Chapter 1 - How We Got Here

You do not need this chapter to do the work. You need it to understand why the work looks the way it does.

Read it once. It explains where these rules came from, why Europe is further along, and what lands in 2027.

Europe Started First

Data protection as a fundamental right

The European Union treats data protection as a fundamental right in its own name — listed separately from privacy in the EU Charter.

That framing has a practical consequence. Europe regulates data protection **horizontally**, across every sector at once, rather than industry by industry.

The General Data Protection Regulation

The General Data Protection Regulation — GDPR — took effect in May 2018. It replaced a directive that had been in place since 1995.

It established machinery that European practitioners have now run for years:

- Documented legal basis for every processing activity
- Purpose limitation and data minimization
- Impact assessments before high-risk processing
- Rules governing automated decisions about people

Two provisions matter most for what follows:

PROVISION	REQUIREMENT
Article 22	Governs automated decisions with significant effects
Article 35	Requires an impact assessment before high-risk processing

Neither mentions artificial intelligence. Both apply to it.

The United Kingdom diverged

After leaving the EU, the UK kept an adapted version of the same regulation. It is enforced by the Information Commissioner's Office – the ICO.

- Similar to the EU version, not identical
- The differences have widened since 2020
- 2025 legislation revised the automated-decision provisions specifically

If your footprint covers both, you track two bodies of law.

The United States Went Sector by Sector

There is no comprehensive federal privacy law in the United States. There are sectoral laws instead.

LAW	COVERS
HIPAA – Health Insurance Portability and Accountability Act	Health information
GLBA – the Privacy Rule and Safeguards Rule	Financial institutions
FERPA – Family Educational Rights and Privacy Act	Education records
COPPA – Children's Online Privacy Protection Act	Children's data
FCRA – Fair Credit Reporting Act	Consumer reporting

Read those names in full and the pattern is obvious. Protection. Privacy. Rights. Safeguards.

All of it is data privacy and data protection law.

The consequence for practitioners: if your background is US compliance, your privacy experience is probably deep in one or two of these and unfamiliar in the rest. That is how the law is structured, not a gap in your training.

Then the States Borrowed From Europe

Beginning with California, US states began passing comprehensive privacy laws. They borrowed heavily from the European model.

Concepts that arrived in recognizably European form:

- Data protection assessments
- Sensitive data categories
- Profiling restrictions
- Purpose limitation
- Consumer rights and opt-outs

A US practitioner encountering a "data protection assessment" requirement in Colorado or Virginia is meeting something European colleagues have completed many times.

This is the bridge. The state privacy laws brought European concepts into US practice. The AI laws now build on those same concepts.

AI-Specific Laws Arrive

The United States moved first, in pieces

LAW	IN FORCE
NYC Local Law 144 – bias audits for hiring tools	2023
Utah AI Policy Act	May 2024
Illinois HB 3773 – AI employment discrimination	Jan 1, 2026
California AB 2013 and SB 53	Jan 1, 2026
Connecticut – private sector obligations	Oct 1, 2026

State laws vary considerably in scope:

- Some reach any organization deploying AI
- Others cover only government use
- Several were narrowed substantially during drafting

Texas is worth knowing as an example. Its Responsible AI Governance Act was cut back significantly before passage, and private-sector obligations are far lighter than the original draft suggested.

You cannot assume a state law applies. You cannot assume it doesn't.

The European Union passed the first comprehensive AI law

The EU AI Act was adopted in 2024. It phases in over several years.

The critical structural point: **it sits on top of the General Data Protection Regulation. It does not replace it.**

A single high-risk deployment in Europe can require both:

- A **Data Protection Impact Assessment** — DPIA — under GDPR Article 35
- A **Fundamental Rights Impact Assessment** — FRIA — under AI Act Article 27

One examines risk to data. The other examines risk to people.

They ask different questions. Neither substitutes for the other.

The Assessment Obligations Are Older Than People Think

One category deserves separating out, because organizations consistently misjudge it.

Most assessment requirements are already live.

REQUIREMENT	MANDATORY SINCE
Data Protection Impact Assessment — EU	May 2018
Data Protection Impact Assessment — UK	May 2018
Virginia data protection assessment	January 2023
Colorado data protection assessment	July 2023
Connecticut data protection assessment	July 2023
Texas data protection assessment	July 2024
California risk assessment requirements	January 2026
Fundamental Rights Impact Assessment	**December 2027**

Only the last one is future work — and it is the one attracting most of the attention.

A note on the UK date. The obligation did not begin at Brexit. It applied from May 2018 while the UK was an EU member state, and carried into UK GDPR when the transition period ended.

An organization planning for December 2027 while carrying an unmet obligation from 2018 has its attention in the wrong place.

Chapter 6 covers what these documents are and what goes into them.

Enforcement Is Already Happening

Not warnings. Decisions, with orders attached.

MATTER	SUBJECT AND PENALTY	WHAT THE REGULATOR ORDERED PRODUCED
Massachusetts – Earnest Operations	AI loan underwriting · 2025 · $2.5M	A written AI governance system. Documented fair-lending testing. Risk assessments. A named oversight team.
Hungary – Budapest Bank	Emotion analysis on calls · 2022 · HUF 250M	A completed impact assessment, a documented legal basis, and demonstrable safeguards – or stop the processing.
Italy – Foodinho	Algorithmic rider scoring · 2021 · €2.6M	Human intervention in algorithmic decisions, and periodic checks on the algorithm's accuracy.
Germany – Berlin bank	Automated credit decision · 2023 · €300K	The specific data, factors, and criteria behind a single automated decision, on request.

Different countries. Different sectors. Different years.

In each case the regulator asked for the same categories of material:

- The assessment
- The documented legal basis
- A record of what the system did and why

None of it could be produced quickly. None of it should be created after the request arrived.

Remediation is different. A regulator expects gaps to be addressed immediately, and starting that work is the right response.

But a document dated after the request answers a different question than the one being asked. The record shows what you did at the time. Remediation shows what you are doing now. Both matter. They are not interchangeable.

What Lands in 2027

Two dates.

January 1, 2027 — the US state cluster

At least ten state requirements take effect on this single date. They cover very different subject matter, which is the point worth noticing.

Automated decision-making

LAW	WHAT IT REQUIRES
Colorado SB 26-189	Notice, explanation, correction, and human review rights where automated technology materially influences consequential decisions
California — CCPA automated decision-making regulations	Risk assessments, pre-use notices, and consumer opt-outs for significant decisions

Colorado's applies to employment, housing, education, lending, insurance, healthcare, and essential government services. Deployers must provide consumer notices, keep compliance records for at least three years, support consumer rights, and offer meaningful human review after some adverse decisions.

The law was signed May 14, 2026, replacing the original after it was repealed. It removes the earlier duty-of-care standard and the annual impact assessment requirement. Developers get a 60-day cure period, which sunsets January 1, 2030.

Comprehensive privacy laws

STATE	NOTE
Louisiana	Applies at $25M revenue, or 75,000 consumers, or 50% of revenue from selling personal information
Oklahoma	New omnibus privacy law

Both follow the familiar consensus framework, with state-specific differences worth reading.

AI in healthcare and insurance

LAW	WHAT IT REQUIRES
Georgia SB 444	Coverage decisions may not be based solely on AI; a clinical peer must participate before an adverse determination

LAW	WHAT IT REQUIRES
Utah SB 319	Insurers must disclose AI use in authorization review; adverse determinations require independent medical judgment

Companion chatbots

LAW	WHAT IT REQUIRES
Washington HB 2225	Disclosure that the bot is not human, crisis protocols for suicidal ideation and self-harm, public reporting of crisis referrals
Oregon SB 1546	Protections for minors interacting with chatbot products

Washington's duty runs to all users, not only minors, and violations are unfair or deceptive acts carrying a private right of action.

Frontier models and content provenance

LAW	WHAT IT REQUIRES
New York RAISE Act	Safety and transparency obligations on large frontier AI developers
California AI Transparency Act – platform provisions	Large platforms must detect and label machine-readable provenance on AI-generated content

What that spread tells you

Ten laws. One date. And no single theme.

- Employment and lending decisions
- Health insurance authorization
- Consumer chatbots
- Frontier model safety
- Content labeling
- Two general privacy frameworks

You cannot scope this by watching one category of law. An organization tracking only "AI hiring rules" misses the insurance provisions. One tracking only privacy misses the chatbot duties.

The obligations attach based on what your tools actually do – not on which regulatory conversation you happen to be following.

December 2, 2027 – EU AI Act high-risk deployers

Obligations for deployers of standalone Annex III high-risk systems become applicable, following the deferral agreed in 2026.

Annex III categories include:

- Employment and workforce management
- Creditworthiness assessment
- Insurance risk assessment
- Education
- Access to essential services

Article 26 deployer obligations attach on that date. For certain deployer categories, Article 27 also requires a Fundamental Rights Impact Assessment before the system is put into use.

The Penalty Structure

INSTRUMENT	BREACH	MAXIMUM
EU AI Act	Prohibited practices	€35M or 7% of turnover
EU AI Act	Deployer and provider duties	€15M or 3% of turnover
GDPR	Serious violations	€20M or 4% of turnover
GDPR	Lesser violations, incl. missing DPIA	€10M or 2% of turnover

The two instruments are independent.

A single deployment can breach both. **The exposure is cumulative, not alternative.**

The Short Version

If you read nothing else in this chapter:

- Europe built the foundation, starting in 2018
- The United States regulated by sector, then by state

- State privacy laws imported the European concepts
- AI-specific laws build on those same concepts
- The EU AI Act sits on top of GDPR, not instead of it
- Enforcement has already begun in several jurisdictions

- At least ten state laws take effect January 1, 2027
- EU AI Act high-risk deployer duties follow on December 2, 2027

This is not a new discipline. It is data protection law, extended to a new technology, arriving faster than most organizations have built the practice to absorb it.

Chapter 2 covers where to start.

Sources

All statutory provisions should be read in their current form. Where a secondary source is cited below, it is named — verify against the primary text before relying on any specific requirement.

European law

- General Data Protection Regulation, Regulation (EU) 2016/679 — Articles 22 and 35
- EU Artificial Intelligence Act, Regulation (EU) 2024/1689 — Articles 26 and 27, with Annex III for the high-risk categories
- UK General Data Protection Regulation and the Data Protection Act 2018
- The December 2027 date for standalone Annex III high-risk deployer obligations reflects the deferral agreed in 2026. Confirm the current date before planning against it.

Penalty thresholds

Each figure in the penalty table comes from one of two provisions. Read both before quoting the numbers.

THRESHOLD	PROVISION
€35M or 7% — prohibited practices	EU AI Act, Article 99
€15M or 3% — deployer and provider obligations	EU AI Act, Article 99
€20M or 4% — serious violations	GDPR, Article 83
€10M or 2% — lesser violations, including a missing assessment	GDPR, Article 83

Both provisions set maximums rather than fixed amounts, and both are expressed as the higher of a fixed sum or a percentage of worldwide annual turnover. The tiers within each Article are more granular than the table above; check which tier applies to a specific breach rather than assuming the headline figure.

United States – laws effective January 1, 2027

LAW	VERIFICATION STATUS
Colorado SB 26-189	Confirmed against the Colorado General Assembly record. Passed May 9, 2026; signed May 14, 2026. See litigation note below.
California – CCPA automated decision-making regulations	Confirmed. Related risk assessment requirements took effect January 1, 2026.
California AI Transparency Act – platform provisions	Confirmed.
New York RAISE Act	Confirmed.
Louisiana Data Privacy Act	Confirmed.
Oklahoma comprehensive privacy law	Confirmed.
Georgia SB 444	Confirmed against the Georgia legislative record.
Utah SB 319	Confirmed.
Washington HB 2225	Confirmed against the Washington State Legislature record.
Oregon SB 1546	Confirmed against the Oregon legislative record. Or. Laws 2026 ch. 85.

United States – laws already in effect

- New York City Local Law 144 of 2021 – automated employment decision tools
- Utah Artificial Intelligence Policy Act
- Illinois HB 3773 – amending the Illinois Human Rights Act
- California AB 2013 and SB 53
- Connecticut – private sector AI obligations, staged from October 1, 2026
- Texas HB 149, Responsible Artificial Intelligence Governance Act

Assessment effective dates

- Data Protection Impact Assessment – GDPR Article 35, applicable from May 25, 2018 in both the EU and the UK
- ~~Virginia Consumer Data Protection Act – effective January 1, 2023~~
- Colorado Privacy Act – effective July 1, 2023

- Connecticut Data Privacy Act — effective July 1, 2023
- Texas Data Privacy and Security Act — effective July 1, 2024
- California risk assessment requirements — effective January 1, 2026
- Fundamental Rights Impact Assessment — EU AI Act Article 27

Enforcement decisions referenced

MATTER	CITATION
Massachusetts — AI loan underwriting, 2025, $2.5M	Massachusetts Office of the Attorney General, *In re Earnest Operations LLC*, Assurance of Discontinuance, July 10, 2025
Hungary — emotion analysis on customer service calls, 2022, HUF 250M	Hungarian National Authority for Data Protection and Freedom of Information (NAIH), Decision NAIH-85-3/2022 (Budapest Bank)
Italy — algorithmic rider scoring, 2021, €2.6M	Garante per la protezione dei dati personali, order no. 234/2021 (Foodinho)
Germany — automated credit decision, Berlin, 2023, €300K	Berlin Commissioner for Data Protection and Freedom of Information (BlnBDI), administrative fine, May 31, 2023 — GDPR Articles 5(1)(a), 15(1)(h), 22(3)

These are publicly reported enforcement actions, described here for information only. Nothing in this chapter is a prediction of any outcome for any organization.

A note on Colorado

The Colorado framework has been unusually unstable, and a reader planning against the January 2027 date should know why.

- The original law, SB 24-205, was enacted in 2024 with an effective date later extended
- A federal court enjoined its enforcement in April 2026 in *xAI v. Weiser*
- The legislature then repealed and reenacted it as SB 26-189
- The replacement narrows the statute substantially — the duty-of-care standard, risk management programs, and annual impact assessments from the original are gone
- The replacement takes effect January 1, 2027, and applies to decisions made on or after that date

Litigation over the framework was ongoing at the time of writing. Confirm the current position before planning around this date specifically.

A standing note on verification

Effective dates move. Colorado's original AI law was repealed and replaced before it took effect. The EU AI Act's high-risk deployer date was deferred by more than a year. Nothing in this chapter should be treated as settled

without checking the current position.

Use and reuse

Free to use, share, and adapt within your organization. Attribution appreciated, not required. No warranty of any kind — this is practitioner guidance, not legal advice, and sufficiency is determined by your counsel.

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Chapter 2 - Where to Begin

You have obligations. Chapter 1 covered where they came from.

This chapter covers where to start, and it begins with better news than most people expect: **you are not starting from nothing.**

This chapter covers:

- Who carries the obligation
- Why this program has a different shape
- What you already have that applies
- Who else in your organization has a stake
- Why the work starts now
- What you need to capture first

Who Carries the Obligation

Start here, because it determines whether this is your problem at all.

When your organization deploys a third-party AI tool, **your organization carries the legal obligation.** The company that built the model does not carry it for you.

What this means practically:

- Vendor terms do not transfer the obligation
- A vendor's System and Organization Controls report — SOC 2 — does not satisfy it
- Providers and deployers have separate duties
- Their duties do not discharge yours

If a regulator opens an inquiry, they ask you. What you knew. What you assessed. What you documented. When.

"We licensed it from an established vendor" is a procurement position. It is not a regulatory one.

Why This Program Has a Different Shape

If you have run any of these, you know how a compliance program works:

- **SOC 2** – System and Organization Controls 2
- **ISO 27001** – the international information security management standard
- **SOX** – internal controls over financial reporting
- **PCI DSS** – Payment Card Industry Data Security Standard

This one is different in one specific way, and it changes your planning.

Other frameworks give you the list

- SOC 2 gives you trust services criteria
- ISO 27001 gives you Annex A
- SOX gives you a control framework

The work is substantial. But you never have to determine what is expected of you. Someone else assembled the register.

Regulation does not

The obligations are:

- Written by legislators across many jurisdictions
- Drafted in statutory language, not control language
- Published in many places, collected in none
- Revised without notice to you

That moves where the difficulty sits

The standard compliance sequence:

1. **Identify** what applies to you
2. **Determine** what you are obligated to do
3. **Produce** the documentation required
4. **Maintain** it as things change

In SOC 2, step one is finished before you start.

Here, **step one is the hardest part of the program.** It is where most organizations stall, and it is why the rest of this manual spends so much time on it.

There is no completion point

No certificate. No annual report. No date on which you are finished for the year.

That is not an oversight in the law. The obligations are continuous, so the program is continuous.

You run it. You do not complete it.

What You Already Have

Most organizations have more relevant material than they realize. Find it before you build anything new.

Existing compliance programs

Overlap between programs is normal in compliance work. One piece of evidence often serves several.

WHAT YOU ALREADY COLLECT	ALSO SERVES
Vendor data handling terms	SOC 2 **and** AI obligations
Vendor risk assessments	Half the AI intake questions
Access reviews	Security **and** AI oversight
Contract reviews	Where processing terms surface
Data inventories	Scope and data-type questions

Two consequences:

- You are not starting from nothing
- You should not collect the same information twice

An existing AI governance intake

Some organizations already have an approval process for new AI tools. If yours does, start there.

What it usually captures well:

- Tool name and vendor
- Requesting team and business justification

- Technical architecture and integration
- IT owner

What it usually misses:

- What the AI actually does to people's data
- Where the affected people live
- Vendor training and retention terms
- Whether staff have been trained
- Whether vulnerable populations are in scope

Most governance intakes were built to route an approval, not to establish legal obligation. They capture the plumbing. They rarely capture the obligation.

A rule on age: treat a record under six months old as a strong starting point that still needs verification. Older than that, re-baseline it. Deployments expand, vendor terms change, and use cases drift from what was approved.

Who owns that process

In most organizations, data privacy and legal own the AI governance intake.

That is useful. It usually means the record you need is one your own function already maintains — not a turf conversation.

Build the Cadence From What You Run

You already have review cycles. Attach this to them rather than inventing a separate rhythm.

The absence of a completion date is disorienting the first time. The answer is to borrow structure from programs that already have one.

- Annual vendor reviews — add the AI questions
- Quarterly access reviews — add AI oversight checks
- Contract renewal — check processing terms
- New vendor onboarding — capture obligations at intake

The earlier this attaches to an existing cycle, the sooner it stops being a project and becomes operations.

Who Else Has a Stake

You may be handling this with a small team, or alone. Either way, you will be asking other people for information.

It helps to know what each of them is actually worried about.

The board — oversight duty Directors must make a good-faith effort to maintain a reasonable monitoring system for material risk. A risk never identified cannot be monitored. Officers carry the same duty. This is corporate law, not regulatory guidance.

The chief executive — reputation For most organizations the penalty is survivable. The publicity is not. Being named by a regulator signals something about how the company is run.

The general counsel — the inquiry The letter arrives on their desk. They answer for the response. At some point they will ask you what the exposure is.

The CFO and audit committee — unquantified exposure An exposure that has not been identified cannot be sized. Reserves, disclosures, and remediation estimates all rest on scoping that has not happened.

IT and security — half your answers They are already involved. They hold much of the information you need in Chapter 3.

Business unit leads — the deployments They own them. These laws attach to their tools, which makes the exposure theirs as well as yours. In practice they are usually helpful. Most people want to know what they are carrying.

Why the Work Starts Now

The reason has little to do with penalties.

A process does not work on day one

You know this from every program you have stood up. Expect several quarters before it produces the results you are looking for.

So the question is not the date

It is whether you meet 2027 with a program that runs, or a compressed effort producing documents nobody trusts.

Those two situations leave different files.

A running program leaves:

- Assessments dated across a period
- Updates recorded when the law changed
- Ownership reassigned as people moved

A compressed effort leaves:

- Documents dated within the same few weeks

Anyone reviewing the file later can tell the difference. That includes regulators.

One practical point on resourcing

Organizations that wait will do this work in the same quarter as everyone else — competing for the same outside counsel.

When You Find a Gap

This will happen. It is worth knowing what it looks like before it does.

How it usually surfaces:

- A regulator or auditor requests something
- A customer questionnaire asks a new question
- A due diligence request lands during a deal

Why the gap existed:

- The work is being done, but proof was never captured
- Nobody knew the requirement applied
- Someone left, and the process left with them

That last one is more common than people expect. A process depending on one person has an expiry date.

What to do:

1. Identify what can be done immediately
2. Look for overlapping processes that captured part of it
3. Check whether another program already does the same work
4. Document the short-term fix and the long-term approach

And one rule that is not negotiable.

You do not hide it. You own it, and you pivot to how it gets addressed.

A gap you found and documented is a different situation from a gap someone else found. The first shows a functioning program. The second raises questions about everything you have not looked at.

What You Capture First

Everything in this manual depends on one thing: **an accurate description of what is actually deployed.**

Not what was approved. Not what the vendor's marketing says. What is running, what it does, and to whom.

That description is called the intake, and it is where the work begins.

A note on vocabulary. In AI governance conversations, this is often referred to as the *deployment use case*. It means the same thing: one tool, used one way, described accurately.

This manual uses "intake" because it names the act of capturing the information rather than the record that results. When you are talking to a governance team, the two terms are interchangeable.

The intake establishes:

- What the tool is and who built it
- What your team actually does with it
- Where the affected people live
- What data it sees
- Whether it makes or influences decisions about people
- How the vendor handles your data
- What controls you already have

Why it comes first: every obligation in every jurisdiction turns on those facts. Get them wrong and everything downstream inherits the error.

An obligation analysis built on inaccurate intake produces a well-formatted, confidently stated, incorrect result. Nothing about the output looks wrong.

Chapter 3 covers how to capture it properly:

- Who in your organization holds each answer
- What to ask them, in language they will understand
- How to tell a solid answer from a confident guess
- Where the information lives when nobody knows it offhand
- What to do when it does not exist anywhere

It is the least interesting chapter in this manual. It is also the one that determines whether anything after it is worth relying on.

A Note on Scope

This manual is written by a compliance practitioner, for compliance practitioners.

- It is not legal advice
- It does not replace your counsel's judgment
- Points requiring a lawyer are identified where they occur

There are also substantial portions of this work that do not require a lawyer, and cannot be scaled if you insist on one. Every compliance function already works within that constraint.

Where the answer is settled, this manual states it plainly.

Where it is not — and in AI regulation there is more unsettled ground than usual — it says so.

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Chapter 3 - What You Are Actually Capturing

Chapter 2 ended on the intake. This chapter covers what goes into it.

Start with the part most organizations get wrong, because getting it wrong makes everything after it wrong too.

This chapter covers:

- Why tools are not the unit of analysis
- What changes an obligation set
- What the intake actually captures
- Who in your organization holds each answer

Tools Are Not the Unit

The instinct is to build a tool inventory. Twenty tools, twenty rows, twenty analyses.

That instinct is wrong, and it will produce an incomplete picture.

Obligations do not attach to software. They attach to what the software does, to whom, and where.

The unit is the deployment

A deployment is three things together:

ELEMENT	QUESTION IT ANSWERS
The tool	What is running
The use case	What your team does with it
The jurisdiction	Where the affected people live

Change any one of the three and you have a different deployment.

Different deployment, different obligation set.

One tool, three use cases

Take a general-purpose AI assistant. Same vendor, same license, same product.

USE CASE	WHAT CHANGES
Drafting internal communications	No decisions about people
Scoring sales leads for follow-up	Now it ranks individuals
Screening job applicants	Now employment law attaches

Three uses of one tool. Three different pictures.

The first may carry very little. The third carries employment law, discrimination exposure, and in several jurisdictions a specific assessment duty.

Counting that as "one tool" hides two of the three.

One use case, three jurisdictions

Now hold the use case still and move the people.

The same applicant-screening deployment reaches:

- Applicants in Illinois
- Applicants in New York City
- Applicants in Germany

Same tool. Same use. Three different obligation sets, because the law follows the people, not your office.

The arithmetic

This is why a tool count understates the work.

- **20 tools** does not mean 20 analyses
- Some tools have one use case
- Others have four or five
- Each use case may reach several jurisdictions

You cannot scope this from a software inventory. You scope it from a deployment inventory, and the second is always larger than the first.

What this means for the intake

One intake per deployment. Not per tool.

If a business lead describes three distinct things their team does with one tool, that is three intakes.

That feels like more work at the start. It is considerably less work than discovering later that your analysis covered one use and your organization is running three.

What Changes an Obligation Set

Beyond the three core elements, a handful of facts move the answer more than anything else.

These are the questions to get right:

- What data does the tool actually see
- Does it make or influence decisions about people
- Who are those people — employees, customers, patients
- Are any of them in a vulnerable group
- What industry are you in
- How does the vendor handle your data

Each of these can change the obligation set on its own. Several of them change it substantially.

Two examples of how much turns on a single answer:

- Biometric data in Illinois carries a private right of action
- Employment context changes which state privacy laws reach the deployment

You do not need to know which law each answer triggers. You need to capture each answer accurately.

Start Wide, Then Narrow

Before the categories, one thing worth knowing: **you will not have every answer on every deployment, and you do not need them to start.**

A small number of answers determine whether the analysis is right at all. The rest determine how precise it is.

The seven that matter most

ANSWER	WHY IT IS LOAD-BEARING
The tool	Identifies what is being assessed
The use case	Different uses carry different obligations
Where the affected people are	Determines which laws reach you at all
Your industry	Determines which sector rules layer on
What data it sees	Drives most obligation triggers
Who it affects — employees or the public	Changes which laws apply entirely
Whether it decides about people	Separates routine processing from consequential

Get these wrong and the analysis is wrong. Get these right and you have something usable, even if the rest is incomplete.

Everything else narrows

The remaining fields rarely change whether an obligation exists. They change how confidently you can rule things in or out.

Without them, the picture stays wide. More frameworks remain possible. More items sit unresolved.

With them, the picture narrows. Frameworks get excluded with a documented reason. Obligations move from "may apply" to "applies" or "does not apply."

What this means practically

A wider result is not a wrong result. It is a less precise one, and it is honest about being less precise.

The failure is not missing information. The failure is **guessing to fill the gap**, which produces a narrow answer built on nothing.

So:

- Capture the seven properly

- Add what else you can
- Mark the rest unresolved
- Refine as answers arrive

A deployment assessed with seven solid answers and six unresolved is in a defensible position. One assessed with thirteen answers where six were invented is not — and nothing about the second one looks worse on the page.

What the Intake Captures

Six categories. Each one exists because it changes the legal picture.

1. The tool

- Product or service name
- Vendor
- Who built the underlying model, if known
- Where it runs

The last two are useful, not essential. If nobody knows, that is a fine answer.

2. The use case

The single most important field.

Describe what your team does with it, what data goes in, and what the output drives.

Be concrete. "We use it for HR" is not a use case. "It screens inbound applications and ranks them for recruiter review" is.

One use per intake. If the description contains "and also," you probably have two.

3. The footprint

Where the affected people live.

Not where your office is. Not where the vendor is hosted. Where the people whose data the tool sees actually are.

This is where surprises show up. Two in particular:

- Customers in states you had not considered
- Any European resident bringing GDPR into scope

A single EU resident is enough. So is one employee working remotely from another state.

4. The data

What the tool sees or uses.

Categories that matter most:

- Personal identifiers
- Health information
- Biometric data used to identify someone
- Financial and employment records
- Behavioral signals
- Precise location

Also: whether the tool combines data from multiple sources. Matching separate data sets together raises the risk profile in several frameworks.

5. The people and the decisions

- Who the tool interacts with or decides about
- Whether it makes or influences decisions about individuals
- Whether it scores, ranks, predicts, or segments people
- Whether it acts without being triggered by a person
- Whether any affected group is in a vulnerable category

On that fourth point. Autonomous processing is usually easy to find, because teams treat it as a feature.

It is often the reason they wanted the tool. "It handles this without anyone touching it" is the business case, and it will be in the justification they wrote to get the purchase approved.

Listen for it in the value proposition:

- Runs automatically
- No manual intervention
- Processes overnight
- Handles it end to end
- Frees the team from

You are rarely extracting this. You are usually just recognizing it.

A note on vocabulary. You will meet the term *ADMT* — automated decision-making technology. It appears in California's rules, in Colorado's law, and in vendor material.

It describes what these questions are getting at: computation standing in for human judgment in decisions that matter to people.

One thing worth knowing. The decision questions and the oversight questions are read together.

A tool that scores applicants with real, case-by-case human review sits differently than the same tool where review is nominal. Same output. Different picture.

That is why "is there human review" is not a yes or no. Someone approving a queue of two hundred recommendations is not exercising judgment on any of them.

Capture what is actually happening. The determination comes later.

Vulnerable groups include patients, children, elderly individuals, people with disabilities, and people in financial vulnerability.

Employees belong in this category too. Not because of who they are, but because of the relationship. An employee cannot freely refuse their employer's processing, and data protection guidance treats that power imbalance as a vulnerability factor.

6. The vendor and your controls

- Whether the vendor trains models on your data
- How long they retain prompts, inputs, and outputs
- Whether they log what is sent to the tool
- What safeguards you already have documented

The last one is where most intakes go soft. Chapter 4 covers how to keep it honest.

Who Holds Each Answer

You will not get the intake from one person. Expect two conversations, sometimes three.

The business lead knows the work

They know what the tool does, what value it provides, what data it works with, and who acts on the output.

They almost always know their data specifically. It is their business.

What they may not be able to tell you: the plumbing. What the tool connects to, where it runs, what else it can reach.

IT knows the infrastructure

Every team adopting an AI tool has an IT person attached to it. Either they recommended it, or they are wiring it into the environment.

That person knows:

- What systems the tool connects to
- What it can access beyond what was described
- Where it is hosted
- Whether it runs on a schedule

You do not have to hunt for who to ask. There is always someone.

The vendor knows their own handling

Training, retention, logging, sub-processors.

Sometimes this is published. Often it is not. Chapter 4 covers how to get it.

Where the split usually falls

QUESTION	WHO ANSWERS
What does the tool do for us	Business lead
What data does it work with	Business lead
What does it connect to	IT
Where does it run	IT
Does it act without being triggered	Either — expect to ask both
Training and retention terms	Vendor, then contract
Do we have a signed data processing agreement	Procurement or legal

That last row is worth noting early. Some of what you need is already in a file somebody else maintains.

The Principle Underneath All of It

You will not get every answer directly. Some you will work out from what people tell you.

That is fine, and it is normal in compliance work. But hold the discipline:

Derive but verify.

Form the hypothesis from what you have been told. Then confirm it. Never assume your reading is correct because it is reasonable.

A confident wrong answer in the intake produces a confident wrong analysis downstream — and nothing about the output will look wrong.

Chapter 4 covers how to run the conversations, where the information lives when nobody knows it offhand, and what to do when it does not exist anywhere.

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Chapter 4 - Getting the Answers

Chapter 3 covered what the intake captures. This chapter covers how to get it.

Most of it comes from two conversations. Some of it comes from a vendor. A useful amount is already sitting in files your organization maintains.

This chapter covers:

- Checking for an existing governance record first
- The conversation with the business lead
- Recognizing when one deployment is actually two
- The conversation with IT
- Getting information out of a vendor
- Keeping the controls answers honest
- Where the information lives
- What to do when it does not exist

Check for an Existing Governance Record First

Before the first conversation, check whether your organization already has an AI approval record for this deployment.

Chapter 2 covered what those records typically give you and typically miss, and the rule on how old is too old to trust without re-baselining. The same read applies here — a governance record is a starting point for the conversations below, not a substitute for them.

If nothing exists yet, the rest of this chapter is where the answers come from instead.

Start With the Business Lead

They requested the tool, or their team uses it. They know what it does and why.

Ask the broad questions first

Do not open with the regulatory question. Open with the work.

The sequence:

1. What does this tool do for your team?
2. How does it do that?
3. What data specifically does it access?
4. Does it perform any tasks without a person triggering it?

The first three establish the picture. The fourth is the one people forget to ask.

Why that order works

You are not asking "does this make decisions about people." That question invites a defensive answer, or a confused one.

You are asking what the tool does. The regulatory answer falls out of the description.

Someone explaining that the tool ranks inbound applications so recruiters know who to call first has told you it influences decisions about people. They did not need to know that was the question.

They know their data

Business leads almost always know what data they work with. It is their business.

What they may not be able to tell you:

- What the tool is connected to
- What it can reach beyond what they described
- Where it runs
- Whether it operates on a schedule

That is not evasion. It is a different job.

Autonomy is usually volunteered

When you ask the fourth question, expect a clear answer more often than not.

Teams treat autonomous processing as a benefit. It is frequently the reason they wanted the tool, and it will appear in the justification written to get the purchase approved.

Phrases that mean yes:

- Runs automatically
- No manual intervention
- Processes overnight
- Handles it end to end

You are usually recognizing this rather than extracting it.

Recognizing Two Use Cases

This comes up constantly, and it is worth watching for.

A team will describe one tool deployment that contains two separate processes. They will not think of it as two things, because to them it is one tool doing their work.

What it sounds like:

- "It drafts the responses, and it also flags which accounts to prioritize"
- "It summarises the call, and then it scores the interaction"
- "It handles scheduling, and it screens the initial applications"

Each of those is two use cases wearing one description.

The tell is usually the word "and." Two verbs describing two different outputs, doing two different things to different people.

Why it matters

Chapter 3 covered the reason: different use, different obligation set. The first half of that sentence may carry very little. The second half may carry a great deal.

An analysis covering the first while your organization is running both leaves a gap nobody knows about.

What to do about it

Recognise it. Say so plainly to the business lead — most people find it obvious once it is pointed out.

How your organization handles the split from there is an internal decision. Some teams run separate intakes immediately. Others record both and sequence them.

That is a process question for your function, not a rule this manual can set for you.

Then Talk to IT

Every team adopting an AI tool has an IT person attached to it. Either they recommended it, or they are connecting it to your environment.

You do not have to find out who to ask. Ask the business lead who supports them.

What IT tells you

- What systems the tool connects to
- What it can access beyond what was described
- Where it is hosted
- Whether it runs on a schedule
- What integrations were configured

The gap this closes

The business lead describes what they put into the tool. IT tells you what the tool can reach.

Those are frequently different. A tool connected to a shared drive, a mailbox, or a customer record system sees considerably more than the person using it thinks about.

Ask specifically: what is this connected to, and what can it read?

Getting Answers From the Vendor

Training, retention, and logging usually come from the vendor. Sometimes it is published. Often it is not.

Look at what you already have first

Before contacting anyone:

- Training materials from implementation
- Kick-off call notes and decks
- The procurement file
- Security review documentation
- The contract and any processing agreement

A surprising amount is already in the building.

Then email the vendor contact directly

If it is not published and not in your files, ask.

Frame the request in a category they already recognize. Third-party vendor due diligence for a security certification is a request every vendor has handled many times. It is also accurate — that information genuinely serves those purposes.

A request framed as unfamiliar gets deprioritized. A request framed as routine gets answered.

On timing: every vendor is different. Some respond in days, some do not respond without escalation. Do not build a schedule around an assumption.

When the vendor answer conflicts with the contract

This happens, and it matters.

The contract takes precedence. What a support contact or salesperson says does not override an executed agreement.

But a conflict is not simply a tiebreaker to resolve. It is a signal to escalate.

Why: if the vendor is describing handling that differs from the agreement, your processing agreement or privacy notice may need updating. That is procurement's and legal's work, not yours.

The sequence:

1. Note the conflict
2. Hand it to procurement
3. Let them clarify with the vendor
4. Take the aligned answer back

Everyone then works from the same record. That is the point.

Keeping the Controls Answers Honest

The controls block asks what safeguards you already have. Signed agreements, transfer mechanisms, privacy notices, retention rules, training, named oversight.

This is where intakes go soft, because everyone wants the answer to be yes.

The follow-up question

Whatever answer you get, ask the same thing:

> **Show me how you can prove that.**

Acceptable answers:

- A document
- A link
- A screenshot of a configured setting
- A named person who owns it

That is not an accusation. It is the next reasonable question, and it is exactly what a regulator or auditor will ask.

What "yes" often means

A policy document exists. That is not the same as the policy covering this deployment.

Check the scope, not just the existence:

- Does the retention schedule cover this tool
- Does the privacy notice describe this processing
- Does the processing agreement name this service
- Is the training completed by the people using it

A retention policy that does not mention this system is not a control for this system.

When you cannot substantiate it

Record it as unresolved.

Unresolved and documented is a defensible position. It shows you looked and you know where the gap is.

A yes you cannot support is not.

Where the Information Lives

When nobody knows an answer offhand, it is usually recorded somewhere.

The deployment

WHAT YOU NEED	WHERE TO LOOK
What the tool does for the team	Business lead; the approval request
Deployment status — pilot, live, inherited	Business lead; IT; change records
Model maker and AI type	Vendor documentation; IT
Connections and integrations	IT; the architecture record
Where it is hosted	IT; the vendor's documentation
Autonomous operation	The business justification; IT configuration

The footprint

This is the one people skip, and it is where the surprises are.

WHAT YOU NEED	WHERE TO LOOK
Where employees are located	HR; payroll
Remote workers in other states or countries	HR; IT device and access records
Where customers or members are located	The customer record system; billing
Where applicants come from	The applicant tracking system
Where support requests originate	The ticketing system
Any European residents at all	All of the above — one is enough

A single EU resident brings the General Data Protection Regulation into scope. So does one employee working remotely from another state, for that state's law.

The data

WHAT YOU NEED	WHERE TO LOOK
Data categories the tool sees	Business lead; IT
Whether it combines data from multiple sources	IT; the integration record
Roughly how many people are affected	Business lead; system user counts; HR
Whether vulnerable groups are in scope	Business lead; the population being served
Revenue and consumer-count thresholds	Finance; legal

The vendor

WHAT YOU NEED	WHERE TO LOOK
Model training on your data	Vendor documentation; the contract
Retention of prompts, inputs, outputs	Vendor documentation; the contract
Whether prompts are logged	Vendor documentation; admin settings
Sub-processors	Vendor trust page; the processing agreement

Your controls

WHAT YOU NEED	WHERE TO LOOK
Signed data processing agreement	Procurement; the contract file
Transfer safeguards	The processing agreement; legal
Privacy notice coverage	The published notice; privacy team
Legal grounds for sensitive data	Privacy team; legal
Retention rules	Records management; the retention schedule

WHAT YOU NEED	WHERE TO LOOK
Prior impact assessments	Privacy team; the governance record
Documented risk classification	Privacy or legal; the governance record
Named owner for oversight	The governance record; business lead
Escalation path	Incident response documentation; security
Staff training records	Learning system; the team's own records

Start with your own files. Ask people second.

When It Does Not Exist Anywhere

This will happen, and it is worth being clear about what it means.

It usually does not mean the information is hard to find. It means **nobody ever captured it.**

- The approval recorded the business justification but not the data categories
- The vendor questionnaire asked about security but not about model training
- The contract was signed but never mapped to the actual deployment

Two things to do

First, record it as unresolved. Not as a no. There is a real difference between "the tool does not do this" and "we cannot currently establish whether it does."

An unknown that is documented can be closed later. An unknown recorded as a no disappears.

Second, fix it going forward. This is the more valuable half.

If a field is missing for this deployment, it will be missing for the next one. Add it to the process that should have captured it:

- Add the data categories to the AI request form
- Add training and retention questions to the vendor questionnaire
- Record hosting region on the asset record
- Capture the affected population at approval, not afterwards

That advice helps whether or not anything else changes. The next deployment arrives with the answers already attached.

The Discipline That Holds It Together

You will work some answers out rather than being told them directly. That is normal.

Derive but verify.

Form the hypothesis from what you have. Then confirm it with the person or the document that can settle it.

Never conclude you are right because your reading is reasonable. Reasonable and wrong looks identical on the page.

Chapter 5 covers what happens once the intake is accurate: determining what the deployment actually obligates you to do.

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Chapter 5 - What the Deployment Obligates You to Do

You have an accurate intake. This chapter covers what happens next.

Fair warning: this is the hardest step in the program, and the manual is going to be honest with you about why.

This chapter covers:

- What an obligation actually is
- The three questions that determine whether one applies
- Why a single deployment triggers several frameworks
- The kinds of obligations you will find
- The discipline that keeps the answer defensible
- Where this stops being your call

What "Obligation" Means Here

People use "does this law apply to us" as if it were one question. It is three, and they resolve in order.

A law can apply to your organization and still create no obligation for a particular deployment. That is not a loophole. It is how the statutes are built.

The Three Questions

1. Does this law reach our organization at all?

Threshold questions. Revenue, headcount, consumer counts, sector, whether you are a covered entity.

Some of these are organization-level facts that have nothing to do with the tool.

Examples of what sits here:

- Revenue and consumer thresholds under state privacy laws
- Whether you are a covered entity or a business associate
- Whether your regulator is a banking agency rather than the Federal Trade Commission
- Whether you have any establishment or activity in a jurisdiction

If the answer is no, nothing downstream matters. Record why and move on.

2. Does it reach this deployment?

The law applies to you. Does it apply to *this* use of *this* tool?

What sits here:

- Whether the affected people are in scope
- Whether the processing activity is one the law covers
- Whether an exclusion applies

The exclusions matter more than people expect. Several state privacy laws exclude individuals acting in an employment context from the definition of consumer. A deployment touching only employees may fall outside those laws entirely — while remaining squarely inside others.

3. Does it trigger the obligation?

The law reaches the deployment. Does this particular duty attach?

What sits here:

- Whether the processing is high-risk
- Whether decisions have significant effects
- Whether sensitive categories are involved
- Whether the tool is classified as high-risk under the applicable framework

This is where the intake earns its keep. Every one of these turns on facts you captured.

Why the order matters

A failure at question one should stop the analysis. If you skip to question three you will find obligations that do not apply to you.

And you should record where each analysis stopped. More on that below.

One Deployment, Several Frameworks

This is the part that surprises people who have run single-framework programs.

An AI deployment is not governed by one law. It sits at an intersection.

A single deployment can be reached by:

LAYER	EXAMPLE
General data protection law	GDPR, UK GDPR
AI-specific law	EU AI Act, state AI statutes
State comprehensive privacy law	Whichever states the people are in
Sector regulation	Health, financial, insurance, education
Employment law	If the deployment touches workers

These do not substitute for one another.

Meeting your obligations under one framework does not discharge the others. A deployment can require an assessment under a data protection law and a separate assessment under an AI law, addressing different questions about the same system.

The practical consequence

You cannot determine obligations by picking the framework you know best and working through it.

You have to work outward from the deployment facts, and check every framework the facts reach.

That is the step nobody can do quickly by hand. Chapter 7 covers how to manage it at volume — prioritization, cadence, and knowing when to bring in help.

The Kinds of Obligations You Will Find

Not every obligation looks the same. It helps to sort them, because they get satisfied differently.

Things you must do

Conduct obligations. The organization has to behave a certain way.

- Provide meaningful human review before certain decisions
- Give people a way to contest an outcome
- Ensure a qualified person participates in specific determinations
- Restrict certain processing without a lawful basis

Things you must produce

Documentation obligations. A specific artifact has to exist.

- An impact assessment, completed before processing begins
- A record of the processing activity
- A risk classification
- A signed agreement with the processor

Timing matters here. Several of these must exist *before* the activity starts. A document created afterward may satisfy the form and not the requirement.

Things you must tell people

Transparency obligations.

- Notice that AI is being used
- Explanation of a decision
- Disclosure of what data is processed and why

Things you must keep

Retention and record obligations.

- Logs, for a defined period
- Compliance records, for a defined period
- Evidence that oversight occurred

Why sorting them helps

Different people close them.

Conduct obligations usually belong to the business or to IT. Documentation obligations usually belong to you. Transparency obligations often sit with legal or marketing. Record obligations sit with whoever owns the system.

A list of obligations that does not distinguish between them is a list nobody can action.

The Discipline: Three Answers, Not Two

This is the most important habit in the chapter.

For every framework you assess, the answer is one of three things:

ANSWER	MEANING
Applies	The obligation attaches. Here is why.
Does not apply	Assessed and excluded. Here is why.
Unresolved	We could not establish this. Here is what is missing.

Never let an unknown resolve silently into "does not apply."

That is the single most common failure in this work, and it is invisible. A framework you never checked and a framework you checked and excluded look identical in a file that only records what applies.

Record the reasoning, not just the conclusion

For each framework, write down:

- What you concluded
- Which fact drove it
- Where the analysis stopped

"Assessed and excluded — employment context, consumer definition does not reach this deployment" is a defensible record.

Silence is not.

Why this matters later

When someone asks what you looked at, the answer should not depend on memory.

An examiner, an auditor, or your own successor should be able to see the shape of the analysis — including the parts that came out negative.

Where This Stops Being Your Call

Some of this is factual. Some of it is judgment, and the judgment belongs to counsel.

Reasonably yours to determine:

- Which frameworks the deployment facts reach
- Which obligations are documented as triggered
- What documentation exists and what does not
- Where the gaps are

Counsel's to determine:

- Whether a contested threshold is met
- Whether an exclusion genuinely applies
- Whether what you have is sufficient
- Whether to proceed with the deployment

Escalate when:

- The applicability question is genuinely close
- Two frameworks appear to conflict
- The standard is unsettled and the exposure is material
- Someone wants to rely on an exclusion you cannot clearly evidence

There is more unsettled ground in AI regulation than in most areas. Several key standards are awaiting interpretation. That is not a reason to stall — it is a reason to document your reasoning clearly, so the position can be defended or revised as the law settles.

An Honest Word About Scale

Everything above is workable for one deployment.

Run the same analysis across every deployment in an organization of any size and the arithmetic changes.

- Each deployment touches several frameworks

- Each framework has its own applicability, scope, and trigger questions
- Each answer has to be recorded with its reasoning
- All of it has to be revisited when the law moves

That is not a knowledge problem. It is a volume problem.

Most organizations do not fail this step because they cannot reason about a statute. They fail it because they cannot reason about four hundred of them, repeatedly, and keep the record current.

Chapter 6 covers what the mandated documents actually are. Chapter 7 covers the discipline for maintaining all of it at volume — prioritizing, keeping the record current, and being honest about what you can and cannot review.

Use and reuse

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Chapter 6 - The Documents the Law Requires

Some obligations are satisfied by behaving a certain way. Others are satisfied only by producing a specific document.

This chapter covers the second kind, because they are the ones organizations most often discover late.

This chapter covers:

- Why these documents exist
- The instruments you will meet
- What each one is actually for
- Why timing matters more than content
- What goes into them, and who supplies it
- Why one deployment can require several

Why These Exist

Regulators cannot inspect every AI system in operation. What they can do is require you to examine it yourself, in writing, before you deploy it.

That is what these documents are. **A structured record that you looked, and what you found.**

Two consequences follow, and both matter.

The document is evidence of the thinking, not a substitute for it. A completed assessment that nobody engaged with satisfies nothing.

The absence of one is itself a finding. In several enforcement decisions the failure was not that a system caused harm. It was that the organization had never assessed whether it might.

The Instruments You Will Meet

Four, broadly. The names vary by jurisdiction, and the acronyms overlap in unhelpful ways.

Data Protection Impact Assessment

Where it comes from: General Data Protection Regulation, Article 35. The UK equivalent applies under UK GDPR.

In force since: May 2018 in both the EU and the UK.

What it examines: risk to data and to the people the data is about.

When it is required: before processing likely to result in high risk. Large-scale sensitive data, systematic monitoring, and automated decisions with significant effects are the common triggers.

Commonly abbreviated: DPIA.

Fundamental Rights Impact Assessment

Where it comes from: EU AI Act, Article 27.

Applies from: December 2027.

What it examines: risk to people's fundamental rights — whether the system treats them fairly, whether it creates disadvantage, and whether affected people can meaningfully challenge an outcome.

Who it applies to: the obligation reaches specific deployer categories rather than every deployer of a high-risk system. The categories are defined in the Article, and whether your organization falls inside them is a question for counsel.

Commonly abbreviated: FRIA.

Worth knowing: the official template contemplated by the Act has not been published. Organizations are building the assessment from the Article's own requirements.

State privacy and data protection assessments

Where they come from: US state comprehensive privacy laws. Virginia, Colorado, Connecticut, Texas and others each have their own version.

In force since: 2023 in the earliest states, with more added each year.

What they examine: processing that presents heightened risk to consumers.

When they are required: common triggers include profiling with significant effects, targeted advertising, sale of personal data, and processing of sensitive categories.

A caution on abbreviations. These are variously called data protection assessments, privacy risk assessments, and data privacy impact assessments depending on the state. **Read the statute rather than the acronym.**

The one that is not an assessment at all

"DPA" is used for both a *data protection assessment* — a document you produce — and a *data processing agreement* — a contract you sign with a vendor.

They are unrelated. One is an internal analysis. The other is a bilateral agreement under a different provision entirely.

Write them out in full. This causes real confusion in real organizations, including between people who should know better.

When These Took Effect

This is the part worth pausing on.

Three of the four have been mandatory for years. Only one is still ahead.

INSTRUMENT	MANDATORY SINCE
Data Protection Impact Assessment — EU	May 2018
Data Protection Impact Assessment — UK	May 2018
Virginia data protection assessment	January 2023
Colorado data protection assessment	July 2023
Connecticut data protection assessment	July 2023
Texas data protection assessment	July 2024
California risk assessment requirements	January 2026
Fundamental Rights Impact Assessment	**December 2027**

A note on the UK date. The obligation did not begin at Brexit. It applied from May 2018 while the UK was an EU member state, and carried over into UK GDPR when the transition period ended in January 2021.

Organizations sometimes assume the UK requirement is newer than it is. It is not.

What that means

If your organization has European exposure, the Data Protection Impact Assessment obligation is eight years old.

If you have consumers in Virginia, Colorado, or Connecticut, the state assessment obligation has been live since 2023.

These are not deadlines to prepare for. They are requirements that already apply.

The Fundamental Rights Impact Assessment is the only one of the four that is genuinely future work — and it is the one attracting most of the attention.

The practical read

An organization focused on the December 2027 date, while carrying an unmet assessment obligation from 2018, has its attention in the wrong place.

Check what is already required before planning for what is coming.

Timing Matters More Than Content

This is the part organizations get wrong most often.

Several of these must exist before the activity begins.

- A Data Protection Impact Assessment is required *before* high-risk processing starts
- A Fundamental Rights Impact Assessment is required *before* the system is put into use

An assessment completed after deployment may satisfy the form. It does not satisfy the requirement, and the date on it says so.

What this means practically

The trigger point is procurement or design, not go-live.

By the time a tool is running, the window for a pre-deployment assessment has closed. You cannot reopen it.

Two implications:

- Your assessment process has to connect to whatever approves new tools
- Deployments already running need a different conversation than new ones

For deployments already live

You cannot backdate. You can document honestly.

A defensible position:

- The assessment was completed on this date
- The deployment predated the process
- Here is what we found and what we changed

Not defensible: an assessment with no date, or a date that quietly implies it existed all along.

Chapter 2 covered this. It applies here in the most concrete way.

What Goes Into Them

Every one of these documents has two halves, and they are supplied by different people.

The regulatory half

- Which framework applies and why
- Which provisions are engaged
- What the statute requires the document to address
- How the instrument must be structured
- What the vendor's published practices are

This half requires knowing the law. It requires nothing about your organization.

The organizational half

- What your controls actually are
- How the data actually flows
- What risk mitigations you have chosen
- Who owns the system
- Whether what you have is adequate

This half requires knowing your organization. No one outside it can supply this.

Why the split matters

The two halves fail differently.

Get the regulatory half wrong and the document addresses the wrong requirements. It can look complete and satisfy nothing.

Get the organizational half wrong and the document is inaccurate about your own operations. That is worse, because it is discoverable.

And the sufficiency judgment — whether what you have is adequate — is neither half. That is counsel's.

Why One Deployment Can Require Several

The instruments are not alternatives.

A high-risk deployment in Europe can require a Data Protection Impact Assessment under data protection law **and** a Fundamental Rights Impact Assessment under AI law. They examine different things.

- One asks what happens to the data
- The other asks what happens to the people

Neither answers the other's question, and completing one does not discharge the other.

And the state assessments multiply by state

A deployment reaching consumers in several states may trigger an assessment obligation in each of them, under each state's own rules.

The arithmetic is the same as Chapter 3. It is not one document per tool. It is a set of documents per deployment, determined by where the people are and what the system does to them.

Where Counsel Comes In

Reasonably yours:

- Identifying which assessments the deployment triggers

- Assembling the factual record
- Completing the organizational half
- Tracking what exists and what does not

Counsel's:

- Whether a contested applicability question is met
- Whether an exclusion applies
- Whether the completed assessment is sufficient
- Sign-off

Escalate early when:

- The applicability question is genuinely close
- The deployment sits in a high-consequence domain
- An exclusion is being relied on that cannot be clearly evidenced
- The assessment surfaces a risk nobody has decided how to handle

That last one is worth watching for. An assessment that identifies a serious risk and records no decision about it is not a finished document.

The Practical Reality

For one deployment, this is manageable. You determine what is required, assemble the record, complete it, and get it reviewed.

Across your AI tool inventory, the arithmetic changes:

- Several deployments per tool
- Several jurisdictions per deployment
- Several instruments per jurisdiction
- Each one dated, each one reviewed
- All of it revisited when the law moves

Organizations rarely fail this because they cannot complete an assessment. They fail because they cannot complete forty of them, keep each one current, and evidence when each was done.

Chapter 7 covers that problem directly.

Use and reuse

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Chapter 7 - Keeping It Current

Every compliance program has a maintenance requirement. This one is no different.

You have run this before. Controls get tested annually. Vendor reviews come around. Policies get refreshed. The work in this chapter is the same shape, applied to a new subject.

This chapter covers:

- What "current" actually requires
- What changes underneath a deployment
- Setting a review rhythm
- What a maintained record looks like
- Prioritizing when you cannot review everything at once

What "Current" Requires

It does not mean reviewing everything constantly. No program works that way.

It means two things:

1. You know when something changed that affects an obligation you documented
2. You can show when you last looked

The second matters more than people expect. A review that found nothing is still a review, and recording it is how the file demonstrates ongoing attention.

A snapshot says what was true on one date.

A maintained record says what was true, what changed, when you checked, and what you did about it.

What Changes Underneath a Deployment

Three things move, and they move independently.

The law

Statutes change on legislative calendars, and they change in both directions.

- Colorado's original AI law was repealed and replaced before it took effect
- The EU AI Act's high-risk deployer date moved from August 2026 to December 2027
- The UK revised its automated decision-making provisions in 2025

An analysis that was correct in March may need revisiting in September. That is expected, not a failure of the original work.

The vendor

Product updates rarely announce the parts that matter to you.

What moves:

- Sub-processors added or replaced
- Retention periods changed with a settings update
- Model versions swapped underneath the product
- Terms revised at renewal
- New features enabled by default

The deployment

This one happens inside the building, which makes it easy to miss.

How deployments drift:

- The use case expands beyond what was assessed
- Another team adopts the same tool for a different purpose
- New data sources are connected
- The tool reaches people in a new location

It arrives as ordinary business progress, not as a request for review.

Setting the Review Rhythm

Two kinds of review. You need both.

Calendar-based

Set an interval and hold it. Annual is standard. Shorten it for higher-risk deployments.

What to check each time:

- Has the use case changed
- Has the vendor changed anything material
- Has the law changed in any location in scope
- Are the documented controls still in place
- Is the named owner still in the role

That last one catches more than people expect. Ownership decays quietly when people change jobs.

Event-based

Some things should prompt a review regardless of the calendar.

TRIGGER	WHY IT MATTERS
New location in scope	Different law may apply
New use case for an existing tool	New deployment, new analysis
Vendor contract renewal	Terms may have changed
New data category added	May cross a threshold
Vendor announces a significant change	Model, hosting, or sub-processor
Regulatory change where you operate	Direct impact
Incident or complaint involving the tool	Obvious

A tool reviewed in January and expanded in February should not wait until the following January.

Attach it to what already runs

You have review cycles. Add these questions to them rather than building a separate process.

- **Vendor risk reviews** – add the AI questions
- **Contract renewal** – check processing terms
- **Access reviews** – confirm oversight ownership is current
- **New vendor onboarding** – capture the intake at the start

The last one is the highest-value change you can make. A deployment captured properly at onboarding never needs reconstructing later.

Watching for Regulatory Change

You are watching the locations your deployments reach.

What you are looking for:

- New laws where your people are
- Amendments to laws you already track
- Effective dates arriving
- Regulator guidance that shifts interpretation
- Enforcement decisions that show how a provision is read

Practical approach:

- Subscribe to updates from the regulators in your main locations
- Follow the law firm alerts covering your sectors
- Set calendar reminders for known effective dates
- Ask counsel what they are watching – they usually have a list

Known dates are the easiest win. December 2027 is on the calendar. Put it there.

What a Maintained Record Looks Like

If someone reviews your file, these are the markers.

Present:

- Assessments dated across a period rather than clustered
- Reviews logged, including ones that found nothing
- Changes recorded when the law moved
- Ownership updated when people changed roles
- Unresolved items still marked unresolved, with what is missing

Record the reviews that found nothing. They are evidence you looked. A gap in the log reads as inattention, whether or not that is what happened.

Keep prior versions

When you update an assessment, keep what it said before.

The question is rarely "what does this document say now." It is "what did you know, and when."

A document that has been overwritten cannot answer that. A versioned one can.

Prioritizing the Work

You will not review everything at the same depth. No program does.

Scope deliberately, and record the reasoning.

Higher attention for deployments that:

- Affect people in consequential domains — employment, lending, healthcare, insurance, housing
- Involve sensitive data categories
- Operate with limited human review
- Reach multiple locations
- Serve vulnerable populations

Lower attention for deployments that:

- Produce internal outputs only
- Touch no personal data
- Have consistent human review before anything reaches a person

The important part is that the reasoning is written down. A deliberate, documented prioritization is a defensible program decision. An undocumented one is indistinguishable from having missed something.

If the volume outgrows the team

That happens, and there are ordinary answers to it.

- Narrow the scope deliberately and record why
- Move the intake earlier so less has to be reconstructed
- Use the reviews already running rather than adding new ones
- Bring in help for the parts that scale badly

What is not an option: describing a monitoring practice you are not actually running. The record has to match what you do.

Where This Leaves You

The four questions are covered.

QUESTION	CHAPTER
What applies to us?	3 and 4
What are we obligated to do?	5
What documents does the law require us to produce?	6
How do we keep them current?	7

The program does not finish. It runs — and what it produces is a record you can hand to anyone who asks.

That record is built before the question arrives, or it is built under it. The two do not look the same, and the difference is visible to anyone reading the file.

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Handout - Data Protection Impact Assessment

Also called: DPIA. In some organizations, a privacy impact assessment or PIA.

Where it comes from: General Data Protection Regulation, Article 35. The UK equivalent applies under UK GDPR.

Mandatory since: May 2018, in both the EU and the UK.

Timing: completed *before* the processing begins.

When It Is Required

Required where processing is likely to result in a high risk to people.

Common triggers:

- Systematic and extensive evaluation of people, including profiling
- Large-scale processing of sensitive categories
- Systematic monitoring of a publicly accessible area
- Automated decisions with legal or similarly significant effects
- Large-scale processing generally
- Combining or matching data from several sources
- Data about vulnerable people, including employees
- Use of new or innovative technology

Rule of thumb: the more of these a deployment hits, the less room there is for judgment about whether one is needed.

If in doubt, do it. A completed assessment that turns out not to have been required costs you time. A missing one that was required is a finding.

What It Has to Cover

Four core elements, plus two consultation requirements.

1. A description of the processing

- What the tool does
- What data it processes
- Why you are doing it
- Who is involved — controller, processor, sub-processors
- How long data is kept
- Where data goes, including transfers

2. Necessity and proportionality

- Why this processing is necessary for the purpose
- Whether a less intrusive option would achieve it
- The lawful basis you are relying on
- The additional condition, if sensitive categories are involved
- How people exercise their rights

3. The risks to people

- What could go wrong for the individuals affected
- How likely it is
- How serious it would be
- Which groups are most exposed

4. The measures addressing those risks

- What safeguards are in place
- What you are changing as a result
- Residual risk after mitigation
- Who owns each measure

Plus: consultation

- **Your data protection officer's advice**, where one is appointed

- **The views of the people affected**, where appropriate

That second one is frequently skipped. Where it is not practical, record why.

Where the Information Probably Lives

WHAT YOU NEED	WHERE TO LOOK
What the tool does	Business lead; the approval request
Data categories	Business lead; IT
Purpose of processing	Business lead; the business justification
Systems it connects to	IT; the integration record
Sub-processors	Vendor trust page; the processing agreement
Retention periods — vendor	Vendor documentation; the contract
Retention periods — yours	Records management; the retention schedule
Transfers outside the EU or UK	The processing agreement; IT hosting records
Transfer safeguards	The processing agreement; legal
Lawful basis	Privacy team; prior assessments for similar processing
Condition for sensitive data	Privacy team; legal
How rights requests are handled	Privacy team; the rights request process
Existing security measures	Information security; the security review
Human review arrangements	Business lead; the process documentation
Named owner	The governance record; business lead
DPO advice	Your DPO — ask directly, and record the response

WHAT YOU NEED	WHERE TO LOOK
Views of affected people	Employee representatives, customer research, or a recorded reason why not

Who Supplies What

You or the privacy function:

- The regulatory framing and structure
- Lawful basis and conditions
- Risk analysis
- Assembling the record

The business:

- What the tool does and why
- Who is affected
- What happens to the output

IT and security:

- Architecture, hosting, connections
- Security measures in place

Legal or counsel:

- Contested applicability questions
- Whether the assessment is sufficient
- Sign-off

Common Gaps

Watch for these – they are the ones that show up repeatedly.

- **No date, or a date after processing started.** The timing is part of the requirement.
- **Necessity treated as a formality.** "We need it because we decided to use it" is not an assessment of necessity.

- **Risks described generically.** "Data breach" is not a risk analysis. What happens to *these* people if *this* system fails?
- **Measures listed as planned, never confirmed.** A mitigation with no owner and no date is not a control.
- **DPO consulted but the advice not recorded.** The consultation is required. So is the record of it.
- **No residual risk stated.** After mitigation, what is left? Someone has to accept it, by name.
- **Never revisited.** An assessment from three years ago for a tool that has changed twice is a snapshot.

One Practical Note

A DPIA is not a form to be completed. It is a record that you examined the processing and reached conclusions.

An assessment that reaches no conclusions, identifies no risks, and prompts no changes has not been done — regardless of how many boxes are filled.

Statutory content requirements should be verified against the current text of the regulation before relying on this handout for a live assessment.

Sources

- General Data Protection Regulation, Regulation (EU) 2016/679 — Article 35, and Articles 35(2) and 35(9) for the consultation requirements
- UK General Data Protection Regulation and the Data Protection Act 2018
- Guidance on identifying high-risk processing is issued by the European Data Protection Board and, in the UK, by the Information Commissioner's Office

The content requirements summarized here are drawn from Article 35(7). Read the current text of the Article before relying on this handout for a live assessment. National supervisory authorities also publish lists of processing operations that do and do not require an assessment in their jurisdiction — check the list for each country in scope.

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Handout - Fundamental Rights Impact Assessment

Also called: FRIA.

Where it comes from: EU AI Act, Article 27.

Applies from: December 2027.

Timing: completed *before* the system is put into use.

Who It Applies To

This is the first thing to establish, and it is narrower than people assume.

The obligation reaches specific categories of deployer rather than everyone deploying a high-risk system. The categories are set out in the Article.

Two questions, in order:

1. Is the system high-risk under the Act
2. Does your organization fall within the deployer categories the Article names

A high-risk system alone does not create the obligation. Both have to be true.

This is a question for counsel. Get the applicability answer before building the assessment, not after.

What It Examines

The Data Protection Impact Assessment asks what happens to the **data**.

The Fundamental Rights Impact Assessment asks what happens to the **people**.

The questions it is built around:

- Which rights could be affected

- What happens to a specific person when the system gets it wrong
- Whether the system creates disadvantage for particular groups
- Who is accountable for the outcomes
- What oversight exists
- Whether affected people can meaningfully challenge a decision

That last one is the heart of it, and it is the question a data protection assessment does not ask.

What It Has to Cover

The deployer's own use, described concretely:

- The processes in which the system will be used
- How long you intend to use it and how often
- The categories of people likely to be affected
- The specific risks of harm to those categories
- The human oversight arrangements in place
- What happens if a risk materialises — governance, escalation, complaint routes

Note the emphasis on *your* deployment. This is not an assessment of the technology in general. It is an assessment of what you are doing with it, to whom, in your context.

A Note on the Template

The official template contemplated by the Act has not been published.

Organizations are building the assessment from the Article's own requirements. Where structured guidance exists, it comes from civil society and human rights bodies rather than the regulator.

Practical consequence: document your reasoning about structure as well as substance. If the format is later standardized, you want a record of why you built it the way you did.

Where the Information Probably Lives

WHAT YOU NEED	WHERE TO LOOK
High-risk classification	Prior governance record; vendor documentation; counsel
Whether you fall in the deployer categories	Counsel
The processes the system supports	Business lead; process documentation
Intended duration and frequency of use	Business lead; the business case
Categories of people affected	Business lead; HR; the customer record system
Whether vulnerable groups are included	Business lead; the population served
Specific risks of harm	Business lead; prior incidents; complaint records
Human oversight arrangements	Business lead; process documentation; IT configuration
Who can override the system	Business lead; the named owner
Escalation route	Incident response documentation; security
Complaint mechanism	Customer service; HR for employee-facing systems
Provider information about the system	Vendor documentation; the instructions for use
Existing data protection assessment	Privacy team — reuse what overlaps

On the last row. Where a Data Protection Impact Assessment already exists for the same deployment, the factual sections overlap substantially. Reuse the description of the processing rather than rewriting it.

The analysis does not overlap. The facts do.

Who Supplies What

Counsel:

- Whether the obligation applies at all

- Whether the assessment is sufficient

You or the privacy function:

- Structure and assembly
- Coordinating the input
- The record

The business:

- What the system does in practice
- Who is affected
- What happens when it is wrong

People closer to those affected:

- Employee representatives, where the system touches workers
- Customer-facing teams, where it touches customers
- Anyone who handles complaints about the current process

That last group is easy to overlook and worth including. The people who field complaints already know how the existing process fails people.

Common Gaps

- **Built before applicability was confirmed.** Establish whether you are in scope first.
- **Assessed the technology rather than the deployment.** The Article asks about your use, in your context.
- **Affected groups described too broadly.** "Customers" is not a category of person likely to be affected. Which customers, in what circumstances?
- **Human oversight described as existing rather than functioning.** Who can override it, do they have the information to do so, and does it happen in practice?
- **No complaint route.** If someone is affected and disagrees, what do they do?
- **Nobody outside the project consulted.** The assessment is about people. Talk to someone who represents them.
- **Treated as a document exercise.** If the assessment surfaces a serious risk and nothing changes, it has not done its job.

One Practical Note

This assessment is not a compliance artifact bolted onto a decision already made.

It is intended to inform whether and how to deploy. An assessment produced after the deployment decision, to justify it, satisfies the form and misses the point — and that is visible to anyone reading it.

Applicability and content requirements should be confirmed against the current text of the Act, and applicability specifically should be confirmed with counsel before relying on this handout.

Sources

- EU Artificial Intelligence Act, Regulation (EU) 2024/1689 — Article 27, with Annex III for the high-risk categories and Article 26 for the wider deployer obligations
- The applicability categories are set out in Article 27(1). Read them directly.
- The template contemplated by Article 27(5) had not been published at the time of writing

Structured methodologies have been published by civil society and human rights organizations rather than by the regulator. Where you rely on one, record which and why.

Applicability specifically should be confirmed with counsel. A high-risk classification alone does not create the obligation.

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Handout - State Data Protection and Privacy Risk Assessments

Also called: data protection assessments, privacy risk assessments, data privacy impact assessments — depending on the state.

Where they come from: US state comprehensive privacy laws. Virginia, Colorado, Connecticut, Texas, California and others each have their own version.

Mandatory since: 2023 in the earliest states, with more added each year.

Timing: before the triggering processing begins.

A Warning on Names

These instruments are not standardized, and the abbreviations collide.

- Some states call it a data protection assessment
- Some call it a privacy risk assessment
- California's regulations use risk assessment language of their own

And separately: "DPA" is used for both a *data protection assessment* — this document — and a *data processing agreement* — a contract with a vendor. They are unrelated.

Read the statute, not the acronym. Write both out in full when you are talking to anyone.

When One Is Required

Triggered by processing that presents heightened risk. The specific triggers vary by state, but the common set is:

- **Targeted advertising**
- **Sale of personal data**
- **Profiling** where it produces legal or similarly significant effects

- **Sensitive data** processing
- Any processing presenting a heightened risk of harm

The important consequence: these attach per state, per processing activity.

A deployment reaching consumers in four states with an assessment obligation may require an assessment in each — under each state's own rules.

What They Have to Cover

The structure is broadly consistent across states. Verify the specific requirements for each state in scope.

Identify the processing

- What the activity is
- What data is involved
- Which consumers are affected

Weigh benefits against risks

This is the distinctive feature of the state assessments. They are explicitly a **balancing exercise**.

- Benefits to your organization
- Benefits to the consumer
- Benefits to other stakeholders and the public
- Risks to the rights of the consumer

Account for context

Several states direct you to consider:

- Whether de-identified data could be used instead
- What the consumer reasonably expects
- The context in which the data was collected
- The relationship between your organization and the consumer

Record the safeguards

- What mitigations reduce the risk

- Whether the benefits, as mitigated, outweigh the risks

Where the Information Probably Lives

WHAT YOU NEED	WHERE TO LOOK
Which states are in scope	The customer record system; HR for employees; billing
Whether processing is a "sale" or "share"	Legal; the vendor contract; ad tech configuration
Whether targeted advertising is involved	Marketing; the ad platform configuration
Whether profiling produces significant effects	Business lead; what the output drives
Sensitive categories in scope	Business lead; IT; the data inventory
Benefits to the organization	The business case; the approval request
Benefits to consumers	Business lead; customer research
Consumer expectations	The privacy notice; customer research; complaint records
Context of collection	The privacy notice; the collection point
Whether de-identified data would work	The business lead and IT together
Existing safeguards	Information security; the privacy team
Retention rules	Records management; the retention schedule
Prior assessments for similar processing	The privacy team

Two Families, Not One Standard

This is what makes state assessments harder than they look.

There is no single US model. There are two — and then variation inside one of them.

The Virginia model

The landscape is dominated by what is usually called the Virginia model — a template set by the Virginia Consumer Data Protection Act in 2021 and since adopted by fifteen or more states.

Colorado, Connecticut, Utah, Indiana, Montana and many others followed Virginia rather than California — partly because Virginia used clear numeric thresholds a business can assess, and partly because it created no private right of action.

California is structurally different

California's framework diverges significantly. It created the only dedicated state privacy enforcement agency, the broadest consumer rights, a limited private right of action for breaches, and the most detailed implementing regulations of any state.

A California assessment is not a Virginia assessment with the state name changed.

And the Virginia-model states still differ from each other

Sharing a template is not sharing requirements.

Sensitive data is where the divergence is widest. Virginia, Colorado, Connecticut, Oregon, Maryland and most newer laws require opt-in consent. Utah and Iowa allow opt-out.

Thresholds vary just as much. Virginia applies at 100,000 consumers. Montana at 50,000. Texas has no revenue threshold and reaches any business not classified as small. Nebraska has no threshold at all.

Even the opt-out rights differ. Most Virginia-model states provide three — sale, targeted advertising, and profiling. Iowa omits the profiling opt-out.

None of that is visible from the fact that they share a lineage.

What This Means Arithmetically

One deployment. One tool, used one way, reaching people in several places.

REQUIREMENT	APPLIES WHEN
Data Protection Impact Assessment — EU	Any EU residents affected
Data Protection Impact Assessment — UK	Any UK residents affected

REQUIREMENT	APPLIES WHEN
Fundamental Rights Impact Assessment	If in scope, from December 2027
State assessment – Virginia	If triggered
State assessment – Colorado	If triggered
State assessment – Connecticut	If triggered
State assessment – Texas	If triggered
State assessment – California	If triggered

Eight documents for one deployment, and that list is not exhaustive. Twenty states have enacted comprehensive privacy laws as of 2026.

They are not copies of each other

- The **European assessment** examines risk to data
- The **Fundamental Rights assessment** examines risk to people's rights
- The **state assessments** are a benefits-versus-risks balancing exercise
- **California** differs structurally from the Virginia-model states
- **Virginia-model states** differ from each other on sensitive data, thresholds, and opt-out scope

You cannot write one and adapt it seven times.

The factual sections overlap – what the tool does, what data it processes, who is affected. Reuse those.

The analysis does not. Each instrument asks a different question, and answering the wrong one produces a document that satisfies nothing.

Working Through It

The exclusions matter as much as the triggers. Several state privacy laws exclude individuals acting in an employment context from the definition of consumer. A deployment touching only employees may fall outside those states entirely – while remaining inside others.

Practical sequence:

1. Establish which states are in scope, from where the people actually are
2. Check the trigger for each state separately
3. Record the states assessed and excluded, with the reason
4. Do not assume requirements are interchangeable

Step three is the one people skip, and it is the one that makes the file defensible.

Who Supplies What

You or the privacy function:

- Determining which states are in scope
- Checking triggers per state
- Structure and assembly
- The balancing analysis

The business:

- What the processing achieves
- Benefits to the organization and to consumers
- What the output drives

Marketing and ad operations:

- Whether targeted advertising is involved
- What data flows to advertising platforms

Legal or counsel:

- Whether an exclusion applies
- Whether a "sale" or "share" is occurring
- Whether the assessment is sufficient

Common Gaps

- **Scoped from where the company operates** rather than where the consumers are.
- **One assessment used for all states** without checking whether the requirements match.

- **Benefits stated, risks not.** It is a balancing exercise. Both sides have to be present.
- **Consumer benefit asserted without basis.** "Consumers benefit from a better experience" is not an analysis.
- **De-identification never considered.** Several states direct you to consider it. Record the conclusion either way.
- **Exclusions relied on but not documented.** If a state is out of scope, write down why.
- **Never revisited when a new state law took effect.** New states are added regularly.

One Practical Note

The state assessments are a balancing exercise, not a risk register.

An assessment that lists risks and mitigations but never states whether the benefits outweigh them has not answered the question the statute asks.

Triggering thresholds and content requirements vary by state and change regularly. Verify against the current text for each state in scope before relying on this handout.

Sources

Statutes referenced

- Virginia Consumer Data Protection Act
- Colorado Privacy Act
- Connecticut Data Privacy Act
- Texas Data Privacy and Security Act
- California Consumer Privacy Act as amended, and the implementing regulations

On the comparative claims

The characterization of a dominant "Virginia model," the count of states with comprehensive privacy laws, the opt-in and opt-out split on sensitive data, and the threshold variations are drawn from published comparative analyses current to 2026. They are directional rather than definitive.

Verify each state's requirements against its own statute before relying on them. Thresholds, sensitive data definitions, exclusions, and cure periods all vary, and several have been amended since enactment.

A standing note

New state laws are enacted every legislative session, and existing ones are amended. Any list of states in this handout is a snapshot.

Use and reuse

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Worksheet - Intake Preparation

Use this before you fill in an intake. Work through it, note where each answer came from, and flag anything you could not establish.

One worksheet per deployment. One tool, used one way. If the description contains "and also," you likely have two.

The Seven That Matter Most

You will not have every answer. You do not need them to start.

These seven determine whether the analysis is right at all. Get them properly, and the worksheet is usable even if the rest is incomplete.

- ☐ The tool
- ☐ The use case
- ☐ Where the affected people are
- ☐ Your industry
- ☐ What data it sees
- ☐ Who it affects — employees or the public
- ☐ Whether it decides about people

Everything else narrows the picture. Without it, more frameworks stay possible and more items sit unresolved. That is a wider result, not a wrong one.

Mark what you cannot establish as unresolved. Do not guess to fill the gap.

Before You Start

Check what already exists:

- ☐ AI approval or governance request for this tool

- ☐ Procurement file
- ☐ Security review documentation
- ☐ Contract and any processing agreement
- ☐ Implementation or kick-off materials

If a governance record exists, note its date: _____

Under six months — use it as a starting point, verify the regulatory fields. Over six months — re-baseline it.

Part 1 — The Deployment

FIELD	ANSWER	SOURCE
Tool name		
Vendor		
Deployment status — evaluating / pilot / live / inherited		
Who built the underlying model		
Where it runs		

Use case — describe in plain language:

What the team does with it, what data goes in, what the output drives.

Check: does this description contain more than one process?

- ☐ Single use — continue
- ☐ Multiple uses — stop, split into separate worksheets

Part 2 — Where the People Are

Not where your office is. Where the affected people live.

POPULATION	LOCATIONS	SOURCE
Employees		
Remote workers		
Customers or members		
Applicants		
Anyone else		

Specific checks:

- ☐ Any employee working remotely from another state or country
- ☐ Any customer or member outside your main operating states
- ☐ Any European resident, even one
- ☐ Any UK resident, even one

Part 3 — The Data

What the tool sees or uses:

- ☐ Names, contact details, identifiers
- ☐ Government or tax identifiers
- ☐ Health information
- ☐ Biometric data used to identify someone
- ☐ Financial account or transaction data
- ☐ Employment records
- ☐ Behavioural signals — usage, patterns, interactions
- ☐ Precise location
- ☐ Children's data
- ☐ Other sensitive categories: _____
- ☐ None of the above

Additional:

QUESTION	ANSWER	SOURCE
Roughly how many people are affected		
Does it combine data from multiple sources		

Part 4 — People and Decisions

QUESTION	ANSWER	SOURCE
Who does it interact with or decide about		
Does it make or influence decisions about individuals		
Does it score, rank, predict, or segment people		
Does it act without a person triggering it		

Vulnerable populations in scope:

- ☐ Patients or healthcare recipients
- ☐ Children under 18
- ☐ Elderly individuals
- ☐ People with disabilities
- ☐ People in financial vulnerability
- ☐ Employees or workers
- ☐ None

Part 5 — The Vendor

Check published documentation first. Email the vendor contact if it is not there.

QUESTION	ANSWER	SOURCE
Does the vendor train models on your data		
How long are prompts, inputs, outputs retained		
Are prompts logged or stored		
Who are the sub-processors		

If the vendor's answer conflicts with the contract:

- ☐ Conflict noted
- ☐ Handed to procurement
- ☐ Aligned answer received

Part 6 – Controls

For every "yes," ask: show me how you can prove that.

Acceptable evidence – a document, a link, a configured setting, or a named owner.

CONTROL	Y / N / UNRESOLVED	EVIDENCE	COVERS THIS DEPLOYMENT?
Signed data processing agreement			
Transfer safeguard for EU or UK data			
Privacy notice covering this use			
Documented legal grounds for sensitive data			
Documented retention rules			
Prior impact assessment			
Documented risk classification			
Meaningful human review			

CONTROL	Y / N / UNRESOLVED	EVIDENCE	COVERS THIS DEPLOYMENT?
Staff training on this tool			
Named owner for oversight			
Escalation path for incidents			

On human review, be specific:

- ☐ Every time, before it affects anyone
- ☐ Sometimes – describe: _____
- ☐ After the fact only
- ☐ None
- ☐ Unresolved

On the last column. A policy existing is not the same as it covering this deployment. Check whether the retention schedule names this system, whether the privacy notice describes this processing, whether the agreement names this service.

Part 7 – Before You Submit

Verification pass:

- ☐ Every answer has a source noted
- ☐ Anything I could not establish is marked unresolved, not "no"
- ☐ I spoke to the business lead
- ☐ I spoke to IT about connections and hosting
- ☐ The use case describes one process, not several
- ☐ The footprint reflects where people are, not where we are
- ☐ Nothing here is an assumption I have not confirmed

Unresolved items – list them:

ITEM	WHAT IS MISSING	WHO CAN CLOSE IT

ITEM	WHAT IS MISSING	WHO CAN CLOSE IT

After You Submit

If any field could not be answered from existing records, fix that going forward:

- ☐ Add the missing fields to the AI request form
- ☐ Add vendor training and retention questions to the vendor questionnaire
- ☐ Record hosting location on the asset record
- ☐ Capture the affected population at approval, not afterwards

The next deployment should arrive with these answers already attached.

Completed by: _____

Date: _____

Deployment reference: _____

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Worksheet - Source Map

Where each answer lives, organized by who holds it.

Use this to plan your conversations. Take the relevant section to each person rather than asking everyone everything.

Start Here — Your Own Files

Check these before you ask anyone anything.

- ☐ AI approval or governance request for this tool
- ☐ Procurement file
- ☐ Security review documentation
- ☐ Contract and any processing agreement
- ☐ Implementation or kick-off materials
- ☐ Prior assessments for similar processing
- ☐ The published privacy notice
- ☐ The retention schedule
- ☐ Asset or system inventory

A surprising amount is already in the building. Asking people for information you already hold costs you credibility the second time you do it.

The Business Lead

What they hold:

QUESTION	NOTES
What the tool does for the team	

QUESTION	NOTES
What data it works with	
What the output drives, and who acts on it	
Who is affected — employees, customers, patients	
Whether it runs without being triggered	Often in the business case
Whether other teams use the same tool	Catches parallel deployments

Also worth asking: who in IT supports you on this? That gets you the next conversation without hunting.

IT and Information Security

What they hold:

QUESTION	NOTES
What systems the tool connects to	
What it can access beyond what was described	Ask specifically
Where it is hosted	Region matters, not just provider
Whether it operates on a schedule	
What integrations were configured	
Existing security measures	
Whether data leaves the region	

The gap this closes: the business lead describes what they put in. IT tells you what the tool can reach.

Human Resources

Often overlooked, and holds the answer to the question people get wrong most.

QUESTION	NOTES
Where employees are located	By state and country
Remote workers outside the main locations	One is enough to change the analysis
Whether the tool touches employee data	
Employee representative arrangements	Needed for some assessments
Training records for this tool	

Sales, Service, and Marketing

Where the customer footprint actually lives.

QUESTION	NOTES
Where customers or members are located	
Where support requests originate	
Where applicants come from	Talent or recruiting
Whether targeted advertising is involved	Marketing or ad operations
What data flows to advertising platforms	

Procurement and Legal

QUESTION	NOTES
Is there a signed data processing agreement	
Does it name this service	Not just this vendor
Transfer safeguards for EU or UK data	
Contract renewal date	Triggers a review
Whether a "sale" or "share" is occurring	Legal call
Any exclusions being relied on	Document the basis

The Privacy Function

If that is not you, it is whoever owns the program.

QUESTION	NOTES
Lawful basis for this processing	
Condition for sensitive categories	
Privacy notice coverage for this use	
Prior impact assessments	Reuse the factual sections
Documented risk classification	
How rights requests are handled	

Records Management

QUESTION	NOTES
Retention rules for this data	
Whether the schedule names this system	Existence is not coverage
Deletion process	

The Vendor

Check published documentation first. See Worksheet 3 for how to ask.

QUESTION	NOTES
Model training on your data	
Retention of prompts, inputs, outputs	
Whether prompts are logged	
Sub-processors	Usually a linked page, not the contract
Hosting region	
Security certifications	

Finance

Only needed for threshold questions.

QUESTION	NOTES
Annual gross revenue	Some state laws turn on this

QUESTION	NOTES
Number of consumers whose data is processed	
Revenue from selling personal information	

When Nobody Has It

Record it as **unresolved**, not as a no. Then fix the process that should have captured it.

MISSING FIELD	WHICH PROCESS SHOULD CAPTURE IT	OWNER

Common fixes:

- ☐ Add data categories to the AI request form
- ☐ Add training and retention questions to the vendor questionnaire
- ☐ Record hosting region on the asset record
- ☐ Capture affected population at approval
- ☐ Ask where users are located at onboarding

The next deployment should arrive with these answers attached.

Use and reuse

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Worksheet - Conversation Guides

Three conversations, in order. Business lead first, IT second, vendor when the published documentation does not answer it.

Conversation 1 – The Business Lead

Goal: understand what the tool does and what the output drives.

Do not open with the regulatory question. Open with the work. The regulatory answer falls out of the description.

The Four Questions

1. What does this tool do for your team?

Let them explain it in their own terms.

2. How does it do that?

This is where the mechanism surfaces.

3. What data specifically does it access?

Business leads usually know this well. It is their business.

4. Does it perform any tasks without a person triggering it?

Expect a clear answer. Teams treat this as a feature.

Follow-Ups Worth Asking

QUESTION	WHY
What happens after the tool produces its output?	Reveals whether it influences decisions
Who sees it, and what do they do next?	Same
Does anyone review before it affects someone?	Human oversight
Every time, or only some cases?	The distinction that matters
Who else in the organization uses this tool?	Catches parallel deployments
Who in IT supports you on this?	Sets up conversation 2

Listen For – Autonomy

These phrases mean the tool acts on its own:

- ☐ Runs automatically
- ☐ No manual intervention
- ☐ Processes overnight
- ☐ Handles it end to end
- ☐ Frees the team from

You are usually recognizing this rather than extracting it.

Listen For – Two Use Cases

The tell is the word "and." Two verbs, two outputs, two different things happening.

Examples:

- "It drafts the responses, and it also flags which accounts to prioritize"
- "It summarises the call, and then it scores the interaction"
- "It handles scheduling, and it screens the initial applications"

If you hear it:

☐ Say so plainly — most people find it obvious once pointed out

☐ Start a second worksheet

☐ Note both here before you forget the second one

Second use case identified: _____

Before You Leave

☐ I can describe the use case in one sentence

☐ I know what the output drives

☐ I know who is affected

☐ I have the name of the IT contact

☐ I asked whether other teams use this tool

Conversation 2 — IT

Goal: find out what the tool is connected to and what it can reach.

Who to ask: the person the business lead named. Every team adopting a tool has one — either they recommended it or they are wiring it in.

The Questions

QUESTION	ANSWER
What is this connected to?	
What can it read that we have not discussed?	
Where is it hosted — which region?	

QUESTION	ANSWER
Does any data leave that region?	
Does it run on a schedule, or only on request?	
What integrations were configured?	
What security measures are in place?	
Who has administrative access?	

The One to Push On

"What can it read?"

The business lead describes what they put into the tool. That is often narrower than what the tool can reach.

Ask specifically about:

- ☐ Shared drives or document repositories
- ☐ Mailboxes
- ☐ Customer or member record systems
- ☐ Ticketing or case systems
- ☐ Calendars
- ☐ Anything else connected at setup

If the Answers Differ From Conversation 1

That is normal and worth recording.

Discrepancy noted: _____

Resolved by: _____

Do not assume either person is wrong. They are answering from different vantage points.

Conversation 3 — The Vendor

Goal: training, retention, logging, and sub-processors.

Only after you have checked published documentation and your own files.

Check First

- ☐ Vendor trust or security page
- ☐ Product documentation
- ☐ Admin console settings
- ☐ The contract and processing agreement
- ☐ Implementation or kick-off materials
- ☐ Security review from procurement

If it is there, you are done. Record where you found it.

If You Have to Ask

Frame the request in a category the vendor recognizes.

A third-party vendor due diligence request is something every vendor has handled many times. A request framed as unfamiliar gets deprioritized.

It is also accurate — this information genuinely serves those purposes.

Draft Email

> **Subject:** Third-party vendor information request — [your organization]
> Hello,
> We are completing third-party vendor due diligence for [tool name], which we use for [brief description].
> Could you confirm the following, or point me to where it is documented:
> 1. Is customer data used to train or improve your models? If so, is opt-out available?
> 2. How long are prompts, inputs, and outputs retained?
> 3. Are prompts logged or stored, and for how long?

4. Who are your sub-processors, and where are they located?
5. In which region is our data processed and stored?

If there is a standard documentation pack that covers these, that would be ideal.
Thank you, [name]

On timing: every vendor is different. Some respond in days. Some need escalation through your account manager. Do not build a schedule around an assumption.

Record What You Get

QUESTION	ANSWER	SOURCE
Model training on our data		
Retention period		
Prompt logging		
Sub-processors		
Processing region		

If the Answer Conflicts With the Contract

The contract takes precedence. What a support contact says does not override an executed agreement.

But a conflict is not just a tiebreaker. It is a signal to escalate.

If the vendor describes handling that differs from the agreement, your processing agreement or privacy notice may need updating. That is procurement's and legal's work.

Sequence:

- ☐ Conflict noted
- ☐ Handed to procurement
- ☐ Procurement clarified with the vendor
- ☐ Aligned answer received and recorded

Everyone then works from the same record.

The Discipline Across All Three

Derive but verify.

You will work some answers out from what people tell you. Form the hypothesis, then confirm it with the person or document that can settle it.

Never conclude you are right because your reading is reasonable. Reasonable and wrong look identical on the page.

Use and reuse

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Worksheet - Controls Evidence

This worksheet exists because controls answers go soft.

Everyone wants the answer to be yes. The question that keeps it honest is the same one an auditor asks:

> **Show me how you can prove that.**

How to Use This

For every control, three things:

1. **The answer** — yes, no, or unresolved
2. **The evidence** — a document, link, setting, or named owner
3. **Whether it covers this deployment** — existence is not coverage

That third column is where most programs fall down. A retention policy that does not name this system is not a control for this system.

Unresolved is a legitimate answer. It shows you looked and knows where the gap is. A yes you cannot substantiate does not.

Acceptable Evidence

- A document you can produce
- A link to the current version
- A screenshot of a configured setting
- A named person who owns it and can confirm

Not acceptable:

- "We have a policy on that"
- "I believe legal handled it"

- "It was part of the procurement process"

Those are leads to follow, not evidence.

The Controls

Agreements and safeguards

CONTROL	Y / N / UNRESOLVED	EVIDENCE	NAMES THIS DEPLOYMENT?
Signed data processing agreement			
Agreement names this specific service			
Transfer safeguard for EU or UK data			
Sub-processor list obtained			

Check on the second row. A master agreement with the vendor is not the same as an agreement covering this product. Vendors frequently have several.

Notice and lawful basis

CONTROL	Y / N / UNRESOLVED	EVIDENCE	COVERS THIS PROCESSING?
Privacy notice published			
Notice describes this processing			
Lawful basis documented			
Condition for sensitive data documented			

Check on the second row. A privacy notice that describes your website does not necessarily describe an AI tool analyzing customer conversations.

Retention and deletion

CONTROL	Y / N / UNRESOLVED	EVIDENCE	NAMES THIS SYSTEM?
Retention rules documented			
Schedule covers this system			
Deletion process exists			
Vendor retention period known			

Assessments and classification

CONTROL	Y / N / UNRESOLVED	EVIDENCE	CURRENT?
Prior impact assessment completed			
Assessment covers this use case			
Risk classification documented			
Assessment dated before processing began			

Check on the last row. Timing is part of several of these requirements. A date after go-live is a finding, not a control.

Human oversight

CONTROL	Y / N / UNRESOLVED	EVIDENCE	ACTUALLY HAPPENS?
Human review before output affects anyone			
Reviewer can override			
Reviewer has information needed to judge			
Named owner for oversight			

CONTROL	Y / N / UNRESOLVED	EVIDENCE	ACTUALLY HAPPENS?
Owner still in the role			

On human review, be specific.

- ☐ Every time, before it affects anyone
- ☐ Sometimes — describe: _____
- ☐ After the fact only
- ☐ None
- ☐ Unresolved

The distinction that matters: someone approving a queue of two hundred recommendations is not exercising judgment on any of them. Record what actually happens, not what the process document says.

Training and escalation

CONTROL	Y / N / UNRESOLVED	EVIDENCE	COVERS THE RIGHT PEOPLE?
Training on this tool delivered			
Completed by people who use it			
Completed by people who oversee it			
Escalation path documented			
Escalation path known to users			

Check on the last row. A documented escalation path nobody has been told about does not function.

The Follow-Up Questions

When someone answers yes, these separate a real control from a remembered one.

ASK	WHAT YOU ARE TESTING
Where is it documented?	Existence
When was it last reviewed?	Currency
Does it name this system?	Coverage
Who owns it?	Accountability
What happens if it fails?	Whether it is operational

None of these are accusatory. They are the next reasonable questions, and they are exactly what a regulator or auditor will ask.

Summary

	COUNT
Controls confirmed with evidence	
Controls in place but not covering this deployment	
Controls unresolved	
Controls absent	

Unresolved items:

CONTROL	WHAT IS MISSING	WHO CAN CLOSE IT	BY WHEN

One Note Before You File This

A control that exists on paper and does not operate is worse than an absent one.

The absent control is a gap you know about. The paper control is a gap you believe is closed — and it will stay believed until someone tests it.

Completed by: _____

Date: _____

Deployment reference: _____

Use and reuse

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Worksheet - Review Triggers and Log

Two kinds of review. You need both.

The calendar catches drift. The trigger list catches everything that does not wait for the calendar.

Part 1 — Setting the Interval

Deployment: _____

Review interval:

- ☐ Annual — standard
- ☐ Every six months — higher risk
- ☐ Quarterly — consequential decisions, sensitive data, limited human review

Reason for the interval chosen:

Record the reasoning. A deliberate, documented interval is a program decision. An undocumented one is indistinguishable from having missed something.

Next scheduled review: _____

Part 2 — What to Check Each Time

CHECK	CHANGED?	NOTES
Has the use case changed		
Has another team adopted the tool		
Has the vendor changed anything material		

CHECK	CHANGED?	NOTES
Have new data sources been connected		
Has the population expanded to a new location		
Has the law changed anywhere in scope		
Are the documented controls still in place		
Is the named owner still in the role		

On the last row. Ownership decays quietly when people change jobs. It catches more problems than people expect.

Part 3 – Event Triggers

These prompt a review regardless of the calendar.

TRIGGER	WHY IT MATTERS	OCCURRED?
New location in scope	Different law may apply	
New use case for an existing tool	New deployment, new analysis	
Vendor contract renewal	Terms may have changed	
New data category added	May cross a threshold	
Vendor announces a material change	Model, hosting, sub-processor	
Regulatory change where you operate	Direct impact	
Effective date arriving	Known in advance – calendar it	
Incident or complaint involving the tool	Obvious	
Team expansion or reorganization	Ownership and access change	
Integration added or removed	Changes what the tool can reach	

A tool reviewed in January and expanded in February should not wait until the following January.

Part 4 — Attaching to Cycles You Already Run

Do not build a separate process. Add these questions to reviews that already happen.

EXISTING CYCLE	WHAT TO ADD	OWNER	ADDED?
Vendor risk review	The AI intake questions		
Contract renewal	Check processing terms		
Access review	Confirm oversight ownership		
New vendor onboarding	Capture the intake at the start		
Annual policy refresh	Check notice still covers this use		

The last row of value is the fourth one. A deployment captured properly at onboarding never needs reconstructing later.

Part 5 — Watching for Regulatory Change

You are watching the locations your deployments reach.

SOURCE	SET UP?	OWNER
Regulator updates for main locations		
Law firm alerts covering your sectors		
Known effective dates on the calendar		
Counsel's watch list — ask them		

Known dates are the easiest win. December 2027 is already on the calendar. Put it there.

Locations currently in scope:

Part 6 — The Review Log

Record every review, including the ones that found nothing.

They are evidence you looked. A gap in the log reads as inattention, whether or not that is what happened.

DATE	REVIEWED BY	TRIGGER	FINDINGS	ACTION TAKEN

Trigger column: scheduled, or name the event.

Findings column: "no change" is a valid and useful entry.

Part 7 — Version Control

When you update an assessment, keep the prior version.

The question is rarely what the document says now. It is what you knew, and when.

VERSION	DATE	WHAT CHANGED	PRIOR VERSION RETAINED?

Part 8 – Prioritizing

You will not review everything at the same depth. No program does.

Higher attention where the deployment:

- ☐ Affects people in consequential domains – employment, lending, healthcare, insurance, housing
- ☐ Involves sensitive data categories
- ☐ Operates with limited human review
- ☐ Reaches multiple locations
- ☐ Serves vulnerable populations

Lower attention where the deployment:

- ☐ Produces internal outputs only
- ☐ Touches no personal data
- ☐ Has consistent human review before anything reaches a person

Prioritization reasoning:

Write the reasoning down. A deliberate, documented prioritization is defensible. An undocumented one is not distinguishable from an oversight.

One Constraint

Describing a monitoring practice you are not actually running is the one option that is not available.

If the volume outgrows what you can review, the answers are ordinary – narrow the scope deliberately and record why, move the intake earlier, use cycles already running, or bring in help.

The record has to match what you do.

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.

Worksheet - Deployment Inventory

One row per deployment. Not per tool.

This is the distinction that determines whether your inventory is accurate or misleading.

A tool used three ways is three rows. A tool used one way in four locations may still be one row, but the locations column has to reflect all four.

Why Tools Are the Wrong Unit

A software inventory tells you what you have licensed.

A deployment inventory tells you what you are responsible for.

They are not the same list, and the second is always longer.

TOOL INVENTORY SAYS	DEPLOYMENT INVENTORY SAYS
We have 20 AI tools	We have 34 deployments
One row per license	One row per tool-and-use
Scoped by procurement	Scoped by what the tools actually do

Before You Start

Where to find deployments you may not know about:

- ☐ The AI approval or governance log
- ☐ Procurement records for the last 24 months
- ☐ Expense reports — departmental software purchases

- ☐ Single sign-on application list
- ☐ Network or endpoint tooling — what is actually running
- ☐ Ask each business unit lead directly
- ☐ Vendor renewal notices

And the one most likely to surface surprises: ask whether any existing tool has recently added AI features. Vendors add them to products you already own, without a new purchase and often without a new approval.

The Inventory

Copy this table and extend as needed. Columns are ordered by what you will need first.

#	TOOL	USE CASE (ONE PER ROW)	BUSINESS OWNER	STATUS	LOCATIONS AFFECTED	PEOPLE AFFECTED	DE AF
1							
2							
3							
4							
5							

Column Notes

Use case — one per row. If the description contains "and also," split it.

Status — evaluating, pilot, live, or inherited. Inherited means it was running before the process existed. Those need a different conversation than new deployments.

Locations affected — where the people are, not where you are. This column is the one most often wrong.

People affected — roughly. A number is fine. Employees, customers, applicants, patients — note which.

Decisions about people — yes, influences, or no. If yes or influences, this deployment needs closer attention.

Sensitive data — yes or no, and which categories if yes.

Assessments required — from the determination. May be several per row.

Assessment status — not started, in progress, complete, or unresolved pending a scope question.

Working Through It

You will not complete this in one pass. Nobody does.

A workable sequence:

1. List every tool you know about
2. For each, ask the business lead how many ways it is used
3. Split into one row per use
4. Fill the locations column — this is where the work is
5. Flag the rows involving decisions about people or sensitive data
6. Work through the flagged rows first

Step four is the one that takes longest and matters most. Most organizations can name their tools. Very few can name where their affected people are without going to ask.

Prioritizing the Backlog

Work through the flagged rows in this order:

PRIORITY	CHARACTERISTICS
First	Consequential decisions, sensitive data, limited human review, multiple locations
Second	Decisions about people, single location, human review present
Third	Personal data, no decisions about individuals
Fourth	No personal data, internal outputs only

Record the prioritization reasoning:

A documented order of work is a program decision. An undocumented one is indistinguishable from having missed something.

Keeping It Current

Add a row when:

- ☐ A new tool is approved
- ☐ An existing tool is used a new way
- ☐ Another team adopts a tool already listed
- ☐ A vendor adds AI features to an existing product

Revisit a row when:

- ☐ The use case changes
- ☐ A new location comes into scope
- ☐ The vendor changes something material
- ☐ The scheduled review date arrives

The fourth trigger in the first list is the one that gets missed. A product you already own becoming an AI product does not arrive as a purchase request.

Summary Position

Useful for reporting upward.

	COUNT
Tools in use	
Deployments identified	
Deployments involving decisions about people	
Deployments involving sensitive data	

	COUNT
Assessments identified as required	
Assessments complete	
Assessments outstanding	
Deployments not yet assessed for scope	

That last row is the honest one. It is the number that tells you how much of the picture you can actually see.

Inventory owner: _____

Last updated: _____

Next full review: _____

Use and reuse

LegisGate · Regulatory Intelligence for AI Deployments · legisgate.com

Written by compliance and audit practitioners who spent decades on the other side of the table.

Version 1.0 · Current as of August 2026 · Verify time-sensitive claims against current sources before relying on them.