

PRIVATE AI INFRASTRUCTURE · SERVICE PROFILE

Private AI

Put AI to work on your most sensitive data — without that data ever leaving your control. You own the model, the answers and the IP.

10,000+ systems delivered for clients in 150+ countries — and we run our own 17-brand company on systems like this. You own everything we build.

19+ YEARS BUILDING	28,000+ CLIENTS SERVED	10,000+ SYSTEMS SHIPPED	200+ INTEGRATIONS BUILT
--	--	---	---

PROVEN ON US

We run AI in production across our own platforms serving 1,000+ businesses. Which means private AI here is operating experience, not a checkbox.

01 · THE COST YOU PAY TODAY, AND THE FIX

WHAT IT COSTS YOU TODAY

Every prompt your team sends to a public AI API takes your data outside your control. Customer records and trade secrets cross borders you may not be allowed to cross. So you block AI — and lose the hours it could save, every month. Or you allow it — and carry a risk your regulator or clients never agreed to.

WHAT YOU GET INSTEAD

You get production AI inside your own environment — private cloud, on-premise or air-gapped. We fine-tune open models on your data, and the result is yours outright. We have run 10,000+ systems for 19+ years. Which means your AI is built to production standard, and your data never leaves your perimeter.

What you get

01

Your Data Stays Home

We deploy in your private cloud, your data centre, or an air-gapped network. Which means nothing leaves your perimeter and no outside provider ever touches your data.

02

A Model You Own

We fine-tune open-weight models on your own data. Which means you own the trained model outright — never rented, shared, or used to train anyone else's system.

03

Answers You Can Verify

AI grounded in your own documents, databases and systems, with sources shown for each answer. Which means your team can check every answer before acting on it.

04

Compliance Built In

Data residency, encryption, role-based access and audit logs are set up from day one. Which means you can show regulators and clients exactly who saw what, and when.

05

AI Inside Your Tools

Secure connections put AI inside your ERP, CRM, helpdesk and internal tools. We maintain 200+ integrations. Which means AI works where your team already works — no new silo.

06

No ML Team Needed

Optional 24/7 monitoring, updates and scaling run by our team. Which means you get reliable production AI without hiring an in-house ML team.

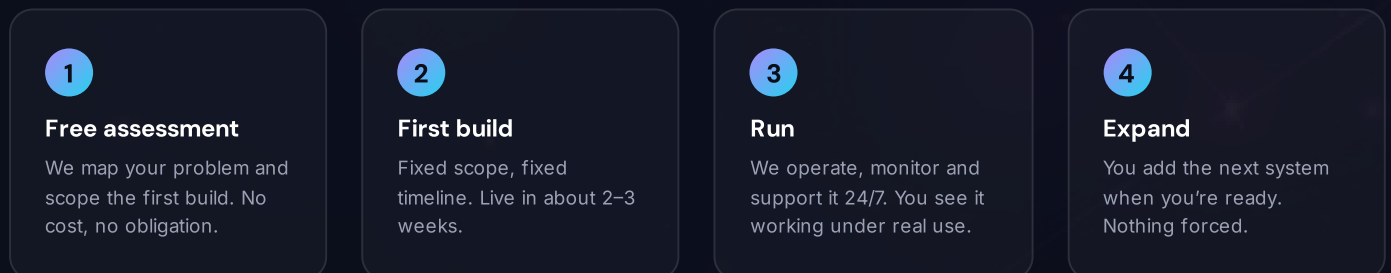
YOU OWN IT | Everything we build — code, data and IP — is yours. Which means no lock-in, ever.

02 · WHAT'S IN SCOPE

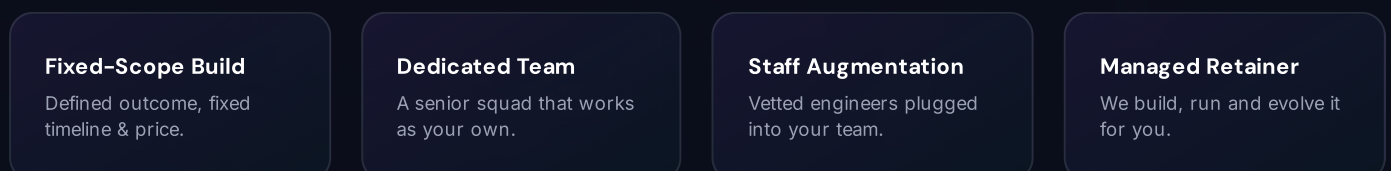
What's in scope

- ✓ Free use-case and data-readiness assessment
- ✓ Model selection and fine-tuning on your data
- ✓ Role-based access control and audit logging
- ✓ Integration with your ERP, CRM and internal tools
- ✓ Private cloud, on-premise or air-gapped deployment
- ✓ AI answers grounded in your own documents (RAG)
- ✓ Data residency and encryption set up for you
- ✓ Optional 24/7 monitoring and managed operations

How you start



Ways to engage



03 · WHAT CHANGES FOR YOU

What you can expect



AI capability for your team with zero loss of control over your data



A trained model you own outright — a business asset, not a rented subscription



Governance you can prove to regulators, auditors and enterprise clients



Routine work shifts to AI — cutting overhead 40–70% where automation fits your workflows

Products we run at scale

ServerAdmin.ai

AI server management

serveradmin.ai

rnaccounts

Cloud ERP & POS

rnaccounts.com

AI Ticket & Chat

AI helpdesk

aiticketandchat.com

Common questions

Where does the AI actually run?

Inside an environment you control — your private cloud, your own data centre, or an air-gapped network. That includes AWS, Azure, GCP or a private VPC. Your data is never sent to a shared public endpoint.

Who owns the model once it is trained?

You do. We fine-tune open-weight models on your data, and the trained model is yours outright. It runs on your infrastructure, and it is never used to train anyone else's system. You keep the model, the data and the IP if we ever part ways.

Can it connect to the systems we already use?

Yes. It connects to your ERP, CRM, helpdesk and internal tools through secure APIs. It can also answer from your live documents and databases. We maintain 200+ integrations across our products, so most common systems are already covered.

How do we start?

With a free assessment of your use case and data — no cost, no obligation. Then we agree a fixed scope, so there is no surprise bill. Your first build goes live in about 2–3 weeks. After launch we run it, and you expand only when it proves itself.

Book your free assessment

Free assessment → fixed-scope first build, live in ~2–3 weeks → run & expand.

No cost, no obligation. Full profile: rapinova.com/resources

Zeeshan Mudassir · Founder

hello@rapinova.com

+971 50 186 3838 · UAE

rapinova.com



WHATSAPP US



RAPINOVA.COM