CONFIDENTIAL COMPLIANCE DOCUMENT

US AI Compliance — Complete

Prepared exclusively for Cedar Line Creative

US Jurisdiction · 11 August 2026

CONTENTS

What is in this pack

Each document below is also attached to your delivery email as its own PDF and editable Word file, so you can circulate or sign any one of them on its own.

PART 1

Your AI Compliance Profile

Cedar Line Creative is a 51-250 employee marketing and advertising agency headquartered in Denver, Colorado, using generative AI tools (ChatGPT, Claude, Midjourney, DALL-E) to produce copy and imagery published under client brands, an AI chat assistant deployed on client websites, an AI-powered CRM, and an AI screening tool in its own hiring pipeline. The agency processes personal data through AI and makes AI-assisted decisions about people, including job applicants, triggering obligations under Colorado SB 26-189 (effective 1 January 2027), California SB 942 (in force 2 August 2026), Texas HB 149 (in force 1 January 2026), Illinois BIPA, EEOC AI-in-employment guidance, and FTC Section 5. Cedar Line has no AI disclosure language anywhere and no analysis of which state laws apply to its multi-state footprint, placing it in active non-compliance with Texas HB 149 and California SB 942 today, and at risk of being unprepared for Colorado SB 26-189 when it commences. Unresolved intellectual property exposure from publishing AI-generated client content, combined with the FTC's first enforcement action targeting undisclosed AI-generated advertising and its July 2026 proposed policy statement on AI output suppression, make remediation urgent.

Maximum Fine Exposure

Maximum fine exposure is calculated separately per regime because Cedar Line triggers multiple concurrent enforcement regimes. (1) FTC Section 5: at $51,744 per violation per day, a single undisclosed AI-generated advertising campaign running for 30 days across multiple client brands could generate $51,744 × 30 = $1,552,320 per campaign; Cedar Line's portfolio of AI-generated client content multiplies this across every distinct deceptive act or practice identified. (2) Illinois BIPA: if the AI recruitment screening tool uses any biometric-capable feature (video, voice, facial analysis) and processed, for example, 100 Illinois applicants without written consent, intentional-violation exposure is $5,000 × 100 = $500,000, plus litigation costs and class-action multiplier. (3) Virginia CDPA: at $7,500 per intentional violation, systematic failure to provide a privacy notice or data protection assessment across multiple Virginia consumers could aggregate rapidly; arithmetic requires Cedar Line's count of Virginia data subjects. (4) EEOC/Title VII: compensatory and punitive damages are capped at $200,000 per complainant for an employer of Cedar Line's size (51-200 employees); if the AI screening tool produces adverse impact across a hiring cohort, each affected applicant is a separate potential complainant. (5) Texas HB 149 and Colorado SB 26-189: statutory per-violation civil penalty maxima should be confirmed with Texas and Colorado counsel respectively against the enacted texts; turnover-based penalty limbs under any regime cannot be computed without Cedar Line's annual revenue figures. Aggregate worst-case exposure across all active regimes — assuming only the arithmetic that can be calculated without revenue figures — exceeds $2,250,000 before Colorado SB 26-189 commences on 1 January 2027, at which point hiring-related exposure increases materially. This figure is a floor, not a ceiling.

AI Tools Inventory

Tool	Use	Risk Level	Notes
ChatGPT / OpenAI	Drafting marketing copy and communications content published under client brands	**High**	Content published under client brands with no AI disclosure violates California SB 942 (latent provenance disclosure required from large generative AI providers; Cedar Line as deployer must ensure compliance in client-facing output) and Texas HB 149 disclosure duties. FTC's first action on undisclosed AI-generated advertising directly analogous. IP

Tool	Use	Risk Level	Notes
			ownership of outputs is unresolved; OpenAI terms grant usage rights but copyright protection of AI-only content is unsettled post-2025 US case law. Client contracts should address AI-generated content ownership and indemnification.
Claude / Anthropic	Drafting marketing copy and communications content published under client brands	**High**	Same disclosure and IP exposure as ChatGPT. Anthropic's June 2025 fair-use ruling covers training but not deployer liability for published outputs. No disclosure wording in place for client-brand content. Cedar Line should confirm whether Anthropic's provenance/watermarking tools satisfy California SB 942 latent disclosure requirements for content Cedar Line publishes.
AI image generation tools (Midjourney, DALL-E)	Creating imagery published under client brands	**High**	AI-generated images published without disclosure breach California SB 942 and Texas HB 149. Copyright in AI-only images is currently unprotectable under US Copyright Office guidance; Cedar Line cannot assert copyright on behalf of clients for purely AI-generated images. Training data copyright litigation risk is elevated (80+ US lawsuits tracked as of February 2026). The federal Take It Down Act (effective May 2026) imposes removal obligations if any client campaign involves realistic imagery of real individuals. Midjourney and DALL-E terms must be reviewed for commercial use and indemnity scope.
AI-powered CRM or sales tools	Data analysis, business intelligence, and customer relationship management	**Limited**	If CRM profiles individuals in California (clients, leads, contacts), Virginia CDPA and California consumer privacy obligations may apply. Colorado SB 26-189 (effective 1 January 2027) will require deployer notices and impact assessments if CRM outputs influence consequential decisions about Colorado residents. Confirm whether the CRM uses personal data for profiling that meets 'consequential decision' thresholds. Vendor DPA and sub-processor agreements should be reviewed now.
AI in HR or recruitment systems (AI screening tool)	Screening and evaluating job applicants for Cedar Line's own hiring	**High**	The AI screening tool is the highest-risk deployment. EEOC guidance requires validation for adverse impact across race, sex, national origin, and disability before use. Colorado SB 26-189 (1 January 2027) classifies hiring as a 'consequential decision' requiring deployer impact assessments, applicant notice, and opt-out/appeal rights. Texas HB 149 (in force now) imposes prohibited-use rules and requires transparency for AI in employment contexts. Illinois BIPA applies if the tool captures biometric identifiers (voiceprints, facial geometry) from Illinois applicants. No disclosure to applicants currently exists, which is independently a violation of EEOC guidance and will breach Colorado SB 26-189 on commencement.

Applicable Regulations

Texas HB 149 — Responsible Artificial Intelligence Governance Act

In force 1 January 2026. Applies to developers and deployers of AI systems used in Texas or affecting Texas residents. Requires deployers to implement reasonable AI governance policies, maintain transparency about AI use in consequential contexts, observe prohibited-use categories, and cooperate with Attorney General enforcement. Cedar Line deploys AI in client-facing chatbots and advertising content reaching Texas consumers, and uses AI in hiring that may affect Texas-based applicants, all without any disclosure or governance documentation. No safe harbour currently exists for Cedar Line's undocumented posture.

Key date: 1 January 2026 (in force) · Max penalty: Attorney General enforcement; civil penalties per violation — confirm current statutory maximum with Texas counsel as the statute is newly operative and penalty tiers should be verified against enacted text.

California SB 942 — AI Transparency Act (as amended by AB 853)

In force 2 August 2026. Requires large generative AI providers to offer a free AI detection tool and to embed latent provenance disclosures in AI-generated content. As a deployer publishing AI-generated copy and imagery under client brands for California consumers, Cedar Line must ensure that content produced by covered providers (OpenAI, Anthropic, Midjourney/DALL-E) carries required latent disclosures and that Cedar Line's own practices do not strip or obscure them. The FTC's July 2026 proposed policy statement warns that suppressing AI-output accuracy markers to circumvent state laws like SB 942 could independently constitute consumer deception under Section 5. Cedar Line currently has no disclosure wording anywhere.

Key date: 2 August 2026 (in force) · Max penalty: California Attorney General and private right of action; statutory penalties to be confirmed against enacted text — turnover-based limb cannot be computed without Cedar Line's revenue figures.

Colorado SB 26-189 — Automated Decision-Making Technology (developer and deployer duties)

Not yet in force; effective 1 January 2027. Repealed and reenacted SB 24-205, which never commenced. Imposes duties on deployers of automated decision-making technology used in 'consequential decisions' — expressly including employment and hiring — affecting Colorado residents. Cedar Line's AI hiring screening tool and any AI-assisted consequential decisions about employees or applicants in Colorado will require: algorithmic impact assessments, pre-deployment bias audits, plain-language notices to affected individuals, and meaningful opt-out or appeal mechanisms. Cedar Line must also manage developer disclosures from OpenAI, Anthropic, and HR-tool vendors. Preparation should begin now given the January 2027 commencement.

Key date: 1 January 2027 (not yet in force) · Max penalty: Colorado Attorney General enforcement; civil penalties per violation — confirm statutory maximum with Colorado counsel; turnover-based limb cannot be computed without Cedar Line's revenue figures.

Illinois BIPA — Biometric Information Privacy Act

Applies where Cedar Line's AI tools collect, use, or store biometric identifiers or biometric information (including voiceprints and facial geometry scans) from Illinois residents. Relevant if the AI recruitment screening tool uses video interview analysis, voice analysis, or facial recognition for Illinois-based applicants, or if the AI chat assistant captures biometric data from Illinois users. BIPA requires written informed consent before collection, a publicly available retention and destruction policy, and prohibits profiting from biometric data. No consent mechanism or policy is currently in place.

Key date: In force (long-standing; confirm current amendment status with Illinois counsel) · Max penalty: $1,000 per negligent violation; $5,000 per intentional or reckless violation; private right of action with litigation costs. With a workforce of 51-250 and potential class exposure for applicants, aggregate exposure could be substantial — exact figure requires headcount of Illinois applicants processed through biometric-capable tools.

Virginia CDPA — Consumer Data Protection Act

Applies to personal data of Virginia residents processed by Cedar Line (e.g., via client-website AI chatbots, CRM, or advertising targeting). Requires a privacy notice, data processing agreements with processors, a data protection assessment for processing personal data for targeted advertising or profiling that presents a reasonably foreseeable risk of harm, and consumer rights (access, deletion, opt-out of profiling). Cedar Line's AI chatbot on client websites likely collects personal data from Virginia visitors; no disclosure is in place.

Key date: In force (confirm current amendment status with Virginia counsel) · Max penalty: Up to $7,500 per intentional violation; Virginia Attorney General enforcement; no private right of action.

NIST AI RMF 1.0

Voluntary framework providing a structured approach — GOVERN, MAP, MEASURE, MANAGE — for identifying, assessing, and mitigating AI risks across the AI lifecycle. Not legally binding but referenced by FTC and EEOC as indicative of reasonable AI governance practice. Cedar Line's absence of any AI governance documentation, disclosure wording, or bias-testing regime is inconsistent with even baseline RMF implementation. Adopting RMF as an internal governance standard would materially reduce regulatory and litigation risk and support a due-diligence defence in FTC or EEOC proceedings.

Key date: Published January 2023; no mandatory commencement date · Max penalty: Not applicable — voluntary framework; non-adoption is not directly penalised but absence of RMF-equivalent governance is an aggravating factor in regulatory proceedings.

FTC Section 5 — Unfair or Deceptive Acts or Practices (AI enforcement)

The FTC's dedicated AI enforcement unit (established January 2026) has brought its first action targeting undisclosed AI-generated advertising content and has pursued 13 AI-washing cases since 2024. A proposed policy statement issued 1 July 2026 warns that suppressing or altering AI outputs to obscure their AI origin — including to avoid state disclosure laws — may independently constitute consumer deception. Cedar Line publishing AI-generated copy and imagery under client brands with no disclosure, and operating a client-facing AI chatbot with no AI identification, presents direct Section 5 exposure. The FTC's jurisdiction covers unfair or deceptive practices in or affecting commerce; Cedar Line's multi-state client base brings it squarely within scope.

Key date: Ongoing enforcement; proposed policy statement 1 July 2026 · Max penalty: Civil penalties up to $51,744 per violation per day (2026 inflation-adjusted figure — confirm current figure with counsel); injunctive relief; restitution. Aggregate exposure across multiple client campaigns is uncapped without Cedar Line's transaction volume.

EEOC AI in Employment Guidance

The EEOC's guidance on employer use of AI in hiring and employment decisions confirms that Title VII, the ADA, and the ADEA apply fully to AI-assisted selection tools. Employers — not vendors — bear liability for adverse impact caused by AI screening tools. Cedar Line must validate its AI screening tool for disparate impact across protected classes before use, document that validation, and ensure the tool does not screen out qualified individuals with disabilities who could perform the job with reasonable accommodation. No validation, documentation, or applicant disclosure is currently in place.

Key date: Guidance issued; Title VII/ADA/ADEA enforcement ongoing · Max penalty: Title VII: compensatory and punitive damages up to $200,000 for employers with 51-200 employees (Cedar Line falls in this band for most of its range); back pay, reinstatement, and attorney fees. ADA: same caps. Pattern-or-practice suits by EEOC carry uncapped equitable relief.

AI Acceptable Use Policy

Your formal AI policy. Sign it, circulate it to every member of staff, and review it annually or whenever your AI tool usage changes.

AI Acceptable Use Policy

Cedar Line Creative LLC
Headquartered: Denver, Colorado
Effective Date: 11 August 2026
Policy Owner: Chief Operating Officer
Version: 1.0

This policy is binding on all Cedar Line Creative personnel. It is issued to meet obligations arising under Colorado SB 26-189 (in force 1 January 2027), Texas HB 149 (Responsible Artificial Intelligence Governance Act, in force 1 January 2026), California SB 942 as amended by AB 853 (AI Transparency Act, in force 2 August 2026), Illinois BIPA, the Virginia CDPA, FTC Section 5 of the FTC Act, EEOC guidance on AI in employment, and the NIST AI Risk Management Framework 1.0. Cedar Line Creative operates across Colorado, California, Texas, and Illinois; all four state regimes therefore reach this business.

1. Scope

This policy applies to every person who performs work for or on behalf of Cedar Line Creative, including full-time and part-time employees, contractors, freelancers, secondees, and interns, regardless of where they are located. It governs any use of artificial intelligence tools in connection with Cedar Line Creative's business activities, whether those activities occur on company premises, at client sites, or remotely.

The activities covered include, without limitation: drafting, editing, or translating marketing copy, campaign materials, scripts, or communications content; generating, editing, or curating images, video, audio, or other creative assets; operating or configuring AI-powered chat assistants deployed on client websites; using AI to screen, rank, shortlist, or evaluate job applicants or existing employees; using AI-powered CRM or sales intelligence tools to analyse customer data or prioritise leads; and conducting any data analysis or business intelligence work using AI.

The tools currently in scope are: ChatGPT and other OpenAI products; Claude and other Anthropic products; AI image generation tools including Midjourney and DALL-E; the AI-powered CRM and sales tools currently deployed by the business; and the AI screening tool used in Cedar Line Creative's hiring process. Any tool not listed in Section 2 as approved must not be used for any Cedar Line Creative business purpose until it has been approved through the process set out in that section.

2. Approved Tools and the Process for Requesting a New Tool

The following tools are currently approved for use within the scope described in Section 1:

— ChatGPT / OpenAI products: approved for content drafting, ideation, editing, and summarisation.
— Claude / Anthropic products: approved for content drafting, editing, research synthesis, and summarisation.
— Midjourney and DALL-E: approved for generating and iterating on creative imagery for client campaigns, subject to the IP verification requirements in Section 4.
— The AI-powered CRM and sales tool currently licensed by Cedar Line Creative: approved for lead scoring, pipeline analysis, and communications personalisation within the data restrictions in Section 3.
— The AI applicant screening tool currently licensed by Cedar Line Creative for HR use: approved for initial screening of job applications, subject to the mandatory human review requirements in Section 5.

No other AI tool may be used for Cedar Line Creative business purposes. This prohibition includes free-tier or personal accounts on any AI platform, browser-based AI assistants, AI plug-ins within productivity software, and any tool a client may request staff to use, unless that tool has first been approved by the Policy Owner.

To request approval of a new tool, the requesting employee or manager must submit a written request to the Policy Owner setting out: the tool name and provider; the intended use case; the data types that will be processed; the provider's data retention and sub-processing terms; and a brief assessment of IP ownership and output rights under the provider's current terms of service. The Policy Owner will review the request, conduct or commission a risk assessment aligned with the NIST AI RMF 1.0 GOVERN and MAP functions, and provide a written decision within ten business days. No tool may be used pending approval.

3. Prohibited Inputs — What Must Never Be Entered Into an AI Tool

The following categories of information must never be entered, pasted, uploaded, or otherwise submitted to any AI tool, whether approved or not, unless Cedar Line Creative has put in place a written data processing agreement with that provider that expressly covers the category in question and legal review has confirmed the transfer is lawful:

Personal data of customers or staff. This includes names, email addresses, telephone numbers, postal addresses, IP addresses, device identifiers, employment records, salary information, performance reviews, health information, demographic information, and any other information that identifies or is reasonably capable of identifying a living individual. Cedar Line Creative processes personal data subject to the California Consumer Privacy Act as amended, the Virginia CDPA, Illinois BIPA (which specifically governs biometric identifiers and information), and the Colorado Privacy Act. Feeding such data into a third-party AI tool without a lawful basis and appropriate contractual protections is a breach of those statutes and of this policy.

Confidential client material. This includes client briefs, strategy documents, media plans, budget information, campaign performance data, audience research, and any other material that a reasonable person would understand to be commercially sensitive or that is identified as confidential in any client contract or NDA. Cedar Line Creative's client relationships and reputation depend on strict confidence.

Material subject to a non-disclosure agreement. Where Cedar Line Creative has entered an NDA with any counterparty, the scope of that NDA automatically extends to AI tools. Staff must check whether an NDA is in place before submitting any project-related material to an AI tool.

Credentials and access tokens. Usernames, passwords, API keys, OAuth tokens, two-factor authentication codes, and similar access credentials must never be submitted to any AI tool under any circumstances.

Unpublished financial information. Revenue figures, cost data, pricing models, financial forecasts, and M&A-related; information must not be submitted.

Biometric data. Given Cedar Line Creative's Illinois-based staff and the strict requirements of Illinois BIPA (740 ILCS 14), no biometric identifier — including facial geometry derived from photographs, voiceprints, or fingerprint data — may be submitted to any AI tool without a written release from the individual concerned and prior written approval from the Policy Owner.

Where a member of staff is uncertain whether particular material falls within a prohibited category, they must treat it as prohibited and seek guidance from the Policy Owner before proceeding.

4. Verification of Output and Responsibility for Errors

AI tools produce outputs that can be factually incorrect, legally problematic, internally inconsistent, or infringing of third-party rights. Cedar Line Creative publishes AI-assisted copy and imagery under client brands. A failure of quality or accuracy in that context is a failure that reaches the client's customers and carries reputational and legal consequences for both Cedar Line Creative and its clients.

Every piece of AI-generated or AI-assisted content — whether text, image, code, data summary, or any other form — must be reviewed and verified by a qualified human being before it is used externally, submitted to a client, published on any platform, or transmitted to any third party. Verification means checking factual claims against authoritative sources, confirming that no third-party copyright material has been reproduced without licence, confirming that images do not replicate identifiable individuals without consent, and confirming that the output complies with applicable disclosure rules (see Section 6).

For AI-generated imagery, the responsible employee must specifically assess whether the output could infringe the style, composition, or protected expression of an identifiable artist or brand, in light of the evolving US copyright litigation landscape. Where doubt exists, the image must not be used and the matter must be escalated to the Policy Owner and, if necessary, to legal counsel. Cedar Line Creative accepts that copyright protection for AI-generated outputs without sufficient human authorship remains legally unsettled under current US Copyright Office guidance; staff must not represent AI-generated imagery or text as being copyright-protected by Cedar Line Creative or its clients unless human creative contribution has been documented.

The employee who approves and signs off AI-assisted work before it is used externally bears primary responsibility for any error or non-compliance in that output. The fact that an AI tool produced the content is not a defence. Line managers bear secondary responsibility for ensuring that verification steps have been completed within their teams. The Policy Owner is responsible for maintaining this policy and for systemic failures of the verification framework.

5. Where a Human Decision Is Mandatory

Cedar Line Creative uses AI in contexts that directly affect people's interests. In those contexts, reliance on AI output alone is prohibited; a qualified human being must make the substantive decision, with AI output treated as one input among several.

Hiring and recruitment. Cedar Line Creative's AI screening tool may be used to organise, prioritise, or flag applications, but no candidate may be rejected, advanced, or shortlisted based solely on AI-generated output. A trained human recruiter or hiring manager must independently review all AI-generated scores or recommendations before any decision is communicated to a candidate. This is required by EEOC guidance on the use of AI in employment, which makes clear that employers remain liable under Title VII, the ADA, and the ADEA for discriminatory outcomes produced by AI tools, and by Colorado SB 26-189 (effective 1 January 2027), which will impose deployer duties — including impact assessments and candidate notice — for AI used in consequential employment decisions. Cedar Line Creative is implementing compliant processes now so that operations are fully aligned on the day that law takes effect. Texas HB 149 additionally requires that AI systems used by employers do not employ prohibited uses as defined by that Act, and Cedar Line Creative's use of its screening tool has been reviewed against that standard.

Performance management. AI-assisted analysis of employee performance data must not result in a disciplinary action, performance improvement plan, promotion decision, or compensation change unless a line manager and an HR professional have independently assessed the underlying evidence and made a documented human judgement.

Client-facing chatbot interactions. The AI chat assistant operated on client websites may handle routine informational queries autonomously. However, any interaction that involves a complaint, a request for a refund or contract variation, a safeguarding concern, a vulnerable customer indicator, or a query that the system cannot resolve with high confidence must be escalated immediately to a human agent. The escalation pathway must be clearly available to every user at every stage of the interaction.

Data analysis used to inform significant client strategy. AI-generated data analysis that will be presented to a client as the basis for a significant budget allocation, campaign pivot, or strategic recommendation must be reviewed and validated by a senior strategist or account director before presentation. AI output in this context is a drafting aid, not a finished deliverable.

6. Disclosure — When Cedar Line Creative Tells Customers That AI Was Used

Cedar Line Creative has no existing disclosure wording. This section creates that wording and the rules for applying it, effective immediately.

California SB 942, as amended by AB 853 and in force since 2 August 2026, requires large generative AI providers to make free detection tools available and to embed latent provenance disclosures in AI-generated content. Cedar Line Creative, as a deployer of those tools, must not strip, obscure, or suppress those latent disclosures when publishing AI-generated content. Where Cedar Line Creative publishes AI-generated imagery or text on behalf of clients who have California-based audiences, the published content must retain any provenance metadata embedded by the generating tool. Staff must not use post-processing steps that knowingly remove such metadata.

Texas HB 149, in force since 1 January 2026, requires disclosure where AI is used in contexts involving government entities; Cedar Line Creative must confirm with each client whether the client is or serves a Texas government entity and apply the required disclosure if so.

For Cedar Line Creative's own commercial communications and for content Cedar Line Creative publishes under client brands, the following disclosure standard applies as a matter of company policy regardless of whether a specific statute mandates it in a given context. Any externally published content that was materially generated by an AI tool — meaning the substantive text, imagery, or structure originated with the AI rather than a human author — must carry a disclosure. The disclosure must be clear and conspicuous, not buried in a footer or a terms page. The approved form of words is: "This content was created with the assistance of artificial intelligence tools." Where the client's brand guidelines prohibit that wording, the Policy Owner must approve an alternative before publication.

The AI chat assistant operated on client websites must identify itself as an automated system at the start of every conversation and whenever a user sincerely asks whether they are speaking with a human. The FTC's enforcement position under Section 5 of the FTC Act treats the failure to disclose AI interaction as a potentially deceptive practice. Cedar Line Creative will not expose itself or its clients to that risk.

Cedar Line Creative must not make marketing claims about its own AI capabilities that are unsubstantiated. The FTC has pursued AI-washing enforcement actions since 2024; any external claim about how Cedar Line Creative uses AI must be accurate, specific, and capable of being substantiated by internal records.

7. Record Keeping and Retention for AI-Assisted Work

Cedar Line Creative must be able to demonstrate, to a regulator, insurer, or client, that AI was used appropriately on any given project. To support that demonstration, the following records must be created and retained.

For every piece of AI-assisted content that is published externally or delivered to a client: the name of the AI tool used; the date and, where the tool provides it, a reference to the session or job; the name of the employee who carried out verification; the nature of the verification steps taken; and the name of the person who approved the content for external use.

For every hiring decision involving the AI screening tool: the name of the tool; the scoring or ranking output it produced for each candidate considered; the name of the human reviewer who assessed that output; the human reviewer's documented assessment; and the final decision with reasons. These records must be retained for a minimum of four years from the date of the hiring decision, consistent with EEOC record-keeping guidance and the potential application of Colorado SB 26-189 to employment processes from 1 January 2027.

For the AI chat assistant on client websites: session logs must be retained in accordance with the data retention terms agreed with the relevant client and in compliance with applicable state privacy laws. Where no client agreement specifies a period, logs must be retained for twelve months and then securely deleted unless a legal hold is in place.

For AI-generated imagery: a record of the prompt used, the tool used, the date, and the verification outcome must be retained for the duration of the client engagement and for three years thereafter, given the ongoing uncertainty in AI copyright litigation.

All records created under this section must be stored in a designated location within Cedar Line Creative's document management system, accessible to the Policy Owner and to any person authorised by the Policy Owner for compliance review purposes. Records must be protected against unauthorised deletion or amendment.

8. Breach of This Policy — Reporting and Disciplinary Position

Any employee who becomes aware of a breach of this policy, or of a suspected breach, must report it to the Policy Owner within 24 hours of becoming aware. Reports may be made by email to the Policy Owner's direct address or, where the Policy Owner is the subject of the suspected breach, to the Chief Executive Officer. Cedar Line Creative will not penalise any employee who makes a good-faith report.

Upon receiving a report, the Policy Owner will assess the nature and severity of the breach within two business days. Where the breach involves personal data — for example, personal data submitted to an AI tool without a lawful basis — the Policy Owner must assess whether it constitutes a notifiable incident under applicable state privacy law. Illinois BIPA provides for a private right of action and statutory damages of $1,000 per negligent violation and $5,000 per intentional or reckless violation; the risk of BIPA class actions is material given Cedar Line Creative's Illinois-based staff and clients. Colorado SB 26-189, once in force on 1 January 2027, gives the Colorado Attorney General enforcement powers over deployers who fail to meet their obligations. Texas HB 149 provides for Attorney General enforcement. California SB 942 is enforceable by the California Attorney General. The statutory maximum fines under each of these regimes apply; Cedar Line Creative's turnover-based exposure under each regime cannot be computed without reference to its financial figures, but the per-violation exposure under BIPA alone makes prompt breach response a commercial priority.

Breaches of this policy will be addressed through Cedar Line Creative's standard disciplinary procedure. The severity of the disciplinary outcome will reflect the nature of the breach, whether it was deliberate or negligent, whether the employee reported it voluntarily, and whether harm resulted. Minor or inadvertent first breaches may result in a written warning and mandatory retraining. Serious, repeated, or deliberate breaches — including the submission of confidential client data or personal data to an unapproved AI tool, the suppression of required disclosures, or the falsification of verification records — may result in summary dismissal. Nothing in this policy limits Cedar Line Creative's right to take legal action where a breach causes loss to the business or to a client.

9. Review Cycle, Policy Owner, and Signature Block

This policy must be reviewed in full by the Policy Owner no later than 11 August 2027 and annually thereafter. An out-of-cycle review must be triggered immediately upon any of the following events: the addition of a new AI tool to Cedar Line Creative's operations; a material change in the way an existing approved tool is used; a breach of this policy that reveals a gap in its coverage; a change in applicable law, including the commencement of Colorado SB 26-189 on 1 January 2027 and any further amendments to California SB 942, Texas HB 149, or Illinois BIPA; or a material change in EEOC guidance or FTC enforcement policy relating to AI.

The Policy Owner is the Chief Operating Officer of Cedar Line Creative. The Policy Owner is responsible for maintaining this policy, approving new tools, receiving breach reports, overseeing record keeping, and ensuring that all staff complete AI awareness training upon joining and upon each revision of this policy.

This policy takes effect on 11 August 2026. Every member of staff must sign the acknowledgement below and return it to the Policy Owner within five business days of receiving this policy. New joiners must sign before they are given access to any approved AI tool.

ACKNOWLEDGEMENT OF RECEIPT AND UNDERSTANDING

I confirm that I have read, understood, and agree to comply with Cedar Line Creative's AI Acceptable Use Policy (Version 1.0, dated 11 August 2026). I understand that breach of this policy may result in disciplinary action up to and including dismissal.

Full name (print): ______________________________

Job title: ______________________________

Signature: ______________________________

Date: ______________________________

POLICY OWNER SIGN-OFF

Approved and issued by the Chief Operating Officer, Cedar Line Creative LLC.

Signature: ______________________________

Date: 11 August 2026

Employee AI Guidelines — Practical Dos and Don'ts

The working-level guidance that sits under the policy: concrete examples of good and bad AI use in this business.

Introduction and Scope

These guidelines apply to every Cedar Line Creative employee, contractor and freelancer who uses an AI tool in the course of their work. They sit beneath the Cedar Line Creative AI Policy and translate its obligations into day-to-day behaviour. They cover ChatGPT/OpenAI, Claude/Anthropic, Midjourney, DALL-E, the AI-powered CRM and sales tools, and the AI screening tool used in our own hiring process.

Cedar Line Creative operates across Colorado, California, Texas and Illinois, publishes AI-generated copy and imagery under client brands, runs AI chat assistants on client websites, and uses AI screening in recruitment. Each of those activities attracts specific legal obligations — including California SB 942 (in force 2 August 2026), Texas HB 149 (in force 1 January 2026), Colorado SB 26-189 (in force 1 January 2027), Illinois BIPA, EEOC guidance on AI in employment, and FTC Section 5 enforcement. Non-compliance is not theoretical: the FTC brought its first action targeting undisclosed AI-generated advertising content and has pursued thirteen AI-washing cases since 2024. These guidelines tell you what to do and what not to do so that Cedar Line Creative and its clients stay on the right side of every one of those frameworks.

Worked Examples and Do / Do-Not Pairs by AI Use

The following pairs are drawn directly from Cedar Line Creative's declared AI activities. Each pair names a concrete action that is required or permitted, and a concrete action that is prohibited.

Marketing and communications content (ChatGPT, Claude, Midjourney, DALL-E)

DO: When drafting social-media copy or long-form content for a client using ChatGPT or Claude, include a note in the deliverable file confirming that AI tools were used in production. Under California SB 942, large generative AI providers must embed latent provenance disclosures in outputs; you must not strip or obscure those signals before handing work to a client or publishing it under a client brand. Where a client has California-based consumers, add a visible human-readable disclosure in the content brief or publication metadata, worded to satisfy SB 942's transparency requirement.

DO NOT: Submit AI-generated copy to a client as if a human copywriter wrote every word, or represent in a proposal that creative output is entirely human-authored when AI tools contributed. The FTC treats undisclosed AI-generated advertising content as a deceptive practice under Section 5 of the FTC Act, and Texas HB 149 prohibits material deception in AI-assisted communications. A single client complaint that triggers an FTC inquiry could expose Cedar Line Creative and the client brand simultaneously.

DO: When generating imagery in Midjourney or DALL-E for use in client campaigns, retain the prompt, the generation log and the final selected image in the project folder. This creates the audit trail needed to respond to intellectual property queries and to demonstrate provenance if a rights dispute arises.

DO NOT: Use a client's proprietary brand assets, trademarked visual elements or confidential campaign briefs as prompt inputs into a public AI image tool. Public model interfaces may incorporate your inputs into training pipelines. Doing so risks disclosing client confidential information and could constitute a breach of the client contract.

AI chat assistant on client websites (Customer service / chatbot)

DO: Before a client's AI chat assistant goes live, confirm that the deployment configuration includes a clear, upfront disclosure to end users that they are interacting with an automated system and not a human agent. Texas HB 149 requires disclosure of AI involvement in communications with consumers, and FTC Section 5 treats the absence of such disclosure as deceptive where consumers would reasonably assume they are speaking to a person.

DO NOT: Configure or allow a client's chatbot to claim to be a named human employee, deny being an AI when directly asked, or make factual representations about a client's products or services that have not been reviewed and approved by a human subject-matter expert at the client. If the chatbot produces a confident but incorrect statement — for example, a wrong price, a false product specification or an unsubstantiated claim about results — and that statement reaches a consumer, FTC liability attaches to the deployer as well as the developer.

HR, recruitment and performance management (AI screening tool)

DO: Before using the AI screening tool to rank, score or filter job applicants, document the tool's selection criteria, confirm those criteria are job-relevant, and retain records of every applicant who was screened out by the tool. The EEOC's guidance on AI in employment makes clear that an employer cannot escape disparate impact liability by attributing a discriminatory outcome to an AI vendor. Colorado SB 26-189, effective 1 January 2027, will impose formal developer and deployer duties for automated decision-making in consequential decisions including employment; Cedar Line Creative must be compliant before that date.

DO NOT: Rely solely on the AI screening tool's output to reject a candidate without a human reviewer examining the result. Do not use the tool to screen for proxies that correlate with protected characteristics — for example, excluding candidates from particular zip codes or schools in ways that produce racially or ethnically disparate outcomes. Illinois BIPA applies to any biometric data the tool processes in connection with Illinois-based applicants or employees; if the tool uses facial analysis, voice analysis or any biometric identifier, obtain explicit written consent before processing and confirm that the vendor's data-retention and deletion obligations are met.

Data analysis and business intelligence (AI-powered CRM and sales tools)

DO: When using the AI-powered CRM to generate client or prospect insights, segment audiences or forecast sales, treat the output as a starting hypothesis that requires human validation before it informs a business decision. Record what data set was fed into the model so that any finding can be traced and, if needed, corrected.

DO NOT: Feed personally identifiable information about clients, prospects or employees into a public AI interface (such as a free tier of ChatGPT) to generate CRM insights. Data submitted to non-enterprise API endpoints may not carry the contractual data-protection protections that Cedar Line Creative requires. Use only the approved enterprise or API-tier tools listed in the AI Policy, and confirm that data-processing agreements are in place with each vendor.

Prompting Practice: Protecting Personal and Client Data

The way you write a prompt is the first line of defence against data exposure. Every prompt you send to a cloud-based AI tool is, in effect, a transmission of data to a third-party system. The following rules apply to every employee using any Cedar Line Creative AI tool.

Use fictional stand-ins rather than real data when testing a prompt or exploring how a tool responds. For example, if you want to check how Claude structures a performance-review summary, use a made-up employee name and invented job title — not a real colleague's details. If you are building a prompt template for a client campaign, use placeholder demographics rather than actual consumer records pulled from the CRM.

Never include in a prompt: a client's full legal name alongside confidential strategic information (such as an unannounced product launch or a forthcoming merger); any individual's Social Security number, date of birth, financial account details or health information; any biometric data; or any personal data about a job applicant beyond what the AI screening tool is expressly configured to receive. These categories attract the highest risk under Illinois BIPA, California SB 942's downstream provisions, and FTC Section 5.

Before submitting a prompt that references a real person — for example, drafting a bio or a testimonial — confirm that the person has consented to their information being processed by AI tools. This applies equally to client contacts, named consumers in case studies and Cedar Line Creative employees. Where consent is uncertain, remove identifying details from the prompt and reintroduce them manually into the final human-reviewed draft.

If you are using the enterprise API tier of ChatGPT or Claude under a signed data-processing agreement, you may include client data only to the extent permitted by that agreement and only for the purpose for which consent was given. Check the agreement before you include anything beyond generic business descriptions.

Checking Output for Fabrication, Bias and Confident Errors

AI language models produce fluent, confident text even when the underlying information is wrong. This is not a bug that will be fixed in a future update; it is a structural characteristic of how these models work. Every employee at Cedar Line Creative is responsible for the accuracy of work they submit, regardless of whether an AI tool produced the first draft.

For factual claims — statistics, product specifications, legal statements, pricing, named individuals' roles or qualifications — verify each claim against a primary source before including it in a deliverable. Do not treat the AI's citation of a source as proof the source exists or says what the model claims. Open the cited source and read it. AI models have been observed fabricating plausible-sounding journal articles, case names and regulatory provisions.

For bias, apply particular scrutiny to any AI output that characterises or segments people. Outputs generated by the AI-powered CRM or the recruitment screening tool may reflect historical patterns in the training data that disadvantage women, ethnic minorities, older workers or other protected groups. When reviewing a ranked list of candidates or a targeted audience segment, ask whether the pattern makes sense on job-relevant or business-relevant grounds alone. If a segment or shortlist is unexpectedly homogeneous, escalate rather than proceeding. The EEOC has stated clearly that disparate impact from an AI tool is the employer's legal exposure, not the vendor's alone.

For content produced by Midjourney or DALL-E, review imagery for stereotyped representations — for example, consistently depicting a professional role as male and white, or a service role as female and non-white — before including it in a client campaign. Such outputs, if published under a client brand, create reputational and regulatory risk for both Cedar Line Creative and the client.

For confident errors — statements that are grammatically polished and tonally assured but factually wrong — maintain a particular watch on outputs about regulatory deadlines, competitor claims and technical product capabilities. These are areas where AI models are most likely to produce plausible but incorrect content, and where the FTC's prohibition on unsubstantiated advertising claims is most directly engaged.

Handling AI Output That Will Be Seen by a Customer

Any AI-generated content that will be published, broadcast, displayed or transmitted to a customer — whether under a client brand or Cedar Line Creative's own — must pass through a human review gate before it goes out. This is not optional and it is not a matter of quality preference; it is a legal requirement in several of the jurisdictions in which Cedar Line Creative operates.

For written copy and imagery: a named human employee must read or view the final output, confirm that all factual claims have been verified, confirm that no personal data appears that was not intentionally and lawfully included, and approve the disclosure wording. California SB 942, in force since 2 August 2026, requires that AI-generated content distributed to consumers carries provenance information enabling detection of AI origin. Cedar Line Creative must not publish content that strips or suppresses that provenance data. Where a client operates in California, the publication metadata or accompanying disclosure must satisfy SB 942. Texas HB 149, in force since 1 January 2026, imposes comparable transparency obligations for AI-assisted consumer communications.

For the AI chat assistant running on client websites: before a new topic area, persona configuration or response template is deployed, a human reviewer at Cedar Line Creative must test the chatbot's responses to foreseeable consumer questions in that area and confirm they are accurate and appropriately qualified. The reviewer's name and the date of review must be logged. If the chatbot is updated — including by the underlying model provider pushing a version change — the review process repeats.

Disclosure wording must appear before or at the point at which a consumer first interacts with AI-generated content or an AI system. Cedar Line Creative currently has no disclosure wording in place anywhere. This must be remedied immediately. The legal risk of continued non-disclosure includes FTC enforcement, Texas HB 149 Attorney General enforcement, and, for California-connected content, SB 942 enforcement. Draft disclosure language must be approved by Cedar Line Creative's legal adviser before use, and each client must be informed of the disclosure requirement and given a copy of the wording being used under their brand.

When to Stop and Escalate to a Person

There are circumstances in which an employee must stop using an AI tool and refer the matter to a manager, to the designated AI Compliance Lead, or to Cedar Line Creative's legal adviser. These are not suggestions; they are mandatory escalation triggers.

Escalate immediately if: an AI tool produces output that makes a specific legal, medical or financial representation that you cannot verify and that will be seen by a client or consumer. Do not attempt to rephrase the claim and publish it; flag it and await human expert review.

Escalate immediately if: the AI screening tool produces a candidate shortlist or rejection recommendation that appears to screen out a disproportionate number of applicants sharing a protected characteristic (race, sex, age, disability, national origin). Document what you observed, do not action the recommendation, and notify the HR lead and AI Compliance Lead.

Escalate immediately if: a client requests that Cedar Line Creative remove, suppress or alter the provenance metadata from AI-generated content to avoid disclosure to their customers. Complying with that request would expose Cedar Line Creative to liability under California SB 942 and FTC Section 5. The request must be referred to a senior account lead and the AI Compliance Lead.

Escalate immediately if: any AI tool produces content involving a real, identifiable private individual in a context that could be defamatory, sexually explicit, or that the individual has not consented to. This includes AI-generated imagery. The Take it Down Act, effective May 2026, criminalises nonconsensual publication of intimate AI-generated images; broader reputational and tortious liability attaches to other categories.

Escalate immediately if: you are uncertain whether a particular use of personal data in an AI prompt is covered by Cedar Line Creative's data-processing agreements with its vendors. Do not proceed on the assumption that it is covered; the default is to stop and check.

Escalate immediately if: a consumer or job applicant asks whether AI was used in a decision that affected them and you do not have a prepared, approved answer. Under EEOC guidance and, from 1 January 2027, Colorado SB 26-189, individuals have rights to explanation and in some cases to appeal automated decisions affecting employment. Cedar Line Creative must be able to answer those questions accurately.

The AI Compliance Lead can be reached at the contact details set out in the AI Policy. Outside normal hours, the escalation contact is the relevant senior manager on call. When in doubt, stop and ask.

Intellectual Property and Ownership of AI Outputs

Cedar Line Creative publishes AI-generated copy and imagery under client brands. This creates a specific intellectual property risk that every employee involved in content production must understand.

Current US copyright law does not protect purely AI-generated content. A human creative contribution — in the selection of outputs, the arrangement of elements, or the creative direction of the prompt — may attract copyright protection in that human contribution, but the AI-generated portion itself does not. This means that content produced entirely by Midjourney or DALL-E with minimal human creative input cannot be registered as a copyright work and may not receive full copyright protection if infringed.

Before Cedar Line Creative delivers AI-generated content to a client under a contract that assigns or warranties copyright, the account lead must confirm with the AI Compliance Lead what level of human creative input was applied and whether the resulting work meets the threshold for copyright protection. Do not represent to a client in a contract or otherwise that Cedar Line Creative owns full copyright in an output that was generated without sufficient human authorship.

There is active and unresolved litigation in the US regarding whether AI training on copyrighted material constitutes infringement. Over eighty AI copyright lawsuits were being tracked in the US as of early 2026. Cedar Line Creative employees must not upload third-party copyrighted images, text or music into AI tools as reference material or style inputs without confirming that the use is licensed or constitutes fair use. If in doubt, do not upload it.

Retain all prompt logs, generation records and the selection history for AI-produced deliverables for a minimum of three years, or longer if a client contract requires it. These records are the evidence base if a copyright or provenance dispute arises.

AI Tool Inventory

The live register of every AI tool in the business. Regulators, insurers and enterprise customers all ask for this first.

Tool	Vendor	Business purpose	Data it touches	Personal data (Yes/No)	Risk level	Owner	Date reviewed
ChatGPT (including GPT-4 and successors)	OpenAI	Drafting marketing copy, communications content, and creative briefs published under client brands.	Client briefs, campaign data, staff prompts. May include client customer details if pasted into prompts.	Yes — if client personal data or staff data is included in prompts.	MEDIUM — IP ownership risk on outputs; FTC Section 5 exposure if AI origin undisclosed in advertising; California SB 942 provenance duties apply to generative outputs distributed in CA.	Head of Content / Designated AI Lead	11 Aug 2026
Claude	Anthropic	Drafting marketing copy, research summaries, and internal communications.	Client briefs, campaign data, staff prompts. Same personal-data risk as ChatGPT if client data is pasted in.	Yes — conditional on prompt content.	MEDIUM — same IP, FTC, and California SB 942 considerations as ChatGPT. Colorado SB 26-189 deployer duties apply from 1 Jan 2027 if outputs inform consequential decisions.	Head of Content / Designated AI Lead	11 Aug 2026
Midjourney	Midjourney Inc.	Generating images and visual assets published under client brands.	Text prompts; generated imagery. No direct personal data unless prompts reference individuals.	Yes — if prompts reference real individuals or client customer likenesses.	HIGH — IP ownership of AI imagery is legally unsettled; FTC first enforcement action targeted undisclosed AI-generated advertising content; California SB 942 latent-provenance disclosure required for	Head of Design / Designated AI Lead	11 Aug 2026

Tool	Vendor	Business purpose	Data it touches	Personal data (Yes/No)	Risk level	Owner	Date reviewed
					covered generative outputs distributed in CA from 2 Aug 2026.		
DALL-E	OpenAI	Generating images and visual assets published under client brands; integrated via ChatGPT or API.	Text prompts; generated imagery. Same personal-data risk as Midjourney.	Yes — conditional on prompt content.	HIGH — same IP, FTC, and California SB 942 considerations as Midjourney.	Head of Design / Designated AI Lead	11 Aug 2026
AI-powered CRM / sales tool	Vendor name TBC — complete on next review	Lead scoring, sales forecasting, client pipeline management, and automated outreach.	Client contact records, sales interaction history, staff sales data. Likely includes names, email addresses, and behavioural data.	Yes — client contacts and prospective customers are data subjects.	MEDIUM-HIGH — if tool scores or ranks individuals it may constitute automated decision-making under Colorado SB 26-189 (in force 1 Jan 2027); Illinois BIPA applies if biometric identifiers are collected; FTC Section 5 if outputs mislead clients about AI involvement.	Head of Sales / Designated AI Lead	11 Aug 2026

Tool	Vendor	Business purpose	Data it touches	Personal data (Yes/No)	Risk level	Owner	Date reviewed
AI HR / recruitment screening tool	Vendor name TBC — complete on next review	Automated screening and ranking of job applicants for Cedar Line Creative's own hiring.	Applicant CVs, names, contact details, employment history, and any protected-class data inferable from profiles.	Yes — applicants are data subjects; sensitive categories likely present.	HIGH — EEOC AI in Employment guidance applies now; Colorado SB 26-189 consequential-decision deployer duties apply from 1 Jan 2027; Illinois BIPA if biometric data collected from IL applicants;	Head of People / Designated AI Lead	11 Aug 2026

Tool	Vendor	Business purpose	Data it touches	Personal data (Yes/No)	Risk level	Owner	Date reviewed
					Texas HB 149 prohibited-use and disclosure duties in force 1 Jan 2026; AI bias risk requires pre-deployment adverse-impact testing.		
AI chat assistant (client websites)	Vendor name TBC — complete on next review	Customer service and chatbot interactions on client-brand websites operated by Cedar Line Creative.	Website visitor inputs, session data, names, queries, and any personal data volunteered in chat.	Yes — end-users are data subjects; California, Texas, and Illinois residents likely among users.	HIGH — California SB 942 requires disclosure that users are interacting with AI (in force 2 Aug 2026); Texas HB 149 government-AI disclosure duty and prohibited uses (in force 1 Jan 2026); FTC Section 5 deception risk if AI persona is undisclosed; Illinois BIPA if voice or facial biometrics are collected.	Head of Digital / Designated AI Lead	11 Aug 2026
[Add new tool here]							
[Add new tool here]							

RISK RATING SCALE. LOW: minimal personal-data exposure, no consequential decisions, outputs not published externally. MEDIUM: personal data processed but limited volume or sensitivity; outputs published but disclosure controls in place. MEDIUM-HIGH: personal data processed at scale or outputs inform business decisions about individuals; partial controls in place. HIGH: personal data of consumers or applicants processed; outputs published under client brands without current disclosure; regulatory duties already in force and currently unmet. REGISTER OWNERSHIP AND REVIEW CADENCE. This register is owned by the Designated AI Lead (to be formally appointed). It must be reviewed in full every six months, and updated within 14 days whenever Cedar Line Creative adopts, retires, or materially changes any AI tool. The next scheduled full review is due by 11 February 2027. VENDOR DETAILS. Three tools are listed without confirmed vendor names (CRM/sales tool, HR screening tool, AI chat assistant). Vendor identity, sub-processor locations, data-processing agreements, and contractual IP terms must be confirmed and entered before the next client or regulator request. APPLICABLE REGULATORY FRAMEWORK. The following instruments directly govern one or more tools in this register. California SB 942 (AI Transparency Act, as amended by AB 853): in force 2 August 2026 — latent provenance disclosure and free detection-tool duties for large generative AI providers; Cedar Line Creative must confirm whether its generative-AI vendors are covered providers and ensure distributed outputs carry required disclosures. Texas HB 149 (Responsible Artificial Intelligence Governance Act): in force 1 January 2026 — prohibited AI uses and disclosure obligations; applies to AI chat deployed on client sites and to HR screening where Texas applicants are assessed. Colorado SB 26-189: in force 1 January 2027 — developer and deployer duties for automated decision-making in consequential decisions; will apply to HR screening tool and potentially the CRM scoring tool once commenced. Illinois BIPA: in force — applies immediately if any tool collects biometric identifiers (fingerprints, voiceprints, facial geometry) from Illinois residents. EEOC AI in Employment guidance: applies now to the HR screening tool; adverse-impact analysis required before deployment and on an ongoing basis. FTC Section 5: applies now — undisclosed AI-generated advertising content and AI-persona chatbots without disclosure are actionable as deceptive practices; Cedar Line Creative currently has no disclosure wording in place, which is the single highest-priority remediation action. IP NOTE. Ownership of AI-generated copy and imagery under client contracts is legally unsettled. Cedar Line Creative should audit client agreements to confirm whether AI-output assignment clauses are present and whether client brand guidelines require disclosure of AI involvement in creative work.

Plain-English Summary of Your Obligations

The whole regulatory picture in language a director can act on without a lawyer sitting next to them.

Plain-English Summary of Obligations — Cedar Line Creative

This document sets out every material AI compliance obligation that applies to Cedar Line Creative as of 11 August 2026. It is written for directors and senior managers who need to act on it, not study it. Legal references are gathered at the end. No obligation has been omitted and no placeholder language has been used.

What Applies to Cedar Line Creative and Why

Cedar Line Creative is a marketing and advertising agency headquartered in Denver, Colorado, with staff and clients in California, Texas and Illinois. The business publishes AI-generated copy and imagery under client brands, operates AI chat assistants on client websites, and uses an AI screening tool in its own hiring process. It uses ChatGPT, Claude, Midjourney, DALL-E, AI-powered CRM tools and an AI HR screening system. Because the business processes personal data through AI, uses AI to make decisions affecting people, and operates commercially across multiple states, several overlapping legal regimes reach it right now.

California's AI Transparency Act has been in force since 2 August 2026. It imposes disclosure obligations on businesses that deploy generative AI to produce content for California consumers, which Cedar Line Creative does whenever it operates AI chat assistants on client websites with California visitors or publishes AI-generated copy aimed at California audiences.

Texas's Responsible Artificial Intelligence Governance Act has been in force since 1 January 2026. It applies to businesses that deploy AI systems affecting Texas residents, including in commercial communications and employment. Cedar Line Creative has Texas staff and Texas clients, so it is in scope.

Colorado's automated decision-making law — SB 26-189 — is not yet in force. It takes effect on 1 January 2027 and will impose developer and deployer duties for AI used in consequential decisions. Cedar Line Creative is headquartered in Colorado and uses an AI hiring tool, so it will be in scope from that date. Preparation is needed now.

Illinois's Biometric Information Privacy Act applies because Cedar Line Creative has Illinois-based staff. If the AI HR screening tool processes any biometric identifiers — including facial geometry from photographs — BIPA applies immediately and requires specific consent and data governance steps.

The FTC's authority under Section 5 of the FTC Act applies nationally. The FTC has an active AI enforcement programme and has already brought actions for undisclosed AI-generated advertising content. Cedar Line Creative publishes AI-generated copy and imagery under client brands with, by its own admission, no disclosure wording anywhere. That is a live enforcement exposure.

The EEOC's guidance on AI in employment applies to Cedar Line Creative's use of an AI screening tool in its own hiring. The guidance holds employers responsible for discriminatory outcomes produced by AI tools even when those tools are supplied by a third party.

Virginia's Consumer Data Protection Act is listed as an applicable regime. Virginia's CDPA imposes obligations on controllers and processors of personal data. Cedar Line Creative processes personal data through its AI chat assistants and CRM tools. The CDPA applies to businesses that process the personal data of at least 100,000 Virginia consumers, or 25,000 consumers where data is sold. Cedar Line Creative should confirm whether it meets these thresholds for Virginia residents specifically. If it does not currently meet them, the obligation should be kept under review as client volumes grow.

NIST AI RMF 1.0 is a voluntary framework, not a law. It does not carry penalties. It is referenced in this pack because adopting its Govern, Map, Measure and Manage structure is the most practical way to demonstrate responsible AI use to clients, insurers and regulators who do ask about governance frameworks.

Obligations With Deadlines

The following obligations are live or have a fixed future date. Each is stated plainly.

CALIFORNIA AI TRANSPARENCY ACT — IN FORCE 2 AUGUST 2026. Cedar Line Creative must ensure that any AI-generated content delivered to California consumers — including copy, imagery and chat responses — carries a disclosure that the content was produced by AI. The Act requires that large generative AI providers embed latent provenance data in outputs, but it also requires deployers such as Cedar Line Creative to ensure disclosures are visible to end users. Because Cedar Line Creative publishes AI-generated material under client brands with no disclosure wording anywhere, it is in breach of this requirement today for any California-facing content. This must be remedied immediately.

TEXAS RESPONSIBLE ARTIFICIAL INTELLIGENCE GOVERNANCE ACT — IN FORCE 1 JANUARY 2026. Texas law prohibits certain AI uses and requires transparency about AI in consumer-facing interactions. Cedar Line Creative's AI chat assistants deployed on client websites that serve Texas residents must disclose that users are interacting with an AI system, not a human. This obligation is already overdue. The Act also requires businesses to maintain documentation of their AI systems and their intended uses. Cedar Line Creative currently has no such documentation in place.

ILLINOIS BIPA — IN FORCE, NO GRACE PERIOD. If the AI HR screening tool Cedar Line Creative uses processes any biometric data from Illinois residents — including face scans from photographs submitted by applicants — the business must: (1) have a written policy on biometric data retention and deletion; (2) obtain informed written consent from each individual before collection; and (3) not profit from or disclose biometric data without consent. There is no deadline to prepare — these requirements apply from the moment biometric data is collected. Cedar Line Creative must assess immediately whether its screening tool captures biometric data and, if it does, implement BIPA compliance before the next Illinois candidate is processed.

FTC SECTION 5 — ONGOING. The FTC requires that AI-generated content in advertising be disclosed, that AI capability claims be substantiated, and that AI outputs not be suppressed or altered in ways that deceive consumers. Publishing AI-generated imagery and copy under client brand names with no disclosure is the precise fact pattern the FTC actioned in its first undisclosed AI-generated advertising enforcement case. Cedar Line Creative must add disclosure language to all AI-generated content it publishes on behalf of clients. This is not a future obligation — it is an existing one.

EEOC AI IN EMPLOYMENT — ONGOING. Cedar Line Creative is responsible for the outcomes of its AI hiring screening tool, regardless of who supplies the software. If the tool screens out candidates of a protected class at a disproportionate rate, Cedar Line Creative bears liability as the employer. The business must obtain the tool vendor's bias testing documentation, conduct its own adverse impact analysis on hiring outcomes, and keep records sufficient to demonstrate non-discriminatory selection.

COLORADO SB 26-189 — IN FORCE 1 JANUARY 2027. From that date, Cedar Line Creative as a deployer of AI in consequential decisions — hiring is the clearest example — must conduct risk assessments, disclose to affected individuals that automated decision-making is being used, and provide a meaningful opportunity to appeal or seek human review. Cedar Line Creative has approximately five months to build these processes. Work must begin now because the operational changes required — appeals workflows, disclosure notices, documented risk assessments — cannot be implemented overnight.

What Happens If Each Obligation Is Missed

These consequences are concrete and current, not theoretical.

California AI Transparency Act breaches are enforced by the California Attorney General. Penalties can be imposed per violation, and a violation can be read as each piece of non-disclosed AI content delivered to a

California consumer. Given the volume of content Cedar Line Creative produces for clients, cumulative exposure is significant. Reputational damage with California-based clients who discover non-compliance is an additional and immediate commercial risk.

Texas Responsible Artificial Intelligence Governance Act enforcement sits with the Texas Attorney General. The Act carries civil penalties. More significantly for Cedar Line Creative, clients in Texas could treat non-compliance as a breach of contract if Cedar Line Creative has represented that it operates lawfully. Loss of a major Texas client relationship is a foreseeable consequence of inaction.

Illinois BIPA has the most severe private litigation exposure of any regime in this pack. Each negligent violation carries a statutory damages award of $1,000; each intentional or reckless violation carries $5,000. These figures are per person, per violation. Illinois courts have allowed class actions under BIPA, and verdicts and settlements in the tens of millions of dollars are on record for companies far smaller than Cedar Line Creative's sector peers. There is no cap on aggregate exposure.

FTC Section 5 enforcement can result in civil penalties, mandatory injunctive relief requiring Cedar Line Creative to change its practices, and public enforcement actions that are disclosed on the FTC's website. Being named in an FTC AI enforcement action would damage Cedar Line Creative's reputation with the agency clients, consumer brands and regulated-sector businesses that make up a marketing agency's typical client base. The FTC has already pursued over thirteen AI-washing cases and its first undisclosed AI advertising action. This is active enforcement, not a theoretical risk.

EEOC and employment discrimination liability can result in individual and class claims under Title VII and analogous state laws. Remedies include back pay, compensatory and punitive damages, and injunctive relief. The EEOC can also investigate and publicise findings. For a business that uses AI in hiring and has no bias monitoring in place, the exposure is open-ended until testing is done.

Colorado SB 26-189 non-compliance from 1 January 2027 carries enforcement by the Colorado Attorney General, with civil penalties available. Because Cedar Line Creative is headquartered in Colorado, it is a local business under local law, which makes enforcement more straightforward for the regulator than pursuing an out-of-state business. Missing the 1 January 2027 deadline is not low-risk.

What Is Already Handled by This Compliance Pack

This compliance pack for Cedar Line Creative addresses the following areas, so that the team working through it does not need to commission separate analysis on each point.

The pack includes an AI Use Register, which documents each AI tool in use — ChatGPT, Claude, Midjourney, DALL-E, the CRM AI tools and the HR screening system — along with the data each processes, the decisions each influences, and the legal basis for each use. This directly supports the documentation requirements under the Texas Act and provides the foundation for Colorado's forthcoming risk assessment obligations.

The pack includes a Disclosure and Transparency Framework, which provides ready-to-deploy disclosure language for AI-generated content, AI chat interactions and AI-assisted hiring. This addresses the California and Texas disclosure obligations and the FTC's undisclosed AI content enforcement risk.

The pack includes a Bias and Fairness Assessment Protocol for the AI HR screening tool, covering the steps Cedar Line Creative must take to obtain vendor documentation, run adverse impact analysis and maintain records — meeting EEOC expectations.

The pack includes a BIPA Compliance Checklist and consent template for Illinois-based staff and applicants, covering the written policy, individual consent and data retention requirements.

The pack includes an Intellectual Property and Ownership Policy for AI outputs, addressing the copyright and ownership risks that arise when AI-generated content is published under client brands.

The pack includes a Colorado SB 26-189 Readiness Plan, laying out the risk assessment structure, human review workflow and individual notice process that must be operational by 1 January 2027.

The NIST AI RMF alignment summary in the pack maps Cedar Line Creative's current practices to the Govern, Map, Measure and Manage functions, giving the business a structured governance record it can provide to clients, insurers or regulators on request.

The Three Things to Do First

These three actions address the highest-probability, highest-consequence exposures. They must be completed before anything else in the pack.

FIRST: Add AI disclosure language to every piece of AI-generated content published on behalf of clients, and to every AI chat assistant Cedar Line Creative operates on client websites. This single step reduces the most immediate FTC, California and Texas enforcement exposure. The Disclosure and Transparency Framework in this pack contains the exact wording. The person responsible for content production and digital delivery should deploy this wording across all active client accounts within five working days. Do not wait for client approval before beginning the process — the legal obligation sits with Cedar Line Creative, not the client, for content Cedar Line Creative produces and deploys.

SECOND: Assess whether the AI HR screening tool processes biometric data from Illinois-based applicants or staff. Contact the tool's vendor today and request confirmation in writing of what data types the tool captures, whether facial geometry or other biometric identifiers are among them, and what the vendor's own BIPA compliance position is. If biometric data is being processed, suspend use of the tool for Illinois individuals until the BIPA consent and policy requirements are met. The BIPA Compliance Checklist in this pack is the working document for this step. BIPA exposure is the only obligation in this pack that carries uncapped per-person statutory damages and a proven class action litigation track record. It is the single highest financial risk Cedar Line Creative faces.

THIRD: Assign a named individual — a director or senior manager — as the responsible owner for AI compliance. That person should review the full compliance pack, confirm completion of the AI Use Register, and set a calendar checkpoint for 1 December 2026 to confirm that the Colorado SB 26-189 readiness steps are on track for the 1 January 2027 deadline. Without a named owner, the pack will not be implemented and the obligations will remain unmet regardless of the quality of the documents.

Legal References

The following references identify the specific laws and guidance documents relied upon in this summary. They are provided here so that Cedar Line Creative's legal advisers or insurers can locate primary sources directly.

California SB 942, the AI Transparency Act, as amended by AB 853, in force 2 August 2026. Source: California Legislature, leginfo.legislature.ca.gov.

Texas HB 149, the Responsible Artificial Intelligence Governance Act, in force 1 January 2026. Source: Texas Legislature, capitol.texas.gov.

Colorado SB 26-189, automated decision-making technology duties for developers and deployers, in force 1 January 2027. Repealed and reenacted SB 24-205, which never commenced. Source: Colorado General Assembly, leg.colorado.gov.

Illinois Biometric Information Privacy Act, 740 ILCS 14. Continuously in force.

FTC Act, Section 5, unfair or deceptive acts or practices. FTC AI Enforcement Unit established January 2026. FTC proposed policy statement on AI output accuracy, 1 July 2026.

EEOC guidance on AI in employment selection and performance management. Continuously applicable under Title VII of the Civil Rights Act and analogous state employment laws.

Virginia Consumer Data Protection Act. Applicable where thresholds of 100,000 Virginia consumers (or 25,000 where data is sold) are met. Cedar Line Creative should confirm threshold status with reference to its own data processing records.

NIST AI Risk Management Framework 1.0. Voluntary framework. No statutory penalties. Referenced for governance structure alignment.

State AI Obligations Assessment

Which state laws reach this business, and what each one requires.

State	Law	Trigger / Nexus Test	Applies to Cedar Line Creative?	Obligations	Deadline
Colorado	SB 26-189 (Automated Decision-Making Technology)	Deployer of AI making 'consequential decisions' affecting Colorado residents. Cedar Line's AI hiring screener affects Colorado-based applicants.	YES — as a deployer of an AI hiring screener that makes consequential decisions (employment) touching Colorado residents.	Conduct impact assessments; provide pre-adverse-action notice and human review right; publish risk management policy; do not discriminate. Developer duties fall on screener vendor.	1 January 2027 — not yet in force. Begin preparation now; assessments must be ready at commencement.
Colorado	SB 26-189 — consumer-facing AI	Same statute; also covers high-risk AI in consumer interactions. AI chatbots on client websites may qualify if they affect consequential decisions.	POSSIBLY — depends on whether client-site chatbots influence decisions (credit, insurance, housing) for Colorado consumers. Confirm chatbot scope.	If in scope: disclose AI use to consumers; provide opt-out or human escalation; maintain incident log; annual impact assessment.	1 January 2027. Scope assessment to be completed before that date.
Illinois	Biometric Information Privacy Act (BIPA)	Collects, captures or possesses biometric identifiers or information of Illinois residents. Triggered by any HR/recruitment AI using facial geometry, voice prints or similar.	YES — AI recruitment screener likely processes biometric data (e.g. video interview facial analysis) of Illinois-based candidates or employees.	Obtain written informed consent before collection; publish retention/destruction policy; no sale of biometric data; comply with five-year or employment-end destruction schedule.	Already in force. Immediate compliance required. Retroactive claims possible; review all past collection.
Illinois	BIPA — client chatbot angle	BIPA applies if Cedar Line's AI chat tools on Illinois client sites capture voice or facial biometrics of Illinois residents.	POSSIBLY — if chatbots use voice recognition or facial features. Text-only chat is not in scope. Confirm technical spec with each AI chat vendor.	If biometrics collected: written consent, published policy, no onward sale. Cedar Line must ensure client-site deployments comply, not only vendor-side.	Already in force. Audit chat tools immediately.
Texas	HB 149 — Responsible Artificial Intelligence Governance Act	Applies to developers and deployers of AI systems that present a 'reasonable risk of harm' and operate in Texas. Cedar Line has Texas clients and staff.	YES — deployer of AI hiring screener (consequential employment decisions) and AI chat tools affecting Texas residents.	Adopt an AI governance policy; conduct risk assessments for high-risk AI; implement safeguards against algorithmic discrimination; maintain documentation available to Attorney General on request.	In force 1 January 2026. Cedar Line is already subject to this law. Immediate remediation required.

State	Law	Trigger / Nexus Test	Applies to Cedar Line Creative?	Obligations	Deadline
Texas	HB 149 — government AI disclosure	Separate obligation: entities using AI in interactions with the public must disclose AI use. Applies to Cedar Line's client-website chatbots used in Texas.	YES — AI chat assistants on client websites interact with Texas consumers. Disclosure is required regardless of client brand.	Display clear, prominent disclosure that interactions are AI-powered before or at point of first contact. Disclosure must not be buried in terms of service.	In force 1 January 2026. No disclosure currently exists — urgent gap.
California	SB 942 — AI Transparency Act (as amended by AB 853)	Applies to 'large generative AI providers'. Cedar Line is a downstream deployer publishing AI-generated content, not the provider (OpenAI, Anthropic, Midjourney). Provider obligations fall on those vendors.	INDIRECT — Cedar Line relies on OpenAI, Anthropic and Midjourney to embed latent provenance disclosures. Cedar Line must verify those vendors comply and must not strip provenance data from outputs.	Do not remove or suppress AI provenance metadata from generated copy or images before publishing under client brands. Confirm vendor compliance with SB 942's detection-tool and latent-disclosure requirements.	In force 2 August 2026. Vendor verification audit required immediately.
California	SB 942 — con sumer-facing disclosure	If Cedar Line itself meets the 'large provider' threshold (position should be confirmed against statutory revenue/volume thresholds), direct obligations apply.	CONFIRM — Cedar Line should verify whether its own generative AI output volume triggers provider-level duties under SB 942. Current indications suggest it is a deployer, not a provider.	If provider threshold met: offer free AI detection tool; embed latent disclosures in all generative outputs; maintain provenance records.	2 August 2026 if threshold met. Confirm threshold status now.
California	California Consumer Privacy Act / CPRA — AI and personal data	Applies to businesses meeting CCPA thresholds that process personal data of California residents. Cedar Line's CRM and HR AI tools process such data.	YES if CCPA thresholds met (annual gross revenue >$25 M, or buying/selling data of 100,000+ consumers, or deriving 50%+ revenue from data sale). Confirm thresholds against Cedar Line financials.	Disclose automated decision-making in privacy notice; honour opt-out of profiling; conduct data protection assessments for high-risk processing; maintain records of processing activities.	Already in force. Confirm threshold applicability; update privacy notice and DPA records immediately if in scope.
Virginia	Consumer Data Protection Act (CDPA)	Applies to businesses that control/process data of 100,000+ Virginia consumers, or 25,000+ consumers where 50%+ revenue derives from data sale. Cedar Line has Virginia clients or staff; confirm volumes.	CONFIRM — Cedar Line's client base and CRM data may encompass Virginia residents. Threshold must be verified against actual data volumes.	If in scope: provide opt-out of profiling for decisions with legal or similarly significant effects; conduct data protection assessments; publish privacy notice disclosing AI-driven profiling.	Already in force. Threshold check required; implement if met.

State	Law	Trigger / Nexus Test	Applies to Cedar Line Creative?	Obligations	Deadline
Federal (FTC)	FTC Act Section 5 — Unfair or Deceptive Acts (AI enforcement)	Applies to any business making commercial AI-related representations that are false or unsubstantiated, or that use AI to deceive consumers. Cedar Line publishes AI-generated copy under client brands.	YES — Cedar Line publishes AI-generated advertising and operates AI chatbots. Both activities are within the FTC's current AI enforcement focus. No disclosure wording currently exists.	Disclose AI-generated content where it constitutes a material representation; do not make unsubstantiated AI performance claims; do not configure AI to suppress accurate outputs. Document substantiation for all AI-related claims.	Already enforceable. FTC AI enforcement unit active since January 2026. Immediate disclosure and substantiation audit required.
Federal (EEOC)	EEOC Guidance — AI in Employment Decisions	Applies to any employer using AI tools that affect hiring, promotion or performance management. Cedar Line uses an AI screening tool for its own recruitment.	YES — Cedar Line is an employer using AI to screen job candidates. EEOC guidance treats disparate impact from AI screening as potential Title VII and ADA violation.	Validate screener for adverse impact across protected classes; retain adverse-impact data; ensure human review of AI-recommended rejections; do not use screener as sole determinative factor.	Guidance already in effect. Adverse-impact analysis of current screener required without delay.

SCOPE NOTE — CEDAR LINE CREATIVE: This assessment covers every jurisdiction where Cedar Line Creative operates or sells: Colorado (HQ), California, Texas and Illinois (declared client and staff locations), Virginia (assessed for CDPA nexus), and applicable federal regimes. Where Virginia CDPA applicability depends on data volumes not provided, Cedar Line must run a threshold check against its CRM and analytics data. Where Colorado SB 26-189 is not yet in force, preparatory obligations (impact assessments, vendor due diligence, policy drafting) should begin now to achieve readiness by 1 January 2027. | BIOMETRIC RISK — ILLINOIS BIPA: BIPA carries a private right of action worth $1,000–$5,000 per violation, per person, per occurrence. Class actions are common. Cedar Line must audit every AI tool in its hiring workflow — including any video-interview analysis, voice-print or facial-geometry processing — and obtain written, informed consent from each affected individual before any further collection. Existing unconsented collection may be actionable retrospectively. | AI-GENERATED CONTENT AND IP: This assessment covers regulatory obligations only. Intellectual property ownership of AI-generated copy and imagery (Midjourney, DALL-E, ChatGPT, Claude) published under client brands is addressed in the IP and Ownership Assessment document within this pack. Cedar Line should not represent AI outputs as wholly human-authored in any client contract or public-facing statement. | DISCLOSURE GAP — URGENT: Cedar Line currently has no AI disclosure wording anywhere. This creates simultaneous exposure under Texas HB 149 (in force), FTC Section 5 (active enforcement), and — for chatbots — potential state consumer protection liability. A single, layered disclosure framework (chatbot initiation notice, content provenance footnote, HR applicant notice) should be implemented as the first remediation action. | COLORADO SB 26-189 STATUS: This law repealed and reenacted SB 24-205, which never commenced. The verified commencement date is 1 January 2027. The briefing reference to SB 24-205 as a separate operative instrument is superseded by the verified date; Cedar Line should plan against SB 26-189 alone. | CALIFORNIA SB 942 STATUS: Delayed from 1 January 2026 to 2 August 2026 by AB 853. The law is now in force. Cedar Line's obligations as a deployer (rather than a large provider) centre on not stripping provenance metadata. Whether Cedar Line meets the provider threshold requires confirmation against statutory criteria not yet published in final form — position should be confirmed with California counsel. | REVIEW CADENCE: This assessment should be reviewed quarterly given the pace of US state AI legislation. The next scheduled review date is 11 November 2026, or earlier if any new state law is enacted affecting Colorado, California, Texas, Illinois or Virginia. | DOCUMENT OWNER: Chief Compliance Officer (or designated AI Governance Lead), Cedar Line Creative. | RATING SCALE FOR APPLICABILITY COLUMN: YES = confirmed applicable on stated facts; POSSIBLY / CONFIRM = nexus depends on volume or technical data Cedar Line must verify; INDIRECT = obligations rest primarily on AI vendors but Cedar Line has secondary duties.

Employment AI Disclosure Notices

The notices given to candidates and staff when AI is used in a decision about them, ready to issue.

These Employment AI Disclosure Notices are ready-to-issue documents for Cedar Line Creative. They address two distinct contexts: (1) candidates applying for roles at Cedar Line Creative, where an AI screening tool is used in the hiring process; and (2) current employees subject to AI-assisted monitoring or performance evaluation. Both notices comply with obligations under Colorado SB 26-189 (deployer duties, effective 1 January 2027 — notices should be issued now to establish practice ahead of commencement), the EEOC's guidance on AI in employment, Illinois BIPA (for staff and candidates in Illinois), Texas HB 149 (in force 1 January 2026), and FTC Section 5 (which prohibits material omissions about AI use in consumer and employment-facing communications). California SB 942 applies to Cedar Line Creative's generative AI outputs but its disclosure obligations are directed at large generative AI providers rather than deployers; nonetheless, the transparency principle is reflected here. Virginia CDPA applies to personal data processing and informs the data description fields. A jurisdiction-specific advance-notice variant is included for Illinois (BIPA requires a written policy and informed consent before any biometric or automated data collection). Notices must be delivered before the relevant process begins, not after. Completed acknowledgement copies must be retained for a minimum of three years or the duration of the employment relationship plus two years, whichever is longer.

NOTICE A — CANDIDATE-FACING NOTICE: USE OF ARTIFICIAL INTELLIGENCE IN RECRUITMENT
Cedar Line Creative | Human Resources

DATE OF ISSUE: [DATE]
ROLE APPLIED FOR: [JOB TITLE]
REQUISITION NUMBER: [REQ NUMBER]
CANDIDATE FULL NAME: [CANDIDATE NAME]
CANDIDATE EMAIL: [CANDIDATE EMAIL]
HIRING LOCATION (state): [STATE]

A1. WHO IS ISSUING THIS NOTICE

Cedar Line Creative ("we", "us") is headquartered at [REGISTERED ADDRESS], Denver, Colorado. We are the deployer of the AI tool described in this notice. Questions may be directed to [HR CONTACT NAME], [HR CONTACT EMAIL], [HR CONTACT PHONE].

A2. THAT AI IS USED AND FOR WHAT PURPOSE

We use an automated screening tool ([TOOL NAME / VENDOR]) at the initial stage of our recruitment process. The tool analyses your application materials to produce a ranked score or shortlisting recommendation. This output is one input into a decision about whether you advance to a human-reviewed stage. It is not the sole basis for any final hiring decision.

A3. WHAT DATA THE TOOL PROCESSES

The tool processes the following categories of personal data that you submit or that we hold:

[ ] CV / résumé text
[ ] Cover letter text
[ ] Application form responses
[ ] Self-reported skills or qualifications
[ ] Publicly available professional profile data (if applicable — specify source: [SOURCE])
[ ] Other: [SPECIFY]

The tool does NOT process biometric identifiers, photographs, or protected-class data as screening inputs. If any additional data category is added, this notice will be reissued before processing begins.

A4. LEGAL BASIS AND APPLICABLE LAW

This notice is issued in compliance with:

— EEOC guidance on AI in employment (2023 and updated 2025): requires transparency about AI use in selection and a mechanism for human review.
— Colorado SB 26-189 (effective 1 January 2027, deployer obligations): requires notice to individuals subject to consequential automated decisions and a meaningful human review option. Cedar Line Creative issues this notice in advance of commencement to establish compliant practice.
— Texas HB 149, Responsible Artificial Intelligence Governance Act (in force 1 January 2026): requires disclosure of AI use in decisions affecting individuals and prohibits certain high-risk AI uses without safeguards.
— Illinois BIPA (740 ILCS 14): if you are located in or applying for a role in Illinois, see Section A8 (Illinois Advance Notice Variant) below.
— FTC Section 5 of the FTC Act: material omissions about AI use in employment-facing communications may constitute an unfair or deceptive act or practice.
— Virginia CDPA: if you are a Virginia resident, you have rights regarding personal data used in automated processing; see Section A6.

A5. HUMAN REVIEW AVAILABLE TO YOU

No offer will be refused, and no candidate will be eliminated from consideration, solely on the basis of an AI-generated score without human review of that output.

If you wish to request human review of the AI screening recommendation before a final decision is made, you must submit your request in writing to [HR CONTACT EMAIL] within [NUMBER] calendar days of receiving this notice or of being notified of an adverse preliminary outcome, whichever is earlier.

Upon receipt of a valid request, a qualified human reviewer ([JOB TITLE OF REVIEWER]) will independently assess your application materials and the AI output. You will receive a written response within [NUMBER] business days.

A6. YOUR DATA RIGHTS

You may request access to, correction of, or deletion of personal data we hold about you by contacting [PRIVACY CONTACT EMAIL]. Virginia CDPA residents may opt out of solely automated decisions that produce legal or similarly significant effects; submit opt-out requests to [PRIVACY CONTACT EMAIL]. We will respond within 45 days.

A7. RECORD-KEEPING

A signed copy of this notice, together with any human-review request and response, will be retained in your candidate file for a minimum of three years from the date of this notice. Completed records may be inspected by regulators including the EEOC, Colorado Attorney General, or Texas Attorney General on request.

A8. ILLINOIS ADVANCE-NOTICE VARIANT (COMPLETE THIS SECTION FOR ILLINOIS CANDIDATES ONLY)

Illinois BIPA (740 ILCS 14/15) requires a written policy and informed written consent before any collection or use of biometric or automated personal data. This section constitutes Cedar Line Creative's required advance disclosure for Illinois candidates.

Specific data collected in Illinois: [SPECIFY ANY BIOMETRIC OR ADDITIONAL DATA]
Retention schedule: Data will be retained for [PERIOD] and destroyed within [PERIOD] of the earlier of (a) your withdrawal, (b) a final hiring decision, or (c) three years from collection.
Consent: Your signed acknowledgement in Section A9 constitutes your informed written consent under BIPA. You may withdraw consent at any time by written notice to [HR CONTACT EMAIL]; withdrawal will not itself disqualify you but will mean the AI screening tool cannot be used and your application will proceed by manual review only.

A9. CANDIDATE ACKNOWLEDGEMENT

I confirm that I have read and understood this notice before the AI screening process began. I understand that AI is used to assist (not replace) human judgment in the hiring process, that I may request human review, and that my personal data will be handled as described above.

Candidate signature: ______________________________
Printed name: ______________________________
Date: ______________________________

For Illinois candidates — BIPA consent: By signing above I also provide informed written consent to the collection and use of personal data as described in Section A8.

Received by (HR representative): ______________________________
Printed name: ______________________________
Date: ______________________________
Staff ID: ______________________________

NOTICE B — EMPLOYEE-FACING NOTICE: USE OF ARTIFICIAL INTELLIGENCE IN MONITORING OR PERFORMANCE EVALUATION
Cedar Line Creative | Human Resources

DATE OF ISSUE: [DATE]
EMPLOYEE FULL NAME: [EMPLOYEE NAME]
EMPLOYEE ID: [EMPLOYEE ID]
DEPARTMENT: [DEPARTMENT]
LINE MANAGER: [MANAGER NAME]
WORK LOCATION (state): [STATE]

B1. WHO IS ISSUING THIS NOTICE

Cedar Line Creative is issuing this notice in its capacity as your employer and as deployer of the AI tools described below. Contact for queries: [HR CONTACT NAME], [HR CONTACT EMAIL], [HR CONTACT PHONE].

B2. THAT AI IS USED AND FOR WHAT PURPOSE

Cedar Line Creative uses AI-powered tools in one or more of the following ways that may affect your employment or produce information about your performance. Tick all that apply to your role:

[ ] AI-powered CRM or sales tools: activity data (calls, emails, pipeline entries) is aggregated and scored to produce performance indicators used in sales reviews.
[ ] AI analysis of communications metadata for workload or productivity assessment.
[ ] AI in HR systems: performance data is processed to generate review inputs, development recommendations, or risk flags.
[ ] Other: [SPECIFY TOOL AND PURPOSE]

The purpose of each active use is to assist managers in making informed, fair assessments — not to replace managerial judgment.

B3. WHAT DATA THE TOOLS PROCESS

Depending on the tools ticked above, the following personal data categories may be processed:

[ ] CRM activity logs (call volumes, email counts, deal values)
[ ] Project management system outputs (task completion rates, deadlines met)
[ ] HR system records (attendance, leave, prior appraisal scores)
[ ] Client satisfaction scores attributable to your account
[ ] Other: [SPECIFY]

Sensitive personal data (health, protected characteristics) is not used as a direct input to any AI evaluation tool. If this changes, a revised notice will be issued before the change takes effect.

B4. LEGAL BASIS AND APPLICABLE LAW

— EEOC AI in Employment guidance: requires that AI evaluation tools be monitored for adverse impact on protected groups and that employees are not subjected to discriminatory automated outputs without recourse.
— Colorado SB 26-189 (effective 1 January 2027): for consequential employment decisions (promotion, demotion, termination, material change in compensation), deployers must provide notice and a human review mechanism. Cedar Line Creative implements this now.
— Texas HB 149 (in force 1 January 2026): disclosure of AI use in decisions affecting individuals is required; prohibited uses include AI that produces discriminatory outcomes in employment without adequate safeguards.
— FTC Section 5: material non-disclosure of AI-generated performance data used to support adverse employment actions may constitute an unfair practice.
— Illinois BIPA: if you are based in Illinois, see Section B8.
— Virginia CDPA: Virginia-based employees have the right to opt out of solely automated consequential decisions; see Section B6.

B5. HUMAN REVIEW AVAILABLE TO YOU

No consequential employment decision (including performance improvement plans, demotion, compensation reduction, or termination) will be made solely on the basis of an AI-generated output without human managerial review of that output and of the underlying data.

If you believe an AI-generated output about you is inaccurate or unfair, you may:

1. Submit a written challenge to [HR CONTACT EMAIL] within [NUMBER] calendar days of being informed of any decision influenced by AI output.
2. Request sight of the AI-generated data used in the decision.
3. Request that a human reviewer ([JOB TITLE]) conduct an independent assessment.

Cedar Line Creative will respond in writing within [NUMBER] business days of receiving a valid challenge.

B6. YOUR DATA RIGHTS

You may request access to personal data held about you in any AI system by contacting [PRIVACY CONTACT EMAIL]. Virginia CDPA: you may opt out of solely automated consequential decisions; submit requests to [PRIVACY CONTACT EMAIL] and we will respond within 45 days. Requests will not be used as grounds for any adverse action.

B7. RECORD-KEEPING

This notice, together with any challenge, response, and human-review record, will be retained in your personnel file for the duration of your employment plus two years, or three years from the date of this notice, whichever is longer. Records are available to regulators on request.

B8. ILLINOIS ADVANCE-NOTICE VARIANT (COMPLETE FOR ILLINOIS-BASED EMPLOYEES ONLY)

Under Illinois BIPA (740 ILCS 14/15), Cedar Line Creative must provide a written policy and obtain informed written consent before using any system that collects biometric or automated personal data about you.

Specific data collected in Illinois: [SPECIFY]
Retention schedule: Data will be retained for [PERIOD] and destroyed within [PERIOD] of the earlier of (a) end of employment or (b) three years from collection, whichever comes first.
Consent: Your signature in Section B9 constitutes informed written BIPA consent. You may withdraw consent by written notice to [HR CONTACT EMAIL]. Withdrawal will result in manual-only performance assessment processes being applied to your role; it will not itself constitute grounds for any adverse employment action.

B9. EMPLOYEE ACKNOWLEDGEMENT

I confirm that I received this notice before any AI monitoring or evaluation process began or, where already underway, that this notice has been provided to me promptly and the human review mechanism is available to me with immediate effect.

Employee signature: ______________________________
Printed name: ______________________________
Date: ______________________________

For Illinois employees — BIPA consent: By signing above I also provide informed written consent as required by Section B8.

Received by (HR representative): ______________________________
Printed name: ______________________________
Date: ______________________________
Staff ID: ______________________________

ADMINISTRATION NOTES FOR CEDAR LINE CREATIVE HR
(Not for issue to candidates or employees)

TIMING REQUIREMENT: Notice A must be issued before the AI screening tool processes a candidate's application. Notice B must be issued before any AI monitoring or evaluation tool is first applied to an employee, or, for existing employees already subject to such tools, as soon as practicable and in any event before the next evaluation cycle.

JURISDICTION ROUTING: Colorado (all staff) — full notices apply now; SB 26-189 deployer duties commence 1 January 2027 but practice should begin immediately. Texas — HB 149 in force; disclosure and human-review obligations are live. Illinois — BIPA advance-notice and consent sections (A8 / B8) are mandatory; do not omit. California — SB 942 AI Transparency Act (in force 2 August 2026) applies to large generative AI providers, not to Cedar Line Creative as a deployer; no separate SB 942 notice is required from Cedar Line Creative in the employment context, though the transparency principle is reflected throughout. Virginia CDPA opt-out rights are noted in both notices for Virginia-resident candidates and employees.

RECORD RETENTION: Retain signed originals (or secure electronic equivalents with audit trail) for the periods stated in A7 and B7 respectively. Store separately from the general recruitment file so that records are producible on regulatory request without disclosing unrelated candidate data.

BIAS MONITORING: Under EEOC guidance and in anticipation of Colorado SB 26-189, Cedar Line Creative must maintain records sufficient to audit the AI screening tool for adverse impact on protected groups. [HR CONTACT NAME] is responsible for commissioning an annual disparate-impact review of hiring outcomes.

VENDOR OBLIGATIONS: Obtain and retain the AI screening tool vendor's bias testing documentation, data processing agreement, and any NIST AI RMF alignment statements. These records substantiate Cedar Line Creative's deployer due-diligence obligations under Colorado SB 26-189 and the EEOC guidance.

Algorithmic Discrimination Prevention Procedures

How you detect and prevent discriminatory outcomes, and what you do when testing finds one.

1. Purpose and Scope

These Algorithmic Discrimination Prevention Procedures set out how Cedar Line Creative identifies, measures, remediates and documents discriminatory outcomes arising from its use of artificial intelligence tools. Cedar Line Creative operates as a marketing and advertising agency headquartered in Denver, Colorado, with staff and clients in California, Texas and Illinois. It uses ChatGPT/OpenAI, Claude/Anthropic, AI image generation tools (Midjourney and DALL-E), AI-powered CRM and sales tools, and an AI-based screening tool in its own recruitment process.

These procedures apply to every AI system that Cedar Line Creative deploys or instructs a vendor to deploy where outputs may affect individuals. That includes the AI hiring screener used for Cedar Line Creative's own recruitment, the AI chat assistant running on client websites, and any AI tool that segments, scores or ranks individuals in a CRM or sales context. AI image and copy generation tools are included to the extent that outputs could reinforce stereotypes that harm protected groups in client-facing communications.

The legal obligations engaged include: the EEOC's guidance on AI in employment selection; FTC Section 5 of the FTC Act as applied to AI-generated claims and outputs; Colorado SB 26-189 (in force 1 January 2027) imposing developer and deployer duties for automated decision-making technology used in consequential decisions; Texas HB 149, the Responsible Artificial Intelligence Governance Act (in force 1 January 2026); Illinois BIPA where biometric data is implicated; California SB 942 as amended by AB 853 (in force 2 August 2026); and Virginia CDPA where applicable to data subjects in that state. NIST AI RMF 1.0 informs the risk management framework underpinning these procedures.

2. Protected Characteristics and How Disparate Impact Is Measured

Cedar Line Creative recognises the following protected characteristics for the purpose of disparate impact analysis: race, colour, national origin, sex (including pregnancy, gender identity and sexual orientation), religion, disability, age (40 and over), and genetic information, as protected under Title VII of the Civil Rights Act, the Age Discrimination in Employment Act, the Americans with Disabilities Act and Executive Order 11246. The EEOC's April 2023 guidance on the use of software, algorithms and artificial intelligence to assess job candidates confirms that these protections apply fully where an AI tool is used as a selection procedure, and Cedar Line Creative's AI hiring screener falls squarely within that scope.

For employment-related AI tools, Cedar Line Creative measures adverse impact using the four-fifths (80%) rule established in the Uniform Guidelines on Employee Selection Procedures. The selection rate for any protected group is divided by the selection rate for the group with the highest selection rate. A ratio below 0.80 constitutes prima facie evidence of adverse impact and triggers the remediation steps in Section 5 of these procedures. Where sample sizes are too small for the four-fifths rule to be statistically reliable (fewer than 30 candidates in a group), Cedar Line Creative will supplement with Fisher's exact test or a chi-square test and document which method was used and why.

For AI tools used in client website chat interactions and CRM segmentation, disparate impact is assessed by examining whether response quality, escalation rates or commercial outcomes differ systematically across demographic groups. Because Cedar Line Creative does not always hold demographic data for website visitors, proxy analysis using zip-code-level census data or platform-reported demographic aggregates will be used where direct data is unavailable. The Legal and Operations lead will document the proxy methodology and its limitations in the testing record.

For AI-generated imagery and copy published under client brands, Cedar Line Creative evaluates whether outputs systematically under-represent, stereotype or exclude identifiable groups defined by race, sex, age or disability.

This is assessed through structured human review using a coded rubric rather than a statistical ratio, and the threshold for action is set out in Section 4.

3. Testing Method, Cadence and Who Performs It

Cedar Line Creative conducts algorithmic discrimination testing at three levels: pre-deployment, periodic scheduled review, and event-triggered review.

Pre-deployment testing is mandatory before any new AI system goes live in a use case that affects individuals. For the AI hiring screener, pre-deployment testing consists of running a structured audit using a representative synthetic candidate dataset that reflects the demographic distribution of Cedar Line Creative's applicant pool. The dataset must include at least 200 synthetic profiles per protected-class group being tested. The Head of People Operations commissions this audit and appoints an internal reviewer who was not involved in selecting the tool. Where Cedar Line Creative lacks the internal statistical capability to perform the audit independently, it will engage an external auditor before the tool is activated.

Scheduled periodic testing is conducted on the following cadence. The AI hiring screener is tested every six months, timed to align with Cedar Line Creative's two main recruitment cycles. The AI chat assistant and CRM tools are tested annually, or following any material update to the underlying model or the prompt configuration used. AI image and copy generation outputs are reviewed quarterly by a two-person panel drawn from the Creative and Legal functions using the rubric described in Section 2.

Event-triggered testing is initiated immediately when: a complaint or grievance alleges discriminatory treatment linked to an AI tool; a vendor notifies Cedar Line Creative of a model update that could affect bias characteristics; a regulator (including the EEOC, FTC, Colorado Attorney General or Texas Attorney General) publishes enforcement action or guidance relevant to a tool in use; or internal monitoring data shows an anomaly in outcome rates exceeding the thresholds in Section 4.

Responsibility for testing sits with the Compliance Lead, who reports to the CEO. For the hiring screener specifically, the Head of People Operations is the operational owner but must route all bias-testing results through the Compliance Lead. External auditors must be instructed in writing and must provide a written report that Cedar Line Creative retains as a compliance record. No AI system may be re-activated after a suspension without written sign-off from the Compliance Lead and, for employment tools, the Head of People Operations.

4. Threshold at Which a Result Requires Action

The following numerical thresholds apply. Any result at or beyond a threshold automatically triggers the remediation process in Section 5; it is not discretionary.

For the AI hiring screener: an adverse impact ratio below 0.80 (the four-fifths rule) for any protected group with a statistically sufficient sample (30 or more candidates) requires action. Where statistical significance testing is used for smaller samples, a p-value of 0.05 or below (two-tailed) requires action regardless of whether the four-fifths ratio is met.

For the AI chat assistant: a difference in satisfactory-resolution rates between any identifiable demographic group and the overall mean of more than 10 percentage points, measured over a rolling 90-day window, requires action. If Cedar Line Creative is using proxy demographic data, the threshold remains 10 percentage points but the Compliance Lead must note in the record that the finding is indicative rather than conclusive.

For CRM and sales AI tools: a difference in positive-outcome rates (for example, lead scoring, offer presentation or discount eligibility) of more than 10 percentage points between demographic groups, or a statistically significant difference at $p \leq 0.05$, requires action.

For AI-generated imagery and copy: if a quarterly panel review finds that 20% or more of sampled outputs across a campaign type contain stereotyped or exclusionary representations of a protected group, that campaign type is suspended pending remediation. A single output that is judged to be severely harmful (for example, content that could constitute discriminatory advertising under the Fair Housing Act or civil rights laws) requires immediate removal regardless of the 20% threshold.

Thresholds are reviewed annually by the Compliance Lead and may be tightened if regulatory guidance or case law indicates a stricter standard applies. Colorado SB 26-189, which comes into force on 1 January 2027, will impose specific deployer obligations regarding consequential decisions; Cedar Line Creative will review these thresholds before that date to ensure alignment.

5. Remediation When a Disparity Is Found, Including Suspending Use

When a testing result meets or exceeds a threshold in Section 4, the following steps are taken in sequence and documented in the Compliance Register.

Step 1 — Immediate notification. The Compliance Lead is notified within 24 hours of the result being available. If the Compliance Lead is unavailable, the CEO assumes this responsibility. For findings involving the hiring screener, the Head of People Operations is also notified immediately.

Step 2 — Suspension decision. Within 48 hours of notification, the Compliance Lead makes a written suspension decision. The default position is suspension of the AI tool for the affected use case. Continued use without suspension may only occur if the Compliance Lead documents, in writing, a specific operational justification and the mitigating controls in place. This exception is narrow; it does not apply to the hiring screener, which must be suspended immediately upon a confirmed threshold breach given the EEOC's direct applicability and the personal, irreversible nature of employment decisions.

Step 3 — Root cause investigation. Within 14 days of suspension, Cedar Line Creative completes a root cause investigation. The investigation considers: whether the disparity originates in the model's training data, the prompt configuration Cedar Line Creative applies, the input data fed to the model, or the post-processing of outputs. The investigation is led by the Compliance Lead with input from the function that owns the tool.

Step 4 — Vendor engagement. Cedar Line Creative notifies the relevant vendor in writing, requests the vendor's own bias-testing results and model documentation, and asks whether a model update or configuration change is available. The vendor questions in Section 7 of these procedures are used as the basis for this dialogue. Vendor responses are retained in the Compliance Register.

Step 5 — Remediation options. Depending on root cause, remediation may include: modifying or replacing the prompt configuration; applying post-processing filters to outputs; adjusting input data to remove proxies for protected characteristics; requesting a different model version from the vendor; supplementing or replacing the AI tool with human review; or permanently discontinuing use of the tool for the affected use case. Remediation chosen must be documented with a rationale.

Step 6 — Re-testing before reactivation. No tool suspended under these procedures is reactivated until re-testing confirms the disparity has been reduced below the applicable threshold. Re-testing uses the same method as the original test. Results are recorded and retained.

Step 7 — Affected individuals. Where a biased hiring screener result may have affected a real candidate's outcome, Cedar Line Creative will re-evaluate that candidate's application using a human reviewer and, where practicable, notify the candidate. Legal advice on notification obligations under applicable state law will be sought before any outreach.

6. Records Kept as Evidence of Reasonable Care

Cedar Line Creative maintains a Compliance Register dedicated to AI bias and discrimination matters. The register is held electronically in a location accessible to the Compliance Lead and CEO but restricted from general staff access. Records are retained for a minimum of five years from the date of creation, consistent with the retention periods implied by EEOC recordkeeping rules for employment selection procedures and the potential limitation periods for civil rights claims.

The following records are created and retained for every AI tool in scope.

A tool inventory entry, created at onboarding, recording: the tool name and vendor, the use case, the categories of individuals affected, the protected characteristics relevant to that use case, the pre-deployment testing outcome, and the date the tool was approved for use.

A testing record for each scheduled or event-triggered test, recording: the date and type of test, the method used (including statistical method and sample size), the results for each protected group tested, whether any threshold was met or exceeded, and the name of the person or firm that conducted the test.

A remediation record for each threshold breach, recording every step taken under Section 5 including dates, decision-makers, written suspension decisions, root cause findings, vendor correspondence, the remediation chosen, re-testing results, and the date of reactivation or permanent discontinuation.

A vendor due diligence record for each vendor, recording the questions asked under Section 7, the vendor's responses, the date of each exchange, and any commitments made by the vendor regarding bias testing or model updates.

A panel review record for each quarterly creative output review, recording the sample of outputs reviewed, the rubric scores, the names of the two reviewers, the overall finding, and any action taken.

These records collectively constitute Cedar Line Creative's evidence of reasonable care for the purpose of the EEOC's guidance, Colorado SB 26-189's forthcoming deployer obligations (in force 1 January 2027), Texas HB 149's governance requirements (in force 1 January 2026), and NIST AI RMF 1.0 MAP, MEASURE and MANAGE function requirements. The Compliance Lead reviews the completeness of the register quarterly and certifies it annually.

7. Vendor Questions That Establish What Testing They Have Done

Before contracting with or renewing a contract with any AI vendor whose tool is used in a context that affects individuals, Cedar Line Creative asks the following questions in writing and requires written responses. Where a vendor refuses to answer or provides materially incomplete answers, this is recorded in the vendor due diligence record and escalated to the CEO before the contract is signed or renewed.

The mandatory vendor questions are as follows.

Bias and fairness testing: What bias and fairness evaluations has the vendor conducted on this model or product prior to its release? Which protected characteristics were tested? What methodology was used (for example, benchmark datasets, red-teaming, disparate impact analysis)? What were the results, and are they available in a published model card, system card or audit report?

Training data: What data was used to train the model? Does the vendor have documentation of the demographic composition or geographic coverage of training data relevant to bias? Has the vendor assessed whether training data over- or under-represents any protected group?

Ongoing monitoring: Does the vendor conduct post-deployment monitoring for disparate outcomes? If so, at what cadence, using what method, and will Cedar Line Creative be notified of adverse findings?

Model updates: When the vendor updates the underlying model, will Cedar Line Creative receive advance notice? What is the vendor's process for re-evaluating bias characteristics following a model update? Does the vendor warrant that a material update will not introduce new disparate impact without re-testing?

Employment-specific questions (for the AI hiring screener vendor only): Has the tool been validated as a selection procedure under the Uniform Guidelines on Employee Selection Procedures? Has an independent adverse impact analysis been conducted? Is the vendor willing to provide technical documentation sufficient for Cedar Line Creative to conduct its own validation study? Does the vendor indemnify or share liability for EEOC or state civil rights enforcement arising from the tool's outputs?

State-law specific questions: For tools used in Illinois, does the vendor comply with BIPA where any biometric data is processed? For tools affecting individuals in California, does the vendor comply with California SB 942 as amended by AB 853 (in force 2 August 2026) regarding provenance disclosure for AI-generated content? For tools used in Texas, does the vendor have documentation of compliance with Texas HB 149 (in force 1 January 2026)?

Vendor responses are retained in the Compliance Register. Where a vendor provides a model card, audit report or third-party assessment, a copy is saved alongside the vendor due diligence record. Cedar Line Creative reviews vendor responses annually or upon contract renewal, whichever is sooner, and re-poses unanswered questions at each review.

8. Governance, Ownership and Review

The Compliance Lead holds overall accountability for these procedures. The Head of People Operations holds operational accountability for the AI hiring screener. The Head of Creative holds operational accountability for AI image and copy generation tools. The Head of Client Services holds operational accountability for the AI chat assistant and CRM tools.

These procedures are reviewed annually, and additionally when: a new AI tool is onboarded; a relevant law or regulatory guidance is issued or amended; a threshold breach occurs; or an external audit recommends a material change. The next scheduled review is no later than 1 November 2026 to incorporate Colorado SB 26-189 obligations ahead of their 1 January 2027 commencement date.

Breaches of these procedures by Cedar Line Creative staff are treated as conduct matters under the staff handbook. Vendor non-compliance with contractual bias-testing obligations is treated as a material breach of contract.

These procedures do not constitute legal advice and have been prepared as an operational compliance document. Where a specific legal question arises — in particular regarding EEOC enforcement, state civil rights liability or the scope of Colorado SB 26-189 once in force — Cedar Line Creative will obtain jurisdiction-specific legal advice at that time.

Consumer-Facing AI Transparency Disclosures

Telling customers when they are dealing with AI, in the wording each applicable law expects.

1. Purpose and Scope

This document sets out the consumer-facing AI transparency disclosures that Cedar Line Creative must implement across all customer and candidate touchpoints where artificial intelligence is used. It is operative as of 11 August 2026 and supersedes any earlier draft. It has been prepared against the specific legal regimes that reach Cedar Line Creative by reason of its headquarters in Colorado, its staff and client operations in California, Texas and Illinois, and its publication of AI-generated content under client brands.

The disclosures in this document are mandatory, not aspirational. Cedar Line Creative currently has no disclosure wording in place anywhere. That gap creates live exposure under California SB 942 (AI Transparency Act, as amended by AB 853, in force 2 August 2026), Texas HB 149 (Responsible Artificial Intelligence Governance Act, in force 1 January 2026), the Illinois AI Video Interview Act, EEOC guidance on AI in employment, and FTC Section 5 of the Federal Trade Commission Act. Colorado SB 26-189 is not yet in force (commencement 1 January 2027) but its deployer-duty framework is noted where it will affect future disclosure obligations. Illinois BIPA applies to biometric data; Cedar Line Creative should confirm with counsel whether its HR screening tool captures biometric identifiers, as BIPA obligations are not triggered unless it does.

All disclosure wording below is ready for implementation. Where a wording block is marked [CLIENT BRAND], Cedar Line Creative must substitute the relevant client's trading name before publication on that client's website or channel.

2. Consumer Touchpoints Using AI — Complete Inventory

Cedar Line Creative operates AI at five identified touchpoints, each of which requires its own disclosure treatment.

Touchpoint 1 — AI-generated marketing copy published under client brands. Cedar Line Creative uses ChatGPT (OpenAI) and Claude (Anthropic) to draft advertising copy, campaign messaging, social media posts, email content and other communications that are then published under client brand names. End-consumers reading that content are the affected persons.

Touchpoint 2 — AI-generated images and synthetic media published under client brands. Cedar Line Creative uses Midjourney and DALL-E to create images, illustrations and other visual assets that appear in client campaigns. These assets are presented as creative work produced for or by the client brand.

Touchpoint 3 — AI chat assistant on client websites. Cedar Line Creative deploys an AI-powered chat assistant on client websites to handle customer enquiries. Visitors to those sites interact with the assistant believing they may be speaking with a human representative of the client brand.

Touchpoint 4 — AI-powered CRM and sales tools. Cedar Line Creative uses AI within its CRM platform to analyse customer data, generate insights and inform sales outreach. This processes personal data of prospects and existing clients and may produce communications directed at named individuals.

Touchpoint 5 — AI screening tool in Cedar Line Creative's own recruitment. Cedar Line Creative uses an AI system to screen job applicants for its own open roles. Applicants are affected individuals subject to AI-assisted decisions about employment.

Each touchpoint is addressed in the sections that follow.

3. Chatbot and Virtual Agent Disclosure — Wording and Requirements

Legal basis: Texas HB 149 prohibits any person from deploying an AI system that impersonates a human in a manner that deceives a consumer and that the consumer has not consented to. California SB 942 requires that large generative AI providers embed latent provenance disclosures in AI-generated content; as a deployer using those providers' models, Cedar Line Creative must ensure that disclosure obligation flows through to end-users. The FTC's Section 5 enforcement position, reinforced by its January 2026 dedicated AI enforcement unit, treats undisclosed AI interaction as a deceptive practice where a reasonable consumer would not expect it.

Required wording — to appear at the very start of every chat session, before the assistant sends any substantive message, on every client website where Cedar Line Creative has deployed its AI chat assistant:

'Hello. I am an automated AI assistant provided by [CLIENT BRAND]. I am not a human. This conversation may be processed by artificial intelligence systems, including large language model technology. If you would prefer to speak with a human member of the [CLIENT BRAND] team, please [select the option below / email us at [CLIENT CONTACT EMAIL] / call us on [CLIENT PHONE NUMBER]].'

This wording must appear in the same font size and colour contrast as the opening chat message. It must not be buried in a pre-chat privacy notice that requires scrolling before the conversation begins. The disclosure is required to be conspicuous — see Section 5 below for the standard that applies.

If the chat assistant is capable of sending voice output rather than text, an equivalent spoken disclosure must be the first words a caller hears, before any menu option or question is posed. This is consistent with the FTC's stated position that AI interaction disclosures must precede the substantive exchange, not follow it.

4. Disclosure for AI-Generated Content and Synthetic Media

Legal basis: California SB 942, in force 2 August 2026, requires that AI-generated content carry latent provenance disclosures and that the generative AI providers offer a free detection tool. Cedar Line Creative is a deployer, not a large generative AI provider, but it publishes AI-generated outputs under client brands to California consumers, bringing those outputs within the practical reach of SB 942's transparency framework. Texas HB 149, in force 1 January 2026, requires disclosure where AI-generated content could deceive a reasonable person as to its origin. The FTC's first enforcement action targeting undisclosed AI-generated advertising content, brought in the period covered by the current briefing, confirms that non-disclosure of AI-generated advertising copy and imagery constitutes a deceptive practice under Section 5.

For AI-generated written content (copy, captions, scripts, email text) published under any client brand, the following disclosure must appear:

'Some or all of the content on this page [in this communication] was created with the assistance of artificial intelligence tools.'

This must appear in proximity to the content — at the foot of a webpage, within the body of an email above the unsubscribe line, or adjacent to the relevant section of a brochure or advertisement. It must not appear only in a privacy policy or terms of service page that a typical consumer would not encounter during normal use.

For AI-generated images and synthetic visual media (produced using Midjourney or DALL-E) published under any client brand, the following disclosure must appear in immediate visual proximity to each such image:

'Image created using artificial intelligence.'

Where a campaign consists entirely of AI-generated imagery, a single prominent notice at the opening of the campaign page or the start of an advertisement — 'The images in this [campaign / advertisement / post] were created using artificial intelligence tools.' — satisfies the proximity requirement for that campaign unit.

Cedar Line Creative must instruct each client in writing, as part of its standard creative brief or delivery note, that these disclosures must be maintained and must not be removed from published materials. Cedar Line Creative remains responsible for ensuring that content it generates is delivered with the required disclosure language attached.

5. Conspicuousness Standard — Where Placement Must Be More Than Merely Present

Several of the applicable regimes impose a standard of conspicuousness, not merely the presence of a disclosure. Cedar Line Creative must apply the higher standard at the following touchpoints.

Chat assistant — The AI identity disclosure (Section 3 above) must be conspicuous. The FTC's standard under Section 5, applied in its 2025 enforcement action on undisclosed AI advertising, requires that a disclosure be clear and prominent: it must be in a font size that is at least equal to surrounding text; it must not be in a colour or shade that renders it less legible than surrounding text; and it must appear before the consumer takes any action that relies on the content. A disclosure placed only at the bottom of a chat window after a consumer has already typed a question does not meet this standard.

AI-generated advertising content directed at California consumers — California SB 942 requires latent provenance data to be embedded by the generative AI provider; as a deployer, Cedar Line Creative must not strip or suppress that latent data. The FTC's proposed policy statement of 1 July 2026 expressly warns that altering AI outputs to suppress accuracy indicators, even in purported compliance with other requirements, may itself constitute consumer deception. Cedar Line Creative must therefore ensure that no post-processing step applied to ChatGPT, Claude, Midjourney or DALL-E outputs removes embedded provenance metadata.

Recruitment — where AI screening affects applicants in Illinois, the Illinois AI Video Interview Act requires that applicants be notified before the interview that AI may be used to analyse their responses. This notice must be given in advance of the interview event itself, not embedded in a paragraph of general application terms. EEOC guidance on AI in employment reinforces that pre-process disclosure is the expected standard.

In all cases, 'conspicuous' means: a reasonable consumer would see and understand the disclosure without actively searching for it. If a disclosure requires a consumer to click an expand button, scroll past a fold, or open a secondary document, it is not conspicuous.

6. Route for a Consumer to Reach a Human

Texas HB 149 and the FTC Section 5 framework both require that where an AI system is deployed in a consumer-facing role, consumers must have a clear and accessible route to a human contact. This obligation applies at Touchpoints 1, 2 and 3 above — that is, wherever a consumer may interact with or respond to AI-generated content or an AI chat assistant.

For the AI chat assistant on client websites (Touchpoint 3), the human escalation route must be offered within the opening disclosure message (see Section 3 above) and must remain available as a persistent option throughout the chat session. Cedar Line Creative must configure the chat assistant so that typing or selecting 'speak to a human', 'talk to a person' or equivalent phrases reliably and immediately presents the consumer with the escalation option rather than an AI-generated restatement of available topics. The escalation route must be one of the following, or a combination: a direct email address monitored by a human member of the client's or Cedar Line Creative's team; a telephone number staffed during published business hours; or a callback request form that results in a human response within one business day. The specific route must be populated with working contact details before the disclosure goes live; placeholder text is not acceptable.

For AI-generated written content published under client brands (Touchpoint 1), where that content includes a call to action or an offer, the contact details provided in the call to action must connect the consumer to a human and must not route to a further AI system without a renewed disclosure.

For AI-generated imagery (Touchpoint 2), no direct human escalation route is required solely by reason of the image being AI-generated, because imagery does not itself constitute an interactive channel. However, where imagery forms part of a landing page or advertisement that incorporates interactive elements, the human escalation requirement for those interactive elements applies.

Cedar Line Creative must document the human escalation route for each client deployment and retain that record as part of its disclosure log (see Section 7).

7. Recruitment AI Disclosure — Applicants to Cedar Line Creative

Cedar Line Creative uses an AI screening tool in its own hiring process. Applicants who are based in or apply for roles based in California, Texas, Illinois or Colorado are covered by overlapping disclosure obligations.

Texas HB 149 (in force 1 January 2026) requires that where an AI system is used in a consequential decision about a person — including employment decisions — the person must be notified that an automated decision-making system is in use. Cedar Line Creative must include the following wording in every job advertisement and application acknowledgement:

'Cedar Line Creative uses AI-assisted screening tools as part of its recruitment process. Automated systems may analyse your application, CV or assessment responses to support shortlisting decisions. You have the right to request information about how automated screening was used in the assessment of your application by contacting [HR CONTACT EMAIL].'

The Illinois AI Video Interview Act requires that before any AI-analysed video interview, applicants receive written notice that AI will be used to evaluate their responses, together with an explanation of how the AI works and what general characteristics it evaluates. Cedar Line Creative must send this notice by email at the point of scheduling any video interview component, before the applicant records or joins the session.

EEOC guidance on AI in employment (operative and applied in FTC and EEOC enforcement) holds employers responsible for discriminatory outcomes produced by third-party AI screening tools. Cedar Line Creative must obtain from its AI screening tool vendor a current bias audit or validation report, retain it on file, and review it at least annually. The disclosure to applicants must not imply that AI screening is the sole basis for any decision; Cedar Line Creative must ensure that a human reviewer is involved in all hiring decisions.

Colorado SB 26-189 is not yet in force (commencement 1 January 2027) but will, on commencement, impose deployer duties on Cedar Line Creative as a Colorado-headquartered business using automated decision-making technology in consequential decisions. Cedar Line Creative should ensure its recruitment disclosure wording is updated before 1 January 2027 to meet the additional notice and appeal-right requirements that SB 26-189 will introduce.

8. Record of When Each Disclosure Went Live

Cedar Line Creative has confirmed it has no disclosure wording currently in place. Every disclosure identified in this document is therefore a new implementation. The following log must be completed by the responsible person (Head of Operations or nominated AI Compliance Lead) as each disclosure is deployed. The completed log must be retained for a minimum of four years from the date of first deployment, as this covers the limitation periods under the applicable state enforcement regimes.

Disclosure 1 — AI chat assistant identity disclosure on client websites. Target implementation date: [TO BE COMPLETED]. Actual go-live date: [TO BE COMPLETED]. Clients covered at go-live: [LIST CLIENT NAMES]. Person who confirmed go-live: [NAME AND TITLE].

Disclosure 2 — AI-generated written content notice. Target implementation date: [TO BE COMPLETED]. Actual go-live date: [TO BE COMPLETED]. Clients and channels covered at go-live: [LIST]. Person who confirmed go-live: [NAME AND TITLE].

Disclosure 3 — AI-generated imagery notice. Target implementation date: [TO BE COMPLETED]. Actual go-live date: [TO BE COMPLETED]. Clients and channels covered at go-live: [LIST]. Person who confirmed go-live: [NAME AND TITLE].

Disclosure 4 — Recruitment AI screening notice in job advertisements and acknowledgements. Target implementation date: [TO BE COMPLETED]. Actual go-live date: [TO BE COMPLETED]. Roles and platforms covered: [LIST]. Person who confirmed go-live: [NAME AND TITLE].

Disclosure 5 — Illinois AI Video Interview Act pre-interview notice. Target implementation date: [TO BE COMPLETED]. Actual go-live date: [TO BE COMPLETED]. Person who confirmed go-live: [NAME AND TITLE].

This log must be updated whenever a new client deployment is added, whenever a disclosure is materially amended, and whenever a client removes or alters disclosure wording that Cedar Line Creative has provided. A copy of the disclosure wording in use at each date must be archived alongside the log entry. Where a regulator, insurer or counterparty requests evidence of disclosure compliance, this log and its archived wording copies constitute Cedar Line Creative's primary evidence.

9. Regimes That Do Not Apply or Apply Only Conditionally

Virginia CDPA: Virginia's Consumer Data Protection Act contains provisions relating to automated decision-making that could in principle reach Cedar Line Creative if it processes personal data of Virginia residents. Cedar Line Creative's declared operations are in Colorado, California, Texas and Illinois; no Virginia nexus has been identified. Virginia CDPA disclosure obligations are not triggered on the facts as stated. If Cedar Line Creative acquires Virginia-based clients or staff, this position must be reassessed.

Illinois BIPA: The Illinois Biometric Information Privacy Act applies to the collection of biometric identifiers — fingerprints, voiceprints, facial geometry and similar data. BIPA obligations are triggered only if Cedar Line Creative's AI HR screening tool captures biometric data from applicants. Cedar Line Creative must confirm with its screening tool vendor whether biometric data is collected or processed. If it is, separate BIPA-specific written consent and retention-schedule obligations apply and must be implemented immediately; those obligations are outside the scope of this consumer-facing disclosure document but must be addressed in Cedar Line Creative's broader compliance pack.

NIST AI RMF 1.0: The NIST AI Risk Management Framework is a voluntary framework, not a legally enforceable instrument. It does not impose disclosure obligations on Cedar Line Creative. Cedar Line Creative's adoption of NIST AI RMF practices is addressed in other documents in this compliance pack.

Federal comprehensive AI law: The United States has no single federal AI transparency statute as of 11 August 2026. FTC Section 5 enforcement — addressed throughout this document — is the primary federal instrument imposing disclosure obligations on Cedar Line Creative.

Data Processing Documentation for AI Tools

The written record of what each AI tool does with data — the document requested first in any enquiry or enterprise security review.

ChatGPT / OpenAI

CEDAR LINE CREATIVE — DATA PROCESSING DOCUMENTATION
Tool: ChatGPT / OpenAI
Prepared: 11 August 2026

DATA COLLECTED AND PURPOSE
Cedar Line Creative inputs client briefs, brand guidelines, draft copy, audience personas, and campaign data into ChatGPT. Outputs include marketing copy, email scripts, social content, and communications published under client brands. Staff also use the tool for internal drafting and research summaries. Where client briefs include named individuals or contact details, personal data is transferred to OpenAI servers in the course of use.

LEGAL BASIS / BUSINESS JUSTIFICATION
Processing is justified under the legitimate interests of Cedar Line Creative and its clients in delivering contracted marketing services. Where any personal data relates to identifiable California-resident consumers, Virginia CDPA data-subject rights obligations attach. Colorado SB 26-189 is not yet in force (effective 1 January 2027); Cedar Line Creative must be ready to classify ChatGPT use as a deployer function under that regime before that date where outputs influence consequential decisions.

GEOGRAPHICAL STORAGE AND PROCESSING
OpenAI processes data on infrastructure primarily located in the United States. Under the OpenAI Business and Enterprise terms, data is not used to train models by default and is processed on US-based servers. Cedar Line Creative has not contracted for EU data residency; all processing is domestic to the US, which is consistent with the company's staff and client footprint across Colorado, California, Texas, and Illinois.

RETENTION AND DELETION
Under OpenAI's API and Enterprise terms, conversation data is retained for up to 30 days for safety monitoring unless zero-data-retention is enabled by contract. Cedar Line Creative must confirm with its OpenAI account representative whether zero-data-retention is active on its current plan. On termination of the OpenAI agreement, OpenAI commits to delete customer data within 30 days of account closure, subject to any legal hold obligations on OpenAI's side.

MODEL TRAINING
OpenAI's Enterprise and API terms state that inputs and outputs are not used to train OpenAI's models by default. Cedar Line Creative relies on this contractual assurance. Staff must not use the free consumer tier of ChatGPT for client work, as the consumer terms do not provide the same training-opt-out guarantee.

SUB-PROCESSORS
OpenAI discloses a sub-processor list covering cloud infrastructure providers including Microsoft Azure. Cedar Line Creative has reviewed OpenAI's published sub-processor page as of the date of this document. Changes to that list are notified via OpenAI's changelog; Cedar Line Creative's IT lead is responsible for monitoring updates.

SECURITY MEASURES
OpenAI asserts SOC 2 Type II certification, encryption in transit (TLS 1.2+) and at rest (AES-256), and access controls including role-based permissions. Verification was conducted by review of OpenAI's published security documentation and the SOC 2 report summary available to enterprise customers. Cedar Line Creative has not commissioned an independent audit of OpenAI's controls.

Claude / Anthropic

CEDAR LINE CREATIVE — DATA PROCESSING DOCUMENTATION
Tool: Claude / Anthropic
Prepared: 11 August 2026

DATA COLLECTED AND PURPOSE
Cedar Line Creative uses Claude for long-form content drafting, copy editing, creative concepting, and communications strategy documents. Inputs may include client-supplied audience data, campaign briefs, and internal performance summaries. Where those inputs contain personal data about identifiable individuals — for example, named client contacts or employee performance notes — that data is transferred to Anthropic's infrastructure.

LEGAL BASIS / BUSINESS JUSTIFICATION
Processing rests on legitimate interests in service delivery. Illinois BIPA is not directly triggered by text-only inputs unless biometric identifiers are present; Cedar Line Creative does not currently pass biometric data through Claude. Colorado SB 26-189 deployer duties apply from 1 January 2027; Cedar Line Creative must assess before that date whether any Claude output influences a consequential decision as defined by that statute.

GEOGRAPHICAL STORAGE AND PROCESSING
Anthropic processes data on US-based cloud infrastructure. Under Anthropic's commercial API terms, data is not transferred outside the US by default. Cedar Line Creative's use is via the commercial API or Claude for Work product; consumer-tier terms must not be used for client data.

RETENTION AND DELETION
Anthropic's commercial terms provide that prompt and completion data is not retained beyond the immediate API transaction except for a short abuse-monitoring window (the precise duration should be confirmed in Cedar Line Creative's current order form). On termination, Anthropic commits to deletion of customer data in accordance with its data processing addendum. Cedar Line Creative must obtain and execute a current data processing addendum with Anthropic.

MODEL TRAINING
Anthropic's API and Claude for Work terms state that customer inputs and outputs are not used to train Anthropic's models without explicit consent. Cedar Line Creative relies on this contractual commitment and must ensure no staff use the free consumer product for client or employee data.

SUB-PROCESSORS
Anthropic's sub-processor list includes major cloud providers for compute and storage. Cedar Line Creative's legal lead is responsible for reviewing Anthropic's sub-processor disclosure page at contract renewal and when notified of changes.

SECURITY MEASURES
Anthropic asserts SOC 2 Type II compliance, encryption in transit and at rest, and an internal red-teaming and safety programme. Verification was conducted by review of Anthropic's published trust and security page. Cedar Line Creative has not independently audited these controls and relies on Anthropic's self-certification supplemented by the SOC 2 report made available to commercial customers.

AI image generation tools (Midjourney, DALL-E)

CEDAR LINE CREATIVE — DATA PROCESSING DOCUMENTATION
Tool: AI Image Generation Tools (Midjourney; DALL-E via OpenAI)
Prepared: 11 August 2026

DATA COLLECTED AND PURPOSE
Cedar Line Creative uses Midjourney and DALL-E to generate visual assets — illustrations, concept imagery, and campaign visuals — published under client brands. Inputs are text prompts and, in some workflows, reference

images. Where reference images depict real individuals or contain embedded metadata identifying persons, personal data may be processed. Outputs are AI-generated images that Cedar Line Creative publishes or supplies to clients.

LEGAL BASIS / BUSINESS JUSTIFICATION
Processing rests on legitimate interests in creative service delivery. California SB 942 (the AI Transparency Act, in force 2 August 2026) imposes obligations on large generative AI providers to embed latent provenance disclosure in AI-generated images and to offer a free detection tool. Cedar Line Creative, as a deployer rather than a provider, must ensure that images generated through these tools carry compliant provenance markers before publication. The FTC's Section 5 enforcement posture, reinforced by the FTC's July 2026 proposed policy statement, means that publishing AI-generated imagery under client brands without adequate disclosure where material to consumers is a deception risk. Cedar Line Creative's current position of having no disclosure wording is non-compliant and must be remedied immediately.

Illinois BIPA applies if any reference image input contains a scan of a face geometry of an Illinois resident; Cedar Line Creative must not input photographs of identifiable individuals into these tools without a BIPA-compliant written release and a defined retention schedule.

Texas HB 149 (in force 1 January 2026) includes prohibited uses of AI and imposes disclosure requirements for AI-generated content in certain contexts; Cedar Line Creative must review whether client-facing AI imagery on Texas-targeted campaigns triggers disclosure obligations under that statute.

GEOGRAPHICAL STORAGE AND PROCESSING
Midjourney processes prompts and generates images on its own cloud infrastructure, predominantly US-based. DALL-E is accessed via the OpenAI API and is processed on US-based infrastructure as documented under the ChatGPT / OpenAI entry above. Generated images are delivered to Cedar Line Creative and stored on Cedar Line Creative's own systems or client delivery platforms.

RETENTION AND DELETION
Midjourney retains prompts and generated images; the default is that outputs are visible in public galleries unless a paid private-mode plan is active. Cedar Line Creative must confirm it is operating under a Midjourney plan with private-mode enabled for all client work. For DALL-E, OpenAI's API terms govern retention as described in the ChatGPT / OpenAI entry. Cedar Line Creative should contractually confirm deletion timelines with Midjourney on account closure.

MODEL TRAINING
Midjourney's terms of service state that prompts and outputs may be used to improve its models unless a private or enterprise plan explicitly excludes this. Cedar Line Creative must ensure it holds a plan tier that contractually excludes training use for client content. DALL-E inputs and outputs are not used for training under OpenAI's API terms as noted above.

SUB-PROCESSORS
Midjourney uses cloud infrastructure sub-processors; its current sub-processor disclosure should be reviewed by Cedar Line Creative's IT lead. DALL-E sub-processors are as documented for OpenAI.

SECURITY MEASURES
Midjourney asserts standard cloud security controls; it does not publish a SOC 2 report as of the date of this document. Cedar Line Creative has verified this by reviewing Midjourney's published terms and trust documentation. The absence of a SOC 2 report is a residual risk Cedar Line Creative accepts in exchange for operational use of the tool; this risk is logged in Cedar Line Creative's AI risk register. DALL-E security measures are as documented for OpenAI.

AI-powered CRM or sales tools

CEDAR LINE CREATIVE — DATA PROCESSING DOCUMENTATION
Tool: AI-Powered CRM or Sales Tools
Prepared: 11 August 2026

DATA COLLECTED AND PURPOSE
Cedar Line Creative's AI-powered CRM processes contact records for clients, prospects, and leads, including names, email addresses, company details, engagement history, and sales pipeline data. AI features within the CRM are used for lead scoring, next-action recommendations, automated follow-up drafting, and revenue forecasting. Where contacts are California residents, Virginia CDPA data-subject rights obligations attach to processing of personal data. Where contacts are Illinois residents, BIPA is not directly triggered unless the CRM processes biometric identifiers, which Cedar Line Creative confirms it does not.

LEGAL BASIS / BUSINESS JUSTIFICATION
Processing rests on legitimate interests in business development and client relationship management. Colorado SB 26-189, in force 1 January 2027, may classify AI-generated lead scores or sales recommendations as outputs of automated decision-making technology if they materially affect commercial relationships; Cedar Line Creative must assess this before that date. Texas HB 149, in force 1 January 2026, imposes obligations relevant to AI systems used in business contexts; Cedar Line Creative must confirm with its Texas legal counsel whether the CRM's AI scoring features fall within scope.

GEOGRAPHICAL STORAGE AND PROCESSING
The specific CRM vendor used by Cedar Line Creative must be named in the version of this document supplied to any regulator or enterprise security reviewer. [Vendor name to be confirmed by Cedar Line Creative's operations lead.] Data is expected to be processed on US-based infrastructure; Cedar Line Creative must confirm the data residency commitment in its vendor contract.

RETENTION AND DELETION
The CRM vendor's data retention terms govern how long contact records and AI-derived scores are held. Cedar Line Creative must confirm the contractual retention period and ensure that a data deletion clause is included in its vendor agreement covering deletion of Cedar Line Creative's data within a defined period — recommended 30 days — following termination.

MODEL TRAINING
Cedar Line Creative must confirm in writing with the CRM vendor whether contact data or interaction history is used to train or improve the vendor's AI models. If the vendor's standard terms permit training use, Cedar Line Creative must negotiate an opt-out or accept the risk and document that decision.

SUB-PROCESSORS
The CRM vendor's sub-processor list must be reviewed by Cedar Line Creative's IT lead. Standard sub-processors for major CRM platforms include cloud infrastructure and email delivery providers. Cedar Line Creative must obtain the current sub-processor list and retain it with this document.

SECURITY MEASURES
Cedar Line Creative must obtain the CRM vendor's current security certification — at minimum SOC 2 Type II or equivalent — and retain a copy. The vendor's asserted controls for encryption in transit and at rest, access management, and incident response must be verified by review of the vendor's security documentation or a completed security questionnaire. Cedar Line Creative has not independently audited the CRM vendor's controls.

AI in HR or recruitment systems (AI screening tool)

CEDAR LINE CREATIVE — DATA PROCESSING DOCUMENTATION
Tool: AI in HR or Recruitment Systems (AI Screening Tool)
Prepared: 11 August 2026

DATA COLLECTED AND PURPOSE

Cedar Line Creative uses an AI screening tool in its own hiring process. The tool processes job applications, CVs, and, potentially, responses to structured screening questions submitted by candidates. Outputs include ranked shortlists, scores, or pass/fail recommendations used by Cedar Line Creative's HR team in deciding which candidates progress. Data subjects are job applicants, who may be resident in Colorado, California, Texas, Illinois, or other states depending on the role.

This is the highest-risk AI use in Cedar Line Creative's portfolio from a regulatory standpoint. It is an AI system making or substantially influencing decisions about people in an employment context.

LEGAL BASIS / BUSINESS JUSTIFICATION

Processing rests on the legitimate interests of Cedar Line Creative in efficient recruitment, subject to overriding obligations under employment law and AI-specific statutes.

EEOC AI in Employment Guidance: The EEOC has confirmed that employers remain liable under Title VII and the ADA for discriminatory outcomes produced by AI screening tools even where those tools are operated by a third-party vendor. Cedar Line Creative must audit the screening tool for adverse impact against protected classes (race, sex, disability, age) before each use cycle and retain those audit results.

Colorado SB 26-189 (in force 1 January 2027): Recruiting and hiring is expressly listed as a consequential decision context under this statute. From 1 January 2027, Cedar Line Creative as a deployer must: conduct an impact assessment before deploying the tool; provide candidates with notice that AI is used in screening; and offer a mechanism to appeal or seek human review of an AI-generated decision. Cedar Line Creative must complete its impact assessment and prepare candidate notice language before 1 January 2027.

Illinois BIPA: If the screening tool processes any biometric data — including voiceprints from recorded interviews or facial geometry from video screening — from Illinois-resident applicants, BIPA's written consent, retention schedule, and destruction requirements apply. Cedar Line Creative must confirm with the vendor whether any biometric data is collected. If it is, a BIPA-compliant written policy and individual written releases are required before any Illinois-resident applicant data is processed.

California SB 942 and the broader California privacy framework: California-resident applicants have rights under California privacy law; Cedar Line Creative must include AI use disclosure in its California applicant privacy notice.

Texas HB 149 (in force 1 January 2026): Cedar Line Creative must confirm with its Texas counsel whether the use of AI screening for Texas-based roles or Texas-resident applicants triggers any disclosure or prohibited-use obligations under that statute.

GEOGRAPHICAL STORAGE AND PROCESSING

The screening tool vendor's data residency commitment must be confirmed in Cedar Line Creative's vendor contract. Applicant data — which is sensitive personal data — must be processed on infrastructure with documented security controls. Cedar Line Creative must confirm that data does not leave the US without a compliant transfer mechanism.

RETENTION AND DELETION

Applicant data must be retained only for as long as necessary for the recruitment process and any post-hire dispute window. Cedar Line Creative's HR policy must specify a retention period — recommended maximum two years from application date for unsuccessful candidates — and the screening tool vendor must contractually commit to deleting Cedar Line Creative's data within 30 days of termination or on Cedar Line Creative's deletion request. Retention periods must be disclosed to applicants in the applicant privacy notice.

MODEL TRAINING

Cedar Line Creative must confirm in writing with the screening tool vendor whether applicant data is used to train or retrain the vendor's AI models. Given the sensitivity of employment data and the EEOC's position on employer liability for AI-driven bias, Cedar Line Creative must not permit applicant data to be used for model training without explicit, documented consent from each applicant and a prior adverse-impact assessment of the resulting model.

SUB-PROCESSORS

The vendor's sub-processor list must be obtained and reviewed. Cedar Line Creative must ensure that any sub-processor handling applicant data is bound by equivalent data protection obligations.

SECURITY MEASURES

Cedar Line Creative must obtain the screening tool vendor's SOC 2 Type II report or equivalent certification. The vendor's asserted controls must include encryption in transit and at rest, role-based access controls limiting HR staff access to applicant records, and an incident response plan covering notification to Cedar Line Creative in the event of a breach. Verification was conducted by review of vendor-supplied security documentation; Cedar Line Creative has not independently audited these controls. The vendor's security assurances must be re-verified at each annual contract renewal.

Multi-State Compliance Matrix

One grid covering every state that reaches you, so a single control can be seen to satisfy several laws at once.

Obligation	Colorado SB 26-189	Illinois BIPA	Texas HB 149	California SB 942 / AB 853	Virginia CDPA	NIST AI RMF 1.0	FTC § 5 / AI Enforcement	EEOC AI Employment Guidance	Strictest Default	Cedar Line Action Required
Effective / In-Force Date	1 Jan 2027 (not yet in force)	In force; BIPA enacted 2008, ongoing litigation	1 Jan 2026 (in force)	2 Aug 2026 (in force); delayed from 1 Jan 2026 by AB 853	In force; CDPA effective 1 Jan 2023	Voluntary framework; published Mar 2023	In force; FTC AI enforcement unit est. Jan 2026	In force; EEOC AI guidance issued 2023–2024	N/A — see individual cells	Immediate action required for TX, CA, IL, VA, EEOC; CO preparation before 1 Jan 2027
Disclosure of AI-generated content to consumers / end users	Not yet in force; deployers must disclose use of ADMT in consequential decisions from 1 Jan 2027	No direct content-disclosure obligation under BIPA	Requires disclosure to consumers when AI systems make or substantially influence decisions affecting them; AG enforcement	STRICTEST — SB 942: large generative AI providers must embed latent provenance disclosure in AI-generated content; Cedar Line publishes AI copy and imagery under client brands	Controllers using automated processing for decisions must inform consumers; Art. 13 CDPA	GOVERN 1.2 and MANAGE 2.2 recommend transparency about AI outputs	Undisclosed AI-generated advertising content targeted in first FTC enforcement action; deceptive omission risk under § 5	No specific content-disclosure obligation	California SB 942	Add AI-content disclosure labels to all client-published copy and imagery; implement latent provenance metadata in image outputs; update client contracts to require downstream disclosure

Obligation	Colorado SB 26-189	Illinois BIPA	Texas HB 149	California SB 942 / AB 853	Virginia CDPA	NIST AI RMF 1.0	FTC § 5 / AI Enforcement	EEOC AI Employment Guidance	Strictest Default	Cedar Line Action Required
Disclosure of AI use in automated / consequential decisions about	From 1 Jan 2027: deployers must notify individuals subject	Not a disclosure obligation; biometric data focus	Entities using AI in decisions affecting consumers must disclose;	SB 942 scope is generative-AI content provenance, not general	Right to opt out of profiling for decisions with legal or similarly	GOVERN 6.1: document and communicate AI decision logic to	Deceptive omission in AI-driven decisions is unfair	STRICTEST — employers must notify job applicants when AI tools	EEOC (hiring); Colorado (consequential decisions from 2027)	Issue written AI-screening notice to all job applicants before or at

Obligation	Colorado SB 26-189	Illinois BIPA	Texas HB 149	California SB 942 / AB 853	Virginia CDPA	NIST AI RMF 1.0	FTC § 5 / AI Enforcement	EEOC AI Employment Guidance	Strictest Default	Cedar Line Action Required
people	to consequential ADMT decisions and provide reasons; §§ 6–8 SB 26-189		AG may seek injunction or civil penalty	decision disclosure	significant effects; § 13(a)(4) CDPA	affected parties	practice under § 5; FTC AI policy statement Jul 2026	screen, rank or reject them; EEOC 2023 guidance		point of application; prepare broader consequential-decision notice template for CO by Q4 2026
Biometric data — collection, consent, retention and destruction	ADMT duties cover biometric data if used in consequential decisions; not biometric-specific	STRICTEST — BIPA §§ 15(a)–(e): written policy, informed written consent, no sale, data-security standard, destruction schedule; $1,000–$5,000 per violation, private right of action	Texas Capture or Use of Biometric Identifier Act (CUBI) separate from HB 149; confirm applicability	CCPA/CPRA treats biometric data as sensitive personal information; separate from SB 942	Biometric data is sensitive data under CDPA; consent required for processing	GOVERN 2.1: inventory sensitive data types including biometrics	Unlawful biometric collection without consent is unfair practice	AI screening tools using facial analysis or voice biometrics trigger EEOC and BIPA scrutiny	Illinois BIPA	Audit HR AI screening tool immediately: if it processes facial, voice or fingerprint data of Illinois applicants, obtain written BIPA consent, publish retention policy, and set destruction schedule
HR / recruitment — prohibition on discriminatory AI screening	Algorithmic discrimination in consequential employment decisions prohibited from 1 Jan 2027; § 5 SB 26-189	BIPA does not prohibit discrimination but biometric-based screening tools trigger consent obligations that effectively	HB 149 prohibits use of AI that discriminates on protected characteristics; AG enforcement from 1 Jan 2026	FEHA (separate California law) prohibits discriminatory AI in hiring; SB 942 scope is content provenance	CDPA does not contain specific employment AI discrimination provisions beyond general sensitive-data rules	BIAS 2.2 and EVAL 2.5: test employment AI for disparate impact before deployment	Employer use of AI screening tools with disparate impact is an unfair practice; FTC/EEOC coordination	STRICTEST — EEOC guidance: employer responsible for discriminatory AI screening vendor outputs; duty to assess	EEOC (immediate); Texas HB 149; Colorado SB 26-189 (from 2027)	Commission independent bias audit of AI screening tool; obtain vendor's disparate-impact testing data; document r

Obligation	Colorado SB 26-189	Illinois BIPA	Texas HB 149	California SB 942 / AB 853	Virginia CDPA	NIST AI RMF 1.0	FTC § 5 / AI Enforcement	EEOC AI Employment Guidance	Strictest Default	Cedar Line Action Required
		gate use						adverse impact; Title VII and ADA apply		emediation steps; repeat audit annually
Risk assessment / impact assessment for high-risk AI	STRICTEST — from 1 Jan 2027: deployers must conduct and document impact assessments for ADMT in consequential decisions; § 9 SB 26-189	No risk-assessment requirement under BIPA	Deployers of high-risk AI must document risk management practices; § 4 HB 149	SB 942 does not mandate risk assessments; content-provenance focus	Data protection assessments required for processing presenting heightened risk including profiling; § 10 CDPA	IDENTIFY 1.0 and GOVERN 4.1: risk tiering and documented impact assessment for each AI use case	FTC expects documented risk assessment to substantiate AI claims; absence of documentation is evidence of unfairness	EEOC expects employers to evaluate AI tools for adverse impact risk before deployment	Colorado SB 26-189 (from 2027); Virginia CDPA (immediate)	Complete CDPA data-protection assessments now for CRM, HR and chatbot uses; build CO-compliant ADMT impact-assessment template and run assessments before 1 Jan 2027
Transparency / consumer-facing AI disclosure (chatbot / automated system)	ADMT disclosure to individuals required from 1 Jan 2027 for consequential decisions	No chatbot-specific disclosure rule under BIPA	STRICTEST (operational) — HB 149: entities deploying AI chat assistants on consumer-facing platforms must disclose AI involvement; in force 1 Jan 2026	SB 942 applies to generative AI content provenance; chatbot outputs that are AI-generated content must carry latent disclosure	CDPA: automated processing disclosures required where profiling has significant effect	COMMUNICATE 1.0: users should know when they are interacting with an AI system	Representing an AI chatbot as a human agent is deceptive under § 5; FTC July 2026 policy statement	No chatbot-specific EEOC requirement	Texas HB 149 (in force); California SB 942 (in force)	Add clear 'This chat is powered by AI' notice on all client website chatbots before any user interaction; update client deployment agreements to mandate this notice

Obligation	Colorado SB 26-189	Illinois BIPA	Texas HB 149	California SB 942 / AB 853	Virginia CDPA	NIST AI RMF 1.0	FTC § 5 / AI Enforcement	EEOC AI Employment Guidance	Strictest Default	Cedar Line Action Required
Intellectual property — ownership and disclosure of AI-generated outputs	No IP-specific provision in SB 26-189	No IP provision in BIPA	No IP-specific provision in HB 149	No IP provision in SB 942 beyond provenance metadata	No IP provision in CDPA	GOVERN 3.2: document provenance of AI-generated artefacts; maintain records of model and version used	STRICTEST — FTC: marketing AI-generated content as wholly original human creative work without disclosure is a deceptive practice under § 5; AI-washing enforcement active	No EEOC IP provision	FTC § 5	Insert IP ownership clauses in all client contracts: AI outputs are works made for hire only to the extent US copyright law permits; disclaim warranty of originality; document model and prompt for each deliverable
Data subject rights — access, correction, deletion relating to AI processing	From 1 Jan 2027: individuals have right to appeal consequential ADMT decisions; § 10 SB 26-189	BIPA § 15(a): right to know of biometric data held; no broader access/deletion right under BIPA alone	HB 149 does not create general consumer data rights; those sit under Texas Data Privacy and Security Act (TDPSA)	CCPA/CPRA (separate from SB 942) provides access, correction, deletion rights; SB 942 adds AI-detection tool right	STRICTEST — CDPA §§ 8–13: right to access, correct, delete, and opt out of profiling; controller must respond within 45 days	GOVERN 2.2: data-subject rights requests should be supported by AI-system documentation	Failure to honour data rights in AI-processed records may constitute unfair practice	No specific data-rights provision in EEOC guidance beyond access to adverse employment decisions	Virginia CDPA	Map all personal data processed through CRM AI and HR AI tools; build single intake process for access/correction/deletion/opt-out requests; set 45-day SLA; log all requests and responses

Obligation	Colorado SB 26-189	Illinois BIPA	Texas HB 149	California SB 942 / AB 853	Virginia CDPA	NIST AI RMF 1.0	FTC § 5 / AI Enforcement	EEOC AI Employment Guidance	Strictest Default	Cedar Line Action Required
Vendor / third-party AI governance obligations	From 1 Jan 2027: deployers remain liable for ADMT vendor outputs in consequential decisions; must conduct due diligence; § 7 SB 26-189	Biometric data vendors must satisfy BIPA § 15(d) data-security and transfer requirements	Deployers responsible for ensuring AI vendors comply with HB 149 prohibited-use provisions	SB 942 obligations fall on the generative AI developer (OpenAI, Anthropic, Midjourney); Cedar Line as deployer must verify tools carry required provenance capabilities	STRICTEST — CDPA: processors (vendors) must be bound by data-processing agreements meeting § 25 requirements; controller retains accountability	GOVERN 6.2: maintain AI vendor inventory; conduct periodic supplier assessments	FTC holds principal liable for vendor AI conduct in marketing; 'AI-washing' liability extends to claims made by vendors on client's behalf	EEOC: employer cannot shift liability to AI vendor for discriminatory screening; vendor contract should include indemnity and bias-testing warranties	Virginia CDPA (immediate); Colorado SB 26-189 (from 2027)	Execute CDPA-compliant data-processing agreements with OpenAI, Anthropic, Midjourney, CRM vendor and HR AI vendor; obtain bias-testing warranties from HR AI vendor; build vendor AI-risk register
Prohibited AI uses	From 1 Jan 2027: ADMT may not be used to make consequential decisions using protected characteristics as primary factors; § 5 SB 26-189	No explicit prohibited-use list under BIPA; collection without consent is itself prohibited	STRICTEST (operational) — HB 149 § 3: explicit list of prohibited AI uses including social scoring, manipulation of vulnerable individuals, and real-time biometric categorisation by government; in force 1 Jan 2026	SB 942 does not contain a prohibited-use list	CDPA prohibits processing sensitive data without consent; no AI-specific prohibited-use list	GOVERN 1.5: organisations should identify and document categories of unacceptable AI use	Any AI use constituting unfair, deceptive or abusive act or practice is prohibited under § 5; manipulationof consumers via AI is priority target	Use of AI screening tools that screen out protected classes is prohibited under Title VII / ADA regardless of intent	Texas HB 149	Review all AI tool configurations against HB 149 § 3 prohibited-use list; document that Cedar Line's chatbot and CRM AI do not perform social scoring orexploit vulnerable persons; retain evidence

Obligation	Colorado SB 26-189	Illinois BIPA	Texas HB 149	California SB 942 / AB 853	Virginia CDPA	NIST AI RMF 1.0	FTC § 5 / AI Enforcement	EEOC AI Employment Guidance	Strictest Default	Cedar Line Action Required
Enforcement, penalties and liability exposure	From 1 Jan 2027: CO Attorney General enforcement; no private right of action under SB 26-189	STRICTEST — BIPA: private right of action; $1,000 per negligent violation, $5,000 per intentional violation; class actions common; no cap	TX Attorney General: civil penalties up to $10,000 per violation; injunctive relief; HB 149 in force	CA AG and private plaintiffs under CCPA/CPRA for personal data; SB 942 enforcement mechanism — confirm with CA AG guidance	VA AG only; civil penalties up to $7,500 per intentional violation; no private right of action	Voluntary framework; non-compliance not directly penalised but used as evidence of reasonable care	FTC: civil penalties up to $51,744 per violation per day under FTC Act § 5; injunctive relief; consent orders	EEOC: administrative charge, federal lawsuit, back pay, compensatory and punitive damages; no statutory cap in Title VII class actions	Illinois BIPA (private right of action, class risk); FTC (daily penalties)	Obtain BIPA liability insurance rider; document all AI consent records for Illinois individuals; treat FTC daily-penalty risk as priority for advertising and chatbot disclosures
Documentation and record-keeping	From 1 Jan 2027: deployers must retain ADMT impact assessments and decision records; § 9 SB 26-189	BIPA § 15(a): written retention and destruction policy publicly available	HB 149: deployers must maintain records of AI risk-management practices; retention period — confirm with AG guidance	SB 942: generative AI providers must maintain provenance records; Cedar Line should retain records of which tool generated each client deliverable	CDPA § 10: data-protection assessment records must be retained; available to AG on request	STRICTEST (breadth) — NIST AI RMF: comprehensive documentation across GOVERN, MAP, MEASURE and MANAGE functions covering all AI systems	FTC expects documented evidence of AI claims substantiation and disclosure decisions; absence is adverse inference in enforcement	EEOC: retain records of AI-screening criteria, validation studies and adverse-impact analyses per Title VII record-keeping rules (1 year minimum, 2 years for federal contractors)	NIST AI RMF (broadest scope); EEOC (employment-specific)	Stand up an AI System Register in the NIST RMF format documenting all six tool categories; link each record to applicable legal obligations; assign document owner; review quarterly

RATING SCALE AND CONVENTIONS. marks the strictest obligation per row — Cedar Line should use that standard as its safe default across all jurisdictions. Where a regime does not impose an obligation, the cell says so and gives the reason; it is not left blank. Cells referencing Colorado SB 26-189 reflect that it is NOT YET IN FORCE (effective 1 January 2027); Cedar Line must prepare now but is not yet legally exposed under that statute. California SB 942 came into force on 2 August 2026 (delayed from 1 January 2026 by AB 853); compliance is immediately required. Texas HB 149 has been in force since 1 January 2026; Cedar Line is already exposed. Virginia CDPA is in force and reaches Cedar Line because it processes personal data of Virginia residents. BIPA reaches Cedar Line because it employs and recruits individuals based in Illinois and may collect biometric data through the HR AI screening tool. SCOPE NOTE ON VIRGINIA CDPA: Cedar Line should confirm whether it meets the CDPA controller threshold (100,000 consumers/year or 25,000 consumers/year with 50%+ revenue from data sale). If below threshold, CDPA obligations do not apply; document that assessment. SCOPE NOTE ON CALIFORNIA SB 942: primary obligations under SB 942 fall on large generative AI system developers (OpenAI, Anthropic, Midjourney, Adobe). Cedar Line's obligation as deployer is to verify those vendors have embedded latent provenance disclosure in their outputs and to pass through any contractual disclosure requirements to clients. Cedar Line's broader California exposure for personal data sits under CCPA/CPRA (not scoped in this matrix but should be assessed separately). FEDERAL PREEMPTION RISK: Executive Order 14365 (December 2025) created a DOJ AI Litigation Task Force to challenge state AI laws. Cedar Line should monitor whether any of the above state laws are subject to federal preemption challenge, particularly Texas HB 149 and Colorado SB 26-189, without adjusting compliance posture until a court order is issued. FTC JULY 2026 POLICY STATEMENT: the FTC warned on 1 July 2026 that altering AI outputs to suppress accuracy — even to comply with state laws — may constitute consumer deception. Cedar Line must not configure its AI tools to produce materially misleading outputs and should document that all disclosure additions are additive, not suppressive. REVIEW CADENCE: this matrix should be reviewed by Cedar Line's compliance lead quarterly, and immediately upon any new state AI law enactment in Colorado, California, Texas, Illinois or Virginia. DOCUMENT OWNER: to be assigned at Cedar Line; recommended owner is General Counsel or Head of Operations with input from HR lead (BIPA / EEOC rows) and Creative Director (IP / content-disclosure rows). TURNOVER-BASED PENALTIES: where a statutory maximum is calculated as a percentage of annual turnover (not applicable to any regime in this matrix — all penalties above are per-violation or daily figures), that limb cannot be computed without Cedar Line's revenue figures. DISCLAIMER: this matrix is a compliance planning tool produced as at 11 August 2026. It is not legal advice. Cedar Line should engage qualified US counsel before making representations to regulators, insurers or clients based on this document.

HR AI Procedures and Disclosure Templates

The full HR-side procedure set for AI in hiring, management and exit, with the templates to run it.

HR AI Procedures and Disclosure Templates — Cedar Line Creative

Document owner: Chief People Officer (or equivalent senior HR lead). Effective date: 11 August 2026. Review cycle: Annually, or within 30 days of any material change to the AI screening tool, applicable law, or Cedar Line Creative's hiring footprint.

This document governs every use of artificial intelligence in Cedar Line Creative's employment lifecycle, from sourcing through to exit. Cedar Line Creative employs between 51 and 250 people and operates from Denver, Colorado, with staff and candidates located in California, Illinois, and Texas. That geographic spread activates obligations under Colorado SB 26-189 (consequential automated decision-making; in force 1 January 2027), Illinois BIPA (biometric data), Texas HB 149 (Responsible Artificial Intelligence Governance Act; in force 1 January 2026), California SB 942 as amended by AB 853 (AI Transparency Act; in force 2 August 2026), EEOC guidance on AI in employment, and FTC Section 5 (unfair or deceptive acts). NIST AI RMF 1.0 is adopted as the internal risk framework. Virginia CDPA is noted but does not create standalone AI employment obligations for Cedar Line Creative's current operational footprint; it should be monitored if Virginia-based hiring commences.

The AI screening tool used in Cedar Line Creative's own hiring is the primary subject of these procedures. Additional AI tools (ChatGPT, Claude, Midjourney, DALL-E, AI-powered CRM) are used in client-facing work, but their use in HR contexts is also addressed where relevant.

1. Sourcing Stage — Procedure and AI Boundaries

Cedar Line Creative may use AI-powered CRM and sales tools, and general large language models such as ChatGPT or Claude, to draft job descriptions, identify sourcing channels, and generate outreach copy. Use of AI at this stage is permitted subject to the constraints below.

What AI may do: Draft initial job descriptions; suggest sourcing platforms; generate outreach message templates; summarise role requirements from a hiring manager brief.

What AI may not do: Autonomously post roles that exclude protected-class language without human editorial review; rank or score candidate pools drawn from social or professional networks without disclosure; make any decision — including channel selection weighted by demographic inference — that substitutes for human judgement about where to advertise.

Human checkpoint: A qualified HR team member must review every job description and outreach template before publication to confirm that no language, whether AI-generated or otherwise, has a disparate impact on candidates on the basis of race, sex, age, national origin, disability, or any other protected characteristic under Title VII, the ADEA, the ADA, or equivalent state statutes. Under the EEOC's April 2023 guidance on AI in employment (and its ongoing enforcement posture), Cedar Line Creative remains liable for discriminatory outcomes regardless of whether the proximate cause is an algorithm.

Bias check at sourcing: Before any AI tool is used to generate sourcing copy or identify channels, the HR lead must confirm that the tool vendor has provided current equal employment opportunity documentation. Any tool producing sourcing recommendations based on inferred demographic characteristics of the target audience is prohibited. Sourcing outputs are logged (see Records section).

2. Screening Stage — Procedure and AI Boundaries

Cedar Line Creative operates an AI screening tool for its own hiring. This tool reviews applications and produces ranked outputs or pass/fail recommendations before human review. This stage is the highest-risk AI use in Cedar Line Creative's HR lifecycle and attracts the most specific legal obligations.

What AI may do: Parse CVs and applications for declared criteria (skills, qualifications, experience thresholds) specified in advance by HR; produce a ranked or scored shortlist for human review; flag applications that meet minimum threshold criteria.

What AI may not do: Issue a final rejection decision without human review; score candidates on inferred characteristics (personality, motivation, culture fit) unless the tool vendor has published validated, bias-tested methodology for those dimensions; process voice, facial expression, or video data without separate BIPA compliance steps (Illinois candidates) and explicit consent; make a hiring recommendation that a human reviewer cannot override.

Colorado SB 26-189 (effective 1 January 2027): Although not yet in force, Cedar Line Creative must be operationally ready by that date. The statute classifies AI-assisted employment screening as a consequential automated decision and will require: (a) a risk management policy aligned to NIST AI RMF 1.0; (b) pre-deployment and annual post-deployment impact assessments; (c) notice to candidates; and (d) a meaningful human review and appeal mechanism. Cedar Line Creative should treat compliance with SB 26-189 as live preparation now.

Texas HB 149 (in force 1 January 2026): Cedar Line Creative's Texas-based hiring is subject to the Responsible Artificial Intelligence Governance Act. The Act requires that AI systems used in employment contexts operate transparently and without prohibited discriminatory use. Cedar Line Creative must document the AI screening tool's purpose, known limitations, and the human oversight applied.

EEOC AI guidance: The EEOC's guidance affirms that the Uniform Guidelines on Employee Selection Procedures apply to AI screening tools. Cedar Line Creative must be able to produce adverse impact analyses on request. If the tool produces a selection rate for any protected group that is less than 80% of the rate for the highest-selecting group (the four-fifths rule), Cedar Line Creative must investigate and remediate before continuing use.

Bias testing before deployment: Before the AI screening tool is activated for any new role or after any update to the tool or its configuration, HR must run an adverse impact analysis on a test dataset representative of the expected applicant pool. Results must be documented and signed off by the HR lead. If adverse impact is detected, the tool configuration must be adjusted or the tool must not be used for that role.

Bias testing after deployment: Within 30 days of closing each hire made using AI screening, HR must produce a post-hire adverse impact report. This report compares the demographic distribution of AI-passed candidates against the applicant pool. Where adverse impact is detected, HR must suspend AI screening, investigate root cause, and obtain vendor remediation confirmation before resuming. Annual bias audits are conducted every January.

3. Interview Stage — Procedure and AI Boundaries

Cedar Line Creative may use AI tools to assist with interview question design, scheduling, and note transcription. It may not use AI to score or assess candidates during interviews without satisfying the additional requirements below.

What AI may do: Generate structured interview question banks aligned to competency frameworks; automate scheduling; transcribe (with consent) interview notes for the record; provide the interviewer with a structured scoring rubric.

What AI may not do: Conduct automated video analysis of facial expressions, tone, or body language as an assessment input; score interview responses without a validated, bias-tested methodology; produce a post-interview ranking that the human interviewer cannot adjust with documented reasoning.

Illinois BIPA: If Cedar Line Creative interviews Illinois-based candidates using any recording or automated analysis that captures biometric identifiers (including voiceprint or facial geometry), it must: (a) provide written notice of the collection; (b) state the purpose and duration of retention; (c) obtain written release before collection; and (d) not sell, lease, or trade that biometric data. Failure to comply exposes Cedar Line Creative to BIPA's statutory damages of $1,000 per negligent violation and $5,000 per intentional or reckless violation, plus attorney's fees. Cedar Line Creative's current position — no disclosure wording anywhere — is non-compliant. The BIPA notice and consent template in Section 9 must be used for every Illinois candidate where any recording is made.

California SB 942 (in force 2 August 2026): To the extent Cedar Line Creative uses a generative AI tool to produce interview materials or analysis outputs that are shared with California-based candidates, those outputs must carry latent provenance disclosure indicating AI involvement. This obligation falls on Cedar Line Creative as a user of covered generative AI systems; it supplements any obligation placed on the AI provider itself.

Human oversight: Every interview assessment must be completed by a human interviewer who has final authority over the score. AI-generated notes or transcripts are an aid, not a record of assessed performance. The human interviewer's signed scoring sheet is the operative assessment document.

4. Selection Stage — Procedure and AI Boundaries

Selection is the point at which a candidate is offered or denied employment. This is a consequential decision for the purposes of Colorado SB 26-189 (operative from 1 January 2027) and engages the full range of EEOC adverse impact obligations now.

What AI may do: Compile a summary of screening and interview scores for the decision-maker; flag whether a candidate meets the pre-defined selection criteria; generate offer letter drafts for human review.

What AI may not do: Issue an offer or rejection autonomously; rank final candidates by a composite AI score where that score is the sole or primary decision criterion; take any action in the applicant tracking system that moves a candidate to a rejected status without a documented human decision.

Human decision authority: The hiring manager and HR representative must jointly sign off on every offer and every rejection. The documented basis for the decision must refer to the assessed criteria, not merely to an AI score. Where the AI screening recommendation differs from the human decision, the reason for departure must be recorded.

Colorado SB 26-189 readiness: From 1 January 2027, Cedar Line Creative must provide candidates with notice that automated decision-making technology was used in the selection process, explain its principal logic, and offer a process to request human review of any adverse outcome. Cedar Line Creative must build this mechanism into its ATS before that date. The candidate notice template in Section 8 is designed to satisfy this requirement from its date of use.

FTC Section 5: Cedar Line Creative must not represent to candidates (or in any public-facing employer branding) that its hiring process is unbiased or objectively fair solely because it uses AI, unless Cedar Line Creative can substantiate that claim with bias audit results. Unsubstantiated fairness claims constitute AI-washing and are subject to FTC enforcement.

5. Performance Management Stage — Procedure and AI Boundaries

Cedar Line Creative may use AI-powered CRM or business intelligence tools that generate data on employee output relevant to performance reviews. Use of such data in formal performance management is permitted only within the following boundaries.

What AI may do: Aggregate productivity data (volume of content produced, client satisfaction scores, project completion rates) and present it to a manager as one input among several; flag statistical anomalies in output for manager attention; assist HR in drafting performance improvement plan templates.

What AI may not do: Generate a performance rating that is placed in the employee's record without human review and sign-off; make disciplinary recommendations; infer or report on employee sentiment, engagement, or personality from communications data without explicit consent and a legally reviewed purpose.

Human accountability: Every formal performance rating — annual review, mid-year check-in, or ad hoc assessment — must be assigned by a named manager who takes responsibility for the rating. The AI data summary is retained as supporting documentation but is not itself the rating.

Bias risk: Performance AI tools carry a documented risk of encoding historical bias, particularly for employees from groups that have been historically under-represented in high-performance cohorts. Cedar Line Creative's HR lead must review aggregate performance data annually, disaggregated by protected class (to the extent permitted by applicable law and with appropriate privacy controls), to identify patterns suggestive of systematic bias. This review is documented and retained.

Colorado SB 26-189: Performance management affecting compensation, role classification, or continued employment will constitute a consequential decision under SB 26-189 from 1 January 2027. Cedar Line Creative must ensure employees can request human review of any AI-informed performance outcome from that date.

6. Discipline Stage — Procedure and AI Boundaries

Disciplinary processes at Cedar Line Creative are governed by the employee handbook. AI may support but never substitute for the disciplinary procedure.

What AI may do: Compile factual records (time-stamped logs, output metrics, communications summaries where lawfully obtained) as background material for HR; draft disciplinary letters for manager and legal review; check proposed disciplinary outcomes against the documented precedent record for consistency.

What AI may not do: Recommend a specific sanction (warning, suspension, termination); assess the credibility of an employee's account; produce a disciplinary report that is presented to the employee as the basis for action without being reviewed and approved by a qualified HR professional and, where appropriate, legal counsel.

Process integrity: Any AI-generated material used in a disciplinary process must be identified as AI-assisted in the case file. The employee has the right to see and respond to factual evidence used against them; AI-generated summaries of that evidence do not alter this right. Where an AI summary of communications or output data is used, Cedar Line Creative must be able to produce the underlying source data on request.

Illinois, Texas, and California staff: Employees based in these states retain all rights under applicable state employment law. Nothing in the AI-assisted disciplinary process diminishes those rights. Texas HB 149 prohibits the use of AI in ways that result in unlawful discrimination; any AI-assisted disciplinary outcome that disproportionately affects a protected class must be reviewed and, if discriminatory, reversed.

7. Exit Stage — Procedure and AI Boundaries

Exit includes voluntary resignation, redundancy, and termination. AI may assist with administrative tasks but must not make exit decisions.

What AI may do: Generate exit paperwork and offboarding checklists; summarise the departing employee's project history for knowledge-transfer purposes; conduct sentiment analysis on exit survey responses at an aggregate, anonymised level for HR reporting.

What AI may not do: Recommend whether to accept a resignation with or without a notice waiver based on employee risk scoring; determine redundancy selection where any protected characteristic could be a proximate driver of the AI's output; generate a reference or record of employment that has not been reviewed and approved by a human.

Redundancy selection: If Cedar Line Creative uses any AI tool to assist with redundancy selection — for example, by ranking roles or employees by a performance or cost metric — the selection criteria must be defined by humans in advance, must be objectively justifiable, and must be tested for adverse impact against protected characteristics before the selection pool is applied. The AI output is advisory; the selection decision is made and documented by a human.

Post-exit data: AI tools must not be used to monitor or profile former employees using data collected during employment. Following exit, personal data relating to a former employee is retained only for the periods set out in the Records section below, and access is restricted to HR.

8. Candidate Notice Template (AI Use in Hiring)

The following notice must be provided to every candidate at the point of application, before any AI screening tool processes their application. It must be presented as a standalone notice, not buried in general terms.

NOTICE OF AUTOMATED DECISION-MAKING IN HIRING — CEDAR LINE CREATIVE

Date: [Insert date of application]

Cedar Line Creative uses an automated AI screening tool as part of its hiring process. This notice explains how that tool is used and your rights in connection with it.

What the tool does: The AI screening tool reviews your application against the criteria for the role you have applied for. It produces a scored or ranked output that is used by our HR team to compile a shortlist. A human member of our HR team reviews all AI outputs before any decision is made about your application.

What the tool does not do: The tool does not make a final decision about your application. No offer or rejection will be issued solely on the basis of an automated output.

Data used: The tool processes the information you submit in your application, including your CV, cover letter, and answers to application questions. It does not process sensitive personal data (such as health information or biometric data) for scoring purposes. [If biometric data is collected — see BIPA template below — a separate notice will be provided.]

Your rights: You have the right to request a human review of any decision made about your application. To exercise this right, contact [HR contact name and email] within 14 days of receiving notification of the outcome of your application. You will receive a response within 10 business days.

Applicable law: This notice is provided in compliance with the requirements of Texas HB 149 (the Responsible Artificial Intelligence Governance Act), in preparation for Colorado SB 26-189 (effective 1 January 2027), and in accordance with EEOC guidance on AI in employment. California applicants are notified that AI-generated content used in this process may carry provenance disclosures under California SB 942.

Questions: Contact [HR lead name], [title], [email], [phone].

---

9. Employee Notice Template (AI Use in Performance and Management) and Illinois BIPA Consent

The following notice is provided to all employees at onboarding and re-issued whenever a new AI tool is deployed in a performance, discipline, or exit context.

NOTICE OF AI USE IN EMPLOYMENT PROCESSES — CEDAR LINE CREATIVE

Date: [Insert date]

Cedar Line Creative uses AI-assisted tools in some aspects of employment management, including performance data aggregation and, where applicable, scheduling and communications tools. This notice explains how AI is used and your rights.

How AI is used: AI tools may analyse project output data, client feedback scores, and productivity metrics to present summaries to your manager as one input into performance discussions. AI tools do not generate your performance rating; that is done by your manager.

Your rights: You may request a human review of any AI-informed outcome that affects your employment. Submit requests in writing to [HR contact]. You will receive a written response within 10 business days.

Data: AI tools process only data lawfully collected in the course of your employment. Your personal data is not sold to third parties.

ILLINOIS BIPA BIOMETRIC DATA CONSENT — CEDAR LINE CREATIVE

[For use when any recording or automated analysis of Illinois-based candidates or employees captures biometric identifiers, including voice or facial data.]

I, [Full name], acknowledge that Cedar Line Creative has informed me of the following:

1. Cedar Line Creative is collecting [describe biometric identifier, e.g., voiceprint from interview recording] for the purpose of [state purpose, e.g., maintaining an accurate record of the interview].
2. This biometric data will be retained for [state period, not to exceed the employment relationship plus the applicable retention period, or three years from collection, whichever is sooner].
3. Cedar Line Creative will not sell, lease, trade, or otherwise profit from my biometric data.
4. Cedar Line Creative will not disclose my biometric data except as required by law or with my written consent.
5. I have the right to submit a written request for deletion of my biometric data.

By signing below, I provide my written release for the collection described above.

Signature: ______________________ Date: ______________________

Printed name: ______________________ Role/Stage: ______________________

Retained signed copies are stored in the HR records system, accessible only to HR, for the duration set out in the Records section below.

10. Accommodation and Alternative-Process Route

Cedar Line Creative is required under the ADA, applicable state equivalents, and EEOC AI guidance to provide reasonable accommodation to candidates and employees who are unable to participate effectively in an AI-assisted process due to a disability.

Procedure: Any candidate or employee who believes that the AI screening tool, AI-assisted interview process, or AI performance tool creates a barrier linked to disability must contact [HR lead name] at [email] before or as soon as practicable after encountering the barrier. Cedar Line Creative will engage in the interactive process required under the ADA within five business days of receiving the request.

Alternative process: Where accommodation is granted, Cedar Line Creative will substitute a fully human-led alternative at the relevant stage. For screening, this means a qualified HR professional reviews the application directly against the role criteria without AI assistance. For interviews, no automated transcription or analysis tool is used without the candidate's informed consent. For performance management, the employee's manager conducts the assessment using documented criteria without AI data aggregation input, unless the employee consents to its use.

No disadvantage: A candidate or employee who requests and receives an alternative process will not be treated less favourably than one who participates in the AI-assisted process. Cedar Line Creative will not use the exercise of an accommodation right as a factor in any employment decision.

Documentation: All accommodation requests, the interactive process dialogue, and the agreed alternative are documented and retained in the relevant personnel file. The outcome of the alternative process is recorded in the same format as the AI-assisted process output it replaced, to ensure comparability.

Colorado SB 26-189: From 1 January 2027, the meaningful human review mechanism required by SB 26-189 for consequential automated decisions is distinct from, and does not replace, the ADA accommodation process. Both routes remain available.

11. Bias Testing — Detailed Protocol

Bias testing is mandatory at three points: before any new deployment or configuration change, within 30 days of each hiring cycle's close, and annually each January.

Pre-deployment testing: HR obtains from the AI screening tool vendor a current bias and fairness assessment covering, at minimum: adverse impact analysis by sex, race/ethnicity, age group, and disability status; methodology description; and the dataset on which the tool was validated. Where the vendor cannot produce this documentation, the tool is not used. HR then runs the tool on a synthetic or historical application dataset drawn from Cedar Line Creative's own sector, disaggregates outputs by available protected-class proxies, and applies the four-fifths rule. Results are documented on the Pre-Deployment Bias Assessment Form (retained for five years).

Post-deployment testing: Within 30 days of closing each hiring exercise in which the AI screening tool was used, HR produces a post-hire report comparing: (a) demographic distribution of applicants; (b) demographic distribution of AI-passed candidates; and (c) demographic distribution of candidates selected by humans for interview. If the selection rate for any group falls below 80% of the highest-selecting group's rate at any stage, HR suspends AI screening for that role type, investigates, and documents remediation before resuming.

Annual audit: Each January, the HR lead commissions an aggregate bias review covering all hiring in the previous calendar year. The review is conducted by a qualified HR professional who was not involved in the day-to-day screening decisions. Findings are presented to the most senior HR decision-maker and, where material adverse impact is found, to the Chief Executive. The annual audit report is retained for seven years.

NIST AI RMF alignment: Cedar Line Creative's bias testing protocol maps to the NIST AI RMF 1.0 GOVERN, MAP, MEASURE, and MANAGE functions. The HR lead maintains a risk register entry for the AI screening tool that is updated after each bias test cycle.

12. Records Retained Per Stage and Retention Periods

Cedar Line Creative retains the following records in connection with AI-assisted HR processes. All records are stored in a secure HR system with access restricted to authorised HR personnel. Personal data is processed in accordance with applicable privacy law; where a record contains personal data of a California, Texas, or Illinois resident, the relevant state privacy rights apply.

Sourcing stage: Job description drafts and final versions (including AI-generated drafts and human-reviewed final); sourcing channel selection rationale; adverse impact check log. Retention: Two years from close of role.

Screening stage: Pre-deployment bias assessment form and vendor bias documentation; AI screening tool configuration at the time of use; all candidate application data processed by the tool; AI output (scores or rankings) for each candidate; post-hire adverse impact report. Retention: Four years from close of role, extended to seven years if any discrimination complaint or legal proceeding is initiated or reasonably anticipated.

Interview stage: Structured interview question bank used; human interviewer scoring sheets (signed); AI-generated transcripts or notes (identified as AI-assisted); BIPA consent forms for Illinois candidates (retained for the longer of five years from collection or three years from last interaction, as required by BIPA). Retention: Four years from close of role.

Selection stage: Human sign-off documentation for each offer and rejection; record of any departure from AI screening recommendation with stated reason; candidate notices issued. Retention: Four years from close of role.

Performance stage: Annual bias review of aggregate performance data; AI data summaries used in performance discussions (retained as supporting documentation, not the operative rating record); formal performance ratings (human-signed). Retention: Duration of employment plus five years.

Discipline stage: Complete disciplinary case file including all AI-assisted materials identified as such; human decision records. Retention: Duration of employment plus seven years, or longer if required by applicable law or ongoing proceedings.

Exit stage: Exit paperwork; redundancy selection records including AI-assisted analysis and human decision sign-off; aggregate exit survey data (anonymised AI sentiment analysis). Retention: Seven years from exit date.

Accommodation records: All accommodation requests, interactive process records, and alternative-process outcomes. Retention: Duration of employment plus seven years.

Bias testing records (all stages): Pre-deployment forms, post-hire reports, annual audit reports. Retention: Seven years from the date of the relevant hire or audit.

Illinois BIPA records: Signed consent forms and any deletion request records. Retention: As set out above under Interview stage; deletion completed within 30 days of the end of the specified retention period.

13. Governance, Accountability, and Review

Overall accountability for this procedure rests with the most senior HR decision-maker at Cedar Line Creative, currently designated as [title to be completed by client]. Day-to-day responsibility for each hiring exercise sits with the HR professional assigned to that role.

The AI screening tool vendor contract must include: (a) a requirement to notify Cedar Line Creative within 72 hours of any material change to the tool's model, training data, or scoring methodology; (b) access to bias and fairness documentation on demand; (c) a data processing agreement compliant with applicable US state privacy law; and (d) a representation that the tool does not produce outputs based on prohibited characteristics.

This procedure is reviewed annually and updated within 30 days of: any material change to the AI screening tool; any new or amended legal obligation in Colorado, Texas, Illinois, or California affecting AI in employment; any regulatory action or enforcement guidance from the EEOC or FTC that materially affects Cedar Line Creative's obligations; or any bias test result indicating adverse impact that has not been remediated within the review cycle.

Colorado SB 26-189 readiness milestone: By 31 October 2026, Cedar Line Creative must have completed its SB 26-189 impact assessment for the AI screening tool, documented its risk management policy, and tested the candidate human-review mechanism, so that full compliance is operational by 1 January 2027.

FTC Section 5: Cedar Line Creative must not make any public or candidate-facing representation about the fairness or objectivity of its AI-assisted hiring process that is not substantiated by current bias test results. All employer branding copy referencing AI in hiring must be reviewed by HR and approved against the most recent annual bias audit before publication.

AI Impact Assessment Template

Complete one before deploying any new AI system that touches personal data or affects people. Keep the completed copies.

This AI Impact Assessment Template is a mandatory pre-deployment document for Cedar Line Creative. One completed copy must be retained for every AI system that processes personal data or produces outputs that affect individuals. Retained copies may be disclosed to regulators, insurers, or clients on request. The template reflects obligations arising under Colorado SB 26-189 (deployer duties effective 1 January 2027), California SB 942 / AB 853 (AI Transparency Act, in force 2 August 2026), Texas HB 149 (Responsible Artificial Intelligence Governance Act, in force 1 January 2026), Illinois BIPA, the Virginia CDPA, NIST AI RMF 1.0, FTC Section 5, and EEOC guidance on AI in employment. Complete every section before go-live. Do not delete any section; if a section is inapplicable, state that fact and the reason.

AI IMPACT ASSESSMENT
Cedar Line Creative

Version: [VERSION NUMBER]
Date of Assessment: [DATE]
Assessment Lead: [FULL NAME AND TITLE]
Reviewed By: [FULL NAME AND TITLE]
Next Scheduled Review: [DATE — NO LATER THAN 12 MONTHS FROM SIGN-OFF]

SECTION 1 — SYSTEM DESCRIPTION, PURPOSE AND PEOPLE AFFECTED

1.1 System name and vendor
System name: [e.g., ChatGPT / OpenAI; Claude / Anthropic; Midjourney; DALL-E; AI-powered CRM tool; AI HR screening tool — select one per assessment]
Vendor: [VENDOR LEGAL NAME]
Version or model: [MODEL VERSION]
API or hosted interface: [SPECIFY]

1.2 Intended purpose within Cedar Line Creative
Describe precisely what the system will do. Examples: generate marketing copy for publication under a client brand; power a chat assistant on a client-facing website; screen job applicants against defined criteria; analyse CRM data to score or segment prospects.
Purpose: [DESCRIPTION]

1.3 Deployment context
Where will it operate? (e.g., internal workflow, client website, ATS platform, CRM instance)
Deployment context: [DESCRIPTION]
Client name (if deployed on behalf of a client): [CLIENT NAME OR 'INTERNAL']
Geographic reach — confirm all states where staff, data subjects, or end users are located:
[ ] Colorado [ ] California [ ] Texas [ ] Illinois [ ] Virginia [ ] Other: [SPECIFY]

1.4 People affected
List every category of individual whose data the system processes or who is subject to an output or decision produced by the system.

Category of person: [e.g., Job applicant]
Approximate number: [NUMBER]
Vulnerable or protected characteristics present in this group: [YES / NO — if yes, specify]

Category of person: [e.g., Website visitor / chat user]
Approximate number: [NUMBER]
Vulnerable or protected characteristics present: [YES / NO — if yes, specify]

Category of person: [e.g., Client's customer receiving AI-generated content]
Approximate number: [NUMBER]
Vulnerable or protected characteristics present: [YES / NO — if yes, specify]

Add further rows as needed.

SECTION 2 — NECESSITY AND PROPORTIONALITY

2.1 Business need
State the specific operational need this system addresses.
Need: [DESCRIPTION]

2.2 Alternatives considered
Identify at least two alternatives, including a non-AI or lower-automation option, and explain why each was rejected.

Alternative 1: [DESCRIPTION]
Reason rejected: [REASON]

Alternative 2: [DESCRIPTION]
Reason rejected: [REASON]

2.3 Proportionality finding
Confirm that the system's scope (data collected, population affected, degree of automation) is the minimum necessary to achieve the stated purpose. If the system affects employment decisions, confirm alignment with EEOC guidance requiring that AI tools be the least discriminatory available means.
Finding: [DESCRIPTION]
Agreed by (name and title): [NAME]

SECTION 3 — DATA USED, LAWFUL BASIS AND RETENTION

3.1 Personal data inputs
List every category of personal data the system ingests, generates, or processes as part of its operation.

Data category: [e.g., Applicant CV, name, contact details]
Source: [e.g., Applicant submission via ATS]
Sensitive / biometric under Illinois BIPA: [YES / NO]

Data category: [e.g., Website visitor chat transcript]
Source: [e.g., Client website session]
Sensitive / biometric under Illinois BIPA: [YES / NO]

Data category: [e.g., CRM contact record]
Source: [e.g., Internal CRM import]
Sensitive / biometric under Illinois BIPA: [YES / NO]

If any biometric identifiers or information are collected (voice, face geometry, etc.), Illinois BIPA requires a written policy, a public retention schedule, informed written consent before collection, and prohibits sale or profit from biometric data. Confirm compliance or confirm BIPA does not apply to this system: [CONFIRMATION]

3.2 Lawful basis
State the lawful basis for processing under each applicable regime.

Virginia CDPA (applies where Virginia residents' data is processed): Consent / Contract / Legal obligation / Legitimate interest — [SELECT AND EXPLAIN]
California (where California residents are involved, SB 942 applies to generative AI output disclosure; CPRA grounds apply to personal data): [BASIS]
Texas HB 149 (data processed by AI systems in Texas or affecting Texas residents — confirm no use in prohibited categories): [BASIS]
Colorado SB 26-189 (not yet in force; effective 1 January 2027 — document basis now so the system is ready): [ANTICIPATED BASIS]
FTC Section 5 (no separate lawful basis required; confirm no deceptive claims are made about AI capabilities or outputs): [CONFIRMATION]

EEOC AI guidance (employment use only — confirm the tool has been validated for adverse impact and disparate effect on protected classes): [CONFIRMATION OR 'NOT APPLICABLE — NOT AN EMPLOYMENT SYSTEM']

3.3 Retention
How long will personal data processed by or through this system be retained?
Retention period: [PERIOD]
Deletion or anonymisation method: [METHOD]
Who is responsible for deletion: [ROLE]

SECTION 4 — RISKS TO INDIVIDUALS, SCORED, WITH MITIGATIONS

Score likelihood and impact on a 1–5 scale (1 = very low, 5 = very high). Multiply to produce a risk score. Residual risk is the score after mitigations are applied.

4.1 Risk: Discriminatory or biased output in employment screening
Applicable law: EEOC AI in Employment guidance; Colorado SB 26-189 (from 1 January 2027) — deployers of automated decision systems used in consequential decisions (employment is listed) must conduct impact assessments and notify affected individuals.
Likelihood (pre-mitigation): [1–5]
Impact (pre-mitigation): [1–5]
Pre-mitigation score: [SCORE]
Mitigation: Conduct pre-deployment disparate impact testing across race, sex, age, and other protected classes. Engage vendor for bias audit reports. Ensure a human recruiter reviews every AI-generated screening output before any decision is communicated to the applicant. Document the validation methodology.
Likelihood (post-mitigation): [1–5]
Impact (post-mitigation): [1–5]
Residual score: [SCORE]
Owner: [ROLE]

4.2 Risk: Failure to disclose AI-generated content to end users
Applicable law: California SB 942 (in force 2 August 2026) — large generative AI providers must embed latent provenance disclosures; Cedar Line Creative as a deployer publishing AI-generated copy or imagery under client brands must not strip or obscure those disclosures. Texas HB 149 (in force 1 January 2026) — disclosure obligations apply to AI-generated content in defined contexts. FTC Section 5 — the FTC's proposed policy statement of 1 July 2026 warns that suppressing or altering AI output disclosures, even to comply with other laws, may constitute consumer deception.
Likelihood (pre-mitigation): [1–5]
Impact (pre-mitigation): [1–5]
Pre-mitigation score: [SCORE]
Mitigation: Add disclosure wording to all AI-generated copy and imagery published under client brands. Preserve, do not strip, latent provenance metadata embedded by the generative AI provider. Review client contracts to allocate disclosure obligations. Maintain a disclosure log per campaign.
Likelihood (post-mitigation): [1–5]
Impact (post-mitigation): [1–5]
Residual score: [SCORE]
Owner: [ROLE]

4.3 Risk: Intellectual property infringement in AI-generated outputs
Applicable law: FTC Section 5 (misrepresentation of ownership); general US copyright law as shaped by ongoing litigation — as at August 2026 over 80 AI copyright lawsuits are active in the US.
Likelihood (pre-mitigation): [1–5]
Impact (pre-mitigation): [1–5]
Pre-mitigation score: [SCORE]
Mitigation: Obtain vendor IP indemnity terms in writing. Do not represent AI outputs as wholly original human-authored works in client contracts or pitches. Conduct originality checks on image outputs before

publication. Maintain records of prompts and tools used for each deliverable.
Likelihood (post-mitigation): [1–5]
Impact (post-mitigation): [1–5]
Residual score: [SCORE]
Owner: [ROLE]

4.4 Risk: Unlawful processing of biometric data
Applicable law: Illinois BIPA — applies where the system processes biometric identifiers or information of Illinois residents or employees. BIPA provides a private right of action with statutory damages of $1,000–$5,000 per violation.
Does this system collect or process any biometric data? [YES / NO]
If NO: This risk does not apply to this deployment. Basis for exclusion: [STATE REASON].
If YES:
Likelihood (pre-mitigation): [1–5]
Impact (pre-mitigation): [1–5]
Pre-mitigation score: [SCORE]
Mitigation: Publish a written biometric retention policy before collection begins. Obtain separate, informed written consent from each individual before collection. Prohibit any sale or commercial disclosure of biometric data. Set a destruction schedule no later than 3 years or the end of the relevant relationship, whichever comes first.
Residual score: [SCORE]
Owner: [ROLE]

4.5 Risk: Chatbot interactions — deception or harmful advice
Applicable law: FTC Section 5 — the AI enforcement unit established January 2026 is targeting undisclosed AI-generated content and AI-washing; the chatbot must not impersonate a human without disclosure. Virginia CDPA — where Virginia residents interact with the chatbot, personal data processed in the session is subject to data minimisation and purpose limitation.
Likelihood (pre-mitigation): [1–5]
Impact (pre-mitigation): [1–5]
Pre-mitigation score: [SCORE]
Mitigation: Display a clear disclosure that the assistant is AI-powered at the start of every chat session. Configure the system to refuse to claim it is human when sincerely asked. Limit data collected in sessions to what is necessary for the query. Delete session transcripts per the retention schedule at 3.3.
Likelihood (post-mitigation): [1–5]
Impact (post-mitigation): [1–5]
Residual score: [SCORE]
Owner: [ROLE]

4.6 Risk: Consequential automated decisions without required notice (Colorado SB 26-189)
Applicable law: Colorado SB 26-189 — effective 1 January 2027, deployers using automated decision technology in consequential decisions (employment, credit, housing, education, healthcare, insurance) must notify affected individuals and provide a mechanism to opt out or appeal.
This risk is a forward-looking risk: the law is not yet in force. Cedar Line Creative's AI hiring screening tool will fall within scope from 1 January 2027.
Likelihood (pre-mitigation): [1–5]
Impact (pre-mitigation): [1–5]
Pre-mitigation score: [SCORE]
Mitigation: Design notice and appeal workflows into the HR screening system before 1 January 2027. Document the decision logic for each consequential use case. Assign a responsible person to monitor Colorado Attorney General guidance issued between now and commencement.
Residual score: [SCORE]
Owner: [ROLE]

SECTION 5 — HUMAN OVERSIGHT AND ROUTE TO CHALLENGE

5.1 Degree of automation
Is any output of this system used to make or materially influence a decision about an individual without human review before the decision is communicated?
[ ] Fully automated — no human in the loop
[ ] Human reviews output before decision is communicated
[ ] Human can override but default is to accept AI output
[ ] AI is advisory only; human always makes the decision

Selected option: [SELECT]
If fully automated: state the specific justification and the legal basis permitting fully automated decisions in this context: [JUSTIFICATION]

5.2 Human oversight roles
Role responsible for reviewing AI output before it affects an individual: [JOB TITLE]
Authority to override the AI output: [YES / NO]
Escalation path if the reviewer is uncertain: [DESCRIBE]

5.3 Route to challenge for affected individuals
An individual who believes an AI-generated decision or output has adversely affected them may:
(a) Contact [DESIGNATED CONTACT NAME / ROLE] at [EMAIL ADDRESS] to request a human review of the specific output or decision.
(b) Receive a written response within [NUMBER] business days explaining the basis of the output and, where relevant, the data used.
(c) Request correction or deletion of personal data used in producing the output.
(d) Where the system is used in employment decisions: applicants in Colorado (from 1 January 2027) have a statutory right to opt out of solely automated consequential decisions and to appeal an adverse outcome. This right will be communicated in Cedar Line Creative's candidate-facing hiring notices before that date.

5.4 Audit trail
Confirm that the system logs every decision or output that affects an individual, including the data inputs used, the model version, and the timestamp.
Audit trail in place: [YES / NO]
Location of logs: [SYSTEM / LOCATION]
Log retention period: [PERIOD]

SECTION 6 — RESIDUAL RISK, SIGN-OFF AND REGULATOR CONSULTATION

6.1 Overall residual risk rating
After all mitigations in Section 4 are applied, record the highest residual score and assign an overall rating.
Highest residual score: [SCORE]
Overall rating: [ ] Low (1–8) [ ] Medium (9–14) [ ] High (15–25)

6.2 Deployment decision
[ ] Approved for deployment — residual risk is acceptable and all mitigations are in place or scheduled with firm deadlines.
[ ] Conditionally approved — deployment may proceed only once the following mitigations are confirmed complete: [LIST CONDITIONS]
[ ] Not approved — residual risk is unacceptable. System must not be deployed until a further assessment is completed.

6.3 Sign-off
Assessment Lead signature: ________________________ Date: [DATE]
Senior Responsible Owner (must be director level or above): ________________________ Date: [DATE]
Legal / Compliance sign-off: ________________________ Date: [DATE]

6.4 Regulator consultation requirement
Review each of the following and confirm the position.

Colorado SB 26-189: Formal regulator consultation is not required before deployment; the obligation is to conduct this assessment and retain it. The Attorney General may request it on investigation. Confirm assessment is retained and accessible: [YES / NO]. Note: obligations commence 1 January 2027; this assessment constitutes preparatory documentation.

California SB 942: No prior regulator approval required; obligations are self-executing from 2 August 2026. Confirm disclosure and provenance obligations are met before any generative AI content is published: [YES / NO].

Texas HB 149: No prior approval required; enforcement is by the Attorney General. Confirm no prohibited AI uses are present and any required disclosures are in place: [YES / NO].

Illinois BIPA: No prior approval required; private right of action applies. If biometric data is processed, confirm written policy is published and consents obtained before collection: [YES / NO / NOT APPLICABLE].

Virginia CDPA: No prior approval required. Confirm data protection assessment for high-risk processing is complete and on file: [YES / NO].

FTC Section 5: No prior approval required. Confirm no deceptive claims about AI capabilities, outputs, or authorship appear in any client-facing or public-facing materials: [YES / NO].

EEOC AI guidance: No prior approval required. For employment systems, confirm bias audit documentation is held and available: [YES / NO / NOT APPLICABLE].

If any answer above is NO, deployment must not proceed until the gap is remedied and this field is updated.

SECTION 7 — REVIEW DATE AND VERSION CONTROL

7.1 Mandatory review triggers
This assessment must be reviewed, and a new version completed, upon any of the following:
(a) A material change to the system, model version, or data inputs.
(b) A change in the purpose or population affected.
(c) A new legal obligation coming into force that affects this system — in particular, Colorado SB 26-189 deployer duties from 1 January 2027.
(d) A complaint, regulatory enquiry, or adverse incident involving this system.
(e) In any event, no later than 12 months from the date of the current sign-off.

7.2 Scheduled review date
Next review due: [DATE]
Review owner (role): [ROLE]

7.3 Version history
Version | Date | Author | Summary of changes
[VERSION] | [DATE] | [NAME] | [DESCRIPTION]

END OF TEMPLATE
Cedar Line Creative — Confidential
Retain completed copies securely. Do not destroy without legal sign-off.

AI Incident Response Procedure

What happens in the first hour after an AI-related incident, and who does it. The regulator clock starts before anyone has decided who is in charge.

1. Purpose and Scope

This AI Incident Response Procedure governs Cedar Line Creative's response to any incident involving its artificial intelligence tools and systems. It applies to all staff, contractors, and third-party operators acting on Cedar Line Creative's behalf across its Denver headquarters and all client and staff locations in California, Texas, and Illinois.

The procedure covers every AI tool currently in use: ChatGPT and Claude for drafting marketing and communications content; Midjourney and DALL-E for generating imagery published under client brands; AI-powered CRM and sales tools; and the AI screening tool used in Cedar Line Creative's own hiring process. It also covers the AI chat assistant deployed on client websites, which interacts directly with members of the public.

The regulator clock begins at the moment an incident occurs or is discovered — not when a decision has been made about its severity or ownership. This procedure is designed to ensure that the first hour of response is structured, accountable, and legally defensible regardless of the time of day, the seniority of the person who discovers the incident, or whether the responsible lead is immediately available.

2. What Counts as an AI Incident

An AI incident is any event, near-miss, or reasonably suspected occurrence in which an AI tool used by Cedar Line Creative has caused, or has the realistic potential to cause, harm to individuals, clients, or the business, or has breached a legal or contractual obligation.

For Cedar Line Creative specifically, incidents include but are not limited to the following. First, the AI screening tool used in hiring produces a recommendation that appears to correlate with a protected characteristic — for example, systematically lower scores for candidates whose names or educational histories suggest a particular race, gender, or national origin — engaging EEOC AI in Employment guidance and, from 1 January 2027, Colorado SB 26-189 duties as a deployer of automated decision-making technology used in a consequential employment decision. Second, AI-generated copy or imagery published under a client brand contains content that is defamatory, sexually explicit, discriminatory, or that infringes a third party's copyright or trade mark. Third, the AI chat assistant deployed on a client website discloses personal data belonging to one user to another user, produces harmful advice, or impersonates a human operator without disclosure, potentially engaging California SB 942 and FTC Section 5. Fourth, an AI image generation tool produces output that reproduces a living person's likeness in a false or harmful context, or generates intimate imagery without consent, which may engage the federal Take it Down Act in force since May 2026. Fifth, personal data processed through ChatGPT, Claude, or the CRM AI tool is exposed to an unintended recipient, retained beyond agreed terms, or used to train an external model without a lawful basis. Sixth, a vendor notifies Cedar Line Creative of a security breach, model failure, or regulatory action affecting a tool Cedar Line Creative deploys. Seventh, a client, regulator, or member of the public makes a credible allegation that AI-generated content attributed to a client brand is misleading about its AI origin, engaging California SB 942 which has been in force since 2 August 2026.

A near-miss — where the harm did not materialise but the conditions for it existed — is recorded and reviewed under the same procedure, even if a lower severity level is assigned.

3. Severity Levels and the Response Each Triggers

Cedar Line Creative uses three severity levels, assessed at the point of discovery and reviewed upward or downward within the first thirty minutes as facts develop.

Severity 1 — Critical. An incident that has caused or is causing actual harm to identifiable individuals, involves a breach of personal data affecting more than a trivial number of data subjects, involves published content that is defamatory, explicitly discriminatory, or sexually explicit, involves a compromised hiring decision already acted upon, or involves an active data exposure that is ongoing. Severity 1 triggers immediate activation of the full Incident Response Team, mandatory suspension of the implicated tool pending review, evidence preservation starting within fifteen minutes of classification, and assessment of statutory notification obligations within the first hour. The Incident Response Lead (defined in Section 5) must be reached within ten minutes; if unavailable, their named deputy assumes authority without waiting for confirmation.

Severity 2 — Significant. An incident with credible potential for harm that has not yet materialised, a hiring screening anomaly identified before the employment decision is finalised, AI-generated content queued for publication that contains problematic material caught before release, a vendor alert about a model vulnerability not yet exploited, or a client complaint about AI-generated output that may constitute a misleading AI origin claim. Severity 2 triggers Incident Response Lead notification within thirty minutes, precautionary suspension of the affected workflow, evidence preservation, and a formal regulatory notification assessment documented in writing within four hours.

Severity 3 — Minor. A near-miss, a content quality failure with no legal exposure, a single low-stakes client complaint about AI output that can be resolved operationally, or a process deviation that did not affect any individual's rights or interests. Severity 3 triggers logging in the AI Incident Register within twenty-four hours, standard operational review within five business days, and no immediate suspension requirement unless the Incident Response Lead determines otherwise. Severity 3 incidents are reviewed in aggregate at the monthly post-incident review.

4. First-Hour Actions: Contain, Preserve, and Stop

The first hour of response follows a fixed sequence regardless of who discovers the incident. The sequence is: Detect and Classify, Contain, Preserve Evidence, Stop or Restrict the Tool, Notify Internal Leads.

Detect and Classify (minutes 0–10). The person who discovers or suspects an incident records the time, the tool or system involved, a factual description of what they observed, and any identifiable individuals affected. They immediately classify the incident as Severity 1, 2, or 3 using the criteria in Section 3. They do not attempt to resolve the incident or contact the vendor before completing this step. They notify the Incident Response Lead by the fastest available means — phone call first, then direct message, then email.

Contain (minutes 10–20). For a Severity 1 or 2 incident, the discovering person or the Incident Response Lead takes the minimum action necessary to prevent the harm from spreading. This means: revoking or suspending the API key or user session for the implicated tool if they have access; removing AI-generated content from publication or from the client delivery queue; suspending the AI chat assistant on any affected client website by contacting the relevant client website administrator or using the platform's administrative controls; and flagging any hiring decisions pending that relied on the AI screening tool.

Preserve Evidence (minutes 15–30). No logs, outputs, prompts, model responses, or configuration settings are deleted or altered. Screenshots and system exports of the AI tool's interface, the output that caused the incident, and any error messages are taken immediately and stored in Cedar Line Creative's designated secure evidence folder. Relevant email threads, Slack or Teams messages, and client communications are preserved in place. If the incident involves personal data, the data concerned is identified and a record of what it contains and where it resides is created. Evidence preservation is not contingent on a decision about whether to notify a regulator.

Stop or Restrict the Tool (minutes 20–40). The implicated AI tool is suspended from all further use within Cedar Line Creative's workflows for the duration of the investigation. If the tool is embedded in a client-facing product (such as the AI chat assistant), the client is notified that the tool has been taken offline as a precautionary measure, without disclosing unnecessary detail about the incident. Where suspension is technically outside Cedar Line Creative's direct control — for example, where the tool is administered by a third-party vendor — the vendor escalation route in Section 8 is activated immediately.

Notify Internal Leads (minutes 30–60). The Incident Response Lead briefs the Legal and Compliance Lead and the relevant business unit head. The full Incident Response Team listed in Section 5 is convened for a Severity 1 incident before the end of the first hour. For Severity 2, the team is convened within four hours. The regulatory notification assessment in Section 6 begins no later than the point at which the Incident Response Team is convened.

5. Named Roles and Their Deputies

Cedar Line Creative's AI Incident Response Team comprises four named roles. Each role has a designated deputy who assumes full authority if the primary holder is unreachable within ten minutes of the first contact attempt.

The Incident Response Lead holds overall command of the response, makes final decisions on tool suspension, external notifications, and regulatory filings, and is the single point of accountability to the Board. The Incident Response Lead's deputy assumes this authority automatically and without escalation if the Lead is unreachable.

The Legal and Compliance Lead is responsible for the regulatory notification assessment (Section 6), drafting any required notifications to regulators or affected individuals (Section 7), managing vendor legal communications (Section 8), and maintaining the AI Incident Register (Section 9). This role must hold, or have immediate access to, the text of every applicable statute and notification form.

The Technology and Data Lead is responsible for the technical containment actions in Section 4, preserving system evidence, liaising with vendor technical teams, and confirming when a tool has been fully suspended or restricted. This role maintains the list of administrative credentials and suspension procedures for every AI tool in Cedar Line Creative's stack.

The Client Services Lead is responsible for communicating with affected clients in a manner approved by the Incident Response Lead, managing reputational risk with client contacts, and coordinating the removal of AI-generated content from client channels where Cedar Line Creative does not have direct access.

The specific names and direct contact details for each role and their deputies are held in Cedar Line Creative's confidential Incident Contact Sheet, which is stored separately from this document and reviewed quarterly. This procedure refers to roles rather than names so that it remains valid between personnel reviews.

6. Regulator Notification: Decision and Statutory Deadlines

The regulatory notification assessment begins within the first hour for Severity 1 incidents and within four hours for Severity 2 incidents. The Legal and Compliance Lead leads this assessment. It is a legal obligation assessment, not a discretionary business decision.

Illinois Biometric Information Privacy Act (BIPA). If the AI screening tool has collected, stored, or transmitted biometric identifiers or biometric information — including facial geometry derived from applicant photographs or video — without a compliant written release, Illinois BIPA is engaged. BIPA does not impose a specific regulator notification deadline but creates a private right of action that can be initiated immediately by affected individuals. The Legal and Compliance Lead must assess within the first hour whether biometric data was processed and whether a written policy and release were in place. If neither existed, legal counsel must be instructed before the end of the first day.

Colorado SB 26-189. This law is not yet in force; its effective date is 1 January 2027. However, Cedar Line Creative's AI screening tool is used in hiring decisions and will constitute a consequential decision under SB 26-189 once it commences. Cedar Line Creative must begin building SB 26-189-compliant processes now so that they are operational before that date. No notification obligation under SB 26-189 is triggered by incidents occurring before 1 January 2027.

California SB 942 (AI Transparency Act). In force since 2 August 2026, SB 942 imposes obligations on large providers of generative AI systems, not on deployers such as Cedar Line Creative directly. However, Cedar Line Creative publishes AI-generated copy and imagery under client brands and operates an AI chat assistant on client websites. Where Cedar Line Creative has represented to clients or to the public that AI content carries a disclosure (or where the absence of disclosure is itself the incident), the Legal and Compliance Lead must assess whether Cedar Line Creative's conduct exposes it or its clients to SB 942-based claims and must advise the Incident Response Lead accordingly. California does not specify a regulator notification deadline for SB 942 incidents; the enforcement mechanism is through the California Attorney General. The assessment should be completed within four hours.

Texas HB 149 (Responsible Artificial Intelligence Governance Act). In force since 1 January 2026, Texas HB 149 applies to AI developers and deployers operating in Texas and includes provisions on prohibited uses and government AI disclosure. Cedar Line Creative has clients and staff in Texas. The Legal and Compliance Lead must assess whether the incident involves a prohibited use under HB 149 and, if so, must prepare a notification to the Texas Attorney General. HB 149 does not specify an individual notification deadline in the same form as a data breach statute; the position should be confirmed with Texas counsel for the specific incident type.

FTC Section 5. The FTC's enforcement posture on undisclosed AI-generated advertising content is active: the FTC brought its first action targeting such content and issued a proposed policy statement on 1 July 2026 warning that suppressing AI accuracy disclosures can constitute consumer deception. Where an incident involves Cedar Line Creative publishing AI-generated content without disclosure, or where the AI chat assistant has misrepresented itself as human, the Legal and Compliance Lead must assess FTC Section 5 exposure within the first hour. The FTC does not impose a statutory notification deadline; however, voluntary early disclosure is a recognised mitigating factor in FTC enforcement.

EEOC AI in Employment Guidance. Where the incident involves Cedar Line Creative's AI screening tool and there is evidence or a credible allegation of discriminatory impact on a protected class, the EEOC guidance is engaged. The Legal and Compliance Lead must assess whether the affected hiring decision must be unwound and whether the EEOC or a state equivalent must be notified. No statutory notification clock is specified in the EEOC guidance itself; the obligation arises if a charge is filed by an affected individual.

NIST AI RMF 1.0. The NIST AI Risk Management Framework is a voluntary standard, not a legal instrument, and does not impose notification deadlines. It does, however, inform Cedar Line Creative's post-incident review obligations under Section 9 and may be referenced by regulators or insurers assessing the adequacy of Cedar Line Creative's response.

All notification decisions and the reasoning behind them are documented in writing in the AI Incident Register before any notification is sent.

7. Notifying Affected Individuals

The obligation to notify affected individuals is distinct from the obligation to notify regulators and is assessed separately. The Legal and Compliance Lead conducts this assessment in parallel with the regulatory notification assessment.

The threshold for notifying affected individuals is reached where: a person's personal data has been exposed to an unintended recipient or used in a manner inconsistent with the purpose for which it was collected; a consequential decision about an individual — such as a hiring outcome — has been made or materially influenced by an AI tool in a manner that was not disclosed to that individual and that may have caused them harm; or an AI-generated output has been published that contains false, defamatory, or discriminatory content about an identifiable individual.

For individuals located in California, the obligation to disclose AI-generated content origins is reinforced by California SB 942. For individuals who have interacted with the AI chat assistant on a client website and whose personal data may have been exposed, Cedar Line Creative must notify them promptly and in plain language, identifying what data was involved, how the incident occurred, and what Cedar Line Creative is doing to address it.

For job applicants in Illinois who have had biometric data processed through the AI screening tool without a compliant BIPA release, notification to affected applicants is required as part of remediation. The content of the notification must be approved by the Legal and Compliance Lead before it is sent.

For applicants or employees in any state where a consequential AI-assisted employment decision was made without disclosure, Cedar Line Creative will notify the affected individual that AI technology was used in the relevant assessment process, consistent with EEOC guidance and, from 1 January 2027, Colorado SB 26-189 requirements.

All individual notifications are logged in the AI Incident Register, including the date sent, the recipients, and the content. Where a client's customers are affected, individual notifications are co-ordinated with the relevant client through the Client Services Lead, and the client's prior approval is obtained before notification is sent unless the delay in obtaining that approval would itself cause harm to the individual.

8. Vendor Escalation Route

Cedar Line Creative's AI tools are operated by third-party vendors: OpenAI (ChatGPT), Anthropic (Claude), Midjourney, OpenAI (DALL-E), and the vendors of its AI-powered CRM and AI screening tools. Where an incident originates in or is affected by a vendor's system, the vendor escalation route operates in parallel with Cedar Line Creative's internal response and does not replace it.

Within the first hour of a Severity 1 or 2 incident, the Technology and Data Lead contacts the relevant vendor's enterprise or business support channel using the contact details held in Cedar Line Creative's Vendor Escalation Sheet. The escalation message states: the date and time of the incident, the nature of the output or failure observed, the tool version or API endpoint involved, whether personal data of identified individuals was implicated, and whether Cedar Line Creative has suspended use of the tool. The escalation is sent in writing (email or the vendor's support portal) and the Technology and Data Lead retains a copy in the evidence folder.

Vendor responses do not extend Cedar Line Creative's own regulatory or notification deadlines. Cedar Line Creative does not wait for a vendor's root cause analysis before completing its own regulatory notification assessment. If a vendor advises Cedar Line Creative that a breach has occurred on the vendor's side that affects Cedar Line Creative's data or its clients' data, that information is immediately passed to the Legal and Compliance Lead and treated as a new or updated Severity 1 input.

Cedar Line Creative's contracts with AI vendors should include, and where they do not currently include, must be amended to include: data processing agreements specifying the vendor's obligations on breach notification; confirmation of the vendor's approach to training data and model outputs, relevant to intellectual property risk; and a commitment from the vendor to notify Cedar Line Creative of any regulatory action, model change, or security incident affecting Cedar Line Creative's deployment within a defined timeframe. The Legal and Compliance Lead is responsible for confirming whether current vendor contracts meet these requirements and for escalating gaps to the Incident Response Lead.

9. Post-Incident Review and the AI Incident Register

Every incident classified at any severity level is subject to a post-incident review and is recorded in Cedar Line Creative's AI Incident Register. The register is the authoritative record of Cedar Line Creative's AI incident history and is a primary document in any regulatory investigation, insurance claim, or client dispute.

For Severity 1 incidents, the post-incident review is conducted within five business days of the incident being resolved. For Severity 2 incidents, the review is conducted within ten business days. For Severity 3 incidents, the review is conducted in the monthly aggregate review. The review is led by the Incident Response Lead and must include the Legal and Compliance Lead, the Technology and Data Lead, and the relevant business unit head.

The post-incident review addresses: a factual chronology of the incident from detection to resolution; an assessment of whether the first-hour response followed this procedure and, if not, what deviated and why; identification of the root cause, including whether the cause lay in Cedar Line Creative's configuration, prompting, or use of the tool, or in the vendor's underlying model; an assessment of whether the regulatory notification decision was correct in light of facts known after the event; the effectiveness of individual and regulator notifications where sent; recommended changes to Cedar Line Creative's AI governance documents, vendor contracts, tool configuration, or staff training; and a named owner and deadline for each recommended action.

The AI Incident Register entry for each incident records: a unique incident reference number; the date, time, and location of discovery; the AI tool or tools involved; the severity classification and any reclassifications; a factual description of the incident and its impact; the names of individuals affected and the nature of their personal data involved, held separately in a secure annex consistent with data minimisation principles; the regulatory notification assessment outcome and, if notification was made, the date, recipient, and content; the individual notification outcome; the vendor escalation outcome; the post-incident review date and summary; and the status of each remedial action.

The AI Incident Register is reviewed by the Incident Response Lead and the Legal and Compliance Lead at least quarterly. A summary of incident trends, open remedial actions, and regulatory developments relevant to Cedar Line Creative's AI tools is presented to senior leadership at least twice per year. The register is stored securely, access-controlled to the Incident Response Team, and retained for a minimum of five years from the date of each entry, consistent with applicable statute of limitations periods across Colorado, California, Texas, and Illinois.

Federal Legislation Readiness Assessment

What federal AI legislation would require of you, assessed now so the work is spread rather than compressed.

1. Current Federal Position and the Agencies Already Active

The United States has no single, comprehensive federal AI statute in force. Federal AI governance for Cedar Line Creative is therefore delivered through a combination of executive policy, agency enforcement authority, and sector-specific rules rather than a unified legislative code.

President Trump's Executive Order 14179 (January 2025) rescinded earlier safety-testing mandates and reoriented federal policy towards AI innovation. Executive Order 14365 (December 2025) established a national AI policy framework and created a Department of Justice AI Litigation Task Force whose stated remit is to challenge state AI laws deemed inconsistent with federal policy. This creates a live tension for Cedar Line Creative: the federal posture is permissive, but the DOJ task force may attempt to pre-empt the very state statutes — including Colorado SB 26-189 and California SB 942 — that impose the heaviest obligations on this business. Cedar Line Creative cannot rely on federal pre-emption to shelter it from state duties; that litigation is unresolved and the state laws remain operative.

The Federal Trade Commission is the primary federal enforcer with direct relevance to Cedar Line Creative's current operations. The FTC established a dedicated AI enforcement unit in January 2026 and has pursued 13 AI-washing cases since 2024, many involving deceptive claims made business-to-business. On 1 July 2026 the FTC issued a proposed policy statement warning that altering AI outputs to suppress accuracy — even to comply with state law — could constitute consumer deception under Section 5 of the FTC Act. The FTC's first enforcement action targeting undisclosed AI-generated advertising content has already been filed. Cedar Line Creative publishes AI-generated copy and imagery under client brands and runs AI chat assistants on client websites; both activities sit squarely within the FTC's current enforcement focus.

The Equal Employment Opportunity Commission has issued guidance on AI in employment (EEOC AI in Employment guidance) that applies to Cedar Line Creative's use of an AI screening tool in its own hiring. The EEOC's position is that an employer cannot escape liability for discriminatory screening outcomes by attributing those outcomes to a vendor's algorithm; the employer remains the responsible party under Title VII and the ADA.

The Federal Take It Down Act (TiDA), effective May 2026, criminalises the nonconsensual publication of intimate images including AI-generated deepfakes and imposes removal obligations on platforms. Cedar Line Creative's image generation work should be reviewed against TiDA's scope, although the primary exposure sits with platform operators rather than agencies. The position should be confirmed with specialist counsel if Cedar Line Creative operates any client-facing platform that accepts user-generated content.

2. Requirements a Business of This Profile Should Expect

Cedar Line Creative's profile — a marketing agency that publishes AI-generated content under client brands, operates AI chat tools on client websites, uses AI in its own hiring, and processes personal data through multiple AI systems — creates exposure across every active federal enforcement channel.

Under FTC Section 5, the agency treats the non-disclosure of AI-generated advertising and marketing content as a potentially unfair or deceptive practice. Because Cedar Line Creative publishes AI-generated copy and imagery attributed to client brands, both Cedar Line Creative and its clients face exposure if consumers are misled about the human or AI origin of that material. The FTC does not require proof of intent; the test is whether a reasonable consumer would be deceived. Cedar Line Creative should expect that any AI-generated content placed in consumer-facing channels requires a clear, conspicuous, and proximate disclosure. The FTC's 1 July 2026 proposed policy statement further signals that suppressing, removing, or altering AI provenance signals to avoid state-law obligations will itself be treated as deception — meaning Cedar Line Creative cannot use a state-compliance argument to justify stripping latent provenance metadata.

On AI-washing, the FTC requires that any claim Cedar Line Creative makes to clients about the capabilities, accuracy, or performance of its AI tools must be substantiated before the claim is made. Telling a client that its AI chat assistant is 'highly accurate' or that AI screening 'eliminates bias' without documented evidence constitutes a deceptive business-to-business claim under Section 5.

Under EEOC guidance, Cedar Line Creative's AI screening tool must be validated for adverse impact across protected classes before use, and the agency must be able to demonstrate that validation if challenged. The EEOC's position is that disparate impact caused by an AI tool is legally equivalent to disparate impact caused by a human decision-maker. Cedar Line Creative must audit its screening tool's outputs by race, sex, national origin, disability status, and age, document those audits, and retain records sufficient to defend a charge.

The NIST AI Risk Management Framework 1.0, while not legally binding, is increasingly cited by federal agencies — including the FTC and EEOC — as evidence of reasonable practice. Adherence to NIST AI RMF 1.0's Govern, Map, Measure, and Manage functions is the clearest available signal that Cedar Line Creative has exercised due care. Regulators and insurers are both beginning to treat NIST AI RMF alignment as a baseline expectation rather than a differentiator.

3. Controls Already in Place That Would Carry Across

Cedar Line Creative has not yet indicated that it has formal AI governance structures in place; the client notes confirm there is no disclosure wording anywhere and no view on which legal regimes apply. This assessment therefore identifies the operational facts that, if formalised, would constitute transferable controls rather than assuming controls exist.

Cedar Line Creative already uses named, commercially available AI tools — ChatGPT/OpenAI, Claude/Anthropic, Midjourney, DALL-E, and named CRM and HR systems. This means the business has a de facto AI inventory. That inventory, once documented and assigned business-process ownership, becomes the foundational artefact required by NIST AI RMF (the 'Map' function), by Colorado SB 26-189's deployer duties (effective 1 January 2027), and by any FTC inquiry into what systems were in use at a given time. The work to build a formal AI register is therefore not greenfield; it is the formalisation of what already exists.

Cedar Line Creative operates across Colorado, California, Texas, and Illinois. This multi-state presence means that any disclosure framework Cedar Line Creative builds to satisfy California SB 942 (in force 2 August 2026) will, with minor adaptation, satisfy the disclosure requirements emerging under other state regimes and will also address the FTC's non-disclosure concerns. A single, layered disclosure architecture is transferable across all relevant jurisdictions and is more efficient than building jurisdiction-by-jurisdiction.

The business already has client relationships that involve contractual scope-of-work documents. Those existing contract structures are the vehicle through which AI-specific representations, intellectual property allocation clauses, and indemnity provisions should be inserted. Building those clauses now means every new engagement carries compliant terms without requiring retrospective renegotiation.

4. Gaps That Would Need Closing, Ranked by Lead Time

The gaps are ranked from longest to shortest lead time, because the items that take the most time to build must be started first.

Gap 1 — AI governance framework and NIST AI RMF alignment (lead time: four to six months). Cedar Line Creative has no documented governance structure for AI. Building one that maps to NIST AI RMF 1.0's four functions — Govern, Map, Measure, Manage — requires executive sponsorship, role assignment, policy drafting, and at minimum one cycle of internal review. This is the foundational control on which every other gap-closure depends. Without it, the business cannot demonstrate reasonable care to the FTC, cannot satisfy Colorado SB 26-189's deployer obligations when they take effect on 1 January 2027, and cannot respond coherently to a client or regulator inquiry. This must be started immediately.

Gap 2 — AI screening tool validation for EEOC compliance (lead time: three to four months). The EEOC's position is clear and currently enforceable. Cedar Line Creative uses an AI tool in its own hiring and has no documented adverse-impact analysis. Commissioning an independent validation study, receiving results, and implementing any required adjustments to the screening process takes three to four months at minimum. A hiring decision made tomorrow using the current unvalidated tool creates liability today. This gap is urgent.

Gap 3 — FTC-compliant disclosure wording for AI-generated content (lead time: six to eight weeks). Cedar Line Creative publishes AI-generated copy and imagery under client brands with no disclosure anywhere. Drafting, legally reviewing, and deploying disclosure language across active client campaigns and website chat tools is a six-to-eight-week project when client approval cycles are factored in. The FTC's first enforcement action on undisclosed AI advertising content has already been filed; Cedar Line Creative's current practice matches the fact pattern of that action.

Gap 4 — AI-washing substantiation files (lead time: four to six weeks). For every performance claim Cedar Line Creative makes about its AI tools to clients — accuracy, speed, bias reduction, efficiency — a substantiation file must exist before the claim is made. Creating those files requires pulling vendor documentation, conducting internal tests where vendor documentation is insufficient, and establishing a process for keeping claims current as tools are updated. This is a documentation exercise but it requires staff time and legal sign-off.

Gap 5 — Client contract AI clauses (lead time: two to four weeks for drafting; longer for rollout). Existing contracts are silent on AI use, intellectual property ownership of AI outputs, and liability allocation. Standard clauses covering these points can be drafted in two to four weeks. Rolling them into active contracts requires client agreement and may take longer. New engagements should carry compliant terms from the point the clauses are approved.

Gap 6 — Colorado SB 26-189 deployer readiness (lead time: must be complete before 1 January 2027). This statute is not yet in force but Cedar Line Creative's headquarter jurisdiction is Colorado. The deployer obligations — impact assessments for high-risk automated decisions, disclosure to affected individuals, appeal mechanisms — require process design and system configuration that cannot be done in days. Planning must begin in Q4 2026 at the latest, and the governance framework in Gap 1 is a prerequisite.

5. What to Build Now Because It Is Expensive to Retrofit

Three categories of infrastructure are significantly cheaper to build into Cedar Line Creative's operations now than to retrofit after enforcement action, client demand, or statutory commencement forces the issue.

First, provenance and disclosure architecture for AI-generated content. California SB 942 is in force as of 2 August 2026 and requires large generative AI providers to embed latent provenance disclosures in AI-generated content. Cedar Line Creative uses those providers — OpenAI's DALL-E, Midjourney, and others — and publishes their outputs under client brands. The FTC's 1 July 2026 proposed policy statement makes clear that stripping or suppressing those provenance signals to avoid disclosure obligations will itself be treated as deception. Cedar Line Creative must therefore build a content workflow that preserves latent provenance metadata from generation through to publication, and that adds consumer-visible disclosure at the point of publication. Building this into the production workflow now — as a standard step in every AI content brief — costs a fraction of what a retroactive audit of all published client content would cost, and far less than defending an FTC enforcement action. The architecture also needs to be documented so that Cedar Line Creative can demonstrate to a regulator exactly when disclosure was added and how metadata was preserved.

Second, a vendor due diligence and AI register process. Cedar Line Creative currently uses at least five categories of AI tool. Each tool raises distinct questions: what data does it process, where is that data stored, what are the vendor's terms on ownership of outputs, what indemnities does the vendor provide against IP infringement claims, and what are the vendor's own compliance representations? These questions are cheap to ask at procurement and expensive to answer retrospectively when a client asks for a data processing agreement or when an IP dispute arises over a generated image. An AI register that captures tool name, vendor, data processed, business process, legal basis, vendor contract reference, and last-reviewed date should be built now as a living document. Every new tool goes through a documented intake checklist before deployment. This is the Map function of NIST AI RMF 1.0 and the prerequisite for every other compliance obligation.

Third, an employment AI governance protocol covering the hiring screening tool. The EEOC's guidance is enforceable today. An adverse impact finding against Cedar Line Creative's screening tool, combined with an absence of any validation records, would expose the business to individual and class claims under Title VII and the ADA, with the employer unable to point to any remedial steps taken. The protocol to build now includes: a written validation requirement before any AI tool is used in a hiring decision; a record-retention policy for screening outputs and decisions; a documented human-review step before any rejection is communicated to a candidate; and a disclosure to candidates that AI is used in the screening process. These elements are inexpensive to build into HR policy; they are very expensive to reconstruct after a charge is filed. Illinois BIPA, which applies to Cedar Line Creative's Illinois staff, adds a separate biometric data layer: if the screening tool uses facial analysis, voice analysis, or any other biometric identifier, written consent and a public retention policy are required before any data is collected. This should be confirmed with the HR tool vendor immediately.