


PRIVATE AI · YOUR DATA · YOUR INFRASTRUCTURE

Your organisation knows the answer. Your employees just cannot find it.

Esii CaptiveAI is an enterprise conversational AI platform that connects to your existing data sources and answers questions from your own documents, databases and emails. It runs on your infrastructure. Your data stays private. No public AI service receives your confidential information. This is enterprise AI on your terms. Your infrastructure, Your rules.

CONNECTS TO YOUR EXISTING DATA SOURCES



Connectors are added on request.

If your data source is not listed, speak to us.

THE PROBLEM

Enterprises have invested years in building knowledge and data assets. The problem is that employees cannot access this knowledge when they need it and traditional tools are not solving them fast enough.

01 Public AI tools like ChatGPT cannot be used with confidential enterprise data.

Employees have already started copying internal documents into ChatGPT to get summaries and answers. This is a data security and compliance risk that most IT and legal teams are aware of but have not yet found a way to address. The need for AI-assisted knowledge access is real. The question is how to meet it without putting sensitive data in a public service.

02 The answers exist somewhere in SharePoint, email or a database. Nobody knows exactly where.

A new employee asks a policy question. They search SharePoint and get fourteen results. They forward the question to three colleagues. Two colleagues reply with conflicting answers from different versions of the same document. The correct answer exists in the organisation. The problem is that finding it takes longer than the work itself.

03 Data in traditional structured databases is only accessible to people who know how to write queries.

A sales manager wants to know which products had the highest return rate in Q3 from a specific region. This answer is in the ERP database. But extracting it requires writing a SQL query, which the sales manager cannot do. They raise a request to IT or analytics. The answer arrives two days later. By then the context has changed.

04 Scanned documents and older files are outside the reach of any search tool.

A significant portion of enterprise knowledge exists in scanned contracts, printed circulars that have been digitised and older PDFs that were never processed for text extraction. These files are stored but not searchable. The information they contain is effectively inaccessible to anyone who does not already know it is there.

05 AI adoption is fragmented and ungoverned across the organisation.

Different teams have adopted different AI tools. One team uses ChatGPT. Another uses Copilot. A third has built its own workaround. There is no single, governed AI experience. There is no audit trail across tools. Management has no visibility into what questions are being asked, what data is being accessed, or whether the answers employees are acting on are accurate.

WHAT Esii CAPTIVEAI DOES

CaptiveAI connects to your organisation's existing data sources and makes that knowledge accessible through a simple conversational interface, in plain English.

Your employees type a question. CaptiveAI searches your connected data sources, retrieves the most relevant content and generates a clear, accurate answer, with a citation pointing to the exact document or database record it drew from. The entire process happens within your own infrastructure.

Retrieval-Augmented Generation

CaptiveAI does not generate answers from general knowledge. Every response is grounded in your own documents and data, retrieved in real time from the sources you have connected. Hallucination risk is minimised because the model is constrained to your actual content.

Natural Language to SQL

Employees can ask questions about data in SQL Server and MySQL databases in plain English. CaptiveAI translates the question into a valid SQL query, executes it against the connected database and returns the result as a readable answer. No query-writing skill is required.

OCR and AI on Scanned Documents

Scanned PDFs, digitised circulars and image-based documents are processed through an OCR and AI pipeline that extracts, structures and indexes the content. These files become fully searchable and answerable through the CaptiveAI interface alongside all other connected sources.

Private Deployment on Your Infrastructure

CaptiveAI is deployed on your own servers or on a private cloud instance. No query, no document and no answer passes through a public AI service. Your data governance, data residency and compliance requirements are fully maintained.

Voice Interaction

Employees can speak their question and receive a spoken response. Esii Captive AI supports voice input and output, making the platform accessible for multilingual workforces, field teams, and mobile users across locations.

Document and Image Generation

Employees can ask for a report, a presentation, a spreadsheet, a data chart, or an AI-generated image in plain English. Esii Captive AI generates a finished, ready-to-use file drawn from your connected data and delivers it in seconds. Outputs include PDF, Word, PowerPoint, Excel, charts, and AI-generated visuals.

Any AI Model, Your Choice

Fully Private Deployment
No data leaves your
Environment

Plain English Queries
No SQL.
No search syntax.

Source-Cited Answers
Every answer shows
its source

THE CRITICAL DIFFERENCE FROM PUBLIC AI TOOLS

ChatGPT and similar tools are trained on public data and generate answers from general knowledge. CaptiveAI generates answers exclusively from your organisation's own data, deployed entirely within your own environment.

Categories	Public AI (ChatGPT, Copilot)	Esii CaptiveAI
AI Model Choice	Typically limited to models provided by the service	Choose from OpenAI GPT, Anthropic Claude, Google Gemini, Meta Llama, xAI Grok, Microsoft Phi, open source models, or bring your own model. Deploy in your private cloud or on premises.
Where Your Data Goes	Requests are processed by the provider's cloud service	Requests are processed within your private infrastructure or enterprise cloud subscription. Data never passes through a public AI endpoint
Model Training	Depends on the provider and enterprise plan	Your prompts and enterprise data are never used to train foundation models, whether deployed on premises or in your private cloud.

Esii CAPTIVEAI - TWO LEVELS OF AI. ONE SEAMLESS EXPERIENCE

Capability	Lightweight AI (Small Language Models)	Advanced AI (Large Language Models)
Document Q&A	✓	✓
Document Summarisation	✓	✓
Knowledge Base Search	✓	✓
Email Drafting	✓	✓
Text Rewriting	✓	✓
Prompt Library	✓	✓
Natural Language Query to Database	✓ (Limited)	✓
Multiple AI Model Selection	✓	✓
Presentation & Report Generation	✓ Not Applicable	✓
Spreadsheet & Chart Generation	✓ Not Applicable	✓
AI Image Generation	✓ Not Applicable	✓
Voice Interaction	✓ Not Applicable	✓
Candidate Screening	✓ Not Applicable	✓
Task Automation Agents	✓ Not Applicable	✓
Agent Builder	✓ Not Applicable	✓ Coming Soon

Esii CaptiveAI operates within your organization's data governance framework.

HOW IT WORKS

From connecting your first data source to answering your first question, CaptiveAI follows a structured process that keeps your data secure and your answers accurate at every step.

1

Your data sources are connected to CaptiveAI through configured connectors

The CaptiveAI administrator connects the relevant data sources: SharePoint libraries, OneDrive folders, SQL Server or MySQL databases, Azure Blob Storage, email accounts (Outlook and Gmail), Amazon S3 buckets and scanned document repositories. Each connector is configured with the appropriate credentials and access scope.

2

Documents and data are indexed and vectorised within your environment

CaptiveAI processes all connected content, including scanned documents through the OCR and AI pipeline, and creates a vector index that allows semantic search across all sources simultaneously. This index lives entirely within your deployment. The AI model used for inference is selected by the organisation from cloud providers including Azure OpenAI, OpenAI, Anthropic, and Google, or from open-source models hosted entirely on your own hardware.

3

An employee asks a question in plain English through the CaptiveAI interface

The employee types their question into the CaptiveAI chat interface, which is accessible through a browser. The interface is designed to feel like a standard chat application. No special syntax, search terms or query language is required. For database questions, the employee simply describes what they want to know. For teams that are new to AI, the built-in guided prompt library for each business function gives employees a ready starting point, pre-written scenarios covering the most common needs of their role, available from the interface from day one.

4

CaptiveAI retrieves the most relevant content from your indexed sources

The question is converted to a vector query and matched against the index. The most relevant document sections, database records or email content are retrieved. For SQL database questions, the natural language query is translated into a valid SQL statement and executed against the connected database. The retrieved content forms the factual basis of the answer.

5

A clear, cited answer is generated from your own data and displayed to the employee

The AI generates a readable answer based only on the retrieved content. The answer includes a citation linking to the specific document, section or database record it drew from. The employee can verify the source directly. If the information does not exist in the connected sources, CaptiveAI says so, rather than generating a plausible-sounding but unsupported answer.

6

Every query and response is logged in the admin dashboard for governance

All interactions are recorded in the admin log with the user's identity, the question asked, the sources accessed and the response generated. Usage analytics, access patterns and query volumes are available from the dashboard. This record supports compliance reviews and helps administrators understand how the system is being used across the organisation.

7

Where configured, AI agents execute multi-step tasks end-to-end (In Active Development)

ESII Captive AI is being extended to go beyond answering questions to executing complete business workflows autonomously. An HR agent sources candidates, screens applications, and shortlists the strongest fits. A sales agent qualifies incoming leads, drafts personalised outreach, and books meetings without manual follow-up. A marketing agent segments audiences and monitors campaign performance in real time. Each agent connects to your existing business systems through purpose-built integration connectors, executes every step of the workflow, and escalates to a human only when a genuine decision is required.

KEY CAPABILITIES

CaptiveAI is built to handle the breadth of data formats and sources that a large enterprise actually works with, not just structured text documents.

Multi-Source RAG

Retrieval-Augmented Generation across all connected sources simultaneously. A question can draw answers from SharePoint, an email thread and a database record.

Natural Language to SQL

Questions about data in SQL Server and MySQL are automatically translated into valid SQL queries. The query is executed and the result is returned as a readable answer. No query-writing skill is required from the employee.

OCR and AI Fusion

Scanned PDFs and image-based documents are processed through an OCR pipeline before AI extraction. All text in scanned files becomes part of the searchable knowledge base, including older documents that were never text-indexed.

Source Citation in Every Answer

Every response includes a direct citation pointing to the document, section or database record that was used to generate it. Employees can verify the answer at source. There is no ambiguity about where the information came from.

Azure AD Single Sign-On

Employees log in to CaptiveAI using their existing Azure Active Directory credentials. No new passwords or accounts are required. Authentication is handled by the same identity system the organisation already manages.

Conversation History

CaptiveAI maintains conversation context across a session. Follow-up questions are understood in the context of earlier turns, allowing employees to drill down into a topic through a natural conversation rather than restating context with each question.

Incremental Index Updates

When new documents are added to connected sources or existing ones are updated, the index is refreshed incrementally. The knowledge base reflects current content without requiring a full re-indexing process.

Admin Dashboard and Analytics

Administrators have a full view of usage patterns, frequently asked questions, source access frequencies and user activity. These insights help identify knowledge gaps and inform decisions about which data sources to connect or expand.

RBAC-Enforced Access

Users can only receive answers from data sources they are authorised to access. CaptiveAI enforces the same access rules that apply to the underlying systems. A finance team member cannot receive answers derived from HR-restricted documents.

Any AI Model, Your Choice

As the AI landscape evolves, your platform evolves with it. ESII Captive AI connects to any model, including Anthropic Claude, OpenAI GPT, Google Gemini, Meta Llama, xAI Grok, Microsoft Phi, or any open-source alternative, running in the cloud or hosted entirely on your own servers. Switch models as better options become available.

Document, Report and Image Generation

Ask for a report, a presentation, a spreadsheet, a data chart, or an AI-generated image. ESII Captive AI generates finished, ready-to-use files and visuals drawn from your connected data and your prompt, delivered in moments. Documents are produced in standard business formats. Images are generated through AI APIs based on your description or your data context.

Agentic Workflows (In Active Development)

Beyond answering questions, ESII Captive AI can be configured to execute complete business workflows end to end, autonomously. Specialist agents for HR, Sales, and Marketing connect to your existing business applications through purpose-built integration connectors, execute every step of the workflow, and escalate to a human only when a genuine decision is required.


Built-In Question Library for Every Department

Most employees do not know what to ask an AI tool on their first day, or their fiftieth. ESII Captive AI removes this barrier entirely. The platform includes a ready-built library of guided question sets organised by business function: HR, Finance, Sales, Operations, Legal, Customer Service, and more. Each set contains the most common, high-value scenarios for that team, pre-written and ready to use. New functions and scenarios can be added at any time as the organisation's needs evolve.

Speak to us about your priority use case.


SUPPORTED DATA SOURCES

CaptiveAI connects to the platforms and storage systems your organisation already relies on. There is no requirement to migrate or restructure existing data.



DOCUMENT FORMATS SUPPORTED

CaptiveAI processes documents across all common enterprise formats without requiring any pre-processing or reformatting.



AUTHENTICATION

Azure Active Directory (SSO)

Employees log in with their existing M365 credentials. No separate account setup required. User identity is verified through Azure AD on every session.

RBAC Inherited from Source Systems

Access permissions from SharePoint, OneDrive and other connected systems are respected. Users cannot receive answers from sources they are not permitted to access.

WHERE CAPTIVEAI MAKES A MEASURABLE DIFFERENCE

The departments that benefit most are those where employees regularly need answers from documents, policies or databases, and where the cost of searching manually is high.

HUMAN RESOURCES

Policy questions answered instantly from the HR handbook and circular archive

Employees ask questions about leave policies, travel reimbursement limits, appraisal timelines and other HR matters through CaptiveAI. The answers are drawn directly from the current HR policy documents in SharePoint. The HR team handles fewer routine queries and employees get accurate answers without waiting for a callback.

SALES AND BUSINESS DEVELOPMENT

Sales teams query CRM and ERP data in plain English to get business insights without raising an IT request

A sales manager asks which accounts have not been contacted in the past 60 days, or which product lines had the highest margin in the previous quarter. CaptiveAI translates these questions into SQL queries against the connected database and returns the results as a readable answer, with no IT involvement required.

LEGAL AND COMPLIANCE

Contract clauses and regulatory obligations retrieved from a large document archive without manual search

Legal teams ask CaptiveAI to find all contracts that contain a specific liability clause, or to summarise the termination conditions across a portfolio of vendor agreements. The answers are drawn from the connected contract repository, with citations to the specific document and clause. Scanned legacy contracts are included through OCR processing.

IT HELPDESK

First-line resolution for common IT queries by connecting CaptiveAI to the knowledge base and ticket history

Common IT support questions, such as VPN configuration steps, software installation procedures and access request processes, are answered directly from the IT knowledge base and past ticket resolutions. First-line agents resolve more queries without escalation, and the response time for common questions drops significantly.

OPERATIONS AND QUALITY

Field teams and plant managers access SOPs and compliance documents through a conversational interface

Operations staff ask questions about process specifications, safety procedures and quality standards from a mobile browser. CaptiveAI retrieves the relevant section from the connected SOP library and displays it with a citation. This is faster and more reliable than navigating a shared drive folder structure from a mobile device.

FINANCE AND ACCOUNTS

Finance teams query transaction data and generate summaries without needing SQL skills or analyst support

Finance managers ask questions like "what was the total vendor spend in Q3 by category" or "show me all invoices above a specified amount that are pending approval". CaptiveAI translates these into database queries, executes them and presents the results in plain English. This reduces the dependency on the analytics team for routine data queries.

SECURITY AND DATA GOVERNANCE

CaptiveAI is built for enterprises where data residency, access control and audit compliance are non-negotiable requirements.

Your Data Never Reaches a Public AI Service

ESII Captive AI processes all queries on your own infrastructure or through your private enterprise cloud subscription on Azure. No query, document, or answer passes through a public AI endpoint.

RBAC Enforced at Every Query

Access permissions from source systems are respected by CaptiveAI. Users can only receive answers from data they are authorised to see. There is no way to use CaptiveAI to access restricted information.

Full Query and Response Logging

Every question asked and every answer generated is logged with the user's identity, the timestamp, the sources accessed and the content of the response. These logs are available for compliance and security reviews.

No Model Training on Your Data, Ever

Whether ESII Captive AI is running on your own servers or calling private enterprise APIs, your data is never used to train, fine-tune, or improve any AI model. The models perform inference only, they read and respond, they do not learn or retain.

Token and Usage Transparency

Every AI interaction records token consumption by user, team, and data source. Administrators have full visibility into usage volumes, costs, and query patterns, enabling governance, budgeting, and optimisation across the organisation.

DEPLOYMENT OPTIONS

CaptiveAI can be deployed in three configurations depending on your infrastructure preferences and data governance requirements.



On-Premise

Deployed on your own servers within your data centre. Fully air-gapped from external networks if required. Recommended for organisations with strict data residency policies or classified data environments.



Hybrid

Core processing on-premise with selected cloud connectors for specific data sources. Suitable for organisations with mixed infrastructure where some data is on-premise and some is in cloud storage.



Private Cloud (Azure/AWS/GCP)

Deployed on a dedicated private instance within your Microsoft Azure, AWS, or Google Cloud subscription, with data residency maintained in your chosen region and consistent RBAC enforcement across all deployment options, making it suitable for organisations regardless of their preferred cloud infrastructure.

COMPLIANCE ALIGNMENT

Azure AD Authentication

ISO 27001 Aligned

Data Residency Controls

RBAC Enforcement

Full Audit Logging

No External Data Transmission

No Model Training on Customer Data

See CaptiveAI answering questions from your own data.

We connect a demonstration instance to a sample of your data and run a live question-and-answer session.

[Book a Demo →](#)

esii.in
info@esii.in
A PC Solutions IP Brand
New Delhi, India