THE RISE OF THE CHIEF AI OFFICER

Mandate, Operating Model and First 100 Days

29 MARCH 2024 • ENTERPRISE AI SERIES



Paul Forrest

PAPER SIX · ENTERPRISE AI, 2023 TO 2024

About this paper

This is the sixth paper in a nine-part series on the enterprise adoption of generative artificial intelligence.

It is published one day after a United States budget office memorandum required federal agencies to designate Chief AI Officers, which makes the role concrete in one large class of organisations and raises the question of what it means everywhere else.

The memorandum and the executive order behind it are quoted from the documents themselves. The operating model, the decision rights and the argument in the final section are mine, and the tags mark the boundary.

HOW TO CITE

Forrest, P. (2024). *The Rise of the Chief AI Officer: Mandate, Operating Model and First 100 Days.* Enterprise AI, 2023 to 2024, Paper Six. Version 1.0, 29 March 2024.

HOW TO READ THE EVIDENCE TAGS

Every substantive claim carries a tag stating what kind of evidence stands behind it. A claim resting on two kinds carries both.

PRIMARY An institutional or official document, cited from the issuing body.

PREPRINT Research posted as a preprint or working paper, not peer-reviewed at this date.

LEGAL A statute, regulation, executive order, legislative proposal or court decision.

AUTHOR My own framework or judgement, proposed and not yet proven.

A NOTE ON DATES

Research for this paper closed on 28 March 2024, the day the memorandum it takes as its starting point was issued. One absence is deliberate. I could find no defensible survey of how many private-sector organisations have appointed a Chief AI Officer, so no such figure appears here.

LEGAL AND REGULATORY NOTICE

Sections one and two describe a government memorandum, an executive order and the progress of European legislation. None of it is legal advice, which sits outside the terms of any engagement with me, and a question about a particular organisation's obligations belongs with its own qualified advisers.

Legal and administrative status is stated precisely. A memorandum binding federal agencies imposes nothing on a private company, and a legislative text adopted at first reading by one institution is not law. What follows depends on both distinctions.

Contents

FRONT MATTER

Abstract and summary of the argument 4

THE PAPER

1 Why the role emerged, and what happened yesterday 5

2 Where the European position stands, precisely 6

3 What the role is for, and where it stops 7

4 A federated model with a small centre 9

5 Decision rights, written down 11

6 The first 100 days 12

7 Most organisations should not create this post 15

END MATTER

Method, evidence and limits 16

Sources considered and set aside 17

About the author 18

References 19

Further sources consulted 20

ABSTRACT AND SUMMARY OF THE ARGUMENT

A government made the role mandatory yesterday, which is not the same as making it necessary

On 28 March 2024 the United States budget office issued a memorandum requiring every federal agency to designate a Chief AI Officer within 60 days, and setting out the seniority, positioning and responsibilities the role carries.[1] LEGAL It is the most detailed public specification of the job that exists, and it is worth reading closely by organisations to which it does not apply.

One correction belongs at the front. The memorandum implements a requirement created five months earlier by executive order, and it says so in its own text.[1,2] LEGAL The temptation will be to present 28 March as the moment the role was invented, and that gets the chronology wrong. The chronology matters because the five months between the two documents contain the public consultation in which the design was contested.

This paper distinguishes the role from the chief information, data, technology and risk officers, proposes a federated operating model with a small centre, sets out decision rights including the one that is usually unassigned, and gives a first 100 days. It also argues that most organisations should not create the post, which is an unfashionable position and the one the evidence supports. AUTHOR

WHAT THIS PAPER ARGUES

The accountability gap is real and the job title is optional. What has to exist is one named person who owns enterprise coherence, and in most organisations that person already exists under a different name.

- The requirement predates the memorandum. The executive order of 30 October 2023 directed that guidance specify the designation of a Chief AI Officer at each agency, and the March memorandum is that guidance.[1,2] LEGAL
- The specification is unusually precise. It asks for senior executive level at the largest agencies, a primary role concerned with artificial intelligence specifically, and regular engagement with the deputy head of the organisation.[1] LEGAL
- The governance board is chaired by the deputy rather than by the officer. The memorandum makes the Chief AI Officer the vice-chair, which places convening authority above the role rather than inside it, and that is the design choice most worth copying.[1] LEGAL
- It binds no private company. The memorandum is guidance to federal agencies, and citing it as evidence that organisations generally need the post is a category error. AUTHOR
- The model should federate, with a small centre. A central office of two to five people owns coherence while delivery accountability stays in the business, because centralising delivery produces a queue and a reputation for obstruction within two quarters. AUTHOR

SECTION 1

Why the role emerged, and what happened yesterday

A memorandum issued on 28 March makes the post mandatory across federal agencies. It implements a requirement created the previous October.

The United States Office of Management and Budget issued a memorandum on 28 March 2024 on advancing governance, innovation and risk management for agency use of artificial intelligence.[1] LEGAL It requires each agency to designate a Chief AI Officer within 60 days of issuance, and it specifies the role in more detail than any publicly available job description.

The memorandum states that it acts as required by the executive order of 30 October 2023, which directed that this guidance specify the designation of such an officer at each agency.[1,2] LEGAL The role was created in October and specified in March. Anyone who treats yesterday as the origin misses the five months in between, which included a public consultation on exactly these questions, opened at the beginning of November and closed in early December.[3] PRIMARY

What the memorandum specifies

1. Designation within 60 days, with officers identified to the budget office and changes notified.[1] LEGAL
2. A primary role concerned with artificial intelligence specifically, as distinct from data or technology issues in general, at the largest agencies. An existing information, data or technology officer may hold it only where they have significant relevant expertise.[1] LEGAL
3. Seniority at senior executive level or equivalent at the largest agencies, and at a stated grade elsewhere.[1] LEGAL
4. Regular engagement with agency leadership, to include the deputy head of the organisation.[1] LEGAL
5. Responsibilities covering senior advice to the agency head, governance and oversight processes, an annual inventory of use cases, a compliance plan, a strategy, advice on resourcing, determination of which uses are safety or rights affecting, and the granting and revocation of waivers.[1] LEGAL
6. A governance board within 60 days at the largest agencies, chaired by the deputy head with the Chief AI Officer as vice-chair, meeting at least twice a year.[1] LEGAL

The board is chaired by the deputy, and the officer is the vice-chair. That design decision is the most transferable thing in the document, and the cheapest to copy. It places convening authority above the role, which means the officer does not have to secure attendance from peers who outrank them. Most private sector equivalents get it wrong.

> WHAT THIS MEMORANDUM DOES NOT DO
>
> It binds federal agencies. It creates no obligation on any private company, and it is not evidence that organisations in general require the post. LEGAL
>
> It is nonetheless the most detailed public specification of the job available, and reading it as a design document is the useful way to use it.

SECTION 2

Where the European position stands, precisely

Adopted at first reading by one institution on 13 March. Not signed, not published, not in force, and imposing nothing on anybody today.

The language in circulation is loose enough to mislead a board. On 13 March 2024 the European Parliament adopted its position at first reading on the proposed artificial intelligence regulation, by 523 votes to 46 with 49 abstentions.[4] LEGAL That is an important step and it is not the same as the measure becoming law.

The sequence to date runs as follows. The Commission proposed the regulation in April 2021. The Council agreed a general approach in December 2022 and the Parliament its negotiating mandate in June 2023. A provisional political agreement was reached in trilogue on 9 December 2023, member state representatives confirmed the text on 2 February 2024, the responsible committees endorsed it on 13 February, and the Parliament adopted its first-reading position on 13 March.[4] LEGAL

SPECIFICATION

What is true of the European regulation on 29 March 2024

- It has been adopted at first reading by the European Parliament.
- It remains subject to finalisation by lawyer-linguists and to formal adoption by the Council of the European Union.
- It has not been signed, has not been published in the Official Journal, and has not entered into force.
- No obligation under it applies to any organisation today, and no application date can be stated with certainty, because the clock runs from a publication that has not happened.

One consequence for anybody drafting this quarter is that the regulation has no official number yet, since numbers are assigned on publication. A board paper citing one cites something that does not exist. Describing the measure as passed into law, signed, ratified or in effect is wrong in each case, and the errors are not interchangeable. LEGAL

One institutional development did complete inside the window. The European AI Office was established by Commission decision on 24 January 2024, taking effect the following month.[5] LEGAL An administrative body now exists, which is a reasonable signal about direction and pace whatever the legislative timetable turns out to be.

SECTION 3

What the role is for, and where it stops

The mandate has seven parts, and the boundary that matters separates owning coherence from owning delivery.

The accountability gap is real and predates the job title. Artificial intelligence work touches strategy, data, technology, risk, legal, human resources and every operating unit, and in most organisations each of those owns a fragment. The result is that nobody can answer how many systems exist, what they cost, what evidence supports them, or who decided.

THE MANDATE, IN SEVEN PARTS



FIGURE 1 What the role owns. The boundary that matters runs between owning coherence and owning delivery, and a mandate that crosses it turns the office into a bottleneck.

My framework, read against the responsibilities set out in the March memorandum. AUTHOR / LEGAL

1. Strategy, which is where the organisation intends capability to matter and where it does not.
2. The portfolio, meaning what is being built, what it is worth, what is duplicated and what should stop.
3. Governance, covering tiering, intake, evaluation standards and the review forum.
4. Value, both the measurement method and the honest reporting of it, failures included.
5. Talent, meaning how capability is distributed across the four populations and where it is thin.
6. The ecosystem of supplier relationships, contractual terms and dependency concentration.

7. Reporting, as a single account of all of the above to the executive and the board.

WHERE THE ROLE ENDS, AND WHO PICKS UP

ADJACENT ROLE	WHAT THEY KEEP	WHAT MOVES, AND WHY
Chief information officer	Infrastructure, integration, service management, and the operational running of deployed systems.	Nothing moves. The artificial intelligence office sets the standard a system must meet before it enters the estate, and does not run the estate.
Chief data officer	Data quality, lineage, stewardship, cataloguing and the governance of data as an asset.	Nothing moves, and the dependency runs one way. An artificial intelligence programme built on ungoverned data will fail for data reasons and be recorded as an AI failure.
Chief technology officer	Engineering standards, build decisions, and technical architecture.	Model selection and evaluation standards move, because they are a risk and coherence question rather than an engineering preference.
Chief risk officer	Risk appetite, the control framework, and acceptance of residual risk above threshold.	Nothing moves. The artificial intelligence office proposes tiering and evidences controls, and does not hold the appetite it is being measured against.
Chief information security officer	Threat model, security controls, incident response and supplier security assurance.	Nothing moves. Prompt injection and model supply chain are added to an existing threat model rather than governed separately.
Business unit leaders	What gets built, what it is worth, and whether their people will use it.	Nothing moves, and this is the boundary most often breached. An office that selects use cases for the business owns outcomes it cannot deliver.

The boundary that gets breached first

Selecting use cases for the business is the most common overreach and the most damaging. It is attractive because the office can see the whole portfolio and the business units cannot, and it fails because the office then owns outcomes it has no means to deliver. Within two quarters the business regards it as a tax and routes around it, which is how these functions acquire their reputation. AUTHOR

SECTION 4

A federated model with a small centre

Two to five people at the centre. Anything larger starts doing the work, and anything doing the work stops governing it.

Three structures are available and two of them are known to fail. A fully centralised team delivering everything becomes a queue. A purely advisory function with no authority produces guidance that nobody follows. The federated arrangement holds a small centre with real decision rights over a defined set of questions, and delivery accountability everywhere else.

A FEDERATED OPERATING MODEL

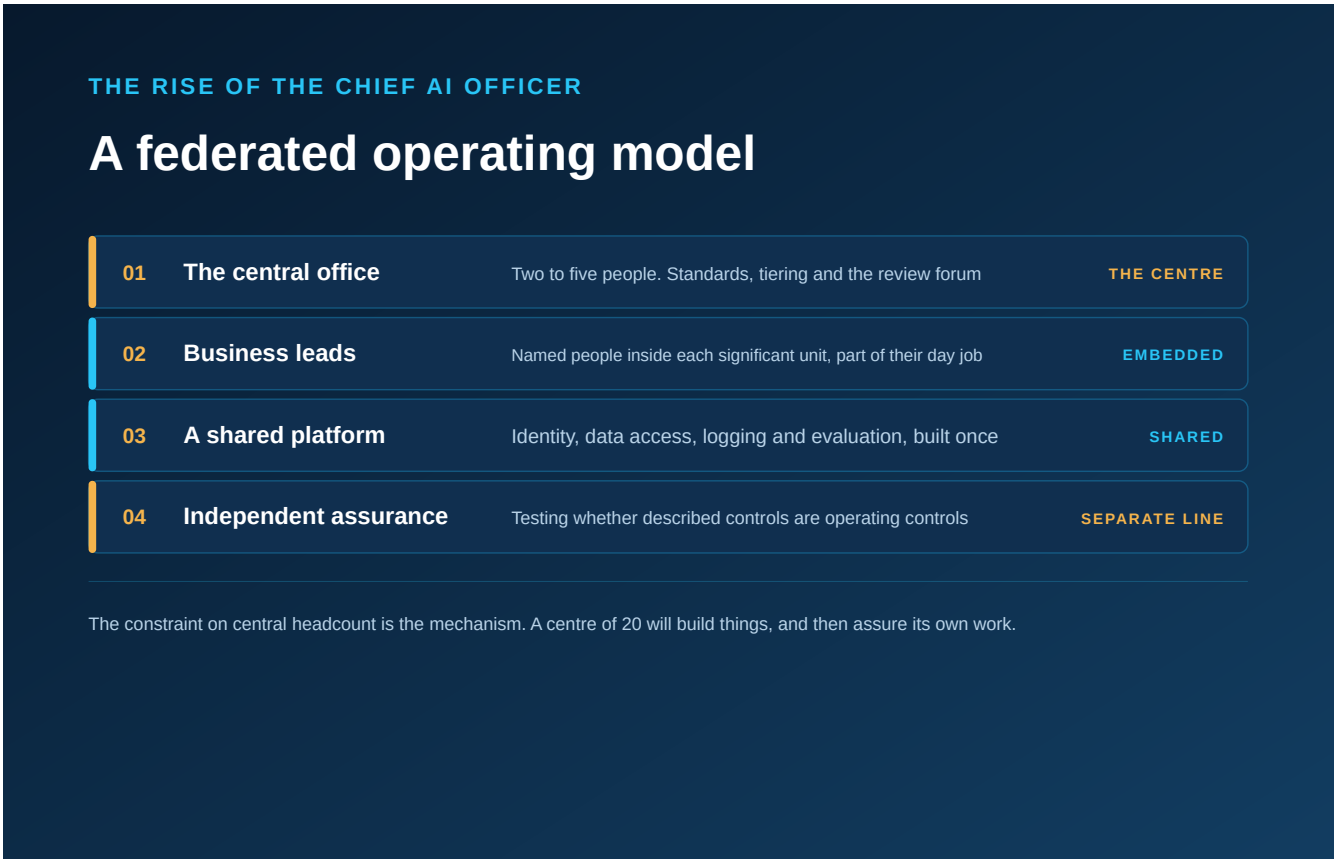


FIGURE 2 A small centre, business leads embedded where the work is, a shared platform and an independent assurance line. The dotted relationships are the ones that fail first.

My framework. AUTHOR

1. The central office is two to five people, and it owns standards, tiering, intake, the review forum, the portfolio view and the reporting. It is deliberately too small to build anything, and that is the point.
2. Business leads are named people inside each significant unit, part-time, accountable to their unit and connected to the centre. They hold the domain knowledge the centre cannot have.
3. A shared platform covers identity, data access, logging, evaluation tooling and the approved model routes, and it is owned by the technology function.

4. Independent assurance tests whether described controls are operating controls, and it reports separately.

Why the centre must stay small

A central team of 20 will build things. It can, and building is more satisfying than governing. It will then be assuring its own work, competing with the units it was created to coordinate, and defending a portfolio it authored. The constraint on headcount is the mechanism that prevents this, and it should be set on purpose, before the budget round decides it by accident. AUTHOR

There is a useful piece of evidence bearing on why coherence is worth owning at all. A field experiment published in September 2023 with 758 consultants found that on tasks inside the system's capability, participants completed 12.2 per cent more tasks and 25.1 per cent faster with materially higher quality, while on a task selected to sit just outside that capability they were 19 percentage points less likely to produce a correct answer.[6] PREPRINT The boundary was invisible to the people working. An organisation without a mechanism for locating it will see the aggregate of both effects and conclude something in the middle, which is true of nothing.

SECTION 5

Decision rights, written down

Seven decisions need owners. Six usually have one somewhere, and the seventh is what stalls deployments.

The exercise that produces most of the value in the first 60 days takes half a day. Put the adjacent officers in a room, list the decisions, and agree in writing who proposes, who funds, who approves, who accepts residual risk, who operates, who monitors and who may stop a running system.

Six of the seven are usually assigned already, written down or not. Acceptance of residual risk is the exception, and its absence explains a large share of the deployments that sit finished and unshipped for a quarter. Everybody can see the risk, nobody is authorised to accept it, and the decision defaults to whoever is willing to proceed without a decision. AUTHOR

The authority to stop

The right to halt a running system is most often missing entirely. It is distinct from the right to approve. A system that passed review in January may behave differently in June because the supplier changed the model. Somebody must be able to suspend it without assembling a committee, and that authority should be named in the operating note before an incident makes the question urgent.

> **A TEST FOR WHETHER DECISION RIGHTS ARE REAL**
>
> Ask three people independently who may accept residual risk on a high-consequence deployment and who may stop one. Where the answers differ, the rights are not assigned however thoroughly they are documented.
>
> Ask also for an instance in the last six months of either authority being used. A right that has never been exercised in an organisation doing this work at any scale is usually a right that nobody believes they hold.

SECTION 6

The first 100 days

Discover, decide, mobilise. Credibility is earned in the first period, and an officer who spends it on a strategy document has spent it badly.

The pressure on arrival is to produce a strategy, because a strategy is what the appointment was announced with. It is the wrong first artefact. Nobody in the organisation can evaluate a strategy for this before they know what is already happening, and the officer who arrives with one signals that they have not looked.

DISCOVER, DECIDE, MOBILISE



FIGURE 3 The first 30 days establish what exists, the next 30 decide what the organisation will and will not do, and the last 40 build the mechanism. Credibility is earned in the first period and spent in the third.

My framework. AUTHOR

THE FIRST 100 DAYS, WITH AN OUTPUT FOR EACH PERIOD

PERIOD	WHAT HAPPENS	WHAT EXISTS AT THE END OF IT
Days 1 to 30	Inventory under amnesty, read every existing policy and assessment, interview the adjacent officers and four business leaders, and find the work already running.	A register of what exists by tier, a list of the three things most likely to cause an incident, and a map of who currently decides what.
Days 31 to 60	Set tiering, publish the intake route, agree decision rights with the adjacent officers in writing, and name what the organisation will not do.	A one-page operating note that the executive committee has approved, and a prohibited list with an owner and a review date.
Days 61 to 100	Stand up the review forum, run three real assessments through it, appoint business leads, and publish the first dashboard.	A forum with minutes, three assessed use cases, named leads in the business, and a board-ready page showing inventory, accepted risks and incidents.

SEVEN MEASURES FOR THE ROLE ITSELF

MEASURE	WHAT IT SHOWS	THE FAILURE IT DETECTS
Governed inventory coverage	Proportion of known systems registered and tiered.	Shadow use, and a register that records only what was easy to find.
Time from intake to decision	Median elapsed days, by tier.	The office becoming a queue, which is the most common way the role fails.
Evaluation coverage	Proportion of high-consequence systems with a current evaluation set.	Controls asserted rather than tested, and evaluations that have gone stale after a model change.
Portfolio value realised	Measured benefit against baseline on deployed workflows, including quality.	Activity reported as value, and pilots that never reached an operational owner.
Incidents and near misses	Count and severity, with near misses recorded separately.	A reporting culture that has stopped functioning. Zero is a warning rather than a result.
Stopped work	Number of proposals declined or halted, with reasons.	A governance function that has never said no, which means it is not governing.
Capability distribution	Where trained and practising people actually sit.	Capability concentrated in the centre, which is the signature of a model that has not federated.

The measure worth watching hardest is time from intake to decision, because it is the leading indicator of the failure mode this role is most prone to. An office whose median decision time is climbing is becoming a queue, and the response is to push more decisions down to the business leads. Adding people at the centre lengthens the queue.

The discovery that changes the plan

In every organisation I have seen attempt this, the first 30 days produce the same surprise in different clothes. There is more happening than anybody knew, it is concentrated in one or two enthusiastic functions, at least one system is connected to data nobody formally approved, and the most valuable work is being done by somebody junior with no mandate. What the officer does about that last point in the first month determines how the role is regarded for the rest of the year. AUTHOR

What to publish at day 60

One page, approved by the executive committee. Tier definitions, the intake route, who decides what, and what the organisation will not do. Not a framework, not a strategy and not a policy suite. A single page that a manager can read and act on beats a document that has to be interpreted, and it can be extended once people are following it.

SECTION 7

Most organisations should not create this post

The accountability gap is real. A new seat on the executive committee is one answer to it and rarely the best one.

This is the section the rest of the paper has been arguing towards, and it sits uncomfortably beside a memorandum issued yesterday making the role compulsory across an entire government. The memorandum applies to federal agencies with statutory duties, public accountability, and use cases that affect rights directly. Those conditions do not describe most commercial organisations. AUTHOR

SPECIFICATION

Conditions under which a dedicated post is justified

- Artificial intelligence is material to the organisation's product rather than to its internal efficiency, so that the risk sits in what customers receive.
- The organisation operates in a sector where a regulator will ask, and where the answer must come from a single accountable person.
- Existing executives are demonstrably unable to hold it, which is tested by asking them.
- There is enough portfolio to coordinate. Below roughly a dozen consequential systems the coordination problem does not yet exist.

Where those conditions do not hold, the alternative is to assign the mandate explicitly to an existing executive, with the seven areas from section three written into their objectives, a named deputy who does the work, and the governance board chaired above them in the manner the memorandum describes. That arrangement costs a fraction as much and fails less often, because it does not require the new officer to build authority from nothing.

The failure this avoids

A new executive post created without the four conditions produces a predictable sequence. The appointee has accountability without authority, the business units treat the office as an approval gate, delivery slows, the office responds by building its own capability to demonstrate value, and within 18 months it is a small competing technology function. I have watched this happen with chief digital officers and with chief transformation officers, and there is nothing about this technology that would make the outcome different.

What cannot be avoided is the underlying requirement. Somebody must be able to say what the organisation is doing, what evidence supports it, and who decided. Whether that person carries a new title is a question about organisational design. Whether they exist at all is not optional, and at the moment in a great many organisations they do not.

END MATTER

Method, evidence and limits

How the sources in this paper were chosen and dated, and what the reader should distrust.

Sources

Primary government instruments cited from their issuing bodies, one field experiment, and one industry survey cited with its method. Commentary on the memorandum is not cited, because the memorandum is public and readable.

Administrative and legal status

The memorandum binds federal agencies and creates no private obligation, which the text states wherever it is cited. The European regulation is described at first-reading adoption by the Parliament, awaiting Council adoption and publication, and the paper deliberately avoids the official citation number because none has been assigned at this date.

The absence of an adoption statistic

I could find no defensible survey of Chief AI Officer prevalence in private-sector organisations for any date before this paper closed. Circulating figures are vendor marketing. None appears here, and a reader encountering such a figure elsewhere should ask for its fieldwork date.

On the one study cited

The field experiment in section five is a Harvard Business School working paper, and it carries the preprint tag for that reason. It is used for the shape of its finding, and the effect sizes are treated as indicative.

What remains open

Whether federated models outperform centralised ones, which is untested. Whether the post correlates with any outcome, which cannot be known while the population of holders is this small and this new. And what the European regulation will finally require, which depends on steps that have not happened.

Where the reader should push back

The seven-part mandate, the boundary table, the federated model with its headcount constraint, the decision rights and the four conditions in section seven are my own. The argument that most organisations should not create the post is a judgement, offered with its reasoning exposed so that a reader who rejects it can say which step failed.

Interest declared

I advise boards on operating models for this technology, and the work described in section six is the kind of work I do. The paper argues against creating the executive post that generates the largest such engagements, which I mention so that the reader can weigh it.

END MATTER

Sources considered and set aside

Four surveys and statements I read and chose not to use. None is a bad source; none answers the question this paper asks.

SOURCE	PUBLISHED	WHY IT WAS SET ASIDE
Deloitte, State of Generative AI in the Enterprise, first quarterly report	January 2024	A self-selected panel of about 2,800 leaders reporting sentiment. It asks who is in charge and gets the answer a survey of senior people always gets.
Foundry, State of the CIO 2024	January 2024	A survey of technology leaders about their own remit. It shows information officers expecting to own this, which is a data point about information officers.
IBM, Global AI Adoption Index 2023	10 January 2024	A vendor-sponsored survey of 8,584 information technology professionals, fielded in November 2023. It measures deployment intentions, and this paper is about accountability.
Council of the European Union, press release on the provisional agreement	9 December 2023	The press release describing the political agreement. Useful for the shape of the deal and not a legal text, and section two uses the Parliament's adopted position instead.

END MATTER

About the author

Paul Forrest advises boards and executive teams on the adoption of new technology in operating businesses, with a practice grounded in delivery rather than in advisory work alone.

His background is in business process reengineering, which he has practised since the 1990s, and in applied natural language processing, which he has worked on since 2014. The scepticism in section seven is drawn from the first of those. Reengineering produced a generation of new executive posts created to own a change nobody else would own, and the ones that worked were almost always the ones where the authority existed before the title did.

He writes about regulation as a board adviser rather than as a lawyer, and on the condition that a claim carries its evidence.

END MATTER

References

Six works cited in the text, each with its original date of publication. None is later than 28 March 2024.

1. **United States Office of Management and Budget.** *M-24-10: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence.* Memorandum to heads of executive departments and agencies.
   PUBLISHED 28 MARCH 2024

2. **The White House.** *Executive Order 14110 on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.* Section 10.1(b) directs the designation of agency Chief AI Officers.
   PUBLISHED 30 OCTOBER 2023

3. **United States Office of Management and Budget.** Draft memorandum released for public comment, with a Federal Register request for comments published 3 November 2023 and comments closing 5 December 2023.
   PUBLISHED 1 NOVEMBER 2023

4. **European Parliament.** Position adopted at first reading on the proposed artificial intelligence regulation, 523 votes to 46 with 49 abstentions, Texts Adopted P9_TA(2024)0138.
   PUBLISHED 13 MARCH 2024

5. **European Commission.** Decision establishing the European AI Office, applying from 21 February 2024.
   PUBLISHED 24 JANUARY 2024

6. **Dell'Acqua, F., McFowland, E., Mollick, E., and others.** *Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality.* Harvard Business School working paper 24-013. 758 consultants, 18 tasks.
   PUBLISHED 18 SEPTEMBER 2023

END MATTER

Further sources consulted

Eight sources I read for this paper and did not cite. They are listed so that the reading behind it can be seen in full, and none is later than 28 March 2024.

- **International Organization for Standardization and International Electrotechnical Commission.** *ISO/IEC 42001, Information technology, Artificial intelligence, Management system.* Clause 5.3 addresses organisational roles, responsibilities and authorities.
  PUBLISHED 18 DECEMBER 2023
- **McKinsey and Company.** *The state of AI in 2023: Generative AI's breakout year.* Cited for the proportion of organisations reporting policies governing employee use.
  PUBLISHED 1 AUGUST 2023
- **National Institute of Standards and Technology.** *Artificial Intelligence Risk Management Framework 1.0.* NIST AI 100-1. The Govern function assigns accountability structures to senior leadership.
  PUBLISHED 26 JANUARY 2023
- **United States Congress.** *AI in Government Act of 2020*, Public Law 116-260, division U, title I, section 104.
  PUBLISHED 27 DECEMBER 2020
- **United States Congress.** *Advancing American AI Act*, Public Law 117-263, division G, title LXXII, subtitle B. Required agency artificial intelligence leads and use case inventories before the officer title existed.
  PUBLISHED 23 DECEMBER 2022
- **United States Department of Defense.** Establishment of the Chief Digital and Artificial Intelligence Office. The earliest federal precedent for a consolidated artificial intelligence leadership office.
  PUBLISHED 1 FEBRUARY 2022
- **Group of Seven.** *Hiroshima Process International Code of Conduct for Organizations Developing Advanced AI Systems.*
  PUBLISHED 30 OCTOBER 2023
- **Countries attending the AI Safety Summit.** *The Bletchley Declaration.*
  PUBLISHED 1 NOVEMBER 2023

WHERE THIS GOES NEXT

The next paper returns to the people doing the work

This paper concerned accountability at the top. The seventh will return to the bottom, where employees have been adopting these tools faster than any programme has been able to channel them, and ask what an organisation does with demand it did not create. It is due at the end of May.

The largest vendor published a workforce survey in May last year and I would expect another around the same time. If it comes, I expect its adoption figure to be larger than most organisations' own telemetry suggests, and the gap between the two is the subject of that paper.

WHAT TO HAVE WRITTEN DOWN BEFORE THE NEXT PAPER ARRIVES

- A written allocation of the seven decisions, including who accepts residual risk and who may stop a running system.
- A one-page operating note approved by the executive committee, rather than a framework in draft.
- A median time from intake to decision, measured rather than estimated.

Paul Forrest · 29 March 2024 · Version 1.0 · Research closed 28 March 2024