EXASCALE COMPUTING

Der Sprung in eine neue Ära des Höchstleistungsrechnens

11.11.2022 I THOMAS LIPPERT, ANDREAS HERTEN, ROBERT SPECK UND DAS GANZE JSC

Baseline

- What kind of CPU does your computer have?
  CPU generation, clock ~~speed~~ rate, number of cores, vector length

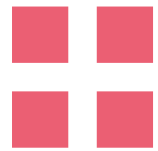

Baseline

- What kind of CPU does your computer have?
  CPU generation, clock ~~speed~~ rate, number of cores, vector length
- How much memory does your computer have?
  Amount of memory, type, links (GB and GB/s)

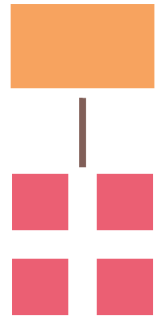

Baseline

- What kind of CPU does your computer have?
  CPU generation, clock ~~speed~~ rate, number of cores, vector length
- How much memory does your computer have?
  Amount of memory, type, links (GB and GB/s)
- What kind of GPU do you have?
  GPU generation, number of cores, watt intake (TDP)

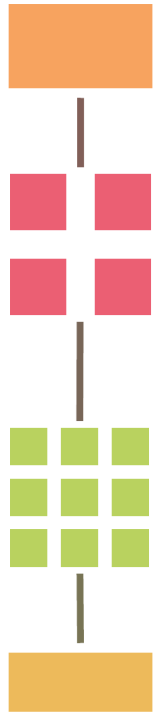

Baseline

- What kind of CPU does your computer have?
  CPU generation, clock ~~speed~~ rate, number of cores, vector length
- How much memory does your computer have?
  Amount of memory, type, links (GB and GB/s)
- What kind of GPU do you have?
  GPU generation, number of cores, watt intake (TDP)
- How fast is your network ?
  Throughput, latency

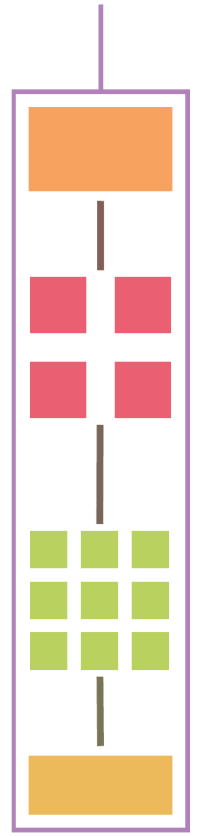

Baseline

- What kind of CPU does your computer have?
  CPU generation, clock ~~speed~~ rate, number of cores, vector length
- How much memory does your computer have?
  Amount of memory, type, links (GB and GB/s)
- What kind of GPU do you have?
  GPU generation, number of cores, watt intake (TDP)
- How fast is your network ?
  Throughput, latency

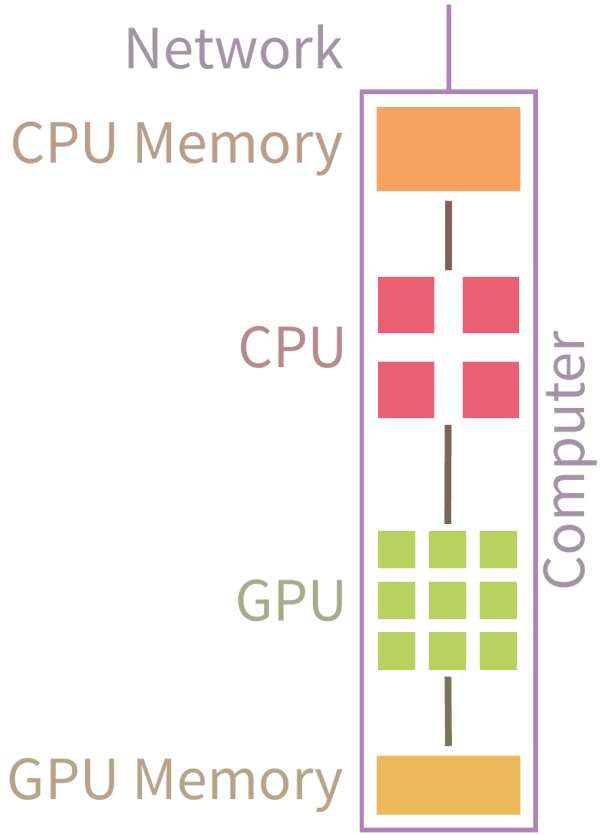

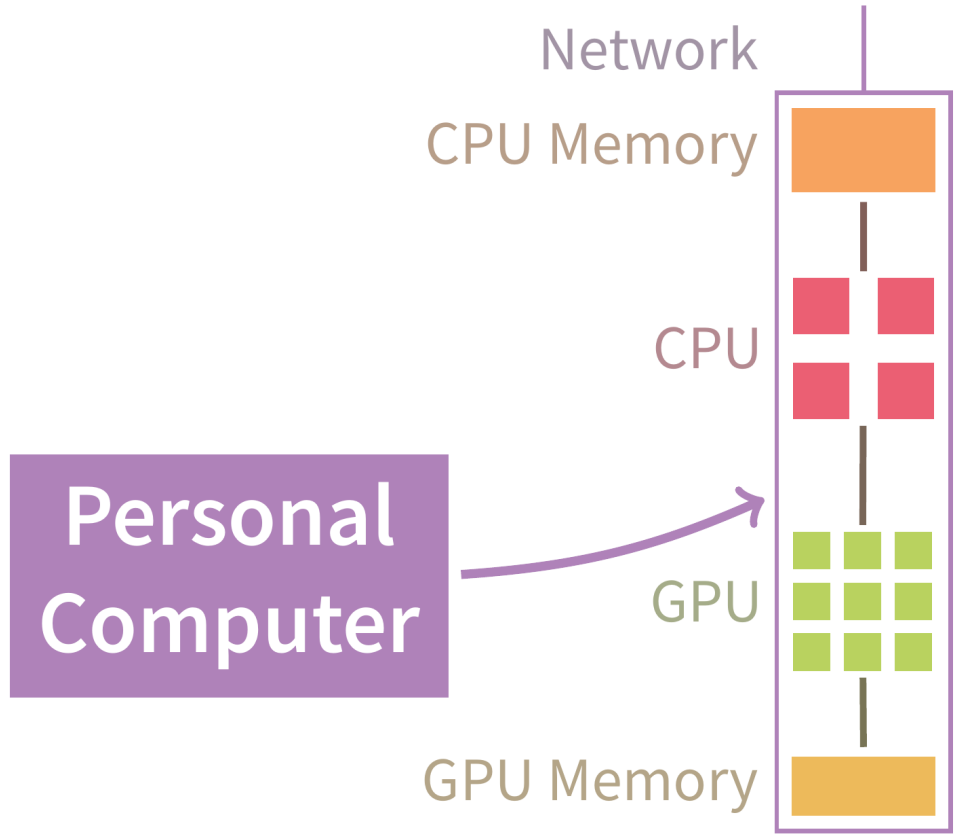
Network
CPU Memory
CPU
Personal Computer
GPU
GPU Memory

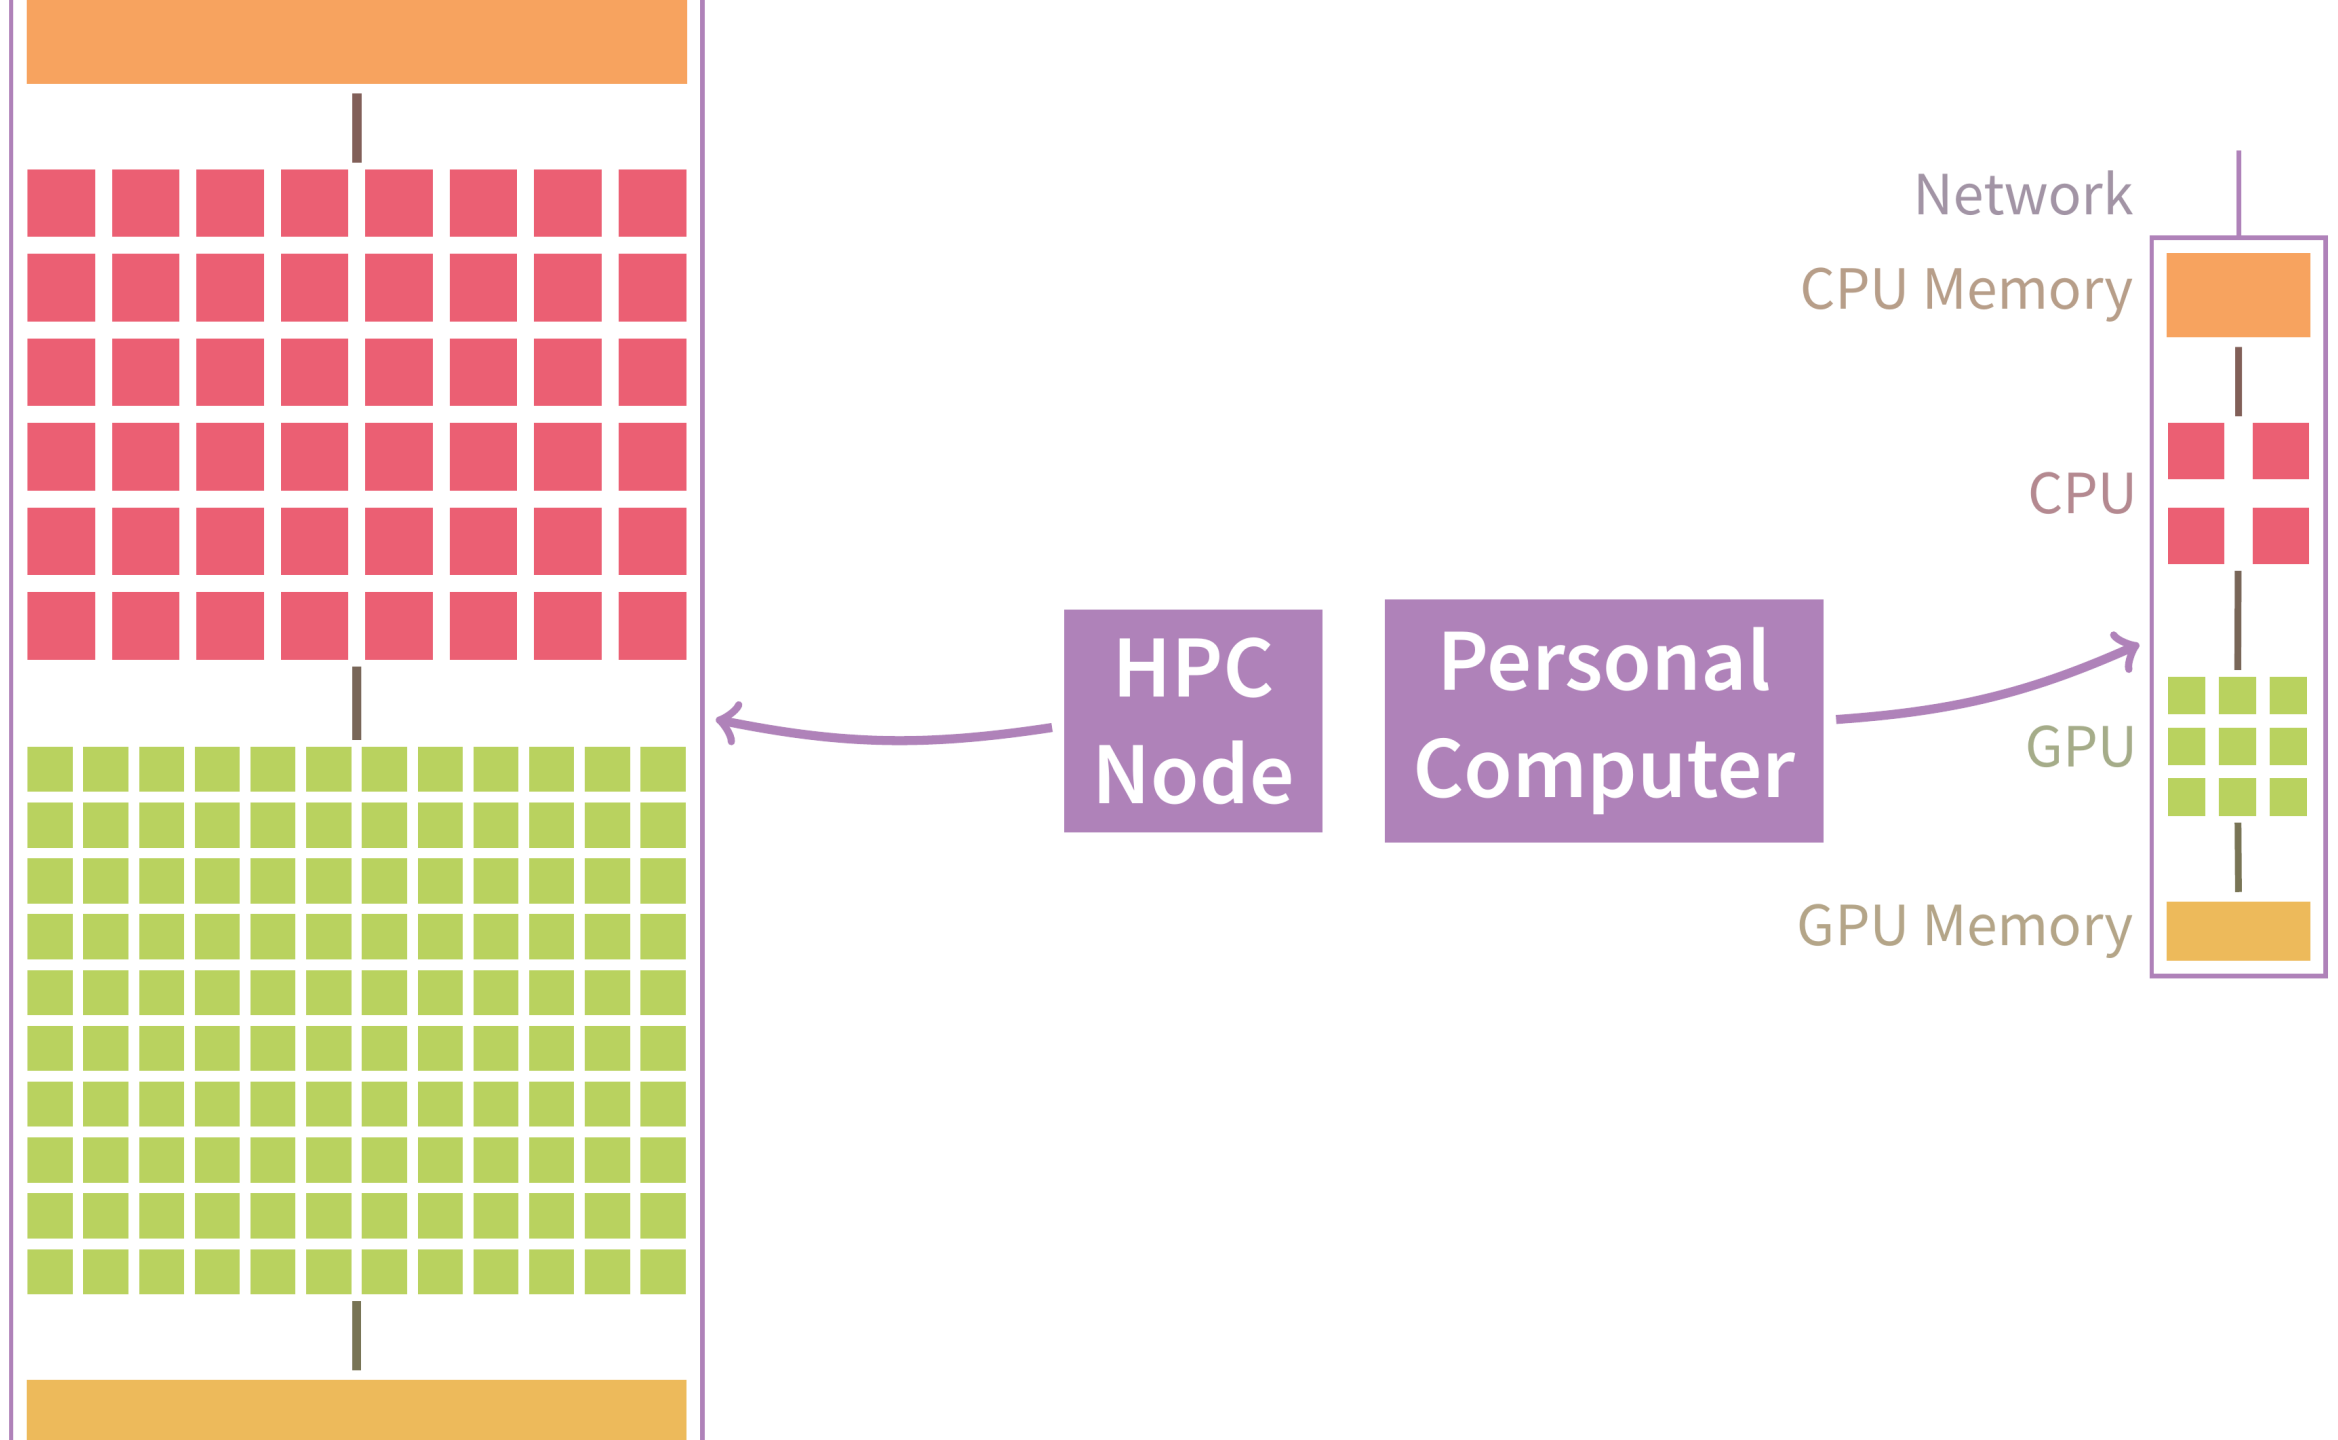
Network
CPU Memory
CPU
GPU
GPU Memory
HPC Node
Personal Computer

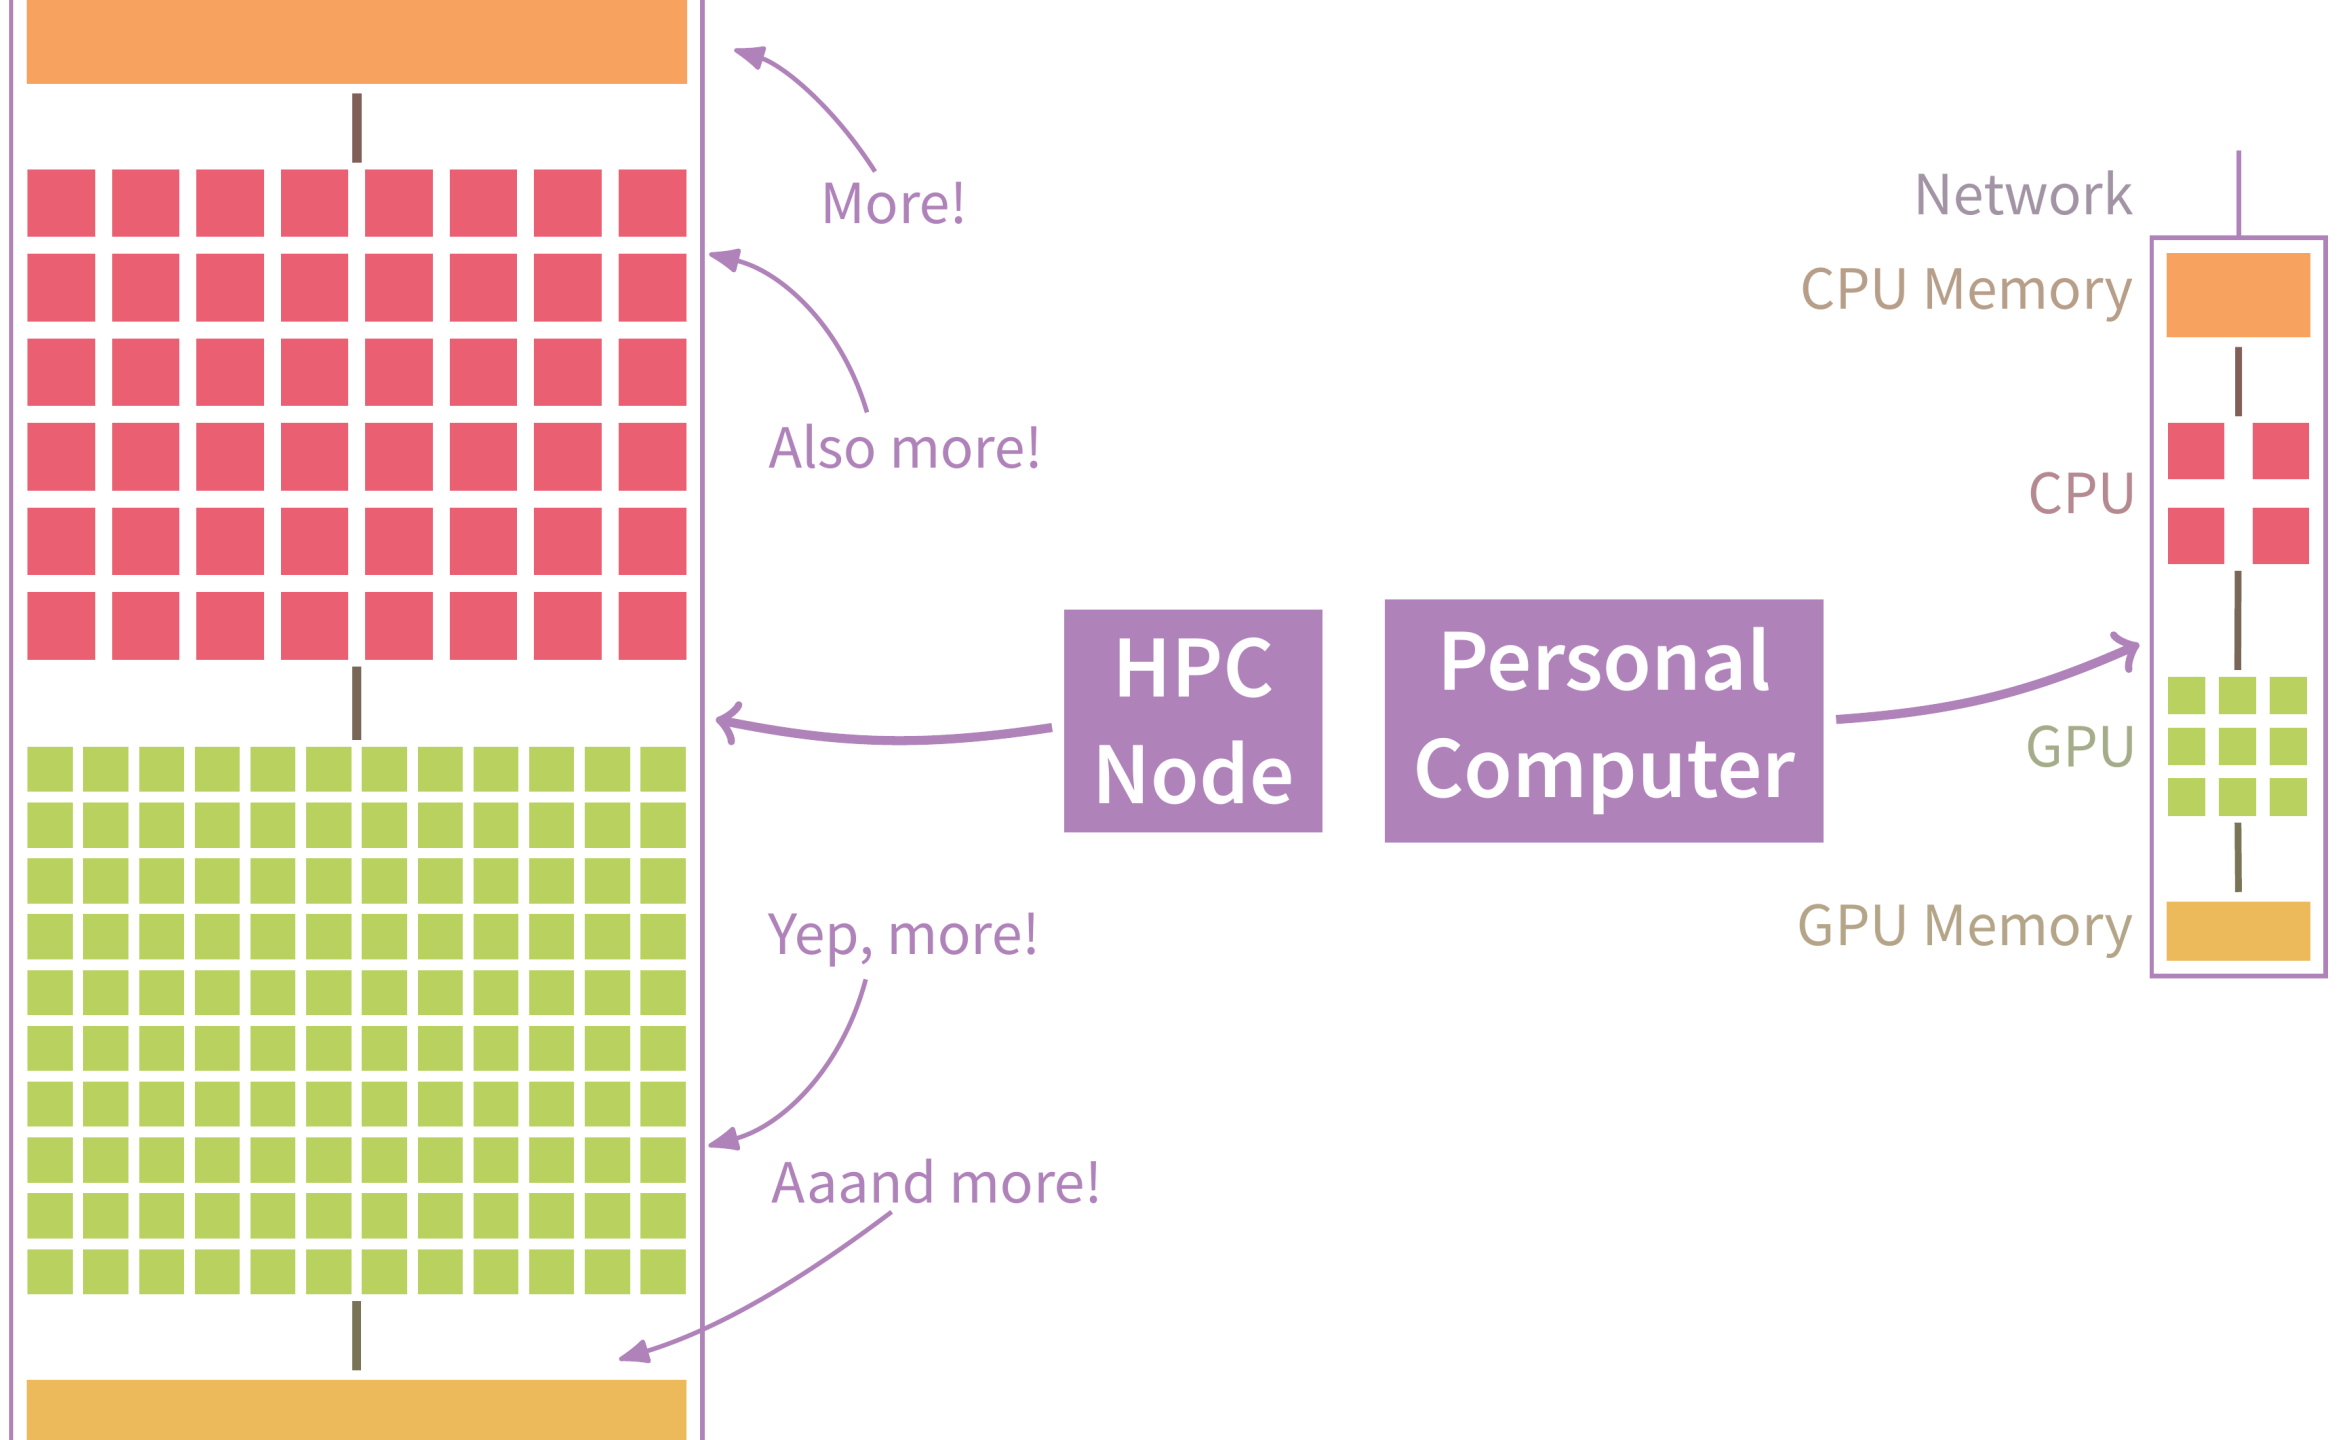
More!
Also more!
HPC Node
Personal Computer
Yep, more!
Aaand more!
Network
CPU Memory
CPU
GPU
GPU Memory

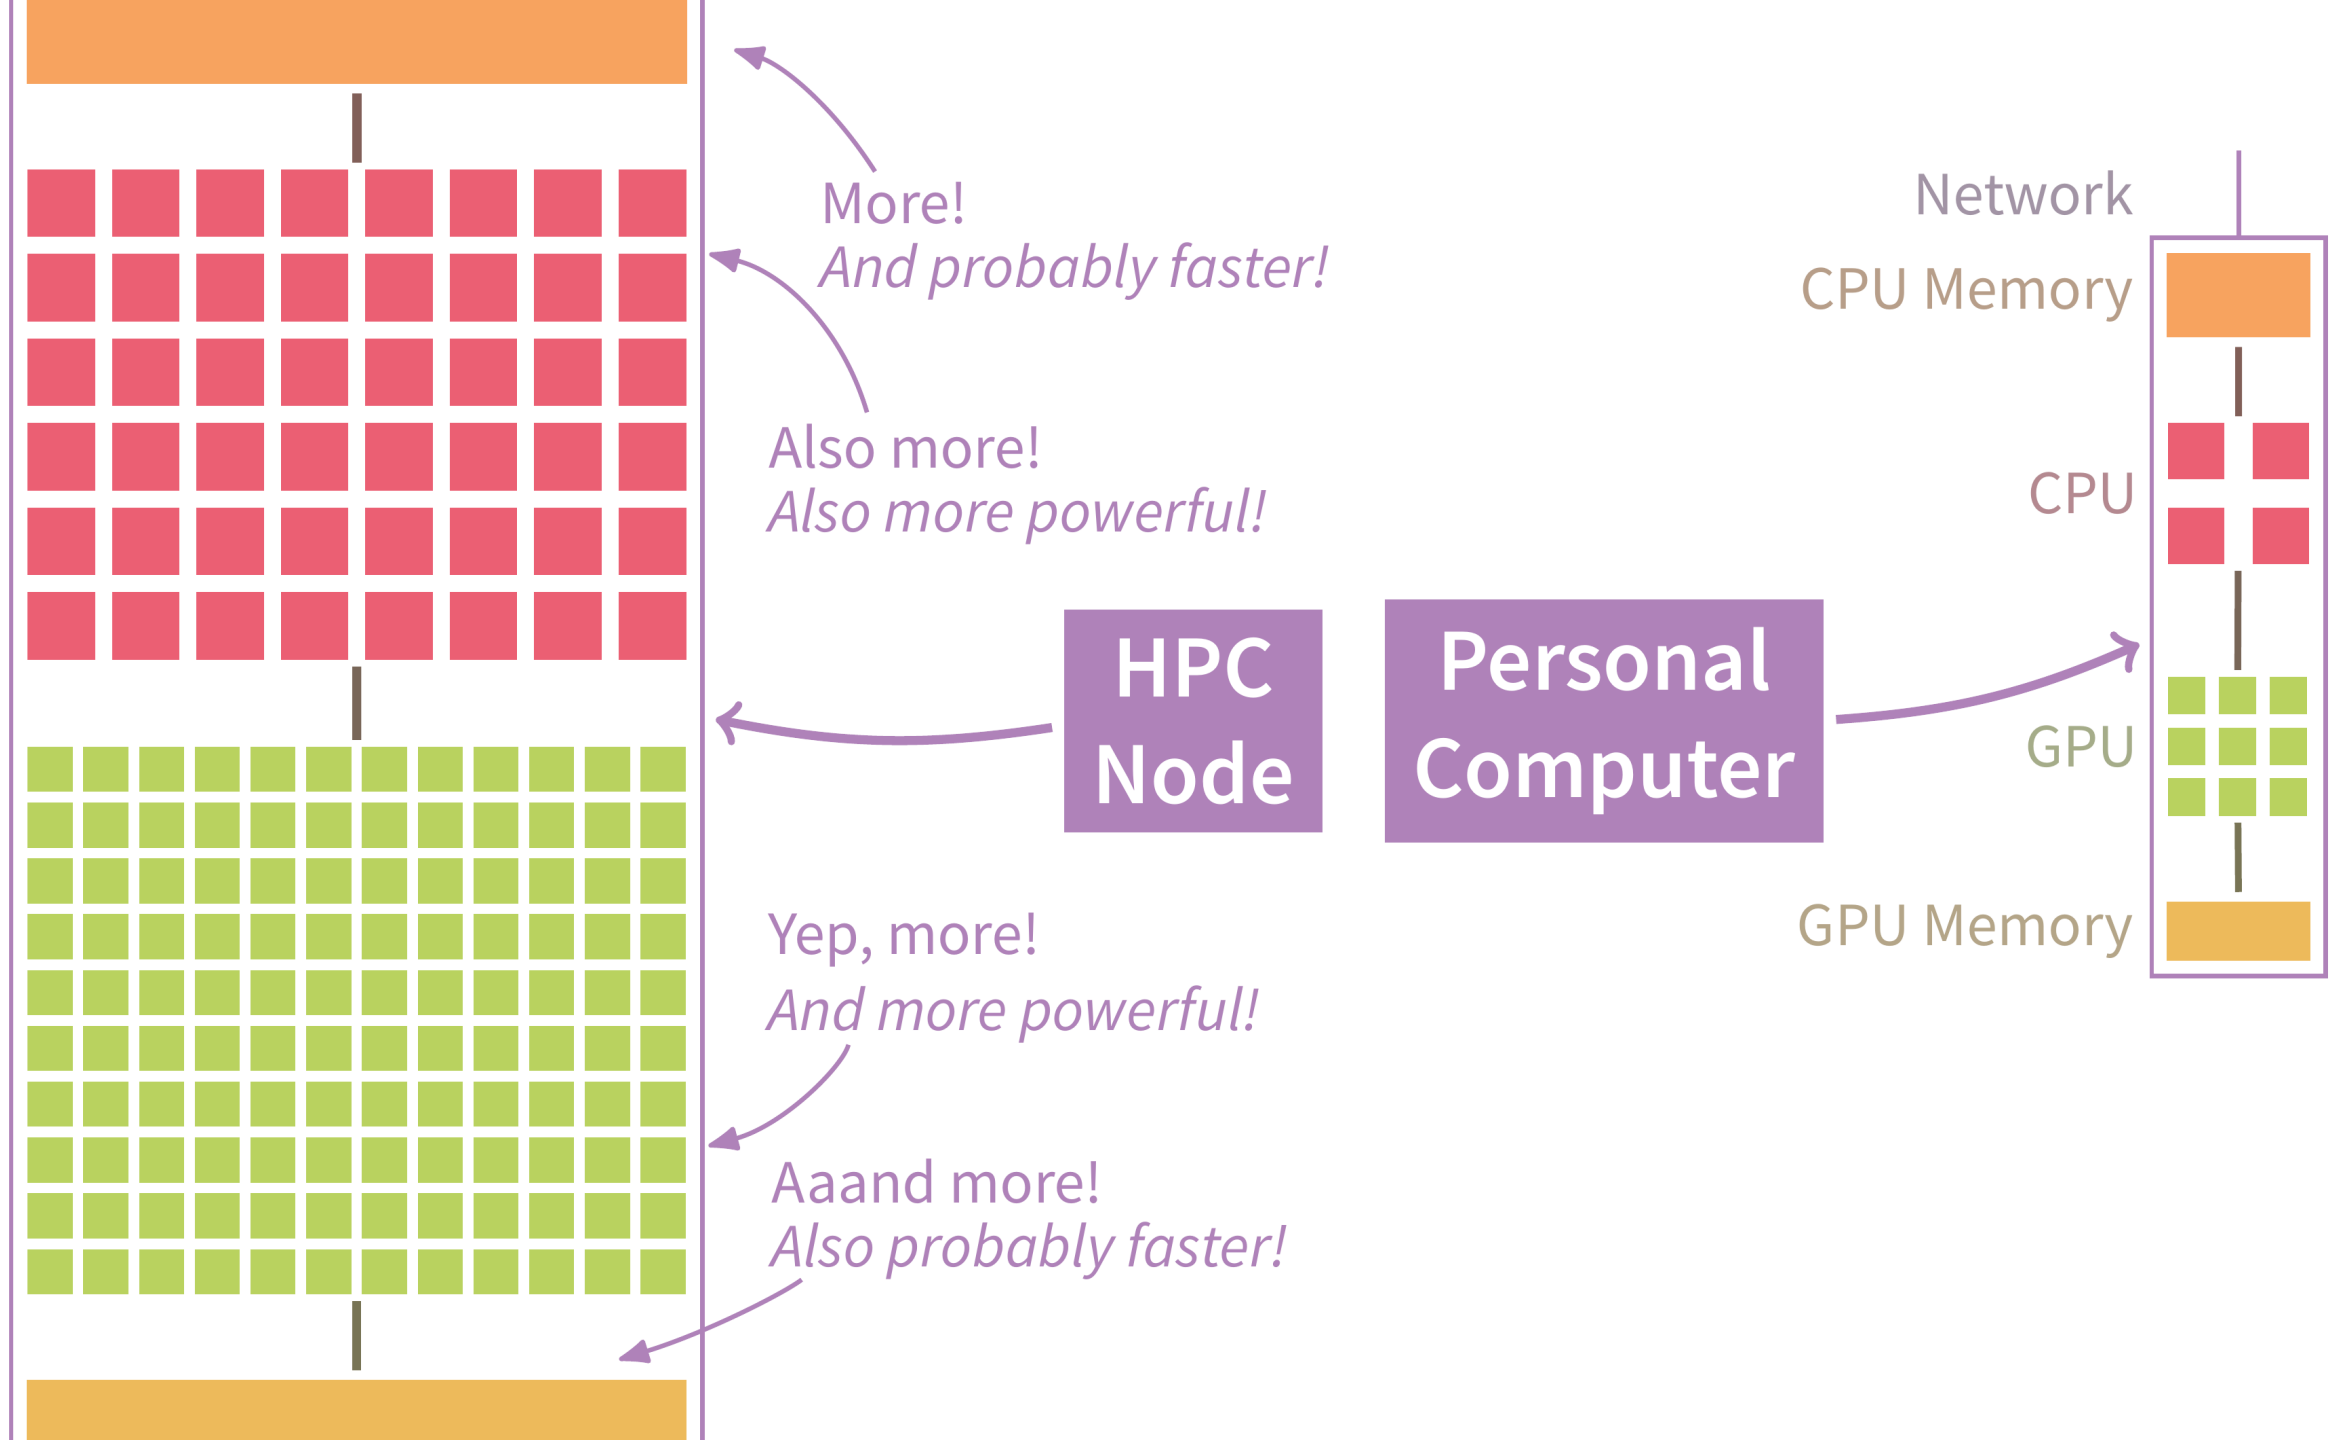
More!
And probably faster!
Also more!
Also more powerful!
HPC Node
Personal Computer
Network
CPU Memory
CPU
GPU
GPU Memory
Yep, more!
And more powerful!
Aaand more!
Also probably faster!

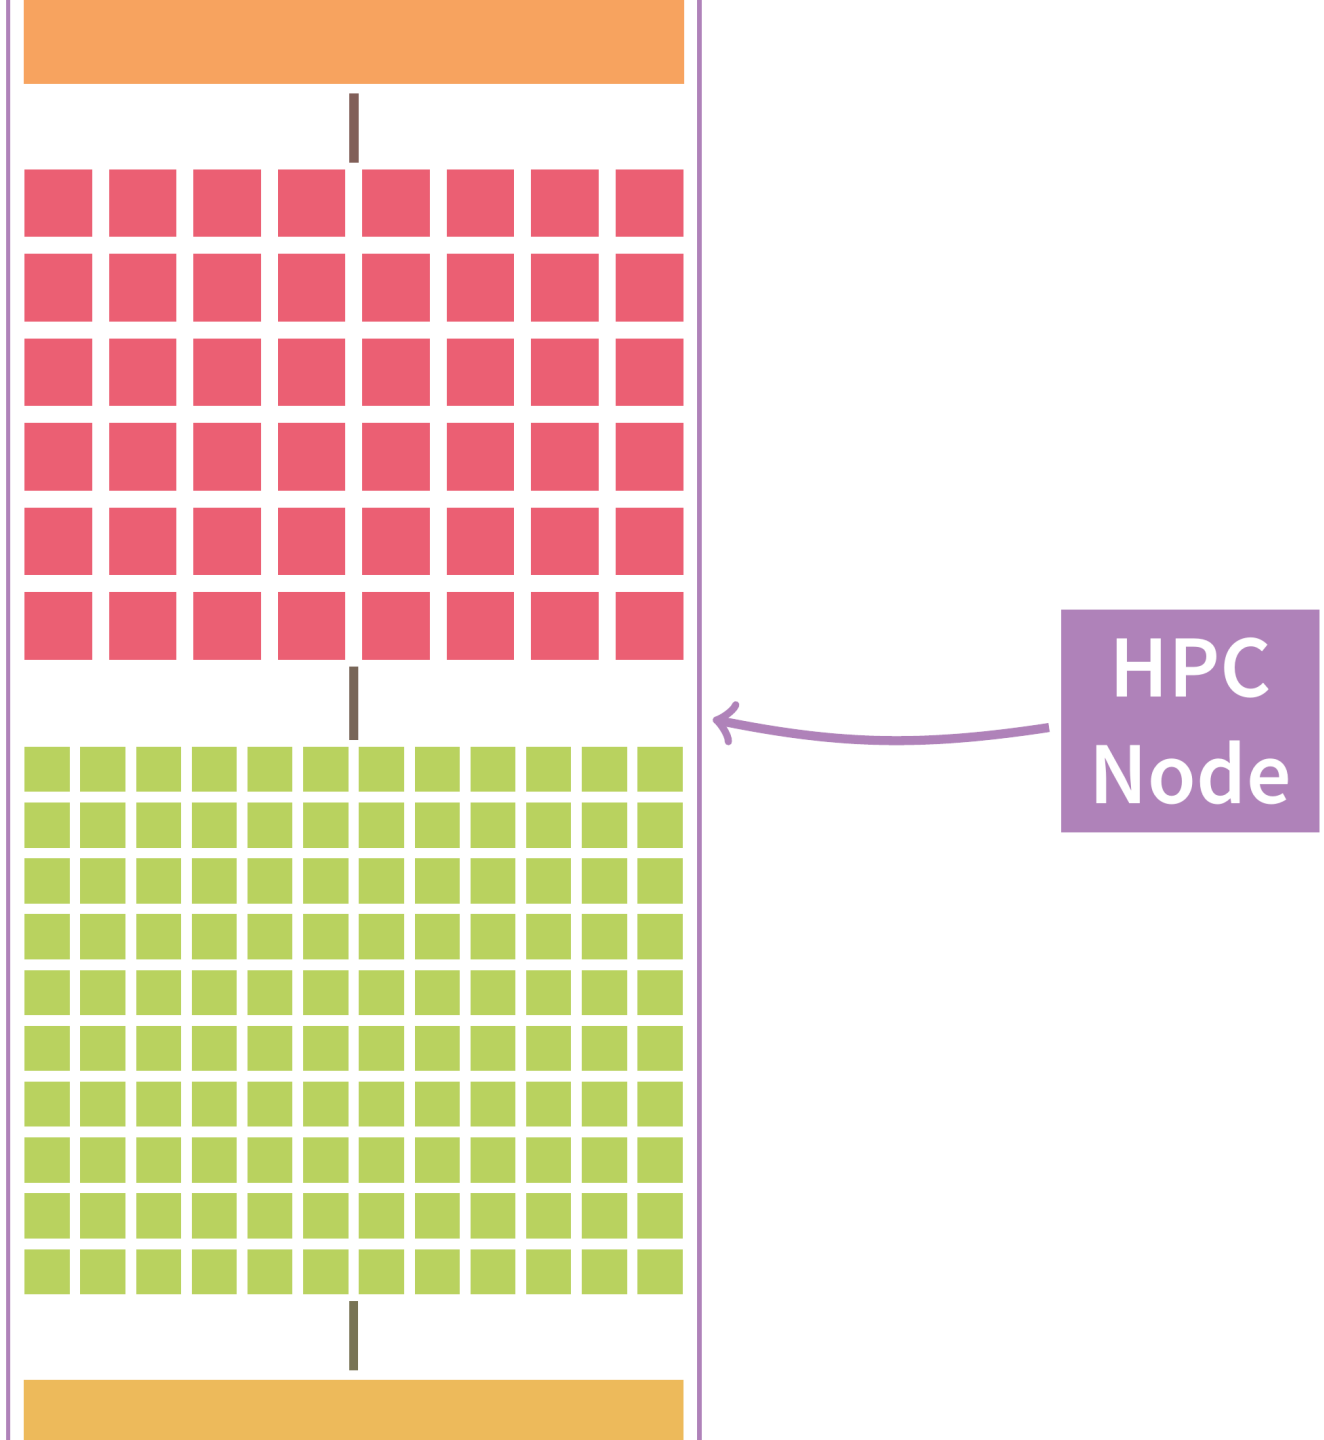
HPC
Node

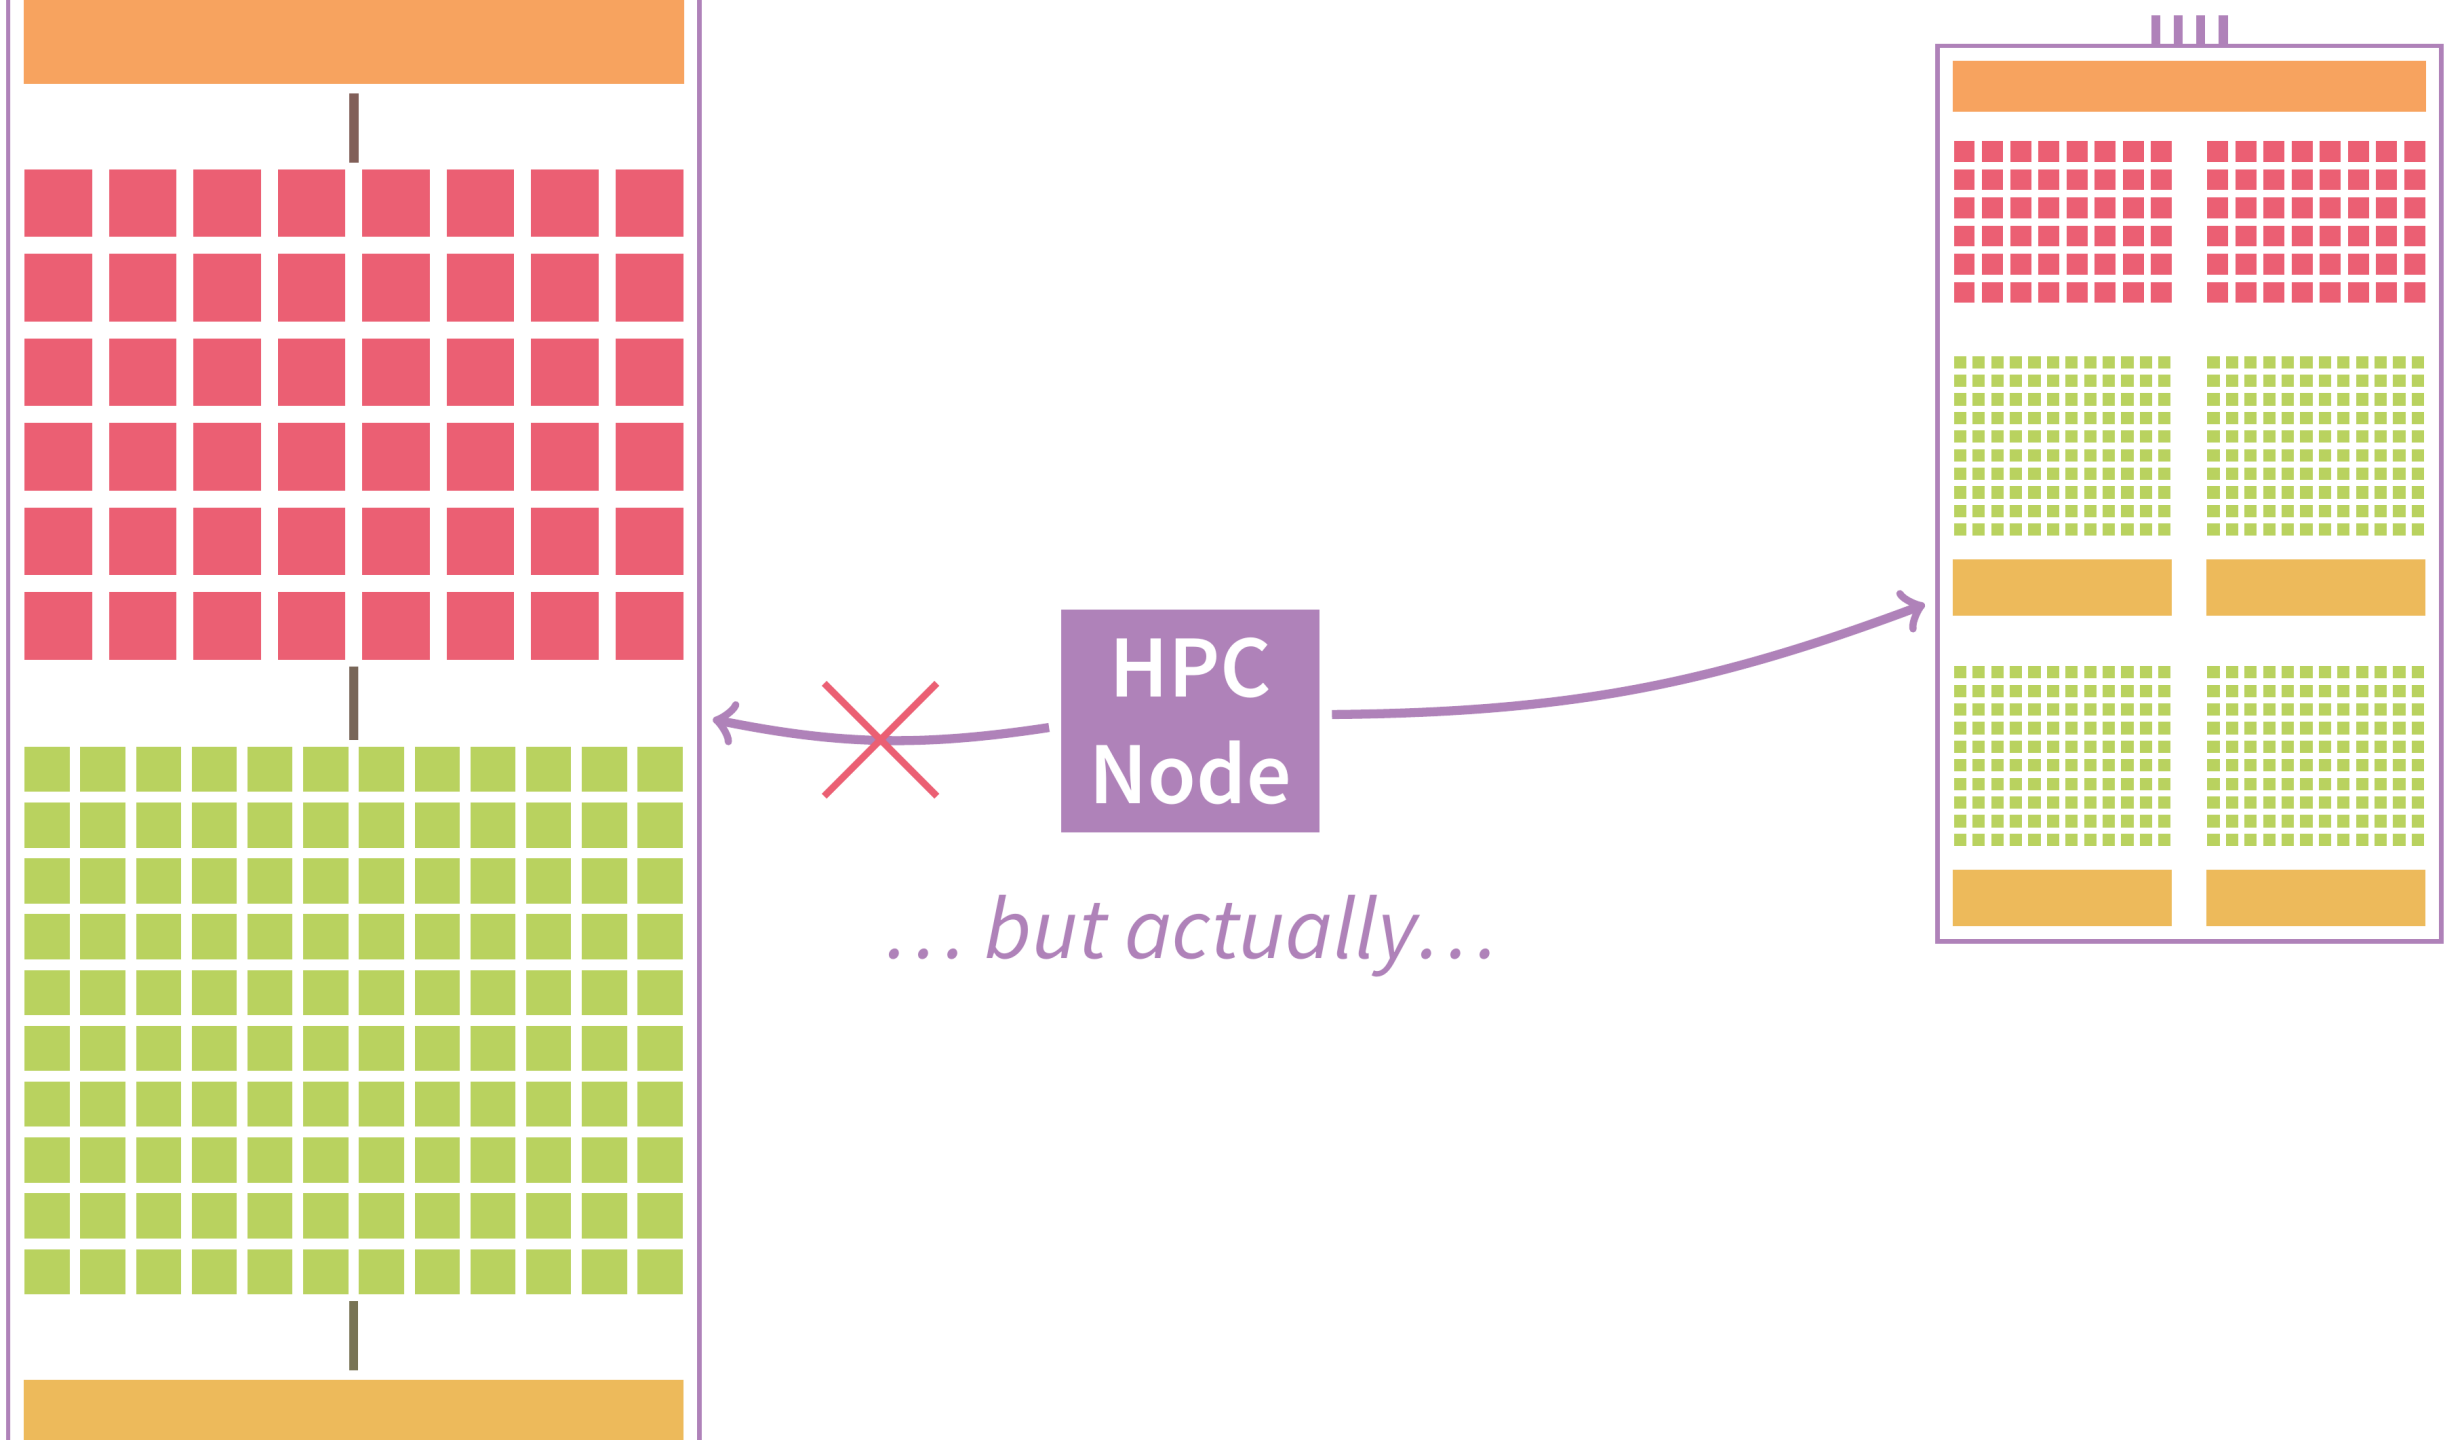
HPC
Node
… but actually…

- Usually, 2 CPUs **sockets**, each with 64 cores; use mostly as one CPU with one memory
- 4 distinct GPUs, connected with each other (600 GB/s)
- 4 network connections, each 200 Gbit/s in each direction (*InfiniBand HDR-200*)

HPC Node

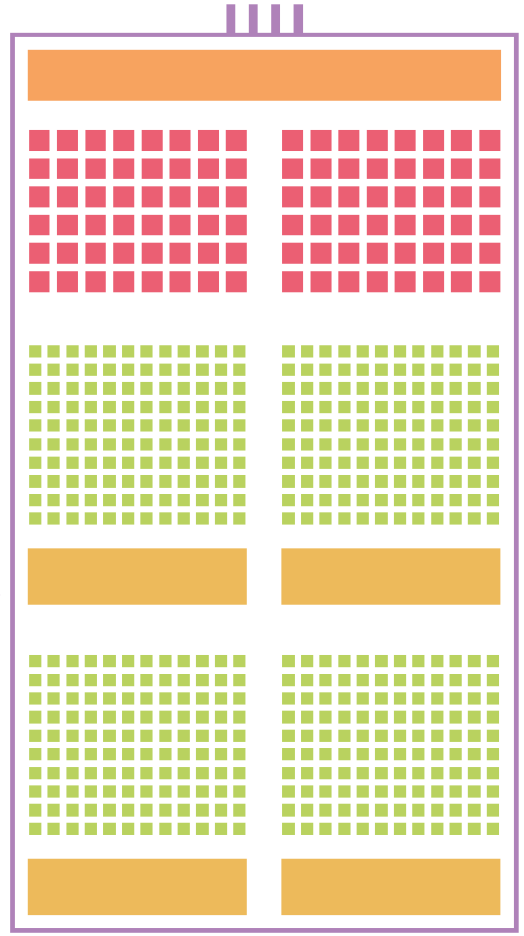

Specs based on JURECA DC, JUWELS

HPC Node

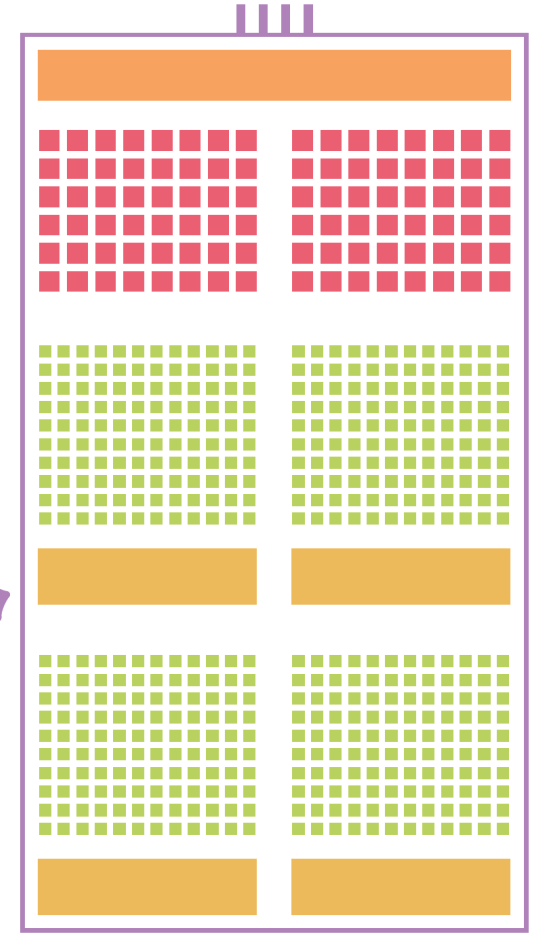

Specs based on JURECA DC, JUWELS

HPC Nodes

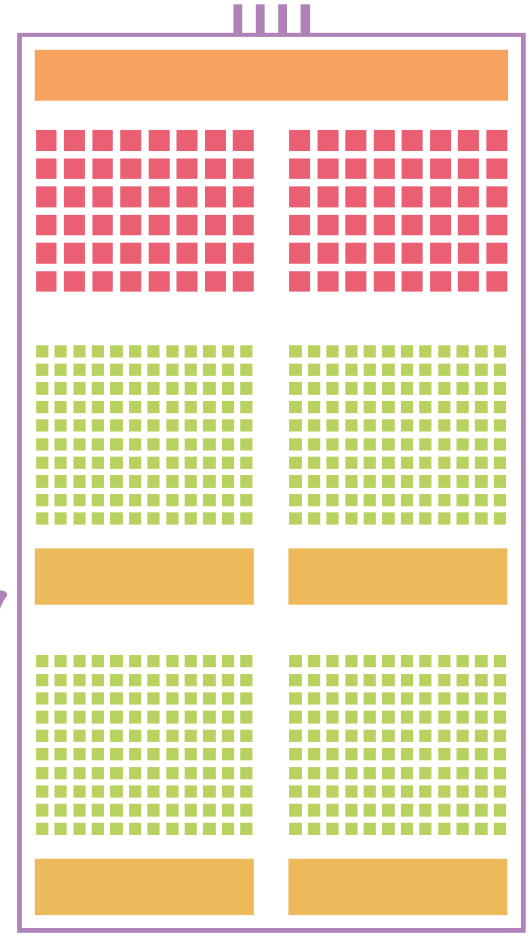

Specs based on JURECA DC, JUWELS

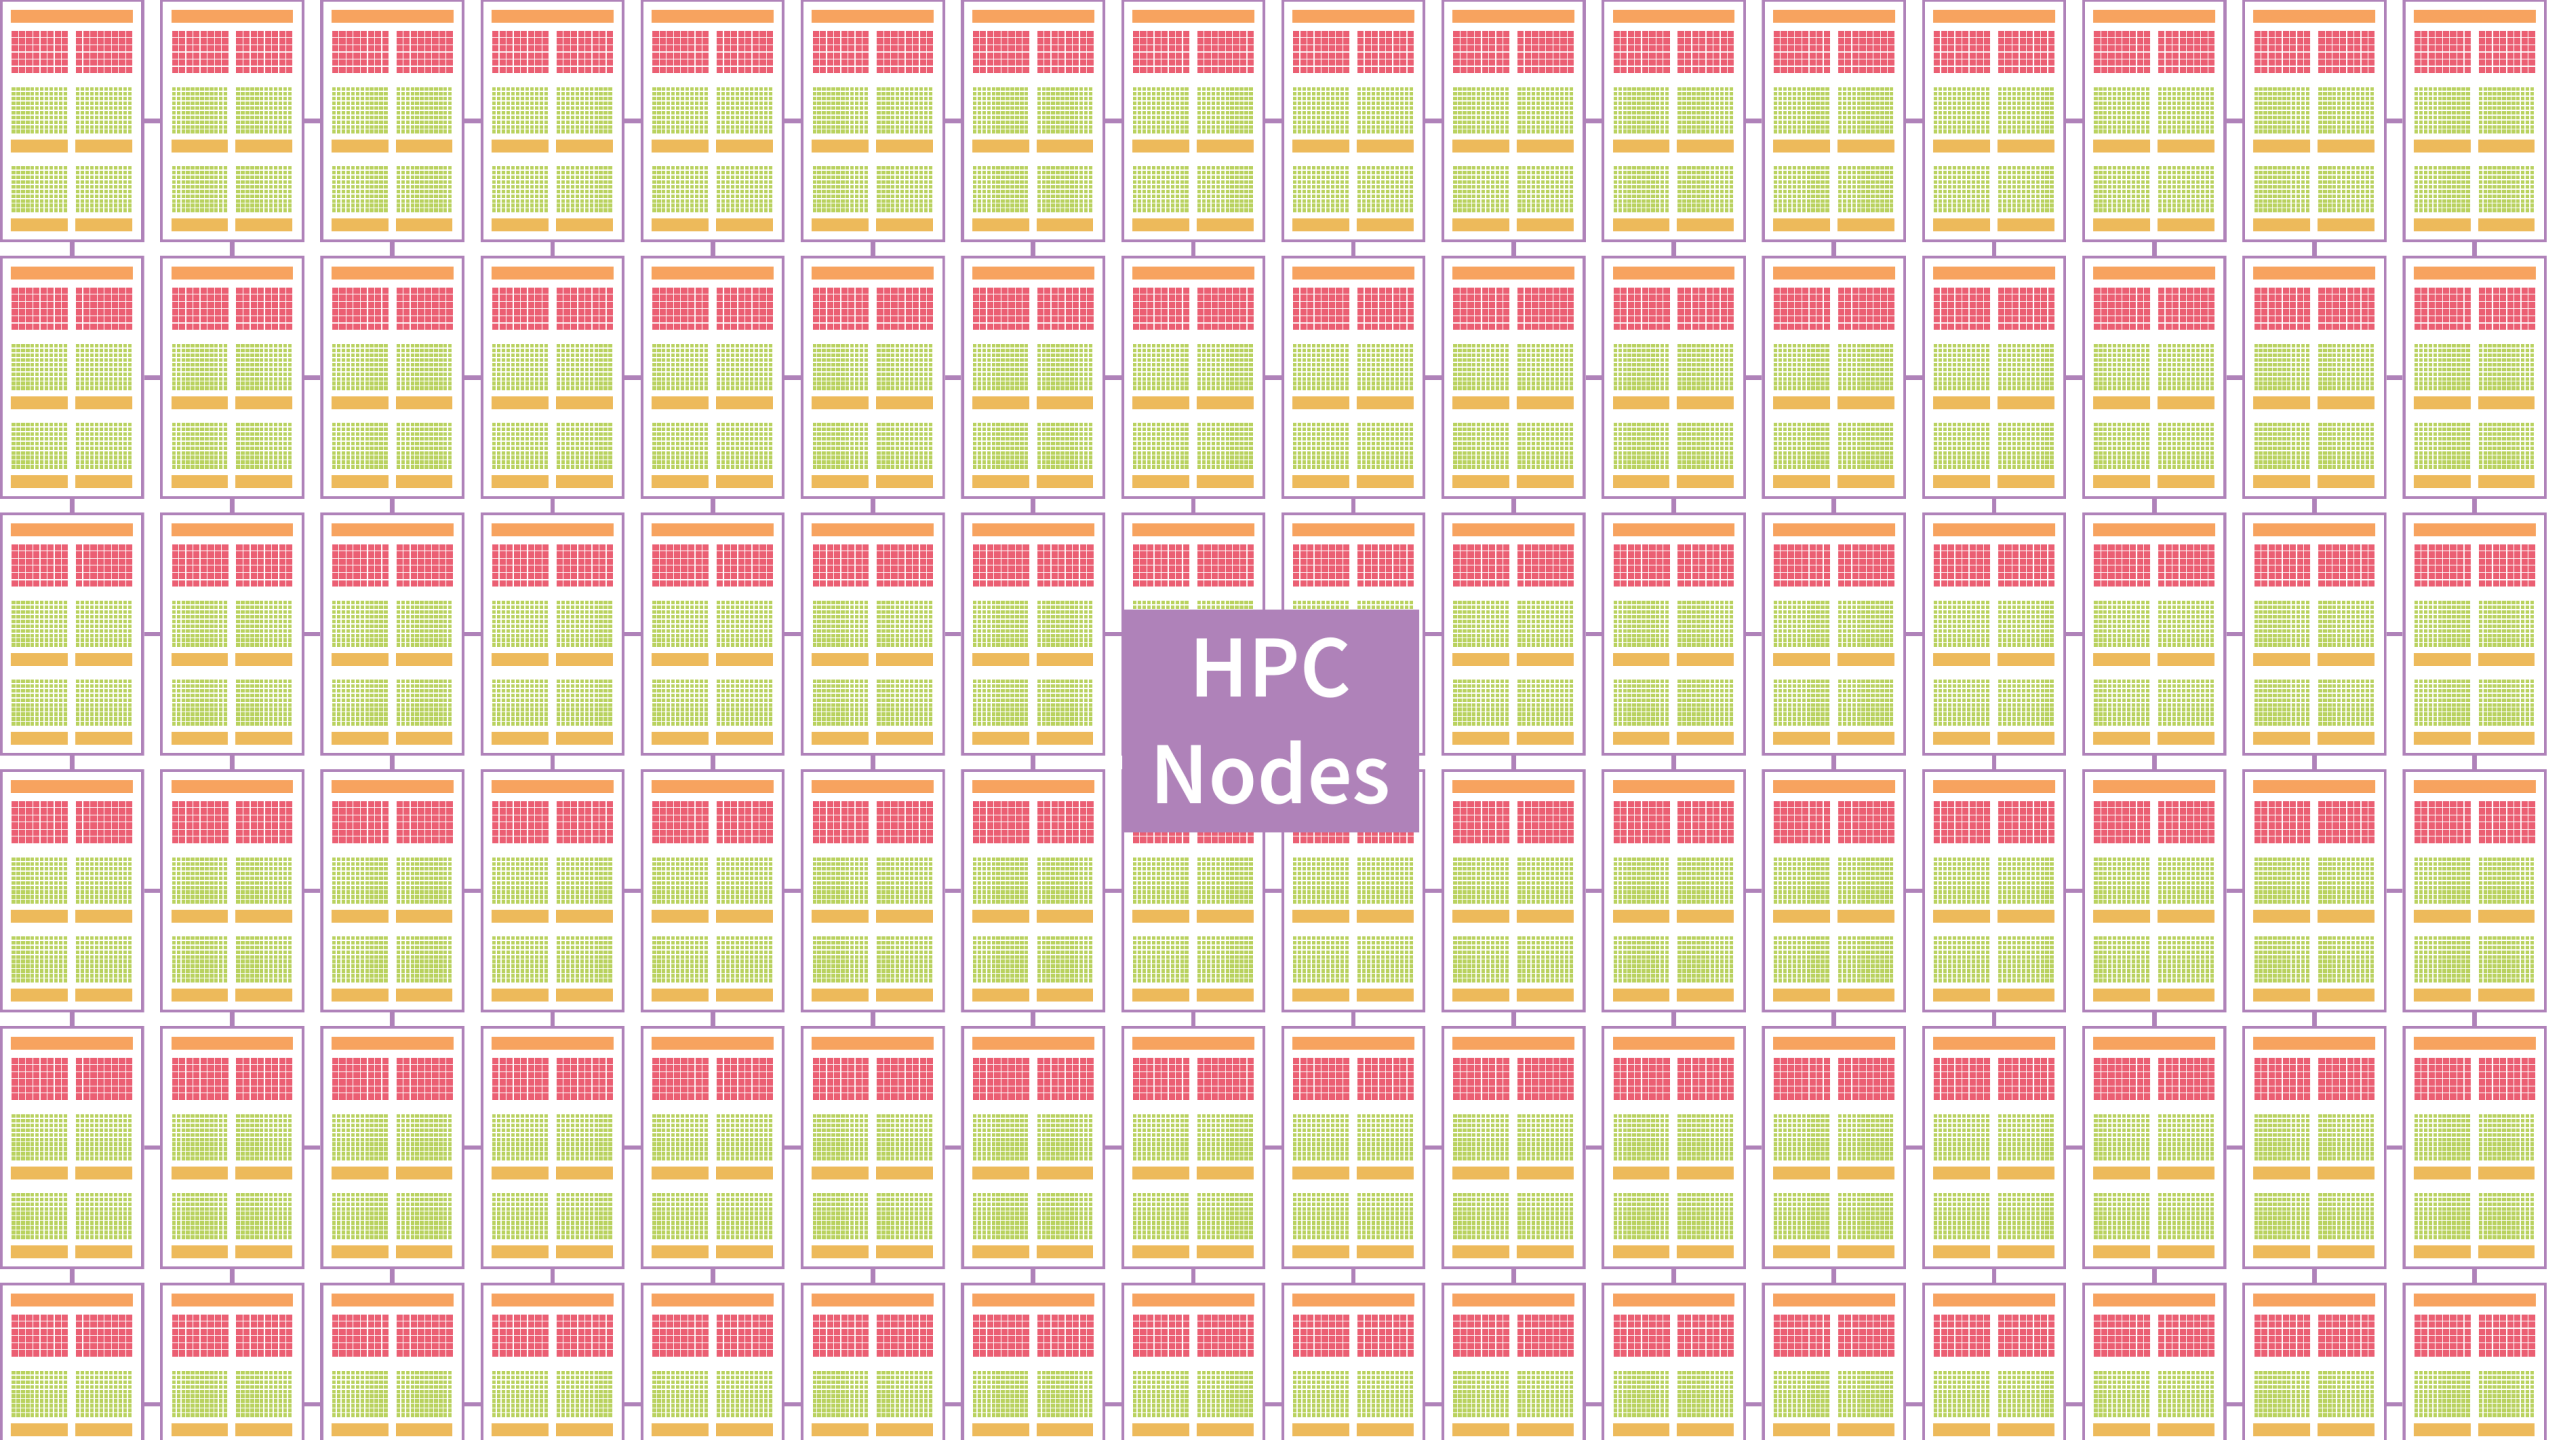
HPC
Nodes

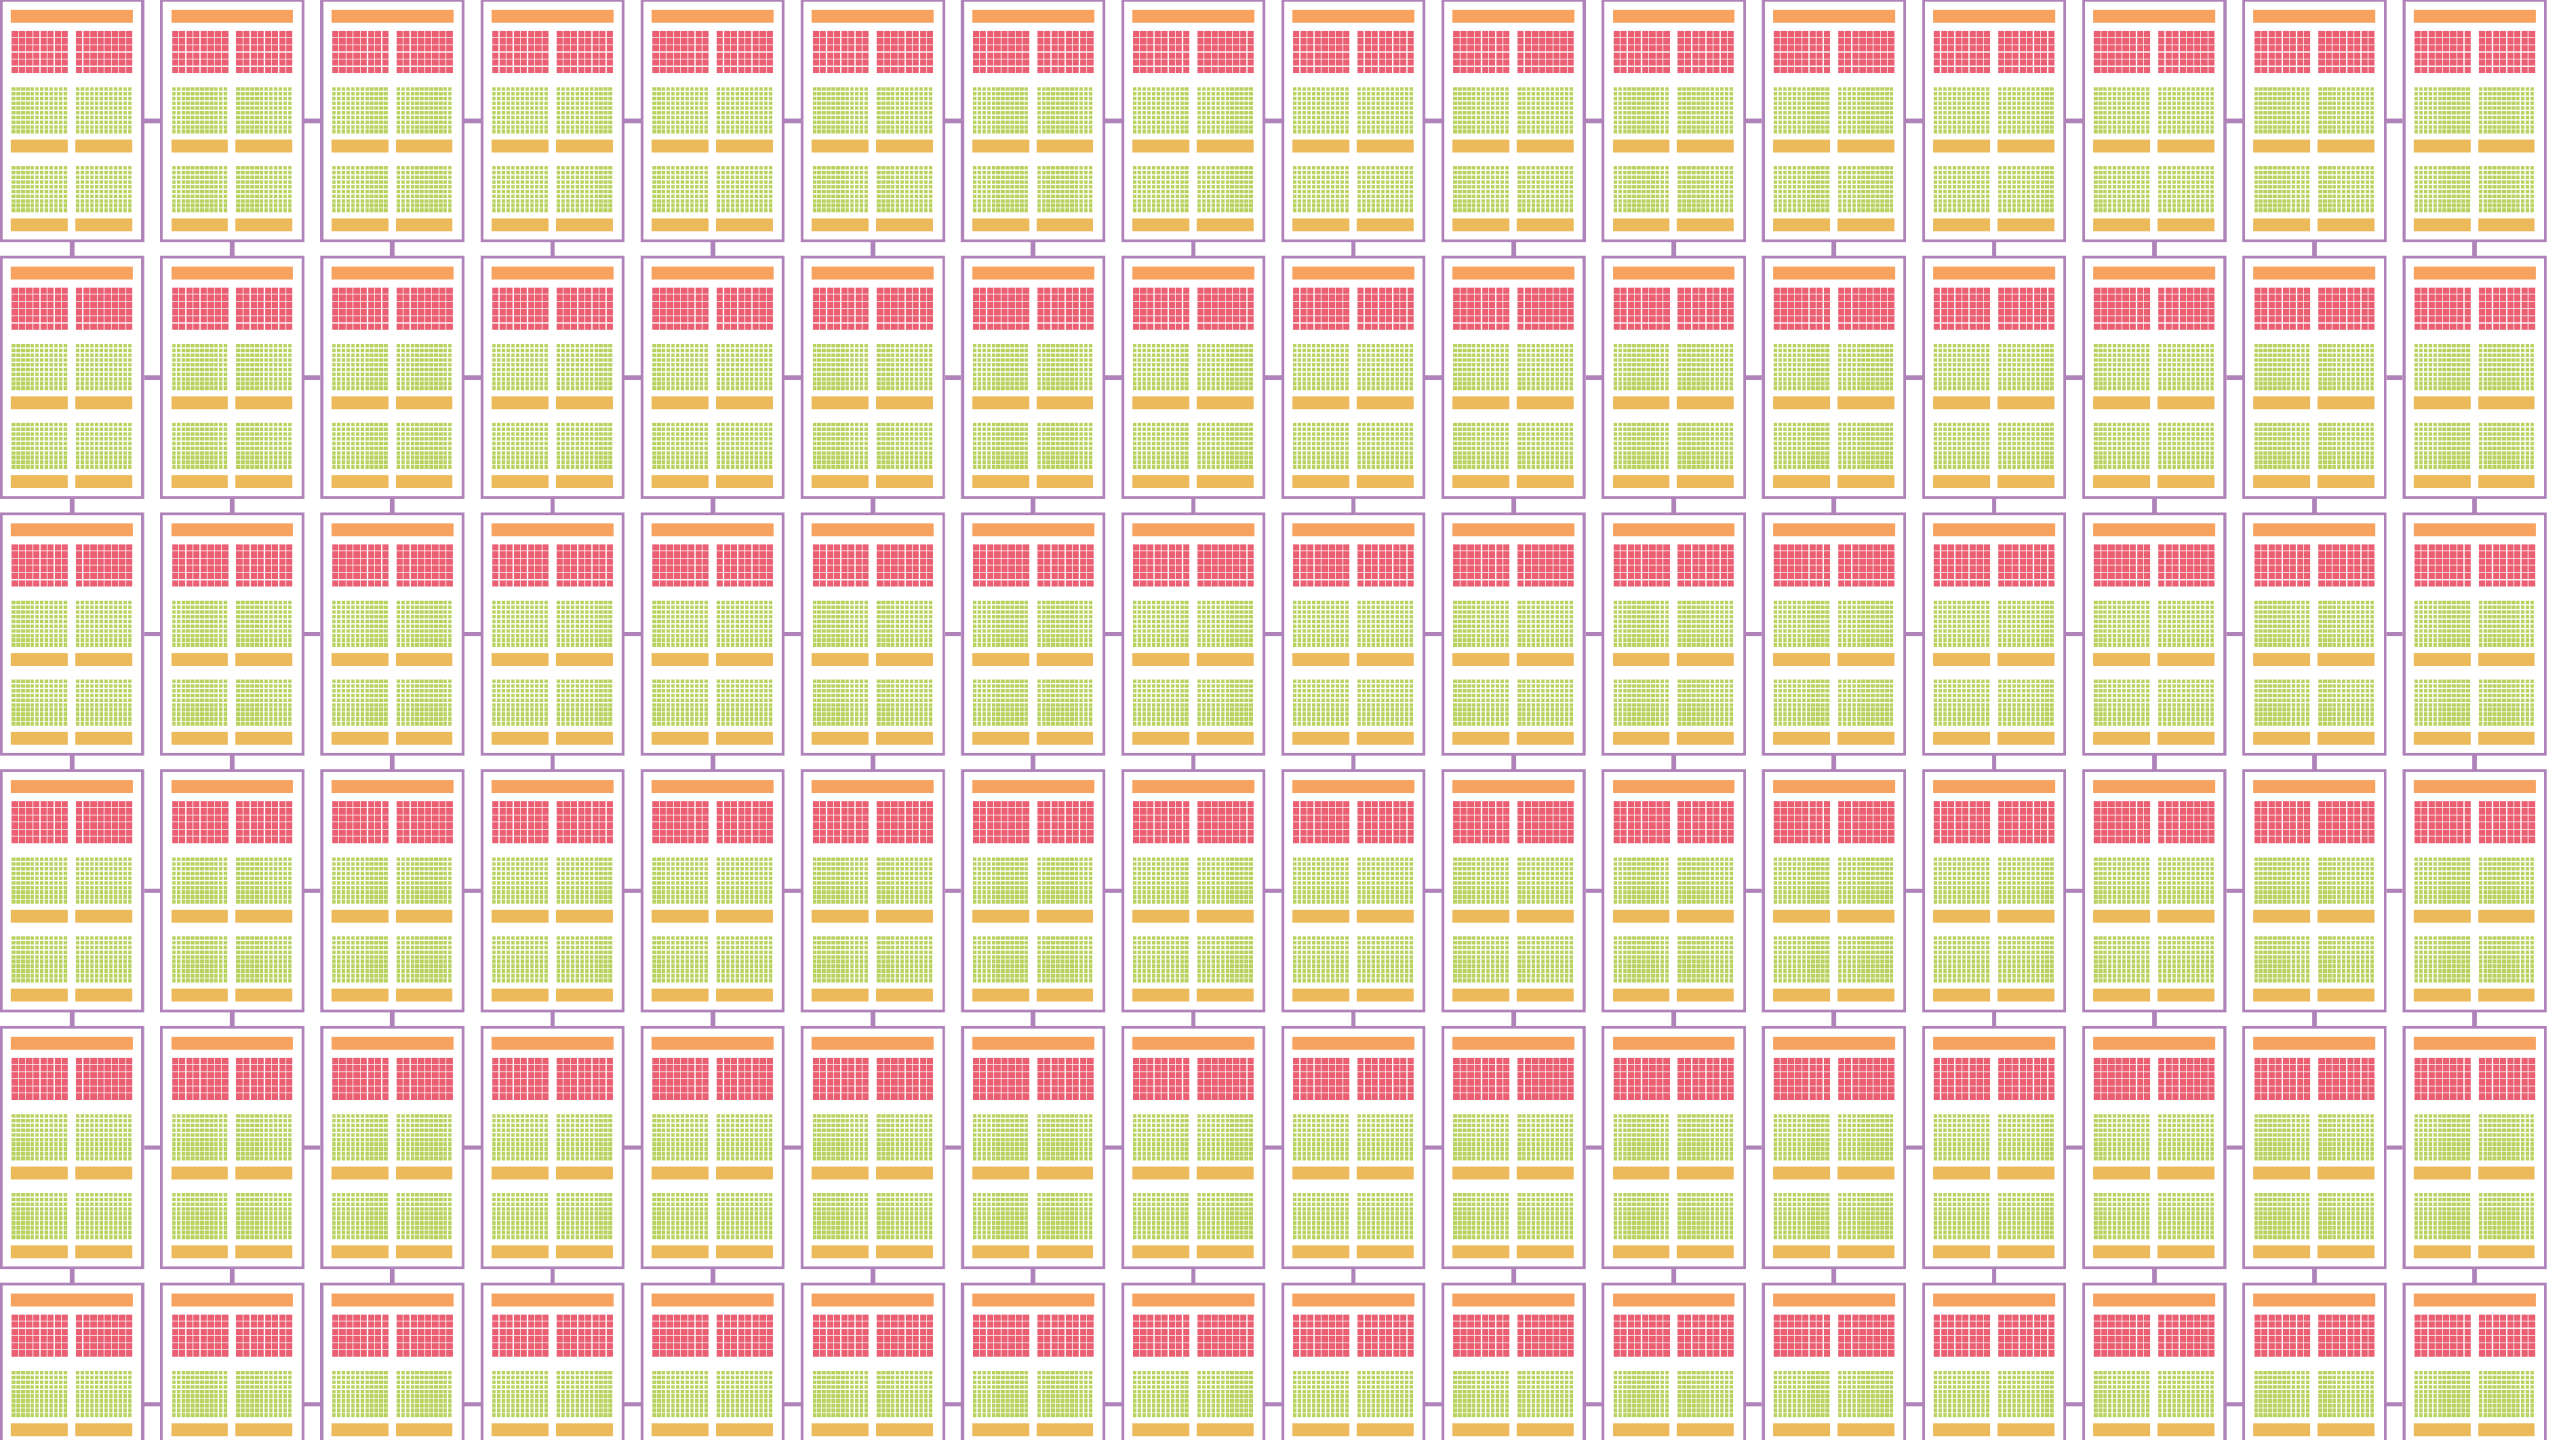

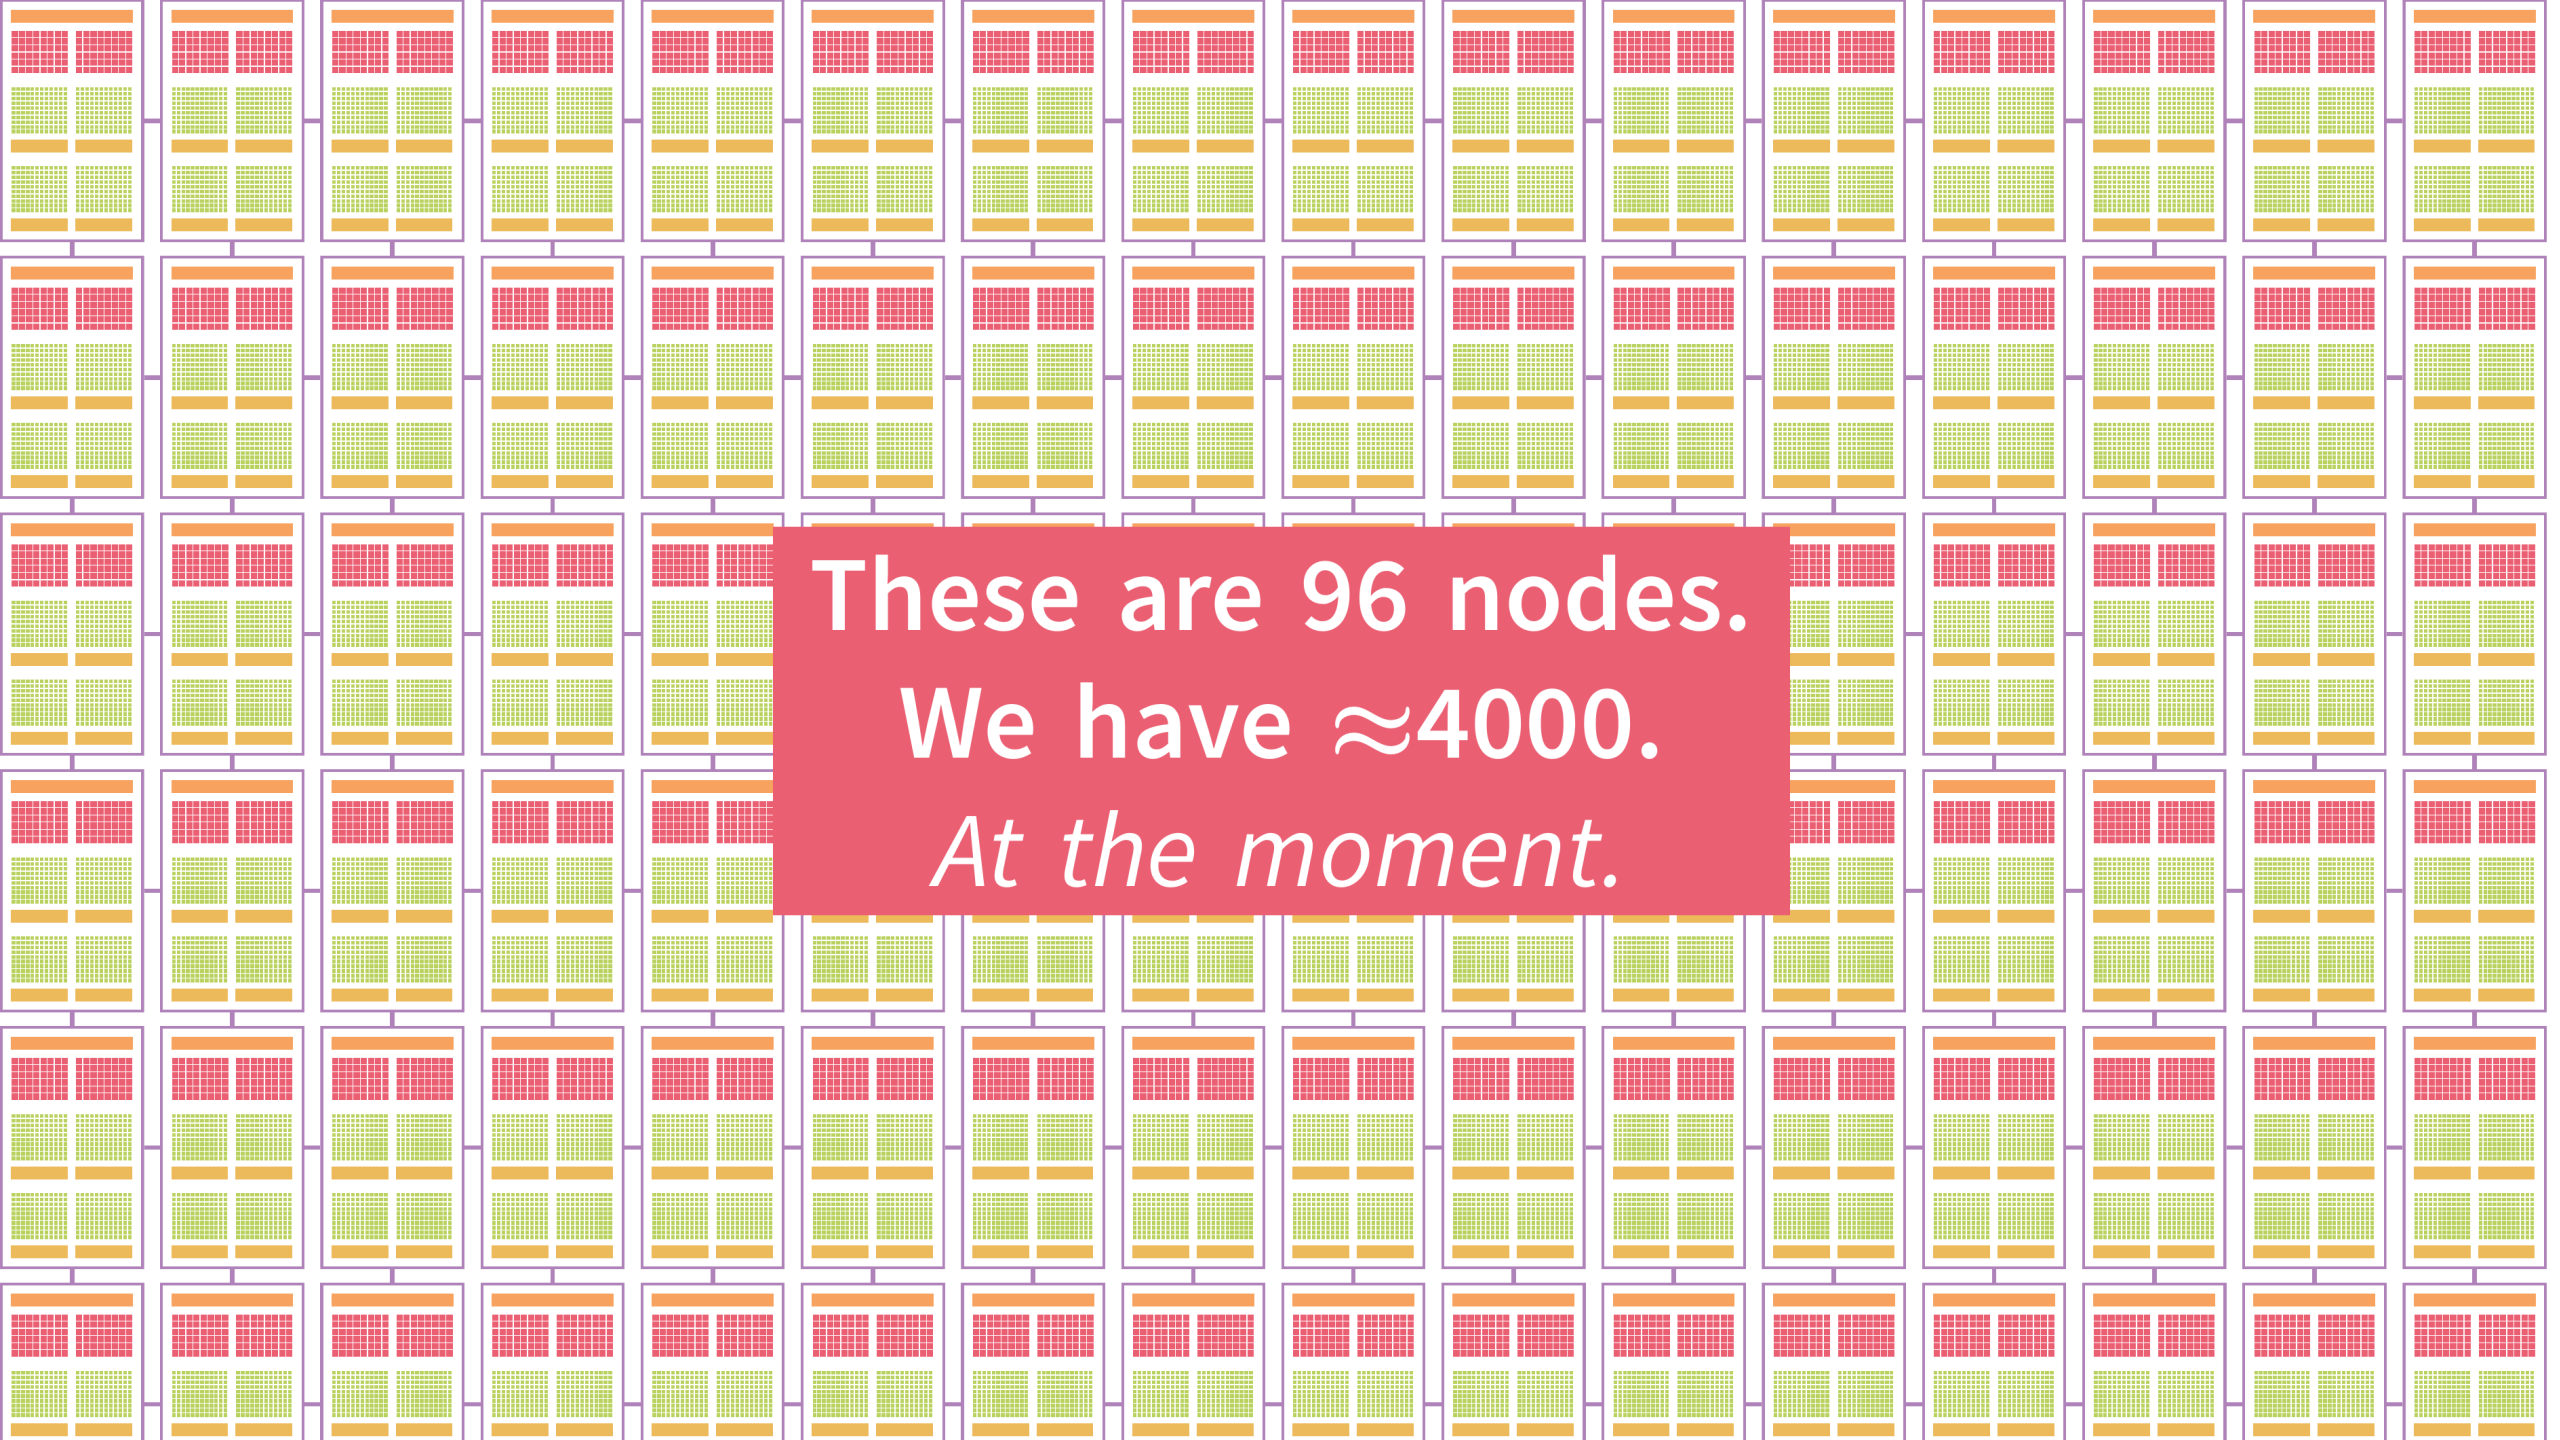
These are 96 nodes.
We have ≈4000.
At the moment.

Supercomputers in Pictures

- CDC 6600 supercomputer
- Around 1965
- First supercomputer
- 3 MFLOP/s
- See Wikipedia for more
- Picture by Control Data Corporation

Supercomputers in Pictures

HPC performance measured in **FLOP**/s.

- Floating-point (like 3.14) operations per second

Supercomputers in Pictures

HPC performance measured in **FLOP**/s.

- Floating-point (like 3.14) operations per second
- Example: Processor with 2 GHz; 10 cores; per core: 2 multiplications and 2 additions (*FMA*) per cycle

$$
\begin{aligned}
& 2 \times 10^9\ 1/\mathrm{s}\ 1/\mathrm{core} * 10\ \mathrm{core} * \\
& * (2+2)\ \text{floating-point operation} \\
=& 2 * 10^9 * 10 * 4\ \text{fl-op}/\mathrm{s} \\
=& 80 * 10^9\ \mathrm{FLOP}/\mathrm{s} \\
=& 80\ \mathrm{GFLOP}/\mathrm{s}
\end{aligned}
$$

Supercomputers in Pictures

- Cray-1 supercomputer
- Around 1978
- Very successful
- 160 MFLOP/s
- Probably pictured at NERSC

Supercomputers in Pictures

- Intel XP/S 140 supercomputer
- Around 1994
- 3680 Intel i860 RISC processors; large-scale parallel system
- 143 GFLOP/s
- Picture by top500.org

Supercomputers in Pictures

- JUGENE supercomputer
- 2008
- 294 912 PowerPC 450 cores; energy-efficient
- 800 TFLOP/s
- Picture by top500.org

Supercomputers in Pictures

- Summit supercomputer
- 2018
- 27 000 GPUs hosted by POWER9 CPUs; first #1 GPU supercomputer
- 200 PFLOP/s
- Picture by Oak Ridge National Lab

Supercomputers in Pictures

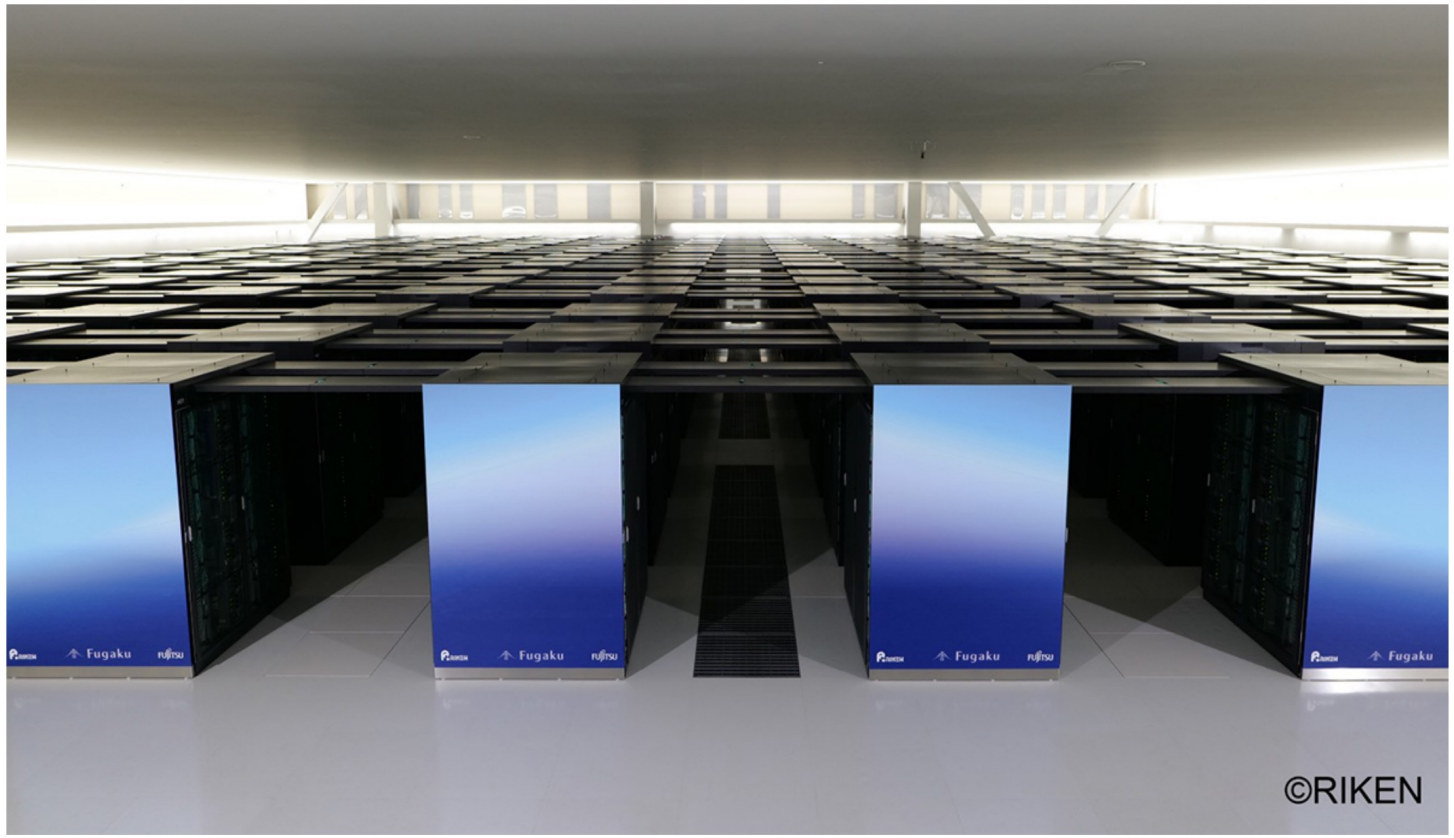

- Fugaku supercomputer
- 2020
- 7 630 848 Arm A64FX cores; #1 supercomputer
- 537 PFLOP/s
- Picture by RIKEN

JUWELS Cluster – Jülich's Scalable System

- 2500 nodes with Intel Xeon CPUs (2×24 cores)
- $46 + 10$ nodes with 4 NVIDIA Tesla V100 cards (16 GB memory)
- 10.4 (CPU) + 1.6 (GPU) PFLOP/s peak performance (Top500: #86)

JUWELS Booster – Scaling Higher!

- 936 nodes with AMD EPYC Rome CPUs (2×24 cores)
- Each with 4 NVIDIA A100 Ampere GPUs (each: *FP64TC:* 19.5, *FP64:* 9.7 TFLOP/s, 40 GB memory)
- InfiniBand DragonFly+ HDR-200 network; 4×200 Gbit/s per node

But, why HPC?

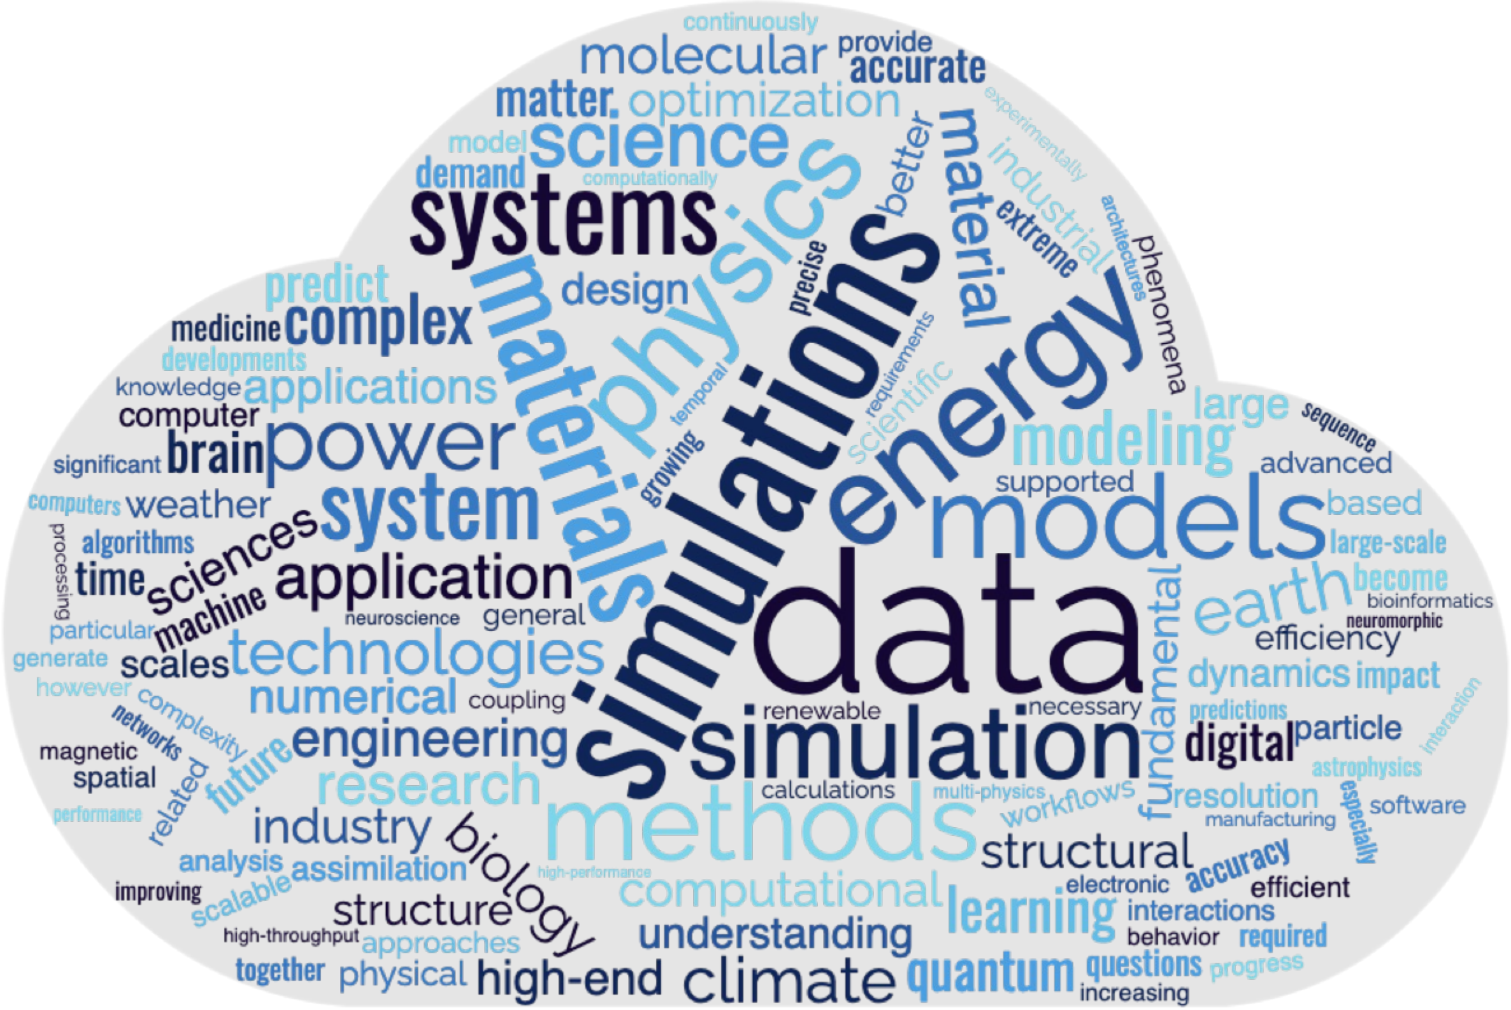
continuously
molecular
provide
accurate
matter
optimization
experimentally
model
science
demand
computationally
systems
physics
better
material
industrial
extreme
architectures
phenomena
predict
design
precise
materials
simulations
energy
medicine
complex
developments
knowledge
applications
requirements
scientific
temporal
computer
growing
large
sequence
brain
power
modeling
significant
advanced
supported
computers
weather
system
models
based
algorithms
sciences
large-scale
processing
time
application
earth
become
machine
data
bioinformatics
neuroscience
general
neuromorphic
particular
efficiency
generate
scales
technologies
fundamental
dynamics
impact
however
numerical
coupling
predictions
complexity
renewable
necessary
interaction
networks
engineering
simulation
digital
particle
magnetic
future
astrophysics
spatial
research
calculations
multi-physics
resolution
software
related
workflows
especially
performance
methods
manufacturing
industry
biology
structural
analysis
accuracy
assimilation
high-performance
electronic
efficient
improving
scalable
computational
learning
interactions
structure
understanding
behavior
required
high-throughput
approaches
questions
together
physical
high-end
climate
quantum
progress
increasing

How do you do HPC?

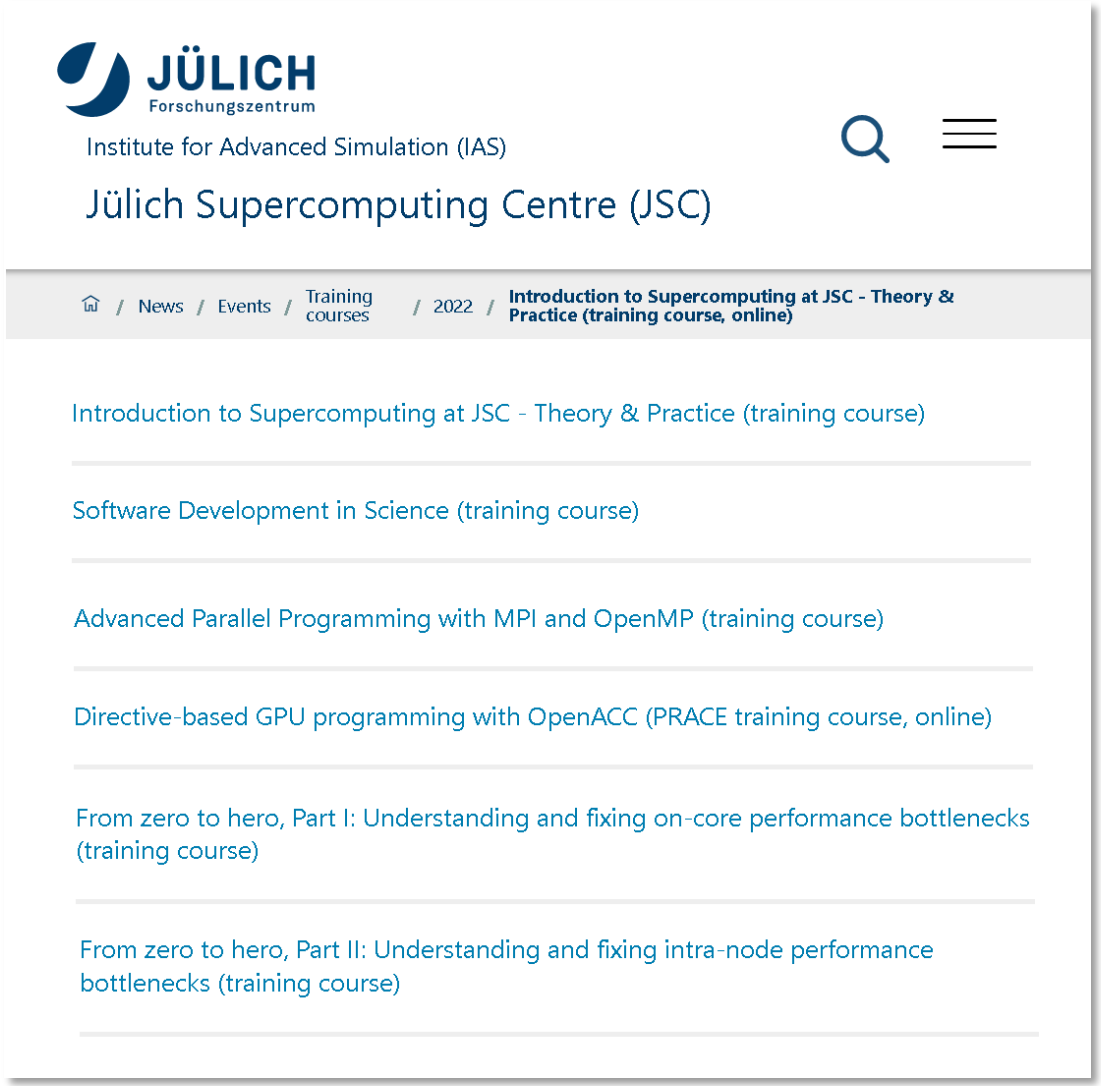

HPC application examples

(non-exhaustive list…)

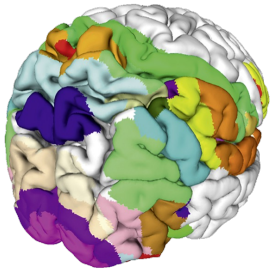

Explore the human brain (“Jülich Brain Atlas”)

- 💡 Massive data sets generated from actual brains

Video 1

https://fz-juelich.sciebo.de/s/WKMSp3FGb9y7Ob6

HPC application examples

(non-exhaustive list…)

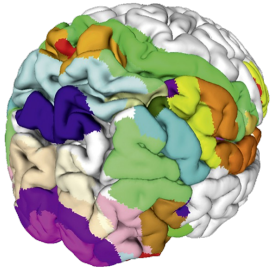

Explore the human brain (“Jülich Brain Atlas”)

- Massive data sets generated from actual brains

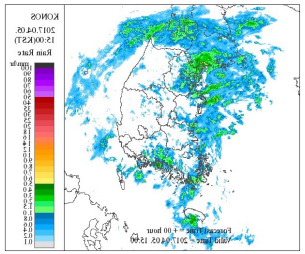

Earth system modelling, here: weather forecasting

- Deep learning on real-world measurements

Video 2

https://fz-juelich.sciebo.de/s/2JxRiTOCsUf1pXv

HPC application examples

(non-exhaustive list…)

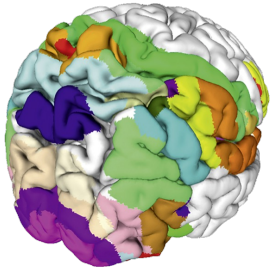

Explore the human brain (“Jülich Brain Atlas”)

- Massive data sets generated from actual brains

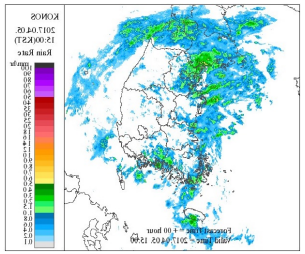

Earth system modelling, here: weather forecasting

- Deep learning on real-world measurements

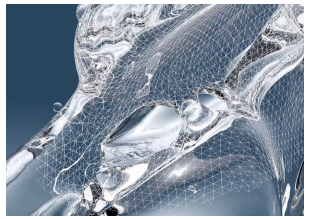

Liquid sheet breakup using direct numerical simulation

- Large-scale 3D simulations of complex physical phenomena

Video 3

- https://fz-juelich.sciebo.de/s/jYLKzcVKQd0UhGQ

But, why Exascale?

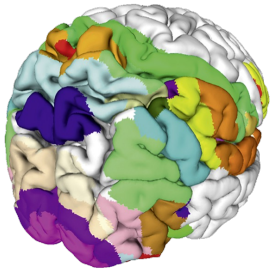

Explore the human brain (“Jülich Brain Atlas”)

- 💡 Massive data sets generated from actual brains
- ⛔ Currently only small parts of the brain in high-res

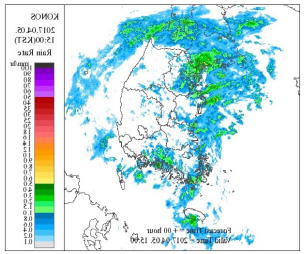

Earth system modelling, here: weather forecasting

- 💡 Deep learning on real-world measurements
- ⛔ What about all the other effects impacting out weather?

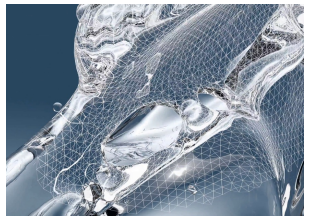

Liquid sheet breakup using direct numerical simulation

- 💡 Large-scale 3D simulations of complex physical phenomena
- ⛔ Shape optimization would need many of these simulations

State of the Art

Picture by *OLCF at ORNL* on Flickr

- Current fastest supercomputer: Frontier at Oak Ridge (USA) with 38 000 AMD MI250X GPUs – 1.102 EFLOP/s; also most energy-efficient!

State of the Art

Picture by *OLCF at ORNL* on Flickr

- Current fastest supercomputer: Frontier at Oak Ridge (USA) with 38 000 AMD MI250X GPUs – 1.102 EFLOP/s; also most energy-efficient!
- 2022: Aurora at Argonne with $>$ 60 000 Intel Ponte Vecchio GPUs – $>$ 2 EFLOP/s
- 2023: El Capitan at Lawrence Livermore with AMD MI300 GPUs – $>$ 2 EFLOP/s

State of the Art

Picture by *OLCF at ORNL* on Flickr

- Current fastest supercomputer: Frontier at Oak Ridge (USA) with 38 000 AMD MI250X GPUs – 1.102 EFLOP/s; also most energy-efficient!
- 2022: Aurora at Argonne with $>$ 60 000 Intel Ponte Vecchio GPUs – $>$ 2 EFLOP/s
- 2023: El Capitan at Lawrence Livermore with AMD MI300 GPUs – $>$ 2 EFLOP/s
- 2023: JUPITER at JSC – $>$ 1 EFLOP/s! GPUs, but details TBD

JUPITER announcement

Breaking news! (on 15.06.2022)

tagesschau

Startseite ▸ Wirtschaft ▸ Technologie ▸ Hochleistungs-Rechner: Supercomputer "Jupiter" kommt nach Jülich

Hochleistungs-Rechner

Supercomputer "Jupiter" komm

Stand: 15.06.2022 16:43 Uhr

Das Forschungszentrum Jülich wird Standort des ersten europäischen Exascale-Computers. "Jupiter" soll die Schallmauer von einer Trillion Rechenoperationen in der Sekunde durchbrechen.

SPIEGEL Netzwelt

»Jupiter«

Jülich bekommt Europas ersten Exascale-Supercomputer

Das Forschungszentrum Jülich bekommt für eine halbe Milliarde Euro einen neuen Vorzeigerechner. Er soll helfen, Fragen zum Klimawandel zu beantworten – mit mehr als einer Trillion Rechenoperationen pro Sekunde.

15.06.2022, 16.52 Uhr

SIGN IN

The Register

{ HPC }

Germany to host Europe's first exascale supercomputer

Jupiter added to HPC solar system

Dan Robinson

Thu 16 Jun 2022 // 07:33 UTC

3

Germany will be the host of the first publicly known European exascale supercomputer, along with four other EU sites getting smaller but still powerful systems, the European High Performance Computing Joint Undertaking (EuroHPC JU) announced this week.

Germany will be the home of Jupiter, the "Joint Undertaking Pioneer for Innovative and Transformative Exascale Research." It should be switched on next year in a specially designed building on the campus of the Forschungszentrum Jülich research centre and operated by the Jülich Supercomputing Centre (JSC), alongside the existing Juwels and Jureca supercomputers.

JÜLICH Forschungszentrum

JUPITER – A modular supercomputer

> 1 EB

Parallel High Capacity Data System

Parallel High Bandwidth Flash Module

High Capacity Backup/Archive System

> 1 EF

GPU

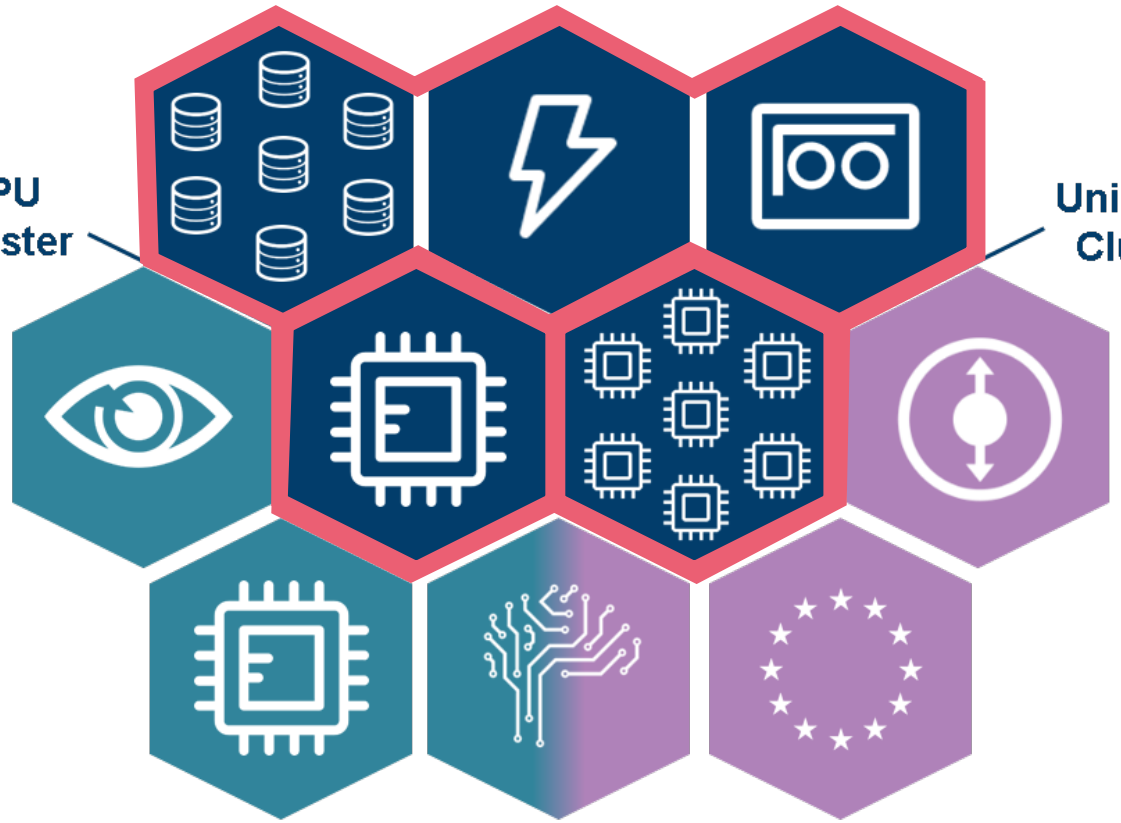

Target >20×

application performance compared to JUWELS Booster

7.5 PB/s

x86 or ARM

- **Basis Configuration**
- **Optional Modules**
- **Future Technology Modules**

JÜLICH Forschungszentrum

JUWELS – The first modular supercomputer

JUWELS Cluster	JUWELS Booster
Intel Xeon (Skylake) processor	AMD EPYC Rome 7402 processor
2,500 compute nodes	3,700 NVIDIA A100 GPUs
InfiniBand EDR network	InfiniBand HDR DragonFly+
~10 PFLOP/s peak (CPU-based)	**~70 PFLOP/s peak** (GPU-based)

JUWELS as blueprint for JUPITER

(modulo specific hardware vendors and generations)

JUWELS connectivity

Welcome to the jungle

Cluster node
Cluster GPU node
Cluster switch
Cluster gateway

Top level switch

Booster node
Booster switch
Booster gateway

JUST-IME node
JUST-IME switch

Service node

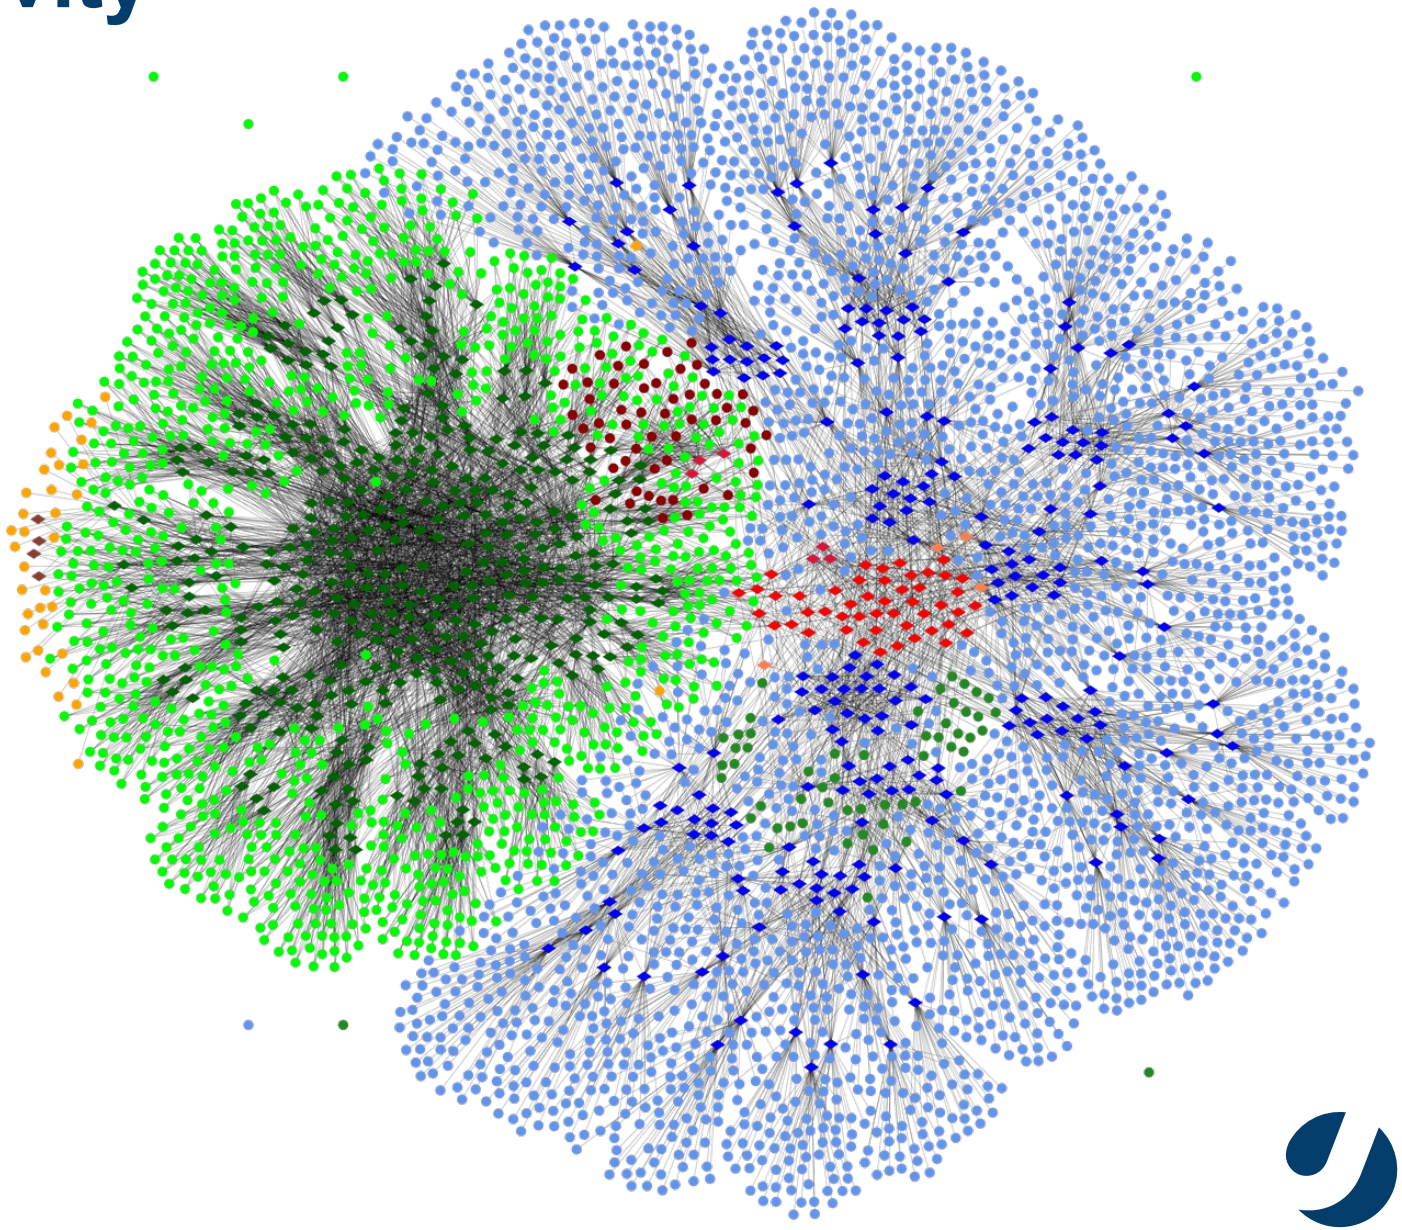

Three take-away messages for tonight

1️⃣ HPC can be the only way to achieve specific **scientific goals**.

2️⃣ Exascale Computing heralds a **new era** for HPC.

3️⃣ Building and using these machines does **not get easier**.

One final comment

> "I would rather have today's algorithms on yesterday's computers than vice versa."
>
> Philippe Toint

Thank You

JÜLICH
Forschungszentrum

JÜLICH
SUPERCOMPUTING
CENTRE

JUWELS vs. JUPITER (expectations)

	JUWELS	JUPITER
Cluster	**CPU:** Intel Xeon Platinum 8168 **GPU:** NVIDIA V100 **HPL perf.:** 10 PFlop/s	**CPU: ?** (AMD, Arm, Intel,…) **GPU: none** **Mem. Bandwidth:** 7.5 PB/s
Booster	**CPU:** AMD Epyc Rome **GPU:** NVIDIA A100 GPUs **Peak:** 70 PFlop/s	**CPU: ?** (AMD, Arm, Intel,…) **GPU: ?** (AMD, Intel, NVIDIA) **Peak: 1 EFlop/s**
Network topology	Fat tree + DragonFly+	**?** (likely some kind of DragonFly)
System access	GCS or PRACE proposals	GCS or EuroHPC proposals
User support	HLST, SDL, ATML, training courses, targeted early access program	**same**

User Support

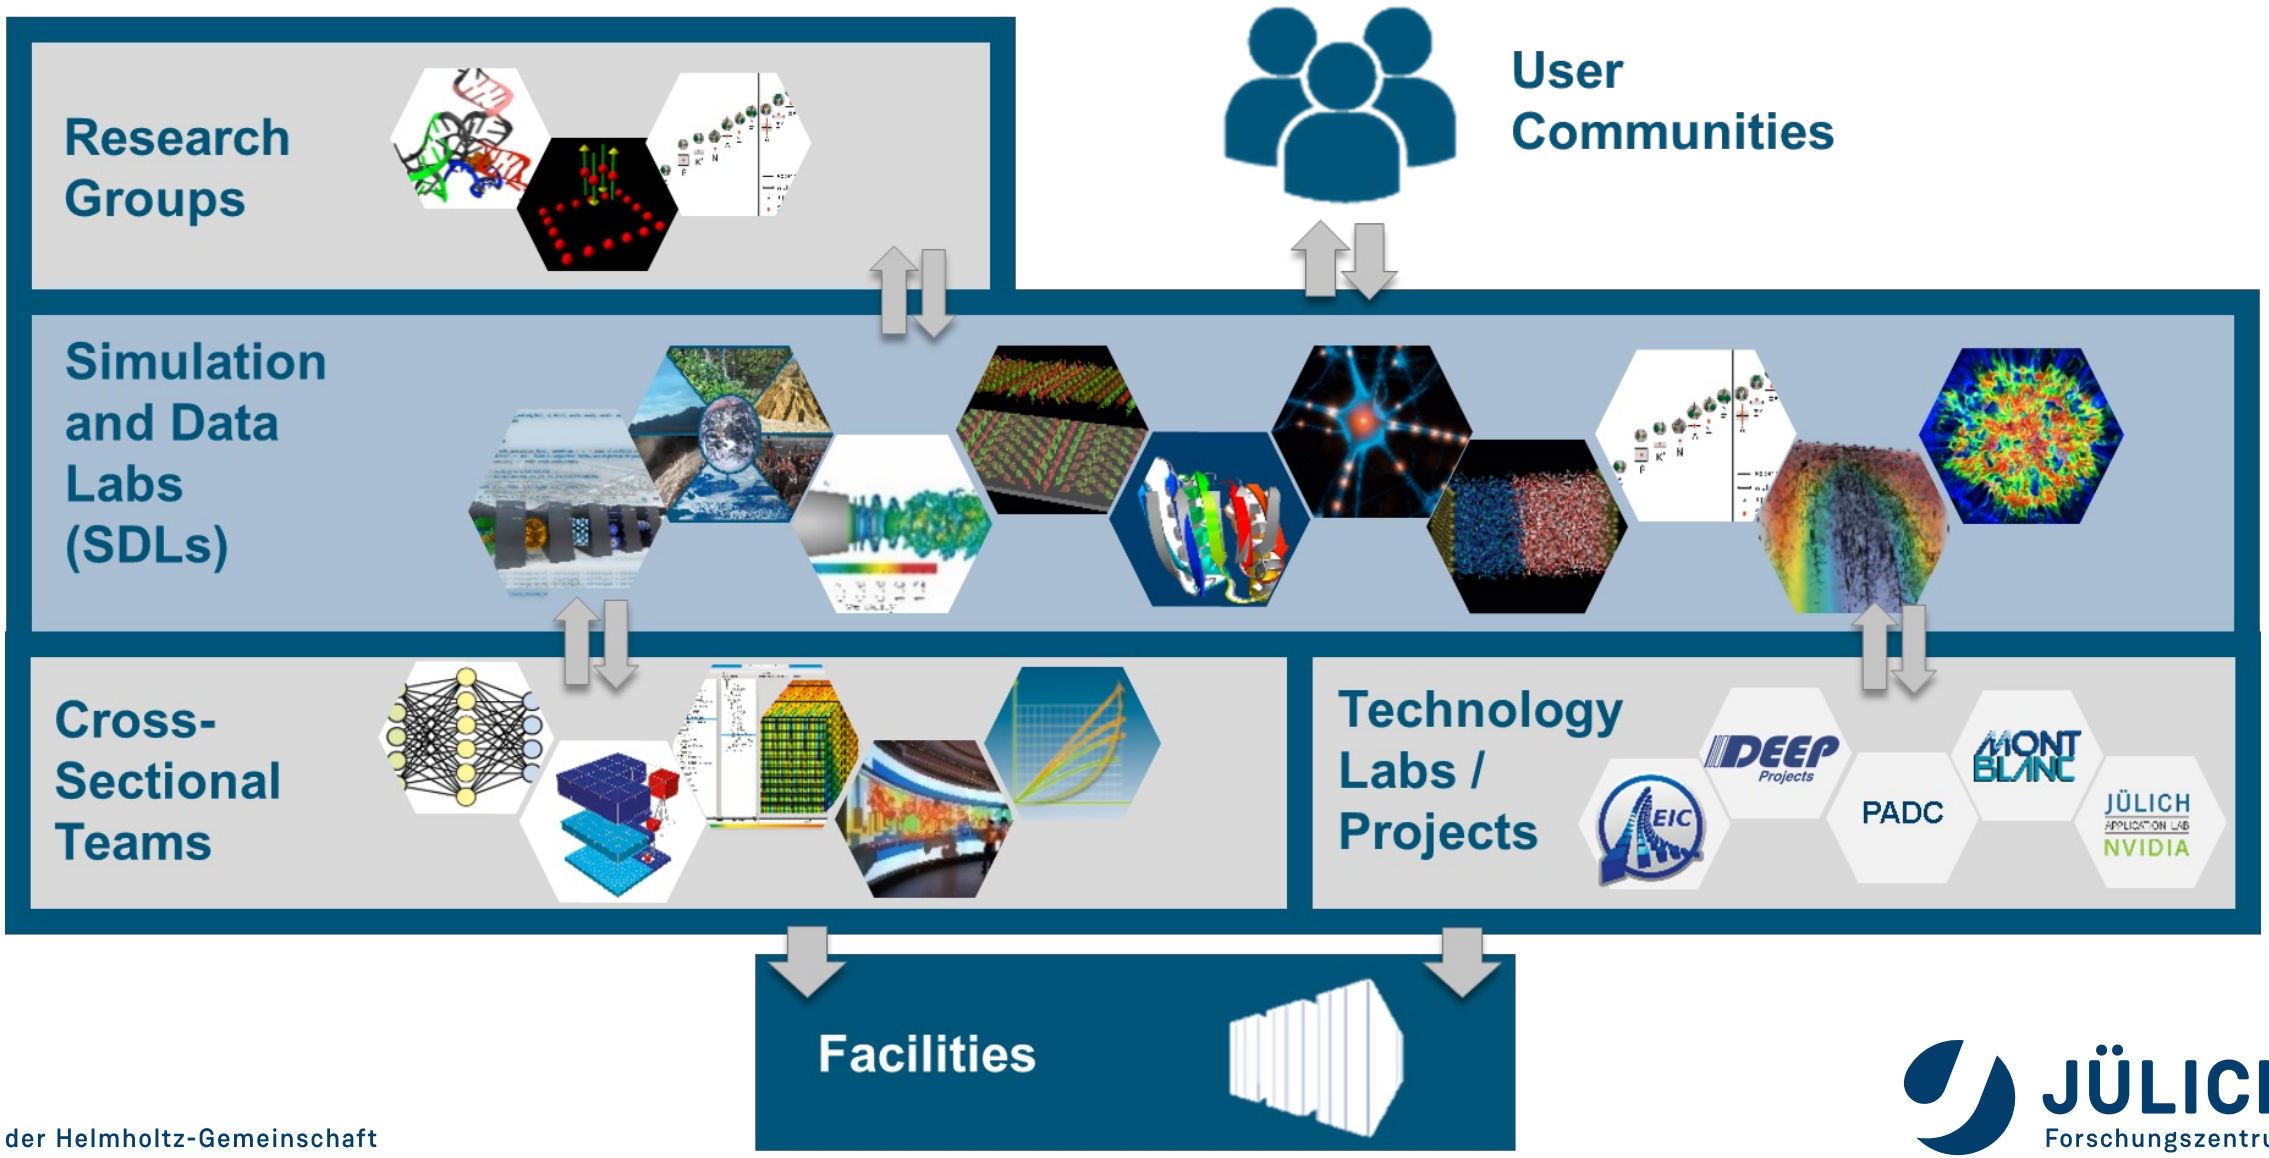

Special programmes, e.g.:
JUWELS Booster Early Access Program

Preparation Timeline

Additional events

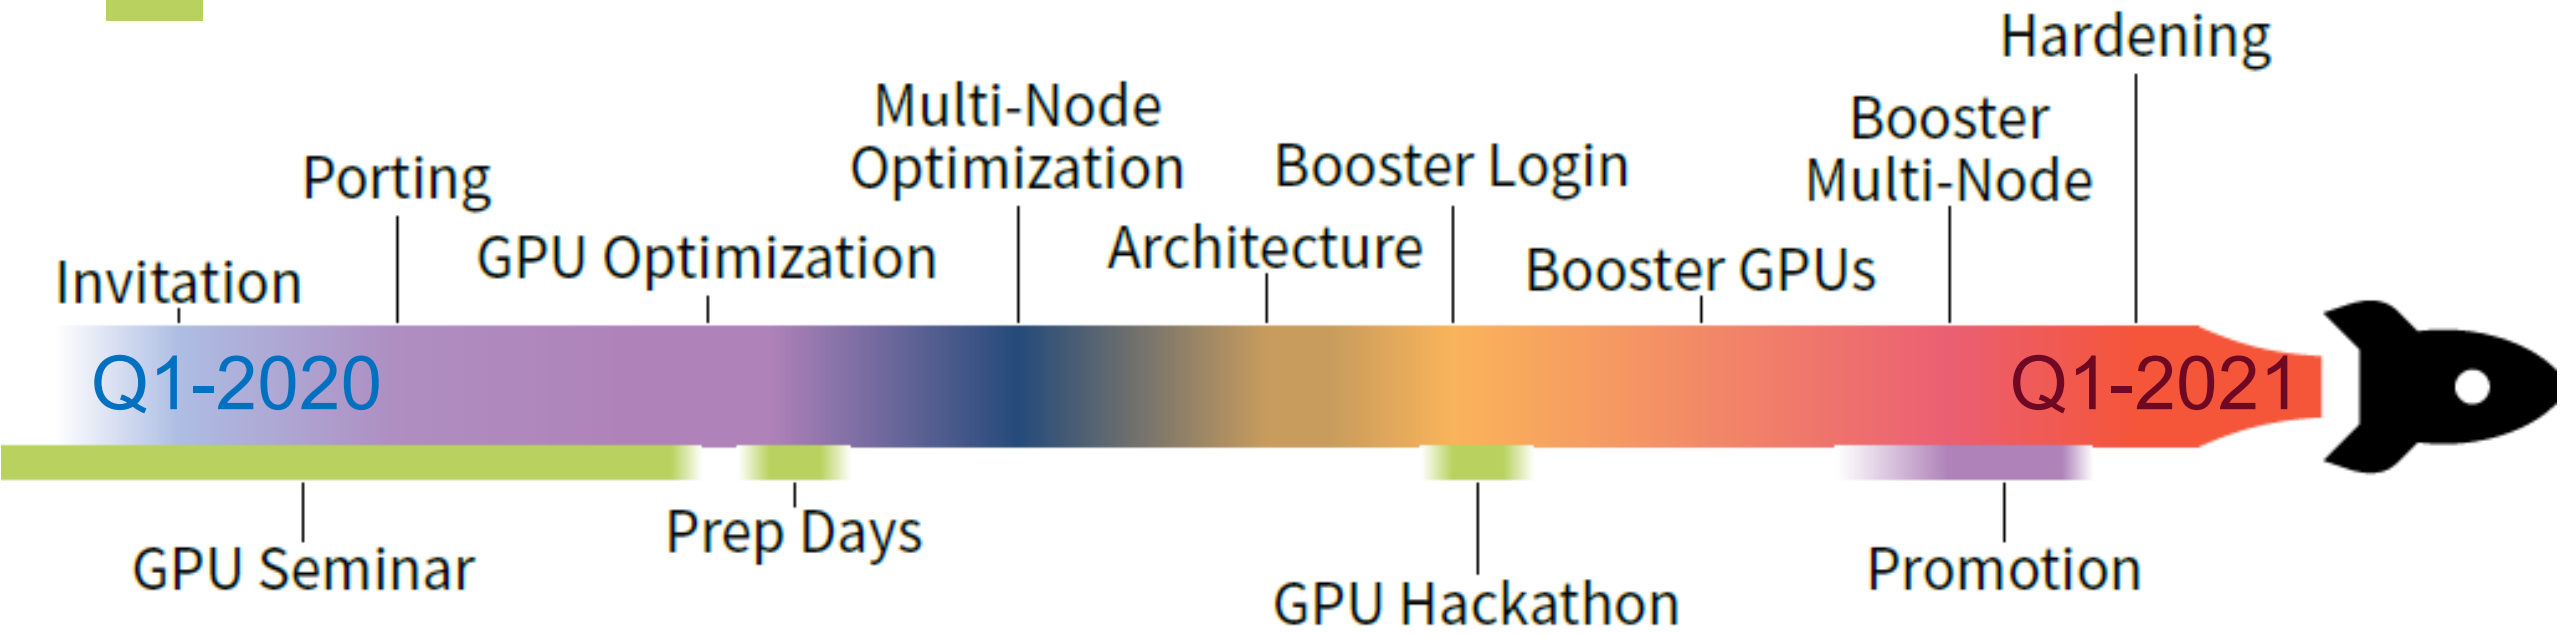

JÜLICH Forschungszentrum