GEMMA ANALYTICS

PROPOSAL

Document Extraction & IXOS Automation

Prepared for

hellomed Group GmbH

2026-04-22

Valid for 30 days

EXECUTIVE SUMMARY

hellomed is a B2B online pharmacy serving care facilities and nursing services across Germany. Two operational bottlenecks are driving this engagement:

1. **Patient onboarding in IXOS** – every new patient must be entered manually, field by field, into Pharmatechnik's IXOS desktop software. No API exists. This consumes significant pharmacist time.
2. **Document data extraction** – incoming documents (Einwilligung, SEPA mandates, prescriptions, medication plans) arrive as scanned PDFs via fax or email. An existing N8N workflow already classifies them with 98% accuracy, but does not yet extract structured data.

This proposal presents **four options** addressing either the first problem alone or both problems together, with different underlying architectures. The options are designed so that hellomed can choose the right balance between speed-to-value, scope, and architectural investment.

All options are created in light of the fact that hellomed holds an approved **Forschungszulage** to build a closed-loop, multi-agent AI system covering the pharmacy's full operational intelligence. Eight specialized agents (A0–A7) are defined in the approved scope, including the OCR and document extraction agent (A1). The work described here is the first concrete step toward that system and the foundation Agent A1 will be built on.

IXOS AUTOMATION STAGES

IXOS patient data input automation can be staged with varying degrees of automation. All options in this proposal cover Stage 1 (MVP). Later stages are outlined here for context.

Stage	Manual Effort	Extraction Method
0 – Current	100% manual input + review	N/A
1 – MVP	20% manual input + review	OCR
2 – Review Only	No manual input, review needed	OCR
3 – Automated	No manual input, reviewed by agent	hellomed OS API

OPTION A – IXOS PATIENT DATA AUTOMATION

IXOS RPA (Stage 1 MVP)

A Python script, executed via PowerShell, automates patient data entry into the IXOS desktop application on Windows. Using OCR to extract patient information from scanned Einwilligung forms and Windows UI accessibility identifiers, the automation interacts with IXOS exactly as a human would – without the manual effort.

The primary goal is to automate 80% of the UI actions (button presses, screen transitions) and data entry (typed characters into form fields). The remaining 20% requires human intervention for data inference, adjustments, or handling missing information.

Scope: IXOS RPA only (no document classification or extraction pipeline)
Best for: Fastest path to reducing manual IXOS data entry

DELIVERABLES

1. **Proof of Concept (POC)** – Python script executed via PowerShell to demonstrate functionality and collect initial feedback. Script-only, no end-user interface.
2. **Minimum Viable Product (MVP)** – Basic user interface allowing the end-user to drag and drop a PDF of a scanned form to trigger the automation process.
3. **Testing and validation** – End-to-end testing with hellomed team on sandbox environment.

ESTIMATED EFFORT

Phase	Est. Duration
POC (script only, no UI)	3–5 days
MVP (includes end-user UI)	2–3 days
Testing and validation	1–2 days
Total	**6–10 days**

PREREQUISITES

- Access to hellomed's environment, specifically an IXOS license
- A sandbox environment is highly desirable for dedicated testing and validation

OPTION B – AGENT A1 ON N8N + IXOS AUTOMATION

This option addresses both problems: document extraction and IXOS automation. Agent A1 is built as an extension of the existing N8N workflow.

MODULE 1 – AGENT A1 INFRASTRUCTURE (N8N)

Document Classification + Extraction on N8N

hellomed's existing N8N workflow already classifies incoming documents with 98% accuracy. Module 1 extends this workflow with two new extraction modules:

- **Einwilligung OCR Extraction** – the patient registration form is automatically read using OCR and AI (Mistral), with automated data completeness and data type validation.
- **SEPA OCR Extraction** – SEPA mandate document data is automatically read using OCR and AI (Mistral), with automated data completeness and data type validation.

From this point forward, a document does not just get classified: it gets read, validated, reviewed, and routed automatically, end to end.

Infrastructure: Existing self-hosted N8N instance
Architecture: Classifier + both extractors as a single N8N workflow with branching logic per document type
Human-in-the-loop: Consolidated review mechanism (agent + human) ensuring data quality before downstream use

What this means in practice:

- Familiar infrastructure, running on the existing self-hosted N8N instance
- The classifier and both extractors are connected as a single N8N workflow with branching logic per document type
- Structured output per document type routed to human-in-the-loop process and then to the correct downstream system (IXOS via Module 2 for Einwilligung; billing process for SEPA)
- Extraction performance data tracked for future improvements

Connection to the agent roadmap: Agent A1 lives entirely in N8N. Additional document types are added as new branches in the same workflow. The transition to a more structured agent architecture can happen incrementally in later phases.

MODULE 2 – IXOS PATIENT DATA AUTOMATION (RPA)

IXOS RPA (Stage 1 MVP)

Once the patient data extracted from Module 1 is confirmed, a Python script automatically enters it into the IXOS desktop application on Windows. Using Windows UI accessibility identifiers, the automation interacts with IXOS exactly as a human would, but without the manual effort.

The primary goal is to automate 80% of the UI actions and data entry. The remaining 20% requires human intervention for data inference, adjustments, or handling missing information.

Note: This module is built in parallel with Module 1.

ESTIMATED EFFORT

Module	Phase	Est. Duration
1	Design phase and UI/UX decisions	2–3 days
1	Agent prototyping	2–4 days
1	Feedback and improvement iterations	4–5 days
2	POC (script only, no UI)	3–5 days
2	Testing and validation	1–2 days

Module 1 -- Agent A1 (N8N)	8–12 days
Module 2 -- IXOS RPA	4–7 days
Total (modules built in parallel)	**12–19 days**

PREREQUISITES

- Access to hellomed's environment, specifically an IXOS license; sandbox environment highly desirable
- Example documents (filled out) for Einwilligung and SEPA in all formats they arrive in (PDF, JPG, etc.)
- Access details for Zoho ticket API, N8N, Mistral
- List of fields (required and optional) to be extracted per document, with field validation info if available (field type, character length, etc.)
- Point of contact for: system owners, validation owners, test users

OPTION C – AGENT A1 ON PYTHON + IXOS AUTOMATION

This option addresses both problems with a Python-native architecture instead of extending N8N. It provides the cleanest path from today's manual process to the full multi-agent system described in the Forschungszulage.

MODULE 1 – AGENT A1 INFRASTRUCTURE (PYTHON)

Document Classification + Extraction on Python

A Python service is built to handle the full Agent A1 logic: document classification, Einwilligung extraction, and SEPA extraction are all consolidated into Python-based agents. The existing N8N classifier logic is replicated and extended in Python, becoming the authoritative classification and extraction layer.

- **Einwilligung OCR Extraction** – automatically read using OCR and AI (Mistral), with automated data completeness and data type validation.
- **SEPA OCR Extraction** – automatically read using OCR and AI (Mistral), with automated data completeness and data type validation.

Infrastructure: New Python service (deployed independently)
Architecture: Python-based agents with full audit trail and logging
Human-in-the-loop: Consolidated review mechanism ensuring data quality before downstream use

What this means in practice:

- Einwilligung and SEPA extractors are implemented as separate Python modules within the same service
- Full audit trail and logging per document and per patient – queryable at any time
- Flexibility to evolve the architecture: by building in Python from day one, hellomed is not locked into a fixed communication pattern between agents. Both approaches are possible:
  - **Chain setup** (like N8N): each agent hands off to the next in a fixed sequence. Simple, predictable, well-suited for Wave 1.
  - **Orchestrator pattern**: a coordinator agent decides at runtime which sub-agent to trigger and when. More flexible, better suited for the full 8-agent vision.

Connection to the agent system roadmap: The Python service built here is the foundation for the agents system. Agents A2–A7 are built on top of it in subsequent phases. This option gives hellomed the cleanest path to the full multi-agent system described in the Forschungszulage.

MODULE 2 – IXOS PATIENT DATA AUTOMATION (RPA)

Same scope as Option B, Module 2. Built in parallel with Module 1.

ESTIMATED EFFORT

Module	Phase	Est. Duration

1	Design phase and UI/UX decisions	2–3 days
1	Agent prototyping	2–4 days
1	Feedback and improvement iterations	2–4 days
1	Deployment	1 day
2	POC (script only, no UI)	3–5 days
2	Testing and validation	1–2 days

Module 1 -- Agent A1 (Python)	7–12 days
Module 2 -- IXOS RPA	4–7 days
Total (modules built in parallel)	**11–19 days**

PREREQUISITES

Same as Option B.

OPTION D – AGENT A1 ON N8N + PYTHON + IXOS AUTOMATION

This option builds Agent A1 in both N8N and Python in parallel, allowing a direct side-by-side comparison before hellomed commits to one architecture for the broader agent rollout.

MODULE 1 – AGENT A1 INFRASTRUCTURE (N8N + PYTHON)

Document Classification + Extraction (Dual Architecture)

Both implementations handle the same document types (Einwilligung and SEPA) with the same extraction logic – making the comparison meaningful and fair.

- Everything from Option B – N8N unified agent workflow
- Everything from Option C – Python Agent A1 service, built in parallel
- hellomed makes the architecture decision for Wave 2 based on real MVP experience, not theory

Infrastructure: Both N8N and Python service
Architecture: Side-by-side comparison of both approaches
Outcome: Data-driven architecture decision for the full agent rollout

MODULE 2 – IXOS PATIENT DATA AUTOMATION (RPA)

Same scope as Option B, Module 2. Built in parallel with Module 1.

ESTIMATED EFFORT

Module	Phase	Est. Duration
1	Design phase and UI/UX decisions	2–3 days
1	Agent prototyping (both architectures)	3–6 days
1	Feedback and improvement iterations	4–5 days
1	Deployment	1 day
2	POC (script only, no UI)	3–5 days
2	Testing and validation	1–2 days

Module 1 -- Agent A1 (N8N + Python)	10–15 days
Module 2 -- IXOS RPA	4–7 days
Total (modules built in parallel)	**14–22 days**

PREREQUISITES

Same as Option B.

OPTION COMPARISON

Dimension	A	B	C	D
Scope	IXOS RPA only	Extraction + RPA	Extraction + RPA	Extraction + RPA
Agent A1 infra	–	N8N	Python	Both
Effort	6–10 days	12–19 days	11–19 days	14–22 days
Forschungszulage alignment	Low	Medium	High	High
Extensibility to A2–A7	N/A	Incremental within N8N	Native foundation	Informed choice
Risk	Low	Low	Medium	Low

INVESTMENT

Our daily rate is **[TBD]** (1 day = 8 hours).

Option	Est. Days	Est. Amount
A – IXOS RPA only	6–10	[TBD]
B – Agent A1 (N8N) + IXOS RPA	12–19	[TBD]
C – Agent A1 (Python) + IXOS RPA	11–19	[TBD]
D – Agent A1 (N8N + Python) + IXOS RPA	14–22	[TBD]

Day estimates reflect the total calendar duration. For Options B–D, Module 2 (IXOS RPA) is built in parallel with Module 1, so total elapsed time equals the longer of the two modules.

All prices net (excl. VAT). Fixed-price items are billed upon completion. Time & materials items are billed monthly based on actual days worked. Travel expenses, if any, are billed at cost.

PREREQUISITES

The following are required before the engagement starts. Items marked with * apply only to Options B–D.

- **IXOS access** – IXOS license in hellomed's environment. A sandbox environment is highly desirable for dedicated testing and validation.
- **Example documents*** – Filled-out Einwilligung and SEPA forms in all formats they arrive in (PDF, JPG, etc.)
- **Access details*** – Zoho ticket API, N8N instance, Mistral API
- **Field specifications*** – List of fields (required and optional) to be extracted per document, with field validation info if available (field type, character length, etc.)
- **Points of contact** – System owners, validation owners, test users

WAYS OF WORKING

Area	Approach
Communication	Shared Slack channel or agreed communication tool
Sync Cadence	Weekly sync call to review progress
Code Review	All work via PRs with review before merge
Documentation	All automation scripts and agent logic documented to handoff standard
Testing	End-to-end validation on sandbox environment with hellomed team

PHASE 2 AND BEYOND

This engagement delivers the foundation for hellomed's multi-agent AI system. Subsequent phases extend the architecture to cover additional document types and agents:

- **Additional extraction modules** – prescriptions, medication plans, and other document types added as new extraction modules (branches in N8N or Python modules, depending on architecture choice)
- **IXOS Stage 2 (Review Only)** – full automation of data entry with human review only
- **IXOS Stage 3 (Automated)** – agent-reviewed data entry via hellomed OS API, removing human review from the loop
- **Agents A2–A7** – remaining agents from the Forschungszulage scope, built on the infrastructure established in this phase

ABOUT GEMMA ANALYTICS

Gemma Analytics is a Berlin-based data consultancy specializing in modern data platform implementation. We work with growth-stage companies to build production-grade data infrastructure – from warehouse setup and pipeline engineering through to KPI dashboards and team enablement.

Our team brings deep experience across the modern data stack, with 70+ completed data platform projects. We operate code-first and leverage AI-assisted tooling to deliver at speed without compromising on quality or maintainability.

Over these projects we have built extensive internal documentation, reusable patterns, and engineering best practices that carry over to every new engagement – accelerating development significantly from day one.

GEMMA
ANALYTICS

Gemma Analytics GmbH – Chausseestraße 17, 10115 Berlin
gemmaanalytics.com