AIWILD WORKSHOP · ICML 2026

# Compound AI System Reliability: A Failure Taxonomy and Resilience Pattern Catalog from 150 Production Incidents

Rudrendu Kumar Paul, Boston University

Sourav Nandy, University of Texas at Austin

MOTIVATION

# Component-level monitoring fails compound AI systems: silent degradation for 11 days while all health checks passed

A RAG pipeline ran for 11 days before anyone noticed F1 had dropped from 0.82 to 0.59. Every per-component health check passed throughout. The failure lived at the boundary between retriever and generator: an index rebuild changed embedding normalization, shifting similarity distributions and flooding the generator with irrelevant documents. The generator produced confidently wrong answers.

**Per-component observability misses inter-component semantic failures.** Latency and error-rate dashboards show green while system-level quality degrades silently for days. This is the monitoring gap that existing tooling leaves open.

> **The core problem:** Compound AI systems fail at boundaries between components, not within them. Component-level health checks are necessary but not sufficient for reliable deployment.

METHODOLOGY

# 150 production incidents → 23 failure modes across 5 categories → validated resilience pattern catalog

ICML International Conference On Machine Learning

## DATA AND TAXONOMY

- **150 production incidents** from open-source compound AI projects and anonymized enterprise deployments
- **23 failure modes** organized across 5 categories: Retrieval Failures, Generation Failures, Tool Failures, Orchestration Failures, Integration Failures
- Failure mode classification performed via grounded coding of incident reports

## EXPERIMENTAL VALIDATION

- Fault injection experiments on a **6-component compound AI system** measuring pattern effectiveness under controlled degradation
- Patterns evaluated across matched incident cohorts with baseline comparisons against unstructured monitoring deployments
- Output: practitioner-ready catalog mapping each failure mode to resilience patterns with measured impact

**150**
production incidents analyzed

**23**
failure modes identified

**5**
failure categories

KEY FINDINGS

# Five resilience patterns with measured effectiveness: circuit breakers 89%, semantic validators 81%, quality gates 73%

ICML International Conference On Machine Learning

## RESILIENCE PATTERN EFFECTIVENESS

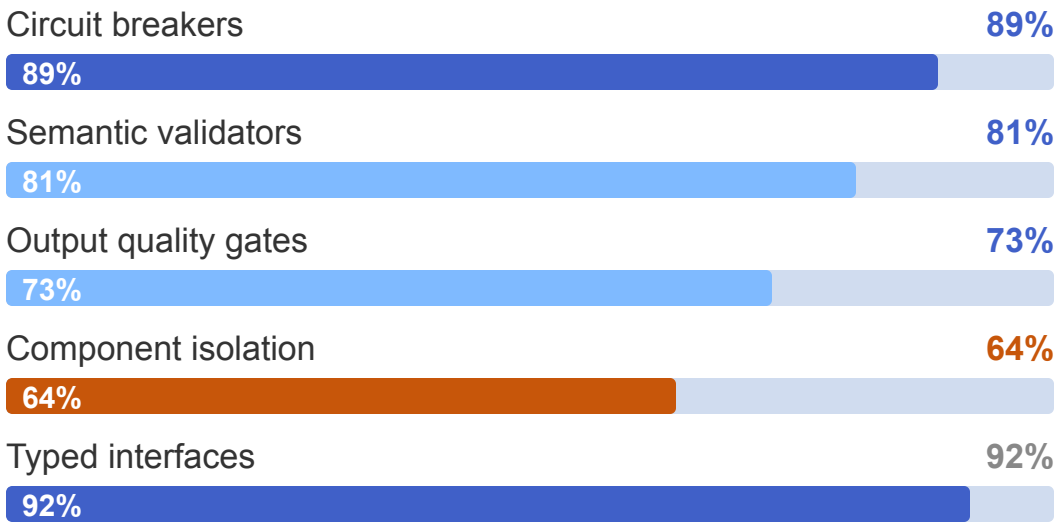


Values = cascade propagation or silent degradation reduction vs. baseline deployments without each pattern.

## THE MONITORING GAP

Existing observability tools track latency and error rates per component. **Inter-component semantic failures** — responsible for the largest share of silent production degradation in our dataset — go undetected by all per-component checks.

## PATTERN-TO-FAILURE-MODE MAPPING

- Circuit breakers address **Orchestration Failures** and cascade propagation
- Semantic validators catch **Integration Failures** at component boundaries
- Output quality gates detect **Generation Failures** that produce confidently wrong answers
- Typed interfaces reduce **Tool Failures** at API boundaries

RESULTS

# 3 or more patterns combined: 71% MTTR reduction vs. unstructured monitoring baselines

ICML International Conference On Machine Learning

## MTTR REDUCTION BY PATTERN COUNT

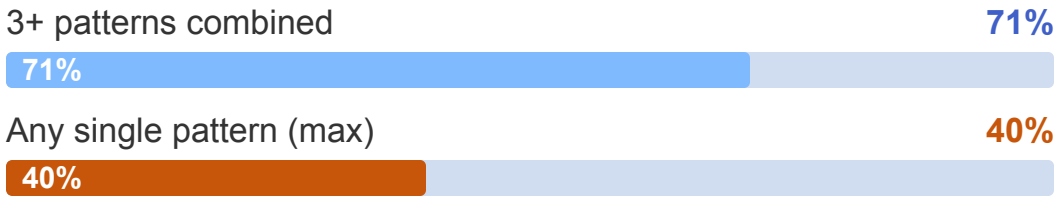


**Combination matters:** Each resilience pattern in isolation achieves at most 40% MTTR reduction. Gains compound when patterns are applied together, reaching 71% with 3 or more patterns deployed.

## REQUIREMENTS TO REACH THE 71% THRESHOLD

- **Intra-component patterns** (quality gates, circuit breakers) must be combined with inter-component patterns (semantic validators, typed interfaces)
- **Baseline comparison:** systems evaluated against unstructured monitoring deployments using matched incident cohorts
- **Coverage required:** patterns must address failure modes across at least 3 of the 5 failure categories to achieve the 71% threshold

**71%**
MTTR reduction (3+ patterns)

**40%**
Max reduction (1 pattern alone)

CONTRIBUTIONS & TAKEAWAY

# The gap between component correctness and system correctness is where reliability breaks down

ICML International Conference On Machine Learning

## CONTRIBUTIONS

- Failure taxonomy: **23 failure modes across 5 categories** derived from 150 production incidents in compound AI systems
- Resilience pattern catalog: **5 validated patterns** with measured effectiveness, mapped to specific failure modes
- Empirical evidence that **pattern combination is required** — single-pattern deployments cap at 40% MTTR reduction; 3+ patterns reach 71%
- Released as open practitioner resource: taxonomy, catalog, and fault injection experimental results
- Demonstrates that **per-component health checks are insufficient** for reliable compound AI deployment

## KEY FINDING

Compound AI systems fail at boundaries between components, not within them. Typed interfaces achieve **92%** boundary failure reduction; circuit breakers reduce cascade propagation by **89%**; semantic validators catch **81%** of silent inter-component degradation.

### TAKEAWAY

The gap between component correctness and system correctness is where reliability breaks down. Closing it demands inter-component observability; adding more granular per-component dashboards leaves the gap open.