Status of HPC integration with fapptainer

Tomas Lindén[1], Gianfranco Sciacca[2], Ievgen Sliusar[3]

HIP[1], UniBE[2], UiO[3]

24.02.2026

NLCG meeting, Oslo and virtual, 24th of February 2026

Contents

Using fapptainer on HPCs

fapptainer is a tool developed by levgen to remove container restrictions in HPC systems: `https://source.coderefinery.org/slu/fapptainer`

CSC cPouta OpenStack VMs

- *ARC CE:s*
  - **snowarc.hip.fi**, used for sshfs submission to **Mahti**
  - **arc2lumi.hip.fi**, used for sshfs submission to **LUMI**
- *Frontier Squid proxy:*
  - unnamed VM

Using fapptainer on HPCs

Un-nesting containers - Fapptainer tool

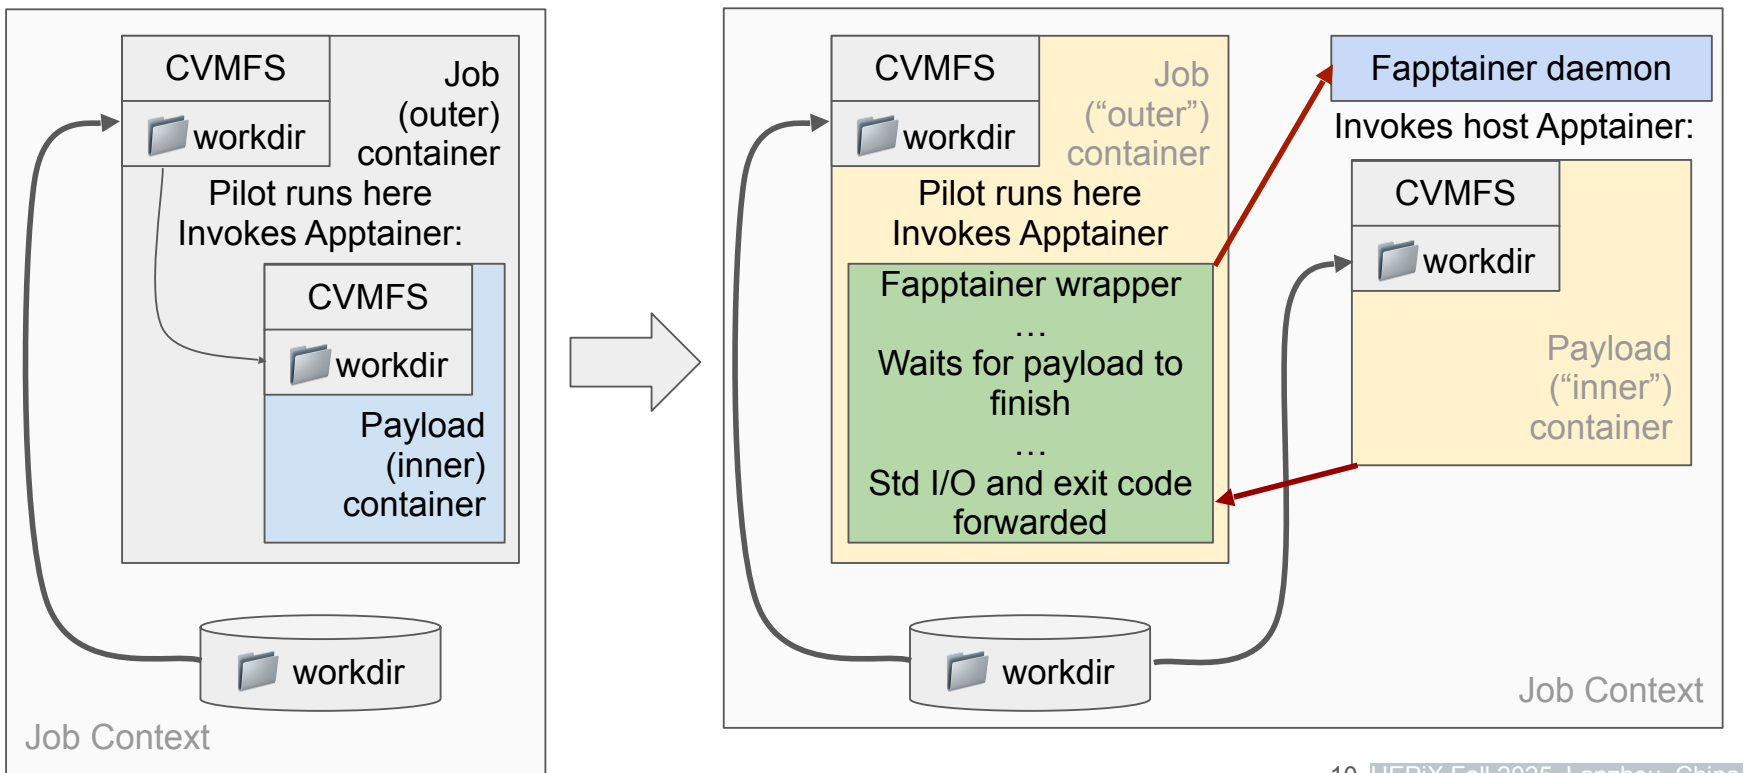

Using fapptainer on HPCs

CSC HPC summary:

		CPU cores	nodes	**disk cores**	disk nodes	disk nodes **RAM/core** (GB)
LUMI	→2027	262 244	2 048	-	-	**2–8**
Puhti	→05/2026	27 280	682	**3760**	94	**4.8–37.5**
Mahti	→08/2026	179 712	1404	**7168**	56	**1.875**
Roihu	04/2026→	186 624	486	**186 624**	486	**2–4**

Using fapptainer on HPCs

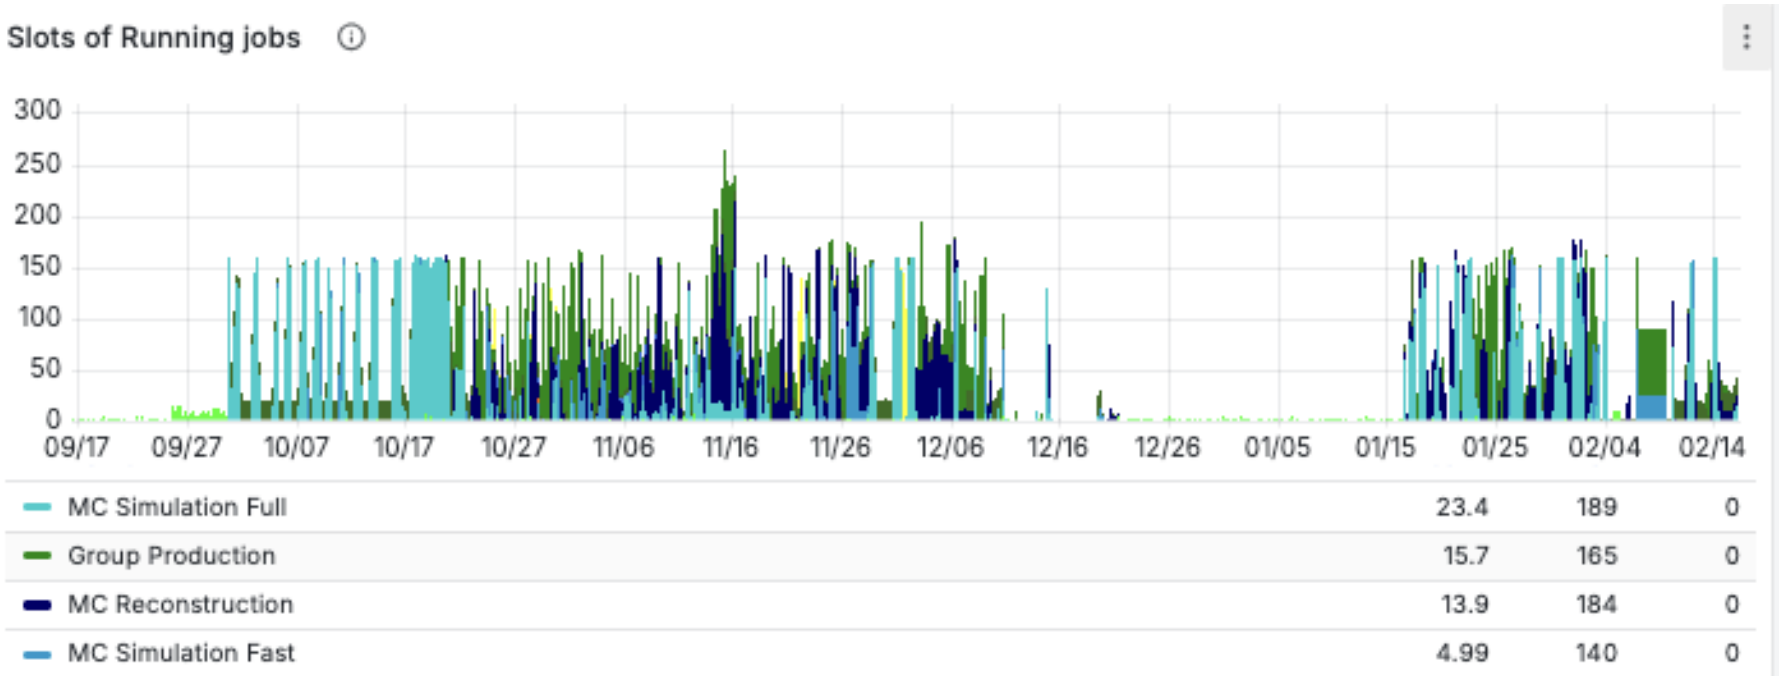

Figure: Timeline of ATLAS jobs on snowarc for the past five months.

Using fapptainer on HPCs

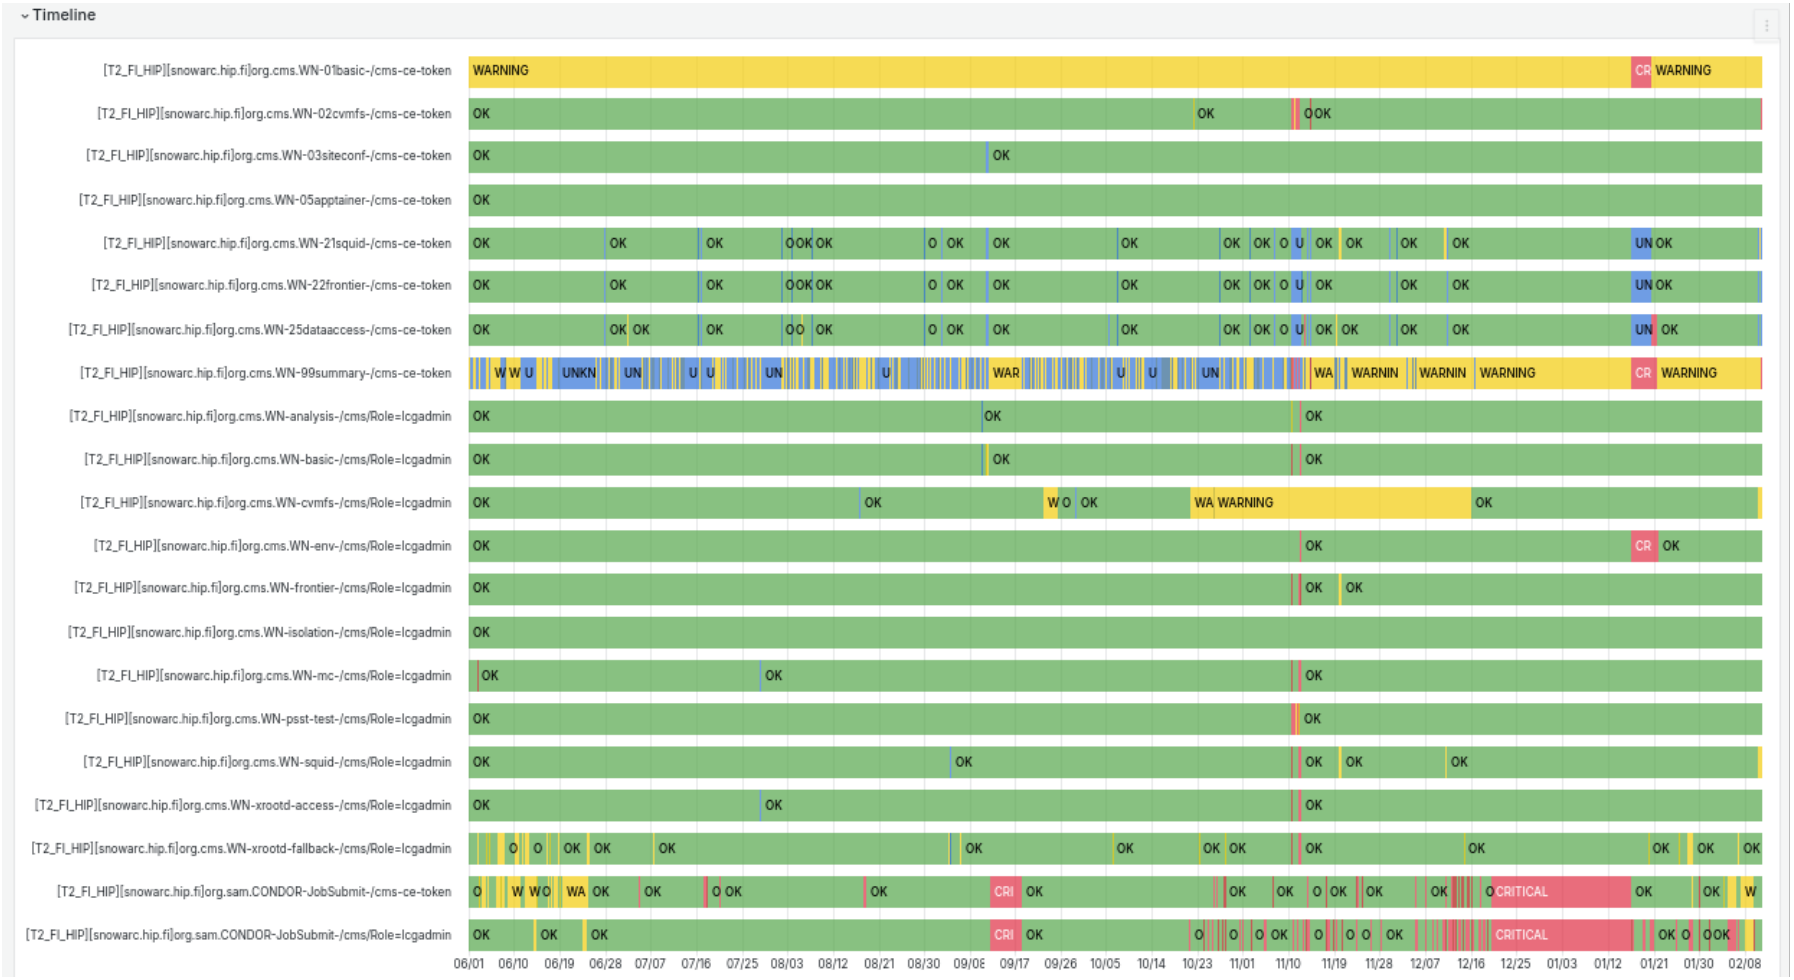

Figure: Timeline of CMS SAM/ETF tests on snowarc from the 1st of June 2025 to the 11th of February 2026.

Using fapptainer on HPCs

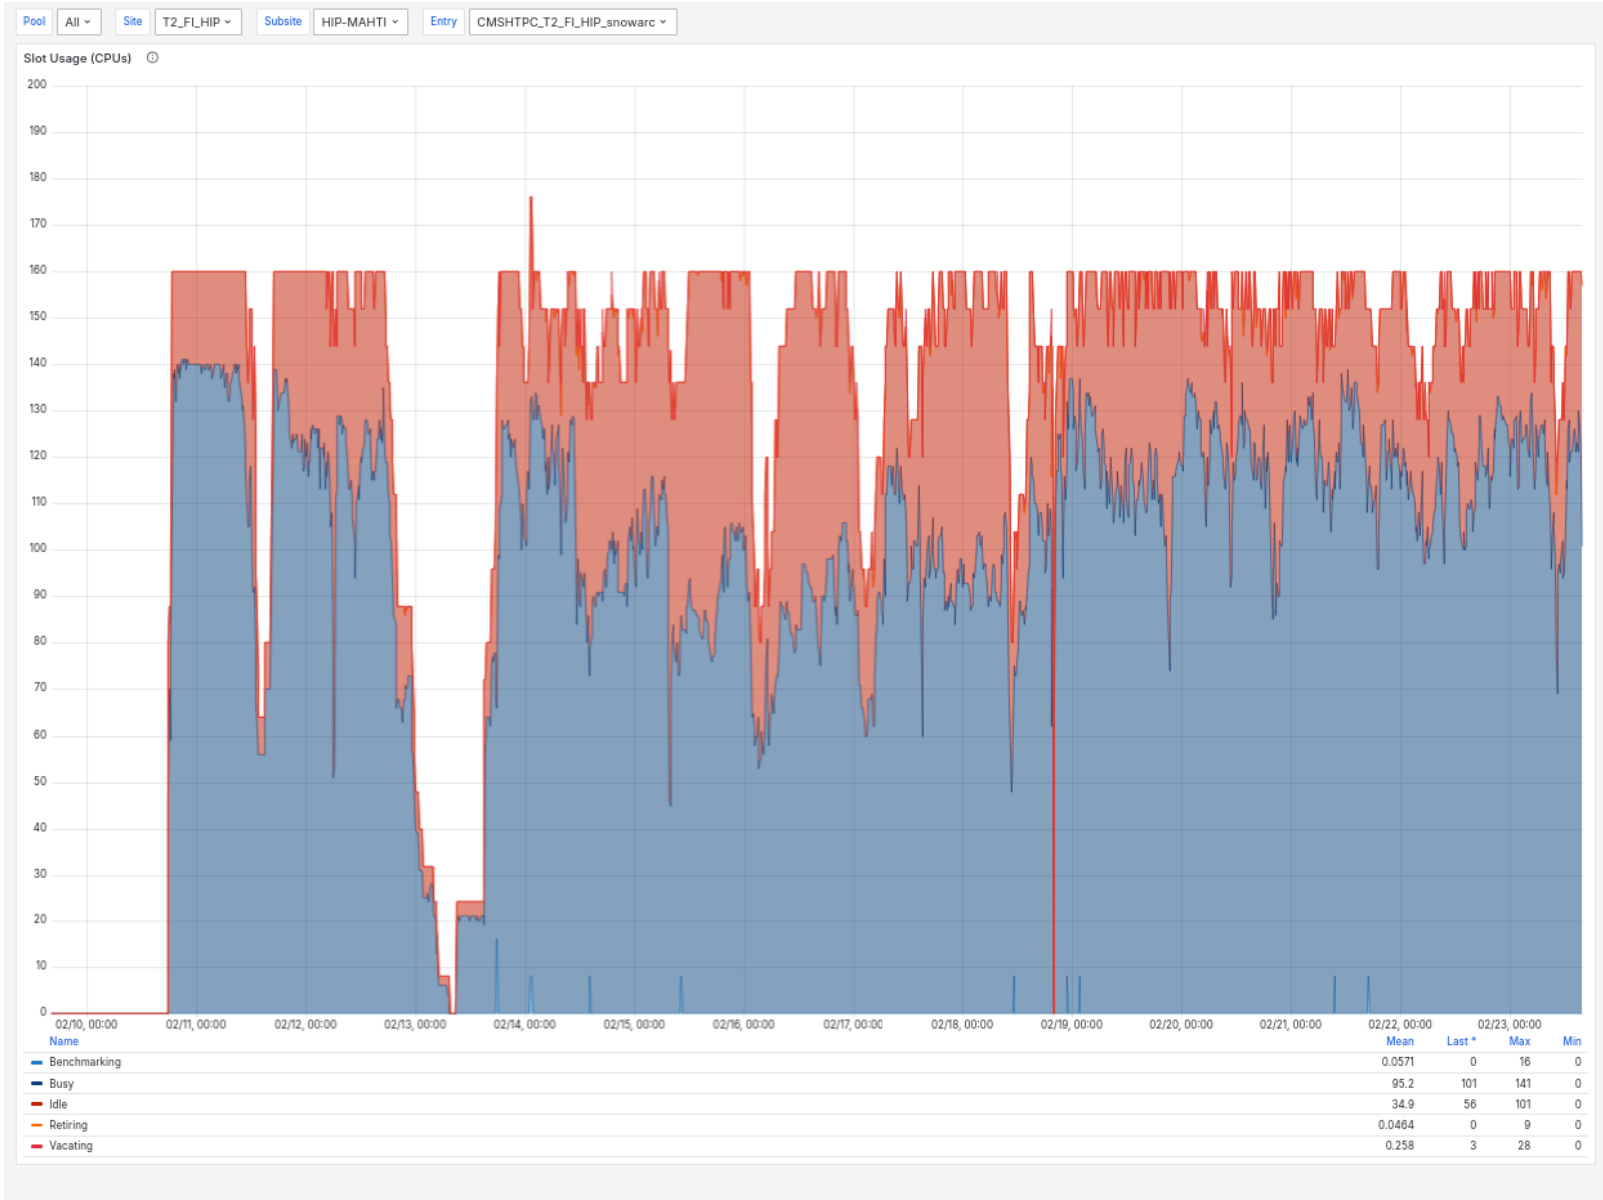

Figure: Timeline of CMS production pilots on snowarc since February 11th 2026.

Using fapptainer on HPCs

- The global CVMFS Alien cache on the Mahti Lustre was changed to a 50 GB job local alien cache on node NVMe storage
  - This solved failures of CMS jobs
- ATLAS & CMS are running production jobs on $\approx$ 160 cores each
- The system is generally stable, but scaling up will stress the (small) VM too much
- The VM could be upgraded or we could use several VMs and try to scale up with simulation jobs only
- Sometimes the sshfs connection goes down, which can fill up the local filesystem on the VM
- Grafana shows Lustre file system errors that could be intrinsic or due to sshfs problems
- After the upgrade to ARC 7.1.1 we lost the RTE configuration
- Memory and CPU usage information of jobs run with fapptainer are missing, so this needs to be fixed if possible
- M. Svatos has run ATLAS jobs on up to 21 kcores simultaneously on LUMI using fapptainer and HyperQueue without ARC data staging

Using fapptainer on HPCs

- We plan to go back to use also LUMI
- Other workflows could be tested
- A common ATLAS & CMS CHEP 2026 abstract was submitted
- There were 49 Roihu pilot usage applications in total and 28 were approved. Unfortunately our application was not approved.

Summary

- Development of fapptainer continues to enable HEP usage of HPCs like Mahti and LUMI
- The missing performance figures should be fixed for ATLAS (prmon) and CMS (HTCondor)
- Improving the sshfs stability would be needed
- LUMI usage is planned to be resumed
- Other workflows could be studied
- Performance and scaling studies would be interesting