ARC-CE based HPC job submission

Michal Svatoš

ATLAS Software & Computing Week

3.-7.2.2025

Categories

ARC-CE based HPCs can be grouped depending on how restricted they are in comparison to standard grid site

1. Vega and other ND HPCs: Push jobs via ARC CE with data transfer, CVMFS available, all workloads
2. Karolina+Barbora: ssh access to batch and FS, stage-in/out from grid storage, wholenode scheduling, no CVMFS but cvmfsexec can run there, multicore production workloads
3. LUMI: ssh access to batch and FS, stage-in/out from grid storage, wholenode/subnode scheduling, no CVMFS and no cvmfsexec (no user namespaces), fat container simulation only, SLES compatible with centos8
4. MN5, SuperMUC: no outbound IP $\rightarrow$ push with data via ARC CE, no CVMF/cvmfsexec, fat container simulation only

Introduction

- Established as a first EuroHPC in Q1 2021
- In production until Q1 2027
- specs
  - 130k cores (260k HT slots) AMD 7H12
  - 240 A100 40GB GPUs
  - 25PB Ceph (HDD/NVMe)
  - 1PB Lustre (NVMe)
  - 300 Gb/s WAN
  - 3M HS06 on CPU

Environment

- ARC-CE + squid on site
- push jobs (ARC-CE orchestrates data transfers from/to dCache at NDGF)
- singularityPro available
- OS: RHEL8
- CVMFS available
- running ATLAS all workloads

Submission system

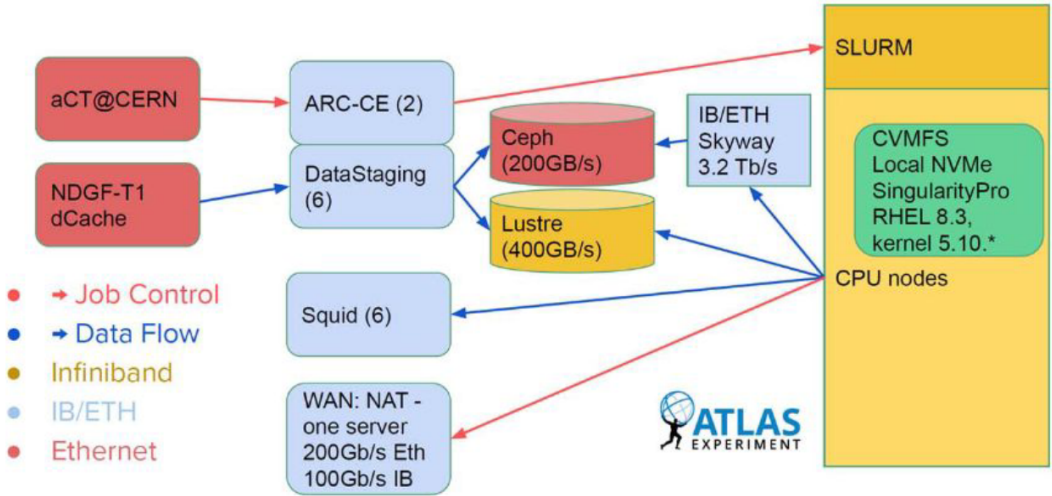

Introduction

- The ATLAS distributed computing is opportunistically using resources of the Czech national HPC center IT4Innovations in Ostrava through Czech Tier2 praguelcg2 in Prague:
- Barbora (2019-present)
  - CPU nodes: 192 WN with 36 cores and 192GB of RAM
  - in production in ATLAS: since 2020
- Karolina (2021-present)
  - CPU nodes: 720 WN with 128 cores and 256GB of RAM
  - in production in ATLAS: since 2021

Environment

- OS: RHEL8/Rocky8
- user namespaces enabled
- outgoing connectivity restricted to few ports (ssh, rsync, http(s))
  - firewall opened to praguelcg2 storage, CERN machines reachable via http proxy (praguelcg2's squid)
- whole-node scheduling (if needed, worker nodes are partitioned using HyperQueue)
- CVMFS not installed, software available locally+atlas-cvmfsexec
- local apptainer installation available
- allocations are not pre-emptable, opportunistic usage is pre-emptable

Usage

- runs multicore production workloads (mostly simulation)

CHEP paper in CDS

Submission system

- the ARC-CE shares storage space with the Lustre via sshfs connection through a login node and communicates with the batch system via ssh connection (through the same login node)
- when the ARC-CE receives a job, it translates the job description into a script that can be run in the batch system, puts necessary files into a folder within sshfs shared area and submits the job via ssh connection to the HyperQueue server running on a login node
- the HyperQueue server collects jobs and when there are enough of them, it submits jobs into the batch system

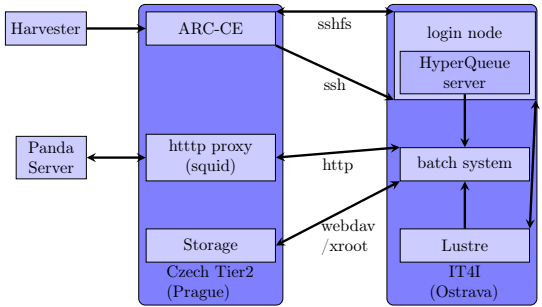

- when the batch job starts, HyperQueue jobs start in it (in sufficient numbers to fill the worker node - if available)
- in each HyperQueue job, pilot wrapper starts, launching the pilot
- pilot contacts panda server through http proxy (Czech Tier2 squid) to receive a payload (as there are only few open ports)
- when it receives the payload, it gets input file from the Czech Tier2 storage via xroot or webdav
- then it starts the calculation
- when the payload finishes, it sends outputs to the Czech Tier2 storage via xroot or webdav
- when this is finished, pilot will request another payload

Introduction

- **LUMI (Large Unified Modern Infrastructure)** is a pan-European pre-exascale supercomputer located in CSC's data center in Kajaani, Finland.
- The supercomputer is hosted by the **LUMI consortium** including eleven European countries (Finland, Belgium, the Czech Republic, Denmark, Estonia, Iceland, the Netherlands, Norway, Poland, Sweden, and Switzerland).
- Half of the LUMI resources belong to the EuroHPC Joint Undertaking, and the other half of the resources belong to the LUMI consortium countries. The shares for each of the countries are allocated according to local considerations and policies, i.e. LUMI is seen and handled as an extension to national resources. Our allocation was granted from the Czech share.
- testing done in 2024

Environment

- we had a small allocation for testing
- difficult environment
  - OS is SUSE
    - grid middleware supported for RHEL only
    - native python3 version is 3.6
  - user namespaces will NEVER be enabled
    - meaning we cannot use cvmfsexec, nested containers, etc.
  - we were able to run because the environment was compatible with centos8 (thanks to Asoka De Silva for all his help)
- outgoing connectivity allowed
- CVMFS not installed, we were running in fat containers
- local apptainer installation not available but singularity is there
- no pre-emption

Environment

- WNs in CPU partition
  - *standard* partition
    - 256 cores (2x64 physical + HT)
    - 256 GB of memory but only 224 GB available for job, i.e. there is 0.875 GB of RAM per core
    - max resources: 512 nodes
    - only whole node submission allowed
  - *small* partition
    - 256 cores (2x64 physical + HT)
    - 512 or 1024 GB of memory
    - max resources: 4 nodes
    - jobs taking only part of the WN can be submitted

Submission system

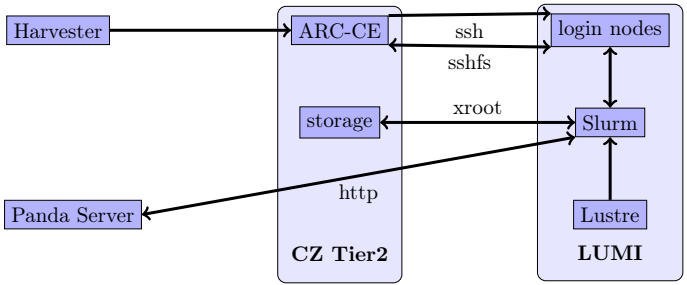

- the ARC-CE (located at CZ Tier2) receives a pilot job, translates the job description into script that can be run in the Slurm batch system, puts necessary files into a folder shared with the HPC via sshfs, and submits the job via ssh connection to a login node
- when the batch job starts, pilot contacts panda server to receive payload job
- if it receives payload job, it gets input file from CZ Tier2 storage, starts the calculation (in software container), and sends outputs to CZ Tier2 storage when the payload finishes
- if the pilot can expect that another payload would finish, it requests it

Usage

- pilot needs to run on bare metal
- stage-in/out runs in EL9 container
- payload runs in fat containers → runs simulation only
- running 8-core jobs on the *small* partition

MareNostrum4/5: Introduction

- MareNostrum4
  - 11.5 Petaflops (166k CPU cores)
  - 15 PB GPFS shared disk storage
- MareNostrum5
  - 730k CPU cores
  - 250 PB GPFS disk storage

MareNostrum4/5: Environment

- no outbound connectivity
- push jobs (ARC-CE orchestrates data transfers)
- running only fat container simulation

MareNostrum4/5: Submission system

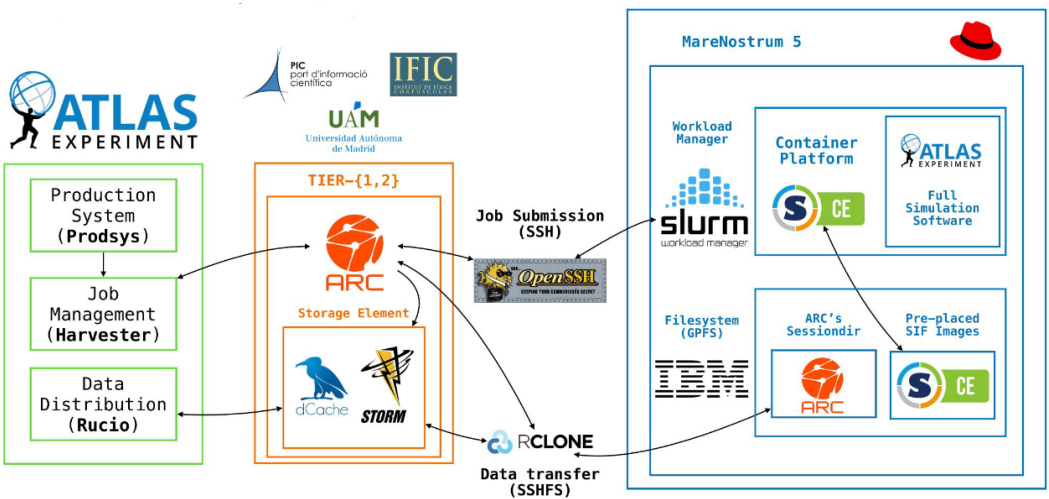

MUC

- no outboundIP so must be push jobs with data input and output handled by the ARC CE
- preempt-only. Wholenode G4 typically <1h, so preempt level acceptable.
- SuperMUC requires now 2FA which complicated ARC CE via ssh approach
- keep ssh channel open using ControlMaster(aka multiplexing)

Summary

- HPCs are providing significant amount of resources to ATLAS