


Communications Has a Shadow AI Problem

When adoption outruns governance, trust becomes an operating risk



ARTEFACT91

Campaign Intelligence + Strategic Communications

artefact91.com

In this briefing

When adoption outruns governance, trust becomes an operating risk

·	Executive Brief
01	Shadow AI Is a Communications Governance Problem
02	The Problem Is Not Only the Tool
03	The Artefact91 Shadow AI Exposure Model
04	Why Communications Needs a Seat in AI Governance
05	Build a Communications AI Control Plane
06	Human-in-the-Loop Is Not a Control Unless the Human Can Check
07	The Communications AI Pre-Mortem
08	The Trust Question Is Proof, Not Intent
09	A 30-Day Communications AI Governance Sprint
10	What Leaders Should Measure
·	Artefact91 Point of View
·	Questions for Leaders
·	Figure 1 — AI Adoption Is Outrunning Governance
·	Sources & Research Base

HOW TO READ IT

A briefing, not a report: the executive brief carries the argument, the numbered sections carry the evidence and the method, and every external figure is attributed to its source.

ARTEFACT91 BRIEFING 13

Executive Brief

AI is now embedded in communications work faster than most organizations have built the rules, training, verification systems and decision rights needed to govern it. The immediate risk is not simply that an AI tool produces a bad sentence. It is that communications teams can move confidential information, unverified claims, stakeholder data and institutional voice through systems the organization cannot fully see or audit.

Fresh evidence makes the governance gap difficult to dismiss. PRSA's September 2026 ethics analysis, drawing on new Delphi research with advanced AI adopters and earlier Global Alliance research, reports pressure to adopt AI rapidly with limited training or ethical guidance; unauthorized tools are being used, confidential or proprietary information may be exposed, and AI-generated material can reach the public without adequate verification. The cited Global Alliance research found that only 39.4% of organizations reported responsible-AI guidelines, 38% reported having none, and 57.5% of PR and communication managers said they had no role in AI governance.

The communications-specific evidence now converges with wider enterprise findings. OneTrust reported in September 2026 that 86% of surveyed organizations experienced at least one AI-related incident in the previous year, while only 27% slowed or paused deployment; 33% reported employees using unapproved AI because approved tools or processes were not available quickly enough. Akamai separately reported that more than 40% of enterprise users had installed AI-powered browser extensions and that more than 6% of chatbot conversations contained sensitive information.

This is a distinct communications risk. The issue is no longer whether communications teams will use AI. It is whether institutions can prove how AI-assisted communications were produced, verified, authorized and protected.

SECTION 01

Shadow AI Is a Communications Governance Problem

Shadow AI is the use of AI systems, agents, browser extensions, plug-ins or workflows outside the organization's approved governance, procurement, security or review processes.

In communications, the exposure is unusually broad because teams routinely handle:

- unreleased announcements and executive remarks;
- crisis facts and legal-sensitive drafts;
- employee and stakeholder information;
- media lists and relationship intelligence;
- campaign strategy and opposition analysis;
- embargoed research and financial information;
- client or partner materials;
- reputation monitoring data;
- institutional credentials and platform access.

A productivity shortcut can therefore become a confidentiality, accuracy, governance or reputation event.

SECTION 02

The Problem Is Not Only the Tool

Organizations often respond by creating a list of approved and prohibited applications. That is necessary but incomplete.

Risk also depends on the task, information and authority involved.

A public brainstorming prompt is not equivalent to uploading an unreleased board paper. Drafting headline options is not equivalent to allowing an agent to publish. Summarizing a public report is not equivalent to generating a factual crisis response from incomplete internal notes.

Artefact91 therefore treats AI communications risk as the interaction of five variables:

DATA — What information enters the system?

DECISION — What institutional judgment is being delegated or influenced?

ACTION — What can the system produce, send, publish or change?

VERIFICATION — Who checks the output, against what evidence?

ACCOUNTABILITY — Who remains answerable when the system is wrong?

SECTION 03

The Artefact91 Shadow AI Exposure Model

Classify AI-enabled communications work across four levels.

LEVEL 1 — ASSIST — Low-consequence tasks using public or non-sensitive information: ideation, formatting, transcription or first-pass synthesis. Human review remains required.

LEVEL 2 — INFORM — AI contributes analysis that may shape communications decisions: audience synthesis, issue monitoring, research summaries, message testing or scenario generation. Source traceability and analytical review are required.

LEVEL 3 — REPRESENT — AI produces material that can be interpreted as the institution speaking: executive drafts, stakeholder responses, media statements, social copy, public FAQs or campaign assets. Fact verification, disclosure rules and accountable approval are required.

LEVEL 4 — ACT — AI or an agent can send, publish, purchase, alter records, contact stakeholders or trigger workflows. Explicit authorization, logging, access controls, monitoring and stop mechanisms are required.

The higher the level, the less acceptable informal experimentation becomes.

SECTION 04

Why Communications Needs a Seat in AI Governance

The finding that 57.5% of PR and communications managers reported no role in AI governance is strategically important.

AI governance is often treated as an IT, security, privacy or legal issue. Those functions are essential, but they do not fully govern institutional meaning.

Communications leaders understand questions such as:

- Could this output be mistaken for an official institutional position?
- What disclosure would a stakeholder reasonably expect?
- Could automation undermine a relationship that depends on human judgment?
- Does the generated claim have evidence strong enough for public use?
- Could a synthetic asset mislead even if it is technically labelled?
- What happens when AI produces inconsistent answers across channels?
- Who corrects the record when an automated system speaks incorrectly?

The communications function should therefore participate in AI governance wherever AI can shape public claims, stakeholder interactions, institutional identity or reputation.

SECTION 05

Build a Communications AI Control Plane

Artefact91 recommends six controls.

INVENTORY — Know which tools, extensions, agents and integrations are actually being used—not only which ones have been purchased.

INFORMATION RULES — Define what categories of information may and may not enter external AI systems. Make the rules task-specific and understandable.

AUTHORITY RULES — Define what AI may draft, recommend, send, publish or execute. Separate assistance from autonomous action.

VERIFICATION — Match the review standard to consequence. Public facts, quotations, statistics, legal claims and crisis information should be checked against authoritative evidence.

TRACEABILITY — For consequential work, preserve enough information to reconstruct the workflow: source material, tool or model, material AI contribution, human reviewer, approval and final version.

ESCALATION — Create a clear route for hallucinations, confidential-data exposure, impersonation, unauthorized publishing, compromised accounts, harmful automation or unexplained agent behaviour.

SECTION 06

Human-in-the-Loop Is Not a Control Unless the Human Can Check

Many organizations describe human review as the solution to AI risk. That phrase can create false assurance.

A reviewer cannot meaningfully verify an output when:

- the source evidence is unavailable;
- the workload makes review superficial;
- the reviewer assumes the model is more accurate than it is;
- the output is outside the reviewer's expertise;
- the system has already taken an external action;
- no audit trail exists;
- the organization has not defined what must be checked.

The control test

Human review becomes a control only when the reviewer has evidence, authority, time and a defined verification standard.

SECTION 07

The Communications AI Pre-Mortem

Before expanding AI use, assume an AI-enabled communications workflow creates a public incident six months from now.

Ask:

- What information leaked?
- Which claim was wrong?
- Which automated action occurred without meaningful approval?
- Which audience felt deceived?
- Which vendor, extension or integration did we not know was in the workflow?
- Which employee believed they were following normal practice?
- What evidence would we need to reconstruct what happened?
- Could we explain our controls publicly without sounding evasive?

Turn the answers into controls before deployment.

SECTION 08

The Trust Question Is Proof, Not Intent

The emerging governance evidence suggests a broader institutional problem: organizations can believe their AI use is responsible without being able to demonstrate why.

For communications, proof means being able to answer:

- What was AI used for?
- What information did it access?
- What could it do?
- Who reviewed the output?
- What evidence supported public claims?
- What was logged?
- What happened when something went wrong?

This is the difference between saying “we use AI responsibly” and having an auditable communications practice.

SECTION 09

A 30-Day Communications AI Governance Sprint

WEEK 1 — DISCOVER — Inventory tools, extensions, agents, recurring use cases and sensitive-data pathways.

WEEK 2 — CLASSIFY — Assign use cases to Assist, Inform, Represent or Act. Identify prohibited information and high-consequence workflows.

WEEK 3 — CONTROL — Define verification, approval, logging, disclosure and escalation requirements for each level.

WEEK 4 — TEST — Run a simulation: an AI-generated public claim is wrong, a confidential document is entered into an unapproved tool, or an agent publishes without expected review. Measure how quickly the organization can establish facts, contain exposure and reconstruct the workflow.

SECTION 10

What Leaders Should Measure

Do not measure AI adoption alone.

Track:

- approved versus observed tools;
- percentage of staff trained for their actual use cases;
- high-consequence workflows with defined approval;
- percentage of consequential outputs with traceable evidence;

- AI-related incidents and near misses;
- time to detect and contain unauthorized AI use;
- percentage of Level 4 actions with explicit authorization and logging;
- recurring reasons employees bypass approved tools.

The metric that explains the rest

Shadow AI often grows where governance makes legitimate work unnecessarily difficult.

ARTEFACT91 BRIEFING 13

Artefact91 Point of View

The communications profession is entering an accountability phase of AI adoption.

The first phase asked: Can AI make communications faster? The next asks: Can the institution prove that faster communications remain accurate, authorized, secure and worthy of trust?

The strategic response is not to ban experimentation. It is to govern consequence.

ARTEFACT91 COMMUNICATIONS AI GOVERNANCE SEQUENCE

DISCOVER → CLASSIFY → AUTHORIZE → VERIFY → TRACE → ESCALATE → LEARN

The goal is not zero AI. It is zero ambiguity about who is accountable when AI touches institutional voice.

ARTEFACT91 BRIEFING 13

Questions for Leaders

Which AI tools are communications staff actually using today?

What information are they putting into them?

Where can AI move from drafting to representing the institution?

Which workflows allow an AI system to act externally?

Could we reconstruct a consequential AI-assisted communication after the fact?

Does communications have a formal role in enterprise AI governance?

Figure 1 — AI Adoption Is Outrunning Governance

EXPOSURE LEVEL	AI ROLE	MINIMUM CONTROL
1 — Assist	Ideate, format, summarize public material	Human review
2 — Inform	Shape analysis or decisions	Source traceability and analytical review
3 — Represent	Speak in institutional voice	Fact verification, accountable approval, disclosure rules
4 — Act	Send, publish, purchase or trigger workflows	Explicit authorization, logging, monitoring, stop mechanisms

Sources & Research Base

PRSA (September 2026), “Ethically Navigating the Waves of AI Integration.” New Delphi research with advanced adopters and discussion of Global Alliance research on responsible-AI guidance and communications governance.

OneTrust (September 2026), AI governance research: 86% of surveyed organizations experienced an AI-related incident; 33% reported unapproved AI use; 98% planned to increase AI-governance technology budgets.

Akamai Technologies (22 September 2026), “Speed, Scale, and Nonhuman Identity: The Agentic Threat Landscape.” Reported more than 40% of enterprise users installing AI-powered browser extensions and more than 6% of chatbot conversations containing sensitive information.

Muck Rack, State of AI in PR 2026, summarized by University of Minnesota: approximately 75% of surveyed communications professionals used generative AI; 51% reported an AI-use policy and 43% reported workplace AI training.

Research Note

Artefact91 analysis and frameworks. External statistics are attributed to their sources. The Shadow AI Exposure Model, Communications AI Control Plane and governance sequence are Artefact91 analytical frameworks and should be adapted to the organization’s legal, security, operational and jurisdictional context.