PHYSTER
TECHNOLOGY

Augmented
Personal
Environment

Bring AI into the organization.

Keep the organization in charge.

APE – Augmented Personal Environment

The control layer for enterprise GenAI.

Generative AI is no longer optional.
But for most organizations, adopting it responsibly is harder than adopting it at all.

Every new model, every new vendor,
every new integration raises the same three questions:

Where does our data go?
Who is accountable for what the model says?
What happens when the rules change again?

APE was built to answer those questions before they become incidents, not after.

It is not another AI tool competing for a place in your workflow.
It is the layer between your organization and every model you choose to use,
keeping the intelligence flexible and the control fixed.

"Let the AI do the work. You run the business."

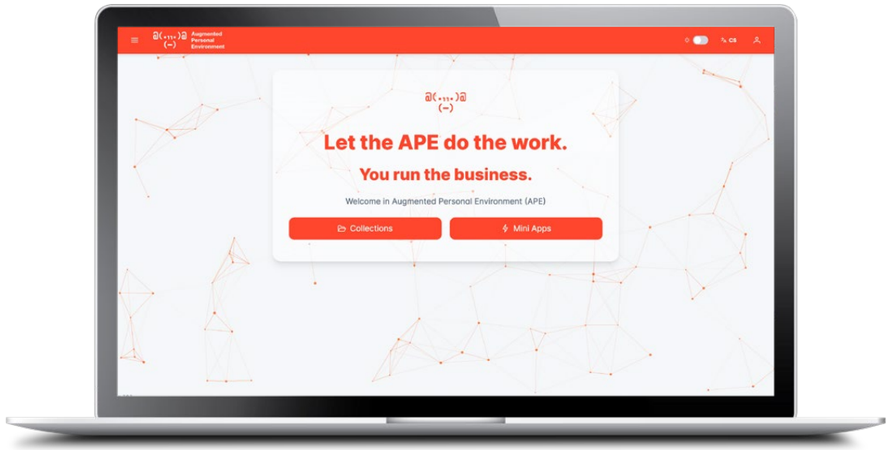

WHY PILOTS DON'T SCALE

The risk isn't AI. It's losing track of it.

Most AI initiatives don't fail because models are weak.
They fail because organizations lose control over data, usage and governance.

1 | DATA

Once a prompt leaves your environment, you no longer control where it goes or who can see it. For a regulated enterprise, that single fact can keep GenAI out of production indefinitely.

2 | PEOPLE

Most teams don't have the in-house expertise to evaluate models, write safe prompts, or track a field that changes every few months. Relying on a handful of specialists doesn't scale.

3 | VENDOR LOCK-IN

The model that solves a problem today may be discontinued, repriced, or outperformed within a quarter. Building deeply around one provider is a decision you will revisit sooner than planned.

4 | COST AND TIME

AI experiments are easy. Scaling them across an organization is not. Without visibility and governance, costs grow faster than value.

None of these problems are solved by a better model.
They are solved by better architecture.

THE PLATFORM

A control layer, not another tool.

APE sits between your organization and the AI models you choose to use. It provides one governed environment for AI across the organization, replacing fragmented experiments with a shared approach to knowledge, access and accountability.

Underneath, APE is model-agnostic by design. Swap a provider, add a new model, or run several side by side, and your people, your data structures, and your audit trail stay exactly where they are.

Above that, every authorized user starts from the same place: their own governed data, their own assistants, and their own record of how AI was used to reach a conclusion.

APE can run inside your own infrastructure, in a private cloud, or as a managed multi-tenant service.

The operating model follows your security policy, not the other way around.

GOVERNANCE BY DESIGN, NOT BY POLICY DOCUMENT

Data stays organized. Access stays deliberate.

Inside APE, knowledge is organized into collections, each with its own owner and access rights. A finance team's knowledge stays separate from HR. Partner-facing assistants don't share the same context as internal ones.

Every collection carries its own audit trail. Questions, answers and source documents remain traceable long after the conversation ends. When a result needs to be reviewed, the evidence is already there.

This is the difference between an organization that has an AI policy, and one that can prove it follows one.

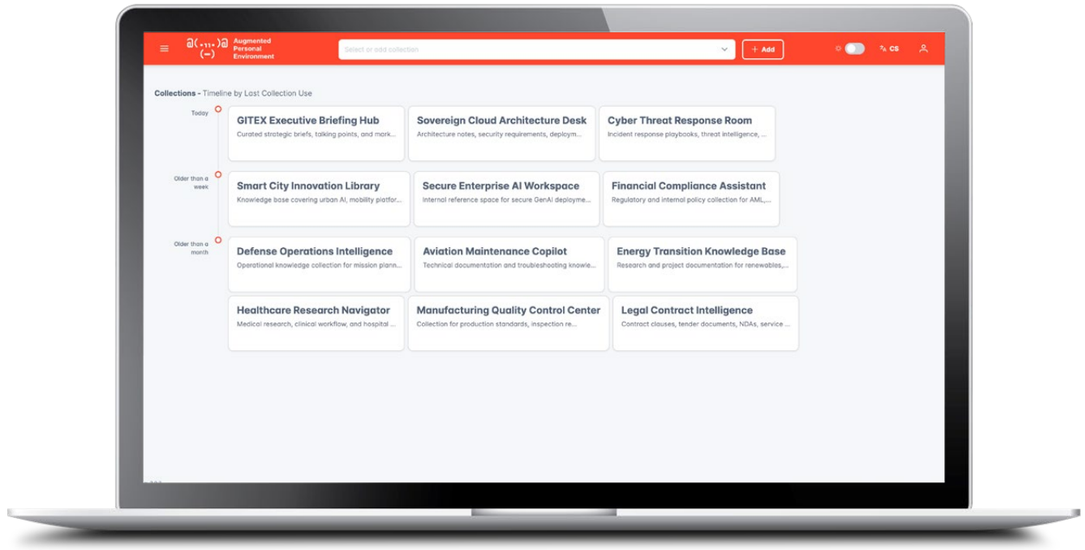

HOW CONTROL IS MAINTAINED

One architecture. Full accountability.

Enterprise documents, databases and applications enter APE through governed entry points. Each collection maintains its own users, permissions and knowledge context. Questions are matched against trusted sources first. Only then are AI models used to generate answers. The model sees only the information required to answer the question. Your infrastructure remains under your control.

The architecture stays consistent.
Across collections.
Across mini apps.
Across models.

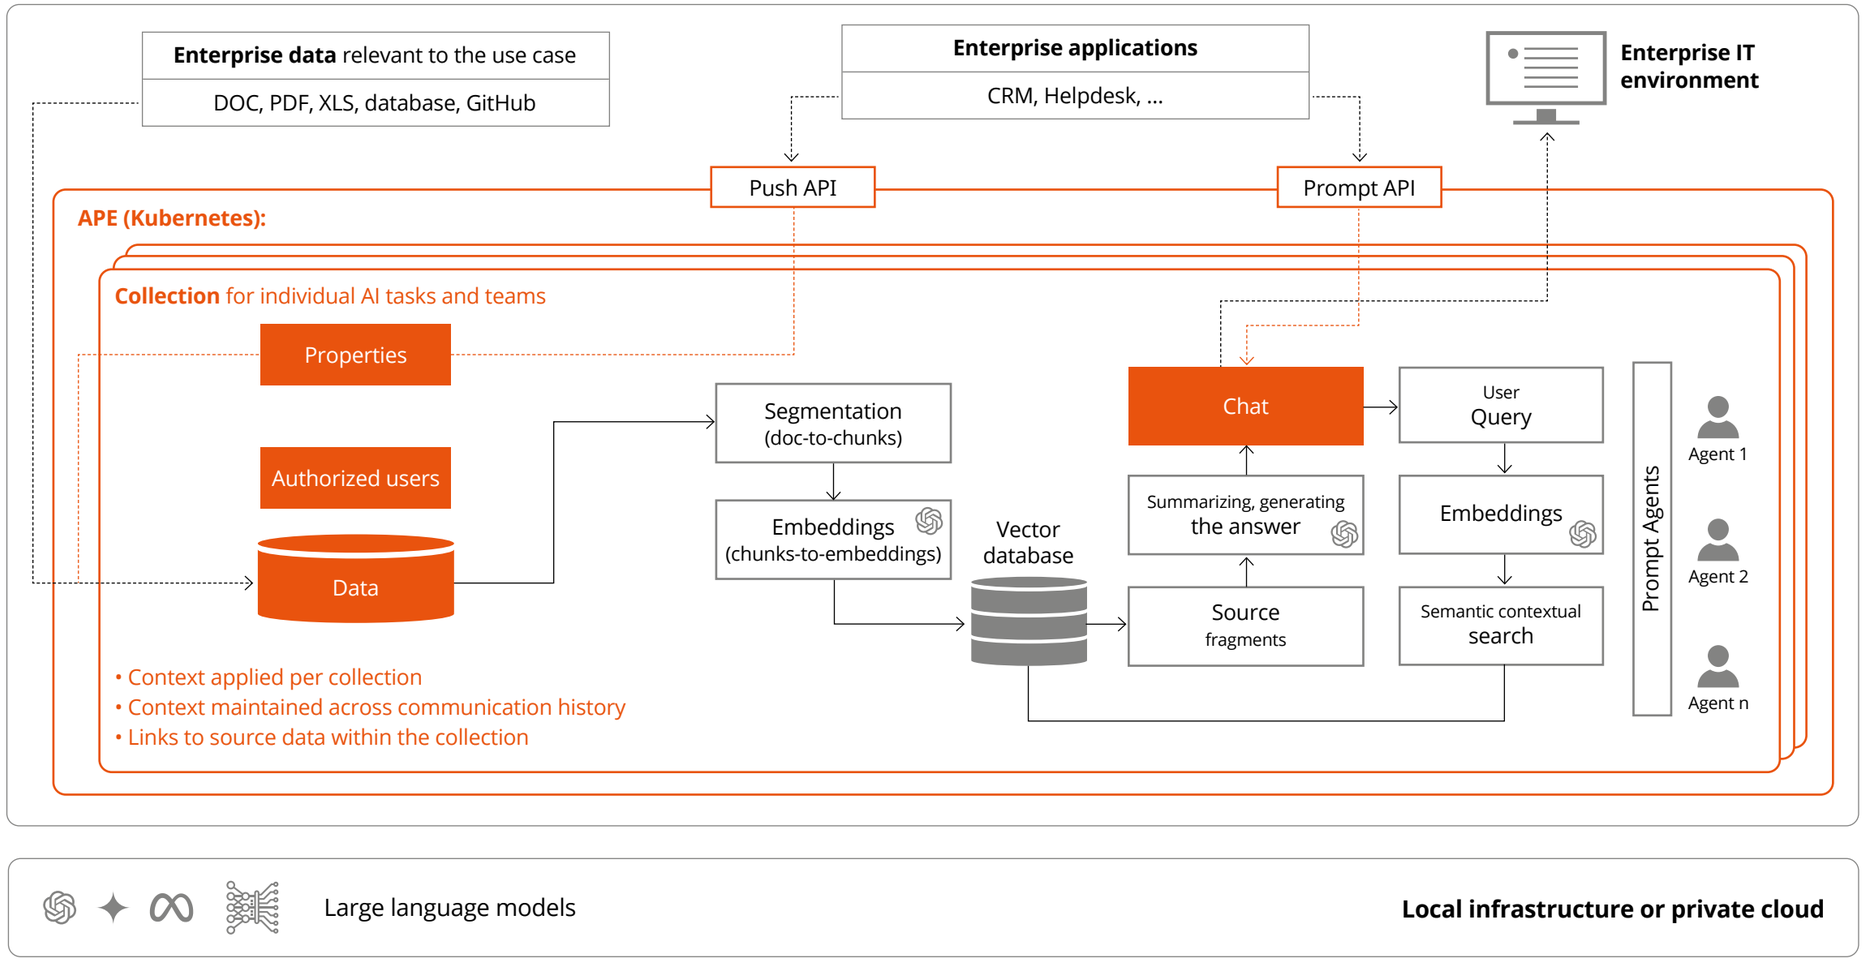

SECURITY

Built for the conversation with your CISO, not around it.

Security teams rarely reject AI because they don't understand it. They reject it because nobody can answer their questions quickly enough. APE is built so those questions already have answers.

ZONED ACCESS

APE operates as a closed ecosystem with defined access rights. Every interaction happens inside a boundary your organization sets: never against the open internet, never with unverified sources.

GROUNDED ANSWERS

Responses are built from your verified collections, not from whatever the model happened to learn elsewhere.

STRUCTURED, AUDITABLE OUTPUT

Responses are returned in structured formats that are easy to log, review and reuse.

FULL COMMUNICATION AUDIT

Every user action, every API call, every exchange with a model is recorded. If something needs to be reconstructed after the fact, it can be.

ENCRYPTION AND IDENTITY

Data stays encrypted. Identity stays under your control. APE integrates with the identity systems you already trust.

Every automated task leaves a record: what ran, on what data, with what result.

Security is easier when the evidence is already there.

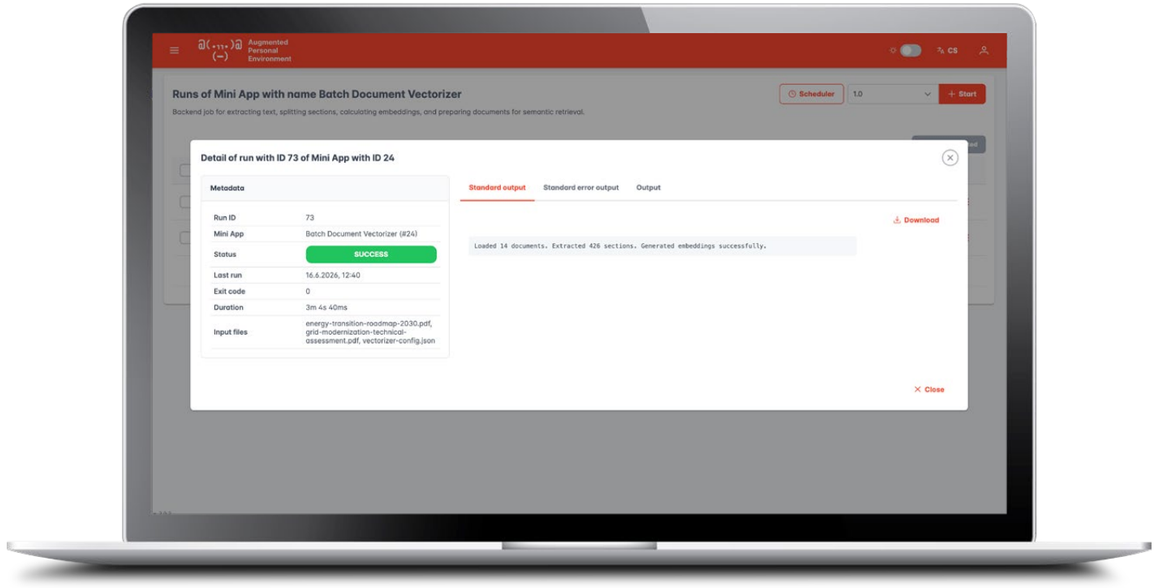

Security concerns should be explicit, not hidden.
APE addresses the risks security teams already understand and document.

PROMPT INJECTION	DATA LEAKAGE
DATA POISONING	JAILBREAKING
UNAUTHORIZED ACTION	

IN PRACTICE

Real use cases. Real data. Real constraints.

SYNTHETIC DATA FOR HEALTHCARE

Generate realistic datasets for testing hospital information systems without exposing patient records.

RESULT: Safe testing without touching production data

DATA TRANSFORMATION ACROSS STANDARDS

Automate conversion between FHIR and CDA formats.

RESULT: Faster interoperability projects and less manual work.

TENDER & RFP ANALYSIS

Analyze hundreds of pages and prepare executive summaries in minutes.

RESULT: Hours of manual review reduced to minutes.

REAL PROJECTS. REAL OUTCOMES.

- Test environments without exposing sensitive data.
- Hours of document review reduced to minutes.
- Interoperability tasks automated instead of rewritten manually.
- Reusable workflows shared across teams.

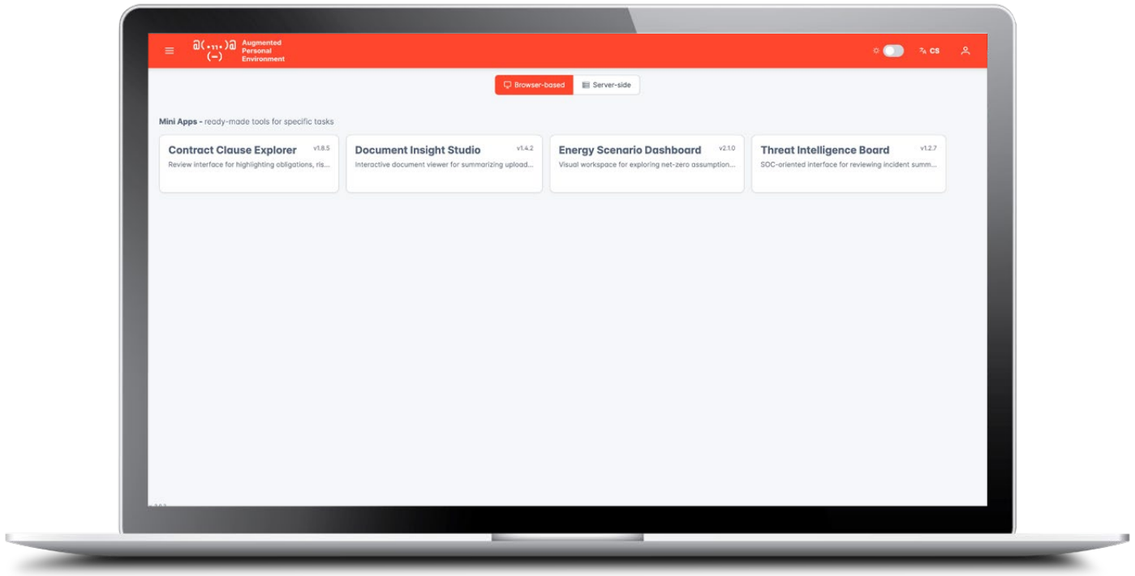

Repeatable AI tasks, defined once and reused by an entire team, without writing a line of code.

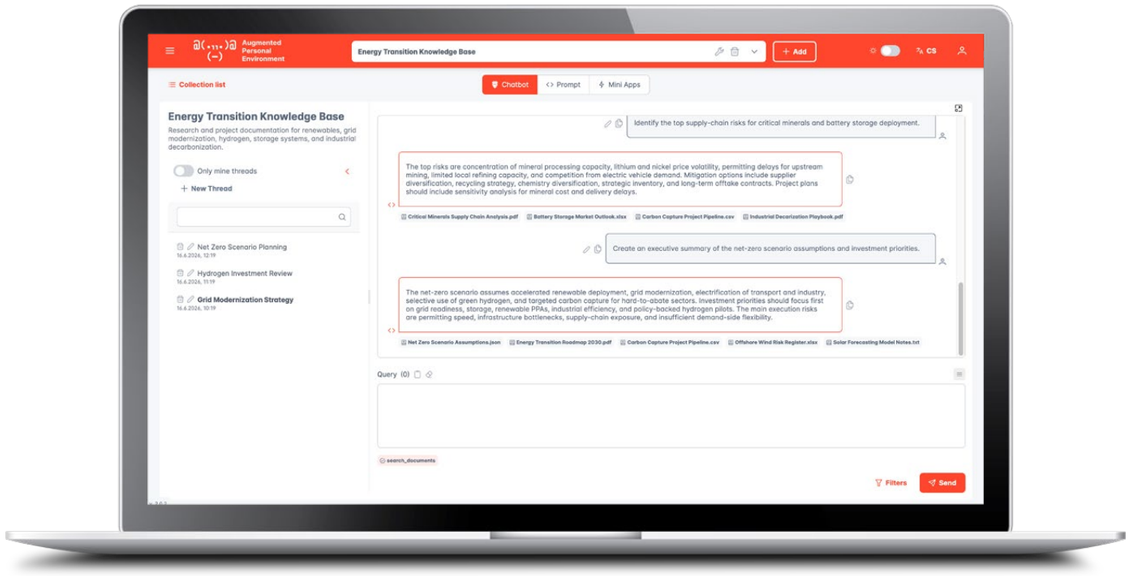

Answers are returned with their source documents attached, so a result can be checked, not just trusted.

DEPLOYMENT

It fits your infrastructure. Not the other way around.

DEPLOYMENT
Run APE on-premise, in your private cloud, or as a managed service.

VENDOR INDEPENDENCE
Switch models without rebuilding your environment.

COST VISIBILITY
Track AI usage by user, team and collection before invoices become surprises.

AI costs should be visible before they become invoices.

AI spending becomes measurable, not mysterious.

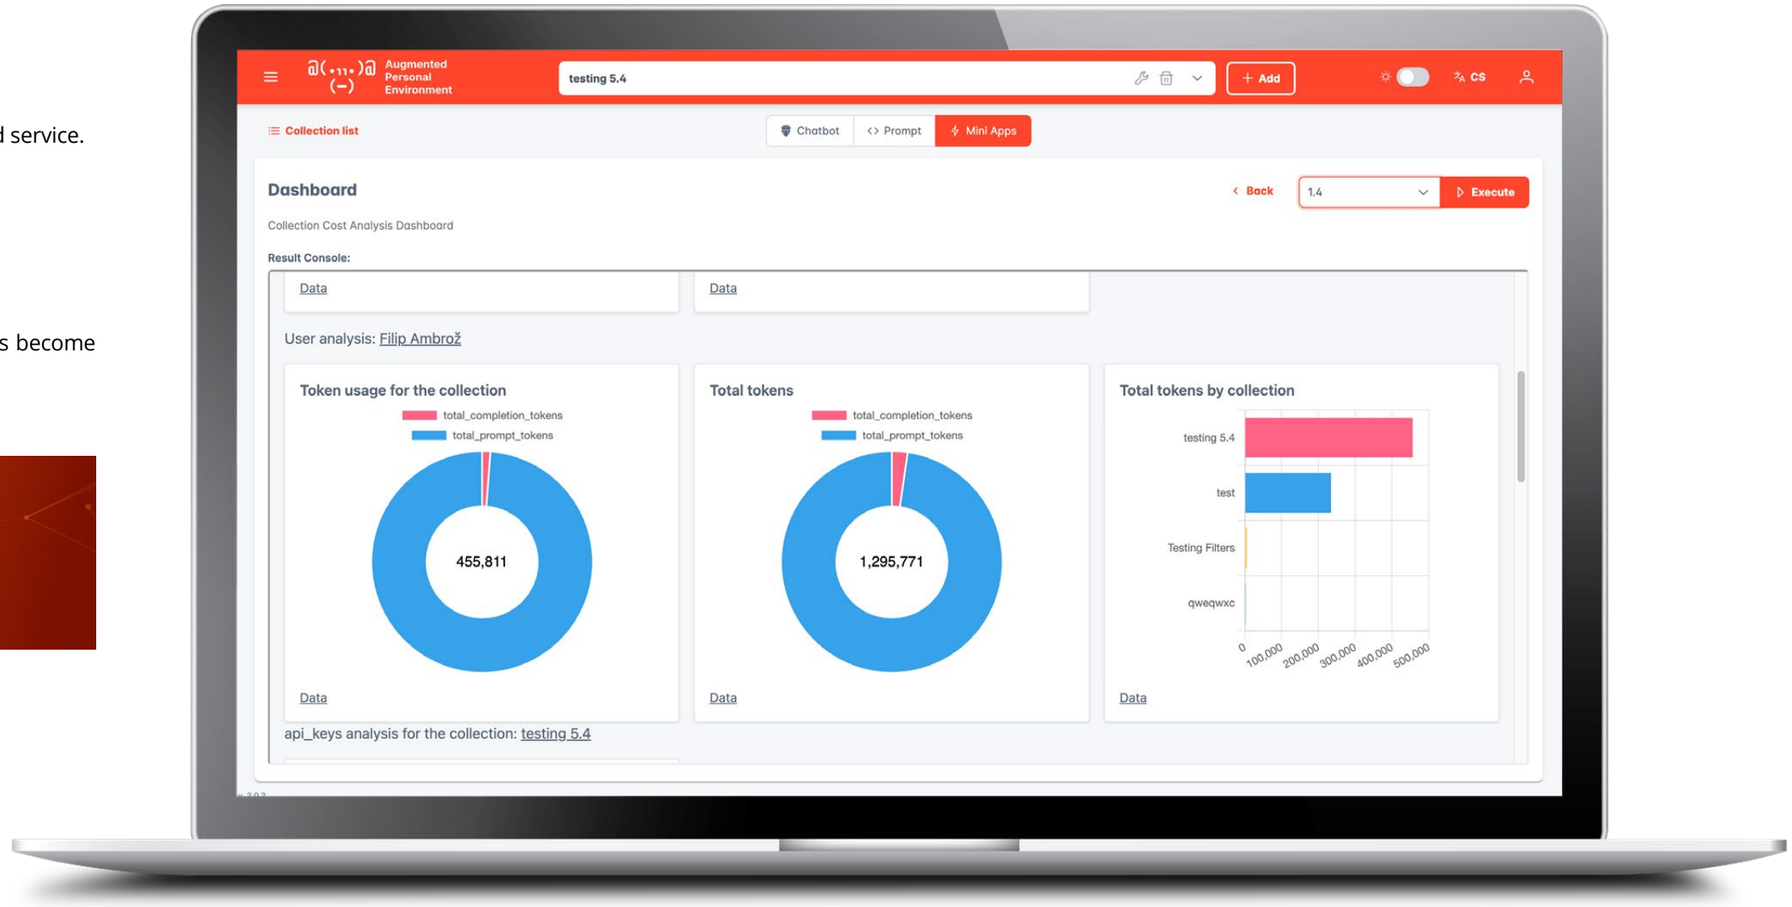

THE TEAM BEHIND THE PLATFORM

Twenty years of systems that couldn‘t afford to fail.

Physter Technology has been building and operating mission-critical systems since 2003. Working in regulated environments taught us that trust, accountability, and long-term reliability matter more than trends. APE is the result of that experience.

2003	50+	100+
Operating since	Engineers and consultants	Completed projects

Trusted by organizations where reliability is non-negotiable.

PHYSTER
TECHNOLOGY

Augmented
Personal
Environment

Ready to explore APE?

physter.com/APE

E-mail: info@physter.com

Augmented Personal Environment

AI built to be trusted