

superluminar

EXAMPLE DELIVERABLE

AI Kickstart

You already know what you want to build. We take that one use case, check the data behind it, and hand you a build-ready plan: architecture, costs, and a clear first step.

PREPARED FOR

Brückner Hausgeräte GmbH

Home-appliance maker · ~850 employees · Stuttgart · already on AWS

The short version

You arrived wanting to automate the data entry from incoming supplier invoices and delivery notes, which your accounts-payable team keys into the ERP by hand. It is a good fit: the technology is well understood, your AWS platform is ready, and the documents already exist. The work is not the extraction itself. It is getting a labelled set to measure against, and a clean way to post results back into the ERP. Both are doable, and this plan shows how.

<p>VERDICT ON THE USE CASE</p> <h2>Worth building</h2> <p>Ready, once two gaps are closed</p>	<p>THE BUILD</p> <h2>Invoice extraction</h2> <p>Textract + Bedrock, human review for low confidence</p>	<p>INDICATIVE FIRST ENGAGEMENT</p> <h2>~€20,000</h2> <p>14 person-days · ~5 weeks · ~€800/mo run</p>
---	---	--

Why it fits

Pulling structured fields from invoices is one of the most proven uses of AI on AWS. You run on AWS in **eu-central-1 (Frankfurt)**, so a compliant build sits on what you already have. The AP team is keen to stop re-keying, and the documents arrive in volume every day.

What stands in the way

Two things, and neither is the model. First, there is no labelled set yet to measure accuracy against, though your ERP already holds the correct values to build one. Second, posting clean data back into the ERP needs a proper integration, not a manual paste. Phase 0 of the build handles both.

What to do first

Start with one document type, supplier invoices, on a real stream. High-confidence extractions post straight through; anything uncertain goes to a person to check. That is roughly **14 person-days**, an indicative **€20,000** one-off, and around **€800/month** to run at modest volume.

The honest part

If the labelled set shows accuracy is not good enough on your messier suppliers, you will hear that before you commit to a full rollout. The MVP exists to prove it on your own invoices.

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

One use case, checked and costed

A quick read on your readiness, an honest check of the data behind your use case, and a build-ready plan.

DELIVERABLE 01

AI readiness snapshot

A quick, high-level read on whether the ground is ready for **this** build, the invoice extraction. Not an organisation-wide assessment, just the three things that decide whether you can start: the platform, the people, and the compliance picture for this use case.

Ready

Platform

You run on AWS in Frankfurt, and every service this build needs (Textract, Bedrock, Step Functions) is available to you today. Nothing to stand up first.

Ready, with us alongside

People

Your ERP and accounts-payable teams know the process inside out. There is no ML-aware engineer in-house yet, so we build alongside them and leave the knowledge with you.

Low burden

Compliance

Invoices carry little personal data, and the use case assists people rather than deciding about them. Standard EU controls apply; nothing high-risk under the AI Act.

Verdict

Ready to build. The one real dependency is the data behind it, which the next page checks in detail. If you want a fuller picture of your readiness across the organisation, and the other use cases worth building, that is what our AI Readiness Workshop is for.

Use-case data check

Whether the data your use case needs actually exists, and what is missing. This is the gate. The extraction itself is not in doubt; these four rows are what decide whether it ships.

What the build needs	State	Detail, and what to do about it
Incoming invoices & delivery notes	Usable	Thousands arrive each month by email and scan, and they are retained. A strong, real stream to build and test on.
Correct field values (ground truth)	Partial	Your ERP already holds the values your team keyed in. Paired with the source documents, that becomes a labelled set to measure accuracy. This pairing is the main piece of Phase 0.
Format & supplier variety	Messy	Hundreds of supplier layouts, and scan quality varies. Manageable, but it is why a measured accuracy bar per supplier segment matters before you trust straight-through posting.
ERP write integration	Blocker	There is no clean, permission-controlled way to post extracted data back into the ERP today. Without it the result is another manual step. Designing this integration is the second piece of Phase 0.

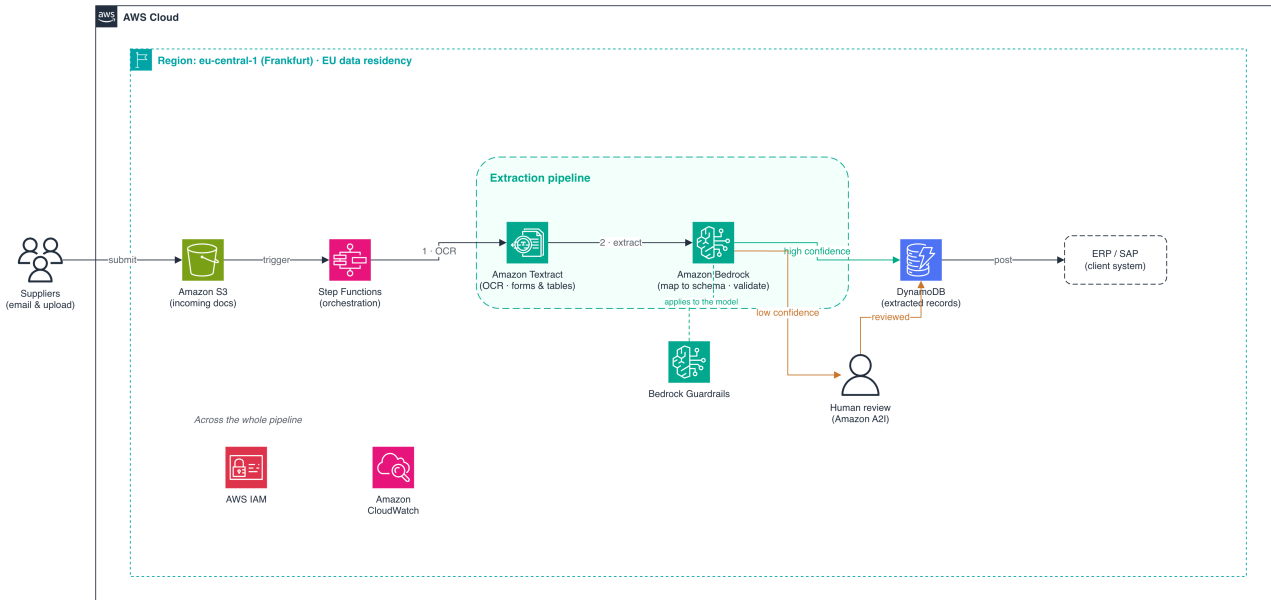
What actually decides this

Not whether AI can read an invoice. It can. The two things that decide success are a labelled set to prove accuracy on your own suppliers, and a clean path to write results into the ERP. Phase 0 delivers both before any extraction goes live.

If, on the labelled set, accuracy on your messier suppliers falls short, the right call is to keep those on human review and automate the clean ones first. You will have the numbers to make that call.

Reference architecture

The extraction pipeline on AWS, designed to sit inside your existing eu-central-1 account. Managed services throughout, with a person in the loop wherever the model is unsure.



Document extraction pipeline on AWS. Editable draw.io source provided alongside this report.

The straight-through path

Supplier documents land in S3. Step Functions runs each one through Amazon Textract for OCR, then Amazon Bedrock maps the result onto your invoice schema and validates it. When the model is confident, the clean record goes to DynamoDB and on to your ERP, with no one touching it.

The exception path

When confidence is low, the document goes to a person to check and correct before it posts. That keeps accuracy honest and gives you the data to widen straight-through over time. At the Bedrock step, Guardrails screen what the model reads and returns, for example catching prompt-injection text hidden inside a document. Across the pipeline, CloudWatch logs every step and IAM keeps access least-privilege. Everything stays in the EU.

DELIVERABLE 04

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 check them against your own numbers.

One-off delivery

Item	Basis	Indicative
Invoice-extraction MVP: build and knowledge transfer	14 person-days × €1,400/day	€19,600
<i>of which: Phase 0 (labelled set + ERP integration)</i>	<i>~5 person-days</i>	<i>€7,000</i>
<i>of which: Phase 1 (pipeline build & review loop)</i>	<i>~9 person-days</i>	<i>€12,600</i>

Day rate shown at €1,400 (illustrative senior EU consulting; typical range €1,200–1,600).

Estimated monthly run cost (modest volume)

Line item	Assumption	Indicative / mo
Amazon Textract (AnalyzeExpense)	~4,000–7,000 pages/mo	€350 – 700
Amazon Bedrock model tokens	Map to schema & validate	€120 – 350
Compute: Step Functions + Lambda	Per-document orchestration	€30 – 70
Storage: S3 + DynamoDB	Documents + extracted records	€20 – 50
Logging: CloudWatch	Standard retention	€20 – 40
Total	Modest-volume run	~€540 – 1,210

Where the cost sits

Textract is the cost driver here, and it scales with pages. If most of your invoices arrive as digital PDFs rather than scans, a lighter Textract mode applies and the run cost drops. We size this against your real document mix rather than assuming the worst case.

Recommended first engagement

The build we would take on with you, laid out in full. This is the decision in front of you.

THE BUILD	EFFORT	TIMELINE	INDICATIVE PRICE
Invoice extraction MVP One document type, real stream	14 days person-days, incl. Phase 0	~5 wks elapsed	~€20k + ~€800/mo run

Scope

- One document type: supplier invoices, on a live stream.
- Texttract for OCR, Bedrock to map to your invoice schema and validate.
- Straight-through posting for high-confidence records; human review for the rest.
- A labelled set and a measured accuracy bar, built with your team.
- A clean, permission-controlled write path into the ERP.

Out of scope for the MVP

- Delivery notes and order confirmations (a fast follow once invoices work).
- Full coverage of every supplier on day one.
- Removing human review entirely. It stays for low-confidence documents.

What you get

- A working pipeline running on real invoices, posting clean data to the ERP.
- A measured straight-through rate and accuracy, so you can decide on wider rollout.
- The labelled set and review loop to extend to other document types yourselves.
- Your team upskilled alongside ours. **You own what we build.**

The decision

Fund a 14-day, roughly €20,000 MVP that proves invoice extraction on your own documents and posts the results into your ERP. Modest run cost, EU-resident, and measured so you can widen it on evidence.