BUTTERFLY Brain

WHITE PAPER · ENTERPRISE MEMORY AND KNOWLEDGE LAYER

What your organization knows, in one memory.

Butterfly Brain is the enterprise memory of the Butterfly platform: it captures decisions, definitions, rules and the context behind past work, organizes them and brings them back at the right moment. Every Butterfly application reads from this layer and writes what it learns back to it.

Captures

Knowledge is collected without manual entry

Organizes

Duplicates are removed, decisions are tagged

Recalls

Ask in natural language, get the answer with its reasoning

Runs in-house

Knowledge never leaves your network

EXECUTIVE SUMMARY

Enterprise memory should not depend on individuals.

Decisions stay in chats, notes in files and context in people's heads. As the team grows, the same questions are asked again and the same work is rediscovered.

Any AI model can answer a general question. Only the organization's own memory knows its working definitions, the decisions that shaped them and the reasoning behind them. Butterfly Brain builds this memory as a knowledge layer.

Brain is not a chatbot or an assistant. It is the foundation layer that other Butterfly applications read from and write to. It captures knowledge while people and agents work, organizes it so it stays clean and serves it through an access layer that applications query for the right context at the right moment. The more the team uses Butterfly, the more Brain knows, and every application built on it benefits.

IN THIS DOCUMENT

- The problem of scattered knowledge and lost context
- Brain's three steps: capture, organize, recall
- One layer, many applications: the journey of a decision
- Applications built on Brain: Presentations Agent, Coding Agents, App Marketplace and SDK
- Which teams it suits
- On-premise deployment, model independence, human approval and frequently asked questions

OUTCOME

What the organization knows is gathered in a single, searchable, permission-based memory; applications and agents work with this context.

THE PROBLEM

Knowledge is scattered, context gets lost.

As the team grows, the same questions are asked again; if the answers are not recorded, every new person gathers the context from scratch.

Fragmented notes

Knowledge is spread across files, chats and people; nobody sees the whole picture.

Repeated work

People say "we solved this before" but there is no record; the work starts from zero.

Knowledge walking out the door

When someone leaves, the context leaves with them; enterprise memory erodes.

AI without context

A model that does not know the organization's rules gives a generic answer; the review burden comes back to the team.

THE ANSWER

Brain captures the knowledge the team produces in the background, organizes it with AI and turns it into a single memory that can be queried. Knowledge belongs to the organization, not to individuals.

HOW IT WORKS

A working memory in three steps.

Brain keeps decisions and their reasoning, the organization's definitions and rules and the context from past work, and links them all in a knowledge graph.

1 · CAPTURE

Automatic collection

Follows work sessions and enterprise knowledge in the background; captures knowledge without manual entry.

2 · ORGANISE

Organized with AI

Removes duplicates, summarizes and tags decisions; the memory stays clean and up to date.

3 · RECALL

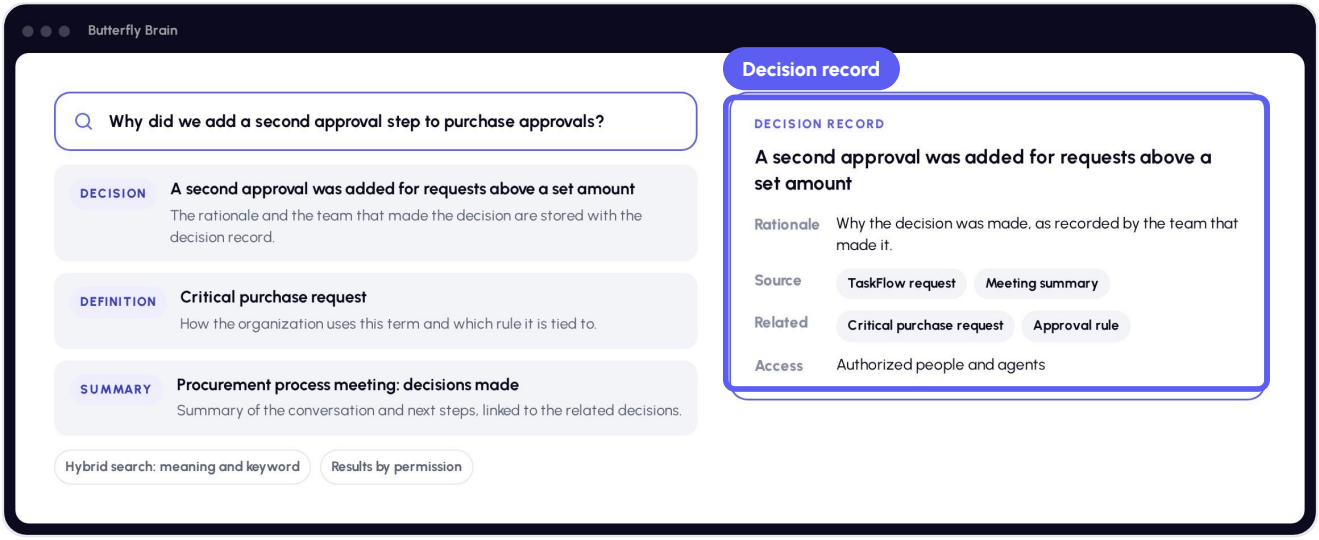
Instant retrieval

Ask in natural language; the relevant records, decisions and summaries come back ranked.

4 · SERVE

Context for every surface

Applications and agents query the right context at the right moment.



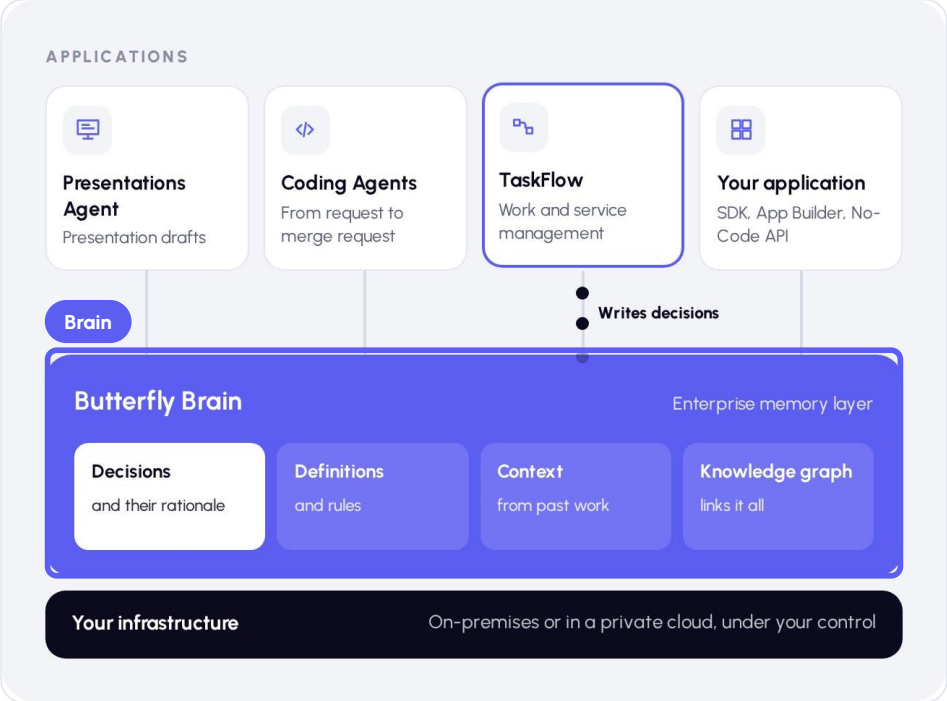
Search and decision record. The question is asked in natural language; decisions, definitions and summaries come back with their reasoning, source and access permissions.

Illustrative screen, sample data.

HOW IT FITS TOGETHER

One layer.
Many applications.

Butterfly is designed in layers, not as a single product. Brain sits beneath the applications and provides context to all of them; each application writes what it learns back to Brain. Everything runs on infrastructure you own.



Knowledge layer. Applications read context from Brain and write new knowledge back; Brain runs on the organization's own infrastructure.

The journey of a decision

STEP	WHAT HAPPENS
1 · TaskFlow	A second approval is added to purchase requests above a certain amount; the decision and its reasoning are written to Brain.
2 · Brain	The decision is linked to the related definition and earlier context; it becomes available to every person and agent with permission.
3 · Presentations Agent	When a presentation is prepared, the new approval rule goes into the slide as the organization's actual decision, not as generic text.

Sample flow; the steps are illustrative.

APPLICATIONS

Applications built on Brain.

Every Butterfly application is useful on its own; with Brain it does better. All of them run on your own infrastructure, work with the AI models you approve and always keep a person in the process.

Presentations Agent

Presentation agent

WHAT IT DOES

Point it at a document, a data set or a project; it prepares a structured, on-brand presentation draft.

WITH BRAIN

Slides reflect the organization's definitions, earlier decisions and tone. Nothing is shared without approval.

Coding Agents

Coding agents

WHAT IT DOES

Assign a request; the agent finds the right repository, makes the change and opens a merge request ready for review.

WITH BRAIN

It loads company context from Brain. Sessions are managed from iOS, Android and macOS over an end-to-end encrypted channel; a person always gives the approval.

App Marketplace and SDK

Your own applications

WHAT IT DOES

The Butterfly SDK offers native clients in eleven languages, including Swift, Kotlin, Go and Python. App Marketplace distributes internal applications with one-tap installation and over-the-air updates.

WITH BRAIN

In-app feedback turns into a work item together with a screenshot and a note.

ITS PLACE IN THE PLATFORM

Because Brain is a shared foundation layer rather than a feature of a single application, value builds up: knowledge captured while classifying a request shapes the next presentation, and a decision recorded in one workflow becomes available to every agent after it.

CAPABILITIES

Capabilities out of the box.

Brain runs on the Butterfly platform; management, permissions and audit come from the platform.

Semantic and keyword search

Hybrid retrieval: understands intent and does not miss exact word matches either.

Decision and summary tracking

"What did we decide, where were we?" is answered from one place.

iOS app

Memory in your pocket: search, read and share from mobile.

Access via MCP and REST

Accessed from within the tools in use; no need to switch to a separate screen.

Automatic organization

A scheduled AI worker merges duplicates and keeps the memory well maintained.

Presentation generation

A shareable slide deck is prepared from the memory.

Why it matters

Faster onboarding

New team members find the context without asking.

Knowledge loss stops

Critical context belongs to the organization, not to individuals; departures do not erase the memory.

The team looks at the same truth

A single source of truth; fewer repeated questions and conflicting decisions.

Better decisions

Past context is at hand instantly; decisions are based on records.

WHICH TEAMS

Teams that want to free enterprise memory from dependence on individuals.

Knowledge-intensive, fast-growing teams

Managers, project and operations teams

WHY

As the team grows, the same questions are asked again; if the answers are not recorded, every new person gathers the context from scratch.

WHAT IT DOES

Gathers decisions, definitions and meeting outcomes in a single searchable memory.

New team members

Human resources and team leads

WHY

If context lives in one person's head, onboarding depends on that person's time.

WHAT IT DOES

New joiners find past decisions and the definitions the organization uses without asking.

Software and IT teams

Engineering and IT management

WHY

If an agent writes code without knowing the organization's rules, the review burden comes back to the team.

WHAT IT DOES

Coding agents work with company context, and every change is merged with human approval.

Four questions to check the fit

TOPIC	QUESTION
Scattered knowledge	Where are decisions and their reasoning kept today: in a shared record, or in chats and e-mails?
Repeated work	Does it happen that someone says "we solved this before" and no record can be found?
Enterprise memory	When an employee leaves, is what they know recorded, or does it leave with them?
New joiners	Does a new joiner learn the organization's definitions and past decisions from a searchable source?

TRUST AND CONTROL

In-house, audited, model-independent.

The value of Brain is that it holds only what your organization knows. That is why it is designed to run on your own infrastructure, be isolated per organization and not be shared with third parties.

Control and oversight

- Human review and approval before anything is published
- Model-independent: use the agents and models you approve
- Knowledge isolated per organization
- Immutable audit log across applications

Deployment and data

- Runs on-premise or in your private cloud
- End-to-end encrypted channels for agent sessions
- Knowledge, code and prompts never leave your network
- No user-level data sharing with third parties

Personal data and KVKK

Meeting outcomes and correspondence may contain personal data. In an on-premise deployment, knowledge is processed and stored on the organization's own infrastructure; who has access and how long knowledge is kept are governed by the organization's own KVKK policy. The Butterfly platform provides identity, the permission model and the audit trail so that this policy can be enforced.

Certificates

CERTIFICATE	SCOPE
ISO/IEC 27001	Information security management system
ISO 9001	Quality management system
Domestic Product Certificate (Yerli Malı Belgesi)	Business process management software (Butterfly), 100% domestic content, class A, high technology
Technological Product Certificate	Butterfly DXP, Ministry of Industry and Technology of the Republic of Türkiye

All certificates and verification links: butterfly.dev/en/kurumsal/sertifikalar

FREQUENTLY ASKED QUESTIONS

Short answers about Brain.

Is Butterfly Brain a chatbot?

No. Brain is the enterprise knowledge layer that Butterfly applications read from and write to. It keeps decisions, definitions and context; applications and agents use this context.

Do we enter knowledge by hand?

There is no need. Brain captures knowledge from work sessions and enterprise sources in the background; AI removes duplicates and summarizes.

How does search work?

You ask in natural language. Hybrid retrieval, which combines semantic and keyword search, returns relevant records ranked by score. It is accessible via the web, the iOS app, MCP and REST.

Where does our data stay?

Brain runs on-premise or in your private cloud. Knowledge, code and prompts never leave your network and are not shared with third parties.

Which AI models does it work with?

Brain is model-independent; you can use the agents and models your organization approves.

How can we evaluate it?

The clearest way is a guided proof of concept (PoC) installed on your infrastructure. While Brain captures your context, Presentations Agent and the coding agents work on real material under your supervision. Our team supports installation, configuration and migration from your current tools.

See Brain and Butterfly applications in your own environment.

With a guided proof of concept installed on your infrastructure, let's assess together how Brain captures your organization's context.

Let's plan a proof of concept · butterfly.dev