



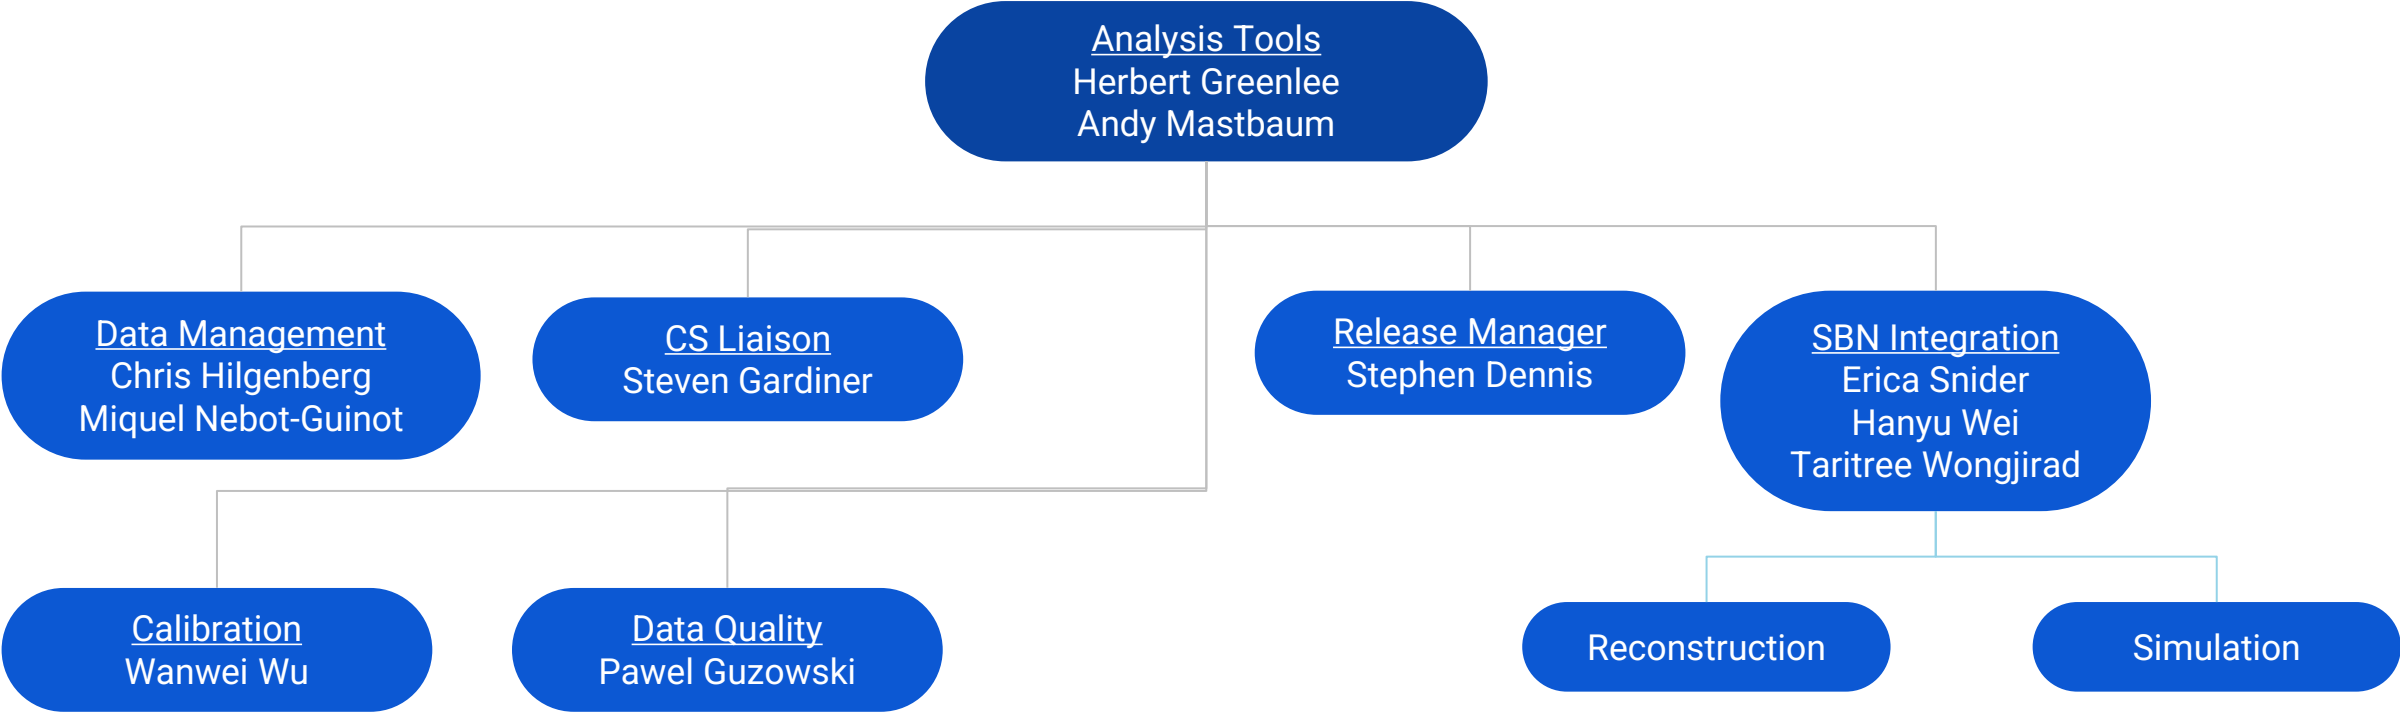
MicroBooNE FCRSG FY23

H. Greenlee

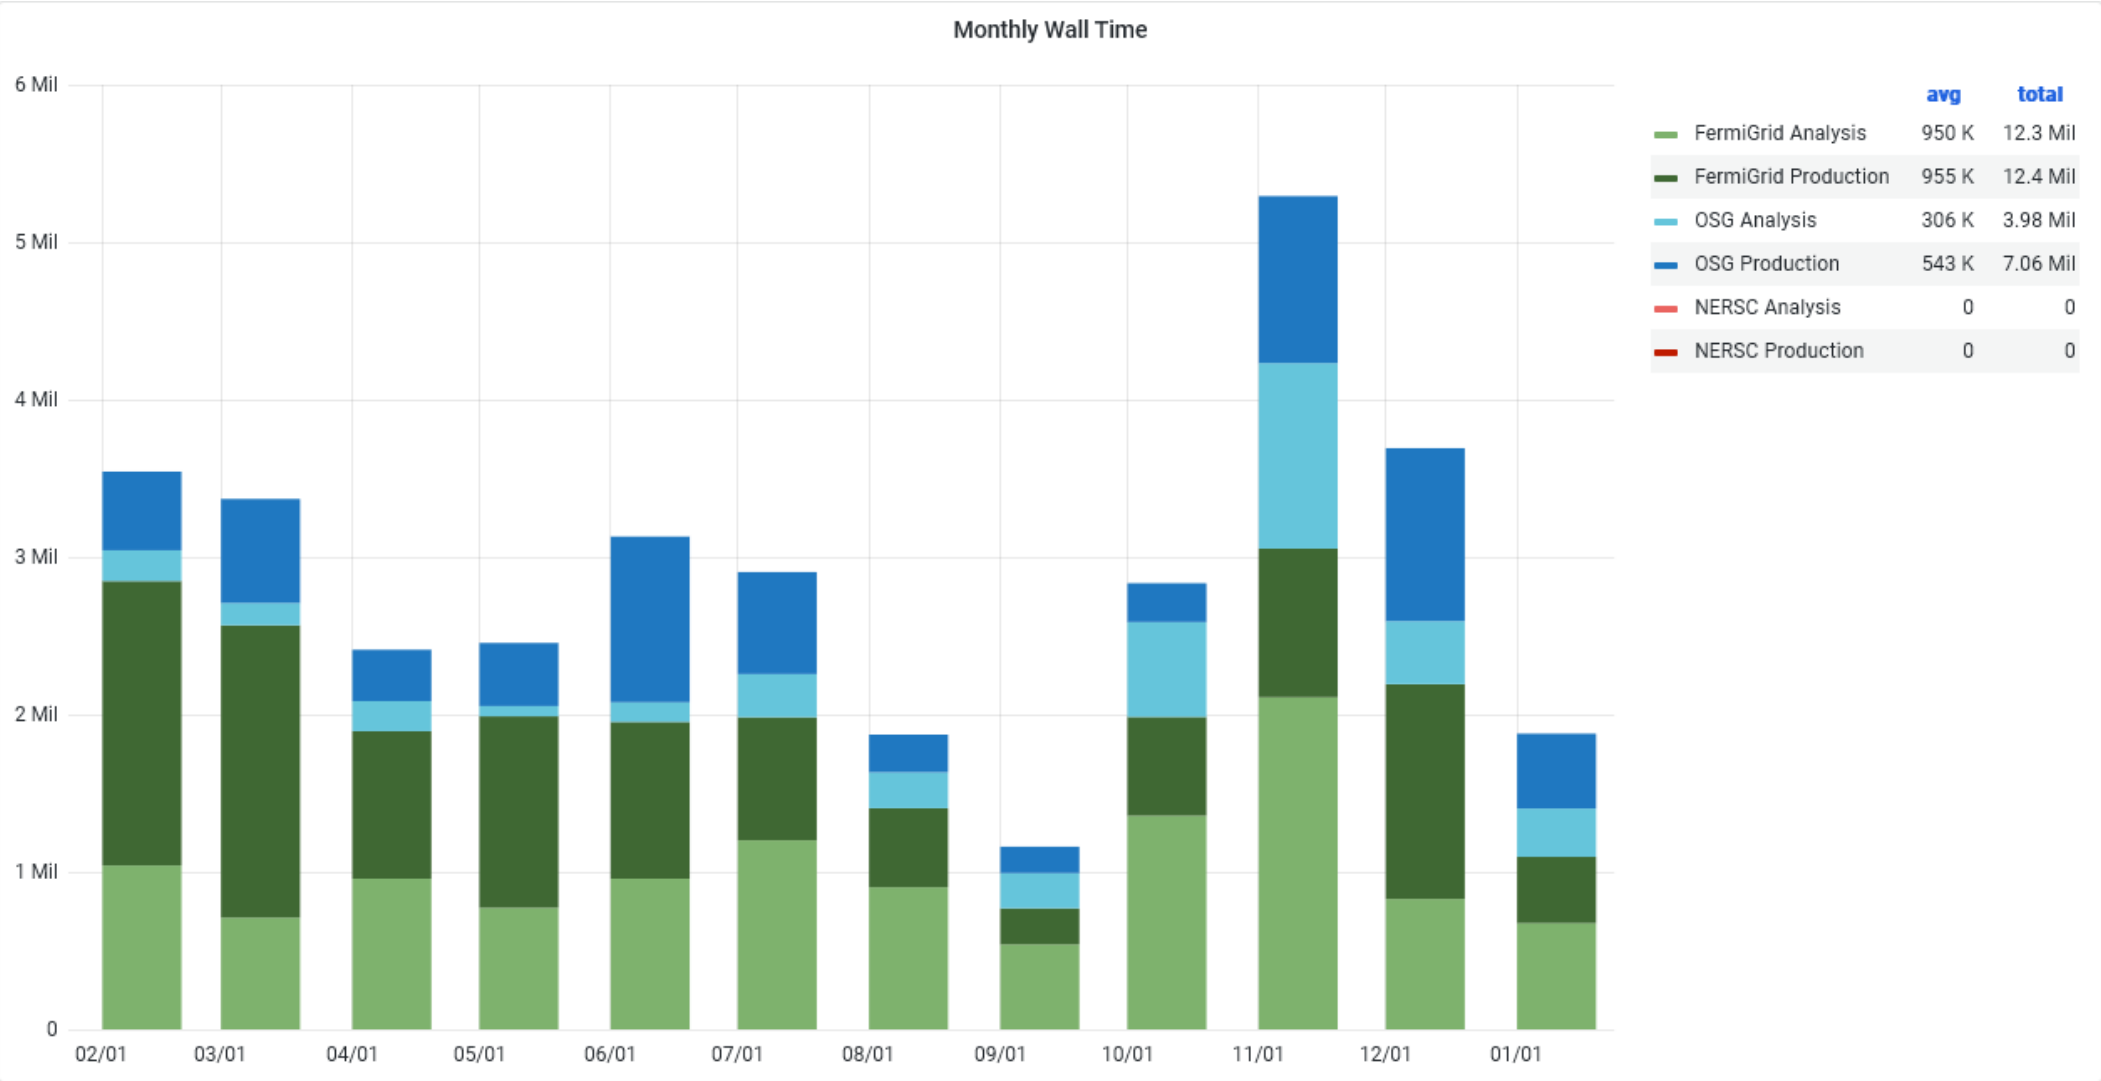
Overview of MicroBooNE Computing in 2023

- MicroBooNE took data for five years from 2015-2020 (known as runs 1-5).
 - MicroBooNE took detector R&D data in 2021.
 - MicroBooNE data taking is complete.
- MicroBooNE is the first large LArTPC experiment to do a complete physics analysis using automated reconstruction.
 - MicroBooNE has produced 54 publications (published and submitted) so far.
- MicroBooNE has been running data processing campaigns (data and MC) using software versions known as MCC9 since Dec., 2018.
 - In 2021, MicroBooNE released its first results on the MiniBooNE BNB Low Energy Excess (LEE), based on runs 1-3.
- MicroBooNE's major goal in 2023 is to complete the analysis of its full five-year dataset.
 - MicroBooNE is planning major production campaigns for run 4-5 data in 2023.


MicroBooNE Organization Chart for Offline Computing



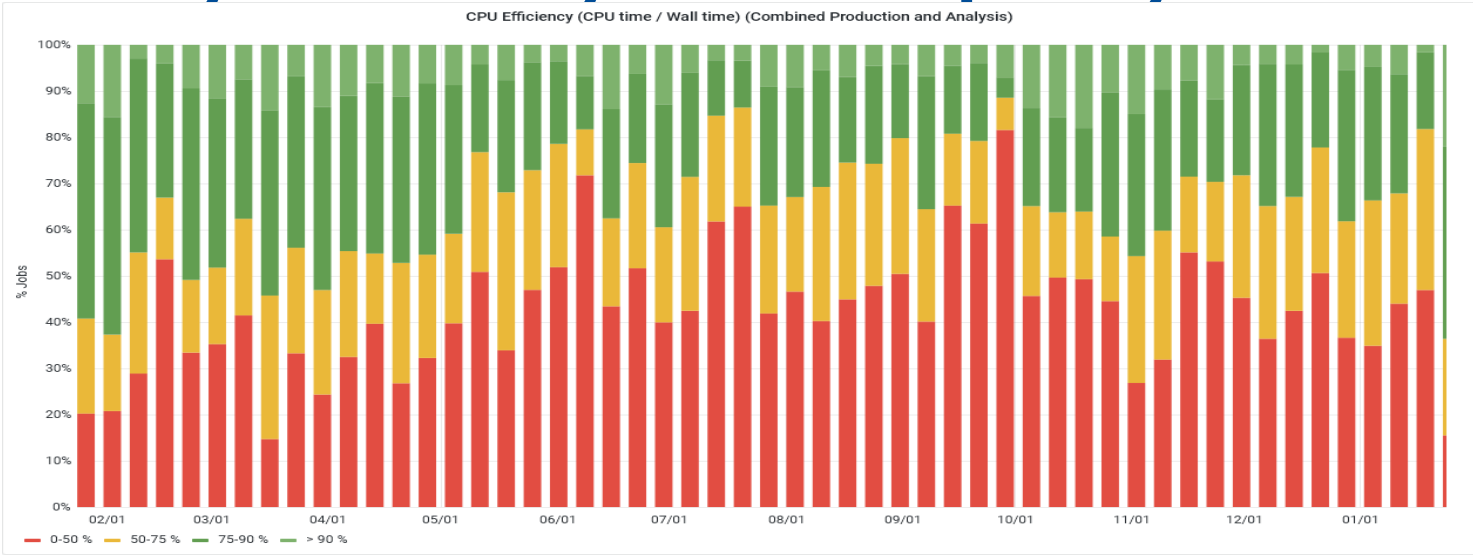
Experiment CPU Usage over the past year



Memory footprint over the past year



CPU and memory efficiency over the past year



What do you want to achieve in computing over the next 5 years?

Goals	Where does the experiment need to contribute	Where does SCD need to contribute
Complete analysis of runs 1-5 data using MCC9	Updated experiment software, calibrations, data quality.	Infrastructure & tools (sam, ifdhc, jobsub, databases, FermiGrid, dCache, enstore).
Improve integration of MicroBooNE software with other SBN experiments.	MicroBooNE experiment software	Support from larsoft and scisoft
Explore improvements in simulation and reconstruction beyond MCC9.	MicroBooNE experiment software	Support from larsoft and scisoft
Develop the ability to run statistical analysis using SBNSFit using HPC resources	Joint with other SBN experiments	General support

Campaign Schedules

	2023	2024	2025	2026	2027
Processing campaigns (start month-end month if known). Include when you expect to be prestaging	<ol style="list-style-type: none"> 1. Run 4/5 BNB data reco2: Feb-Aug 2. Run 4/5 BNB overlay MC: Apr-Dec 3. Run 4/5 NUMI data reco1: Jan-Jun 4. Run 4/5 NUMI data reco2: Mar-Sep 5. Run 4/5 NUMI overlay MC: Apr-Dec 6. Stage remaining runs 1-3 external unbiased data to production pool: Feb-Apr 	<ol style="list-style-type: none"> 1. New analyses requiring exotic or nonstandard MC signal samples. 			
Storage + CPU estimates (call out any special resource needs if known, e.g. HPC or GPU). Include amount(s) to be prestaged and file families, in addition to space needed for new outputs.	<ol style="list-style-type: none"> 1. 50 TB, 5.0e6 MWC @ 4 GB, input from disk. 2. 80 TB, 6.0e6 MWC @ 4 GB, input from disk 3. 100 TB, 4.0e6 MWC @ 4 GB, 500 TB input from tape (data_raw, 40% complete) 4. 100 TB, 10.0e6 MWC @ 4 GB, 100 TB input from disk 5. 80 TB, 6.0e6 MWC @ 4 GB, input from disk 6. 0 MWC, 520 TB input from tape (data_raw). 	<ol style="list-style-type: none"> 1. 10 TB, 1.0e6 MWC, input from disk. 			
Conference or result targets (month if known)	First combined BNB+NUMI oscillation results (runs 1-3).	First oscillation results using full runs 1-5 dataset (summer, Neutrino 2024 conference)		Combined SBN 3-detector oscillation results	

MWC = Hours * GB

Campaign Resources

- All MicroBooNE workflows that run on FermiGrid are also submitted to OSG opportunistic resources.
- MicroBooNE has access to 1000 high-memory batch slots and 600 TB of storage at the Bern OSG site.
 - We are planning on using this resource for the run 4/5 BNB reco2 campaign.
- MicroBooNE also has priority access at the Manchester OSG site (no storage).
 - Available for all campaigns.

CPU @ Fermilab Prediction Going Forward and Accuracy of Your Predictions [units of Million (1 CPU, 2GB) wall hours per CY]

	2019	2020	2021	2022	2023	2024	2025	2026	2027
Requested	20	54	43	30	30	25	20	10	
Actual Used	29.1	40.6	30.0	26.4	N/A	N/A	N/A	N/A	N/A
Efficiency	145%	75%	70%	88%	N/A	N/A	N/A	N/A	N/A



CPU – non-FNAL HTC Resources Going Forward and Accuracy of Your Predictions [units of Million (1 CPU, 2GB) wall hours per CY]

	2019	2020	2021	2022	2023	2024	2025	2026	2027
Requested	0	0	0	0					
Actual Used (OSG)	3.6	6.3	6.8	10.7	N/A	N/A	N/A	N/A	N/A
Efficiency	%	%	%	%	N/A	N/A	N/A	N/A	N/A

CPU – HPC Resources Going Forward and Accuracy of Your Predictions [units of Million (1 CPU, 2GB) wall hours per CY]

	2019	2020	2021	2022	2023	2024	2025	2026	2027
Requested									
Actual Used					N/A	N/A	N/A	N/A	N/A
Efficiency	%	%	%	%	N/A	N/A	N/A	N/A	N/A



CPU – GPU Resources Going Forward and Accuracy of Your Predictions [units of Million (1 CPU, 2GB) wall hours per CY]

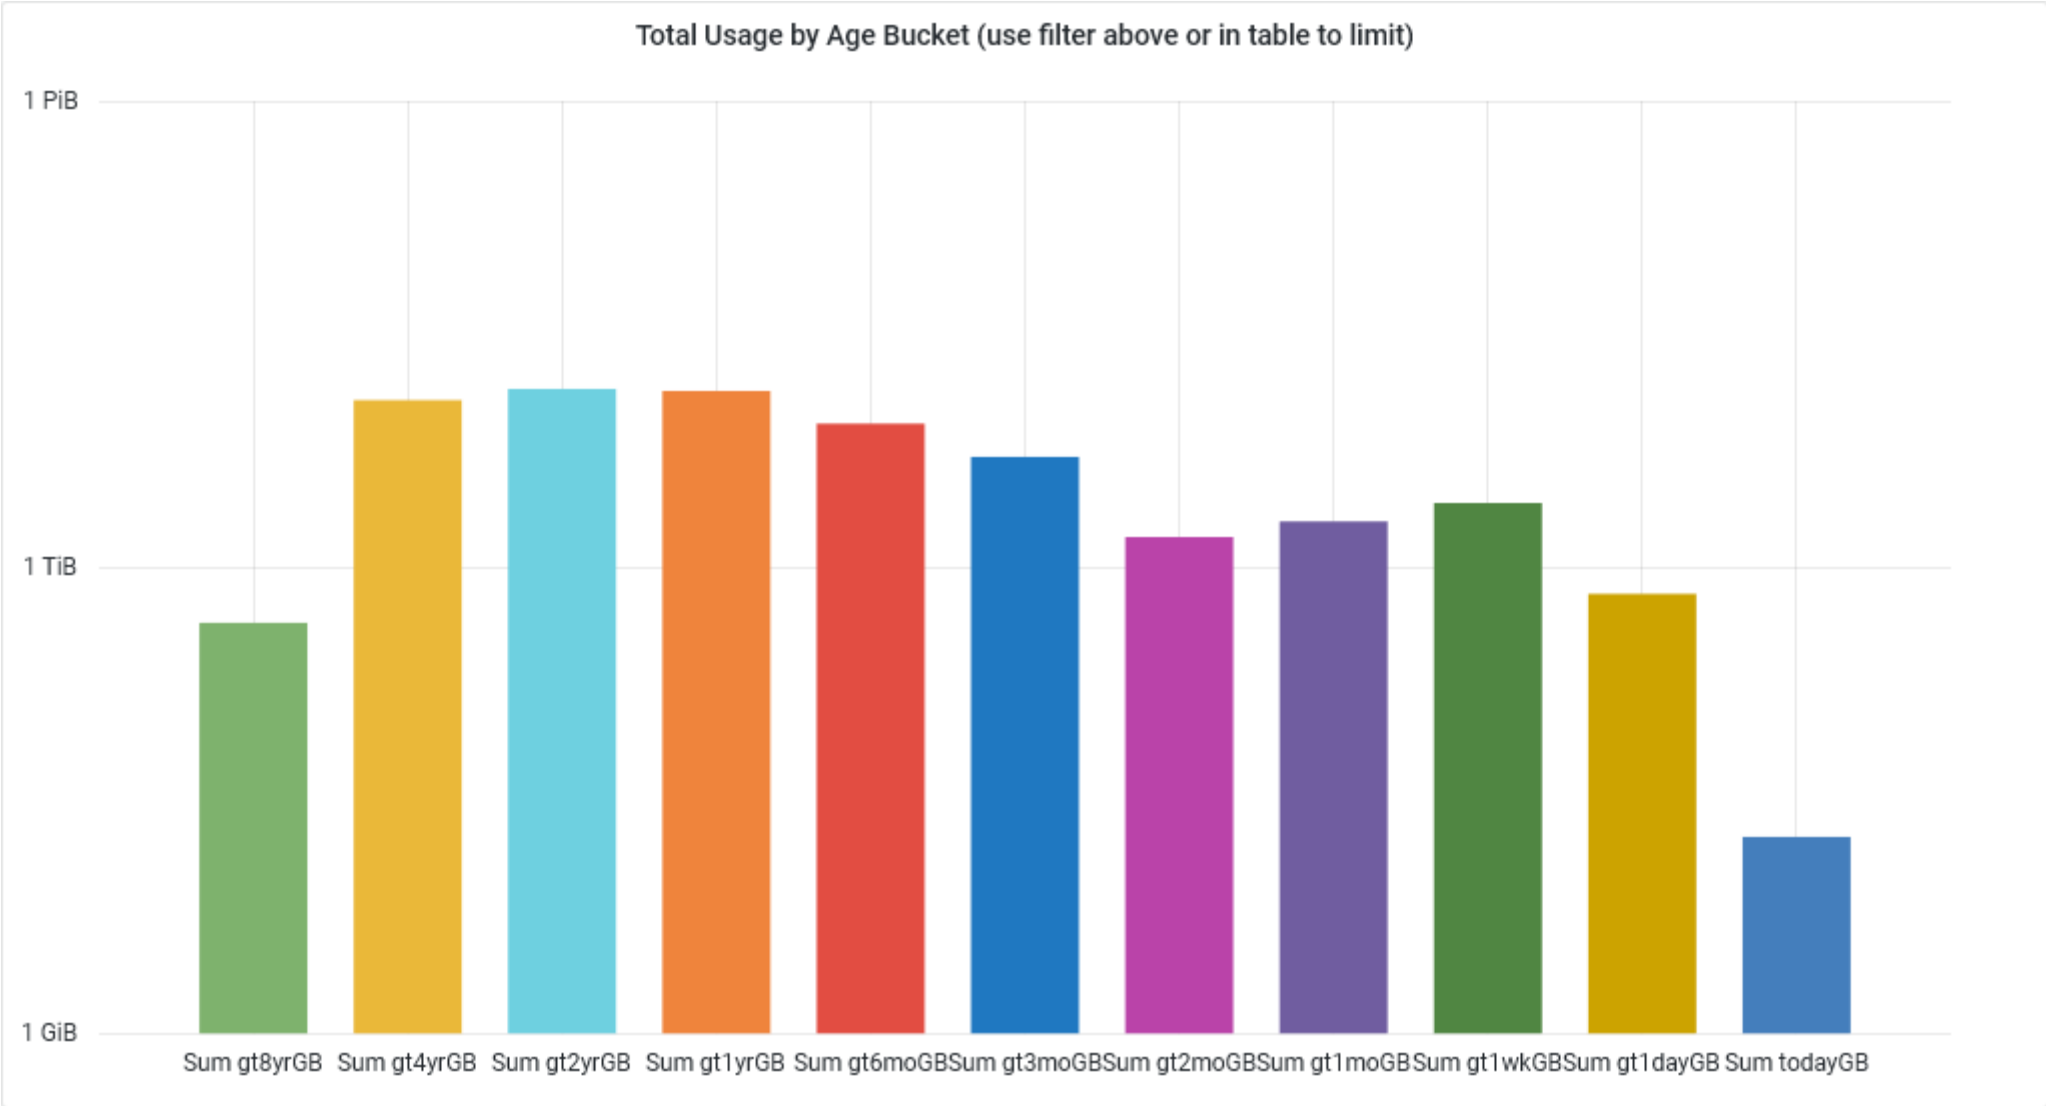
	2019	2020	2021	2022	2023	2024	2025	2026	2027
Requested									
Actual Used					N/A	N/A	N/A	N/A	N/A
Efficiency	%	%	%	%	N/A	N/A	N/A	N/A	N/A



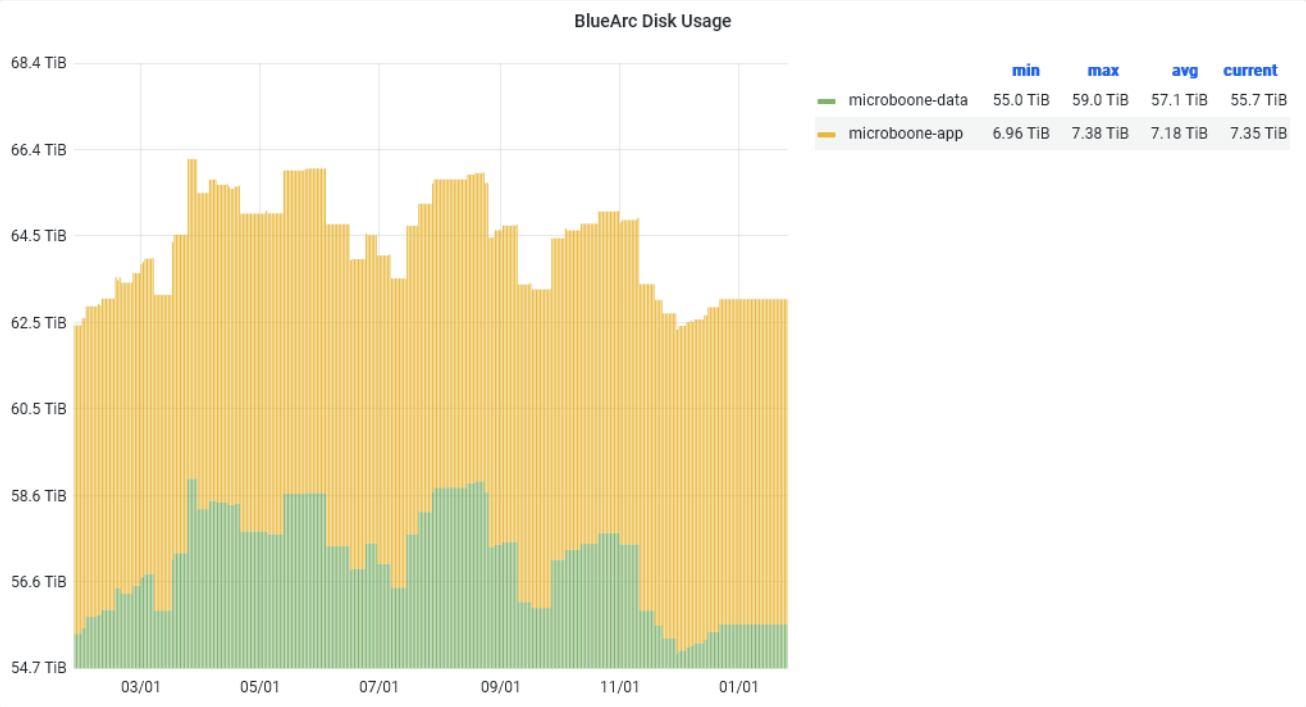
Comment on HPC and GPU Resources

- MicroBooNE has an allocation of 12.8M hours at the ANL/Theta HPC facility.
 - Previous attempts to run standard simulation and reconstruction workflows on HPC were not very successful.
- Current attempts to use HPC are focusing on compute-intensive statistical analysis using the SBNFit program.
 - MicroBooNE may request time from Fermilab Intensity Frontier/NERCS allocation (not sure how much yet).
 - May also be possible to access HPC resources at ANL/Theta.
- MicroBooNE is not requesting GPU resources from Fermilab.

Age of files in NAS



NAS Usage and Projections




MicroBooNE is not asking for an increase in NAS/Ceph. We just want to keep what we have.

	App	Data
2022	7.5 GiB	60 GiB
2023	7.5 GiB	60 GiB
2024	7.5 GiB	60 GiB
2025	7.5 GiB	60 GiB
2026	7.5 GiB	60 GiB
2027	7.5 GiB	60 GiB



dCache Usage and Predictions (in TB)




MicroBooNE Dedicated pools:

- Analysis pool 137 TiB (150 TB).
 - Used 96%.
- Production pool 1362 TiB (1500 TB).
 - New in 2022.
 - Used 41%.
 - MicroBooNE will be expanding use in 2023 (see slide 8).
- Bern OSG site 600 TB.
 - Used 73%.

	Analysis (Persistent)	Production (Persistent)
Current	137 TiB	1362 TiB
2022	137 TiB	1362 TiB
2023	137 TiB	1362 TiB
2024	137 TiB	1362 TiB
2025	137 TiB	1362 TiB
2026	137 TiB	1362 TiB
2027	137 TiB	1362 TiB

Tape usage and predictions (in TB)



	Total Added By End of Year
At end 2021	31.8 PB (29.4 PiB) Total
2022	25.7 PB Total +1.4 PB added (via sam) -6.4 PB deleted
2023	+2.5 PB requested
2024	+1.5 PB requested
2025	+1.5 PB requested
2026	
2027	



Data Lifetimes

- In 2022, MicroBooNE deleted 6.4 PB of MCC8 and earlier data and MC.
 - As promised in last year's FCRSG.
- Remaining data on tape is mainly raw data and MCC9 processed data and MC.
 - No further large deletions are planned.

Analysis Facility Use

- At this time, MicroBooNE is not planning on being a big user of EAF.
 - Could possibly change if EAF becomes popular among SBN experiments.

Token Authentication

- The transition from `jobsub_client` to `jobsub_lite` hasn't been totally smooth.
 - We were testing `jobsub_lite` and targeting Feb. 15 (today) as our main deadline.
 - We were caught off guard by updates of `sam_web_client` on Feb. 1 and POMS on Feb. 8.
- Token authentication future updates and unfinished business.
 - We are aware that SCD is working on a method to capture and store `jobsub_lite` log files from condor.
 - This will be necessary to integrate `jobsub_lite` into `fifemon` batch job monitoring.
 - `Fifemon` monitoring is important to MicroBooNE (both for production and ordinary analyzers).
 - We request that batch job submission using `jobsub_client` not be turned off until `fifemon` integration is available.
 - We believe additional updates of `sam_web_client` will be necessary to enable token authentication for `xrootd` file delivery.
 - MicroBooNE will also have to update our MCC9 `xrootd` client.
 - We haven't been given much information about this so far.