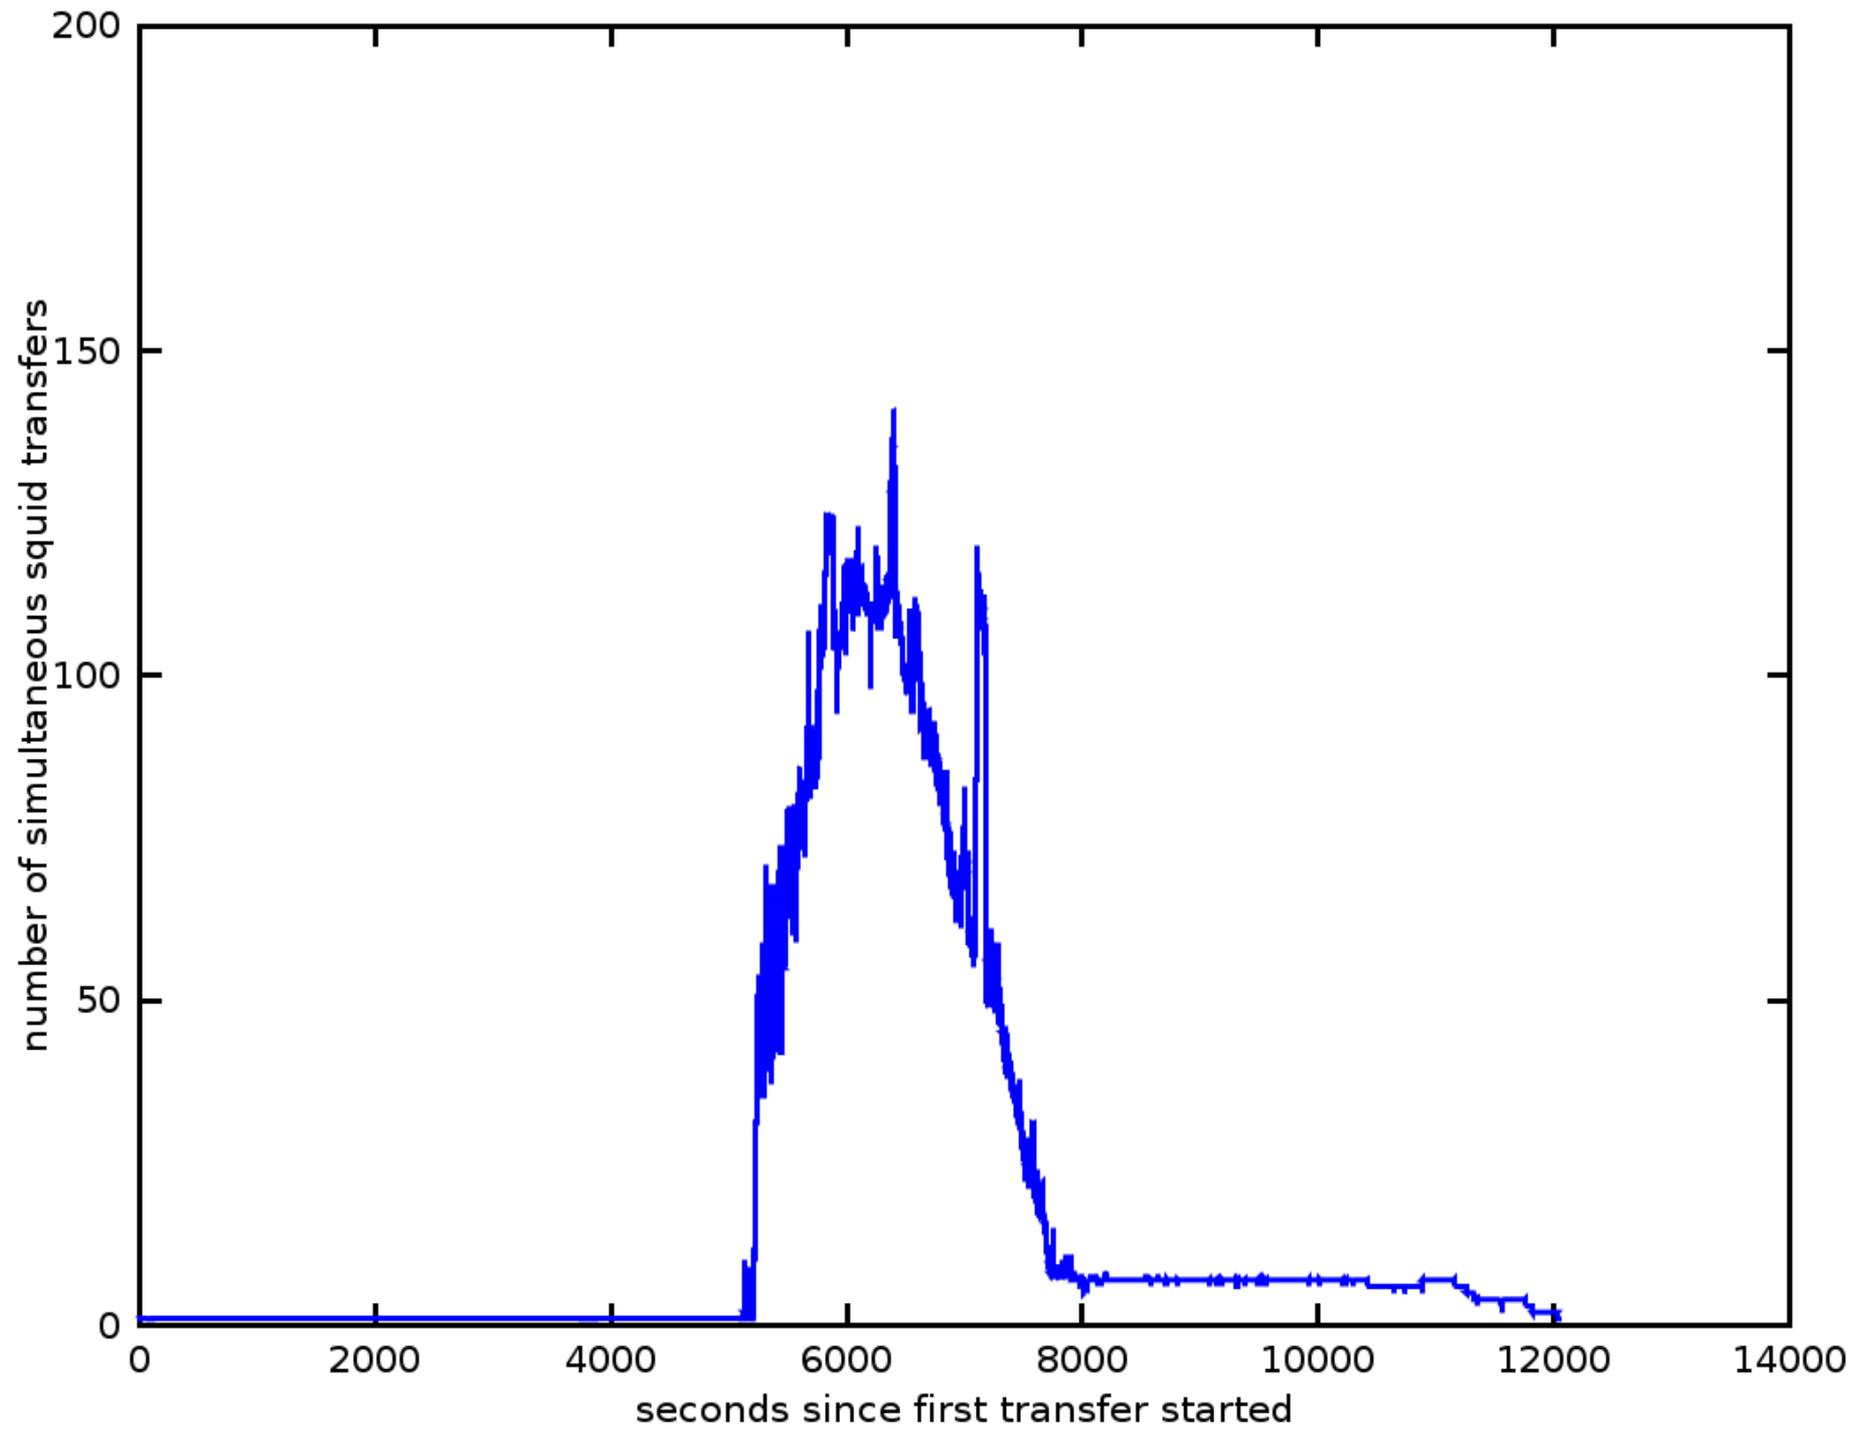
Dave Dykstra

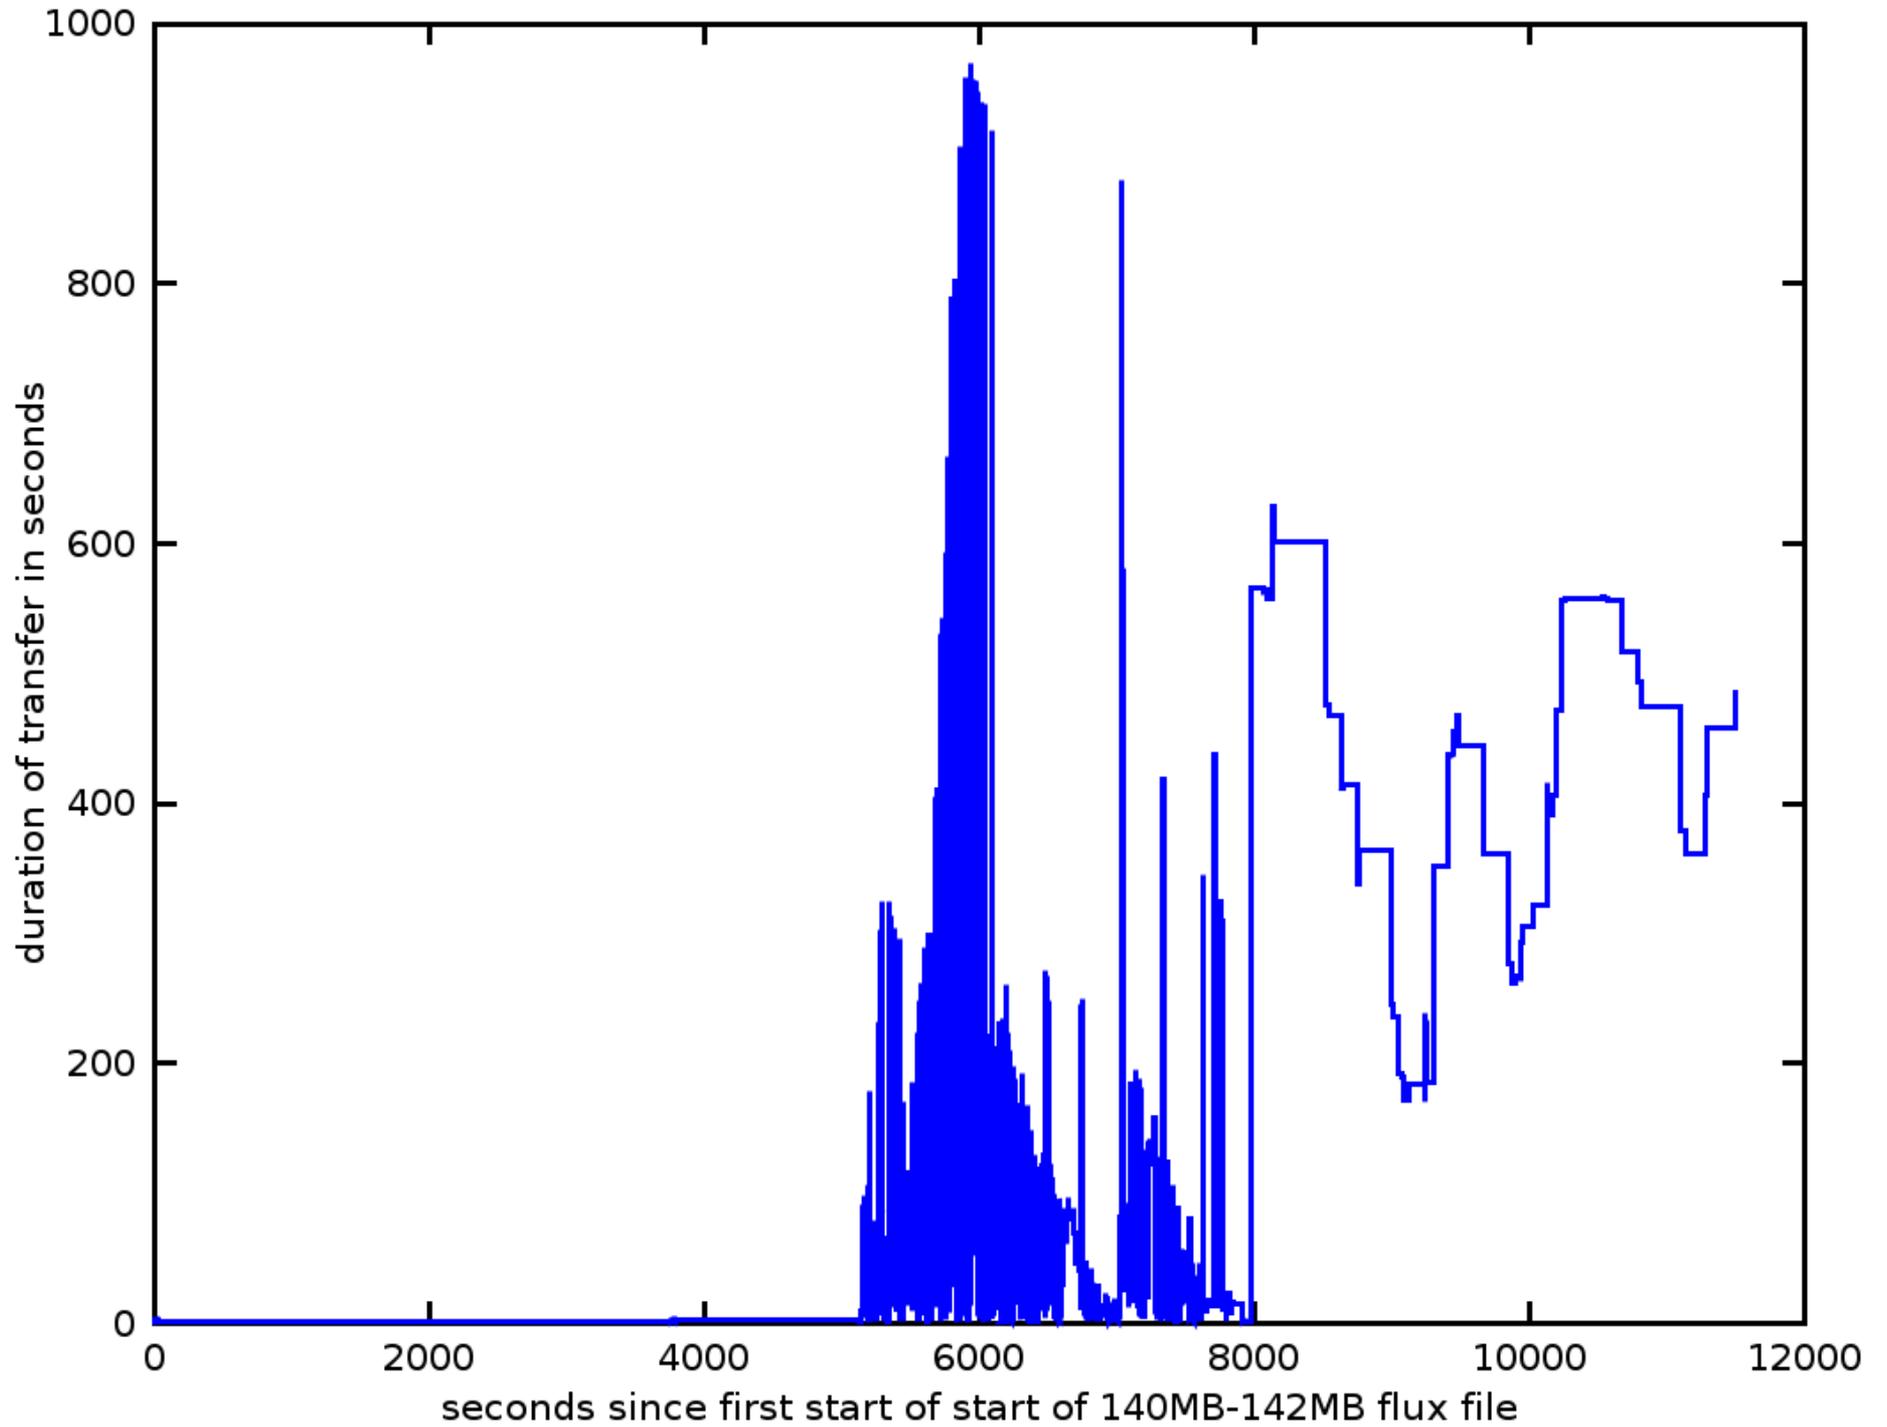
[dwd@fnal.gov](mailto:dwd@fnal.gov)

# NOVA GENIE CVMFS test

- Ran 128 jobs on FNAL CMS cluster March 21
  - 1.8GB of ~140MB flux files read per job
  - 4 2Gbit/s squids used
- Jobs landed on 87 different nodes
- Total read: 230GB of ~140MB files and 50GB other size
- Download times of ~140MB files as high as 16 minutes
  - Average time 87 seconds, or 1.6MB/s
- Peak combined transfer rate as high as 3Gbit/s
- High transfer rate spread out over 50 minutes
- Max of 138 >1 sec transfer in given second







# Conclusion

- GENIE jobs run from CVMFS significantly strain squids and the data loading is very inefficient
- The squids were unable to reach their full network capacity, probably limited by disk accesses because there was too much data to fit into filesystem buffers
  - (each squid machine was running two squids so data would be cached in two different places; two machines had 16GB RAM and two had 32GB)