V1.4 / 19-Oct-2026 / © InsideMind.ai / admin@insidemind.ai

A Safer Path to Enterprise AI

Turn organisational knowledge into immediate business value - without compromising security, control or accountability.



InsideMind.ai gives organisations a practical and controlled path to enterprise AI. It transforms approved internal knowledge into cited answers, decision support, proactive intelligence and governed automation while protecting confidential data, proprietary expertise and operational control.

Deployed on dedicated infrastructure, InsideMind can deliver immediate value through familiar, plain-English interaction without lengthy implementations, extensive staff retraining or dependence on public AI services.

Private AI Models • Permission-aware knowledge • Agentic monitoring • Internal app creation • Offline operation

EXECUTIVE OVERVIEW

Safely move beyond the AI hype

Every technology provider claims AI will transform business. The challenge for executive teams is determining where it can produce measurable value without creating unacceptable cyber, privacy, governance or operational risk.

Uncontrolled AI adoption can expose confidential information, distribute organisational knowledge across external platforms and introduce automated decisions that are difficult to govern or audit. But delaying adoption can also mean slower decisions, missed opportunities and competitors moving ahead.

InsideMind.ai provides a controlled path forward. It brings approved organisational knowledge, local AI, decision support, proactive monitoring and automation together in a private platform owned and governed by the organisation.

What InsideMind delivers

- **Immediate insight** - Staff receive cited answers, reports and decision support from approved organisational knowledge.
- **Protected knowledge** - Confidential data, intellectual property and proprietary expertise remain within an organisation-controlled environment.
- **Governed AI activity** - User and agent access is controlled through permissions, approval rules and clearly defined operating boundaries.
- **Auditable actions** - Data sources, tools, approvals, actions and outputs can be logged and reviewed.
- **Earlier awareness** - Background agents can monitor approved information for risks, exceptions, savings and commercial opportunities.
- **Faster adoption** - Staff use a familiar browser and plain English, minimising training, disruption and implementation time.
- **Long-term flexibility** - Models and capabilities can evolve without surrendering organisational knowledge or rebuilding around one AI provider.

The result is a practical enterprise AI capability that delivers value now, becomes more useful over time and remains under the organisation's control.

ENTERPRISE KNOWLEDGE

Turn Organisational Knowledge into Action

InsideMind.ai creates a governed knowledge centre from information the organisation already owns. Staff ask questions in plain English and receive answers grounded in approved internal sources, with citations that allow the evidence to be checked.

Capture information without creating more work

- **Simple contribution**. Upload files, use batch import, forward approved emails or monitor authorised shared folders and repositories.
- **Broad format support**. PDF, Word, Excel, PowerPoint, CSV, text, Markdown, HTML, images, OCR scans, email content, source code, JSON and XML.
- **Secure file ingestion**. Scheduled exports can be placed into a protected file-drop area for validation and ingestion without opening live access into core applications.
- **Automatic processing**. Metadata extraction, tagging, chunking, indexing, duplicate detection, version tracking and audit records keep the knowledge base usable.

A memory that explains what changed and why

The **Organisation Memory Timeline** connects historical documents, decisions, contracts, policies, projects, suppliers and customer relationships. Users can explore not only what the organisation knows now, but how that knowledge developed over time.

Knowledge capability	Business value
Policy and contract history	Compare versions, track changes and understand the source of current obligations.
Decision history	Retrieve supporting documents, meeting records and the context behind earlier decisions.
Customer and supplier history	Summarise relationships, proposals, issues and commercial activity across time.
Staff knowledge retention	Preserve institutional knowledge when roles change or experienced employees leave.

Enterprise Charter

An organisation-specific Enterprise Charter can provide the values, operating context, regulatory obligations, stakeholder relationships, ESG goals, culture and communication principles that guide the AI's responses and decision support. It helps InsideMind.ai interpret knowledge in the context of how the organisation actually operates and communicates.

FROM INSIGHT TO ACTION

Agentic assistance & monitoring

InsideMind.ai helps staff analyse organisational knowledge, prepare decision-ready outputs and continuously monitor approved data for material developments.

Everyday AI assistant

Capability	What it delivers
Questions and research	Answer organisation-specific questions, compare documents and trace conclusions back to source material.
Decision support	Prepare executive briefings, options, summaries and supporting evidence using the knowledge permissioned to the individual user.
Reports and writing	Generate reports, meeting preparation, policy explanations, translations and structured drafts.
Reasoning support	Explore relationships across documents and data while keeping a human responsible for consequential decisions.

Background agents and scheduled intelligence

Authorised agents can run approved checks and workflows on a schedule. They monitor selected business information, compare it with relevant organisational knowledge and alert the appropriate user group when attention may be required.

- **Opportunity detection**. Surface potential revenue, efficiency and cost-reduction opportunities.
- **Risk awareness**. Identify relevant changes, exceptions, gaps or emerging issues in approved data.
- **Targeted alerts**. Deliver evidence-backed notifications to the right teams rather than broadcasting sensitive information.
- **Governed action**. Respect tool permissions, data access, approval requirements and audit controls before any action is taken.

InsideMind.ai supports people rather than bypassing accountability: sensitive actions can require explicit approval, and every workflow can be scoped, logged and reviewed.

APPLICATIONS AND AUTOMATION

Build practical business tools around the way the organisation works

InsideMind provides a secure browser-based environment for building internal applications, dashboards and automated workflows around the organisation's own requirements, terminology and operating processes.

Internal Apps Workspace

Capability	What it delivers
Approval and operational workflows	Digitise repeatable processes with permissions, routing and status visibility.
Dashboards and reporting tools	Bring approved information into decision-focused operational views.
Employee and customer portals	Create secure interfaces for internal services or selected external use cases.
Forms and booking systems	Replace manual collection and coordination with structured applications.
Inventory, CRM and HR utilities	Develop focused tools that complement existing enterprise platforms.

AI Software Development Workspace

- **Develop and improve software**. Generate applications, analyse existing codebases, refactor code and create documentation.
- **Work across the full stack**. Support API development, database design, testing, debugging, Docker, Git and deployment automation.
- **Keep code under control**. Store source code securely and apply role-based access within the enterprise environment.
- **Choose the deployment destination**. Host internal applications on the appliance or deploy approved public-facing solutions to external platforms.

Workflow automation

A visual workflow layer can coordinate document processing, email workflows, scheduled jobs, approvals, notifications, file handling, report generation and controlled integrations. The organisation decides what the AI may do autonomously, what requires approval and what remains advisory only.

ENTERPRISE SECURITY

Enterprise AI Under Organisational Control

InsideMind.ai is deployed inside the enterprise environment on infrastructure dedicated to the organisation. Knowledge, models and processing can remain on-site, reducing exposure to public AI services and enabling operation with no mandatory cloud dependency.

Control layer	How it protects the organisation
Identity and access	Verified users, role-based permissions, department workspaces and user groups control who may access knowledge, tools and applications.
Data-depth control	Groups can receive permitted summaries or analysis without exposing restricted source records, identities or transaction-level detail.
Action governance	Tool permissions and approval rules control what users and agents may execute.
Encryption and infrastructure	HTTPS, SSH key authentication, firewall integration, encrypted backups and optional air-gapped deployment protect the environment.
Auditing and testing	Audit logs, controlled external audit access and managed cyber-security testing support assurance and regulatory review.
Support boundary	On-site installation and physical update options allow support without permanent remote access into the appliance.

Controlled knowledge ingestion

For high-security deployments, enterprise applications do not need live connections into InsideMind.ai. Systems can export approved files to a protected file-drop disk. The appliance validates and imports those files on a schedule, keeping knowledge current while maintaining a clear security boundary.

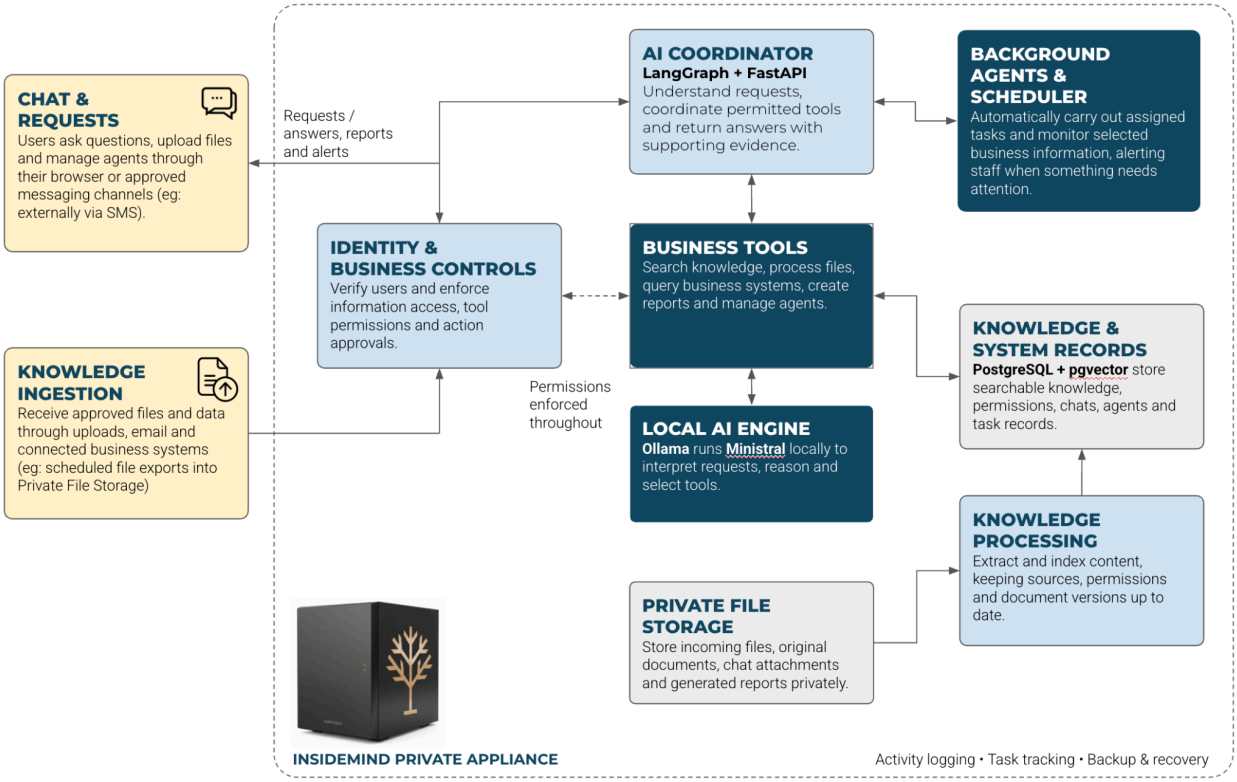
Offline operation with restricted external access

The appliance is designed to operate entirely within the enterprise intranet without an internet connection. Where required, optional GSM interfaces can provide restricted SMS or cellular access. Only responses expressly approved for those channels leave the secure environment; source documents and sensitive underlying records remain on-site.

TECHNICAL ARCHITECTURE

How InsideMind Works Securely

The platform separates user access, identity controls, AI reasoning, business tools, knowledge processing, storage and automation. Permissions are enforced across the full request path rather than applied only at the user interface.



How a request is handled

1. **Request**. A user asks a question, uploads a file or invokes an approved workflow through the browser or permitted messaging channel.
2. **Control**. Identity, information-access, tool and action permissions are checked.
3. **Coordination**. The AI coordinator interprets the request and selects only the permitted tools and knowledge sources.
4. **Local intelligence**. Local models reason over the request while business tools search knowledge, process files, query approved data and prepare outputs.
5. **Evidence and record**. The response returns with supporting evidence where available, while activity, task and audit records are retained.

The current reference architecture uses a local model runtime, an orchestration layer, PostgreSQL with vector search, private file storage and scheduled background agents.

DEPARTMENT VALUE

A governed platform, enterprise-wide

InsideMind.ai gives each team access to the depth of information and tools it is authorised to use, while allowing the organisation to apply one consistent governance model across departments.

Team	Representative outcomes
Executive leadership	Decision briefs, organisation-wide context, risk and opportunity alerts, and rapid access to supporting evidence.
Finance	Contract and report analysis, spending patterns, cost opportunities, policy interpretation and recurring management reporting.
Operations	Procedure lookup, maintenance knowledge, incident history, operational dashboards and workflow automation.
Sales and customer teams	Proposal history, customer context, pricing knowledge, account preparation and opportunity identification.
Human resources	Policy access, onboarding, training knowledge, controlled employee support and consistent process guidance.
Legal risk and compliance	Version-aware policy and contract search, evidence retrieval, audit support and monitoring of approved obligations.
Technology teams	Local AI infrastructure, code assistance, internal app development, workflow integration and controlled model management.

Across the organisation, InsideMind reduces time spent finding information, improves the consistency of decisions, identifies risks and opportunities earlier, and preserves institutional knowledge as a permanent enterprise asset.

RESILIENCE AND CONTINUITY

Keep knowledge available through disruption

Local processing reduces reliance on internet connectivity and external cloud services. Backup, recovery and continuity options are designed to preserve both the underlying knowledge and the ability to use it when normal operating conditions are interrupted.

Built-in recovery capability

Capability	What it delivers
Automated and incremental backup	Protect configuration, knowledge, models and applications according to scheduled policies.
Snapshots and verification	Support point-in-time recovery and confirm that backups remain usable.
Appliance recovery	Restore the operating environment, knowledge centre and internal applications following hardware or site disruption.
UPS and resilient hardware options	Tailor storage, RAID, power protection, GPU capacity and cluster design to operational requirements.

InsideMind Knowledge Vault New Zealand

For organisations requiring an exceptional level of geopolitical and infrastructure resilience, InsideMind.ai proposes an optional sovereign continuity service based in New Zealand. Encrypted backup media can be physically collected, transported under a verifiable chain of custody and stored in a secure facility geographically separated from the primary operating region.

- **Verified physical custody**. Collections, transfers and storage events are documented for client verification.
- **Rapid recovery path**. A replacement appliance at the New Zealand facility can restore the latest verified backup under the client's existing access restrictions.
- **Alternative access**. Authorised users may reconnect through secure VPN where internet service remains available, with cellular data and restricted SMS options where supported.
- **Knowledge preservation**. The service is intended to preserve the enterprise's accumulated knowledge and restore a restricted AI capability after exceptional disruption.

Important: The Knowledge Vault is a proposed optional continuity service. Its final scope, facility controls, jurisdictions, service levels and assurance framework would be defined with participating clients and specialist advisers.

DEPLOYMENT AND OPERATION



A turnkey appliance sized to the organisation

InsideMind.ai can be delivered as a workstation, rack-mounted server, multi-GPU system or supported cluster. Hardware, storage, model selection and operating controls are tailored to user numbers, document volumes, required response times, application workloads and resilience objectives.

Included platform capabilities

- Secure browser interface
- Enterprise knowledge platform
- AI assistant and model routing
- Software development and Apps Workspaces
- Workflow automation and background agents
- Administration and monitoring console
- Backup and recovery

Hardware options

- Enterprise SSD and RAID
- Multi-GPU configurations
- High-capacity storage
- UPS integration
- Cluster support

Implementation path

Stage	Purpose
1 Discovery and governance	Define priority use cases, information boundaries, the Enterprise Charter, user groups, approvals and success measures.
2 Sizing and solution design	Confirm hardware, model, storage, ingestion, security, integration, backup and deployment requirements.
3 Installation and knowledge migration	Install on-site, configure controls and ingest an agreed initial knowledge set.
4 Pilot and validation	Test priority workflows with selected teams, tune results and verify security, evidence quality and user experience.
5 Adoption and expansion	Train users, introduce approved agents and applications, and extend the knowledge environment in controlled stages.

SUPPORT AND LIFECYCLE

Control without being left to manage complexity alone

InsideMind.ai can be operated by the organisation's technology team, supported jointly, or delivered with a tailored managed service. The support model is designed around the security posture of each deployment.

Capability	What it delivers
Installation and migration	On-site deployment, data migration, initial configuration and user setup.
Training and adoption	Administrator onboarding, staff training and use-case development.
Software and model lifecycle	Software updates, model upgrades, performance optimisation and industry-specific data refreshes.
Monitoring	System health, storage, CPU, GPU, services and backup status through the administration console.
Secure service options	Physical support and update processes for environments that do not permit persistent remote access.
Enterprise support agreements	Tailored coverage, response arrangements and operating responsibilities.

Move from interest to a defined enterprise use case

InsideMind works with business, technology, security and governance leaders to identify where private AI can deliver the greatest immediate value.

The initial discovery process defines priority use cases, information boundaries, security requirements, agent permissions, hardware needs and measurable pilot outcomes. InsideMind then delivers a controlled implementation that can be tested with selected teams before wider deployment.

Own the infrastructure. Protect the knowledge. Put organisational intelligence to work.