

*superluminar*

EXAMPLE DELIVERABLE

# AI Readiness Workshop

A straight answer on AI: where you stand today, which use cases are worth building, and the first build to start now, with the costs mapped out.

PREPARED FOR

## **Nordmann Logistik GmbH**

Mid-sized logistics · ~750 employees · Hamburg · already on AWS

# The short version

You are in a good position to build, though not everywhere you might hope. Your AWS foundation is strong and your teams are willing. The gaps are your data and your in-house AI skills. One use case is worth building now. One that looks attractive is not ready yet, and this report explains why.

<p>OVERALL READINESS</p> <p><b>Build-ready, with caveats</b></p> <p>Strong on cloud; weak on data &amp; skills</p>	<p>FIRST BUILD TO START NOW</p> <p><b>Support assistant</b></p> <p>Document-grounded Q&amp;A for the support team</p>	<p>INDICATIVE FIRST ENGAGEMENT</p> <p><b>~€21,000</b></p> <p>15 person-days · 5-6 weeks · ~€700/mo run</p>
--	---	--

## Where you stand

You already run on AWS in **eu-central-1 (Frankfurt)**, so the platform for a compliant AI build is in place. You have a large body of operational documents (shipping papers, customs forms, SOPs) and a support team answering the same questions repeatedly. What you do not yet have is that content in one accessible, clean place, or an ML-aware engineer and a named data owner in-house. Both are fixable, and the report shows how.

## What is worth building

A **document-grounded support assistant** is the clear first build: high value, high feasibility, and it runs on data you already own. A **route-planning helper** sounds appealing, but the data to support it is not there yet. Building it now would be guesswork, so we recommend against it for now.

## What to do first

Start with a small MVP of the support assistant: roughly **15 person-days** of delivery, an indicative **€21,000** one-off, and an estimated **~€700/month** to run at low volume. It uses Amazon Bedrock with a Knowledge Base over your documents in S3, stays inside the EU, and applies Bedrock Guardrails for PII and safety from day one.

### Data readiness is the gate

A use case with no usable data behind it will not work, however good it sounds. We check this first, so you do not spend the budget and find out later.

*All figures in this document are illustrative ballparks for an example client and are not a quote. They show the shape and order of magnitude of a real engagement.*

## OUTCOME 01


# A clear picture of where you stand

An honest read on your readiness, your regulatory exposure, and the capabilities you need in-house.

## DELIVERABLE 01

### Readiness scorecard

Five dimensions, scored 0–5. The pattern matters more than any single number: a strong platform sitting on top of scattered data and thin AI skills.



Readiness across five dimensions (0 = not in place, 5 = strong). Illustrative.

**Data (2.5).** You have plenty of it, which is the good news. The problem is access and consistency: documents live across shared drives, email and a document store, in mixed formats and languages, with no single source of truth. This is the main thing between you and a working assistant, and the focus of Phase 0.

**Cloud and infrastructure (4.0).** Your strongest dimension. You are already on AWS in Frankfurt, with the networking, accounts and identity in place. A Bedrock-based build slots into what you have rather than asking you to migrate.

**Governance and compliance (3.0).** GDPR practice is solid and you have a DPO. What is missing is AI-specific governance: how you classify use cases under the EU AI Act, who signs off on a model going live, and how you review outputs. The register that follows is the starting point.

**Skills (2.0).** Your engineers are capable on cloud and operations, but there is no ML-aware engineer, no formal data owner, and no evaluation practice. This is the gap most likely to slow you down. It is why we work alongside your team and transfer what we build rather than handing over something only we can run.

**Organisational appetite (3.5).** Leadership wants a clear answer and the support team is open to help with repetitive work. Appetite is real but pragmatic: people want to see it work on their own documents before they trust it.

## Risk & compliance register

Built for European data realities. GDPR posture, EU AI Act classification per candidate use case, and where your data lives. We treat these as first-class concerns.

### EU AI Act classification

Candidate use case	Risk class	Why, and what it obliges you to do
Support assistant (document Q&A)	Limited	Interacts with people, so transparency applies: tell users they are talking to an AI and that answers may be imperfect. Keep a human escalation path. Not high-risk: it informs staff and does not decide anything about a person.
Customs-form data extraction	Limited	Assists a human who checks the result before submission. Log extractions and keep human sign-off. Risk stays low precisely because a person remains in the loop.
Internal SOP assistant	Minimal	Internal knowledge lookup over your own procedures. Standard data-handling applies; no specific AI Act obligations beyond good practice.
Route-planning helper	Minimal	Operational optimisation with no personal data and no automated decisions about people. Minimal obligations. The constraint here is data, not law (see Deliverable 05).

None of your candidate use cases fall under the AI Act's high-risk Annex III categories (such as biometric identification, employment decisions, or critical-infrastructure safety). That keeps your compliance burden proportionate. Worth confirming with your DPO before go-live.

### GDPR posture & data residency

Concern	Position for the recommended build
Data residency	All processing and storage stays in <b>eu-central-1 (Frankfurt)</b> . Documents in S3, the Knowledge Base, the vector store and the Bedrock calls are all EU-region.
Model training on your data	Amazon Bedrock does not use your prompts or documents to train the foundation models, and content is not shared with model providers. This is contractual, not just configuration.
Personal data in documents	Shipping and customs papers contain names and addresses. Bedrock Guardrails redact PII in prompts and responses; access to source documents stays least-privilege via IAM.
Processor agreement	Covered by the AWS Data Processing Addendum. Record the assistant in your processing activities and run a short DPIA before rollout. It should be light, given the limited-risk class.
Sovereignty	If full EU operational sovereignty becomes a requirement, the same architecture maps onto the AWS European Sovereign Cloud without redesign. Not needed for the MVP.

## Skills & organisational readiness

The capabilities to have in-house to get the most from what gets built. Our aim is to leave your team able to run and extend this without us.

Capability / role	Status	What it covers, and the gap to close
Cloud platform (AWS)	Have	Accounts, networking, identity and operations are in place and well run. The build sits on top of this.
ML-aware engineer	Gap	Someone who understands retrieval, prompts, evaluation and the failure modes of LLMs. The most useful person to hire or upskill. We pair with this person during delivery so the knowledge stays with you.
Data owner	Partial	Operations leads own the documents in practice but not formally. Name one accountable owner for the assistant's corpus: quality, access and freshness.
AI governance lead	Gap	Your DPO covers GDPR; extend that remit to AI Act classification, sign-off and output review. Can be a part-time responsibility, not a new headcount.
Product owner	Partial	A support-team lead who decides what "good enough to ship" means and owns the eval set. Partly there; make it explicit.
Evaluation / MLOps practice	Gap	A repeatable way to measure answer quality and catch regressions. We set up the first version with you during the MVP so it is not an afterthought.

### Where the gaps are

The two real gaps are an ML-aware engineer and an evaluation practice. Neither blocks the first build: we cover both during delivery and hand them to your team. They do matter for whether you can run and grow this yourself afterwards.

## OUTCOME 02

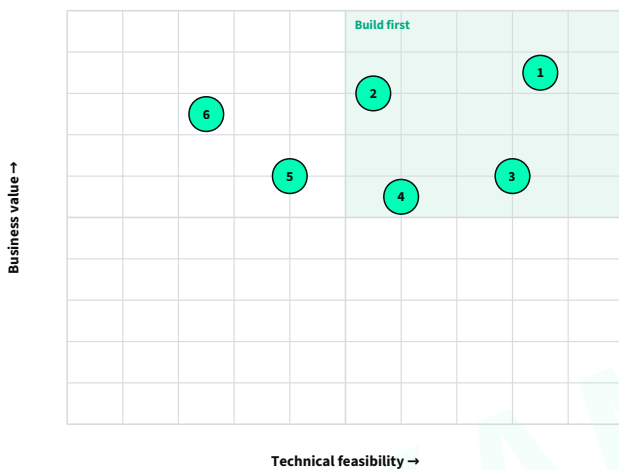
# The use cases worth building

Every candidate scored on value and feasibility, ranked so the strongest opportunities are obvious, and gated on whether the data actually exists.

## DELIVERABLE 04

### Prioritised use-case shortlist

Six candidates, scored on business value and technical feasibility. The top-right is where to build first. Rank reflects both axes and, decisively, whether usable data exists today.



#	Use case	Verdict
1	Support-ticket deflection via document Q&A	Build first
2	Customs-form data extraction	Strong second
3	Internal SOP assistant	Easy win; reuses build 1
4	Invoice / PO matching	Later
5	Shipment ETA prediction	Needs data work
6	Route-planning helper	Not ready: data gap

Value vs feasibility. Numbers match the ranked table. Illustrative.

Use cases 1 and 3 share most of their plumbing (a document-grounded assistant over your own content), so shipping the support assistant brings the SOP assistant almost for free. Use case 6 scores well on value but low on feasibility for one reason: the data is not there. That is covered next.

## Data-readiness gap analysis

For the top candidates: what data exists, what is missing or messy, and the one constraint that actually blocks the build. This is the gate.

### Use case 1: Support assistant *(recommended build)*

Data	State	Detail & what it needs
SOPs & policies	Partial	Exist as PDFs and Word files, but spread across drives and somewhat out of date. Needs collecting into one place and a light freshness pass.
Past support tickets	Usable	Available in the helpdesk with question and resolution. A strong grounding source and an eval set, once exported.
Shipping & customs docs	Partial	Plenty of them, but they contain PII and vary in format. Usable with redaction and basic structure on ingest.
Access & permissions	Blocker	No single service account or clean permission model to read across these stores. This is the real blocker: an access problem, not an AI one.

#### What blocks the build

It is not the model or the cloud. It is **document access**. Once a permission-cleared, PII-handled corpus exists in S3, the assistant is a routine build. That is what Phase 0 delivers, and why we do not start Phase 1 until it is done.

### Use case 6: Route-planning helper *(why we say "not yet")*

The value is real, but the data to support it is not. Historical route, timing and cost data is incomplete and inconsistent, telematics feeds are not retained in a usable form, and there is no clean record of what "a good route" looked like in the past to learn from. Building now would mean modelling on guesswork. The honest call is to fix the data foundation first, or accept this is a later project. It is cheaper to say no now than to spend the budget finding out.

## OUTCOME 03

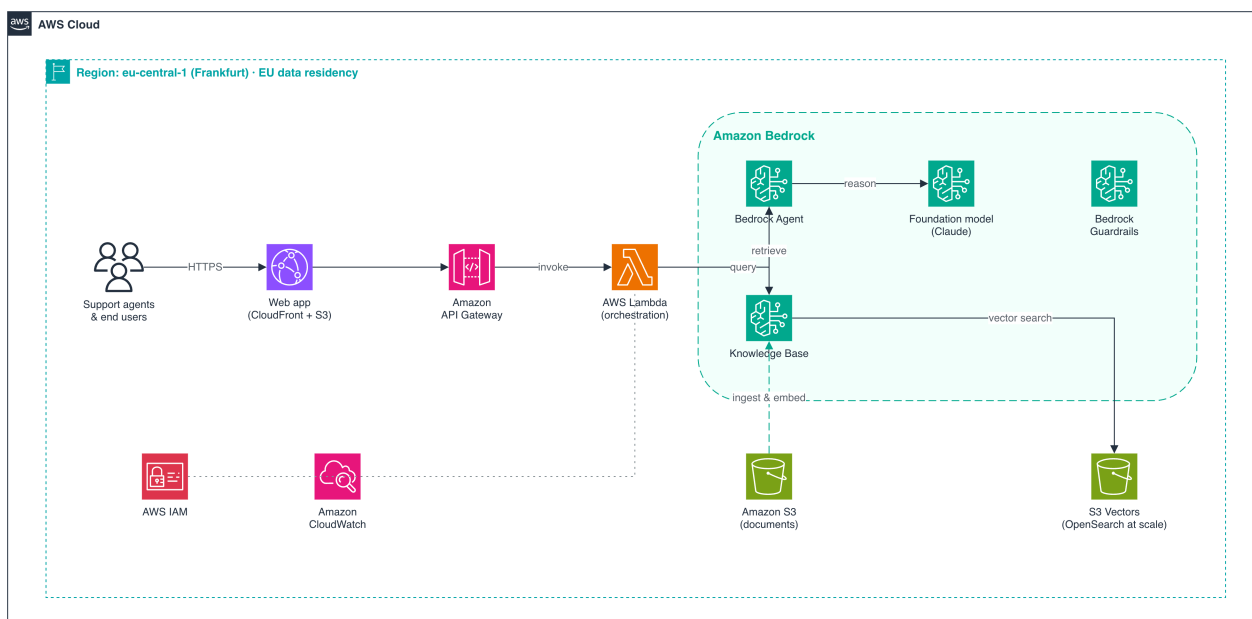
# Your first build, ready to start

A reference architecture on your own platform, the costs to deliver and run it, a phased plan gated on evidence, and the concrete first engagement.

## DELIVERABLE 06

### Reference architecture

The document-grounded support assistant on AWS, designed to sit inside your existing eu-central-1 account. Managed services throughout, so your team runs it without operating model infrastructure.



Retrieval-augmented support assistant on Amazon Bedrock. Editable draw.io source provided alongside this report.

### Request path

A support agent asks a question in a lightweight web app. API Gateway and Lambda pass it to a Bedrock Agent, which retrieves the relevant passages from the Knowledge Base, then has the foundation model (Claude on Bedrock) answer grounded in those passages, with a citation back to the source document.

### Knowledge path

Your documents sit in S3. The Bedrock Knowledge Base embeds them into a vector store (Amazon S3 Vectors by default at this volume, or OpenSearch Serverless at scale) and keeps them in sync. Guardrails handle PII and safety, CloudWatch logs every call, and IAM keeps access least-privilege. Everything stays in the EU.

## Indicative cost estimate

What it costs to deliver, and what it costs to run. Realistic ballparks in EUR, with the assumptions stated so you can sanity-check them against your own numbers.

### One-off delivery

Item	Basis	Indicative
MVP support assistant: build and knowledge transfer	15 person-days × €1,400/day	<b>€21,000</b>
<i>of which: Phase 0 data preparation</i>	<i>~5 person-days</i>	<i>€7,000</i>
<i>of which: Phase 1 MVP build &amp; eval</i>	<i>~10 person-days</i>	<i>€14,000</i>

Day rate shown at €1,400 (illustrative senior EU consulting; typical range €1,200–1,600). We embed with your team, so part of this spend is capability you keep.

### Estimated monthly run cost (low volume)

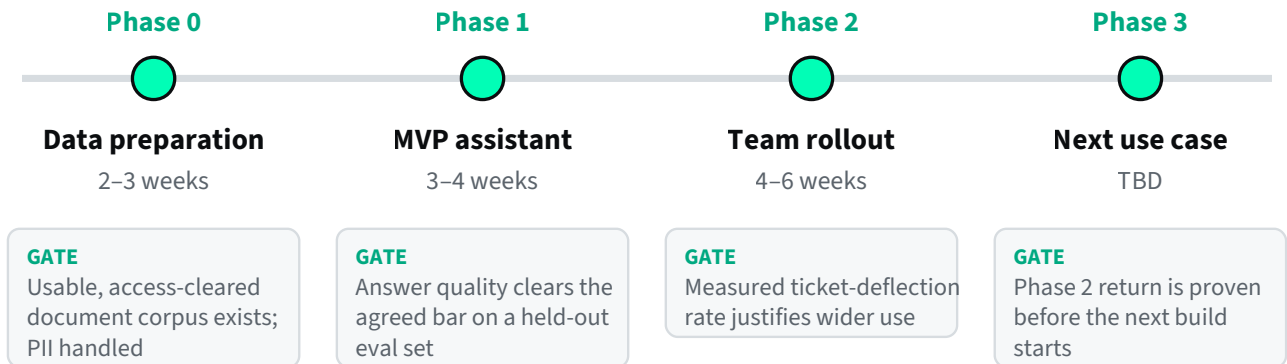
Line item	Assumption	Indicative / mo
Bedrock model tokens	~3,000 assistant answers/mo, RAG context	€300 – 800
Vector store: Amazon S3 Vectors	Low-volume serverless default	€40 – 120
Compute: Lambda + API Gateway	Low request volume	€30 – 60
Storage and sync: S3 + Knowledge Base	Tens of GB of documents	€20 – 40
Logging: CloudWatch	Standard retention	€20 – 40
<b>Total</b>	<b>Low-volume run</b>	<b>~€410 – 1,060</b>

#### On the vector store

For an MVP at this volume we default to **Amazon S3 Vectors**, which keeps the vector store at roughly €40–120 a month. OpenSearch Serverless suits higher query volume or tighter latency, but it starts around €350 a month, so it is not the right default here. We size this with you rather than reaching for the most expensive option.

## Phased roadmap

MVP first, then expansion. Each phase is gated on evidence. You only fund the next step once the previous one has earned it.



Each gate is a go/no-go. If a gate is not met, you stop and fix it before moving on. Illustrative timings.

Phase	What happens	The gate to clear before moving on
0 • Data prep	Collect documents into S3, clear access, redact PII, build the first eval set.	A usable, permission-cleared corpus exists. <i>No corpus means no Phase 1.</i>
1 • MVP	Build the assistant on Bedrock, ground it in the corpus, test against the eval set.	Answer quality clears the agreed bar on held-out questions.
2 • Rollout	Put it in front of the support team, measure real use and deflection.	Measured deflection justifies wider rollout and run cost.
3 • Next use case	Extend to customs-form extraction or the SOP assistant, reusing the platform.	Phase 2 return is proven before new spend starts.

## Recommended first engagement

The first build, laid out in full. This is the decision in front of you.

THE BUILD	EFFORT	TIMELINE	INDICATIVE PRICE
<p><b>Support assistant MVP</b></p> <p>Document-grounded Q&amp;A on Bedrock</p>	<p><b>15 days</b></p> <p>person-days, incl. data prep</p>	<p><b>5–6 wks</b></p> <p>elapsed, Phase 0 + Phase 1</p>	<p><b>~€21k</b></p> <p>+ ~€700/mo run</p>

### Scope

- One use case: a support assistant answering from your own documents.
- A starter corpus: ~50 core SOPs plus exported ticket history, German and English.
- Bedrock Knowledge Base + Agent, Guardrails for PII, all in eu-central-1.
- A first evaluation set and quality bar, set with your team.
- A lightweight internal web app for the support team to use it.

### Explicitly out of scope

- Customer-facing deployment (internal first).
- Route planning and ETA prediction: blocked on data.
- Any use case without a usable data source today.

### What you get

- A working assistant your support team uses on real questions.
- Evidence (measured answer quality and early deflection) to decide on rollout.
- The platform and patterns to add the next use case yourselves.
- Your team upskilled alongside ours. **You own what we build.**

#### The decision

Fund a 15-day, roughly €21,000 MVP that tests, on your own documents, whether a support assistant earns its place. It is low cost, stays in the EU, and is gated so you can stop early if the evidence is not there.

**superluminar**

*Not a vendor, a sparring partner.  
superluminar · AWS Advanced Consulting Partner · Hamburg*