Beyond Observability: Domain-Aware, Agentic Operations

Proactive monitoring, intelligent decisioning, and controlled action for modern engineering teams.

REFERENCE ARCHITECTURE

SRE & PLATFORM ENGINEERING

AgentiAILLC

contact@azentiq.ai

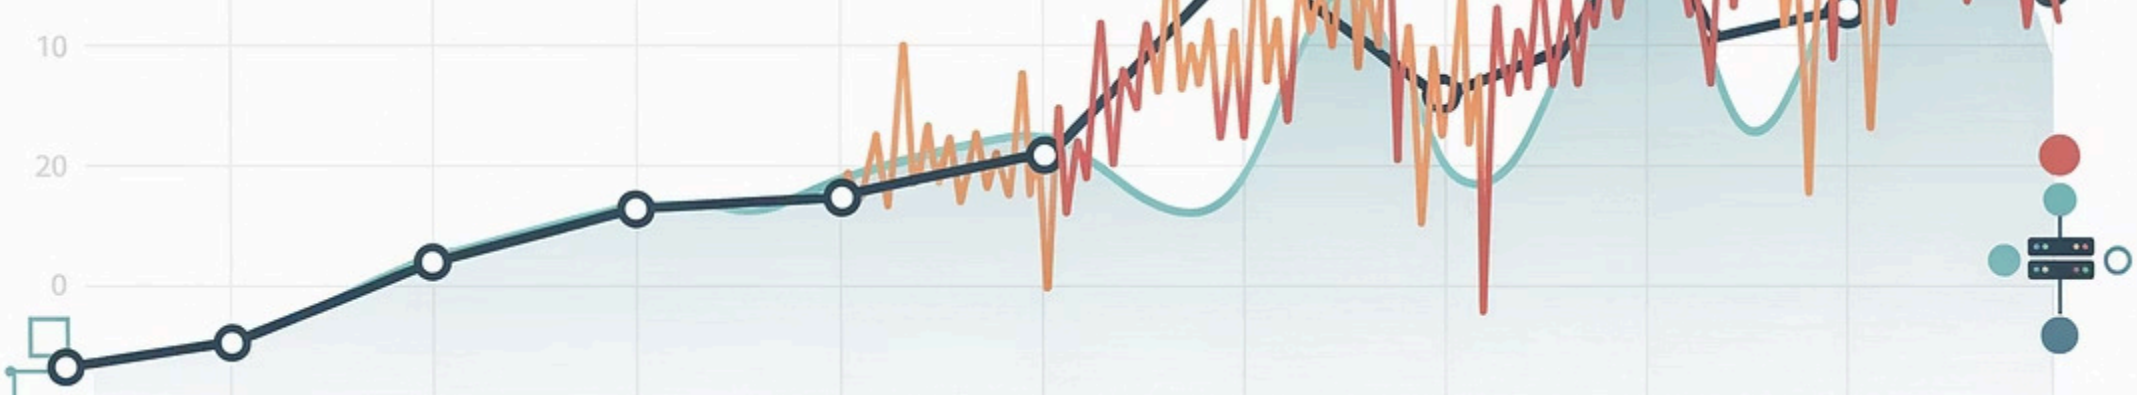

Systems Don't Fail Dramatically — They Degrade

Most incidents begin not with a bang, but with a slow accumulation of signals: latency creeps upward, error rates spike only in one region, a dependency starts timing out. A deployment introduces a subtle regression that won't surface in infra metrics for hours.

Meanwhile, customers experience failed checkouts, stalled insurance claims, and delayed bookings long before any infrastructure alarm reaches a critical threshold. The gap between **what monitoring sees** and **what customers feel** is where most reliability risk lives.

> Traditional monitoring excels at detection. Modern operations demands decisioning — and in many cases, **controlled action**.

AgentiAILLC

contact@azentiq.ai

The Fundamental Shift in Operations

This is not about adding more dashboards. It's about changing what observability is fundamentally *for* — shifting the operating model from passive visualization to active intelligence.

From	To
• Dashboards that visualize	• Intelligence systems that decide
• Alerts that notify humans	• Signals that trigger workflows
• Manual correlation across tools	• Agent-driven correlation across evidence
• Runbooks as documents	• Runbooks as controlled, executable actions
• Infrastructure health	• **Business + domain health**

AgentiAILLC

contact@azentiq.ai

The Core Architecture: Six Layers

A scalable agentic operations design is built on six well-defined layers — each with a distinct responsibility, designed to work together as a cohesive intelligence and action system.

1. Signal Collection

Technical, change, and domain signals ingested in real time

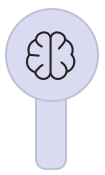

2. Operational Intelligence

Centralized "Ops Brain" — signals become searchable and correlatable

3. Reporting Layer

Real-time operational, analytical trend, and forensic investigation views

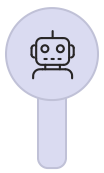

4. Agent Layer

Specialized agents: Triage, Root Cause, Domain Insight, Remediation, Reporter

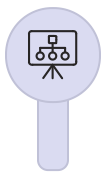

5. Workflow Orchestration

Repeatable, engineered workflows replace improvised incident scrambles

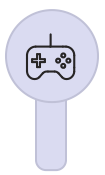

6. Action Layer

Approved runbooks as controlled APIs — technical and domain actions

AgentiAILLC

contact@azentiq.ai

Signal Collection & The Ops Brain

Three Classes of Signals

- **Technical:** logs, metrics, traces, latency, errors, saturation
- **Change:** deploys, config updates, feature flags, rule changes
- **Domain:** order placed, payment authorized, claim submitted; conversion rate, success rate, SLA breaches, stuck transactions

The Operational Intelligence Layer

The "Ops Brain" is a centralized system where all signals converge — becoming searchable, correlatable, and aggregatable in near real time. It answers the questions that matter most during an incident:

- "What changed right before the degradation?"
- "Is impact isolated to one region, tenant, or segment?"
- "Which dependency is the common factor?"
- "Which domain outcomes are failing, and how badly?"

A fast **read model** of rolled-up health snapshots ensures both humans and agents can instantly perceive the state of the world — without scanning raw telemetry every time.

AgentiAILLC

contact@azentiq.ai

The Agent Layer: Specialized Roles, Not One Mega-Agent

The most resilient pattern deploys a coordinated **agent team**, where each agent has a precise, bounded role. This is how you build "AI operating" rather than "AI guessing."

Triage Agent

Determines severity and scope; identifies impacted services, regions, and tenants; produces an evidence-backed "fact pack"

AgentiAILLC

contact@azentiq.ai

Root Cause Agent

Correlates technical, change, and domain signals; identifies the "first bad moment"; generates hypotheses with confidence scores

AgentiAILLC

contact@azentiq.ai

Domain Insight Agent

Translates system signals into domain impact; quantifies failed transactions, backlog growth, revenue risk, and SLA breach probability

AgentiAILLC

contact@azentiq.ai

Remediation Agent

Proposes and executes safe actions through controlled interfaces; runs post-action verification; escalates when risk or ambiguity is high

AgentiAILLC

contact@azentiq.ai

Reporter Agent

Posts stakeholder updates, opens and updates incident tickets, drafts postmortem timelines and recommended next actions

AgentiAILLC

contact@azentiq.ai

Domain-Aware Monitoring: The Game Changer

When domain data enters the intelligence layer, you can detect issues earlier and respond with precision that pure infra signals can never provide.

Events vs. States

Events tell you something happened. **States** tell you something is *stuck*. Domain state monitoring is where proactive operations shines:

- Transactions stuck beyond a processing threshold
- Backlog aging rising faster than normal baselines
- Success rates dropping for one segment or region
- SLA risk increasing even when infra metrics look "fine"

Business-Native Correlation

With domain signals in the intelligence layer, you can ask questions that map directly to business outcomes:

- "Is the conversion drop correlated with a recent flag change?"
- "Are failures isolated to a single provider or payment route?"
- "Is customer impact concentrated in a specific tier?"
- "Are we losing outcomes, or just seeing noisy infrastructure logs?"

This correlation capability transforms operations from infrastructure-centric to **outcome-centric**.

AgentiAILLC

contact@azentiq.ai

Proactive Operations: From Response to Prevention

Once domain signals are embedded in the intelligence layer, the system can shift from reactive incident response to proactive impact avoidance — running automated workflows triggered by conditions, not crises.

1

Post-Change Workflows

Every deploy or config change automatically triggers a health evaluation comparing pre- and post-change domain and technical signals.

2

Anomaly Workflows

Unusual deviations from baseline trigger investigation workflows before customers are impacted — catching regressions in minutes, not hours.

3

Trend Workflows

Slow degradation — memory leaks, growing queue lag, creeping latency — triggers early mitigation, not a 2 AM page.

4

Recurrence Workflows

"This pattern matches a known incident class" triggers known-safe action paths, turning past incidents into automated institutional knowledge.

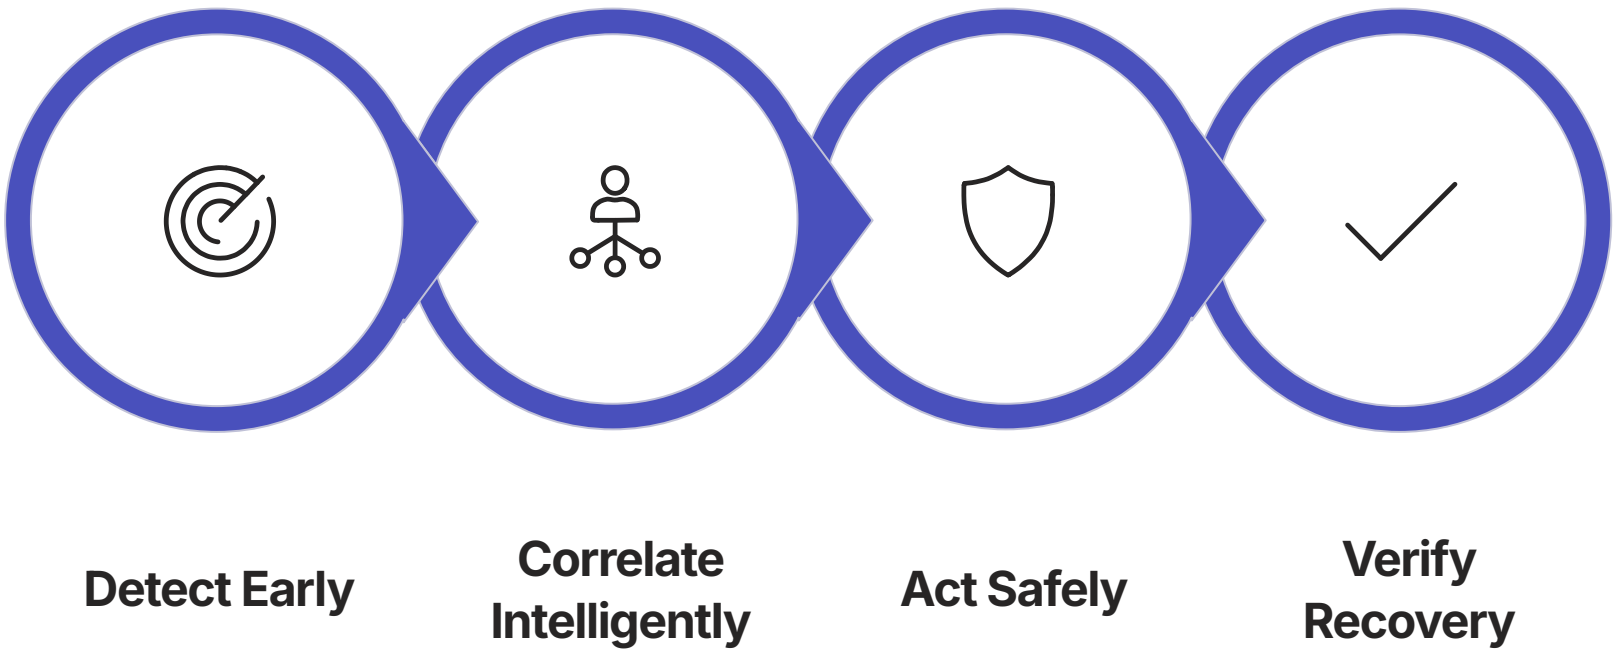

Over time, the system evolves from incident responder to impact avoider — compounding reliability gains with every workflow executed.

AgentiAILLC

contact@azentiq.ai

Guardrails: The Non-Negotiables

Agentic operations introduces real power — and real risk. Speed without control is chaos. Control without speed is stagnation. Strong guardrails are what allow both to coexist.

Safety Constraints

- **Least-privilege access** — agents operate only within their defined scope
- **Approved interfaces only** — actions invoked exclusively through controlled runbook APIs
- **Approval gates** for high-risk actions: rollbacks, traffic shifts, policy changes
- **Post-action verification loops** — every action must prove recovery before closing

Accountability & Ownership

- **Full audit trail**: evidence → decision → action, immutably logged
- **Clear ownership boundaries**: on-call approvals vs. domain owner approvals, explicitly defined
- **Escalation paths**: agents escalate to humans when risk or ambiguity exceeds defined thresholds
- **No improvised actions**: every automated step is a pre-approved, versioned workflow

AgentiAILLC

contact@azentiq.ai

What Success Looks Like

When this architecture is implemented well, the impact is measurable — across both engineering and business outcomes. Operations becomes a **scalable, engineered capability**, not a heroic scramble.

↓ MTTD

Faster Detection

Automated evidence packs and domain-first signals cut mean time to detect dramatically

↓ MTTR

Faster Resolution

Workflow-driven response replaces improvised scrambles, compressing time to recovery

↑ Signal

Better Signal Quality

Domain-correlated alerting reduces noise and surfaces only actionable, high-confidence signals

↓ Harm

Reduced Customer Impact

Domain-first mitigation protects outcomes before customers feel the failure

> Dashboards are necessary — but not sufficient. The next leap in reliability is building systems that can **observe, reason, and act safely**, with domain outcomes as the primary truth. When telemetry becomes an intelligence layer and alerts become workflows, operations becomes a scalable capability that protects both systems and the business they serve.

REFERENCE ARCHITECTURE | AGENTIC OPS | DOMAIN-AWARE RELIABILITY

AgentiAILLC

contact@azentiq.ai