Fermilab

Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

DarkSide-50 (E-1000) FY15 and FY16 Computing Needs

Stephen Pordes

Scientific Computing Portfolio Management Team (SC-PMT) Review

Mar 5 2015

The DarkSide-50 Collaboration
7 countries, 31 institutions (15 U.S.)

IHEP

APC, Université Paris Diderot
IPHC, Université de Strasbourg

INFN LNGS
Università degli Studi Genova
Università degli Studi Milano
Università degli Studi Federico II Napoli
Università degli Studi Perugia
Università degli Studi Roma Tre
Università degli Studi Cagliari

Jagiellonian University

KINR, NAS Ukraine

Joint Institute for Nuclear Research – Dubna
Lomonosov Moscow State University
Kurchatov Institute – Moscow
Saint Petersburg Nuclear Physics Institute

Augustana, Black Hills State, Fermilab, LLNL, PNNL
Princeton, SLAC, Temple University
University of Arkansas, UCLA, University of Chicago
University of Hawaii, University of Houston
University of Massachusetts, Virginia Tech

People & Responsibilities at Fermilab

• Trigger (FPGA based)	Boris Baldin (PD - EED)
• DAQ (artdaq)	Kurt Biery (SCD)
• Argon system	Cary Kendziora (PD - MED)
• Offline	Ken Herner (SCD)
• Post doc	Yann Guardincerri (PD - AD - DMI)
• Old person	Stephen Pordes (ND)

The Gran Sasso from the East

Fermilab

Schematic of DarkSide-50 Detector

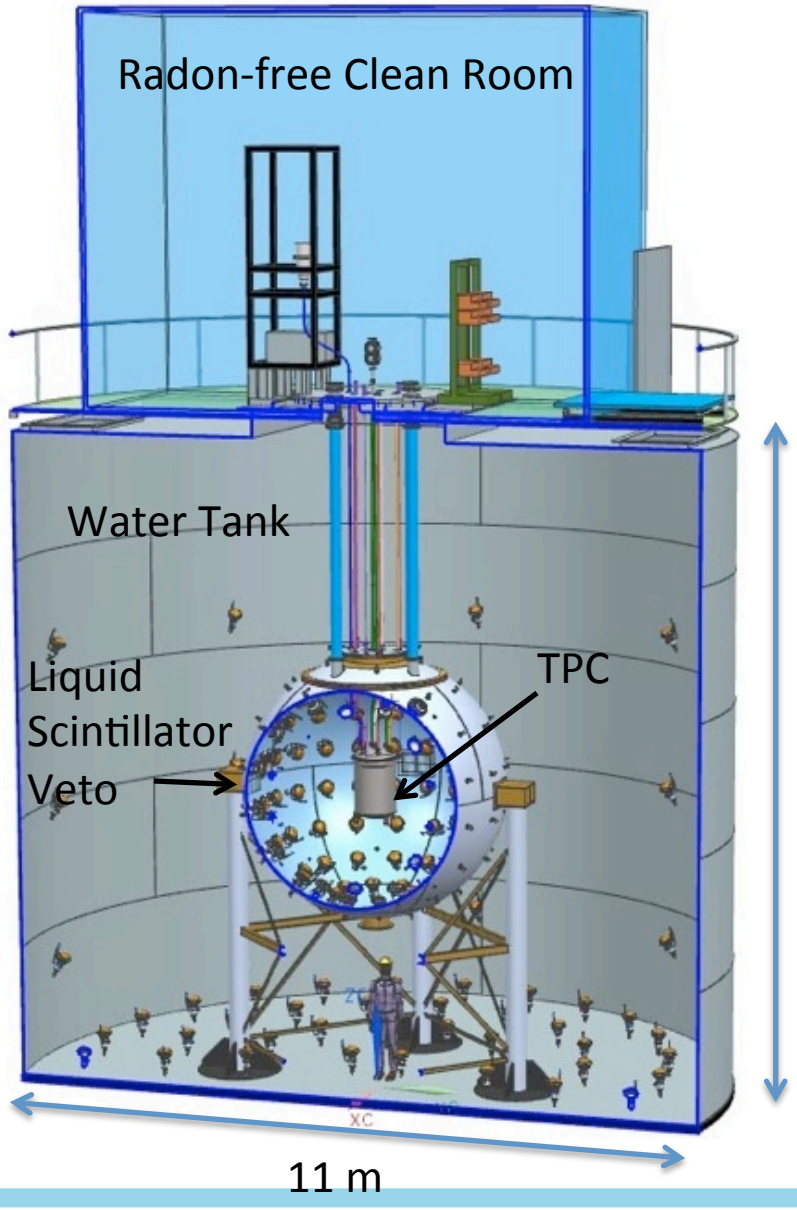

DarkSide-50 TPC Schematic

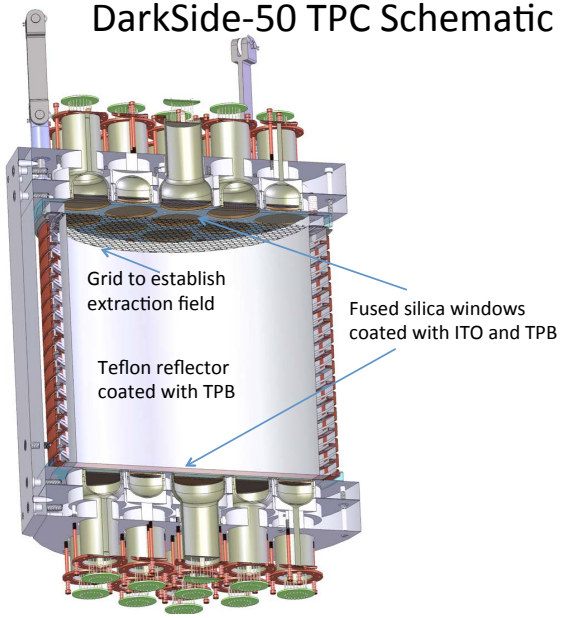

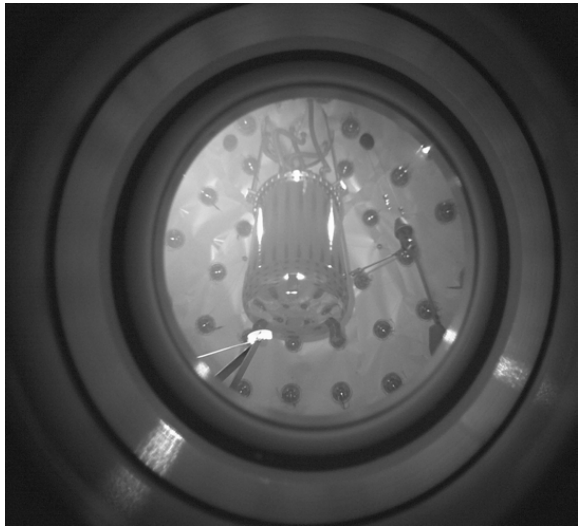

TPC
in
LSV

Argon produces scintillation light and allows long (meters) free electron drift.

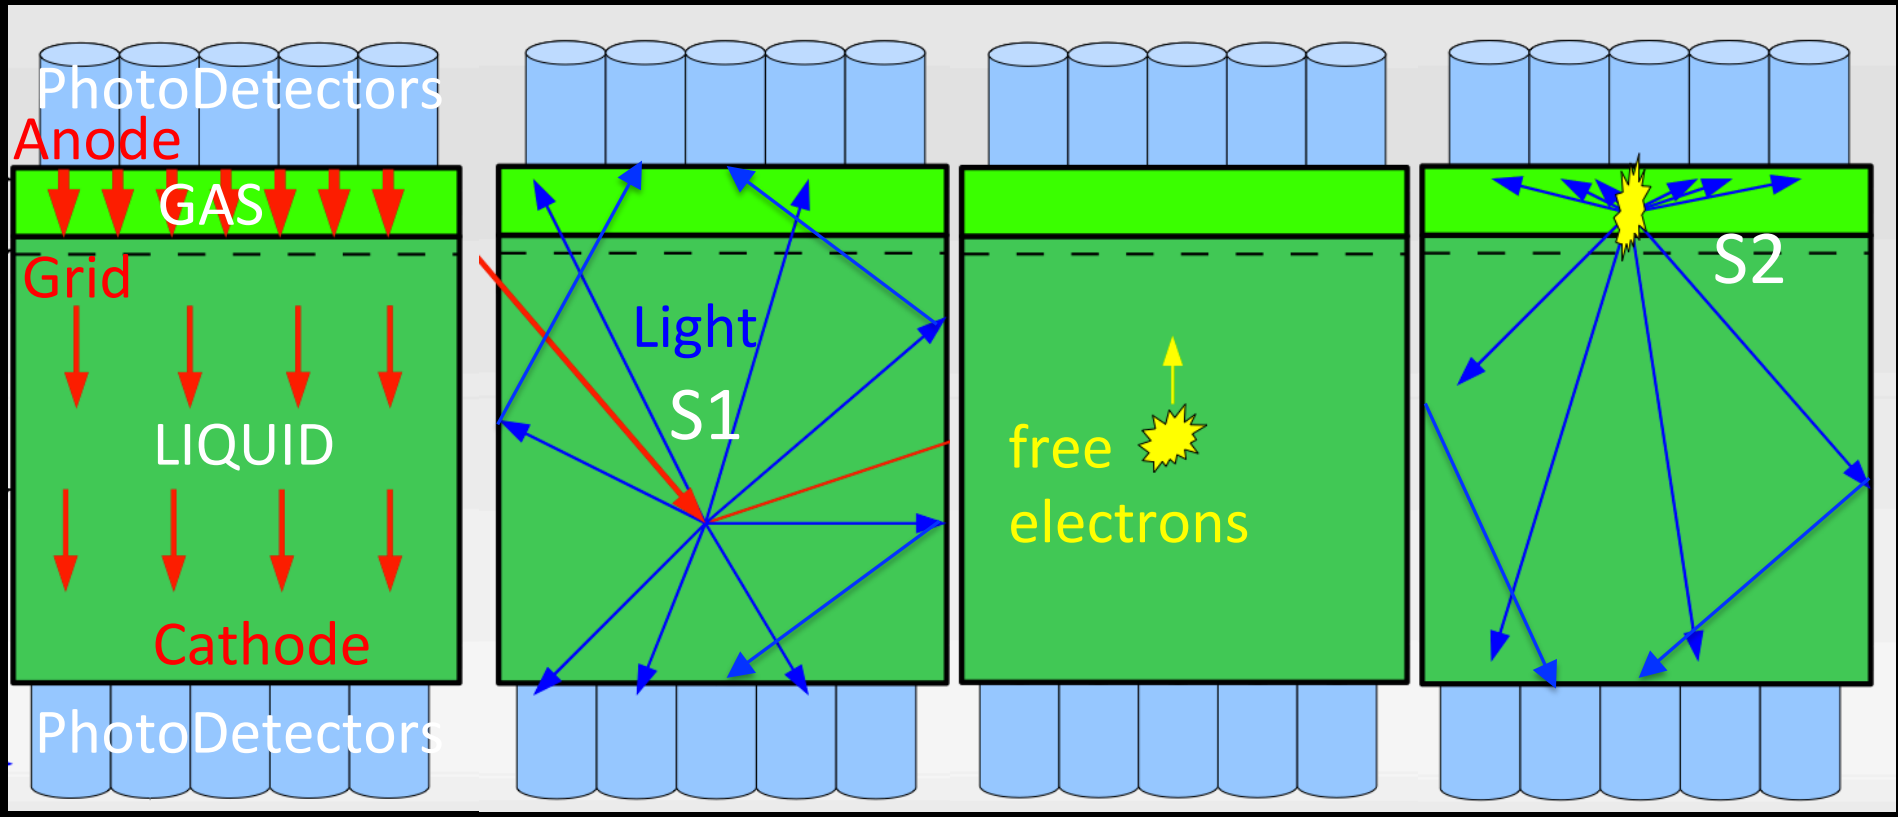

Liquid Phase;
Gas Pocket above;
Electric Field in both;
0.2 kV/cm (Liquid),
3.5 kV/cm (gas)

WIMP interacts;
Nucleus Recoils
Argon produces
light (S1) and
free electrons

Electrons drift
under E field
are extracted
by field of
3 .5 kV/cm

into gas region
where they produce
secondary
fluorescence (S2),
proportional to
number of e's.

Typical event waveform

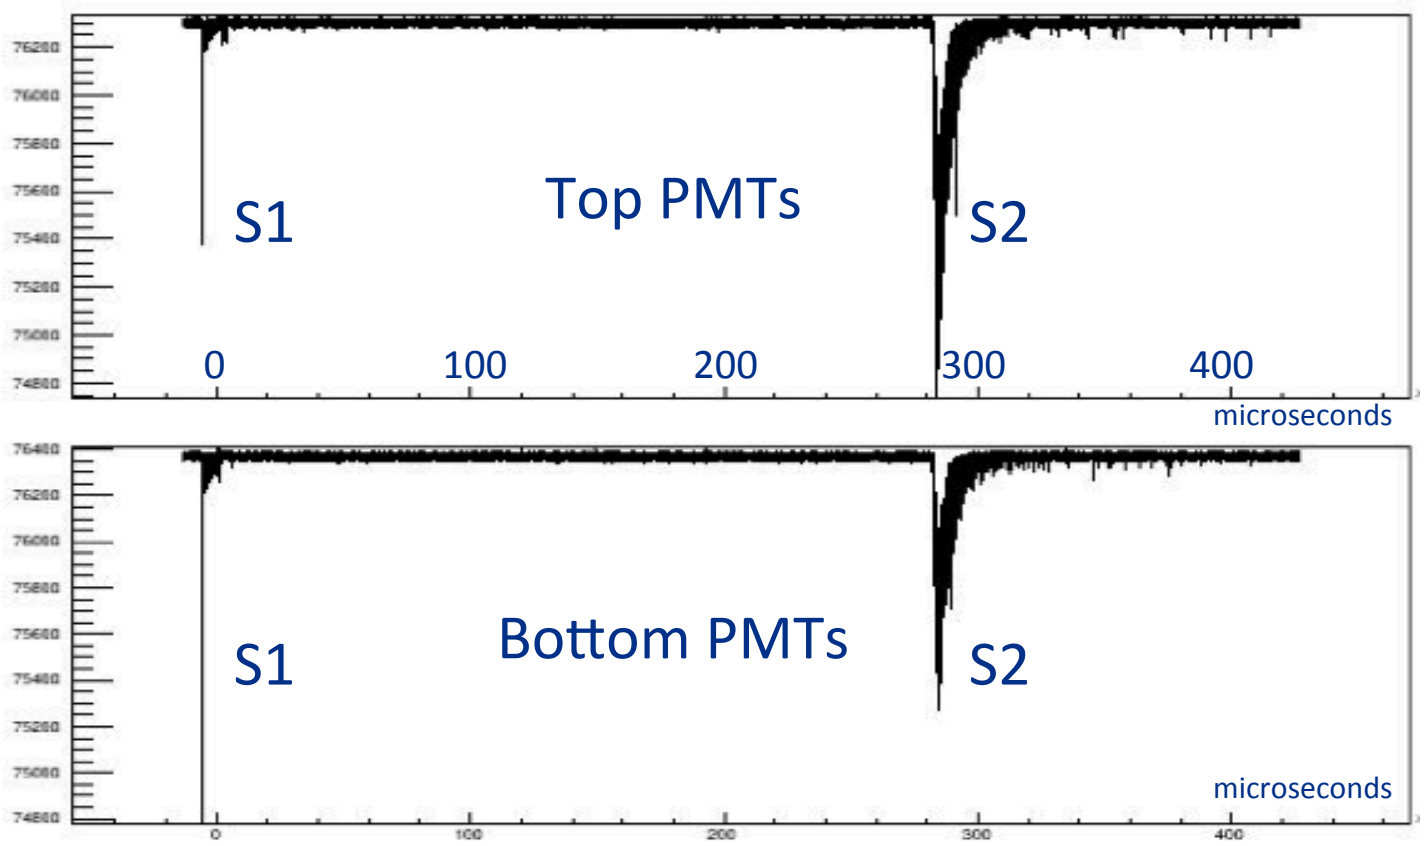

Event data = 440 us x 38 x 250 MHz x 2 bytes ≈ **8.5 x 10^6 bytes** + some

Event rate from 39Ar decay ~ 50 Bq; trigger rate in ROI ~ **15 Hz** ..

DAQ readout rate ≈ 120 MB/s (fully live at this rate) – can go x 3 faster

Huffman compression factor ~ 5 =**> 25 MB/s** to disk at LNGS, and **dCache** at FNAL

**** **with underground (low-radioactivity) argon, trigger rate ~ 3 Hz or less** *****

Fermilab

Scientific Goals for FY15 and FY16

- Exploit atmospheric argon for further rejection studies and detector optimization studies (2 months)
- **Run with underground (low-radioactivity) argon for dark matter search from ~ April 2015 onwards.**
- Occasional system calibrations for argon technology
- **Aim to produce interim result this (2015) year**
- **Have requested to DOE to run till mid 2018 (3 yr run)**
- **LNGS Scientific committee has asked for report on running with underground argon at April 28th 2015 meeting**

Large Scale or out of ordinary computing needed to complete these goals – in general

- It might not be expected that a dark matter search for a few events/year would produce 100's of Terabytes of data
- While in absolute (or compared to other projects) our needs are not extraordinary, they are significant.
- Our needs are driven by the need to **establish rejection at the level of 10^9**, which prima facie requires a few 10^9 events.
- We have done much of this study while waiting for the low-radioactivity argon. (We had imagined doing the study later in the experiment run.) This rejection-study data set is a major part of the explanation for the fact that we are using **550 TB** ~ 3 times the storage we requested.

DarkSide-50 Data Acquisition Schematic at LNGS

Clean Room above detector

Control Room ~ 100 ft away

CAEN digitizers V1720 & V1724

FR = data fragment receiver

EB = event builder

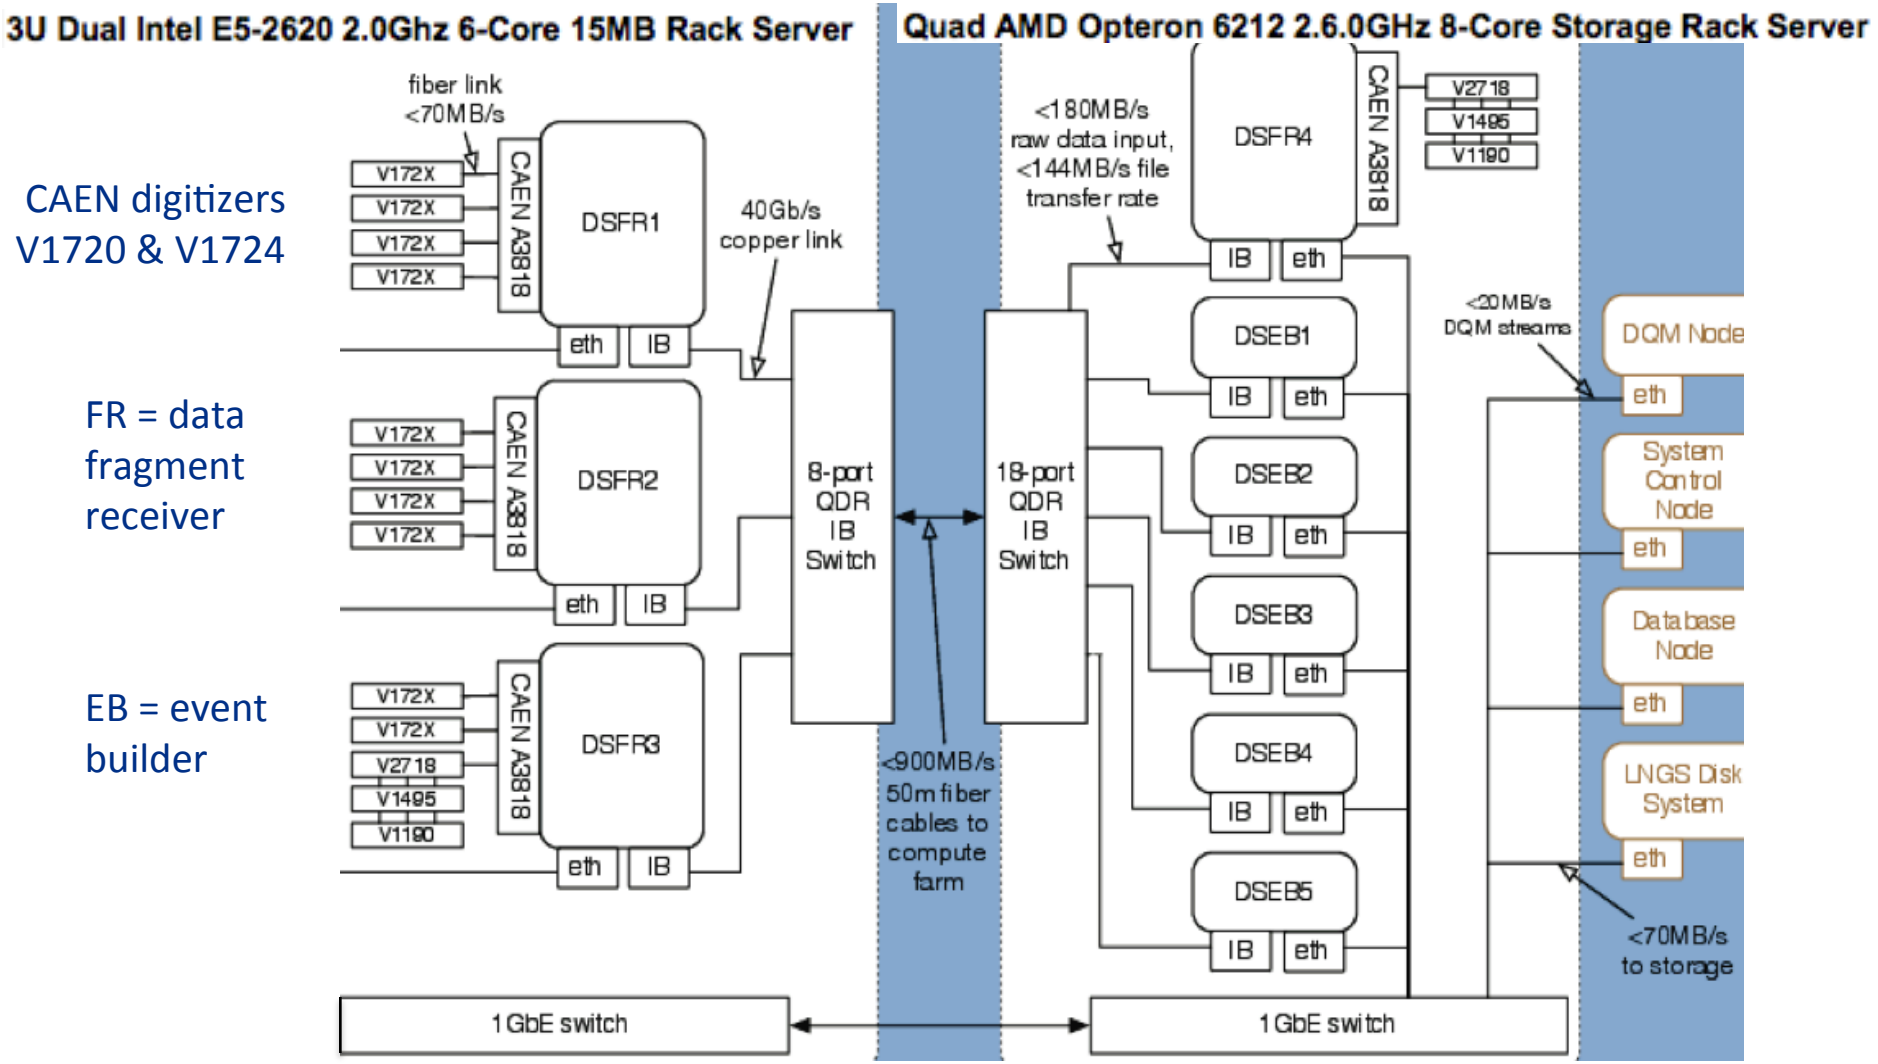

Fermilab

Did we meet our FY14 Scientific Goals?

- No – we had planned to run with underground (low radioactivity) argon and the argon was not available
- Yes – we took data with atmospheric argon and demonstrated excellent rejection against electromagnetic (gamma-ray and electron) background

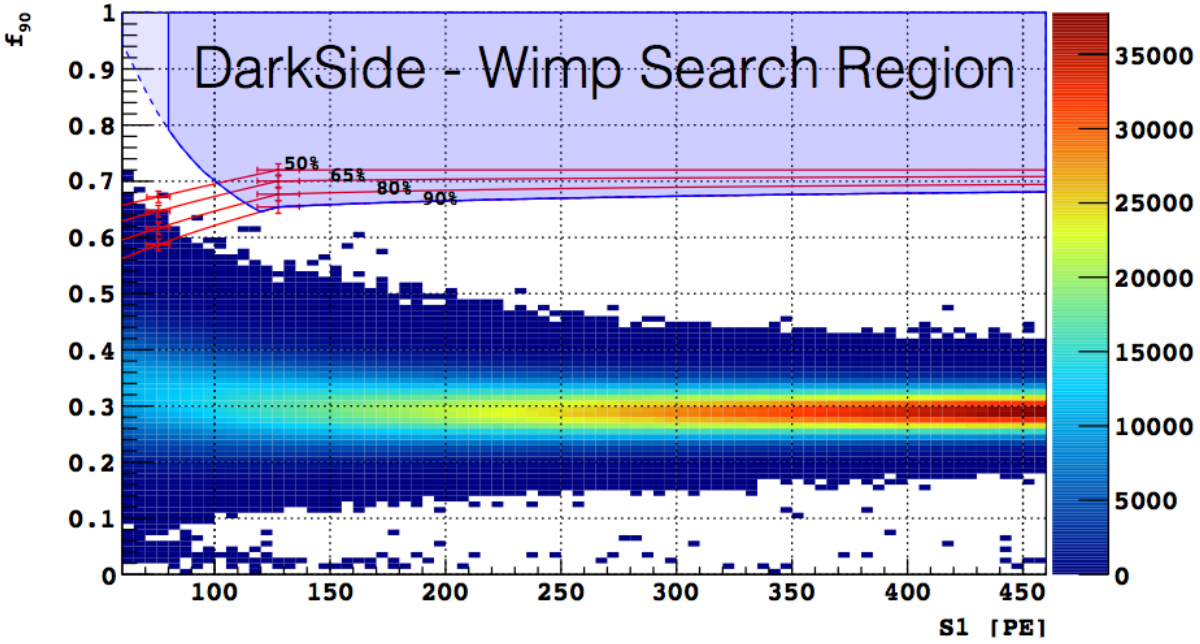

The scatter-plot shows two quantities for all the events recorded in this DarkSide-50 run. The x-axis is the energy of the event - the brightness of the first light flash. The y-axis is essentially the inverse of the duration of the flash (in time) - shorter pulses correspond to larger values of f90. The big splash of color is from the radioactivity of the argon itself. A single WIMP-like signal would give a blue square in the top shaded region.The red lines show the acceptance for WIMP events.

Did we meet our FY14 Scientific Goals?

Compared to last year's requests, the consequences of doing rejection studies early were:

- used more storage ~ 400 TB than requested 100 to 200 TB *
- stressed the Data Acquisition system in its infancy

In practice, the SCD was able to accommodate our needs, and we established an efficient running environment for the experiment. **

* part of the increased storage required came from our SCENE experiment results (Phys.Rev. D 88 (092006) 2013) showing that we needed to run at lower drift-fields to maintain the light-yield for the WIMP signal; lower drift-field means longer drift-times and therefore larger event sizes.

** I recognize that we need to pay for the storage, and funds are available so to do.

DarkSide-50 Data Distribution & Reconstruction Scheme Present

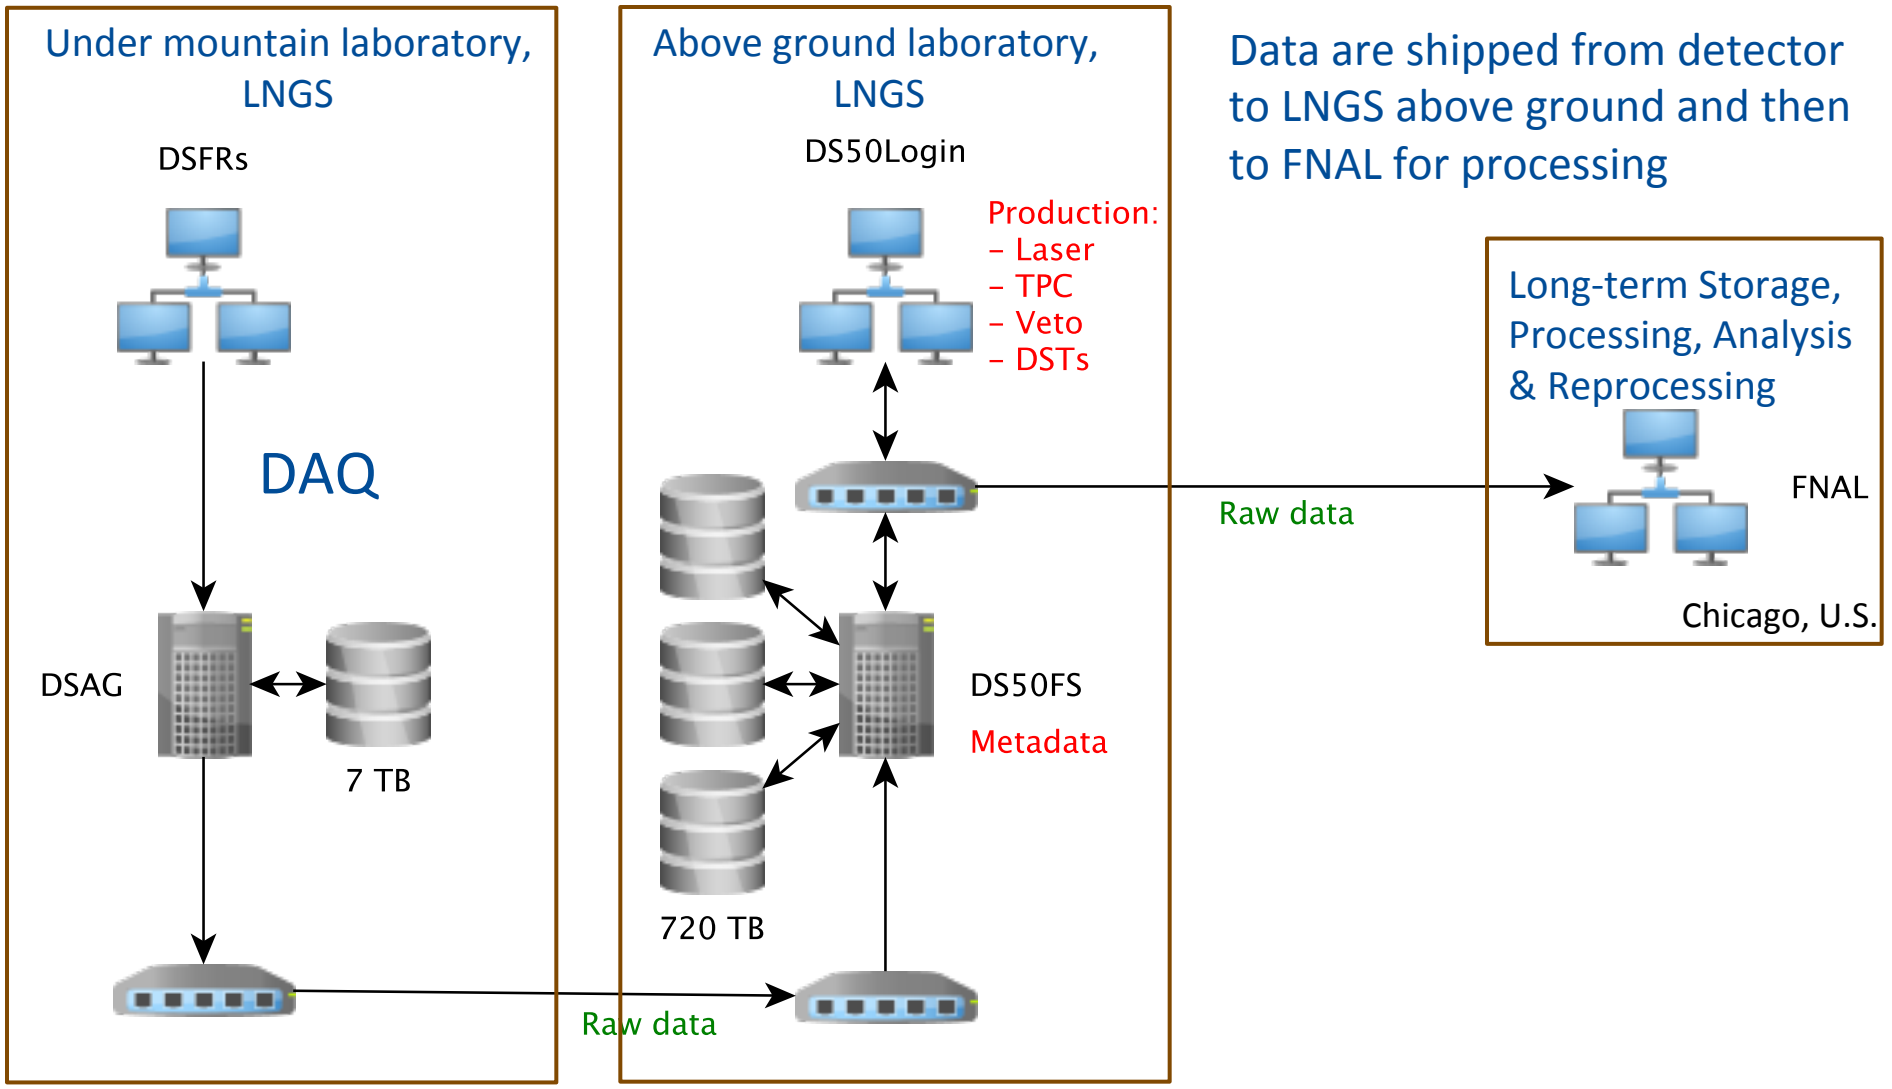

Data are shipped from detector to LNGS above ground and then to FNAL for processing

Fermilab

DarkSide-50 Data Distribution Scheme (to be implemented by summer 2015)

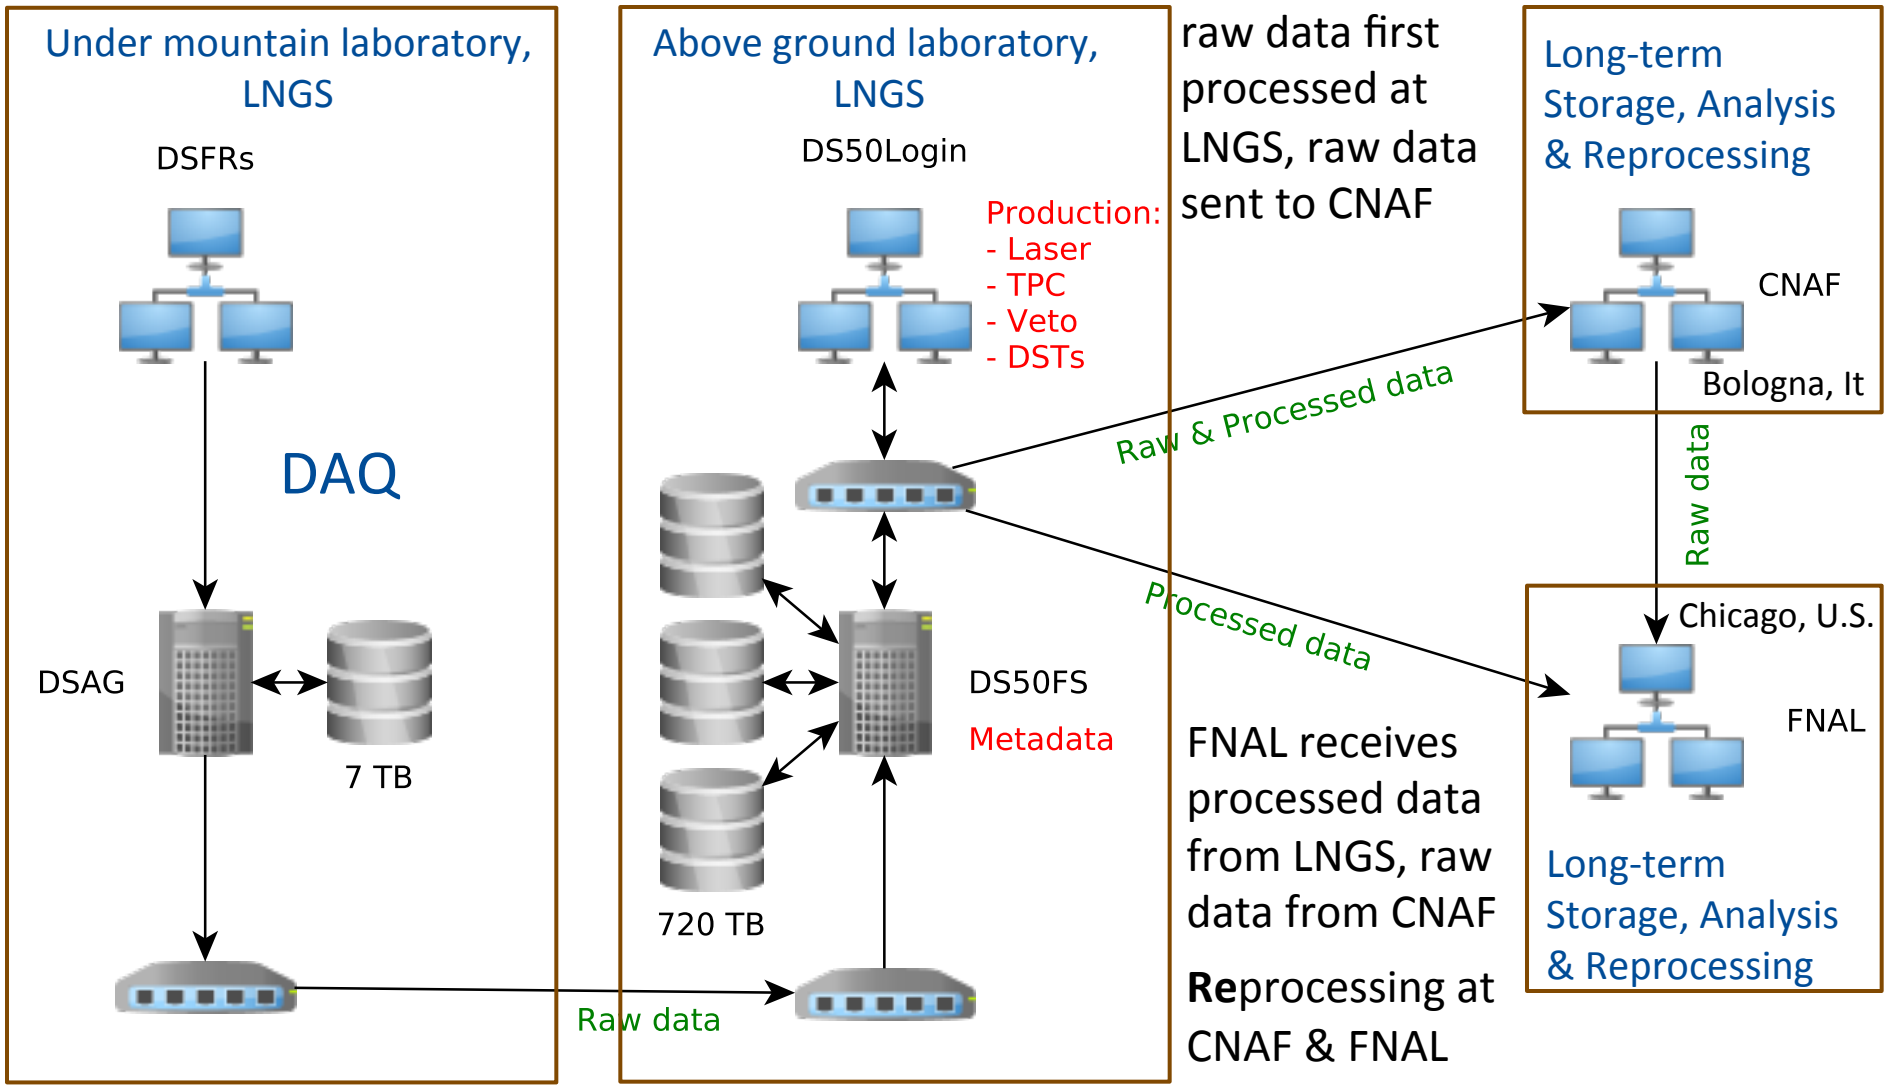

Needs from Service Areas

Table of Service Areas being used

Services	Offerings	DarkSide
DAQ and Controls	DAQ	In Production
Grid and Cloud	Fermigrid	In Production
Grid and Cloud	Fermicloud	
Grid and Cloud	Gratia	In Production
Grid and Cloud	Job Sub	In Production
Grid and Cloud	FIFEMON	In Production
Grid and Cloud	OSG Enabled	Testing
Grid and Cloud	Amazon	
high performance computing		
Scientific Data Storage and Access	dcache/ enstore	In Production
	Gridftp	In Production
Scientific Data Management	IFDH	In Production
Scientific Data Management	SAM Web	In Production
Scientific Data Management	FTS	In Production
Scientific Frameworks	Software Framework (art)	In Production
Scientific Frameworks	Continuous Integration - Orchestration	
physics and detector simulation	Genie	
physics and detector simulation	geant4	In Production
Scientific Software	Software Processing (LArSoft)	
Scientific Databases	IF Beam	
Scientific Databases	Custom Databases	
Databases	mysql	
Databases	postgres	YES
experiment specifc services	Production Operations	
networking		
Scientific Computing Systems	CVMFS	In Production
Scientific Computing Systems	Interactive machine in GPCF	In Production
Scientific Computing Systems	Experiment Control Rooms	NO
Scientific Computing Systems	Continuous integration - middleware	NO
Scientific Collaboration Tools	Redmine	In Production
Scientific Collaboration Tools	CVS/ Subversion/ Git	In Production
Scientific Collaboration Tools	Electronic Log	Custom
Scientific Collaboration Tools	ups/upd	In Production
Central Web Hosting		
Scientific Document Management	doc dB	In Production
Video Conferencing		
Futures	Federated Data Management	xrootd
Futures	High Throughput Analysis Facilities	

Fermilab

Spreadsheet for DarkSide-50

Service	Service Offering	Offering details	Definition of row or column	notes below	(if below threshold, then enter "Yes" if small request)	(eg previous 3 months, to be supplied)	(allocation and/or utilization of resources, to be supplied)	immediate reallocation of resources.	level then requires gradual reallocation of resources.	then requires purchases or reallocation in FY15.	then requires purchases or reallocation in FY16.	Use footnotes if necessary.
SCIENTIFIC COMPUTING SYSTEMS												
	Batch Worker Nodes	FermiGrid yearly integral	# CPU-hours	1	100,000	340,000 hrs = past year	same as prior	same as column G	300 khrs	500 khrs	550 khrs	A
	(assume all CPUs equivalent)	FermiGrid peak integral	# CPU-hours	2	any	40000 hrs in 1 week	same as prior	same as column G	same as column I	same as column J	same as column J	B
		FermiGrid peak duration	# of hours	2	any	48 hrs	same as prior	same as column G	same as column I	same as column J	same as column J	B
		FermiGrid peak count	# of peak periods	2	any	4	same as prior	same as column G	same as column I	same as column J	same as column J	C
		OSG opportunistic yearly integral	# CPU-hours	1	100,000	100 hrs (not vislble)	same as prior	same as column G	same as column I	same as column J	same as column J	
		OSG opportunistic peak integral	# CPU-hours	2	any	n/a	same as prior	same as column G	same as column I	same as column J	same as column J	
		OSG opportunistic peak duration	# of hours	2	any	n/a	same as prior	same as column G	same as column I	same as column J	same as column J	
		OSG opportunistic peak count	# of peak periods	2	any	n/a	same as prior	same as column G	same as column I	same as column J	same as column J	
		External dedicated yearly integral	# CPU-hours	3	100,000	n/a	same as prior	same as column G	same as column I	same as column J	same as column J	
		Large Memory or Multi-CPU	Describe needs in Comments		any	n/a	same as prior	same as column G	same as column I	same as column J	same as column J	
	Server & Storage Support	Static Interactive Service	# of Static Services (eg VMs)	4	any	1	same as prior	same as column G	may add 1 -> 2	same as column J	same as column J	D
		Other Static Services	# of Static Services	4	any	n/a	same as prior	same as column G	same as column I	same as column J	same as column J	
		Dynamic Services, average	# of Dynamic Services	5	2	n/a	same as prior	same as column G	same as column I	same as column J	same as column J	
		Dynamic Services, peak	# of Dynamic Services	2	10	n/a	same as prior	same as column G	same as column I	same as column J	same as column J	
		cvmfs Service	Repository (Yes or No)		Yes	Yes	same as prior	same as column G	same as column I	same as column J	same as column J	
		Build & Release Service	Use facility (Yes or No)	6	Yes	n/a	same as prior	same as column G	same as column I	same as column J	same as column J	
		Database Service	Specify type(s), numbers	7	any	1 postgres replica	same as prior	same as column G	same as column I	same as column J	same as column J	
		Other Disk Service (specify)	Servers with attached disk	8	any	1	same as prior	same as column G	same as column I	same as column J	same as column J	
SCIENTIFIC DATA STORAGE & ACCESS												
	dCache	Shared RW	Cache disk storage (TB)	1	20	520	520	100 more ->620	50 more -> 670	150 more -> 820	150 more -> 970	E
		Shared RW lifetime	Cache disk desired lifetime (days)	1	10	30	30	30	60	same as column J	same as column J	
		Shared Scratch	Cache disk storage (TB)	1	20	5	5	7	50	same as column J	same as column J	
		Shared Scratch lifetime	Cache disk desired lifetime (days)	1	10	100	100	100	100	same as column J	same as column J	
		Dedicated Write	Cache disk storage (TB)	1	any	0	0	0	100 ?	same as column J	same as column J	F
	enstore	New/additional capacity	Tape media (TB)	1	25	520	520	100 more ->620	50 more -> 670	150 more -> 820	150 more -> 970	
NETWORKED STORAGE												
	NAS/BlueArc	*-app	Dedicated NAS (TB)	1	any	1 TB	same as prior					
		*-data	Dedicated NAS (TB)	1	any	25 TB	same as prior					
		*-prod	Dedicated NAS (TB)	1	any	n/a	same as prior					
		*-ana	Dedicated NAS (TB)	1	any	n/a	same as prior					
NETWORK SERVICES												
	Physical Infrastructure	DAQ LAN bandwidth	Dedicated for DAQ		any 10GE	n/a						
		LAN bandwidth	Specific to experiment		any 10GE	n/a						
	WAN Infrastructure	DAQ WAN bandwidth	Dedicated for DAQ		any	n/a						
		WAN bandwidth	Specific to experiment		average > 2 Gb/s	need 50 MB/s	same	same	less data/ may need higher inst. rate	same as column J	same as column J	G
		Dedicated WAN circuits	Dedicated for experiment		any	n/a						

Supplied info

Requested info

Notes:

1 If a wide range, enter single number of best estimate and make note in Comments column
2 If peaks expected to be reached > 10% of time, make note in Comments column
3 Remote sites other than OSG, note in Comments column
4 Assume a static server (eg GPCF) is 4 cores, 12 GB memory
5 Assume a dynamic server (eg FermiCloud) is 1 core, 2 GB memory
6 Build and Release facility is in progress. State Yes if planning to use central facility. Add Comment on frequency
7 This row is for the database servers; the database services are below.
8 This row is for dedicated servers with attached disk. Specify the number of such. Add Comment on capacity.

Comments:

A Based on experience - includes 2 reprocessing of fulll data-set and 200k hrs simulation; reprocessing goes as 100khrs (existing) + 25khrs * (yr-2014) (incremental); first term is calibration, second is search-mode data
B Peak comes from simulations -
C 2 Reprocessing, 1 Simulation - expect first reconstruction to be done in Italy
D may purchase 2nd server depending on load - not certain
E data rate in search mode should reduce starting 4-2015; calibraition mode remains but totals shold be less than in 2014.
F may find an advantage in a dedicated write disk for data from Italy
G 50MB/s will keep up with data in most circumstances - source in Italy may change from LNGS to CNAF in mid to late 2015

Physical Resources Summary:

550 TB so far; total will grow by 150 TB/yr (**growth rate small due to UAr**)

800k to 1.3M CPU-hrs/yr in future for all types of jobs

1 private server DS50srv01 – may go to 2

1 gpvm DS50gpvm01 (less powerful used when DS50srv01 busy)

Fermilab

Processing and Simulation Details

- Data (re)processing
  - Expect 2, possibly 3, passes per year
  - If all data reprocessed, jobs would access about 500 TB
    - Some might be done at CNAF
  - Currently jobs process a single ~7 GB file (could be changed if needed) and takes about 1 hour
  - Figure 70k CPU-hours per pass as upper limit (full dataset)
  - Reco output about 32x smaller; 15 TB for entire dataset per pass
- Simulation
  - Typically done in bursts; about 50k CPU hours each
  - Expect upper limit of 16 such bursts per year, 800k CPU-hrs total
  - Working w/developers to get code running offsite
- Total: 70k x 3 (data) + 800k (sim) + 200k (contingency) = 1.3 M

Needs from Service Areas – 1 of 4

- Network from Italy (LNGS and CNAF)
- Data storage
- CPU-hrs for reconstruction and analysis
- Data management and data storage management
- Support for use of OSG services
- DAQ
- Support for our server
- Docdb & ReadyTalk

    quantities for data storage and cpu in spreadsheet

Needs from Service Areas - 2 of 4

Key people/functions

- **Data management and data storage management**
- **Support for use of OSG services**

  Our liaison, Ken Herner, has been and is crucial; we would very much like his continued support.

- **Data Acquisition**

  We are (one of) the first users of artdaq. Kurt Biery and his group are key to this. The system is operating rather well and/but we have some continuing requests.

Needs from Service Areas - 3 of 4

Key people/functions

- **FIFE Project**

  "Darkside-50 is the first Astro experiment to be onboarded to many of the FIFE services. FIFE has made the process of setting up the computing infrastructure much easier and means DS-50 is following best practices from the beginning."

Needs from Service Areas (4 of 4)

- Grid and Cloud Computing
  - Infrastructure established. Will use if necessary
- Networked Storage
  - BlueArc - keep as is
- Physics and Detector Simulation
  - Geant4 (not yet)
- Scientific Data Management
  - SAM – cataloging but not yet exploiting for file delivery (planned)
- Scientific Data Storage and Access
  - Enstore (tape), dCache (tape-backed and scratch - yes
- Scientific Frameworks
  - srt, ART, personnel needs - current level

TSW/EOP Status and Plans

- Are your TSWs signed and up to date?

  I believe we have an MOU, possibly the last MOU.
- If not, do they need revision?

  I cannot answer that

Future Directions (Challenges and R&D)

- Will your SOPs change significantly in the future (new phase of the experiment, new running conditions, etc.)?

  Data rate should reduce in a few months; simulations may increase; see the numbers on the spread sheets
- Are future R&D projects in the pipeline?

  DarkSide did not get G2 approval but was given encouragement to continue to develop the technology. The collaboration is preparing a next generation proposal to LNGS– which would involve modest DOE support. The FNAL involvement is not clear -DAQ would be a key part to be discussed at the time (fall 2015).
- Are additional resources going to be required to meet them?

  We would expect that the effort for a future detector would be similar in scale to getting the first version of artdaq going.

Additional Projects/Comments

- Anything else you want to say that was not covered in previous slides.

Fermilab SCD and CCD are absolutely key to DarkSide-50 – both for Online and Offline.

The success of the DAQ has been key to present results (and the challenge to the storage).

A dark matter experiment may not be expected to have order 1 PB of `data' and required 500k CPU-hrs of processing/year. The search data will probably be much less than 100 TB; the calibration and rejection data to prove that one can reject at the 10^9 level dominate our technical requirements

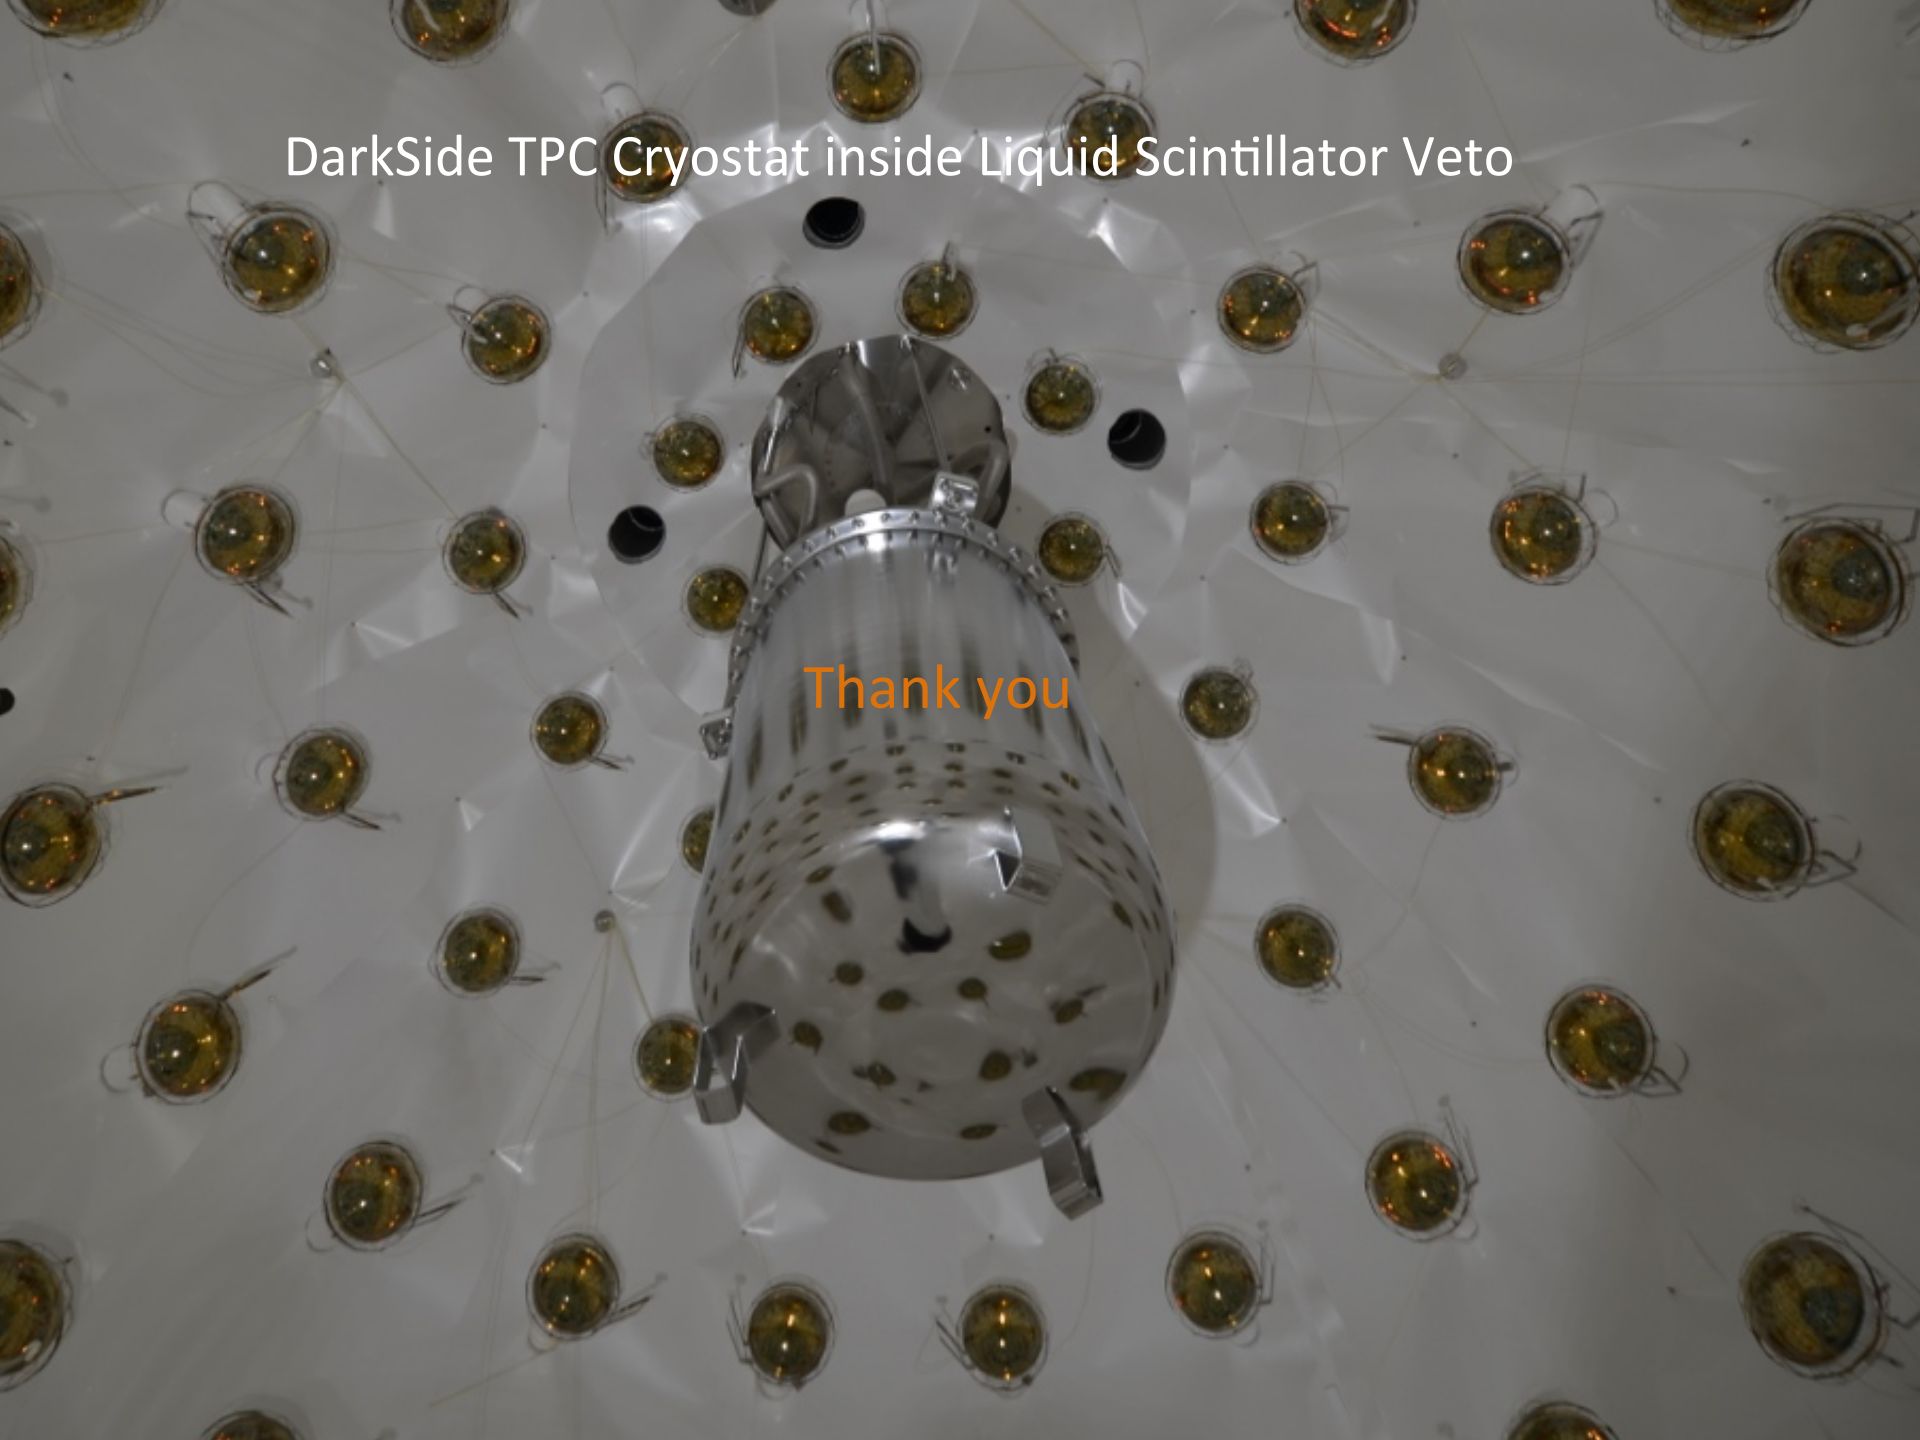
DarkSide TPC Cryostat inside Liquid Scintillator Veto
Thank you

extras

Latest Usage Plots

Storage Group	ArchivePools	DESPools	EmptyPools	ExpDbWritePools	FermigridVolPools	LQCDPools	MinervaWritePools	NovaWritePools	PublicScratchPools	RawDataWritePools	SfaPoolGroup	readWritePools	Total
nova	0	0	0	0	0	0	0	38,998	93,520	0	0	1,463,236	1,595,754
darkside	0	0	0	0	0	0	0	0	23,336	0	0	574,381	597,717
uboone	0	0	0	0	0	0	0	0	337,917	0	0	21,136	359,053
minerva	0	0	0	0	0	0	14,499	0	48,360	0	0	183,799	246,658

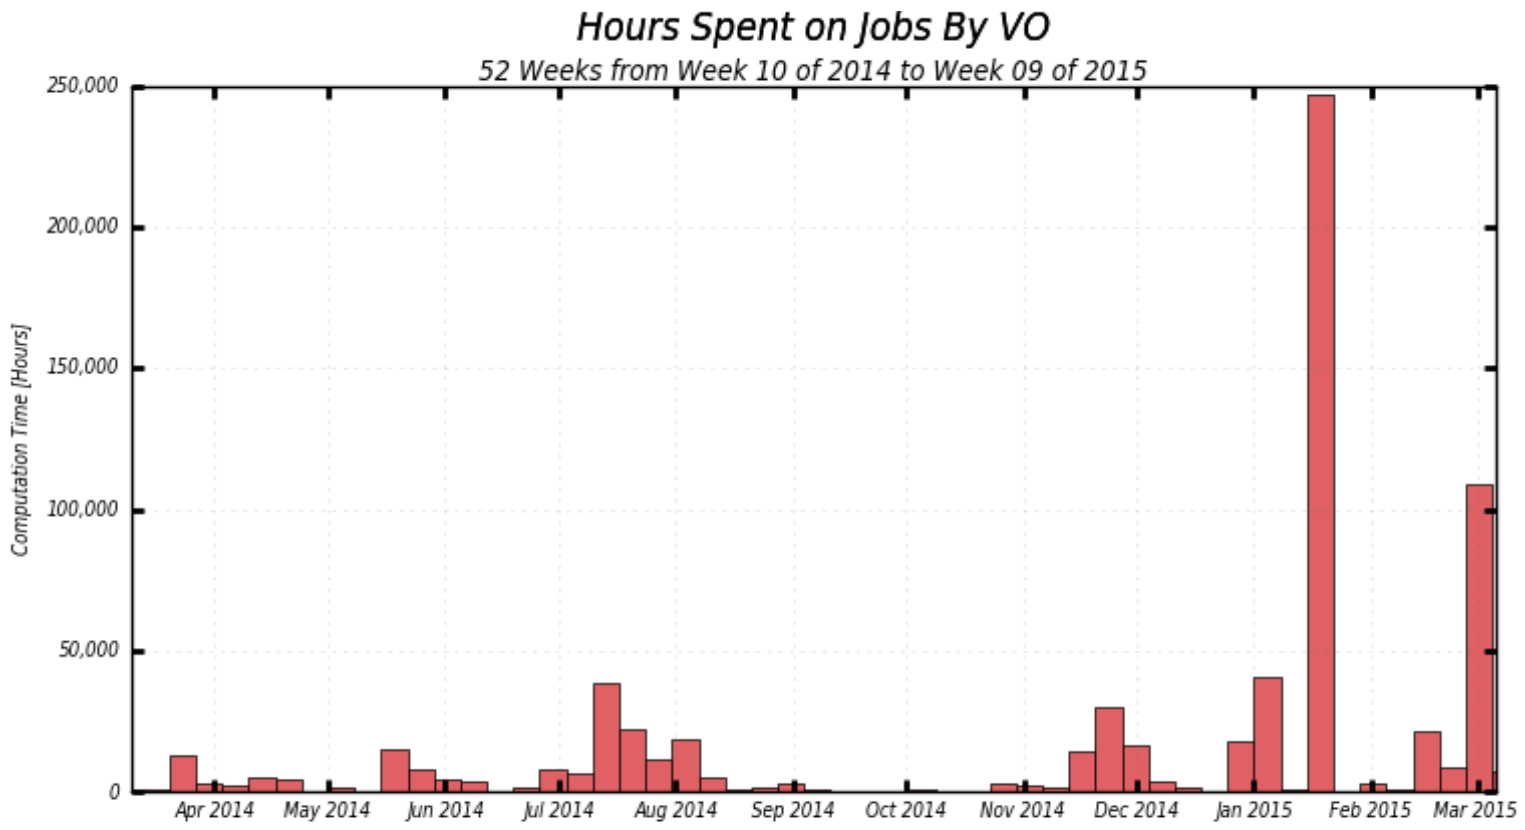


DarkSide-50 Inside Detector under final assembly
back - ups

DarkSide- 50
Liquid Scintillator Sphere inside the Water Tank


DarkSide-50 Water tank and liquid scintillator vessel

S1 signal shapes for

Electron event →

Nuclear recoil event ↓

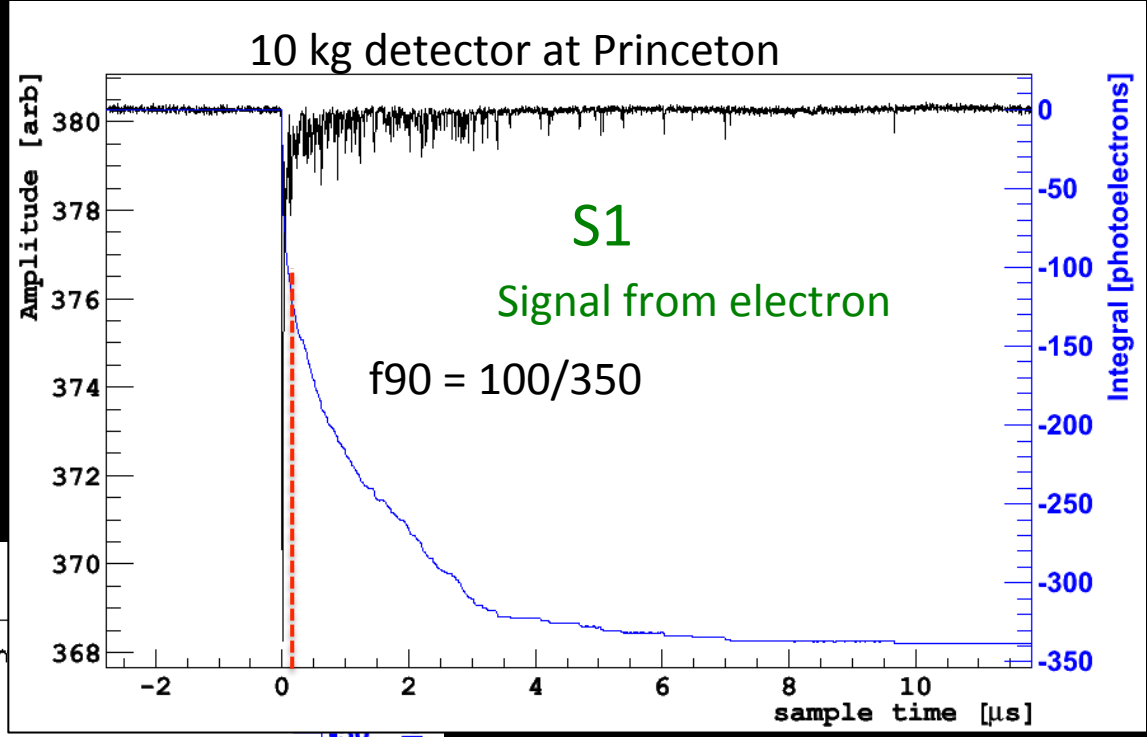

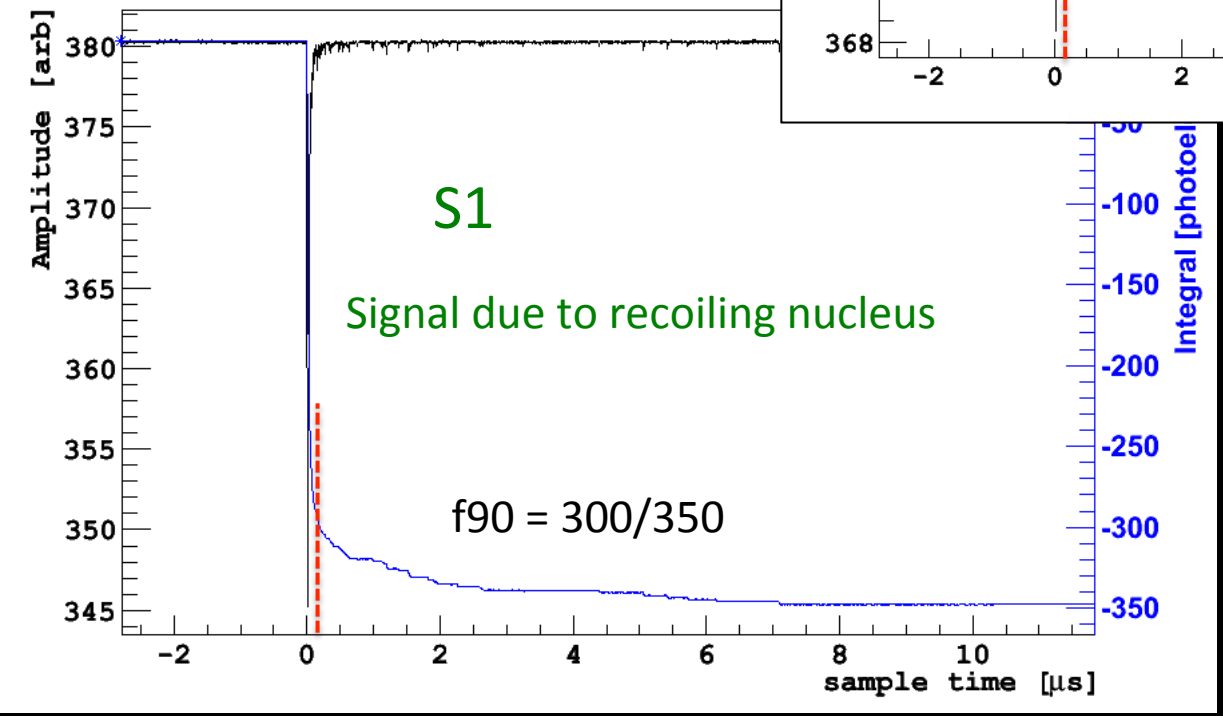

Define **f90** =

Fraction of integral in first 90 ns

S1 & S2 signals

Electron event →

Nuclear recoil event

↓

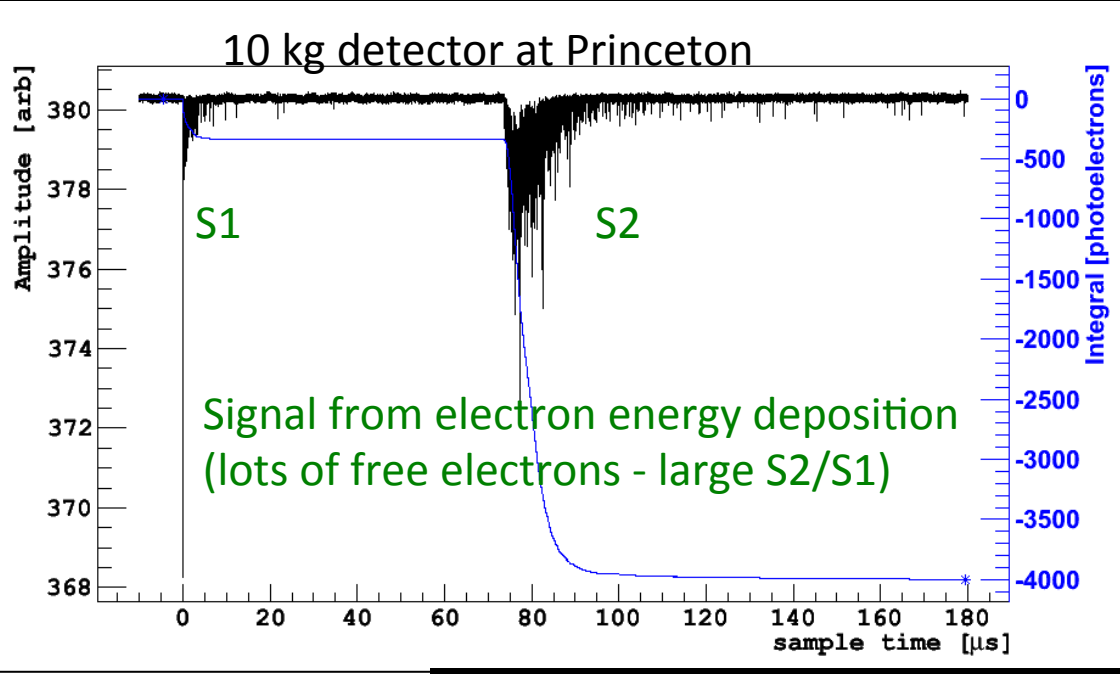

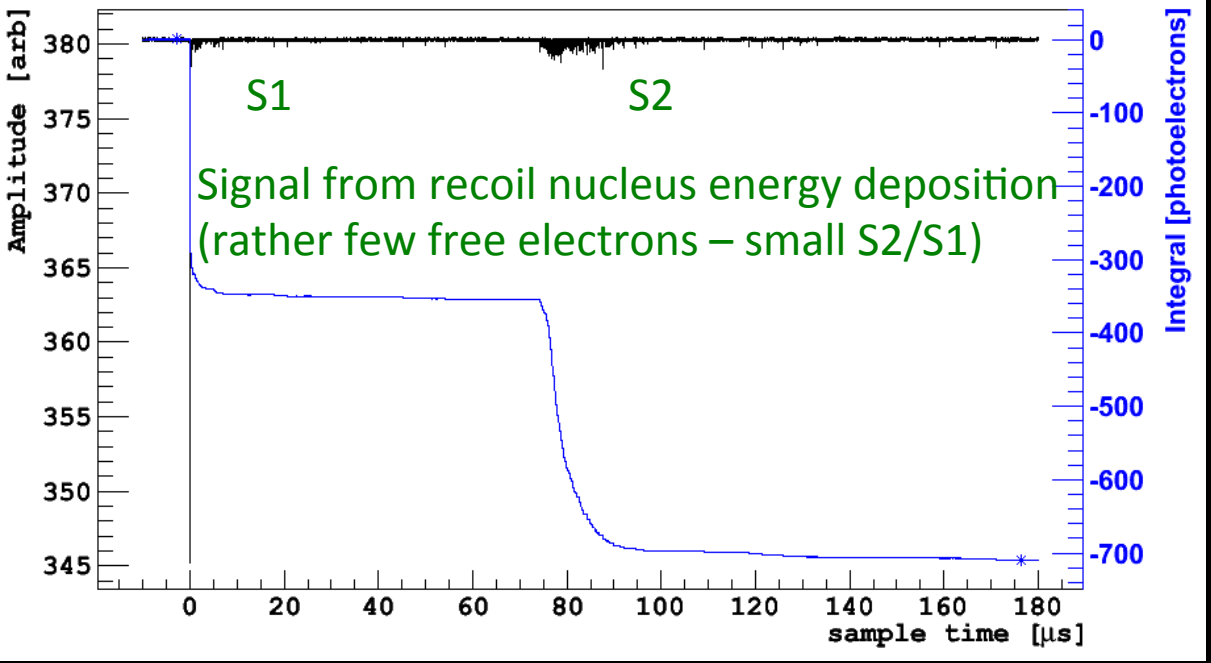

Use **log(S2/S1)**
as discriminant

Needs from Service Areas

Services	Offerings	DarkSide
DAQ and Controls	DAQ	In Production
Grid and Cloud	Fermigrid	In Production
Grid and Cloud	Fermicloud	
Grid and Cloud	Gratia	In Production
Grid and Cloud	Job Sub	In Production
Grid and Cloud	FIFEMON	In Production
Grid and Cloud	OSG Enabled	Testing
Grid and Cloud	Amazon	
high performance computing		
Scientific Data Storage and Access	dcache/ enstore	In Production
	Gridftp	In Production
Scientific Data Management	IFDH	In Production
Scientific Data Management	SAM Web	In Production
Scientific Data Management	FTS	In Production
Scientific Frameworks	Software Framework (art)	In Production
Scientific Frameworks	Continuous Integration - Orchestration	
physics and detector simulation	Genie	
physics and detector simulation	geant4	In Production
Scientific Software	Software Processing (LArSoft)	
Scientific Databases	IF Beam	
Scientific Databases	Custom Databases	
Databases	mysql	
Databases	postgres	YES
experiment specifc services	Production Operations	
networking		
Scientific Computing Systems	CVMFS	In Production
Scientific Computing Systems	Interactive machine in GPCF	In Production
Scientific Computing Systems	Experiment Control Rooms	NO
Scientific Computing Systems	Continuous integration - middleware	NO
Scientific Collaboration Tools	Redmine	In Production
Scientific Collaboration Tools	CVS/ Subversion/ Git	In Production
Scientific Collaboration Tools	Electronic Log	Custom
Scientific Collaboration Tools	ups/upd	In Production
Central Web Hosting		
Scientific Document Management	doc dB	In Production
Video Conferencing		
Futures	Federated Data Management	xrootd
Futures	High Throughput Analysis Facilities	

spare