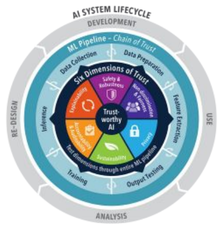


# Developing High-Performance Computing Agents for Resource Management

Ruchi Honale, Paul Brenner, Scott Hampton
The University of Notre Dame

CRC CENTER FOR RESEARCH COMPUTING
UNIVERSITY OF NOTRE DAME

## Introduction

High-Performance Computing (HPC) systems are vital for advanced research but managing them is complex and requires specialized knowledge. Automating routine HPC support tasks such as job scheduling, error analysis and script generation can drastically improve research productivity and system efficiency. This project develops Large Language Model (LLM)-powered command-line agents that analyze HPC job logs, generate job submission scripts, and assist users with context-aware responses.

## Objectives

1. Design LLM-powered command-line agents to streamline:
   a. HPC user support
   b. Job script generation
   c. HPC job error analysis
2. Evaluate the performance of multiple LLMs (Gemini, GPT-3.5 Turbo, GPT-4, GPT-4.1, GPT-4o)
3. Enable actionable, context-rich log summarization
4. Enhance system understanding for both users and administrators

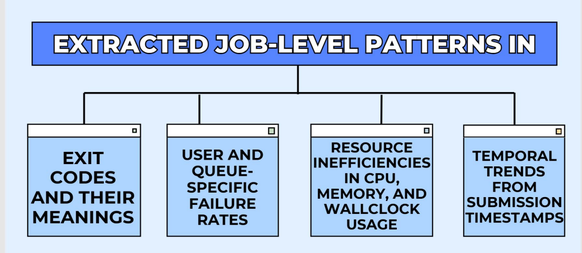


## Methods

### HPC User Support Agent

- Extracted 60 representative support issues from a corpus of 31,700 CRC user emails
- Augmented this dataset with CRC documentation and structured internal knowledge
- Used **Google Gemini** to develop agents capable of accurate, context-aware, and tonally appropriate responses

### Job Script Validation

- Submitted required program to Gemini
- Generated scheduler scripts using the support agent
- These scheduler scripts were successfully validated on CRC's clusters

### Error Analysis Agent

- Parsed **1.5 years** of Job Scheduler logs (~9 GB, 1.7 million entries)
- Isolated & categorized **8,002 failed jobs**
- Designed multi-level summarization prompts for **GPT-4**, **GPT-3.5 Turbo**, and **GPT-4o**

## Results

- Gemini effectively addressed both common and novel HPC support questions, and successfully generated runnable job scripts for a variety of job configurations on CRC systems.
- GPT-4o analyzed 10,000 log chunks, delivering the most comprehensive HPC error analysis in this project. It outperformed GPT-3.5 Turbo in both speed and reliability for large-scale summarization.
- GPT-4o demonstrated a strong grasp of HPC resource behavior—accurately interpreting differences between ru_maxrss and maxvmem, and highlighting nuanced inefficiencies in memory and CPU usage.
- Frequent failures were linked to specific exit codes. Summaries revealed repeated errors for specific users and queues, and flagged systemic over-provisioning where requested resources far exceeded actual usage.

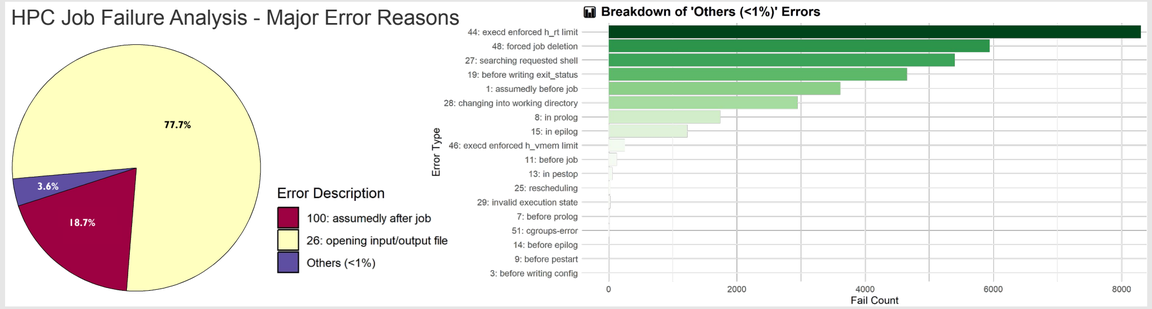


## Conclusion

This project demonstrates the feasibility and impact of integrating LLMs into HPC workflows. From auto-generating accurate job scripts to synthesizing large-scale job logs, the agents reduce manual intervention and reveal operational insights. Future work focuses on refining the resource manager module for proactive failure prediction.

## Discussion

LLMs show strong potential to augment HPC systems:

- Automating query response saves support team's time, ensuring they only face the toughest challenge
- Identifying trends in failed jobs improves operational planning
- Getting curated scripts makes it easier for researchers to deal with HPC giving more time for scientific inquiry

Limitations:

- Summarization is only as effective as the context provided in the prompts
- Real-time use demands careful prompt optimization and chunking strategies to minimize token usage

## Acknowledgements

HPC Team at ND CRC and my mentors Prof. Paul Brenner and Prof. Scott Hampton for their constant help and guidance. Thanks to University of Notre Dame - iSURE program, ND Global and MKSSS Cummins College of Engineering for Women, Pune for this wonderful opportunity.