AI agents inside the operator

What they can read, what they can reach, and who decides

IN BRIEF

01 Operators have put AI to work at scale. Care assistants handle millions of contacts, and agents now work beside engineers in the network. In most published cases an agent recommends and a person makes the change.

02 An operator's work runs through its support systems and its network. No public catalogue of ready-made connections holds them, and the vendors who sell them do not publish their interfaces.

03 The industry does. TM Forum's Open APIs and CAMARA's network APIs are published for anyone to read, so the shape of an operator's interfaces is known before the first meeting.

04 An operator needs one place where every tool an agent can call is listed, every call is checked against policy, anything that would change the network or a customer's account waits for a person, and the record of it all can be put in front of an auditor.

05 All of this can run on infrastructure the operator controls, in the country its law names. Several operators already say theirs does. And the platform an operator runs for itself is one it can sell.

IN THIS PAPER

01 Where operators are today 2
02 In nobody's catalogue, but published 3
03 Three doors, one gate 4
04 What an agent may do 6
05 What regulators already ask 7
06 The operator as a seller 8
07 Seven questions before an agent touches the network 9
08 Where to start 9

PART 01

Where operators are today

2 million
contacts a month taken by True's assistant in Thailand

1 million
network devices behind NTT Docomo's maintenance agents

1 hour to minutes
for a radio network fix at Deutsche Telekom

100
AI use cases reported by Telkomsel

Customer care is at scale. True, in Thailand, says its assistant takes nearly 2 million contacts a month and handles nearly 90% of them itself, with people taking the rest. Telstra's internal assistant serves more than 8,000 front-line staff. Telkomsel reports more than 100 AI use cases across its network, care, sales and IT.

In the network, most agents recommend and an engineer acts. NTT Docomo put an agent system into commercial use in February 2026. It analyses data from more than one million network devices and presents recommended actions to maintenance engineers, and Docomo says it has cut the response time for complex failures by more than half. Telia keeps people in the loop for automated execution until its engineers trust the results. At AT&T a person oversees the chain of agents, with required checkpoints.

A few act on the live network. Deutsche Telekom's radio network agent finds anomalies and starts corrective actions by itself, and work that took about an hour now takes a few minutes. The lead of its next programme has said that for every critical change a person approves before it runs. Rakuten Mobile runs energy saving in its radio network with no person in the loop, and reports around 20% less energy used.

Some run it on their own infrastructure, and say so. SoftBank trains and runs its telecom model in its own data centres in Japan. SK Telecom says it designs and operates its own AI infrastructure, and runs an assistant with over 10 million monthly users on its own model. Telenor runs a sovereign AI factory in Norway.

Full autonomy is claimed one task at a time. The industry's own scale for network autonomy runs from zero to five. Every claim of the fourth level that we opened is for a single task, such as saving energy in the radio network or handling one kind of fault.

The caution is the operators' own. True says an assistant must never offer a price plan that does not exist.

So the question for an operator is not whether agents work. It is what each one may change, who approves it, and whether that holds.

PART 02

In nobody’s catalogue, but published

327
servers in Docker’s public catalogue on 19 November 2026

0
of them an operator’s support system or network controller

113
public API files given to the platform’s import, all accepted

43%
of the tools a standard file makes are for systems, not agents

The Model Context Protocol, MCP, has become the common way to give an agent a tool. A system is wrapped once as an MCP server, and any agent can use it.

The largest public catalogue of these servers is Docker’s. We read every entry. On 19 November 2026 it held 327 servers.

Not one is an operator’s support system, billing system, inventory or network controller. What the catalogue does hold is the layer around the network: servers for metrics, logs, databases and the clusters that network functions run on. That is useful, and it is the edge.

The vendors are moving faster here than in other industries. Every large vendor of support and network systems has launched agents of its own, sold inside its product, and several now say their systems speak MCP. Few say in writing that their agents can run on the operator’s own infrastructure. An operator with five vendors gets five sets of agents, five places where the rules are set, and no single list.

There is a second fact underneath this one. **Vendors do not publish their API descriptions.** Of about 60 products we surveyed, two had a description anyone can download.

What telecom has instead is two public standards. TM Forum publishes its Open APIs, more than a hundred of them, for the business and operations side: customers, orders, bills, trouble tickets, alarms, inventory. CAMARA, an open-source project of the Linux Foundation, publishes the APIs of the network itself, the ones behind the GSMA’s Open Gateway: whether a SIM was changed recently, where a device is, what quality a connection can be given. Both are published as OpenAPI files under an open licence.

We tested what that is worth. On 21 November 2026 we gave Katonic’s import 113 of these files as published: 91 from 25 TM Forum APIs, and 22 from the latest release of 14 CAMARA APIs. It accepted all 113.

Three things follow from reading them. A standard file is a contract, not a system: none of the TM Forum files says how to sign in, and each leaves the address blank, because the operator’s own gateway supplies both. The files carry more than an agent should have: 43% of the tools they produce exist for one system to notify another, so somebody has to choose. And the network’s APIs ask for a stricter sign-in than a shared password: a token the caller signs with its own private key, valid for five minutes at most.

So the route to an operator’s systems does not run through a catalogue. It runs through the operator’s own gateway, in a shape the industry has already published.

PART 03

Three doors, one gate

An operator's work has documents at the edge, systems of record at the centre and the operator's own rules over both. An agent platform needs a door to each.

The first door is for documents. A contract, a tender, a customer's letter, a recording of a call. On Katonic a user gives the agent the files in the conversation, and a recording is transcribed by a speech model the operator serves. The standing documents, the runbooks, equipment manuals, tariffs, product terms and the complaints code, are indexed once as knowledge, with each reader's permissions kept, so an agent answers from the document and cites the section.

One rule belongs here. **Where the runbook or the record is silent, the agent says so.** An answer about a procedure says when the runbook does not cover the fault. A draft review lists what the tickets do not show. An agent that fills a gap in an incident's timeline has written something nobody observed.

The second door is the catalogue, for the edge. Katonic's catalogue of MCP servers is drawn from Docker's and carries Katonic's own additions. A server is deployed as a container inside the operator's own cluster. This is how an agent reaches metrics, logs, a reporting database, or a web search for a regulator's notices. Where a vendor's system offers MCP itself, it comes in through this door too.

One caution applies to any catalogue. A container that runs inside the operator may still call a service outside it. A search or mapping server must never be given a subscriber's number or location, and that has to be a policy on the tool, not a line in a prompt.

The third door is the operator's own APIs, for the centre. The operator gives Katonic the description of an interface it already runs: the file its own system produces, or the public file of the standard it follows, with its own gateway's address. The operations it chooses become tools an agent can call.

- The description is checked before anything is created. The operator sees every operation in it, and any warnings.
- The operator ticks the operations an agent may have. The ones that exist for systems to notify each other are left out.
- Each tool is rated by what it does: reading is low risk, changing is medium, deleting is high.
- A tool can be tried with real values, or as a dry run that shows the exact request and sends nothing.
- Each connection has its own sign-in, including the signed token the network's APIs expect, a limit on calls per minute, and a time limit.
- A document can travel with a call: a bill fetched as a file, or a contract attached to a ticket.

No agent talks to network equipment. Where a change to the network is involved, the agent prepares it and an engineer makes it through the change process.

All three doors lead to one gate. A document, a tool from a catalogue server and a tool from an imported interface sit under the same policies, with the same record.

THREE DOORS, ONE GATE

Door 01	Door 02	Door 03	All three
Documents.	**The catalogue.**	**Your own APIs.**	**One gate.**

An agent builder sees tools. The operator's security and risk teams see one inventory.

PART 03 · CONTINUED

Six ways into an operator's systems

Most of what moves inside a network does not move through a modern API. Configuration, alarms, counters and call records travel over protocols of their own, and as files. The work of giving an agent a way in is done in products the operator already owns, and it follows a small number of patterns.

	THE SITUATION	WHAT THE OPERATOR DOES	WHAT THE AGENT PLATFORM DOES
01	**A support system or a network controller already offers MCP**	Switches it on and issues a credential	**Connects to it, and sets a policy on each tool**
02	**A system follows a TM Forum Open API**	Gives its own gateway's address and sign-in for that API	**Imports the public file, with the operations the operator chooses**
03	**The network's own APIs follow CAMARA**	Registers the platform as a client, with a public key	**Imports the released file, and signs in with a signed token**
04	**A system produces its own API description**	Exports the file from its own installation	**Imports the file**
05	**Equipment or a system that speaks only a network protocol, files or a message bus**	Reaches it through the controller, the mediation platform or the data platform it already runs	**Imports that product's file, or reads the data platform through a database server**
06	**A system with only a screen, such as a partner's or a regulator's portal**	Records the steps in a screen-automation product	**Connects to that product's tool, for reading only**

Two things follow. Nothing is replaced: the agent comes in through the same gateway as every other consumer of the operator's interfaces. And the oldest parts of the network are usually behind pattern 5.

That work is the operator's, and it is worth doing once for every consumer, not for agents alone.

PART 04

What an agent may do

Reaching a system is half of it. The other half is deciding what an agent may do there, and being able to show the decision afterwards.

Every call is checked before it runs. On Katonic the order is fixed: is this user allowed at all, do the inputs carry anything that looks like an injection attack, and what does the operator's policy say about this tool.

A policy gives one of four answers. Allow it. Allow it and record it. Hold it until a person approves. Block it. A policy can cover one tool, every tool on a server, or every tool, and it can apply to the whole company or to one team. Where two policies meet, the stricter one wins.

A tool that would change something waits. When an agent is published, any tool that changes something and has no policy of its own is set to wait for approval. That is how nothing reaches the network or a customer's account unapproved: the ticket update, the credit and the notice each wait for the person whose job it is. An operator can relax that for a tool. It has to decide to, and the decision is on record.

The platform errs towards holding. Some look-ups, such as asking the network whether a SIM was changed recently, are sent in the form a change would take. They wait for approval too, until the operator gives that tool its own policy.

No agent changes the network. The operators that let software act alone in the network did it one narrow, reversible task at a time, after long trials. An agent platform should not make that choice for them.

An agent gets the least it needs. Tools are granted to each agent and policies set for each team, so a billing agent has no tool that reads location, and a network agent has none that reads a bill. Secrecy of communications is a rule about what may be read, not only about where it is stored.

If the policy cannot be read, the call is refused.

The agent acts as the person using it. Katonic passes the signed-in user's identity with each call. The value comes from the platform, not from the model, so an agent cannot be talked into acting as someone else. Where the operator's interface enforces it, a care agent's assistant sees the customers that care agent may see.

Every call leaves a record: who asked, which agent, which tool, what the policy decided and who approved.

There is always a way to stop. An agent, a tool or a whole server can be switched off at once.

A POLICY GIVES ONE OF FOUR ANSWERS

Answer 01	Answer 02	Answer 03	Answer 04
Allow it.	**Allow it and record it.**	**Hold it until a person approves.**	**Block it.**

A fault in the control must not become permission.

PART 05

What regulators already ask

No regulator we read has written a binding rule for AI agents in a telecom operator. Two have published guidance on agents for organisations of every kind, and it reads like section 4 of this paper. Everywhere else, agents fall under texts an operator already lives with: secrecy of communications, network security, data protection and consumer rules.

Across the texts we read, the same expectations recur. A person approves an action that is sensitive or cannot be undone, and the approval is enforced by the system. An agent gets the least access its task needs. People are told when they are dealing with AI. The content of communications, traffic data and location are confidential. Software that can change the network sits under the operator's security duties, its suppliers included. A record is kept.

What two regulators already say about agents

Singapore's Infocomm Media Development Authority publishes a Model AI Governance Framework for Agentic AI, now in version 1.5. It is voluntary, and it is written for every industry.

- Organisations should "define significant checkpoints or action boundaries that require human approval, especially before sensitive actions are executed".
- Its examples of actions that cannot be undone include "sending communications" and "making payments".
- Agents should work under "least-privilege policies" that give them "only the minimum tools and data access needed".
- Approval should be enforced "through system-level controls where possible", and not by guardrails in a prompt, "which may be bypassed".

Hong Kong's privacy regulator, in guidance of 25 August 2026, asks organisations to give an agent "the minimum access rights necessary to complete the tasks at hand", to keep a person in the loop for decisions likely to have a significant impact on individuals, and to build in traceability.

MARKET BY MARKET

Europe. The law here is written down, and some of it is specific to telecom. The ePrivacy Directive requires member states to "ensure the confidentiality of communications and the related traffic data". Traffic data must be erased or made anonymous when it is no longer needed to carry the communication, and location data may be processed only when it is made anonymous or with the user's consent. For an agent, that is a rule about what it may read. The data protection regulation gives a person "the right not to be subject to a decision based solely on automated processing" where the decision has a legal or similarly significant effect, and a right "to obtain human intervention". The network security directive known as NIS2 covers providers of public electronic communications networks: their management bodies approve the security measures and "can be held liable", the measures extend to "supply chain security", and a significant incident needs an early warning "within 24 hours".

The EU's AI Act. The rules for high-risk systems were moved in July 2026. They apply from 2 December 2027 for the uses the Act lists, and from 2 August 2028 for AI built into regulated products. The Act lists "safety components in the management and operation of critical digital infrastructure", and explains them as systems that "directly protect the physical integrity" of the infrastructure. On that wording an agent that drafts, checks or recommends is not one, though a lawyer should decide any single case. An agent that scores a person's credit is listed, and one used to detect fraud is not. A system that talks to people has to be built so that they know they are dealing with AI.

PART 05 · CONTINUED

The United Kingdom. A network provider "must take such measures as are appropriate and proportionate" against security compromises, which the law defines to include "any unauthorised access to, interference with or exploitation of the network" and "anything that compromises the confidentiality of signals". An agent with a tool that can change the network is inside that duty.

The Gulf, Asia, Australia and New Zealand. Singapore and Hong Kong are set out above. We could not open the telecom regulators' own texts on AI, on customer identity or on where data may be held for the Gulf states, Australia, New Zealand, or the rest of Asia, and an operator there should take them from its regulator. Data protection laws in many of these markets set conditions on sending personal data abroad, and we state none of them here.

So running agents on the operator's own infrastructure does not take away the duty to keep communications confidential, to secure the network, or to keep a person on a decision that matters. It takes away the questions about where the infrastructure is, who else shares it, and which country's law reaches the data.

And the controls in section 4 are the evidence these texts ask for: a list of agents and their tools, a named person who approves, the least access for each agent, a record of each decision, and a way to stop.

PART 06

The operator as a seller

An operator is the one kind of company in this series that can also sell what it has built. It owns data centres, a brand its market trusts and a sales force that already calls on businesses across the country, and those customers are asking where their AI may run. Several operators now sell AI capacity and agents from infrastructure in their own country.

Katonic's platform is built for that as well. The same platform an operator runs for its own agents can carry the operator's brand, with a console for creating customer organisations, keeping each one's data and models apart, allocating GPUs and billing. ePLDT, the ICT arm of the Philippines' largest telco, selected Katonic to power Pilipinas AI, a sovereign AI cloud hosted in the Philippines.

The two sides help each other. An operator that has run agents on its own care desk and in its own operations centre, under its own approval rules, sells from experience. This paper is about the first side. The second is described at katonic.ai/ai-cloud.

The platform an operator runs for itself is one it can put its own name on.

PART 07

Seven questions before an agent touches the network

These are worth asking of any agent platform, ours included.

01	**Is there one list of every agent in the company, of every tool each one can call, and of which tools change something?**
02	**Who approves before anything changes in the network or on a customer's account, and where is that written?**
03	**Is that approval enforced by the platform, or only asked for in a prompt?**
04	**Whose identity does the agent act under when it calls a system?**
05	**Which agents can read call records, traffic data or location, and why does each one need to?**
06	**Where does the model run, and does customer or network data leave the country on any request?**
07	**Can every call be shown to an auditor, with the decision that allowed it?**

PART 08

Where to start

Pick work where reading is slow, the decision stays with a person, and no customer is waiting on the agent.

THREE PLACES THAT FIT

Runbook questions

Engineers ask what the operator's own runbook says and get the answer with the section cited, and "the runbook is silent" where it is. No customer data is involved, so it can run on day one.

Post-incident review

The agent sets out the timeline of an incident from the tickets and the alarms and drafts the review in the operator's own format. The incident manager owns the findings.

Change request review

The agent checks each change request against the operator's own change rules and lists what is missing or in conflict. The change board approves, and an engineer makes the change.

PART 08 · CONTINUED

Then, in order

01 Name the documents and the systems the agent needs, and which of the six ways each system is reached.

02 Decide the policy for every tool, and who approves each kind of change, before the agent is built.

03 Build it with the engineers and staff who do the work today.

04 Run it on invented data, and then beside the manual process.

05 Measure it against that process: time, and how often a person corrected it.

06 Go live for one team and one domain.

A fuller list of 32 agents across nine areas is published with this paper.

About Katonic AI

Katonic AI is the enterprise platform for serving models and building, running and governing AI agents. It runs on your machines, with your models, under your administrators. Katonic is listed in the Red Hat Ecosystem Catalog and runs on Red Hat OpenShift. ePLDT, the ICT arm of the Philippines' largest telco, selected Katonic to power Pilipinas AI, a sovereign AI cloud hosted in the Philippines. The platform is in production at a top-tier bank in Asia, air-gapped, inside the bank's own perimeter, and it has run in production for about 1,100 staff of a life insurer in Asia, on the insurer's own cluster, inside its network. It was founded in Sydney in 2020 by Prem Naraindas, with its global hub in Dubai and teams in Sydney, Mumbai and Dubai.

KATONIC.AI

SOURCES

All read on 21 November 2026.

01 True Corporation, article on its assistant Mari. true.th

02 Telstra, "Telstra's AI transformation", 20 April 2026. telstra.com.au

03 Telkomsel, news release, 20 March 2026. telkomsel.com

04 NTT DOCOMO, press release, 25 February 2026. docomo.ne.jp

05 TelecomTV, interview with Telia, recorded October 2026, and article by AT&T's chief data officer, 14 November 2025.

06 Deutsche Telekom, media information, 11 November 2025. telekom.com

07 RCR Wireless News, report on Deutsche Telekom, 21 October 2026.

08 Rakuten Mobile, news release, 19 February 2026.

09 SoftBank Corp., press release, 29 October 2025. softbank.jp

10 SK Telecom, newsroom article, 26 August 2026. news.sktelecom.com

11 Telenor, announcement, 11 June 2025. telenor.com

12 Docker MCP registry, 327 servers, read in full on 19 November 2026.

13 TM Forum, Open APIs, and its public repositories of API definitions. tmforum.org

14 CAMARA project, the Linux Foundation, API definitions and the Security and Interoperability Profile.

15 Infocomm Media Development Authority, Singapore, Model AI Governance Framework for Agentic AI, version 1.5.

16 Privacy Commissioner for Personal Data, Hong Kong, media statement of 25 August 2026.

17 Directive 2002/58/EC on privacy and electronic communications, Articles 5, 6 and 9.

18 Regulation (EU) 2016/679, the General Data Protection Regulation, Article 22.

19 Directive (EU) 2022/2555, NIS2, Articles 20, 21 and 23.

20 Regulation (EU) 2024/1689, the AI Act, Article 50, Annex III and recital 55, and Regulation (EU) 2026/1744 of 8 July 2026.

21 Communications Act 2003, United Kingdom, section 105A. legislation.gov.uk