PhD Research Progress 2026

Candidate: Haitham Hussein Jodha AL SATAI

Topic: From Telemetry Allocation to Multi-Level Adaptive Tracing

Supervisor: Prof. Foutse Khomh, Ph.D., and Prof. Naser Ezzati-Jivan, Ph.D.

Department: Computer Engineering

Thursday, July 9, 2026

The Observability Dilemma

Rich telemetry helps diagnosis. Full telemetry is expensive.

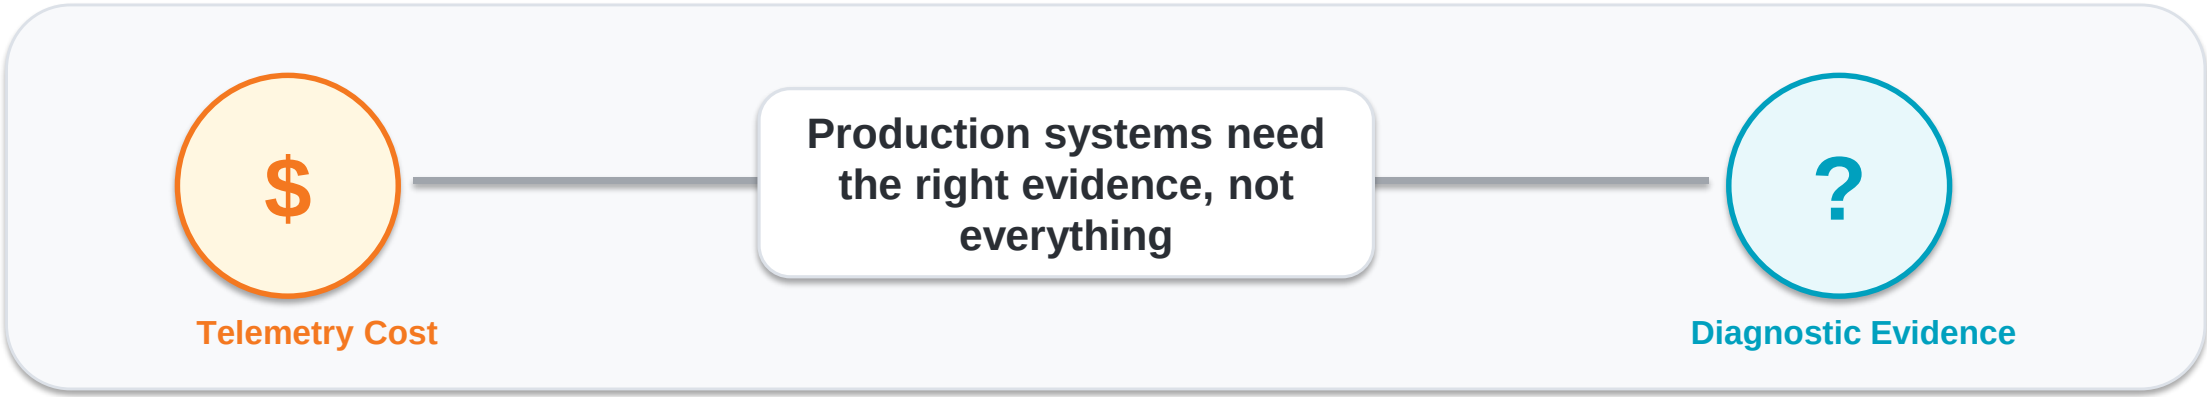

Too little telemetry

Evidence disappears before diagnosis

Too much telemetry

Storage, compute, network, and operational cost

Research challenge

Escalate only when suspicion justifies cost

Research question: When should a system escalate telemetry, and how far?

First Step: Allocating Telemetry Under Budget

Before studying when to escalate, I studied what telemetry to keep under cost constraints.

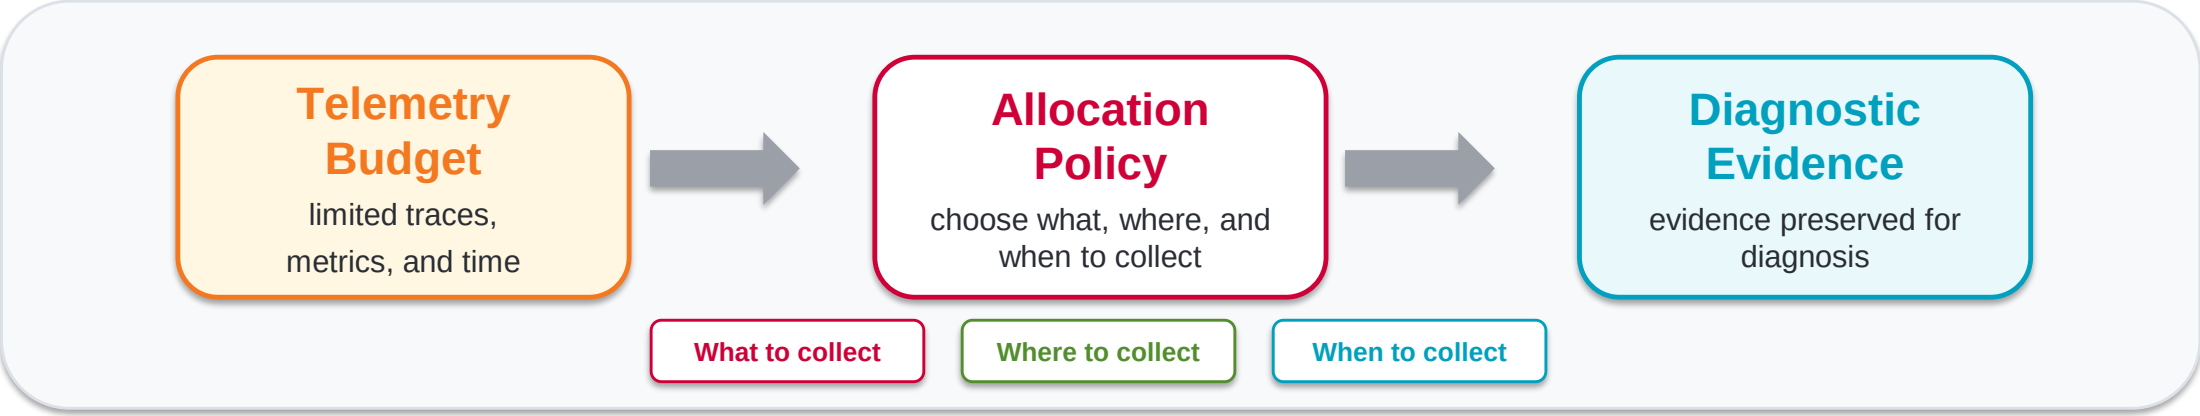

Starting point

Full tracing is costly; naive sampling can discard useful abnormal traces.

Research question

With a limited telemetry budget, what evidence should be kept for diagnosis?

First direction

Diagnosis-aware allocation across traces, metrics, services, and time.

Key idea: observability is an allocation problem, not just a fixed sampling rate.

First Paper Attempt: Constrained Telemetry Allocation

The first empirical stage became a submission and clarified the next research question.

Paper 1

A Few Good Traces

Constrained telemetry allocation for microservice diagnosis

What I studied

Reduced trace budgets
Trace/metric allocation
Post-alert diagnosis

Outcome

Submitted to ASE 2026
Not accepted
Useful empirical foundation

What it taught me

Clarified limitations
Sharpened the next question
Moved toward escalation

Lesson: deciding what to keep is important, but the next question is when telemetry should increase before evidence is lost.

Shift in Research Question

From “what telemetry should we keep?” to “when should telemetry increase?”

First stage

What should we keep?

Post-alert setting

Limited telemetry budget

Allocation across traces, metrics, services, and time

Goal: preserve diagnostic evidence

earlier in the pipeline

Next stage

When should telemetry increase?

Minimal always-on signals

Detect suspiciousness early

Escalate before evidence is lost

Goal: collect richer evidence only when justified

The work moved from post-alert allocation to early suspiciousness-triggered escalation.

Current Direction: Multi-Level Adaptive Tracing

From a binary tracing switch to gradual telemetry escalation.

Binary view: minimal telemetry ➡ full tracing | Multi-level view: escalate step by step as suspicion and cost justify it

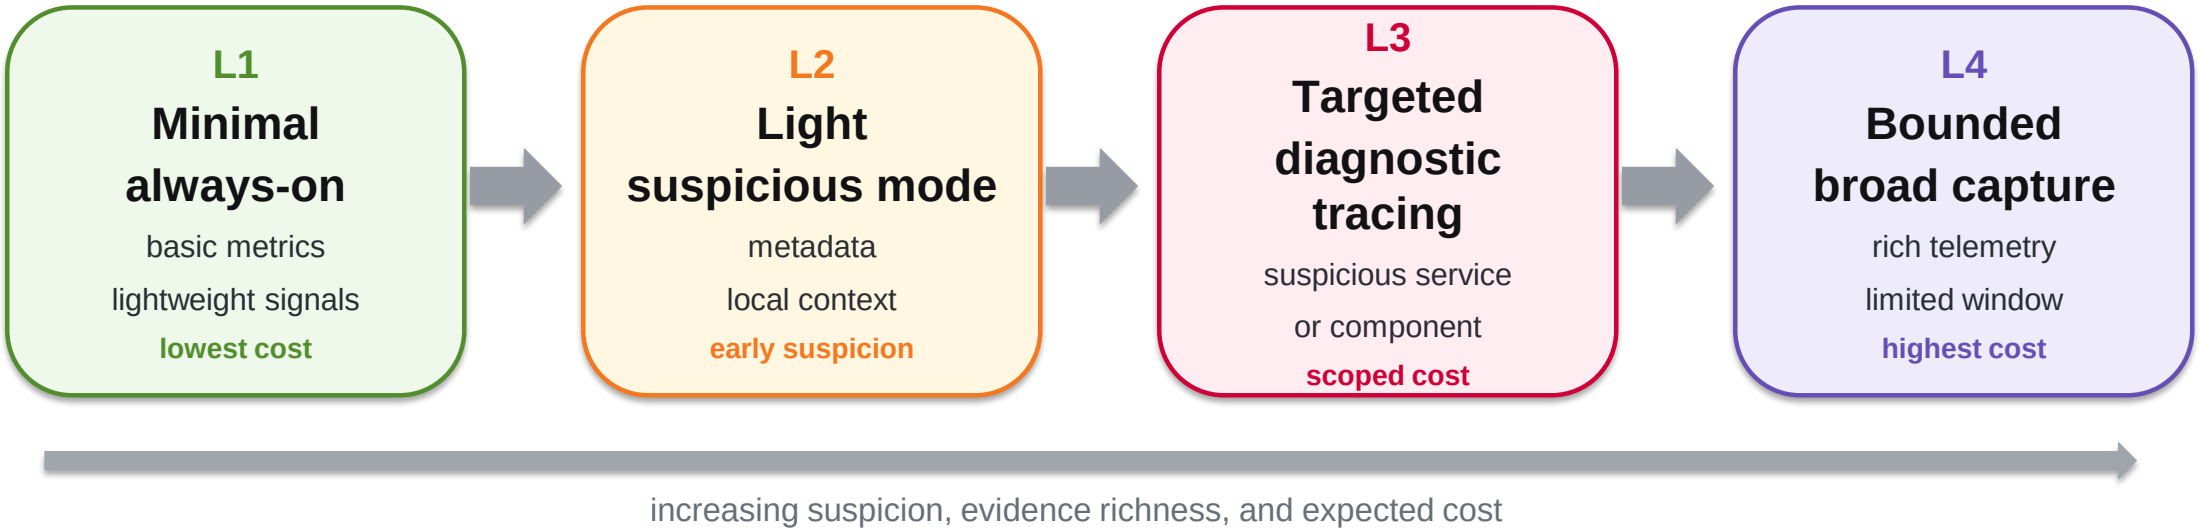

Key idea: the question is not only whether to trace, but how far to escalate.

Current ICSE Submission

A level-aware study of adaptive tracing under suspicion-cost tradeoffs.

ICSE 2027

Escalate, But How Far?

Multi-Level Adaptive Tracing Under Suspicion-Cost Tradeoffs

Core claim

Binary adaptive tracing hides important tradeoffs.

Primary evidence

GAIA + AIOps-Dataset with supporting settings

Contribution

A level-aware framework and evaluation approach

Positioning

Tradeoffs, not a claim that one policy always wins

GAIA: distributed microservices | **AIOps:** service-level multimodal telemetry | **Supporting:** host KPI, AgentTelemetry, modern methods

The paper's goal is not to show that multi-level tracing always wins, but to expose the tradeoffs that binary tracing hides.

Main Findings: Tradeoffs, Not Universal Wins

Multi-level tracing changes where cost and evidence appear.

GAIA

Favorable low-level warning tradeoff

Earlier Level >=2 suspicious evidence

Less high-cost Level 4 capture

But lower early Level >=4 coverage

Aiops-Dataset

Stronger cost-control tradeoff

Much less broad multimodal capture

Lower clean weighted cost

But early evidence drops compared with binary

More early broad capture <-> **Lower telemetry cost**

Multi-level tracing is not a magic win; it exposes an effectiveness-cost frontier.

MLAT-Bench Artifact

A reusable benchmark core for evaluating telemetry escalation policies.

Paper evaluation problem -> Reusable benchmark core -> Policy evaluation and regeneration

Purpose

Make multi-level adaptive tracing evaluation reusable and reproducible.

What it evaluates

Binary, multi-level, and modality-aware escalation policies.

Metrics

Level >=2, >=3, >=4 outcomes plus clean-period cost.

Scope

Built over GAIA and Aiops-derived streams; no raw telemetry redistribution.

Includes code, schemas, split rules, policy interfaces, tests, validation summaries, and regeneration documentation.

MLAT-Bench turns the paper's evaluation problem into a reusable artifact.

Current Status and Next Steps

Where the work stands after my first six months.

Current status

- Started PhD in January 2026
- ASE paper submitted; not accepted
- RV paper submitted
- ICSE 2027 paper submitted
- MLAT-Bench artifact prepared

Next steps

- Learn from review feedback
- Strengthen benchmark and reproducibility
- Refine multi-level adaptive tracing policies
- Extend evaluation to broader settings
- Build the next PhD study

The first six months produced a clear research direction, concrete submissions, and a reusable artifact.

Takeaway

Adaptive observability is a control problem under telemetry cost.

My PhD direction

Preserving diagnostic evidence under telemetry cost through adaptive, level-aware observability.

Allocation

What evidence should we keep?

Escalation

When should telemetry increase?

Control

How far should the system escalate?

The key question is not only whether to trace, but how far to escalate.

Questions?

Thank you for your attention.