GEMMA ANALYTICS

PROPOSAL

AI Initiative 2.0

Prepared for

Argos Security GmbH

2026-06-19

Table of Contents

EXECUTIVE SUMMARY

This document is structured as a matrix with two dimensions:

1. **Platform Options** (Section 1) – Four distinct tools that can be adopted independently or in combination, depending on which teams and use cases you want to address.
2. **Engagement Levels** (Section 2) – A modular engagement arc from platform setup through enablement, workshops, building, and roadmapping. Each level can be combined with any platform option – you choose the depth that fits.

Section 3 maps which engagement levels apply to which platform options, so you can see the full option space at a glance and assemble the right package for your situation.

PLATFORM OPTIONS

Four distinct tools that address different audiences and use cases. These are not mutually exclusive – they can be adopted independently or combined based on which teams and workflows Argos wants to address first.

OPTION A: LANGDOCK ROLLOUT

Personal tool – knowledge workers & back-office

Langdock -- Enterprise AI Workspace

A centrally managed, GDPR-compliant ChatGPT-like interface with projects, collaboration, and small agents. Model-agnostic – switch freely between Claude, GPT, Gemini and others without vendor lock-in.

Compliance: ISO 27001:2022, SOC 2 Type II, EU data hosting – built for GDPR from the ground up
Governance: Centralized admin: who uses what, usage statistics, permissions, audit logs
Integrations: Connects to company tools (HubSpot, Box, Confluence, etc.) to ground AI in company-specific data
Agents: Low-code agent building accessible to non-technical staff

Best for: Company-wide rollout, replacing shadow AI, centralized control, multi-model flexibility.

OPTION B: GDPR-COMPLIANT CLAUDE COWORK ROLLOUT

Personal tool – knowledge workers & back-office

Claude CoWork -- AI Desktop Companion

A desktop app that lets users delegate tasks to Claude directly on their local machine – files, browser, computer use. Offers guided walkthroughs and batch actions in an accessible UI. Requires pairing with AWS Bedrock or Google Vertex for GDPR-compliant, EU-hosted inference.

Model: Single-model (Claude only) – Bedrock/Vertex infrastructure supports other models but CoWork itself uses Claude
Compliance: Via Bedrock/Vertex setup for EU-hosted inference
Governance: Enterprise settings and admin console via Bedrock/Vertex – not a full admin platform like Langdock
Integrations: Via MCP (Model Context Protocol)

Best for: Individual users doing document-heavy, research, or operations work.

OPTION C: GDPR-COMPLIANT CLAUDE CODE ROLLOUT

Personal tool – specifically for engineers

Claude Code -- Agentic Coding Tool

A terminal/IDE-native agentic coding tool that lets engineers delegate entire development tasks to Claude – writing code, refactoring, debugging, managing git workflows, and interacting with internal tooling. Requires pairing with AWS Bedrock or Google Vertex for GDPR-compliant, EU-hosted inference.

Model: Single-model (Claude only) – same Bedrock/Vertex infrastructure as CoWork
Compliance: Via Bedrock/Vertex setup for EU-hosted inference
Governance: Enterprise settings and admin console via Bedrock/Vertex
Integrations: Via MCP – operates directly in the engineer's environment, can read, write, and modify files and repos autonomously

Best for: Engineers and technical SOC staff who want to delegate coding, automation scripting, and tooling tasks.

OPTION D: N8N SETUP

Company-wide tool – automated AI agents & workflows

n8n -- Workflow Automation & AI Agents

An open-source, fair-code workflow automation platform that combines visual/low-code building with custom code (JavaScript/Python) to create autonomous AI agents and automated pipelines. Unlike the other three options, this is backend infrastructure for company-wide automation – not a personal productivity tool.

Model: Model-agnostic – connects via API to any provider. GDPR-compliant inference still requires routing through Bedrock/Vertex (EU regions) or equivalent
Compliance: Can be self-hosted on Argos's own infrastructure – data never leaves the environment
Integrations: 400+ native integrations + custom nodes via CLI/SDK for internal APIs and proprietary systems

Best for: Automating repetitive, multi-step company processes – alert triage, incident enrichment, reporting pipelines.

REQUIREMENTS VS. PLATFORM FIT

Requirement	Langdock	CoWork	Code	n8n
Model-agnostic	✓ Multi-model	Claude only *	Claude only *	✓ Any via API **
EU data hosting	✓ Built-in	✓ Bedrock/ Vertex	✓ Bedrock/ Vertex	✓ Self-hosted
Replace shadow AI	✓ Primary fit	✓ Via rollout	Engineers only	Backend only
Secure doc upload	✓	✓	For engineers	–
Low-code agents	✓	–	–	✓ Visual builder
Centralized admin	✓ Full admin	✓ Bedrock/ Vertex	✓ Bedrock/ Vertex	✓ Self-managed
Tool integrations	✓ Native + API	✓ Via MCP	✓ Via MCP	✓ Native
SOC automation	Limited	–	Advanced	✓ Primary fit
Dev productivity	–	–	✓ Primary fit	–

* Bedrock/Vertex (part of the setup) offer other models, but CoWork and Claude Code specifically only use Claude.
** n8n connects to any model provider via API, but GDPR-compliant LLM inference still requires routing through Bedrock/ Vertex (EU regions) or equivalent.

ENGAGEMENT LEVELS

A modular engagement arc. Each level builds on the previous one, but Argos chooses where to start and how deep to go. These levels apply across all platform options.

LEVEL 0: PLATFORM FOUNDATION

Platform Foundation -- Technical Setup & Governance

Set up the chosen platform(s): accounts, configuration, security, EU data hosting. Define the platform governance framework – roles, permissions, usage policies, and audit structure. This becomes the foundation for all subsequent governance.

Output: Platform(s) live and accessible; governance framework documented and ready for team rollout.

LEVEL 1: ENABLEMENT – AI LITERACY & INSPIRATION

Enablement -- Leadership Session & Broader Team Rollout

Equip leaders and teams with the understanding and confidence to evaluate, sponsor, and govern AI use cases. Not tool training – inspiration-led with live demos on the real platform.

Leadership session (team leads, department heads):
- Define what "AI-augmented teams" means in practice for a SOC team lead vs. a consulting manager vs. back-office
- How to assess use-case proposals (what to say yes to, what to defer)
- Present the platform governance framework from Level 0

Broader team enablement (all employees):
- Support to prepare and structure enablement sessions across shifts and departments
- Inspiration-led with live demos – build a shared, accurate mental model of what current AI can and can't do

Output: Calibrated expectations across the org, shadow AI displaced, leaders ready to drive the workshop.

LEVEL 2: AI OPPORTUNITY WORKSHOP

Workshop -- Cross-Functional Use-Case Discovery

Surface, cluster, and prioritize use cases across selected departments (SOC, IR, Consulting, OT, Back-office). Use cases are scored on an Impact × Effort grid and mapped to the most appropriate platform option(s).

High impact / low effort → Quick Wins
High impact / high effort → Strategic Projects
Low value / not viable → parked and documented

Feasibility reflects the actual state of platform integrations from Level 0.

Output: One prioritized backlog that drives the rest of the engagement, mapped to platform options.

LEVEL 3: BUILD & CHAMPION USE CASES

Build & Champion -- Live Use Cases in Production

Real use cases, live in production, delivering value. Not proofs of concept – production-bound from day one.

3a – Build Clinic: People bring a real weekly task (ideally, a quick win from the workshop in Level 2) and leave with a working assistant.
3b – Champion Use Cases: An internal taskforce supported by Gemma takes top-backlog use cases to working assistants deployed in production. Support level scoped as a separate engagement at the beginning of the project.

Selection criteria: High impact × low effort, buildable now, representative (one per team), safe to deploy

Governance: During this phase, the per-agent governance framework is built against real builds – extending platform governance from Level 0 into a reusable per-project checklist covering data & access, actions & fail-safes, ownership, testing, auditability, and tiering. Each completed checklist becomes an audit artifact.

Output: Assistants live in production, internal champions, first KPI data, reusable governance framework battle-tested.

LEVEL 4: ROADMAP & MEASUREMENT

Roadmap & Measurement -- Handover & Mid-Term Plan

Handover and mid-term planning. The engagement ends with Argos self-sufficient.

Roadmap: High-impact/high-effort backlog items sequenced into a mid-term plan (up to automated SOC agents). Grounded in what the champion use cases proved buildable.

Governance framework: Refined and handed over – a reusable template Argos applies to every new agent or automation, independently.

KPI / reporting framework: Time saved per process, automated tasks, error reduction, user acceptance – with defined measurement intervals. Baseline set from champion use cases.

Output: Mid-term roadmap, governance framework handed over, KPI/reporting framework live.

PLATFORM × ENGAGEMENT MATRIX

Each platform option can be combined with any subset of engagement levels. This matrix shows which levels are relevant for each option – pick the platforms that fit, then choose how deep to go.

Engagement Level	Langdock	CoWork	Code	n8n
0 – Platform Foundation	Tenant, models, governance	Bedrock/ Vertex + config	Bedrock/ Vertex + config	Self-hosted deploy
1 – Enablement	All-hands demos	User onboarding	Engineer onboarding	Light, tool intro
2 – Workshop	✓	✓	✓	✓
3 – Build & Champion	Assistants & agents	Workflow templates	Dev tooling & scripts	Automated pipelines
4 – Roadmap	✓	✓	✓	✓

ABOUT GEMMA ANALYTICS

Gemma Analytics is a Berlin-based data and AI consultancy. We work with companies that value speed, transparency, and engineering quality over corporate overhead – building modern data infrastructure, AI-powered workflows, and team enablement.

Our team brings deep experience across data platforms and AI tooling, with 70+ completed 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 delivery significantly from day one.

CONTACT

Bijan Soltani

Founder & CEO

bijan.soltani@gemmaanalytics.com

+49 179 7438926

GEMMA
ANALYTICS

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