

ISACA · SEPTEMBER 2026

The road to Sovereign and Safe AI

Insights on governance, compliance and security



Harmenjan Sijtsma
CEO, EU GPT


Vadim Chepegin
CTO, EU GPT

From ASN GPT to EU GPT



ASN BANK · 2023–2025

Internal AI platform



EU GPT B.V. · TODAY

Sovereign European AI platform

OVER 20 YEARS OF AI DELIVERY



Our experienced team, for Europe

Fourteen people, 75 years of combined hands-on AI experience.

MANAGEMENT AND COMMERCIAL



Harmenjan

CEO, AI lead



David

CSO, interim director



Vadim

CTO, all-round tech



Benito

Sales, entrepreneur



Sandra

Marketing strategist



Simran

Marketing, media

ENGINEERING AND DESIGN



Steven

AI guru, >10 years



Stef

Engineering lead South



Peter

Front-end engineering



Niels

Cloud engineer, architect



Roshna

Data engineer



Fako

Engineering lead North



Lev

AI student



Aya

UX, intern HvA

ADVISORS



Bart

Experienced entrepreneur, Forward8



Michel

Experienced executive, Forward8

Contents

01

Where do we stand?

The evolution of AI, from chess computers to ChatGPT

02

Lessons from practice: ASN Bank

From the first AI tools to what worked

03

Jurisdiction and sovereign AI

Shadow AI, the risk of US law and the European answer

04

EU GPT in practice and outlook

Features, examples, time savings and roadmap

05

Controls and security

The CTO perspective, presented by Vadim Chepegin

PART 1 OF 5

01

Where do we stand?

The evolution of AI, from chess computers to ChatGPT

01 Where we stand

02 ASN Bank lessons

03 Sovereign AI

04 Practice & outlook


05 Controls

The evolution of AI

?

WARM-UP QUESTION

**When was the first
AI model built?**



Up to 2000: ever more computing power

?

WARM-UP QUESTION

**What was the name
of the chess
computer?**




Since 2022: smart AI models

?

WARM-UP QUESTION

**How long did
ChatGPT take to
reach 100 million
users?**



PART 2 OF 5

02

Lessons from practice: ASN Bank

From the first AI tools to what worked

01 Where we stand


02 ASN Bank lessons

03 Sovereign AI


04 Practice & outlook

05 Controls

Examples of AI tooling at ASN Bank



Transaction
categorisation



Customer feedback
categorisation



AI assistant

Which is GenAI and which is traditional AI?

Generative AI vs traditional AI

What are the differences?

Traditional AI

Collect your own data and build models to predict or recognise patterns.

EXAMPLES

Statistics · Analytics · Machine learning

Generative AI

Data has already been collected and models already built to generate new content such as text, audio, images and video.

EXAMPLES

LLMs · ChatGPT · Copilot

Challenges and what worked at ASN Bank

CHALLENGE 1 OF 2

1

Lack of knowledge and trust

How do we build AI in this context?



What worked at ASN Bank?

Discuss in pairs: how would you approach this?

Challenges and what worked at ASN Bank

CHALLENGE 1 OF 2

1

Lack of knowledge and trust

How do we build knowledge of and trust in AI?




WHAT WORKED

Make it safe, then make it easy

- Central AI team for coordination and education
- Develop an AI policy
- Quick win with risk colleagues
- Start with a secure AI assistant for everyone
- Champions who drive adoption per team

Retrospective at ASN Bank



Three types of AI applications



Challenges and what worked at ASN Bank

CHALLENGE 2 OF 2

2

Too many AI tools

How do we keep the number of AI tools under control?



?

What worked at ASN Bank?

Challenges and what worked at ASN Bank

CHALLENGE 2 OF 2

2

Too many AI tools

How do we keep the number of AI tools under control?




WHAT WORKED

Ask: is this really needed?

- Business case, reuse?
- Risk assessment of every new AI application
- Central register, visible to all colleagues

Define the ambition per department, but start small



PART 3 OF 5

03

Jurisdiction and sovereign AI

Shadow AI, the risk of US law and the European answer

01 Where we stand

02 ASN Bank lessons

03 Sovereign AI

04 Practice & outlook

05 Controls

QUESTION

What share of AI users bring their own AI tools to work?

A 21 percent

B 53 percent

C 78 percent

ANSWER

What share of AI users bring their own AI tools to work?

A 21 percent

B 53 percent

C 78 percent

Correct

Source: Microsoft & LinkedIn, Work Trend Index 2024 (31,000 knowledge workers, 31 countries)

Critical organisations are stuck

THE PRESSURE

30–50%

**of working time goes
to administration**



THE BOTTLENECK

**The AI time
saving lies in
files with
sensitive
data**

And that is exactly where
US AI cannot be used



THE RESULT

**Manual work
or Shadow AI**

The US government can demand your data

US LAW · ALREADY IN FORCE

CLOUD Act

Providers must hand over data, even when it sits in Europe.

FISA 702

Agencies may collect data on non-US persons. No case-by-case review.

Export controls and sanctions

Access can be cut off overnight.

EU LAW · ON ITS WAY

1

Location

Data and infrastructure in the EU

General use

2

Enhanced safeguards

EU staff, EUCS, external audit

Legal and judiciary

3

High sovereignty

No third-country control

Government · Education · Healthcare · Financial

4

Full control

Grip on the entire software chain

Defence and critical infrastructure

Sovereignty shifts from a position to a procurement requirement.


Source: Cloud and AI Development Act (CADA), four sovereignty levels

European AI, such as EU GPT, solves this


LAYER	US PROVIDERS	EU GPT
Tools and agents	ONE BLOCK UNDER US LAW Connections to your systems	Our own layer Agents in your own environment EU
Application	ChatGPT · Copilot · Gemini	Our own layer EU GPT B.V., Amsterdam EU
Data	<i>Your prompts, documents and case files</i>	Customer stays the owner
Model	OpenAI · Anthropic · Google	Open weights (US, CN, EU) Models do not save any data EU
Compute and cloud	NVIDIA · AWS · Azure · Google Cloud	Cloud (Scaleway) French company, security at scale EU

EU GPT keeps your data safe from US law

Open-source models practically at the same level



Enough depth in open-source models



Six times not the easy way

For every component a US service was the obvious choice. Each time we chose European, self-built or open source.

6 / 6

components under our own
control, on European infrastructure

One subprocessor

Scaleway, Paris, 100% European
owned

Authentication

Auth0, Okta, Entra ID

✓ Open source, self-hosted per environment

Language models

OpenAI, Anthropic, Google

✓ Open weights on European cloud

Document processing

Google Document AI, AWS

✓ Open source, within our own environment

Speech to text

US speech APIs

✓ Open model, self-hosted in Europe

Spell checking

Microsoft Editor, Google


✓ Open source plus our own model

Error tracking

Sentry (US)

✓ European alternative, self-hosted

Mistral and SafeGPT use US cloud



Microsoft Inc. • Cloud Infrastructure

Products: Le Chat, La Plateforme, Mistral Code



Google LLC • Cloud Infrastructure

Products: Le Chat, La Plateforme (US API)



Stripe Inc. • Payment and Billing

Products: All Products under a paid subscription



Get Lago Corp. • Billing

Products: All Products under a paid subscription



Blackforest Labs Inc. • Image Generation

Products: Le Chat, Agents API



Brave Software, Inc. • Web Search

Products: Le Chat, Agents API

PART 4 OF 5

04

EU GPT in practice and outlook

Features, examples, time savings and roadmap

01 Where we stand

02 ASN Bank lessons

03 Sovereign AI

04 Practice & outlook

05 Controls

Immediate value with European AI, such as EU GPT



Web search

Web and your own documents



AI Agents

Take over recurring work




Citations

Mandatory source references, built with UvA



Output in PPT and PDF

Straight to presentation or report



Voice AI

Reports and minutes, also offline




Vibe-coding

Build software in plain language



API

Embed in your own portals and systems



Everything runs in Europe

Scaleway, France · open-weight models

Citations

ANSWER IN EU GPT

3. Maximum inkomen voor aanspraak

Je kunt geen zorgtoeslag krijgen als:

- Je eigen rendementsgrondslag meer bedraagt dan €146.011 of
- Bij een partner: het gezamenlijke rendementsgrondslag meer is dan €184.633 ^[7]

Deze bedragen worden jaarlijks aangepast via een ministeriële regeling ^[7].

THE SOURCE, DIRECTLY AVAILABLE IN EU GPT

zorgwet.txt · zorgwet.txt

zorgwet.txt-7 · focused

Artikel 3 1. In afwijking van artikel 7, derde lid, van de Algemene wet inkomensafhankelijke regelingen, bestaat geen aanspraak op een zorgtoeslag indien de rendementsgrondslag, bedoeld in artikel 5.3 van de Wet inkomstenbelasting 2001, van de verzekerde aan het begin van het berekeningsjaar meer bedraagt dan € 146.011 of, indien de verzekerde het gehele berekeningsjaar dezelfde partner heeft, de gezamenlijke rendementsgrondslag, bedoeld in artikel 5.3, van de Wet inkomstenbelasting 2001, van de verzekerde en zijn partner aan het begin van het berekeningsjaar meer bedraagt dan € 184.633.

Example: homework checker for NMI

a $\frac{25}{17} + \frac{3}{4} = \frac{28}{21}$

b $25 \cdot 25 = 625$

Opgave a: $\frac{25}{17} + \frac{3}{4} = \frac{28}{21}$ ✗ FOUT

Opgave b: $25 \cdot 25 = 625$ ✓ CORRECT

Opgave a: $\frac{25}{17} + \frac{3}{4} = \frac{28}{21}$ ✗ FOUT

Deze uitwerking is **niet correct**. Het lijkt erop dat de tellers en noemers apart zijn opgeteld ($25+3=28$ en $17+4=21$), maar zo werken breuken niet.

Juiste uitwerking:

• Gelijksnamig maken (noemer: $17 \times 4 = 68$)

• $\frac{25}{17} = \frac{25 \times 4}{17 \times 4} = \frac{100}{68}$

• $\frac{3}{4} = \frac{3 \times 17}{4 \times 17} = \frac{51}{68}$

• Som: $\frac{100}{68} + \frac{51}{68} = \frac{151}{68}$

Antwoord: $\frac{151}{68}$ (of ongeveer 2,22)

Opgave b: $25 \cdot 25 = 625$ ✓ CORRECT

Dit klopt! $25^2 = 625$

Example: fitness schedule at UvA (USC)

Weekly Overview

Upper Body – Day 1 (Monday):

Chest-press, Seated-row, Lateral-raise,
Biceps-curl, Triceps-pushdown.

Upper Body – Day 2 (Thursday):

Incline-press, Cable-row, Front-raise,
Hammer-curl, Rope-pushdown.

Lower Body – Day 1 (Tuesday):

Leg-press, Leg-curl, Leg-extension,
Calf-raise, Plank.

Lower Body – Day 2 (Friday):

Goblet-squat, Romanian-deadlift,
Hip-thrust, Walking-lunge, Side-plank.

Value from European AI



Healthcare

30–45

minutes per shift

Voice reporting, answers
from your own protocols



Client files



Local government

1.5

hours per interview (approx.)

Interview reports for social
support and youth care,
grant information



Social support and youth
care files



Legal

1–2

hours per contract review

Contract review against
your own model clauses



Litigation files



Education

4–6

hours per teacher per week

Marking with feedback,
creating teaching material



Student files

Indicative time savings. They can only be realised safely with European AI.

Product roadmap

Six steps in product depth. The first two are live, the next two follow within three months.

TODAY

First knowledge features

Chat
Web search
Agents
Citations

Live

New knowledge features

Voice AI notes
PPT and PDF

Just live

WHERE EU GPT IS HEADING

Knowledge retention

Internal memory

In 2 months

Compliance

Regtech

In 3 months

Workflows

AI automation

Later

Business models


Model training

Later


Compliance hub: compliance through use

From regulation to a score and a go / no-go per answer or action


Starting point




Compliance hub



Output



Optional



PART 5 OF 5

05

Controls and security

The CTO perspective, presented by Vadim Chepegin

01 Where we stand

02 ASN Bank lessons

03 Sovereign AI

04 Practice & outlook

05 Controls

What would you need to trust this AI?

A practical question for the people who assess technology, controls and evidence.

“Prepare a management briefing using our files and what we learned last month.”

One request. Several data boundaries. A consequential output.



Company sources



Retained context




Accountable use

DISCUSS

What would you inspect first?

From assistant to organisational intelligence

Our direction: useful answers that carry context, boundaries and responsibility.



The value: less reconstruction, more relevant answers and controlled follow-through.

Sovereignty is a chain of controls

European infrastructure is the foundation. Control must continue through the workflow.



Infrastructure

Where processing happens
EU location and operating dependencies



Identity & isolation

Whose data is in scope
Tenant, project and user boundaries



Policy enforcement

What may happen next
Purpose, permissions and approvals




Evidence

How we demonstrate it
Decisions, tests and accountable review

Infrastructure + boundaries + enforceable decisions + evidence


Our foundation today

The operating architecture described in our current platform baseline.



Current architecture: where controls apply

Identity, application checks and operating boundaries in the existing baseline.



Secrets remain outside application code; universal tenant enforcement is a hardening priority.

What is already implemented today

Six concrete elements of the platform's operating foundation.

1 EU-hosted environments

Current environments and services run on Scaleway.

2 Environment separation

TEST, acceptance and production are operationally distinct.

3 Project access control

Membership and project roles shape workspace access.

4 Resource ownership checks

Files and conversations are linked to user context.

5 Identity and API access

Keycloak authentication and revocable API keys.

6 Operational baseline

Secrets outside code; logging, monitoring and recovery controls.

A working foundation today, with defined priorities for further hardening.

A control claim needs inspectable evidence

Start with the baseline; make the verification question explicit.

CONTROL IN THE EXISTING BASELINE	EVIDENCE TO INSPECT
EU-hosted environments	Processing locations and dependency inventory
TEST / ACC / PROD separation	Configuration boundaries and access tests
Roles, ownership and API access	Positive and negative authorisation tests
Operational controls	Restore results, access reviews and log coverage

Implementation, test results and independent assurance are distinct claims.

The next engineering work is explicit

Our current limitations define concrete priorities for strengthening the platform.

1

Organisation boundary

Enforce the boundary consistently across services.

2

Authorisation semantics

Apply the same rules across all access paths.

3

Mandatory policy checks

Put a common decision layer before sensitive actions.

4


Decision evidence

Link policy, decision, action and outcome.

Some controls still depend on configuration and operational discipline.

Shared meaning makes AI inspectable

A research foundation for explicit concepts, relationships and modular systems.



Define meaning → preserve provenance → support review


[Aroyo, De Bra & Chepegin \(2004\) · Semantic Web-based Adaptive Hypermedia](#)

[Sangers et al., including Chepegin \(2013\) · Semantic Web service discovery using NLP techniques](#)

Research foundations for shared semantics and language matching; today's architecture is a new application.

GraphRAG: connect evidence, then answer

Proposed knowledge workflow: reviewed relationships plus permission-filtered source passages.



DISCUSS

Where must permissions apply to derived knowledge?

Useful memory has an owner and a scope

Retain context deliberately; make correction, expiry and deletion part of the design.



Company knowledge

Approved terms and reusable facts

Stewarded for authorised users



Project memory

Decisions, constraints and findings

Bound to project membership



Personal memory

Preferences and working context

User-visible and controllable

Source + scope + purpose + review state + lifetime

Promotion into shared knowledge is an explicit decision. Retained context does not require model training.

An employee changes teams. What should follow?

A small organisational change tests the whole knowledge and memory lifecycle.



**Same person.
New context.**

Access to Project A is revoked.

Personal preference

“Keep my briefings concise.”

Restricted finding

A fact extracted from Project A.

Previous output


A briefing already saved or exported.

DISCUSS

What stays, what is revoked, and what evidence proves it?

Target: one governed knowledge-to-action flow


Identity and context constrain retrieval, memory and execution throughout.



Evidence across the flow: source → policy version → decision → action → outcome

Target architecture: three layers of control


Context, execution and governance work together across the complete AI workflow.



The governance layer applies across context, retrieval, memory and execution.

The policy decision point makes intent explicit


A proposed control contract between AI orchestration and sensitive execution.



Record the request, rule version, reason, approval, action and outcome.

EU-SAGA: from requirements to enforceable rules

A proposed path from interpreted obligations to testable technical controls.



Legal interpretation remains accountable human work; enforcement must be tested.

EU-SAGA: the rule bundle and decision record

Make the regulatory-to-technical interface precise enough to review and test.



EU law + sector requirements + customer policy → reviewed, testable controls

Contact

Questions afterwards? Feel free to get in touch directly.



Harmenjan Sijtsma

CEO, EU GPT B.V.

📞 +31 6 10 556 905

✉️ harmsijtsma@eugpt.nl

🌐 eugpt.nl



EU GPT B.V. · Chamber of Commerce 97997250 · Startup Village Amsterdam