EURO EXA
FUNDED BY THE EUROPEAN UNION

Building Computing and Data Centres for ExaScale in the EU

Presenter – Peter Hopton
Dissemination Lead, EuroEXA
Founder, Iceotope
http://thecoolingguy.me

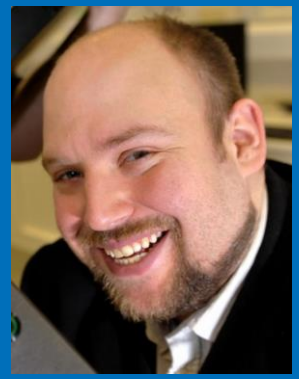

Project ID: 754337

EURO EXA

The EU Route to Exa-Scale

- 2015 – First H2020 Exa-Scale Projects (Subject area focus)
- 2017 – First Co-Design Projects
- ≅ 2019/20 – First ExaScale Demonstrators
- ≅ 2023/24 – First ExaScale Machine Turned On

Project ID: 754337

ExaScale, what do we Expect?

1. 1 Billion Billion (10^18) FLOPs or equivalent
2. Approx. €500m per system
3. 20MW-60MW of Power
4. EU expected to procure 2 systems, one from EU only Consortia
5. EU spending substantially in NRE to build an Ecosystem to Deliver
6. Time-scales of 2023 for switch-on
7. 2 €100m “Pre Exa-Scale machines” and (maybe) €40m “Exa-Scale Demonstrators” deployed in the meantime.

Project ID: 754337

What is EuroEXA?

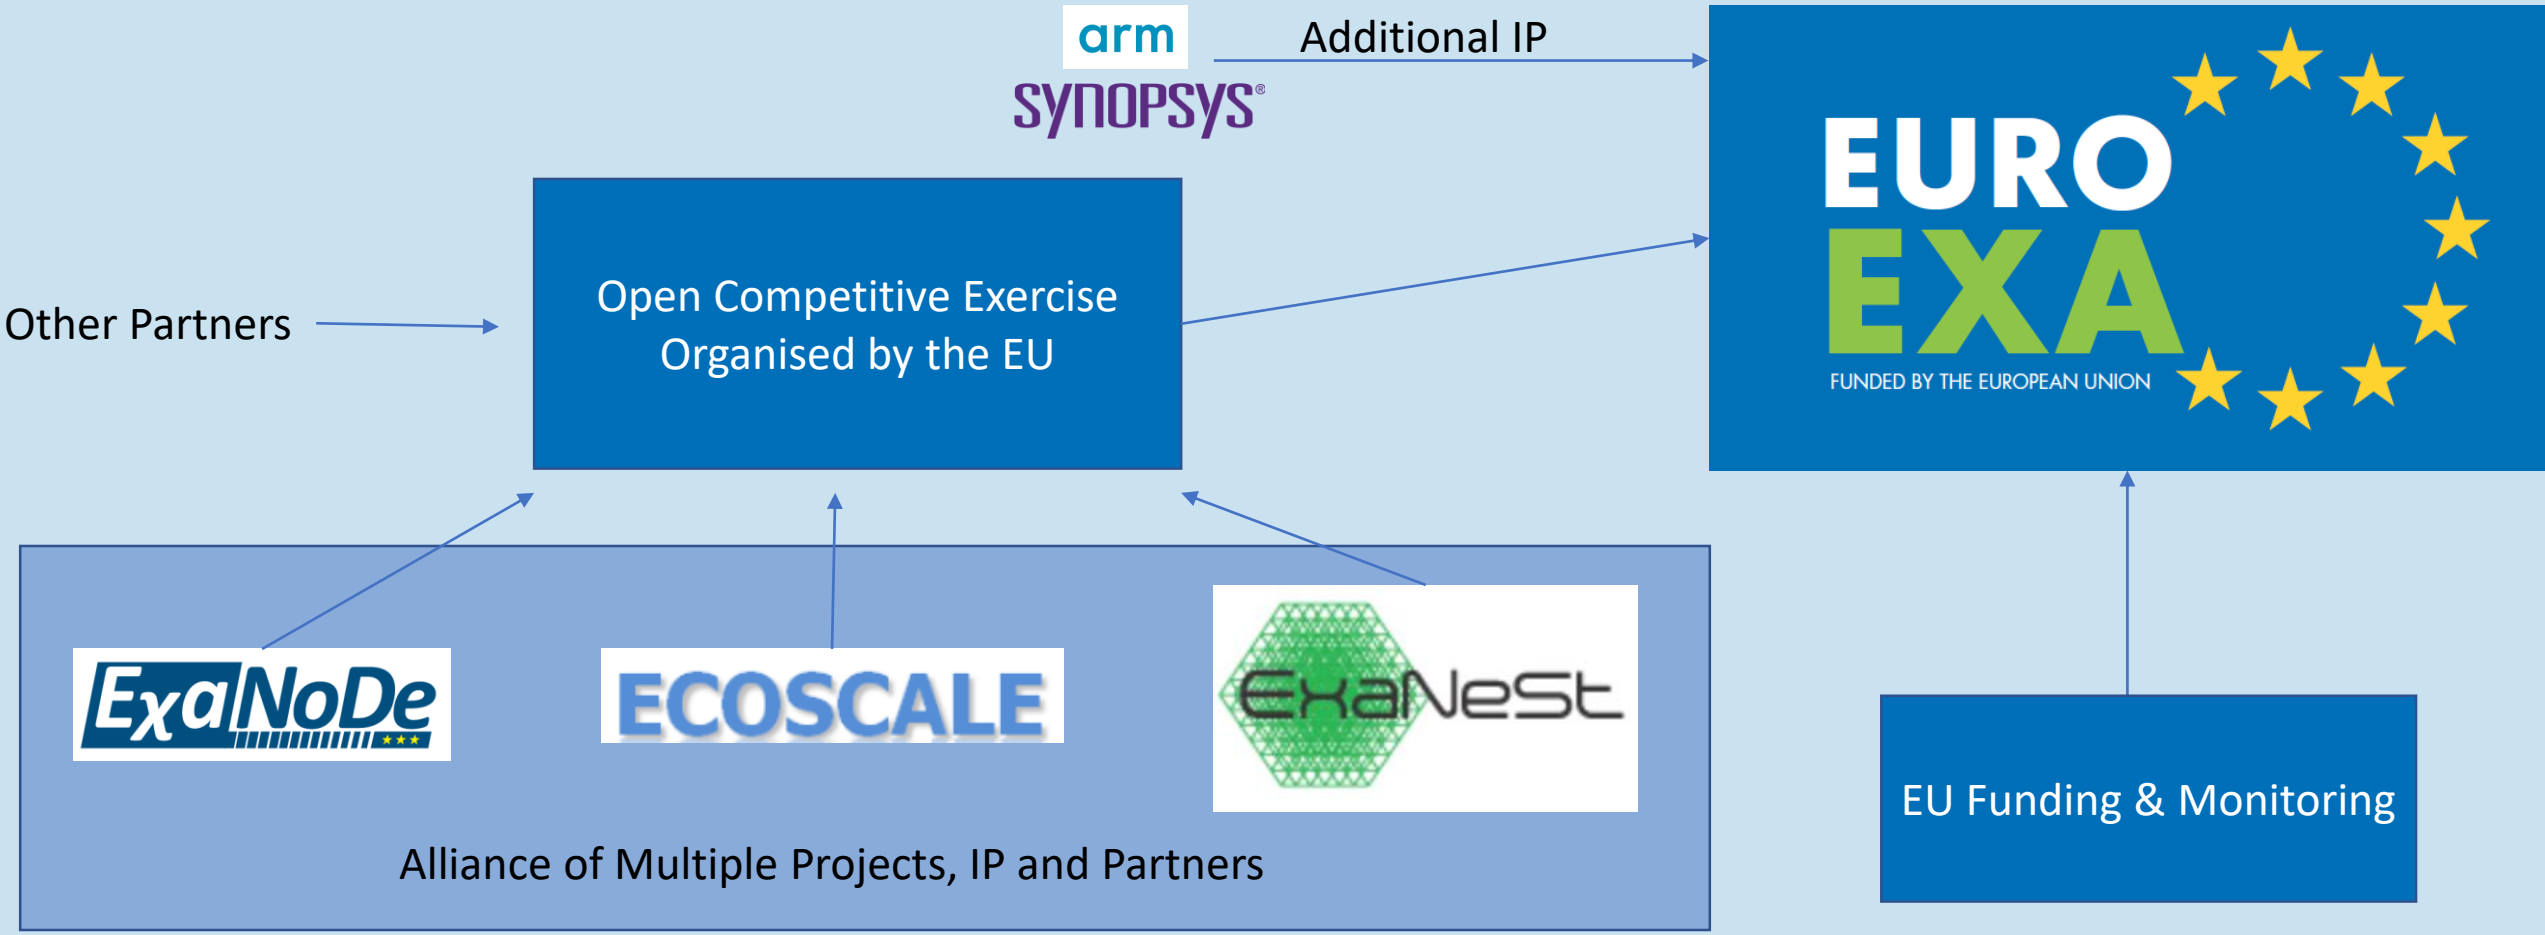

Project ID: 754337

EuroExa: European co-design for exascale applications

The Horizon 2020 EuroExa project proposes a ground-breaking design for mind-blowing results: over four times more performance and four times more energy efficiency than today's High-Performance Computing platforms.

Who Have Been Selected?

Commercial Partners

ICEOTOPE

SYNELIXIS

arm

MAXELER Technologies
Maximum Performance Computing

ZeroPoint Technologies

Academic/Gov. Partners

ICCS

MANCHESTER 1824
The University of Manchester

INAF ISTITUTO NAZIONALE DI ASTROFISICA

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

imec

FORTH
FOUNDATION FOR RESEARCH AND TECHNOLOGY - HELLAS

Fraunhofer

Hartree Centre
Science & Technology Facilities Council

ECMWF

NEURASMUS
A European Master in Neuroscience

INFN
Istituto Nazionale di Fisica Nucleare

Supporters

XILINX

SYNOPSYS®
Silicon to Software™

SKA
SQUARE KILOMETRE ARRAY

Met Office

cnag
centre nacional d'anàlisi genòmica
centro nacional de análisis genómico

janssen
PHARMACEUTICAL COMPANIES OF Johnson & Johnson

NYU School of Medicine
NYU LANGONE MEDICAL CENTER

CRS4
IDEAS BECOME LIFE

Schneider Electric

Noldus
information Technology

SOLVAY
asking more from chemistry®

Project ID: 754337

EuroEXA: co-design, demonstration and evaluation using exascale-class apps

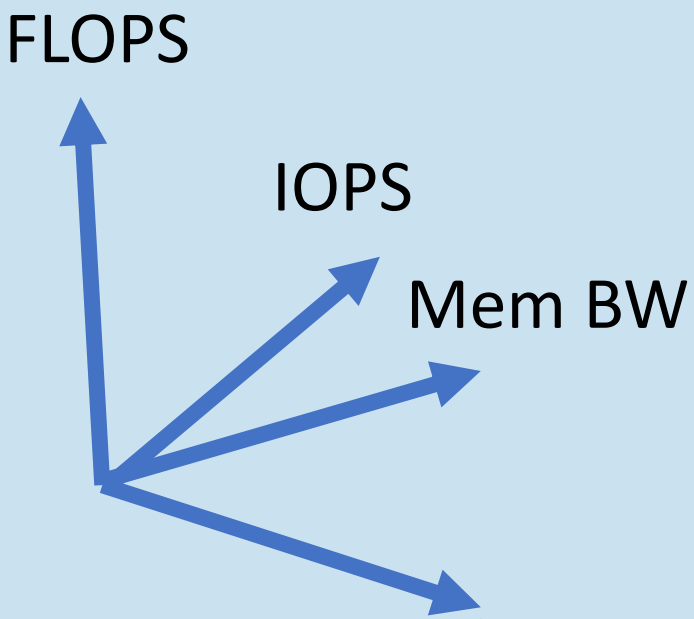

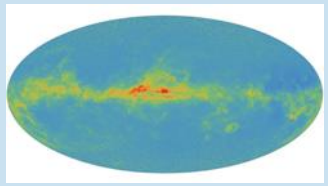
AVU-GSR

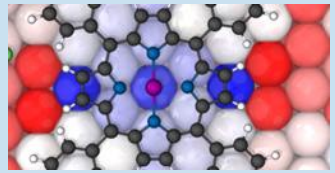
Quantum Espresso

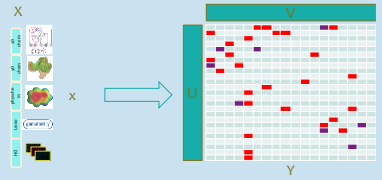
SMURFF

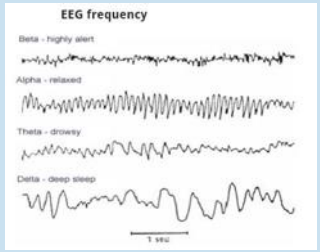

Neuromarketing

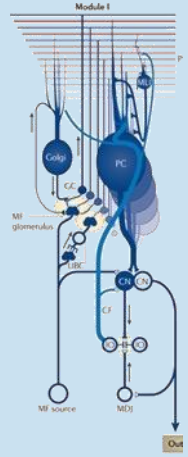

InfOli

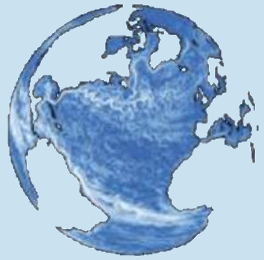
NEMO


Astronomy
image classification

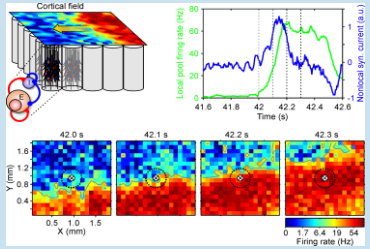

NEST/DPSNN

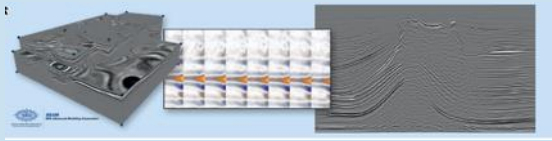
FRTM

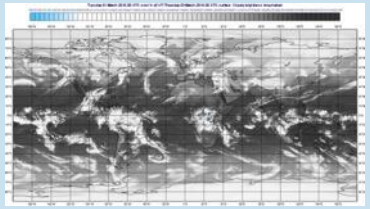
IFS

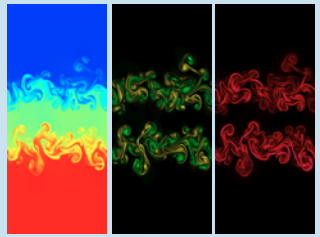
LBM

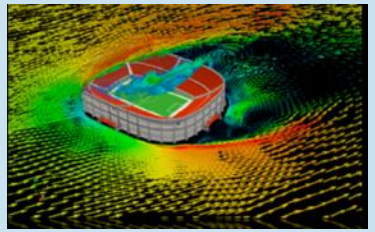
Alya


GADGET

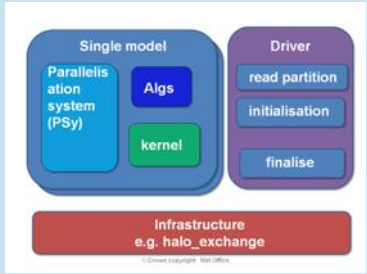

LFRic

Why these 14 applications?

- 10 strong European application partners
- 14 applications covering three important application domains
  - Climate and weather (LFRic, NEMO, IFS)
  - Physics and energy (LBM, Alya, GADGET, AVU-GSR, FRTM, Astronomy image classification, Quantum E)
  - Life science and bioinformatics (NEST, Neuromarketing, InfOli, SMURFF)
- These domains will require exascale computing in near future
- ALYA, GADGET, NEMO and Quantum ESPRESSO part of PRACE UEABS
- Applications for porting, demonstration, evaluation and co-design

Big Challenges For Exa-Scale

1. Programming models & Applications
2. Resilience and Reliability
3. Supply Chain
4. Scalability
5. Latency/Cooling/Density/Distance
6. **Energy Effectiveness**

Project ID: 754337

Energy is King

- Dr Moore's Law – The observation that the number of transistors in a dense integrated circuit doubles about every two years
- Dr Koomey's Law – Describes a long-term trend in the history of computing hardware. The number of computations per joule of energy dissipated has been doubling approximately every 1.57 years.
- Dr Shannon & Dr Hartley - Describes bandwidth with respect to signal power and noise of a channel.

EURO EXA
FUNDED BY THE EUROPEAN UNION

Dr Moore

The observation that the number of transistors in a dense integrated circuit doubles about every two years

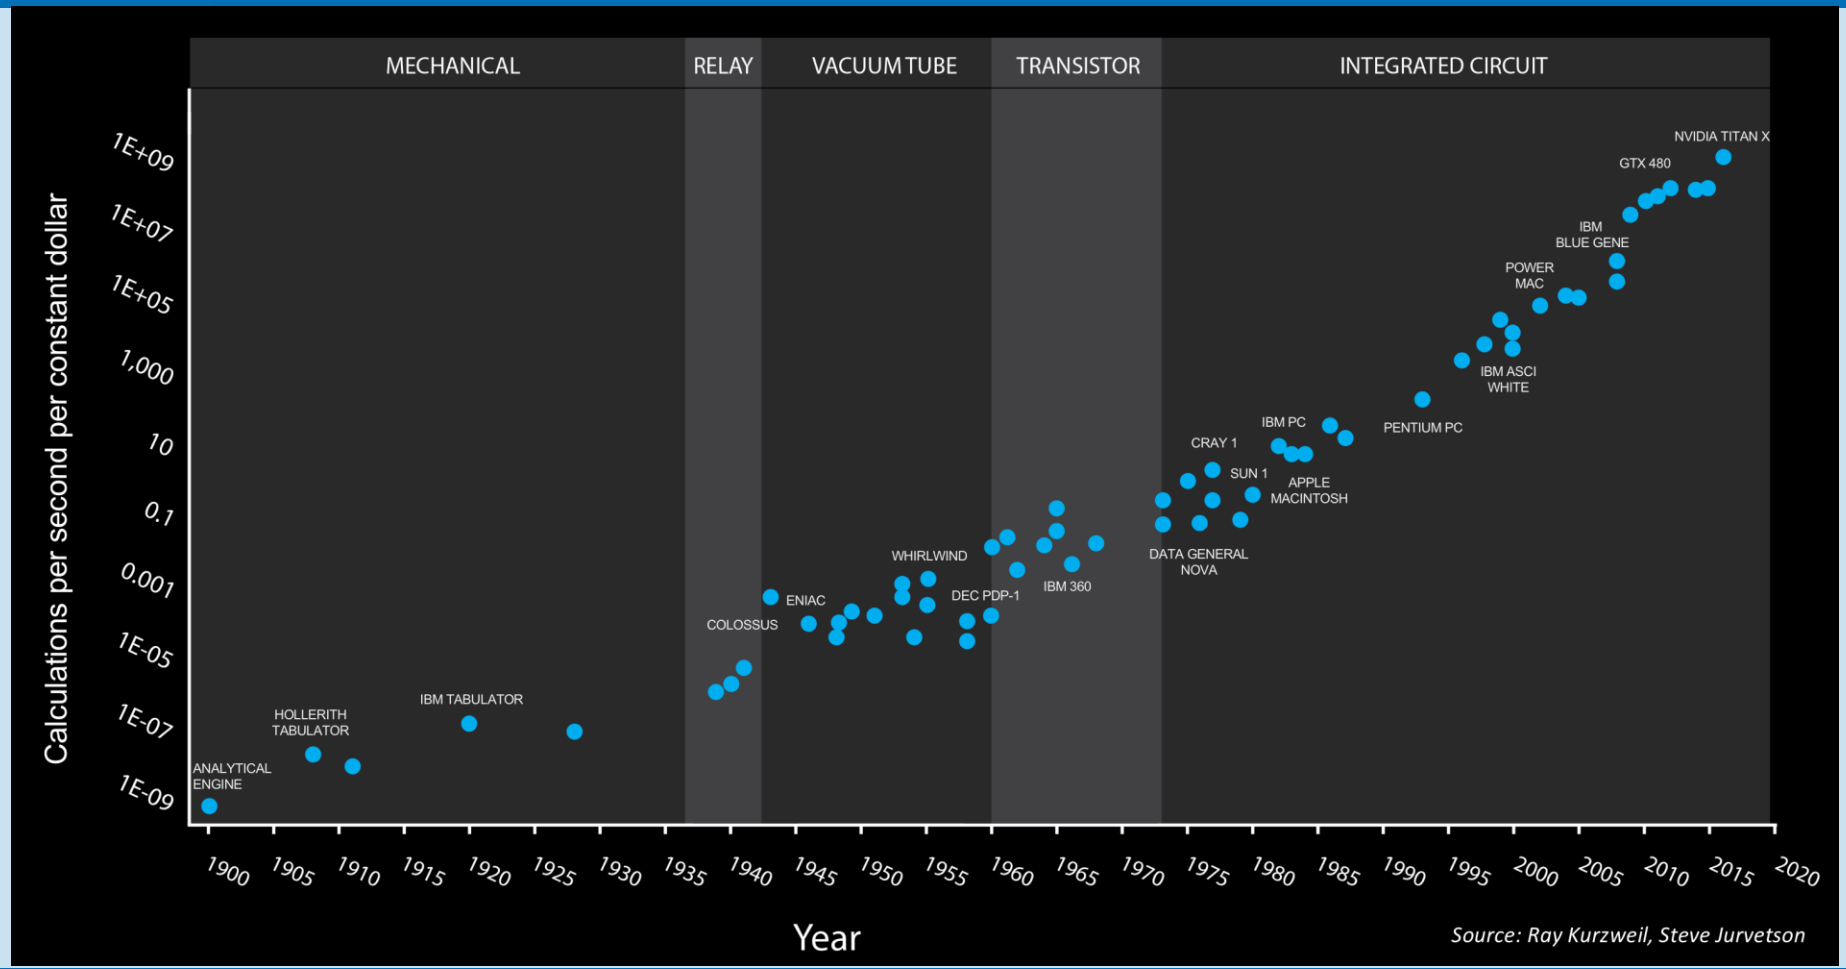

Dr Koomey–

Describes a long-term trend in the history of computing hardware. The number of computations per joule of energy dissipated has been doubling approximately every 1.57 years.

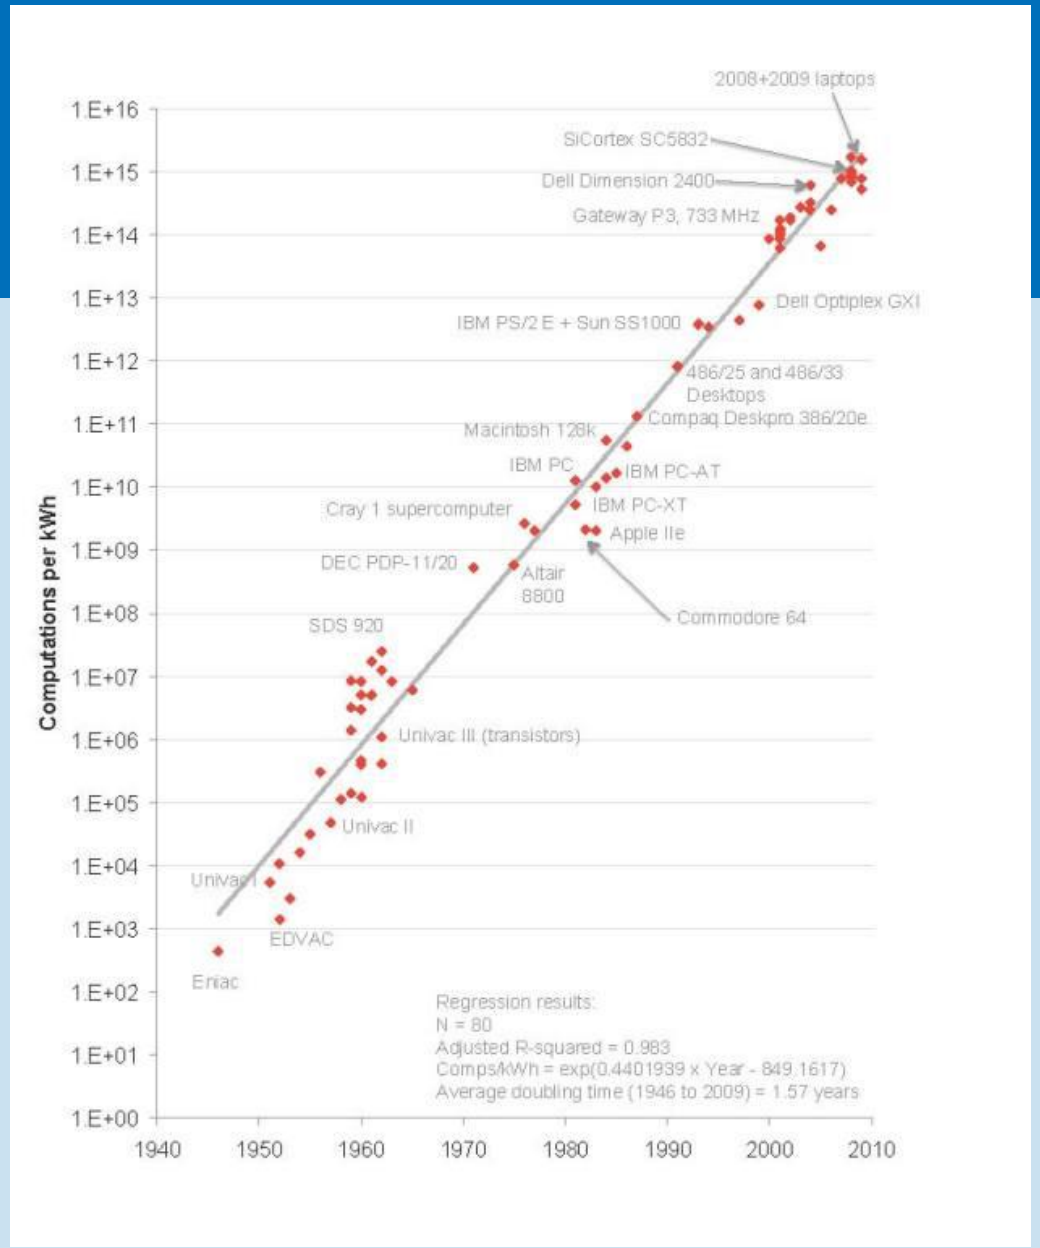

Dr Shannon & Dr Hartley–

1928/1948 Bell Labs

Describes bandwidth with respect to signal power and noise of a channel.

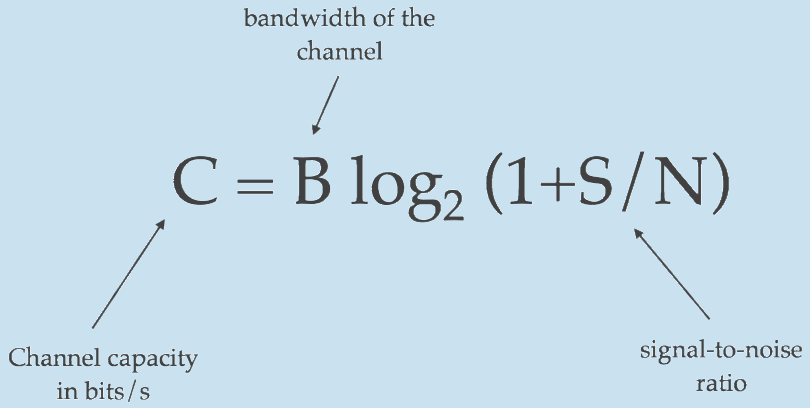

OR, it can be re-presented as the **minimum amount of energy** required to transmit a bit of data over a channel with an environment of noise.

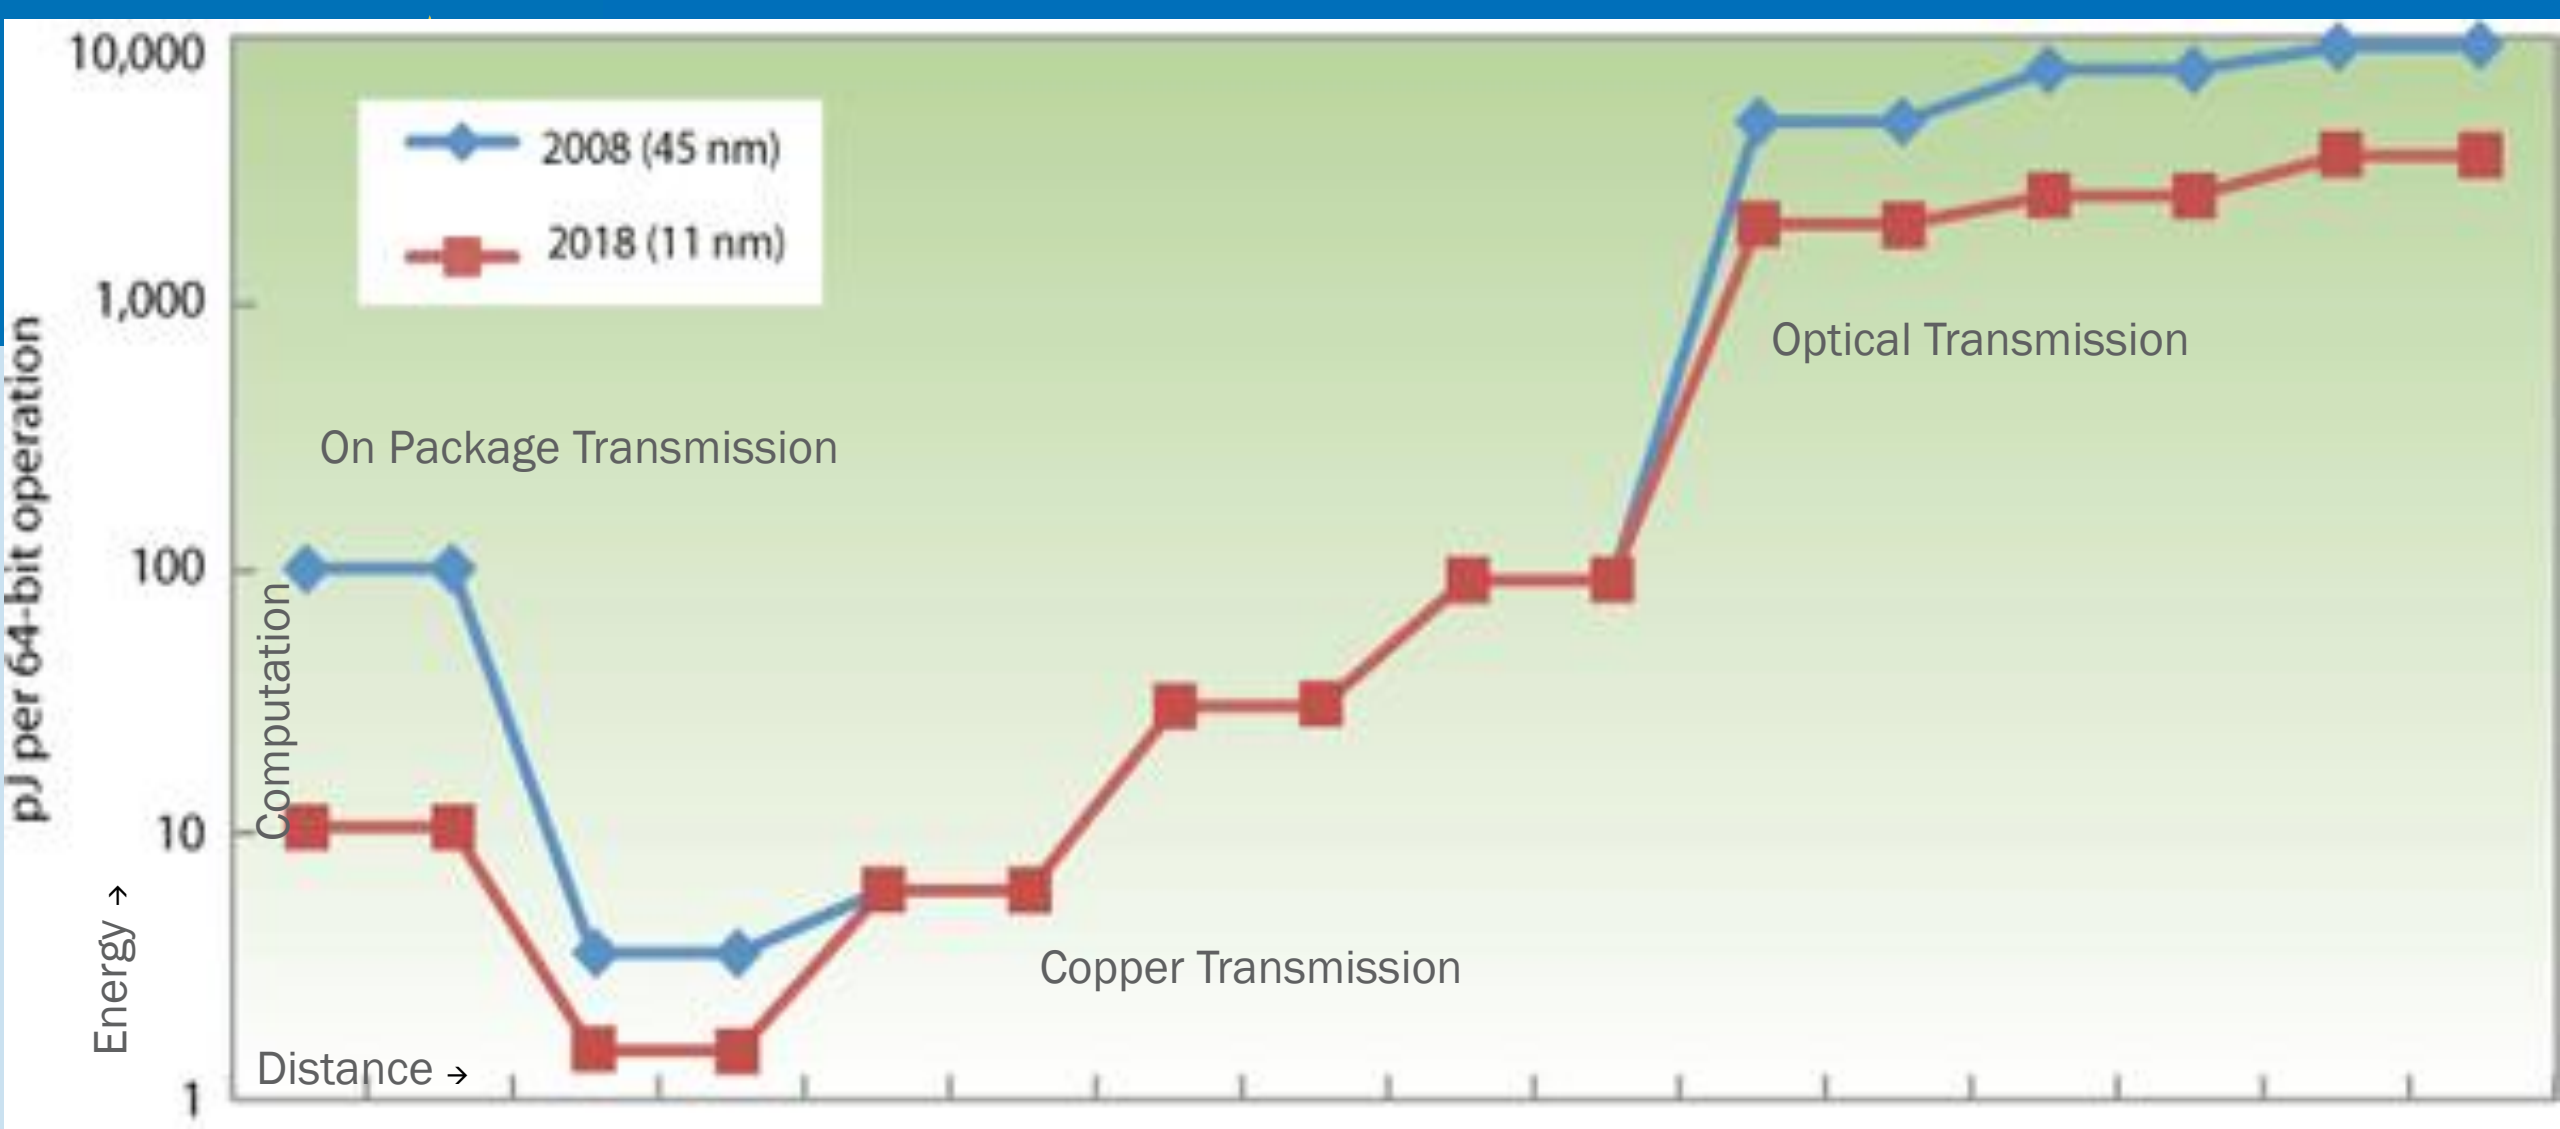
10,000
1,000
100
10
1
pJ per 64-bit operation
2008 (45 nm)
2018 (11 nm)
On Package Transmission
Optical Transmission
Copper Transmission
Computation
Energy →
Distance →

Redefining High Density….

- Not 10kW per Cabinet,
- But 60MW in a <40m Diameter
- Leading to 200kW in a conventional Cabinet footprint
- Requires a rethink…

Co-Design; The Next 3 Years

- 3 Testbeds will be deployed to test the co-design of selected technologies
  - **Testbed 1** will be installed in early 2018, providing 50 nodes of new technology for software development.
  - **Testbed 2** will be installed in early 2019, providing 500 nodes and new infrastructure technologies to test scaling
  - **Testbed 3** will be installed in 2020, to test new node and processor technologies that will ultimately deliver Exa-Scale

Project ID: 754337

Architectural Features

- UniMem
- PGAS (Partitioned Global Address Space)
- Reprogrammable Accelerators with Partial Reconfiguration
- Flexible, High Speed, Low Latency Network (INFN/UoM)

System architecture and technology: Testbed 1 compute node

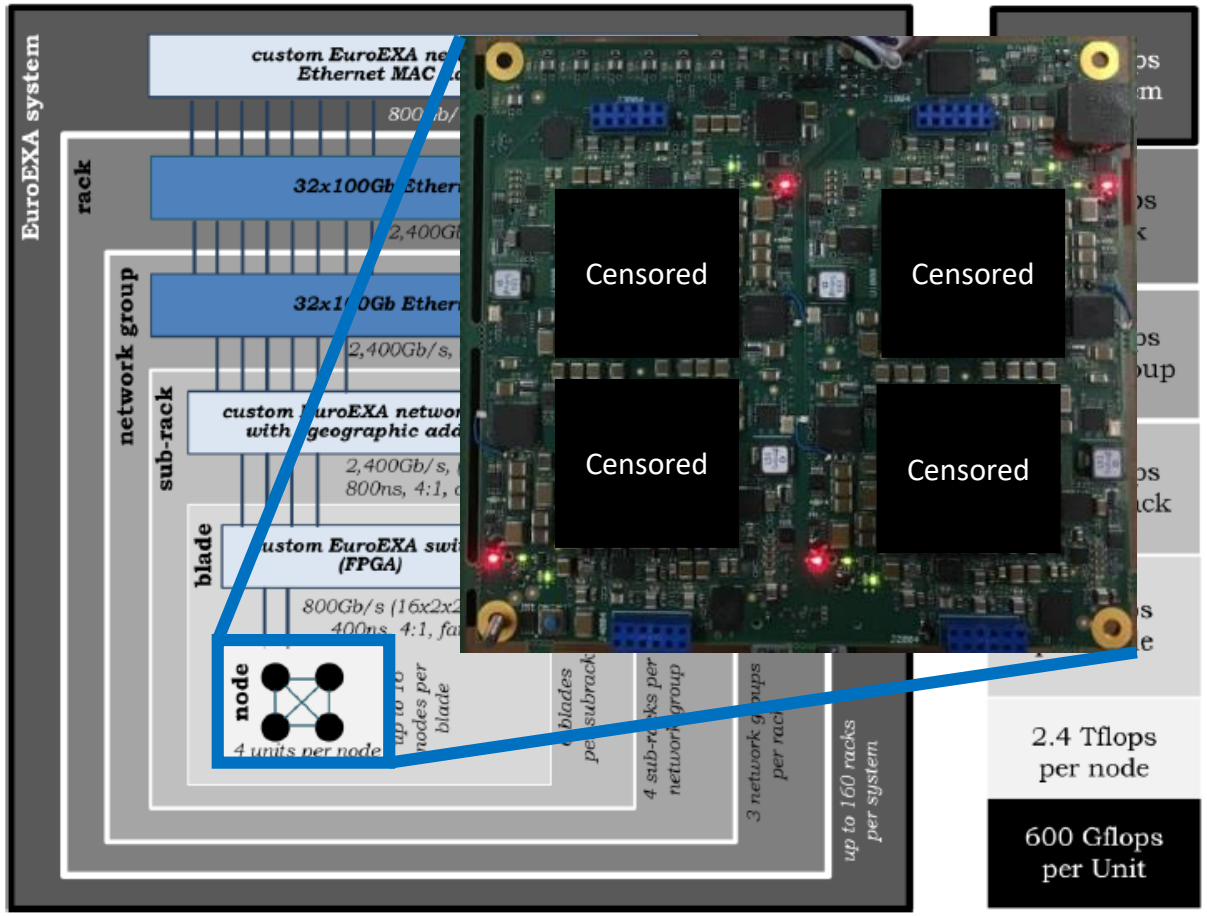

Initial technology from FORTH (QFDB)

- 12 cm x 13 cm
- 4 ARM Processors and 4 FPGA Accelerators
- M.2 SSD
- 4 x SODIMMs + Onboard RAM
- Daughterboard style
- 160 Gb/s of I/O

Technology – (Codesign)DB

Original plan: QFDB

Four Xilinx Zynq UltraScale+ ZU9P

New proposal: CRDB

Xilinx Zynq UltraScale+ ZU9P for interconnect

Xilinx Zynq UltraScale+ VU9P for acceleration

Advantages of CRDB

- More FPGA resources to help compute heavy applications
- Higher I/O bandwidth and memory bandwidth
- HPC applications already demonstrated on VU9P
  - Maxeler: BQCD, NEMO, QE, SpecFEM3D
- Compatibility with Maxeler and Amazon AWS EC2 F1
- Lower proportion of compute FPGA used for interconnect
- Removes heterogeneity among compute FPGAs

Technology - Blades

Technology from Iceotope:

- 16 COM Express Extended Nodes
- 1oU Chassis
- High Speed Switch
- 3.2kW per oU
- Total Liquid Cooling technology
- 48v DC distribution
- Hot water out, chiller-less operation

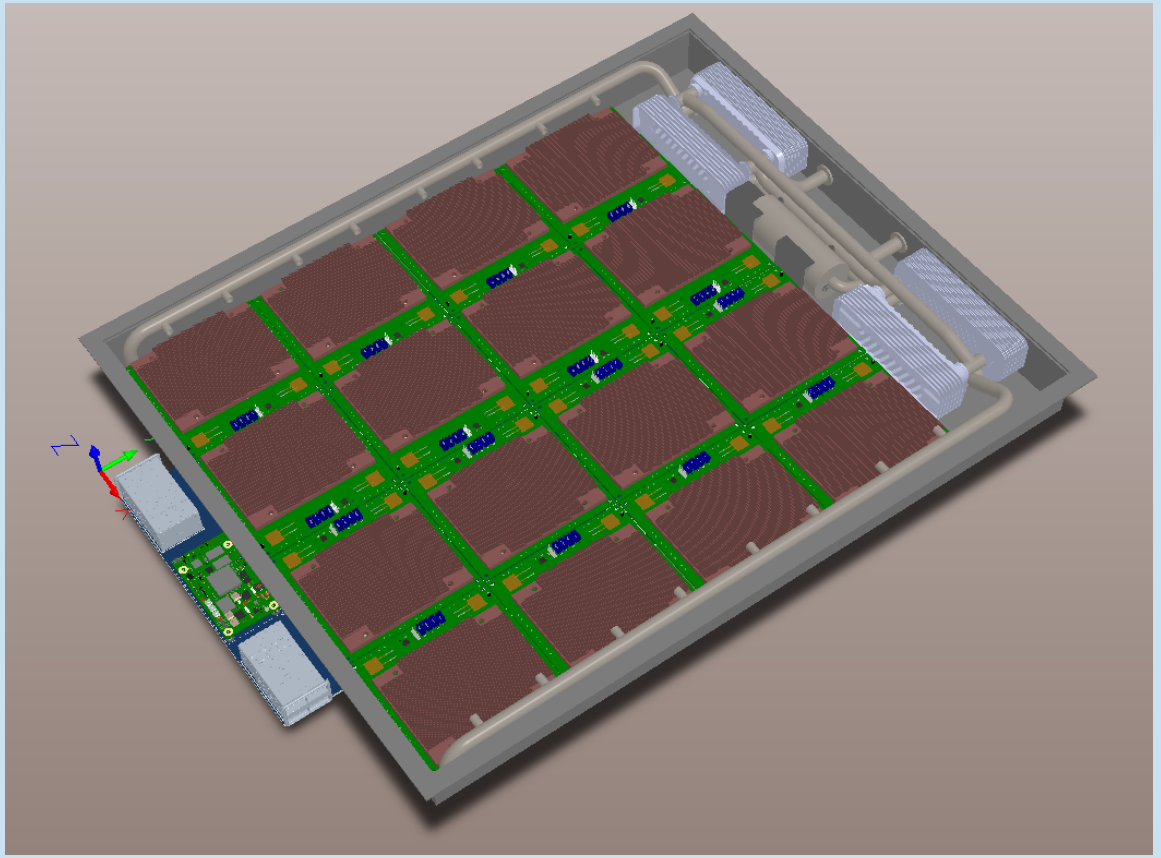

Project ID: 754337

Liquid Cooling 2MW Boxes

- To Truly Scale:
  - We Needed Zero Airflow = immersion
  - But immersion tanks are usually space inefficient
  - The project used Iceotope technology “immersion without submersion”

Turning it up to 11

- >100kW a cabinet needs liquid cooling solutions
- OCP OpenRack *is perfect for liquid cooling* as it reduces wasted space and maximises density
- Reducing the depth, enables dense facilities
- Distributed DC power enables consolidation of PSUs – also saving space

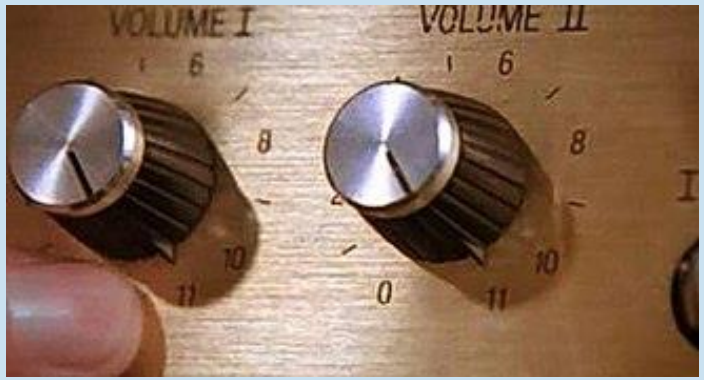

https://www.youtube.com/watch?v=4xgx4k83zzc

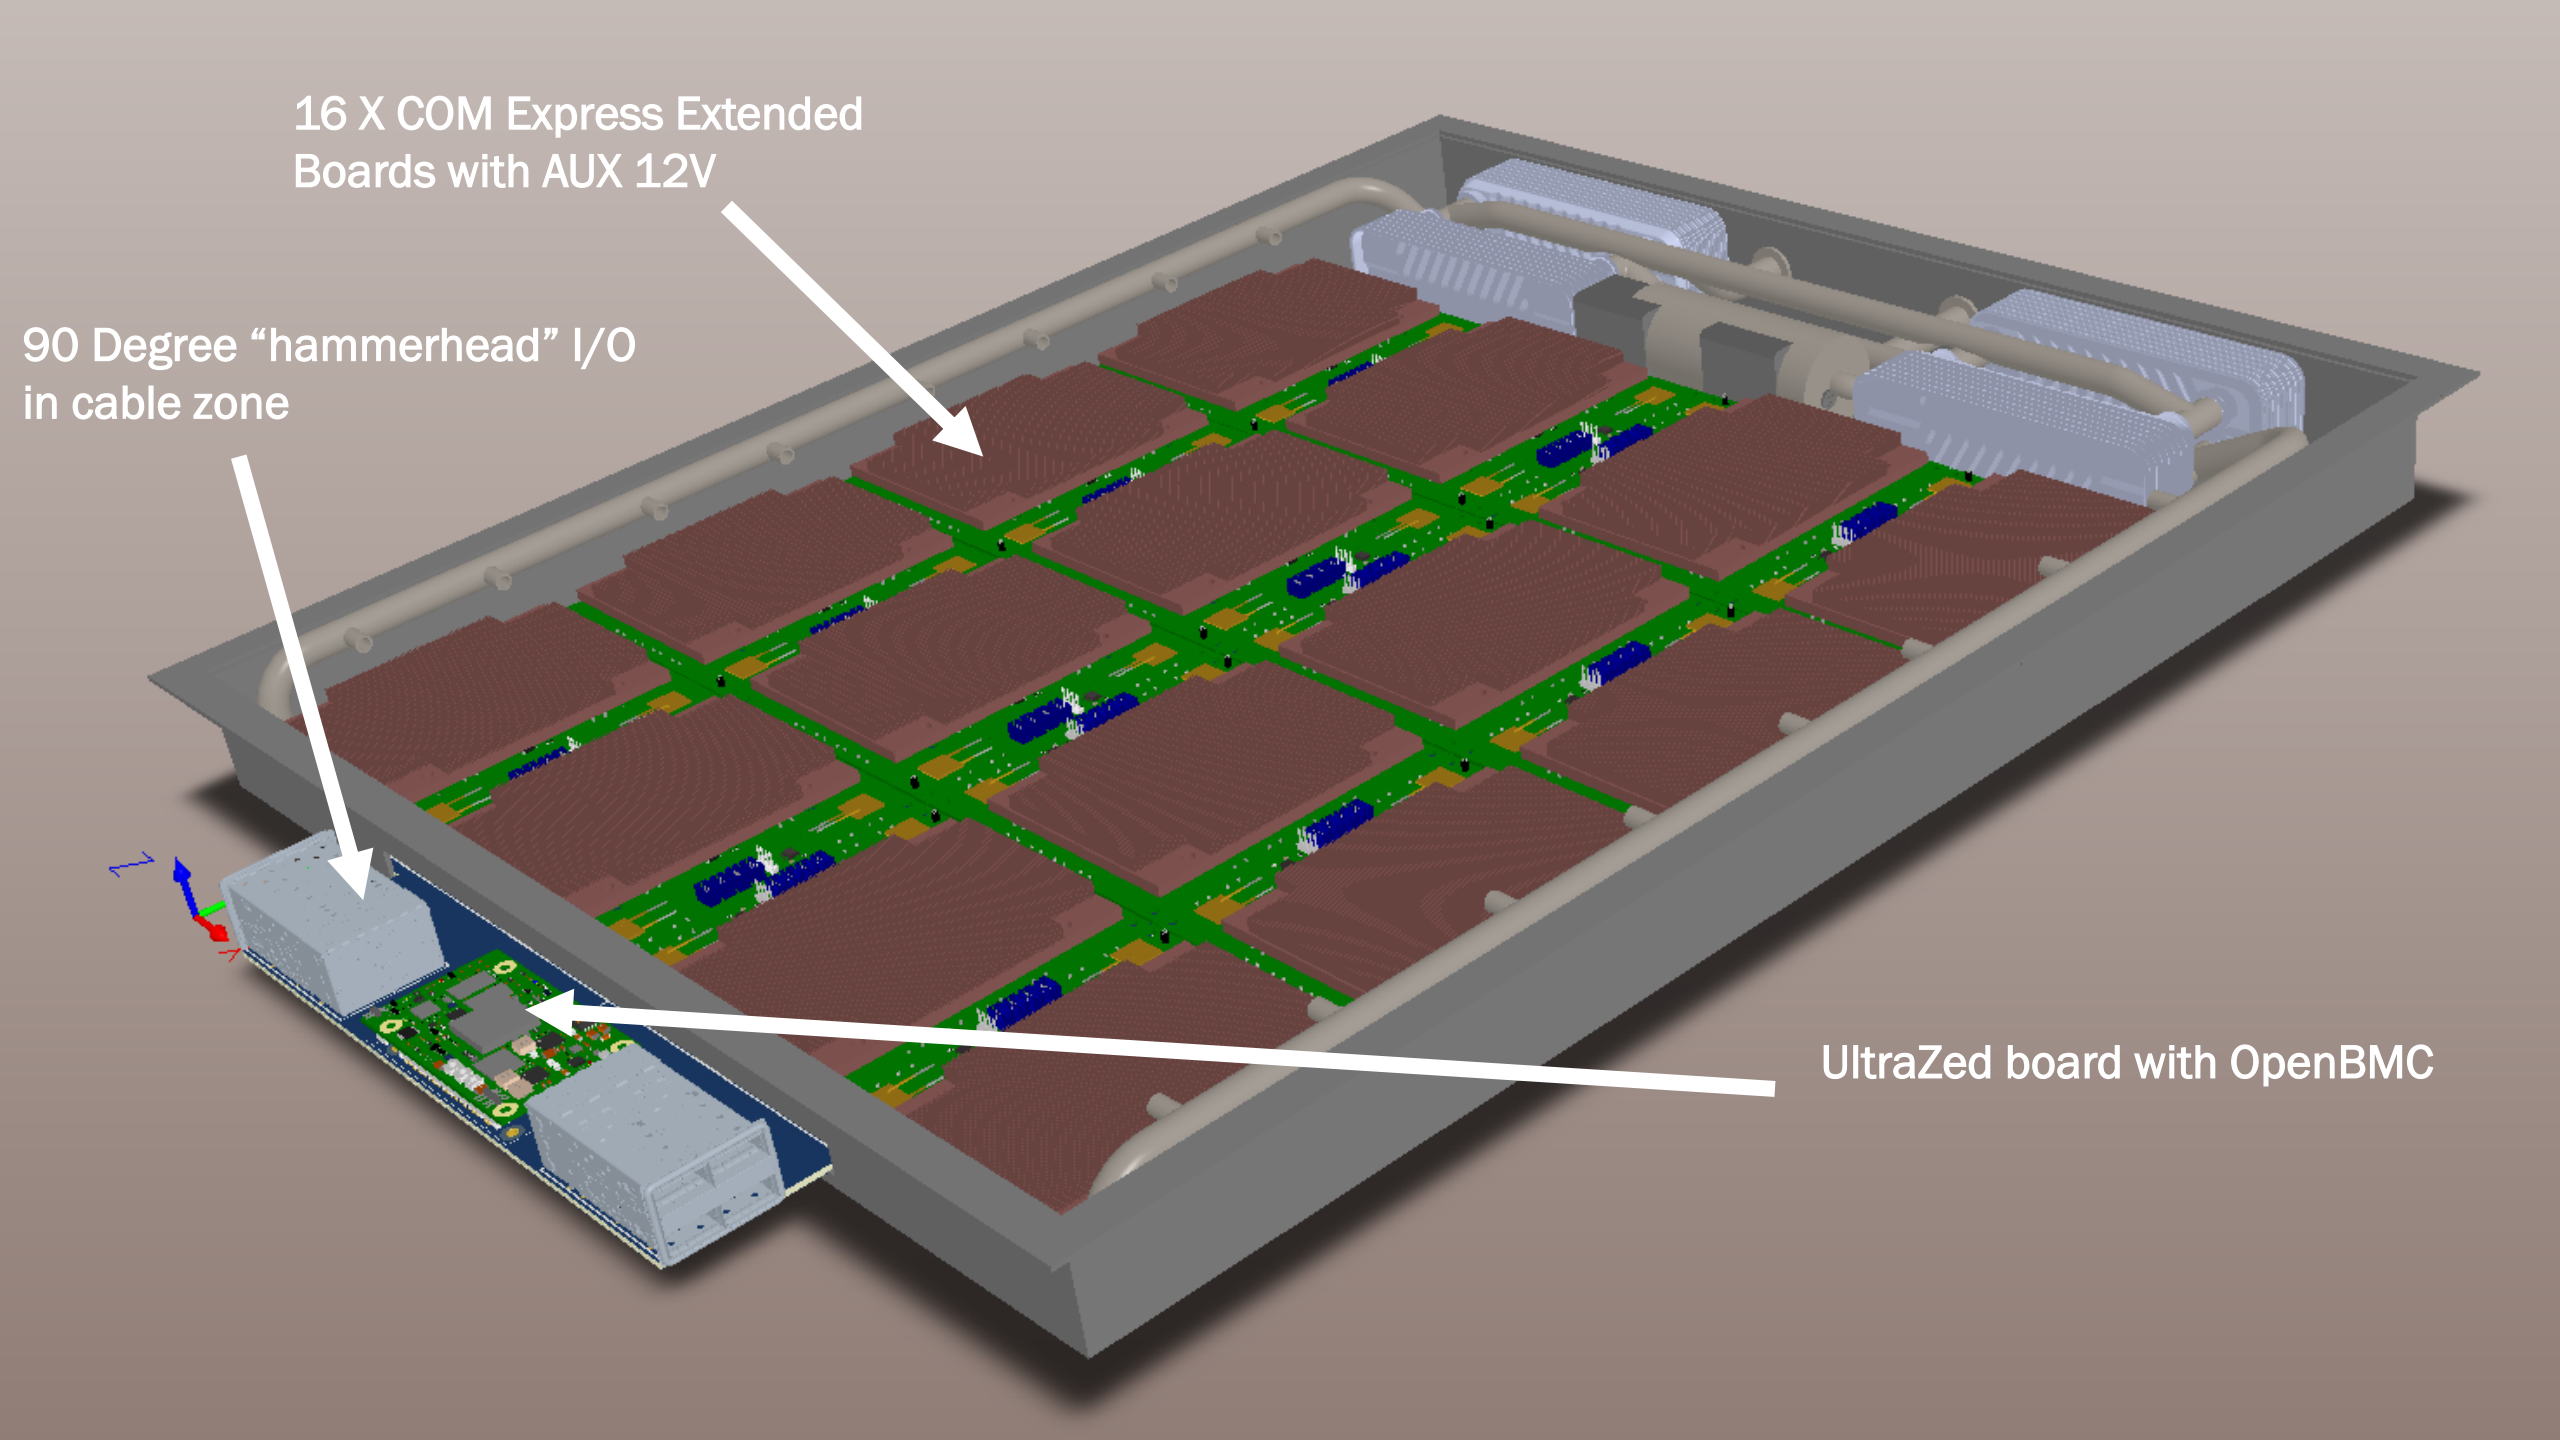
16 X COM Express Extended
Boards with AUX 12V
90 Degree “hammerhead” I/O
in cable zone
UltraZed board with OpenBMC

EURO EXA

FUNDED BY THE EUROPEAN UNION

60MW in 35m Diameter

10 40ft High-Cube ISOs per layer

3 Layers

2MW per ISO

Approx. 110kW per Rack

Prototype will feature 1 container, initially with 3 live Racks

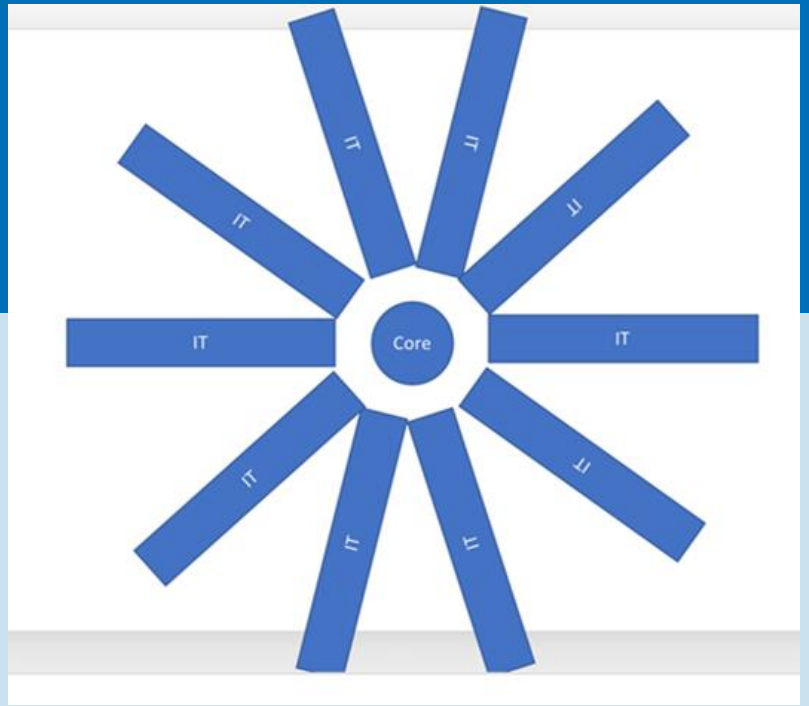

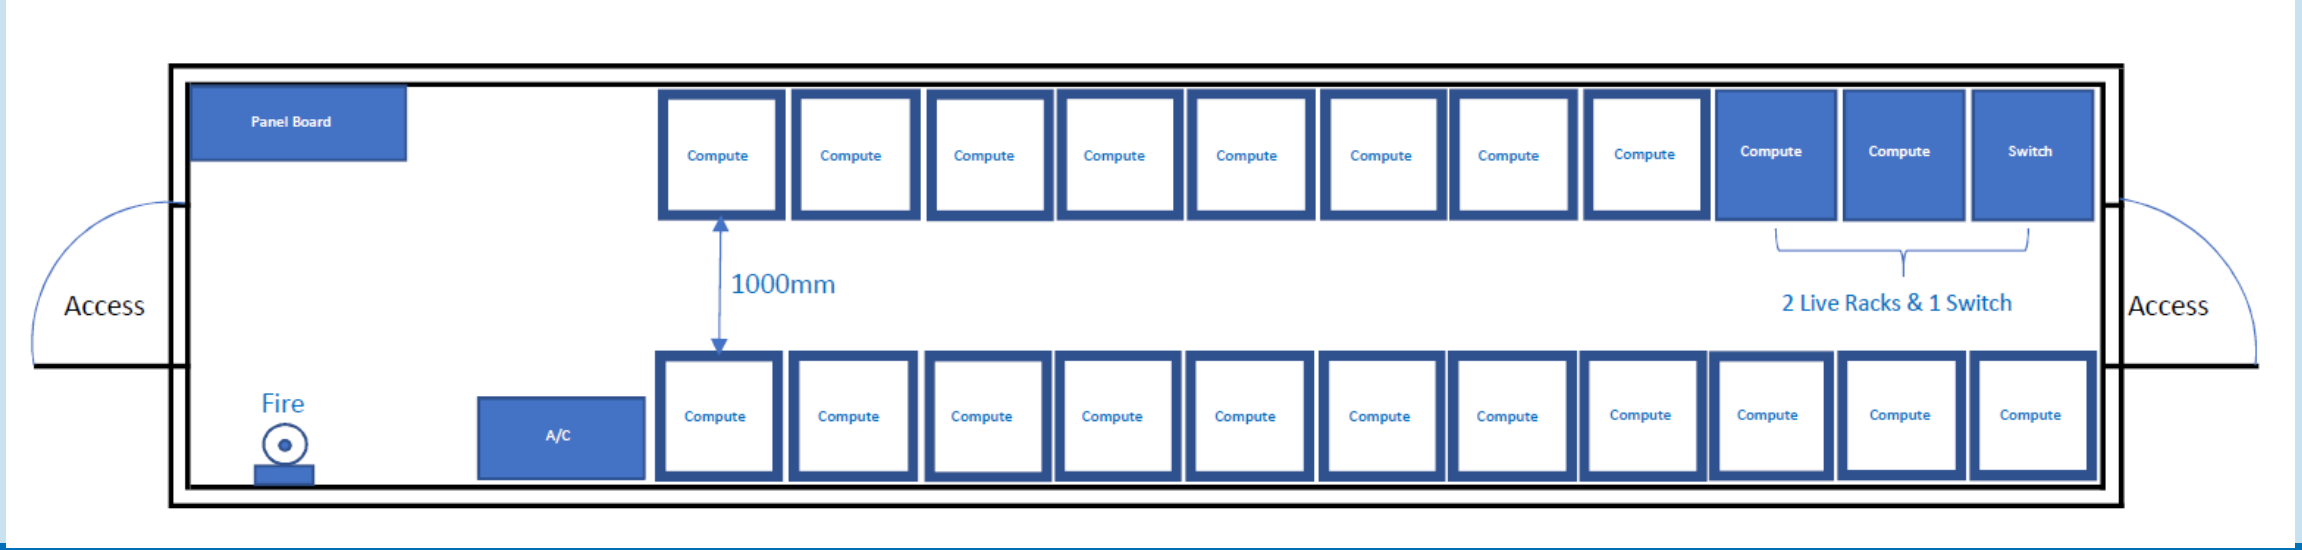

EURO EXA

Liquid Cooling Objectives

- 45C Inlet “Chiller-less” ASHRAE W4/W5
- Ultra-Low Risk Server Coolant System
  - PFPE Coolants, small quantity, non conductive, low dielectric constant.
  - Low Pressure Drop Heat Exchange
- Ultra-Low Risk Facility Coolant System
  - Sub Atmospheric Leak Prevention System

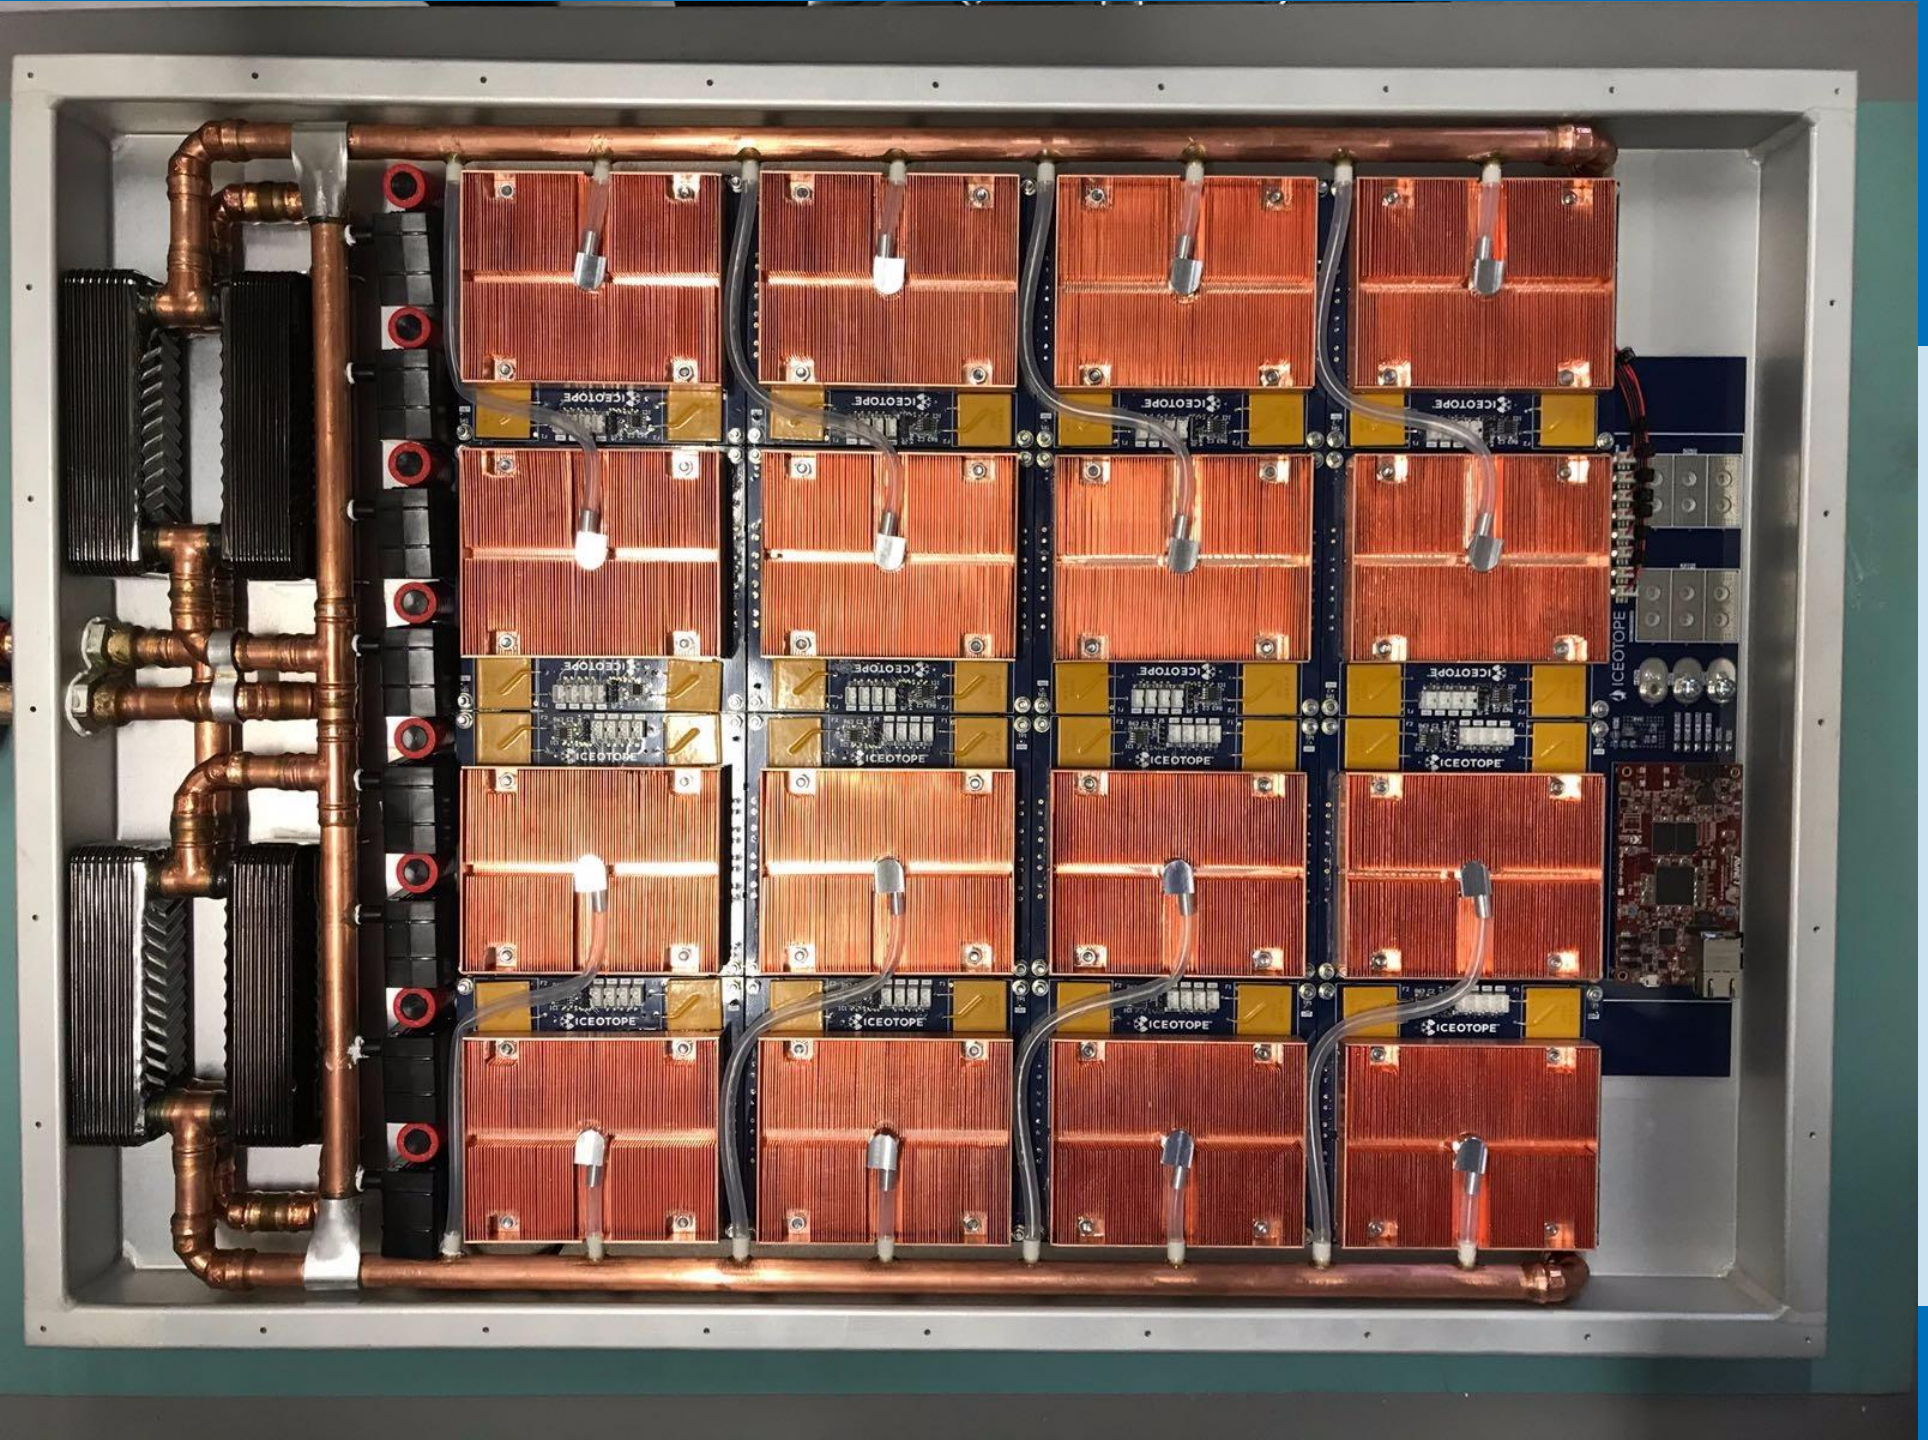

- TED – Test Rig
- Thermal-proxy Extra-high Density
- 4kW in 1oU test capacity
- 16 Little Modules

EURO EXA
FUNDED BY THE EUROPEAN UNION

- Immersion without Submersion
- Achieved:
  47C inlet
  55C outlet
  @3.2kW/OU

Rack-Centre, Lagos Nigeria

Many Thanks
Please Connect at thecoolingguy.me

Presenter – Peter Hopton
Dissemination Lead, EuroEXA
Founder, Iceotope
http://thecoolingguy.me

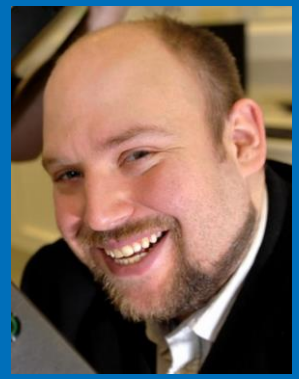

Project ID: 754337