WHERE TO START WITH AI · NO. 6

Triaging and routing IT tickets

An assistant for triaging internal support tickets, rolled out in two stages: first it suggests a triage for the technician, then it routes approved standard requests on its own.

Stage 1 · the technician decides, with AI support

Stage 2 · AI routes standard cases; support handles exceptions

Illustrative case: fictitious company; all figures are approximate.

WHY START HERE

It meets all four tests of a good first workflow

High volume

Hundreds of tickets a month across about thirty categories.

Clear owner

The IT organization, which already maintains the service catalog.

Clear output

A ticket correctly categorized and assigned, with a likely fix.

Known pain point

For support staff and employees alike.

Written rules and articles. The assistant works from a service catalog and a knowledge base the company already has. The risk does not go away. It becomes more visible, measurable, and manageable once the categories, the scope, and the escalation rules are defined.

Illustrative case: fictitious company; all figures are approximate.

ON ONE CONDITION

Low risk, but only within a strict scope

1. **A written service catalog**
   current categories, priorities, and teams.

2. **No admin rights**
   and no account actions taken by AI.

3. **Password resets and standard access requests**
   through the existing secure workflows.

4. **Incidents, security issues, and ambiguous cases**
   routed to technicians.

> Outside this scope, it becomes a different use case with a different risk profile.

Illustrative case: fictitious company; all figures are approximate.

THE WORKFLOW AS IT RUNS TODAY

Six steps, all done by hand

Example Industries Inc. · 900 employees · fictitious company

1	**Ticket opened by the employee** via portal, email, or phone	
2	**First-line triage** category, priority, team	● Overloaded shared queue
3	**Search for a fix** knowledge base, past tickets	● Scattered or outdated articles
4	**Assignment to a team** devices, network, apps, access	● Tickets reassigned several times
5	**Handling by a technician** often a request seen before	● The same fix explained again
6	**Closure and filing** reply to the employee	● Categories hard to analyze

≈ 1,500
tickets a month

≈ 40%
are simple, recurring requests

≈ 1 day
until a simple request is picked up

● friction point · illustrative figures

Illustrative case: fictitious company; all figures are approximate.

BEFORE YOU START

Get the catalog in order

The work starts by clarifying the service catalog: categories, the priority matrix, the teams responsible, and the knowledge base articles. Each article has an owner and a version date; duplicate and outdated articles are retired. The tool comes after.

1
Categories and priority matrix

2
Current articles, no duplicates

3
One owner per article

4
Dated versions

> **The tool comes after.**

STAGE 1 · HUMAN IN THE LOOP

The assistant suggests a triage; the technician decides

The ticket comes in
through the same support tool as today.

↓

The assistant suggests a triage
category, priority, team, and cited article.

↓

The technician decides
confirms, corrects, or reassigns.

↓

Every correction is logged
category by category.

What this stage produces: a category-by-category measure of how accurate the suggested triage is.

Estimated duration: 6 to 10 weeks.

STAGE 2 · THE ASSISTANT ROUTES STANDARD REQUESTS

It routes or answers standard requests; support handles the exceptions

THE ASSISTANT ROUTES OR ANSWERS WHEN

- The request falls in an approved category
- A current article or a secure workflow applies
- The team and priority follow from written rules

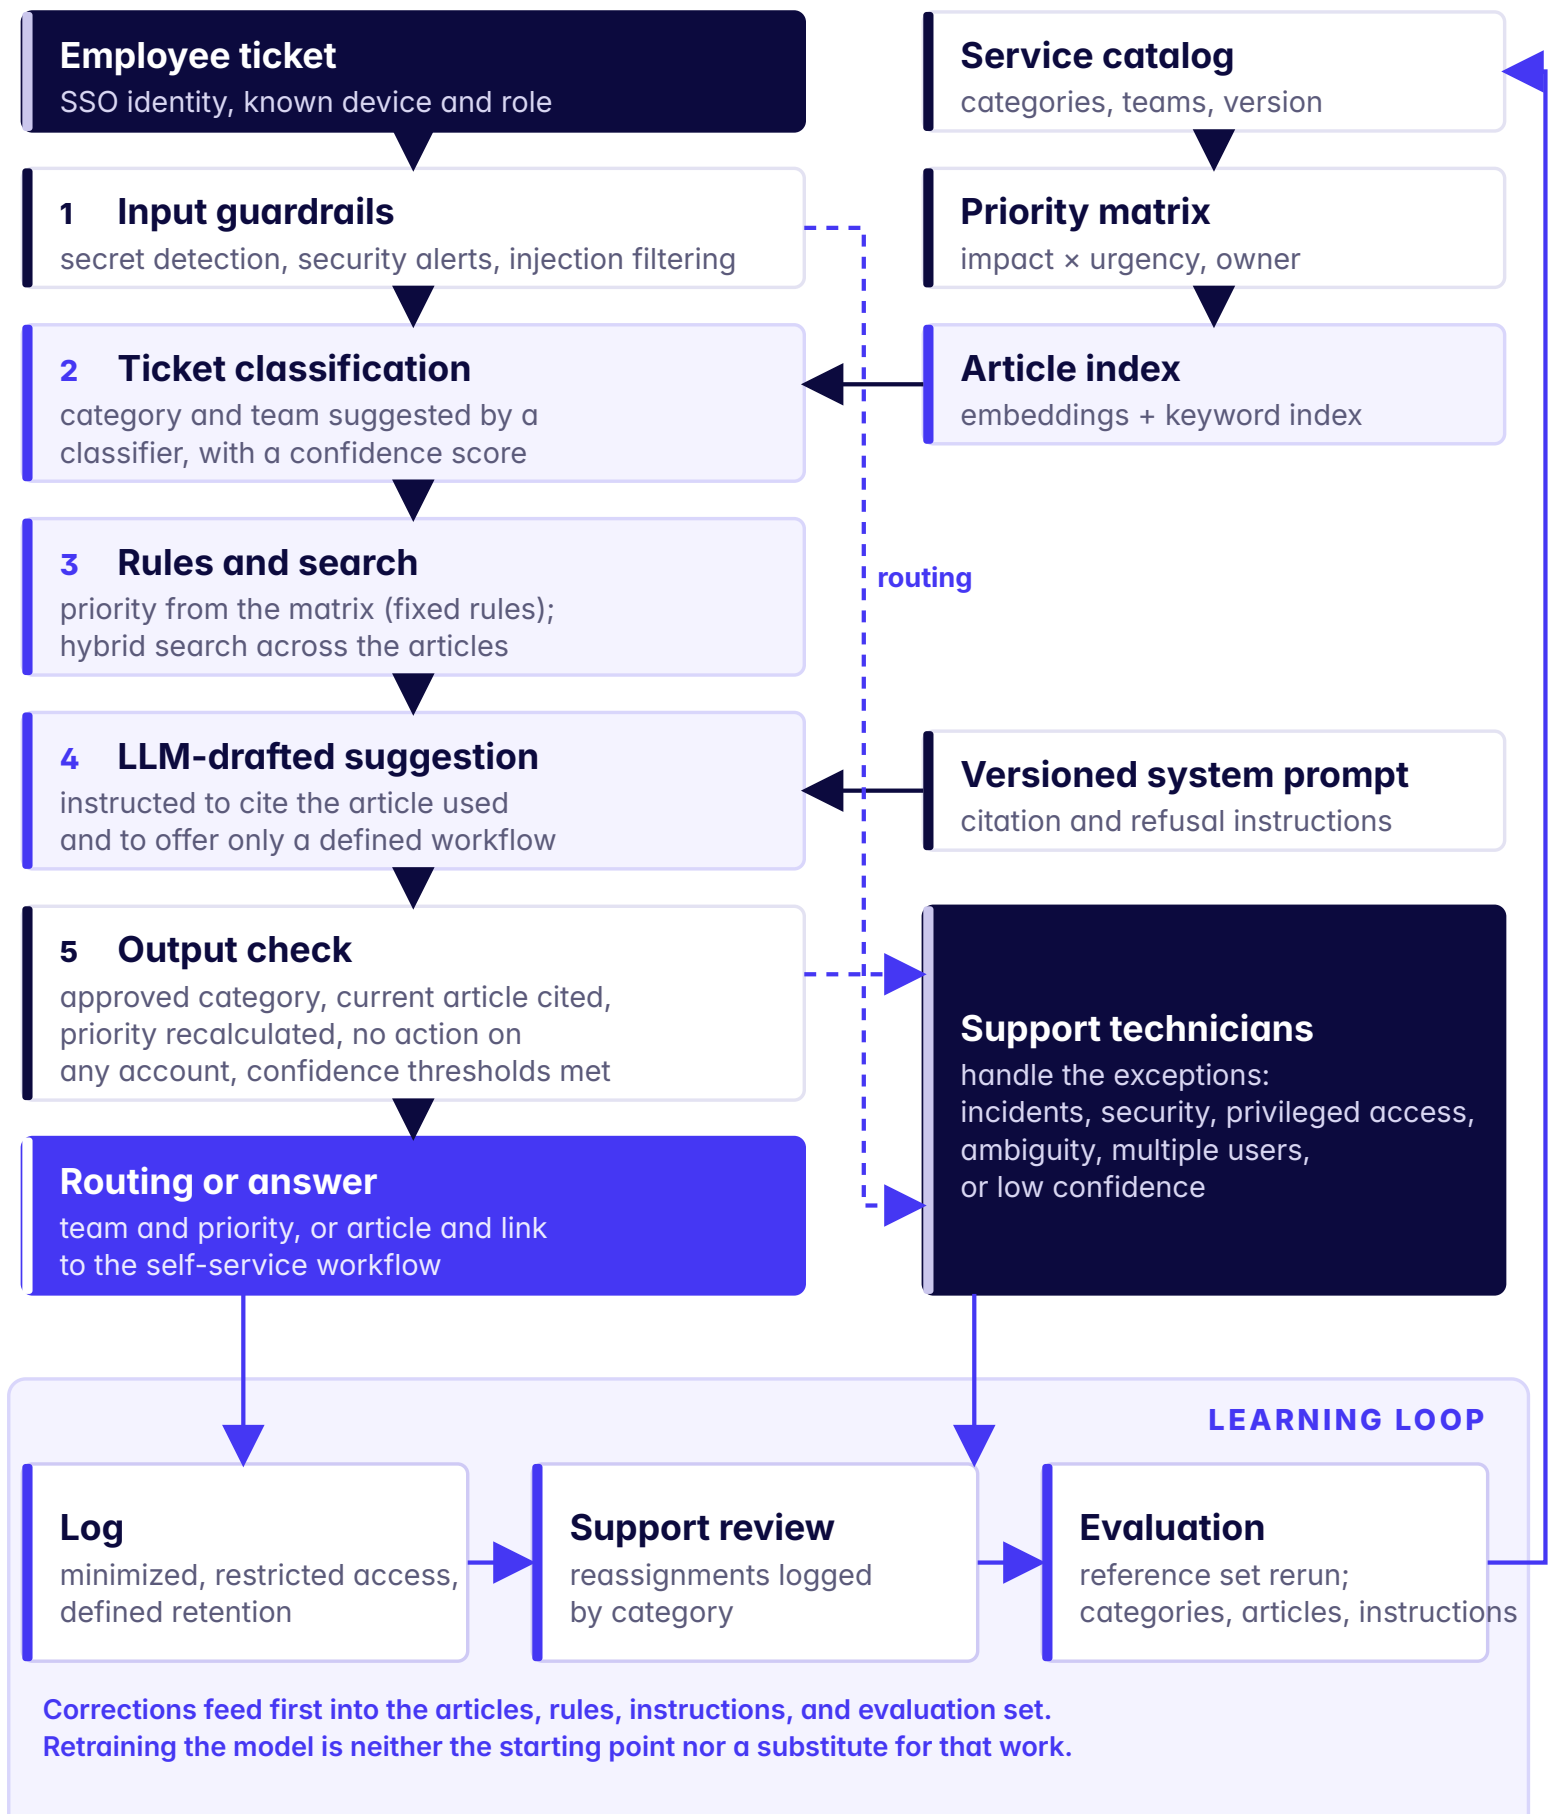
IT ESCALATES TO SUPPORT WHEN

- The request is an incident, urgent, or ambiguous
- The article is missing, conflicting, or outdated
- It involves security, privileged access, a business-critical device, multiple users, or similar

> Every week, support reviews a sample of routed tickets and keeps the articles up to date.

UNDER THE HOOD · CLASSIFICATION, RULES, SEARCH

Reference architecture: how the system stays under control



Illustrative case: fictitious company; all figures are approximate.

UNDER THE HOOD · CONTROLS AND SECURITY

The controls required before any autonomous answer

1 **Evaluation**
A reference set of about 300 past tickets, minimized, each with the expected category, team, and priority. Metrics: triage accuracy, reassignment rate, citation accuracy, and the rate at which exceptions are detected and correctly routed. Rerun whenever the model, the catalog, or the instructions change.

2 **Observability**
Every ticket processed is logged: category, priority, suggested article, confidence score, routing decision. Since tickets can contain personal data or secrets, the log is minimized, accessible only to authorized staff, and kept for a defined period.

3 **Identity and access**
SSO (SAML or OIDC). The assistant has no admin rights: password resets and access requests go through the existing workflows, including MFA and approvals.

4 **Data**
Hosting and data transfers verified for the chosen solution; contractual terms barring use of the data for model training; encryption; log retention and access defined; documented Swiss FADP assessment before launch.

5 **Application security**
Prompt injection defenses for ticket text; attachments, if analyzed, are extracted in an isolated environment. Only four outcomes: suggest, route, answer, or escalate. No command is ever run on any system.

> Three options: the AI module in your existing service desk tool, a managed classification and search platform, or a custom build (classifier, rules engine, model API). The choice comes after the catalog audit, not before.

Illustrative case: fictitious company; all figures are approximate.

MOVING FROM STAGE 1 TO STAGE 2

A decision based on evidence, category by category

The move to stage 2 should not be decided in advance. It depends on what support observes during stage 1: the accuracy of the suggested triage, the corrections needed, and the system's ability to recognize tickets that fall outside its scope.

WHAT SUPPORT OBSERVES

- Accuracy of the triage and the cited article
- Reassignments needed, by category
- Exceptions correctly routed
- No significant incident within the tested scope
- Current articles, each with an owner

> Printer requests could move to stage 2 while application access stays in stage 1.

Illustrative case: fictitious company; all figures are approximate.

THE GOAL, AND HOW TO MEASURE IT

The goal is not to replace technicians. It is to cut response time on simple requests, stop the same tickets from being triaged and reassigned by hand, and focus support time on incidents that call for diagnosis and judgment.

Time until a simple request is picked up

≈ 1 day → **minutes**

today (illustrative) — illustrative stage 2 target

Also tracked: the share of simple requests routed or answered without a technician, **30–40%** (illustrative target)

Data protection	Swiss FADP: employee notice, minimized data, and hosting and transfers clarified before launch.
Governance	The IT organization owns the outcome and reviews a sample of routed tickets every week.

Illustrative case: fictitious company; all figures are approximate.

In this example, AI does not replace the technicians' diagnosis.

It steps in first where a clear category and a current article are enough to route the request.

So the useful question is not "Which tool should we buy?" It is "Which repetitive work can we have AI assist with, measure, and then delegate under clear controls?"

> **What repetitive work in your company already meets these conditions?**

Sergio Castagna Consulting.
sergiocastagna.com · info@sergiocastagna.com · Geneva

Sergio Castagna Consulting is a member of the OpenAI Partner Network and a registered partner in the Google Cloud Partner Network. The tool comes after the problem is defined.