
CERN
NexTGen
Next Generation Triggers

Next Generation Triggers: WP 1.1 Status and Plans

Raul Chiorescu, Ricardo Rocha

CERN IT

NexTGen
Next Generation Triggers

Task Recap

This task will focus on **designing, procuring, deploying and operating the computing infrastructure (hardware and software) and platforms** required to support the common tasks in WP1 (hardware-aware neural network training workflows and next-generation physics simulations) and the specific activities in WP2 and WP3.

Contacts

Homepage: https://ngt-wp1-1.docs.cern.ch/

Mailing List: ngt-wp1-task1-1@cern.ch

Mattermost Channel: #Task 1.1

Weekly Meetings (Fri 2pm): Indico Category, Agenda and Minutes

Hannes Hansen
(September 24)

Raulian Chiorescu
(June 24)

Ricardo Rocha

Jessy Sobreiro
(September 24)

Amine Lahouel
(January 25)

2024 Review

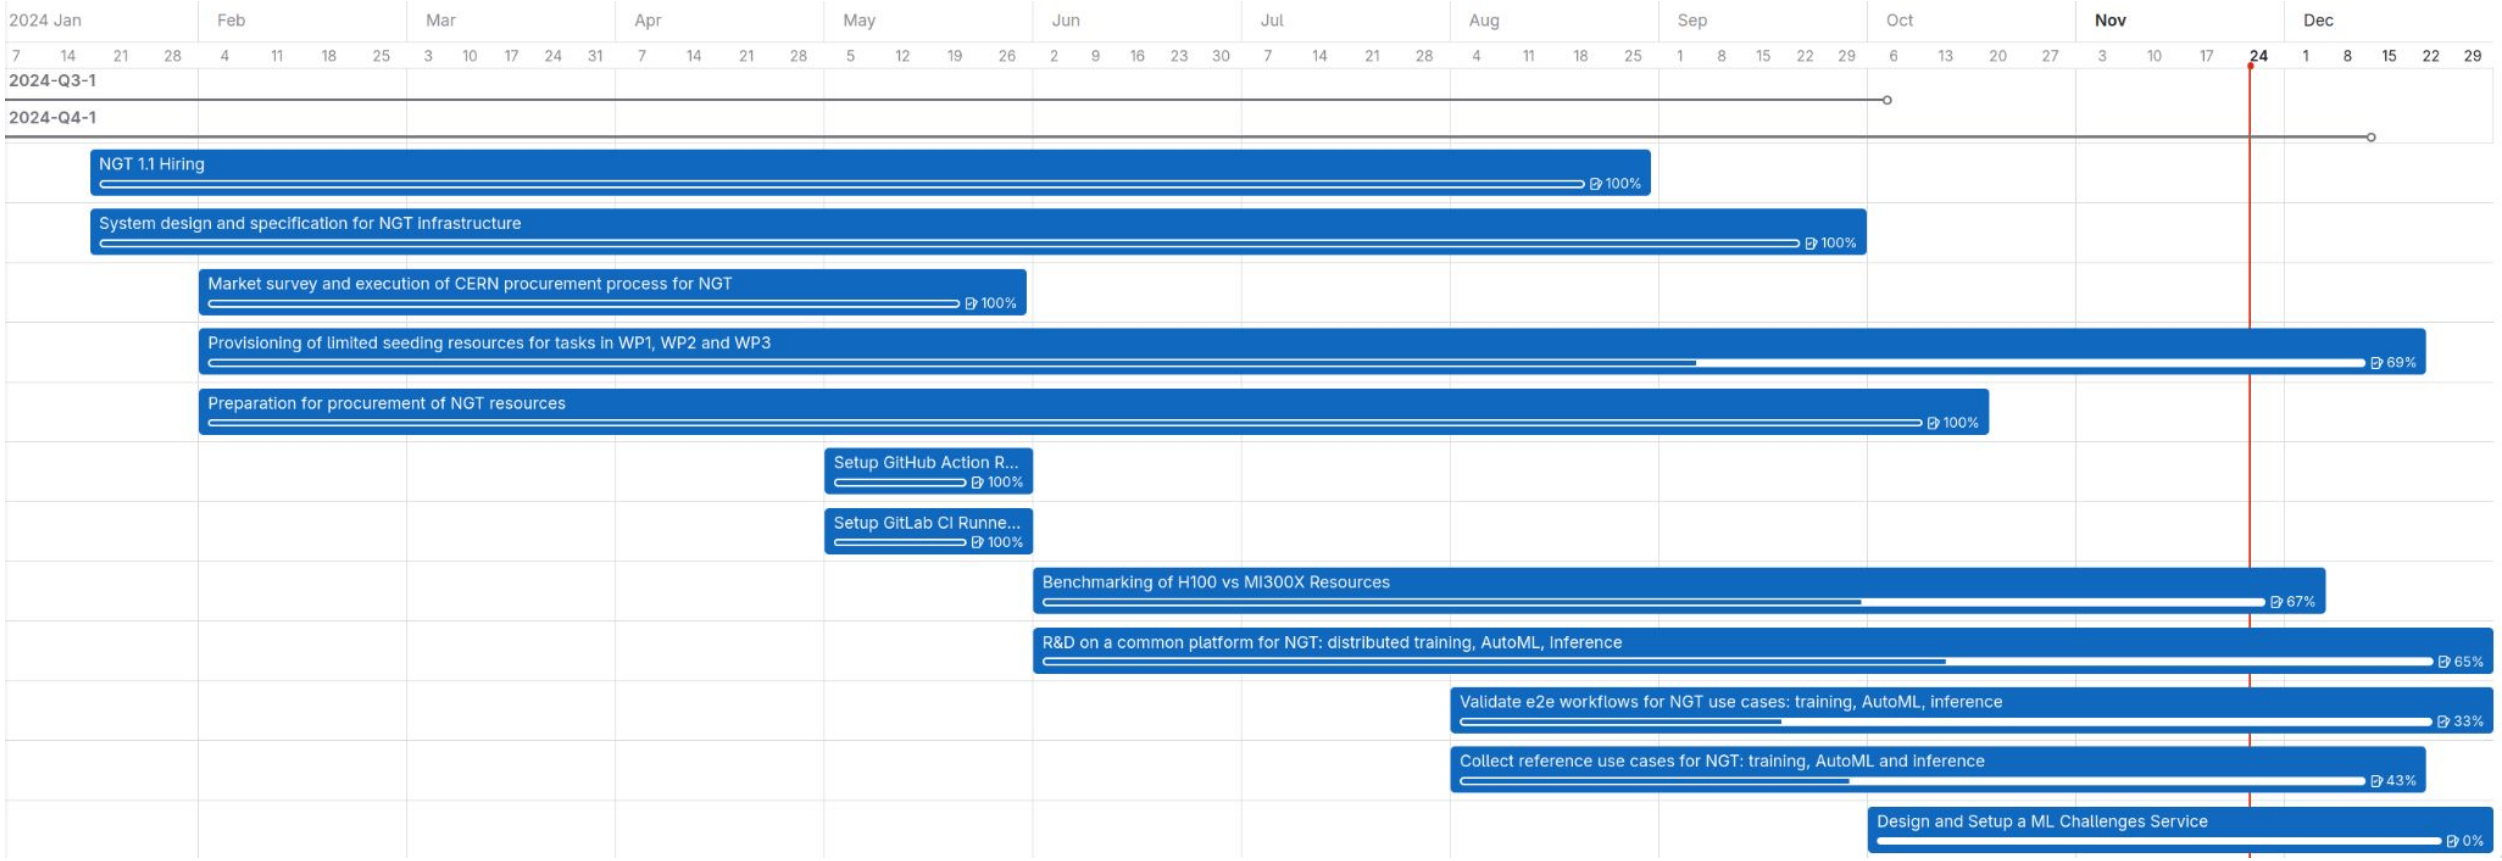
2024 Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
Nov
Dec
2024-Q3-1
2024-Q4-1
NGT 1.1 Hiring
100%
System design and specification for NGT Infrastructure
100%
Market survey and execution of CERN procurement process for NGT
100%
Provisioning of limited seeding resources for tasks in WP1, WP2 and WP3
69%
Preparation for procurement of NGT resources
100%
Setup GitHub Action R...
100%
Setup GitLab CI Runne...
100%
Benchmarking of H100 vs MI300X Resources
67%
R&D on a common platform for NGT: distributed training, AutoML, Inference
65%
Validate e2e workflows for NGT use cases: training, AutoML, inference
33%
Collect reference use cases for NGT: training, AutoML and inference
43%
Design and Setup a ML Challenges Service
0%

Resources: Procurement

Long process to collect requirements and agree on a set of specifications

Key dates

May 2024: Market survey for bulk Nvidia H100 NVL

September 2024: Bids received for bulk Nvidia H100 NVL

October 2024: All hardware specifications completed

November 2024: Orders out for bulk Nvidia H100 NVL resources, quote requests for others

Current estimate for hardware arrival: Q2 2025

Resources: Hardware Specification

Description	Nodes	Resources	Cores	RAM	Network	Status
Nvidia H100 188GB NVL NVLink	12	8x GPUs	192	2TB	100G	Ordered (Nov 24)
Nvidia H100 188GB NVL NVLink	6	4x GPUs	96	2TB	25G + 4x200G IB	Bids Received (Nov 24)
Nvidia L40S 48GB w/ RoCEv2	7	4x GPUs	128	768GB	200G	Quote Request (Nov 24)
AMD MI300X HBM3 192GB w/RoCEv2	2	8x GPUs	128	768GB	200G	Quote Request (Nov 24)
AMD Radeon PRO W7900 48GB w/RoCEv2	6	4x GPUs	128	768GB	200G	Quote Request (Nov 24)
x86 Large Memory	2		256	1TB	100G	Available (Jun 24)
x86 Large CPU Core Count	2		2x256	1TB	100G	Quote Request
x86 High Single Core Performance	2		128	768GB	100G	Quote Request

Resources: Seeding

A survey was sent to collect needs for seeding resources prior to the on-premises hardware arrival

Currently available set of resources, still with direct assignment to individuals / teams

Reminder: possibility to access a larger number of Nvidia A10 GPUs

Node	Type	Status
ngt-highmem-001	High Memory CPU Node	Assigned (June 2024)
ngt-highmem-002	High Memory CPU Node	Assigned (June 2024)
ngt-a100-001	x1 Nvidia A100 (VM)	Assigned (September 2024)
ngt-h100-001	x8 Nvidia H100 (BM)	Assigned (October 2024)
10.78.64.2	x8 Nvidia H100 (OCI)	Unassigned

Benchmarks: Use Cases

A survey was sent to collect different project use cases, to be used for HW benchmarking

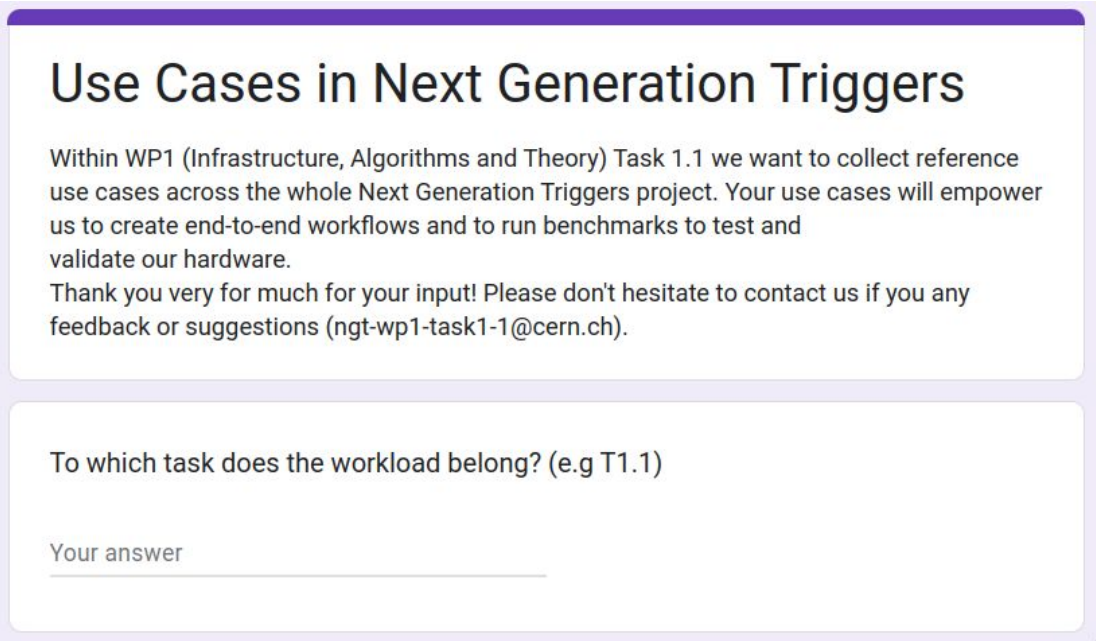

Use Cases in Next Generation Triggers

Within WP1 (Infrastructure, Algorithms and Theory) Task 1.1 we want to collect reference use cases across the whole Next Generation Triggers project. Your use cases will empower us to create end-to-end workflows and to run benchmarks to test and validate our hardware.
Thank you very for much for your input! Please don't hesitate to contact us if you any feedback or suggestions (ngt-wp1-task1-1@cern.ch).

To which task does the workload belong? (e.g T1.1)

Your answer

What resources do you need? (CPU, GPU, FPGA; number and type)

Your answer

In case you have running workloads, what resources do you currently use? (e.g. university clusters, external HPC centers, Tier-1-2-3, ...)

Your answer

Do you have a link to a repository (e.g. GitHub), documentation, paper or any other website that we can use to run the workload and reproduce the results?

Your answer

Benchmarks: Use Cases

A survey was sent to collect different project use cases, to be used for HW benchmarking

Task	Goal	HW Requirements	Contacts
1.2	Optimization of algorithms on FPGAs	Mostly CPU, ~10 GPUs, FPGAs	
1.5	Simulation of Lattice Quantum Field Theory		
2.4	Track Reconstruction for ATLAS L1 Trigger		Benjamin Huth
3.1.1	ML for Full Event Particle Reconstruction	96 GPUs for HP tuning	Joseph Paata, Eric Wulff
3.1.1	CMS reconstruction on GPUs, in particular HLT online selection	2x 128 CPUs + 8x Nvidia L40	
3.7	ML for Search BSM		
CMS	ML for Particle Reconstruction		Jan Kieseler, Philipp Zehetner

Work In Progress

Platform: Services (Current)

First two services deployed and available to the whole project

GitLab CI Runners with GPUs

Integrated into the shared runners in the CERN GitLab CI

```
job:
  tags:
    - k8s-gpu
  image: rochaporto/gpu_burn # overrides the default image.
  script:
    - nvidia-smi
    - cd /app
    - ./gpu_burn 120
```

Platform: Services (Current)

First two services deployed and available to the whole project

GitHub Action Runners with GPUs

Requires explicit setup of the GitHub repo/org, check here for instructions

```yaml
name: self-hosted-workflow
on:
  workflow_dispatch:

jobs:
  test-gpu:
    runs-on: kops-github-arc-scale-set-ngt-gpu
    container: registry.cern.ch/docker.io/rochaporto/gpu_burn
    steps:
      - name: nvidia-smi
        run: |
          nvidia-smi
      - name: gpuburn
        run: |
          /app/gpu_burn 60
```

Platform: Survey

A survey was sent to collect platform and hardware access requirements

19 responses from almost all tasks in the project: **thank you!**

Full results presented in the last WP1.1 weekly meeting

Requirements: Access to NGT Resources

WP1 (Infrastructure, Algorithms and Theory) Task 1.1 is tasked with provisioning and offering access to resources in NGT.

This form tries to collect requirements that will help us understand how to best setup the system for shared, efficient access to these resources.

Thank you very for much for your input! Please don't hesitate to contact us if you any feedback or suggestions (ngt-wp1-task1-1@cern.ch).

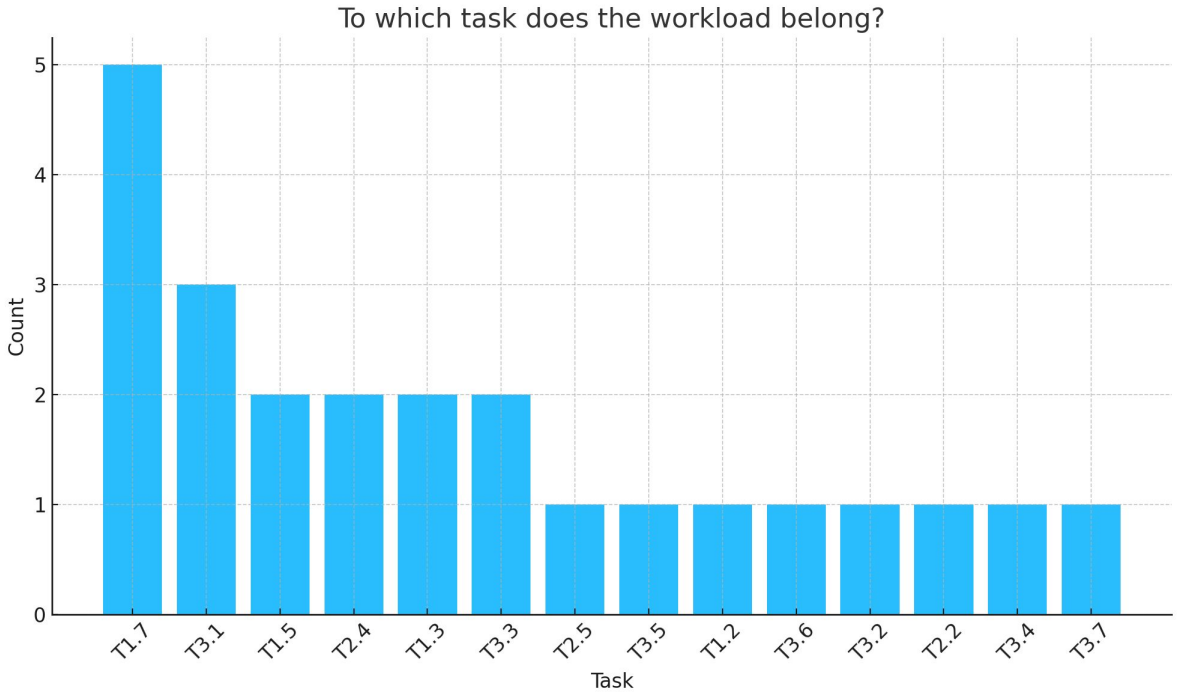

Platform: Survey

What's the expected load / usage pattern? Spiky would be anything that gives an overall usage <30%, full would mean >80% expected usage.

18 responses

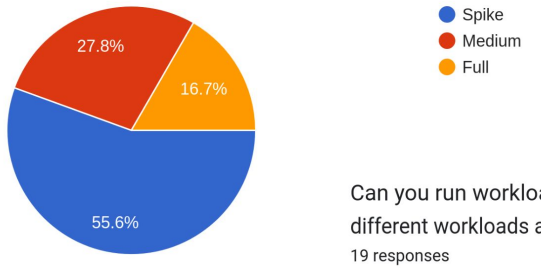

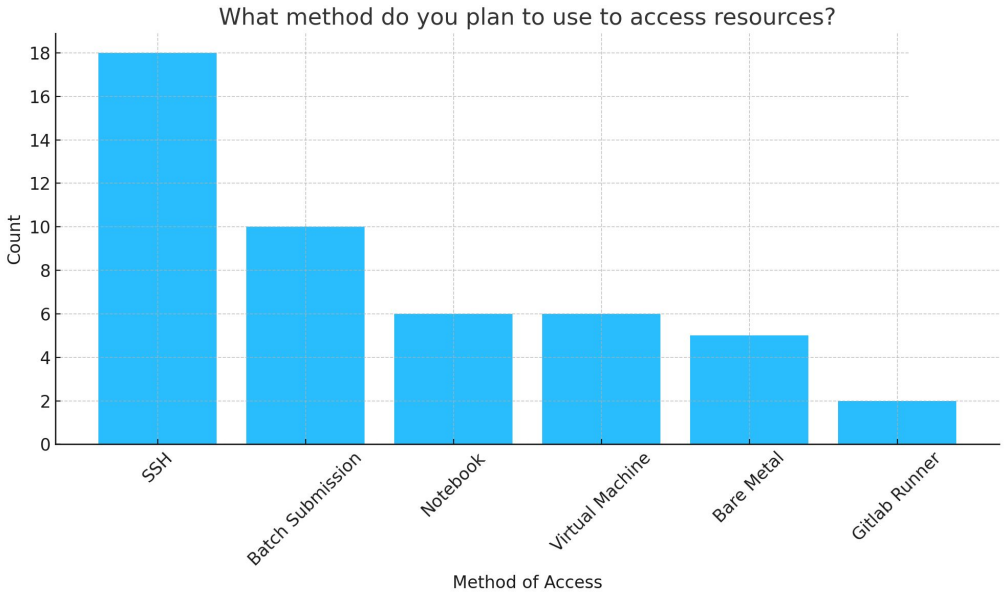

Can you run workloads containerized? Containerization facilitate sharing of the nodes among different workloads and users. When answering no pl...ith other users, if possible (see question below).

19 responses

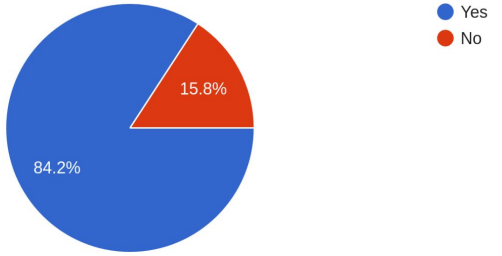

Would your workloads survive automated culling of resources periodically? The goal being to increase overall resource usage.

19 responses

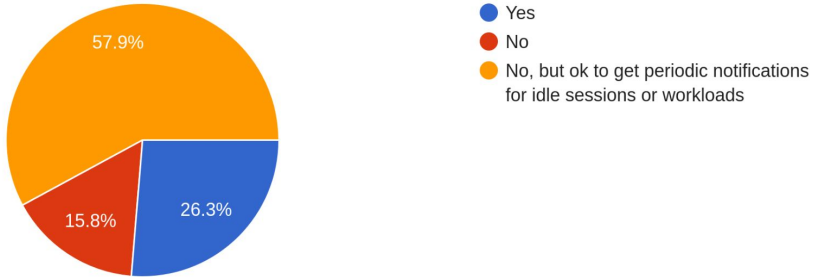

Summary & Next Steps

2024 focused on getting the task up and running, hiring

Significant Achievements

- Hardware specification and on-premises procurement set to be completed by end of 2024
- Initial set of seeding resources available
- Initial set of platform services available (GPUs in GitLab CI, GitHub Actions)

2025 work and planning

- Install and configure all on-premises resources and onboard cloud resources
- Complete use case collection and setup benchmark automation
- Deliver a common platform for shared access to project resources
- Deliver a MLOps platform covering the full ML lifecycle