

Risk tiers for worker-built apps

<https://prompt-to-prototype.com/operating-model/risk-tiers/> · as of 2026-10-22

Every worker-built app sits in one of four risk tiers, Sandbox, Team, Important or Critical, set by the highest answer to six questions about its data, users, internet reach, dependence, decisions about people and actions. Two floors apply at every tier, and an owner can request a documented downgrade when compensating controls remove the reason for a higher answer. This page is the only full statement of the model and the [gate matrix](#), which you can also [download as a CSV](#); every other page links here.

Who this is for

IT managers, security reviewers and app owners who classify apps, and builders who want to know what their answers mean. You get the whole tier model, what each tier requires at every lifecycle stage, and worked cases. Reading time: about 25 minutes; the matrix is for reference.

Jump to: [the four tiers](#) · [the six questions](#) · [the gate matrix](#) · [downgrades](#)

"Worker-built app" here includes automations, scheduled jobs, workflow-tool flows with AI steps, and agents built in SaaS agent builders or coding tools. They're tiered, registered, owned and retired like any other app, however they were built; for the search term, see [what vibe coding means for a business](#).

The four tiers

Tier	In one line	Owner and support level
Sandbox	A personal, time-boxed experiment. Only you use it, on public data or your own work files with no Confidential or Restricted data, and it's private by default	The builder; self-supported
Team	Shared inside the business (one team, several teams or company-wide), on internal data. If it stops, the people who use it work around it	A business app owner; team-supported (builder plus a champion)
Important	Outside users or a business process rely on it, or it holds Confidential data (most personal data), or it writes to a shared system of record	A business app owner with a named backup; assisted support
Critical	If it fails or errs, harm reaches customers, money, legal obligations or people's livelihoods	A business app owner, plus IT or engineering as co-owner (graduation); IT-supported

Use the tier names alone, in prose, tables and the [app registry](#). Support uses **support levels** (self-supported, team-supported, assisted, IT-supported), described on the [support model](#) page.

Don't: number the tiers

Don't label tiers with numbers, levels or "Low/Medium/High". In IT service management and continuity planning the first-numbered tier usually holds the *most* critical systems, and help desks use the same number for first-line support, so a number gets misread in both directions. If a spreadsheet export needs a sort order, add a hidden rank column (1 = Sandbox through 4 = Critical) that readers never see.

The six tier questions

Not legal advice

Not legal advice. This page summarizes laws as of October 2026 for a general audience. It was written with an AI legal-research persona and hasn't been reviewed by a lawyer for your situation. Ask your counsel.

Answer each question for what the app **can** do with the data, reach and permissions it holds, not for what it's meant to do. For an agent or automation, answer for every permission it has been granted. Each answer maps to exactly one tier.

- Scope the token first. Give the app a token or credential scoped to what it needs (for example, read-only) wherever the tool supports it. If the tool can't scope it, answer for everything the token can do, then check [standard downgrade 3](#).
- Floor breaches: if an answer is higher only because the app breaks a floor by mistake, apply the [floor-breach rule](#).

An answer marked **(legal)** is one whose tier a legal trigger sets. No downgrade ever goes below a (legal) answer. The "Why" column says which kind of law sits behind it, or why this framework sets the tier higher than the law alone would.

Q1. Data

What's the most sensitive data the app reads, stores or sends? (Classes are defined on [data handling](#). Until that page is rewritten, use the four class definitions in the answers below: its current text is older and classes customer personal data differently.)

Answer	Tier	Why
(a) Public data, or your own work files that contain nothing listed in (c) or (d). "Your own work files" means files you created yourself, or exports you're entitled to use for your own job, that you don't share with anyone and that nobody else relies on. Receipts and statements for your own expenses are own work files even when they show a truncated card number	Sandbox	Below the legal floor for personal data, by design: the approval gate and Floor 1 carry the duty (see how the legal triggers map)

Answer	Tier	Why
(b) Internal: non-public business information, including information that identifies people only by business contact details (name, job title, work email, work phone, employer), including the business contact details of customer, prospective-customer and supplier staff, or by staff work-administration details (role, team, work location, shifts, leave dates but not the reason or type of leave, desk or room bookings, start date, equipment, training completion). The prices on individual invoices, purchase orders and quotes are (b)	Team	(legal) where the information identifies people: privacy laws apply to any information about identifiable people. Non-personal internal data sits here for operational reasons
(c) Confidential: any other information about identifiable people, including everything beyond business contact details about customers, consumers, applicants and patients, and staff HR records (pay, benefits, performance, discipline, leave reasons and leave types, home contact details). This includes the sensitive categories (health or insurance-claim data, bank account numbers, truncated card numbers other than on your own expense receipts, government ID numbers, biometrics, precise location, children's data) and commercially sensitive data (price lists, margins, contracts and their negotiated terms, supplier bank details)	Important	(legal) for the sensitive categories, which carry sector and sensitive-data duties. Truncated card numbers are here for a security reason, not a legal one: breach laws and sensitive-data definitions don't cover them. For the rest, the legal floor is Team; this framework sets Important because these records carry privacy-law duties and are the data in the documented exposure failures (AP1) ⁴
(d) Restricted: other people's sign-in credentials (passwords, PINs, security answers, account access codes) that the app collects or stores; full payment card numbers or card security codes; or patient or health-plan-member health information, if your business is a health-care provider, a health plan, or a vendor handling that information for one (your privacy officer can confirm)	Critical	The legal floor is Important. This framework sets Critical for a security reason: these classes need controls only IT can run

- **Consent gate (a gate, not a tier).** If the app collects **biometric identifiers** (fingerprints, face or voice prints), **data about anyone under 18**, or **consumers' health information or other sensitive-category data listed above**, other than Q1(d) health information, **route to counsel before collecting**. Consent-before-collection laws apply: Illinois's biometric privacy law requires a written release first¹⁴, and Texas and Washington have their own biometric laws; the federal children's online privacy rule requires verifiable parental consent for children under 13¹⁵, and some state laws protect teens too; Washington's My Health My Data Act requires consent before collecting consumer health data¹⁶, and Nevada has a similar law; and most state privacy laws require opt-in consent for consumers' sensitive data.
- The app's own keys and tokens aren't data for this question, but the data they can read is: answer Q1 for everything the token can read. The tokens themselves belong in a secrets manager (never-do 2 of [the five never-dos](#)) and in the registry (M3 below).
- Card data has a better alternative. Rather than holding card numbers, use a tokenized or hosted-payment pattern from your payment provider, so card numbers never reach the app and it doesn't answer Q1(d). Card data brings the connected systems into the card industry's security standard (PCI DSS v4.0.1), a contractual duty¹⁸.

Q2. Users

Who uses it or relies on its output? People who only *receive messages* from it (reminders, notifications, emails, replies) are answered under Q6, not here.

Answer	Tier	Why
(a) Only you	Sandbox	
(b1) People in one team	Team	
(b2) Several teams or the whole company, all inside the business. If that app holds Confidential data or a business process depends on it, Q1(c) or Q4(c) makes it Important; Q2 doesn't need to	Team	Operational. Wider internal use alone doesn't raise the risk; the data and dependence questions catch the apps that do
(c) Anyone outside the business (customers, applicants, suppliers, partners, the public)	Important	(legal) Outside users bring FTC Act claims, accessibility and disclosure duties. This framework sets Important, above the legal Team floor, because external users nearly always add personal data

External users and external messages both give **Important** (Q2(c), Q6(c)), so the two questions never give conflicting tiers for the same fact. People who only receive messages are counted once, under Q6: colleagues who get reminders are Q6(b), not a Q2 answer, and outside recipients are Q6(c).

Q3. Reach

Can someone reach it from the internet without company sign-in?

Answer	Tier	Why
(a) No. It's reachable only after company sign-in (SSO), or only inside the company network	Sandbox	
(b) Yes: a public URL, "anyone with the link" sharing, a public form, or publicly readable storage or database	Important , and it may run only on an approved public-facing hosting pattern (Floor 2)	Public-by-default data is the root cause of the largest documented exposures (AP1) ⁴

Q4. Dependence

What happens if it stops, or gives wrong results, for a day?

Answer	Tier	Why
(a) Only your own work is inconvenienced	Sandbox	
(b) The people who use it work around it	Team	
(c) A business process stops or produces wrong results, or its output feeds a system of record, a report, a reconciliation or a control, but a person re-performs that output or checks it against the system of record before it's used	Important	Operational

Answer	Tier	Why
(d1) Its output is posted to the general ledger or financial statements, or is itself the evidence for an account reconciliation or a control your auditors test, <i>without being independently recomputed</i>	Critical	(legal) Internal control over financial reporting under SOX for SEC registrants ¹⁹ ; audit and lender covenants for private firms
(d2) It directly produces regulatory or legal filings, safety-related actions, or customer transactions with no person checking each output before it takes effect	Critical	Operational

Output a person re-performs or checks against the system of record before use is (c), not (d1). If you're not sure, ask your finance controller before you build: "Is this app's output posted, or kept as reconciliation or control evidence, without anyone recomputing it or checking it against the system of record?" Yes is (d1); no is (c).

Q5. Decisions about people

Does its output decide, rank, score, recommend or flag anything, or produce a score, ranking, flag, summary or recommendation that a person relies on to decide, about a person's job (hiring, firing, promotion, performance rating), pay, discipline, credit, insurance, housing, education, health care, or access to public benefits?

Answer	Tier	Why
(a) No	Sandbox	
(b) Yes, and every outcome is reviewed before it takes effect by a person who meets the meaningful-human-review test	Important , plus the notice, explanation and review controls on apps that make decisions about people	(legal) Automated-decision and anti-discrimination laws (see the mapping table)
(c) Yes, and an outcome can take effect before that review happens	Critical , plus graduation	(legal) The "solely automated" decision tests

This applies whether or not the app uses AI. Calculating an amount under a fixed, documented rule (a commission plan, a published rate card) isn't a decision about a person; answer Q4 and Q6 for it instead.

Q6. Actions

What can it do without a person approving each action?

Answer	Tier
(a) Nothing beyond itself: it reads, calculates, shows, or drafts for a person to act on	Sandbox
(b) It writes to its own data store, or sends messages to colleagues	Team
(c) It writes to a shared system of record (CRM, ERP, HR system, ticketing), or sends automated messages to people outside the business	Important
(d) It moves or commits money, deletes records, changes records in bulk (more than the change cap, which defaults to 100 records per run, a hypothesis in A-028), or takes any action it can't undo	Critical

Don't: tier the app by how it was built

"It was made with AI, so it's high risk" and "it's only a spreadsheet, so it's low risk" are both wrong. The tier comes from what the app does: its data, users, reach, dependence, decisions and actions ([principle 2](#) of the principles behind this framework). In North America, building with AI triggers almost no AI-specific law; the duties follow the app's effects (see [US and Canadian laws](#)). The UK National Cyber Security Centre's June 2026 guidance makes the same point about security: "The risk isn't in using AI. The risk is not applying the right safeguards when the stakes are high."³

Meaningful human review

The reviewer (1) knows how to interpret the app's output, (2) reviews that output, (3) considers other relevant information, and (4) has the authority to change the decision. If any of the four is missing, the review doesn't count, and Q5 is answered (c).

This four-part test is this site's test, adapted from the three-element definition of "human involvement" in the California Privacy Protection Agency's automated decision-making regulations (Cal. Code Regs. tit. 11 §7001(e)(1)(A)–(C))⁶. The phrase "meaningful human review" comes from Colorado SB 26-189⁹. The California definition is the best-specified one in North America, and the test maps onto the GDPR, Quebec, Colorado and UK tests.

California discipline and termination decisions: add corroboration. For a decision to discipline or terminate a worker in California, from 1 Jul 2027 (SB 947's operative date; re-verify before relying on it), part (3) also requires the reviewer to **corroborate** the app's output with the underlying data or other evidence (supervisor evaluations, personnel files, work product, peer reviews, witness interviews), and to record where that evidence is kept. A review without corroboration fails part (3), so Q5 is (c). SB 947 also bars using output that can't be corroborated, and the employee gets a written notice after use¹³. The duty covers discipline and termination only, not hiring. This site's rule is deliberately broader than the statute's own trigger: SB 947 applies when an employer primarily relies on an automated system's output, but this site requires corroboration in every review of such a decision.

The jurisdiction flag

Does the app handle information about people in Quebec, the EU or the UK, or about staff who work in the EU or UK, or is an AI feature's output used in the EU? The EU AI Act reaches providers and deployers outside the EU whose system's output is used there (Art. 2(1)(c))²³. If you don't know, answer yes. If yes, **route to counsel before building**.

Quebec's Law 25 requires a privacy impact assessment for any project to develop an information system that handles personal information (s. 3.3), and another before communicating it outside Quebec, including to a US-hosted AI service (s. 17)¹⁰. The flag never changes the tier; it adds a counsel gate.

Being the *provider of a high-risk AI system* under the EU AI Act isn't a tier trigger. It's rare for worker-built apps, so it's handled by this counsel route and by the [EU AI Act page](#), which recommends graduation if counsel finds the business is a high-risk provider.

The audit-scope flag

A separate routing flag, asked at intake beside the six questions. Lead with the concrete test, because builders can't answer "is it in SOC 2 scope?":

- Does the app produce records used to show a control works? For example **access reviews, change approvals, incident logs, vendor reviews, or onboarding and offboarding records**.
- Does it help deliver a service covered by a SOC 2, ISO 27001 or similar attestation your business gives customers?
- Does it hold **privileged access**: admin rights, or bulk read or write access, to a system that's in audit scope?
- Does it touch financial reporting: data that flows into the ledger, financial statements or a financial control?

If you're not sure, ask your compliance or security lead. Any yes means the app is **routed to the compliance owner** for change-management and evidence requirements. **The flag doesn't change the tier.** The compliance owner's change-management and evidence requirements apply at whatever tier the app is, Team included; an audit-scope app gets no minimum tier, because a floor would only over-tier it. SOC 2 is an attestation against criteria that customer contracts ask for, not a law, and audit scope isn't a risk to people, so it's a gate, not a tier. Output posted to the ledger or kept as audited control evidence without recomputation is still answered under Q4(d1).

If your business has no compliance function, see the fallback under the compliance owner on [roles and decision rights](#).

First question for your compliance owner: which worker-built apps fall inside the system boundary of our SOC 2 or ISO 27001 scope, and does our system description need updating?

The AI modifiers

These aren't extra questions. They add controls or set a minimum tier.

- **M1. AI feature** (the app calls an AI model at runtime): the tier is unchanged, so a personal experiment whose answers are all (a) stays in Sandbox. **Approval gate (every tier):** if the app sends any data to a model or other service that isn't approved **for that data class**, the app is blocked until that service is approved (vendor-contract duties; see [approve tools](#)). In Sandbox, the approval gate and Floor 1 are the only M1 requirements. From Team up, add the AI-feature controls: an eval set with a pass threshold, an AI-feature review, the provider's data terms, and a disclosure check (see [AI features](#)). A tool's built-in AI features (an AI formula in a spreadsheet, an AI step in a workflow tool) count as M1 when they process the app's data at runtime; AI used only to write the app doesn't. The approval gate is satisfied when that built-in feature is already on the [approved-tools list](#) for the data class involved; it needs no separate approval. A built-in feature that isn't listed for that data class is blocked like any other unapproved service.
- **M2. Customer-facing AI:** model output reaches a customer, applicant or member of the public *without a person approving each output* → **Critical**. A draft that a person reads and sends is not customer-facing AI.

- **M3. Autonomous agent or automation** (it acts **without a person approving each action**: it runs on a schedule or trigger, or chains tool calls unattended, under a person's identity or a service identity): the tier is unchanged, but Q1–Q6 are answered for every granted permission. Its identity, OAuth grants and tokens are recorded in the registry under the app owner's role, and they're revoked at retirement or when the owner changes. An interactive coding assistant in a builder's editor, where a person approves each action, isn't itself a worker-built app and isn't M3; the app it helps build is tiered as usual (see [coding agents](#)). Worked through in [agents and automations](#).

The two floors (every tier, no exceptions, no downgrade)

1. **The trifecta rule.** No agent or AI feature may combine private data, untrusted content and a way to send data out (the "lethal trifecta")⁵. Remove one leg before it runs, in Sandbox too (AP6).
2. **No internet reach without company sign-in below Important.** Sandbox and Team apps are private by default. An app reachable from the internet without company sign-in must be approved at Important or Critical on an approved public-facing pattern. A builder tool that can't keep a Sandbox app private isn't approved for Sandbox use (AP1).

[The five never-dos](#) also apply at every tier. Never-do 5 ("Never turn on permission-bypass modes in coding agents") includes classifier-approved modes: a mode in which a classifier or model, not a person, approves each action (often called "auto" mode) isn't human approval, so it counts as a bypass mode. It's allowed only in a disposable sandbox with no credentials to shared systems and no real data.

The floor-breach rule

Sometimes an answer is higher only because the app breaks a floor by mistake. For example, a Team app's storage bucket is publicly readable, which makes Q3 (b) as found. In that case:

1. Contain the breach as soon as it's found (for Floor 2, close the public access), and handle any exposure as an incident ([incident response](#)).
2. Record in the registry the as-found answer, the breach, and a fix-by date no more than 5 working days after the breach was found.
3. Tier the app on its answers as fixed.
4. If the fix isn't confirmed by the fix-by date (for Floor 2, by a clean exposure scan), the as-found answer applies, with every gate of that tier, until the fix is confirmed.

The rule applies only to breaches of the two floors. It doesn't apply when the owner wants to keep the higher answer (for example, to keep a public link), or when an answer comes from something other than a floor, such as an over-broad token or a never-do.

The highest-answer rule

An app's tier is the **highest** tier that any answer or modifier gives it. Never add, average or score answers. One Critical answer makes the app Critical, unless an owner-signed downgrade applies. Re-answer the six questions whenever the app gains users, data, reach, permissions or an AI feature, and at every scheduled review. This is

the high-water-mark principle of NIST's FIPS 199, section 3: when a system holds several types of information, its impact level is the highest of them¹. FIPS 199 dates from 2004 and is still the US federal categorization standard, because the logic doesn't age.

Don't: sum or average the answers

A points score lets a severe factor get averaged away: an app that moves money but has few users and no personal data "scores" medium. The page this one replaces used exactly that kind of five-factor sum. Take the highest answer instead¹.

Don't: accept a self-declared tier

"It's low risk" without the six answers isn't a tier. Risk-calibrated review only works if the low-risk label is checked; a label nobody checks is an opinion, and it's the first thing to drift when a builder is in a hurry. Every tier comes from the six answers, recorded in the registry, and the enablement lead checks them at [intake](#) during triage.

Sandbox rules

- Personal and time-boxed. The default is 90 days (A-029, a hypothesis). On day 80 the builder gets a notice. On day 90 an unrenewed sandbox app is quarantined (access paused, data kept), and 30 days later it's deleted unless it has been registered at a higher tier.
- Register the sandbox workspace once per builder (automatically from the platform where possible), not each app. App-level registration starts the moment any answer rises above Sandbox: the first share, the first data beyond a Q1(a) answer, the first internet reach, or the first write to anything shared.
- Counsel gates in Sandbox: the six questions and the flags aren't answered for each sandbox app, so the builder attests once, at workspace registration, that they'll route to counsel before any sandbox app collects consent-gate data or handles information that would make the jurisdiction flag yes.
- Approved tools only, on business accounts. Both floors and the five never-dos apply.

How the legal triggers map onto the questions

The legal research behind this model proposed seven yes/no legal questions and a jurisdiction flag, each setting a minimum tier. They're folded into the six questions so intake stays short. The principle holds: in North America, building with AI triggers almost no AI-specific law, so intake asks about effects, not about AI (see [US and Canadian laws](#); the EU AI Act is the main exception, through the jurisdiction flag).

Question	What the law cares about	Examples of laws it reflects (as of October 2026)
Q1(b), (c)	Information about identifiable people	State privacy laws, including the CCPA (which covers staff and applicant data); Canada's PIPEDA; Quebec Law 25
Q1(c) sensitive categories, Q1(d), consent gate	Health, financial, card, ID, biometric, children's and credential data	HIPAA ¹⁷ ; the GLBA Safeguards Rule; PCI DSS (contractual) ¹⁸ ; Illinois BIPA ¹⁴ , COPPA ¹⁵ ; Washington's My Health My Data Act ¹⁶ ; state breach-notification laws
Q2(c), Q6(c)	People outside the business	FTC Act claims about the app; accessibility duties; AI-disclosure laws
Q4(d1)	Financial reporting	SOX internal control over financial reporting (registrants) ¹⁹ ; audit and lender covenants (private firms)
Q5(b), (c)	Decisions about people	CCPA automated decision-making rules (compliance from 1 Jan 2027) ⁶ ; California's FEHA automated-decision system regulations (in force since 1 Oct 2025, for California employers with five or more employees; keep automated-decision data for four years) ⁷ ; Colorado SB 26-189 (from 1 Jan 2027) ⁹ ; NYC Local Law 144 ¹¹ ; Illinois HB 3773 (since 1 Jan 2026) ¹² ; California SB 947 (from 1 Jul 2027) ¹³ ; Ontario's rule that a public job posting must disclose AI used to screen, assess or select applicants (postings from 1 Jan 2026, employers with 25 or more employees) ⁸ ; FCRA ²⁰ and ECOA/Regulation B adverse-action notices ²¹ ; Title VII and the ADA; Quebec Law 25 s. 12.1 ¹⁰ ; GDPR Art. 22 where EU people are involved ²²
M1 approval gate	Sending data to an AI service	Vendor-contract and service-provider terms
Jurisdiction flag	Quebec, EU and UK people; AI output used in the EU	Quebec Law 25 ¹⁰ ; GDPR Art. 3 ²² ; the EU AI Act, Art. 2(1)(c) for reach (its high-risk rules are in Chapter III; see the EU AI Act page) ²³ ; the UK Data (Use and Access) Act

Five stated departures from the legal floors. This model departs from the legal research's floors in five places.

Higher: Q1(c) (personal data beyond business contact details and staff work-administration details), Q1(d) (Restricted data) and Q2(c) (any external user), for the security or operational reasons stated in each row. **Lower:** Q1(a) (the builder's own files in a private sandbox) and M1 (a personal AI experiment on an approved service), because the approval gate and Floor 1 carry the duty the legal floors reflect. A legal floor of Team for every AI feature would end sandbox experimentation with AI features (A-002).

Three gates that aren't tiers: the [jurisdiction flag](#) and the Q1 consent gate route to counsel; the [audit-scope flag](#) routes to the compliance owner. None of them changes the tier. The counsel gates hold the build (or, for the consent gate, the collection) until counsel has answered; the audit-scope flag holds go-live until the compliance owner has set the app's requirements. Who counsel is, what to send and how fast to expect an answer are on [the counsel route](#). For the law itself, see [US and Canadian laws](#). The tiers are a way to route review; they don't by themselves meet any legal duty.

The downgrade path (owner-signed and documented)

When a single answer is the only reason for the higher tier, and compensating controls remove that reason, the app owner may ask for one tier lower:

1. Record it as a Decision Record in the registry: the answer being offset, the compensating controls, how each control is checked, and a review date no more than six months away.
2. It's signed by the app owner *and* the approver for the higher tier. From Critical, that's the IT/platform owner and the security reviewer. From Important, it's the enablement lead or the security reviewer. (Standard downgrade 3 has its own signers.) If one person holds both Critical signer roles, see [when one person holds several roles](#).
3. The downgrade is one tier at most, and never below the tier any other answer sets. (Standard downgrade 3 is the only exception to the one-tier limit.)
4. It isn't allowed for the two floors, M2 or the M1 approval gate, and it never goes below a floor marked **(legal)**. Downgrades offset only the answers this model sets above the legal floor for operational or security reasons: Q4(c), Q4(d2), Q6, and Q1(d)'s step from Important to Critical.
5. It lapses automatically if a compensating control fails or the review date passes, and the app returns to the higher tier.
6. The registry reports the share of apps on a downgrade (A-043). If more than about 15% of Important and Critical answers are downgraded, the questions are wrong and get revised. Standard downgrade 3 is reported as a separate count, not in this share. This mirrors control tailoring in NIST SP 800-53B, where baselines are tailored with documented reasons².

Standard downgrades

These three are pre-agreed; anything else needs the full process above.

- **Standard downgrade 1:** bulk writes with a per-run change cap, a pre-run snapshot, a tested restore and a reviewed dry-run diff (Q6(d) → Important).
- **Standard downgrade 2:** health information in Q1(d) on a platform your HIPAA security official has approved for it, with a business associate agreement with every vendor that touches it, the app entered in the HIPAA risk analysis, role-based minimum-necessary access, and audit logging on (Q1(d) → Important)¹⁷.
- **Standard downgrade 3 (tokens the tool can't scope):** an app that has no AI feature (M1) and isn't an agent or automation (M3), whose code calls no write or delete endpoint, which holds the token in a secrets manager, and which passes the technical check, answers Q6 for the actions its code calls, not for the token's full rights. It may drop more than one tier, but never below the tier any other answer sets. It's signed by the app owner and the technical reviewer who did the check. It lapses (rule 5) as soon as the code calls a write or delete endpoint, or the app gains an AI feature or starts running as an agent or automation. It offsets Q6 only, never Q1.
 - **Where the token is held.** A hosting platform's server-side secret or environment-variable store also counts as "a secrets manager" here, if all four hold: the value is never sent to the browser or bundled into client code; access to it is limited to the app's admins; it can be rotated without a code change; and the app record names where it lives (`secrets_location` on the [app record](#)). A secret in source code, client-side configuration or a shared document never qualifies.

Card data and credentials have no standard downgrade.

Don't: downgrade without a signed Decision Record

A downgrade agreed in a chat thread is an exception nobody can find in six months. Without the record, the review date and the second signature, a downgrade becomes a routine bypass. NIST's tailoring guidance expects tailoring decisions to be documented with their rationale². The builder never signs their own downgrade; see [roles and decision rights](#).

Bounded exceptions

An exception changes **when or how a gate is met**, never whether it applies: for example, a runbook due 30 days after go-live, or a negative access test run by hand because the platform can't automate it. It's the only exception path; anything else goes through the downgrade path or not at all. An exception is allowed only with all four of:

1. A named approver: the person accountable for changing a gate's timing or method in [the RACI](#) (the enablement lead for Team and Important apps; the IT/platform owner, with the security reviewer, for Critical apps). The builder never approves their own exception.
2. A written reason: why the gate can't be met as written, and what covers the risk meanwhile.
3. An expiry of 90 days or less. On expiry the gate applies as written. A further exception needs a new approval and a new reason; it isn't renewed automatically.
4. An entry in the app record in the [app registry](#): the gate, the approver, the reason and the expiry.

No exceptions to the floors or the legal triggers. An exception never applies to the two floors, the five never-dos, M2, the M1 approval gate, any answer marked (legal), the counsel gates (the jurisdiction flag and the consent gate), or data-owner sign-off, and it never changes a tier or an answer.

Critical means co-ownership, with a capacity check

A Critical app keeps its business app owner, and IT or engineering co-owns it through [graduation](#). If IT can't take a Critical app within the graduation target, the app's scope is capped (no new users, data or permissions) until it can (A-033). We expect fewer than 10% of registered apps to land in Critical (A-028), and 35% or less to land in Important or Critical (A-045). If more do, review the questions before adding IT or review capacity.

Hypothesis to test

We expect the share of registered apps in Important or above to be 35% or less with this wording (A-045). The model was revised in October 2026 because a review estimated that the earlier wording put 60–70% of registered apps there, a reviewer's estimate rather than a measurement. Test it by reporting the tier mix from the registry each quarter of your pilot.

Gate matrix

What each tier needs at each lifecycle stage. This is the only copy on the site; other pages link to a stage anchor (for example [#gate-matrix-verify](#)). Download it as [CSV](#). A "+" means "everything in the previous column, plus".

The **business review** (Important and above) is done by any colleague who knows the process. It covers the logic, the data used and whether the access matches the intent. The **technical check** is an automated exposure scan and the data-rule negative tests, plus a 30–45 minute check against a fixed list (data rules or row-level security, no secrets in the client, sharing settings, sign-in). It's done by a **technical reviewer**: a platform or IT engineer, or a champion who has been trained and listed as a technical reviewer. The target wait is 5 working days (A-042). If no technical reviewer is available, the app stays at Team scope. How to run both is on [business review and technical check](#); the approved hosting patterns are on [hosting](#).

Intake

Gate	Sandbox	Team	Important	Critical
Register	sandbox workspace only	app record	app record, plus data-owner sign-off for Confidential or Restricted data	+ a graduation plan
Answer the six questions and the flags	not needed until any answer rises above Sandbox	yes, at intake	yes, at intake	yes, at intake
Counsel gates	attested once at sandbox workspace registration: route to counsel before collecting consent-gate data (biometric identifiers, data about anyone under 18, consumers' health or other sensitive data) or handling jurisdiction-flag information	consent gate or jurisdiction flag yes: route to counsel before collecting or building	same	same
Audit-scope flag	—	yes: route to the compliance owner	same	same
Kill criteria and retire-by date	90-day time-box	set at intake	set at intake	set at intake

Build

Gate	Sandbox	Team	Important	Critical
Tools and accounts	approved tools only, on business accounts	same	same	same, on IT's standard platform
Floors and never-dos	both floors and the five never-dos	same	same	same
AI feature (M1)	approval gate for the data class, and Floor 1	+ eval set with a pass threshold, AI-feature review, provider data terms, disclosure check	same	same; M2 makes the app Critical

Gate	Sandbox	Team	Important	Critical
Agents and automations (M3)	answer for every granted permission	+ identity, grants and tokens recorded under the app owner's role	same	same

Verify

Gate	Sandbox	Team	Important	Critical
Tests	—	the builder's own test cases pass before every release	+ negative access tests	+ independent tests
Review	verification self-check (recommended)	verification self-check, recorded	+ business review + technical check; + accessibility (WCAG 2.2 AA) and AI-disclosure check if people outside the business use it (Q2(c)) or receive messages it sends or drafts	+ security review

Deploy

Gate	Sandbox	Team	Important	Critical
Hosting	private workspace	approved pattern, SSO	approved pattern, SSO, deny-by-default data rules, separate production	IT's standard platform
Internet reach (Floor 2)	not allowed	not allowed	only on an approved public-facing pattern	only on an approved public-facing pattern
Documentation	none required	one-page app record	+ runbook	+ IT's standard documentation
Go-live	—	go-live checklist items for Team	checklist items for Important	checklist items for Critical

Operate

Gate	Sandbox	Team	Important	Critical
Support level	self-supported	team-supported	assisted: a named backup who has run the app, help-desk routing from the registry, a runbook	IT-supported, normal service levels
Scheduled review	at 90 days	every 12 months	every 6 months	every 6 months, plus quarterly access review
Graduation	—	when a trigger fires	when a trigger fires	required: IT or engineering co-owns; scope capped if IT can't take it in time

Gate	Sandbox	Team	Important	Critical
Downgrade in force	—	re-checked by its review date (six months at most)	same	same
Exception in force	—	a bounded exception to a gate's timing or method: named approver, written reason and app-record entry; expires in 90 days or less, then the gate applies as written	same	same
Floor breach	contain, record the as-found answer and a fix-by date (5 working days at most); as-found answer applies if not fixed in time	same	same	same

Change

Gate	Sandbox	Team	Important	Critical
Re-tier	register the app when any answer rises above Sandbox	re-answer the six questions when the app gains users, data, reach, permissions or an AI feature; when its output starts feeding the ledger or control evidence without recomputation (Q4(d1)); when a feature starts ranking, scoring or flagging people (Q5); when the app owner changes; or when an agent's grants change	same	same
Checks per release	—	the tier's tests and review, in version history you can roll back; tests run in full, and only the human review may focus on what changed until a full re-review trigger fires	same	same, through IT's change process
Audit scope	—	if the audit-scope flag is yes: the compliance owner's change-management and evidence requirements	same	same

Per-release review depth. The matrix sets *which* reviews apply to a release. Some always run in full: the builder's own tests on every release from Team, the negative access tests from Important, and the automated part of the technical check (the exposure scan and the data-rule negative tests). Only the human business review and the reading of the diff or settings may focus on what changed, until a **full re-review trigger** fires: a tier change, a new data class, a new AI feature or agent grant, new external users, a new owner, or 12 months since the last full review. Then every review runs in full.

See [change a worker-built app safely](#).

Retire

Gate	Sandbox	Team	Important	Critical
Trigger	day 90 unrenewed: quarantine (access paused, data kept); deleted 30 days later unless registered at a higher tier	retire-by date passed, unused, owner gone with no successor, or replaced	same	same, with IT as co-owner

Gate	Sandbox	Team	Important	Critical
Revoke	the app's access and data	+ identity, OAuth grants and tokens (M3)	same	same
Record	—	registry status updated	same	same

See [retire a worker-built app safely](#) for data disposal, retention and legal holds.

The matrix by tier

Sandbox

- Intake: register the sandbox workspace once, with the one-time counsel-gate attestation; 90-day time-box.
- Build: approved tools on business accounts; both floors; the five never-dos; the M1 approval gate for the data class.
- Verify: verification self-check recommended.
- Deploy: private workspace only; no internet reach; no documentation required.
- Operate: self-supported; review at 90 days.
- Change: register the app as soon as any answer rises above Sandbox.
- Retire: quarantine on day 90 if unrenewed; delete 30 days later.

Team

- Intake: app record; six questions and flags answered; kill criteria and retire-by date.
- Build: approved tools; both floors; the never-dos; AI-feature controls if M1; agent grants recorded if M3.
- Verify: the builder's own test cases pass before every release; verification self-check recorded.
- Deploy: approved pattern with SSO; no internet reach; one-page app record.
- Operate: team-supported (builder plus a champion); review every 12 months.
- Change: re-tier on new users, data, reach, permissions or an AI feature, ledger or control-evidence output (Q4(d1)), scoring or flagging people (Q5), a new owner or changed agent grants; tests run in full on every release.
- Retire: revoke access, grants and tokens; update the registry.

Important

- Intake: app record plus data-owner sign-off for Confidential or Restricted data; counsel and compliance gates as they apply.
- Build: as Team.
- Verify: + negative access tests; business review and technical check; accessibility and AI-disclosure check if people outside the business use it or receive what it sends or drafts.
- Deploy: approved pattern, SSO, deny-by-default data rules, separate production; internet reach only on an approved public-facing pattern; + runbook.
- Operate: assisted support (a named backup who has run the app, help-desk routing, a runbook); review every 6 months.
- Change: as Team; downgrades re-checked by their review date, and exceptions by their expiry.
- Retire: as Team.

Critical

- Intake: as Important (including data-owner sign-off for Confidential or Restricted data), plus a graduation plan.
- Build: IT's standard platform.
- Verify: + independent tests and a security review.
- Deploy: IT's standard platform and documentation.
- Operate: IT-supported at normal service levels; review every 6 months plus a quarterly access review; IT or engineering co-owns.
- Change: through IT's change process.
- Retire: as Team, with IT as co-owner.

Agents and automations

M3 means you answer the six questions for every permission an agent or automation has been granted, not for the job it was built to do.

Illustrative composite

An operations analyst builds a flow in a SaaS agent builder. Every 15 minutes it reads the shared supplier-inquiries mailbox, uses an AI step to classify each email, and opens a ticket. It runs under the analyst's own OAuth consent, which granted full mailbox access, including send and delete.

As granted: Q1(c), because the mailbox holds supplier bank details. Q2(b1): only the operations team uses the tickets. Q3(a). Q4(b). Q5(a). Q6(c) for writing to the ticketing system, and Q6(d) because the grant can delete mail. M1 applies, and the trifecta is complete: private data, untrusted inbound email, and a way to send data out. **Result: Critical, and Floor 1 is broken, so it can't run.** The floor-breach rule doesn't lower this tier: Critical comes from the over-broad grant (Q6(d)), not from a floor, and a Floor 1 breach means the flow stops until a leg is removed.

As rebuilt: a service identity recorded under the app owner's role, with read-only mailbox access, ticket creation only, and no send. The trifecta loses its send leg, so Floor 1 holds. Q6 drops to (c), and Q1(c) remains. M1 needs an AI service approved for Confidential data, plus the Team-and-up AI-feature controls.

Result: Important. The grants are revoked when the owner changes or the flow is retired (see [when the builder leaves](#)).

For the agent-specific controls, see [MCP and agents](#).

Worked test cases

Use these to calibrate whoever triages intake. The field guides must match these tiers. All rows are illustrative composites.

App (illustrative composite)	Answers that set the tier	Tier
Invoice matching , week 1: the analyst's own CSV exports from the ERP (which they're entitled to use for their job), with invoice and PO prices but no supplier bank details, used only by them, in a private workspace	every answer (a): the exports are own work files, and invoice prices are not (c)	Sandbox
Invoice matching after teammates adopt it: internal invoice data with invoice and PO prices, no bank details	Q1(b) (invoice prices are (b)), Q2(b1), Q6(b)	Team
Invoice matching once it holds supplier bank details and pre-sorts exceptions for the AP clerks, while the ERP still performs the three-way match the auditors test	Q1(c) supplier bank details; Q2(b2) AP plus the controller's team; Q4(c): the clerks work each exception in the ERP, which re-performs the match, so the output isn't control evidence	Important. It never edits supplier master data, and bank-detail changes need dual approval in the ERP
Invoice matching if the controller retired the ERP match and relied on the app as the documented control	Q4(d1) (legal): its output becomes the evidence for a control the auditors test, with no independent recomputation; no downgrade is possible	Critical , plus graduation. The guide records the decision not to do this
HR onboarding tracker as first specified, holding scanned ID documents, ID-check results and home addresses	Q1(c) government ID numbers, HR records, home contact details	Important
The tracker after minimization: start dates, equipment and training completion only; ID documents, ID-check results and home addresses stay in the HR system	Q1(b) staff work-administration details; Q2(b2) HR plus hiring managers; Q4(b): if it stops, HR works from the HR system; Q6(b)	Team (minimization lowers the tier)
The tracker's proposed AI flag of new hires "at risk of failing probation", which managers would see before deciding	Q5(c) (legal) as proposed: nothing makes each manager's review meet all four parts of the meaningful-human-review test, so the review doesn't count (with a documented review that meets all four, and SB 947 corroboration for California staff, it would be Q5(b), Important); M1; the employment-AI duties apply (notice, a CCPA risk assessment for inferring employee traits, Illinois HB 3773, California's FEHA regulations)	Critical as proposed, with a counsel gate. Declined as a worker-built feature and moved to an HR-system project
Ticket triage with AI drafts: customer tickets, the IT operations team, a person sends each reply	Q1(c) customer personal data; Q2(b1) the IT operations team (customers only receive replies a person sends, so they're not a Q2 answer, but the accessibility and AI-disclosure check applies because outside people receive what it drafts); Q6(c) it writes classifications to the ticketing system; M1; M2 not triggered; trifecta removed by design	Important
Ticket triage, if it sent replies automatically	M2	Critical

App (illustrative composite)	Answers that set the tier	Tier
Sales-ops CRM sync as first designed: a nightly bulk overwrite by an agent-built job on the analyst's admin token	Q1(c): the admin token can read every CRM record, including price lists and negotiated contract terms; Q6(d) bulk changes without per-record approval	Critical
CRM sync as rebuilt: a scoped service credential that reads and writes only contacts' business contact fields, two fields written, a 100-record cap per run, a pre-run snapshot, a tested restore, a reviewed dry-run diff	Q1(b) business contact details of customer staff (the service credential is the app's own token, so not Q1(d), and it can read only those fields); standard downgrade 1 of Q6(d); Q6(c) remains; jurisdiction flag answered yes because contacts' locations are unknown (counsel gate, not a tier)	Important (downgraded, owner-signed)
Facilities reminder scheduler : staff names and desk bookings, the facilities team, reminders to colleagues	Q1(b) staff work-administration details, Q2(b1) the facilities team (staff who get reminders are Q6, not Q2), Q6(b); Q3(b) as found, from a public bucket that broke Floor 2, so the floor-breach rule applies: the bucket is closed, and the as-found answer and fix-by date are recorded	Team as fixed. It was never registered
An event sign-up form for external business guests on an approved forms tool, behind no sign-in, collecting only name, employer and work email	Q1(b) business contact details; Q2(c) (so the accessibility check applies); Q3(b) on an approved public-facing pattern	Important
A commission calculator under a fixed plan, used by sales ops and finance; finance checks each figure against the CRM bookings before posting the accrual	Q1(c) staff pay; Q2(b2); Q4(c): checked against the system of record before posting, so not (d1); Q5 doesn't apply (fixed rule)	Important (if finance posted its figures unchecked, Q4(d1) would make it Critical)
A company-wide desk-booking app: staff names and desk bookings, everyone books their own desk	Q1(b), Q2(b2), Q4(b), Q6(b)	Team
A team leave calendar: leave dates only, no reasons or leave types	Q1(b) leave dates, Q2(b1), Q6(b)	Team
A personal receipt-reading expense helper: the builder's own receipts (with truncated card numbers), output to their own sheet	Q1(a) own expense receipts; Q2(a); Q6(a); M1 if the reading is done by an AI service (approval gate and Floor 1 only)	Sandbox
A personal dashboard that reads one SaaS tool's project data through a personal API token the tool can't scope (the token can also delete)	Q1(b), Q2(a); Q6(d) for the token's full rights; standard downgrade 3 (no AI feature, not an agent, code calls only read endpoints, token in a secrets manager, technical check passed) makes Q6 (a)	Team with standard downgrade 3 (Critical without it)

When to re-tier

Re-answer the six questions and both flags when any of these happens:

- the app gains users, data, internet reach, permissions or an AI feature;
- its output starts being posted to the ledger or kept as control evidence without being independently recomputed (Q4(d1));
- a feature starts ranking, scoring or flagging people, or producing a summary or recommendation someone relies on to decide about them (Q5);
- a Sandbox app's answer to any question rises above (a): for example it's first shared, first uses data beyond a Q1(a) answer, is first reachable from the internet, first writes to anything shared, or first sends messages;
- a scheduled review falls due (see [Operate](#));
- a downgrade's compensating control fails or its review date passes (the app returns to the higher tier automatically);
- a floor breach isn't fixed by its fix-by date (the as-found answer applies);
- the app owner changes, or an agent's grants change.

Resolving disputes

Most disputes are about one answer, not the tier. Settle the answer and the tier follows.

1. Re-read the exact wording of the disputed answer together. Most disputes end here.
2. While it's unresolved, record the higher answer. Don't start work that the lower tier would allow and the higher tier wouldn't.
3. Escalate to the approver for the higher tier (the enablement lead or the security reviewer for Important; the IT/platform owner and the security reviewer for Critical) and aim to decide within five working days.
4. If the question's wording caused the dispute, log it. Repeated disputes on the same wording are a signal to revise the question through a Decision Record, not to interpret it differently each time.

Don't: rewrite the questions in your own forms

Paraphrased questions drift, and drifted questions stop mapping to exactly one tier. Copy the wording on this page into your intake form and registry. Change it only by recording a decision.

Questions to ask your counsel

1. Which of our data is regulated in a way that sets Q1(d), and do we collect biometric identifiers, children's or teens' data, or health or other sensitive consumer data outside HIPAA anywhere, so that consent is needed before collecting?
2. Does any decision our apps support fall under an automated-decision or employment-AI rule (the CCPA rules, California's FEHA regulations, Colorado SB 26-189, NYC Local Law 144, Illinois HB 3773, California SB 947, Ontario's job-posting disclosure rule)?
3. Does our business count as a HIPAA covered entity or business associate for any line of business, and who is our security official?
4. Which of our financial processes and controls do our auditors test (the Q4(d1) list), and which worker-built apps fall inside our SOC 2 or ISO 27001 system boundary?

5. Do our customer, staff or contact lists include Quebec, EU or UK residents, and is any AI feature's output used in the EU?

Record the decision

Record your adoption of this model once, as a program-level Decision Record, using the [decision record template](#).

Illustrative composite

Title: Adopt the four-tier model for worker-built apps · **Status:** accepted · **Date:** YYYY-MM-DD · **Owner:** enablement lead (a role) · **Sponsor:** the executive sponsor named in the leadership charter

Context. Staff build apps and automations with AI tools. We have no shared way to decide what each one needs.

Decision. Use the four tiers, the six tier questions, the jurisdiction and audit-scope flags, the consent gate, the two floors, the floor-breach rule and the downgrade path exactly as written on this page. Downgrades are signed by the roles named in the downgrade path.

Alternatives considered. A points-based risk score (rejected: severe factors get averaged away). Three tiers (rejected: no place for shared apps on internal data that don't need assisted support). No tiers, with every shared app reviewed by IT (rejected: slow intake drives unsanctioned use). Do nothing (rejected: apps already exist with no owner or tier).

Consequences. Every shared app gets a tier, an owner and a review date. We measure the tier mix, the share of Critical apps and downgrades quarterly.

Review by: three months from the date above, or sooner if more than about 10% of registered apps are Critical, more than about 35% are Important or Critical, or more than about 15% of Important and Critical answers are downgraded.

Assumptions and evidence

- **A-018** (Medium): we expect four tiers set by the highest of six answers to be enough to route support, review and graduation; test this by counting apps that don't fit a tier, and tier disputes per month, in your pilot.
- **A-045** (Low): we expect 35% or less of registered apps to land in Important or above with this wording; test this from the registry's tier mix each quarter.
- **A-008** (Low): we expect builders to recognize a consequential decision about a person when asked in plain words; test this by auditing a sample of Q5(a) answers.
- **A-028** (Low): we expect fewer than 10% of registered apps to land in Critical, and the 100-record change cap to be workable; test this from registry counts each quarter.
- **A-029** (Low): we expect a 90-day sandbox time-box to be about right; test this by tracking how many sandbox apps are renewed rather than registered or retired.

- **A-033** (Low): we expect IT to absorb graduations at the rate the Critical triggers fire; test this by watching the graduation queue for two quarters.
- **A-042** (Low): we expect technical-check capacity to keep up; test this by measuring the median wait against five working days.
- **A-043** (Low): we expect owner-signed downgrades to stay rare and honest; test this by reporting the downgrade share (and the standard downgrade 3 count separately) and checking for lapsed controls.
- **A-002** (Low): we expect builders to keep using an approved sandbox for AI experiments when it isn't gated by a Team floor; test this by comparing sandbox use with unsanctioned AI use.

Evidence: published government policy and security-agency guidance, vendor guidance on support tiers, a durable federal categorization standard, and statutes and regulations; no controlled study of tier models for worker-built apps. Tiering by impact, with any-of triggers, appears in Saskatchewan's May 2026 low-code policy²⁴ and in vendor guidance on tiered support²⁵. Facts on this page are as of October 2026.

Related

- [Get a worker-built app approved](#): the intake form that asks these questions.
- [Check an AI-built app before release](#): how the verify gates work in practice.
- [Graduate a worker-built app to IT](#): what happens when an app is Critical.
- [Apps that make decisions about people](#): the controls behind Q5.
- [Risk tier workshop for existing apps](#): tier a batch of discovered apps in 90 minutes.
- [Can I use company data in an AI app?](#): the data classes behind Q1.

-
1. NIST, "FIPS 199: Standards for Security Categorization of Federal Information and Information Systems", Feb 2004, section 3 (high-water mark). <https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.199.pdf> [primary], accessed 2026-10-21.
 2. NIST, "SP 800-53B: Control Baselines for Information Systems and Organizations", Sep 2020 (minor release Aug 2025); cited for tailoring only. <https://csrc.nist.gov/pubs/sp/800/53/b/upd1/final> [primary], accessed 2026-10-21.
 3. UK National Cyber Security Centre, blog post on the "vibe coding spectrum" approach to AI-assisted software development, 18 Jun 2026. <https://www.ncsc.gov.uk/blogs/the-vibe-coding-spectrum-approach-to-ai-assisted-software-development> [regulator], accessed 2026-10-21.
 4. UpGuard research on Power Apps portals that exposed about 38 million records to anonymous users because table permissions were off, 23 Aug 2021. <https://upguard.com/breaches/power-apps> [vendor], accessed 2026-10-21.
 5. Simon Willison, blog post defining the "lethal trifecta" for AI agents, 16 Jun 2025. <https://simonwillison.net/2025/Jun/16/the-lethal-trifecta/> [research], accessed 2026-10-21.
 6. California Privacy Protection Agency, approved text of the CCPA regulations on automated decision-making technology, risk assessments and cybersecurity audits (effective 1 Jan 2026; ADMT compliance by 1 Jan 2027), §7001(e). https://coppa.ca.gov/regulations/pdf/ccpa_updates_cyber_risk_admt_appr_text.pdf [regulator], accessed 2026-10-21.
 7. California Civil Rights Council, release on approval of regulations on employment discrimination and automated-decision systems (effective 1 Oct 2025; four-year retention of automated-decision data), 30 Jun 2025. <https://calcivilrights.ca.gov/2025/06/30/civil-rights-council-secures-approval-for-regulations-to-protect-against-employment-discrimination-related-to-artificial-intelligence/> [regulator], accessed 2026-10-21.
 8. Government of Ontario, "Your guide to the Employment Standards Act: requirements related to publicly advertised job postings" (AI disclosure for postings from 1 Jan 2026; O. Reg. 476/24). <https://www.ontario.ca/document/your-guide-employment-standards-act-0/requirements-related-publicly-advertised-job-postings> [regulator], accessed 2026-10-21.

9. Colorado General Assembly, SB 26-189 (repeals and reenacts SB 24-205; effective 1 Jan 2027). <https://leg.colorado.gov/bills/sb26-189> [statute], accessed 2026-10-21.
10. Quebec, Act respecting the protection of personal information in the private sector, CQLR c. P-39.1, ss. 3.3, 12.1 and 17. <https://www.legisquebec.gouv.qc.ca/en/document/cs/P-39.1> [statute], accessed 2026-10-21.
11. NYC Department of Consumer and Worker Protection, "Automated Employment Decision Tools". <https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page> [regulator], accessed 2026-10-21.
12. Illinois Public Act 103-0804 (HB 3773), amending the Illinois Human Rights Act, effective 1 Jan 2026. <https://www.ilga.gov/legislation/publicacts/fulltext.asp?Name=103-0804> [statute], accessed 2026-10-21.
13. California SB 947 (reported as Stats. 2026, ch. 859), Labor Code §§1520–1526.7. https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=202520260SB947 [statute; the page refused our tools, so the scope, the 1 Jul 2027 operative date and the chapter number come from law-firm summaries only, secondary], accessed 2026-10-21.
14. Illinois Biometric Information Privacy Act, 740 ILCS 14. <https://www.ilga.gov/legislation/ilcs/ilcs3.asp?ActID=3004&ChapterID=57> [statute], accessed 2026-10-21.
15. Children's Online Privacy Protection Rule, 16 C.F.R. Part 312 (amended Apr 2025; compliance by 22 Apr 2026). <https://www.law.cornell.edu/cfr/text/16/part-312> [statute], accessed 2026-10-21.
16. Washington My Health My Data Act, RCW 19.373 (consent before collecting consumer health data, RCW 19.373.030). <https://app.leg.wa.gov/RCW/default.aspx?cite=19.373.030> [statute], accessed 2026-10-21.
17. HIPAA Security Rule administrative safeguards, 45 C.F.R. §164.308 (risk analysis; security official). <https://www.law.cornell.edu/cfr/text/45/164.308> [statute], accessed 2026-10-21.
18. PCI Security Standards Council, document library (PCI DSS v4.0.1). https://www.pcisecuritystandards.org/document_library/ [primary], accessed 2026-10-21.
19. Sarbanes-Oxley Act §404, 15 U.S.C. §7262. <https://uscode.house.gov/view.xhtml?req=granuleid:USC-prelim-title15-section7262&num=0&edition=prelim> [statute], accessed 2026-10-21.
20. Fair Credit Reporting Act, 15 U.S.C. §1681m (adverse-action notices). <https://uscode.house.gov/view.xhtml?req=granuleid:USC-prelim-title15-section1681m&num=0&edition=prelim> [statute], accessed 2026-10-21.
21. CFPB, Regulation B, 12 C.F.R. §1002.9 (statement of specific reasons). <https://www.consumerfinance.gov/rules-policy/regulations/1002/9/> [regulator], accessed 2026-10-21.
22. Regulation (EU) 2016/679 (GDPR), Arts. 3 and 22. <https://eur-lex.europa.eu/eli/reg/2016/679/oj> [statute], accessed 2026-10-21.
23. Regulation (EU) 2024/1689 (EU AI Act), Art. 2(1)(c) (providers and deployers in third countries whose system's output is used in the EU). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj> [statute], accessed 2026-10-21.
24. Government of Saskatchewan, "Low-Code/No-Code Development and Maintenance Policy", May 2026. <https://taskroom.saskatchewan.ca/-/media/project/taskroom/documents/policy/lowcodenocodedevelopmentandmaintenancepolicy.pdf> [primary], accessed 2026-10-21.
25. Microsoft, "Support strategy for solutions" (Power Platform adoption guidance), May 2025. <https://learn.microsoft.com/en-us/power-platform/guidance/adoption/support-strategy-solutions> [vendor], accessed 2026-10-21.

Prompt to Prototype, Risk tiers for worker-built apps, <https://prompt-to-prototype.com/operating-model/risk-tiers/>, as of 2026-10-22. Licensed [CC BY-SA 4.0](#).