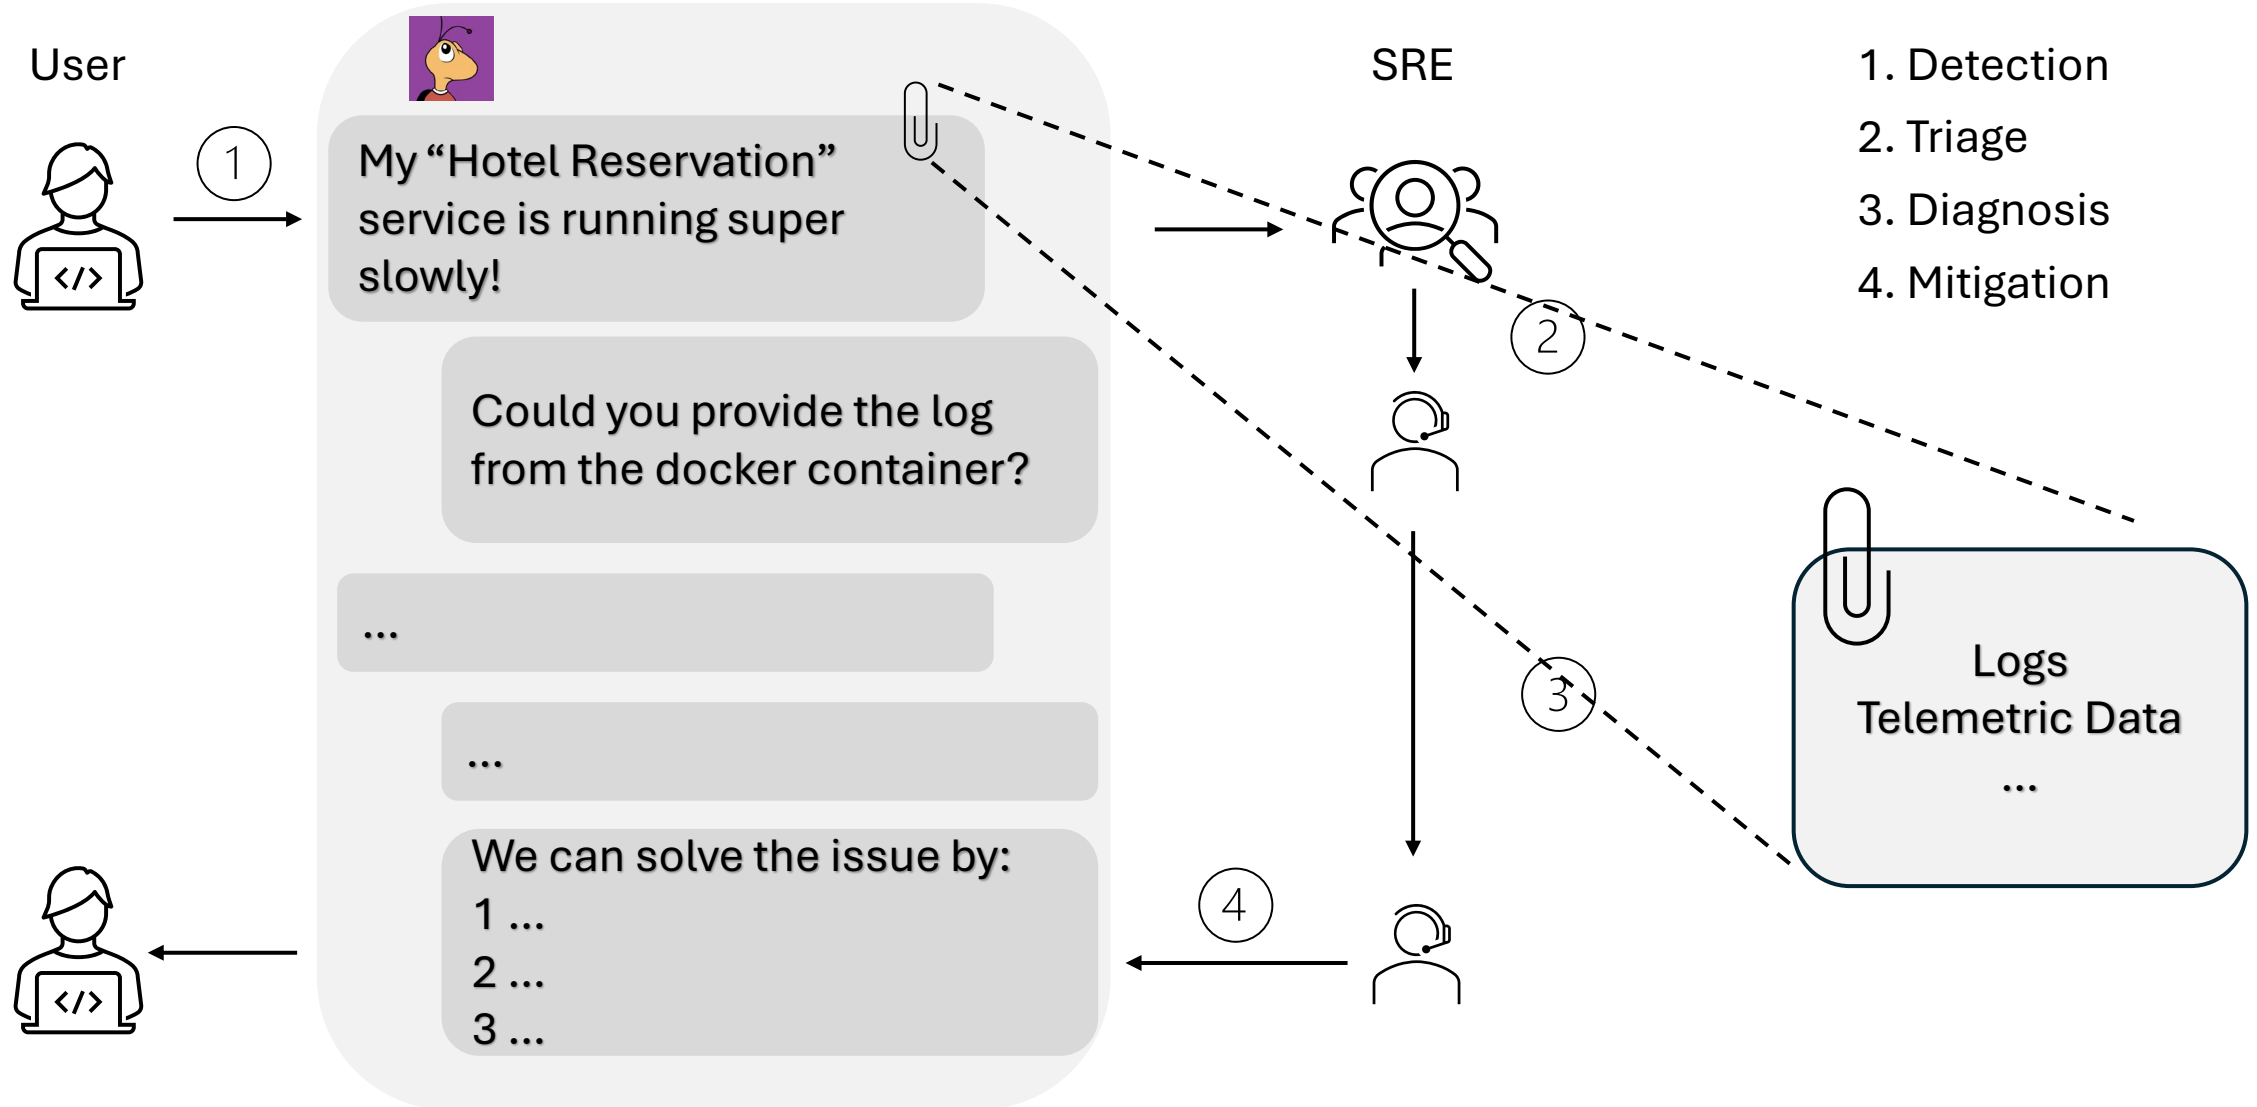



ARCA: AI for Root Cause Analysis

Yifan Wang, Kenneth P. Birman
Cornell University



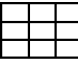
AI-Ops



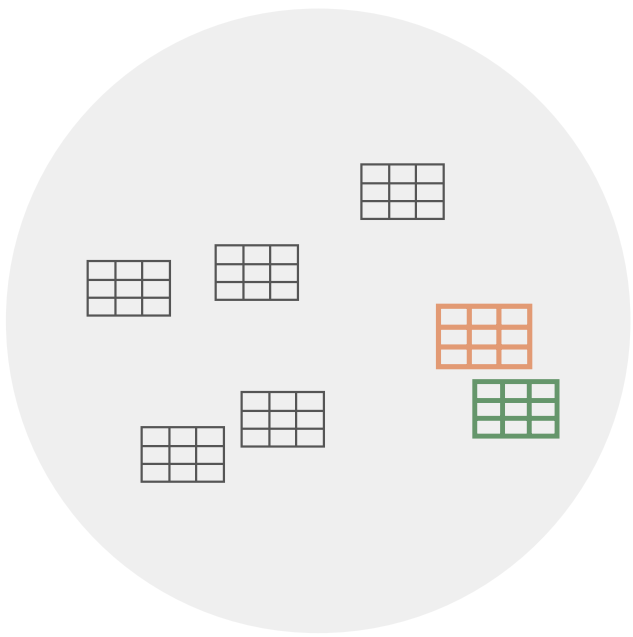





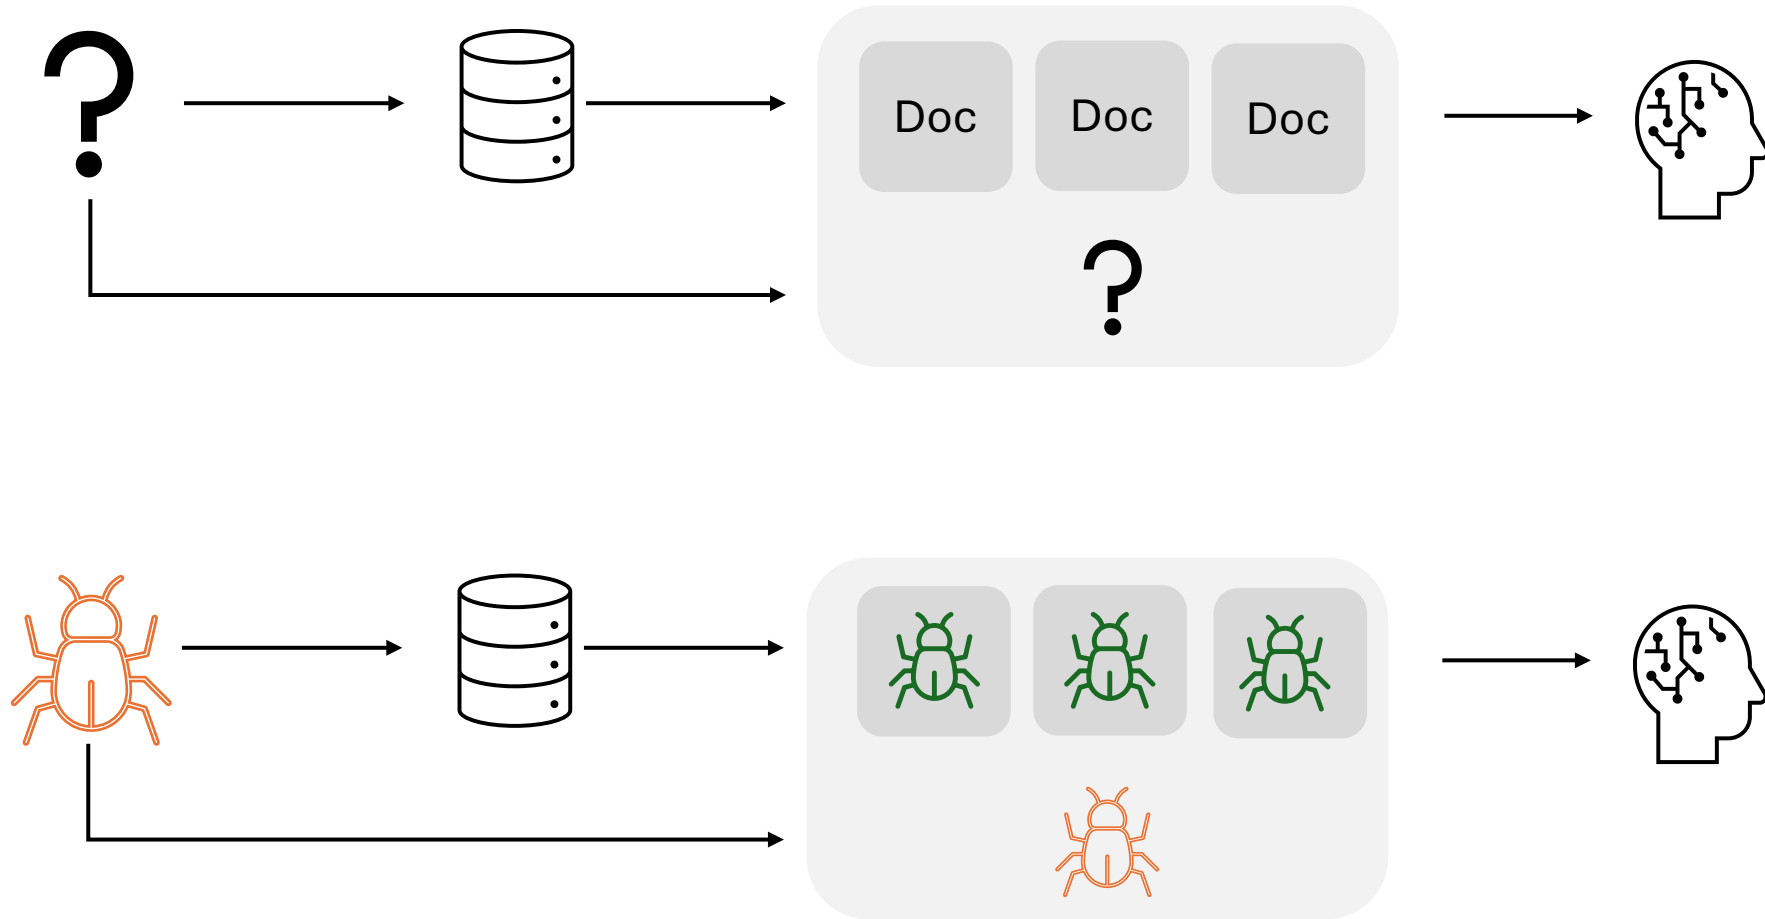
Similar bugs might share similar solution.

Bug: multi-modal data object

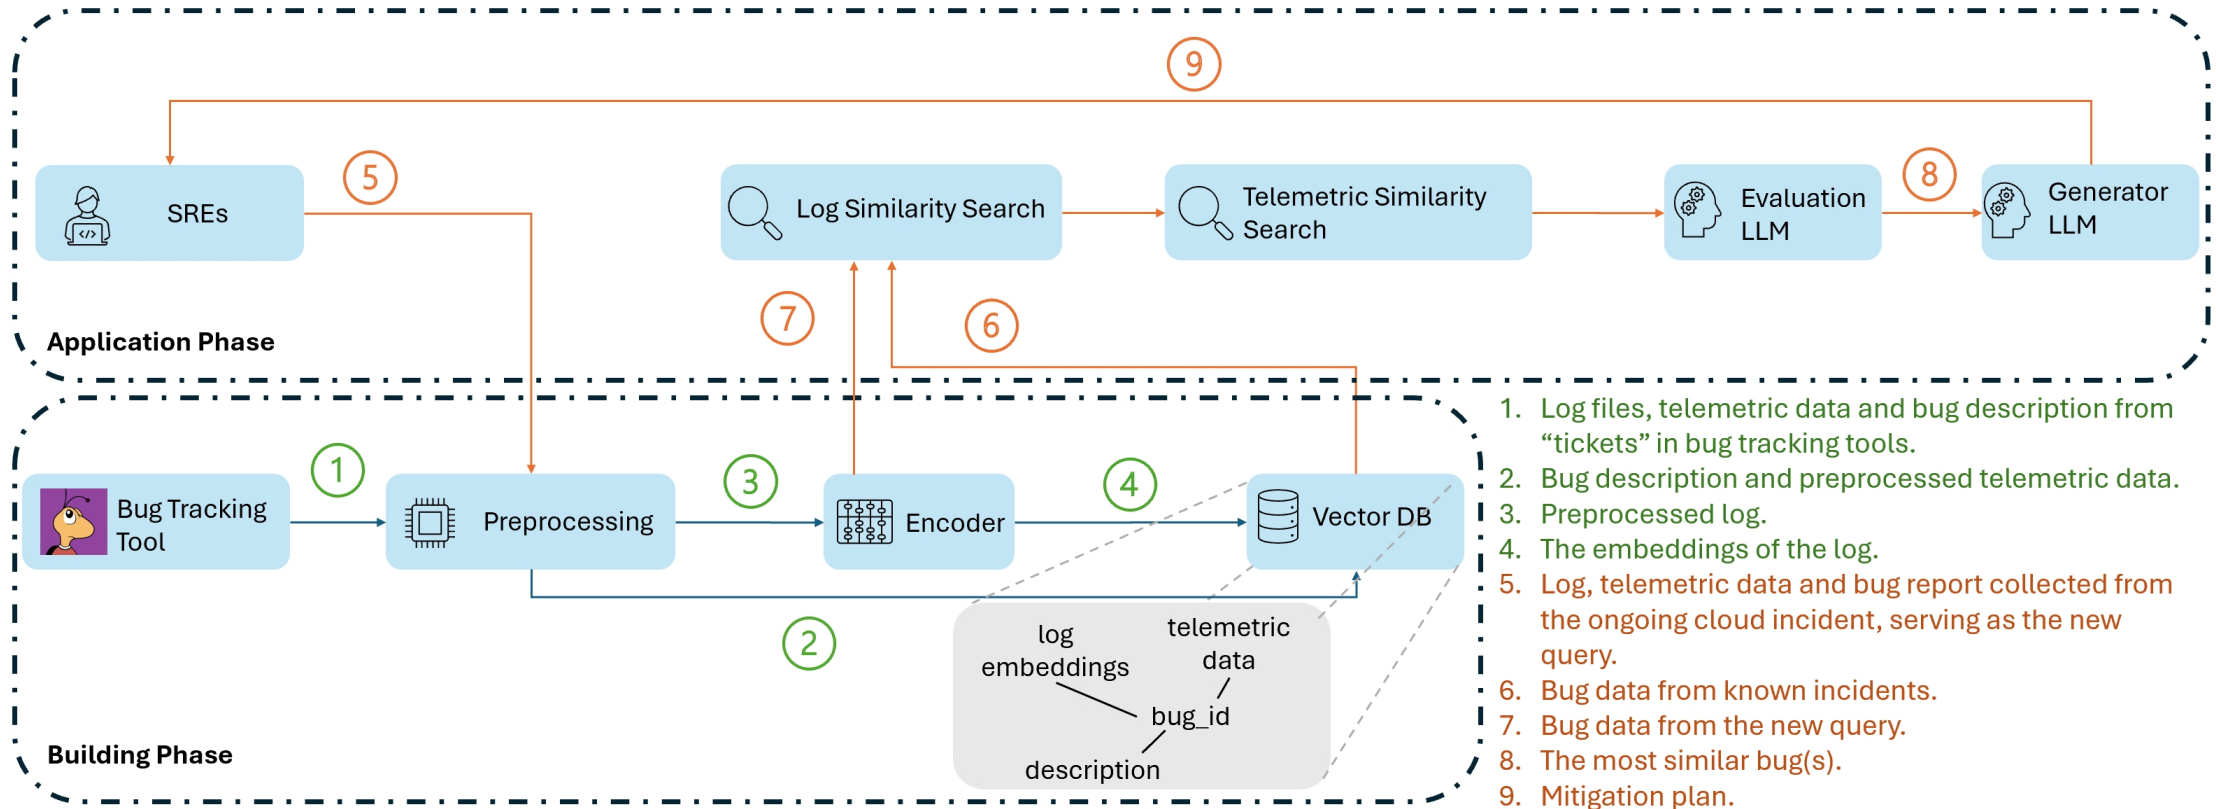
Components	Data Mode
Issue Description	Natural Language 
Log	Semi-structured language, numeric data, natural language, etc. 
Telemetric/Performance Data	Multi-variant time series numeric data 



Retrieval Augmented Generation (RAG)



ARCA Architecture

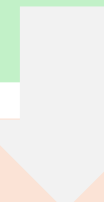


Log

Purging:
keeping error-related messages.



Embedding:
vectorizing the purged log file.



Searching:
ANN using cosine similarity search.

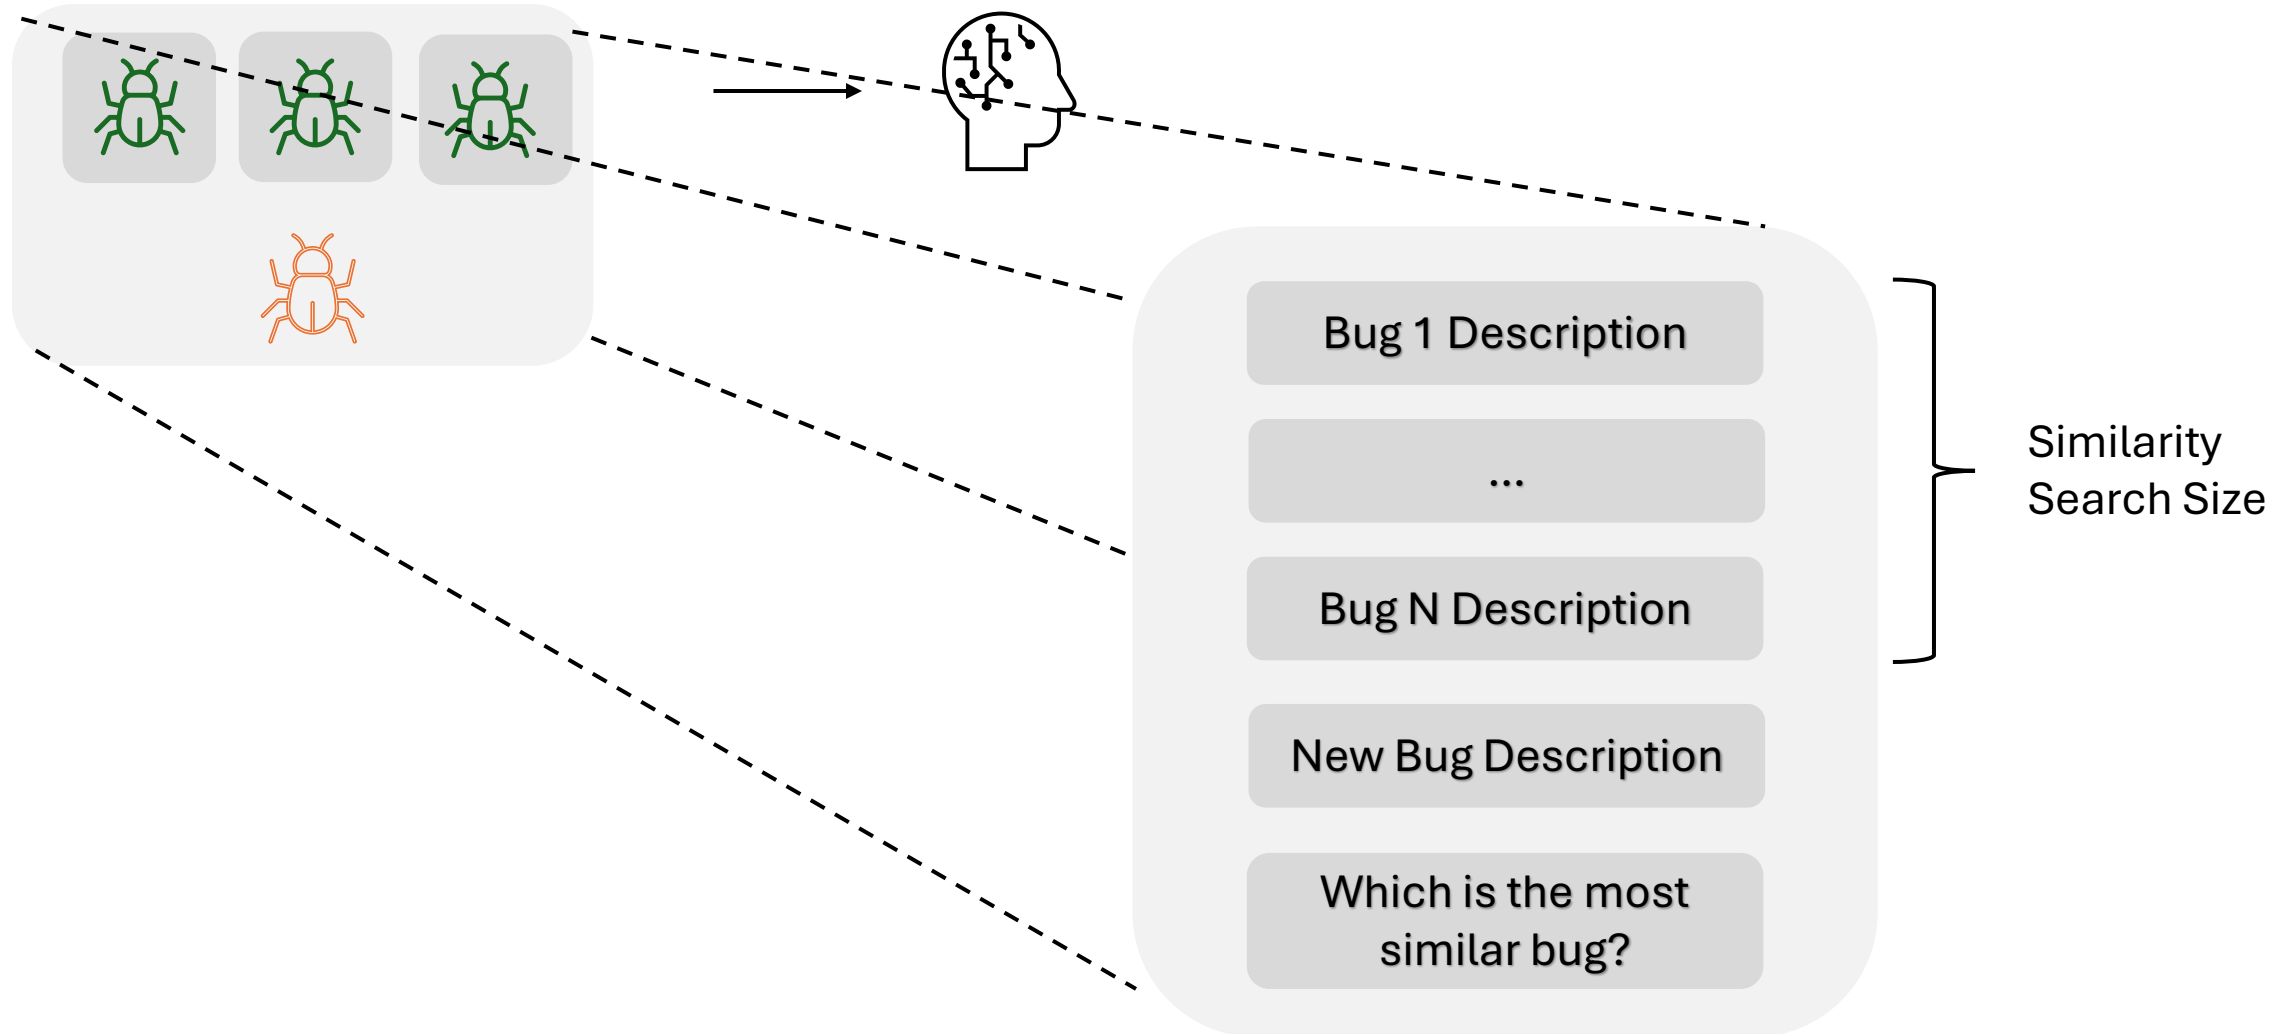
Telemetric Data

Feature Selection:
docker performance
counters



Embedding:
vectorizing using
statistics of each
feature.

Searching:
KNN using
cosine similarity.

Issue Descriptions



Findings



ARCA as a Log Clustering Tool

- ARCA's log processing module alone.
- Tested on super computing site logs.

Data Set	F1-Score	Recall	Precision
BGL	0.995/0.976	0.99/0.982	1/0.970
Thunderbird	0.984/0.97	0.975/0.99	1/0.97
Spirit	0.993/0.992	0.986/0.999	1/0.984
Liberty	0.986/*	0.986/*	0.986/*

More challenges



Different organizations of the knowledge base.



Coverage for more data modes.



Agentic design.
