

{intradiegetic

Building the next communication OS

# The Human × AI Toolkit for Corporate Communications

THINK

SPEAK

DECIDE

EXIST

A practical operating guide for preserving judgment, voice, and trust while redesigning communications work with AI.

Prepared by Intradiegetic · May 2026



## Paul Floren

Founder & CEO,  
Intradiegetic

AI gives communications leaders a chance to do more than accelerate existing workflows. It creates an opening to redesign the function itself: how narrative is built, how decisions are supported, how teams coordinate, and how value is measured.

That redesign only works if capability and governance evolve together. The objective is not automated output for its own sake. It is a communications operating model that is more coherent, more strategic, and more resilient under pressure.

This toolkit is a practical companion to that work: a way to allocate tasks between humans and AI, to set the human checkpoints that protect voice and trust, and to sequence implementation with the discipline that senior stakeholders expect.

## How to use this document

Read it once end to end to align on principles. Then use the tools at the back — the work allocation map, the human checkpoints, the four operating workflows, the governance kit, and the templates — as the starting point for your own internal standard. Adapt the language to your organization. Keep the discipline.

## SECTION 01

# Why Human × AI now.

Communications functions do not need more AI tools. They need an operating model for where machines accelerate work and where human judgment remains sovereign.

# The terrain has shifted

Most communications functions now have generative AI in the building. Drafting, summarization, translation, image and video generation, research, monitoring and synthesis are all available at the press of a button. The marginal cost of producing a competent piece of content has collapsed.

What has not collapsed is the cost of getting communication wrong. Misjudged tone in a crisis, an executive statement that does not match the strategy, a regulatory disclosure that drifts from the legal position, a proactive narrative that lands flat with employees — these still cost trust, value, and time. The economics of mistakes have, if anything, intensified.

## The wrong question, and the right one

Many programs start with: how do we automate more? A more useful question is: where do we want human judgment to be sovereign, and how do we build an operating model that protects it while letting machines do the rest?

This is what Intradiegetic calls building architecture, not tools. It is also what makes the difference, in practice, between communications teams that compound advantage with AI and those that quietly erode their own credibility.

### What this toolkit gives you

- A shared principle for splitting work between humans and AI.
- A four-layer operating model — Think, Speak, Decide, Exist.
- A work allocation map with four zones.
- Human checkpoints for high-consequence communication.
- Four ready-to-adapt operating workflows.
- A governance kit and templates for immediate use.

The aim is not novelty. It is a communications operating model that senior leaders, legal counsel, and the board can recognize as serious.

## SECTION 02

# The Human × AI principle.

AI handles scale, pattern, drafting, routing, and memory. Humans own intent, meaning, ethics, accountability, and final judgment.

# One sentence to remember

Machines accelerate the work. Humans own the meaning.

This is the spine of the operating model. Every workflow, policy, and tool decision should clarify which side of that line it sits on. Most failures in AI-assisted communications come from blurring that boundary — either by asking humans to do work that the machine can do better, or by letting the machine occupy ground that only humans should hold.

## AI DOMAIN

### Scale

Producing variants, translations, summaries, re-cuts.

### Pattern recognition

Detecting trends, anomalies, sentiment shifts in stakeholder signals.

### Drafting

First-pass copy, structure, alternative framings.

### Routing

Sorting incoming requests, classifying issues, triaging queues.

### Memory

Holding institutional context, prior statements, brand history.

## HUMAN DOMAIN

### Intent

Deciding why we are saying anything at all.

### Meaning

Interpreting what stakeholders actually need to hear.

### Ethics & trade-offs

Naming the values at stake when they conflict.

### Accountability

Owning the consequences of what is said and not said.

### Final judgment

The last yes or no, especially under pressure.

### Relationship calls

Reading the room with a regulator, a journalist, a CEO, a workforce.

## How to apply the principle

Each time you design a workflow or evaluate a tool, ask three questions: (1) What part of this is genuinely scale or pattern work? (2) What part requires judgment, intent, or relationship? (3) Where is the human signature on the final output, and is it still honest? If a workflow erases the human signature on something that demanded one, redesign the workflow.

## SECTION 03

# The four-layer communication OS.

Think, Speak, Decide, Exist. Each layer has a human role, an AI role, a failure mode if unmanaged, and a practical intervention.

# Why a layered model

Communications is too often described as content production. That framing leads to AI deployments that optimize the wrong thing. The Intradiegetic framework treats communications as a four-layer system: how the organization thinks, how it speaks, how it decides, and how it exists in the perception of stakeholders. Each layer has its own human and machine roles, and its own failure modes.

<b>THINK</b> Sense-making and narrative architecture	
<b>HUMAN ROLE</b> Frame the strategic story. Decide what is and is not the point. Hold the long-term narrative arc.	<b>AI ROLE</b> Synthesise inputs, surface patterns across signals, draft alternative framings, maintain context memory.
<b>FAILURE MODE</b> AI-generated narratives that are coherent but strategically empty — smooth prose, no point of view.	<b>INTERVENTION</b> Brief AI with the strategic question, not the deliverable. Require a one-line thesis from a named human before any external draft is produced.
<b>SPEAK</b> Voice, language, and content	
<b>HUMAN ROLE</b> Own tone, register, and edge. Approve anything that goes out in the name of an executive or the institution.	<b>AI ROLE</b> Produce drafts, variants, translations, accessibility versions, and channel adaptations at scale.
<b>FAILURE MODE</b> Voice drift — outputs that sound like a generic AI rather than the organization. Trust erodes quietly.	<b>INTERVENTION</b> A maintained voice profile, brand glossary, and a human review checkpoint for any externally facing copy of consequence.
<b>DECIDE</b> Choices, escalations, and accountability	
<b>HUMAN ROLE</b> Choose what to say, what to delay, what to refuse. Carry the consequence.	<b>AI ROLE</b> Lay out options, prior precedents, likely reactions, and missing information. Never autonomously commit the organization.
<b>FAILURE MODE</b> Decision laundering — the team treats the AI's recommended option as a decision, removing accountability from named humans.	<b>INTERVENTION</b> Every consequential decision has a named owner, a written rationale, and a recorded refusal option.

<b>EXIST</b> Presence in stakeholder perception	
<p><b>HUMAN ROLE</b> Tend relationships. Calibrate presence over time. Read what the audience actually feels.</p>	<p><b>AI ROLE</b> Continuously monitor stakeholder signals, measure narrative alignment, surface drift in perception.</p>
<p><b>FAILURE MODE</b> Performative omnipresence — high output, declining resonance. The brand is everywhere and felt nowhere.</p>	<p><b>INTERVENTION</b> Set a deliberate presence cadence. Measure resonance and trust, not just volume. Cut content rather than dilute meaning.</p>

## SECTION 04

# The work allocation map.

Four zones for splitting work between humans and AI: Automate, Augment, Escalate, Prohibit. Use it as a triage tool, not a one-off audit.

## Two axes, four zones

Plot any communications task on two dimensions: consequence if it goes wrong (low to high) and need for human judgment (low to high). The four resulting zones tell you how to staff and govern the task. The same task can sit in different zones in different contexts — the matrix is a discipline, not a label.

<p><b>ESCALATE</b></p> <p>High consequence, high judgment. AI prepares the case; a named human (and often a small group) makes the call.</p> <p>e.g. CEO statement on a crisis; regulatory filing; M&amp;A communications; redundancy announcement.</p>	<p><b>PROHIBIT / PROTECT</b></p> <p>High consequence, but the task is human by definition. AI is kept out of the production loop.</p> <p>e.g. Reading the room with the board; a one-on-one apology; sensitive whistleblower handling; employee bereavement.</p>
<p><b>AUTOMATE</b></p> <p>Low consequence, low judgment. AI runs the task; a human spot-checks outputs and watches for drift.</p> <p>e.g. Repetitive newsletter assembly; tagging and routing inbound mentions; transcription; meta-data and accessibility versions.</p>	<p><b>AUGMENT</b></p> <p>Low consequence, high judgment. AI accelerates a human; the human decides what ships.</p> <p>e.g. Drafting an internal update; option-generation for a campaign; research synthesis for a brief; first-pass Q&amp;A prep.</p>

< Lower consequence | Higher consequence >

Need for human judgment increases with depth of zone

### How to use the map

- Run a 30-minute audit per quarter: list current AI use cases and place each one in a zone. Anything mismatched (AUTOMATE work being done by humans, ESCALATE work being done quietly by AI) gets reassigned.
- Use the map in any new tool conversation: the salesperson's slide ends; the team places the use case on the map. If a tool is positioned as automation for an ESCALATE task, decline.
- Treat zone assignments as living. A merger announcement that was AUGMENT during planning becomes ESCALATE the morning of the leak.

## SECTION 05

# Human checkpoints.

The minimum points where a named human must look, sign, or refuse. Skip these and the operating model fails — quietly at first, then publicly.

## Seven non-negotiable checkpoints

These are the minimum human-in-the-loop gates for high-consequence communications. They are not optional, and they are not the same person every time. Each gate names who reviews, what they look for, and what they may refuse.

Domain	Reviewer	Looking for	Right to refuse
<b>Executive voice</b>	Named comms lead + the executive themselves on consequential statements.	Drift from prior positions, voice mismatch, claims unsupported by record.	Any line not recognizable as the named executive's own.
<b>Crisis</b>	Crisis lead, legal counsel, CEO or delegated principal.	Factual accuracy, sequencing, accountability language, victim sensitivity.	AI-drafted apologies; speculative cause attribution; premature commitments.
<b>Legal &amp; regulatory</b>	General Counsel or designated regulatory counsel.	Disclosure obligations, jurisdictional differences, privilege risk.	Anything that creates undisclosed liability or contradicts a filing.
<b>Employee comms</b>	HR / People lead with named comms partner.	Tone, sequencing relative to managers, sensitivity to affected groups.	AI-only delivery of news that affects livelihoods or status.
<b>Investor / financial</b>	CFO or IR lead with comms and legal.	Material non-public information, forward-looking statement framing, consistency with disclosures.	Any number, projection, or characterisation not sourced from approved materials.
<b>Public statements</b>	Comms lead with the named spokesperson.	Position, evidence, reversibility, second-order effects.	Statements that cannot be defended on the record by a named human.
<b>Sensitive stakeholders</b>	Relationship owner (regulator, journalist, partner, community).	History of the relationship, prior commitments, relational tone.	Outreach drafted without input from the relationship owner.

Where a checkpoint feels burdensome, ask whether the workflow has placed AI in the wrong zone. The fix is almost never to weaken the checkpoint.

## SECTION 06

# Operating workflows.

Four practical workflows the team can adopt this quarter: executive briefing, issue and crisis preparation, content system, stakeholder intelligence loop.

# How to read these workflows

Each workflow lists inputs the system relies on, AI tasks that run automatically or on demand, human decisions that cannot be delegated, the output the workflow produces, and the governance check that must accompany it. Adapt the names to your organization; do not skip the structure.

01

## Executive briefing

<b>PURPOSE</b>	Prepare a senior leader for a meeting, interview, or internal moment in less time, with more rigour.
<b>INPUTS</b>	Calendar item; counterpart profile; prior public statements; current strategic priorities; recent stakeholder signals; internal rumour register.
<b>AI TASKS</b>	Draft a one-page briefing covering counterpart, agenda, likely questions, prior commitments, sensitive topics, suggested framings, and three risks. Maintain version history.
<b>HUMAN DECISIONS</b>	Strategist sets intent ('what does success look like in this room?'). Comms lead approves framings and refuses any not in the executive's voice. Executive owns final positions.
<b>OUTPUT</b>	One-page briefing, three-line top-of-mind summary, and a list of 'do not say' items.
<b>GOVERNANCE CHECK</b>	Source list with timestamps; redaction of any privileged or MNPI material; logged refusals.

02

## Issue & crisis preparation

<b>PURPOSE</b>	Move faster in the first 60 minutes without losing factual accuracy or accountability.
<b>INPUTS</b>	Issue intake template; signal monitoring; precedent library; stakeholder map; legal and regulatory constraints.
<b>AI TASKS</b>	Match the issue to precedents, draft holding statement options, produce stakeholder-by-stakeholder cuts, simulate Q&A, monitor real-time response.
<b>HUMAN DECISIONS</b>	Crisis lead confirms facts. Legal clears statements. CEO or delegated principal owns positions. Comms lead owns tone and sequencing.
<b>OUTPUT</b>	Tiered response pack: holding line, full statement, internal note, regulator notification, FAQ.
<b>GOVERNANCE CHECK</b>	Time-stamped decision log; named owners on each artefact; post-incident review within five working days.

03

## Content system

<b>PURPOSE</b>	Run a high-volume, on-brand content engine with one voice and full traceability.
<b>INPUTS</b>	Strategy narrative; voice profile; brand glossary; channel calendar; performance data.
<b>AI TASKS</b>	Generate drafts, channel adaptations, translations, image and video variants, accessibility versions; flag voice drift against profile.
<b>HUMAN DECISIONS</b>	Editor approves anything externally facing. Subject experts verify any factual claim. Brand lead arbitrates voice. Performance lead decides what gets cut.
<b>OUTPUT</b>	Approved content packages by channel, with a single source of record for the underlying message.
<b>GOVERNANCE CHECK</b>	Voice-drift score logged; sources cited per asset; archived approvals attached to each piece.

04

## Stakeholder intelligence loop

<b>PURPOSE</b>	Replace dashboard theatre with a usable read on how stakeholders are actually moving.
<b>INPUTS</b>	Owned channels, earned coverage, social signal, policy and regulator chatter, employee listening, customer comms.
<b>AI TASKS</b>	Continuous classification, sentiment and topic shifts, narrative alignment scoring, weak-signal detection, weekly synthesis.
<b>HUMAN DECISIONS</b>	Analyst interprets shifts; comms lead decides what is worth a response; executives decide what to act on.
<b>OUTPUT</b>	Weekly intelligence note (one page), monthly narrative review, ad-hoc alerts when thresholds are crossed.
<b>GOVERNANCE CHECK</b>	Source list and methodology; documented thresholds; named owner for any action triggered.

## SECTION 07

# Governance kit.

Seven policy primitives that protect voice, accountability, and trust without grinding the work to a halt.

# Seven primitives

These are the building blocks of a governance regime that scales. They are deliberately small. Heavy regimes are ignored; this set, written and enforced, is enough for most organizations to operate safely.

<b>Data boundary</b>	Define which inputs may be sent to which models, and which must never leave the perimeter. Cover personal data, MNPI, privileged material, security-sensitive information, and third-party confidential content.
<b>Source hierarchy</b>	Rank source types from authoritative (filings, official policies, signed-off briefs) to provisional (model output, ungated web). AI outputs cannot be cited as sources for AI outputs.
<b>Approval matrix</b>	Match content type to required approvers and minimum review window. Crisis exceptions are written, not improvised.
<b>Audit trail</b>	Every consequential AI-assisted output keeps a record: prompt or brief, model and version, sources, human approvers, time stamps. Retained at least as long as the content itself.
<b>Model and tool registry</b>	A maintained list of approved models and tools with intended use, data boundary, owner, last review date, and known limitations. Anything not on the list is not in production.
<b>Prompt and brief standard</b>	All AI-generated work is preceded by a written brief: objective, audience, constraints, sources, definition of done. No standing 'open prompts' for consequential work.
<b>Escalation rules</b>	Pre-agreed triggers that move work from AUGMENT to ESCALATE: legal exposure, executive voice, stakeholder sensitivity, factual uncertainty, novel use case.

These primitives are written, owned, and reviewed quarterly. They sit alongside the organization's existing data, security, legal and HR policies — they do not replace them.

## SECTION 08

# Templates & checklists.

Ready-to-adapt artefacts for use this week: AI brief, human review, risk triage, and a 30/60/90 day implementation roadmap.

## AI brief template

Use this brief before initiating any AI-assisted work that will go anywhere near a stakeholder. If a section cannot be filled in, the work is not yet ready for the model.

<b>Objective</b>	What outcome are we trying to produce, in one sentence.
<b>Audience</b>	Primary and secondary readers, what they already know, what they care about.
<b>Voice</b>	Whose voice carries this, and the three tonal anchors (e.g. plainspoken, decisive, warm).
<b>Constraints</b>	Length, channel, format, must-include and must-not-say lines, regulatory limits.
<b>Sources</b>	Specific approved sources by name. Anything outside this list requires explicit approval.
<b>Risks</b>	Top three things that could go wrong, and how this brief reduces them.
<b>Definition of done</b>	What the finished output must contain to be approved.
<b>Owner &amp; reviewers</b>	Named human owner. Named reviewer(s). Named approver.

## Human review checklist

Run this checklist on any AI-assisted output before approval. Anything uncertain promotes the item to a higher checkpoint.

<input type="checkbox"/>	Voice. Does this read like the named human or institution, or like a generic AI?
<input type="checkbox"/>	Facts. Is every factual claim sourced to an approved record?
<input type="checkbox"/>	Position. Is the position consistent with prior public commitments and current strategy?
<input type="checkbox"/>	Audience. Does the language meet the audience where they actually are?
<input type="checkbox"/>	Reversibility. If we are wrong, how hard is this to walk back?
<input type="checkbox"/>	Second-order effects. Who else hears this, and what do they conclude?
<input type="checkbox"/>	Refusal. Is there anything here we should refuse to say at all?
<input type="checkbox"/>	Accountability. Is there a named human willing to defend this on the record?

## Risk triage checklist

Use this when a piece of work could plausibly go to ESCALATE. Score each row low / medium / high. Two or more 'high' scores routes the work to a senior reviewer before any external move.

Domain	Low	Medium	High
<b>Regulatory and legal exposure</b>	No regulator interest; routine.	Adjacent to filings or live matters.	Direct disclosure or litigation risk.
<b>Executive or board exposure</b>	No leader is named or quoted.	Leader is referenced in passing.	Statement carries an executive signature.
<b>Stakeholder sensitivity</b>	No protected or vulnerable group affected.	Affects an employee group, partner or community.	Affects a sensitive group, victim, or relationship under strain.
<b>Factual certainty</b>	All facts traceable to approved sources.	Some inference required; sources mixed.	Active uncertainty or contested facts.
<b>Reversibility</b>	Easily corrected or withdrawn.	Correction possible but with cost.	Hard to walk back without further damage.
<b>Novelty</b>	Routine task with strong precedent.	Variation on a known pattern.	First time the team has done this.

## 30 / 60 / 90 day implementation roadmap

An order of operations for a communications team adopting this toolkit. Tighten or extend the windows to suit your context.

DAY 0–30	DAY 31–60	DAY 61–90
<p><b>Set the foundation</b></p> <ul style="list-style-type: none"> <li>• Run a workshop on the Human × AI principle and the four-layer OS with the senior team.</li> <li>• Inventory existing AI use cases and place each on the work allocation map.</li> <li>• Publish the seven human checkpoints; assign named reviewers by domain.</li> <li>• Stand up the model and tool registry; remove anything not approved.</li> </ul>	<p><b>Operationalize</b></p> <ul style="list-style-type: none"> <li>• Pilot the four operating workflows with one team each; capture friction.</li> <li>• Adopt the AI brief template and human review checklist as standard.</li> <li>• Establish the audit trail standard; back-fill for the past quarter.</li> <li>• Train comms, legal, HR, and IR partners on the escalation rules.</li> </ul>	<p><b>Scale and review</b></p> <ul style="list-style-type: none"> <li>• Roll the workflows out across the function; retire workarounds.</li> <li>• Run the first quarterly governance review; update the registry and primitives.</li> <li>• Publish a one-page internal standard so the operating model is portable.</li> <li>• Set the next quarter's priorities for AUGMENT to ESCALATE migration.</li> </ul>

# The objective is not to remove people from communications. It is to put human judgment back where it matters most.

Intradiegetic builds the operating layer through which organizations sense change, align people, and act at speed.

THINK

SPEAK

DECIDE

EXIST

[intradiegetic.com](https://intradiegetic.com) · [intradiegetic.com/reports](https://intradiegetic.com/reports)