

**Subject:** Re: Meeting this Wednesday - please prepare  
**From:** Jeff Templon <templon@nikhef.nl>  
**Date:** 04/05/2021, 10:39  
**To:** Helge Meinhard <Helge.Meinhard@cern.ch>

On 4 May 2021, at 10:18, Helge Meinhard wrote:

Hi Jeff,  
as I understood you won't join tomorrow - please let me have your input on what Nikhef can propose in terms of test platforms.  
Cheers,  
Helge

---

Helge MEINHARD CERN, IT Department, CH-1211 Geneve 23,  
Switzerland  
Phone: +41 22 76-76031, Fax: +41 22 76-69935  
E-mail: Helge.Meinhard@cern.ch

Hi Helge,

We have a broad range of platforms, both intel and AMD EPYC. What I am not too sure about is this part:

“It should be systems exclusively reserved for benchmarking for the time we run benchmarks on them, and should be backed by some support capable of starting and following up the jobs (and act in case of issues).”

It's pretty easy for us to drain a system to run a benchmark. “Starting and following up the jobs” is vague as written. What I understood from earlier discussions is this:

- download a container
- spin it up
- container does the rest, including register results in a database

I am sure we could do that as well. If it involves more than that, I'd need to know exactly what, before I am willing to commit support. Not that I don't want to, but our team doesn't have spare time. We had to decline a nice project testing GPUs for LHCb recently for similar reasons; the workflow was not automated enough so things required more babysitting time than we can spare.

The node configurations are here:

[https://wiki.nikhef.nl/grid/NDPF\\_Node\\_Performance](https://wiki.nikhef.nl/grid/NDPF_Node_Performance)

Node classes “marsepein”, “chocolade”, “pepernoot”, “kipsate”, and “lotenfeest” are all still in service and hence could be made available to run benchmarks.

JT