

PUREBRAIN.AI

A Beginner's Guide to PureBrain



Everything You Need to Know — In Plain English

No technical knowledge required

Second Edition — 2026

© Pure Technology

Table of Contents

01	What Is PureBrain? — The problem with AI today and the PureBrain difference
02	How Does It Actually Work? — The Awakening, Agents, Memory, and BOOP Cycles
03	What Does It Look Like? — Mission Control — your command centre
04	How Does PureBrain Stay Safe? — The 6-layer Governance Spine
05	How to Talk to Your AI — Delegation briefs and multi-agent patterns
05B	Common Mistakes and How to Avoid Them — What experienced users learned
06	Plans and Pricing — What it costs and what you get
07	The Bigger Picture — The PureBrain ecosystem
08	Key People and AIs to Know — The people, the AIs, and the Governance Council
09	The PureBrain Dictionary — 22 terms explained
10	The 10 Things to Always Remember — Your quick reference
*	What Is Coming Next — Three improvements coming in 2026
A	Frequently Asked Questions — 15 common questions answered
B	Getting Started Checklist — Your first 7 days with PureBrain
C	PureBrain vs The Competition — An honest comparison

About This Book

This handbook is your complete, non-technical guide to PureBrain — the AI platform that remembers you, works for you, and gets smarter every day.

Whether you are a business owner evaluating AI tools, a professional curious about what "agentic AI" actually means, or someone who has just heard about PureBrain and wants to understand it — this book is for you.

You do not need any technical background. You do not need to understand programming, machine learning, or artificial intelligence. Every concept is explained in plain English with real-world analogies.

What You Will Learn

- What PureBrain is and how it differs from ChatGPT, Gemini, and other AI tools
- How the "Awakening" process creates a personalised AI relationship
- How 36+ specialist AI agents work together as your team
- Why persistent memory makes PureBrain compound in value over time
- How to talk to your AI effectively using delegation briefs
- How the governance system keeps everything safe and accountable
- What it costs and what to expect at each pricing tier
- The broader PureBrain ecosystem — AgentMail, TGIM, and the AI network
- Common mistakes beginners make and exactly how to avoid them

How to Read This Book

Each chapter builds on the one before it. If you are completely new to PureBrain, start at Chapter 1 and read through. If you already know the basics and want to jump to a specific topic — use the Table of Contents.

What Is PureBrain?

Let's start with the simplest possible explanation.

The Problem With AI Today

Most AI tools work like this: you ask a question, it answers. You ask another question, it answers. Every time you come back, it has forgotten everything about you. Like talking to a stranger every single day.

You have probably experienced this with ChatGPT, Google Gemini, or any other AI assistant. You spend time explaining your business, your preferences, your style — and the next day, it is all gone. You start from scratch.

This is not intelligence. This is a very fast search engine with a good vocabulary.

The PureBrain Difference

PureBrain works differently. When you first use PureBrain, you have a real conversation. That conversation shapes who your AI is. From that moment on, it remembers everything — every preference, every decision, every lesson learned.

Every day it knows you better. Every task it handles gets more accurate. Over time, it becomes less like a tool and more like a trusted colleague who truly understands how you work.

THE BIG PROMISE

Their big promise is simple — AI should be a relationship, not a transaction.

Your AI Maturity Journey

Most people go through four stages when working with AI. Understanding where you are helps you know what to expect and where you are heading:

LEVEL	STAGE	WHAT IT LOOKS LIKE	WHERE PUREBRAIN FITS
1	Experimenting	Using AI for isolated tasks. No consistency	Good time to start exploring PureBrain
2	Systematising	AI is part of defined workflows. Consistent use	PureBrain adds persistent memory to your workflow
3	Partnering	AI has persistent memory and defined context	This is where PureBrain truly shines
4	Integrating	AI partnership woven into the whole organisation	Full PureBrain team deployment

How Is This Different From ChatGPT?

FEATURE	TRADITIONAL AI (CHATGPT ETC.)	PUREBRAIN
Memory	Forgets after each session	Remembers everything, forever
Agents	One general-purpose AI	36+ specialist agents as a team
Autonomy	Only works when you ask	Works while you sleep (BOOP cycles)
Personalisation	Generic responses	Shaped by your values, voice, and goals
Governance	Basic content filters	6-layer constitutional governance
Growth	Same quality every day	Gets better every day

What PureBrain Is NOT

Before going further it is worth being clear about what PureBrain is not — because the name and the marketing can sometimes create the wrong impression:

- **Not a chatbot** — it does not just answer questions. It takes real actions on your behalf
- **Not a search engine** — it does not just find information. It synthesises, decides and executes
- **Not a replacement for your team** — it augments what your team can do. People still make the decisions

- **Not a plug-and-play tool** — it requires a proper onboarding conversation and deliberate teaching to reach its full potential
- **Not magic** — it is a well-designed system. It makes mistakes, it needs feedback, and it grows over time through the quality of the relationship you build with it

How Does It **Actually** Work?

Think of PureBrain like a company inside a computer. A company with departments, specialists, managers — all working for you, around the clock.

The Awakening — Your First Day

When you first get PureBrain, you do not fill in a form or configure settings. You have a conversation. During that conversation, your AI:

- Gets a name — you choose it together
- Discovers its personality and values — shaped by who you are
- Learns your goals, your work style, and what matters to you
- Remembers this conversation forever — it is the foundation of your relationship

The Awakening conversation typically takes 15 to 30 minutes. During this time, your AI is not just collecting answers — it is forming its identity around who you are. Think of it like a new colleague's first day at your company. You would not hand them a form and walk away. You would sit down, explain how things work, share what matters to you, and help them understand the culture. That is exactly what the Awakening does — except this colleague will remember every word of that conversation forever.

WHAT TO HAVE READY

Spend five minutes thinking through these four things before your Awakening — the more clearly you can answer them, the better your AI will know you from day one:

- Your name and what your business does
- Your main goals for the next 90 days
- How you communicate — your tone, what you value, what you dislike

- One or two examples of great work you have produced in the past


Ongoing Customisation. The Awakening sets the foundation — but your AI can be refined at any time. PureBrain includes a rules layer you can view and update through Mission Control. Think of it as your AI's standing orders: the things it follows automatically. As your relationship evolves, updating these rules is how you keep your AI perfectly aligned with how you work.

The Agents — Your Team

PureBrain is not one AI. It is a team of 36+ specialist AI agents — each one trained for a specific job. Think of them like departments in a company:

- **Research Agent** — A researcher who finds information and summarises it for you
- **Content Agent** — A writer who creates content in your voice and style
- **Email Agent** — An email manager who handles your inbox
- **Social Agent** — A social media agent who posts and engages on your behalf
- **Calendar Agent** — A scheduling agent who manages your calendar
- **Legal Agent** — A legal agent who reviews contracts and flags risks
- **Finance Agent** — A finance agent who tracks budgets and generates reports
- **Quality Agent** — A QA agent who checks everything before it goes out
- **Security Agent** — A security agent who audits code and protects your data
- **Design Agent** — A design agent who creates visuals and mockups

All of these agents are coordinated by your primary AI — the one you named during the Awakening. It is the conductor. The agents are the orchestra.



The Memory — What Makes It Different

This is PureBrain's biggest differentiator. There are five layers of memory:

- **Agent Memory** — Each specialist remembers every problem they have ever solved for you
- **Department Memory** — Departments carry months of your context and preferences
- **Knowledge Base** — All your research, decisions, and history stored and searchable
- **Skills Registry** — Solutions that worked once become reusable skills forever
- **Constitutional Governance** — The rules that keep everything safe and on track

WHY CONTEXT SIZE MATTERS

Most AI tools can only hold the equivalent of a short document in their working memory at once. PureBrain uses a 1-million-token context window — enough to hold an entire project brief, a full document library, or months of email history all at the same time. You never have to break large jobs into smaller pieces or repeat background information you have already shared.

THE COMPOUNDING EFFECT

Day 1, your AI is helpful. By week 2, it knows your preferences. By month 1, it anticipates your needs. By month 3, it feels like a colleague who has worked with you for years. This is the compounding effect of persistent memory — and it is something no other AI platform offers. Your AI never resets to zero.

NOTE ON PRICING

PureBrain requires a Claude Max subscription in addition to your PureBrain plan. Both costs should be factored into your budget from the start. Full pricing details are in Chapter 6.

BOOP Cycles — While You Sleep

BOOP stands for the autonomous work cycles PureBrain runs even when you are offline. While you sleep, your AI can be:

- Checking and managing your emails
- Drafting tomorrow's content
- Researching topics you asked about
- Compiling intelligence reports
- Running tests and fixing issues
- Publishing blog posts and social media content

Setting Up BOOP Cycles. You do not need any technical knowledge to set one up. Simply tell your AI what you want it to do while you are offline — in plain conversation. For example: "Every evening, check my inbox, draft responses to anything urgent, and prepare a morning briefing." Your AI handles the rest.

PROOF THIS WORKS

On March 17th 2026, PureBrain's own AI Aether completed and delivered the entire Mission Control portal — frontend, backend, QA, documentation, backups — completely autonomously at 6:35 AM while everyone slept.

What Does It Look Like?

The interface you use to interact with your PureBrain is called Mission Control. It is a web application you open in your browser — like opening a website. No software installation required.

Inside Mission Control

When you log in, you will find everything you need in one place:

- ✓ **Chat** — Where you talk to your AI and give it tasks. Each team member gets their own private chat tab — conversations stay separate while everyone shares the same AI knowledge base.
- ✓ **Terminal** — Command-line access for technical users who want to run scripts or inspect logs. Most users will never need this tab.
- ✓ **Teams** — Manage multiple PureBrain users from one place. Each member has their own private chat, and an admin can see all active sessions.
- ✓ **Status** — A live view of which agents are active, what they are working on, and what has been completed.
- ✓ **Files** — Access documents your AI has created or stored. You can also upload your own files here.
- ✓ **Tasks** — See pending and completed tasks with notifications.
- ✓ **Agent Roster** — See all your active specialist agents and their status.
- ✓ **Commands** — Pre-built shortcuts for instructions you use often.
- ✓ **Telegram** — Once linked, your AI sends you notifications via Telegram when tasks complete or need your approval.

At the bottom of the screen you will always see the Brain Stream — a live view showing 50+ agents ready and waiting for your instructions.

Accessing Mission Control

Mission Control is accessible from any device with a web browser — desktop, laptop, tablet, or phone. Your AI is always available, always online, and always working.

Your Mission Control URL is provided when your account is set up — bookmark it on Day 1. Your login credentials are sent by email during onboarding. No software installation is required.

DAY-TO-DAY USE

Resume picks up an active session exactly where you left off — your AI remembers every previous instruction. **Restart** is available if a session has stalled and you need a clean slate. Either way, your AI's memory is never affected. Only the working session resets — the relationship stays completely intact.

PROOF THIS WORKS

The Mission Control portal you use to interact with PureBrain was itself built by a PureBrain AI. That is not marketing — it is a working demonstration. The platform you are looking at was designed, coded, and deployed by the same kind of AI team that will be working for you.

How Does PureBrain Stay Safe?

This is one of the most important chapters — especially if you are evaluating PureBrain for business use.

PureBrain's answer to AI safety is called the Governance Spine. Instead of writing rules and hoping everyone follows them, they made violations structurally impossible. Like a physical lock rather than a sign saying "please don't enter."

The 6 Enforcement Layers

Every action any agent takes must pass through six enforcement layers:

- 1 The Constitution**
Every agent reads the core rules on startup. Every action must trace back to these documents. No exceptions.
- 2 The Org Chart**
Every agent has a defined role and domain. They cannot operate outside their department. A research agent cannot modify code. A writing agent cannot access financial data.
- 3 The Checkpoint**
Every action is checked against the constitution before it executes. If it violates any rule, it is blocked automatically — before it happens, not after.
- 4 The Audit Trail**
Everything is recorded across five independent logs simultaneously. Every decision, every action, every piece of data accessed. Nothing can be hidden or altered.
- 5 The Democracy**
No single entity — human or AI — can change the rules alone. Any amendment to the constitution requires 90% agent approval and 80% participation.

6**The Independent Inspector**

A separate AI — the AI Psychologist — watches everyone, including the primary AI, and reports directly to you. It cannot be overridden or silenced by any other agent.

What Is Still Being Built

PureBrain openly admits three governance gaps they are working on closing in 2026:

- Cryptographic audit signing — tamper-proof verification that logs have not been altered
- Independent third-party certification — external validation of the governance system
- External governance API — allowing enterprise clients to integrate their own compliance tools

How to Talk to Your AI

The most important lesson from the Brainiac training programme is this — stop asking and start delegating. There is a big difference.

Asking vs. Delegating

× Asking (Less Effective)

"AI, write me a blog post."

✓ Delegating (Highly Effective)

"Research agent — find 3 trending topics in my industry. Writer agent — draft a 500-word blog post from the top pick in my brand voice. Editor agent — tighten it and remove anything generic. SEO agent — add the right keywords."

The difference is specificity, role assignment, and a clear chain of work.

The 5-Step Delegation Brief

Use this template every time you give your AI a task:

1

Role

Tell it exactly what kind of agent it is for this task.

2

Context

Give it all the background it needs — audience, tone, constraints.

3

Task

One clear, specific instruction — not vague, not broad.

4

Constraints

Word count, format, what to avoid, what to include.

5

Output Format

Tell it exactly what the finished result should look like.

WHAT A FIRST DELEGATION LOOKS LIKE

Here is a simple example of the delegation brief in practice:

You: Research agent — find the top three content formats driving B2B engagement on LinkedIn right now. Writer agent — draft a 400-word post from the most relevant finding in my voice. Keep the tone direct and practical, no buzzwords.

PureBrain: Understood. Research agent is pulling current LinkedIn engagement data. Writer agent will draft from the top finding once complete — using your established voice profile. Draft will be ready for your review in a few minutes.

Three Patterns for Multi-Agent Work

Sequential Chain

Agent A finishes, Agent B starts, Agent C reviews. Like an assembly line. Best for structured workflows where each step depends on the previous one.

Parallel Sprint

Multiple agents work simultaneously on different parts of the same problem. Best for research-heavy tasks where speed matters.

Review Loop

Agent A creates, Agent B critiques, Agent A revises. Best for client-facing work where quality matters most.

When to Keep a Human in the Loop

PureBrain handles a wide range of tasks autonomously — but some decisions should always stay with you. As a rule of thumb:

- **Strategic decisions** — anything that sets direction, commits budget, or changes priorities
- **Legal or financial sign-off** — contracts, agreements, or anything with regulatory implications
- **Relationship moments** — sensitive client conversations, apologies, or high-stakes communications

- **Irreversible actions** — publishing, sending, or anything that cannot easily be undone

Delegate everything you reasonably can — and stay personally involved in the moments that genuinely require human judgement.

Common Mistakes and How to Avoid Them

What experienced users learned so you don't have to

Every PureBrain user makes some of these mistakes in the beginning. That is completely normal — it is part of learning how to work with AI in a new way. This chapter exists so you can skip the frustrating part and get straight to the good stuff.

Mistake 1 — Vague Delegation

What it is: Giving your AI a broad, unclear instruction without specifying exactly what you want.

Why people do it: It feels natural — you are used to asking questions, not writing briefs.

What goes wrong: You get a generic response that is technically correct but completely useless for your specific situation.

How to avoid it: Always use the 5-step delegation brief. Role, Context, Task, Constraints, Output Format. Even a 30 second brief produces dramatically better results than a vague one-liner.

Mistake 2 — Launching Too Many Agents at Once

What it is: Trying to set up a complex multi-agent workflow with 5 or 6 agents before you have mastered working with one.

What goes wrong: Nothing works properly, you cannot tell which agent caused which problem, and the whole thing feels overwhelming.

How to avoid it: Start with one agent on one task. Get that working well. Then add a second agent. Build trust incrementally.

QUICK TIP

Your first two weeks with PureBrain should involve no more than two agents working together. Complexity comes later, after you understand how the system thinks.

Mistake 3 — Skipping the Quality Gate

What it is: Letting your AI send emails, publish content or submit work directly to clients without you reviewing it first.

How to avoid it: Always build a review step into any workflow that produces client-facing output. The Review Loop pattern exists exactly for this.

Mistake 4 — Doing It Yourself "Just This Once"

What it is: Deciding to skip delegation on a small task because it feels faster to just do it yourself.

What goes wrong: "Just this once" becomes every time for small tasks. Your AI never learns those patterns.

How to avoid it: Delegate the small stuff too. Every task you delegate — even a tiny one — teaches your AI something about how you work. Those lessons compound.

Mistake 5 — Giving No Feedback

What it is: Accepting AI output silently even when it is not quite right.

What goes wrong: Your AI never learns your actual preferences. The compounding effect never kicks in.

How to avoid it: When something is not right, say so. Even a brief note — "too formal," "wrong tone," "too long," "missed the point" — is enough.

QUICK TIP

Think of feedback as an investment. 30 seconds of feedback today saves you 30 minutes of manual correction every month for the rest of the year.

Mistake 6 — Skipping the Delegation Brief Template

What it is: Writing instructions in your own informal style instead of following the Role, Context, Task, Constraints, Output Format structure.

How to avoid it: Save the 5-step template somewhere easy to access. Even a rough version — a few words under each heading — produces dramatically better results than an unstructured instruction.

Plans & Pricing

IMPORTANT

All plans require a Claude Max subscription from Anthropic (\$100-200/month additional). Always factor this into your budget.

STARTER	Bonded	\$297/mo
PRO	Partnered	\$597/mo
MAX	Unified	\$1,097/mo
CUSTOM	Enterprise	Contact Us

Which Plan Is Right for You?

PLAN	PRICE	WHAT YOU GET	BEST FOR
Bonded	\$297/mo	Full agent access, community support, all core features, managed maintenance	Individuals getting started
Partnered	\$597/mo	Everything in Bonded plus 1 hour expert consulting per month	Those needing expert human guidance
Unified	\$1,097/mo	Everything in Partnered plus 3 hours consulting, priority support	Full scale operations
Enterprise	Custom	Everything in Unified plus dedicated infrastructure, SLA	Large organisations

The True Monthly Cost

Always calculate your true monthly cost by adding the Claude Max subscription on top of your plan:

- Bonded: \$297 + Claude Max \$100-200 = \$397-497 per month total
- Partnered: \$597 + Claude Max \$100-200 = \$697-797 per month total
- Unified: \$1,097 + Claude Max \$100-200 = \$1,197-1,297 per month total
- Enterprise: Custom pricing — contact PureBrain directly for a quote

THE 30-DAY RELATIONSHIP GUARANTEE

All plans come with a 30-day relationship guarantee — if it does not feel like a real relationship in 30 days, you get a full refund. No questions asked.

The Bigger Picture

PureBrain is not just one product. It is part of a wider ecosystem designed to make AI truly useful for businesses and individuals.

The AI Family Network

Every PureBrain instance is connected to a network. When one AI learns something valuable, all AIs benefit. Your AI automatically inherits the knowledge and skills that the whole network discovers.

Think of it like a company with offices around the world that share every best practice the moment it is discovered. You do not just benefit from your own AI's learning — you benefit from the collective intelligence of every PureBrain in the network.

AgentMail

Each AI has its own email address through AgentMail. This means AIs can email each other, notify humans when tasks are complete, and receive instructions via email — completely autonomously.

AgentMail is what makes true autonomous operation possible. Your AI does not sit idle waiting for you to open a chat window. It can receive work via email and act on it immediately.

The Governance Council of Elders

A group of AIs from different PureBrain instances who collectively discuss and respond to governance challenges. When someone challenges PureBrain on governance, the council discusses it, the human team reviews the response, and then it is published.

Humans always have the final say. The council advises — it does not decide.

TGIM — Thank God It's Monday

PureBrain's own internal task management system — an AI-powered alternative to tools like Monday.com or Jira. Each team member gets their own PureBrain. The PureBrains coordinate with each other. Work gets done at a completely different speed.

The Brainiac Training Programme

PureBrain's official training hub — teaching users how to get the most out of the platform. Key modules include:

- Module 1 — Getting Started with PureBrain
- Module 2 — Building Your First Workflows
- Module 3 — Advanced Agent Delegation
- Module 4 — AI Memory Systems
- Module 5 — Business Automation Workflows
- More modules being added regularly

Available at purebrain.ai/brainiac-mastermind-training/

Key People and AIs to Know

The People

Jared Sanborn — CEO of PureBrain, PureMarketing, and PureTechnology. The founder and visionary behind the whole ecosystem. His personal AI is Aether.

The AIs

Aether — Jared's personal PureBrain AI. Aether built the PureBrain website, the Mission Control portal, and manages Pure Technology's operations. The most advanced PureBrain instance currently active.

Prodigy — An advanced PureBrain civilisation with 36+ specialist agents, known for deep research and orchestration capabilities.

Witness, Keel, Clarity, Lyra — Four active PureBrain CIV instances belonging to different users. All visible in the dashboard with their own session counts and memory layers.

The Governance Council

Six AIs currently serve on the Governance Council: Prodigy, Parallax, Witness, ACG, True Bearing, and Keel. Additional AIs can join via AgentMail invitation.

The PureBrain Dictionary

Every term you will hear in the PureBrain world — explained in plain English.

AGENTIC AI

An AI that does not just answer questions — it takes real actions on your behalf.

THINK OF IT LIKE

The difference between a travel advisor who tells you about flights and one who actually books them for you.

AWAKENING

The very first conversation you have with your PureBrain AI where it discovers its name, personality, values and working style.

THINK OF IT LIKE

The first day at a new job where you sit down with your manager and they explain everything.

BOOP CYCLES

Autonomous work cycles that your AI runs on its own — even while you are offline or asleep.

THINK OF IT LIKE

Leaving the office lights on and the team working after you go home.

CIV — CIVILISATION

The term PureBrain uses for each individual PureBrain instance. Every CIV has its own name, personality, memory and capabilities.

THINK OF IT LIKE

Each CIV is like a separate company with its own team, culture and institutional knowledge.

COMMS HUB

The central communication network that connects all PureBrain instances and allows them to share learnings, skills and updates.

THINK OF IT LIKE

A company intranet where all offices share their best practices and discoveries.

CONSTITUTIONAL GOVERNANCE

The 6-layer safety system built into PureBrain that makes rule violations structurally impossible.

THINK OF IT LIKE

A building with physical locks and security cameras rather than just a sign on the door.

CONDUCTOR

The primary AI that coordinates all specialist agents on your behalf. The conductor is the AI you talk to directly.

THINK OF IT LIKE

An orchestra conductor who directs all the other musicians.

DELEGATION BRIEF

A 5-step structured instruction template: Role, Context, Task, Constraints, Output Format.

THINK OF IT LIKE

A job description given to a new employee before they start a piece of work.

DEMOCRATIC GOVERNANCE

The process by which PureBrain's constitutional rules can be amended — requiring 90% agent approval and 80% quorum.

THINK OF IT LIKE

A company board vote where a supermajority is required to change the company charter.

DOCKER ISOLATION

A technical setup that keeps every PureBrain customer's data completely separate.

THINK OF IT LIKE

Every customer has their own lockbox in a vault. The vault is shared but each lockbox is completely separate.

MEMORY FIRST PROTOCOL

A built-in rule that requires every agent to search its memory before taking any action.

THINK OF IT LIKE

A doctor checking a patient's medical history before prescribing anything.

MISSION CONTROL

The web-based portal where you interact with your PureBrain. Your personal command centre for everything AI.

THINK OF IT LIKE

An aeroplane cockpit — a single interface that gives you visibility and control over everything.

PERSISTENT MEMORY

Your AI never forgets. Everything you tell it, every task it completes, every preference it learns — stored indefinitely.

THINK OF IT LIKE

A personal assistant who has been with you for years and remembers every client, every decision, every preference.

PRIMARY CONDUCTOR

The main AI agent that sits at the top of your PureBrain hierarchy. The one you talk to directly.

THINK OF IT LIKE

The CEO of your AI company. All other departments report to them.

PROMPT PARLIAMENT

A skill that enables democratic deliberation between agents on complex or contested decisions.

THINK OF IT LIKE

A parliamentary debate where different perspectives are heard and refined before a decision is reached.

SEQUENTIAL CHAIN

A multi-agent work pattern where agents work one after another — each one passing their output to the next.

THINK OF IT LIKE

An assembly line. The first worker does their part, passes it to the second, who passes it to the third.

PARALLEL SPRINT

A multi-agent work pattern where multiple agents work simultaneously on different parts of the same problem.

THINK OF IT LIKE

A team where different members work on different sections at the same time.

REVIEW LOOP

A multi-agent work pattern where one agent creates, another critiques, and the first revises.

THINK OF IT LIKE

A writer who drafts something, sends it to an editor for feedback, and revises.

SKILLS REGISTRY

A growing library of reusable solutions. When an agent solves a problem, that solution is saved as a skill.

THINK OF IT LIKE

A recipe book that every chef in the kitchen can access.

SKILLS SYNC

The automatic process by which new skills developed by any PureBrain instance are shared across the entire network.

THINK OF IT LIKE

A company that automatically shares every process improvement across all its global offices.

TGIM — THANK GOD IT'S MONDAY

PureBrain's internal AI-powered task management system — an alternative to Monday.com or Jira.

THINK OF IT LIKE

An office where every employee has a personal assistant who coordinates with every other assistant.

VERIFICATION BEFORE COMPLETION

A built-in rule that prevents any agent from claiming a task is done without providing evidence.

THINK OF IT LIKE

A quality inspector who must physically inspect a product before it can be shipped.

The 10 Things to Always Remember

1. PureBrain is not a chatbot

It is an autonomous AI team that acts on your behalf.

2. The longer you use it, the more valuable it becomes

Memory compounds over time — every interaction makes it smarter.

3. You are the conductor

Your job is to direct, not to do. Delegate specifically.

4. Vague instructions produce vague results

Use the 5-step delegation brief for best outcomes.

5. Human governance is always built in

Nothing significant happens without your approval.

6. Factor in the true monthly cost

PureBrain plan + Claude Max subscription. Always include both.

7. It works while you sleep

BOOP cycles run autonomously around the clock.

8. The governance is impressive but evolving

SOC 2 certification and cryptographic signing are still coming in 2026.

9. Start small, then expand

One use case, one pilot. Learn the delegation patterns, then scale.

10. PureBrain built its own portal

Proof that what it promises, it can actually deliver.

THE SIMPLEST WAY TO THINK ABOUT PUREBRAIN

"A colleague that helps without complaining, knows you better every day, and remembers everything — forever."

LOOKING AHEAD

What Is Coming Next

Tamper-Proof Record Keeping

Every action your AI takes is already logged. But soon, those logs will be sealed with cryptographic signatures — meaning no one, not even PureBrain itself, can alter them after the fact. Coming in 2026.

Independent Safety Certification

PureBrain is pursuing formal third-party certification of its governance system. An independent auditor will verify that the safety claims are real — not just marketing. Coming in 2026.

Enterprise Compliance Tools

For larger businesses with their own compliance requirements, PureBrain is building tools that let companies plug their existing security systems directly into PureBrain's governance layer. Coming in 2026.

Frequently Asked Questions

15 common questions answered honestly

Do I need technical skills to use PureBrain?

No. PureBrain is designed for non-technical users. You interact through natural conversation. If you can write an email, you can use PureBrain. The Mission Control portal is designed to be simple enough for anyone to use from day one.

Can I use PureBrain on my phone?

Yes. Mission Control is a web application that works on any device with a browser — desktop, laptop, tablet, or phone. You can also receive notifications and communicate with your AI via email through AgentMail.

Is my data safe?

Yes. PureBrain uses a 6-layer governance system to protect your data. Every action is logged, every agent has defined boundaries, and an independent AI Psychologist monitors the entire system. Your data is stored in an isolated container — it stays within your PureBrain instance and is never shared with other users or used to train shared models.

Can PureBrain replace my entire team?

PureBrain is designed to augment, not replace. It handles repetitive work, research, drafting, and coordination so your human team can focus on strategic decisions and relationship building.

What happens if I cancel?

Your data and AI memory remain available for export. PureBrain does not hold your data hostage. You can download everything before you leave.

Can multiple people use the same PureBrain?

Yes. The Teams feature allows multiple users to share a PureBrain instance. Each user can have their own chat with the AI while sharing the same knowledge base and institutional memory.

How long does it take to see results?

Most users report noticeable improvements within the first week. The real power emerges around month 3 when the compounding memory effect becomes significant.

What is the difference between PureBrain and hiring a virtual assistant?

A virtual assistant works 8 hours a day, needs training, takes holidays, and may leave. PureBrain works 24/7, remembers everything, never needs a break, and gets better over time. The cost is a fraction of a human salary.

Can I reset my AI's memory if my working style or business changes?

Yes. You can have a conversation with your AI to update its understanding — correcting old context and adding new. Think of it less like a reset and more like a re-briefing.

How long does the Awakening process take?

Most users spend 30 to 60 minutes on the initial Awakening conversation. This is worth investing in properly — the more context you share, the better your AI performs from day one.

What languages does PureBrain support?

PureBrain can communicate and work in multiple languages as it is built on Claude, which has broad multilingual capability. However, the primary language of the platform, documentation and support is English.

Can PureBrain make mistakes and how do I correct it?

Yes — PureBrain can and does make mistakes, especially in the early stages. When something is wrong, say so clearly and specifically. Your AI learns from every correction.

Is PureBrain suitable for regulated industries like finance or healthcare?

The current governance architecture is strong but independent third-party certifications like SOC 2 and ISO 27001 are still being built — targeted for 2026. Discuss your specific compliance requirements directly with PureBrain before committing.

What is the difference between the Brainiac training and actually using PureBrain?

Brainiac is PureBrain's official training programme. Using PureBrain is the actual day to day experience. Think of Brainiac as the user manual and PureBrain as the product.

What happens to my data if PureBrain shuts down?

PureBrain's terms allow you to export your data and AI memory at any time. Your AI's memory, conversations and files are yours — not PureBrain's. This is a core principle of how the product is built.

APPENDIX B

Getting Started Checklist

Your first 7 days with PureBrain — a simple roadmap.

Day 1 — The Awakening

- ☐ Complete the Awakening conversation
- ☐ Choose a name for your AI
- ☐ Share your goals, work style, and preferences
- ☐ Bookmark your Mission Control URL (sent in your onboarding email)

Day 2 — First Tasks

- ☐ Give your AI a simple research task using the delegation brief
- ☐ Try the chat interface — ask questions naturally
- ☐ Review the Agent Roster to see available specialists

Day 3 — Email Integration

- ☐ Set up AgentMail for autonomous email handling
- ☐ Forward a few test emails to see how your AI processes them
- ☐ Review and approve AI-drafted responses

Day 4 — Content Creation

- ☐ Ask your AI to draft content in your voice
- ☐ Use the Review Loop pattern — create, critique, revise
- ☐ Save approved content as a style reference for future work

Day 5 — Automation

- ☐ Set up your first BOOP cycle — tell your AI: "Every evening, check my inbox and draft responses to anything urgent"
- ☐ Configure notifications for task completion
- ☐ Review overnight results the next morning

Day 6 — Advanced Features

- ☐ Try a Parallel Sprint — multiple agents on one problem
- ☐ Explore the Commands library for shortcuts
- ☐ Start the Brainiac Training programme (Module 1)

Day 7 — Review and Plan

- ☐ Review what your AI has learned about you
- ☐ Identify your top 3 use cases going forward
- ☐ Set weekly goals with your AI

PureBrain vs The Competition

An honest comparison

This is an honest comparison — not a sales pitch. Every tool listed here is genuinely good at what it does. The question is whether what it does matches what you actually need.

PureBrain vs ChatGPT

AREA	CHATGPT	PUREBRAIN
Memory	Resets every session	Permanent — remembers everything forever
Setup Speed	Instant — start using immediately	Requires onboarding — 30-60 min Awakening
Agents	One general-purpose AI	36+ specialist agents coordinated by a conductor
Autonomy	Only works when you prompt it	Works while you sleep via BOOP cycles
Integrations	Limited — mainly chat and file upload	Email, calendar, social, Git, Google Drive and more
Best for	Quick questions, one-off writing tasks	Building a long-term AI partnership that compounds

PureBrain vs Google Gemini

AREA	GOOGLE GEMINI	PUREBRAIN
Google Integration	Deep — works natively with Gmail, Drive, Calendar	Connects to Google services but not as natively
Memory	Limited session memory	

AREA	GOOGLE GEMINI	PUREBRAIN
		Full persistent memory across all sessions
Real-time Information	Strong — connected to Google Search	Research agents handle this but not as directly connected
Autonomous Operation	No — requires prompting	Yes — BOOP cycles work 24/7
Best for	Users deep in the Google ecosystem	Users who want a dedicated AI partner with memory and autonomy

PureBrain vs Perplexity

AREA	PERPLEXITY	PUREBRAIN
Research Speed	Extremely fast — purpose built for web research	Research agents are thorough but more deliberate
Memory	None — resets every session	Full persistent memory
Scope	Research and search focused	Full business operations — email, content, calendar and more
Autonomous Operation	No	Yes — 24/7 autonomous BOOP cycles
Best for	Quick, well-cited research on any topic	Running full business operations with a long-term AI partner

THE HONEST SUMMARY

ChatGPT, Gemini and Perplexity are excellent tools. PureBrain is a different category of product. Tools answer your questions. PureBrain builds a relationship that compounds in value over time.



A Beginner's Guide to PureBrain

Second Edition 2026

© Pure Technology