

Use these templates as starting points and adapt them to your specific context and brand voice. **These are frameworks, not scripts.**

From [Guardrails: How To Embrace AI in Your Business Without Damaging Your Brand](#) by Craig McDonogh.

CUSTOMER AI DISCLOSURE LANGUAGE

For AI-assisted content

Use when using human-reviewed and edited AI-generated output

"This [article/description/analysis] was created with AI assistance and reviewed by our [team/editorial staff] to ensure accuracy and quality."

What makes this work: transparent about AI involvement without overstating it. The human review step is the reassurance customers actually need.

For AI-generated content

Use when AI produced the output with post-hoc review rather than active editing

"This [article/description/analysis] was generated by AI based on [data sources/parameters]. While we review all AI-generated content, please [contact us/report issues] if you notice errors."

What makes this work: honest about the process; gives customers a clear action if something is wrong.

For AI-powered features

Use in product UI, onboarding flows, and feature announcements

"This feature uses AI to [specific function]. You can [opt out/adjust settings/use alternative] if you prefer."

What makes this work: names the function specifically; always provides an exit path.

For AI customer service

Use at the start of any AI-handled customer interaction

"You're currently chatting with our AI assistant. For complex issues or if you'd prefer human assistance, [say 'agent'/click here/press 0]."

What makes this work: discloses AI status upfront; makes escalation to a human obvious and easy.

For AI-powered recommendations

Use in any personalization, product recommendation, or content suggestion context

"These recommendations are generated by AI based on [criteria]. They may not reflect all your preferences or needs. [Here's how to adjust them]."

What makes this work: sets honest expectations; gives customers agency over their experience.

CRISIS STATEMENT FRAMEWORKS

Initial acknowledgment (First 24 hours)

Publish within hours of becoming aware - do not wait for a full investigation

"We have become aware of [specific issue] affecting [who]. We take this very seriously and have taken immediate [specific action]. We are investigating the full scope and will provide updates as we learn more. [Contact information for affected parties]."

What makes this work: specific issue identification, explicit immediate action, commitment to updates, and support for affected parties. Silence in the first 24 hours is almost always the wrong choice.

Complete response (After investigation)

Publish once root cause is confirmed - typically 48-72 hours after initial acknowledgment

"Our investigation has confirmed that [specific problem] occurred, affecting [number/type] of [customers/interactions]. This happened because [specific cause]. We apologize to everyone affected.

We have taken the following immediate actions: [specific steps]. To prevent recurrence, we are implementing [specific changes].

For those affected, we are [specific remediation]. [Contact information and timeline].

We recognize this represents a failure of [our testing/our oversight/our judgment], and we are committed to doing better."

What makes this work: specific problem and cause, actual actions taken, meaningful remediation, and ownership of failure. Vague language here destroys the credibility that the initial acknowledgment built.

Bias incident statement

Use when an AI system has produced disparate outcomes across demographic groups

"Our [system name] produced disparate outcomes for [affected groups], resulting in [specific harm]. This is unacceptable and contrary to our values.

We have immediately [suspended the system/implemented manual review/taken specific action] while we investigate root causes and develop corrective measures.

We sincerely apologize to everyone affected and are committed to [specific timeline] for resolution. [Here's what affected parties should do]."

What makes this work: acknowledges specific bias and harm; explicit, immediate action; a timeline commitment; guidance for affected parties.

Privacy breach notification

Use for required breach notifications - check applicable state and federal timing requirements

"We are writing to inform you that between [dates], data about your account was [exposed/accessed] due to [specific issue]. The affected data included [specific types].

We have secured the system and are investigating the incident. We are offering [credit monitoring/specific assistance] to all affected individuals.

[Here's what you should do]. [Here's how to contact us]. [Here's our timeline for updates].

We sincerely regret this incident and are taking steps to prevent recurrence, including [specific actions]."

What makes this work: specific timeframe and data types, clear security actions, meaningful assistance offered, and action guidance.

EMPLOYEE POLICY COMMUNICATION

Introducing AI governance

Use for all-staff communications when rolling out a new governance program

"We're implementing governance for AI use across the organization. This isn't about limiting innovation - it's about using AI responsibly and protecting our reputation and our customers.

Here's what this means for you: [specific requirements for their role]. Here's why this matters: [specific risks in their context]. Here's where to get help: [specific resources and contacts].

We're all learning how to use AI effectively and responsibly. These policies will evolve as we learn. Your feedback on what's working and what isn't helps us improve."

What makes this work: frames governance as enabling rather than limiting; provides role-specific guidance; acknowledges the learning process. Governance that feels like punishment gets worked around.

AI tool approval process

Use when communicating the vendor review and approval workflow to staff

"Before using new AI tools for [customer-facing work/content creation/decision-making], please complete the AI assessment template [link] and get approval from [specific role].

This takes [realistic time estimate] and helps us identify potential risks before they become problems. Most tools get approved quickly — the process helps us understand what we're using, rather than blocking useful tools.

If you're unsure whether something needs approval, err on the side of asking. [Contact information]."

What makes this work: clear scope, realistic expectations, explanation of purpose without being defensive, and encouragement of questions.

AI incident reporting

Post in internal channels and include in onboarding for all roles using AI tools

"If you notice AI systems producing concerning results - including bias, errors, inappropriate content, privacy issues, or anything else that seems wrong - please report it to [contact] immediately.

This isn't about blame. Early identification of problems helps us fix them before they escalate. You're helping protect our customers and our reputation.

What to report: [specific examples]. How to report: [simple process]. What happens next: [realistic timeline and process]."

What makes this work: no-blame framing, clear examples of what to report, a simple process, and it explains the outcome. Blame culture kills early warning systems.

VENDOR COMMUNICATION

AI vendor due diligence request

Send at the start of any new vendor evaluation or significant product change

"As part of our AI governance process, we need additional information about [specific AI system/feature]. Please provide: [Specific list of due diligence questions relevant to this vendor]

We need this information by [realistic date] to complete our assessment. This is standard due diligence we're conducting across all AI vendors, not a specific concern about your system.

[Contact] can answer questions about what we need and why."

What makes this work: clear information request, reasonable deadline, framed as a standard process rather than an accusation. Adversarial opening communications create adversarial vendor relationships.

Vendor performance issue

Use when a vendor's AI system is underperforming or creating risk - send before escalating

"We've identified concerns with [specific AI system] performance: [particular issues with data/examples]. This affects [specific stakeholders] and creates [specific risks].

We need [specific actions] by [realistic date]. Can we schedule a call [date/time] to discuss the resolution?

We value our partnership and want to resolve this collaboratively. Please let us know your proposed approach."

What makes this work: specific problem identification with evidence, clear needed actions, and a collaborative tone while being firm.

PARTNER COMMUNICATION

Partner AI disclosure notice

Send before introducing AI into any workflow that involves shared data or joint deliverables

"We're introducing [AI tool/system] into [shared workflow/campaign name] starting [date]. Here's what this means for our partnership:

[Data involved: describe which shared data the AI system will access or confirm it accesses none.]

[Output ownership: who owns AI-generated outputs under your current agreement, or flag if this needs to be addressed.]

All AI-assisted work will be reviewed by [role/team] before it reaches you or your customers — our approval process is unchanged.

[Any change to how your team works, or confirmation that there is none.]

We wanted to let you know before we proceed, not after. If you have questions or concerns before [go-live date], please reach out to [contact]. We'll proceed on [date] unless we hear from you."

What makes this work: treats partners as stakeholders rather than afterthoughts. The go-live date with a window to respond creates accountability without pressure.

Partner incident notification

Send within hours of identifying an incident affecting a partner — never wait until you have all the answers

"I'm writing to let you know about an issue with [system name] that has affected [brief description]. We identified this at approximately [time] on [date].

What happened: [plain-language description. If investigation is incomplete: 'We are still investigating and will have more information by [date/time].']

What we've done so far: [specific immediate actions — e.g., system suspended, affected content removed, review team assembled.]

What we need from you right now: [specific ask — e.g., 'Please hold the [deliverable] sent on [date] and don't distribute it until further notice.' Or: 'No action needed from you at this stage.']

We will send a full incident report, including root cause and remediation plan, by [date — no more than 72 hours from now]. Your direct contact for this incident is [name, phone, email]."

What makes this work: prompt, honest notification preserves the relationship even when the news is bad. An honest 'we don't know everything yet' is always better than silence.