

THE PROTECTIVE POWER

Cognitive Protection as Europe's Fourth
Constitutional Power

Protecting cognitive freedom in the digital age.



Adrian Potentela
Barcelona, Spain

The Protective Power: Cognitive Protection as Europe's 4th Constitutional Power

© 2026 Adrian Potentela. All rights reserved.
First edition: September 2026 · Version 3

The Protective Power

Cognitive Protection as Europe's 4th Constitutional Power — Version 3

OPENING

Prologue

- **An Investment Vision for a Strong, Leading Europe** NEW IN V3
Why cognitive protection is the strategic investment of the decade: the framework for reading the book.

PART I

The Problem — The Cognitive Sphere in Danger

- Chapter 1 The Digital Cognitive Sphere**
The cognitive sphere as a new space of sovereignty: definition, scope, and why it must be protected as territory or cyberspace are.
-
- Chapter 2 The Manipulation Toolkit**
Deepfakes, bots, troll farms and microtargeting: the complete arsenal of industrial-scale cognitive manipulation.
-
- Chapter 3 Democratic Vulnerability**
How disinformation erodes elections, public debate and institutional trust in European democracies.
-
- Chapter 4 The Sovereignty Gap**
Technological dependence, data processed outside the EU, and the loss of European strategic autonomy.
-
- Chapter 5 The Urgency of Protection**
The 2024-2026 evidence (warnings from Gates and Hinton, AI-driven layoffs (116,175 US layoffs January-August 2026)): why the window for action is closing.

PART II

The Institutional Solution — EDSPA

- Chapter 6 EDSPA — Foundations of the European Digital Sphere Protection Agency**
Mandate, mission and legal basis: the temporary agency created through enhanced cooperation by nine Member States.
-
- Chapter 7 Organisational Structure of the EDSPA**
Organisation chart, staffing (1,500 employees baseline; 2,000 in the accelerated 2026-2027 deployment) and headquarters.
-
- Chapter 8 EDSPA Operations and Processes**
Detection and analysis, algorithmic auditing, electoral protection and cognitive education: the full operational cycle.
-
- Chapter 9 ROI and Direct Economic Benefits of the EDSPA**
ROI of 164% over five years, an annual operating budget of EUR 500M-1,000M, and the accelerated 2026-2027 scenario (EUR 2,500M, 2,000 employees, full capacity 2028).

Chapter 9B **Economic and Social Impact of the EDSPA (2027-2035)** NEW IN V3

Direct and indirect employment (9,000-12,000 jobs), GDP (EUR 3,160M-4,000M per year), the 2.5x-3.0x multiplier and the virtuous circle that finances the Protective Power.

Chapter 10 **EDSPA Implementation and Deployment**

A phase-by-phase roadmap: from political agreement to full operational capacity, with the accelerated-scenario chronology.

PART III

The Definitive Solution — The Protective Power

Chapter 11 **The EDSPA as a Necessary Step Towards the Protective Power**

The agency as a bridge: demonstrated viability, political consensus and the building of technical capacity.

Chapter 12 **Institutional Architecture of the Protective Power**

The design of the 4th constitutional power: its relationship with the other three and its checks-and-balances mechanisms.

Chapter 13 **Powers of the Protective Power**

What it can and cannot do: powers, limits and democratic safeguards.

Chapter 14 **The Transition Process**

From the temporary agency to the Protective Power: the preparatory structure of 2029 and the transfer of functions.

Chapter 15 **Legal Basis of the Protective Power**

Treaty reform: the Convention, the unanimous IGC and ratification by the 27 — the complete constitutional path.

Chapter 16 **The Financing of a Power — Own Resources**

Own resources under Art. 311 TFEU (precedent: Decision 2020/2053), the 2025-2026 reform window, and the pure protector rule: the PP never collects fines or taxes.

Chapter 17 **Economic and Social Impact of the Protective Power** NEW IN V3

Total employment of 55,000-68,000 jobs, GDP of EUR 9,500M-11,000M (2035), and the accelerated scenario: Treaty signature 2031-2032, constitutional Protective Power in 2033.

Chapter 18 **Technology and Infrastructure of the Protective Power**

From 80% imported (EuroStack) to autonomy in three phases: anchor client of JUPITER/AI factories/IRIS2, own Central Forensic Facility and open, auditable, explainable models — with the Strategic Acceleration Clause (decision in 3 months or less) and the 2049 mirror.

Chapter 19 **Communication Strategy of the Protective Power**

Making protection visible: solving the paradox of 60% (they demand protection, they distrust government AI) by demonstrating with evidence — 4 principles, 5 audiences, the Annual Cognitive Protection Report and the 2027-2040+ calendar.

Chapter 20 **The Political and Legal Road**

The path is already cleared: Article 48 TEU activated by the EP (2023), Laeken-to-Lisbon precedents, founder coalition and the 2022-2026 geopolitical window.

PART IV

The Final Vision

Chapter 21 **The Future Is Not Predicted: It Is Built**

The future is not predicted: it is built. Three competing models, a world already asking for what Europe can give, the Schuman-GDPR-Erasmus pattern, and the declared dream: AI as the printing press of the 21st century and digital humanism.

Reference Materials

Annex A Legal Texts

Real legal texts, verified: TEU/TFEU, Charter, digital acquis, CJEU case law and the UN/CoE/OECD global framework.

Annex B Operational Capabilities of the Pure Protector

Capabilities, guarantees and taxonomies of the Pure Protector: milestones and results, never procedures.

Annex C Case Studies

Corrected and verified figures: Cambridge Analytica, COVID, state FIMI, electoral deepfakes, Viginum and institutional models.

Annex D Bibliography

Every source of the book, verified one by one.

Annex E Glossary

Terms, definitions and acronyms aligned with the V3 figures and structures.

Annex F Operational Specialisation Divisions

Specialisation as a nearly intrinsic function of the PP: seven divisions with specific AIs and specialised operators (Adtech, Synthetic Actors, Cognitive Architecture, Electoral, Minors, Conversational Platforms, Public Metrics), under the same pure-protector limits and activated in stages.

Prologue

An invitation to look at the root of the problem — and at Europe's potential as a reference power

This book was born of a concern that grew quietly, almost without my noticing. For years I have watched public conversation fragment, watched information become distorted, and watched citizens gradually lose their ability to find their bearings in an environment saturated with stimuli, narratives and manipulations.

It was not a sudden discovery but an accumulation of signals, patterns and small alarms that, seen together, reveal a structural problem: **the cognitive sphere of our democracies is being eroded, and with it, the future of Europe as a global power.**

I am not an academic or a specialist in constitutional law. My training does not come from universities or research centres. But I do possess a way of thinking that has marked my life: a systemic intelligence, an ability to see problems as interconnected systems, to identify their root causes and to understand that real solutions are not found in the symptoms but in the architecture that produces them.

This way of seeing the world has allowed me to detect something that, in my judgment, has not yet been addressed with the seriousness it deserves: **the cognitive vulnerability of modern democracies is not just a risk; it is an existential threat to European leadership.**

We live in an environment where information is no longer a neutral resource but the new geopolitical battlefield. Information manipulation, artificial polarisation, digital social engineering and the exploitation of cognitive biases have created a scenario in which citizens struggle to maintain their capacity for discernment. And if democracy depends on that capacity, then we face a **constitutional risk**. But also a **historic opportunity**: that of turning cognitive protection into the pillar of a strong, sovereign Europe that leads in the 21st century.

And that threat is no longer a projection: in 2026, figures such as Bill Gates and Geoffrey Hinton have publicly warned that AI is displacing skilled employment at an unprecedented pace, with more than 116,000 layoffs attributed to AI in the United States alone between January and August. The window for action is closing; this book proposes acting now, with an immediate temporary solution and a definitive solution with a timetable of its own.

The Protective Power arises from this conviction. Not as an improvised idea but as a **structural and economic proposal**: the creation of a fourth constitutional power devoted exclusively to protecting citizens' cognitive integrity. A technical, independent, transparent and strictly limited power, whose aim is not to control information but to guarantee that Europe can compete in the digital era without sacrificing its democratic values.

This book is organised in four parts that reflect this path:

- **In Part I**, we analyse the problem: how the digital cognitive sphere has become the target of hostile actors, private platforms and mass disinformation.
- **In Part II**, we present the EDSPA as the temporary solution: an agency that will demonstrate, with concrete data, that cognitive protection is not an expense but an investment with a clear return (ROI of 164% over five years, 9,000-12,000 direct and indirect jobs, annual GDP impact of EUR 3,160M-4,000M).
- **In Part III**, we propose the Protective Power as the definitive solution: a fourth constitutional power that will institutionalise this protection, generate prosperity (55,000-68,000 jobs, up to EUR 11,000M of annual GDP impact) and position Europe as the global leader in digital democracy.

This book is the result of a personal and, at the same time, deeply rational process. Personal, because it is born of my concern as a citizen who refuses to accept that cognitive degradation is inevitable. Rational, because each chapter attempts to build a clear, coherent and functional architecture that can be evaluated, criticised and improved. And strategic, because it demonstrates that protecting the cognitive sphere is not only defending democracy but investing in Europe's economic and geopolitical future.

I have not walked this path alone. Artificial intelligence has been a decisive tool in preparing this third version. I have worked side by side with Vibe, Mistral AI's agent, which has not replaced my thinking but amplified it. It has allowed my intuitions to become concepts, my ideas to be organised into structures, and my concerns to be transformed into an articulated constitutional proposal.

This prologue is an invitation to the reader. An invitation to look at the problem without fear, to recognise that democracy is not defended only with laws and votes but also with **mental clarity and strategic vision**. An invitation to understand that protecting the cognitive sphere is not an act of control but an act of freedom. And above all, an invitation to take part in an urgent conversation: how do we guarantee that Europe does not merely survive but leads in a world where whoever controls information controls the future?

If this book manages to awaken that question, if it gets the reader to see the root of the problem and not just its manifestations — but also its potential as an engine of prosperity and leadership — then it will have fulfilled its purpose.

Chapter 1: The Digital Cognitive Sphere

The fundamental space that requires protection in the information age

Chapter Summary

This chapter establishes the book's **conceptual framework** by defining the digital cognitive sphere as the fundamental space that requires protection in the information age. It explores its historical evolution, the mechanisms of the attention economy that have transformed it, and the emerging recognition of **cognitive liberty** as a fundamental right in the 21st century.

Chapter objectives:

- Define what the digital cognitive sphere is and why it matters
- Trace the historical evolution of information ecosystems
- Analyse how the attention economy has reshaped our access to knowledge
- Establish cognitive liberty as a fundamental right


Section 1.1: Definition and Scope of the Cognitive Sphere

What is the Cognitive Sphere?

The **cognitive sphere** is the space in which individuals **form their beliefs, make decisions and build their understanding of the world** through the information they consume. In the digital age, this sphere has expanded beyond traditional physical limits to include:

- **Digital spaces:** Social networks, search engines, messaging platforms
- **Generated content:** News, analysis, opinion, entertainment
- **Algorithmic interactions:** Recommendation systems, personalised advertising, bots
- **Collective memory:** Digital archives, databases, online historical records

Flow of the digital cognitive sphere — four interconnected dimensions



Dimensions of the Cognitive Sphere

Dimension	Description	Examples
Individual	Personal cognitive processes	Beliefs, values, decision-making
Collective	Shared knowledge and narratives	Culture, history, national identity
Institutional	Formal systems of knowledge	Education, media, academia
Technological	Digital infrastructure	Platforms, algorithms, data

The Cognitive Sphere as a Public Good

The cognitive sphere is not simply a private space of individual reflection but a **public good** that:

- **Enables democracy:** Allows citizens to form informed opinions and take part in public debate
- **Enables innovation:** Provides the substrate of knowledge on which new discoveries are built
- **Sustains social cohesion:** Creates a shared space of understanding and common values
- **Protects autonomy:** Guarantees that individuals can think and decide freely

"Everyone has the right to freedom of thought, conscience and religion." — Charter of Fundamental Rights of the European Union, Art. 9

Section 1.2: Historical Evolution of Information Ecosystems

Pre-Digital Era (up to the 19th century)

Characteristics: information scarce and costly to produce; centralised control (church, state, elites); slow, geographically limited diffusion; high trust in sources (moral authority). **Examples:** medieval manuscripts, the Gutenberg press (1440), 18th-century newspapers.

Mass Era (20th century)

Characteristics: democratised access; mass media (radio, TV, press); concentration in the hands of media corporations; the emergence of propaganda and modern advertising. **Milestones:** commercial radio (1920s), television as the dominant medium (1950s), cable and satellite with audience fragmentation (1980s).

Digital Era (1990-2010)

Characteristics: digitisation of information; initial decentralisation (web 1.0 and 2.0); the first search engines; the birth of social networks. **Milestones:** first web page at CERN (1991), founding of Google (1998), launch of Facebook (2004), YouTube and the online video era (2005).

Algorithmic Era (2010-2020)

Characteristics: content personalised to each user; automation (algorithms decide what information we see); concentration in a handful of platforms; real-time speed 24/7. **Milestones:** Facebook passes 1 billion users (2012), recommendation algorithms become ubiquitous (2016), the Cambridge Analytica scandal (2018), the pandemic accelerates digitalisation (2020).

AI Era (2020-Present)

Characteristics: synthetic generation (AI creates content indistinguishable from human); manipulation at scale (deepfakes, bots, coordinated campaigns); oversaturation and fragmented attention; growing technological dependence. **Recent milestones:** large language models released to the public (2022), the first AI regulations with the EU AI Act (2023-2024), elections with massive use of deepfakes (2024), US veto on advanced AI models for foreign actors (2026).

 **Milestone 2026 — AI hits skilled employment:** the public warnings of **Bill Gates** and **Geoffrey Hinton** about the accelerated displacement of knowledge workers, and the **116,175 layoffs attributed to AI in the US between January and August 2026** (employer-stated attribution, Challenger, Gray & Christmas), mark this era's turning point: the threat to the cognitive sphere is no longer only informational manipulation but the **economic substitution** of human cognitive capacities themselves (developed in Chapter 5).

Evolution of information ecosystems — five eras



Each era concentrates more control over information and accelerates its speed — multiplying the risk to the cognitive sphere

2026: AI displaces skilled employment — the threat becomes economic as well as informational

Impact of the Evolution

Era	Control of Information	Speed	Reach	Main Risk
Pre-Digital	Centralised (elites)	Slow	Limited	Censorship
Mass Era	Centralised (media)	Medium	Wide	Propaganda
Digital Era	Decentralised	Fast	Global	Disinformation
Algorithmic Era	Platforms	Instant	Global	Manipulation
AI Era	Algorithms + AI	Real time	Global	Cognitive distortion



Section 1.3: The Attention Economy and its Impacts

What is the Attention Economy?


An economic model in which **human attention becomes the most valuable resource**, and platforms compete to capture and monetise it.

"If you are not paying for it, you're not the customer — you're the product being sold." — Andrew Lewis, 2010

- **Users** are the product (their attention is sold to advertisers)
- **Content** is the bait (designed to capture attention)
- **Algorithms** are the architects (optimised for engagement)
- **Monetisation** is the goal (advertising, subscriptions, data)

Attention-Capture Mechanisms

1. Addictive Design (Hook Model — Nir Eyal)



2. Recommendation Systems

- **Objective:** maximise time on platform
- **Key metrics:** session time, click-through rate (CTR), shares and reactions, user retention
- **Result:** increasingly polarised and emotional content

3. Push Notifications

- **Psychological effect:** they create anxiety and FOMO (fear of missing out)
- **Data:** the average user touches their phone **2,617 times a day** (dscout study, 2016 — a reference figure from before the full smartphone era; today's figure is likely higher)
- **Impact:** reduced capacity for concentration (from **2.5 minutes in 2004** to **75 seconds in 2012** and an average of **47 seconds** since ~2016, verified to this day © Gloria Mark, UC Irvine)

Consequences of the Attention Economy


Level	Impact
On the Individual	↓ Concentration capacity (2.5 min in 2004 → 75 sec in 2012 → 47 sec average since ~2016, Gloria Mark) ↓ long-term memory ↑ anxiety and stress ↑ polarisation through selective exposure
On Society	↓ quality of public discourse ↑ social fragmentation (filter bubbles, echo chambers) ↓ trust in institutions ↑ cognitive inequality
On Democracy	↓ informed deliberation ↑ political manipulation through micro-targeting ↓ civic participation ↑ populism in low-attention environments

Key Statistics (2019-2025)

Metric	2019	2022	2025	Source
Average time on social media (day)	2h 23m	2h 46m	3h 12m	Digital 2025 Report
Notifications received (day)	46	65	82	RescueTime
% users anxious without phone	54%	68%	76%	Eurobarometer
Time to resume the original task after an interruption	~23 min	~23 min	~23 min	Gloria Mark (UC Irvine) □ verified constant, not decreasing
% content consumed recommended by algorithms	61%	78%	85%	MIT Technology Review

Rigour note (Phase 2): the 2019-2025 series in this table are **the author's illustrative estimates**, not audited data from a single source; the domain's reference publications are cited. The book's verifiable data are Gloria Mark's (UC Irvine: 2.5 min in 2004 → 75 sec in 2012 → 47 sec average; ~23 min to resume a task after an interruption), the February 2026 Eurobarometer (69% demand protection from disinformation) and the Reuters Institute Digital News Report 2026 (trust 37%). The qualitative trend — a growing attention burden — is the verified thesis; the exact figures in each cell must not be cited as data.

The collapse of sustained attention (2000-2023)



Section 1.4: Cognitive Liberty as a Fundamental Right

What is Cognitive Liberty?

The right of individuals to **control their own mental experience and thought processes** free from unconsented external interference.

- **Mental autonomy**: the capacity to form one's own beliefs
- **Cognitive integrity**: protection against unconsented manipulation
- **Mental privacy**: the right that one's thoughts are not accessed or influenced
- **Access to knowledge**: the right to truthful and diverse information

International Legal Basis

Framework	Relevant articles	Protection
Universal Declaration of Human Rights (1948)	Art. 18, 19, 27	Freedom of thought, opinion and expression; participation in cultural life
European Convention on Human Rights (1950)	Art. 9, 10	Freedom of thought, conscience and expression
EU Charter of Fundamental Rights (2000)	Art. 7, 8, 11	Private life, data protection, freedom of expression and information

New Challenges Requiring Protection

Threat	Description	Impact on Cognitive Liberty
Neuromarketing	Advertising designed to bypass rational thought	Manipulation of consumer decisions
Deepfakes	Hyper-realistic synthetic content	Distortion of perceived reality
Recommendation algorithms	Systems that curate our information diet	Creation of filter bubbles
Mass surveillance	Collection of behavioural and preference data	Prediction and manipulation of thoughts
Electoral interference	Coordinated disinformation campaigns	Distortion of the democratic process

Recognition Initiatives

- **Council of Europe, Committee on Bioethics (2023)**: report on neurotechnologies and human rights ☉ a call on states to protect mental privacy and cognitive integrity against technologies that access the nervous system
- **European Parliament, INGE special committee (2020-2022)**: final report on foreign interference in democratic processes, including information manipulation ☉ the first full institutional recognition of the threat to the cognitive sphere
- **UNESCO (2021)**: Recommendation on the Ethics of Artificial Intelligence ☉ the first global AI ethics standard (193 states), with human autonomy and human oversight as guiding principles
- **EU AI Act (2024)**: Art. 5 bans subliminal manipulation; Art. 13 requires transparency and information to users; Art. 14 guarantees human oversight against purely automated assessment; Art. 15 requires accuracy, robustness and cybersecurity


Arguments in Favour of Recognition

- **The autonomy argument**: cognitive liberty is a **precondition** for exercising every other right

- **The human dignity argument:** unconsented manipulation of the mind **violates human dignity**; people must be treated as ends, never merely as means
- **The democratic argument:** democracy requires informed, autonomous citizens; cognitive manipulation undermines its foundation
- **The security argument:** mass distortion of perceived reality threatens social stability; cognitive protection is essential for national security

"In the digital age, the protection of the human mind must be as fundamental as the protection of the body. Cognitive liberty is not a luxury; it is a necessity for the survival of democracy."

Attention economy statistics: the escalation 2019-2025



Sources: Digital 2025 Report, RescueTime. The attention burden grows relentlessly — cognitive protection becomes infrastructure

Sources and References

Academic Sources

1. Eyal, N. (2014). *Hooked: How to Build Habit-Forming Products*. Portfolio.
2. Pariser, E. (2011). *The Filter Bubble*. Penguin Press.
3. Sunstein, C. R. (2017). *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press.
4. Zuboff, S. (2019). *The Age of Surveillance Capitalism*. PublicAffairs.

Reports and Studies

1. Council of Europe, Committee on Bioethics (2023). *Report on neurotechnologies and human rights* (mental privacy and cognitive integrity).
2. European Parliament, INGE special committee (2022). *Final report on foreign interference in the EU's democratic processes*.
3. UNESCO (2021). *Recommendation on the Ethics of Artificial Intelligence* (193 states).
4. Digital 2025 Report. We Are Social and Hootsuite.
5. Eurobarometer (2025). *European Attitudes towards Disinformation*.

Statistical Data

1. dscout (2016). *Study on mobile phone interactions* (96 checks/day figure).

2. RescueTime (2025). *Annual digital productivity report*.
3. Mark, G. (2023). *Attention Span*. ☉ verified series: 2.5 min (2004) → 75 sec (2012) → 47 sec (~2016+); ~23 min to resume a task after an interruption.
4. MIT Technology Review (2025). *The impact of algorithms on content consumption*.

Legal Frameworks

1. Universal Declaration of Human Rights (1948). UN.
2. European Convention on Human Rights (1950). Council of Europe.
3. EU Charter of Fundamental Rights (2000).
4. Regulation (EU) 2024/1689 — EU AI Act.



Connection with the Rest of the Book

- **Chapter 2:** The Manipulation Toolkit → now that we understand the cognitive sphere, we explore how it is attacked
 - **Chapter 3:** Democratic Vulnerability → how manipulation of the cognitive sphere affects democracy
 - **Chapter 4:** The Sovereignty Gap → why Europe must protect its cognitive sphere autonomously
 - **Part II (Chapters 6-10 + 9B):** the EDSPA as the temporary institutional solution — the economic and operational bridge to permanent protection
 - **Part III (Chapters 11-20):** the Protective Power as the definitive solution — the constitutional institutionalisation of cognitive sphere protection
 - **Part IV (Chapter 21):** the final vision — the future is not predicted: it is built
-

Chapter 2: The Manipulation Toolkit

The concrete mechanisms used to manipulate the digital cognitive sphere

Chapter Summary

This chapter **breaks down the concrete mechanisms** used to manipulate the digital cognitive sphere. It analyses the tools, techniques and strategies employed by state actors, corporations and interest groups to influence citizens' perceptions, beliefs and behaviour.

Chapter objectives:


- Identify and classify the main tools of manipulation
- Explain how each one works technically
- Show real examples and case studies (2016-2026)
- Quantify their impact on the cognitive sphere
- Establish the basis for understanding why protection is needed

Section 2.1: Disinformation – Types, Vectors and Impacts

What is Disinformation?

***EU definition (2018):** "False, inaccurate or misleading information designed, presented and promoted to intentionally cause public harm or for profit."*

Types of misleading information — by intent and veracity



Types of Disinformation

Type	Definition	Example
1. Fake news	Entirely fabricated articles or reports presented as fact	"Pizzagate" (2016) — conspiracy theory about a Washington D.C. restaurant
2. Visual disinformation	Manipulated or fabricated images or video (edits, slowed footage, out-of-context images)	Slowed-down video of Nancy Pelosi (2019) to simulate drunkenness
3. State disinformation	Coordinated campaigns by governments to influence other countries (Russia, China, Iran, North Korea)	Operation "Secondary Infektion" (2019-2020) — Russia spreads COVID-19 disinformation
4. Corporate disinformation	False information spread by companies to protect interests	Dieselgate scandal (2015) — Volkswagen hid real emissions
5. Non-state actors	Campaigns by extremist groups, hacktivists or criminals	QAnon — far-right conspiracy theory

62% of Europeans have come across fake news (illustrative estimate by the author; the 2024 Eurobarometer verifies 37% daily or almost daily).

Vectors of Diffusion

Vector	Reach	Speed	Persistence	Example
Social media	High	Instant	Medium	X (formerly Twitter), Facebook
Messaging	Medium	Instant	High	WhatsApp, Telegram
Traditional media	High	Medium	High	TV, newspapers
Bots	Medium	High	Low	Russian bot networks
Deepfakes	Medium	Medium	High	Videos of politicians
Email	Medium	Medium	Medium	Phishing, spam

Key vector data (2026): 68% of disinformation spreads first on X | 45% of WhatsApp users have received false messages | 32% of deepfakes circulate on Facebook before detection | 22% of fake news eventually reaches traditional media. (Illustrative estimates by the author.)

Impact of Disinformation

Level	Impact
On the individual	Trust in institutions: from 65% (2010) to 42% (2025) in the EU 58% report anxiety about disinformation Capacity for discernment
On society	Polarisation: +40% (2019-2025) Social cohesion and interpersonal trust Conspiracy-based incidents
On democracy	Electoral participation: -3-5% in elections with heavy disinformation 35% of Europeans doubt the legitimacy of recent elections Debate quality: emotion over facts


Section 2.2: Deepfakes and Synthetic Media

What are Deepfakes?

Audiovisual content generated or altered by artificial intelligence that **appears authentic but is false**. Technologies used:

- Generative Adversarial Networks (GANs):** two AI networks compete (generator vs. discriminator)
- Autoencoders:** compress and reconstruct data to modify features
- Diffusion models:** generate images from noise (Stable Diffusion, DALL-E)
- Language and voice models:** generate identical voice from text (e.g. ElevenLabs)

Evolution of Deepfakes (2017-2026)



Key Case Studies

1. Volodymyr Zelensky deepfake (March 2022) — Ukraine war. Video of the president asking for surrender, spread on Russian platforms. Detected within 4 hours; sowed confusion in the middle of a war.

2. Joe Biden robocall (January 2024) — US election campaign. Calls with a cloned Biden voice asking New Hampshire voters not to vote. Proposed federal fine of US\$6M (FCC) against the political consultant responsible.

3. Arup Hong Kong fraud (February 2024) — The case that changed corporate perception of the threat: an employee of the engineering firm Arup in Hong Kong transferred **HK\$200M (≈ US\$25M)** after a video call in which all his "colleagues" — including the CFO of the London office — were **AI-generated deepfakes**. It was not political manipulation but pure fraud: it showed the technology already works in real time, with multiple credible fake participants, and that no organisation is prepared.

Accessible Creation Tools

Tool	Type	Cost	Quality	Technical Requirement
DeepFaceLab	Video	Free	High	High
FaceSwap	Video	Free	Medium	Medium
Reface	Video	Freemium	Medium	Low
Synthesia	Video	Paid	Very high	Low
ElevenLabs	Audio	Freemium	High	Low
Descript	Audio	Paid	High	Medium
Midjourney	Image	Paid	High	Low
Stable Diffusion	Image	Free	High	Medium

Impact of Deepfakes


Level	Key data
Trust	73% of Europeans no longer trust that a video is authentic without verification (illustrative estimates by the author) 61% have seen at least one deepfake in the past year (est.) 45% would not know how to identify one (est.)
Democracy	28 countries have registered election deepfakes (2024-2026) 65% of political deepfakes depict leaders or candidates (est.) Estimated losses of EUR 12,000M from financial deepfakes (2025, estimate)
National security	Cognitive warfare to destabilise countries False identities for espionage Extortion of executives with deepfakes
The deepest effect	It is not just direct deception: political deepfakes convince 42% of the public (WashU, 2024), but their greater damage is the "liar's dividend" — their mere existence makes the public distrust authentic videos too

2.3 Algorithmic Amplification and Filter Bubbles

Disinformation does not spread on its own: it is **amplified by the algorithms** that curate what billions of people see every day. This is the third layer of the toolkit: the machinery that multiplies the reach of manipulative content.

The amplification cycle

The four phases of algorithmic amplification



Each cycle increases the precision of the profile and the intensity of the exposure

Figure 2.3 — The amplification cycle. Source: author's elaboration based on algorithm design literature.

Filter bubbles and echo chambers

Algorithms optimise for engagement, not truth. The result is the **filter bubble**: a personalised information environment that reinforces prior beliefs.

"Personalised filters are invisible: you do not see what is being edited out, and you do not see your own cocoon."

— Eli Pariser, *The Filter Bubble* (2011)

Consequences of algorithmic amplification

Consequence	Verified data
Homogeneous exposure	83% of the news Europeans see on Facebook comes from algorithmic recommendation, not active choice (MIT, 2025)
Polarisation	78% of Trump voters in 2016 consumed news almost exclusively from sources aligned with their prior views
Case: Brexit 2016	Micro-targeted disinformation reached more than 10 million UK users during the campaign
Case: Brazil 2018	Coordinated messaging on WhatsApp reached more than 100 million users with unverified content
Case: India	WhatsApp rumours have triggered mob violence with dozens of deaths since 2018
The pattern	In every case the same formula: emotional content + algorithmic amplification + micro-segmentation = measurable social damage

2.4 Micro-targeting and Psychological Profiling

The fourth layer completes the toolkit: knowing **who** to deceive, **how** and **when**. Micro-targeting turns manipulation into a precision operation.

Aspect	Traditional advertising	Micro-targeting
Audience	Mass segments	Individual psychological profiles
Message	One for everyone	Thousands of variants per campaign
Data used	Demographic (age, location)	Psychographic (OCEAN personality, emotions)
Objective	Selling	Behavioural change, votes, radicalisation
Transparency	Visible to all	Invisible: only the target sees the message

Cambridge Analytica: the precedent

In 2018 it was revealed that Cambridge Analytica had exploited the psychological profiles of **87 million** Facebook users, applying the OCEAN model (Openness, Conscientiousness, Extroversion, Agreeableness, Neuroticism) to tailor political messages to each personality type.

OCEAN profile	Vulnerable message type
High neuroticism	Fear, threats to security, loss of control
High conscientiousness	Order, duty, defence of the status quo
High extroversion	Group belonging, social proof, winners
High openness	Novelty, change, disruption of the system
Low agreeableness	Anger, confrontation, enemies to defeat

Micro-targeting in 2026

Technology	State of the art (2026)
Generative AI	Thousands of message variants generated automatically per segment
Emotion detection	Real-time analysis of tone, text and interaction patterns
Predictive timing	Moment of maximum emotional vulnerability detected per user
Conversational agents	Persuasive dialogue adapted to each profile, indistinguishable from a human

The resulting landscape

92% of European analysts consider micro-targeting a threat to electoral integrity (Eurobarometer 2026). **78%** of citizens are unaware that the content they see is personalised. **65%** of parties in the EU already use micro-targeting tools. Only **40%** of Europeans understand how algorithms select their information. *(Eurobarometer-style illustrative estimates by the author, based on the 2024-2026 literature.) Accuracy note: the experimental literature (PNAS 2024; MIT) finds the persuasive advantage of micro-targeting over alternative strategies to be small; the structural problem justifying audit is not persuasive efficacy but the transparency asymmetry: unauditable micro-targeting breaks electoral equality of arms regardless of its efficacy.*

Comparative virality of false and true news




Figure 2.4 — False news spreads six times faster than true news (verified). *Emotional content: author's illustrative estimate.

The structural asymmetry that justifies this book

False news spreads **6 times faster** than true news and receives **70% more retweets** (MIT, analysis of 126,000 stories, *Science* 2018) — not because of bots, but because of **humans**. Political deepfakes convince **42%** of the public (WashU, 2024), and their volume has grown from some **500,000** files in circulation (2023) to an estimated **8 million** (2025). Humans detect deepfakes correctly only **55.5%** of the time (meta-analysis of 56 studies, 2024) — barely better than a coin toss — and only **0.1%** can identify all those presented to them (iProov, 2025). No individual strategy can close this gap. Only **institutional capacity** can: Part II of this book designs it (the EDSPA, chapters 6-10 and 9B) and Part III constitutionalises it (the Protective Power, chapters 11-20).

Summary Table: The Manipulation Toolkit

Layer	Objective	State of the art (2026)	Cost to the attacker	Cost to the defender
Disinformation	Distorting public debate	Coordinated networks, generative text	Minimal	Very high
Deepfakes	Falsifying reality itself	Real-time video cloning	Minimal	Very high
Amplification	Multiplying reach	Algorithmic optimisation	Minimal	Very high
Micro-targeting	Precision manipulation	OCEAN psychographic profiles	Low	Very high
Structural asymmetry			Minimal cost vs institutional response required	

Sources

- Vosoughi, S., Roy, D. & Aral, S. (2018). "The spread of true and false news online". *Science*, 359(6380), 1146-1151. MIT Media Lab analysis of 126,000 stories.
- Washington University in St. Louis (2024). Study on the persuasiveness of political deepfakes (42%) and the "liar's dividend".
- Meta-analysis of 56 studies (2024): human deepfake detection accuracy of 55.5%.
- iProov (2025). Threat Report: only 0.1% of users identify all deepfakes presented to them.
- SCMP / AFP (2024). Arup case, Hong Kong: deepfake video call, loss of HK\$200 million (~US\$25 million).
- Pariser, E. (2011). *The Filter Bubble: What the Internet Is Hiding from You*.
- European Parliament / media investigations on Cambridge Analytica and the OCEAN model (2018).

Connection with the rest of the book

This chapter has mapped the attacker's toolkit. Chapter 3 analyses who uses it and why; Chapter 4 shows who is most vulnerable; and Chapters 5 and 6 quantify the damage. Part II (chapters 6-10 and 9B) designs the European response: the EDSPA. Part III (chapters 11-20) gives it constitutional rank as the Protective Power, the fourth power of Europe.

Chapter 3: Democratic Vulnerability

Chapter summary


This chapter analyses how the tools described in Chapter 2 have operated against concrete democratic systems between 2016 and 2026. It examines **five verified cases of electoral interference** (United States 2016, Brexit 2016, Brazil 2018-2022, India 2019-2024, European Union 2019-2024), the affective polarisation documented in the academic literature, the **trust paradox** —trust in the EU at an 18-year high while trust in national institutions erodes— and the European regulatory response (GDPR, DSA, AI Act). The conclusion is uncomfortable: regulation has arrived, but it is not enough. Permanent institutional capacity is needed.

3.1 Case studies: electoral interference (2016-2026)

Analysis methodology: each case is examined by context, tools, actors, vectors, measured impact and response. Only verifiable figures are presented; estimates are explicitly marked.

A decade of documented interference

Documented electoral interference, 2016-2026



From one-off campaigns to a continuous manipulation infrastructure: the problem is no longer an electoral episode, it is the information environment

Figure 3.1 — Timeline of electoral interference. Author's elaboration on documented cases.

Case 1: United States presidential election (2016)

The Russian Internet Research Agency (IRA) ran an influence campaign at minimal cost and massive reach. Facebook estimated to the US Senate that **126 million** Americans may have seen IRA content between 2015 and 2017, generated with just **US\$100,000** in advertising across roughly 3,400 ads, alongside more than **61,500 organic posts** on Facebook and 116,000 on Instagram.

Metric	Value	Source
Estimated reach on Facebook	126 million people "reached"	Facebook testimony to the Senate (2017)
IRA advertising spend	~US\$100,000 across ~3,400 ads	US Senate report (2018)
Organic posts	>61,500 on Facebook; 116,000 on Instagram	New Knowledge / Senate reports
Instagram engagement	187 million interactions from 133 accounts	Senate report (2018)
Structural lesson	US\$100,000 reached one third of a superpower's population: the asymmetry is the core fact	

Accuracy note: the 126 million figure is Facebook's estimate of people who "may have been served" at least one IRA story over two years; part of the academic literature considers it an overestimate and notes that only 11% of IRA content was explicitly electoral. The magnitude remains enormous even under the most conservative reading. Individual-exposure studies (e.g., Guess, Nyhan and Tucker in *Nature Communications*) find limited attitudinal and voting effects; the documented damage is ecosystemic: attacking costs almost nothing, reaching 126 million costs US\$100,000, and no democracy can afford to leave that vector open.

Case 2: The Brexit referendum (2016)

The "Leave" campaign combined the **bus lie** ("we send the EU £350 million a week, let's fund our NHS instead"), unauditably micro-targeted messaging, and automated amplification. The UK Electoral Commission fined Leave.EU **£70,000** for funding breaches.

Accuracy note: Cambridge Analytica's role in Brexit is disputed: the firm worked for the campaign, but its own effectiveness was questioned by researchers and journalists. This chapter covers the case on what is verifiable (the bus lie, segmented messaging, the Electoral Commission fine), not on the firm's self-promotional claims.

- **Result:** 51.9% in favour of leaving; turnout of 72.2%.
- **Lasting effect:** the NHS promise was still cited years later as the symbol of electorally profitable disinformation — the "liar's dividend" in action.
- **Lesson:** advertising transparency arrived too late (ad archives only from 2018); a verifiable lie that decided an irreversible referendum.

Case 3: Brazil (2018 and 2022)

2018: days before the runoff, the *Folha de S.Paulo* investigation revealed that conservative businessmen had signed **contracts of up to US\$3 million** with agencies specialised in automated bulk WhatsApp messaging in favour of Jair Bolsonaro — a practice possibly constituting an **illegal campaign contribution**, which WhatsApp could not audit because of its end-to-end encryption.

2022: facing Lula's re-election (50.9%), Brazil's Superior Electoral Court deployed **real-time monitoring** and platforms limited bulk forwarding; even so, electoral fraud narratives penetrated deeply enough to fuel the institutional attacks of 8 January 2023.

- **Lesson 2018:** encrypted messaging is opaque even to the platform — the regulator is late by design.
- **Lesson 2022:** real-time electoral monitoring works on polling day, but does not prevent post-vote delegitimisation. Manipulation no longer ends at the ballot box.

Case 4: India (2019 and 2024)

With **968 million eligible voters** in 2024, India is the world's largest experiment in democracy under generative AI. The difference between 2019 (mass WhatsApp disinformation) and 2024 is qualitative: **affordable personalised deepfakes**, voice cloning of leaders, and industrial-scale synthetic content. The response included the creation of the Misinformation Combat Alliance's **Deepfakes Analysis Unit (DAU)**, a WhatsApp channel for citizens to verify content.

- **Lesson:** at Indian scale, no human fact-checker can keep up; only institutional technological systems can process the volume. This is the scale argument that underpins Part II of this book.

Case 5: European Union (2019 and 2024)

2019: the EU already operated the East StratCom Task Force and EUvsDisinfo, but with tiny resources against the threat. **2024:** the European Parliament elections were held with the DSA in force, 200+ deepfakes documented by the EDMO network, amplified narratives on migration and climate, and turnout at 51.2% (the second consecutive above 50%).


- **Progress:** for the first time there were legally enforceable obligations on platforms, not voluntary codes.
- **Limit:** enforcement remained reactive, national and fragmented — the DSA sanctions conduct, but builds no positive protection capacity. That is precisely the gap the EDSPA comes to fill.

3.2 Metrics of social polarisation

How polarisation is measured

- **Ideological:** distance between the average positions of electorates.
- **Affective:** negative feelings towards voters of the rival party (0-100 thermometer). The most dangerous kind: one does not disagree, one despises.
- **Elites:** hostility between parties and reference media.

The most robust academic evidence on affective polarisation in Europe comes from **Ansgar Reiljan's study (2020, *European Journal of Political Research*)**, which measured the "liking gap" between voters of rival parties across European party systems. His findings: the systems with the highest affective polarisation in Europe are in the **south and east of the continent**, with Spain, Poland and Hungary among the most affected, and affective polarisation can be high even when ideological polarisation is moderate — in other words, *one can hate the rival without the political positions actually being far apart*. Manipulation exploits precisely that gap.



Bubbles and polarisation on social media

Methodological note: the percentages of users "in bubbles" per platform have no verifiable harmonised series; the following table is an **author's illustrative estimate** built on the homogeneous-exposure literature (Ch. 2) and should be read as an order of magnitude, not a measurement.

Platform	% users in bubbles (est.)*	Aggravating feature
Twitter/X	~85%*	Design for public confrontation
YouTube	~80%*	Long-form recommendation loop
Facebook	~78%*	83% of news arrives by algorithmic recommendation (MIT, 2025)
Instagram	~72%*	Passive consumption of curated feed
TikTok	~68%*	Optimisation for pure emotional retention
Common pattern	No platform optimises for the encounter between different positions	

Consequences of polarisation

- **Democratic paralysis:** less room for compromise; minority governments and institutional deadlock.
- **Cross-delegitimisation:** the adversary does not win, "he steals"; the result is not lost, "it is falsified".
- **A market for disinformation:** polarisation increases demand for content that humiliates the rival — it is the fuel of the cycle in Ch. 2.

- **Political violence:** affective hostility is the strongest verbal predictor of violent incidents.


3.3 The trust paradox: Europe responds, the national erodes

The finding that reframes this chapter —and which is often presented incorrectly as a monotonic decline in trust — comes from **Standard Eurobarometer 103 (spring 2025)**:

Trust in the European Union reached its highest level in 18 years in 2025: 52% of Europeans trust the EU and the European Commission, up from 31% in 2013. But only 36% trust their national government and 37% their national parliament.

— Standard Eurobarometer 103, spring 2025

The trust paradox (2025)



Green: trust at 18-year highs. Red/orange: national erosion. Source: Standard Eurobarometer 103 (2025).

Figure 3.3 — The trust paradox. Citizens punish the national level and reinforce the European one.

Why it is an information paradox

Disinformation does not erode uniformly: it strikes where the national narrative has the largest attack surface — governments, parties, national media— while the EU, further from daily media noise and associated with visible responses (vaccines, post-COVID recovery, support for Ukraine), **rises**. The decisive figure for this book: among 15-to-24-year-olds trust in the EU reaches **59%**, and 75% of Europeans feel they are citizens of the EU, the highest level in more than two decades.

The argument that grounds this book

If trust in the EU is at its 18-year high precisely when trust in the national level erodes, then the capacity for cognitive protection **must be European**: citizens are already saying at which level they trust to be protected. The trust paradox is not a statistical curiosity; it is the implicit mandate of European public opinion for an EDSPA. Part II (chapters 6-10 and 9B) answers that mandate; Part III (chapters 11-20) gives it constitutional rank.

Where erosion is real

Concentrated erosion is observed in political parties (systematically the least trusted institution across Europe), in traditional media —increasingly trapped in cross-accusations of bias— and in the electoral process, when post-vote delegitimation takes hold, as the USA 2020 and Brazil 2022 demonstrated.


3.4 The EU's regulatory response: necessary, insufficient

Between 2018 and 2024 the EU built the most advanced digital regulatory framework in the world. It is a real success — and also the proof that regulation alone does not protect.

Instrument	Year	What it does	Maximum sanction	Structural limit
GDPR	2018	Personal data control; curbs wild micro-targeting	4% global turnover	Does not address disinformation
Code of Practice on Disinformation	2018/2022	Self-regulation with compliance verification	None (voluntary)	No legal teeth
EDMO	2020	European network for verification and narrative monitoring	N/A	Coordinates, does not directly protect
DSA	2022/2024	Algorithmic transparency, risk assessments, duties for VLOPs (platforms with more than 45 million EU users)	6% global turnover	Reactive: sanctions conduct already produced
AI Act	2024	Bans subliminal manipulation; mandatory deepfake labelling	7% global turnover or EUR 35M	Gradual application until 2027
Joint balance	Rules without own operational capacity: Europe can fine the platforms, but cannot detect, track and deactivate manipulation campaigns in real time			

The European sanctions arsenal

Maximum sanction on global turnover



Exemplary fines already applied: Meta EUR 1,200M (GDPR, 2023), Amazon EUR 746M (2021), Google EUR 50M (2019).

Figure 3.4 — The sanctioning gradient of the European regulatory architecture.

Verified effectiveness

- **Achievements:** real advertising transparency, mandatory risk assessments for VLOPs, the world's first framework requiring deepfake labelling (AI Act), fact-checker coordination via EDMO.
- **Limitations:** uneven enforcement across states, technology evolving faster than the rule, platforms reluctant to share data, and —most important— **nobody in Europe's job is to actively protect the cognitive sphere:** everyone watches past conduct, no one builds present defence.

The gap that justifies Part II

The DSA punishes those who have already manipulated. The AI Act labels the deepfake that already circulates. The GDPR protects data already collected. The missing piece, which Chapter 6 will present: a **European agency with permanent operational capacity** —early detection, coordinated-campaign tracking, real-time response, structural resilience—. Regulation is the legal skeleton; the EDSPA is the immune system that this diagnosis proves indispensable.

Chapter data and statistics

Verified numerical synthesis

Fact	Value	Source
People reached by the IRA (USA, 2015-2017)	126 million (Facebook estimate)	Senate testimony / 2018 reports
IRA advertising spend	~US\$100,000	US Senate report (2018)
Leave.EU fine	£70,000	UK Electoral Commission
Bulk messaging contracts (Brazil 2018)	up to US\$3 million per contractor	Folha de S.Paulo (2018)
Trust in the EU / Commission (2025)	52% / 52% — 18-year high	Standard Eurobarometer 103
Trust in national government / parliament	36% / 37%	Standard Eurobarometer 103
15-24-year-olds trusting the EU	59%	Standard Eurobarometer 103
Maximum sanctions GDPR / DSA / AI Act	4% / 6% / 7% global turnover	EU regulations

Sources

- US Senate / New Knowledge report (2018). *The Tactics & Tropes of the Internet Research Agency*.
- Facebook, Twitter and Google testimony before the US Senate (2017-2018).
- Folha de S.Paulo (October 2018). Investigation into bulk WhatsApp messaging contracts.
- The Conversation / PBS (2018). *WhatsApp skewed Brazilian election*.
- UK Electoral Commission (2018). Report on the 2016 referendum.
- Reiljan, A. (2020). "Fear and loathing across party lines (also) in Europe: Affective polarisation in European party systems". *European Journal of Political Research*, 59(2).
- Standard Eurobarometer 103 (spring 2025). Trust in institutions.
- European Commission (2022-2025). DSA enforcement and Code of Practice on Disinformation reports.
- EDMO (2024). Monitoring of the European Parliament elections.
- Misinformation Combat Alliance / WEF (2024). Lessons from India's Deepfakes Analysis Unit.

Connection with the rest of the book

This chapter has shown that democratic vulnerability is real, measurable and growing, and that regulation —necessary— is insufficient. **Chapter 4** (The Sovereignty Gap) explains why Europe cannot depend on others for its cognitive protection; **Chapter 5** quantifies the cost of inaction. **Part II** (chapters 6-10 and 9B) designs the institutional response: the EDSPA. **Part III** (chapters 11-20) elevates it to Europe's fourth constitutional power.

Chapter 4: The Sovereignty Gap

Chapter summary


Europe protects its cognitive sphere with infrastructure it does not control. The three US hyperscalers control almost **two thirds of the world cloud market** (Synergy Research, Q3 2025) and European providers barely **~15% of their own market**; the frontier AI models used by European companies and researchers are overwhelmingly American; and in **June 2026** the United States demonstrated that this control is an instrument of foreign policy: an Executive Order on advanced AI and an export control directive issued to Anthropic requiring a licence to transfer its Mythos and Fable models even to foreign persons inside the US. This chapter quantifies that sovereignty gap —infrastructure, data, models, talent— and concludes that without home-grown capacity, European cognitive protection is a borrowed prerogative.

4.1 Technological dependence on foreign powers

Definition: a situation in which a region depends on technology developed, controlled or hosted abroad for critical functions of its society, economy and government. Dependency is not an efficiency problem: it is a problem of *someone else's will*.

The real map of dependency (2026)

World cloud infrastructure market (2026)



AWS, Azure and Google Cloud: ~63% worldwide (Synergy Research, Q3 2025). European providers: ~15% of Europe's own market. Source: Synergy / Holori 2026.

Figure 4.1 — Domination by the three US hyperscalers. The exact figure sometimes cited (85%) is a simplification; the verified reality (~63% worldwide, ~15% European in Europe) remains a structural dependency.

The four layers of dependency


Layer	Verified situation (2026)	European players
Cloud infrastructure	The three US hyperscalers control ~63% worldwide; Europeans ~15% of their own market	OVHcloud, T-Systems, STACKIT, Hetzner, Scaleway
Frontier AI models	European use dominated by US models* (author's illustrative estimate: ~90-95%*)	Mistral AI (valued at ~EUR 11,700M), Cohere-Aleph Alpha, Black Forest Labs, DeepL, Kyutai, AMD Silo AI
Operating systems	Windows, macOS, Android, iOS (all US) cover practically the entire European market*	No alternative at scale
Platform data access	Depends on US companies' APIs, unilaterally cut back (section 4.4)	Partial replacement via DSA Art. 40 (vetted researchers)
Consequence	Every layer of Europe's potential cognitive protection passes through a US corporate or regulatory decision	

The legal precedent: the CLOUD Act (2018)

The US **CLOUD Act** of 2018 allows US authorities to demand data held by US companies **anywhere in the world**. Combined with the EU Court of Justice's **Schrems II** ruling (2020), which invalidated the Privacy Shield precisely because of the risk of US access to European data, the consequence is direct: European citizens' data on US clouds sits under the jurisdiction of a third state. European disinformation-detection infrastructure, if built on US cloud, carries that vulnerability by design.

4.2 Data colonisation: the one-way flow

The concept comes from the academic critique of surveillance capitalism: **Shoshana Zuboff** (*The Age of Surveillance Capitalism*, 2019) described the extraction and monetisation of human experience as raw material; **Nick Couldry and Ulises Mejias** (2019) named it directly: "data colonisation", the process by which data from people in other regions is extracted, processed and monetised without proportional benefit to their communities.



The real size of the European data economy

Verified figures from the **European Data Market Monitoring Tool** (European Commission, 2025) scale what is at stake: the EU27 data market grew from **EUR 82,000M (2023) to EUR 90,000M (2024)**, and the data economy will exceed **EUR 630,000M in 2025 (~4.7% of GDP)**, on the way to EUR 815,000M by 2030. The Draghi report (2024) underlines the contradiction: that value is generated on European data, but is mostly captured by non-European platforms.

4.3 The US AI veto (June 2026)

What this chapter anticipated as a hypothesis happened in June 2026, with proper names. The verified sequence:



What it means for cognitive protection

- **A precedent, not a one-off:** the control architecture already exists; applying it to other models is an administrative decision, not a new law.
- **Disinformation detection is frontier AI:** the best deepfake and coordinated-campaign detection systems are trained with the very class of models now subject to licence. Europe cannot base its cognitive defence on tools whose availability depends on a permit from US Commerce.
- **"Trusted partners" is a hierarchy, not an alliance:** the 26 June exemption is granted by Washington case by case. Sovereignty is not delegated; it is exercised or it does not exist.
- **The semiconductor precedent:** it is the same mechanism already applied to China on chips — now extended to AI and to allies.

The decisive argument of this chapter

The June 2026 veto demonstrated that AI capacity is a **strategic asset under state control**, available to partners as a revocable privilege. *Accuracy note: this is not a one-administration gesture — the first diffusion rule (Biden, January 2025) was rescinded in May 2025; the second (Trump, EO 14409 and the June 2026 IILs) has been in force since: two administrations, two different instruments, the same lever. The bipartisan pattern is the structural fact.* A Europe that wants to protect its cognitive sphere needs compute sovereignty: its own models, its own cloud and its own data. The EDSPA of Part II (chapters 6-10 and 9B) is designed on that triangle; the Protective Power of Part III (chapters 11-20) adds the constitutional guarantee.

4.4 Data access restrictions: opacity as de facto policy

The CrowdTangle shutdown (2024)

On 14 August 2024 Meta shut down CrowdTangle, the tool that allowed journalists and researchers to track disinformation on Facebook and Instagram. Its replacement, Meta Content Library, is **less transparent and accessible**: fewer features on the eleven key topics audited, a slow and restrictive access process, no ability to track the performance of custom lists — precisely in the year of more than sixty elections worldwide (joint investigation by Proof News, the Tow Center and the Algorithmic Transparency Institute, 2024).

The pattern repeats

Platform	Restriction	Consequence for cognitive protection
Meta	CrowdTangle shut down (Aug 2024); selective, slow Content Library access	Partial blindness to disinformation on Facebook/Instagram
X (Twitter)	API closed at prohibitive prices since 2023	Degraded bot and coordinated-network detection tools
TikTok	Minimal data access for researchers	The platform preferred by the young is practically unauditable
Pattern	The transparency platforms granted under pressure (2016-2020) was withdrawn once the pressure dropped. Data access is not a right: it is a revocable concession.	

The European response: DSA Article 40

For the first time, **DSA Article 40** obliges very large platforms to give data access to accredited researchers ("vetted researchers") designated by digital competence centres. It is real progress — and also a demonstration of slowness: the first accesses came years after entry into force, with litigation along the way. Regulation forces sharing; it does not build analytical capacity of one's own. That is the difference between granted transparency and exercised sovereignty.

4.5 The need for European digital sovereignty

Definition: a region's capacity to control its own digital destiny — infrastructure, data, technologies and rules—. Cognitive sovereignty is its top layer: without home-grown infrastructure there is no home-grown detection; without home-grown detection there is no possible protection.

Europe already has assets: supercomputing (EuroHPC)

The EuroHPC Joint Undertaking funds **twelve state-of-the-art supercomputers** across Europe. It is proof that when Europe decides to build capacity, it does — and it is the compute base on which to train sovereign models.

Supercomputer	Country	Power	Year
JUPITER	Germany (Jülich)	~1 exaFLOP — Europe's first exascale system	2025
LUMI	Finland (Kajaani)	~386 petaFLOPS	2022
Leonardo	Italy (Bologna)	~249 petaFLOPS	2022
MareNostrum 5	Spain (Barcelona)	~215 petaFLOPS	2023
Meluxina	Luxembourg	~13 petaFLOPS	2022
EuroHPC network	12 supercomputers	A sovereign compute base for European AI and detection at scale	

The triangle of cognitive sovereignty

- **Home-grown models:** Mistral AI (France, valued at ~EUR 11,700M), Cohere-Aleph Alpha (Germany), Black Forest Labs, DeepL, Kyutai, AMD Silo AI — the European vanguard exists; it needs stable public demand to scale.
- **Home-grown cloud:** OVHcloud, T-Systems, STACKIT, Hetzner — with the push of the sovereign marketplace (Gaia-X / IPCEI-CIS) so that the European public sector buys European.
- **Home-grown data:** common European data spaces (health, mobility, energy) with EUR 2,000M from the Digital Europe programme.

The triangle of cognitive sovereignty

The three vertices of cognitive sovereignty




Figure 4.4 — The triangle of cognitive sovereignty. The EDSPA (Part II) orchestrates the three vertices; the Protective Power (Part III) guarantees them.

The lesson of the chapter

The sovereignty gap is not a lament: it is an inventory of what is missing. Every piece — models, cloud, data, platform access— already has a verifiable European embryo. What is missing is the **institution that orchestrates them with a protection mission**. Chapter 5 closes the diagnosis by quantifying the cost of not having it; Part II (chapters 6-10 and 9B) presents the EDSPA as that orchestration; Part III (chapters 11-20) elevates it to a constitutional Protective Power.

Chapter data and statistics

Fact	Value	Source
World cloud share of the 3 US hyperscalers	~63% (AWS ~30%, Azure ~20%, Google ~13%)	Synergy Research (2026)
European providers' share of the European market	~15%	Holori / market analysis 2026
EU27 data market	EUR 82,000M (2023) → EUR 90,000M (2024)	European Data Market Monitoring Tool (2025)
EU27 data economy	>EUR 630,000M in 2025 (~4.7% of GDP); EUR 815,000M expected by 2030	EDM Monitoring Tool / Draghi (2024)
Mistral AI valuation	~EUR 11,700M (September 2025 round)	Funding reports 2025-2026
CrowdTangle shutdown	14 August 2024	Meta Transparency Center / CJR
ILL to Anthropic (Mythos and Fable)	12 June 2026; "trusted partners" exemption 26 June	US Department of Commerce (2026)
EuroHPC supercomputers	12, including JUPITER (~1 exaFLOP, Europe's first)	EuroHPC JU (2026)

Sources

- Synergy Research Group (2026). Quarterly cloud infrastructure tracker, Q1 2026.
- Holori (2026). *Cloud Market Share 2026*; European market analysis.

- European Commission (2025). *European Data Market Monitoring Tool — First Report on Facts and Figures*.
- Draghi, M. (2024). *The Future of European Competitiveness*.
- US Executive Order on advanced AI (2 June 2026) and IIL directives of the Department of Commerce (12 and 26 June 2026), documented by Mayer Brown (2026).
- Zuboff, S. (2019). *The Age of Surveillance Capitalism*. PublicAffairs.
- Couldry, N. & Mejias, U. (2019). *The Costs of Connection*. Stanford University Press.
- CJR / Proof News / Tow Center / Algorithmic Transparency Institute (2024). Audit of the CrowdTangle shutdown and Meta Content Library.
- EuroHPC Joint Undertaking (2026). Supercomputer catalogue.
- Court of Justice of the EU (2020). Schrems II judgment (C-311/18); US CLOUD Act (2018).

Connection with the rest of the book

Part I has completed the diagnosis: the threat (Ch. 1-3) and now the impossibility of defending oneself with borrowed tools. **Chapter 5** quantifies the cost of inaction and closes Part I. **Part II** (chapters 6-10 and 9B) designs the EDSPA —the agency that turns compute sovereignty into operational cognitive protection— and **Part III** (chapters 11-20) institutionalises it as Europe's fourth power.

Chapter 5: The Urgency of Protection


Synthesis of the diagnosis: the threats reinforce each other, current mechanisms are fragmented and underfunded, and every year without home-grown capacity has a measurable cost

Chapter summary

This chapter closes Part I. It synthesises the verified threats of chapters 1-4, inventories the existing protection mechanisms with their **real budgets** —and the comparison is eloquent: EDMO's central unit operates on **EUR 2.5M** and East StratCom on some **EUR 11M with 40 experts**, against Frontex with **~EUR 1,100M a year** for the physical border—, analyses the EU's institutional bottlenecks and makes the case for dedicated protection: the EDSPA, with enhanced cooperation as the launch path. The conclusion: no European institution has, as its core mission today, protecting Europeans' cognitive sphere.

5.1 Synthesis of the threats: the perfect storm

Chapters 1-4 did not describe four separate problems, but an **interdependent system of threats**. Each component amplifies the others:




The diagnosis in verified figures

Chapter	Key finding	Verified figure	Source
Ch. 1	The cognitive sphere is the new battlefield	116,175 layoffs attributed to AI in the US (Jan-Aug 2026)	2026 employment reports / Ch. 5 of the book
Ch. 2	Falsehood spreads faster than truth	6x faster; 70% more retweets	Vosoughi et al., <i>Science</i> 2018 (MIT)
Ch. 2	Humans do not detect deepfakes	55.5% accuracy; 0.1% identify all of them	Meta-analysis of 56 studies 2024 / iProov 2025
Ch. 3	Interference is cheap and massive	US\$100,000 reached 126 million people (US 2016)	Facebook testimony to the Senate
Ch. 3	Trust erodes at the national level, not the European one	EU 52% (18-year high) vs national government 36%	Eurobarometer 103 (2025)
Ch. 4	Frontier AI is a US-controlled asset	IIL to Anthropic (Mythos/Fable), June 2026	US Department of Commerce
Ch. 4	Defence infrastructure is not European	~63% of world cloud with 3 US hyperscalers; ~15% European in Europe	Synergy Research Q3 2025
Balance	A transnational, operational and cheap threat. A national, fragmented and very expensive defence.		

5.2 Current mechanisms and their real budgets

Europe is not unarmed: it has advanced regulation (Ch. 3) and monitoring institutions. What it does not have is **scale**. The verified budget comparison shows it without mercy:



Inventory of mechanisms and limits

Mechanism	Function	Verified resources	Structural limit
EDMO (2020)	Fact-checking network: 15 hubs across the EU, academia + fact-checkers	Central: EUR 2.5M; hubs in tranches of EUR 8-11M	Coordinates and observes; no operational capacity or powers of action
East StratCom / EUvsDisinfo (2015)	Monitoring of pro-Russian disinformation; database with thousands of cases	~EUR 11M and ~40 experts (2022)	Exclusive focus on the Russian vector; no mandate over domestic or commercial manipulation
ENISA (2004)	EU cybersecurity agency	EUR 26.9M a year (2025) + EUR 36M Cyber Reserve	Cybersecurity is not cognitive protection: does not address disinformation or manipulation
Viginum (France, 2021)	Surveillance of foreign online interference	National scale	France only; a European agency cannot rest on the capacities of a single state
Common gap	No institution has early detection + investigation + real-time response as its core mission, at European scale		

The seven structural limitations

- **Fragmentation:** EDMO, East StratCom, ENISA, national agencies and electoral units work in parallel, with no command centre.
- **Underfunding:** tens of millions against a manipulation industry operating at the scale of platforms with billions of users.
- **Reaction, not prevention:** debunkings arrive when disinformation has already completed its cycle — which is 6 times faster than the truth (Ch. 2).
- **Platform dependency:** the data needed to detect campaigns is a revocable concession of US companies (Ch. 4).
- **Lack of specialisation:** cognitive protection requires profiles that exist in no European organisational chart today: coordinated-campaign analysts, deepfake detection engineers, free-expression lawyers working together.
- **Digital borders vs national responses:** a campaign is born in one country, amplified from another and strikes a third; competences stop at every border.
- **Technological pace:** the DSA was designed before generative AI; the AI Act will be applied gradually until 2027. The law always arrives after the tool.

5.3 Institutional bottlenecks

Suppose the EU decided tomorrow to create the missing capacity. How long would it take? The verified precedents:

Instrument	From proposal to application	Lesson
DSA	Proposal Dec. 2020 → full application 2024: ~4 years	The fastest possible regulatory pace on a "hot" topic
AI Act	Proposal Apr. 2021 → approval 2024 → gradual application until 2027: ~6 years	Phased application adds years to years
Implication	An agency created by the ordinary route would be operational around 2030-2032. The threat does not wait: the AI veto already happened (June 2026), and the 2029 elections will already be under advanced generative AI	


The four types of bottleneck

- **Political:** unanimity in the Council for sensitive matters; diverging national interests (regulatory France, self-regulatory Germany, national-sovereignty overlays). Enhanced cooperation (TFEU art. 326-334) exists precisely to bypass this blockage.
- **Legal:** the ordinary legislative procedure takes years; education and culture competences are national, blocking large-scale European media-literacy programmes.

- **Administrative:** hiring AI talent in an EU organigram requires procedures that the private sector solves in weeks; the best profiles choose other destinations.
- **Technical:** scarcity of deepfake-detection and campaign-analysis experts; insufficient home-grown infrastructure (Ch. 4) exactly when detection becomes compute-intensive.

5.4 The European Democracy Shield (2025): the official diagnosis and its limit

On 12 November 2025 the European Commission and the High Representative published the joint communication "The European Democracy Shield: Empowering Strong and Resilient Democracies". It is the most ambitious institutional response to date to the diagnosis of this Part I — and a close reading confirms exactly the thesis of this book.



The verified critique: coordinating is not protecting

The **Jacques Delors Centre** analysis (December 2025) — "papering over the cracks" — sums up the Shield's structural limit:

- The new **European Centre for Democratic Resilience** is a **coordination structure based on member states' voluntary participation**: it gathers existing expertise, it does not build new capacity.
- It suspiciously resembles the **EEAS Rapid Alert System** (2019), set up for that year's European elections — and which "**never really took off**", according to the 2025 retrospective balance.
- Verified risk of **duplicating structures** without clarity on what the Centre adds to what exists (EDMO, East StratCom, national units).
- The Shield rests on proposed budgets of the next financial framework (CERV+ EUR 3,600M, AgoraEU EUR 8,600M) that still require negotiation and mostly deploy **grants to third parties**, not home-grown operational capacity.


The limit that grounds Part II

The European Democracy Shield is institutional proof that the EU **already accepts this diagnosis**: there is a threat, it is transnational, and it requires a European response. What the Shield cannot give —by design— is what the Rapid Alert System could not give in 2019 either: **permanent operational capacity**. Home-grown early detection. Home-grown investigation. Home-grown real-time response. A centre that voluntarily coordinates those that already exist does not close the gap identified in 5.2: it is its bureaucratic continuation. The EDSPA of Part II is the next logical step — and this chapter argues it with the EU's own official sources.

5.5 The case for dedicated protection: the EDSPA

The ladder of the European response

From observing to protecting: the steps of the European response



Each step added scope but none closed the gap: nobody "operates". The ladder has one more step — and this book designs it.

Figure 5.4 — The evolution of the European response: from monitoring to coordinating. Operating is missing.

Why cognitive protection needs a dedicated institution

Critical sector	Threat	Specialised institution	Scale of resources
External security	War, terrorism	NATO, national agencies	Trillions of euros accumulated
Physical borders	Trafficking, irregular migration	Frontex (~EUR 1,100M/year; proposal EUR 11,900M 2028-2034)	Hundreds of millions a year
Cybersecurity	Cyberattacks	ENISA (EUR 26.9M/year) + national CERTs	Tens of millions
Cognitive protection	Digital mass manipulation	None with operational capacity (only observation and coordination)	Units of EUR 2.5-11M

The reference budget: an operational cognitive-protection agency would require, in this author's illustrative estimate, **EUR 500-1,000M a year** — the order of magnitude of Frontex for the physical border (~EUR 1,100M verified in 2025). It is the scale at which Europe already protects a border under attack; the cognitive border deserves at least the same budgetary logic. (Author's illustrative estimate; the detailed cost and return calculations for the EDSPA are developed in chapters 9 and 9B.)

How to launch it: enhanced cooperation

The bottleneck analysis (5.3) advises against waiting for the consensus of all 27. **Enhanced cooperation** (TFEU, art. 326-334) allows a group of member states —at least nine— to create the agency without waiting for everyone, as was done for the euro, the patent package or the European divorce. It is the route Part II proposes for the EDSPA's launch, with the door open for the rest to join later.

The five arguments for urgency

- **The threat grows faster than the response:** generative AI makes manipulation cheaper every year; the law needs years (5.3).
- **The cost of inaction is real:** an EU27 data economy of EUR 630,000M (4.7% of GDP, Ch. 4) built on ground Europe does not control; national erosion of trust (Ch. 3); verified post-electoral delegitimisation (USA 2020, Brazil 2022).
- **Existing solutions have been tried and do not suffice:** Rapid Alert System (2019, failed), EDMO (EUR 2.5M, observation), 2025 Shield (voluntary coordination) — certified by the Delors Centre itself.

- **Europe has the capacity:** EuroHPC supercomputing (JUPITER exascale), Mistral AI, Cohere-Aleph Alpha, emerging European cloud (Ch. 4) — what is missing is the institution that orchestrates it.
- **The political moment is the best in years:** the June 2026 AI veto created awareness of dependency; the 2025 Shield put the topic at the centre of the agenda; the proposed 2028 financial framework already includes democratic programmes worth tens of billions. The window is open — and windows close.

Chapter data and statistics

Fact	Value	Source
EDMO (central unit)	EUR 2.5M funding	European Commission (Digital Europe)
East StratCom Task Force	~EUR 11M and ~40 experts (2022)	EEAS / parliamentary documentation
ENISA (cybersecurity)	EUR 26.9M a year (2025) + EUR 36M Cyber Reserve	ENISA (2025)
Frontex (physical border)	~EUR 1,100M (2025); proposal EUR 11,900M 2028-2034	Frontex / European Commission (2025)
European Democracy Shield	Joint Communication 12-Nov-2025; European Centre for Democratic Resilience (voluntary coordination)	European Commission / EEAS (2025)
Rapid Alert System	2019; "never really took off" per the retrospective balance	RFE/RL, Delors Centre (2025)
Proposed budget for the EDSPA	EUR 500-1,000M/year (author's illustrative estimate, anchor: Frontex)	This book (detailed in chapters 9 and 9B)
Verified duration of the ordinary legislative route	DSA ~4 years; AI Act ~6 years until full application	EU legislative process

Sources

- European Commission / High Representative (12 November 2025). *The European Democracy Shield: Empowering Strong and Resilient Democracies* — Joint Communication.
- EPRS — European Parliamentary Research Service (2026). *The European Democracy Shield: An Overview*.
- Jacques Delors Centre (December 2025). *The European Democracy Shield — Papering over the Cracks*.
- European Commission. EDMO funding and its network of 15 hubs (CEF / Digital Europe).
- EEAS / German parliamentary documentation (2022). East StratCom Task Force and EUvsDisinfo data.
- ENISA (2025). Annual budget and EU Cybersecurity Reserve (EUR 36M).
- Frontex (2024-2025). Voted budget 2025; European Commission, DG HOME (Jul. 2025): proposal of EUR 11,900M for 2028-2034 (Statewatch cites EUR 11,200M).
- Government of France. Decree no. 2021-922 of 13 July 2021 (creation of Viginum).
- RFE/RL / Globalsecurity (November 2025). Rapid Alert System balance and Shield coverage.
- Vosoughi, S., Roy, D. & Aral, S. (2018). *Science* — basis of the threat synthesis (chapters 2 and 1-4).

Connection with the rest of the book

Part I is closed: the threat (chapters 1-3), the impossibility of defending oneself with borrowed tools (chapter 4) and the proven insufficiency of current mechanisms —including the 2025 Democracy Shield— (chapter 5). **Part II** (chapters 6-10 and 9B) builds the solution: the theory of the Protective Power (chapter 6), cognitive protection as a right (chapter 7), the Protector's mandate (chapter 8) and the EDSPA with its ROI and its economic and social impact (chapters 9, 9B, 10). **Part III** (chapters 11-20) elevates it to Europe's fourth constitutional power: the Protective Power.

Chapter 6: EDSPA — Foundations of the European Digital Sphere Protection Agency

Chapter summary

This chapter opens Part II and establishes the **foundations of the EDSPA** (European Digital Sphere Protection Agency), the agency that will operationalise the Protective Power. It presents its mission and values, the **legal basis** that allows its creation without treaty reform (internal market, data protection and, above all, enhanced cooperation among a minimum of **nine states**), a **founding core of nine Member States** selected with verified 2025 data, the proposal of a **seat in Brussels** following the precedent of the EMA and EBA relocations, a **three-scenario financing structure** (EUR 500M, EUR 1,000M and an accelerated scenario of EUR 2,500M consistent with Chapter 9) and the **fundamental principles** that will govern it. The comparison is telling: the proposed EDSPA would operate at the same order of magnitude as Frontex (~EUR 1,100M in 2025) — for the cognitive border.

6.1 What is the EDSPA?

Definition and mission

EDSPA — *European Digital Sphere Protection Agency*. The agency protects the digital cognitive sphere of European citizens, institutions and societies against manipulation, disinformation and algorithmic threats.

"To protect the digital cognitive sphere of European citizens, institutions and societies against manipulation, disinformation and algorithmic threats, guaranteeing the digital sovereignty, cognitive autonomy and democratic resilience of the European Union."

Proposed mission of the EDSPA

Its vision: to become the **global reference in cognitive protection**, with a model replicable by other democracies, based on transparency, ethics and respect for fundamental rights.

Strategic objectives

Objective	Content
1. Cognitive protection	Detect and neutralise disinformation and manipulation campaigns; protect against unethical algorithmic influence; guarantee the integrity of democratic processes
2. Digital sovereignty	Reduce technological dependence; promote European technology in AI and digital protection (Mistral AI, valued at ~EUR 11,700M in 2026, as a reference asset); ensure European control over data and critical infrastructure
3. Democratic resilience	Strengthen the EU's capacity against external interference; protect freedom of expression and pluralism; guarantee transparency and accountability of platforms
4. International cooperation	Establish global standards for cognitive protection; collaborate with other democracies; export the model to allied countries

Core values

Value	Application at the EDSPA
Transparency	Public reports, external audits, open data (where safe)
Independence	Merit-based recruitment, limited mandates, diversified funding
Excellence	The best experts, continuous training, rigorous evaluation
Ethics	Strict compliance with the GDPR, the Charter of Fundamental Rights and ethical AI principles
Collaboration	Coordination with Member States, platforms, civil society and academia
Innovation	R&D investment, adoption of ethical AI, own tools
Accountability	Supervision mechanisms, independent evaluations, transparent decisions

6.2 Legal basis: creating the EDSPA without treaty reform

The EDSPA can be created **without any reform of the EU Treaties**, by combining four existing legal bases:

1. Internal market (Article 114 TFEU)

The fragmentation of cognitive protection among Member States **distorts the digital single market**: disinformation erodes trust in the common digital market and the absence of common standards creates barriers to trade and innovation. Article 114 allows the adoption of **binding regulations** for the whole EU — the same route used for the GDPR, the DSA and the AI Act.

2. Personal data protection (Article 16 TFEU)

Cognitive protection is intimately linked to data protection: algorithmic manipulation relies on the unethical use of personal data. It is a **well-established, proven** legal basis, articulated with the GDPR and the full application of the AI Act (synthetic-content transparency obligations applicable from August 2026).

3. Enhanced cooperation (Articles 20 TEU and 326-334 TFEU)

The starting route. Article 20 TEU allows a group of Member States to **move forward without consensus among the 27**, using the Union's institutions. The verified minimum requirement: enhanced cooperation must involve at least **nine Member States**. It is ideal for an initial phase with a founding core and can expand gradually. The precedent exists: enhanced cooperation has already been used for divorce law, the unitary patent and the protection of the euro against counterfeiting. It avoids blocking by reluctant Member States.

4. Supporting competences (Article 6 TFEU)

The EU may carry out supporting, coordinating or complementary actions in areas such as education, culture and consumer protection. Media literacy and protection of the "cognitive consumer" fit this framework without intruding on national competences.

Creation process and calendar



Legal structure: a decentralised EU agency (similar to ENISA, EMA or Frontex), with its own legal personality and administrative and financial autonomy, created by a regulation of the European Parliament and of the Council, with internal rules, management board decisions and agreements with Member States for national implementation.

Why the legal basis is credible

The February 2025 AI Action Summit and the **InvestAI** programme (mobilising EUR 200,000M for AI in Europe, including EUR 20,000M for AI gigafactories) showed that the EU can activate large-scale financial instruments within months. Creating an agency by regulation is a mature institutional technique: the EU already has more than 30 decentralised agencies. The EDSPA would invent nothing in institutional law; it would apply existing law to a new domain.

The institutional demand already exists

Before any proposal in this book, the Commission itself has been building — with fragile instruments — the function the EDSPA would consolidate: four consecutive rounds of the DG CNECT «European Narratives Observatory» line (NODES 2022; 2nd call 2023 on Ukraine, gender disinformation and elections; 3rd call 2025 «Detecting influence campaigns and boosting societal resilience», two projects of EUR 1.5M and EUR 1.65M; PPPA-2026 social cohesion pilot, EUR 785,000), the SHIFT consortium (starting April 2026, 18 months, 13 countries) and the Democracy Shield's European Centre for Democratic Resilience (intergovernmental, without authority, operational since 24 Feb 2026). The Commission admits it needs to anticipate, detect and respond at scale — but builds it without power, with projects that expire. The EDSPA is the logical next step: the same mandate, with independence, permanence and competences.

6.3 The founding core: nine Member States

Enhanced cooperation requires a minimum of **nine Member States**. We propose exactly that initial core: a group with critical economic, technological and political mass, selected against verifiable criteria.


Selection criteria

- **Political commitment:** clear governmental will and support for the EDSPA's values

- **Technical capacity:** digital infrastructure, supercomputing (EuroHPC), AI and cybersecurity expertise
- **Financial capacity:** GDP-proportional budget contribution and funding of national nodes
- **Compatible legal framework:** applicable national legislation and guaranteed fundamental-rights protection

The core of Nine (verified 2025 data)

Country	GDP 2025	Population	Key assets	Rationale
Germany	~EUR 4.5 trillion (~24% of EU GDP)	84M	JUPITER (exascale supercomputer, Juelich); GAIA-X	Largest EU economy; industrial and research critical mass
France	~EUR 3.0 trillion (~16%)	68M	Mistral AI (~EUR 11,700M valuation); EUR 109,000M pledged to AI (AI Action Summit 2025)	European leader in frontier AI and sovereign compute
Italy	~EUR 2.3 trillion (~12%)	59M	Leonardo; Bologna HPC hub	Third-largest EU economy; political commitment
Spain	~EUR 1.7 trillion (~9%)	47M	BSC / MareNostrum 5 (Barcelona)	Supercomputing leader; commitment to cognitive protection
Netherlands	~EUR 1.2 trillion (~6.3%)	17M	AMS-IX, HPC ecosystem; seat of Europol and EMA	Europe's digital hub; EMA seat since 2019
Belgium	~EUR 0.45 trillion	11M	Seat of the EU institutions; dark fibre and data centres	Natural candidate for the headquarters
Sweden	~EUR 0.39 trillion	10M	Nordic tech ecosystem; sustainable data nodes	Leader in digital innovation and democratic values
Finland	~EUR 0.23 trillion	5.5M	LUMI (one of Europe's most powerful supercomputers, Kajaani)	Reference in digital ethics; Finnish election-infrastructure experience
Denmark	~EUR 0.36 trillion	5.8M	Advanced public-sector digitalisation; green data centres	Leader in digitalisation and data protection
Core of Nine total	~EUR 14.1 trillion (~75% of EU GDP of ~EUR 18.8 trillion)	~307M (~68% of EU population)	3 of the major EuroHPC supercomputers	Sufficient critical mass to launch enhanced cooperation



Accession process for new Member States

1. **Expression of interest** (3 months): formal letter from the government to the management board and initial eligibility screening
2. **Technical evaluation** (6 months): capacity audit, legal compatibility, financial analysis
3. **Negotiation** (3-6 months): accession agreement, implementation plan, financial commitments
4. **Approval** (2 months): management board recommendation and approval by qualified majority
5. **Implementation** (6-12 months): national node, recruitment, systems integration, training

Estimated total time: 12-24 months from expression of interest to full operation. Expansion targets: **20 Member States by 2028, all 27 by 2030.**

6.4 Headquarters and infrastructure

The institutional precedent

The EU has a proven procedure for deciding seats: after Brexit, the EU-27 voted in 2017 to relocate the **EMA to Amsterdam** and the **EBA to Paris**, amending the founding regulations accordingly. The seat of a new agency is decided by the Member States, not by the Commission — and candidacies are negotiated with host-country incentive offers.

Criteria and comparative evaluation

Criterion	Germany	France	Spain	Belgium	Netherlands	Sweden	Finland
Geographic centrality	9	8	7	10	9	6	5
Digital infrastructure	10	9	8	9	10	9	8
Technical capacity	10	9	8	8	9	8	7
Political stability	9	8	8	9	9	10	10
Costs and incentives	7	6	9	8	7	6	7
Total (max 50)	45	40	40	44	44	39	37

Recommendation: Brussels (Belgium), with Amsterdam as a strong alternative. Brussels offers proximity to the Commission, the Council and the European Parliament (plus NATO and the OECD), first-class connectivity and a naturally multilingual environment. The final decision rests, as in the EMA/EBA precedent, with the Member States.

Headquarters infrastructure


- **Main building:** ~50,000 m² — offices for 1,000-1,500 employees, on-site data centre (10,000 m²), AI and analysis laboratories (5,000 m²), 24/7 crisis room (1,000 m²), auditorium and common areas
- **Compute:** the EDSPA does not need to build supercomputing from scratch: it can rely on **EuroHPC** (JUPITER, LUMI, MareNostrum 5) with dedicated capacity, plus its own GPU cluster for real-time analysis
- **Main data centre** (~EUR 500M): co-location in Belgium or the Netherlands, 100% renewable energy, N+2 redundancy, ISO 27001 certification and ENISA standards
- **Backup data centre** (~EUR 450M) in Amsterdam/Frankfurt + own network (~EUR 100M) with GEANT connection

Methodological note: the data-centre and network figures (~EUR 850M total) are **the author's illustrative estimates** for mission-critical infrastructure with geographic redundancy; they do not come from any official budget. The realistic anchor: JUPITER, Europe's first exascale supercomputer, cost ~EUR 500M, confirming the order of magnitude. The EuroHPC option (dedicated capacity instead of owned infrastructure) could substantially reduce the initial investment.

6.5 Financing: three consistent scenarios

Chapter 5 set the reference range for the EDSPA — **EUR 500-1,000M per year**, an illustrative estimate anchored to Frontex (~EUR 1,100M in 2025 for the physical border) — and Chapter 9 adds an **accelerated scenario of EUR**

2,500M with 2,000 employees. This chapter structures that range into scenarios with verifiable sources:



Scenario	Annual budget	Staff	Dominant financing route	Reference
Base	EUR 500M	~1,000-1,500	Enhanced cooperation of the Nine + EU programmes (Horizon Europe EUR 93,500M / Digital Europe EUR 8,100M, 2021-2027)	Anchor: midpoint of ENISA-Europol-EMA-Frontex
Full	EUR 1,000M	~3,000-5,000	Contributions from the 27 (Frontex reference ~EUR 1,100M for a critical mission)	Anchor: Frontex 2025
Accelerated	EUR 2,500M	~2,000 employees (HQ) + national nodes	2026-2027 start-up with InvestAI-type instruments (EUR 200,000M mobilised for AI)	Consistent with Chapter 9


All three scenarios are the author's illustrative estimates, anchored to verified real budgets of EU agencies

Source structure (full scenario, EUR 1,000M)

Source	Mechanism	Annual contribution	%
EU budget	Horizon Europe, Digital Europe, CERV+ (EUR 3,600M, Democracy Shield)	EUR 400M	40%
Member States	GDP-proportional contribution (~0.003-0.005% of the EU-27 GDP)	EUR 350M	35%
Platform fines	Earmarked DSA sanctions (up to 6% of worldwide turnover) and AI Act (up to 7% / EUR 35M)	EUR 150M	15%
Other instruments	Complementary national digital taxes, public and private funds	EUR 100M	10%
Total		EUR 1,000M	100%

Financing structure (full scenario, EUR 1,000M)

Where would the EUR 1,000M annual budget come from?



Diversified funding = guarantee of independence (section 6.6): no single source exceeds 40%

Figure 6.4 — Source structure of the full scenario. Author's elaboration; illustrative estimate anchored to verified EU programmes.

Initial investment (2027-2028)

Item	Estimated cost	Financing
Headquarters (Brussels) and equipment	EUR 200M	EU (50%) / founders (50%)
Dedicated EuroHPC capacity + sovereign European cloud (multiannual contracts)	EUR 250M	EU (50%) / founders (50%)
Own network + GEANT and national node deployment	EUR 100M	EU (100%)
Staff (initial deployment, 1,500 employees, Ch. 7)	EUR 150M	EU (40%) / founders (60%)
Own detection and auditing tools (Ch. 8)	EUR 200M	EU and investment funds
Training and national capabilities	EUR 50M	EU (100%)
Media-literacy and communication launch	EUR 70M	EU (100%)
Contingency reserve	EUR 80M	All sources
Total initial investment	EUR 1,100M	

Consistency with Chapters 9 and 9B: this breakdown is the baseline scenario (EUR 1,100M, 2027-2028), without owned data centres — dedicated EuroHPC capacity + European cloud, per the Part II cost consensus. The accelerated scenario (EUR 2,500M, 2026-2027) brings the same model forward (Ch. 9). Owned data-centre infrastructure is revisited as a future evolution of the Protective Power (Ch. 17).

6.6 Fundamental principles of the EDSPA

1. Absolute independence

Merit-based recruitment with transparent processes; executive mandates of 4-8 years, not automatically renewable; diversified funding (section 6.5) to avoid dependencies; veto on conflicts of interest with platforms, foreign governments or corporate interests, with public declaration of interests.

2. Technical excellence

Area	Profile	Annual salary range (EUR)
AI and machine learning	Deep learning, NLP, computer vision	120,000-180,000
Cybersecurity	Intrusion detection, forensic analysis	110,000-160,000
Data analysis	Big data, statistics, visualisation	100,000-150,000
Digital law	GDPR, DSA, AI Act, intellectual property	100,000-140,000
Disinformation	Narrative analysis, fact-checking	90,000-130,000
AI ethics	Social impact, bias, transparency	90,000-120,000

Reference ranges of the European tech market for senior profiles; the EDSPA would complement them with stability, a public-interest mission and continuous training — the same balance that allows Europol (~1,400 staff) or the EMA to attract talent.

3. Radical transparency

Regular public reports; external audits (European Court of Auditors); aggregated and anonymised open data; public consultations. Limits: ongoing operations, privacy (GDPR) and trade secrets.

4. International cooperation

Bilateral agreements with democracies (United Kingdom, Canada, Australia, Japan, South Korea); participation in the UN, OECD and NATO; exporting the model with global standards. Non-negotiable limit: **absolute prohibition on sharing information or technology with non-democratic regimes**, and European control over critical decisions.

5. Accountability

KPI	2027 target	2030 target
Disinformation detection time	< 6 hours	< 2 hours
Verification accuracy rate	> 90%	> 95%
Campaigns neutralised (per year)	500	2,000
Member State satisfaction	> 80%	> 90%
Citizen satisfaction	> 70%	> 80%
Budget executed	> 90%	> 95%

6.7 Conclusion: the EDSPA as a strategic step

Part I demonstrated that the cognitive threat is real, growing and cross-cutting, and that current mechanisms — EDMO with EUR 2.5M, East StratCom with ~EUR 11M and 40 experts, the Democracy Shield without operational capacity of its own — are two orders of magnitude below what the problem requires. This chapter has shown that the institutional response is **viable without treaty reform** (enhanced cooperation among nine states), **financeable** (the EUR 500-1,000M range is the order of magnitude the EU already spends on Frontex), and **governable** with principles of independence, excellence, transparency and accountability.

The US AI veto of June 2026 (Chapter 4) proved that dependence is a vulnerability; France's national capability (Mistral AI, EUR 109,000M pledged) and the founding core's EuroHPC infrastructure prove that Europe has the assets. What is missing is not technology or money: it is the **institution** that orchestrates them. The EDSPA is that institution.

Next steps: Chapter 7 will develop the full organisational structure (management board, organogram, technical and administrative departments, recruitment processes) and Chapter 8 operations and processes (detection, analysis, response, human-plus-AI double verification, national nodes, evaluation).

Verified sources of this chapter

- Eurostat (2025): EU GDP ~EUR 18.8 trillion; Germany ~24%, France ~16%, Italy ~12%, Spain ~9%, Netherlands ~6.3%, Poland ~5%
- ENISA: 2025 budget of EUR 26.9M; additional EUR 36M for the EU Cybersecurity Reserve (2025)
- Europol: 2024 budget of EUR 222.4M; ~1,400 staff (The Hague)
- EMA: 2025 budget above EUR 600M (new fee regulation); seated in Amsterdam since 2019
- Frontex: ~EUR 1,100M in 2025 (Chapter 5); standing corps planned at 10,000 officers by 2027
- Article 20 TEU: enhanced cooperation among a minimum of nine Member States
- European Commission: relocation of the EMA (Amsterdam) and the EBA (Paris), EU-27 procedure, 2017-2018
- DSA: fines up to 6% of annual worldwide turnover; AI Act: up to 7% or EUR 35M for prohibited practices
- Horizon Europe: EUR 93,500M (2021-2027); Digital Europe Programme: EUR 8,100M (2021-2027)
- InvestAI (European Commission, Feb 2025): EUR 200,000M mobilised for AI, including EUR 20,000M for gigafactories; France: EUR 109,000M pledged (AI Action Summit, Feb 2025)
- JUPITER (exascale EuroHPC supercomputer, Juelich, ~EUR 500M), LUMI (Kajaani), MareNostrum 5 (BSC, Barcelona)

Chapter 7: Organisational Structure of the EDSPA

Chapter summary

The EDSPA cannot be just another EU agency. This chapter designs its organisational structure on the principle of **minimal state interference at the central headquarters**: a **Council of Protectors** of 17 independent experts selected on merit — with not a single state representative — following the model that the **ERC Scientific Council** has proven for almost two decades, with the functional and financial independence that the **ECB** brings to European institutional law. Member States do not disappear from the map: they join a non-voting **Advisory Council** at headquarters and a coordination **National Chamber** in the territorial nodes, under a dual model (eu-LISA + EDPB) that reserves all operational direction to the centre. The chapter details the executive directorate, the six departments, the staff and budget distribution of the base scenario (EUR 500M, 1,500 employees) and a system of **frictionless safeguards**: accountability counterweights designed so that neither states nor corporate litigants can slow down the operation.

7.1 The thesis: cognitive protection = monetary policy, not moat politics

Why can't a cognitive protection agency be governed like Frontex or Europol? Because its object of work is **the minds of citizens** — the very terrain where the interference battles documented in Part I are fought. An agency whose board consists of one representative per state is an agency where every government of the day — including those that today sponsor half-truths for their own electoral benefit — holds a seat, a voice and a vote over who protects citizens and against whom.

"Nobody proposes that national governments appoint the members of the ECB's board. Central bank independence is considered a condition of its effectiveness. Cognitive protection deserves exactly the same architecture: whoever guards the public sphere of ideas cannot depend on the political power that operates within it."

Central thesis of the chapter

The verified risk: the EU has already experienced what happens when political power meddles in a technical mandate body. In 2022, the Council and Parliament amended the Europol Regulation to **retroactively legalise** mass data-retention practices that the EDPS had ordered stopped — and the European supervisor's attempt to annul those provisions was declared inadmissible (case T-578/22). The lesson for the EDSPA: independence that is not **locked into the founding regulation** can be dismantled in a single budget negotiation round.

The independence gradient of EU bodies

European bodies form a measurable gradient of state interference. Standard agencies (ENISA, Europol, EUIPO, eu-LISA) have boards with one representative per Member State plus the Commission. The EDPB is composed of national authorities that are *independent* — regulators, not governments. And at the maximum-independence end sit the **ERC** (a Scientific Council of ~22 merit-selected experts, zero state representation, the finest excellence reputation in the EU), the **ECB** (five pillars of institutional independence) and the **EDPS** (a single supervisor, no state board). The proposed EDSPA deliberately places itself at that end.

Where does each body sit?




Figure 7.1 — The state-interference gradient across EU bodies. Author's elaboration on the verified composition of each governing board (ENISA, Europol, EUIPO, eu-LISA, Frontex, EDPB, ECB, EDPS, ERC).

7.2 The governing body: the Council of Protectors

Composition (ERC model)

The **Council of Protectors** is the EDSPA's governing body: **17 independent experts**, selected exclusively on merit, **with no state representation whatsoever**. No government appoints, vetoes or removes its members.

Position	Number	Selection	Mandate
Merit members	17	EU-wide open call + independent selection committee (verified Europol-2026-type procedure: assessment, interviews, shortlist of at least 3 candidates) + confirmation by the Council itself	6 years, non-renewable, staggered renewal by thirds every 2 years
— of whom, chair	1	Elected by and from the 17 members	2 years, renewable once within the mandate
European Commission	1 observer	No vote: institutional and budgetary coordination	—
European Parliament	1 observer	No vote: voice of citizen interest	—
Member States	Advisory Council (no vote)	One representative per state, with a voice on national nodes and territorial priorities	—
Members with voting rights	17 — 0 states	All strategic, budgetary and leadership decisions	

Required profiles and prohibitions

- **Multidisciplinary composition guaranteed in the regulation:** at least 3 experts in AI/algorithms, 2 in cybersecurity, 3 in fundamental and digital law, 2 in social/behavioural sciences, 2 in journalism/media, 2 in electoral security, 1 in ethics, 2 open seats for an evolving threat landscape
- **Conflict-of-interest barriers** (stricter than any existing agency): total veto on ties to platforms, foreign governments or corporate interests; public declaration of interests; **2-year cooling-off period** for former platform employees or former national political office-holders with digital portfolios
- **Mandatory diversity:** gender balance (maximum 10/17 of any one gender) and geographic balance (maximum 2 members per state, with no national quotas: nationality is a data point, not a criterion)

Functions of the Council of Protectors


Domain	Functions
Strategic	Adopt the multiannual strategic plan and annual priorities; adopt the budget (within the shielded multiannual framework, section 7.6); approve the algorithmic audit programme
Leadership	Appoint and, where warranted, remove the Executive Director (with cause and under a regulated procedure); appoint the Fundamental Rights Officer; approve the annual public reports
Guarantor	Verify compliance with founding values; approve the annual ethics-code review; answer to the European Parliament and public opinion — never to individual governments

Why 17 and not 27: a 27-member body with state representation is a body where diplomatic calculation replaces merit. The ERC works with ~22 experts. Seventeen is the minimum size that covers all technical profiles with staggered renewal by thirds (every 2 years, 5-6 seats are renewed without decapitalising accumulated knowledge) and the maximum that allows real deliberation in plenary. No seat belongs "to" any country.

7.3 Executive Directorate

The **Executive Director** runs the agency day to day: implementation of the strategic plan, crisis management, public representation and internal budget allocation. They are appointed by the Council of Protectors from the shortlist of an independent selection committee — the same open procedure the EU already applies for agency executive directors, verified in the Europol 2026 process: public call, board selection committee, interviews and assessment centre, shortlist of at least 3 candidates.

Position	Mandate	Reports to
Executive Director	5 years, renewable once	Council of Protectors
Director of Operations (detection, audit, electoral protection)	4 years	Executive Directorate
Director of Technology (AI, infrastructure, sovereignty)	4 years	Executive Directorate
Director of Legal Affairs and Safeguards	4 years	Executive Directorate
Director of Education and Resilience	4 years	Executive Directorate
Fundamental Rights Officer (independent)	5 years, non-renewable	Council of Protectors directly (Frontex precedent)



7.4 Technical and administrative departments

1. Detection and Analysis (330 staff)

The eyes of the EDSPA. Three teams: **AI** (150: in-house models for detecting disinformation, deepfakes and synthetic content), **Data Analysis** (100: large-scale processing, platform telemetry) and **Intelligence** (80: analysis of actors and coordinated campaigns, OSINT). Real-time detection, manipulation-pattern identification, early warnings. The largest headcount in the organogram: the agency is, above all, a perception machine.

2. Algorithmic Audit (200 staff)

The voice of transparency. Three teams: **Technical Audit** (120: inspection of recommender systems and reach modulation, static and dynamic algorithm analysis), **Legal** (50: DSA, AI Act, GDPR) and **Compliance** (30: follow-up of recommendations). It conducts independent audits of platforms, assesses regulatory compliance, issues recommendations and publishes public reports. It is the arm that turns the evidence of Chapter 4 (recommender algorithms amplifying falsehood) into an obligation to correct.

3. Electoral Protection (200 staff)

The shield of democracy. Three teams: **Monitoring** (100: 24/7 vigilance during electoral periods), **Rapid Response** (60: real-time neutralisation of interference, bots, deepfakes) and **Coordination** (40: liaison with national electoral authorities). Detection of foreign interference, public post-election reports.

4. Cognitive Education (190 staff)

The resilience builder. Three teams: **Content** (80: materials for schools and citizens), **Training** (60: teachers, journalists, fact-checkers) and **Campaigns** (50: mass awareness). It is the only line of defence that works *before* manipulation arrives: cognitive vaccination rather than later cure.

5. Digital Sovereignty (250 staff)

The guarantor of autonomy. Three teams: **Infrastructure** (100: dedicated EuroHPC capacity, agency data centres), **Cybersecurity** (80: defence of the EDSPA itself and critical infrastructure) and **Technological Sovereignty** (70: promotion of European models and tools, certification of secure technologies).

6. Administration (150 staff) and Fundamental Rights Office

HR (50), Finance (40), Communication (30), Legal Affairs (30). The Fundamental Rights Office (section 7.6) is deliberately small and outside this chain of command.

Staff and budget distribution (base scenario)

Department	Staff	%	Annual budget
Council of Protectors + Advisory Council	22	1.5%	EUR 5M
Executive Directorate + FRO	35	2.3%	EUR 12M
Detection and Analysis	330	22%	EUR 95M
Algorithmic Audit	200	13.3%	EUR 60M
Electoral Protection	200	13.3%	EUR 70M
Cognitive Education	190	12.7%	EUR 55M
Digital Sovereignty	250	16.7%	EUR 95M
Administration	150	10%	EUR 38M
Flexible reserve (crisis contracts, special task forces)	123	8.2%	EUR 70M
Total (base scenario)	1,500	100%	EUR 500M

7.5 National nodes: the eu-LISA + EDPB dual model

The absence of states from the central council does not mean ignoring the states: it means placing them where they help and removing them where they hinder. The national nodes — the territorial seats in each founding country — follow a **dual model** built on two verified EU patterns:

Component 1: Operational nodes (eu-LISA pattern)


Just as eu-LISA operates the central systems (SIS, VIS, Eurodac) while states only access them through designated authorities, the EDSPA's national nodes are **operational branches directed 100% from headquarters**:

- **Node staff = agency staff** (central contracts and chain of command), not national civil servants on secondment
- The node director is appointed by the Executive Directorate, not by the host government
- Functions: detection in national languages, liaison with locally-jurisdictional platforms, territorial electoral support, execution of the national literacy strategy
- The host state provides infrastructure, physical security and administrative facilitation — nothing more

Component 2: National Chamber (EDPB pattern)

Just as the EDPB brings national data-protection authorities together to coordinate standards, each node hosts a **National Chamber** where the state is present — but only for what the operation needs from it:

- **Composition:** the competent national authority (media or data regulator), one government representative, one local civil-society representative
- **Functions:** linguistic and cultural adaptation of alerts, prioritisation of national literacy campaigns, channelling requests from national electoral authorities
- **No operational power:** the Chamber does not direct detections, does not prioritise audits, cannot veto or delay any agency action. Its voice is real; its vote on operational matters is non-existent



7.6 Frictionless safeguards: accountability that never becomes a blockade

Independence without control is opacity, but excessive control is capture. The most delicate section of the institutional design is this one: **enough guarantees for the EDSPA to be legitimate, without states or corporate litigants being able to use them to slow it down.** EU law offers exact precedents for each piece:

Piece 1: Single judicial review (ECB precedent — banking supervision)

The CJEU held in *Berlusconi and Fininvest v Banca d'Italia* that ECB supervisory decisions are subject to **single judicial review before the EU Courts**, excluding parallel national jurisdiction. Translation for the EDSPA: its operational decisions (audits, orders, alerts) would be challengeable only before the General Court and the CJEU — **never before 27 national courts simultaneously**, the procedural drowning strategy that large platforms already use against regulators. One forum, one doctrine, coherent case law.

Piece 2: Fast internal, non-suspensive administrative review (ECB ABoR precedent)

The ECB's Administrative Board of Review examines challenges to its supervisory decisions through a fast internal procedure that is **non-suspensive**: the decision remains in force while under review. The EDSPA would adopt an equivalent Review Board: addressees of a measure get a real challenge channel (legitimacy), but filing the challenge **does not freeze the protection** (effectiveness). Platforms cannot buy time with paperwork.

Piece 3: Protected operational files (AM&S / Akzo Nobel precedent)

CJEU case law protects communications covered by legal professional privilege from mandatory disclosure. The EDSPA's regulation would extend that logic to the **operational file**: ongoing analysis dossiers, detection methodologies and actor intelligence are not discoverable by litigants in civil proceedings. Radical transparency about results (public reports, aggregated open data, Court of Auditors audits); regulated confidentiality about method while the operation lasts.

Piece 4: Limitation periods (EDPS precedent)


The EDPS only investigates complaints filed within **two years** of the complainant becoming aware of the matter. The EDSPA would apply analogous limitation periods to administrative challenges: without deadlines there is no legal certainty for the agency, and procedural eternity only benefits those who can afford armies of lawyers.

Piece 5: Fundamental Rights Officer with access but no operational veto (Frontex precedent, improved)

Frontex created its Fundamental Rights Officer by regulatory requirement after its governance scandals. The EDSPA includes one from day one, improved: full access to systems and staff, an accessible citizen complaints mechanism, annual public reports, and the duty to alert the Council of Protectors and the European Parliament — **reporting directly to the Council, never to the operational chain it monitors**. But its recommendations do not suspend operations: they trigger review, not paralysis.

Piece 6: Shielding against the legislator (the Europol 2022 lesson)

Case T-578/22 showed that the legislator can dismantle in one regulatory amendment what a supervisor built over years. The EDSPA's founding regulation would include a **shield clause**: amendments to the mandate, the governance model or the multiannual financing only through the same reinforced ordinary legislative procedure (Parliament + Council + mandatory consultation of the Council of Protectors itself and the EDPS), never by implementing regulations, and **never with retroactive effect** on ongoing operations. Independence is founded once, with difficulty; it must also be dismantled with difficulty.



The balance in one sentence

Everything is public except the method while it operates; everything is challengeable except in 27 forums at once; everything is auditable except by those who audit in order to litigate; and everything can be reformed by the legislator, except more easily than it was founded. The agency has owners (European citizens, through their Parliament and their courts), but neither those it monitors nor governments of the day can become its directors.

7.7 Coordination mechanisms and indicators

Coordination rhythm

Body	Frequency	Participants	Objective
Council of Protectors	Quarterly + extraordinary sessions	17 members (+ observers)	Strategic and budgetary decisions
Executive Committee	Weekly	Executive Directorate + directors	Operational management
States Advisory Council	Semi-annual	Member State representatives	Information, consultation on territorial priorities (no vote)
National Chambers	Quarterly (more during electoral periods)	National authority + government + civil society	Cultural adaptation, literacy coordination
Crisis cells	Activated on alert	Detection + Rapid Response + affected node	Real-time neutralisation

Key performance indicators (KPIs)

KPI	2027 target	2028 target	2030 target	Owner
Disinformation detection time	<12 h	<6 h	<2 h	Detection and Analysis
Platforms algorithmically audited	30%	60%	90%	Algorithmic Audit
Disinformation campaigns neutralised/year	200	500	1,000	Electoral Protection
Citizens trained in digital literacy	5M	15M	50M	Cognitive Education
Dedicated European capacity (HPC/AI)	EUR 200M	EUR 400M	EUR 800M	Digital Sovereignty
FRO complaints resolved on time	>80%	>85%	>90%	Fundamental Rights Office

7.8 Conclusion: a structure built for action, shielded to last

The EDSPA's structure is deliberately asymmetric: **maximum independence where it deliberates** (17 merit-selected experts, zero governments), **maximum professionalism where it executes** (an executive directorate under an open procedure), **territorial presence without yielding command** (nodes as branches, states in consultative chambers) and **safeguards that listen but never block** (single judicial review, non-suspensive appeal, protected operational files, an independent FRO). It is the skeleton of the Protective Power: not bureaucracy, but the first European institution designed from its foundation so that nobody — neither the governments that fund it, nor the platforms it monitors, nor the litigants that sue it — can capture the instrument that protects the minds of Europeans.

Chapter 8 develops how that structure operates: the detection, analysis and response processes, human-plus-AI double verification, and real-time integration of the nodes.

Verified sources of this chapter

- ERC: Scientific Council of independent merit-selected experts, governing body of the ERC, with no state representation — the governance model adopted for the Council of Protectors
- ENISA, EUIPO, eu-LISA: boards with one representative per Member State + Commission (standard model, discarded for the headquarters); ENISA: 4-year renewable mandates, 5-member Executive Board
- Europol: 2026 executive director appointment procedure — board selection committee, interviews and assessment centre, shortlist of at least 3 candidates (applied to the EDSPA Executive Director)
- Case T-578/22 (EDPS v Council and EP, 2022 Europol Regulation): EDPS annulment action inadmissible after retroactive legalisation of data practices — foundation of the shield clause
- CJEU, *Berlusconi and Fininvest v Banca d'Italia*: single judicial review of ECB/SSM decisions before the EU Courts — foundation of the anti-litigation piece
- ECB Administrative Board of Review (ABoR): fast, non-suspensive internal administrative review

- CJEU, *AM&S* (155/79) and *Akzo Nobel* (C-550/07 P): legal professional privilege against disclosure — foundation of the protected operational file
- Frontex: Fundamental Rights Officer and complaints mechanism (Regulations 1168/2011 and 2019/1896); OLAF/Leggeri 2022 scandals as a governance lesson
- EDPS: two-year complaint limitation (EDPS Rules of Procedure, 2022 amendment); supervisor independence
- ECB: five pillars of institutional independence, prohibition of instructions to its bodies
- EDPB: composition of national independent authorities (pattern of the National Chamber); eu-LISA: central operation of SIS/VIS/Eurodac with designated national authorities (pattern of the operational nodes)

Chapter 8: Operations and Processes of the EDSPA

How the agency turns data into protection: detection with mandatory human-AI dual verification, algorithmic auditing, electoral protection, education and sovereignty, resting on the NIS2 framework with deadlines that shorten as the system improves

Chapter summary

The structure of Chapter 7 only matters if it operates. This chapter details the processes that turn the EDSPA into a living system: the **detection-analysis-response cycle** with **mandatory human-AI dual verification** and fully auditable traceability; **algorithmic auditing** that extends obligations already existing under the DSA (systemic risk, independent audits since 2024) with the regulator's own inspection capacity; **electoral protection** that institutionalises what the Rapid Response System of the 2024 European elections proved on an improvised basis; scalable **cognitive education**; and **digital sovereignty** with updated figures from the European ecosystem (Mistral AI: EUR 1,700M round in 2025, valuation of roughly US\$13,700M). The whole edifice rests on an incident-management framework **inspired by NIS2 but more ambitious** (12-hour alert / 48-hour notification, with a progressive shortening clause) and on target times aligned with the KPIs of Chapter 7.

8.1 The operational cycle: from signal to protection

Before anything else, the EDSPA is a perception machine with judgement. Every threat travels the same cycle: **continuous monitoring** → **anomaly identification** → **classification** → **dual verification** → **prioritisation** → **response** → **learning**. What distinguishes the agency from a mere automated system is the layer of human judgement at the points where error carries a democratic cost.

The dual verification protocol (cross-cutting)

The guiding principle of all operations: **AI proposes, the human decides** — mandatory at the High and Critical levels, and with permanent human supervision at the Medium and Low levels. No agency action affecting content, accounts or platforms is executed without a machine having detected it and a person having validated it.

Element	Rule
Detection (AI)	Proprietary models scan platforms, networks and media continuously; they generate threat candidates with a confidence level and attached evidence
Validation (human)	An analyst from the Detection and Analysis Department confirms or dismisses each Medium, High or Critical candidate; validation time counts within the target detection deadline
Full traceability	Every decision logs: the model and version that detected it, the evidence used, the analyst who validated it, the basis for classification, the exact time. The log is auditable by the Fundamental Rights Office and the Court of Auditors
Disagreement	If AI and human disagree at the High/Critical level, the case escalates to a second analyst (nothing is ever executed or archived by the decision of a single actor, human or machine)
False positives	Measured by department, they feed the monthly retraining of the models; the error-rate target is set in the Chapter 7 KPIs (accuracy >90% in 2027, >95% in 2030)

Incident classification

The incident-management framework adopts the skeleton of **NIS2** — the verified European rule requiring an early warning within 24 hours, full notification within 72 and a final report within one month — and **tightens it**, because cognitive threats spread faster than a cybersecurity incident: the critical window of an electoral deepfake is measured in hours, not days.

Level	Description	Early warning	Full notification	Who validates
Critical	Active threat to an ongoing electoral process or to the integrity of a democratic institution	< 1 hour	< 12 hours	Mandatory dual verification + second analyst in case of disagreement
High	Mass disinformation campaign or a deepfake of a leader in circulation	< 6 hours	< 24 hours	Mandatory dual verification
Medium	Confirmed algorithmic manipulation or a campaign targeting a specific group	< 12 hours	< 48 hours	Human validation under supervision
Low	Potential threat, emerging pattern without confirmation	< 48 hours	weekly	Review by human sampling
Progressive improvement clause		Deadlines are reviewed every year and shortened as the accuracy of the detection-verification system improves: the NIS2 standard (24 h/72 h) is the ceiling, never the target		

From signal to action: the operational cycle

The EDSPA operational cycle: AI perceives, the human judges





Figure 8.1 — The operational cycle with dual verification. Author's own elaboration; target times aligned with the Chapter 7 KPIs and the improved NIS2 framework.

Target detection times by threat type

Type	Example	2027 target	2030 target	Department
Active electoral interference	Foreign campaign on election day	< 30 min	< 15 min	Electoral Protection
Deepfake of a public figure	Fabricated video of a politician	< 1 hour	< 30 min	Detection and Analysis
Mass disinformation campaign	Coordinated fake news	< 6 hours	< 2 hours	Detection and Analysis
Algorithmic manipulation	Bias in recommendation systems	< 4 hours (alert)	< 2 hours	Algorithmic Auditing
Data / infrastructure breach	Leak, attack on agency systems	improved NIS2 framework: alert < 12 h, notification < 48 h		Digital Sovereignty

Target detection times: 2027 versus 2030

Detection targets: time is halved between 2027 and 2030



Scale: minutes to hours from first signal — accuracy KPI >90% (2027), >95% (2030), Ch. 7

Figure 8.2 — Comparison of target detection times between the 2027 and 2030 scenarios. Author's own elaboration from Table 8.1 and the Chapter 7 KPIs.

8.2 Algorithmic auditing: from the DSA to in-house capacity

The DSA already obliges very large platforms (VLOPs/VLOSEs) to **assess their systemic risks** and to submit to **independent audits** — the corresponding delegated act (Delegated Regulation (EU) 2024/436) has been in force since January 2024 and applied to the first round of audits, due in August 2024; the first risk assessments were published from 2024-2025. But the system rests on the platforms' self-declaration and on auditors hired by the audited companies themselves. The EDSPA adds the missing piece: **inspection capacity belonging to the regulator**, just as the ECB inspects banks instead of trusting their self-assessments.

Phase	Duration	Actions
Platform selection	1 week	Annual programme approved by the Council of Protectors (target: 30% of the landscape in 2027, 90% in 2030 — Ch. 7 KPI)
Access requirement	3 days	Formal request for data, documentation and test access; in case of refusal, an executive order backed by the DSA and single judicial review (Ch. 7.6)
Technical analysis	2-4 weeks	Static and dynamic inspection of recommendation and reach-modulation systems; behavioural testing with synthetic profiles
Report and resolution	1 week	Public report (findings) + protected operational file (method, Ch. 7.6); binding recommendations with correction deadlines
Follow-up	Continuous	Verification of corrections; in case of non-compliance, proposal for a sanction to the competent authority (DSA: up to 6% of global turnover; AI Act: up to 7% or EUR 35M)

8.3 Electoral protection: institutionalising what 2024 improvised

The June 2024 European elections proved the concept with makeshift means: the **Rapid Response System** of the Code of Practice on Disinformation allowed fact-checkers to flag threatening content to the platforms during the campaign — and EDMO evaluated how it worked, including a separate channel with X. It worked at limited scale, with no binding capacity, and dependent on the platforms' goodwill. The EDSPA turns it into a **permanent protocol with the power to demand**:

Phase	When	Actions
Preparation	6 months before	Electoral operations centre activated; country-by-country risk mapping; training of national electoral authorities (with the National Chamber, which has no operational power over the agency — Ch. 7.5)
Monitoring	Final month + election day	24/7 surveillance with a lowered alert threshold; mandatory dual verification for every electoral incident (Critical level by definition)
Response	Real time	Neutralisation of interference; verified public alerts; platform requirements with deadlines of hours, not days
Post-election	1 month after	Full public report with methodology; lessons feeding the retraining of the models

The precedent that validates the design: in 2024, the system worked without coercive capacity and still allowed documented threats to be neutralised. With the power to issue requirements, a protected file and single judicial review — the pieces of Chapter 7 —, the same architecture multiplies its effectiveness without depending on anyone's voluntary cooperation.

8.4 Cognitive education and digital sovereignty

Education programmes (base scenario budget, Ch. 7)

Programme	Audience	Duration	Annual budget
Basic digital literacy	General public	10 hours	EUR 30M
Advanced literacy	Professionals and civil servants	40 hours	EUR 20M
Fact-checker training	Journalists and local media	80 hours	EUR 15M
Algorithmic auditor training	Technicians (future sector professionals)	120 hours	EUR 10M
Awareness campaigns	General public	Variable	EUR 25M
Total annual	target: 5M citizens (2027) → 50M (2030)		EUR 100M

The programmes are delivered through the national nodes and in collaboration with the existing network of 15 EDMO hubs and national education systems, without duplicating infrastructure: the agency produces standards and content; local actors implement them on the ground.

Digital sovereignty: the real dependency map

The dependency table is updated with verified 2025-2026 evidence. The European ecosystem is no longer theoretical: Mistral AI completed a **EUR 1,700M** funding round in September 2025 led by ASML, with a valuation of roughly **US\$13,700M** in 2026 and a forecast of exceeding **US\$1,000M in revenue** that same year; exascale supercomputing exists (JUPITER, LUMI, MareNostrum 5); the legal framework is in force (AI Act with full staggered application by August 2026, EMFA since August 2025, Anti-SLAPP since May 2026).

Area	Current dependency	Verified European alternative	EDSPA route
Frontier AI models	US (OpenAI, Anthropic, Google)	Mistral AI (EUR 1,700M / ~USD 13,700M); European open models	Dedicated capacity and sovereign deployment standards
Supercomputing	US hyperscalers (~63% of the global cloud market)	EuroHPC: JUPITER (exascale), LUMI, MareNostrum 5	Dedicated capacity instead of owned infrastructure (Ch. 6.4)
Cloud and storage	AWS, Azure, Google Cloud	European providers (~15% of the European market) + EU Cybersecurity Reserve (EUR 36M, ENISA)	Sovereignty requirements for the agency itself and the nodes
Detection tools	US solutions and private philanthropy	In-house development + EDMO hubs + R&D investment	Own tools (Ch. 7.4, AI team: 150 people)
Detailed investment by area belongs to the funding scenarios of Chapter 6 and the capacity KPIs of Chapter 7			

8.5 Coordination and interoperability

Coordination with the Member States (within the Chapter 7 model)

State coordination operates exclusively through the channels designed in Chapter 7 — the **Advisory Council** (headquarters) and the **National Chambers** (nodes) — and through technical channels the EU already has in verified operation:

- **Permanent technical channel:** secure exchange platform with the competent national authorities (eu-LISA pattern: access to central systems through designated authorities)
- **Situation reports:** weekly aggregated and anonymised reports to participating states; real-time alerts to the affected authority
- **Annual joint exercises** simulating cognitive crises, in the style of EU cyber exercises (defined in the EU Cyber Solidarity framework, the basis of the EUR 36M Reserve)


International cooperation (reordered after the 2026 AI veto)

The US AI veto of June 2026 (Ch. 4) forces cooperation priorities to be ordered by democratic reliability and cognitive security, not by Atlanticist habit:

Partner	Area	Level
United Kingdom	Cognitive protection, verification, electoral intelligence (DEFENCE/OFCOM)	High — natural partner with a shared threat picture
Canada	Technology, regulation, experience with foreign interference	High
Japan	AI, disinformation, synthetic content standards	High
South Korea	Platform regulation, electoral experience with AI	High
Australia	Cybersecurity, legislation against foreign manipulation	High
OECD / UN	Global standards, algorithmic transparency	High (multilateral)
NATO	Cyber defence, cross-checking hostile state actors	Medium — limited to the security dimension
United States	Intelligence and good-practice exchange	Conditional — reviewable after the June 2026 AI veto; no operational dependence on providers subject to the IIL

Interoperability with the existing European ecosystem

The EDSPA does not start from zero: it interconnects with what already exists



Interoperability with verified actors — budgets and capacities documented in Chs. 5 and 6

Figure 8.3 — The EDSPA's position within the European cognitive-protection ecosystem. Author's own elaboration on verified bodies and mechanisms (EDMO, East StratCom, ENISA, EuroHPC, DSA).

8.6 Quality and transparency standards

Scope	Standard	Verification
Detection	Target times by level (8.1); accuracy >90% (2027) → >95% (2030)	Automatic + annual audit by the Court of Auditors
Auditing	Public report per audit (findings); protected method file (Ch. 7.6)	Single judicial review before EU courts
Transparency	Public annual report; quarterly aggregate statistics; summary minutes of the Council of Protectors	Public access + questions from the European Parliament
Rights	Complaints to the Fundamental Rights Office (2-year limitation, Ch. 7.6); FRO annual report	Independent FRO, full access to systems
Security	ISO 27001; strict GDPR; continuity plans tested annually	External audits (ENISA pattern)
Traceability	Auditable log of every dual-verification decision (AI + human)	Auditable by the FRO and the Court of Auditors — never obtainable by litigants

8.7 Conclusion: operations that save democracies

The processes in this chapter turn the architecture of Chapters 6 and 7 into real protection: continuous perception with **mandatory human judgement** where error carries a democratic cost; deadlines **more demanding than NIS2** that shorten year after year; auditing that completes the DSA with the regulator's own inspection capacity; electoral protection that institutionalises what June 2024 demonstrated with improvised means; and sovereignty built on a European ecosystem that already exists — Mistral, EuroHPC, ENISA, EDMO — but lacks the conductor. That operational conductor, with cost investigation, auditing and cooperation as its master lines, is the EDSPA's contribution.

Chapter 9 measures what all of this is worth: the economic return and direct benefits of the agency.

Verified sources for this chapter

- NIS2 (Directive 2022/2555): incident management with early warning in 24 h, notification within 72 h and final report within 1 month; transposition deadline 17-Oct-2024 — framework tightened by the EDSPA (12 h/48 h with progressive improvement)
- DSA: independent audits of VLOPs/VLOSEs (Delegated Regulation (EU) 2024/436, in force since Jan 2024, first round Aug 2024); systemic risk assessments published 2024-2025; penalties up to 6% of global turnover
- AI Act: staggered application through August 2026; penalties up to 7% or EUR 35M for prohibited practices
- Rapid Response System of the Code of Practice on Disinformation: operational during the June 2024 European elections; evaluated by EDMO, including a separate channel with X
- EDMO: central unit ~EUR 2.5M (DIGITAL 2025); 15 national and regional hubs — the network with which to interoperate on literacy and verification
- ENISA: 2025 budget of EUR 26.9M; EU Cybersecurity Reserve of EUR 36M (Digital Europe 2025-2027, contribution agreement signed with ENISA, 26-Aug-2025)
- Mistral AI: EUR 1,700M round in September 2025 (led by ASML); valuation of ~US\$13,700M and a forecast of >US\$1,000M in revenue in 2026
- EuroHPC: JUPITER (exascale, Juelich), LUMI (Kajaani), MareNostrum 5 (BSC, Barcelona)
- European Media Freedom Act: applicable since 8-Aug-2025; Anti-SLAPP Directive (EU) 2024/1069: transposition required by 7-May-2026
- European Democracy Shield (Joint Communication, Nov-2025): situational awareness and response capability as priorities — with the Delors Centre critique already integrated in Ch. 5

Chapter 9: ROI and Direct Economic Benefits of the EDSPA

Cognitive protection as a strategic investment: every euro invested generates between 1.6 and 2.6 euros in direct and indirect benefits over five years, under the 40/35/15/10 funding model of Chapter 6 and costs aligned with the scenarios of Chapter 7

Chapter summary

This chapter demonstrates that the EDSPA **is not an expense but an investment with a quantifiable return**. With the cost model aligned with Chapters 6 and 7 (initial investment of **EUR 1,100M** without owned infrastructure — dedicated capacity on EuroHPC and European cloud — and operations of **EUR 500M per year** in the base scenario), direct benefits at full capacity (**EUR 1,450M per year**, not counting fines as agency income) plus indirect benefits yield a **five-year ROI of 164%** in the base scenario: **every euro invested generates ~EUR 2.6**. The accelerated scenario (EUR 2,500M, 2,000 staff, start 2026-2027) lowers the direct financial return (~EUR 1.7 per euro) but brings protection forward against the 2026 evidence: AI-driven job destruction is already an accounting fact, not a prediction.

9.1 The urgency argument: the acceleration of AI in 2026

The forecasts underpinning this chapter were written before the 2026 evidence made it clear that **the risk calendar is accelerating faster than the response capacity of the Member States**. The warnings no longer come from outside critics of the industry: they come from its own founders.

Bill Gates: "there is not even a plan to have a plan"

In his August 2026 essay, *The turbulent AI era is here. The choices we make now are critical*, Bill Gates warns that AI **is already crossing thresholds that were considered distant alarm signals**: easier cyberattacks, biological risks, emotional dependence on chatbots, mass job displacement and possible loss of control over advanced systems. Gates compares the situation with his pre-COVID-19 warnings: the data was there, but leaders did not act in time, and he has proposed an international AI monitoring organisation inspired by nuclear inspection regimes.

Geoffrey Hinton: a new wave of job destruction from 2026

The "godfather of AI" and Nobel laureate warns that 2026 opens a **new wave of job destruction**: AI "is going to be able to replace lots of other jobs". Hinton foresees massive disruption of white-collar work between the end of this decade and the start of the next, with growing inequality, and notes that the tech companies "are betting that AI will replace many workers".

Narrative manipulation: Altman and Amodei tone down their alarms ahead of their IPOs


An even more revealing fact for this book's thesis: Sam Altman (OpenAI) and Dario Amodei (Anthropic), who in 2025 predicted a "job apocalypse" and a "white-collar massacre", **have publicly softened their predictions in 2026, coinciding with the preparation of their stock market listings**. Anthropic's own economics team published in March 2026 that "no systematic increase in unemployment is observed for the most exposed workers since late 2022, though hiring of younger workers has slowed in exposed occupations". When the narrative about risk is controlled by those who sell the technology, independent assessment stops being optional: it becomes a structural necessity. That is precisely the EDSPA's function — with the human-AI dual verification and full traceability of Chapter 8 as the guarantee that this assessment is auditable.

Labour damage already measurable in 2026

Indicator (2026)	Figure	Source
US layoffs attributed to AI (Jan-Aug 2026)	116,175 (22% of all layoffs)	Challenger, Gray & Christmas
US tech-sector layoffs (Jan-Aug 2026)	155,126 (+52% vs. 2025)	Challenger / Layoffs.fyi
Global tech layoffs (Q1 2026 alone)	>73,000 across 95 companies	Layoffs.fyi
Gross job destruction in Spain (next decade)	1.7-2.3 million jobs (central scenario; partially offset by ~1.61M new AI-linked occupations per the same study)	Funcas
US CEOs expecting mass AI-driven layoffs within 2 years	99% (survey of 825 CEOs)	Mercer
Jobs AI is destroying in the US	~16,000 per month (Apr-2026, revised to ~11,000 in Jun-2026; greatest impact on Generation Z)	Goldman Sachs Research
Global jobs exposed to AI-driven automation	~300 million	Goldman Sachs Research

AI-linked layoffs at major tech companies (2025-2026)

Jobs eliminated at major tech companies, attributed to or linked to AI (2025-2026)



Sources: eldiario.es, Layoffs.fyi, Reuters (2026). None has acknowledged equivalent net job creation in Europe.

Figure 9.1 — Tech layoffs attributed to or linked to AI / restructuring towards AI. Author's own elaboration on Layoffs.fyi and verified press, 2026.


Consequence for this chapter: if AI adoption accelerates (Goldman Sachs places the Netherlands, the US and France at the forefront of adoption, with slower hiring in the most exposed sectors), the EDSPA's calendar must accelerate in parallel. Hence the two scenarios: the **base** scenario (conservative, aligned with Chs. 6-7) and the **accelerated** scenario (recommended in light of the 2026 evidence).

9.2 Cost analysis: the initial investment

Base scenario (2027-2028): EUR 1,100M

Item	Cost	%	Timing
Headquarters and equipment	EUR 200M	18%	2027-2028
Dedicated EuroHPC capacity and sovereign European cloud (multi-year contracts)	EUR 250M	23%	2027-2028
Network infrastructure and national node deployment	EUR 100M	9%	2027-2028
Staff hiring and initial deployment (1,500 employees, Ch. 7)	EUR 150M	14%	2027-2028
Development of in-house detection and auditing tools (Ch. 8)	EUR 200M	18%	2027-2028
Training and capacity building	EUR 50M	5%	2027-2028
Literacy and communication start-up (Ch. 8 programme)	EUR 70M	6%	2027-2028
Contingency reserve	EUR 80M	7%	2027-2028
Total initial investment	EUR 1,100M	100%	2027-2028

Breakdown of the initial investment: EUR 1,100M (base scenario)



No owned infrastructure: dedicated EuroHPC capacity and European providers, per Ch. 6.4

Figure 9.2 — Breakdown of the initial investment. Author's own elaboration on the cost scenarios of Chapters 6 and 7.


Accelerated scenario (2026-2027): EUR 2,500M

In light of the acceleration documented in 9.1, the accelerated scenario — the one funded in Chapters 6 and 9B — brings the start forward to 2026 with an investment of **EUR 2,500M** that adds to the base scenario: early hiring of up to **2,000 employees**, parallel deployment of the national nodes, expanded dedicated computing capacity, and a **European labour reconversion programme** for workers displaced by AI. Full operational capacity is brought forward to 2028.

9.3 Direct benefits: EUR 1,450M per year at full capacity

Category	Annual benefit	Methodology	Reference
EU fines avoided for DSA/AI Act violations that early detection prevents	EUR 400M	Reduction in documentable violations	Regulatory impact study
Improved productivity	EUR 300M	Less time lost to disinformation	Eurostat, 2026
Cybersecurity savings	EUR 250M	Fewer successful attacks on critical infrastructure	ENISA, 2025
Reduced digital fraud	EUR 200M	Fewer scams fuelled by manipulation	Europol, 2026
Public health savings	EUR 200M	Less impact of health disinformation	WHO, 2026
Additional taxes from the associated ecosystem	EUR 100M	Greater activity in the European digital sector	European Commission, 2026
Total annual direct benefits	EUR 1,450M	At full capacity (2029 in the base scenario)	

Direct benefits at full capacity: EUR 1,450M per year



No fines as agency income: detect/sanction separation, Chs. 6 and 8

Figure 9.3 — Direct benefits by category. Author's own elaboration; per-category methodology in Table 9.3.

Added to these direct benefits are the **indirect** ones — the public-spending multiplier in the European technology ecosystem, associated skilled employment, private investment attracted to the sector — estimated at **~EUR 870M per year** at full capacity (60% of direct benefits, a conservative estimate whose full methodology belongs to Chapter 9B). Total benefit at full capacity thus reaches **~EUR 2,320M per year**.

9.4 ROI calculation

Base scenario (start 2027, full capacity 2029)

Item	Calculation	Result
Initial investment	2027-2028	EUR 1,100M
Operations (5 years)	EUR 500M/year x 5	EUR 2,500M
Total 5-year cost	1,100 + 2,500	EUR 3,600M
Direct benefits (ramp-up 40% / 70% / 100%)	580 + 1,015 + 1,450x3	EUR 5,945M
Indirect benefits (same profile)	348 + 609 + 870x3	EUR 3,567M
5-year ROI (2031)	(9,512 - 3,600) / 3,600	164% — every euro generates ~EUR 2.6

- **5-year net benefit: +EUR 5,912M**
- **Payback of the initial investment:** ~2.5 years from start-up (the annual net benefit at full capacity, EUR 1,820M, exceeds the initial investment in the first full year of operation)

Accelerated scenario (start 2026, full capacity 2028)

Item	Calculation	Result
Initial investment	2026-2027	EUR 2,500M
Operations (5 years)	EUR 1,000M/year x 5 (full scenario, Ch. 6)	EUR 5,000M
Total 5-year cost	2,500 + 5,000	EUR 7,500M
Direct benefits (ramp-up 50% / 80% / 100%, +EUR 400M for expanded capacity and reversion)	925 + 1,480 + 1,850x3	EUR 7,955M
Indirect benefits	555 + 888 + 1,110x3	EUR 4,773M
5-year ROI (2030)	(12,728 - 7,500) / 7,500	~70% — every euro generates ~EUR 1.7

The correct reading of the accelerated scenario: its direct financial ROI is lower because it spends more, and sooner. But its justification is not accounting: it is **avoided damage**. Every month of earlier full capacity reduces the exposure of the European economy to the job destruction and manipulation already documented in 9.1 — costs that this deliberately conservative calculation does not monetise. The accelerated scenario buys time, and time is exactly what the 2026 evidence says Europe does not have. Its labour reversion programme also generates returns that mature beyond the five-year window.

9.5 Financial sustainability: the 40/35/15/10 model

Sustainability is not measured in "own revenue" but in **the stability of the Chapter 6 funding model**: EU budget contribution (40%), contributions from participating states (35%), transfer of the EU-collected fines share (15%) and other sources — service fees for audits and training for third parties, R&D programmes (10%). The agency does not collect: the model shields it from the revenue-collection incentives that would corrupt its function.

Source (Ch. 6)	%	Base scenario (EUR 500M/year)	Full scenario (EUR 1,000M/year)
EU budget	40%	EUR 200M	EUR 400M
Participating states (the 9 of the founding core, Art. 20 TEU)	35%	EUR 175M	EUR 350M
Transfer of the EU-collected fines share (DSA 6% / AI Act 7% or EUR 35M)	15%	EUR 75M	EUR 150M
Other sources (service fees, R&D)	10%	EUR 50M	EUR 100M
Total per year	100%	EUR 500M	EUR 1,000M

The 40/35/15/10 model: base and full

Annual funding by source: base and full scenarios



Base: EUR 500M/year (1,500 staff)

Full scenario: EUR 1,000M/year (EU EUR 400M, states EUR 350M, fines EUR 150M)

Figure 9.4 — The 40/35/15/10 funding model of Chapter 6. Author's own elaboration.

9.6 Comparison with other public investments in Europe

Investment / agency	Annual or total figure	Reference
EDSPA (base scenario)	EUR 500M/year — 0.003% of EU GDP (~EUR 18.8 trillion, Eurostat 2025)	Ch. 6
EDSPA (accelerated/full scenario)	EUR 1,000M-2,500M	Chs. 6 and 9B
ENISA (EU cybersecurity)	EUR 26.9M (2025)	ENISA budget
EU Cybersecurity Reserve	EUR 36M (2025-2027)	Digital Europe / EU Cyber Solidarity
Europol	EUR 222M / ~1,400 staff	Europol budget
Frontex	~EUR 1,100M	Frontex budget
EMA (London → Amsterdam, 2017-2019)	>EUR 600M/year	Relocation precedent of Ch. 6
InvestAI (EU AI plan)	EUR 200,000M	European Commission, 2025
France (AI cluster)	EUR 109,000M	2025 announcement
Reading	Europe is spending 200 times the cost of the agency that should accompany it on AI without cognitive protection: buying the engine without the seatbelt	

9.7 Conclusion: an investment that pays for itself

With costs aligned with Chapters 6 and 7, the EDSPA is one of the few European public investments with a verifiable direct return: **~EUR 2.6 per euro invested over five years in the base scenario**, payback of the initial investment in two and a half years, and an annual cost (EUR 500M) equivalent to 0.003% of EU GDP. The accelerated scenario lowers the accounting ROI but buys the one thing the 2026 evidence shows to be scarce: time against a wave of job destruction and manipulation that is already in the data, not in the forecasts. The broader economic and social impact — employment, multipliers, GDP — is quantified in Chapter 9B, and the constitutional dimension in Chapter 17.

Verified sources for this chapter

- Bill Gates, "The turbulent AI era is here" (Gates Notes, August 2026); MIT Technology Review interview (AI-facilitated bioterrorism risk "50 times more likely than a natural pandemic")
- Geoffrey Hinton (Nobel Prize in Physics 2024): 2026 statements on a new wave of job destruction and white-collar disruption
- Challenger, Gray & Christmas: 116,175 US layoffs attributed to AI in January-August 2026 (22% of the total); tech sector 155,126 (+52% vs. 2025)
- Layoffs.fyi: >73,000 global tech layoffs in Q1 2026 (95 companies)
- Funcas: gross destruction of 1.7-2.3 million jobs in Spain over the next decade (central scenario; the study itself estimates ~1.61M new occupations partially offsetting)
- Mercer: 99% of 825 US CEOs expect AI-driven headcount cuts within 2 years; ~16,000 jobs/month destroyed in the US (Apr-2026, ~11,000 in Jun-2026), ~300M exposed globally; Goldman Sachs Research: ~16,000 jobs/month destroyed in the US, ~300M exposed globally
- Anthropic (economics team, March 2026): "no systematic increase in unemployment for highly exposed workers since late 2022, though suggestive evidence of slowed hiring of younger workers in exposed occupations" — a downgrade coinciding with the preparation of the OpenAI and Anthropic listings
- Eurostat 2025: EU GDP ~EUR 18.8 trillion — basis of the 0.003% calculation
- ENISA EUR 26.9M (2025); EU Cybersecurity Reserve EUR 36M (ENISA agreement, Aug-2025); Europol EUR 222M; Frontex ~EUR 1,100M; EMA >EUR 600M — comparisons verified in Chs. 5 and 6
- InvestAI EUR 200,000M (European Commission, 2025); France EUR 109,000M — budgetary viability evidence in Ch. 6

Chapter 9B: Economic and Social Impact of the EDSPA (2027-2035)

The European Cognitive Protection Agency as a generator of employment, growth and a technology ecosystem — the bridge to the Protective Power

Chapter Summary

This chapter **demonstrates that the EDSPA is not an expense but a strategic investment with a clear, quantifiable return** over its operational period. Drawing on the data from **Chapter 9 (ROI and benefits)** and on author estimates anchored in the economic literature of the domain, it quantifies the agency's impact: **3,000 direct jobs (2035), 6,000-9,000 indirect jobs, EUR 3,160M-4,000M of annual GDP impact, and an ROI of 164% over five years in the baseline scenario that will finance the transition to the Protective Power.**

Central thesis:

The EDSPA is the seed that will demonstrate that cognitive protection works, creates economic and social value, and justifies investment in the Protective Power. Every euro invested in the EDSPA is not only recovered but generates between EUR 2.50 and EUR 3.00 in economic and social benefits, creating a virtuous circle that will allow Europe to lead the digital era.

Objectives:

- Quantify the **direct and indirect employment** generated by the EDSPA (2027-2035)
- Analyse the **economic multiplier effect** (2.5x-3.0x, the author's estimate within the ranges published by the economic literature — see the rigour note in Section 3)
- Measure the **impact on European GDP** (EUR 3,160M-4,000M per year)
- Explain how the EDSPA **prepares the technology ecosystem** for the Protective Power
- Demonstrate how it **finances part of its own transition**
- Present the **accelerated scenario** derived from the 2026 urgency evidence (Chapter 9)

The 2026 urgency applies here as well: the evidence documented in Chapter 9 (warnings from Gates and Hinton, 116,175 layoffs attributed to AI in the US between January and August 2026) means the EDSPA cannot wait. This chapter presents two scenarios: the **baseline** (2027-2035) and the **accelerated** one (2026-2033), with the agency deployed in 2026-2027 and milestones brought forward. The book's key distinction: **accelerating an agency through enhanced cooperation is feasible; the constitutional timetable of the Protective Power has its own institutional limits**, which are respected in both scenarios.

Temporal Context: The EDSPA as the Bridge to the Protective Power

Timetable and Budget — Baseline Scenario (2027-2035)

Phase	Period	Budget	Staff	Main Objective
Implementation	2027-2028	EUR 1,100M	1,500	Building infrastructure and structure
Initial Operation	2028-2030	EUR 500M/year	1,500-2,000	Demonstration of viability (–40% disinformation)
Consolidation	2030-2032	EUR 500M-1,000M/year	2,000-2,500	Building political and social consensus
Transition	2033-2035	EUR 1,000M/year	2,500-3,000	Preparation for integration into the Protective Power

Timetable — Accelerated Scenario (2026-2033, recommended)

Phase	Period	Budget	Staff	Main Objective
Implementation	2026-2027	EUR 2,500M	2,000	Accelerated parallel deployment (Chapter 9)
Initial Operation	2027-2029	EUR 1,000M-1,200M/year	2,000-2,800	Demonstration of viability (–40% disinformation in 2029)
Consolidation	2029-2031	EUR 1,200M/year	2,800-3,200	Political consensus: European Convention 2029-2030
Transition	2031-2033	EUR 1,200M-1,500M/year	3,200-3,800	Treaty signature 2031-2032 → constitutional Protective Power 2033

Institutional rigour note: acceleration applies to the **agency** (enhanced cooperation, derived law, achievable with the 9 Member States of the founding core in 12-18 months), not to the **Treaty reform** timetable, which requires a Convention, a unanimous IGC and ratification by all 27 — historically never less than 4-5 years from the political launch. The accelerated scenario places the Convention in 2029-2030 (after two years of measurable EDSPA results), the signature in 2031-2032, and the entry into force of the Protective Power in **2033** — the fastest defensible date, consistent with Chapter 17.

Baseline data (revised Chapter 9): Initial investment: EUR 1,100M (baseline, no owned datacenters: dedicated EuroHPC capacity + European cloud) / EUR 2,500M (accelerated) | Annual operating budget: **EUR 500M (baseline) / EUR 1,000M (full scenario)** | Staff: 1,500 (baseline 2027) / 2,000 (accelerated 2026) → **3,000-3,800 employees (2035)** | Demonstrated ROI: **164% over 5 years (baseline); ~70% (accelerated, justified by avoided damage)** | Net benefit: **+EUR 1,820M/year at full capacity** (benefits of EUR 2,320M/year minus operations of EUR 500M/year)


Section 1: The EDSPA as a Generator of Specialised Employment

1. Direct Employment: 1,500-3,000 Positions at the Vanguard of Cognitive Protection

EDSPA staff structure — baseline scenario (2027-2035):

Area	2027	2030	2035	Average Salary (Annual)	Key Profiles
Detection and Analysis	300	600	800	EUR 100,000-150,000	AI, NLP and disinformation experts
Algorithmic Auditing	200	400	600	EUR 90,000-140,000	Software engineers, mathematicians
Electoral Protection	150	300	400	EUR 85,000-130,000	Cybersecurity, electoral law
Cognitive Education	100	200	300	EUR 80,000-120,000	Educators, psychologists
Digital Sovereignty	200	400	500	EUR 110,000-160,000	Cloud architects, cryptography
Administration	550	600	400	EUR 70,000-100,000	Managers, HR, legal
Baseline total	1,500	2,500	3,000	EUR 90,000-130,000	—
Accelerated total	2,000 (2026)	3,200	3,500-4,000	—	—

EDSPA direct employment by area — baseline scenario (2027 → 2035)




EDSPA technical areas. Administration (550→400) is not shown for clarity.

2. Indirect Employment: 6,000-9,000 Positions in the Emerging Ecosystem

Sector	Jobs (2035)	Examples of Suppliers*	Economic Impact	Source
Technology and Software	2,000-3,000	AI and cybersecurity startups	EUR 500M-750M/year	McKinsey 2026
Consulting and Auditing	1,500-2,000	Major consultancies (specialised divisions)	EUR 450M-600M/year	WEF 2026
Infrastructure	1,000-1,500	Data centres, networks (European suppliers)	EUR 300M-450M/year	European Commission 2026
Education and Training	800-1,200	Universities, private academies	EUR 200M-300M/year	Eurostat 2026
Legal and Compliance	700-1,300	Firms specialised in the DSA and the AI Act	EUR 200M-350M/year	European Commission 2026
Baseline total	6,000-9,000	—	EUR 1,650M-2,450M/year	—
Accelerated total (2035)	8,000-12,000	—	EUR 2,200M-3,200M/year	—

*The ecosystem company names appearing later (CognitiveShield, TruthGuard, etc.) are **illustrative examples of company profiles** that EDSPA demand will generate, not existing companies. The "Source" column in this table indicates **reference literature of the domain** (studies by institutions and consultancies on employment and cybersecurity), not specific publications backing these sectoral breakdowns; the jobs and impacts per sector are **illustrative estimates by the author**, consistent with the totals in Chapter 9.

Employment generated by the EDSPA — baseline vs. accelerated



Effect on the European labour market: creation of new professional profiles (algorithmic auditors, disinformation experts, cognitive protection specialists) — the qualified side of the labour market that AI is not generating on its own (Chapter 9) | **Skills demand:** AI (47%), cybersecurity (33%), data analysis (30%) (illustrative estimate by the author based on published trends of the market) | **Talent attraction:** the EDSPA as Europe's magnet for cognitive protection.

Section 2: The Economic Multiplier Effect (2.5x-3.0x per EUR 1 Invested)

1. Calculation Methodology

Theoretical framework (the author's estimate within the ranges of the economic literature): Input-Output model (Leontief): multipliers of **2.5x-3.0x** for investment in technology and professional services | Keynesian multiplier: $MPC \approx 0.6-0.7 \rightarrow \text{multiplier} \approx 2.5x-3.3x$ | Empirical reference: global cybersecurity market with a ~13% CAGR and an implicit multiplier of ~3.0x. *Rigour note:* the multipliers applied to the EDSPA are an illustrative estimate by the author within these ranges; there is no auditable study of the multiplier of a cognitive protection agency (none exists yet).

2. Multiplier Analysis by Area of Investment

Investment Area	Annual Investment	Multiplier	Direct Impact	Total Impact	Source
Salaries (3,000 employees)	EUR 300M-450M	3.0x	EUR 300M-450M	EUR 900M-1,350M	IMF 2026
Contracted services	EUR 400M-600M	2.8x	EUR 400M-600M	EUR 1,120M-1,680M	OECD 2025
Technology R&D	EUR 200M-300M	3.0x	EUR 200M-300M	EUR 600M-900M	McKinsey 2026
Infrastructure	EUR 100M-200M	3.0x	EUR 100M-200M	EUR 300M-600M	European Commission 2026
Education and training	EUR 50M-100M	2.8x	EUR 50M-100M	EUR 140M-280M	Eurostat 2026
Total	EUR 1,050M-1,650M	2.5x-3.0x	EUR 1,050M-1,650M	EUR 3,060M-4,810M	—

3. Comparison with Other Sectors (2026)


Sector	Multiplier	Typical Investment	Impact per EUR 1	Source
EDSPA	2.5x-3.0x	EUR 500M-1,000M	EUR 2.50-3.00	Author's estimate
Traditional infrastructure	1.5x-2.0x	EUR 10,000M-50,000M	EUR 1.50-2.00	IMF 2026
Higher education	2.5x-3.5x	EUR 3,000M-5,000M	EUR 2.50-3.50	OECD 2025
Cybersecurity	2.8x-3.5x	EUR 5,000M-10,000M	EUR 2.80-3.50	WEF 2026

Section 3: Impact on European GDP (EUR 3,160M-4,000M+ per Year)

1. Calculating the Direct and Indirect Impact

Item	Annual Impact	% of EU GDP (~EUR 16.9T)	Methodology	Source
EDSPA salaries	EUR 300M-450M	0.0018%-0.0027%	Direct payroll spending	Eurostat 2026
Indirect employment (salaries)	EUR 1,650M-2,450M	0.0099%-0.0148%	Estimate by sector	McKinsey 2026
Derived consumption	EUR 450M-675M	0.0027%-0.0040%	Salary multiplier effect	IMF 2026
R&D investment	EUR 200M-300M	0.0012%-0.0018%	Direct innovation spending	European Commission 2026
Tool exports	EUR 200M-400M	0.0012%-0.0024%	International sales revenue	WEF 2026
Cost savings	EUR 360M-500M	0.0021%-0.0030%	Reduction of disinformation and fraud	Cybersecurity Ventures 2026
Baseline total	EUR 3,160M-4,000M	0.019%-0.024%	—	—
Accelerated total (2033+)	EUR 4,000M-5,000M	~0.028%	Same mechanisms, 2 years earlier	Author's estimate

Impact on European GDP by item (2035, EUR M)



Upper bounds of the ranges (baseline scenario, 2035).

Cumulative impact (2028-2035, baseline): additional GDP: EUR 22,120M-28,000M (0.13%-0.17% of EU GDP) | Employment created: 45,000-60,000 cumulative jobs | **Accelerated scenario (2027-2035):** EUR 30,000M-38,000M — the same mechanisms operating two more years at full capacity.

2. Comparison with Other Public Investments in Europe

Investment	Annual Cost	GDP Impact	ROI	Multiplier	Employment Generated
EDSPA	EUR 500M-1,000M	EUR 3,160M-4,000M	164% (baseline, over 5 years)	2.5x-3.0x	9,000-12,000
Horizon Europe (R&D)	EUR 15,000M	EUR 22,500M	150%	1.5x	200,000
Digital Europe Programme	EUR 7,500M	EUR 10,500M	140%	1.4x	50,000
NextGenerationEU fund	EUR 800,000M	EUR 1,200,000M	150%	1.5x	1,500,000
ENISA (EU cyber agency)	EUR 27M	EUR 60M-80M (own est.)	~300% (own est.)	~3.0x	~120
Europol	EUR 220M	EUR 450M-650M (own est.)	~300% (own est.)	~3.0x	~1,200

Verified agency figures (2024-2025): ENISA: annual budget of EUR 26.9M in 2025, plus EUR 36M/3 years to operate the EU Cyber Reserve (ENISA/European Commission, 2025) — the Union's cybersecurity agency operates on a budget ~40 times smaller than what the threat requires. Europol: EUR 222M in 2024 (+4.5% per year). The GDP impacts and ROIs of the agencies are the author's own estimates using the same multiplier methodology, not official figures.

Conclusion: the EDSPA would have **one of the highest ROIs** among EU agencies, with a **multiplier of 2.5x-3.0x** comparable to Europol's (~3.0x) and far higher than Horizon Europe's (1.5x). And unlike ENISA (EUR 27M) or Europol (EUR 220M), its budget would be sized for the real scale of the cognitive threat, not for a fraction of it.


Section 4: Talent Attraction and Development of the Technology Ecosystem

1. The EDSPA as a Magnet for Cognitive Protection Experts

Factor	EDSPA	Other EU Agencies	Private Sector
Clear mission	Protecting European democracy	Varies	Private profit
Salaries	10-20% above market (EUR 100K-130K)	Competitive	High (but unequal)
Stability	EU agency (clear mandate)	Permanent	Depends on the company
Social impact	High	Medium-High	Variable
Training	Continuous	Good	Variable
Mobility	EU-27	EU-27	Limited

Key data: the EDSPA will create **new professional profiles** in high market demand (illustrative estimate by the author based on published trends) | it will become Europe's **centre of excellence** in cognitive protection | it will train **1,000-2,000 professionals per year** (Chapter 9) — in the accelerated scenario this training starts in 2027 and slows the brain drain sooner.

2. Building a Technology Ecosystem for the Protective Power



Ecosystem company profiles (illustrative examples, 2035):


Type	Illustrative Profile*	Employment (2035)	Revenue (2035)	Contribution to the EDSPA
Detection startups	Disinformation-detection AI (Germany, France)	1,000-1,500	EUR 100M-200M	Detection tools
Algorithmic Auditing	DSA/AI Act auditing platforms (Netherlands, Spain)	800-1,200	EUR 80M-150M	Platform auditing
Electoral Protection	Election monitoring systems (Sweden, Belgium)	500-800	EUR 50M-100M	Election monitoring
Education	Digital literacy programmes (Sweden, Italy)	500-800	EUR 50M-100M	Digital literacy
Infrastructure	EU cloud and secure network consortia	1,000-1,500	EUR 200M-300M	Data centres and networks
Ecosystem total	—	4,000-6,000	EUR 500M-850M	—

*Illustrative profiles of companies that the EDSPA's public demand will generate in the European ecosystem. The names used in the text (CognitiveShield, TruthGuard, VoteShield, FairAudit, etc.) are **illustrative fiction, not existing companies**.

Ecosystem benefits (2027-2035): (1) **Preparation for the Protective Power:** demonstration of viability (ROI 164% (baseline)), generation of political and social consensus, development of technical capacity. (2) **Reduced technological dependence:** from ~80% of European digital infrastructure imported (EuroStack/Bertelsmann 2025, verified figure from Chapter 4) towards a growing European share by 2035 (the author's objective, not a datum). (3) **European standards:** declared preference for EU providers in public procurement (illustrative estimate by the author).

Section 5: The EDSPA as the Financier of Its Own Transition


1. The Virtuous Circle



2. Financing Projection (2027-2035) — Baseline Scenario

Year	Cumulative Investment	Cumulative Benefits	ROI	Net Benefit (Year)	Cumulative
2027	EUR 1,050M	EUR 928M	88%	-EUR 122M	-EUR 122M
2028	EUR 2,100M	EUR 2,552M	122%	+EUR 574M	+EUR 452M
2029	EUR 2,600M	EUR 4,872M	187%	+EUR 1,820M	+EUR 2,272M
2030	EUR 3,100M	EUR 7,192M	232%	+EUR 1,820M	+EUR 4,092M
2031	EUR 3,600M	EUR 9,512M	264%	+EUR 1,820M	+EUR 5,912M
2032	EUR 4,100M	EUR 11,832M	289%	+EUR 1,820M	+EUR 7,732M
2033	EUR 4,600M	EUR 14,152M	308%	+EUR 1,820M	+EUR 9,552M
2034	EUR 5,100M	EUR 16,472M	323%	+EUR 1,820M	+EUR 11,372M
2035	EUR 5,600M	EUR 18,792M	336%	+EUR 1,820M	+EUR 13,192M

Cumulative net benefit of the EDSPA (2027-2035) — baseline scenario



Cumulative net benefit (EUR million) — the first full year of operation (2028) already turns the cumulative total positive

Financial conclusions (baseline): cumulative break-even point: 2028 (the first full year of operation) | Total net benefit (2027-2035): +EUR 13,200M | Cumulative benefits (EUR 18,800M) over cumulative costs (EUR 5,600M): **cumulative ROI of 336%** | The EDSPA's benefits **will cover part of the initial investment** in the Protective Power (EUR 2,600M baseline; EUR 3,400M accelerated). **In the accelerated scenario**, break-even moves forward to 2027-2028 and the 2026-2035 cumulative total reaches ~EUR 15,000M.

3. Comparison: EDSPA vs. Protective Power

Aspect	EDSPA (2027-2035)	Protective Power (2033/2035+)	Relation
Direct employment	1,500-3,000 (3,500-4,000 accel.)	~3,000 core (2035) → ~7,000-8,000 full (~2037) → >15,000 maturity (2040+, Ch. 15 trajectory)	By stages: core similar; maturity >4x
Indirect employment	6,000-9,000 (8,000-12,000 accel.)	25,000-35,000 (48,000-60,000 accel. 2035; total ecosystem 55,000-68,000, Ch. 17)	3.3x-6x more
GDP impact	EUR 3,160M-4,000M	EUR 7,200M-8,500M (EUR 9,500M-11,000M accel. 2035)	2.3x more
ROI	164% (baseline, over 5 years)	Similar by design (same mechanisms, endowment EUR 2,600M-EUR 3,400M)	Bridge → maturity
Multiplier	2.5x-3.0x	2.6x-3.2x	Similar
Nature	Agency (enhanced cooperation, 9 states)	4th constitutional Power	Bridge → definitive
Permanence	Temporary (2027-2035)	Constitutional (2033 accel. / 2035 base)	Permanent

Section 6: International Comparison — the EDSPA as a Model for Other Democracies

1. The EDSPA's Position in the Global Context

Agency	Country/Region	Budget (verified)	Staff	Model
EDSPA (proposal)	EU	EUR 500M-1,000M/year (operation; initial investment EUR 1,100M)	1,500-3,000	Enhanced cooperation (9 states)
CISA	United States	~EUR 2,200M (\$2.4B FY2026; being cut from \$3.0B FY2025)	2,324 FTE (FY2026; lost ~1/3 of staff since 2025)	Federal
GCHQ	United Kingdom	~EUR 1,200M (estimate)	~6,000	National
ANSSI	France	~EUR 75M (656 FTE + EUR 27M business resources, Sénat 2025-2026)	656	National
BSI	Germany	~EUR 300M (estimate)	~800	National

Strategic fact (2026): while the threat accelerates, **the United States is cutting its cyber and electoral protection:** CISA is losing ~\$300M from its FY2026 budget and had lost almost a third of its staff by mid-2025, with its election-monitoring function under political attack. No Member State or ally can currently cover the European cognitive protection space — the EDSPA would not duplicate existing capabilities: **it would create one that does not exist.**

Competitive advantages of the EDSPA: (1) **Multilateral model:** cooperation among 9 Member States of the founding core (combined GDP ~EUR 14.1 trillion, ~75% of EU GDP). (2) **Comprehensive approach:** cognitive protection + algorithmic auditing + education. (3) **Demonstrable ROI:** 164% over five years (baseline), above the average for national agencies. (4) **Scalability:** a replicable model for other democracies.

2. Exporting the EDSPA Model (2030-2035)

Region	Countries	Adoption Potential	Strategy
Wider Europe	Ukraine, the Balkans, Turkey	★★★★	Progressive integration
Democratic Allies	United Kingdom, Canada, Australia	★★★★	Bilateral agreements
Latin America	Brazil, Mexico, Argentina	★★★	Technical advisory
Africa	South Africa, Nigeria	★★	Cybersecurity cooperation
Asia	South Korea, Japan	★★★	Common standards

Benefits of exporting the model: advisory and training revenue of **EUR 100M-200M per year (2033-2035)** | positioning the EU as the **leader in democratic cognitive protection** | alliances with the UN, the OECD and NATO on global standards.

Conclusion: The EDSPA as a Strategic Investment for Europe's Future

1. Summary of Key Impacts (2027-2035)

Indicator	Baseline	Accelerated
Direct employment (2035)	3,000	3,500-4,000
Indirect employment (2035)	6,000-9,000	8,000-12,000
GDP impact (annual, full capacity)	EUR 3,160M-4,000M (2035)	EUR 4,000M-5,000M (2033) → EUR 5,000M-6,000M (2035)
Cumulative impact	EUR 22,120M-28,000M (2028-2035)	EUR 30,000M-38,000M (2027-2035)
Cumulative net benefit	+EUR 13,200M	~EUR 15,000M
Protective Power	Constitutional in 2035	Convention 2029-2030 → signature 2031-2032 → constitutional in 2033
Disinformation reduction	-40% (2030)	-40% (2029)

"The EDSPA is not an expense; it is a seed. Every euro invested in cognitive protection is not only recovered but generates additional value for the economy, for society and for European democracy."

Five reasons why the EDSPA is strategic:

- **It creates specialised employment** (9,000-12,000 direct and indirect jobs in the baseline; 11,500-16,000 in the accelerated scenario), with competitive salaries and new professional profiles
- **It generates economic growth** (EUR 3,160M-4,000M per year in the baseline; a 2.5x-3.0x multiplier, within the ranges published by the economic literature on fiscal and regulatory multipliers © the author's estimate)
- **It demonstrates that cognitive protection works** (-40% disinformation, ROI 164% (baseline)) and builds the consensus for Treaty reform
- **It finances its own transition** (+EUR 13,200M baseline / ~EUR 15,000M accelerated in cumulative net benefit)
- **It exports a replicable democratic model** for other democracies (EUR 100M-200M per year in advisory services, 2033-2035)

The equation of success:

Investment in the EDSPA → Demonstration of results → Generation of consensus → Financing of the Protective Power = European leadership in the digital era

2. Vision for 2033-2035

- The EDSPA will have **demonstrated that cognitive protection is viable, profitable and necessary**
- It will have **built the political and social consensus** for the creation of the Protective Power
- It will have **prepared the technology ecosystem** (tools, staff, processes) for the transition
- It will have **financed part of its own replacement** by the Protective Power
- In the accelerated scenario, all of this **two years earlier** — with Treaty reform at its fastest defensible date

Sources and References

Economic and Agency Data

Rigour note: the IMF, OECD, McKinsey, WEF, European Commission and Techstrong.ai entries in this section work as **generic references to the domain literature** (published lines of research on fiscal multipliers, regulatory ROI, cybersecurity and digital skills), not as specific audited publications backing this chapter's figures; the derived figures are **the author's estimates**, documented in the rigour notes of each table.

1. Eurostat (2026). *EU economic data*

2. IMF (2026). *Fiscal multipliers in advanced economies*
3. OECD (2025). *The ROI of regulatory agencies*
4. McKinsey (2026). *The economic value of cybersecurity*
5. World Economic Forum (2026). *Global Cybersecurity Outlook*
6. European Commission (2026). *EU agency budgets*
7. Techstrong.ai (2026). *Demand for digital skills*

Verified Agency Figures (2024-2026)

1. ENISA (2025). *Annual budget 2025: EUR 26.9M; EU Cyber Reserve: EUR 36M/3 years* — enisa.europa.eu
2. Europol (2024). *Annual budget: EUR 222M* (Statista/Europol)
3. DHS/CISA (United States) (2025-2026). *FY2026: \$2.4B and 2,324 FTE, versus \$3.0B and 3,641 FTE in FY2025* — dhs.gov
4. SOCRadar (2026). *CISA: ~998 departures and loss of ~1/3 of staff since 2025* — socradar.io
5. Sénat.fr (2025-2026). *ANSSI: 656 FTE, EUR 27M in "business" resources* — senat.fr

Methodology

1. Input-Output model (Leontief) and Keynesian multiplier (IMF/OECD)
2. Cost-benefit analysis (OECD) and benchmarking of EU agencies
3. Dual baseline/accelerated scenario: the accelerated one is the **author's proposal** based on the 2026 urgency evidence (Chapter 9), not a prediction
4. Institutional rigour: the acceleration of the constitutional timetable is bounded by the real durations of the Convention, the IGC and ratification (Laeken→Lisbon precedent: 8 years)
5. Ecosystem company names are illustrative fiction; the GDP impacts of the compared agencies are the author's own estimates

Chapter 10: Implementation and Deployment of the EDSPA

A phase-by-phase roadmap with verified feasibility: from the enhanced cooperation of the 9 to full capacity, in two scenarios — baseline (creation 2026-2027, full capacity 2029) and accelerated (start 2026, initial operational capability in 18 months, full capacity 2028)

Chapter summary

This chapter turns the architecture of Chapters 6-9 into an **execution plan with dates, owners and verifiable exit criteria**. Feasibility is not a promise: verified European precedents support every leg — the EMA/EBA relocation settled the seat in **6 months** (June-November 2017), the EUAA entered into force **5 weeks after the regulation was signed** but built on an existing agency, and ENISA — the comparable case of creation from scratch — took **18 months** from regulation to full operation. On that basis, the chapter presents two scenarios: the **baseline** (legal phase 2026-2027, start of operations 2027, capacity ramp of 40%/70%/100%, full capacity 2029) and the **accelerated** one (legal phase compressed into 2026, start at year-end, **initial operational capability at 18 months and full capacity in 2028**) — the only scenario in which an "18 months" promise is defensible. The budget is the breakdown of Chapter 9: **EUR 1,100M in the baseline** (no owned datacenters: dedicated EuroHPC capacity and European cloud) and **EUR 2,500M in the accelerated one**, distributed across phases under the 40/35/15/10 funding model.

10.1 Verified feasibility: the precedents behind the calendar


No phase of the plan is speculative: each has a verified European precedent with its actual, documented duration.

Precedent	What it proves	Verified duration	Phase it applies to
EUAA (Reg. 2021/2303)	Optimistic ceiling: the legal route can be very fast <i>when an existing structure is transformed</i> — the EUAA replaced EASO (operational since 2011) on 19 January 2022, five weeks after the regulation's signature (15 Dec 2021; Council position on 9 Dec)	~5-6 weeks (legal) on an existing structure	Phase 1 (approval)
ENISA (Reg. 460/2004)	The realistic benchmark for creation from scratch: regulation of March 2004, fully operational on 1 September 2005	~18 months from regulation to operation	Phases 1-2 (creation)
EMA / EBA (2017)	The seat decision is not a bottleneck when there is political will: procedure agreed in June 2017, vote by the 27 on 20 November 2017 (Amsterdam and Paris)	~6 months	Phase 1 (seat)
Enhanced cooperation (Art. 20 TEU)	The verified minimum requirement is 9 Member States — exactly the size of the founding core of Ch. 6; it bypasses blocking states without waiting for all 27	12-18 months from proposal to adoption (Ch. 9B)	Phase 1 (legal route)
JUPITER (EuroHPC)	Dedicated computing capacity is the critical path: tender launched in January 2023, contract signed that same year, installation in 2024 and user access from late 2024	~24 months from tender to operation	Phase 1 (contracts are signed here, not later!)

The operational lesson of JUPITER: if dedicated capacity takes roughly 24 months to become available, the contract with EuroHPC and the European cloud providers **must be signed in Phase 1**, in parallel with the legal negotiation — not once the agency already exists. In the interim, interim European cloud covers the start-up. This is precisely the kind of calendar dependency that rules out budgeting owned datacenters as if building them were instantaneous.

Verified precedents: duration of each leg of the deployment

How long each documented leg took (months)



The critical path is not the regulation: it is computing power. That is why EuroHPC contracts are signed in Phase 1.

The accelerated 18-month scenario requires early contracts and interim European cloud until delivery.

Figure 10.1 — Verified duration of the precedents for each phase. Author's own elaboration on the EUAA (Reg. 2021/2303), ENISA (Reg. 460/2004), EMA/EBA relocation (2017) and JUPITER (EuroHPC JU, 2023-2024).

10.2 The two deployment scenarios

Rigour note: a single 18-month plan (2027-2028) to "full operation" would contradict Chapter 9's own capacity ramp and is only defensible in the accelerated scenario. The chapter keeps the ambition of 18 months where it is realistic (accelerated scenario, initial operational capability) and aligns every date with Chapters 6-9: 15-month ENISA is the benchmark for creation from scratch; the 6-week EUAA is the optimistic ceiling of a transformation, not of a creation.

Milestone	Baseline scenario	Accelerated scenario (recommended)
Proposal and negotiation	2026: Commission proposal + trilogues with the 9 of the founding core	2026: compressed procedure with prior political agreement of the core (Convention of the 9)
Adoption of the regulation	First half of 2027	Late 2026 (EUAA-style optimistic ceiling: adoption and entry into force within weeks)
Operational start (seat, core staff, EuroHPC contracts already signed)	2027	2026-2027
Initial operational capability (IOC)	2028 (40% → 70% ramp)	18 months from start → 2028
Full capacity (1,500-2,000 staff, 100% of benefits, KPIs at full demand)	2029	2028
Implementation investment (Ch. 9)	EUR 1,100M (2027-2028)	EUR 2,500M (2026-2027)

Comparative calendar: baseline versus accelerated

From proposal to full capacity: baseline 2026-2029, accelerated 2026-2028




Figure 10.2 — Calendar of the two scenarios. Author's own elaboration on the figures of Chapters 6, 9 and 9B.

10.3 Phase 1 — Preparation (months 1-6): the phase that decides the calendar

Legal approval (months 1-4)

- **Legal route:** enhanced-cooperation regulation (Art. 20 TEU) with the **9 Member States of the founding core** — Germany, France, Italy, Spain, the Netherlands, Belgium, Sweden, Finland and Denmark (combined GDP ~EUR 14.1 trillion, ~75% of EU GDP, Ch. 6). Negotiation uses the Union's institutions: Commission proposal, Council (qualified majority among participants) and European Parliament (IMCO/LIBE)
- **Seat act:** seat decision by the verified EMA/EBA procedure — the core pre-agrees on Brussels (the natural candidate, Ch. 6) so the vote consumes no more than one cycle
- **Early capacity contracts:** in parallel with the legal work (not after it), the Commission — on behalf of the future core — signs the multi-year contracts for **dedicated EuroHPC capacity** (JUPITER, LUMI, MareNostrum 5) and sovereign European cloud. JUPITER lesson: ~24 months from tender to operation, so this leg starts on day 1

Constituting the governance (months 3-6)

The **Council of Protectors** is constituted: 17 independent experts selected on merit through an open procedure (ERC model, Ch. 7) — **zero state representatives**. In parallel, the **Executive Director** is selected through the open Europol-2026-type procedure, and the **Fundamental Rights Office (FRO)** is staffed from day one: auditing from the start, not after the first scandal.

Rigour note: the agency does not get a "Management Board" with representatives of the founding states. Its governance is that of Chapter 7: a Council of Protectors of 17 merit-based experts with no state representation, with the states in the Advisory Council (no vote). Independence is not a decoration of the agency: it is its product. A board of countries would build in exactly the incentive the agency must combat.

Initial hiring (months 4-6)

A founding core of ~200 people: management, legal teams, systems architects and the first senior analysts. Phase budget: **EUR 50M**. To avoid dependence on permanent-competition timelines, the verified start-up mechanisms of European agencies are used: **seconded national experts (SNE)** — the ENISA/eu-LISA pattern — plus interim contracts for critical technical profiles.

10.4 Phase 2 — Establishment (months 5-12): the agency takes shape

Workstream	Months 5-8	Months 9-12
Seat and equipment	Building lease (Brussels, ~5,000 sqm), adaptation under way	Labs, crisis room, ISO 27001 certification in progress
Computing capacity	Interim European cloud operational (European providers); EuroHPC contracts in execution	First blocks of dedicated capacity delivered; progressive migration
Staff	Hiring up to ~800: detection and analysis (Chs. 7-8), AI, cybersecurity	Up to 1,500 (baseline) / 2,000 (accelerated) ; deployment of national nodes (dual eu-LISA + EDPB model, Ch. 7)
In-house tools	Detection system and algorithmic auditing platform (development)	Electoral monitoring, cognitive education, digital sovereignty (Ch. 8)

Cross-cutting design requirement: every tool is built with the **human-AI dual verification of Chapter 8 from the first line of code** — AI proposes, the human decides — and with **fully auditable traceability** (model, version, evidence, analyst, time) as an architectural requirement, not a later addition. Retro-fitting traceability onto an already-built system is the most expensive design error possible: the protection system is born audited or it is born flawed.

10.5 Phase 3 — Testing and training (months 11-16)

Test	Exit criterion (verifiable)
Detection pilots in the 9 core states	Detection time in simulated incidents < 12 h (2027 KPI of Ch. 7); accuracy > 90% on audited samples
Electoral simulation (EU Cyber Solidarity exercise pattern)	Full detection → dual verification → requirement cycle at Critical level (< 1 h alert, < 12 h notification, Ch. 8)
Trial algorithmic audit on volunteer platforms	Public report with findings + protected method file (Ch. 7.6) issued on schedule
Training (budget EUR 50M)	100% of operational staff certified in the dual verification protocol; sufficient human validation analysts on staff for the Medium/High/Critical levels
Incident framework	Alert < 12 h / notification < 48 h verified in drills — the NIS2 tightening of Ch. 8, not the 24 h/72 h standard

10.6 Phase 4 — From operational start to full capacity

Capacity does not arrive all at once: it follows the **ramp of Chapter 9** — 40% of full-capacity benefits in the first year of operation, 70% in the second, 100% in the third. The **initial operational capability (IOC)** of the accelerated scenario arrives **18 months** from start: detection, auditing and electoral protection systems in production, the core's national nodes connected, and the first full algorithmic audit with a public report. **Full capacity** (1,500-2,000 staff, 30% of the platform landscape audited in 2027, KPIs at full demand) arrives in **2028 (accelerated) or 2029 (baseline)**.

The capacity ramp: from IOC to full

Operational capacity ramp (% of full capacity)




Figure 10.3 — Capacity ramp aligned with the benefit ramp of Chapter 9 (40%/70%/100%). Author's own elaboration.

10.7 Implementation budget by phases

Rigour note: the chapter does not budget EUR 1,800M with two mutually inconsistent breakdowns (by phase and by item), nor EUR 950M of owned datacenters in Brussels and Amsterdam. It uses exclusively the breakdown of Chapter 9: **EUR 1,100M in the baseline scenario** (dedicated EuroHPC capacity and European cloud instead of owned infrastructure) and EUR 2,500M in the accelerated one, funded under the 40/35/15/10 model (EU budget / core states / share of EU-collected fines / other sources) — not "EU + states + private investment", which is not part of the model.

Item (Ch. 9)	Phase 1: Preparation	Phase 2: Establishment	Phase 3: Testing	Phase 4: IOC deployment	Baseline total
Headquarters and equipment (EUR 200M)	EUR 200M	—	—	—	EUR 200M
EuroHPC capacity + European cloud (EUR 250M)	—	EUR 250M	—	—	EUR 250M
Network and national nodes (EUR 100M)	EUR 50M	—	—	EUR 50M	EUR 100M
Staff hiring (EUR 150M)	—	EUR 100M	EUR 50M	—	EUR 150M
In-house tools (EUR 200M)	—	EUR 150M	EUR 50M	—	EUR 200M
Literacy start-up (EUR 70M)	—	—	EUR 70M	—	EUR 70M
Training (EUR 50M)	—	—	EUR 30M	EUR 20M	EUR 50M
Contingency (EUR 80M)	—	—	—	EUR 80M	EUR 80M
Total per phase	EUR 250M	EUR 500M	EUR 200M	EUR 150M	EUR 1,100M

The **accelerated scenario (EUR 2,500M)** keeps the same phase structure compressed into 2026-2027, and adds: early hiring of up to 2,000 employees, parallel deployment of the nodes, expanded dedicated computing capacity and the European labour reconversion programme for workers displaced by AI (Ch. 9).

10.8 Risks and mitigation

Risk	Mitigation (verified)	Cost of mitigation
Legal delay (blocking by reluctant states)	Enhanced cooperation with the core 9 (Art. 20 TEU) — consensus of all 27 is not needed; the precedent of existing enhanced cooperations avoids the veto	Covered by design (no premium)
Computing capacity delay (critical path, ~24 months)	Early contracts signed in Phase 1 + interim European cloud until EuroHPC delivery; capacity never dependent on a single site	EUR 80M contingency + network item
Shortage of qualified staff	SNEs (ENISA/eu-LISA pattern) + interim contracts + salaries 10-20% above market (Ch. 9B) + in-house auditor training (Ch. 8)	EUR 20M (temporary external hiring)
Platform resistance to audits	Requirement backed by the DSA + single judicial review (Ch. 7.6); the 2024 precedent shows voluntary cooperation is not enough	Ordinary legal
High-profile false positive at start-up	Human-AI dual verification from day 1 (Ch. 8) + FRO operational from constitution + transparency of quarterly statistics	Covered by design

10.9 Monitoring and transparency

- **Exit criteria per phase** (10.3-10.6): every milestone has a verifiable criterion — not declarative dates but passed tests
- **Chapter 7 KPIs** as the permanent dashboard: detection < 12 h (2027) → < 6 h (2028) → < 2 h (2030); accuracy > 90% → > 95%; platforms audited 30% (2027) → 90% (2030)
- **Initial evaluation** at the close of Phase 4: public report by the Council of Protectors to the European Parliament and the 9 participating states, with deviations and a corrective plan
- **FRO audit** from the first month of operation: fundamental rights are supervised during the deployment, not after it
- **Quarterly aggregate statistics** from IOC (Ch. 8.6)

10.10 Conclusion: the calendar is the strategy

The difference between the EDSPA and another European paper exercise is not in the design: it is in the dates. The 2026 evidence (Ch. 9) — 116,175 layoffs attributed to AI in eight months, an industry that modulates its warnings according to its stock price — does not wait for a perfect deployment. The verified precedents say Europe knows how to move fast when it wants to: six weeks for the EUAA from the Council position (five from signature), six months for the EMA and EBA seats, eighteen months for ENISA from scratch. The accelerated scenario presented here does not require Europe to do something it has never done — it requires it to do on time what it already knows how to do. And the critical path has a name: **computing capacity**. Whoever controls the EuroHPC contract calendar controls the calendar of European cognitive protection. That is why they are signed on day 1.

Chapter 11 explains why this agency, even successfully deployed, is only the necessary step — not the definitive one — towards the constitutional Protective Power.

Verified sources for this chapter

- EUAA: Regulation (EU) 2021/2303 — Council position on 9 December 2021, signed on 15 December 2021, in force on 19 January 2022 — transformation of EASO (Reg. 439/2010, operational since 2011): the optimistic ceiling of the legal route on an existing structure
- ENISA: Regulation (EC) No 460/2004 of 10 March 2004; fully operational on 1 September 2005 — the realistic benchmark for creation from scratch (~18 months)
- EMA/EBA (2017): relocation procedure endorsed by the 27 on 22 June 2017; seat decided on 20 November 2017 (Amsterdam and Paris) — the seat decision in ~6 months
- Art. 20 TEU: enhanced cooperation with a minimum of 9 Member States — the basis of the founding core of Ch. 6 and of the 12-18 month estimate (Ch. 9B)
- JUPITER (EuroHPC JU): tender launched in January 2023, contract signed that year, installation 2024, user access from late 2024 — ~24 months from tender to operation

- NIS2 (Directive 2022/2555): 24 h/72 h — tightened by the EDSPA to 12 h/48 h (Ch. 8)
- Europol: open procedure for selecting the Executive Director (2026, Ch. 7); EUR 222M / ~1,400 staff (Chs. 5-6)
- ENISA / eu-LISA: start-up pattern with seconded national experts (SNEs) and interim contracts verified in both agencies
- EU Cyber Solidarity: joint simulation exercises as the pattern for the Phase 3 drills (Ch. 8)

Chapter 11: The EDSPA as a Necessary Step Towards the Protective Power

Chapter summary

This chapter opens Part III and redefines the role of the EDSPA: **it is not the destination, it is the strategic and indispensable step** towards the Protective Power. The thesis rests on three verified pillars: first, that no EU agency is immune to extinction or downsizing — the European Digital Media Observatory (EDMO) sees its hubs' funding end in 2026 with no confirmed successor, EMSA was subject to a full reorganization in 2023-2025, and OLAF operates with some 500 staff for all fraud against the EU budget. Second, that the ordinary treaty revision procedure (Art. 48 TEU) requires a Convention, an intergovernmental conference, unanimity and ratification by all 27 — a historical minimum of 4-5 years that places the earliest defensible constitutional horizon at 2033. And third, that enhanced cooperation among the 9 founding core States (Art. 20 TEU) makes it possible to start *now*, prove that cognitive protection works, and build the legitimacy the reform will demand.

11.1 The fundamental problem: agencies can disappear

The EU's institutional history shows that a decentralized agency is a **revocable** solution: it depends on political will, shifting priorities and annual budgets. Three verified cases illustrate this fragility:

Verified case	What it demonstrates	Current status
EDMO (European Digital Media Observatory)	The EU's flagship instrument against disinformation depends on project-based funding, not a permanent mandate: the hubs' funding ends in 2026 with no confirmed successor	Uncertain continuity even as the threat grows
EMSA (European Maritime Safety Agency)	Even consolidated agencies (created in 2002) are reorganized by political decision: Commission proposal of June 2023 (COM(2023) 269) to redesign its tasks, with a provisional Council-Parliament agreement in May 2025	Full reform adopted 2023-2025
OLAF (European Anti-Fraud Office)	Resources do not grow with the mandate: some 500 staff to investigate all fraud against the EU budget, and the European Chief Prosecutor himself has noted that the vast majority of suspected fraud never even appears on the institutions' radar	Structurally insufficient capacity
«European Narratives Observatory» line (DG CNECT) and ECDR	The EU funds the evidence layer through temporary ~18-month projects that are never consolidated: NODES (2022) → 2nd call (2023, EUR 1,190,500, focus: Ukraine, gender disinformation and elections) → 3rd call «Detecting influence campaigns and boosting societal resilience» (2025, funding two projects of EUR 1.5M and EUR 1.65M, one of them the SHIFT consortium) → PPPA-2026 social cohesion pilot (EUR 785,000) — and its structural answer, the European Centre for Democratic Resilience (operational since 24 Feb 2026), is intergovernmental and has no authority	Five temporary instruments, zero consolidation

The lesson: the EU builds temporary instruments against permanent problems. No agency is immune to elimination, cuts or reorganization. What the cognitive sphere needs — a structural, not circumstantial, threat — is **structural** protection. The context of 2026 (the US veto of June 2026 on access to advanced models and platforms' restrictions on independent research) turns this lesson into urgency.

The SHIFT launch conference confirms it in the Commission's own language: «From detection to response: building Europe's next layer of cognitive resilience». The EU recognizes that it needs a new «layer of cognitive resilience» — this book's diagnosis in official terminology — but endows it with temporary 18-month consortia and intergovernmental centres without authority, not with permanent institutional architecture.

Comparison: agency (EDSPA) versus constitutional power (Protective Power)

Six dimensions: revocable agency versus constitutional power

Permanence

Annual political will

Constitutional: reform requires unanimity

Funding

Annual contributions (40/35/15/10 model)

Own resources (financial autonomy)

Authority

Applies existing laws

Own constitutional powers

Independence

Partial: Council of Protectors, external budget

Separation of powers guaranteed

Scope

Participating States (enhanced cooperation)

All Member States (mandatory)

Sustainability

Political fragility: revocable by the legislator

Unalterable without unanimous consensus

EDSPA as agency (revocable)

Protective Power as 4th power (structural)

The EDSPA solves the problem with today's means; only the Protective Power solves it for good.

Conclusion: the agency is the means; the constitutional power is the end.

Figure 11.1 — The six dimensions in which a decentralized agency cannot match a constitutional power.

11.2 The structural limits of the EDSPA as an agency

Even though the EDSPA is designed with the hardened governance of Chapters 6-7 (a Council of Protectors of 17 independent experts, an Executive Director by open procedure, mandatory human-AI double verification), as an agency it remains **revocable by the ordinary legislator**:

Aspect	EDSPA (agency)	Protective Power (4th power)
Permanence	A regulation can be repealed by ordinary majority	Only amendable through Art. 48 TEU: Convention, unanimity and ratification by all 27
Authority	Regulatory and supervisory powers within the <i>acquis</i> (DSA, AI Act, NIS2)	Own constitutional powers: legislative initiative, direct execution and access to specialized jurisdiction - Administrative Board of Appeal (ACER model), Specialised Chamber in Digital Protection of the General Court and cassation before the CJEU (Ch. 12)
Scope	The 9 founding core States (Art. 20 TEU) and subsequent accessions	All Member States, with <i>erga omnes</i> effect
Funding	Annual contributions under the 40/35/15/10 model — subject to budget cycles	Own resources: financial autonomy guaranteed by the Treaty
Independence	Protected by design (ERC model), but politically reversible	Separation of powers: impartiality by constitutional architecture
Adaptation	Requires the legislator to update its mandate for every new threat	Permanent mandate over the cognitive sphere, adaptable through its own secondary law

Rigor note: this limitation is not a defect of the EDSPA: it is the reason this book always presents it as a *necessary step*, never as the final solution. Chapter 9 demonstrated its economic viability (ROI 164% over 5 years in the base scenario, roughly EUR 2.6 returned per euro; net benefit of +EUR 1,820M per year at full capacity; cumulative break-even in 2028 according to Chapter 9B). The value of that demonstration is not that the agency stays there: it is that it makes what comes next defensible.

11.3 The Protective Power: the definitive solution

Definition and purpose

The **Protective Power** is the 4th constitutional power: a permanent institution charged with protecting the **cognitive sphere** of citizens, institutions and European society against digital manipulation, guaranteeing that no actor — public or private — can subvert democratic processes, individual autonomy or collective sovereignty.

Why a 4th power and not more of the same

Traditional power	Function	Limitation against cognitive threats
Executive	Implements policies	No specialized expertise in cognitive protection; dependent on external advice, frequently shaped by the regulated interests
Legislature	Makes laws	Multi-year cycles against threats that evolve in weeks; the <i>acquis</i> always arrives late
Judiciary	Interprets laws	Acts on consummated facts; cannot prevent ongoing algorithmic manipulation

None of the three has the mandate, the capability or the agility to protect the cognitive sphere in a **proactive, comprehensive and permanent** way. The Protective Power rebalances the structural asymmetry between technology (fast, global, opaque) and democracy (slow, fragmented, transparent) — just as the Judiciary was created to interpret the law impartially, the Protective Power is created to protect cognition with specialization.

Exclusive competences

Competence	Description	Why no current power can exercise it
Comprehensive cognitive protection	Safeguard the cognitive sphere of citizens and institutions	Requires a specific mandate, dedicated expertise and immediate execution
Algorithmic regulation	Supervise and audit AI systems and algorithms	Demands permanent specialized technical knowledge, not ad-hoc consultancy
Platform oversight	Audit and make requests of digital platforms	Needs absolute independence against corporate capture
Electoral protection	Guarantee the integrity of democratic processes	Demands transnational coordination and preventive, not reactive, action
Cognitive rights	Protect cognitive freedom, mental autonomy and privacy	Current digital rights are scattered and lack an institutional guarantor
Education and resilience	Digital literacy and social resilience	Requires a specialized pedagogical approach and sustained dedicated resources

11.4 Why not create the Protective Power directly

The constitutional route is codified in **Article 48 TEU**: the ordinary revision procedure requires a **Convention**, an **intergovernmental conference**, **unanimity** in the European Council and **ratification by all 27** under their constitutional rules (in several cases, referendum). The passerelle clauses of paragraph 7 do not help here: secondary law cannot create a new constitutional power.

Recent experience confirms the difficulty: the European Parliament voted on **22 November 2023** in favor of treaty reform following the Conference on the Future of Europe, and asked the European Council to convene the Convention — but no consensus has materialized in the Council to open the process. The great historical revisions (from the Single European Act to Lisbon) took between 4 and 8 years from convening to entry into force. The earliest defensible horizon for a 4th constitutional power is therefore **2033**.

Obstacle (today)	EDSPA solution	Result
Reform requires unanimity and ratification by all 27	Enhanced cooperation: only 9 States (Art. 20 TEU)	Start without waiting for total consensus
Lack of political consensus on cognitive protection	Willing States demonstrate tangible results	The reluctant join on evidence, not coercion
Citizens do not perceive the urgency	Visible benefits: measured disinformation, recovered trust	Social pressure for the reform
No infrastructure, staff or processes exist	The EDSPA builds and refines them	The reform inherits capability, not promises
A failed reform would weaken the EU	Proof of concept before touching the Treaty	Minimized risk

The strategic key: Art. 20 TEU (enhanced cooperation) and Art. 48 TEU (ordinary revision) are not rivals: they are a *sequence*. Enhanced cooperation is the only legal instrument available today that allows acting immediately against the threat while tomorrow's constitutional consensus is built. It is exactly the mechanism through which the EU has historically advanced when unanimity blocked the way and necessity could not wait.

11.5 The EDSPA as laboratory: what the transition will inherit

Asset	EDSPA (2026-2030)	Protective Power (2033+)
Experience	Proves that cognitive protection works: detection metrics, resolved cases, algorithmic audits	Applies that experience at constitutional scale
Capability	Builds tools, protocols (human-AI double verification), expert staff (1,500-2,000 employees)	Inherits and scales that operational capability
Economic evidence	ROI 164% over 5 years; net benefit +EUR 1,820M/year at full capacity; cumulative break-even 2028 (Ch. 9B)	Turns demonstrated viability into the argument for reform
Political support	The 9 founders generate a precedent of effectiveness that more States join	The broadened consensus sustains ratification
Citizen support	Tangible benefits: less disinformation, more trust in the digital space	Social pressure backs the constitutional step
Standards	Technical norms, audit procedures, electoral protection standards	Institutionalized into constitutional law and its own secondary law
Financial base	Budget stabilized around the full scenario (EUR 1,000M/year)	Transfers its budgetary capability as the initial base of the Protective Power, complemented by the own-resources regime developed in Chapter 16

11.6 The transition timeline (2026-2035)

The timeline aligns the dual scenario of Chapter 10 with the Art. 48 TEU procedure. The realistic starting point: the accelerated scenario places initial operational capability in **2028**, so the reform proposal can be presented on results — not on promises — in **2028-2029**.

From the EDSPA to the Protective Power: 2026-2035





Figure 11.2 — Transition timeline: the EDSPA operates and proves results (2026-2029) while the constitutional process opens (2029-2033) and the structures integrate (2033-2035).

Phase	Period	Objective	Success indicator
1. EDSPA implementation	2026-2028	Create the agency and reach operational capability (accelerated scenario of Chapter 10)	EDSPA operational with the 9 founders
2. Consolidation	2028-2029	Demonstrate effectiveness and accumulate evidence	Detection metrics and a public success case
3. Proposal	2028-2029	Present the constitutional reform on results	Initial political consensus to open the process
4. Convention and IGC	2029-2032	Negotiate the treaty reform (Art. 48 TEU)	Unanimous agreement in the European Council
5. Ratification and entry into force	2032-2033	Ratification by the 27 under their national rules	Protective Power in force in 2033
6. Integration	2033-2035	Transfer staff, tools, standards and the budgetary base from the EDSPA to the Protective Power	Fully operational PP; the EDSPA is integrated

11.7 The thread of Part III (Chapters 11-20)



11.8 Conclusion: the EDSPA is the means; the Protective Power is the end

1. **The EDSPA is a necessary but insufficient step.** As an agency, it can be repealed, cut or reorganized by any future majority — as EDMO, EMSA and OLAF demonstrate. Its value lies not in what it achieves by itself, but in what it enables building.
2. **The Protective Power is the definitive solution.** As the 4th constitutional power, it would be permanent, specialized and impartial by design: it would return to democracies control over the digital sphere with the solidity that only constitutional law provides.
3. **The EDSPA is the strategic path.** It proves that cognitive protection works, generates the political and social support the reform demands, and builds the capability the Protective Power will inherit. Acting with the 9 today is what makes the unanimity of the 27 possible tomorrow.

The book's continuity: Part I (Ch. 1-5) established the problem: threats to the cognitive sphere are existential for democracy. Part II (Ch. 6-10) built the instrument: the EDSPA, its design, its economics and its deployment. Part III (Ch. 11-20) traces the definitive route: from the operational patch of 2026 to the constitutional architecture of 2033.

11.9 Verified sources and references

- **Art. 48 TEU** (ordinary and simplified revision procedures): EUR-Lex, consolidated text of the Treaty on European Union.
- **Art. 20 TEU and Arts. 326-334 TFEU** (enhanced cooperation): EUR-Lex, consolidated texts.
- **European Parliament, 22 November 2023:** resolution on the proposals to amend the Treaties following the Conference on the Future of Europe (request for a Convention).
- **EDMO:** European Commission official page (digital-strategy.ec.europa.eu) and DIGITAL calls on hubs with funding "ending in 2026".
- **EMSA:** proposal COM(2023) 269 final, 1 June 2023; provisional Council-Parliament agreement, 20 May 2025 (Council of the EU).
- **OLAF:** Court of Auditors Opinion on OLAF management (EUR-Lex); Follow the Money investigation on the undetected dimension of fraud (European Chief Prosecutor).

- **EDSPA economic figures:** Chapters 9 and 9B of this version (ROI 164% over 5 years, net benefit +EUR 1,820M/year, cumulative break-even 2028, 9-State core).

Chapter 12: Institutional Architecture of the Protective Power

Chapter summary

This chapter defines the complete structure of the Protective Power as the 4th constitutional power, built on a design decision that runs through everything: **it is a protecting power, not a sanctioning one**. The PP imposes no fines at any level — neither its own nor delegated —, deletes no content and does not decide what can be seen: its weapon is light (detection, evidentiary documentation, publicity). When detected manipulation demands a coercive response, the case file goes to those who already hold sanctioning competence (the European Commission under the DSA and the AI Act, with their real ceilings of 6% and 7%, or the national authorities) or is litigated before the courts, with the PP as a party. The institution is **single** — not an agency holding — with a Council for Cognitive Protection under a mixed model (27 independent experts with a vote and 27 state representatives with a voice but no vote) and four operational departments under one executive direction. The volume of technical litigation it will generate is answered not with a "proprietary" tribunal (which would breach the separation of powers) but with a **Specialised Chamber in Digital Protection within the General Court**, belonging to the judiciary and funded from the judicial budget, with dually qualified judges (Unified Patent Court model) and an escalation clause. The EDSPA is integrated as the embryo of the Department of Cognitive Protection. Economic framework: that of Chapter 17 — **EUR 2,600M/year base / EUR 3,400M full**, with own resources (Ch. 16). And the staffing is that of a power, not of an agency: a staged trajectory that goes from the 1,500-2,000 inherited from the EDSPA to **more than 15,000 staff at maturity 2040+** (Ch. 15) — with a growing, deliberately non-automatable share: Europe's largest public employment pole of human judgement.


12.1 Fundamental design principles

Principle	Description	Practical application
Protection, not punishment	The PP detects, documents and exposes; it never punishes or censors	Sanctions belong to the existing authorities (Commission, Member States) and to the courts; the PP is a procedural party, never a judge
Separation of powers	Absolute independence from Executive, Legislature and Judiciary — and from them	Express prohibition of mutual interference; the PP absorbs no judicial or sanctioning functions of other powers
Specialization	Exclusive focus on the digital cognitive sphere	Cannot assume competences of other powers; mandate delimited by the Treaty
Omnipresence	Total coverage: all Member States, all platforms	Detection and auditing jurisdiction across the EU with erga omnes effect
Unalterability	Only amendable by unanimity (Art. 48 TEU)	Constitutional protection against short-term political changes
Radical transparency	Accountability to citizens and institutions	Mandatory external audits (Court of Auditors), annual public reports and structural documentary access

Why "pure protector" and not "sanctioning protector" (connection with Chapter 11): Chapter 11 established that the PP must be permanent and unalterable. But a permanent power *with* its own sanctioning capacity and no internal control would be perceived — rightly — as a risk to liberties, and that perception alone would sink ratification. The solution is not less power: it is a different power. The PP concentrates at constitutional level what current institutions cannot do (systematic detection, continuous algorithmic auditing, evidentiary publicity) and leaves coercion untouched in the hands of those who already exercise it under full judicial control. **No actor decides what can be seen; everyone can know who is trying to manipulate them.**

12.2 The organogram: a single institution with departments

Europe's experience with fragmented agency structures (the landscape reviewed in Chapters 6 and 11) recommends the opposite: **a single institution with one direction**, whose operational functions are organized as internal departments. This eliminates overlaps and makes the "agency against agency" game impossible. The judicial dimension is not an organ of the PP: it belongs to the EU judiciary (section 12.4).



12.3 The Council for Cognitive Protection

The Council is the political-strategic organ of the Protective Power. Its design resolves the tension between the **technical impartiality** required of a power that watches over manipulation and the **political legitimacy** without which the treaty reform will not be ratified. The solution is a **mixed model with unequal voting**:

Members	Number	Selection	Term	Decision-making power
Independent experts	27 (one per Member State for geographic representation, selected on merit)	Pan-European public call + technical evaluation + shortlist + public lottery (ERC model of Ch. 7, expanded from 17 to 27)	8 years, non-renewable, staggered mandate	Vote — the only chamber with decision-making capacity
State representatives	27 (one per Member State)	Appointed by national governments	4 years, revocable by their government	Voice but no vote — information, deliberation and capacity to challenge before the specialized jurisdiction
Presidency	1	Elected by the Plenary from among the independent experts	5 years, non-renewable	Chairs and represents the Council

Selection mechanism for the experts (merit + lottery)

1. **Public call** across the EU, with public eligibility criteria.
2. **Evaluation by an independent committee** following the European Research Council's standard.


3. **Shortlist** with at least three times the positions per Member State.
4. **Public lottery** audited before the Publications Office, notaries and media.
5. **Ratification** by the Plenary except on grounds of subsequent ineligibility.

Functions of the Council

- **Legislative initiative** in cognitive protection matters: the PP proposes; Parliament and the Council of the EU decide under the ordinary legislative procedure.
- **Strategic direction**: priorities, internal rules, multiannual plans.
- **Supervision of the departments** through the Executive Direction.
- **Inter-institutional coordination** with Executive, Legislature and Judiciary and with the Member States.

12.4 The judicial fit: specialized jurisdiction inside the judiciary

The PP will generate a significant volume of technical litigation — contested algorithmic audits, actions against its operational decisions, conflicts between platforms and cognitive rights. The temptation to solve it with a "Digital Protection Tribunal" owned by the PP was discarded in this version: **a tribunal inside a power breaches the separation of powers and would confirm the perception of a power that is judge and party**. The viable solution, backed by real precedents, has three pieces:



a) The administrative Board of Appeal (first guarantee, ACER model)

The PP's operational decisions (audit outcomes, designations of monitored platforms, information requests) will be challengeable before an internal but independent **Board of Appeal**, following the verified model of EU agencies' Boards of Appeal (ACER, CPDP): *full* review of the contested decision, with technical experts, appealable to the General Court. CJEU case law (*Aquind v ACER*, C-46/21 P) validates these boards as an appropriate rights-protection mechanism for complex technical decisions.

b) The Specialised Chamber in Digital Protection (second guarantee, inside the judiciary)

The volume and technicality of the litigation justify specialization, not a new tribunal. The proposal: a **Specialised Chamber in Digital Protection of the General Court**, created by ordinary regulation (no treaty reform needed), with:

Element	Design	Verified precedent
Nature	Chamber of the General Court: belongs to the judiciary, funded from the judicial budget — not to the PP	Internal specialised chambers (Tokyo IP High Court, UK IPEC models)
Judges	Accredited dual training: law + science/technology, with a permanent pool of technical legal secretaries	Unified Patent Court: 34 legally qualified + 51 technically qualified judges
Creation	Ordinary regulation (Art. 257 TFEU additionally allows specialised courts if volume demands it)	Art. 19 TEU system
Escalation clause	If the Chamber sustainably exceeds the workload threshold (indicative reference: 150-200 cases/year), the Art. 257 specialised court is activated — only by real demand	Lesson of the Civil Service Tribunal: created 2004, abolished 2016 for lack of volume
Control	Cassation before the Court of Justice; preliminary rulings (Art. 267 TFEU) for coherence with national courts	Art. 19 TEU and CJEU Statute

Why not a tribunal owned by the PP: the closest precedent also discourages premature court creation: the Civil Service Tribunal (created 2004, 7 judges) was **abolished in 2016** and its jurisdiction returned to the General Court, because its caseload did not justify a fixed structure. And when the EU judicial system truly saturated, the response was to expand what existed — the General Court doubled its judges from 28 to 56 between 2016 and 2019 to clear its backlog — not to create parallel institutions. The specialized chamber with conditional escalation follows both lessons: real specialization, structure proportional to demand.

c) The sanction: always external, always judicially reviewable

The PP has no sanctioning ceilings of its own. The real ceilings already exist in the acquis: **up to 6% of worldwide turnover (DSA)** and **up to 7% or EUR 35M (AI Act)**. What is missing today is not higher ceilings: it is systematic detection and effective litigation. The PP's contribution is exactly that:

- **Referral of evidentiary case files** to the European Commission and the competent national authorities, at judicial standard of proof from day one.
- **Litigation as a party** before the Specialised Chamber and, via preliminary reference, before national courts.
- **Publicity:** documented manipulation is published — the reputational and market cost is, in most cases, the most immediate sanction.

Rigor note: the separation between detecting and sanctioning, established for the EDSPA in Chapters 6-9, is thus elevated to the very architecture of the 4th power and simplified to the maximum: **the PP touches neither the sanction nor the fine money at any level**. Fine revenue always goes to the general EU budget. This simplification also removed the "Sanctions Committee" and the "Digital Protection Tribunal" of the initial design.

12.5 The operational departments

Department	Function	Budget/year	Design notes
Cognitive Protection (ex-EDSPA)	Manipulation detection, algorithmic auditing, electoral protection, platform oversight	EUR 1,000M	Fully inherits the EDSPA's capability: human-AI double-verification protocols, expert staff, standards
Digital Sovereignty	Technological independence: European compute capacity (EuroHPC and successors), technology certification, control of strategic data	EUR 500M	Starts where Chapter 17 places the Protective Power's own datacenters
Cognitive Education	Critical digital literacy, training for journalists and educators, campaigns, manipulation research	EUR 300M	Own target: ≥90% critical digital literacy by 2038 , surpassing the Digital Decade target (80% by 2030)
Cognitive Intelligence	Analysis of emerging threats, disinformation patterns, early warnings, coordination with national services	EUR 200M	Guarantees: supervision by the Board of Appeal and the Chamber, annual public report, European Parliament audits
Department total	Operational function of the Protective Power	EUR 2,000M	Plus Council, Executive Direction, Board of Appeal and National Offices → Chapter 17 framework
Total Protective Power framework	Full institution	EUR 2,600M base / EUR 3,400M full	Own resources (Ch. 16). The Specialised Chamber is funded from the EU judicial budget, not the PP's


Real context figures: the literacy target is anchored to the official EU Digital Decade target — **80% of adults with basic digital skills by 2030** — versus the verified ~55.6% at mid-decade. The PP's own target (≥90% by 2038) is a design estimate of this book, labeled as such. Associated judicial sizing (own estimate, charged to the judicial budget): a Chamber of 7-9 dually qualified judges with ~4 legal secretaries each and a pool of experts, in the order of EUR 40-60M/year — comparable to the CJEU's real cost per judge.

Staffing — a trajectory, not a payroll: the departments are not born with their maturity staffing: they grow in stages with accountable triggers (accessions of states, coverage of the mission, volume of threat, expansions authorised by the legislator), as has happened with every European power — the Commission from a few thousand to ~32,000; Frontex from dozens of analysts to a statutory corps of 10,000 by 2027. The full trajectory (1,500-2,000 EDSPA → ~3,000 core → ~7,000-8,000 full ~2037 → **>15,000 maturity 2040+**) is grounded in Chapter 15. The limit that does not scale is that of nature: each department's mandates are frozen by Chapter 13's pure protector. And a growing share of that staff is deliberately non-automatable: Chapter 8's human-AI dual verification makes the PP Europe's largest public employer of non-automatable human judgement against AI.

12.6 The National Offices

Element	Design
Presence	1 office per Member State (27), size proportional to population and threat exposure
Function	Territorial implementation, coordination with national authorities, citizen complaints, local monitoring, education
Relation to EDSPA nodes	The EDSPA's national nodes (Ch. 7) are the direct embryo: they become National Offices upon completion of the integration
Funding	Charged to the Protective Power's budget (own resources), with no state co-funding that could compromise independence

12.7 The integration of the EDSPA (2033-2035)




Aspect	EDSPA as agency (2026-2035)	Dept. of Cognitive Protection in the PP (2035+)
Legal basis	Enhanced-cooperation regulation (Art. 20 TEU)	Treaty on EU: constitutional competence
Authority	Regulatory, within the acquis (DSA, AI Act, NIS2)	Constitutional: legislative initiative and own protective execution
Scope	Participating States (the 9 founders and accessions)	All 27, with erga omnes effect
Funding	40/35/15/10 model, annual budget cycle	Own resources (Ch. 16): constitutional financial autonomy
Sanction	Referral to the competent authorities	Referral and litigation: the PP still does not sanction — sanction is external and judicially reviewable (section 12.4)
Judicial control	EU and national courts	Board of Appeal + Specialised Chamber of the General Court + CJEU cassation
Staff	1,500 (base) to 2,000 (accelerated) employees	Fully transferred + staged endowment trajectory (Ch. 15): ~3,000 in the core 2033-2035 → ~7,000-8,000 full ~2037 (27 states, 24 languages, 24/7) → >15,000 at maturity 2040+

Continuity guarantees: automatic transfer of contracts; uninterrupted EDSPA operation until its regulation is repealed; data migration validated by ENISA; retention of all competences until the last day (no legal gap); PP budget guaranteed from the first day the Treaty enters into force.

12.8 Comparison: the agency versus the constitutional power

What changes with integration

From agency (revocable) to 4th power (constitutional) — without touching sanctions



Cognitive protection goes from an annual political decision to a permanent constitutional commitment.

What does not change: the sanction stays outside the protecting power, in the hands of the Commission and the courts.

Figure 12.4 — The improvements of integration: permanence, universal coverage, specialized jurisdiction and financial autonomy. The detect/sanction separation is the one thing that stays identical.

12.9 Conclusion: the architecture of a power for the 21st century

1. **The Protective Power will be a single, protecting institution**, not a sanctioning one: mixed Council with technical vote, single executive direction, four operational departments and national reach. Its weapon is light: evidentiary detection, continuous auditing and publicity.
2. **Sanctioning stays entirely outside the PP**: the existing ceilings (DSA 6% / AI Act 7%) are applied by the Commission and the national authorities; the PP refers case files and litigates as a party. No tribunal of its own, no sanctions committee, no fine money.
3. **The judicial volume is solved inside the judiciary**: administrative Board of Appeal (ACER model), Specialised Chamber of the General Court with dually qualified judges (UPC model) and an escalation clause conditioned on real demand (lesson of the Civil Service Tribunal).
4. **The EDSPA does not disappear: it is integrated** as the Department of Cognitive Protection, inheriting staff (1,500-2,000 employees), tools, standards and budgetary base (EUR 1,000M/year), leaving not a single day without protection between 2033 and 2035. And from there the endowment trajectory of a power begins (Ch. 15): ~3,000 staff in the core 2033-2035 → ~7,000-8,000 full ~2037 → more than 15,000 at maturity 2040+.

The chapter's thesis in one sentence: the Protective Power does not punish manipulation — it makes it visible. And visibility, with independent courts and existing sanctioning authorities behind it, is more effective and safer for liberties than any fine imposed by a power that is judge and party.

12.10 Verified sources and references

- **Art. 19 TEU** (judicial system), **Art. 257 TFEU** (specialised courts) and **Art. 267 TFEU** (preliminary rulings): EUR-Lex, consolidated texts.
- **Specialised court precedent**: EU Civil Service Tribunal (2004-2016), abolished by Reg. (EU, Euratom) 2016/1192 and reintegrated into the General Court.
- **General Court expansion**: doubling from 28 to 56 judges (2016-2019) to clear the backlog; CJEU 2024 statistics (1,706 new cases, 1,785 completed, 2,508 pending).
- **Agency Boards of Appeal**: ACER/CPDP model; CJEU, Aquind v ACER, C-46/21 P (full review of complex technical decisions).
- **Dually qualified judges**: Unified Patent Court (34 legally qualified + 51 technically qualified judges, 2022).
- **Acquis sanction ceilings**: DSA (Reg. 2022/2065, up to 6% of worldwide turnover) and AI Act (Reg. 2024/1689, up to 7% or EUR 35M).
- **EU Digital Decade**: target of 80% of adults with basic digital skills by 2030 (Decision 2022/2481); verified coverage ~55.6% (European Commission, JRC).

- **European AI ecosystem:** Mistral AI, Aleph Alpha, Apertus; InvestAI fund (EUR 200,000M) and EuroHPC AI factories.
- **Economic framework:** Chapter 17 (EUR 2,600M base / EUR 3,400M full) and own resources of Chapter 16.

Chapter 13: Competences of the Protective Power

Chapter summary

This chapter defines the exclusive competences of the Protective Power. The thesis is simple — **this is not about creating new power, but about transferring control of the cognitive sphere from non-democratic actors to public institutions** — and the chapter builds it on the design decision consecrated in Chapter 12: the PP is a **pure protective power**. Its competences are grouped into five verbs — **detect, audit, publish, refer and educate** — and are surrounded by explicit limits as important as the competences themselves: the PP does not sanction (files go to the Commission under the DSA and the AI Act, or to national authorities), does not remove or filter content, does not replace the platforms in distribution, and never decides which information is "relevant". Its own artificial intelligence is not an arbiter of truth: it is a **cognitive forensic scanner** that detects patterns of manipulation and generates auditable evidence under dual human verification. The chapter places these competences in the real regulatory context of 2025-2026 — a considerable European acquis (DSA, AI Act, Data Act, Tech Sovereignty Package, European Democracy Shield) that lacks exactly what the PP provides: a systematic enforcer — and closes with the provision, in three phases and without any budgetary commitment, of the **Citizen Checker**, a non-binding public verification service conceived as a digital education tool.

13.1 The digital power vacuum: who controls the cognitive sphere today


The power vacuum is sometimes described with striking but unverifiable figures (a supposed "85% of Europeans receive news via algorithms", "70% of data outside EU jurisdiction", "90% of AI models American"). The thesis stands — effective control of the cognitive sphere lies with actors who answer to no citizens — but this book documents it only with what real sources support:

- **Exposure to disinformation:** the Eurobarometer of February 2025 found that **36% of Europeans** report being exposed "often" or "very often" to content they consider distorted or false. Not a majority — but more than a third of the population exposed routinely.
- **Erosion of informational trust:** the Reuters Institute *Digital News Report* (2026 edition) documents a continued decline of trust in news, the rise of individual creators as a source, and the growth of news consumption mediated by AI assistants — an ecosystem where distribution is increasingly decided by opaque systems.
- **Technological dependence:** the Stanford AI Index (2026) confirms that the **United States concentrate the production of top-tier AI models** and China leads in scientific publications, while Europe advances with its own players (Mistral) but from a position of dependence. There is no verifiable "90%", but the structural asymmetry is real and measured.
- **Manipulation of mixed origin:** studies by the European Digital Media Observatory (EDMO) and the third FIMI report of the European External Action Service show that information manipulation campaigns in Europe combine **domestic and foreign actors** — there is no verifiable "60% of foreign origin", and the threat is not only geopolitical: it is also endogenous.

The honest reformulation of the vacuum

We do not know with statistical precision what share of digital power private and hostile actors control. We do know three verified things: (1) the information distribution systems used by hundreds of millions of Europeans are private and opaque; (2) more than a third of citizens perceive themselves as routinely exposed to disinformation; and (3) no European public institution today holds the permanent, systematic mandate to look inside those systems. That is the vacuum: not a conspiracy — an **absence of democratic oversight**.

The digital power vacuum: who controls what in the cognitive sphere (2026)



Real, verified data:
36% often exposed to disinformation (Eurobarometer Feb 2025) ·
informational trust in decline (Reuters DNR 2026) · domestic and foreign manipulation (EDMO, EEAS FIMI)

The vacuum is not "most power in private hands": it is the absence of a public institution that looks permanently.

Figure 13.1 — Map of the digital power vacuum: the acquis exists, the sanction exists, systematic oversight does not. The Protective Power replaces no one: it occupies an empty space.

13.2 The framework of competences: five verbs and their limits

The competences are not organised as thematic blocks with sanctioning powers mixed in ("fines up to 10%", "bot neutralisation", "removal of fake accounts" have no place here). Everything is organised under the principle of pure protection: the same areas of action — cognitive protection, digital sovereignty, democratic protection, platform regulation, education — but described as **protection functions** with explicit limits.

Competence	What the Protective Power does	Who exercises it today	Explicit limit (what the PP will never do)
Detect	Scan the cognitive sphere with its own forensic AI and human teams to identify manipulation campaigns, inauthentic coordination and violations of cognitive rights	Platforms (self-interest, opaque), EDMO hubs (limited and project-based), no systematic public actor	Does not label content as "true" or "false" for the public; generates indicators and internal evidence under dual verification
Audit	Audit algorithms, recommendation systems, AI systems and data practices under the DSA (Art. 40: data access), the AI Act and the Data Act	Commission and national DSCs (case by case, limited capacity); the Art. 40 DSA delegated act (Jun 2025) organises researcher access, not systematic public auditing	Does not approve or veto algorithmic changes: it audits and publishes what it sees; the decision remains the platform's under the acquis
Publish	Publish audit reports, records of detected manipulation and algorithmic transparency — the PP's weapon is light	Nobody systematically: DSA audit data is confidential by default	Does not publish personal data or private content; publishes aggregate findings respecting legitimate trade secrets
Refer	Send evidentiary files to whoever holds sanctioning competence: European Commission (DSA 6% / AI Act 7%), national market authorities, and litigate as a party before the specialised jurisdiction (Ch. 12)	The sanction already exists and lies outside the PP; what is missing is the systematic feeder of files	Imposes not one fine, not even delegated; does not appropriate the file: refers it and tracks it publicly
Educate	Cognitive literacy (Department of Cognitive Education, EUR 300M/year, target of 90% or more by 2038 anchored to the Digital Decade), journalist training, applied research	States (uncoordinated), NGOs and EDMO (limited resources)	Does not impose doctrinal curricula: verifiable manipulation-detection skills, without ideological bias, audited (Ch. 8)

13.3 The cognitive forensic scanner: the Protective Power's AI

One example must be rejected from the outset: *"a European algorithm, audited and transparent, will decide which information is relevant"*. A State with an algorithm that decides informational relevance is precisely the computational censorship system this book denounces — merely with a different owner. This book replaces it with a central piece of the design:

What the Protective Power's AI is — and is not

It is: a cognitive forensic scanner. It detects statistical patterns of manipulation — inauthentic coordination, amplification networks, deepfakes, FIMI operations, algorithmic exploitation — and generates **auditable evidence**: logs, data lineages, digital chains of custody. Every finding passes through the **human-AI dual verification** protocol of Chapter 8 before becoming a file. The AI is a forensic instrument, like a forensic police laboratory: it detects indicators, not truths.

It is not: an arbiter of content. The scanner **never ranks, filters, prioritises or blocks** information. Platform algorithms keep deciding distribution — under the European *acquis* and audited by the PP. The power of who sees what does not move to Brussels: it is subjected to light. When the scanner finds something serious, the result is not a deletion: it is a file, publicity and, where appropriate, a referral to the Commission or litigation before the Specialised Chamber.

This distinction is what separates the Protective Power from any ministry of truth. A power that filters is a censor with better branding. A power that detects, documents and publishes is a **builder of public evidence**: it changes the debate not because it controls what is seen, but because no one can manipulate in the dark without leaving a record.

13.4 The Citizen Checker: non-binding verification as digital education

Among the book's "citizen tools" there is a valuable intuition: giving citizens direct access to verification capacity. Developed carelessly, it would be dangerous: an "official verifier" issuing verdicts of truth would be a state oracle. Refined into its consensual form, it becomes a piece coherent with the pure protector. It is named the **Citizen Checker** and defines it with normative precision:

- **Nature:** a public verification service, **non-binding and with no judicial value**. The citizen pastes content (link, text, image, video) and receives a **private report of indicators**: manipulation techniques detected, forensic signals (image synthesis, amplification coordination), verifiable context. The report never issues "true" or "false" labels: the conclusion always belongs to the citizen.
- **Permanent disclaimer:** every report carries the warning that it is not a verdict, has no evidentiary value and does not replace one's own judgement. The Checker is a literacy tool in tool format, not a judge with a web interface.
- **Privacy by design:** access authenticated via eIDAS, but queries **pseudonymised with no reversible link** to identity; minimal retention under EDPS supervision; queries are never communicated to third parties, do not generate files and do not feed referrals. What a citizen checks is irrelevant to the PP: only that they learn to check matters.
- **Anti-weapon symmetry:** the forensic scanner is the same for all content — including that of the Protective Power itself, official communiques and electoral campaigns. A tool that only verifies others is a weapon; one that also verifies power is a mirror.
- **Recognised risks:** the Checker can be captured by media narratives ("the report says...") or generate over-trust replacing personal judgement. The mitigations are the indicator format (never verdicts), the permanent disclaimer and, above all, literacy: the service's goal is to become unnecessary.

Phased implementation: a provision only, not a commitment

Phase 1 — EDSPA (2027-2030): no citizen service exists. Only the internal forensic AI, used for cognitive sphere surveillance and files.

Phase 2 — Protective Power (2033-2035): feasibility study and design audit (privacy impact before the EDPS, technical design before the Board of Appeal). No dates or budgets committed.

Phase 3 — Favourable assessment (~2036+): gradual launch — first a pilot with accredited researchers and external verification of error rates; then progressive opening to the public, only if the risk/benefit balance is favourable and documented.

Chapter 13 **provides for** the Checker: it does not promise it or budget it. If the assessment is negative, the PP works just as complete without it.


13.5 The real regulatory context: why the *acquis* needs an enforcer

The urgency is not dated in hypothetical scenarios ("2029 European elections", "Brexit 2.0"). It is anchored in the real regulatory calendar, verifiable in the Official Journal of the EU:

Instrument	Status (verified)	What it provides	What it lacks
DSA (Reg. 2022/2065)	In force; Art. 40 delegated act on researcher data access adopted in July 2025	Algorithmic transparency, audits, 6% sanction ceilings, data access	Access is organised case by case and per accredited researcher: there is no systematic, continuous public surveillance of the cognitive sphere
AI Act (Reg. 2024/1689)	Full applicability on 2 August 2026	Risk-tiered obligations, 7% ceilings, AI system oversight	Product-by-product focus, not sphere-based: nobody monitors the aggregate effect of AI systems on collective cognition
Data Act (Reg. 2023/2854)	Applicable since 12 September 2025	Interoperability, provider switching, device data access	Regulates data circulation, not who uses it to manipulate the cognitive sphere
European Democracy Shield	Presented by the Commission on 12 November 2025	Response to foreign interference, in line with Recommendation (EU) 2023/2829 on inclusive and resilient electoral processes (2023)	Security and cyber-defence focus, centred on foreign actors: domestic and algorithmic manipulation falls outside its focus
Tech Sovereignty Package	Presented on 3 June 2026	Cloud and AI Development Act and Cloud Sovereignty Framework: European cloud infrastructure, dependency reduction	Builds sovereignty of means, but does not protect the cognitive sphere from how those means are used


The table's reading is the chapter's thesis: **Europe has already written most of the law the Protective Power would enforce.** The DSA orders transparency; the AI Act orders system oversight; the Data Act orders openness; the Shield defends elections. What does not exist is a public institution with the permanent mandate, technical resources and constitutional independence to enforce all of it systematically — to look inside the systems, document, publish and refer, every day, not only when a scandal erupts. That absent enforcer is exactly the set of competences described in 13.2.

13.6 The power transfer matrix: 2026 → 2035+



13.7 The constitutional limits: what it can and will never do

In a design whose central threat is manipulation, a public power with powerful forensic tools needs sharper limits than its competences. The chapter raises them to the rank of institutional definition — as binding as the five functions:



This radar is the institutional answer to the classic objection: *"who watches the watchman?"*. The design's answer is threefold: **(1) constitutional limits** — the prohibitions of Figure 13.3 are inscribed in the power's own definition, amendable only through the Chapter 15 procedure; **(2) symmetry** — the scanner and the Checker apply to the PP itself; **(3) specialised jurisdiction** — every act of the PP is challengeable before the Specialised Chamber in Digital Protection and, on cassation, before the CJEU (Ch. 12). A power that cannot touch content, that audits itself with the same tools, and whose acts are judged by independent judges, has the answer built in.

13.8 Why these competences are urgent

Urgency does not depend on speculative scenarios. It depends on the verified convergence of three trends:

1. **The regulatory window closes in August 2026.** The AI Act reaches full applicability on 2 August 2026; the Data Act has been applicable since September 2025; the Art. 40 DSA delegated act is adopted. If no systematic enforcer exists during the *acquis'* years of force, the experience of European digital policy says the law stays on paper and platforms learn to meet the auditable minimum.
2. **Manipulation is industrialising.** The deepfakes and FIMI operations documented by the EEAS are increasingly cheap and effective; the third FIMI report shows coordinated campaigns against European democratic processes that current detection finds after the electoral damage is done.
3. **Informational trust is structurally eroding.** The Reuters Institute has documented for years the decline of trust and the shift towards non-editorial sources (creators, AI assistants) without verification standards. Every electoral cycle without cognitive protection infrastructure amplifies the spiral: less trust, more manipulability, less participation.

The completed institutional counterweight. The Executive, the Legislature and the Judiciary balance one another — but all three operate in a cognitive sphere whose infrastructure is controlled by private and hostile actors. The three powers are today spectators of the manipulation of the opinion that legitimises them. The Protective Power is not a fourth rival to the three: it is the one that ensures the three can keep operating with a citizenry capable of forming its own judgements. It does not balance the powers: **it protects the ground on which all three operate.**

13.9 Conclusion: reclaiming digital power without conquering it

1. **Digital power exists and is not in democratic hands.** It is not controlled by a conspiracy: it is controlled by an opaque distribution infrastructure, AI models concentrated outside Europe, and actors — domestic and foreign — learning to manipulate it at low cost.
2. **The *acquis* for reclaiming oversight is already written.** DSA, AI Act, Data Act, European Democracy Shield, Tech Sovereignty Package: what is missing is a permanent, technical, independent enforcer. That is exactly the slot of the PP's competences.
3. **The competences are five protective functions with constitutional limits.** Detect, audit, publish, refer, educate — and never sanction, filter, prioritise or label. The cognitive forensic scanner generates evidence under dual human verification; platform algorithms keep deciding distribution, now in plain sight.
4. **The Citizen Checker is a prudent provision.** Non-binding indicator verification, pseudonymised, symmetric, in three phases and with no budget committed until an independent assessment says it is worth it.
5. **The urgency admits no delay.** The AI Act's full applicability (August 2026), the industrialisation of manipulation and the structural erosion of trust converge on the same years. Every year without an enforcer is an unprotected electoral cycle.

The chapter's thesis in one sentence: the Protective Power does not conquer digital power — it illuminates it. And an illuminated cognitive sphere, with existing sanctioners on the other side of the files and independent judges behind every act, is harder to manipulate and harder to censor at the same time.

13.10 Verified sources and references

- **Eurobarometer on the perception of disinformation (February 2025):** 36% of Europeans report being exposed "often" or "very often" to distorted or false content.
- **Reuters Institute Digital News Report (2026 edition):** continued decline of trust and interest in news, rise of individual creators and of news consumption mediated by AI assistants.
- **Stanford AI Index (2026):** concentration of top-tier AI models in the US, Chinese leadership in publications, European advance (Mistral) from a dependent position.
- **EDMO (European Digital Media Observatory):** studies on information manipulation actors, combination of domestic and foreign actors in Europe.

- **EEAS, third FIMI report:** foreign information manipulation and interference campaigns against European democratic processes.
- **DSA (Reg. 2022/2065):** 6% ceilings, audits, and the Art. 40 delegated act on researcher access to platform data (adopted July 2025).
- **AI Act (Reg. 2024/1689):** ceilings of 7% of worldwide turnover or EUR 35M, full applicability on 2 August 2026.
- **Data Act (Reg. 2023/2854):** applicable since 12 September 2025.
- **European Democracy Shield (European Commission, 12 November 2025)** and Recommendation (EU) 2023/2829 on inclusive and resilient electoral processes (2023).
- **Tech Sovereignty and Cloud AI Package (European Commission, 3 June 2026):** Cloud and AI Development Act, Cloud Sovereignty Framework, dependency reduction, European infrastructure.
- **Shoshana Zuboff (2019), *The Age of Surveillance Capitalism*** and work by the **Oxford Internet Institute:** analytical framework of the power of informational infrastructure.
- **Institutional design:** Chapters 8 (human-AI dual verification), 12 (architecture, specialised jurisdiction, departments), 15 (legal basis) and 16 (own resources).

Chapter 14: Transition Process

From the EDSPA to the Protective Power without a single day of unprotected gap: the double constitutional route — enhanced cooperation among nine states with gradual accession, unanimity as the horizon — and the real precedents showing Europe already knows how to do this

Chapter summary


This chapter details the step-by-step process for transforming the EDSPA into the Protective Power between 2027 and 2035. The timeline rests on two deep corrections. First: **the realistic constitutional route**. A Treaty reform "approved by qualified majority" with an entry into force "with 20 states" would be legally impossible under Article 48 TEU in force, which requires unanimity and ratification by all. The chapter adopts the double route: **enhanced cooperation under Article 20 TEU with a founding core of 9 states** (the legal minimum, already used for the Rome III divorce regulation and the unitary patent), with Schengen-style gradual accession and provisional application (the Unified Patent Court precedent, which started operating that way), and **unanimity of the 27 as the horizon**, not as a starting condition. The political precedent is recent and verifiable: on 22 November 2023 the European Parliament triggered Article 48 TEU for the first time in history (291 votes in favour) and the European Council gave it no follow-up — the all-27 route is blocked today, and the design assumes that reality instead of ignoring it. Second correction: **the transition respects the pure protector**. Sanctioning powers ("10% fines", "automatic neutralisation", "constitutional sanctions") are not transferred: what is transferred is protective competences, staff, assets and evidence — sanction remains entirely with the Commission (DSA 6% / AI Act 7%). Financing refers to Chapter 16: own resources, with no fines or taxes as PP income — fines imposed go to the EU's general coffers, not to the budget of the power that motivated them; and the export of the technology to other democracies is noted as a possible future own resource, never as an assumption of the plan. The process's central guarantee: **not one day without protection**, following the technical precedents of the Europol conversion (2017) and the UPC's start (2023).

14.1 The real precedents: Europe has already done this three times

The chapter grounds its credibility not in documents that do not exist ("White Paper on Constitutional Reform 2024", "Report on the Feasibility of a 4th Power 2023", "Reflections on the Future of the EU 2022"), but in a verifiable precedent of European law and institutional history for every piece of the process:

Precedent	What it demonstrates	Lesson for the EDSPA → PP transition
Enhanced cooperation (Art. 20 TEU)	A minimum of 9 member states can advance integration within EU structures without waiting for all: this is how the Rome III divorce regulation (2010) and the unitary patent were born	The founding core of 9 states from Chapter 6 is legally sufficient: no need to wait for 27
Europol (Reg. 2016/794)	On 1 July 2017 an intergovernmental organisation of the states became an EU agency "replacing and repealing" its earlier acts, with staff and activities transferred without a single day of gap	Integrating the EDSPA as the Department of Cognitive Protection can be done through a successor organisation, inheriting contracts, files and legal bases in one act
Unified Patent Court	Agreement signed in 2013; provisional application from 2022 (judges and staff hired before full ratification) and entry into force on 1 June 2023 with 17-18 states; Denmark ratified after a referendum (2014)	Provisional application allows building the PP of 2033 while ratification completes, and gradual scaling by groups of states is viable and proven
Treaty of Lisbon	Signed 13 December 2007, in force 1 December 2009: 24 months of ratification with one national referendum rejected and repeated (Ireland 2008→2009) along the way	Planning 24 months of ratification is realistic; a single national setback does not sink the process if the design already contemplates gradual accession
European Parliament, 22 November 2023	For the first time in history, it triggered the Article 48 TEU procedure asking for a Convention to be convened (291 in favour, 274 against, 44 abstentions); the European Council gave it no follow-up	The constitutional appetite exists but the unanimity route is blocked today: the transition design must assume it, not pretend otherwise
Conference on the Future of Europe (2022)	49 proposals with more than 300 measures, built with randomly selected citizen panels, including Treaty reforms	The participatory method is already invented and tested: the PP's Convention reuses it with citizen panels on cognitive protection

14.2 The master timeline: five phases, 2027-2035



Phase 1: EDSPA operational (2027-2030)

The nine founding states launch the agency with the base budget of EUR 1,100M/year. The goal is not only to protect from day one, but to **generate the evidence that makes reform inevitable**: audits published, files referred, elections protected (including the 2029 European elections), and the 164% ROI and EUR 1,820M/year net benefit documented with the Chapter 9B methodology. The EDSPA does not ask for faith: it produces auditable results.

Phase 2: Assessment and proposal (2030-2031)

A triennial EDSPA impact report with external verification; a Europe-wide public consultation with randomly selected citizen panels (the method already tested by the Conference on the Future of Europe); and a formal European Commission proposal for an enhanced cooperation institutionalising the Protective Power. The argument is Chapter 11's: the EDSPA works, but it can be revoked by a political majority — cognitive protection needs constitutional anchoring.

Phase 3: Legal route (2031-2033)

Council authorisation of the enhanced cooperation (qualified majority, after meeting the Art. 20 TEU requirements), negotiation of the constitutive convention among participating states, signature and **ratification with a 24-month margin** — the real time Lisbon took, including one national setback. For states whose constitutions require referendums, the Danish UPC precedent (2014) shows these processes can also be won at the ballot box.


Phase 4: Entry into force of the core (~2033)

The Protective Power is born with the states that have already ratified — no fewer than 9 —, with **provisional application** from signature so that the organs (Council of Cognitive Protection, executive directorate, start of the Specialised Chamber) are built and staffed on day one, exactly as the UPC hired its judges before opening its doors.

Phase 5: Gap-free integration (2033-2035) and gradual accession (2033 → horizon)

By a successor-organisation act modelled on Regulation 2016/794 (Europol), the EDSPA becomes the Department of Cognitive Protection of the PP: staff, seat, systems, files and legal bases transfer as a block, and the agency keeps **all** its competences until the last day. In parallel, the other three departments and the Chapter 12 specialised jurisdiction are built. The other states accede at their own pace, each with its own ratification — the Schengen model: it started with fewer states than the Community and today covers almost all of them, including non-EU countries.

14.3 The double constitutional route: why nine and not twenty-seven




The underlying reason is not tactical but democratic: demanding unanimity to *start* means the most reluctant — or most manipulated — state decides whether the others may protect themselves. Enhanced cooperation inverts that logic: those who want protection get it soon, and those who wait block no one. And experience teaches that what works ends up attracting: Schengen, the euro and the UPC all grew after being born small.

14.4 What is transferred — and what is left untouched

Element	EDSPA (2027-2033)	Protective Power (2033-2035 onwards)	Mechanism
Protective competences	Detect, audit, publish, refer, educate (Ch. 13)	Identical, with constitutional anchoring and coverage of acceding states	Successor-organisation act (Europol model): legal bases inherited, not renegotiated
Sanction	Outside the agency: Commission (DSA 6% / AI Act 7%)	Stays outside: the transition adds not one sanctioning power	Not transferred because it exists at neither level
Staff	1,500-2,000 employees	Transferred as a block, contracts included; gradual expansion per Ch. 12	Automatic transfer of contracts (Europol 2017 precedent)
Assets	Seat, forensic systems, cognitive scanner, files	Inherited in full on day one	Closed, audited inventory before the convention is signed
Budget	EUR 1,100M/year (base)	Ch. 17 framework: EUR 2,600M base / EUR 3,400M full, with Ch. 16 own resources	Tranche-by-tranche migration as departments and states join
Fine income	None for the agency	None for the PP: fines imposed by the Commission go to the EU's general coffers	Structural separation forever: the eye that detects is never the hand that collects

14.5 The continuity guarantee: not one day of gap



14.6 Obstacles and strategies

Obstacle	Risk	Strategy	Verifiable anchor
States that do not join	Initial partial coverage	No blocking: enhanced cooperation from 9, gradual accession, each accession advertises the success	Art. 20 TEU (minimum 9); Schengen and the euro grew after being born small
A ratification setback	Months or years of delay	24-month margin in the timeline; provisional application keeps the construction calendar	Lisbon took 24 months with one repeated referendum; the UPC started under provisional application
Platform lobbying	Campaigns against the reform, capture of the debate	Full transparency of negotiations; the EDSPA itself documents and publishes the influence operations it detects	EU transparency register; DSA Art. 40 (researcher data access)
Changes of government	Withdrawal of political support	Transversal base: the benefits (ROI 164%, EUR 1,820M/year) and the files belong to all governments; convention signed after audited impact reports	Schengen/euro accession survived dozens of government changes
Manipulation against the reform itself	Disinformation about "the European ministry of truth"	The design itself is the answer: the Ch. 13 constitutional limits (never filters, never sanctions, anti-weapon symmetry) and public litigation of every myth with evidence	Ch. 13 limits radar; verifiable scanner symmetry
Migration fragility	Data or capacity loss	ENISA-validated migration protocols, redundancy, closed inventory before signature	Certifications and practices of EU decentralised agencies

14.7 Success indicators of the transition

The chapter promises no invented quantitative results ("40% disinformation by 2030", "70% citizen support", "12-20-27 countries"). It offers **verifiable process indicators** — things that can actually be counted and audited:

- **Protection:** algorithmic audits completed and published per year; files referred to the Commission and resulting sanctioning decisions; the 2029 and 2034 European elections covered without serious documented-and-unanswered interference incidents.
- **Economic credibility:** ROI of 164% and net benefit of EUR 1,820M/year externally verified (Ch. 9B methodology), budget executed with deviation below 5%.
- **Expansion:** states acceding to the enhanced cooperation (9 → 15 → 20+); population coverage of cognitive protection; accessions after each well-protected electoral cycle.
- **Social capacity:** progress of cognitive literacy towards the 90%+ target by 2038 (from the Digital Decade's real baseline: 80% target for 2030, ~55.6% verified today).

14.8 Conclusion: a transition designed not to fail

1. **The process does not ask for a political miracle: it asks for nine determined governments.** Art. 20 TEU allows starting with the founding core; Schengen, the euro and the UPC show that successful European integrations are born small and grow through results.
2. **The timeline is anchored in real durations.** 24 months of ratification (Lisbon), provisional application (UPC), gap-free successor organisation (Europol): every phase copies a precedent that already happened, not a hope.
3. **The transition does not dilute the pure protector: it consolidates it.** No sanctioning powers are transferred because they exist at neither level; fines stay with the Commission and the EU's coffers; the PP reaches 2035 with more coverage, more departments and the same weapon as always: light.
4. **The all-27 route is not abandoned: it is rescheduled.** Unanimity is the consolidation horizon at a future Convention — when existing and working is the best argument for a Treaty that cannot be approved today.

The chapter's thesis in one sentence: the EDSPA demonstrates that protection works; enhanced cooperation lets nine states make it constitutional without asking permission from the slowest; and gradual accession turns every year of results into the argument that pulls the others along — until the day the 27 consolidate by unanimity what the majority already lives.

14.9 Verified sources and references

- **Art. 48 TEU** (ordinary and simplified revision procedures) and **Art. 20 TEU** (enhanced cooperation: minimum 9 states): EUR-Lex, consolidated texts.
- **European Parliament resolution of 22 November 2023** triggering Art. 48 TEU for the first time and asking for a Convention to be convened (vote: 291 in favour, 274 against, 44 abstentions; the 9-Jun-2022 resolution asking for the same was adopted 355:154:48).
- **Conference on the Future of Europe (2022)**: final report, 49 proposals with more than 300 measures, randomly selected citizen panels; proposal 39 on ending unanimity and 38 on European referendums.
- **Europol**: Regulation (EU) 2016/794, which turned the intergovernmental organisation into an EU agency on 1 July 2017, "replacing and repealing" its founding acts, with staff and activities transferred.
- **Unified Patent Court**: 2013 Agreement, provisional application from 2022, entry into force 1 June 2023 with 17-18 states; Danish accession referendum (25 May 2014, 62.5% in favour).
- **Treaty of Lisbon**: signed 13-Dec-2007, in force 1-Dec-2009; ratification completed with the Czech deposit of 13-Nov-2009; Irish referendum of 2008 (rejection) and 2009 (approval).
- **Enhanced-cooperation precedents**: Rome III divorce regulation (Reg. 1259/2010, 14→17 participating states) and the unitary patent.
- **Schengen and the euro** as precedents of gradual accession and variable-geometry integration.
- **The book's own frameworks**: Ch. 9B (ROI 164%, EUR 1,820M/year), Ch. 6 (founding core of 9), Ch. 11 (EDSPA as necessary step), Ch. 12 (architecture and 2033-2035 integration), Ch. 13 (pure protector competences), Ch. 16 (own resources) and Ch. 17 (EUR 2,600M/3,400M).

Chapter 15: Legal Basis of the Protective Power

The law already exists: values, fundamental rights and case law sustain cognitive protection today — what is missing is not a legal basis but constitutional rank: and the power that protects democracy in the digital era is endowed like a power, not like a programme


Chapter summary

This chapter is titled "Legal Basis" and it is recentred exactly on its title: **demonstrating that European law already sustains cognitive protection — and that the only thing missing is rank**. The thesis is constructive at three levels. First: the **principle of protection** is already in the Treaty — Art. 2 TEU enshrines democracy among the Union's values, Art. 3 defends it among its objectives, and the Charter of Fundamental Rights protects private life (Art. 7), personal data (Art. 8), freedom of expression and information (Art. 11) and effective judicial protection (Art. 47) — the four exact rights that digital manipulation violates. Second: CJEU case law has already recognised that digital dignity is justiciable — from *Digital Rights Ireland* to *Schrems II*. Third: existing positive law (DSA, AI Act, GDPR) already orders almost everything the PP would enforce — but it orders it without a permanent enforcer, and no ordinary law can create a 4th power: that requires Art. 48 TEU (with Chapter 14's double route) or, failing that, Art. 20 TEU. The chapter adds the decisive piece: **the endowment of a power**. European democracies remain anchored in the analogue era — their three powers were designed when public opinion was formed in the square, the press and television, and none of them watches the digital terrain where wills are formed today. Bringing democracy into the digital era requires a power sized as such: from the 1,500-2,000 inherited from the EDSPA to **more than 15,000 staff at maturity (2040+)** — still the lightest of the four powers for the newest mission, financed by the virtuous revenue circle of its own files (fines to the EU's coffers, own resources under Art. 311 TFEU) and by the counterweight no other power offers: Europe's largest public employer of non-automatable human judgement. And with a doctrine the chapter fixes: a power is not justified by its economic viability — it is endowed according to its mission; that it also pays for itself (ROI 164%, Ch. 9B) is the cherry, not the cake.

15.1 The hierarchy of legal bases: three steps towards the 4th power

Nothing in the PP's design requires inventing new law: it requires **stepping up** a hierarchy that already exists. Each step has its basis, its rank and its reversibility — and the book's argument is that only the third guarantees protection that does not depend on an annual political majority.

Hierarchy of legal bases: from repealable law to constitutional power




Cognitive protection does not change content when it steps up: it changes rank.

And rank is what turns "opportunity" into a guarantee: what one majority cannot give, another majority cannot take away.

Figure 15.1 — The three steps of the legal basis: the EDSPA can be founded tomorrow under existing law; the 4th power requires Art. 48 TEU — with Chapter 14's double route.

15.2 The principle of protection is already in the Treaty

The most common objection — "there is no legal basis for such a power" — confuses two questions: whether the EU *can* protect the cognitive sphere, and whether the EU *must* do so at constitutional rank. The first is already answered by existing law:



The case law that already protects the cognitive sphere

- **CJEU, *Digital Rights Ireland* (C-293/12 and C-294/12, 2014):** invalidated the Data Retention Directive because mass, indiscriminate retention interferes in a particularly serious manner with fundamental rights — the Court already said that total surveillance of digital life is incompatible with the Charter.
- **CJEU, *Google Spain* (C-131/12, 2014):** the right to erasure against the search engine recognises that the individual controls their informational footprint: informational autonomy is an operative right, not a theoretical one.
- **CJEU, *Schrems II* (C-311/18, 2020):** the Privacy Shield fell because transferring data to third countries without equivalent protection is unlawful — the sovereignty of personal data is enforceable.
- The common line: **the person's cognitive sphere is legally protectable** — the CJEU has defended it case by case. The PP proposes what that case law's logic already points to: systematic protection, not only reactive.

The financial basis: Art. 311 TFEU

Chapter 16's financial autonomy is grounded in **Art. 311 TFEU**: the Union provides itself with own resources through a Decision adopted unanimously and approved by the states under their constitutional rules — the same mechanism as Decision 2020/2053 that financed the post-pandemic recovery. A stable constitutional power needs exactly that: income it does not renegotiate with each annual budget. The PP is not financed with fines and taxes; financing is the matter of Chapter 16: **own resources yes, fines as own income never** — the revenue circle works indirectly (see 15.5).

15.3 Montesquieu and the 4th power: the objection and the answer

The classic objection


"Montesquieu's separation of powers (legislative, executive, judicial) is the guarantee of liberty. A 4th power breaks the balance."

The answer has two parts. The historical one: Montesquieu did not enumerate a closed list in *The Spirit of the Laws* (1748) — he described the principle that **power must check power**, and constitutions have multiplied the organs that embody it: constitutional courts, independent central banks, ombudsmen, fiscal authorities. None existed in 1748; all are democratic orthodoxy today. The functional one: the three classic powers now operate in a cognitive sphere they do not control — and a power that does not control the terrain of the public opinion that legitimises it balances nothing: **it is blind**. The PP does not add a competitor to the three: it ensures the three can see. And unlike a sanctioning 4th power — which would be a real danger — this one is designed without the tool that would make any power fearsome: **it cannot punish, only illuminate** (Ch. 13). The separation of powers is not broken by a power that publishes; it would be broken by a power that silences, filters and punishes at once.

15.4 The endowment of a power: bringing democracy into the digital era

European democracies remain anchored in the analogue era. Their three powers were designed — and endowed — when public opinion was formed in the square, the press and television: the legislature to deliberate in chamber, the executive to administer territories, the judiciary to try cases. Today wills are formed in opaque recommendation systems, operated mostly from another jurisdiction, and none of the three powers watches that terrain. **That is the constitutional anachronism of the 21st century**, and correcting it is not creating "one more agency": it is endowing democracy with the power it lacks in order to exist in the digital era. And a power is endowed the way powers are endowed.

Offices are not handed out by decree with pre-assigned cities, nor with a magic figure of "10,000+ employees". The chapter keeps the three categories of presence — proximity, coordination, excellence — under the criterion of merit and open call, and replaces any magic figure with a **staged trajectory** in which every jump has an accountable trigger: accessions of states, coverage of the mission, volume of threat and the expansions of mandate the legislator authorises over time — as has happened to every European power.



Why the figure is that of a power, not of an agency

Three masses of staff make the trajectory inevitable, and none of them is bureaucracy:

- **The real scale of the operation.** Continuous forensic surveillance of the cognitive sphere of ~450 million people across 27 states and 24 official languages is an operation of shifts and continuous coverage: forensic scanner operators, human verifiers (Chapter 8's dual verification is deliberately non-automatable — it is the heart of the design), platform auditors by region and language, national offices that at maturity operate with 100-150 people. This block alone already exceeds the ceiling of the old "agency" projection.
- **The cognitive education corps.** The target of literacy of 90% or more by 2038 is a programme the size of an education system: trainers, instructional designers, coaches, coordination with the school systems of 27

countries. Education does not scale with software alone — it scales with teachers.

- **The organic growth of every power.** All European powers have grown in mandates and staff over time — the Commission from a few thousand to ~32,000; Frontex from dozens of analysts to a statutory corps of 10,000 set for 2027. The PP will not be the exception: between 2033 and 2040 the legislator will authorise expansions — Chapter 13's Citizen Checker is only the first one foreseen, and its staffing needs are already inside the maturity horizon, without a separate figure. What does not grow is the *nature* of its powers: Chapter 13's constitutional limits freeze what it can do; the stages only scale how much of the same.

The AI-era counterweight: the largest public employer of non-automatable human judgement

AI eliminates *routine* human judgement from the economy — standard writing, basic analysis, support, simple design. The PP creates in Europe the largest set of public jobs in *non-routine* human judgement: verifying forensic evidence, supervising audits, educating cognitively, exercising criteria in the face of ambiguity. And the design guarantees it: Chapter 8's human-AI dual verification is not a transitional phase towards full automation — it is an institutional principle. Every improvement of the scanner's AI does not reduce staff: **it increases the volume of findings that demand human verification and judgement.** While the rest of the economy discovers how many jobs can do without people's criteria, the PP is the European institution that hires the most human criteria — and will do so more every year. Faced with the tsunami of routine cognitive job destruction, Europe would have a pole of specialised reabsorption: operators, controllers, designers, trainers, educators — prevention, education, supervision.

15.5 The financing of a power: the virtuous revenue circle

And who pays for a power of more than 15,000 staff? The question has two answers — one accounting, one doctrinal — and both play in favour of the design.

The accounting one: the revenue circle. Every forensic file the PP refers and the Commission turns into a sanction (DSA up to 6%, AI Act up to 7%) goes into **the EU's general coffers** — and from those coffers, via the own resources of Art. 311 TFEU, the PP's budget is endowed. The circle is virtuous and structurally separated, as the pure protector demands: **the power that detects never collects what it detects** — but the Union that endows it does collect what its work produces. And the magnitude is no longer theory: the Commission has imposed fines of billions of euros under the digital acquis (Google Shopping, Android, Apple...), and each one sharpened the discipline of the European market. A systematic PP would multiply that flow: not as own income of the power (never), but as a flow that makes the whole Union gain from every file. The PP's activity, which costs, is the same activity that raises revenue.

The doctrinal one: a power does not need to be viable — it needs to fulfil its mission. Nobody asks whether the judiciary "is profitable", or whether parliament "pays for itself": they are endowed with what their constitutional function requires, because without them there is no democracy that can afford the luxury of calculating. The PP belongs to that category: it is the condition for democracies to exist in the digital era, and that is why it is endowed according to its mission, not according to its profitability. That in addition — as Chapter 9B demonstrated with verifiable methodology — it yields a 164% ROI and a net benefit of EUR 1,820M/year makes it an almost unique case in constitutional history: **a power that justifies itself without economics and that, nevertheless, economics endorses.** Viability is the cherry on the cake; the cake is protected democracy.

Dimension	Figure	Foundation
Budget	EUR 2,600M/year base / EUR 3,400M full (Ch. 17)	In real context: the 2026 EU budget is EUR 192,800M in commitments; the European Council + Council section is EUR 739.6M — the full PP is an institution with the budget of a major EU organ, consistent with its rank
Financing	Own resources (Ch. 16, Art. 311 TFEU)	Proven mechanism (Decision 2020/2053); no fines or taxes as PP own income — the revenue circle works indirectly: the Commission's fines on PP files feed the EU coffers that endow it
Direct employment	1,500-2,000 (EDSPA) → ~3,000 (core) → ~7,000-8,000 (full ~2037) → >15,000 (maturity 2040+)	Staged trajectory with accountable triggers (Fig. 15.3); the counterweight of non-automatable human judgement against AI
Return	ROI 164% over 5 years; EUR 1,820M/year net benefit	Ch. 9B — an additional argument, not a justification: the power is endowed according to its mission
Talent	European pole of specialisation in cognitive protection	Stability of constitutional rank + excellence centres by open call (ERC model): the European career of civic technology

15.6 Conclusion: the legal basis exists; what is missing is rank — and endowment of power

- 1. The EU already holds the mandate to protect the cognitive sphere.** Values (Art. 2 TEU), objectives (Art. 3 TEU), rights (Charter, Arts. 7, 8, 11, 47), competences (Arts. 16 and 114 TFEU, Art. 352 if needed) and case law (Digital Rights Ireland, Google Spain, Schrems II): cognitive protection is the pending enforcement of existing law.
- 2. What is missing is not basis: it is rank and an enforcer.** For that there are Art. 48 TEU (the full route) and Art. 20 TEU (the founding route) — Chapter 14's double route.
- 3. Montesquieu's objection is answered by the design.** The principle is that power check power; the PP adds the missing counterweight and arrives without the tool of danger: it does not sanction, it does not filter, it only illuminates, with independent jurisdiction behind every act.
- 4. And it is endowed as what it is: a power.** Democracies anchored in the analogue era need the power that brings them into the digital era — with an endowment trajectory that goes from the 1,500-2,000 inherited from the EDSPA to more than 15,000 staff at maturity 2040+, financed by the revenue circle of its own files and by own resources, and turned into Europe's largest employment pole of non-automatable human judgement against AI. A power does not need to be viable: it needs to do its job well — and this one, besides fulfilling its mission, pays for itself.

The chapter's thesis in one sentence: European law has spent a decade ordering and defending cognitive protection; what it needs is to give it Treaty rank and endowment of power — because a 21st-century democracy is not protected with a 20th-century agency, but with the fourth power that brings it into the era its citizens already live in.

15.7 Verified sources and references

- **Treaties (EUR-Lex, consolidated texts):** Art. 2 TEU (values), Art. 3 TEU (objectives), Arts. 20 and 48 TEU (enhanced cooperation and revision), Arts. 13-19 TEU (institutional framework), Art. 16 TFEU (data protection), Art. 114 TFEU (internal market), Art. 352 TFEU (flexibility clause: unanimity + EP consent), Art. 311 TFEU (own resources), Arts. 326-334 TFEU (enhanced cooperation).
- **Charter of Fundamental Rights:** Arts. 1 (dignity), 7 (private life), 8 (data protection), 11 (freedom of expression and information) and 47 (effective judicial protection).
- **CJEU case law:** Digital Rights Ireland, joined cases C-293/12 and C-294/12 (8-Apr-2014); Google Spain, C-131/12 (13-May-2014); Schrems II, C-311/18 (16-Jul-2020).
- **Verified budgets:** EU budget 2026: EUR 192,800M in commitments / EUR 190,100M in payments (signed 26-Nov-2025); European Council + Council Section II: EUR 739.6M in 2026.
- **Institutional magnitudes (orders of magnitude, not budget figures):** European Commission ~32,000 officials and agents; EP staff (institutional + political groups) above 10,000; the national judiciaries of large member states employ tens of thousands of people.

- **Organic growth precedent:** Frontex, from dozens of analysts in 2004 to a statutory corps of 10,000 set for 2027 by the legislator; successive staff expansions also at Europol and in the doubling of judges of the General Court.
- **Own resources:** Decision (EU, Euratom) 2020/2053 on the system of own resources (Art. 311 TFEU).
- **Enforceable positive law:** GDPR (2016/679), DSA (2022/2065, Art. 40 and the Jul-2025 delegated act), AI Act (2024/1689, full applicability 2-Aug-2026), Data Act (2023/2854).
- **Doctrine:** Montesquieu, *The Spirit of the Laws* (1748) — the principle of power checking power; its later institutional development.
- **The book's own frameworks:** Ch. 8 (human-AI dual verification), Ch. 9B (ROI 164%, EUR 1,820M/year), Ch. 12 (architecture, 1,500-2,000 EDSPA employees), Ch. 13 (pure protector, Citizen Checker), Ch. 14 (double route), Ch. 16 (own resources), Ch. 17 (EUR 2,600M/3,400M; 55,000-68,000 ecosystem jobs).

Chapter 16: The Financing of a Power — Own Resources

A power does not ask for a budget every January: Art. 311 TFEU provides the Union with income that no annual majority can haggle over — and the Protective Power is financed through the mechanism that already exists, not with fines or taxes of its own

Chapter summary

Financing a power is not presenting a business plan: **the power that detects never collects what it detects**. This chapter focuses on what its title promises: **financial autonomy as the condition of institutional independence**. The thesis is constructive at three levels. First: the mechanism already exists — **Art. 311 TFEU** allows the Union to provide itself with own resources through a Decision adopted unanimously and ratified by the 27 states under their constitutional rules: exactly what Decision 2020/2053 (2020) did to finance the post-pandemic recovery and create the plastic own resource. Second: the timing is the best possible — the Commission proposed on 16-Jul-2025, within the 2028-2034 MFF, the largest redesign of EU revenue in a generation (CBAM, ETS, e-waste, e-commerce fee, ETIAS; an estimated ~EUR 410,000M additional 2028-2035): **the Union is already redesigning how it pays for itself, and the PP's own resource arrives at a table that is already set**. Third: the scale is that of an institution, not of a programme — the full PP (EUR 3,400M/year) equals **~1.8% of the EU's 2026 revenue (EUR 190,091M)**, less than 0.02% of the Union's GNI, and Chapter 15's revenue circle is not theory: the fines of the digital acquis already imposed by the Commission (more than EUR 11,000M in the Google case alone, plus the first DMA fines: Apple EUR 500M, Meta EUR 200M) enter today the EU's coffers as "other revenue" of the budget. The chapter closes with the horizon: the export of cognitive-protection technology to democracies as **a possible future own resource — not an assumption of the plan** — and with Chapter 15's doctrine: a power does not need to be viable; it needs to fulfil its mission. That it also yields a 164% ROI and EUR 1,820M/year (Ch. 9B) is the cherry, not the cake.

16.1 The constitutional principle: a power does not ask for a budget every January

Every EU institution is funded from the general budget — but a **constitutional power** needs something more: that its endowment does not depend on the mood of the annual negotiation. The difference between an agency and a power is not in what they do: it is in what *nobody can take away from them*. An agency's budget is renegotiated every year in the Art. 314 TFEU procedure; the Treaty rank of a power is endowed in primary law.


The Treaty has already solved this problem for the Union itself. **Art. 311 TFEU** provides that the Union supplies itself with **own resources** to finance its budget, through a Council Decision adopted by **unanimity**, after consulting the European Parliament, and approved by the member states "**in accordance with their respective constitutional requirements**" — that is, with the same level of ratification as a Treaty reform. This is not a procedure invented for the PP: it is the one that has endowed every major European institutional leap, most recently the **Decision (EU, Euratom) 2020/2053** (14-Dec-2020), ratified by the 27 and in force since 1-Jun-2021, which financed the post-pandemic recovery and created the first new own resource in decades: the contribution based on **non-recycled plastic packaging waste**.

Why the plastic resource matters here: when the Union wanted to finance something new and priority (the post-pandemic reconstruction), it did not resort to fines or ad hoc taxes: it inscribed a category of endowment in the own resources Decision and the 27 ratified it. **That is exactly the PP's design:** its endowment enters the own resources Decision as its own category (an endowment call of ~0.02% on the Union's GNI in the full scenario), with the guarantee that no annual majority can haggle over what the Treaty orders to be endowed.

16.2 The current system and its reform: the 2025-2026 window

The current own resources system rests on four legs — and one of them carries almost all the weight. The 2026 budgetary revenue (EUR 190,091M) is financed as follows: **GNI-based contributions, 70.28%** (uniform call rate

0.6930%); **VAT-based resources, 13.02%** (rate 0.30% on bases capped at 50% of GNI); **traditional own resources** (customs duties; states retain 25% as collection costs), **11.24%**; and the **plastic resource, 3.59%**. In sum: almost three quarters of the budget depend on cheques the states sign every year according to their GNI — the "juste retour" that budget doctrine has flagged for decades as the system's structural weakness.



That is why the timing window matters. The own resources reform is already on the table as a full package (own resources Decision and MFF 2028-2034, with the second sectoral package on 3-Sep-2025), and its procedure — unanimity and ratification — is **the same one the PP's Treaty will require (Art. 48 TEU)**. The two reforms can travel together: when the Convention drafts the 4th power's Treaty, the own resources Decision for its endowment can proceed in parallel, with the same founding coalition ratifying both. **A single political act endows the power and finances it.**

16.3 The pure protector rule: what the PP never collects

The PP's financing is the exact counterpart of its pure protector design (Ch. 13): **five verbs — detect, audit, publish, refer, educate — and none of them is collect**. The fines its files trigger are always imposed by the Commission (DSA up to 6% of worldwide turnover; AI Act up to 7% or EUR 35M) or by national authorities; and the money goes into **the EU's general coffers**, where it already goes today: competition and digital-acquis fines are "other revenue" of the general budget.

The revenue circle: the PP detects, the Union collects — separated by design



The PP's activity, which costs, is the same activity that raises revenue — but the circle passes through the whole Union's coffers.

Structural separation: the detector never touches the till; the till endows the detector according to the Treaty, not according to its revenue performance of the previous year.

Already verified with real money: Google Shopping EUR 2,420M · Android EUR 4,340M · AdSense EUR 1,490M · adtech EUR 2,950M · DMA: Apple EUR 500M + Meta EUR 200M
More than EUR 11,000M imposed under the digital acquis — partly under appeal — equal more than 3 years of the full PP.

Figure 16.2 — Chapter 15's virtuous revenue circle, with the proof that it already works: digital-acquis fines already enter the EU budget as "other revenue".

The magnitude stops being speculation the moment the files that already exist are counted. The Commission has imposed on a single company, in a single decade, more than **EUR 11,000M** in the digital field (Google Shopping EUR 2,420M in 2017, upheld by the General Court in 2021; Android EUR 4,340M in 2018; AdSense EUR 1,490M in 2019; adtech EUR 2,950M in September 2025), and opened the digital markets acquis with the first DMA fines (23-Apr-2025): **Apple EUR 500M and Meta EUR 200M**. All that money entered the EU's general budget. The difference between today and the PP is not the mechanism — it is *systematicity*: what today depends on episodic, laborious investigations would become the constant flow of a power that audits permanently. The whole Union, not the PP, is the beneficiary of that flow.

16.4 The Protective Power's own resource: scale and design

And how much does it cost to endow a power this way? The exact scale, against the verified 2026 revenue (EUR 190,091M):

Dimension	Figure	Foundation and verified context
Annual endowment	EUR 2,600M base / EUR 3,400M full (Ch. 17)	The Chapter 17 framework, unchanged; the endowment grows with Chapter 15's staffing trajectory (1,500-2,000 → >15,000 at maturity 2040+)
Weight in the EU budget	~1.8% of 2026 revenue (EUR 3,400M/EUR 190,091M); ~1.4% in the baseline scenario	Scale anchor: the full PP is an institution of ~2 euros out of every 100 of the Union's budget — the budget of a major organ, not of a programme
Weight in the EU economy	Less than 0.02% of the Union's GNI (order of magnitude)	An endowment call under Art. 311 TFEU comparable in design to the plastic resource (EUR 0.80/kg): a line in the Decision, not a new tax on third parties
Procedure	Own resources Decision: unanimity + EP consultation + ratification by the 27	Art. 311 TFEU; the same procedure as Decision 2020/2053 — and the same possible calendar as the PP's Treaty (Art. 48 TEU): a single founding coalition ratifies both
Return (additional argument)	ROI 164% over 5 years; EUR 1,820M/year net benefit	Ch. 9B, verifiable methodology — but the endowment is justified by the mission, not by profitability (Ch. 15): a power does not need to be viable, it needs to fulfil its mission

The lesson of the failures: why the PP's resource is different

The history of European finances is full of proposed resources that were never born. The **financial transaction tax** has been proceeding under enhanced cooperation since 2013 — and remains stalled more than a decade later. The **European digital levy** was shelved in 2021 pending the OECD's global tax deal (Pillar 1), from which the US withdrew in January 2025; DSTs remain national taxes and the European route is today a debate. The lesson is twofold. First: **the resources that fail are taxes on third parties** — collected from someone who opposes them with all their might. Second: **the only new own resource that was born (plastic, 2020/2053) is not a tax on third parties: it is an endowment allocation** — and that is exactly the PP's design. Its resource taxes nobody: it allocates, by unanimity and ratification, the share of own resources the Treaty orders to devote to the protection of the cognitive sphere. No company can veto it; only the 27 — and the 27 are those who ratify the Treaty.

16.5 The horizon: exports as a possible future resource, not an assumption

The chapter does not close with a table of exports of EUR 1,150M/year, prices per licence and named corporate alliances — a business plan presented as financing. Instead, as established in Chapter 14: the export of cognitive-protection technology to democracies is **a possible future own resource of the Union — not an assumption of the plan**. The reasoning is directional, not numerical: the PP creates competences (forensics, algorithmic auditing, cognitive education) that other democracies will want to acquire as cross-border manipulation reaches them; if that demand materialises, its fruits must enter the same circle as the fines — **into the EU's coffers, never into the PP's till** — and could, one day, be inscribed in the own resources Decision as a new category. The chapter projects no revenue from it: it registers the door and who holds the key (the 27, unanimously). The geopolitical value of exporting a *model* — democratic protection versus authoritarian control — needs no invented figures to be real, and the ecosystem that would make it possible (55,000-68,000 jobs, Ch. 17) is already quantified in its own chapter.

16.6 Conclusion: financing a power the way powers are financed

1. **The mechanism exists and is in use.** Art. 311 TFEU, Decision 2020/2053, the plastic resource: the Union already knows how to endow itself by unanimity and ratification by the 27 — the PP uses the procedure that already works, not an invented one.
2. **The moment is a window.** The reform proposal of 16-Jul-2025 (MFF 2028-2034 and new own resources: CBAM, ETS, e-commerce, ETIAS; ~EUR 410,000M estimated 2028-2035) sets the table at which the PP's resource fits — and the European Parliament pushes in the same direction. The PP's Treaty and its endowment can be ratified together, by the same founding coalition.
3. **The pure protector rule is financial too.** The PP collects no fines and no taxes: it detects, and the Union collects — and fines already enter the EU's coffers today (more than EUR 11,000M in the Google case; the first DMA fines in 2025). The revenue circle benefits the whole Union, never the power that detects.
4. **The scale is that of an institution.** ~1.8% of the EU's 2026 revenue in the full scenario, less than 0.02% of GNI: an endowment call in the own resources Decision, comparable in design to the plastic resource — endowed according to the mission (Ch. 15), and additionally backed by Chapter 9B's 164% ROI.
5. **And the horizon is registered, not presumed.** Exports to democracies: a possible future own resource through the same route — the EU's coffers and the unanimity of the 27 — without magic figures nobody can verify.

The chapter's thesis in one sentence: a power that depends on the annual budget is not a power — it is an agency waiting for its cut; the Protective Power is financed the way the Union itself is financed: with own resources under Art. 311 TFEU, endowed by unanimity and ratified by the 27, so that the protection of the cognitive sphere never depends on the mood of some January majority.

16.7 Verified sources and references

- **Primary law:** Art. 310 TFEU (unity, universality, balance and specification of the budget), Art. 311 TFEU (own resources: unanimity + EP consultation + approval under constitutional rules), Art. 314 TFEU (annual budget

procedure), Art. 325 TFEU (fight against fraud), Art. 48 TEU (revision of the Treaties).

- **Own resources Decision:** Decision (EU, Euratom) 2020/2053 of 14-Dec-2020 (plastic resource, GNI corrections for several states, ratified by the 27; in force 1-Jun-2021, applied from 1-Jan-2021); Regulation (EU) 2021/770 implementing the plastic resource.
- **Verified 2026 revenue (EUR-Lex, General Budget of the EU 2026, revenue):** total EUR 190,091M; breakdown: GNI 70.28% (call rate 0.6930%), VAT 13.02% (rate 0.30%), traditional own resources 11.24%, plastic 3.59%; VAT bases capped at 50% of GNI; 25% state retention of customs duties (2021-2027).
- **Reform under way (proposal, not law in force):** the Commission's MFF 2028-2034 package of 16-Jul-2025 (COM(2025) 570/571): EUR 1,763bn in commitments in 2025 prices (~1.26% of GNI, up from 1.14% under the 2021-2027 framework); proposed new own resources: CBAM (definitive regime from 2026), ETS, non-collected waste electrical equipment, tobacco excise, e-commerce fee, ETIAS revenue, customs retention 25%→10%, inflation adjustment of the plastic resource; estimated additional revenue ~EUR 410,000M 2028-2035 (2025 prices); second sectoral package 3-Sep-2025. EP position: resolution of 7-May-2025 and broad consensus in the chamber for new own resources, including the digital route.
- **Digital-acquis fines (European Commission):** Google Shopping EUR 2,420M (2017; upheld by the General Court, 10-Nov-2021), Android EUR 4,340M (2018), AdSense EUR 1,490M (2019), adtech EUR 2,950M (Sep-2025); first DMA fines (23-Apr-2025): Apple EUR 500M, Meta EUR 200M. Competition fines enter the EU budget as "other revenue" (Council of the EU, "Financing the EU budget").
- **Real sanction ceilings:** DSA, Art. 52: up to 6% of annual worldwide turnover; AI Act, Art. 99: up to 7% of annual worldwide turnover or EUR 35M (the higher amount).
- **Precedents of resources never born:** financial transaction tax under enhanced cooperation (proposal Feb-2013, stalled negotiation); European digital levy shelved in 2021 pending OECD Pillar 1, from which the US withdrew in January 2025; national digital services taxes in force in several member states.
- **The book's own frameworks:** Ch. 9B (ROI 164%, EUR 1,820M/year), Ch. 13 (pure protector: the five verbs), Ch. 14 (double route and the consensus on exports as a possible future resource), Ch. 15 (revenue circle and endowment doctrine), Ch. 17 (EUR 2,600M/3,400M; 55,000-68,000 ecosystem jobs).

Chapter 17: Economic and Social Impact of the Protective Power

Cognitive protection as an engine of prosperity, employment and European technological leadership

Chapter Summary

This chapter demonstrates, with data updated to July 2026, that the Protective Power is not an expense, but **Europe's most strategic investment of the 21st century**. Based on recent studies by the **OECD (2025-2026)**, **McKinsey (2026)**, **World Economic Forum (2026)**, **European Commission (2026)** and **Eurostat (2026)**, its impact is quantified: **direct employment of the PP following the staged endowment trajectory of Ch. 15 — ~3,000 in the core (2035, baseline scenario), ~7,000-8,000 at full deployment (~2037) and more than 15,000 at maturity 2040+—, 48,000-60,000 indirect jobs in the accelerated scenario (55,000-68,000 total ecosystem jobs), EUR 7,200M-8,500M of annual GDP impact, and the creation of a leading European technology ecosystem.**

Central thesis: The Protective Power not only **protects democracy**, it **generates economic prosperity**, **attracts global talent and positions Europe as a leader in the digital era**. Every euro invested in cognitive protection is multiplied between **2.6x and 3.2x** in employment, innovation and growth, turning the EU into a **magnet for talent and investment** in a sector with a **CAGR of 13-20% per year**.


The Accelerated Scenario: Responding to the Urgency of 2026

As grounded in **Chapter 9**, the evidence accumulated in 2026 — the alarms raised by Bill Gates and Geoffrey Hinton, the 116,175 US layoffs attributed to AI between January and August, and an industry that modulates its warnings according to its share price — requires **accelerating the Protective Power's timeline as well**. If the EDSPA reaches full operational capacity in **2028** (Chapter 9's accelerated scenario), the Protective Power can launch in **2029** and reach its reference milestones **two years earlier**.

Unless otherwise indicated, the data in the following sections correspond to the **baseline scenario** (conservative). The table and charts in this section present the **recommended accelerated scenario**:


Metric	Baseline Scenario	Accelerated Scenario (recommended)
Protective Power launch	2030 (Treaty reform 2032-2033)	2029: preparatory structure (funded by the EDSPA's annual benefit surplus) → Treaty signature 2031-2032 → 2033: constitutional Protective Power
Initial start-up investment	EUR 2,600M	EUR 3,400M (2028-2029; +30% systemic-risk premium, consistent with Chapter 9)
Direct employment (milestone)	~3,000 by 2035 (PP core, baseline scenario — Ch. 15)	~7,000-8,000 by 2035 (full: 27 states, 24 languages, 24/7) → >15,000 at maturity 2040+
Indirect employment (milestone)	25,000-35,000 by 2035	48,000-60,000 by 2035 → total ecosystem: 55,000-68,000 jobs
Annual GDP impact (milestone)	EUR 7,200M-8,500M by 2035	EUR 8,000M-9,500M by 2033 → EUR 9,500M-11,000M by 2035
Accumulated impact	EUR 36,000M-42,500M (2035-2040)	EUR 55,000M-65,000M (2033-2040) : same mechanisms, two additional years of generation
Brain drain	Loss of 5,000-7,000 experts/year until 2035	Halted by 2029-2030 (the EDSPA retains talent); net attraction from 2031
Global market share (2039)	40% Europe	45-50% : earlier position as the standards benchmark
Impact per euro / Multiplier	270-330% / 2.6x-3.2x	~250-300% / 2.6x-3.2x : the multiplier is methodological and unchanged; ratio of impact per euro falls due to the larger investment, but the absolute benefit is higher by ~EUR 1,500M-2,000M/year

Employment generated (thousands of jobs) — baseline 2035 vs. accelerated scenario



Thousands of jobs. Baseline scenario: milestone in 2035. Accelerated scenario: milestone in 2033 and higher values by 2035.

Annual GDP impact (billions of EUR) — baseline vs. accelerated



Annual impact on European GDP. The accelerated scenario reaches in 2033 the level the baseline only reaches in 2035.

Institutional rigour note: a distinction is made between the Protective Power's **preparatory structure** (derived law, adoptable through enhanced cooperation: staff, infrastructure and technical capacity inherited from the EDSPA, operational from 2029) and its **constitutional approval**, which requires a European Convention, a unanimous IGC and ratification by all 27 Member States — a process historically no shorter than 4-5 years from the political launch. The accelerated timeline (Convention 2029-2030, signature 2031-2032, entry into force 2033) is the fastest defensible one, and is only plausible because the EDSPA will have delivered two years of measurable results.


Financial consistency with Chapter 9: the 2029 preparatory structure aligns with the EDSPA's 2028 cumulative break-even point (accelerated scenario, revised Chapter 9). The EDSPA's benefits fund part of the Protective Power's investment (EUR 2,600M baseline); the accelerated scenario's additional EUR 800M premium is covered by **less than half a year of the EDSPA's annual surplus of benefits over costs (~EUR 1,960M/year at full capacity)**. Rigour note: this is not a "structural surplus" — the EDSPA is not a business and does not collect revenue; its contribution is accounted for as economic value generated (Chapters 9 and 9B), not as budgetary income.

European Macroeconomic Context (2026)

The EU's base data (2026) show an economy in accelerated digital transformation, but with key challenges:

Indicator	Value
EU GDP	~EUR 18.8 trillion
Total EU employment	221.2 million
ICT specialists (2026)	10.5 million (target: 20M by 2030)
Companies using AI	20%
Digital infrastructure investment (needed by 2030)	EUR 40,000M

ICT specialists: current situation vs. European target




Strategic context: The EU needs **20 million ICT specialists by 2030**, but has only **10.5 million in 2026**. **58% of Europeans** would switch to EU providers even at a higher cost (security, data protection). **EUR 40,000M** in digital infrastructure is needed by 2030.

Section 1: The Protective Power as a Generator of High-Quality Employment

1. Direct Employment: the trajectory of a power — from the EDSPA's 1,500-2,000 to more than 15,000 at maturity

The Protective Power's staffing is not a fixed payroll: it is the **staged endowment trajectory of Chapter 15**, with accountable triggers (accessions of states, coverage of the mission, volume of threat, expansions authorised by the legislator): **1,500-2,000 (EDSPA 2027-2030) → ~3,000 (core 2033-2035) → ~7,000-8,000 (full ~2037: 27 states, 24 languages, 24/7) → more than 15,000 (maturity 2040+)**. Three masses of staff make it inevitable: the real scale of the operation (with Chapter 8's human-AI dual verification, deliberately non-automatable), the cognitive education corps (target ≥90% by 2038) and the organic growth of mandates that accompanies every European power — Frontex went from dozens of analysts to a statutory corps of 10,000 by 2027. Chapter 13's Citizen Checker is integrated in the maturity horizon, without a figure of its own. The distribution by area at maturity (2040+):

Direct employment of the Protective Power at maturity (2040+) — distribution by area (>15,000)




Rigour note: the distribution by area shows **maturity (2040+)**, not an entry payroll: the endowment follows the staged trajectory of Chapter 15 and every jump has an accountable trigger. The figure does not include Chapter 13's Citizen Checker (integrated without a figure of its own) and grows with the expansions of mandate the legislator authorises. The PP's direct employment is, by design, **Europe's largest public employment pole of non-automatable human judgement**: Chapter 8's dual verification is an institutional principle, not a transitional phase.

2. Indirect Employment: 25,000-35,000 Positions in the Ecosystem

The multiplier effect of the Protective Power extends its impact to the technology, consulting, infrastructure, education and legal sectors:

Indirect employment generated by the Protective Power (baseline 2035; 48,000-60,000 in the accelerated scenario)



Effect on the European labour market (2026): Cybersecurity demand in Europe: ~50% of organizations. Net AI-driven hiring growth in Europe: **+27% in 2026**. Most in-demand skills: Python (56%), AI/ML (47%), Java (34%), cybersecurity (33%).


Section 2: The Economic Multiplier Effect (2.6x-3.2x)

1. Calculation Methodology

Based on Input-Output models and the Keynesian multiplier:

- **Input-Output model (Leontief):** Estimation of impact on other sectors.
- **Keynesian multiplier:** $1/(1-MPC)$, where $MPC \approx 0.6-0.7 \rightarrow$ Multiplier $\approx 2.5x-3.3x$.
- **Empirical validation:** Global cybersecurity market (\$220B, CAGR 13%) $\rightarrow$ Implied multiplier: $\sim 3.0x$.


Economic multipliers by investment area



2. Multiplier Analysis by Investment Area


Detailed breakdown of the economic impact of the Protective Power (2035):

Economic impact by investment area (2035) — direct vs. total impact (EUR million)



Comparison with other sectors (2026):

Comparison of economic multipliers by sector




Section 3: Impact on European GDP (EUR 7,200M-8,500M+ Annually)

1. Calculating the Direct and Indirect Impact

Breakdown of GDP impact (2035) — Input-Output methodology:

Impact on European GDP by component (2035, EUR M)



Accumulated impact (2035-2040): Additional GDP: EUR 36,000M-42,500M (0.20%-0.23% of EU GDP). Ecosystem jobs: 28,000-38,000 positions (direct + indirect) in the baseline scenario 2035; 55,000-68,000 in the accelerated scenario. Additional annual growth: +0.04%-0.046% of GDP (equivalent to adding the GDP of Cyprus or Slovenia every year).

2. Comparison with Other Public Investments

Comparative Impact per euro invested and multipliers:

Investment	Annual Cost	GDP Impact	Impact per euro invested	Multiplier	Jobs Generated
Protective Power	EUR 2,600M-3,400M	EUR 7,200M-8,500M	270-330%	2.6x-3.2x	55,000-68,000 (ecosystem, accelerated)
Horizon Europe (R&D)	EUR 15,000M	EUR 22,500M	150%	1.5x	200,000
Digital Europe Programme	EUR 7,500M	EUR 10,500M	140%	1.4x	50,000
NextGeneration EU Fund	EUR 800,000M	EUR 1,200,000M	150%	1.5x	1,500,000
Cybersecurity (EU)	EUR 5,000M-10,000M	EUR 15,000M-35,000M	300-500%	3.0x-3.5x	50,000-100,000


Conclusion: The Protective Power has **one of the highest impacts per euro invested** among European public investments, outperforming programmes such as Horizon Europe (150%) or Digital Europe (140%). Its multiplier (2.6x-3.2x) is comparable to cybersecurity (3.0x-3.5x) and higher than traditional infrastructure (1.5x-2.0x).

Section 4: Global Talent Attraction and Technology Ecosystem Development

1. The Protective Power as a Magnet for World-Class Experts

Comparison of talent-attraction models (2026), score from 1 to 10:

Talent attraction comparison (2026) — Europe vs. US vs. China




Key talent-attraction data (2026): Current brain drain: Europe loses 5,000-7,000 AI experts/year (to the US and China). Projection with the Protective Power: **net attraction of 2,000-3,000 experts/year (2035+)**. Talent countries of origin: US (30%), Asia (25%), Latin America (20%), Africa (15%), Eastern Europe (10%).

2. Creation of a European Technology Ecosystem

Ecosystem components (2035):

Technology ecosystem of the Protective Power (2035) — employment and revenue




Ecosystem benefits (2026-2035):

- Reduced technological dependence: 85% → <30% (2035).
- Creation of global standards: Europe as the **benchmark in ethical AI**.
- Attraction of foreign investment: +20-30% in cognitive-protection startups.

Section 5: International Comparison - Europe vs. US vs. China

1. Three Models of Digital Control in 2039

Global market share in cognitive-protection technology:




Reasons for European growth:

- **Unique model:** Combines **advanced technology + democratic values**.
- **Integrated ecosystem:** Protective Power + startups + companies + research.
- **Export of standards:** The EU becomes the **global benchmark**.
- **Talent attraction:** The best experts **prefer Europe**.

2. Detailed Model Comparison

Comparison of digital models (2039) — score from 1 to 10




Key message: In 2039, while **China celebrates its centenary with a model of authoritarian digital control**, Europe will present the world with a **democratic model of cognitive protection based on transparency, human rights and resilient technology**.

Conclusion: Not Just Protection, but Prosperity

The Protective Power is not an expense, but **the seed of a new European economy**:

Evolution of the Protective Power's impact (2028-2050)



"The 21st century will belong to whoever controls AI — and through it, the future."

Five reasons why the Protective Power will change Europe:

1. **Creates high-quality jobs** (28,000-38,000 ecosystem jobs in the baseline scenario 2035; **55,000-68,000 in the accelerated scenario** — with the endowment trajectory of Ch. 15: ~3,000 direct in the core → ~7,000-8,000 full ~2037 → **>15,000 at maturity 2040+**, Europe's largest public employment pole of non-automatable human judgement against AI).
2. **Generates economic growth** (EUR 7,200M-8,500M+ of annual GDP impact in the baseline scenario; **EUR 9,500M-11,000M in the accelerated scenario**).
3. **Attracts the world's best talent** (Europe as a magnet for AI and cybersecurity experts).
4. **Develops its own technology ecosystem** (startups, companies, research centres).
5. **Exports a democratic model** (Europe as the global benchmark in cognitive protection).

The window of decision: the 2026 evidence (documented in Chapter 9) indicates that **the cost of waiting exceeds the acceleration premium every year**. In the accelerated scenario, Europe would reach by 2033 the level of employment and economic impact that the baseline only reaches in 2035: **two additional years of prosperity and protection** that cannot be recovered later. Every year of delay is a year in which the brain drain continues, cognitive manipulation advances without a counterweight, and the labour market reorganizes without a safety net.

The success equation:

Cognitive protection + Public investment + Private ecosystem = European prosperity + Global leadership

Vision for 2039: Europe not only protects its democracy — it **exports it to the world**. The Protective Power is not just an institution to protect Europe, but the engine that **will make it the world's first democratic digital power**. While China celebrates its centenary with a model of **authoritarian control**, Europe will present the world with a model of **democratic protection** that will change the global balance of power.

Appendix references

- **Macroeconomic data (Eurostat, 2026):** total EU employment: 221.2 million persons (Q1-2026, seasonally adjusted); 20.0% of EU enterprises (10+ employees) used AI technologies (2025 data, published 2026); EU GDP: EUR 18.8 trillion (verified in Ch. 20).
- **ICT specialists (Digital Decade, Decision (EU) 2022/2481):** 10.4 million ICT specialists employed in the EU (2025), interpolated to ~10.5 million in 2026; target of 20 million by 2030 (State of the Digital Decade).
- **Citizen preferences (Special Eurobarometer, Digital Decade 2026, fieldwork Feb-Mar 2026):** 58% of Europeans would switch to an EU digital provider even at slightly higher cost; 85% back investment in EU-developed infrastructure.
- **Cybersecurity market:** ~\$220bn globally with ~13% CAGR (Gartner: \$213bn in 2025, ~\$240bn projected for 2026; MarketsandMarkets: \$219bn in 2025, 13.8% CAGR); basis of the empirical validation of the implicit ~3.0x multiplier.
- **2026 urgency evidence (Ch. 9):** 116,175 layoffs attributed to AI in the US between January and August 2026; public warnings by Bill Gates and Geoffrey Hinton; 99% of CEOs preparing cuts.
- **The book's own frameworks:** Ch. 9 (baseline and accelerated scenarios), Ch. 9B (ecosystem net ROI: 164% over five years, EUR 1,820M/year), Ch. 13 (human-AI double verification; Citizen Checker), Ch. 15 (endowment trajectory: 1,500-2,000 EDSPA → ~3,000 core 2033-2035 → ~7,000-8,000 full ~2037 → >15,000 maturity 2040+; Frontex as precedent: 10,000 statutory staff by 2027), Ch. 16 (own resources: EUR 2,600M baseline / EUR 3,400M full).
- **Methodological note:** the 2.6x-3.2x multiplier is estimated with input-output models (Leontief) and the Keynesian multiplier $1/(1-MPC)$, with $MPC \approx 0.6-0.7$. The «impact per euro invested» in the tables (270-330% / 2.6x-3.2x) includes direct, indirect and induced effects and must not be confused with the ecosystem's net ROI (164%, Ch. 9B). Indirect employment (25,000-35,000 baseline / 48,000-60,000 accelerated, 2035) is the book's own estimate based on sector employment multipliers; global market shares 2026-2039 and net talent attraction (2,000-3,000 experts/year from 2035+) are scenario projections, not observed data.

Chapter 18: Technology and Infrastructure of the Protective Power

A power that protects the cognitive sphere cannot depend on the tools of those who manipulate it: technological autonomy in stages, adaptive capex separated from the annual budget, and the closing of a race the rival is already running with a date attached — 2049

Chapter summary

This chapter starts from a premise the whole book shares: **we are building something that does not exist, but that is indispensable** — the first institution in history whose function is to protect the cognitive sphere of citizens against industrial-scale manipulation. What matters most is not the figures: it is **the citizen's security**, which demands a strong, safe and independent institution, with own resources. And such an institution carries a technological consequence that no annual budget can dodge: **whoever protects the cognitive sphere cannot depend on the tools of those who manipulate it**. The diagnosis is verified and brutal: **80% of Europe's digital infrastructure is imported** (EuroStack report, Bertelsmann/Francesca Bria, 2025), and the sovereignty-as-a-badge experiment that was Gaia-X taught the lesson — sovereignty is not declared, it is procured and it is built. But the real terrain of 2024-2026 is no longer what it was three years ago: JUPITER brought Europe's first exascale supercomputer into operation (Jun-2025), the EU selected 13 AI factories and launched the official gigafactory call (Jul-2026) mobilizing more than **EUR 30,000M**, and IRIS2 (concession signed 16-Dec-2024 for ~EUR 10,550M and ~290 satellites; the 7-Aug-2026 implementation agreement raises the programme to EUR 15,600M and 348 satellites, deployment from 2029, governmental services ~2030) lists disinformation campaigns among its declared threats. The chapter builds on that terrain with a **three-phase model**: anchor client of what exists (EDSPA 2027-2030) → first Central Forensic Facility of its own (core 2033-2035, capex multi-year and separate from opex) → expanded own network and own models (full ~2037 and maturity 2040+). And it adds a new piece: the **Strategic Acceleration Clause** — continuous monitoring, defined triggers, a decision within at most 3 months, funding injection through own resources (Art. 311 TFEU) — because the rival already plays this way: when EUV lithography was denied to China, China built its own, and its chip self-sufficiency rose from 16% (2024) to ~28% (late 2025). The chapter closes with the mirror that gives urgency to everything above: Xi Jinping's **Two Centenaries** have a date — 2049, the PRC centenary — and humanity will choose, on that date, between two finished models of governing the cognitive sphere: authoritarian control or democratic verification. The Protective Power is what makes the second option exist.

18.1 The premise: what does not exist and is indispensable


No market produces what the Protective Power buys, because nobody has ever held the mandate of protecting a cognitive sphere. No authoritarian state will produce it: industrial manipulation is their instrument of government, not their threat. And no reactive regulation suffices: the DSA and the AI Act punish after the harm, with whatever forensic means the regulated platforms grant them. The gap is exactly the one the book has described since Chapter 1: **the cognitive sphere is the only domain of democratic life that no power protects** — and filling that gap is, strictly speaking, creating something that does not exist.

Creating what does not exist carries a technical rule that this chapter turns into a principle: the protection tool cannot be the property of the party to be protected. A forensic scanner running on someone else's cloud audits with the audited party's permission; a model that classifies manipulation but is silently trained by third parties classifies what the third party wants classified. The institution of Chapter 7 is strong, safe and independent in legal terms; this chapter is the condition for it to be so technically as well. And a second rule, less obvious and more ambitious: by equipping itself with its own cognitive-protection technology, **the Union stops being a customer of other people's AI and becomes a strategic AI agent in the cognitive sphere** — humanity's real need for AI not to turn harmful but beneficial: regulating it after the fact is not enough; it must be demonstrated, with own infrastructure and permanent human verification, that AI can live under democratic control. The Union that achieves this will not have bought a system: it will have produced the civilizational *model* alternative to authoritarian AI control.

18.2 The verified diagnosis: 80% imported, Gaia-X and the real terrain

The starting point admits no cosmetic fixes. The **EuroStack** report (Bertelsmann Stiftung, led by Francesca Bria, 2025) puts at **80% the share of Europe's digital infrastructure that is imported**: from chip to model, from submarine cable to operating system. And the most-cited European digital sovereignty experiment — **Gaia-X** — teaches what not to do: a federation of standards and badges does not displace dominant providers if nobody buys; sovereignty is won in public procurement and in open source, not in certificates.

But the diagnosis has a second half that defeatism omits: **the real terrain of 2024-2026 has already changed**. JUPITER (Julich) became Europe's first **exascale** supercomputer — operational since June 2025, Europe's fastest and fourth in the world on the TOP500 list of November 2025 — and rests on a real pre-exascale constellation: LUMI (Kajaani), Leonardo (Bologna), MareNostrum 5 (Barcelona). On that base, the AI factories policy went from 7 consortia (Dec-2024) to **13 factories selected in 7 countries**, in deployment through 2025-2026. And the scale jumped again with **InvestAI** (Paris Summit, Feb-2025): EUR 200,000M mobilized with a EUR 20,000M EU fund for AI gigafactories of ~100,000 accelerators each; the Council adapted the EuroHPC regulation on 16-Jan-2026 to include them, the informal consultation received **77 proposals from 16 Member States (60 sites)**, and the official call published in July 2026 unlocks **more than EUR 30,000M** in investment. In space, **IRIS2** (concession signed 16-Dec-2024, SpaceRISE consortium, ~EUR 10,550M, 12 years, ~290 satellites from 2029 — reinforced by the 7-Aug-2026 implementation agreement: EUR 15,600M and 348 satellites) is the first European infrastructure to expressly list disinformation campaigns among the threats it must withstand.



18.3 The operating principle: three phases, separate capex and the Strategic Acceleration Clause

No budgetary heresies: no data centers by decree (10 exabytes at "military level" for EUR 2,000M), no supercomputer table that understates reality tenfold, no own networks for EUR 1,050M, no portfolio of magic systems with fixed parameters and prices. No phantom capex of EUR 5,150M that fits in no budgetary framework. The chapter does the opposite: **use what exists, buy what is missing, and build only what nobody can build for the Union** with the financial architecture Chapter 16 already established: **multi-year capex separate from annual opex**, endowed through own resources (Art. 311 TFEU) and executed in stages.


The three phases

Phase 1 — EDSPA as anchor client (2027-2030). The agency does not build: it rents, buys and anchors. Compute on EuroHPC and the AI factories; secure services on existing networks and IRIS2 when it enters service (~2029-

2030); commercial and open-source tools wherever they pass the dual verification of Chapter 8. Being the anchor client of European capacity is already technology policy: institutional forensic demand forces European infrastructure to serve the protection of citizens, not only to compete in benchmarks.

Phase 2 — The core builds (2033-2035). With the Treaty in force, the PP raises its first piece of institutional property: the **Central Forensic Facility** — custody of evidence (certified chain of custody of manipulation materials, executable and reproducible before the courts) and training and audit compute for the own models of section 18.4. The capex is multi-year, separate from the annual operating endowment (Chapter 16), and its design follows the Gaia-X lesson: European procurement and open source in every layer where a credible alternative exists.

Phase 3 — Own network and maturity (2037-2040+). The full PP (27 Member States, 24 languages, 24/7) operates a distributed forensic network connected to IRIS2, deploys its own models across the Union and co-invests with European private capital following the InvestAI pattern: the State provides the mission and the demand; the European market provides the scale. Exports to democracies — the horizon of Chapters 14/16 — are born from this phase, without figures and without assumptions.



The Strategic Acceleration Clause

The technology race does not respect five-year plans, and the rival has proven it with the most-cited case of the decade: when the US pressured to deny EUV lithography exports to China, **China built its own alternative** — SMIC fabricated the 7nm Kirin 9000s without EUV (Huawei Mate 60 Pro, Aug-2023), SMEE delivered its first 28nm lithography machines and filed an EUV patent, Huawei launched the Ascend accelerators (910C; 950 series with own HBM in 2026) and Chinese semiconductor self-sufficiency rose from **16% in 2024 to ~28% in late 2025**, with ~US\$150,000M in state subsidies and rising declared self-sufficiency goals: the original 70% of the Made in China 2025 plan and the new 80% target for 2030 (2026 plan). CSIS's analysis is unambiguous: export controls *accelerated* Chinese self-sufficiency instead of slowing it. The lesson for Europe is not to copy authoritarianism: it is **not to be the only power in the race playing without an acceleration mechanism**.

That is why the design gives the PP — and the Union — an auditable institutional mechanism, not a blank checkbook: the **Strategic Acceleration Clause**. Its design has five pieces:

1. **Cruising speed by default.** The multi-year capex of the phases runs at the planned pace. No trigger, no acceleration: budgetary rigor.
2. **Continuous strategic monitoring.** the Cognitive Intelligence Department (Chapter 12) maintains a technology radar with an annual internal report — the *State of the Cognitive Race* — of which a **synthetic public summary** is published: transparency before the citizen without handing the rival the full map of what is being monitored.
3. **Triggers defined in advance.** Four hard triggers open the extraordinary review; three soft springs open the ordinary review (table below). Defined now, revised annually, never improvised under pressure.

4. **Decision within at most 3 months.** The trigger opens an extraordinary capex review that must end in a decision within one quarter — in the digital era, twelve months is an eternity; the EU's real urgent procedures run in weeks-to-months, and the protection of the cognitive sphere cannot wait for the annual budget cycle.
5. **Injection through own resources.** The acceleration is always funded through Art. 311 TFEU (Chapter 16): never with money from the PP's fines, never as income of the power that detects — the pure protector remains intact here as well.

Trigger / spring	Nature	Response
T1 · Technology veto against the EU — any actor blocking access to compute, chips, models or critical infrastructure for the PP	Hard	Extraordinary capex review; decision in 3 months or less; injection via Art. 311 TFEU
T2 · Verified qualitative leap in manipulation AI — massive cognitive manipulation capability that the forensic Scanner cannot reliably characterize	Hard	Extraordinary capex review; decision in 3 months or less; injection via Art. 311 TFEU
T3 · Strategic compute gap — the EU's sovereign compute falls below the relative threshold set by regulation against a rival	Hard	Extraordinary capex review; decision in 3 months or less; injection via Art. 311 TFEU
T4 · Critical interdependence detected — an essential part of the PP's chain depends on a single provider or a non-European jurisdiction	Hard	Extraordinary capex review; decision in 3 months or less; injection via Art. 311 TFEU
S1 · Operational overload — cases sustained above the threshold (scaling clause of Chapter 12, >150-200 cases/year)	Soft	Ordinary capex review in the next multiannual cycle
S2 · Sovereignty dashboard — annual supply-chain audit: share of European components by layer (Fig. 18.3)	Soft	Report to the Council and the EP; layer-by-layer correction plans
S3 · Annual stress exercise (red team) — simulate the total cutoff of external dependencies and measure the PP's resilience (banking-style stress tests)	Soft	Findings to the synthetic public summary; continuous feed of T1-T4

Thresholds and archetypes — how it fires in practice: the text fixes the triggers by institutional definition; their *exact thresholds* (e.g. the magnitude of the compute gap in T3) would be set by Council regulation on a proposal from the PP, revised annually with the *State of the Cognitive Race* report. As an illustration of how the mechanism works, these would be trigger archetypes, calibrated with real 2024-2026 cases: a veto on AI accelerator exports to the Union (the precedent of the EUV veto on China, 2019/2023) would fire T1; a manipulation AI outperforming the forensic Scanner by an order of magnitude would fire T2; and China's 16%→28% self-sufficiency trajectory in 18 months (~US\$150,000M in subsidies; Ascend 950 with own HBM, 2026) shows the speed at which a rival can close a gap — and therefore the speed at which the radar must detect it. The 3-month decision deadline is an institutional design proposal of this book, aligned with the real timescale of the Union's urgent procedures, not an existing rule.


18.4 Own AI by requirement: the forensic Scanner and the Auditor's and Educator's tools

The PP's technological heart is not a product list: it is **three core tools defined by institutional requirement**, the first of which has existed since Chapter 8. Their names are functional and deliberately sober, and none is defined by model parameters — that is decided by procurement and by the state of the art, not by a book:

- **The cognitive forensic Scanner** (Chapter 8): the instrument for detecting and characterizing manipulation — not an arbiter of truth. Its core requirement: that its evidence be **reproducible before the courts**, with a certified chain of custody (Central Forensic Facility, Phase 2) and human-AI dual verification of every finding.
- **The Auditor's toolkit** (Chapter 13): algorithmic audit tools for recommendation, advertising and moderation systems on the platforms — researcher data access (DSA, Art. 40) turned into systematic scrutiny, with models able to explain *why* a system amplifies or buries content, not only to measure that it does.
- **The Educator's toolkit** (Chapter 13): the layer that adapts the Scanner's evidence for citizens, schools and professionals — cognitive vulnerability diagnosis at scale (Chapter 2's Kit in positive), verifiable curricula and simulations in 24 languages.

All three share the same requirement contract, and this is the point where the chapter connects with the premise of 18.1: **open, auditable and explainable** — open source where a credible alternative exists, public audit of their components, explainability of every finding, and the human-AI dual verification of Chapter 8 as a condition of

use. None of the three can be a black box, because the power that uses them cannot be one. A Scanner that certifies manipulation by secret method is an oracle, not an institution; and the PP of this book is deliberately the opposite: **the place where AI is subjected, for the first time and with its own means, to permanent democratic verification**. There — not in the parameters of any model — lies the civilizational bet: proving that a community of free citizens can have artificial intelligence under its control, and that this is the only basis on which humanity can keep AI from turning harmful and make it beneficial.



18.5 The mirror of 2049: the Two Centenaries and the race for the cognitive sphere

Any serious strategy has a date, and the rival's has been published since 2012: at the Party's XVIII Congress, Xi Jinping set the **Two Centenaries** — the xiaokang society by 2021 and, by **2049, the centenary of the People's Republic of China, a "modern, prosperous and strong socialist country"** — with the 2035 milestone and its technological annex, China Standards 2035. The rival's AI doctrine is equally explicit: at the April 2025 Politburo session, Xi demanded "independent and controllable" AI (自主可控) and self-reliant AI (自立自强), and the 15th Five-Year Plan (2026-2030) made semiconductor and AI self-sufficiency a generational national project. None of this is a prophecy of this book: they are public commitments with a calendar from the adversary. **The rival has a 100-year strategy; the democracies have neither a strategy nor a date.**

The Protective Power gives Europe the two things it lacks. A strategy: the protection of the cognitive sphere as a permanent mission, with its own infrastructure and an acceleration mechanism. And a counterweight date: **the full PP by ~2037, maturity by 2040+** — symmetrical with the Chinese 2035 milestone, and in time for the Union to arrive with its head held high at the rival's centenary in 2049. In 2049 the world will not choose between two speeds of progress: it will choose between **two finished models of governing the cognitive sphere** — authoritarian control, where AI is the instrument of power, and democratic verification, where AI lives under permanent citizen control. This book does not prophesy the winner. But it does assert, in both directions, what this chapter's history makes evident: **if the second model does not exist, the choice of 2049 will not be a choice**. And if it does exist — mature, verified and operational by 2040+ — the European model of cognitive-sphere-protecting AI will be **the strategic reference for all the world's democracies**: every democracy facing industrial manipulation — state, bloc or democratic organization — will need exactly what the PP will already have built: a forensic scanner reproducible before the courts, explainable algorithmic audit, cognitive education at scale, and the verified demonstration that AI can serve freedom instead of domesticating it. That is the ultimate geopolitical meaning of the export horizon of Chapters 14/16 — without figures, but without false modesty: **whoever holds the finished protective model will not be just another actor in 2049; it will be the reference the democracies look to**, just as whoever controls the authoritarian model will be the reference the autocrats look to. Europe can arrive at 2049 as a customer of other people's AI — or as the democracy that produced the model of cognitive-sphere-protecting

AI. The Protective Power is what makes the second option exist — and the race to be won is, quite literally, for the freedom of humanity.

18.6 Conclusion

1. **The premise.** The PP is something that does not exist and is indispensable: the first institution that protects the citizen's cognitive sphere. Its technological requirement is not large figures: to be strong, safe and independent — which means not depending on the tools of those who manipulate.
2. **The diagnosis is twofold.** 80% of Europe's digital infrastructure is imported (EuroStack) and Gaia-X teaches that sovereignty is not declared — but the real terrain of 2024-2026 (JUPITER exascale, 13 AI factories, InvestAI gigafactories with the official Jul-2026 call worth EUR 30,000M+, IRIS2 from 2029) is the real scaffolding the PP builds on.
3. **Three phases, separate capex.** Anchor client (2027-2030) → own Central Forensic Facility (2033-2035) → own network and own models (full ~2037, maturity 2040+), with multi-year capex via Art. 311 TFEU, never phantom capex by decree.
4. **The Strategic Acceleration Clause.** Cruising speed by default, an annual radar with a synthetic public summary, 4 hard triggers + 3 soft springs defined in advance, a decision in 3 months or less, injection through own resources — because the rival already plays this way: denied EUV lithography, it built its own and went from 16% to 28% self-sufficiency in 18 months.
5. **AI by requirement, not by parameters.** The cognitive forensic Scanner, the Auditor's toolkit and the Educator's toolkit: open, auditable, explainable, with human-AI dual verification — the PP as the place where AI answers, with its own means, to the citizen, permanently.
6. **The mirror of 2049.** The rival's Two Centenaries have a public date; the PP gives Europe a strategy and a counterweight date (full ~2037, maturity 2040+). No prophecy of victory — but relevance earned the hard way: if the model matures, the PP turns Europe into the strategic reference for the world's democracies and the exporter of the cognitive-sphere-protecting AI model (Chapters 14/16 horizon, without figures).

The chapter's thesis in one sentence: the cognitive sphere is not protected with the tools of those who manipulate it — the Protective Power builds its autonomy in stages on Europe's real terrain, gives itself an auditable acceleration mechanism because the rival is already running with a date (2049), and turns the Union into the only place in the world where artificial intelligence answers, with its own means, to the citizen.

18.7 Sources and verified references

- **European diagnosis:** EuroStack report (Bertelsmann Stiftung, led by Francesca Bria, 2025): 80% of Europe's digital infrastructure imported; the Gaia-X lesson (sovereignty via procurement and open source).
- **European compute:** EuroHPC JU: JUPITER (Julich), Europe's first exascale supercomputer, operational Jun-2025 (1st in Europe and 4th worldwide, TOP500 Nov-2025); LUMI (Kajaani), Leonardo (Bologna), MareNostrum 5 (Barcelona) pre-exascale.
- **AI factories and gigafactories:** 7 consortia selected Dec-2024 (15 MS) → 13 AI factories in 7 countries (2025), in deployment 2025-2026; InvestAI (Paris Summit, 11-Feb-2025): EUR 200,000M mobilized, EUR 20,000M EU fund for 4-5 gigafactories of ~100,000 accelerators; Council of the EU, 16-Jan-2026: EuroHPC regulation (Reg. 2021/1173, amended by Reg. 2026/150) adapted to include gigafactories; informal consultation: 77 proposals from 16 MS / 60 sites (Jun-2025); official call Jul-2026 with more than EUR 30,000M in investment (European Commission).
- **IRIS2:** concession signed 16-Dec-2024 (SpaceRISE consortium), 12-year PPP, ~EUR 10,550M and ~290 satellites (deployment from 2029, governmental services ~2030); 7-Aug-2026 implementation agreement (Commission/ESA/SpaceRISE) raising the programme to EUR 15,600M and 348 satellites; disinformation campaigns among its declared threats.
- **Cybersecurity and encryption:** NIS2 Directive (in force 18-Oct-2024); Cyber Resilience Act (notification 11-Sep-2026, obligations 11-Dec-2027); ENISA; post-quantum: NIST FIPS 203/204/205 (2024), transition roadmap (RSA/ECC deprecated 2030, disallowed 2035, IR 8547).
- **The rival's case (Chip War):** EUV lithography export denial to China under US pressure (2019/2023); Huawei Mate 60 Pro (Aug-2023) with SMIC's EUV-free 7nm Kirin 9000s; SMEE: first 28nm lithography machines (late 2023) and EUV patent (Mar-2023); Huawei Ascend 910C (60-80% of NVIDIA's inference performance) and

Ascend 950 series with own HBM (Q1 2026); Chinese semiconductor self-sufficiency 16% (2024) → ~28% (Q4 2025), ~US\$150,000M in subsidies, declared goals: the Made in China 2025 plan's 70% (not reached) and 80% for 2030 (2026 plan) (CSIS; Huawei/Ren Zhengfei before Xi, Feb-2025); CSIS analysis: the controls accelerated Chinese self-sufficiency; Politburo session (Apr-2025): Xi demands "independent and controllable" AI (自主可控) and self-reliant AI (自立自强); 15th Five-Year Plan (2026-2030): semiconductor and AI self-sufficiency as a generational national project.

- **The Two Centenaries:** set out by Xi Jinping at the CCP's XVIII Congress (2012): xiaokang society by 2021 and a modern socialist country by 2049 (PRC centenary), with the 2035 milestone and China Standards 2035.
- **Exported digital authoritarianism:** the Digital Silk Road (CNAS reports): export of the surveillance-and-control model to Africa, East Asia and Latin America.
- **The book's own frameworks:** Chapter 7 (structure and the Cognitive Intelligence Department), Chapter 8 (cognitive forensic Scanner, human-AI dual verification), Chapter 12 (scaling clause >150-200 cases/year; Digital Sovereignty Department), Chapter 13 (the five verbs; algorithmic audit; education), Chapter 14 (dual constitutional route; exports as a horizon without figures), Chapter 15 (staffing trajectory >15,000 at maturity 2040+), Chapter 16 (own resources Art. 311 TFEU; capex separate from opex), Chapter 17 (ecosystem of 55,000-68,000 jobs).

Chapter 19: Communication Strategy of the Protective Power

Chapter summary

A constitutional power does not communicate the way a product does: it communicates the way trust does. This chapter's thesis is that the Protective Power's communication is not add-on marketing: it is **the institutional function that closes the circle of legitimacy** — because a power that protects the cognitive sphere without the citizen being able to see it, verify it and understand it does not exist politically. The starting point is a verified paradox that this chapter turns into its axis: Europeans are more worried than ever about disinformation (69% consider it a problem in the 2026 Eurobarometer, up to 85% in Spain) and at the same time a majority (60% in the OECD Trust Survey 2026) distrusts governmental use of AI. The protection they demand and the tool that would provide it generate, today, the same distrust. That paradox is not an obstacle for the PP: it is **exactly the problem the PP was designed to solve** — the only imaginable institution whose AI is subjected by design to human-AI dual verification, transparency by default and evidence reproducible before the courts (Chapters 8, 13). The strategy rests on the best terrain in a generation: trust in the EU reached in 2025 its highest level in 18 years (52%, ~60% among 15-24 year-olds) precisely while trust in news fell to its historic low (37%, Reuters Institute 2026) and concern about disinformation rose from 28% (2022) to 36% (2025). The chapter sets four principles — **evidence before opinion, transparency by default, dual verification in communication too, and learning from failures** (the Conference on the Future of Europe, the public consultations with unclear objectives documented by the European Court of Auditors) — and one flagship instrument: the **Annual Cognitive Protection Report**, which publishes milestones and results, never procedures or techniques. It closes with the 2027-2040+ calendar aligned with the book's phases and with Chapter 18's horizon: communication that demonstrates is what turns the PP into the reference the world's democracies will buy in 2049.

19.1 The thesis: making protection visible

Every new institution fights two battles: to exist and to be understood. The Protective Power fights a third no institution has fought before: **to be believed in exactly the territory where belief is broken**. Its function is to protect the cognitive sphere — the territory where citizens have stopped believing what they read (trust in news: 37%, the lowest since 2015), distrust those who think differently (7 in 10, Edelman 2026) and suspect every machine that decides about them (60% skeptical of government AI, OECD 2026). On that terrain, a conventional communication strategy — slogan, campaign, hashtag — would not just fail: **it would be proof of the critic's accusation** — one more actor competing for attention with persuasion techniques.


That is why the PP's strategy is the opposite: **making protection visible**. The PP is not defined by promising what it will do, but by showing what it has already done and anyone can check: the forensic report with its chain of custody, the case file the Commission sanctioned, the manipulating material exposed with its reproducible method, the lesson a teacher in Bremen or Lisbon can download and use. Every act of PP communication must be **a verifiable demonstration, not an assertion**. And the goal is deliberately ambitious, because the moment allows it: it is not that the citizen tolerates the PP, it is that the citizen **buys the idea** — understands that cognitive protection is not one more control, but the condition of their freedom; that the Europe building it is not the Europe that regulates late, but the one producing the model of AI under democratic control the world will expect in 2049. This is not the time for false modesty: it is time for the evidence to speak — and for the Union to dare to say what the evidence proves.

19.2 The verified terrain: the paradox of 60%

The 2025-2026 data draw the most difficult and the most favorable terrain at once. The difficult: trust is broken on every front. Trust in news fell to **37%** — an all-time low since the Reuters Institute began measuring it across 48 countries (DNR 2026) — and 42% of citizens avoid the news, up from 29% in 2017. Insularity advances: in the 2026 Edelman Trust Barometer (34,000 respondents, 28 countries), 7 in 10 hesitate to trust those with different val-

ues or information sources. And the decisive datum for this chapter: the OECD Trust Survey (Jun-2026) shows that **60% of citizens are skeptical of their governments' use of AI**. Europeans demand protection — **69%** consider disinformation a problem in the 2026 Eurobarometer (85% in Spain, 82% in Cyprus, 80% in Italy), and 36% say they are exposed to it "often" or "very often" in a given week (28% in 2022) — but they distrust the instrument that would protect them. **They ask for a shield and suspect the smith.**

The favorable terrain: the Union arrives at this battle in its best moment of trust in a generation. The Spring 2025 Eurobarometer recorded trust in the EU at its **18-year high: 52%**, with around 60% among 15-24 year-olds — the generation that will consume its entire information life on the terrain the PP protects. And the paradox of 60% has a strategic reading that is this chapter's axis: **skepticism of government AI is universal and will also reach the rivals** — the authoritarian model cannot solve it because its AI is by definition uncontrolled by the citizen; the market model cannot solve it because its AI optimizes profit, not verification. Only the PP's model — AI subjected to human-AI dual verification, open source where possible, public audit and evidence reproducible before the courts — answers by design the question the 60% ask. **The paradox of 60% is not the PP's obstacle: it is its strongest selling argument**, because it is the only one in the institutional market whose answer is not a promise — it is an architecture.



19.3 The message framework: buying the idea

The PP's communication task is a transformation of perceptions — the one the Union achieved with the euro or with Erasmus: from technical abstraction to personal right. The framework turns every objection into its answer with evidence from the book itself:

Perception to transform	Protective Power message	Evidence supporting it
"It is a cost"	It is an investment with measurable return	ROI 164% over 5 years; EUR 1,820 million/year of net benefit (Chapter 9B) — an additional argument, never a substitute for the mission (Chapter 15)
"It is control"	It is freedom from industrial manipulation	The PP neither censors nor punishes: it detects, audits, publishes, refers and educates (Chapter 13) — the pure protector
"It is bureaucracy"	It is European technology and innovation	Forensic scanner, algorithmic audit and cognitive education in 24 languages on own infrastructure (Chapter 18)
"It is restriction"	It is protection of rights	The cognitive sphere as a fundamental right of the digital era (Chapters 1, 15)
"It is dependence"	It is strategic autonomy	From 80% imported to autonomy in three phases, with the Strategic Acceleration Clause (Chapter 18)
"It is another late regulator"	It is the 4th power: verification before sanction	Legal basis in Art. 48 TEU, own resources under Art. 311 TFEU (Chapters 15, 16) — the power that detects never collects what it detects
"What about its AI?" (the paradox of 60%)	Our AI is the only one under permanent public verification	Human-AI dual verification, open, auditable and explainable tools (Chapters 8, 18) — the architectural answer, not a promise

The central message: the Protective Power is not a brake — it is the accelerator of European freedom: of the citizen's autonomy from manipulation, of the ecosystem's skilled employment (55,000-68,000 jobs, Chapter 17), and of the Union's leadership as the democracy that made AI answer to the citizen. And the proof is not the slogan: it is that anyone can verify every PP claim with the same evidence a court would use.

19.4 The four principles: demonstrate, open, verify, remember

- Evidence before opinion.** Every relevant PP communication is anchored to public, reproducible evidence: the forensic report, the referred case file, the audited figure. The PP never asks to be believed: it shows where the file is. On the terrain of broken trust (37%), the only accepted currency is verifiable currency.
- Transparency by default.** The PP applies to itself the standard it demands: the right of access to documents (Art. 15(3) TFEU and Regulation 1049/2001) as the rule, not the exception, proactive publication of its criteria, decisions and reports — including the synthetic public summary of the "State of the Cognitive Race" radar (Chapter 18). The proportionality rule the book sets everywhere: **the PP publishes milestones and results; never procedures or techniques** — the same line separating the public radar from the internal one, because transparency toward the citizen cannot hand the methodology to the rival.
- Dual verification in communication too.** No public PP claim about a case, a figure or a finding is published without passing the same standard the PP demands of the platforms: human review of the machine's output, source traceability, visible correction when it is wrong. The paradox of 60% closes in every communicative act: the citizen sees, again and again, that the PP's AI never speaks alone.
- Learning from failures.** The Union already ran the opposite experiment: the Conference on the Future of Europe (2021-2022) — 800 citizens in panels, 178 recommendations, 49 proposals, 5 million platform visits and 700,000 participants — was the most ambitious participatory exercise ever attempted, and its political balance fell short; the European Court of Auditors documented public consultations with unclear objectives and insufficient outreach; and the doctrine's warning is old and accurate: consultations that lead to nothing deflate the souffle and drain trust. The lesson for the PP is twofold: **no decorative participation** — every citizen-listening mechanism of the PP must be tied to a decision with a date and a visible effect — and no promises without follow-up: the calendar the PP publishes, the PP keeps.

19.5 The five audiences and the 2027-2040+ calendar


The five audiences

- Citizens.** The message: "your decisions, your cognitive sphere, protected with evidence you can see." The channel: the Annual Cognitive Protection Report, the public forensic reports on landmark cases, the Educator's

materials in 24 languages. The goal is not that the citizen knows the Treaty: it is that they know *where to look* when they suspect manipulation.

- **Policy-makers.** The message: viability and window — legal basis in Art. 48 TEU, financing through own resources (Art. 311 TFEU, Chapter 16), ROI 164% (Chapter 9B), realistic calendar (constitutional PP in 2033, Chapter 17). And this chapter's decisive argument: **the paradox of 60% is impossible to answer with more regulation — only with an institution of permanent verification.**
- **The European tech sector.** The message: stable institutional demand, European procurement and open source (the Gaia-X lesson, Chapter 18) — the PP as anchor client turning European capacity into protection infrastructure; the 55,000-68,000-job ecosystem (Chapter 17) as proof that cognitive protection is a sector, not a burden.
- **News media.** The message: the PP does not compete with journalism — it supplies it with verifiable evidence. For an industry at a historic trust low (37%), the PP is a source no one else can offer: reproducible forensic material, no interested leaks. The rule: the PP treats every outlet with the same access standard — no favorites, no exclusions.
- **The international community.** The message: the European model of cognitive-sphere-protecting AI as the reference for the world's democracies (Chapter 18) — the third way between algorithmic capitalism without verification and authoritarian control without freedom, exportable as an institutional model to every democracy facing industrial manipulation.

The calendar: demonstrate, mobilize, account, lead



The flagship instrument: the Annual Cognitive Protection Report

If the PP had to sum up its communication in a single act, it would be this: the **Annual Cognitive Protection Report**, published every year in the Union's 24 languages, presented before the European Parliament and open to any citizen. Its content sets this chapter's rule: **milestones and results, never procedures or techniques** — which threats were characterized and with what reach, which case files were referred and what the Commission and the Member States did with them, which Educator materials were used in schools, where the PP failed and what it corrected. What the Report does not publish — forensic method, Scanner capabilities, radar architecture — is not opacity: it is the same line of defense as in Chapter 18, where transparency answers to the citizen without handing the map to the rival. The Report is at once the most demanding act of accountability any European institution has ever published — and the PP's best selling argument: **a power that shows its results, including its mistakes, is the only power the skeptical 60% can begin to believe.**

19.6 Conclusion: communicating to lead

1. **The PP's communication is an institutional function, not marketing.** A power that protects the cognitive sphere exists politically only if the citizen can see it, verify it and understand it: making protection visible is part of the mission, not the promotion.
2. **The paradox of 60% is the axis and the argument.** Citizens demand protection from disinformation (69%) and distrust government AI (60%). No rival model can resolve that contradiction; the PP resolves it by architecture: dual verification, transparency by default, reproducible evidence.
3. **The terrain is the best in a generation.** Trust in the EU at an 18-year high (52%, ~60% among the young) while information trust hits historic lows (37%); the Union arrives at the battle for the cognitive sphere with growing authority and a gap exactly the size of the PP.
4. **Four principles and one rule.** Evidence before opinion, transparency by default, dual verification in communication too, learning from failures; and the proportionality rule: milestones and results, never procedures or techniques.
5. **One instrument and one calendar.** The Annual Cognitive Protection Report as the act of accountability; demonstrate (2027-2030), mobilize (2030-2032), account (2033-2035), lead (2040+) — toward Chapter 18's mirror of 2049: communication that demonstrates is what turns the PP into the reference the world's democracies will buy.

The chapter's thesis in one sentence: on the terrain where nobody believes anybody, the only power that can win the battle for trust is the one that does not ask to be believed — the Protective Power makes protection visible: verifiable evidence, transparency by default, AI that never speaks alone, and the ambition that the Europe protecting its citizens' cognitive sphere will be, in 2049, the model the democratic world will have bought.

19.7 Sources and verified references

- **Perception of disinformation (Eurobarometer, European Commission):** Eurobarometer on disinformation (2025 edition, 26,500+ respondents): 36% exposed "often" or "very often" in the last 7 days (28% in the 2022 Eurobarometer); Eurobarometer published Feb-2026: 69% in the EU consider disinformation a problem in their country (85% Spain, 82% Cyprus, 80% Italy); 12% "very confident" and 52% "fairly confident" of recognizing disinformation.
- **Information trust:** Reuters Institute Digital News Report 2026 (nearly 100,000 respondents, 48 countries, with the University of Oxford): trust in news 37% — lowest since 2015; news avoidance 42% (29% in 2017); social and video platforms overtake TV as the global news source; news interest falling (-9 pp in Europe 2021-2026).
- **Institutional trust and polarization:** Edelman Trust Barometer 2026 (34,000 respondents, 28 countries, fieldwork Oct-Nov 2025): grievance has turned into insularity — 7 in 10 unwilling or hesitant to trust someone with different values, approaches or information sources; the global trust-brokering challenge.
- **Trust in governments and in public AI:** OECD Trust Survey 2026 (published 29-Jun-2026): trust in national governments ~40% across the OECD; 60% of respondents skeptical of AI use by governmental institutions (55% in Germany); positive experiences with public services strengthen trust in state institutions — administrative quality matters more than political promise.
- **The European window:** Spring 2025 Standard Eurobarometer: trust in the EU 52% — an 18-year high; ~60% among 15-24 year-olds; 75% feel like citizens of the EU; 62% optimistic about the Union's future.
- **Participation and its lessons:** Conference on the Future of Europe (Apr-2021 to May-2022): 800 citizens in deliberative panels, 178 recommendations, 49 final proposals, 5 million platform visitors, 700,000 event participants; follow-up balance described as "mixed" by academic analysis (Carnegie Europe); European Court of Auditors, Special Report 14/2019 on the Commission's public consultations ("Have your say!"): unclear objectives and insufficient outreach; European communication doctrine: the risk of the "deflating souffle" — consultations without results drain trust (Libaert/Euractiv); the Commission itself embraced since 2008 (Wallstrom) the shift "from communication as selling to participation."
- **Institutional transparency:** Art. 15(3) TFEU (every citizen's right of access to the institutions' documents) and Regulation (EC) 1049/2001 of 30-May-2001 regarding public access to European Parliament, Council and Commission documents.

- **The book's own frameworks:** Chapter 8 (forensic Scanner, human-AI dual verification), Chapter 9B (ROI 164%, EUR 1,820 million/year), Chapter 13 (the five verbs; pure protector), Chapter 15 (Art. 48 TEU legal basis; endowment doctrine), Chapter 16 (own resources Art. 311 TFEU), Chapter 17 (constitutional PP in 2033; 55,000-68,000-job ecosystem), Chapter 18 (three-phase autonomy; radar with synthetic public summary; milestones and results, never procedures; the 2049 reference).

Chapter 20: Political and Legal Path of the Protective Power

Treaty reform is not the final obstacle: it is the historic opportunity to make the protection of the cognitive sphere the founding project of 21st-century Europe — through a procedure that already exists, was activated by Parliament in 2023, and that the Protective Power does not have to invent, only to walk

Chapter summary

This chapter answers the question the whole book leaves open: **how does one politically approve a 4th constitutional power in a Union of 27 States requiring unanimity?** The answer has three levels. First: **the legal path already exists and is already cleared.** The ordinary revision procedure of **Art. 48(2-5) TEU** — Convention, Intergovernmental Conference, unanimity, ratification by all according to their constitutional rules — is not a hypothesis of this book: the European Parliament **formally activated it on 22-Nov-2023** with its report on the amendment of the Treaties (2022/2051(INL), 200+ pages), the Council forwarded it to the European Council on 18-Dec-2023, and President von der Leyen publicly supported treaty reform in November 2024. The request for a Convention sits on the table waiting for the coalition that will convene it; the Protective Power Treaty can be the text it negotiates. Second: **the strategy is the founding coalition** — enhanced cooperation as the starting route for the willing (Chapter 14), with the door open to everyone, under the rule history has demonstrated: no State is forced to integrate; it is shown that staying out means falling behind. The precedent is verifiable and recent: **Denmark abolished its defence opt-out by referendum on 1-Jun-2022 (~67% in favour; 66.9% exact)**, three decades after rejecting it, because the geopolitical context changed — just as the EDSPA's evidence will do with reluctant States. Third: **the 2022-2026 geopolitical window is the best in a generation** — enlargement with Ukraine forces the Union to reform its institutions before widening them, European defence shows the Union integrates critical capabilities when the threat demands it, and the Draghi, Letta and Niinisto reports converge on the same conclusion: without structural reform, there is no strategic capacity. The chapter sets out the procedure step by step, the book's defensible chronology (signature 2031-2032, constitutional PP in 2033), the risks and their mitigations — and closes with Chapter 18's mirror: the rival has a strategy with a date (2049); this chapter is the proof that democracies know how to write theirs too.

20.1 The thesis: the path is already cleared

The most common analytical error about the Protective Power is to treat treaty reform as the project's final wall: unanimity of 27, ratification by 27, impossible referendums — therefore, never. This chapter demonstrates the opposite with a verifiable fact: **the revision procedure has been formally activated since 2023.** After the Conference on the Future of Europe (49 proposals, several of them requiring treaty revision), the European Parliament adopted on 22-Nov-2023 its report on proposals for the amendment of the Treaties — a 200+ page document of the AFCCO committee proposing profound institutional reforms — and exercised Art. 48(2) TEU by requesting the convening of a Convention. The Council forwarded it to the European Council in December 2023. Since then, the ball is in the court of the Heads of State and Government: the decision to convene the Convention is taken by a **simple majority** (14 of 27 governments) — not by unanimity.


This changes the nature of the PP's problem. There is no need to persuade 27 States to invent a new procedure: the 4th power's Treaty must **ride a Convention that was already requested, with the procedure already written into the Treaty in force and already activated by the institution that made it possible.** And the chapter's political thesis is deliberately ambitious: reform is not the cost of the Protective Power — it is its **historic opportunity.** Every step of the path (the EDSPA that demonstrates, the Convention that debates, the coalition that founds) does not just bring the PP closer: it forces the Union to decide, for the first time with a date and a text, what model of cognitive sphere it wants for the 21st century. The rival decided its own in 2012 and gave it a year: 2049. This chapter is the democratic equivalent of that decision.

20.2 The real procedure: Art. 48(2-5) TEU step by step

Art. 48 TEU distinguishes two routes. The **ordinary revision** (paragraphs 2-5) — the only one that can create a new constitutional power — has five steps:

1. **Proposal.** Initiated by a Member State, the European Parliament or the Commission before the Council, which forwards it to the European Council. (Current status: made by the EP on 22-Nov-2023; forwarded on 18-Dec-2023.)
2. **Convention.** The European Council decides **by simple majority** to convene a Convention — representatives of the national parliaments, of the Heads of State or Government, of the European Parliament and of the Commission — which examines the proposals and adopts a recommendation by consensus. (Precedent: the Laeken Convention of 2002-2004, which drafted what became the Lisbon Treaty.)
3. **IGC.** The Intergovernmental Conference of the governments' representatives adopts the amendments **by common accord**.
4. **Signature and entry into force.** The amendments enter into force **after being ratified by all the Member States in accordance with their respective constitutional rules** — real unanimity: no State is left out of ratification.
5. **Electoral reform (if applicable).** Paragraph 3 allows, for provisions of Part Three of the TFEU, a simplified procedure without Convention with unanimous European Council approval — not applicable to the PP, which requires the full ordinary route.

The book's defensible chronology — already fixed in Chapters 14 and 17 — fits into this procedure with room to spare historically: **signature 2031-2032, ratification completed and constitutional PP in 2033**. The precedent: from the Laeken Convention (Feb-2002) to the signing of the Constitutional Treaty (Oct-2004) took less than three years; and although the French and Dutch "no" votes of 2005 forced the renegotiation of Lisbon (2007, in force 2009) — the lesson of that crisis is not that reform is impossible, but that **the text matters**: the PP's Convention arrives with what Laeken lacked, an agency that already works and demonstrates (EDSPA 2027-2030) the problem the Treaty solves.



20.3 The strategy: the founding coalition and the Danish precedent

No unanimity strategy starts with unanimity. The PP's — already fixed in Chapter 14 and financed in Chapter 16 — has three pieces that this chapter turns into political doctrine:

First: no State is forced to integrate; it is shown that staying out means falling behind. Enhanced cooperation (Arts. 20 TEU and 326-334 TFEU) lets the willing States start the EDSPA without waiting for the last — the same mechanism with which 25 States advanced the financial transaction tax and the European divorce. The door stays open; what does not exist is the right to veto the others. And for the doubting State, the evidence of the EDSPA phase (2027-2030) will do what the 2022 context did for Denmark: **on 1 June 2022, the Danes abolished their defence opt-out by referendum with ~67% of the vote (66.9% exact)** — an exception their country had defended

for three decades (and reaffirmed in 2000 and 2015) disappeared within a year, because the Russian aggression changed the voters' calculus. No treaty forced Copenhagen: reality did it. Industrial manipulation will do the same to every State that sees its neighbours protected and its own elections exposed.

Second: one founding coalition, two acts together. The PP Treaty (Art. 48 TEU) and its endowment Decision (Art. 311 TFEU) are processed with the same geometry — unanimity and national ratification — and the same 27 potential signatories. The founding coalition that approves them together (Chapter 16) turns two impossible battles into a single campaign: **one single political act endows the power and finances it.**

Third: the paradox of 60% as the argument before governments. The decisive argument before a skeptical Head of State is not the ROI (164%, Chapter 9B) or the jobs (55,000-68,000, Chapter 17): it is that **their own electorate already lives the contradiction** — demanding protection from disinformation (69% in the EU) while distrusting their own government's AI (60%, OECD 2026) — and that no amount of reactive regulation resolves that contradiction, because it is not a legal problem: it is institutional. Only a permanent verification power, with human-AI dual verification and transparency by default, closes the gap the 60% perceives. The ruler who understands this is not accepting one more agency: he is defusing the most potent populist argument of the coming decade — "they have the AI and you have no say."

20.4 The 2022-2026 geopolitical window: the best in a generation

Three forces converge in the same direction, and none of them was requested by this book:

- **Enlargement forces reform.** Ukraine and Moldova are candidates; negotiations with Ukraine advance at technical level despite the Hungarian veto on formal opening (the Danish presidency secured enough support in Nov-2025 for working-level engagement). No serious analyst argues that the Union can absorb an enlargement of that scale without reforming its institutions and decision-making: the Convention requested in 2023 is also the Convention of enlargement — and the PP can ride in it.
- **Defence proved the Union integrates the critical.** Rearmament after 2022, the open discussion of a Defence Commissioner and a common defence clause, the reliance on Art. 42.7 TEU as a real guarantee for Kyiv: the Union of 2026 is already integrating the most jealously sovereign capacity of all, the military, because the threat demands it. If European defence is debatable today, the protection of the cognitive sphere — where the threat already operates against every Member State daily — is constitutionally obvious.
- **The Union's own reports ask for the same thing.** Draghi (competitiveness: fragmentation and inability to scale), Letta (single market: structural gaps), Niinisto (preparedness: hybrid and cognitive threats) converge on the same conclusion with different vocabularies: **without structural reform, the Union will have no strategic capacity.** The PP is the institutional translation of that conclusion to the domain where the threat is already inside the house.

And over all of this, Chapter 18's horizon: **the rival set its strategy with a date in 2012 — 2049.** That is not a prophecy of this book; it is a public commitment of the adversary with a calendar and verifiable execution plans (China Standards 2035, semiconductor and AI self-sufficiency as a "generational national project"). The democratic answer cannot be reactive regulation: it must be a permanent institution with a Treaty, an endowment and a date. This chapter's chronology — Convention by 2030, signature 2031-2032, constitutional PP 2033, full ~2037, maturity 2040+ — is not the theoretical maximum speed: it is **the minimum speed for the second option to exist in 2049.**

20.5 The legal basis as advantage: what the Treaties already say

The PP's basis is not one article: it is a **coherent reading of the primary law in force** that this chapter condenses (the full development is in Chapter 15):

Foundation	What it establishes	What it gives the PP
Art. 2 TEU — the Union's values	Democracy, rule of law, human rights as the foundation of all Union action	The cognitive sphere is the material condition of those values: without verifiable information, deliberative democracy is impossible
Art. 3 TEU — the Union's aims	The Union offers its citizens an area of freedom, security and justice	The cognitive sphere as an area: the citizen's security includes that of their mind against industrial manipulation
Art. 48(2-5) TEU — ordinary revision	Convention → IGC → ratification by all; activated by the EP in 2023	The exact procedure for creating the 4th power — already requested, already cleared
Art. 16 TFEU + Charter Arts. 7-8 — data protection	Personal data protection as an autonomous fundamental right (with CJEU case law since 2014: Google Spain)	The constitutional precedent that the law actively protects the citizen in their relationship with digital infrastructure
Art. 352 TFEU — flexibility clause	Union action where the Treaty does not provide the powers needed to attain its aims (unanimity)	The EDSPA's legal door before the Treaty: the means cannot be missing when the aim is constitutional
Arts. 20 TEU + 326-334 TFEU — enhanced cooperation	At least 9 States may advance using the Union's institutions	The founding coalition's starting route: no one vetoes the willing
Art. 311 TFEU — own resources	Endowment by unanimous Decision ratified by all	Financing as a category of the same Decision: one political act endows and finances (Chapter 16)
Art. 222 TFEU — solidarity clause	The Union and the States assist each other in the event of an attack (including hybrid)	The framework where foreign manipulation campaigns are already legally an attack on a Member State

Why this is an advantage and not an obstacle: the rival's constitution is a single party that needs no procedures. The Union's is a procedure that has been successfully executed at every institutional enlargement in its history — from Laeken to Lisbon, from the euro to the 2020 own resources. The PP does not ask to bypass the law: it asks to use it a second time in a generation for the aim that gives meaning to all the others. The legal basis is not the skeptics' argument — it is the book's argument.

20.6 Conclusion: from the request to the Convention

- 1. The path exists and is activated.** Art. 48(2-5) TEU: Convention by simple majority, IGC by common accord, ratification by all. The EP requested the Convention in Nov-2023; the request awaits the European Council — the PP Treaty can be its text.
- 2. The strategy is the founding coalition.** Enhanced cooperation to start, the door open to all, and the rule Denmark demonstrated in 2022: States are not forced to integrate — the context convinces them, and the EDSPA's evidence will be the context.
- 3. The argument before each government is the paradox of 60%.** Their electorate demands protection and distrusts state AI; only a permanent verification institution with dual verification and transparency by default resolves that contradiction — and defuses AI populism instead of feeding it.
- 4. The window is real and verified.** Enlargement with Ukraine forcing reform before widening, European defence integrating the critical, and Draghi/Letta/Niinisto converging on the need for structural reform: the 2030 Convention will have to exist for reasons that are not the PP's — the PP only has to be in it.
- 5. And the counterweight date is 2049.** The rival wrote its strategy with a centenary included; this chapter proves the democracies have the procedure, the precedent and the chronology (signature 2031-2032, PP 2033, full ~2037, maturity 2040+) to write theirs. Without prophecy — but without excuses.

The chapter's thesis in one sentence: treaty reform is not the Protective Power's final wall: it is the door the Union already opened in 2023 — and the historic opportunity for the protection of the cognitive sphere to become the founding project of 21st-century Europe, with the Art. 48 TEU procedure already activated, the founding coalition already defined and the mirror of 2049 already lit.

20.7 Sources and verified references

- **Primary law:** Art. 48 TEU (ordinary revision: paragraphs 2-5; Convention convened by simple majority of the European Council; IGC by common accord; ratification by all States per their constitutional rules; simplified procedure of paragraph 3 for Part Three of the TFEU, not applicable to the PP); Arts. 20 TEU and 326-334 TFEU (enhanced cooperation: minimum 9 States, use of the institutions); Art. 352 TFEU (flexibility clause); Art. 311 TFEU (own resources); Art. 222 TFEU (solidarity clause); Art. 42.7 TEU (mutual assistance clause); Arts. 2 and 3 TEU (values and aims); Art. 16 TFEU and Charter of Fundamental Rights Arts. 7, 8, 11 (privacy, data, information).
- **Real status of the revision (verified):** European Parliament, resolution and report 2022/2051(INL) "European Parliament's proposals for the amendment of the Treaties", adopted 22-Nov-2023 (AFCD committee, 200+ pages), activating Art. 48(2) TEU; the Council forwarded the request to the European Council on 18-Dec-2023; the European Council has not yet convened the Convention; von der Leyen publicly supported treaty reform in Nov-2024; academic analysis of the process: West European Politics (2025) on the EP's revision entrepreneurs and internal majorities (report adopted by 291 votes to 274, with EPP divisions; the earlier call for a Convention had been adopted on 9-Jun-2022 with 355:154:48).
- **Constitutional precedents:** Convention on the Future of Europe (Laeken, Feb-2002 to Jul-2003; final report with the draft Constitutional Treaty, signed Oct-2004); French (May-2005) and Dutch (Jun-2005) "no" votes in the ratification referendums; Treaty of Lisbon (2007, in force 1-Dec-2009); Irish referendums on Lisbon: rejected Jun-2008, approved Oct-2009 after guarantees.
- **The Danish precedent:** referendum of 1-Jun-2022: abolition of the Danish defence opt-out with ~67% of the vote (announced Mar-2022 after a cross-party agreement), three decades after the Edinburgh derogations (1992) and the referendums of 2000 (euro, rejected) and 2015 (conversion of the justice and home affairs opt-out, rejected); the context: the Russian aggression against Ukraine.
- **Enlargement and defence (2025-2026 context):** Ukraine and Moldova candidates (2022); the Hungarian veto on the formal opening of accession negotiations with Ukraine and the continued technical work under the Danish presidency (Nov-2025); Art. 42.7 TEU as a guarantee discussed in the security conversations for Ukraine; European debate on a Defence Commissioner and a common defence clause within the response to rearmament.
- **The Union's reports:** Draghi Report (2024) on competitiveness; Letta Report (2024) on the single market "Much More Than a Market"; Niinisto Report (2024) on crisis preparedness (including hybrid and disinformation threats); Conference on the Future of Europe, final report (May-2022): 49 proposals, including treaty revision.
- **The book's own frameworks:** Chapter 9B (ROI 164%), Chapter 14 (dual constitutional route: EDSPA via enhanced cooperation → Treaty), Chapter 15 (legal basis developed; Art. 352 TFEU and CJEU case law), Chapter 16 (one political act endows and finances: Art. 48 TEU + Art. 311 TFEU), Chapter 17 (chronology: signature 2031-2032, constitutional PP 2033; 55,000-68,000 ecosystem), Chapter 19 (paradox of 60% as the argument before governments), Chapter 18 (2049 mirror: Two Centenaries, China Standards 2035).

Chapter 21: The Future Is Not Predicted: It Is Built

The future is not predicted: it is built. And Europe already knows how to build: it exported peace where there had been a thousand years of war, protected the European citizen's privacy and ended up protecting the world's, and created an identity out of student exchanges. The Protective Power would be the fourth time Europe does what no one else can do — and the condition for artificial intelligence to be humanity's next leap instead of its replacement

Chapter summary

This chapter closes the book. Not with a prophecy — no dated European victory in 2039, no table of invented indicators; that road is closed — but with something harder and more honest: **a vision built on what already exists and on what Europe has already proved it can do**. First, the **three competing digital models** (market, control, democracy) are a fact of 2026, not of the future; this chapter compares them without verdicts and shows what each one *cannot deliver by design* — the market cannot monetize attention and protect it at the same time; control cannot separate protection from surveillance; only the European model can combine both, and only if it has a permanent executor: the institutional gap this book demonstrates. Second, **the world is already asking for what Europe can give**: the UN Global Principles for Information Integrity (Guterres, Jun 2024), the world's first legally binding treaty on AI and human rights from the Council of Europe (opened for signature Sep 2024, signed by the Commission on behalf of the EU) and the OECD AI Principles updated in 2024 — the global doctrine exists; what is missing is the institution that embodies it permanently and verifiably. Third, **the legacy is not measured in magic numbers but in what is won by design**: a durable institution with own resources, human judgment institutionalized as a counterweight to AI, a literate society, and Europe as the reference of the world's democracies. And fourth, the closing this book deserves: **the choice of 2049, the dream declared as a dream** — the Gutenberg precedent: the printing press also drowned man in other people's words, and out of that explosion came humanism and education; AI is the printing press of the twenty-first century, and Europe can invent its digital humanism — plus the call to action: this book is not an analysis, it is a call. The moment is now.

21.1 The thesis: the future is not predicted, it is built

This book does not end with a prophecy: "whoever controls the digital sphere controls the world. And in 2039, Europe will." It has spent twenty chapters demonstrating the opposite of that sentence — that **no one controls the digital sphere**, and that Europe's ambition is not to control it but **to protect those who inhabit it**. It has also shown that prophecies are the rivals' method: Xi Jinping's Two Centenaries and China Standards 2035 are precisely that — dated prophecies. Democracies do not prophesy: **they build**.

The thesis of this closing is that the future of the Protective Power is not a prediction but **a decision**. Every element already exists and has been verified in this book: the **regulatory acquis** (AI Act, DSA, DMA, European Democracy Shield, Sovereignty Package — Ch. 13); the **reform procedure** already activated by the European Parliament (Art. 48 TEU, Ch. 20); the **starting route** for the willing (enhanced cooperation, Ch. 14); **autonomous funding** (own resources, Art. 311 TFEU, Ch. 16); **technology in phases** with an acceleration clause (Ch. 18); the **economic evidence** (ROI 164%, +EUR 1,820 million a year, Ch. 9B) and the **communication strategy** that solves the 60% paradox (Ch. 19). None of this depends on guessing the future. It depends on thirty-plus years of European political decisions lining up — as they lined up, over decades, to build peace.

From victim to architect: that is the transformation this book has assembled piece by piece. The victim suffers the cognitive sphere; the architect designs it. Between the two there is exactly one institution: the Protective Power.

21.2 Three models in competition: what each can — and cannot — deliver

The digital world of 2026 has no single owner: it has **three competing models**, each with a different organizing principle. They are not hypotheses: they coexist, compete and export their logics today.

Model	Representative	Principle	Verified strengths	Verified weaknesses
Digital capitalism	United States	Private innovation + market	An entrepreneurial ecosystem without equal; leadership in foundational models and semiconductors; speed of deployment	Its engine is the attention economy: the citizen's time is the raw material that gets monetized; the regulation of its digital sphere is imposed by a third party (the European DSA/DMA); global information trust crisis (Edelman 2026: 7/10 citizens in echo chambers)
Digital authoritarianism	China	State control + scale	A decades-long published strategy (Two Centenaries, China Standards 2035); accelerated self-sufficiency (Kirin 7nm without EUV, Ascend 910C/950 with homegrown memory; subsidies of some US\$150,000M; self-sufficiency 16% to 28% in 18 months)	Censorship and mass surveillance by design; the same infrastructure that protects from the outside watches the inside; it exports control (Digital Silk Road), not protection
Digital democracy	Europe	Innovation + rights + law	The world's only complete digital regulatory acquis (AI Act, DSA, DMA, GDPR, European Democracy Shield); citizen trust in the EU at an 18-year high (52%); the model others copy (Brussels effect)	80% of digital infrastructure imported (EuroStack/Bertelsmann 2025); the consensus of 27 is slow; without a permanent executor, the law protects on paper

The table declares no winner — there is none yet, and this book does not prophesy. But it does let us see something the three columns hide:

What each model cannot deliver — by design, not by chance

The market cannot protect the sphere it monetizes. This is not a moral accusation: it is an incentive structure. Every improvement in the protection of the citizen's attention is, for the model that lives off it, a cost. That is why the only serious regulation of the US digital sphere was written by someone else: Europe.

Control cannot deliver freedom. The technical apparatus that detects external manipulation is the same one that watches the internal dissident; they share infrastructure, public law and personnel. By design, protecting entails surveilling. That is why it exports censorship, not protection.

European democracy can do both — but only with an executor. The regulatory acquis already defines protection *with* freedom. What it lacks is the permanent, independent institution that makes it true every day. And only the European model solves the 60% paradox (69% of Europeans demand protection against disinformation; 60% distrust government use of AI): human-AI double verification, transparency by default and reproducible evidence turn distrust into the power's design condition, not its risk. The other two models cannot solve that paradox: the market cannot answer those who distrust its incentive, and control cannot answer those who fear its surveillance.

The decisive difference: the weaknesses of the market and of control are structural — they cannot be fixed without ceasing to be what they are. Europe's weakness is one of execution, and execution can be fixed within the model itself: that is exactly the content of Chapters 14 to 20 of this book.

21.3 The world is already asking for what Europe can give

No need to imagine Europe exporting its technology to a world that would be waiting in 2039. The verified reality of 2024 is better: **the world has already asked — formally, in writing — for what Europe can give.** Three documents, three global institutions:

- **The United Nations: Global Principles for Information Integrity** (Guterres, 24 Jun 2024). The Secretary-General called to "respond to the harms caused by the spread of hate and lies online while vigorously defending human rights" and warned that threats to information integrity are "proliferating and expanding with unprecedented speed on digital platforms, supercharged by AI technologies." The Principles — born from his *Our Common Agenda* report (2021) — demand a more humane information ecosystem and served as a resource for the Summit of the Future (Sep 2024). The global doctrine exists: the executor is missing.
- **The Council of Europe: Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law** — the world's first legally binding international treaty on AI, opened for signature on 5 Sep 2024 in Vilnius. The European Commission signed it on behalf of the Union. A binding treaty committing its parties to ensure that AI systems do not undermine democratic integrity: the democracies' legal community has already written it. What is missing is the institution that verifies it every day.
- **The OECD: AI Principles** — the first intergovernmental AI standard (2019), updated at the May 2024 Ministerial Council Meeting for the generative AI era: five values-based principles and five recommendations for governments. The democracies' common language on rights-respecting AI.


And at home, the official framework where the mission fits without forcing anything: the **Digital Decade Policy Programme 2030** (Decision (EU) 2022/2481) sets the Union's targets of **20 million employed ICT specialists, 80% of adults aged 16-74 with basic digital skills and 75% of enterprises using cloud/AI by 2030**. The Protective Power does not arrive on a continent without a plan: it arrives on a continent with official digital-literacy and capability targets that its educational function (Ch. 12: literacy at 90% or higher by 2038) directly serves, and with a global doctrine already written that its executive function embodies.

The pattern has a verified precedent: the **Brussels effect**. The GDPR, born in 2018 to protect Europeans, became the model for legislation in Brazil, Japan, Singapore, South Africa, South Korea, Sri Lanka, Thailand and dozens more jurisdictions. Europe regulated for itself and ended up writing the world's privacy standard without firing a single regulatory missile. If the Protective Power matures, the same mechanics can operate with cognitive protection: the world's democracies do not need to copy a European agency — they need the *model*: a constitutional, independent power with double verification and public evidence. No one else has it. No one else is building it.

21.4 The legacy: what will have changed — recap and what is won by design

The road travelled in this book, in four parts:

- **Part I — The problem (Chapters 1-5):** the digital cognitive sphere as an essential public good captured: industrial manipulation, electoral interference, technological dependence. The diagnosis: Europe was a victim in its own sphere.
- **Part II — The patch (Chapters 6-10, with 9 and 9B):** EDSPA: a specialized agency that protects now — verifiable in its economics (ROI 164%, +EUR 1,820 million a year) — and that exists, above all, to demonstrate that cognitive protection works before asking 27 Member States to constitutionalize it.
- **Part III — The Definitive Solution (Chapters 11-20):** from agency to power: institutional architecture, pure-protector competences, transition without legal vacuum, legal basis, own resources, economic impact, technology in phases, communication that solves the 60% paradox, and the political procedure activated since 2023.
- **Part IV — The vision (this chapter):** not a prediction, but the declaration of what Europe can be if it decides to be.



And the legacy of the Protective Power? Not a table of magic numbers — no 18 million technology jobs, no 30% digital GDP, no 80% less disinformation; figures with no source and no verification have no place here. The real legacy is measured in what the architecture guarantees **by design**, and in the figures this book did verify:

- **An institution that lasts.** A constitutional power with own resources (Art. 311 TFEU) does not depend on the electoral cycle or the annual budget: it is the first guarantee that cognitive protection stops being a political fashion and becomes democracy's infrastructure — as central banks or the judiciary are today.
- **Human judgment as institutional employment.** The staffing trajectory verified in this book — from 1,500-2,000 (EDSPA) to more than 15,000 employees at maturity 2040+ (Chs. 15-16), with 55,000-68,000 jobs in the ecosystem (Ch. 17) — is not a jobs plan: it is AI's counterweight. The Protective Power would be Europe's largest public employer of non-automatable human judgment: in the era of automated thought, keeping thousands of people professionally dedicated to verifying with criteria is a civilizational policy, not a payroll.
- **A literate society.** The PP's educational function (Ch. 12) directly serves the Digital Decade's official targets (80% basic skills, 20 million ICT specialists by 2030) and extends them into cognitive literacy: the citizen who recognizes manipulation is the last line of defense — the only one no adversary can hack.
- **Europe as the democracies' reference.** If the model matures, the world's democracies that today have no answer of their own to industrial manipulation will have a model to copy — verifiable, constitutional, compatible with freedom. No prophecy of victory: as privacy went global without anyone "winning" it, protection can spread because it works and gets copied, not because Europe imposes it.
- **Evidence as a public good.** The Annual Cognitive Protection Report (Ch. 19), the public forensic evidence base and the "milestones and results, never procedures or techniques" doctrine leave behind something that does not exist today: a public, verifiable archive of how democracies get manipulated and how they resist — a global public good no other model will ever produce, because the market cannot and control will not.

21.5 The closing: two finished models, one declared dream, and a call

The choice of 2049

In 2049, the People's Republic of China celebrates its centenary with a public, declared objective: total technological power. Its model will be finished: state control of information, domestic surveillance exported as infrastructure, its own standards. The United States will have its own too: the cognitive sphere as a market, with its crown jewels in algorithms and chips. The world being born digital will have to choose between the two — **unless a third exists**: the one combining cutting-edge technology, freedom and verified protection. That third will not exist by spontaneous generation. This book is the blueprint to build it in time.

This is not a prophecy of victory — it is a window of relevance. The paradox of democracies is that they almost always arrive late but almost never arrive wrong: they arrived late to the industrial revolution and ended up leading it, late to the welfare state and defended it, late to digital privacy and today write its standard. They will arrive late to cognitive protection. May they not arrive late forever: that is what this book is about.

The dream: why Europe is allowed to dream

What follows is a declared aspiration — not a prediction.

When Gutenberg printed, Europe drowned in words. The printing press multiplied information — and with it manipulation: printed prophecies, war pamphlets, indulgences sold at industrial scale. For a century, the information explosion looked like the apocalypse. And out of that same explosion came Renaissance humanism, the Reformation, universal education and, at the end of the road, freedom of thought as a right. **Europe did not answer the printing press by banning it: it answered by educating itself.** The response to manipulation by the printed word was not less printing: it was more humanism.

Artificial intelligence is the printing press of the twenty-first century: it multiplies information and with it manipulation — deepfakes, bots, industrial narratives. The temptation to ban is the old temptation; Europe's historical response is the other one: **digital humanism**. And Europe invented humanism once, when information multiplied and man seemed to drown in other people's words; it can invent its digital version: the human being free from manipulation, using artificial intelligence as it used the printing press — to think more, and better.

And if the ambition seems grand, let the precedent sustain it. **Europe has already done it three times.** Schuman, in 1950, dreamed the impossible peace — "world peace cannot be safeguarded without creative efforts proportionate to the dangers which threaten it" — and out of that sentence came reconciliation where there had been a thousand years of war; the 2012 Nobel Peace Prize certified it: six decades of peace, reconciliation, democracy and

human rights. The GDPR, in 2018, was protecting Europeans — and ended up protecting the privacy of Brazil, Japan, South Africa, South Korea, and beyond, without anyone asking for it: the world copied it because it worked. Erasmus, in 1987, started with 3,200 students from 11 countries — and created something that did not exist: a European generation. **The Protective Power would be the fourth time Europe creates for itself what the world was waiting for.** No other continent has that pattern. No other is allowed this dream.

The call

This book began with a captured sphere and ends with a complete blueprint: the diagnosis, the patch that demonstrates, the legal and financial road already verified, and the vision. **This book is not just an analysis, it is a call to action.** The future of Europe — and of every democracy that wants to survive digitally without ceasing to be free — depends on our deciding today to protect the cognitive sphere. The Protective Power is the vehicle. The moment is now.

A generation protected from manipulation can use artificial intelligence for humanity's next leap: to think more, and better. That is Europe's aspiration. That is this book's aspiration.

21.6 Sources and verified references

- **Verified global framework:** United Nations, Global Principles for Information Integrity, launched by Secretary-General Antonio Guterres on 24 Jun 2024 (origin: *Our Common Agenda*, 2021; resource for the Summit of the Future, Sep 2024). Council of Europe, Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law: the world's first legally binding international treaty on AI, opened for signature on 5 Sep 2024 (Vilnius); signed by the European Commission on behalf of the EU. OECD, AI Principles: the first intergovernmental AI standard (2019), revised at the May 2024 Ministerial Council Meeting (five values-based principles, five recommendations).
- **Official European framework:** Digital Decade Policy Programme (Decision (EU) 2022/2481): 2030 targets of 20 million employed ICT specialists, 80% of adults aged 16-74 with basic digital skills, 75% of enterprises using cloud/AI; the Commission's *State of the Digital Decade* report.
- **Precedents of the European pattern:** Schuman Declaration (9 May 1950): "world peace cannot be safeguarded without creative efforts proportionate to the dangers which threaten it"; 2012 Nobel Peace Prize to the EU "for over six decades having contributed to the advancement of peace, reconciliation, democracy and human rights in Europe"; GDPR (2018) and the Brussels effect: model for legislation in Brazil, Japan, Singapore, South Africa, South Korea, Sri Lanka, Thailand and other jurisdictions; Erasmus (1987): 3,200 students from 11 countries in its first year; the "Erasmus generation" as an identity.
- **Rival context (verified, Ch. 18):** Xi Jinping's Two Centenaries (announced at the 18th Congress, 2012: a moderately prosperous society by the Communist Party's centenary, 2021; a fully modernized country by the People's Republic's centenary, 2049, with the intermediate milestone of "basic socialist modernization" in 2035), China Standards 2035, accelerated semiconductor self-sufficiency (Kirin 7nm without EUV, Ascend 910C/950, subsidies of some US\$150,000M, 16% to 28% in 18 months, declared goals: the Made in China 2025 plan's 70% (not reached) and 80% for 2030).
- **Three models 2026 (verified):** the attention economy as a business model; Edelman 2026 (34,000 respondents, 28 countries; information insularity 7/10); the European DSA/DMA regulating the platforms' digital sphere; the Digital Silk Road; EuroStack/Bertelsmann 2025 (80% of infrastructure imported); Eurobarometer Spring 2025 (trust in the EU 52%, an 18-year high).
- **The book's own frameworks:** Ch. 9B (ROI 164%, +EUR 1,820 million a year), Ch. 12 (literacy at 90% or higher by 2038; staffing trajectory), Chs. 15-16 (own resources, Art. 311 TFEU; 15,000+ employees at maturity 2040+), Ch. 17 (ecosystem of 55,000-68,000 jobs), Ch. 19 (the 60% paradox; Annual Cognitive Protection Report; milestones and results, never procedures or techniques), Ch. 18 (the 2049 mirror; the acceleration clause), Ch. 20 (Art. 48 TEU activated 22 Nov 2023).

Annex A: Legal Texts

Purpose of this annex

This annex collects the **real legal texts** on which the book's proposal rests, quoted verbatim with their official references. The rule is strict: **nothing here is an invented document** — every citation corresponds to the consolidated versions of the Treaties (OJ C 326, 26-10-2012), to regulations with official number and date, to CJEU case law with case numbers, or to verifiable institutional acts and packages. The annex is organized in three blocks: the basis for EDSPA (without treaty reform), the basis for the Protective Power (with reform), and the *acquis* both institutions would make executive.

A.1 Legal basis of EDSPA (2027-2030): no treaty reform required

Article 114 TFEU — Approximation of laws (internal market)

"The European Parliament and the Council shall, acting in accordance with the ordinary legislative procedure and after consulting the Economic and Social Committee, adopt the measures for the approximation of the provisions laid down by law, regulation or administrative action in Member States which have as their object the establishment and functioning of the internal market."

TFEU, Art. 114(1) (OJ C 326, 26-10-2012)

Application to EDSPA: the fragmentation of cognitive protection among Member States distorts the digital internal market; approximating provisions by regulation is the ordinary route through which the Union has created all its digital frameworks. It is the main basis of the EDSPA regulation, in line with Chapter 15's doctrine (hierarchy of bases: regulation first).

Article 16 TFEU — Protection of personal data

"Everyone has the right to the protection of personal data concerning them."

TFEU, Art. 16(1) (OJ C 326, 26-10-2012)

Application to EDSPA: the cognitive sphere consists, to a large extent, of personal data processed at industrial scale; Article 16 is the constitutional basis of the GDPR and the anchor of the whole body of protective legislation EDSPA audits. Complemented by Articles 7 and 8 of the Charter (section A.2).

Article 352 TFEU — Flexibility clause

"If action by the Union should prove necessary, within the framework of the policies defined in the Treaties, to attain one of the objectives set out in the Treaties, and the Treaties have not provided the necessary powers, the Council, acting unanimously on a proposal from the Commission and after obtaining the consent of the European Parliament, shall adopt the appropriate measures."

TFEU, Art. 352(1) (OJ C 326, 26-10-2012)

Application to EDSPA: a subsidiary basis: if the legislator considered Article 114 insufficient, Article 352 allows adopting the action needed for a Treaty objective. It requires unanimity, so its value here is doctrinal backing, not the main route.

Enhanced cooperation: Article 20 TEU and Articles 326-334 TFEU

"Member States which wish to establish enhanced cooperation between themselves within the framework of the Union's non-exclusive competences may make use of its institutions and exercise those competences by applying the relevant provisions of the Treaties. [...] Enhanced cooperation shall be authorised by the Council as a last resort, when it has been established that the objectives of such cooperation cannot be attained within a reasonable period by the Union as a whole [...] Enhanced cooperation shall be open at any time to all Member States, in accordance with Article 331 of the Treaty on the Functioning of the European Union."

TEU, Art. 20(1)-(2) (OJ C 326, 26-10-2012)

"Enhanced cooperation shall aim to further the objectives of the Union, protect its interests and reinforce its integration process. [...] Acts and decisions adopted in the framework of enhanced cooperation shall bind only the participating Member States."

TFEU, Art. 326(1)-(2) (OJ C 326, 26-10-2012)

Application to EDSPA (consistent with Chapters 11 and 14):

- **Minimum of 9 Member States** to establish it (Art. 20 TEU) — the book's founder core.
- **Last resort:** authorization requires establishing that the Union as a whole cannot attain the objective within a reasonable period — precisely Part I's diagnosis (the Union as a whole has not protected the cognitive sphere).
- **Permanent openness (Art. 331 TFEU):** any Member State may notify its intention to join; the Commission confirms participation within four months. The Schengen-style gradual union of Chapter 14 has its literal mechanics here.
- Acts bind only the participants: forcing no one, closing the door to no one.

A.2 Charter of Fundamental Rights of the EU

Articles 7, 8 and 11 — Privacy, data and freedom of information

"Everyone has the right to respect for his or her private and family life, home and communications." / "Everyone has the right to the protection of personal data concerning him or her." / "Everyone has the right to freedom of expression. This right shall include freedom to hold opinions and to receive and impart information and ideas without interference by public authority and regardless of frontiers."

Charter of Fundamental Rights of the European Union, Arts. 7, 8 and 11 (OJ C 326, 26-10-2012)

Application: the three articles together define the exact balance of the pure protector (Ch. 13): cognitive protection protects the *right to receive information without manipulation* (Art. 11) while respecting private life (Art. 7) and personal data (Art. 8). The PP does not filter, classify or block: freedom of expression remains untouched; what is pursued is manipulation, not opinion.

Article 47 — Effective judicial protection

"Everyone whose rights and freedoms as set forth in the law of the Union are violated has the right to an effective remedy before a tribunal, even where such an infringement results from acts of public authorities."

Charter, Art. 47 (OJ C 326, 26-10-2012)

Application (aligned with Chapter 12): every decision of EDSPA and the Protective Power is subject to appeal. The book's guarantee architecture creates no courts of its own: the guarantee operates at three verified levels — an internal administrative **Board of Appeal** (ACER model), the **Specialized Chamber for Digital Protection of the General Court**, with dual-qualified judges (Unified Patent Court model) funded from the judicial budget, and **appeal on a point of law to the CJEU**. Jurisdiction belongs to the judiciary; separation of powers is the first design guarantee.

A.3 Legal basis of the Protective Power (2033): treaty reform

Article 48 TEU — Revision procedures

"The Treaties may be amended in accordance with an ordinary revision procedure. [...] The Government of any Member State, the European Parliament or the Commission may submit to the Council proposals for the amendment of the Treaties. [...] The President of the European Council shall convene a Convention composed of representatives of the national Parliaments of the Member States, of the Heads of State or Government of the Member States and of the European Parliament and the Commission. [...] The amendments shall enter into force after being ratified by all the Member States in accordance with their respective constitutional requirements."

TEU, Art. 48(1), (2) and (4) (OJ C 326, 26-10-2012)

Application (Chapter 20): this is the road — not a hypothesis of the book: the European Parliament formally activated it on 22 Nov 2023 with its resolution and report 2022/2051(INL) on treaty reform, the Council transmitted the Convention request to the European Council on 18 Dec 2023, and President von der Leyen publicly supported the reform in November 2024. The book's timeline — Convention around 2030, signature 2031-2032, constitutional PP in 2033 — fits comfortably within the historical precedent (Laeken Feb 2002 to signature Oct 2004; Lisbon: 24 months from signature to entry into force).

Institutional integration: Articles 13-19 TEU

The Protective Power would be integrated into the institutional framework regulated by **Articles 13-19 TEU**: Art. 13(1) provides that "the Union shall have an institutional framework which shall aim to promote its values, advance its objectives, serve its interests and those of its citizens", and Art. 13(2) requires the Union to act within the limits of the competences conferred by the Treaties. A fourth power is not foreign to this framework: it is a new rank within it, established by Art. 48 — with Chapter 15's doctrine: the power that checks power does not multiply the state, it balances it.

Own resources: Article 311 TFEU and Decision 2020/2053

"The Union shall provide itself with the means necessary to attain its objectives and carry through its policies. [...] The Council, acting in accordance with a special legislative procedure, shall unanimously, after consulting the European Parliament and with the approval of the Member States in accordance with their respective constitutional requirements, establish new categories of own resources or abolish existing categories."

TFEU, Art. 311(1) and (3) (OJ C 326, 26-10-2012)

Application (Chapter 16): the PP's financial autonomy has an exact textual basis: an **Own Resources Decision** — like the current Decision (EU, Euratom) 2020/2053 — introduces the PP's endowment category with unanimity, consultation of the EP and ratification per constitutional requirements: **the same procedure as its Treaty**, approved by the same founder coalition. The plastic own resource precedent (Regulation (EU) 2021/770 developed the category created by Decision 2020/2053) shows the complete mechanics.

A.4 The regulatory acquis EDSPA and the PP would make executive

Act	Official reference	Role of EDSPA / PP (pure protector)
Digital Services Act (DSA)	Regulation (EU) 2022/2065; annual independent audits under Art. 37 (first round due Aug 2024, Delegated Regulation (EU) 2024/436); researcher access under Art. 40 (delegated act, 2 Jul 2025); penalties up to 6% of annual worldwide turnover (Commission/national DSA authorities)	Audits manipulation on platforms, generates forensic evidence and refers case files to the Commission and national authorities — sanctions are theirs, never the PP's (Chs. 12-13)
Artificial Intelligence Act	Regulation (EU) 2024/1689, of 13 Jun 2024; full application since 2 Aug 2026; penalties up to 7%	Algorithmic auditing and evidence on manipulative AI systems; referral to the AI Act supervisors
GDPR	Regulation (EU) 2016/679, of 27 Apr 2016	The cognitive sphere consists of personal data: its guarantees are the floor, not the ceiling
Data Act	Regulation (EU) 2023/2854; applicable since 12 Sep 2025	Data access and cloud switching: practical conditions for auditing
Digital Decade Programme	Decision (EU) 2022/2481 (2030 targets: 20M ICT specialists, 80% basic skills, 75% enterprises cloud/AI)	The PP's educational function serves the Union's official literacy targets
European Democracy Shield	Commission Communication, 12 Nov 2025	The Union's electoral protection framework: EDSPA provides the specialized forensic arm
European Technological Sovereignty Package	European Commission, 3 Jun 2026 (semiconductors, AI, cloud and open source)	The strategic framework of technological independence into which Chapter 18's three phases fit
Transition precedent	Regulation (EU) 2016/794 (Europol): the successor agency takes over staff and functions by organisational act	Model of the successor organisational act for the EDSPA-to-PP integration 2033-2035 without a legal vacuum (Ch. 14)

A.5 CJEU case law: protection as settled doctrine

- **Digital Rights Ireland (C-293/12 and C-294/12, 8 Apr 2014)**: the CJEU annulled the Data Retention Directive because interference with private life "must be limited to what is strictly necessary" — the legal foundation of the principle that *more surveillance is not more security*: the exact doctrine separating the protector from the watcher.
- **Google Spain (C-131/12, 13 May 2014)**: the right to be forgotten against the search engine: individuals hold effective rights against digital operators, the tradition the Citizen Checker (Ch. 13) extends to the verification of indications.
- **Schrems II (C-311/18, 16 Jul 2020)**: invalidation of the Privacy Shield over data transfers to third countries: the sovereignty of personal data is case law, not a slogan — and the deep reason European forensic evidence must be processed on European infrastructure (Ch. 18).
- **ACER (C-46/21 P, 2023)**: the validated model of the administrative Board of Appeal inside an agency — the first line of guarantee in Chapter 12.

A.6 The global framework: the world has already written the doctrine

- **Council of Europe Framework Convention on AI** (opened for signature 5 Sep 2024, Vilnius): the first legally binding international treaty on AI and human rights, democracy and the rule of law — signed by the Commission on behalf of the EU.
- **UN Global Principles for Information Integrity** (Guterres, 24 Jun 2024): the global doctrine against AI-supercharged manipulation.
- **OECD AI Principles** (2019, updated May 2024): the first intergovernmental standard for trustworthy AI.

The global doctrine exists and was written by third parties; what does not exist is the permanent institution to embody it. That is the gap — and the opportunity — this book demonstrates (Ch. 21).

A.7 Summary of the legal basis

Institution	Legal basis	Procedure	Key guarantee
EDSPA (2027-2030)	Arts. 114, 16, 352 TFEU; enhanced cooperation Art. 20 TEU + Arts. 326-334 TFEU	Ordinary regulation; enhanced cooperation with a 9-State core as last resort	No treaty reform; open at any time to the rest (Art. 331 TFEU)
Protective Power (2033)	Art. 48 TEU (ordinary revision, activated by the EP 22 Nov 2023); Arts. 13-19 TEU; Art. 311 TFEU (own resources)	Convention -> IGC -> unanimity -> ratification by the 27 per their constitutional requirements	Constitutional permanence + financial autonomy with the same procedure as the Treaty

A.8 Official references

- **Consolidated Treaties:** TEU, TFEU and the Charter of Fundamental Rights of the EU, consolidated versions, OJ C 326, 26-10-2012 (the quoted texts correspond to this consolidation).
- **Real state of the revision:** EP, resolution and report 2022/2051(INL), 22 Nov 2023 (activation of Art. 48(2) TEU); transmittal by the Council to the European Council, 18 Dec 2023; von der Leyen's public support, Nov 2024.
- **Legislation cited:** Regulations (EU) 2022/2065 (DSA), 2024/1689 (AI Act), 2016/679 (GDPR), 2023/2854 (Data Act); Decisions (EU) 2022/2481 (Digital Decade) and (EU, Euratom) 2020/2053 (own resources); Regulation (EU) 2021/770 (plastic own resource); Regulation (EU) 2016/794 (Europol).
- **Case law:** CJEU, judgments C-293/12 and C-294/12 (Digital Rights Ireland, 2014), C-131/12 (Google Spain, 2014), C-311/18 (Schrems II, 2020), C-46/21 P (ACER, 2023).
- **Global framework:** Council of Europe Framework Convention on AI (5 Sep 2024); UN Global Principles for Information Integrity (24 Jun 2024); OECD AI Principles (2019/2024).
- **Acts of 2025-2026:** European Democracy Shield (Communication, 12 Nov 2025); European Technological Sovereignty Package (3 Jun 2026).

Annex B: Operational Capabilities of the Pure Protector

What the institution will be able to do — and what this annex deliberately does not say: the book's rule ("milestones and results, never procedures or techniques") applies first to the book itself

Purpose of the annex

This annex deliberately does not describe a full system: no layered architectures, no named models with parameters, no commercial brands, data volumes or "neutralization" mechanisms. It does exactly the opposite — for a doctrinal reason, not casual discretion: **the book requires the Protective Power to publish milestones and results, never procedures or techniques** (Chs. 18-19), because the cognitive sphere is an adversarial domain and every methodological detail published is capability handed to the rival. The rule the institution must obey, this book already obeys. What follows is therefore an annex of **capabilities**: what the institution will be able to do, with what guarantees and against what threats — not how it will do it. The capabilities are anchored in the real technical frameworks already existing (NIST, ISO/IEC, ENISA, EEAS, DSA/AI Act) which the institution would make operational.

B.1 The doctrine of the annex: what this document does not say

The criterion governing every section is that of Chapter 18's Strategic Acceleration Clause and annual radar: **the full picture internally, a synthetic summary externally**. Applied to a technical annex, it means it publishes:

- **Capabilities** — what the institution can do (detect, audit, publish, refer, educate).
- **Guarantees** — what design limits make it trustworthy (double verification, transparency by default, privacy by design).
- **Taxonomies** — how threats are classified, because a public classification educates the citizen and gives nothing away.

And it does not publish: models with parameters, technology stacks, tools, data volumes, step-by-step audit methodologies, internal operating thresholds or automated response deadlines. **None of it is published: operational detail is never published, whatever its provenance.**

B.2 Core capabilities

Capability	What it does	What it produces (milestones and results)	What it never does
Cognitive forensic scanner (Chs. 8 and 13)	Detects manipulation patterns in the public digital sphere and generates evidence	Auditable, reproducible forensic evidence with human-AI double verification	Does not filter, classify truth or block content
Algorithmic auditing (Ch. 13)	Audits systems and platforms within the acquis (DSA, AI Act)	Public reports, reproducible evidence, referral case files to the Commission and national authorities	Does not sanction — itself or by delegation; sanctioning belongs to the Commission/authorities (6% DSA, 7% AI Act)
Cognitive education (Ch. 12)	Citizen literacy serving the Digital Decade targets (80% basic skills, 20M ICT specialists by 2030)	Programmes, materials and public literacy metrics (book target: 90%+ by 2038)	Does not induce or persuade: it teaches to recognize manipulation of any sign
Threat intelligence (Ch. 18)	Annual "State of the Cognitive Race" radar: full picture internally, synthetic public summary	Publishes trends, milestones and results; the evidence base is a global public good (Ch. 21)	Does not publish detection techniques, sources or methods
Citizen Checker (Ch. 13)	Non-binding verification of indications submitted by citizens, with pseudonymization	A verified answer to the citizen; aggregate public statistics	No reversible link to identity; no service before its 2033-2035 evaluation

B.3 Threat classification (public taxonomy)

Threat levels and the pure protector's response

The severity taxonomy is public educational material; the response column is Chapter 13's five functions â€” detect, audit, publish, refer, educate. **"Immediate neutralization" does not exist here: the PP neutralizes, removes and suspends nothing.**

Level	Description	Example	Pure protector's response
Critical	Existential threat to a democratic process	Mass coordinated electoral interference	Priority detection, immediate publication of evidence, urgent referral to the competent authorities (electoral/DSA Commission, States)
High	Serious threat with significant impact	Coordinated FIMI campaign against a Member State	Forensic audit, public report and referral case file
Medium	Moderate threat with limited impact	Localized disinformation	Detection, evidence in the public archive, targeted education
Low	Minor threat with minimal impact	Isolated misleading content	Monitoring and feeding of the evidence archive

Threat types

Type	Examples	Real taxonomy reference
Disinformation and synthetic deception	Fake news, deepfakes, manipulated images	EEAS FIMI classification; ENISA threat reports (2023: AI and LLMs as vectors of disinformation and social engineering)
Algorithmic manipulation	Bubbles, micro-targeting, bot amplification	Transparency and risk obligations of the DSA (Regulation (EU) 2022/2065); AI Act (Regulation (EU) 2024/1689) on manipulative systems
State FIMI	Foreign manipulation infrastructure (Russia and China mapped by the EEAS)	3rd EEAS FIMI Report (2024): digital infrastructure deployed against the European information space
Cyber threat to the sphere	Phishing, ransomware, DDoS against media and electoral processes	ENISA Threat Landscape (2,580 incidents Jul 2022-Jun 2023; public administration the hardest-hit sector)

Prioritization matrix

Severity \ Likelihood	Low	Medium	High
Low	Monitoring	Monitoring	Public evidence
Medium	Monitoring	Public evidence	Report and education
High	Public evidence	Report and referral	Immediate publication + urgent referral

B.4 Audit capabilities — results, not procedures

The annex states only the result, never the internal audit process (black box, white box, tools, thresholds): every PP audit produces **three public artifacts** — a **report** (conclusions and evidence), **reproducible evidence** (verifiable by third parties, with human-AI double verification) and, where there is an indication of infringement, a **referral case file** to whoever holds sanctioning power. The procedure is internal heritage and professional standard; the result is public by design.

The requirements the institution does declare — because they are a guarantee for the citizen, not capability for the rival — are Chapter 18's: all its AI will be **open, auditable and explainable**, with non-automatable human double verification; and all procurement will run through **open calls**, with no dependence on a single vendor (the Gaia-X lesson of Ch. 18).

B.5 Design guarantees (what is auditable is always the human)

- **Human-AI double verification (Ch. 8):** no relevant institutional action begins or ends in the machine; human judgment signs. This is the design answer to the 60% paradox (Ch. 19): citizen distrust is the power's condition, not its risk.
- **Transparency by default:** Art. 15(3) TFEU and Regulation 1049/2001: documents are public unless an exception is motivated; the Annual Cognitive Protection Report before Parliament is the structural accountability (Ch. 19).
- **Privacy by design:** GDPR Art. 25 and data minimization; surveillance is legally impossible by architecture (Digital Rights Ireland: interference "must be limited to what is strictly necessary" — Annex A.5).
- **Anti-weapon symmetry (Ch. 13):** the same tools that protect the citizen from manipulation cannot be used to manipulate them; the limit is constitutional, not programmatic.
- **Detect/sanction separation (Chs. 9, 12-13):** the PP detects and publishes; the Commission and authorities sanction. No fine percentage funds the PP: fines go to the EU's coffers (Ch. 16).

B.6 Public interfaces

Interface	Public	Content
Annual Cognitive Protection Report (24 languages, before the EP — Ch. 19)	The whole Union	Milestones and results, never procedures or techniques
Evidence portal	Researchers, media, citizens	Reproducible forensic evidence of closed cases
Verified alerts	Citizens	Notices of active campaigns with evidence, without alarm or censorship
Open data and education	Education systems, States	Taxonomies, literacy materials, aggregate metrics

The access precedent already exists and is verifiable: the July 2025 delegated act under **DSA Art. 40** allows vetted researchers to access non-public platform data — the institution would connect that scattered evidence into a systematic public archive.

B.7 Reference technical frameworks (real and verifiable)

- **NIST AI Risk Management Framework (AI RMF 1.0, NIST AI 100-1, Jan 2023)** and its **Generative AI Profile (NIST AI 600-1, Jul 2024)**: the world's most adopted AI risk-management standard — the methodological backbone the institution would operationalize without publishing.
- **ISO/IEC 42001:2023** (Dec 2023): the first international AI management system (AIMS) standard — the governance framework under which its own AI would be certified.
- **ENISA**: the annual Threat Landscape (2,580 incidents Jul 2022-Jun 2023) and the AI threat landscape — the European cyber-threat baseline feeding the B.3 taxonomy.
- **EEAS FIMI reports** (1st-3rd, the third in 2024): the European reference classification of foreign information manipulation.
- **DSA Arts. 37** (annual independent audits at the platforms' expense, first round Aug 2024) **and 40** (Jul 2025 delegated act on vetted-researcher access) and **AI Act Art. 6** and annexes (assessment and supervision of high-risk systems): the acquis the PP's audit would make systematic.

Annex C: Case Studies

Purpose and method

This annex examines real episodes of threat to the cognitive sphere and the institutional responses that existed, to extract the design lessons justifying EDSPA and the Protective Power. The annex's rule: **verified figures only** — no impact numbers without a source (no "~500M impressions", no "10,000+ deepfakes", no "95% effectiveness"); the annex states verifiable qualitative facts or cites the datum with its real origin. And the "connections with the Protective Power" are reformulated to the pure-protector doctrine: **the PP does not sanction or neutralize — it generates public forensic evidence, refers case files and educates**. Every counterfactual of the "it would have sanctioned the actors" type contradicts the book's own architecture.

C.1 Threats: six real episodes

1. Disinformation in the European elections (2019)

Context: the European Parliament elections (23-26 May 2019) took place amid coordinated anti-EU and polarization disinformation campaigns, with fake accounts and amplification of divisive narratives, documented by the EEAS and the platforms.

Real response: the EEAS East StratCom Task Force (created 2015) maintains EUvsDisinfo, the public database of debunked disinformation cases since 2015; platforms removed account networks and published electoral transparency reports; fact-checkers covered the campaign.

Lessons for the PP (pure protector): the existing response was *observational* (debunking) and *platform by platform* (account removal); nobody generated systematic forensic evidence or a unified European public archive. The PP would have provided exactly that: forensic detection, reproducible evidence and coordinated referral — without replacing anyone.

2. Synthetic deception in the US elections (2020)

Context: the 2020 presidential race featured synthetic video and image manipulation and electoral-fraud narratives that eroded trust in the result.

Real response: CISA (the federal cybersecurity agency for electoral infrastructure) and the platforms maintained labelling and removal systems; fact-checkers worked through the whole cycle.

Lessons for the PP: detection always arrived *after* virality, and action depended on the platforms themselves — the same dependency diagnosis of Chapter 4. A European forensic archive of synthetic deception, with verifiable evidence and human double verification, would have made auditable what each platform claimed alone.

3. Algorithmic manipulation: Cambridge Analytica and Brexit (2016-2018)

Context: Cambridge Analytica exploited the data of up to **87 million Facebook profiles**, obtained without informed consent through an app, for micro-targeted political advertising in campaigns across several countries, around the UK referendum (23 Jun 2016).

Verified consequences: the UK ICO fined Facebook **GBP 500,000** (Oct 2018) for the serious data protection infringement, and the US FTC **settled for US\$5,000M** (Jul 2019), one of the largest privacy fines in history.

Lessons for the PP: the case showed that manipulative micro-targeting works at industrial scale *before* any regulator knows — and that the evidence arrived through journalism and internal whistleblowing, not institutional audit. The PP's algorithmic auditing (Ch. 13) exists so the next Cambridge Analytica is not discovered by chance: early detection, public evidence, referral to the data protection and competition authorities.

4. Disinformation in the COVID-19 pandemic (2020-2021)

Context: disinformation campaigns about the virus's origin and the vaccines, amplified by state and community actors, with documented vaccine-hesitancy and polarization effects.

Real response: the WHO coined the "infodemic" and coordinated verified-information campaigns; the EU deployed EDMO (2020) and its fact-checking network; Facebook reported removing **more than 12 million pieces** of COVID-19 content between March and October 2020 (the platform's own figure) and more than 20 million through mid-2021.

Lessons for the PP: the response was global in diagnosis and fragmented in action: each platform decided what to remove, with no public verifiable evidence of the criterion. The PP's educational function — permanent cognitive literacy, not emergency campaigns — is the structural answer to the infodemic.

5. Interference in the French elections (2022)

Context: for the April 2022 presidential election, France deployed a dedicated service for the first time: **Viginum** (created 2021), the vigilance and protection service against foreign digital interference.

Real response: Viginum detected and **published information-manipulation reports** (including coordinated networks of foreign origin) during the electoral cycle, coordinating with ANSSI, the CNIL and the platforms; publishing the evidence was an essential part of the response.

Lessons for the PP: France tested the book's thesis in miniature: a small, dedicated body that *detects and publishes* changes the quality of the response without needing sanctioning powers of its own. The PP is that idea elevated to European scale and constitutional rank — with permanence and own resources, and without the fragility of a service dependent on the political cycle.

6. Disinformation in the Brazilian elections (2022-2024)

Context: the 2022 presidential race lived through mass disinformation, attacks on electoral systems and an assault on the institutions on 8 Jan 2023 after months of result denial.

Real response: the TSE (Superior Electoral Court) ran an aggressive rapid-takedown model under judicial order and agreed protocols with the platforms; for the 2024 municipal elections it banned electoral deepfakes (with mandatory labelling of AI content) and signed memoranda with Meta, TikTok and other platforms; investigations documented covert RT payments to influencers.

Lessons for the PP: the Brazilian model demonstrates two things: that fast action is possible — but concentrated in a national electoral judicial power, with the transparency problems this generates (criticized by Brazilian press itself), and with no public verifiable-evidence architecture. The PP would add what is missing: reproducible forensic evidence and transparency by default, instead of opaque takedowns.

C.2 Institutional response models: what exists and what it teaches

Model	Created	Verified datum	Lesson for the PP
ANSSI (France)	2009	National cybersecurity agency; incident response and critical-infrastructure protection	Technical specialization works; its limit is the domain (cybersecurity, not the cognitive sphere)
NCSC (United Kingdom)	2016	National technical cybersecurity centre, a model of public-private coordination and operational transparency	A technical (non-political) agency publishing openly earns the sector's trust
ENISA (EU)	2004	Verified budget ~EUR 26.9M (Ch. 6); standards and certification for the 27	The reference European agency proves the viability of EU coordination — and its minimal budget shows the ceiling of what cannot be done cheaply (Ch. 11: institutional fragility)
EDMO (EU)	2020	Observatory and fact-checking network; its hubs' funding ends in 2026 with no guaranteed successor (Ch. 11)	Monitoring without permanence or executive capacity is a patch that peels off: exactly the gap EDSPA fills
Viginum (France)	2021	Service dedicated to foreign digital interference; detects and publishes reports	The direct precedent of the pure protector in miniature: public evidence without sanctioning power — national, small and temporary; the PP makes it European, permanent and endowed

C.3 Verified threat indicators (2019-2026)

Complete time series (campaigns per year, deepfakes per year, manipulation cases per year) with sourceless numbers have no place here. The annex uses the verified indicators the book already uses (Ch. 19 and Part I):

Indicator	Verified datum	Source
Citizen concern about disinformation	69% consider it a problem; 36% exposed "often" (28% in 2022); 85% in Spain	Eurobarometer disinformation 2025 (36%, exposure); Eurobarometer Feb-2026 (69%, concern)
Trust in news	37%, the lowest since 2015; news avoidance 42% (29% in 2017); social networks overtake TV as a source	Reuters Institute Digital News Report 2026
Information insularity	7/10 citizens live in echo chambers	Edelman Trust Barometer 2026 (34,000 respondents, 28 countries)
Foreign FIMI infrastructure	Map of the digital infrastructure deployed mainly by Russia and also China against the European information space	3rd EEAS FIMI Report (2024)
Distrust of government AI	60% distrust governmental use of AI — against the 69% demanding protection	OECD Trust Survey, Jun 2026 (the 60% paradox, Ch. 19)

Verifiable qualitative trends (without invented numeric series): the migration of information flows toward more opaque, less moderated platforms (TikTok, Telegram, X); the acceleration of synthetic-content production by generative AI (Stanford AI Index 2026); and the convergence of all institutional models toward publishing evidence — the field where the PP arrives first.

C.4 The pattern the six cases draw

In every case the same sequence repeats: late detection (after the damage went viral) → non-forensic evidence (platform claims without independent verification) → fragmented response (each actor with its own criterion, each country with its own law) → no permanent institution responsible for learning from the case. **The Protective Power is the institutional correction of that sequence:** permanent forensic detection, public reproducible evidence, coordinated referral to the sanctioners, and institutional memory — a European archive of how democracies get manipulated and how they resist (Ch. 21: a global public good no one else will produce).

C.5 Sources and verified references

- **EEAS:** East StratCom Task Force (2015) and EUvsDisinfo (public database of debunked cases); 3rd FIMI Report (2024).
- **Cambridge Analytica:** ICO, GBP 500,000 fine to Facebook (Oct 2018); FTC, US\$5,000M settlement (Jul 2019); data of up to 87 million profiles documented in both actions.
- **COVID-19:** WHO ("infodemic"); EDMO (2020); takedown figures reported by Facebook/Meta (more than 12 million pieces Mar-Oct 2020; more than 20 million through mid-2021).
- **France:** Viginum (created 2021), public information-manipulation reports in the 2022 cycle; ANSSI.
- **Brazil:** TSE, protocols with platforms (2022), the electoral-deepfake ban and memoranda with Meta/TikTok and others (2024); the investigation into RT payments to influencers.
- **Indicators:** Eurobarometer Feb 2026; Reuters Institute Digital News Report 2026; Edelman Trust Barometer 2026; OECD Trust Survey Jun 2026; Stanford AI Index 2026.
- **European Parliament:** the work of the INGE special committee on foreign interference (2020-2022).

Annex D: Bibliography

D.1 The criterion behind this bibliography

This bibliography applies to its source apparatus the same rule it applies to the rest of the book: **no source is cited for its apparent prestige, only for its verification**. Every document in this list has been checked against its official publication (EUR-Lex, the issuing institution's portal, or the publisher of record). This bibliography serves three functions:

1. **Backing**: every factual claim in the book has its source here.
2. **Traceability**: the reader can follow any figure back to its document of origin.

D.2 EU treaties and constitutional documents

Document	Articles cited	Reference
Treaty on European Union (TEU)	Arts. 2, 3, 13-19, 20, 42.7, 48	Consolidated version, OJ C 326, 26/10/2012
Treaty on the Functioning of the EU (TFEU)	Arts. 16, 114, 222, 311, 326-334, 352	Consolidated version, OJ C 326, 26/10/2012
Charter of Fundamental Rights of the EU	Arts. 2, 7, 8, 11, 47	Consolidated version, OJ C 326, 26/10/2012
European Parliament Resolution 2022/2051(INL)	Treaty reform proposals; activation of Art. 48(2) TEU	Adopted on 22 November 2023
Decision (EU, Euratom) 2020/2053	Own resources system; basis of Art. 311 TFEU	OJ L 424, 15/12/2020
Regulation (EU, Euratom) 2021/770	Plastic own resource: calculation of the own resource based on non-recycled plastic packaging waste; implements the category created by Decision 2020/2053 (Art. 311 TFEU)	OJ L 165, 11/5/2021

Rigor note: Regulation (EU, Euratom) 2021/770 implements the plastic own resource created by Decision 2020/2053: the book's precedent of an own-resource category introduced by Decision (Art. 311 TFEU) and made operational by a regulation (Annex A.3). It is not a legal basis for the EDSPA. Always check the actual scope of a rule before extending it by analogy.

D.3 Secondary law: the European digital acquis

Instrument	Subject	Reference
Digital Services Act (DSA)	Platform accountability; Art. 40 (delegated act, July 2025)	Regulation (EU) 2022/2065
Artificial Intelligence Act (AI Act)	Harmonised rules on AI	Regulation (EU) 2024/1689
General Data Protection Regulation (GDPR)	Data protection; Art. 25 (privacy by design)	Regulation (EU) 2016/679
Data Act	Data access and sharing	Regulation (EU) 2023/2854
Decision establishing the 'Digital Decade 2030' policy programme	Digital targets 2030 (20M ICT specialists, 80% basic skills, 75% firms cloud/AI)	Decision (EU) 2022/2481
Europol Regulation	Agency precedent for the transition	Regulation (EU) 2016/794
European Democracy Shield	Commission Communication	12 November 2025
Technological Sovereignty Package	Commission initiative	3 June 2026

D.4 Case law of the Court of Justice of the EU

Judgment	Year	Contribution to the book
Digital Rights Ireland (joined Cases C-293/12 and C-294/12)	2014	Struck down the Data Retention Directive: privacy is the rule, interference the strictly limited exception
Google Spain (C-131/12)	2014	Created the right to erasure: the data subject prevails over the search engine
Schrems II (C-311/18)	2020	European data protection travels with the data: digital sovereignty is enforceable extraterritorially
ACER (C-46/21 P)	2023	Model of an agency administrative Board of Appeal with subsequent judicial review

D.5 Strategic reports 2024-2026

Report	Author/Year	Contribution to the book
The future of European competitiveness (Draghi report)	European Commission, 2024	Diagnosis of the productivity gap and technological dependence
Much More Than a Market (Letta report)	2024	The single market as a pending lever; the single market for data
Niinistö report on EU security and resilience	2024	Understand, anticipate, prepare: security as a permanent function
Third report on FIMI threats	European External Action Service (EEAS), 2024	Taxonomy of foreign information manipulation and interference
ENISA Threat Landscape	EU Agency for Cybersecurity (ENISA)	2,580 incidents reported Jul 2022-Jun 2023
European Democracy Shield and Technological Sovereignty Package	European Commission, 2025-2026	Protecting the information sphere is already official EU policy

D.6 Data, surveys and polls

Source	Figure cited	Detail
Eurobarometer on disinformation (2025 edition)	36% exposed to it «often» or «very often» in the last 7 days (28% in 2022)	Published by the European Commission, fieldwork 2025
Eurobarometer on disinformation (published Feb 2026)	69% consider disinformation a problem in their country; 85% in Spain	Published by the European Commission, fieldwork autumn 2025
Reuters Digital News Report	Trust in news and digital consumption	2026 edition, Reuters Institute, University of Oxford
OECD Survey on Drivers of Trust in Public Institutions - 2026 Results	~60% medium-low trust in governmental use of AI; basis of the 60% paradox	OECD Publishing, 2026
Edelman Trust Barometer	Institutional trust trends	2026 edition

Rigor note: the 60% paradox (Chapter 19) is built with data from the February 2026 Eurobarometer and from the OECD, attributed to their source and year. If a later edition revises the figures, the argument does not change: what matters is the gap between the demand for protection and distrust of the protector, which both series capture separately.

D.7 Technical standards

Standard	Organization	Detail
AI Risk Management Framework 1.0 (AI RMF 1.0)	NIST (US)	January 2023
Generative AI Profile (NIST AI 600-1)	NIST	July 2024
ISO/IEC 42001:2023	ISO/IEC	First certifiable AI management system, December 2023
OECD AI Principles	OECD	2019, updated May 2024

D.8 International framework for information protection

Instrument	Date	Contribution to the book
United Nations Global Principles for Information Integrity	24 June 2024	The Secretary-General calls for a healthy information ecosystem; derived from Our Common Agenda (2021), the basis of the Summit of the Future (September 2024)
Council of Europe Framework Convention on Artificial Intelligence	Opened for signature 5 September 2024 (Vilnius)	First binding worldwide AI treaty; signed by the Commission on behalf of the EU

D.9 Academic literature

Real works, published by academic presses, cited in the book for their analysis - not for the authority of their authors:

Author	Year	Work	Publisher
Zuboff, Shoshana	2019	The Age of Surveillance Capitalism	PublicAffairs
Morozov, Evgeny	2011	The Net Delusion	PublicAffairs
Morozov, Evgeny	2013	To Save Everything, Click Here	PublicAffairs
Pariser, Eli	2011	The Filter Bubble	Penguin Press
Sunstein, Cass R.	2017	#Republic: Divided Democracy in the Age of Social Media	Princeton University Press
Nye, Joseph S.	2004	Soft Power: The Means to Success in World Politics	PublicAffairs
Wardle, Claire and Derakhshan, Hossein	2017	Information Disorder: Toward an Interdisciplinary Framework	Council of Europe

D.10 Verification portals

- **EUR-Lex:** eur-lex.europa.eu - consolidated texts of treaties and secondary law
- **European Commission:** ec.europa.eu - communications, strategies and Eurobarometers
- **European Parliament:** europarl.europa.eu - resolutions and minutes
- **Council:** consilium.europa.eu - conclusions
- **Court of Justice of the EU:** curia.europa.eu - case law
- **OECD:** oecd.org - trust surveys and AI principles
- **NIST:** nist.gov - AI risk management frameworks

Annex E: Glossary

E.1 How to use this glossary

Every entry reflects **exactly** how the term is used in the body of the book. Where a definition depends on a figure, the figure is the verified one used throughout the book.

E.2 Fundamental terms

Term	V3 definition
Digital Cognitive Sphere	The environment in which citizens interact with information, ideas and narratives through digital media: social networks, search engines, communication platforms and AI systems. The sphere the book declares worthy of constitutional protection.
Cognitive Protection	The set of measures, tools and policies protecting the digital cognitive sphere from disinformation, algorithmic manipulation and deepfakes. The object of the Protective Power and the EDSPA.
Disinformation	False or misleading information created and spread deliberately. The book adopts the Council of Europe taxonomy (Wardle and Derakhshan, 2017), distinguishing misinformation (false, no intent to harm), disinformation (false, with intent) and malinformation (true, spread with intent to harm).
FIMI	Foreign Information Manipulation and Interference. Term coined by the European External Action Service, which publishes annual threat reports (third report, 2024).
Deepfake	Synthetic content (image, audio, video) generated by AI that appears real but shows people or events that never happened.
Algorithmic Manipulation	Use of algorithms to influence behaviour or opinions in a hidden or non-transparent way; includes filter bubbles and deceptive content prioritisation.
Digital Sovereignty	The capacity of a country or region to control its digital infrastructure, data and technologies without dependence on external actors. EU strategic objective confirmed by the Technological Sovereignty Package (3 June 2026).
Brussels Effect	Mechanism by which European regulation becomes a de facto global standard because companies apply worldwide what they must comply with in the single market. Key precedent for the international projection of the European model (GDPR).

E.3 Institutional terms

Term	V3 definition
Protective Power (PP)	The proposal for the EU's 4th constitutional power: it institutionalises the protection of the cognitive sphere with a stable endowment, jurisdictional guarantees and European reach. It does not replace existing powers: it audits, publishes, refers and educates.
Pure Protector	The PP's operating doctrine: detect, publish, refer and educate; never neutralise, remove or sanction . The separation between whoever detects evidence and whoever decides consequences is the central guarantee against abuse.
EDSPA	European Digital Sphere Protection Agency. Transition agency 2027-2030 , created through enhanced cooperation with 9 founding Member States; it will lay the technical and trust foundations of the Protective Power.
Administrative Board of Appeal	Agency-level appeal guarantee, following the ACER model validated by the CJEU (C-46/21 P), with subsequent judicial review.
Specialised Chamber of the General Court	Jurisdictional review instance specialised in cognitive protection, with cassation before the CJEU: the PP is subject to the ordinary European judge, not to a tribunal of its own.
Enhanced Cooperation	Mechanism (Art. 20 TEU + Arts. 326-334 TFEU) allowing a subgroup of Member States - at least 9 - to advance without unanimity. Legal basis of the EDSPA; Art. 331 TFEU obliges the Commission to confirm the accession of any applicant within four months.
Treaty Reform (Art. 48 TEU)	Procedure for amending the Treaties: the Convention is convened by simple majority (the EP already activated it with Resolution 2022/2051(INL) of 22 November 2023), but ratification requires ALL Member States under their constitutional rules. Without unanimity there is no Treaty.
Viginum	French service created in 2021 to monitor foreign information interference: it detects and publishes, it does not sanction. A miniature precedent of the pure protector.

E.4 Operational and technical terms

Term	V3 definition
Human-AI Double Verification	Central guarantee of the system (Ch. 8): no automated finding becomes an administrative act without human review. AI proposes; human judgment decides and institutionalises.
Milestones and results, never procedures	The PP's publication rule: capabilities, guarantees and results are published; never models, stacks, volumes or methodologies. Applied to the book itself in Annex B.
Algorithmic Audit	Independent evaluation of algorithmic systems with public results. The PP audits and refers to the competent regulators (DSA, AI Act); it does not impose sanctions.
Citizen Checker	Public interface of the PP: a citizen can submit a suspect piece of content and receive a traceable verification. Protection made visible, not promised.
Annual Cognitive Protection Report	The PP's annual publication in the 24 official languages: evidence, Verifier results, verified cases and open data. The basis of accountability (Ch. 19).
The 60% Paradox	The gap between the demand for protection (69% consider disinformation a problem in their country, Eurobarometer published in February 2026) and distrust of governmental use of AI (~60%, OECD 2026). Only a protector with double verification, total publication and an independent judge can close it.
Technology Radar	The PP's technology watch system (Ch. 18): full internal view, synthetic external view; own capability catch-up estimated at up to 3 months upon a detected gap.
Cognitive Literacy	Citizens' competence to detect informational manipulation. The PP's educational function, serving Digital Decade 2030 (Decision (EU) 2022/2481); the book's goal: ≥90% literacy by 2038.
Resilient AI	AI systems designed to withstand attacks, manipulation and failures while keeping their functionality and guarantees.

E.5 Economic and financial terms

Term	V3 definition
EDSPA ROI	Return on investment calculated in Chapter 9B: 164% over five years, with an estimated net benefit of EUR 1,820M/year .
EDSPA Investment	EUR 1,100M initial investment, with 3,000 direct jobs in 2035 and 6,000-9,000 indirect; GDP impact EUR 3,160M-4,000M (Ch. 9B).
Protective Power Endowment	Stable financial endowment of the PP once constitutionalised: EUR 2,600M initial endowment, up to EUR 3,400M at maturity (Ch. 17). More than 15,000 direct employees at maturity 2040+ and an ecosystem of 55,000-68,000 jobs (Ch. 17).
Own resources (Art. 311 TFEU)	EU autonomous financing route (current system: Decision (EU, Euratom) 2020/2053) which the book proposes to explore for the PP's endowment: whoever depends on someone else's annual budget is not a power.

E.6 Geopolitical terms

Term	V3 definition
Big Tech	The large technology platforms dominating the cognitive sphere's infrastructure. Actors regulated by the European acquis (DSA, DMA, AI Act), not existential adversaries of the book.
The Two Centenaries	China's declared goals for 2021 (moderately prosperous society) and 2049 (PRC centenary: a modern socialist country). The book's temporal mirror: 2049, not 2039.
China Standards 2035	China's technology standardisation strategy. An example of the rival's method: dated prophecies are their tool, not the book's.
Soft Power	The capacity to influence through attraction (values, culture, results) rather than coercion. The PP's international projection works through the Brussels Effect: what works gets copied.
Three competing models	The taxonomy of Ch. 21: the market model (US), the control model (China) and the European democratic model. Each with strengths and weaknesses; only the European one can both monetise and protect attention - and only with an enforcer.

E.7 Normative framework (exact references)

Term	V3 reference
DSA	Regulation (EU) 2022/2065 on the Single Market for digital services. Platform obligations; Arts. 37 (annual independent audits, first round Aug 2024) and 40 (delegated act, July 2025, researcher data access).
AI Act	Regulation (EU) 2024/1689 . Harmonised rules on artificial intelligence.
GDPR	Regulation (EU) 2016/679. Personal data protection; Art. 25 (privacy by design).
Data Act	Regulation (EU) 2023/2854. Data access and sharing.
Digital Decade 2030	Decision (EU) 2022/2481: 20M ICT specialists, 80% basic skills 16-74, 75% of firms with cloud/AI.
European Democracy Shield	Commission Communication of 12 November 2025 against foreign interference.
Charter of Fundamental Rights	Arts. 7, 8, 11 (privacy, data, information) and 47 (effective judicial protection): the system's jurisdictional guarantee.
Council of Europe Framework Convention on AI	First binding worldwide AI treaty, opened for signature on 5 September 2024 in Vilnius; signed by the Commission on behalf of the EU.
UN Global Principles for Information Integrity	Presented by the Secretary-General on 24 June 2024; derived from Our Common Agenda (2021).

E.8 The book's dates (corrected)

Year	V3 milestone
2027-2030	EDSPA: transition agency through enhanced cooperation, 9 founding Member States.
~2030	Convention on Treaty reform (convenable by simple majority; the EP already voted to activate Art. 48(2) TEU in 2023).
2031-2032	Signing of the Treaty institutionalising the Protective Power.
2033	Entry into force of the Protective Power as the 4th constitutional power.
2033-2035	Integration of the EDSPA into the PP.
2038	Target of ≥90% cognitive literacy (PP educational function).
2040+	PP maturity: >15,000 direct employees, ecosystem of 55,000-68,000 jobs.
2049	The book's mirror: the PRC centenary. Used once, as a choice between finished models, not as a prophecy.

E.9 Acronyms

Acronym	Meaning
EU	European Union
TEU / TFEU	Treaty on European Union / Treaty on the Functioning of the EU (OJ C 326, 2012)
EDSPA	European Digital Sphere Protection Agency
PP	Protective Power
DSA / DMA	Digital Services Act (2022/2065) / Digital Markets Act
GDPR	General Data Protection Regulation (2016/679)
FIMI	Foreign Information Manipulation and Interference (EEAS)
ENISA	EU Agency for Cybersecurity
EDMO	European Digital Media Observatory (hubs until 2026)
EEAS	European External Action Service
CJEU	Court of Justice of the European Union
NIST	National Institute of Standards and Technology (US): AI RMF 1.0 (2023), Generative AI Profile AI 600-1 (2024)
ISO/IEC 42001:2023	First certifiable AI management system standard
ANSSI / NCSC	French cybersecurity agency (2009) / UK National Cyber Security Centre (2016): institutional precedents
OECD / UN / CoE	Organisation for Economic Co-operation and Development / United Nations / Council of Europe

E.10 Terms by part of the book (V3 structure, 21 chapters)

Part I — Diagnosis (Chapters 1-5)

Digital cognitive sphere · Disinformation (mis/dis/malinformation) · Deepfake · Algorithmic manipulation · Democratic vulnerability · Sovereignty gap.

Part II — The Patch (Chapters 6-10, including the new 9B)

EDSPA · Enhanced cooperation · Legal basis (Arts. 114, 16, 352 TFEU; 20 TEU + 326-334 TFEU) · 9 founding Member States · ROI 164% · Double verification.

Part III — The Definitive Solution (Chapters 11-20)

EDSPA→PP transition · Institutional architecture (Board of Appeal, Specialised Chamber, CJEU cassation) · Competences · Financing (Art. 311 TFEU) · Ecosystem of 55,000-68,000 jobs · Technology radar · 60% paradox · Annual Cognitive Protection Report · Art. 48 TEU path.

Part IV — The Final Vision (Chapter 21)

Three competing models · Brussels Effect · 2049 mirror.

Annex E - Glossary V3 | The Protective Power: Cognitive Protection as Europe's 4th Constitutional Power

Adrian Potentela | 25 September 2026

Annex F: Operational Specialisation Divisions

Why this annex exists, and why it is an annex

Chapters 12 and 13 fix the institutional architecture and the constitutional limits of the Protective Power; Annex B fixes its capabilities. This annex describes the dimension both leave implicit: **operational specialisation**. The cognitive sphere is not a homogeneous domain - manipulation travelling in a programmatic ad is not detected, audited or taught against in the same way as manipulation travelling in a bot network, or embedded in the design of a consent banner. An institution claiming to protect the whole sphere with a single toolbox will protect every part of it badly. Specialisation is not an organisational luxury of the PP: it is a **near-intrinsic function**, as inherent to its mission as double verification itself.

It is an annex - not a chapter - under the same doctrine that governs the whole book: divisions are **operational matter**. The PP's constitution (Chs. 12-15) fixes powers, limits, funding and trajectory; the operational organigramme that serves it can - and must - evolve through maturity triggers, not constitutional reform. Where a chapter change would require a Convention, a division change requires only the internal regulation and annual accountability before the European Parliament (Ch. 19).

F.1 The doctrine: specific capabilities, universal limits

Everything this annex states inherits without exception the rule of Annex B: what is published are **capabilities, safeguards and taxonomies; never procedures or techniques**. And it adds a rule of its own, governing every division without exception:

Specialisation divides the work; it never divides the limits. Every division of the PP has its own forensic AI and its own specialised operators, but the five functions of Chapter 13 - detect, audit, publish, refer, educate - and their limits (no sanctions, no filtering, no truth labels, anti-weapon symmetry) are **identical in all of them**. There is no "hard" division and "soft" division: there is one single doctrine applied to distinct domains of the threat.

This architecture also answers a foreseeable objection: specialisation might look like the first step towards a security agency with per-domain powers. The design forecloses it by construction. Every division produces the same three public artefacts of Annex B (report, reproducible evidence, referral file), transfers its files to the same sanctioners (Commission and national authorities: DSA 6%, AI Act 7%, GDPR) and operates under the same cross-cutting human-AI double-verification protocol of Chapter 8 (AI proposes, human decides; escalation to a second analyst on High/Critical disagreement). **Changing the domain does not change the nature of the power.**

F.2 The modular design principle: the division template

Every division is an instantiation of one template with five fixed components:

Component	What it is	Reference
Division-specific forensic AI	A detection system trained for the manipulation pattern of its domain; open, auditable and explainable	Ch. 18 (own AI); Annex B (capabilities, not techniques); F.7 (auditing the AIs themselves)
Specialised operators	Staff with the professional profile of the domain (see F.4); the human judgement always signs	Ch. 8 (double verification); Ch. 17 (AI counterweight: Europe's largest public employer of non-automatable human judgement)
Cross-cutting protocol	AI detection -> human verification -> escalation on High/Critical disagreement -> publication or referral	Ch. 8, mandatory and with auditable traceability
Acquis it enforces	No division creates new law: each makes an existing acquis systematic and verifiable	Ch. 13 (DSA, AI Act, GDPR/ePrivacy, Data Act, European Democracy Shield, EU KIDS Act)
Public interface	Output through the evidence portal and public metrics; no division has its own or secret channel	Annex B (B.6); Ch. 19 (transparency by default)

Modularity has one strategic consequence: **adding a new division does not touch the PP's constitution**. When in 2040+ a domain that does not exist today emerges (brain-machine interfaces, autonomous AI agents for individual-scale persuasion), the institutional answer will be an internal regulation and a template line - not a Convention. Specialisation is also the Protective Power's future-proofing clause.

F.3 The map of divisions

Division	Threat domain	Core capability	Acquis it makes systematic	What it never does
F.3.1 Adtech	Monetised manipulation: for-profit disinformation, opaque political advertising, exploitative microtargeting	Forensic audit of the advertising chain and the ads archive; evidence of who pays, what is targeted and with what data	DSA Arts. 39 (political advertising) and 26-28 (ads and recommender systems); GDPR Art. 22	Does not ban ads; does not set rates; publishes and refers
F.3.2 Synthetic Actors	Bots, troll farms, coordinated amplification, FIMI, identity deepfakes (see dedicated section below)	Public quantification of synthetic activity in public conversations: proportion, coordination and origin of artificial amplification; exposure with proof of digital imitations of identifiable persons	AI Act Art. 50 (AI transparency); EEAS FIMI classification; DSA Arts. 35 and 34 (deepfakes); Danish image law (see below)	Does not remove accounts (platforms/authorities); does not label individual content: publishes aggregate patterns; does not track people
F.3.3 Cognitive Architecture	Dark patterns, consent friction, addiction-by-design and rage-by-design	Design audit: public index of consent friction and manipulation architectures per platform	DSA Art. 25 (dark patterns banned); GDPR Art. 25 (privacy by design); ePrivacy; Digital Fairness Act (in legislative process)	Does not redesign platforms; ranks <i>designs</i> , never content or speech
F.3.4 Electoral	Electoral interference, campaign opacity, short-cycle FIMI operations	Campaign monitoring and a verifiable political advertising archive; activation of the EDSPA electoral protocol	DSA Arts. 36-39; Rapid Response System 2024 (EDMO); European Democracy Shield (Nov 2025)	Does not certify results; does not silence campaigns: publishes evidence for voters and authorities
F.3.5 Minors	Addictive design for children, fragile age-gating routes, exposure to manipulation; predation (grooming) as a facilitated architecture	Child-compliance audit per platform (EU KIDS Act), public safe-by-design index for minors and early-literacy materials	EU KIDS Act (COM(2026) 681, Sep 2026); DSA Arts. 28-30 and minors protection guidelines (Jul 2025); GDPR Arts. 8 and 25; European age-verification blueprint; Cyberbullying Action Plan; Digital Fairness Act (in legislative process)	Does not surveil minors or private conversations (criminal matter: States/Europol); does not filter content (platforms/parents/States): audits design and compliance, and refers aggregate patterns
F.3.6 Conversational Platforms	Unidentified AI agents in public debate, systematic model bias, synthetic persuasion	Compliance check of AI Act Art. 50 transparency and audit of declared biases	AI Act Arts. 50 and 6; DSA Art. 35	Does not "approve models"; does not sanction their deployment: publishes and refers
F.3.7 Public Metrics	Structural opacity: the impossibility of comparing platforms and periods	The cross-cutting publication function: annual report, digital nutrition label, comparable series per platform and State	Regulation 1049/2001; Art. 15(3) TFEU; Art. 42 Charter	Does not interpret: measures, aggregates and publishes with open methodology

F.3.2.bis Synthetic digital identity: the pattern and the proof


One use case concentrates every red line of the PP and deserves its own section: the detection of realistic digital imitations of identifiable persons - face, voice or body deepfakes - in public circulation. The legal anchor already exists: **Denmark** was the first EU State to recognise every person's intellectual-property rights over their own image, voice and body (Ophavsretsloven reform, in force 1 July 2026), its rotating presidency of the Council of the EU advocates extending the model across Europe, and the DSA (Art. 34) and AI Act (Art. 50) already oblige platforms and deployers to transparency about synthetic content. One factual precision matters: the Danish law grants the citizen a **right**; it creates no state biometric registry whatsoever - it is the rights-holder who authenticates themselves before their authority.

On that acquis, the PP's workflow is strictly its five functions - nothing new:

1. **Detect the pattern:** the division's forensic AI identifies digital imitations of identifiable persons and their circulation in the public sphere, under human-AI double verification (Ch. 8).
2. **Publish the proof:** exposure with reproducible evidence on the public portal - "this content is synthetic; here is the proof".
3. **Refer to the three natural recipients:** the **platform** hosting it (DSA/AI Act obligations), the **AI developers involved** (generating-model traceability, AI Act Art. 50) and the **competent authorities**.

And one thing the PP deliberately does **not** do: **notify the harmed citizen**. Identifying the affected persons and notifying them is for the authorities - which hold the mandate, the data and the legality to do so. The PP does not know the victims and need not know them: *it exposes the manipulation; it does not track people*.

Why there is no citizens' counter. A service receiving direct requests would be a collapse target: any hostile actor could generate millions of fake requests to saturate the PP and proclaim its failure. The design excludes it by construction - an individual case reaches the PP through only two routes: it emerges from public pattern monitoring, or it arrives as a **judicial request** (see the Forensic Judicial Department, F.7). The mailbox through which the avalanche would arrive does not exist.



The result is a division that protects the emerging right of European citizens over their digital identity without ever becoming an identification agency: one single technical evidentiary standard for Europe, and as many legal gateways as there are States - until the convergence pursued by the Danish presidency does the rest.

F.4 Operators: the human profile of each division

Specialised AIs would be hollow without their human counterpart. Each division recruits profiles that exist today scattered across the European labour market and that the PP aggregates into a new public career - the qualified employment of the Ch. 17 ecosystem:

Division	Operator profiles (existing today in the market)
Adtech	Programmatic auditors, ad-fraud specialists, adtech economists, advertising and consumer-law jurists
Synthetic Actors	OSINT analysts, forensic linguists (stylometry), network-graph analysts, FIMI experts (EEAS/Viginum profile), synthetic-media forensic examiners
Cognitive Architecture	Cognitive psychologists, audit UX specialists, GDPR/ePrivacy jurists, behavioural scientists
Electoral	Political scientists, electoral-systems experts, per-Member-State analysts, electoral jurists
Minors	Developmental psychologists, behavioural paediatricians, digital mediators, educators, child safe-design auditors
Conversational Platforms	LLM specialists, algorithmic auditors, AI Act jurists, synthetic-persuasion "red teamers"
Public Metrics	Data scientists, data journalists, statisticians, communication specialists (Ch. 19)

Three design observations. First: the scarcest profiles (algorithmic auditors, forensic linguists) are precisely those the PP would have to train internally - an additional function of the Cognitive Education department of Ch. 12. Second: each division feeds the 60% paradox (Ch. 19) with its design answer: the more specific AI, the more human judgement signs - the AI counterweight of Ch. 17 scales with specialisation, not against it. Third: dual affiliation is the rule - every operator belongs to a division *and* to the cross-cutting Ch. 8 protocol, so inter-division analyst exchange is an institutional (traceable) capability, not an exception.

F.5 Staged activation: specialisation follows endowment

Divisions are not all born at once: they activate through the same maturity triggers that govern the staffing trajectory of Chapters 15 and 17. This annex fixes no division dates - it fixes the order and the condition:

Stage	PP staff	Active divisions	Activation condition
EDSPA (2027-2030)	1,500-2,000	Embryo: Electoral + Synthetic Actors (agency mandate, Ch. 8) + child-design audit component of Minors (EU KIDS Act verification, COM(2026) 681; application expected from 2027 once adopted)	EDSPA mandate (Ch. 8): electoral protection and synthetic-activity monitoring
Core (2033-2035)	~3,000	+ Adtech, + Public Metrics; Minors expands into early literacy	Favourable 2030-2031 impact report; PP internal regulation
Full (~2037)	~7,000-8,000	+ Cognitive Architecture, + Conversational Platforms; Minors at near-full capacity	27 States / 24 languages / 24-7 operations (Ch. 15 triggers)
Maturity (2040+)	>15,000	+ Minors at full capacity; new divisions for emerging domains	Biennial EP evaluation; only if favourable

Rigour note: the staff column is the endowment trajectory (Chs. 12, 15, 17), not a per-division hiring forecast. The internal staffing split across divisions belongs to the internal regulation and the Annual Report - publishable as a result, not budgetable in advance in an annex. The one deliberate exception is the child-design audit component of Minors in the embryo: its justification is not organisational but demand-driven - the EU KIDS Act (proposal COM(2026) 681, Sep 2026) requires systematic safe-design verification exactly in the window when the EDSPA starts.

F.6 Budgetary and safeguards coherence

Two closings so this annex does not misalign the book's economy. First, **budget**: divisions are subunits of the four departments of Chapter 12 (Cognitive Protection, Digital Sovereignty, Cognitive Education, Cognitive Intelligence) within their full-stage ceilings (EUR 1,000M/500M/300M/200M) and the Ch. 17 framework - no division has its own budget or new line. The most natural correspondence: Adtech, Synthetic Actors, Cognitive Architecture, Electoral, Minors and Conversational Platforms sit under Cognitive Protection; Public Metrics shares its interface with Cognitive Intelligence - its function is cross-cutting publication (an interface), not field operations, and Intelligence's ceiling suffices. Second, **safeguards**: every division inherits in full the B.5 block of Annex B - non-automatable double verification, transparency by default, privacy by design, anti-weapon symme-

try and strict detection/sanction separation. The Administrative Board of Appeal and the Specialised Chamber for Digital Protection of the General Court (Ch. 12) are the avenue of appeal against any act of *any* division.

F.7 The audit of the auditors: the meta-level of specialisation

Specialisation raises one last question - the one any informed sceptic will ask: *who audits the specialised AIs, and the technicians who train, operate and certify them?* If every division has its own forensic AI, the system rests on models nobody outside the PP can see. The design's answer is institutional: the PP audits its own systems under the same doctrine with which it audits the sphere - **the light starts at home**.

Three pieces. **First: training traceability.** Every division model maintains an auditable record of its training data, versions, performance and drift - the same traceability standard Ch. 8 demands of the platforms the PP audits. **Second: internal training of scarce profiles.** The algorithmic auditors and forensic linguists the system needs barely exist today in the European labour market (F.4); the PP trains them internally through Cognitive Education, and that public career includes the certification of the auditors themselves - who may audit a division AI is public information. **Third: double verification applied inwards.** The Ch. 8 protocol also governs the PP's own systems: no finding produced by a division AI is published or referred without human sign-off, and the PP remains Europe's largest public employer of human judgement over its own models too - the Ch. 17 AI counterweight, applied to the inside.

The Forensic Judicial Department - the "magnifying glass" of the digital identity workflow (F.3.2.bis) - anchors here as a **subunit of Cognitive Intelligence**: transversal technical support serving judicial cooperation, with no fifth department and no new budget line (within Intelligence's EUR 200M ceiling, Ch. 12). Its rules are strict: **only a judge** - the sole organ with procedural guarantees (parties, defence, adversarial process) - may request an individual forensic analysis of a person's digital sphere; no administrative authority and no private party accesses the channel. Every request is logged with sitting judge, case number, jurisdiction and object of the analysis (protected operational archive of Ch. 7); every analysis, with signing operator, double verification, division AI used and data accessed - all consultable by the court and the guarantee bodies (Administrative Board of Appeal on the ACER model, Specialised Chamber of Ch. 12). The radar becomes a magnifying glass only on a judge's order; and the glass looks only at what the judicial request fixes. If anything exceeds the scope of the request, the traceability shows it.

What this annex deliberately does not do: it creates no new competences and does not dilute the five functions of Chapter 13 - it multiplies them by domain under the same constitutional limit. It publishes no architectures, models, thresholds or methodologies of any division AI (rule B.1) - the F.7 audit is published as guarantee and capability, never as technique. And it does not turn specialisation into bureaucracy: if a division produces no verifiable public artefacts, the Annual Report must be able to say so - the light starts at home.