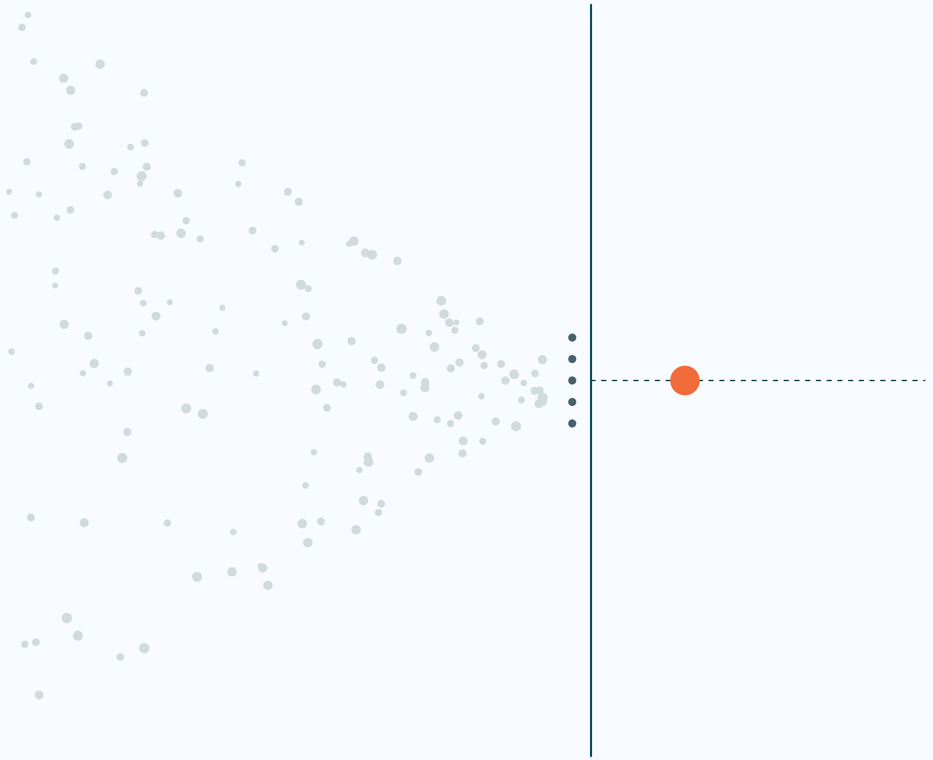
LEADING MARKETING TEAMS IN AN AI-FIRST ERA

Content Engine

The core of the human-in-the-loop marketing OS: a content engine that runs research, production, distribution and analytics with AI, while keeping people in charge.



SEBASTIAN HESSE · MARKETING COMMUNICATIONS & CAMPAIGN ARCHITECT

NOVEMBER 2026

Contents

3 The playbook on one page

PART 1

Why content needs a system

5 The problem

8 The marketing OS as a framework

PART 2

How the content engine is governed

12 Topic planning

16 Standards

19 Approval and ownership

23 AI rules

25 The editorial conference

PART 3

How the content engine works

28 The content engine's working rules

30 The architecture

31 Context

36 Research and pick

38 Production and approval

39 Distribution and posting

40 Analytics and reading

PART 4

How a team makes it work

42 Owners in the team

44 Roles and adoption

46 Tools

48 Pitfalls

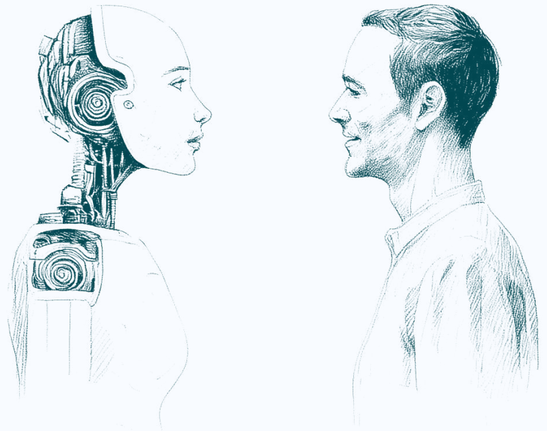
50 The first 90 days

54 Templates

59 References

62 About the author

Strategic content needs governance first and automation second.



Marketing teams of every size publish on more channels than ever, and too little of what they make lands. The cause is rarely effort. It is how the work is organised: disciplines in silos, plans built around channels, new tools without a shared way of working, and everyone with their own AI chatbot. Strategic content starts with an effective team and a solid brief, not with a tool.

This playbook describes a model for organising that work: the marketing operating system (marketing OS). This framework combines a shared workspace and a shared context as its base with a governance that sets three principles for everything built on it. The machine prepares, a named person decides, and every approval has an owner. Every module draws on the same shared context. One module of the marketing OS is the content engine, which this playbook describes in detail. The model starts from one conviction: an AI-first team needs clear governance before it needs more AI.

Governance sets the principles: who decides what, and when. For the content engine, it plans from the company's objectives down to the channel and writes down the standards every piece of content has to meet. It names an owner for each type of risk, hands the final release to the chief on duty, and settles who makes corrections after publication, if needed. It sets clear rules for the use of AI, and it brings all of this together in one editorial conference each week.

The content engine, as the first module, does the preparation: it researches, drafts, prepares each piece for its channels and measures the results. At four checkpoints (pick, approve, post and read) a person takes the decision, the principle this playbook calls • human-in-the-loop. Context, not the AI model, decides the quality of every draft.

In the content engine, these principles become concrete at the checkpoints, in the approval owners and in the context: every checkpoint and every part of the context has a named owner, and some decisions are never automated.

A team can start within about 90 days: governance first, then the shared workspace, then the content engine, one station at a time, with the whole team involved.

For a team lead, the proposed marketing OS offers a way to give a team direction and alignment, clear decision rights and fewer bottlenecks. For a practitioner, it is a blueprint that can be rebuilt and expanded with a variety of tools.

GLOSSARY

• HUMAN-IN-THE-LOOP

A principle from control engineering: a person stays inside an automated process and decides at fixed points. In this playbook: the machine prepares, a named person decides (page 7).

CHIEF ON DUTY

A role taken from newspaper newsrooms: the person who runs the weekly editorial conference, sets priorities and gives the final release before anything goes online (pages 13, 20).

EDITORIAL CONFERENCE

The team's weekly meeting, modelled on a newspaper's editorial conference. It looks back at last week's work and sets topics, priorities and owners for the week ahead (pages 25–26).

SIDE NOTE

Appendix A provides the templates to begin with. Every claim rests on research or on practice, with the sources in the footnotes.

The playbook follows the same order: why content needs a system, the governance of the content engine, the content engine at work, and how a team runs it. Readers in a hurry can follow the action titles alone.



PART 1

Why content needs a system

5 The problem

8 The marketing OS as a framework

Silos produce output that rarely works together.



Working in silos

More channels, more formats, but rarely more people. Marketing teams are at full capacity, yet the results often fall short of the effort. The bottleneck is usually not the team's skills or motivation but the way the work is organised.

In many teams, marketing, PR, social media, events and design plan side by side. Each discipline keeps its own plan and its own channels. This is working in silos. The question in the room becomes "What could we post on this channel?" instead of "What do we have to say, and to whom?" The message breaks apart into channel messages, and the job of coordinating the chaos stays with a few reliable people.[1]

I have often seen this in my own teams: five disciplines, five plans, one audience.

Why silos form

Silos are not a sign of a bad team. The journalist and anthropologist Gillian Tett sees them as an inescapable and necessary consequence of organisational scale and complexity. Specialists need room to do good work. But silos also harden over time. In Tett's words, "people become very incentivized and invested to defend existing borders".[2] Size makes it worse, not better: every additional person and channel adds handovers, and every handover is a chance to lose the point.

Why structure and tools alone do not fix it

The obvious fix is structural: put everyone in one room, or draw a new organisation chart. On its own, that is not enough. A meta-analysis in the *Journal of Marketing*, which pooled the results of many earlier studies, found that integration across functions pays off more when it happens at team level and when customer information is shared.[3] The study looked at new products, not content, but the lesson carries over: what connects silos is shared work and shared information at team level, not a new chart or a shared floor alone.

New tools rarely fix this either. A shared tool introduced without a shared way of working tends to be used as a personal one: people carry over the habits they already have.[4] Generative AI adds more silos to every team. Each person opens their own chatbot, writes their own prompts and keeps their own context. Nothing connects the drafts to a plan, a voice or a decision.

And the work starts in the wrong place. Content begins with a free slot in a channel plan or a request from someone senior, not with an analysis of what the audience needs and what the company has to say.

[1] Cross, R., Rebele, R. and Grant, A. (2016) 'Collaborative overload', *Harvard Business Review*, 94(1), pp. 74–79. Available at: https://hbr.org/2016/01/collaborative-overload (Accessed: 2 Nov 2026).

[2] Tett, quoted in Barnett, P. (2015) 'Breaking bad barriers', *strategy+business*, 20 Oct. Available at: https://www.strategy-business.com/article/00382 (Accessed: 2 Nov 2026).

[3] Troy, L.C., Hirunyawipada, T. and Paswan, A.K. (2008) 'Cross-functional integration and new product success: an empirical investigation of the findings', *Journal of Marketing*, 72(6), pp. 132–146. Available at: https://doi.org/10.1509/jmkg.72.6.132 (Accessed: 2 Nov 2026).

[4] Orlikowski, W.J. (1992) *Learning from Notes: organizational issues in groupware implementation*. Working Paper 3428-92. Cambridge, MA: MIT Sloan School of Management. Available at: https://dspace.mit.edu/handle/1721.1/175292 (Accessed: 2 Nov 2026).

Without a clear brief, more output leads to more corrections.



What a missing brief costs

Briefs translate strategic goals into work assignments. When each discipline plans for itself in its silo, each writes its own brief, but there is no brief for the shared goal. And clear briefs are what teams neglect most often, although they are the foundation. In a global study of more than 1,700 marketers and agency staff, 80% of marketers said they write good briefs; only 10% of agencies agreed. Respondents estimated that a third of marketing budgets is wasted on poor briefs and misdirected work.[5]

Project research points the same way. In a survey of nearly 10,000 project professionals, projects that defined their success criteria before work started reached a net project success score of 41, against 20 for projects that did not. The score is the share of projects their stakeholders rate a success minus the share they rate a failure.[6] Without agreed criteria, every piece is judged by whoever sees it last, by their own taste.

AI makes a missing brief more expensive. It produces drafts faster than any team can check them, and without a brief there is nothing to check them against. Outside the tasks AI handles well, people who rely on it make more mistakes, not fewer (page 28).[7] Every unbriefed draft becomes a review, and the reviews pile up with the few people who decide.

In my own team

It showed up in three ways:

- **Duplicate work**, because the briefs upfront were poor.
- **Last-minute requests** that could have been planned.
- **Approval bottlenecks**, mostly for complex content and client case studies waiting for sign-off at the top.

What a clear brief gives a team

Beyond efficiency and effectiveness, a clear brief also gives a team stability. Google's study of its own teams found five dynamics that set effective teams apart.[8] The most important was psychological safety: a shared belief that the team is safe for interpersonal risk-taking.[9] Another was structure and clarity: knowing what is expected and how to get there.[8] In my experience, the two depend on each other. When the brief says what the content must achieve and who decides, people can question it, admit a mistake or turn down a last-minute request without putting their standing at risk.

None of this is a tool problem. It needs a shared way of deciding what to say, who decides and when. The principle behind it is • human-in-the-loop.

[5] Madigan, M. (2021) 'The BetterBriefs Project finds 33% of marketing budgets wasted due to poor briefs', *B&T*, 13 Oct. Available at: https://www.bandt.com.au/the-betterbriefs-project-finds-33-of-marketing-budgets-wasted-due-to-poor-briefs/ (Accessed: 2 Nov 2026).

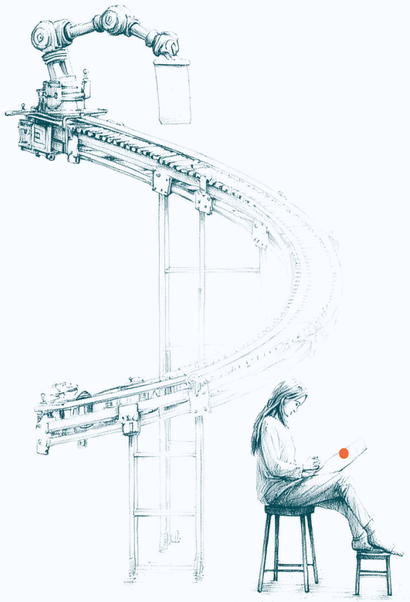
[6] Project Management Institute (2024) *Maximizing project success*. Newtown Square, PA: Project Management Institute, pp. 2, 14, 20. Available at: https://www.pmi.org/-/media/pmi/documents/public/pdf/learning/thought-leadership/project%5Fsuccess%5Freport%5F2024.pdf (Accessed: 2 Nov 2026).

[7] Dell'Acqua, F., McFowland III, E., Mollick, E., Lifshitz, H., Kellogg, K.C., Rajendran, S., Krayer, L., Candelon, F. and Lakhani, K.R. (2026) 'Navigating the jagged technological frontier: field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality', *Organization Science*, 37(2), pp. 403–423. Available at: https://www.hbs.edu/ris/Publication%20Files/dell-acqua-et-al-2026-navigating-the-jagged-technological-frontier_5c589c8c-fbb5-458f-b285-c944746cd717.pdf (Accessed: 2 Nov 2026).

[8] re:Work Editorial Team (2015) 'Understand team effectiveness', *re:Work*. Google. Available at: https://rework.withgoogle.com/intl/en/guides/understand-team-effectiveness (Accessed: 2 Nov 2026).

[9] Edmondson, A. (1999) 'Psychological safety and learning behavior in work teams', *Administrative Science Quarterly*, 44(2), pp. 350–383. Available at: https://dash.harvard.edu/entities/publication/13a7b031-0fdd-45ec-a7e0-2b80e2bc679f (Accessed: 2 Nov 2026).

• Human-in-the-loop closes the gap: named owners decide at fixed points.



Silos, missing briefs and personal chatbots share one pattern: work moves on without anyone able to step in at the right moment. Work in silos often means that each discipline decides for itself, but rarely for the whole team. A missing brief means nobody decided what a piece should achieve. A free slot in a channel plan lets the calendar decide instead of the team. A draft from a personal chatbot, written without a brief, lets the machine decide what to say. The answer this playbook gives is a principle with a long history: • human-in-the-loop.

Origin of the term

The term comes from control engineering. It appears in publications from the 1950s, in research on control systems such as missile fire control, where the human operator was treated as one component of a control loop. From the early 1990s, machine learning took it up for people who train an algorithm. In the current debate about AI, it is mostly used in the sense of human oversight.[10] Google's guidance for AI agents puts it plainly: every agent needs a well-defined human controller, and critical or irreversible actions need explicit human confirmation, "ensuring the user remains in the loop".[11]

What it means in this playbook

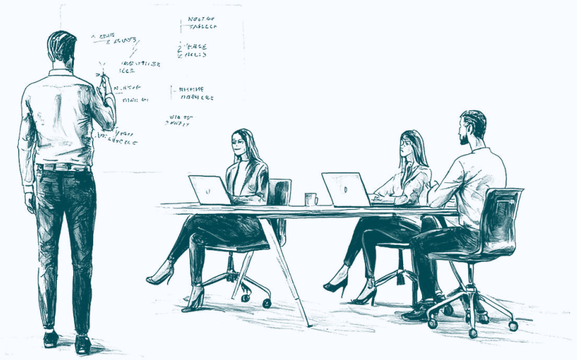
In this playbook, • human-in-the-loop means something specific. The machine prepares, and at fixed points a named person decides. Nothing moves from one step to the next without that decision. Each decision has more than one possible outcome, such as approve, change, send back, reject or park.

That makes • human-in-the-loop more than a safety net for AI. It is a leadership principle: it defines who decides what, when and on what basis. It applies to the team (Part 2) and, once AI is involved, to the machine as well (Part 3). Even without AI, every piece runs through this loop: propose a topic, brief, draft, release, review. At every handover, a named owner decides. That is why governance comes first (page 8), inside the marketing OS that holds team and machine together (page 9).

[10] Anderson, M.M. and Fort, K. (2022) 'Human where? A new scale defining human involvement in technology communities from an ethical standpoint', *International Review of Information Ethics*, 31(1). Available at: https://informationethics.ca/index.php/irie/article/download/477/449/699 (Accessed: 2 Nov 2026).

[11] Díaz, S., Kern, C. and Olive, K. (2025) *Google's approach for secure AI agents: an introduction*. Google, p. 13. Available at: https://services.google.com/fh/files/misc/googles_approach_for_secure_ai_agents.pdf (Accessed: 2 Nov 2026).

An AI-first team needs clear governance.



Research, drafting, adapting content to channels and reading results: in all of these tasks, AI saves time and extends what a team can do. Many companies now call themselves AI-first.

The phrase became widely known in 2016, when Google's chief executive, Sundar Pichai, wrote: "We will move from mobile first to an AI first world."[12] A former chief of research at Gartner describes its common use: an AI-first company prioritises "the use of AI to accomplish anything, rather than relying on established processes, systems, or tools". The same source adds that AI-first does not mean AI-only.[13]

I agree with both points. A team that does not think AI into its processes will fall behind. And AI-first does not mean AI-only. But the term is often read as "fewer people". This playbook reads it differently.

Being AI-first means considering AI for every process. It does not start with AI, but with clear governance. Without it, teams fall back into silos: everyone uses their own chatbot, and the drafts rest on nothing the team has agreed.

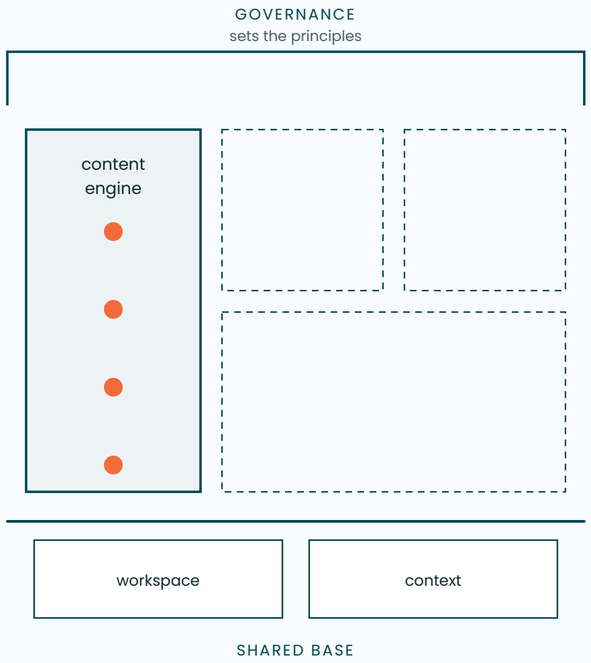
I have often seen this in my own teams. Neither project management tools nor individual AI chatbots solved the problem. What was missing was governance: a defined way of working that settles who decides what, and when.

So every task raises one question: should the machine prepare it, should a person decide it, or both? Governance answers that question before any AI is switched on. At every decision point, the answer is • human-in-the-loop (page 7): the machine prepares, a named person decides. That is why this playbook starts with governance. Automation comes last.

[12] Wheatley, M. (2016) 'Google's Sundar Pichai eyes a shift from "mobile first" to "AI first"', *SiliconANGLE*, 29 Apr. Available at: https://siliconangle.com/2016/04/29/googles-sundar-pichai-eyes-a-shift-from-mobile-first-to-ai-first/ (Accessed: 2 Nov 2026).

[13] Simpson, J. (2024) 'What does it mean to go AI-first?', *Nordic APIs*, 12 Jun. Available at: https://nordicapis.com/what-does-it-mean-to-go-ai-first/ (Accessed: 2 Nov 2026).

The marketing OS has three levels: a shared base, one governance, many modules.



The marketing OS, short for marketing operating system, is the framework this playbook describes. It has three levels.

The base is shared by everyone. It is a common workspace in which the whole team plans and works, and a common context that every piece of work draws on: the personas, the positioning, the tone of voice, the facts and the lessons learned. In my team, the workspace covered all marketing work, not only content production: projects, campaigns, tasks, content and design work, all linked to each other. Predefined content pillars, the standing themes a company covers in depth, set the frame (page 12). Everyone saw the same picture: what a task belonged to, who was working on it and what was still open.

The governance sets the principles for everything built on the base. Three of them hold for every module: the machine prepares, a named person decides; every approval has an owner, chosen by the type of risk; and every module draws on the same shared context.

The modules carry the actual work. Each module handles one kind of work and breaks the principles down into its own rules: who decides what, at which point and by when. The content engine is one of them, and the one this playbook describes in detail (Part 3).

IN PRACTICE

The first thing we automated was not content but transparency. Every team member estimated the hours for their own tasks. Our Notion board then showed this effort for each task in every project and campaign. That kept the workload visible and helped us prioritise and share tasks more evenly across the team.

Before an AI workflow adds output, the team needs to see what it can carry. A visible workload lets people say no with evidence, and lets the editorial conference prioritise content tasks. For a team, this means security rather than control, because a visible workload is a resource. The job demands–resources model describes how resources such as feedback and room to decide buffer the strain of high demands.[14] There will always be ad-hoc requests. But whoever knows the team's capacity can manage them better.

[14] Bakker, A.B. and Demerouti, E. (2007) 'The Job Demands-Resources model: state of the art', *Journal of Managerial Psychology*, 22(3), pp. 309–328. Available at: https://doi.org/10.1108/02683940710733115 (Accessed: 2 Nov 2026).

Principles stay in place, while modules and tools can change.

Why governance stays separate

The governance of the marketing OS is deliberately kept apart from the rules of each module. There are two reasons: stability and discipline.

Stability comes first. Tools change faster than the way a team decides. When an AI model, an automation or a database is replaced, the principles, the approval owners and the content pillars stay the same.

Modules change too. A new module needs no new principles. It breaks the existing ones down into its own rules by answering three questions. Where does it stop, and who decides at that point? Which risks are there, and who owns them? What can it take from the context the first module has built up? That is how the modules interlock.

Discipline is the second reason. Every module answers to the same governance. A module without rules only produces more output. Without a clear brief, there are more drafts and more correction loops. The number of people who decide stays the same.

People and AI also do not get better results just because they work together. A meta-analysis of more than a hundred experiments found that human–AI combinations, on average, performed worse than the better of the two alone. Gains appeared mainly in tasks that involve creating content.[15] What matters is who does what. So the roles have to be designed.

Where the principles become concrete

In the content engine, the principles become concrete in three places.

The first place is the checkpoints. When the content engine needs a decision, it stops and hands over to a person: to pick a topic, approve a draft, post it and read the results. Posting stays entirely in human hands.

The second is the context. The content engine draws on the shared context in the base: the knowledge base and facts, the audience, the voice and key messages, and the lessons learned. Each part of it has an owner (pages 35, 42).

The third is the approval owners. There is one for each type of risk (page 20).

[15] Vaccaro, M., Almaatouq, A. and Malone, T. (2024) 'When combinations of humans and AI are useful: a systematic review and meta-analysis', *Nature Human Behaviour*, 8(12), pp. 2293–2303. Available at: https://doi.org/10.1038/s41562-024-02024-1 (Accessed: 2 Nov 2026).

GLOSSARY

KNOWLEDGE BASE

The team's shared collection of everything worth keeping, such as articles, newsletters, reports and videos, each with a summary, a category and a rating (page 32).

SIDE NOTE

The following parts show the marketing OS through one module, the content engine: Part 2 how it is governed, Part 3 how it works, Part 4 how a team starts it. The governance was used in a team of ten. The full content engine came later; today I run it for my own content.



PART 2

How the content engine is governed

12 Topic planning

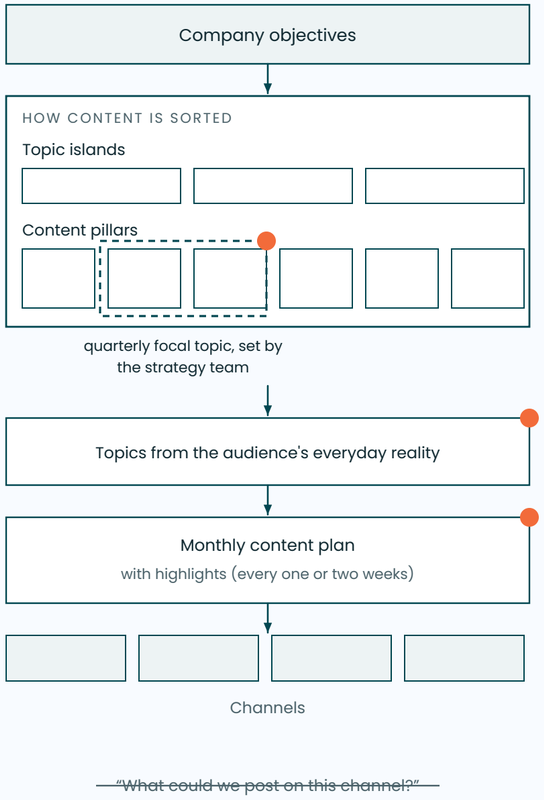
16 Standards

19 Approval and ownership

23 AI rules

25 The editorial conference

Topic planning starts with the company's objectives.



This part sets the rules the team agrees before any AI is switched on. It starts with the question of what to say.

The planning cascade

The planning starts above marketing, with the company's objectives. Below them sit the topic islands: the few broad fields a company works in and wants to be known for. Each island holds several content pillars, the standing themes that go into depth and set the frame for every campaign. A software company for small firms might have data security as one island, with backups, access rights and compliance as pillars.

Once a quarter, based on the objectives, the strategy team (page 13) sets the focus for the next three months: the quarterly focal topic. A focal topic can sit within one pillar or run across several, and it can start product-specific, such as a launch. From it, the team derives topics from the audience's everyday reality, and each piece connects to one of them. Then it splits the quarter into months and weeks, so that it can coordinate the editorial plan and stay consistent. The team decides how often a big highlight comes: a small team does well with one every two weeks, to avoid overstretching; a larger team can carry one every week.

Only at the end does the team decide which channels carry each piece.

Why the order matters

This order puts a few shared goals ahead of each channel's own. Goal-setting research has shown over decades that specific, challenging goals lead to higher performance than vague intentions.[16] The effect holds for teams too. A meta-analysis of group goal studies found that specific, difficult group goals raised group performance considerably. Individual goals inside a group seemed to help only when they were aimed at the person's contribution to the group. Goals aimed at individual success hurt the group's result.[17]

In research terms, this is exactly the problem of working in silos. A channel owner who optimises "their" channel works towards an individual goal. The cascade replaces those goals with a few shared ones that every discipline can see and work towards.

Islands and pillars also help search. Google's own guidance notes that a site organised in a logical way can help search engines and users understand how its pages relate.[18] Topic islands give that structure: one island, several pillars, many pieces that link to each other.

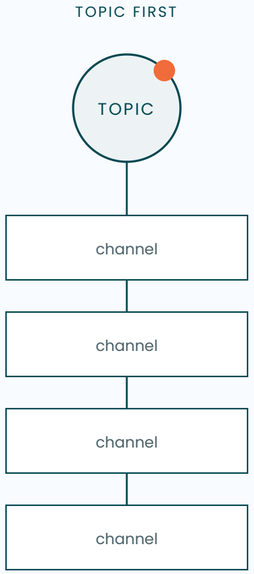
In my team, the question "What could we post on this channel?" stopped coming up.

[16] Locke, E.A. and Latham, G.P. (2002) 'Building a practically useful theory of goal setting and task motivation: a 35-year odyssey', *American Psychologist*, 57(9), pp. 705–717. Available at: https://doi.org/10.1037/0003-066X.57.9.705 (Accessed: 2 Nov 2026).

[17] Kleingeld, A., van Mierlo, H. and Arends, L. (2011) 'The effect of goal setting on group performance: a meta-analysis', *Journal of Applied Psychology*, 96(6), pp. 1289–1304. Available at: https://doi.org/10.1037/a0024315 (Accessed: 2 Nov 2026).

[18] Google (2025) *Search Engine Optimization (SEO) Starter Guide*. Google Search Central, last updated 10 Dec 2025. Available at: https://developers.google.com/search/docs/fundamentals/seo-starter-guide (Accessed: 2 Nov 2026).

Choose the topic first, then the channel.



An order of planning only holds if someone keeps it. The corporate newsroom model gives that job to clear roles and a fixed process.

The corporate newsroom model

Many marketing plans are built around channels. Each channel has a calendar, each calendar has empty slots, and someone has to fill them.

The corporate newsroom model turns this around. It was developed by Prof. Christoph Moss and grew out of his doctoral research on how newspaper editorial teams are organised.[19] Its core rule is short: first the topic, then the channel.[20] A marketing team must first decide what it has to say and to whom. Only then does it decide which channels carry the message, in which format and when.

Roles and process

The model splits the two on purpose, with separate responsibilities for topics and for media.[19] Topic management stays close to the business and knows what matters. Media management knows the channels and adapts the material to each of them. A chief on duty connects the two, sets the priorities and runs the editorial conference. A strategy team at the top steers the topics and has the final say on content in case of conflict.[19]

Behind these roles sits a standard process. A topic is logged, assessed and planned in the editorial conference, turned into a base document, approved, adapted to each channel, published, and finally analysed against the goals set at the start.[21]

How this playbook adapts the model

A small team does not need the full organisation chart. What matters is the principle. Moss puts it in one sentence that holds for every team below corporate scale: setting up a newsroom is not a question of a department's size, but of attitude.[20] A team that cannot fill every box can still think in topics before channels. That is a decision about how the team works, not about how many people it has.

This playbook builds on the model and changes it in four ways. It applies the model to the whole marketing team and its project management, not only to communications. It treats the desks as functions people take on, not as departments. It splits approval by type of risk and gives the final release to the chief on duty (page 20). And it adds what the model does not cover: a content engine in which the machine prepares and a named person decides (Part 3).

[19] Mediamoss (no date) *Der Newsroom beginnt im Kopf*. Available at: https://www.mediamoss.me/das-newsroom-modell (Accessed: 2 Nov 2026).

[20] Moss, C. (no date) 'Corporate Newsroom: Zuerst das Thema, dann der Kanal', *GABAL Magazin*. Available at: https://www.gabal-magazin.de/marketing/corporate-newsroom/ (Accessed: 2 Nov 2026). Own translation of: "Zuerst das Thema, dann der Kanal" and "Der Aufbau ist keine Frage der Größe einer Abteilung, sondern eine Frage der Haltung."

[21] Becker, N. (2026) 'Der Standardprozess im Corporate Newsroom', *Mediamoss-Blog*, 21 May. Available at: https://mediamossblog.com/2026/05/21/der-standardprozess-im-corporate-newsroom/ (Accessed: 2 Nov 2026).

GLOSSARY

SUBJECT MATTER EXPERTS

Specialists inside the company who know a topic in depth, such as engineers, analysts or consultants. In this playbook, they approve facts and complex content (page 20).

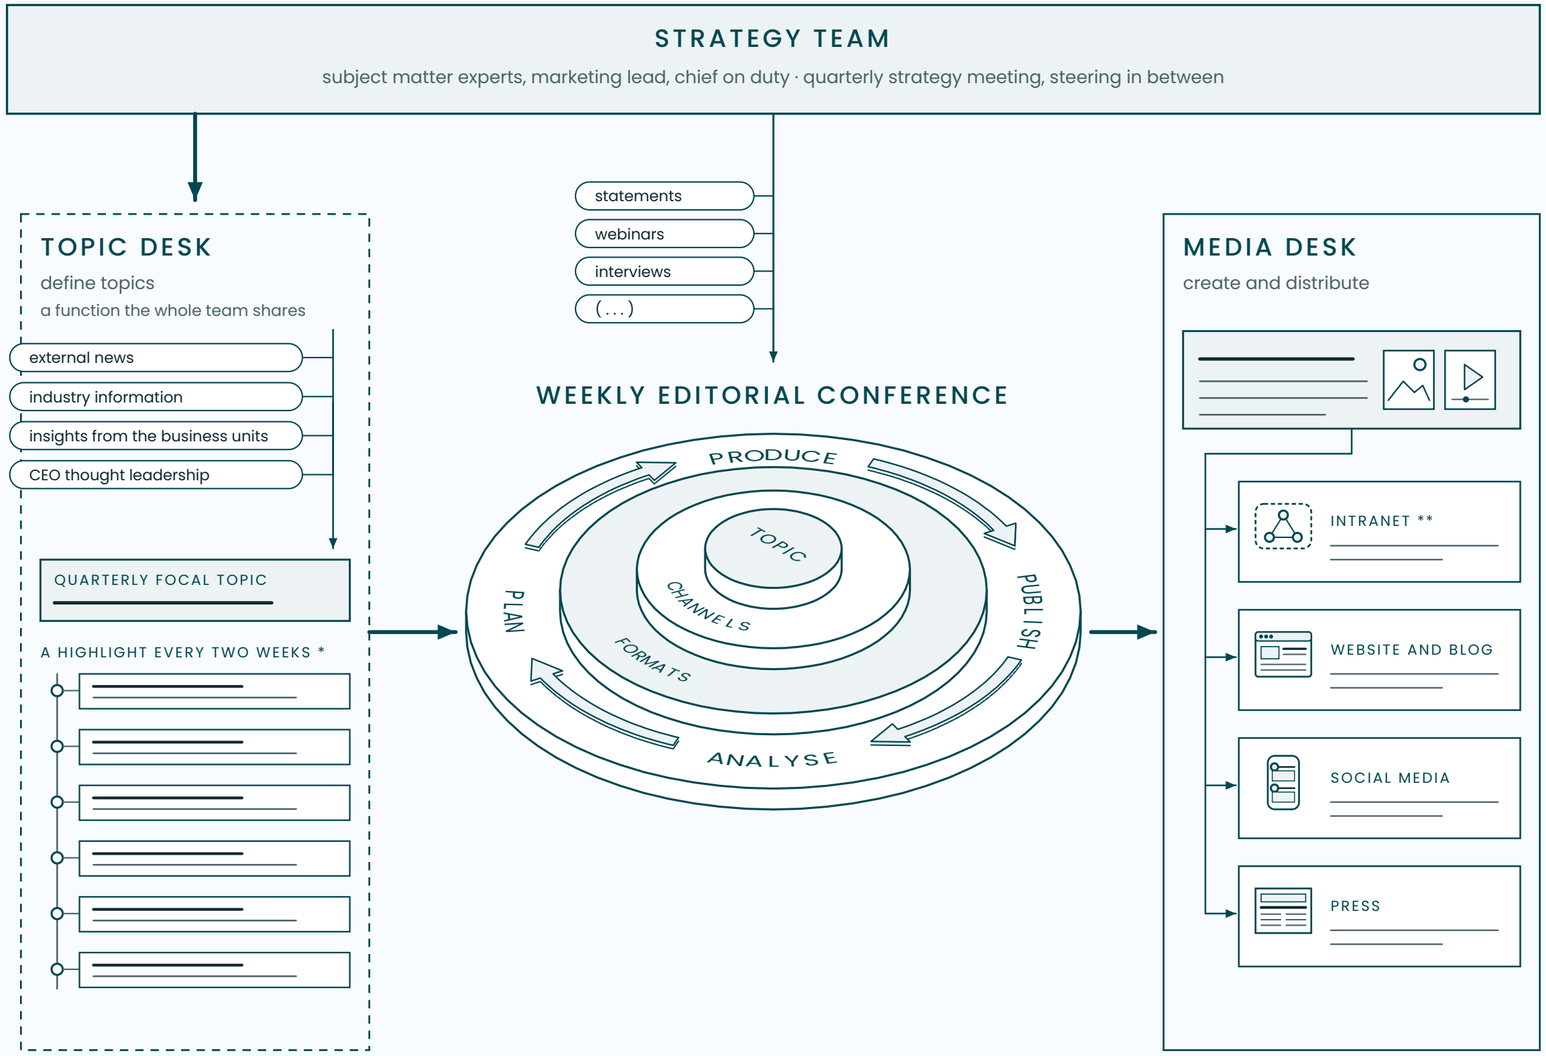
BASE DOCUMENT

In the corporate newsroom model, the core text on a topic: topic management writes it, the business unit approves it, and media management adapts it to each channel.

IN PRACTICE

At Purpose Green, the model covered the whole marketing operation: campaigns, tasks, design work and workload, all planned in one shared workspace for the whole team.

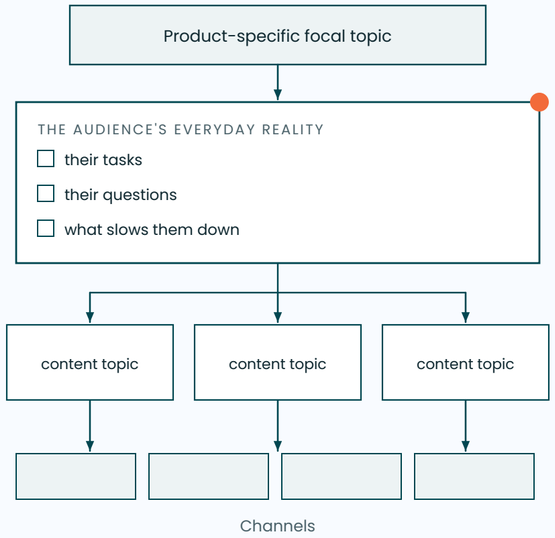
In a team of ten, I was communications lead and chief on duty. The set-up had three levels. A strategy team of subject matter experts, the marketing lead and me met once a quarter and steered in between. There was no dedicated topic desk. Topic first was not a role but a way of thinking the whole team shared. Everyone contributed topics, from external news and industry information, from insights in the business units and from the CEO's thought leadership. From these, the strategy team set a quarterly focal topic, and, depending on strategy and resources, the team derived a new highlight every two weeks. A media desk created the content and distributed it first on the intranet, then on the website and blog, on social media and to the press. The weekly editorial conference connected this topic work with the media desk, in a cycle of plan, produce, publish and analyse.



* depending on strategy and resources

** before publishing externally, always make sure employees (especially sales teams) and close stakeholders are informed

Let the product answer the questions your audience asks.



From product to topic

A company's quarterly focal topic often starts with the product: a launch, a new module, a new service. Taken on its own, such a topic invites product messages: features, modules, prices. Content built that way speaks about the company, not to the reader.

So the team breaks the topic down. It starts from the everyday reality of the people the product is for: what they are responsible for, what slows them down, what they have to decide. Their questions become the topics. Each question becomes its own piece, in the format that answers it best: a guide, a data analysis, a checklist, a live talk. One channel serves as the hub, and the other channels are built around it.

The product stays in the story, but as the answer, not as the headline.

How the team takes it on

A planning model like this works best when it does not arrive from the top. Research on organisational change finds that people who receive information about a change and take part in the decision are more open to it.[22] So present the model to the team before it is rolled out: the problems the team has raised next to the proposed solution. Ask for the team's go first, and only then build the shared workspace.

A planning model that the team has chosen is one the team will use. Adoption and the first month follow this path (pages 45, 51).

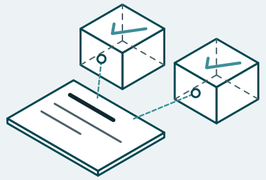
IN PRACTICE

At Purpose Green, one quarter's focal topic was the launch of a new software platform for ESG (environmental, social and governance) and renovation management in real estate. We addressed ESG managers at mid-sized property owners, who wanted one tool for the ESG performance of a whole portfolio. Their questions became our topics. How do I turn a portfolio analysis into a renovation roadmap? What will rising CO_2 prices cost my buildings, and who carries those costs? Which funding programmes make a renovation pay off? Each question became its own piece: a guide, a data analysis of expected CO_2 costs, a funding checklist and a live talk on the economics of renovating existing buildings. The blog was the hub; press, events and other channels were built around it.

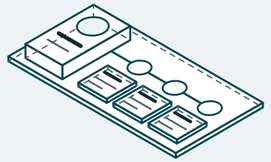
The planning model did not arrive from the top either: the team itself had raised the problems. I presented the corporate newsroom model to the team, problem and solution side by side, and asked for their go first. Only then did I build the shared workspace, with the go-ahead from management.

[22] Wanberg, C.R. and Banas, J.T. (2000) 'Predictors and outcomes of openness to changes in a reorganizing workplace', *Journal of Applied Psychology*, 85(1), pp. 132–142. Available at: https://doi.org/10.1037/0021-9010.85.1.132 (Accessed: 2 Nov 2026).

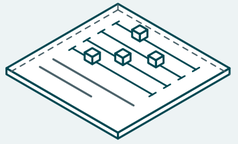
FACTS & SOURCES
"Is it true?"
OWNER SUBJECT MATTER EXPERTS



KEY MESSAGE MAP
"What do we say?"
OWNER COMMUNICATIONS LEAD



TONE OF VOICE
"How do we sound?"
OWNER COMMUNICATIONS LEAD



Create a single source of truth.

Newsrooms do not leave quality to chance. They write down what every story has to meet. The German press code, for example, puts truthfulness first and care second.[23] A marketing team needs the same: standards that are written down, known to everyone and the same for every piece.

Without them, every reviewer checks against their own taste. Approval turns into a debate about wording, and in the end the most senior person decides what is right. With written standards, approval becomes a check: does the piece meet what we agreed?

Three standards cover most of what a content team needs: care with facts and sources, a key message map, and a tone of voice. This page covers the first. The next page covers the other two.

Care starts with the facts. Every claim is checked against a primary source, such as the original study, report or statement, rather than against a summary of it. Figures, dates and names are checked against a list of verified facts that the subject matter experts maintain. The same list records what must not be used, such as outdated figures or claims the company cannot prove.

Sources need a rule as well. In the content engine described in this playbook, a topic is only accepted when at least two independent organisations back it (page 36). The rule is simple enough to apply under time pressure and strict enough to keep vendor blogs from becoming the evidence.

With AI, care matters even more. A model can write a confident sentence around a wrong figure. So a person checks every fact in an AI draft (pages 23–24).

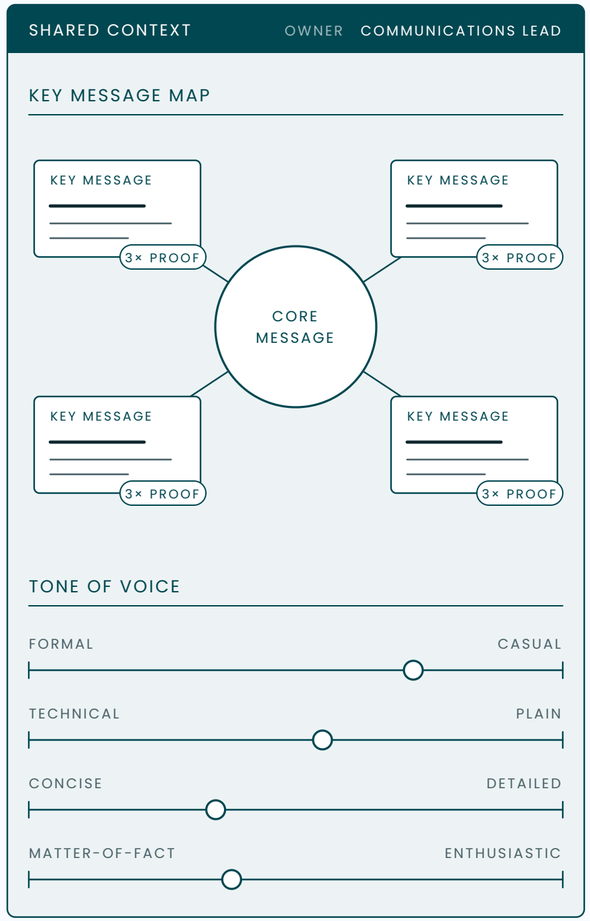
[23] Deutscher Presserat (2025) *Publizistische Grundsätze (Pressekodex)*. Version of 19 Mar 2025. Available at: https://www.presserat.de/pressekodex.html (Accessed: 2 Nov 2026).

GLOSSARY

KEY MESSAGE MAP

One overview of what the company says again and again: a core message at the centre, supporting messages grouped by pillars, each with its proof (page 17, template A4).

Be consistent in your messages.



Key message map

The key message map describes what the company says again and again. At its centre stands the core message: in two lines, what the company is and offers, and how customers benefit from that. Every other message is derived from it.

Around it are supporting key messages, each with the proof points that back it. They are grouped by pillars. A pillar is a theme the company keeps returning to: one of its main offers, or a benefit the offer brings to the customer, to their community or industry, or to society. A simple pillar may need one key message; a complex one may need three or four. In the full map, each pillar has a short heading, a one-sentence claim and its own set of messages (Appendix A4). These message pillars are not the content pillars of topic planning (page 12). Content pillars sort what the team writes about; message pillars hold what the company wants to be remembered for.

For each audience, a short note says which message matters most. The map builds on the positioning brief and the personas (Appendix A1, A2). Appendix A4 gives a template.

Tone of voice

The tone of voice describes how the company says it. It works best as a few clear axes. The Nielsen Norman Group describes four: formal or casual, serious or funny, respectful or irreverent, matter-of-fact or enthusiastic.[24] They are a starting point, not a rule. A team keeps the axes that matter for its content and replaces the others. A company that explains complex products may need "technical or plain" and "concise or detailed" more than "serious or funny", and respect goes without saying. For each axis, the guide sets where the company stands, with dos and don'ts and two or three example sentences. Where a brand guide exists, the voice builds on it. Appendix A3 gives a template.

The communications lead owns both.

[24] Moran, K. (2016a) 'The four dimensions of tone of voice', *Nielsen Norman Group*, 17 Jul. Available at: https://www.nngroup.com/articles/tone-of-voice-dimensions/ (Accessed: 2 Nov 2026).

Consistency builds trust, so store each standard only once.

Why consistency matters

Tone and repetition both have measurable effects. In an online survey of 100 users, the Nielsen Norman Group showed the same message in different tones. The tone changed how trustworthy users found a company, and trust was the strongest driver of whether they would recommend it to a friend: it explained more than half of the differences in those ratings.[25] A tone that shifts from piece to piece therefore puts that trust at risk. Repetition works too: a meta-analysis of 51 studies found that people rate statements they have heard before as more likely to be true.[26] An industry survey points the same way for brand messages. In 2019, the communications agency Edelman asked 8,000 people in eight markets how much they trusted a brand message. Among those who had not come across it before, 13% trusted it strongly. Among those who had already met it in three different ways, such as advertising, media reports or friends, it was 52%; after six, 87%.[27] That is a reason to repeat the key messages, and a reason to check them against the verified facts first (page 16).

The key message map and the tone of voice only work if everyone uses the same version. They are stored once, in the shared context of the marketing OS, and nowhere else. People read them before they write and check against them before they approve. The AI model drafts from the same files. This is the third principle of the marketing OS at work in the content engine: one shared context, and the same yardstick for people and machine (page 9).

IN PRACTICE

In my team, as communications lead, I maintained the tone of voice guide and the key message map and made them available to everyone who worked with content. Other units used them too, to bring their presentations in line with the company's wording.

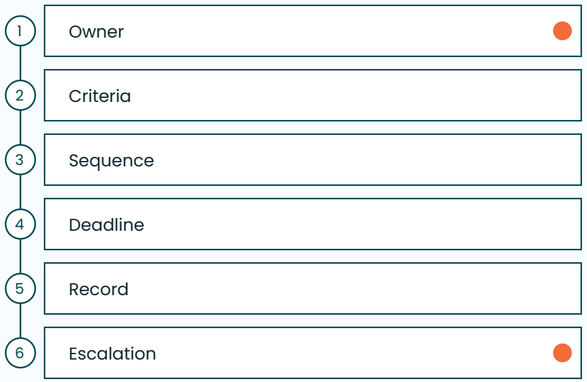
[25] Moran, K. (2016b) 'The impact of tone of voice on users' brand perception', *Nielsen Norman Group*, 7 Aug. Available at: https://www.nngroup.com/articles/tone-voice-users/ (Accessed: 2 Nov 2026).

[26] Dechêne, A., Stahl, C., Hansen, J. and Wänke, M. (2010) 'The truth about the truth: a meta-analytic review of the truth effect', *Personality and Social Psychology Review*, 14(2), pp. 238–257. Available at: https://doi.org/10.1177/1088868309352251 (Accessed: 2 Nov 2026).

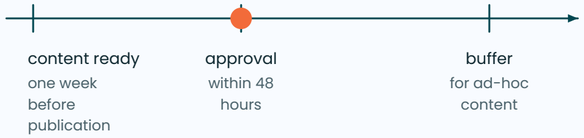
[27] Edelman (2019) *2019 Edelman Trust Barometer Special Report: In Brands We Trust?* Available at: https://www.edelman.com/sites/g/files/aatuss191/files/2019-06/2019_edelman_trust_barometer_special_report_in_brands_we_trust.pdf (Accessed: 1 Nov 2026).

A clear approval process rests on six building blocks.

SIX BUILDING BLOCKS



TIMELINE



Every approval has an owner: the marketing OS sets this principle for every module. For content, it looks like this.

Approval is where governance becomes real. It is also where work most often stalls. In my team, that was complex content, such as political topics around regulation or case studies from client projects, waiting for sign-off at the top.

Whatever the size of the team, an approval process needs six building blocks:

- **An owner** for each type of check, named by role.
- **Criteria:** what is checked, and the one question the check answers.
- **A sequence:** who checks first, and who gives the final release (next page).
- **A deadline** for every approval.
- **A record** of every decision, visible to the whole team.
- **An escalation rule** for when an approval does not come in time.

Without these, approval is a favour someone does when they find the time. With them, it is a step in the process that everyone can plan around.

Two safeguards keep the building blocks honest. The first comes from information security and is called separation of duties: no one should hold enough rights to misuse a system on their own.[28] For content, this means that whoever writes a piece does not approve it alone.

The second is timing. Research on accountability finds that it prompts careful, self-critical thinking mainly under one condition: people know before they decide that they will have to justify the decision. The audience must be one they see as legitimate and whose views they do not know. Even then, the effect on the quality of judgement varies. When accountability comes only afterwards, it tends to produce self-justification rather than better thinking.[29] So the reviewer is named when the work is planned, not when the draft is finished. The author knows from the start who will check what.

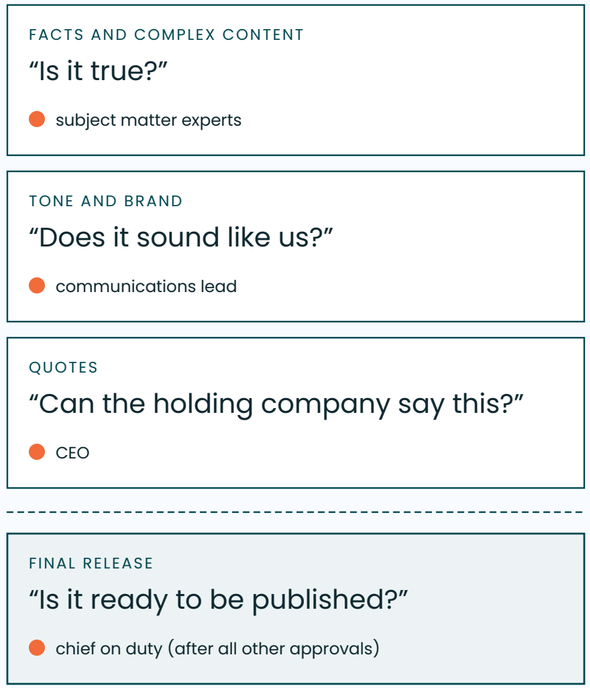
IN PRACTICE

In my team, a piece was ideally ready a week before its publication date and then only waited for approval. Within that week, our goal was to give every approval within 48 hours. This buffer is what left room for ad-hoc content that had to be created and approved at once. Every approval was recorded in the shared workspace. We set up the approval process as an automation in Notion: as soon as a content piece was approved, its author received a notification, and its label in the content calendar changed to "to be published".

[28] National Institute of Standards and Technology (no date) 'Separation of duty', *Computer Security Resource Center glossary*. Available at: https://csrc.nist.gov/glossary/term/separation_of_duty (Accessed: 2 Nov 2026).

[29] Lerner, J.S. and Tetlock, P.E. (1999) 'Accounting for the effects of accountability', *Psychological Bulletin*, 125(2), pp. 255–275. Available at: https://doi.org/10.1037/0033-2909.125.2.255 (Accessed: 2 Nov 2026).

Split approval by type of risk, and trust each owner.



The building blocks say what an approval needs. This page shows who gives it.

Owners by type of risk

Not every check asks the same question, so not every check needs the same person. In my team, approval was split by type of risk:

- **Facts and complex or scientific content.** The question: "Is it true?" The owner: subject matter experts, and only they.
- **Tone of voice and brand.** The question: "Does it sound like us?" The owner: the communications lead.
- **Quotes.** The question: "Can leadership say this?" The owner: the CEO.
- **Final release.** The question: "Is it ready to be published?" The owner: the chief on duty, always, after all other approvals.

As at a rocket launch, every position gives its go only for its own area, and one lead gives the final go.[30]

IN PRACTICE

At Purpose Green, as communications lead, I approved everything about language. Content and facts were approved by the subject matter experts, quotes by the CEO. As chief on duty, I gave the final release before anything went online.

Here the playbook departs from the corporate newsroom model by Prof. Christoph Moss (page 13). Its standard process has a single approval point, where the business unit approves the base document.[31] For a team that publishes science-heavy content, one approval point would have mixed questions that need different people to answer them. So approval here is split by type of risk, and each owner answers only the question they can judge. In the model, the strategy team also has the final say on content in case of conflict. Here the work is split: the strategy team decides the quarterly focal topics, and the chief on duty releases each piece.

Trust

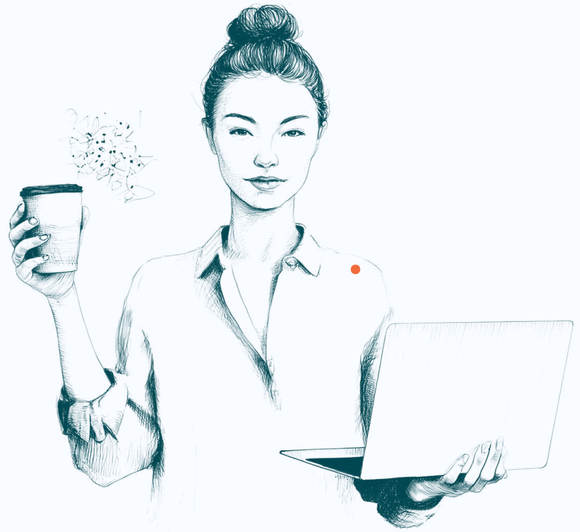
Trust research explains why the split works. Trust means being willing to rely on someone without being able to monitor or control them. And the ability that trust rests on is domain-specific: high ability at one task does not necessarily imply high ability at another.[32] A subject matter expert is the right person to judge a fact, but not necessarily a headline. A CEO is the right person to judge a quote, but not the details of a regulation. So once an owner has answered their question, nobody reopens it. The chief on duty checks that the piece works as a whole, not whether the facts are right.

[30] NASA (2005) *Launch Services Program Pre-Launch Readiness Reviews*. NASA Policy Directive NPD 8610.24C, revalidated 2 Apr 2025 with Change 3. Washington, DC: NASA, section 5.d(6). Available at: https://explorers.larc.nasa.gov/APSMEX26/SMEX/pdf_files/N_PD_8610_024C_chg3.pdf (Accessed: 3 Nov 2026).

[31] Becker, N. (2026) 'Der Standardprozess im Corporate Newsroom', *Mediamoss-Blog*, 21 May. Available at: https://mediamossblog.com/2026/05/21/der-standardprozess-im-corporate-newsroom/ (Accessed: 2 Nov 2026).

[32] Mayer, R.C., Davis, J.H. and Schoorman, F.D. (1995) 'An integrative model of organizational trust', *Academy of Management Review*, 20(3), pp. 709–734. Available at: https://doi.org/10.5465/amr.1995.9508080335 (Accessed: 2 Nov 2026).

Trusted owners keep approvals moving.



Not every approval belongs at the top

The lesson I would pass on to any team: do not place every approval at the top. When each type of risk has its own owner, leadership signs off only what only leadership can speak for, and pieces stop waiting in one inbox. Everything else needs owners who are trusted to decide. And however many specialists have approved a piece, one role releases it: the chief on duty, who sees the piece as a whole before it goes online.

This is not only a question of speed. Writing in *Harvard Business Review*, Paul Rogers and Marcia W. Blenko of Bain & Company observed that "decisions routinely get stuck inside the organization" when it is unclear who decides.[33]

Research on empowerment points the same way. A meta-analysis links psychological empowerment with higher task performance, and team empowerment with higher team performance.[34] Another finds that leaders who hand over responsibility see better performance and more creativity, in individuals and in teams, partly because, individually, people trust them more.[35] These findings are correlational, but they point in one direction.

Clear ownership also protects against a familiar trap. A meta-analysis links unclear roles with lower performance.[36] When "everyone" approves, nobody is accountable.

The team sets its priorities in the weekly editorial conference (pages 25–26). In the content engine, the same owners decide at the approval checkpoint (page 42).

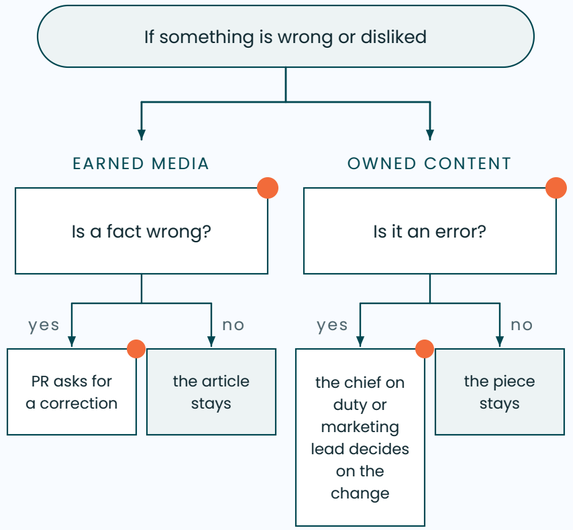
[33] Rogers, P. and Blenko, M.W. (2006) 'Who has the D? How clear decision roles enhance organizational performance', *Harvard Business Review*, Jan. Available at: https://hbr.org/2006/01/who-has-the-d-how-clear-decision-roles-enhance-organizational-performance (Accessed: 2 Nov 2026).

[34] Seibert, S.E., Wang, G. and Courtright, S.H. (2011) 'Antecedents and consequences of psychological and team empowerment in organizations: a meta-analytic review', *Journal of Applied Psychology*, 96(5), pp. 981–1003. Available at: https://doi.org/10.1037/a0022676 (Accessed: 2 Nov 2026).

[35] Lee, A., Willis, S. and Tian, A.W. (2018) 'Empowering leadership: a meta-analytic examination of incremental contribution, mediation, and moderation', *Journal of Organizational Behavior*, 39(3), pp. 306–325. Available at: https://doi.org/10.1002/job.2220 (Accessed: 2 Nov 2026).

[36] Gilboa, S., Shirom, A., Fried, Y. and Cooper, C. (2008) 'A meta-analysis of work demand stressors and job performance: examining main and moderating effects', *Personnel Psychology*, 61(2), pp. 227–271. Available at: https://eprints.lancs.ac.uk/id/eprint/45119 (Accessed: 2 Nov 2026).

After publication, changes need a clear owner too.



Approval ends with the release. Governance does not. Once a piece is out, the question is no longer who approves it, but who may change it. Without a rule, the answer is whoever is most senior, and every published piece stays open for revision.

Earned media

The first distinction is who owns the content. Communications professionals sort media by exactly that question. The PESO model, created by Gini Dietrich, distinguishes paid, earned, shared and owned media, and its point is that the four reinforce each other.[37] This page looks at two of them: earned and owned media. Earned media, such as an article by a journalist, belongs to the newsroom that published it. It is the journalist's work, not the company's text. Newsrooms correct their own mistakes: the German press code obliges a publication to correct reports that turn out to be false.[38] So if a fact is wrong, PR can and should ask for a correction. If a sentence is merely not to someone's taste, the article stays as it is. Asking a newsroom to change it puts its independence in question and the relationship PR has built at risk. Published is published.

Owned content

Owned content is different: the company's own posts, blog articles and newsletters. Here, corrections are possible, and errors should be fixed. Everyone in the team can and should point them out. But a published piece should not be revised again and again. Each change needs a reason, not a preference.

Who decides

That is why the decision has an owner. The chief on duty or the marketing lead decides whether a published piece on an owned channel is changed, not the most senior person who reads it. The CEO approves their own quotes before publication (page 20).

IN PRACTICE

Several times in my career, a CEO asked whether we could call the journalist and have an article changed, because one sentence was not quite to their taste. The distinction between earned and owned media gave a clear answer: the article belongs to the newsroom, which corrects factual errors, not matters of taste.

[37] Spin Sucks (no date) *The PESO Model®: the MarComm operating system*. Available at: https://spinsucks.com/peso-model/ (Accessed: 2 Nov 2026).

[38] Deutscher Presserat (2025) *Publizistische Grundsätze (Pressekodex)*. Version of 19 Mar 2025. Available at: https://www.presserat.de/pressekodex.html (Accessed: 2 Nov 2026).

GLOSSARY

PAID, EARNED, SHARED AND OWNED MEDIA (PESO)

Paid: advertising. Earned: coverage by others, such as journalists. Shared: social media. Owned: the company's own channels, such as its website, blog or newsletter.

Use AI openly, and have a person check everything before release.



Why openness

A marketing team that wants to be AI-first has to use AI, openly and every day. Banning it does not stop it. It only moves the use into private chatbots, where nobody sees the prompts, the sources or the drafts. That is working in silos, in its newest form.

What dpa decided

The news agency dpa took the same view when it set out its own rules for AI. Its first rule is openness: the agency uses AI and is open to using more of it. Its second rule sets the condition: "dpa only uses AI under human supervision."[39] That is • human-in-the-loop, written down by a newsroom.

Five rules turn that condition into practice for a marketing team.

Nothing goes out unchecked

The first rule is that nothing goes out unchecked. Every draft an AI model produces passes the people responsible for it: first the author, then the owners of each check (page 20). The AI model prepares; it never publishes.

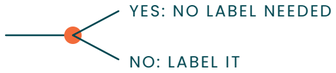
Sources before speed

The second rule is sources before speed. The biggest risks of AI in content are unchecked sources and invented facts. A model can phrase a wrong figure as confidently as a right one. So every fact in an AI draft is checked against the list of verified facts and against its primary source, as for any other draft (page 16).

[39] Raabe, S. (2023) 'Offen, verantwortungsvoll und transparent – Die Guidelines der dpa für Künstliche Intelligenz', *dpa Innovation*, 3 Apr, updated 10 Jul 2024. Available at: https://innovation.dpa.com/2023/04/03/kuenstliche-intelligenz-fuenf-guidelines-der-dpa/ (Accessed: 2 Nov 2026).

Label AI by responsibility, share one context, protect personal data.

5 RULES FOR AI USE

1. Nothing goes out unchecked.
2. Sources before speed.
3. Label AI by responsibility.
   Did a person review it and take responsibility?
   

4. One shared context.
5. Protect personal data.

Label AI by responsibility

The third rule is about labelling. Today, almost every piece of content is made with some help from AI: a first draft, a summary, a headline. Labelling every piece would say little. What matters is whether a person checked the piece and takes responsibility for it.

What the law says

European law takes the same line. Under the EU AI Act, AI-generated text that is published to inform the public on matters of public interest has to be disclosed. The obligation does not apply where the content "has undergone a process of human review or editorial control" and a person or company "holds editorial responsibility for the publication".[40] The checkpoints in this playbook are exactly that kind of review. This is not legal advice; teams should check what applies to them.

Where I draw the line

Marketing teams have to set their own guardrails for the use of AI. Labelling everything makes little sense, but I would always label content produced 100% by AI, with no human involved.

One shared context

The fourth rule is one shared context. Prompts, the tone of voice, the key messages and the verified facts live in the shared context of the marketing OS, not in anyone's private chat history. Everyone, and every AI model, works from the same standards (pages 16–18).

No personal data

The fifth rule: customer data and other personal data do not go into AI tools. Only tools that the company has approved are used, with business accounts whose data is not used to train the model.

IN PRACTICE

In my team, AI was used openly; we were even asked to use it, to work more efficiently. We never published AI content without a person's approval: the communications lead or a subject matter expert checked it, and the chief on duty released it. Customer data and personal data stayed out of AI tools, and nobody used a private account. We only used business accounts whose data was not used for training.

[40] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), Article 50(4). *Official Journal of the European Union*, L series, 2024/1689, 12 Jul 2024. Available at: http://data.europa.eu/eli/reg/2024/1689/oj (Accessed: 29 Oct 2026).

Plan the week together with the whole team.



Governance needs one place where decisions are made, together and at a fixed time. In my team, that place was the weekly editorial conference, every Monday morning. The whole team took part. It worked like the editorial conference of a large daily newspaper, only once a week instead of every morning. It was not a board that approved from above.

Four inputs came together: topics from the business units, external news monitoring, the quarterly focal topics and the company's objectives. As chief on duty, I ran the conference. We matched what came in against the focal topics and set priorities by the objectives. The workload view in the shared workspace showed what the team could realistically take on. A person, not the system, assigned the work.

The conference always started by looking back. What went well last week, and what did not? The question covered the content and the collaboration. We did not go through every result, but discussed the pieces that had worked especially well or especially poorly.

One example: at times the design team was overloaded. Marketing and content were not its only stakeholders, and too many requests came in at once. We named this in the conference and took the design team's capacity into account in the planning.

The outcome was always the same: a concrete plan for the week. Everyone left the conference knowing what they had to deliver, and by when.

Weekly meetings pay off when the team shares and reviews the work.



"Not another meeting" is a fair objection. Research gives a more precise answer: the time spent in meetings is not the problem as such. Whether a lot of meeting time weighed on well-being depended, in two surveys of full-time employees, on three things: how much the work depended on others, how good the meetings were and how much people strove to get their work done.[41]

A content team is interdependent by nature. A campaign needs text, design, channels and subject matter experts at the same time. That is exactly where a good meeting pays off, and where a poor one costs the most.

The review as a debrief

Looking back each week turns the conference into a structured debrief. A meta-analysis of debrief studies found that debriefs improved effectiveness markedly, for teams and for individuals, in simulated and in real settings. Alignment made them more effective, and facilitation and structure probably help as well.[42] Our design team, overloaded at times, showed why: a problem that is named every week can be planned around before it turns into conflict.

Psychological safety

The look back only works if people say what did not work. That takes psychological safety: a shared belief in the team that it is safe to take interpersonal risks. In a field study of work teams, psychological safety went with more learning behaviour, and learning behaviour was linked to team performance.[43]

In a conference, this means that anyone can call a topic weak, a claim unproven or a deadline unrealistic, without paying for it. The chief on duty sets that tone.

Who decides what

The conference is where the team decides priorities and who does what; the owners decide approvals. In the content engine, the conference picks the topics and reads the results (page 42).

[41] Rogelberg, S.G., Leach, D.J., Warr, P.B. and Burnfield, J.L. (2006) '"Not another meeting!" Are meeting time demands related to employee well-being?', *Journal of Applied Psychology*, 91(1), pp. 83–96. Available at: https://doi.org/10.1037/0021-9010.91.1.83 (Accessed: 2 Nov 2026).

[42] Tannenbaum, S.I. and Cerasoli, C.P. (2013) 'Do team and individual debriefs enhance performance? A meta-analysis', *Human Factors*, 55(1), pp. 231–245. Available at: https://doi.org/10.1177/0018720812448394 (Accessed: 2 Nov 2026).

[43] Edmondson, A. (1999) 'Psychological safety and learning behavior in work teams', *Administrative Science Quarterly*, 44(2), pp. 350–383. Available at: https://dash.harvard.edu/entities/publication/13a7b031-0fdd-45ec-a7e0-2b80e2bc679f (Accessed: 2 Nov 2026).



PART 3

How the content engine works

28 The content engine's working rules

30 The architecture

31 Context

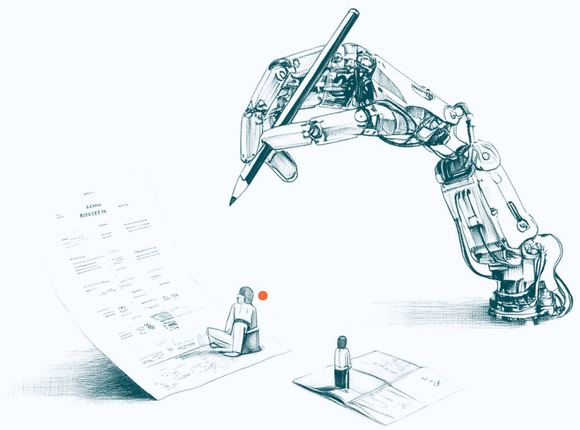
36 Research and pick

38 Production and approval

39 Distribution and posting

40 Analytics and reading

The machine prepares the work, and a person decides.



Three principles hold for every module: the machine prepares, a named person decides; every approval has an owner; and all modules share one context. This part shows how the content engine applies them every day. Three working rules guide it: two restate the first and third principles, and one more adds what a running system needs, learning from every decision. The technology can change; the rules should not.

The first rule

The first rule is the one the content engine is built around: machines prepare, people decide. The machine researches topics, sorts sources, drafts texts and gathers the results. A person decides at four checkpoints: which idea to pick, whether a draft is approved, when and where it is posted, and what the results mean. Research, writing and publishing are separate stations, and a person decides between one and the next. Within a station, the machine may check its own work, such as a review step that filters weak proposals. Those checks prepare the decision; they do not replace it.

Why the line is drawn there

Why draw the line there? A final decision should always be made by a person, not by a machine, because a person answers for what the team publishes.

A model cannot be held to account for a wrong figure or a harmful claim; the person whose name is on the piece can.

What the research says

Research on automation gives this line a frame. A classic model describes four functions that can each be automated to a different degree: gathering information, analysing it, choosing a decision and carrying it out.[44] The content engine automates the first two heavily. The choice stays with a person.

There is another reason to leave the decision with a person. In a field experiment with consultants, AI improved work on tasks inside its capabilities, and made people more likely to get it wrong on a task outside them, although the tasks were of similar difficulty.[45] The boundary is not visible from the outside, so someone has to judge every output.

That does not make AI a poor tool. In an experiment with 453 professionals doing writing tasks, access to an AI chatbot cut the time taken by 40% and raised the quality of the output by 18%.[46] It is an excellent first-draft machine. It is not a publisher.

[44] Parasuraman, R., Sheridan, T.B. and Wickens, C.D. (2000) 'A model for types and levels of human interaction with automation', *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans*, 30(3), pp. 286–297. Available at: https://doi.org/10.1109/3468.844354 (Accessed: 2 Nov 2026).

[45] Dell'Acqua, F., McFowland III, E., Mollick, E., Lifshitz, H., Kellogg, K.C., Rajendran, S., Krayer, L., Candelon, F. and Lakhani, K.R. (2026) 'Navigating the jagged technological frontier: field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality', *Organization Science*, 37(2), pp. 403–423. Available at: https://www.hbs.edu/ris/Publication%20Files/dell-acqua-et-al-2026-navigating-the-jagged-technological-frontier_5c589c8c-fbb5-458f-b285-c944746cd717.pdf (Accessed: 2 Nov 2026).

[46] Noy, S. and Zhang, W. (2023) 'Experimental evidence on the productivity effects of generative artificial intelligence', *Science*, 381(6654), pp. 187–192. Available at: https://doi.org/10.1126/science.adh2586 (Accessed: 2 Nov 2026).

Context comes before content, and every decision teaches the system.



The second rule

The second rule is context before content. The content engine does not write from a blank prompt. It drafts from stored context: the tone of voice and the verified facts, the knowledge collected over time, what is known about the audience, and the lessons from earlier pieces. Every claim in a draft carries its source, so the person who approves it can check it. This context, not the AI model, decides the quality of a draft (page 31).

The third rule

The third rule is that the system learns from decisions. Every pick, every rejection, every note on a draft and every result is recorded and flows back into the context. Over time, the content engine prepares better because people have decided. The numbers inform the next decision. They do not take it.

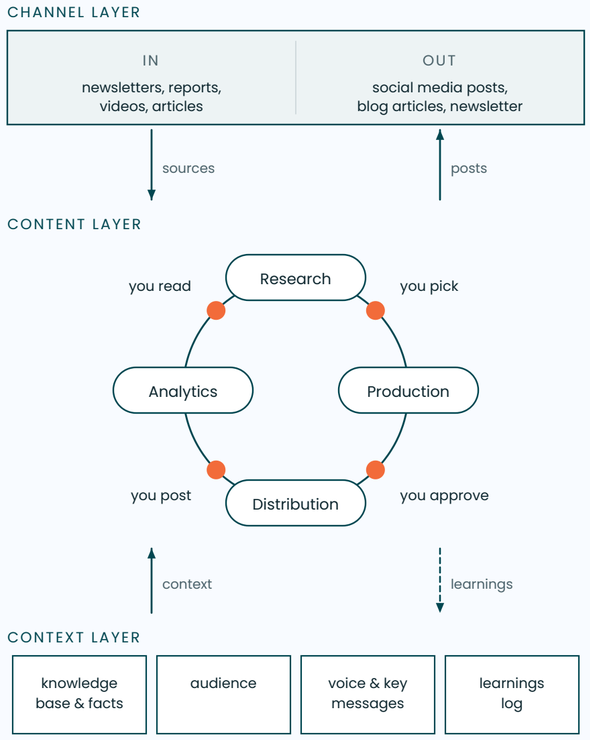
Trust that fits

Together, the three rules aim at one thing: trust that fits. People can trust a system too much or too little. The trust they place in it should match what it can actually do. Visible sources, recorded decisions and checkable drafts make that possible.

Who holds the checkpoints

In the content engine described here, I hold all four checkpoints myself. That follows from what it is: a system that works for me and my own content. In a team, the checkpoints would go to different owners: pick to the editorial conference, approval to the owners by type of risk and the final release to the chief on duty, posting to the channel owners, reading to the editorial conference, led by the chief on duty, in the weekly look back (page 42). Each of them is a defined team role (page 44).

The content engine works on three layers.



The content engine has three layers, stacked on one axis. Each has a single job.

The top layer is the channel layer: where things come in and where they go out. In come newsletters, reports, videos and articles. Out go social media posts, blog articles and a newsletter.

The middle layer is the content layer. This is where the work happens, in four stations on a loop. Research finds starting points for relevant content. Production turns them into drafts in the right voice. In distribution, a person publishes the content. Analytics measures what landed.

The bottom layer is the context layer: four knowledge stores that keep learning. The knowledge base and facts hold internal and external sources, that is, everything the team and the automated intake collect, each item with a summary, a category and a rating (page 32), plus the facts verified from them. The audience store holds who it is for. Voice and key messages hold how you sound and what you say. The learnings log holds what you learned.

Four flows connect the layers. Sources come in from the channels. Context feeds every draft. Finished posts go out. And learnings flow back from analytics into the context, where they update the dos and don'ts.

Between the four stations sit four checkpoints: you pick, you approve, you post, you read. The loop never runs on by itself. On every turn, it passes a person four times.

The picture shows the content engine as I run it for my own content. Every part of it is a role, not a product. The knowledge stores can be any shared database; the stations can run on any automation layer with any AI model.

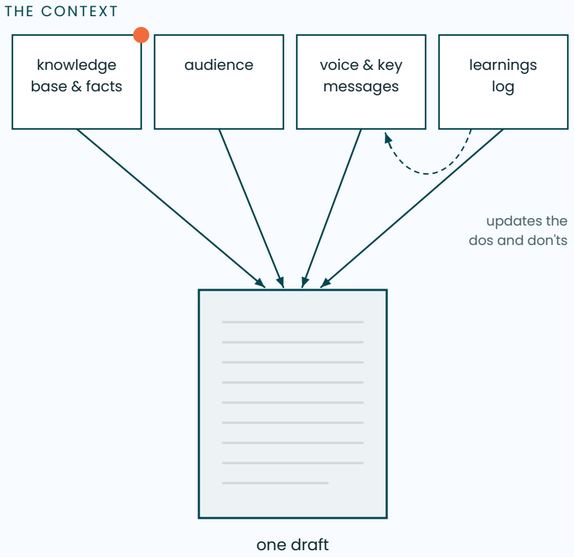
The following pages show the context layer first, which every station draws on, then the four stations one at a time: what the machine prepares, what the person decides, and why the checkpoint sits where it does.

GLOSSARY

LEARNINGS LOG

A running record of what the team learns from its results. Its lessons update the dos and don'ts that feed every new draft (page 40).

The AI model writes the draft, and the context decides its quality.



Give an AI model a blank prompt and it writes a text that could come from anyone. Give the same model good context and it writes a text that sounds like you, uses facts you can stand behind and speaks to the people you want to reach. What makes a draft yours is not the model. It is what the model is given.

That is why the content engine treats context as its most valuable part. AI models change quickly, and the content engine is built so that one can be swapped for another. The context stays, grows and gets better with every decision.

The four stores

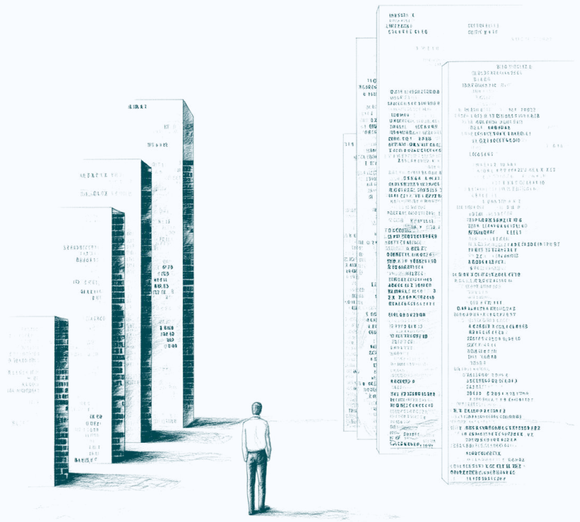
The **knowledge base and facts** hold what you know: internal and external sources such as articles, studies and reports, each with a summary, a category and your own rating or note. From these sources comes a list of verified facts. That list also records what must not be used, such as a figure that turned out to be wrong. A good facts store remembers mistakes as carefully as truths. The knowledge base also keeps a record of the topics already covered, so the content engine does not repeat itself.

The **audience store** holds who you write for: personas and the positioning brief (page 34).

Voice and key messages hold how you sound and what you say: the tone of voice, the key message map and the dos and don'ts.

The **learnings log** holds what you learned: the results and decisions from analytics, which update the dos and don'ts.

A shared knowledge base gives the topic desk a memory.



Topics come from what a team reads, hears and learns: a study in the trade press, a newsletter, a question from a customer, an insight from a business unit. Without one shared place, that knowledge stays in private inboxes and bookmarks. It is working in silos again, one step before any content is written.

The knowledge base is that shared place. Of the four context stores, it is the only one the outside world flows into every day. Because it records which items became topics, it also remembers what the team has said. That makes it the topic desk's memory. Its sources also feed the list of verified facts (page 16).

Two ways in

Items reach the knowledge base in two ways. The first is people. Every member of the team can add something worth knowing, as long as it fits the topic islands and content pillars (page 12) and comes with a concrete, reliable source. No source, no entry.

The second is the machine. In a content engine, an automated intake sits in front of the knowledge base. It takes in selected newsletters, plus articles, videos and documents that someone has saved by hand. For each item, it fetches the full text, flags what it could not read in full, writes a short summary, assigns a category and removes duplicates.

What people add

The machine sorts; people judge. A team member rates an item and adds a note on why it matters. That note is the most valuable field in the knowledge base. Rated and annotated items are proposed first when their category is next on the plan, so a single note steers what the content engine researches later (page 36).

How it feeds the topic desk

Every working day, the knowledge base produces a short digest of its newest and highest-rated items. It goes to research and is collected for the weekly editorial conference, where topics are set (pages 25–26). When an item becomes a topic, it is linked to the piece in the editorial record. So the knowledge base also shows what has already been used, and nothing is told twice by accident.

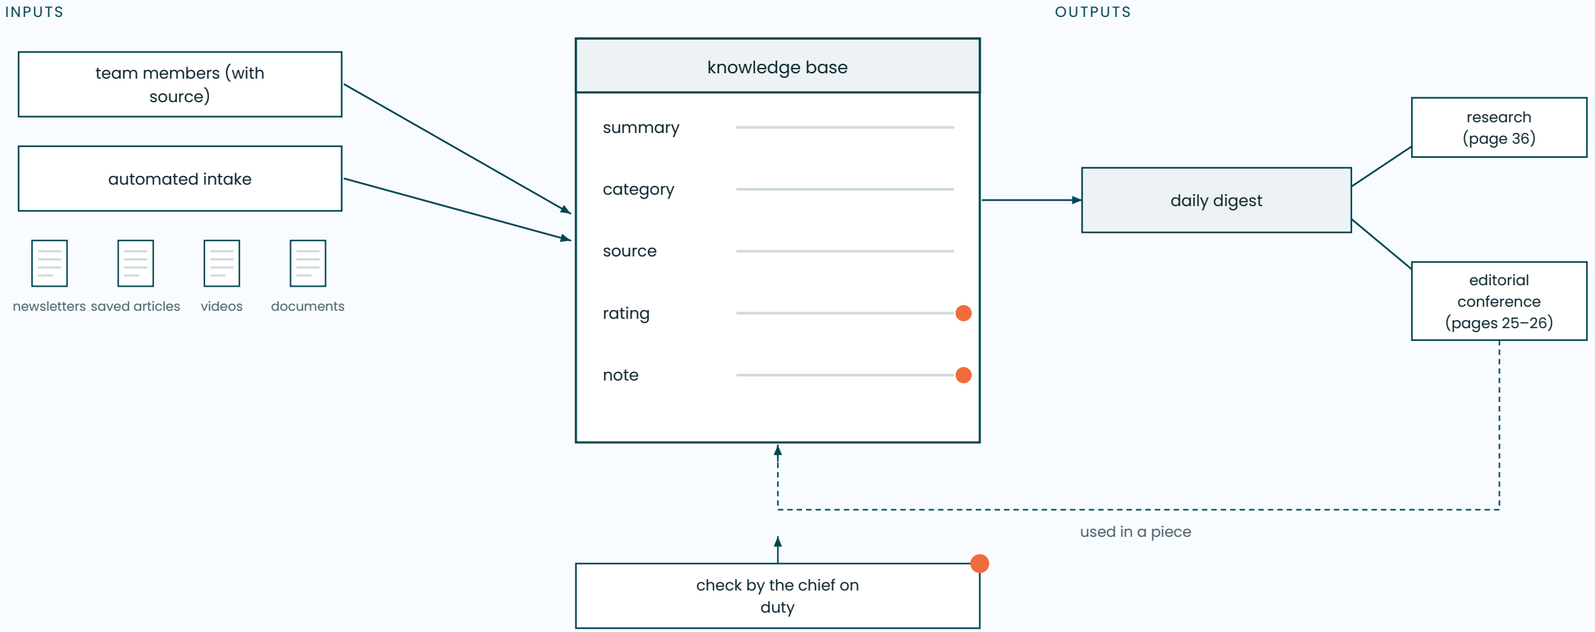
Who maintains and reviews it

Collecting is open to everyone; deciding is not. The whole marketing team maintains the knowledge base, and the chief on duty checks it: whether the categories still fit, whether the sources hold and whether an item is ready to become a topic. The rule of two independent sources applies to every item that becomes a topic.

IN PRACTICE

At Purpose Green, the team worked from two sources: a project database that held everything about every project and campaign, and a knowledge base. Everyone in the team added to the knowledge base; as chief on duty, I checked it.

For my own content engine, I built the intake as an automation. Selected newsletters arrive by email; n8n resolves the links in them and fetches each article in full. I can save single articles and videos by hand, also from my phone; for videos, the transcript is fetched. PDFs I drop into a folder are summarised by Gemini. Everything lands in one database in Notion, with a summary, a category and space for my rating and note.



Know who you write for.





The audience store answers two questions. Who are we writing for? And what do we want them to remember about us?

Personas

Personas answer the first. A persona describes a typical person in the target group as if you knew them: their role and responsibilities, their working day, the questions they ask, their pain points and goals, what convinces them, and which media they use first and which second. It also includes personal details, such as an age range, a place and a life situation, because these give the team a feel for the person. Research on personas finds that their greatest value lies in providing a shared basis for communication.[47] Everyone in the team, and every station of the content engine, writes for the same person.

Positioning

Positioning answers the second question. A positioning brief states whom the company is there for, which problem it solves, what sets it apart and the proof for that, the core message (set out in full in the key message map, page 17), and what it does not say.

Where the content comes from

Where does the content of a persona come from? Large companies run market segmentation studies. Most start-ups and small teams cannot afford them, and they do not have to wait for them. A realistic persona combines the team's experience and gut feeling with what is already there: customer conversations, notes from sales and service, existing records and collected feedback. The same research describes a lighter persona effort by two people that relied solely on existing user research; its personas were less detailed.[47]

What matters is honesty about the basis. Early persona efforts failed when the characters were not believable, because they were "obviously designed by committee (not based on data)" or because their relationship to data was not clear.[47] So every persona should state what it is based on, and be revised when new conversations or results come in. Creating personas also makes the team's assumptions about its audience explicit,[47] and an explicit assumption can be checked.

Appendix A provides templates for both: a persona profile and a positioning brief.

[47] Pruitt, J. and Grudin, J. (2003) 'Personas: practice and theory', in *Proceedings of the 2003 Conference on Designing for User Experiences (DUX '03)*. New York: ACM, pp. 1–15. Available at: https://www.microsoft.com/en-us/research/wp-content/uploads/2017/01/personas-practice-and-theory.pdf (Accessed: 2 Nov 2026).

Good context reduces errors but does not prevent them.

Research supports the idea that context improves AI writing. AI models that retrieved information from an external knowledge source before they wrote produced more specific and more factual text than a comparable model that relied only on what it had learned in training.[48]

But context is not a guarantee. An audit of AI search engines in 2023 found that only about half of the generated sentences were fully supported by the citations attached to them, and that about a quarter of the citations did not support their sentence.[49] A survey of research on AI hallucinations also discusses the limits of retrieval as a remedy.[50] The figures describe the systems of that time, not today's error rate. The lesson holds: a source next to a sentence is not yet proof. That is why every draft still passes a person who checks it (page 38).

Context also has owners. Alongside the checkpoints and the approval owners, it is the third place where governance meets the content engine. In a team, the communications lead owns voice, key messages and positioning. Subject matter experts own the facts. Whoever works closest to customers keeps the personas current. Context is curated, not just collected.

And context has to be maintained. Every decision and every result flows back into it. A store that nobody looks after goes stale, and the drafts go stale with it. Maintaining the context is part of the work, not overhead.

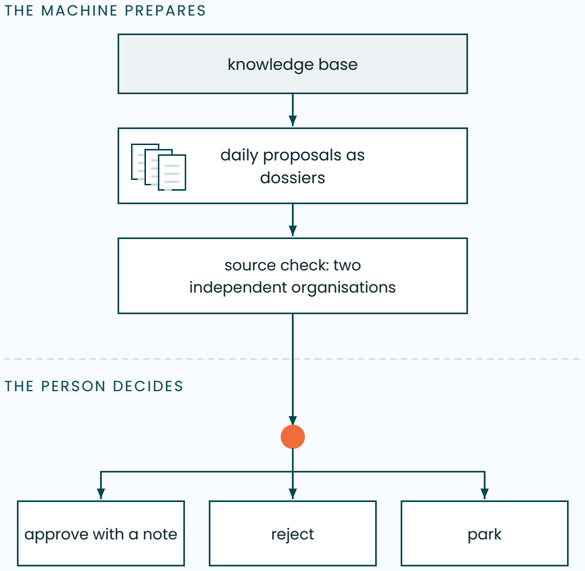
With the context in place, the stations can draw on it. The next pages take them one at a time, from research to analytics.

[48] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S. and Kiela, D. (2020) 'Retrieval-augmented generation for knowledge-intensive NLP tasks', in *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*. Available at: https://arxiv.org/abs/2005.11401 (Accessed: 2 Nov 2026).

[49] Liu, N.F., Zhang, T. and Liang, P. (2023) 'Evaluating verifiability in generative search engines', in *Findings of the Association for Computational Linguistics: EMNLP 2023*. Singapore: Association for Computational Linguistics, pp. 7001–7025. Available at: https://doi.org/10.18653/v1/2023.findings-emnlp.467 (Accessed: 2 Nov 2026).

[50] Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B. and Liu, T. (2025) 'A survey on hallucination in large language models: principles, taxonomy, challenges, and open questions', *ACM Transactions on Information Systems*, 43(2), pp. 1–55. Available at: https://arxiv.org/abs/2311.05232 (Accessed: 2 Nov 2026).

Research proposes topics, and a person picks.



What the machine prepares

Everything worth keeping goes into the knowledge base first: articles, newsletters, reports, videos and documents. Each item gets a short summary, a category and a relevance rating (page 32).

Every working day, the research station starts from the knowledge base: its daily digest of the newest items, with the ones I have rated or annotated first. From there, it proposes a small set of topics. Each proposal comes as a dossier: what triggered it, the key facts, the sources, a possible angle and a counter-argument. Categories rotate by weekday, so the mix stays broad.

Before anything reaches me, a strict review step checks the proposals. A topic is rejected unless at least two independent organisations back it, meaning sources with no commercial stake in the claim. That rule came from experience: the research step kept surfacing agency and tool-vendor blogs as sources. A source rule had to be built into the process.

What the person decides

I read the dossiers and decide: approve a topic with my own note, reject it, or park it for later. The note matters, because it tells the next station what my angle is.

I can also steer research ahead of time: a personal rating or note on an item in the knowledge base gives it priority in later proposals.

Why the checkpoint sits here: picking is the most strategic decision in the loop. Everything after it carries out what was chosen. Research on AI-assisted decisions finds that people rely less blindly on AI suggestions when the design makes them think before they accept, even though they like such designs less.[51] A dossier that includes a counter-argument is built on the same idea. In a team, the topic proposals collect over the week, and the next editorial conference picks from them (page 42).

IN PRACTICE

The In practice boxes on research, production and analytics come from my time leading communications at Purpose Green, where people ran the same checkpoints with less automation. There, we often responded to current coverage. When press reports voiced doubts or fears among property owners, we answered with numbers, data and facts. We also did myth-busting: we showed owners in concrete terms what could really happen to them. Whenever something caught our eye (a question in a customer email, a report in the daily press or a trade magazine), we noted it in the knowledge base and took it up in our content when the time was right.

[51] Buçinca, Z., Malaya, M.B. and Gajos, K.Z. (2021) 'To trust or to think: cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making', *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), pp. 1–21. Available at: https://doi.org/10.1145/3449287 (Accessed: 2 Nov 2026).

Knowledge base

Filled automatically from newsletters and a form; the team adds by hand.

All items | By category | Category filter | Check: other | Pipeline

Title	Type	Sender	Author	Date	Category	Summary
Lorem ipsum dolor sit amet	ARTICLE	no-reply@example-news.test	[Author]	29 October 2026	AI models and market	Lorem ipsum dolo
Consectetur adipiscing elit	ARTICLE	newsletter@example-mag.test	[Author]	29 October 2026	AI models and market	Sed do eiusmod te
Sed do eiusmod tempor	VIDEO	Added by hand		29 October 2026	Social media	Ut labore et dolore
Incididunt ut labore et	ARTICLE	newsletter@example-mag.test	[Author]	29 October 2026	AI models and market	Magna aliqua ut e
Ut enim ad minim veniam	PDF	Added by hand		29 October 2026	Other	
Quis nostrud exercitation	ARTICLE	team@example-review.test	[Author]	29 October 2026	Leadership and work	Ullamco laboris nis
Duis aute irure dolor	ARTICLE	team@example-review.test	[Author]	29 October 2026	AI models and market	In reprehenderit in
Voluptate velit esse	PDF	Added by hand		29 October 2026	Brand and content	Cillum dolore eu fu
Excepteur sint occaecat	ARTICLE	no-reply@example-news.test	[Author]	29 October 2026	Social media	Cupidatat non pro
Sunt in culpa qui officia	ARTICLE	newsletter@example-mag.test	[Author]	29 October 2026	Advertising	Deserunt mollit ani
Nemo enim ipsam voluptatem	ARTICLE	no-reply@example-news.test	[Author]	29 October 2026	Tech and platforms	Quia voluptas sit
Neque porro quisquam est	VIDEO	Added by hand		29 October 2026	Social media	Qui dolorem ipsum

ITEM TYPES — article (scraped text) — video (link, transcript fetched) — PDF file (summarised)

KNOWLEDGE BASE · ARTICLE RECORD

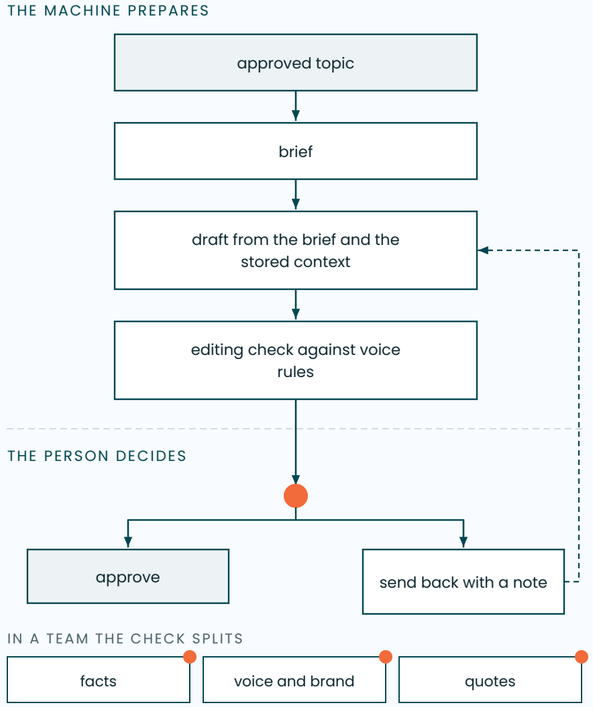
Lorem ipsum dolor sit amet: consectetur adipiscing elit

Sender	no-reply@example-news.test
Author	[Author]
Date	29 October 2026
Category	AI models and market
Summary	Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore.
Reading time (min)	2
Post status	Empty
My note	Empty
PDF	Empty
Source	example-news.test
Editorial	Empty
Relevance	Empty
Tags	Lorem, Ipsum
Language	EN
Status	OK
URL	example-news.test/lorem/ipsum-dolor/

ARTICLE TEXT (SCRAPED)

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure dolor in reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur. Excepteur sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt.

No draft without a brief, and no post without approval.



What the machine prepares

When I approve a topic, the content engine turns it into a brief: the topic, the angle, the key facts and my note. This is the clear brief every piece needs, now as a fixed step. No draft starts without one. The brief carries the sources the topic came from, many of them items from the knowledge base (page 32). So production builds on what was collected and selected, not on a new search.

The writing station then drafts from the brief and the stored context (verified facts, audience and positioning, tone of voice, key messages, dos and don'ts), and nothing else. It does not search the open web at this point, so it does not bring in new sources. It can still get facts wrong, which is why every draft is approved by a person. An editing step checks the draft against the voice rules. The sources travel with the draft, so every claim can be traced back.

What the person decides

I read the draft and decide: approve it, or send it back with a note on what has to change. The revision comes back to the same record, next to the earlier version, so the history of the decision stays visible.

Why the checkpoint sits here: this is where the risk of a confident AI error is highest. A wrong fact, a claim that goes too far or a tone that does not fit would go out under my name. In a team, this checkpoint splits by type of risk (page 20): facts to subject matter experts, voice and brand to the communications lead, quotes to the CEO.

IN PRACTICE

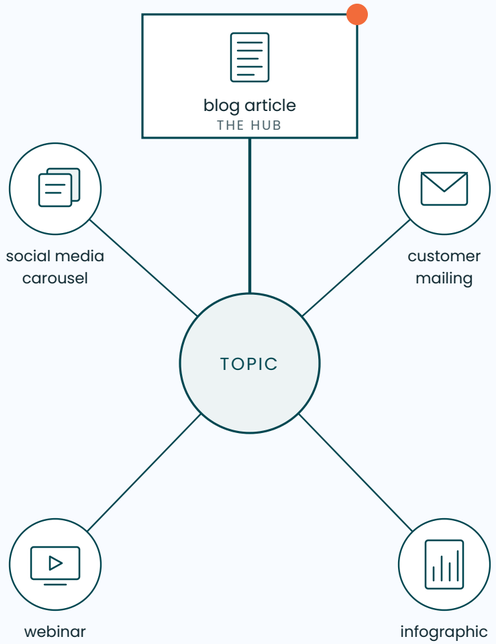
Building on what we had collected, we did data-driven PR, for example an analysis of what the new EU emissions trading for buildings (ETS-2) would cost property owners in Germany's largest cities. It helped make a broad public aware of energy-efficient renovation. Our analyses appeared in national media such as *Handelsblatt* and *Der Tagesspiegel*, and our CEO was quoted in many outlets. We used the same idea on our website: a CO_2 calculator that showed property owners how rising CO_2 prices would drive up their heating costs. Our argument always followed the same order: insulate the building first, then look at the heating, and only then ask whether a heat pump pays off. The calculator made the need for renovation visible in figures, gave owners orientation, and worked as a lead magnet at the same time.

GLOSSARY

DATA-DRIVEN PR

Press work built on the company's own data or analyses, offered to journalists as news in its own right.

Publishing is a person's job alone.



What the machine prepares

At this station the machine prepares nothing of its own, and that is deliberate. Distribution stays in the hands of a person. The approved draft waits in the editorial record with its planned publication date. The calendar shows what is due and when.

Scheduling tools for social media can publish a post at a set time. They help, but they do not change who is responsible: the post goes out under a human owner who has checked it before it is scheduled. In my own content engine, I decided against automated scheduling. Every automated step is another moving part that can fail, and another place where something goes out that nobody looked at again.

One topic, many formats

A blog works well as the hub (page 15). From one approved piece, the team derives other formats for other audiences: a webinar, an infographic, a social media carousel that explains the key point, a customer mailing. Each builds on the same approved facts and sources. One topic then multiplies into content that fits each audience.

Most marketing teams already use these channels. In a 2024 survey of 980 B2B marketers, 89% used organic social media, 84% a blog on their own website and 71% email newsletters. Yet only 37% to 42% said each of these channels produced their best results, and 37% named repurposing content as a challenge.[52] A topic can also travel beyond the team's own channels. A study of an online lending marketplace found that mentions in blogs and online communities drove sales, and also helped bring about coverage in traditional media.[53]

What the person decides

I post the piece myself, at the time and in the place that fit. Then I record the link to the published post in its editorial record. That one link connects the post to everything before it: the topic, the sources, the brief, the drafts and the decisions.

Why the checkpoint sits here: publishing is the moment a piece leaves the team's control. A last look before it goes out costs a minute. A post that should not have gone out costs far more. In a team, posting belongs to the owner of the channel, who knows the channel and its timing, after the chief on duty's final release (page 20); sometimes the chief on duty posts directly.

[52] Stahl, S. (2024) 'B2B content marketing benchmarks, budgets, and trends: outlook for 2025', *Content Marketing Institute*, 9 Oct. Available at: https://contentmarketinginstitute.com/b2b-research/b2b-content-marketing-trends-research-2025 (Accessed: 2 Nov 2026).

[53] Stephen, A.T. and Galak, J. (2012) 'The effects of traditional and social earned media on sales: a study of a microlending marketplace', *Journal of Marketing Research*, 49(5), pp. 624–639. Available at: https://doi.org/10.1509/jmr.09.0401 (Accessed: 2 Nov 2026).

Analytics closes the loop, and a person reads the results.



What the machine prepares

At regular intervals, the results of published posts are imported into the content engine. Each result is matched to its editorial record through the post link recorded at publication. The content engine then shows what landed, by category and by format, over time.

This station came late. The first version of the content engine optimised throughput: more ideas, more drafts. Nobody could say which of them landed. Measurement was added to close the loop.

What the person reads

I read the results and decide what they mean. Does the research focus need to change? Is a format not working, or does it simply need more time? The numbers do not steer topic selection automatically. They support my decision, and I make it. If the research focus changes, I change the categories the research station works from (page 36). What I learn about writing goes into the learnings log, and from there into the dos and don'ts that feed every new draft.

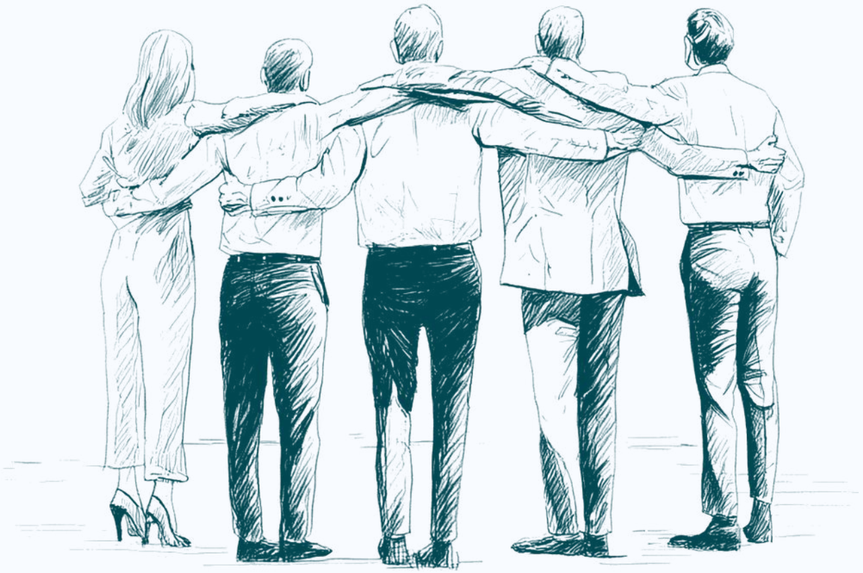
Why the checkpoint sits here: results need interpretation. A weak number can mean a weak topic, a bad moment or a format that is still new. Only a person who knows the context can tell the difference.

In a team, the editorial conference analyses the content results in its weekly look back (page 26).

One honest limit: the loop needs a little human discipline. The system cannot learn from a post whose link nobody recorded.

IN PRACTICE

At Purpose Green, our analysis showed a clear pattern: the pieces built on our own data were the ones the media picked up and spread most widely. The numbers showed the pattern; what to do about it was our decision. We refocused our PR on data-driven pieces (page 38) and added targeted thought leadership from our CEO and our subject matter experts.



PART 4

How a team makes it work

42 Owners in the team

44 Roles and adoption

46 Tools

48 Pitfalls

50 The first 90 days

Every checkpoint in the content engine has a named owner.

CHECKPOINT	OWNER IN A TEAM
pick	● editorial conference ● chief on duty (ad hoc)
approve	● subject matter experts ● communications lead or marketing lead ● CEO ● chief on duty (final release)
post	● channel owner ● chief on duty (sometimes)
read	● editorial conference ● chief on duty (lead)
context	● communications lead ● subject matter experts ● closest to customers

Part 2 described the governance of the content engine, Part 3 the content engine at work. This part shows how a team runs them.

In the content engine as I run it, I hold every checkpoint myself. In a team, governance decides who holds which. Governance sets the owners. On this page, they meet the content engine's checkpoints.

Pick. The editorial conference picks the topics, once a week. The content engine's proposals collect over the week and go into the conference together with the topics from the business units, news monitoring and the quarterly focal topics. All of them are measured against the same objectives (page 12). When a topic cannot wait for the next conference, the chief on duty decides.

Approve. Approval is split by risk (page 20). Subject matter experts approve facts and complex content. The communications lead approves voice and brand. The CEO approves quotes. The chief on duty gives the final release before anything goes online. The reviewers are named when the piece is planned, and the goal is to approve within 48 hours, for pieces that are ready a week ahead (page 19).

Post. After the chief on duty's release, the owner of the channel publishes, because they know its timing and its audience; sometimes the chief on duty publishes directly (page 39).

Read. In its weekly look back, the editorial conference, led by the chief on duty, discusses the pieces that worked especially well or poorly (page 26). The team decides what they mean and whether the focus changes.

The context has owners too (page 35). In addition, the chief on duty checks the knowledge base (page 32) and keeps the learnings log, updated in the weekly look back.

Seven decisions are never automated.

Some decisions stay with people, whatever the tools can do. In this model, seven are never automated:

- deciding what to say (the strategy team for the quarterly focal topics; the editorial conference for the week; the chief on duty for ad-hoc topics);
- approving facts, voice and quotes (subject matter experts, the communications lead, the CEO), and the final release (the chief on duty);
- publishing (the channel owner, sometimes the chief on duty);
- changing published content (the chief on duty or the marketing lead, page 22);
- interpreting the results (the editorial conference, led by the chief on duty, in its weekly look back);
- assigning work to people (the chief on duty);
- changing the rules in the context, such as the dos and don'ts or the list of verified facts (the owner of each part of the context).

Each of them has a named owner. That matters more than it seems.

People tend to over-rely on automated suggestions. A systematic review of 74 studies, mostly from clinical decision support, found that this automation bias is influenced by experience, trust, workload, task complexity and time pressure. Among the measures that reduce it is emphasising the user's accountability.[54] The studies come from healthcare, so the transfer to marketing is an analogy. But the pattern is familiar: a checkpoint that belongs to everyone becomes a click-through.

A checkpoint is also a real decision only if it has more than one outcome. At every checkpoint, the owner can approve, change, send back, reject or park (page 7). A step where the only option is "OK" is not a checkpoint. It is a formality.

Each decision belongs to a role, and a team adopts this way of working best by shaping it (pages 44–45).

[54] Goddard, K., Roudsari, A. and Wyatt, J.C. (2012) 'Automation bias: a systematic review of frequency, effect mediators, and mitigators', *Journal of the American Medical Informatics Association*, 19(1), pp. 121–127. Available at: https://doi.org/10.1136/amiajnl-2011-000089 (Accessed: 2 Nov 2026).



The team runs the system, not the other way round.

The content engine changes people's work. For every role, the share of writing from a blank page goes down and the share of deciding goes up. So the first question for a team is: who decides what? The choice of tool comes after.

- **Chief on duty:** runs the weekly editorial conference, decides on ad-hoc topics, gives the final release before anything goes online, keeps the weekly plan and the workload view.
- **Communications lead:** approves voice and brand; owns voice, key messages and positioning in the context.
- **Marketing lead:** decides, like the chief on duty, on changes to published content (page 22).
- **Subject matter experts:** bring topics from the business; approve facts and complex content; own the list of verified facts.
- **Strategy team:** sets the quarterly focal topics from the company objectives (page 12).
- **CEO:** approves quotes. Leadership also sets the company objectives that the planning cascade starts from.
- **Channel owners:** publish, know the timing of their channels and bring the results into the weekly look back.
- **Content and design team:** turns approved drafts into finished pieces. Its capacity is part of the planning.
- **System owner:** looks after the content engine itself: its inputs, its rules and what happens when a step fails. Every automation needs an owner, a clear input and a stop rule.

With the content engine, approvals become planned work with a deadline, not interruptions. That load belongs in the workload view like any other task. Otherwise the people who approve most, often subject matter experts whose main job lies elsewhere, carry an extra load that nobody sees.

GLOSSARY

STOP RULE

A set condition under which an automation stops and hands over to a person, for example when an input is missing or a check fails.

Teams adopt a shared system they understand and can shape.



IN PRACTICE

A word on what this page is based on. At Purpose Green, the shared workspace and the governance ran with the whole team. The content engine came later and so far runs for my own content. What follows on adoption is what I learned when the team took on the workspace and the governance, and what I recommend for adding the content engine.

One database, many views

Not everyone took to the shared workspace straight away. For many in my team, the tool was new and hard to understand. The biggest challenge at the start was that different people wanted to work with different views: a simple list, a table, a timeline chart.

It worked once I made one rule completely clear: everyone can adapt their own view, but everyone works in the same database. Nobody creates their own tables. Alongside the standard team view, I built a personal dashboard with each team member, one to one. Everyone had their own area, and everyone could see their individual contribution and accountability. From then on, it became a real collaboration.

Ways in for everyone

New team members, and in part people from outside the marketing team, were introduced through the basic idea: campaigns, with their linked tasks, content and design. Whoever understands that structure can work with the system.

We also built entry points for people who did not work in the system every day. A marketing briefing form created new projects directly in the database: a clear brief, built into the system. A separate view of the shared database showed everyone which topics were in focus, so unit leads always knew what marketing was working on.

For the content engine, I would carry over what worked: one shared set of rules, and everyone sees their own part in it. In practice: explain the problem and the idea before the tool. Show every person their checkpoint and their context store. Start with one content stream, not all at once. Adjust the rules in the weekly look back. People who are informed about a change and take part in the decisions are more open to it.[55]

The system can then be built with any tools that cover every function (pages 46–47).

[55] Wanberg, C.R. and Banas, J.T. (2000) 'Predictors and outcomes of openness to changes in a reorganizing workplace', *Journal of Applied Psychology*, 85(1), pp. 132–142. Available at: https://doi.org/10.1037/0021-9010.85.1.132 (Accessed: 2 Nov 2026).



shared workspace | knowledge base and context stores | AI model | automation layer | shared file store | results view | forms | notifications

Any stack will do, as long as every function is covered.

Everything in this playbook is a function, not a product. That is why the model can run on very different stacks. The question is not which tools to buy, but whether every function is covered. The names below are examples from the market, not recommendations.

- **Shared workspace** for campaigns, tasks, content, design and the editorial record: a database-style tool with linked records and personal views, such as Notion, Airtable, Monday or Asana.
- **Knowledge base and context stores:** the same workspace or a separate database, plus documents for voice, key messages, positioning, facts and personas.
- **AI model** for research, drafting and editing: any capable model, such as Gemini, Claude, ChatGPT or Mistral.
- **Automation layer** that moves data between the steps, starts the next step, checks outputs and reports errors: a workflow tool such as n8n, Make, Zapier or Power Automate. It follows rules; it does not need to be intelligent.
- **Shared file store** for the handover between steps, such as Google Drive, SharePoint or Dropbox.
- **Results view:** a spreadsheet or reporting tool, such as Google Sheets, Excel or Looker Studio.
- **Forms** for briefs and new sources.
- **Notifications** that tell an owner a checkpoint is waiting: team chat such as Slack or Microsoft Teams, or email.

IN PRACTICE

My own content engine runs on a shared workspace in Notion, which holds the editorial record and the knowledge base. Gemini does the research, evaluates the incoming newsletters and drafts the content. n8n is the automation layer: it moves the data between the steps, checks the outputs and reports errors. The handover files sit in Google Drive, the results in Google Sheets, and Slack tells me when a checkpoint is waiting for me. Any of these could be replaced by another tool from the same category without changing the model.

Connect the tools with clear rules, and set up the governance first.

1. Use each tool for what it does well
2. Define exactly what comes out of every AI step
3. Keep your data where it already is
4. Change automations with a fixed routine
5. Start with the smallest version that closes the loop

Building the content engine taught me five rules for choosing and connecting the tools.

Use each tool for what it does well. My first plan put the results dashboard into the workspace tool, which soon reached its limits. The dashboard moved to a spreadsheet, and the workspace now links to it.

Define exactly what comes out of every AI step. From the outside, an AI step can be a black box. A fixed output format with clear markers lets the automation check the result and lets a person read it.

Keep your data where it already is. In my content engine, the AI model runs inside the office suite where the working files live, so they stay there.

Change automations with a fixed routine: export, change, check, import. A convenient shortcut once silently dropped the access credentials of an automation.

Start with the smallest version that closes the loop, and add from there.

The order matters: governance first, tools last. Set up the editorial conference and the approval rules, then give the work one shared place, then fill the context, and only then add the automation, one station at a time. The first 90 days follow that path, one phase per month (pages 50–53).

First, the next two pages cover the pitfalls I met along the way, with people and with automation (pages 48–49).

Most pitfalls in a marketing OS are about people, not tools.



Every team that sets up a shared marketing OS meets similar pitfalls. Most of them are not technical problems. Here are four, each with a way to avoid it.

1. The tool comes before the rule. *The pitfall:* A new shared tool arrives, and everyone carries on as before, with their own lists, their own tables and their own way of seeing the work. The shared system turns into one more personal tool.[56] *How to avoid it:* Set one rule before the tool goes live: everyone may shape their own view, but everyone works in the same database. Explain the basic structure first: campaigns, linked to their tasks, content and design (page 45).

2. Work starts without a brief. *The pitfall:* A request arrives as a single sentence; the team starts, guesses and starts again. Duplicate work and last-minute requests follow. *How to avoid it:* Make the brief the entry point. No project starts without an objective, an audience, a key message, success criteria, an owner and a deadline. Poor briefs waste time and budget (page 6). A briefing form that creates the project directly in the shared workspace makes this the easy way (page 45).

3. Every approval goes to the top. *The pitfall:* Leadership approves everything, including facts and wording. Complex content waits, and leadership becomes the bottleneck. *How to avoid it:* Give each type of risk its own owner: facts to subject matter experts, voice and brand to the communications lead, quotes to the CEO. Keep leadership sign-off for what only leadership can say (page 21).

4. Shared teams are planned as if they had unlimited time. *The pitfall:* Design and other shared teams serve several stakeholders. When each stakeholder plans on its own, the shared team ends up overloaded. *How to avoid it:* Put the shared team's capacity into the weekly planning, and look at it again in every weekly look back (pages 25–26).

[56] Orlikowski, W.J. (1992) *Learning from Notes: organizational issues in groupware implementation*. Working Paper 3428-92. Cambridge, MA: MIT Sloan School of Management. Available at: https://dspace.mit.edu/handle/1721.1/175292 (Accessed: 2 Nov 2026).

The content engine's pitfalls start when people trust automation too much.



The content engine brings its own pitfalls. Most of them share one root: trusting automation more than it deserves.

1. Optimising for output, not impact. *The pitfall:* The content engine produces more ideas and more drafts, and it feels like progress. But nobody can say what landed. *How to avoid it:* Build measurement into the loop before automating more, and let a person read the results (page 40).

2. Trusting the AI's sources. *The pitfall:* AI research finds what is loud, not what is true: marketing blogs, vendor pages, claims repeated until they look like facts. *How to avoid it:* Build source rules into the process, for example at least two independent organisations per topic, and keep a person as the final check (page 36).

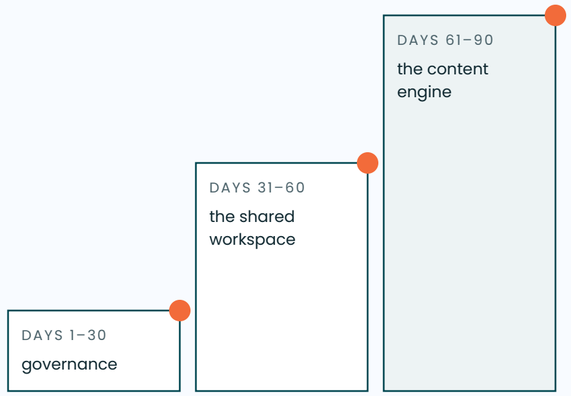
3. Automations without an owner or a stop rule. *The pitfall:* An automation runs unattended and does something nobody intended. *How to avoid it:* Give every automation an owner, a clear input and a stop rule (page 44).

IN PRACTICE

One of my AI steps was set to watch the wrong folder. It read its own output as new input and ran in an endless loop. Since then, input and output have been kept strictly apart, and a guard step checks what comes in.

4. Automating the last step. *The pitfall:* Because it is possible, the final step is automated as well, and posts go out on a schedule that nobody looks at again. *How to avoid it:* Keep the last look human. Not everything that can be automated should be (page 39).

Start with how you decide, and automate last.



each phase ends with a milestone

There is no big bang. The order of the steps is the message of this playbook: decide how you decide before you automate anything. A team that starts with the tools carries old habits into them and trusts automation too much (pages 48–49). A team that starts with the governance gets a way of working that the tools can later support.

IN PRACTICE

At Purpose Green, it took about six weeks from presenting the model to a weekly routine that worked. I presented the model and drove the implementation; the team shaped the rules and built their personal dashboard views with me (page 45).

In another team, this does not have to be the work of one person. The more of the team takes part in building the system, the more of the team will carry it.

The path has three phases of about a month each:

- **Days 1–30: governance.** How the team decides what to say, who decides and when.
- **Days 31–60: the shared workspace.** Where the work becomes visible.
- **Days 61–90: the content engine, one station at a time.** What the machine prepares.

Each phase ends with a milestone that the team can check, not just a date in a project plan. If a milestone is not reached, the next phase waits. Moving on before the team works with what is there only adds a new layer to an unfinished one.

Three habits hold through all phases: involve the whole team; start small and close the loop before adding more; and look back every week (page 26).

Days 1–30: build the governance together with the team.

Start with the problem, not with the model. Ask the team where the work gets stuck today: duplicate work, last-minute requests, approvals that wait (page 6). Then show the model as one possible answer, problem and solution side by side, and develop the marketing OS together. Every voice matters and should be heard. The people who will work in the system know best where it hurts, and people who take part in a change are more open to it.[57]

Then set the frame:

- Derive the topic islands and their content pillars from the company objectives, and let the strategy team set the quarterly focal topics for the next quarter (page 12).
- Start the weekly editorial conference with the whole team, including the look back at the previous week (pages 25–26).
- Assign approval owners by type of risk, name the chief on duty for the final release, and agree on a deadline for approvals (pages 19–22).
- Write down the standards: verified facts, key messages and tone of voice (pages 16–18; templates in Appendix A3, A4), with a first version of the personas and the positioning brief (Appendix A1, A2).
- Agree who decides on changes after publication (page 22).
- Agree the AI rules with the team (pages 23–24).
- Make the brief the entry point: no work starts without an objective, an audience, a key message, success criteria, an owner and a deadline (page 48).

Who is involved: the whole marketing team; leadership for the company objectives and the go-ahead; subject matter experts for their role in approval.

What to watch for: the temptation to start with a tool. In this phase, the rules matter more than the software. A shared document for the plan and a fixed slot for the weekly conference are enough to begin with.

Milestone: the first weekly plan that everyone knows: who delivers what, by when.

Check question: can every team member say which focal topic their current work serves?

[57] Wanberg, C.R. and Banas, J.T. (2000) 'Predictors and outcomes of openness to changes in a reorganizing workplace', *Journal of Applied Psychology*, 85(1), pp. 132–142. Available at: https://doi.org/10.1037/0021-9010.85.1.132 (Accessed: 2 Nov 2026).

Days 31–60: make all the work visible in one place.

Once the team decides together, give it one place to work. Set up one database, with campaigns linked to their tasks, content and design (page 9). Agree on one rule before anyone starts: everyone may shape their own view, but everyone works in the same database, and nobody creates their own tables.

Then build three things on top of it:

- **The workload view,** so that the editorial conference sets priorities against real capacity, including that of shared teams such as design (pages 25, 48). A person, not the system, assigns the work.
- **Personal views or dashboards,** built with each team member, so that everyone sees their own area and their own contribution (page 45).
- **Entry points for people outside the team:** a briefing form that creates a new project directly, and a view of the topics in focus, so that unit leads know what marketing is working on (page 45).

Who is involved: the whole team, each person for their own view; one person who sets up and looks after the structure.

What to watch for: old habits in the new tool. If a team member keeps a separate list "just for now", the shared system is already losing ground (page 48).

Milestone: every piece of work in the team is visible in one place, and the weekly conference plans from it.

Check question: does anyone still keep work outside the shared database?

Days 61–90: add the content engine one station at a time.

Start with the context, not the automation. Before the content engine drafts anything, fill the four stores (page 31). Most of the context already exists. From phase 1, the tone of voice and the key message map go into the voice-and-key-messages store, the verified facts into the knowledge base and facts. Refine the personas and the positioning brief with what the team has learned since phase 1. Add what is missing: a first knowledge base, and an empty learnings log that fills once the first results come in. Start the knowledge base with what the team already reads: the newsletters it subscribes to and the articles it shares with each other (page 32).

Then add the stations in the order of the loop:

- **Research and pick first.** Let the content engine propose topics, and take its proposals into the weekly editorial conference (pages 36, 42–43). This first station already shows whether the context is good enough.
- **Production and approval next,** with the owners from phase 1 (page 38).
- **Posting stays manual** (page 39).
- **Measurement last in the loop,** feeding the weekly look back (page 40). Before you automate further, use what it shows to fine-tune the strategy, so the content engine does not run faster in the wrong direction.

Name a system owner before the first automation runs: someone who knows its inputs and its stop rule (page 44).

Who is involved: the system owner; the communications lead and the subject matter experts for the context; the whole team in the weekly editorial conference.

What to watch for: the temptation to automate everything at once. One station at a time keeps errors visible; a step that fails unnoticed repeats its mistake (page 49).

Milestone: the first piece that has gone through the whole loop, from proposal to result, with a person at every checkpoint.

Check question: can the team name, for every station, who decides at its checkpoint?

After 90 days. The work is not finished; it has started. Read the results in the weekly look back, and measure what lands before you automate more (pages 40, 49). Adjust the rules where they do not fit. Add the next content stream only when the first one runs.

And keep the whole team involved. The system belongs to the people who use it.

The principles sit at the level of the marketing OS, not in a single module (page 10). That is why the governance stays the same when the tools change.

A persona keeps the whole team writing for the same reader.

HOW TO USE THIS TEMPLATE

A persona describes one typical person in your target group as if you knew them. Fill in one page per persona; most teams need two to four. Use what you already have: your team's experience, customer conversations, notes from sales and service, existing records and collected feedback. You do not need a large research budget, but you do need to say what each persona is based on, and to revise it when you learn more (page 34). Store the finished personas in the audience store of the context, so that every brief and every draft can use them.

PERSONA

Name and role in one line (e.g. "the portfolio manager")

Primary or secondary target group

◯ primary ◯ secondary

Picture (photo or illustration, optional) [photo]

PERSONAL DETAILS (to get a feel for the person)

Age range

Place (city, region)

Life situation

Education and background

WORK

Job title and type of organisation

Responsibilities and the decisions they own

Short description: who they are and what their working day looks like

Characteristics of this group

Examples of organisations in this group

BASED ON

Sources for this persona (tick all that apply): team experience · customer conversations · sales and service notes · existing records · collected feedback · surveys or research

What the media ratings are based on

Owner of this persona

Last updated

WHAT DRIVES THEM

Questions they ask in their everyday work

Pain points

Goals and wishes in this context

What convinces them (proof, sources, senders they trust)

Four words that describe them (e.g. analytical, efficiency-driven)

MEDIA

MEDIA TYPE	PRIMARY OR SECONDARY	RELEVANCE (1–5)
Daily news and online news		◇ ◇ ◇ ◇ ◇
Trade press		◇ ◇ ◇ ◇ ◇
Social media (which platforms)		◇ ◇ ◇ ◇ ◇
Newsletters		◇ ◇ ◇ ◇ ◇
Podcasts and video		◇ ◇ ◇ ◇ ◇
Events		◇ ◇ ◇ ◇ ◇
Other		◇ ◇ ◇ ◇ ◇

A positioning brief says what you stand for, and what you avoid.

HOW TO USE THIS TEMPLATE

The positioning brief answers one question: what do we want the people we write for to remember about us? Keep it to one page. Write it together with the people who will use it, and let the communications lead own it (pages 35, 42). Every brief and every draft is checked against it: does this piece fit what we stand for? Revise it when the company's objectives or offers change, not every time a campaign starts.

POSITIONING BRIEF

For whom (link the personas from A1)

Proof for that (facts, examples, references; link to the verified facts list)

The problem we solve for them

Core message (link to the key message map, A4)

What sets us apart

What we do not say (claims, words, topics to avoid)

Owner

Last updated

A tone of voice sheet shows how the brand sounds.

HOW TO USE THIS TEMPLATE

The tone of voice sheet describes how the company sounds, whoever writes and whichever channel carries the piece. Keep it to one page. Choose four axes that matter for your content: start from the four axes from Nielsen Norman Group (page 17), keep what fits and replace the rest. The table shows one example set. Place the brand on each axis, then make it concrete with words, dos and don'ts, and example sentences. The voice stays the same everywhere; the tone may shift a little by channel, and the sheet says how far. Let the communications lead own it, store it in the shared context, and check every draft against it, whether a person or an AI model wrote it (page 17). This page is a starting point. It can be extended, for example in a shared spreadsheet such as Google Sheets or a similar tool, to hold more dos and don'ts, longer word lists or example sentences for each channel.

THE BRAND IN THREE WORDS

Three words that describe how we sound

What we never want to sound like

FOUR AXES (example set; mark one point on each scale)

POSITION			WHY
Formal ↔ casual			
Formal	○ 1 ○ 2 ○ 3 ○ 4 ○ 5	casual	
Technical ↔ plain			
Technical	○ 1 ○ 2 ○ 3 ○ 4 ○ 5	plain	
Concise ↔ detailed			
Concise	○ 1 ○ 2 ○ 3 ○ 4 ○ 5	detailed	
Matter-of-fact ↔ enthusiastic			
Matter-of-fact	○ 1 ○ 2 ○ 3 ○ 4 ○ 5	enthusiastic	

IN PRACTICE

Dos (three to five)

Don'ts (three to five)

Words and phrases we use

Words and phrases we avoid

Example sentence: before

Example sentence: after

How far the tone may shift by channel (e.g. social media, press, website)

OWNERSHIP

Based on (brand guide, positioning brief A2)

Owner

Last updated

A key message map keeps every piece on the same core message.

HOW TO USE THIS TEMPLATE

The key message map holds what the company says again and again. Start in the middle and work outwards. First, write the core message in two lines: what the company is and offers, and how customers benefit from that. Then add the offering pillars above it, one for each main offer, each with a few concrete messages. Below it, add the benefit pillars: what the offer changes for the customer, for their community or industry, and for society. Give every message its proof, and link it to the verified facts. The map builds on the positioning brief (A2) and the personas (A1). The communications lead owns it; the whole team, and other units, use it (page 18).

OFFERING PILLARS (what we offer; two to four pillars)

Pillar 1	Pillar 2	Pillar 3
MESSAGE	MESSAGE	MESSAGE
PROOF	PROOF	PROOF

CORE MESSAGE

What we are and offer: "[Company] is …"	
Why it matters: "… so that [customer] can …"	

BENEFIT PILLARS (why it matters; one row per level)

For the customer	For their community or industry	For society
BENEFIT THEME	BENEFIT THEME	BENEFIT THEME
SUPPORTING MESSAGES	SUPPORTING MESSAGES	SUPPORTING MESSAGES
PROOF	PROOF	PROOF

A key message map keeps every piece on the same core message.

USE

Which messages matter most for each persona (A1)

What we do not say (from A2)

Owner

Last updated

Every source in this playbook is listed in alphabetical order.

This list contains every external source cited in the footnotes, in the Harvard style (Cite Them Right). Each online source was opened and checked on the date given in brackets. Where a publisher's page could not be opened, the reference gives the DOI, the permanent identifier of the publication.

Anderson, M.M. and Fort, K. (2022) 'Human where? A new scale defining human involvement in technology communities from an ethical standpoint', *International Review of Information Ethics*, 31(1). Available at: https://informationethics.ca/index.php/irie/article/download/477/449/699 (Accessed: 2 Nov 2026).

Bakker, A.B. and Demerouti, E. (2007) 'The Job Demands-Resources model: state of the art', *Journal of Managerial Psychology*, 22(3), pp. 309–328. Available at: https://doi.org/10.1108/02683940710733115 (Accessed: 2 Nov 2026).

Barnett, P. (2015) 'Breaking bad barriers', *strategy+business*, 20 Oct. Available at: https://www.strategy-business.com/article/00382 (Accessed: 2 Nov 2026).

Becker, N. (2026) 'Der Standardprozess im Corporate Newsroom', *Mediamoss-Blog*, 21 May. Available at: https://mediamossblog.com/2026/05/21/der-standardprozess-im-corporate-newsroom/ (Accessed: 2 Nov 2026).

Buçinca, Z., Malaya, M.B. and Gajos, K.Z. (2021) 'To trust or to think: cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making', *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), pp. 1–21. Available at: https://doi.org/10.1145/3449287 (Accessed: 2 Nov 2026).

Center for Creative Leadership (no date) *How leadership works: connecting exponential potential with collective progress* [white paper]. Available at: https://ccl.org/wp-content/uploads/2024/03/how-leadership-works-DAC-ccl-white-paper.pdf (Accessed: 2 Nov 2026).

Cross, R., Rebele, R. and Grant, A. (2016) 'Collaborative overload', *Harvard Business Review*, 94(1), pp. 74–79. Available at: https://hbr.org/2016/01/collaborative-overload (Accessed: 2 Nov 2026).

Dechêne, A., Stahl, C., Hansen, J. and Wänke, M. (2010) 'The truth about the truth: a meta-analytic review of the truth effect', *Personality and Social Psychology Review*, 14(2), pp. 238–257. Available at: https://doi.org/10.1177/1088868309352251 (Accessed: 2 Nov 2026).

Dell'Acqua, F., McFowland III, E., Mollick, E., Lifshitz, H., Kellogg, K.C., Rajendran, S., Krayer, L., Candelon, F. and Lakhani, K.R. (2026) 'Navigating the jagged technological frontier: field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality', *Organization Science*, 37(2), pp. 403–423. Available at: https://www.hbs.edu/ris/Publication%20Files/dell-acqua-et-al-2026-navigating-the-jagged-technological-frontier_5c589c8c-fbb5-458f-b285-c944746cd717.pdf (Accessed: 2 Nov 2026).

Deutscher Presserat (2025) *Publizistische Grundsätze (Pressekodex)*. Version of 19 Mar 2025. Available at: https://www.presserat.de/pressekodex.html (Accessed: 2 Nov 2026).

Díaz, S., Kern, C. and Olive, K. (2025) *Google's approach for secure AI agents: an introduction*. Google. Available at: https://services.google.com/fh/files/misc/googles_approach_for_secure_ai_agents.pdf (Accessed: 2 Nov 2026).

Edelman (2019) *2019 Edelman Trust Barometer Special Report: In Brands We Trust?* Available at: https://www.edelman.com/sites/g/files/aatuss191/files/2019-06/2019_edelman_trust_barometer_special_report_in_brands_we_trust.pdf (Accessed: 1 Nov 2026).

Edmondson, A. (1999) 'Psychological safety and learning behavior in work teams', *Administrative Science Quarterly*, 44(2), pp. 350–383. Available at: https://dash.harvard.edu/entities/publication/13a7b031-0fdd-45ec-a7e0-2b80e2bc679f (Accessed: 2 Nov 2026).

Gilboa, S., Shirom, A., Fried, Y. and Cooper, C. (2008) 'A meta-analysis of work demand stressors and job performance: examining main and moderating effects', *Personnel Psychology*, 61(2), pp. 227–271. Available at: https://eprints.lancs.ac.uk/id/eprint/45119 (Accessed: 2 Nov 2026).

Goddard, K., Roudsari, A. and Wyatt, J.C. (2012) 'Automation bias: a systematic review of frequency, effect mediators, and mitigators', *Journal of the American Medical Informatics Association*, 19(1), pp. 121–127. Available at: https://doi.org/10.1136/amiajnl-2011-000089 (Accessed: 2 Nov 2026).

Google (2025) *Search Engine Optimization (SEO) Starter Guide*. Google Search Central, last updated 10 Dec 2025. Available at: https://developers.google.com/search/docs/fundamentals/seo-starter-guide (Accessed: 2 Nov 2026).

Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B. and Liu, T. (2025) 'A survey on hallucination in large language models: principles, taxonomy, challenges, and open questions', *ACM Transactions on Information Systems*, 43(2), pp. 1–55. Available at: https://arxiv.org/abs/2311.05232 (Accessed: 2 Nov 2026).

Kleingeld, A., van Mierlo, H. and Arends, L. (2011) 'The effect of goal setting on group performance: a meta-analysis', *Journal of Applied Psychology*, 96(6), pp. 1289–1304. Available at: https://doi.org/10.1037/a0024315 (Accessed: 2 Nov 2026).

Lee, A., Willis, S. and Tian, A.W. (2018) 'Empowering leadership: a meta-analytic examination of incremental contribution, mediation, and moderation', *Journal of Organizational Behavior*, 39(3), pp. 306–325. Available at: https://doi.org/10.1002/job.2220 (Accessed: 2 Nov 2026).

Lerner, J.S. and Tetlock, P.E. (1999) 'Accounting for the effects of accountability', *Psychological Bulletin*, 125(2), pp. 255–275. Available at: https://doi.org/10.1037/0033-2909.125.2.255 (Accessed: 2 Nov 2026).

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S. and Kiela, D. (2020) 'Retrieval-augmented generation for knowledge-intensive NLP tasks', in *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*. Available at: https://arxiv.org/abs/2005.11401 (Accessed: 2 Nov 2026).

Liu, N.F., Zhang, T. and Liang, P. (2023) 'Evaluating verifiability in generative search engines', in *Findings of the Association for Computational Linguistics: EMNLP 2023*. Singapore: Association for Computational Linguistics, pp. 7001–7025. Available at: https://doi.org/10.18653/v1/2023.findings-emnlp.467 (Accessed: 2 Nov 2026).

Locke, E.A. and Latham, G.P. (2002) 'Building a practically useful theory of goal setting and task motivation: a 35-year odyssey', *American Psychologist*, 57(9), pp. 705–717. Available at: https://doi.org/10.1037/0003-066X.57.9.705 (Accessed: 2 Nov 2026).

Madigan, M. (2021) 'The BetterBriefs Project finds 33% of marketing budgets wasted due to poor briefs', *B&T*, 13 Oct. Available at: https://www.bandt.com.au/the-betterbriefs-project-finds-33-of-marketing-budgets-wasted-due-to-poor-briefs/ (Accessed: 2 Nov 2026).

Mayer, R.C., Davis, J.H. and Schoorman, F.D. (1995) 'An integrative model of organizational trust', *Academy of Management Review*, 20(3), pp. 709–734. Available at: https://doi.org/10.5465/amr.1995.9508080335 (Accessed: 2 Nov 2026).

Mediamoss (no date) *Der Newsroom beginnt im Kopf*. Available at: https://www.mediamoss.me/das-newsroom-modell (Accessed: 2 Nov 2026).

Moran, K. (2016a) 'The four dimensions of tone of voice', *Nielsen Norman Group*, 17 Jul. Available at: https://www.nngroup.com/articles/tone-of-voice-dimensions/ (Accessed: 2 Nov 2026).

Moran, K. (2016b) 'The impact of tone of voice on users' brand perception', *Nielsen Norman Group*, 7 Aug. Available at: https://www.nngroup.com/articles/tone-voice-users/ (Accessed: 2 Nov 2026).

Moss, C. (no date) 'Corporate Newsroom: Zuerst das Thema, dann der Kanal', *GABAL Magazin*. Available at: https://www.gabal-magazin.de/marketing/corporate-newsroom/ (Accessed: 2 Nov 2026).

NASA (2005) *Launch Services Program Pre-Launch Readiness Reviews*. NASA Policy Directive NPD 8610.24C, revalidated 2 Apr 2025 with Change 3. Washington, DC: NASA. Available at: https://explorers.larc.nasa.gov/APSMEX26/SMEX/pdf_files/N_PD_8610_024C_chg3.pdf (Accessed: 3 Nov 2026).

National Institute of Standards and Technology (no date) 'Separation of duty', *Computer Security Resource Center glossary*. Available at: https://csrc.nist.gov/glossary/term/separation_of_duty (Accessed: 2 Nov 2026).

Noy, S. and Zhang, W. (2023) 'Experimental evidence on the productivity effects of generative artificial intelligence', *Science*, 381(6654), pp. 187–192. Available at: https://doi.org/10.1126/science.adh2586 (Accessed: 2 Nov 2026).

Orlikowski, W.J. (1992) *Learning from Notes: organizational issues in groupware implementation*. Working Paper 3428-92. Cambridge, MA: MIT Sloan School of Management. Available at: https://dspace.mit.edu/handle/1721.1/175292 (Accessed: 2 Nov 2026).

Parasuraman, R., Sheridan, T.B. and Wickens, C.D. (2000) 'A model for types and levels of human interaction with automation', *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans*, 30(3), pp. 286–297. Available at: https://doi.org/10.1109/3468.844354 (Accessed: 2 Nov 2026).

Project Management Institute (2024) *Maximizing project success*. Newtown Square, PA: Project Management Institute. Available at: https://www.pmi.org/-/media/pmi/documents/public/pdf/learning/thought-leadership/project%5Fsuccess%5Freport%5F2024.pdf (Accessed: 2 Nov 2026).

Pruitt, J. and Grudin, J. (2003) 'Personas: practice and theory', in *Proceedings of the 2003 Conference on Designing for User Experiences (DUX '03)*. New York: ACM, pp. 1–15. Available at: https://www.microsoft.com/en-us/research/wp-content/uploads/2017/01/personas-practice-and-theory.pdf (Accessed: 2 Nov 2026).

Raabe, S. (2023) 'Offen, verantwortungsvoll und transparent – Die Guidelines der dpa für Künstliche Intelligenz', *dpa Innovation*, 3 Apr, updated 10 Jul 2024. Available at: https://innovation.dpa.com/2023/04/03/kuenstliche-intelligenz-fuenf-guidelines-der-dpa/ (Accessed: 2 Nov 2026).

Radical Candor (no date) *Our approach: Kim Scott's feedback framework*. Available at: https://www.radicalcandor.com/our-approach/ (Accessed: 2 Nov 2026).

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), Article 50(4). *Official Journal of the European Union*, L series, 2024/1689, 12 Jul 2024. Available at: http://data.europa.eu/eli/reg/2024/1689/oj (Accessed: 29 Oct 2026).

re:Work Editorial Team (2015) 'Understand team effectiveness', *re:Work*. Google. Available at: https://rework.withgoogle.com/intl/en/guides/understand-team-effectiveness (Accessed: 2 Nov 2026).

References

Rogelberg, S.G., Leach, D.J., Warr, P.B. and Burnfield, J.L. (2006) '"Not another meeting!" Are meeting time demands related to employee well-being?', *Journal of Applied Psychology*, 91(1), pp. 83–96. Available at: https://doi.org/10.1037/0021-9010.91.1.83 (Accessed: 2 Nov 2026).

Rogers, P. and Blenko, M.W. (2006) 'Who has the D? How clear decision roles enhance organizational performance', *Harvard Business Review*, Jan. Available at: https://hbr.org/2006/01/who-has-the-d-how-clear-decision-roles-enhance-organizational-performance (Accessed: 2 Nov 2026).

Seibert, S.E., Wang, G. and Courtright, S.H. (2011) 'Antecedents and consequences of psychological and team empowerment in organizations: a meta-analytic review', *Journal of Applied Psychology*, 96(5), pp. 981–1003. Available at: https://doi.org/10.1037/a0022676 (Accessed: 2 Nov 2026).

Simpson, J. (2024) 'What does it mean to go AI-first?', *Nordic APIs*, 12 Jun. Available at: https://nordicapis.com/what-does-it-mean-to-go-ai-first/ (Accessed: 2 Nov 2026).

Spin Sucks (no date) *The PESO Model®: the MarComm operating system*. Available at: https://spinsucks.com/peso-model/ (Accessed: 2 Nov 2026).

Stahl, S. (2024) 'B2B content marketing benchmarks, budgets, and trends: outlook for 2025', *Content Marketing Institute*, 9 Oct. Available at: https://contentmarketinginstitute.com/b2b-research/b2b-content-marketing-trends-research-2025 (Accessed: 2 Nov 2026).

Stephen, A.T. and Galak, J. (2012) 'The effects of traditional and social earned media on sales: a study of a microlending marketplace', *Journal of Marketing Research*, 49(5), pp. 624–639. Available at: https://doi.org/10.1509/jmr.09.0401 (Accessed: 2 Nov 2026).

Tannenbaum, S.I. and Cerasoli, C.P. (2013) 'Do team and individual debriefs enhance performance? A meta-analysis', *Human Factors*, 55(1), pp. 231–245. Available at: https://doi.org/10.1177/0018720812448394 (Accessed: 2 Nov 2026).

Troy, L.C., Hirunyawipada, T. and Paswan, A.K. (2008) 'Cross-functional integration and new product success: an empirical investigation of the findings', *Journal of Marketing*, 72(6), pp. 132–146. Available at: https://doi.org/10.1509/jmkg.72.6.132 (Accessed: 2 Nov 2026).

Vaccaro, M., Almaatouq, A. and Malone, T. (2024) 'When combinations of humans and AI are useful: a systematic review and meta-analysis', *Nature Human Behaviour*, 8(12), pp. 2293–2303. Available at: https://doi.org/10.1038/s41562-024-02024-1 (Accessed: 2 Nov 2026).

Wanberg, C.R. and Banas, J.T. (2000) 'Predictors and outcomes of openness to changes in a reorganizing workplace', *Journal of Applied Psychology*, 85(1), pp. 132–142. Available at: https://doi.org/10.1037/0021-9010.85.1.132 (Accessed: 2 Nov 2026).

Wheatley, M. (2016) 'Google's Sundar Pichai eyes a shift from "mobile first" to "AI first"', *SiliconANGLE*, 29 Apr. Available at: https://siliconangle.com/2016/04/29/googles-sundar-pichai-eyes-a-shift-from-mobile-first-to-ai-first/ (Accessed: 2 Nov 2026).

Leadership comes from direction, alignment and commitment.



Sebastian Hesse
Marketing Communications & Campaign Architect

What 15 years taught me

I have worked in marketing communications and PR for more than 15 years: at Weber Shandwick, in my own practice Big Bang & Whisper, and in corporate and campaign leadership. My work spans fintech, climate tech and proptech, new mobility and agency work. The themes I keep coming back to are integrated campaigns, the marketing OS and, in recent years, AI in marketing operations.

Across those years, one lesson kept recurring. Projects rarely fail for lack of ideas or effort. They fail because nobody wrote a clear brief, or because nobody managed the process around the work (pages 5–6). That is a leadership problem before it is a tool problem, and it is why this playbook starts with governance.

At Purpose Green, I introduced the corporate newsroom model and the shared marketing OS with a team of ten, as communications lead and chief on duty. For half a year, I also led the marketing team as interim marketing lead. Today I run the content engine described in this playbook for my own content.

Leadership and the system

My answer to that lesson is a leadership style that is clear and empathetic at the same time. It is close to what Kim Scott calls Radical Candor: caring personally and challenging directly, at once.[58] Clear means a brief, a goal and honest feedback, given early. Empathetic means that people can work on their own without losing sight of the goal, and dare to raise a problem before it grows. That takes psychological safety.[59]

To make this concrete, I use a simple formula from the Center for Creative Leadership: direction, alignment and commitment. Leadership shows in its outcomes: a group that agrees on its overall goals, coordinates its work and feels mutually responsible for it. Leaders set the tone and make space for this, but everyone helps shape all three.[60]

This is where a marketing OS helps. It supports all three. Direction, because every piece traces back to a focal topic and an overall marketing goal. Alignment, because everyone works from one plan and one set of rules. Commitment, because every checkpoint has a named owner who sees their own contribution. And when owners and deadlines are visible, nobody has to chase approvals, and a delay shows up as a fact, not as blame.

[58] Radical Candor (no date) *Our approach: Kim Scott's feedback framework*. Available at: https://www.radicalcandor.com/our-approach/ (Accessed: 2 Nov 2026).

[59] Edmondson, A. (1999) 'Psychological safety and learning behavior in work teams', *Administrative Science Quarterly*, 44(2), pp. 350–383. Available at: https://dash.harvard.edu/entities/publication/13a7b031-0fdd-45ec-a7e0-2b80e2bc679f (Accessed: 2 Nov 2026).

[60] Center for Creative Leadership (no date) *How leadership works: connecting exponential potential with collective progress* [white paper]. Available at: https://ccl.org/wp-content/uploads/2024/03/how-leadership-works-DAC-ccl-white-paper.pdf (Accessed: 2 Nov 2026).

CONTACT

linkedin.com/in/sebastianhesse
bigbangandwhisper.com