

Planned Career Break: Family Caregiving

Applied Generative AI for Internal Communications: A Practitioner's Self-Directed Evaluation

When I led Chevron's digital-fluency program, I watched a workforce of 700+ legal professionals move from caution to responsible AI adoption; and I helped the global Law Function enterprise get ahead of ChatGPT with sanctioned tooling and usage guidelines. That work convinced me that generative AI would reshape the internal communications discipline.

So, when I stepped back from contract work to assist my mother after my father's passing, I didn't stop my professional development. I've spent my time stress-testing ChatGPT, Claude, Copilot, Gemini, and NotebookLM against the real deliverables of internal communications teams—leadership communications, change messaging, editorial QA, long-document aggregation, governance synthesis—to answer one practical question: which model, for which task, with which prompts? The result: a working playbook for faster drafting, consistent executive voice, and responsible, governed adoption.

See my white paper below: *From Prompts to Playbooks: Building Repeatable AI Applications for Research, Content Development, and Editorial Productivity in Internal Communications*

Self-directed research

I wanted to answer the crucial questions internal communications professionals actually have about AI adoption:

- Which model produces the cleanest **first draft** for a given content type, and which produces the cleanest final one?
- Which holds a **leader's voice** most consistently across a long series of messages?
- Which can be trusted with **long, sensitive source material**—governance briefs, policy documents, board materials—without losing the thread or inventing detail (hallucinating)?
- Where does each model **break**, and what prompt patterns reliably prevent it?
- Which gains the most from **tool integration** (e.g., Microsoft 365, Google Workspace) versus the raw capability of the model itself?

What I bring to a team

I'm an internal communications leader who has done the hands-on work of evaluating applied generative AI against the real deliverables of the discipline, and I offer the expertise to translate that into faster production, stronger governance, and a responsible adoption path. I've been on the enterprise side of getting ahead of this technology once already. I'd do it again, with the benefit of structured testing behind me.

Let's talk about what that looks like for your team.

From Prompts to Playbooks

Building Repeatable AI Applications for Research, Content Development, and Editorial Productivity in Internal Communications

A practical framework for internal communications leaders moving from ad hoc experimentation to a governed, scalable operating model

A White Paper

June 2026

Peter P. Speliopoulos
peter.p.speliopoulos@gmail.com

Contents

- Contents 2
- Executive Summary 3
- Why Internal Communications Is Positioned to Win..... 3
 - The work is structured beneath the surface 3
 - The inputs already exist as assets 4
 - The bottleneck is throughput, not ideas 4
- A Framework for Repeatable Applications 4
 - Three levels of maturity 5
- Pillar 1: Research..... 5
 - Synthesizing employee listening data 5
 - Audience and stakeholder analysis..... 5
 - Message pre-mortems and testing..... 6
 - Benchmarking and background briefs 6
- Pillar 2: Content Development 6
 - The create-once, adapt-everywhere application..... 6
 - Narrative architecture and message houses 7
 - Register and audience adaptation 7
 - FAQs, Q&As, and anticipated objections 7
 - Localization and accessibility starters..... 7
- Pillar 3: Editorial Productivity 7
 - Style-guide and consistency enforcement 7
 - Length and format transformation 8
 - Headlines, subject lines, and variants..... 8
 - Plain-language and accessibility review 8
 - Summarization and repurposing 8
- The Operating Model: Making It Stick..... 8
 - 1. A managed application library 8
 - 2. Grounding as a discipline 9
 - 3. Clear ownership and human-in-the-loop review..... 9
 - 4. Governance, disclosure, and risk 9
 - Measuring what matters 9
- A Practical Roadmap 9
 - Start where the friction is highest..... 10
- Conclusion 10

Executive Summary

The internal communications discipline has quietly become one of the most AI-ready functions in the modern enterprise. The work is language-intensive, pattern-rich, and repetitive in structure, even when it varies in subject: a change announcement, a leadership message, a town-hall script, and a benefits explainer differ in content but share a scheme. That underlying regularity is precisely what large language models exploit well.

Yet most communications teams still use AI the way an individual uses a search engine: one prompt at a time, with results that depend on the skill and approach of the person who happens to be at the keyboard. The value created evaporates the moment the browser tab closes. This paper argues that the real opportunity is not better prompting but **repeatable applications**, durable, governed, reusable assets that encode a team's standards, voice, and judgment so that quality becomes a property of the system rather than the individual.

Presented here is a model for designing these applications across the three areas where internal communications spends most of its effort—**research, content development, and editorial productivity**—and a practical operating model for building, governing, and measuring them. The goal is a communications function that produces work faster without sounding like a machine produced it, and that gets better every time it is used.

The central shift

Stop treating AI as a clever assistant you re-brief each morning. Start treating it as a library of repeatable applications, each one a documented, owned, reusable asset that captures how your team does a specific job well.

Why Internal Communications Is Positioned to Win

Three characteristics make internal communications unusually well suited to systematic AI adoption.

The work is structured beneath the surface

While communications outputs are tailored to specific audiences and channels, they follow recognizable forms. A program launch almost always needs a leader email, an intranet article, a manager toolkit, a set of FAQs, and a short channel teaser. Acknowledging that you are not facing a fresh, custom set of outputs each time but rather the same predictable package over and over, you can build one repeatable application that generates the entire package from a single brief, instead of treating each launch as a blank-page exercise and rebuilding the five pieces by hand. The recurrence is what justifies the investment in building the application, because the cost of designing it once is repaid across every future instance of that same bundle.

The inputs already exist as assets

Most teams already maintain the raw material that makes AI output good: a style guide, a tone-of-voice document, a messaging house (umbrella statement, core themes, proof points), brand principles, and an archive of approved content. These are exactly the materials that, when supplied to a model as grounding, move output from generic to on brand. The asset is sitting in a shared drive or, preferably, a digital asset library, waiting to be operationalized.

The bottleneck is throughput, not ideas

Communications teams are rarely short of things to say; they are short of the hours to draft, adapt, localize, and proof everything at the standard their stakeholders expect. AI's comparative advantage—first drafts, variants, summaries, and adaptations—maps directly to where the team loses the most time.

A Framework for Repeatable Applications

A repeatable application is more than a saved prompt. It is a small, documented system that anyone on the team can run to get a reliable result. Every well-designed application shares the same six components.

Component	What it defines	Example
Trigger	The recurring situation that calls the application into use	A new policy needs to be announced across channels
Inputs	The brief and source material the user supplies each time	Policy summary, effective date, audience, key dates
Grounding	The fixed reference assets the application always consults	Style guide, tone-of-voice document, messaging house, past examples
Instruction set	The encoded logic, the reusable prompt, role, and rules	Structure, mandatory sections, reading level, what to avoid
Human review	The defined checkpoint where a person owns the output	Communications lead edits and approves before anything ships
Feedback loop	How edits and outcomes improve the application over time	Editor's changes are folded back into the instruction set

The discipline of writing these six components down is what converts a lucky prompt into an asset the whole team can rely on. It also makes the application portable: it can move from a chat interface today to a custom assistant, a project space, or an integrated workflow tomorrow without losing the institutional knowledge baked into it.

Design principle: capture the judgment, not just the task

The most valuable thing communicators do is exercise judgment, what to lead with, what to leave out, how to pitch the tone. A well-constructed application encodes that judgment in its instruction set and grounding so the whole team inherits it.

Three levels of maturity

Teams typically progress through three stages. Naming them helps a function position itself and plan the next move.

Level	Characteristic	What it produces
1. Ad hoc	Individuals prompt freely; nothing is shared or documented	Inconsistent quality; value lost when people leave
2. Templated	Shared prompt library with agreed upon structures and grounding	Consistent first drafts; faster onboarding; repeatable results
3. Integrated	Applications connect to systems, data, and review workflows	End-to-end acceleration; measurable cycle-time gains

Most teams should aim deliberately for Level 2 and treat it as the foundation. It captures the majority of the value, requires no specialist tooling, and is the prerequisite for any responsible move to Level 3.

Pillar 1: Research

Research is the work of understanding before communicating: who the audience is, what they already feel, what has been said before, and what good looks like elsewhere. It is also the work most often compressed under deadline pressure. Repeatable applications let teams do more of it faster, without cutting the corner that usually gets cut.

Synthesizing employee listening data

Engagement surveys, pulse checks, and open-text feedback generate volumes of qualitative data that few teams have time to read in full. A repeatable application, grounded in the survey instrument and the organization's priorities, can cluster verbatim comments into themes, surface representative quotes, and flag shifts since the last cycle. The communicator still interprets and decides; the application removes the manual labor of getting to the point where interpretation can begin.

Audience and stakeholder analysis

A reusable audience-profiling application takes a population—frontline staff, senior engineers, a newly acquired business unit—and produces a structured profile: likely concerns, information needs, preferred channels, and language sensitivities. Because the application always works to the same template, profiles become comparable across audiences and reusable across campaigns.

Message pre-mortems and testing

Before a sensitive message ships, an application can roleplay the audience and surface how a draft might be misread, what questions it will provoke, and where trust might fracture. This is not a substitute for real employee testing, but it catches the obvious failure modes early and cheaply, when they are still easy to fix.

Benchmarking and background briefs

For leadership communications, a grounded application can assemble a tight background brief—context, prior commitments made on a topic, and relevant external developments—in a consistent format that a communications lead can verify and refine. The repeatability matters: every brief looks the same, so leaders learn to read them quickly.

Guardrail for research applications

AI-synthesized research is a starting point for human judgment, never a verdict. Treat themes, profiles, and predictions as hypotheses to validate, and never let a model’s summary stand in for reading what real employees actually said on a high-stakes issue.

Pillar 2: Content Development

Content development is where AI’s value is most visible and most often misused. Used badly, it produces fluent, generic copy that erodes the credibility internal communications spends years building. Used as a set of well-grounded repeatable applications, it accelerates drafting while protecting voice.

The create-once, adapt-everywhere application

This is the single highest-value application most teams can build. From one approved source—a brief, a leadership message, or a finished article—the application generates the full channel bundle: intranet copy, an email version, a digital-signage teaser, a manager talking-points sheet, and a short message on a collaboration platform. Each output respects the length, register, and conventions of its channel because those conventions are encoded in the application’s instruction set.

From one source...	...the application produces	Calibrated for
Approved announcement	Intranet article	Depth, links, scannability
Approved announcement	Email to all employees	Concision, single clear action
Approved announcement	Manager toolkit	Cascade talking points, anticipated questions
Approved announcement	Channel teaser	One sentence, drives click-through

Narrative architecture and message houses

An application can draft a first-pass messaging house—from a brief, core narrative, supporting pillars, proof points, and audience-specific framing—in the team’s standard format. The communicator refines it but starts from structure rather than a blank page. Because the output format is fixed, message houses become consistent artifacts the whole team can build on.

Register and audience adaptation

The same content must often reach entry-level hires, cross-functional leadership teams, a warehouse floor, and a research division. A reusable adaptation application shifts register, vocabulary, and assumed context for a named audience while preserving the core facts and intent, a task that is mechanical enough to delegate and sensitive enough to always review.

FAQs, Q&As, and anticipated objections

From a policy or announcement, an application can generate a first draft of the questions employees will actually ask, including the uncomfortable ones, and structured answer stubs. This turns one of the most tedious and most skipped tasks into a 10-minute review-and-refine exercise.

Localization and accessibility starters

For global teams, applications can produce first-draft translations and culturally adapted versions for human linguists to finalize and can rewrite dense copy into plain language at a target reading level. Neither replaces specialist review; both remove the slow first step.

Voice is non-negotiable

Ground every content application in your tone-of-voice document and a handful of genuinely on-brand examples. Without grounding, a model defaults to a flat, corporate average that any reader can now recognize as machine-written, the fastest way to lose internal trust.

Pillar 3: Editorial Productivity

Editorial work—the polishing, checking, tightening, and reshaping that turns a draft into something publishable—consumes a large and largely invisible share of a communications team’s time. It is also highly rule-based, which makes it ideal territory for repeatable applications.

Style-guide and consistency enforcement

An application grounded in the corporate style guide can check a draft against the team’s published rules, terminology, capitalization, inclusive-language standards, formatting conventions, and return a marked-up list of issues. Unlike a generic grammar tool, it enforces your standards, not a vendor’s defaults.

Length and format transformation

Communicators constantly need the same content at different lengths: a 600-word article condensed to a 90-word email message, a 40-word teaser, and a one-line push notification. A transformation application produces all variants at once, each preserving the essential message, ready for a quick human check.

Headlines, subject lines, and variants

Open rates and click-throughs turn on small wording choices. An application can generate a structured set of subject-line and headline options against defined criteria, clarity, length, and tone, giving the editor real choices to test rather than a single take to accept or reject.

Plain-language and accessibility review

An application can assess reading level, flag jargon and acronyms, and propose plainer alternatives, an accessibility discipline that is easy to endorse and hard to sustain manually under deadline. Building it into the editorial workflow makes inclusion a default rather than an aspiration.

Summarization and repurposing

Long-form assets—a strategy document, a town-hall transcript, a detailed policy—can be turned into the derivative pieces a campaign needs: a summary, a set of key messages, a social-style internal post, and a manager briefing. The application does the extraction; the editor does the judgment.

Editorial task	Repeatable application	Human owns
Proofing to house style	Style-guide checker grounded in the corporate guide	Final editorial decisions
Resizing content	Multi-length transformer	Choice of which version ships
Generating options	Headline and subject-line generator	Selection and testing strategy
Plain-language pass	Reading-level and jargon reviewer	Voice and nuance
Repurposing long-form	Derivative-asset extractor	Accuracy and emphasis

The Operating Model: Making It Stick

Applications create lasting value only inside an operating model that maintains them. Four elements turn a collection of clever prompts into a durable capability.

1. A managed application library

Maintain a single, shared, version-controlled home for every approved application, with its six components documented and a named owner. This is the difference between Level 1 and Level 2 maturity. The library is the asset; individual prompts are merely its contents.

2. Grounding as a discipline

The quality of every application depends on what you feed it. Curate the grounding set, style guide, tone of voice, messaging house, and a small collection of genuinely exemplary past work, and keep it current. Stale or generic grounding produces stale or generic output, regardless of how good the instruction set is.

3. Clear ownership and human-in-the-loop review

Every application names who may run it and who must approve its output before anything reaches an employee. AI accelerates the draft; accountability for what ships stays unambiguously human. This is a commitment to both quality control and trust.

4. Governance, disclosure, and risk

Codify the rules of use: what data may and may not be entered into which tools, how confidential or pre-announcement information is handled, when AI involvement is disclosed, and how factual accuracy is verified. Internal communications is often the function that explains AI policy to the wider organization, so its own practice should be exemplary.

A note on trust

The internal communications discipline trades on credibility. The moment employees suspect that they are reading undifferentiated machine output, engagement falls. The entire purpose of grounding, review, and ownership is to ensure AI raises your output's quality and never flattens its humanity.

Measuring what matters

To justify and steer the investment, measure a small set of indicators rather than vanity metrics.

Dimension	Indicator	Why it matters
Efficiency	Cycle time from brief to approved draft	Captures the core throughput gain
Adoption	Share of eligible work run through the library	Reveals whether the capability is real or abstract
Quality	Editor revision depth; stakeholder satisfaction	Guards against speed at the cost of standards
Reach	Output volume per channel without added headcount	Shows capacity created, not just time saved

A Practical Roadmap

A team can move from ad hoc use to a working library in a single quarter without specialist tooling. The sequence below front-loads the highest-value, lowest-risk applications and builds the governance habit early.

Phase	Focus	Build this first
Weeks 1–2: Foundation	Assemble grounding assets; agree on governance basics and the review rule	A documented grounding set and a one-page usage policy
Weeks 3–6: First applications	Build two or three high-value applications and document their six components	Create once, adapt everywhere; FAQ generator; style-guide checker
Weeks 7–9: Library	Stand up the shared library; assign owners; train the team	A version-controlled library every communicator can run
Weeks 10–12: Measure and refine	Capture baseline metrics; fold editor feedback into instruction sets	A simple dashboard and a quarterly review rhythm

Start where the friction is highest

You are looking for the work that has both high frequency (you do it constantly) and high pain (each instance is tedious, slow, or error-prone). A task that is painful but rare is not worth systematizing first; neither is one that is frequent but already painless. The ones worth building first sit where both are high.

The equation is simple. Choose your first applications by multiplying frequency by pain: the tasks done most often that hurt most each time. For nearly every team, that points to the same place: the channel-adaptation bundle (the create-once, adapt-everywhere application) and the FAQ generator (scoring high on frequency because most announcements and policies need an FAQ). Early, visible wins on familiar pain build the credibility needed to take the team to Level 3. You do not need a lengthy diagnostic to find your first build. The math almost always lands in the same place, so begin with those two.

Conclusion

The communications teams that benefit most from AI will not be those that prompt most cleverly. They will be those that treat AI as an opportunity to **systematize their craft**, to write down how they do their best work, and to encode it in applications the whole team can run, govern, and improve.

This reframing changes the nature of the investment in AI. Time spent designing a repeatable application is not time spent on a single deliverable; it is time spent building an asset that pays out across every future deliverable of that type. Research gets deeper because synthesis is faster. Content stays on voice because voice is grounded in the system. Editorial standards rise because they are enforced consistently rather than sporadically.

The function that emerges is not a smaller team doing the same work faster. It is a team with the capacity to do the strategic, relational, and judgment-heavy work that only communicators can do, because the repeatable work now runs on rails they built and own. That is the destination this

paper points toward: **from prompts to playbooks, and from playbooks to a durable communications capability.**