


Computing and Common Processing



Daniel Ruterbories
University of Rochester
17 October 2016

Charge Questions




- Question 1(D): The model for data processing and analysis including the budget and effort required
- Question 1(E): SCD and collaboration resources
- Question 4: Are there robust plans for data processing and data analysis? Have adequate resources from the collaboration been identified for data analysis to meet the set goals?

Outline



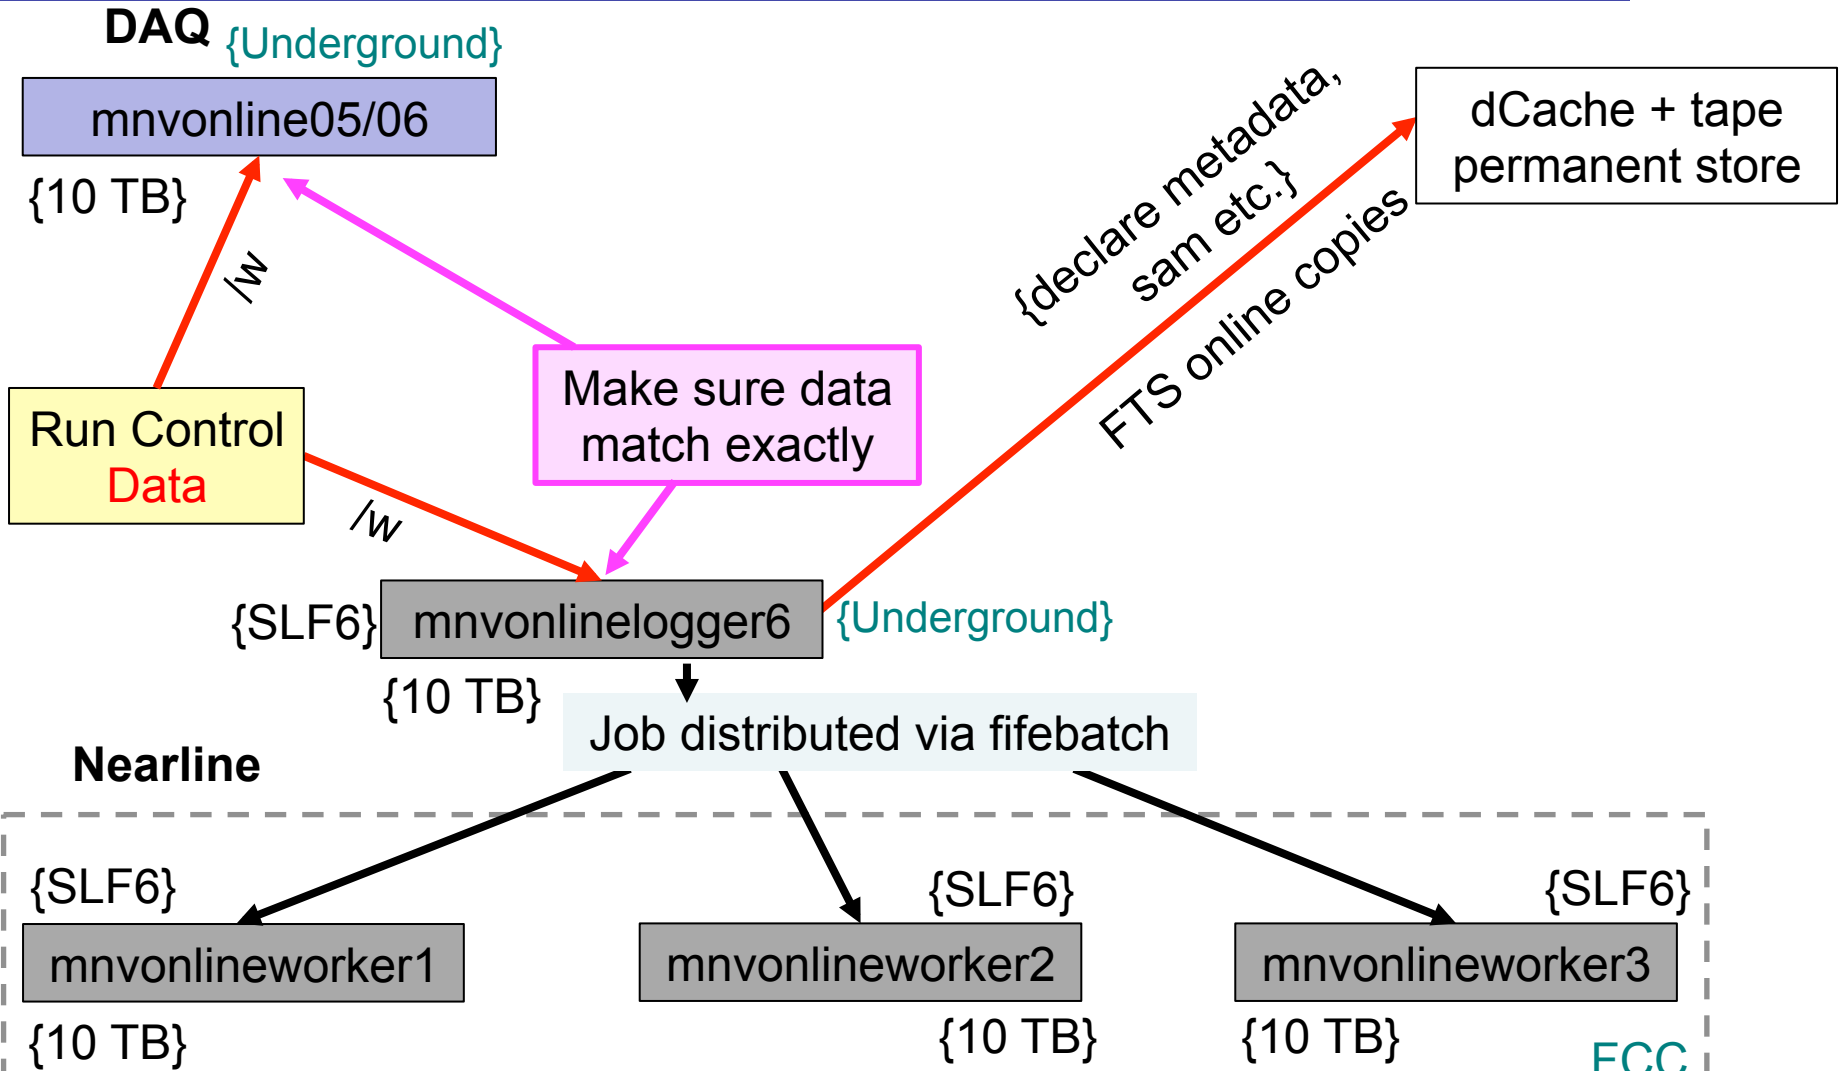

- Models for online/nearline analysis
- Models for offline data analysis
- SCD and collaboration resources
- What resources are necessary to execute the data analysis plan

Charge Question

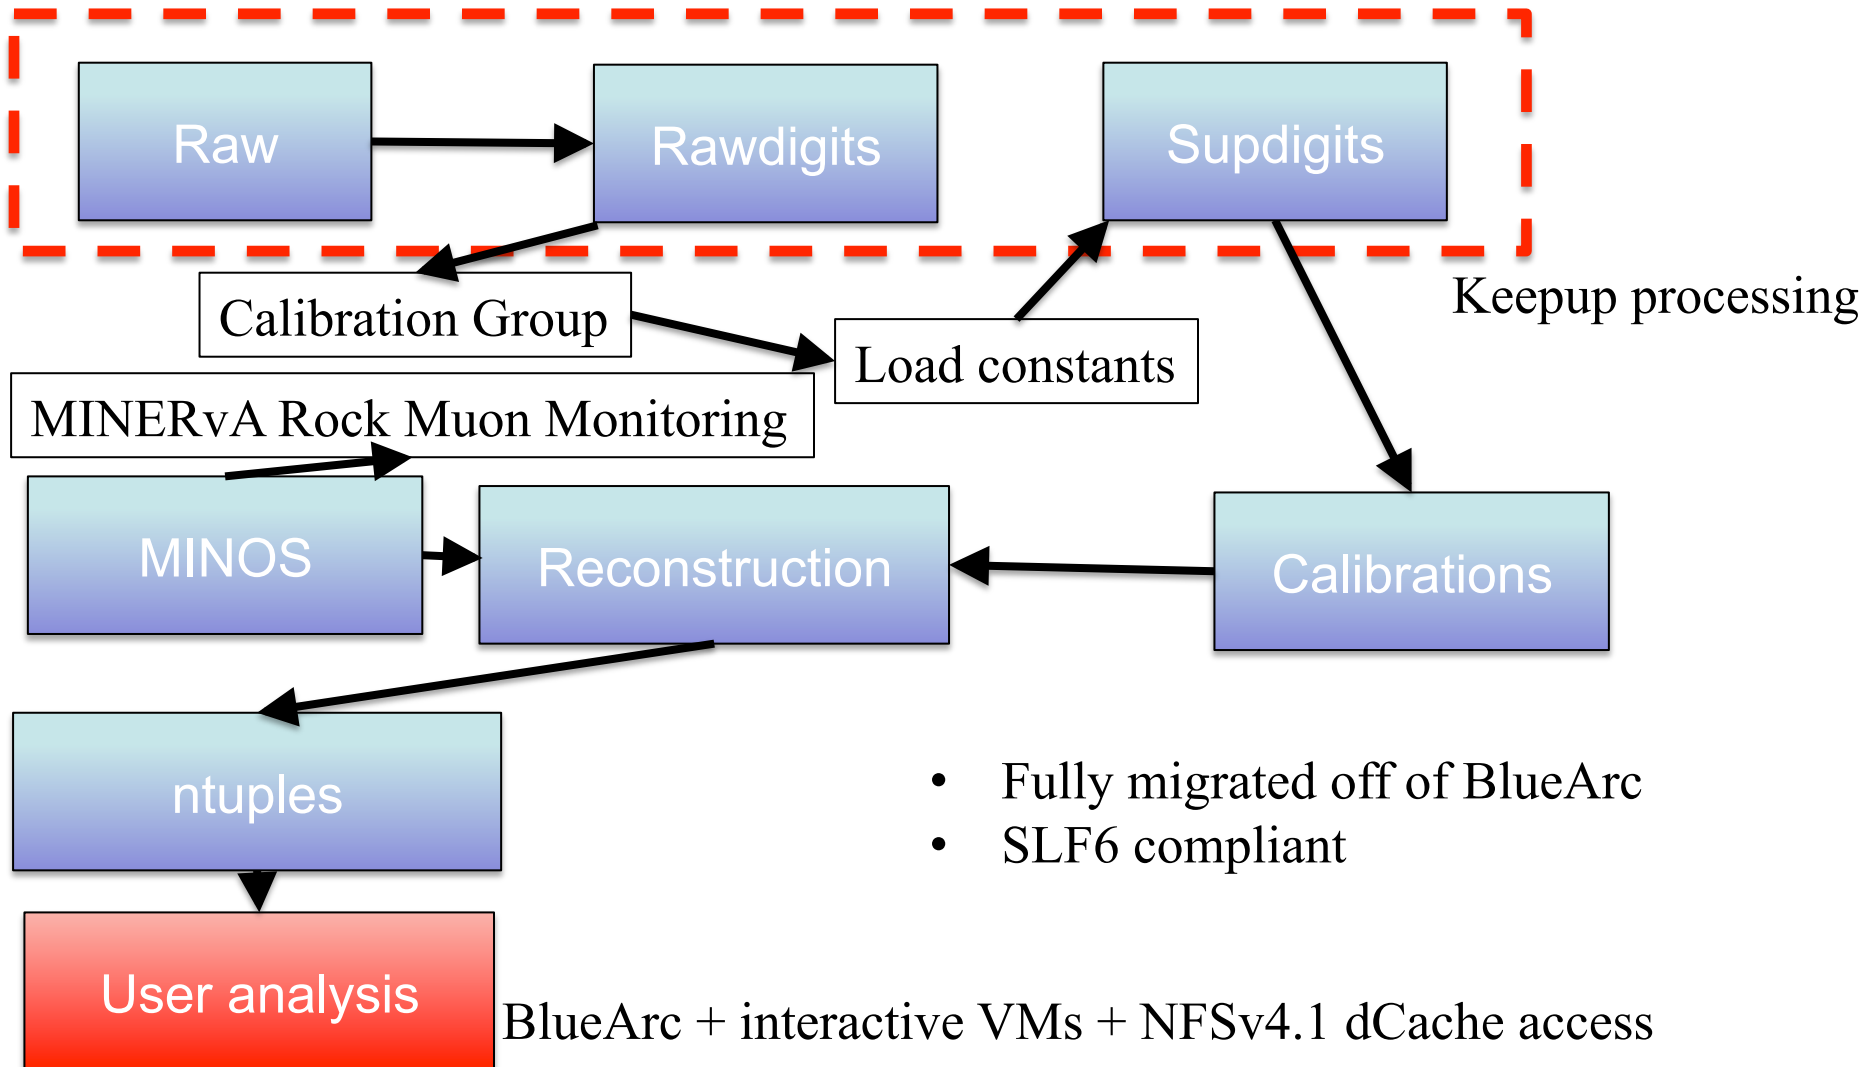


- Question 1(D): The model for data processing and analysis including the budget and effort required
- Question 1(E): SCD and collaboration resources

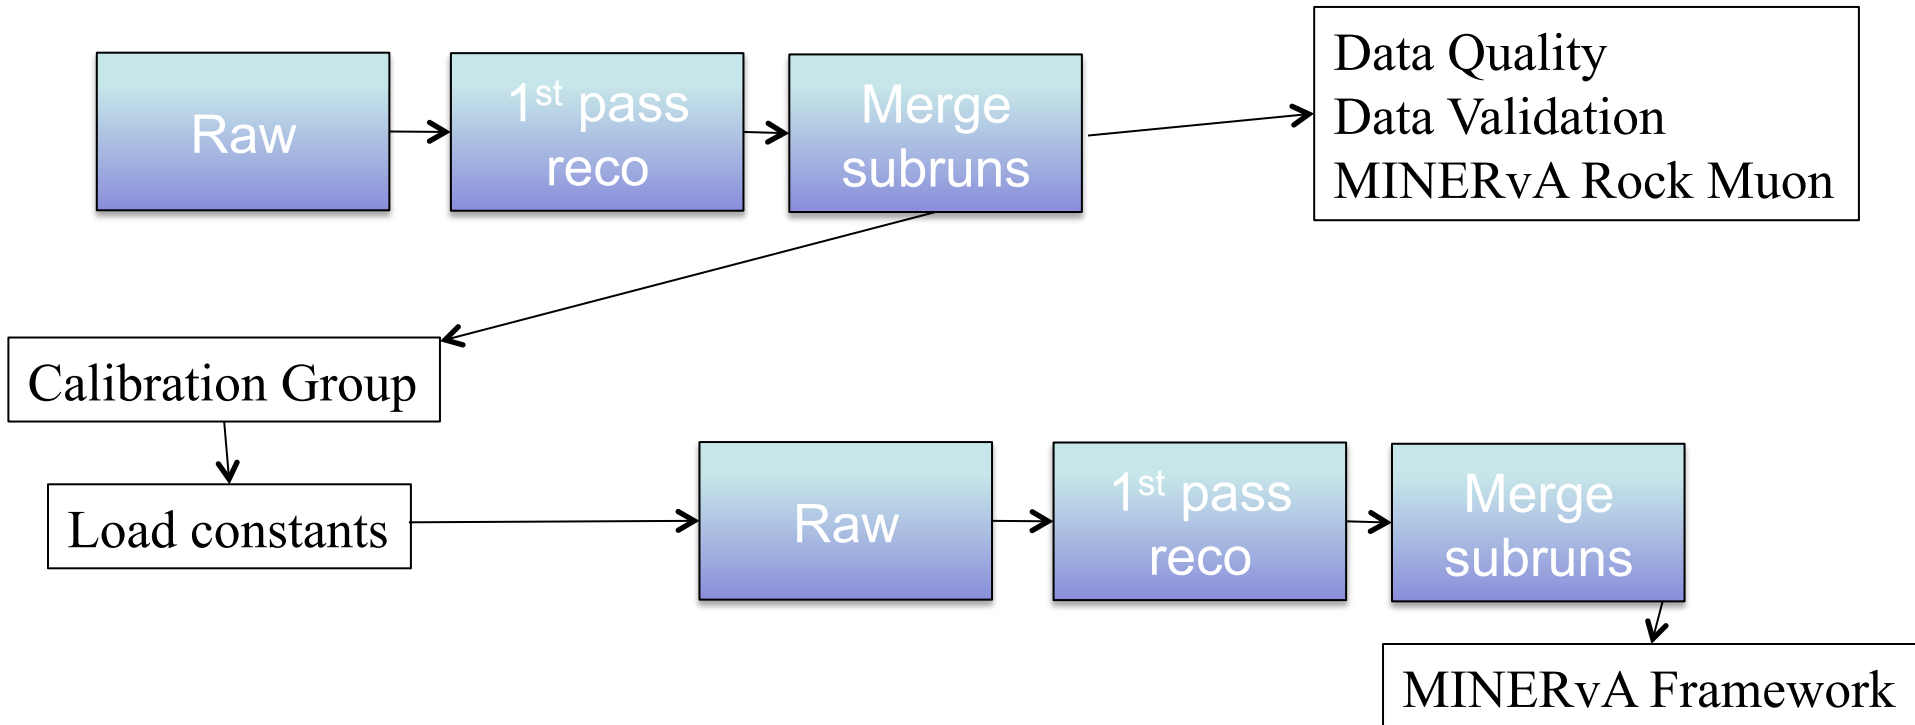

Future Nearline



MINERvA Data Model

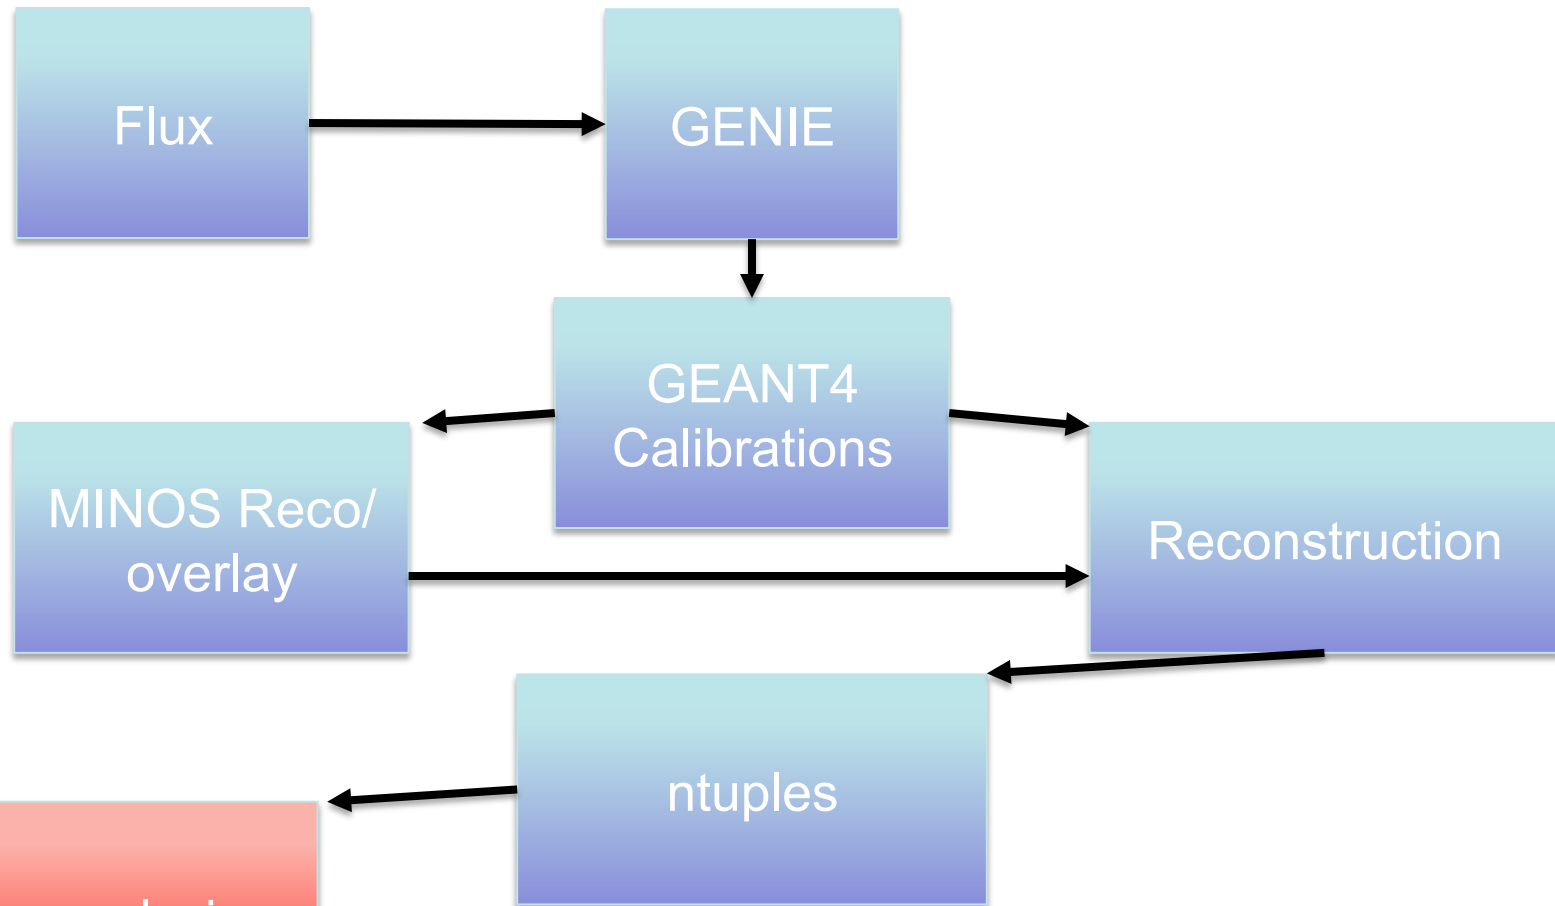



MINOS Data Model



- Used to be MINOS collaborators who managed and ran jobs – expertise has been transferred to MINERvAfs
- Physical machines are currently SLF5
 - Tested this process on SLF6
- Migration off BlueArc nearly fully tested

MC Generation Model



BlueArc + interactive VMs+ NFSv4.1 dCache access

Computing and collaboration resources



- MINERvA Collaboration Resources

- Computing Infrastructure [1.0 FTE]

- Daniel Ruterbories (Lead) U. Rochester Postdoc
- Philip Rodrigues U. Mississippi Professor

- Production [0.7 FTE]

- Dipak Rimal (Lead) U. Florida Postdoc
- Tejin Cai U. Rochester Graduate Student

- Software Release [0.5 FTE]

- Trung Le (Lead) Tufts Postdoc

- Keepup MINERvA [0.25 FTE]

- Anushree Ghosh (Lead) USM Postdoc
- Heidi Schellman Oregon State U. Professor


- Keepup MINOS [0.2 FTE]

- Dipak Rimal (Lead) to be handed off U. Florida Postdoc

- MINOS data quality / validation [0.2 FTE]


- TBD who maintains scripts but not evaluate the plots.

Computing and collaboration resources




- SCD Resources
 - Liaison for efficient communication – Gabe Perdue
 - Grid and Cloud computing
 - Scientific Collaboration Tools
 - Scientific Computing System
 - Scientific Data Management
 - Scientific Data Storage and Access
 - Scientific Databases
 - Scientific Software
 - Simulation Software

Computing and collaboration resources



- Core CD Resources
 - Authentication and Directory
 - Central Web Hosting
 - Database Hosting
 - Desktop Services
 - Fermilab (Data Center) Facilities
 - Network Services
 - Networked Storage Hosting
 - Service Desk

Budget - media



Media Type - Tape	TB written in the past year	Adjusted TB*	Cost
tape raw binary	16	18.67	\$560
tape supp / cal digits	15.5	18.08	\$543
tape reco	23.2	27.07	\$812
tape simulation	418.3	488.02	\$14,641
Media Type - Disk	Total TB allocated (private) / in-use (public)		
persistent dCache	193		\$5,629.17
raw data write dCache	132		\$3,850.00
read/write pools dCache	470		\$13,708.33
public scratch dCache	76		\$2,216.67
Total media cost			\$41,959.17

EOP Table 2

Tape adjusted to 14 month period.

Disk is actual usage or private allocation

Budget – CPU Time




Marginal cost of computing	1 year CPU hours (millions)	CPU Hours**	Cost
data processing	3.9	4.55	\$45,500
simulation processing	11.7	13.65	\$136,500

** Adjusted to 14 months

EOP Table 2

Budget – SCD Personnel and total



Computing Labor	Neutrino Experiments (total FTE)	MINERvA (total FTE)	Neutrino Experiments (Incremental)	MINERvA (Incremental)
Software Development and Support	17.92	0.93	6.72	0.64
Operations	16.73	2.5	4	0.55
Facilities	22.4	2.45	4.5	0.53
Total	57.05	5.88	15.22	1.72

Total computing and media cost	\$223,959
SCD Support Services (5.88 FTE-year, 14 month period)***	\$1,715,000
Incremental Support Services (1.72 FTE-year, 14 month period)	\$501,667
Grand total (incremental)	\$725,625

*** FTE estimate based off of DOE operations review, from service desk tickets. Source M. Votava, A. Lyon, S. Fuess


EOP Table 2

Charge Question



- Question 4: Are there robust plans for data processing and data analysis? Have adequate resources from the collaboration been identified for data analysis to meet the set goals?

Computing Plan



- To meet the analysis goals we need the following processing campaigns per year
 - 1-2 small 1x data POT MC samples
 - 1 large 10x data POT MC sample
- In addition continued processing for data monitoring and calibrations

Computing Resource Usage




- 1 full processing @ 10E20 POT (10x data stat.)
 - Disk space usage is 1.5PB
 - 750TB is direct input into user analysis (reco files)
 - Remaining files typically reside on tape only
 - CPU Usage is 6M CPU-hours
 - This is JUST processing time
 - Efficiency corrected is 7.5M CPU-hours (80% efficiency)
- Small passes
 - Disk Space usage 150TB
 - CPU Usage efficiency corrected 750k CPU hours

CPU Time Allocation




- We are given a quota of 1800 “slots”
 - Equivalent to ~16M CPU hours
- CPU time → real time is 175 days
- Add overhead for book keeping and resubmissions 25%
 - Catalogue, transfer, audit files
- Allocation left for users:
 - 7M CPU Hours
- Historically users are ~60% of our usage, but in the ME era they will not process the full dataset every time as was typically done in LE

CPU Time Allocation




- MINERvA production is 1 technical step and a testing phase away from being OSG enabled
 - Plan is to push production, when possible, to OSG
 - Timeline is end of year
- Bottom line – MINERvA can fit within the allocated time on GPGrid completely, **but** will be OSG enabled.

Disk Space



- Currently have 2 large storage solutions for users
 - 240 TB BlueArc (local access only), 200 TB dCache
- Plan is to reduce BlueArc space, utilizing dCache more.

Databases




- Conditions database – no foreseeable issue
- MINOS replica database – static in size, just need it up
- MINOS data – no foreseeable issue in growth

Interactive Machines



- Currently have 5 at our disposal
- With longer processing times and many active analyses we may need to explore the possibility of a few more

Conclusions



- Through the yearly reviews and regular communication SCD understands our computing needs and helps us meet our scientific goals
 - Spokespersons and the head of SCD
 - Offline coordinator and SCD experts
 - Collaboration and our SCD Liason – Gabe Perdue
- MINERvA is always consulting with SCD about new products and ways to adopt them
- With MINERvA adding OSG capability we will continue our contribution to the local effort in understanding techniques that benefit the lab's experiments.