



CORTEX • SUPPORT INTELLIGENCE ENGINE

Every other AI support platform peaks on day one. Cortex peaks on day one *thousand*.

The self-hosted Support Intelligence Engine for engineering-led manufacturers — trained on your products, running on your infrastructure, learning from every case it touches.

SELF-HOSTED

TRAINED ON YOUR DATA

YOURS FOREVER

95.8%

CITATION RATE

5.0%

FABRICATION RATE

83.3%

RECALL @ 5

0

EXTERNAL AI CALLS

Measured on self-hosted infrastructure — no customer data ever leaves the building.

THE PROBLEM CORTEX WAS BUILT TO KILL

Customer support is a quiet, compounding tax on an engineering business.

£40,000–£100,000 given away every year

Unbilled engineering time bleeds into free support — untracked, unreconciled and invisible on the P&L.

Support data in the wrong cloud is a liability

Configurations, failure modes and NDA-protected detail should never enter a multi-tenant AI platform.

Your best engineers stuck in triage

The knowledge that resolves cases lives in their heads — and walks out of the building the day they leave.

Patterns nobody connects

The same component fails across systems and customers for months before anyone joins the dots.

FOUR WAYS CORTEX PAYS FOR ITSELF

01

Cost centre → profit centre

Every support hour tracked against your free-support policy. The moment a case becomes billable, Cortex generates a professional three-option conversion offer.

02

Liberate your R&D

Known issues resolve without an engineer; only the genuinely novel escalates. Tribal knowledge becomes organisational capability — and your senior people get their week back.

03

Data that never leaves

Every byte of inference runs locally. Air-gappable on every tier — the only AI support platform architecturally able to serve the defence supply chain under DEFCON 658.

04

The learning flywheel

Cortex learns from your company knowledge and your team's corrections. Where other platforms plateau on day one, Cortex compounds. Deal with it once — learnt for life.

A manufacturer handling 500 cases a year, converting one in five previously-unbilled cases at an average of £600, recovers £60,000 in year one — more than platform, implementation and maintenance combined.

ONE ENGINE • SEVEN STAGES • END TO END

From first contact to resolution — *and it learns.*

01 • INTAKE Every channel, one record Each case normalised into a single unified record — nothing falls through the cracks.	02 • CLASSIFY Right team, first time Support, sales and info separated — sales enquiries never clog the queue.	03 • TRIAGE Right case, right owner Priority scored on contract, severity, warranty and repeat-failure signals.	04 • RESOLVE The 80% already solved Validated fixes proposed, policy enforced, the customer reply drafted.	05 • DETECT Patterns, not one-offs Failures clustered across unit, part and component; systemic alerts raised.	06 • IMPROVE Sharper every ticket Every closed case feeds the knowledge base — the platform compounds.	07 • SATISFY CSAT, visible & actioned Post-closure surveys read for sentiment; unhappy customers escalated.
---	--	---	--	--	--	---

CAPABILITIES — BUILT, IN PRODUCTION

01

Support that bills itself

The commercial policy engine turns the awkward "this isn't covered any more" conversation into a professional, on-brand offer — and tracks every conversion for your commercial analytics.

02

Pattern detection → R&D intelligence

Recurring failures surface as structured Product Improvement Recommendations for engineering — with failure timeline, support-cost impact and the at-risk installed base attached.

03

Equipment-aware knowledge

Knowledge organised around your equipment, not flat folders. Answers arrive at product, sub-assembly and component level — specific to the machine in front of you, never generic.

04

A glass box, not a black box

Every recommendation cites its source; if it can't, it flags inference rather than asserting fact. Every setting is yours to tune — and the AI model itself is swappable, so there is no lock-in at the AI layer.

05

Just ask Cortex

A permission-aware assistant over your live data. Ask how to do something and it walks you through it — onboarding becomes a conversation, and a new hire is productive on day one.

06

In the field — even off-grid

A native mobile app puts the whole operation in an engineer's pocket: voice field reports that transcribe on-device, and full case access with no signal at all — basement, shielded lab or offshore.

EIGHT CHANNELS IN Email • web portal • distributor portal • agent entry • field voice & video • mobile • forwarded email • API

CRM & BUSINESS SYSTEMS

Salesforce • HubSpot • Zoho • Dynamics • SAP

Your customers and installed base, pulled in as live context on every case.

TICKETING & INBOUND

Jira • Zendesk • Slack • Teams • webhook

Tickets and messages become structured cases automatically.

DOCUMENTS & KNOWLEDGE

SharePoint • network file share

Point Cortex at your manuals; it ingests, validates and indexes them — read-only at the source.

IDENTITY

SAML 2.0 • OpenID Connect

Your people sign in with the identity provider you already run — Entra ID, Okta.

Cortex reads *from* your systems — it never needs write-access into them.

DEFENSIBLE PERFORMANCE

Published numbers. *Measured*, not asserted.

Cortex publishes its accuracy against a locked, open benchmark — a 120-case hand-authored corpus across ten real manufacturing failure clusters, a held-out evaluation split and bootstrap confidence intervals, measured on the same hardware we ship. The full, reproducible methodology is available on request.

METRIC	RESULT	WHAT IT MEANS
Citation rate	95.8%	How often a recommendation is grounded in a real, clickable source document.
Fabrication rate	5.0%	How rarely the AI asserts something it cannot ground — far below typical LLM baselines.
Recall @ 5	83.3%	When the right answer exists in your knowledge base, how often it lands in the AI's top five.

Supporting figures: **Hit@3 of 60.4%**, and **110 of 120 cases** resolved with a cited, grounded answer. For context, NVIDIA's published 2026 production-RAG baseline tops out at **64.8%** page-level retrieval — Cortex's 83.3% clears it comfortably, on self-hosted, open-source infrastructure with zero external AI calls. Scale validated separately at **1,000 cases** with sustained throughput and no degradation — the accuracy you measure on day one is the accuracy you keep at volume.

AN ENGINEERING CLAIM, NOT A MARKETING ONE

Every retrieval change ships only after a falsifiable experiment against a recorded gold set — and we publish the full table, including the changes that didn't work. Proven in production at a UK scientific-equipment manufacturer running three subsidiaries on a single platform.

SOVEREIGN BY DESIGN

Your data, your AI, *your walls*.

Enterprise software with a formal security architecture and verifiable controls — 4,000+ automated tests on every code change and a production-approved security risk score.

- ▶ **Dedicated single-tenant deployment.** Your data, your AI model and your knowledge on your own instance — not a shared cloud. Everything stays inside your perimeter.
- ▶ **Data sovereignty.** 100% self-hosted inference — no OpenAI, Anthropic, Google, Azure or Cohere. Outbound AI calls are blocked at the network layer; air-gappable on every tier.
- ▶ **Encryption & access.** AES-256-GCM at rest, TLS 1.3 in transit; SSO via SAML 2.0 and OpenID Connect, multi-factor authentication and granular role-based access — 8 roles, 32 permissions.
- ▶ **Audit & oversight.** Every action logged immutably with a seven-year default retention — and a human approves every customer-facing output. No autonomous decisions.

GDPR

ISO 27001-ALIGNED

DEFCON 658 · CSMV4

CYBER ESSENTIALS

EU AI ACT

One-click GDPR right-to-erasure with full cascade; the appliance tier runs fully air-gapped for defence and export-controlled work.

THE COMPARISON

Why Cortex *wins*.

✓ NATIVE ± PARTIAL / ADD-ON — NOT AVAILABLE

CAPABILITY	CORTEX	CLOUD HELPDESK ¹	SELF-HOSTED ITSM ²	ENTERPRISE FSM ³	AI-NATIVE ⁴
AI inference on customer infrastructure	✓	—	—	—	—
Air-gappable / DEFCON 658 ready	✓	—	—	—	—
Equipment-centric data model (BOMs, serials, warranty)	✓	—	±	✓	✓
Commercial policy enforcement (auto billing detection)	✓	—	—	±	—
Automated quote / conversion generation	✓	—	—	—	—
Pattern detection → R&D recommendations	✓	—	—	±	±
Three-entity channel model (distributor vs end customer)	✓	—	—	±	—
BOM-linked knowledge base (hierarchical retrieval)	✓	—	—	±	—
AI learns from your product data & compounds over time	✓	±	—	±	±
Citation engine — every recommendation traces to source	✓	—	—	—	±
Swappable LLM — no lock-in at the AI layer	✓	—	—	—	—
Offline-capable field access (works with no connectivity)	✓	—	—	±	—
Built-in AI assistant over your own data, self-hosted	✓	±	—	±	±
Published retrieval benchmark with full methodology	✓	—	—	—	—
Customer owns the knowledge base	✓	—	✓	—	—
Perpetual licence option	✓	—	✓	—	—
Priced for £2m–£200m manufacturers	✓	±	✓	—	—

¹ Zendesk, Freshdesk, Intercom • ² HaloITSM, ManageEngine, SysAid, TOPdesk • ³ ServiceMax, IFS, SAP FSM, Salesforce FSL • ⁴ Aquant, Neuron7, Mavenid

THE ONE LINE THAT ENDS THE COMPARISON

Every other AI support platform requires your customer data to go to an external AI company. Cortex is the only one where everything — data, AI, knowledge — stays on your hardware.

DEPLOY IT YOUR WAY — AND OWN IT

Appliance

ON-PREMISES • PERPETUAL

Pre-configured, customer-owned hardware. Air-gappable, no internet required. For defence, NDA-heavy and regulated work.

Managed VPS

DEDICATED • PERPETUAL

A dedicated UK/EU instance, Ascentis-managed. For engineering SMEs wanting full data control without the hardware.

SaaS

HOSTED • SUBSCRIPTION

Ascentis UK/EU cloud, logically isolated. The fastest time-to-value — with full data portability if you move tiers.

The same software and the same guarantees on all three tiers. Perpetual licensing means you can **own** Cortex outright rather than rent it — and a typical deployment pays for itself inside the first year on recovered support revenue alone. Indicative pricing and a tailored ROI model are available on request.

BUILT FOR ENGINEERING-LED MANUFACTURERS • £2M–£200M REVENUE

Scientific instruments • test & measurement • semiconductor equipment • industrial machinery • medical devices • precision engineering • defence subcontractors • process equipment

See it on *your* data. Twenty minutes is all it takes.

BOOK A 20-MINUTE WALKTHROUGH • SECURITY OVERVIEW UNDER NDA • BENCHMARK METHODOLOGY ON REQUEST

info@ascentis-ai.com • ascentis-ai.com