

p3s/DQM 2018 summary and update

Maxim Potekhin

DUNE S&C Meeting
01/11/2019

Quick summary of 2018

- According to the TPC experts, the DQM system proved extremely useful
- Operation of the p3s and DQM services (module the infrastructure with which they interacted) has largely been smooth since the start of commissioning in summer of 2018 and continuing into now, with one serious outage happening due to an unexpected automatic software update on the VM (which was promptly fixed)
- We did experience periodic problems with two elements of the CERN infrastructure - EOS and HTCondor - which for the most part were successfully resolved on a reasonable time scale via the service ticket system
 - failures are be expected in an environment of such scale and complexity
 - but there were also genuine hard failures which took a while to resolve
- There was a good engagement with the DRA team and other experts, which allowed timely deployment of payloads, software updates and packaging
- Our design which calls for self-describing DQM output data proved very successful, with commissioning and updates of software happening with little or no changes to the server code or operation
 - a JSON file is required from each DQM job which describes the data products and their type, along with optional textual information that's propagated to the UI

The platform

- p3s and DQM were hosted on two separate VMs in the CERN OpenStack cloud
- boths were backed by a single PostgreSQL RDBMS also running on a VM
- this environment proved reliable amd required little intervention for 1.5 years since it became active
- configuration of the Apache servers is stored in the DUNE repo so the setup is reproducible by an operator with a medium level of expertise
- access to services for users outside of CERN was via ssh tunnels which worked OK for most people (and it certainly worked well for me all the time)

A few numbers

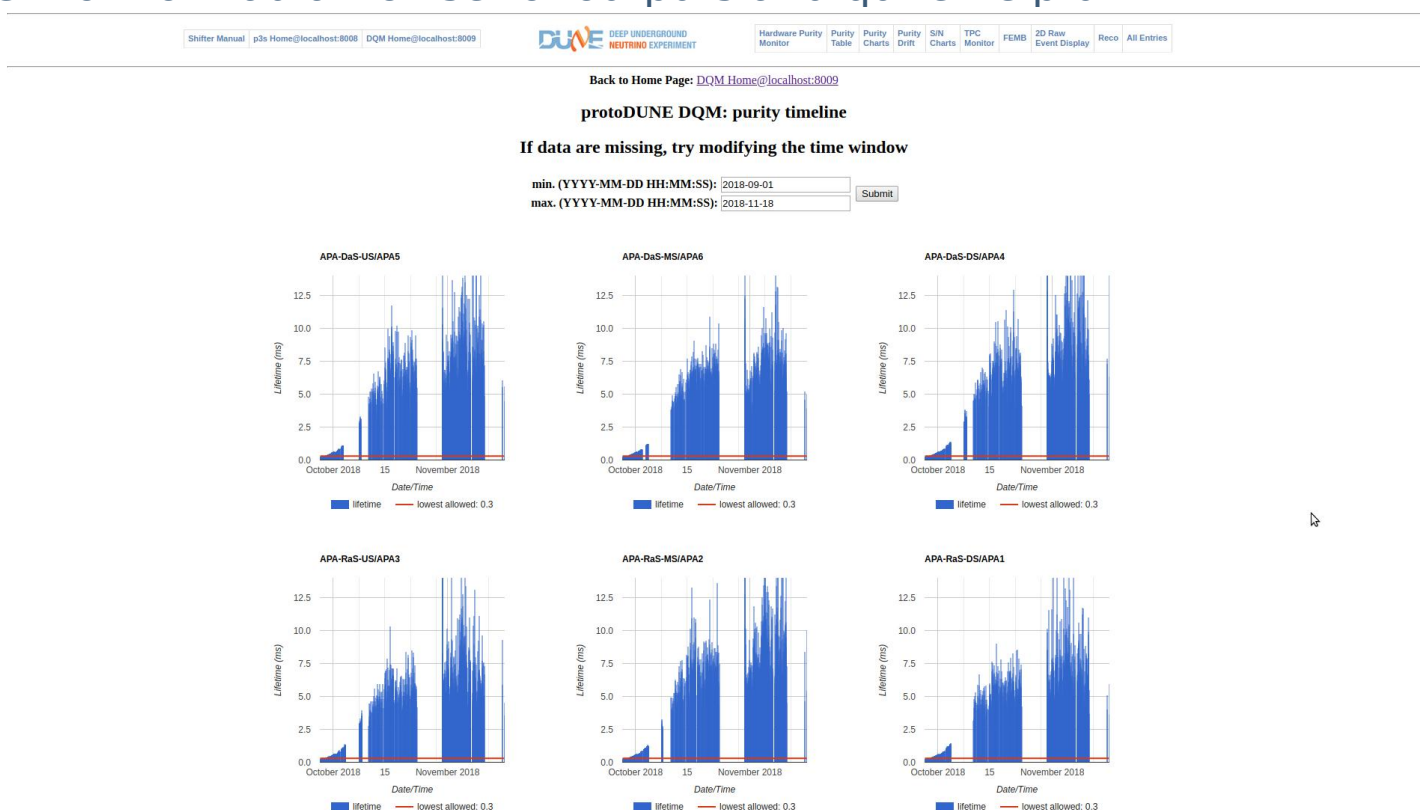
- A total of ~45k jobs ran from the start of commissioning, with about 32k during the beam time
- p3s was receiving between 10 and 20% of the total raw data stream at various times during the run
- Maintenance and orchestration jobs which required periodic execution were managed using "acrontab" at CERN (kerberized cron)
- Scan for new data was performed every 15 minutes i.e. fresh batches of jobs would be created 4 times an hour
- Depending on the type of the job the execution time was anywhere between 10 min and a few hours
- p3s jobs produced 10^5 files totaling about 2TB
- We typically ran ~100 jobs (or less) at a time, which was below our initial estimates

Output data

- The DQM data produced in 2018 proved useful in troubleshooting, there has been a request to keep it available via the DQM service for an extended period of time
- best solution would be to migrate the data and clone the service at FNAL
- ...at the same sourcing 2TB+ of disk to be attached to an Apache server proved non-trivial
- Separately, we always had plans to transfer the DQM data products to FNAL and commit it to tape etc, but never got around to do that
- ...so for now it continues to occupy a fairly trivial fraction of our EOS allocation at CERN and is transparently available through the DQM service
- The very large number of directories and files made the POSIX/FUSE access to the data and its management increasingly problematic - we didn't plan for this to become a semi-permanent storage so made no provision to divide the data in manageable chunks
- ... since then, the server code has been updated and sharding of the data has started, with no interruption of service

Payloads

- The 2D event display (D.Adams) was the first dramatic result we got from DQM
- ...but purity graphs and tables were instrumental in the commissioning of the detector
- the TPC monitor had a rich set of outputs and quite helpful



Outlook for 2019

- The older input data has been removed from EOS to make space for new data
- Log directories has all been cleared (condor, job, pilot)
- Currently job submission is disabled until a green light is received from the DRA and DAQ teams (e.g. there were pulser runs recently which has little or no DQM value, so we must not record them)
- The types of payload irrelevant to the cosmic ray data taking have been removed from the UI to make it more neat (apa2, crt, beam etc)
- New types of payloads can be added. If interested, please send a request. Compliance with the self-describing data will be appreciated
- Anticipate gradual updates of the code for better maintainability. Will consider containerization of the services for quick deployment and/or recovery. PostgreSQL replication?