

DUNE DC24 Recap

Doug Benjamin

BNL

3-Mar-24

Acknowledgements

All the credit for these slides come from the following people (team effort) -

Brandon White, Jake Calcutt, Aaron Higuera, Elisabetta Pennacchio, Heidi Schellman, Mike Kirby, Steve Timm, Ken Herner, Andrew McNab, Chris Brew and DB and all folks operating DUNE sites

Mistakes are mine.

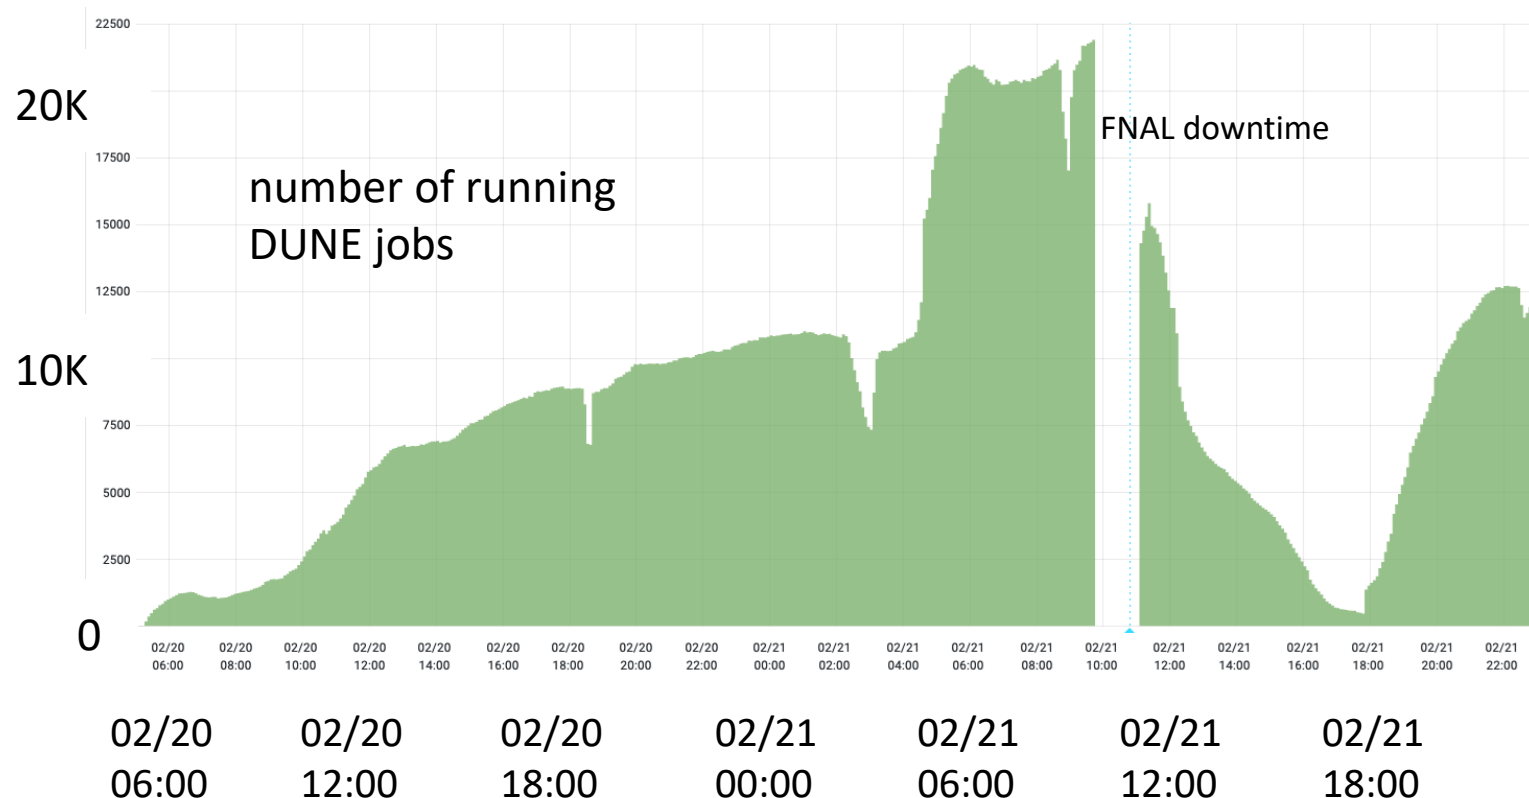


Initial Test Plans

- Test 1 - “FD” Raw Data to archival storage
 - Simulate the archival of 25% of the raw data rate from the Far Detector
 - translates to 1 GB/s from SURF (BNL as stand-in) to FNAL
 - replicate that “FD” raw data to archival storage facilities around the world
 - replicate the “FD” raw data to disk storage elements around the world for prompt access from compute elements
- Test 2 - “FD” Raw Data keep up processing
 - Maintain continuous processing workload at distributed sites commensurate with 25% “FD” raw data rate (1 GB/s)
 - Utilize compute resources across sites in Europe and North America
 - Match the locality of jobs with locality of data at nearby RSEs
- Test 3 - SuperNova Raw Data rapid transfer & processing
 - 3.5 GB/s SURF (BNL) to FNAL to NERSC

Keep up processing scale testing

- Used FD1-HD FD2-VD MC step 2 full reconstruction production processing using justIN
 - More I/O intensive than normal keep up processing will be – 1 GB – 2 GB output (4 files – 1 big 3 small)

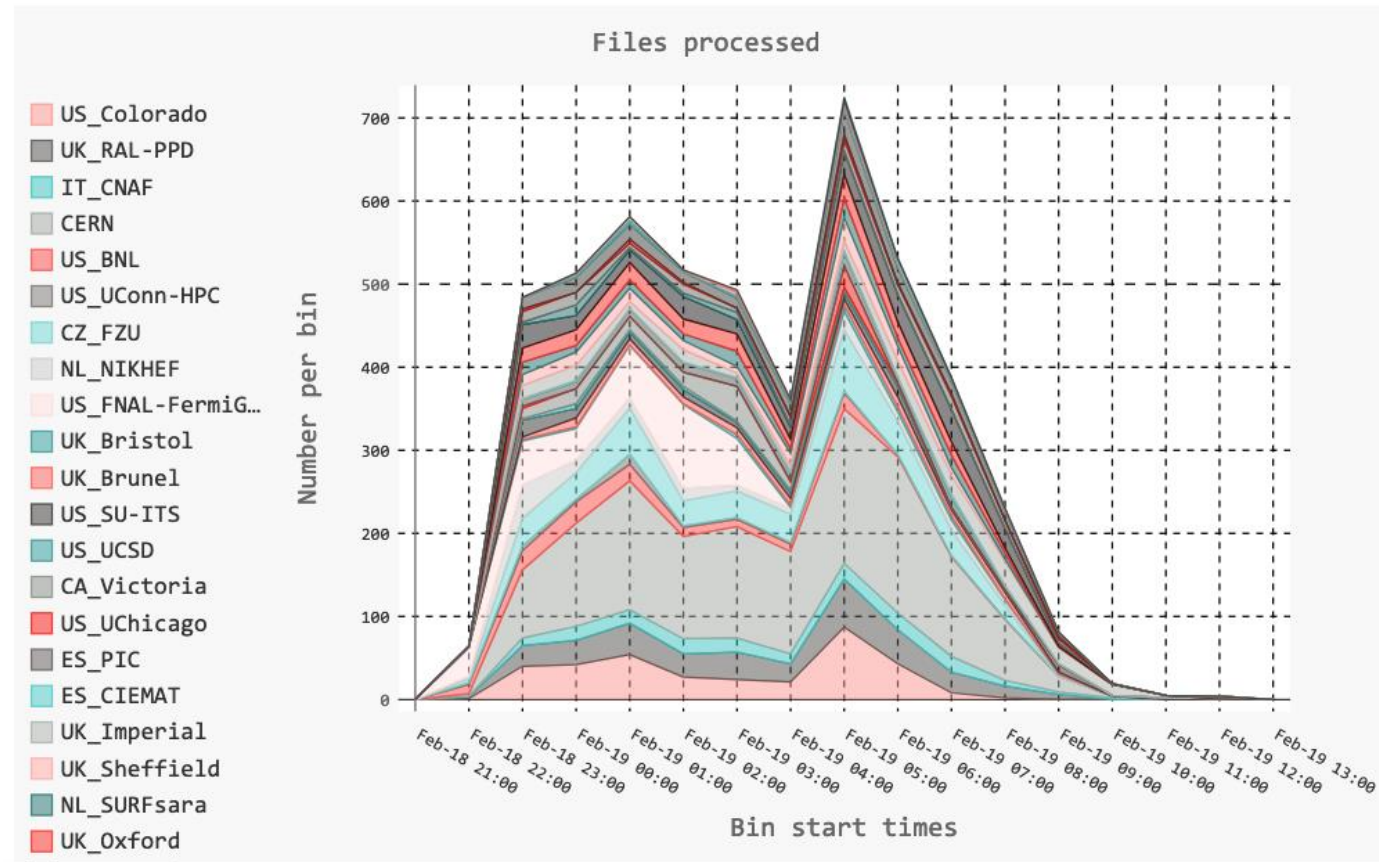


Lessons learned so far

- Rucio/MetaCat scale to 16k jobs
 - FNAL DB admins to increase # Rucio DB connections to 500
- JustIN schedd at RAL need lots of memory and disk space
 - Increased to RAM 60 GB, disk 250 GB
- multiple schedd VM's at RAL
 - will retest
- Needed to tune File Transfer Service configuration settings for better throughput of output files back to FNAL
 - some channels clogged with slow transfers - due to WLCG Data Challenge 24 (DC24) interference
 - Tuning complete

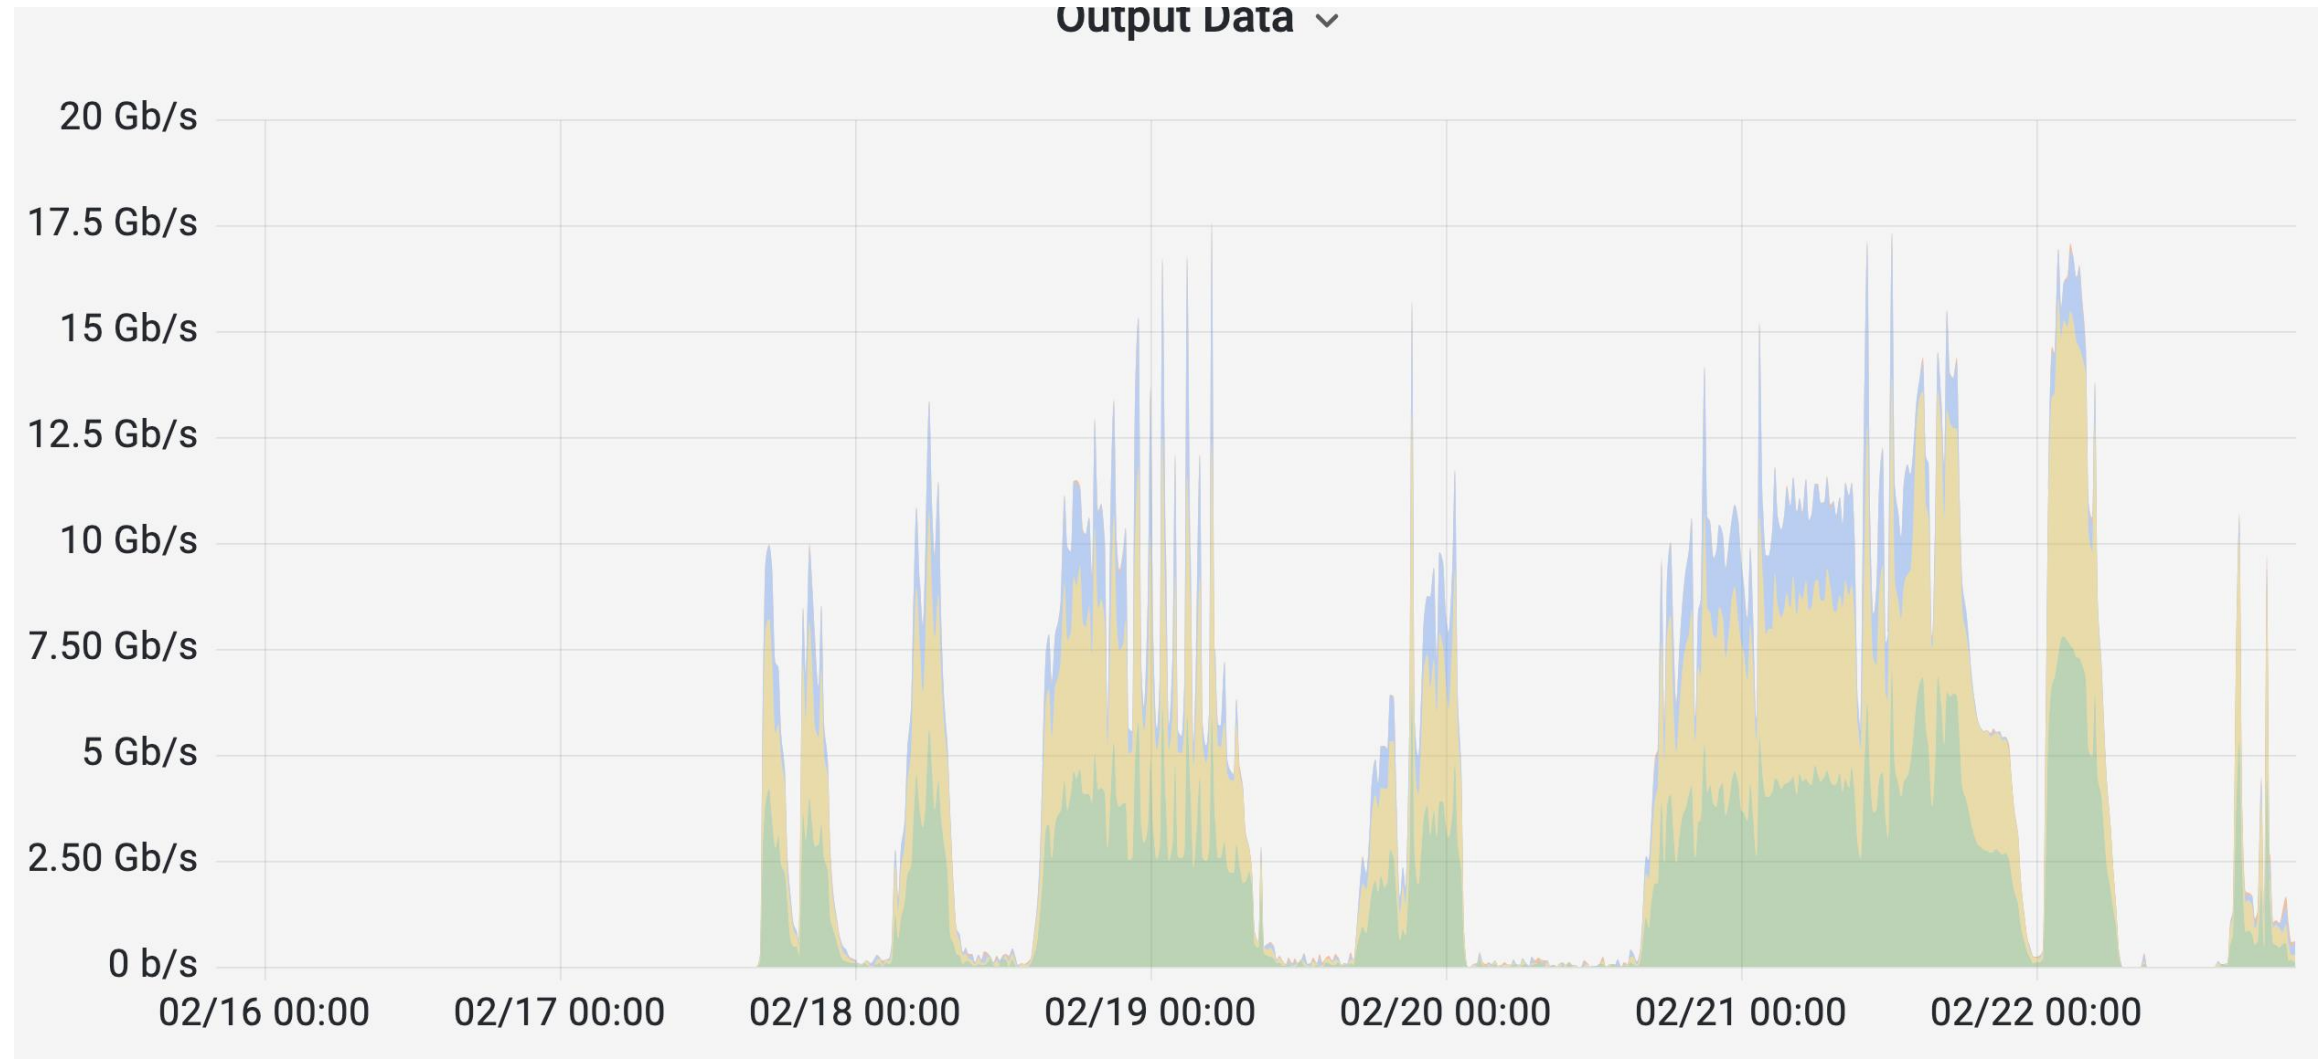
- ~15k jobs running simultaneously to keep up with ProtoDUNE data taking this Spring 2024.

Example justIN workflow task



File processing rate across different Compute sites in US, Canada and Europe

Data streaming from PIC to European sites



Conclusions

- Were not able to run the SNB or DC running tests at rate desired because we were occupied with the keep up processing test
- Discovered important limitations and weaknesses in our current setup for keep up processing.
 - Initial Rucio mis-configuration – lead to many TPC transfers using xrootd instead of davs. (many timeouts at RAL)
 - DC24 extremely helpful as a stress test for our processing system. DUNE used MC production reconstruction (reco2) as part of the DC24 activity.
 - Already applying lessons learned immediately improve our workflow system justIN
 - Identified services that need to be hardened.
- Planning to perform another dress rehearsal ahead protoDUNE HD data taking this spring