

FINALE

TNO innovation
for life

**Futureproof AI
Impact Acceleration
Challenge 2026**



Opening

Our program for today

13.30 – 14.00

Event opening

Karen van der Zanden your host,
introducing TNO Futureproof AI

ECOSYSTEM BUILDING

14.00 – 14.30

Pitches 1-2

First pitches on sustainability by
design & energy measurement

ENERGY & AI

14.45 – 15.45

Pitches 3-6

Sustainability by design & energy
measurement

ROI & SOVEREIGNITY

16.00 – 17.00

Insights & closing

The future of compute & jury
feedback on pitches

INSIGHTS & NEXT STEPS




WE ASK YOU TODAY

Connect, contribute and collaborate

Let's build the ecosystem together. Today.



Let's build the ecosystem together. Today.



Why the Futureproof AI challenge

Sovereignty

A moment of reckoning has arrived for the West.

Our culture has fallen into shallow consumerism while abandoning national purpose. Too few in Silicon Valley have asked what ought to be built — and why.

We did.

We built Palantir to ensure America's future, not to tinker at the margins. On the factory floor, in the operating room, across the battlefield — we build to dominate.

Join us.

ALEXANDER C. KARP
CEO & CO-FOUNDER
PALANTIR TECHNOLOGIES INC.



PALANTIR.COM
PALANTIR TECHNOLOGIES
INC.



A blue-tinted photograph of a server room. Rows of server racks are visible, filled with various electronic components. The floor is a light-colored, reflective surface. The overall atmosphere is high-tech and industrial.

**SpaceX is spending
\$2.8 billion to buy
Gas Turbines for its
datacenters**



United Nations: AI's environmental cost threaten water, land and climate


Protecting European Values




**Future
prosperity and
earning power**



Balancing the ROI of AI







**Together
we go
further...**

The Challenge program

From shared challenges, to shared, validated solutions

1



Real, shared challenges

Each working group anchors to a live issue with stakeholders behind it.

FOUR CHALLENGES

2



Trusted orchestration

Neutral, mission-aligned, and credible orchestration for every party

SHARED EFFORT FOR IMPACT

3




Complementary partners

Research depth, market access, and execution muscle working as one team.

RESEARCH → APPLICATION

4



Repeatable method

Time-boxed agile sprints, a common vocabulary to move in ~~years~~ months

OPEN INNOVATION

From shared challenges, to shared, validated solutions

1



Real, shared challenges

Each working group anchors to a live issue with stakeholders behind it.

THREE CHALLENGES

2



Trusted orchestration

Neutral, mission-aligned, and credible orchestration for every party

SHARED EFFORT FOR IMPACT

3




Complementary partners

Research depth, market access, and execution muscle working as one team.

RESEARCH → APPLICATION

4



Repeatable method

Time-boxed agile sprints, a common vocabulary to move in ~~years~~ months

OPEN INNOVATION



Energy



ROI



Sovereignty

From shared challenges, to shared, validated solutions

1



Real, shared challenges

Each working group anchors to a live issue with stakeholders behind it.

THREE CHALLENGES

2



Trusted orchestration

Neutral, mission-aligned, and credible orchestration for every party

SHARED EFFORT FOR IMPACT

3




Complementary partners

Research depth, market access, and execution muscle working as one team.

RESEARCH → APPLICATION

4



Repeatable method

Time-boxed agile sprints, a common vocabulary to move in ~~years~~ months

OPEN INNOVATION



Trusted
knowledge
institute



Experienced
open
innovator

From shared challenges, to shared, validated solutions

1



Real, shared challenges

Each working group anchors to a live issue with stakeholders behind it.

THREE CHALLENGES

2



Trusted orchestration

Neutral, mission-aligned, and credible orchestration for every party

SHARED EFFORT FOR IMPACT

3




Complementary partners

Research depth, market access, and execution muscle working as one team.

RESEARCH → APPLICATION

4



Repeatable method

Time-boxed agile sprints, a common vocabulary to move in ~~years~~ months

OPEN INNOVATION



Knowledge
partners



Public
companies



Private
companies

From shared challenges, to shared, validated solutions

1



Real, shared challenges

Each working group anchors to a live issue with stakeholders behind it.

THREE CHALLENGES

2



Trusted orchestration

Neutral, mission-aligned, and credible orchestration for every party

SHARED EFFORT FOR IMPACT

3




Complementary partners

Research depth, market access, and execution muscle working as one team.

RESEARCH → APPLICATION

4







Repeatable method

Time-boxed agile sprints, a common vocabulary to move in ~~years~~ months


OPEN INNOVATION

Challenge
shaping



Ecosystem
pitch

Building an Ecosystem



GPT NL



Our jury members



**René
Steenvoorden**

Voorzitter Raad van
Bestuur
UWV



**Tanja
Hulst**

Programme Manager
Digital Economy
Provincie Zuid-
Holland



**Bert
Voorbraak**

Algemeen
Directeur
Logius



**Fleur
de Groot**

Area Lead
Mortgages
ING



**Stefan
Groot**

Head of Data, AI &
Agentic Tech
ABN AMRO




**Muriël
Serrurier
Schepper**

Programma manager
AI
De Nederlandsche
Bank



**Hillen
Oost**

Toekomst-bestendige
& Digitale Economie
RVO



AI Rating Model

Wybren van der Zee



A

B

C

D

E

F

AI Rating Model

Wiebren van der Zee


Truly **green** AI




is not a contradiction, it's a **design discipline.**

Right now, AI is not truly green.

It's truly **dirty**.




Resources



Footprint



Trend



Impact

A

B

C

D

E

F



Which ~~fridge~~ AI would you choose?



Can we apply GenAI?



Yes, but how **green**?




Green? What you mean?



Well AI Impact differs hugely..



Oh like that, minimum **Rating C!**



The Green AI Deal.

With a **carrot** and a **stick**.



Subsidies & credits.

Favour and support greener AI built and operated in Europe.



Carbon tax & offsets.

Linking AI emissions into the EU regime EED, ETS, CBAM, CE.

Recently discussed with...



Sustainable IT Experts

“We need this.”



European Commission Officers

“Let’s talk further”



Rijksoverheid

Local Legislator

“Interesting!”


Why select green AI?



Support the transformation
to Net Zero



Credibility on your
Green Targets



Avoid Carbon Cost

AI Sustainability Rating

Time & Resources

*for positioning and adoption in your resp. organizations SDLC
(software development lifecycle)*

Sustainability of AI

added as new non-functional

ALONGSIDE THE CLASSICS (and other newcomers)

● Confidentiality


● Integrity

● Availability

● Privacy

● Security

● Sovereignty



**What you do makes a difference,
and you have to decide what kind
of difference you want to make.**

Jane Goodall

The Framework

Drivers

- Energy/Token
- Carbon/Token
- Carbon/Model
- Water/Token ^{*1}
- Carbon/Hardware
- Circularity/Hardware ^{*1}
- Carbon Credits/Rating
- Carbon Tax/Rating

^{*1} = Optional for first iteration/under construction

GenAI-Sustainability Rating – Linking into EED/ETS/CBAM/CE						
Assessing/scoring characteristics of GenAI						
• CO2eq Location based measured (Local Grid Intensity) • Next rating level when all criteria are met (Asset circularity and Water stress are optional at initial stage) • Norms to be reassessed over time, Values indicative for concept at this stage • Incentivized via Bonus/Malus regime						
Sustainability Rating	Inference phase		Training phase - Carbon/Model	Embodied - Carbon/Hardware	Circularity - Minerals	Water - Stress
	Energy/Token	Carbon/Token				
A	< 0,001 Joule	< 1 Gram	< 1 Ton	< 10 Kg	> 90% Circular	< 0,000001
B	< 0,01 Joule	< 10 Gram	< 10 Ton	< 50 Kg	> 75% Circular	< 0,00001
C	< 0,1 Joule	< 20 Gram	< 25 Ton	< 100 Kg	> 50% Circular	< 0,0001
D	< 1 Joule	< 30 Gram	< 50 Ton	< 250 Kg	> 25% Circular	< 0,001
E	< 2 Joule	< 40 Gram	< 100 Ton	< 500 Kg	> 10% Circular	< 0,01
F	> 2 Joule Or no Data from supplier	> 40 Gram Or no Data from supplier	> 100 Ton Or no Data from supplier	> 500 Kg Or no Data from supplier	< 10% Circular Or no Data from supplier	> 0,01 Or no Data from supplier
Remarks	Energy usage per Token (Joule/Token)	Emissions (GHG) per Token (CO2eq/Token)	Emissions (GHG) per Model Post + Pre Training (CO2eq/Model)	LCA/PCF: Comp./Store./Transp. GHG: Manufacturing and decom. (CO2eq/TA Assets (avg))	GCP/CTI based Link with CBMA + Circ. economy Act.	Water use as availability at infer. Location (ml/Token)
EED EU Circular Economy ETS CBAM CE						

Wielonen van der Zee - April 2026
Version 0.9/1

Why together?

*Alone you may go **faster**.*

*Together we go **further**,
and have more **impact**.*



Help us make **AI**
worth its **energy.**

Join the pilot →

Co-author the standard

Let's talk

Scan me



Wiebren van der Zee

ABN Amro –
Sustainable IT



Pepijn van der Spek

Rijksorganisatie voor
Ontwikkeling,
Digitalisering en
Innovatie



Pim Scheele

ABN AMRO
Data Scientist



Kevin Gong

Logius



Justus Bogner

VU Amsterdam

“An IT Transformation that does not lead to a reduced environmental impact, is not the right Transformation”. 2021 WvdZee

Truly **green** AI



is not a contradiction, it's a **design discipline.**



Energy Measurement

The world is on fire.





The world is on **fire**.

And one of the
fastest-growing
causes is **AI**.

AI has a big climate impact




We don't know
how big

AI has a big climate impact



We don't know
how big



We know
we're not in control

Presenting

TrAIce 

A decision layer for the true cost of AI.

TrAIce_

Last month

4 models

All providers

Retail Banking

All data centers

Upload data

Run simulation

Ask AI about this data...

Q

Usage & emissions over time

Energy per period - one line per selected model - 15% margin shaded

Tokens used

357.5k

Prompts

2.8k

Active users


385

Cost

€744

Avg. latency

723 ms



AI-powered suggestions

2 of 5 - swipe to review



Lift your energy label from D to B over the next 12 months

Proactive

Impact on productivity & time-to-market

Units of work delivered per period - vs. pre-AI baseline


Units

Cost / unit


+15%

vs. start of window





guessing is expensive,



guessing is expensive, insight is not!

**Productivity is rising,
but control is not keeping pace.**



A decision layer for senior management



What if a **smaller model** can do the job?



What if we **forecast impact** before scaling a use case?



What if we **compare suppliers** on real business workflows?

**A decision layer for senior management
for real-time decision making**

Last month 4 models All providers All departments All data centers

Upload data

Run simulation

Ask AI about this data...



Usage & emissions over time

Energy per period · one line per selected model · ±5% margin shaded

Energy (kWh)

Emissions

Cost (€)

±5%

How is this measured?

Tokens used

1.63M

⚡ Prompts

12.7k

👤 Active users


493

€ Cost

€3.4k

🕒 Avg. latency

726 ms



Impact on productivity & time-to-market ⓘ


Units of work delivered per period · vs. pre-AI baseline

Units

Cost / unit

+30%

vs. start of window



Unit: u

AI-powered suggestions ⓘ

Proactive

1 of 4 · swipe to review



Use Caveman prompts to cut input tokens

Most prompts carry 30–45% filler the model still pays for. A pre-processor that strips them down to blunt, caveman-style instructions cuts input tokens by $\sim\frac{1}{3}$ with no quality loss.

EST. SAVINGS / MONTH $\approx$ €1195

2850 kWh energy

× Skip

✓ Accept

Not one vendor.

A Dutch coalition.



Rijksdienst voor Ondernemend
Nederland



UNIVERSITY
OF TWENTE.



**Real enterprise
use cases**



**Scientific
rigour**



**Vendor-neutral
orchestration**

Join us.

- **Join our coalition**
- **Run pilots with real organizational data.**




Contact us

**Embed carbon impact of AI into governance and
decision-making.**

Let's control the impact



Let's control the impact



Guessing is expensive, insight is not!

Thank you



RVO
Eward Bartlema



UWV
Tuan Nguyen



UWV
Dima Altai



UWV
Peter de Rooij



Try out Tr**AI**ce yourself!



DNB
Ping Lin



ING
Effi Bennekers



ING
Karen Coolen



ING
Harold Weaver



UvA
Shashikant Ilager



UT-Twente
Fernando Castor



TNO
Aditya Ganesh

With support from

TNO: Dejan Zamurovic, Karen van der Zanden, Lieke Dom


nlmtd: Anna Wendt Spies, Nicholas Morgan

TIME
FOR
BREAK!



ROI of AI

Xander Mol





For people
and work



Less pressure

More time
for people



Shorter waiting times

People helped
sooner



Faster decisions

Efficient and
consistent



Happier doctors

More focus,
less admin



UWV AI Assistant

Online

Good morning!
How can I help with this assessment?



Summarize case
Extract key facts and history



Draft assessment
Create structured report



Check consistency
Validate information



Suggest next steps
Recommendations

Assessment draft

B I [List Icon] [List Icon] [List Icon]

Summary

[Text lines]

Functional limitations

- [Text line]
- [Text line]
- [Text line]

Conclusion

[Text lines]



Report structured
Ready to review

What we must watch



Bias
Unfair outcomes
possible



Privacy
Sensitive data
at risk



Errors
Hallucinations
and mistakes



Accountability
Unclear
responsibility



Automation bias
Over-reliance
risk



AI **empowers** our experts.

Responsible use **protects** the people we serve.



For people
and work



More pressure

Overwhelmed
teams



Longer delays

People left
waiting




Slower decisions

Inefficient and
inconsistent



Unhappy doctors

Burnout, errors
and turnover



What we must watch




Data breach
Sensitive data
compromised



System outage
Services
disrupted



Operational risk
Processes
paralyzed



Reputational damage
Loss of trust
and credibility



Compliance risk
Fines and legal
consequences



AI without safeguards puts our mission at risk.
Irresponsible use can harm the people we serve.



AI support for better human decisions



Better Decisions



Human control



AI speed

AI Decision Making Tool Box



Assessment Assistant



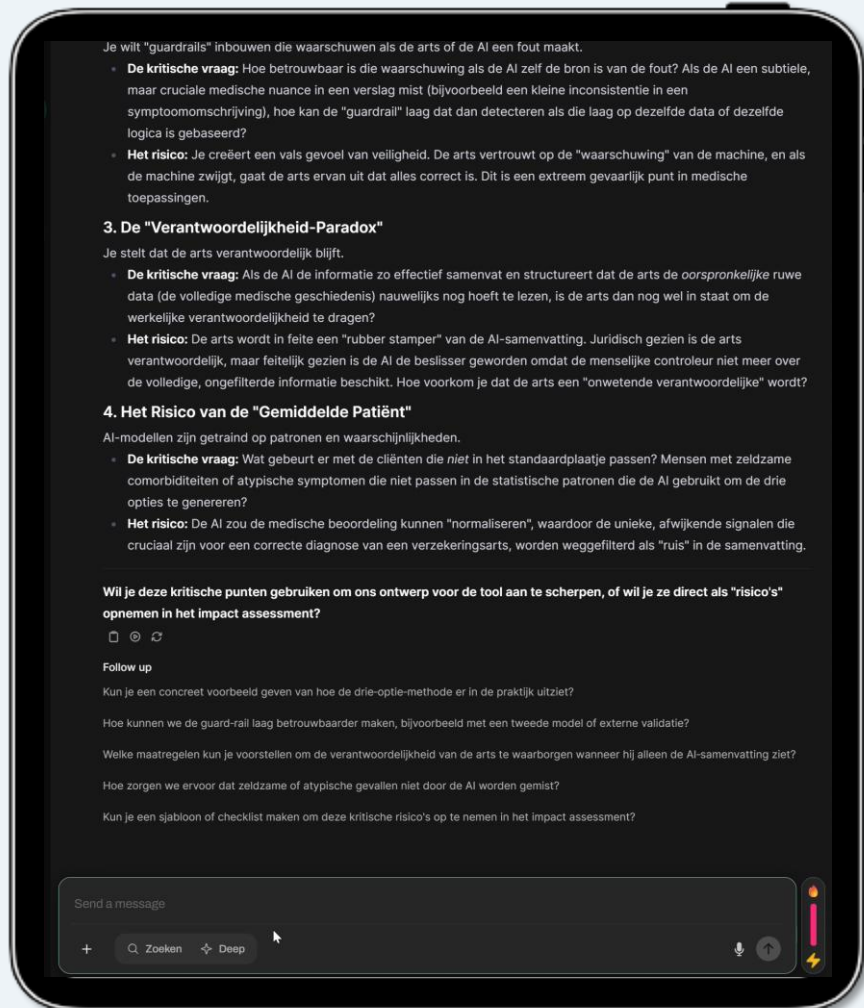
Agentic Stakeholder Dialogue



Decision Dashboard



Assessment Assistant



Interactive dialogue guidance



Company wide validation




Evidence-Based context



Historical cross check & online enrichment



Agentic Stakeholder Dialogue



**Extensive scenario
description and specific
requirements**

**Detailed actor traits and
knowledge**

**Moderator style and final
configuration**

**Results
Full transcript and summary
of dialogue**



Decision Dashboard



Recommendations



Risk Analysis



Decision Memory

Toprisico's voor besluitvorming

Gerangschikt op combinatie van kans, impact en mate waarin mitigatie al concreet is.

Accuracy & over-reliance

Onjuiste of te stellige samenvattingen kunnen doorwerken in medische beoordeling.

hoog

Reputatie & publiek vertrouwen

Medische AI binnen een publieke uitvoerder vraagt sterke uitlegbaarheid en transparantie.

hoog

Vendor lock-in

Afhankelijkheid van model, cloud, contract, monitoring tooling en roadmap leverancier.

middel

Damages & aansprakelijkheid

Schade door foutieve ondersteuning, onvoldoende logging of onduidelijke verantwoordelijkheid.

middel+

Privacy & dataminimalisatie

Medische gegevens vereisen strikte grondslag, logging, toegangsbeheer en bewaartermijnen.

hoog



We recommend:
use AI!

I USE AI
DO YOU?

WE USE
OUR OWN
PRODUCT!

And that makes sense.
We are, after all,
a robot.

The solution for
EVERYTHING!



SMARTER • FASTER • EASIER

💡 Ideas
💬 Texts
📊 Analysis
✍️ Creation
★ Inspiration

24/7
WORKING FOR
YOU!

Why it works

"This is exactly what we are missing."



Validated in Practice

- UWV
- Logius
- ABN AMRO
- TNO



Recognised Pain

- Fragmentation
- Endless discussion
- Lack of overview



Strong Enthusiasm

Different organisations,
same conclusion.

Why us



**Frank
Kaptein**



**Rachel
Kuijlenburg**



Logius
Ministerie van Binnenlandse Zaken en
Koninkrijksrelaties



**Tineke
Blom**



ABN·AMRO



**Hans
Stokking**


TNO innovation
for life



**Xander
Mol**




What we ask



Co-creation



Pilot partners



Launch

Today's Schedule

8:00 AM	On time	✓
8:30 AM	On time	✓
9:00 AM	On time	✓
9:30 AM	On time	✓
10:00 AM	On time	✓
10:30 AM	On time	✓

😊 All clear – No wait

Better care.
Less wait.
Happier together. ❤️

AI Triage

🤖 Symptoms assessed
No urgent risks

✓ Scheduled today

AI Summary

📄 ✓ History
✓ Medications
✓ Allergies

✓ Summary ready

Smart Scheduling

📅 Optimal schedule
No gaps
Reduced wait time

✓ All good

Auto Notes

📝 Visit notes
captured
automatically

✓ Saved



Follow-up Reminders

🔔 Reminders sent
to patient

✓ Done

Care Insights

📊 Population trends
Better outcomes

✓ On track

🏥 Check-in
Easy

🕒 Seen on time
Every time

👥 Better outcomes
Together

AI Assistant

🤖 Patient Summary

- 47 y/o • Hypertension
- Well controlled
- No medication changes

Suggested Next Steps

- ✓ Continue current plan
- ✓ Lifestyle coaching
- ✓ Follow-up in 3 months

Care Insights

BP trend improving
Risk score: Low

A line graph showing a positive trend in blood pressure over time, with a green line rising from left to right.

AI support for better human decisions

- Xander Mol & Frank Kaptein | UWV
- Tineke Blom | ABN AMRO
- Rachel Kuijlenburg | Logius
- Hans Stokking | TNO



Logius
Ministerie van Binnenlandse Zaken en
Koninkrijksrelaties



ABN·AMRO



**Sector efficient
SLMs**



Question to the audience

**Was moving to US cloud
a mistake for Europe?**

Yes, or no?

Question to the audience

Is your organization using US Cloud for AI today?






A thriving European AI ecosystem.
**Sustainable, sovereign, and running on our
own ground.**

Three commitments behind it



Sustainable

Smaller models, a fraction of the energy and carbon.



On-prem

Your data and models stay within your own walls.



Dutch/EU cloud

Sovereign infrastructure, under European jurisdiction.

Why the status quo can't hold



Vendor lock-in

Lack of control & bargaining power



Obscure emissions reporting

Lack of clear disclosure



Regulatory pressure

SEAL will force companies to act

This isn't just our view.

Digital dependence of the
financial sector

risks, resilience and European autonomy

DeNederlandscheBank

EUROSYSTEM



AFM

US tech firms successfully lobbied EU to
keep datacentre emissions secret

Legally questionable confidentiality clause adopted almost
word for word from demands of Microsoft and trade groups

Central bank calls to reduce our digital dependence

Digital dependence of the financial sector

risks, resilience and European autonomy

DeNederlandscheBank

EUROSYSTEEM



And people discovered anti-sustainability lobbying

US tech firms lobbied to keep datacentre emissions out of EU regulation, the very data European buyers need to manage their footprint.


US tech firms successfully lobbied EU to keep datacentre emissions secret

Legally questionable confidentiality clause adopted almost word for word from demands of Microsoft and trade groups



Smaller, fine-tuned models win where it counts

Small fine-tuned LLMs beat large proprietary models on the tasks you actually run.



Open-sourced LLMs vs. proprietary small transformers on task accuracy — across encoder-only, decoder-only and encoder-decoder architectures.

Alternatives already exist, we build on, not against, them



“Een verantwoord alternatief op bestaande LLMs”
(A responsible alternative to existing LLMs)

A Dutch government-backed initiative to build a responsible, sovereign alternative to existing LLMs, a natural collaboration partner.



“Do it all with Mistral.”

A European open-weight model family already used as a foundation for sector fine-tuning across the EU.

You can join!



Collaboration initiative

After this challenge is over TNO, GPT-NL & partners will continue to collaborate to finetune language models. Join this initiative!



Fine-tuning factory blueprint

A documented “way of working”, an (AI) factory blueprint for fine-tuning that other sectors and organizations can adopt.

We can wait and see, or we can act now.

“

*Because we don't just care and hope, we care and act.
Swiftly, sustainably & collaboratively.*

Let's build sovereign AI together

Join us to turn Dutch and European data into secure, sustainable AI models for the sectors that need them most.



Scan me!



Wiebren van der Zee

ABN Amro,
Sustainable IT



Justus Bogner
VU Amsterdam



Pim Scheele

ABN AMRO,
Data Scientist



Pepijn van der Spek

Rijksorganisatie voor
Ontwikkeling,
Digitalisering en
Innovatie



Kevin Gong

Logius



AI Sovereignty quick scan

Colette Wibaut




AI SOVEREIGNTY QUICKSCAN



Marco Alberti • Siebrand Romeyn • Colette Wibaut • Marianne Schimmel
Marlena van Ooijen • David Kok • Hristo Valev

TNO



How are we doing on AI Sovereignty?



V+ ANALYSE

Europa stippelt snel een nieuwe koers uit: eigen tech eerst

Zowel Den Haag als Brussel nam de afgelopen week maatregelen om van het Amerikaanse techinfuus af te komen. Europa laat na jaren eindelijk zijn tanden zien, al zijn niet alle plannen even doortastend of effectief.



Frank Rensen

is techredacteur.

29 mei 2026, 14:09

What if we *can* be decisive and effective?

We create the ability to quickly assess whether an AI system aligns with an organization's sovereignty goals:

accelerating sharp choices

Every AI system introduces sovereignty risks. No practical tool exists to assess them.



Existing instruments focus on **cloud strategy**, security and IT portfolio.

Assessing sovereignty requires **new expertise**.

Every AI system introduces sovereignty risks. No practical tool exists to assess them.



Existing instruments focus on **cloud strategy**, security and IT portfolio.

Assessing sovereignty requires **new expertise**.

You need to know where you stand

Every AI system introduces sovereignty risks. No practical tool exists to assess them.



Existing instruments focus on **cloud strategy**, security and IT portfolio.

Assessing sovereignty requires **new expertise**.

You need to know where you stand

Without assessment capability:

- Vendor lock-in
- Loss of strategic control
- Increase of migration costs
- Unjustifiable decisions
- AI consumers, instead of stewards

Back at it, Colette

Claude Fable 5 is currently unavailable.

[Learn more](#)



Paste a doc, an email, or a question to get started



Haiku 4.5 Extended ▾



Becoming decisive and effective in AI sovereignty



Understand

Gain **overview** of the relevant components of your AI system.



Measure

A structured, repeatable **assessment**.



Improve

Actionable next steps on risks and opportunities.

Your quickscans (6)

Overview of all completed and ongoing quickscans.

+ Scan new AI system

Microsoft Copilot

Microsoft, United States

Legal

Operational

Economic

Data sovereignty

Biggest risk: AI model

Orientation

Siebrand Romeyn

17 12/05/2026

Proton Lumo

Proton AG, multiple open-wei...

Legal

Operational

Economic

Data sovereignty

Biggest risk: AI model

Orientation

Siebrand Romeyn

17 07/04/2026

ChatGPT Enterprise

OpenAI, United States

Legal

Operational

Economic

Data sovereignty

Biggest risk: Data sovereignty

Orientation

Siebrand Romeyn

17 22/04/2026

Mistral Le Chat

Mistral AI, France (EU)

Legal

Operational

Economic

Data sovereignty

Biggest risk: AI model infrastructure

Orientation

Siebrand Romeyn

17 06/05/2026

GPT-NL

TNO / SURF, Netherlands

Legal

Operational

Economic

Data sovereignty

Biggest risk: Operational

Orientation

Siebrand Romeyn

17 19/05/2026

DeepL

DeepL SE, Germany (EU)

Legal

Operational

Economic

Data sovereignty

Biggest risk: AI model

Orientation

Siebrand Romeyn

17 28/05/2026

GENERAL DETAILS

Microsoft Copilot · Orientation

 Edit details

✓ Components — ✓ Design choices — ③ **Sovereignty matrix** — ④ Mitigations

Step 3: Matrix deep-dive

Assess each of the four components from step 1 (columns) on the four sovereignty dimensions (rows). Answer a short question for each cell below. The colour in the matrix updates automatically based on your assessment.

	Application	App-infra	AI model	AI-model-infra
A. Legal ⓘ	— A1	— A2	— A3	— A4
B. Operational ⓘ	— B1	— B2	— B3	— B4
C. Economic ⓘ	— C1	— C2	— C3	— C4
D. Data sovereignty ⓘ	— D1	— D2	— D3	— D4

✓ Low risk ⓘ Medium ⚠ High risk ⓘ Unknown — Not filled in

[← Back to step 4](#)

QUICKSCAN RESULTS

Microsoft Copilot

Status: **Orientation**Date: **2026-05-12**Author: **Siebrand Romeyn**Team: **Digitaal Innovatielab (DIL)**

Design choices

Open-source application

Is the application open-source?

No

Microsoft 365 Copilot is a proprietary, closed-source product. Routing logic and system prompts are not public.

Open-weights AI model

Is the AI model open-weights?

No

Built on proprietary models from OpenAI and Microsoft. The weights are not available and model routing is not configurable.

EEA application

Was the application developed in the EEA?

No

Microsoft is based in the United States; core development and operation sit outside the EEA.

Exit strategy

Is there a feasible exit strategy?

Partial

Data is exportable, but deep Microsoft 365 integration makes switching costly.

Discover sovereignty assessments from organisations

Explore shared AI system assessments, risks, dependencies and mitigation strategies, compiled by public and private organisations.

Search

- Popular:
- Microsoft Copilot
- ChatGPT
- Gemini
- AWS Bedrock
- Mistral
- Proton Lumo

1.240+

Shared scans

318

Organisations

94

AI systems assessed

INDUSTRY

Government	412
Defence & Security	87
Financial services	198

AI applications

7 applications, crowdsourced sovereignty assessments

Sort: Last updated

Proton Lumo Proton AG, Privacy-focused assistant Medium	Microsoft 365 Copilot Microsoft, Productivity assistant High
--	---

A bottom-up approach that moves *with* the dynamic field



AI-specific & system-level

Accounting for different **layers** of AI, from organizational to societal impact.

A bottom-up approach that moves *with* the dynamic field



AI-specific & system-level

Accounting for different **layers** of **AI**, from organizational to societal impact.



Flexible & future-proof

Moving with **fast-evolving** technology and sovereignty risks.

A bottom-up approach that moves *with* the dynamic field




AI-specific & system-level

Accounting for different **layers** of **AI**, from organizational to societal impact.



Flexible & future-proof

Moving with **fast-evolving** technology and sovereignty risks.



Collaborative ecosystem

Shared knowledge base of quickscans.

Let's build, smartly.

1 Developing MVP

- Research-backed
- Robust
- Locally deployable tool

2 Pilot

- Fit for purpose
- Success criteria to go/no-go
- Scale after value-proof

Current coalition

TNO

RVO

UWV

Logius

ABN AMRO

- Methodology experts
- Use case owners
- Development capacity
- Governance participants

This is an **open invitation**
we **define roles together**.

Current coalition

TNO

RVO

UWV

Logius

ABN AMRO

- Methodology experts
- Use case owners
- Development capacity
- Governance participants

This is an **open invitation**
we **define roles together**.

As a partner

- Pilot
- Steering group seat
- Internal capability to assess your AI systems
- Access to shared knowledge base

A public building block,
not commercial software.

A shared responsibility for a better future

1

We safeguard European values

A job **close** to our
daily work.

2

We shape the solution together

Commitment in a
wicked problem.

3

The window is now

Building solution **now**,
setting the **standard**.

Let's *be* decisive and effective

Align your AI systems with your sovereignty ambitions —
accelerate sharp choices.

Time for a substitution!



We're taking this cup home.



AI SOVEREIGNTY QUICKSCAN

Thank you

Try it
yourself!



Siebrand Romeyn
(RVO)



Marianne Schimmel
(UWV)



Marlena van Ooijen
(Logius)



Marco Alberti
(ABN AMRO)




Hristo Valev
(TNO)



Colette Wibaut
(TNO)



David Kok
(UWV)



AI Sovereign roadmap



DO YOU REALLY CONTROL YOUR OWN AI?

AI: Someone else's design, agenda and geopolitical leverage



AI: Someone else's design, agenda and geopolitical leverage

Europe's Growing Fear: How Trump Might Use U.S. Tech Dominance Against It

To comply with a Trump executive order, Microsoft recently helped suspend the email account of an International Criminal Court prosecutor in the Netherlands who was investigating Israel for war crimes.

NY Times 20-6-25

AI: Someone else's design, agenda and geopolitical leverage

Europe's Growing Fear: How Trump Might Use U.S. Tech Dominance Against It

To comply with a Trump executive order, Microsoft recently helped suspend the email account of an International Criminal Court prosecutor in the Netherlands who was investigating Israel for war crimes.

NY Times 20-6-25

HOME / CYBERSECURITY

Microsoft Accused Of Sharing Dutch Officials' Data with U.S. Government

Microsoft turned over Dutch regulators' emails and names to U.S. Congress under CLOUD Act, targeting EU tech law enforcers



Rex Freiburger

May 22, 2026 · 2 min read



AI: Someone else's design, agenda and geopolitical leverage

Europe's Growing Fear: How Trump Might Use U.S. Tech Dominance Against It

To comply with a Trump executive order, Microsoft recently helped suspend the email account of an International Criminal Court prosecutor in the Netherlands who was investigating Israel for war crimes.

NY Times 20-6-25

US's Anthropic order exposes EU's AI dependency

Washington's export controls on Anthropic spark renewed calls for Europe to accelerate development of its own cutting-edge AI models.

Politico 14-6-26

HOME / CYBERSECURITY

Microsoft Accused Of Sharing Dutch Officials' Data with U.S. Government

Microsoft turned over Dutch regulators' emails and names to U.S. Congress under CLOUD Act, targeting EU tech law enforcers



Rex Freiburger

May 22, 2026 · 2 min read



AI: Someone else's design, agenda and geopolitical leverage

Europe's Growing Fear: How Trump Might Use U.S. Tech Dominance Against It

To comply with a Trump executive order, Microsoft recently helped suspend the email account of an International Criminal Court prosecutor in the Netherlands who was investigating Israel for war crimes.

NY Times 20-6-25

US's Anthropic order exposes EU's AI dependency

Washington's export controls on Anthropic spark renewed calls for Europe to accelerate development of its own cutting-edge AI models.

Politico 14-6-26

HOME / CYBERSECURITY

Microsoft Accused Of Sharing Dutch Officials' Data with U.S. Government

Microsoft turned over Dutch regulators' emails and names to U.S. Congress under CLOUD Act, targeting EU tech law enforcers



Rex Freiburger May 22, 2026 • 2 min read



Pentagon CTO urges Anthropic to 'cross the Rubicon' on military AI use cases amid ethics dispute

His comments come as the Pentagon is locked in a high-stakes dispute with Anthropic about the U.S. military's use of the startup's Claude AI model in real-world operations.

BY BRANDI VINCENT • FEBRUARY 19, 2026



INDEPENDENCE



RESILIENCE



CHOICE





EVERYONE KNOWS



We need more control



We need more sovereignty



The gap between goals and execution
is where sovereignty gets paralysed.




NO-ONE KNOWS



Where to start



What is really (not) needed



1 MODEL SOVEREIGNTY

1A Model

Who created the model, who owns its weights, and what kind of model is it.

1B Model Infrastructure

On what hardware the model runs for inference & training — where it sits and who owns it.

• L3 Structural Risk

US Proprietary Closed Model

GPT-4o, Claude 3.5, Gemini, US-aligned training, opaque weights, vendor-controlled lifecycle.

Superseded by any higher level

• L3 Structural Risk

US Hyperscaler Default GPU

Azure / AWS / GCP default GPU SKUs. Full CLOUD Act exposure on model weights and runtime traffic.

Superseded by any higher level

• L3 Quick Win

EU Proprietary Model API

Mistral commercial tier or Aleph Alpha Pharia API. EU-trained, GDPR-native, closed weights.

Superseded by L2 for Tier 2+ (not considered Change for Tier 2)

• L3 Quick Win

US Hyperscaler in EU Region + EU Data Boundary

Azure/AWS EU region pinning + EU Data Boundary contract. Compute remains US-jurisdictional.

Superseded by L2 EU sovereignty (GPU)

• L3 Achievable

HIGHEST PRACTICAL

• L3 Achievable

HIGHEST PRACTICAL

2 DATA SOVEREIGNTY

2A Data

Where the data physically sits and which jurisdiction can legally compel access.

2B Data Security & Access

Who controls access — IAM, DPAs, opt-out, customer-managed keys, zero trust.

• L3 Structural Risk

US Cloud Default Storage

Azure Blob / S3 in US region. Data may enter US training pipelines. CLOUD Act exposure.

Superseded by any higher level

• L3 Structural Risk

Vendor-Managed Keys + Default IAM

Provider holds the keys. Default tenant IAM. No DPA, no training opt-out enforced.

Superseded by any higher level

• L3 Quick Win

US Cloud + EU Data Boundary

Azure/AWS EU region + EU Data Boundary contract. Physical location EU, jurisdiction still US.

Superseded by L2 EU sovereignty (storage)

• L3 Achievable — 4-8 wks

BYOK in EU HSM + DPA + Training Opt-Out

Customer-managed keys (Azure Key Vault HSM EU / AWS KMS XKS). Signed DPA. Training opt-out flag set.

Advisory for Tier 2+ only

• L3 Achievable

HIGHEST PRACTICAL

• L3 Achievable

HIGHEST PRACTICAL

3 APPLICATION SOVEREIGNTY

3A Application

Who built the app that fuses model + data — incl. orchestration & guardrails.

3B Application Infrastructure

On what hardware the application itself runs — where it sits and who owns it.

• L3 Structural Risk

US SaaS Application

ChatGPT Enterprise, Microsoft 365 Copilot, Google Gemini for Workspace. App, model, data — all US-owned.

Superseded by any higher level

• L3 Structural Risk

US PaaS / Hyperscaler-Managed

Azure App Service / AWS Lambda / Vercel default. App runtime under US jurisdiction.

Superseded by any higher level

• L3 Structural Risk

Cloud-Native Orchestration + Vendor Guardrails

Azure Content Safety + Prompt Shield + LangChain calling OpenAI. All logic in US cloud.

Superseded by L2 (storage, app logic)

• L3 Achievable (containerised)

Containers on US Hyperscaler + EU Data Boundary

AKS / EKS in EU region + EU Data Boundary contract. Physical EU, jurisdiction still US.

Superseded by L2 EU sovereignty (virtual)

Click for analysis & instructions →

• L3 Achievable — 2-5 wks

HIGHEST PRACTICAL

• L3 Achievable

HIGHEST PRACTICAL

DATA · LEVEL 3 — HIGHEST PRACTICAL

EU-Sovereign Storage + EU Vector DB

Data Sovereignty Track

OVHcloud Object Storage or Hetzner Storage Box for blob/object data, plus Qdrant (Berlin) or Weaviate (Amsterdam) for vectors. EU-incorporated providers, full EU jurisdictional reach. The recommended baseline for RAG and agentic systems. To reach this state as a company: Migrate raw documents to EU-owned storage, deploy Qdrant or Weaviate in the EU, and regenerate embeddings without leaving EU jurisdiction. Bind the EU storage migration to customer-managed keys, DPA coverage, and training opt-out controls from the start.

CONCRETE NEXT STEPS

Action-oriented checklist — named tools, no narrative.

- 1 Migrate: `rsync s3:// → ovhcloud://` or `scaleway://` with `--s3-sse-c` encryption in transit; verify no plaintext intermediate copies in source or transit buckets
- 2 CRITICAL: migrate embedding generation to EU-hosted model (multilingual-e5-large or bge-multilingual on mi3) — EU storage with US-generated embeddings is partial sovereignty only
- 3 Deploy: Qdrant Cloud (EU region) or self-hosted Qdrant on mi3 Kubernetes cluster (Helm: helm.qdrant.io)
- 4 Alternatively: Weaviate (Amsterdam HQ) self-hosted via Helm on ai3
- 5 Configure: embedding pipeline to use EU embedding model (sentence-transformers on mi3) — not OpenAI/Azure embeddings
- 6 Validate: signed data residency report from provider confirming EU jurisdiction
- 7 Decommission: delete US-region copies after 30-day verification window with data integrity check

EU CLOUD SOVEREIGNTY FRAMEWORK

8 dimensions: Strategic Control · Legal/Jurisdictional · Data/AI · Operational · Supply Chain · Technology · Security/Compliance · Environmental. Grey bar = previous level · Green +N = improvement vs previous level.

Strategic Control		4/5	+2
Legal/Jurisdictional		5/5	+3
Data/AI Sovereignty		5/5	+2
Operational		4/5	+1
Supply Chain		4/5	+1
Technology		5/5	+3
Security/Compliance		4/5	+1
Environmental		4/5	+1

THE FUTURE SOCIETY (TFS) FRAMEWORK

5 EU strategic pillars for AI sovereignty. Delta vs previous level shown as +N.

Economic Competitiveness		4/5	+2
Resilience & Preparedness		5/5	+3
Security & Defence		4/5	+2
European Values & Identity		5/5	+2
Foreign Relations		4/5	+1

● L3 Achievable

HIGHEST PRACTICAL

EU Open-Weight Model

Mistral-Large-2, EuroLLM-9B, Codestral, Apache 2.0 weights you own and can serve anywhere.

RECOMMENDED FOR: **Chatbot**

RAG Agentic Coding

● L3 Achievable

HIGHEST PRACTICAL

EU Sovereign GPU Cloud

OVHcloud H100 / STACKIT / Scaleway / Hetzner. EU-headquartered. No CLOUD Act. Inference + training on vLLM.

RECOMMENDED FOR: **Chatbot**

RAG Agentic Coding

● L3 Achievable

HIGHEST PRACTICAL

EU-Sovereign Storage + EU Vector DB

OVHcloud / Hetzner object storage + Qdrant (Berlin) or Weaviate (Amsterdam). Full EU jurisdiction.

RECOMMENDED FOR: **RAG**

Agentic

● L3 Achievable

HIGHEST PRACTICAL

Zero-Trust RBAC

(mTLS), customer-managed keys in EU HSM, RBAC/ABAC, immutable audit logs.

RECOMMENDED FOR: **Chatbot**

RAG Agentic Coding

● L3 Achievable

HIGHEST PRACTICAL

Router + LLM Gateway

Langfuse (self-hosted). No vendor lock-in.

RECOMMENDED FOR: **Coding**

Agentic

● L3 Achievable

HIGHEST PRACTICAL

Full EU Jurisdiction

Everything under EU jurisdiction.

RECOMMENDED FOR: **Chatbot**

RAG Agentic Coding

Management Pitch Builder

Select your current level by clicking nodes in the dependency map. The pitch considers all selected nodes, cross-row dependencies, and the next coordinated move across the stack.

Clear

Generate pitch

MODEL

L3 - EU Open-Weight Model

MODEL INFRASTRUCTURE

L3 - EU Sovereign GPU Cloud

DATA

L3 - EU Sovereign Storage + EU
Vector DB

DATA SECURITY

L4 - Air-Gapped HSM + Offline Key
Custody

APPLICATION

L3 - Model-Agnostic OSS + Basic
Sovereign Guardrails


APPLICATION INFRASTRUCTURE

L4 - On-Prem Air-Gapped App
Hosting

CHOOSE
TARGET
STATE

-
GENERATE
DECISION
MEMO

Co-built. Validated. Shared.



The economics compound



Every deployment is more validation



Legitimacy can't
be copied



EVERYONE KNOWS



We need more control



We need more sovereignty



The gap between goals and execution
is where sovereignty gets paralysed.



NO-ONE KNOWS



Where to start



What is really (not) needed





EVERYONE KNOWS



We need more control



We need more sovereignty

	APPLICATION		MODEL		DATA	
	APPLICATION	APP INFRASTRUCTURE	MODEL	MODEL INFRASTRUCTURE	DATA	SECURITY & ACCESS
SEAL 0	US SaaS application	US managed orchestration	US frontier model	US hyperscaler GPU	US cloud storage	No controls (IAM, audit, protection)
SEAL 1	EU-region deployment	EU-region orchestration	EU model API	EU-region GPU instances	EU data residency	DPA with all vendors. Role-based access.
SEAL 2	Self-hosted EU app	Self-hosted orchestration	Self-hosted open model	EU sovereign cloud GPU	EU sovereign storage + vector DB	Customer-managed keys
SEAL 3	Fully open-source app	Sovereign orchestration stack	Fine-tuned EU model	Dedicated EU hardware	On-prem data + full encryption	Zero trust architecture. EU-only IAM.
SEAL 4	Custom sovereign app	100% EU orchestration stack	100% EU model supply chain	On-prem EU chips	Full data ownership (on-prem)	Complete audit ownership. No foreign jurisdictional exposure.



goals and execution
The gap between goals and execution is where sovereignty gets paralysed.



NO-ONE KNOWS




Where to start




What is really (not) needed

OUR OFFERING



Co-built. Validated. Shared.



**SOVEREIGNTY CANNOT
BE ACHIEVED ALONE**

CONTRIBUTE SECTOR AND SUBJECT EXPERTS

6+ ORGS – 0.5 FTE EACH



GEOPOLITICAL INCENTIVES

Strengthen Europe's strategic autonomy
in a multipolar world.



MORAL INCENTIVES

Build AI that reflects our values and
serves people and society.



FINANCIAL INCENTIVES

Unlock economic growth, innovation
and long-term competitiveness.



GEOPOLITICAL INCENTIVES

Strengthen Europe's strategic autonomy
in a multipolar world.



MORAL INCENTIVES

Build AI that reflects our values and
serves people and society.



FINANCIAL INCENTIVES

Unlock economic growth, innovation
and long-term competitiveness.



JOIN THE COALITION – The Time is Now



TOGETHER, WE BUILD A EUROPEAN AI ECOSYSTEM THAT IS **RESILIENT**, **TRUSTED** AND **COMPETITIVE**.



STOP DEBATING SOVEREIGNTY. START BUILDING IT.

The Modular **Sovereign AI** Roadmap



GEOPOLITICAL INCENTIVES

Strengthen Europe's strategic autonomy
in a multipolar world.



MORAL INCENTIVES

Build AI that reflects our values and
serves people and society.



FINANCIAL INCENTIVES

Unlock economic growth, innovation
and long-term competitiveness.



JOIN THE COALITION – The Time is Now



TOGETHER, WE BUILD A EUROPEAN AI ECOSYSTEM THAT IS **RESILIENT**, **TRUSTED** AND **COMPETITIVE**.

SPEAK TO OUR TEAM & LET'S BUILD THIS TOGETHER



Jeroen Franse
ABN AMRO



Joey van de Pasch
DNB



Erik Slingerland
Hogeschool Utrecht



Puck Stok
UvA



Joao F. Ferreira Goncalves
Erasmus University



Ana Oprea
UvA



Attila Popping
UWV



Michele Murgia
AI-Hub Noord Nederland



Dejan Zamurovic
TNO



Sovereign AI Roadmap Prototype
<https://cvds-nhlstenden.com/sovereignai.html>
(best on pc)



T I M E

F O R A

B R E A K !

Lotte de Groen

TNO innovation
for life



**Building the
future together**



Truly green AI

not a contradiction, it's a design discipline.

Let's control the impact

Processing is expensive, insight is not!

Solution

Smaller, fine-tuned models win where it counts

Small fine-tuned LLMs beat large proprietary models on the tasks you actually run.

The chart displays task accuracy (Y-axis, 0.0 to 0.7) for various models (X-axis) across three architectures: Encoder-Only Model, Decoder-Only Model, and Encoder-Decoder Model. The models are categorized by size (M) and type (Rule-based, ESD, RD, ZS, Raw/unscaled). The chart shows that smaller, fine-tuned models (e.g., Llama 3 (7B), Llama 3 (8B)) generally achieve higher task accuracy than larger proprietary models (e.g., Proprietary Decoder Only, Proprietary Encoder-Decoder).

Architecture	Model	Model Size (M)	Task Accuracy
Encoder-Only Model	Start (OOB)	~100	~0.35
	Start (Fine-tuned)	~100	~0.45
	Proprietary Encoder Only	~100	~0.55
Decoder-Only Model	Llama 3 (OOB)	~7B	~0.65
	Llama 3 (Fine-tuned)	~8B	~0.65
	Proprietary Decoder Only	~100	~0.65
Encoder-Decoder Model	Flan-T5 (OOB)	~100	~0.55
	Flan-T5 (Fine-tuned)	~100	~0.65
	Proprietary Encoder-Decoder	~100	~0.65

Legend: Rule-based (Blue), ESD, RD (Light Blue), ZS (Orange), Raw/unscaled (Red).


Small fine-tuned transformers on task accuracy — across encoder-only, decoder-only and encoder-decoder architectures.


The screenshot shows the 'AI Systems Sovereignty Quickscan' interface. At the top, there are tabs for 'Overview', 'Shared scans', and 'New'. The main heading is 'Your quickscans (5)', with a subtitle 'Overview of AI-completed and ongoing quickscans.' A search bar is present with the text 'Search by name, supplier or type...' and a button 'Scan new AI system'.

The interface displays five quickscan cards, each representing an AI system. Each card includes the system name, a risk level indicator (High, Medium, Low), a table of risk scores, and a 'Biggest risk' label. The risk scores are represented by a 3x3 grid of colored squares (green for low, orange for medium, red for high).

System Name	Risk Level	Legal	Operational	Economic	Data sovereignty	Biggest risk
Microsoft Copilot	High risk	Green	Green	Green	Red	AI model
Proton Luma	High risk	Green	Green	Green	Red	AI model
ChatGPT Enterprise	Medium risk	Green	Green	Green	Red	Data sovereignty
Mistral Le Chat	Medium risk	Green	Green	Green	Red	AI model
GPT-4o	Low risk	Green	Green	Green	Green	AI model
DeepL	Low risk	Green	Green	Green	Green	AI model

A circular inset in the bottom left corner shows a portrait of a woman with dark hair, smiling.







Impact


Research


Effort























FINALE
Futureproof AI
June 2026

TNO innovation
for life

Thank you

