



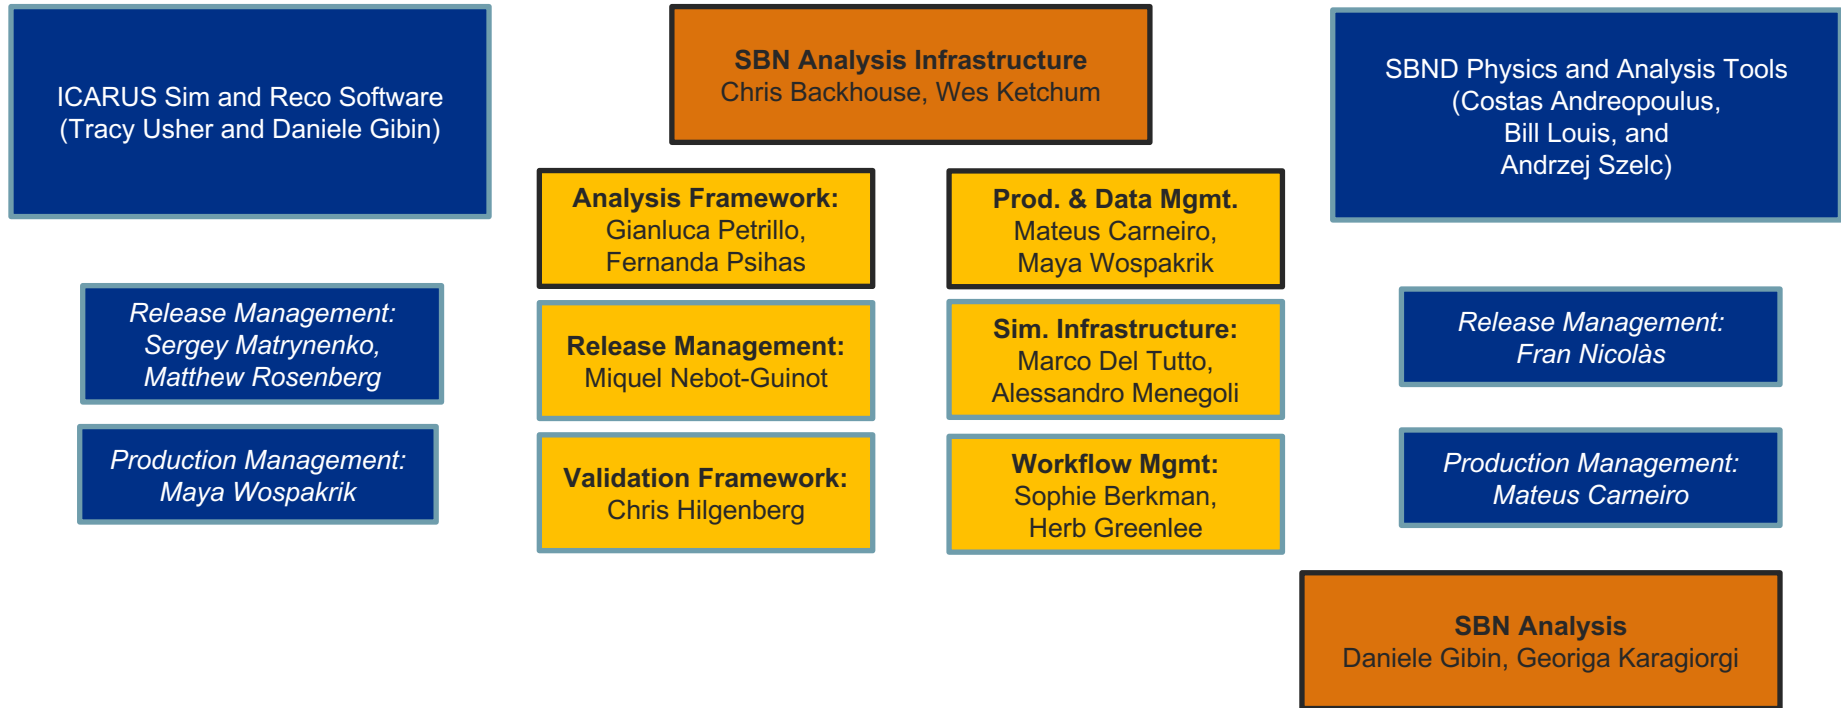
SBN FCRSG 2022: Resource Request

Chris Backhouse^a, Wesley Ketchum^b, and Joseph Zennamo^b

(a) UCL, (b) Fermilab

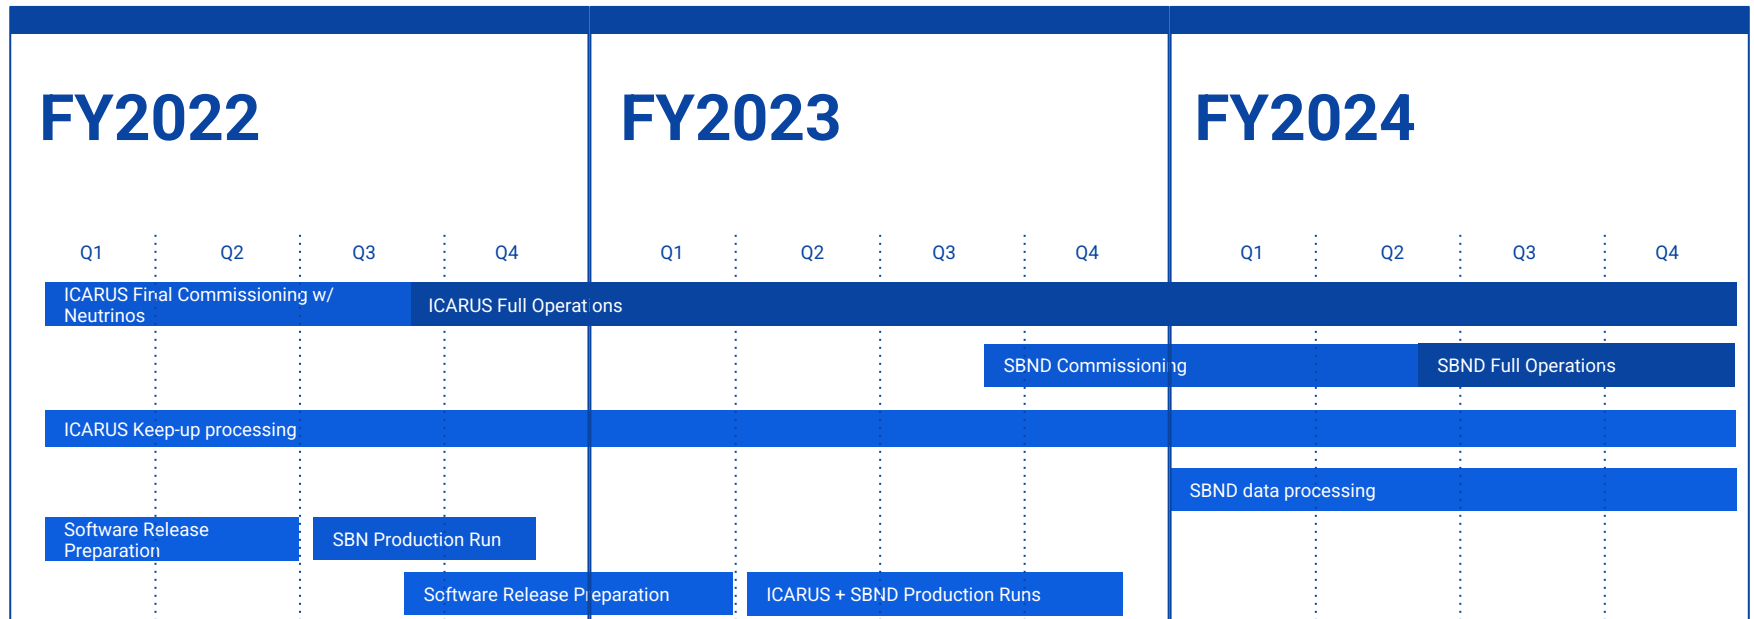
01 June 2022

Organization for Offline Computing



- Each experiment has an existing software and computing structure
- Common *SBN Analysis Infrastructure* group to organize common needs and efforts
 - SBN numbers are presented jointly, though with breakdown by individual detector available
 - Work hand-in-hand with each detector's groups and common SBN Analysis group

Overview Timeline




- ICARUS trigger commissioning complete; final cosmic-ray overburden being installed now → physics run
 - Have been taking neutrino data this whole year; various stages of detector commissioning
- SBND to enter commissioning in FY2023-2024 and move (hopefully rapidly) to full operations
- Major data production campaigns:
 - Ongoing first 'large-scale' campaign'
 - Expect two major production campaigns for FY2023, with a goal to begin one as early as October 2022


Updates in computing model

- No major overhaul of ‘default’ computing model since last year
 - Our computing continues to be ‘Fermilab-centric’, though we have added data backup storage at CNAF and pursuing developments to use at HPC
- Detector data production workflows rely heavily on inline (“keep-up”) processing of data
 - Reduces data volume, but there is a risk for re-processing
- Data production campaigns
 - Assume two per year, but out-years assume only one major simulation production
 - Require campaigns to be ~90 days in order to be useful
 - Assume 10x statistics for major simulation samples, and have factored in increased POT/year estimates based on recent beam performance
- Data lifetimes:
 - Raw (detector) data: permanent
 - Derived data: 1-2 years

ICARUS CPU usage over last year




CPU Efficiency (CPU time / Wall time) (Combined Production and Analysis)




ICARUS Memory usage over last year



Memory Usage (Combined Production and Analysis)





Memory Efficiency (Usage/Request) (Combined Production and Analysis)



SBND CPU usage over last year



SBND Memory usage over last year



CPU and Memory usage notes

- Memory usage and efficiency a continual concern
 - **Getting to less than 2 GB will require significant work**
 - ICARUS full cosmics simulation ~5 GB → can be offset in future by use of data-driven cosmic backgrounds, but at cost of I/O
 - SBND simulation needs 2-4 GB
 - Already have off-loaded optical simulation from look-up library to semi-analytic (CPU-heavy) model
 - Need to realistically assess capability of < 2 GB goal
- Across SBN, CPU efficiency often limited by data I/O
 - Need improved workflow management and resources to improve/work around this
 - Production efficiency ***significantly*** improved this year with pushing consecutive stages to run on same worker node and **making use of SBN data pool**
 - More later...

SBN CPU usage history and prediction


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

	CY2018	CY2019	CY2020	CY2021	CY2022	FY2023	FY2024
Requested	9	13.2	16.5	19.5	27	42.8	45.3
Actual Used	0.67	4.46	7.7	11.4	15	N/A	N/A
Efficiency					~80-90%		

CPU Adaptations Moving Forward

- We have used OSG resources, though often limited in availability due to high memory requirements
 - Will hope to use more as we improve workflows
- We are interested in onboarding to HEPCloud, but has not been highest priority for us
 - And, importantly, no mechanism for funding yet identified
- Significant work in SBN on use of HPC resources
 - SciDAC projects on optimizing signal-processing and hit-finding in reconstruction and final fits for HPC → demonstration tests with ICARUS data workflows ongoing
 - Previous UK efforts on containerized SBND cosmic simulation on ANL Theta → achieved additional processing capacity in recent production
- Heterogeneous computing use currently not fully-factored in, though use of GPUs for ML training/inferences is being discussed

ICARUS and SBND dCache usage




NOTE! SBN Data Pool Double-Counted!

SBN Tape usage

ICARUS


Tape Used



	min	max	avg ^	current
Fermilab Enstore:StorageGroup:icarus	5.909 PiB	10.272 PiB	7.881 PiB	10.272 PiB

SBND

Tape Used



	min	max	avg ^	current
Fermilab Enstore:StorageGroup:sbnd	826.49 TiB	872.82 TiB	853.03 TiB	872.82 TiB

SBN Tape usage and predictions

Units of PB

	Total Added / Total Footprint By End of Year
Current	11.2
Additional FY2022 end	+3.9 / 15.1
FY2023**	+0.6 / 15.7
FY2024	+15.6 / 31.25

**Assumes aggressive data retirement, see next slide

Some notes on data usage

- We continue to focus on driving down data volume needs in SBN
 - Previously saw large gains in SBND sim event sizes
 - ICARUS reconstruction workflows target dropping of raw data from reco files
- Need to formally assess current data on tape for retirement
 - Bulk is ICARUS commissioning data, which was always planned to have limited lifetime
 - FY2023 additional added raw data volume is 13.2 PB → goal would be to offset that increase by retiring most current ICARUS commissioning data
- We are using Rucio for data transfers to CNAF, but yet to fully integrate into data management system
 - Critically need SCD support for this

SBN dCache Requests

- Anticipate needs to increase in persistent space to support increasing analysis activities
- Requesting ~1 PB write volume to support SBND data-taking in FY2023
- Received 2 PB allocation last year for “SBN data pool”
 - Target was for output of production: derived data, commonly used for analysis work, with medium-term lifetimes
 - Has had a *major* improvement on production efficiency and reduction in tape usage

	Analysis (Persistent)	Other Dedicated (RW)
Current	209 TB (actual)	3000 TB (actual)
FY2022	250 TB	3000 TB
FY2023	300 TB	4000 TB + ?
FY2024	300 TB	4000 TB + ?


More notes on data usage

- A major computing problem for SBN continues to be data I/O
 - Estimate that in FY2023, we will need peak rates of 3.22 TB/hr read, and 3.95 TB/hr write (simultaneously) to support data-taking and production campaigns
 - For reference, our steady-state raw data write rate is ~ 0.7 TB/hr
 - By FY2025, this increases to 7.62 TB/hr and 7.34 TB/hr, respectively
 - This is *after* aggressive (but hopefully realistic) improvements on cutting down data usage
- **We consider SBN data pool to have been a major success and, as we discussed when we requested the resource, we will want to pursue how to expand**
 - Still assessing how to optimize use patterns (using ongoing production as a test point), but would like to minimize using tape as a ~ 2 -yr-lifetime data cache
 - Total derived data is ~ 8 PB in FY2023 \rightarrow upper-limit scale for what would be used
 - Will plan to prepare presentation for SPPM to outline proposal
 - Aiming for later this summer

SBN NAS usage

ICARUS


BlueArc Disk Usage



	min	max	avg	current
icarus-data	5.080 TiB	8.044 TiB	6.637 TiB	8.044 TiB
icarus-app	2.961 TiB	3.990 TiB	3.540 TiB	3.990 TiB

SBND

BlueArc Disk Usage



	min	max	avg	current
lar1nd-data	22.45 TiB	25.00 TiB	23.84 TiB	24.58 TiB
lar1nd-app	4.22 TiB	4.88 TiB	4.62 TiB	4.75 TiB

SBN NAS usage and predictions

ICARUS

	App	Data
current	4.0	25
FY2023	5.0	30
FY2024	5.0	30

Units
of TB

SBND

	App	Data
current	4.5	25
FY2023	5.0	30
FY2024	5.0	30

- Needed to support areas for software development, tests, and analysis work
- Modest increases support needs for onboarding new analyzers
 - Effort to help keep builds small, tools for cleaning up larsoft builds, etc. are often biggest drivers of efficient use (besides general cleanup)
- [Astute observers will note we have not actually increased these over last years, and we try to continue to cleanup rather than ask for more ...]

General concerns

- LArSoft / framework support
 - We rely heavily (and contribute back to) LArSoft and the LArSoft ecosystem (e.g. *mrbs*, *art*, etc.)
 - Improvements in CPU efficiency, memory usage, production and validation systems, HPC deployment, etc. rely on SCD support even when those activities have significant collaboration effort/leadership
 - We are worried about maintaining the level of support needed, especially when effort goes into new developments
 - Going from “realistic demonstrator” to “production-stable” infrastructure generally takes time, and general leads to dependence on “feature-frozen” software far longer than anticipated

Backup

ICARUS CPU usage prediction

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

	CY2018	CY2019	CY2020	CY2021	CY2022	FY2023	FY2024
Requested	0.5	4.8	12	15	20	17.9	16.7
Actual Used	0.5	1.6	3.5	6.5	10.5*	N/A	N/A
Efficiency					~80-90%		

*extrapolation to end of CY

SBND CPU usage prediction

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

	CY2018	CY2019	CY2020	CY2021	CY2022	FY2023	FY2024
Requested	8.5	8.4	4.5	4.5	7	24.9	28.6
Actual Used	0.2	2.9	2.2	4.9	4.5*	N/A	N/A
Efficiency					~90-95%		

*extrapolation to end of CY

ICARUS Tape usage and predictions

Units of PB

	Total Added / Total Footprint By End of Year
Current	10.3
Additional FY2022 end	+3.9 / 14.2
FY2023**	-0.2 / 14.0
FY2024	+10.3 / 24.3

****Assumes aggressive data retirement of commissioning data**

SBND Tape usage and predictions

Units of PB

	Total Added / Total Footprint By End of Year
Current	0.9
Additional FY2022 end	+0 / 0.9
FY2023	+0.8 / 1.7
FY2024	+5.3 / 6.9

Some 2021 Slides

General notes/concerns

- ICARUS and SBND continue to make use of common tools for computing, and want to continue to do so but will need ongoing support
 - Software builds and releases: LArSoft, spack transition, continuous integration/validation tools
 - Strategic plan to use Rucio for data management / data transfers at risk due to lack of support
 - Missed first critical deadlines, working to avoid missing others
- Working towards improving our computing model ensuring it continues to meet both joint analysis and individual analysis needs
 - Putting a focus on optimization of production workflows, in particular use of HPC in computing production given significant recent progress

Goals over the next three years

Goals	Where does the experiment need to contribute	Where does SCD need to contribute
Transition of ICARUS from commissioning to operations now, and in future for SBND	Detector commissioning, computing model, sensible physics plan	Support for optimization of onsite data transfers and data storage, processing of data
Deployment of Rucio for improved data management, including automated data transfers (CNAF, other sites)	Collaboration support/effort, expertise on external sites and interfaces to external computing	Core development work, documentation, support
Increased usage of HPC resources, integrating into standard production workflows	Collaboration support/effort, algorithm development, workflow development	Expertise for development, support and expertise for transitioning to production
Large-scale production for single and multi-detector physics analyses	Software, data management, workflow development	(LArSoft) software support, profiling and optimization expertise, including on data storage and management

ICARUS CPU usage prediction

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

	2018	2019	2020	2021	2022	2023
Requested				14.3	20.8	29.1
Actual Used				N/A	N/A	N/A
Efficiency	%	%	%	N/A	N/A	N/A

SBND CPU usage prediction

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

	2018	2019	2020	2021	2022	2023
Requested				4.0	11.8	19.0
Actual Used				N/A	N/A	N/A
Efficiency	%	%	%	N/A	N/A	N/A

ICARUS Tape usage and predictions

Units of PB

	Total Added / Total Footprint By End of Year
At end 2020	2.8
2021	+10.9 / 13.7
2022	+4.0 / 17.7
2023	+6.5 / 24.2

SBND Tape usage and predictions

Units of PB

	Total Added / Total Footprint By End of Year
At end 2020	0.4
2021	+0.6 / 1.0
2022	+3.6 / 4.6
2023	+2.8 / 7.4