



Guide on Artificial Intelligence-based Projects Licensing and Risk Management Basics

Natalia Zhuravleva

Outline and purpose of the document: This Guide on AI Licensing and Risk Management Basics is part of the Classical Plus IMPAC3T-IP toolkit. It helps researchers and Knowledge Exchange (KE) / Technology Transfer (TT) Offices without former experience with AI-based projects to plan, protect, license, and govern AI-enabled outputs through the full life cycle - from early choices and data governance to IP protection, EU compliance, commercialization, and post-market duties. The guide gives a practical, step-by-step framework to understand rights in the AI stack (data, code, models, and outputs), choose fit-for-purpose licenses, and manage regulatory and operational risks so projects can reach the EU and global market lawfully and sustainably.

Target readership: Researchers, interdisciplinary project teams, spin-out founders, and KE/TT professionals involved in creating, evaluating, licensing, or deploying AI components; research managers and Higher Education Institutions (HEI) administrators supporting these activities; and partner SMEs or public bodies working with academic AI assets.

Level: The document assumes a solid research or technical background but no prior expertise in EU AI regulation. Readers should seek legal, data-protection, or regulatory advice where needed to ensure compliance with institutional policies and applicable laws.

Focus: Clear guidance on (i) business and deployment models; (ii) data governance and provenance; (iii) ownership and contributor agreements; (iv) EU AI Act risk tiers and conformity steps; (v) IP protection across patents, copyright, database rights, and trade secrets; (vi) open-source and hybrid licensing; (vii) key contract clauses and compatibility checks; (viii) investment, due diligence, and risk mitigation; and (ix) lifecycle and post-market duties (monitoring, incident response, and sunset).

Scope: To give the reader understanding of potential issues and complications of work with commercialization and licensing of AI-based research and projects in the EU-context; to explain complex legal and procedural topics in a practical, accessible way using checklists, decision points, and concise explanations tailored to academic-to-market transitions.

Disclaimer: This guide is not intended to be exhaustive. It provides an overview of issues that may arise when managing IP, licensing, compliance, and risk in AI projects. IMPAC3T-IP and project partners accept no responsibility for outcomes of its use. It is not a substitute for qualified professional legal, IP, data protection, or regulatory advice.

Index

1.	Preface	1
1.1.	Outline and Purpose of the Document.....	1
1.2.	Target Readership	1
1.3.	Terms.....	1
2.	Understand the factors influencing the process.....	3
2.1.	Know the EU AI Product Landscape	3
2.2.	Consider Other Drivers (sector, funding, risk)	4
3.	Prepare Before You Build or Share	4
3.1.	Map Your Risk Tier under the EU AI Act	4
3.2.	Define Your Business and Deployment Model	5
3.3.	Set Up Data Governance and Access Rights	6
3.4.	Agree Ownership and Contributor Terms.....	7
3.5.	Use Retrieval-Augmented Generation (RAG) lawfully	8
4.	Protect What You Create	8
4.1.	Plan for Patents	9
4.2.	Secure Copyright and Database Rights.....	10
4.3.	Keep Trade Secrets	11
4.4.	Choose Open-Source or Hybrid Licensing.....	12
4.4.1.	Open Ecosystem primer	13
4.5.	Consider Defensive Publishing	14
5.	Navigate EU Rules and Act on Ethics	15
5.1.	GDPR + ePrivacy	15
5.2.	Respect Data Sovereignty and Access	16
5.3.	Embed Responsible-AI Practices	17
5.4.	Be ready for Real-world risk challenges.....	18
6.	Choose and Structure Your Licensing Deal.....	19
6.1.	Pick a License Archetype.....	19

6.2.	Draft the Key Clauses.....	20
6.3.	Check Upstream License Compatibility.....	21
6.4.	Plan Exit, Hand-back, and Escrow	22
7.	Prepare for Investments.....	23
7.1.	Triage Due-Diligence Hot Spots	23
7.2.	Show Compliance in Valuation.....	24
7.3.	Apply Risk Mitigation Measures.....	24
8.	Operate, Monitor, and Improve Post-Launch.....	25
8.1.	Watch for Model Drift and Re-train.....	25
8.2.	Handle Incidents and Recalls	26
8.3.	Plan for Sunsets and Transitions	27
9.	Annex 1: Regulatory documents and useful guides.....	29
10.	Annex 2: Popular AI models & their licenses	35
11.	Annex 3: Open-Source License Compatibility	36
12.	Annex 4: Useful contacts.....	38

1. Preface

1.1. Outline and Purpose of the Document

This document is a concise guide for individuals who want to understand key steps for creating and licensing products with artificial intelligence components in the European Union. It follows a life-cycle structure: first outlining early strategic choices, then explaining how to secure intellectual-property rights and meet regulatory duties, and finally covering commercialization, risk management, and long-term maintenance. By walking through each phase in order, the guide helps the reader spot legal and technical issues at the right moment and build an AI product that can reach the European market smoothly and lawfully.

The guide is organized as a practical end-to-end path. It opens with process factors that shape AI licensing decisions, then moves to pre-start checks (business model, data, ownership, risk tiering). Next, it explains IP protection routes, followed by the EU legislative and ethical framework and the concrete actions this implies. The middle chapters cover commercialization and licensing models with key contract clauses and compatibility checks, then investment and risk management. The final chapter addresses lifecycle and post-market duties (monitoring, incidents, sunset). Annexes provide core regulatory documents, examples of popular AI models and their licenses, an open-source license compatibility map, and useful contacts.

1.2. Target Readership

This guide is written for researchers and interdisciplinary project teams who aim to commercialize their work and need a clear view of how intellectual-property rights and licensing work when artificial-intelligence components are involved, as well as for Technology Transfer and Knowledge Exchange professionals who are new to the processes of licensing products and technologies with artificial-intelligence components. It assumes a solid technical or research background but no prior expertise in patent law and AI-based products licensing.

1.3. Terms

Term	Description
AI Model	The mathematical structure (e.g., neural network) that turns input data into outputs after training.
Algorithm	A fixed set of rules or steps the model follows to learn or make decisions.
Bias	Systematic errors in data or models that give unfair results for certain groups.
Conformity Assessment	Formal procedure required under the EU AI Act to show a high-risk AI system meets safety and transparency rules.
Copyleft License	An open-source license that forces anyone who distributes modified code to share it under the same license (for example, General Public License GPL).

Copyright	Legal right that gives its owner the exclusive legal right to copy, distribute, adapt, and reproduce an original creative work (code, text, images, audio), usually for a limited time. Can be limited based on public interest considerations. Copyright does not normally protect ideas, methods, or purely machine-generated material.
Data Card / Datasheet	A provenance summary for datasets: motivation and collection context, composition, consent and TDM status, preprocessing/labeling, allowed uses/licensing, distribution, and maintenance plan.
Data Governance	Policies and processes that make sure data is collected, stored, and used lawfully and securely.
Data Sovereignty	Principle that data stays under the legal control of the country or region where it is stored or processed.
Dual Licensing	Offering the same software or model under two different licenses (for example, General Public License GPL for research and a paid proprietary license for commercial users).
Fair Use / Quotation Exception	Limited permission in some jurisdictions to use small parts of a work without a license for criticism, teaching, or research; scope differs by country. In the EU, limited use without a license relies on specific exceptions (e.g., quotation, teaching, research) defined by national law implementing EU Directives. In the US, “fair use” is a broader, case-by-case doctrine. Do not treat “fair use” as an EU concept; provide EU exception mapping instead.
Generative AI	Artificial intelligence systems that can create new content (text, images, audio) rather than just classify existing data.
Inference	The run-time phase when a trained model processes new data and produces a result.
License	Contract that states how you may use, modify, or share intangible assets.
Model Weights	Numerical parameters learned during training; they define how strongly each input feature affects the output.
Model Card	A short transparency record for a model: intended use and limits, data sources, training procedures, performance metrics, safety/fairness notes, and contact issues.
Open-Source Software (OSS)	Code released under licenses that allow anyone to use, see, and modify it, often with conditions on redistribution.
Open-Source AI Model	Models or weight files released under licenses that permit reuse/modification (subject to license terms). “Open” here refers to license terms, not to absence of obligations (e.g., attribution, share-alike, or RAIL restrictions).
Patent	Exclusive right granted for a technical invention that provides a new, non-obvious solution to a technical problem.
Personal Data	Any information that can identify a living person, subject to General Data Protection Regulation (GDPR) and similar privacy laws.

Post-Market Monitoring	Ongoing checks and incident reporting required after an AI-based product is launched, especially for high-risk systems.
Retrieval-Augmented Generation (RAG)	An architecture where prompts are enriched with retrieved documents from a controlled knowledge base; licensing, database right, and TDM rules apply to the retrieved corpus.
Trade Secret	Valuable business information is kept confidential to maintain a competitive edge (in the document context, for example, model weights or training methods).
Training Data	The examples are fed into a model so it can learn patterns and make predictions.

2. Understand the factors influencing the process

Read this section to learn more about the framework which will structure your AI-related projects and products development.

2.1. Know the EU AI Product Landscape

The European Union has built a layered set of rules that any commercial AI product must be utilized. At the core sits the [EU AI Act](#), which classifies systems into *Unacceptable*, *High*-, *Limited*-, or *Minimal-Risk* tiers (see more in the following sections). High-risk systems - covering medical, finance, Human resources (HR), critical infrastructure, and similar fields - require conformity assessment, quality-management system, Conformité Européenne (CE) marking, and ongoing incident reporting. Limited-Risk tools (e.g., chat-bots) need only transparency notices, while Minimal-Risk tools face no special duties beyond general law.

Around this act there are several data-centric laws. [GDPR](#) and the [ePrivacy Directive](#) govern personal data collection, consent, and users' rights such as erasure and objection to profiling. The [Data Governance Act](#) and the [Data Act](#) extend control to industrial data, mandating data-sharing frameworks and safeguards on cross-border transfers. Together they demand a legal basis for every data set, technical measures for deletion on request, and contractual clauses that keep data within approved jurisdictions or under Standard Contractual Clauses.

There are sector-specific rules that add further layers. Health AI must meet the [Medical Devices Regulation](#); financial tools fall under the [Digital Operational Resilience Act](#) (DORA) and existing banking directives; autonomous-driving software intersects with [UNECE vehicle regulations](#); dual-use or defense projects trigger [EU export-control rules](#). [NIS 2](#) widens cyber-security duties for critical service providers, and all AI products must comply with the [Cyber Resilience Act](#) once it enters force (secure-by-design, vulnerability reporting).

Beyond binding law, the EU promotes [ethical guidelines for Trustworthy AI](#): lawfulness, transparency, fairness, accountability, and human oversight. While not hard law, these principles are increasingly referenced in procurement and investment checklists, making them de facto market standards.

See the [Navigate EU Rules and Act on Ethics](#) section below for more details.

2.2. Consider Other Drivers (sector, funding, risk)

Consider other drivers before picking up a licensing path. Check sector rules and founder terms. Review your organization's know-how and trade-secret rules (confidentiality, disclosure approval, publication embargoes) and your institution/funding organization for KE/TTO/IP policies. Check data governance: classification, provenance, legal basis, data protection risks, retention, access, and sharing/transfer. Make sure these do not conflict with partner agreements/collaboration agreements/funding agreements or your protection choices.

Your deployment model also changes the picture: it affects which IP rights you must secure, what license type fits, and your liability if the system fails or infringes a third party. See Section 3 for the list of common business and deployment models. Also consider basic procurement duties, required security/safety standards in your domain, and any cross-border data or export limits that may apply.

3. Prepare Before You Build or Share

3.1. Map Your Risk Tier under the EU AI Act

The [EU AI Act](#) divides AI systems into four risk categories, each with its own rules:

- **Unacceptable** **Risk**
These systems are outright banned. Examples include social scoring by public authorities, real-time biometric surveillance in public spaces, and AI for manipulative or exploitative advertising. You cannot place these products on the EU market or put them into service.
- **High** **Risk**
Before launching, you must complete a conformity assessment, implement a quality-management and risk-management system, prepare detailed technical documentation, obtain CE marking, and register in the EU database. Post-market monitoring and incident reporting are also required. This category includes systems used in safety- or rights-critical areas, such as medical-diagnosis tools, credit-scoring algorithms, recruitment software, and components of critical infrastructure.
- **Limited** **Risk**
Systems included into this category are not forbidden but must meet transparency obligations. You need clear user notices (for example, "you are talking to an AI") and guidance on how to detect or report misuse. No full conformity assessment is needed, but these disclosures must appear in the user interface and documentation. Applications like chat-bots, deep-fakes or emotion-recognition tools fall into this category.
- **Minimal** **or** **No** **Risk**
Minimal or no risk systems remain subject to general EU law (consumer protection,

product safety, data protection) but face no extra AI-specific rules. Ordinary software tools with incidental AI components - spell-checkers, basic recommendation engines, or simple image filters - do not trigger specific AI Act duties.

Action checklist

- To determine your tier, consult [Annex III](#), perform a self-classification exercise, and document the result. This classification then drives your compliance roadmap - defining which audits, labels, or user-notices you must implement.

See more details in the [Regulatory & Ethical Framework](#) section below.

3.2. Define Your Business and Deployment Model

Start by fixing your business model (how value is created and paid for) and your deployment model (where the system runs and who operates it). These two choices shape almost every later decision. They decide which assets you will ship (code, model weights, datasets, prompts) and which you will host; who is the operator of record; how data moves for training, fine-tuning, and inference; and which standards and approvals you must meet. They also drive the IP stack you need to secure (patents, copyright, database rights, trade secrets), the licenses you draft (scope, field-of-use, redistribution, attribution), and the liability and compliance split across you, customers, and partners.

Choosing the model early reduces rework. It allows you to time publications and patent filings, decide what stays a trade secret, select open-source, commercial, or hybrid terms that truly fit the delivery, plan for updates, monitoring, and incident response, and set clear indemnities, warranties, and SLAs/DPAs that match who operates the system. It also enhances partner discussions and investor due diligence, as there is a clear map of assets, rights, duties, and risks.

Here is a non-exhaustive list of possible business models you can consider:

- **Embedded device (edge AI in hardware, for example smart sensor, medical scanner)**
 - IP scope: firmware, model weights, and hardware schematics need clear ownership and patent clearance; updates may require re-certification.
 - Liability: product-safety law applies; vendors face strict or even no-fault liability for defects or high-risk AI under the EU AI Act.
- **On-premise enterprise software (licensed to run in the customer's data center)**
 - IP scope: source code and models are delivered to the client, so escrow, reverse-engineering clauses, and copyleft compliance are critical.
 - Liability: customer controls infrastructure, but suppliers still bear warranty for performance and security patches.
- **B2B SaaS / AI-as-a-Service API**
 - IP scope: provider keeps all code and weights; users get output rights only. Service Level Agreements must cover uptime and data handling.
 - Liability: data-protection duties stay with both parties; failure to meet SLA or data-breach can trigger contractual penalties.
- **White-label platform for channel partners**
 - IP scope: branding elements are licensed out; need clauses on how partners may modify UI or retrain models.

- Liability: partner is the “face” to end-users, so contracts must pass through indemnities and compliance duties.
- **B2C mobile or web app (consumer generative-AI service)**
 - IP scope: user-generated content terms must state who owns outputs and whether they feed future training.
 - Liability: consumer law rules on transparency, unfair terms, and minors; possible copyright claims over generated media.
- **Internal decision-support tool (no external users)**
 - IP scope: easier - IP stays in-house but checks licenses for any open-source components.
 - Liability: mostly internal, yet regulators may still audit if tool influences safety-critical choices (e.g. clinical decisions).
- **Consultancy / bespoke model-building service**
 - IP scope: clarify who owns the trained model, fine-tunes, and derivative datasets; often split between client (weights) and provider (underlying framework).
 - Liability: professional services liability for errors; confidentiality and trade-secret clauses protect both sides.

3.3. Set Up Data Governance and Access Rights

You need a solid data-governance plan to tell where every bit of data comes from, on what legal basis you hold it, and how you can prove that later. Decide all of this before you start collecting or buying datasets, because the rules differ for research and for commercial release. If you gather personal data under a “research only” consent or license, you may have to delete or re-license it before you can sell the product. In some cases, there might be a way to base your final project on another set of data than one used for the research, but such approach might be too complicated and risky, thus, undesired by potential collaborators and investors.

Addressing all these questions up front, followed by recording and updating all the answers in a special data-governance file, helps you to avoid most last-minute license gaps, and gives all commercialization stakeholders (including collaborators and potential investors) the clear compliance story they expect.

Action checklist

- Source mapping and licenses – List each dataset, its owner, the exact license text, and any “text-and-data-mining opt-out”. Check that permissions cover *commercial use*, not just academic study.
- Dual-track consent – When you collect data from people, use a consent form (or other legal ground) that already includes future commercialization; otherwise, you will need to re-contact every subject.
- Provenance logs – Keep a live log that links training runs to specific data versions. This proves lawful origin, helps you trace errors, and supports later audits under the EU AI Act.

- Opt-out and deletion rights – Build a workflow that can locate and delete an individual's data - and, where feasible, the related embeddings or fine-tunes - within GDPR time limits.
- Updates and re-training – Document how you will re-train the model after any large deletion or new data intake, so outputs stay accurate and compliant.
- Retention and sunset rules – Set clear timelines for how long raw data and processed logs are kept, and how they will be wiped when no longer needed.
- Cross-border storage – Note where each dataset is physically stored; if any server sits outside the EU, add Standard Contractual Clauses or move the data to an EU-based cloud.

3.4. Agree Ownership and Contributor Terms

List the chain-of-title for every asset (including source code, model weights, datasets, and generated outputs) with the names of the people person/organizations that created and now own each item. Link this list to signed agreements so rights flow without gaps from individual contributors to the university/spin-out/licensee. Pay special attention to datasets to regulate data use, data management, and exploitation of the result (mind that different countries have different concepts about this).

Every contributor (staff, students, visiting researchers, freelancers) should sign an IP assignment that:

- transfers all present and future rights in their work to the project lead or holding entity;
- waives or licenses any moral rights if local law requires;
- confirms there are no conflicting obligations to previous employers or funders;
- set a publication embargo (to let researchers publish theses or papers after a set delay, for example, 6–12 months, while patents are filed, trade secrets are put in place; IP strategy is finalized).

Action checklist

- Background and side-ground IP – list any pre-existing code or data brought to the project and state what are the conditions of use (including main terms on which it is licensed or assigned)
- Improvement rights – decide who owns future fine-tunes, retraining runs, or derivative datasets and whether royalties apply.
- Open-source contributions – require written approval before releasing project code under an external license; attach an approved license matrix.
- Warranties and indemnities – ensure each contributor's work is original and indemnifies the project against hidden third-party claims.
- Exit and hand-back – outline how rights return or transfer if a partner leaves the consortium, ensuring no asset gets stranded.

3.5. Use Retrieval-Augmented Generation (RAG) lawfully

RAG lets you ground answers in your own corpus without re-training the base model. Treat the corpus as a licensed knowledge base:

- License check – confirm the corpus permits Text and Data Mining (TDM). In the EU, Articles 3–4 of the CDSM Directive allow TDM on content you have lawful access to unless the rightsholder opted out; “research-only” data may not allow commercial use.
- Database right – bulk extraction from EU databases may trigger the sui generis right even if individual items are public domain.
- Confidentiality – if the corpus contains trade secrets or personal data, add access controls, minimization, and opt-out deletion paths.
- Outputs – keep logs of retrieved chunks for auditability and for take-down workflows.

Action checklist

- Keep a TDM register (source URL, license/opt-out, purpose).
- Add “no training / RAG-only” clauses where needed in supplier contracts.
- Show retrieved context in UI for transparency in high-risk uses.
- Re-license content or replace it before commercialization if the current terms are research-only.

4. Protect What You Create

Your AI-Act risk tier (see [Section 3.1](#)) and deployment model ([Section 3.2](#)) decide which IP route and license shape will actually work. High-risk uses often require tighter documentation, auditability, and post-market controls; that, in turn, affects whether you keep weights as trade secrets, publish defensively, or file patents. This section should be read with your risk tier and deployment model in mind.

Intellectual assets rights (intellectual property rights) are the main tools that let you control, share, or sell an AI product. They also shape how investors value the project and how easily a competitor can copy your work. Because an AI stack contains many layers - datasets, source code, model weights, generated outputs - no single right gives full cover. You might combine several: patents for technical inventions, copyright and database rights for code and data, trade secrets for confidential know-how (weights, training models), and open-source or hybrid licenses to encourage collaboration while keeping commercial options open.

What is typically patentable in AI

- Novel model architecture is tied to a technical effect (e.g., reduced bandwidth/latency/energy on-device).
- Data selection/augmentation pipelines that improve sensor or network performance in a measurable way.
- Training procedures that control hardware or memory layout to achieve deterministic technical savings.
- Deployment adaptations that transform a model for constrained targets (DSP, MCU, edge accelerator).
- Post-deployment processing that stabilizes predictions in a safety-critical loop (e.g., medical or automotive).

See [EPO guidance](#) and AI case law news.

The next subsections explain each of these protection routes, show when they apply, and warn about typical pitfalls (for example, patent-eligibility limits for “pure” algorithms or copyleft clauses that can force you to reveal source code).

4.1. Plan for Patents

Technical solutions (technical application or technical implementation) can be patented, not abstract math or algorithms. Pure algorithms (e.g., network topologies, loss functions, training tricks) are usually treated as mathematical methods and excluded “as such.” Under the European Patent Office (EPO) approach, AI and machine learning (ML) claims are patent-eligible when the method is implemented in, and measurably improves, a technical system (e.g., cuts bandwidth, storage, latency, energy) or improves the application (control of a machine, quality of a phone call, quality of medical diagnostics).

To illustrate “improvement”, say what technical resources are saved (bandwidth, memory, CPU, energy) and where in the system it happens (client, network, server, device). Show that the effect is causal and repeatable over the whole claim scope (not a mere by-product of better accuracy). Tie the gain to a technical objective (e.g., fewer bytes sent; fewer writes to disk, lower latency) with measurable thresholds or control logic.

To show “implementation”, go beyond math. Specify how the method is realized in the system: data structures, memory layout or message format; control signals or feedback loops; scheduling, caching, or rate-limiting; placement of computation (edge vs. server) and data-flow constraints; interfaces and failure handling.

See the cases from the Boards of Appeal:

- T 0702/20 (sparsely connected hierarchical NN): benefits were only “within the computer” (e.g., storage/implementation effects). The claim failed ([EPO](#), [Appleyard Lees](#), [Novagraaf](#)).
- T 0183/21 (recommender system with feedback control): the Board upheld claims because the method reduced network bandwidth and server storage; savings were framed as solving a technical problem in a real system ([EPO](#), [Lexology](#), [BARDEHLE PAGENBERG](#)).

Action checklist

- Define a specific technical problem your model solves - bandwidth, energy use, sensor accuracy, safety margin, etc.
- Show cause and effect: link algorithmic steps to that problem; avoid vague “better performance” claims.
- Provide quantitative evidence or at least plausible calculations of the improvement.
- Describe the implementation context - hardware set-up, data flows, control loops, signal paths.
- Highlight novel hardware cooperation (edge devices, specialized accelerators) enabled by the model.
- Add fallback claims that tie the AI to different technical effects or subsystems in case examiners narrow your main claim.
- Run a pre-filing review with specialists (the EU IP Helpdesk offers free guidance) to check the draft clears the “technical effect” bar in all target jurisdictions.

4.2. Secure Copyright and Database Rights

Software and some data can be protected by copyright. To be protected, data and datasets should be proved to be based on a certain level of creativity. That means that machine generated data, as well as individual data items resulting from direct collection (personal data items, pharmaceutical data items, physical measurement items, experimental data) cannot be protected by copyright. Other types of individual data with certain creativity component items (images, text, audio) can be protected by copyright. Compilations of the machine generated data can also be protected by copyright if a certain level of creativity can be proved.

At the same time, the method of data collection can be protected, as it may involve an element of creativity.

Also, in the EU, data collection can also be protected by the [sui generis database right](#) if making or checking the database took substantial investment (that right also blocks copying a “substantial part” of the database). The Directive on the Digital Single Market ([DSM Directive](#), Article 4) lets anyone get access to the text-and-data mining content (TDM content) to which they have lawful access if rightsholders didn't opt out of this. That means that even if single items are public domain, the database right may still restrict bulk extraction. Remember that the EU's sui generis database right has no U.S. equivalent; in the U.S., databases are protected mainly via copyright in selection/arrangement and by contract. When exporting models or datasets, fill gaps with contracts and trade-secret controls; for background, see the U.S. Copyright Office's report on [database protection](#).

In practice, record each dataset's owner, license, any opt-out, and the exact attribution wording in your contract with the financing organization (for example, the project's Horizon Europe DMP template).

Mind that training your system on performance recordings can trigger performers' and producers' neighboring rights, even if the composition itself is public domain. Check and clear

any such rights before launch: see the [WIPO WPPT](#) and EU rules in the [InfoSoc Directive 2001/29/EC](#) and [2006/115/EC](#).

If initial dataset used in the project was protected by any license, mind that even after heavy preprocessing (tokenising, normalising, vectorising) which creates an “adapted” or derivative dataset, it stays tied to the original license where downstream commercial use can be not allowed (see Creative Commons definition of “[Adapted Material](#)”). Keep a change-log of every transformation and re-confirm permissions in your data management plan.

Training scripts, pipelines, and serving wrappers are protected as “literary works” under the EU [Software Directive 2009/24/EC](#). Track authorship and licenses, and watch for copyleft triggers: the [GPLv3](#) activates on distribution of binaries; the [AGPLv3](#) adds a network-use disclosure duty.

Model weights rarely reflect human creative expression, so copyright is unclear, so, they can be treated as trade secrets under the EU [Trade Secrets Directive](#) (use confidential marking, access controls, NDAs). If you distribute weights, you can use licenses that restrict certain uses (for example, [OpenRAIL-M](#)); note that software law still preserves limited decompilation rights for interoperability (see Art. 6 of the Software Directive [2009/24/EC](#)).

When it comes to prompts, short prompts are usually too small to be protected; a carefully crafted prompt may be protected if it shows the author’s “own intellectual creation” (CJEU standard from [Infopaq](#)). Outputs follow the same rule: where a human selects/edits/curates the result, protection can arise; fully machine-generated material without meaningful human authorship is generally outside copyright (see the U.S. Copyright Office’s [AI guidance](#)). Make these rules explicit in your terms (e.g., [OpenAI Terms - Ownership of output](#)).

In many EU countries and the UK, authors retain moral rights: to be named and to object to derogatory treatment. Proper contributor agreements can be used to include appropriate waivers or irrevocable licenses so publication and downstream licensing stay smooth (see UK guidance on [moral rights](#)).

Action checklist

- Maintain a provenance register for every dataset, code module, and transformation.
- Use contributor agreements that transfer copyright and any database right - and include moral rights waivers.
- Mark model weights as confidential trade secrets; control access and logging.
- Draft clear terms and conditions on ownership and reuse of prompts and outputs, including re-training rights.
- Re-audit all licenses and rights just before each commercial release to catch research-only content that could block the product.

4.3. Keep Trade Secrets

A trade secret is knowledge that remains valuable only while it is not publicly known. For an AI project, this often includes unreleased datasets, data-cleaning scripts, hyper-parameter schedules, training pipelines, and the final model weights. To gain legal protection under the EU Trade Secrets Directive (and similar rules worldwide) you must show that the information is (i) commercially valuable, (ii) reasonably protected, and (iii) clearly identified as confidential.

Mind the difference between Patent and Trade-Secret: Choose a patent when the invention is visible - or easy to copy - once you publish a paper, present at a conference, or ship a product. A patent gives up to 20 years of exclusive rights and clear enforcement power, but you must file *before* any public disclosure, or the idea loses novelty. Use a trade secret approach when the value sits in information that stays hidden behind an application program interphase (API) or on secure servers, such as training pipelines or tuned model weights; protection starts immediately, can last forever, and costs less, yet it survives only while strict NDAs, access controls, and audit logs are in place. In practice: if rivals can reverse-engineer the technology from your release, patent early; if they cannot see or extract it, keep it confidential and invest in solid secret-management routines.

The measures below strike a balance: strong enough to convince a court that “reasonable steps” were taken, yet light enough not to disturb collaboration or day-to-day research.

Action checklist

- Keep a spreadsheet or vault record that lists every secret item, its holder, its business value, and review dates; update it after each new training run or data clean-up.
- Use layered access control – combine role-based permissions in the code repo with hardware keys or secure enclaves for the production weights; grant evaluation access through time-limited demo licenses.
- Require partners, cloud providers, and beta testers to apply *equivalent* controls to the trade secrets and allow audit of their logs; add liquidated damages for leaks.
- Ship inference modules as containers with encrypted weights or compile-time obfuscation; watermark models (e.g. trigger phrases, parameter tagging) so misappropriation can be traced.
- Ensure clear marking - label files “CONFIDENTIAL - TRADE SECRET” and run employee training so staff know what may be shared and what must stay inside the project.
- Build automated off-boarding steps: revoke credentials, collect hardware tokens, confirm secure deletion, and sign off in a checklist.
- Schedule periodic housekeeping - six- or twelve-month reviews to reassess the commercial value of each secret and retire those that no longer need protection.

4.4. Choose Open-Source or Hybrid Licensing

Open-source licenses can let you share code or models to speed up adoption yet still shape how downstream users may exploit them. [MIT](#) and [Apache 2.0](#) are *permissive*: anyone may copy, modify, and ship your work - even inside closed products - provided they keep the copyright notice (and, for Apache 2.0, the patent grant and NOTICE file). [GPL v3](#) is *copyleft*:

if someone distributes a product containing GPL code, the entire combined work must be released under the GPL, forcing disclosure of source code and derived model weights. Before releasing, check that every upstream component is license-compatible; a single GPL dependency can “infect” the whole stack and block proprietary commercialization.

For more detailed consideration of the Open Licenses see “Guidelines and Considerations for Open Science” as part of this IMPAC3T-IP Toolbox.

If you need both research and commercial revenue, adopt a dual-license or hybrid model. Typical pattern: publish under GPL (or a [Responsible AI License](#), RAIL) for academic use and offer a paid proprietary license that waives copyleft or use-restrictions for industry customers. This keeps community contributions flowing while generating income and letting corporate partners embed the technology into closed systems.

[RAIL](#) clauses add ethical use restrictions (e.g. “no facial recognition for surveillance”), plugging a gap in classic OSS terms that stay silent on downstream purposes. Because such field-of-use limits break with Open-Source Initiative definitions, code under RAIL is not true “open source”; list the license clearly to avoid confusion and confirm compliance monitoring duties in your contracts.

Action checklist

- Map all inbound licenses and confirm compatibility (e.g. MIT + Apache 2.0 - fine; GPL + proprietary - no).
- Decide the outbound strategy: permissive for visibility, copyleft for reciprocity, hybrid for revenue.
- Add a contributor-license agreement so rights flow back, and dual licensing stays legally clean.
- State attribution, patent, and ethical-use terms in a top-level LICENCE and NOTICE file.
- Document how license obligations (e.g. providing source code, passing notices) will be met when the product ships.

4.4.1. Open Ecosystem primer

Modern AI stacks mix Open Data (often under Creative Commons), Open-source software, and “open-weight” models. Treat each layer separately:

- Open Data – CC BY/CC BY-SA require attribution/Share-Alike; CC BY-NC bars for commercial use.
- Open-source software – permissive (MIT/Apache-2.0) vs. copyleft (GPL/AGPL/LGPL) duties; check SBOM before shipping.
- Open weights / Open models – license terms vary (e.g., Llama Community License vs. Apache-2.0 vs. OpenRAIL-M). Some allow commercial use; others restrict fields of use.

AI Act interplay

- Open-source AI models may benefit from transparency reliefs in some contexts, but deployers still carry out duties under the AI Act.
- SMEs – the Act provides lighter support measures and access to regulatory/legal sandboxes; use them to de-risk early pilots.

Action checklist

- | |
|--|
| <ul style="list-style-type: none"> <input type="checkbox"/> Map “open” inputs in a single SBOM/SBDOM and align outbound license choices. <input type="checkbox"/> Pass through attribution/Share-Alike in customer docs where required. <input type="checkbox"/> Record whether “open” status affects your AI-Act transparency strategy and model-card content. |
|--|

4.5. Consider Defensive Publishing

Sometimes the best way to protect freedom to operate is not to patent but to publish the idea first, turning it into **Prior Art** that blocks anyone else from claiming exclusive rights. Such defensive publishing suits AI features that are easy to replicate, too expensive to patent in every market, or offer only a short-term edge. By placing a clear, **time-stamped** description on a public platform - [arXiv](#), [Zenodo](#), a peer-reviewed journal, or the company blog - you put examiners and rivals on notice: the concept is already in the public domain, so any later patent application lacks novelty.

Action checklist

- | |
|---|
| <ul style="list-style-type: none"> <input type="checkbox"/> Check internal patent plans first. Once you publish, you lose the chance to file a valid patent in most jurisdictions (see absolute-novelty rule). <input type="checkbox"/> Disclose enough detail. Explain the technical problem, the algorithm or workflow that solves it, and at least one concrete implementation example; vague marketing copy will not stand as prior art. <input type="checkbox"/> Add searchable metadata. Use clear titles, keywords, and DOIs so patent examiners can find the document during prior-art searches. <input type="checkbox"/> Timestamp and archive. Post on platforms that provide immutable records and keep a signed PDF in your IP folder. <input type="checkbox"/> Update when improved. If you iterate the method, publish brief addenda so the public disclosure stays ahead of would-be filers. <input type="checkbox"/> Tell partners and investors. Note the defensive publication in the IP register, so due-diligence teams understand why no patent was filed and how freedom to operate is preserved. |
|---|

5. Navigate EU Rules and Act on Ethics

Besides the EU AI Act discussed in Section [Know the EU AI Product Landscape](#), there are other regulations you need to follow.

5.1. GDPR + ePrivacy

Under EU law, personal data and electronic communications are protected by two complementary frameworks:

GDPR (General Data Protection Regulation)

- Scope: Any information relating to an identified or identifiable person (data subject), including names, IP addresses, or behavioral profiles created by AI.
- Lawful basis: You must pick and document one basis for each processing activity (e.g. user consent, contract performance, legitimate interest). Research-only consent does not cover commercial use - plan early if you intend to monetize.
- Data Protection Impact Assessment ([DPIA](#)): Mandatory for high-risk AI (profiling, automated decisions with legal or similarly significant effects). A DPIA must describe the process, assess necessity and proportionality, identify risks, and list mitigating measures.
- Data subject rights: Provide transparent notices (Articles 13–14), and enable rights to access, correction, erasure, restriction, portability, and objection to automated decision-making (Article 22).
- Records & governance: Maintain a Record of Processing Activities ([Roopa](#)), implement “privacy by design” (minimize data collection, pseudonymize where possible), appoint a Data Protection Officer if required, and report breaches within 72 hours.

[ePrivacy Directive](#) (and forthcoming Regulation)

- Electronic communications confidentiality: Interception or surveillance of communications metadata (call logs, IP traffic) is forbidden without user consent or strong legal basis.
- Cookies & trackers: Any non-strictly-necessary cookies (e.g. analytics, personalization) require clear prior consent from the user; consent must be freely given, specific, informed, and revocable.
- Direct marketing: Unsolicited communications by email, SMS, or automated calls need opt-in consent; existing-customer rules may allow some contacts but must still offer an easy opt-out.

Note that GDPR regulations are a large topic, and consulting with more dedicated materials or certified professional engagement is recommended.

Automated decision-making and profiling

Separate your standard rights handling (access, erasure, portability) from [Article 22](#) GDPR cases where decisions are based solely on automated processing with legal or similarly significant effects. For such ADM (Automated decision-making), provide a human-review channel and explain logic, significance, and envisaged consequences in plain language.

Action checklist

- Map all personal data flows and communications metadata in your system.
- Choose and document a lawful basis for each data use, separating research from commercial tracks.
- Carry out a DPIA for profiling or automated decision of AI and adopt its risk-mitigation measures.
- Update privacy notices to cover automated decision-making and data subject rights.
- Implement a consent mechanism for cookies and in-app trackers; log and manage consent records.
- Establish processes for data-subject requests (access, erasure, portability, objection).
- Encrypt personal data in transit and at rest; apply pseudonymization to minimize exposure.
- Sign detailed data-processing agreements with any processors; ensure they meet GDPR and ePrivacy standards.
- Appoint or consult a [Data Protection Officer](#) and set up breach-reporting procedures.
- Review and refresh these measures at least annually or when you add new AI features.

5.2. Respect Data Sovereignty and Access

Data sovereignty means keeping information under the laws of the country or region where it is stored or processed. Inside the EU, this starts with GDPR and the new Data Governance Act and EU Data Act. They let people and firms insist that personal or industrial data stay in the European Economic Area (EEA) or move abroad only with strong safeguards. Several Member States add tighter sector rules - France demands Healthcare Data certified hosting (HDS-certified hosting) for health data, Germany's public sector requires 'Bundesamt für Sicherheit in der Informationstechnik Cloud Computing Criteria Catalogue Compliance' (BSI C5-compliant clouds) with local key control, and Spain's banking regulator wants critical data in EU data centers. Defense contracts often impose their own localization clauses.

To handle cross-border transfers you need a written plan:

- Map every location where raw data, backups, logs, and model weights sit.
- Choose a legal ground: use GDPR Standard Contractual Clauses, Binding Corporate Rules, or (for the US) the EU-US Data Privacy Framework. Add the UK IDTA if you send data to or from Britain.

- Assess foreign-law risks such as access under the US CLOUD Act or China’s PIPL and record them in a short memo.
- Set fallback rules so you can move or delete data within 30 days if a destination loses “adequacy” status.

Technical choices make the legal plan real:

- Pick an EU-hosted cloud (for example, AWS Europe-Central (Frankfurt), Azure EU Data Boundary, Google Cloud EU-West), or fully European providers (for example, OVHcloud or T-Systems). If possible, prefer “sovereign-cloud” offers that keep admin keys inside the EEA and follow [GAIA-X](#) initiative (initiative to develop a plan to potentially develop a federated secure [data infrastructure](#) for [Europe](#)) or the coming EUCS label (European Union Cloud Cybersecurity Certification Scheme, a proposed voluntary certification for cloud services designed to harmonize security across the EU).
- Encrypt in transit and at rest with customer-managed keys stored in an EU hardware-security module; enable geo-fencing so data cannot leave approved regions.
- Segment workloads: keep sensitive tables in EU-only virtual networks or on-prem edge nodes, push only anonymized features to global storage.
- Automate exit scripts to back up, checksum, redeploy, and securely wipe data if you must switch regions or providers.

Finally, write all legal clauses and technical controls into supplier and license agreements, and attach a two-page “technical-and-organizational measures” annex. This joined-up approach shows regulators, customers, and investors exactly how the project keeps data under European jurisdiction and ready for compliant scaling.

5.3. Embed Responsible-AI Practices

Responsible AI means designing, building, and using AI systems in a way that is fair, transparent, safe, and socially beneficial. Key dimensions are:

- **Bias** & **Fairness**
AI models can learn unfair patterns from biased data and then make decisions that hurt certain groups. Regulations like the EU’s Equality Directives and the upcoming AI Act require you to assess and mitigate bias. Best practices include using fairness toolkits (for example IBM’s AI Fairness 360), testing models on balanced datasets, and documenting results in a “bias audit” report.
- **Explainability** & **Transparency**
Users and regulators must understand how AI reaches its conclusions. The EU AI Act demands “appropriate transparency” for high-risk systems. You can meet this by publishing Model cards or Datasheets for datasets (if possible), providing user-friendly explanations in your interface, and keeping logs of key decisions for auditors.
- **Safety & Robustness**
AI must behave reliably under normal and unexpected conditions. Standards such as ISO/IEC TR 24028 (trustworthiness) and the UK’s BS 8611 guide on ethical risk both

cover robustness. Techniques include adversarial testing, red-team exercises, and fail-safe design (for instance, defaulting to a safe output when confidence is low).

- **Societal & Environmental Impact**
Think about how your AI affects jobs, privacy, and the environment. The OECD AI Principles and UNESCO's Recommendation on the Ethics of AI call for human oversight and sustainability. You should carry out an impact assessment that considers data privacy, energy use, and downstream effects on communities.

Minimum contents for transparency artefacts

- Model Card – model name/version; intended/Out-of-scope uses; training data summary; metrics (incl. bias/robustness); evaluation protocol; known limitations; contact & update policy.
- Data Card – motivation & ownership; collection period & method (incl. TDM status); composition; preprocessing/labeling; license/consents; allowed uses; distribution & maintenance.

Voluntary Agreements & Codes

- [EU Code of Conduct on AI](#) – a practical guidance for trustworthy AI, which is in development.
- [IEEE Ethically Aligned Design](#) – a global framework for ethical priorities.
- [Partnership on AI](#) – multi-stakeholder best-practice sharing.

Action checklist

- Run a bias impact assessment on your training and test sets, document findings and mitigation steps.
- Create model cards and datasheets explaining data sources, intended use, and known limitations.
- Perform adversarial and stress testing, record failure modes, and design fallback plans.
- Conduct a Responsible-AI impact assessment covering social, privacy, and environmental risks.
- Adopt a voluntary code (EU Code of Conduct or IEEE guideline) and assign a senior lead to oversee compliance.
- Schedule periodic reviews (typically every 6–12 months) to re-evaluate fairness, transparency, and safety as models evolve.

5.4. Be ready for Real-world risk challenges

Some risks cannot be fully eliminated and must be managed transparently: [residual bias](#) in complex data, limited explainability of deep models ([XAI](#)), trade-offs between precision and interpretability, scraping exposure when curating datasets, and security risks tied to open components.

Action checklist

- Document the chosen trade-offs, publish user-facing guidance, and include measurable guard-rails (confidence thresholds, fallback behaviors, human-in-the-loop).

6. Choose and Structure Your Licensing Deal

Turning a research-grade AI asset into a sustainable product is more than a technical exercise; it is a business design task. The right commercialization and licensing model decides who may use the technology, under which terms, and how revenue flows back to the research team and institution. It must balance four factors:

- Market reach – how widely and quickly the solution needs to spread.
- Control & compliance – the level of oversight required to meet regulatory duties and protect reputation.
- Investment & support load – resources needed for hosting, updates and customer success.
- Return on investment – royalties, equity or service fees that fund further research.

This section maps the main routes - spin-out creation, out-licensing, Software-as-a-Service, evaluation agreements, dual licensing and consortium models - showing when each makes sense, the associated risks, and how to align them with the IP and compliance strategy laid out in earlier chapters. Subsequent subsections give practical checklists and term-sheet tips for each path.

6.1. Pick a License Archetype

Licensing is not one-size-fits-all. Below are the most common archetypes or models you can mix and adapt. Pick the lightest model that still meets your goals for reach, control and return.

- **Evaluation** / **research-only** **license**
Short, no-fee permit for internal tests, benchmarks or proofs of concept.
Pick when: you need feedback, demos or data from potential partners before a full deal.
- **Non-exclusive** **field** **license**
Many firms may use the technology, each inside a defined market, product line or territory.
Pick when: broad adoption is more valuable than tight control; good for niche tools sold to multiple OEMs (Original Equipment Manufacturers).
- **Exclusive license (full or field-limited)**
One partner gains sole rights - often with minimum-sales clauses and milestone payments.
Pick when: you want a single, committed investor to fund scale-up, or to avoid channel conflict.

- **Open-source dual license**
Core code released under a permissive license such as [Apache-2.0](#); commercial users must take a paid license for extra features or support.
Pick when: community adoption and peer review speed R&D, but you still need a revenue path.
- **Copyleft Community license**
Source shared under a reciprocal license such as [GNU GPL v3](#) or model weights under [CC BY 4.0](#).
Pick when: you want openness and downstream sharing to drive public benefits; less fit for proprietary addons.
- **Software-as-a-Service (SaaS) access license**
Users pay subscription fees for hosted access; the code and models stay on your cloud.
Pick when: you need usage data, quick updates, and strict IP control; higher support overhead.
- **Usage-based or royalty license**
Fees linked to API calls, seats, devices shipped or revenue share.
Pick when: customer volumes are hard to forecast, or you expect rapid growth but little upfront cash.
- **Joint development / co-ownership agreement**
Partner funds new features and gains shared IP rights; future licenses pay royalties to both parties.
Pick when: bespoke work or domain expertise is vital, and both sides bring key assets.

Use these patterns as building blocks. Combine clauses (e.g. evaluation → non-exclusive → exclusive) as the project matures but avoid over-engineering early deals - extra restrictions raise legal costs and slow adoption.

6.2. Draft the Key Clauses

A well-balanced license sets clear ground rules for both parties and prevents expensive surprises later. Below are the six clauses that deserve extra care in every AI product deal.

- **Scope of use**
spells out *who* may use the software or model, *where*, *how often* and for *which purposes*. Typical controls cover number of seats, cloud regions, sublicense rights, and whether the customer may retrain or fine-tune the model. Keep the clause tight enough to protect core IP, but not so narrow that it blocks legitimate deployment plans.
- **Sublicensing**
decides *if and when* the licensee can pass rights to affiliates, integrators or customers. If you allow sublicensing, require written approval, flow-down of all obligations, and transparent royalty reporting; silence means it is forbidden by default.

- **Update & maintenance rights**
define what upgrades, patches and new model versions the licensee receives, how fast they arrive, and whether extra fees apply. Link update delivery to the system's risk tier (e.g., security patches within 30 days for high-risk use) and set clear end-of-support dates to avoid disputes.
- **Service levels (SLA)**
for hosted or API access, commit to uptime, response times and data-recovery windows, plus remedies such as service credits when targets are missed. Align the SLA with any promises you made under the EU AI Act or NIS 2 so contractual and regulatory duties match.
- **Liability caps**
limit each party's financial exposure (often to a multiple of fees paid) while carving out unlimited liability for intentional misconduct, IP infringement or personal-injury claims. Recent litigation shows vendors can avoid nine-figure losses when caps are drafted clearly.
- **Indemnities**
the licensor normally defends and pays third-party claims that AI infringes patents, copyright or trade secrets. Narrow the duty to IP claims, reserve the right to replace or withdraw infringing components, and require prompt notice and cooperation from the customer. Decide whether indemnity costs sit inside or outside the liability cap.

Draft these clauses in plain language, tie up obligations to the product's risk level, and review them at every major release. Careful wording here protects revenue, reputation, and research budgets alike.

6.3. Check Upstream License Compatibility

Before you sign any deal, make sure the license you want to grant **does not clash with the licenses that cover the code, models and datasets you rely on**. A missed conflict can force a recall, block investment, or trigger legal claims.

Action checklist

- Build a clean Software & Data Bill of Materials (SBOM/SBDM). Run a scanner such as [SPDX-compatible](#) tools, [FOSSA](#) or [ScanCode](#) to list every direct and transitive component with its license ID. Track versions to catch later swaps and security fixes.
- Classify licenses by duty level - group them as permissive (e.g. MIT, Apache-2.0), weak-copyleft (LGPL), strong-copyleft (GPL-3.0), non-commercial or share-alike (CC BY-NC, CC BY-SA) so you see at a glance which ones can mix.
- Spot red-flag combos early. For example, GPL code pulled into a closed-source SaaS, CC BY-NC data used in a paid model, or CC BY-SA images embedded in a proprietary UI. Replace, re-license, or segregate the asset before launching.

- ❑ Map obligations to your outbound terms - if you must give attribution, share source, or include the license text, build these steps into your release script and customer docs. Check that your liability and indemnity clauses line up with any upstream disclaimers.
- ❑ Check jurisdiction and patent rules - some licenses (e.g. Apache-2.0) include patent grants; others (GPL-2.0) do not. Align this with your patent strategy and target markets.
- ❑ Review data privacy and consent - verify that personal data in training sets has valid consent for the commercial uses planned, and that any removal or opt-out duties are covered in your support playbook.
- ❑ Record decisions in a compatibility log - keep a short note of each conflict found, the fix applied, and the legal sign-off. Store it next to the SBOM/SBDOM for audits.
- ❑ Re-run the scan before each major release - new dependencies arrive through updates and retraining; automate the check in CI so surprises surface fast.

6.4. Plan Exit, Hand-back, and Escrow

Even the best license needs a clear “way out”. Good exit terms show customers how they can keep their systems running if the relationship ends, and they shield researchers from unexpected liabilities.

Action checklist

- ❑ Termination triggers - state cause (breach, non-payment, compliance failure) and *convenience* rights, with a 30- to 90-day cure window. See guidance in the UK’s [Digital Contracting Playbook](#) for wording tips.
- ❑ Wind-down license - let the customer run the last delivered version for a fixed period (e.g. 12 months) to avoid service breaks while migrating.
- ❑ Data-privacy wrap-up - return or delete personal data as promised; provide written confirmation and audit logs within 30 days.
- ❑ Hand-back package - on exit, deliver current source code, model weights, docs and SBOM/SBDOM; certify secure deletion of customer data in line with the GDPR “right to erasure” ([Article 17 GDPR](#)).
- ❑ Software escrow deposit - place the full build environment with an independent agent (for example, [ESCode](#) or [Iron Mountain](#)). Update the deposit at every major release. Trigger release if the vendor enters insolvency (see the [EU Insolvency Regulation 2015/848](#)), fails to give agreed support for > 60 days, or misses a critical security patch window. On release, grant the customer a non-exclusive, royalty-free license to maintain and patch internally.
- ❑ Survival clauses - keep confidentiality, IP ownership, liability caps and indemnities alive after termination; specify governing law and dispute venue.

7. Prepare for Investments

Moving an AI asset from lab to market needs two things at the same time: capital to grow and controls that protect that capital. This chapter helps choose a funding route while keeping technical, legal, and market risks within agreed limits.

Why this matters: investors release money only when they see a clear risk register and a credible plan to reduce those risks over time; regulators - especially under the coming EU AI Act - expect risk-based governance across the full lifecycle, not just at launch.

What counts as investment: you can mix seed grants, proof-of-concept funds, equity rounds, revenue-share deals, and non-dilutive options (some accelerators provide this). Each option has different timing, dilution, reporting, and control requirements, so pick the one that matches your stage and evidence.

What counts as risk management: identify, score, and treat risks in a simple, repeatable way, aligned with recognised frameworks (e.g., [ISO 31000](#)). In practice, track five buckets that investors care about most - technical, regulatory, IP, cyber-security, and market demand - and show how each one will move from “high/unknown” to “acceptable” with evidence.

How the pieces fit: tie capital milestones to risk-reduction gates. For example, completing a clinical validation study, passing a security audit, or clearing a freedom-to-operate review can trigger the next tranche. Add safety nets where they help: insurance for specific liabilities, escrow for critical code or weights, staged royalties in licenses, and covenant clauses that require ongoing compliance.

The subsections that follow provide short checklists and term-sheet tips so you can match an investor’s appetite with a disciplined, proportional risk-control plan.

7.1. Triage Due-Diligence Hot Spots

7.1 Due-diligence hot spots

Investors, corporate buyers and grant panels will dig into a small set of hot-spot issues before they release funds or sign a term sheet. Get these areas ready first, and you will speed up the deal, cut legal costs, and avoid last-minute price drops.

Action checklist

- IP chain-of-title – show signed assignments for every contributor, and patent filings or invention disclosures that match the roadmap.
- License compatibility – proves that upstream open-source and data licenses can live with your outbound terms; attach the compatibility log you built in earlier sections.
- Data provenance & privacy – provide data collection records, consent proofs, and a GDPR impact assessment; flag any datasets under opt-out or research-only terms.
- Regulatory risk tier – state the product’s classification under the EU AI Act or other sector rules (medical, financial), plus the gap-closing plan if still in pre-market testing.

- ❑ Model robustness & safety – share red-team reports and alignment scores (as described earlier), along with mitigation steps and pass-fail gates.
- ❑ Cyber-security posture – map controls to [ISO / IEC 27001](#) or equivalent; include the last external pen-test letter and vulnerability-management policy.
- ❑ Operational resilience – show your disaster-recovery runbook, the escrow deposit receipt, and cloud service-level records that meet [NIS 2](#) duties.
- ❑ Financial assumptions – connect usage metrics (API calls, GPU hours) to unit economics and cash-flow forecasts; stress-test the plan against a 30 % cost increase in compute or data fees.
- ❑ Team & governance – list key people and decision processes; include any Responsible-AI or ethics board charter.
- ❑ Environmental impact – provide a carbon-emissions estimate for model training and inference, following the [Green Software Foundation SCI](#).

7.2. Show Compliance in Valuation

Absence of compliance might result in severe cash burn and increase of time-to-market.

Action checklist

- ❑ Map compliance tier to direct costs - conformity assessment, notified-body fees and technical-file preparation for high-risk AI. Continuous log retention, incident reporting and patch SLAs under NIS 2 and [Cyber Resilience Act](#) (CRA).
- ❑ Translate delays into discount - each extra quarter to secure certification pushes revenue out and raises the investor's discount rate; factor this when setting milestones and pre-money valuations.
- ❑ Quantify capitalized OPEX - ongoing audits, security tooling and privacy staff add recurring spend that lowers EBITDA multiples; show how automation or shared platforms keep these costs flat as revenue grows.
- ❑ Model "tail-risk" scenarios - assign probability and impact to a GDPR fine or AI-Act market withdrawal; investors will haircut valuations if the mitigation fund looks thin.
- ❑ Leverage compliance as an asset - proving early alignment with ISO 27001 and the Eu AI Act can justify premium pricing or faster enterprise sales cycles - add that upside back into the valuation story.
- ❑ Use staged funding to match risk burn-down - link tranche releases to passing key audits (e.g., security penetration test, EU AIAct conformity review) so investors see risk reduced in lock-step with their capital outlay.
- ❑ Track accrued certification fees, escrow deposits and insurance premiums.

7.3. Apply Risk Mitigation Measures

Good risk management is about reducing the chance and the cost of something going wrong. Follow the principles in [ISO 31000](#) and map every action to the risk tier you declared under the [EU AI Act](#).

Action checklist

- Rescope the product early – drop or gate features that push the system into a higher-risk category.
- Add clear usage limits, liability caps and export-control notices so downstream users cannot create unplanned legal exposure.
- Transfer residual risk with insurance – take out cyber-security and IP-infringement cover.
- Build a contingency buffer – reserve part of the raise for recalls, breach of notifications or emergency model retraining.
- Keep an incident-response playbook ready – align with [NIS 2](#) 24-hour reporting; rehearse twice a year and log lessons learned.
- Monitor, patch, repeat – track vulnerabilities in third-party code, apply security patches within 30 days, and re-scan the SBOM/SBDOM before every major release to stop risk creeping back in.
- Review and refresh** – update this playbook whenever new regulations appear.

8. Operate, Monitor, and Improve Post-Launch

Launching an AI product starts with a regulated service life. From that day forward you must keep the system safe, lawful and secure under evolving rules such as the EU AI Act, NIS 2 and CRA. This section shows how to:

- track performance and safety drift;
- meet short incident-reporting deadlines;
- refresh approvals after major updates; and
- retire a system responsibly at end-of-life.

The following sections give only the new routines and artefacts you need after first release, building on the technical, security and risk frameworks you set up in earlier chapters.

8.1. Watch for Model Drift and Re-train

Once the model is live, you need a tight loop that spots drift, decides when to re-train, and records every new version. The practices below meet the continuous-monitoring duty in the EU AI Act while following [NIST](#) and [MLOps](#) guidance.

Action checklist

- Set a performance baseline and alert thresholds. Log the validation metrics of the launch model and agree on statistical thresholds that will trigger investigation.

- ❑ Automated drift detection. Use a monitoring stack or cloud service that compares live inputs and outputs with the baseline. Fine-tune alert thresholds to cut false positives.
- ❑ Define re-training triggers. Retrain when any of these occur:
 - ❑ drift alert persists for > N windows;
 - ❑ model KPIs fall below the service-level floor in the SLA;
 - ❑ major data-generating process change (new sensor, taxonomy, regulation);
 - ❑ scheduled cadence (quarterly or after M new samples) for low-volume domains. Continuous monitoring and periodic re-validation are core requirements for high-risk AI under the EU AI Act.
- ❑ Track data, code, and model versions together. Store large artefacts in a tool, such as [DVC](#) so each Git commit points to the exact dataset and weight file.
- ❑ Gate new versions. Before deployment, run the full safety and robustness suite (see earlier chapters) and update the technical file. If the model's intended purpose or risk tier changes, redo the conformity assessment before release.
- ❑ Log and archive. Keep drift metrics, retrain decisions and new weight hashes for ≥ 10 years (high-risk systems) to meet the AI Act's record-keeping duty and to support audits.

8.2. Handle Incidents and Recalls

When things go wrong, speed and transparency keep users safe and protect your license revenue. EU law now sets hard clocks:

- High-risk AI – report any “serious incident” to the national market-surveillance authority *within 15 days*, or *2 days* for widespread harm, under [AI Act Art. 73](#).
- Essential / important entities – send an *early warning in 24 h* and a full incident notice in 72 h per [NIS 2 Art. 23](#).
- Personal-data breaches – notify the DPA in $\leq 72 h$ (and affected persons “without undue delay”) under [GDPR Art. 33–34](#).

Practical playbook

- Detect & triage - *Route every critical alert through your 24 × 7 on-call rotations*. Apply the severity matrix from ISO/IEC 27035 or [ENISA](#)'s cyber-crisis guide to decide whether the clock starts.
- Contain & collect evidence - Isolate affected components, snapshot logs, and keep hash-signed copies for forensic review. Follow the steps in [ISO/IEC 30111](#) to avoid destroying evidence while patching.
- Communicate openly - Send a plain-language advisory to users explaining the issue, workaround, and patch path. Coordinate public statements with the authority if the incident is likely to be public.

- Post-mortem & preventative action - Within 30 days complete a root-cause analysis, log in to the risk register, and schedule red-team tests to verify that similar vectors are closed.

8.3. Plan for Sunsets and Transitions

Build a well-planned sunset to avoid security gaps, stranded users, and legal fines. Build the steps below into your product's roadmap, so the end-of-life (EOL) is a routine rollout, not an emergency scramble.

Core steps for all AI systems

- Declare an EOL window up front. Put the planned end-of-support date in the license or SLA and review it after every major release.
- Freeze and archive the final build. Lock source, weights, SBOM and docs; store a signed snapshot for the required retention period.
- Give users at least six months' notice. Share a migration guide, last-patch timetable, and help-desk contact.
- Erase or return data safely. Follow the wipe methods in [NIST SP 800-88 Rev. 1](#).
- Release the escrow deposit or fallback repo. Hand over build assets (see § 6.4) so customers can self-maintain after EOL.

Extra steps for high-risk AI under the EU AI Act

Follow this EU-wide deregistration checklist when you sunset a system that has been registered as high-risk:

1. Confirm classification - Check if the use-case is in Annex III or captured by the rules in [Article 6](#). If a self-assessment under Art 6 §3 shows *no significant risk*, keep that memo and stop here.
2. Update the EU High-Risk AI Database - Change your entry to Status = "withdrawn" as required by [Article 49](#) and Annex VIII; the database itself is set up under [Article 71](#). Save the receipt ID.
3. Archive & retain evidence - Keep the full technical file for 10 years ([Article 18](#)) and any automatically generated logs for ≥ 6 months while you control them ([Article 19](#)).
4. Stay available for questions - Be ready to answer competent-authority queries during the same 10-year window, as required by [Article 21](#).
5. Bundle the paperwork - Store together: the risk-assessment memo, final database receipt, any Article 20 notices, and your retention schedule.

Timeline These deregistration duties start to apply on **2 August 2026**, 24 months after the Act's entry into force. See the official [implementation schedule](#).

9. Annex 1: Regulatory documents and useful guides

A. Regulatory documents, standards & official references

- **Official AI Act in EUR-Lex** – authoritative OJ version, multilingual.
<https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
 - **EU AI Act portal** – full, searchable text of the Regulation plus article-by-article navigation.
<https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
 - **Art. 6 AI Act** – legal test for deciding whether an AI system is high-risk.
https://eur-lex.europa.eu/eli/reg/2024/1689/oj#art_6
 - **Art. 18 AI Act** – ten-year retention requirement for technical documentation.
https://eur-lex.europa.eu/eli/reg/2024/1689/oj#art_18
 - **Art. 19 AI Act** – minimum six-month log-keeping duty for high-risk systems.
https://eur-lex.europa.eu/eli/reg/2024/1689/oj#art_19
 - **Art. 20 AI Act** – obligations to take corrective action and notify deployers or authorities.
https://eur-lex.europa.eu/eli/reg/2024/1689/oj#art_20
 - **Art. 21 AI Act** – co-operation duties with market-surveillance authorities.
https://eur-lex.europa.eu/eli/reg/2024/1689/oj#art_21
 - **Art. 49 AI Act** – mandatory registration of high-risk AI in the EU database.
https://eur-lex.europa.eu/eli/reg/2024/1689/oj#art_49
 - **Art. 71 AI Act** – legal basis for the EU High-Risk AI Database itself.
https://eur-lex.europa.eu/eli/reg/2024/1689/oj#art_71
 - **Art. 72 & 73 AI Act** – post-market monitoring and serious incident reporting.
https://eur-lex.europa.eu/eli/reg/2024/1689/oj#art_72 • https://eur-lex.europa.eu/eli/reg/2024/1689/oj#art_73
 - **Art. 79 AI Act** – rules on withdrawal and recall of non-compliant AI products.
https://eur-lex.europa.eu/eli/reg/2024/1689/oj#art_79
 - **Annex III AI Act** – exhaustive list of use-cases automatically considered high-risk.
https://eur-lex.europa.eu/eli/reg/2024/1689/oj#anx_III
 - **AI Act implementation timeline** .
https://eur-lex.europa.eu/eli/reg/2024/1689/oj#art_113

- **GDPR consolidated text** – the 2016/679 Regulation with search & cross-links.
<https://gdpr-info.eu/>
 - **GDPR Art. 17** – the “right to erasure” (“right to be forgotten”).
<https://gdpr-info.eu/art-17-gdpr/>
 - **GDPR Arts. 33–34** – breach notification duties to authorities and data subjects.
<https://gdpr-info.eu/art-33-gdpr/>
 - **Irish DPC DPIA guide** – national authority’s practical handbook for data-protection impact assessments.
<http://www.dataprotection.ie/en/organisations/know-your-obligations/data-protection-impact-assessments>
- **ePrivacy Directive 2002/58/EC** – EU rules on cookies, traffic and location data.
<https://eur-lex.europa.eu/eli/dir/2002/58/oj/eng>
- **Data Governance Act page (EU)** – overview, factsheets and final legal text for 2022/868.
<https://digital-strategy.ec.europa.eu/en/policies/data-governance-act>
- **Data Act page (EU)** – resources on the forthcoming horizontal data-sharing Regulation.
<https://digital-strategy.ec.europa.eu/en/policies/data-act>
- **Cyber Resilience Act overview** – Commission explainer for the draft product-security Regulation.
<https://digital-strategy.ec.europa.eu/en/policies/cyber-resilience-act>
 - **Cyber Resilience Act proposal text** – COM/2022/454 in EUR-Lex.
<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52022PC0454>
- **NIS 2 Directive text** – official publication of Directive (EU) 2022/2555.
<https://eur-lex.europa.eu/eli/dir/2022/2555/oj>
 - **NIS 2 information hub** – scope, requirements and FAQs on the 2022 Directive.
<https://digital-strategy.ec.europa.eu/en/policies/nis-directive2>
 - **NIS 2 Article 23** – exact time limits for incident notification.
https://www.nis-2-directive.com/NIS_2_Directive_Article_23.html
- **DORA (Digital Operational Resilience Act)** – supervisory package for ICT risk in finance.
https://www.eiopa.europa.eu/digital-operational-resilience-act-dora_en
- **EU Medical Devices Regulation 2017/745** – rules for software as a medical device.
<https://eur-lex.europa.eu/eli/reg/2017/745/oj/eng>

- **UNECE vehicle regulations** – homologation and safety standards for automotive products.
<https://unece.org/transport/vehicle-regulations>
- **EU dual-use export-control guidance** – licensing rules for sensitive technologies.
https://policy.trade.ec.europa.eu/help-exporters-and-importers/exporting-dual-use-items_en
- **CE-marking explainer** – Commission page on conformity marking for products.
https://single-market-economy.ec.europa.eu/single-market/ce-marking_en
- **ISO/IEC 27001** – international standard for information security management systems.
<https://www.iso.org/standard/27001.html>
- **ISO 31000** – principles and framework for enterprise risk management.
<https://www.iso.org/standard/65694.html>
- **ISO/IEC 30111** – guidance on handling and disclosing software vulnerabilities.
<https://www.iso.org/standard/69725.html>
- **ENISA study: NIS investments 2024** – survey data on cybersecurity spending under NIS.
https://www.enisa.europa.eu/sites/default/files/2024-11/CSPA%20-%20NIS%20Investments%20-%202024_0.pdf
- **ENISA study: cyber-crisis management** – best-practice handbook for large-scale incidents.
<https://www.enisa.europa.eu/sites/default/files/2024-11/ENISA%20Study%20Best%20Practices%20Cyber%20Crisis%20Management.pdf>
- **Ethics Guidelines for Trustworthy AI** – non-binding EU high-level ethics framework (2019).
<https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>
- **EPO “absolute novelty” rule** – patentability guidance in the European Patent Convention.
https://www.epo.org/en/legal/guide-epc/2023/ga_c3_3_1.html
- **EPO news in focus for Artificial intelligence**
<https://www.epo.org/en/news-events/in-focus/ict/artificial-intelligence>

B. Examples, methodologies, licenses & practical guides

- **MIT License** – permissive open-source license allowing reuse with attribution.
<https://opensource.org/license/mit>
- **Apache 2.0 License** – permissive license with explicit patent grant.
<https://www.apache.org/licenses/LICENSE-2.0.txt>

- **GNU GPL v3** – copyleft license requiring derivative works to stay open-source.
<https://www.gnu.org/licenses/gpl-3.0.en.html>
- **Creative Commons BY 4.0** – lets others share/remix your work with attribution.
<https://creativecommons.org/licenses/by/4.0/>
- **RAIL (Responsible AI License)** – model license adding AI-specific usage restrictions.
<https://www.licenses.ai/>
- **Model Cards** – Google’s template for transparent AI model documentation.
<https://modelcards.withgoogle.com/>
- **Datasheets for Datasets** – Microsoft Research template for dataset provenance and ethics.
<https://www.microsoft.com/en-us/research/project/datasheets-for-datasets/>
- **Fairlearn** – open-source toolkit for fairness assessment and mitigation.
<https://fairlearn.org/>
- **OWASP Top 10** – community list of the most critical web-app security risks.
<https://owasp.org/www-project-top-ten/>
- **OWASP GenAI Red-Teaming Guide** – practical adversarial-testing playbook for generative AI.
<https://genai.owasp.org/>
- **CVE database** – searchable catalogue of publicly disclosed vulnerabilities.
<https://cve.mitre.org/>
- **NIST TEVV programme** – U.S. framework for AI test, evaluation, validation & verification.
<https://www.nist.gov/ai-test-evaluation-validation-and-verification-tevv>
- **OpenChain SBOM specification** – standard process for creating software bills of materials.
<https://www.openchainproject.org/>
- **SPDX license list** – machine-readable identifiers and metadata for software licenses.
<https://spdx.org/licenses/>
- **Semantic Versioning 2.0** – widely used rules for version numbers and change signaling.
<https://semver.org/>
- **DVC use-case: versioning data & models** – tutorial on managing large datasets in Git-style workflows.
<https://dvc.org/doc/use-cases/versioning-data-and-models>

- **Google Cloud Vertex AI best practices** – reference architecture for production ML on GCP.
<https://cloud.google.com/architecture/ml-on-gcp-best-practices>
- **BigQuery model-skew monitoring** – blog guide to detecting drift in deployed ML models.
<https://cloud.google.com/blog/products/data-analytics/monitor-ml-model-skew-and-drift-in-bigquery/>
- **AWS MLOps best practices** – overview of pipeline design and continuous delivery for ML.
<https://aws.amazon.com/what-is/mlops/>
- **NCSC cyber-insurance guidance** – UK national-security agency’s advice on buying cyber cover.
<https://www.ncsc.gov.uk/guidance/cyber-insurance-guidance>
- **CloudEagle SLA guide** – checklist of clauses to track in SaaS service-level agreements.
<https://www.cloudeagle.ai/blogs/service-level-agreements>
- **Law Insider software-maintenance clause** – real contract language example for maintenance obligations.
<https://www.lawinsider.com/clause/software-maintenance>
- **ContractsCounsel sublicense agreement template** – sample contract for sublicensing IP.
<https://www.contractsounsel.com/t/us/sublicense-agreement>
- **BLG practical guide to software licensing** – 40-page PDF covering models, audits and compliance.
https://www.blg.com/-/media/insights/documents/software-licensing_practical-guide_dec2022.pdf
- **Salesforce “What is SaaS?”** – plain-language explanation of the Software-as-a-Service model.
<https://www.salesforce.com/saas/>
- **Reverera software-license models** – industry blog detailing subscription, usage and value-based licenses.
<https://www.reverera.com/blog/software-monetization/software-licensing-models-types/>
- **Tidelift dual-licensing overview** – pros and cons of releasing OSS under two licenses.
<https://www.tidelift.com/blog/what-is-dual-licensing>
- **arXiv** – preprint repository for rapid dissemination of research papers.
<https://arxiv.org/>

- **Zenodo** – CERN-hosted platform for archiving datasets, code and grey literature with DOIs.
<https://zenodo.org/>
- **Defensive publication (Wikipedia)** – explains how disclosing an invention prevents later patents.
https://en.wikipedia.org/wiki/Defensive_publication
- **OWASP Multi-licensing article** – wiki page on distributing software under more than one license.
<https://en.wikipedia.org/wiki/Multi-licensing>
- **Linux Foundation Open Compliance Program** – best practices for OSS due-diligence processes.
<https://compliance.linuxfoundation.org/developers/process/>
- **WIRED analysis of 2024 CrowdStrike outage** – investigative piece on a cloud security failure's ripple effects.
<https://www.wired.com/story/crowdstrike-outage-microsoft-delta-lawsuits-analysis>
- **Reuters 2024 article on U.S. cloud-AI reporting rule** – news on proposed mandatory disclosures for advanced AI in the cloud.
<https://www.reuters.com/technology/us-proposes-requiring-reporting-advanced-ai-cloud-providers-2024-09-09/>

10. Annex 2: Popular AI models & their licenses

Model (latest major version)	Provider	License / Terms	Core permissions & limits
GPT-4o	OpenAI	OpenAI Model License (API only)	Use outputs commercially; no access to weights; rate-limited; must not break policy on disallowed content.
Gemini 1.5 Pro	Google DeepMind	Google Generative AI Terms	API/SaaS; no redistribution of model or weights; user owns outputs unless prohibited content.
Claude 3 Opus	Anthropic	Anthropic Terms of Service	API; outputs licensed to user; safety filter may refuse prompts; no reverse engineering.
Llama 3	Meta AI	Llama 3 Community License	Free for research & commercial use up to 700 million MAU; redistribution allowed with license notice; no harmful use.
Mistral 8×22B	Mistral AI	Apache 2.0	Fully permissive; weights and code may be modified and sold; must keep NOTICE & patent clause.
Mixtral 8×7B Instruct	Mistral AI	Apache 2.0	Same as above; popular for on-device inference.
Stable Diffusion 3	Stability AI	CreativeML-Open RAIL-M	Permissive plus ethical constraints (no sexual minors, political persuasion, etc.); must pass on RAIL notice.
DALL·E 3	OpenAI	OpenAI Model License	SaaS/API; user owns image outputs; policy bars deepfake or biometric misuse.
Whisper (speech-to-text)	OpenAI	MIT	Totally permissive; redistribution and commercial use allowed; credit required.
BERT-base	Google	Apache 2.0	Permissive; often embedded in search and chatbots.
T5 v1.1	Google	Apache 2.0	Permissive; suited for text-generation fine-tuning.

PaLM-2 API	Google	Google Generative AI Terms	API only; no weight access; usage limits and content policy.
BioGPT	Microsoft Research	MIT	Permissive; domain-specific (biomedical).
Code Llama	Meta AI	Llama Community License	Same rules as Llama 3; optimized for code completion.

11. Annex 3: Open-Source License Compatibility

(Row license → combined work distributed under Column license)

From\To	MIT	BSD-3	Apache-2.0	GPLv2	GPLv3	LGPLv2.1	LGPLv3	MPL-2.0	AGPLv3	CC0
MIT	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
BSD-3-Clause	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Apache-2.0	X	X	✓	X	✓	X	✓	~	✓	X
GPLv2	X	X	X	✓	X	X	X	X	X	X
GPLv3	X	X	X	X	✓	X	X	X	✓	X
LGPL v2.1	X	X	X	✓	X	✓	X	X	X	X
LGPL v3	X	X	X	X	✓	X	✓	X	✓	X
MPL-2.0	X	X	X	✓	✓	✓	✓	✓	✓	X
AGPLv3	X	X	X	X	X	X	X	X	✓	X
CC0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Legend

- ✓ compatible – Code under the row license can be combined and redistributed under the column license with no extra conditions.
- ~ conditional / dual-license needed – Possible, but only if you also keep the row license (e.g. dual-licensing Apache-2.0 with MPL-2.0).
- X incompatible – Cannot redistribute under the column license without removing or relicensing the row-license code.

Compatibility with Proprietary Products

License	Proprietary use – internal or SaaS (no code distribution)	Proprietary binary distribution (closed-source installables / firmware)
MIT	✓	✓
BSD-3-Clause	✓	✓
Apache-2.0	✓	✓ ¹
GPL v2	✓	✗
GPL v3	✓	✗
LGPL v2.1	✓	~ ²
LGPL v3	✓	~ ²
MPL-2.0	✓	~ ³
AGPL v3	~ ⁴	✗
CC0	✓	✓

Legend

- ✓ **compatible** – You can use the OSS code in this proprietary scenario with only the normal notice obligations.
- ~ **conditional** – Possible, but extra terms apply (e.g. supplying object files, keeping copylefted files open).
- ✗ **incompatible** – Combining the code and distributing the proprietary product is impossible without relicensing or open-sourcing.

Footnotes on the conditions

1. **Apache-2.0** → **Proprietary binary** – Allowed, but you must keep the NOTICE file and patent grant text with your product docs.
2. **LGPL (both versions)** – Closed binaries are fine *if* you **dynamically link** or provide object files that let users relink the LGPL library; otherwise, you must open-source the whole work.
3. **MPL-2.0** – You may keep your proprietary code closed, but any files containing MPL-licensed code (and your changes to them) must be shipped in source form.
4. **AGPL v3** → **SaaS / internal** – Merely running AGPL code privately is fine, but the moment end-users interact with the software over a network you must offer them the full source under the AGPL.

Key takeaway: Permissive licenses (MIT, BSD, Apache, CC0) integrate cleanly into proprietary products, whereas strong copyleft licenses (GPL, AGPL) never do for redistributed

binaries; weak/file-level copyleft (LGPL, MPL) sit in the middle and are workable with some engineering discipline.

12. Annex 4: Useful contacts

Contact	What they offer	Link
EU IP Helpdesk	Free first-line advice on patents, trademarks, licensing, and EU funding rules for SMEs and researchers.	EU IP Helpdesk
European Patent Office (EPO)	Patent search (Epicene), filing services, examiner guidelines, training on AI patentability.	EPO
Irish Patents Office (IPOI)	National filings, fee schedule, and search support in Ireland.	IPOI
German Patent & Trade Mark Office (DPMA)	National patents, utility models, and designs; AI examination practice notes.	DPMA
French INPI	Filing portal, prior-art searches, and start-up IP vouchers.	INPI
Spanish OEPM	Patent and trade-mark services plus regional innovation grants.	OEPM
UK Intellectual Property Office (UKIPO)	UK patents, designs, and the “AI patent guidelines” consultation papers.	UKIPO
WIPO Directory of IP Offices	Quick links to every other national office worldwide.	WIPO directory
Horizon Europe – Digital, Industry & Space	Large collaborative grants for advanced AI R&D and pilots.	Horizon Europe
EIC Accelerator	Equity + grant funding (€0.5–15 m) for deep-tech start-ups bringing AI to market.	EIC Accelerator
Digital Europe Programme	Supports data spaces, testing facilities, and skills projects for trustworthy AI.	Digital Europe
Innovate UK Smart Grants	Competitive funding for UK-based technology businesses, including AI.	Innovate UK

Enterprise Ireland Commercialization Fund	Grants to spin out AI technologies from Irish research teams.	Commercialization Fund
NANDO Notified-Body Database	List of accredited bodies that can perform AI-Act conformity assessments.	NANDO
IEEE Standards Association	Working groups and standards (e.g., BS 8611, ISO/IEC 24028) for ethical and safe AI.	IEEE SA
Partnership on AI	Multi-stakeholder forum sharing best practice on responsible AI deployment.	Partnership on AI
European Commission – AI Act Service Desk & Explorer	European Commission – AI Act Service Desk & Explorer	https://ai-act-service-desk.ec.europa.eu/en

info@impac3tip.eu

www.impac3tip.eu



Funded by
the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or HADEA. Neither the European Union nor the granting authority can be held responsible for them.