

AI and cybersecurity: trust in real-world use

Technologies, AI & cybersecurity / 8 pages / 3 October 2026

Proposed decision: link each significant technology claim to a defined use, evidence and follow-up responsibility.

Central risk

A company can properly describe a functionality and suggest that it guarantees a broader result. The terms "autonomous", "secure" or "reliable" then become implicit commitments on errors, data and human work. The committee must examine what the public understands, beyond what the team thinks they have said.

Recommended decision

Create a common process between product, security, legal, sales and communication. For sensitive uses, an announcement must specify the context of evaluation, known limits and expected supervision. The communication must have a right of alert when a formulation exceeds the available evidence, without becoming the sole arbitrator of technical risk.

Three immediate decisions

Identify the person responsible for public statements about the product. Set situations that require a review of content or campaign suspension. Validate a response process that works even when the cause of an incident remains unknown. Commercial launch and explanation capacity must move forward together.

DECISION POINT

If a user takes our promise literally, are we ready to stand behind the use they infer from it?

Moving from demonstration to actual use

The success of a demonstration alone does not describe the robustness of a service in working situations.

Name the promise precisely

A system can assist a task, propose an answer or make a decision in a limited framework. These three situations do not involve the same responsibilities; therefore, a presentation must describe who is acting, on what data, with which possibility of control and under what conditions the user should interrupt the process. The innovation vocabulary should not erase this allocation of responsibilities.

Read the evidence with its context

An evaluation result depends on a dataset, a method, a version of the model and a usage scenario. It is necessary to explain what has been tested and what has not been tested. The voluntary NIST risk management framework, together with its generative AI profile, provides benchmarks for identifying and managing risks [1]. It does not constitute automatic product certification.

Integrate chain of dependencies

The proposed experience sometimes depends on a model provider, cloud, external data and third party integrations. However, reputation often focuses on the visible brand. The committee must know the dependencies that change the quality of service, availability or usage conditions. Communication must be able to explain them at a useful level, without publishing details that would undermine security.

DECISION POINT

A known limitation, but difficult to find, remains a poorly communicated limitation.

Put promises under control

Governance should focus on concrete formulations, not just on a general responsible AI charter.

A four-entry register

For each important statement, keep the published sentence, the test supporting it, the associated limits and the owner of its update. The register covers product pages, sales arguments, interviews and documents intended for partners. An amendment to the product must trigger a review of the dependent statements.

Use proportionate review

All sentences do not require the same circuit. A stable factual description can follow a common validation. A statement of performance, safety or autonomy requires the advice of the competent officials. The communications department formalizes the possible consequences of an ambiguity; the committee decides on acceptable risk and the resources to reduce it.

Avoid two pitfalls

The first is to make the limits so discreet that the main message contradicts them. The second is to accumulate incomprehensible caveats to the point of no longer explaining usage. The right level of accuracy allows a reasonably informed user to choose how to use the service and when to request human intervention.

Claim	Control question	Expected answer
Performance	Which test and version?	A contextualised, reproducible result
Security	What scope is covered?	Explicit measures and limits
Autonomy	Who approves or stops the process?	An identifiable responsibility

DECISION POINT

A change in evidence must be able to change the sentence, even after a launch.

Preparing for incidents that undermine trust

The following exercises are fictitious. They are used to test decisions, responsibilities and initial responses.

An erroneous result becomes public

A user publishes an incorrect response produced by the service. The team must distinguish between the existence of the case, its reproducibility and its possible extension. Recognizing a documented example does not mean immediately finding a general failure. Conversely, merely referring to a terms-of-use clause leaves unanswered the question of actual experience and available protections.

Unavailability reveals dependence

An essential supplier is experiencing a breakdown. The brand must communicate what its users can do, the functionality affected and the time of the next update. It does not automatically repeat the vendor recovery schedule as a firm commitment. The committee determines the continuity of service and the level of transparency on dependency.

Suspected unauthorised access becomes public

Security leads the investigation; authorized teams evaluate notification obligations. Communication prepares messages consistent with verified facts and protection instructions. It clearly distinguishes unavailability, compromise and data leakage: one does not establish the others. New information must be able to correct a first response without erasing its evolution.

DECISION POINT

Separate confirmed facts, points under analysis and action to be taken for the persons concerned.

Turn expertise into credible communication

The visibility of an expert becomes useful when it improves understanding of a decision or risk.

Executive communication about choices

The leader explains the uses chosen, the development decisions and the limits that he agrees to make visible. Experts describe the methods and results. Security officials specify what can be shared without harming protection. This complementarity avoids asking a single person to represent every area of expertise.

Media relations that withstand difficult questions

Prepare a file to distinguish between available capabilities, experimental features and development objectives. Propose demonstrations whose conditions are known. In case of critical questions, the spokesperson must be able to answer on the test method and indicate what he does not know yet. A precise correction is more useful than a general promise of reliability.

Associate teams with the explanation

Sales teams, support and partners are reputational interfaces. They need examples of usage, limitations and a circuit to bring up misinterpretations. If an objection often comes back, the product or its explanation must be examined, rather than merely refining the sales response. This loop makes communication a tool for understanding real uses.

DECISION POINT

A sheet "what the service does, doesn't do and requires from the user", linked to its version.

GEO as a precision discipline

The challenge is to make the expertise understandable in search engines and AI responses, without promising a mastery of their results.

Clarify entities and uses

Describe separately the company, its products and their versions. Linking each capacity to a use, documentation and an interlocutor. A page that overlaps institutional discourse, catalog and superlative makes it difficult to distinguish. Direct definitions and answers must remain consistent with the conditions of service and technical documents.

Publish usable evidence

Evaluations gain in specifying protocol, date and limits. Public documentation must be accessible and important content available in text form. Google indicates that sound search practices remain relevant to its AI features [2]. No mark-up or editorial format guarantees an inclusion or citation. GEO does not therefore authorize a promise of guaranteed ranking.

Organize observation

Follow questions of comparison, understanding and risk in the languages actually used by customers. Keep answers and their sources to compare equivalent observations. A variation may come from the question, service or its evolution; it should not be attributed automatically to a modification of the site. Prioritize errors with real consequence rather than the simple volume of mentions.

DECISION POINT

When an AI describes our product, does it also reproduce the limits necessary for appropriate use?

90 days to connect product and reputation

The proposed programme focuses on the statements made and situations of breach of trust.

First month - Delimit

Select priority products and uses. Compare published promises with technical documentation, testing and feedback from support. Identify vocabulary gaps and unclear limits. Identify an owner for each critical claim and define a rule of review when product, supplier or usage context changes.

Second month - Test

Organize an exercise bringing together product, security, legal, support and communication. To introduce successively a new fact, a journalist's request and a customer's question. Check that the circuit produces an understandable answer without waiting for an impossible certainty. Correct message templates, documentation and responsibilities that slow down the decision.

Third Month - Install Routine

Update priority content and train the people who use it. Test a stable set of search and AI questions. The committee examines resolved gaps, unresolved risks and resources. The objective is a sustainable mechanism for correction, not a communication operation limited to launching a new service.

DECISION POINT

Support, sales and the manager can explain the same ability with the same limitations.

The trust dashboard

Monitoring combines quality of evidence, public experience and ability to correct.

Indicators to be selected

Measure the share of documented critical statements, the time to update after product evolution and recurrent misunderstandings raised by support. For GEO, distinguish presence, accuracy and quality of sources. Calculation bases must be constant and protocol changes clarified.

A signal that requires a decision

A high level of satisfaction can coexist with an uncommon use case with serious consequences. The committee must see severe cases, detection limits and disagreements between teams.

Communication should not dramatize an isolated example or hide a material risk behind an average.

The decision to formalize

Validate the owners of promises, alert thresholds and the ability to suspend a formulation that has become inaccurate. Set up a joint review with product and security. Request a return on the corrections made and the reasons why certain requests for stronger claims have been delayed.

Technological confidence is built when the explanation, evidence and experience remain aligned.

Sources and reference points

[1] NIST — AI Risk Management Framework / Generative AI Profile

[2] Google Search Central — AI features and your website

Analysis and recommendations by Belief System. Scenarios are illustrative. Sources provide context; they do not constitute endorsement of our recommendations.