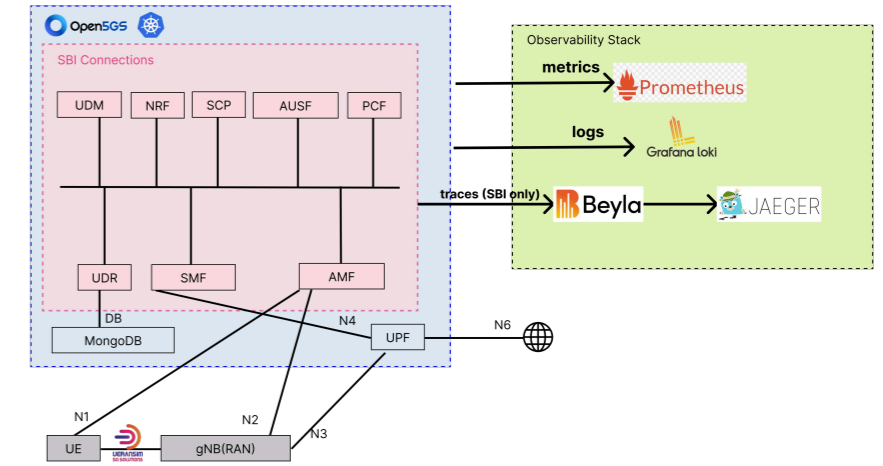


Introduction and Background

- Problem:** Fault diagnosis in cloud-native 5G Cores usually relies on a single telemetry modality (either metrics, logs, or traces). We ask what cross-modal correlations reveal across fault signals.
- Background:** 5G Core Network Functions (NFs) run as interdependent microservices, so faults propagate across modalities. Multi-modal telemetry is known to help microservice diagnosis (DiagMLP, 2025), but to our knowledge no study characterizes what cross-modal correlations reveal across fault types in a 5G Core and no public dataset captures all three modalities under controlled 5G-Core fault injection, so we built one.

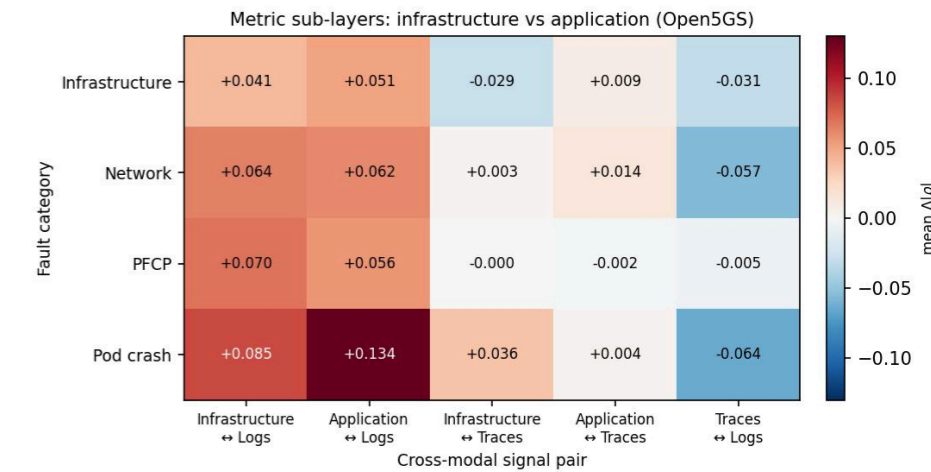
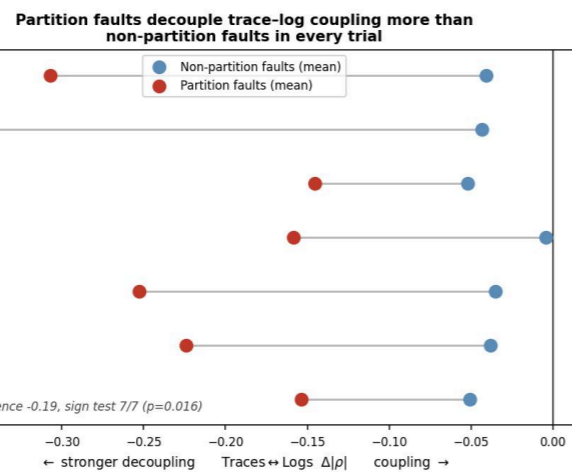
Methodology

- Testbed:** Open5GS 5G Core on Kubernetes · 10 components (9 NFs + MongoDB) · 22 fault scenarios · 7 independent trials
- Built pipeline:** bin all 3 modalities into 10-s windows → z-score vs pre-fault baseline → Spearman correlation per pair → $\Delta|\rho| = |\rho_{\text{during}}| - |\rho_{\text{pre}}|$
- Two feature views:** coupling $\Delta|\rho|$ (characterization) + raw per-signal stats (classification baseline)
- Faults:** 22 operational (infrastructure / network / PCF / pod-crash) + 6 security attacks; 7 operational trials, 3 security trials



Key Findings

- Interface partitions collapse trace↔log coupling, in all 7 trials
- Visible only in the relationship; no single modality shows it
- Each fault category has a distinct cross-modal signature
- Pod crash couples to logs via application KPIs (+0.134), a procedure failure, not just resource stress
- Metrics↔traces are close to 0 everywhere → the trace blindspot



Limitations

- Security results are preliminary (3 trials, one hardware platform used for data collection, deterministic injection)
- Does not generalise to unseen fault types (leave-one-fault-out ≈ 23%, below baseline)
- Traces limited to the SBI (eBPF, per-NF latency); single small-scale testbed

Conclusion

- Blindspots are fault-specific & structural; traces miss non-SBI faults; logs miss valid-request attacks (NRF SBI flood)
- Cross-modal coupling reveals relational signatures invisible to single signals (partition decoupling)
- Metrics dominate classification on the provided dataset → multi-modal's value is coverage & characterization (richer incident context), not accuracy

Research Question

How can multi-modal correlation of observability signals (logs, metrics, and traces) characterize faults and expose where single-modality analysis is blind?

SQ1: How can a method be developed to identify and align cross-modal correlations across logs, metrics, and traces within an existing observability dataset?

SQ2: How do time alignment strategies resolve sampling rates mismatches between continuous and event-driven signals?

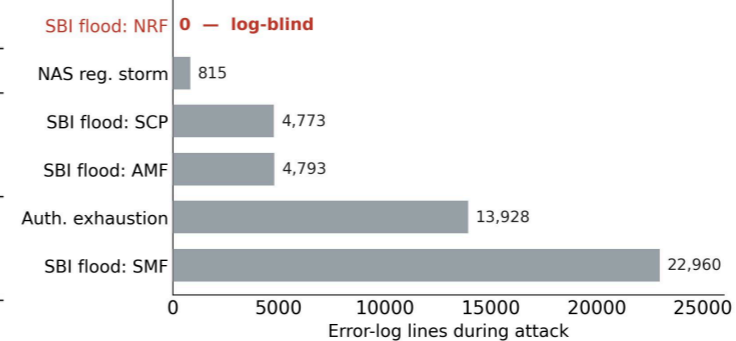
SQ3: What temporal or contextual patterns emerge across the combined signals during different fault types?

SQ4: How does the correlated multi-modal analysis compare to isolated single-modality analysis in fault discriminability?

Modality subset	#feat	LOOCV		LOTO	
		Acc.	F1	Acc.	F1
Majority baseline	—	31.8	0.12	31.8	0.12
Metrics	248	83.4	0.85	84.3	0.86
Logs	60	57.5	0.57	60.7	0.60
Traces	71	72.7	0.72	74.7	0.73
Metrics + Logs	308	82.9	0.84	83.5	0.85
Metrics + Traces	319	85.4	0.87	86.2	0.87
Logs + Traces	131	72.1	0.71	73.6	0.71
All modalities	379	85.3	0.86	85.5	0.86

- Random Forest on raw per-signal features; leave-one-out (LOOCV) and leave-one-trial-out (LOTO) accuracy.
- Metrics alone 84% ≈ best modality pair result 86%, thus fusion barely helps

Security blindspot: the NRF flood emits zero error logs



- Valid-request attack: the NRF answers the flood normally → zero error logs (every other attack logs thousands)
- A log-based detector is blind to it; visible only via metrics (NRF resource exhaustion) + traces (SBI span skyrockets)