From Chaos to Clarity: Log-based Kernel Panic Root Cause Analysis for Large-Scale Cloud Services

Tianyu Cui, Yang Zhang, Shenglin Zhang, Xin Wu, Yicheng Sui, Liangyan Peng, Yuhe Ji, Feng Wang,Changchang Liu, Zeyu Che, Xiaozhou Liu, Yongqian Sun, Yu Zhang

Frontiers of Computer Science, DOI: 10.1007/s11704-025-50788-0

Problems & Ideas

- Problems of conventional Log-based RCA matching approaches:
  - Kernel panic logs are massive and only a tiny portion indicates real faults.
  - Strong noise overwhelms critical panic signals.
  - Long-range interdependencies exist across logs, making RCA difficult.
- Ideas: A two-stage RCA framework for kernel panic named LogSage.

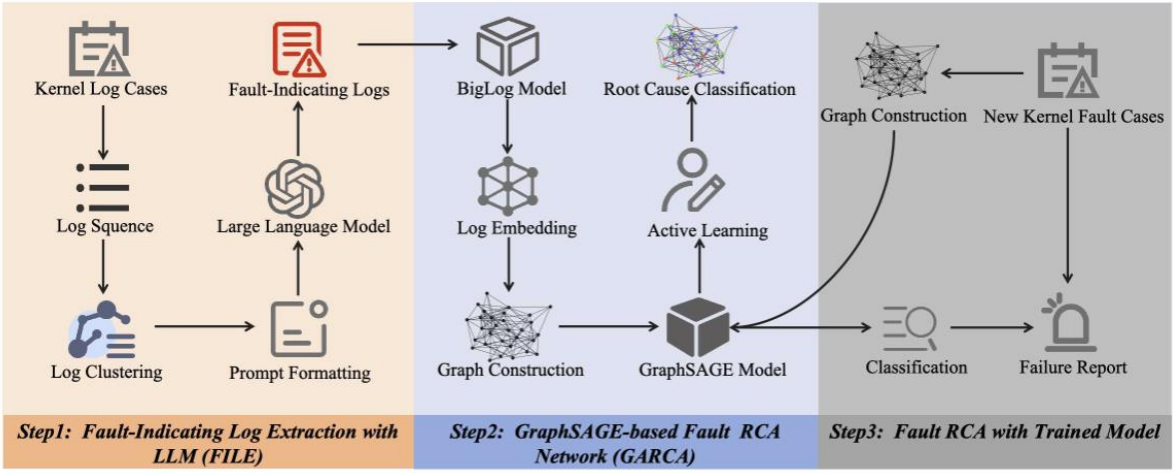

The framework consists of three steps: (1)Fault-indicating Log Extraction (FILE) using unsupervised clustering and LLM summarization, (2) GraphSage-based Fault RCA (GARCA) via GraphSAGE and active learning, and (3) Fault RCA with Trained Model.

Main Contributions

- Contributions:
  - An LLM-enhanced mechanism for extracting fault-indicating kernel panic logs, combining clustering and structured summarization to significantly reduce noise and improve interpretability.
  - A graph-based RCA framework augmented with active learning, leveraging BigLog embeddings and GraphSAGE to model long-range inter-log dependencies with high label efficiency.
  - Extensive experiments on three real-world datasets, and successfully deployed in ByteDance's cloud infrastructure, assisting engineers in real-world kernel panic diagnosis at scale.

Table VI: RCA performance comparison across three datasets

Method	Dataset 1				Dataset 2				Dataset 3			
	P	F	R	T	P	F	R	T	P	F	R	T
LogCluster	42.6	38.3	36.9	0.03	40.1	35.9	34.5	**0.01**	58.7	54.4	53.1	**0.01**
LogRule	65.9	61.5	60.3	**0.02**	63.2	59.8	58.1	0.02	64.0	59.2	58.6	0.02
LogKG	81.3	76.7	75.4	2.40	79.9	75.0	74.3	2.45	80.2	76.2	75.0	2.42
LogPrompt	74.7	70.1	68.9	4.26	73.5	69.3	67.6	4.24	75.6	71.0	69.2	4.25
LogSage	**92.4**	**92.2**	**91.8**	3.21	**94.1**	**95.3**	**95.9**	3.08	**96.7**	**96.3**	**95.9**	3.10

F1 scores: 92.2%, 95.3%, 96.3% across three real-world datasets.Up to 20% improvement over strongest baselines. Deployed in ByteDance's cloud infrastructure for 6+ months.