

Cloud Infra & Cloud Native Ecosystem

<https://indico.global/event/18402/>

Sep 23rd 2026, Ricardo Rocha, CERN

Who am I



Lead Platforms Infrastructure @ CERN

Kubernetes & Cloud Native, Machine Learning

Cloud Native Computing Foundation (CNCF)

Previously in the Technical Oversight Committee (TOC)

Currently chair of the End User Technical Advisory Board (TAB)

Member of the Governing Board (GB)

Motivation(s)

Strong need to simplify and streamline infrastructure management

A growing need for distributed, heterogeneous resources (CPUs, GPUs, TPUs, ...)


A trend for those resources to be available remotely, more than on-premises

Consolidation of scarce resources among multiple services

Consolidation of multiple workloads types, interactive computing, batch or ML



```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx-deployment
  labels:
    app: nginx
spec:
  replicas: 3
  selector:
    matchLabels:
      app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
        - name: nginx
          image: nginx:1.14.2
          ports:
            - containerPort: 80
```




Declarative API

Automatic reconciliation

Standardization across multiple providers

A very large ecosystem, a very large community

723	3	17	19	24	77	583
TOTAL	ACADEMIC	GOLD	PLATINUM	NONPROFIT	END USER SUPPORTER ...	SILVER
(0.41%)	(2.35%)	(2.63%)	(3.32%)	(10.65%)	(80.64%)	



CNCF projects are the foundation of cloud and AI native computing

As part of the **Linux Foundation**, we provide support, oversight and direction for fast-growing, **cloud native** projects, including Kubernetes, Envoy, and Prometheus.

39

GRADUATED
PROJECTS


38

INCUBATING
PROJECTS

151

SANDBOX
PROJECTS

ALL PROJECTS





Database

Streaming & Messaging

Application Definition & Image Build

Continuous Integration & Delivery



Scheduling & Orchestration

Coordination & Service Discovery

Remote Procedure Call

Service Proxy

API Gateway


Service Mesh






Cloud Native Storage

Container Runtime

Cloud Native Network







Sustainability

A lot of the governance hard work is handled in the foundation

A good stamp on project maturity ... but not enough

e.g. At CERN we use several sandbox projects in production

We adapt our internal due diligence accordingly



Sustainability

A lot of the governance hard work is handled in the foundation

A good stamp on project maturity ... but not enough

e.g. At CERN we use several sandbox projects in production

We adapt our internal due diligence accordingly

☐	REPOSITORY	ACTIVITY	PR OVERVIEW	LOTTERY FACTOR	OSSF SCORE	CONTRIBUTOR CONFIDENCE	CONTRIBUTORS	LAST 180 DAYS
☐	 helm @helm	↗ High	202 PRs ↑ 50%	Low	9.5/10	20%	 +116	
☐	 etcd @etcd-io	↗ High	591 PRs ↑ 46%	Low	8.9/10	11%	 +68	
☐	 containerd @containerd	↗ High	457 PRs ↑ 42%	Low	8.8/10	18%	 +138	
☐	 prometheus @prometheus	↗ High	725 PRs ↑ 44%	Low	7.2/10	17%	 +189	
☐	 cilium @cilium	↗ High	1431 PRs ↑ 28%	Low	7.2/10	27%	 +256	
☐	 argo-cd @argoproj	↗ High	1304 PRs ↑ 45%	Low	8.4/10	25%	 +361	
☐	 coredns @coredns	↗ High	208 PRs ↑ 51%	Low	8.1/10	15%	 +42	
☐	 kubernetes @kubernetes	↗ High	2329 PRs ↑ 44%	Low	7.2/10	15%	+690	

Sustainability

A lot of the governance hard work is handled in the foundation

A good stamp on project maturity ... but not enough

e.g. At CERN we use several sandbox projects in production

We adapt our internal due diligence accordingly

REPOSITORY	ACTIVITY	PR OVERVIEW	LOTTERY FACTOR	OSSF SCORE	C	REPOSITORY	ACTIVITY	PR OVERVIEW	LOTTERY FACTOR	OSSF SCORE	CONTRIBUTOR CONFIDENCE	CONTRIBUTORS	LAST 180 DAYS
 helm @helm	 High	202 PRs  ↑ 50%	Low	9.5/10	20	 ingress-nginx @kubernetes	 High	763 PRs  ↑ 49%	Very High	6.9/10	18%	 +111	
 etcd @etcd-io	 High	591 PRs  ↑ 46%	Low	8.9/10	11	 ContainerSSH @ContainerSSH	 Medium	7 PRs  ↑ 50%	Very High	5.3/10	10%	-	
 containerd @containerd	 High	457 PRs  ↑ 42%	Low	8.8/10	18	 node-feature-gates @kubernetes-	 High	197 PRs  ↑ 51%	Very High	6.9/10	30%	 +17	
 prometheus @prometheus	 High	725 PRs  ↑ 44%	Low	7.2/10	17	 csi-driver-nfs @kubernetes-	 High	83 PRs  ↑ 50%	Very High	7.6/10	31%	 +8	
 cilium @cilium	 High	1431 PRs  ↑ 28%	Low	7.2/10	27	 gpu-operator @NVIDIA	 High	263 PRs  ↑ 50%	Very High	6.2/10	17%	 +21	
 argo-cd @argoproj	 High	1304 PRs  ↑ 45%	Low	8.4/10	25	 stargz-snapshotter @containerd	 High	169 PRs  ↑ 47%	Very High	6.6/10	17%	 +3	
 coredns @coredns	 High	208 PRs  ↑ 51%	Low	8.1/10	15								
 kubernetes @kubernetes	 High	2329 PRs  ↑ 44%	Low	7.2/10	15								

November 2024


OK


FOLLOW

REACT

Project	Status	Criticality	Exit Complexity	Maturity	Lottery Factor	Comments
helm	OK	High	High	Graduated	Low	Some kustomize usage but mostly helm across our deployments
etcd	OK	High	High	Graduated	Low	Troubled times with maintainership in 2023. CNCF came forward to fix them
containerd	OK	High	Medium	Graduated	Low	Possible to move to a different container runtime (if ever needed)
prometheus	OK	High	High	Graduated	Low	De facto standard. Separation between the format and the tool (in principle)
cilium	OK	Low	Low	Graduated	Low	Criticality can change if we start using cluster mesh extensively. Not yet the case
argo-cd	OK	Medium	High	Graduated	Low	Tendency for criticality to increase. Project lock-in but healthy vendor diversity
coredns	OK	High	Medium	Graduated	Low	Building block. Barely noticeable but a critical component
kubernetes	OK	High	High	Graduated	Low	The core of the stack

Quick History





70 TB Dataset



OpenStack Magnum

25000 Kubernetes Jobs



Job Results



**Interactive
Visualization**


Aggregation

Kubecon Higgs Analysis: **70TB in 5min**


Throughput up to **32Gb per VM**

In memory storage used to handle this


Peak at close to **200GB/s**



Network Traffic Received



High Throughput



70 TB Dataset

70 TB of Physics Data ~25000 Files

200+ Core K8s Clusters

Make Pull!

Event Data

Background


Summary Data

Real Data

Google Cloud


Kubernetes

CloudNativeCon



Aggregation


A flagship project: Next Generation Triggers



Next Generation Triggers

Innovative computing technologies for data acquisition and processing for the High-Luminosity Large Hadron Collider and beyond


Our research




ATLAS and CMS experts in NextGen develop new workflows and data processing techniques to increase the sensitivity of future particle physics triggers



NextGen develops and benchmarks software and techniques to exploit accelerated computing architectures keeping in mind cost and energy efficiency



NextGen investigates the use of AI technologies to improve the experiments' physics impact leveraging massive data throughput pipelines and large-scale models



Theoretical physicists and software engineers in NextGen look at new event generators techniques to improve simulation and detection of exotic signatures




NextGen works with the High-Energy Physics community and computer science experts to provide new generations of scientists with the right skills for tomorrow's challenges

<https://nextgentriggers.web.cern.ch>



WP1: Infrastructure, Algorithms and Theory

Researchers in NextGen are making use of neural networks optimisation, quantum-inspired algorithms, high-performance computing, and field-programmable gate array (FPGA) techniques to improve the theoretical modelling and optimise their tools and instruments for the search of ultra rare events.



ATLAS and CMS experts in NextGen develop new workflows and data processing techniques to increase the sensitivity of future particle physics triggers



NextGen develops and benchmarks software and techniques to exploit accelerated computing architectures keeping in mind cost and energy efficiency



NextGen of AI tech the impact data through larg




WP4: Education Programmes and Outreach

A dedicated work package aims at providing a unique multi-disciplinary training programme for the researchers involved in the project and the community at large. It also design targeted events and conferences for the broader community of physicists and interested scientists and interested in the fi



h



WP2: Enhancing the ATLAS Trigger and Data Acquisition

This work package focuses on an enhancement of the already ambitious upgrade of the ATLAS experiment's trigger and data acquisition system for the High Luminosity Phase of the LHC (HL-LHC) scheduled to start in 2029. Novel approaches to trigger event selection will extend the ATLAS physics potential, state-of-the-art Machine Learning techniques.



WP3: Rethinking the CMS Real-Time Data Processing

This work package aims to rethink the CMS data acquisition system allowing the CMS physics program to operate over all the collisions produced by the LHC. This is achieved through the High-Level Trigger (HLT) Real-time Reconstruction Revolution (R³) and a novel L1-trigger scouting stream.

<https://nextgentriggers.web.cern.ch>

Nvidia	AMD	FPGA	CPU		
Description	Nodes	Cores	RAM	GPU	Network
Nvidia H100 188GB NVL	12	192	1.5TB	8	100G
Nvidia H100 80GB SXM	6	64	1TB	4	400G
Nvidia L40S 48GB	7	128	1.5TB	4	200G
Nvidia RTX6000 96GB	2	96	720GB	4	10G

Nvidia	AMD	FPGA	CPU
Description	Nodes	Cores	RAM
AMD Instinct MI300X 192GB	2	128	1.5TB
AMD Radeon Pro W7900 48GB	6	128	770GB

					CPU Intensive Node
Nvidia	AMD	FPGA	CPU		
Description	Nodes	Cores	RAM	FPGA	Network
Xilinx AUP PYNQ Z2	2	48	256GB	15	100G

Nvidia	AMD	FPGA	CPU			
Description	Nodes	Cores	RAM	Network	node/type	
High Core Count	2	384	1.5 TB	100G	HIGHCORE-384	
High Core Frequency	2	64	770 GB	100G	HIGHFREQUENCY-3.8	
CPU Intensive Node	4	192	1.5 TB	100G	CPU-INTENSIVE	

Nvidia	AMD	FPGA	CPU		
Description	Nodes	Cores	RAM	GPU	Network
Nvidia H100 188GB NVL	12	192	1.5TB	8	100G
Nvidia H100 80GB SXM	6	64	1TB	4	400G
Nvidia L40S 48GB					
Nvidia RTX6000 96GB					
Nvidia	AMD	FPGA	CPU	GPU	Network
Description	Nodes	Cores	RAM	GPU	Network
AMD Instinct MI300X 192G	2	48	256GB	15	100G
AMD Radeon Pro W7900 41GB					
Nvidia	AMD	FPGA	CPU	GPU	Network
Description	Nodes	Cores	RAM	GPU	Network
Xilinx AUP PYNQ Z2	2	48	256GB	15	100G


GPU	CPU	RAM	/scratch	nvidia.com/gpu.product
Nvidia H100 NVL 94GB	23	180 GB	450 GB	NVIDIA-H100-NVL
Nvidia H100 NVL 47GB	11	90 GB	225 GB	NVIDIA-H100-NVL
Nvidia H100 NVL 12GB	3	25 GB	60 GB	NVIDIA-H100-NVL
Nvidia H100 SXM 80GB	15	240 GB	450 GB	NVIDIA-H100-80GB-HBM3
Nvidia L40S 48GB	31	360 GB	380 GB	NVIDIA-L40S
Nvidia RTX6000 96GB	22	150 GB	250 GB	NVIDIA-RTX-PRO-6000-Blackwell-Server-Ed


node/type
HIGHCORE-384
HIGHFREQUENCY-3.8
CPU-INTENSIVE

mlflow™



ContainerSSH
Launch containers on demand





Interactive Sessions

The example below defines an interactive session.

You can customize it to the image, resources and services you would like to use - in particular the number and type of GPU, with minor differences for Nvidia and AMD.

Nvidia AMD

session.yaml

```
apiVersion: v1
kind: Pod
metadata:
  name: session-1
  labels:
    home: "eos"
    mount-eos: "true"
    mount-cvmfs: "true"
    mount-afs: "true"
spec:
  containers:
  - name: container
    image: registry.cern.ch/ngt/lxplus-like:9
    command: ["/usr/bin/tini", "--", "sleep", "infinity"]
    resources:
      limits:
        nvidia.com/gpu: 1
    nodeSelector:
      nvidia.com/gpu.product: NVIDIA-H100-NVL
```

Info

Change home: eos with home: afs in the example above to have

Pass here the model of the GPU you would like to use. [Check this page](#) for the complete list of available flavors and the equivalent value for this field, as well as the total numbers available in the cluster.

You might also want to check our [monitoring page](#) for real time

Table of contents

[Access Request](#)

[Client Setup](#)

Interactive Sessions

[Session Access](#)

[SSH](#)

[VSCode](#)

[kubectl](#)

[Machine Learning](#)

[Distributed Processing](#)

[MPI](#)

[Storage](#)

[Monitoring](#)

[Storage](#)

Access is available to CVMFS, AFS and EOS to all workloads.

[Remote Access](#)

Tip

Your EOS home directory is the same as your CERNBOX. You can rely on this to retrieve files from your workloads.

In addition workloads have cluster local `/scratch` and `/shared` areas automatically mounted.

Path	Description	Storage
/scratch	Data is cleaned after session termination	According to type and number of GPUs
/shared	Data is shared among different sessions and persisted after session termination. Not currently available for GitLab/GitHub runners	500GB per user

Interactive Sessions

The example below defines an interactive session.

You can customize it to the image, reso
number and type of GPU, with minor dif

NvidiaAMD

session.yaml

```
apiVersion: v1
kind: Pod
metadata:
  name: session-1
  labels:
    home: "eos"
    mount-eos: "true"
    mount-cvmfs: "true"
    mount-afs: "true"
spec:
  containers:
    - name: container
      image: registry.cern.ch/ngt/
      command: ["/usr/bin/tini", "
      resources:
        limits:
          nvidia.com/gpu: 1
      nodeSelector:
        nvidia.com/gpu.product: NVIDIA
```

Info

Change home: eos with home: afs in the example above to have

Image	Description
registry.cern.ch/ngt/lxplus-like:9	lxplus-gpu like environment, based on alma9, with similar tools
registry.cern.ch/ngt/lxplus-like:8	lxplus8-gpu like environment, based on alma8, with similar tools
registry.cern.ch/ngt/lxplus-like:9-hls4ml-vivado2022	lxplus-like:9 environment with hls4ml, vivado 2022.2 and relevant ML frameworks/packages
registry.cern.ch/ngt/lxplus-like:9-hls4ml-vivado2023	lxplus-like:9 environment with hls4ml, vivado 2023.2 and relevant ML frameworks/packages
registry.cern.ch/ngt/lxplus-like:9-hls4ml-vivado2024	lxplus-like:9 environment with hls4ml, vivado 2024.1 and relevant ML frameworks/packages
registry.cern.ch/ngt/lxplus-like:9-hls4ml-vivado2025	lxplus-like:9 environment with hls4ml, vivado 2025.2 and relevant ML frameworks/packages
registry.cern.ch/ngt/pytorch:2.3.1	PyTorch environment
registry.cern.ch/ngt/tensorflow:2.16.2	TensorFlow v2.16.2 environment

this page for the complete list of available flavors and the equivalent value for this field, as well as the total numbers available in the cluster.

You might also want to check our [monitoring page](#) for real time

Table of contents

Access Request

Client Setup

to all workloads.

IOX. You can rely on this to retrieve files from your workloads.

branch and /shared areas automatically mounted.


Storage	
ination	According to type and number of GPUs
ssions and persisted after available for GitLab/GitHub	500GB per user

Distributed Processing & Pipelines



Available directly or via pipelines / workflows

Training Operators, Ray, Spark on Analytix

Shared workflows, pre-built for common scenarios




Examples: HEPScore23, hls4ml synthesis, (soon) model quantization




```

apiVersion: kubeflow.org/v1
kind: MPIJob
metadata:
  name: mpi-devel +
spec:
  slotsPerWorker: 1 +
  runPolicy:
    cleanPodPolicy: Running +
  mpiReplicaSpecs:
    Launcher:
      restartPolicy: OnFailure
      replicas: 1
      template:
        metadata:
          labels:
            home: "eos" +
            mount-eos: "true" +
            mount-cvmfs: "true" +
            mount-afs: "true" +
        spec:
          containers:
            - image: registry.cern.ch/mlops/mpi:openmpi-5.0.8 +
              name: launcher
              command: ["/usr/bin/tini", "--", "sleep", "infinity"]

```



S



Worker:

```

replicas: 2 +
template:
  metadata:
    labels:
      home: "eos" +
      mount-eos: "true" +
      mount-cvmfs: "true" +
      mount-afs: "true" +
  spec:
    containers:
      - image: registry.cern.ch/mlops/mpi:openmpi-5.0.8 +
        name: worker
        command: ["/usr/bin/tini", "--", "sleep", "infinity"]
        resources:
          limits:
            nvidia.com/gpu: 1 +
    nodeSelector:
      nvidia.com/gpu.product: NVIDIA-A100-PCIE-40GB +

```




Infrastructure

CPU, memory and scratch space automatically assigned per flavor

Improved workload isolation (NUMA awareness, CPU / GPU affinities, CPU pinning)

Up to 30% gains depending on the workload




Infrastructure

CPU, memory and scratch space automatically assigned per flavor

Improved workload isolation (NUMA awareness, CPU / GPU affinities, CPU pinning)

Up to 30% gains depending on the workload





Infrastructure

CPU, memory and scratch space automatically assigned per flavor

Improved workload isolation (NUMA awareness, CPU / GPU affinities, CPU pinning)

Up to 30% gains depending on the workload

Support for low latency interconnects (Infiniband, RoCEv2)



NGT MPI bandwidth benchmark with different multi-node interconnects - Infiniband, RDMA over Converged Ethernet (RoCEv2), etc.



Notebooks

 **Filter** Enter property name or value

+ New Notebook



ContainerSSH



```

fcardeo@rochaxali ~/$ ssh session-afs-011@ngt.cern.ch
hepix: W: Your home directory is 96% full. This may
hepix: W: cause problems, especially when logging in.
hepix: W: Please remove unwanted files, or go to
hepix: W: http://cern.ch/account to increase your quota.
rbritoda@session-afs-011:~$ whoami
rbritoda
rbritoda@session-afs-011:~$ sudo whoami
root
rbritoda@session-afs-011:~$ nvidia-smi
Tue May 27 23:44:02 2025
+-----+
| NVIDIA-SMI 570.124.06                  Driver Version: 570.124.06      CUDA Version: 12.8     |
+-----+-----+-----+-----+-----+-----+
| GPU   Name                               Persistence-M   Bus-Id        Disp.A    Volatile Uncorr. ECC  |
|  Name   Temp      Perf                Pwr:Usage/Cap       Memory-Usage   GPU-Util    Compute M. |
|                               Pwr:Usage/Cap       Memory-Usage   GPU-Util    Compute M. |
|-----+-----+-----+-----+-----+-----+-----+
|  0  NVIDIA H100 NVL                      Off               00000000:01:00.0 Off          0
| N/A   45C      P0                        64W / 400W       1MiB / 95830MiB      0%          Default
|                               Pwr:Usage/Cap       Memory-Usage   GPU-Util    Compute M. |
|                               Pwr:Usage/Cap       Memory-Usage   GPU-Util    Compute M. |
+-----+-----+-----+-----+-----+-----+
|                               Pwr:Usage/Cap       Memory-Usage   GPU-Util    Compute M. |
|                               Pwr:Usage/Cap       Memory-Usage   GPU-Util    Compute M. |
+-----+-----+-----+-----+-----+-----+
|                               Pwr:Usage/Cap       Memory-Usage   GPU-Util    Compute M. |
|                               Pwr:Usage/Cap       Memory-Usage   GPU-Util    Compute M. |
+-----+-----+-----+-----+-----+-----+
+-----+
| Processes:
| GPU   GI    CI          PID    Type    Process name                        GPU Memory
|      ID    ID                                     name                                Usage
|-----+-----+-----+-----+-----+-----+
| No running processes found
+-----+
rbritoda@session-afs-011:~$

```

Demo

Lessons Learned

A paradigm shift (workloads, not nodes) with a significant learning curve

Upstream first, community first (and that takes effort and time)

Internal due diligence still needed, not a one size fits all

In the era of the GPU... cloud is no longer on-demand, though that might change

Hiring and onboarding new people greatly improved

The same stack now usable for services, scientific computing and ML workloads

Q & A