Transparent extension of INFN-T1 with heterogeneous computing architectures

Stefano Dal Pra
INFN CNAF, v.le B. Pichat 6/2 - 40127 Bologna, IT

Daniele Spiga
INFN Sezione di Perugia, Via Alessandro Pascoli, 06123 Perugia, IT

Tommaso Boccali
INFN Sezione di Pisa, Largo B. Pontecorvo 3, 56127 Pisa, IT

Lorenzo Rinaldi
Bologna University and INFN Sezione di Bologna, v.le B. Pichat 6/2 - 40127 Bologna, IT

E-mail: `stefano.dalpra@cnaf.infn.it`

Abstract. The INFN-CNAF Tier-1 is engaged for years in a continuous effort to integrate its computing centre with more typologies of computing resources. In particular, the challenge of providing opportunistic access to nonstandard CPU architectures, such as PowerPC, ARM or hardware accelerators (GPUs) has been actively exploited. In this work, we describe a solution to transparently integrate access to ppc64 and GPUs. This solution has been tested to transparently extend the INFN-T1 Grid computing centre with Power9 based machines and V100 GPUs from the Marconi 100 HPC cluster managed by CINECA, and eventually used for the next CINECA machine, Leonardo. We also discuss further possible improvements and how this will meet requirements and future plans for the new tecnopolo centre, where the CNAF Tier-1 will be hosted soon.

1. Introduction

The INFN-CNAF Tier-1 [1] computing centre has been serving as one of the main resources for the High Energy Physics (HEP) community. As the scale and complexity of scientific computing continue to increase, the Tier-1 has continuously investigated ways to incorporate various computing resources to meet the growing demand. Over the years, the centre has made significant progress at integrating its computing facility with diverse types of computing devices, ranging from opportunistic resources instantiated on Cloud Instances [2], or nodes made available on high-performance computing (HPC)systems [3] and data centre extensions [4]. These works highlight the innovative approaches that have been taken to enhance the centre's capabilities, ultimately providing researchers with the tools they need to conduct cutting-edge research in High Energy Physics.

The HEP computing infrastructure relies primarily on the x86_64 architecture. However, there have been recent collaborations between the World LHC Computing Grid (WLCG) [5] community and HPC centres that offer computing power through PowerPC architecture based machines. This development has spurred an investigation into methods for making these resources accessible through the standard procedures adopted in the WLCG community.

This requires a dual effort: on one hand, experiments must modify and verify their software stack to function with an alternative computing architecture. On the other hand, Grid resource providers, such as the INFN-CNAF Tier-1, must devise solutions to make these new resources as easily accessible as possible for the community. In this work we specifically focus on the latter: we address the challenge of enabling new external and heterogeneous resources to join the HTCondor batch pool managed by the INFN-CNAF Tier-1 and define convenient methods for Grid submitters to use these resources.

2. Integrating additional resources

The integration of heterogeneous resources requires careful consideration of several factors. When integrating pledged resources with opportunistic ones, attention must be paid to details such as the usage policy, as the quality of service may no longer be homogeneous through the pool. Hardware differences must also be taken into account, such as the possibility of onboard GPUs on some nodes. Network configurations, including bandwidths, ACLs, and routing, may vary between different resources. Storage may not be uniformly available, or might be lacking dedicated capabilities to rely on. Additionally, run time environments may be less manageable compared to pledged nodes. Finally, fair-share policies must be adapted to ensure that opportunistic resources do not unduly impact the balance of pledged resources. We now consider a specific use case to clarify these considerations.

2.1. Use case: integrating opportunistic resources from the Marconi100 HPC cluster at CINECA

In order to integrate with the Marconi100 HPC cluster [6] (M100) at CINECA [7], a IBM Power9 architecture with NVIDIA Volta GPUs , we developed a custom Python script that submits jobs to the Slurm batch system managed by Cineca from a Submit Node, running as a regular unprivileged user. These jobs are queued to run on the opportunistic M100 machines. When started, the Slurm job launches a Singularity container that, in turn, instantiates an HTCondor STARTD daemon. This STARTD is equipped with proper IDTOKEN credentials, enabling it to join the HTCondor pool at CNAF. HTCondor-CEs at INFN-CNAF Tier-1 are then properly configured so that jobs submitted at CNAF and requiring or accepting to run on M100 resources (which have PowerPC architecture) are properly routed and queued to HTCondor for execution in the M100 nodes.

2.1.1. Requirements There are a few requirements here to consider. Firstly, jobs usually require access to a CVMFS [8] filesystem that provides the needed libraries and software dependencies. Fortunately, CINECA already supports this technology, making it a straightforward requirement to fulfill. Secondly, the remote STARTD must be able to connect to the HTCondor [9] Collector of the INFN-CNAF pool via the 9618 tcp port of the Connection Control Broker (CCB). This requires CINECA to allow outbound connectivity from the HPC nodes towards CNAF networks and adapt CNAF firewall rules to accept incoming connections from these machines. It's worth noting that the CCB may potentially act as a bottleneck, but traffic through this component is minimal, as it's only used for cluster management and no storage data traffic is expected to flow through the CCB. In the case when not even outbound connectivity is available, a "network path" from the computing node to the Collector can be using tsocks. This possibility has been exploited as described in [10]

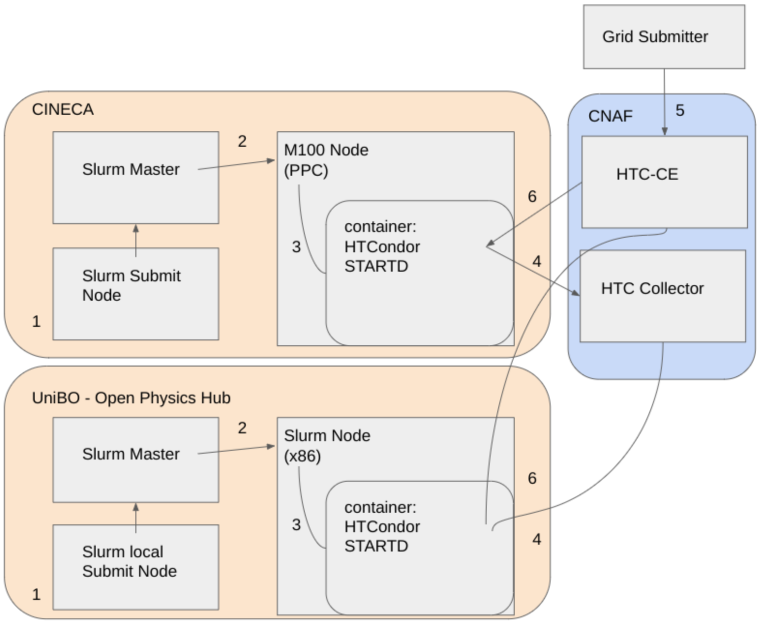

Figure 1: Two groups of external Computing Node resources join the CNAF HTCondor pool: PowerPC machines from the CINECA M100 cluster, and traditional x86_64 Virtual machines from the UniBO - Open Physics Hub Cloud Instance. The Grid submitter specifies the group of choice (see Listing 1) to match an entry of the `JobRouter` in the HTC-CE (see Listing 2). The external Compute Nodes are instantiated as jobs scheduled by the remote batch system. At start, these run a singularity container which in turn executes an HTCondor STARTD, which finally joins the CNAF pool and becomes ready to execute jobs submitted at CNAF. More external groups of resources can be added in the same way.

2.2. Quality of Service

Finally, more differences come from the expected QoS: even though these resources appear as PowerPC machines, the maximum run time is bounded to that of the underlying Slurm job, which is in this case capped at 24 hours. Submitters need to take this constraint into account and submit suitable jobs for these resources. To mitigate the impact of this time limit, it has been decided to provide these machines as whole nodes, meaning that once a job is started, it can utilize all available CPUs and memory. This approach has already been utilized successfully in previous works, such as in [4].

3. Integration with the HTCondor pool at CNAF

A test job submitted to a HTCondor-CE at CNAF, matches with a JobRouter entry defined for the target group of external resources (M100 in this case). This assumes that a custom classad "WantRoute" specified the desired resource type in the job submit parameters. Listing 1 below represents a simple example submit file to request M100 resources, and Listing 2 shows how the JobRouter is configured in the HTCondor-CE to accordingly match these requests.

This has been tested to transparently extend the INFN-T1 Grid computing centre with Power9 based machines and V100 GPUs from the Marconi 100 HPC cluster managed by CINECA Standard resources from the x86_64 HPC farm of Open Physics Hub (OPH, [11]) project managed by Physics and Astronomy Department of the Bologna University.

The early integration of the system with the CMS distributed system has been tested, including the embedding of manual pilots from CMS in the setup that we developed to verify the entire chain described above.

The final integration will be discussed with the CMS Submission Infrastructure, specifically regarding the submission of pilots through CMS GlideinWMS [12].

```
[sdalpra@ui-htc ~]$ cat ce_m100_p9.sub
universe = vanilla
use_scitokens = true
+Owner = undefined
+WantRoute = "cms_m100"
executable = p308/htcp308.p9
output = htcp308_$(ClusterId).$(ProcId).out
error = htcp308_$(ClusterId).$(ProcId).err
log = htcp308_$(ClusterId).$(ProcId).log
arguments = "0␣0␣1␣2001"
queue 1
```

Listing 1: Example submit file for M100 resources. The `+WantRoute` custom attribute must be agreed with the administrators of the HTCondor pool and declares the desired group of computing resources. Its value must match with the configuration of the `JobRouter` (see Listing 2)

The ATLAS experiment workflow has been fully tested for x86_64 on the OPH farm, on several nodes managed by a Slurm scheduler and with a custom CVMFS configured with CVMFS-exec. Since no official ATLAS software build is available for PowerPC, a partial test on M100 has successfully been performed by running a ROOT-based physics analisys workflow (in this specific example we run the same analysis code used in [13]). Remote data were accessed via XRootD, through a dedicated XRootD-proxy server within the CNAF network, in forwarding mode.

After early successful tests with CMS and ATLAS, the PowerPC CPU architecture has been made available to INFN-CNAF Tier-1 user communities having suitable payloads. Moreover, further early tests with other kind of resources, such as ARM machines and other traditional x86_64 (from Intel or AMD) machines coming from Cloud instances have confirmed the idea that this model can be used to seamlessly integrate the HTCondor pool at CNAF with several heterogeneous groups of external resources. Fig. 1 represents the first two of these, with M100 being of course the more interesting in terms of potentially available resources and because of the different CPU architecture and availability of V100 GPUs.

This model proved to enable transparent extension of the CNAF-INFN Tier-1 pool to external computing resources of different typologies (x86, ppc, arm, GPUs; all of these have been tested).

```
JOB_ROUTER_ROUTE_cms_m100 @=jrt
REQUIREMENTS (WantRoute =?= "cms_m100") &&\
  (AuthTokenIssuer =?= "https://cms-auth.web.cern.ch/" &&\
   AuthTokenSubject =?= "78f275d5-bb1a-4b2d-9956-f82316a8482e")
 UNIVERSE VANILLA
 SET Requirements (TARGET.Arch =?= "ppc64le")
@jrt

JOB_ROUTER_ROUTE_atlas_m100 @=jrt
REQUIREMENTS (WantRoute =?= "atlas_unibo") &&\
 (AuthTokenIssuer =?= "https://atlas-auth.web.cern.ch/" &&\
  AuthTokenSubject =?= "ccff569b-bda5-45ac-91bb-44c8e198a385")
 UNIVERSE VANILLA
 SET Requirements (TARGET.Arch =?= "x86\_64")
@jrt
```

Listing 2: Example JobRouter rules for M100 and Unibo Cloud. The value for the `+WantRoute` custom attribute specified by the Grid submitter (see Listing 1) determines which rule is applied.

4. Future work and upcoming objectives

The main objective of the described work is to consolidate the model in order to being production quality for the exploitation of opportunistic resources at Leonardo machine [14] at CINECA. From the technical perspective this will require a further development in order to dynamically submit slurm jobs de on the requirements of the HTCondor jobs queued at CNAF. We plan also to apply the very same approach to all the Cloud/HPC machines where we will be allowed to perform integration tests, and eventually production jobs: we consider this as a viable solution for scientific communities already supported at Tier1 CNAF to bring their own resources, i.e. resources acquired via grants (a.g., PRACE grants or commercial cloud grants) or external funding (a.g. from collaboration with industry), such as worker nodes from external batch systems, for opportunistic usage, while keeping their usual submission model unchanged (no need to explicitly add a target for job submission). Another extension under discussion at the time of writing is the usage of the described model to dinamically instantiate HTCondor worker nodes on the cloud in order for a given scientific community to absorb temporary peak of resource request. An important security aspect for the proposed model is related to the token distribution: every external STARTD instance needs a valid IDTOKEN in order to join the CNAF pool, and this must be kept confidential enough to not become available to untrusted users. We currently adopt dedicated IDTOKENS for each group of external resources and work is needed to enforce and improve the reliability of the IDTOKENS used for this purpose.

5. Conclusion

In this document, we have discussed the extension of the CNAF-INFN Tier-1 computing farm with additional opportunistic resources with PowerPC architecture from the Marconi100 HPC cluster. We have presented the method used to integrate these resources into the existing HTCondor pool at CNAF, also discussing the requirements and limitations due to differences in terms of available network and storage facilities and maximum runtime. Finally, we have discussed ongoing work to improve the reliability of the IDTOKENS used for token distribution and the expansion of this model to include other kinds of resources, such as ARM and traditional x86_64 machines from Cloud instances.

6. Acknowledgements
This work is (partially) supported by ICSC – Centro Nazionale di Ricerca in High Performance Computing, Big Data and Quantum Computing, funded by European Union – NextGenerationEU.

References
[1] Dell'Agnello L, Boccali T, Cesini D, Chiarelli L, Chierici A, Dal Pra S, De Girolamo D, Falabella A, Fattibene E, Maron G *et al.* 2019 INFN tier–1: a distributed site *EPJ Web of Conferences* vol 214 (EDP Sciences) p 08002
[2] Dal Pra S, Ciaschini V, Dell'Agnello L, Chierici A, De Girolamo D, Sapunenko V, Boccali T and Italiano A 2016 Elastic CNAF data center extension via opportunistic resources *International Symposium on Grids and Clouds* vol 13
[3] Boccali T, Sapunenko V, Cesini D, Valassi A, Dal Pra S, Stagni F, Spiga D, Bozzi C, Zani S, Doria A *et al.* 2021 *PoS* 003
[4] Boccali T, Dal Pra S, Spiga D, Ciangottini D, Zani S, Bozzi C, De Salvo A, Valassi A, Noferini F, dell'Agnello L *et al.* 2020 Extension of the INFN tier-1 on a HPC system *EPJ Web of Conferences* vol 245 (EDP Sciences) p 09009
[5] Shiers J 2007 *Computer physics communications* **177** 219–223
[6] 2022 Marconi 100 , CINECA - URL `https://www.hpc.cineca.it/hardware/marconi100`
[7] 2022 CINECA - URL `https://www.cineca.it/en`
[8] Buncic P, Sanchez C A, Blomer J, Franco L, Harutyunian A, Mato P and Yao Y 2010 *Journal of Physics: Conference Series* **219** 042003 URL `https://dx.doi.org/10.1088/1742-6596/219/4/042003`
[9] Thain D, Tannenbaum T and Livny M 2005 *Concurrency - Practice and Experience* **17** 323–356
[10] Mariotti M, Spiga D and Boccali T 2021 *ISGC 2021* 22–26
[11] 2022 Open physics hub - department of physics and astronomy - URL `https://site.unibo.it/openphysicshub/en`
[12] Balcas J *et al.* 2015 *J. Phys. Conf. Ser.* **664** 062031
[13] The ATLAS Collaboration 2021 *Eur.Phys.J.C* **81** 218
[14] 2022 Leonardo HPC - CINECA - URL `https://leonardo-supercomputer.cineca.eu`