ADMIRE
malleable data solutions for HPC

Adaptive multi-tier intelligent data manager for Exascale

EuroHPC
Joint Undertaking

admire-eurohpc.eu

ADMIRE Users Day

Metric Proxy: Enabling real-time measurement at Supercomputer Scale

Jean-Baptiste Besnard, ParaTools SAS

December 12th 2023.

Barcelona Supercomputing Center

Grant Agreement number: 956748 — ADMIRE — H2020-JTI-EuroHPC-2019-1

Performance Monitoring For Malleability

- **Malleability** is about adapting the payload to external constraints to maximise machine throughput:
  - Optimize computation
  - Minimize wait-time
  - Maximize machine utilization
  - Lower Power
- It is a multi-criterion process, and therefore it requires a wide-range of monitoring capabilities to feed the various models.
- This motivated a general approach for monitoring in ADMIRE with two main challenges:
  - Need for **real-time** data (malleability is temporal)
  - Need for **machine-wide** metrics
  - Need for **per-program** models

} **Challenging !**

Example: Moldability

Efficiency of Rodinia LU Benchmark at various scales and problem sizes

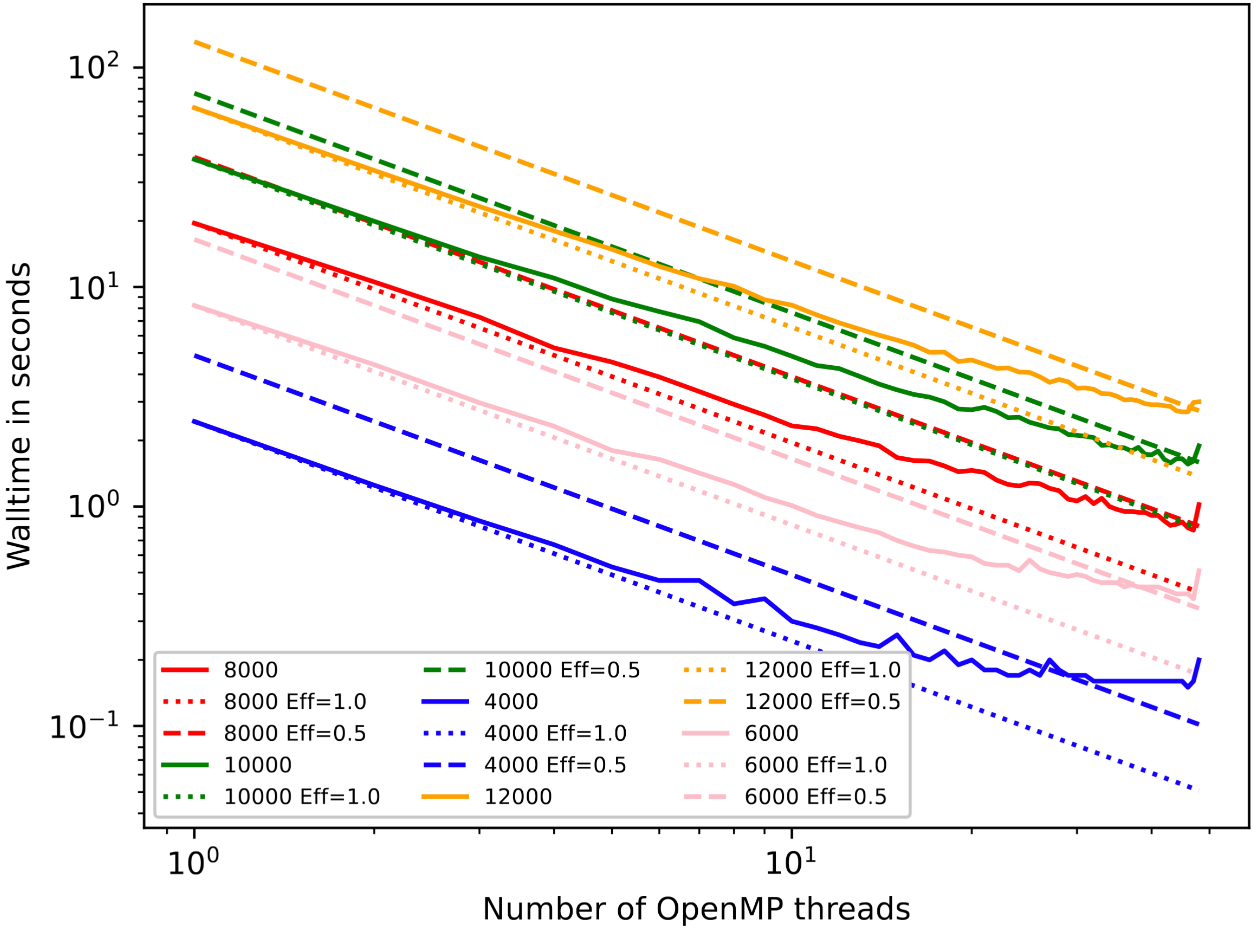

It is a special case of *malleability*, called *moldability*, or more straightforwardly:

« choosing the right configuration at program start ».

Recall of the ADMIRE Feedback Loop

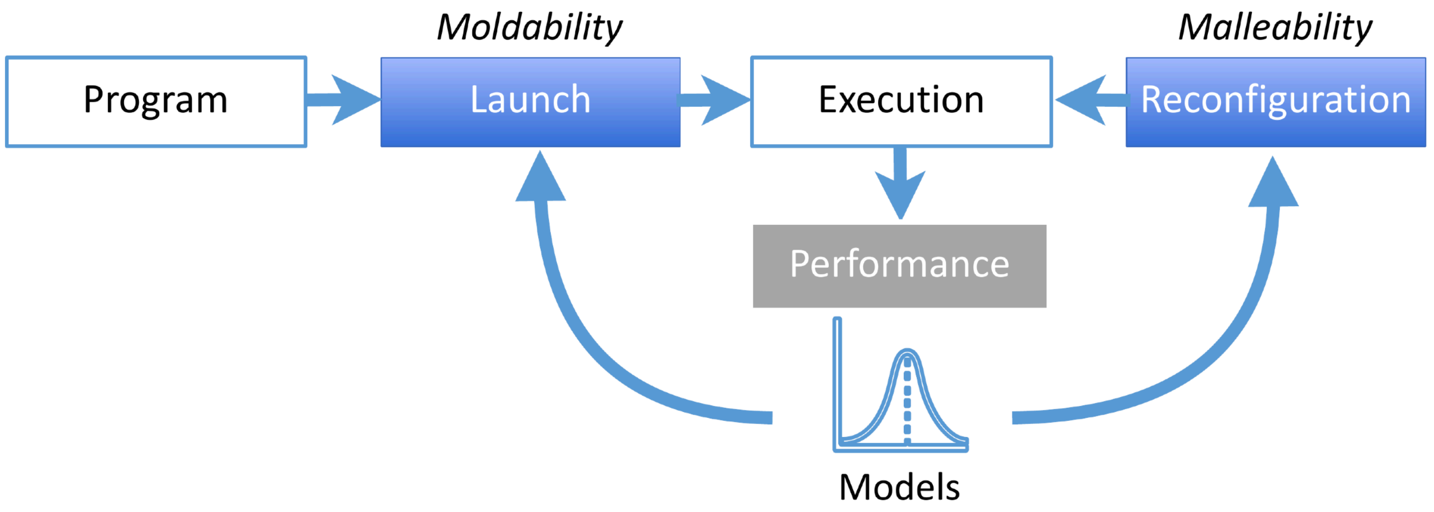

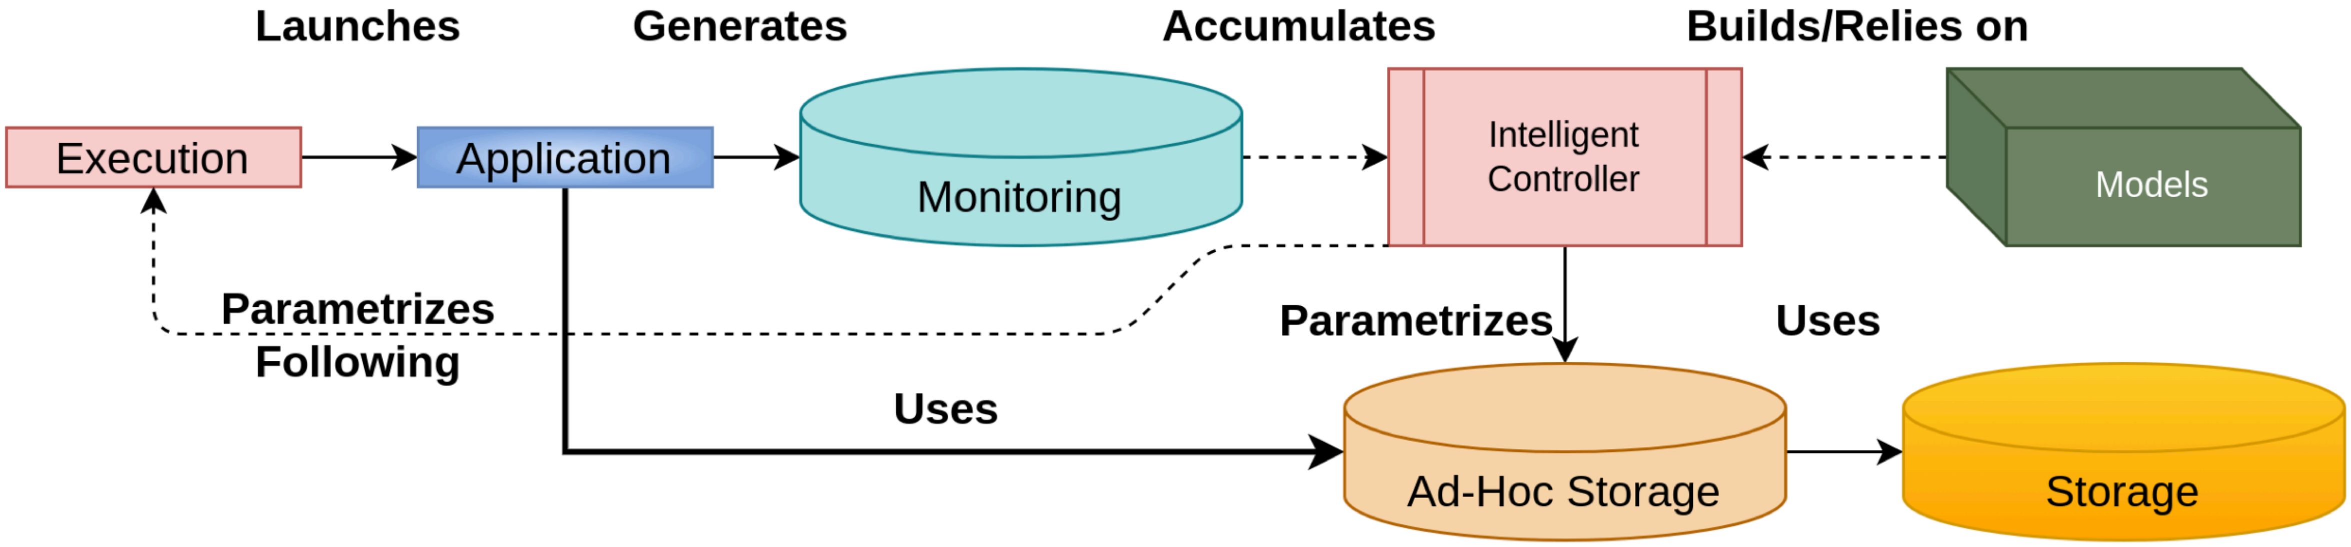

Recall of the ADMIRE Feedback Loop

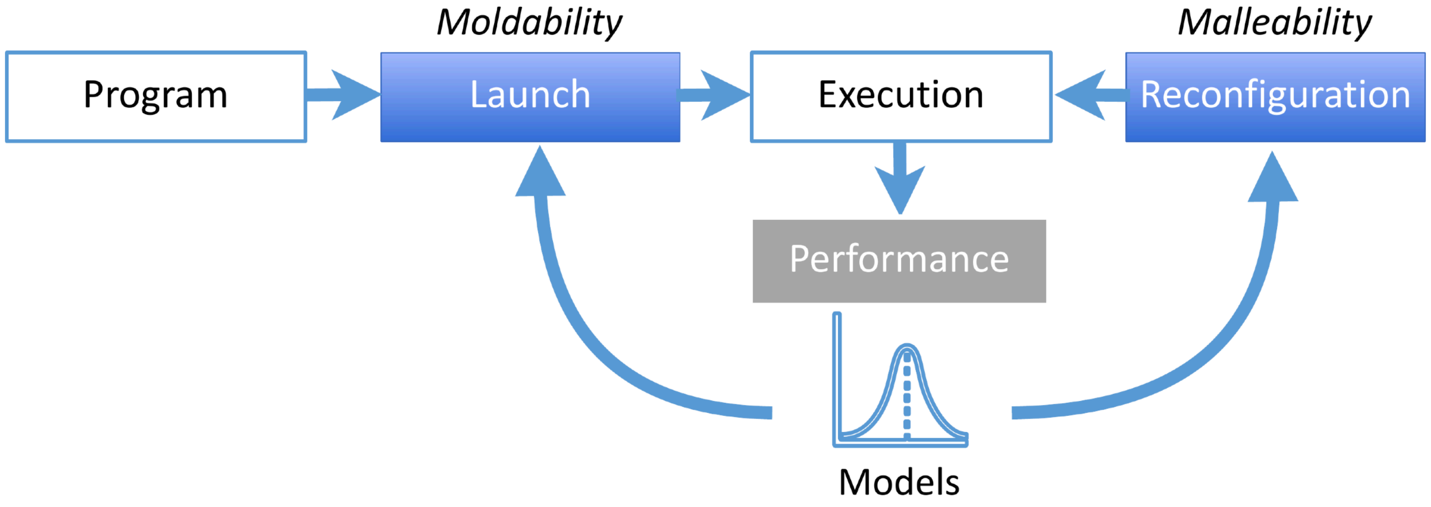

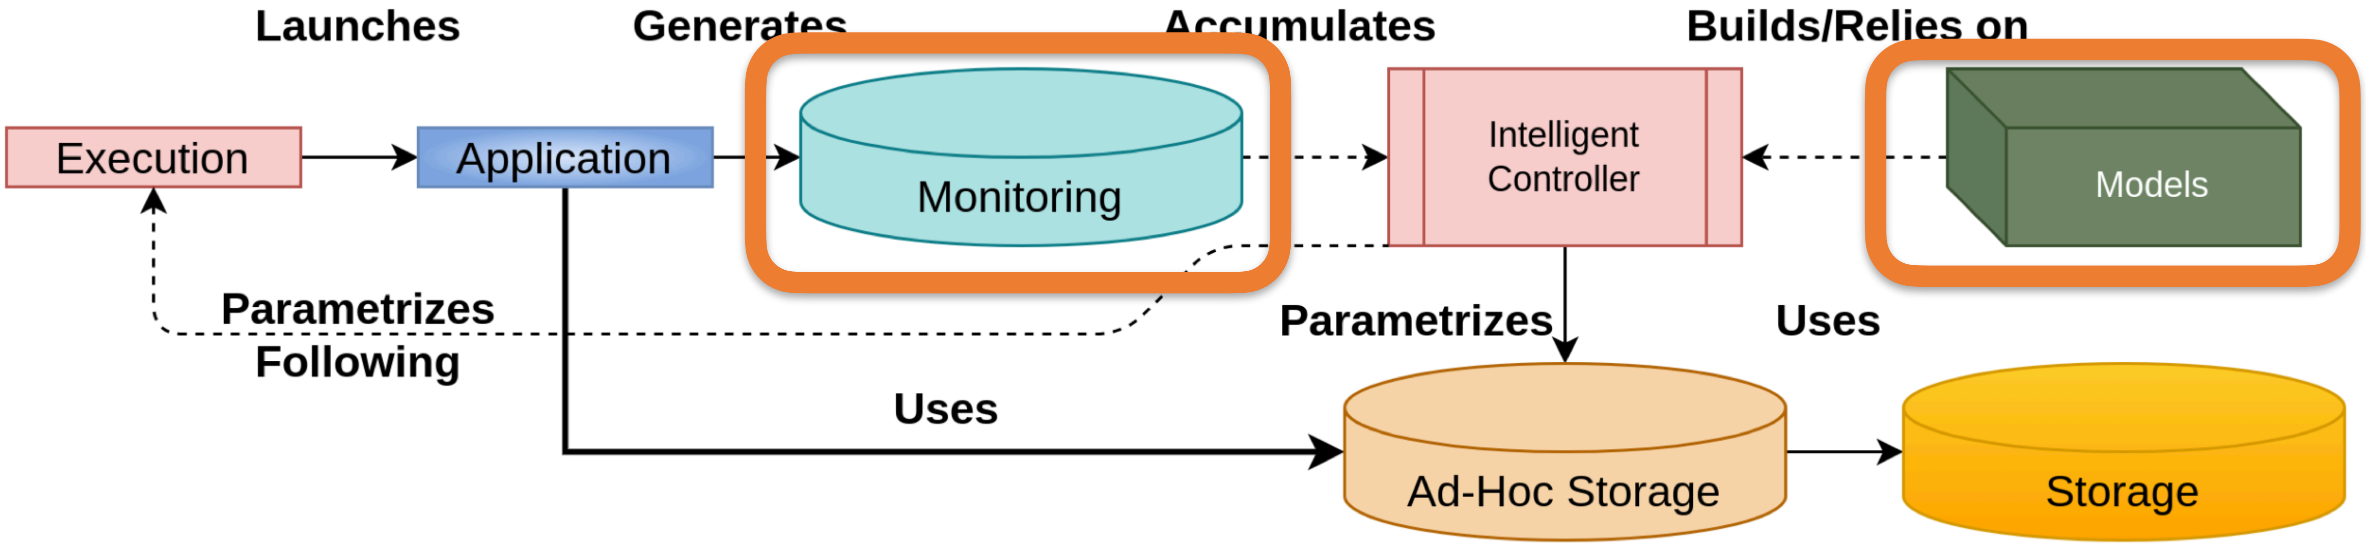

Choosing the Right Measurement Granularity

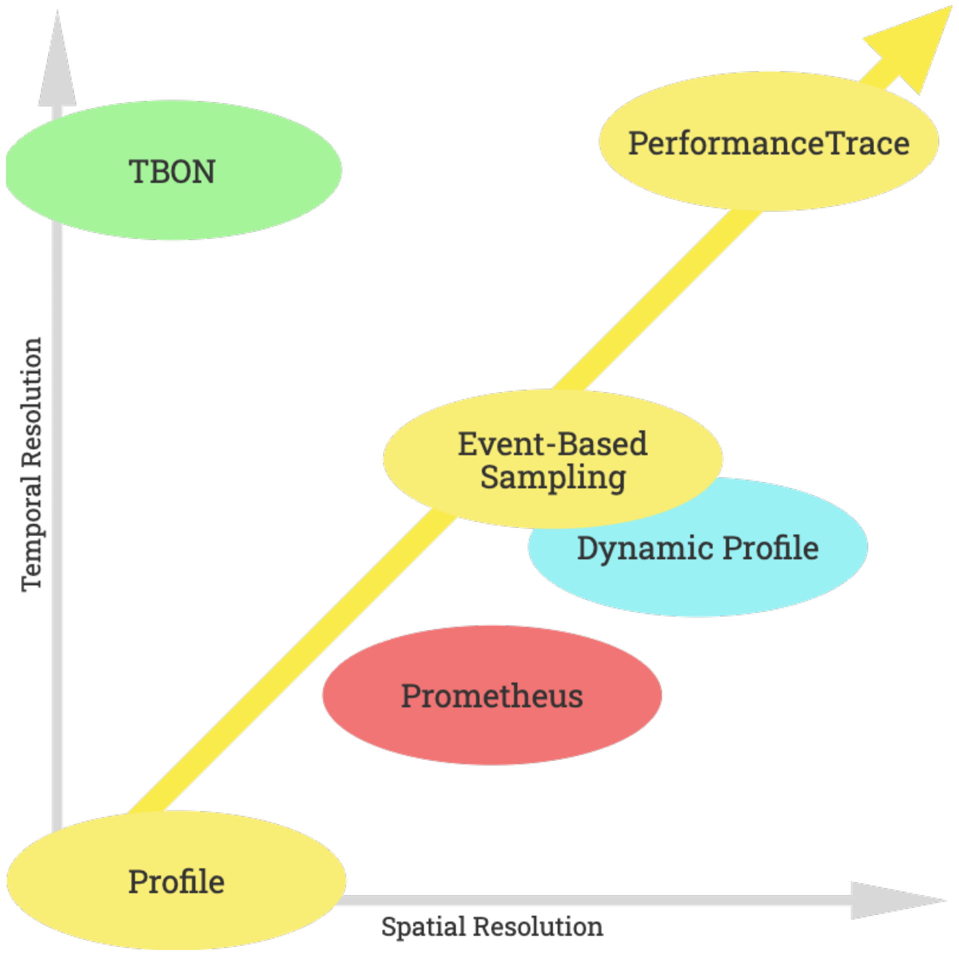

Performance measurement is always a compromise between Verbosity and measurement / storage overhead.

Choosing the Right Measurement Granularity

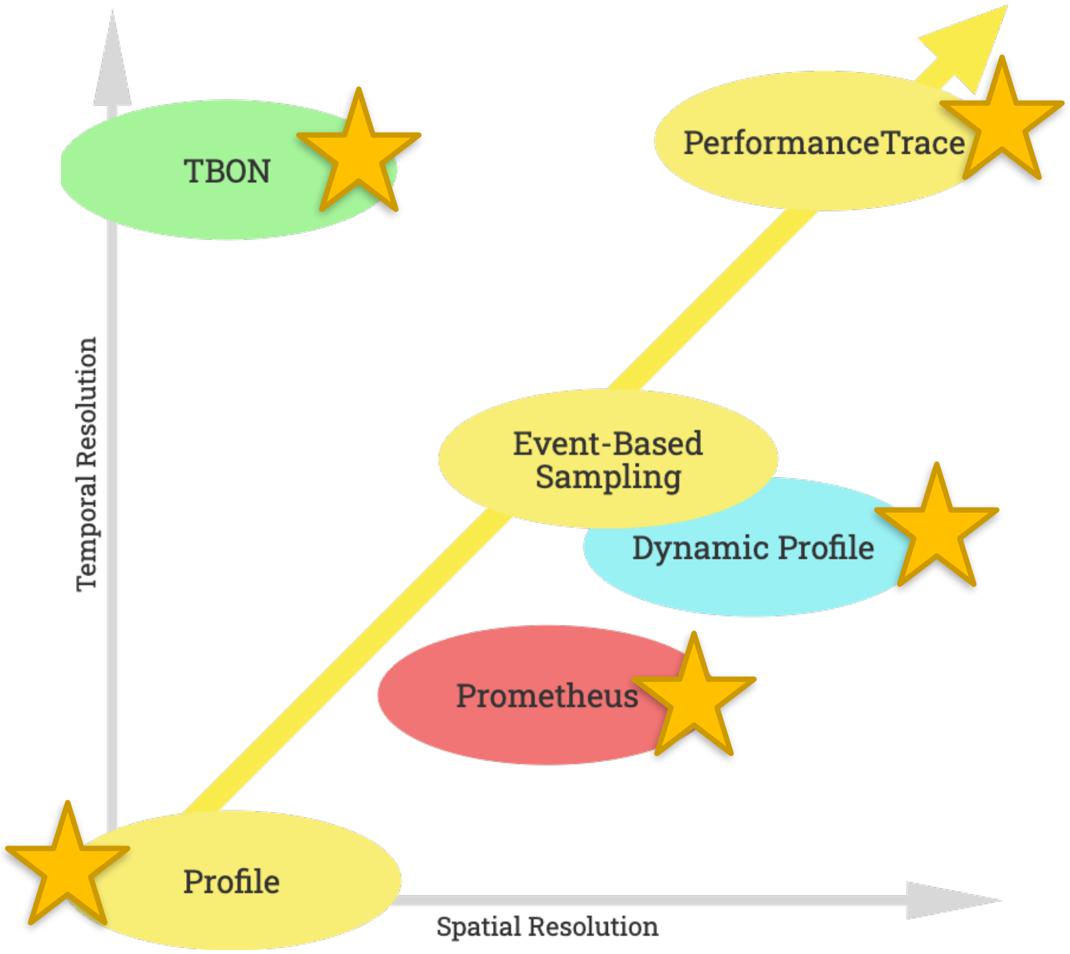

- TBON: for real-time reduction of performance data using LIMITLESS
- Resampled performance traces: for temporal series
- Profiles to describe each run
- Prometheus storage for historization
- Real-time summative profiles (a.k.a snapshots) for current state

Metric Proxy Architecture

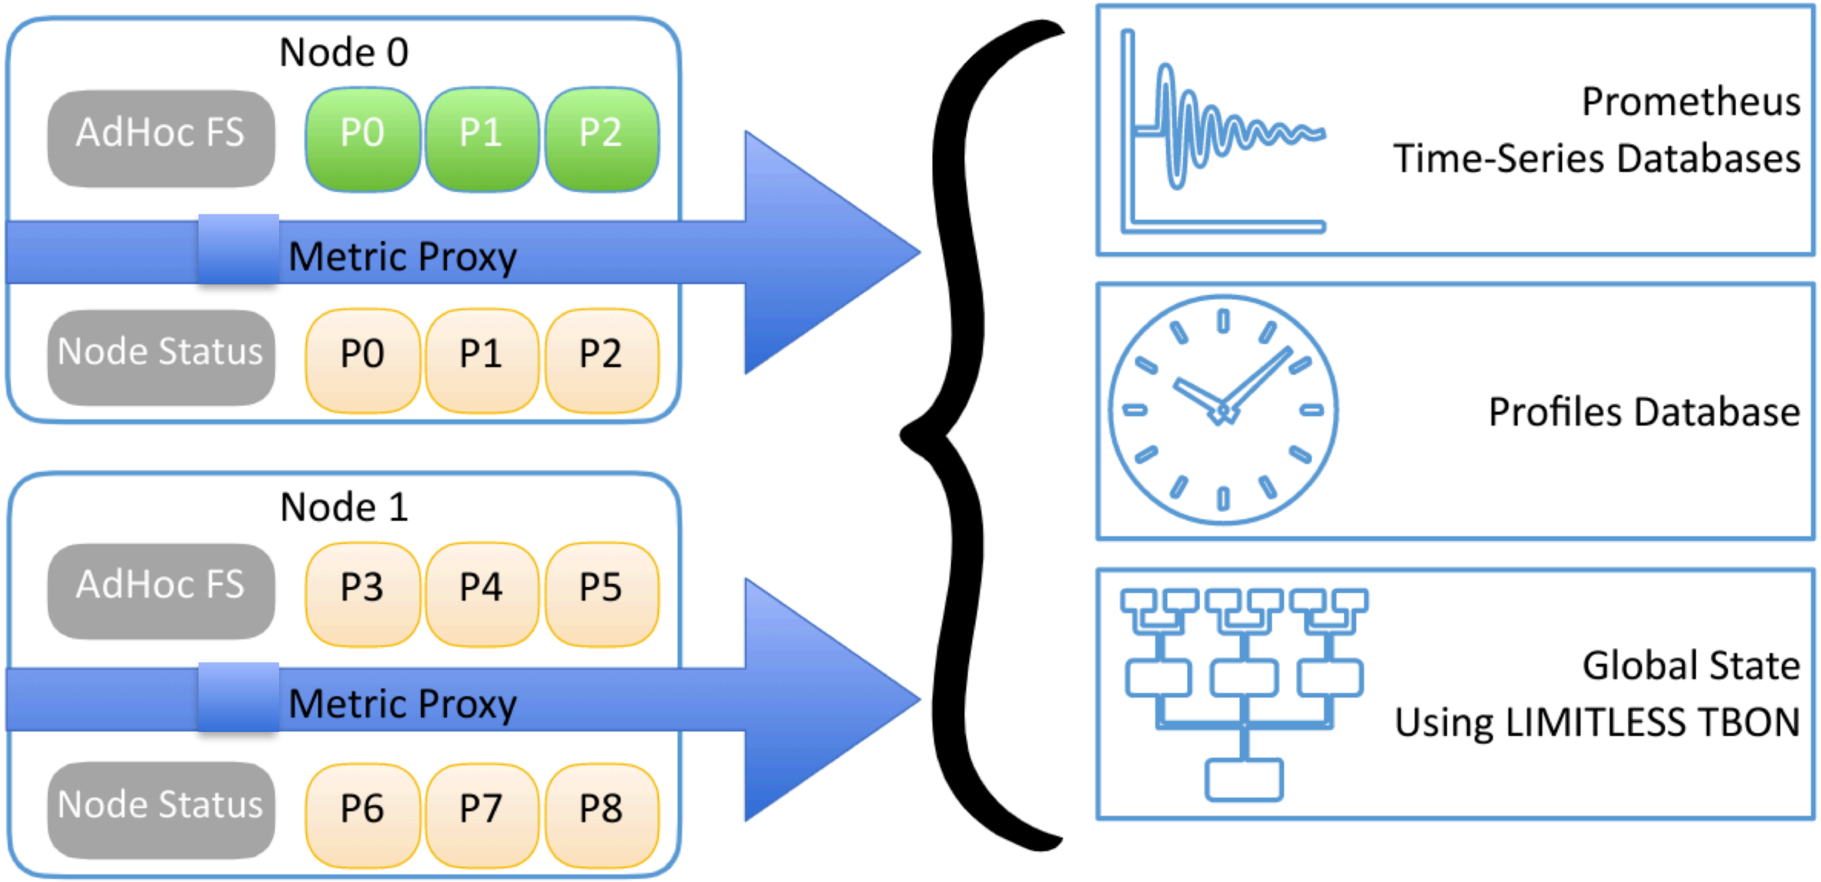

Metric Proxy Architecture

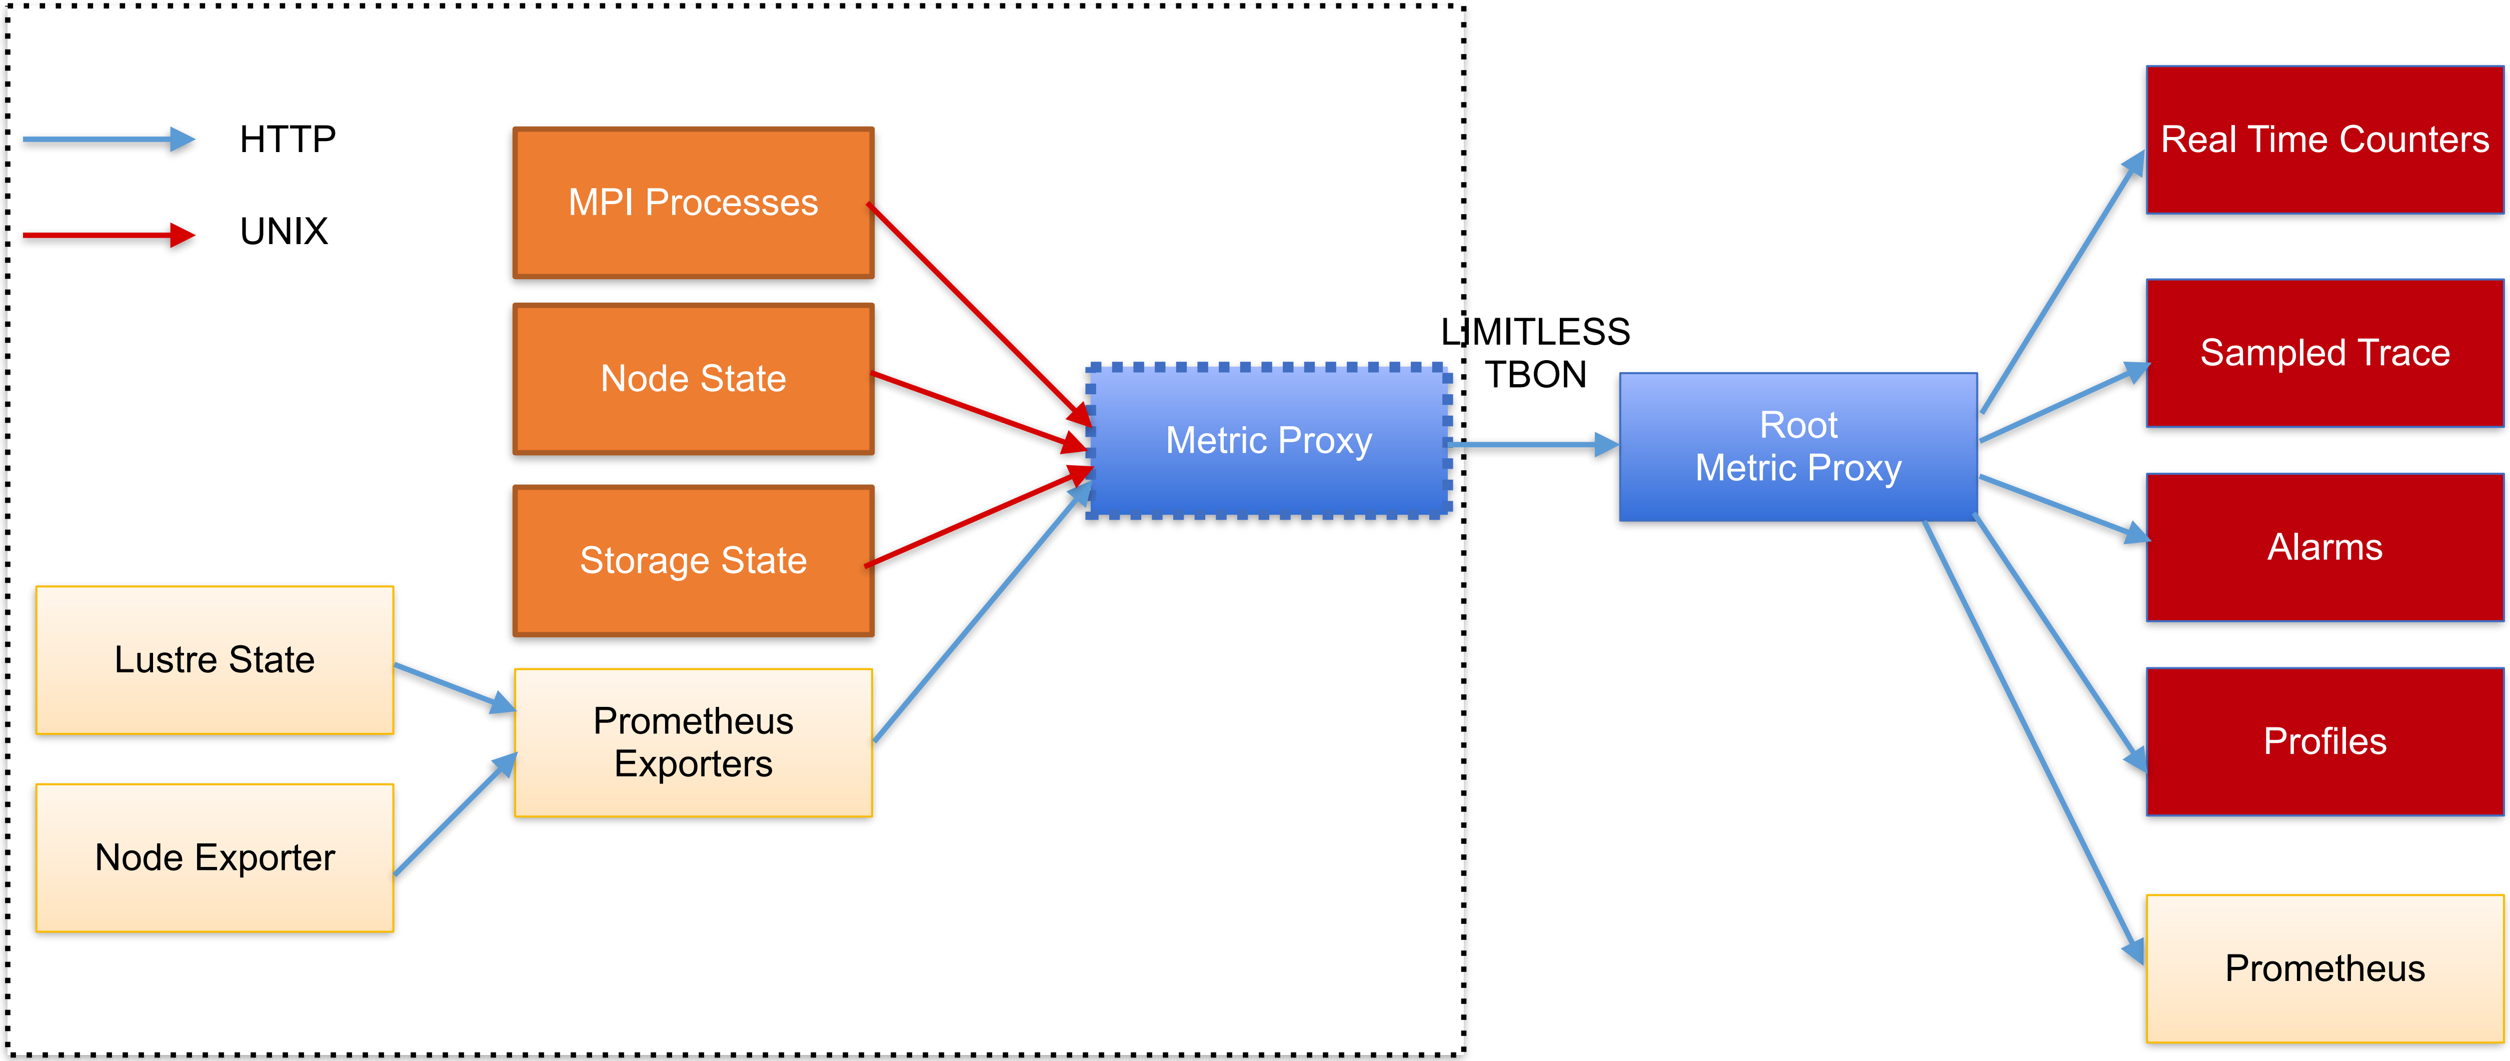

Metric Proxy Architecture

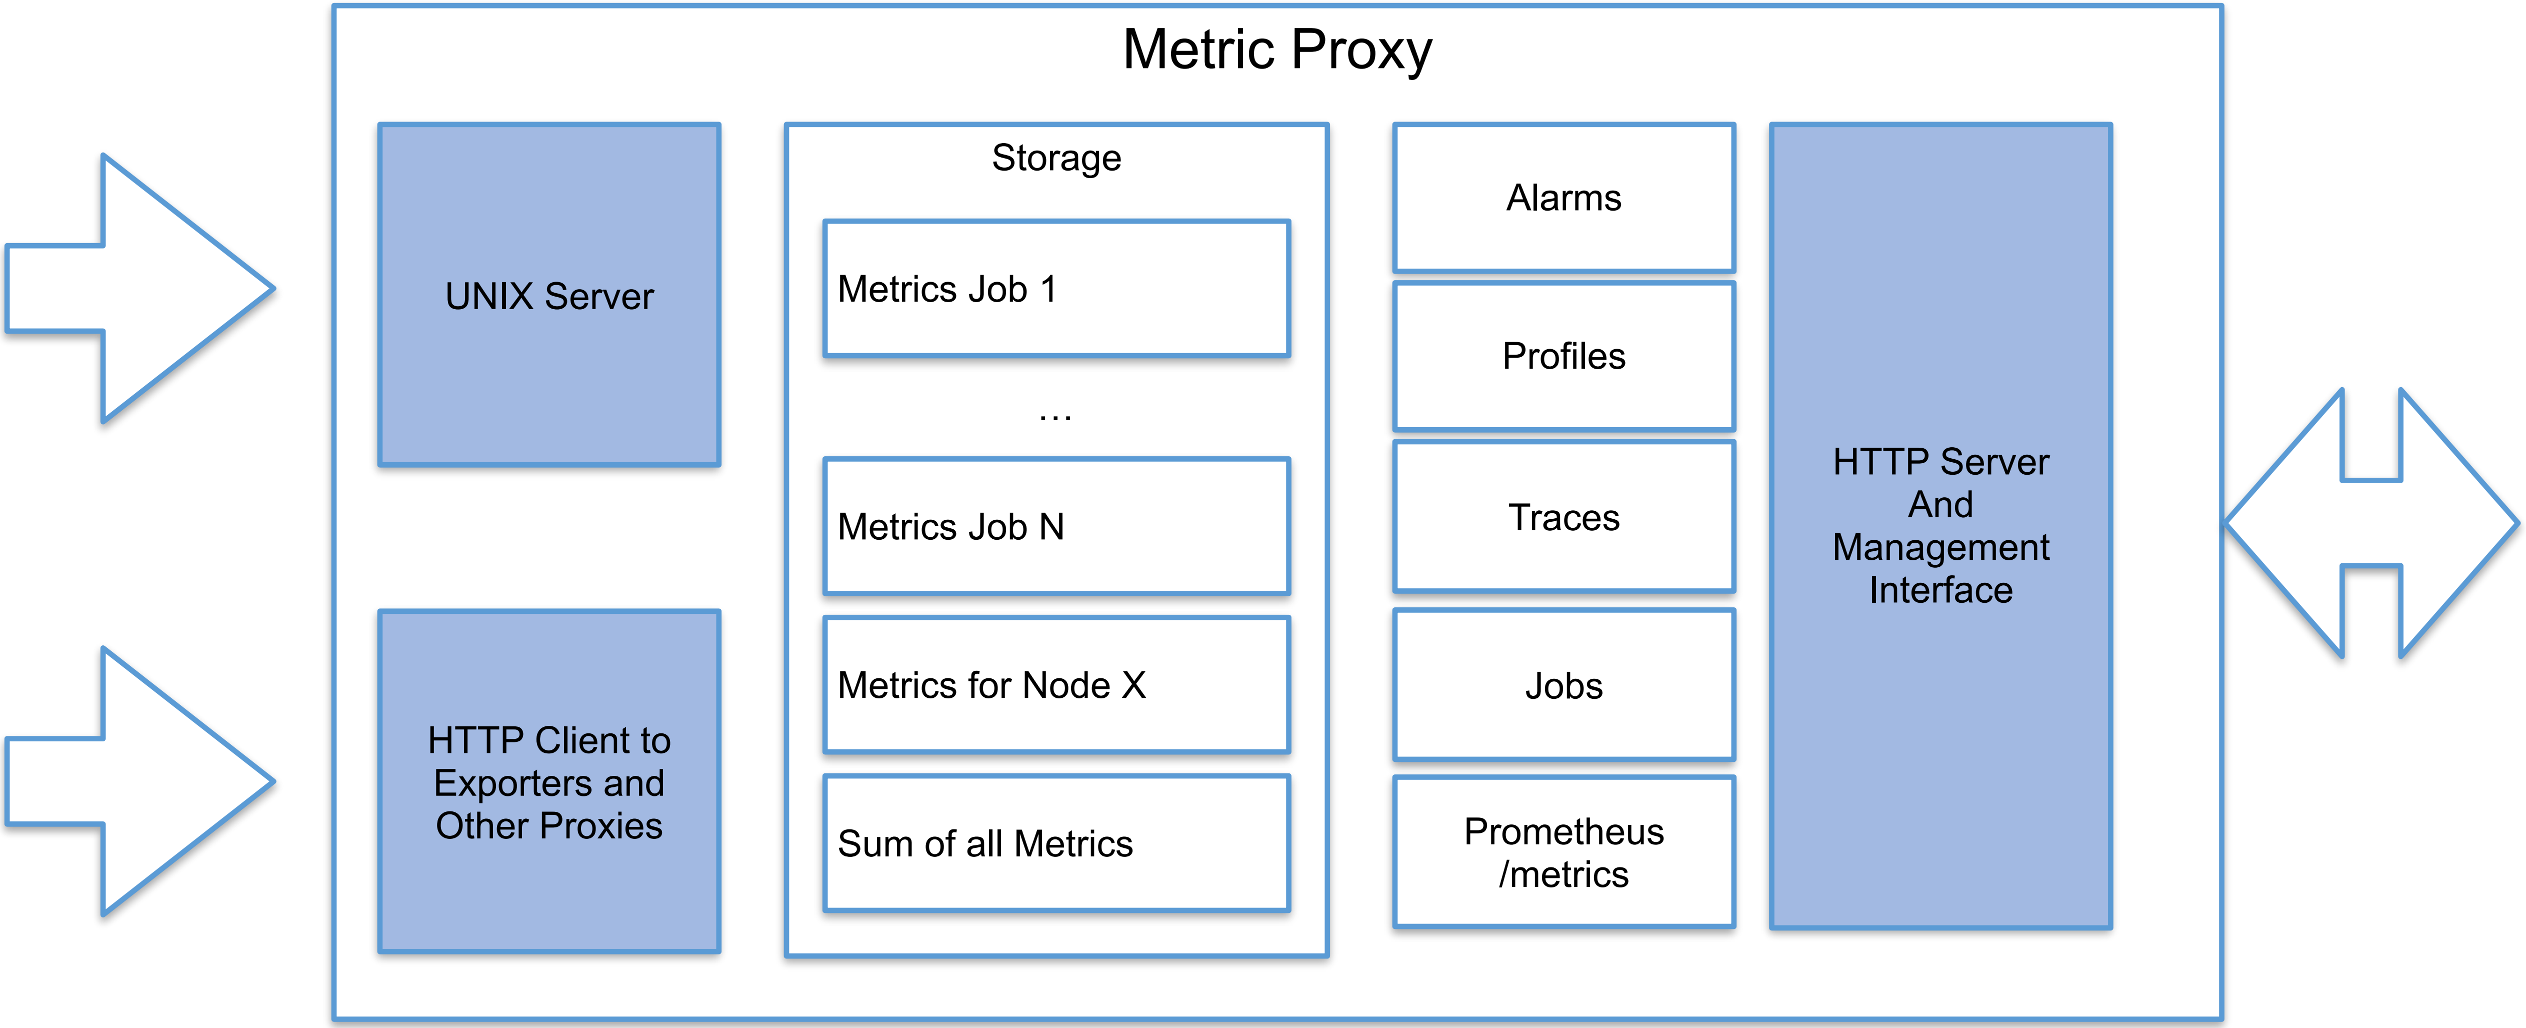

Metric Proxy

UNIX Server

HTTP Client to Exporters and Other Proxies

Storage

Metrics Job 1

…

Metrics Job N

Metrics for Node X

Sum of all Metrics

Alarms

Profiles

Traces

Jobs

Prometheus /metrics

HTTP Server And Management Interface

Handling of Transitive Clients and Jobs

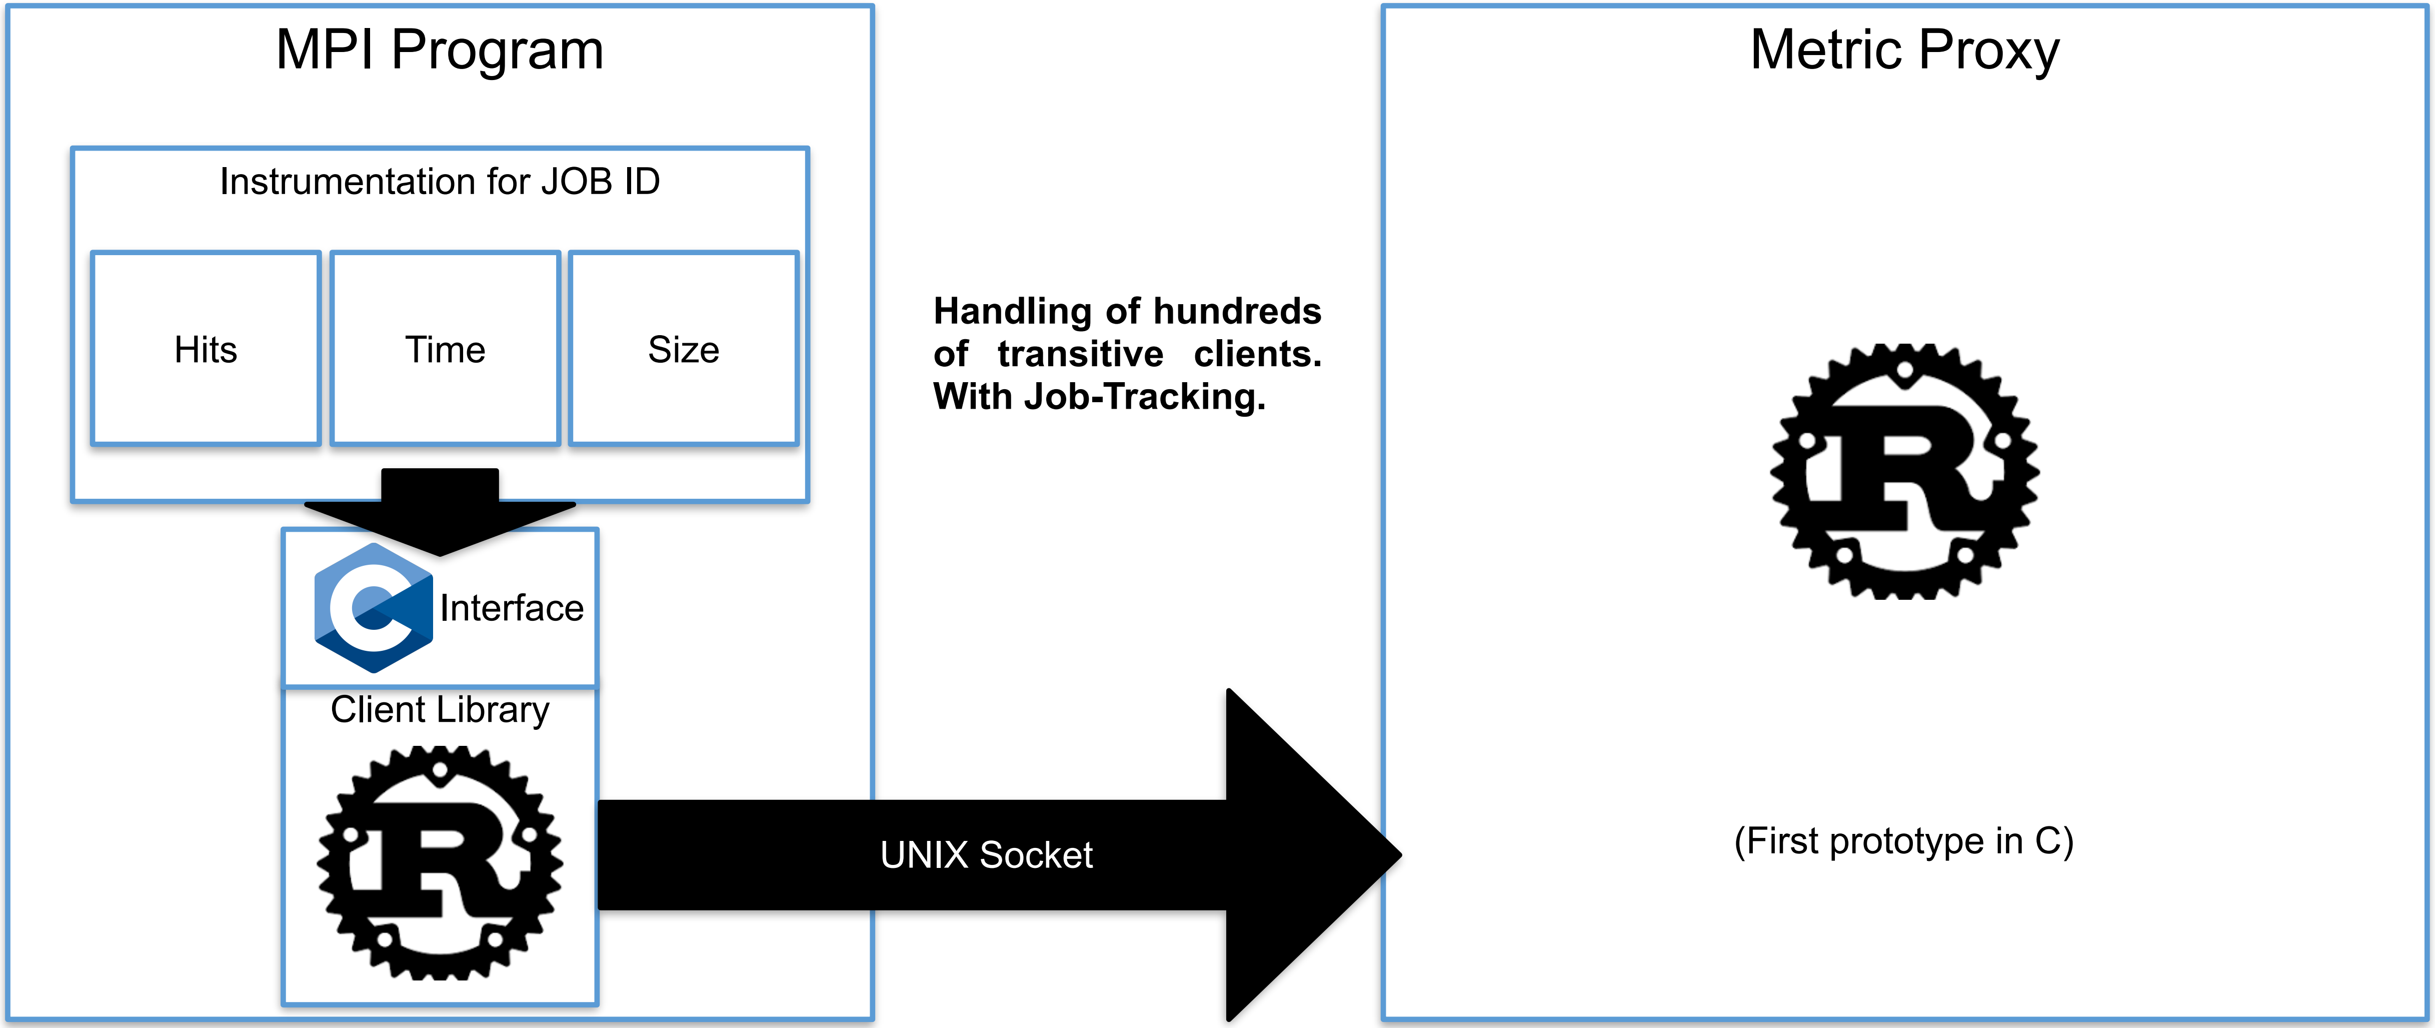

Handling of hundreds of transitive clients. With Job-Tracking.

Metric Proxy : Interfaces

Each metric proxy on each node provides an HTTP endpoint on port 1337 by default.

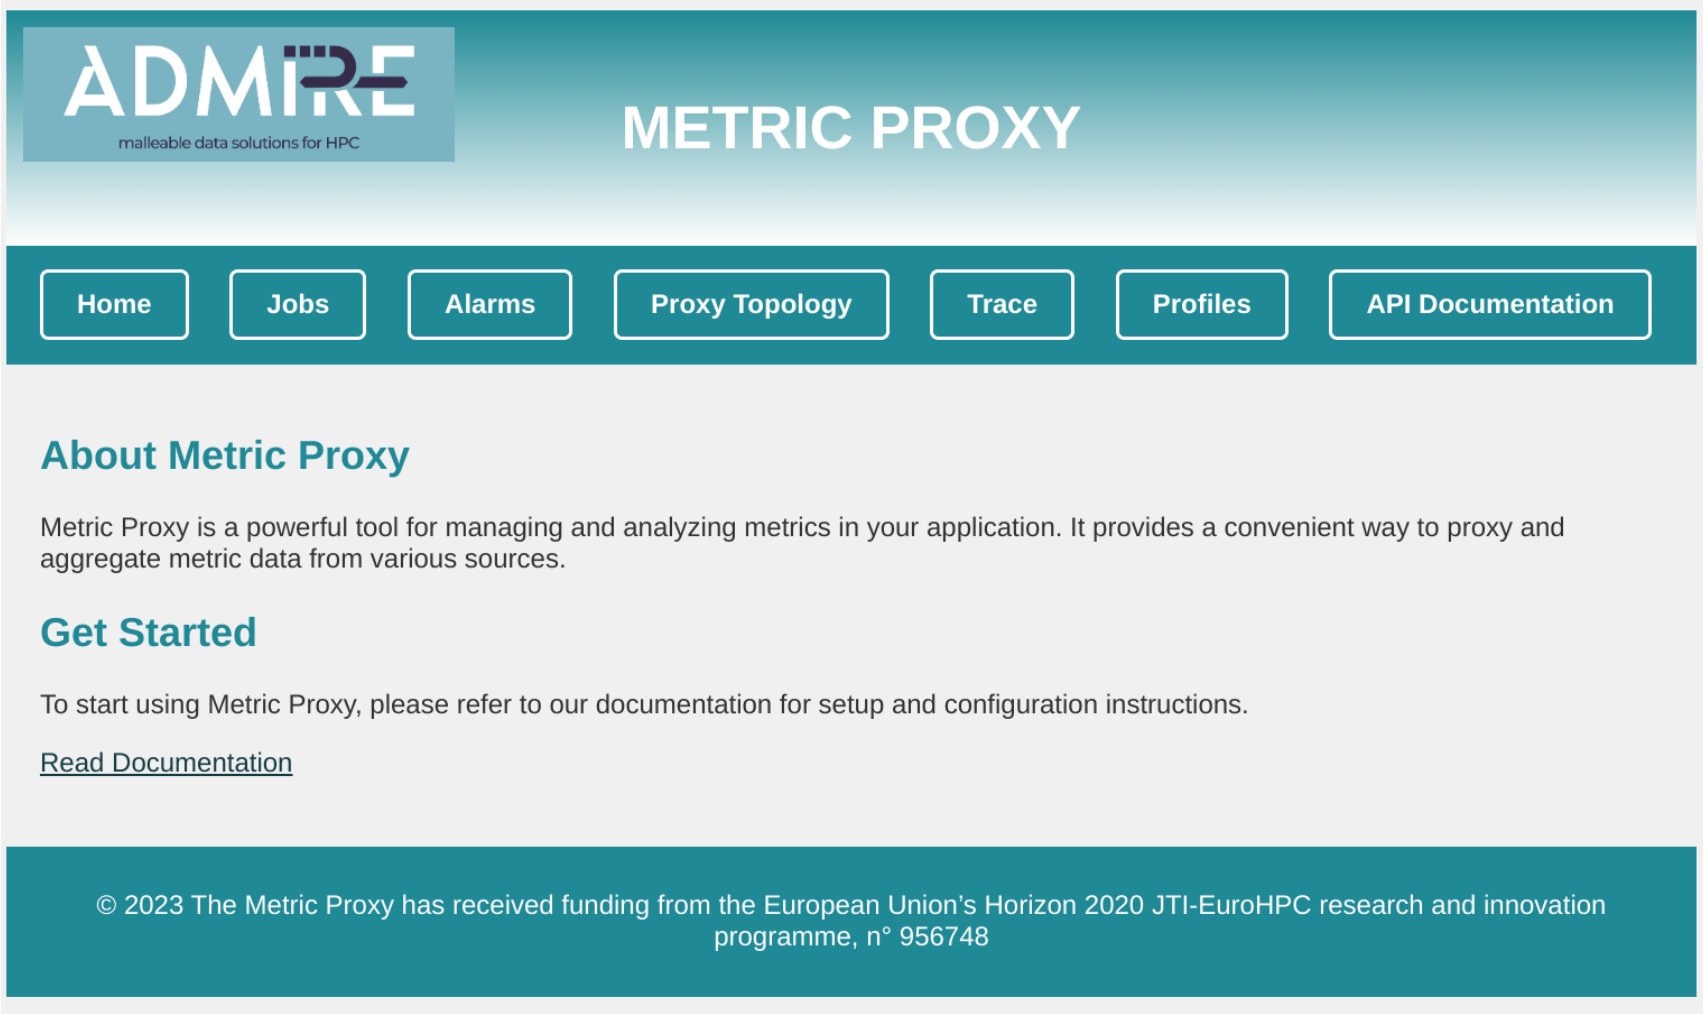

About Metric Proxy

Metric Proxy is a powerful tool for managing and analyzing metrics in your application. It provides a convenient way to proxy and aggregate metric data from various sources.

Get Started

To start using Metric Proxy, please refer to our documentation for setup and configuration instructions.

Read Documentation

© 2023 The Metric Proxy has received funding from the European Union's Horizon 2020 JTI-EuroHPC research and innovation programme, n° 956748

Proxy Topology

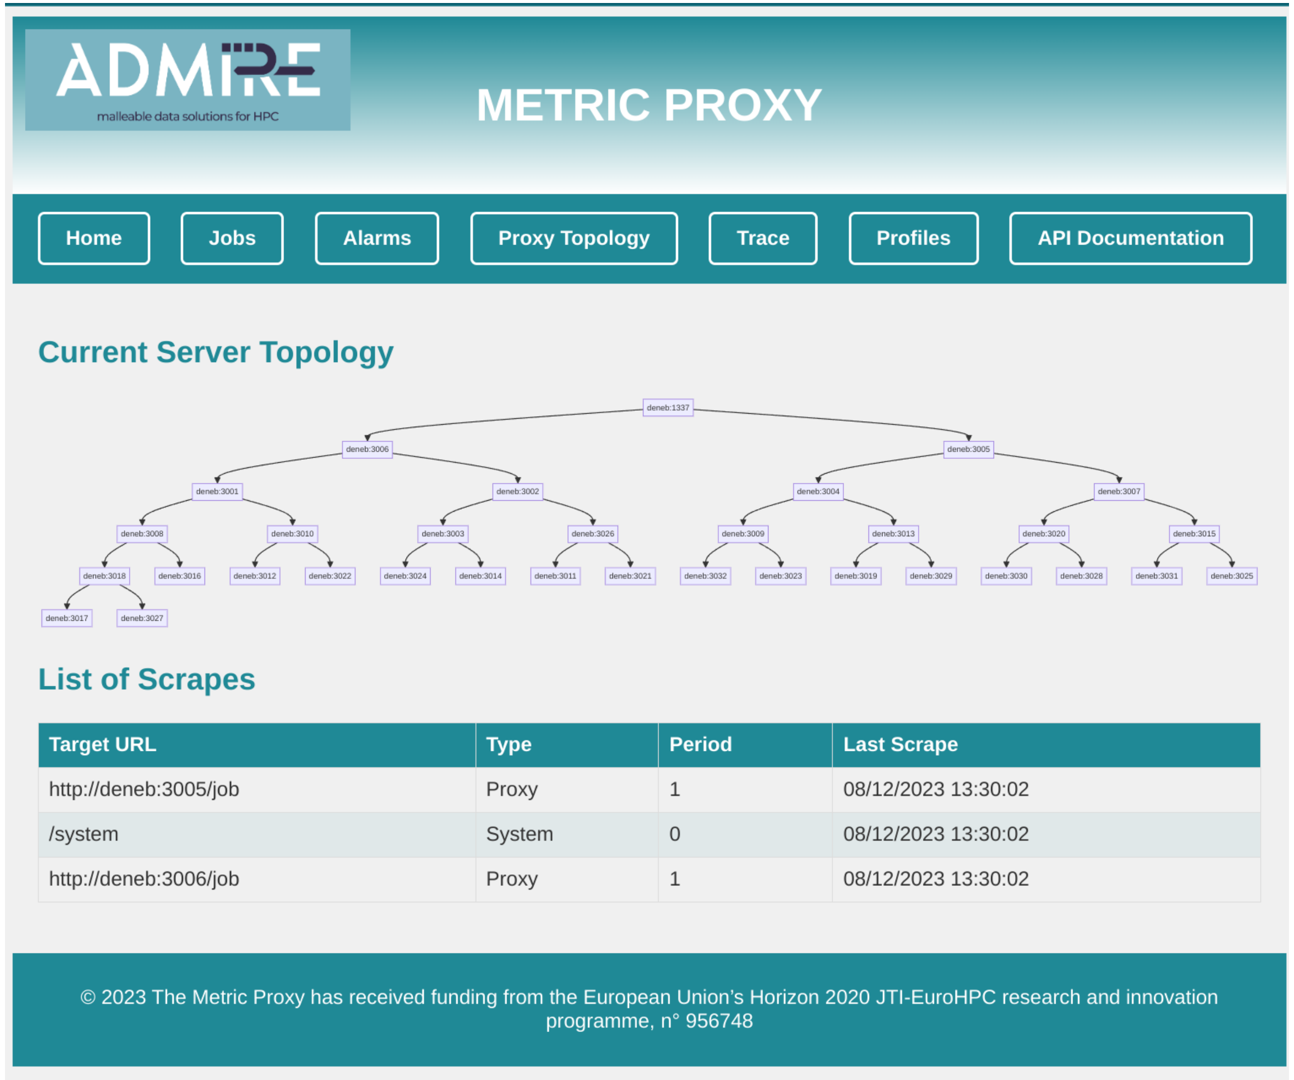

METRIC PROXY

Current Server Topology

List of Scrapes

Target URL	Type	Period	Last Scrape
http://deneb:3005/job	Proxy	1	08/12/2023 13:30:02
/system	System	0	08/12/2023 13:30:02
http://deneb:3006/job	Proxy	1	08/12/2023 13:30:02

© 2023 The Metric Proxy has received funding from the European Union's Horizon 2020 JTI-EuroHPC research and innovation programme, n° 956748

The proxy reduction tree is built automatically by « pivoting » the nodes on a root server which then returns the address of one of the proxy. Here an example with 32 nodes, seen from the root.

} **Scrapes**

Real-Time Job Tracking

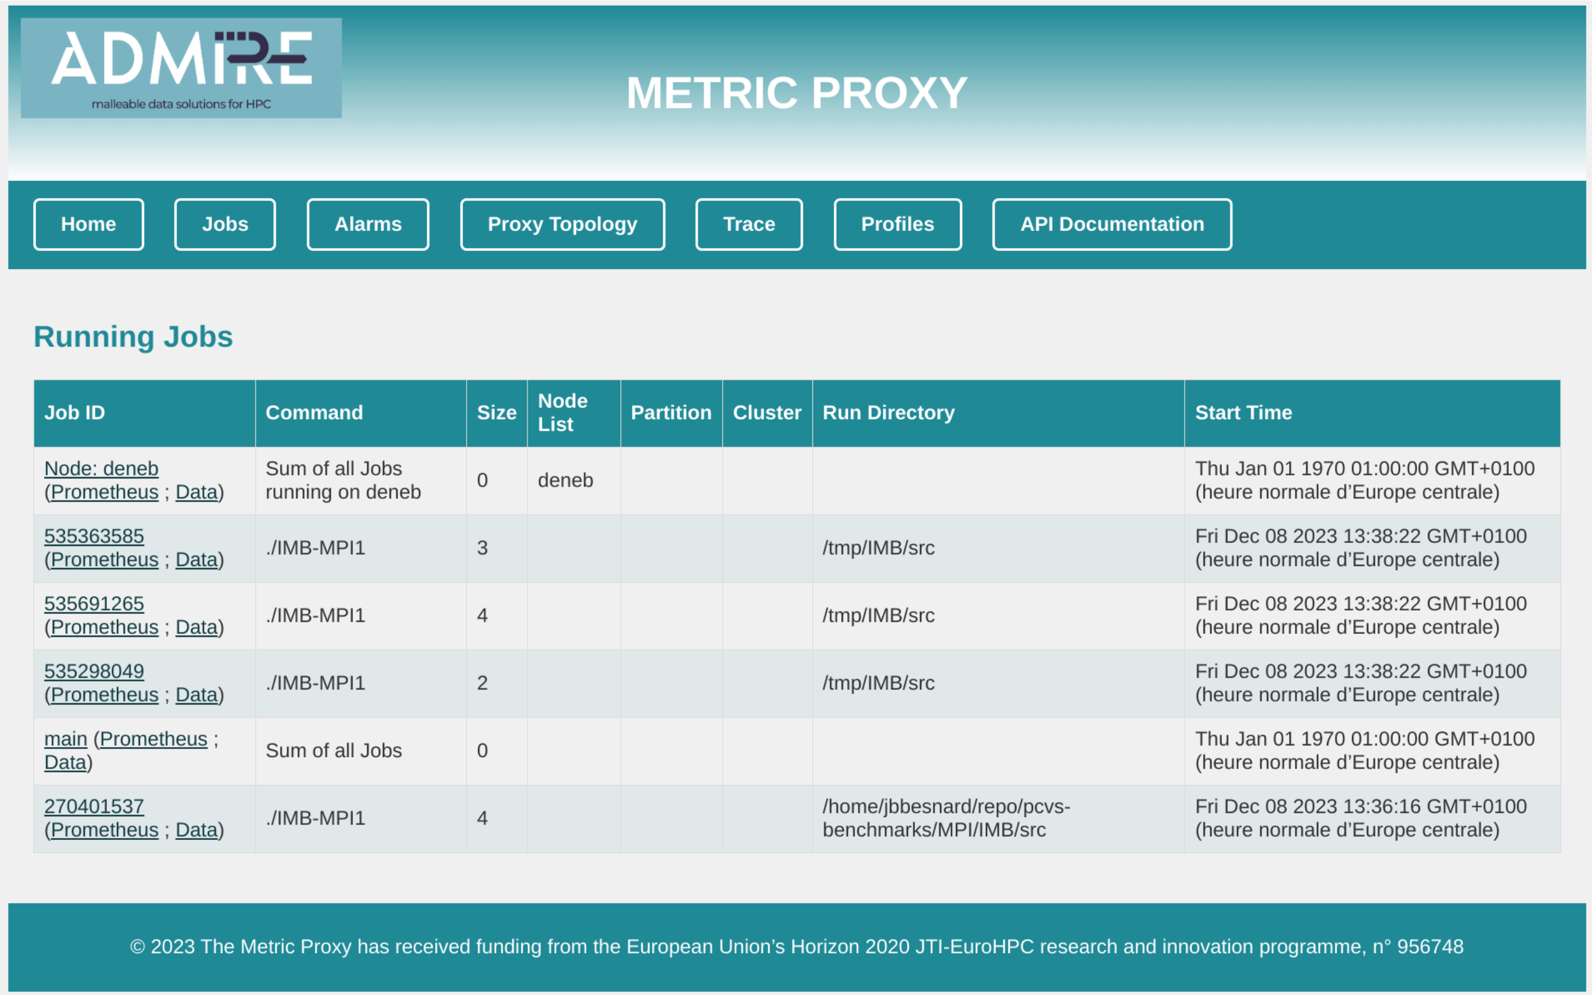

METRIC PROXY

Running Jobs

Job ID	Command	Size	Node List	Partition	Cluster	Run Directory	Start Time
Node: deneb (Prometheus ; Data)	Sum of all Jobs running on deneb	0	deneb				Thu Jan 01 1970 01:00:00 GMT+0100 (heure normale d'Europe centrale)
535363585 (Prometheus ; Data)	./IMB-MPI1	3				/tmp/IMB/src	Fri Dec 08 2023 13:38:22 GMT+0100 (heure normale d'Europe centrale)
535691265 (Prometheus ; Data)	./IMB-MPI1	4				/tmp/IMB/src	Fri Dec 08 2023 13:38:22 GMT+0100 (heure normale d'Europe centrale)
535298049 (Prometheus ; Data)	./IMB-MPI1	2				/tmp/IMB/src	Fri Dec 08 2023 13:38:22 GMT+0100 (heure normale d'Europe centrale)
main (Prometheus ; Data)	Sum of all Jobs	0					Thu Jan 01 1970 01:00:00 GMT+0100 (heure normale d'Europe centrale)
270401537 (Prometheus ; Data)	./IMB-MPI1	4				/home/jbbesnard/repo/pcvs-benchmarks/MPI/IMB/src	Fri Dec 08 2023 13:36:16 GMT+0100 (heure normale d'Europe centrale)

© 2023 The Metric Proxy has received funding from the European Union's Horizon 2020 JTI-EuroHPC research and innovation programme, n° 956748

Real-Time Monitoring at Scale

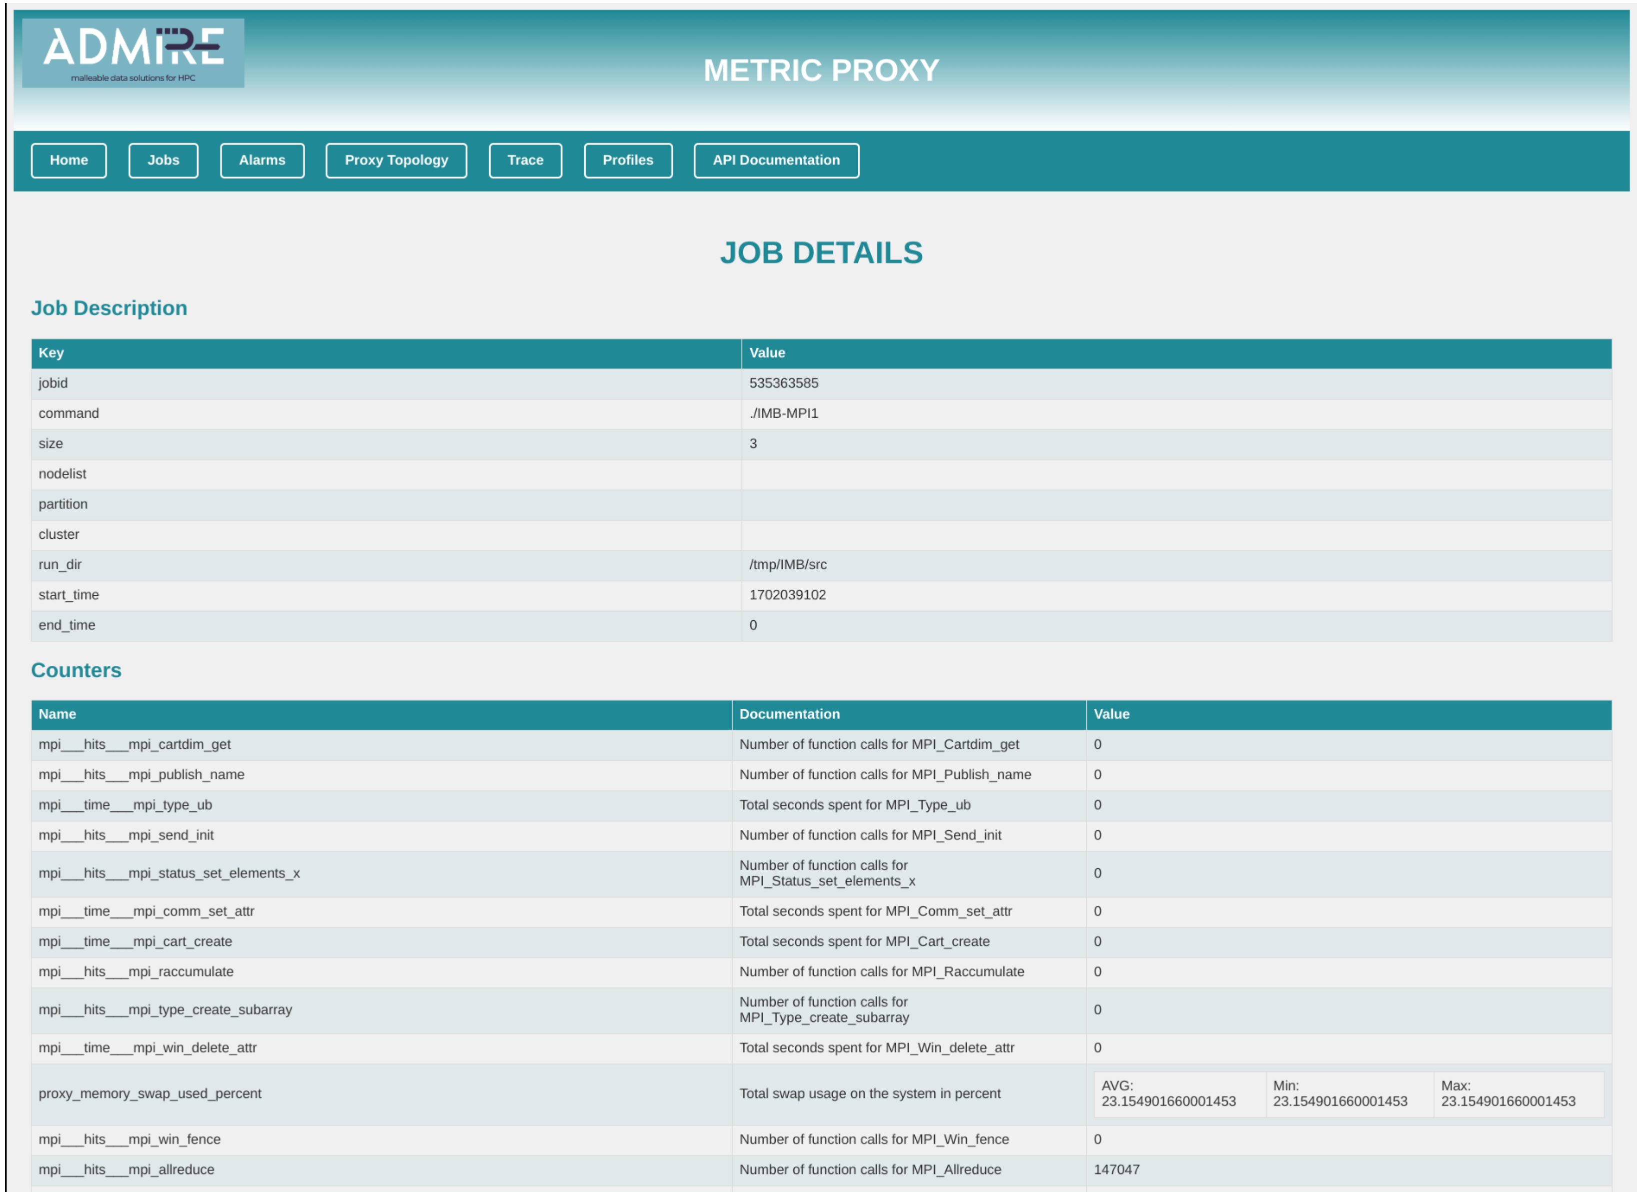

METRIC PROXY

JOB DETAILS

Job Description

Key	Value
jobid	535363585
command	./IMB-MPI1
size	3
nodelist	
partition	
cluster	
run_dir	/tmp/IMB/src
start_time	1702039102
end_time	0

Counters

Name	Documentation	Value
mpi___hits___mpi_cartdim_get	Number of function calls for MPI_Cartdim_get	0
mpi___hits___mpi_publish_name	Number of function calls for MPI_Publish_name	0
mpi___time___mpi_type_ub	Total seconds spent for MPI_Type_ub	0
mpi___hits___mpi_send_init	Number of function calls for MPI_Send_init	0
mpi___hits___mpi_status_set_elements_x	Number of function calls for MPI_Status_set_elements_x	0
mpi___time___mpi_comm_set_attr	Total seconds spent for MPI_Comm_set_attr	0
mpi___time___mpi_cart_create	Total seconds spent for MPI_Cart_create	0
mpi___hits___mpi_raccumulate	Number of function calls for MPI_Raccumulate	0
mpi___hits___mpi_type_create_subarray	Number of function calls for MPI_Type_create_subarray	0
mpi___time___mpi_win_delete_attr	Total seconds spent for MPI_Win_delete_attr	0
proxy_memory_swap_used_percent	Total swap usage on the system in percent	AVG: 23.154901660001453 Min: 23.154901660001453 Max: 23.154901660001453
mpi___hits___mpi_win_fence	Number of function calls for MPI_Win_fence	0
mpi___hits___mpi_allreduce	Number of function calls for MPI_Allreduce	147047

Real-Time Job Tracing

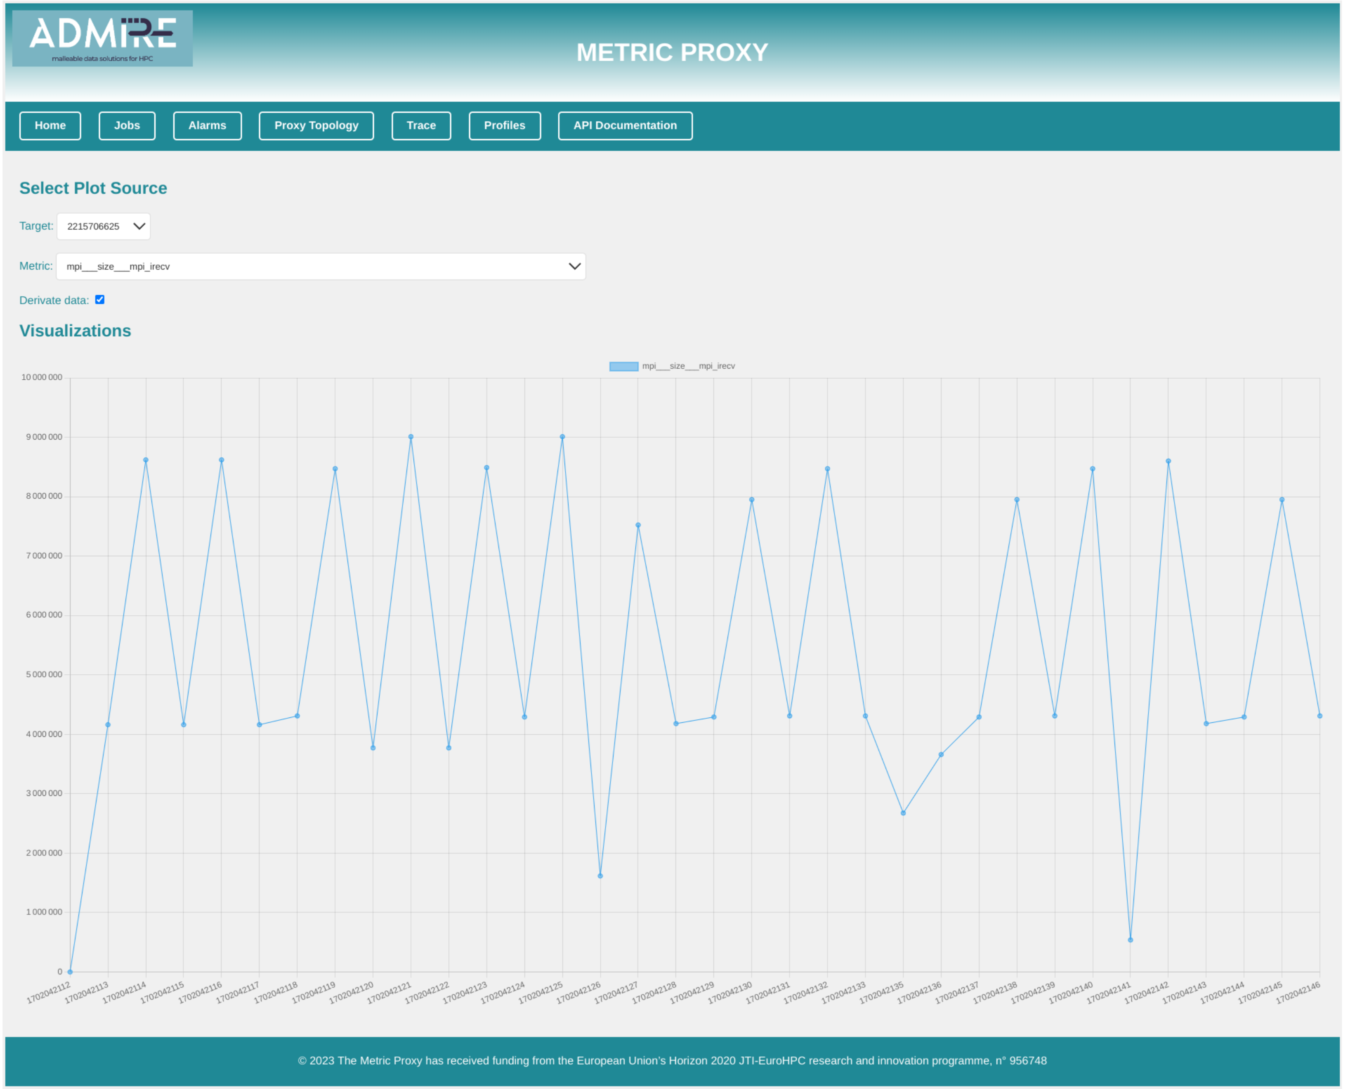

Select Plot Source

Target: 2215706625

Metric: mpi___size___mpi_irecv

Derivate data: ☑

Visualizations

© 2023 The Metric Proxy has received funding from the European Union's Horizon 2020 JTI-EuroHPC research and innovation programme, n° 956748

Fixed size of 32MB maximum per job, filled with sample every one second and slowing down by a factor 2 on resampling.

This maintains full-range traces with dynamically decreasing resolution and bounded size.

Real-Time Job Tracing

Period 1

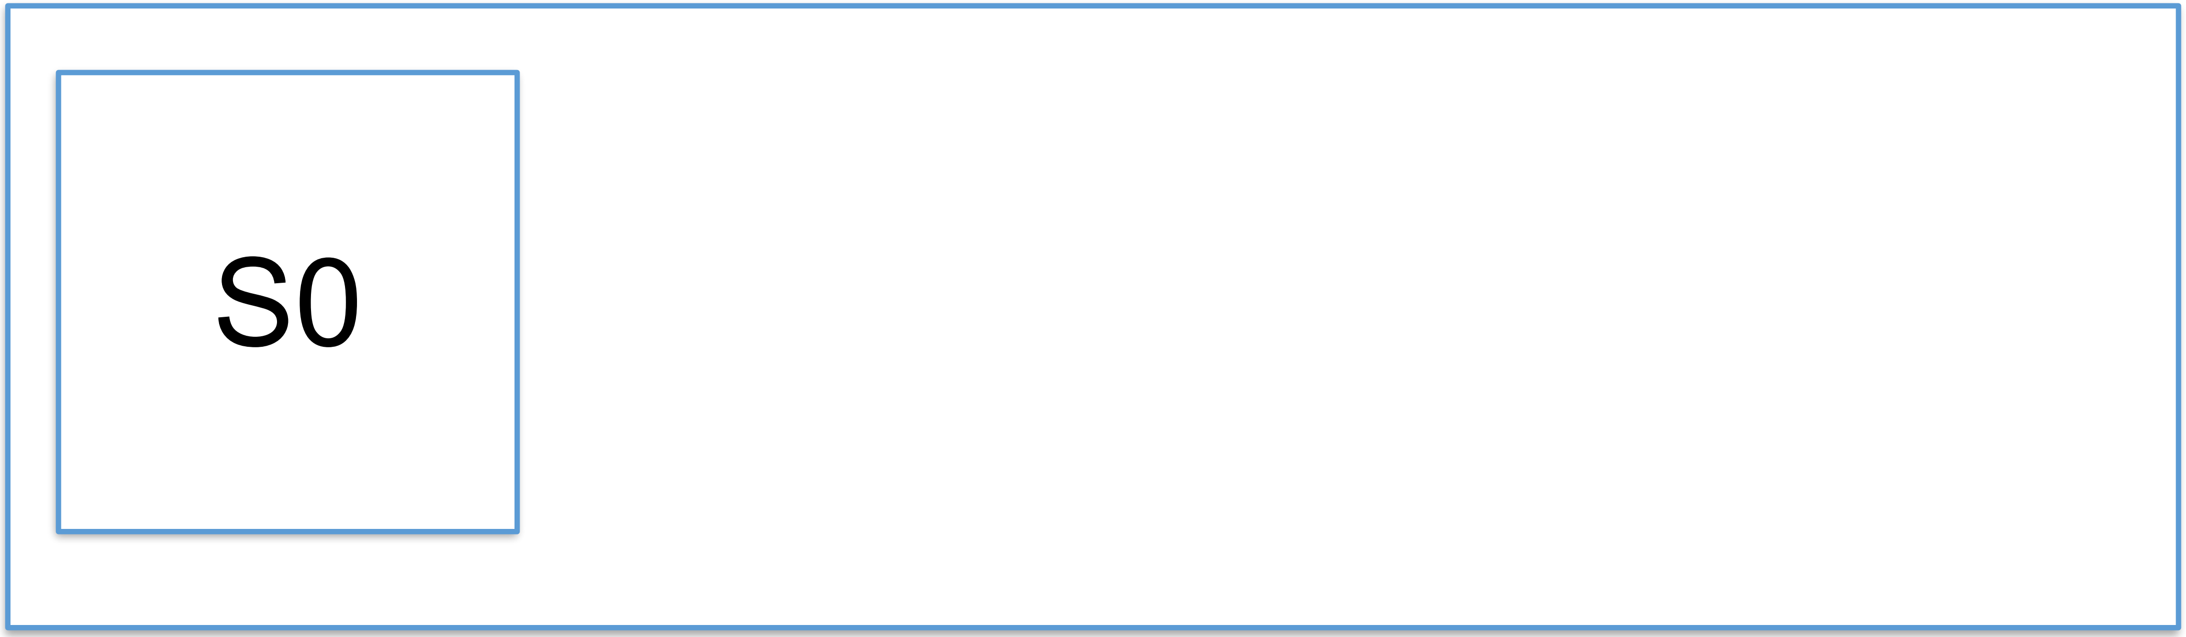

S0

Real-Time Job Tracing

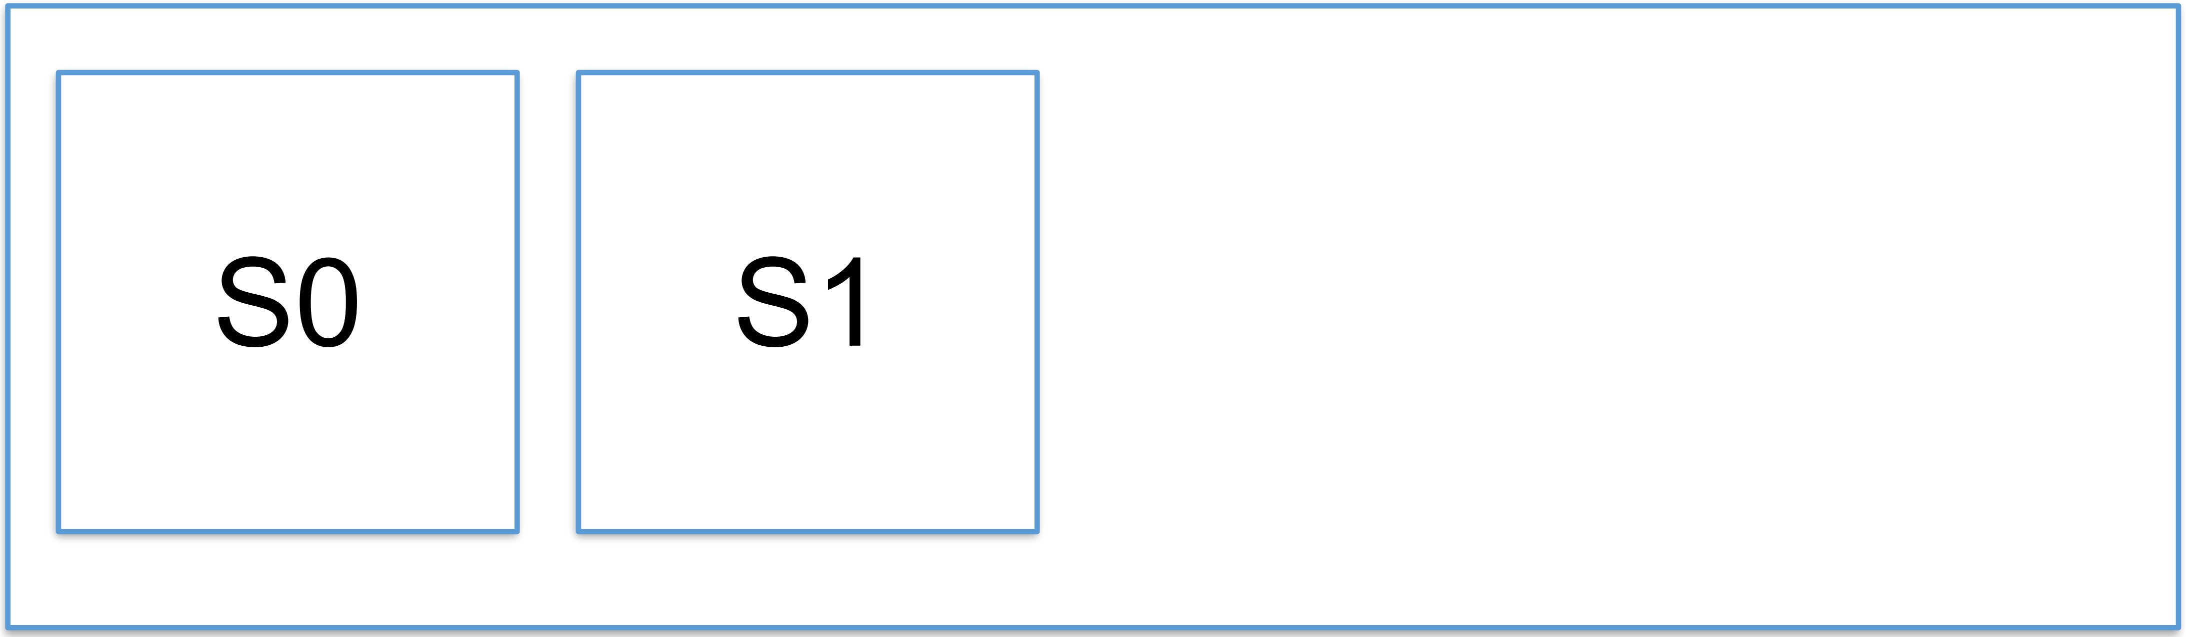

S0

S1

Real-Time Job Tracing

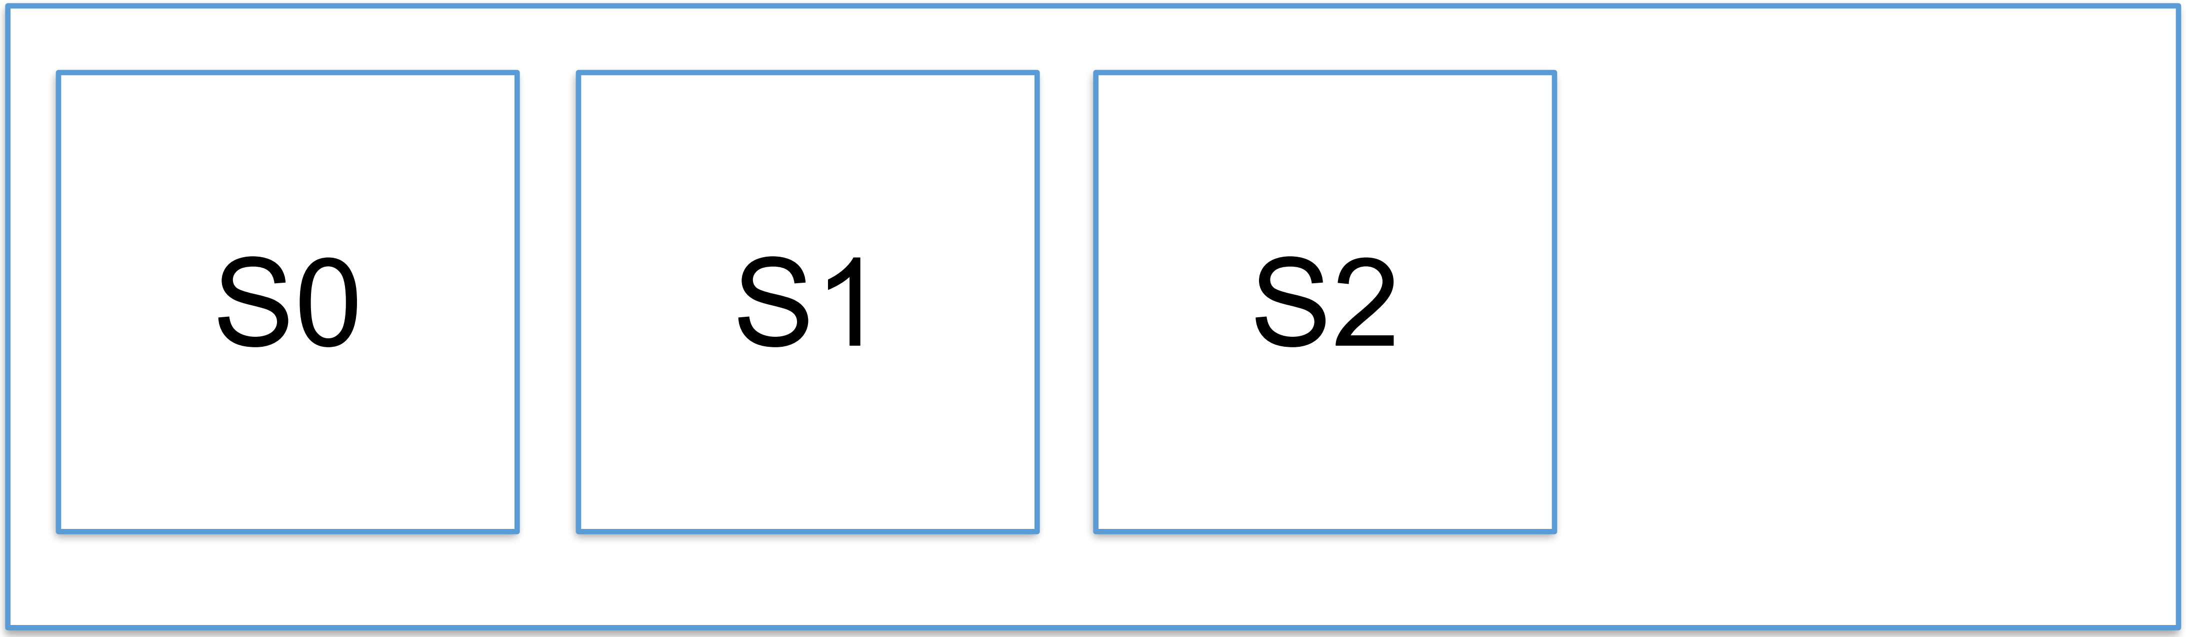

Real-Time Job Tracing

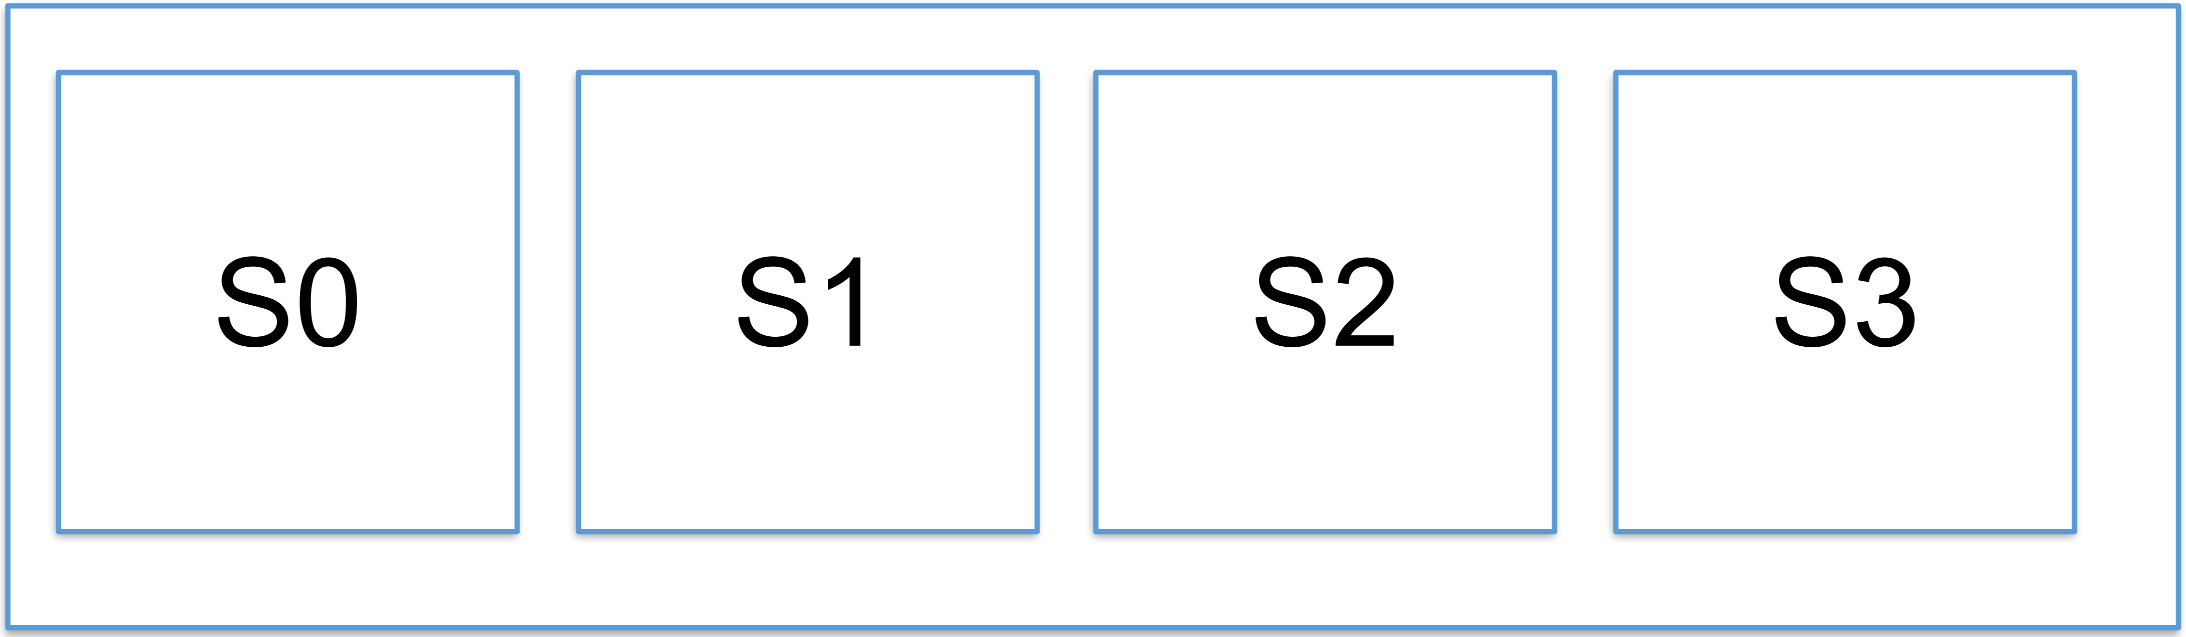

Real-Time Job Tracing

Period 2

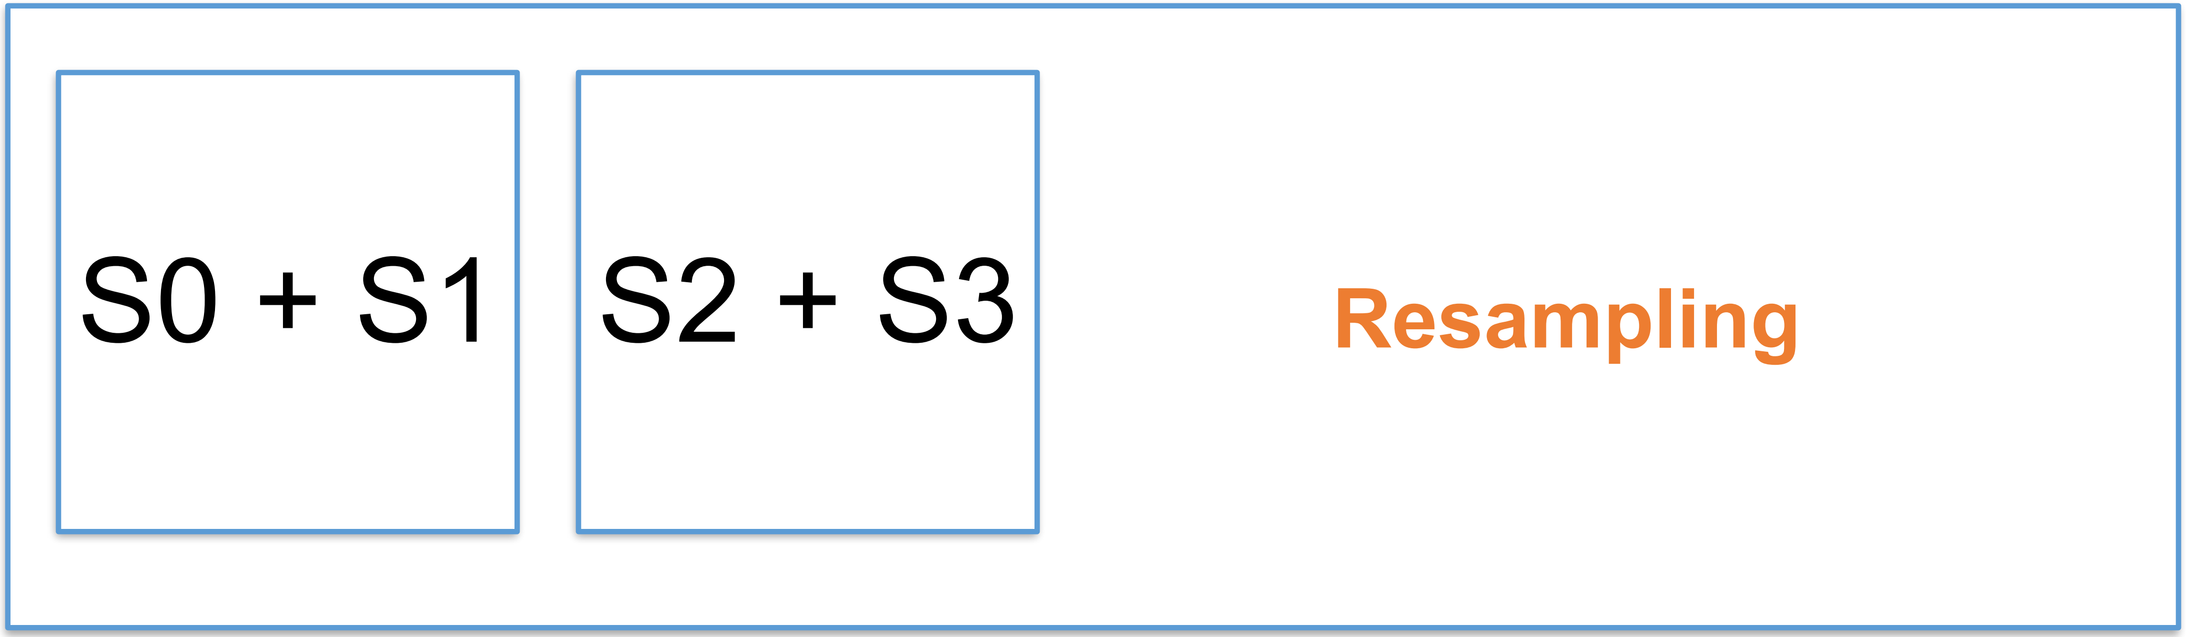

S0 + S1

S2 + S3

Resampling

Real-Time Job Tracing

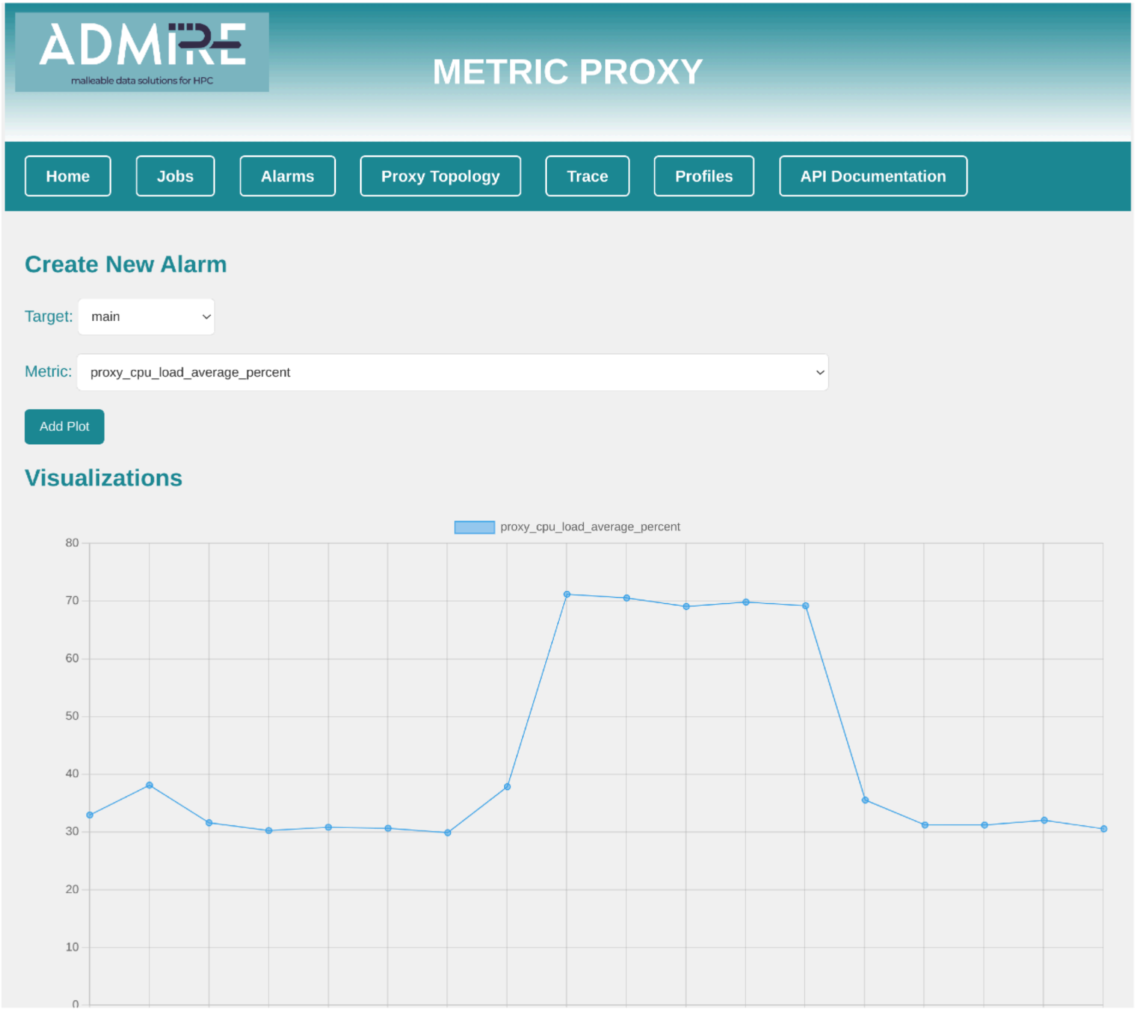

Create New Alarm

Target: main

Metric: proxy_cpu_load_average_percent

Visualizations

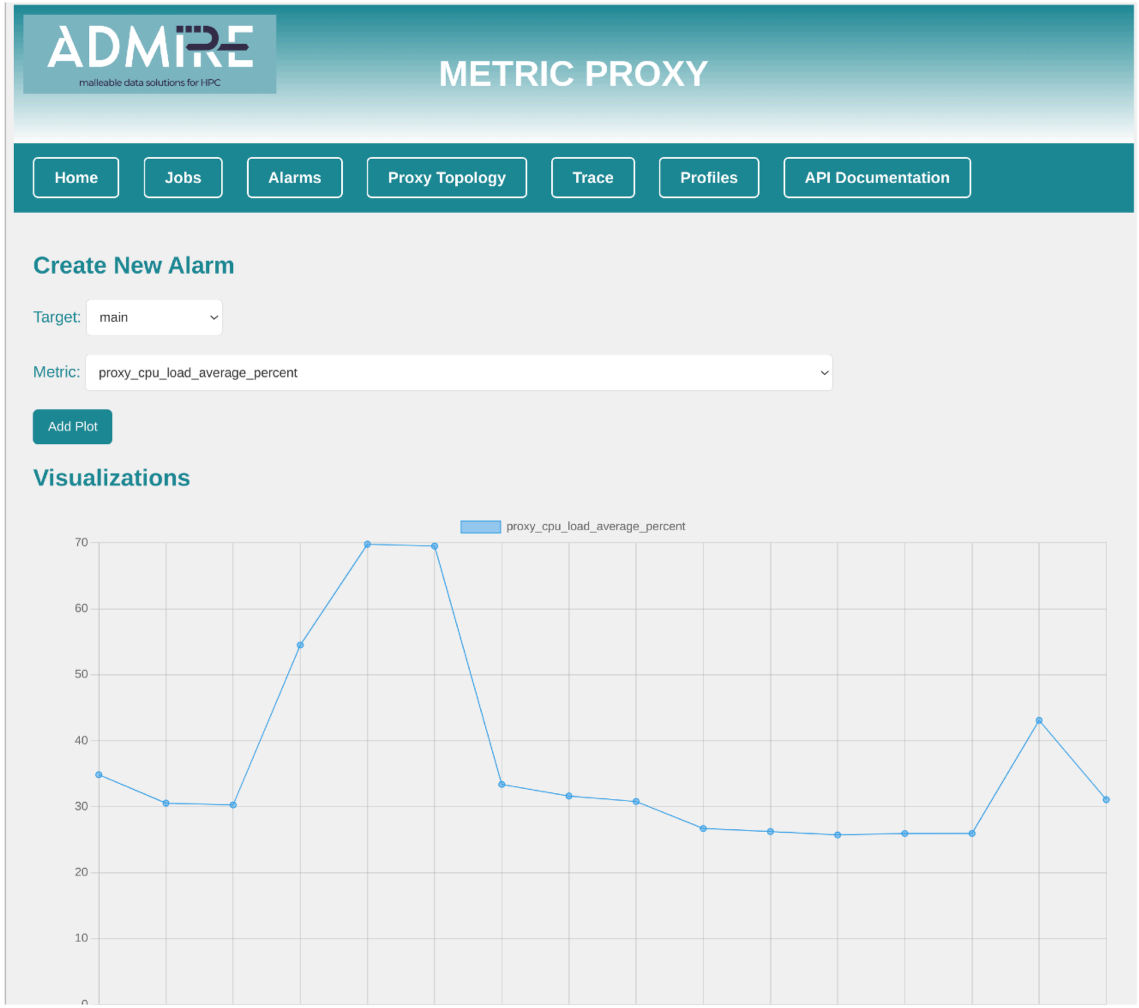

Create New Alarm

Target: main

Metric: proxy_cpu_load_average_percent

Visualizations

32 KB Limit

Real-Time Job Tracing

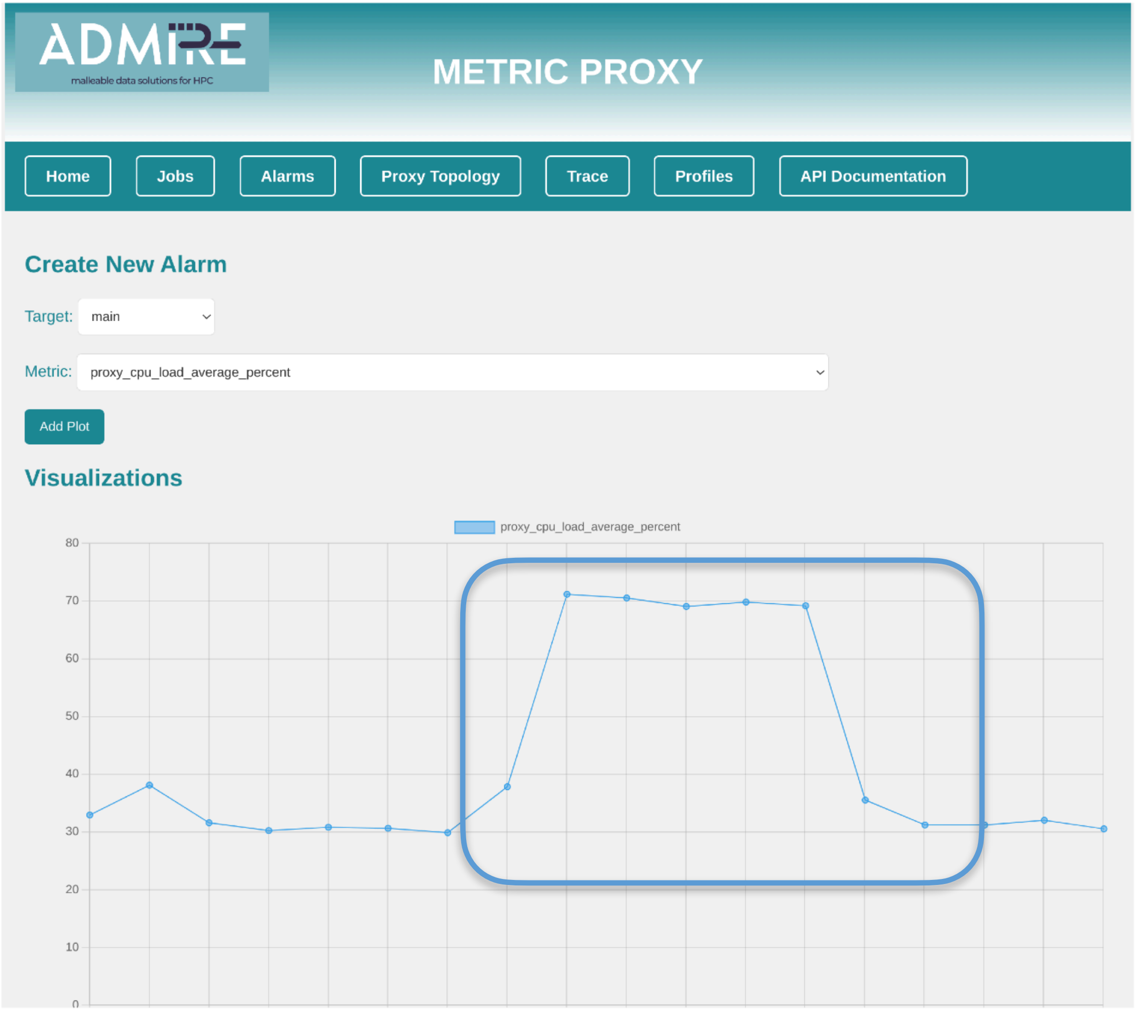

Create New Alarm

Target: main

Metric: proxy_cpu_load_average_percent

Visualizations

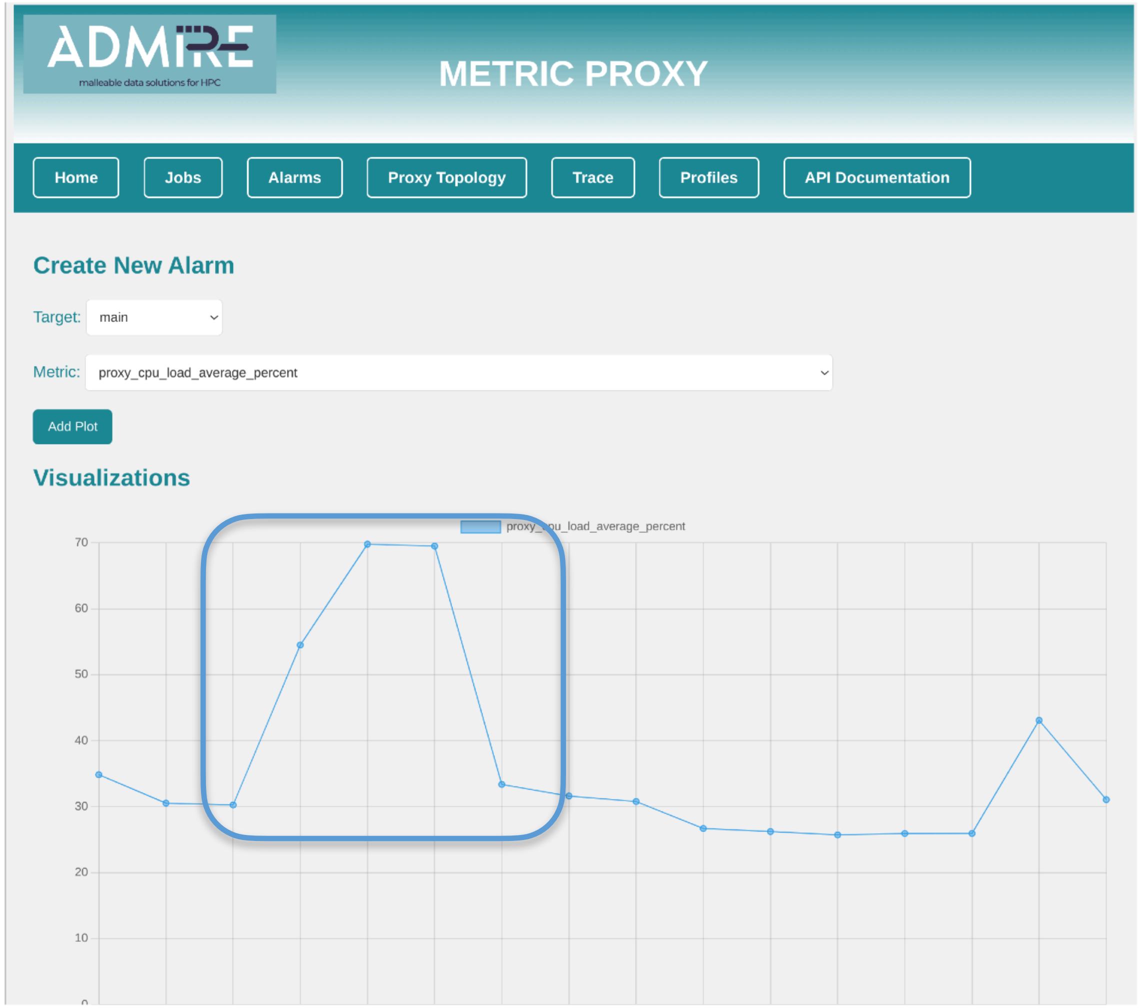

Create New Alarm

Target: main

Metric: proxy_cpu_load_average_percent

Visualizations

32 KB Limit

Job Profile Store

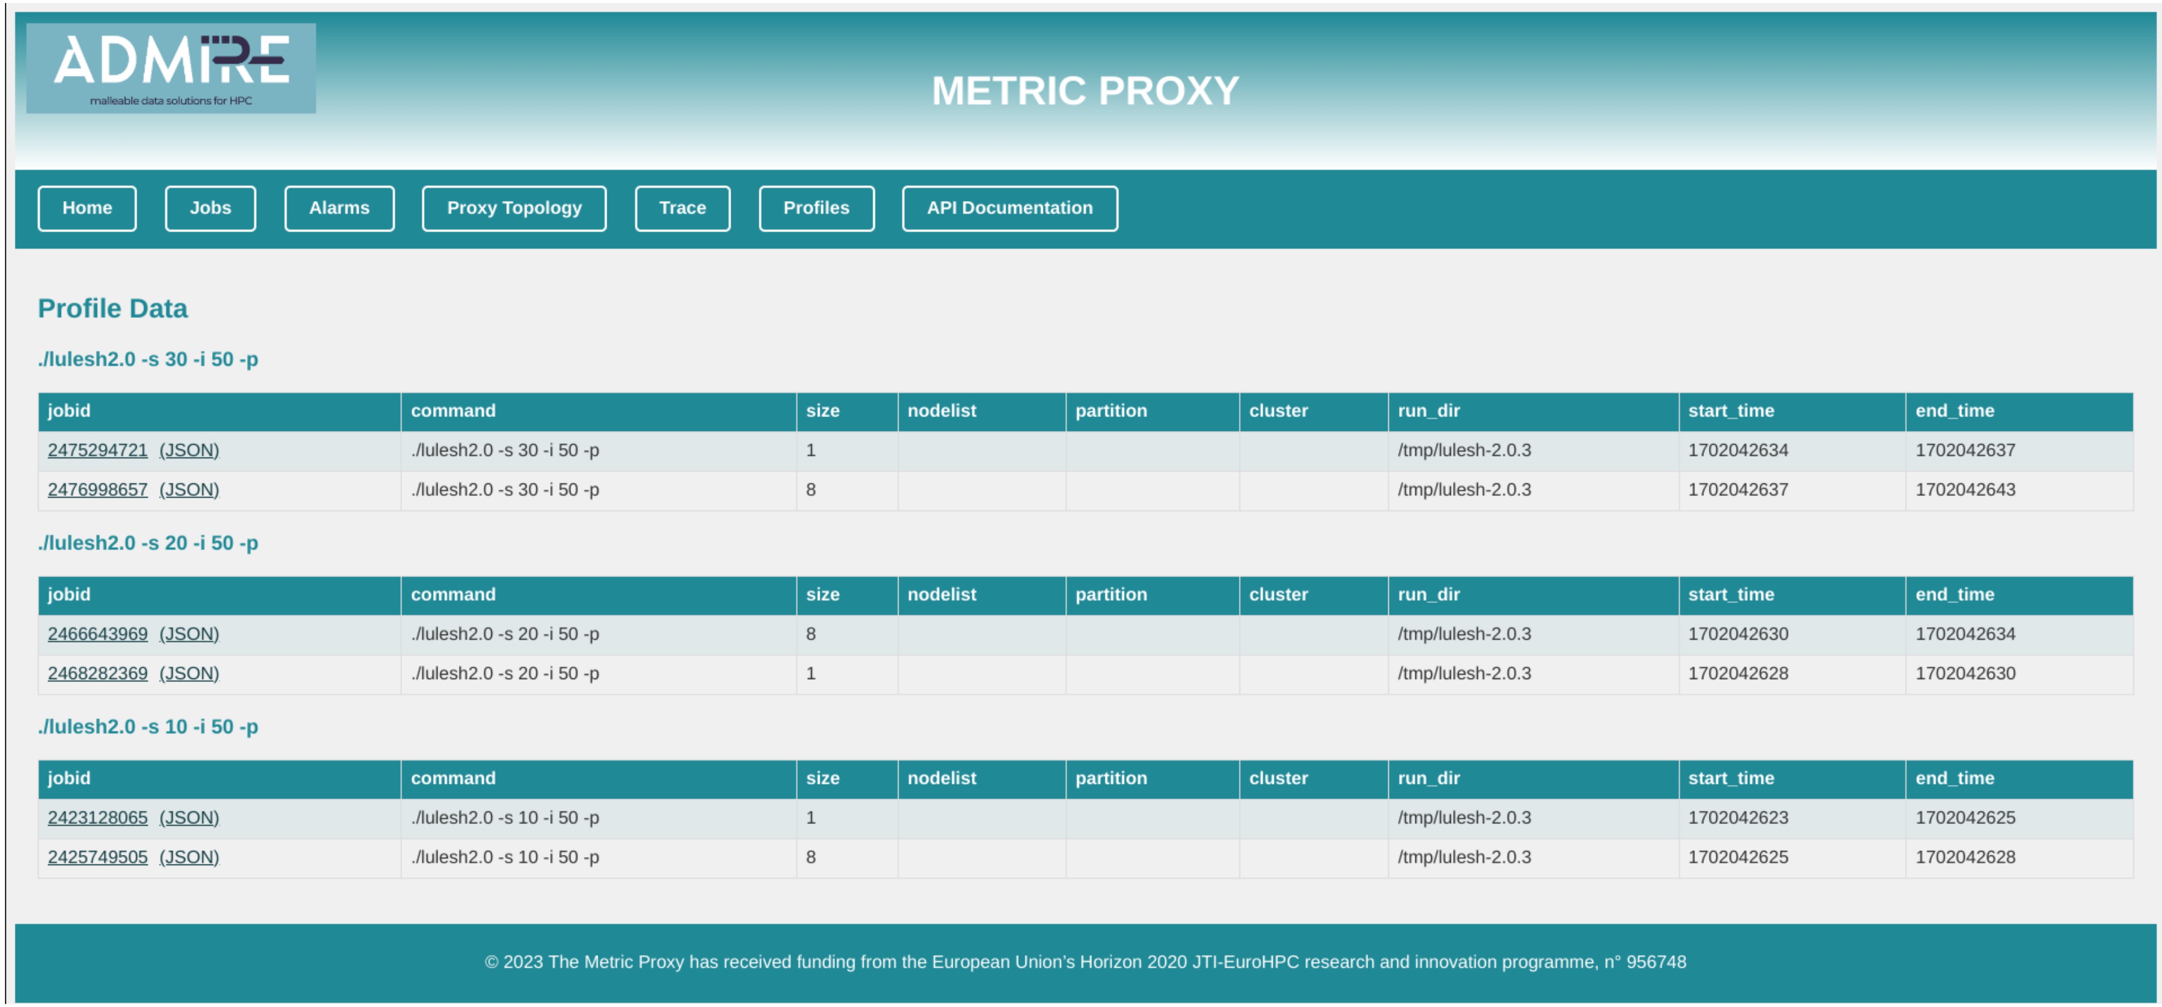

METRIC PROXY

Profile Data

./lulesh2.0 -s 30 -i 50 -p

jobid	command	size	nodelist	partition	cluster	run_dir	start_time	end_time
2475294721 (JSON)	./lulesh2.0 -s 30 -i 50 -p	1				/tmp/lulesh-2.0.3	1702042634	1702042637
2476998657 (JSON)	./lulesh2.0 -s 30 -i 50 -p	8				/tmp/lulesh-2.0.3	1702042637	1702042643

./lulesh2.0 -s 20 -i 50 -p

jobid	command	size	nodelist	partition	cluster	run_dir	start_time	end_time
2466643969 (JSON)	./lulesh2.0 -s 20 -i 50 -p	8				/tmp/lulesh-2.0.3	1702042630	1702042634
2468282369 (JSON)	./lulesh2.0 -s 20 -i 50 -p	1				/tmp/lulesh-2.0.3	1702042628	1702042630

./lulesh2.0 -s 10 -i 50 -p

jobid	command	size	nodelist	partition	cluster	run_dir	start_time	end_time
2423128065 (JSON)	./lulesh2.0 -s 10 -i 50 -p	1				/tmp/lulesh-2.0.3	1702042623	1702042625
2425749505 (JSON)	./lulesh2.0 -s 10 -i 50 -p	8				/tmp/lulesh-2.0.3	1702042625	1702042628

© 2023 The Metric Proxy has received funding from the European Union's Horizon 2020 JTI-EuroHPC research and innovation programme, n° 956748

Profile Modelling with Extra-P

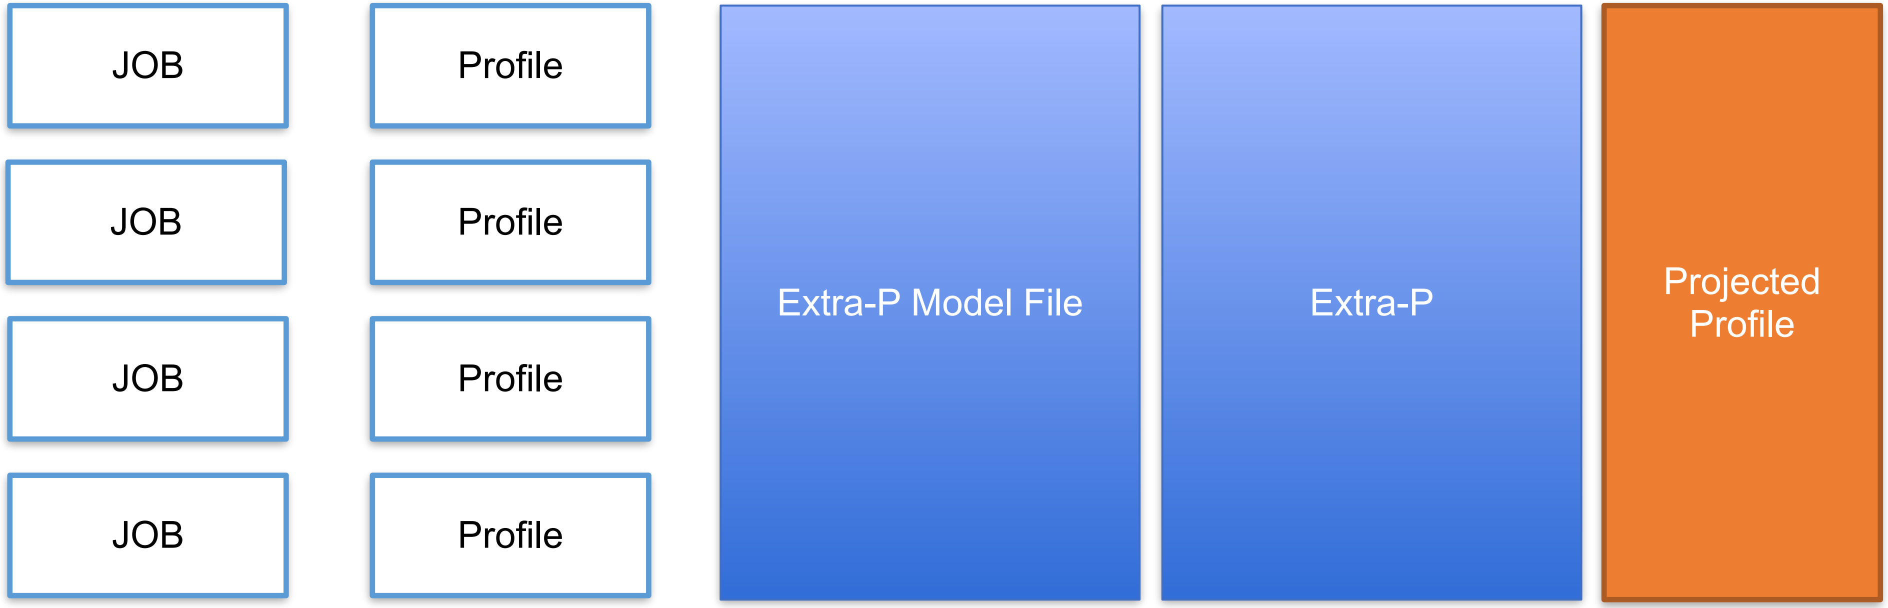

Complete Prometheus Integration

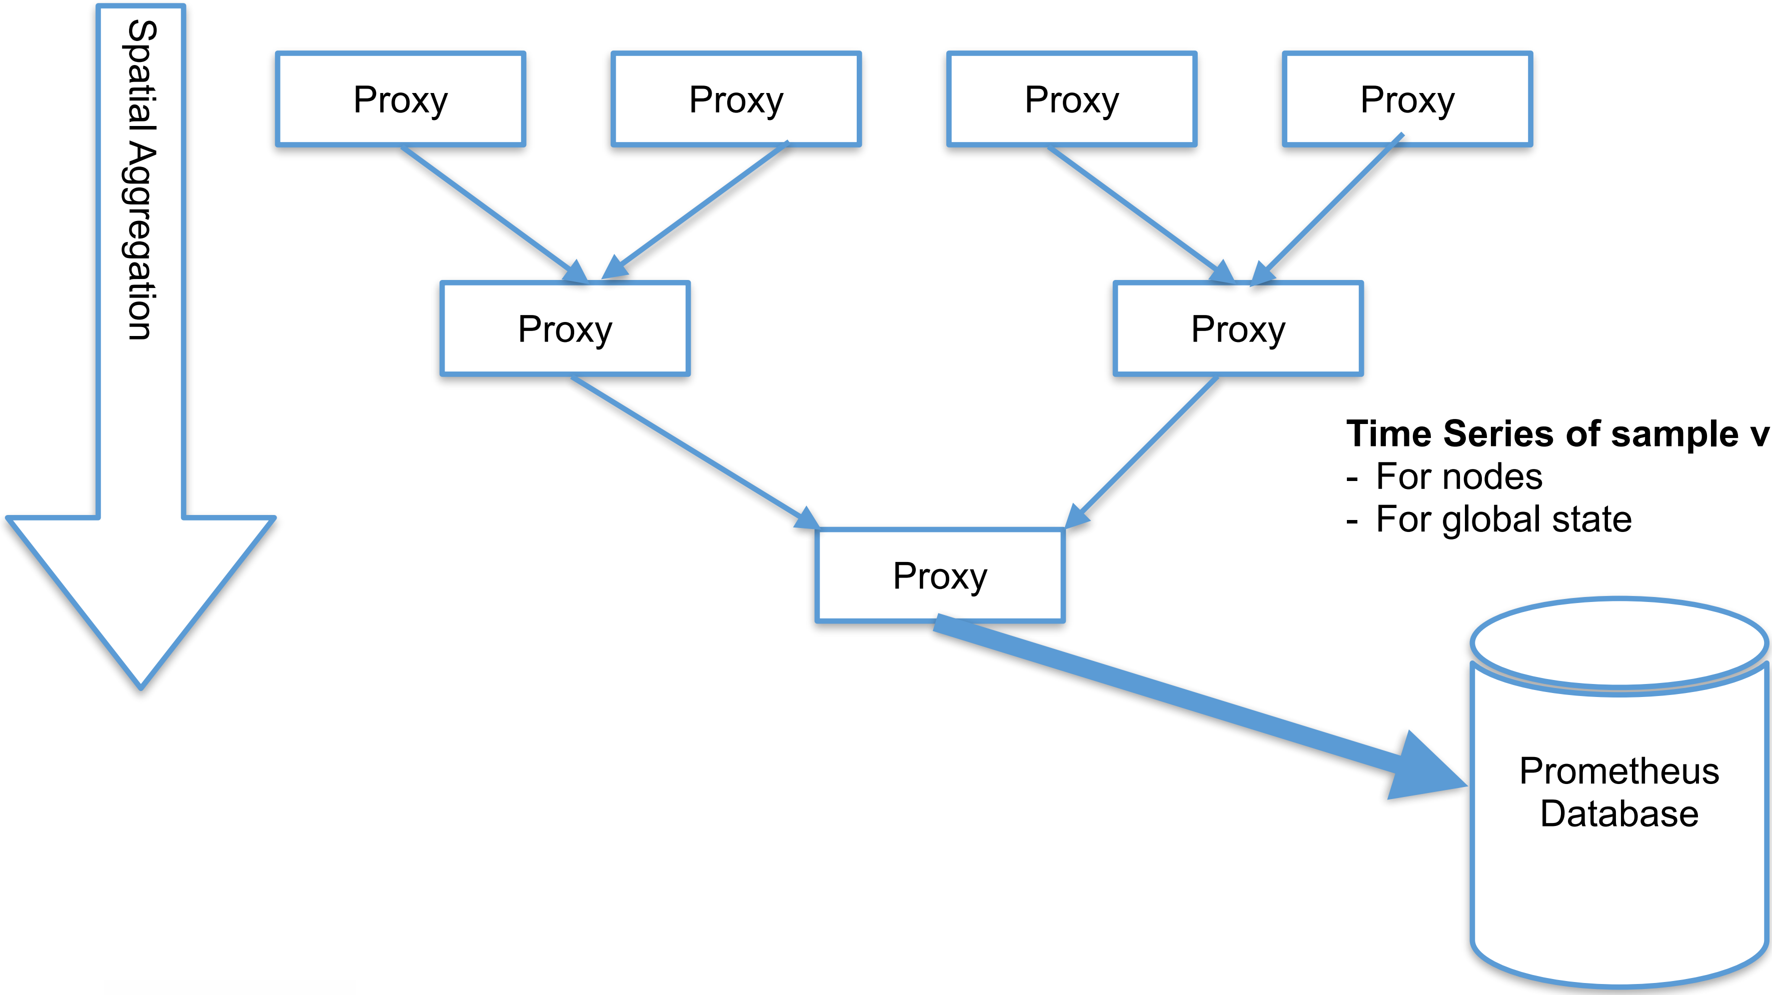

Time Series of sample values are accumulated over time:

- For nodes
- For global state

Sample Prometheus Outputs (Integrated View)

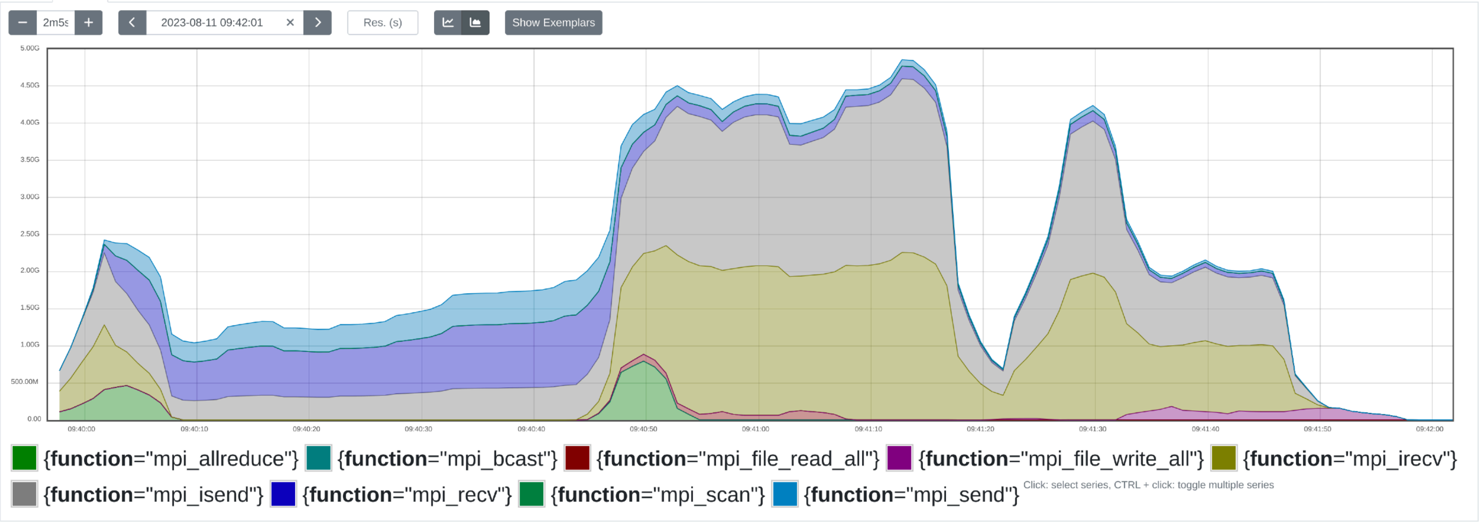

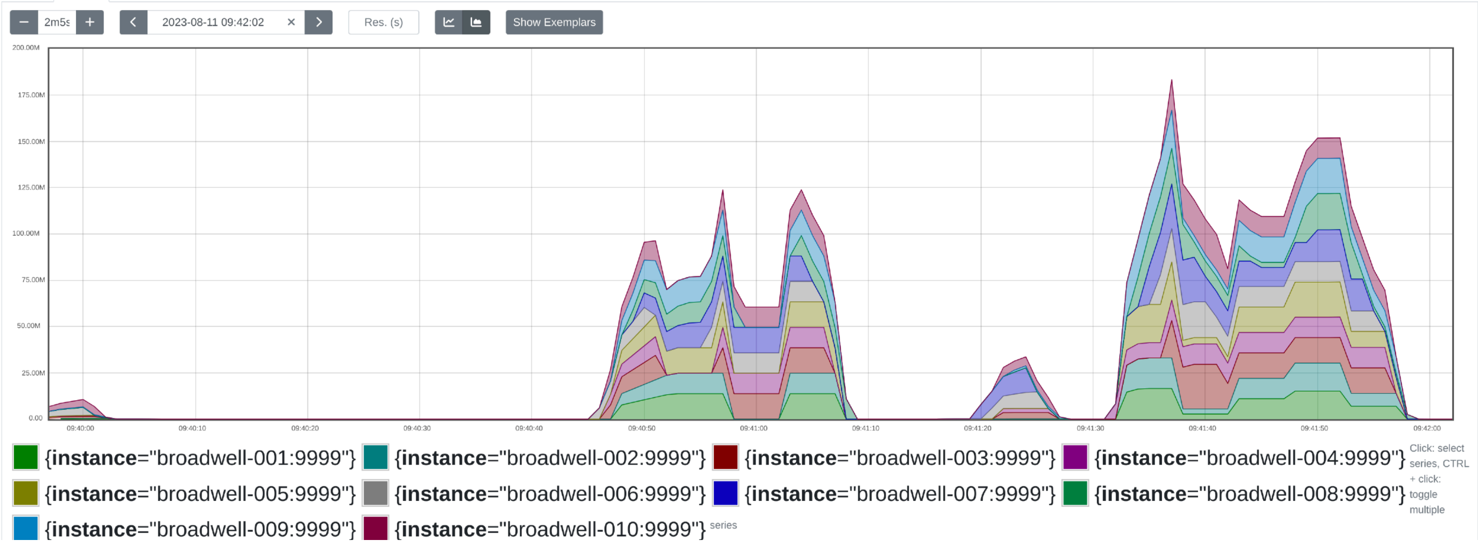
−
2m5s
+
<
2023-08-11 09:42:02
×
>
Res. (s)
Show Exemplars
200.00M
175.00M
150.00M
125.00M
100.00M
75.00M
50.00M
25.00M
0.00
09:40:00
09:40:10
09:40:20
09:40:30
09:40:40
09:40:50
09:41:00
09:41:10
09:41:20
09:41:30
09:41:40
09:41:50
09:42:00
🟩 {instance="broadwell-001:9999"}
🟩 {instance="broadwell-002:9999"}
🟥 {instance="broadwell-003:9999"}
🟪 {instance="broadwell-004:9999"}
🟨 {instance="broadwell-005:9999"}
⬜ {instance="broadwell-006:9999"}
🟦 {instance="broadwell-007:9999"}
🟩 {instance="broadwell-008:9999"}
🟦 {instance="broadwell-009:9999"}
🟥 {instance="broadwell-010:9999"}
Click: select series, CTRL + click: toggle multiple series

Sample Grafana Output

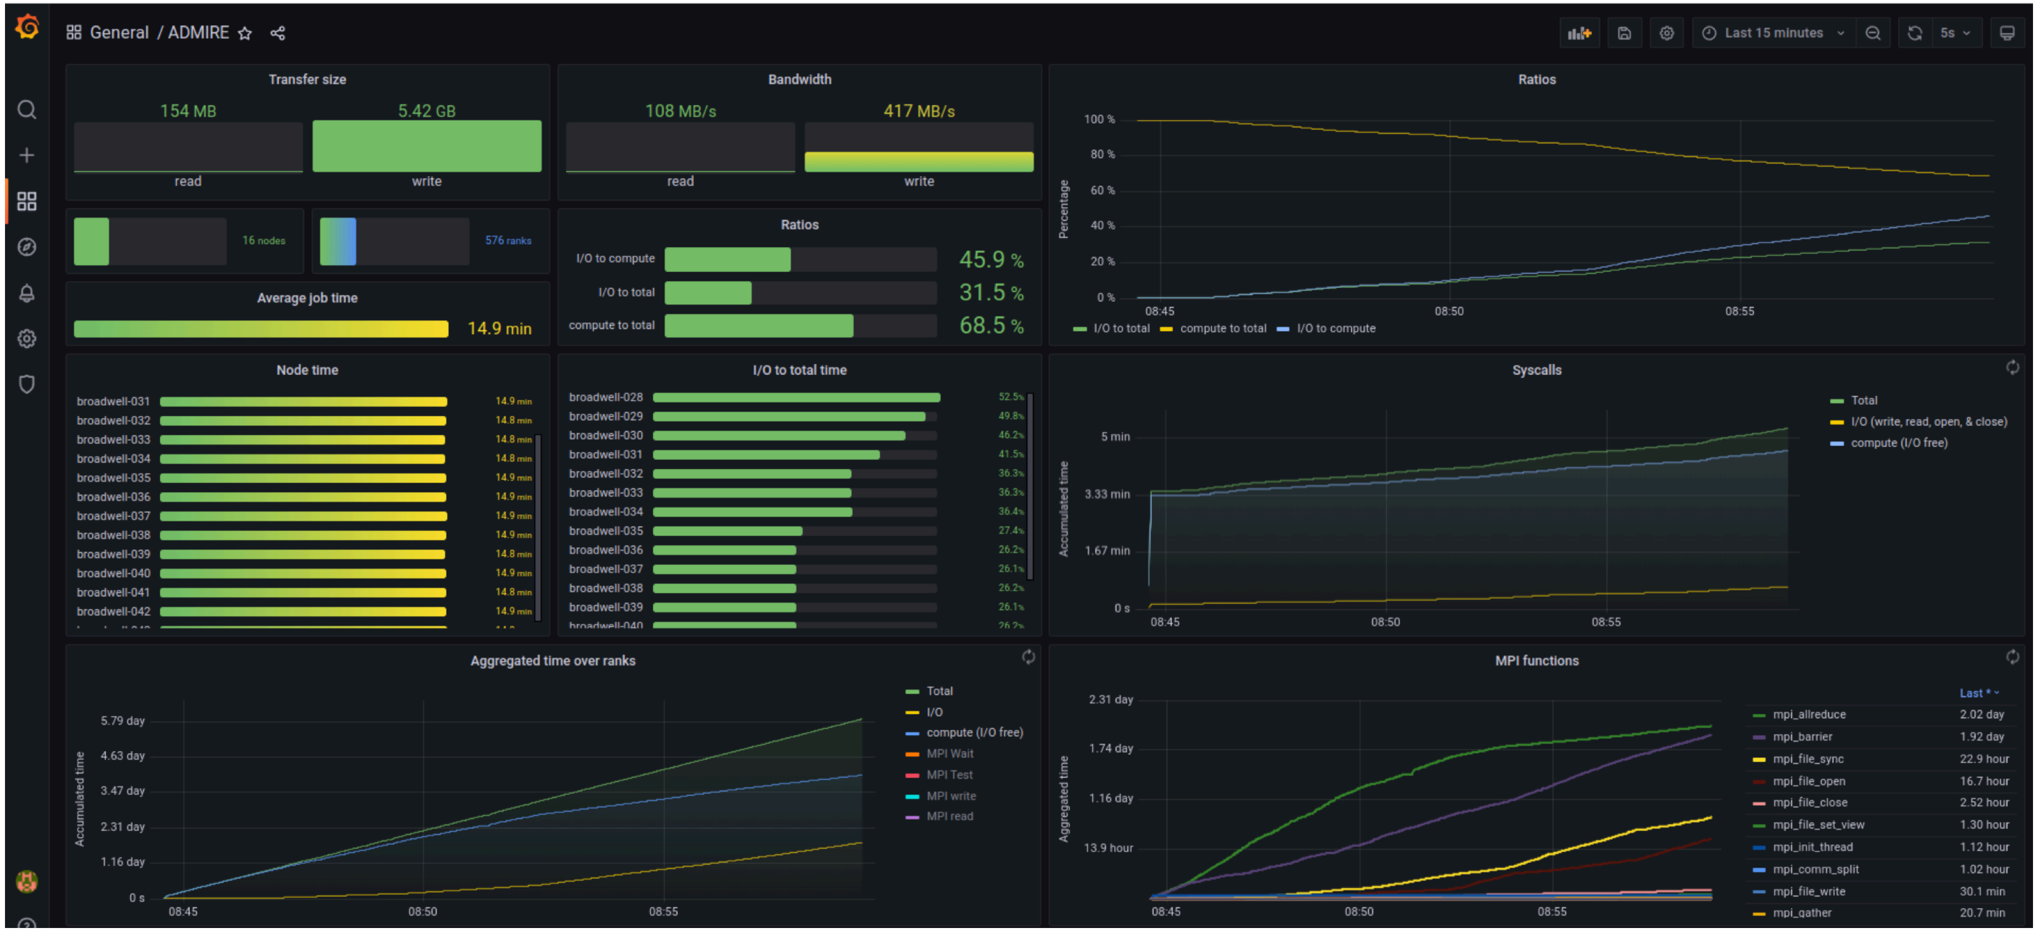

Node time

broadwell-031	14.9 min
broadwell-032	14.8 min
broadwell-033	14.8 min
broadwell-034	14.8 min
broadwell-035	14.9 min
broadwell-036	14.9 min
broadwell-037	14.9 min
broadwell-038	14.9 min
broadwell-039	14.8 min
broadwell-040	14.9 min
broadwell-041	14.8 min
broadwell-042	14.9 min

I/O to total time

broadwell-028	52.5%
broadwell-029	49.8%
broadwell-030	46.2%
broadwell-031	41.5%
broadwell-032	36.3%
broadwell-033	36.3%
broadwell-034	36.4%
broadwell-035	27.4%
broadwell-036	26.2%
broadwell-037	26.1%
broadwell-038	26.2%
broadwell-039	26.1%
broadwell-040	26.2%

MPI functions

	Last *
mpi_allreduce	2.02 day
mpi_barrier	1.92 day
mpi_file_sync	22.9 hour
mpi_file_open	16.7 hour
mpi_file_close	2.52 hour
mpi_file_set_view	1.30 hour
mpi_init_thread	1.12 hour
mpi_comm_split	1.02 hour
mpi_file_write	30.1 min
mpi_gather	20.7 min

Overhead Assessment (on V1)

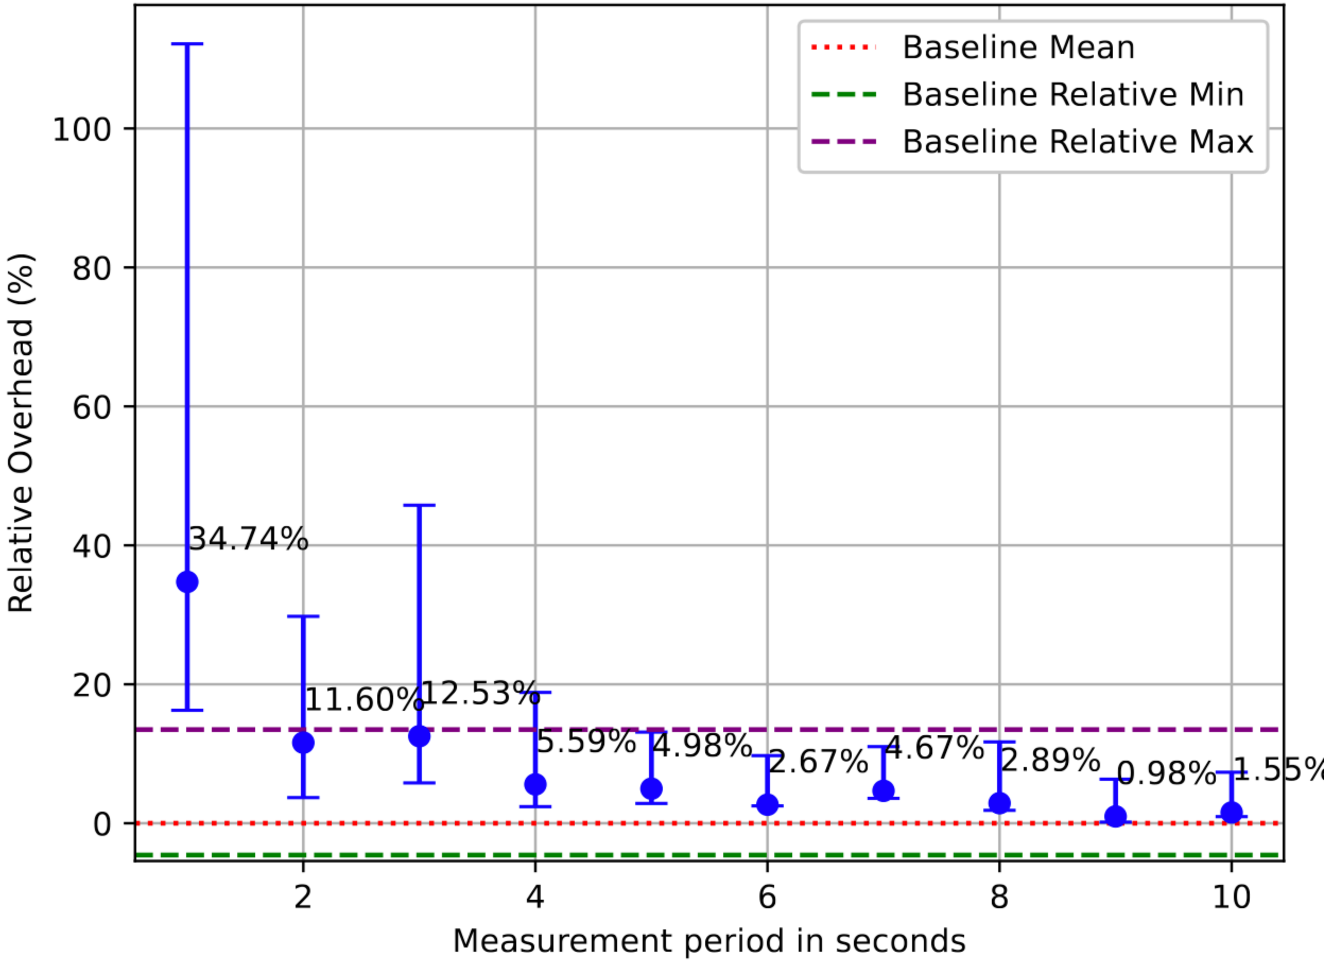

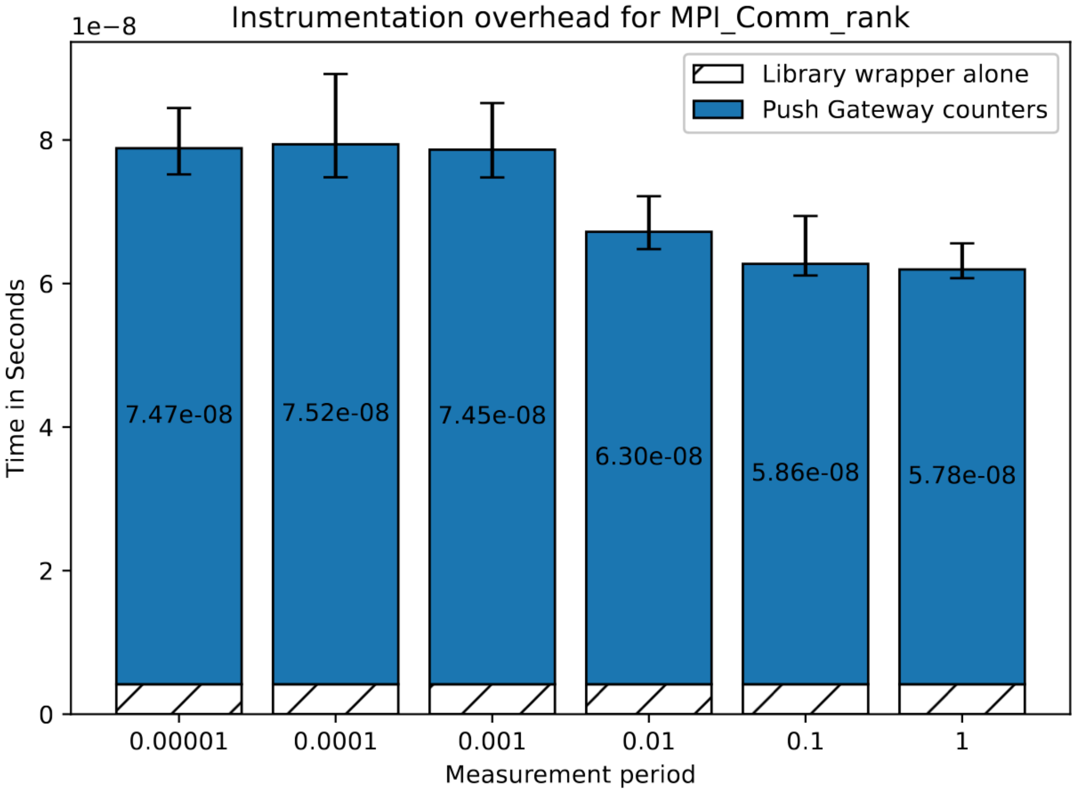

Source Code & Demo

Try it at:

http://github.com/besnardjb/proxy_v2/

Demo

Conclusion

- **We presented the ADMIRE Metric proxy**
  - **Implements an aggregating prometheus push gateway (and more)**
  - **Made in Rust**
- **We have machine wide monitoring capabilities:**
  - **Profiles (per node, per job)**
  - **Traces(per job, per node)**
  - **Machine wide state in real time (0.5 sec resolution)**
- **We work on modelling capabilities thanks to Extra-P and FTIO (WIP)**
- **We will make a first official release in the near future after more testing at scale on the ADMIRE testing supercomputer (thanks to UNITO)**
- **Code available at: http://github.com/besnardjb/proxy_v2/**

Conclusion / Questions

Try it at:

http://github.com/besnardjb/proxy_v2/

ADMIRE

malleable data solutions for HPC

Adaptive multi-tier intelligent data manager for Exascale

EuroHPC Joint Undertaking

admire-eurohpc.eu

ADMIRE Users Day

Metric Proxy: Enabling real-time measurement at Supercomputer Scale

Jean-Baptiste Besnard, ParaTools SAS

December 12th 2023.

Barcelona Supercomputing Center

Grant Agreement number: 956748 — ADMIRE — H2020-JTI-EuroHPC-2019-1