Prospects of FAIR computing

Computing at FAIR

What makes it so fascinating?

- FAIR houses **broad and divers physics communities** from atomic to particle physics, from theory to experiment, and very internationally oriented.
- Computational activities are **very diverse:** from notebooks to HPCs and HTCs, etc..
- **New developments** in hard- (GPU, ARM, QC,…), soft-ware (ML/AI), in data processing (triggerless readouts) and management (F.A.I.R., federated infrastructures).
- Computing support **centrally organised** & understaffed compared to HEP communities!
- **FAIR is not a smaller copy of CERN**, more subtle.

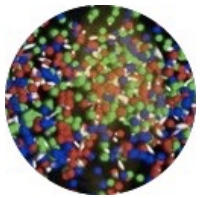

Hot and Dense Nuclear Matter

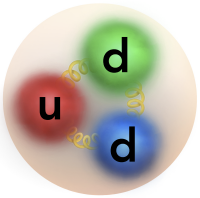

Hadrons

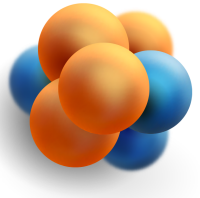

Atomic Nucleus

Nuclei in the Cosmos

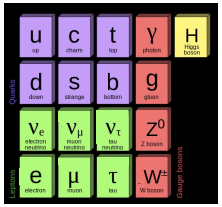

Fundamental Interactions

The objective

…yes a conceptual computing design for FAIR…

- … focussed towards research IT, hence not enterprise IT!
- … with a clear and coherent *vision* for FAIR computing
- … *supported* by relevant stakeholders including you!
- … with a description of *requirements* based on best *estimates*
- … with commonly defined *criteria*
- … FAIR players*: **APPA**,**CBM**,**HADES**,**NUSTAR**,**PANDA**,**THEORY**,**BEAM**
- … considering FAIR scenarios: **FS(+)** and **MSV**

*ALICE uses large fraction of computing resources & strong connections with scientific IT@GSI, but considered as “outside” activity, different funding scheme, etc.

The process

…to derive to FAIR computing model

- Central = the input and advice of the research lines, data management, and GSI-IT
- Collect info from FAIR Phase Zero
- Follow-up trends and strategies
- Participation EU communities, such as EOSC, NUPECC, JENA, …
- Regular bi-weekly meetings
- Deliver Conceptual Design Report

GSI IT:

Mohammad Al-Turany - m.al-turany@gsi.de
Thorsten Kollegger - T.Kollegger@gsi.de

THEORY:

Thomas Neff - t.neff@gsi.de

CBM:

Volker Friese - v.friese@gsi.de

HADES:

Jochen Markert - j.markert@gsi.de

PANDA:

Tobias Stockmanns - t.stockmanns@fz-juelich.de

APPA:

Shahab Sanjari - s.sanjari@gsi.de

NUSTAR:

Haik Simon - H.Simon@gsi.de
Bastian Loeher - B.Loeher@gsi.de
Stephane Pietri - S.Pietri@gsi.de

BEAM PHYSICS:

Adrian Oeftiger - a.oeftiger@gsi.de

RESEARCH DATA MANAGEMENT:

Andrew Mistry - A.K.Mistry@gsi.de

COORDINATION:

Johan Messchendorp - j.messchendorp@gsi.de

Classification

..as an important input to computing model

- **Computing (HEPSpec06)**
  - Online & offline data processing
  - Monte Carlo simulations
  - Theoretical models and simulations
- **Data storage (TB/year)**
  - Raw data
  - Derived data
  - Simulation data
- **Bandwidths (GB/s)**
  - Data rates from experiment to GreenCube
  - …to permanent storage
  - Peak & average rates

Classification

..as an important input to computing model

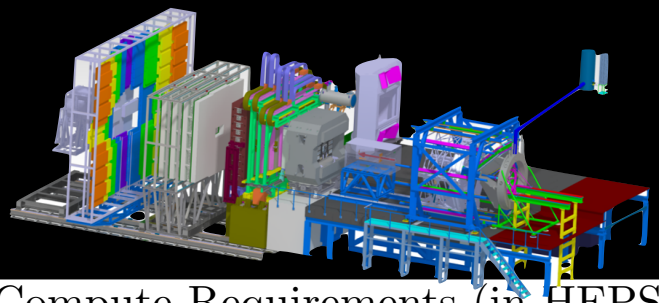

- **Computing (HEPSpec06)**
  - Online & offline data processing
  - Monte Carlo simulations
  - Theoretical models and simulations
- **Data storage (TB/year)**
  - Raw data
  - Derived data
  - Simulation data
- **Bandwidths (GB/s)**
  - Data rates from experiment to GreenCube
  - …to permanent storage
  - Peak & average rates

Table 20: CBM Compute Requirements (in HEPS

		Compute Class	FS+	MSV
Entry Cluster (Exp clusters)		I.a	6,000	6,000
		I.b	0	0
		I.c	0	0
		I.d	0	0
Offline computing		II.a	1,000,000	1,000,000
Online computing		II.b	980,000	980,000

Shared compute clusters

1 physical core (Intel E5-2680v4@2.4GHz) ~22 HEPSpec06

Classification

..as an important input to computing model

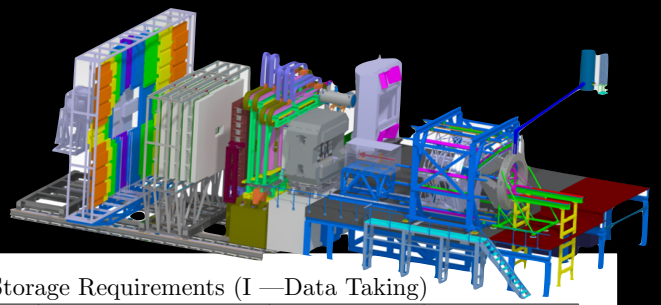

- **Computing (HEPSpec06)**
  - Online & offline data processing
  - Monte Carlo simulations
  - Theoretical models and simulations
- **Data storage (TB/year)**
  - Raw data
  - Derived data
  - Simulation data
- **Bandwidths (GB/s)**
  - Data rates from experiment to GreenCub
  - …to permanent storage
  - Peak & average rates

Table 21: CBM Storage Requirements (I —Data Taking)

		FS+	MSV
Experiment to GreenCube/RZ1	#fibers	120	120
	Bandwidth (GB/s)	400	400
Bandwidth to permanent storage	Peak (GB/s)	8	8
	Average (GB/s)	5.4	5.4
Permanent storage/year (TB/year)		18,000	18,000
Additional disk storage (TB)		11,000	11,000

Table 22: CBM Storage Requirements (II —Processing)

		FS+	MSV
Raw Data	TB/year	18,000	18,000
	#years	2	2
	Bandwidth (MB/s)	0	0
Simulation	TB/year	9,000	9,000
	#years	4	4
	Bandwidth (MB/s)	0	0
Derived data	TB/year	4,000	4,000
	#years	5	5
	Bandwidth (MB/s)	0	0

FAIR *Phase Zero*

Looking back: compute@Virgo

- Virgo statistics integrated from 1/2021- 3/2023
  ~82 kcores years
- Involves about 50 kcores
  ~1 MHSP06
- FAIR Phase Zero compute:
  ~60%
  ~0.5 MHSP06
- FAIR Phase Zero max disk usage: ~10 PB (2023)

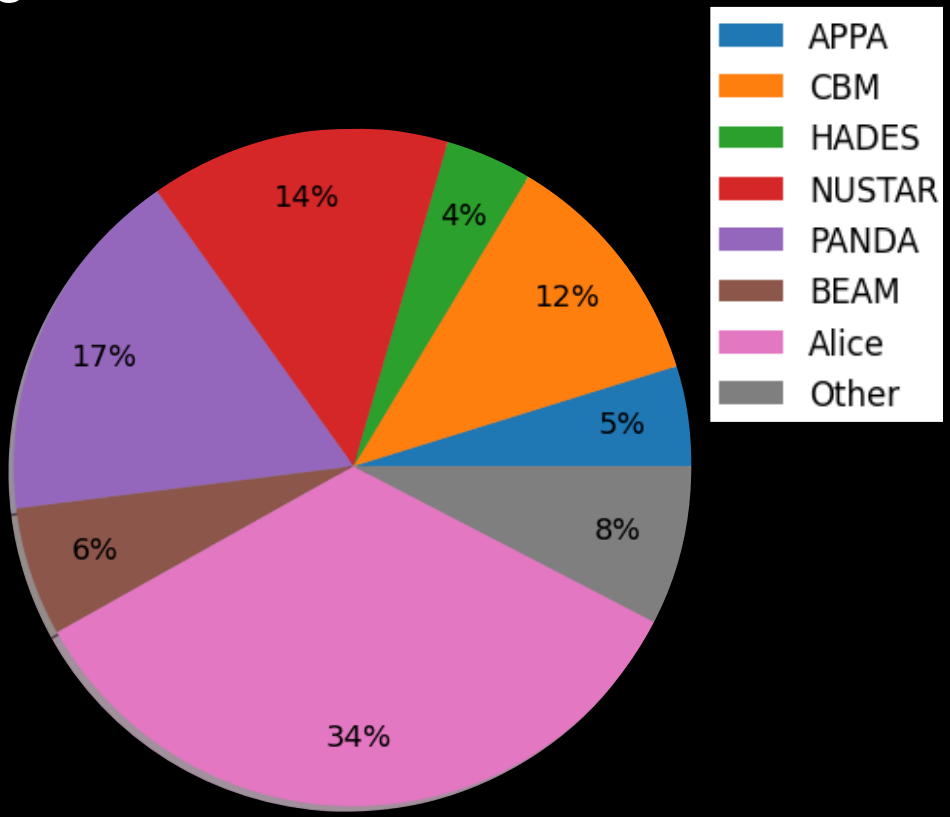


“FAIR2028”

FAIR2028

Compute requirements for *FS+*

- FS would be *without* CBM
- Other non-FAIR activities not accounted for (e.g. ALICE)
- All pillars with strong computing requirements, CBM largest online requirements on shared system
- The offline computations have a dominant data-independent component for NUSTAR & PANDA

*Note: theory requirements are included

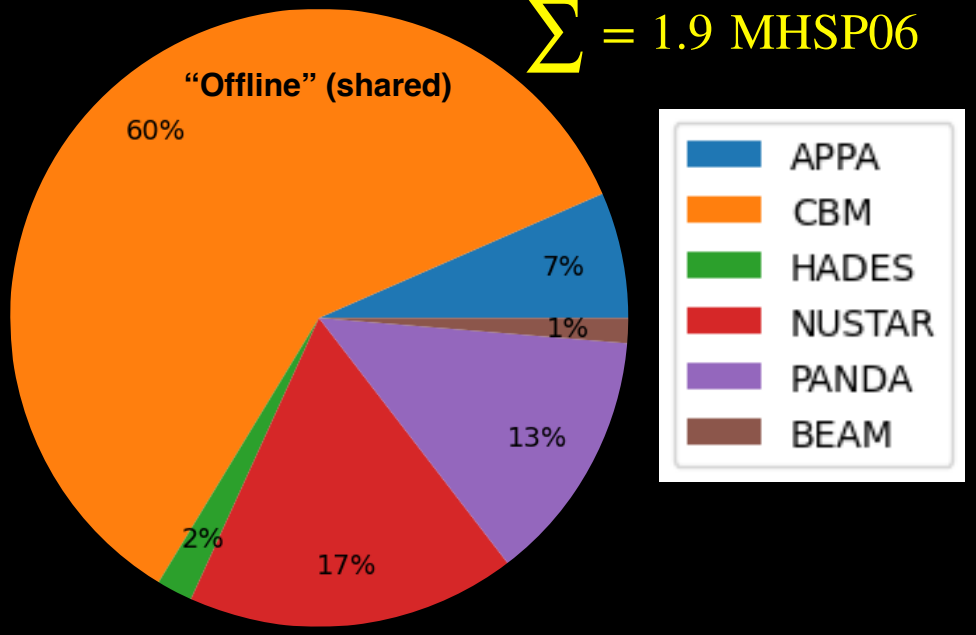

$\sum$ = 1.9 MHSP06

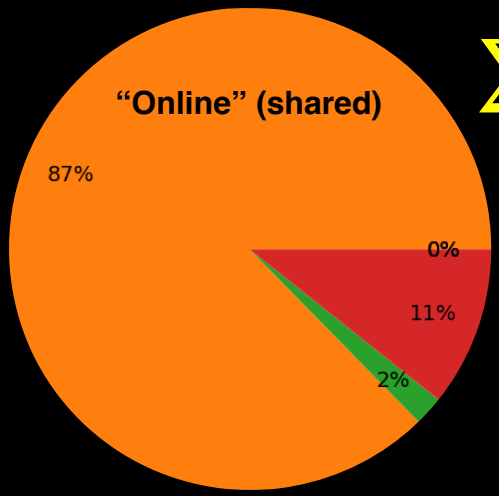

$\sum$ = 1.1 MHSP06

FAIR2032

Compute requirements for *MSVc*

- Required compute resource scale up by factor 2 with respect to FS+
- Other non-FAIR activities not accounted for (e.g. ALICE)
- CBM & PANDA largest contribution for both on- and offline computations
- The offline computations have a dominant data-independent component for NUSTAR

*Note: theory requirements are included

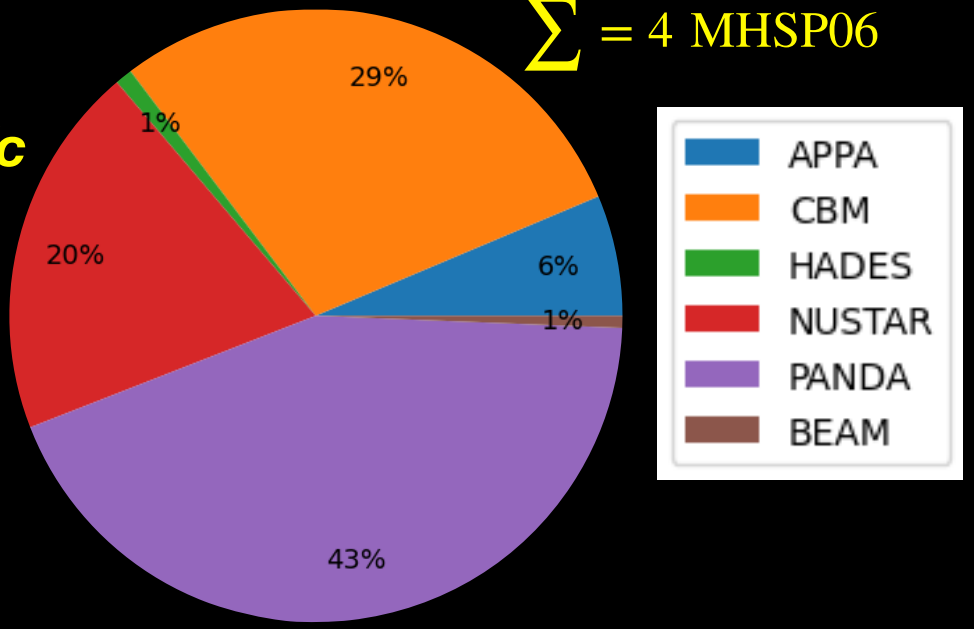

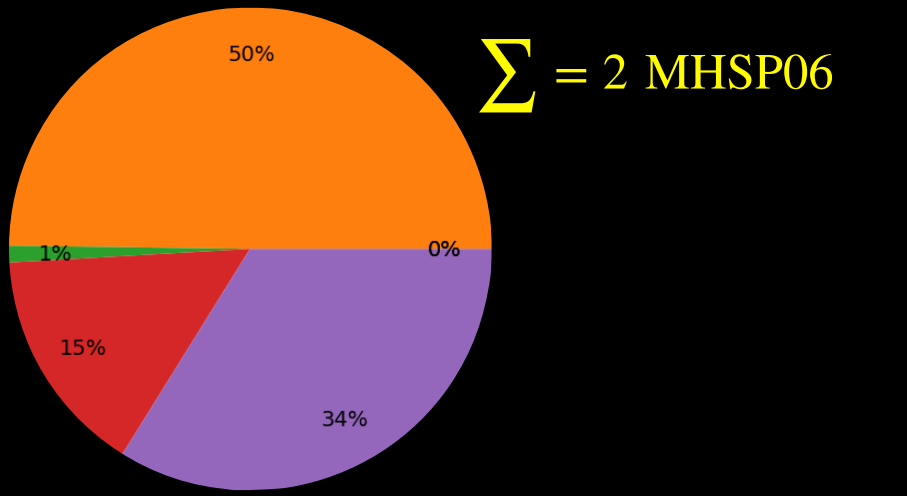

FAIR2028/32

Storage requirements

- *Excluding* pledges for Alice and other non-FAIR related storage
- The legal policy of keeping at least a copy of data elsewhere not included
- IT purchases 1-2 years before

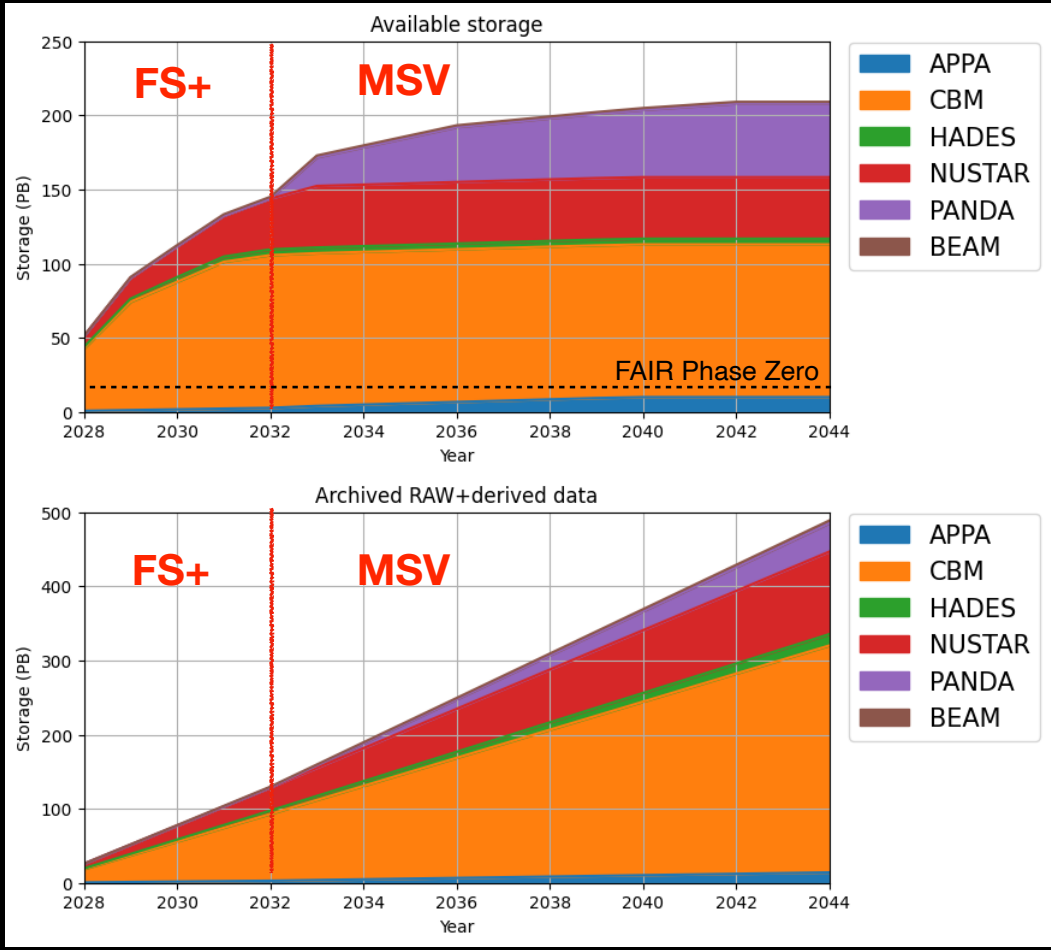

FAIR2028

Perspectives for a nominal FS+ year

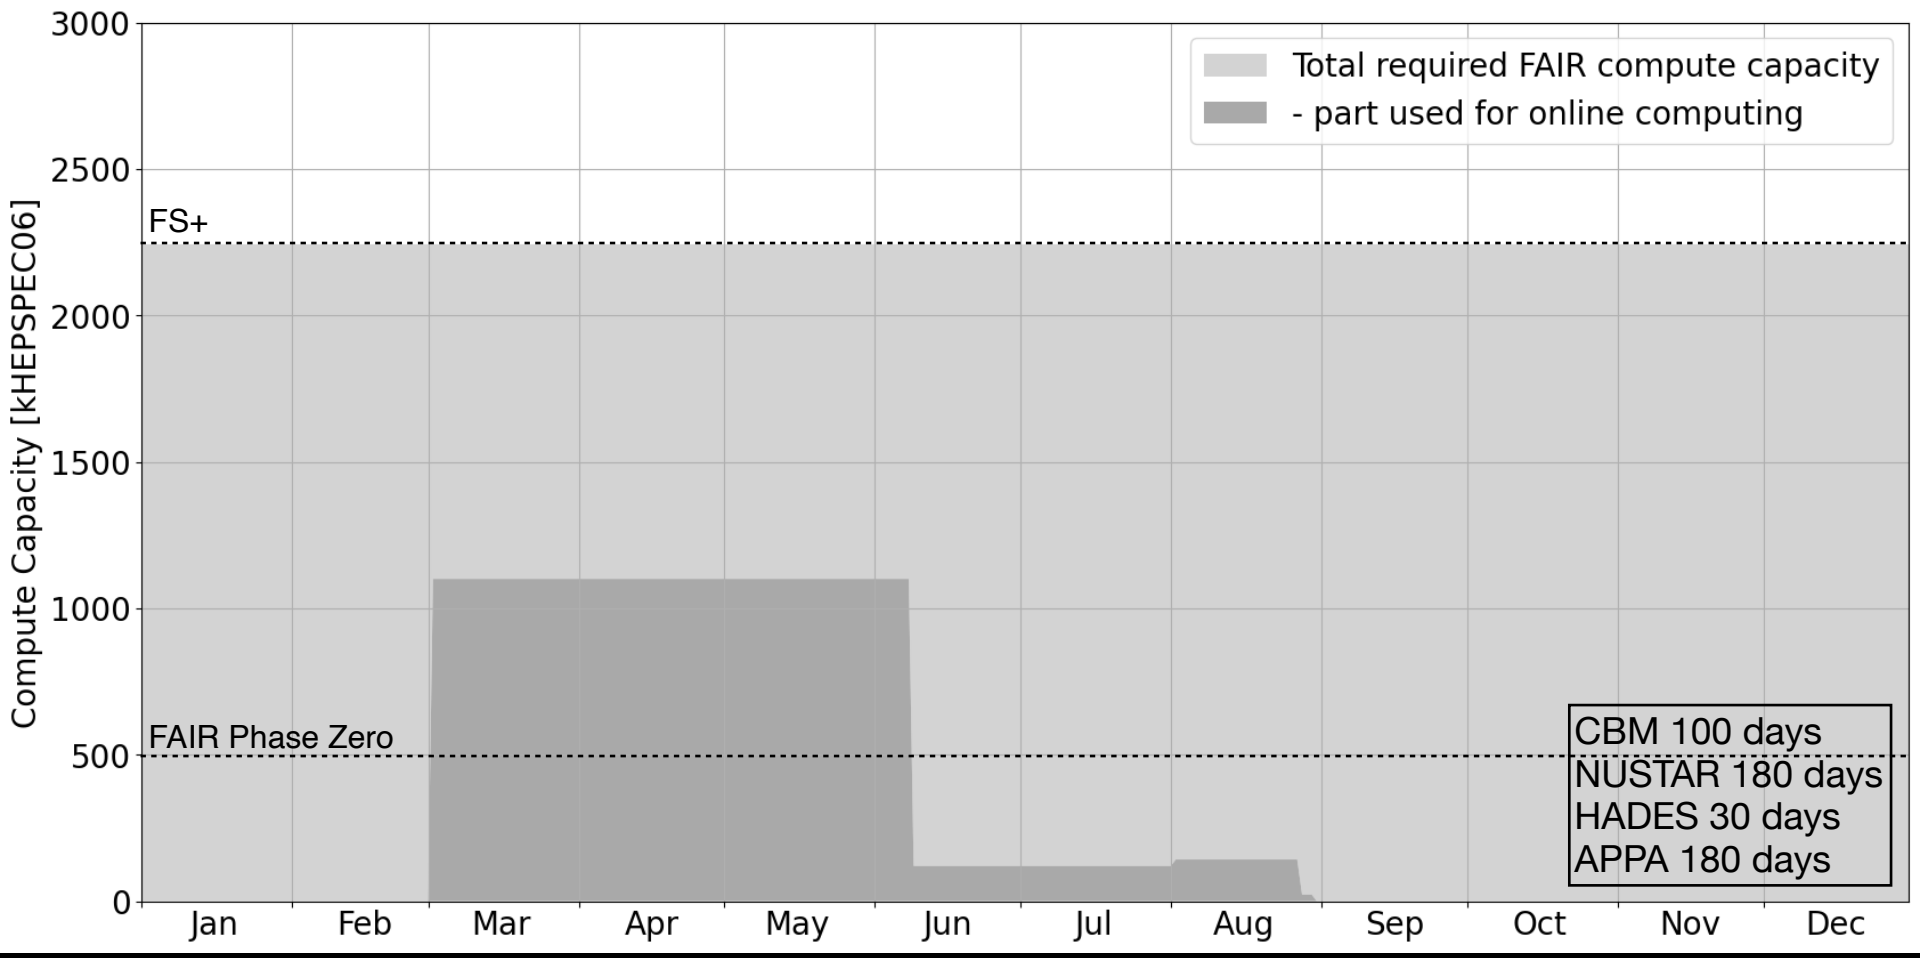

FAIR2028
Perspectives for a nominal FS+ year

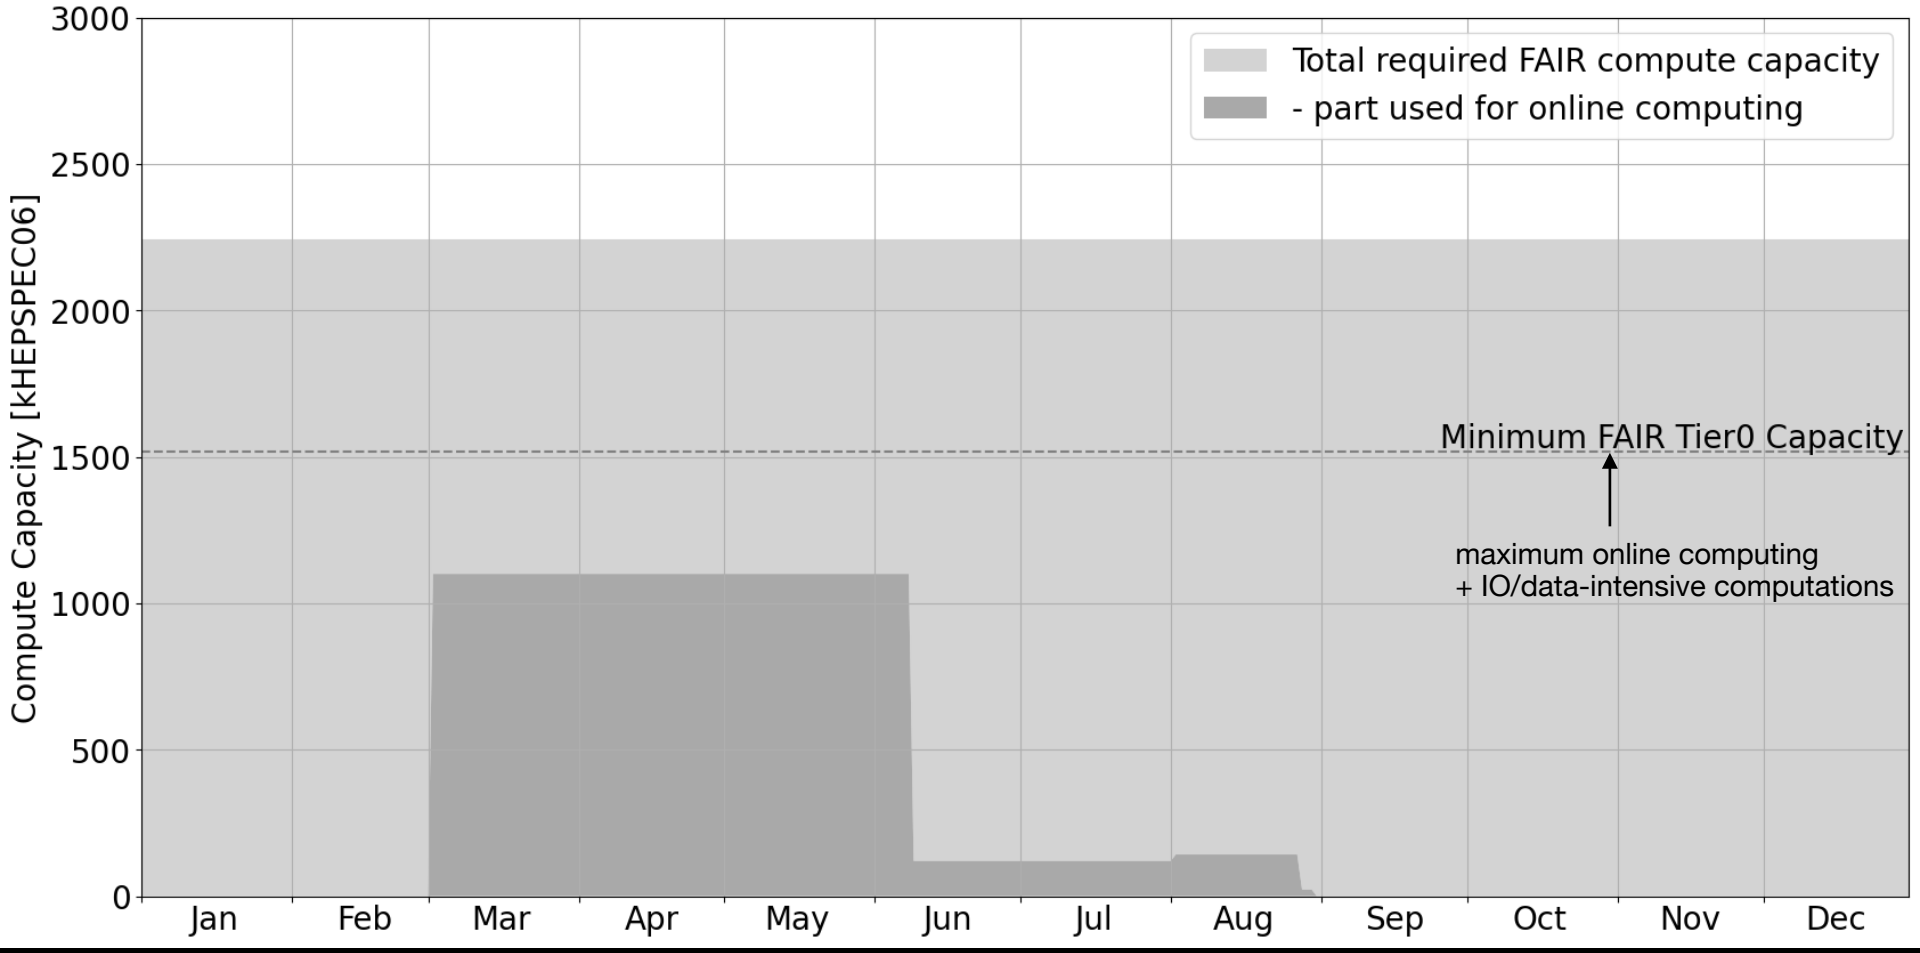

FAIR2028
Perspectives for a nominal FS+ year

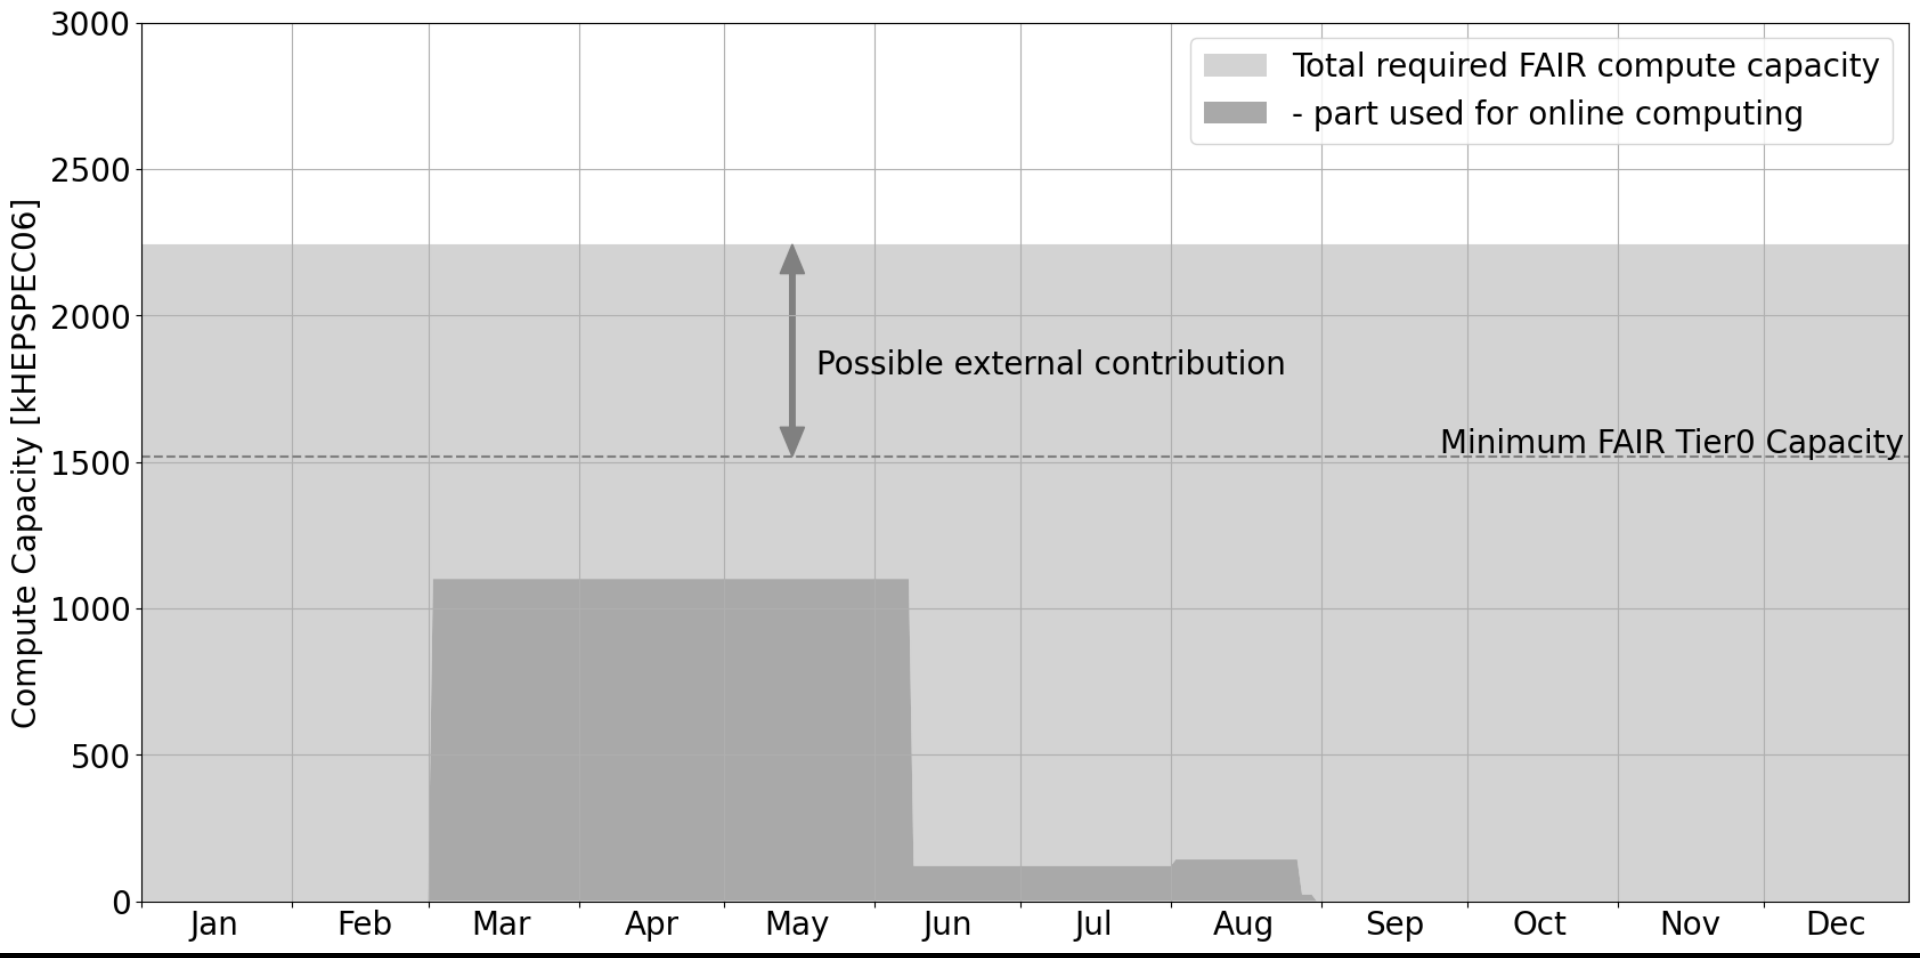

FAIR2032

Perspectives for a nominal MSV year - PANDA CBM parallel

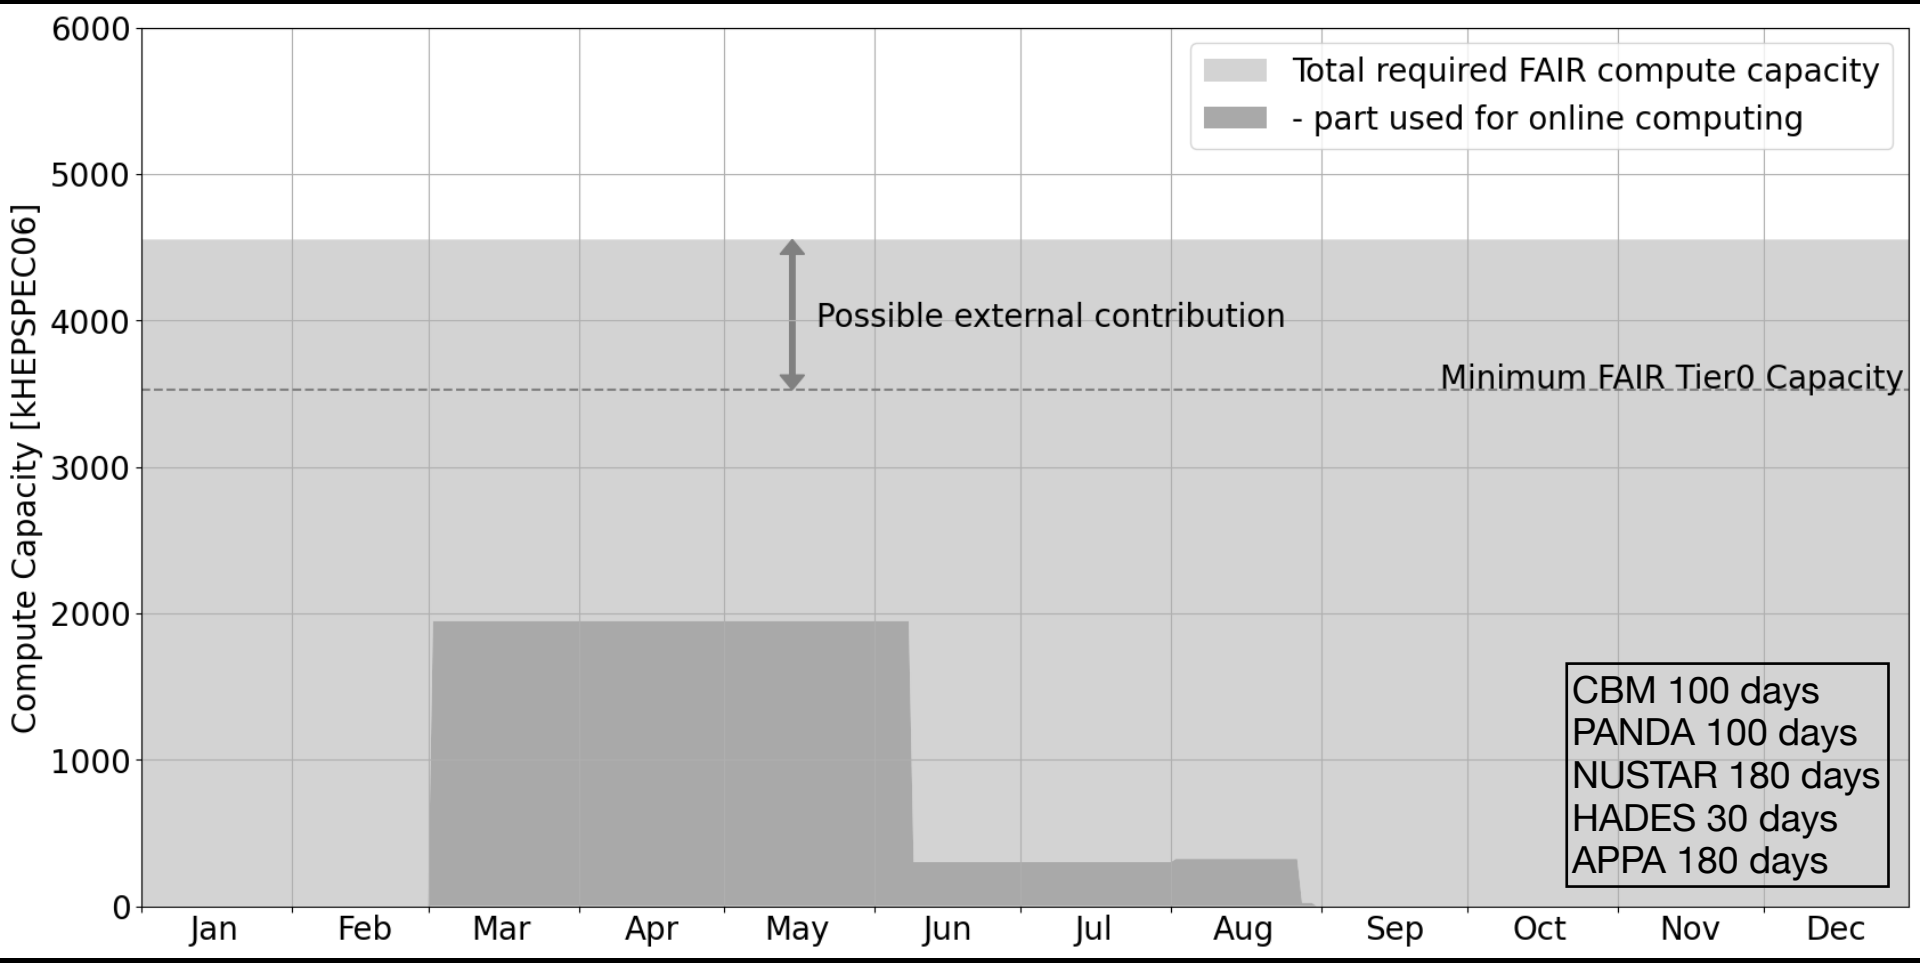

FAIR2032

Perspectives for a nominal MSV year - PANDA CBM parallel

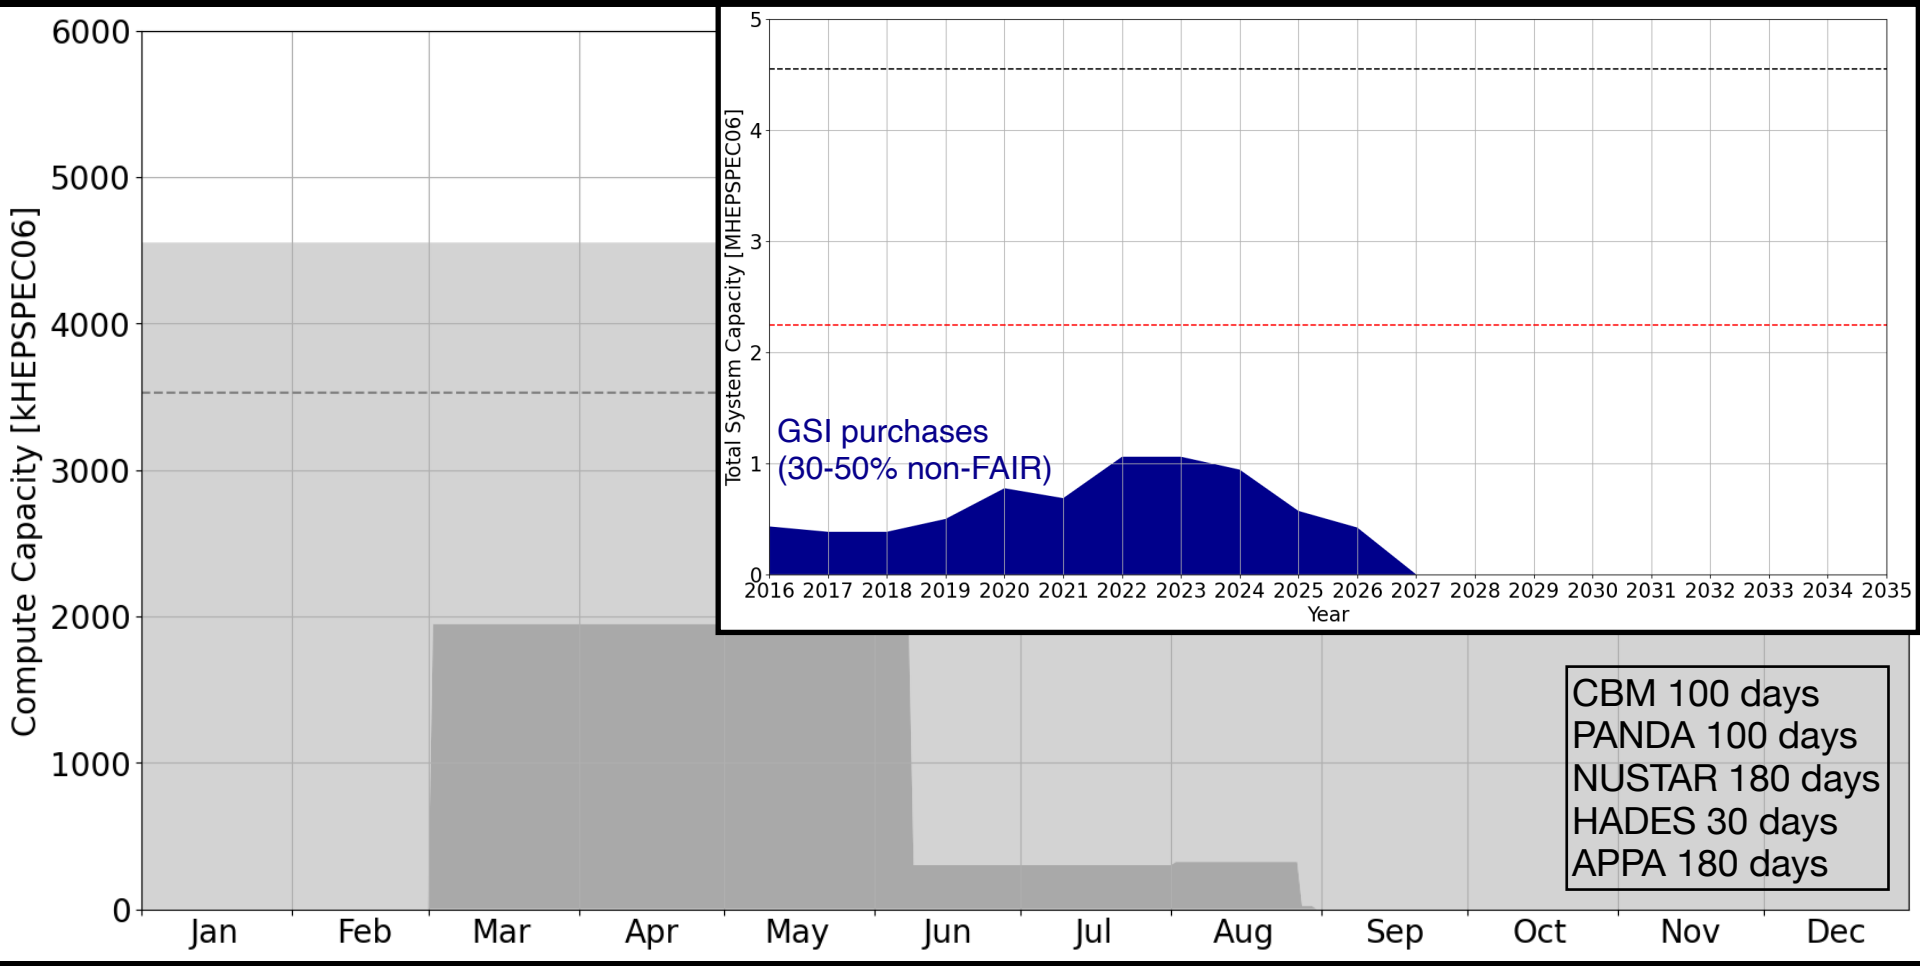

FAIR2032

Perspectives for a nominal MSV year - PANDA CBM parallel

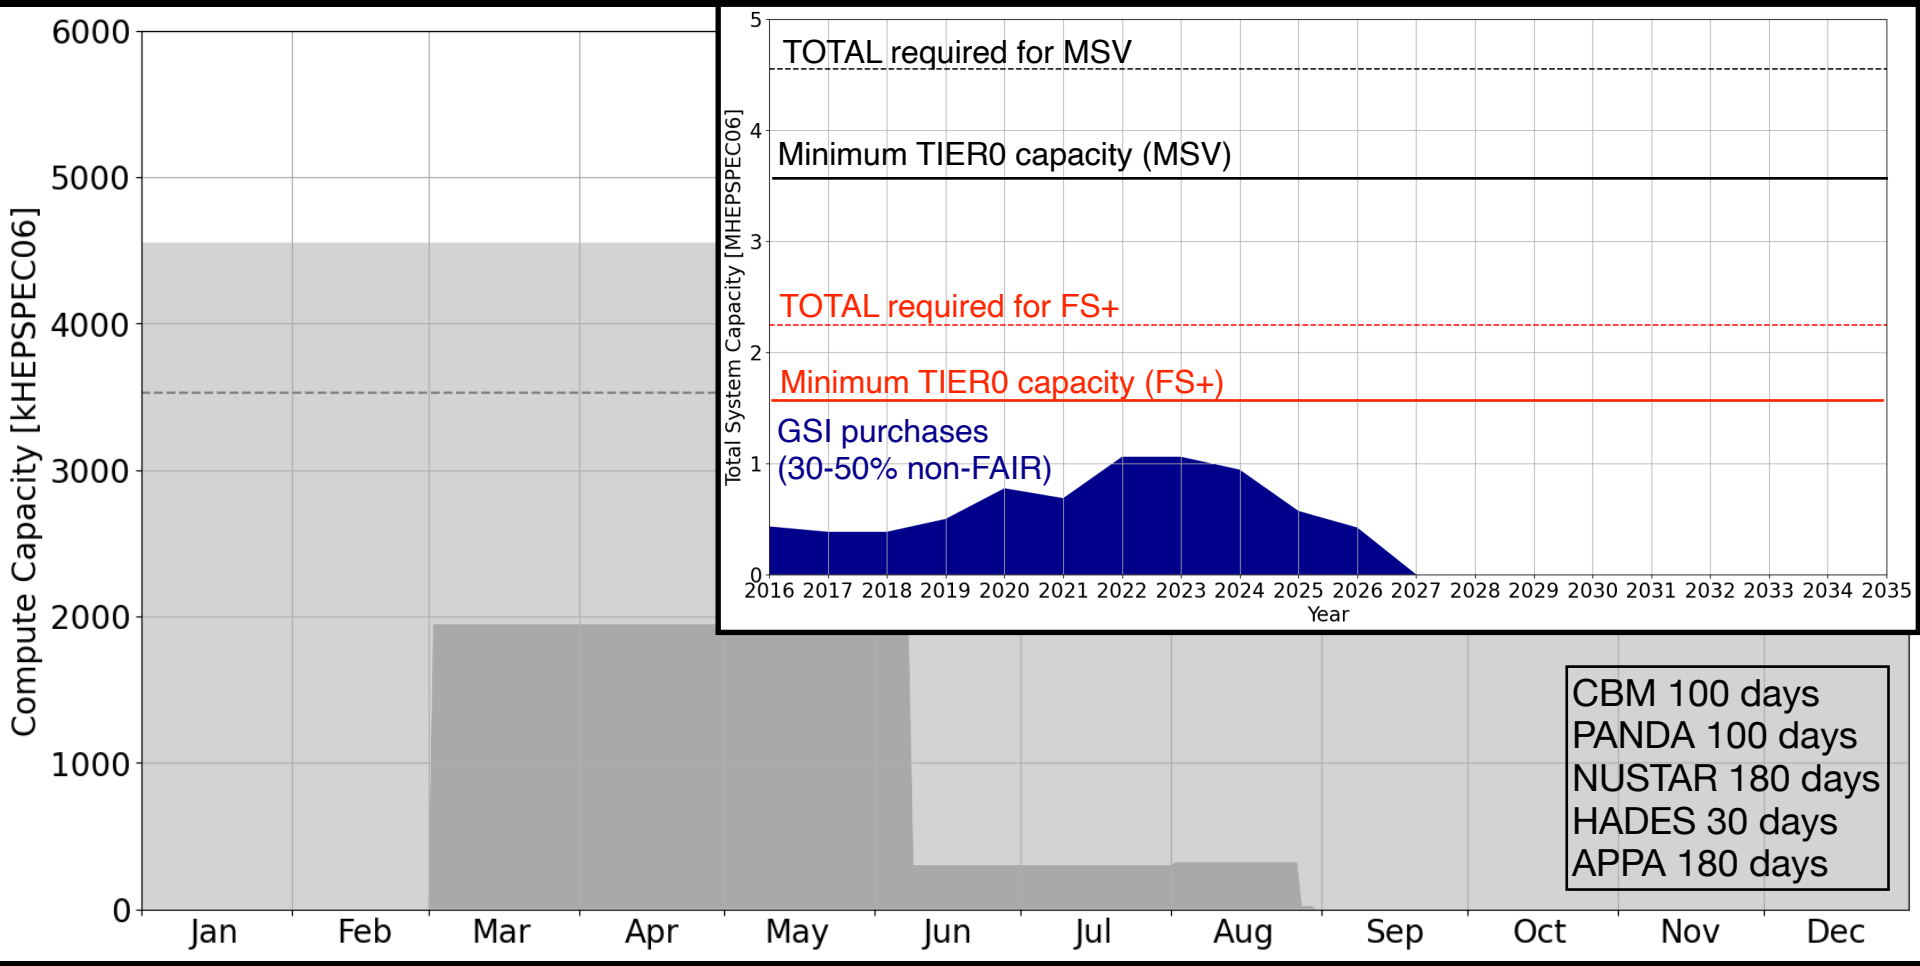

FAIR2032

Perspectives for a nominal MSV year - PANDA CBM parallel

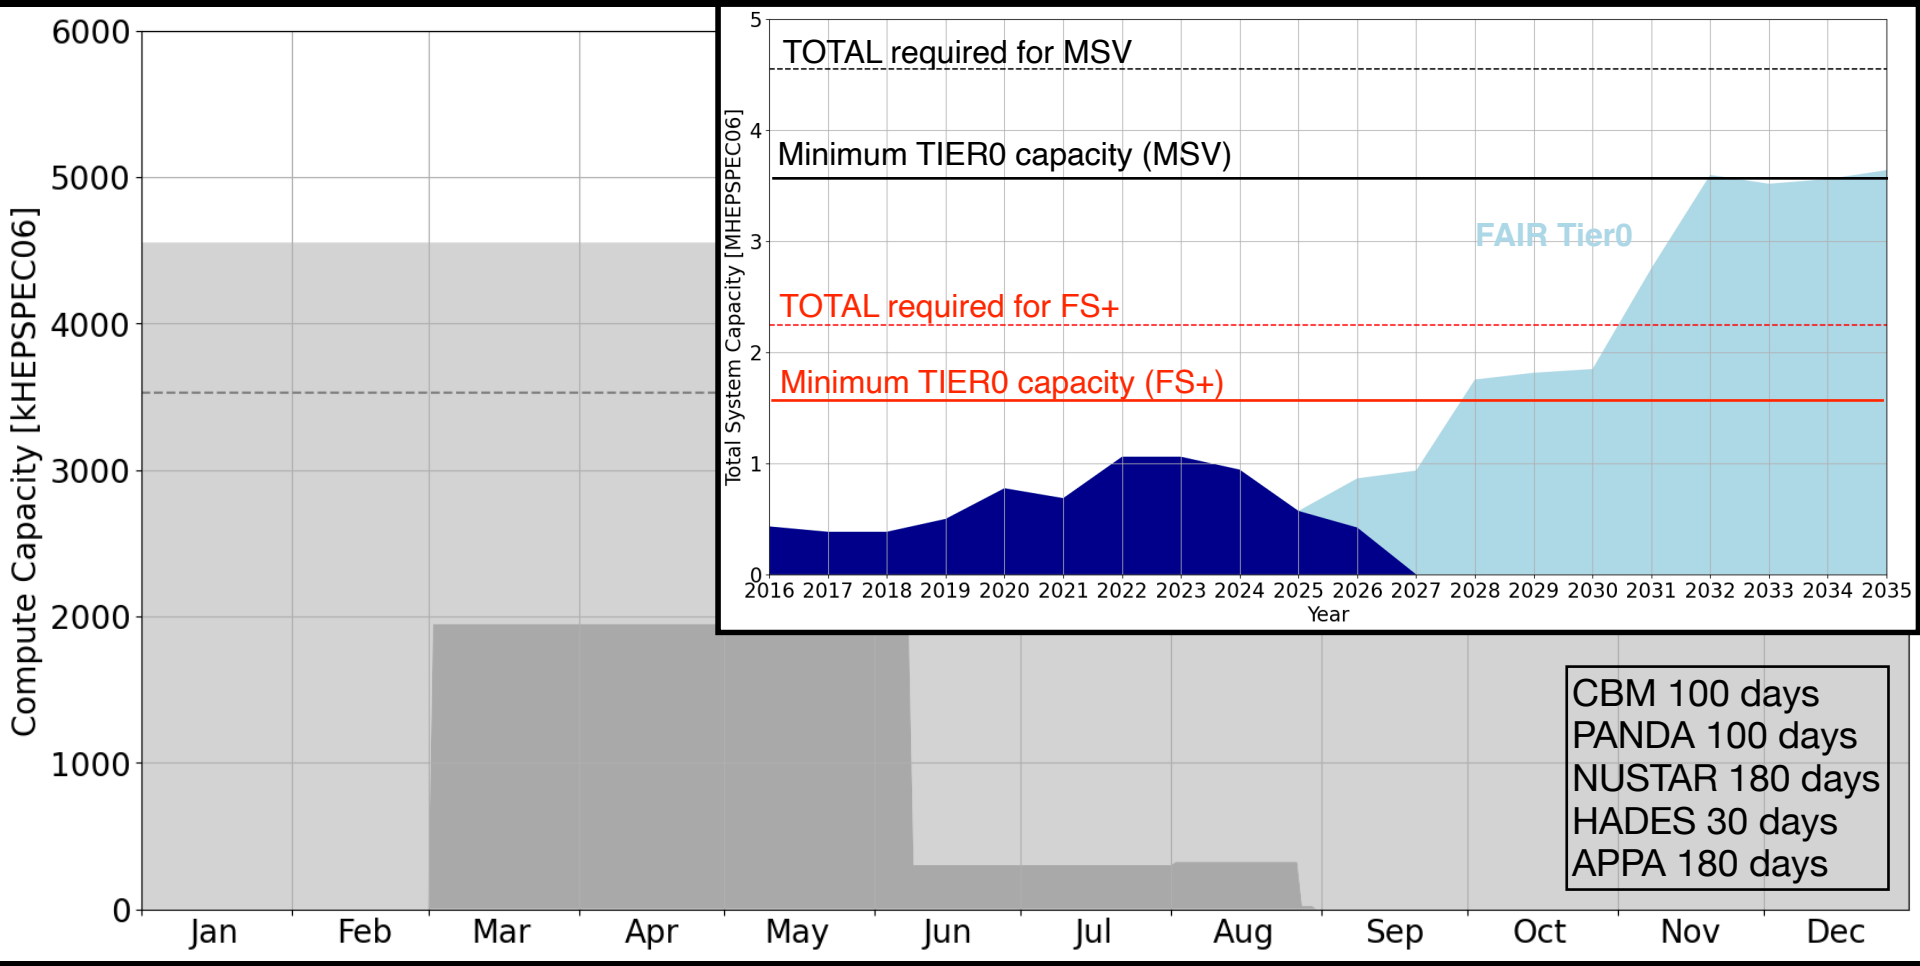


FAIR


FAIR
Findable
Accessible
Interoperable
Reusable

- Involved in ESCAPE, PUNCH4NDFI, EuroLabs
- Observer in EOSC

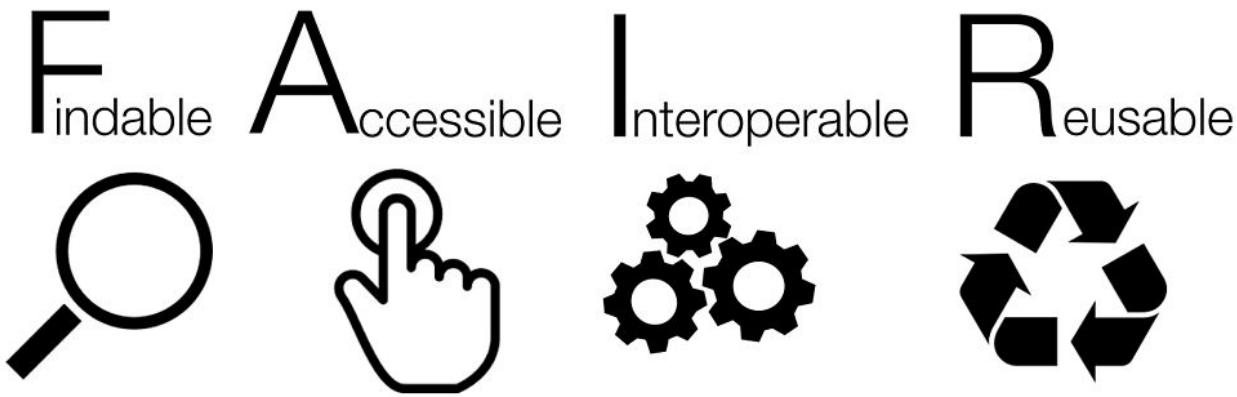

FAIR

- Involved in ESCAPE, PUNCH4NDFI, EuroLabs
- Observer in EOSC

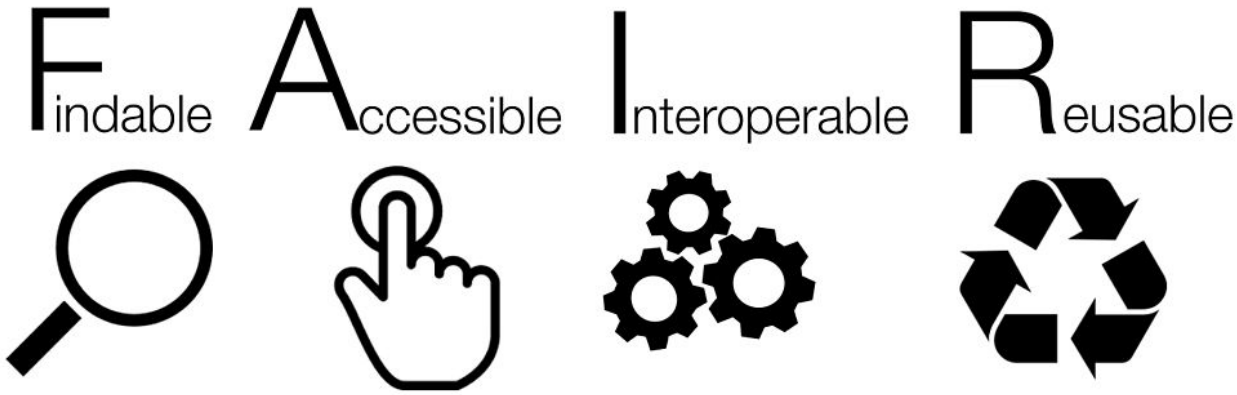

- FAIR will be aligned with the open-science policies of GSI (Andrew Mistry et al.)
- F.A.I.R. policies to minimum basics that actually can be effectively realised!

FAIR principle for FAIR CDR?

…a couple of *fair* conceptual requirements

FAIR principle for FAIR CDR?

…a couple of *fair* conceptual requirements

Accessible

- Data and software produced/dedicated for FAIR communities and publications centrally stored.
- Common & “user-friendly” *interface* to store and retrieve data.
- https access (http->file), token-based authentication, xrootd frontend+lustre backend, eduGAIN.
- Level of openness (what? whom?) defined by each collaborations
- Technology followed-up within ESCAPE, e.g. open-source scientific software and service repository (OSSR) and implementation of “data lake”.

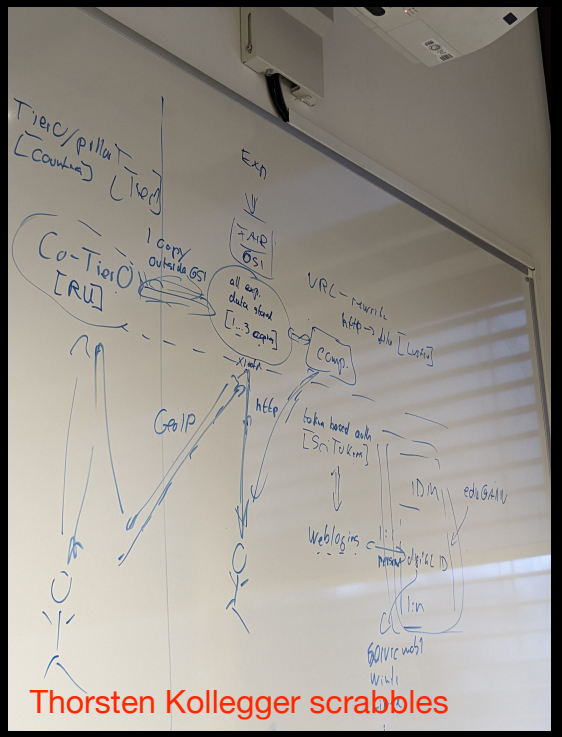

Thorsten Kollegger scrabbles

FAIR principle for FAIR CDR?

…a couple of *fair* conceptual requirements

Findable

- Centrally orchestrated storage and access of data essential to enable the data/software to become findable.
- Usage of Persistent IDentifiers (PID), Digital Object Identifiers (DOI).
- Coupling to GATE.

Accessible

- Data and software produced/dedicated for FAIR communities and publications centrally stored.
- Common & “user-friendly” *interface* to store and retrieve data.
- https access (http->file), token-based authentication, xrootd frontend+lustre backend, eduGAIN.
- Level of openness (what? whom?) defined by each collaborations
- Technology followed-up within ESCAPE, e.g. open-source scientific software and service repository (OSSR) and implementation of “data lake”.

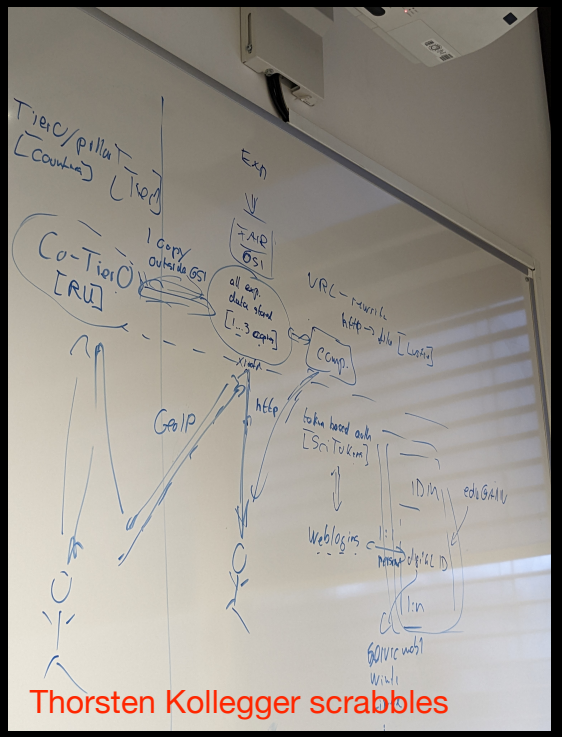

Thorsten Kollegger scrabbles

FAIR principle for FAIR CDR?

…a couple of *fair* conceptual requirements

Findable

- Centrally orchestrated storage and access of data essential to enable the data/software to become findable.
- Usage of Persistent IDentifiers (PID), Digital Object Identifiers (DOI).
- Coupling to GATE.

Accessible

- Data and software produced/dedicated for FAIR communities and publications centrally stored.
- Common & “user-friendly” *interface* to store and retrieve data.
- https access (http->file), token-based authentication, xrootd frontend+lustre backend, eduGAIN.
- Level of openness (what? whom?) defined by each collaborations
- Technology followed-up within ESCAPE, e.g. open-source scientific software and service repository (OSSR) and implementation of “data lake”.

Interoperable

- Agree on metadata format etc.
- FAIR-produced data operable with other research fields? Recommendation: do not make it a general policy, finer research-specific granularity (ESCAPE) and identify use cases.

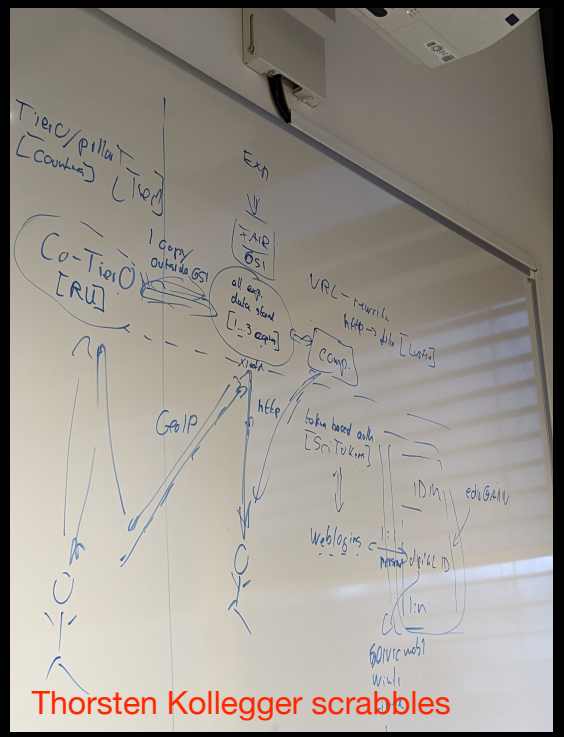

Thorsten Kollegger scrabbles

FAIR principle for FAIR CDR?

…a couple of *fair* conceptual requirements

Findable

- Centrally orchestrated storage and access of data essential to enable the data/software to become findable.
- Usage of Persistent IDentifiers (PID), Digital Object Identifiers (DOI).
- Coupling to GATE.

Accessible

- Data and software produced/dedicated for FAIR communities and publications centrally stored.
- Common & “user-friendly” *interface* to store and retrieve data.
- https access (http->file), token-based authentication, xrootd frontend+lustre backend, eduGAIN.
- Level of openness (what? whom?) defined by each collaborations
- Technology followed-up within ESCAPE, e.g. open-source scientific software and service repository (OSSR) and implementation of “data lake”.

Interoperable

- Agree on metadata format etc.
- FAIR-produced data operable with other research fields? Recommendation: do not make it a general policy, finer research-specific granularity (ESCAPE) and identify use cases.

Reusable

- Follows naturally once “FAI” policies are in place.

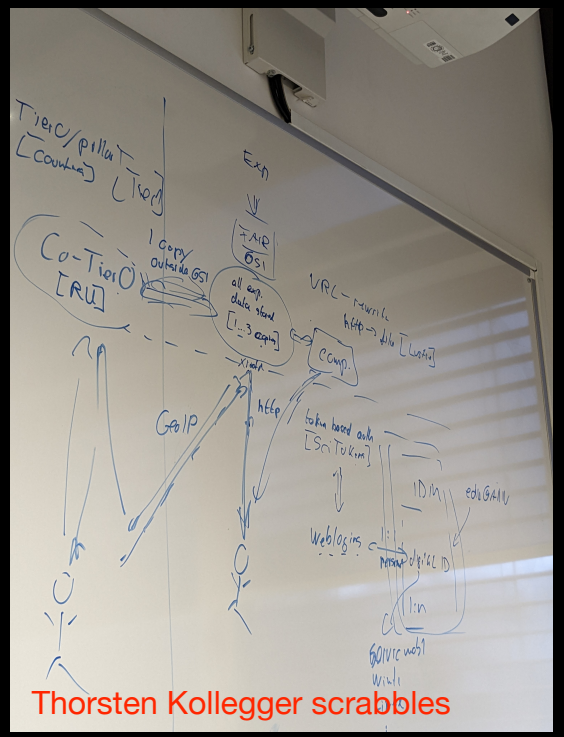

Thorsten Kollegger scrabbles

FAIR-IT support vs research lines responsibilities

…sensitive topic, hence very relevant!

FAIR-IT support vs research lines responsibilities

…sensitive topic, hence very relevant!

FAIR-IT support for research

FAIR-IT support vs research lines responsibilities

…sensitive topic, hence very relevant!

FAIR-IT support for research

- Responsibility: "*at the end of the fibres from the experiment*".
- Define and setup interfaces between experiment/user and compute/storage.
- Promote as much as reasonably acceptable common interfaces, hard/software infrastructures etc.
- Provide VMs, cloud service to minimise "idle" computers.
- Support commonly-used services/frameworks, e.g. Fairroot, FairMQ, CDash, Gitlab, …
- Maintain a strong local scientifically-based IT team well integrated within the various experiments with network/ interface to experts outside FAIR (f.e. GEANT, ROOT, …).

R&D aspects to investigate/follow-up

…that potentially reduce costs, provide more physics output…

- Evaluate applications of ML/AI, *e.g.* smart experiment control and (online) event processing, smart simulations.
- Deployment/benchmarking of algorithms on accelerator cards, ARM, FPGA, QC (long term), …
- Developments in ESCAPE, PUNCH4NFDI, *i.e.* domain-specific initiatives.

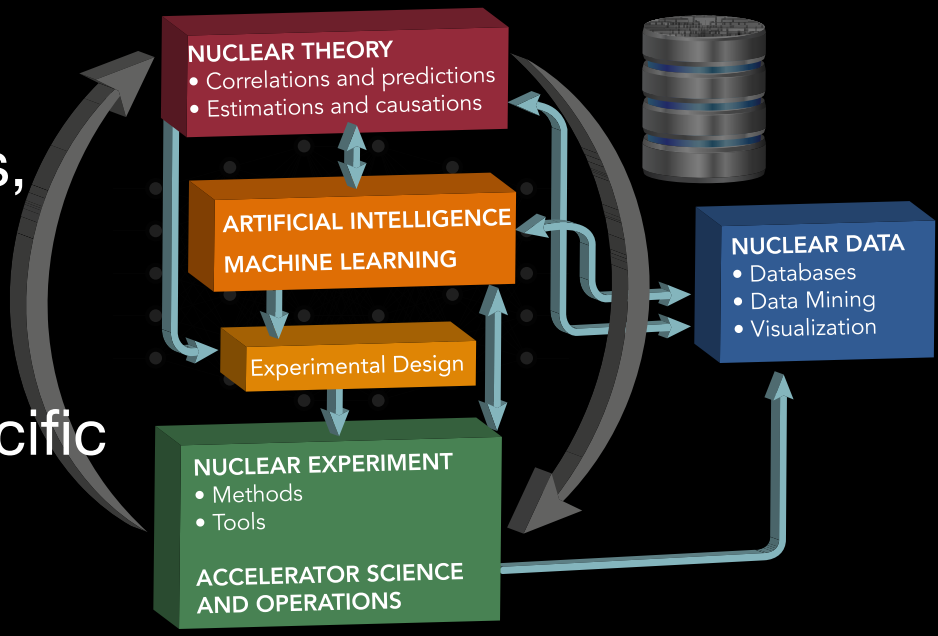

FAIR computing model in a *nutshell*

- *Grid-like distributed* computing is *dead*, long-live federated computing among large centers.
- Effective resource sharing at FAIR TIER-0 center, accounting for most of the data-driven computations.
- Federated storage and computing with 'local' centers using Teralink network & commonly used standards.
- 'Centralised' orchestration —> most suited to incorporate F.A.I.R. principle for our diverse community, introduce and to minimize the operational overhead.
- Containerised approaches and other virtualisation methods for flexible compute operations serving diverse community & optimise usage.
- Data access using http, possibly with xrootd frontend, lustre backend; AAI using widely accepted standards, weblogin, token-based, eduGAIN.

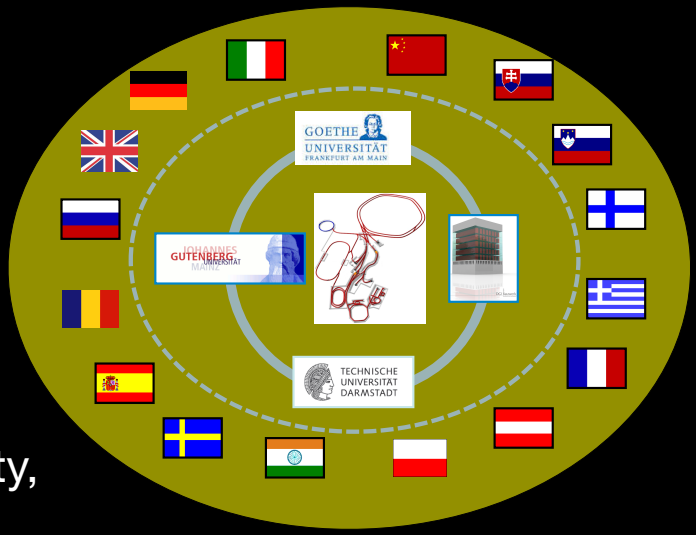

FAIR computing model in a *nutshell*

- *Grid-like distributed* computing is *dead*, long-live federated computing among large centers.
- Effective resource sharing at FAIR TIER-0 center, accounting for most of the data-driven computations.
- Federated storage and computing with 'local' centers using Teralink network & commonly used standards.
- 'Centralised' orchestration —> most suited to incorporate F.A.I.R. principle for our diverse community, introduce and to minimize the operational overhead.
- Containerised approaches and other virtualisation methods for flexible compute operations serving diverse community & optimise usage.
- Data access using http, possibly with xrootd frontend, lustre backend; AAI using widely accepted standards, weblogin, token-based, eduGAIN.

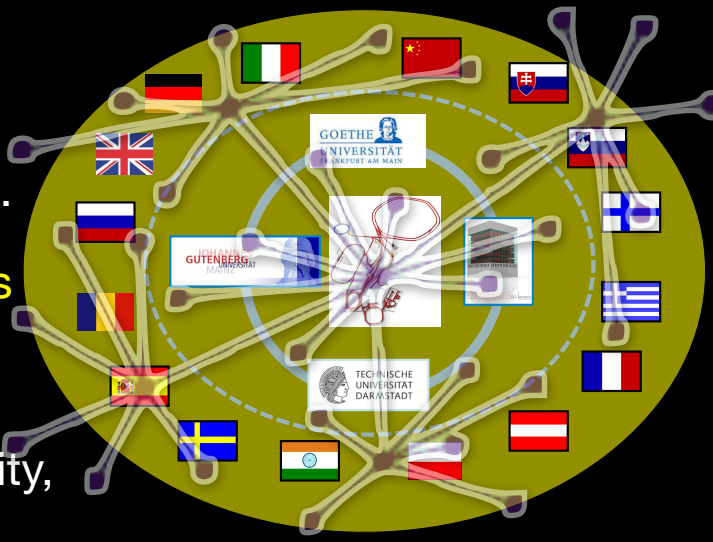

Backup material

“federated” computing - who ordered that?

- Federation: the act of uniting smaller or more localised entities to create a larger entity for mutual benefit, with agreed mixture of common policies and local autonomy.
- Consolidation: the act of reducing the number of entities by dissolution of existing ones and creation of a single larger entity.

“federated” computing - who ordered that?

- Federation allows composition of new solutions out of existing investment (but you can only rearrange the building blocks if you still have the building blocks).
- Federation enables decision making to be devolved “down” the hierarchy to where it best sits, improving choices and protecting against domination of one community or voice to the detriment of the rest.
- Federation can empower communities in a way that consolidation does not. All these elements become particularly important as the scale grows.
- Federation encourages diversity, of ideas, solutions, and people. It can protect against “group think” and stagnation, and can provide resilience against single points of failure – both geographical and technological.
- Federation enables low risk evaluation and testing of “future” technologies, in particular where they are driven by specific well motivated communities that would otherwise be overlooked or dismissed by a large scale operation with a consolidated approach.
- Federation allows smaller operations to benefit from the full scale of the federation. E.g., security, identity management, accounting and allocation; but also in the building of larger communities to share ideas and solutions.
- Federation allows leveraging of local resources that otherwise would not be available.

“federated” computing - who ordered that?

Recognise that communities need gateway projects that worry about the (evolving) complexity for them.

Share software and services where it makes sense to do so (minimises costs; maximises quality)

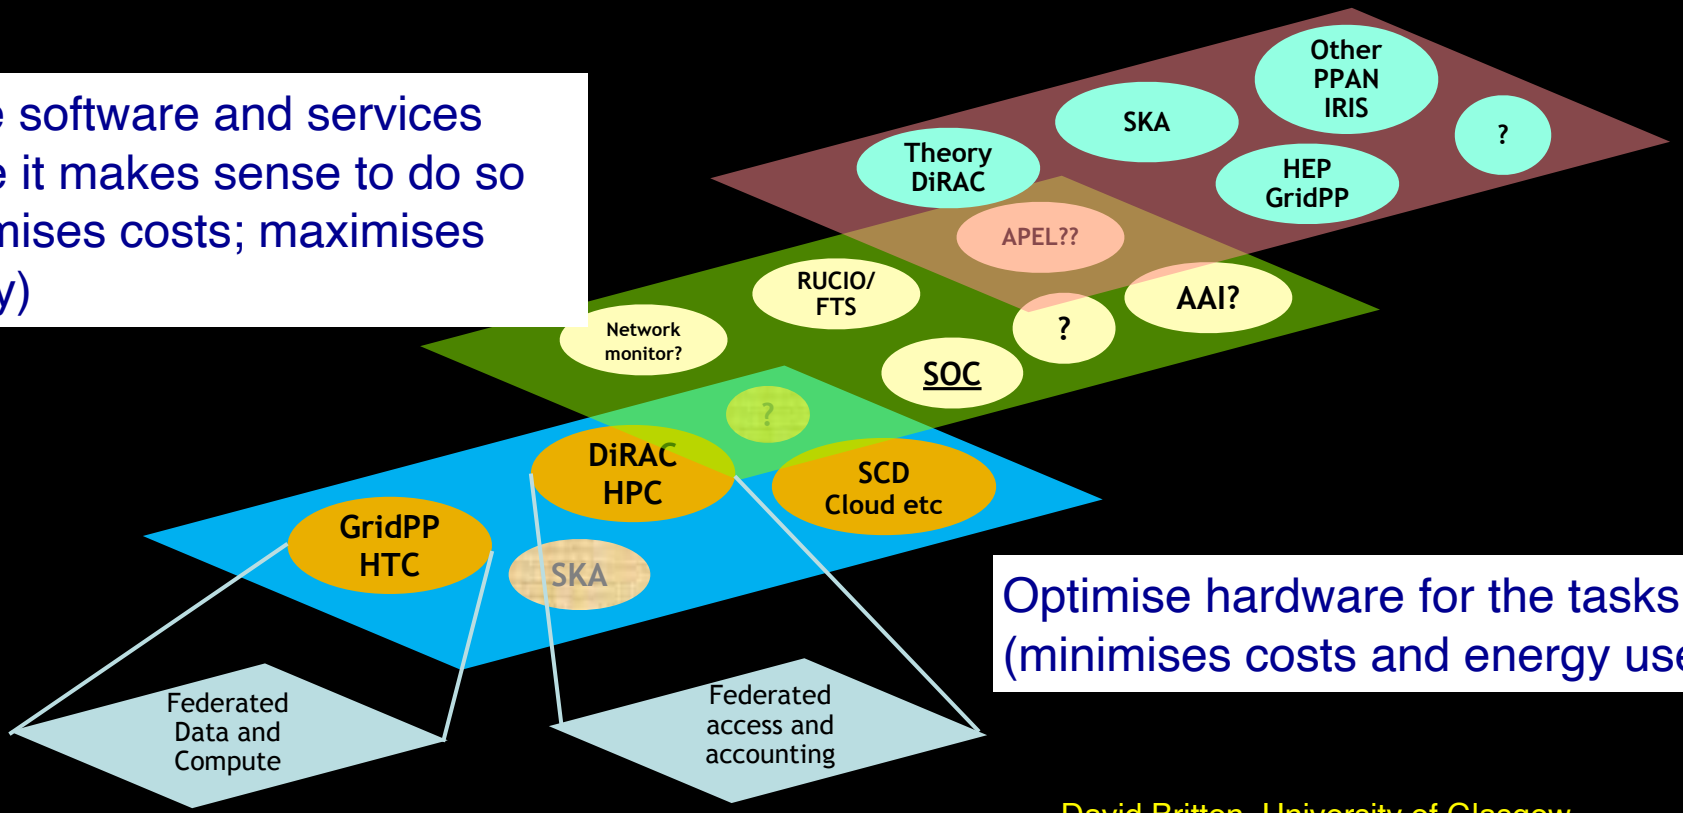

Optimise hardware for the tasks (minimises costs and energy use)

FAIR2028

Perspectives for a nominal MSV year - PANDA CBM sequential

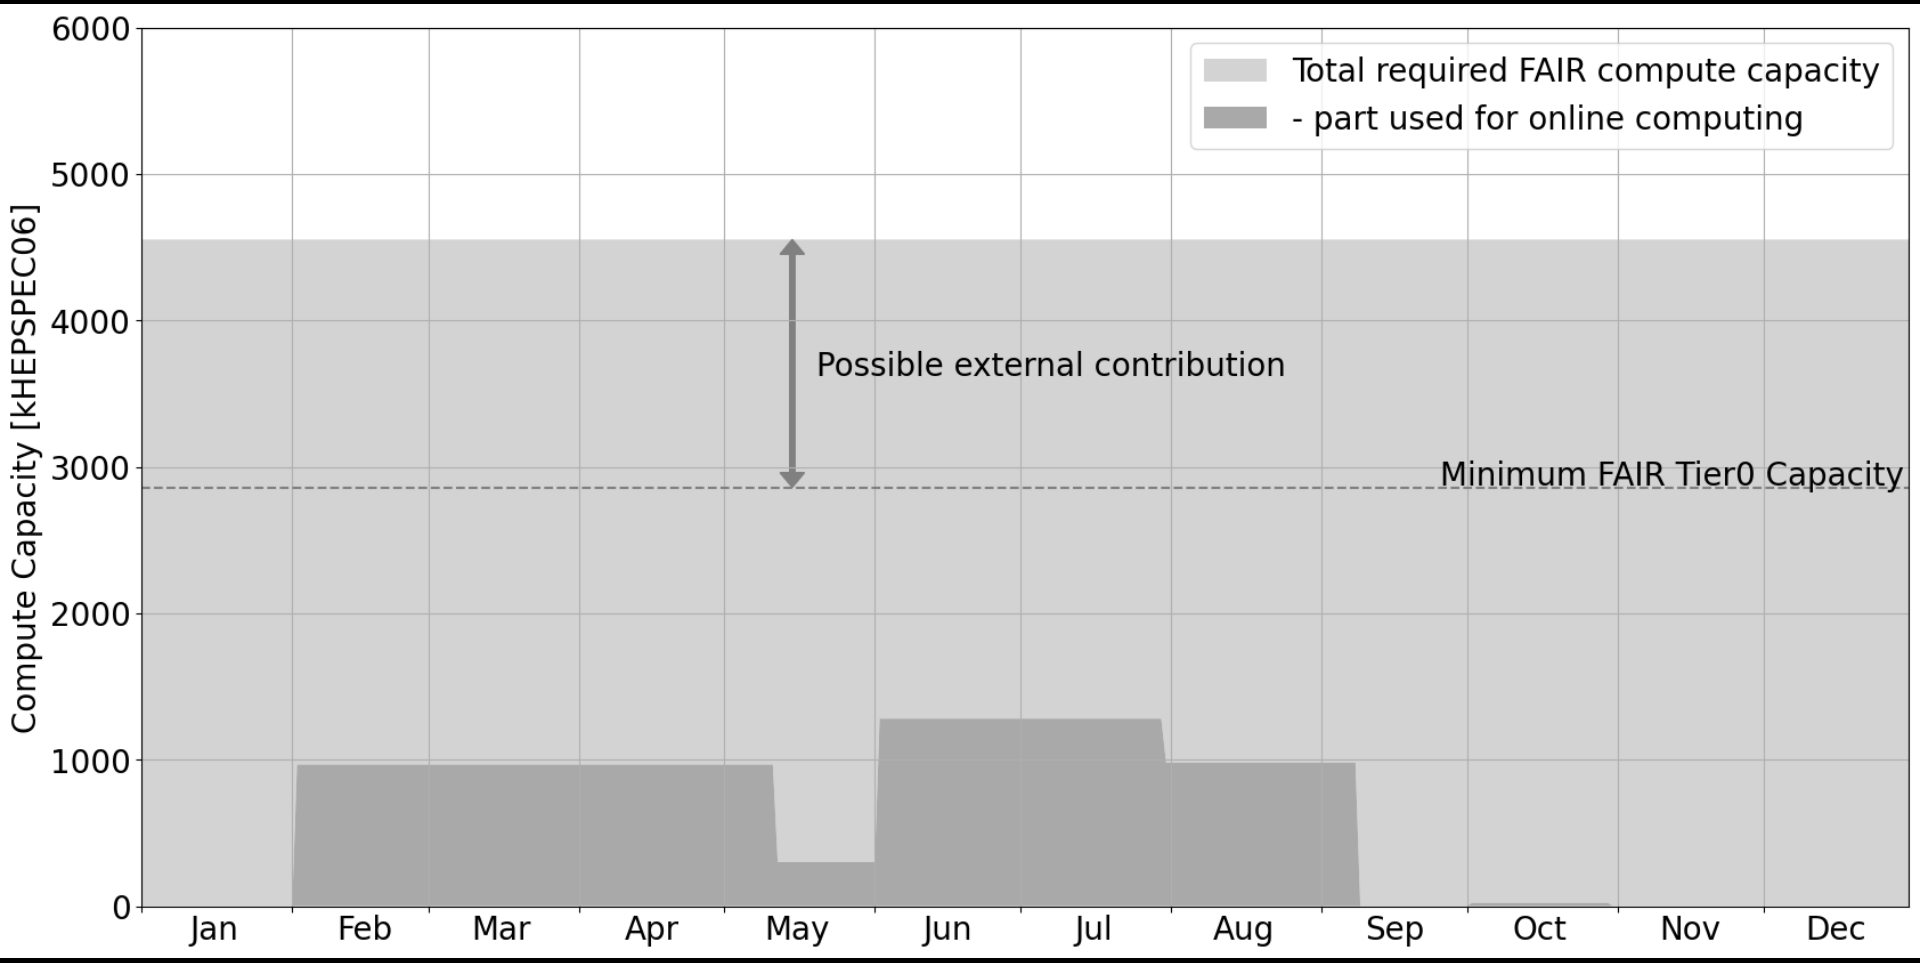

Role of ML and AI in nuclear physics

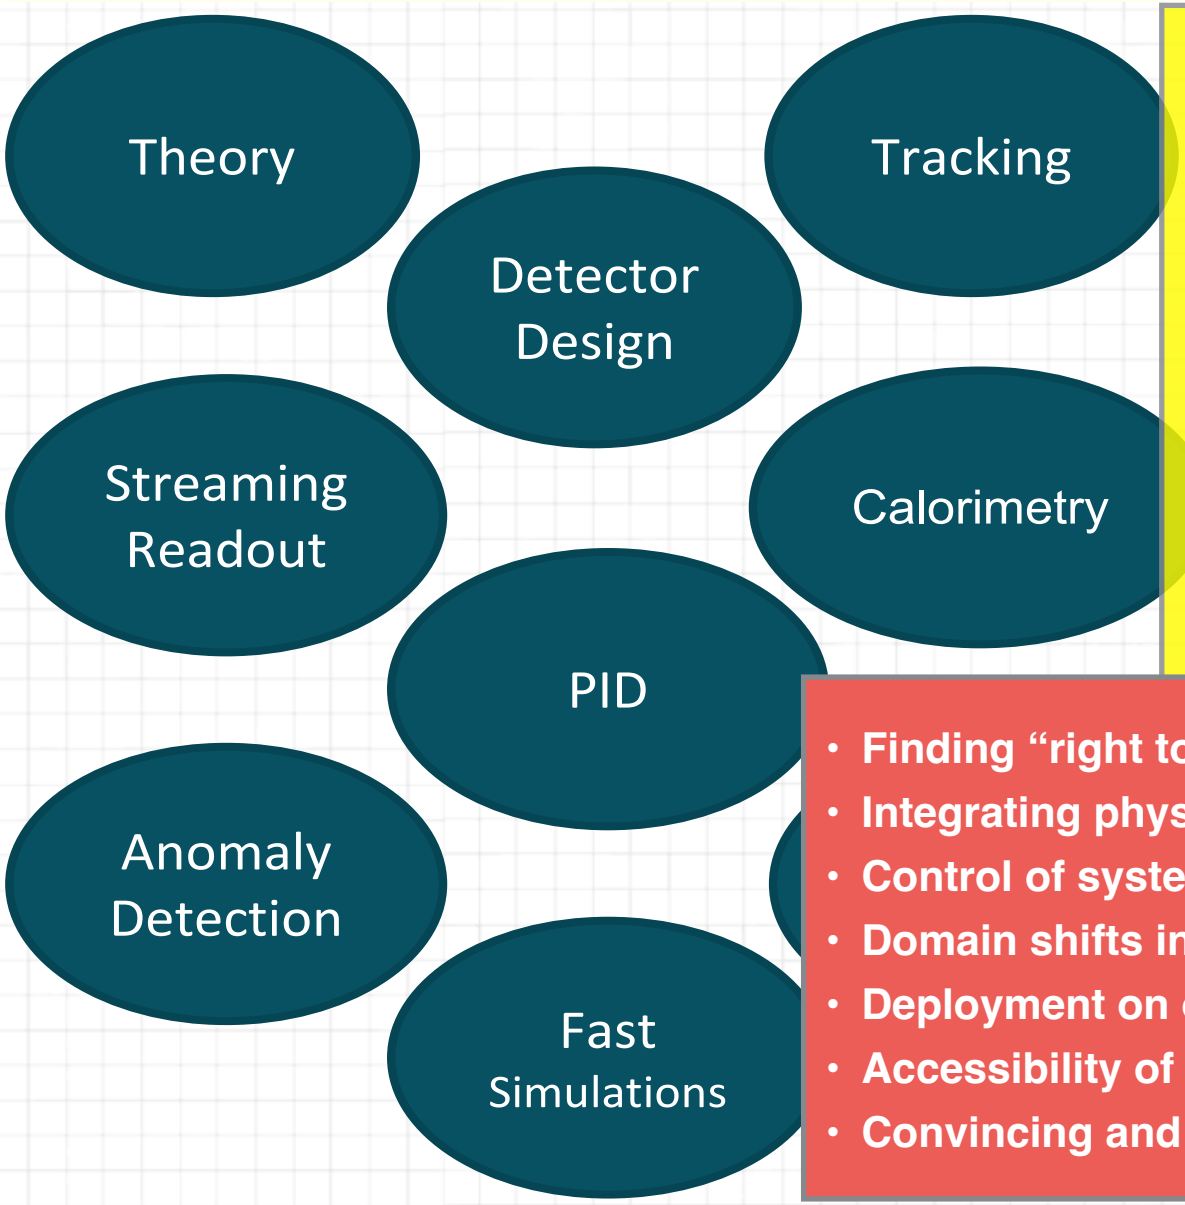

- (Variational) Auto Encoders
- Artificial Neural Networks
- Bayesian Model Averaging/Mixing
- Bayesian Optimisation
- Bayesian Neural Networks
- Convolutional Neural Networks
- Ensemble Methods & Boosting
- Generative Adversarial Networks
- Gaussian Processes
- *k*-Nearest Neighbours
- Kernel Regression
- Logistic Regression
- Long Short-Term Memory
- Principal Component Analysis

Challenges!

- **Finding "right tool for the right job"!**
- **Integrating physics knowledge.**
- **Control of systematic errors.**
- **Domain shifts in supervised learning.**
- **Deployment on online/embedded architectures.**
- **Accessibility of model/trained data.**
- **Convincing and involving community.**

Towards an International Network

For Multiphysics Modelling, Machine learning and Model-based Control in Accelerator Sciences and Technologies

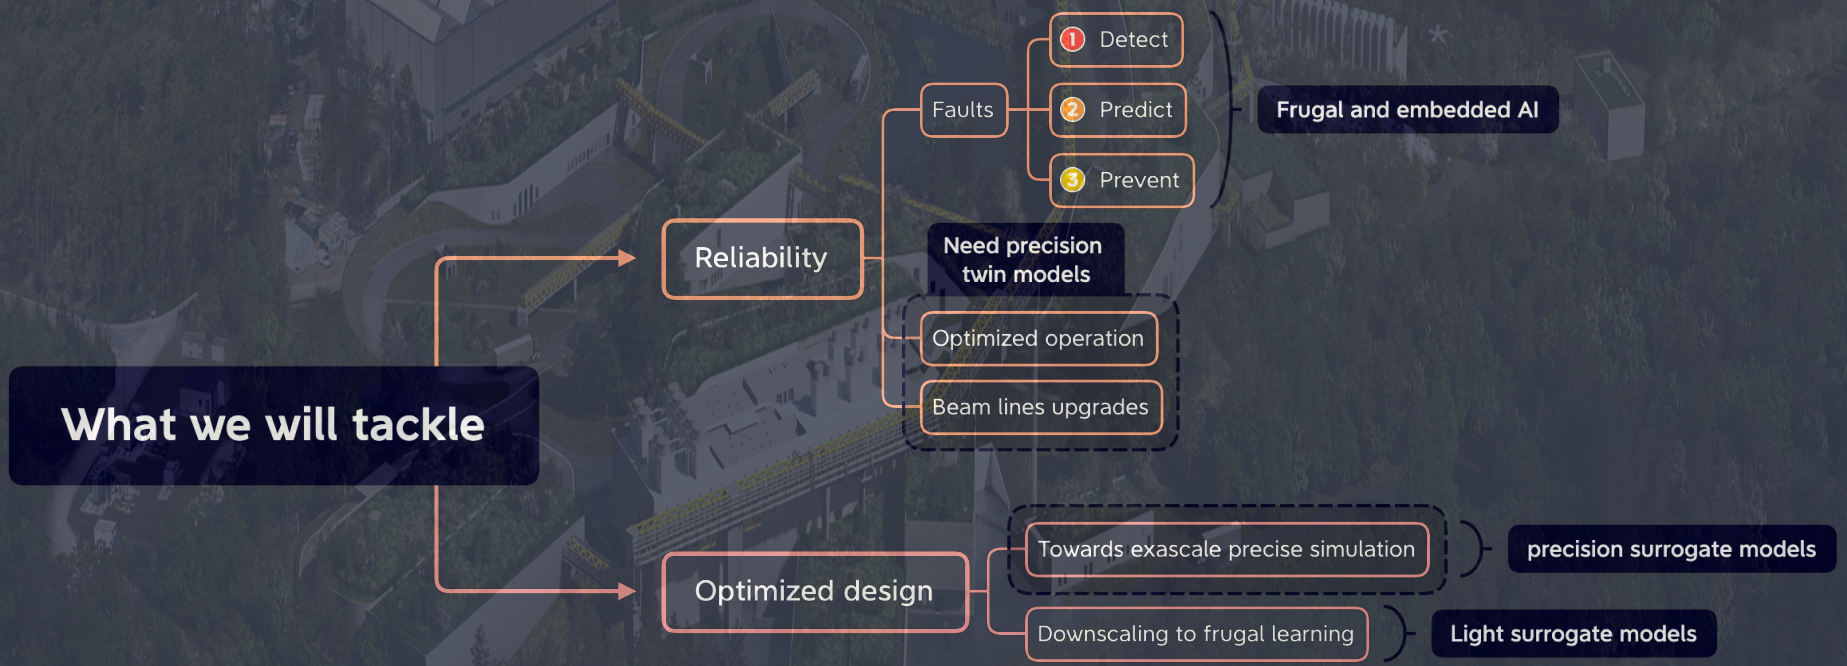

Adrian Oeftiger (GSI)

InM4CAST