MARLUC

CASE STUDY · MARLUC SYNOPSIS

Federal Employment Agency generative AI assistants

IMPLEMENTED 2024–2026 · GOVERNMENT AND PUBLIC SERVICES · GERMANY · BAND D · CLAIMED ONLY

VERSION 1.0 · PUBLISHED 15 NOVEMBER 2026 · UPDATED 15 NOVEMBER 2026 · FACTS CHECKED 15 NOVEMBER 2026

Germany's public employment agency signed a four-year framework agreement with Aleph Alpha in 2024 and, by its own slides, unlocked generative AI assistants for 80,000 staff in February 2026. No measured results from the rollout are published.

ABOUT THIS SYNOPSIS

This is MarLuc's summary and grading of publicly available sources, for information only. It is not investment, legal, financial or professional advice, and should not be relied on for any decision. Readers must validate it against the original sources and, where appropriate, challenge its conclusions.

Facts and figures are as reported by the sources listed below, which may be incomplete, self-reported or out of date. Check the original source before relying on any figure.

The impact band, proof marker and SEAM view are MarLuc's opinion. The organisation is not a MarLuc client and has not reviewed or endorsed this synopsis. Names and trademarks belong to their owners.

NOTICE WORDING LAST UPDATED 16 NOVEMBER 2026

AT A GLANCE

AI used: Generative AI assistants for staff (text, summaries, knowledge search) on a platform with an Aleph Alpha framework agreement, with locally run models in an agent pilot.

Implemented: 2024–2026. Grade: band D, Task support, Claimed only.

Headline result, as reported: Generative AI assistants unlocked for 80,000 staff in February 2026, by the agency's own account; no measured results published

1 · The case

The facts, as reported by the 5 sources listed at the end. Where a source is silent, the field says "Not reported".

Organisation

Bundesagentur für Arbeit (Federal Employment Agency). Germany's public employment agency, described by its chair as Europe's largest social authority. Over 115,000 staff in 2024, by press reporting; about 150 local agencies and 300 jobcentres set up with local authorities. Canvas block: 4. Operations and process. Also 5. People and skills, 6. Data and technology.

Business challenge

The challenge. The agency says many staff will retire in the coming years. In October 2024 press reports put the figure at about 32,000 by 2032; its chief information officer told CIO.de that over 35 percent of colleagues will leave by 2032. He said recruiting that many people again is not realistic.

How it was addressed. In October 2024 the agency confirmed a four-year framework agreement with Aleph Alpha worth about €19 million, after testing models from several suppliers. It said the deciding factors were that it could run the technology itself and that Aleph Alpha showed best how its AI reaches results. Materna integrates the AI applications into the agency's specialist procedures. The agency has the option to run the infrastructure in its own data centre or a cloud. Its slides describe a text assistant (summaries, translations, drafting help, preparing advice conversations) and a knowledge assistant, built on a hybrid infrastructure of on-premises systems and a secured cloud from August 2025 and unlocked for 80,000 users from 23 February 2026. The slides do not name the model supplier for these assistants. A separate pilot, described by Capgemini and an agency author, uses only locally run models for agents that draft Jira tickets, with a person giving final approval. [1] [2] [3] [4] [5]

Who implemented it

Aleph Alpha (Model or cloud provider): Supplies the PhariaAI platform and models under a four-year framework agreement.
Materna (Systems integrator): Integrates the AI applications into the agency's specialist procedures and processes, per CIO.de.
Capgemini (Consultancy): Worked with agency architects on the Jira ticket agent pilot.
In-house (In-house): Bundesagentur für Arbeit IT: Developed an AI platform with its own developers and runs the infrastructure on premises and in a secured cloud.
The slides do not say which model supplier powers the assistants unlocked in 2026.

Technology (SEAM pillar E)

SEAM types: T6 Open-source / locally hosted models, T4 RAG / vector search over your own data, T7 AI agents. Deployment: Mixture. Vendors and products: Aleph Alpha, Materna, PhariaAI (Aleph Alpha), BAKIRA research assistant, GenAI text assistant and knowledge assistant, Locally run models in the agent pilot: Aleph Alpha, Llama, Mistral, Qwen.

Where it was built in: Staff-facing assistants for text, summaries, translation and knowledge search, integrated into the agency's specialist procedures by Materna; a separate pilot connects locally run agents to the agency's own Jira instance.

What changed

Before: Not reported. After: by the agency's slides, staff can use generative AI assistants for text work and knowledge search. The chief information officer said AI would neither endanger jobs nor decide cases on its own. [1] [2] [3] [4] [5]

What it cost

About €19 million over four years: the volume of the Aleph Alpha framework agreement, as confirmed by an agency spokeswoman, which is a contract ceiling and not reported spend. Other costs: Not reported.

What it saved

Not reported for the 2026 rollout. In 2024 the agency expected time savings of over 30% when producing notices after first pilots, per VDI's ingenieur.de; no measured result is reported since. [1] [2] [3] [4] [5]

Governance and control

The agency's slides describe a data ethics board that decides on AI use cases, a data ethics guideline with seven principles, a risk classification for each use case and an EU AI Act timeline, and say there are no prohibited AI practices in the agency. These are the agency's own statements. Accountable owner for the assistants and any monitoring of answer quality: Not reported.

S2 · Accountable owner · Claimed
The agency's slides show a data ethics board that decides on AI use cases and approves the guideline; its chief information officer is the named executive in the press. A single owner for the assistants is not reported. [2] [3]

A1 · Safe use · Claimed
The slides describe a data ethics guideline of seven principles and a risk assessment of each use case through its life cycle, and say there are no prohibited AI practices in the agency. [3]

A2 · Human in the loop · Claimed
The chief information officer said AI would not decide cases on its own; in the Jira pilot a person gives final approval of tickets. [1] [4]

A3 · Monitoring and assurance · Claimed
A reviewer agent checks pilot tickets for consistency and duplicates, and the slides describe an iterative risk assessment along the life cycle; ongoing monitoring of the 2026 assistants is not reported. [3] [4]

A4 · Vendor AI · Claimed
The agency says it tested models from several suppliers and chose Aleph Alpha because it could operate the technology itself and the supplier showed best how its AI reaches results. [1]

A5 · Agent controls · Claimed
The pilot's agents run in a shielded pilot environment on local infrastructure, with a person releasing tickets; controls for the wider assistants are not reported. [4]

A6 · Standards and regulation · Claimed
The slides set out an EU AI Act implementation timeline, risk classification of its AI use cases, and note that most are low or limited risk; no regulator finding is reported. [3]

What worked, what did not, what is unproven

WORKED

– Reach: the assistants were unlocked for 80,000 users by February 2026, by the agency's slides.

– In the Jira pilot, ticket creation worked smoothly in most cases, except for very large requests, by its authors' account.

– Control of the technology: the agency said it chose Aleph Alpha partly because it could run the technology itself.

DID NOT

– In the Jira pilot, a 26-page request exceeded the 8,192-token limit (8,732 tokens) until its first pages were trimmed.

– Running only local models means regular model updates take more effort, by the pilot authors' account.

UNPROVEN

– Whether the over 30% time saving expected in 2024 for producing notices was reached.

– Which model supplier powers the assistants unlocked in 2026.

– How much the assistants are used day to day, beyond the number of users unlocked.

Timeline

Oct 2024 The agency confirms a four-year framework agreement with Aleph Alpha worth about €19 million [1]

Dec 2024 CIO.de reports Materna will integrate the AI applications into the agency's specialist procedures and processes [2]

Apr 2025 Implementation of the BAKIRA research assistant starts [3]

Jul 2025 Testing of BAKIRA begins [3]

Aug 2025 Build of a hybrid infrastructure, on-premises plus secured cloud, begins [3]

Nov 2025 Rollout of the assistants begins [3]

Feb 2026 Assistants unlocked for 80,000 users [3]

2 · Grade and SEAM assessment

MarLuc's opinion, made on the evidence above. How cases are graded: marluc.co.uk/case-studies/how-we-rate

Level of change and impact

Claimed: Not applicable. Grade: D, Task support. Proof: Claimed only. Level of change: Task support.

Task support (Level 1, band D). The assistants support individual staff and decisions stay with people; no change to process, structure or cost is reported.

The reported uses are text, summary, translation and knowledge-search assistants for individual staff, with people still responsible for decisions. Time may be saved, but no change to process, structure or cost is reported. The agency says it is not using AI to decide cases.

The SEAM view and lessons learnt

The agency chose a supplier it could operate itself and set up its own data ethics board and risk classes before scaling. What is missing is evidence: the sources give user numbers and plans, and the 30% saving was an expectation in 2024. For a larger organisation, the transferable point is to publish a baseline and a measured result alongside rollout numbers.

1. Choose a model you can run yourself when the data is sensitive.
2. Assign each AI use case a risk class and a governing body before release.

Evidence

Score 1 of 4. Self-reported: the agency's own slides and a staff-co-authored article, with contract terms confirmed by an agency spokeswoman to trade press; no results published.

Evaluation method. Self-reported organisation figures and press reporting: the agency's own presentation and a staff-co-authored article, with the contract confirmed by an agency spokeswoman; no evaluation of results published.

Sources disagree. Reports differ on retirements: ingenieur.de reports about 32,000 of over 115,000 staff by 2032, while the chief information officer told CIO.de that over 35 percent will leave by 2032. They also differ on notices: ingenieur.de describes AI as partly automating the production of notices, while CIO.de says AI is not meant to create such notices.

Status. Ongoing: assistants unlocked for 80,000 users from 23 February 2026, by the agency's slides; the framework agreement runs for four years from 2024.

3 · Original sources

Read these in full. To challenge a fact or a grade, use the contact form at marluc.co.uk or the link on https://marluc.co.uk/case-studies/bundesagentur-fuer-arbeit-genai-assistants

[1] ingenieur.de (VDI): KI trainiert Behördensprache und zieht bei der Arbeitsagentur ein
https://www.ingenieur.de/technik/fachbereiche/kuenstliche-intelligenz/ki-trainiert-behoerdensprache-und-zieht-bei-der-arbeitsagentur-ein/

[2] CIO.de: GenAI: KI-Basiertes Wissensmanagement für die Bundesagentur für Arbeit
https://cio.de/article/3701792/ki-basiertes-wissensmanagement-fuer-die-bundesagentur-fuer-arbeit.html

[3] Bundesagentur für Arbeit: Künstliche Intelligenz in der BA (presentation slides, hosted by ver.di)
https://arbeitsmarkt-und-sozialpolitik.verdi.de/++file++69aaccd1bda665ad1484b047/download/__PPP%20AMP-Fachtagung%2026022026_Adrian%20Kraus.pdf

[4] Capgemini Germany: KI-Agenten für die Verwaltung: Wie die Bundesagentur für Arbeit große Sprachmodelle einsetzt (co-authored with an agency employee)
https://www.capgemini.com/de-de/news/kundenprojekte/ki-agenten-fuer-die-verwaltung/

[5] heise Schwarz Digits (sponsored): Öffentliche Verwaltung – modern und digital (interview with the agency's chair)
https://www.heise.de/solutions/schwarz-digits/bundesagentur_fuer_arbeit/

Bundesagentur für Arbeit (Federal Employment Agency) website: https://www.arbeitsagentur.de

Version 1.0 · Published 15 November 2026 · Updated 15 November 2026 · Facts checked against sources 15 November 2026 · Downloaded 26 November 2026. This PDF may be out of date: the latest version, and a way to report an error or a better source, is at https://marluc.co.uk/case-studies/bundesagentur-fuer-arbeit-genai-assistants

VERSION HISTORY

v1.0 15 November 2026 First published