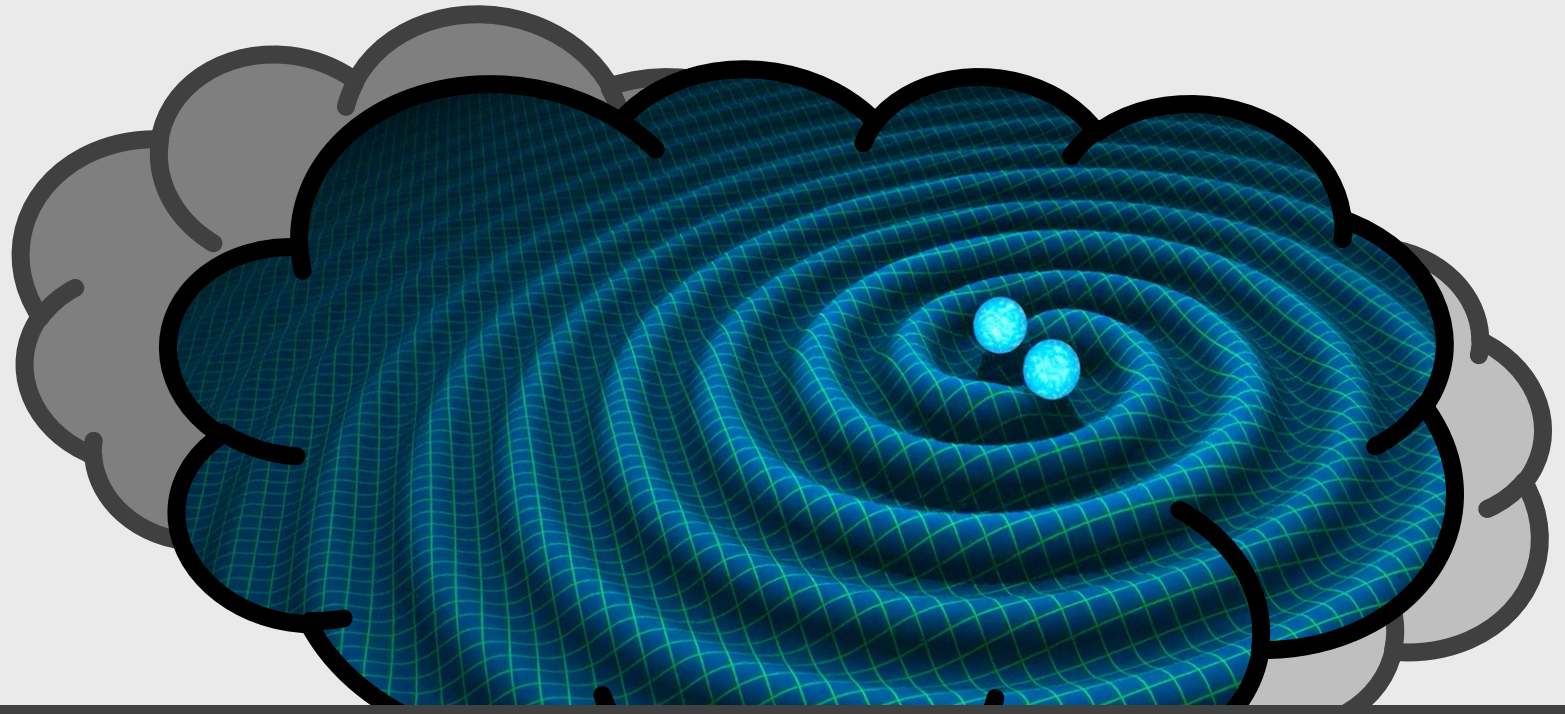

E-INFRASTRUCTURE BOARD

Stefano Bagnasco, INFN

for the ET-EIB

ET: Scienza e Tecnologia in Italia

Assisi | Feb 23, 2024

SETTNG THE SCALE

- Raw interferometer data don't grow much with increasing instrument sensitivity
  - Current detectors write o(2PB)/year of raw data per detector
  - In ET we expect about few tens of PB of raw data per year (baseline 6-interferometer design, more control channels,...)
  - No big deal today, piece of cake by 2035

- However, the amount of useful scientific information encoded in the data does grow a lot
  - And the computing power needed to wring it out (mostly from CBC Parameter Estimation)
  - Larger template banks, longer templates to fit in memory, overlapping events,...
  - ...as much as possible in low-latency
  - Accurately estimating the computing power needs is itself a difficult task

THE MANDATORY BOXES-AND-ARROWS SLIDE

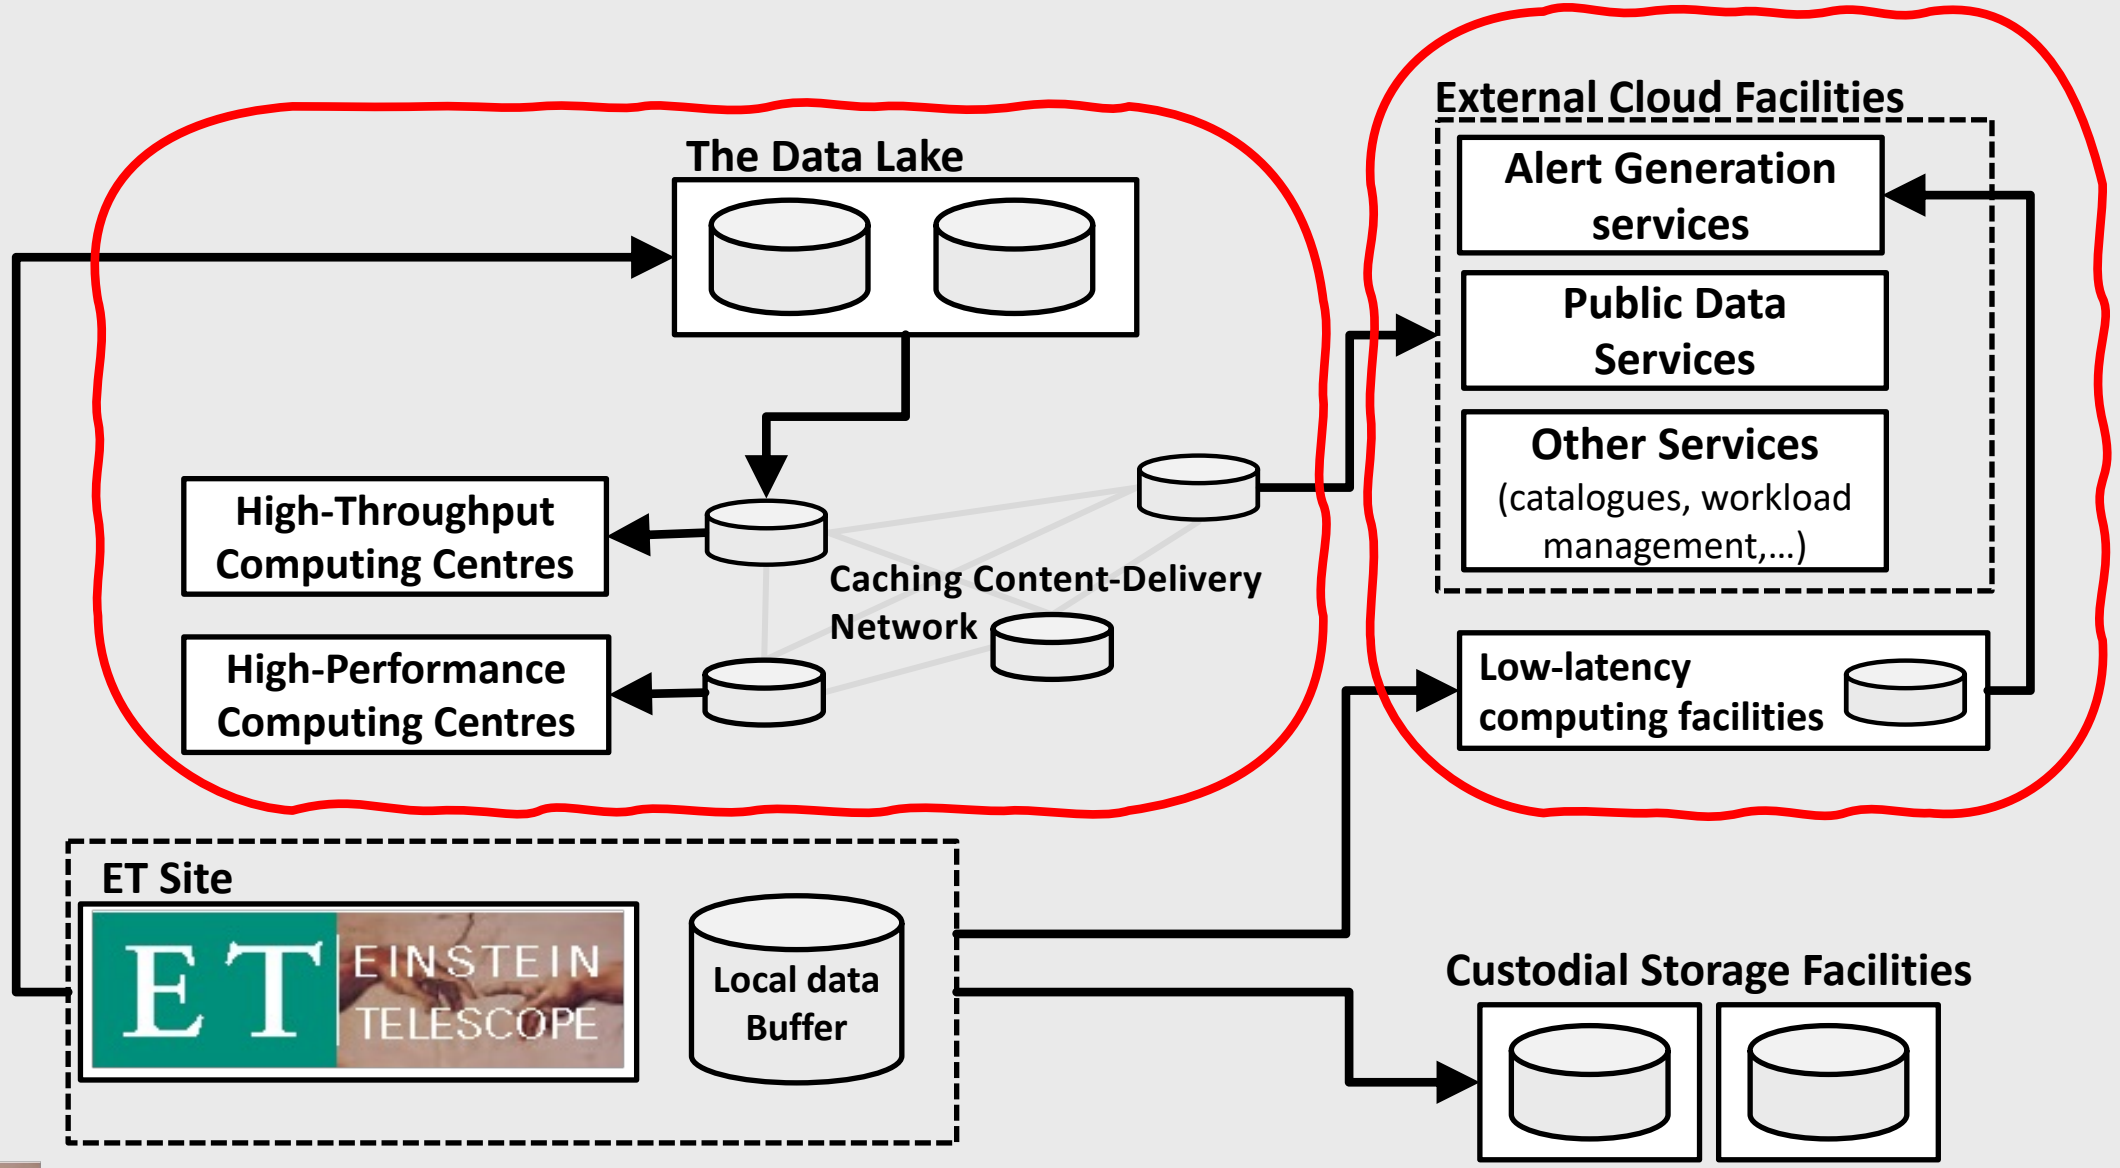

Figure from the ET ESFRI proposal

THE ECOSYSTEM (AND SOME QUESTIONS)

- First and foremost, other 3G facilities (and the 2G network)
  - CE (and LISA!)
  - Will there be an equivalent of the IGWN common computing infrastructure?
- Several EM and astroparticle initiatives starting in the same time frame
  - CTA, SKA, KM3Net, ELT in the ESFRI roadmap, and many more
  - They will all have stringent low-latency alert requirements, as producers or consumers (or both)
  - High rates will imply extreme automation in the generation and selection of triggers, and sophisticated scheduling algorithms
  - Will there be a MM-specific (possibly virtual) shared infrastructure like the WLCG?
  - The architecture of the next LL alert distribution system is being defined now!
- The EU is building the European Open Science Cloud
  - “Scientific Computing in the Digital Continuum”
  - How concrete will it be in 2035?

SHOPPING LIST

Data transfer and storage: safely and efficiently transfer all data to custodial storage and processing centres, including low-latency transfers;

Software packaging and distribution: manage software lifecycle, and make packages available ubiquitously;

Computing power: provide and manage computing resources (HTC and HPC) for the processing of data, in all computing domains;

Data distribution: make data available to worker nodes in computing centres anywhere, and possibly also to single workstations, including support to public releases of data;

High-availability service management: provide a platform for running the collaboration's services (e.g. alert generation services, event databases,...)

Data cataloguing and bookkeeping: organise all data and metadata and provide querying and discovering capabilities;

Job lifecycle management: provide a uniform job submission and runtime environment to research groups;

High-level workload management: keep a database of all jobs and allow the enforcement of priorities and scheduling strategies; provide support for organized large-scale data processing campaigns;

Monitoring and accounting: monitor local and distributed computing, checking performance and looking for issues, and provide reliable accounting both at the user/job and site level;

Authentication, Authorisation and Identity management: provide consistent AAI across all domains and activities.

Collaboration services: provide tools for efficient collaboration management, coordination, and outreach (e.g. document repositories, collaborative tools, administrative databases, communications,...)

THE COMPUTING MODEL

- The overall architecture of the e-Infrastructure, either as a single integrated system or as a few separate systems (e.g. instrument control and DAQ, low-latency, and offline)
- A documented way of evaluating the required computing power and storage space from the evolving scientific program of the collaboration
- Estimates of the involved costs and growth timelines
- A description of the data flows, with estimates for the needed network performances
- A description of the User Experience and workflows for relevant activities
- A description of the tools chosen or to be developed to provide all the required functionalities (foundation libraries, frameworks, middleware,...)
- Separate "Work Breakdown Structure" and "Implementation Plan" documents

THE ET E-INFRASTRUCTURE BOARD

Chairs: Patrice Verdier (IJCLab), SB (INFN-Torino)

Division 1: Software, frameworks, and data challenge support
Chair: Andres Tanasijczuk (UCLouvain)
OSB Liaison: John Veitch (U. Glasgow)

Division 2: Services and Collaboration Support
Chair: Antonella Bozzi (EGO)

Division 3: Computing and data model, Resource Estimation
Chair: Gonzalo Merino (PIC)

Division 4: Multimessenger alerts infrastructure
Chair: Steven Schramm (U. Geneva)

TTG: Technology Tracking working Group
Chair: Sara Vallero (INFN-Torino)

EIB IN ETMD

EIB Resp: Verdier Bagnasco 44 (FRTE: 11.785)

- DIV 01 - Software, frameworks and data-challenge support 3 (FRTE: 0.2)
- DIV 02 - Services and collaboration support 3 (FRTE: 0.485)
- DIV 03 - Computing and data model, resource estimation 1 (FRTE: 0.05)
- DIV 04 - Multimessenger-alert infrastructure 5 (FRTE: 0.55)
- TTG - Technology-tracking working group 5 (FRTE: 1.4)

ET-PP (INFRA-DEV) WP8 DELIVERABLES

D8.1: Computing and Data Requirements (M18): Documentation of the inputs on the computing and data requirements received during the process;

D8.2: Computing and Data Model (M42): Final version of 8.1;

D8.3: Data Access Implementation Guidelines (M48): A document describing how to implement the policy for the storage and the access to the ET data, according to the data model.

- Technologies for multimessenger astronomy together with CTA, KM3NeT, Virgo.
  - Lots of different things synergic with different ETIC labs
- Not funded in first round of INFRA-Tech call two years ago, rewritten and will be resubmitted
  - Deadline March 14
- Computing WP6: “Efficient and sustainable computing”
  - WP leader: Steven Schramm (U. Geneva)
  - Preliminary WP budget: o(2M) EU + o(600k) CH

Task 6.1: efficient data processing

- Early robust processing = less to process later = better energy efficiency
- Mostly supervised ML to enable fast/real-time data processing to enhance MM event identification
- Includes common framework covering full ML lifecycle, supporting cross-fertilisation / knowledge transfer

Task 6.2: sustainable large-scale computing
Involved: CNAF (Daniele Cesini), INFN-Torino/CTLab (SB)

- How to sustainably scale computing to handle large MM event rates and mitigate energy/carbon costs
- Work with large computing centres to study how to scale-up computing for large MM event rates Together with academic partners to bridge the gap between RIs and computing centres

Task 6.3: multimessenger alert tools

Involved: INFN-PG (Giuseppe Greco, task leader)

- How to ensure different research infrastructures can communicate effectively
- Common alert formats, brokers, databases, etc - all while ensuring alerts follow FAIR principles
- Inherently requires cross-RI interactions and joint study/development

Plan: use MDCs to evaluate community expectations and test prototypes

- First MDC: come as you are, we just provide data distribution (through IGWN)
- Second MDC: provide more tools and collect feedback
- Iterate

- First MDC took forever to start
- Communication not very good – difficult to gather information
- Working on it…

COLLABORATION SERVICES

- ET Member's Database
  - Developed by Gary Hemmings at EGO
  - Derived from the Virgo VMD
- AAI services
  - Identity and Access Management, SSO etc.
  - Activity bootstrapped (again…) yesterday, led by Michel Jouvin (U. Paris Saclay), support by CNAF (Francesco Giacomini et al.) and EGO (Stefano Cortese and Gary Hemming)
  - See next slide
- DRESS Voting system support
  - I am doing it – any volounteers?
- More services provided by EGO
  - GitLab, web, wiki, TDS

IAM & SSO

- System based on widely-adopted tools and standards, plus our own flavour
  - INDIGO IAM as core component
  - ETMD as authoritative membership source
  - SAML and OIDC services as identity providers (home institutions, ORCID, EGO Active Directory,...)
- Need some (not much) integration work
  - Definition of the precise workflow
  - Interface between IAM and ETMD
  - Definition of role and group structure
- Next step: hosting site choice
  - Tier-1 service level
  - Availability of support FTEs
  - Availability of operational experience with IAM
  - Natural candidates: EGO, CNAF
  - Less-natural candidates: CC-IN2P3, BSC, KIT, ?

EIB and ET-PP WP8 workshop

Geneva, Oct 26-17

- Requirements collection from ISB & OSB
- ET-PP D8.1 draft sent to ET-PP management last week
- “Living” document that will be updated while constraints and quantitative requirements become clear.

ET EINSTEIN TELESCOPE

Preparatory Phase for the Einstein Telescope Gravitational Wave Observatory

Deliverable 8.1

Computing and Data Requirements

Lead beneficiary: UNIGE
Delivery Date: 29 February 2024
Dissemination level: public
Version: 1.0

This project has received funding from the European Commission Framework Programme Horizon Europe Coordination and Support action under grant agreement 101079696.

EXTERNAL RELATIONS

- ESCAPE Open Collaboration
  - ET is a member, with representatives in the managing bodies
  - Manages the OSCARS calls (see later)
  - Provide feedback for the R&D strategy and join the thematic WGs
- JENA (Joint ECFA/NuPECC/APPEC) computing WG
  - Trying to be present (as GW in general) in the thematic working groups:
    - HTC, WLCG and HPC (HPC)
    - Software and Heterogeneous Architectures (Software)
    - Federate Data Management, Virtual Research Environments and FAIR/Open Data (Data)
    - Machine Learning and Artificial Intelligence (AI)
    - Training, Dissemination, Education (TDE)
- WLCG
  - Through Virgo membership for the time being
- IGWN
  - Makes sense to be a part of it? Yes, but.

WORKFLOW EVALUATION KITS

- ESCAPE Datalake “own” deployment for data distribution
  - Replace OSDF (CVMFS + StashCache)
- IAM-based AAI
  - This might actually come into production immediately thereafter
- ESCAPE Datalake + VRE interactive data analysis
  - Used for the tutorial on Tuesday
- OSDF + INFNCloud interactive data analysis
  - Work in progress
- More to come!

OSCARS CALLS

- Small (10—250 kEUR) projects
- First call opening in March, second November (TBC)

1. Multi-RI DataLakes (UCLouvain/INFN-Torino)
2. Multi-RI VREs (U. Geneva / IJCLab)
3. Additional functionality in the jupyter-rucio extensio (PIC/?) and demonstration through ET MDC workflows.
4. Cloud interfaces for scientific workflows (PIC/?)
5. More?

TRAINING AND RECRUITING

- We need special professional profiles
  - Something between physical science and computer engineering
  - Not exactly “pipeline developers”, not exacly “system architects”
- Such personpower is difficult to find
  - Skilled personpower for computing activities is scarce
  - Hard to train and keep, hard to hire: we are not competitive with industry!
- This is not a problem for the GW community only
  - And neither limited to the EU
- Working on this:
  - Participating in JENA WG about training
  - Trying to interface to HSF Training working group
  - National PhD program in “Advanced technologies for physics and astrophysics” hosted by UNIPD
  - One grant funded by INFN in Torino on “Advanced computing systems for GW research”

SPARE SLIDES

1/10TH OF AN LHC EXPERIMENT

- As of 2024, the computing needs of the entire GW network are roughly o(10%) of an LHC experiment of today

- In ET the event rate will be 10^3 -10^4 times the current one
  - Analysis of the “golden” events (EM counterparts, high SNR or “special” events) would already be within reach using current technologies
    - O(500) events per year = 12.5MHS06-y per year, the same order of magnitude of a LHC experiment in Run 4
    - Target: 1/10th of an LHC experiment in Run 4

E-INFRASTRUCTURE BOARD MANDATE

«...to design, create and operate an evolving, efficient and functional e-infrastructure environment at a reasonable cost for the collaboration. Initially the focus will be the development of a Computing Model for the ET».

- Prepare a plan of the studies and activities that need to be undertaken for the development of the ET computing.
- Propose a computing model and its updates to the collaboration.

Out of scope: actual science code, physics and engineering tools

THREE COMPUTING DOMAINS

On-site infrastructure

Online

- Data acquisition and pre-processing
- Instrument control
- Environmental monitoring
- …

Plain old HTC (and some HPC)

Offline

- Deep searches
- Offline parameter estimation
- Template bank generation, NR
- …

Here's the fun

Low-latency

- Candidate search
- Sky localization
- LL parameter estimation
- Alert generation and distribution

DIVISION 1: SOFTWARE, FRAMEWORKS, AND DATA CHALLENGE SUPPORT

Define the software frameworks for ET computing workflows, the middleware for infrastructure, workload and data management. Develop software quality best practices and support their adoption with training and enforcement policies. Support code development in all computing domains. Provide computing support for mock-data challenges.

- Collaborate with OSB and ISB to define the data formats (both internal and for public release) and organized data processing workflows
- Support the development of the tools for the operation of the telescope
- Coordinate the development of common infrastructural tools and frameworks for the data-analysis, both offline and in low-latency
- Support the operation of large-scale computing campaigns
- Develop policies and best practices to ensure software quality, and encourage/enforce their adoption
- Organize a continuous training programme for both developers and users

DIVISION 2: SERVICES AND COLLABORATION SUPPORT

Define and provide all the IT services needed for the administrative management of the Collaboration. Define and provide all the IT services needed for communication and collaboration within the Collaboration and outside.

- Provide collaborative tools for communication within the collaboration and to the outside
- Coordinate the operation of the collaborative and administrative tools for the management of the collaboration
- Define and provide a future-proof federated AAI infrastructure for the collaboration

Develop the Einstein Telescope Computing Model. Provide a running estimate of the computing resources needed for all computing domains.

- Develop a Work Breakdown Structure for the early stages of the preparation of the Computing Model and Cost Estimates
- Collaborate with OSB to define the initial activities to evaluate actual computing needs
- Liaise with the Numerical Relativity community

Design and develop the infrastructure needed for multi-messenger triggers management and distribution. Follow the development of software tools for low-latency computing.

- Coordinate the development of the tools for the low-latency alert generation and management system
- Participate in the technical development of the alert distribution infrastructure, by liaising with the wider astrophysical community

Follow the evolution of hardware and software computing technologies. Organize regular occasions for inter-division updates.

- A transversal working group, coordinating TT activities across all four divisions.
  - Artificial Intelligence and Machine Learning
  - GPUs and HPC, FPGA and fancier architectures such as TPUs
  - Middleware tools and technologies
  - Quantum computing!