Infrastructure Monitoring and Dynamic Job Matching for ALICE Grid Using Site Sonar

Kalana Wijethunga
CERN, University of Moratuwa

Alice T-1/T-2 Workshop | 16-18 April | Seoul

Overview

- Optimized job matching in Grid middleware domain
- Currently, we can match jobs based on few attributes only
  - CPU Cores
  - Memory
  - Disk Space
- What if we need more attributes?
  - Specific GPU model
  - Specific CPU Architecture and generation
  - Software Version
- We developed an enhanced approach to match jobs to the most suitable worker node

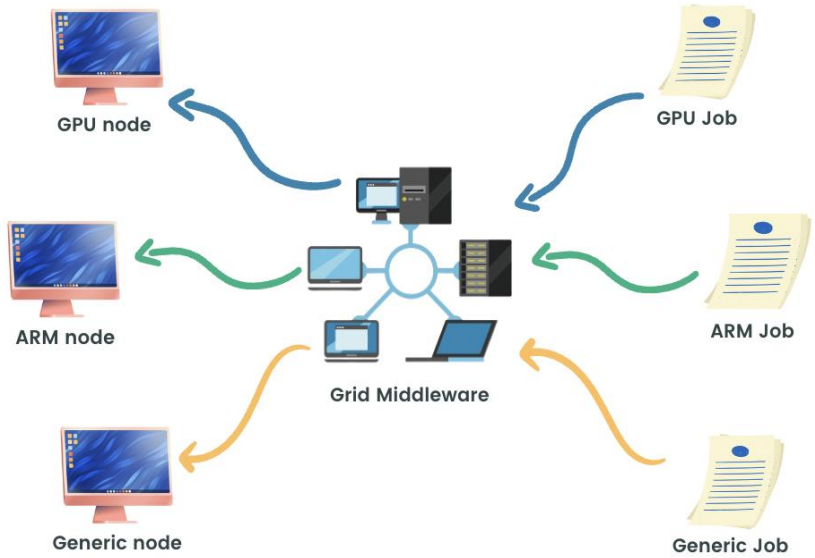

Why?

- User requirements are getting more diverse and granular
  - New workloads optimized for ARM (aarch64)
  - Coprocessing on different nodes with GPUs
- Resources with special infrastructures are limited
  - Only a few Grid sites have GPUs (<5 in ALICE Grid)
  - Only few sites are ARM based etc. (only 1 in ALICE Grid)
- Available resources are not optimally used
  - Above sites run all jobs at the moment
  - Special infrastructure might be kept idle

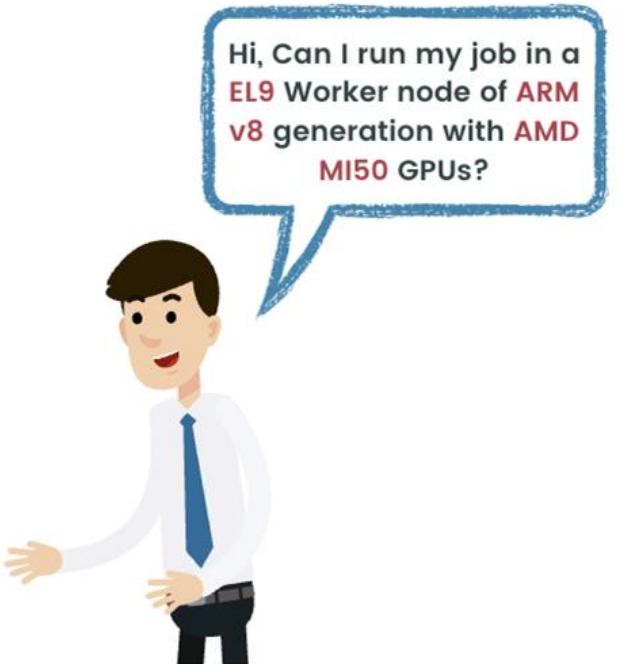

Plan

- Using Site Sonar - Infrastructure monitoring tool for ALICE Grid
- Can monitor any infrastructure attribute of any worker node in the Grid
- Provide a complete picture of available infrastructure
    - Centrally stored and visualized with ELK

220+	36	11	118
Parameters/node	Probes	Operating Systems	CPU Models

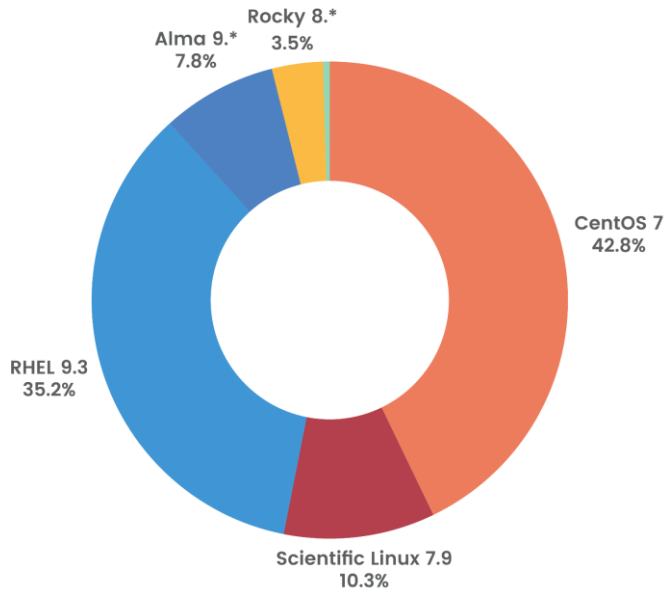

Operating system distribution of ALICE Computing Grid (2024/04/06)

Motivation

It is important to

- Ensure the Grid sites are compatible with the software versions required by payloads
- Ensure correct configuration of the individual site nodes
- Identify and isolate sites and individual nodes with abnormal configuration and behaviour
- Alert the system administrators and provide debugging information
- Goal: Have a full picture of the current status of the Grid and optimize for efficient use

Issues with existing systems

- Data pull model
  - A central service running monitoring probes on individual nodes is not scalable, resource intensive and presents a single point of failure
- Agent installation
  - Most infrastructure monitoring systems require installing custom software agents on Grid sites which is not favoured by Site administrators
- Low flexibility
  - Monitoring systems do not allow collection of unstructured data, hence it does not allow collecting arbitrary data
  - No post data filtering
- Low extensibility
  - Multiple steps and releases needed to add more metrics

Solution

A new Grid Infrastructure Monitoring Tool called **“Site Sonar”** that monitors the infrastructure attributes of individual worker nodes in the Grid, consisting of a :

a. Data Collection framework that is
    i. Flexible to change data structures on demand
    ii. Easy to add new data collection probes easily
    iii. Improving Job Matching functionality using collected data
b. Data Visualization framework that
    i. Allows post data filtering
    ii. Provides no-code visualizations

Site Sonar Design

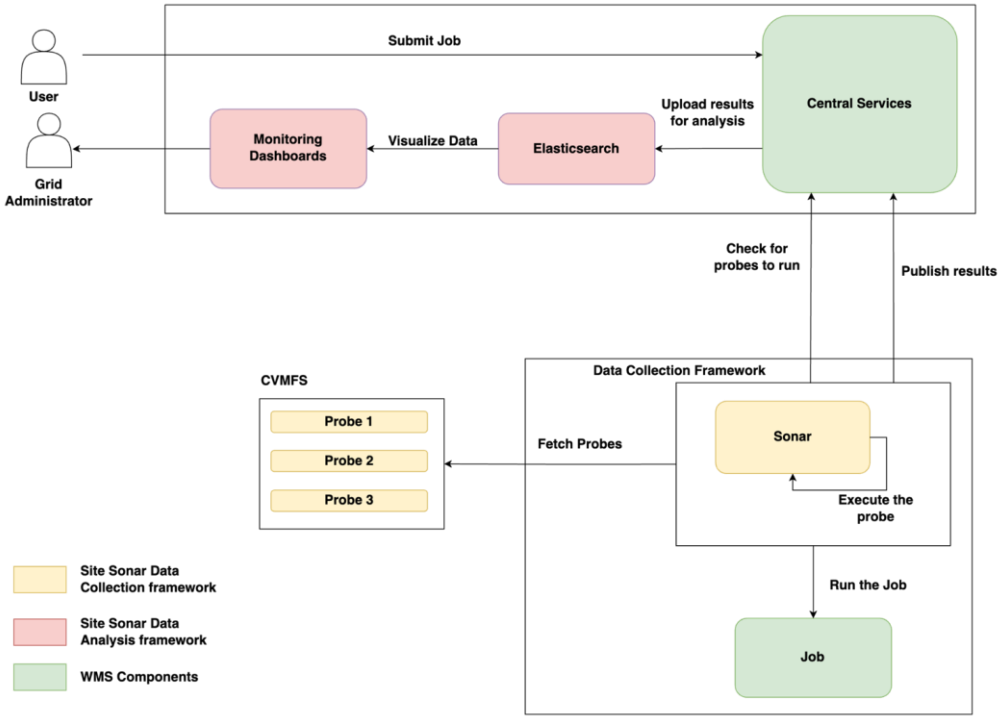

Key features

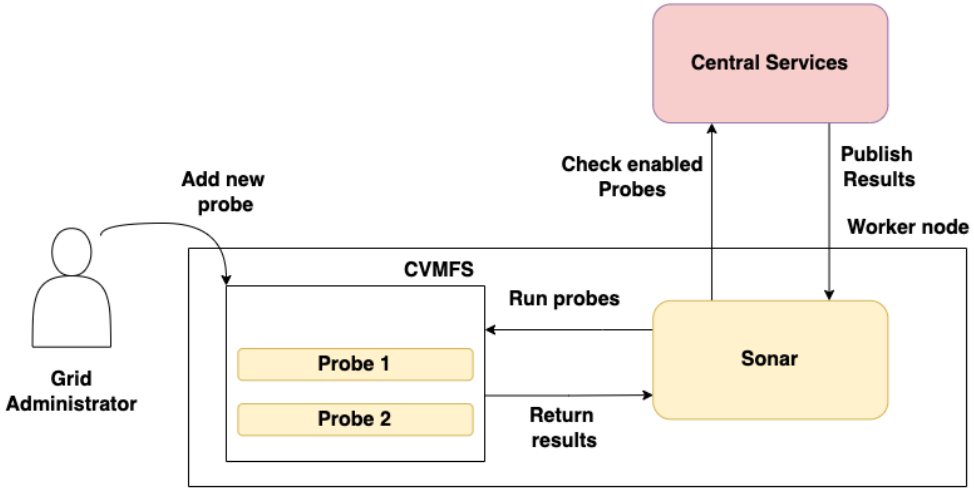


```
{ "addr": "188.184.162.27",
  "hostname": "b7s11p0950.cern.ch",
  "last_updated": 1669942802,
  "ce_name": "CERN",
  "uname": {
      "UNAME": "Linux b7s11p0950.cern.ch"
  },
  "cpu_info": {
      "CPU_cpu_cores": 16,
      "CPU_model_name": "Intel(R) Xeon(R) Silver"
  },
  "home": {
      "HOME": "/pool/condor/dir_156198"
  },
  "os": {
      "OS_PRETTY_NAME": "CentOS Linux 7 (Core)"
  }}
```


Easily Extensible - Can add or remove new tests to collect metrics without any code changes

Highly Flexible - Can collect any data from a node, easily change data structure and type of data

Grid Monitoring system

Site Sonar - ALICE Grid Infrastructure Monitoring Tool

Period: 1 week | Choose Site: % | Total nodes: 10225

Add filter + | Apply

os: OS_PRETTY_NAME = CentOS Linux 7 (Core) (4571)

Site Name	Matching	Not Matching	Matching Percentge	Total
Birmingham	0	36	0.0%	36
Bratislava	49	0	100.0%	49
Catania	5	0	100.0%	
CBPF	61	0	100.0%	
CCIN2P3	549	0	100.0%	
CERN	589	3589	14.0%	
Cibinong	0	2	0.0%	
Clermont	104	0	100.0%	

Available at :
https://alimonitor.cern.ch/sitesonar

Results for the site CERN

Site Summary : CERN

Test Name : os
Test Key : OS_PRETTY_NAME
Test Value : CentOS Linux 7 (Core)
Matching Node Count : 589
Not Matching Node Count : 3589

Hostname	Value	Last Updated
alict-wn-030.ct.infn.it	CentOS Linux 7 (Core)	2024/04/11 at 5:36 PM
b7s02p9848.cern.ch	CentOS Linux 7 (Core)	2024/04/11 at 3:48 PM
b7s04p3712.cern.ch	CentOS Linux 7 (Core)	2024/04/11 at 6:39 PM
b7s03p3850.cern.ch	CentOS Linux 7 (Core)	2024/04/12 at 2:53 AM
b7s02p4264.cern.ch	CentOS Linux 7 (Core)	2024/04/12 at 9:08 AM
b7s01p5227.cern.ch	CentOS Linux 7 (Core)	2024/04/12 at 1:21 AM
b7s02p6768.cern.ch	CentOS Linux 7 (Core)	2024/04/11 at 2:03 PM
b7s03p3399.cern.ch	CentOS Linux 7 (Core)	2024/04/11 at 10:24 PM

Grid Monitoring system

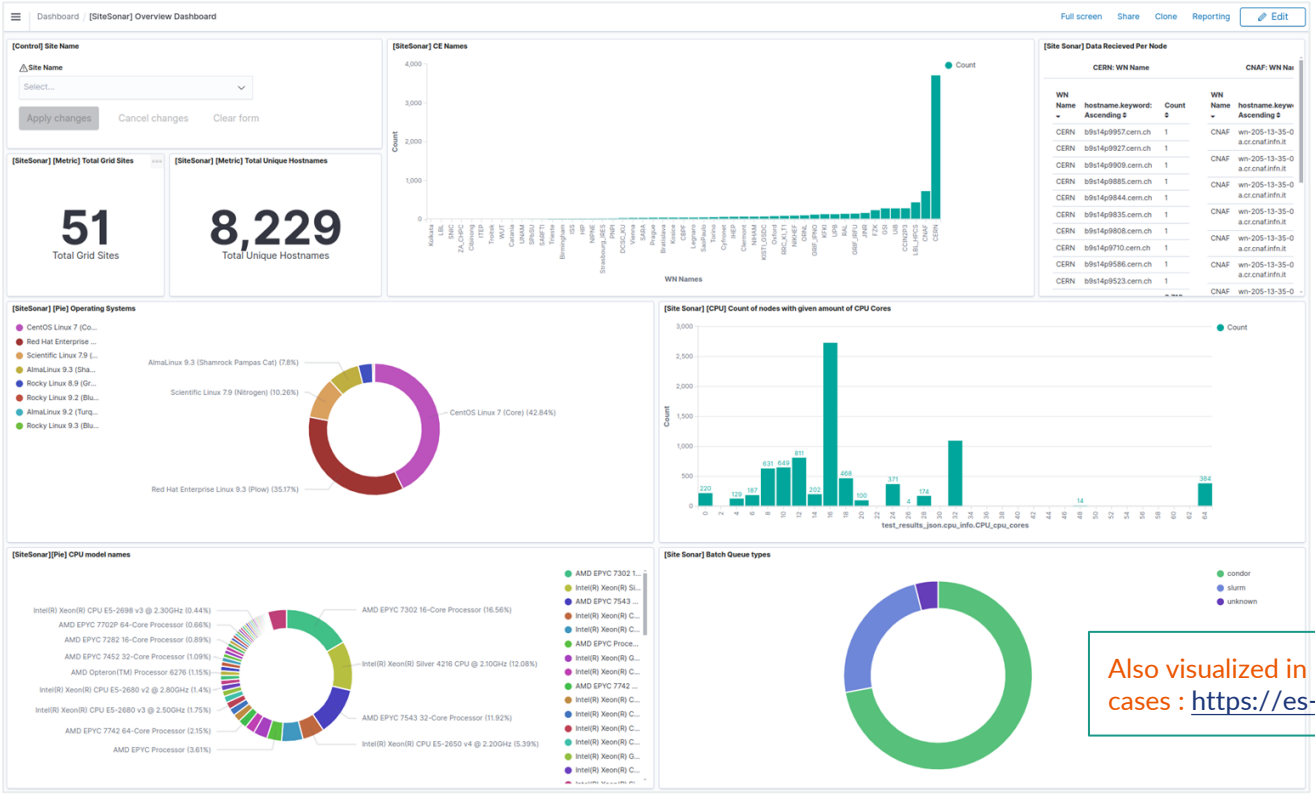

Time Period :
06/04/2024 - 07/04/2024

Also visualized in Kibana for complex use cases : https://es-alicecs.cern.ch/

Findings

- Sites running CentOS 6
  - CentOS 6 reached EoL and CentOS 7+ is necessary to have singularity support by default
  - Site Sonar was used to identify the Sites still running CentOS 6
  - Outcome: Whole Grid was upgraded to CentOS 7+ after notifying the admins
- Reusing hostnames in nodes
  - A unique identifier is used to identify nodes in the Grid
  - Site Sonar identified that this is not properly configured in some sites leading to monitoring errors
  - Outcome: All sites were correctly configured to have proper identifiers
- Intermittent Singularity failures
  - Singularity was intermittently failing on CentOS 7+ sites
  - Site Sonar was used as a debugging tool to identify configuration errors that cause this
  - Outcome: Correct Singularity support in the Grid was observed after corrective actions

Sites Moving away from CentOS

Operating System	coverage%
CentOS Linux 7.9.2009 (Core)	83.18%
Scientific Linux 7.9 (Nitrogen)	10.63%
AlmaLinux 8.7	6.19%

TABLE 1 : Operating system distribution of ALICE Grid as of 2023-03-01

Operating System	coverage%
CentOS Linux 7.9.2009 (Core)	42.84%
RedHat Enterprise Linux 9.3(Plow)	35.17%
Scientific Linux 7.9 (Nitrogen)	10.26%

TABLE 2 : Operating system distribution of ALICE Grid as of 2024-04-06

Associated projects

- CPU Oversubscription
  - Goal : Devise a dynamic scheduling strategy that identifies the idle resources in the nodes that are executing jobs and oversubscribe the node to run extra jobs
  - Usage: requires up-to-date information about the CPU cores, memory and other hardware information collected through Site Sonar
- CPU/ Memory limits crawling (see Marta's presentation)
  - Goal: Anticipate OOM killings derived from exceeding resource limits allocated to jobs
  - Usage: Collect the resource limits that are imposed to our allocated CPU and memory resources
- Profiling and Duration Estimation of the MonteCarlo Jobs
  - Goal : Improve the scheduling algorithms for the CPU-intensive jobs (MonteCarlo simulations) based on their run history
  - Usage: Combines data extracted from job traces (time spent by the job, site and hostname) with information extracted from SiteSonar such as CPU model and CPU flags

Introducing Dynamic Job Matching

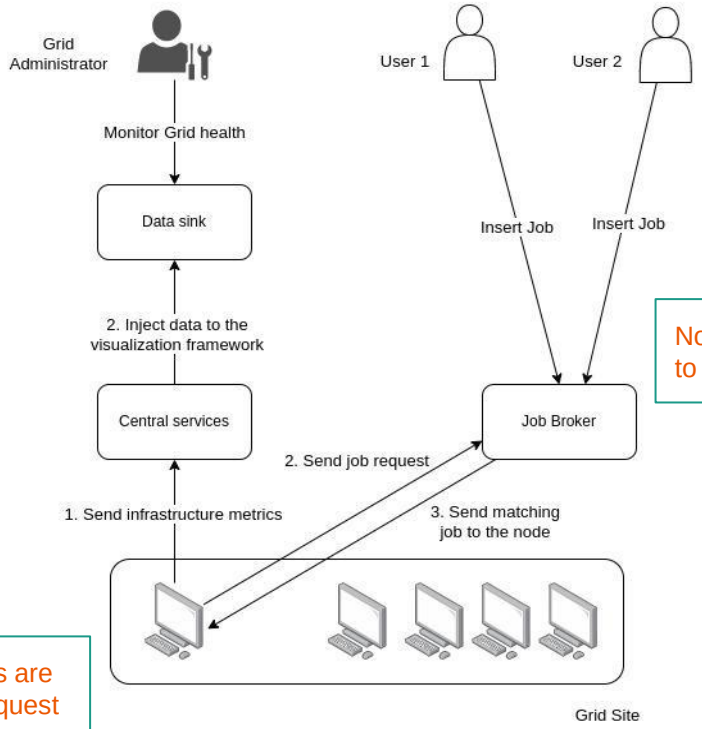

Node infrastructure metrics are used to in the job matching process

Collected infrastructure metrics are injected to the job matching request

Result

- Site Sonar integrated into the ALICE Grid middleware JAliEn
- Built-in monitoring to node infrastructure with JAliEn
- Job matching possible based on any infrastructure attribute of the worker node
- No code update necessary for JAliEn for this

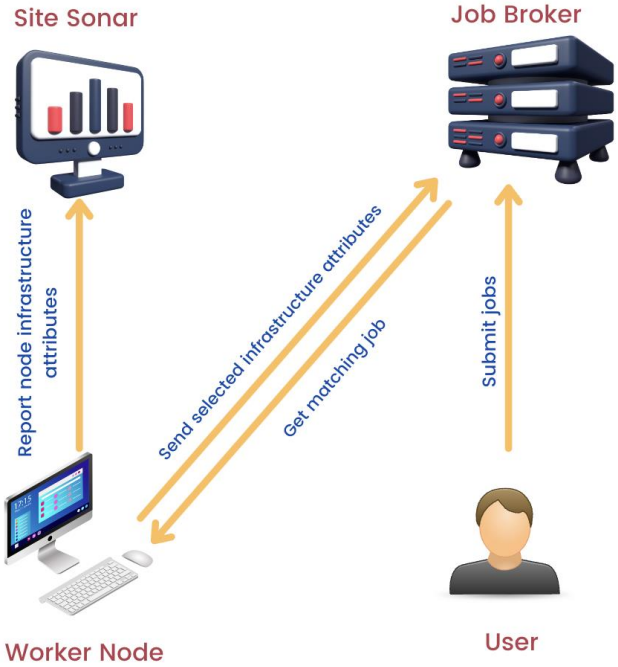

Conclusion

- Work focused on
  - Developing a flexible and extensible Grid infrastructure monitoring framework
  - Introducing an infrastructure aware **dynamic** job matching architecture
- Limitations in existing systems
  - Not extensible - Need code updates or roll outs to add new metrics
  - Not flexible - Cannot collect unplanned data
  - Not easy to visualize - Need to code new monitoring plots to do additional analysis
  - No association with job matching - Cannot use the collected data in job matching process
- Developed Site Sonar as the solution
  - Used in production in ALICE Computing Grid - monitoring all Grid sites and nodes
  - Integrated with JAliEn Job broker for job matching
  - Allows introducing dynamic constraints in job matching process

Future Work

- Allowing Sites to impose constraints on which jobs it would accept
- Identifying anomalies in monitoring data
- Correlating anomalies with Grid failures to identify misconfigurations

Thank You

kalana.wijethunga@cern.ch,

kalana.16@cse.mrt.ac.lk