#### 



#### NOvA FCRSG 2022

#### NOvA @ Neutrino 2022





Presented this morning

https://neutrino2022.org/program/detail program?day=3

# Experiment Organization Chart for Offline Computing



Computing Coordinator serves as CS Liaison (GD)

Production group consists of several "shifters" that handle submissions weekly GENIE/GEANT/CRY development coordinated by Detector Systematics group

#### NOvA

ML development coordinated within Reco & DL sub-group

#### **Important Dates to Remember**

end of '23 Q1 goal dependent on prestaging speed



Next major production campaign ("prod6") result of reco R&D; late '23 or early '24 (no good estimates) NOvA

# CPU - Experiment Usage Over the Last Year



Security changes severely hampered offsite submission ability; ~ 1 year to recover.

NERSC: ~ 90 Million hours; not shown on FIFEmon

@ALCF used ~ 160 GPU node hours (tried using OSG, not very successful: availability/quality too low) NOvA

#### Memory Footprint Over the Last Year



Analysis-dominated. For analyzers, not fruitful to optimize

## CPU and Memory Efficiency Over the Last Year



CPU efficiency is good. Production scripting and support for users is strong

Analysis-dominated. For analyzers, not fruitful to optimize

NOvA



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

	2018	2019	2020	2021	2022	2023
Requested	23.17	31.60	29	33	35	37
Actual Used	28.80	25.94	40.5	41.4	12.5 through April	N/A
Efficiency	58%	77%	67%	74%	80% through April	N/A

## **CPU** Adaptations Going Forward

Freight train scheme in place for filtering raw data; access once and never again.

Unique data transfer challenge between FNAL and ALCF

HEPCloud jobs at NERSC interest for HPC GPUs notably with Perlmutter coming online and successes already shown, including at ALCF.

But SCD support needed

With large-scale GPU availability, opens up additional avenues for R&D reconstruction work planned.

## Disk: dCache Usage and Predictions (in TB)

PIB			min	max ~	avg	current			
		<ul> <li>Fermilab Public dCache:StorageGroup:nova_readWritePools</li> </ul>	22 118	986 TIB	600 TB	551 TIB			
nn		<ul> <li>Fermilab Public dCache:StorageGroup:nova_NovaAnalysisPools</li> </ul>	249 118	399 LIB	335 118	388 118			
P10		<ul> <li>Fermilab Public dCache:StorageGroup:nova_PublicScratchPools</li> </ul>	11 TiB	136 TiB	39 TiB	31 TiB			
A Later .		<ul> <li>Fermilab Public dCache:StorageGroup:nova_SlowNovaAnalysisPools</li> </ul>	20 TiB	49 TiB	33 TiB				Other
		<ul> <li>Fermilab Public dCache:StorageGroup:nova_NovaPrestagePools</li> </ul>	5 TiB	45 TiB	26 TiB	45 TiB		Analysis	
PB		<ul> <li>Fermilab Public dCache:StorageGroup:nova_NovaWritePools</li> </ul>	38 TiB	38 TiB	38 TiB	38 TiB			Dedicated
		<ul> <li>Fermilab Public dCache:StorageGroup:nova_CdfWritePools</li> </ul>	77 MiB	20 TiB				(Persistent)	
		<ul> <li>Fermilab Public dCache:StorageGroup:nova_Geant4ReadWritePools</li> </ul>				5 TIB		(* *********,	(Write)
PIB		<ul> <li>Fermilab Public dCache:StorageGroup:nova_SlowReadWritePools</li> </ul>				5 TiB			· · · /
	A	<ul> <li>Fermilab Public dCache:StorageGroup:nova_HsmWritePools</li> </ul>				4 TiB			
		<ul> <li>Fermilab Public dCache:StorageGroup:nova_HsmReadPools</li> </ul>				2 TiB			
		<ul> <li>Femilab Public dCache:Storagedrouprova_SlowPublicStoratchPublic</li> </ul>	65 GiB	337 GiB	70 GiB	337 GIB	Current	396 TB (actual)	1191 TB (actual)
							2022	+100 TB	

Total r/w (tape backed): ? TB Total scratch: 136 TB Total persistent: 396 TB Total other: ? TB

NOvA

We regularly must monitor persistent as we often approach the 396 TB allocation

An additional 100 TB would cover our next production run.

Using persistent to ensure analysis ntuples are always cached.

We carefully manage allocation with pinning prestaging datasets as needed, and then unpinning Adding small amounts each year doesn't make much difference. We currently are able to manage our current allocation with careful monitoring.

Additional space would alleviate some mgmt.

Cosmic filtering started at ALCF. Processing entire cosmic dataset in finite time, w/o huge waste of tape.

SCD supporting efficient pre-staging and transfer to ANL.

This will reduce tape usage in the future, but hard to predict at present.

	Total Added By End of Year
At end 2021	+3 PB (actual)
2022	+5 PB
2023	+4 PB
2024	+3 PB

# Tape - Usage and Predictions (in PB)



NOvA

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



**NOvA** 



Emergency-driven audits (human resource are thin) Desire for tools for archiving 12

#### Age of files in NAS



#### **Data Lifetimes**

We have never created a data set with a designated lifetime.

- We do not have plans to delete datasets; requires considerable care and effort to ensure that crucial, irreproducible data is not lost.
- There are datasets that we would not copy forward to new media (several Pb).
- We have explored what we could leave behind / delete.
- There is roughly 1.8 Pb of data from design and prototyping stages of experiment
- Ultimately, without provided tools, deleting datasets is not practical

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

NOvA-SCD Workshop held in October 2021: <u>https://indico.fnal.gov/event/51238/</u>

Transition from UPS to Spack

Scisoft packaging support, SRT build system support

SAM longevity

- MetaCat/Rucio transition not projected to have sufficient effort available for a safe transition
- Not all same functionality currently exists

Authentication

• Transition from certificates to tokens; support for FTS configurations needed

Linux Distros

• SL7 currently used for all of its online/offline systems. EOL May 24

#### Anything else?