What Data Do We Have and Where Is it?

Steven Timm

DUNE CRAB

20 Dec. 2024

A word about SAM

- No analysis or production files have been declared to SAM since early 2024.
  - We started using the extended capabilities of MetaCat which are not backward compatible to SAM.
- Starting now, new raw data files will also not be declared to SAM.
- Any SAM locations of raw data that was taken in 2024 are wrong.
- We have not turned off SAM4Users yet but are working on its replacement.
- We have already brought forward all files that were known to SAM into MetaCat
- We still have to transfer some of the location information of legacy files from SAM into Rucio
- SAM will go totally read-only in summer of 2025 and go away completely one year later.

Data placement strategy (pg. 142 of CDR)

- (pg. 142 of CDR). Table 6.8 and 6.9
- 2 copies of raw data “physics” retained indefinitely on tape and for short time on disk.
- DAQ commissioning data “test” retained on disk 6 months
- Detector commissioning data “study” retained on disk 6 months and on tape 5 years
- 2 most recent versions of reprocessed data available on disk at any given time.
- 2 most recent versions of full-reconstructed MC available on disk at any given time
- Plan has evolved since TDR..Latest version always available in DUNE DOCDB
- https://docs.dunescience.org/cgi-bin/sso/ShowDocument?docid=31722

Current plan: (considering only disk)

- https://docs.dunescience.org/cgi-bin/sso/RetrieveFile?docid=31722&filename=CCB-Report-2025-v5.pdf&version=9. (see table 3)
- Have now moved raw “physics” data to have 2-3 year stay on disk
- Raw “test” data 6 months on disk
- Trigger primitives 6 months on disk
- Reco-data 2 years on disk x 2 copies x 2 versions
- Reco-sim 2 years on disk x 1.5 copies x 2 versions
- Analysis-data 5 years on disk x 2 copies x 2 versions. (CAF files etc., small)
- Analysis-sim 5 years on disk x 2 copies x 2 versions. (CAF files etc. small)

Current Raw Data Sets

- Hd-protodune—taken between May-December 2024. All currently on disk.
  - 6308 TB total—1787 TB of that is trigger primitives, rest is TPC data
    - Need a real physics case to analyze these trigger primitives otherwise they are candidates for rolling off to tape.
  - Of the TPC data 2938TB of that is physics taken during beam, 1577TB of cosmics
    - Work in progress between ProtoDUNE conveners and database people to make datasets organized by beam momentum and good run / bad run.
  - Clock of 6 months and 3 years starts as of Dec 6 2024.
- ProtoDUNE-SP: Taken mostly in 2018 with cosmics as late as 2021.
  - 4717TB total. Still on disk only the good runs from the beam runs. 474 TB at RAL_ECHO
- ProtoDUNE-DP: 409 TB of charge readout, 55TB of photon detector readout. All still on disk mostly at IN2P3
- ProtoDUNE-VD. 120TB thus far—expect to take 5 PB
- 2x2-Minerva, 13TB total

Coldbox raw data

- Vd-coldbox (new DAQ) 127TB
- Vd-coldbox-bottom and vd-coldbox-top (old DAQ) 58 + 339 TB respectively
- Hd-coldbox 85 TB
- (retention type “study”) Most of above still on disk.

Current Reconstructed Data Sets

- hd-protodune-det-reco: output of keepup processing. 746 TB. All on disk at Fermilab
- pdsp_det_reco: 257TB of PDSPProd4, split between QMUL, LANCASTER, MANCHESTER, RAL_ECHO and RAL-PP. Some duplication.
- Protodune-dp full reco: 12 TB

Current Full-Reconstructed MC

- pdsp_mc_reco: 1764TB split between QMUL, LANCASTER, MANCHESTER, RAL_ECHO, RAL-PP (campaign PDSPProd4a). More still getting made.
- hd-protodune: 540 TB (campaign pdhd_mc_2024a). At Fermilab
- Fardet-hd: 232TB long baseline + 180TB atmospheric nu (campaign fd_mc_2023a) at MANCHESTER
  - ~150TB of reco1 (hit-reconstructed) that can come off disk
- Fardet-vd: 150 TB long baseline (campaign fd_mc_2023a). At SURFSARA
- Fardet low energy: in progress—expected to be 1350TB
- Still to be generated—a replacement hd-protodune and vd-protodune + ???

Transient output files from production

- 305TB tied up right now in temporary per-RSE datasets used for catching production JustIN output
- Has been more in recent months—depends on how much recent production has been done.
- Should plan that up to 7-8% of space will be used for this at any given time (current level 3%).

Overall RSE usage (TB) Dec 20

RSE	RuleUsage	RucioUsage	StorageUsage	DuneproLimit	StorageLimit	Change (R.U)	Free
BNL	6	10	3	830	831	-71%	828
FNAL_DISK	1574	1613	1613	2100	2674	-2%	1061
CA_SFU	4	4	4	200	200	NEW	196
PRAGUE	550	609	630	1126	1161	-19%	531
NIKHEF	1005	1098	1098	1143	1143	-0.2%	45
SURFSARA	456	614	614	781	756	-6%	142
GLASGOW	495	495	495	500	500	+99%	5
LANCASTER	452	462	462	700	700	-0%	238
MANCHESTER	603	603	603	1000	1077	+9%	474
QMUL	495	641	614	1100	1100	28%	486
RAL-PP	420	558	575	1000	1000	-7%	425
RAL_ECHO	718	868	865	1050	1000	-6%	135
CNAF	8	16	8	300	550	+0.9%	542
CCIN2P3	1003	1062	1061	1100	1181	+0.2%	120
PIC	1185	1209	1368	1300	1439	-1%	71
CERN(new)	4220	4247	4247	4500	4500	+9%	253

Summary by type of data

Data type	Data Tier	Retention	Size(TB)
ProtoDUNE Raw	raw	physics	5772
ProtoDUNE trigprim	trigprim	test	1787
Coldbox raw	raw	study	609
Monte Carlo	Full-reconstructed	physics	2866
Transient files	miscellanous	n/a	305
Duplicates of above			~1500
Free space		n/a	3467

Will try to get all this in a pretty pie chart in a couple of months.

How Much Disk Space Do We Have:

- 3467 TB free
- 19812 TB total.
- Long-term datasets shown in previous slides account for 11869TB
  - In general we only have one copy of reco (MC or data) on disk where there should be two.
- In FY2025 expect O(5PB) from ProtoDUNE Vertical Drift
- Plus 1.3PB for Low Energy fardet-hd and vd production now in process.
- Above figure of total and free disk space does not count the 7.5 PB of disk in front of Fermilab tape-backed dCache which currently is the bulk of the US disk.
  - Will be shifting 5PB of the 7.5PB of disk in front of tape-backed Fermilab dCache to persistent to follow CMS model that disk is disk, both in pledges and reality.

Plan—have 3.8PB free and need 6+

- Remove hitreco MC files from disk save O(300TB)
- Remove duplications of full-reconstructed PDSPProd4 data / PDSPProd4a MC O(200TB)
- Garbage collection of files with no rules. O(1PB)
- Talk to protodune-dp analyzers what are they using, roll unused stuff off of disk
- Ditto with coldbox stuff
- Split the trigger primitives off into separate data sets, roll them off of disk.
- Shift more disk at Fermilab into Rucio-managed disk.
- This will make room for all that we know is coming.
  - But in the wild right now we only have one version on disk, not two.
  - For example, should have PDSPProd2 and PDSPProd4 both still on disk.
- Hoping for a disk buy at Fermilab this year to help.

BACKUP

Overall RSE usage (TB) Oct 25

RSE	RuleUsage	RucioUsage	StorageUsage	DuneproLimit	StorageLimit	Change (R.U)
BNL	597	756	749	830	831	+46%
FNAL_DISK	1634	1670	1670	2100	2674	+44%
PRAGUE	776	834	856	1126	1161	+2%
NIKHEF	1008	1102	1102	1143	1143	+27%
SURFSARA	456	614	614	781	756	-6%
GLASGOW	495	495	495	500	500	+99%
LANCASTER	452	462	462	700	700	-0%
MANCHESTER	603	603	603	1000	1077	-35%
QMUL (down)	178	334	307	1100	1100	0%
RAL-PP	499	637	654	1000	1000	+18%
RAL_ECHO	783	933	929	1050	1000	-5%
CNAF	3	11	3	300	550	+2%
CCIN2P3	1000	1074	1065	1100	1181	+1%
PIC	1200	1224	1222	1300	1439	+35%
CERN(new)	3826	3852	3852	4500	4500	+23%