



BOHIOSYSTEM

• SERVICE DATASHEET • BOHIOSYSTEM

Artificial Intelligence

We apply AI where it produces a measurable result: assistants, document extraction, classification, and decision support, with human control where it matters.

DOCUMENT

BS-SVC-09

VERSION

1.0

DATE

August 1, 2026

PREPARED BY

BohioSystem

BohioSystem

bohiosystem.dev • Confidential

THE SERVICE

The useful question isn't "Where do we use AI?" It's "What task consumes our hours and tolerates a margin of error?" That's where a project worth doing begins.

We design and implement artificial intelligence solutions for real processes: assistants that answer using the company's documentation, data extraction from documents, automatic classification, and decision-support systems.

We work with commercial and open models depending on the case, and we always define what happens when the model gets it wrong. An AI system with no plan for error isn't finished.

Measurable

WITH A DEFINED
SUCCESS CRITERION

RAG

ANSWERS WITH A CITED
SOURCE

Human

IN THE LOOP WHERE IT
MATTERS

Your data

NEVER TRAINS THIRD-
PARTY MODELS

01

What we solve

Where AI genuinely adds value, and where it only generates cost.

COMMON SITUATION	WHAT IT MEANS FOR THE BUSINESS
Knowledge is scattered	Manuals, policies, and procedures spread across folders and emails. No one can find the answer, and the same question gets asked again every week.
Manual capture from documents	Someone types data from invoices, contracts, or forms into a system. It is slow, expensive, and guaranteed to contain human error.
Repetitive classification	Every incoming request has to be read and routed by hand, with delays and inconsistent judgment depending on who reads it.
Overloaded support team	The same handful of questions consumes most of the support team's time.
AI was tried and didn't work	A test was run without a defined success criterion, it produced approximate answers, and it was abandoned without knowing whether the problem was the approach or the model.
Fear of exposing information	There is interest, but no one knows what happens to the data sent to an outside provider.

02

What the service includes

01 Assistants over proprietary documentation

An assistant that answers using the company's own information and cites where it came from, not general knowledge from the internet.

Document ingestion and chunking · Semantic search · Answers with cited source · Per-user permission control · Continuous corpus updates

02 Document extraction

Converting unstructured documents into data usable by the company's systems.

Invoices and receipts · Contracts and clauses · Forms and identification documents · Rule-based validation · Review queue for low-confidence results

03 Classification and routing

Automatically sorting what comes in: emails, requests, tickets, or products.

Automatic categorization · Urgency detection · Routing to the correct department · Sentiment analysis · Learning from human correction

04 Conversational assistants

Web and messaging support that resolves common requests and hands off to a person when appropriate.

Web and WhatsApp · Queries to internal systems · Data capture and scheduling · Handoff to a person with context · Conversation logging

05 Decision support

Predictive models over historical data, when there is enough volume for them to be reliable.

Demand forecasting · Anomaly detection · Customer segmentation · Churn risk · Automatic prioritization

06 AI governance

The rules of use: which data can leave the company, who reviews what, and how answer quality is measured.

Acceptable use policy · Data classification · Continuous quality evaluation · Logging and audit · Model-error response plan

03

How we work

In short stages, with something verifiable at the end of each one.

PHASE	WHAT HAPPENS	WHAT YOU KEEP
01 · Discovery	We understand the business, the current state and the real constraint.	Diagnosis and agreed scope.
02 · Design	We define the solution, its risks and its execution plan.	Architecture and milestone plan.
03 · Build	We implement in verifiable increments, not in a single delivery.	Functional progress per milestone.
04 · Go-live	Controlled deployment, testing and knowledge transfer.	Solution in production and team trained.
05 · Evolution	Operation, measurement and continuous improvement under a service agreement.	Periodic report and living roadmap.

04

What you receive

Solution

- Deployed and integrated system.
- End-user interface.
- Connection to data sources.
- Administration panel.

Governance

- Data and privacy policy.
- Log of queries and responses.
- Permission matrix.
- Usage cost estimation and control.

Evaluation

- Proprietary test case set.
- Accuracy and coverage metrics.
- Defined confidence threshold.
- Procedure for incorrect responses.

Training

- User training.
- Usage best-practices guide.
- Administrator training.
- Technical documentation.

05

Technologies we work with

We choose the tool for the problem, not by trend or by agreement with a vendor.

CLAUDE

OPENAI

LLAMA

MISTRAL

LANGCHAIN

LLAMAINDEX

PGVECTOR

QDRANT

PYTHON

FASTAPI

WHISPER

OCR

OLLAMA

N8N

MCP

06

Engagement models

MODEL	HOW IT WORKS	WHEN IT FITS
Fixed project	Defined scope and deliverables, agreed price and timeline.	Bounded needs with a clear objective.
Monthly retainer	Reserved team capacity month to month.	Continuous evolution and sustained support.
Dedicated team	Profiles assigned exclusively under your direction.	Long programs or high workload.
Hour bank	Consumable hours with agreed priority.	Variable demand or one-off work.

07

Our commitments

COMMITMENT	HOW IT IS MET
Success criterion before we start	We define what counts as an acceptable result and how it is measured, before writing a single line of code.
Your data does not train third-party models	We use configurations and providers that exclude your information from training, documented in writing.
Answers with a source	Document assistants cite the document and section each answer comes from, so it can be verified.
Human where it matters	In decisions with legal, financial, or clinical consequences, the system proposes and a person approves.
Controlled cost per use	Consumption is estimated before deployment, with limits and alerts in place to avoid billing surprises.

WHEN WE SAY NO

When the process is not defined, when there is not enough data, or when the tolerable margin of error is zero and there is no reasonable way to review the output. In those cases, automating with conventional rules is usually cheaper, faster, and more reliable than a model.

08

Next step

GET STARTED

A one-hour session, at no cost, to review your current situation and define together the minimum scope to start with.

CONTACT

contacto@bohiosystem.dev

WEB

bohiosystem.dev

DOCUMENT

BS-SVC-09 · v1.0