DUNE Offline Computing Model Calculations for 2024

H. Schellman for the Computing Consortium

February 11, 2024– version 2

1 Introduction

This is an annual projection for DUNE CPU and storage needs intended for use at the Computing Contributions Board meeting in February 2024. It projects needs for 2024 to 2028.

The overall computing model and 2023 projections for DUNE were described in chapters 6-13 of the recent (Oct. 2022) DUNE Conceptual Design Report [1]. This document provides updates on resource needs for 2024-2025.

The 2024 projection is done using codes at: https://github.com/DUNE/CCB-data/tree/CCB-Feb24/Numbers-2024/ from parameters stored in a json and csv file. We use CPU and storage sizes derived from protoDUNE and simulation experience and apply them to projected numbers of events from the various DUNE detectors.

This note summarizes pledges, usage and projected need for the CCB.

One Line Executive Summary: Because the 2023 request assumed ProtoDUNE-2 running in 2023, the requests for 2024 are similar to the 2023 requests.

Changes since the last report [2, 3] include:

- Delay in the start of ProtoDUNE-2 running at CERN until mid-2024. This moves the need for storage/CPU into mid 2024. As the 2023 estimates were made assuming a 2023 run, most of the necessary resources are already in place.
- Removal of the memory weighted core concept as it caused too much confusion and most available cores now have more than 2000 MB available.
- Switch to HS23-Years as opposed to Core-Years as the main CPU reporting system.
- Revisions to near-term requests based on the 2023 experience including a hold on tape requests from the collaboration during protoDUNE activities.
- Changes to splits (Table 2) between regions in disk and CPU provision.

2 Model Summary

Resource requests are based on a processing and storage model that separated different phases of the detectors and includes 4 classifications of activity.

- Data - described by number of events, CPU time/event, storage/event
- Simulation - described by number of events, CPU time/event, storage/event
- Test - described by PB of storage, assumed not to consume large amounts of CPU
- Analysis - described by a scaling factor relative to Reconstruction and Simulation CPU.

Table 1 summarizes the large detectors and their abbreviations. Smaller prototypes that generate less than 1 PB of data are not explicitly included.

Abbrev.	Detector	Running time
PDSP	ProtoDUNE Single Phase	2018-2021
PDDP	ProtoDUNE Dual Phase	2018-2021
PDHD	ProtoDUNE-2 Horizonal Drift	2024-2026
PDVD	ProtoDUNE-2 Vertical Drift	2024-2026
2x2	Prototype NDLAr Detector	2024-2026
HD	Far Detector Horizonal Drift	2028-
VD	Far Detector Vertical Drift	2028-
NDLAr + TMS	Near Detector Liquid Argon + Muon System	2030-
SAND	On Axis Near Detector	2030-

Table 1: Detector descriptions.

These inputs are then used to calculate storage and CPU needs.

Disk storage policies are designed to optimize access and minimize CPU inefficiencies due to network transfer speeds. In particular, all recent data samples are planned to have at least one disk copy to avoid the need to access tape. As raw data processing is not I/O bound, one disk copy has been found to be sufficient. Analysis of reconstructed samples is I/O bound so multiple copies located closer to CPU are desirable.

Each type of data has a storage retention policy which includes lifetimes on disk and tape and number of copies on disk and tape. For example, raw data has a very long tape retention policy and 2 tape copies, with a 1 year stay of 1 copy on disk. Recent simulated and reconstructed data samples typically have 1 tape copy, 2 disk copies and disk retention times of 1.5-2 years to allow fast analysis. We assume one processing campaign per year, with all real data reprocessed each campaign. Simulation is assumed to be redone as well but older samples are not reprocessed. An extended retention time after the end of data taking for the last version of sim/reco is added to allow for extended data analysis.

The model currently does not include any disk headroom allowance for data movement. This will be added in the 2025 version of the model.

CPU calculations are now made in HS23 units, with results from 2023 reported in both Core and HS23 units for reference. Calculations for the protoDUNE's and DUNE far detector are based on existing reconstruction and simulation experience. We have much less experience with Near Detector codes so those estimates are less accurate.

We assume protoDUNE running in 2024-2025 with startup of DUNE FD commissioning in 2027 (Test stream) and data taking in 2029.

		FNAL:	CERN	Global
Disk				
	Raw	0.50	0.50	0.00
	Sim:	0.40	0.10	0.50
	Reco-Data:	0.40	0.10	0.50
	Test:	0.50	0.50	0.00
Tape				
	Raw:	0.50	0.50	0.00
	Rest"	1.00	0.00	0.00
CPU				
	All:	0.40	0.10	0.50

Table 2: Proposed division between FNAL/CERN/Global for storage and CPU in the near term, until $\sim$2028, when FD replaces ProtoDUNE as the primary source of experimental data. The tape division is not yet finalized as we work on integration of Global tape archives. In the long run, Global sites are expected to take over some of the tape provision currently provided by CERN.

3 Pledge Requests

Pledge requests for 2024 are similar to those for 2023. Similar because the impacts of ProtoDUNE running were already included in the 2023 request.

These pledge requests include only general resources available to the DUNE collaboration, not the additional local resources available to collaborators at their own institutions.

The near-term request does not include much Near Detector simulation. That is currently being done via special allocations at NERSC and ANL and not yet included in our accounting.

The model for proposed pledges for 2024 is detailed in sections: Disk(4.1.1), Tape(4.2.3) and CPU(4.3.3) below. A summary of requests and actual use for 2023 is shown in Table 3 and the requests for 2024 are shown in Table 4

	Disk (PB)	Actual	Tape(PB)	Actual (PB)	CPU (Core-years)	Actual
Model	25.80	–	45.5	–	15,169	–
Request						
FNAL	7.80	9.80	36.2	27.1	3,792	2,844
CERN	2.60	4.02	9.2	5.7	3,792	140
Global	15.40	9.76	0.1	0.1	7,585	1, 524
Total	25.80	23.58	45.5	32.9	15,169	4,515

Table 3: Requests and actual usage from the previous year (2023). CPU and disk utilization were low due to the delay in ProtoDUNE running. Table 6 shows the details of disk use by site. CPU usage is detailed in Table 7. In this table CPU is in units of unweighted core-years as that was the metric used prior to 2024.

	Disk (PB)	Tape(PB)	CPU (kHS23-years)	CPU (Core-years)
Model	22.2	40.0	75.4	6856
Request				
FNAL	9.5	32.3	30.2	2742
CERN	4.7	7.7	7.5	686
Global	8.0	–	37.7	3428
Total	22.2	40.0	75.4	6856

Table 4: Requests for 2024. The disk requests reflect the different data types and the proposed splits from Table 2. They do not include the normal headroom of 5-10%. Tape pledges reflect the dominant use of CERN and FNAL for archival storage of data. CPU pledges are in units of kHS23-years with Core-years provided for comparison to 2023.

4 Supporting Materials

Generally, raw data are stored on tape at both CERN and FNAL. Simulation and reconstructed data have one tape copy at Fermilab and recent reconstructed and simulated samples have one (or two) disk copies with one at Fermilab and one in Europe.

CERN and FNAL have special responsibilities for archival data storage and for disk space for raw data while contributions from other collaborating institutions are aggregated under the heading Global, which includes US sites outside of FNAL. The proposed near-term (until FD data taking) split between FNAL, CERN and Global contributions is shown in Table 2. The pledges proposed here deviate slightly from those numbers with larger contributions from FNAL due to resources already in place.

4.1 Disk

Table 6 summarizes the disk utilization reported by sites via [4], supplemented by rucio reports. Some sites, notably TIFR, are not yet fully integrated so do not show up in the rucio reports.

Figure 1 summarize the cumulative disk needs and requests projected by our model. These numbers are used to generate the request for 2024. They are divided into the two host laboratories (CERN and FNAL) and Global, which includes contributions from the rest of the collaboration, including BNL and NERSC in the US.

Table 3 summarized the pledges from previous years compared to the actual amounts allocated shown in Table 5 . Because of the delay in ProtoDUNE running, disk that had been planned for ProtoDUNE data and simulation was not fully used.

Country	Allocated (TB)	used (TB)	account lim (TB)	rse usage (TB)	percent used
CERN	4020	2121	3000	1957	53
CZ	1126	561	1000	561	50
ES	720	192	1000	192	27
FNAL	9830	7009	9830	12239	71
FR	537	180	500	180	33
IN	750	0	0	0	0
IT	300	0	300	0	0
NL	1899	551	1750	556	29
UK	4303	3250	3950	3149	76
US	850	14	850	2	2
Total	24334	13877	22180	18836	57

Table 5: Disk allocations and usage across countries at the end of 2023. These numbers are derived from usage reports, rucio reports and from cross-checks with individual sites on 2024-02-01. The percentages are Used/Allocation.

4.1.1 Request Summary

The total disk request for 2024 is 22.2PB. This is consistent with the current allocated space of 24.3 PB. Because 5-10% headroom is needed, we request that the current allocated 24.3 PB be retained for 2024.

Country	RSE	Allocation (TB)	Used (TB)	account lim (TB)	rse usage (TB)	percent used
CERN	CERN PDUNE EOS	3020	1626	2000	1504	54
CERN	DUNE CERN EOS	1000	495	1000	454	49
CZ	PRAGUE	1126	561	1000	561	50
ES	DUNE ES PIC	720	192	1000	192	27
FNAL	DUNE US FNAL DISK STAGE	1000	809	1000	809	81
FNAL	FNAL DCACHE	8780	6200	8780	11429	71
FNAL	FNAL DCACHE STAGING	50	0	50	0	0
FR	DUNE FR CCIN2P3 DISK	537	180	500	180	33
IN	DUNE IN TIFR	750	0	0	0	0
IT	DUNE IT INFN CNAF	300	0	300	0	0
NL	NIKHEF	1143	546	1000	550	48
NL	SURFSARA	756	5	750	6	1
UK	DUNE UK LANCASTER CEPH	500	485	500	485	97
UK	MANCHESTER	1078	1033	1000	976	96
UK	QMUL	1100	651	1000	770	59
UK	RAL-PP	625	116	400	99	19
UK	RAL ECHO	1000	965	1050	819	96
US	DUNE US BNL SDCC	830	13	830	2	2
US	T3 US NERSC	20	1	20	1	3

Table 6: Disk allocations and usage across sites. These numbers are derived from usage reports, rucio reports and from cross-checks with individual sites on 2024-02-01. The percentages are Used/Allocation.

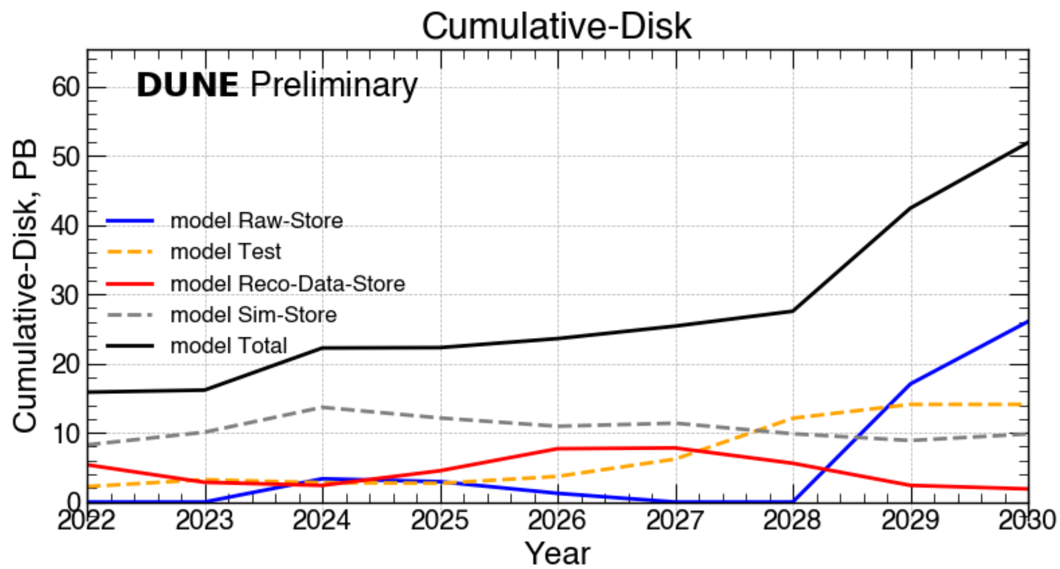

Cumulative-Disk	2022	2023	2024	2025	2026	2027	2028	2029	2030
Raw-Store(PB)	0.0	0.0	3.4	2.9	1.3	0.0	0.0	17.1	26.1
Test(PB)	2.2	3.2	2.8	2.7	3.7	6.2	12.1	14.1	14.1
Reco-Data-Store(PB)	5.4	2.8	2.4	4.5	7.7	7.8	5.6	2.4	1.9
Sim-Store(PB)	8.2	10.1	13.7	12.1	10.9	11.4	9.9	8.9	9.8
Total(PB)	15.9	16.2	22.2	22.3	23.6	25.4	27.6	42.4	51.

Cumulative-Disk	2022	2023	2024	2025	2026	2027	2028	2029	2030
Global(PB)	6.8	6.5	8.0	8.3	9.3	9.6	7.7	15.5	15.8
FNAL(PB)	6.6	6.8	9.5	9.5	9.9	10.8	12.2	26.9	36.1
CERN(PB)	2.5	2.9	4.7	4.5	4.3	5.0	7.6	0.0	0.0
Total(PB)	15.9	16.2	22.2	22.3	23.6	25.4	27.6	42.4	51.

Figure 1: Cumulative Disk needs in PB. Includes data lifetimes. The top table shows the source of the data while the bottom table shows the proposed split using the fractions from Table 2 and a modified version which reflects the disk already in place at FNAL and CERN, thus reducing the Global request.

4.2 Tape

4.2.1 Current Status

DUNE currently has ∼27.1 PB of data[5] on tape at Fermilab and 5.7 PB of raw protoDUNE data at CERN[4].

4.2.2 Model Projections

Figure 2 summarizes the cumulative tape needs projected by the model. These numbers are used to generate the requests for 2024. They are divided into the two host laboratories (CERN and FNAL). The UK and the IN2P3 have offered ~ 3 PB of tape archive but it has not yet been smoothly integrated into our data flow. We therefore substantially reduce our request for resources at those sites to ~ 100 TB per site, to allow testing of integration, with an increased request anticipated future years. .

The increase in 2024-2025 is mainly storage of raw and processed data from the upcoming runs of PDHD and PDVD.

The exact division between FNAL, CERN and Global once DUNE starts taking FD data in 2029 is not yet defined.

4.2.3 Current Status and Request

We anticipate needing 40.0 PB of tape (an increase of 7 PB from 2023) to accommodate the ProtoDUNE run 2 data and increased simulation.

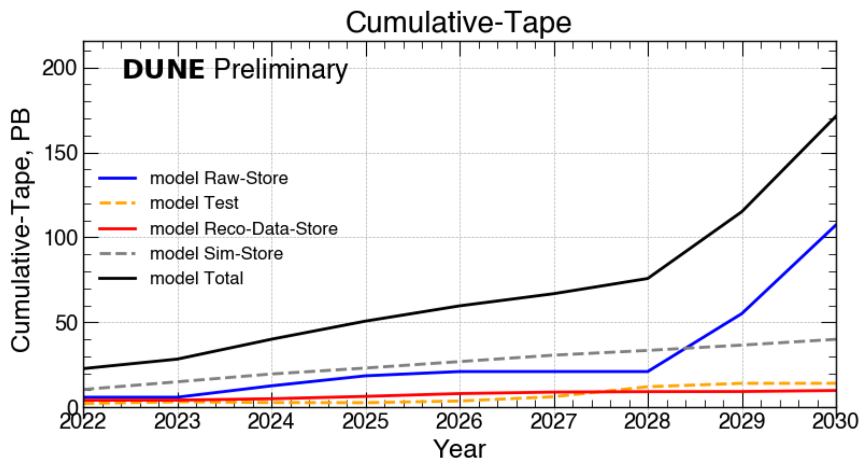

Cumulative-Tape	2022	2023	2024	2025	2026	2027	2028	2029	2030
Raw-Store(PB)	5.9	5.9	12.6	18.5	21.0	21.0	21.0	55.1	107.3
Test(PB)	2.2	3.2	2.8	2.7	3.7	6.2	12.1	14.1	14.1
Reco-Data-Store(PB)	4.2	4.2	5.0	6.4	8.1	8.9	9.2	9.3	9.9
Sim-Store(PB)	10.5	15.0	19.6	23.1	26.9	30.7	33.5	36.6	40.0
Total(PB)	22.8	28.3	40.0	50.7	59.7	66.8	75.8	115.1	171.

Cumulative-Tape	2022	2023	2024	2025	2026	2027	2028	2029	2030
Global(PB)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	43.8	58.8
FNAL(PB)	18.7	23.8	32.3	40.1	47.3	53.2	59.2	57.6	85.6
CERN(PB)	4.1	4.6	7.7	10.6	12.4	13.6	16.6	13.8	26.8
Total(PB)	22.8	28.3	40.0	50.7	59.7	66.8	75.8	115.1	171.

Figure 2: Cumulative Tape needs from the model in PB, includes data lifetimes. The top table shows the origin of the data while the bottom table shows the proposed split. Global contributions are set low in 2024 and grow thereafter as more tape archives are integrated. The exact division between FNAL, CERN and Global once DUNE starts taking FD data in 2029 is not yet defined.

4.3 CPU

4.3.1 Current Status

Tables 7 and 8 show the CPU used by country in calendar 2023. Due to the delay in ProtoDUNE running CPU use was dominated by code development and analysis of the older ProtoDUNE data. Total CPU use was substantially lower than projected.

Site	Production	Analysis	NoMARS	MARS	Total
BR	26	9	36	0	36
CA	7	14	21	0	21
CERN	57	83	140	0	140
CH	2	0	2	0	2
CZ	51	48	99	0	99
ES	19	25	45	0	45
FR	34	24	59	0	59
IN	5	33	38	0	38
IT	0	7	7	0	7
NL	110	88	199	0	199
RU	4	19	24	0	24
UK	287	397	684	0	684
US FNAL	156	1940	2096	751	2848
US OSG	214	86	301	0	301
undefined	0	8	9	0	9
Total	979	2788	3768	751	4520

Table 7: CPU utilization in Core Years for calendar 2023 divided by use case. This is for comparison with the previous pledging system. Production includes official reconstruction and simulation. Analysis is user analysis of data. MARS is beamline simulations performed at Fermilab. NoMARS sums just Production and Analysis.

4.3.2 Model calculation

The wall-time estimates in the model are created by estimating the number of simulated and raw events taken and then scaling by the measured CPU time per event on a gpvm corrected to wall-time by the estimated efficiency (default 70%). Each data type also has an additional "Analysis" factor of 0.5 to 1.0 to account for analysis after the original reconstruction.

The numbers for PD and FD are based on substantial experience with the earlier ProtoDUNE runs and large-scale simulation campaigns. The numbers for ND are much more uncertain and await experience with simulation and reconstruction of results from the 2x2 demonstrator which will run later in 2024 at Fermilab. The model shows ProtoDUNE dominating through the 2024-2025 runs and subsequent analysis period with testing and simulation for the DUNE near and far detectors taking over circa 2028.

Site	Production	Analysis	NoMARS	MARS	Total
BR	0.296	0.102	0.398	0.000	0.398
CA	0.085	0.156	0.241	0.000	0.241
CERN	0.627	0.921	1.548	0.000	1.548
CH	0.030	0.003	0.033	0.000	0.033
CZ	0.562	0.530	1.092	0.000	1.092
ES	0.218	0.281	0.499	0.000	0.499
FR	0.384	0.270	0.654	0.000	0.654
IN	0.060	0.368	0.428	0.000	0.428
IT	0.000	0.083	0.083	0.000	0.083
NL	1.215	0.975	2.190	0.000	2.190
RU	0.052	0.217	0.270	0.000	0.270
UK	3.158	4.373	7.531	0.001	7.532
US FNAL	1.718	21.346	23.063	8.265	31.328
US OSG	2.364	0.955	3.319	0.001	3.320
undefined	0.007	0.097	0.104	0.000	0.104
Total	10.778	30.676	41.454	8.267	49.721

Table 8: CPU utilization in kHS23-Years for calendar 2023 divided by use case. Production includes official reconstruction and simulation. Analysis is user analysis of data. MARS is beamline simulations performed at Fermilab. NoMARS sums just Production and Analysis.

Figure 3 shows the projected wall-time (HS23-yrs) need projections through 2030. They are divided into the two host laboratories (CERN and FNAL) and Global, which includes contributions from the rest of the collaboration, including OSG, BNL and NERSC in the US.

4.3.3 CPU Request

The advent of ProtoDUNE-2 running in 2024 and ramp-up of simulation for the FD and ND will lead to somewhat increased needs for CPU resources. For 2024, we are requesting 75.4 kHS23-Yrs.

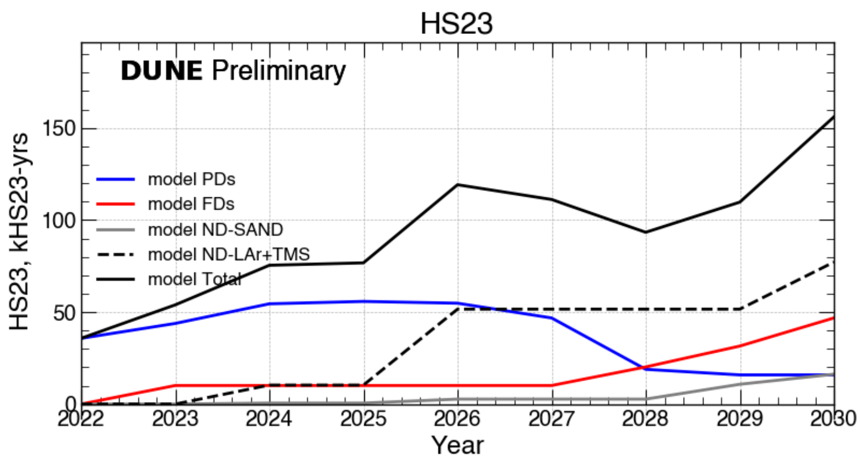

HS23	2022	2023	2024	2025	2026	2027	2028	2029	2030
PDs (kHS23-yrs)	35	43	54	55	54	46	18	15	15
FDs (kHS23-yrs)	0	10	10	10	10	10	20	31	46
ND-SAND (kHS23-yrs)	0	0	0	0	2	2	2	10	16
ND-LAr+TMS (kHS23-yrs)	0	0	10	10	51	51	51	51	77
FNAL (kHS23-yrs)	14	21	30	30	47	44	37	54	78
CERN (kHS23-yrs)	3	5	7	7	11	11	9	0	0
Global (kHS23-yrs)	17	26	37	38	59	55	46	54	78
Total (kHS23-yrs)	35	53	75	76	119	111	93	109	156

Figure 3: Model of CPU wall-time needs through 2030 in kHS23-yr units. The top 4 lines show division by detector, the bottom 4 show the proposed division between FNAL, CERN and Global based on the fractions in Table 2.

References

[1] **DUNE** Collaboration, A. Abed Abud *et al.*, "DUNE Offline Computing Conceptual Design Report," arXiv:2210.15665 [physics.data-an].

[2] H. Schellman, "DUNE Offline Computing Model Calculations for 2023." https://docs.dunescience.org/cgi-bin/sso/ShowDocument?docid=27192.

[3] P. Clarke, "DUNE Computing Contributions Board Minutes for 2023." https://docs.dunescience.org/cgi-bin/sso/ShowDocument?docid=27462.

[4] Yuan, Wenlong, "DUNE Rucio Storage Monitoring Site." https://dune.monitoring.edi.scotgrid.ac.uk/app/dashboards#/view/7eb1cea0-ca5e-11ea-b9a5-15b75a959b33?_g=(filters:!(),refreshInterval:(pause:!t,value:0),time:(from:now-1d,to:now)).

[5] Timm, Steven, "DUNE Tape usage at Fermilab." https://landscape.fnal.gov/monitor/d/r6UDhH-iz/fcrsg-prep?var-expName=DUNE&var-graphiteQuery=%7Bdune%2CDUNE%2Cmarslbne%2Clbne%7D&var-esQuery=%28dune+OR+DUNE+OR+marslbne+OR+lbne%29&var-NASQuery=%28dune*+OR+DUNE*+OR+marslbne*+OR+lbne*%29&var-dCacheQuery=%28dune%7CDUNE%7Cmarslbne%7Clbne%29&var-enstoreQuery=%28dune%7CDUNE%7Cmarslbne%7Clbne%29&orgId=1&viewPanel=8.

A Model Details

This appendix shows the parameters used in the model and plots of all the input and derived quantities as a function of time.

Resource needs for reconstructed data for a given year are based on the number of events produced over the previous "Reprocess" years. For ProtoDUNEs that is 2-4 years.

Simulation resource needs are instead calculated based on a number of simulation events each year. The assumption is that new software versions imply resimulation.

Disk and tape lifetimes for different data types are specified as well as the desirable number of copies.

The splits parameters make CERN responsible for raw data until 2027 with the collaboration taking over after that point.

`Detectors`: Detectors included in the calculation = `['SP', 'PDHD', 'DP', 'PDVD', 'HD', 'VD', 'ND-SAND', 'ND-LAr+TMS']`
`Cap`: Cap on Raw data/year in PB = 30
`Base-Memory`: MB of memory per slot assumed as the average = 2000
`MaxYear`: Plot until year = 2030
`MinYear`: Plot starting with year = 2022
`Reprocess`: Number of years of data reprocessed when doing a new pass = `'SP': 3, 'DP': 3, 'PDHD': 3, 'PDVD': 3, 'PDs': 3, 'VD': 100, 'HD': 100, 'FDs': 100, 'ND-SAND': 100, 'ND-LAr+TMS': 100`
`AnalysisExtend`: Years analysis continues after last reco/sim = 3
`PatternFraction`: Fraction of time taken in pattern recognition = `'SP': 0.7, 'PDHD': 0.7, 'DP': 0.7, 'PDVD': 0.7, 'PDs': 0.7, 'HD': 0.1, 'VD': 0.1, 'FDs': 0.1, 'ND-SAND': 0.9, 'ND-LAr+TMS': 0.9, 'MARS': 0`
`TapeLifetimes`: Number of years kept on tape = `'Raw-Store': 100, 'Test': 1.0, 'Reco-Data-Store': 15, 'Sim-Store': 15`
`DiskLifetimes`: Number of years kept on disk = `'Raw-Store': 1, 'Test': 1.0, 'Reco-Data-Store': 3, 'Sim-Store': 2`
`TapeCopies`: Number of copies kept on tape = `'Raw-Store': 2, 'Test': 1, 'Reco-Data-Store': 1, 'Sim-Store': 1`
`DiskCopies`: Number of copies kept on disk = `'Raw-Store': 1, 'Test': 1, 'Reco-Data-Store': 2, 'Sim-Store': 1.5`
`PerYear`: Number of reprocessing done per year = `'Raw-Store': 1, 'Test': 1, 'Reco-Data-Store': 1, 'Sim-Store': 1, 'Events': 1, 'Sim-Events': 1, 'Reco-Data-CPU': 1, 'Sim-CPU': 1, 'Analysis': 1, 'Analysis-CPU': 1, 'Reco-Data-GPU': 1, 'Sim-GPU': 1`
`Cores`: Description of cores, efficiency and speed relative to 2020 vintage = `'Efficiency': 0.7, '2020Units': 1`
`kHEPSPEC06PerCPU`: kHEPSPEC06 per core assumed = `0.011`
`SplitsYear`: Year CERN no longer responsible for disk or tape = 2029
`SplitsEarly`: Division between FNAL/CERN/Global for storage until SplitsYear = `'Tape': 'Raw-Store':`

'FNAL': 0.5, 'CERN': 0.5, 'Global': 0.0, 'Sim-Store': 'FNAL': 1.0, 'CERN': 0.0, 'Global': 0.0, 'Reco-Data-Store': 'FNAL': 1.0, 'CERN': 0.0, 'Global': 0.0, 'Test': 'FNAL': 0.5, 'CERN': 0.5, 'Global': 0.0, 'Disk': 'Raw-Store': 'FNAL': 0.5, 'CERN': 0.5, 'Global': 0.0, 'Sim-Store': 'FNAL': 0.4, 'CERN': 0.1, 'Global': 0.5, 'Reco-Data-Store': 'FNAL': 0.4, 'CERN': 0.1, 'Global': 0.5, 'Test': 'FNAL': 0.5, 'CERN': 0.5, 'Global': 0.0, 'CPU': 'CPU': 'FNAL': 0.4, 'CERN': 0.1, 'Global': 0.5

SplitsLater: Division between FNAL/CERN/Global for storage after SplitsYear = 'Tape': 'Raw-Store': 'FNAL': 0.5, 'CERN': 0.25, 'Global': 0.25, 'Sim-Store': 'FNAL': 0.5, 'CERN': 0.0, 'Global': 0.5, 'Reco-Data-Store': 'FNAL': 0.5, 'CERN': 0.0, 'Global': 0.5, 'Test': 'FNAL': 0.5, 'CERN': 0.0, 'Global': 0.5, 'Disk': 'Raw-Store': 'FNAL': 1.0, 'CERN': 0.0, 'Global': 0.0, 'Sim-Store': 'FNAL': 0.25, 'CERN': 0.0, 'Global': 0.75, 'Reco-Data-S 'FNAL': 0.25, 'CERN': 0.0, 'Global': 0.75, 'Test': 'FNAL': 0.5, 'CERN': 0.0, 'Global': 0.5, 'CPU': 'CPU': 'FNAL': 0.5, 'CERN': 0.0, 'Global': 0.5

filename: Input configuration file = NearTerm_2024-02-05-2040.json