

AI OPERATIONS LAYER FOR CHECKMK

A PRODUCT BY LYNXMIND

Pardinus AI Suite.

An AI operations layer embedded inside Checkmk. **Chat** talks to your infrastructure, **RCA** turns alerts into structured root-cause reports, and every token is metered — one plugin, your data, your premises.

CHAT · ROOT CAUSE ANALYSIS · COST INTELLIGENCE · TOKEN METERING

LYNXMIND

ENGINEERING, DELIVERED

Alerts are loud. Root causes are quiet.

Signal buried under acknowledgements, downtimes and copy-pasted runbooks.

Operations teams running Checkmk handle thousands of state changes a week. The answer to "why did this break?" still lives inside someone's head — not in the tooling.

73%

OF INCIDENTS

Diagnosed by tribal knowledge, not tooling — average across mid-market enterprise ops teams.

4.2h

MEAN TIME TO RESOLVE

Spent on triage and evidence-gathering before anyone proposes a fix.

1 in 6

ALERTS REPEAT

Because the previous resolution wasn't documented where the next operator can find it.

Two products. One AI operations layer.

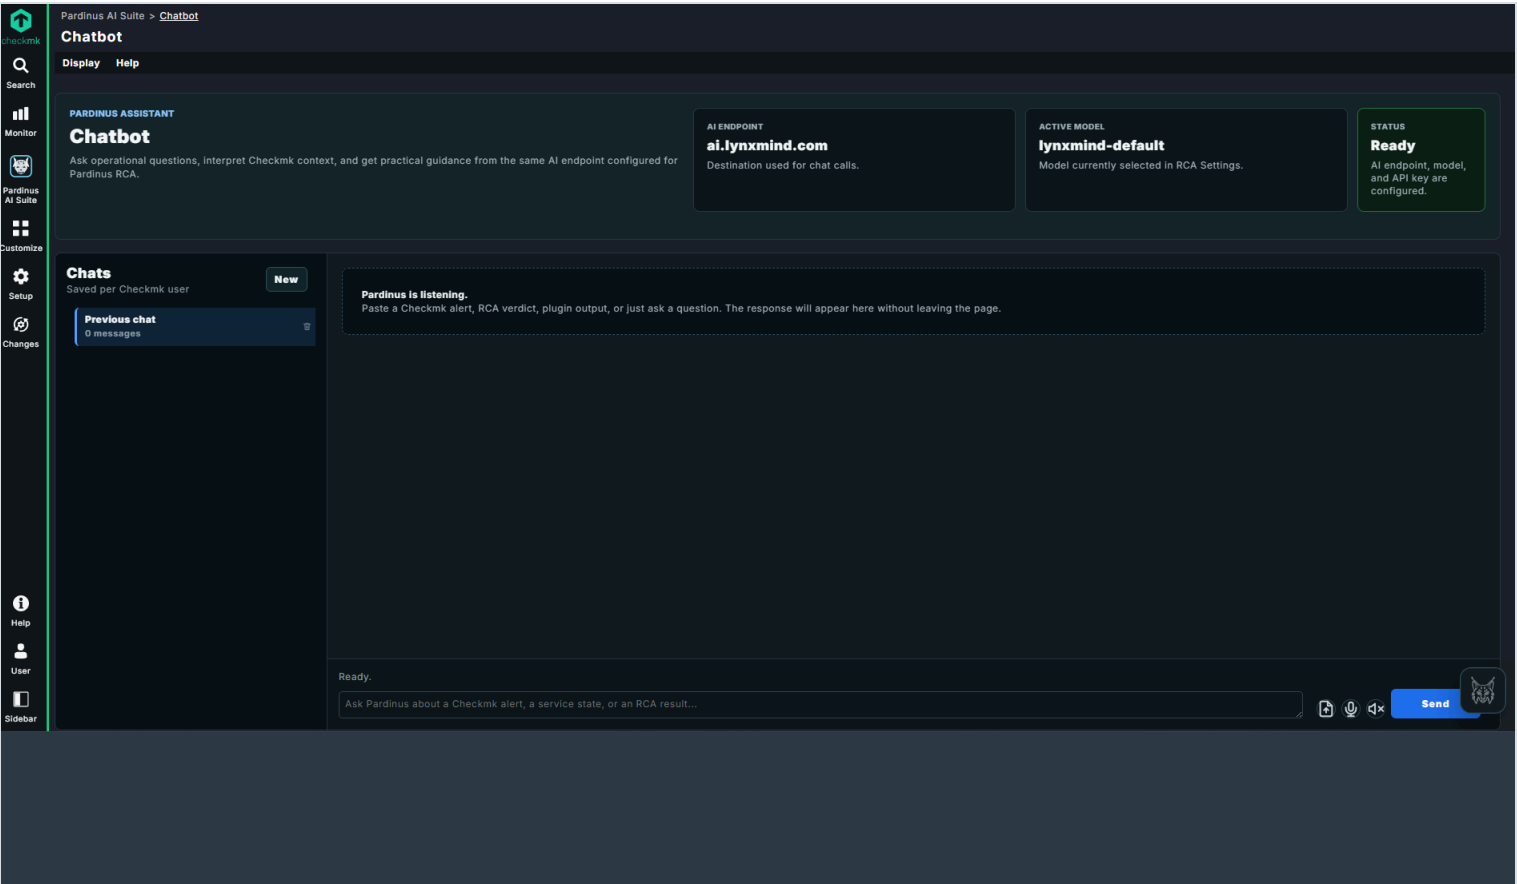
One shared private AI stack,
one audit trail, one MKP plugin.

PRODUCT 01

CONVERSATIONAL

Pardinus Chat

An AI operations assistant embedded in the Checkmk sidebar. Speaks English & Portuguese, across three modes: **Task** (act on the REST API), **Report** (investigate & synthesise), **Wiki** (learn the platform).



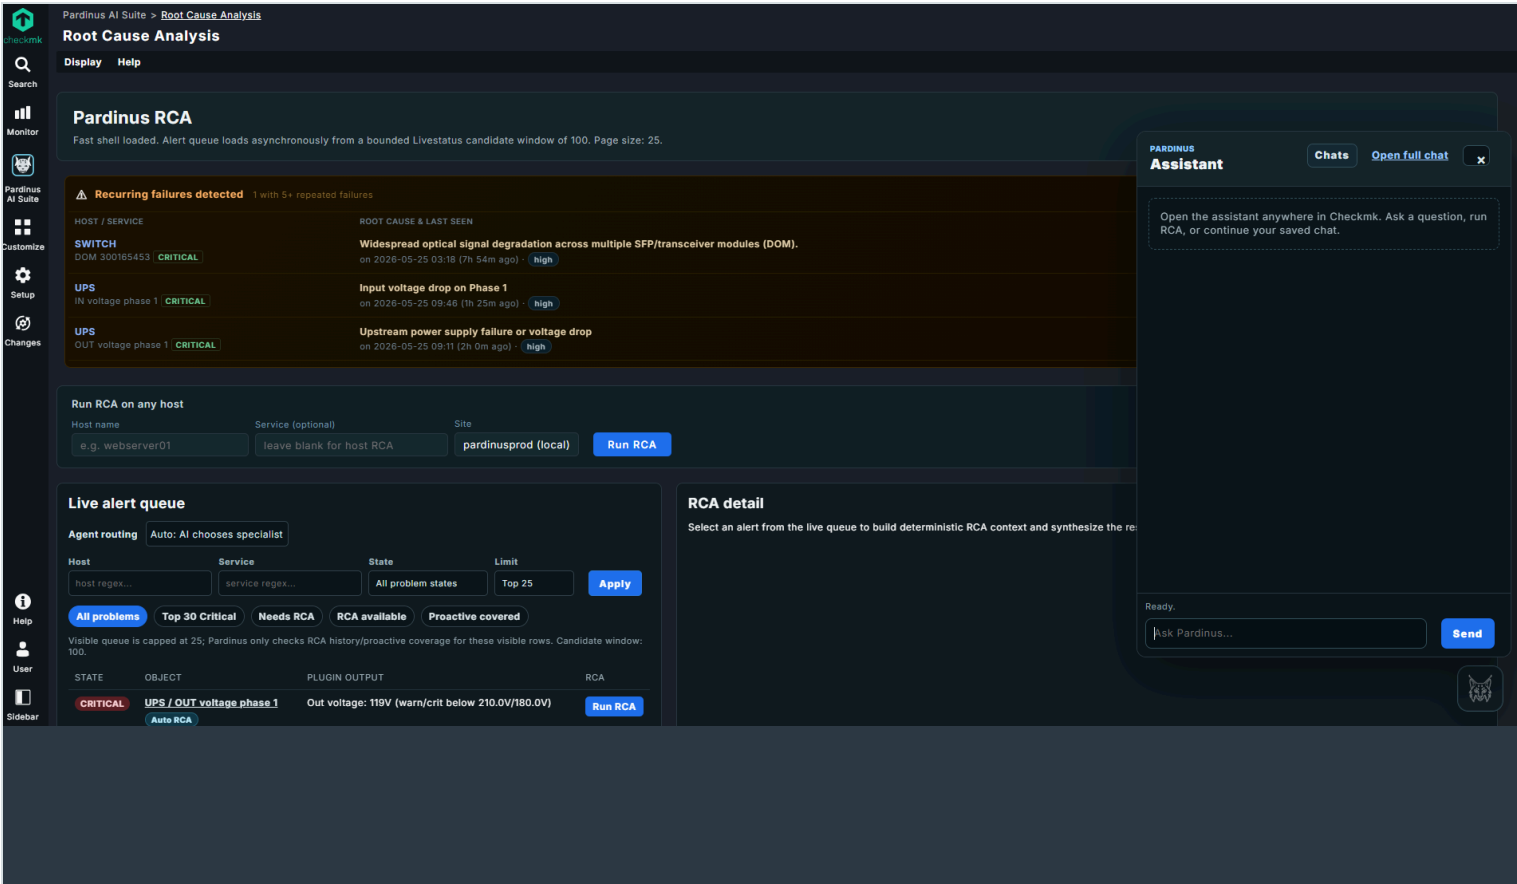
- ◆ Voice input + TTS replies
- ◆ Confirms destructive ops
- ◆ Per-user chat history
- ◆ Role-based access

PRODUCT 02

DETERMINISTIC

Pardinus RCA

A dedicated module that runs root-cause analysis on any host or service. Evidence is built deterministically from Liveness & the REST API; the LLM only explains it. Structured, auditable JSON output.




- ◆ One-click "Run RCA"
- ◆ Confidence + alternatives
- ◆ Specialist agents per domain
- ◆ Auditable JSON post-mortems

DETERMINISTIC-FIRST PIPELINE

Evidence is built by code.
The LLM only explains it.

No RCA decision from raw alerts —
every package is assembled before any token leaves the site.



STRICT RULE If evidence is missing, **missing_data** is reported — the model is forbidden from inventing it.


OUTPUT root_cause · confidence · evidence_used · timeline · recommended_actions · alternatives

MULTI-HOST REASONING

200 alerts. One root. Found algorithmically.

Storm vs cascade, told apart
before the LLM writes a word.

When failure cascades across dependent hosts, Pardinus walks the dependency chain, isolates the originating host with a quantified confidence score, and classifies every event by its role.



EVERY EVENT IS CLASSIFIED

- **root_cause** — the originating failure
- **symptom** — downstream consequence
- **precursor** — early warning sign
- **context** — supporting evidence

STORM VS CASCADE

Pardinus tells coincidental simultaneous incidents apart from a single cascade — and runs RCA only on the **root candidate**, with confidence scored **50–90%**.

MODULAR BY DESIGN

Clean separation. Replaceable parts.

UI	GUI extension Sidebar "AI Operations → RCA Analysis". Host/service action "Run RCA". MKP-packaged.	CHECKMK PLUGINS
API	Backend job controller Accepts host_name, service, site, trigger. Job model: queued · running · completed · failed.	REST
RCA	Deterministic engine Context builder, timeline builder, evidence ranker. Pure Python. Fully tested.	CORE
DATA	Checkmk integration Livestatus client (Unix socket) + versioned REST client. No core-table writes.	READ-ONLY
LLM	Provider abstraction Local · OpenAI-compatible · Azure OpenAI · customer-owned. Swappable per-tenant.	PLUGGABLE
DB	Persistence PostgreSQL preferred, in-memory fallback. Jobs and results outside Checkmk core tables.	ISOLATED

Read-only against Checkmk.
Nothing written to core tables.

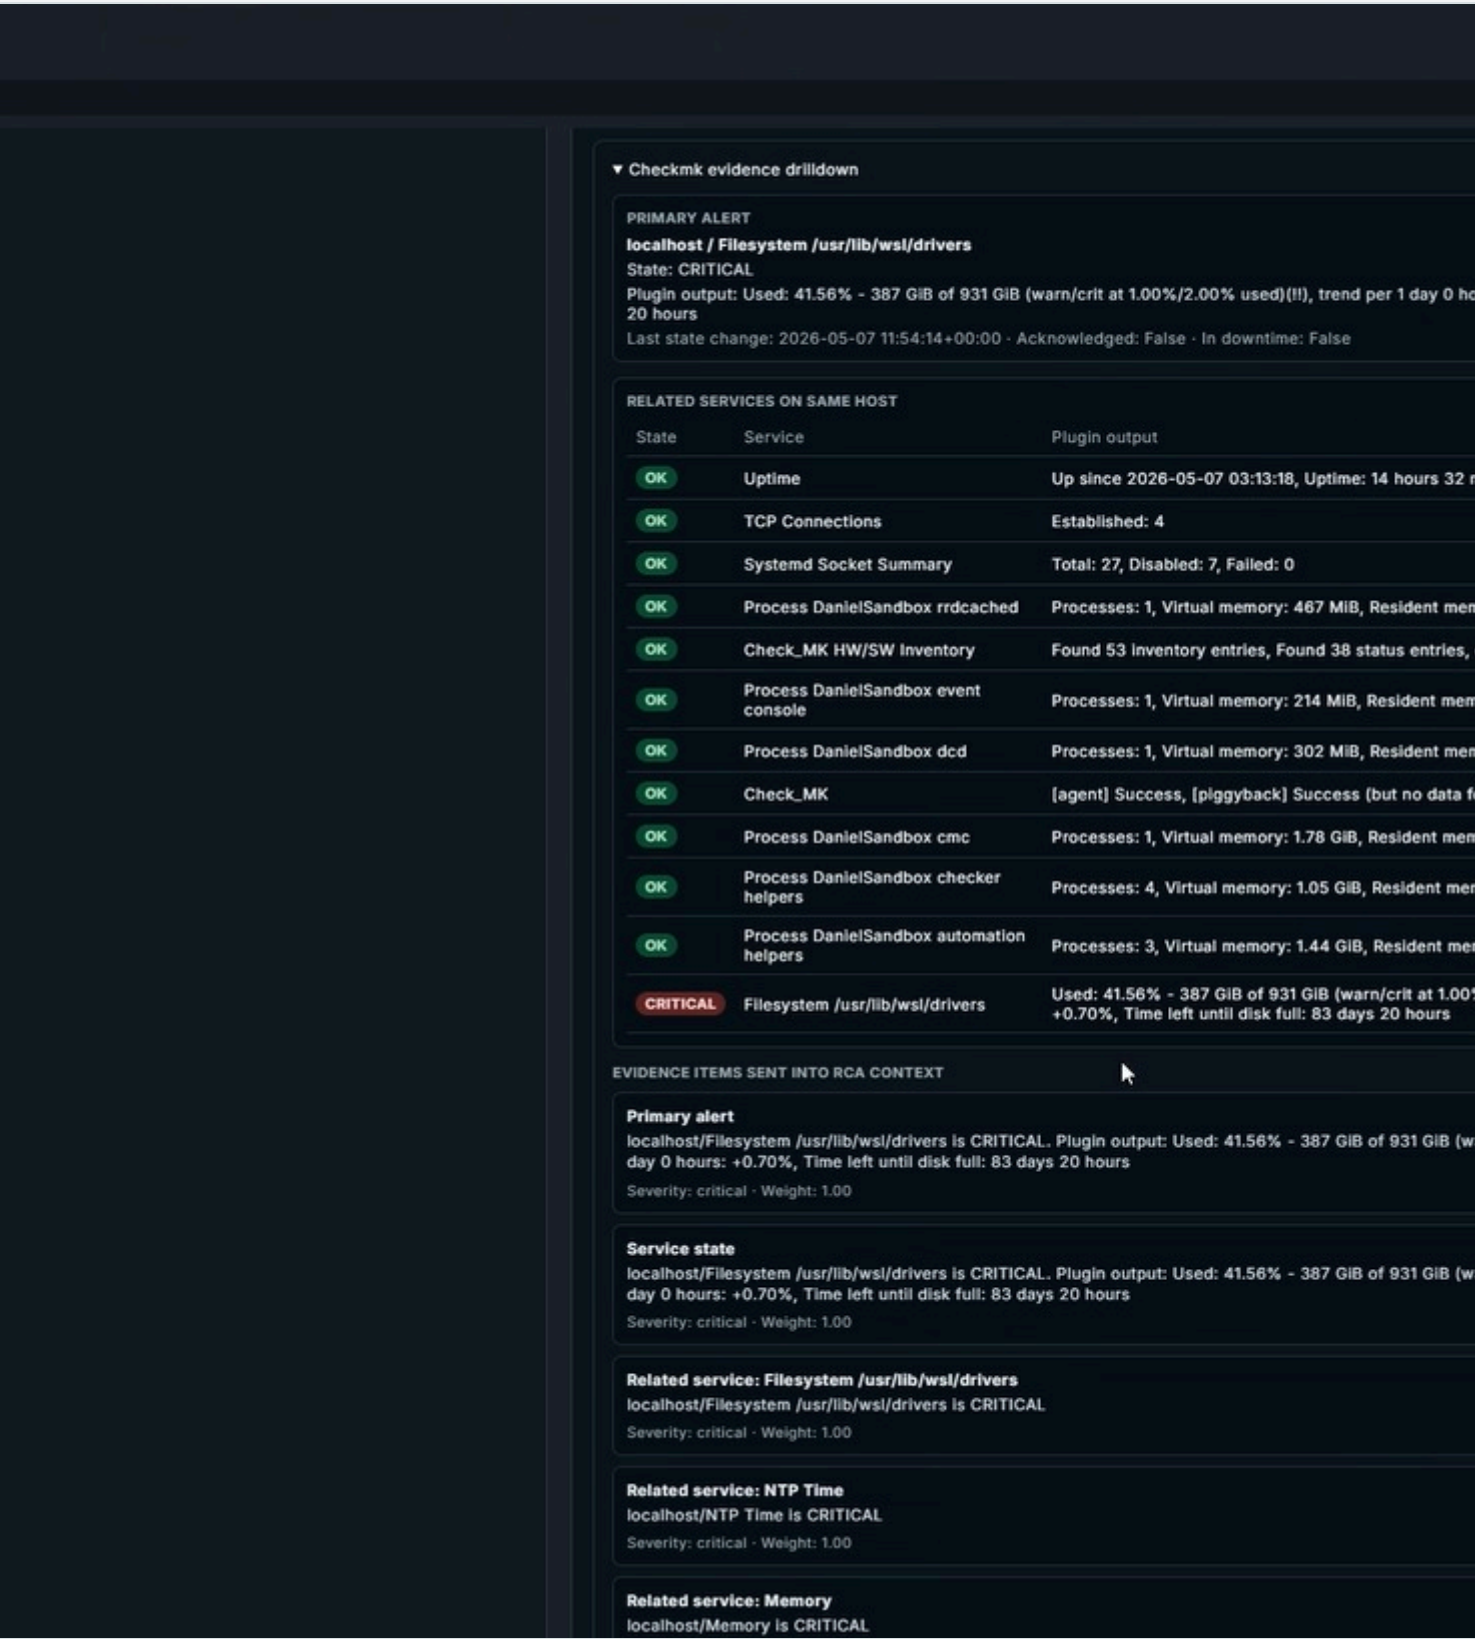
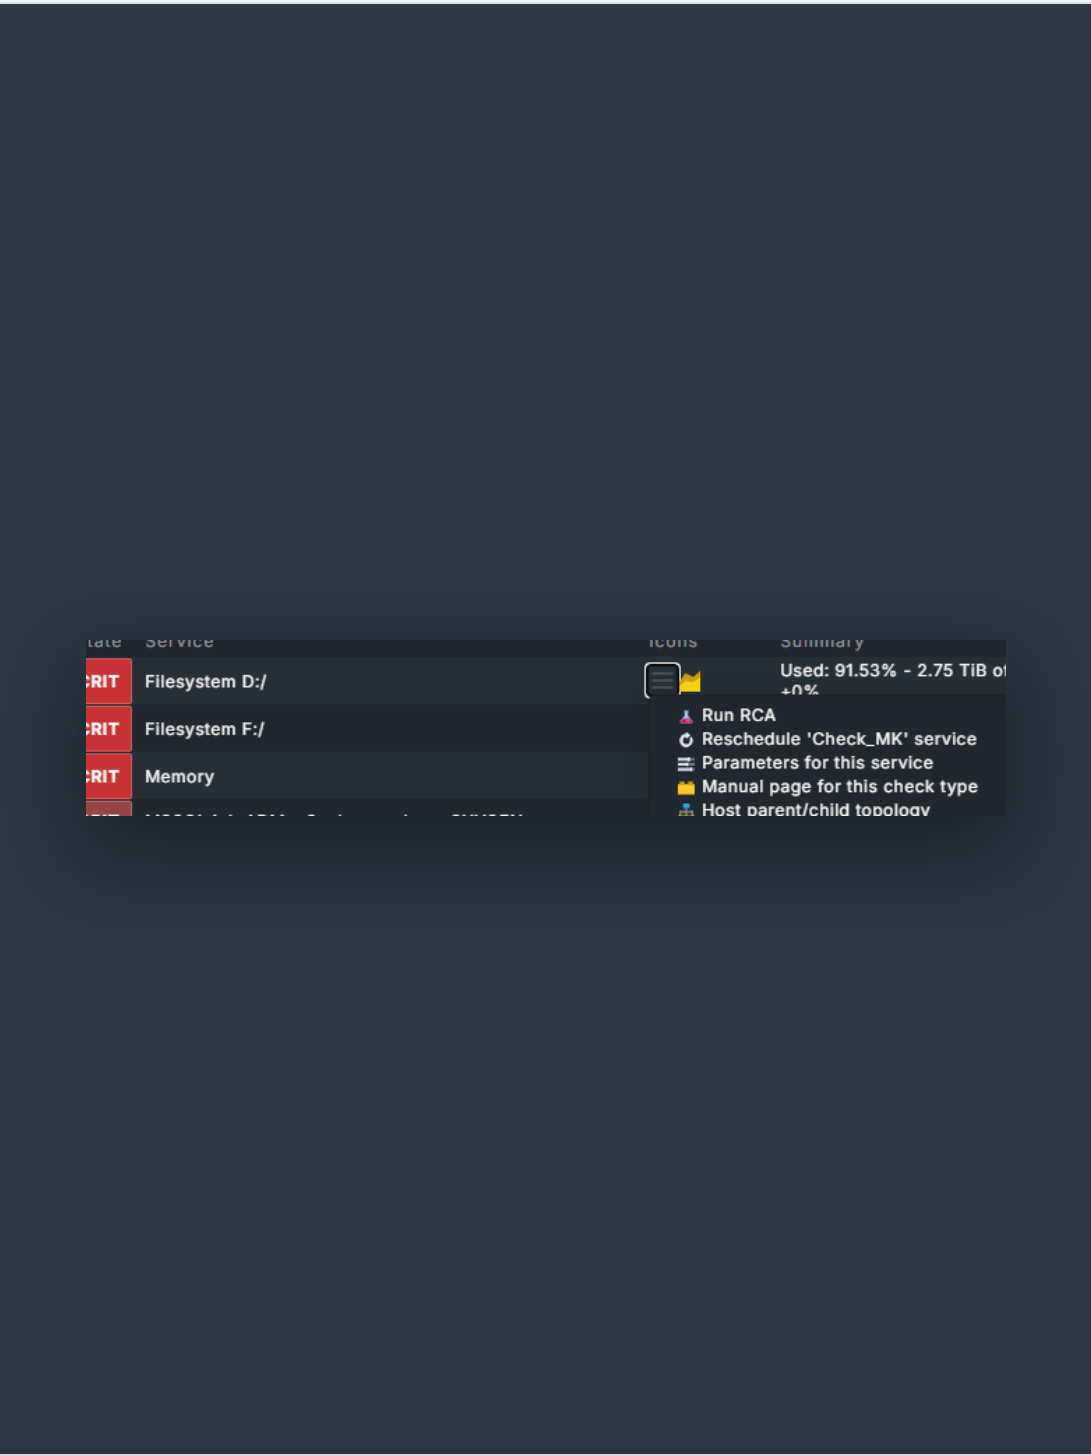


Fig. 05 — Evidence package as collected from Checkmk Livestatus

Three surfaces, one workflow.

01 · TRIGGER

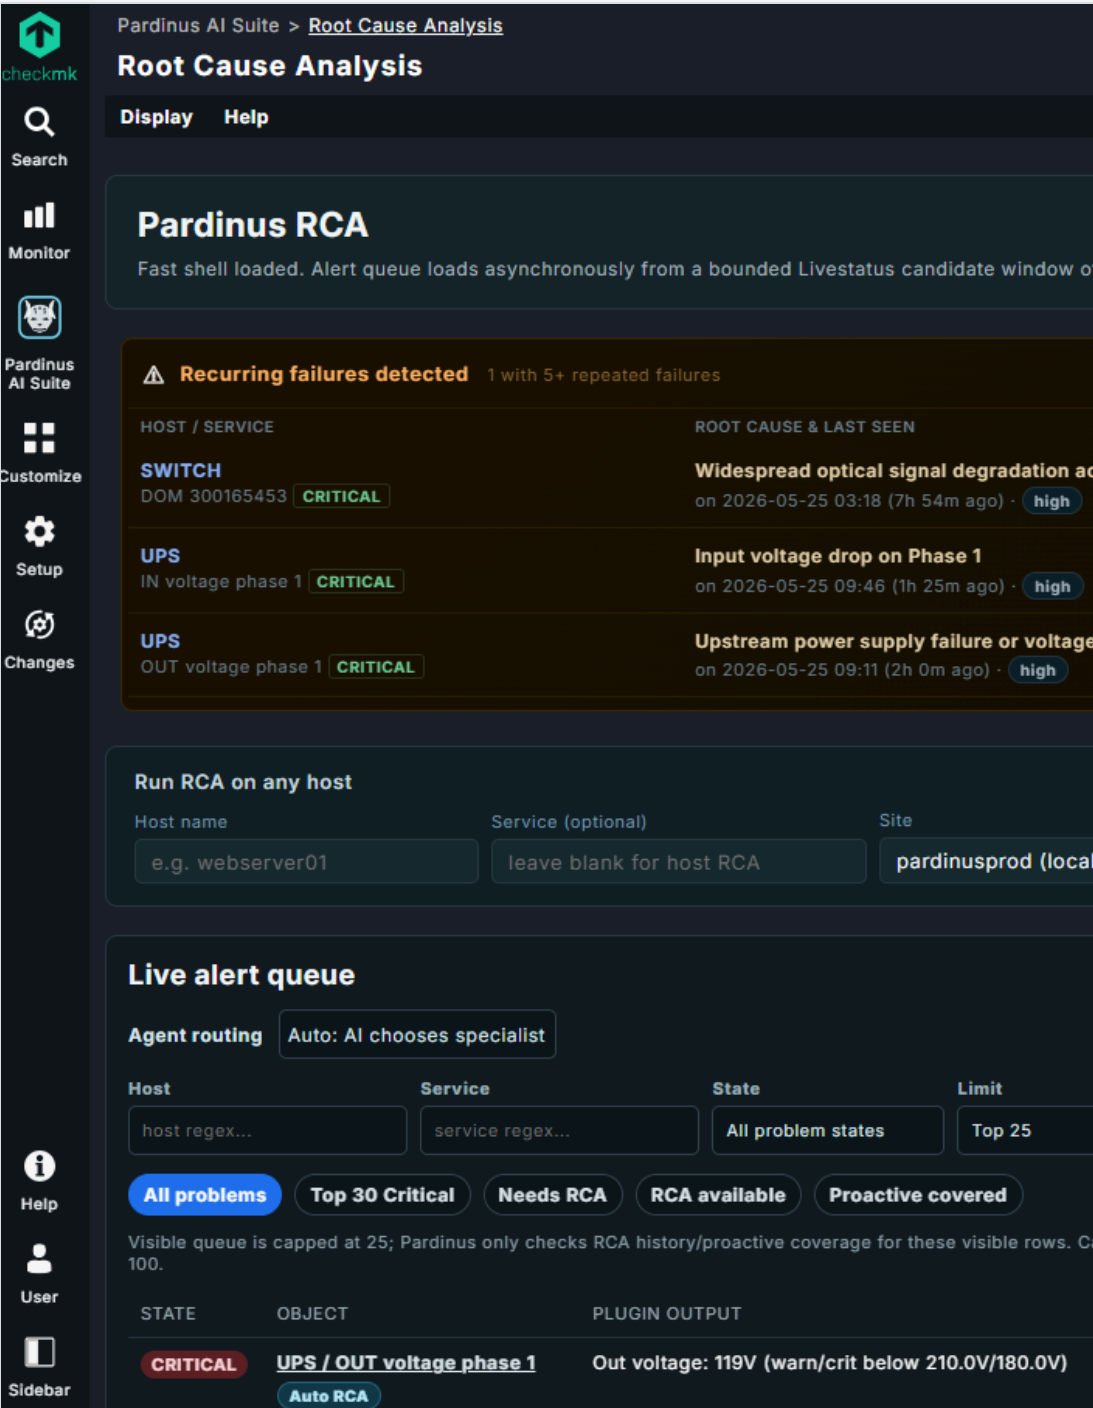
Action menu on any service



One click on a host or service opens the RCA job page.

02 · JOB

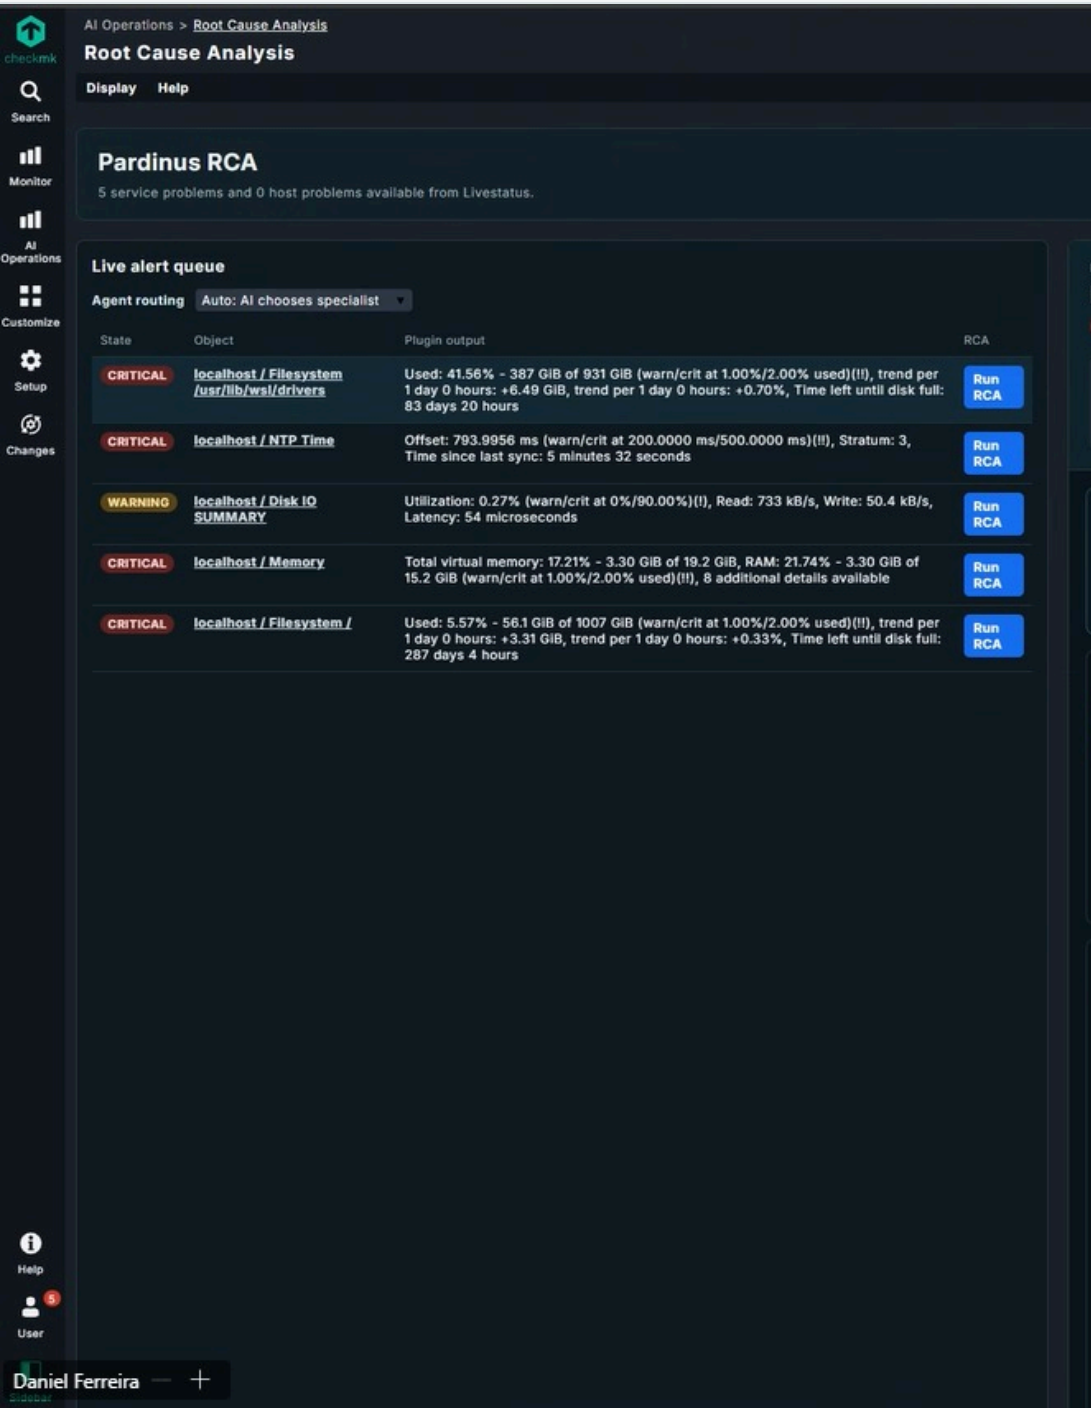
Live alert queue



Queued, running, completed, failed. Recurring failures clustered automatically.

03 · RESULT

Structured RCA report



Root cause, confidence, timeline, evidence, recommended actions, raw JSON.

No new tooling to learn – it lives where operators already work.

ENTERPRISE-GRADE CONTROLS

Your data. Your rules.

Configure the endpoint, redact before it leaves, and cap what the LLM ever sees.

Pardinus RCA ships with first-class data-governance controls — configure your AI endpoint, redact sensitive values before they leave the site, and set hard limits on what the model receives.

GUARDRAIL

Redact sensitive values

IP addresses, tokens, secrets stripped pre-flight.

GUARDRAIL

Deterministic fallback

If AI fails, evidence-only output.

GUARDRAIL

Provider test first

Endpoint must return valid JSON before RCA.

GUARDRAIL

Confidence threshold

Below 0.75? The operator is warned.

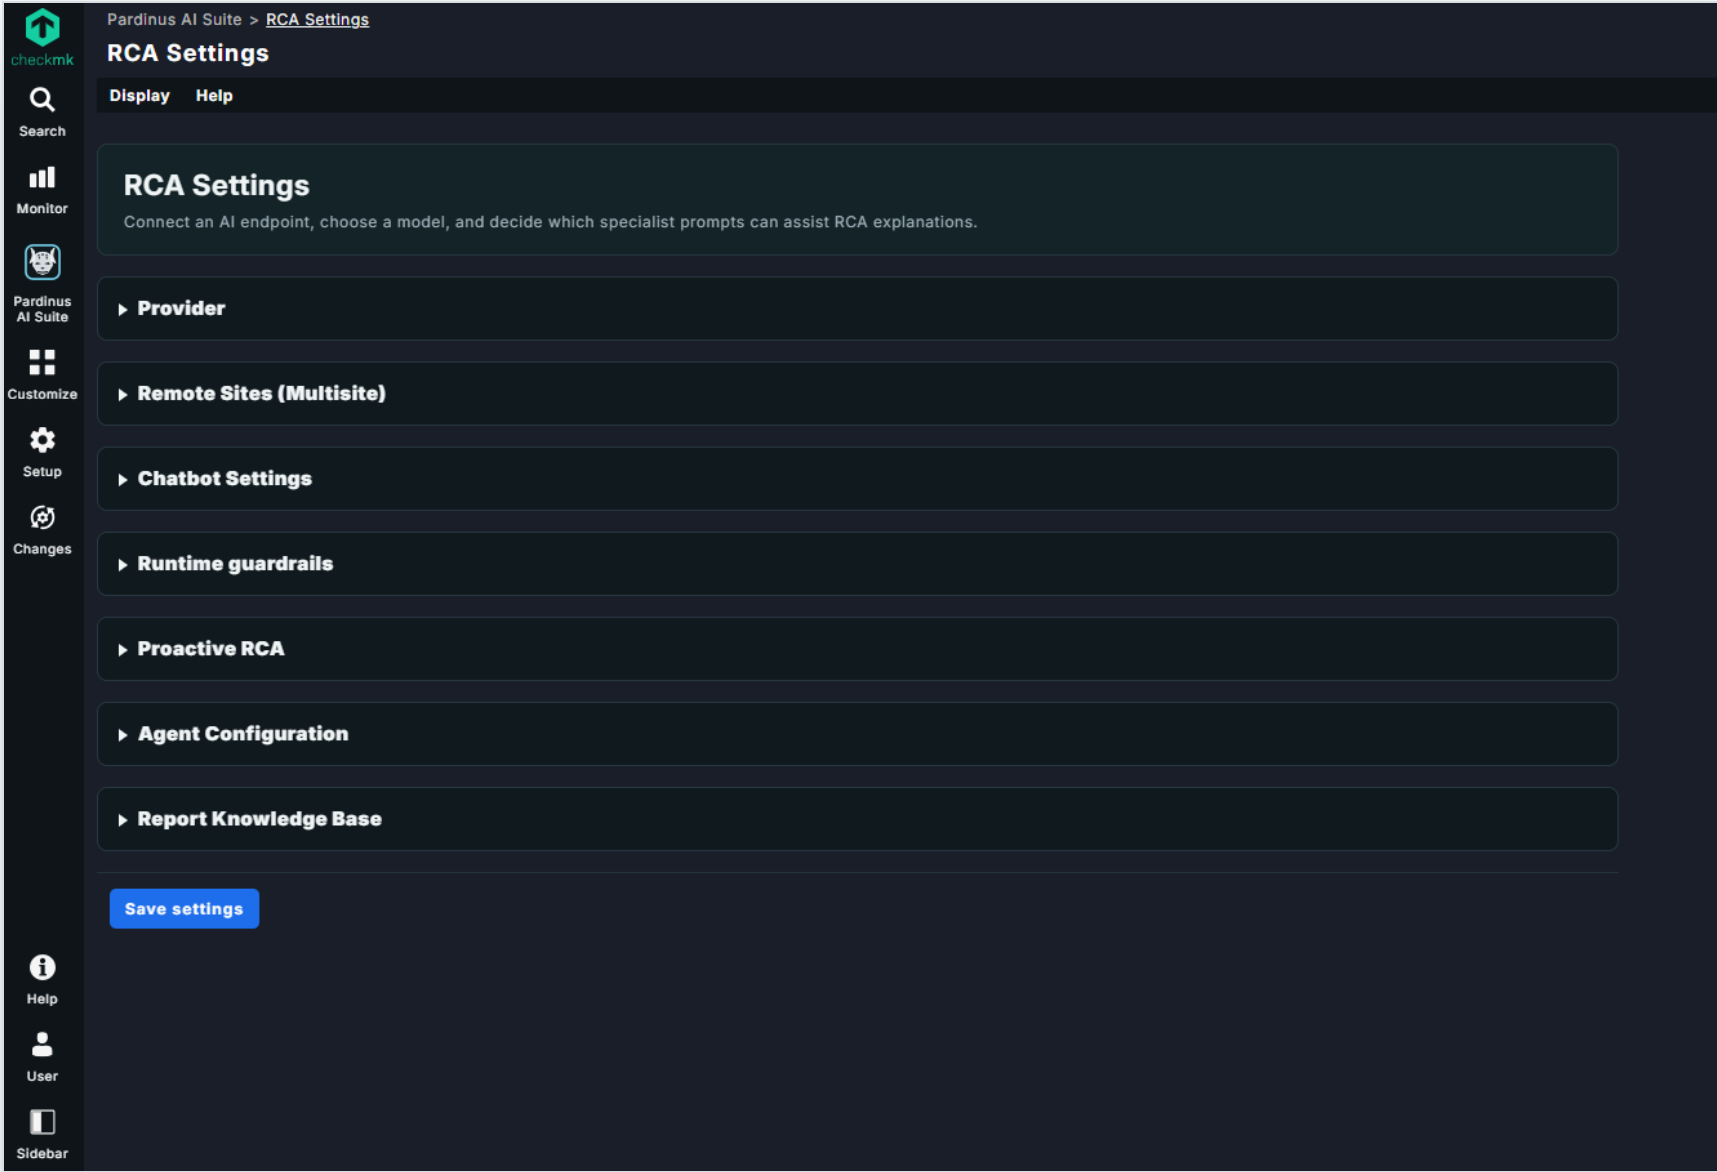


Fig. 07 — RCA Settings panel inside Checkmk

TALK TO YOUR INFRASTRUCTURE

Three modes of intelligence.

TASK Act

Calls the Checkmk REST API on the user's behalf. Confirms destructive ops. Creates hosts, moves folders, schedules downtimes, manages users.

"Create a downtime for all hosts in Lisboa for the next 2 hours."

REPORT Understand

Autonomously investigates: host status, service health, performance metrics. Synthesises a structured answer in plain language.

"How are my Linux machines performing today?"

WIKI Learn

Answers Checkmk concept & configuration questions with session context. New operators ramp without reading the manual.

"What is a host tag group?"

ALSO
Voice input · per-user history · dark mode · multilingual (EN/PT) · tamper-evident audit trace on every interaction.

English & Portuguese. No API docs, no command syntax — just conversation.




Fig. 08 — Pardinus Chatbot · dedicated page inside Checkmk

RIGHT EXPERT FOR THE RIGHT ALERT

7 specialists. Auto-routed.

PRESET

Linux / Unix

cpu · memory · filesystem · systemd

PRESET

Windows

service · eventlog · wmi · hyper-v

PRESET

Networking

ping · tcp · dns · interface · ntp

PRESET

Database

mysql · postgres · oracle · mssql

PRESET

SAP

hana · netweaver · basis · abap

PRESET

Middleware

jvm · tomcat · kafka · rabbitmq

PRESET

Cloud

aws · azure · gcp · k8s

CUSTOM

Your own

unlimited · import/export

ROUTING DECISION
label match → keyword match → **AI routing** → fallback. The reason is logged on every RCA.

Lynxmind-managed prompts you don't maintain —
plus unlimited custom agents.

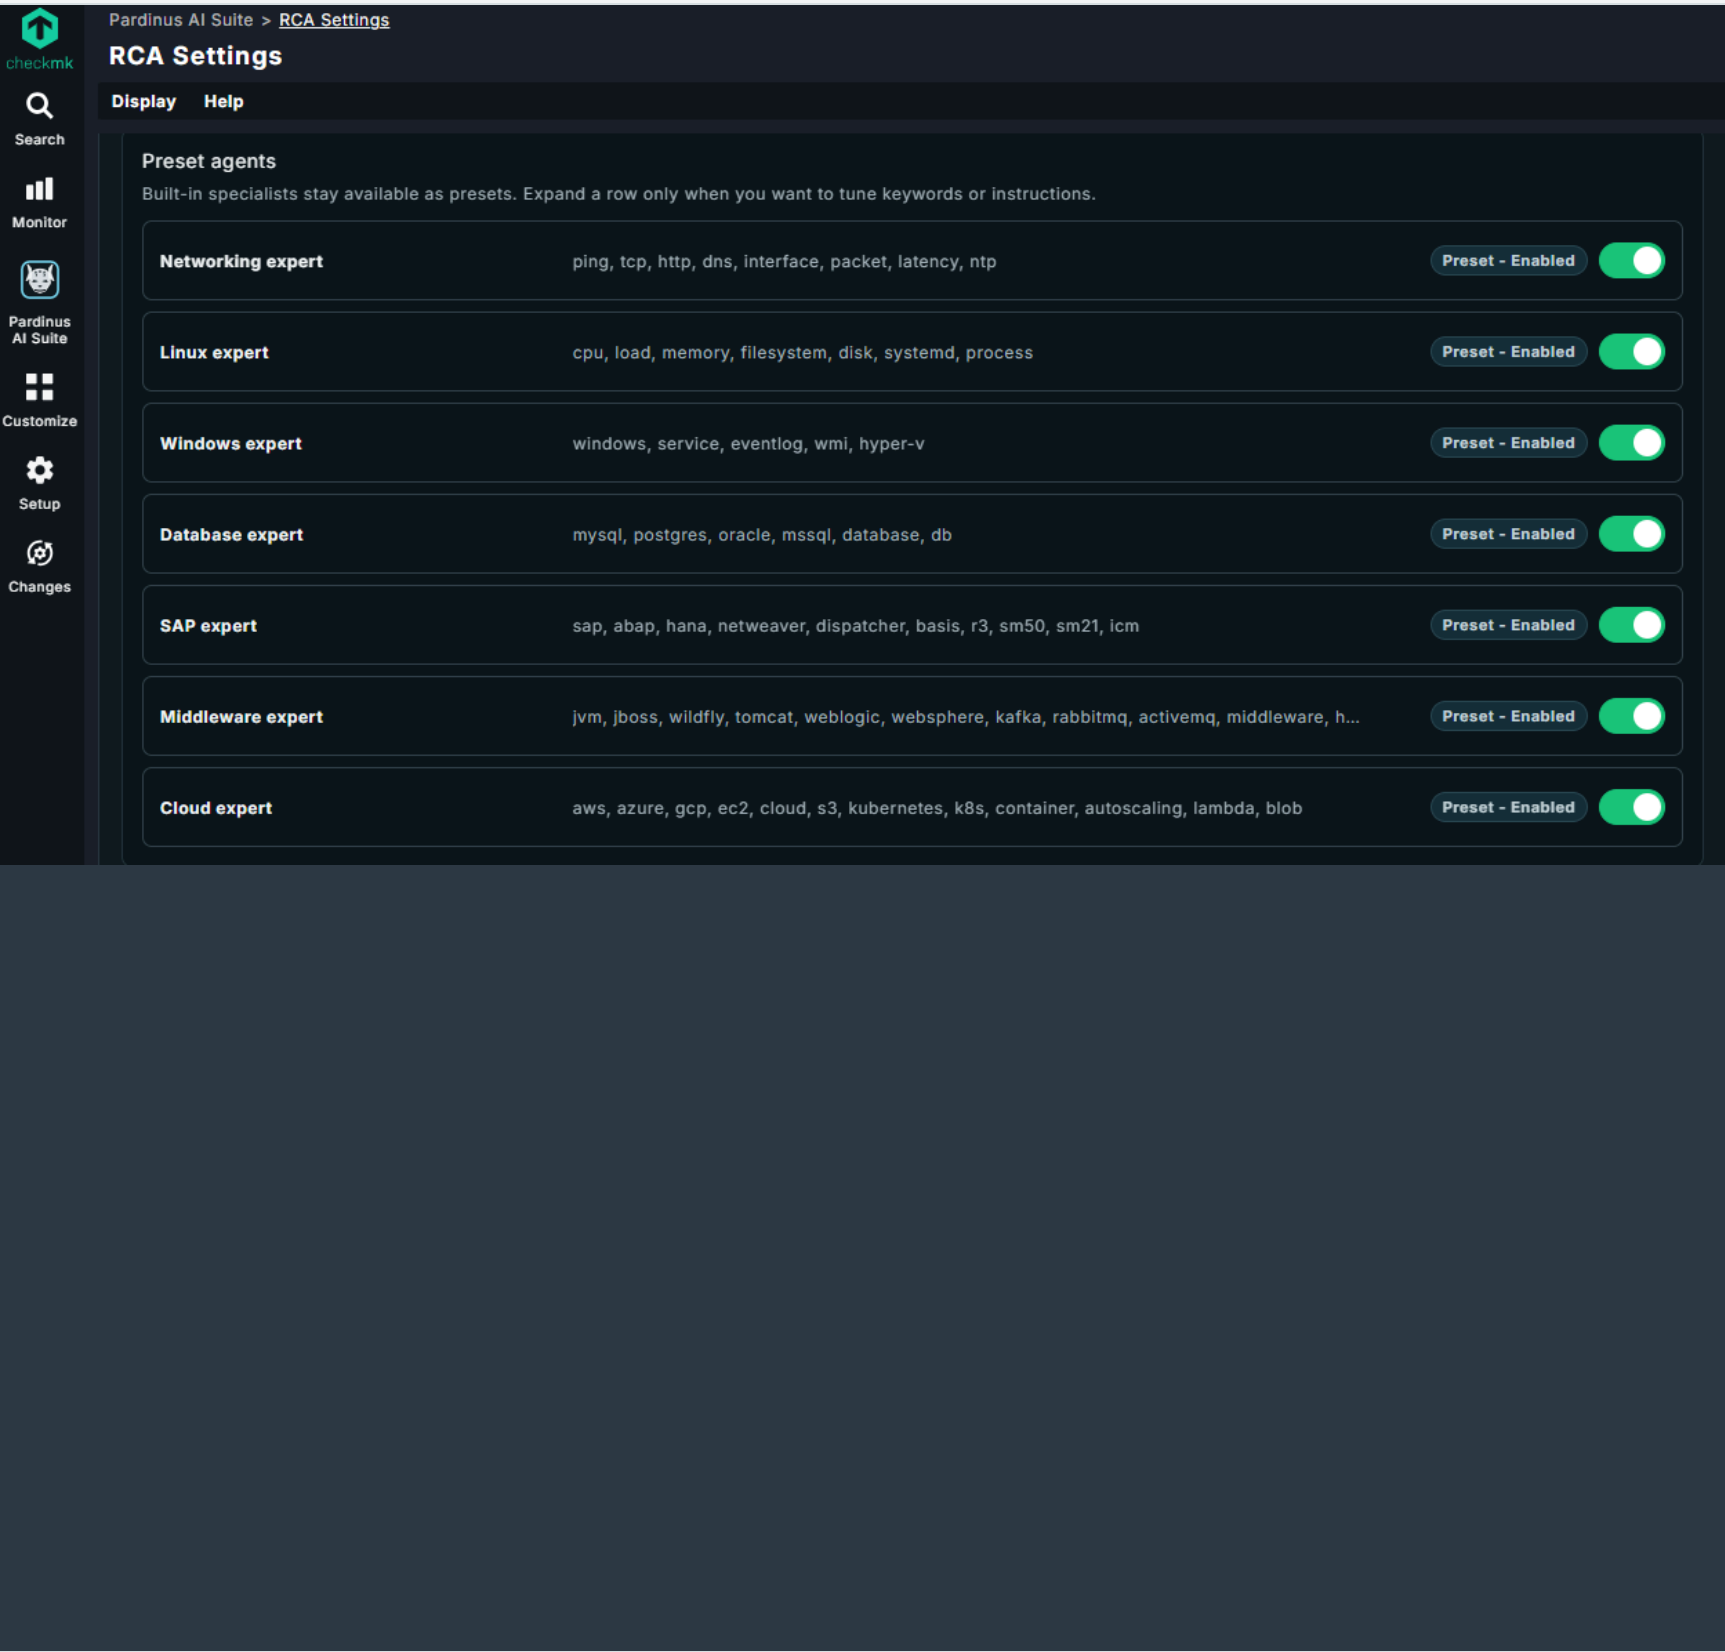


Fig. 09 — Preset specialists · Lynxmind-managed prompts

BEYOND PASSIVE ANALYSIS

It doesn't just read alerts.
It investigates.

Bounded by hard budgets, gated by human approval, fully audited.

Opt-in agentic RCA lets specialist agents run live, read-only diagnostics in a ReAct loop — observe, reason, act — feeding findings back as evidence.

REACT INVESTIGATION LOOP

observe

 →

reason

 →

act · read-only tool

 →

evidence

 ↺

SUPERVISOR · CHAIN

A central supervisor delegates to specialists in sequence.

SUPERVISOR · SWARM

Agents hand off directly to one another as needed.

SAFETY RAILS · ALWAYS ON

✓ Allowlisted, read-only tools

✓ Dry-run mode

Budgets: depth · tools · time · tokens

Every diagnostic traced & audited

APPROVAL GATE

auto

selective

full · human sign-off

ACTIVE DIAGNOSTICS · LIVE, READ-ONLY, AUDITED

SSH Linux OS

WinRM Windows

PostgreSQL db

Oracle db

SAP RFC netweaver

SNMP devices

Vendor CLI netcli · network devices

NETWORKING RCA · VENDOR AUTO-DETECT

Detects the device family from sysObjectID / sysDescr, then runs deep SNMP + vendor-CLI reads — interface counters, BGP/OSPF state, sensors.

Cisco

Juniper

Arista

Fortinet

Palo Alto

CREDENTIALS

Resolved from the Checkmk Password Store (host → folder → default). Agentless SSH via a read-only account. **Never stored by Pardinus.**

One click to an answer. Or a fix to the wrong rule.

Read-only, permission-gated, audited —
from a single button, not a full investigation.

QUICK DIAGNOSE

Green button, right on the alert

The Live Alert Queue renders a labelled button on every alert whose service matches a mapping — the operator sees exactly what will run before clicking. One bound read-only command, output plus a short summary.

web01 · CPU utilization

CRIT · 6m

▶ Top CPU processes

A EXECUTAR: TOP_CPU · READ-ONLY

PID	USER	%CPU	COMMAND
2841	app	97.3	java -Xmx8g ...
1120	root	2.1	checkmk_agent

◆ Escalate to a full RCA on the same host — one click.

CHECKMK RULE ADVISOR

When the threshold is the problem

The metric is within its expected range but the configured levels are too tight. Instead of chasing a fault that isn't there, the report points at the ruleset to review — with the current levels and a draft change note.

Likely false positive — threshold_misconfig finding raised

RULESET TO REVIEW

Memory usage levels

CURRENT LEVELS

WARN 80%

CRIT 90%

observed 83%

◆ Opens the ruleset in place, scoped to host / service.

Help where you already are. Smarter every week.

Read-only, explanation path only —
never a write action.

EXPLAIN WITH AI

Every inline-help box

A small button on every Checkmk help box explains that setting in plain language — what it controls, when to change it, sensible values — using the field label and active section as context.

context · label + help + tab

ASK AI · DASHBOARDS

An operator summary

A button beside the native Share on every dashboard reads the dashlet data — host/service stats, Top-10 tables, filesystems — and returns what deserves attention. It reads the tables, not the graph images.

reads · tabular data

SELF-LEARNING REPORT KB

Faster, cheaper recurring queries

After every report the model grades the query 1–5 and auto-saves the good ones. Future questions get the matching search hint injected — so recurring lookups get faster and cost fewer tokens over time.

grade ≥ 4 · auto-saved

ALWAYS-ON, NEVER NOISY

RCA fires before the on-call wakes up.

Rate-limited, flapping-aware, storm-protected —
AI usage stays bounded even when production isn't.

01 · RULES

Match what matters

Per-rule filtering on host, service, severity, plugin output, labels and tags. Cooldown windows prevent duplicate analysis.

host regex	service regex
severity	labels / tags

02 · RATE LIMITS

Bounded cost

Global and per-site caps, hourly and daily. AI consumption can never spike beyond the budget you set.

Global / hour	20
Global / day	100
Per-site / hour	10
Per-site / day	50

03 · STORM PROTECTION

One root cause, not 200.

When a storm hits, Pardinus picks the **earliest-state-change** alert as the probable root cause and runs RCA only on that one. Everything else is held with a reason.

PLUS

Flapping suppression · hold log with reason · amber storm banner on the alert queue

EVERY TOKEN IS METERED

You see the cost.
Before it shows up.

Provider-neutral metering, RCA-reuse savings,
model-level breakdowns · 24h / 7d / 30d.

Token Usage is a first-class page inside the suite — the same data that powers the commercial model later in this deck.

TRACKED

1.3M

tokens · per site, last 7 days

AI QUERIES

170

across RCA + chatbot

SAVED BY REUSE

465.6k

15 repeated alerts · zero extra cost

AVG LATENCY

22.9s

fresh RCA synthesis only

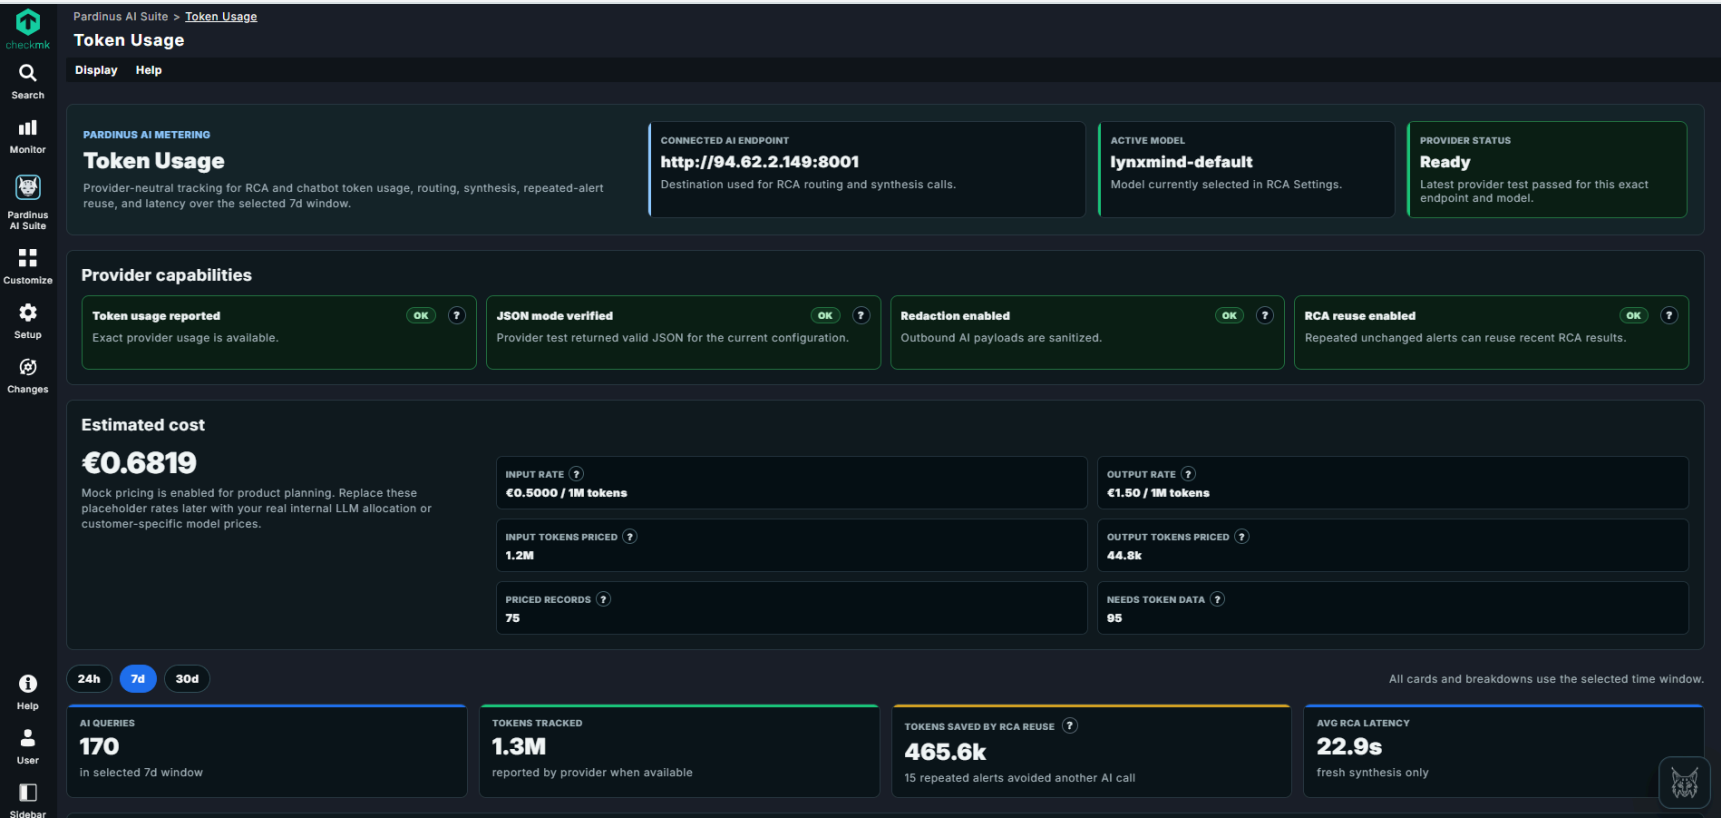


Fig. 11 — Token Usage dashboard · provider-neutral metering

OPEN ARCHITECTURE

API-first. Wire it into anything.

Trigger RCA from your incident pipeline, fetch reports into your tools — no Checkmk UI.

REST ENDPOINTS

POST	/api/v1/rca/run	trigger an RCA job
GET	/api/v1/rca/{job_id}	poll status + result
GET	/api/v1/rca/{job_id}/report.json	structured report
GET	/api/v1/rca/{job_id}/report.md	markdown for tickets
POST	/api/v1/chat/query	conversational ops
HOOK	webhooks · job.completed	push, don't poll

DROP-IN FOR THE TOOLS YOU ALREADY RUN

n8n	PagerDuty	ServiceNow	Jira
Slack	MS Teams	Opsgenie	Zapier
Ansible	Webhooks	Email / SMTP	your stack

● ● ● TRIGGER RCA FROM ANYWHERE

```
curl -X POST https://checkmk.lynxmind.ai/site/pardinus/api/v1/rca/run \
-H "Authorization: Bearer $PARDINUS_TOKEN" \
-H "Content-Type: application/json" \
-d '{
  "host_name": "db01",
  "service_description": "CPU load",
  "trigger_source": "n8n",
  "callback_url": "https://n8n.lynxmind.ai/webhook/rca-done"
}'

→ 202 Accepted { "job_id": "rca_a8f3...", "status": "queued" }
```

EXAMPLE · N8N WORKFLOW IN PRODUCTION

Auto-attach RCA to every new incident

Checkmk CRIT

 →

n8n trigger

 →

POST /rca/run

 →

await webhook

 →

create ServiceNow + RCA attachment

TIME TO SET UP

~15 min

LINES OF CODE

0

AUTH

Bearer

CHANGE INTELLIGENCE

It knows when you ship.

Every deploy, config change or release becomes context for the next spike.

Pipe CI/CD events — or any third-party signal — into Pardinus through an inbound webhook, so a symptom spike is traced straight back to the change behind it.

POST

/pardinus/events

CI/CD pipelines

config changes

feature flags

any 3rd-party event

PARDINUS · EVENT CORRELATION

DEPLOYMENT RECEIVED

app-x v2.4.1

via inbound webhook · 14:02 · source: CI/CD

CORRELATED SIGNALS · POST-DEPLOY

HTTP sessions

+320%

User experience · p95

4.8s

Memory usage

94%

Correlated to deploy app-x v2.4.1 — not a coincidental incident.

Infra is red. But is anyone actually affected?

Each is a deterministic correlation finding, computed before the LLM writes.

01 · CHECKMK BI

Business blast radius

Maps the alert to Checkmk BI aggregations — which business services this host underpins, and how far the failure reaches.

```
finding · business_impact
```

02 · ROBOTMK

Real user-facing impact

A failing synthetic test is proof a real journey is broken — and the failing step often names the broken dependency outright.

```
finding · synthetic_impact
```

03 · OPENTELEMETRY · RED

Rate • Errors • Duration

Application RED signals from OTel metrics in Checkmk — ties an error or latency spike straight to a recent deploy via the Event Console.

```
finding · red_signal_impact
```

Four things nothing else gives you.

Native, deterministic, hardened, auditable.

01

Native to Checkmk.

Not a side dashboard. Lives inside the menu, the action bar and the views the operator already uses. MKP-packaged, version-pinned per site.

02

Deterministic before AI.

Evidence is built by code, not by the model. The LLM only narrates a verified package — so confidence scores are real, not vibes.

03

Licensed & hardened.

Per-customer licensing, validated fully offline — no phone-home. Enforced end to end and fail-closed, with a hardened, tamper-resistant build.

04

Auditable output.

Every RCA emits typed JSON: evidence used, alternatives considered, missing data, timeline. Compliance and post-mortem ready.

WHERE IT EARNS ITS KEEP

From 3am pages to weekly reviews.

Six workflows, one structured source of truth.

A — ON-CALL

First-responder triage

RCA fires automatically on critical state. The on-call engineer wakes up to a structured report, not a wall of plugin output.

B — TIER-1

Service-desk escalation

L1 attaches the RCA report to the ticket before escalating. L2 gets evidence, not "it's broken".

C — CAPACITY

Recurring incident clustering

Pardinus surfaces patterns across past RCAs — "this CPU spike has happened 12 times on db01 in 30 days".

D — COMPLIANCE

Post-mortem evidence

Every RCA is auditable JSON. Feeds directly into ITIL post-incident reviews and SOC 2 evidence collection.

E — MSP

Multi-tenant RCA

Per-customer LLM provider, per-customer billing. MSPs can resell RCA as a managed service line.

F — KNOWLEDGE

Living runbook

Every resolved RCA becomes searchable context for the next one. Tribal knowledge becomes structured data.

OPEN BY DESIGN

Your model. Your sites. Your rules.

SUPPORTED LLM PROVIDERS · ALL PRODUCTION

OpenAI	gpt-4o · gpt-4-turbo	• PROD
Azure OpenAI	enterprise deployment	• PROD
Ollama	LLaMA · Mixtral · local	• PROD
vLLM	standalone inference server	• PROD
Open WebUI	self-hosted OpenAI wrapper	• PROD
Any OpenAI-compatible	your endpoint, your keys	BYO

Any OpenAI-compatible endpoint,
every Checkmk site you operate.

DISTRIBUTED MONITORING

RCA across every site

Cross-site analysis over Livestatus — Unix socket and TCP. A per-site collector reaches remote sites; the Remote Sites panel lives in Settings.

- livestatus · unix
- livestatus · tcp
- per-site collector

PRIVACY GUARDRAIL

Redacted before the LLM

IPs, hashes, UUIDs and secrets are masked pre-flight — whichever provider you point at. Your sensitive data never reaches the model.

Annual subscription. Sized to your estate.

Priced by monitored Checkmk services.
You bring your own LLM endpoint.

WHAT'S INCLUDED

Everything, in one license.

- ◆ Pardinus Chat — three modes
- ◆ Specialist agents + custom agents
- ◆ Integration API + webhooks
- ◆ Updates & support included
- ◆ Pardinus RCA — deterministic engine
- ◆ Proactive RCA + storm protection
- ◆ Cost-intelligence dashboard
- ◆ Unlimited operators

BRING YOUR OWN LLM

Data sovereignty by default.

from €3,200

PER YEAR · SIZED TO YOUR CHECKMK

- ◆ Any OpenAI-compatible endpoint · local, on-prem, or your cloud
- ◆ Data never leaves your premises
- ◆ Fits regulated & air-gapped environments

SIZED TO YOUR CHECKMK

Pricing scales with you. Same model as Checkmk.

Annual subscription by monitored services.
Bring your own OpenAI-compatible endpoint.

MONITORED SERVICES	ANNUAL SUBSCRIPTION · BYO LLM ENDPOINT	TYPICAL ESTATE
7,000	€1,600 per year	Small team · single site
12,000	€2,400 per year	Growing infrastructure
18,000 POPULAR	€3,200 per year	Mid-market · ~3,000 hosts
30,000	€4,700 per year	Large enterprise
60,000	€7,800 per year	Multi-datacenter
100,000	€11,500 per year	Global estate
200,000	€19,200 per year	Service provider · MSP

* > 200,000 services – custom quote. Multi-site, multi-tenant, and MSP partner discounts on request.

ROI · 18,000-SERVICE SITE · 12-MONTH HORIZON

A typical mid-market customer.

~3,000 hosts · ~40 deep RCAs/month ·
auto-triage on every CRIT · daily chatbot ops.

COST — ANNUAL

License · Pardinus Suite (BYO LLM)	€3,200
AI endpoint	customer-owned
Onboarding (one-off)	€2,500
Year 1 total	€5,700

VALUE — ANNUAL

MTTR reduction	–38%
Engineer hours saved	~720 h
Engineer hour value	€85
Time savings	€61,200
Avoided major-incident escalation	€18,000
Net value · Year 1	€73,500

PAYBACK PERIOD

≈ 1 month payback · 13.9× ROI

● EARLY ACCESS

30% off your first-year license — founding customers only.

Sign before 30 June 2026 and lock the discount for the full initial term. Pilot fee fully credited.

YEAR-1 LICENSE

€3,200

€2,240

LET'S RUN THE FIRST RCA ON YOUR STACK

Stop guessing. Start **knowing.**

30-day pilot on your Checkmk site. We deploy the MKP, run RCA on your top 10 noisiest services, and hand you a written report. Pilot fee credited against your year-one license.

TALK TO US

vendor@lynxmind.com

TRY THE DEMO

Reach out for demos

PILOT TIMELINE

1 min to deploy • measure immediately

A PRODUCT BY **LYNXMIND**

© 2026 LYNXMIND · PARDINUS IS A LYNXMIND PRODUCT