

Operational Intelligence for Computing Operations, enabled by Generative AI

Tuesday 16 September 2025 17:35 (5 minutes)

Abstract

This proposal outlines a plan to relaunch the Operational Intelligence (OpInt) initiative, leveraging recent advances in Generative AI and AI Agent (AIA) technology to address the escalating complexity of distributed computing operations at CERN. While previous OpInt efforts demonstrated the value of data-driven insights, the landscape has now fundamentally shifted. AIAs offer a transformative opportunity to automate and simplify a wide range of operational tasks, moving beyond traditional analytics to create truly proactive and intelligent systems. The project will focus on two primary areas: enhancing distributed computing operations and improving on-site computing intelligence.

Background

The management of the Worldwide LHC Computing Grid (WLCG), a vast, heterogeneous, and distributed infrastructure, still requires significant manual intervention from developers, shifters, and site administrators. This is a substantial operational cost and a source of potential delays. Our past work on Operational Intelligence, documented in *“Operational Intelligence for Distributed Computing Systems for Exascale Science”* [1] and *“Preparing Distributed Computing Operations for the HL-LHC Era With Operational Intelligence”* [2], laid the groundwork by demonstrating the potential of machine learning to analyze operational data. However, the current evolution of AI, particularly the emergence of large language models and autonomous agents, presents a unique and timely opportunity to restart this work.

Proposed Solutions and Use Cases

We will focus on a phased approach, with initial efforts concentrated on the following high-impact use cases:

- **Automated Diagnostics for Data Transfers and Site Operations:** Create AIAs to monitor system logs and data transfer services (e.g., FTS). These agents would:
 - **Predict Failures:** Analyze log patterns from FTS transfers to predict potential failures before they occur.
 - **Streamline Troubleshooting:** When an error is detected, the AIA would generate a preliminary diagnostic report, correlating the error with known issues (e.g., misconfigured batch nodes, network issues, recent software updates etc.), thereby simplifying the lives of shifters and site administrators. A generative AI could also help operators, pointing them to the correct direction and providing all relevant information from various sources in a user-friendly centralized environment.
- **GGUS Ticketing Intelligence:** Develop an AIA to assist with the Global Grid User Support (GGUS) ticketing system. The agent would:
 - **Proactively Assist Users:** When a user submits a ticket, the AIA would analyze the problem description and automatically suggest similar, previously resolved issues, helping users find solutions more quickly.
 - **Automate Diagnostics for Shifters/Admins:** For new tickets, the AIA would analyze the ticket text and historical data, hinting at potential root causes. For example, it could correlate a specific error message with a past configuration change or an outage at a particular site.
- **On-site Intelligence:** Deploy AIAs at computing sites to correlate experiment-specific errors with underlying infrastructure issues. For example, an AIA could monitor job failures from a specific experiment at a given site and flag if a particular batch node is consistently misconfigured or broken. Another valuable example could be the anomaly detection of Storage (e.g. EOS) logs, this will enhance the precision of alarms, reduce false positives, and provide operators with actionable, early-warning insights into potential failures.

References

1. Operational Intelligence for Distributed Computing Systems for Exascale Science —<https://doi.org/10.1051/epjconf/202024503017>
2. Preparing Distributed Computing Operations for the HL-LHC Era With Operational Intelligence —
<https://doi.org/10.3389/fdata.2021.753409>

Working area

Area 6: Large Language Models-based assistants

If Other, please specify

Operations simplification and automation

CERN group/ Experiment

IT-TC and IT-SD

Project goals

The primary goal is to fundamentally simplify and automate computing operations, reducing the need for human intervention and increasing the efficiency and resilience of the WLCG. We aim to achieve this by developing and deploying a suite of AI Agents (AIA) that can understand, analyze, and act upon operational data in real time

Is this an already ongoing activity?

No

Indicative hardware resources needs

Available person power

0 now, 0.5 in a few months

Additional person power request

2.5

Timeline

Timeline and Resources We envision this as a multi-year project with a phased implementation. The initial phase (6-12 months) would focus on a proof-of-concept for the Automated Diagnostics for Data Transfers use case and Site Intelligence use-case for Anomaly Detection in Storage, requiring dedicated effort from a small team of AI/ML engineers and computing experts. Subsequent phases would scale the development to other use cases, with timelines dependent on resource allocation.

Authors: DI GIROLAMO, Alessandro (CERN); Dr ARSUAGA RIOS, Maria (CERN); PAPARRIGOPOULOS, Panos (CERN)

Presenters: DI GIROLAMO, Alessandro (CERN); Dr ARSUAGA RIOS, Maria (CERN); PAPARRIGOPOULOS, Panos (CERN)

Session Classification: Large Language Models-based assistants

Track Classification: 6. Large Language Models-based assistants