



Site Sonar - A Flexible and Extensible Infrastructure Monitoring Tool for ALICE Computing Grid

CHEP 2023



Authors: Kalana Wijethunga, Maksim Storetvedt, Costin Grigoras, Latchezar Betev, Maarten Litmaath,
Gayashan Amarasinghe, Indika Perera
Presented by: Maksim Storetvedt



Content




1. Introduction
2. Research Problem
3. Literature review
4. Outcomes
5. Design & Implementation

Introduction

Introduction

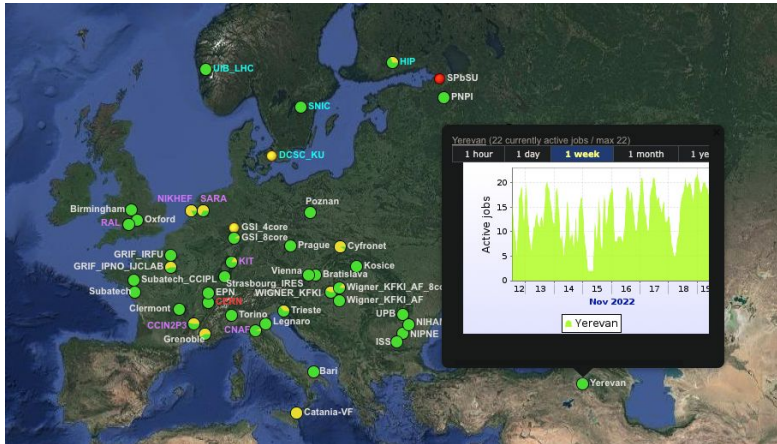
Grid Computing -

- Comprised of individual computing sites distributed across the globe
- Connected by a middleware and a robust network
- A computing site can be a HTC or a HPC facility with attached storage
- Resulting infrastructure is a distributed batch system with individual elements hidden from the end user



Grid Computing

Introduction ctd..



Grid monitoring -

- Measure and publish state of resource at a particular point in time
- Filtered and aggregated to provide a full overview of the Grid
- Used by administrators to monitor health and efficiency of the Grid
- Used by end users to follow the payload processing status

Research Problem

Research Problem



- Grid is **heterogeneous**
 - Different hardware, software, packages, locations, configurations, architectures
- Sometimes hard to **predict the job behaviour**
 - Payload successfully executed on one site may fail or behave differently on a different site
- Requires a more **granular monitoring** and understanding of the Grid sites and sometimes **individual Grid nodes**

Motivation



It is important to

- Ensure the Grid sites are **compatible with the software** versions required by payloads
- Ensure a **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**
- Survey the entirety of the Grid sites and nodes and **provide statistical analysis** of various parameters
- Goal: have a **full picture** of the current status of the Grid and optimize for efficient use

Problem Statement



Develop an extensible framework to identify node configurations, monitor the grid, and visualize Grid infrastructure information

Literature Review

Issues with existing systems

- Data pull model
 - A central server 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 agent installations 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

Summary

	MonAlisa [1]	GridIce [2]	Paryavekshanam[3]	MONIT [4]	SiteSonar
Focus on infrastructure monitoring	✗	✓	✓	✓	✓
Push model	✓	✗	✗	✓	✓
No agent installation on sites	✗	✗	✓	✗	✓
Allow collecting unstructured data	✗	✗	✗	✓	✓
Easily extensible	✗	✗	✗	✗	✓
Can act upon alarming information	✗	✗	✗	✓	✓



Outcomes

Outcomes



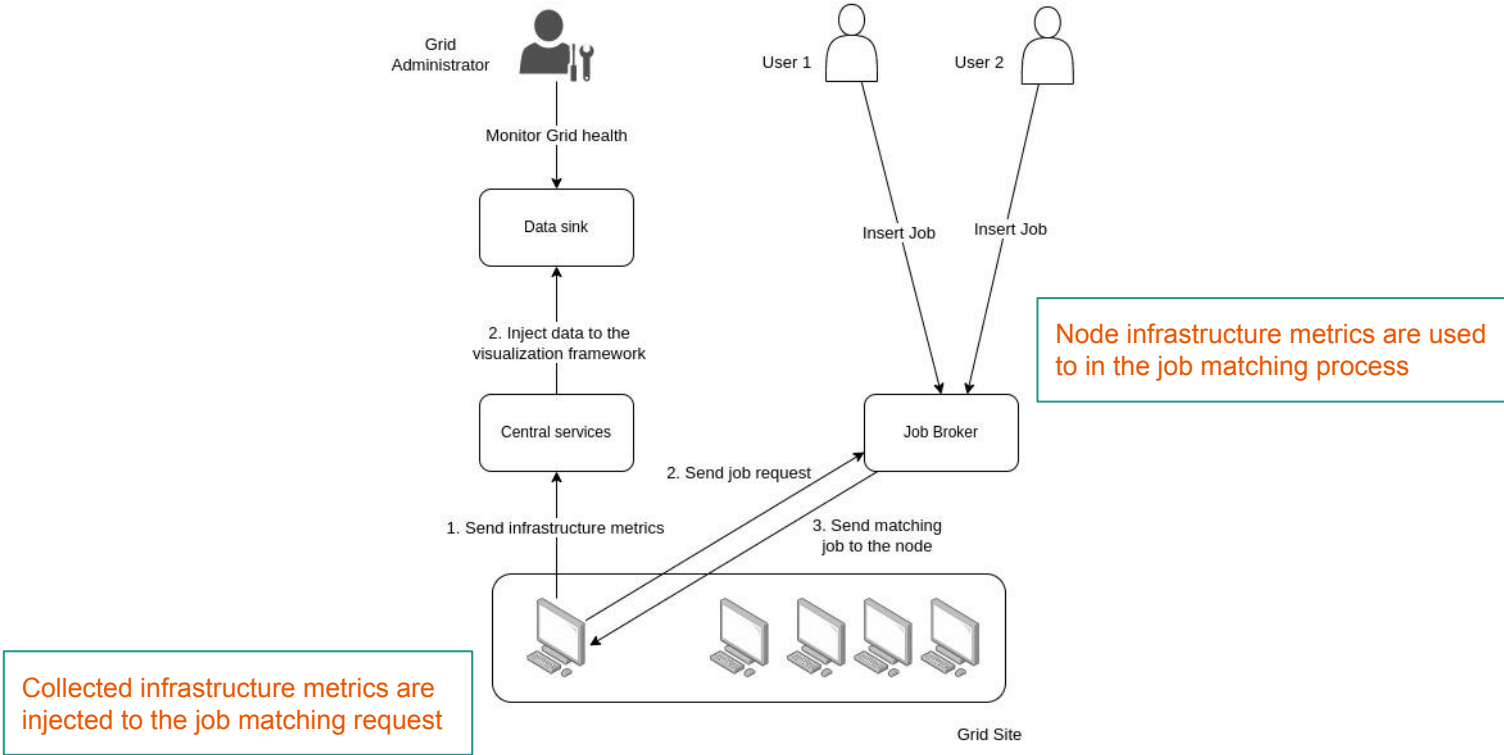
A new Grid Infrastructure Monitoring Tool called “**Site Sonar**” that provides a methodology to identify the capabilities of individual worker nodes in a distributed computing Grid

consisting of a :

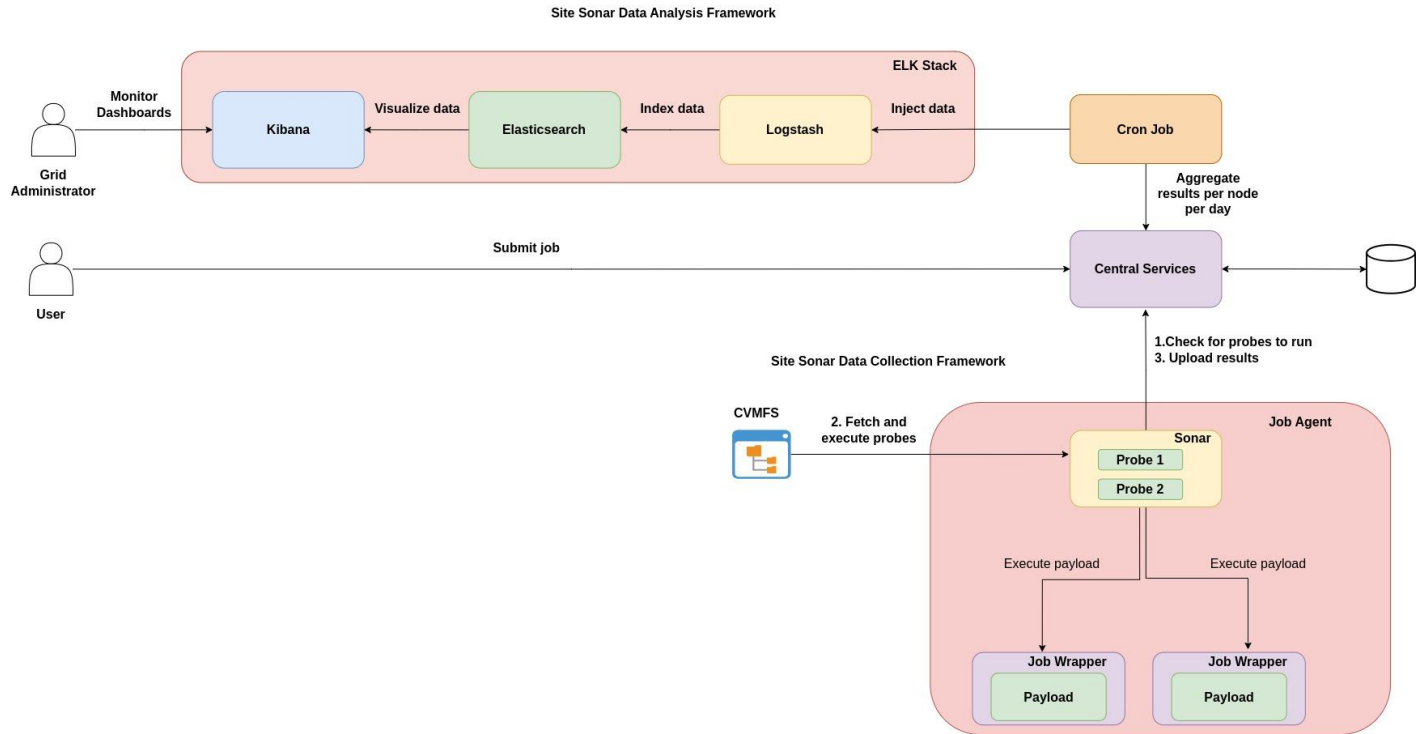
- a. Data Collection framework that is
 - i. Flexible to change data structures on demand
 - ii. Easy to add new data collection probes easily
 - iii. Improving Job Matching functionality using collected data
- b. Data Visualization framework that
 - i. Allows post data filtering
 - ii. Provides no-code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

Design & Implementation

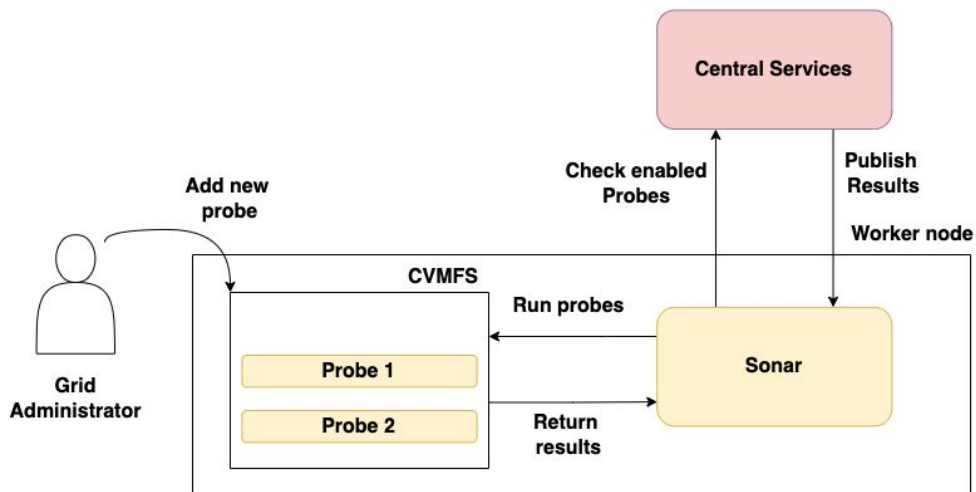
Overview of the solution



Detailed Design



Key features

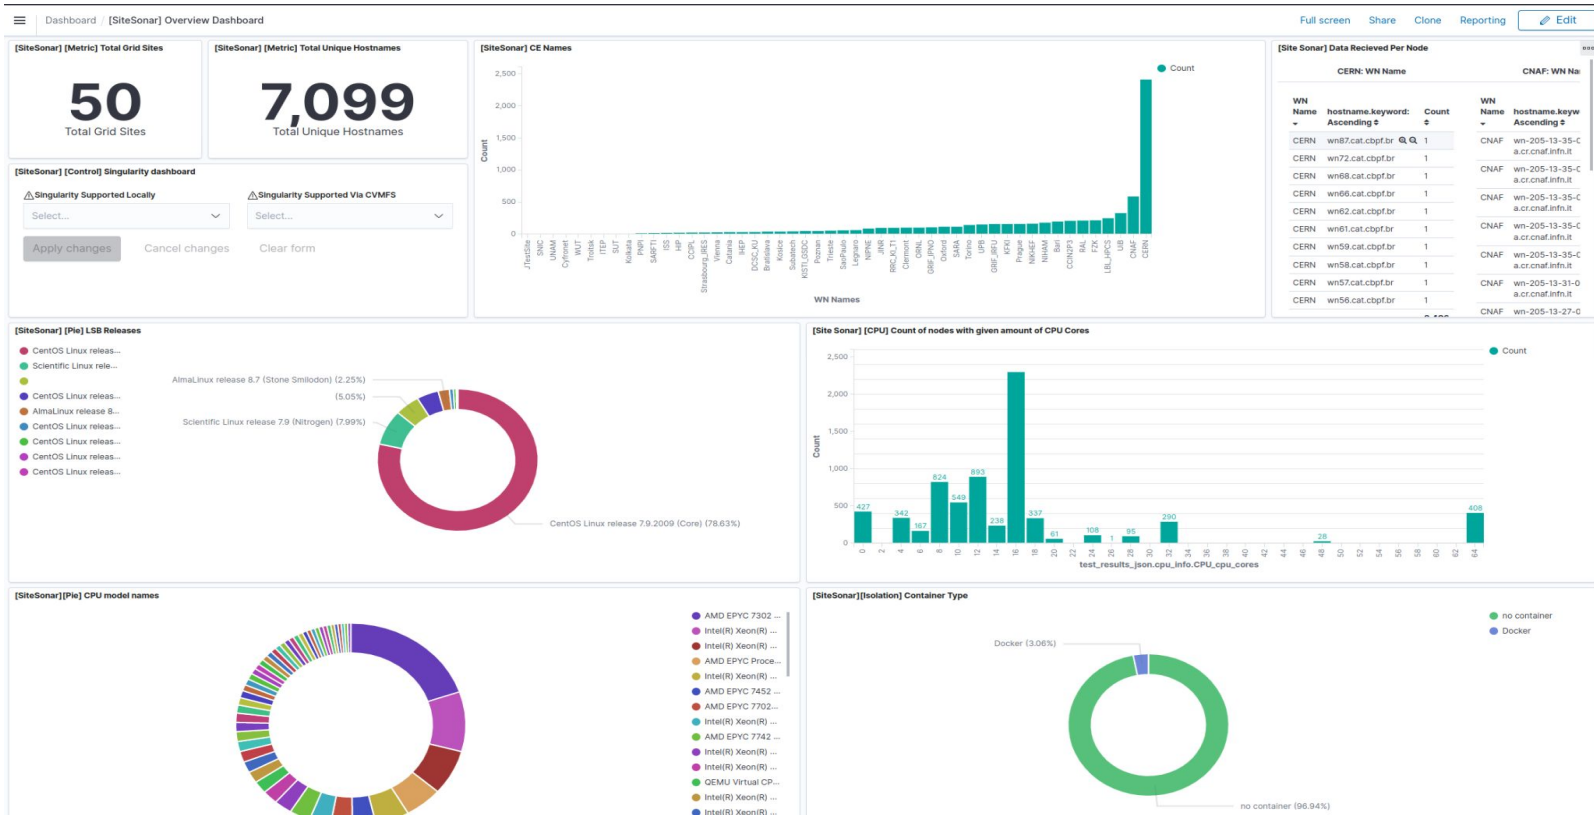


Extensible - Can add or remove new probes to collect metrics without any code changes

```
{ "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)"  
  }  
}
```

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

Grid Monitoring system



Thank You Questions?

email: kalana.wijethunga@cern.ch

References



- [1] I. Legrand et al., "MonALISA: An agent based, dynamic service system to monitor, control and optimize distributed systems", Computer Physics Communications, vol. 180, no. 12, pp. 2472-2498, 2009. Available: [10.1016/j.cpc.2009.08.003](https://doi.org/10.1016/j.cpc.2009.08.003).
- [2] S. Andreozzi et al., "Next steps in the evolution of GridICE: a monitoring tool for grid systems", Journal of Physics: Conference Series, vol. 119, no. 6, p. 062010, 2008. Available: [10.1088/1742-6596/119/6/062010](https://doi.org/10.1088/1742-6596/119/6/062010).
- [3] K. Prasad, H. Gupta, N. Mangala, C. Subrata, H. Deepika and P. Rao, "Challenges of monitoring tool for operational indian national grid GARUDA", 2013 National Conference on Parallel Computing Technologies (PARCOMPTECH), 2013. Available: [10.1109/parcomptech.2013.6621396](https://doi.org/10.1109/parcomptech.2013.6621396)
- [4] A. Aimar et al., "MONIT: Monitoring the CERN Data Centres and the WLCG Infrastructure", EPJ Web of Conferences, vol. 214, p. 08031, 2019. Available: [10.1051/epjconf/201921408031](https://doi.org/10.1051/epjconf/201921408031).