



KubeCon



CloudNativeCon

Europe 2026

#KubeCon #CloudNativeCon



NextGen



Collisions in the Dark: Illuminating the 95% of Kubeflow You Can't See

Amine Lahouel, Laura Llinares - CERN



CERN and the Collisions




KubeCon




CloudNativeCon


Europe 2026




Next Generation Triggers



Reconstruction of Particle Positions




Anomaly Detection for New Physics



Magic Computing Platform

Kubecon Talk



One shared pool: multi-vendor GPUs, FPGAs, varied CPU profiles, fast interconnects – across on-prem, cloud, and HPC

Ref Architecture



The "Superposition" of User Requirements



KubeCon




CloudNativeCon

Europe 2026

Researchers experimenting with heterogeneous workloads – having different needs and problems to solve

MLOps team have limited by-default observability on the workloads



The Observability Gap

Standard k8s stack

Node metrics

CPU, memory, disk, network

Pod metrics

state, restarts, resource requests & usage

Cluster metrics

deployment status, replica counts, object health

Dashboards

Cluster, node and pod level dashboards

Infrastructure and CPU-centric

Meanwhile, users ask ...




Data Scientist

Why is my training so slow ?



Ops

How much energy did you tell me all those GPUs consume ?



Management

How efficient are those GPUs being used ?

Needed stack

GPU telemetry

Power & Energy

Energy attribution

Capacity insights

Smart alerting

Platform & ML-workloads-centric

The "Dark Matter" of Metrics



Platform Usage

Active users, sessions, notebook count, job submissions over time



GPU Telemetry

Utilisation, memory, temperature, frequency across all accelerator vendors



Resource Availability

Real-time view of free vs. allocated GPUs, CPUs, and memory per node pool



Power Consumption

Hardware-level energy draw at node and accelerator level



Sustainability

CO₂-equivalent emissions per cluster and per individual workload



Idle Resource Detection

Zombie jobs, unused allocations, hoarder namespaces

Architecting the Light




KubeCon



CloudNativeCon

Europe 2026



Machine Exporter




KubeCon



CloudNativeCon

Europe 2026



kube-state-metrics

Metric	Description
<code>kube_pod_container_resource_requests</code>	Resource (CPU/memory/GPU) requested by container
<code>kube_node_spec_taint</code>	Taints applied to a node
<code>kube_namespace_labels</code>	Labels attached to a namespace
<code>kube_pod_status_phase</code>	Current lifecycle phase of a pod
<code>kube_pod_info</code>	General pod metadata (node, namespace, UID)
<code>kube_pod_start_time</code>	Timestamp when the pod started
<code>kube_node_spec_unschedulable</code>	Whether a node is cordoned/unschedulable
<code>kube_node_status_allocatable</code>	Allocatable resources on a node
<code>kube_pod_created</code>	Timestamp when the pod was created
<code>kube_pod_container_resource_limits</code>	Resource limits set on a container
<code>kube_node_info</code>	Node metadata (kernel, OS, container runtime)
<code>kube_pod_status_scheduled</code>	Whether a pod has been scheduled to a node

node_exporter

Metric	Description
<code>node_hwmon_temp_celsius</code>	Hardware monitor temperature sensor readings (°C)
<code>node_memory_MemTotal_bytes</code>	Total physical memory on the node (bytes)
<code>node_memory_MemAvailable_bytes</code>	Available memory on the node (bytes)
<code>node_cpu_scaling_frequency_hertz</code>	Current CPU scaling (dynamic) frequency (Hz)

cAdvisor

Metric	Description
<code>container_cpu_usage_seconds_total</code>	Cumulative CPU time consumed by container (s)
<code>container_network_receive_bytes_total</code>	Cumulative bytes received by container over network
<code>container_network_transmit_bytes_total</code>	Cumulative bytes transmitted by container over network
<code>container_fs_writes_bytes_total</code>	Cumulative bytes written to filesystem by container
<code>container_fs_reads_bytes_total</code>	Cumulative bytes read from filesystem by container
<code>container_memory_working_set_bytes</code>	Working set memory in use by container (bytes)
<code>container_cpu_core_frequency_hertz</code>	Current CPU core frequency seen by container (Hz)

GPU Exporters




KubeCon



CloudNativeCon

Europe 2026



DCGM (NVIDIA)

Metric	Description
<code>DCGM_FI_PROF_GR_ENGINE_ACTIVE</code>	GPU compute engine active fraction (0–1)
<code>DCGM_FI_DEV_POWER_USAGE</code>	Instantaneous power draw per GPU (W)
<code>DCGM_FI_DEV_GPU_TEMP</code>	GPU die temperature (°C)
<code>DCGM_FI_DEV_GPU_UTIL</code>	GPU utilization percentage
<code>DCGM_FI_DEV_FB_USED</code>	Framebuffer (VRAM) used (MiB)
<code>DCGM_FI_DEV_FB_TOTAL</code>	Total framebuffer (VRAM) capacity (MiB)
<code>DCGM_FI_DRIVER_VERSION</code>	NVIDIA driver version (info metric)
<code>DCGM_FI_DEV_SM_CLOCK</code>	Streaming multiprocessor clock speed (MHz)

AMD GPU

Metric	Description
<code>amd_gpu_gfx_activity</code>	GPU graphics engine utilization (0–1)
<code>amd_gpu_power_usage</code>	Instantaneous power draw per GPU (W)
<code>amd_gpu_used_vram</code>	VRAM used (bytes)
<code>amd_gpu_total_vram</code>	Total VRAM capacity (bytes)
<code>amd_gpu_junction_temperature</code>	GPU junction (hotspot) temperature (°C)
<code>amd_gpu_clock</code>	Current GPU clock speed (MHz)

Power Probing




KubeCon



CloudNativeCon

Europe 2026



Power Probing




KubeCon



CloudNativeCon

Europe 2026



Handling the Data Deluge



KubeCon



CloudNativeCon

Europe 2026



count[{"__name__=~".+"}]

Table

Graph

<

Evaluation time

>

{ }

2325832

```
prometheus:
  prometheusSpec:
    retention: 30d
    retentionSize: 325GB
    storageSpec:
      volumeClaimTemplate:
        spec:
          storageClassName: ssd
          accessModes: ["ReadWriteOnce"]
          resources:
            requests:
              storage: 362Gi
    replicaExternalLabelName: "__replica__"
    additionalRemoteWrite:
      - url: "https://example.com/api/v1/push"
        writeRelabelConfigs:
          # Drop high-cardinality non-interesting metrics
          - action: keep
            sourceLabels: [ __name__ ]
            regex: "^DCGM_.*|^amd.*|^kepler_node_.*|kepler_container_cpu_joules_total"
```

Debugging the Time Dilation



KubeCon



CloudNativeCon

Europe 2026

Demo time

User/Workload Analytics: Provides namespace-filtered views for individual developers to monitor their specific deployments. This view balances resource efficiency (allocated vs. actual usage) with performance profiling (GPU/CPU/RAM saturation) and power consumption, allowing users to independently debug bottlenecks and optimize job performance.

Debugging the Time Dilation



Debugging the Time Dilation




KubeCon



CloudNativeCon

Europe 2026



Managing the Chaos

Demo time

Cluster Overview: Tracks aggregate utilization (CPU, GPU, RAM, Network, Thermals) and node-level health. It highlights idle resources and historical trends to guide capacity planning.

Managing the Chaos



Managing the Chaos





KubeCon



CloudNativeCon

Europe 2026



Towards a Sustainable Future




KubeCon



CloudNativeCon

Europe 2026

Sustainability Tracking: A dedicated dashboard for CO₂-equivalent emissions, offering transparency into the carbon footprint at both the cluster and individual workload levels.



Challenges and Outlook

New hardware with old-ish kernel
New features challenging our assumptions
MIG partitions have their particularitis

Kepler on the VMs
Visualisation for Kubeflow service-level metrics (Pipelines, Notebooks,
Training Operators, etc ...)

Lighting the Way for Everyone

Configuring each component is tedious and error-prone

Weeks of effort to build from scratch



*From Zero to
Observability in 5
minutes*



Upstream Issue



The Kubeflow Observability Manifest

A simple recipe that deploys

- A sane Prometheus config
- Needed ServiceMonitors
- Pre-built Grafana Dashboards
- Kepler & DCGM/AMD configurations

Words of Wisdom



KubeCon



CloudNativeCon

Europe 2026



Don't let your users operate in the dark,
give them the visibility they deserve

Power-aware AI is no longer optional,
energy metrics belong in every ML platform



Last Slide, Promise!



KubeCon




CloudNativeCon

Europe 2026



Get involved, the issue is open and ready for contributions

What's Next ?





KubeCon



CloudNativeCon

Europe 2026

Thanks! Questions?





KubeCon



CloudNativeCon

Europe 2026

