Nikhef

Nikhef SEP panel visit November 2023

Physics Data Processing, Engineering and Computing

accelerating 'time to results' through computing and collaboration

David Groep
Roel Aaij
Ronald Starink
Karol Popławski
Osama Karkout

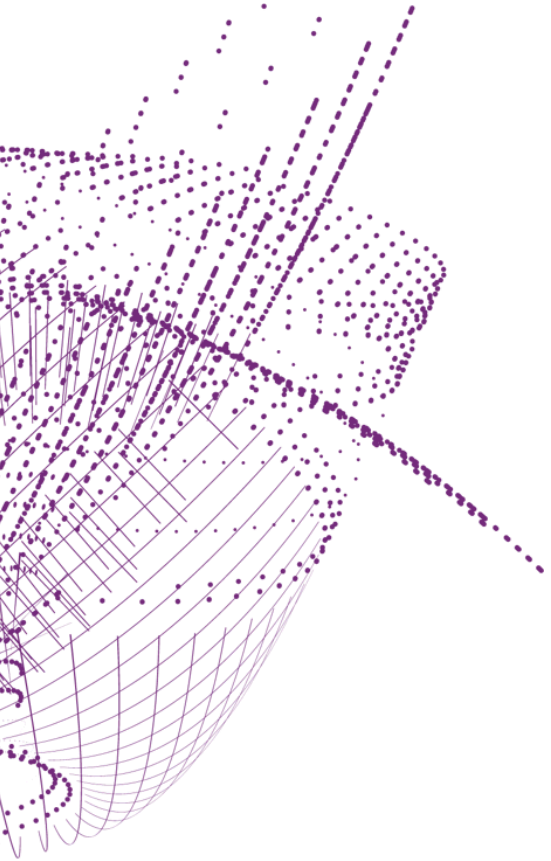

David Groep | PL Physics Data Processing

Physics Data Processing & ‘CT-PDP’ engineering

PDP: Computing as research and instrumentation

validated through our real-life applications

Physics Data Processing programme lines

1. **infrastructure, network & systems research**
   - **building** ‘research IT facilities’ through co-design & development
   - big data science innovation: research **next gen IT infrastructure**

2. **infrastructure for trusted collaboration**
   - **trust and identity** for enabling communities
   - managing complexity of **collaboration mechanisms**
   - **securing** the infrastructure of our **open science** cloud

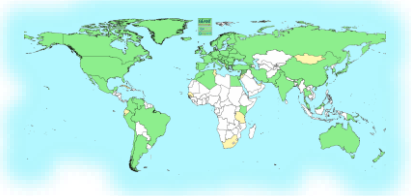

3. **algorithmic design patterns** - *more in Roel’s introduction*
   - **GPU accelerated** computing, **Quantum Computing**, AI and **Machine Learning**

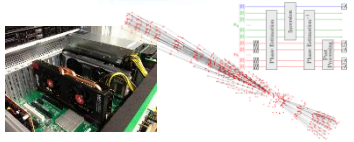

Images: H234b data centre; IGTF trust map; plofkip.nikhef.nl; Davide Nicotra, Maastricht and Nikhef, arXiv:2308.00619v2 quant-ph for JInst.

Physics Data Processing and CT-PDP engineering effort

2.5 staff (**David Groep**, **Roel Aaij**, Jeff Templon)

1 postdoc (Maarten van Veghel)

~ 11 engineers: DevOps, research software, RDM, Collaboration / Trust & Identity
with organic embedding in the Computer Technology engineering group

With wide range of activity leads, including for example **LHCb RTA Reco convener**, **SURF innovation expert group chair**, **EOSC** Security Coordinator, **AEGIS** Trust and Identity policy lead, **Interoperable Global Trust Federation** chair, Dutch **National Infrastructure Executive**, lead for the **Thematic DCC Natural and Engineering Sciences**, and members in board & committees for the (global) e-Infrastructure landscape: GEANT GCC, PC-GWI, CieDO…

Infrastructure for Research

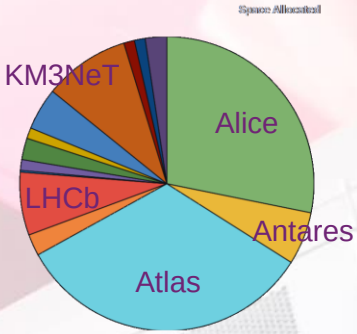

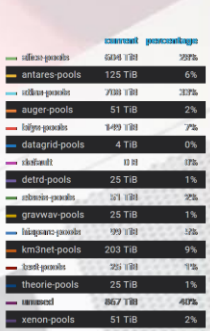

High-Through Compute (HTC) + HT Storage

- **National e-Infrastructure** coordinated by SURF
  LHC NL-T1, IGWN, KM3NeT, Xenon, DUNE, WeNMR, MinE ...
- **'Stoomboot' local analysis facility** + IGWN cluster & 'submit node'

~ 12 000 cores (total), 13 PByte storage – very competitive *w.r.t.* commercial cloud

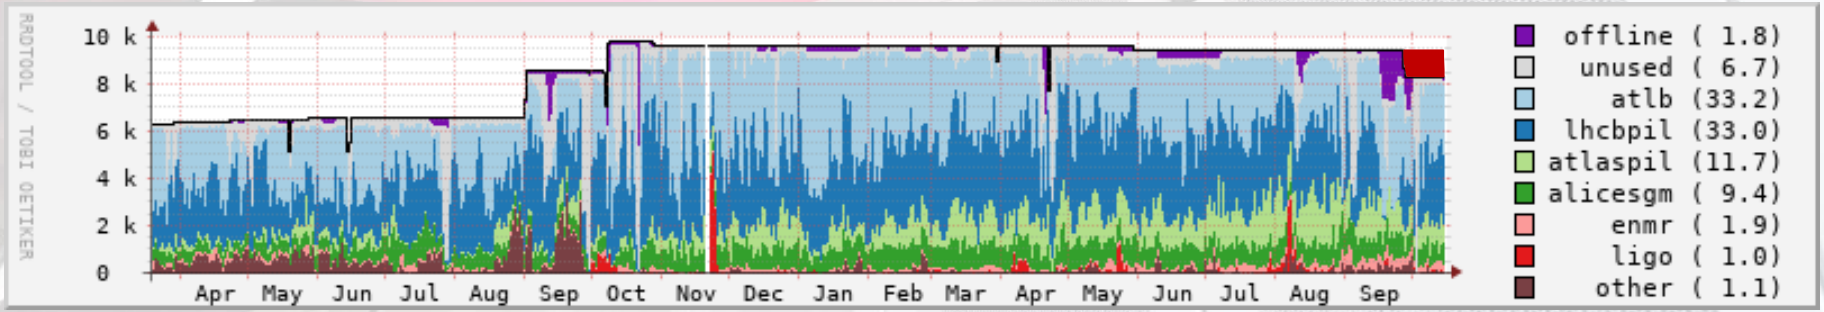

Occupancy: NDPF DNI processing facility in the period March 2021 .. October 2022. Top-right: storage capacity allocated in DNI Nikhef segment

Common solutions are essential for our ‘national’ facility

Alignment of common e-Infrastructure
and shared use by experiments
(LHC, GW, KM3NeT, DUNE, Xenon, ...)

- *common solutions,* since bespoke systems for each experiment do not scale for Nikhef (or NL)
- efficient sharing of both hardware and DevOps effort
- synergy with other domains helps sustainable funding

IGWN use of NL DNI infrastructure at NIKHEF & SURF 4-11 April, 2023

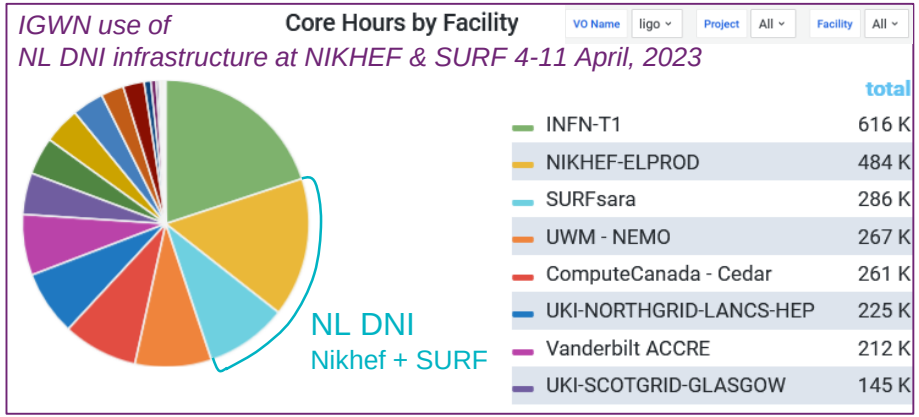

Continuation of our long-term strategy - from EU DataGrid in 2000 onwards:

- drives collaborative efforts we work with for identity management and **common protocols globally**: ACCESS-CI and CILogon (US) – key players for LIGO, DUNE, and US-ATLAS
- common **processing framework** development (e.g. together with KM3NeT in its INFRADEV project)

Data on Tue April 11th 2023 from the OSG accounting for the LIGO VO for past week (with also SURFsara fully in production)
https://gracc.opensciencegrid.org/d/9u1-Q3vVz/cpu-payload-jobs?orgId=1&var-ReportableVOName=ligo&var-Project=All&var-Facility=All&var-Probe=All&var-interval=1d&from=1680566400000&to=1681257600000

Innovation *on* infrastructure

- Network-to-systems integration
- Storage throughput & parallelism
- Systems integration design and tuning

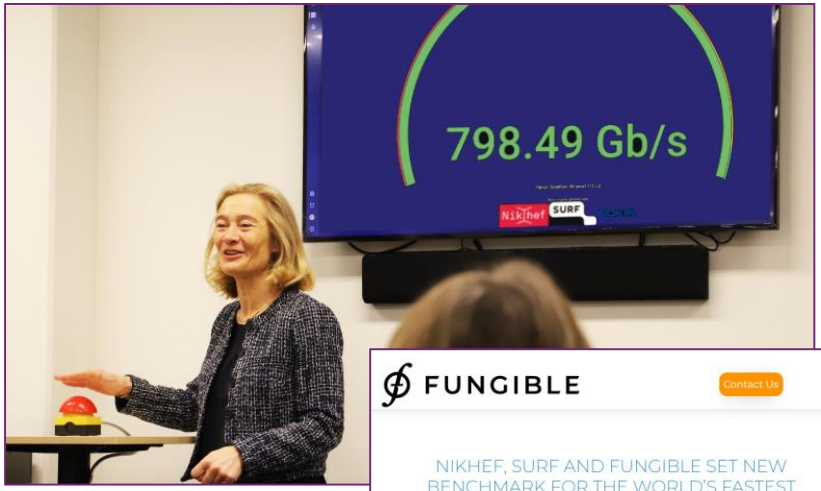

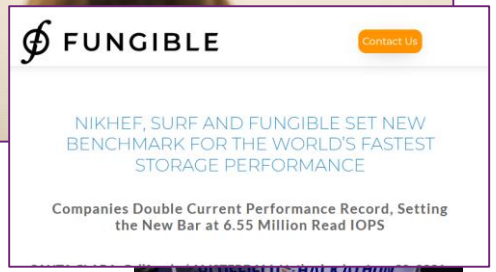

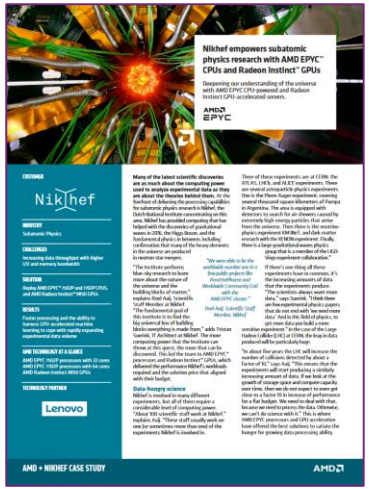

- early **engineering engagement with vendors** to build us suitable systems
- **co-design** of our national HPC systems ('Snellius')
- **data-intensive compute** with DPUs, or on-NIC FPGAs?
- networks > 800Gbps, >1 Bpps (today: 400G to CERN)

Image: Minister of Economic Affairs M. Adriaansens launched the Innovation Hub with Nikhef, SURF, Nokia and NL-ix, January 2023. Composite image from https://www.surf.nl/nieuws/minister-adriaansens-lanceert-testomgeving-voor-supersnelle-netwerktechnologie; Bluefield Hackathon by Nvidia/Mellanox; *abbreviations*: **DPU**: Data Processing Unit; **on-NIC FPGAs:** on-network interface card field programmable gate arrays; **pps**: packets per second

Our science data flows are somebody else's DDoS attack

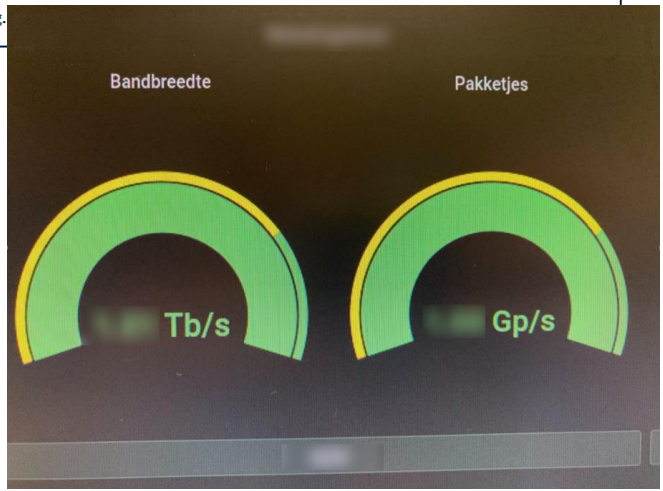

Image sources: belastingdienst.nl, rws.nl, nu.nl, werkentegennederland.nl

Infrastructure for Collaboration

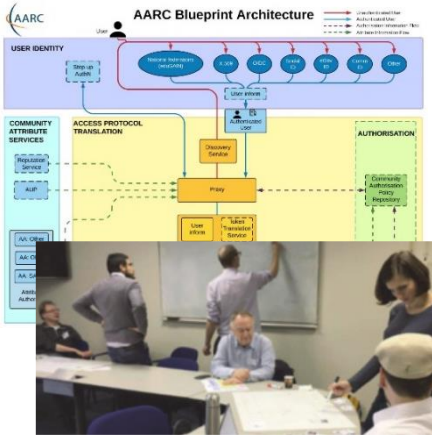

Target impactful areas in architectures for ‘AAI’ & ‘OpSec’

- **authentication & authorization for research collaboration**
  - **AARC** project & community: GEANT Framework projects, R&E federation, identity and credentialing services, EOSC Future, …
  - recently awarded: **EOSC Core** security and **AARC-TREE**
  - **policy frameworks for interoperability** for data protection and global seamless service access
  - continuous **technical evolution** driving IGWN, WLCG in line with AARC and global AAI architecture

- embedding data **processing needs** of our experiments in the (EOSC) landscape
  - **EOSC** Interoperability Framework: Security Baseline, AAI Architecture
  - **eduGAIN** Operational Security for the global R&E inter-federation service
  - **EGI** Advanced computing for research federation**, GEANT community**

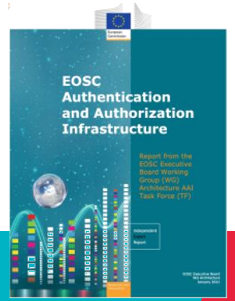

AAI: Authentication and Authorization Infrastructure; **OpSec**: Operational Security (incident response); **AARC**: Authentication and Authorisation for Research Collaboration community and projects**; EOSC**: European Open Science Cloud; **EGI** and **GEANT**: pan-European e-Infrastructure and network collaborations; **IGWN**: International Gravitational Waves Observatory Network;

Collaboration: Research Data Management beyond ‘FA’

FAIR for **live data**, in large volumes, from ‘FA’ towards the ‘I’ and ‘R’

- not *that* many disciplines with really **voluminous data**
  - so nationally join forces with those who do: ASTRON (SRCnet), KNMI (earth observation, seismology), *&c*
- work with those who care about ***software*** to bring data to life on ***infrastructure***
  - NLeSC, 4TU.RD/TUDelft, CWI, and with those who ensure the *infrastructure*: SURF
- for our own analyses and the local (R&D) experiments we work towards *continuous deposition* of **re-usable** data and software:
  - **Thematic Digital Competence Centre** for the Natural and Engineering Sciences
  - **co-develop Djehuty RDM** repository software link ‘Stoomboot’ analysis cluster storage

PDP strategic projects mechanism

Join initiatives and projects that

- *strengthen* the strategic areas
- ensure *continuity* of research and infrastructure

project pathways include
SURF innovation, GN5-*, AARC-TREE, EOSC Core, LHC4D (planned), …

Public partner R&D engagement
AMD, Nokia, Nvidia/MLNX, NL-ix, …
Dutch national government

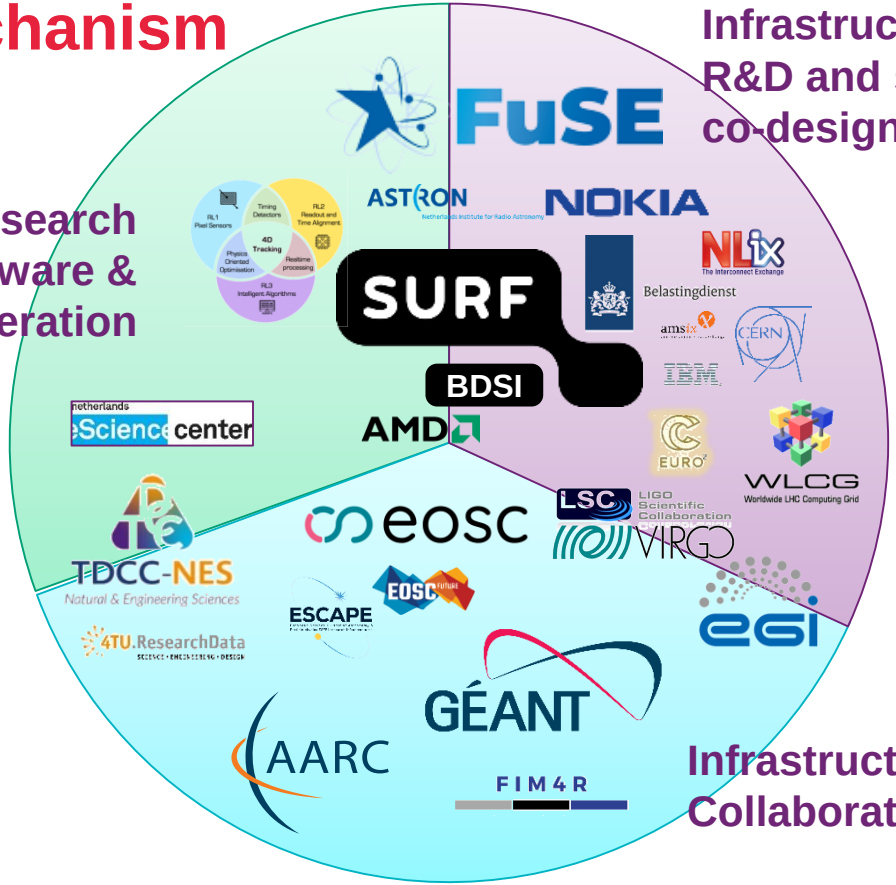

Sustained infrastructure for advanced computing?

Data processing: a **persistent need** for all our experiments. And ‘we’ are not alone!

- many disciplines: ESCAPE (our experiments plus astronomy), bio-informatics, health, SSH, …

  need infrastructure to exploit the collected data with **long-term ICT capabilities**

- project-based funding for ‘upgrades’ of infrastructure not the appropriate way of funding **persistent requirements** … but only thing we have at the moment:

timeline: ICES-KIS, BIG and GWI projects for computing capacity, by submission date

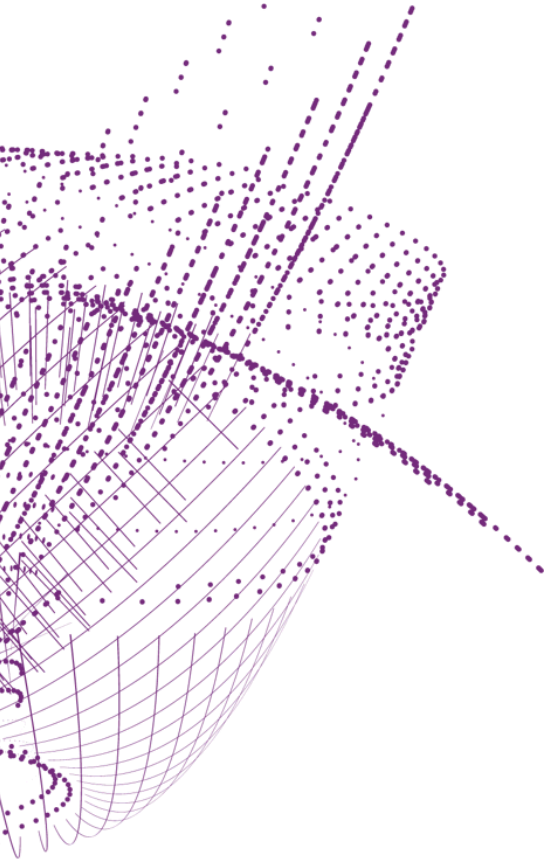

Roel Aaij | Senior Scientific Researcher PDP

Algorithms and Acceleration

Efficient and Scalable Computing For HEP

- CPU-only is not going to be affordable;
  In other words: compute accelerators offer more physics for less money
- Most mature compute accelerators: GPUs

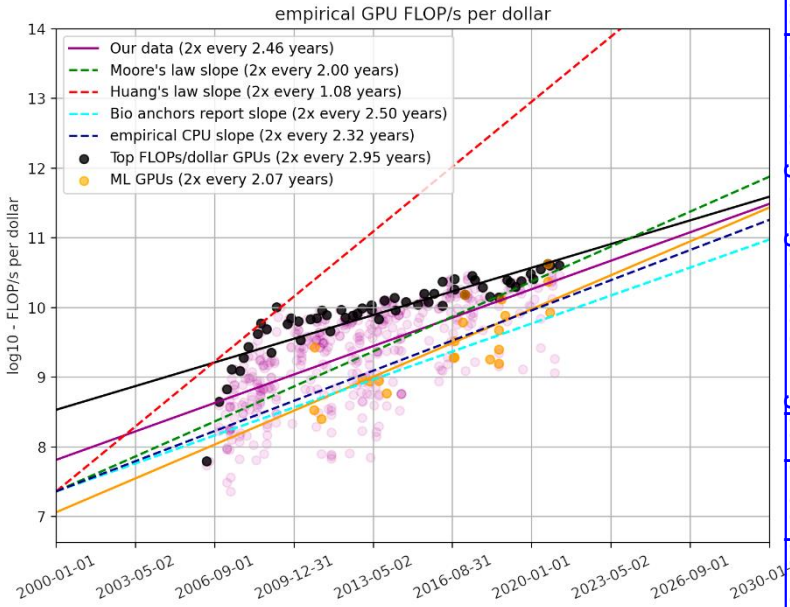

https://epochai.org/blog/trends-in-gpu-price-performance

- How to program?
- How to optimize at algorithm and system level?
- How to integrate in frameworks and infrastructure?
- How to do this efficiently in terms of people's time?
- How to maintain?
- What will future (compute) accelerators look like?
- How to provide career perspective to people who focus on (accelerated) software?

Efficient and Scalable Computing For HEP

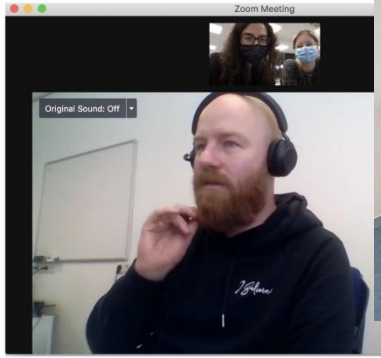

- LHCb's first-stage GPU trigger (Allen)
  - Bespoke application with all-custom kernels
  - 4 TB/s of detector data on 400 GPUs
  - Nominal luminosity next year
  - Focused on integration (DAQ, LHCb stack, etc.)
- - From idea to R&D to production in 5 years
- With NLeSC: fast ML inference
  - Using standard format (ONNX) and libraries
- User (software) support for GPUs (AMD and NVIDIA)
- FASTER: computing for HL-LHC & '4D' reconstruction

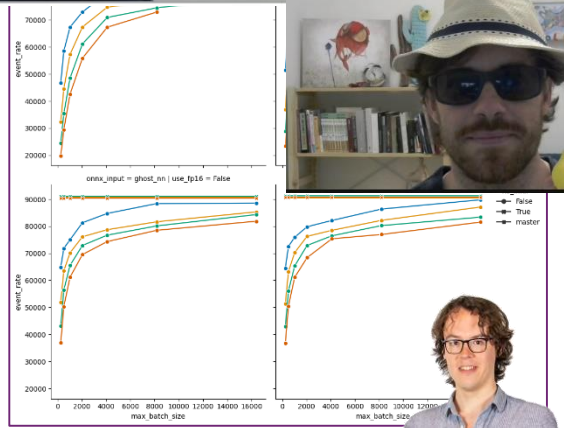

https://github.com/LHC-NLeSC/run-allen-run

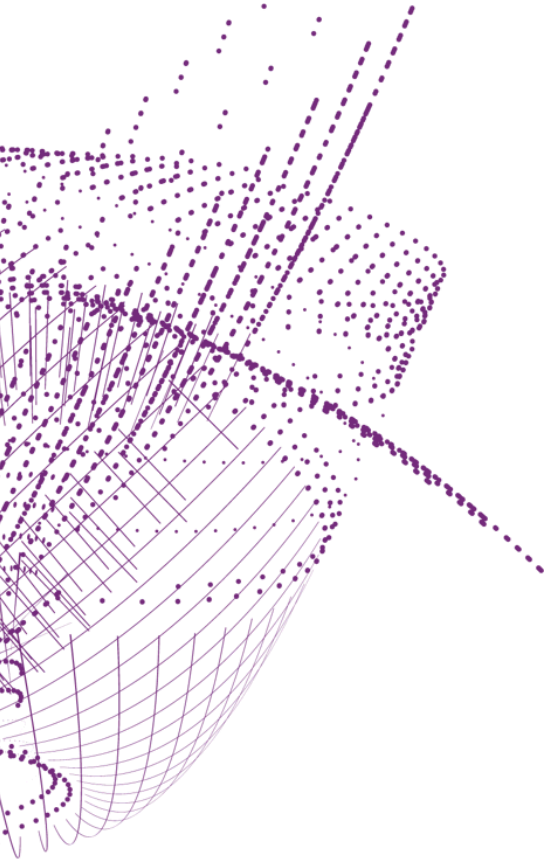

Ronald Starink | TGL Computing Technology

Computing group at Nikhef

CT organisation

CT-B: system administration & service desk

- General support for ICT, end user support
- Staff: 7

CT-PDP: dedicated support for the PDP programme

- Software & infrastructure innovation
- Staff: 11
- (→ presentations DG and RA)

CT-PO: support for projects by experiments

- Software engineering: slow controls, data acquisition, analysis framework support
- Staff: 6
- (→ presentation KP)

Challenges

- General: recruitment in a competitive labour market
- CT-B: balance flexibility ↔ standardization
- CT-PO: matching resource with experiments' needs (time, expertise)

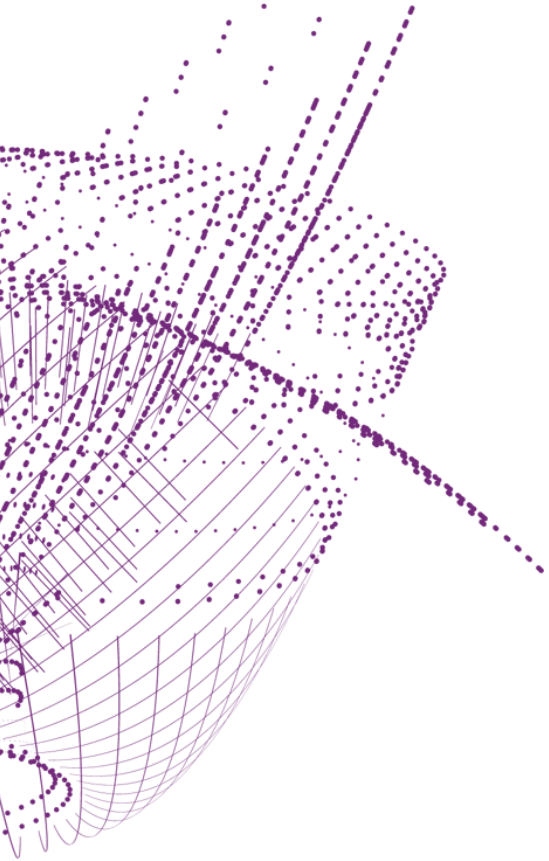

Karol Popławski | Software Engineer CT-PO

Project engineering & support

KP – Software, Controls and more

Barrel Alignment	DAQ of ~5800 channels
MDT DCS Module	1200 chambers: T, B-field, electronics monitoring & configuration
SciFi tracker	FEE configuration with DB, granularity up to 1.5mln thresholds
Scrum Master	MT, ET, R&D

CT-PO

Henk	FELIX
	MDM firmware
Ton	SciFi FEE calibration
	PTOLEMY
Kostis	Future 4D tracking with White Rabbit

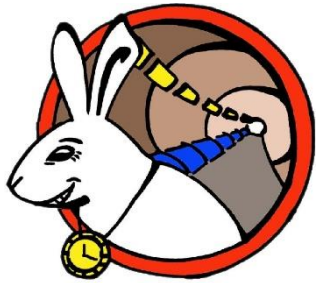

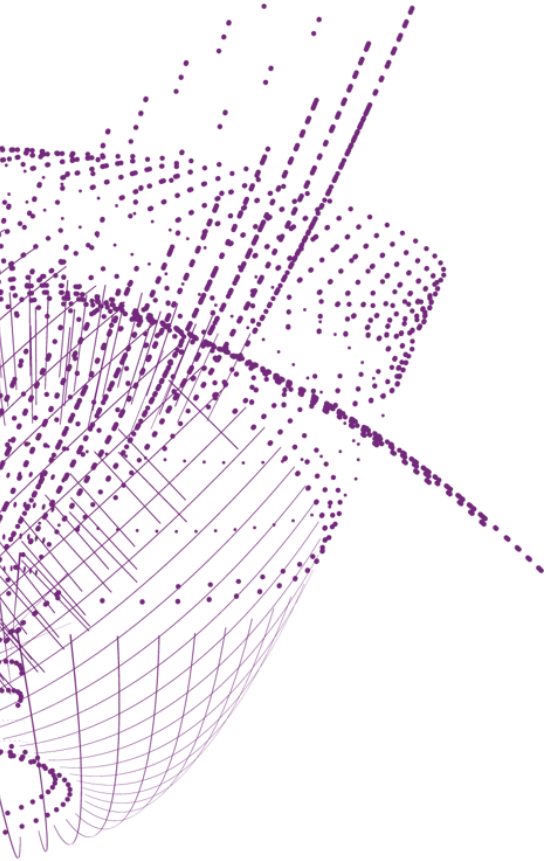

Osama Karkout | happy ATLAS & stoomboot user

How I use stoomboot

Data Analysis

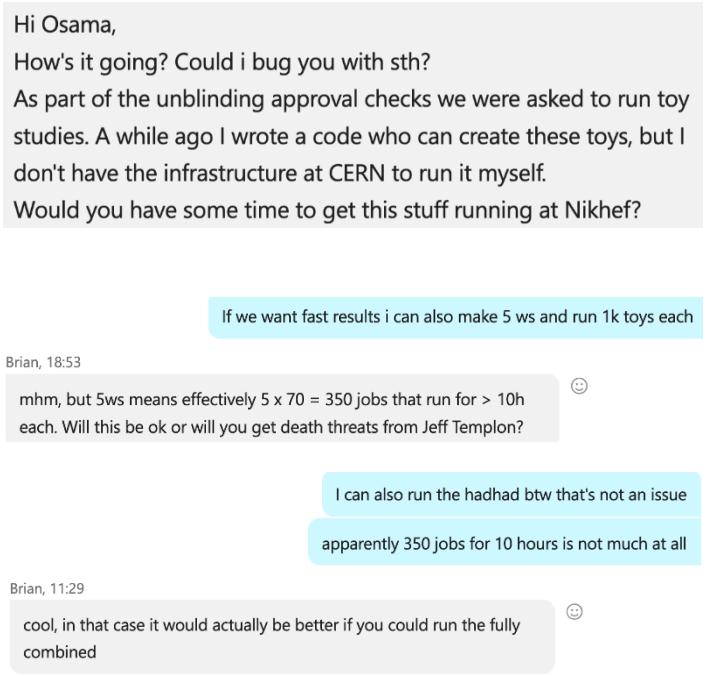

designed for data analysis:
economic & efficient

- /project/atlas (3Tb total space, 4Tb since this week)
  - Backed up daily: reliable, but expensive storage
  - NFS disk: slow (max speed 30% of network speed)
  - *Usage: only for code and sensitive information, not for bulk ntuples*
- /data/atlas (40Tb total space)
  - No backup: cheap larger-volume storage
  - NFS disk: slow (max speed 30% of network speed)
  - *Usage: Intended for bulk data that is not intensively analyzed*
- /dcache/atlas (350Tb total space (? TBC)
  - High Performance File system (masquarading as network file system)
  - Only accessible from stoomboot and Tier-1 computing facilities
  - Files can *not* be modified once written (but can be deleted, and recreated)
  - Not suitable, nor efficient for small files
  - *Usage: for storage of and intensive usages of ntuples, dAODs etc*

Neural Network training: ATLAS GNNs and Transformers

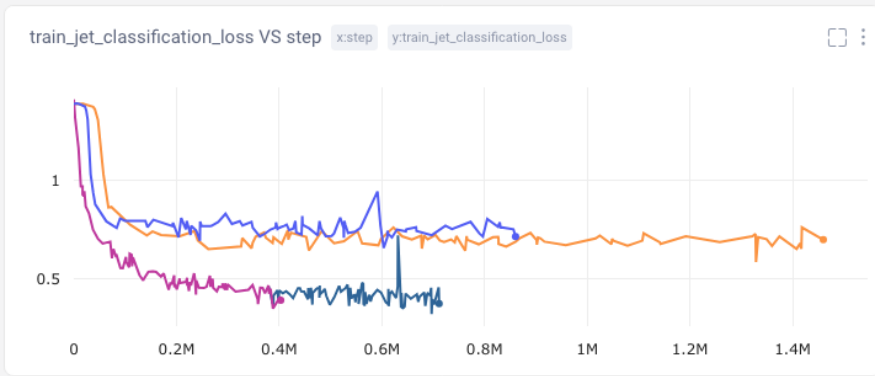

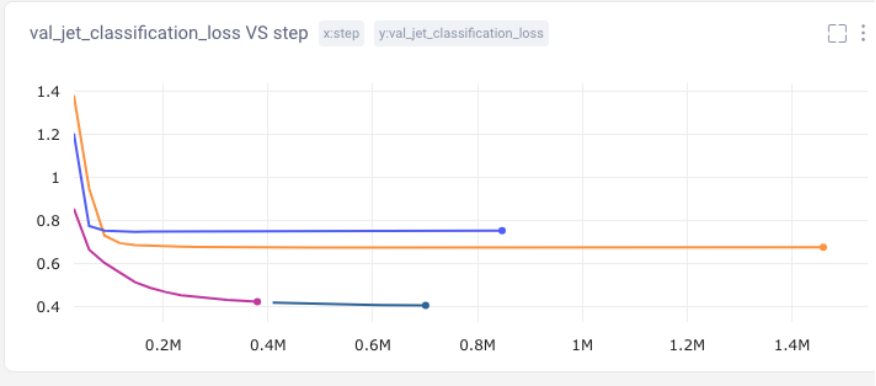

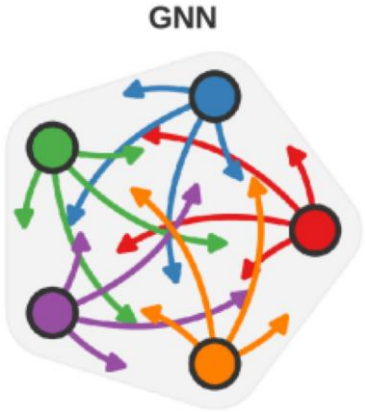


```
qsub -q gpu-nv -d /project/atlas/users/okarkout/salt/june_fork SubjetXbb_sub.sh
qsub -q gpu-nv -d /project/atlas/users/okarkout/salt/june_fork sub.sh
qsub -q gpu-nv -l pvmem=210gb -d /project/atlas/users/okarkout/salt/june_fork GN2X_bench_sub.sh

qstat -u okarkout
```

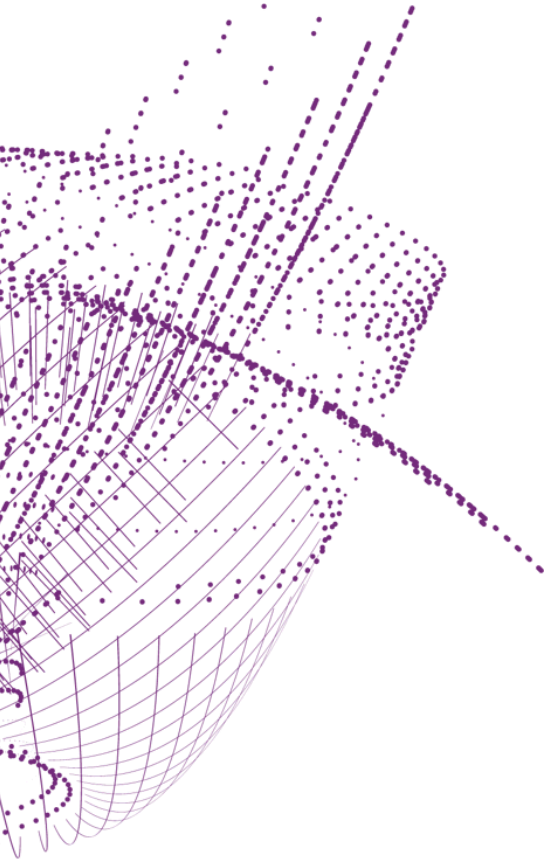

Arjen van Rijn
David Groep

Data centre visit and networks

Our datacenter is all about connectivity

PeeringDB

Search here for a network, IX, or facility.

Advanced Search Legacy Search

NIKHEF Amsterdam

Networks 315 Local Exchanges 20

Organization	Nikhef
Also Known As	
Long Name	
Company Website	http://www.nikhefhousing.nl
Address 1	Science Park 105
Address 2	
Floor	
Suite	
Location	Amsterdam, NH, 1098 XG
Country Code	NL
Continental Region	Europe

Networks

Peer Name A-Z

1-IX Route Servers
1Key
23M GmbH
2AT B.V.
2Hip Consultancy
31173 Services
A1 Internet B.V.
A2B Internet
Accenture Backbone

#7 on the worldwide peering list

Name	CLLI	City	State	Networks
Digital Realty Frankfurt FRA1-16 (Gold Sponsor)	FRNKGE	Frankfurt	Hessen	530
Digital Realty	-	DE	60314	
Telehouse - London (Docklands North) (Gold Sponsor)	LONDEN	London	-	464
Telehouse - Global Data Centers	-	GB	E14 2AA	
Equinix DC1-DC15, DC21 - Ashburn	ASBNVA	Ashburn	VA	444
Equinix, Inc.	703-723	US	20147-6205	
Equinix SP4 - São Paulo	-	Barueri	-	435
Equinix, Inc.	-	BR	06455000	
Equinix SG1 - Singapore	-	Singapore	-	428
Equinix, Inc.	-	SG	139964	
Equinix FR5 - Frankfurt, KleyerStrasse	FRNKGE	Frankfurt	-	400
Equinix, Inc.	-	DE	60326	
NIKHEF Amsterdam	AMSTNL	Amsterdam	NH	315
Nikhef	-	NL	1098 XG	
Telehouse - Paris 2 (Voltaire - Léon Frot) (Gold Sponsor)	PARSFR	Paris	-	309
Telehouse - Global Data Centers	-	FR	75011	

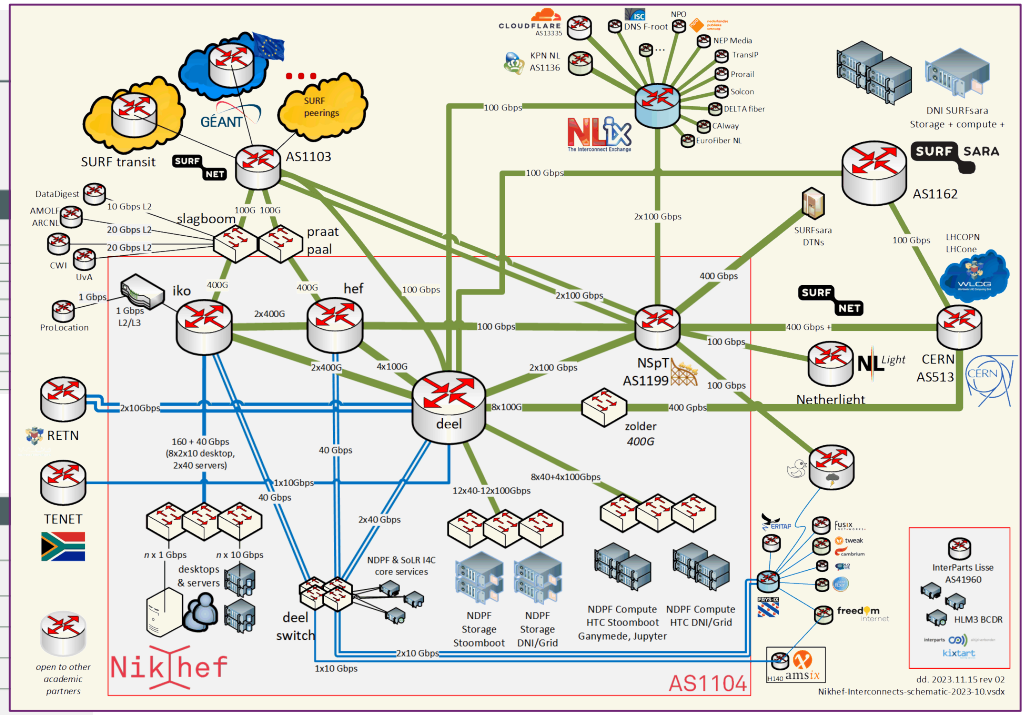

Nikhef

Summary data of our datacentre

Maximum power: 2,7 MW

Power Usage Efficiency (PUE): 1,3

Residual heat used for:
- Nikhef building
- Amsterdam University College
- Student housing

Room	Purpose	# racks max	IT power average per rack (kW)	IT power total (kW)	IT power currently (kW)
H234b	Scientific computing	47	6,25	300	180
H140	Nikhef Housing	282	1,9	660	528
H142	Nikhef Housing (extension)	112	4,0	540	10

Financial figures (annually):

- Turnover currently ~5 M€; will grow to about 6,5 M€;
- Running costs will grow to 3 – 3,5 M€ (energy costs!);
- Depreciation of extension: 1 M€ annually until end 2028;
- Net result: 2 – 3 M€

Because we can … does not mean it’s the scalable way ☺

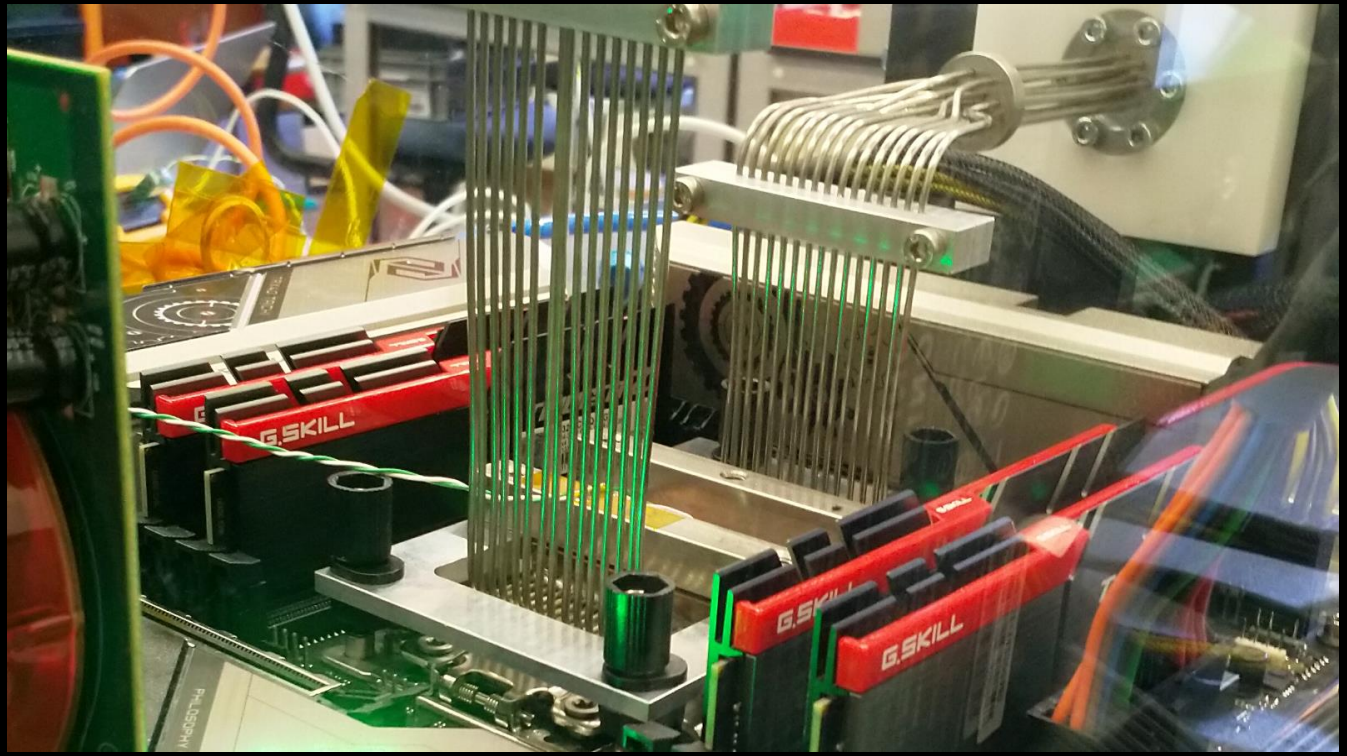

LCO2 cooling of an AMD Ryzen Threadripper 3970X [56.38 °C] at 4600.1MHz processor (~1.5x nominal speed) sustained, using the Nikhef LCO2 test bench system (https://hwbot.org/submission/4539341) - (Krista de Roo en Tristan Suerink)

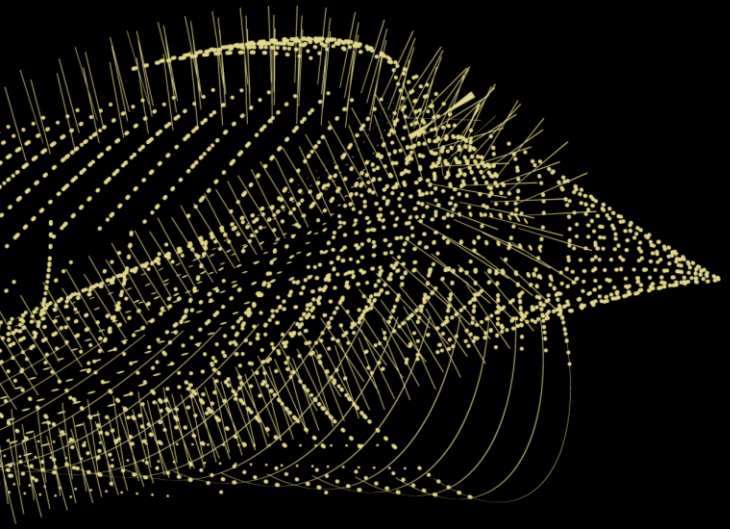

Supplementary materials

Backup – Balign

ATLAS Barrel Alignment

Hardware Setup (1)

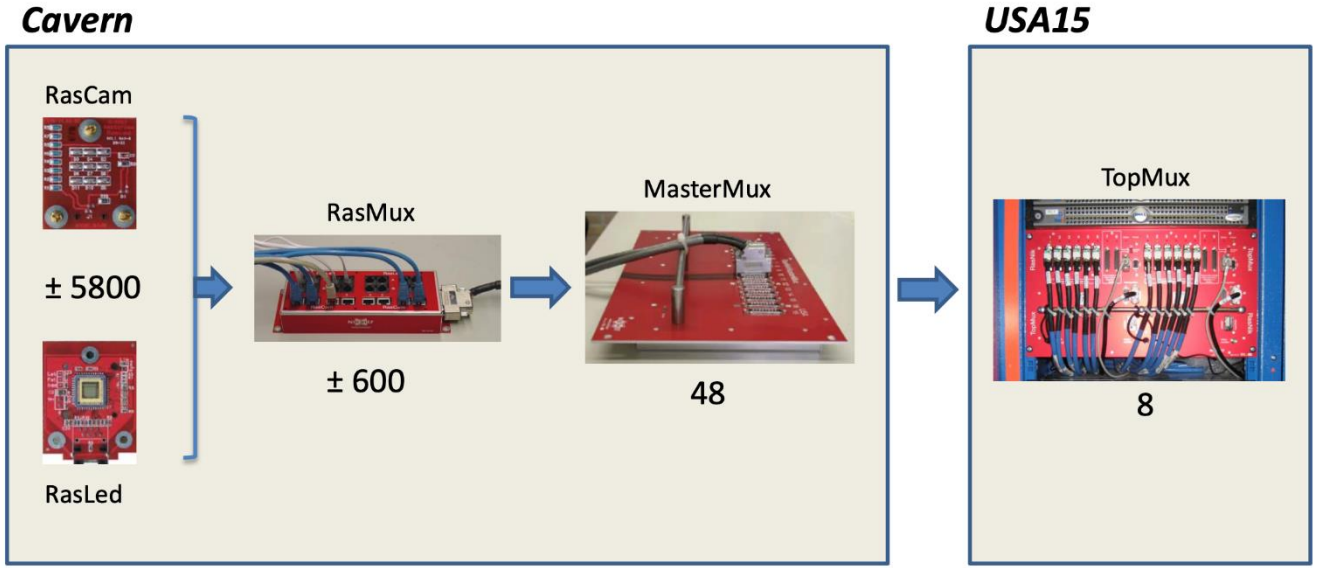

Source: Future Barrel Alignment by Robert Hart

Backup – MDM

1,171 chambers with total 354,240 tubes (3 cm diameter, 0.85-6.5 m long)

Tube resolution 80 μm

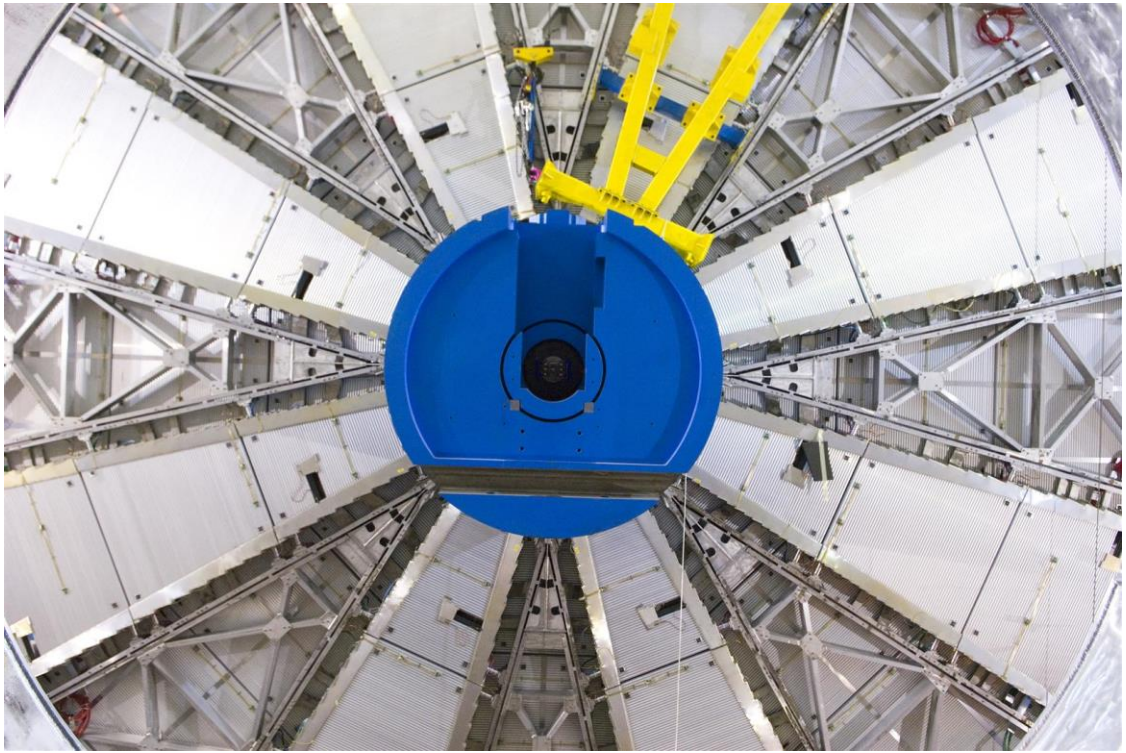

Source: https://atlas.cern/Discover/Detector/Muon-Spectrometer

Backup – SciFi

SciFi - Channels

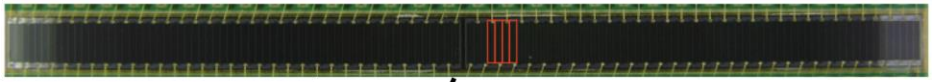

128-channel SiPM (Hamamatsu)

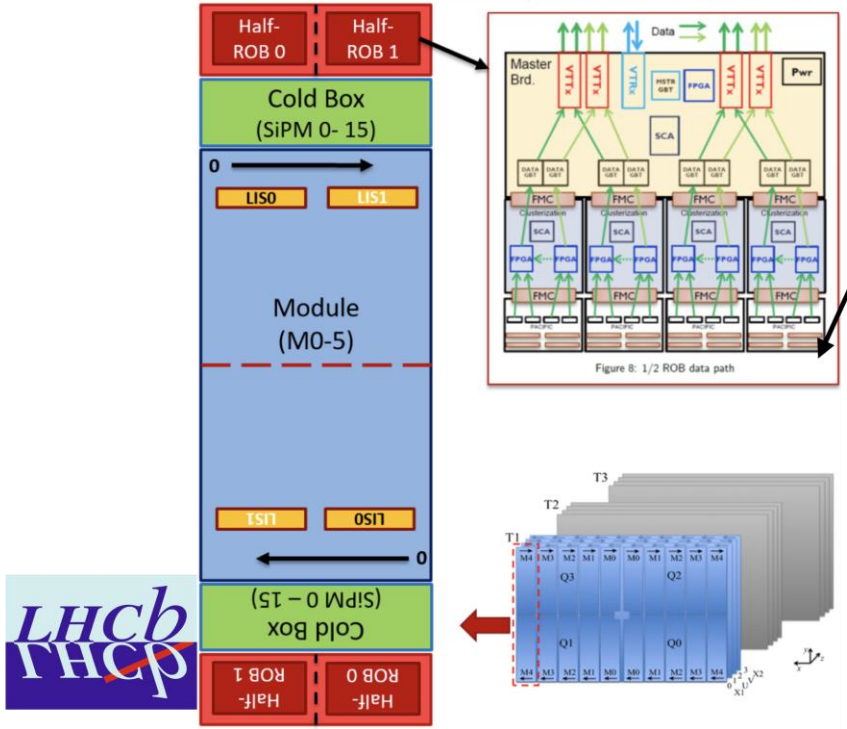

Figure 8: 1/2 ROB data path

512 HalfROBs:

- 512 MBs
- 2048 CBs
- 2048 PBs
- 8192 ASIC chips on PBs
- e.g. 524288 SiPM channels with 3 thresholds

Source: SciFi Collaboration