ADITI
CONSULTING



AI-Powered Operations

A SOLUTION GUIDE

Detect, diagnose, and resolve repeat incidents across your existing stack, with governed automation that reduces MTTR, alert noise, and engineering toil.

The problem: capacity, not visibility

Modern teams can see more than ever, but they still spend too much time deciding what matters, finding the root cause, and repeating the same recovery steps.

Alerts arrive from separate monitoring tools. Tickets lack context.



Critical knowledge sits with a small number of experienced engineers. Incident volume grows faster than the team that must handle it, manual triage stretches MTTR, and senior engineers spend time restoring service instead of improving it.

The cost is measurable. A 10-person SRE team at $200K fully loaded represents $2M a year in operations capacity. At a 60% firefighting load, $1.2M of that goes to reactive work, before any downtime or SLA cost. Every 10-point reduction in firefighting returns $200K of capacity a year.

CORE PROBLEM

The real constraint is not effort. It is operational capacity.

The **solution**: complete the loop from alert to action

AI-Powered Operations is a services-led execution layer that works across the tools already in place: monitoring, observability, ITSM, and automation.

It connects signals, applies your approved runbooks, and gives the right work to a supervisor agent, a task agent, or a human expert.



Your tools stay.

The model completes the loop between them.

Ingest	Classify	Execute	Learn
Pull alerts, logs, metrics, traces, and ticket context from the current stack	Correlate signals, assess severity, and choose the right response path	Run approved remediation steps or route the issue with full context	Use outcomes and human decisions to improve confidence and coverage

Automation is tiered and earned. Agents begin by observing and recommending, move to approve-then-act for production changes with reliable rollback, and auto-resolve only high-volume, stable, low-risk patterns within firm limits. High-impact changes, novel incidents, and security or compliance decisions remain human-led.

Governance is part of the design

Six controls make the model usable in production.

Each agent receives only the access its task requires. The response path changes when confidence is weak. Sensitive actions wait for an accountable person. Every write action has a tested rollback.

Every input, decision, action, and outcome is recorded for audit. Named client roles own policies, runbooks, and production promotion. An agent never has more authority than the operating policy it serves.

Proof, with source and scope

70%+ Alert noise reduction, public telecom case study

Up to **95%** MTTR reduction for selected incidents (Buyer's Guide)

40–60% Auto-resolution for repeat patterns (planning range)

100+ Applications supported in the public case study

Approved materials also describe targeted incidents moving from hours to minutes.

Each figure carries its own source and scope; figures from different sources describe different measures and are never combined into one promise.

Results depend on scope, baseline quality, access, workflow repeatability, and approval requirements. Your program is measured against your own baseline.

What the **engagement** looks like

Pilot

About 4-6 weeks. A bounded scope of 3-5 high-volume incident patterns, measured against an agreed baseline

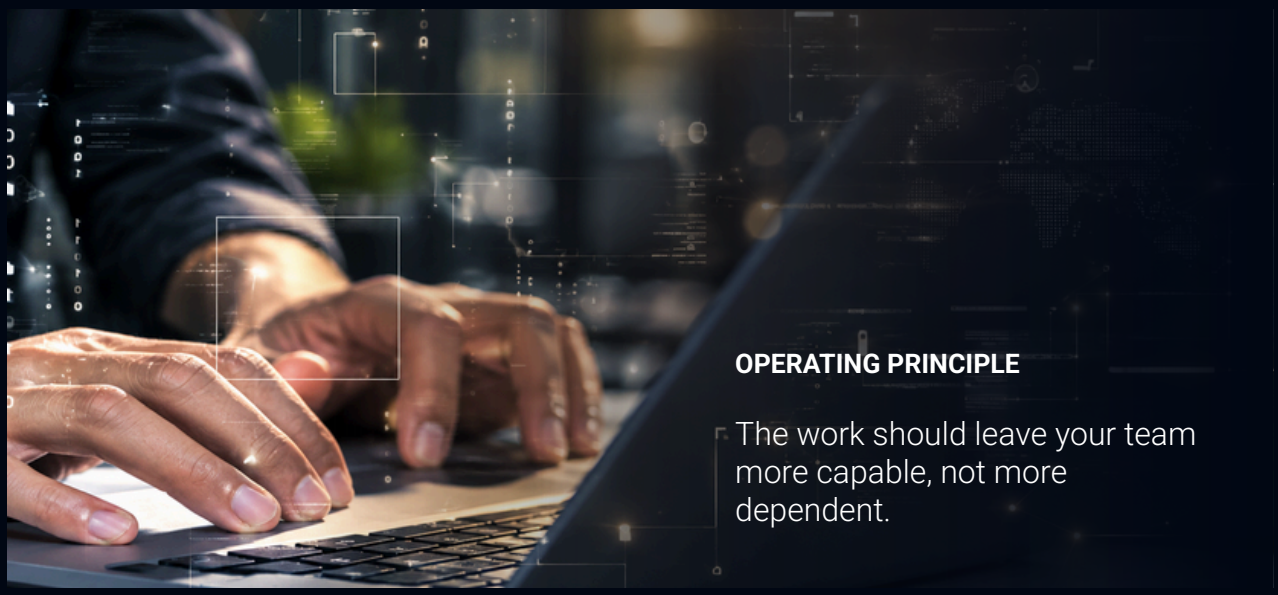
Deployment

About 8-12 weeks including production rollout, training, governance, and knowledge transfer

Ongoing capability

Quarterly review, new agents, tuning, and broader automation coverage as evidence supports it

The engagement leaves behind working assets and operating discipline: an incident taxonomy, integration design, supervisor policy and approval rules, a task-agent catalog built on your approved runbooks, an operational dashboard, and full capability handoff. You remain in control of policies, access, approvals, and exceptions throughout.



OPERATING PRINCIPLE

The work should leave your team more capable, not more dependent.



ADITI
CONSULTING

THE NEXT STEP

Two bounded ways to begin. A readiness assessment baselines your incident lifecycle and returns a prioritized automation map. A scoped pilot conversation defines a 4-6 week pilot: the scope, the measures, the guardrails, and the decision the pilot should enable.

Aditi Consulting is a digital engineering services company. We help clients design, build, and operate the systems and processes that support measurable business outcomes.

www.aditiconsulting.com
+1 (425) 305-5091
solutions@aditiconsulting.com

© 2026 Aditi Consulting. All rights reserved.