



# Integrating AI into Data-Center Operations

Dr Marco Lorusso, INFN-CNAF



# Context

- The increasing **complexity** and **scale** of modern data centers generate operational **environments** where the ability to **detect anomalies, predict failures**, and **improve resource** usage is becoming **critically important**.
- Recent advances in **machine learning** and **artificial intelligence** offer **powerful techniques** for extracting actionable insights from **heterogeneous monitoring data**
- In this contribution, we explore an **AI-driven approach** that can be applied to a specific **large-scale computing** facility hosted at **INFN-CNAF** in Bologna.

# Data Complexity

- **Multiple Data sources**

- Logs, metrics, monitoring data, and system outputs

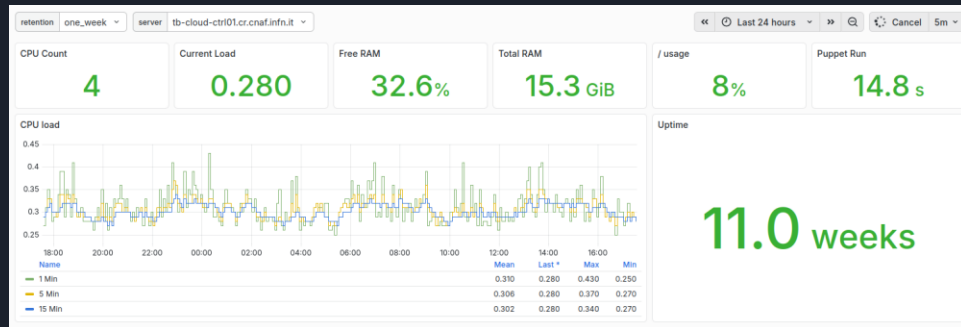
- **Heterogeneous formats**

- Structured (JSON), semi-structured (logs), and time-series metrics

- **Large data volumes**

- High and continuously growing data amounts

- **Processing Complexity**



```
index_prefix: farming-cloud message: 2026-03-09 17:35:58.184 3748685 INFO neutron.plugins.ml2.drivers.openvswitch.agent.ovs_neutron_agent [None req-85805fa5-1004-4edb-bbcb-175a15e5a28e - - - - -] Agent rpc_loop - iteration:1483907 started tags: farm-cloud log_global_request_id: None log_user_id: - log_project_domain: - log_module: neutron.plugins.ml2.drivers.openvswitch.agent.ovs_neutron_agent log_domain: - log_level: INFO log_local_request_id: req-85805fa5-1004-4edb-bbcb-175a15e5a28e log_msg: Agent rpc_loop - iteration:1483907 started log_projectname: - log_file_path: /var/log/neutron/openvswitch-agent.log log_project_id: - log_username: - log_pid: 3748685 ecs.version: 8.0.0 @version: 1 host.name: tb-cloud-net01.cr.cnaf.infn.it service: neutron, openvswitch-agent @timestamp: Mar 9, 2026 @ 17:35:58.184

index_prefix: farming-cloud message: 2026-03-09 17:35:56.952 520833 INFO neutron.plugins.ml2.drivers.openvswitch.agent.ovs_neutron_agent [None req-7465a99-df25-461d-82b6-57ea04502eff - - - - -] Agent rpc_loop - iteration:1483966 started tags: farm-cloud log_global_request_id: None log_user_id: - log_project_domain: - log_module: neutron.plugins.ml2.drivers.openvswitch.agent.ovs_neutron_agent log_domain: - log_level: INFO log_local_request_id: req-7465a99-df25-461d-82b6-57ea04502eff log_msg: Agent rpc_loop - iteration:1483966 started log_projectname: - log_file_path: /var/log/neutron/openvswitch-agent.log log_project_id: - log_username: - log_pid: 520833 ecs.version: 8.0.0 @version: 1 host.name: tb-cloud-net01.cr.cnaf.infn.it service: neutron, openvswitch-agent @timestamp: Mar 9, 2026 @ 17:35:56.952

index_prefix: farming-cloud message: Mar 9 17:35:56 tb-cloud-ctrl01 nova-novncproxy[2648265]: 2026-03-09 17:35:56.435 2648265 INFO nova.console.websOCKETproxy [-] handler exception: [SYS] unknown error (.ssl.c:2651)#03[0m tags: farm-cloud log_file_path: /var/log/messages log_level: ERROR log_msg: [-] handler exception: [SYS] unknown error (.ssl.c:2651)#03[0m log_pid: 2648265 log_module: nova.console.websOCKETproxy ecs.version: 8.0.0 @version: 1 host.name: tb-cloud-ctrl01.cr.cnaf.infn.it service: messages, nova-novncproxy @timestamp: Mar 9, 2026 @ 17:35:56.435 _id: sbp005w09_VzdxmZkY01 _type: - _index: farming-cloud-2026.03.09 _score: -

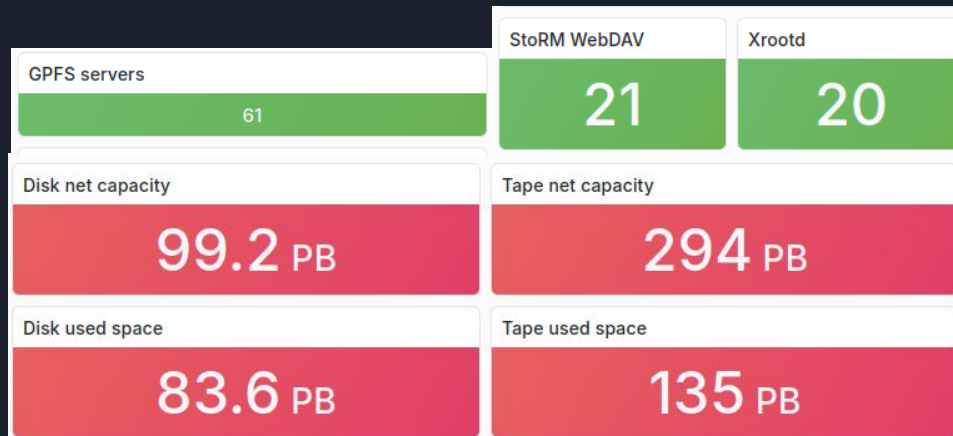
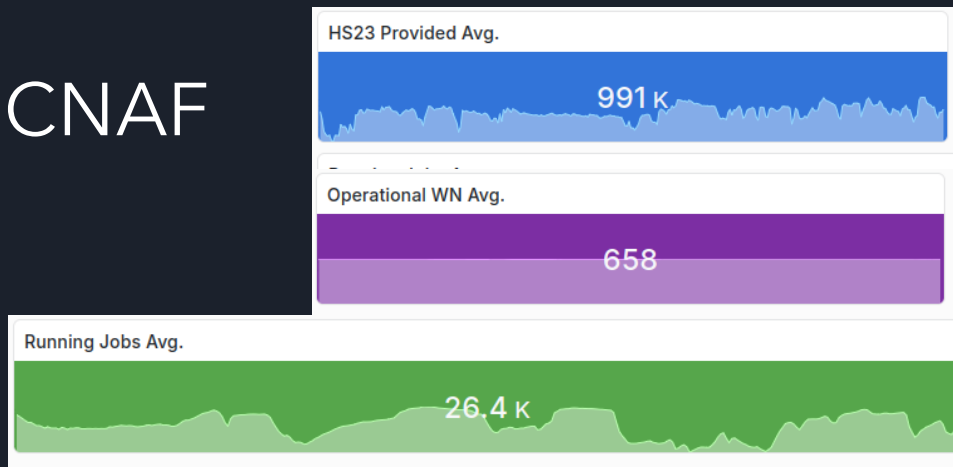
container_image.name: rancher/nginx-ingress-controller:v1.14.3-hardened1 index_prefix: farming-k8s @version: 1 kubernetes.daemonset.name: rke2-ingress-nginx-controller kubernetes.pod.name: rke2-ingress-nginx-controller-h6c97 kubernetes.namespace: kube-system kubernetes.node.name: rke2-agent-1.cloudcnaf kubernetes.container.name: rke2-ingress-nginx-controller tags: k8s, rke2-prod event.timezone: +00:00 message: 131.154.192.148 - - [09/Mar/2026:16:35:56 +0000] "GET /login HTTP/1.1" 200 6947 "-" "Sensu-HTTP-Check" 178 0.003 [ctao-ctao-lam-login-service-8080] [] 10.42.6.43:8080 6947 0.004 200 bbae911be09508e978c565add0a70e35 @timestamp: Mar 9, 2026 @ 17:35:56.326 _id: YRp005w09_VzdxmZkZ01 _type: - _index: farming-k8s-2026.03.09 _score: -
```

```
{
  "_index": "farming-cloud-2026.03.09",
  "_id": "aRp005w09_VzdxmZ1TGI",
  "_version": 1,
  "_score": null,
  "_source": {
    "index_prefix": "farming-cloud",
    "input": {},
    "message": "2026-03-09 17:35:58.184 3748685 INFO neutron.plugins.ml2.drivers.openvswitch.agent.ovs_neutron_agent [None req-85805fa5-1004-4edb-bbcb-175a15e5a28e - - - - -] Agent rpc_loop - iteration:1483907 started",
    "tags": [
      "farm-cloud"
    ],
    "log": {

```



# Data Complexity @ CNAF

- **Large-Scale Infrastructure @ CNAF**
  - CERN Tier-1 data center with large **computing** and **storage** capacity
- **High Data Volume**
  - Petabyte-scale storage and thousands of running jobs
  - Metrics, logs, and monitoring data from hosts and services
- **Analytical Challenge**
  - Large and diverse datasets require novel approaches for analysis



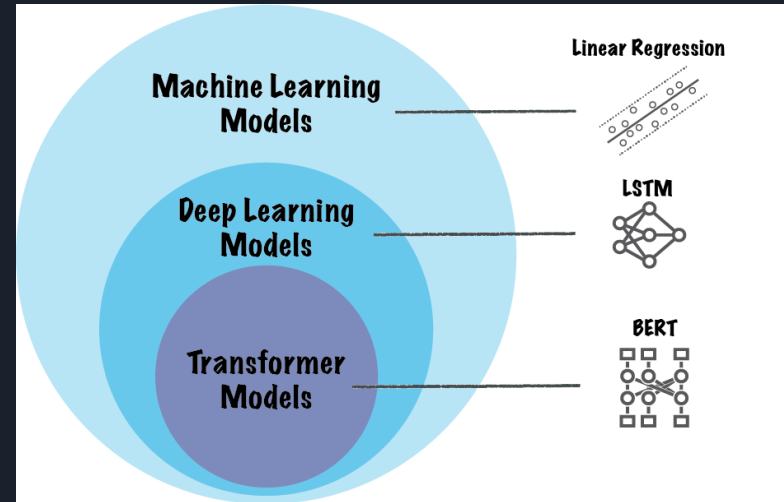
# Different ML Approaches for AD

- **Autoencoder**: trained to **reconstruct** their **inputs**:
  - The **encoder compress** the input into a smaller latent space.
  - The **decoder reconstructs** from the compressed representation.
- **Recurrent Neural Network**: designed for processing sequential data
  - the **output** of a neuron at **one time step** is **fed back** as input at the **next time step**
- **Transformer**: based on the **multi-head attention mechanism**:
  - designed to model relationships **between elements in a set or sequence**.



# Transformer

- Traditional ML models often impose **fixed structural assumptions**, such as **locality** or **sequential processing**;
- Transformers instead rely on **attention mechanisms** to learn **relationships directly from the data**;
  - hosts in a distributed system do not have a meaningful spatial ordering;
- Self-attention allows each element in an input representation to **interact with all other elements**;
  - The model learns which elements of the inputs are **relevant to each other**, enabling **flexible modeling** of complex interaction patterns.



# Transformer Architecture

A typical Transformer layer consists of two main components:

## 1. Self-attention:

- exchanges information between elements

```
attn = tf.keras.layers.MultiHeadAttention(num_heads, key_dim= lat_dim //  
                                         num_heads, name = "multiheadattention") (e, e)
```

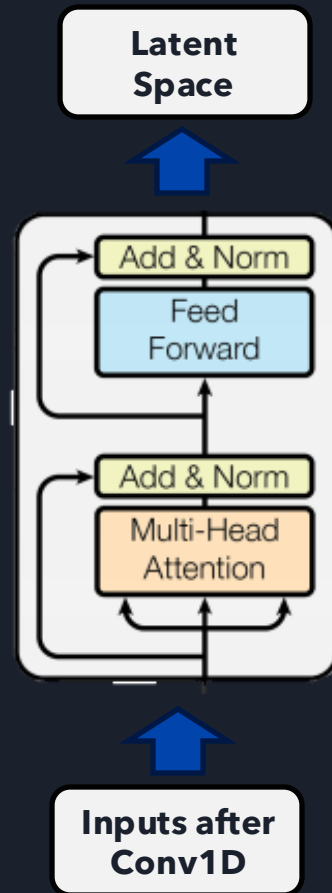
```
ee = tf.keras.layers.Add(name = "attn_resid") ([e, attn])  
ee = tf.keras.layers.LayerNormalization(name = "attn_ln") (ee)
```

## 2. Feed-forward transformation:

- independently refines each element's representation

```
ffn = tf.keras.layers.Dense(lat_dim * ff_trans_dens, activation = //  
                             "relu", name = "ffn1") (ee)  
ffn = tf.keras.layers.Dense(lat_dim, activation = "relu", name = //  
                             "ffn1") (ffn)
```

```
eee = tf.keras.layers.Add(name = "ffn_resid") ([ee, ffn])  
eee = tf.keras.layers.LayerNormalization(name = "fnn_ln") (eee)
```





# Data Overview

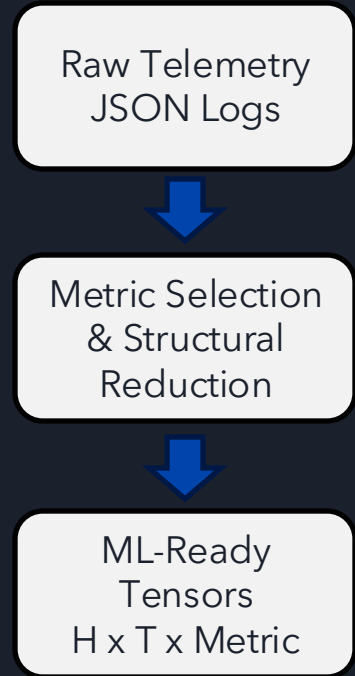
Operational telemetry collected from services managed by the **CNAF Farming Group**, which operates the computing clusters supporting distributed scientific workloads and large-scale data processing.

The data is a **multivariate time series** across **multiple hosts**, which include:

- CPU usage and load
- memory utilization and disk I/O statistics
- process-level resource metrics

# Data Preparation and Processing

- Starting from multi-gigabyte JSON logs with redundant structure
- For Anomaly Detection data must be processed to:
  - Isolate informative system-level metric
  - Remove noisy and low-signal process-level measurements
  - Remove missing values and inconsistent reporting intervals
- The resulting Dataset contains:
  - **4086** (check number) **time** snapshots for **304** compute **nodes** with **36** system-level **metrics** each = a timeseries with ~11k features each step. (~44M input values)







# Signal injection

- Data coming from nodes in production
  - Real anomalies are rare
- To evaluate the detection method, anomalies are **synthetically injected** into validation data by simulating realistic events affecting random hosts and metrics:
  - CPU usage spikes
  - Sustained workload shifts
  - Disk I/O bursts
  - Memory leaks
  - Node outages
  - Abnormal process behaviour

# Training v Test dataset



Time series **with** anomalies



Time series **without** anomalies

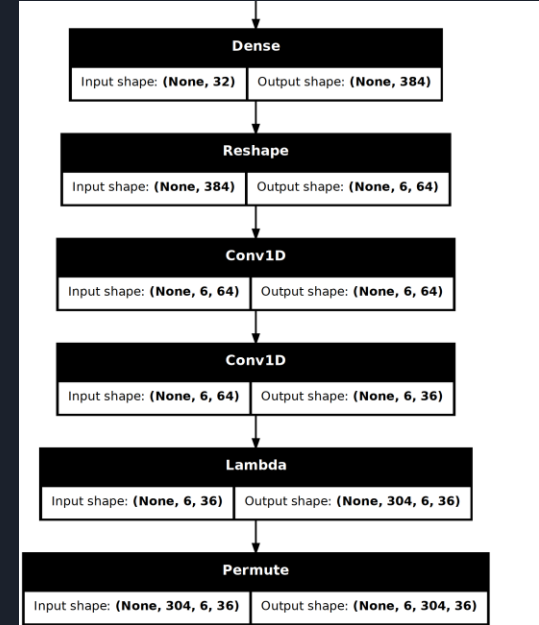
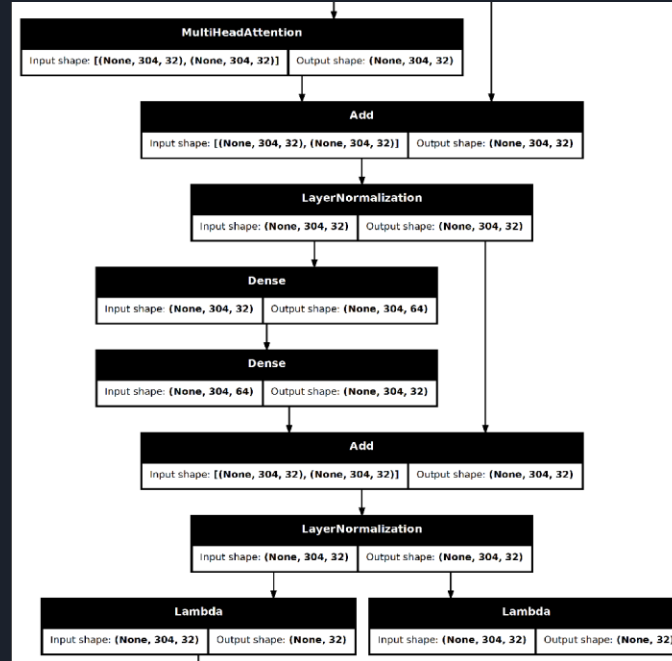
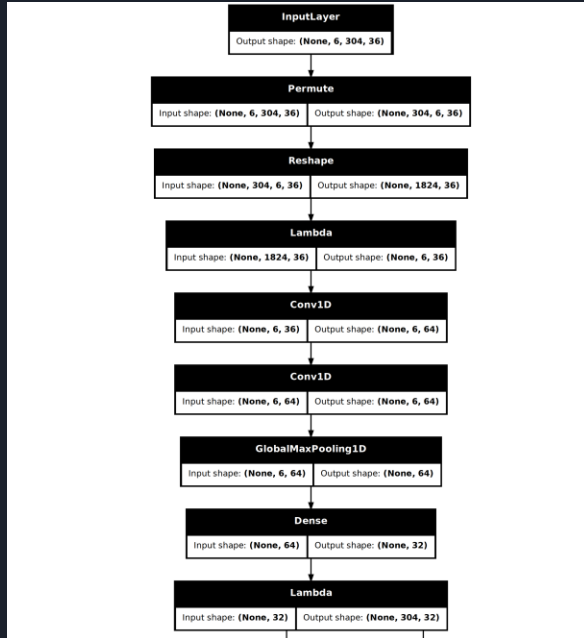


# Model Architecture

- **Input representation:** telemetry windows arranged as **time steps** × **hosts** × **features**;
- **Encoder:**
  - Per-host **temporal Conv1D layers** extract short-term patterns
  - **Global max pooling** produces a compact host representation
- Cross-host modeling:
  - Host embeddings processed with **multi-head self-attention**
  - Residual and feed-forward layers refine the contextual representation
  - A **global latent vector** summarizes the system state
- Decoder:
  - Host embeddings expanded and passed through **temporal Conv1D layers**
  - Reconstructs the original telemetry window

# Model Architecture


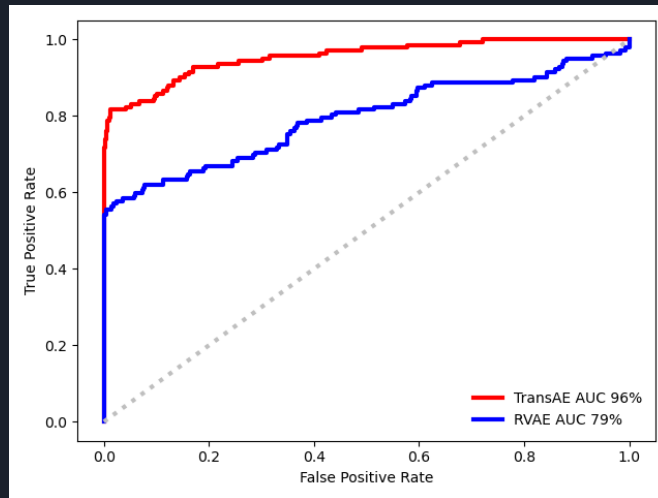
# parameters: 87,524



# Results

- Injected anomalies provide controlled ground truth labels.
- Anomaly score:
  - Compute reconstruction error **per host**
  - Final score = **average error of the top-k most anomalous hosts**
- This allows quantitative evaluation of the anomaly detection model, e.g. using a ROC curve, its AUC and F1-score:

- $$F1 = \frac{2 \times TP}{2 \times TP + FP + FN}$$





# Conclusion

- **Autoencoders** provide a powerful tool for unsupervised anomaly detection in multivariate telemetry;
- **Attention mechanisms** enable modeling dependencies between hosts, allowing to detect cluster-wide behaviour;

## Future Directions:

- **Improve model robustness** by incorporating additional feature engineering;
- Create **more complex and generalized architectures** for modeling temporal and cross-host dependencies;
- Develop a **dedicate data pipeline** for ingestion and preprocessing, enabling **systematic experimentation** and facilitating **deployment in the future**.

