

AGENT-DRIVEN INTELLIGENT MONITORING

Hybrid Reasoning for Fault Detection

Supervisor:
Prof. Andrea Omicini

Candidate:
Davide Di Marco

Co-supervisors:
Prof. Giovanni Ciatto
Ing. Fabio Capriati


Academic Year 2025–2026



The Idea, Motivation and Context


Modern IT systems are increasingly distributed, dynamic, and difficult to manage.

Typical monitoring process



Problem Statement

Traditional monitoring often stops at anomaly detection.



Goals of the Work



Support Explainable Diagnosis




Automate Fault Investigation using
AI



Go Beyond Anomaly Detection


Create structured diagnostic process able to transform monitoring data into **useful information** for the **technical operator**.

Development Approach and Verification




Proposed Architecture - Components

The system is composed by different independent component.



Monitoring and Anomaly Detection

Metrics and logs are collected continuously and stored in OpenSearch that trigger anomaly.




Feature breakdown Anomaly occurrences



Agents, Tools and Knowledge Base

Specialised agents Team cooperate to investigate incidents and retrieve shared system knowledge or evidence when needed.



Evaluation Approach

The prototype was evaluated in a controlled environment.



*Can the proposed architecture
autonomously investigate anomalous
events and **produce evidence-based**
diagnostic information?*

5

Monitored Machines

40

Diagnostic runs

2

Scenarios

4

Configuration Profile

CPU Spike

RAM Leak

Available evidence and knowledge

- **Metrics**
Resource usage and latency
- **Logs**
Application and system events
- **Runtime state**
Processes, containers, network checks
- **Knowledge Base**
All available documentation of the system

Hardware Environment

Windows host · **RTX 4060 8 GB** · Local Ollama inference

Main Results

The evaluation measured diagnostic quality and execution time across four reasoning profiles.

40 diagnostic runs · 10 runs per profile · Useful = correct or partial diagnosis

Quality / Time Comparison

Reasoning profile	Correct	Partial	Useful	Mean time	Evidence coverage
Baseline (ReACT step=6)	60%	10%	70%	224.84 s	90.0%
More reasoning (ReACT step=10)	50%	30%	80%	272.44 s	100.0%
Different model	0%	0%	0%	35.95 s	66.7%
Higher LLM temperature	30%	20%	50%	313.64 s	76.7%

Conclusions and Future Work

The prototype demonstrates that anomaly detection can be extended into an autonomous diagnostic process.

Main conclusions

Efficiency

less manual investigation

Operator Support

evidence, logs and guidance

Feasibility

end-to-end prototype

The system helps the technical operator by organising evidence and supporting remediation.

Future Work

Hardware

more powerful GPU
higher parallelism

Reasoning

stronger LLM models
prompt refinement

Evaluation

more faults
broader environments

The project is **also being considered by other companies** and is currently in an **experimental and evaluation phase**.

VIRTUAL OPERATIONS TEAM


Agent team activity

Specialised agents are presented as the technical roles they emulate inside an IT operations team.

TEAM STATE

IDLE

0 incidents with active agent context



Thank you!