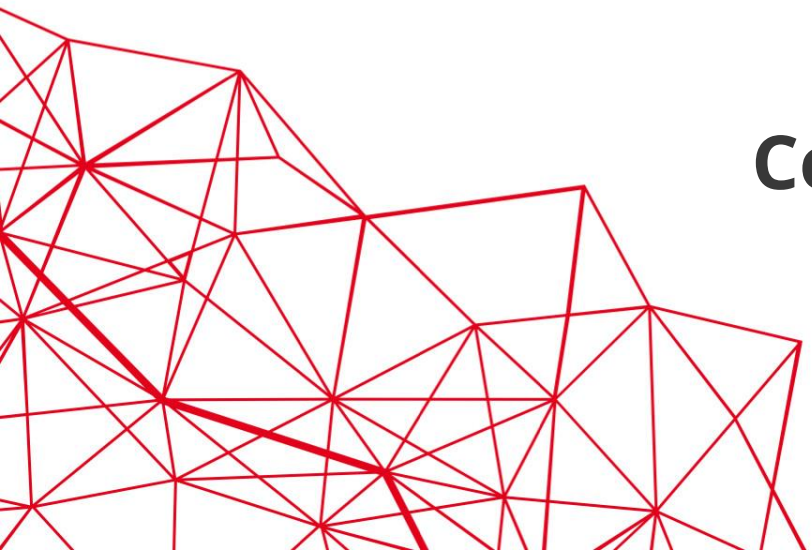

REINVENTING HPC

ISC High Performance

MAY 12 – 16, 2024 | HAMBURG, GERMANY

Compute Energy & Emissions Monitoring Stack (CEEMS)

Mahendra Paipuri
Research Engineer, IDRIS, CNRS

5th ISC HPC International Workshop on Monitoring & Operational Data Analytics
16th May 2024

Jean Zay - HPE/ATOS
123.6 PFLOPS

CPU Partition

Intel 6248
720 nodes
28 800 cores
192 GiB/node
2.3 PFLOPS

V100 Partition

Intel 6248
265 nodes with 4 Nvidia V100 32 GiB
131 nodes with 4 Nvidia V100 16 GiB
31 nodes with 8 Nvidia V100 32 GiB
1832 GPUs
15.6 PFLOPS

A100 Partition

AMD Milan
52 nodes with 8 Nvidia A100 80 GiB
416 GPUs.
8.1 PFLOPS

H100 Partition

Intel Sapphire Rapids
364 nodes with 4 Nvidia H100 80 GiB
1452 GPUs
87.6 PFLOPS
Currently being commissioned

Estimating Power: Top-Down

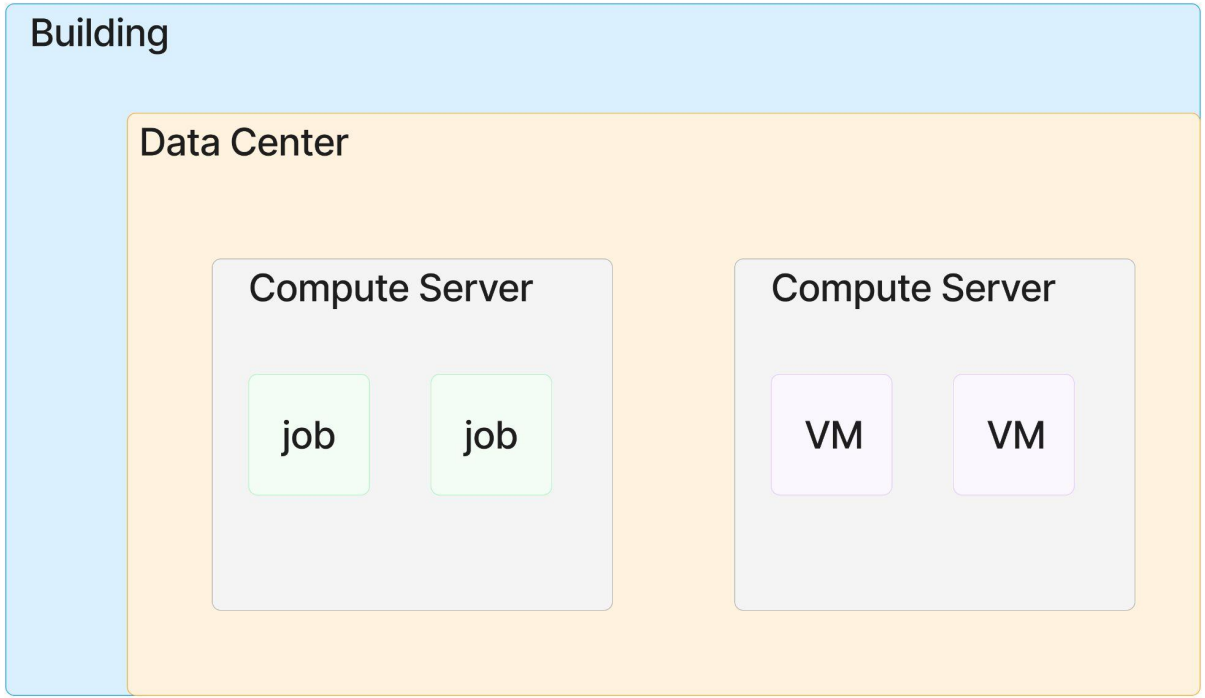

Estimating Power: Bottom-Up

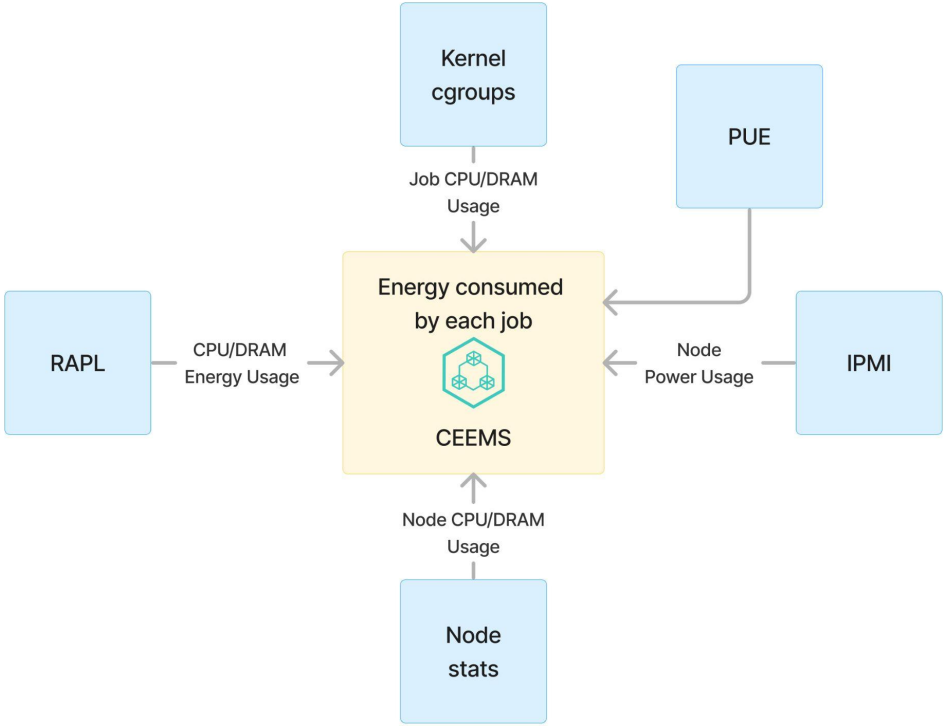

Core Architecture

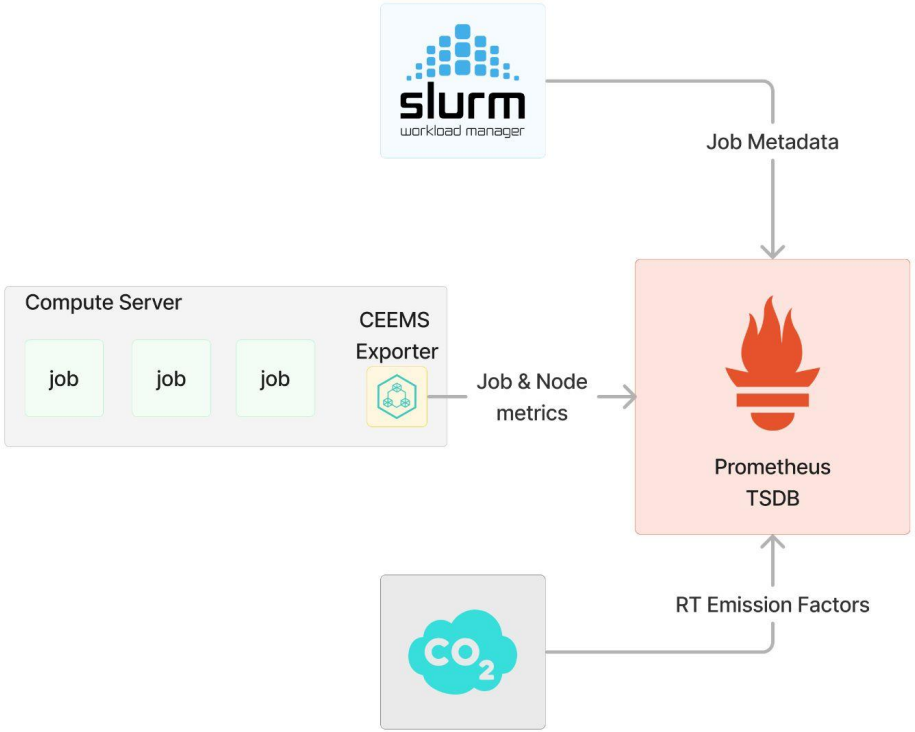

Real time Emission Factors

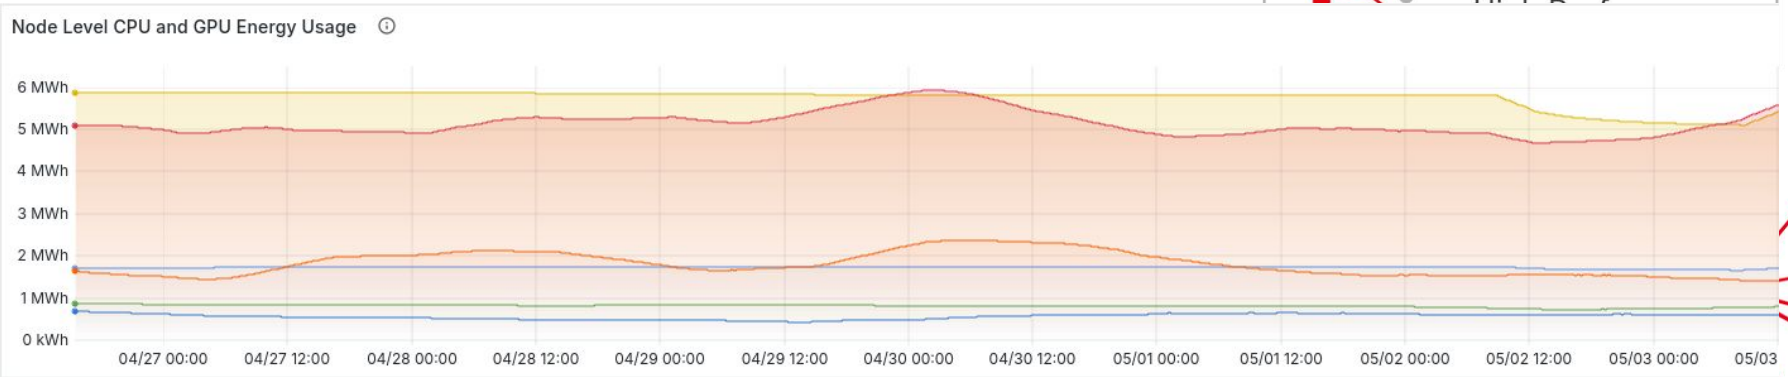

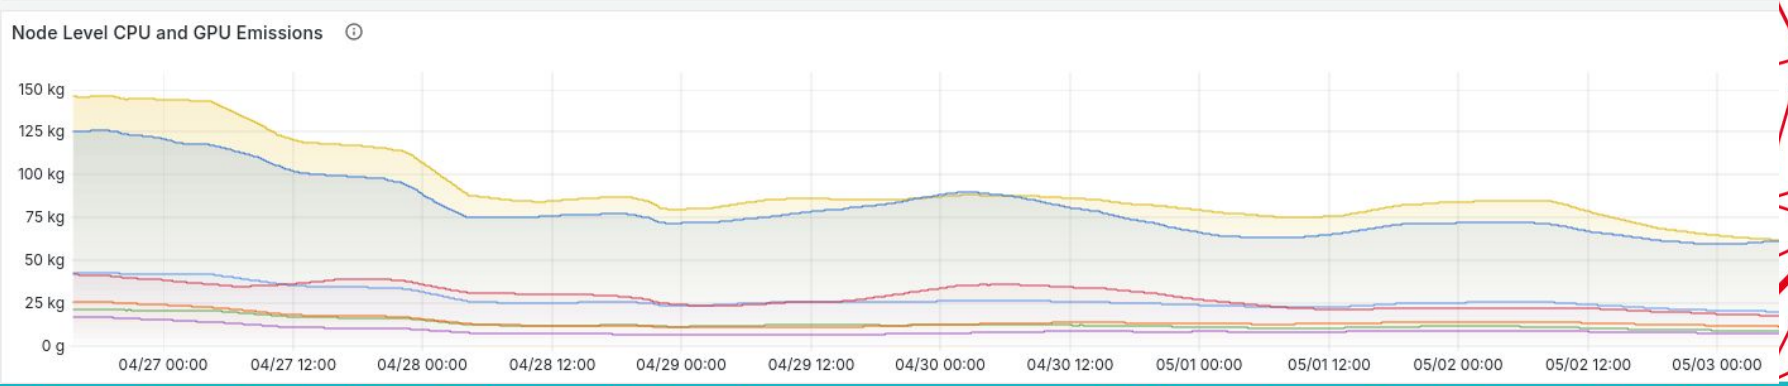

Real time Emission Factors

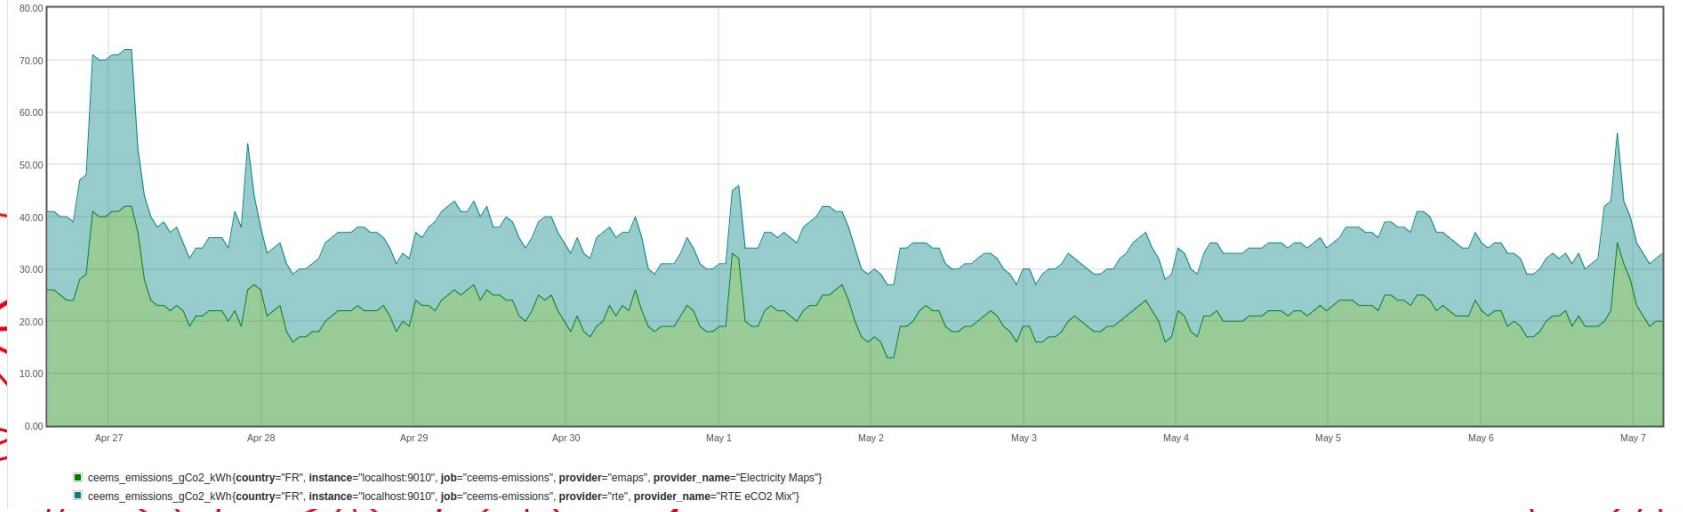

Core Architecture

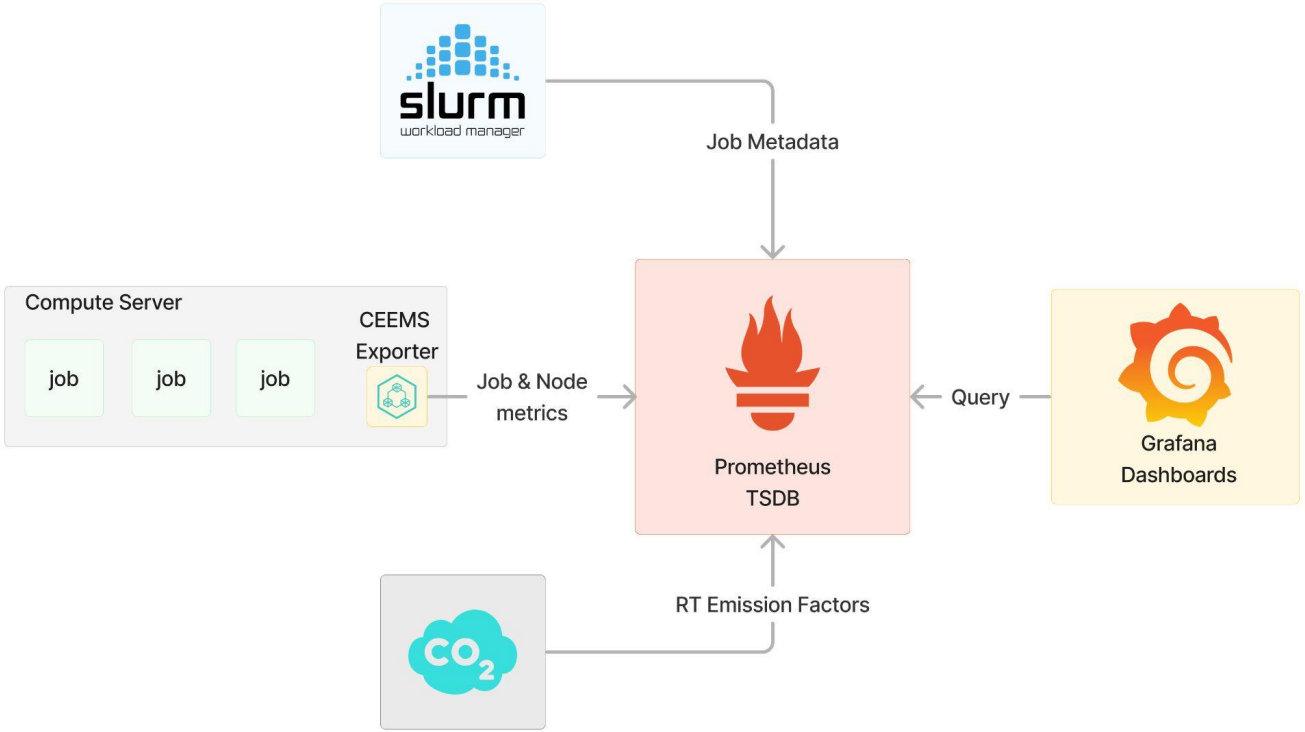

User dashboards

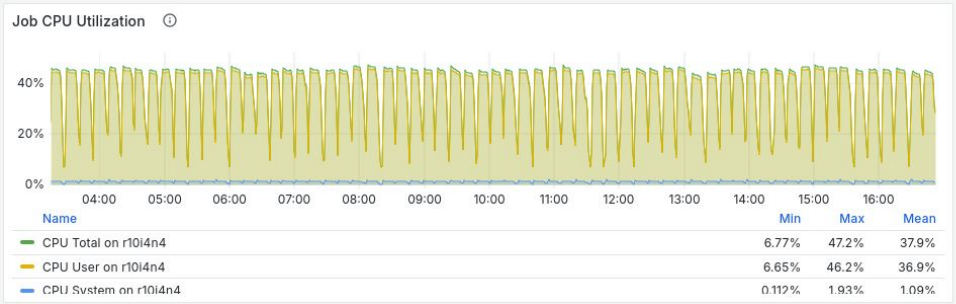

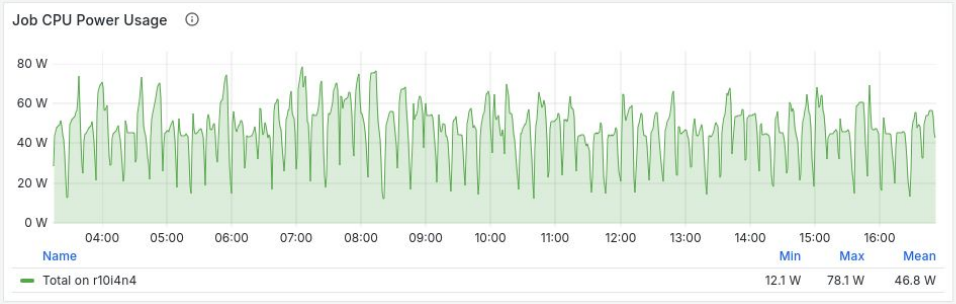

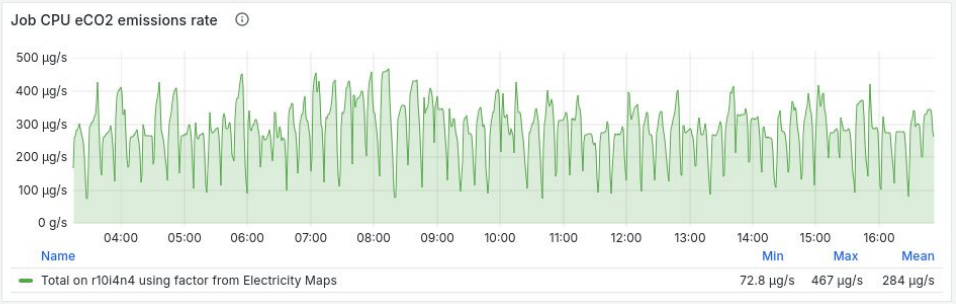

User dashboards

JobID	Start Time	End Time	Elapsed	State	CPU Usage (%)	CPU Energy (kWh	CPU Emissions (g	GPU Usage (%)	GPU Energy (kWh	GPU Emissions (g
1735381	2024-05-01T16:...	2024-05-01T17:...	00:20:04	TIMEOUT	12.4	0.228	3.88	84.3	0.803	13.7
1735377	2024-05-01T16:...	2024-05-01T16:...	00:20:12	TIMEOUT	11.2	0.209	3.55	72.3	0.762	13.0
1735217	2024-05-01T16:...	2024-05-01T16:...	00:20:06	TIMEOUT	11.7	0.0762	1.29	79.3	0.745	12.7
1735216	2024-05-01T16:...	2024-05-01T16:...	00:20:37	TIMEOUT	12.1	0.204	3.46	82.8	0.769	13.1
1735210	2024-05-01T16:...	2024-05-01T16:...	00:20:06	TIMEOUT	10.0	0.221	3.76	64.6	0.759	12.9
1735194	2024-05-01T16:...	2024-05-01T16:...	00:20:09	TIMEOUT	8.35	0.223	3.80	57.5	0.775	13.2
1735177	2024-05-01T16:...	2024-05-01T16:...	00:05:18	CANCELLED by ...	9.07	0.0529	0.899	56.7	0.144	2.45
1735176	2024-05-01T15:...	2024-05-01T16:...	00:10:26	FAILED	8.82	0.107	1.82	53.8	0.313	5.32
1733909	2024-05-01T13:...	2024-05-01T13:...	00:20:19	TIMEOUT	11.7	0.237	3.80	80.3	0.780	12.5
1705935	2024-04-30T01...	2024-04-30T01...	00:04:22	COMPLETED	9.56	0.00569	0.0853	56	0.0118	0.178
1705934	2024-04-30T01...	2024-04-30T01...	00:06:03	COMPLETED	9.78	0.00728	0.109	100	0.0330	0.495
1705840	2024-04-30T0...	2024-04-30T01...	00:04:16	COMPLETED	10.3	0.00549	0.0824	87.5	0.0218	0.327
1705834	2024-04-30T0...	2024-04-30T01...	00:06:09	COMPLETED	4.91	0.00926	0.139	28.5	0.0244	0.366
1705810	2024-04-30T0...	2024-04-30T0...	00:05:30	COMPLETED	10.8	0.00826	0.124	100	0.0353	0.530
1705809	2024-04-30T0...	2024-04-30T0...	00:05:42	COMPLETED	10.6	0.00891	0.134	95	0.0278	0.418
1705808	2024-04-30T0...	2024-04-30T0...	00:04:16	COMPLETED	9.60	0.00600	0.0899	55.8	0.0146	0.219
1705807	2024-04-30T0...	2024-04-30T01...	00:19:41	COMPLETED	10.3	0.0300	0.450	70.0	0.0898	1.35
1705806	2024-04-30T0...	2024-04-30T01...	00:40:20	TIMEOUT	10.4	0.0639	0.959	95.0	0.213	3.20
1705805	2024-04-30T0...	2024-04-30T01...	00:26:08	COMPLETED	11.2	0.0497	0.746	97.0	0.150	2.25
1704678	2024-04-30T0...	2024-04-30T0...	00:40:07	TIMEOUT	11.1	0.0635	0.952	50.9	0.131	1.97
1704677	2024-04-30T0...	2024-04-30T0...	00:09:46	COMPLETED	11.2	0.0129	0.193	100.0	0.0620	0.930
1694324	2024-04-30T01...	2024-04-30T01...	00:05:36	COMPLETED	8.26	0.00281	0.0422	62.1	0.0130	0.195
1694323	2024-04-30T0...	2024-04-30T01...	00:05:40	COMPLETED	5.43	0.00685	0.103	28.0	0.0159	0.238

‹ 1 2 3 4 5 › 1 - 23 of 110 rows

Operator dashboards

What more can we do?

Explore eBPF to get IO and network stats

eBPF proved to be a powerful framework in observability

Support Openstack

CEEMS has been designed to be modular and resource manager agnostic

Thank you

mahendra.paipuri@cnrs.fr

CEEMS Resources:

GitHub Repository

Docs