AI Operations

AIOps

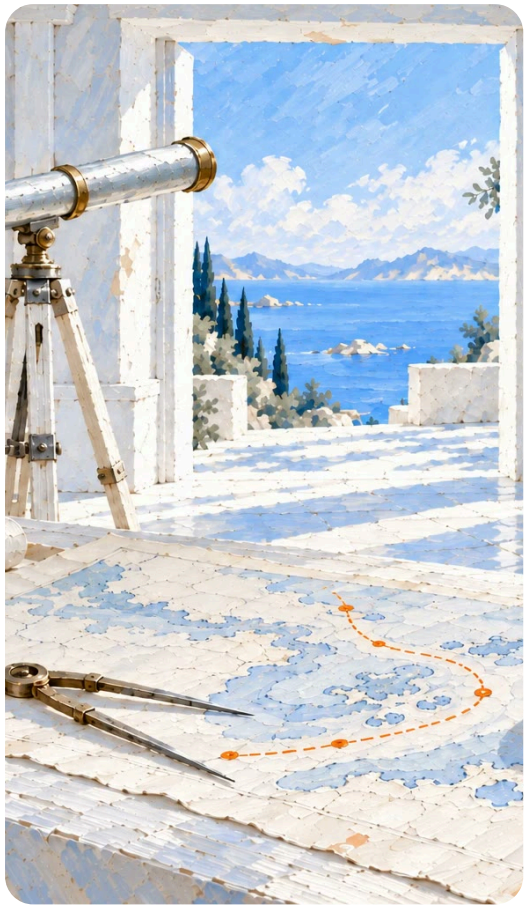
AI-Powered Failure Prediction & Response

Flags early signs of failure with the evidence,
and acts only **within the scope your company approves**.



This brochure doesn’t start with a feature list. It starts with the seven concerns IT operations teams actually raise when systems fail.

Each concern is followed by its solution, so you can read them in order.



Introduction

Introduction	Seven concerns operations teams face when systems fail	p.03
Solution	What AIOps is, and where it solves each concern	p.04

Seven concerns & solutions

Concern 01	We always learn about failures after our users do	p.05
Concern 02	So many alarms that the real cause gets buried	p.07
Concern 03	Finding the cause takes longer than the fix	p.09
Concern 04	Response know-how lives with just a few people	p.11
Concern 05	Letting AI touch production servers feels risky	p.12
Concern 06	No clear basis for trusting the AI’s judgment	p.14
Concern 07	Worried we’ll have to replace our monitoring tools	p.16

Use cases & adoption

Use cases	How incident response changes with AIOps	p.17
Adoption	Start in Approval mode, then expand step by step	p.19
Contact	Demo video · Contact sales	p.20

What we often hear

The seven concerns operations teams raise when systems fail

Your monitoring tools are in place. So why do failures still get fixed by people working through the night?

These are the stories we hear again and again from teams running large estates with small staffs.



Click the concern that weighs on you most to jump straight to how AIOps solves it.

Concern 01 “We always hear about failures from our users.” Early warning →

Concern 02 “Hundreds of alarms go off. Where do we even start?” Alarm triage →

Concern 03 “Finding the cause takes longer than the fix.” Root cause →

Concern 04 “Only one person really knows how to fix that.” Know-how →

Concern 05 “AI touching production servers? That worries us more.” Safe actions →

Concern 06 “What if the AI is wrong? No evidence, no trust.” Evidence →

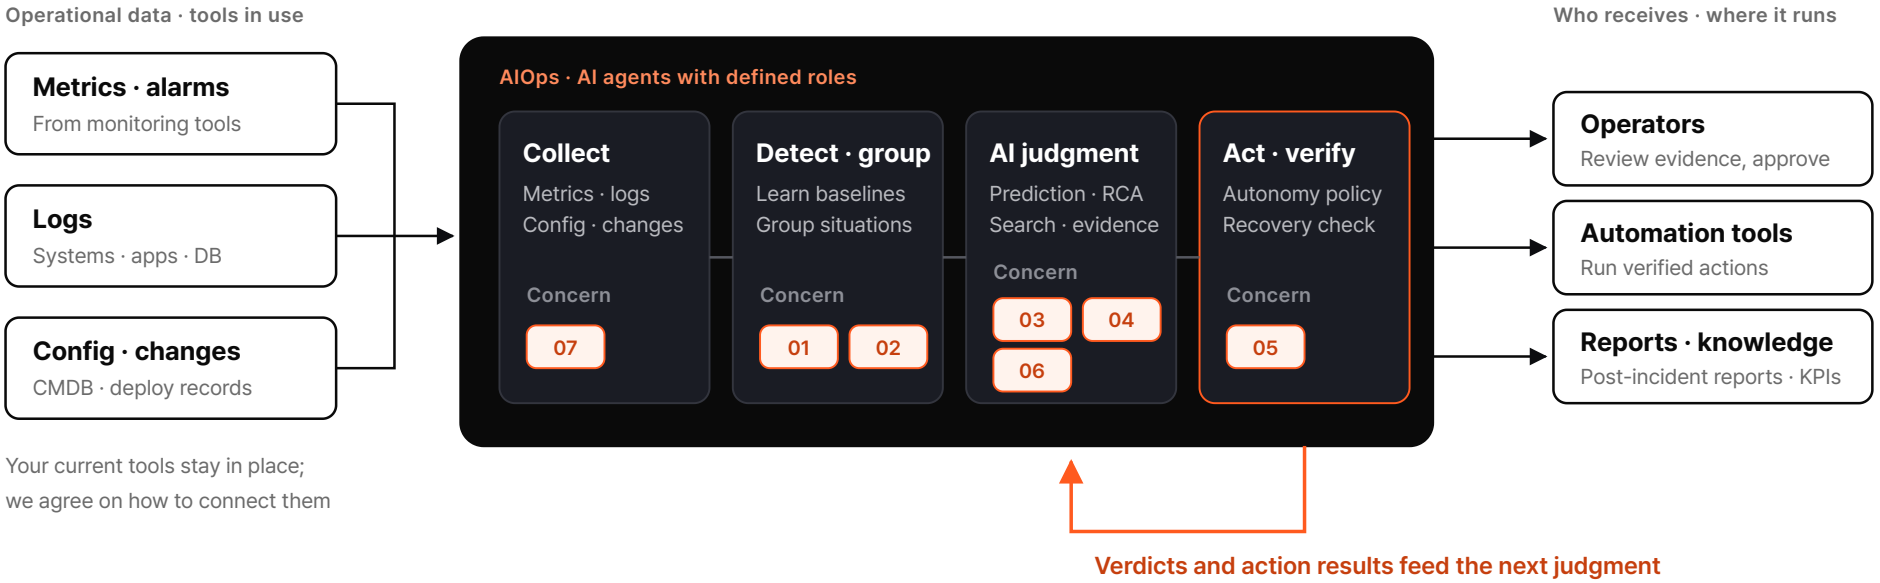
Concern 07 “Do we have to replace all our monitoring tools?” Existing tools →

One flow

AIOps solves these concerns in one connected flow.

AI-Powered Failure Prediction & Response

Adds AI judgment and action on top of the monitoring you already use: it warns you before failures, and when one happens, carries it from root cause analysis through action and recovery check.



Click a number under “Concern” to jump to that concern.

Concern 01 · Reactive response



We always find out about an outage when a user calls us.

The story we hear most often from heads of infrastructure operations.

What happens on the ground

Alarms fire only after the line is crossed

Disk, connections and memory fill up over days, but fixed thresholds only speak up once they're crossed.

The signs were there, but no one saw them

No one can review trends across hundreds of systems every day. The graphs get reopened only after the outage.

Late detection means longer downtime

The gap between when the anomaly began and when the response started turns directly into service downtime.

The hidden worry "That we'll learn about the next outage only after it hits."

See how AIOps solves it →

Solution 01 · Before the threshold

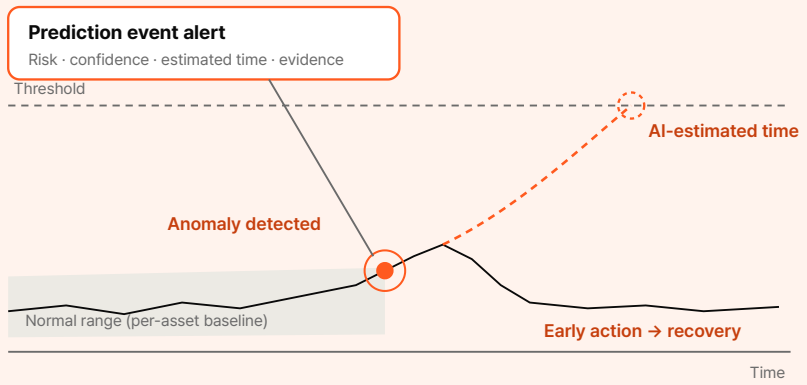
Alerts come while metrics are heading toward the line, not after they cross it.

Today



Alarms fire only after the threshold is crossed; the earlier signs slip by.

With AIOps



Detects the moment metrics leave their normal range, buying time to act before a failure.

A normal range for every asset

Each system's usual behavior becomes its baseline, and deviations are scored by how far they stray.

Trends, logs and history together

Rising trends are read alongside related logs and similar past failures, and AI judges whether they will lead to a failure.

Risk, confidence, expected time

Every prediction comes with its risk level, confidence and the time AI estimates it will occur.

Confirm, then act early

Once an operator confirms a prediction, pre-verified actions run before the failure.

What you gain

- ✓ Respond before failures occur
- ✓ No more checking graphs every day
- ✓ Planned actions, not emergency fixes

Concern 02 · Alarm floods



When hundreds of alarms go off, we don’t know where to look.

Something NOC and on-call engineers all tell us.

What happens on the ground

One cause turns into dozens of alarms

When a single DB slows down, alarms pour in all at once from the WAS, web and batch layers that depend on it.

Critical alarms get lost in the noise

Once people get used to repeated warnings, they start to miss the alarms that really matter.

Several people chase the same failure

Each engineer is working a different alarm, so efforts overlap and gaps open up.

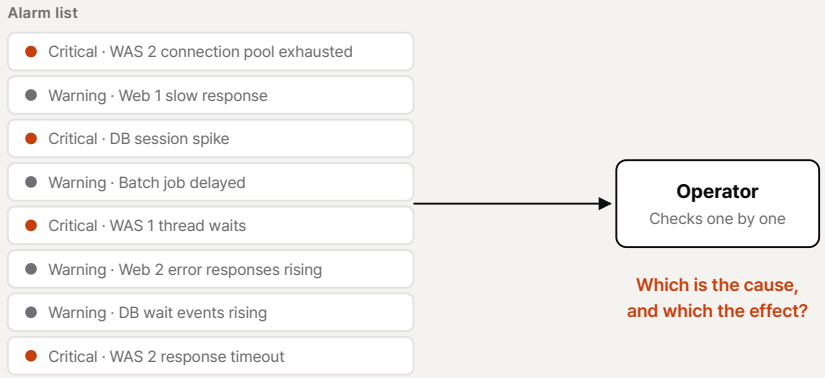
The hidden worry “Missing the real signal because we’re worn out by alerts.”

See how AIOps solves it →

Solution 02 · Situations, not alarms

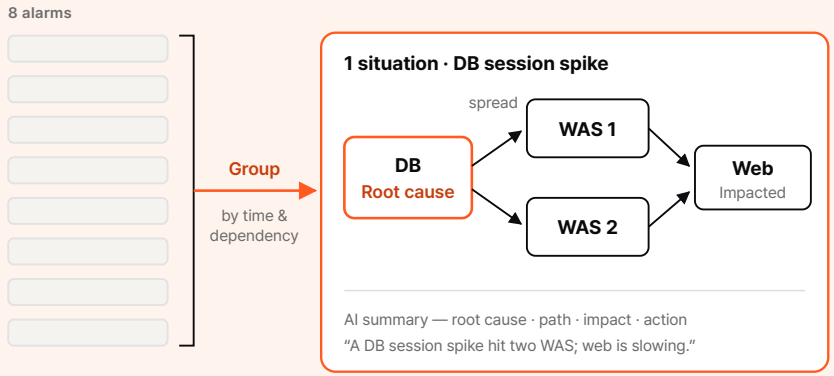
Scattered alarms are grouped into a single situation built around the root cause.

Today



People must untangle alarms where causes and effects are mixed together.

With AIOps



You see one situation, grouped around a single cause. (Sample screen)

Grouped by time and dependency

Signals from interdependent assets in the same time window become one situation.

Grouping follows rules

What gets grouped is decided by defined rules, not AI guesswork, so the same situation always yields the same result.

Shared gear alone stays apart

When shared storage or network gear is the sole overlap, signals aren’t grouped, so unrelated failures are less likely to mix.

One-paragraph summary

When two or more signals are grouped, AI summarizes the cause, propagation path, impact and recommended action.

What you gain

- ✓ Check a few situations, not alarm lists
- ✓ Tell causes from effects right from the start
- ✓ Every team sees the same situation on one screen

Concern 03 · Root cause analysis



Finding the cause takes longer than fixing the problem.

Something we hear at every incident review.

What happens on the ground

Cross-checking across tools

Metrics live here, logs there, change history somewhere else. Engineers jump between screens, lining up timestamps.

No idea how far it has spread

Even once the failure point is found, working out which services are affected means spreading out the system diagrams.

Missing what changed just before

A deployment or config change right before the failure is often the cause, but those records live in another system.

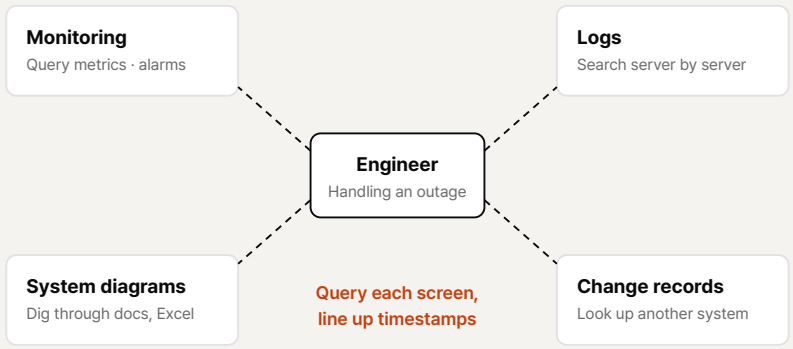
The hidden worry "The service staying down while the analysis drags on."

See how AIOps solves it →

Solution 03 · Cause, impact and response on one screen

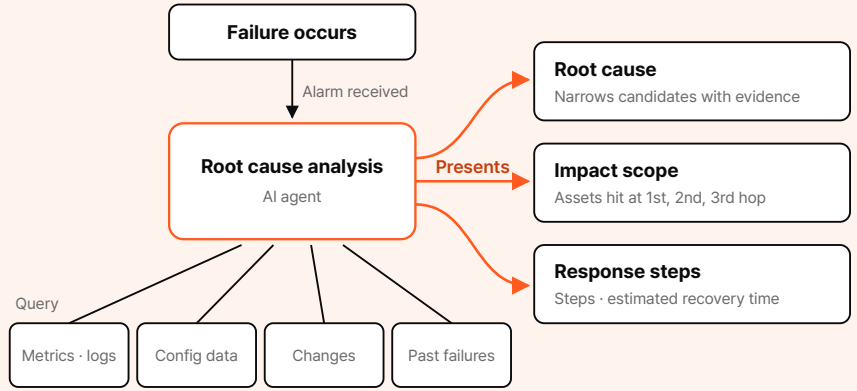
When a failure hits, the cause, impact scope and response steps come up together on one screen.

Today



Finding the cause depends on people's lookups and memory.

With AIOps



The agent gathers scattered evidence and presents cause, impact and response at once.

Narrowing down to the root cause

Combines metrics and logs to narrow cause candidates beyond the symptoms, and proposes step-by-step response with an estimated recovery time.

Impact scope as a map

Follows dependencies in your configuration data to show assets affected at the first, second and third hop, at a glance.

Recent changes checked early

Cross-checks change and deployment history just before the failure to see whether a recent change is the cause.

Similar past failures

Finds past failures that unfolded in a similar sequence, so proven responses can be reused.

What you gain

- ✓ Whoever is on call starts from the same evidence
- ✓ Notify affected services right away
- ✓ Spend the time on recovery, not on hunting for causes

Concern 04 · Know-how

In the end, only one person knows how to fix that one.



The records exist, but no one can find them
Incident reports and runbooks pile up, but the one line you need in a hurry is hard to find.

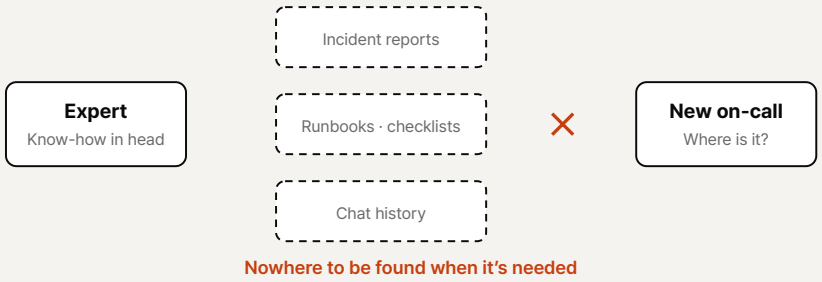
When people leave, the know-how goes too
When an expert moves on, the fixes they knew disappear with them.

The hidden worry
"Response quality wavering every time the person in charge changes."

Solution 04 · Know-how as company knowledge

It grows with every response, and anyone can put it to use.

Today



Know-how is tied to people and scattered documents.

With AIOps



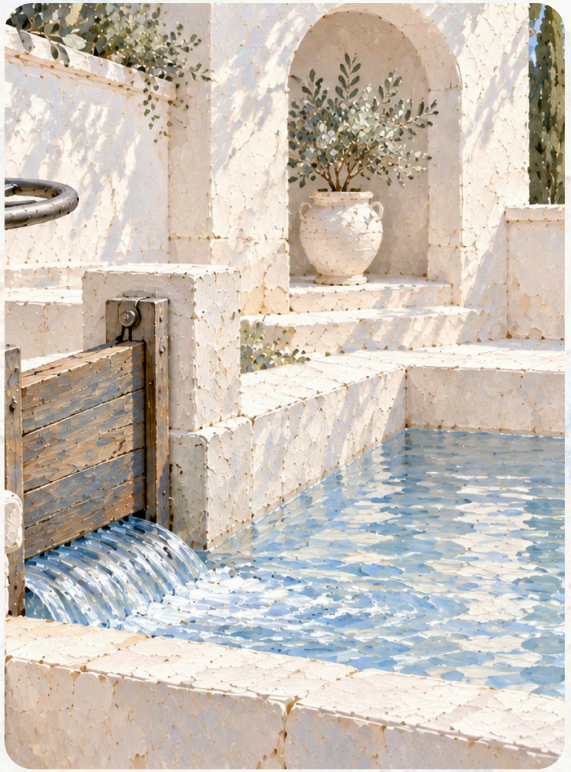
Every resolved incident adds to your company's knowledge.

What you gain

- ✓ Respond to the same standard, expert or not
- ✓ Response methods stay when people change

Back to contents

Concern 05 · Automated actions



AI working directly on our production servers? That worries me even more.

The question operations leaders ask right before signing off on automation.

What happens on the ground

Manual work at 3 a.m. invites mistakes

When a tired on-call engineer types commands by hand, the recovery itself can turn into another outage.

But handing it over is unnerving

There's no telling what automation with no visible reasoning will do to production servers.

Running it doesn't mean it's fixed

Even when a script reports success, someone still has to check whether the service has actually recovered.

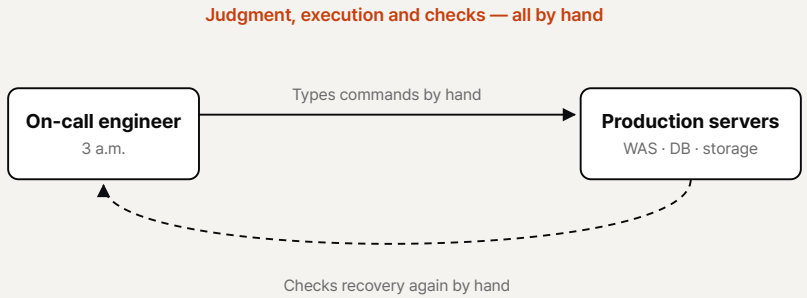
The hidden worry **"Taking the blame when automation makes an incident worse."**

See how AIOps solves it →

Solution 05 · Controlled automation

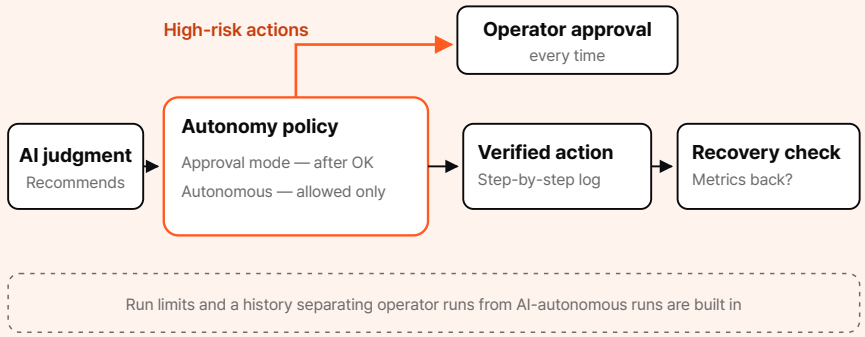
AI acts only within the limits your people set, and checks that the service has actually recovered.

Today



A setup where mistakes and missed checks happen easily.

With AIOps



Your organization decides what to automate, and results are confirmed by metrics.

Start in Approval mode

At the outset, operators approve every action; autonomy then expands, beginning with the action types that have earned trust.

High-risk: always approved

Each action has a risk level and an auto-run setting, so high-risk work needs human approval even in Autonomous mode.

Only verified procedures run

AI doesn't invent new commands. It picks from remediation playbooks verified in advance.

Results are verified

Success means the target metrics actually recovered after the action, not just that the run finished.

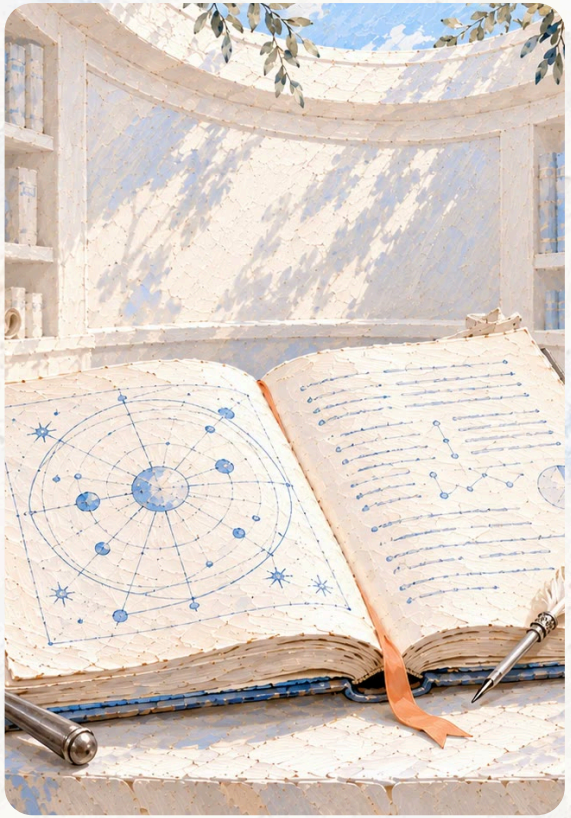
What you gain

- ✓ The same procedure, nights and weekends included
- ✓ Fewer mistakes from typing commands by hand
- ✓ Run history makes accountability clear

Concern 06 · Trusting AI judgment

What if the AI is wrong? We can’t trust a call with no evidence behind it.

What operations teams evaluating AI are most wary of.



What happens on the ground

No idea why it decided that

When AI only reports a result, operators have no way to check whether it’s right.

Wrong calls keep repeating

If the false alarms you flag never shape the next judgment, all you get is one more alarm.

What AI did stays invisible

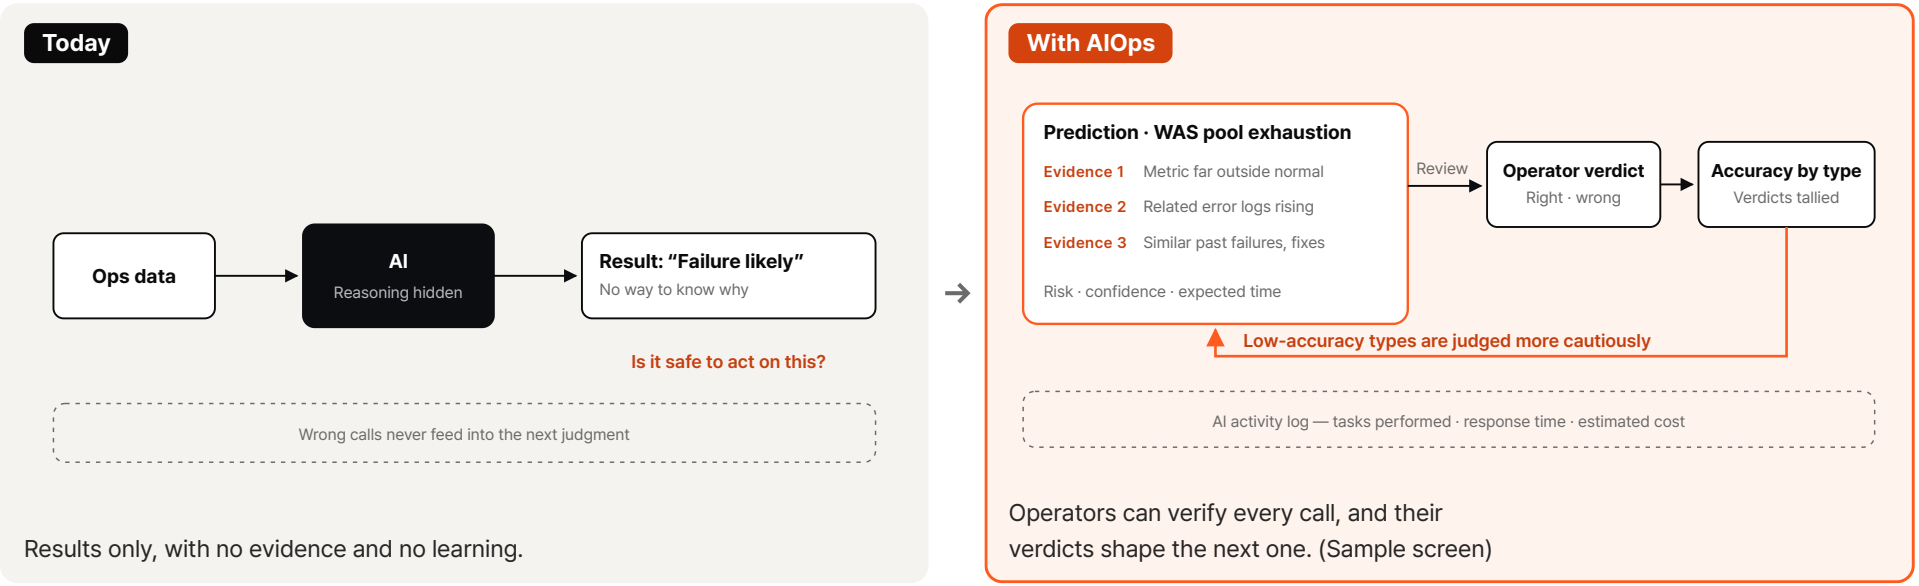
It’s hard to see what AI handled, how much of it, and what it cost.

The hidden worry “Handing operations over to decisions no one can explain.”

See how AIOps solves it →

Solution 06 · Explainable judgment

Every judgment comes with evidence, and operators' verdicts shape the next one.



Evidence with every judgment

Each prediction comes with how far metrics strayed from normal, related logs and similar past failures.

Verdicts build into accuracy

Operators' right/wrong verdicts are tallied as accuracy by type, and low-accuracy types are judged more cautiously.

Separate roles for rules and AI

Signal grouping, which must always give the same result, runs on rules; explanation and prediction go to AI.

AI's own work is logged

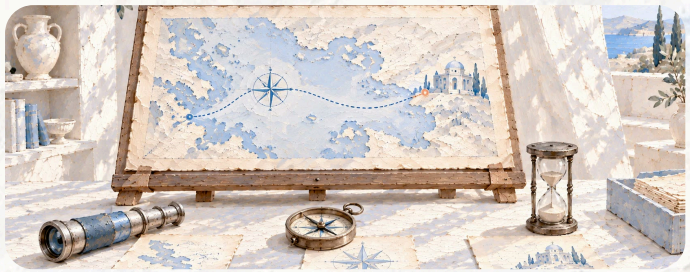
AI tasks, response times and estimated costs are recorded. Operations Copilot only analyzes and recommends; it never executes.

What you gain

- ✓ Judgments operators can verify for themselves
- ✓ Predictions that learn from operator verdicts
- ✓ AI usage and estimated costs in plain view

Concern 07 · Existing tools

Do we have to replace all the monitoring tools we use today?



We've already invested in plenty of tools

Every monitoring, logging and automation tool represents time and money already spent.

Data is scattered across tools

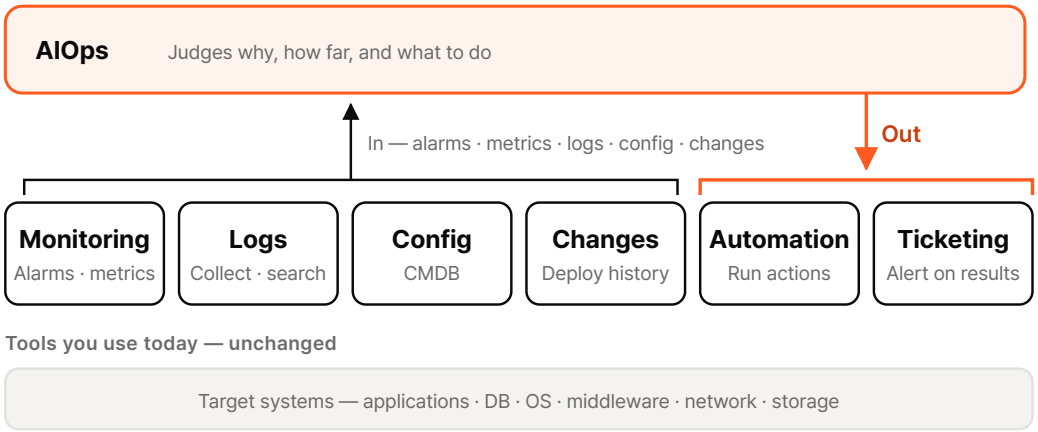
Metrics, logs, configuration data and change history all sit apart, so it's hard to see the whole picture at once.

The hidden worry

"Ending up with yet another tool to run."

Solution 07 · Add, don't replace

Keep the tools you have, and add judgment on top.



Monitoring tells you what crossed a threshold; AIOps works out why, how far it reaches and what to do.

Standard connections

Connections are designed around standard methods such as webhooks, APIs and syslog.

Scattered data, one place

Metrics, logs, configuration data and change history are organized on one common basis.

Scope defined together

Connection targets and methods are finalized during the current-state review.

What you gain

- ✓ Keep using the tools you've invested in
- ✓ Scattered data comes together in one judgment

Back to contents

Use case examples

With AIOps in place, incident response changes like this.

DB operations · Capacity management

Know well before storage fills up

What you gain — capacity issues resolved **during business hours, with planned actions**.

Today	What AIOps handles	What your team does
Tablespace usage creeps up until it crosses the threshold, and only then does an emergency expansion begin.	Flags a trend toward saturation against the normal range, with related logs and history as evidence.	Review the evidence and approve; expansion and recovery check follow the set procedure.

AIOps capabilities used: Solution 01 Early warning · Solution 06 Evidence · Solution 05 Safe actions

Connects to: DB metrics · logs · Incident history · Remediation playbooks

Service operations · Cascading failures

One situation, not dozens of alarms

What you gain — a response order set **with causes and effects separated from the start**.

Today	What AIOps handles	What your team does
A DB session spike sets off WAS, web and batch alarms at once, and each engineer chases a different one.	Groups same-window, dependent signals into one situation, flagging the root-cause DB and the affected WAS.	Teams work from that one situation and focus together on fixing the cause.

AIOps capabilities used: Solution 02 Situation view · Solution 03 Impact scope · Solution 07 Tool integration

Connects to: Monitoring alarms · Config · dependencies · Logs

These are illustrative use cases. Actual scope and impact are designed together with you, around your systems and the way you operate.

Use case examples (continued)

The same approach, across other work.

Infrastructure ops · Storage failure

See how far it spread without unrolling diagrams

What you gain — **incident notices and recovery start at the same time**.

Today
When a storage path fails, engineers dig through diagrams and documents to find each affected server and service.

→

What AIOps handles
Shows the root cause and affected assets hop by hop, with recent changes, past failures and response steps.

→

What your team does
Notify the affected services, then approve and run the proposed recovery action.

AIOps capabilities used: Solution 03 **Root cause** · Solution 04 **Knowledge reuse** · Solution 05 **Safe actions**

Connects to: **Config data (CMDB)** · **Change · deploy history** · **Storage metrics · logs**

Operations management · Reporting

No more writing reports from scratch

What you gain — **with every incident resolved,** your team's knowledge and evidence grow.

Today
After recovery, someone has to write up the timeline, cause and actions, and track prevention measures separately.

→

What AIOps handles
After recovery, drafts a post-incident report, saves it as knowledge, and logs accuracy and recovery time.

→

What your team does
Review and finalize the draft, then use the metrics to decide which failure types to automate next.

AIOps capabilities used: Solution 04 **Knowledge base** · Solution 06 **Verdict loop** · Solution 05 **Autonomy scope**

Connects to: **Incident history** · **Post-incident reports** · **Performance metrics**

These are illustrative use cases. Actual scope and impact are designed together with you, around your systems and the way you operate.

Adoption steps

You don’t hand everything to AI on day one. Start in Approval mode, verify, then expand.

1

Review the current state together

We review your operations tools and asset landscape, then define connection points and scope.

What you gain

Where to connect, and which failure types to tackle up front

2

Start in Approval mode

Apply prediction and root cause analysis to core assets; every action runs only after operator approval. Approval verdicts reveal accuracy by type.

What you gain

Your own metrics showing whether AI judgment holds up for you

3

Autonomy for low-risk actions

Once a type’s accuracy is stable, let its low-risk actions run automatically. High-risk actions still need approval.

What you gain

Early relief for night and weekend on-call

4

Expand targets and scope

Guided by recovery-check results and accumulated knowledge, widen the assets and automation in scope.

What you gain

Growth guided by prediction accuracy and recovery time on your own metrics

For a first meeting, a few of these are enough

Operations tools in use

A list of your monitoring, logging, automation and ticketing tools

Recent recurring failures

Two or three failure types that have recurred in recent months

Asset and configuration data

Server and equipment lists with dependencies (CMDB or Excel)

Standard remediation steps

The runbooks and check or remediation scripts you use today

Start light with SaaS

Define the asset scale and connection scope, then start small with core assets.

Dedicated environment on request

We agree on what data goes to AI and a customer-dedicated cloud configuration that meets your security requirements.

Approval mode by default

Every action starts with operator approval, and the scope of autonomy can be adjusted at any time.

Talk to us

Start by telling us your challenges.

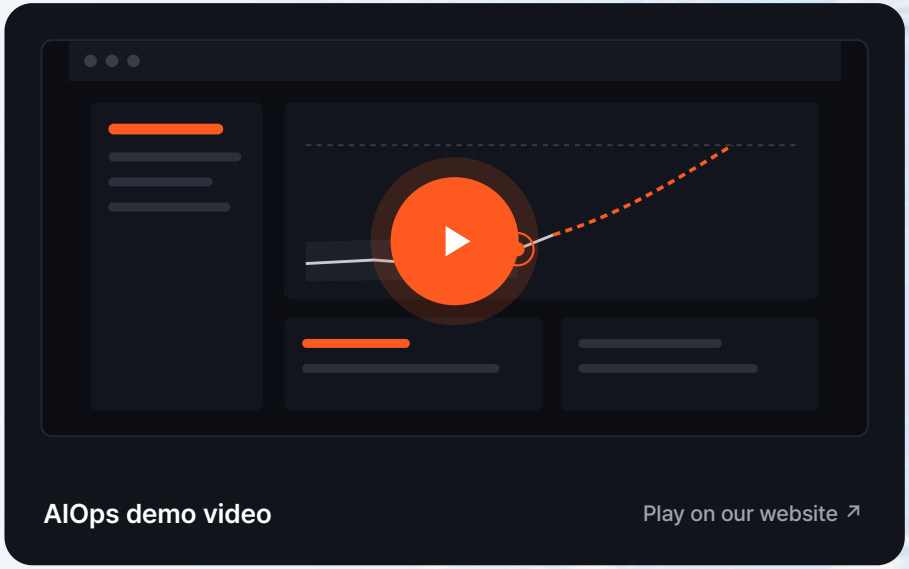
Tell us about one failure that keeps coming back, and we'll show you in a demo how AIOps would flag a similar failure early and take it through approval and remediation. A few items from the list on page 19 are all we need. Want to look around first? Watch the demo video.

Contact sales →

▶ Watch the demo

When you reach out, tell us your most pressing concern

Early warning | Alarm triage | Root cause | Know-how | Safe actions | Evidence | Existing tools



Email
sales@vicsnc.com

Website
vicsnc.com