CINECA

Monitoring and anomaly detection in CINECA's supercomputing facility

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Daniele Cesarini

Cineca

Holistic Monitoring

CINECA

Fine Grain Power and Performance Measurements:

- Verify and classify node performance (In spec / out of spec behaviour, Miss configuration, Aging and wear out)
- Detect security hazards
- Predictive maintenance

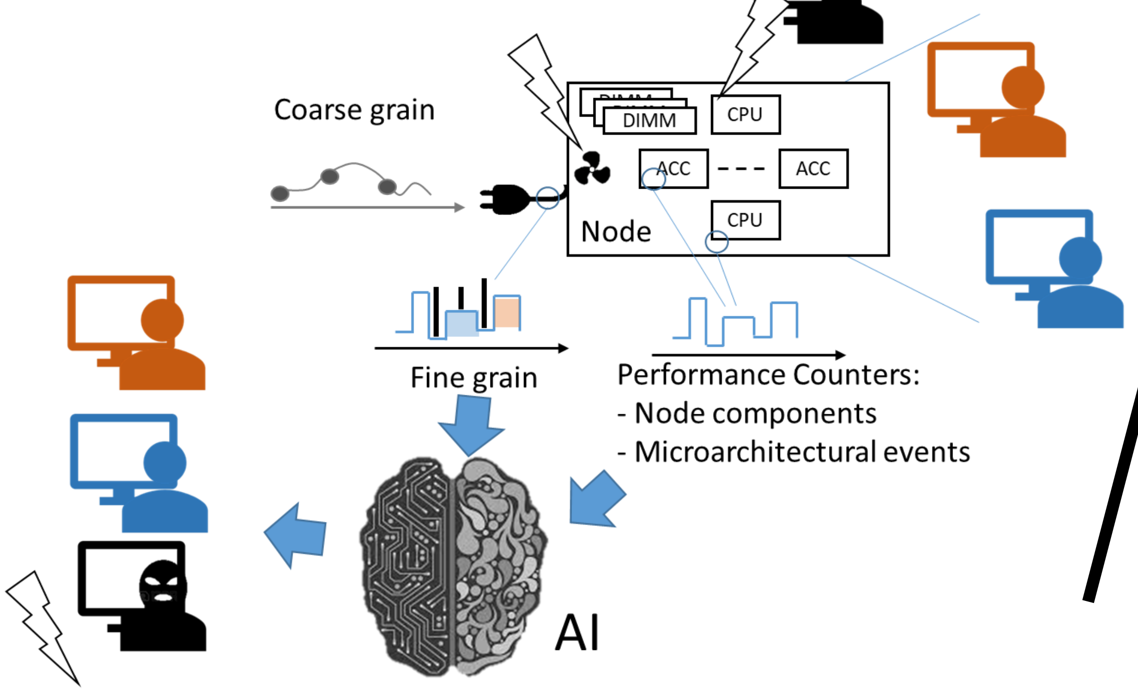

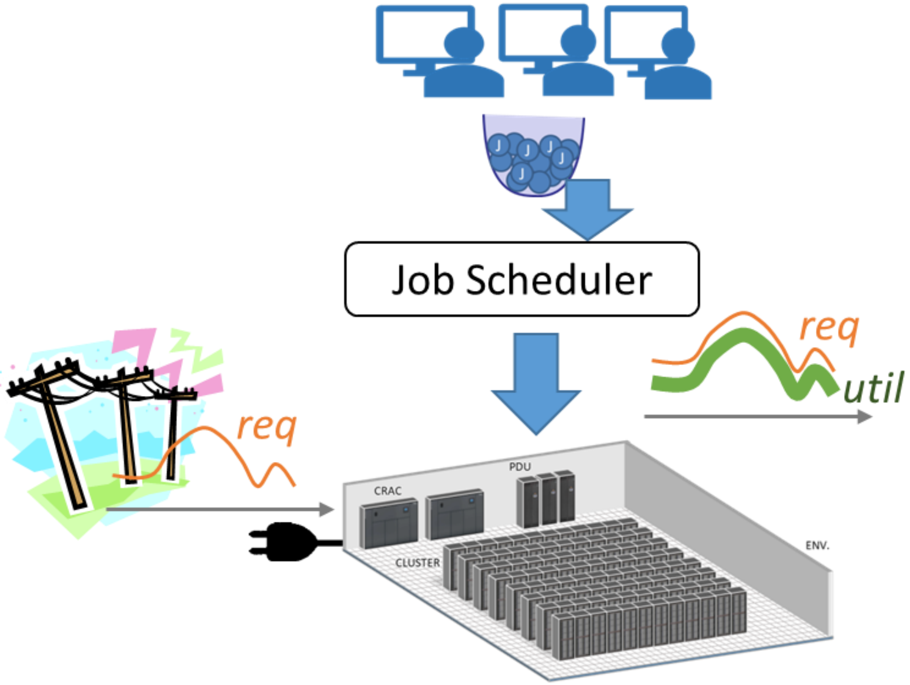

System Power Capping

- New Installations, Grid SLA, Power Shortage, Natural Disasters
- Ensures operating power below a maximum power consumption level

A. Libri et al., “pAElla: Edge AI-Based Real-Time Malware Detection in Data Centers”, JIOT 2020
A. Borghesi et al, “A semisupervised autoencoder-based approach for anomaly detection in high performance computing systems”, EAAI 2019

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Continuous and holistic monitoring of Datacentres

CINECA

A Datacentre (DC) is a large industrial plant

- ~10-20MWatts, >100M€ CAPEX, >1Msensors w. complex relationship
- Three organizations handles data: User support, System administrators, and facility manager.
- Goal: Holisticly optimize DC operations, detect and predict anomalies

ExaMon: Framework for holistic monitoring of a large plant

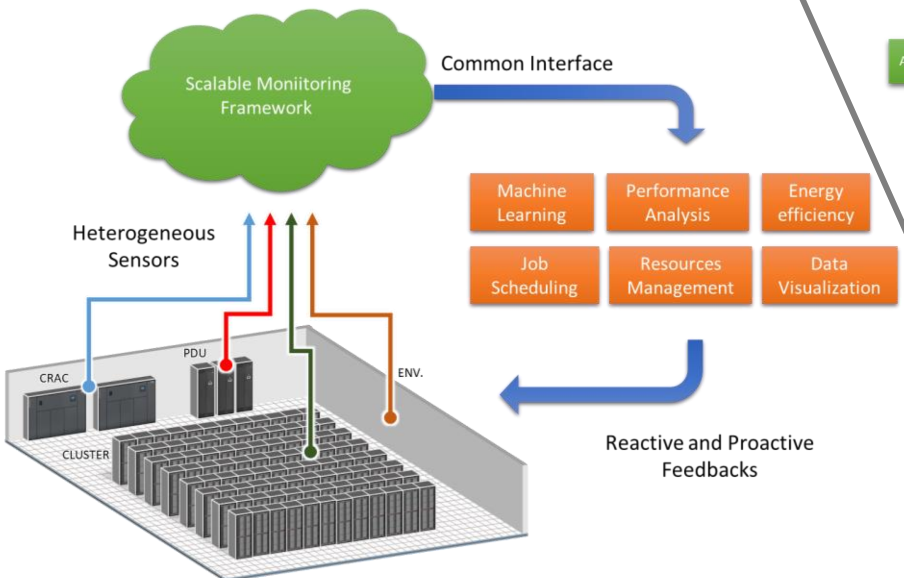

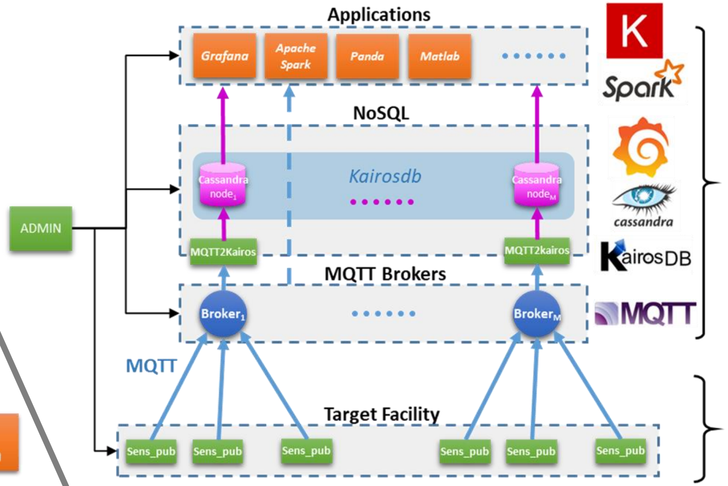

Front-end

- **MQTT** Brokers
- Data **Visualization**
- **NoSQL** Storage
- **Big Data** Analytics

Back-end

- MQTT–enabled sensor collectors

EXAMON

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ExaMon: Exascale Monitoring Framework

ExaMon (Exascale Monitoring) is a data collection and analysis platform oriented to the management of big data.

- Distinctive features:
    - distributed and horizontal scalability
    - heterogeneous data management
    - support both streaming and batch mode
    - SQL-like interface for data access
- Target usage:
    - resource monitoring and alerting
    - large-scale data analytics
    - machine learning and artificial intelligence based applications
- Examples:
    - real time anomaly detection
    - predictive maintenance
    - efficient resource and energy management
    - digital twins

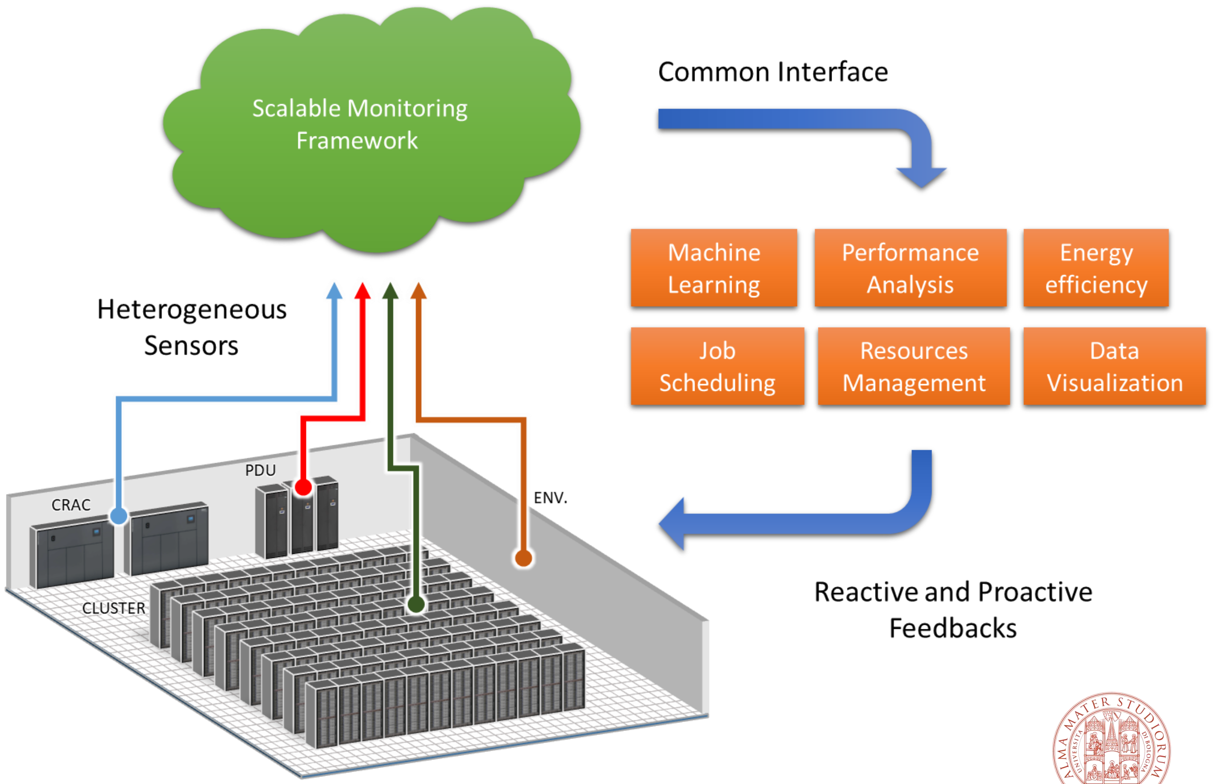

ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

ExaMon: Exascale Monitoring Framework

CINECA

ExaMon was primarily developed at UNIBO, in collaboration with CINECA, for the collection and analysis of HPC node and facility data.

- Early 2015
  - Galileo 2015
  - Marconi (BWL, KNL) 2016
  - Marconi (SKL) 2017
  - D.A.V.I.D.E. 2017-2020
  - Marconi 100 2020
  - Galileo 100 2021
- Current deploy (2023): Marconi (SKL), Marconi100, Galileo100
  - Nodes: ~4830
  - DB size: ~34TB (on-line)
  - >1Million unique sensors

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ExaMon Architecture

ExaMon implements a horizontally scalable architecture using the latest open source technologies.

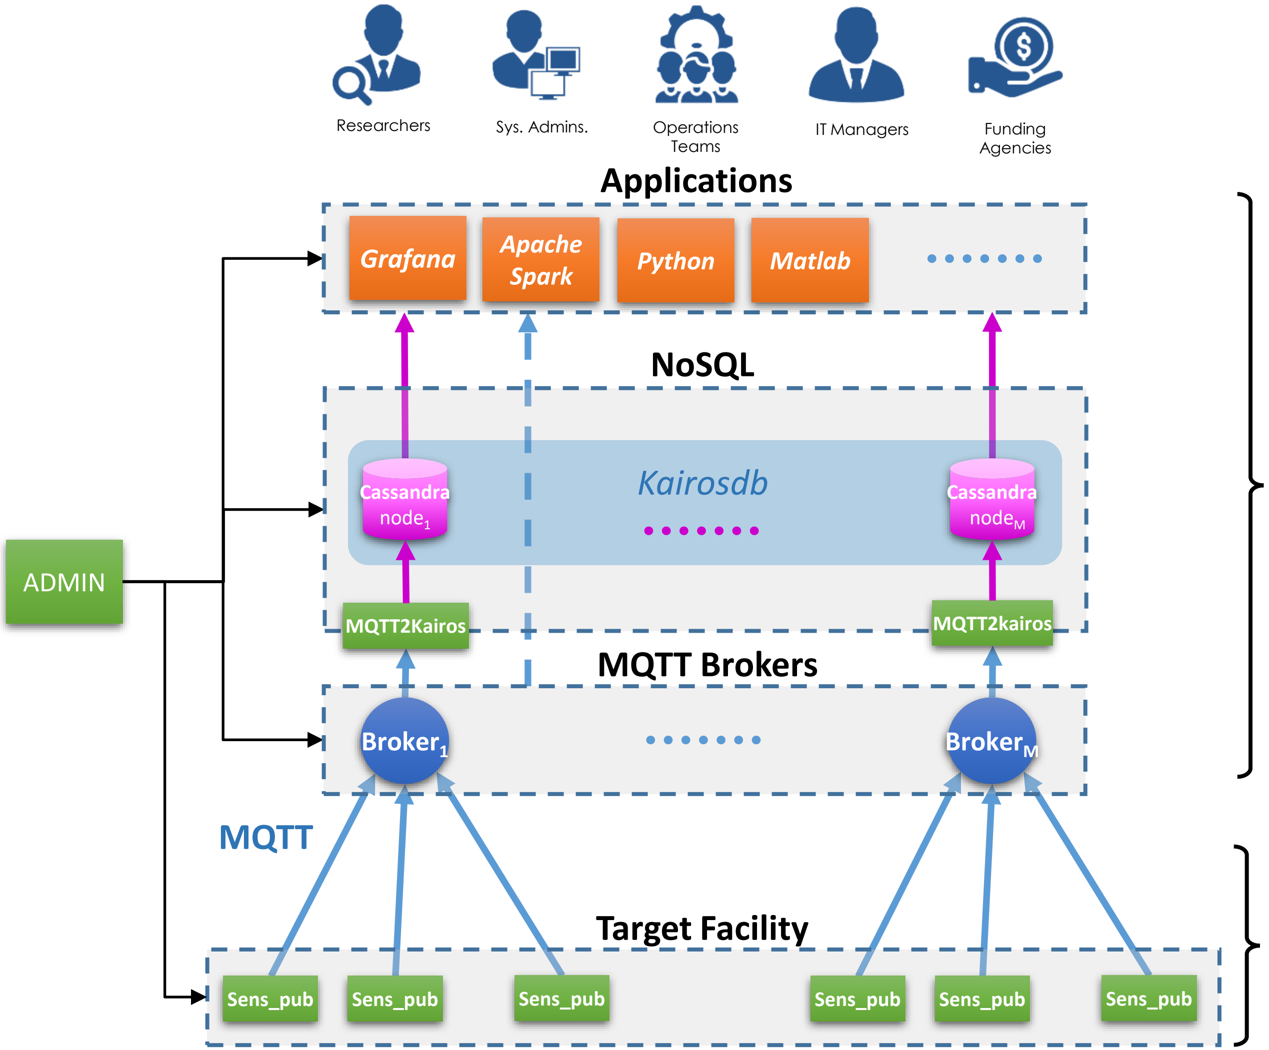

Front-end

- **MQTT** Brokers
- Data **Visualization**
- **NoSQL** Storage
- **Big Data** Analytics

Back-end

- **MQTT**–enabled sensor collectors

MQTT: MQ Telemetry Transport

- **Lightweight** message queueing and transport protocol
- Developed by IBM and Eurotech
- Well suited for **low resource demanding** scenarios like M2M, WSN and **IoT** applications
- Basic features:
  - **PubSub** model
  - Async communication protocol (**messages**)
  - **Low overhead** packet (2 bytes header)
  - **QoS** (3 levels)

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Examon MQTT Formats -Topic

CINECA

Topic format details:

<key>/<value>/<key>/<value>/ … /**plugin**/<plugin_name>/**chnl**/<data|cmd>/<key>/<value>/ … /<metric_name>

Sensor location
Free hierarchical sequence of key/value couples to locate sensors. At least one couple should be present.

Plugin name and channel type
(mandatory)
Channel type:

- *data*: metrics values sent by the plugin
- *cmd*: commands sent to the plugin

Specific plugin info
Additional free key/value couples to add custom plugin attributes (optional)

Metric name

Plugin data

Example:

TOPIC:
`org/cineca/room/n/plugin/sensortag/chnl/data/sensorid/05/temp`
PAYLOAD:
`25;1488908453.000`

MQTT: Comm. Example

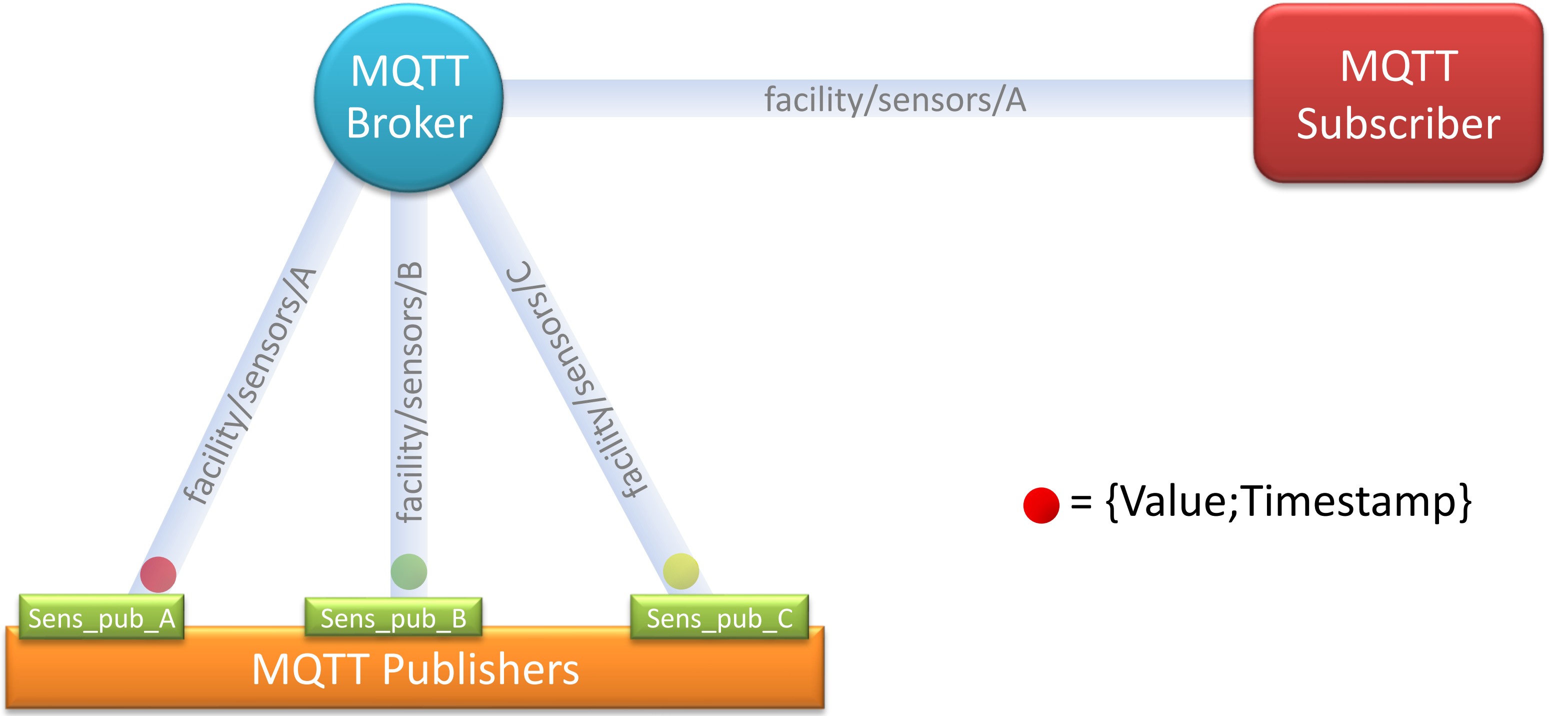

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

NoSQL Storage: MQTT2Kairosdb

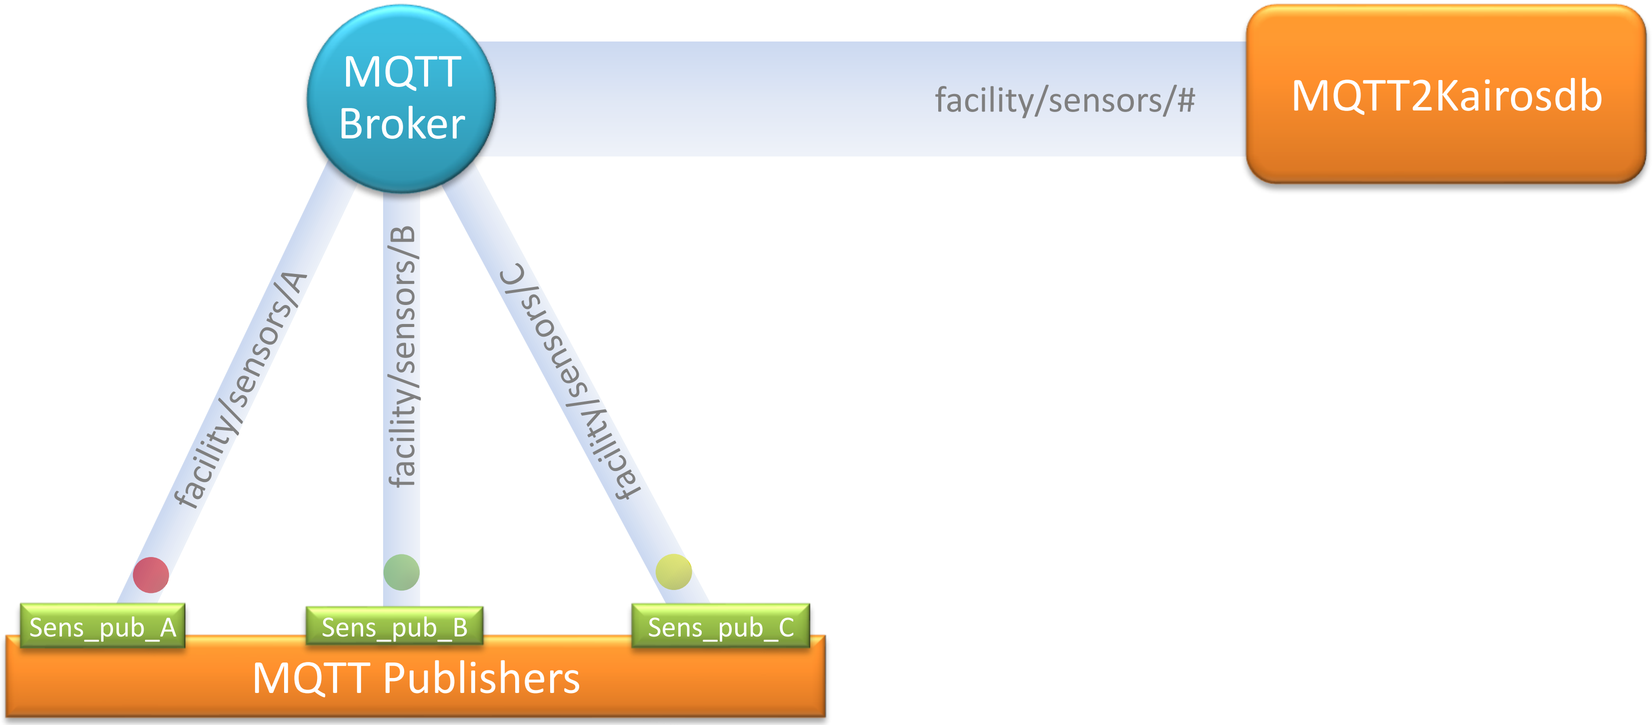

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

CINECA

Examon Analytics: Batch

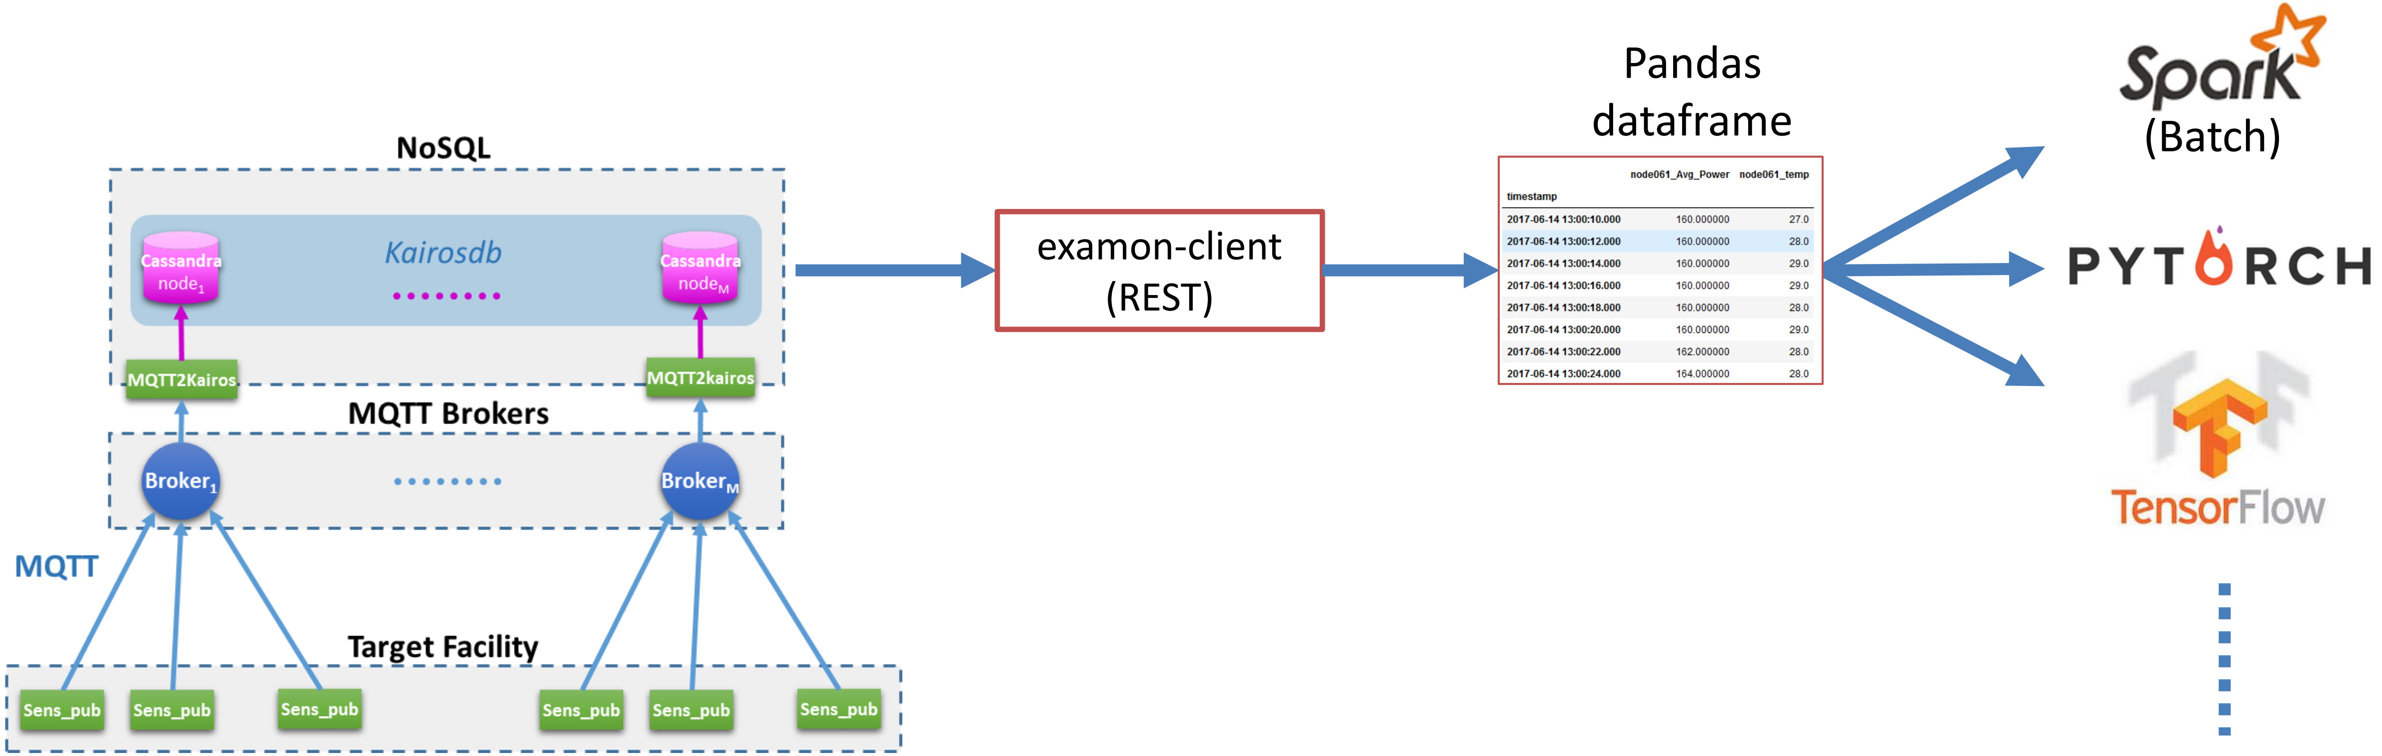

timestamp	node061_Avg_Power	node061_temp
2017-06-14 13:00:10.000	160.000000	27.0
2017-06-14 13:00:12.000	160.000000	28.0
2017-06-14 13:00:14.000	160.000000	29.0
2017-06-14 13:00:16.000	160.000000	29.0
2017-06-14 13:00:18.000	160.000000	28.0
2017-06-14 13:00:20.000	160.000000	29.0
2017-06-14 13:00:22.000	162.000000	28.0
2017-06-14 13:00:24.000	164.000000	28.0

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Examon Analytics: Streaming

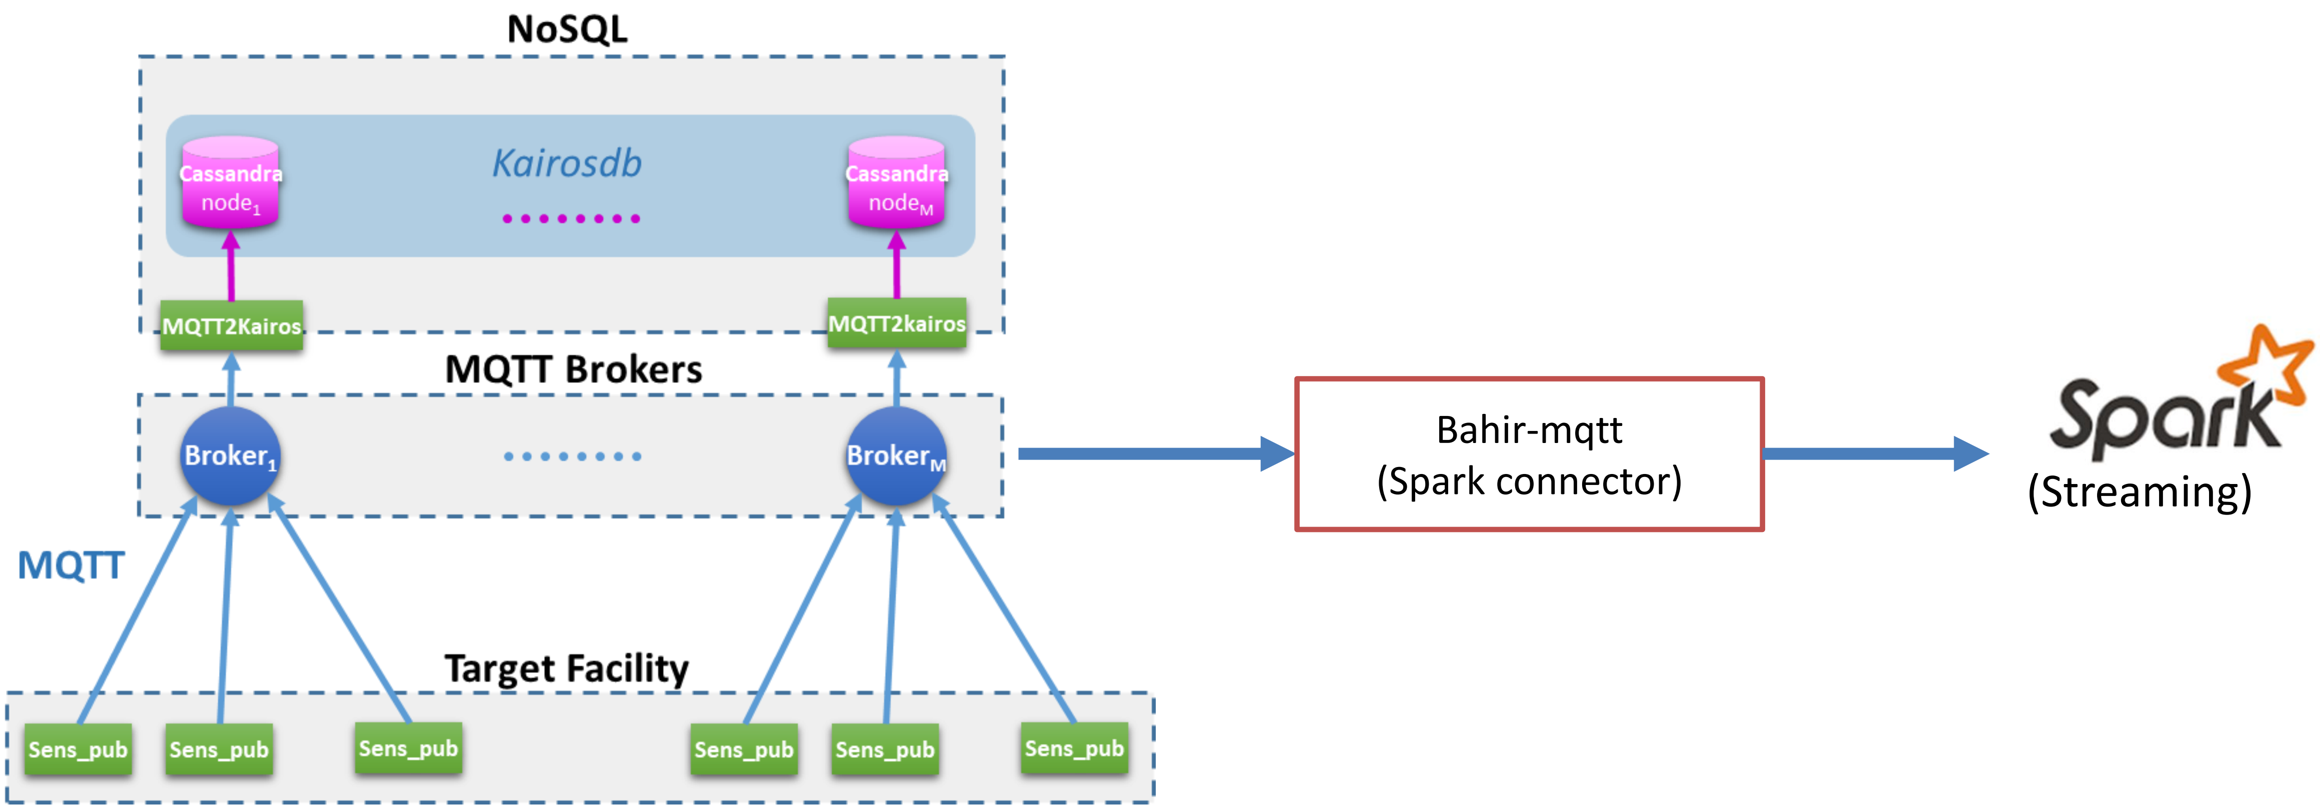

Data Sources

A simple and scalable data model enables ExaMon to collect a large variety of data from heterogeneous sources.

- Node data
  - IPMI (power, inlet air temp, …)
  - OS (CPU, memory, load, …)
  - Users activity (Batch scheduler)
  - Events (Nagios)
- Room data
  - Cooling (CRAC/H, chillers, pumps, …)
  - Power (PDUs, electrical panels, …)
- External environment
  - Temperature, humidity, pressure, …

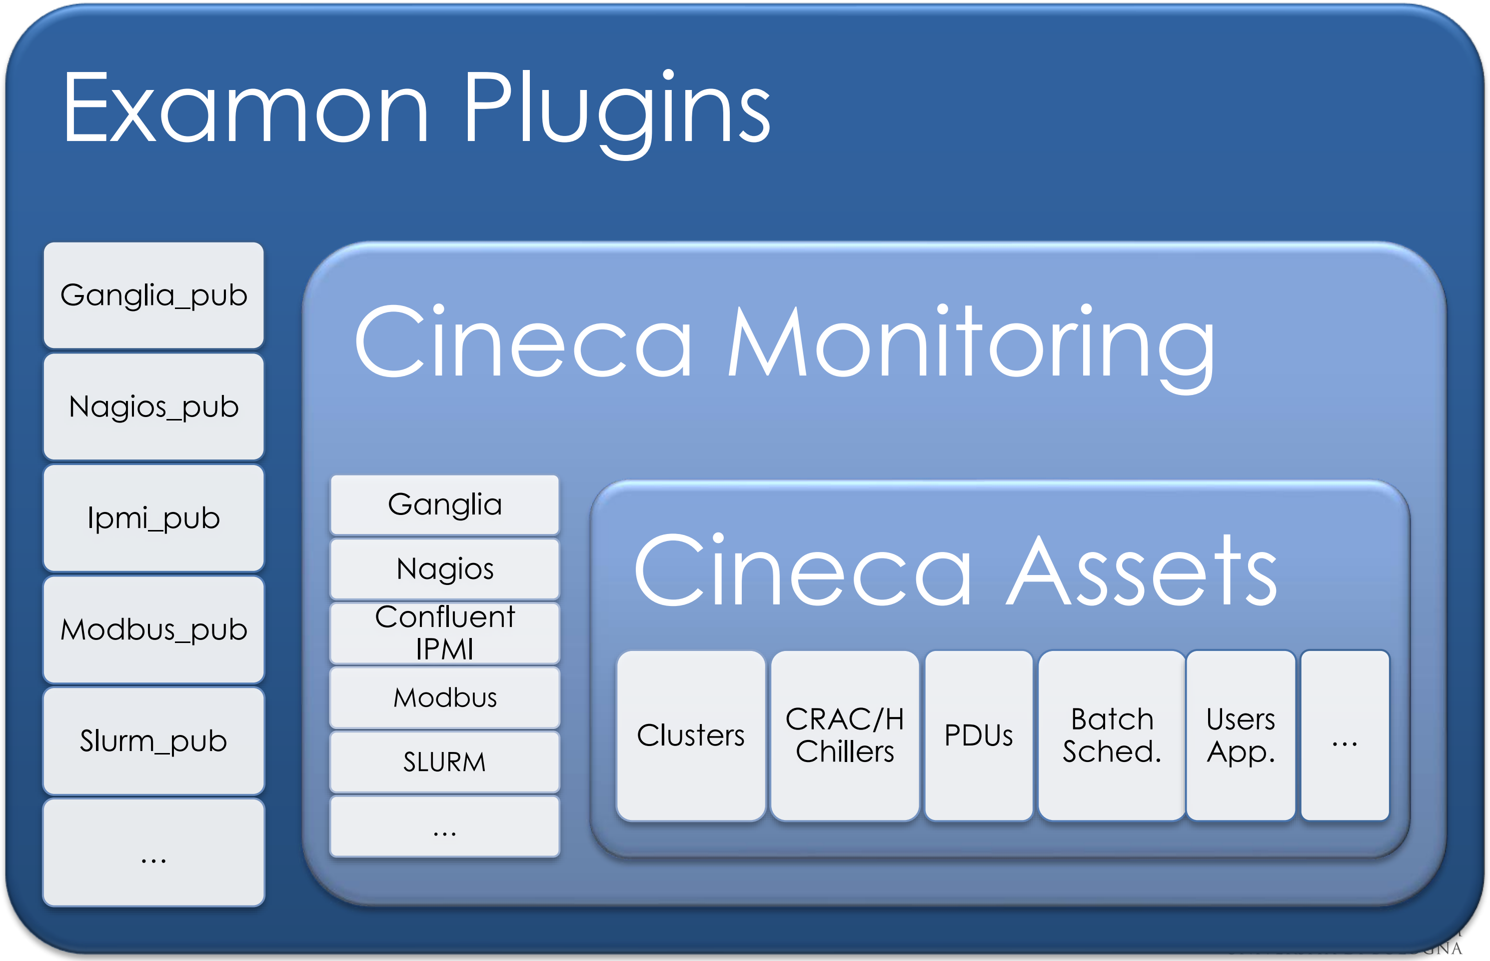

Examon@CINECA: Current Setup

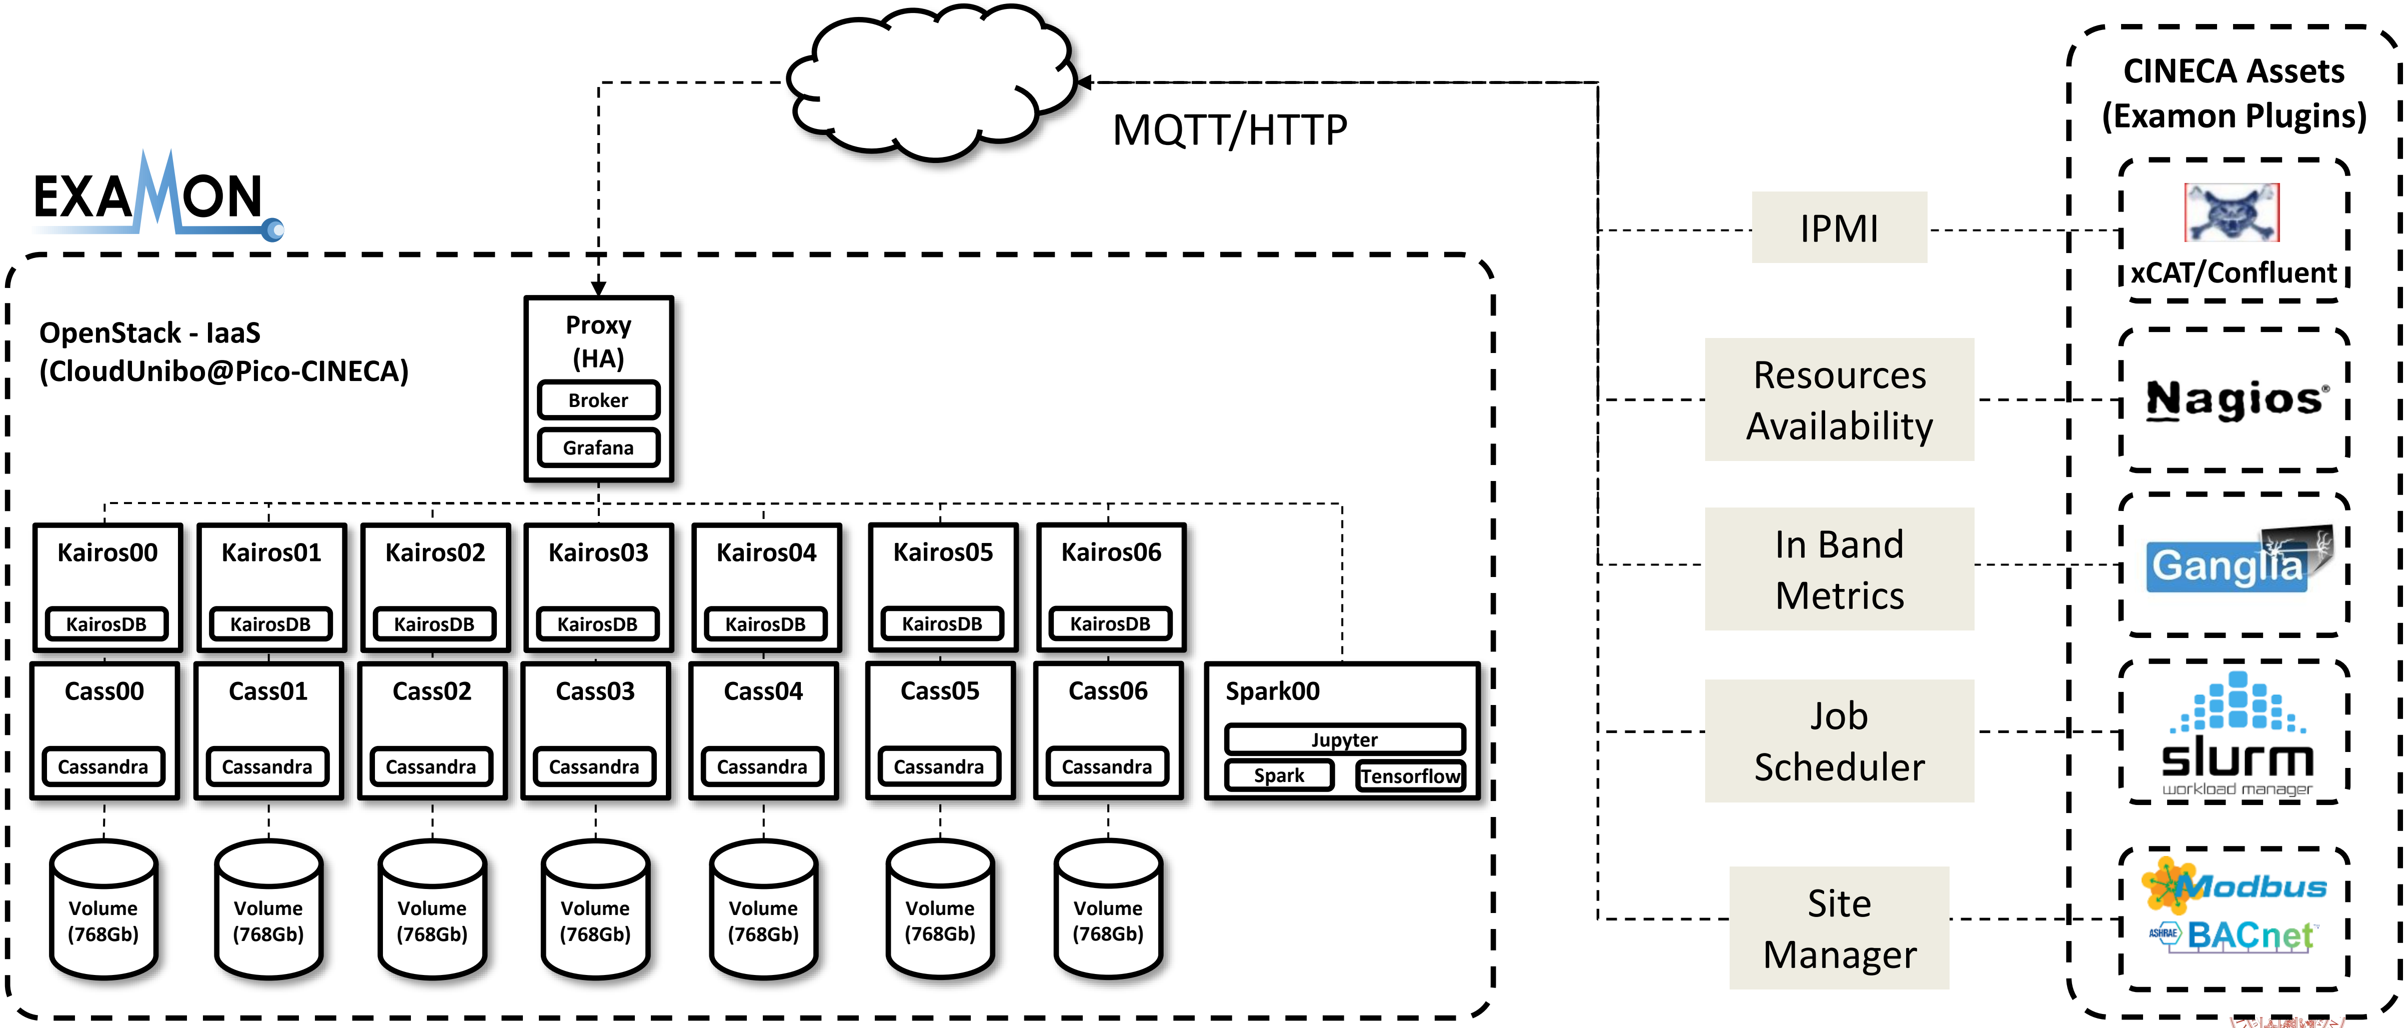

CINECA

Data Analytics and Visualization

ExaMon aims to deliver near real-time data visualization and analytics, as well as advanced batch processing utilizing cutting-edge technologies

- Visualization
  - Using **Grafana** enables dynamic and interactive dashboards, as well as real-time alerting and notifications
- Analytics + visualization
  - Takes advantage of **ANSI/SQL** capabilities and real-time dashboarding capabilities with **BI** tools like **Redash** and **Superset** to manage complex analytics and deliver dynamic insights in near real-time
- Big data analytics
  - Enables data processing using tools like **Jupyter Notebooks**, **Apache Spark**, **Dask** and others, providing a versatile and flexible platform for **historical** data analysis and exploration.
- CLI tools
  - Provides the flexibility to deliver the same results through customizable command line utilities (**CLI**) tools

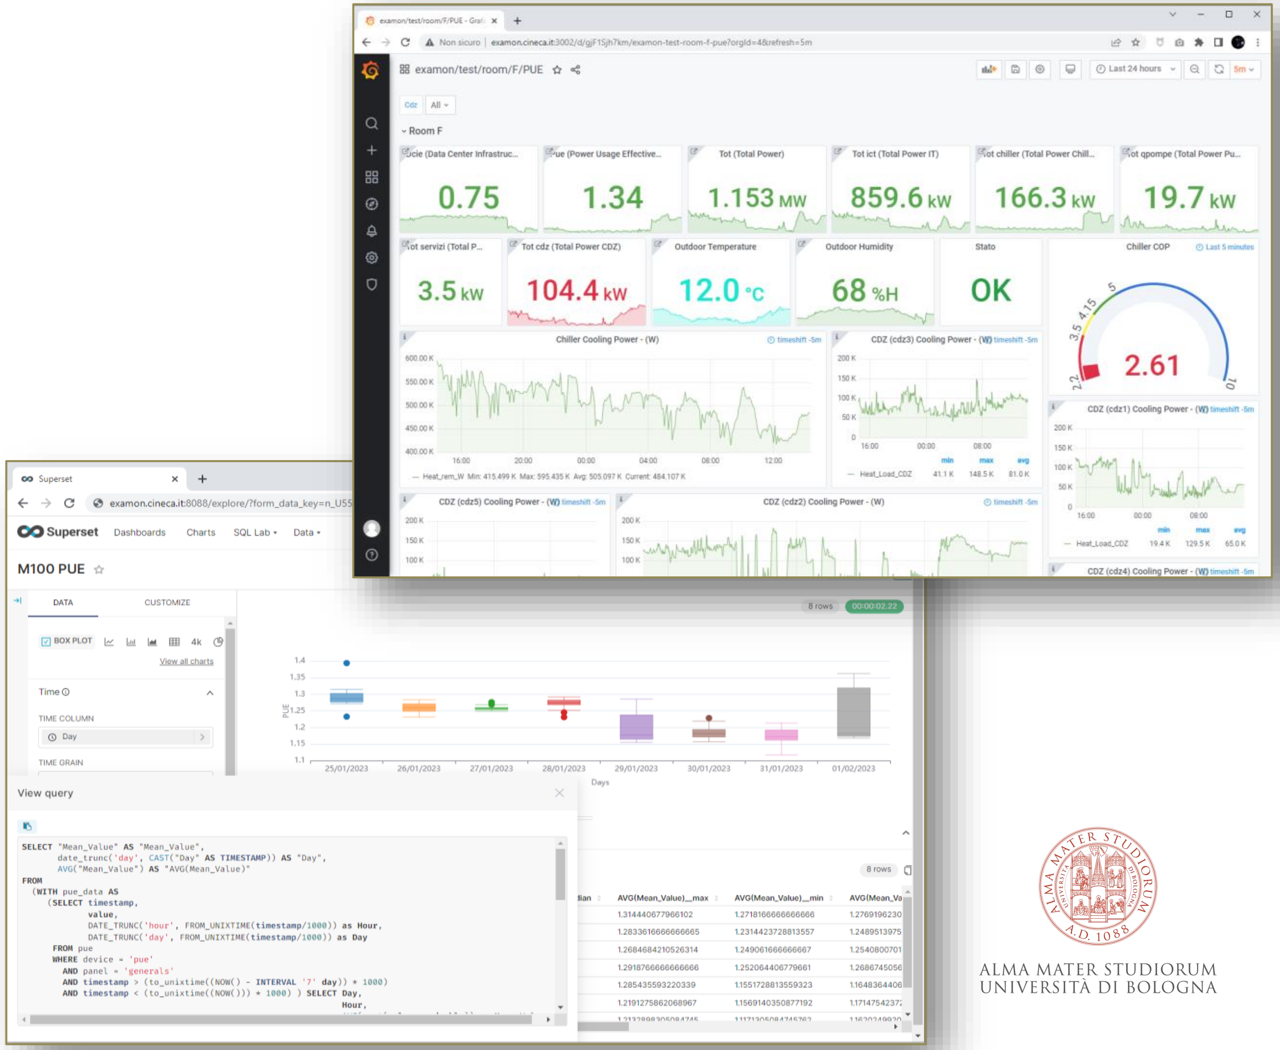

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Use Case: Cineca Clusters Availability

CINECA

Working jointly with Cineca on a real use case, we created an example of how ExaMon can increase productivity and improve the service

- Goal
  - Creation of a real time dashboard for visualizing the state of the clusters useful for both users and operators
- Before
  - Cluster status and availability was previously managed through **cron jobs** and **administrative tools**, limiting access to operators only.
  - Reports for the users were generated on a **monthly** basis.
- Now
  - Ad hoc dashboards powered by ExaMon data give operators and users the same metrics as before plus several new ones with a **10-second refresh rate**.

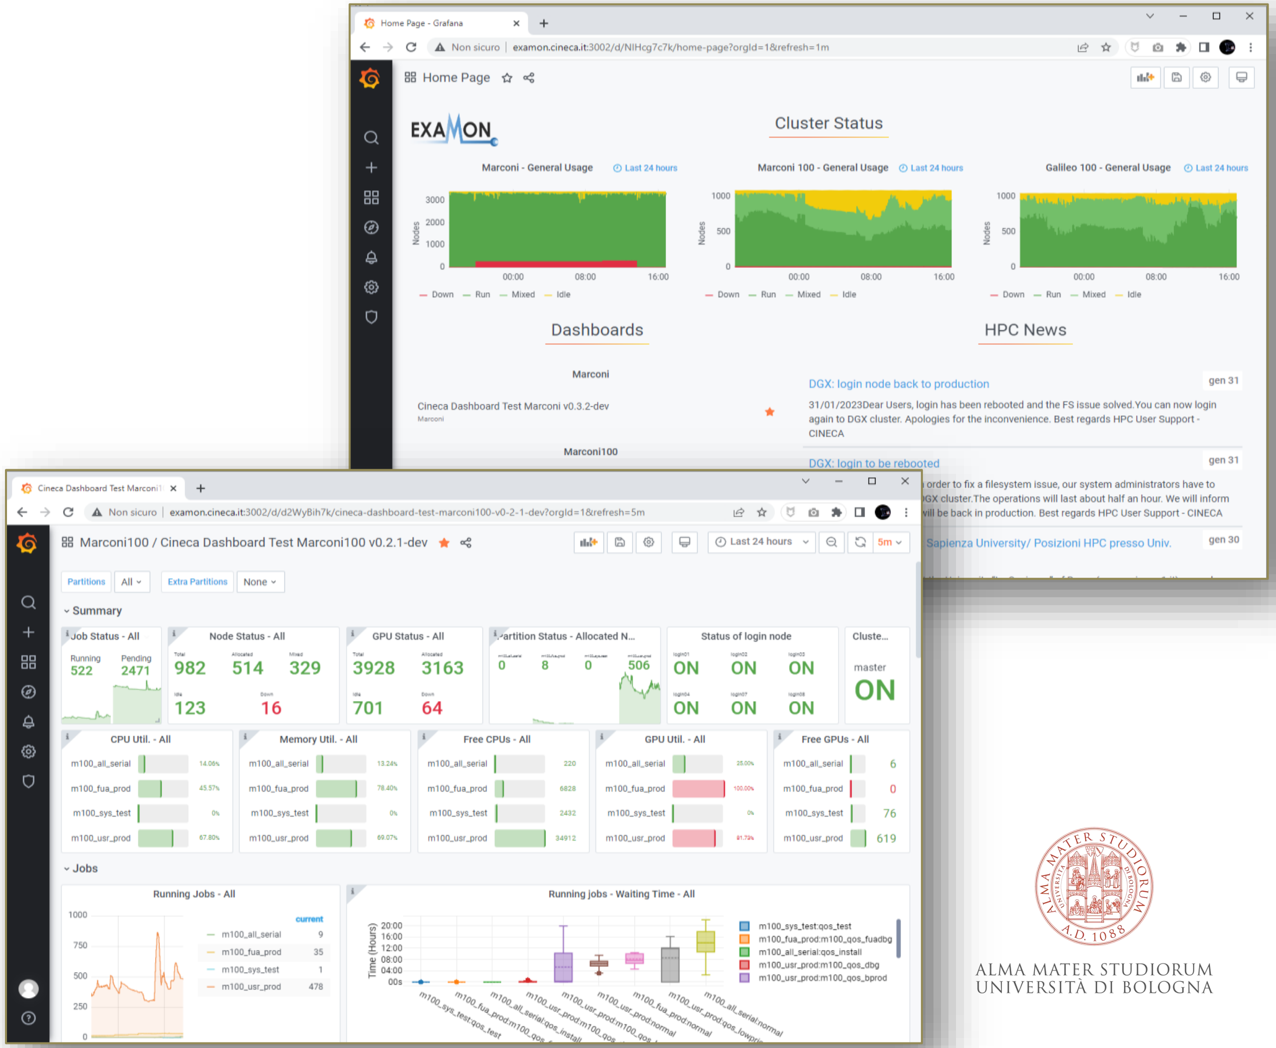

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Use Case: Cluster Digital Twin

Using 3D visualization tool linked to the real time data provided by ExaMon can bring several benefits

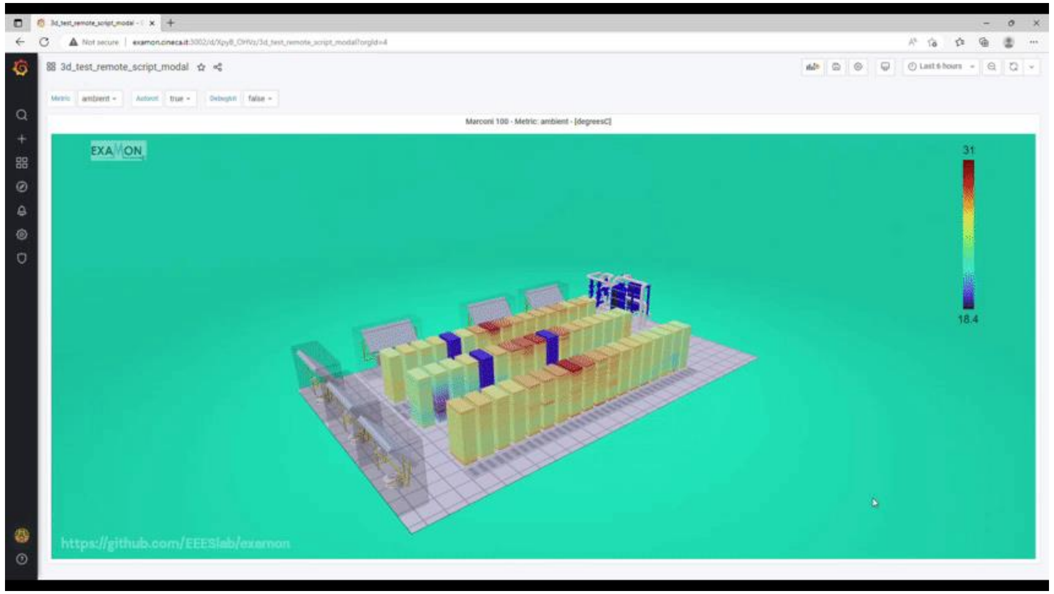

- Visualization and Analysis
  - Helps identify and understand events and behaviors in relation to the location of objects.
  - Enables **XR** (VR/AR/MR) applications

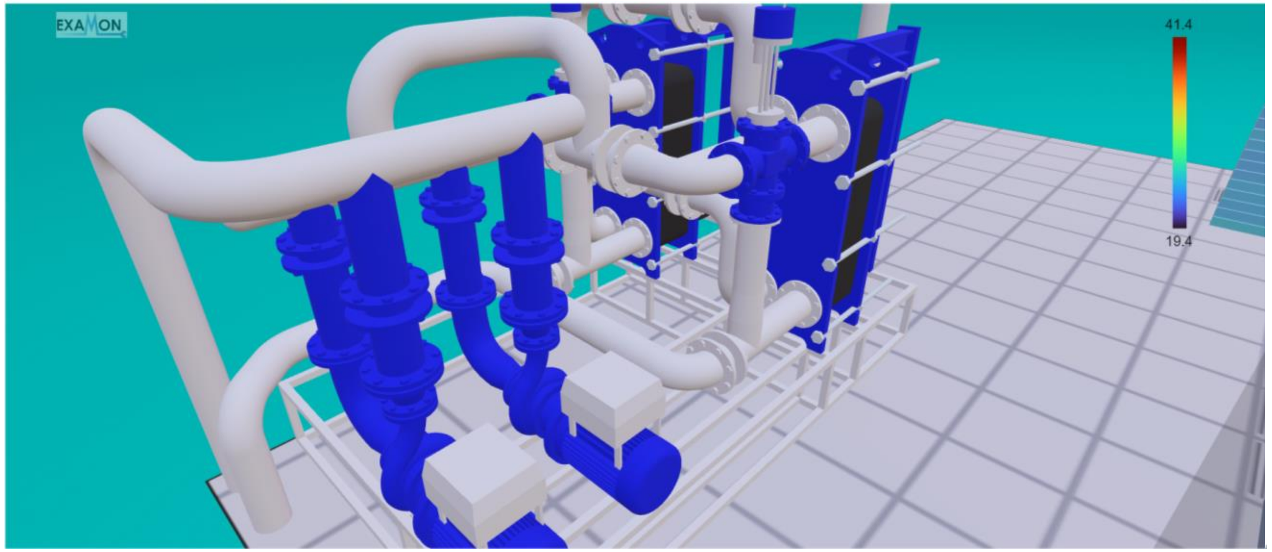

- Improved collaboration
  - Visualizing data and issues in a common and familiar visual representation enables better decision-making through improved communication and collaboration.

Examon Demo

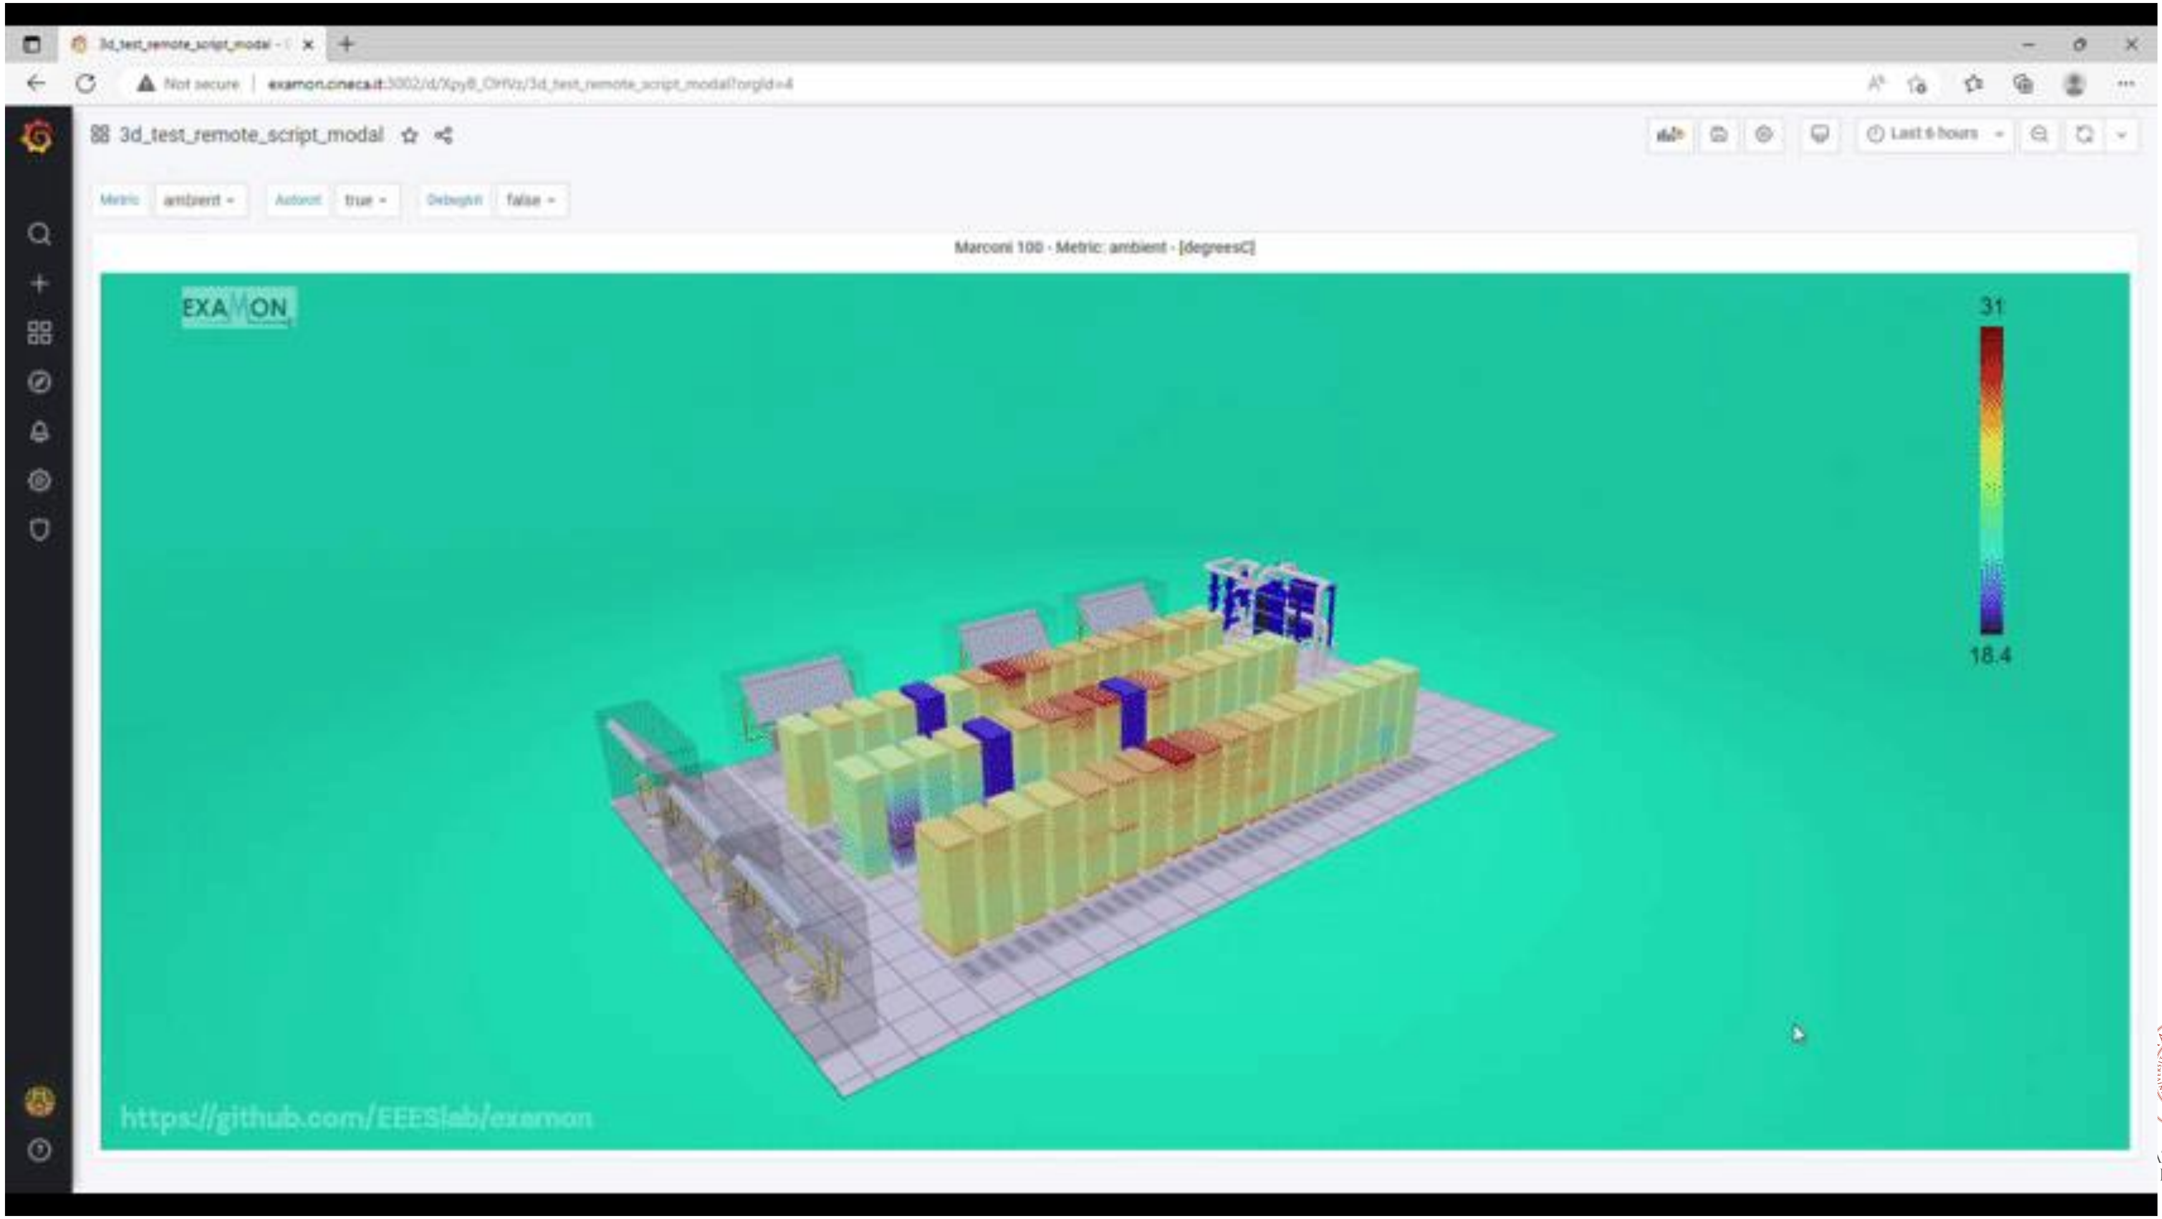

3d_test_remote_script_modal
Marconi 100 - Metric: ambient - [degreesC]
EXAMON
31
18.4
https://github.com/EEESlab/examon

Use Case: Cluster Digital Twin

PoC#1: Data center room power and thermal analysis

The overall view of the **heat generation and dissipation process** allows for a qualitative and immediate evaluation of the current cooling strategy.

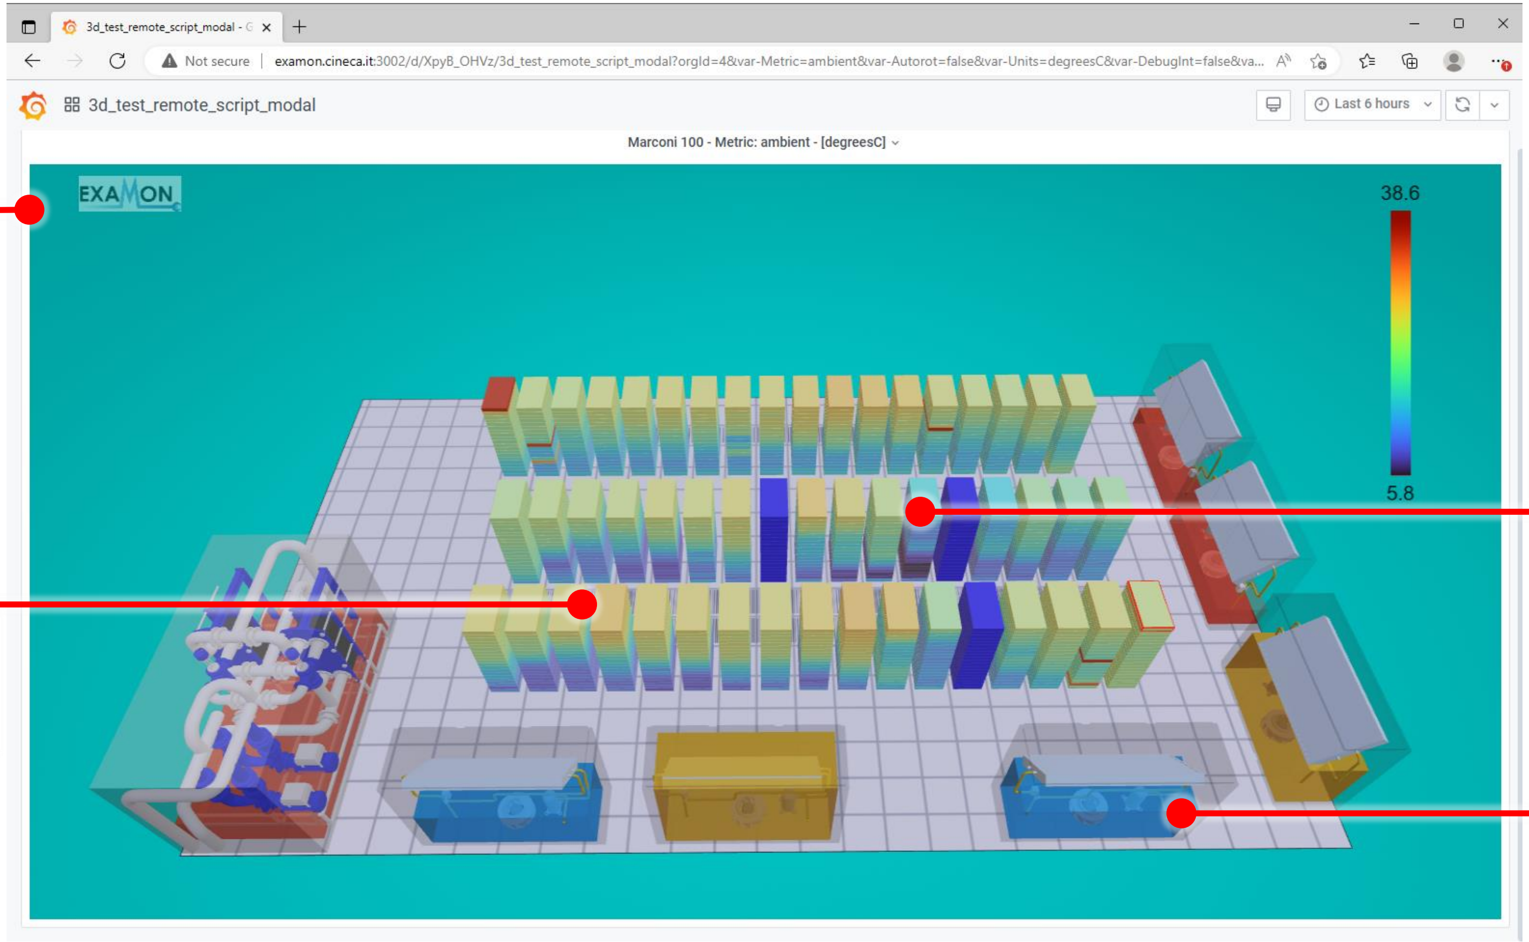

Hot spots, thermal imbalances, and cooling efficiency can be quickly investigated to improve the overall system.

Node **power** consumption and **inlet temperatures** can be mapped and aggregated by node or by rack (average, total, max, ...).

3D **widgets** can display complex metrics such as CRAC and chiller **cooling load** (bar height) and **efficiency** (color).

EXAMON + AI: Anomaly detection in a datacenter

ExaMon@2021:

- Deployed on CINECA Datacentre since 2015
- Monitoring Operation, Facility, ICT and Users:
  >1M sensors, DB: 7TB online, 12GBs/Day, 21KSa/s
- Flexible dashboard for User Support, Admin and Facility managers

ExaMon + AI → **Anomaly Detection & Anticipation!**

Idea: use DL to extract normal behaviour and relationship from the monitored sensors.

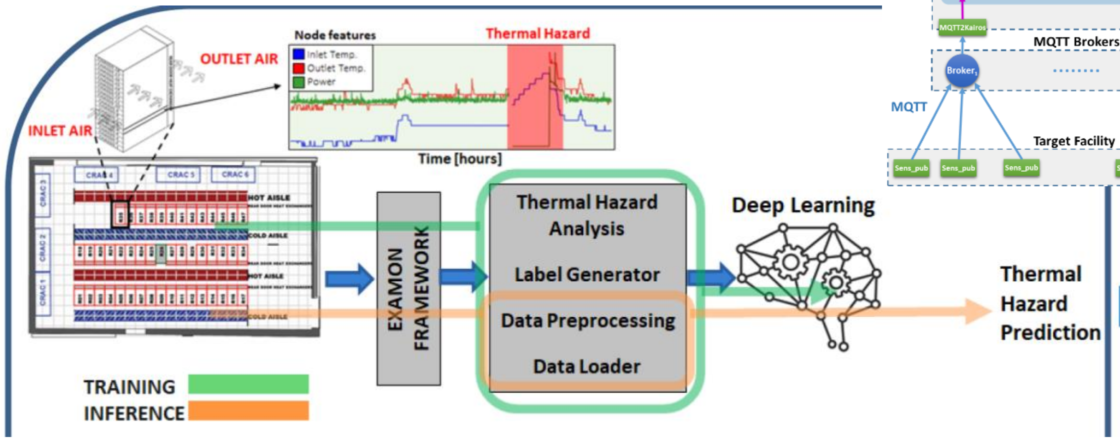

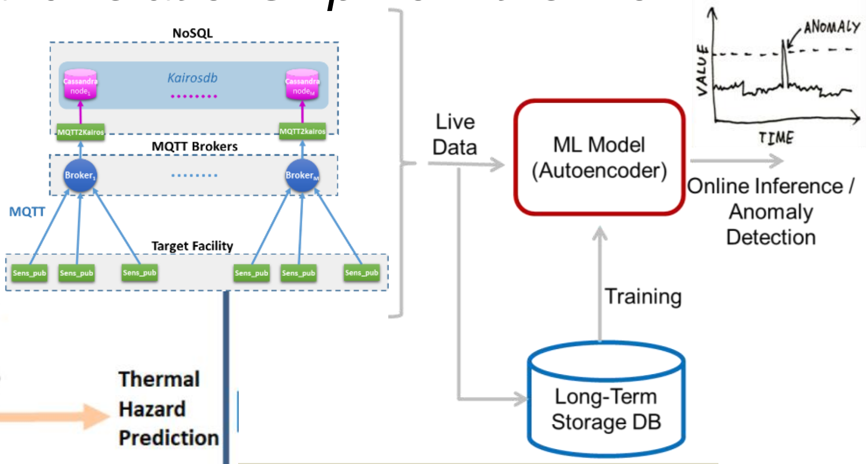

Detect thermal hazards and cooling shortage

Detect node's miss-configurations & anomalies

Use-case #1 - Datacenter Automation (Anomaly Detection & Anticipation)

- Detect anomalies/faults in a HPC system
- Hundreds/thousands of possible sources:
  - HW components that malfunction, breakages, misconfigurations, intruders, etc.
- Strong incentive to automatize the detection process
  - Downtime are *very* expensive
  - It's better to identify a problem as soon as possible

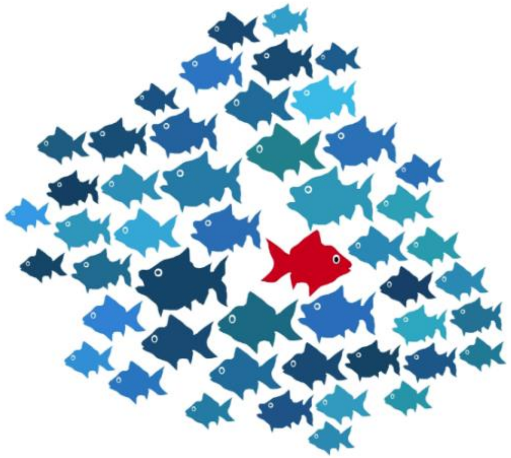

- **Solution: DL models that can distinguish anomalies from normal situations**

A. Borghesi et al., “Anomaly Detection using Autoencoders in High Performance Computing Systems”, AAAI'19
A. Borghesi et al, “Online anomaly detection in hpc systems”, AICAS'19
A. Borghesi et al., "A semisupervised autoencoder-based approach for anomaly detection in high performance computing systems", EAAI 2019

ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

Anomaly Detection General Scheme

CINECA

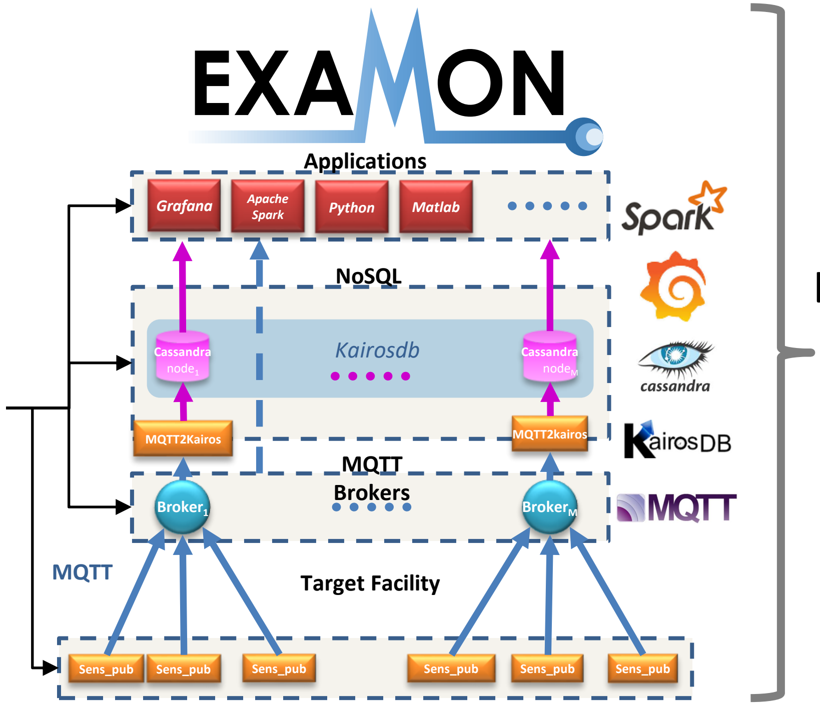

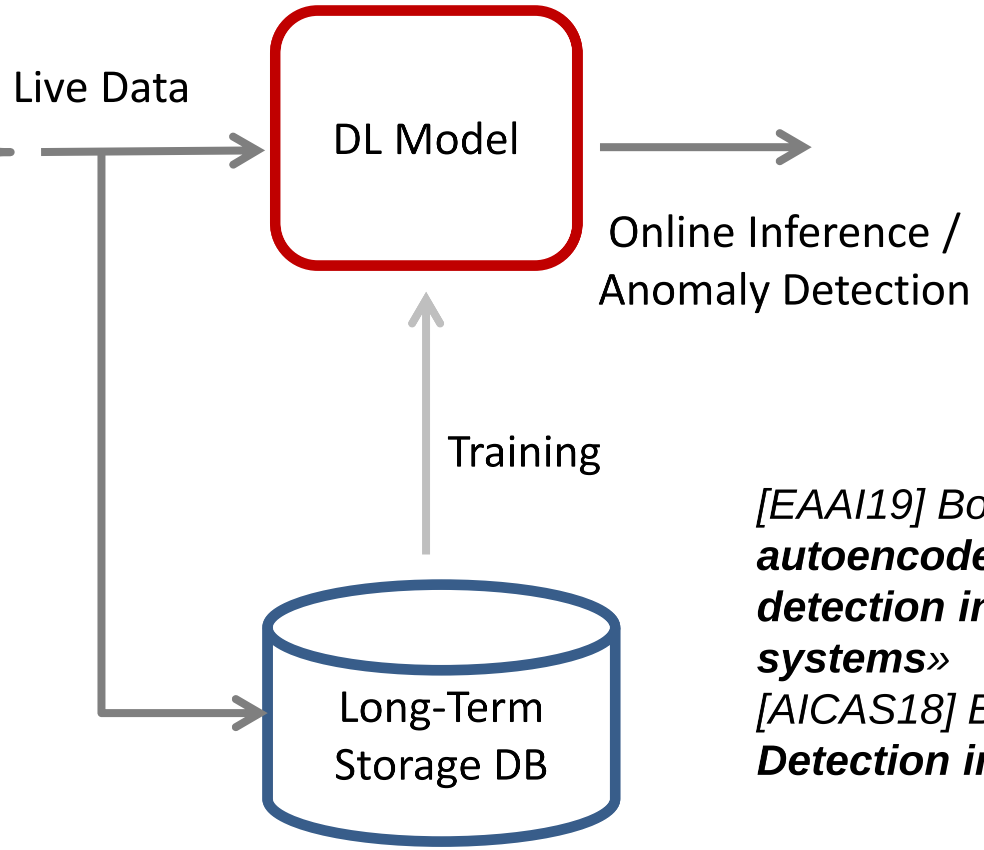

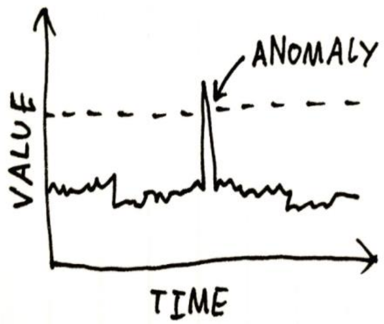

*[EAAI19] Borghesi et. al «**A semisupervised autoencoder-based approach for anomaly detection in high performance computing systems**»*
*[AICAS18] Borghesi et al. «**Online Anomaly Detection in HPC Systems**»*

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Anomaly Detection - Semi-Supervised

- Deep Learning models for anomaly detection based on autoencoder networks, a semi-supervised approach → very few labels are required!
- IDEA: train a set of autoencoders with the normal behaviour of each supercomputing node and use the reconstruction error to detect anomalies

Threshold T:

- Reconstruction Error > T → anomaly
- Reconstruction Error <= T → normal behaviour

Classification Accuracy: 90%-97%

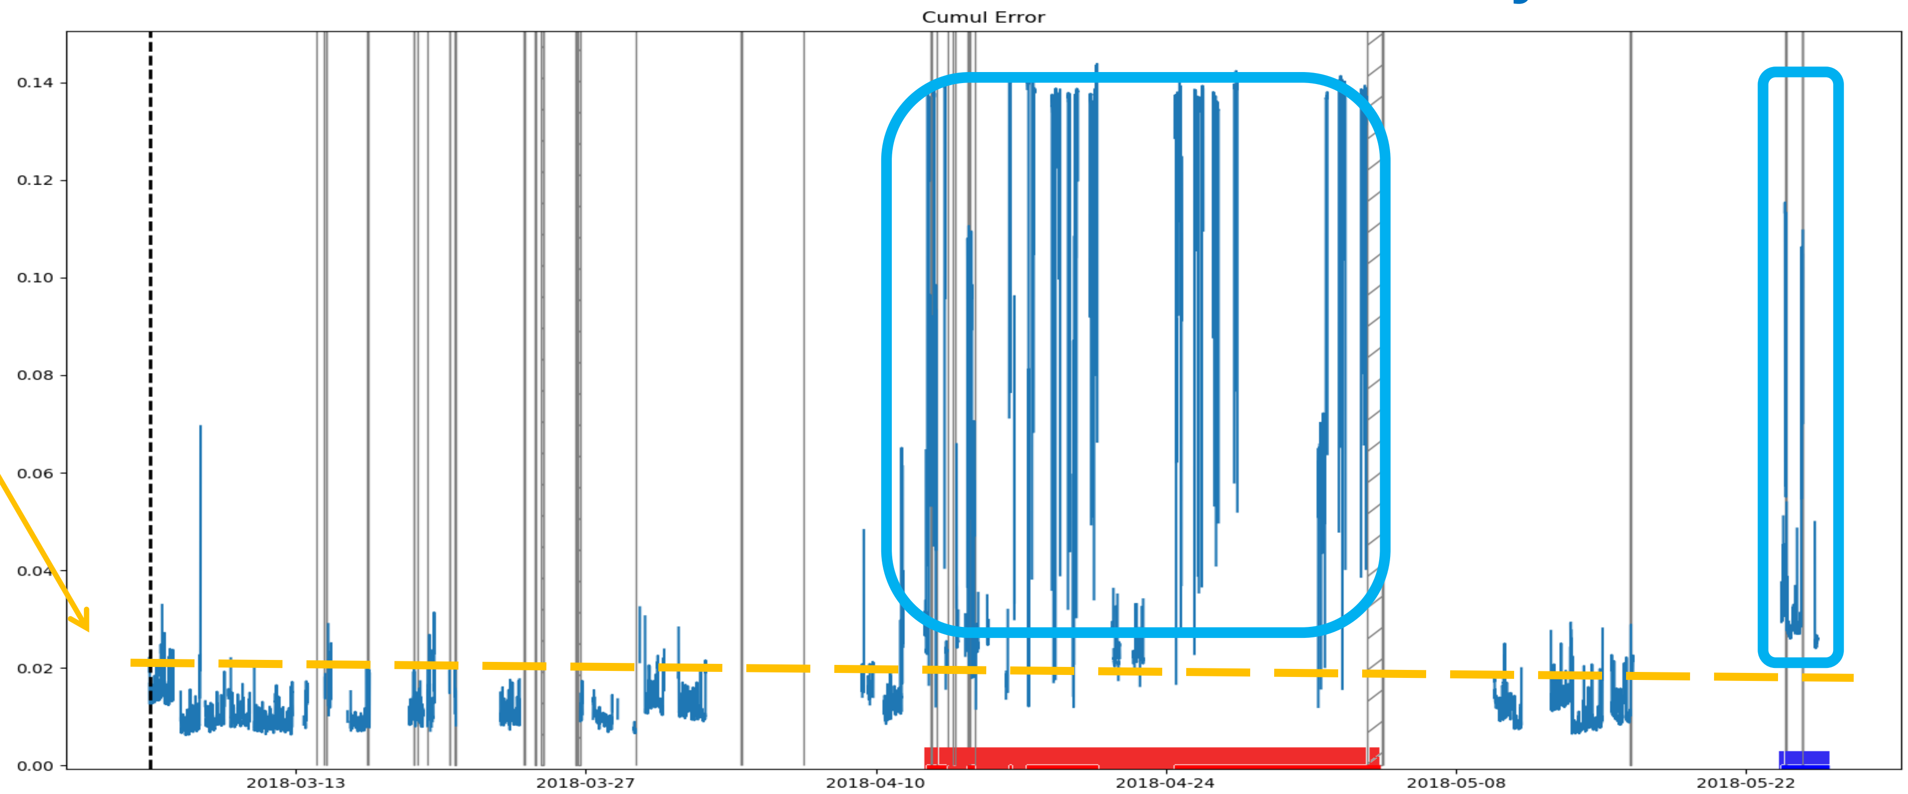

Anomaly Detection - Semi-Supervised & Supervised

- If we have labels describing the state of the datacenter, we can opt for a supervised approach
- However, the optimal results are obtained combining both semi-supervised and supervised approaches!
- We exploit best of both worlds → leading to anomaly prediction
  - The model realize that something is off even *before* the change in label

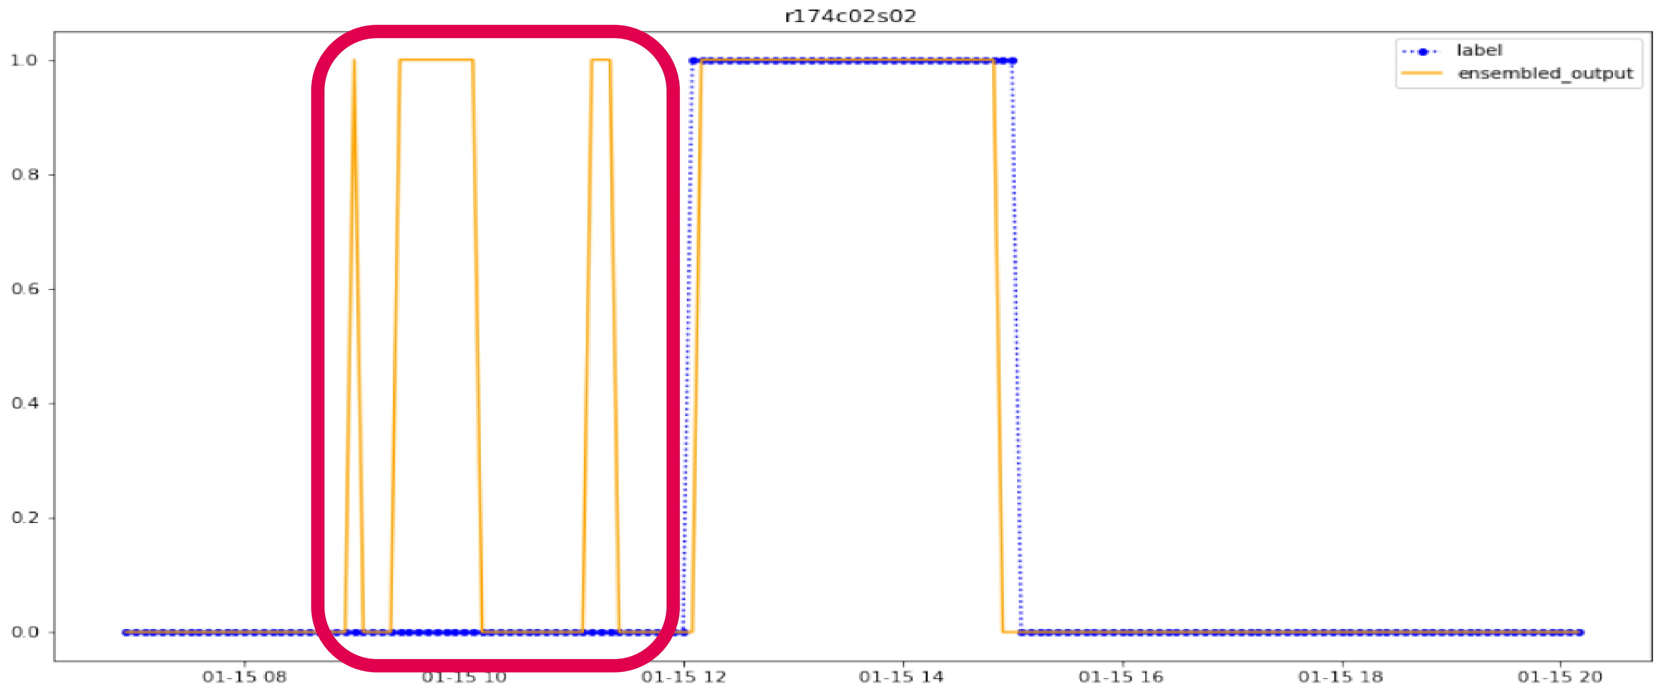

The anomaly signal (yellow line) anticipates the actual label (blue line)

- Average anticipation: 50-40 minutes

A. Borghesi et al., "Anomaly Detection and Anticipation in High Performance Computing Systems", TPDS 2021

Autoencoders

CINECA

SoA for anomaly detection: semi-supervised training of autoencoders

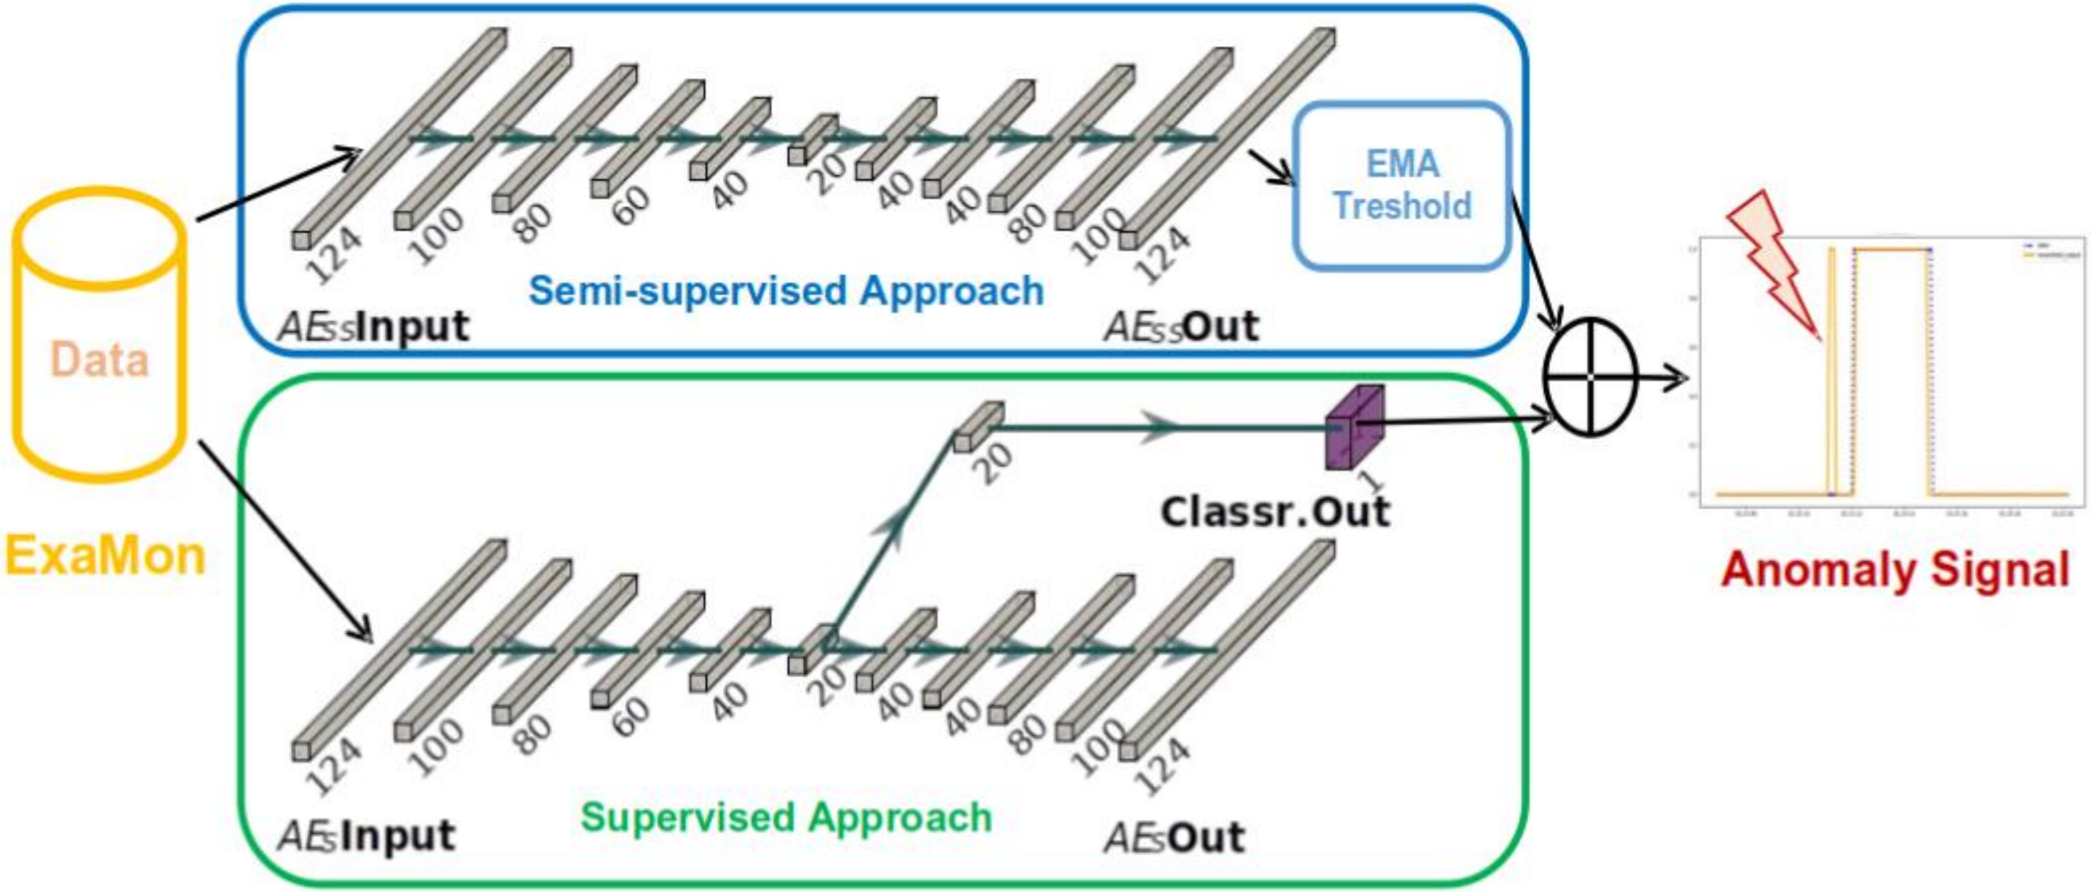

A. Borghesi, M. Molan, M. Milano and A. Bartolini, "Anomaly Detection and Anticipation in High Performance Computing Systems," in *IEEE Transactions on Parallel and Distributed Systems*, vol. 33, no. 4, pp. 739-750, 1 April 2022, doi: 10.1109/TPDS.2021.3082802.

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

WHY UNSUPERVISED ANOMALY DETECTION?

CINECA

- Semi-supervised anomaly detection: models trained only on normal operation

- Need for (accurate) information about downtimes (anomaly timestamps)

- Difficult to deploy - accurate downtime information is **not** always available

- Motivation: train on **all data including anomalies**

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Unsupervised anomaly detection: RUAD

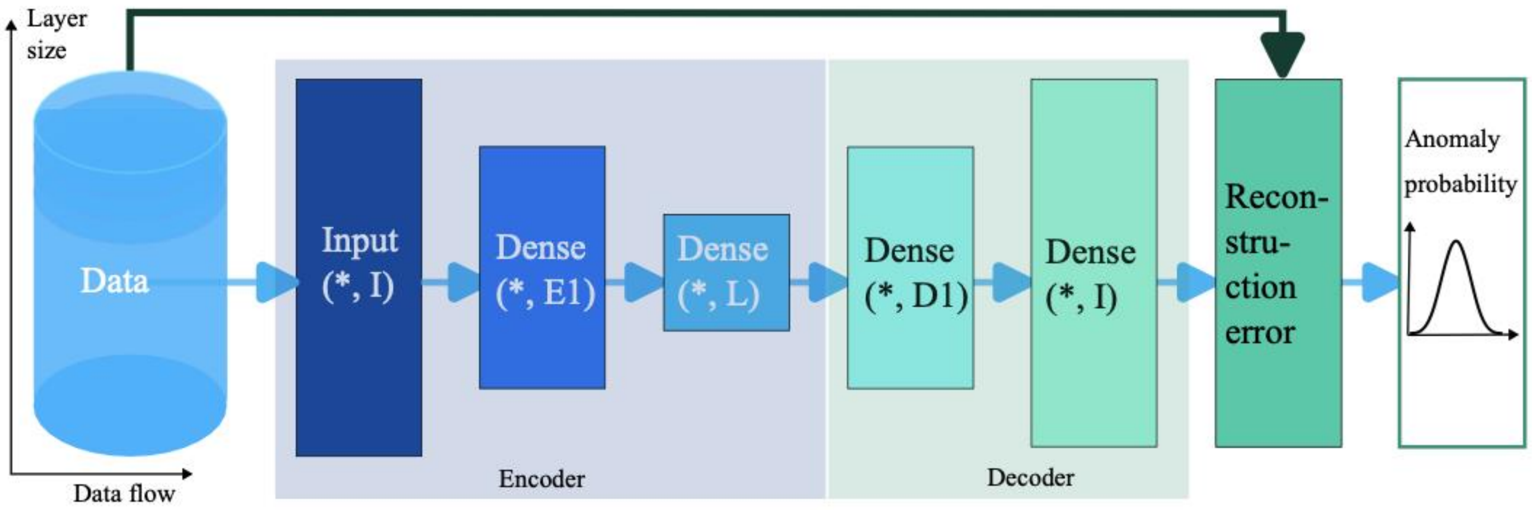

(a) Structure of baseline model - the dense autoencoder.

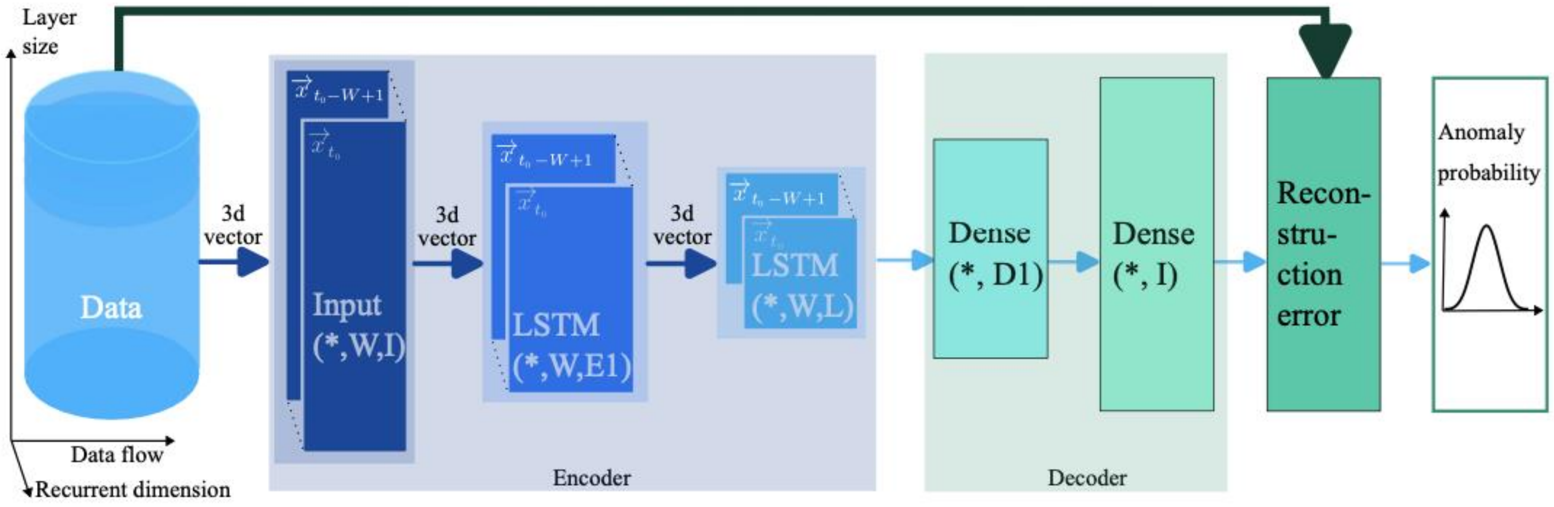

(b) Structure of the proposed RUAD model consisting of the LSTM encoder and dense decoder.

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

RESULTS ON A COMPLETE HPC SYSTEM MARCONI 100

CINECA

Method	Combined ROC score
EXP	0.4276
CLU	0.5478
$DENSE_{semi}$	0.7470
$DENSE_{un}$	0.7344

Method	Combined ROC score			
Sequence length	5	10	20	40
$RUAD_{semi}$	0.7632	0.7582	0.7602	0.7446
$RUAD$	**0.7651**	**0.7672**	**0.7655**	**0.7473**

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Recurrent unsupervised anomaly detection: RUAD

CINECA

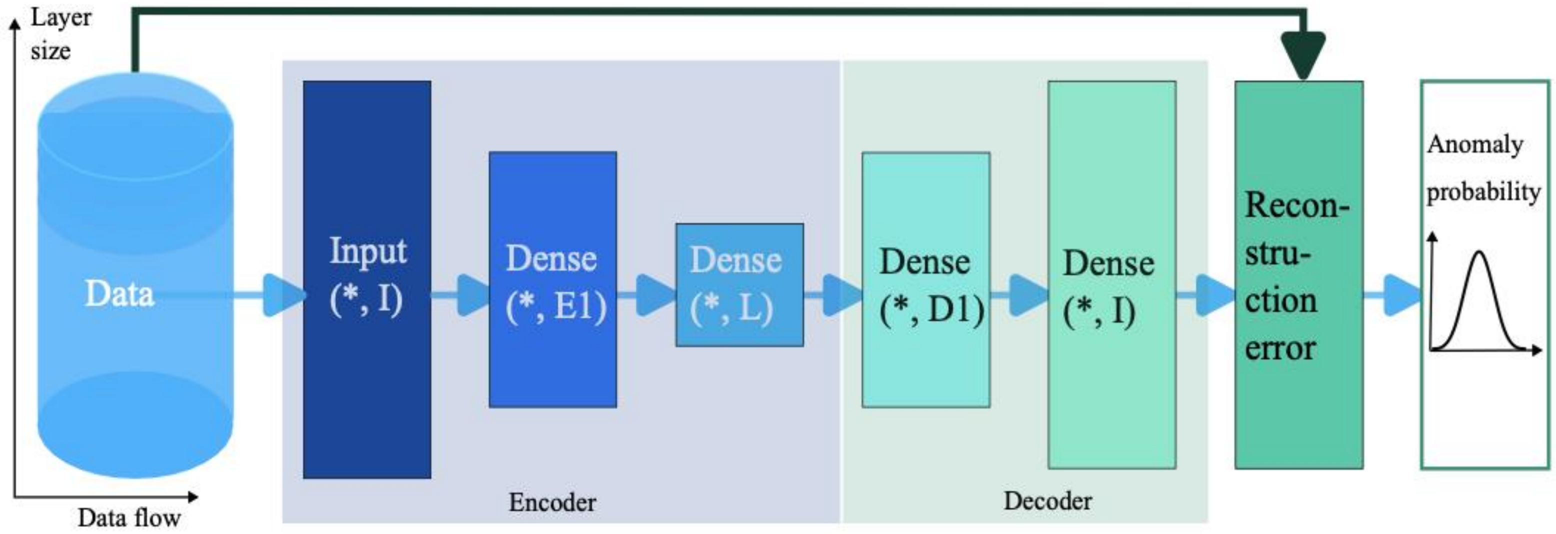

(a) Structure of baseline model - the dense autoencoder.

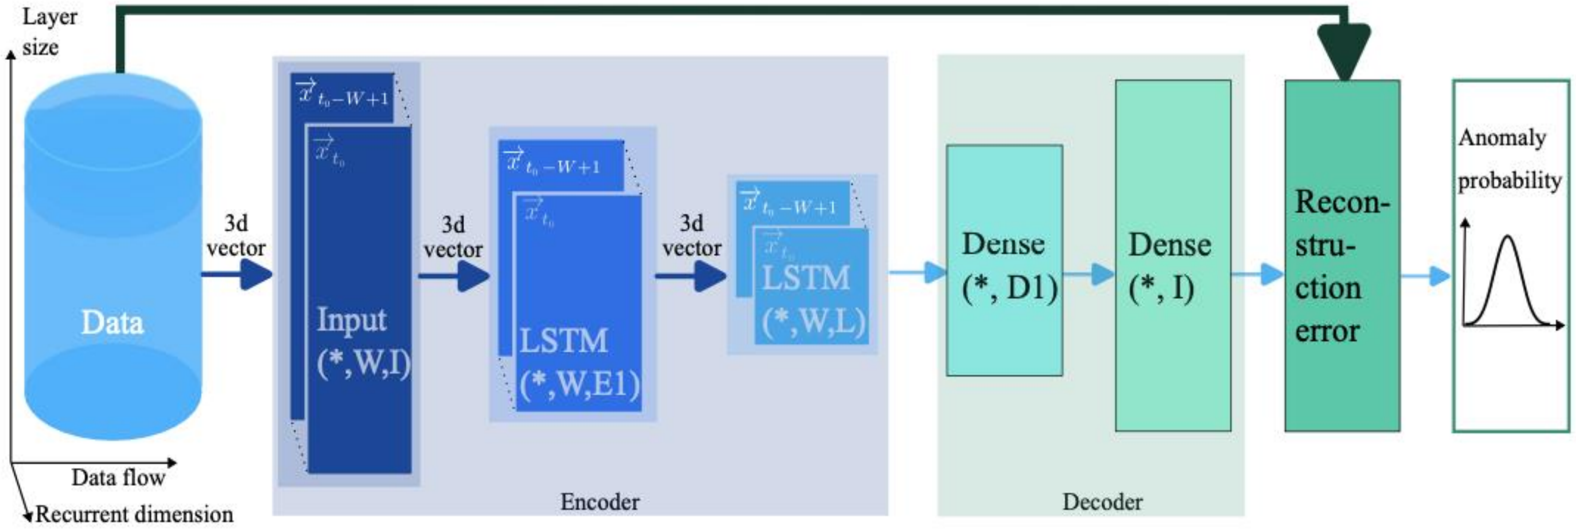

(b) Structure of the proposed RUAD model consisting of the LSTM encoder and dense decoder.

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

EXAMON applications – anomaly detection

Automated anomaly/fault detection via ML models

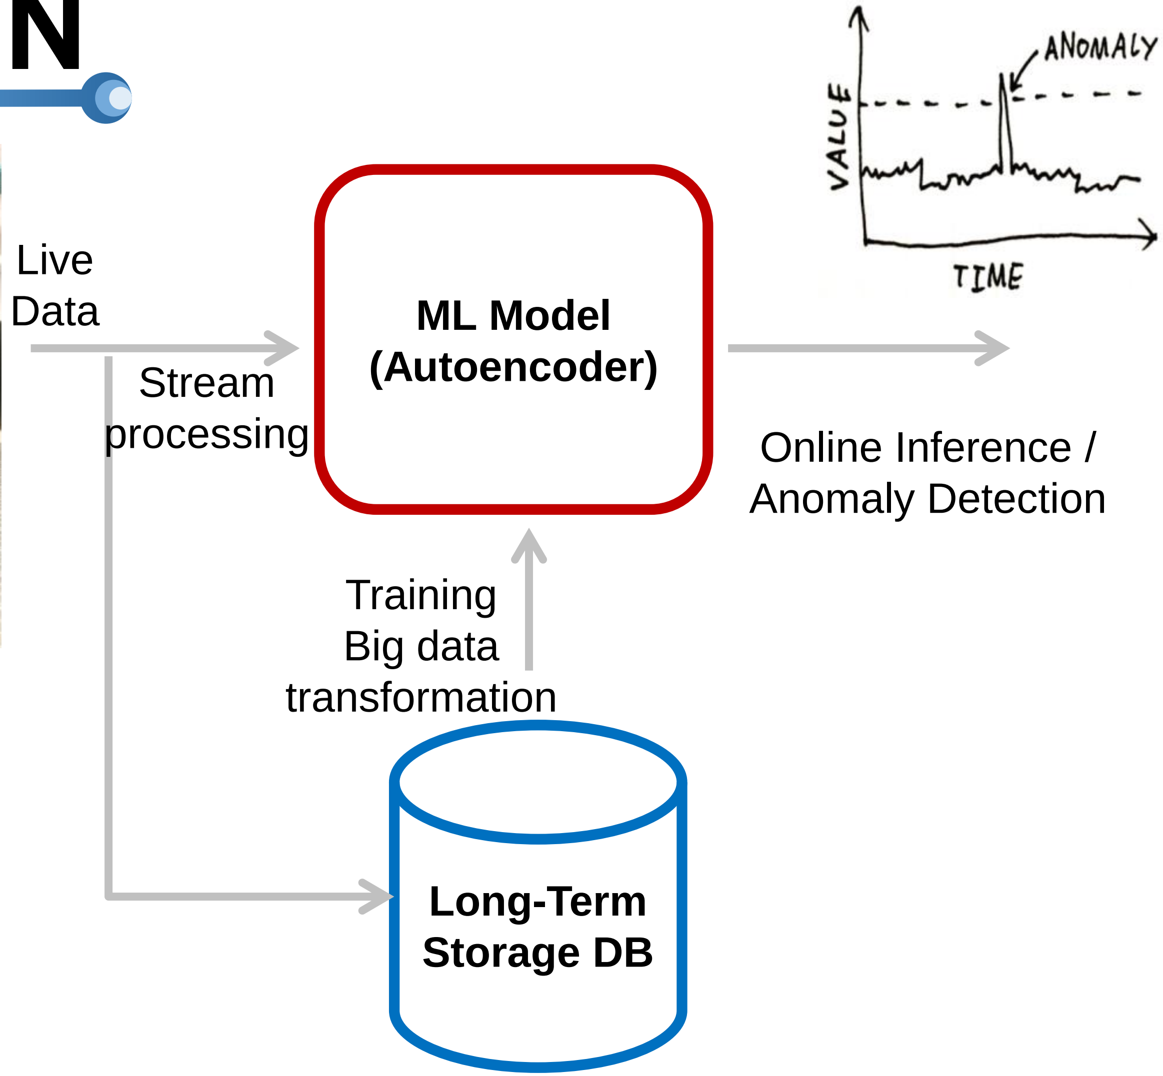

*[EAAI19] Borghesi et. al «**A semisupervised autoencoder-based approach for anomaly detection in high performance computing systems**»*
*[AICAS18] Borghesi et al. «**Online Anomaly Detection in HPC Systems**»*

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Thermal Hazard Prediction

CINECA

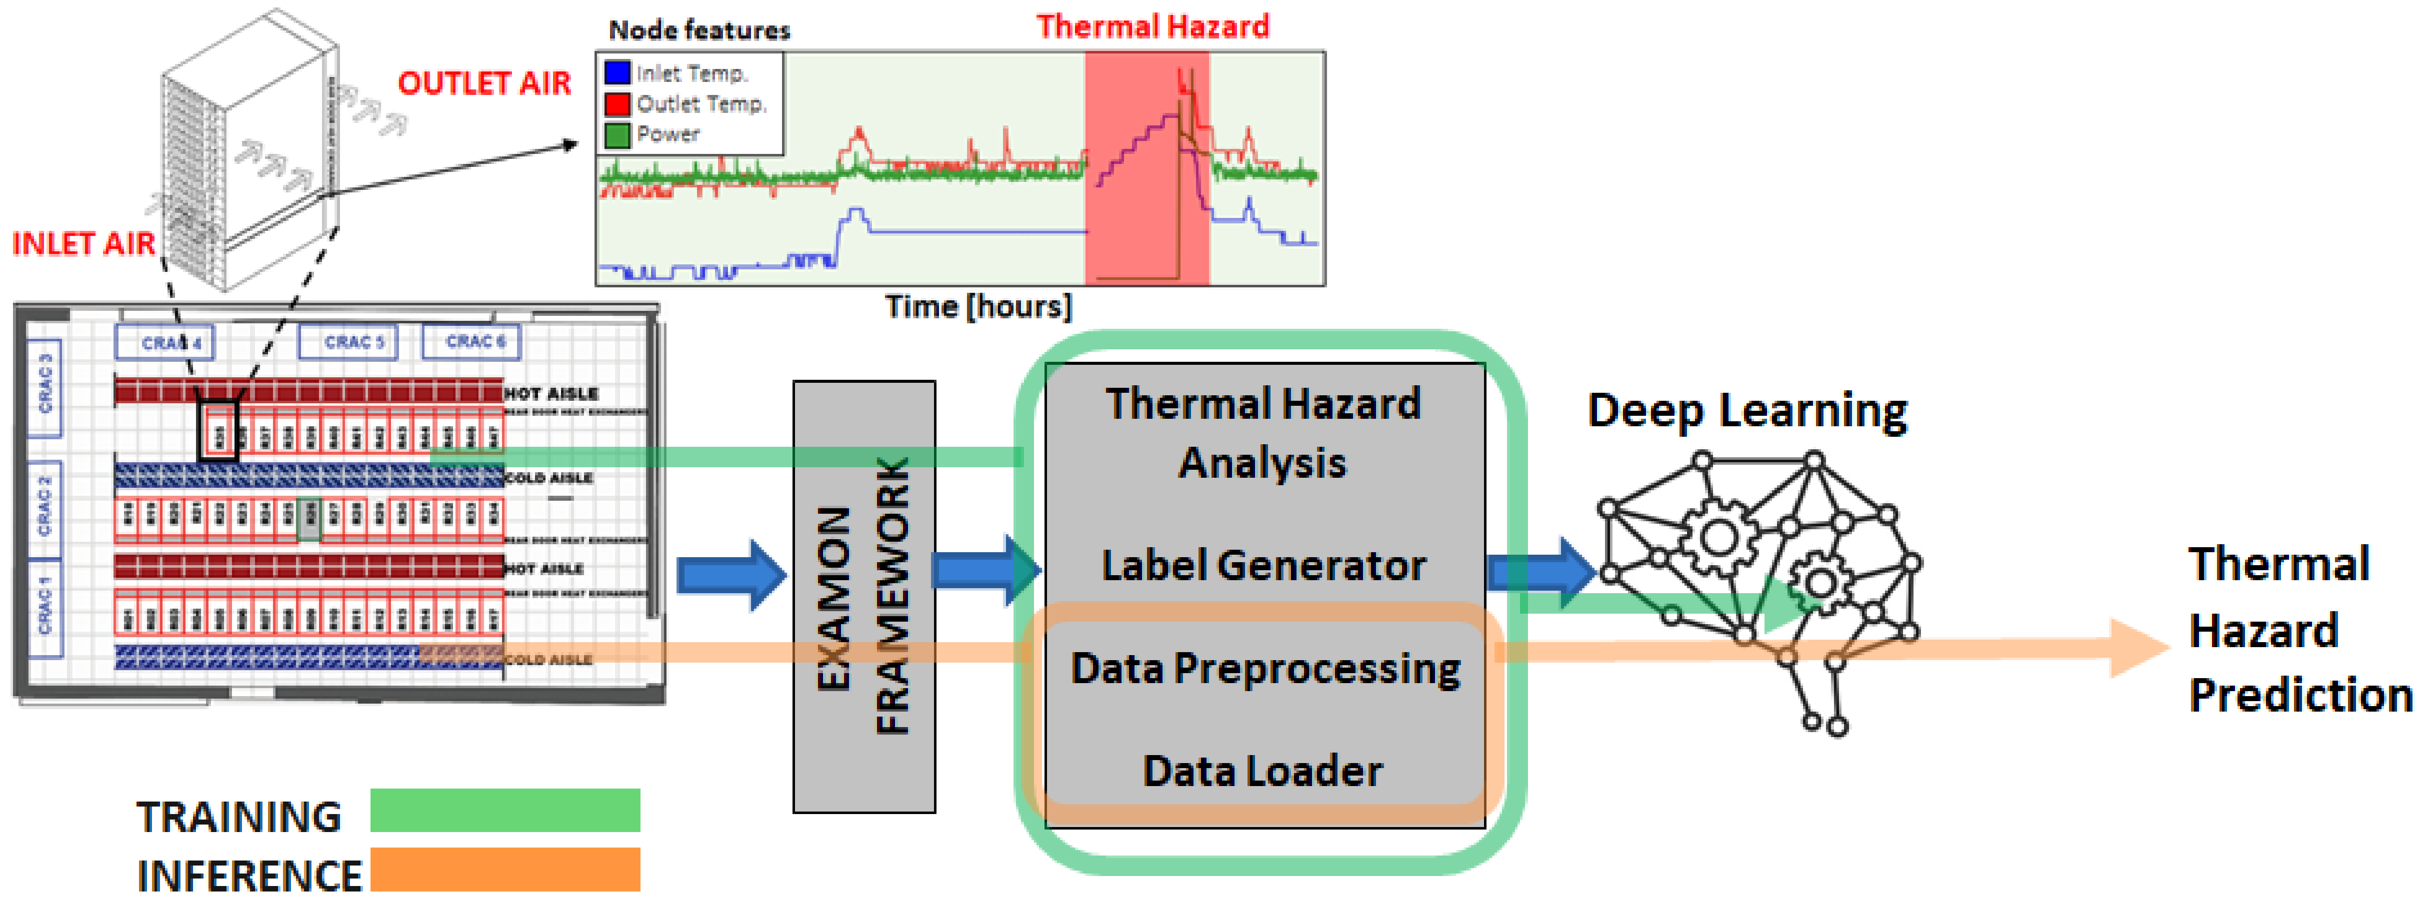

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Job power prediction

CINECA

1. Machine Learning models to predict the power consumption of HPC applications
2. Slurm Custom Extensions to schedule jobs based on their power
3. Interacts with power management

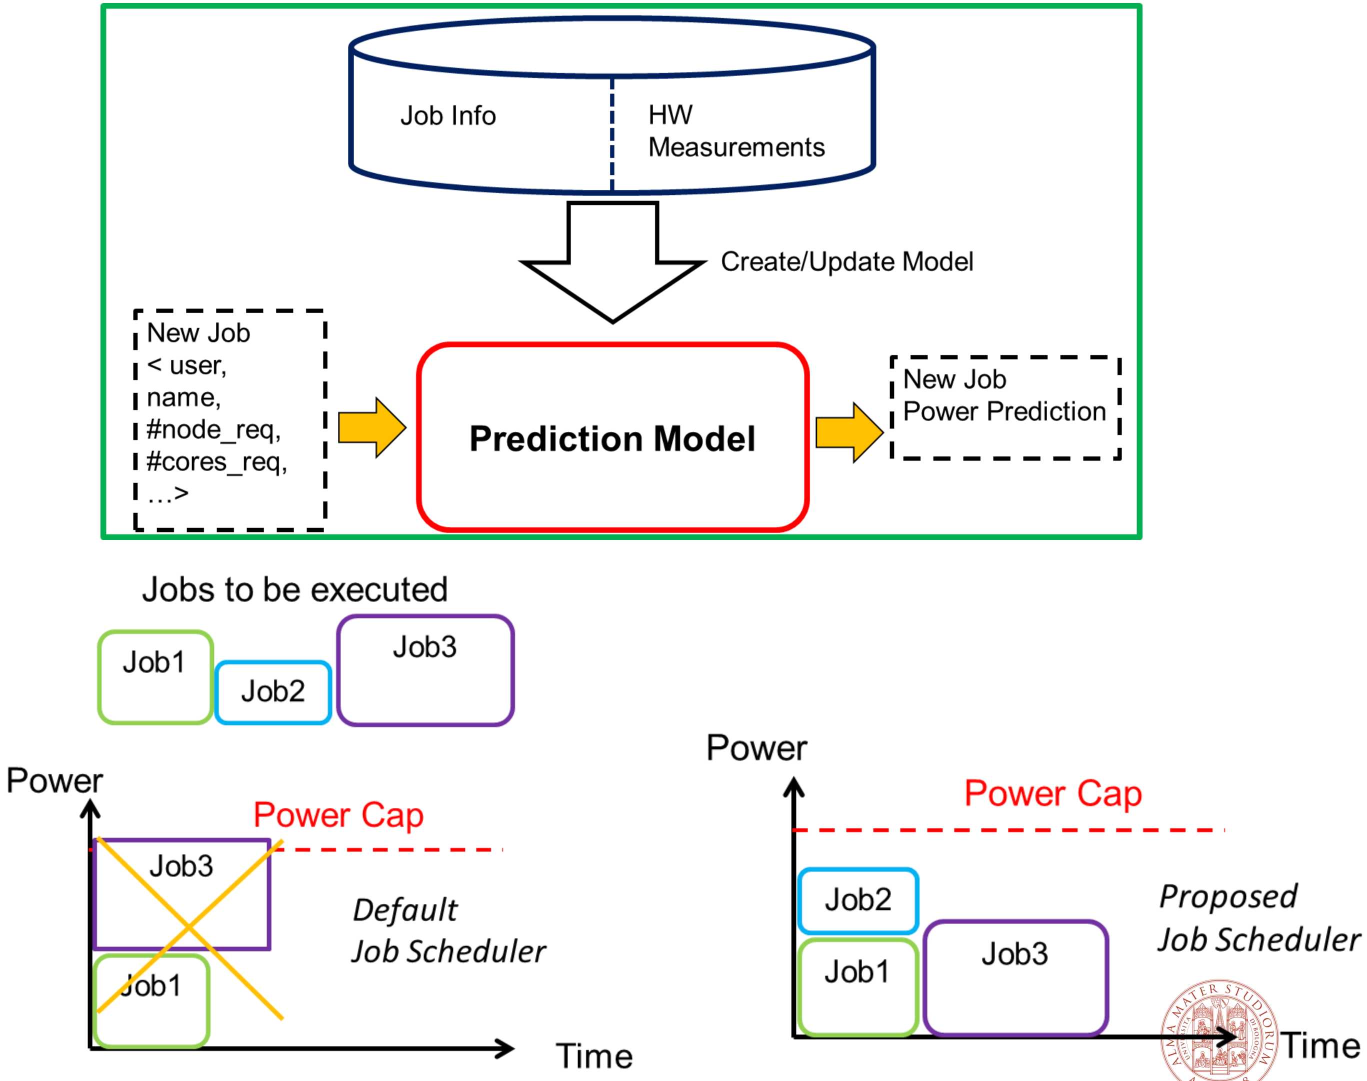

ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

Job power prediction

1. Machine Learning models to predict the power consumption of HPC applications
2. Slurm Custom Extensions to schedule jobs based on their power
3. Interacts with power management

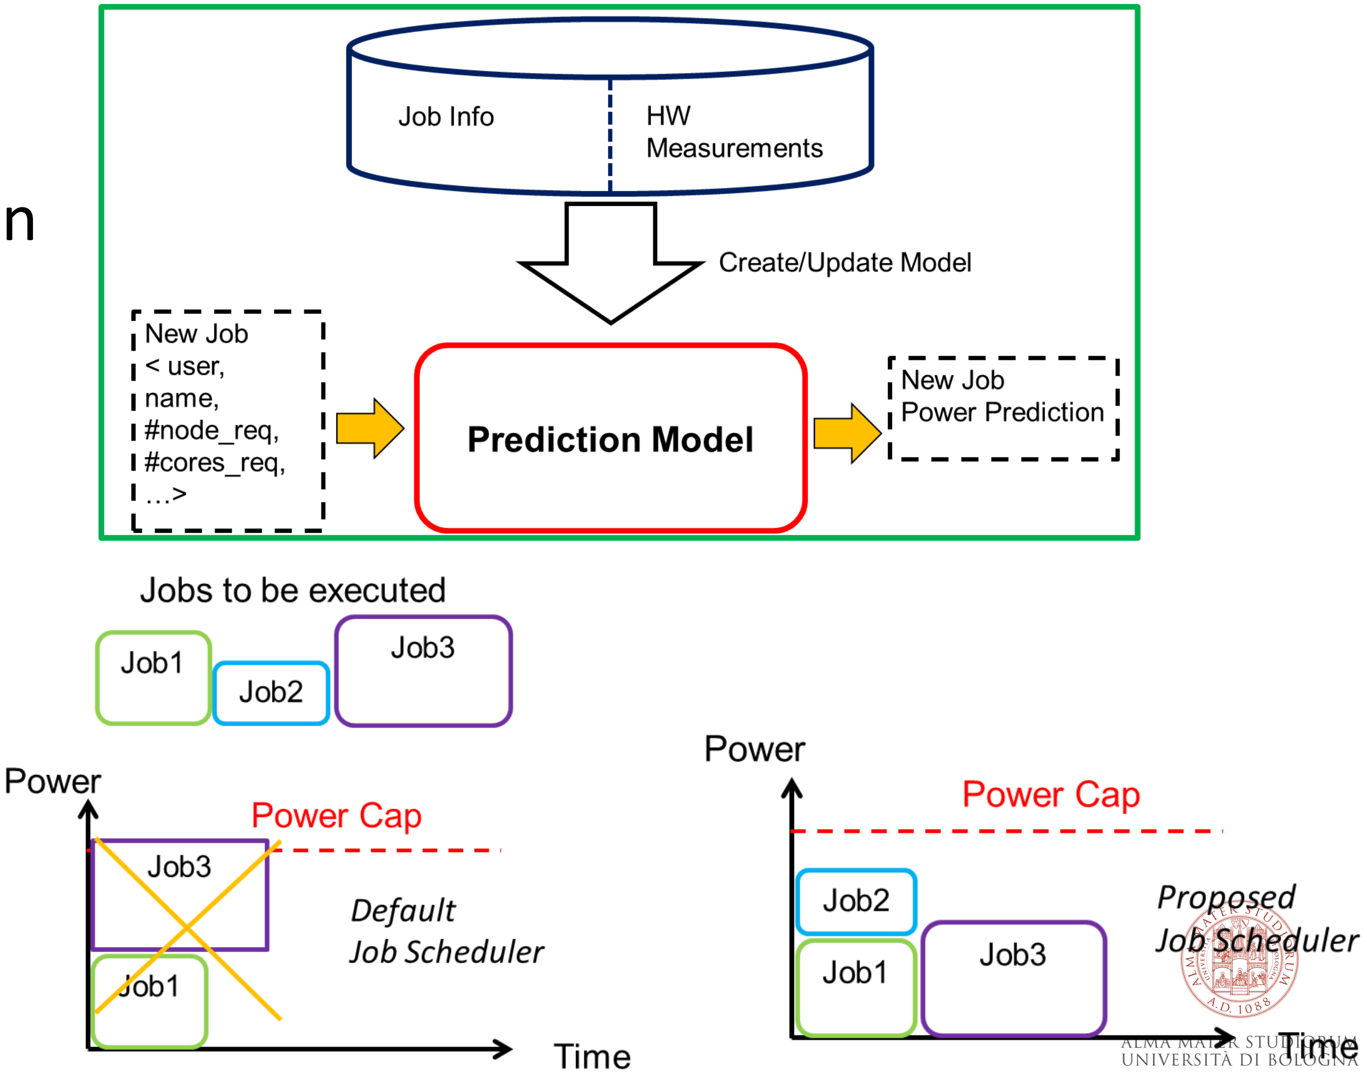

ExaData – open dataset – just released

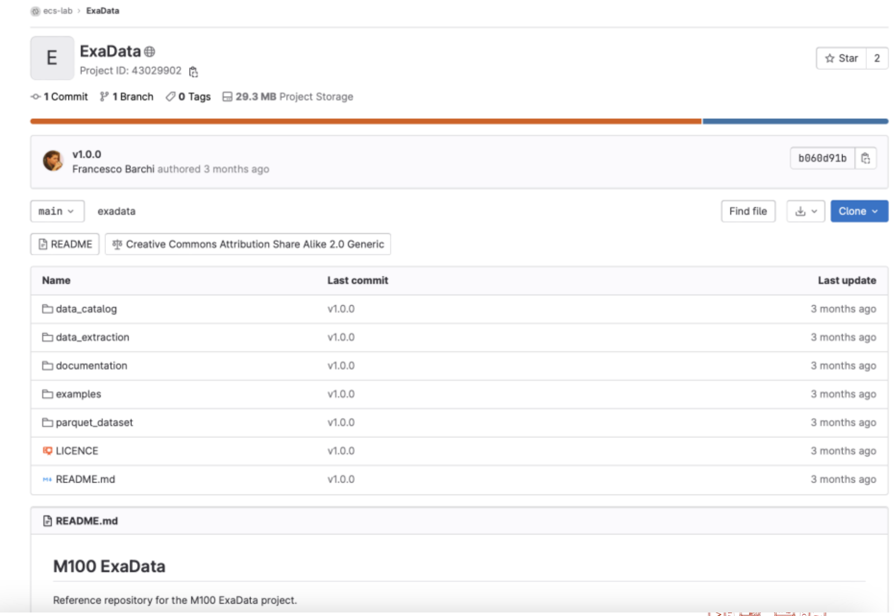

scientific **data**

Explore content | **About the journal** | **Publish with us**

nature > scientific data > data descriptors > article

Data Descriptor | Open Access | Published: 18 May 2023

M100 ExaData: a data collection campaign on the CINECA's Marconi100 Tier-0 supercomputer

Andrea Borghesi, Carmine Di Santi, Martin Molan, Mohsen Seyedkazemi Ardebili, Alessio Mauri, Massimiliano Guarrasi, Daniela Galetti, Mirko Cestari, Francesco Barchi, Luca Benini, Francesco Beneventi & Andrea Bartolini

Scientific Data **10**, Article number: 288 (2023) | Cite this article

1 Altmetric | Metrics

https://www.nature.com/articles/s41597-023-02174-3

https://gitlab.com/ecs-lab/exadata

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ExaData description

- 31 months of data
- 573 metrics, 980+ nodes, approx. 50 TB uncompressed
- Vertiv, Schneider, IPMI, Ganglia, Logics, Weather, Nagios, SLURM, Job table
- Hardware data, system monitoring data, external information
- Different sampling granularities (from seconds) to minutes
- Zenodo + Nature Dataset

Plugin	#Metrics	#Plugin-specific columns	Description
Vertiv	25	1	Mainly collects data from the air-conditioning units (CDZ) located in room F (Marconi 100) of Cineca. The plugin uses the RESTful API interface available on the individual devices to extract the most interesting metrics.
Schneider	164	1	Dedicated data collector designed to acquire data from an industrial PLC by accessing its HMI module (from Schneider Electric). The PLC controls the valves and pumps of the liquid cooling circuit (RDHx) of Marconi 100. It consists of two (redundant) twin systems controllable by two identical HMI panels, Q101 and Q102.The ExaMon plugin extracts and stores all the metrics available on both panels.
IPMI	104	1	Collects all the sensor data provided by the OOB management interface (BMC) of cluster nodes.
Ganglia	177	1	Connects to the Ganglia server (gmond), collects and translates the data payload (XML) to the ExaMon data model.
Logics	37	2	Data collection system already installed at Cineca. It is specialized for collecting power consumption data from equipment in the different rooms, typically using multimeters that communicate via Modbus protocol. The ExaMon plugin dedicated to collecting this data interfaces to the Logics database (RDBMS) via its REST API. NOTE: Since the translation process is fully automated, the same inconsistencies present in the original db may result in the ExaMon database: e.g., metric names in the Italian language, units of measure as metric name, etc.
Weather	10	0	Collects all the weather data related to the Cineca facility location (Casalecchio di Reno) using an online open weather service (https://openweathermap.org).
Nagios	1	5	Interfaces with a Nagios extension developed by CINECA called "Hnagios", collects and translates the data payload to the ExaMon data model.
SLURM	54	4	Collects aggregated data from the SLURM server; these information is gathered through ad hoc scripts created by CINECA system administrators.
Job table	1	89	Collects information regarding the jobs executed on the cluster (and store in the SLURM database); the information collected are those provided by users at submission time.

https://www.nature.com/articles/s41597-023-02174-3

https://gitlab.com/ecs-lab/exadata

Dataset

CINECA

Metrics/day (per plugin)

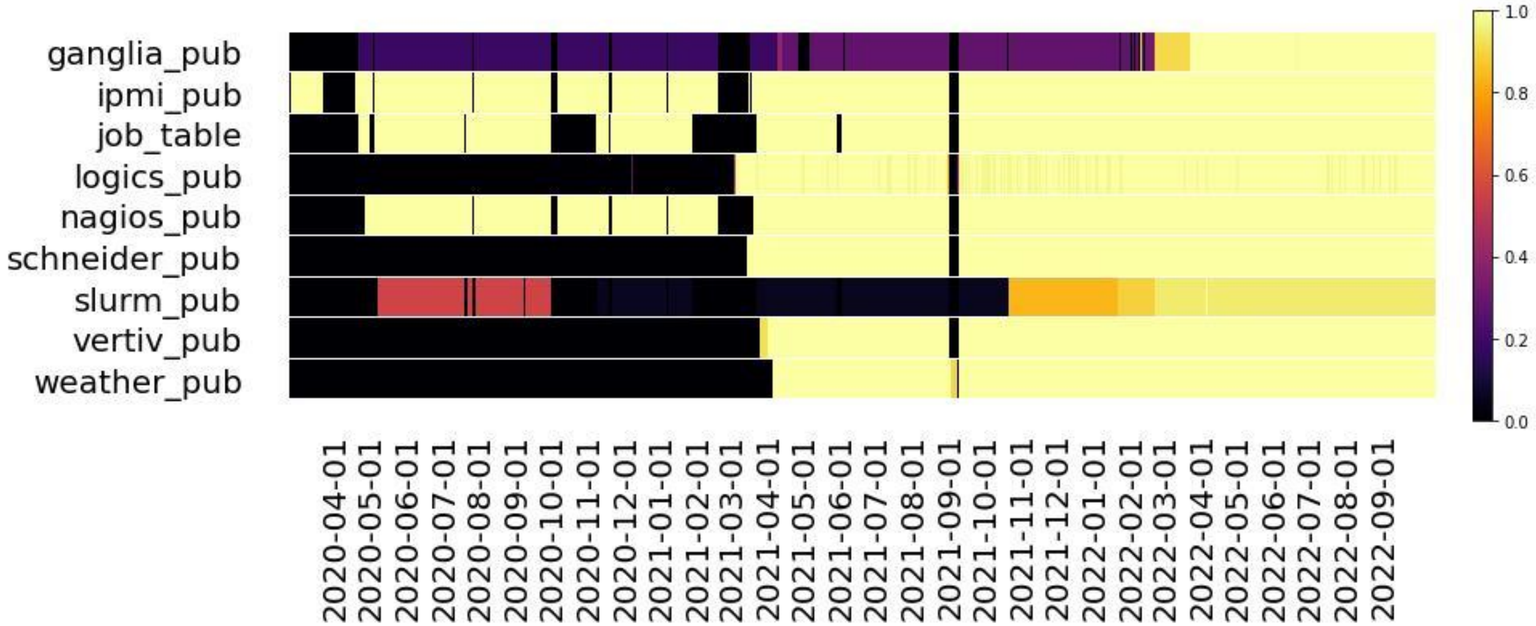

Samples/day (per plugin)

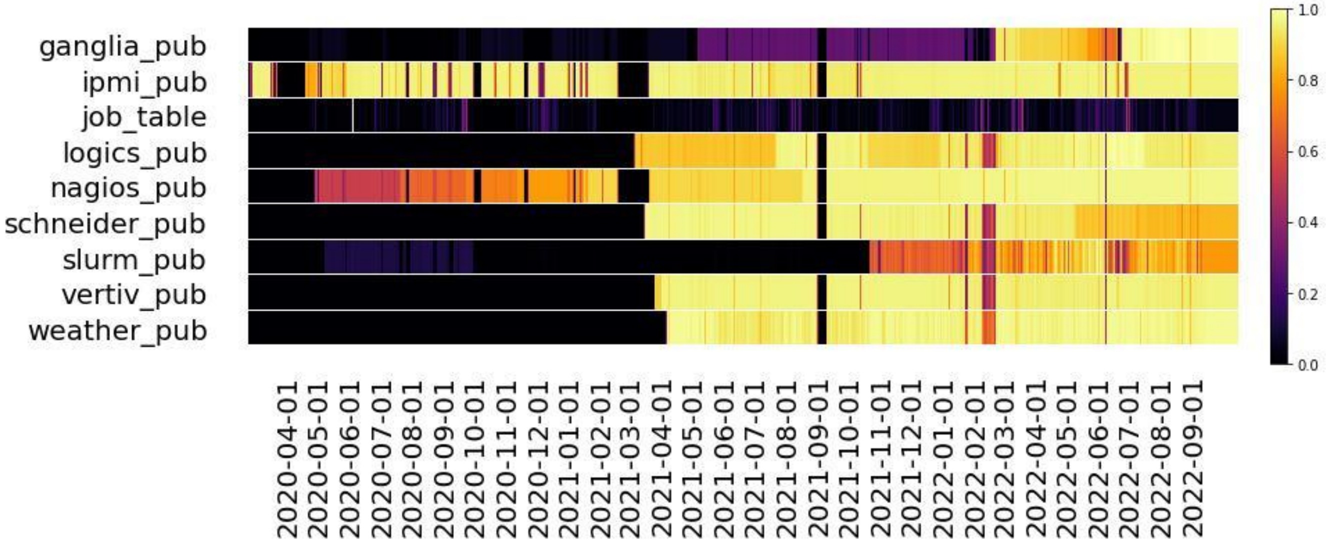

https://gitlab.com/ecs-lab/exadata

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Sample metric

CINECA

Repository structure

- `documentation` : descriptions of the plugins (ExaMon), including related metadata (metrics, tags); spatial distribution of racks.
- `examples` : applications of the dataset described in the "technical validation" part of the paper.
- `parquet_dataset/csv_to_parquet` : scripts used to produce the final Parquet dataset, starting from the extracted CSVs (ExaMon).
- `parquet_dataset/query_tool` : simple tool to load a slice of the dataset into a Pandas DataFrame.
- `parquet_dataset/node_aggregated_data` : scripts to create the anomaly detection dataset.
- `data_extraction` : scripts used to extract the data from ExaMon (in CSV format).
- `data_catalog` : data catalog.

References

Main technologies:

- Parquet (2.6): https://parquet.apache.org/
- PyArrow (9.0.0): https://arrow.apache.org/
- Pandas (1.5.1): https://pandas.pydata.org/

metric	value	year_month	plugin	node	timestamp
Gpu0_gpu_temp	37	22-Jul	ganglia_pub	982	2022-07-05 00:00:01+00:00
Gpu0_gpu_temp	37	22-Jul	ganglia_pub	982	2022-07-05 00:00:12+00:00
Gpu0_gpu_temp	36	22-Jul	ganglia_pub	982	2022-07-05 00:00:23+00:00
Gpu0_gpu_temp	36	22-Jul	ganglia_pub	982	2022-07-05 00:00:33+00:00
Gpu0_gpu_temp	36	22-Jul	ganglia_pub	982	2022-07-05 00:00:44+00:00
Gpu0_gpu_temp	36	22-Jul	ganglia_pub	982	2022-07-05 00:00:54+00:00
Gpu0_gpu_temp	36	22-Jul	ganglia_pub	982	2022-07-05 00:01:05+00:00
Gpu0_gpu_temp	36	22-Jul	ganglia_pub	982	2022-07-05 00:01:16+00:00
Gpu0_gpu_temp	36	22-Jul	ganglia_pub	982	2022-07-05 00:01:27+00:00
Gpu0_gpu_temp	36	22-Jul	ganglia_pub	982	2022-07-05 00:01:37+00:00
Gpu0_gpu_temp	36	22-Jul	ganglia_pub	982	2022-07-05 00:01:48+00:00
Gpu0_gpu_temp	36	22-Jul	ganglia_pub	982	2022-07-05 00:01:59+00:00
Gpu0_gpu_temp	36	22-Jul	ganglia_pub	982	2022-07-05 00:02:10+00:00

https://www.nature.com/articles/s41597-023-02174-3