§ THE PLAYBOOK SERIES · FREE · NO FORM, NO FOLLOW-UP SEQUENCE

# The AI playbook for *IT Operations*

*How to take the ticket win without believing the vendor's deflection number.*

IT ops has the clearest AI economics of any function and the widest gap between claimed and actual results. This playbook covers what containment really looks like, what happens to L1, why your ITSM vendor is selling you both the disease and the cure, and a 90-day path that survives contact with reality.

> *Gartner's verdict on AI in infrastructure and operations: 28% of use cases fully succeed, 20% fail outright, and the rest stall somewhere in the middle. The difference is rarely the model. It is usually the knowledge base nobody wanted to maintain.*

**28%**

of AI use cases in I&O fully succeed; 20% fail outright.

GARTNER, APRIL 2026

**41.2%**

median enterprise automated-resolution rate - against vendor claims of 70-80%.

ZENDESK BENCHMARKS, 2026

**9,000 → 5,000**

Salesforce support headcount after AI took half its interactions.

BENIOFF / CNBC, SEPTEMBER 2025

INSIDE THIS PLAYBOOK

PART ONE

# What actually *changes.*

The evidenced win in IT ops is boringly specific: high-volume, low-complexity tickets. Password resets, access requests, software provisioning. IBM's AskIT cut calls and chats by 74% - a real number, built over roughly eight years of unglamorous work on knowledge, data and process. Not a 90-day miracle, and not a chatbot bolted onto a mess.

Meanwhile the vendors have made their position clear. ServiceNow is now openly selling the automation of the L1 role, and paid $2.85bn for Moveworks to own the deflection layer. Your ITSM supplier is monetising the elimination of the seats it also charges you for. Worth remembering at renewal.

01

### L1 tickets get contained, genuinely

Best-in-class deflection runs at 55-65% for the routine tier. But note the honest median is 41%, and only around 14% of issues reach full self-service resolution. Deflection and resolution are different words because they are different things.

02

### Incident response gets an assistant, not a replacement

What works in production: alert correlation, noise reduction, automated timelines, drafted postmortems. Autonomous remediation remains rare, and everyone credible keeps a human in the loop. AIOps promised more for a decade; Gartner quietly retired the category name in 2025.

03

### The knowledge base becomes the product

Knowledge quality determines deflection performance more than model choice does. The unfashionable work - articles, ownership, hygiene - is now the highest-leverage engineering in the department.

04

### Average ticket difficulty goes up

AI skims the easy tickets, so everything reaching a human is harder. Your L2 and L3 people feel busier, not freer. Plan capacity and morale for that, because the dashboard will insist things improved.

05

### Outsourcing gets repriced

HCL's own CEO forecasts 3-5% 'AI deflation' on revenue; clients are reopening outsourcing contracts within 24 months rather than waiting for renewal. If your provider's costs are falling and your invoice is not, you know where the saving went.

06

### Autonomy stays on a leash

Agents that touch infrastructure need scoped permissions, approval gates and audit trails - the governed-worker model. Anything else is an outage with a countdown timer.

PART TWO - THE HONEST BIT

# What it costs *you.*

The headcount story here is real and public. Salesforce took support from 9,000 people to 5,000 - Benioff's phrase was 'I need less heads'. The US Bureau of Labor Statistics now forecasts computer support roles declining 3% over the decade, explicitly citing automation - the first official forecast to pin an IT occupation decline on AI.

And yet: Gartner predicts half the companies that cut headcount citing AI will be rehiring for similar work under new job titles by 2027. Klarna already did, publicly. The honest plan is neither 'nobody goes' nor 'everyone goes' - it is knowing which tickets a machine resolves before you resize anything.

01

## L1 shrinks - after resolution is proven

The routine tier gets smaller. The mistake is sequencing: cut on the vendor's deflection claim and you rediscover the difference between deflected and resolved through your escalation queue, like Klarna did through its NPS.

02

## The ITSM licence paradox lands on your budget

Now Assist is priced at a 50-60% uplift plus consumption fees - a premium paid to the platform to reduce the seats you pay the platform for. Do that maths before you sign, and put seat reduction in the same business case.

03

## Outsourced contracts need reopening

People-and-time pricing breaks when AI shrinks the team. Reopen the contract this year, or fund your provider's margin expansion out of habit.

04

## The remaining humans need a new deal

Harder tickets, incident command, agent supervision, knowledge stewardship - genuinely more senior work. Pay and title should follow, or the people capable of it will do it somewhere else.

05

## Some AIOps investment is sunk

If your monitoring platform delivered alert deduplication dressed up as intelligence, say so, stop the roadmap sunk-cost, and put the money into knowledge and automation that resolves things.

06

## The vendor count is a warning in itself

Gartner reckons only around 130 of the thousands of 'agentic' vendors are real; the rest are agent-washing. Your procurement process is now a security control.

PART THREE

# What good looks *like.*

Good IT ops on the other side of this measures resolution, not deflection. Machines own the routine tier with proper guardrails; humans own incidents, judgement and the knowledge that feeds the machine. Contracts - internal and external - reflect the new cost of the work rather than the old shape of the team.

01

### Containment with an honest denominator

The routine tier runs through AI with resolution measured end-to-end, including the tickets that bounce back. The number reported upward is the one a sceptical auditor would accept.

02

### Humans on incidents and judgement

Incident commanders, escalation engineers, agent supervisors and knowledge owners - a smaller team doing distinctly harder work, with the career path and pay to match.

03

### Governed agents, scoped blast radius

Named jobs, scoped permissions, approval gates, audit trails and a definition of done. If an agent can touch production, someone can show you exactly what it is allowed to break.

04

### Contracts that track the new economics

ITSM licences, outsourcing agreements and internal chargebacks all repriced against what the work now costs - reviewed annually, because the economics will not sit still.

THE CHANGE, ACROSS ALL FIVE PILLARS

| PEOPLE | PROCESSES | DATA | PLATFORMS | PARTNERS |
|---|---|---|---|---|
| L1 resized only after resolution is proven, with a named senior path for the people who stay. | Escalation, approval gates and incident command redesigned around the agent, not bolted on after it. | The knowledge base treated as the product - owned, audited, and current enough to trust. | ITSM licences and AI uplifts priced against seat reductions, in the same contract. | Outsourcing agreements reopened to follow the new cost of a ticket, not the old shape of the team. |

PART FOUR

# The first *90 days.*

One ticket category, contained properly and measured honestly. IBM took eight years; you are buying the earned patterns, not the timeline - but you still cannot skip the sequence.

*I*

WEEKS 1-2

### Baseline the ticket reality

Volume, categories, true cost per ticket, current resolution rates - and an honest audit of knowledge-base coverage for your top ten categories. This audit predicts your result better than any vendor demo.

*II*

WEEKS 3-6

### Contain one category in production

Pick the highest-volume routine category. Fix its knowledge articles first, then stand the agent up with escalation paths and approval gates. Measure resolution, not deflection.

*III*

WEEKS 7-10

### Verify, then touch the money

Check escalation queues, reopen rates and user satisfaction against the baseline. If it holds, open the commercial conversations: seat counts, the outsourcing contract, the licence uplift.

*IV*

WEEKS 11-13

### Widen, and reshape the team openly

Take the next two categories. Publish the honest role map: what shrinks, what the senior path looks like, who gets trained into knowledge and agent-supervision roles. Before the rumour mill drafts it for you.

THE DAY-90 SCORECARD · WHAT YOU SHOULD BE ABLE TO SHOW

- → End-to-end resolution rate for the contained category - not deflection
- → Escalation and reopen rates against the week-two baseline
- → Knowledge-base coverage of the top ten ticket categories
- → User satisfaction for AI-handled versus human-handled tickets
- → One commercial conversation opened: seats, uplift or the outsourcing contract

PART FIVE

# The traps - and six questions worth an *honest hour.*

| | |
|---|---|
| 01<br>**Believing the vendor's number**<br>Claimed deflection of 70-80% against a real-world median of 41% is not a rounding error. Ask for resolution rates, reference customers, and the definition of 'resolved'. Watch them flinch. | 02<br>**Cutting before resolving**<br>Headcount reductions based on projected containment are how you end up rehiring at contractor rates with an apology. Klarna has kindly run this experiment for everyone. |
| 03<br>**Skipping the knowledge work**<br>An agent on top of a stale knowledge base is a very fast way to give wrong answers politely. The boring work is the work. | 04<br>**Paying twice at renewal**<br>Signing the AI uplift while keeping every seat is the vendor's dream outcome. Make seat reduction a condition of the same contract, in writing. |
| 05<br>**Unscoped agents in production**<br>An agent with broad permissions and no approval gate is an incident report that has not happened yet. Blast radius is a design input, not a lesson. | 06<br>**Forgetting the humans who remain**<br>The team that stays gets harder work and more responsibility. Treat that as a promotion problem to solve, not a productivity gain to bank quietly. |

THE SELF-ASSESSMENT · MORE THAN TWO UNCOMFORTABLE ANSWERS IS AN AGENDA, NOT A CRISIS

- → Do you know your real resolution rate - not deflection - for AI-handled tickets?
- → Would your knowledge base survive an audit of its top ten ticket categories?
- → What does your ITSM renewal cost with the AI uplift and without the seat reduction?
- → When did you last reopen the outsourcing contract - and who benefits from the delay?
- → Which agents can touch production, and what exactly are they allowed to break?
- → What is the promotion path for the L1 people you intend to keep?

START WITH THE WORKSHOP

## Book your *Catalyst Workshop.*

1-2 executive days to turn this playbook into your first move, the case for it, and a 90-day mobilisation plan. hello@veriteer.com

HELLO@VERITEER.COM →