protoDUNE-SP Data Quality Monitoring

Maxim Potekhin (BNL)

ProtoDUNE-SP Data Exploitation Readiness Review@FNAL
May 10th 2018

BROOKHAVEN NATIONAL LABORATORY

DUNE

Overview

- The focus of this talk is mainly on infrastructure implemented for the support of the Data Quality Monitoring (DQM) in protoDUNE-SP
- Motivations for DQM and prompt processing
- Requirements
- System design
- Interfaces
- Deployment and operation
- What we learned in the two Data Challenges
- Remaining work items

* more technical material can be found in the "Backup Slides" section

Motivations for DQM and prompt processing

- Goal: Provide actionable information to the shifters regarding detector performance within minutes (or perhaps tens of minutes) from the time the data is taken
- The Online Monitor has some of the more basic functionality similar to Data Quality Monitoring but some of the tasks are not compatible with its mode of operation
- Many experiments have "express streams" (also referred to as "nearline" or "prompt processing systems")

Online Monitoring vs Prompt Processing.

Online Monitor	DQM/Prompt Processing
Strong coupling to DAQ	No coupling to DAQ
Some fraction of full data rate	~1% of full data rate
Fixed/limited amount of CPU	Scalable CPU resources
Dedicated Hardware	Facility Hardware
DAQ network	Facility Network
Immediate (sec)	Prompt (min)
User access strictly controlled	More relaxed access for DUNE
Workflow Mgt: artDAQ	Graph-based DAG mgt
Software testing and updates tightly controlled	Software can be tested/updated at any time with no impact on data taking

protoDUNE-SP data flow

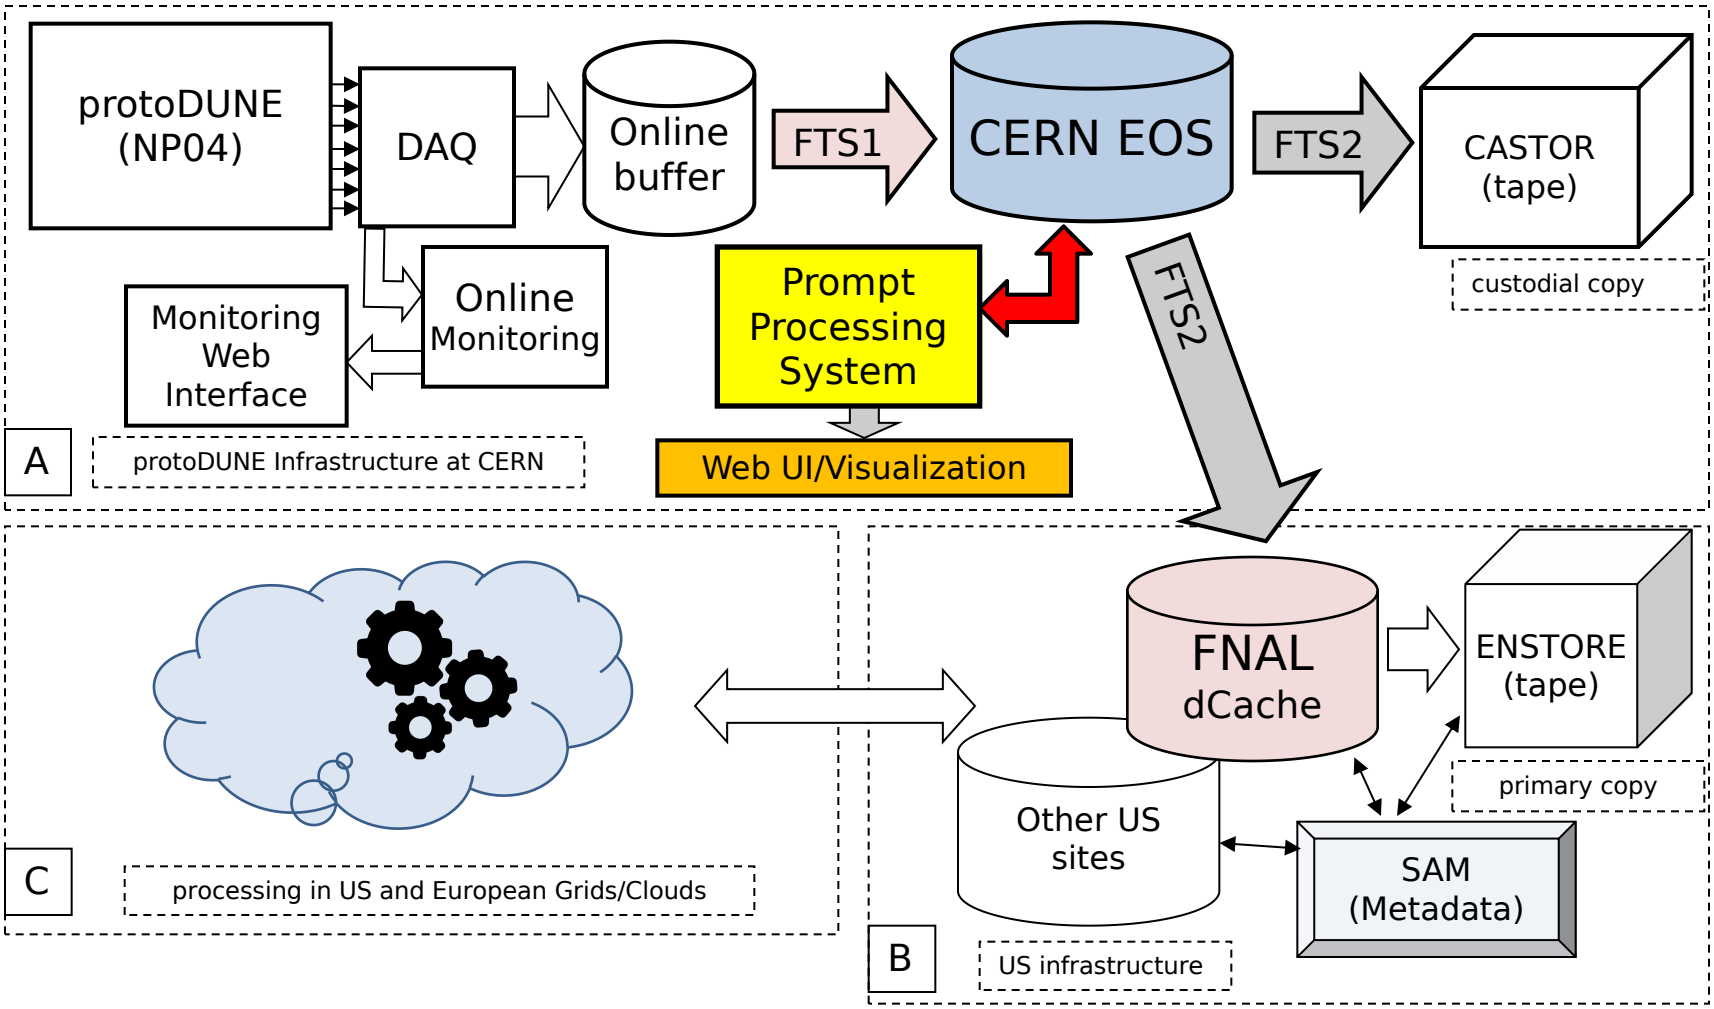

The protoDUNE-SP prompt processing system

- The **p**rotoDUNE-SP **p**rompt **p**rocessing **s**ystem (**p3s**) is needed to support DQM, running a variety of DQM payloads on a fraction of the data already recorded on disk, turnaround time of O(10min)
- Basic requirements for p3s
  - maximal simplicity of deployment and maintenance, resource flexibility
  - automation
  - monitoring capabilities to manage and track execution
  - efficient presentation layer for users' access to the DQM data products

p3s design

- ...see backup slides
- In a nutshell, it is a server-client architecture with HTTP communication between the components
- p3s is based on the concept of the "pilot framework"
  - minimizes the latency of job execution
- version control using git (GitHub)

p3s pilot framework (conceptual)

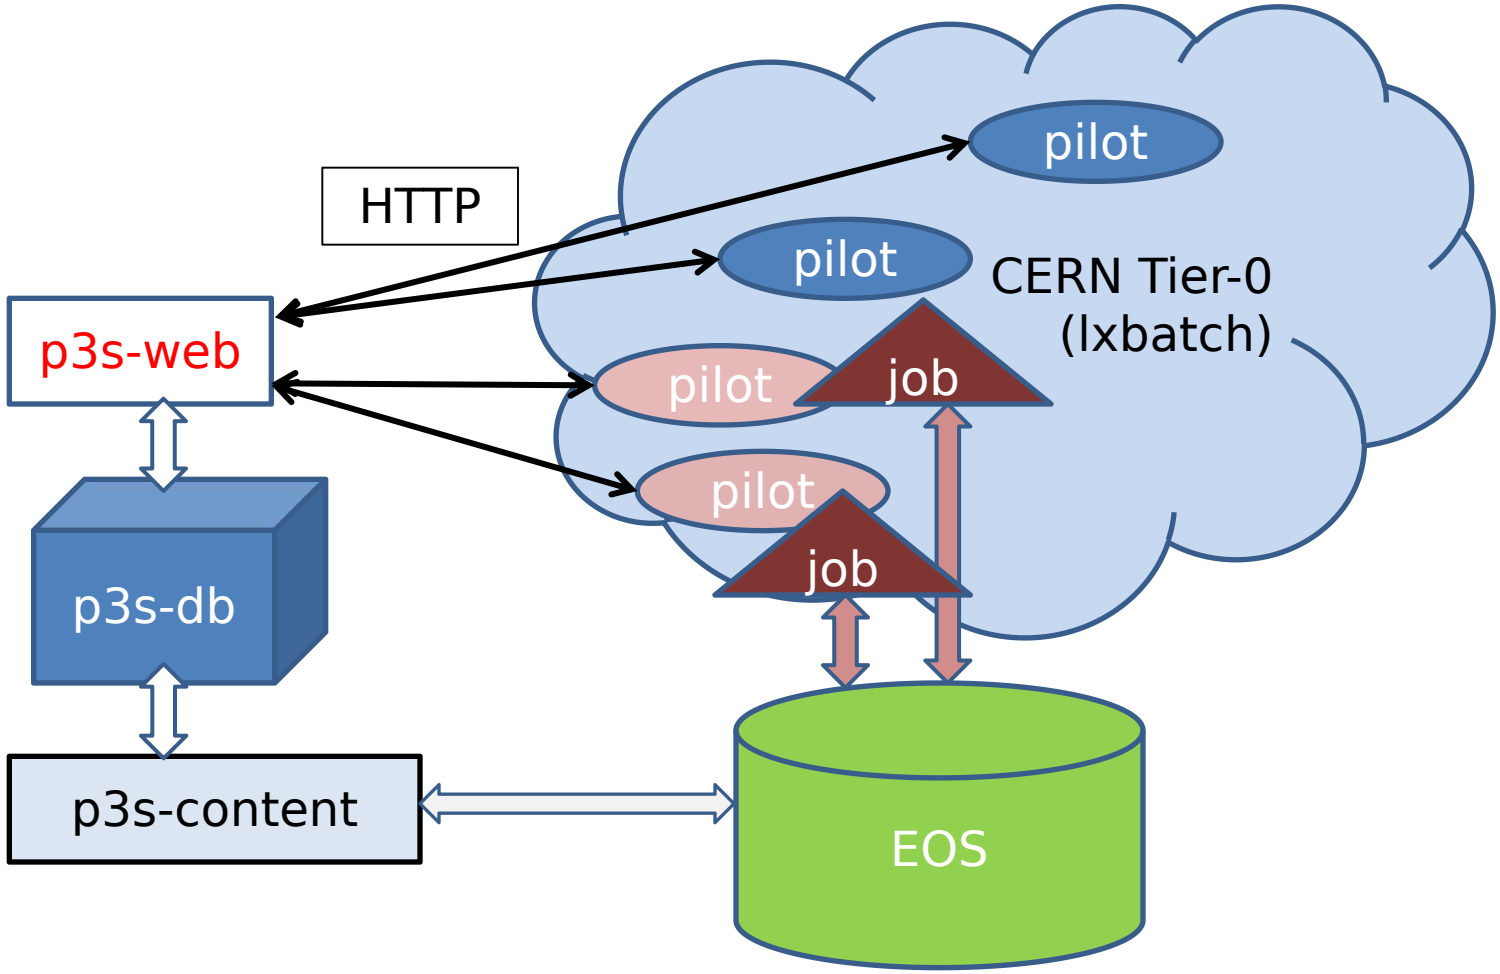

p3s Jobs and Workflows

- Jobs are submitted as records to the p3s database by *interactive or automated* clients
  - effectively a queue
- The state of each job is updated (e.g. from "defined" to "running" to "finished") under the management of a pilot, reported to the server
- Jobs are assigned UUIDs
- p3s supports DAG-type workflows

p3s: an example of Job Description


```
[
  {
    "name":     "EvDisp:Main",
    "timeout":  "1000",
    "jobtype":  "evdisp",
    "payload":  "/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/evdisp/evdisp_main.sh",
    "priority": "1",
    "state":    "defined",
    "env":      {
      "DUNETPCVER":"v06_69_00",
      "DUNETPCQUAL":"e15:prof",
      "P3S_NEVENTS":"5",
      "P3S_LAR_SETUP":"/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/lar_setup_2.sh",
      "P3S_FCL":"/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/evdisp/evdisp_current.fcl",
      "P3S_INPUT_DIR":"/eos/experiment/neutplatform/protodune/np04tier0/p3s/input/",
      "P3S_INPUT_FILE":"dummy_to_be_replaced",
      "P3S_OUTPUT_DIR":"/eos/experiment/neutplatform/protodune/np04tier0/p3s/output/",
      "P3S_EVDISP_DIR":"/eos/experiment/neutplatform/protodune/np04tier0/p3s/evdisp/",
      "P3S_USED_DIR":"/eos/experiment/neutplatform/protodune/np04tier0/p3s/used/",
      "P3S_OUTPUT_FILE":"evdisp.root"}
  }
]
```


Software version

Component reuse

- ...please see backup slides
- the idea is to leverage standard existing frameworks and packages and minimize own development

CPU

- Tested operation with 1000 concurrent jobs executed in p3s over a period of time (utilizing CERN lxbatch service)
- Need to balance available CERN resources to fit within DUNE allocation
- p3s ran with 300 pilots in Data Challenge 1 and with 600 pilots in Data Challenge 2 (to be adjusted once the payload software is finalized)

Hosting p3s services on VMs in CERN OpenStack

- p3s-web: the workload managment and monitoring server (Django+Apache)
- p3s-content: presentation service (Django+Apache)
- p3s-db: the database server (PostgreSQL)

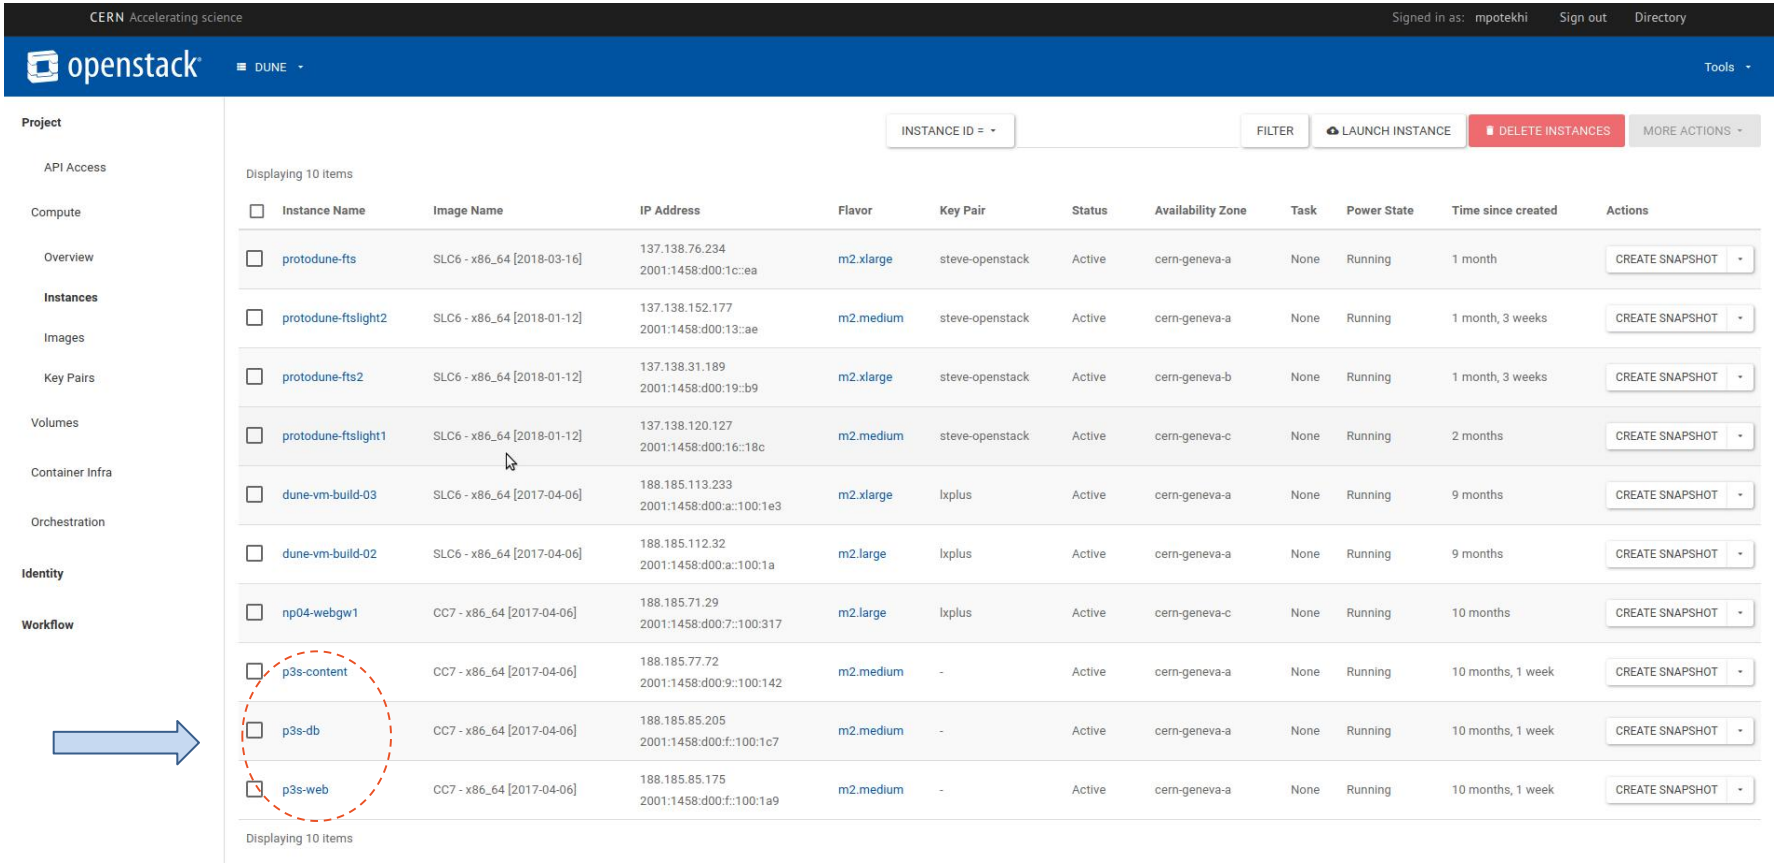

The p3s dashboard and the DQM section of the Grafana monitor

Neutrino PLATFORM

p3s home @"p3s-web.cern.ch" | DQM @"p3s-content" | 10:50:19 04/10/18 CET

Jobs | Pilots | Workflows | DAGs | Data | Data Types | Services | Sites

p3s running on server "p3s-web.cern.ch" at 10:50:19 04/10/18 CET

Summary

Object	Number
Pilots: total idle running stopped TO	400 122 274 0 0
Jobs: total defined running finished TO	739 2 274 463 0
Workflows: total	2
Datasets: total	528
Jobs finished as reported by pilots	438

Jobs

State	1min	10min	1hr	2hrs	24hrs
Defined	130	289	527	581	739
Started	131	287	525	579	737
Stopped	4	15	251	307	463
pilotTO	0	0	0	0	0

System Status

Attribute	Value
Uptime	83 days, 6:05:23.110000
Load	0.87 0.49 0.38 1/184 9667
Sites	lxvm
Data location	/eos/experiment/neutplatform/protodune/np04tier0/p3s

Users: gechrist,hschellm,jingbow,mspanu,mxp,np04dqm,tjunk

Pilot injection

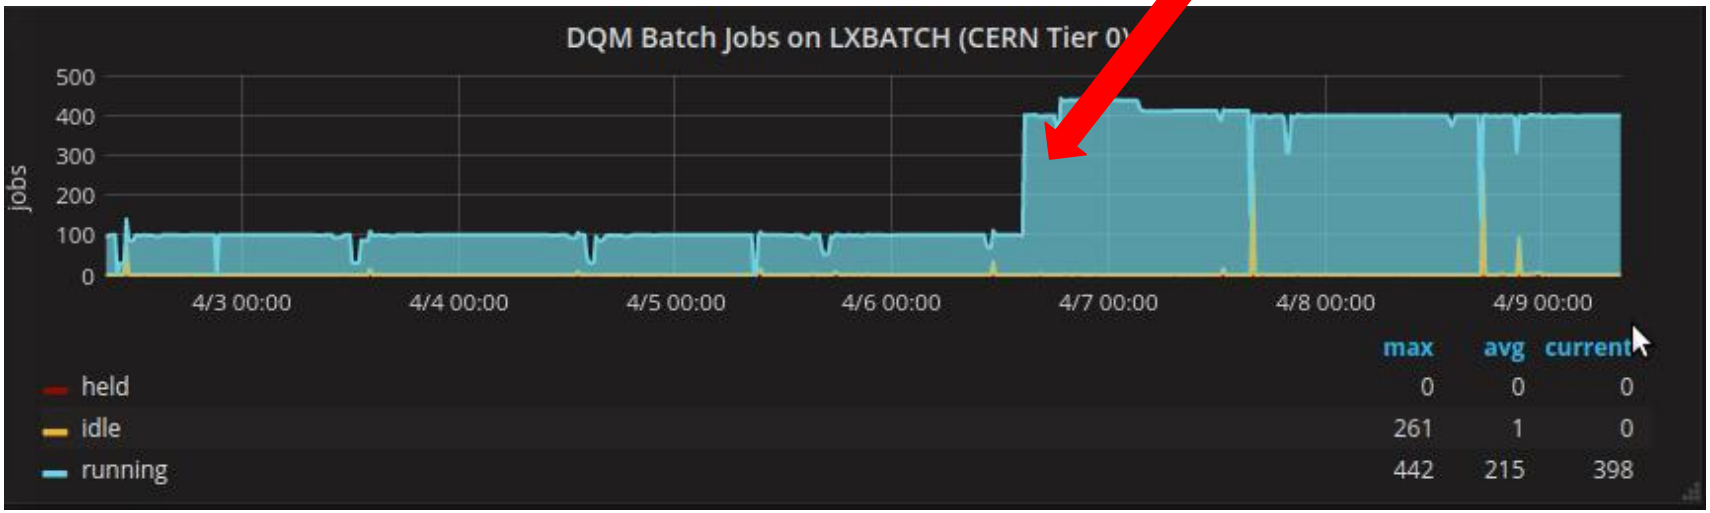

BROOKHAVEN NATIONAL LABORATORY DUNE

The p3s job monitoring page

DEEP UNDERGROUND NEUTRINO EXPERIMENT

p3s home @"p3s-web.cern.ch" | DQM @"p3s-content" | 05/07/18 21:30:18 CET

Jobs | Pilots | Workflows | DAGs | Data | Data Types | Services | Sites

job : Total in DB:14, selected:14 on server "p3s-web.cern.ch"

Job States:
- All
- Template
- Defined
- Running
- Finished
- Pilot Timed Out
- Time Limit

User: All **Type:** All **# per page:** 25 Submit

ID	Uuid	User	Site	Host	Name	Wfuuid	type	Payload	Params	Pri.	t.limit	State	defined	started	stopped	Pid	Errcode	Directive
30609	6a4baa72-4fad-11e8-bd99-fa163e69d62c	gechrist	lxvm	b62d79dd7e.cern.ch	Monitor:Main	—	monitor	/afs/cern.ch/work/g/gechrist/public/monitor_main.sh	—	1	1000	finished	20180504 17:11:26	20180504 17:11:26	20180504 17:14:47	76	0	—
30608	cb398984-4fa9-11e8-8e95-fa163e69d62c	gechrist	lxvm	b616c5b70d.cern.ch	Monitor:Main	—	monitor	/afs/cern.ch/work/g/gechrist/public/monitor_main.sh	—	1	1000	finished	20180504 16:45:31	20180504 16:45:31	20180504 16:50:32	76	1	—
30607	b927cbb4-4fa6-11e8-98ea-fa163e69d62c	gechrist	lxvm	b6e1da905a.cern.ch	Monitor:Main	—	monitor	/afs/cern.ch/work/g/gechrist/public/monitor_main.sh	—	1	1000	finished	20180504 16:23:32	20180504 16:23:33	20180504 16:30:19	76	1	—
30606	6a9fd2e8-4d80-11e8-aada-fa163e953a6b	mxp	lxvm	b67dec01e4.cern.ch	Monitor:Main	—	monitor	/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/monitor/monitor_main.sh	—	1	1000	finished	20180501 22:44:17	20180501 22:44:18	20180501 22:54:24	76	1	—
30605	4e7e2662-4d7d-11e8-8217-fa163e953a6b	mxp	lxvm	b62a6ccd95.cern.ch	Monitor:Main	—	monitor	/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/monitor/monitor_main.sh	—	1	1000	finished	20180501 22:22:01	20180501 22:22:02	20180501 22:30:25	76	0	—
30604	304c7cca-4d7d-11e8-bb15-fa163e953a6b	mxp	lxvm	b600350b92.cern.ch	Monitor:Main	—	monitor	/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/monitor/monitor_main.sh	—	1	1000	finished	20180501 22:21:11	20180501 22:21:12	20180501 22:32:55	76	1	—
30603	7d265aee-4d6d-11e8-9b59-fa163e953a6b	mxp	lxvm	b63b14eaa1.cern.ch	Monitor:Main	—	monitor	/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/monitor/monitor_main.sh	—	1	1000	finished	20180501 20:28:48	20180501 20:28:48	20180501 20:35:31	76	0	—
30602	6854d9dc-4d64-11e8-9c0f-fa163e953a6b	mxp	lxvm	b632a7bdb1.cern.ch	Monitor:Main	—	monitor	/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/monitor/monitor_main.sh	—	1	1000	finished	20180501 19:23:47	20180501 19:23:48	20180501 19:30:32	76	0	—
30601	0667d0c4-472b-11e8-9c7c-fa163e28fa79	np04dqm	lxvm	b65d289b81.cern.ch	HitDump	—	hitdumptest	/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/test/hitdumptest.sh	—	1	1000	finished	20180423 21:17:55	20180423 21:17:56	20180423 21:39:40	76	0	—
30600	0667c750-472b-11e8-9c7c-fa163e28fa79	np04dqm	lxvm	b6f2fc57ca.cern.ch	HitDump	—	hitdumptest	/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/test/hitdumptest.sh	—	1	1000	finished	20180423 21:17:55	20180423 21:17:55	20180423 21:44:42	76	0	—
30599	bdd93732-4728-11e8-b931-fa163e28fa79	np04dqm	lxvm	b60246ba3e.cern.ch	EvDisp:Main	—	evdisp	/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/evdisp/evdisp_main.sh	—	1	1000	finished	20180423 21:01:34	20180423 21:01:34	20180423 21:43:21	76	0	—
30598	6eb4dc26-4726-11e8-ab32-fa163e28fa79	np04dqm	lxvm	b6f80bf120.cern.ch	HitDump	—	hitdumptest	/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/test/hitdumptest.sh	—	1	1000	finished	20180423 20:45:02	20180423 20:45:05	20180423 21:13:32	76	0	—
30597	43711cb0-442d-11e8-8006-fa163eb16b15	np04dqm	lxvm	b6ee0d0f54.cern.ch	EvDisp:Main	—	evdisp	/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/evdisp/evdisp_main.sh	—	1	1000	finished	20180420 01:56:23	20180420 01:56:23	20180420 02:11:26	76	0	—
30596	78e05438-43d8-11e8-b825-fa163e2b45e6	mspanu	lxvm	b62590704d.cern.ch	HitDump	—	hitdumptest	/afs/cern.ch/user/n/np04dqm/public/p3s/scripts/test/hitdumptest.sh	—	1	1000	finished	20180419 15:49:25	20180419 15:49:25	20180419 15:56:08	76	0	—

Current DQM payloads

- "TPC Monitor" (includes the Photon Detector)
- Event Display + Data Preparation
- Purity Monitor
- BI Monitor (currently in a rough prototype stage)
- Currently all are LArSoft apps, this simplies the setup which is common

Notes:

- Software is provisioned to the worker nodes via CVMFS
- The list is not final and certain applications are in the works
- p3s is designed to make it easy for the operators to add new payload jobs and workflows is this becomes necessary during activation, commissioning and data taking
- High degree of compatibility between OM and DQM, some software has been successfully ported

Job detail in the p3s monitor

Neutrino PLATFORM

p3s home @"p3s-web.cern.ch" | DQM @"p3s-content" | 04/13/18 11:10:25

Jobs | Pilots | Workflows | DAGs | Data | Data Types | Services | Sites

job: 552728ba-3ee6-11e8-9cf3-fa163e4f89e6

Attribute	Value
priority	1
ts_sto	April 13, 2018, 9:01 a.m.
user	np04dqm
wfuuid	—
id	30594
env	{"P3S_OUTPUT_DIR": "/eos/experiment/neutplatform/protodune/np04tier0/p3s/output/", "P3S_INPUT_FILE": "mcc10_dc2_p5GeV_cosmics_3ms_sce_9_20180112T163817_merged0_65881_1523590783.root", "P3S_FCL": "/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/evdisp/evdisp_current.fcl", "DUNETPCQUAL": "e15:prof", "P3S_EVDISP_DIR": "/eos/experiment/neutplatform/protodune/np04tier0/p3s/evdisp/", "DUNETPCVER": "v06_69_00", "P3S_INPUT_DIR": "/eos/experiment/neutplatform/protodune/np04tier0/p3s/input/", "P3S_LAR_SETUP": "/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/lar_setup_2.sh", "P3S_USED_DIR": "/eos/experiment/neutplatform/protodune/np04tier0/p3s/used/", "P3S_NEVENTS": "5", "P3S_OUTPUT_FILE": "evdisp.root"}
ts_def	April 13, 2018, 8:46 a.m.
directive	—
ts_sta	April 13, 2018, 8:46 a.m.
payload	/afs/cern.ch/user/n/np04dqm/public/p3s/p3s/inputs/larsoft/evdisp/evdisp_main.sh
site	lxvm
params	—
ts_dis	April 13, 2018, 8:46 a.m.
p_uuid	fb832956-3e1f-11e8-8bd4-02163e00b02f
timelimit	1000
uuid	552728ba-3ee6-11e8-9cf3-fa163e4f89e6
errcode	0
host	b635daa0d5.cern.ch
jobtype	evdisp
pid	8523
name	EvDisp:Main
state	finished

BROOKHAVEN NATIONAL LABORATORY DUNE

DQM payload output on the "p3s-content" pages

DEEP UNDERGROUND NEUTRINO EXPERIMENT

Home Purity Event Display Images Event Display Prototype Interface

protoDUNE DQM: Purity Monitor

 # per page: 25 Submit

ID	Run	TPC	Ts	LifeTime	Error	Count
27947	4649	10	04/13/2018 1:46 p.m.	2.69	0.13	91
27946	4649	9	04/13/2018 1:46 p.m.	2.72	0.12	96
27945	4649	6	04/13/2018 1:46 p.m.	2.58	0.08	106
27944	4649	5	04/13/2018 1:46 p.m.	2.6	0.08	123
27943	4649	2	04/13/2018 1:46 p.m.	2.66	0.16	73
27942	4649	1	04/13/2018 1:46 p.m.	2.79	0.17	95
27941	4648	10	04/13/2018 1:20 p.m.	2.69	0.13	91
27940	4648	9	04/13/2018 1:20 p.m.	2.72	0.12	96
27939	4648	6	04/13/2018 1:20 p.m.	2.58	0.08	106
27938	4648	5	04/13/2018 1:20 p.m.	2.6	0.08	123
27937	4648	2	04/13/2018 1:20 p.m.	2.66	0.16	73
27936	4648	1	04/13/2018 1:20 p.m.	2.79	0.17	95
27935	4647	10	04/13/2018 11:54 a.m.	2.69	0.16	79
27934	4647	9	04/13/2018 11:54 a.m.	2.59	0.17	81
27933	4647	6	04/13/2018 11:54 a.m.	2.4	0.1	86
27932	4647	5	04/13/2018 11:54 a.m.	2.66	0.14	84
27931	4647	2	04/13/2018 11:54 a.m.	2.61	0.15	79
27930	4647	1	04/13/2018 11:54 a.m.	2.98	0.16	81
27929	4646	10	04/13/2018 11:51 a.m.	2.65	0.16	80
27928	4646	9	04/13/2018 11:51 a.m.	2.66	0.12	87
27927	4646	6	04/13/2018 11:51 a.m.	2.3	0.13	77
27926	4646	5	04/13/2018 11:51 a.m.	2.5	0.11	85
27925	4646	2	04/13/2018 11:51 a.m.	2.99	0.16	74
27924	4646	1	04/13/2018 11:51 a.m.	2.52	0.17	71
27923	4645	10	04/13/2018 11:46 a.m.	2.59	0.19	72

Next Page 1 of 637

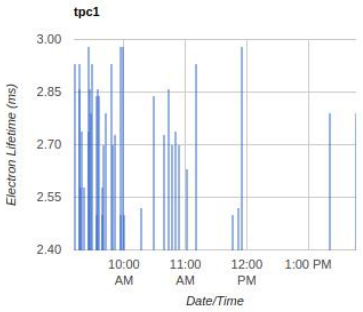

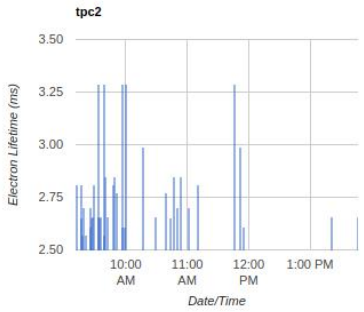

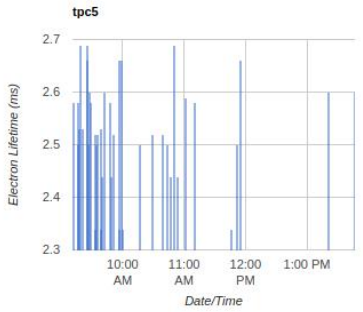

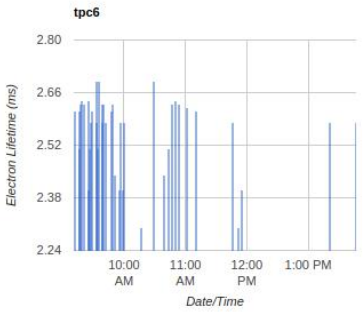

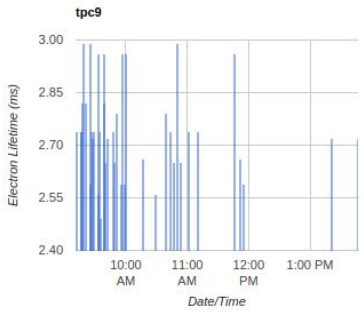

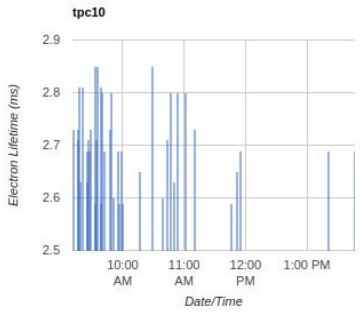

BROOKHAVEN NATIONAL LABORATORY DUNE

DQM Event Display + Data Preparation (a prototype)

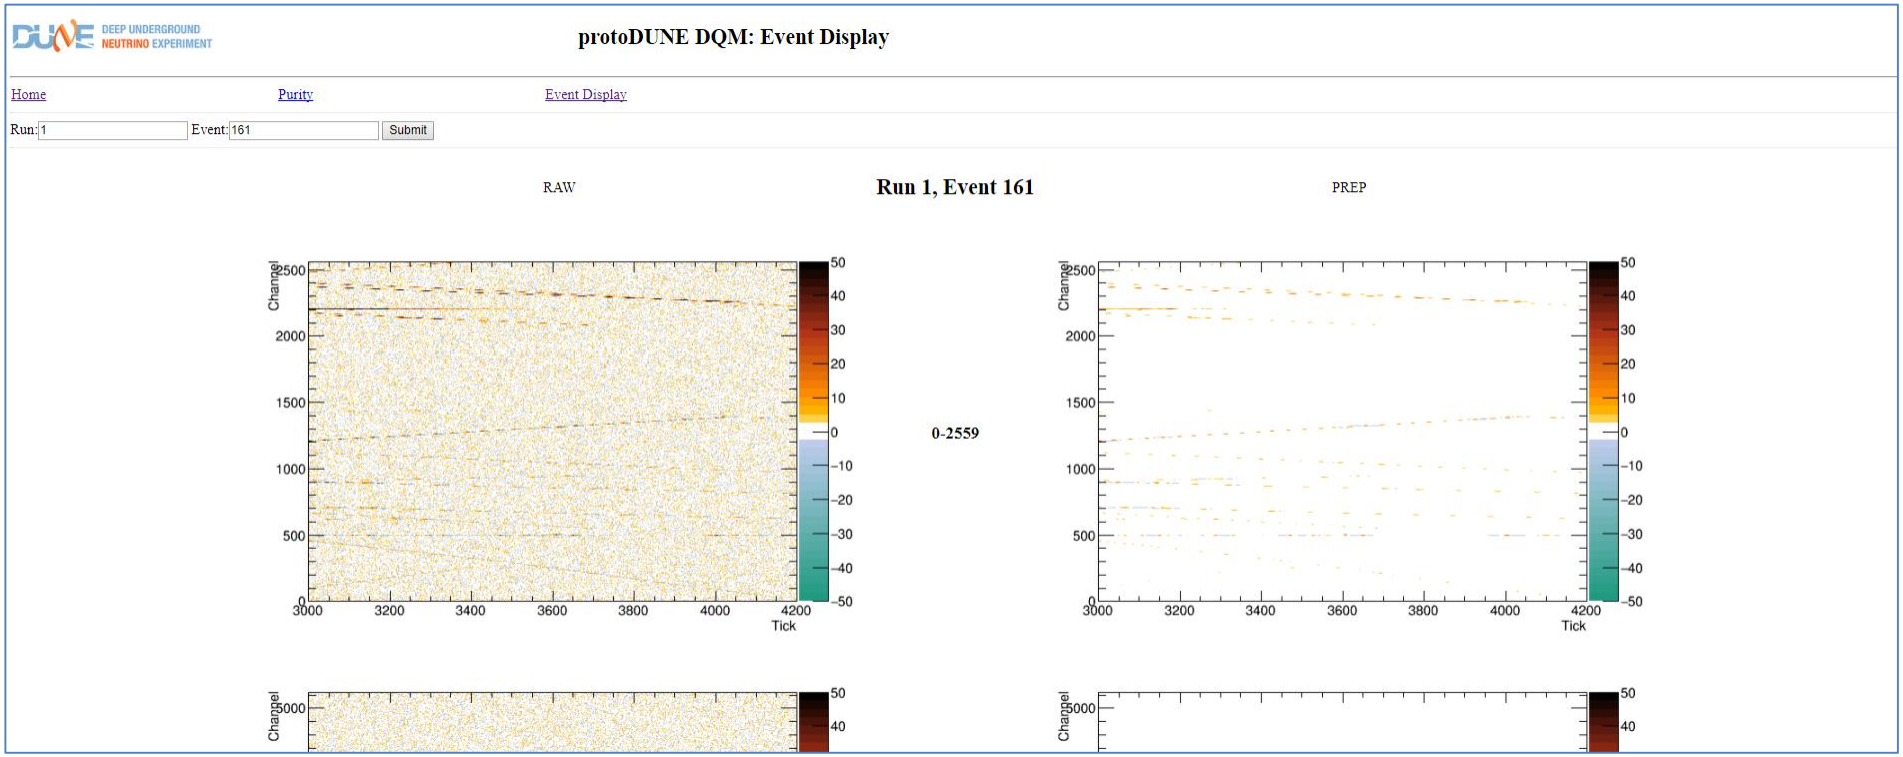

DQM "TPC Monitor" application (histograms produced in p3s, UI integration is work in progress)

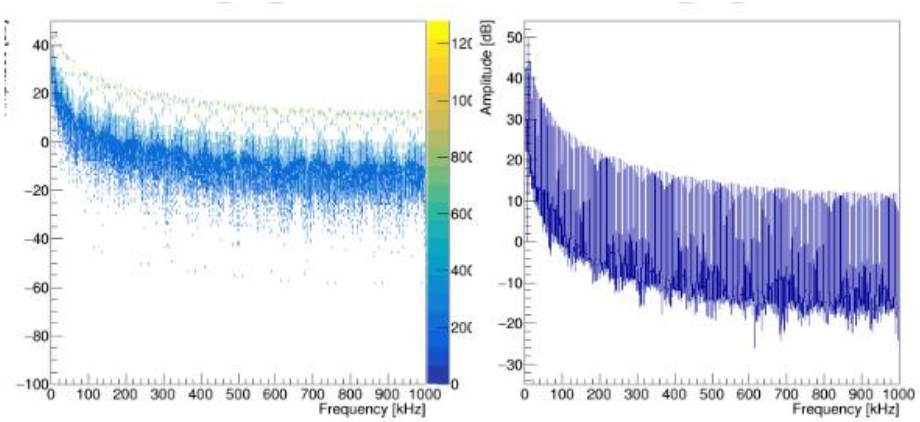

Persistent_FFT_Fiber#3

Profiled_FFT_Fiber#3

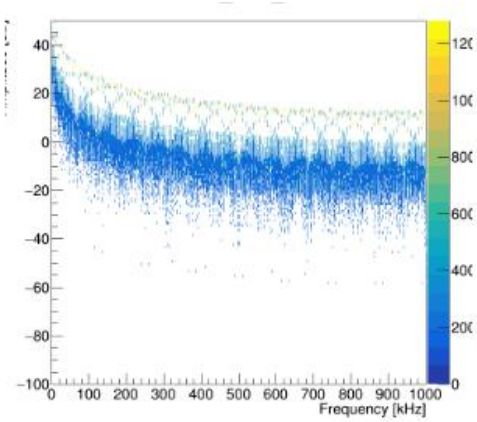

Persistent_FFT_Fiber#4

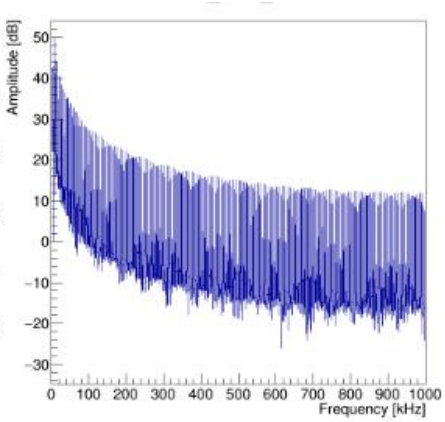

Profiled_FFT_Fiber#4

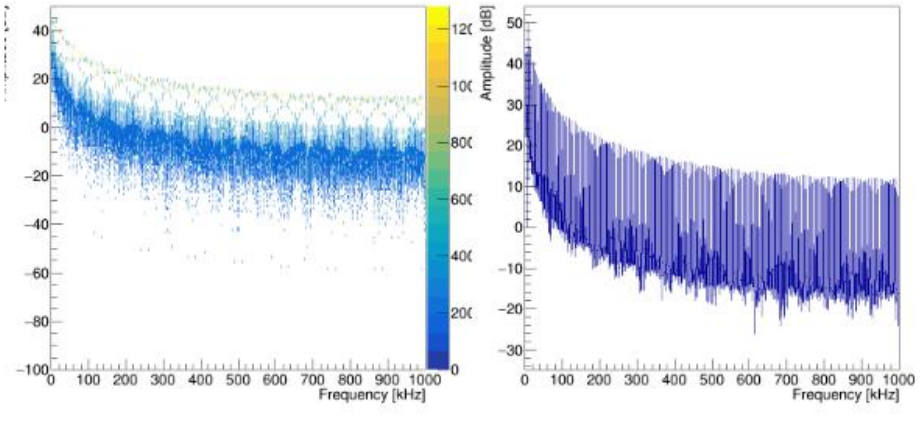

Persistent_FFT_Fiber#6

Profiled_FFT_Fiber#6

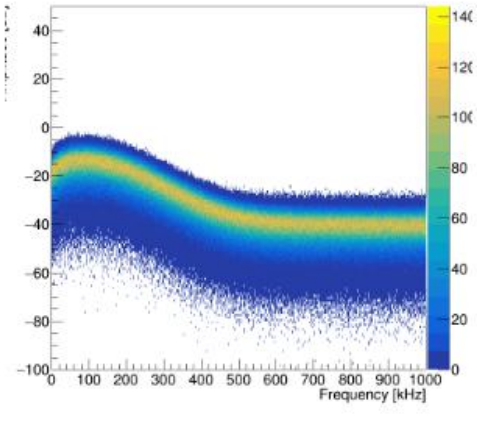

Persistent_FFT_Fiber#7

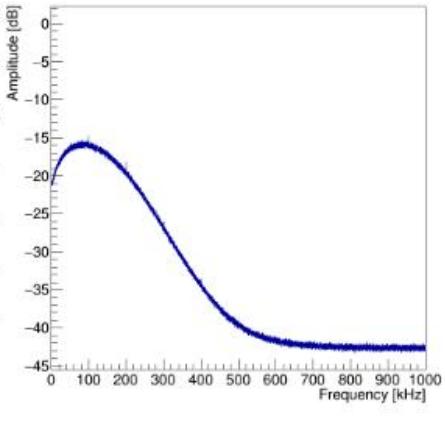

Profiled_FFT_Fiber#7

BROOKHAVEN NATIONAL LABORATORY

DUNE

Deployment

- Services on OpenStack: standard installation of Python, Django, Apache, PostgreSQL and a few packages
- Network configuration/firewall/SELinux
- Client software is ready to use for any DUNE member
- Storage
  - CERN EOS for I/O, with initial reliance on FUSE interface (a POSIX-like layer)
  - CERN AFS for local software deployment and HTCondor log files
- a designated "inbox" where a predefined portion of the data is copied by an instance of F-FTS
- one or more "outbox" folders for output data

Operation in 2017 - Spring'18

- Operating continuously for about a year with core services running in a stable manner, used to test DQM payloads
- A few types of cron jobs are active using the CERN distributed "acrontab" (services)
- Two data challenges were conducted in the past 6 months and they will be summarized in a separate report during this review

Services and the service log

- p3s persists reports from its services in a database (service log)
- helpful in finding errors and reporting them to CERN ITD e.g. HTCondor
- can add any service due to a simple API

DEEP UNDERGROUND NEUTRINO EXPERIMENT

p3s home @"p3s-web.cern.ch" DQM @"p3s-content" 05/08/18 01:09:46 CET

Jobs | Pilots | Workflows | DAGs | Data | Data Types | Services | Sites

service : Total in DB:948, selected:948 on server "p3s-web.cern.ch"

Service: All # per page: 25 Submit

ID	Name	Time	Info
3106	pcalc	May 8, 2018, 1:07 a.m.	Required pilots:100, detected:99, submit 1
3105	purge	May 8, 2018, 1:05 a.m.	purge object:"pilot" in state:"timeout" deletions:13
3104	pcalc	May 8, 2018, 12:57 a.m.	Required pilots:100, detected:87, submit 13
3103	TO	May 8, 2018, 12:51 a.m.	{"N": "13", "TO": {"143731.9": "90027", "143731.13": "89949", "143731.11": "89993", "143731.12": "89965", "143731.0": "89968", "143731.18": "89987", "143731.1": "89946", "143731.7": "89994", "143731.3": "90025", "143731.2": "89990", "143731.6": "89998", "143731.10": "89990", "143731.5": "89961"}}
3102	purge	May 7, 2018, 10:05 p.m.	purge object:"pilot" in state:"timeout" deletions:13
3101	TO	May 7, 2018, 10:01 p.m.	{"N": "5", "TO": {"143733.8": "78377", "143732.1": "14239", "143733.24": "78322", "143733.34": "16004", "143733.40": "78202"}}
3100	TO	May 7, 2018, 9:51 p.m.	{"N": "8", "TO": {"143733.41": "78126", "143732.23": "37878", "143732.25": "78619", "143733.6": "78102", "143733.42": "78094", "143733.32": "78136", "143732.10": "38107", "143732.11": "38119"}}
3099	purge	May 7, 2018, 9:46 p.m.	purge object:"pilot" in state:"timeout" deletions:3
3098	pcalc	May 7, 2018, 9:37 p.m.	8 idling condor jobs found, exiting
3097	TO	May 7, 2018, 9:31 p.m.	{"N": "3", "TO": {"143733.9": "13074", "143732.17": "13073", "143732.13": "13103"}}
3096	purge	May 7, 2018, 9:25 p.m.	purge object:"pilot" in state:"timeout" deletions:4
3095	TO	May 7, 2018, 9:21 p.m.	{"N": "4", "TO": {"143734.0": "75514", "143731.15": "13797", "143731.17": "77449", "143731.8": "77335"}}
3094	pcalc	May 7, 2018, 9:07 p.m.	error: SCHEDD_ADDRESS_FILE.
3093	pcalc	May 7, 2018, 8:57 p.m.	error: SCHEDD_ADDRESS_FILE.
3092	purge	May 7, 2018, 6:05 p.m.	purge object:"pilot" in state:"timeout" deletions:5
3091	TO	May 7, 2018, 5:51 p.m.	{"N": "5", "TO": {"143732.1": "23639", "143732.13": "23165", "143731.16": "64669", "143733.9": "63300", "143732.17": "23163"}}
3090	purge	May 7, 2018, 5:45 p.m.	purge object:"pilot" in state:"timeout" deletions:5
3089	TO	May 7, 2018, 5:41 p.m.	{"N": "4", "TO": {"143733.2": "62735", "143733.35": "62516", "143733.18": "63175", "143731.14": "63959"}}
3088	pcalc	May 7, 2018, 5:37 p.m.	error:
3087	TO	May 7, 2018, 5:31 p.m.	{"N": "1", "TO": {"143733.34": "62572"}}
3086	pcalc	May 7, 2018, 5:27 p.m.	5 idling condor jobs found, exiting
3085	purge	May 7, 2018, 5:25 p.m.	purge object:"pilot" in state:"timeout" deletions:2
3084	TO	May 7, 2018, 5:21 p.m.	{"N": "2", "TO": {"143731.15": "62969", "143731.4": "62953"}}
3083	pcalc	May 7, 2018, 4:57 p.m.	error:
3082	pcalc	May 7, 2018, 4:47 p.m.	error:

Next Page 1 of 38

Check the pilot lifetime

BROOKHAVEN NATIONAL LABORATORY DUNE

Data Challenges (DCs)

- The two data challenges took place in Nov. 2017 and Apr. 2018 with teams working at both CERN and FNAL, instrumental for us achieving readiness
- ...contained components for "keep up processing" (offline) and Data Quality Monitoring, which was running continuously consuming data delivered to it by F-FTS
- Utilized both MC data as well as real data from the Cold Box test

Infrastructure issues identified in Data Challenges

- DC1:
  - AFS timeouts
  - premature termination of pilots due to a bug in the HTCondor configuration (fixed!)
  - occasional slowness when using EOS FUSE CLI commands
- DC2:
  - a new bug in EOS (unreadable files), fixes by CERN experts are work in progress
  - increased failure rate with large files when writing files through EOS FUSE mount
  - occasional HTCondor "shadow exception" errors

- post-DC2:
  - HTCondor services non-reponsive for some period of time
  - ...due to general high load on the servers machines and misconfigured jobs

Mitigation

- EOS FUSE:
  - migrate away from FUSE wherever possible
  - stage the data using xrdcp (this does not exclude errors but still more stable)
  - harden the scripts to better handle errors
  - *...this currently is work in progress*
- HTCondor:
  - there is little that can be done in case of a genuine outage apart from escalation of the issue with the CERN ITD services
  - an additional alarm would be helpful to quicker identify these occurrences (currently detection is done by consulting the "service log DB" which is a part of p3s)

Error Detection/Recovery Procedures for shifters

- General:
  - look at the *service log* of p3s already to detect HTCondor submission problems
  - this is the spot that the shifters can routinely watch once an hour
  - some of EOS failures are less obvious and probes need to be added, reporting to the same service log
  - completed jobs record the error code in the job monitor DB, and if that is not zero than log files can be examined
  - all of the above requires training for shifters
  - CERN service ticket system is pretty good and we have a working relationship with CERN ITD, and a liaison
- EOS:
  - missing outputs are a clear sign, but require expertise
  - ...see the comment above
- HTCondor:
  - p3s job logs record a variety of runtime HTCondor failures
- Improvements:
  - create alarms

DQM payload development cycle and software deployment

- Software provisioning is done via CVMFS
- ...but we also require local builds and started testing at CERN
- It is assumed that individual DQM payload developers are responsible for curating DQM outputs (histograms, diagrams, tables), bug fixes, enhancements etc
  - Bruce Baller is the lead
- Participation of other DRA team members is highly useful
- p3s shifters are responsible for day to day operation, system monitoring and responding to alarms or anomalies

DQM/p3s operations

- In a steady state, a dedicated p3s shifter duty is not necessary, ~10% FTE availability is needed on a 24/7 basis during the operation
  - but we expect things to be hectic during commissioning so more effort will be required around August
- First p3s tutorial was held at CERN during DC1 and achieved its goals with participants running LArSoft jobs
- Documentation was subsequently improved
- We anticipate that we'll need at least 3 (and preferably 4-5 for redundancy) trained personnel to insure adequate coverage of p3s
- 1-2 weeks of hands-on experience (part-time) will likely be required to achieve proficiency
  - goal is to be able to reliably run and re-run both established and ad hoc payloads
  - local builds
  - add a few alarms, improve logs and write up the shift instructions
- In addition, once the presentation layer is finalized we'll arrange a tutorial for DRA/DQM experts to help them navigate DQM outputs

Further work items

- Tagging/Cataloging DQM output
- p3s-content Web interface additions and improvements
- Training shifters

Summary of the timeline and milestones

- Jan 2017: a p3s prototype operational
- Apr 2017: deployment on the Neutrino Platform Cluster
- Jun 2017: migration of services to OpenStack
- Aug 2017: prototype DQM payloads tested (Purity Monitor and Event Display)
- Nov 2017: DC1 with 3 types of DQM payloads, scalability test
- Dec 2017: documentation rewrite
- Jan 2018: migration of services to the production account
- Apr 2018: DC2 with Purity, Display, Monitor, BI(*)
- May 2018: Improvements in DQM scripts to mitigate infrastructure issues
- Jun 2018: Better presentation layer for the TPC monitor, Ev. Disp. etc
- Jun 2018: Addition of S/N
- Jul 2018: DC2.1(?)
- Jul 2018: automated transport of DQM outputs to FNAL
- Jul 2018: train shifters
- Aug 2018: BI integration
- Aug 2018: debugging and adjustments during the commissioning
- Sep 2018: operations, data taking

Backup Slides

Documentation

- User-level documentation created and maintained on GitHub:
  - https://github.com/DUNE/p3s/tree/master/documents
- Prior Documentation:
  - DocDB **1811:** "*Prompt Processing System Requirements*"
  - DocDB **1861:** The outline of the design of the *p*rotoDUNE *p*rompt *p*rocessing *s*ystem (**p3s**)

Setting up the LArSoft (duneTPC) environment

On an either interactive or batch node at CERN:

```
source /cvmfs/dune.opensciencegrid.org/products/dune/setup_dune.sh
setup dunetpc ${DUNETPCVER} -q ${DUNETPCQUAL}
```

Considerations for reuse of existing systems

- A few exsiting "express stream" systems were considered and found impractical to adopt because of high degree of their coupling to the respective experiment infrastructure
- Existing large scale Workload Management Systems are powerful but an apparent overkill and carry substantial deployment and maintenance costs
- On the other hand, a simpler assembly of scripts that could automate DQM functionality but won't afford the user an efficient UI for either monitoring of job execution or access to the DQM data products; keeping track of the state of objects w/o a database is problematic
- p3s fills the gap between these different domains
- we are leveraging the CERN distributed storage to streamline data handling, and straightforward interface to the Tier-0 batch system (HTCondor) to achieve overall simplicity of the design

p3s design

- Server-client architecture, with a few available clients performing various functions. The server is a Web service (HTTP interface).
- p3s is based on the concept of the "pilot framework"
  - the pilot is a client running on a WN and managing jobs
  - the pilot is an agent deployed to the worker nodes that orchestrates the execution of the payload jobs
  - approach successfully used in systems such as PanDA and Dirac
  - pilots can run on ad hoc clusters or large facilities such as CERN Tier-0, tested on both
- Once activated, the pilot job sends a request to the p3s server in order to be assign a job for execution
  - the server fetches a job from its queue and matches it to the pilot
  - the pilot's lifetime is substantially longer than the typical execution time of DQM payloads, so a single pilot will serially process a large number of DQM jobs before termination (reaching its time limit)
  - since it operates in a live batch slot, the time for job dispatch is extremely short which provides the necessary responsiveness of the system

p3s on GitHub: https://github.com/DUNE/p3s

protoDUNE prompt processing system (p3s)

1,151 commits | **1** branch | **0** releases | **1** contributor

Branch: **master** | New pull request | Find file | Clone or download

buddhasystem Add version 6_75 — Latest commit `ca683db` 7 days ago

clients	more debug	22 days ago
configuration	lower the number of pilots a this stage of testing	22 days ago
display	Kill the debug print	8 days ago
documents	updated the JOB document	2 months ago
images	A pretty thorough rework of the landing page, added and versioned images	6 months ago
inputs	Add version 6_75	7 days ago
promptproc	Lowercase in table	22 days ago
sandbox	Continue the cleanup and doc update	4 months ago
tools	add Tom's histogramming script	7 days ago
.gitattributes	Added option to the job client, tweaked the larsoft test	a year ago
.gitignore	Restoring migrations	6 months ago
README.md	Stubbed out the "clients" document	4 months ago

README.md

Component reuse

- The goal is to minimize the amount and complexity of the application code
- This is achieved by using industry-standard, proven components
  - Apache Web server
  - Django Web application framework and helper packages, extensive use of Django template mechanism
  - PostgreSQL for the database service
  - Standard JSON and XML parsers
  - *Google Charts* for dynamic graphics
- Standard HTCondor interface for automatic submission of p3s pilots to lxbatch (a service script on top of HTCondor CLI)
- Web UI: purposely minimalistic but functional
- For data movement the capabilities of F-FTS are being leveraged

p3s: workflow support

- p3s supports workflows described as DAGs
- a standard XML schema (GraphML - developed for graphs) is used, supported by third-party apps
- parsing of XML comes for free with NetworkX package (Python)
- workflows are created using prefab DAGs as templates
- both classes are persisted in the DB as lists of nodes and edges
- only basic testing done up to this point

```
    <data key="d103">noop</data>
    <data key="d104">0</data>
    <data key="d105">0</data>
  </node>

  <node id="filter">
    <data key="d101">/bin/grep -c TEST $INPTXT > $OUTTXT</data>
    <data key="d102">{"P3S_EXECMODE":"SHELL"}</data>
    <data key="d103">filter</data>
    <data key="d104">7777</data>
    <data key="d105">4</data>
  </node>

  <node id="NOOP2">
    <data key="d101"></data>
    <data key="d102">{}</data>
    <data key="d103">noop</data>
    <data key="d104">0</data>
    <data key="d105">0</data>
  </node>

  <!-- ++++++ DATA ++++++ -->
  <edge source="NOOP1" target="filter">
    <data key="d1">input.txt</data>
    <data key="d2">/home/maxim/p3sdata/</data>
    <data key="d3">TXT</data>
    <data key="d4">INPTXT</data>
    <data key="d5">Simulated input</data>
  </edge>

  <edge source="filter" target="NOOP2">
    <data key="d1">output.txt</data>
    <data key="d2">/home/maxim/p3sdata/</data>
    <data key="d3">TXT</data>
    <data key="d4">OUTTXT</data>
maxim@serenity:~/projects/p3s/inputs$
```