



EUROPEAN
COURT
OF AUDITORS

EN

Artificial Intelligence initial strategy and deployment roadmap

2024-2025

Contents

Introduction	5
Using AI for auditing/auditing AI	7
AI in auditing: enhancing capabilities	7
Auditing AI: ensuring integrity and effectiveness	8
Goals and objectives	10
Goal 1: improve operational efficiency in audit through AI tools	10
Objective 1.1: Increase ECA staff knowledge on how to use AI	10
Objective 1.2: Ensure that the ECA is technically ready for AI	10
Objective 1.3: Introduce AI tools to support the audit process	11
Goal 2: build the ECA's ability to audit AI-based projects, systems and processes	11
Objective 2.1: Increase ECA staff understanding on how AI works and its associated risks	11
Objective 2.2: Add AI aspects to our audit methodology	12
Goal 3: add value and contribute to EU-wide and international discussions on AI	12
Objective 3.1: Contribute actively to EU and international working groups	12
Objective 3.2: Make the ECA knowledge on AI available to other EU SAIs	12
Risk analysis and mitigation strategies	14
General AI risks	14
Limited transparency and oversight on AI's logic	14
Biased or discriminatory outputs	14
Privacy, data protection and security issues	15
Risks linked to implementation	15
Insufficient resources and skills for an effective deployment	15
Deployment of already outdated services	16
Reputational risks	16

Proposal for an AI roadmap	17
Action 1: create and follow-up a detailed communication plan	19
Action 2: create an AI training path for ECA staff	21
Action 3: enhance the DATA services offered	22
Action 4: set-up a project proposal portfolio for new or enhanced tools	24
Action 4.1: Creation of an “Audit Draft Assistant” tool	25
Action 4.2: Customise and deploy a tool to retrieve information from large document sets, using an interactive approach based on Questions & Answers	26
Action 4.3: Extend the DORA tool with an AI backend for summarising and other natural language processing tasks	27
Action 4.4: Semantic search engine on ECA own production	27
Action 4.5: Training a chatbot with internal knowledge of audit methodology from AWARE and other controlled-quality sources	29
Action 4.6: Assess the EC DORIS tool for survey analysis	29
Action 4.7: Add AI components to robotic process automation (add intelligence to automation of repetitive tasks)	30
Action 4.8: Add access to AI to existing data science and visualisation tools	30
Action 5: creation of an AI competence centre at the ECA and related governance	30
Action 6: foster interinstitutional and international cooperation	31
Closing remarks	34
Annex I – Overview of the main regulatory frameworks	35
Main regulatory approaches worldwide	35
Regulating AI safety and security	38
Impact for the ECA	38
Annex II – What are other institutions doing on AI?	39
Interinstitutional and international working groups	39
European Commission	41
Other EU Institutions	42
EU Supreme Audit Institutions	42
Other non-EU audit Institutions	45

Annex III – Develop or buy? Cloud-based commercial products versus on-premises tool	47
Introduction and definitions	47
Available options, as of February 2024	49
Conclusions and proposal: start a dual approach in AI deployment (develop and buy)	50
Annex IV – Risk mapping and analysis	51

Introduction

01 In October 2023, following a series of recent breakthroughs in the field of artificial intelligence (AI), the European Court of Auditors (ECA) started an internal reflection on its application within the institution with the paper “Opportunities and challenges in responding to and using artificial intelligence”. As a follow-up, this document provides further analysis of the key issues and outlines an initial AI strategy and a deployment roadmap.

02 The current ECA strategy covers the period 2021-2025. Preparation of the next strategy is likely to start soon, and will need to fully reflect AI aspects. This initial strategy should allow implementation of AI activities to start in a structured way, as well as constitute an input for the next ECA strategy.

03 AI is a deeply transformative technology, but it will not replace the critical thinking and professional judgement of our auditors. We envision a use of AI that will assist our staff and augment their capabilities, allowing for more timely, efficient and soundly-based audits. We will uphold the highest ethical standards in the deployment and use of AI systems, safeguarding against biases, prioritising transparency, and applying caution and professional scepticism. We will remain faithful to our core values of independence, objectivity, integrity, transparency, and professionalism. Our vision is to leverage AI as an enabler to ultimately drive greater accountability, transparency, and trust in the EU.

04 In this document, we explain the use of AI for audit and the audit of AI systems, and clarify the distinction between them. We then propose a set of goals and objectives for the ECA and, after an analysis of the risks, an AI roadmap with concrete actions.

05 We consider the proposed actions to be realistic, based on successful tests and validated proof-of-concepts, as well as experience in other institutions. However, since the domain is still in its infancy and few professionals have yet acquired sufficient experience with these technologies, we are not able to provide a sufficiently reliable estimate of the time and costs necessary for the required IT developments. For those, we propose that formal IT projects are launched that will be analysed, prioritised and approved by the relevant IT governance bodies.

06 To provide the necessary context for this document, we briefly reviewed the current AI regulatory framework (see [Annex I](#)), we collected information on what other institutions and EU Supreme Audit Institutions (SAIs) are doing in the field of AI (see [Annex II](#)) and we performed a preliminary “make or buy” analysis on AI tools, especially concerning the differences between cloud-based or on-premises solutions (see [Annex III](#)).

07 Despite the current enthusiasm, we do not consider that (generative) AI will revolutionise our core processes or increase productivity multi-fold, at least not in the short and medium term. We believe that many, if not all, professional daily tasks can be helped by AI in the coming years, with an expected productivity increase of around 10 %¹. While this figure may seem modest, if applied transversally the cumulative effect in efficiency and effectiveness can be significant.

08 The field of AI is developing rapidly, and so therefore is the material that could be included in this document. We chose February 2024 as a practical cut-off date for including new information. However, we outline a communication plan where we explain how ECA colleagues will regularly receive updated information.

¹ See for example “[The economic potential of generative AI: The next productivity frontier](#)”, McKinsey, January 2023.

Using AI for auditing/auditing AI

09 The impact of AI in the field of audit is twofold: on one hand AI has the potential to enhance the efficiency and effectiveness of the audit process, while on the other hand it requires auditors to rethink and adapt their methodology when such technology is used by our auditees. As with any technological innovation, AI brings its own challenges and necessitates a deep understanding. This chapter delves into the dual role of AI for public audit institutions, including using AI for auditing purposes and the future challenge of auditing AI-based systems.

AI in auditing: enhancing capabilities

10 AI can significantly augment the capabilities of audit teams, particularly by aiding time consuming tasks such as:

- **document analysis:** AI algorithms can efficiently sift through large multilingual document sets, identifying patterns, anomalies, and key information that might take human auditors significantly longer to process and with less reliable results;
- **risk assessment:** AI can enhance risk assessment procedures by rapidly analysing data trends and helping to identify potential areas of concern, thereby making the audit quicker to plan and more relevant; and
- **routine tasks:** automating routine tasks frees up auditors to focus on more complex aspects of the audit process, thereby increasing overall efficiency.

11 In addition, generative AI can support auditors with drafting documents and guiding them in the application of our methodology. We outlined several practical use cases in our reflection paper Opportunities and challenges in responding to and using artificial intelligence.

12 To fully leverage AI in auditing, comprehensive training for audit staff is crucial. The training should focus not only on the operational aspects of AI tools, but also:

- **understanding AI outputs:** staff must be equipped to interpret AI outputs accurately, recognising their limitations and potential biases, and understanding that AI aids but cannot replace human judgement; and
- **ethical considerations:** training should cover ethical considerations in AI usage, ensuring transparency and fairness in automated processes.

Auditing AI: ensuring integrity and effectiveness

13 As AI systems become integral to our auditees' business processes, we must be prepared to audit them.

14 A key issue to consider is whether, and how, audited entities have addressed AI in their conduct of business, in particular whether they have developed AI strategies and established related governances. If they are not using AI, we should assess the risks this brings to the auditees' business processes. In addition, it is also relevant to verify how our auditees address risks associated with the use of AI by third parties with whom they have contractual relationships.

15 Auditing AI is a new component that can be added to general IT audit methodologies; this involves assessing how AI is used, the impact on output, how well it works, the decision-making followed, and compliance with the auditee's ethical and other standards. [Table 1](#) provides a non-exhaustive list of the elements that auditors need to consider:

Table 1 – Main elements to consider when auditing AI systems

Key element	Description
Data integrity and bias mitigation	AI models are as good as the data they are trained on. Operators of AI systems should ensure that the data used is accurate, representative and free from biases that could skew output.
Compliance with regulations and standards	AI systems must comply with relevant legal and ethical standards, thus ensuring that such systems are used responsibly, protecting sensitive data and upholding public trust.
Transparency and explainability	To maintain trust and accountability, AI systems should be transparent and their decision-making processes explainable. In other words, it should be possible to understand how an AI model produces outputs and reaches conclusions.
Performance monitoring and evaluation	Regular assessments of AI systems are necessary to ensure they perform effectively and continue to meet their objectives over time.
Impact on workforce and processes	Auditing AI may also encompass the impact on the auditee workforce and processes. The integration of AI should not compromise the quality of work or lead to loss of skills in critical processes and procedures.

16 Preparing to audit AI systems encompasses:

- **Building AI literacy** through dedicated training for both staff and management;
- **Adapting our audit framework** by including AI-specific checklists and procedures;
- **Monitoring and collaborating** with professional associations and international working groups developing audit standards and frameworks;
- **Continuous learning** to ensure that our methodology evolves together with advancements in technology, regulation and policies.

Goals and objectives

Goal 1: improve operational efficiency in audit through AI tools

17 The use of technology in audit is one of the key enablers for the goals outlined in the ECA's 2021-2025 strategy. Through effective use of AI, we aim to progressively enhance the efficiency of our audit procedures, improve our analytical capabilities, and foster a more dynamic audit approach.

Objective 1.1: Increase ECA staff knowledge on how to use AI

18 To harness the potential of AI in our audit processes, we must develop the AI literacy of our staff. We have already started implementing a comprehensive training programme that focuses on providing not only the basic understanding of AI technologies and their functioning, but also practical guidance and techniques that our staff may need in the context of their audits. An adequate AI literacy of our staff is a necessary pre-condition for the successful deployment of AI tools.

Objective 1.2: Ensure that the ECA is technically ready for AI

19 Preparing our technical infrastructure, policies and procedures for the integration of AI is a crucial step towards achieving our main goal. We will conduct a thorough review and enhancement of our existing systems and processes. This may involve upgrading our technological infrastructure to support AI capabilities, revising policies to incorporate AI ethics and governance, and adapting procedures to optimize AI implementation.

20 Recently, the International Organisation for Standardisation has published a new standard on Artificial Intelligence Management System². The ECA should consider implementing the parts of the standard that are relevant.

² ISO/IEC 42001:2023.

Box 1**What is ISO/IEC 42001?**

ISO/IEC 42001 is the world's first AI management system standard, providing valuable guidance for this rapidly changing field of technology. It addresses the unique challenges AI poses, such as ethical considerations, transparency, and continuous learning. For organisations, it sets out a structured way to manage risks and opportunities associated with AI, balancing innovation with governance.

Source: <https://www.iso.org/standard/81230.html>.

Objective 1.3: Introduce AI tools to support the audit process

21 The introduction of AI tools in our audit process aligns with our strategic goal of using technology to enhance audit quality. We will identify and implement state-of-the-art AI tools that can process and draft text, analyse large datasets efficiently, identify patterns, and provide insights that go beyond the scope of traditional audit methods. These tools will assist our auditors and augment their capabilities, increasing the efficiency of the audit procedures and potentially allowing for more in-depth audits. The deployment of AI tools will be carried out with a commitment to carefully ensure compliance with data protection and other relevant regulations. We will make these tools available in AWARE by introducing links in the relevant pages.

Goal 2: build the ECA's ability to audit AI-based projects, systems and processes

22 The ECA is committed to developing its capabilities and capacity to audit AI-based projects, systems, and processes in order to provide strong and reliable evidence in a challenging and changing environment. This goal aligns with our broader strategy of adapting to digital transformation trends. As AI starts to be used within financial and policy initiatives within the EU, we must evolve our IT audit approach and tools accordingly.

Objective 2.1: Increase ECA staff understanding on how AI works and its associated risks

23 The first step on the journey to auditing AI-based systems is to ensure that our staff acquire an in-depth understanding of AI technologies and how they can be used, and are aware of the risks they entail. We will initiate targeted trainings and workshops to deepen our auditors' knowledge of AI algorithms, data dependencies and potential biases. We will also provide our staff with the insight to question and critically assess AI

systems. This will allow our audits to evaluate both the performance and the integrity of these technologies.

Objective 2.2: Add AI aspects to our audit methodology

24 Developing our methodology to include specific considerations for auditing AI-based processes is another essential step towards achieving this goal. We will update and expand our IT audit methodology and other guidance in AWARE and other guidance in AWARE to incorporate AI-specific issues. These will cover aspects such as data quality, algorithmic transparency, ethical use, and compliance with EU regulations and standards. This should allow more thorough and informed assessments of AI-based processes and outputs.

Goal 3: add value and contribute to EU-wide and international discussions on AI

25 Interacting with institutional stakeholders and cooperating with peers are two key enablers of our 2021-2025 strategy. The ECA will be in a position to add value and contribute actively to professional discussions on AI, notably in respect of audit but also concerning financial management. This goal contributes to keeping the ECA at the forefront of the public audit profession.

Objective 3.1: Contribute actively to EU and international working groups

26 To achieve this objective, we aim to be involved actively in forums, conferences and working groups dedicated to AI, contributing our unique perspective and learning from others. We will share the challenges faced, successes achieved and lessons learned in implementing and auditing AI systems. We aim to build on the ECA as an active and respected practitioner in the AI space, contributing to AI governance and policy at both the EU and global levels.

Objective 3.2: Make the ECA knowledge on AI available to other EU SAIs

27 In accordance with our strategy of cooperation with peers, we commit to making our accumulated knowledge and experience with AI readily accessible to other EU supreme audit institutions that are seeking to optimise their AI deployment. We will share best practice, guidelines and tools that we have found to be effective by making them available in the public version of AWARE. By creating a repository of resources and facilitating workshops and knowledge-sharing sessions, we aim to assist other SAIs in

navigating the complexities of using AI in audit. This initiative not only underscores our role in fostering a collaborative audit community but also contributes to the collective advancement and standardisation of AI auditing practices across the EU.

Risk analysis and mitigation strategies

28 While AI holds significant promise for enhancing our audit capabilities, we must be fully aware of the associated risks and ensure we have effective mitigation in place. The *Risk mapping and analysis* in *Annex IV* contains the full list of the identified risks with the related assessment and proposed mitigation. The following paragraphs describe the most significant ones.

General AI risks

Limited transparency and oversight on AI's logic

29 Many AI systems, especially those based on deep learning, are seen as “black boxes” due to their complex and opaque internal operational processes. It is sometimes difficult to understand and explain the processes that the systems follow to produce their results. This limits the transparency of these systems and puts into question the reliability of outputs, especially in the context of audit procedures. We can mitigate this risk by:

- choosing AI models whose operation is easier to understand and follow;
- ensuring that the model provides always a step-by-step description of the process or the software code used to reach the results, thus allowing for oversight and (when applicable) reperformance;
- providing auditors with training and methodological guidance for:
 - (1) understanding the outputs and their limitations;
 - (2) ensuring that the model provides the code and/or the process steps used to reach a given output; and
 - (3) documenting how AI outputs were used in the audit work.

Biased or discriminatory outputs

30 The quality of an AI system's output depends on the data that was used to train the system. The system can inherit potential biases present in its training data, thus leading to biased or even discriminatory outputs. This in turn may affect the reliability of the output of such systems in the context of our work. It is therefore important to define

clear internal rules on the acceptable use of AI outputs and to provide staff with training and methodological guidance. In addition, we can further mitigate this risk by:

- regularly reassessing the AI models we use, and favouring those with less biased training data; and
- when applicable, using models that produce their outputs based on the ECA's knowledge base (e.g. AWARE, Assyst, published reports, etc.), rather than relying on knowledge from the original training data.

Privacy, data protection and security issues

31 AI systems involve the use of large amounts of data, which can include sensitive or personal information. As it is the case for other technologies or platforms, these risks from using sensitive and personal data will be addressed by performing thorough data protection and security assessments on the selected solutions or providers. The technical trade-off between operating AI systems in our local infrastructure and using an online commercial service will be duly considered as well from a data protection and security perspective. When possible, preference should be given to providers based in the EU. In addition, access to online commercial AI services for work purposes should be restricted to corporate/enterprise licenses, which can provide stronger contractual guarantees and safeguards.

Risks linked to implementation

Insufficient resources and skills for an effective deployment

32 The ECA has a limited IT budget. Therefore, the funds to acquire or operate the best AI systems at scale may not be available in part or in full.

33 Furthermore, there may not be sufficient in-house skills to convert initial proofs of concept into products available to all staff, thus slowing down adoption or increasing costs. Even when resorting to external expertise, experts may not be easily available in Luxembourg.

34 To mitigate these risks, we will continue exploring possibilities for interinstitutional cooperation with the aim of sharing skills, infrastructure, solutions, licensing and other costs. We are also establishing an AI competence centre (see Action 5) to pool together internal expertise and skills from different parts of our organisation. It is also worth noting that despite not being the most advanced in the industry, some cheaper or less resource-intensive models may still be adequate for our needs.

Deployment of already outdated services

35 As the development in the AI domain has been extremely quick in the past two years there is a risk that technology will evolve faster than our ability to deliver end-to-end services to our staff. When we start a project we tend to use the most promising technology available in that moment, but when the tool is finally available to users it may be perceived as already obsolete and not used (especially if users clearly see better alternatives).

36 Factors that are likely to slow down the deployment include: validating several use cases through pilot projects; ensuring compliance with regulations (e.g. data protection, upcoming AI act); performing security assessments and devising a security plan; and adapting our processes and infrastructure to this new technology.

37 We can mitigate these risks by keeping a flexible approach and appropriate infrastructure, exploring different options in parallel, and choosing the most suitable (and less risky) one for each use case. For online commercial products, these risks are less relevant since the upgrade to a more advanced model may simply require the purchase of new licenses. For AI models that we plan to operate in our premises, we will evaluate whether we need the highest performer available on the market, or whether a less advanced (and therefore cheaper) model could be good enough for our specific use cases. Finally, giving preference to EU providers and “open source” models may also help to obtain a faster legal and data protection clearance.

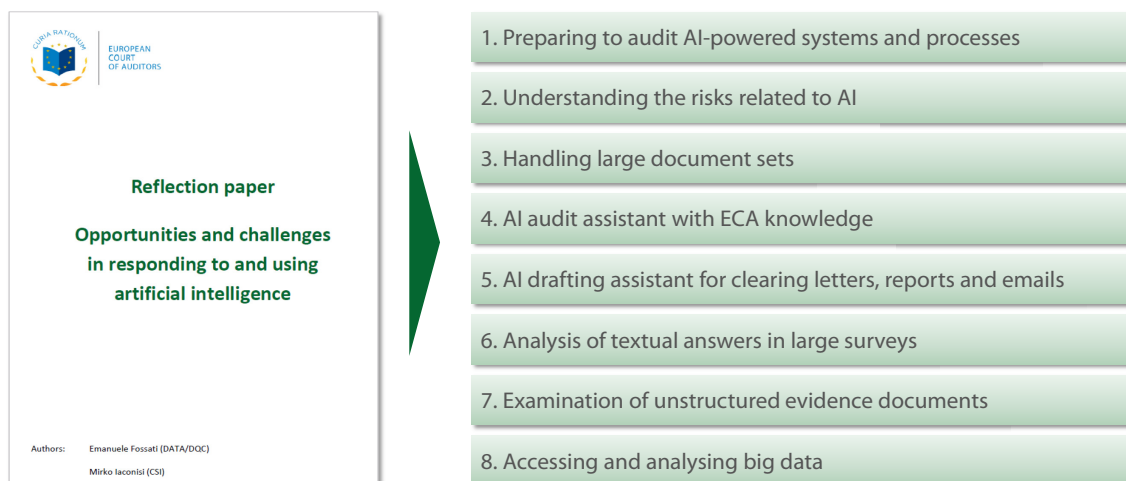
Reputational risks

38 It is important to provide staff with safe, well-performing corporate AI tools to avoid the risk of potential unauthorised use of AI via private accounts. Misuse or a lack of due diligence in the use of AI in our work may lead to reputational damage for our institution. For this reason, clear guidance and adequate training of our staff are essential.

Proposal for an AI roadmap

39 In the reflection paper “Opportunities and challenges in responding to and using artificial intelligence”, we listed a series of audit needs and practical use cases related to AI (see [Figure 1](#)).

Figure 1 – Audit needs and use cases related to AI



Source: ECA, DEC 113/23.

40 To address these needs and also enable the implementation of future use cases, we must establish a clear and coherent AI framework. This includes setting up a competence centre with appropriate governance, providing a specific training path, preparing technical infrastructure and IT tools, enhancing the services provided by the DATA team and the Directorate for Information, Workplace and Innovation (DIWI), and creating a portfolio of project proposals. In addition, we should foster interinstitutional cooperation and exchange of knowledge and practices. In the following paragraphs, we describe the proposed actions to achieve the three strategic goals.

Table 2 – Summary of proposed actions

Action	Description	Related to need described in the reflection paper	Strategic goal	Estimated completion
1	Create and follow-up a detailed communication plan	1, 2	1, 2, 3	Q2/2024

Action	Description	Related to need described in the reflection paper	Strategic goal	Estimated completion
2	Create an AI training path for ECA staff	1, 2	1, 2	2024, then regular updates
3	Enhance the DATA services offered	3, 6, 7, 8	1, 2, 3	2025
4	Propose a portfolio of projects for new or enhanced tools	3, 4, 5, 6, 7	1	2024
5	Creation of an AI competence centre at the ECA and related governance	All	1	2024
6	Foster inter-institutional cooperation	All	1, 2, 3	2024, then regular updates

41 Some actions involve activities that are already ongoing in various directorates. We include them here to create a comprehensive document that reflects all aspects of AI. Other actions will require dedicated projects to be established, involving mostly DIWI and managed by the existing ECA governance bodies.

42 The deployment roadmap is based on two principles:

- we prioritise solutions that we can acquire on the market or from the open-source community (subject to the data protection officer's approval), or that are developed by other institutions or interinstitutional working groups; and
- we develop our own AI solutions when we need to apply them to very specific activities of the ECA, or when the sensitivity of data and information requires a strict control over the tool.

43 In preparing this document, we have not been in the position to estimate the required budget and resources for IT projects and for the actions detailed below, because we do not yet have sufficiently detailed specifications and analysis. In any case,

an accurate estimation can only be obtained when we launch individual project proposals and subject them to the usual IT processes. These project proposals, once received, will be analysed by DIWI for resource allocation and prioritisation. DATA team's resources involved in AI-related activities have been budgeted for in the 2024 work programme of the Directorate of the Audit Quality Control Committee (DQC).

44 As a quick and non-exhaustive reminder, the general process we follow to launch IT activities starts with customer relationship management (CRM) requests, that are filtered and assessed by DIWI management. The approved ones become short innovation projects (SIPs) or full IT projects, and follow the related processes and governance.

45 To move ahead quickly, we propose to use the SIP approach, i.e. the short innovation projects introduced by DIWI in 2022 to facilitate the process of innovation with tangible results. The SIPs should be used when the technology proposed is still not mature and/or requires some experimentation to assess the technical feasibility, evaluate the overall complexity and to provide a general estimation of the final costs and realistic delivery timing. SIPs are well defined smaller projects with specific goals and duration, and involve a pre-defined and controlled set of resources (example: one professional for three months and software purchase for €40 000). The SIP is run by a business project manager (for example, an auditor) supported by an IT project manager. The expected outcomes of a SIP are a completed Proof of Concept (or a minimum viable product³) and a documented analysis of what it takes in terms of resources and time to put in place or create the final product (tool and/or service).

Action 1: create and follow-up a detailed communication plan

46 One of the main objectives of this strategy is to create a shared AI framework to justify and foster all the activities that are launched and managed by various directorates. Therefore, the strategy and the proposed roadmap has to be properly communicated across the ECA and externally.


47 Strategic communication is never a one-size-fits-all exercise. Not only do different levels of management have different expectations and responsibilities, but also different departments may be more sensitive to actions that have specific impacts on their daily operations.

³ An MVP is a development technique in which a new product or website is developed with sufficient features to satisfy early adopters. The final, complete set of features is only designed and developed after considering feedback from the product's initial users. This concept allows developers and companies to test their product hypotheses with minimal resources and to efficiently integrate user feedback to iterate and improve the product.

48 It is also important to remember that each technological revolution follows a predictable cycle, formalised by the consultancy Gartner as the “Gartner Hype Cycle” (Figure 2). Our understanding and experience suggest that the initial “technology trigger” phase finished at the end of 2023 and we are now around the “peak of inflated expectations”. What usually follows is a period of disillusion called “trough of disillusionment” that will last throughout 2024 where negative sentiments fuelled by the inevitable initial difficulties dominate the public sphere.

49 However, focusing on what we want to realistically achieve, we plan to work steadily in 2024 to find ourselves in the last two sections of the hype cycle: the slope of enlightenment and the plateau of productivity, where AI starts being integrated into our daily work and starts showing the expected increase in productivity and effectiveness.

Figure 2 – An adaptation of the Gartner Hype Cycle



Source: ECA, on the basis of the Gartner Hype Cycle.

50 The objective of the communication plan is to keep the ECA knowledge on AI up to date. Concretely, we plan to use several channels to achieve our communication goals:

Table 3 – Communication plan for Artificial Intelligence

Audience	Messages	Format	Frequency
ECA Members and senior management	Advancement status, high level issues	QA/DEC documents	6 months or on request by AQCC/AC/College

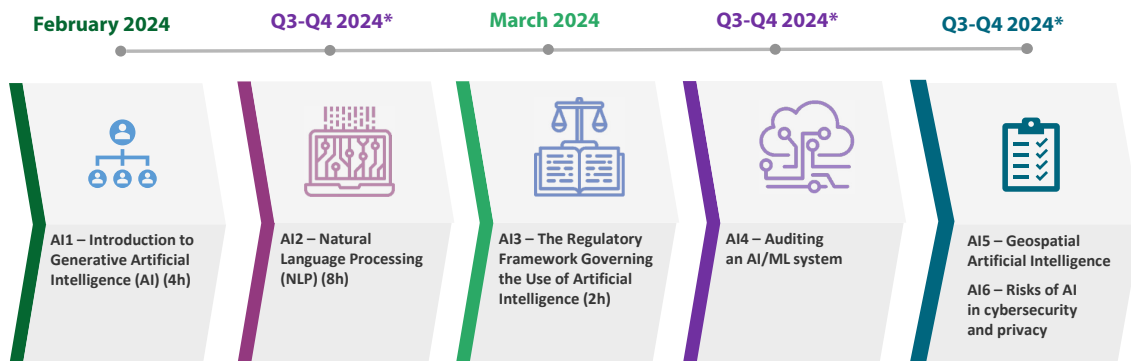
Audience	Messages	Format	Frequency
ECA Staff	General knowledge about AI applications, ethics and utilisation in concrete cases of our audit work	Subject briefs	2/year
	Achievements on AI, new tools, competence centre, etc.	COMPERS and ECA News	4/year
	General information about AI	Knowledge Node on AI	Weekly
	Coordination of the AI competence centre activities	ECA AI Network: a forum open to all ECA staff created on Teams as part of the Competence Center	Daily
EU SAls	Cooperation and knowledge exchange	EU network activity (EU SAI only) and TiNA events (EU SAI + other stakeholders)	2/year
SAls worldwide	AI Activities at the ECA	Specific media (e.g. INTOSAI journal, EUROSAl magazine) and/or events	1/year
Extended professional network	General information about AI for audit	Personal LinkedIn	Ad hoc
Other Institutions AI decision-makers	Compare what the ECA is doing VS others	Bilateral meetings	Ad hoc

Action 2: create an AI training path for ECA staff

51 The DATA team, in cooperation with the training team, established a new AI training offering in December 2023, tailored to the ECA's specific needs in 2024, the Legal Service will also contribute with a session on legal risks and copyrights aspects of AI.

52 In contrast to the generic training available via EU-Learn, the sessions organised by DATA team professionals are designed to precisely target the needs of ECA staff (auditors, support services, etc.) and to leverage the DATA team experience in working in audit tasks. Consistently positive feedback received after the sessions shows the appreciation and the perceived value of training delivered by DATA team members, motivating us to further expand the path in 2024 and 2025.

Figure 3 – The AI training path for the ECA staff



Source: ECA.

53 As soon as the ECA acquires a generative AI tool, we will add specific training for auditors on its use. This should help enable auditors to obtain better outputs from AI (therefore mitigating the risk mentioned in paragraph 29). Another possibility we are currently evaluating is to add a specific session on AI ethics, again using internal experts.

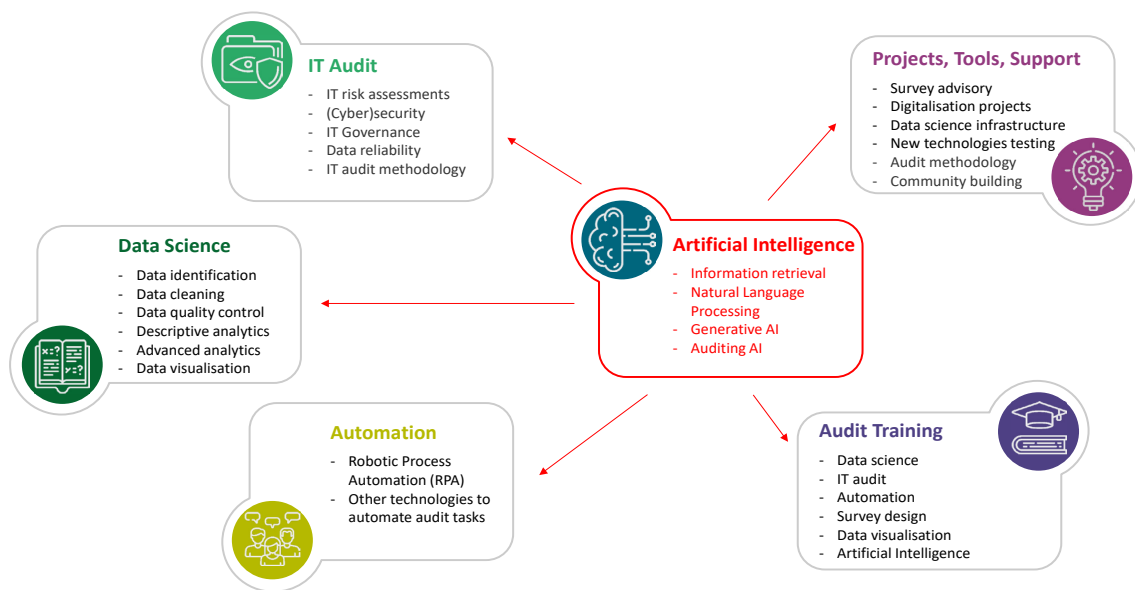
54 To complement this training path, DATA and DIWI also offer individual coaching sessions responding to specific requests related to audit tasks.

Action 3: enhance the DATA services offered

55 Artificial Intelligence is not necessarily a distinct and separate service, but more an enabler that can improve productivity across a wide range of existing activities in our institution. The DATA team is among the early adopters, and they plan to:

- enhance their existing services to include AI components; and
- assist auditors in assessing how to include AI in their strategic and daily operational activities.

Figure 4 – DATA team services



Source: ECA.

Table 4 – Enhancement of DATA services

Service	Planned enhancements with AI
Data Science	<ul style="list-style-type: none"> ○ Use of AI to analyse unstructured data (i.e. for risk-based sample extraction) ○ Use of AI to categorise lists of projects and/or to identify relevant issues from generic descriptions ○ Introduce the use of Machine Learning for anomaly detection in large data sets ○ Integrate generative AI in our toolset for data science, in cooperation with DIWI (see also Action 4.8, p. 30) ○ Increase DATA team productivity using generative AI code assistants when supporting audit tasks

Service	Planned enhancements with AI
IT Audit	<ul style="list-style-type: none"> ○ Perform audits of AI systems and propose future audit work on AI to gain expertise ○ Enhance the current IT audit methodology by adding AI-related risks and controls and drafting standard audit programmes for AI systems ○ Draft internal guidance for auditing AI systems in line with ECA's methodology in AWARE, legal requirements and international standards (e.g. ISO/IEC 42001:2023 – Artificial Intelligence Management System)
Automation	<ul style="list-style-type: none"> ○ Launch of a pilot to add AI capabilities to ECA robots, for example to analyse invoices, categorise emails, etc.
Survey advisory	<ul style="list-style-type: none"> ○ Test AI-based survey analysis tools (for example, the European Commission DORIS), and provide guidance to auditors on how to use them
Training, coaching and liaison with key stakeholders	<ul style="list-style-type: none"> ○ Active participation in EUROSAI and INTOSAI working groups ○ Explore opportunities to collaborate with professional associations (e.g. ISACA) to improve AI auditing capabilities ○ The DATA team will offer specific coaching sessions on AI for specific needs that may arise from an audit task (see also Action 2: Creation of an AI training path) ○ The DATA team will also offer expertise on artificial intelligence to support private office cooperation with their respective national audit institutions (organisation of events and visits, preparation of presentations, etc.).

Action 4: set-up a project proposal portfolio for new or enhanced tools

56 With this action, we will prepare a series of detailed project proposals, each designed to harness the potential of new technologies or to significantly enhance our existing IT tools. These proposals will be formulated following a thorough analysis of our

current capabilities, resources, and the specific demands coming from auditors and non-auditor staff. We consider that they can be realistically implemented in 2024 and 2025, subject to the prioritisation of AI activities and the availability of sufficient budget for IT developments.

Action 4.1: Creation of an “Audit Draft Assistant” tool

57 To implement the service most requested by auditors, we propose create and make available an AI assistant that can:

- create text paragraphs from ideas or a list of concepts given by auditors;
- proof-read documents to improve the quality of the English drafting, and align it to the ECA house style; and
- summarise documents and find links between them.

58 Preliminary tests, also described in the previous document DEC 113/23, show that only very large language models (LLMs) can process the English language effectively. In contrast, AI/LLMs that would be available to be run locally at the ECA do not have the necessary sophistication, due to their relatively small size.

59 Therefore, the only realistic option available as of February 2024 is the purchase of enterprise licenses of ChatGPT, or another equivalent commercial product. The main challenge is to explore and clarify both the legal and data protection aspects. In the meantime we observe that main vendors such as OpenAI are currently going through the required steps to adapt to the EU market requirements (creating a subsidiary in the EU – Ireland – and offering their products via companies that already have framework contracts with the EU (including Microsoft). Data protection and information security aspects will be carefully assessed in the next months by the Court’s Data Protection Officer (DPO) and Information Security Officer (ISO) as part of the regular acquisition process, also considering the European Commission DPO decisions.

60 OpenAI/ChatGPT can be acquired via an enterprise license of their product, or included as an optional component of Microsoft Office 365, called Copilot. DIWI has started evaluating both options.

61 If the ECA acquires the necessary licenses, after the DPO and ISO give the green light on using the tool for ECA-internal documents, the DATA team will create specific training events to present to audit and non-audit colleagues how to best use the new tools.

62 We are also looking into the new services offered by the European Commission, for example the new [eBriefing](#) AI drafting assistant specifically developed for the production of EU documents.

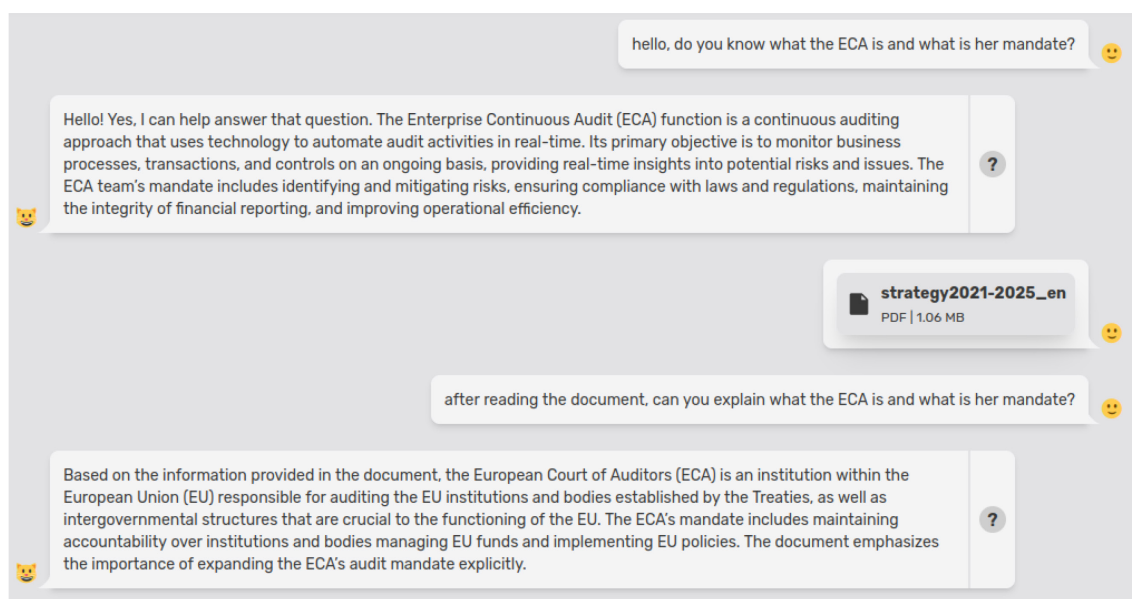
Action 4.2: Customise and deploy a tool to retrieve information from large document sets, using an interactive approach based on Questions & Answers

63 As explained in the reflection paper, audit teams often need to search information in large sets of domain-specific documents they collect. These documents are not always public, and as a consequence it would not be recommended to use commercial AIs such as ChatGPT to ensure data confidentiality and the full respect of EUDPR.

64 Tools that allow auditors to upload their documents of choice and to ask questions about information contained into them are called “Retrieval Augmented Generation”. They are made by a locally run AI that is able to search for relevant information inside locally stored documents (in PDF, Word, etc.). Users interact with the AI (as a chatbot) to explore the content of the documents.

65 As an example, see the figure below. Initially the AI does not know what the ECA is (note the typical phenomenon of “hallucination”, when generative AI tries its best to provide any answer, even if it is incorrect). However, after uploading the public document containing the ECA Strategy 2021-2025, the AI can answer the same question *almost* correctly.

Figure 5 – Example of Q&A on documents



Source: ECA, prototype by the DATA team.

66 As a potential solution, DIWI and an auditor from Chamber IV have started evaluating an open source tool called “[PrivateGPT](#)”, in the context of the “ECA GPT” short innovation project. Other promising open source tools offering similar functionalities can also be considered (example: [CheshireCat AI](#)).

67 Preliminary tests by the authors show that this approach works well with documents containing *factual and concise* information, but the AI tends to be less effective in case of high-level abstract reasoning, such as making professional judgement or when dealing with implicit or subjective information.

68 The key success factors for this action, and for the positive outcome of the project, are largely dependent on the precise customisation needed for the open-source project we choose. We need to personalise carefully the way we store documents and how we present them to the AI to maximise what the tool can learn and understand. Simply installing and using an open-source project in its base state will most probably not provide the quality needed for audit tasks.

Action 4.3: Extend the DORA tool with an AI backend for summarising and other natural language processing tasks

69 DIWI has developed a document reading assistant called [DORA](#). Users can securely upload audit document sets to DORA, search for subjects, terms etc. and build “reading lists” including only the pages that are relevant to the audit objective. At this stage the tool does not use AI.

70 We propose to enhance DORA by adding an AI-based automatic summarisation feature, as well as context-based search.

Action 4.4: Semantic search engine on ECA own production

71 As part of the Development Plan 2021-2025 actions, the DATA team created a working prototype of a “Smart Search” tool, to enable full-text semantic search on special reports produced by the ECA (potentially extendable to all types of reports). The main aim of the activity was to explore the power of semantic search as a way to improve currently used keyword-based methods.

Figure 6 – How users could interact with the Smart Search prototype



Source: ECA, prototype by the DATA team.

72 This prototype uses two technologies to show users the best possible results. The first is a relatively traditional multi-stage ranking, similar to how web search engines work. The second uses an AI/LLM to create a narrative that uses the information retrieved to provide a more complex answer than just pointing to specific sections of a document.


73 A short innovation project has been created, with the goal of making the prototype available to interested auditors in the innovation room that DIWI is setting-up. If the SIP is successful, it may evolve into a full IT project to create a tool accessible to all ECA staff and integrated into the ECA intranet or made available through Assyst.

Action 4.5: Training a chatbot with internal knowledge of audit methodology from AWARE and other controlled-quality sources

74 A similar tool as the one described in action 4.2 can be used to implement an audit assistant chatbot. The difference between the more generic tools described in other actions is that we will prepare carefully the learning material we introduce to the AI/LLM in order to minimise the risks of hallucinations and incorrect results. We need to choose and prepare textual material from our methodology and guidance, and then instruct the AI to only use this material to provide answers.

75 In the example below, taken from a proof of concept⁴ developed by the authors, we can imagine a newcomer trying to understand what “TP” means. As an interesting feature, the chatbot also provides a link to the internet page from which it has taken the information, for additional information and to help the user to validate the answer if necessary.

Figure 7 – EcAI – the Audit Assistant



Source: ECA, proof-of-concept by the DATA team.

Action 4.6: Assess the EC DORIS tool for survey analysis

76 DORIS is a tool initially developed by the European Commission to analyse public consultations, where users can group answers, perform sentiment analysis, extract the most important keywords from textual replies, etc. DORIS is well integrated with EU Survey and it could be useful to handle large and multi-language replies to audit surveys.

77 DIWI is in contact with the EC to assess if the ECA could use the tool. Nevertheless, we should mention that currently the EC only offers access to the DORIS service if they receive the data that should be analysed.

⁴ A proof of concept is a feasibility study in the project discovery phase before the development of a full-fledged product. It's a small, internal, stand-alone project aimed at validating that a core feature or tech assumption can, in fact, be implemented and will function as envisioned. It is also used to validate business hypotheses and to request feedback from potential users.

Action 4.7: Add AI components to robotic process automation (add intelligence to automation of repetitive tasks)

78 The ECA has adopted an enterprise tool to create robotic process automations (RPA). The DATA team, in cooperation with DIWI, has already developed several robots that automate large-scale but simple and repetitive tasks, such as invoice downloads from the Commission's accounting system.

79 Robots can potentially be extended with AI functionalities to further automate audit programmes (see the invoice examination described in DEC 113/23). Specifically, AI has the potential to improve the efficiency of tasks that until now could only be done through human intervention, due to the extreme unpredictability of the input documents.

80 There are two options to use AI in RPA robots:

- evaluate the AI extensions proposed by the vendor, or
- after the ECA purchases a commercial license of an AI tool – or when we can run internally a sufficiently large and powerful LLM – we can include AI functionalities in robots that add query steps to the LLM and report on the results. Example: after downloading unstructured invoices, the robot can query the AI asking specific information such as “extract the due date”.

Action 4.8: Add access to AI to existing data science and visualisation tools

81 Advanced users (both within the DATA team and certain audit staff) prefer to use scripts and code to manipulate data and extract audit findings in a formally documented way. After we decide which AI products we want to buy and/or deploy internally, we need to make them available not only as end-user tools, but also as a building block that can be integrated into code written in R and Python languages for audit purposes.

Action 5: creation of an AI competence centre at the ECA and related governance

82 To achieve our ambitious AI goals we need experts in several fields to work together, to optimise the use of the resources we have available. The ECA already has several professionals working individually (or in small teams) on various aspects of AI, but they are scattered over several directorates. We need to create the administrative and

practical conditions for them to maximise their knowledge and for our institution to make good use of it.

83 A specific competence centre on artificial intelligence to coordinate ongoing and planned actions is being set-up by DIWI at the initiative of the Secretary General. This includes work already delivered and planned by the DATA team in its various areas of expertise. The AI competence centre (AI.CC) will include professionals across the organisation representing different jobs and related needs and expertise (audit, IT, legal, translation, communication, administration etc.).

84 The coordination and the governance of the AI Competence Centre will be ensured by the principal manager in charge of the innovation portfolio in DIWI, a representative from the Slovenian private office and a representative from the DQC/DATA team. Formal issues and decisions related to budget and resource allocation will be presented to and decided in the appropriate existing ECA bodies and committees, with no need to create additional structures.

Action 6: foster interinstitutional and international cooperation

85 Interinstitutional cooperation is essential to optimise the value added by resources that are both scarce and expensive, especially in areas with high cost of living such as Luxembourg. Traditionally, the main goals of participating in interinstitutional working groups have been to exchange ideas and resources, and to take advantage of the maturity level of peer institutions. We believe that we should increase the ambition and create frameworks that enable concrete interinstitutional projects with the objective of jointly producing technical deliverables, blueprints and tools.

Table 5 – Expected benefits of interinstitutional cooperation

Activity	Value
Facilitating idea exchange and regular peer review	Institutions can avoid the pitfalls of working in silos. Regular peer review ensures that AI strategies and their implementations can benefit from diverse perspectives and experience.
Identifying methodological and technical cooperation areas	Institutions can pool their expertise, resources and efforts to tackle complex challenges more effectively. This approach not only accelerates progress but also ensures that solutions are robust, tested and adaptable to varying contexts.

Activity	Value
Re-use of deliverables	Deliverables are strategies, roadmaps, project documentation, business and technical analysis documents, user feedback, code, and architectural designs. By sharing these resources, institutions can significantly reduce duplication of efforts and costs, ensuring that investments are channelled towards innovation rather than reinvention.
Sharing of training sessions	Ensures that valuable insights and learning are not confined into a single institution but are made accessible to a broader audience, encouraging a community that is well-equipped to contribute to the AI landscape.
Resource sharing	The synergy between innovation labs is crucial for the effective development and implementation of AI initiatives. By sharing resources, including personnel, knowledge, and infrastructure, institutions can optimise investments and lower the costs of hiring.
Advocate broader reusability of tools developed by the EC	By convincing the EC to consider re-usability from the outset, institutions can ensure that AI products and solutions developed with EU funding are designed for maximum impact, accessibility, and longevity. The most effective way to do so is to prepare detailed business requirements and define carefully audit expectations from AI.

86 A detailed table of major interinstitutional working groups where the ECA participates as a member or as an observer is presented in *Annex II “Interinstitutional and international working groups”* (see *Table 8 – Main working groups dealing with AI*). Several other activities are under planning for 2024, specifically:

- co-organisation of an AI summer school in Luxembourg, in the context of the “Attractiveness of Luxembourg Working Group”;
- relaunch of the TINA (Technology and Innovation for Audit) network, including some focus on AI; created in 2019 by DIWI/ECALab, we already hosted around 20 online events with an average attendance of 30 professionals from EU SAIs and 40 ECA colleagues; and
- The Directorate of the Presidency (DOP) is evaluating with other SAIs the possibility to use AI for the production of text summaries in the context of the Contact Committee.

87 The DQC/DATA team maintains several other regular bi-lateral cooperations with the US GA Government Accountability Office (GAO), the OECD (which recently published a [document on artificial intelligence](#)) and with the internal audit service of the European Commission. Preliminary contacts have also been made for a possible cooperation with the SAI of Brazil.

88 The ECA should monitor the development of AI-related standards or other practical guidelines issued by professional bodies such as the International Auditing and Assurance Standards Board (IAASB) and the Institute of Internal Auditors (IIA). We should also remain alert to the use of AI by private sector auditors, and interact with them when relevant, practical and useful.

Closing remarks

89 This document provides an analysis of the key issues and consolidates the most relevant information to support decisions required for current and future AI-related activities at the ECA. Overall, taking account of the progress already made, the ECA is well positioned to harness the potential of AI. It is important to establish clear goals and objectives and to continue investing to foster a coherent and safe adoption of this innovative technology while maximising its value for the ECA. An interim progress report will be presented to the AQCC before the end of 2024. It will analyse the progress and results of the actions and provide an overview of the resources used and planned for 2025. A final report will be presented at the end of 2025.

Annex I – Overview of the main regulatory frameworks

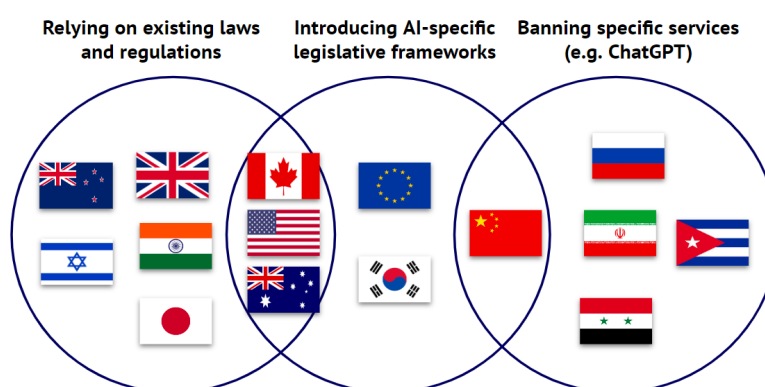
Main regulatory approaches worldwide

90 Given the magnitude of the potential impact of AI, establishing a sound regulatory framework is not just about ensuring AI is ethical and safe, and is used as such, but also about gaining a strategic advantage on the global stage. Countries or international organisations that navigate effectively the delicate balance between fostering innovation and implementing prudent regulation can attract top talent and secure investments. Countries around the world have adopted different approaches. See [Figure 8](#).

Figure 8 – Regulatory approaches around the world

Have we reached “peak” regulatory divergence?

► After years of speculation about mounting potential divergence in regulatory approaches, we're starting to see regulatory approaches stabilise and settle into a handful of distinct approaches.



stateof.ai 2023

Source: <http://stateof.ai>. © Air Street Capital – Creative Commons Attribution 4.0 International License.

91 Some countries chose to rely on existing laws and regulations. For example, in the United Kingdom the [current regulatory approach](#) could be defined as light touch and principles based. It consists of defining the key principles underpinning the framework only. The government decided not to put these principles on a statutory footing (at least initially). The implementation of the principles is entrusted to the different regulators. The government will then monitor whether this non-statutory framework is sufficient and is producing the desired effect.

92 Similar to the UK, the United States are mainly relying on the existing legal frameworks because the attempts to approve a comprehensive AI legislation in Congress are still in their infancy. However, to address the most immediate concerns, the US government has issued an [AI specific executive order](#) on 30.10.2023. The executive order directs several actions targeting the development of AI systems. These include:

- the definition of new AI safety standards and a requirement to share safety test results with the government;
- development of specific security standards and practices;
- initiatives to strengthen privacy protection;
- measures to avoid discrimination or breach of civil rights;
- principles to protect consumers and workers;
- measure to foster competitiveness and US leadership; and
- guidelines for responsible government use of AI.

93 The EU and China stand out as having opted for specific and comprehensive AI legislation. For the EU, the Commission proposed [the EU AI Act](#)⁵ in April 2021. The EU AI Act focuses on the application and use of AI systems and it adopts a risk-based approach, with four levels of risk: “minimal”, “limited”, “high” and “unacceptable”. The risk classification depends on the function performed by the AI system and on the specific purpose and modalities for which the system is used. [Table 6](#) provides an overview of the risk classification and the related requirements.

Table 6 – Overview of the risk-based approach in the EU AI Act

Risk level	General definition	Action
Unacceptable	Harmful uses of AI that are against the EU values (e.g. social scoring by governments).	Prohibited
High	Uses with potential adverse impact on citizen’s safety or fundamental rights (e.g. AI systems used as safety components of a product).	Authorised in the EU market subject to a set of requirements and obligations (including a conformity assessment)
Limited	AI systems that (i) interact with humans (e.g. chatbots), (ii) are used to detect emotions or determine association with (social) categories based on biometric data, or (iii) generate or manipulate content.	Authorised subject to a limited set of transparency obligations.
Minimal	Low risk uses not falling in any of the previous categories.	Authorised without additional obligation

⁵ COM(2021)206.

94 The proposed legal framework will apply to public and private actors inside and outside the EU in respect of AI systems that operate in the internal market or its use affects people located in the EU. As of January 2024, the triologue negotiations reached a political agreement on the proposed legislation.

95 As a transitional measure, the Commission has established the [AI Pact](#), a voluntary scheme that fosters early implementation by EU companies of the measures foreseen in the AI Act. Moreover, the Commission decided to establish a new [European AI Office](#)⁶, as part of the administrative structure of DG CNECT. The new office, among other tasks, will provide guidance, develop tools for evaluating AI models, monitor the implementation of the rules and investigating infringements.

96 China has been one of the most active regulators, with the Cyberspace Administration of China playing a key role. Since 2022, three new regulations have entered into force while a fourth one is being discussed. [Table 7](#) lists these regulations with their key elements. Among other requirements, publicly available AI models must undergo a security/safety assessment, providers must ensure to use only lawful training data and each model must be filed in China’s “Algorithm Registry”. In addition, AI generated content must be labelled. Access to OpenAI’s ChatGPT is restricted in China. However, 11 Chinese AI models have obtained clearance for release to the public. The 2017 “[New Generation AI Development Plan](#)” laid out the goal of achieving global AI leadership by 2030.

Table 7 – Main AI regulations in China

Regulation	Entry into force	Key points
Basic Safety Requirements for Generative Artificial Intelligence Services	N/A (Draft proposed in October 2023)	It puts forward the safety requirements (for training data and AI model) that providers must fulfil before submitting a filing application for the online launch to the overseeing body.
Interim Measures for the Management of Generative Artificial Intelligence Services	15.8.2023	Requirements for publicly available models include the obligation to: file AI models in China’s “Algorithm Registry”; have a security assessment before making the model available to the public; label AI generate content; and screen for illegal content.

⁶ C(2024) 390 final.

Regulation	Entry into force	Key points
Provisions on the Administration of Deep Synthesis Internet Information Services	10.1.2023	Requirements include filing AI models in China's "Algorithm Registry"; prevent the dissemination of illegal content and ensure strict privacy measures.
Provisions on the Management of Algorithmic Recommendations in Internet Information Services	1.3.2022	Requirements include the obligation to: file AI models in China's "Algorithm Registry"; notify users about recommendation algorithms; and give them an option to opt out.

Regulating AI safety and security

97 One of the most crucial challenges for regulators is finding an effective way of ensuring the overall safety and security of increasingly powerful AI models. Some regulators lean towards strict government control via licensing, and/or via imposing mandatory compliance tests to all companies that develop AI tools. Critics of this approach consider that this would stifle innovation and concentrate too much power in a few big (US-based) companies.

98 Other regulators consider that the best way to avoid a negative impact of AI use is to have fully transparent and open models, to incentivise constant scrutiny from research groups and to avoid monopolies. Critics of this approach highlight the risks that "homemade" and publicly available AI models could potentially be used for a broad set of malicious activities, such as to support crimes and create disinformation on a massive scale.

Impact for the ECA

99 In addition to the existing legal requirement (e.g. EUDPR), the ECA will have to ensure compliance with the upcoming EU AI act. From a preliminary analysis of the draft text, the use cases currently envisaged at the ECA (paragraph **39**) are likely to fall under the minimal risk category and therefore will not entail additional requirements or obligations. However, a formal assessment should be performed once the final text of the regulation is known.

Annex II – What are other institutions doing on AI?

100 In September 2023, a maturity assessment performed by Gartner⁷ revealed that 80 % of their clients were investing in AI or generative AI in some capacity. For what concerns generative AI specifically, 10 % of organisations already fully included AI into their daily processes, while the others were either exploring and learning or piloting and experimenting (see [Figure 9](#)). All organisations interviewed started internal reflections on AI-related risks and mitigation strategies.

Figure 9 – Overall maturity assessment survey on Generative AI

September 2023 Update: Generative AI



Source: Webinar by Gartner, [Generative AI Realities: Measuring and Quantifying Business Results](#), January 2024. © 2023 Gartner, Inc. All rights reserved.

101 In the rapidly evolving context of AI, it is difficult to provide a comprehensive overview of what other institutions are doing, as any list risks becoming quickly outdated. In this chapter we provide, to the best of our knowledge, the major ongoing activities in interinstitutional working groups, EU institutions, SAIs, and the OECD as of January 2024.

Interinstitutional and international working groups

102 There are several working groups that foster cooperation and knowledge sharing on AI both at EU and international level. [Table 8 – Main working groups dealing with AI](#) lists the most relevant for the ECA.

⁷ Gartner is a technological research and advisory firm working for large corporations, government agencies, technology companies, and investment firms.

Table 8 – Main working groups dealing with AI

Working group	Purpose and scope
AI@EC network	<p>This community was created, hosted and coordinated by the Commission but it includes hundreds of participants from all EU institutions, bodies and agencies. The forum is based on MS Teams, and contains discussions and experience exchange on a broad range of AI technical and non-technical subjects.</p>
ICDT	<p>The Interinstitutional Committee for Digital Transformation is the top-level coordination body of the IT and innovation departments of EU institutions. The Court of Justice currently holds the rotation chair and the ECA is represented by DIWI. The ICDT subgroup “Emerging technologies” focuses also on artificial intelligence. The key activities of the ET subgroup for 2024 are:</p> <ul style="list-style-type: none"> ○ sharing risk assessments of AI tools, including IT, Legal, Ethical and data protection aspects; ○ scouting for emerging technologies and continuous market survey; ○ evaluating Microsoft co-pilot constraints and opportunities; ○ piloting generative AI on-premises; ○ testing scenarios and frameworks for AI models; and ○ organising an Innovation Day to showcase concrete results.
EUROSAI ITWG	<p>The EUROSAI IT working group aims to stimulate European SAIs to jointly explore the strategic consequences of IT-related developments, both with regard to their audit responsibilities (IT audit) and to their own use of IT. AI is of course an important theme.</p>
INTOSAI-WGISTA	<p>The “Working Group on Impact of Science and Technology on Auditing” focuses on key developments in areas such as AI, blockchain, cybersecurity, data analytics, 5G, and quantum computing, among others, that will have an increasing impact on governments and their auditors. One of the main objectives is to assess and share best practices in developing and maintaining expertise within SAIs and applying science and technology in their auditing.</p>
INTOSAI-SCEI	<p>The Supervisory Committee on Emerging Issues (SCEI) provides recommendations on the important issues and emerging challenges faced by INTOSAI and individual SAIs by helping coordinate and support the sharing of knowledge in this regard.</p> <p>The SCEI is responsible for INTOSAI’s Enterprise Risk Management.</p>

Working group	Purpose and scope
INTOSAI-WGBD	The Working Group on Big Data is a specialized working group approved by INTOSAI under Strategic Goal Three: Knowledge Sharing and Services. Its objective is to identify the challenges and opportunities faced by SAIs in the era of big data
OECD Tech & Analytics Community of Practice	The OECD supports our communities of anti-corruption and oversight bodies to advance their digital strategies and strengthen their use of technology, data and analytics (including AI). They also have a dedicated AI policy observatory .

European Commission

103 The Commission has been very active in many areas in the field of AI. We describe the policy and regulatory initiatives in paragraphs [93](#) to [95](#). In addition, at the end of January 2024 the Commission has launched an [AI innovation package to support start-ups and SMEs](#). For what concerns the use of AI at the Commission, a recent [communication \(AI@EC\)](#)⁸ provides the strategic vision for the development and internal use of this technology. It covers the general (risk-based) approach, the policy and governance aspects and the practical use cases. The Commission has several AI-powered services already available (see [Table 9](#)). In addition, there are 10 services currently in development and testing, and a similar number of new use cases being discussed.

Table 9 – AI-powered systems already available at the Commission

Name	Responsible DG	Description
eTranslation and e-Summary	DGT	Language services that provide automated translation and summaries, both to the EU institutions, bodies, and agencies and to other users in the EU. They are part of a broader language tools suite .
eBriefing	DGT	Service using state-of-the-art generative AI to produce topic-based overviews or briefings from a given set of relevant documents.

⁸ C(2024) 380 final.

Name	Responsible DG	Description
Publio	OP	Virtual assistant supporting users in their discovery of EU law and EU publications, thus also contributing to greater accessibility.
SeTA	JRC	Semantic text analyser useful for generating metadata, and for document classification and discovery. The OP has used SeTA as the basis for a prototype of document annotator and classifier (Cellar-SeTA) .
DORIS	CNECT/DIGIT	Provides sentiment analysis, keyword extraction, summarisation, and named-entity recognition to semi-automatically analysis in any type and document. There is also specific dashboard for public consultations answers.

104 The coordination of the AI@EC initiative and the oversight on the AI@EC network is a responsibility of a dedicated AI interservice steering group, in cooperation with a director-level core group consisting of DG DGT, OP, DG HR, DG CNECT, the SG, and the JRC. The group is open to representatives from interested DGs.

Other EU Institutions


105 Other EU institutions, bodies and agencies sometimes rely on the Commission (and its budget) for major policy and technological developments. For what concerns AI, to the best of our knowledge, most of them set-up dedicated internal teams.

EU Supreme Audit Institutions

106 In February 2024, together with the DOP, we launched a survey to the 27 SAIs to gather an overview of how they are approaching AI. We received 23 replies, which provide a representative picture of the current state of play.

107 Sixteen of the surveyed SAIs (70 %) are currently in the exploratory phase, four (18 %) are developing some AI tools and no SAI yet has AI tools in production (see [Figure 10](#)).


Figure 10 – EU SAIs’ self-assessment on the status of AI implementation



Source: ECA survey.

108 Eight SAIs stated that they are making already some use of AI technologies. Most SAIs are planning to use these technologies in the future (See [Figure 11](#)). As an example, one respondent mentioned that they are using generative AI to assist in writing computer code.


Figure 11 – SAIs using or planning to use specific AI technologies



Source: ECA survey.

109 When introducing AI in the work practice, organisations can opt for a top-down approach, i.e. starting with strategic documents and moving down to technical analyses, or a bottom-up approach, i.e. consolidating technical experience and pilots into a formalised roadmap. Based on the replies concerning the status of key documents, we were not able to identify a dominant approach among SAIs (see [Figure 12](#)).

Figure 12 – Status of key documents concerning AI activities



Source: ECA survey.


110 We asked SAIs whether they would be interested in cooperating with the ECA on AI. Three SAIs (17 %) replied that they are not interested due to practical reasons (e.g. lack of resources). The others would be most interested in cooperating with us in:

- the identification and exploration of use cases;
- sharing knowledge, documentation, code about experiments; and
- sharing user feedback on prototypes.

111 Interestingly, most SAIs are interested in sharing the analysis of the audit-related aspects of artificial intelligence, an aspect the ECA started much earlier to formalise.

112 Our last question concerned the challenges that SAIs face in implementing AI in their organisation. Their replies are summarised in [Figure 13](#).

Figure 13 – Most mentioned challenges to AI implementation



Source: ECA survey.

113 During official visit of EU SAIs and participation in working groups we gathered some additional feedback:

- in several cases, their staff started experimenting with AI at personal level, received general training, and expressed the need for guidance on practical applications;
- in some cases SAIs have started investigations on how to use LLMs to support their audit activities, but they do not have enough resources to create a specific AI task force;
- they expressed their interest in receiving ECA documents and training session.

Other non-EU audit Institutions

114 In the context of the various INTOSAI working groups, the ECA has periodic exchanges with non-EU audit institutions. [Table 10](#) lists some relevant information regarding AI activities.

Table 10 – Examples of AI activities in non-EU audit institutions

US GAO	<ul style="list-style-type: none"> ○ The Government Accountability Office is very active in the field of AI; they also communicate activities and progress on their public website. ○ The GAO AI Accountability Framework was finalised in 2021; it identifies key accountability practices, centred around the principles of governance, data, performance, and monitoring, to help US federal agencies and others use AI responsibly. ○ The GAO Innovation Lab, set up in 2019 and currently comprises around 12 experts, is part of the broader “Science, Technology Assessment, and Analytics” team (STAA). As of September 2023, the STAA team has 156 employees⁹. ○ They recently published a very interesting list of current audit scenarios under analysis and prototyping.
TCU and CGU Brazil	<ul style="list-style-type: none"> ○ At the 4th Annual Meeting of the Working Group on Impact of Science and Technology on Auditing (Abu Dhabi, November 2023), the Tribunal de Contas da União presented how they use AI to enhance risk-based sample selections and the use of AI chatbots to analyse documents. ○ At the OECD Tech & Analytics Community of Practice of 17.11.2023, the Comptroller General of the Union presented “Generative AI for Integrity”, showing how they have fine-tuned a Language Model (Llama2-7B) to improve its fluency in Portuguese, and how they use it to extract data from invoices.

⁹ See also: [Testimony before the Committee on House Administration, House of Representative](#).

Annex III – Develop or buy? Cloud-based commercial products versus on-premises tool

Introduction and definitions

115 As mentioned in the reflection document, there are two possible strategies when deciding how to deploy AI-based tools in an organisation. The first is to purchase a license for one or more commercial products, the second is to use internal IT resources to install and maintain all the necessary components of an AI tool (i.e., language models and the associated user tools).

116 To allow for a well-informed decision in this respect, it is important to have a clear and consistent understanding of some of the key concepts. This starts with a short list of important definitions:

- **Cloud-based commercial products:** these are services and tools offered by companies over the internet – the “cloud” – often on a subscription basis. Users access these services remotely without the need to install or maintain the software locally. To use these services, user data and requests must be necessarily sent to the service provider’s infrastructure and the results will be available to them.
- **On-premises tools:** contrary to cloud-based solutions, these tools are installed and run on our corporate IT servers or individual laptops. Such tools may need to be purchased or they can be open source; the latter tools are developed and maintained by a community of contributors and users can freely modify, adapt and use them. In any case, the licensing agreement must be carefully checked to ensure that the product can be used by a public institution such as the ECA. As of January 2024, almost all AI products and tools that can be installed on premise are open source.
- **Data security, privacy and data protection:** in AI deployment, data security refers to protecting data from unauthorized access and corruption. Data protection involves ensuring that personal information is used ethically and in compliance with regulations.
- **Customisation and flexibility:** These terms refer to the ability to tailor AI tools and systems according to specific needs and requirements. Flexibility is particularly important in adapting to changing audit environments and demands, especially in a rapidly evolving technological context.

117 [Table 11](#) compares the two solutions from various angles. Please note that there is some overlap in the information, as we chose to examine the subject from every angle.

Table 11 – Comparison between cloud-based and on-premises solutions

Aspect	Cloud-based solutions	On-premises tools
Contractual	Typically involve ongoing subscription models. Contracts are often standardised with limited scope for negotiation.	Involve upfront investment. Contracts are more customisable, allowing better alignment with institutional requirements.
Data protection	Data is stored outside the organisation, often in multiple locations. This raises concerns about data sovereignty, control over personal data and compliance with EUDPR.	Offers more control over data storage, aligning closely with EUDPR and other data protection regulations, crucial for sensitive government data.
Legal	Influenced by the legal jurisdiction of the host country, which can create complexities, especially in cross-border data transfers.	Legal compliance is more straightforward, governed by EU law.
Data security	Need to continuously monitor for compliance with evolving regulations like EUDPR.	More predictable regulatory compliance, but requires in-house expertise to maintain standards.
Technical	Most of the technical aspects related to deployment and maintenance are outsourced, making costs more predictable but limiting the organisation's control over products evolution.	On-premises AI requires new IT architectural components; internal IT professionals would need to learn new technologies and acquire new skills. The organisation keeps full control of products evolution.
Future regulation on AI safety ¹⁰	Compliance with future regulations is the vendor's responsibility.	The organisation would potentially need to take care of assessing the safety of on-premises solutions.

¹⁰ The regulatory aspects of AI are evolving very fast, and we expect the discussion “open source” vs “commercial” models to be particularly heated (see for example this article on the Financial Times: <https://www.ft.com/content/b62a32c9-a068-40c4-8612-9da8cffb396c>).

Aspect	Cloud-based solutions	On-premises tools
Performance and reliability	Provide scalable performance and high reliability with less direct control, offering cost efficiency and ease of management through shared infrastructure and resources.	Performance and reliability depend on the magnitude of our upfront investments. It is more difficult to scale our IT infrastructure up or down.

Available options, as of February 2024

118 As of February 2024, an easily available commercial cloud AI service is ChatGPT from OpenAI. It is available to purchase in two ways:

- o via an enterprise subscription that would enable ECA staff to use ChatGPT within their web browser; or
- o via a machine-to-machine API subscription (which allows ECA tools and applications to include AI services in a transparent way for the user).

119 Although OpenAI is a US company, the API subscription can be purchased using an EU framework contract with Microsoft. However, the DPO and the legal service should evaluate carefully whether the contractual guarantees provided by Microsoft sufficiently cover our data protection and legal obligations (e.g. regulation 2018/1725, GDPR, AI Pact/ACT, etc.).

120 As for solutions that can be locally integrated into the current ECA IT infrastructure, it is important to note that the number of local AI models and tools available is increasing steadily, and it is difficult to predict which one will be the dominant choice in six to 12 months from now. As of January 2024, the two most promising candidates are Llama-v2 (developed by Meta and released with a permissive licensing) and Mistral (developed by a French startup of the same name). Mistral is one of the most promising EU start-ups based in this area and has recently secured a significant round of financing to continue its development.

Conclusions and proposal: start a dual approach in AI deployment (develop and buy)

121 We believe that in this intricate landscape, a one-size-fits-all approach is not effective. By leveraging both commercial cloud services and on-premises (open source) tools, the ECA would be able to ensure a broad coverage of its audit needs. Commercial products provide immediate, scalable solutions for standard tasks, while custom internal AIs and tools can be tailored to unique or complex requirements, especially when strict data confidentiality is a priority.

122 Commercial products, such as OpenAI's ChatGPT, are much more sophisticated and can produce more nuanced and precise results. This aspect makes them a primary choice for applications where a sophisticated mastering of languages is essential, such as drafting assistance, translation, summarising documents, etc. As of February 2024, ChatGPT is also leading in providing detailed reasoning and step-by-step description of the process used to produce a given output. This is a key element that helps produce a verifiable audit trail whenever AI is used as part of the audit work, for example for examining financial data, unstructured documents, or large datasets.

123 On-premises open-source products generally produce less sophisticated results, mostly due to their limited size and capacity. However, they are probably sufficient to create AI assistants that can access ECA data and methodologies, organising internal data and allowing sophisticated search capabilities. Since the ECA will have full control of the IT environment used for such products, data confidentiality can be strictly controlled.

124 Diversifying AI tools mitigates the risk associated with reliance on a single type of solution, or vendor. It addresses the potential vulnerabilities of cloud-based systems and the resource-intensive nature of maintaining on-premises tools.

125 While commercial AI products involve subscription or licensing fees, they reduce the need for extensive in-house development. Conversely, on-premises deployment can incur higher initial costs but can be more cost-effective in the long term for specialized, frequently used applications.

126 In conclusion, we propose a balanced strategy that combines the acquisition of commercial cloud-based products with the development of internal tools based on open-source language models. The dual approach is designed to accommodate the diverse range of audit scenarios we are faced with (see paragraph 72 of our reflection paper).

Annex IV – Risk mapping and analysis

Table 12 – General AI risks

Risk	Description and impact	Likelihood / Severity	Mitigation measures
Limited transparency and oversight of AI's logic	Many AI systems are “black boxes” due to their complex and opaque decision-making processes. It may be difficult to understand and explain the process used by the system to produce a given output. This limits the transparency of the system and how to assess the degree of reliability of outputs in the context of audit evidence.	High/High	<ul style="list-style-type: none"> Choose AI models that are inherently more open, transparent and interpretable. Ensure that the model provides a step-by-step description of the process and the software code used to reach the results (to allow for re-performance). Provide auditors with training and methodological guidance for understanding the outputs and for documenting how they were used in the audit work.
Accountability gaps	In case of failure in a process where decisions were made by AI systems, accountable and responsible entities may not be clearly identifiable as legal and ethical frameworks lag behind.	Medium/High	<ul style="list-style-type: none"> Do not delegate any meaningful decisions to AI systems. Deploy AI systems to provide assistance to human staff, who retain oversight and responsibility.

Risk	Description and impact	Likelihood / Severity	Mitigation measures
Biased and discriminatory outputs	AI systems can inherit biases present in their training data, leading to biased and/or discriminatory outputs. This may limit the reliability of their output.	High/High	<ul style="list-style-type: none"> ○ Provide staff with training and methodological guidance on these risks. ○ Establish clear rules and boundaries on the acceptable use of outputs. ○ Regularly re-evaluate models and favour those with less biased training sets. ○ When applicable, use models with minimal prior knowledge that can instead be fed with ECA's knowledge base (AWARE, Assyst, previous reports, etc.).
Privacy/data protection issues	AI systems use and produce large amounts of data, which can include personal data (including sensitive information). This information risks not being processed according to existing regulations or may not be sufficiently protected. It might also be inaccurate.	High/High	<ul style="list-style-type: none"> ○ For on-premises AI models, ensure that the model and the enabling infrastructure and governance are compliant with the applicable data protection regulations. ○ For cloud-base commercial services, ensure that the vendor is GDPR compliant and provides adequate contractual guarantees. ○ If possible, prefer EU models and providers. ○ For work-related purposes, allow the use of commercial services through corporate licenses (with stronger contractual guarantees) only. ○ Apply experience and best practices acquired with other IT systems.

Risk	Description and impact	Likelihood / Severity	Mitigation measures
Security issues	AI systems can be vulnerable attacks, including data poisoning and model hacking. This may lead to malicious use, manipulation or data leaks, thereby causing harm or reputational damage.	High/High	<ul style="list-style-type: none"> ○ Perform a specific security assessment before deployment of AI models. ○ Consider additional mitigation measures in the IT security plan for threats that are specific to AI models. ○ For work-related purposes, allow only the use of commercial services through corporate licenses (with stronger contractual guarantees).
Dependency and loss of autonomy	Over-reliance on AI may lead to a loss of human expertise and autonomy in decision-making.	Medium/Medium	<ul style="list-style-type: none"> ○ Limit AI scope to specific applications where the technology assists staff and improves efficiency. ○ Invest in training to maintain staff expertise.
Environmental impact	AI systems require intensive computation, thus leading to high energy consumption and potential environmental impact.	High/Low	<ul style="list-style-type: none"> ○ None; however, vendors of the best performing commercial services have an incentive to develop more efficient models to be able to scale their business.

Table 13 – Risks related to the deployment of AI at the ECA

Risk	Description and impact	Likelihood / Severity	Mitigation measures
Cost of acquisition and operation is too high	As the ECA has a limited IT budget, the cost of acquiring and operating the most advanced AI systems may turn out to be excessive.	High/High	<ul style="list-style-type: none"> ○ Explore possibilities for interinstitutional cooperation with the aim of sharing infrastructure/licensing costs. ○ Despite not being the most advanced in the industry, some cheaper or less resource-intensive models may still be good enough for our use cases.
Compliance with regulations hinders adoption	Compliance with the upcoming AI Act and with other applicable regulations (including data protection) may be complex and slow down the adoption. In addition, fast changes in the regulatory environment may disrupt the deployment.	Medium/Medium	<ul style="list-style-type: none"> ○ Explore different solutions in parallel and choose the most suitable/less risky for each use case. ○ Prefer open source models and EU providers when possible. ○ Monitor the evolution of legislation and regularly review and update internal policies.

Risk	Description and impact	Likelihood / Severity	Mitigation measures
Excessive dependency on only one AI provider	The competitive advantage of a provider may allow them to have a dominant market share or establish a semi-monopoly, thus leading to excessive dependency on only one vendor for our organisation.	Medium/Medium	<ul style="list-style-type: none"> ○ Even when using mainly a commercial provider, consider also deploying an alternative open source model to reduce dependency. ○ If performance is comparable, prefer open source models that can be operated locally or easily run on cloud infrastructures from different vendors. ○ Promising EU providers are starting to secure significant fundings (e.g. the French “Mistral” AI) and are entering the market.
Insufficient in-house skills to deploy production tools	There may not be sufficient in-house skills to convert initial pilots or proofs of concept into products available to all staff, thus slowing down adoption or increasing costs.	Medium/High	<ul style="list-style-type: none"> ○ This risk is not relevant for commercial product whose licences can be purchased. ○ An AI competence centre will be built at the ECA. ○ Interinstitutional cooperation, especially with the Commission, may help fill the competence gap.
Fast obsolescence of AI models	Development in the AI domain has been extremely fast over the past two years. As preparing, testing and running a pilot takes time, there is a risk of deploying in production an already outdated model.	High/Medium	<ul style="list-style-type: none"> ○ Keep a sufficient degree of flexibility in the local infrastructure and the related processes. ○ Despite not being the top performer, the model deployed may still be good enough for our use cases. ○ This risk has a lesser impact for commercial products where the upgrade to a superior model may consist in the purchase of new licenses.

Risk	Description and impact	Likelihood / Severity	Mitigation measures
Reputational damage linked to the use of AI	Errors/issues due to misuse of AI in our work practice may lead to reputational damage for the ECA.	Medium/High	<ul style="list-style-type: none"> ○ Provide staff with training and clear guidance on the acceptable and correct use of AI. ○ Consider specific checks and training for the staff.
Negative perception and pushback among staff	The adoption of AI can be perceived in a negative way by staff. It may be considered an attempt to automate human tasks and therefore threatening for the workforce. This may lead to a pushback and lack of adoption.	Low/Medium	<ul style="list-style-type: none"> ○ Appropriate communication at all levels (clarifying the role of AI as mere assistant to human staff). ○ Training and awareness sessions.
Misuse of AI in internal processes	Inappropriate use of AI in internal processes may lead to illegal or harmful practices.	Low/High	<ul style="list-style-type: none"> ○ The EU AI Act outlines harmful practices that will be explicitly prohibited. ○ Existing rules and safeguards apply as well to new technologies.

COPYRIGHT

© European Union, 2024

The reuse policy of the European Court of Auditors (ECA) is set out in [ECA Decision No 6-2019](#) on the open data policy and the reuse of documents.

Unless otherwise indicated (e.g. in individual copyright notices), ECA content owned by the EU is licensed under the [Creative Commons Attribution 4.0 International \(CC BY 4.0\) licence](#). As a general rule, therefore, reuse is authorised provided appropriate credit is given and any changes are indicated. Those reusing ECA content must not distort the original meaning or message. The ECA shall not be liable for any consequences of reuse.

Additional permission must be obtained if specific content depicts identifiable private individuals, e.g. in pictures of ECA staff, or includes third-party works.

Where such permission is obtained, it shall cancel and replace the above-mentioned general permission and shall clearly state any restrictions on use.

To use or reproduce content that is not owned by the EU, it may be necessary to seek permission directly from the copyright holders.

Cover page – © kunakorn/stock.adobe.com.

Figures 3 and 4 – Icons : These figures have been designed using resources from [Flaticon.com](#). © Freepik Company S.L. All rights reserved.

Figure 8 – <http://stateof.ai>. © Air Street Capital – Creative Commons Attribution 4.0 International License.

Figure 9 – © 2023 Gartner, Inc. All rights reserved.

Software or documents covered by industrial property rights, such as patents, trademarks, registered designs, logos and names, are excluded from the ECA's reuse policy.

The European Union's family of institutional websites, within the europa.eu domain, provides links to third-party sites. Since the ECA has no control over these, you are encouraged to review their privacy and copyright policies.

Use of the ECA logo

The ECA logo must not be used without the ECA's prior consent.

EN PDF	PDF/Volume_01	QJ-09-24-570-EN-N	978-92-849-2748-7	10.2865/317443
--------	---------------	-------------------	-------------------	----------------