



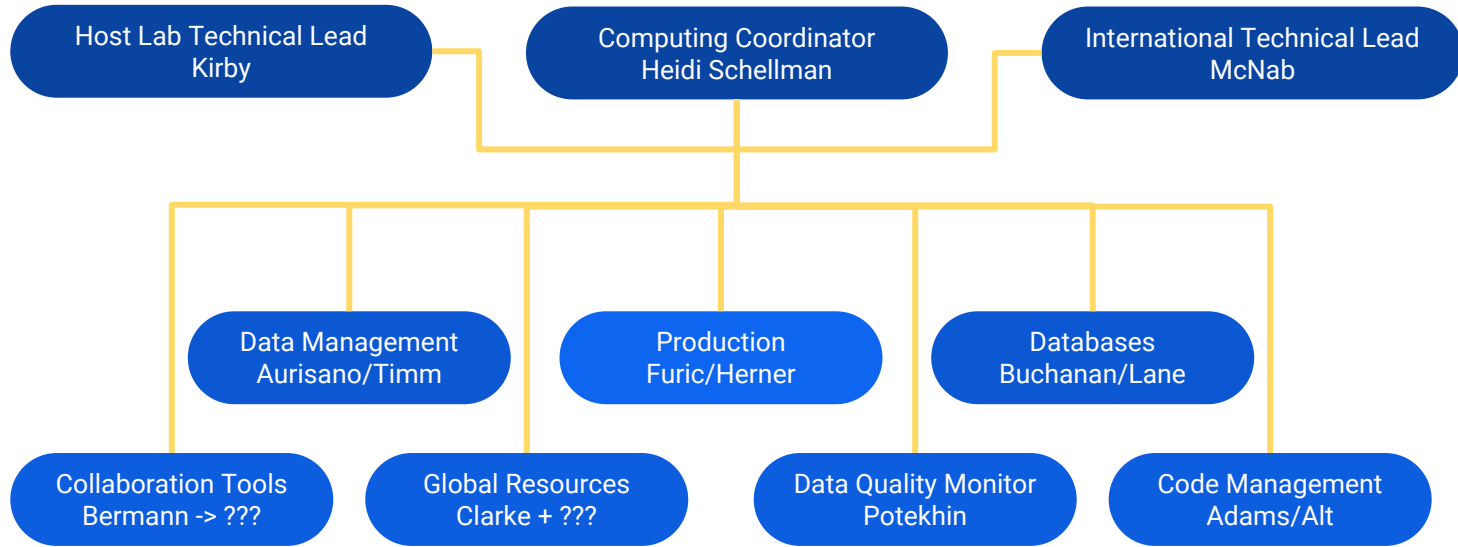
## DUNE SCPMT 2019

Work in progress - delete these comments when complete and ready for presentation

Needed:

- Experiment updates
  - SCD: all 2018 info has been updated (this comment can be deleted for presentation)
-

# Experiment Organization Chart for Offline Computing



# Physics drivers 2019-2021

Support for the Far Detector TDR's in 2019

Support for the Near Detector TDR's in 2019-2020

Computing CDR in late 2019-early 2020

Commissioning and cosmic running for Protodune DP (Dual Phase)


Cosmic running and studies for Protodune SP (Single Phase)

Reconstruction with improved algorithms for DP and SP

Test beam runs for DP and SP in 2021 if approved by the SPC



# Important Dates to Remember




▲ LOREM

# DUNE CPU - Experiment Usage Over the Last Year

FermiGrid total capacity ~ 160M hours / year, Average used ~ 10M hours / month

Total wall hours = 14.0 M FermiGrid: 10.9 M OSG: 2.9 M NERSC: 0.1 M	Total CPU hours = 10.2 M FermiGrid: 8.2 M OSG: 1.9 M NERSC: 0.1 M	Avg CPU Efficiency on FermiGrid 75%
--	--	--



Core weighted slot hours

# > 50% of 800,000 jobs in Oct-Dec were run offsite

Visualize / New Visualization (unsaved)

Save Share Refresh < ⌚ October 1st 2018, 00:00:00.000 to December 31st 2018, 23:59:59.999 >

Owner:dunepro AND Cmd:\*fife\_wrap\*

Uses lucene query syntax 🔍

Add a filter +

fifebatch-history-\*

Data Options

metrics


▶ Slice Size

Count


buckets

▶ Split Sli... MachineAttrGLIDEIN\_Site0:  
Descending



Add sub-buckets



# Memory Footprint Over the Last Year



# CPU and Memory Efficiency Over the Last Year





# FermiGrid CPU

## Prediction Going Forward and Accuracy of Prior Predictions

Units of Million (1 CPU, 2GB) wall hours per CY


		Used in 2017	Used in 2018	Used in 2019	Used in 2020	Used in 2021
<b>Requested</b>	2 years prior 1 year prior Same year	10.90 11.40 7.00	24.40 29.00 9.25	29.00 28.62 20.00	12.45 20.00	30.00
<b>Actual Used</b>	(SCD provides)	8.18	10.9		Fermigrid #'s are dominated by analysis. Production/simulation assumed to run mainly offsite.	
<b>CPU Usage Efficiency</b>	(SCD provides)	70%	75%	80%		

# CPU Adaptations Going Forward

*How can DUNE use OSG/HPC/HEPCloud going forward?*

- Already using > 50% offsite computing
- OSG/HEPCloud should be integrated reasonably easily as we use common tools.
- HPC still requires significant R+D - probably most useful for simulation, ML training and hypothesis testing as:
  - Expt. data processing requires IO, fast turnaround, db access and frequent code updates
  - Analysis is IO intensive and involves 100's of unique users.
- NOvA experience indicates training ML algorithms and final hypothesis testing is well suited to HPC's
- Parallelization of geant would be a big step forward

# Disk: dCache Usage and Predictions (in TB)




	min	max	avg	current
Fermilab Public dCache:StorageGroup:dune_protoDUNEPools	0 GiB	1.058 PiB	210 TiB	53 TiB
Fermilab Public dCache:StorageGroup:dune_readWritePools	0 GiB	713 TiB	215 TiB	390 TiB
Fermilab Public dCache:StorageGroup:dune_PublicScratchPools	0 GiB	575 TiB	168 TiB	168 TiB
Fermilab Public dCache:StorageGroup:dune_DuneAnalysisPools	0 GiB	180 TiB	160 TiB	163 TiB
Fermilab Public dCache:StorageGroup:dune_Geant4ReadWritePools	0 GiB	14 TiB	4 TiB	14 TiB
Fermilab Public dCache:StorageGroup:dune_SlowReadWritePools	0 GiB	14 TiB	4 TiB	14 TiB
Fermilab Public dCache:StorageGroup:lbne_DuneAnalysisPools	0 GiB	6 TiB	6 TiB	6 TiB
Fermilab Public dCache:StorageGroup:lbne_readWritePools	0 GiB	2 TiB	570 GiB	0 GiB
Fermilab Public dCache:StorageGroup:dune_SlowPublicScratchPools	0 GiB	2 TiB	712 GiB	1 TiB
Fermilab Public dCache:StorageGroup:dune_readPools	0 GiB	995 GiB	12 GiB	0 GiB
Fermilab Public dCache:StorageGroup:lbne_SlowReadWritePools	0 GiB	25 GiB	1 GiB	0 GiB
Fermilab Public dCache:StorageGroup:dune_SimonsPools	0 GiB	8 GiB	4 GiB	0 GiB

Request (Avg/Peak)	Current Total	2019	2020	2021
R/W	5,587 Available	6000	6000	10000
Scratch	2,084 Available	(add here)	(add here)	(add here)
Analysis (Persistent)	200 Dedicated	400	400	800
Other	1,136 Dedicated	(add here)	(add here)	(add here)

protoDUNE raw data

# Disk: NAS Usage and Predictions (in TB Units)



	min	max	
/lbne/data	15.77 TiB	20.04 TiB	17.69
/dune/data	15.77 TiB	20.04 TiB	17.69
/lbne/data2	8.96 TiB	13.19 TiB	9.91
/dune/data2	8.96 TiB	13.19 TiB	9.91
/lbne/app	3.18 TiB	6.37 TiB	4.52
/dune/app	3.18 TiB	6.37 TiB	4.52
/grid/fermiapp (dune)	72 GiB	119 GiB	101
/grid/fermiapp (lbne)	78 GiB	78 GiB	78
/grid/fermiapp (marslbne)	18 GiB	28 GiB	21
/grid/data (dune)	622 MiB	622 MiB	622


Ignore /lbne entries  
(alias mount point)

	Current (TB)	2019	2020	2021
<b>App</b>	8.0	10	15	20
<b>Data</b>	60.0	50	50	50

We assume we migrate off fermiapp and that # of users doubles in 2019 as analysis of protoDUNE starts.

# Tape Prediction Going Forward and Accuracy of Prior Predictions

## Units of TB



		Used in 2017	Used in 2018	Used in 2019	Used in 2020	Used in 2021
<b>Increment Requested</b>	2 years prior	100	-	7,500	393	20,000
	1 year prior	500	15,000	4,913	4000	
	Same year	1,000	2,394	8000		
<b>Increment Used</b>	(SCD provides)	307	4,935	Large use in 2021 is long test beam runs for both SP and DP		
<b>Integral Used</b>			5,442			

# Data Lifetimes

## Dataset preservation strategy

- All raw data needs to be retained
- We plan to keep two recent reconstruction passes on tape and dispose of older versions (or not migrate to new media).
- Reconstruction passes are written to unique file families to facilitate later removal. Removals is based on code version, not creation date.

# What Do You Want to Achieve in Computing Over Next Three Years

<b>Goals</b>	<b>Where does the experiment need to contribute</b>	<b>Where does SCD need to contribute</b>
Larsoft ability to run on extremely large events in finite memory	Define problem - adapt event format	SCD mods for Art/larsoft to allow processing of event fragments
Global access to DUNE computing resources	Provide external storage/CPU resources	Work with WLCG and others on uniform authentication/access methods
Geant parallelization	Part of intl. efforts	This goes beyond SCD to worldwide
Global batch systems	Evaluate/mod solutions	Evaluate/mod solutions
Improve Storage Latencies	Provide offsite storage	SCD support for storage - 1 example = better user info on status of files

# Anything else?

We anticipate a doubling of users in 2019 as protoDUNE data becomes available for analysis. **This likely means a need for doubling of the # of interactive machines and the /app/ space.**

While we do have other sites, we expect the overall usage to scale up and the large FNAL-based community to grow.