

ICARUS Analysis

Production and Resource Management

Tracy Usher (SLAC)
ICARUS Collaboration Meeting
October 14, 2024

Overview

- Continuation of a discussion started at the recent SBN Workshop
 - See Giuseppe's presentation [here](#) (much here borrowed from that!)
- General idea here is to briefly present some baseline information:
 - Overview of Release Management
 - Overview of FNAL resources
 - Overview of the Production system
- Then move to a list of potential discussion topics
- Bottom line: we need to have a discussion on how to solve current problems/issues in order to successfully achieve future analyses

Release Management

Brief Reminder of ICARUS Software

- Detailed presentation at last collaboration meeting [here](#)
- Hierarchical structure to code:
 - Dependencies →LArSoft →SBN code →ICARUS code
- ICARUS maintains four code repositories in github
 - icarus_signal_processing, icarusalg, icarusutil, icaruscode
- Generally manage the code along branches
 - Code development meant to be solely along the main develop branch
 - Production running meant to be on production branches which are “maintained”
- The main develop branch is meant to be frequently updated to keep pace with LArSoft “keep up” releases

Brief Reminder of Github tools

- Not only does Github provide a nice code repository, it also...
 - Provides Wiki pages for documenting our code ([current icaruscode wiki](#))
 - Provides “Pull Requests” for integrating code updates into branches
 - With extensive ability to document, test, etc.
 - Provides “Issues” which are useful for
 - Reporting and tracking bugs in the code
 - Future planning (see Projects)
 - Provides “Projects” for overall management
 - Can track issues, pull requests, set timelines, etc.
 - Powerful tool which is underutilized by us
- The bottom line:
 - Github provides a rich set of tools for managing our code and releases
 - We do not need to supplement with external tools, rather we need to fully exploit the provided tools!

SBNSoftware / icaruscode

Code Issues 37 Pull requests 11 Discussions Actions Projects Wiki Security Insights Settings

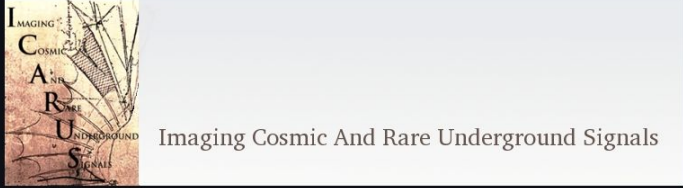
icaruscode Public Edit Pins Unwatch 18 Fork 31 Star 6

develop 233 Branches 251 Tags Go to file Add file Code

mmrosenberg updated tags/dependencies for weekly release 21cbe35 · last week 5,105 Commits

doc	Revert "Revert "ReleaseV09 44 00""	2 years ago
fcl	Merge pull request #726 from yeonjaej/v09_87_00_simple...	last week
icaruscode	Merge pull request #727 from PetrilloAtWork/feature/gp_P...	last week
scripts	Removed obsolete overburden helper scripts	2 weeks ago
test	Move to using the run 2 stage 1 fhid file	last year
ups	updated tags/dependencies for weekly release	last week
.gitignore	Added gitignore rules to prevent unwated commits	4 years ago
CMakeLists.txt	updated tags/dependencies for weekly release	last week
README.md	Update README.md	2 years ago

README



ICARUS simulation and reconstruction software

The documentation of ICARUS is spread mainly between:

- ICARUS wiki in GitHub at https://sbnssoftware.github.io/icaruscode_wiki/Wiki/; there is also an [older Fermilab](#)

About

Main/top level repository for ICARUS specific code

- Readme
- Activity
- Custom properties
- 6 stars
- 18 watching
- 31 forks

Report repository

Releases 26

Production Release Latest
on Mar 29, 2023

+ 25 releases

Packages

No packages published
[Publish your first package](#)

Contributors 51

+ 37 contributors

Languages

C++	92.7%	Python	1.8%
CMake	1.8%	Jsonnet	1.6%
C	1.2%	Shell	0.5%
Other		0.4%	

SBNSoftware / Projects / ICARUS 2024A Production Release

ICARUS 2024A Production Release

Repository View Board View Timeline + New view

Filter by keyword or by field

Title	Assignees	Status	Linked pull requests	Reviewers	Repository
SBNSoftware/icarus_signal_processing 1					
1 Updating coherent noise subtraction #22	SFBayLaser	Merged			SBNSoftware/icarus_signal_processing
+ repo:icarus_signal_processing Search issues and pull requests, create a new issue, or add multiple items					
SBNSoftware/icarusalg 13					
2 Fixed geometry unit test configurations #82	PetrilloAtWork	Merged		SFBayLaser	SBNSoftware/icarusalg
3 Modify DriftVelFudgeFactor to reflect measurements #81		Merged		brucehoward-physi...	SBNSoftware/icarusalg
4 shift southwest modules to avoid overlap with southeast #80		Merged		aheggest, biswaran...	SBNSoftware/icarusalg
5 Deactivated TrackTimeInterval test #79	PetrilloAtWork	Merged		mrosenberg and S...	SBNSoftware/icarusalg
6 Optical detector simulation and reconstruction parameter tests [1/2] #78	PetrilloAtWork	Merged		ameneqol, cfarnese...	SBNSoftware/icarusalg
7 Feature/cerati icarus newg4 dev #74	cerati	Merged		ameneqol, jzenna...	SBNSoftware/icarusalg
8 Various updates to icarusalg algorithms #73	PetrilloAtWork	Merged		SFBayLaser	SBNSoftware/icarusalg
9 Fixed typo in documentation #72	PetrilloAtWork	Merged			SBNSoftware/icarusalg
10 TimeIntervalConfig: utility to specify time intervals in FHiCL configuration #71	PetrilloAtWork	Merged			SBNSoftware/icarusalg
11 Renamed GSL header (gsl_util → util) #70	PetrilloAtWork	Merged		SFBayLaser	SBNSoftware/icarusalg
12 Utility to cross art associations #69	PetrilloAtWork	Merged		brucehoward-physics	SBNSoftware/icarusalg
13 Update to gallery code [1/2] #65	PetrilloAtWork	Merged		SFBayLaser	SBNSoftware/icarusalg
14 Turn off InterPlane offset #64	gputnam	Merged		brucehoward-physics...	SBNSoftware/icarusalg
+ repo:icarusalg Search issues and pull requests, create a new issue, or add multiple items					
SBNSoftware/icaruscode 80					
15 Stage 1 slimming #677	jzenna	Work in progress			SBNSoftware/icaruscode
16 Light Fix #678	PetrilloAtWork	Todo			SBNSoftware/icaruscode
17 Stage1 BDT Retraining #679	acampani and bru...	Todo			SBNSoftware/icaruscode
18 Stage0 calibration #682	mmoonney	Todo			SBNSoftware/icaruscode
19 Updated GENIE systematics #683	sjgardiner	Todo			SBNSoftware/icaruscode
20 Move to Refactored LArG4 #684	bear-is-asleep, cer...	Completed			SBNSoftware/icaruscode

General Comments on Release Management

- Matthew Rosenberg has done excellent job as release manager!
 - But he is leaving us and we need to a replacement ASAP!
- Github provides excellent tools for managing our releases
 - Code repository, Pull Requests, Branching, Issue Tracking, Project Management, Wiki, etc.
 - We need to make better use of them!
- In the past we have not been as disciplined in production releases as we need to be
 - We did not have an official branch for production this year
 - We have had issues with using different code releases to run simulations
 - A recent example used different detector simulation fcl files which led to significant differences when comparing the two samples
- We need to correct this going forward (see discussion)

Resources

Current FNAL Resources

- Disk Space
 - 2 PB in data pool 1, 1 PB in data pool 2 - shared with SBND
- Tape Allocation
 - Currently ~36 PB of which we have used ~28 PB (with some caveats)
- FermiGrid Allocation
 - 2000 “slots”
 - 1 CPU with 2GB memory
 - Slots allocated by memory needed, e.g. 8 GB memory request requires 4 slots
 - 50% of allocation given priority to the “icaruspro” account
 - “Keep up” data processing, production jobs
 - Can exceed quota if FermiGrid is otherwise idle
 - Can run “opportunistically” on the Open Science Grid


Monitoring

There is a nice summary grafana page [here](#)

CPU Wall clock time

- “Production” is from production team
- “Analysis” is from users


Monthly Wall Time



Disk Space Usage

- Orange is pool 1 (2 PB quota)
- Blue is pool 1 (1 PB quota)


Disk Space Used by Pool Group



Tape Usage

- As of this month just over 28 PB!

Tape Used



Other Resources

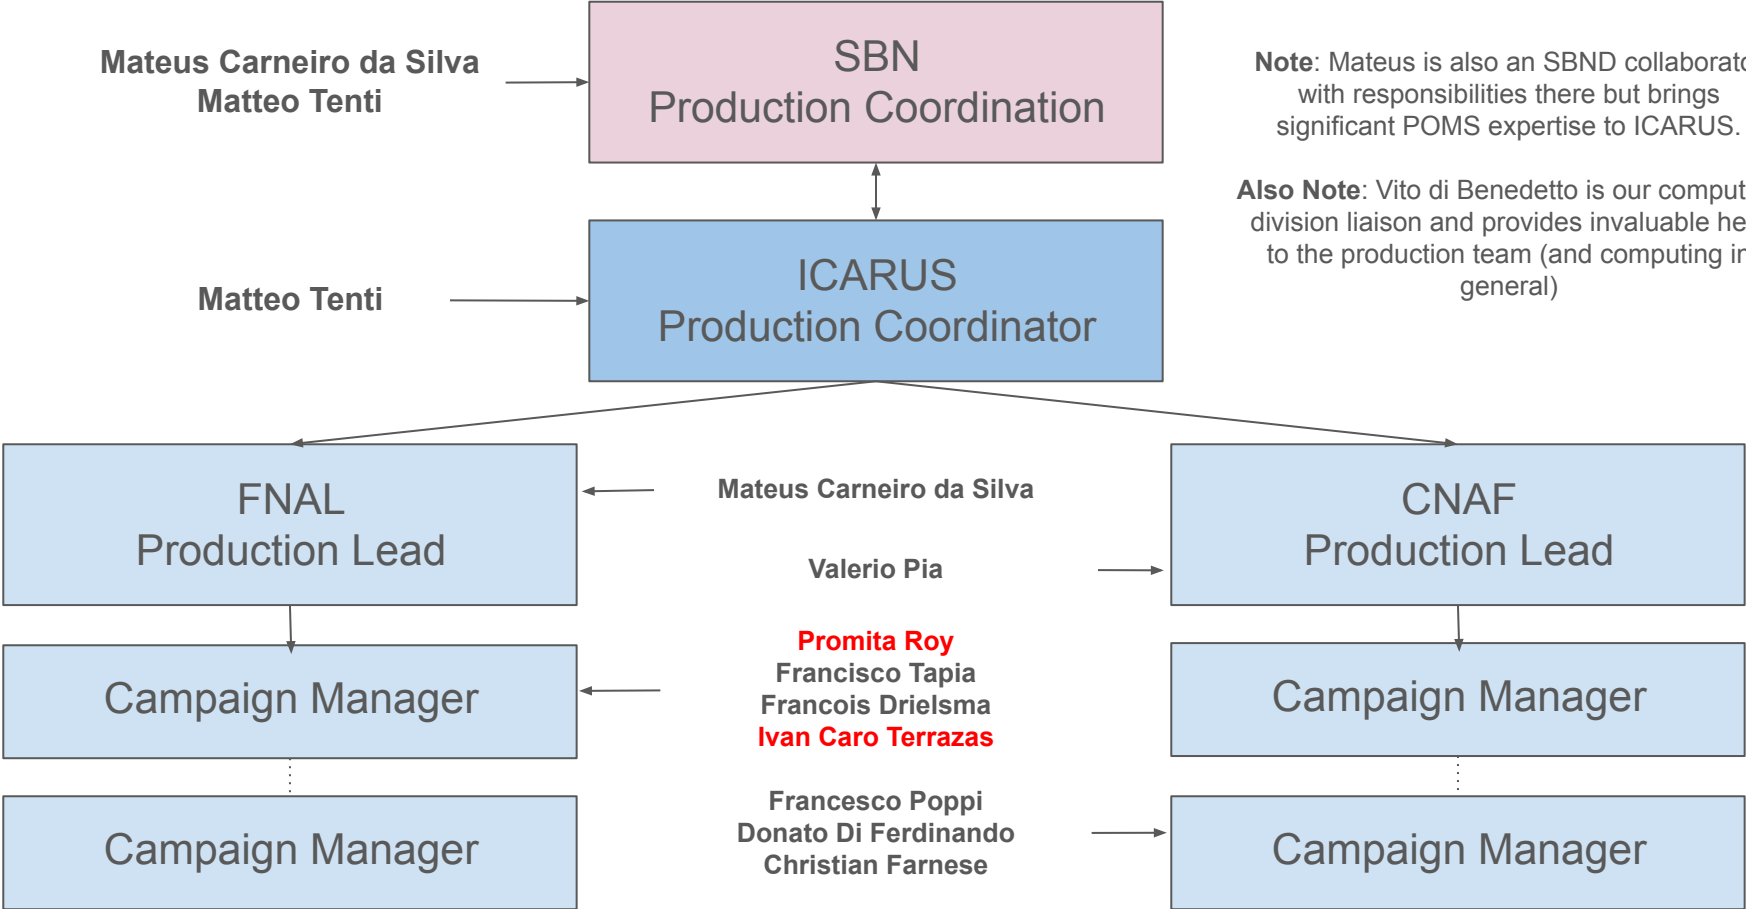
- CNAF
 - Current disk space allocation is 1.6 PB
 - Using just under 1.3 PB now (including Run 1 data on disk)
 - Tape Allocation is 4 PB to keep a second copy of ICARUS beam data
 - Using just under 2.5 PB now
 - Computing Allocation is 3.5M CPU hours
 - Generally, CNAF can be utilized for smaller scale and/or specialized productions
- SLAC
 - Using GPUs there to process SPINE reconstruction
- NERSC
 - Have standing allocation for CPU time but not disk space
 - Need demonstration of ability to run icaruscode at this site?
- Polaris
 - U Chicago has an allocation for processing including 3 PB of disk nominally for SBND
 - But could be considered SBN resource?

General Comments on Resources

- Data on disk at FNAL
 - Data pool 1 (2 PB) has ~150 TB free, data pool 2 (1 PB) has ~300 TB free
 - Breakdown:
 - 0.8 PB for recompressed Run 2 raw data with the rest taken by
 - Stage0 (ICARUS)/Reco0(SBND) and Stage1/Reco1 artroot files needed by analyzers
 - CAFs/FlatCAFs, Calibration tuples and SBND commissioning data
 - **Shared with SBND yet we are dominating usage!**
 - Estimate the Run 3 raw data (already compressed) at ~0.5 PB
- Primary problem:
 - Currently ICARUS simulation Stage 1 artroot output ~300 MB/event
 - A 10M v+cosmics sample would fill entire disk allocation!
- How to address?
 - Can the Stage 1 files be “slimmed” enough to work?
 - Can we convert to CAFs as final output of production?

Production

Production Team Organization



Note: Mateus is also an SBND collaborator with responsibilities there but brings significant POMS expertise to ICARUS.

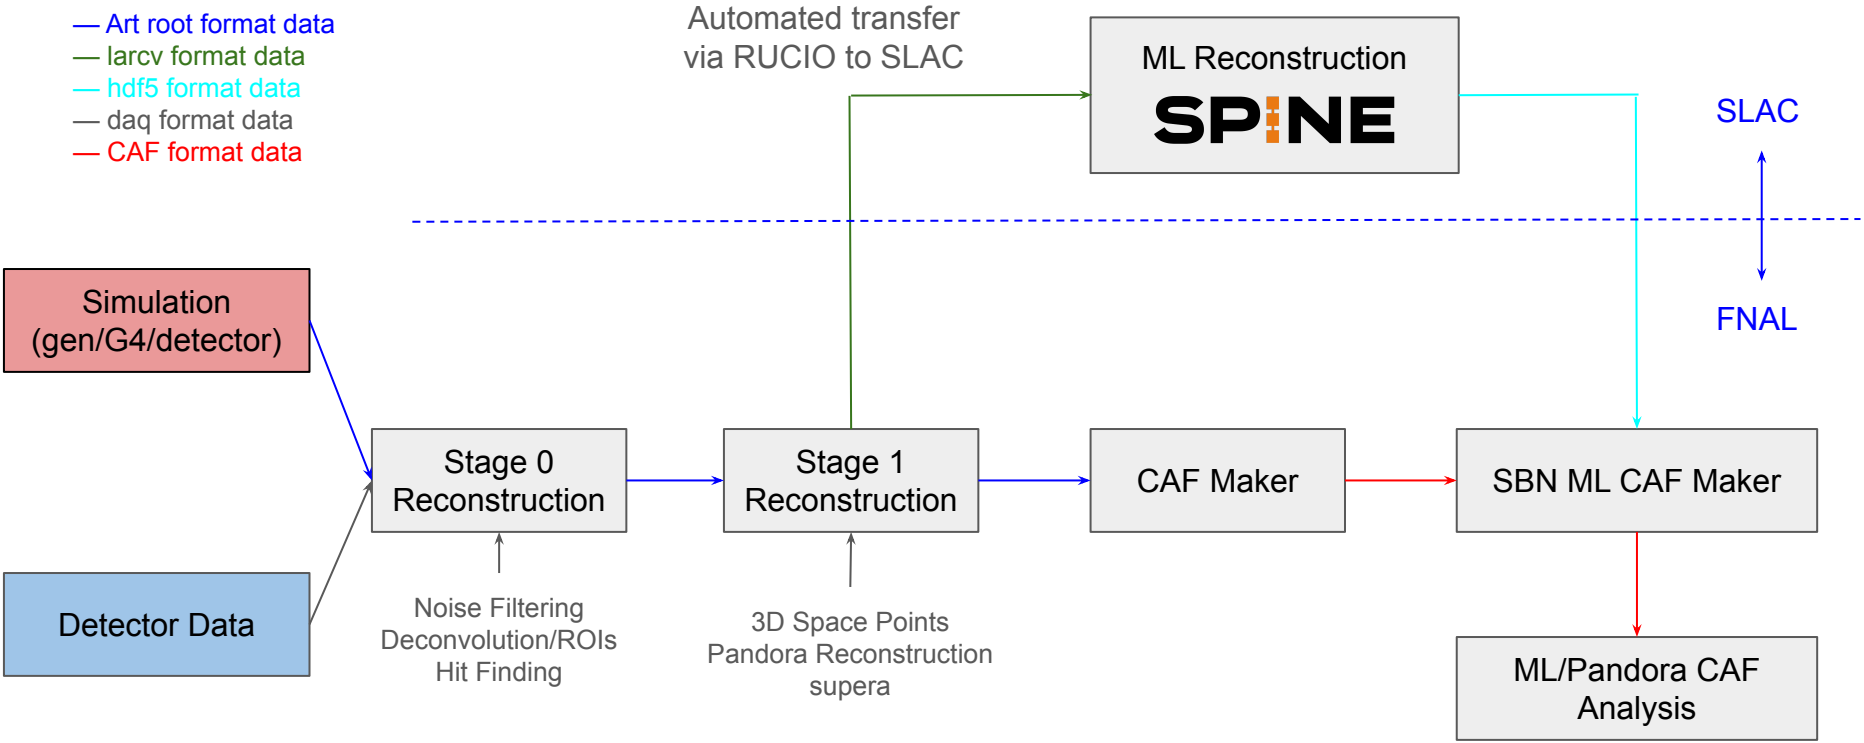
Also Note: Vito di Benedetto is our computing division liaison and provides invaluable help to the production team (and computing in general)

A Few Tools Used By Production Team

- POMS (Production Operations Management System)
 - A software package allowing “production teams and analysis groups across multiple Fermilab experiments to launch, modify and monitor large scale campaigns of related Monte Carlo or data processing jobs”
 - It can be fairly complex to operate... but is generally stable now
- RUCIO
 - Transfer of daq data to CNAF for archiving to tape
 - Transfer of larcv format data to SLAC for running ML processing on GPUs
- Github
 - Hosts wiki pages that collaboration members can access to submit and view requests
 - Hosts Project page for tracking various requests, issues, etc.

Basic Workflow for Production

- Art root format data
- larcv format data
- hdf5 format data
- daq format data
- CAF format data



Credit to Justin Mueller for diagram concept

SBNSoftware / Projects / Production Requests

Production Requests

Backlog Team capacity Current iteration Roadmap My items + New view

Assignees

- francois-drielsma 16
- icaromx 1
- PetrilloAtWork 1
- PromitaRoy 12
- No Assignees 2

Filter by keyword or by field

Title	Status	Size	Estimate	Iteration	Start date	End date
P0 4 (Estimate: 0)						
1 request #29 #4	Assigned	L				
2 request #54 #24	Completed	S				
3 request #51 #7	Completed	XS				
4 request #53 #23	Completed	S				
+ Add Item						
P1 24 (Estimate: 0)						
5 request #62 #32	Completed	M				
6 request #75 #45	Assigned	M				
7 request #56 #27	Assigned	M				
8 request #59 #30	Assigned	M				
9 request #58 #29	Completed	M				
10 request #61 #31	Assigned	L				
11 request #63 #33	Completed	M				
12 request #64 #34	Completed	M				
13 request #65 #35	Completed	M				
14 request #66 #36	Completed	M				
15 request #67 #37	Completed	M				
16 request #68 #38	Completed	M				
17 request #69 #39	Completed	M				
18 request #70 #40	Completed	M				
19 request #71 #41	Completed	M				
20 request #72 #42	Completed	M				
21 request #73 #43	Completed	M				
22 request #13 #2	Completed	S				
23 request #14 #3	Completed	S				

The production team have also set up a production project in the SBN github workspace.

This can be used to see the status of production releases...

For example you can see from the status column which requests have been assigned, completed, etc.

Clicking on a specific release gives more information about that production request:

SBNSoftware / Projects / Production Requests

Production Requests

Backlog Team capacity Current iteration Roadmap My items + New view

Assignees

- francois-drielsma 16
- icaromx 1
- PetrilloAtWork 1
- PromitaRoy 12
- No Assignees 2

Filter by keyword or by field

Title

P0 4 (Estimate: 0)

- request #29 #4
- request #54 #24
- request #51 #7
- request #53 #23

+ Add Item

P1 24 (Estimate: 0)

- request #62 #32
- request #75 #45
- request #56 #27
- request #59 #30
- request #58 #29
- request #61 #31
- request #63 #33
- request #64 #34
- request #65 #35
- request #66 #36
- request #67 #37
- request #68 #38
- request #69 #39
- request #70 #40
- request #71 #41
- request #72 #42
- request #73 #43
- request #13 #2
- request #14 #3

Open request #29 #4

SBNSoftware/icarus-production

mt82 on Jun 25

Check with @mattfc the status of the renaming for the other streams before starting the campaigns for those streams. In the meanwhile @PromitaRoy setups the...

PromitaRoy on Jun 26

Run 2 reprocessing for offbeamajo

mt82 on Jul 9

@PromitaRoy is waiting for the renaming of the files by @mattfc

mt82 moved this from Assigned to Pending in Production Requests on July 7

PromitaRoy on Aug 30

1. Was running test productions to fix the RUCIO thing last week, and I got handful of errors. Amongst all the errors, I could address all of them, (solved them eventually) with Vito's help
2. just two of the errors I could not address as those were due to the fact that RUCIO was not up and running on SLAC.
3. mentioned this to Francois, he raised a ticket and it was solved on last thursday and the files for my test production got transferred to SLAC successfully. Following that I started running the actual Run2 reprocessing production campaign, but this week again I had been getting the same errors since monday which is related to RUCIO not running on SLAC.
4. So, wrote to Daniel and he informed me that SLAC people are doing some debugging tests with RUCIO and that's the reason RUCIO is turned off.
5. This means RUCIO will get turned ON and OFF till the time they do tests with it.
6. So, I think I can only run RUN2 reprocessing campaign when SLAC people are done doing all tests and RUCIO is up and running properly.
7. Latest update: Daniel told me they will not shut down ours anymore. I will run a test campaign again to confirm it.
8. I am done running the staged for all the datastreams, the next step is to run stage1_caf_larcv which needs RUCIO to run properly on SLAC.

mt82 on Sep 12

Dan confirmed that we can keep the larcv files in Fermilab and then he will transfer them to SLAC. So, I will start doing this by

For example, selecting "request #29" shows the summary of the Run 2 data reprocessing.

Note this production processing has encountered a number of issues and you can read a recent summary from Promita in the ongoing comments.

This is obviously a useful way to track progress and keep track of issues!

General Comments on Production

- The Production Team are well organized and utilize an excellent set of tools to enable production requests, track progress, track issues, etc.
- There have been issues encountered this past year impacting production:
 - POMS has had setup, AL9 upgrade, support, etc., issues in last year
 - Hopefully crossing the line now
 - The RUCIO connection to SLAC was having stability problems (better now)
 - Disk space constraints impacted campaigns (and continue to be an ongoing problem)
 - Many problems require expertise in highly specialized areas to diagnose and address issues
 - etc.
- Given constraints production team has done good job to get a lot processed!
- Biggest issue currently is need for more people to get involved here
 - In particular, we will lose Promita (who has shouldered much of the load) at the end of year

Discussion?

An Unordered Incomplete List of Discussion Topics I

- General need to staff Release Management and Production
 - See Bob Wilson's presentation on Wednesday
- Release Management
 - Need to do better job of branching our code for analysis
 - Need to do better job of documenting our releases
- Production cycles
 - Currently large production cycles are not tied to regular timelines, tend to run “when ready”
 - Typically try to set a target date but let that slide as we address issues
 - This can result in smaller “ad-hoc” productions from different releases
 - Proposal from Giuseppe to move to a regular production cycle system
 - Regular schedule of release planning, CI validation, then data/MC production
 - If code integration issues then leave out to next cycle
 - Can track things with Github tools

An Unordered Incomplete List of Discussion Topics II

- Disk Space Management
 - Do we want to store Run 2 and Run 3 raw data on disk to facilitate reprocessing?
 - For how long?
 - Do we need to keep Stage 1 output on disk?
 - Yes: how can we slim to get to acceptable output size?
 - No: Are CAFs ready to be the final analysis product?
 - Available disk space is shared with SBND
 - Should we consider utilizing potential external resources?
 - Polaris, NERSC, SLAC, etc.
 - Potential data transfer issues? (e.g. RUCIO uptime at SLAC)
 - Data access issues if offsite?
 - Coordination with the SBN AI group critical
 - In particular we need common solutions with SBND