Programmation et optimisation sur les architectures hétérogènes du mésocentre MUST

Pierre Aubert

MUST : Hardware

MUST

MUST : Hardware

Computing : ~7500 cores + 24 GPUs
Storage : 7 PB

MUST : Hardware

MUST

Computing : ~7500 cores + 24 GPUs
Storage : 7 PB

- 316 CPU computing units
- 158 blade servers
- 8000 computing tasks in parallel

MUST : Hardware

MUST

Computing : ~7500 cores + 24 GPUs
Storage : 7 PB

- 316 CPU computing units
- 158 blade servers
- 8000 computing tasks in parallel

- 18 servers (cloud, container, virtualization)
- 2x20 Gb/s external connectivity

MUST : Production

Average Usage : 92-96%

MUST : Production

Average Usage : 92-96%

WLCG : Wolrdwide LHC Computing Grid

WLCG
Worldwide LHC Computing Grid

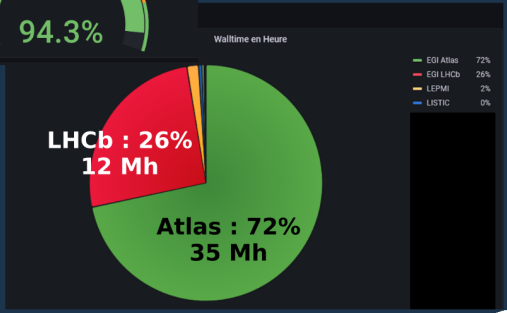

MUST : Data Challenge

MUST : Data Challenge

CERN
T0

MUST : Data Challenge

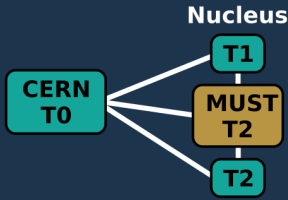

MUST : Data Challenge

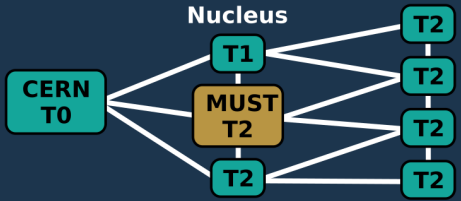

MUST : Data Challenge

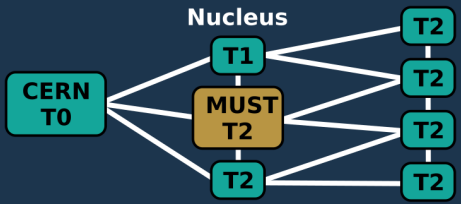

Data Challenge 2024 : 40 Gb/s
between Nucleus -> Nucleus

MUST : IDEFICS

R&D with local companies -> digital transition

MUST : IDEFICS

R&D with local companies -> digital transition

ADTP

Quality Check

MUST : IDEFICS

R&D with local companies -> digital transition

ADTP

Quality Check

Heliocity

Anomaly Detection

MUST : IDEFICS

R&D with local companies -> digital transition

ADTP

Quality Check

Heliocity

Anomaly Detection

Skiply

Predictive Maintenance (batteries)

MUST : IDEFICS

R&D with local companies -> digital transition

ADTP

Quality Check

Heliocity

Anomaly Detection

Skiply

Predictive Maintenance (batteries)

32 prospects meet

MUST : IDEFICS

R&D with local companies -> digital transition

ADTP

Quality Check

Heliocity

Anomaly Detection

Skiply

Predictive Maintenance (batteries)

32 prospects meet

Various Project :
- **Work-study intern**
- **Willingness to respond to MINALOGIC calls for tenders**
- **CIFRE Thesis**

MUST : Simulation

Galactic cosmic-ray sources

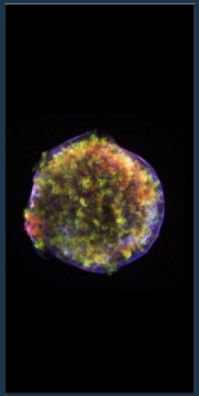

SNRs

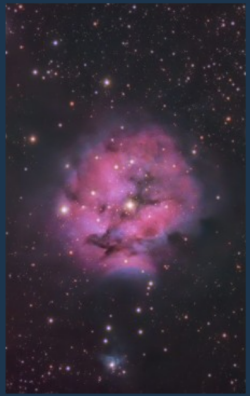

Star Clusters

Pulsar Wind Nebulae

Coliding Wind Binaries

Photostellar Jets Microquasars

and others !

MUST : Simulation

Yoann Génolini

MUST : Simulation

Yoann Génolini

Simulation in C++17 + nvc++

MUST : Simulation

Yoann Génolini

Simulation in C++17 + nvc++

Yoann's computer (with tbb, 8 threads 2.4 GHz)
49152 particules/140 secondes → 354 part/seconde → **gain = 1**

MUST : Simulation

Yoann Génolini

Simulation in C++17 + nvc++

GPU P6000
3145728 particules/44 min→ 1191 part/seconde → **gain = 3.4**

Yoann's computer (with tbb, 8 threads 2.4 GHz)
49152 particules/140 secondes → 354 part/seconde → **gain = 1**

MUST : Simulation

Yoann Génolini

Simulation in C++17 + nvc++

GPU V100
3145728 particules/14 min→ 5242 part/seconde → **gain = 15**

GPU P6000
3145728 particules/44 min→ 1191 part/seconde → **gain = 3.4**

Yoann's computer (with tbb, 8 threads 2.4 GHz)
49152 particules/140 secondes → 354 part/seconde → **gain = 1**

MUST : Simulation

Yoann Génolini

Simulation in C++17 + nvc++

GPU A100
40*3145728 particules/38 min → 55188 part/seconde → **gain = 155**

GPU V100
3145728 particules/14 min→ 5242 part/seconde → **gain = 15**

GPU P6000
3145728 particules/44 min→ 1191 part/seconde → **gain = 3.4**

Yoann's computer (with tbb, 8 threads 2.4 GHz)
49152 particules/140 secondes → 354 part/seconde → **gain = 1**

MUST : GammaLearn

MUST : GammaLearn

LAPP
Laboratoire d'Annecy de Physique des Particules

LAPP

LISTIC

learn

UNIVERSITÉ SAVOIE MONT BLANC

cta
cherenkov telescope array

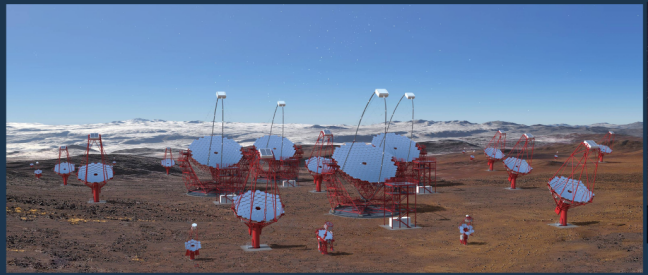

LST-1 prototype

MUST : GammaLearn

LAPP
Laboratoire d'Annecy de Physique des Particules

LISTIC

learn

UNIVERSITÉ SAVOIE MONT BLANC

cta
cherenkov telescope array

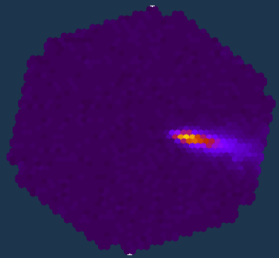

Particle :
- **Type**
- **Energy**
- **Direction**

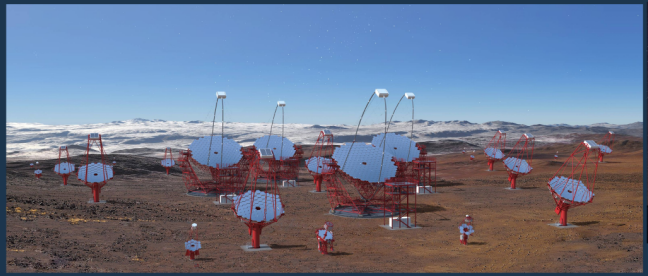

LST-1 prototype

MUST : Formation

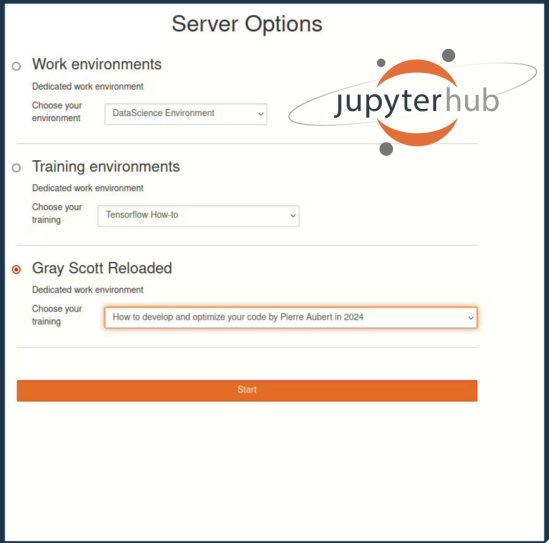

MUST : Formation

Kubernetes

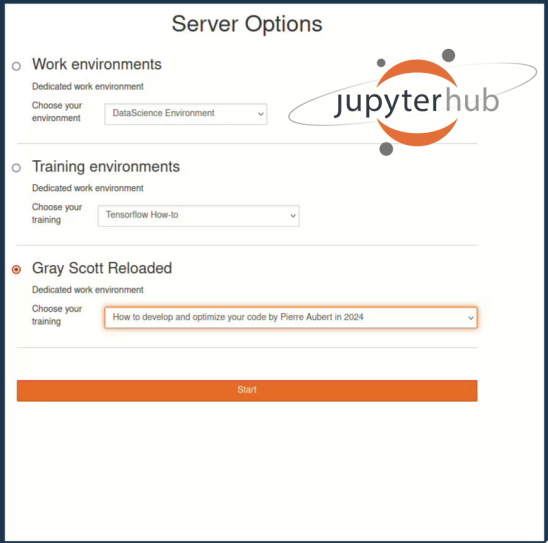

MUST : Formation

Kubernetes

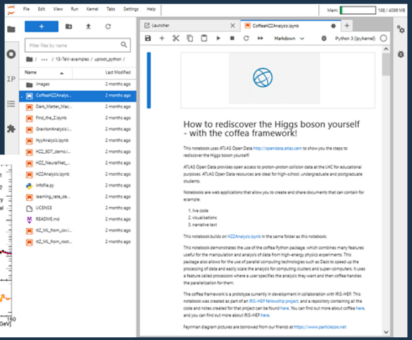

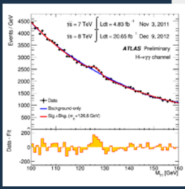

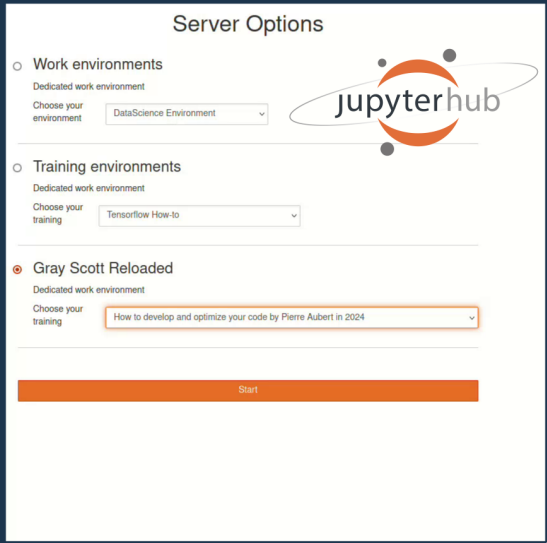

MUST : Formation

Kubernetes

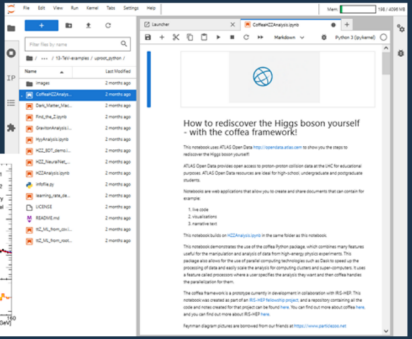

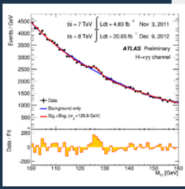

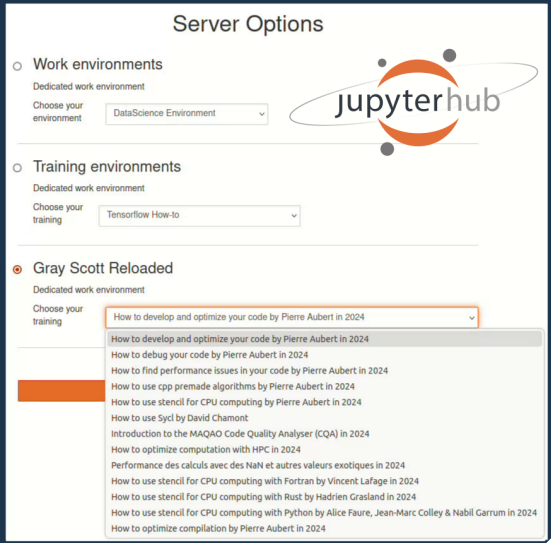

Reprises

Reprises

IN2P3 transverse project

~ **20** people

Reprises

IN2P3 transverse project

~ **20** people

Reprises

IN2P3 transverse project

~ **20** people

Performance tests

Valgrind MAQAO Perf

Precision

Cadna

Container

apptainer.org docker

C++ Fortran

Reprises

IN2P3 transverse project

~ **20** people

Performance tests

Valgrind MAQAO Perf

Precision

Cadna

Container

apptainer.org docker

Fortran

Xtensor

alpaka

HPX V1.2 DRAFT

Kokkos

NVIDIA CUDA

NVIDIA HPC DEV

OpenVINO

OpenCL

OpenACC More Science Less Programming

Reprises

IN2P3 transverse project

~ **20** people

Performance tests

Valgrind MAQAO Perf

Precision

Cadna

Container

apptainer.org docker

Fortran

Xtensor alpaka

HPX V1.2 DRAFT

Kokkos

Vulkan 1.2 FORGING AHEAD

NVIDIA CUDA

NVIDIA HPC DEV

OpenVINO

OpenCL OpenACC More Science Less Programming

CPU

GPU

FPGA

multi-core
many-core

TERATEC / EuroCC / EuroHPC

TERATEC / EuroCC / EuroHPC

- **Association**
- **Valorisation** / **Dissemination** of Knowledge :
  Schools / **Seminars** / **Interviews**
- **Bridge** between **Academy** and **Industry**

TERATEC / EuroCC / EuroHPC

France

- **Association**
- **Valorisation** / **Dissemination** of Knowledge :
  Schools / **Seminars** / **Interviews**
- **Bridge** between **Academy** and **Industry**

TERATEC / EuroCC / EuroHPC

EuroCC (Catalog of HPC/HTC/HPDA formations)

France

- **Association**
- **Valorisation** / **Dissemination** of Knowledge :
  Schools / **Seminars** / **Interviews**
- **Bridge** between **Academy** and **Industry**

TERATEC / EuroCC / EuroHPC

EuroCC (Catalog of HPC/HTC/HPDA formations)

France

Teratec

- **Association**
- **Valorisation** / **Dissemination** of Knowledge :
  Schools / **Seminars** / **Interviews**
- **Bridge** between **Academy** and **Industry**

+33 Countries

TERATEC / EuroCC / EuroHPC

EuroHPC

EuroCC (Catalog of HPC/HTC/HPDA formations)

France

Teratec

- **Association**
- **Valorisation** / **Dissemination** of Knowledge :
  Schools / **Seminars** / **Interviews**
- **Bridge** between **Academy** and **Industry**

+33 Countries

School computing example

Gray Scott reaction (a chemistry game of life)

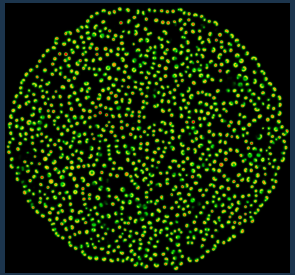

School computing example

Gray Scott reaction (a chemistry game of life)

$$U + 2V \longrightarrow 3V$$
$$V \longrightarrow P$$

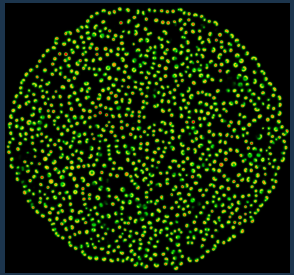

School computing example

Gray Scott reaction (a chemistry game of life)

Computing :

$$\begin{aligned} U + 2V &\longrightarrow 3V \\ V &\longrightarrow P \end{aligned}$$

$$\begin{aligned} \frac{\partial u}{\partial t} &= r_u \nabla^2 u - uv^2 + f_r \times (1 - u) \\ \frac{\partial v}{\partial t} &= r_v \nabla^2 v + uv^2 - (f_r + k_r) \times v \end{aligned}$$

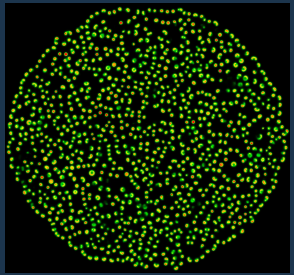

- u and v are concentration of product **U** and **V**
- r_u and r_v diffusion rate of **U** and **V**
- k_r (**Kill Rate**), conversion rate from **V** to **P**
- f_r (**Feed Rate**), speed of process which feed **U** and kills **V** and **P**
- $\nabla^2 u$ and $\nabla^2 v$ are différence of space concentration between current cell and its neighbours

School computing example

Gray Scott reaction (a chemistry game of life)

Computing :

$$\begin{aligned} U + 2V &\longrightarrow 3V \\ V &\longrightarrow P \end{aligned}$$

$$\begin{aligned} \frac{\partial u}{\partial t} &= r_u \nabla^2 u - uv^2 + f_r \times (1 - u) \\ \frac{\partial v}{\partial t} &= r_v \nabla^2 v + uv^2 - (f_r + k_r) \times v \end{aligned}$$

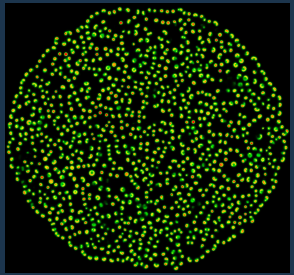

- u and v are concentration of product **U** and **V**
- r_u and r_v diffusion rate of **U** and **V**
- k_r (**Kill Rate**), conversion rate from **V** to **P**
- f_r (**Feed Rate**), speed of process which feed **U** and kills **V** and **P**
- $\nabla^2 u$ and $\nabla^2 v$ are différence of space concentration between current cell and its neighbours

- Easy to understand
- Not so easy for the compiler
- Possibility of high speed up

Languages / Technos

Languages / Technos

Languages / Technos

Fortran

Languages / Technos

Fortran

Languages / Technos

Fortran

Languages / Technos

Fortran

Languages / Technos

Fortran

Languages / Technos

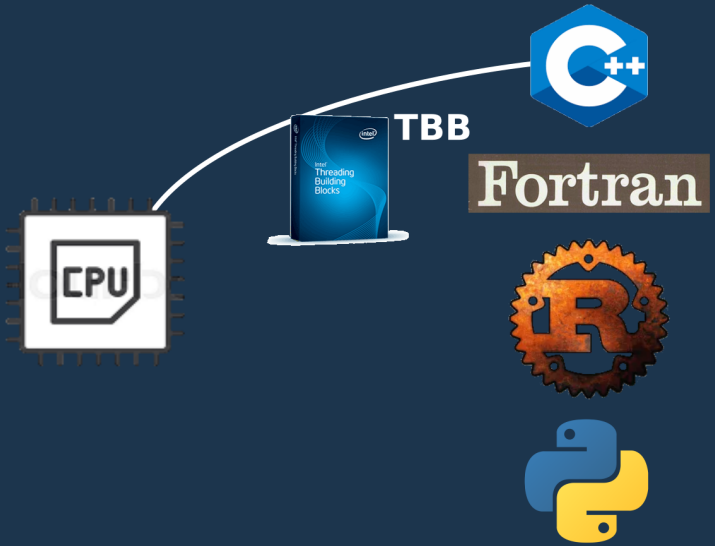

Languages / Technos

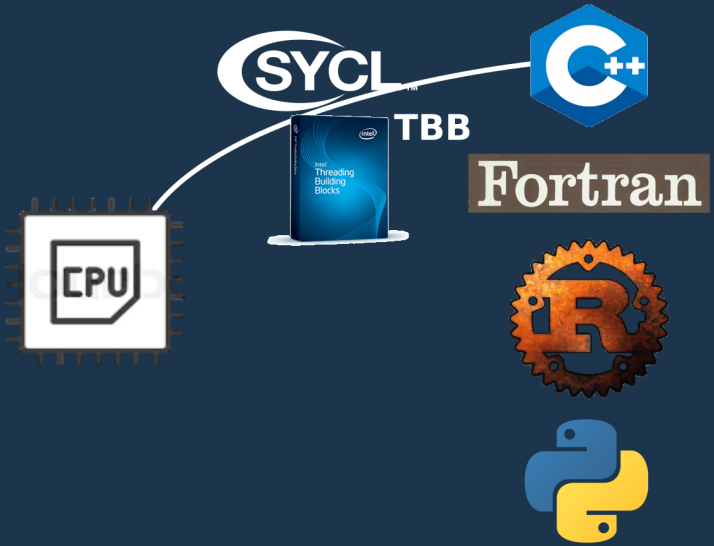

Languages / Technos

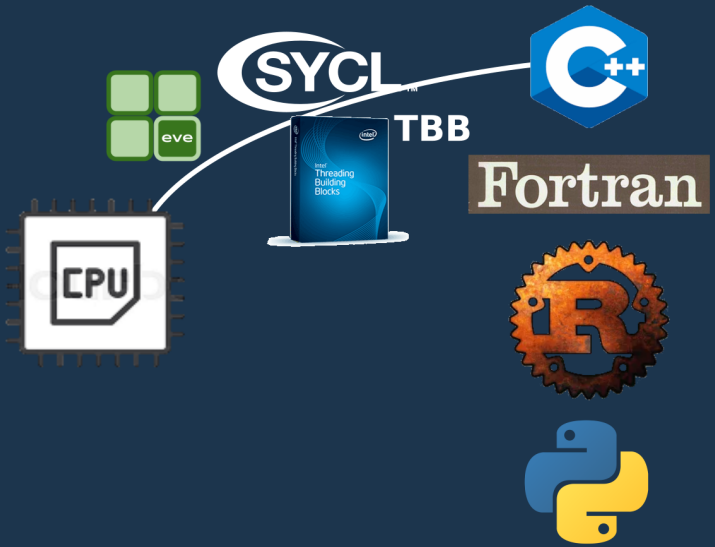

Languages / Technos

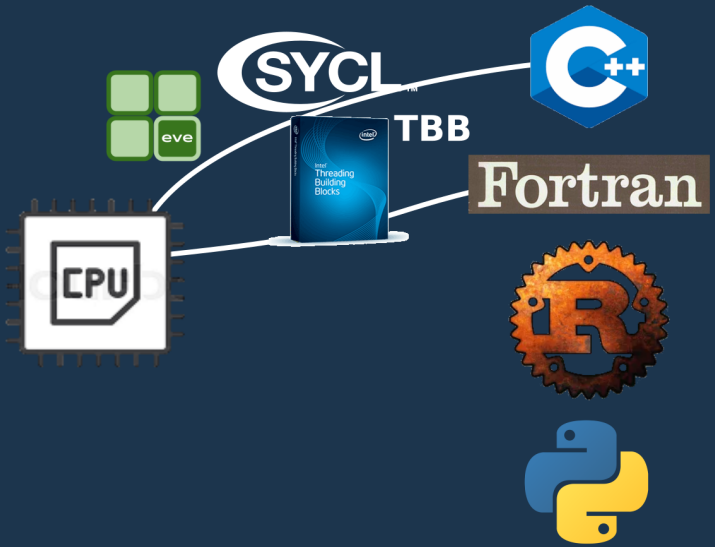

Languages / Technos

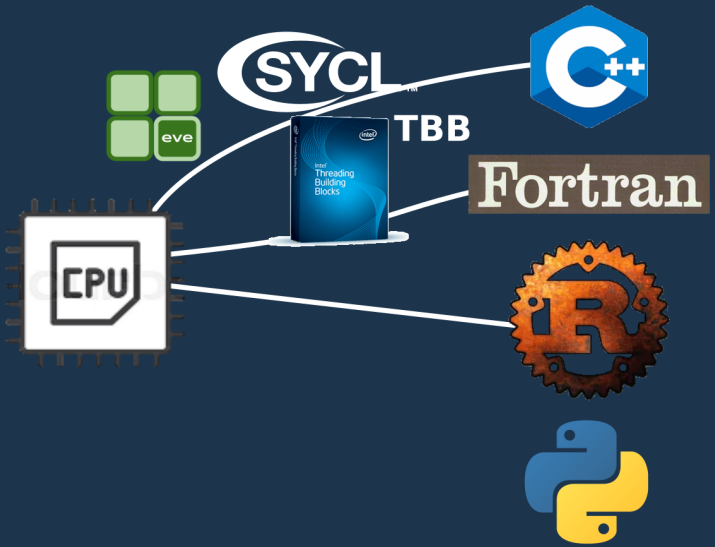

Languages / Technos

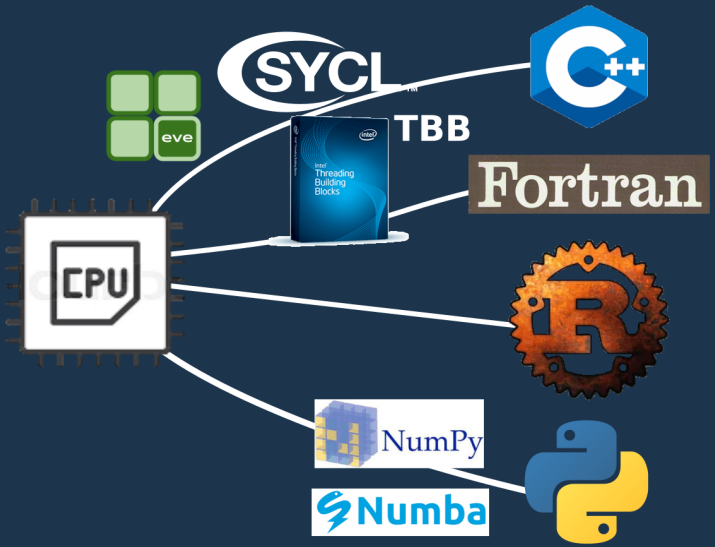

Languages / Technos

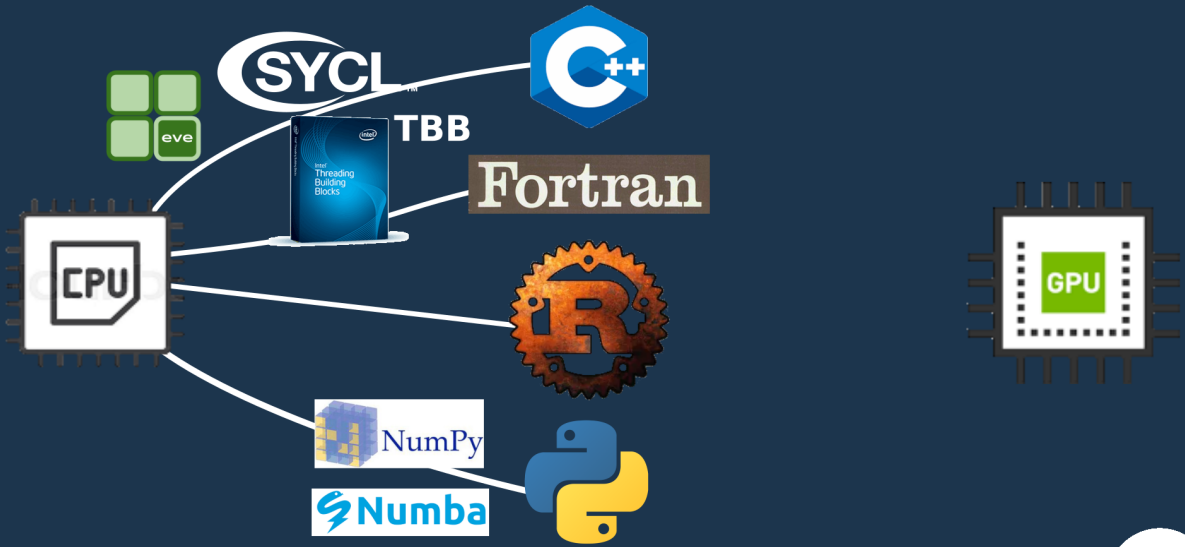

Languages / Technos

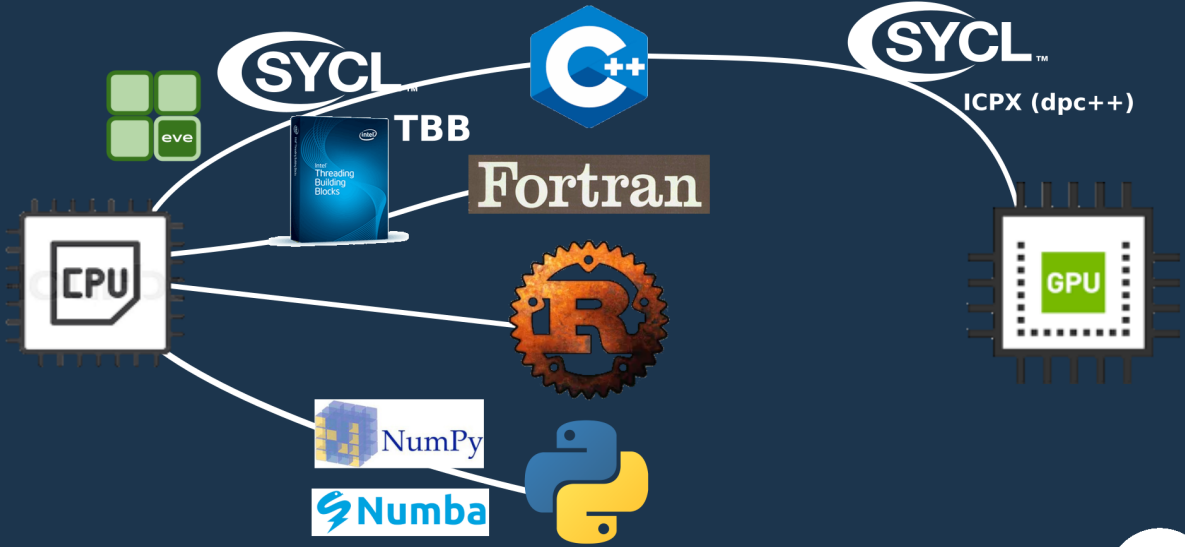

Languages / Technos

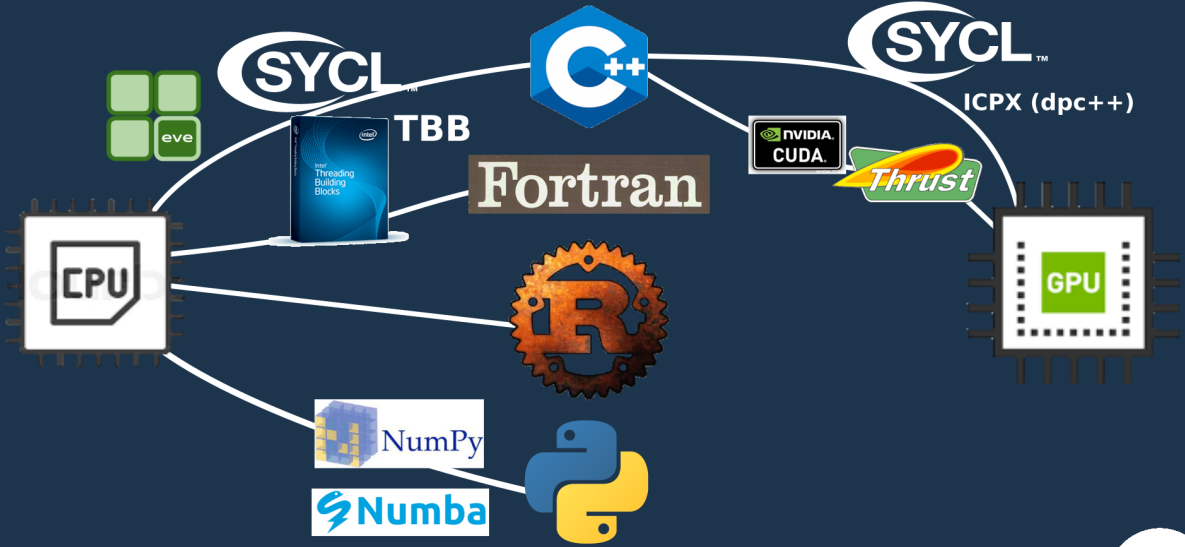

Languages / Technos

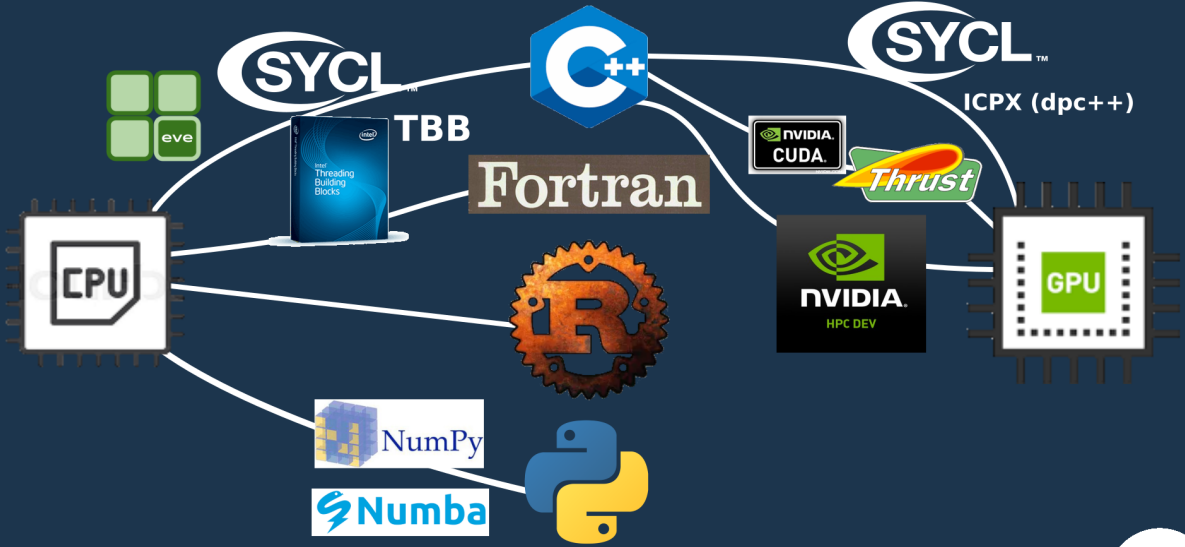

Languages / Technos

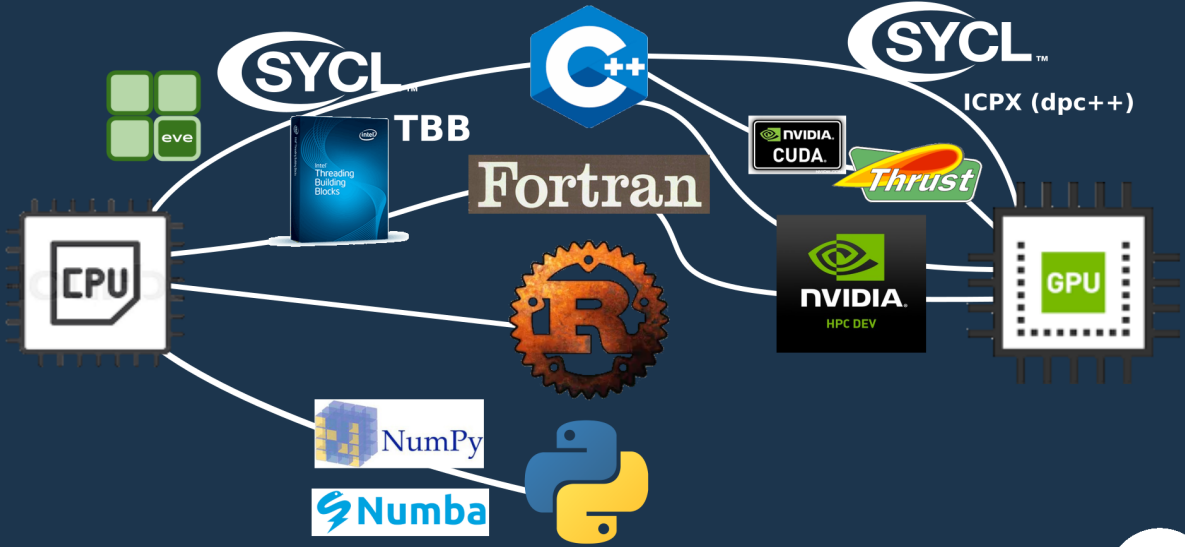

Languages / Technos

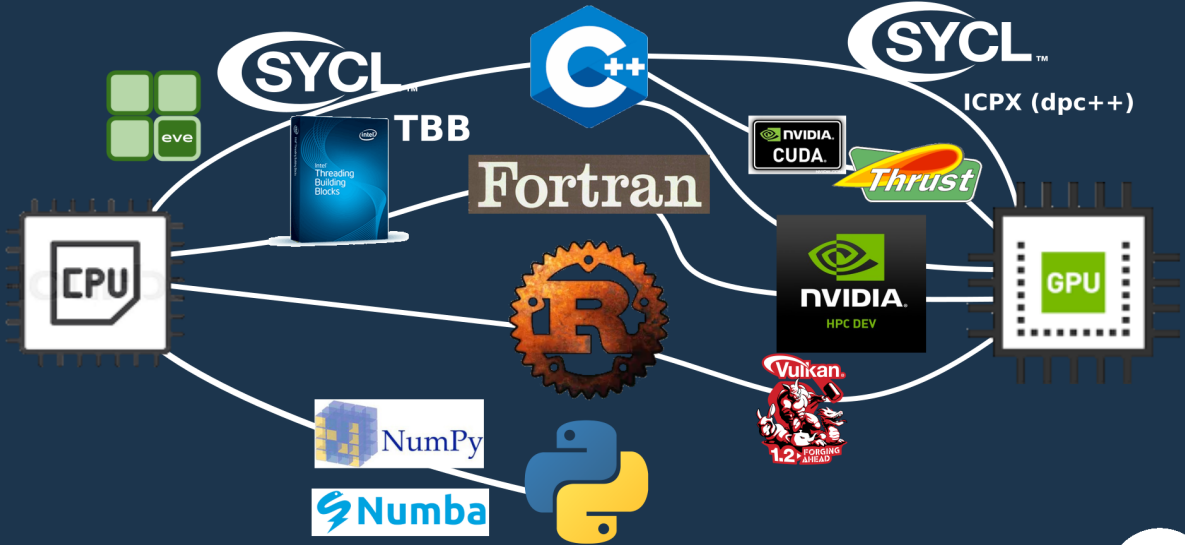

Languages / Technos

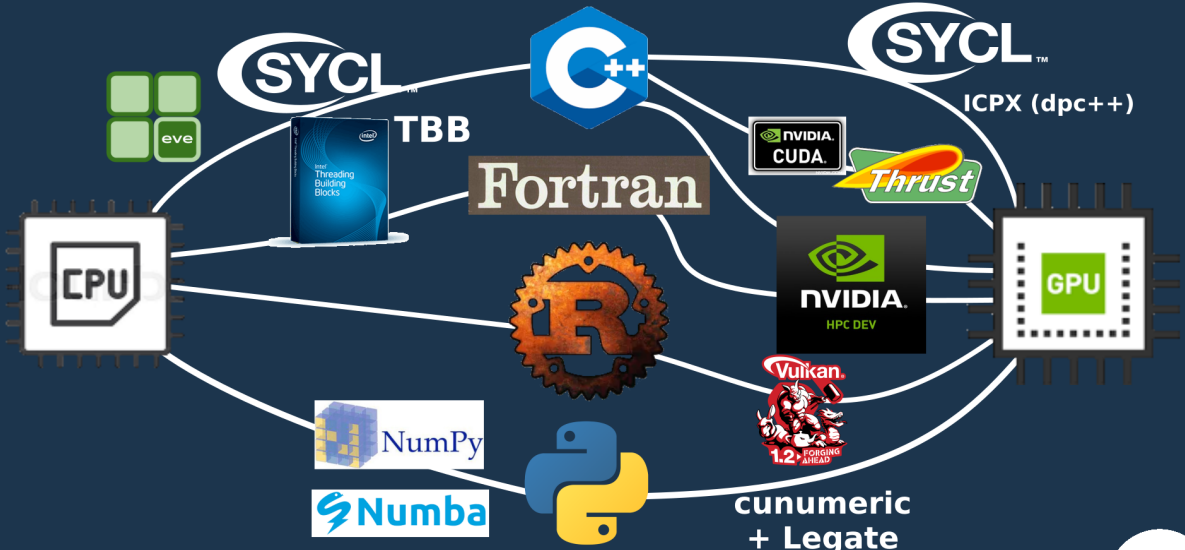

Languages / Technos

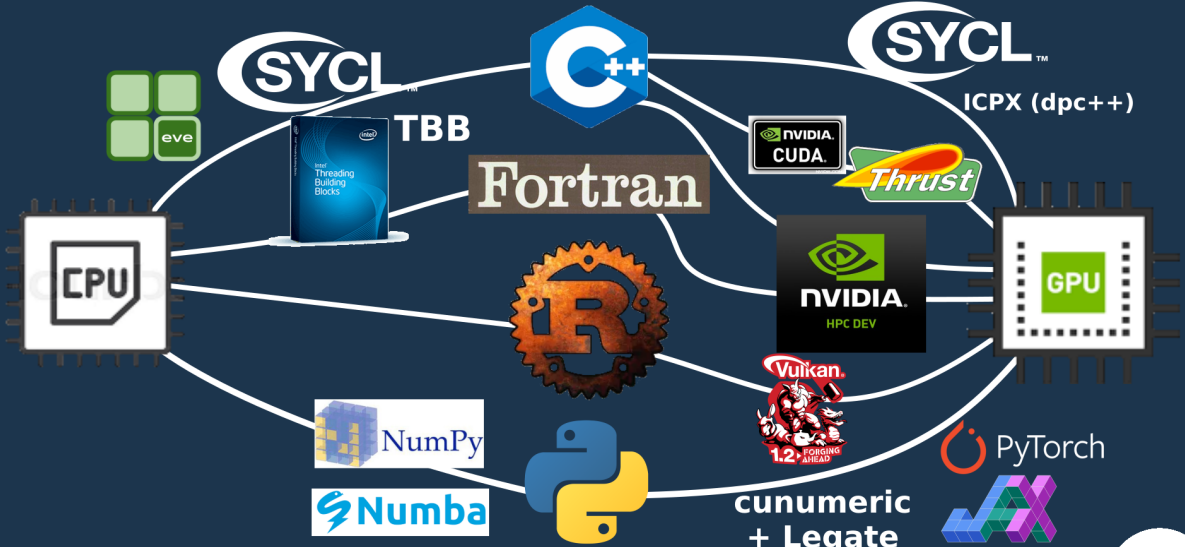

Tools

Tools

Performance tests

Perf

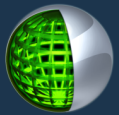

NSight

Tools

Performance tests

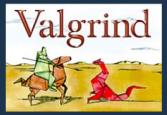

Perf

MAQAO

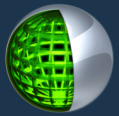

NSight

Memory Profiler

MATL

NUMA Prof

Tools

Performance tests

Perf

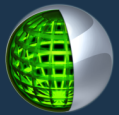

NSight

Memory Profiler

MATL

NUMA Prof

Compilers

Tools

Performance tests

Perf

MAQAO

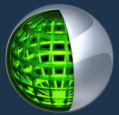

NSight

Compilers

LLVM : CLang / CLang ++ / FLang

Memory Profiler

MATL

NUMA Prof

Tools

Performance tests

Perf

MAQAO

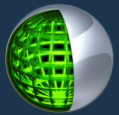

NSight

Compilers

LLVM : CLang / CLang ++ / FLang

GNU : GCC / G++ / GFortran

Memory Profiler

MATL

NUMA Prof

Tools

Performance tests

Perf

MAQAO

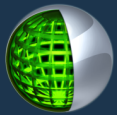

NSight

Memory Profiler

MATL

NUMA Prof

Compilers

LLVM : CLang / CLang ++ / FLang

GNU : GCC / G++ / GFortran

nvc++ / nvfortran

intel

icpx / dpc++

Tools

Performance tests

Perf

MAQAO

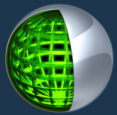

NSight

Compilers

LLVM : CLang / CLang ++ / FLang

GNU : GCC / G++ / GFortran

nvc++ / nvfortran

intel

icpx / dpc++

Precision

Memory Profiler

MATL

NUMA Prof

Tools

Performance tests

Perf

MAQAO

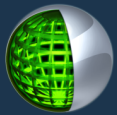

NSight

Compilers

LLVM : CLang / CLang ++ / FLang

GNU : GCC / G++ / GFortran

nvc++ / nvfortran

intel

icpx / dpc++

Precision

Memory Profiler

MATL

NUMA Prof

Tools

Performance tests

Perf

MAQAO

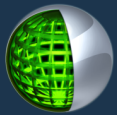

NSight

Compilers

LLVM : CLang / CLang ++ / FLang

GNU : GCC / G++ / GFortran

nvc++ / nvfortran

intel

icpx / dpc++

Precision

Memory Profiler

MATL

NUMA Prof

Tools

Performance tests

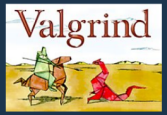

Perf

MAQAO

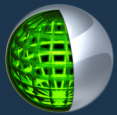

NSight

Compilers

LLVM : CLang / CLang ++ / FLang

GNU : GCC / G++ / GFortran

nvc++ / nvfortran

intel

icpx / dpc++

Precision

Memory Profiler

MATL

NUMA Prof

git

Gitlab

Multi-site diffusion

Multi-site diffusion

Multi-site diffusion

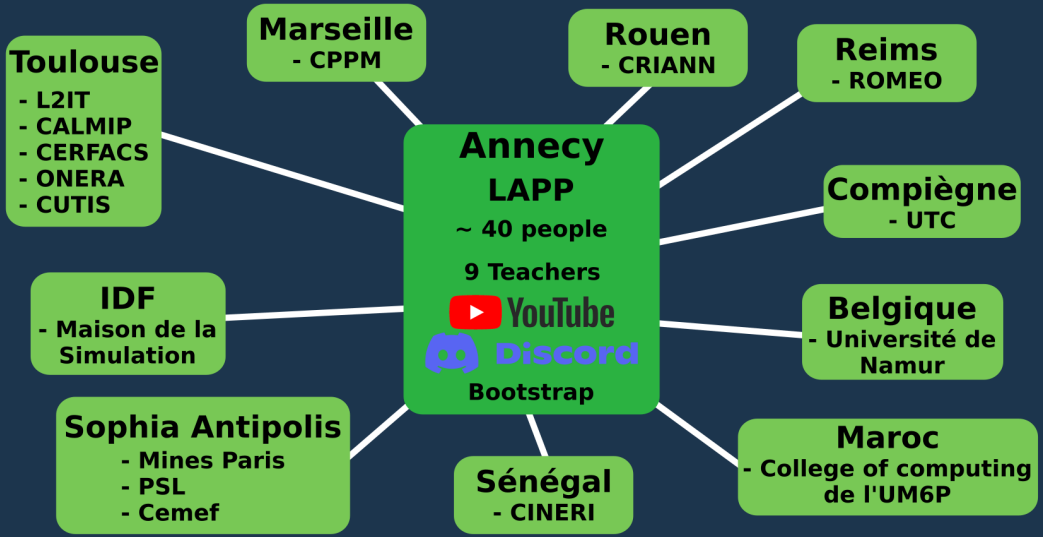

Following Lectures

Following Lectures

Following Lectures

Following Lectures

Following Lectures

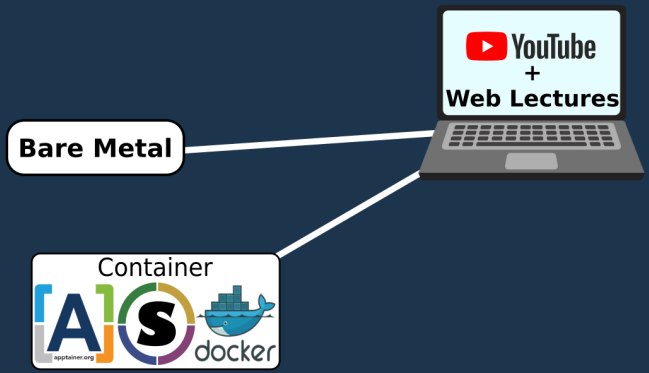

Following Lectures

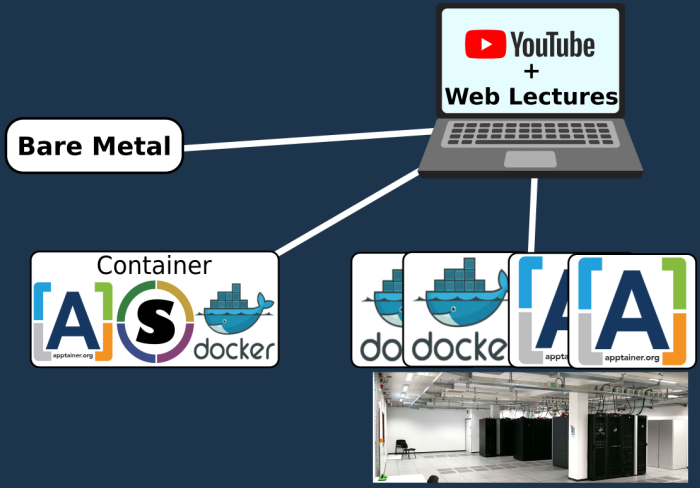

Following Lectures

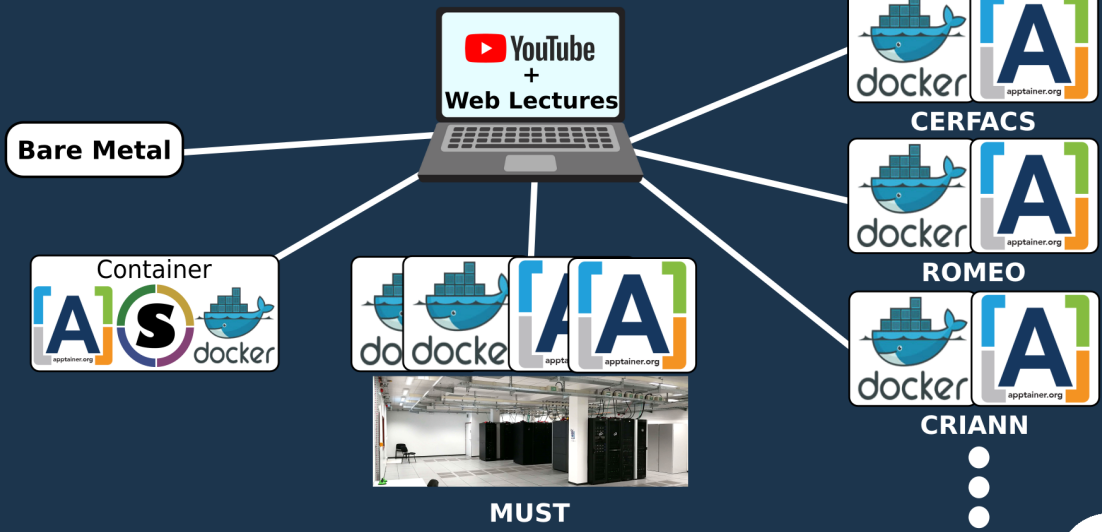

Following Lectures

Hardware :
- **CPU : Intel, AMD**
- **GPU : NVidia** CC >= 6, **AMD**

YouTube
+
Web Lectures

Bare Metal

Container

MUST

CERFACS

ROMEO

CRIANN

Sponsors

CODE RECKONS

Science to the CORE

intel

EVIDEN

an atos business

Sponsors

Computing center access

Engineers helping at bootstrap day

Statistics

- **9** teachers
- **18** webinars (almost **weekly** from **January** to **June**)
- **3** connection modes
- **12** Satellittes Sites over **4** countries (**France**, **Belgium**, **Morocco**, **Senegal**)
- **642** Attendees
  - **~400** people following school on **YouTube**
  - **~200** people following school on **Discord** (**530 Discord accounts** created)
- **~50** Containers-lectures compatible with **Docker**, **Apptainer**, **Podman**, **Kubernetes** and **OpenStack**
- **Lectures available on laptop/remote server/HT condor/slurm (Intel / AMD)**
- **60h** of streamed lecture after editing (127 GB)
- **1 bootstrap day** (Annecy + Toulouse)
- **1 extra day** (hands on + discussion) @ Annecy

Statistics

- **9** teachers
- **18** webinars (almost **weekly** from **January** to **June**)
- **3** connection modes
- **12** Satellittes Sites over **4** countries (**France**, **Belgium**, **Morocco**, **Senegal**)
- **642** Attendees
  - **~400** people following school on **YouTube**
  - **~200** people following school on **Discord** (**530 Discord accounts** created)
- **~50** Containers-lectures compatible with **Docker**, **Apptainer**, **Podman**, **Kubernetes** and **OpenStack**
- **Lectures available on laptop/remote server/HT condor/slurm (Intel / AMD)**
- **60h** of streamed lecture after editing (127 GB)
- **1 bootstrap day** (Annecy + Toulouse)
- **1 extra day** (hands on + discussion) @ Annecy

Next school : 23 june - 04 July 2025