# Security agents on Katonic: the list

## 35 agents across nine areas, for the security operations centre and the security team around it

A security department's work runs through its alert and log platform, its case records, its vulnerability scanner, its directory, its cloud consoles, its threat sources, and its own playbooks and policies. Every agent reaches them in one of four ways, all through the same tool gateway and the same approval rules:

| | HOW EACH AGENT CONNECTS |
|---|---|
| 01 | **A file the user uploads** in the conversation: an alert export, a reported email saved as a file, a scan report, a threat report, a customer's questionnaire. |
| 02 | **A knowledge source.** Playbooks, security policies, the asset list, past incident reviews, detection rules and approved questionnaire answers are indexed in the Knowledge Engine, with each reader's permissions kept. |
| 03 | **An imported API.** The organisation gives Katonic the description of an API it already runs, and the operations it chooses become tools. This is the route to systems of record: the alert platform, the case system, the scanner, the directory, a public vulnerability catalogue. |
| 04 | **An MCP server from the catalogue,** deployed inside the organisation's own cluster. This is the route to tickets, on-call alerts, metrics and logs, and chat channels. |

System names are left out on purpose. The organisation supplies the address and the sign-in. Nothing here is a tested integration with a named product.

THE NINE AREAS


| | | |
|---|---|---|
| 01 | Alerts | 2 |
| 02 | Reported emails | 2 |
| 03 | Incidents | 3 |
| 04 | Vulnerabilities | 3 |
| 05 | Identity | 4 |
| 06 | Cloud and configuration | 4 |
| 07 | Threat intelligence and detection | 5 |
| 08 | Assurance | 5 |
| 09 | Awareness and reporting | 6 |

AREA 01

## Alerts

| AGENT | WHAT IT DOES | READS FROM | WHERE IT STOPS FOR A PERSON |
|---|---|---|---|
| **Alert Triage Summary Agent** | For one alert, gathers what the playbook asks for: the asset, its owner, the account, earlier alerts on both, and proposes the severity the playbook gives, with the record behind each line | The alert through an imported API or an MCP server; the playbook and the asset list as a knowledge source | An analyst confirms the severity and decides the next step. The agent closes nothing and contains nothing |
| **Indicator Lookup Agent** | Looks up the addresses, domains and file hashes in an alert in the threat sources the organisation uses, and reports what each source says and when | The alert; threat sources through imported APIs | The analyst judges what the results mean |
| **Related Alert Check Agent** | Lists other alerts in a period that share an asset, an account or an indicator with this one | The alert platform through an imported API, read operations only | The analyst decides whether they are one case |
| **Shift Handover Summary Agent** | Prepares the handover between shifts: open cases, what was done on each, and what is waiting on someone | The case system and the on-call channel | The outgoing lead corrects it and hands over |

AREA 02

## Reported emails

| AGENT | WHAT IT DOES | READS FROM | WHERE IT STOPS FOR A PERSON |
|---|---|---|---|
| **Reported Email Triage Agent** | Reads an email a colleague reported, as text: the headers, the sender, the links as written and the attachment names, and lists what the playbook treats as signs of fraud. It opens no link and no attachment | The reported email as a file or through an imported API; the playbook as a knowledge source | An analyst decides what the email is and what to do. The agent deletes nothing and blocks nothing |
| **Similar Email List Agent** | Lists other mailboxes that received an email with the same sender, subject or link in a period | The mail system through the organisation's own API, read operations only | The analyst decides whether to remove them, and an administrator does |
| **Reporter Reply Draft Agent** | Drafts the reply to the colleague who reported an email, from the analyst's decision and the organisation's templates | The analyst's decision; the templates | The analyst approves and sends |
| **Reported Email Summary Agent** | Summarises the emails reported in a period: how many, of what kind as the analysts classed them, and which departments were targeted | The case system through a reporting database | The security lead reviews the summary |

AREA 03

## Incidents

| AGENT | WHAT IT DOES | READS FROM | WHERE IT STOPS FOR A PERSON |
|---|---|---|---|
| **Incident Response Agent** | Reads an alert, pulls the related metrics and logs, and drafts an incident summary with the runbook step that fits | The on-call alert, metrics and logs through MCP servers; runbooks as a knowledge source | An engineer carries out any runbook step |
| **Incident Timeline Agent** | Builds the timeline of an incident from the case record, the channel and the logs the analyst names, with the source of each entry | The case record, the channel export and log extracts | The incident lead corrects and owns the timeline |
| **Incident Record Draft Agent** | Drafts the incident record in the organisation's format: what was seen, what was done and by whom, what is known and not known | The case record and the timeline; the record template | The incident lead completes and signs the record |
| **Post-Incident Review Draft Agent** | Drafts the review after an incident: what happened, what worked, what did not, and proposed actions, each tied to its evidence | The incident record, the timeline and past reviews | The review owner and the team agree the causes and the actions. The agent names no person as the cause |

AREA 04

## Vulnerabilities

| AGENT | WHAT IT DOES | READS FROM | WHERE IT STOPS FOR A PERSON |
|---|---|---|---|
| **Vulnerability Prioritisation Preparation Agent** | For a scan result, sets beside each finding what the organisation's policy weighs: how critical the asset is, whether it faces the internet, and whether public catalogues record the flaw as exploited, and proposes the order the policy gives | The scan result as a file or through an imported API; the asset list and the policy as a knowledge source; public catalogues through imported APIs | The vulnerability manager sets the order. System owners schedule the fixes through the change process |
| **Advisory Impact Agent** | For a new security advisory, lists the organisation's systems that run the affected product and version | The advisory as a file or through an imported API; the asset and software inventory | The system owners schedule the fix |
| **Patch Exception Check Agent** | Lists systems past their patch deadline and the exceptions on record, and marks exceptions that have expired | The patch report; the exception register | The vulnerability manager follows up with each owner |
| **Scan Result Summary Agent** | Summarises a scan for the people who must act: findings by owner and by age, what is new since the last scan and what was fixed | Two scan results as files or through an imported API | The vulnerability manager issues the summary |

AREA 05

## Identity

| AGENT | WHAT IT DOES | READS FROM | WHERE IT STOPS FOR A PERSON |
|---|---|---|---|
| **Sign-in Investigation Summary Agent** | For an account an analyst names in an open case, lists its sign-ins, the places and devices they came from, and the changes made to the account in the period | The directory's sign-in and audit records through the organisation's own API, read operations only | The analyst draws the conclusions and decides. The agent disables nothing and draws no conclusion about the person |
| **Privileged Access Review Preparation Agent** | Lists who holds administrative rights, marks the ones the role rules do not explain, and prepares the sheet for the reviewer | The directory through the organisation's own API; role rules as a knowledge source | The system owner keeps or removes each right |
| **Dormant Account List Agent** | Lists accounts with no sign-in for longer than the organisation's own period, and service accounts with no named owner | The directory through the organisation's own API | The identity team decides what to disable |
| **Access Change Summary Agent** | Summarises the changes to privileged groups and roles in a period, with who made each change | The directory's audit records | The security lead reviews the summary |

AREA 06

## Cloud and configuration

| AGENT | WHAT IT DOES | READS FROM | WHERE IT STOPS FOR A PERSON |
|---|---|---|---|
| **Configuration Finding Summary Agent** | Groups the configuration findings of a cloud account by owner and by kind, with the resource behind each | The findings as a file or through an imported API | The resource owners fix their findings through the change process |
| **Baseline Comparison Agent** | Compares a system's configuration with the organisation's own baseline and lists each setting that departs from it | The configuration as a file; the baseline as a knowledge source | The system owner decides what to change or to record as an exception |
| **Exposed Resource List Agent** | Lists the resources that can be reached from the internet, with the owner and the reason on record for each | The cloud inventory through an imported API; the register of approved exposures | The cloud security lead reviews the ones with no reason on record |

AREA 07

## Threat intelligence and detection

| AGENT | WHAT IT DOES | READS FROM | WHERE IT STOPS FOR A PERSON |
|---|---|---|---|
| **Threat Report Summary Agent** | Summarises a threat report and lists which of the products and sectors it names appear in the organisation's own estate | The report as an uploaded file; the asset and software inventory | The threat intelligence analyst judges what applies |
| **Detection Coverage Check Agent** | For the techniques a report names, lists which of the organisation's detection rules claim to cover each, and the techniques no rule names | The report; the detection rules as a knowledge source | The detection engineer decides what to build |
| **Detection Rule Draft Agent** | Drafts a detection rule in the organisation's rule language from an engineer's description, as a file, with the test cases it should pass | The description; existing rules as a knowledge source | The detection engineer tests, corrects and deploys the rule. The agent changes no rule |
| **Rule Tuning Report Agent** | Lists the detection rules that raised the most alerts closed as harmless in a period, with examples | The alert platform through a reporting database | The detection engineer decides what to tune |

AREA 08

## Assurance

| AGENT | WHAT IT DOES | READS FROM | WHERE IT STOPS FOR A PERSON |
|---|---|---|---|
| **Security Questionnaire Answer Draft Agent** | Drafts the organisation's answers to a customer's security questionnaire from its own approved answers and policies, and marks each question with no approved answer | The questionnaire as an uploaded file; the approved answers and policies as a knowledge source | The security lead approves and sends the answers |
| **Test Finding Follow-up Agent** | Lists the findings of a penetration test that are overdue, or marked fixed without a retest on file | The test report and the findings register | The security lead follows up with each owner |
| **Security Exception Register Check Agent** | Lists security exceptions that have expired, have no owner, or have no compensating control on record | The exception register through an imported API or as a file | The security lead reviews each exception |
| **Security Audit Evidence Agent** | For a security control under audit, gathers the evidence the control asks for and lists what is missing | The control list as a knowledge source; tickets, logs and reviews | The control owner checks the evidence and hands it to the auditor |

AREA 09

# Awareness and reporting

| AGENT | WHAT IT DOES | READS FROM | WHERE IT STOPS FOR A PERSON |
|---|---|---|---|
| **Knowledge Agent** | Answers from the organisation's own document stores with the passage cited, and keeps each reader's permissions | Security policies and playbooks as a knowledge source | It informs. The reader checks the source |
| **Awareness Note Draft Agent** | Drafts a short awareness note for staff from a recent case the security team chooses, with names and details removed as the team directs | The case summary the team supplies; past notes | The security lead approves and issues the note |
| **Security Report Agent** | Prepares the security report for a period: cases by kind and severity, time to respond, open vulnerabilities by age, with the records behind each figure | The case system and the scanner through a reporting database | The head of security reviews and issues the report |
| **Notification Deadline Check Agent** | For an incident of a given kind, lists the notifications the organisation's own procedure names, to whom and within what time | The incident record; the notification procedure as a knowledge source | The compliance officer decides whether the incident must be notified. The agent decides nothing and notifies nobody |

## WHAT EVERY AGENT HAS IN COMMON

01 **It runs on the organisation's own infrastructure or private cloud,** with a model the organisation serves. Logs, alerts and incident records stay there.

02 **Every tool call passes the organisation's policy.** A policy can allow a tool, record its use, hold it for a person's approval, or block it, for the whole organisation or for one team.

03 **A tool that changes something waits for approval** unless the organisation has decided otherwise for that tool. That is how no host is isolated, no account disabled, no address blocked and no alert closed on an agent's word.

04 **Read and change are separate tools.** When the organisation imports an API, it chooses which operations become tools. An agent that only needs to read is given the read operations and nothing else.

05 **Each statement shows its source.** An agent that summarises or checks gives the alert, the log line or the record, and says when it found nothing. It does not fill a gap from memory.

06 **The agent acts as the person using it.** A knowledge source keeps each reader's permissions, so a colleague does not see through an agent a case they could not open themselves.

07 **Every call is recorded,** with who asked, which tool, and what was decided, and an agent, a tool or a whole server can be switched off at once.

WHERE TO START

## Pick work where reading is slow, the decision stays with an analyst, and nothing is contained or closed on the agent's word.

THREE PLACES THAT FIT

### Knowledge Agent on the playbooks

Analysts ask what the playbook says and get the answer with the passage cited. No system is touched, so it can run on day one.

### Alert Triage Summary

The agent gathers what the playbook asks for around one alert. An analyst decides.

### Vulnerability Prioritisation Preparation

The agent sets the policy's factors beside each finding. The vulnerability manager sets the order.

One team, one kind of alert, invented data first. Then beside the manual process, measured against it.

# Each one prepares the work. A named analyst decides.

**About Katonic AI**

Katonic AI builds the platform that organisations use to build, run and govern AI agents on their own infrastructure: in their own data centre, in a private or sovereign cloud, or with no connection to the internet at all. The platform is in production with governments, banks, telecom operators and utilities. Katonic is certified to ISO 27001.

KATONIC.AI