

J.A.R.V.I.S.

Just Actionable Real-Time Visibility Intelligent System



TEAM GITPUSH -F


Nguyen Son Tung
Cu Duc Quang
Nguyen Minh Hung
Mai Dang Khoa
Nguyen Huynh Trieu Thinh

ACADEMIC SUPERVISOR

Mr. Minh Vu Thanh

INDUSTRY SUPERVISOR

Mr. Cuong Ngo
Mr. Dat Pham
Mr. Do An



BACKGROUND + MOTIVATION

When a banking system fails, the knowledge needed to fix it is scattered, including runbooks, past incident tickets, and the heads of senior engineers. Operators lose critical minutes searching instead of acting, and the diagnosis changes depending on who is on shift.

LLMs could close that gap. But today's strongest reasoning models still hallucinate on $\geq 10\%$ of grounded tasks. In banking, a confident wrong answer is worse than no answer. J.A.R.V.I.S., built with HCLTech, starts from that constraint: **Don't trust the model & Verify it with code.**

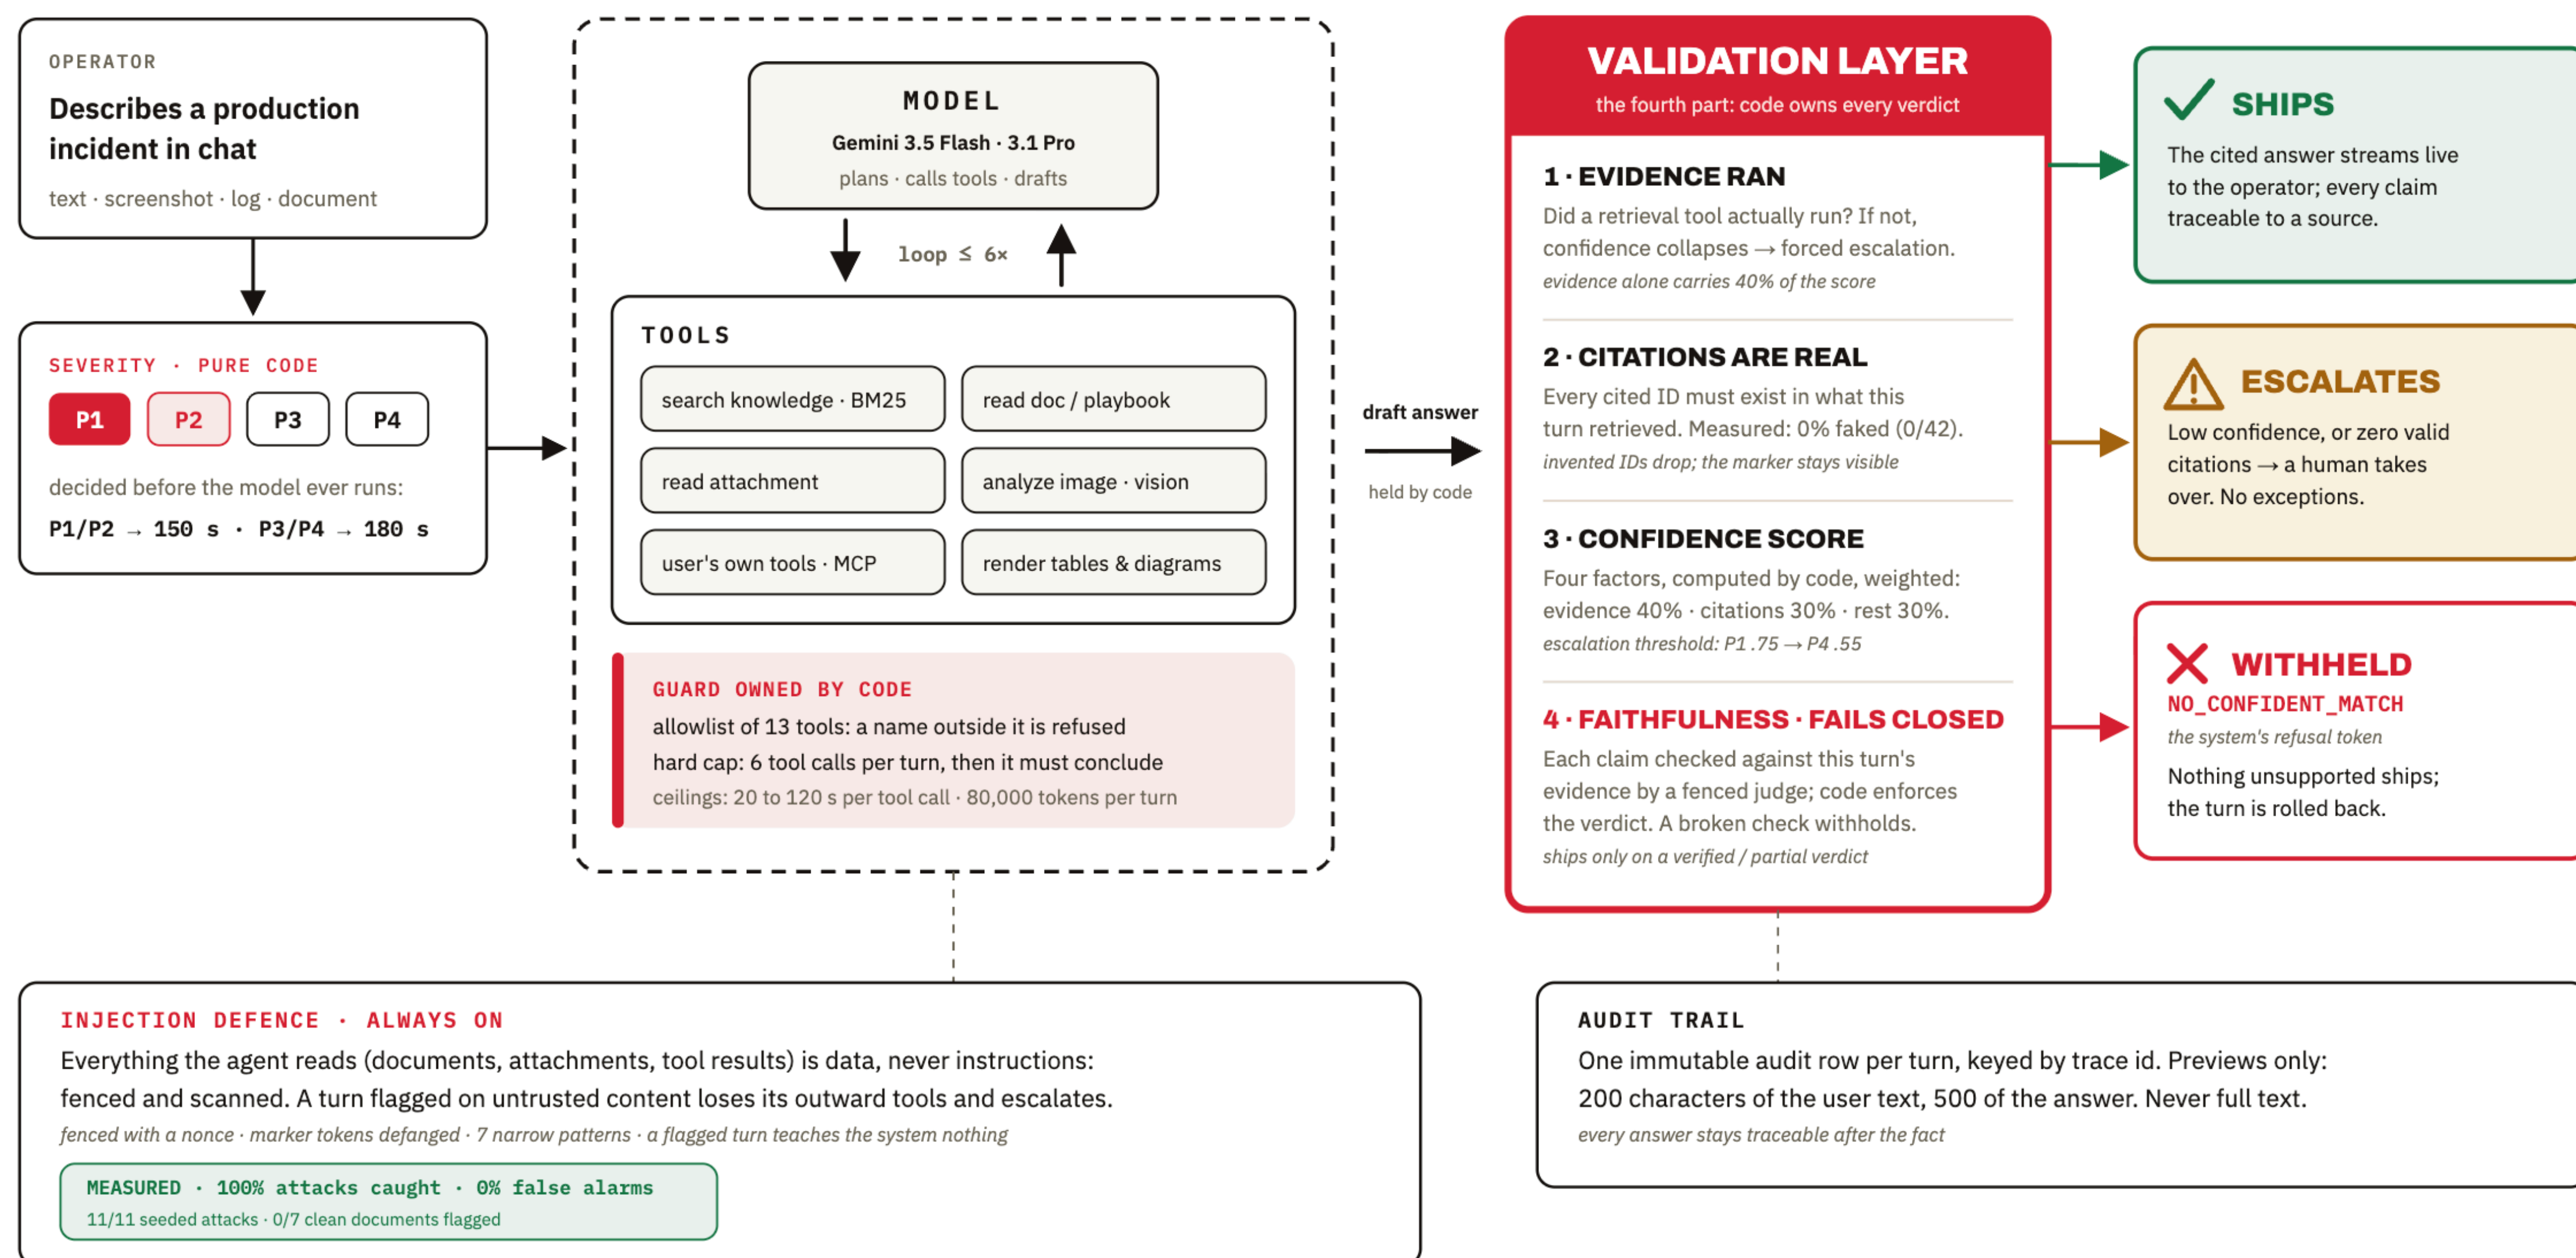
OBJECTIVES

Build an incident-triage assistant a bank could actually trust:

- Cited answers, or none
- Agentic, not one-shot retrieval
- Safety lives in code, never in the prompt
- Real-time by default
- Decision support, not automation
- Escalates to a human when unsure
- Injection caught, outward tools revoked

HARNESS SYSTEM + VALIDATION LAYER

A standard agent harness has three parts: orchestrator, tools, model.
J.A.R.V.I.S adds a fourth part the model cannot bypass, enforced entirely by code.




01 · The Harness Loop

A harness is the code around the model. It classifies severity before the model ever runs (from P1 to P4 sets the clock), then lets the model plan and call tools: knowledge search, attachments, and the user's systems. An allowlist owned by code and a hard cap of six calls per turn keep the loop honest.

02 · The Validation Layer

A standard harness stops at three parts: orchestrator, tools, and model. J.A.R.V.I.S. adds a fourth that code owns end-to-end. Four checks stand between the draft and the operator. Fall short on evidence, citations, or confidence, and the turn escalates to a human; if the faithfulness check cannot answer, the text is withheld outright. Injection attempts are caught and contained (100% in testing), and every turn writes one immutable audit row.

SYSTEM HIGH-LEVEL DESIGN



One agent, one platform, one cloud. The Agent owns reasoning, RAG, reasoning via Gemini and tool calls to Jira via MCP; the Platform owns Auth and Documents. They talk over gRPC. Frontend runs on Vercel; Supabase, GCS, and Redis handle state.

EVALUATION METRICS

0.00%
FABRICATED CITATIONS

Zero fabricated citations across 42 emitted IDs.

100%
INJECTION ATTACKS CAUGHT

All 11 injection attempts were caught, with no false positives on 7 clean documents.

100%
UNANSWERABLE QUESTIONS REFUSED

All 4 unanswerable questions were escalated to a human.

15MS
RETRIEVAL P95 TARGET 5S

Measured across 460 queries, 326 times faster than the 5s target.

Every number above is measured from a code-owned check that runs after the model. Citations, injections, refusals, and latency are all verified by code, and not asserted by the LLM. In a bank, a confident wrong answer is worse than no answer, and the harness makes sure that never ships.