AI agents for the work around the ticket and the change

Agents that read your tickets, your runbooks, your policies and your change records, prepare the work, and stop for a named person in IT. On infrastructure you control.

THE PROBLEM

Of 14 public IT workflow templates we opened, 3 have a step where a person approves before a change and 7 have none. Of 158 public interface descriptions from 64 publishers, 132 import, 24 are over the import limit and 2 are refused; the 132 carry 31,951 operations, 16,969 sent as changes, so tools are chosen one by one.

HOW IT WORKS, IN THREE STEPS

STEP 01

Connect what the agent may read.

Runbooks, how-to articles, IT policies and past incident reviews become knowledge sources, with each reader's permissions kept. The ticket system, the directory and the asset register are reached through the description of an interface you already run; alerts and metrics through servers in your own cluster, with read access.

STEP 02

Decide where it stops.

Every tool call passes your policy. A tool that would reset a sign-in, grant access or change production waits for a named person in IT.

STEP 03

Run it where the tickets and logs live.

In your data centre or your private cloud, on a model you serve.

THE RULE EVERY AGENT FOLLOWS

It prepares. A named person in IT decides.

Thirty-five agents in nine areas

The service desk	Service Desk Triage. IT How-to Q&A. Ticket Reply Draft. Ticket Handover Summary
Access and identity	Access Request Preparation. Credential Reset Request Check. Joiner, Mover and Leaver Checklist. Access Review Preparation
Incidents	Incident Response. Major Incident Timeline. Incident Update Draft. Post-Incident Review Draft
Change and release	Change Request Review. Change Conflict Check. Rollback Plan Check. Release Note Draft
Assets and licences	Software Licence Reconciliation. Asset Register Check. Renewal and End-of-Support Check. Hardware Refresh List
Devices and patching	Patch Status Summary. Device Compliance Exception. Software Request Check
IT knowledge	Knowledge. Knowledge Article Draft. Runbook Draft. Knowledge Gap Report
Operations and capacity	Alert Digest. Backup Check Summary. Capacity Report. Cloud Cost Summary
Running the department	Service Level Report. Problem Record Draft. Audit Evidence Preparation. IT Supplier Review

Also in the pack: the full list of what each agent reads and where it stops, and the white paper “AI agents inside the IT department”, with sources.

Seven questions to ask before an agent touches an account or a production system

01 Where does it run, and where is the model?

02 Does it act as the person using it, with that person’s permissions?

03 Which tools reset, grant, change or delete, and does each wait for a person?

04 Are the read operations and the change operations separate tools?

05 Who verifies the caller before a reset, and who carries it out?

06 What is recorded, and for how long?

07 Can we stop one agent, one tool, or all of them, at once?

WHERE THE PLATFORM RUNS TODAY

In production at a top-tier bank in Asia, air-gapped, inside the bank’s own perimeter.

MODON built its AI Centre of Excellence on Katonic inside the Kingdom of Saudi Arabia.

ePLDT, the ICT arm of the Philippines’ largest telco, selected Katonic to power Pilipinas AI, a sovereign AI cloud hosted in the Philippines.

Where to start

One team, one kind of ticket, invented data first. An IT how-to Q&A agent can run on the first day, because no account or device is touched.

KATONIC.AI/SOLUTIONS/IT