

Overview Session

DUNE and Computing in the 2026 Era (Globalization, Schellman)

- Sign up to the Computing Consortia: talk to Heidi.
- Computing CDR timeline: in a year from now, maybe less.
- How to deal with the increased rate by a supernovae?: It will generate around 30-140TB per supernovae. DAQ people would like to deal with that 'in place' with local storage provided by double buffer and then wonder how to ship that Big chunk of data.

DUNE and Computing in the 2026 Era (DUNE neutrino computing future, Norman)

- Actual detector data from NOVA is changing the landscape in the octant fits.
- Trigger on supernovae: two types of trigger, internal trigger by the DAQ and external trigger by a SNEWS alert. The former is on the fly, the later has to relay in a ~1.5 min alert so we have to be able to read back this time (actually ~1min). Optical input from telescope observation would be too late.
- What would be the Δt ($t_0 + \Delta t$ readout window)?: depending on the model, 1 sec for the burst plus 20-30s.
- Plans for a third level trigger?: Not as default.
- Hit dropping?: Store the hit information and drop the raw data. Not done in actual neutrino experiments (they keep the hit basic information) maybe not suitable for DUNE due to the size of the data to store.
- Is the amount of resources (GPUs) needed for CVN for training or for reconstruction?: You need them for both.

Overview: Support Models

FNAL Computing Support (Margaret Votava, Elizabeth Sexton-Kennedy)

Q. What is auxiliary data handling?

- Files are small and repeatable, not code, flux files from generators
- In LHCb use CVMFS (e.g. conditions database for LHCb)
- Typically files 10/100M each low sharing rate amongst jobs
- CMS pileup handling stream using xrootd

Q. Big picture question - how going to work through merging infrastructures where people have developed a lot of effort?

- Various tools have various strengths: ATLAS has scaled twice as much as CMS, CMS strength is complexity that it can handle, DIRAC used more than more experiment (Belle2, LHCb) and has pluggable design
- On LHC time to develop through LS2, pause after ProtoDUNE
- WMS discussion not covered by HSF

Q. Conversation dominated by collider experiments. How do we get DUNE into those discussions?

- Funding bodies not just considering big experiments (ATLAS, CMS), involve rest of community - that is what the HSF is all about
- *Andy McNab*: Suggest using WLCG structure - WLCG gives access to distributed site
- Next meeting in March (*when? where?*) is combined HSF, WLCG and OSG

Not captured - neutrino workshop week after thanksgiving at CERN (?)

UK Support for DUNE (David Britton)

Q. Where does Workflow management fall (in the diagram on slide 9)?

- Trend (from HSF) to try and extract tools into common community tools (e.g DIRAC) which could then be pushed into operations layer (C)

Q. Comment on slide 13

- DUNE resources (slide 13) supplied manually by Andrew and Heidi
- *Pete Clarke*: Useful to get semi-serious ongoing estimate of computing and storage resource numbers in order to apply funding. We have quarterly resource meeting in advance, ask what you need in next 6 months and look ahead
- *Pete Clarke*: If you say what resources you want in 2020 does not mean have to use them, can be deallocated with ongoing process

Q. Elizabeth Sexton-Kennedy (LSK): Chronically under funding area of application software development. In US motivation to move to more common tools to create free energy to cover job "A", In the UK is that possible?

- Not really, make efficiencies in "B" or "C" does need more resources in "A" depending on funding patterns/ringfences
- *Pete Clarke*: Short answer - fixed funding line for PP, increased understanding in STFC that experiments but need full computing infrastructure and support
- *David Britton*: Can't build another GridPP per project so opportunity to support new communities using same model
- (UK funding discussion redacted!)

Q. How does LSST fit in to this model?

- *Bob Mann*: Using IRIS, enough overlap with other computing models and just starting hands on experiment using the Grid. Not an LHC experiment don't have the expertise and not obvious have all the effort to exploit all infrastructure
- IRIS includes SKA, Eucild and LSST
- *David Britton*: Effort on experiment side is the recurring problem - wasted effort from GridPP e.g one RA as experiment liaison gets another job
- Level A, B not been funded (slide 9)

- *Steve Jones*: GridPP have made case study of LSST collaboration:
<https://www.gridpp.ac.uk/users/case-studies/lst/>

Heidi: should we get a calendar of deadlines for the major contributors. Sure would help me keep track.

Data Management

DUNE Data Management Overview (Robert Illingworth)

Site usage slide

- Time for running out space in order 3-4 days
- *Daniela Bauer*: Just written in tender for more storage at IC

Site connection matrix

- RAL statement on <5GB issue: *"I was able to convince the CERN Davix developers that not being able to transfer files over 5GB was a critical problem (for SKA and DUNE). They have now implemented a version of Davix that supports this (<https://its.cern.ch/jira/browse/DMC-1090>) which took 5 days to write. This is now being tested and should allow Echo S3 to be properly used by DUNE."*
- Edinburgh/QMUL only have xrootd and get "checksums not being supported" errors

Was too late for Rucio to be used on ProtoDUNE, have two separate systems

Q. Heidi: How will SAM replica catalogue and RUCIO be integrated

- Fully replace SAM replica catalogue with RUCIO. Steps described in slide 18.

Q. Heidi: If RUCIO is going to replace SAM, will other experiments in FermiLab do the same?

- Yes but probably only do the phase 1 (slide 18).

Q. Heidi: If using Rucio as a replacement do you need to get 4-5 FNAL experiments on board and work with partners?

- Two phases - nova stay on phase 1 on hybrid system (?)

Q. Massimo: What is the the difference in the CERN and FNAL FTS instance?

- FNAL FTS - DUNE's data upload system written for getting Novas raw data into SAM. Both near and far detectors cataloged into SAM, then feature creep as deployed to offline systems
- Plan in replace FNAL FTS that something talks to Rucio
- Using CERN FTS already
- Using Rucio all transfers are done for FTS3
- Something else still has to do the metadata part

Q. *George Beckett: Is the development work being fed back into Rucio development?*

- Yes already started feeding back changes and fixed
- CMS applying new features
- Aim to have one common Rucio core

Q. *Why have you committed to Rucio?*

- Obvious SAM going to need improvements to be useful in long term
- Going to do disruptive may as well try different product, and at same time CMS also considering Rucio seems liked a good plan to consider common solution
- DUNE are not tied to CMS decision

Rucio Overview (Cedric Serfon)

Q. *George Beckett: What is "scope"? A user community?*

- From ATLAS has one scope per project (e.g data16,data17) and one for each single user
- Different scope helps partition DB
- Can choose to ignore the scope

Q. *Raja: What is the smallest unit you work with? e.g. LHCb run job on group of files how would this translate to Rucio?*

- Dataset - group all files into dataset that it your unit
- Yes can extract single file
- Units of measure: file, dataset and container

Q. *Raja: From sysadmin perspective how do you manage namespaces in Rucio? Can you tell that "/atlas/blah1/test" can be deleted*

- Deletion done automatically from Rucio when file is not needed agent not needed at site as this is all done centrally

Q. *Raja: Tools to work with dark data?*

- Yes some mechanism to compare site data dump and file catalog
- "site consistency check" based on xroot

Q. *George Beckett: File catalog scaling for Exabytes of data?*

- Partitioned namespace by project and one file catalog per scope
- Everything is hosted at CERN, but choice of backend (Oracle, Postgres)
- At CERN using Oracle (which we rely on for HA), have horizontal scaling

Q. *If DUNE using Rucio which well integrates with Panda does it follow that DUNE will be using Panda?*

- Do we provide new workflow management system?

- *LSK*: tried to partner with Panda but failed to get off the ground
- CMS adopting Rucio but not using Panda
- *LSK*: Can Rucio can be integrated with DIRAC?
- *Andy McNab*: depends how exposed to API already, if started from scratch easier (??)
- *Heidi*: DUNE much simpler better suited for smaller experiments. Look at lightweight implementation of Rucio/Panda - does not have to be running in 2 years time (more like 6 years). Get experience of LHC, how would do if starting from scratch
- *Massimo*: In Run 3 LHCb will have the same kind of data complexity as ATLAS
- *LSK*: Argue that DUNE is not going to be simpler. From Andrew's talk complexity is more than Run1 at the LHC. e.g machine learning at scale (which we have not done well as the LHC yet). ensemble of applications is harder to do
- *Andy McNab*: Additional requirements more than a matching problem? Or complicated application problem? How far could you go with Rucio+Condor stack - does it run out of steam? Or do you have to write a lot of DUNE specific code in house?
- *LSK*: Application layer all taking advantage of new tools. Pushing condor devs to higher levels of the stack (e.g. CondorCE)

Data Management: Outlook and Discussion

Data Management Evolution for WLCG (Sam Skipsey)

Weakness in SRM was that the overhead was too great. Experiments tried to work around it. Model became less rigid. LFC disliked. Drivers of change; slow deprecations of technologies (e.g. x509, gridftp, srm...) Hence DOMA.

Liz Sexton Kennedy. Where does Panda fit in this?

Sam: Rucio knows where the data is.

LSK: people were upset on x509. Quicker or faster.

Sam: Both. Various trade-offs between x509 (security) tokens (freer, but not safe enough).

Massimo: After run 1, we thought srm must expire. Run 2, and it is still there. Decision made to scrap it then.

Sam: Could not move until FTS3. Something like SRM required for tape.

Massimo: For ATLAS - only queries for quota use SRM.

Sam: Other VOs need more

Massimo: It's a general trend.

Sam: lcg-utils still in use.

Data Caching (Teng Li)

Steven Timm: StashCash is similar to xCache.

Mrs. Sexton-Kennedy/Teng: Uses CVMFS for namespace.

Mrs. Sexton-Kennedy/Teng: Uses xroot protocol. Could use http protocol.

Federation and Data Lakes (Rob Currie)

Chris Brew: For a definition of data lakes, CMS has all parts; AAA Dynamic Data Placements. Hierarchical within and out of UK.

Rob Currie: Yes, but not VO agnostic. Good example.

Raja: LHCb is also location agnostic.

Eric: what is missing? Data is still maintained at a particular place. Datalakes should be more amorphous.

Heidi: Imagine DUNE and MicroBoone. They could negotiate together. Small experiment might want to dynamically share space. Is Rucio able to deal with dynamic allocations like that? In Fermilab, space is shared, complex but more possibilities.

Liz Sexton-Kennedy: CMS, no such thing. When sharing, we still have quotas, except for cache. The orchestration needs to be right.

Heidi: This maybe a use-case that might come up.

Liz Sexton-Kennedy: We need to balance between amounts of resource, since we don't want to thrash the cache.

Margaret to Liz: How will DUNE storage converge?

Liz: Stash Cache (ATLAS/CMS) underneath xroot anyway. People involved in OSG (Brian) also in WLCG. We hope they converge. But nobody wants to change in mid run. Fermilab may need some help to converge from WLCG due to costs, politics etc. Authorisation could be a hitch.

Heidi: DUNE may wish to join in DOMA.

Sam: Brian Bockelman is common point of connection. He did all the SciToken stuff.

Distributed Computing

What does DUNE need (Steven Timm)

Non-negligible rate of cosmics even 1 mile underground. Rely on it for calibrations and so record as much as possible.

Slide 11 : Mostly from experience in OSG

Liz : Are you talking about the problems we do not know about?

Steve : Problem is shortage of manpower to find out the problems we have and their solutions

Pete : Some of the problem is because it is hugely problematic to commission a “machine” like the HPC machines which can provide some of the needed resources. Could we force the machines to conform to how we want to run them?

Steve : Defer to Liz

Liz : US Government wants all the different facilities to be unique. We have been working with some of the sites to configure them so that they are usable by us. E.g. Argonne and the I/O for the machines there. We also do not know the architecture of the machines that are going to be procured in three years.

Pete : Maybe we can get a list of our requirements to the sites. Maybe the architecture will not be important then?

Liz : And then the file systems also change.

Steve T : Once work is done to onboard one experiment it can easily be leveraged for others

Steve Jones : Good for tea time discussion ...

DUNE Production infrastructure (Ken Herner)

Ken : Focus on here and now and what improvements we can make now.

Slide 5 : Fermilab dominates processing partly because sites do not run jobs for more than 24 - 36 hours. Talking more about it later

Slide 6 :

- 10 - 15minutes per event
- 40 - a few hundred events per file
- At 100 events =~ 1500 minutes which is already >1day for processing the file.
- At 3.5GB RAM and 36 hours, good for >~99% of jobs. Submit recovery jobs for the rest.

Discussion :

Liz : If the data federation is set up in Europe, will we still need to send data across the atlantic?

Ken : Depends on what we actually need to do

Liz : There are a number of steps in the production job. Is there a merge that needs to be done at any point.

Ken : Not necessarily. But depends on what is actually needed.

Liz : Final file size aim?

Ken : 1-8 GB

Raja : Can we see what is going on from the point of view of jobs?

Ken : Yes. Go to the POMS monitoring.

UK Tier 1 (Rob Appleyard)

No request as yet from DUNE for tape from RAL.

Discussion

Steven Timm : Is the stock kernel stable for CEPH?

Rob : We are running stock kernel only (quite new though)

Andy McNab: Certificate thing is a little scary. Means that other sites should install an RPM to work with a site that is not CERN

Rob : Details (with a lot of detail) on back slide.

Rob Illingworth : Made the trusting no longer mandatory for transfers, but used if possible. Possibly a problem on the dCache end.

UK Grid T2 Site Perspective (Steve Jones)

Discussion

Steven Timm: A big thanks to all the very helpful Tier-2 support people.

Heidi : Should try to include the Tier-2 sites that are not part of DUNE in the consortium so that they are also recognised.

Pete : A huge thanks for all of the Tier-2 people who have fixed problems without asking for funding. Cannot invite people to join the consortium without giving them money. Everything is GridPP - who are the only people who will pledge money anyway.

Heidi : US model is slightly different with institutes also having a little more money to work with.

Daniela Bauer From Slack : @Ste Jones @Margaret Votava These kinds of jobs would have to go on IRIS resources once we get them. I prepared an LZ proposal to route their "high memory" (~12 GB/core) to run on IRIS clouds once they come up. The target is April 2019. So if a GridPP supported (that's the important bit here) VO has specific requirements that would be the route to go to try and get them. The person to ask about this is Jeremy Yates and it would have to go to the IRIS DB (I think that stands for "delivery board", but it's just the IRIS term for management).

Distributed Computing II

Global Monitoring (Tanya Levshina)

From dept responsible for glide-in-glite.

Q&A:

- To access these resources you need a Fermilab account. Computing providers need an authentication model which allows for external monitoring.
- LHC had similar problems originally. Are there plans to address this?
- Monitoring contains user information. A policy decision needs to be made here.
- What's the point if admins can't see it?
- Only need a fermilab account.
- Lightweight accounts like at CERN should work.
- Can add trusted people who don't have a fermilab account.
- GOCDB has similar data, but DUNE not in GOCDB
- Request: How does DUNE join GOCDB, need 2 EU states to sponsor the experiment.

Coffee machine starts...

Coffee machine ends...

- Getting non-DUNE members to use this needs documentation somewhere. This needs to be developed.
- Remote access for computing resources is different to site access which should require simple access
- Can the access be granted from GOCDB.
- All of the people are here, everyone here has access to the monitoring stuff.
- Security people at fermilab may stand in the way of this.

GLIDEinWMS (Steven Timm)

Condor predates the grid, may outlive it.

NB: Similar to pilot model

Slide9:

- CRIC? acronym from CMS. ATLAS version has been adopted by WLCG.
- Site config still manually configured not relying on things such as BDII

Q&A:

How is security done without proxied?

- There are proxies but not user generated

How many simultaneous jobs on grid are expected?

- 100k CMS at any time, intensity frontier still not sure. ~2-8k

Is this a DUNE requirement or what is running?

How many jobs will DUNE run on the grid?

- 56hr 1million jobs from similar experiment on HPC
- MC calculations would take months on the grid. This was scaled so that it was quicker on HPC.
- Datarate is similar to protoDUNE
- This takes 5-10k cores.
- This excludes supernova. Supernova may suddenly spin up a huge number of jobs which need to be processed quickly. (low latency)

Global Auth. (Tanya Levshina)

Slide 12:

- Is this in production?
- Working at some sites but depends on the resource as to what authentication is needed.

Q&A:

- Quick observation. Fermilab account is home institution when considering federated framework.
- Some home institutes can't/won't offer proper auth so users rely on falling back to CERN/Fermilab for authentication

This is working for eduroam?

- IT services see the point on eduroam which is not domain specific.
- A physics/science authentication system is not used by everyone so IT services won't want to support it

As other sciences need big compute will perceptions of this change?

- This is still a long distance away in time
- UKRI is supposed to improve this in the UK. This unfortunately is detached from researchers in other departments.
- There are also good units with bad depts and vice versa which makes federated authentication difficult.
- The Diamond light sources have something like this which allows users to move around, maybe they have something useful.

Is this just a single-sign-in?

- Yes and No.

Proxies and accounts separated in LHC jobs, is there a problem with compromised proxies?

- No these are all automatic short-lived proxies so no risk to killing proxy if there are problems.

HEP Cloud (Steven Timm)

Q&A:

no comments

Workload Management

DUNE Requirements (Heidi Schellman)

Q: What's the common requirements of DUNE computing regarding hardware? To what extent does DUNE fit in the box we would have (for example with OpenMPI). What will DUNE do regarding hit finding, tracking with machine learning?

A: DUNE will do hit finding with machine learning but currently not tracking. There're some limits for hit finding now due to the scale.

CMS experiences (Eric Vaandering)

Q: Regarding recovery, how to decide whether to recover the job?

A: The job monitor monitors the jobs to see for example if they are running out of CPU time and memory and if yes switch the parameters and resubmit the jobs. Currently it's done manually (looking at the plots based on the experiences). We are looking into machine learning solutions.

LHCb experiences (Andrew McNab)

Q: What are other experiments do with DIRAC?

A: For production, many small experiments, like Belle2, CTA... See more in Daniela's talk

Q: Are these experiments running their own instances? How would they do the hotfix?

A: For some of them yes. Like XXXDirac with their own specific extensions. They could submit pull request to fix bugs (mostly be released in the next release)

Q: Can you give me an example of LHCb specific extensions of Dirac?

A: The production system is LHCb specific.

Q: Do you have to use WMS to use DIRAC?

A: No. WMS use the pilot system underneath. Experiments can have their own implementation.

Q: Do people frequently use the web-portal or is there a CLI?

A: People use CLI more.

DIRAC WMS (Daniela Bauer)

Q: re priority. How would user decide the priority of their jobs?

A: We wish this could be improved. This is a requested feature of experiments

Q: What's the time scale/manpower grant of the DIRAC integration with RUCIO?

A: 1-2 years funded by IRIS. But could be longer if this is going to be used by multiple VOs.

Q: Could DUNE use the DIRAC WMS in Imperial? Will there be compatible issues?

A: Yes but we don't have enough manpower to do DUNE specific developments. One problem regarding compatibility is the metadata catalogue does not support multiple VO at the moment, that's also what we've requested.

Slide 7: In Ganga, users could specify the backend to submit jobs. That's how to submit DIRAC jobs via Ganga.

VAC/Vcycle (Andrew McNab)

Q: Can the containers run in the openstack instances?

A: Yes if you want to.

DUNE tried to use the CERN openstack, but since it charges DUNE too much as it prevents DUNE from sharing resources with other VOs. From 2016-2017 DUNE moved to HTcondor which does very similar thing and things get improved very much. Vcycle probably provide a solution as VMs of different VOs could use the same tenancy.

Kubernetes handles containers while openstack handles VMs.

FNAL Intensity Frontier (Robert Illingworth)

Q: The current solution of managing DUNE user jobs?

A: Currently simply running the users' submission scripts.

Suggest to use Ganga that supports various backends and is able to track the jobs.

Collaboration and Wrapup

LSK chairing

Common Solutions

- Similar solutions Margaret's chart from Day 1 - DUNE can select components for DUNE needs from WLCG options
- Trend from single experiment context to multi-experiment context pushing into infrastructure layer (e.g. Rucio)
- AMC: But not being moved into infrastructure, moved into collaboration between experiments
- LSK: Who is being paid to maintain generic DIRAC?
- AMC; No effort mostly from experiments (LHCb, ILC, BelleII)
- LSK: Want this to happen, but not been funded. Now need manpower to achieve this. Push experiment work into GridPP layer then frees experiment to work on experiment problems
- Fermilab have to do offload common solutions

- ST: But offload to OSG and then OSG drops the project what when?
- PC: need to find low hanging fruit, does not happen by edict. Data transport, data lakes getting to convergence. WMS a bit separate
- LSK: Why? smaller experiments using DIRAC where does scale enter
- PC: Funding flat, should it be pushed to GridPP layer, IRIS is paying for multiple roles (maintenance, development)
- MV: How are WLCG going to succeed on common solutions? (e.g Panda, DIRAC across the board)

Workload Management

- LSK: WLCG/OSF workshop combining forces. When wrote community whitepaper covered lot of topic but could not tackle workload management
- DIRAC, GlideINWMS, POMS multi-VO capable. Why do some succeed and some don't? Developers or development model - open source project - do they take contributions from other places (compare with BigPanda)
- KH: DIRAC and RUCIO started by one experiment, good job of making it general and then got taken up. GlideinWMS generic from the start. Some of it luck rather than dev model

AuthN

- HS: DUNE use several methods - consider authentication. Different authorities (US, Europe)
- LSK: There is WLCG AI team there could be a single solution
- MV: Do US CMS people do you need a CERN account?
- EV: But what kind of CERN account? Full or lightweight. Issues with Tier-2 people getting credentials and access to documentation
- LSK: Chronic problem, too much of barrier for relative award. CERN credentials SSO gets access to lot of material transparently - more convenient to use the CERN cert (as opposed to OSG)

Workload revisited

- LSK: HEPcloud science gateway for FNAL. If just want service don't care if use Glidein internal
- EV: pulls from global pool, this would have to be DIRAC
- DB: VAC sites do pull
- HS: DIRAC used by one experiment, does not mean all experiment use it
- AM: Inject DIRAC pilots into condor pool for Glidein to run. [details not captured]..it is possible to do..it is just how you provision the pilots
- LSK: could use interoperability job from UK over to US.

Solution choices

- PC: Assess what you need and what is available, cannot make a decision early on based on assumption.

- LSK: Does DUNE want to commission an assessment?
- HS: Yes, diversity of sites in UK mean DUNE has to be nimble to run under several configurations. DUNE has strong European contingent, incentive to work on convergence
- LSK: first target interoperability,
- AM: How far up the stack do you push the federation later. e.g. decision to go to DIRAC or HEPcloud
- ST: HEPcloud provisioned 8 UK sites in a month
- AM: Cannot have bespoke instructions depending on location
- AN: When experiments are smaller have more cross-pollination of solutions across experiments
- AM: WLCG ops meeting superset EGI, OSG? Are they different?
- LSK: In CHEP 2006(?) aim of federation of grid
- PC: GridPP will provide resources as DUNE require so no particular angle - can be agnostic (within reason) UK member of DUNE have lots of experience of DIRAC

AAI

- DB: From site admin dealing with CERN bureaucracy exactly the same as FNAL - CERN is not gold standard has pros and cons
- CB: Just about scalable with CERN, first big experiment with FNAL leading so deal with two major labs. Get AAI working with 2 major labs
- AM: How many host labs does this scale consider Belle, astronomers
- LSK: Jump to N labs is scary. Would like WLCG AI group solution for everybody
- CERN uses ServiceNOW (Microsoft product) (as do FNAL) how you support SSO and registration. Can possible share DB, then redundancy to offer it as a service to Astronomy

Other

- CB: As a person new to DUNE. What is the big problem I could work on?
- HS: How deal with immense events? Object stores? Root vs HDF5?
- LSK: HPC IO workshop at Argonne - ROOT get rid of TFile, then ROOT write directly to object store
- ST: How to get DUNE people in these meetings?
- RN: Is there a tutorial how to run an analysis?

Slack Transcript

Eric Vaandering [9:08 AM]

<https://indico.fnal.gov/event/18094/session/9/contribution/12/material/slides/0.pdf> are the slides being discussed in the pre-meeting Rucio discussion (edited)

Agenda: <https://indico.fnal.gov/event/18094/other-view?view=standard> (edited)

Cedric Serfon [9:08 AM]

https://docs.google.com/presentation/d/1n_OaT5jPEyGpiN8Y-84Mm5FFYb4cSggzQtDVzp0SEiM/edit#slide=id.p

Ste Jones [2:46 PM]

Here is that write-up that GridPP made about our work with the LSST experiment:

<https://www.gridpp.ac.uk/users/case-studies/lst/>

Margaret Votava [3:09 PM]

what certificate can i use to access?

awashbro [3:11 PM]

@Ste Jones do you need to be in the GridPP VO for that page?

Margaret Votava [3:23 PM]

it failed at first, but doesn't now - someone must have fixed it.

DES jobs that run at FNAL can need very large footprints - 10s of Gbytes. did the jobs run on gridpp have the same requirements?

Ste Jones [3:50 PM]

It's the sort of thing that should be public...

Batch jobs that run on T2s at GridPP (such as Liverpool) generally fall into two categories.

Single core:2GB, or Eight core:16 GB. There is also a class of cloud jobs, that Andrew McNab will discuss in his talk on VAC.

@awashbro it looks like it is working now....

timmm [3:54 PM]

Margaret was speaking of the amount of scratch disk needed by DES jobs, not the RAM

And it's not only disk space it's IOPS too

Ste Jones [3:59 PM]

Oops. I can say about L'pool. We generally have 1 TB drives in (say) a 24 slots worker node, which would be ~ 40 GB scratch (less some margin for OS, swap, whatever...) On the other hand, it might be tight on some of our nodes with a different config. But 10GB does not sound totally huge, esp. as you may not get all those sort of jobs on the same system at once. I'd wonder about copying such a file in and out if it was needed, though. It's probably fine if it's "working storage".

Re: IOPS. Yes ... it slows things down.

Margaret Votava [4:28 PM]

no - i was speaking about ram. we have special worker nodes for DES jobs.

Ste Jones [4:45 PM]

Dark Energy Survey... I see. Sometimes, job profiles crop up that do not instantly fit neatly with GridPP's "standard" configs. One size may not fit all. I can dream up answers to this, to make the system "elastic", so to speak. But the present usual baseline is to pick a profile from a predetermined list; and there are two types in the list at present; Single core:2GB, or Eight core:16 GB.

Any colour you like as long as it is black, springs to mind. And that worked for a bit, at least....

Daniela Bauer [10:33 AM]

@Ste Jones @Margaret Votava These kinds of jobs would have to go on IRIS resources once we get them. I prepared an LZ proposal to route their "high memory" (~12 GB/core) to run on IRIS clouds once they come up. The target is April 2019. So if a GridPP supported (that's the important bit here) VO has specific requirements that would be the route to go to try and get them. The person to ask about this is Jeremy Yates and it would have to go to the IRIS DB (I think that stands for "delivery board", but it's just the IRIS term for management).

awashbro [11:30 AM]

@Tanya on ML approaches for distributed computing you might be interested in some related approaches from ATLAS (presented last week) - I have added them to the shared folder <https://goo.gl/zV4CP4>

Raja Nandakumar [12:04 PM]

Not clear to me the need for 1Million simultaneous jobs. How critical is this? Andrew Norman seemed quite emphatic that it is a minimum benchmark (unless I misunderstood something here)

Tanya [1:33 PM]

@Daniela Bauer - some of Fermilab dashboards are open for public. You could try <https://fifemon.fnal.gov/d/000000102/fife-summary?orgId=1> and (as well as <https://fifemon.fnal.gov/d/000000109/dcache?orgId=1>) let me know if it doesn't work for you.

Daniela Bauer [1:36 PM]

@Tanya I can't find anything about my (or any) UK site on this though ?

@illingwo can see

http://wlcg-sam-cms.cern.ch/templates/ember/#/plot?profile=CMS_CRITICAL&sites=T2_UK_London_IC (over https with a certificate) even though he is not a member of CMS. This is not considered confidential.

This is kind of what we are used to.

Tanya [1:39 PM]

nope - this is required fermi account, but if you send me the list of trusted DNs I could try to get an access

as I said it is policy decision that i cannot make; also with GDPR in place it is even more difficult if to figure out if listing person's DN and unname/gname and VO affiliation is legal

e.g

[http://wlcg-sam-cms.cern.ch/dashboard/request.py/metricOutput?host=ceprod05.grid.hep.ph.ic.ac.uk&metricfqan=org.cms.WN-env%20\(/cms/Role_lcgadmin\)&time=2018-10-09T12:26:47Z](http://wlcg-sam-cms.cern.ch/dashboard/request.py/metricOutput?host=ceprod05.grid.hep.ph.ic.ac.uk&metricfqan=org.cms.WN-env%20(/cms/Role_lcgadmin)&time=2018-10-09T12:26:47Z) - has a proxy of Andrea Sciaba and all his affiliations - is it ok?

Daniela Bauer [1:47 PM]

Why not ? As Robert pointed out, as soon as you log into a FNAL computer you get the DOE notice that nothing you do on this computer is private. My DN is /C=UK/O=eScience/OU=Imperial/L=Physics/CN=daniela bauer

Tanya [1:49 PM]

but you are not login at fermilab computer to get this info - it should then have the same disclaimer

we need to talk to our security team

Daniela Bauer [1:51 PM]

Those are nagios tests, really they should run with a robot cert, this is just CMS being lazy.

Daniela Bauer [1:58 PM]

Does your AUP not cover people having access to user data for monitoring and accounting ? When users join a VO, usually it is explained to them that they are using other people's computers and those people have a right to see what's running on their machines. I realize the WLCG page is probably a bit too open, but then there's the "it's paid from taxes, we should be able to see what is done with them" argument.

Tanya [2:01 PM]

we don't sign AUP because we sign Fermilab computing policy. If you just want to see sites, you could use OSG Gracc:

check this

[https://gracc.opensciencegrid.org/dashboard/db/site-summary?orgId=1&var-interval=\\${__auto_interval}&var-site=UKI-LT2-Brunel&var-site=UKI-LT2-IC-HEP&var-site=UKI-LT2-QMUL&var-site=UKI-LT2-RHUL&var-site=UKI-NORTHGRID-LANCS-HEP&var-site=UKI-NORTHGRID-LIV-HEP&var-site=UKI-NORTHGRID-MAN-HEP&var-site=UKI-NORTHGRID-SHEF-HEP&var-site=UKI-SCOTGRID-ECDF&var-site=UKI-SCOTGRID-GLASGOW&var-site=UKI-SOUTHGRID-BRIS-HEP&var-site=UKI-SOUTHGRID-OX-HEP&var-site=UKI-SOUTHGRID-RALPP&var-type=Payload](https://gracc.opensciencegrid.org/dashboard/db/site-summary?orgId=1&var-interval=${__auto_interval}&var-site=UKI-LT2-Brunel&var-site=UKI-LT2-IC-HEP&var-site=UKI-LT2-QMUL&var-site=UKI-LT2-RHUL&var-site=UKI-NORTHGRID-LANCS-HEP&var-site=UKI-NORTHGRID-LIV-HEP&var-site=UKI-NORTHGRID-MAN-HEP&var-site=UKI-NORTHGRID-SHEF-HEP&var-site=UKI-SCOTGRID-ECDF&var-site=UKI-SCOTGRID-GLASGOW&var-site=UKI-SOUTHGRID-BRIS-HEP&var-site=UKI-SOUTHGRID-OX-HEP&var-site=UKI-SOUTHGRID-RALPP&var-type=Payload)

Heidi Schellman [4:45 PM]

DUNE Computing wiki https://wiki.dunescience.org/wiki/DUNE_Computing has a getting started part. at the bottom

https://wiki.dunescience.org/wiki/Accessing_ProtoDUNE-SP_Cold_Box_Test_Data_in_LArSoft ask us if you didn't hear the read only access info. Or ask Raja....

Sam Skipsey [5:19 PM]

For Heidi: <https://indico.cern.ch/category/10360/> DOMA indico category

The top level DOMA list is wlcg-doma@cern.ch

Message Input