Digital Research Alliance of Canada | Alliance de recherche numérique du Canada

Supporting Distributed Subatomic Physics Computing in Canada

Leslie Groer
University of Toronto
for the
Subatomic Physics National Team
October 19, 2023

UNIVERSITY OF TORONTO

SciNet

Some Ancient History (~2000-2007)

INNOVATION
Canada Foundation for Innovation
Fondation canadienne pour l'innovation

- Mostly individual project funding from Canadian Foundation for Innovation, Provinces and Universities
- Departmental clusters running distributed computing for HEP at FNAL, DESY and other locations
- Natural Sciences and Engineering Research Council of Canada (NSERC) and other funding for personnel
- Start building WLCG Tier-2 resources and Canadian WLCG Cloud
  - ALBERTA-LCG2
  - SFU-LCG2
  - TORONTO-LCG2
  - Umontreal-LCG2
  - VICTORIA-LCG2
  - Australia-ATLAS
- Pledging 5% of ATLAS Tier-2 capacity
- TRIUMF Tier-1 established in 2007 (10% of ATLAS)
- Operating with 1Gbps lightpaths

Compute Canada National Platforms Era (2008-2017)

- Push by funding agencies to move to national systems, professionally run, rather than small individual departmental systems
- CFI National Platform Funds 1 announced Dec 2006 (CA$60M)
- Seeds of Compute Canada established and provincial matches procured
  - finally incorporated in 2012
- First national systems installed end 2008
- 5 official Canadian WLCG Tier-2 sites established or migrated
  - CA-SCINET-T2 (Toronto)
  - CA-MCGILL-CLUMEQ-T2 (Montreal)
  - SFU-LCG2 (Vancouver)
  - CA-VICTORIA-WESTGRID-T2 (Victoria)
  - CA-ALBERTA-WESTGRID-T2 (Edmonton)
  - + IAAS & Australia-Nectar cloud resources (2012, Tier-3)

CA-EAST Federation

CA-WEST Federation

All dCache-based except McGill (StoRM on GPFS) and Austrialia (DPM)

- 10Gbps VRF layer supported by CANARIE for LHCONE

compute canada | calcul canada

Compute Canada National Platforms II (2017-2022)

- Consolidate national sites to:
  - 3 general purpose batch computing clusters (SFU, Waterloo, McGill)
  - One high-performance cluster – dragonfly IB (Toronto)
  - One cloud computing site (Victoria)
- Consolidate to 3 official Canadian WLCG Tier-2 sites
  - CA-WATERLOO-T2 (Waterloo) } CA-EAST Federation
  - CA-SFU-T2 (Vancouver)
  - CA-VICTORIA-WESTGRID-T2 (Victoria) } CA-WEST Federation
  - IAAS (mostly Victoria opportunistic cloud resources – Tier-3)
- Also moved the TRIUMF Tier-1 to the SFU datacentre
- Establish 100Gbps backbones with CANARIE and NRENs
- Canadian Tier-2 resources reliably provide about as much compute to ATLAS as TRIUMF Tier-1 (large “burst” capacity on occasions)

Digital Research Alliance of Canada (2022→)

- New federal funding established March 2018 by the Ministry under the Department for Innovation, Science and Economic Development Canada (ISED) with CA$572M
- Digital Research Alliance of Canada (DRAC) created 2019
  and officially took over from Compute Canada Federation April 2022

- March 2023 - Approval from ISED for the federal CN$224M for DRI
  - CA$120M for ARC renewal
- Working on matching funds from provinces and sign-off on hosting institutions
- Plan is to issue RFP's in next few months with equipment delivered and installed by 2024Q3-Q4
- At least two years delay in funding replacements at most sites so critical
- Potential detours and sticking points along the way….

Subatomic Physics National Team

- Team formed in April 2016 under Compute Canada
- Team Mandate
  - "…responsible for all aspects of procurement, deployment and support of distributed data handling, analysis and processing systems and software for the large international experimental subatomic physics community and similar collaborations"
- Added support over the years for new experiments and collaborations
- Face-to Face meetings every couple of years

Tier-2 Support	**TRIUMF Tier-1 & ATLAS Support**
Leslie Groer (UofT)	Reda Tafirout
Sergey Chelsky (SFU)	Asoka De Silva
Lixin Liu (SFU)	Rajan Devbhandari
Simon Nderitu (McGill)	Fernando Fernandez Galindo
Erming Pei (UAlberta)	Vitaliy Kondratenko
Ryan Taylor (UVic)	Simon Liu
[Bryan Caron]	Di Qing
[Neil Knecht]	Yun-Ha Shin
HEPnet Canada (Victoria)	Andrew Wong
Randy Sobie	[Rodney Walker]
Tristan Sullivan	[Denice Deatrich]
Marcus Ebert	**ATLAS-Canada PI**
[Ian Gable]	Isabel Trigger
[Rolf Seuster]	[Doug Gingrich]
	[Reda Tafirout]
[] – former member	[Michel Vetterli]

SPNT in the Canadian DRI space

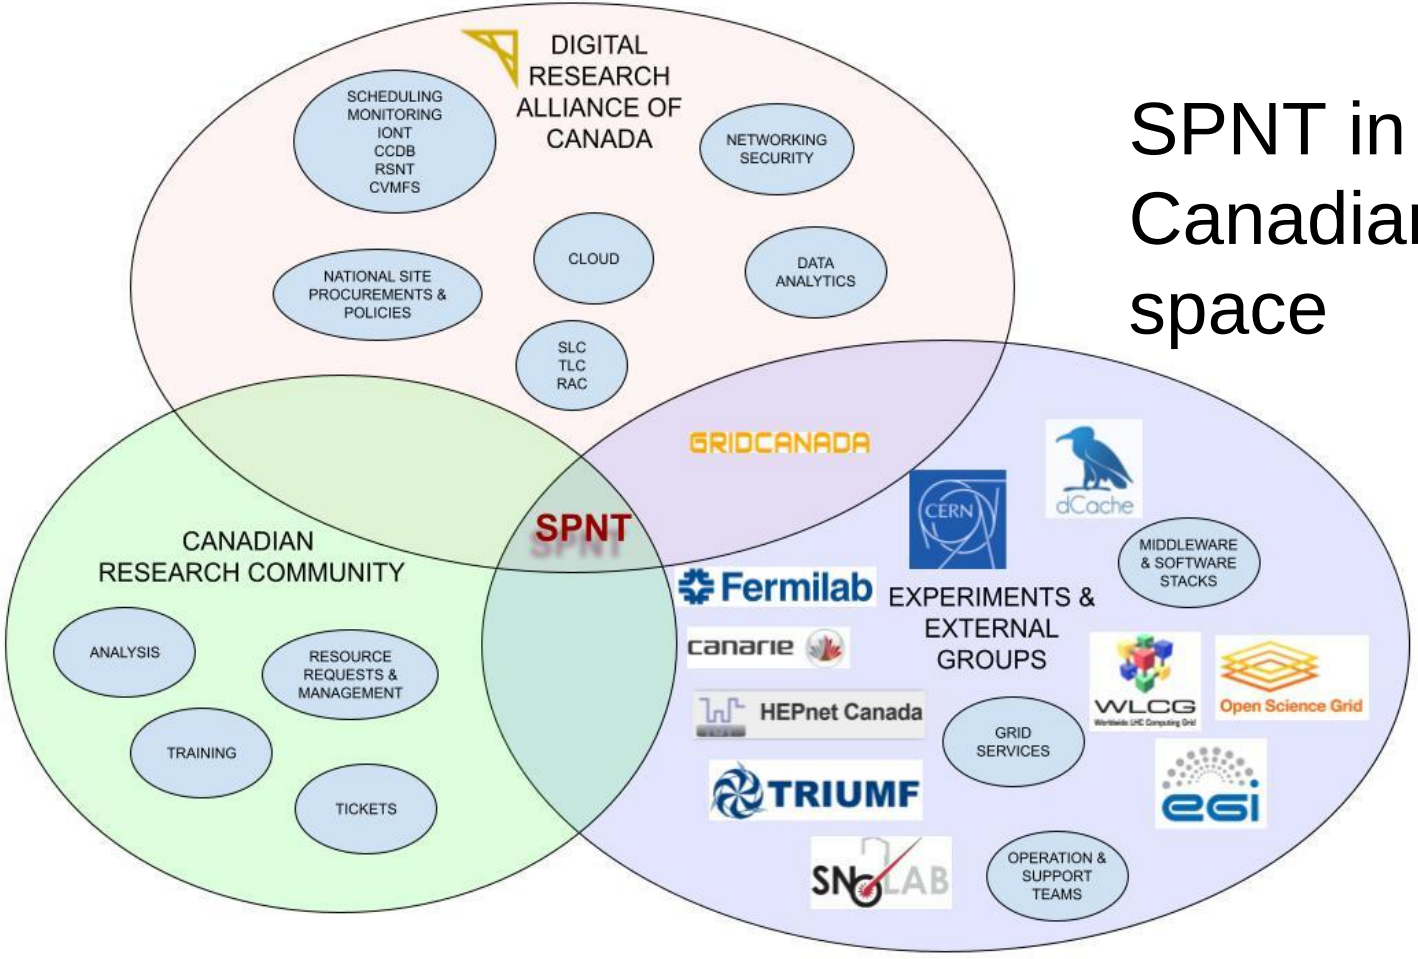

Current Matrix of Services

	Grid	Glideins OSG	Cloud	GPU	dCache	Project	Nearline (tape)
ATLAS	X		X		X		
Belle II	X		X		X		
CHIME						X	X
DEAP						X	X
DUNE	X		X		(X)	X	
GlueX/CLAS12		X					
IceCube		X		X			
IGWN (LIGO)		X		X			
SNO+	X				X		X
SuperCDMS			X			X	
T2K/HyperK				X	X		(X)

Resource Allocation Process

- For larger projects, PI's submit requests for resources
  - Technical and scientific review process
- Request usually good for 3 years, but allocations decided each year based on availability

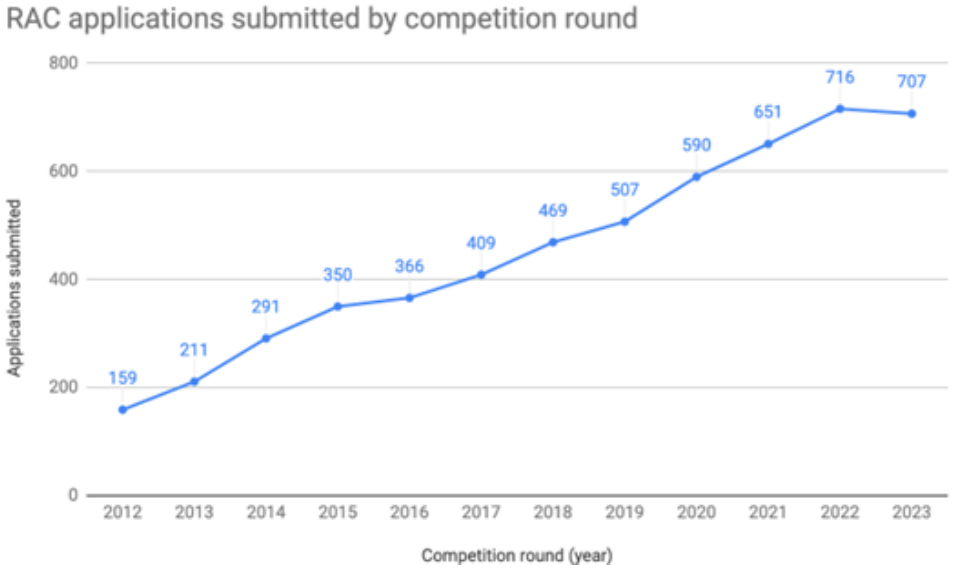

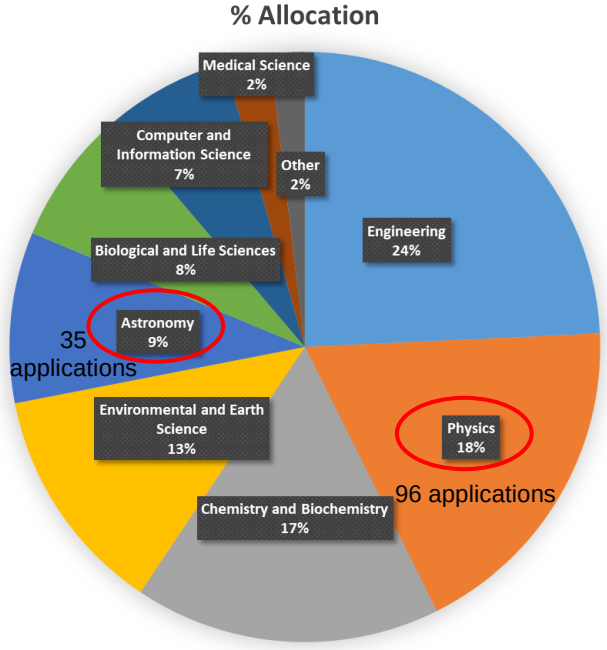

Total Resource Trends

CPU-core growth of about 50% in 10 years
Demand increased by a factor of 3!

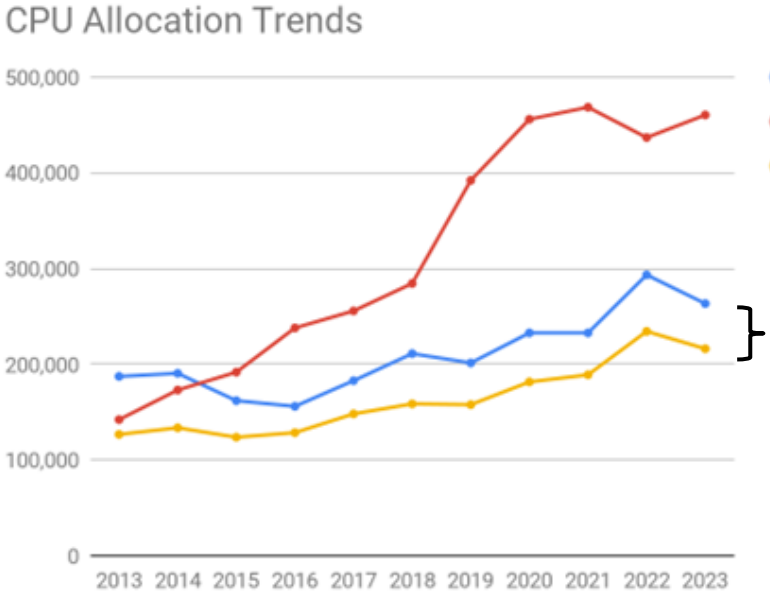

Default usage
~20%

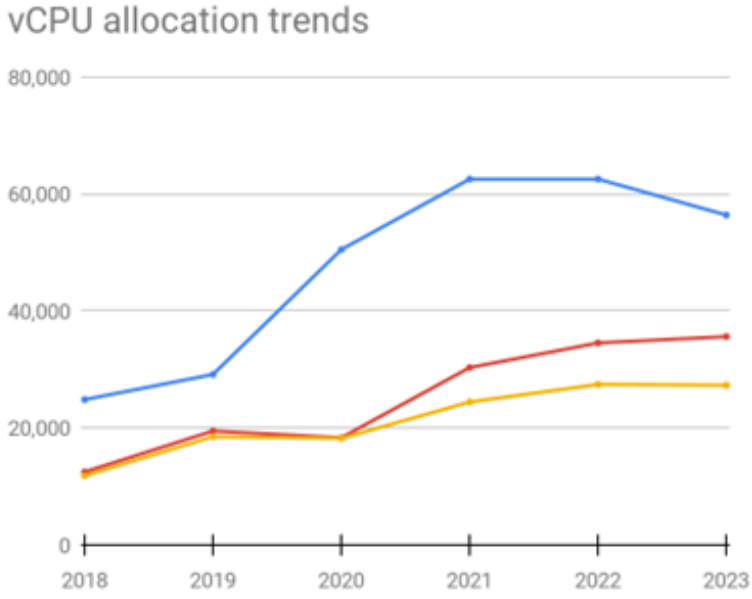

Subatomic Physics & Astronomy Distributed Computing projects account for about 13.5% of CPU-core resources allocated in 2023

Allocating about 150PB of storage in various forms

HEP and related CPU Requests Projection (2022-2026)

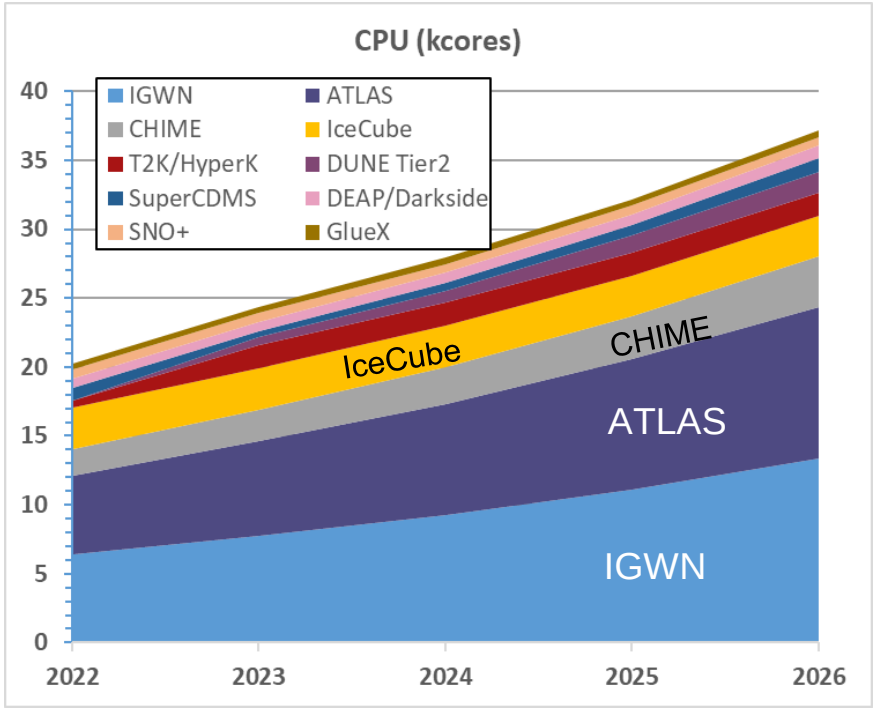

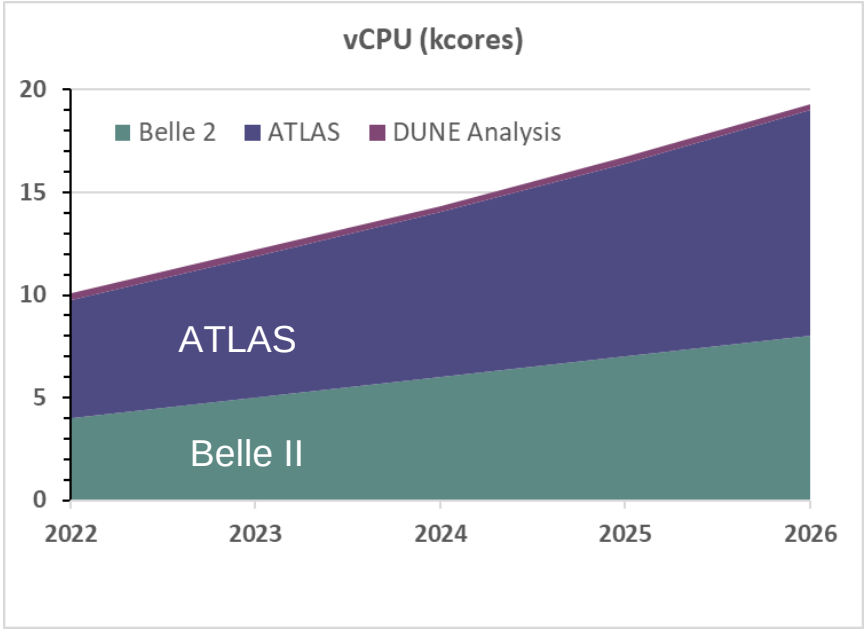

- CPU and vCPU both ~double from 2022
- GPU – IGWN 250, HyperK 150, IceCube 100

dCache Storage Requests Projection (2022-2026)

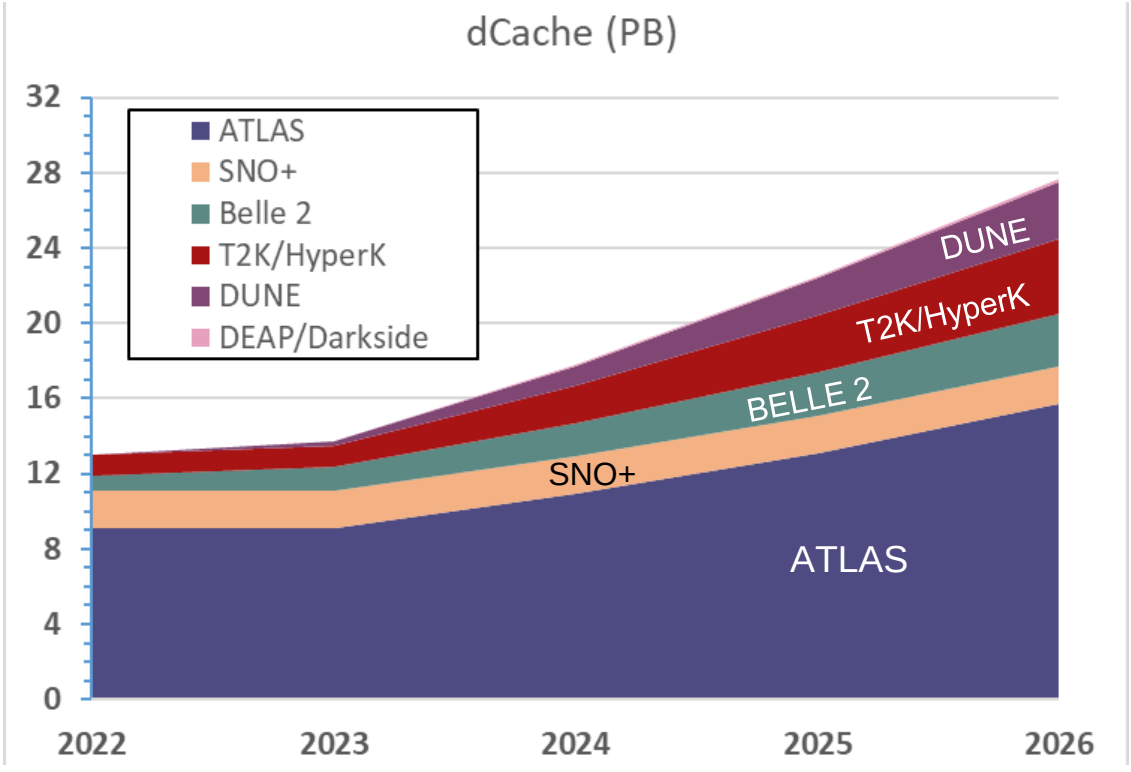

Essentially double by 2026

HEP and related Storage Requests Projection (2022-2026)

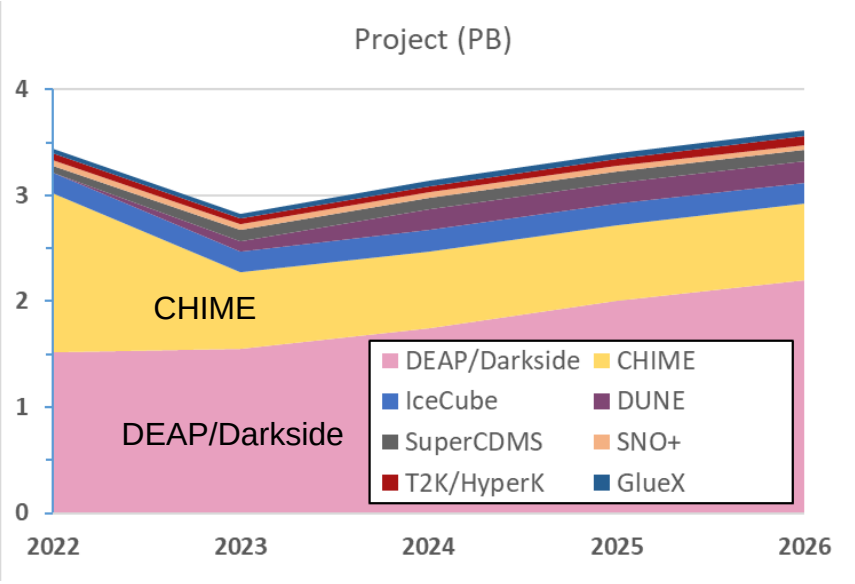

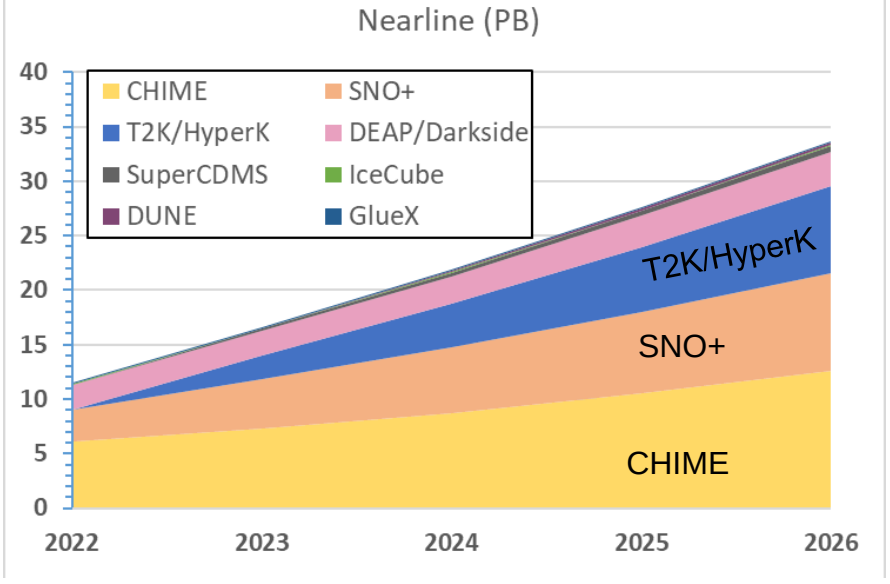

/project space somewhat under control but /nearline (tape) triples from 2022

DRAC available storage ~tripled over the last 5 years

Deployment Scenarios

- Deploying and maintaining grid HTC services on large HPC clusters is always a challenge and somewhat at odds with HPC
- Grid services deployed on a variety of hardware over the years but tending to VMs on various platforms with promise of HA (not always met…)
- Despite push for commonalities and consistency, each site still has unique history and deployment strategies and hardware
  - xCAT, Puppet, oneSIS, Kubernetes
  - Ethernet, Infiniband, OmniPath, Relationship with NREN, internal firewalls and routing
  - Sites share experience and scripts but no common config management tools
  - Common Account database, LDAP and Scheduling (SLURM) and software stack developed over last few years has helped
- Including HEP_OSlibs in the compute node images was always controversial so containers have been a lifesaver, as well as CVMFS

Benefits to Canadian DRI

- HEP pioneered use of X.509 grid certificates & globus file transfers
  - Still using Globus Connect for lots of projects
- Strong positive experience with CVMFS led to common adoption for Compute Canada software stack which is now exported worldwide
  - /cvmfs/soft.computecanada.ca
- Early adopters of PerfSONAR network monitoring
  - Adopted by Compute Canada national sites & CANARIE
- Pioneered Singularity containers on clusters which are now standard
- Strong support from CANARIE & NRENs for backbones and supporting evolving transportation protocols (lightpaths, VRFs, MPLS)
- Deploying Grid services in HA environment informed other deployments
- Work by UVic team on Cloud Scheduler and other deployment strategies used for other projects

Some Past Challenges and Issues

- Have had to deal with sites that have diskless nodes **
  - Grid jobs at scale stress shared filesystems enormously as meta-operations are expensive (inode creation and deletions, finds & disk usage queries, etc)
  - Solution: Buy or borrow localdisk!
- Not all sites have WAN access open from cluster
  - Bifurcate cluster**
  - WAN lease scheme using agents
  - Will be experimenting soon with what we can do via proxy services
- Jobs/pilots that don't match their stated batch requirements (i.e. size and length) are very inefficient to schedule on large clusters with 1000's of users and get severely dinged in accounting and priority scheduling
- Not all sites allow/trust `cvmfsexec` needed to access authenticated repos
  - running with user caches has worked

** SciNet: Lessons Learned from Building a Power-efficient Top-20 System and Data Centre, *J. Phys.: Conf. Ser. (2010) Vol.256, No.1*. http://iopscience.iop.org/article/10.1088/1742-6596/256/1/012026/meta

Some Recent Challenges and Opportunities

- Supporting Waterloo site with no SPNT member native to the hosting collaboration so limited root access
  - Also dropped off LHCONE a few years ago and no way back at the moment
- Network providers are all IPv6-ready but sites have never made this a priority for full deployment so still languishing
- Maintaining older hardware especially dCache storage past 5-6 year warranty life is a challenge
  - Most sites relying on DELL DSS7000 chasses (90 x 8TB drives)
  - Replace a drive every ~ 2-4 weeks per chassis
- Migrating from virtualized batch system on Cloud at Uvic to the first deployment of a true cloud-native k8s WLCG Tier-2 site**

** see Ryan Taylor's Talk this morning

Interactive Analysis Platforms & Support

- Limited Tier-3 access on the sites
  - Setup only at SFU for the moment; bandwidth for dCache NFSv4.1 limited
- Access to interactive sessions or Jupyter notebooks supported but rather limited and no auto-scaling of resources
- Proto-analysis facility built on Cloud resources for DUNE was very time consuming and did not leverage most of the existing infrastructure
  - DRAC folks looking at PAAS but seems orthogonal to the needs of HEP community (mostly used for teaching and small-scale RDM activities)
- General awareness of HSF and other activities but no clear direction of where to put in effort and possible synergies between experiments
  - How to tie these to current allocation projects especially storage, A&A schemes
- Expect lots of discussion around this in next few years but need resources to build or leverage solutions

Ongoing Concerns & Projects

- Funding cycles tend to be stretched longer than we would like and adds a lot of uncertainty for MOU commitments etc
- Some services being redeployed in AL9 at some sites
- Redeploy services at Waterloo after issues with older CEPH cluster underpinning OpenStack VMs
- Prototyping EOS on CEPH for Victoria for new deployment as cloud-native storage element
- Prototype Computing Tier-3 at McGill
- Rudimentary support for moving data off dCache to tape at SFU but totally manual currently
- Would like to improve and standardize our monitoring and alert services
- Utilize gitlab at least for some configuration tracking & management
- Rebuilding Grid Canada x.509 Certificate Authority

Future & Long Term Support

- Eagerly awaiting new funding and “refreshing” systems
  - Next year (or two?) a little rocky
  - Likely less commonality in hardware purchases for dCache storage, increasing deployment and operational burden
- Work already under way in the DRAC for new capital for 2025-2030 period
  - Coincident with expected resource “inflection point” with HL-LHC, SKA and other projects
- Already discussing 400G needs with CANARIE and NRENs
- Optimistic about the long-term prospects but are seeing strains in the national resource allocation models for the biggest projects and how these get funded
- Growing issue between HPC sites and HTC needs with emphasis on security considerations and limiting connectivity
  - Rolling out MFA (DUO) for all users this year
  - May need workarounds to not impact workflows
- Supporting Interactive Analysis activities well still very much an open question
- Discussions with community about deploying a Canadian Rucio service

Some (Obvious?) Lessons Learned

- Embedding of folks wearing HEP hats critical for long-term success
- Establish working partnerships with the communities you are supporting
- Maintain good relations with national and regional network providers
- Avoid or minimize reliance on large shared filesystems (if possible)
- Running separate dCache instances per project tends to be a more reliable and effortless solution
- Include test equipment/platforms in initial purchases if possible as tough to test migrations or upgrades on operating platforms

- Funding cycles are always longer than expected so plan accordingly, manage expectations and keep lots of duct tape around!

Canadian National DRI: An (incomplete) history