

The Content *System*

Frameworks, formats, checklists, and the lifecycle that carries an idea from a note to a published post

Built from the books that explain why things spread, the processes real agencies and media houses run on, and the numbers behind the eight platforms.

Read it in three layers, use it in one

Most content guides fail in the same way. They give you a list of post ideas and call it a strategy. This one is built in three layers, and each layer answers a different question.

- 01 **Strategy.** What are we here for, on which platforms, about which topics, and how often. Parts 1 and 2.
- 02 **Craft.** Why anything spreads, why anything is remembered, and how each format is actually built. Parts 3 and 4.
- 03 **Process.** How a piece of content moves from an idea to a published, measured, reused asset. Part 5.

Every framework here is attributed. When a claim comes from a book, the book is named. When a number comes from a study, the study, its year, and its sample are named. Where the popular version of a framework is wrong, the correction is on the page, because a playbook that repeats a misquote teaches the misquote.

A framework you cannot attribute is a slogan.

What this is not. It is not a list of trending audio. It is not a template pack. It is not a substitute for having something to say. Platform specs move fast, so treat every number in Part 4 as current for now and verify before you publish against it.

VERDICT Strategy decides. Craft persuades. Process ships.

Three questions, nine stages, one gate

The system runs on three questions asked in order. Skip one and the work downstream gets expensive.

- 01 Who is this for, and what should it move? Strategy sets the objective, the platforms, the pillars, and the cadence.
- 02 Why would anyone pass this on or remember it? Craft applies the share and stick frameworks, the hook, and the format.
- 03 How does it get made, checked, published, and reused? Process runs the nine-stage lifecycle with one quality gate before anything ships.

TABLE • THE NINE STAGES

#	STAGE	THE QUESTION IT ANSWERS
1	Intake	What are we being asked for, and why now?
2	Brief	What is the one idea, the one audience, the one action?
3	Ideation	What are the angles, and which one is not obvious?
4	Production	Who makes it, in what format, by when?
5	Review and QA	Does it pass voice, hook, clarity, and spec?
6	Publish	Right platform, right time, right link placement
7	Repurpose	What else does this one asset become?
8	Measure	Did it move the metric this stage owns?
9	Learn and archive	What do we keep, kill, and reuse?

VERDICT One gate, stage five, and nothing ships around it.

Two definitions of content, on purpose

People argue about content without defining it, so the argument never ends. Here are two definitions, announced as two.

For the non-marketer: content is anything you publish that earns attention you did not pay for. If you paid a platform to put it in front of someone, that is advertising. Content borrows attention on the strength of the thing itself.

For the practitioner: content is any published asset that moves a person one step along the funnel, from not knowing you exist, to knowing, to believing, to buying, to coming back. Every piece should be able to name the step it owns. A piece that cannot name its step is decoration.

If you cannot say which step it moves, it is not content. It is output.

Three terms used throughout, defined once. A **pillar** is a recurring topic the brand owns. A **content type** is one of the six purposes a post can serve. A **format** is the physical shape of the post: a reel, a carousel, a thread. Pillar is what you talk about. Type is why. Format is how.

VERDICT Pillar is what. Type is why. Format is how.

How content plans actually die

Here is the normal way. We open the accounts. We agree on three posts a week. Someone designs a template. We post the office photos, the quote cards, the product announcements, a holiday greeting for every occasion on the calendar. Three months in, we sit in a review and ask why nothing moved.

Nothing moved because no one wrote down what was supposed to move. In the 2026 Content Marketing Institute and MarketingProfs research, 97% of B2B marketers said they have a content strategy. Only 12% said they exceeded their goals. The single most-named challenge, at 40%, was creating content that prompts a desired action. Having a strategy and having a strategy that works are two different claims, and the survey lets people make the first one.

97%

say they have a content strategy

12%

say they exceeded their content goals

The same research found what actually moved the needle for teams whose results improved: 74% pointed to refining the strategy, ahead of 51% who pointed to new technology and 40% who pointed to restructuring the team. Note the honest limit on this: it is a self-reported survey of 1,015 marketers, sponsored by a software vendor, and it shows correlation, not proof.

VERDICT Nobody audits a plan nobody wrote down.

Pick the metric before you pick the platform

Strategy starts with a ranked goal, not a list of channels. Rank these four for the next two quarters: awareness, authority, leads, sales. Ranked, not weighted, because a tie means no decision was made.

Then give each funnel stage exactly one metric it owns. One, so that a debate about performance has a single answer.

TABLE • ONE METRIC PER STAGE

FUNNEL STAGE	WHAT IT DOES	THE ONE METRIC
Discovery	Reaches people who do not know you	Reach and new followers
Authority	Earns belief from people who do	Watch time, saves, profile visits
Relationship	Turns an audience into a community	Replies, DMs, community joins
Conversion	Asks for the action	Leads, signups, sales

Avinash Kaushik published the cleanest version of this thinking in 2013 as See, Think, Do, later extended to See, Think, Do, Care. His stages are defined by intent, not demographics: See is the largest qualified audience with no commercial intent, Think shows weak intent, Do shows strong intent, Care is the repeat customer. His KPI ladder climbs from impressions, to micro-conversions, to conversion rate, to lifetime value. The fourth stage was originally called Coddle, which tells you frameworks get edited too.

VERDICT Four stages. Four metrics. No shared credit.

Fewer platforms, each with a named job

Capacity is the constraint that decides this, not ambition. A one-person team cannot run eight platforms, and eight thin accounts lose to two good ones. Assign every platform you keep a single funnel role, and drop any platform you cannot name a role for.

TABLE • THE EIGHT PLATFORMS AND WHAT EACH IS FOR

PLATFORM	FUNNEL ROLE	PRIMARY REACH ENGINE	CADENCE SWEET SPOT
TikTok	Discovery	Short video, Photo Mode	2 to 5 per week
Instagram	Discovery and relationship	Reels for reach, carousels for engagement	3 to 5 Reels, 3 to 5 feed per week
YouTube	Authority	Long-form plus Shorts	1 long-form per week, 3 to 5 Shorts
LinkedIn	Authority	Text, document carousels, native video	2 to 5 per week
X	Relationship	Text, native video, replies	3 to 5 per day plus replies
Threads	Relationship	Text posts and replies	About 1 per day plus replies
Facebook	Community and local	Reels and Groups	1 to 2 per day
Pinterest	Discovery with buy intent	Fresh keyworded vertical Pins	5 to 10 fresh Pins per week

Two of these are structurally different. Pinterest is a search engine with a Pin lifespan measured in months, so its content compounds instead of decaying. Facebook's organic Page reach is structurally low, so Reels, Groups, or paid are the only real levers there.

Then name the owned channel. Social rents attention. An email list or a Group banks it. A strategy without an owned channel is a strategy that pays rent forever.

VERDICT Every platform needs a job or a goodbye.

Three to five topics you intend to own

Pillars are the recurring subjects the brand returns to. Three to five is the working range. Fewer and the account is one note. More and no algorithm ever learns what you are about, because platforms now read topical meaning rather than hashtags.

Marcus Sheridan's answer to what those topics should be is the most commercially useful one in print. In *They Ask, You Answer* (Wiley, 2017), he named five

categories buyers actually search for, drawn from his own pool company: cost and price, problems, versus and comparisons, reviews, and best in class. Most brands avoid all five, because all five feel like they help the buyer more than the seller. That is exactly why they work. Sheridan and IMPACT have since rebranded the system as Endless Customers, but the five topics are unchanged.

TABLE • THE BIG 5, AND WHAT EACH ONE IS

TOPIC	WHAT YOU PUBLISH
Cost and price	What it costs, and what drives the number up or down
Problems	The drawbacks and the risks, said plainly
Versus and comparisons	You against the alternatives, including doing nothing
Reviews	Balanced assessments, including who this is not for
Best in class	The honest best options in the category

The structural companion is the topic cluster, published by HubSpot Research in 2017: one broad pillar page as the hub, with cluster pieces linking back to it and to each other. Their sizing rule is useful even off-site: a pillar should be broad enough to sit over 20 to 30 pieces. Treat it as a structure, not as evidence. HubSpot published no sample size and sells the tool.

VERDICT Own five topics. Rent none of them.

Make one deep thing, then cut it up

The most common production failure is treating eight platforms as eight separate jobs. The fix is old and well documented. Make one deep asset on a rhythm, then derive everything else from it.

Google's own YouTube Creator Playbook for Brands named three tiers. Hero is the push moment, the launch or the tentpole, and a brand may have only a few in a year.

Hub is the regular content that gives a fresh angle on the audience's interests, staggered across the year. Hygiene is the always-on pull content: what your audience is actively searching for, the tutorials, the how-to, the answers. The industry now says Help instead of Hygiene, and Google itself drifted to that name, so both are in circulation.

Hero pushes hard. Hub pushes weekly. Hygiene pulls all year.

A correction worth carrying, because this error is everywhere. The 70/20/10 content split does not appear anywhere in that playbook, which contains no percentage allocation at all. It is Coca-Cola's, from Content 2020, presented by Jonathan Mildenhall: 70% low risk bread and butter, 20% innovating on something that worked, 10% brand-new ideas that could fail. Two useful frameworks, two different companies. Citing "Google's 70-20-10 Hero-Hub-Help model" fuses them and is wrong.

The operational version is Gary Vaynerchuk's Content Model, published in 2018 with the subtitle "how I make 30+ pieces of content from a single keynote." Six steps: establish pillar content, create micro content, distribute both, take community insights, make community-driven micro content, distribute a second round. Treat the 30-plus number as his own claim rather than independent research, and the model as sound.

VERDICT One hero a week beats seven orphans.

Post to your capacity, not to a chart

Brands are posting less and thinking more. In the 2026 Rival IQ and Quid benchmark report, the median brand posted 3.69 times a week on Instagram and 3.0 times a week on Facebook, both the lowest in six years. The sample is around 2,700 brands across 18 industries.

Frequency still matters, but not in the way people assume. Buffer's analysis of more than 52 million posts across 10 platforms

through December 2025 found two things at once: accounts that went silent for a week consistently underperformed their own baseline growth, and reach per post declines as frequency climbs. So consistency protects you and volume taxes you. Buffer's own caveat applies: this is Buffer users only, skewing to solo creators and small teams, not a full view of any platform.

Set the cadence you can hold for 12 weeks with the people you actually have. Then hold it. Under-commit on purpose, because a calendar nobody can staff is abandoned by week three, and abandonment costs more than a lower number would have.

3.69

median Instagram posts per week, the lowest in six years

CHECKLIST • STRATEGY GATE

- | | |
|--|--|
| 01 The goal is ranked, not listed, and the top goal is written in one sentence. | 04 Three to five pillars are named, and at least two of them are uncomfortable to publish. |
| 02 Every funnel stage has exactly one metric, and someone owns it. | 05 One hero asset is named, with a rhythm, and its derivatives are planned. |
| 03 Every platform kept has a named funnel role. Every platform without one is dropped. | 06 One owned channel is named for banking the audience. |
| | 07 The weekly cadence fits the hours the team actually has. |

VERDICT A cadence you cannot hold is a plan to quit.

Six purposes, and every post serves exactly one

Audiences follow accounts that give before they ask. Six types cover everything a brand can publish. Assign every idea exactly one primary type. If you cannot name the type, the idea is not sharp yet.

TABLE • THE SIX TYPES

TYPE	WHAT IT DOES	WHAT BELONGS IN IT
Entertainment	Earns attention and shares	Memes, trends, throwbacks, puzzles, contests, light native video
Inspiration	Makes people feel and believe	Quotes, striking facts and images, client and personal triumph stories
Education	Builds authority and trust	Tips, case studies, research, FAQs, free resources, live training
Conversation	Turns an audience into a community	Questions, polls, fill-in-the-blanks, ask-for-advice, caption this
Connection	Makes the brand human	Behind the scenes, previews, team and vendor features, real life
Promotion	Asks for the sale or the signup	Offers, webinars, testimonials, lead magnets, product posts

Conversation is the odd one out. It is a category, and it is also a layer that sits on top of the other five. Ending almost every post with a real question is a habit, and it is the cheapest lift on this page. Buffer's 2026 analysis found creators who replied to comments saw engagement lifts of 42% on Threads, 30% on LinkedIn, and 21% on Instagram. The reply is content too.

VERDICT Name the type, or the post has no job.

Four posts give, one post asks

The balance is where most brand feeds fail, in one of two directions: all promotion, or all education with nothing human in it. Start from these shares and tune them to the goal.

TABLE • DEFAULT MIX, AND WHEN TO TURN EACH ONE UP

TYPE	DEFAULT SHARE	TURN IT UP WHEN
Education	About 30%	Authority or search is the goal
Entertainment	About 20%	Growth and reach are the goal
Inspiration	About 15%	Building a movement or a personal brand
Connection	About 15%	Trust matters, or a launch is coming
Conversation	About 10%, plus a layer on everything	Engagement has gone flat
Promotion	About 10%	A live launch or an offer window

The rule that rarely bends: promotion stays a minority. Four posts of every five give value or build the relationship, and the fifth asks. During a launch you can push promotion higher for a window you declare in advance, then return to the baseline. A launch window with no end date is not a launch, it is a new personality.

Two failure modes to check for by name. Promotion above roughly 20% outside a declared window. And any type sitting at zero for weeks, which is usually Connection or Entertainment, because those are the two that feel least like work and get cut first.

VERDICT Promotion is a guest. It does not move in.

Most of your content is talking to people who will not buy this year

The strongest argument for weighting content toward the top of the funnel comes from Professor John Dawes at the Ehrenberg-Bass Institute, working with the LinkedIn B2B Institute. His estimate: up to 95% of business buyers are not in the market for a given category at any one time. The arithmetic is simple. If buyers replace a provider roughly every five years, about 20% are in market in a year and about 5% in a quarter.

You are mostly building memory, not demand.

The honest limit: 95-5 is modelled from category purchase cycles, not measured directly, and it fits long-cycle services like banking, software, and legal. Do not apply it to fast-moving consumer categories.

TABLE • WHICH TYPE SERVES WHICH STAGE

STAGE	TYPES THAT CARRY IT	WHERE IT LIVES
Discovery	Entertainment, Inspiration, broad Education	TikTok, Reels, Shorts, Pinterest
Authority	Education, Inspiration by proof	YouTube, LinkedIn
Relationship	Connection, Conversation	X, Threads, Stories, Groups
Conversion	Promotion, plus Education that de-risks the buy	Newsletter, Groups, owned channels

The oldest published map of this is the Content Marketing Matrix by Dave Chaffey and Dan Bosomworth at Smart Insights, 2012: awareness to purchase on one axis, emotional to rational on the other, with formats placed by both. It is still the clearest one-page picture of the idea.

VERDICT Build memory now. Harvest demand later.

Ten minutes a month, on the last 30 posts

Run this on the last 30 published posts, not on the plan. Plans are optimistic. Feeds are evidence.

CHECKLIST • MIX AUDIT

- 01 Tag each of the last 30 posts with exactly one of the six types. Count them.
- 02 Compare the counts against the target shares. Flag any gap over 10 points.
- 03 Is promotion above 20% with no declared launch window? Cut it back.
- 04 Is any type at zero? Name it and schedule two next week.
- 05 How many posts ended with a real question or an invitation to reply? If fewer than half, that is the cheapest fix on this page.
- 06 Are the pillars rotating, or has one pillar taken over three weeks in a row?
- 07 Which three posts got the most saves and sends? What type were they? Do more of that type, not more of that exact post.
- 08 Which three underperformed? Was it the idea, the hook, or the format? Write the answer down, because the answer is the input for next month.

VERDICT Audit the feed you have, not the plan you wrote.

Contagious gave sharing six named drivers

Jonah Berger's *Contagious* (Simon and Schuster, 2013) is the standard reference on why things travel. His framework is STEPPS, and his own gloss on each element is worth quoting rather than paraphrasing.

- 01 **Social currency.** People talk about things that make them look good rather than bad. Give the viewer a line they can repeat and take credit for.
- 02 **Triggers.** Top of mind, tip of tongue. Tie the idea to something the audience meets on a schedule, so the trigger replays the content for you.
- 03 **Emotion.** When we care, we share. High-arousal feeling travels, mild agreement does not.
- 04 **Public.** When we can see other people doing something, we imitate it. Show the visible artifact, not the private result.
- 05 **Practical value.** News you can use. One takeaway worth forwarding to a specific person, because the forward is the share.
- 06 **Stories.** Information travels inside narrative. The story is the carrier, the idea is the cargo.

Two working rules. Never engineer all six, because a piece optimised for everything spreads for nothing. Pick the two or three the topic naturally owns and build them into the beats rather than into extra length.

A small note on accuracy. The Wharton transcript of Berger's talk renders the E as "ease for emotion," a transcription slip that has spread into secondary blogs. The element is Emotion.

VERDICT Pick two triggers. Build them into the beats.

Arousal beats positivity, and the study has limits

The academic paper under Contagious is Berger and Milkman, "What Makes Online Content Viral," Journal of Marketing Research, April 2012, volume 49, issue 2, pages 192 to 205. Its finding is quoted loosely more often than it is quoted correctly.



Content that evokes high-arousal positive emotion, awe, or high-arousal negative emotion, anger or anxiety, is more viral. Content that evokes low-arousal emotion, such as sadness, is less viral.

Read what that says. The mechanism is arousal, not positivity. Anger and anxiety travel. Sadness does not, despite feeling stronger. Positive content is more viral than negative content overall, but valence alone does not explain it, and the popular summary "positive content goes viral" drops the actual finding.

Now the limit, which matters if you are going to cite it in a client deck. The dataset was 6,956 New York Times articles that appeared on the home page between August and November 2008, with making the most-emailed list as the outcome being measured. That is email forwarding of news over three months in 2008. It is not a study of TikTok, Instagram, or any feed algorithm, and citing it as one is a stretch the paper does not support.

TABLE • WHAT TO BUILD FOR, BY EMOTION

BUILD FOR

Awe, amusement, anger, anxiety, surprise

Feelings with energy behind them

AVOID BUILDING FOR

Sadness, mild contentment, vague positivity

Feelings that make people sit still

VERDICT Aim for energy, not for mood.

Made to Stick gave memory six tests

Chip Heath and Dan Heath published *Made to Stick* with Random House in 2007. Where Berger explains transmission, the Heaths explain retention. Their framework is SUCCEsS, six principles with a trailing lowercase s for spelling. Run your core claim and your closing line through all six.

- 01 **Simple.** One core idea, said in one line. Not dumbed down, stripped down.
- 02 **Unexpected.** Break a pattern the audience assumed. Surprise buys attention, mystery holds it.
- 03 **Concrete.** Sensory and specific. "Clay everywhere, people laughing" beats "an engaging experience."
- 04 **Credible.** Real, checkable detail. "12 couples in 24 hours" beats "it blew up."
- 05 **Emotional.** Make them feel the stake for one person. One named case beats a statistic about hundreds.
- 06 **Stories.** The carrier again. People remember what they can retell.

Two accuracy notes, since this framework is misquoted constantly. The mnemonic the Heaths actually wrote is "a Simple Unexpected Concrete Credentialed Emotional Story," using Credentialed to make the acronym work, while the principle itself is Credibility. And the spelling is SUCCEsS, with the trailing lowercase s. Six principles, not seven.

VERDICT Simple, unexpected, concrete. The rest is polish.

You cannot hear your own content

In 1990, Elizabeth Newton ran a study at Stanford for her psychology doctorate. She split people into tappers and listeners. Tappers picked a well-known song and tapped the rhythm on a table. Listeners guessed the song.

50%

what tappers predicted listeners would guess

The tappers heard the melody in their heads while they tapped. The listeners heard knocking. That gap is the curse of knowledge, and it is the single most reliable reason expert content underperforms. You cannot un-know your subject, so you cannot hear how your explanation lands.

2.5%

what listeners actually guessed, 3 songs out of 120

The Heath brothers put this at the centre of Made to Stick, and it is the practical case for simplifying. The target to write against: a sharp 13-year-old follows it on the first pass. In English that is roughly an eighth-grade reading level for the body and sixth for hooks. Simple is not simplistic. You are lowering the effort, not the thinking.

TABLE • THE LOAD MARKERS TO HUNT

MARKER	THE FIX
Words over three syllables with a plain substitute	Use the word you would say to a friend
Sentences over about 25 words	Break into one idea per sentence
More than one nested clause	Cut the nest, keep the claim
Passive constructions that hide who acted	Name who did it
Jargon the audience does not already use	Swap it, or define it once and move on

Two cautions so this does not fight the rest of the craft. Do not flatten every sentence to the same short length, because rhythm is what keeps people reading. And do not remove a technical term your audience knows and expects, because that reads as talking down.

VERDICT Predicted 50, delivered 2.5. Assume you are the tapper.

People share what says something about them

The 2025 research updates the emotion story. In PNAS Nexus, Cosme and colleagues ran four studies with 2,559 participants and 18,780 observations across the United States and the Netherlands, testing what makes people pass news along.

Two things predicted sharing. Self-relevance, meaning "this says something about me," and social relevance, meaning

"this is useful to someone I know." Both held at p below 0.001. The size of the effect changed by channel: for broadcast sharing, posting publicly, self-relevance led. For direct sharing, sending to one person, social relevance dominated. And when the researchers reframed content to raise self or social relevance, sharing intentions and behaviour went up, which makes this causal rather than correlational.

0.36

the strongest effect, social relevance
on direct sharing

The practical translation. Before publishing, ask the two questions the research isolates. Does this let the reader say something about who they are? And can the reader name one specific person who needs it? The second question is the one most content fails. "Useful to marketers" is not a person. "Useful to the person on your team who keeps rewriting the same brief" is.

This lines up with what Instagram itself now says it ranks on. In January 2025, Adam Mosseri named the top three signals as watch time, likes, and sends, and added that likes matter more for reach inside your followers while sends matter more for reach beyond them. Sends are the direct-share behaviour the research just explained.

A hook has one job and four ways to fail

The hook does not sell the piece. It helps the right person decide to keep going, and it does that with two things only: topic clarity within a second or two, and on-

target curiosity, meaning the viewer believes this is for them and a question is left open.

Every weak hook fails in one of exactly four ways. Diagnose which, then apply the matching fix.

TABLE • THE FOUR HOOK KILLERS

KILLER	WHAT IT LOOKS LIKE	THE FIX
Delay	The topic arrives after the greeting, the setup, the "let me explain"	Delete everything before the topic. The first noun phrase carries the subject
Confusion	The line needs a second pass to parse	Fewer words, simpler words, active voice, one idea per sentence
Irrelevance	Clear, but framed around you instead of the viewer	Swap me and my for you and your. Frame around a pain they already feel
Disinterest	Clear and relevant, and still skippable	Name the belief they hold, put your claim against it. The gap is the curiosity

The deadline changes by platform, the disease does not. A reel has roughly a 2-second scroll cliff, and on TikTok closer to 1.5. A LinkedIn post has the "see more" fold at about 140 characters on mobile. A YouTube video has two hooks: the title and thumbnail earn the click, then the first 15 to 30 seconds confirm it. A carousel has slide one, at nine words or fewer. A newsletter has the subject line.

Banned openers, because the wrapper words never survive a rewrite: "let me explain," "let's talk about," "story time" with no story in the same breath, "nobody talks about this" without naming the thing, "wait till you see this," and any call to action used as an opener. A call to action closes. It never opens.

VERDICT Clear in two seconds, or it is not a hook.

The hook earns two seconds. Story earns the rest

Four techniques do most of the work, and the first is the engine of the other three.

- 01 **The dance.** Alternate context and conflict. The connector between beats should be "but" or "therefore," never "and then." The South Park writers' rule, and it is the difference between a story and a list. Conflict opens a loop, context closes it, then opens the next one.
- 02 **Direction.** Write the last line first. Make it the line someone would share if they heard nothing else. In short vertical video the piece loops, so the last line is secretly the setup for the first line on replay.
- 03 **The lens.** A topic is not enough, because dozens of people cover the same topic. The lens is the angle only you would take. Push past the obvious one: invert the villain, find the asset hiding in the failure, jump to the second-order effect, switch the point of view.
- 04 **Rhythm.** Sentences of the same length drone. Write every sentence on its own line and look at where the lines end. A jagged edge means the lengths vary. A straight edge means the reader is about to leave.

Then the rule that outranks all four. Use real material. Invented founder stories, borrowed numbers, and imagined events destroy the exact trust this content exists to build. Structure is a writer's job. The truth is the client's, and it has to be collected before drafting, not filled in during.

For spoken formats there is one hard constraint: English delivery runs about 2 to 2.5 words a second. A 30-second script is 60 to 75 words. A 60-second script is 120 to 150. Arabic packs more meaning per word, so budget roughly 10 to 15% fewer words for the same runtime. Write to the runtime, then cut beats rather than the lens or the closing line.

VERDICT But and therefore. Never and then.

Nine books, one line each, correctly attributed

TABLE • THE REFERENCE SHELF

FRAMEWORK	SOURCE	WHAT IT GIVES YOU
STEPPS	Jonah Berger, Contagious, 2013	Six reasons people share
SUCCESS	Chip and Dan Heath, Made to Stick, 2007	Six tests for whether an idea is remembered
SB7	Donald Miller, Building a StoryBrand, 2017	The customer is the hero, the brand is the guide with a plan
The Hook Model	Nir Eyal, Hooked, 2014	Trigger, action, variable reward, investment
Purple Cow	Seth Godin, 2003	Remarkable is built into the thing, not added by the campaign
Seven principles of influence	Robert Cialdini, Influence, 1984, expanded 2021	Reciprocation, liking, social proof, authority, scarcity, commitment and consistency, unity
Three rules of epidemics	Malcolm Gladwell, The Tipping Point, 2000	Law of the few, stickiness factor, power of context
Content Inc.	Joe Pulizzi, 2015, revised 2021	Sweet spot, tilt, base, audience, revenue, diversify, sell or go big
The Big 5	Marcus Sheridan, They Ask You Answer, 2017	The five topics buyers search for and sellers avoid

Three honesty notes. Gladwell's law of the few is contested: Duncan Watts' network research found influential hubs were not necessary for transmission, and Gladwell revisited the material himself in 2024. Cialdini's unity appeared first in Pre-Suasion in 2016, not in the 2021 Influence. And Content Inc. has six steps in the 2015 edition and seven in 2021.

CHECKLIST • SHARE AND STICK GATE

- 01

Which two or three STEPPS triggers does this piece actually own? Name them.
- 02
- Does the core claim pass simple, unexpected, and concrete?
- 03
- Is there a real, checkable number or a named case in it?

04

05

06

07

VERDICT Cite the edition, or do not cite the book.

One job, one format, one idea

Format is a decision, not a habit. Pick it from the job the piece has to do, then obey that format's rules. The most common mistake is running a teaching idea as a video because video is what the team is used to.

TABLE • PICK THE FORMAT FROM THE JOB

THE JOB	THE FORMAT	WHERE IT LIVES
Reach people who have never heard of you	Short vertical video	TikTok, Reels, Shorts
Teach a framework or a list of steps	Carousel	Instagram, LinkedIn document
Break down an argument in sequence	Thread	X, Threads
Land one opinion fast	Short text post	X, LinkedIn, Threads
Announce, or say one thing beautifully	Single image	Instagram, Facebook, X
Show the process, ask the audience	Story set	Instagram, Facebook
Own a topic in depth and be searchable	Long-form video	YouTube
Be found months later by someone planning	Pin	Pinterest
Make the full argument in writing	Article or newsletter	LinkedIn, blog, email

One rule cuts across all nine. One piece carries one idea. If it needs two ideas, it is two pieces. The second idea is not depth, it is the reason the first idea did not land.

A second rule, on cross-posting. Never post one asset unchanged everywhere. The wrong aspect ratio, a competitor's watermark, or a tone that jars on the next platform will cost you reach on all of them. Reshape, do not repost. Part 5, stage 7, is how.

VERDICT Two ideas in one post is two failed posts.

The reach format, and the least forgiving one

Reels, TikToks, and Shorts are the only formats that reliably put a brand in front of people who do not follow it. They are also where a weak first second costs everything, because the audience is scrolling, not browsing.

Build it in three parts. The hook, from zero to three seconds, leads with the payoff or the tension, written as both a spoken line and on-

screen text. The body holds attention with a reason to stay every few seconds: a turn, a proof, a visual change, a "here is the catch." The payoff lands the point, makes one ask, and ends on a question rather than trailing off.

Write it as timed beats, because a script written as paragraphs gets edited as paragraphs. Each beat carries a timecode, the spoken line, the on-screen text, and the shot.

TABLE • THE SAME SCRIPT, TUNED PER PLATFORM

PLATFORM	IDEAL LENGTH	TONE AND EDIT	WATCH FOR
TikTok	21 to 34s	Low-fi, sound-on, native, trend-aware	Hook cliff is about 1.5s. Keep text out of the top 150px and bottom 250px
Instagram Reels	15 to 90s, up to 3 min	Cleaner, captioned, made for Instagram	No visible TikTok watermark. Design for mute
YouTube Shorts	Up to 3 min	Punchy, top of funnel	A separate system from long-form, and a weak funnel into it
Facebook Reels	60 to 90s	Relatable, captioned	The only reliable non-follower reach on Facebook
X and Threads video	15s to 2:20, and 15 to 60s	Native and tight	Highest-reward format on X. On Threads, pair it with a question

Two production facts worth designing around. Roughly half the audience watches on mute, so captions are not an accessibility extra, they are the script. And on Instagram, Metricool's 2026 social media study, covering 39.7 million posts from about 1.06 million accounts, found Reels reach fell 35% year on year, which means the format still wins on reach and the bar for the hook keeps rising.

VERDICT Half of them cannot hear you. Write for that half.

The teaching format, and the most saved one

A carousel teaches one thing, tells one story, or makes one argument in steps. Write the one-line spine before the first slide, because a carousel without a spine becomes a list of nice slides that nobody finishes.

The structure is four parts. Slide one is the hook slide, bold and specific, readable at a glance and on mute. Then value slides, one idea each, headline plus a line or two. Then a

proof slide, a case, a number, a before and after. Then the final slide with exactly one call to action.

Slide two carries more weight than most people give it. Per Adam Mosseri, when someone sees a carousel and does not swipe, Instagram often gives that carousel a second chance and opens it on the second frame. So slide one earns the stop, and slide two earns the second serve.

TABLE • CAROUSEL BY PLATFORM

PLATFORM	FRAMES	SIZE	NOTES
Instagram	Up to 20	4:5, 1080x1350	Saves and sends are the currency. Design for mute
LinkedIn	8 to 12 ideal, up to 300	1:1 or 4:5 PDF	Uploaded as a document. Best organic format on the platform
TikTok Photo Mode	Up to 35	9:16	Set to a trending sound. Keep text clear of the UI bands
Pinterest	Up to 5	2:3, 1000x1500	Keyword title and description. Steps or product angles
Facebook album	Flexible	1:1 or 4:5	Albums lead on shares

9x

more saves for Instagram carousels than single images

The numbers behind the format. Metricool's 2026 Instagram study found carousels earn nine times more saves than single-image posts. On LinkedIn, Socialinsider's analysis of 1.3 million posts from 16,645 company pages put native document posts at the top of the format table at 7.00% engagement by impressions, ahead of multi-image at 6.45%, video at 6.00%, image at 5.30%, and text at 4.50%.

VERDICT Slide one stops them. Slide two gets the retry.

The two text formats, and where each one breaks

A thread breaks one argument across linked posts. Post one has to work completely on its own, because most people will only ever see that one. It lands the topic, the promise, and the tension, and it never opens with "a thread"

as its own line. Then one idea per post, each skimmable alone, each pulling forward with a small cliffhanger. Then a closer that recaps in one or two lines, makes one ask, and ends on a question.

A tight seven-post thread beats a padded fifteen.

Link placement decides reach. On X the external link goes in the final post or the first reply, never mid-thread. On Threads, build it as a self-reply chain, which keeps the thread alive and signals conversation.

A short text post is one idea, three beats: the hook, the idea, the conversation layer. The conversation layer is a real question or an honest disagreement, not "like if you agree." The rules differ enough per platform that copying one post across all four is the fastest way to underperform on three of them.

TABLE • SHORT TEXT POST BY PLATFORM

PLATFORM	LENGTH TARGET	SHAPE	LINKS AND TAGS
X	280 characters, longer only if earned	Fast, opinionated, native. An image lifts reach	Link in the first reply, never the body. Hashtags 0 to 2
LinkedIn	Hook inside the first 140 characters, body up to 3,000	First person, a lesson or a real question. Post when you can reply for an hour	Link in the first comment. Hashtags 3 to 5
Threads	Under 500 characters	Lighter and curious. A soft observation beats a hot take	One topic tag, not a stack
Facebook	40 to 80 characters for a status	Plain, community, comment-driving	Naked links get the lowest reach

One number that should change who posts. Metricool's 2026 LinkedIn study, covering 673,658 posts from 63,108 accounts, found personal profiles get 63% higher engagement than company pages. If the founder will post, the founder should post.

VERDICT Post one carries the thread. Write it last.

The announcement format and the intimacy format

A single image post carries one idea that does not need teaching or motion. An announcement, a quote, a product, one beautiful frame. The caption does the work: the first line is the hook and it has to live inside the roughly 125 to 140 characters shown before "more" on Instagram and Facebook. Instagram captions run to 2,200 characters, and every line after the first has to earn its place.

One design trap. Instagram feed posts default to portrait 4:5 at 1080×1350, but profile tiles render at 3:4, so anything important near the edges gets cropped in the grid. Design at 4:5 and keep the subject safe.

Stories are the most informal surface a brand has, and the voice runs warmer there even for a formal brand. Three to six frames is the working range, at 9:16, roughly 60 seconds per segment, gone in 24 hours.

TABLE • THE FOUR FRAME TYPES IN A STORY SET

FRAME	ITS JOB
Hook frame	Stop the tap-away in frame one. On-screen text carries it, because Stories are watched on mute too
Content frames	One beat each. A step, a peek, a point, a proof
Interaction frame	At least one response sticker: poll, this-or-that, emoji slider, question box, quiz
Link or CTA frame	The link sticker, and a plain line saying what waits on the other side

The interaction frame is the point of the format, not a decoration on it. Sticker responses signal the algorithm and open direct messages, and a direct message is the closest thing on social to a conversation with a lead. Keep text and stickers clear of the top and bottom interface bands.

VERDICT Stories without a sticker are a slideshow.

Package first, then script

On YouTube the title and thumbnail are not promotion for the video, they are part of the video, and they get designed together before a word of script is written. They should not repeat each other. They should combine into one promise.

Write 3 to 5 title options. Titles run to 100 characters with about 60 visible. Curiosity plus clarity, and never a promise the video does not pay off, because YouTube penalises a high click-through rate paired with low satisfaction. The thumbnail carries one legible idea or one emotion at 1280×720.

The script has four parts. A cold open in the first 30 seconds that states what the viewer gets and why it matters, with no slow logo intro. A body in named chapters, each a clean sub-promise, with a pattern interrupt between them: a shot change, a cut to b-roll, a question, an example. A payoff that delivers the promised result plainly. And end-screen routing to the single most relevant next video plus one call to action.

The end screen is the session-time lever. Do not skip it.

The algorithm to write against, in YouTube's own framing, predicts per viewer the probability of watching and of being satisfied, then optimises total time on the platform rather than one video's views. Which is why routing to the next video is a ranking behaviour, not a courtesy. YouTube's own guidance on packaging is worth quoting: changing your title and thumbnail can get you more views, but do not change what is working.

Specs and rhythm. Long-form is 16:9 at 1080p or better, and the growth-length window is 6 to 20 minutes. Description runs to 5,000 characters with the first 150 doing the search and click work. Chapters as timestamps. About one long-form a week is a healthy creator cadence, and 2 to 4 a month is a fine brand cadence. Batch-produce.

VERDICT Title and thumbnail first. Always.

The three formats with a shelf life

Pinterest is a search engine that looks like a feed. People arrive planning and buying, not scrolling. Keywords go in three places and only three: the title, the first 100 characters of the description, and the board name. The title runs to 100 characters and leads with the keyword phrase. Not clever, findable. Freshness is a ranking factor, so re-pinning the same image is deprioritised: make two or three visually distinct Pins per destination, at 2:3 and 1000×1500, and publish seasonal work 30 to 60 days early. A Pin surfaces for months, which is the only place in social where content compounds instead of decaying.

An article makes the full argument. Find the spine first, one core argument or one promise, because an article that says three things says

nothing. Headline, then an opening with no runway, then a body in skimmable sections with one point and one piece of evidence each, then a takeaway and one action. The LinkedIn article sweet spot is 800 to 2,000 words, it is indexable, and it stays on the profile permanently.

A newsletter is an article with an inbox deadline. The subject line earns the open and the first line earns the read. Preview text extends the subject, it never repeats it. One primary call to action, not five, because five calls to action is a menu and a menu gets closed.

TABLE • THE METADATA LAYER, PER SURFACE

SURFACE	WHAT DECIDES PERFORMANCE
YouTube title	100 characters, about 60 visible, designed with the thumbnail
Caption first line	Must live before the "more" cut, about 125 to 140 characters
YouTube description	First 150 characters carry search and click. Chapters as timestamps
Subject line	The hook. Preview text extends it
Alt text	Descriptive and factual, never keyword-stuffed

On hashtags, one line settles most arguments. They are a topic signal now, not a reach lever: 0 to 2 on X, 3 to 5 on Instagram, TikTok and LinkedIn, one topic tag on Threads, and on Pinterest the keywords in the title, description, and board name do the whole job.

VERDICT Pinterest compounds. Feeds decay. Plan accordingly.

The numbers, verified July 2026

TABLE • SPECS AT A GLANCE

PLATFORM	TEXT LIMIT	IMAGE OR VIDEO SPEC	HASHTAGS
X	280, or 25,000 on Premium	Video 140s free, best 15s to 2:20. Up to 4 images	0 to 2
Instagram	Caption 2,200, 125 shown	Feed 4:5 1080×1350. Reels 9:16 1080×1920, to 3 min. Carousel to 20	3 to 5
YouTube	Title 100, description 5,000	Long-form 16:9 1080p+. Shorts 9:16 up to 3 min. Thumbnail 1280×720	Low weight
TikTok	Caption 4,000	9:16 1080×1920, ideal 21 to 34s. Photo Mode up to 35 images	3 to 5
LinkedIn	Post 3,000, fold at 140	Document to 300 pages, 8 to 12 ideal. Video 3s to 15 min	3 to 5
Facebook	Status best at 40 to 80	Reels 9:16, 60 to 90s. Feed 1:1, 4:5, or 1.91:1 link	Minimal
Pinterest	Title 100, description 800	Pin 2:3 1000×1500. Video Pin 9:16, best 6 to 15s	Not used as tags
Threads	500 characters	Up to 20 images. Video up to 5 min, best 15 to 60s	One topic tag

CHECKLIST • FORMAT GATE

- 01

One idea, one format, and the format matches the job.
- 02

The hook clears the platform's deadline, in the platform's medium.
- 03

Ratio and dimensions correct for this surface, and no watermark from another platform.
- 04

Captions burned in wherever the audience watches on mute, and all text clear of the interface safe zones.
- 05

Link in the right place: first comment on LinkedIn, first reply on X, never mid-thread.
- 06

Hashtag count matches the platform, and every tag is a topic signal.
- 07

One call to action, and it is the last thing, not the first.

METHOD NOTE • Specs shift. Verify against the platform before you quote these numbers to a client.

VERDICT Right idea in the wrong ratio is a wasted idea.

Nine stages, and most teams only run four

Most teams run idea, make, post, and hope. The four they skip are the four that compound: the brief, the review gate, repurposing, and the learning loop. The Content Marketing Institute frames the whole thing as content operations, built on three pillars, people, processes, and

technology, with seven components: intake, analysis, create, manage, distribute, repurpose, measure. The version below splits create from review, because in practice the review gate is where quality is either enforced or quietly abandoned.

TABLE • THE NINE STAGES, AND WHAT EACH ONE OWNS

#	STAGE	OWNER QUESTION	OUTPUT ARTIFACT
1	Intake	What is being asked for, and why now?	A logged request with a business reason
2	Brief	One idea, one audience, one action	A signed brief
3	Ideation	Which angle, and is it the obvious one?	5 to 10 angles, one chosen
4	Production	Who makes it, in what format, by when?	A first draft with visuals specified
5	Review and QA	Does it pass voice, hook, clarity, spec?	An approved file plus a QA report
6	Publish	Right platform, right time, right link	A live post
7	Repurpose	What else does this become?	5 to 10 derivatives, staggered
8	Measure	Did it move the stage's one metric?	A number against a benchmark
9	Learn and archive	What do we keep, kill, and reuse?	A decision, and a filed asset

Content operations borrows its oldest model from records management. AIIM, the standards body, defines the enterprise content lifecycle in five verbs: capture, manage, store, preserve,

deliver. The lesson worth taking from it is stage 9. Content teams almost never preserve or retire anything, so year two rebuilds what year one already made.

VERDICT The four skipped stages are the four that compound.

Every request gets logged before it gets made

The most expensive content is the piece that gets made because someone asked in a corridor. Intake is a small discipline that kills that: every request enters through one door, gets logged, and answers four questions before anyone opens a design file.

- 01 What is the business reason, in one sentence? If the answer is "we have not posted in a while," that is a cadence problem, not a content request.
- 02 Which funnel stage does it serve, and therefore which metric does it own?
- 03 Is there an existing asset that already does this? Check stage 9 before building anything.
- 04 What is the real deadline, and what is driving it?

This is where workflow friction shows up in the data. In the Content Marketing Institute's 2025-outlook research, 33% of B2B marketers named managing workflow and content approval processes as a challenge, down from 41% the year before. It is a solvable problem, and the solution is upstream: an intake question answered on day one prevents a rewrite on day nine.

Requests without a business reason are hobbies with a deadline.

Two intake rules worth enforcing. A request with no named owner does not proceed. And a request that arrives inside another team's review cycle joins the next cycle, because a queue that anyone can jump is not a queue.

VERDICT One door in. Everything else is a favour.

The oldest artifact in advertising, and still the one that decides quality

The brief is where the thinking happens, and the industry has known it for sixty years. Stephen King introduced the T-Plan at J. Walter Thompson in June 1964, and with Stanley Pollitt at BMP founded account planning in 1968. The modern creative brief is what that discipline produced.

The nearest thing to an industry standard is The Client Brief, published in 2003 by the IPA, ISBA, MCCA and PRCA. Its eight headings beat most templates sold today.

- 01 **Project management.** Date, name, type, job number, brand, teams.
- 02 Where are we now?
- 03 Where do we want to be?
- 04 What are we doing to get there?
- 05 Who do we need to talk to?
- 06 How will we know we have arrived?
- 07 **Practicalities.** Deliverables, budget, timing, mandates, legal and identity rules.
- 08 **Approvals.** Who signs off, who wrote the brief, who approves the work.

VERDICT Eight headings. Two of them decide everything else.

Four boxes for a post, nine fields for the file

For a single post that is too much, so use the shorter tool. Saatchi and Saatchi's device is the OIIC, named in public by their London chief strategy officer Richard Huntington: objective, issue, insight, creative challenge. Four boxes, and the insight box decides the work.

TABLE • THE ONE-PAGE CONTENT BRIEF

FIELD	THE TEST IT HAS TO PASS
Objective	One sentence, one metric, a number if there is one
Audience	A person, not a segment. Named behaviour, not demographics
The one idea	One sentence. If it takes two, it is two pieces
The insight	Something true that the audience has not heard said out loud
Type, stage, format	One of six types, one of four stages, format chosen from the job
The one action	What exactly should they do, and where does it lead
Mandatories and approver	Anything that cannot move, and the one name who signs off

VERDICT Empty insight box means the brief is not done.

Find the angle nobody else would take

Ideas do not arrive, they get mined. Eight veins produce almost everything worth publishing, and the first three are the ones teams underuse.

- 01 Customer language and the questions sales gets asked every week.
- 02 Search: what people type into TikTok, YouTube, Pinterest, and Google about this category.
- 03 Wins and proof: results, before and after, client outcomes with permission.
- 04 Behind the scenes: the process, the mistakes, the tools, the room.
- 05 Contrarian takes: the thing your industry agrees on that you do not.
- 06 Trends and formats: a current shape, filled with your own substance.
- 07 Personal stories: real ones, from real notes.
- 08 The audience's own words back to them: polls, replies, comment questions.

Then push past the first angle. The obvious lens on a topic is the one twenty other accounts already used. Four moves that produce a sharper one: invert the villain, because the obvious villain is rarely the real one. Find the asset hiding in the failure. Jump to the second-order effect nobody is connecting. Switch the point of view to the customer, the loser, or the thing itself.

Two documented processes are worth borrowing here, and both are commonly misattributed. IDEO's own three phases, published in its Field Guide to Human-Centered Design in 2015, are inspiration, ideation, implementation. The five-

stage empathise, define, ideate, prototype, test sequence is Stanford d.school's, and the d.school explicitly calls them modes rather than steps, noting they can be used in various orders. Getting that right matters when you are citing it to a client.

For a time-boxed version, the design sprint developed by Jake Knapp at Google in 2010 and brought to GV in 2012 runs five named days: map the problem, sketch competing solutions, decide and turn ideas into a testable hypothesis, build a high-fidelity prototype, test it with real people. A content team can run the same week for a new format or a new series.

VERDICT The obvious angle is already taken. Take the second one.

Batch the making, separate the roles

Production fails for two reasons: the team makes one piece at a time, and the same person writes, designs, and approves. Fix both structurally.

Batch by mode, not by post. Write all the week's scripts in one session, film all the talking-head pieces in one session, edit in another, design in another. Switching between writing and designing five times a week costs more than any tool saves.

Separate the roles even when one person holds several of them, because the roles have different jobs. Writer produces the words. Designer or editor produces the frames. Reviewer applies the gate at stage 5. Publisher owns the platform mechanics and the timing. Owner answers for the metric at stage 8. One person can hold four hats, as long as they do not wear two at once.

TABLE • WHAT "DONE" MEANS PER FORMAT, BEFORE IT ENTERS REVIEW

FORMAT	DONE MEANS
Short video	Timed beat script, hook written as spoken line and on-screen text, shot list, caption, sound choice
Carousel	Per-slide headline and support copy, image direction per slide, caption, hashtags, alt text
Thread	Numbered posts each inside the character limit, link placement noted
Short text post	Final text per platform, posting window, link placement
Long-form video	3 to 5 title options, thumbnail concept, chaptered script, description, end-screen plan
Article	Headline plus 2 to 3 alternates, full piece with subheads, one call to action
Pin	2 to 3 variants, keyword-led title, front-loaded description, board name

One rule on visuals that keeps briefs honest. Nobody hands over a generated image and calls it a design. Every visual leaves production as two things: a generation prompt, and a human brief with exact dimensions, safe zones,

copy to set, typefaces, palette, and logo placement. Brand colours and fonts are never invented or approximated. If a value is unknown, it stays a labelled placeholder until someone fills it in.

VERDICT Batch the mode, not the post.

One gate, and it has no authority over the maker

The most useful review model in creative work is Pixar's Braintrust, described by Ed Catmull in Harvard Business Review in 2008. A group of directors watches a film in progress and says

what is not working. The two rules that make it function: there are no mandatory notes, and the group has no authority. The director decides what to do with the advice.

Candour without authority. That is the whole design.

That structure matters more than the checklist, because a review board with authority turns into an approval queue, and approval queues produce work that is safe rather than good. Pixar pairs it with dailies, showing incomplete work to the whole crew, which normalises being seen mid-draft.

Then run the fixed checks. This is the gate that every piece passes through before it ships, and it is deliberately mechanical so that judgement is spent on the work rather than on remembering the list.

CHECKLIST • THE QA GATE

- 01 **Voice.** Does it sound like the brand, in the right register for this channel?
- 02 **Hook.** Does the opener clear delay, confusion, irrelevance, and disinterest, inside this platform's deadline?
- 03 **Clarity.** Would a sharp 13-year-old follow it on the first pass? Flag any sentence over 25 words or any jargon the audience does not use.
- 04 **Human check.** Strip the machine tells: hollow reframes, significance inflation, category nouns where an instance belongs, adjectives with no number behind them, three-item lists where the real count is two or four.
- 05 **Specificity.** Is there at least one real number, named case, or lived detail? Nothing invented.
- 06 **Spec.** Ratio, dimensions, character limits, hashtag count, link placement, watermark, captions, safe zones.
- 07 **Type and funnel fit.** Does it still serve the type and stage the brief assigned, or has it drifted into a sales post?
- 08 **Conversation layer.** Does it end with something a person can reply to?
- 09 **One action.** One call to action, in the right place.

The output of the gate is two things: the corrected file, and a short report marking each check as pass, fixed, or flagged. Fixed means the reviewer changed it. Flagged means

only the owner can decide. A reviewer who silently rewrites intent is not reviewing, they are rewriting.

VERDICT Fix what is broken. Flag what is a choice.

The first hour is part of the post

Publishing is not uploading. Three things decide how far a finished piece travels: when it goes out, where the link sits, and whether anyone is there to reply.

On LinkedIn the first 60 to 90 minutes largely set the total reach, because the platform tests the post on the most-engaged part of your network and expands on early meaningful comments. So the posting time is not "when the audience is online," it is "when the author can

sit with it for an hour." LinkedIn's own engineering paper, published in February 2026, names the two behaviours its feed ranker optimises: long dwell, meaning time on a post past a threshold, and contribution, meaning a like, comment, or share.

Link placement is a reach decision on three platforms. First comment on LinkedIn. First reply on X. Never mid-thread. On Facebook, naked links get the lowest reach of any format.

TABLE • PUBLISHING RULES THAT COST REACH WHEN BROKEN

PLATFORM	THE RULE
LinkedIn	Post when you can reply for the first hour. Link in the first comment
X	Link in the first reply. Space posts out, since a per-author cap means a flood cannibalises itself
Instagram	Design for sends. A DM share is the main path to non-followers
TikTok	Keyword the caption, the on-screen text, and the spoken words, because search is real there
Threads	Reply to your own post to keep it alive. Reach is largely follower-independent
Pinterest	Publish seasonal content 30 to 60 days early

One thing worth planning around. LinkedIn's ranker is trained on dwell and contribution, both of which are behaviours a real reader produces and a pod does not. Practitioner reporting through early 2026 points the same way:

recycled posts and "comment to agree" bait are being priced down. Treat that reporting as directional, and treat the engineering paper as the source.

VERDICT Publish when you can reply, not when you are free.

Produce once, publish many, reshape every time

One hero asset holds a week or more of content if it is reshaped per platform rather than copy-pasted. Start by pulling out the atoms: the core argument, the best self-contained moments, the teachable steps, the quotes and numbers, and the story.

TABLE • WHERE EACH ATOM GOES

ATOM	BECOMES	PLATFORM
Best moments	Vertical clips	Reels, TikTok, Shorts
Key points or steps	A carousel	Instagram, LinkedIn
Core argument	A thread and a question	X, Threads
A sharp line	A short post	X, LinkedIn, Threads
Quotes, stills, numbers	Pins	Pinterest
The whole piece	An article or newsletter	LinkedIn, blog, email

Reshape the tone, not only the size. Low-fi and sound-on for TikTok. Polished and captioned for Reels. Professional framing for LinkedIn. A conversational question to open on Threads. Strip any source watermark and rebuild natively, because watermarked cross-posts get suppressed on every platform that detects them.

Then stagger the release. Lead with the hero, then the discovery clips, then the authority pieces, then the owned-channel capture. Dumping ten derivatives in one day competes with itself.

The documented version of this is Gary Vaynerchuk's Content Model, 2018: pillar content, micro content, distribute both, harvest community

insight, make community-driven micro content, distribute again. His subtitle claims 30-plus pieces from a single keynote, which is his own number rather than independent research, and the shape of the model holds regardless.

For a sense of how seriously large operators take distribution measurement, BuzzFeed built and published a system called Pound, short for Process for Optimizing and Understanding Network Diffusion, in 2015. It tracked how a story spread from one sharer to the next, across networks and into private messages, and its finding was that content spreads through many interconnected sharing trees rather than one, so a piece should be judged on its whole downstream cascade rather than its first post.

VERDICT One asset, ten shapes, zero copy-paste.

Compare like with like, or do not compare

The most common measurement error in social reporting is comparing two engagement rates that use different denominators. Some vendors divide engagement by followers. Others divide by reach. Both are correct. They answer

different questions, per-reach numbers run several times higher, and there is no conversion between them because the two sets come from different vendors and different customer bases.

TABLE • TWO DIFFERENT QUESTIONS, TWO DIFFERENT NUMBERS

DENOMINATOR	WHAT IT TELLS YOU	TYPICAL 2025 MEDIANS
Per follower	Audience health, and whether your own people respond	TikTok 2.01%, Instagram 0.30%, YouTube 0.21%, X 0.03% (Rival IQ and Quid, 2026 edition, about 2,700 brands)
Per reach	Content quality, and whether the people served responded	LinkedIn 6.2%, Facebook 5.6%, Instagram 5.46%, TikTok 4.6%, Pinterest 4.0%, Threads 3.6%, X 2.5% (Buffer, 2025 data, 52M+ posts)

State the limit whenever you quote either. Every one of these benchmark sources is a social media tool vendor reporting on its own customer base, not a random sample of any platform. Rival IQ and Quid skew to large brands and include an influencer category. Buffer skews to solo creators and small teams. Metricool has the largest sample at about 1.06 million accounts and is the most defensible for year-on-year direction. And watch the labels: a report published in 2026 almost always contains 2025 data. Write the data year, not the edition year.

Measure against the stage's one metric from page 6, then against your own last 30 posts, then against an industry benchmark, in that order. Your own trailing median is the only benchmark that controls for your audience.

For a documented example of measurement built into publishing instead of added at the end, The Washington Post's Bandito tool tests headline, blurb, and thumbnail combinations on live traffic and automatically promotes the winner. Built in-house in 2016. Most teams cannot run that, and most teams can run its principle: test one variable, on the same audience, and let the number settle the argument.

VERDICT Your own median beats any industry benchmark.

The stage everyone skips, and the reason year two costs the same as year one

Two rituals close the loop. Both are cheap, and both are the first thing dropped when a team gets busy.

The first is the postmortem, in Pixar's format: ask each group to list the top five things they would do again and the top five they would not. The forced count matters. Five and five stops the meeting from becoming either a celebration or a blame session, and it produces a written artifact rather than a feeling.

The second is the archive. Content teams almost never retire or catalogue anything, so the same explainer gets rebuilt every eighteen months by someone who did not know it existed. AIIM's lifecycle names preserve and deliver as distinct stages for exactly this reason. In practice you need three shelves: evergreen assets that can be re-cut or reposted with a refresh, reference assets that inform new work without being republished, and retired assets that are wrong now and should stop circulating.

TABLE • THE MONTHLY REVIEW, IN FOUR QUESTIONS

QUESTION	WHAT YOU DO WITH THE ANSWER
Which pillar and which content type over-performed?	Shift next month's mix toward it, by a few points, not a rewrite
Which type sat at zero?	Schedule two of them next month
Which hero asset produced the most derivatives that worked?	Make more of that hero shape
What did we rebuild that already existed?	It goes in the archive index today

One caution against over-engineering this. Netflix's own engineering team describes its productions as snowflake workflows, noting that no two productions share the same workflow, and

organises around functions rather than one rigid pipeline. A lifecycle is a checklist for what must happen, not a claim that every piece takes the same path.

VERDICT Five to repeat, five to never. Written down.

Five roles and one status ladder

The roles below are jobs, not job titles. In a team of one, one person holds all five and switches between them deliberately. In a team of six, every stage has a name attached.

TABLE • THE FIVE ROLES

ROLE	OWNS	HANDS OFF TO
Strategist	Objective, platforms, pillars, mix, cadence	Editor
Editor	Intake, brief, angle selection, the calendar	Maker
Maker	Script, copy, design, edit, visual direction	Reviewer
Reviewer	The QA gate at stage 5	Publisher
Publisher	Platform mechanics, timing, link placement, replies in the first hour	Strategist, with the numbers

Then one status ladder, used everywhere, so that a glance at the calendar answers where anything is.

01 **Idea.** Logged with a pillar and a type. Nothing made yet.

02 **Drafted.** First version exists, not checked.

03 **Ready.** Passed the QA gate, visuals attached, waiting on a slot.

04 **Scheduled.** Slot booked, publisher assigned.

05 **Posted.** Live, with the reply window covered.

The calendar carries one row per slot with the same columns every time: date, platform, format, content type, pillar, working title, visual needed, repurposed from, status, owner. It is boring on

purpose. A calendar that changes shape every month cannot be audited, and an unauditable calendar quietly becomes a wish list.

An over-full calendar is a calendar that gets abandoned.

VERDICT Five roles, five statuses, one row per post.

Nearly everyone uses it. Almost nobody reports better performance

Two sets of numbers, from 2025, that should shape how a team uses AI in this lifecycle.

In the Content Marketing Institute's 2026 research, 95% of B2B marketers said their organisations use AI-powered applications and 89% use content-creation tools. Among those using content-creation AI, 87% reported better

productivity and 80% better operational efficiency. Then the number that matters: only 39% said content performance improved. 34% saw no change and 5% said performance got worse, with the rest unsure or saying it was too early to tell. On quality, a separate question, 12% said it decreased.

87%

say AI improved productivity

39%

say AI improved content performance

The discovery side agrees. Graphite classified 65,000 article URLs sampled from Common Crawl and found that as of May 2025 that sample was about 52% human-written and 48% AI-generated, yet among top-ranking Google articles 86% were human-authored and among ChatGPT and Perplexity citations 82% were. Ahrefs, looking at 900,000 newly created pages in April 2025, found the honest version of the saturation story: pure AI pages were 2.5%, pure human 25.8%, and mixed 71.7%. AI assistance is near-universal. AI authorship is rare, and it does not rank.

Both studies rest on imperfect AI detectors and show correlation rather than cause, so hold them loosely. The working rule they support is not

loose at all. Use AI where the lifecycle is mechanical: intake logging, transcription, first-pass repurposing, variant generation, spec checking, tagging the archive. Keep the human where the value is: the insight in the brief, the lens, the real number, the story only this company can tell, and the final read.

Disclosure is now a platform rule, not a preference. YouTube requires creators to disclose realistic altered or synthetic content, Meta labels AI content across Facebook, Instagram, and Threads using embedded metadata signals, and TikTok says it has labelled over 1.3 billion videos and is applying invisible watermarking so labels survive re-editing.

VERDICT Automate the pipeline. Never automate the point.

Everything on one page, in the order it happens

STRATEGY, ONCE A QUARTER

- 01 The goal is ranked, and the top goal fits in one sentence.
- 02 Every funnel stage has one metric and one owner.
- 03 Every platform has a named funnel role. Platforms without one are dropped.
- 04 Three to five pillars, and at least two of them are uncomfortable to publish.
- 05 One hero asset with a rhythm, and one owned channel.
- 06 A weekly cadence the team can hold for 12 weeks.

THE PIECE, EVERY TIME

- 07 Logged through intake with a business reason, not a corridor request.
- 08 Brief has one idea, one audience, one action, and a filled insight box.
- 09 The angle is not the obvious one. Name the lens.
- 10 Assigned exactly one content type and one funnel stage.
- 11 Format chosen from the job, not from habit.
- 12 Two or three STEPPS triggers named on purpose.
- 13 Core claim passes simple, unexpected, concrete, and carries one real number or named case.

- 14 Hook clears delay, confusion, irrelevance, and disinterest inside the platform's deadline.
- 15 A sharp 13-year-old follows it on the first pass.
- 16 Machine tells stripped: hollow reframes, inflated significance, category nouns, adjectives with no number.
- 17 Spec correct: ratio, limits, hashtags, link placement, captions, safe zones, no foreign watermark.
- 18 One call to action, and it closes rather than opens.
- 19 Ends with something a person can reply to.
- 20 Published when someone can answer replies for the first hour.

AFTER, EVERY MONTH

- 21 Derivatives produced and staggered, not dumped.
- 22 Measured against the stage metric, then against your own trailing median, then against a benchmark, with the data year stated.
- 23 Mix audited on the last 30 published posts, not the plan.
- 24 Five things to repeat and five to never, written down.
- 25 Anything rebuilt that already existed goes into the archive index today.

VERDICT Twenty-five checks. None of them require budget.

Stand up the system in four weeks

Do not try to run all nine stages from day one. Stand them up in this order, because each week produces the input the next week needs.

TABLE • WEEK BY WEEK

WEEK	WHAT YOU DO	WHAT EXISTS AT THE END OF IT
1	Rank the goal. Assign one metric per stage. Cut platforms without a role. Name 3 to 5 pillars	A one-page strategy anyone on the team can recite
2	Choose the hero asset and its rhythm. Write the one-page brief template. Mine 20 to 30 angles from the eight veins	A brief template and an idea bank
3	Build the calendar to your real capacity with the standard columns and the five-status ladder. Produce the first hero and batch its derivatives	Four weeks scheduled, one hero made
4	Run the QA gate on everything before it ships. Publish with the reply window covered. Set the monthly review date	A live cadence and a review in the diary

Three things to resist in month one. Do not add a platform, because the system has not been tested on the ones you have. Do not raise the cadence, because week 4 is when the real capacity shows itself. And do not skip the QA gate to hit a posting date, because the first time it gets skipped is the moment it stops being a gate.

By day 30 you should be able to answer four questions without opening a file: what are we here for, what do we talk about, how often do we publish, and who says yes. If any of the four takes more than a sentence, the strategy page is not finished yet.

VERDICT Four weeks to a system. Then hold it.

Every framework, named. Every number, sourced

Books and research. Jonah Berger, *Contagious*, Simon and Schuster, 2013. Berger and Milkman, "What Makes Online Content Viral," *Journal of Marketing Research*, 2012, 49(2), 192 to 205. Chip and Dan Heath, *Made to Stick*, Random House, 2007, including Elizabeth Newton's 1990 Stanford doctoral study. Cosme and colleagues, *PNAS Nexus*, 2025, 4(2) pgaf019. Donald Miller, *Building a StoryBrand*, 2017. Nir Eyal, *Hooked*, 2014. Seth Godin, *Purple Cow*, 2003. Robert Cialdini, *Influence*, 1984, expanded 2021. Malcolm Gladwell, *The Tipping Point*, 2000. Joe Pulizzi, *Content Inc.*, 2015 and 2021. Marcus Sheridan, *They Ask You Answer*, Wiley, 2017. Ann Handley, *Everybody Writes*, Wiley, 2014 and 2022.

Process and agency sources. The Client Brief, IPA, ISBA, MCCA and PRCA, 2003. Richard Huntington on the Saatchi and Saatchi OIIC, 2008. The 4A's record of Stephen King's T-Plan at J. Walter Thompson, 1964. IDEO.org, *The Field Guide to Human-Centered Design*, 2015. Stanford d.school design thinking process guide. GV, *The Design Sprint*. Ed Catmull, "How Pixar Fosters Collective Creativity," *Harvard Business Review*, September 2008. BuzzFeed on Pound, 2015. The Washington Post on Bandito, 2016. Netflix Technology Blog on production media management. Content Marketing Institute on content operations, 2026. AIMM on the content lifecycle.

Strategy frameworks. Avinash Kaushik, *See Think Do and See Think Do Care*, Occam's Razor, 2013 onward. HubSpot Research, Mimi An, "Topic Clusters," 2017. Smart Insights, Chaffey and Bosomworth, *The Content Marketing Matrix*, 2012. Ehrenberg-Bass Institute, John Dawes, the 95-5 rule. Google, *The YouTube Creator Playbook for Brands*, on Hero, Hub and Hygiene. The Coca-Cola Company, *Content* 2020, 2011, on the 70/20/10 split. Gary Vaynerchuk, *The GaryVee Content Model*, 2018, *Document Don't Create*, 2016, and the \$1.80 strategy, 2017.

Platform and benchmark data. Content Marketing Institute and MarketingProfs B2B research, 2024, 2025 and 2026 editions. Rival IQ and Quid, 2026 Social Media Industry Benchmark Report. Buffer, *The State of Social Media Engagement 2026*. Socialinsider social media and LinkedIn benchmarks. Metricool 2026 studies of Instagram, LinkedIn, and social media overall. Adam Mosseri on Instagram ranking signals, January 2025. LinkedIn engineering, arXiv 2602.12354, February 2026. YouTube Help on search and discovery and on AI disclosure. TikTok Transparency Center. Meta on labelling AI content. Graphite via Axios, 2025. Ahrefs, 2025.

METHOD NOTE

Three limits, stated plainly. Every social benchmark quoted here comes from a tool vendor reporting on its own customers, not a random sample of any platform, so treat all of them as directional. A report published in 2026 almost always contains 2025 data; the data year is what matters. And the AI-content studies rely on imperfect detectors and show correlation, not cause. Platform specs were verified in July 2026 and will move.

We build the *department*, not the deck

BMD, Building Marketing Department, is a boutique marketing consulting practice led by Redha Alayesh. The work is hands-on marketing leadership for companies that need a real marketing function stood up or fixed, rather than another campaign delivered on top of an empty structure.

The track record this guide comes out of 40-plus marketing departments worked with, across sectors including agriculture, real estate, and luxury gifting. The frameworks in Part 3 are from the books. The lifecycle in Part 5 is from the agencies, the media houses, and from watching which stage a marketing team quietly drops first when it gets busy. It is almost always the brief, and then the review gate.

Redha Alayesh • r@redhaalayesh.com • Building Marketing Department

Disclosure, since this guide argues for a system that a consultancy sells. Every framework here is public, attributed, and usable without hiring anyone. Most teams do not need a consultant to run the nine stages. They need to stop skipping four of them.

Stop asking why nothing moved. Decide what should move.