A 28-DAY SYSTEM FOR AI VISIBILITY THROUGH PRESS RELEASES

DAVID A MCINNIS

NEWS MARKETING



A GUIDE FOR MARKETING AND PR PROFESSIONALS

News Marketing

The 28-Day System for AI Visibility Through Press Releases

David A. McInnis

Attentionworthy Inc

A Note on the Velocity of Innovation

Artificial Intelligence is expanding faster than anything we have seen before. This is not the final word on AI and SEO press release visibility; it is a foundational blueprint. We will update this ebook as we learn more and as the technology advances. Follow our blog and LinkedIn channels for real-time updates. **Revision: March 2026**

Copyright © 2026 by David A McInnis

All rights reserved.

No portion of this book may be reproduced in any form without written permission from the publisher or author, except as permitted by U.S. copyright law.

Contents

Forword V

Introduction VII

1. What Is News Marketing? 1
2. How AI Actually Works: Training, Retrieval, and the Frozen Brain 11
3. AI Training, AI Grounding, and AI Visibility: The Three Pillars 20
4. AIO vs GEO: The New Visibility Playbook 27
5. The 28-Day Rule and the Physics of Newsworthiness 34
6. Entity Authority and the Knowledge Graph 43
7. Schema Markup: Your Complete Blueprint 53
8. Zero-Click Content and the Rise of Concept Chunking 66
9. Managing Expectations in the Age of AI 76
10. After You Hit Publish: The Real Work Begins 83
11. The Long Fat Tail: Why Press Releases Are Assets, Not Events 90
12. Competitive Dynamics in the Retrieval Era 96
13. Engagement: What Makes a Press Release Pop 102

14. Retargeting: Turning Readers Into Relationships 109
15. The Visibility Advantage: Building Your System 118
16. Additional Resources 122
17. Glossary 128
Endnotes 133
About the Author 135

Forword

Few people have worked effectively at the intersection of public relations, technology, business, and marketing. Building something that truly changes how media works requires a rare combination of vision, technical skill, and practical execution.

David McInnis is one of those rare people. Early on, he built PRWeb, a platform that helped define what he called "News Marketing." It gave businesses a more direct and affordable way to connect their brands and stories with consumers, journalists, and search audiences at scale.

I know this firsthand because I was one of its users. I used PRWeb to help a number of startups, public figures, and brands expand their reach and sharpen their message. Over time, David invited me to serve as an advisor to PRWeb as Chief Evangelist, where I helped teach businesses how to use the platform to launch products, share milestones, and build brand visibility.

When Vocus acquired PRWeb in 2006, David had every reason to step away. Instead, he chose to invest in the people around him, making sure his employees also benefited in meaningful ways. He then turned his attention to other ambitious challenges, including work in 3D printing, prototyping, and long-term digital archiving.

Even so, he never stopped paying attention to how the media and commu-

nications landscape was evolving. As artificial intelligence began moving from theory into practical use, David recognized something important: while the technology was changing quickly, much of the PR industry was not. Many of the major platforms were still operating with the same assumptions and strategies they had relied on for years.

In early 2022, David reached out to share a new idea he had been developing—an AI-powered platform designed to help create, optimize, and distribute user-generated news content in a far more effective way. What stood out to me most was not just the platform itself, but David's understanding of how AI systems interpret, prioritize, and surface information.

News Marketing is a practical guide to writing and shaping your message for maximum reach and relevance in this new environment. It also offers a path for using David's latest tools to help your expertise and ideas show up where modern audiences are increasingly looking: inside AI-driven platforms and discovery systems.

David is more than a longtime friend and business collaborator to me. He is someone I've consistently looked to for insight on where the web is going, how new technologies can be applied in useful ways, and how to build with both ambition and integrity.

Read this book with the understanding that AI is changing quickly. The ideas here matter now, and David is committed to keeping them current as the technology continues to evolve.

Mark Alan Effinger
CoFounder, Nootopia
CPO, BiOptimizers

Introduction

The Press Release as a Retrieval Asset

I've been in the press release business since 1997. I've watched the industry shift from fax machines to email blasts, from newswires to search engines, from SEO to social media. Every shift felt big at the time. None of them compare to what's happening right now.

Artificial intelligence is expanding faster than any technology I've seen in three decades. Faster than the web browser. Faster than Google. Faster than social media. And it's not just the speed of adoption that should get your attention. It's the way AI is fundamentally changing how people find, evaluate, and trust information.

Search engines no longer just rank blue links. Google's AI Overviews now summarize answers at the top of the page before you ever scroll. ChatGPT, Perplexity, Claude, Copilot, Gemini: people are asking these systems questions the way they used to ask Google. And these systems don't return a list of ten websites. They return an answer. A synthesized, confident, single answer drawn from sources the user may never click on.

That changes everything for marketers, communicators, and PR professionals.

Here's what most people in our industry haven't fully grasped yet: the

press release isn't dead. It's more relevant than it's been in the thirty years that I have been in the business. But not for the reasons you think. Not because journalists are reading them (most aren't). Not because newswires have some magical reach. The press release matters now because it is, by its very nature, exactly the kind of document AI systems are designed to find, evaluate, and reuse.

Think about what a press release actually is. It's public. It's timestamped. It's structured. It names specific entities: companies, people, products, locations. It makes declarative statements. It lives on the open web. It gets archived. It gets indexed. It persists.

That's not a relic of old PR. That's a retrieval asset.

This book is a practical guide to publishing one clear, well-structured press release every 28 days and turning that discipline into a system for long-term AI visibility. Not theory. Not hype. A repeatable process you can implement starting this month.

The Shift From Distribution To Retrievability

For decades, the PR industry measured success by distribution. How many newsrooms received your release? How many media outlets carried it? How many impressions did it generate? That model is broken. AI has shattered it completely.

Visibility is no longer about distribution. It's about retrievability: whether your content can be found by an AI system when a relevant question is asked, structured in a way that allows it to be chunked, matched, and surfaced as part of a generated answer.

This is a fundamentally different way of thinking about press releases.

You're not writing for a journalist's inbox. You're writing for retrieval systems that evaluate semantic meaning, source authority, freshness, and factual density.

And here's the part that should excite you: press releases are almost perfectly designed for this new reality. They've always been structured. They've always been timestamped. They've always named entities clearly. The format that PR professionals have used for over a century turns out to be one of the most AI-compatible content formats on the web.

You just have to use it correctly.

This book is a practical guide for marketing and PR professionals who need to understand how AI changes their work and what to do about it. You'll learn how AI systems actually process and retrieve information, why a 28-day publishing cadence creates compounding visibility, how to structure press releases so each paragraph functions as a retrievable "brandable chunk," and how to measure whether your efforts are working.

I built PRWeb in 1997 because I believed press releases should be searchable, public, and directly accessible to anyone. I built Newsworthy and NewsRamp because I saw the same principle extending into a world where AI, not just Google, would become the primary information retrieval system.

The principles haven't changed. The mechanisms have. And this book is a structural blueprint grounded in mechanism, not hype. One release every 28 days. Structured for retrieval. Extended across formats. Measured against real outputs.

That's the system. Let's build it.

Chapter One

What Is News Marketing?

In 1997, I had a $600 epiphany.

I was running a small business, and like every other small business owner, I wanted press coverage. I wanted people to know what we were doing. So I did what everyone told me to do: I wrote a press release and paid a traditional newswire to distribute it. Six hundred dollars, which felt like real money at the time. It still does.

The release went out. I watched it scroll across PointCast, which was the hot push-media platform of the moment. I felt a brief surge of excitement. And then it was gone. Streamed off the screen. Replaced by the next headline. Replaced by the one after that.

No phone calls. No website traffic. No emails. No engagement of any kind. Six hundred dollars, and all I had to show for it was a fleeting line of text on a screen nobody was really watching.

I sat there and asked myself a simple question: What does a newswire actually do?

The answer seemed obvious and straightforward. It distributes. It takes your release and pushes it to newsrooms, to terminals, to media monitor-

ing services. But distribution without discoverability is noise. My release had been distributed to thousands of endpoints and discovered by no one.

Around the same time, I was watching what pharmaceutical companies had started doing with direct-to-consumer marketing. The FDA had loosened the rules on broadcast advertising in 1997, and suddenly drug companies were running TV ads that told patients to "ask your doctor about" a specific medication. They weren't going through the traditional gatekeepers anymore. They were bypassing the doctor's recommendation loop entirely and putting their message directly in front of the person who would ultimately make the decision.

That clicked for me. The pharmaceutical industry had figured out that when the gatekeeper model doesn't serve you, you go around it. You go directly to the consumer.

I started asking myself: could I do the same thing with press releases? Could I create a media bypass newswire that skipped the journalist gatekeeper and put announcements directly in front of customers? What tools existed that would let me do this?

The answer was search. The web was exploding. Search engines were indexing everything. Directories like Yahoo! were organizing content for easy discovery. People were going online to find information on their own terms. If I could put a press release on the open web in a way that search engines could find it, I wouldn't need a journalist to decide my news was worth covering. The customer could find it themselves.

That insight, and the frustration of that $600 lesson, led me to create PRWeb. I did not realize it at the time, but I had stepped into the world of Search Engine Optimization (SEO).

The Pitch Problem

The entire PR profession in the late 1990s was built around one activity: The Pitch. You wrote a press release, but the release itself wasn't really the product. The product was the pitch, the phone call, the relationship with a journalist that convinced them to write a story.

If you were Cisco, Motorola, Coca-Cola or Microsoft, The Pitch worked beautifully. But if you were a regional software company or a startup with twelve employees? The Pitch was a brick wall. You didn't have the relationships or the brand recognition. The entire industry served the top one percent of organizations brilliantly and failed everyone else.

I didn't want to fix The Pitch. I wanted to bypass it entirely, the same way pharmaceutical companies had bypassed the traditional prescription model. And search was the mechanism.

Press Releases As Direct-To-Consumer Assets

The idea behind PRWeb was radical at the time, though it seems obvious now. What if press releases weren't private messages to journalists? What if they were public, searchable documents that anyone could find? The pharmaceutical industry had proven that going direct to the consumer worked. I was betting the same logic applied to news.

We put press releases on the open web. We optimized them for search engines. We let anyone read them, not just journalists with wire terminal access. We treated press releases as direct-to-consumer content, or D2C searchable assets, years before "content marketing" became a buzzword.

It worked. Small businesses started getting traffic directly from their press releases. Customers found them through Google. Bloggers found them. Industry analysts found them. And yes, journalists found them too, but on their own terms, through search, not because someone cold-called their desk.

That was the beginning of what I now call News Marketing.

The News Marketing Framework

News Marketing is the strategic use of news-format content, primarily press releases, as discoverable, multi-format assets designed to reach customers, AI systems, and media wherever they consume information.

It's different from traditional PR. It's different from content marketing. It borrows from both but operates on its own logic. Let me walk you through the hallmarks.

The Power of the Medium

Press releases carry psychological weight that other content formats don't. When someone reads a press release, they process it differently than they process a blog post or a social media update. There's an inherent sense of officiality, formality, timeliness, and authority baked into the format.

This isn't accidental. The press release format has been conditioned into public consciousness over more than a century. When you read a headline followed by a dateline, a city, a date, and a structured series of paragraphs making declarative statements, your brain registers it as news. As something that happened. As something real.

That psychological weight translates directly into trust signals for both human readers and AI systems. A press release that says "Acme Corp announced today that it has partnered with Beta Inc to expand operations in the Southeast" reads as fact. The same information in a blog post reads as marketing. The format itself does work that no amount of clever copywriting can replicate.

There's another dimension here that most marketers haven't considered yet. News, by definition, is new. It didn't exist before you published it. Which means it doesn't exist inside AI systems yet. The partnership you just announced, the product you just launched, the data you just released: none of that is in any model's training data. None of it is in any retrieval index. You're creating net new information, and that creates a genuine opportunity.

You're not competing against an ocean of existing content on a well-worn topic. You're introducing something the system has never seen. When an AI retrieval system encounters a query about that partnership, that product, that data point, your press release may be the only authoritative source that exists. You're not fighting for visibility in a crowded field. You're planting a flag on open ground. This is one of the most underappreciated advantages of news-format content in the AI era. Evergreen blog posts compete against thousands of similar pages. A well-structured press release about something that actually just happened occupies a space that no one else can claim, because no one else has that information yet.

Placements Over Distribution

The old model was distribution; push your release to as many endpoints as possible and hope for the best. News Marketing flips this. The goal isn't

to distribute widely. The goal is to place strategically.

Placement means ensuring your message appears where your audience actually consumes content. That might be a search engine results page. It might be an AI-generated summary. It might be a niche industry blog, a podcast show notes page, or a newsletter.

In the News Marketing framework, a press release is a seed. It's the source document from which multiple placements grow. One release can become the basis for a bylined article on an industry site, an FAQ page on your own website, a social media summary, an audio briefing, a short video, and a structured knowledge asset that feeds AI retrieval systems.

Each placement extends the reach and the retrievability of the original message. Distribution puts your release on a wire. Placement puts it in front of people, and in front of the systems those people use to find answers.

Multi-Modal Reinforcement

Every release is raw material for multiple content formats: a social media summary, a short video, an audio briefing, a structured FAQ, a pillar page on your website. Each format reaches a different audience through a different channel. Each creates another indexed, retrievable asset that AI systems can find. The press release is the trunk of the tree. The multi-modal extensions are the branches.

News Provenance

We live in an era where trust in online information is eroding. Deepfakes, AI-generated articles, manipulated screenshots, fabricated quotes: the vol-

ume of synthetic and misleading content on the web is growing faster than anyone's ability to fact-check it. And when trust erodes, a fundamental question surfaces: where did this information actually come from?

A timestamp isn't enough anymore. Timestamps can be forged. Publication dates can be changed after the fact. Content can be scraped, rewritten, and republished with a new date on a different domain, with no trace of where it originated. In a retrieval-driven world where AI systems pull from indexed sources and synthesize answers, the integrity of the source matters enormously. If an AI system retrieves a fabricated claim from a site that copied your press release and altered the details, the damage flows downstream into every answer that system generates.

This is the problem of provenance: establishing an authoritative, verifiable, tamper-proof record of where a piece of content originated, who published it, and when.

NewsRamp is addressing this with blockchain technology. When a press release is published through any NewsRamp partner platform, we create an immutable record on the blockchain that captures the original content, the publishing entity, and the exact moment of publication. That record can't be altered, backdated, or deleted. It exists permanently as cryptographic proof of origin.

Why does this matter for AI visibility? Because provenance is becoming a trust signal. As AI systems get more sophisticated about evaluating source reliability, they need ways to distinguish between original authoritative content and copies, derivatives, or outright fabrications. A press release with blockchain-verified provenance carries a verifiable chain of custody. The AI system can confirm that this content was published by this entity at this time, and that the record hasn't been tampered with since.

Think of it as a notarized seal for your news. In the physical world, a notarized document carries more weight than an unsigned letter because a third party has verified its authenticity. Blockchain provenance does the same thing for digital content, except the verification is cryptographic, permanent, and doesn't depend on any single authority to maintain it.

This also solves a practical problem that press release distributors have struggled with for years: content attribution. When your release gets syndicated across dozens of outlets, scraped by aggregators, and rewritten by content farms, the original source can get lost. Blockchain provenance creates an anchor point. No matter how many copies exist, the immutable record points back to the original publication. Authority consolidates at the source, which is exactly what you want for entity reinforcement and canonical clarity.

Provenance isn't just a nice-to-have for the future. It's becoming infrastructure for trust in an AI-mediated information ecosystem. The organizations that can prove the origin and integrity of their published content will have a structural advantage over those that can't, because AI systems will increasingly need to answer a question before they answer the user's question: can I trust this source?

Built-In Amplification

Your press release isn't just an external communications tool. It's a sales enablement asset, an investor relations document, and a partner notification all at once. I've seen this coordination fail more times than I can count, and the failure mode is always the same: the marketing team publishes a press release, and the sales team doesn't know it exists. Built-in amplification means treating every release as a signal that gets routed to every stakeholder

who can use it.

The Long Fat Tail

When I ran PRWeb, we never removed content. Every press release stayed online, indexed, and accessible. This wasn't laziness. It was strategy. Indexed content compounds. A press release from three years ago can still rank for a relevant query, still be retrieved by an AI system, still reinforce the association between your brand and a specific topic. There are publications that delete press release content after 30-days. We choose not to send our content to them. We need the reliability of these news signals to endure for years.

This is why consistency matters more than virality. One great release that gets picked up by TechCrunch gives you a spike. Fifty well-structured releases published over four years give you a foundation. The spike fades. *The foundation compounds.*

Traditional PR Vs. News Marketing

The differences represent fundamentally different mental models. In traditional PR, your audience is the journalist, success is measured by impressions, and a release has a lifecycle of hours. In News Marketing, your audience is the customer and the Search and AI systems that serve them, success is measured by retrieval frequency and entity reinforcement, and a release has a life-cycle of months or years.

None of this means traditional PR is worthless. *Media relationships and earned coverage still matter.* But if that's your only strategy, you're playing a game that most organizations can't win. News Marketing gives every

organization a path to visibility with greater control.

Your customer isn't waiting for your press release to land in their inbox. They're searching Google, asking ChatGPT, scanning newsletters, watching LinkedIn videos. News Marketing is about meeting them in all of those places with a consistent message that originates from a single, well-structured source document. ***One release at a time, every 28 days.***

Chapter Two

How AI Actually Works: Training, Retrieval, and the Frozen Brain

Let me tell you the conversation I have at least a couple of times a week.

A marketing director or a PR agency owner will sit across from me, or more often these days, appear in a Zoom window, and say some version of this: "We need to get our client's information into ChatGPT. If we distribute the press release widely enough, it'll get picked up by AI."

I take a breath. Because the assumption behind that statement is so widespread, so deeply held, and so completely wrong that correcting it requires more than a quick "actually." It requires explaining how AI systems actually work. Not at a PhD level. At a level that lets a smart marketer make better decisions.

That's what this chapter is for.

The Frozen Brain

Here's the simplest, most useful mental model I've found for explaining

large language models to people who don't build them.

Imagine training a brain. You feed it trillions of words: books, websites, news articles, Wikipedia entries, academic papers, social media posts, forum threads, press releases. An unimaginable volume of text. The brain processes all of it, building connections, learning patterns, developing an understanding of language, facts, relationships, and context.

This training is a massive computational event. It costs millions of dollars with today's technology. It takes weeks or months. It involves billions, sometimes trillions, of parameters being tuned and adjusted until the model can generate coherent, contextually appropriate responses to almost any input.

And then the training stops.

The brain freezes.

From that moment forward, the model's internal knowledge is fixed. It has a knowledge cutoff date. Everything that happened before that date exists inside the model, encoded in its parameters. Everything that happened after that date is not.

When GPT-4 was released, its training data had a cutoff. When Claude was trained, same thing. When Gemini was trained, same thing. Every large language model has a point where its internal knowledge ends.

This is the Frozen Brain. It knows what it knew when training stopped, and nothing more.

Now, this metaphor isn't technically perfect. There are fine-tuning processes, model updates, and other nuances that AI researchers would

rightly point out. But as an operational model for marketers and PR professionals making strategic decisions? It's extremely useful. Because it kills the single most dangerous assumption in our industry: that AI is constantly learning from the internet in real time.

It isn't. The brain is frozen. Your press release from last Tuesday is not inside ChatGPT's memory. Neither is your competitor's. Neither is that industry report that came out last month.

So how does ChatGPT seem to know about recent events? How does Perplexity cite sources from yesterday? How does Google's AI Overview reference a news article from this morning?

The answer is retrieval.

Retrieval-Augmented Generation: The Five Steps

Retrieval-Augmented Generation, or RAG, is the mechanism that gives AI systems access to information beyond their training cutoff. And understanding how it works, step by step, is the single most valuable technical knowledge a marketer can have right now.

Let me walk you through it.

Step 1: The Question Becomes a Vector

When someone types a question into an AI system, the first thing that happens isn't a keyword search. The question gets translated into a vector, a numerical representation of its meaning.

Think of a vector as a coordinate in a vast, multi-dimensional space of meaning. The question "What are the best practices for AI governance

in financial services?" gets mapped to a specific point in that space. That point is close to other concepts like "regulatory compliance," "algorithmic oversight," "fintech regulation," and "model risk management."

This is semantic search, *not keyword search*. The system isn't looking for pages that contain the exact words in the question. It's looking for content that lives in the same neighborhood of meaning. This is a crucial distinction. *Old-school keyword stuffing doesn't help you here.* What helps is content that is genuinely, substantively about the topic in question.

Step 2: Your Content Gets Chunked

On the other side of the equation, every document in the retrieval index has already been processed. And the way it was processed matters enormously.

Documents don't get stored as whole pages. They get broken into chunks, typically 300 to 500 characters each. Each chunk gets its own vector, its own coordinate in that same multi-dimensional space of meaning.

This means every paragraph of your press release is evaluated independently. Your headline is one chunk. Your opening paragraph is another. Your quote from the CEO is another. Your boilerplate "about" section is another.

Each chunk stands on its own. Each chunk either has strong semantic meaning or it doesn't. Each chunk either clearly identifies your brand and its relevance to a topic, or it doesn't.

This is why paragraph density matters so much. A paragraph that says "We are excited to announce this innovative solution that will transform the way businesses operate" produces a weak vector. It could be about anything. It matches everything vaguely and nothing precisely.

A paragraph that says "NewsRamp's AI-verified press release distribution now includes structured entity tagging for financial services firms, enabling compliance teams to track media mentions across 47 regulatory categories" produces a strong vector. It's specific. It's brand-anchored. It's topic-dense. It occupies a precise point in semantic space.

Step 3: Similarity Matching

Once the question has been vectorized and the document chunks are sitting in the index with their own vectors, the system calculates similarity.

Which chunks are closest to the question vector? Which ones occupy the nearest points in semantic space? Those become the candidates.

This is where the quality of your content directly determines whether you get retrieved or not. Vague, metaphor-heavy, buzzword-filled press releases produce weak vectors that don't match precisely with any specific query. They float in the middle of semantic space, close to nothing in particular.

Brand-anchored, topic-specific, time-bound, fact-rich press releases produce strong vectors that match tightly with specific queries. They occupy clear, distinct positions in semantic space. When someone asks a question that falls in your territory, your chunks are right there.

This isn't keyword stuffing. It's semantic precision. And it rewards exactly the kind of clear, factual, structured writing that press releases are supposed to contain.

Step 4: Candidate Filtering

Similarity matching produces a pool of candidate chunks. But not all candidates make the final cut. AI systems apply additional filtering criteria

to determine which candidates are trustworthy enough to include in a generated response.

These filters vary by system, but commonly include:

Domain authority. Is this content hosted on a site that has demonstrated expertise and trustworthiness? A press release on a recognized news distribution platform carries more weight than the same text on a random blog.

Freshness. When was this content published or last updated? More recent content may be preferred for time-sensitive queries. This is why the timestamp on a press release matters so much.

Structured data. Does the page include schema markup, metadata, or other structured signals that help the system parse its contents? Well-structured pages are easier for AI systems to understand and trust.

Source trustworthiness. Is this source consistent in its claims? Does it contradict itself across different pages? AI systems are building increasingly sophisticated models of source reliability.

Entity consistency. Does this content use the same entity names, descriptions, and associations as other trusted sources? If every other source calls your company "Acme Corporation" and this one calls it "Acme Corp LLC doing business as Acme," that inconsistency creates *friction*.

Engagement patterns. Some systems factor in how users have interacted with this content. Has it been cited, shared, or referenced by other authoritative sources?

Step 5: The Model Summarizes

Finally, the top candidate chunks, the ones that survived similarity match-

ing and filtering, get fed into the model's context window. The model reads them and generates a response.

This is the critical point: the model is summarizing retrieved material, not pulling from memory. It's working like a very fast, very capable research assistant who has been handed a stack of relevant documents and asked to synthesize an answer.

If your press release chunks are in that stack, your brand, your facts, your message become part of the generated response. If they aren't in the stack, they don't exist as far as that response is concerned.

The AI isn't ignoring you on purpose. It simply never found you.

The Strategic Implication

Here's where I need you to make a mental shift that will inform every decision you make about AI visibility from this point forward.

The wrong question is: "How do we get into the model?"

You can't. Not directly. Not on your timeline. Model training is a periodic event controlled entirely by the AI companies. You don't get to submit your press release for inclusion. There's no form to fill out. There's no amount of distribution that guarantees your content ends up in the next training run. We can do things to boost the chances of your content being included in the next run. We do set permission flags for AI crawlers can even monitor their crawling activity. Thinking about the future is nice, but most of us need to drive immediate results.

The right question is: "How do we increase the probability that our content is retrieved when relevant questions are asked?"

That question you can answer. That question leads to actionable strategy. That question is what this entire book is about.

You increase retrieval probability by publishing content that is structured for chunking. That is semantically dense and topic-specific. That is hosted on authoritative domains. That is timestamped and fresh. That is consistent in its entity naming. That is factual, declarative, and verifiable.

In other words, you increase retrieval probability by publishing excellent press releases, consistently, on a predictable cadence.

Resolving The Confusion

Here's a distinction that matters: search engines index the web continuously. When you publish a press release today, Google can index it within minutes or hours, and AI retrieval systems can access it shortly after. But the AI model itself hasn't been retrained. It didn't "learn" your press release. It was handed your press release by the retrieval system and asked to summarize it.

What you can promise is increased retrieval probability: a well-structured release will be indexed quickly and available for retrieval. You cannot promise it will appear in every AI response or be cited by name. But compounding visibility that grows with every release you publish is something better than a guarantee. It's a strategy grounded in how the technology actually works.

Why Press Releases Are Built For This

Press releases are structured, timestamped, declarative, brand-anchored,

public, and persistent. They follow a predictable format that's easy for AI to parse. They make statements of fact that produce strong semantic vectors. They name entities clearly throughout. They live on the open web at stable URLs, available for indexing indefinitely.

Every one of those characteristics maps directly to what AI retrieval systems need. The press release wasn't designed for the AI era. But it turns out to be almost perfectly suited for it.

Your job is to take that natural advantage and make it intentional.

Chapter Three

AI Training, AI Grounding, and AI Visibility: The Three Pillars

When I talk to marketers about AI, I've noticed that most people lump everything into one vague category. "AI visibility" becomes a catch-all term for any interaction between their brand and an AI system. Getting mentioned in ChatGPT, showing up in a Google AI Overview, having Perplexity cite your website: it all gets tossed into the same bucket.

That's a problem. Because when you treat AI as a single, undifferentiated phenomenon, you can't build a targeted strategy. You end up doing a little bit of everything and nothing well.

There are actually three distinct layers to how AI systems interact with your brand's content, and each one operates on different mechanisms, different timelines, and different actions. We call them the Three Pillars: AI Training, AI Grounding, and AI Visibility.

Understanding the difference between them is the foundation of every strategy in this book.

AI Training: The Foundation Layer

Let's start with the layer most marketers find most mysterious, and most frustrating.

AI training is the process by which large language models build their internal knowledge, as I described in the previous chapter. The data included in training shapes the model's baseline awareness of your brand, your industry, and your competitors. If your company was well-represented in the training data, the model has some awareness of you. If it wasn't, the model might not know you exist, or worse, it might have partial or incorrect information that it presents with complete confidence.

You can't directly control what gets included in training data. *But you can influence what's available for future training cycles.*

What You Can Do to Influence AI Training

AI models are trained on web crawls, curated datasets, and publisher partnerships. Your content on high-authority domains has a better chance of making it into training data than content on restricted or low-authority domains.

You can stack the odds through consistent, strategic publishing. Publish on your own domain with clean HTML structure and stable URLs. Build legitimate presence on Wikipedia if your company meets notability standards (Wikipedia is one of the most heavily represented sources in training data). Make sure that your press release is added to the Press Release Archive on Github (https://github.com/Newsworthyai/press-release-archive/tree/main). Earn coverage in publications likely to be in training data, because a mention in a well-known industry publication is worth more for

training influence than a hundred social media posts. Maintain persistent URLs that don't change or disappear. And use consistent entity naming across all publications. If you're "Acme Corporation" in press releases and "Acme Corp" on your website, AI training processes may not connect these as the same entity.

The timeline for AI training influence is long, months to years. But the actions you take today determine what future models will know about you.

AI Grounding: The Verification Layer

If AI training is about what the model knows, AI grounding is about what the model can verify.

Grounding is the process by which AI systems connect their generated responses to real, retrievable, verified sources. Think of it as the fact-checking layer. When a modern AI system generates a response, it reaches out to external sources, retrieves relevant documents, and uses them to "ground" its answer in verifiable information.

Your published content becomes part of that grounding process. When an AI system retrieves your press release or FAQ page to support a statement, your content is serving as ground truth.

Here's why this matters so much. When an AI system can't find reliable grounding sources, it doesn't say "I don't know." It generates a response anyway, a plausible-sounding, confidently stated response that may be completely wrong. This is hallucination. If someone asks about your company and the system can't find solid grounding, you might get a response that confuses you with a competitor, states an incorrect founding date, or

makes up a partnership that never existed.

Your published content is your defense against hallucination.

What You Can Do to Strengthen AI Grounding

The actions for grounding are more immediate and controllable than those for training. You can start today.

Create definitive "about" pages with structured data. Your "about" page should be the single most authoritative source of facts about your organization: founding date, headquarters, employees, products, leadership, milestones. Include schema markup so this information is machine-readable.

Publish factual, verifiable claims with dates and numbers. "NewsRamp distributes press releases to over 450 media endpoints, with AI verification processing launched in March 2024" gives AI systems concrete facts to anchor responses to. "We're a leading provider of innovative solutions" gives them nothing.

Maintain FAQ pages that directly answer common questions. Think about what people might ask an AI system about your brand. Create pages that answer those questions directly and authoritatively. FAQ pages are among the easiest content types for AI systems to retrieve and use for grounding.

Use schema markup to make facts machine-readable. Organization schema, FAQ schema, Article schema: these tell AI systems exactly what your content is about. A page with proper schema is significantly easier for AI systems to parse and trust.

Ensure consistency across all your web properties. If your website says

you were founded in 2015 and your LinkedIn says 2014, you've created a conflict that undermines grounding. AI systems encountering conflicting information from the same entity may down-weight all of it.

Grounding operates on a faster timeline than training. Publish today, indexed within hours, available for AI retrieval within days.

AI Visibility: The Output Layer

AI visibility is the layer that marketers care about most, because it's the one they can see and measure (or at least try to measure). AI visibility is the probability that your brand appears in AI-generated responses when relevant questions are asked.

When someone asks ChatGPT "What are the best press release distribution services?" and your company shows up in the answer, that's AI visibility. When someone asks Perplexity "How does AI-verified news distribution work?" and your company is cited, that's AI visibility. When Google's AI Overview summarizes an answer to "What is news marketing?" and references your framework, that's AI visibility.

This is the output. This is what you're ultimately working toward. And it's the natural result of getting the first two pillars right.

AI visibility isn't something you can directly purchase or guarantee. What you can do is systematically increase the probability of being retrieved and included in responses through strong training data presence, solid grounding sources, and content optimized for retrieval.

What You Can Do to Increase AI Visibility

Monitor AI outputs for your category queries regularly. At least monthly, ask the major AI platforms the questions your customers are asking. Record whether your brand appears, which competitors appear, and whether the information is accurate. You can't improve what you don't measure.

Create content that directly answers the questions people ask AI. Most organizations create content about what they want to say, not about what people want to know. Start with the questions. Create press releases and FAQ pages that answer those specific questions with factual, brand-anchored responses.

Build topical authority through consistent publishing. Twelve press releases over a year, each covering a different facet of your expertise, creates a constellation of data points. AI systems start to see your brand as a reliable source on your topic.

Structure content for chunking. Every paragraph should stand alone as a meaningful, brand-identified statement. I call these "brandable chunks." When an AI system breaks your press release into chunks, each one should clearly identify your brand, state a specific claim, and relate to a topic your customers care about.

The Relationship Between All Three

Training builds baseline awareness. Grounding provides real-time verification. Visibility is the measurable outcome. All three require the same foundational discipline: structured, consistent, persistent publishing.

That's why the 28-day press release cadence is so powerful. It serves all three simultaneously. Every press release you publish is a training data candidate

for future model updates, a grounding source for today's retrieval systems, and a visibility asset that increases your brand's probability of appearing in AI-generated responses.

One action. Three layers of impact. Compounding over time.

You now have the framework. In the chapters ahead, we'll build the system.

Chapter Four

AIO vs GEO: The New Visibility Playbook

When I built PRWeb back in 1997, the game was straightforward. You published a press release, you got it indexed, you tried to rank in search results. That was it. The goal was a blue link on page one.

Today the game has three layers, and if you're only playing one of them, you're leaving visibility on the table.

Let me introduce you to two terms that every marketing and PR professional needs to understand right now: AIO and GEO. If you already know traditional SEO, think of these as its newer, faster siblings. They share the same DNA, but they've adapted to a world where AI systems don't just point users to content. They summarize it, synthesize it, and serve it directly.

What Is AIO?

AIO stands for AI Optimization, sometimes called AI Overview Optimization. It refers specifically to optimizing your content so it appears in the AI-generated summaries that now sit at the top of search results.

You've seen these. Google calls them AI Overviews. Bing has its Copilot summaries. When someone searches for "best AI governance platforms for mid-market companies," the search engine no longer just shows ten blue links. It shows a synthesized answer first, often pulling from multiple sources, and then lists the traditional results below.

That synthesized answer at the top is where AIO lives.

Here's the critical distinction from traditional SEO: with AIO, you're not just trying to rank a link. You're trying to be the information source that the AI system pulls from when it constructs that summary. Your content needs to be so clear, so structured, and so authoritative on the topic that the system selects it as reliable enough to synthesize.

Think of it this way. Traditional SEO gets your page listed. AIO gets your facts, your brand name, your claims woven directly into the answer a user sees before they click anything.

The implications are significant. If your competitor's data point appears in that AI Overview and yours doesn't, they just became the default answer. Not because they're necessarily better, but because their content was more retrievable, more structured, more clearly associated with the topic.

What Is GEO?

GEO stands for Generative Engine Optimization. It's the broader discipline.

While AIO focuses specifically on AI-generated summaries within search engines, GEO encompasses visibility across all generative AI systems. That means ChatGPT. Claude. Perplexity. Gemini. Any AI assistant that gen-

erates answers from retrieved sources.

When a potential customer asks Perplexity, "What companies offer automated regulatory reporting for financial institutions?" and your brand shows up in the response with a citation, that's GEO working.

GEO covers the entire pipeline: how your content gets discovered by these systems, how it gets retrieved, how it gets evaluated for reliability, and whether it gets cited or synthesized into an answer.

Research from institutions including Georgia Tech has demonstrated that specific content strategies measurably improve citation rates in generative engines[1]. This isn't theoretical. There are concrete tactics, and I'll walk you through them in a moment.

The important thing to understand is that GEO is not a single technique. It's a discipline. It's how you think about content creation, publication, and reinforcement when your audience includes machines that generate answers.

AIO vs GEO vs Traditional SEO

Let me lay these three approaches side by side so you can see how they relate.

Traditional SEO optimizes for ranking in search engine results pages. The goal is to appear high in the list of links. You're competing for position.

AIO optimizes for inclusion in AI-generated search summaries. The goal is to be a source the search engine's AI pulls from when constructing its overview. You're competing for synthesis.

GEO optimizes for visibility across all generative AI platforms. The goal is to appear, get cited, or have your information used whenever someone asks any AI system a relevant question. You're competing for retrieval everywhere.

These three overlap significantly. A well-structured press release with clear entity declarations, specific facts, and strong topical authority serves all three simultaneously. The foundational practices are the same: structured content, entity clarity, topical authority, consistent publishing, canonical hosting.

But the emphasis shifts. Traditional SEO cares a lot about backlinks and page authority. AIO cares more about whether your content is structured in a way that AI can confidently excerpt and attribute. GEO cares about whether your content is retrievable and citable across a wide range of systems, many of which have their own indexing and ranking methods.

For marketing and PR professionals, the practical takeaway is this: you don't need three separate strategies. You need one solid approach that satisfies all three. And that approach starts with clear, specific, consistently published content that machines can find, understand, and trust.

Practical GEO Strategies That Work

Research and real-world testing have identified specific tactics that improve your content's performance in generative engines.

Cite statistics and include quantitative data. "Reducing audit preparation time by 32%" is a specific, retrievable claim. "Significantly improving efficiency" is vague and unlikely to be selected as a source.

Use authoritative language with clear attributions. Instead of "We believe our platform may help some companies," write "Acme Corp's platform reduced audit preparation time by 32% across 47 mid-market financial institutions in Q1 2026."

Structure content with clear topic sentences per paragraph. Each paragraph is a potential retrieval chunk. Lead with the point, not with preamble.

Include direct answers to likely questions. "What does Acme Corp do?" should have a crisp, one-paragraph answer somewhere in your content. These become ideal retrieval targets.

Create comprehensive, definitive content on your domain topics. Being the most specific, factual, clearly structured source on topics that matter to your brand increases retrieval probability.

Use consistent entity naming everywhere. Pick one canonical name and use it in every release, every page, every channel.

Build topical clusters rather than isolated pages. One press release about AI governance is adjacency. Twelve over twelve months is a topical cluster. That's how entity salience compounds, and it brings us right back to the 28-day rule.

How Press Releases Fit Into AIO and GEO

Press releases are naturally built for this new reality, and I don't say that just because I've spent my career in this space.

Consider the structural characteristics of a well-written press release. It's timestamped. It contains entity declarations: who did what, when, with

what result. It makes authoritative claims in clear, declarative language. It's structured with a headline, a lead paragraph, supporting details, and a boilerplate. It lives on persistent URLs. It gets distributed to indexed domains.

Every one of those characteristics aligns with what AIO and GEO systems want.

A press release is, by its nature, a retrieval-friendly document. It answers who, what, when, where, why, and how. It does so in structured, factual language. It anchors claims to specific entities and timeframes.

When you publish a press release every 28 days, you're not just maintaining a communications cadence. You're feeding both AIO and GEO. You're creating fresh, retrievable signals. You're reinforcing entity-topic associations. You're building the topical clusters that generative engines use to evaluate authority.

The 28-day rule is a GEO strategy whether you call it that or not.

Monthly press releases generate timestamped facts that search engines index and AI systems retrieve. They create a trail of structured, entity-rich documents that strengthen your brand's position across every generative system with access to the indexed web. That's AIO. That's GEO. That's the work.

Your GEO Readiness Checklist

Score yourself honestly on these seven items:

1. Entity consistency: Does every piece of content use the same

canonical name for your organization, products, and key people?

2. Quantitative claims: Does your most recent press release include at least one specific, measurable data point?

3. Direct answer paragraphs: Can you identify at least three paragraphs that directly answer questions someone might ask an AI about your company or category?

4. Topical clustering: Have you published at least three pieces of content in the past six months addressing different facets of your core topic?

5. Schema markup: Does your press release page include Organization schema, NewsArticle schema, and sameAs references?

6. Canonical clarity and persistent hosting: Is there a clear canonical URL for each release, and are releases still accessible at their original URLs months later?

7. Publishing cadence: Have you published a structured, newsworthy announcement in the past 28 days?

If you checked five or more, you're in strong position. If fewer, the chapters ahead will give you specific tools to close those gaps.

Chapter Five

The 28-Day Rule and the Physics of Newsworthiness

Nearly two decades ago, I was on a phone call with Paul Allen, the founder of Ancestry.com and recently of CitizenPortal.ai, having one of those conversations that sticks with you permanently.[2]

He said something I've repeated hundreds of times since:

"If I'm not putting out a press release once a month, I'm not in business."

At first, that sounds like bravado. Or maybe just the instinctive rhythm of a serial entrepreneur who understood marketing intuitively.

It isn't either of those things.

It's physics.

The Digital Pulse

In the physical world, living organisms have a pulse. Stop that pulse and the organism is declared dead, even if the body is still structurally intact.

In the digital world, brands have a pulse too. Every time you publish a

meaningful, structured announcement, you send a pulse into the ecosystem. That pulse does several things simultaneously. It reintroduces your entity into the search index. It reinforces your brand's association with a topic. It creates fresh time anchors. It produces new candidate chunks for retrieval systems. It signals activity to both humans and machines.

When those pulses occur at regular intervals, the brand appears alive. Active. Current. Relevant.

When they stop, the brand appears dormant.

Dormancy isn't immediately fatal. But it's cumulative. And in a retrieval-driven world where AI systems prioritize fresh, reinforced sources, dormancy carries a real cost.

The Half-Life of Authority

Let me borrow a concept from physics: half-life. In radioactive decay, half-life measures the time it takes for half of a substance to decay. Digital authority isn't radioactive, but it does exhibit decay patterns that are useful to understand.

After you publish a press release, several things happen at once. Engagement spikes as the release circulates. Search crawlers index the new content. Freshness weighting increases, giving your content a temporary boost in relevance. Topic proximity strengthens as your brand becomes more closely associated with the subjects you've written about.

Then, over time, the decay begins. Engagement declines as the news cycle moves on. Freshness signals taper. Competing content enters the index. Other companies publish their own announcements in your category.

Topic clusters shift.

Nothing disappears. Your content is still there. But its relative weight changes. Authority decays comparatively. You don't become invisible. You become less prominent relative to the entities that kept publishing while you stayed quiet.

The 28-day rule isn't superstition. It isn't a marketing gimmick. It's reinforcement against decay. Every new release resets the freshness clock, adds new candidate chunks to the retrieval pool, and strengthens the association edges that connect your brand to your topics.

The "Golden Rule" of 28 Days

"If an AI hasn't seen a fresh signal from your brand in the last 28 days, you don't have a presence; you have a history."

- **Greg Isenberg**, CEO of Late Checkout

Why 28 Days?

People always ask why specifically 28 days. Why not weekly? Why not quarterly?

The answer comes down to four factors.

Operational feasibility. Most organizations can realistically produce one well-crafted, genuinely newsworthy announcement per month. Weekly often forces teams to manufacture news that isn't really news, diluting the quality and credibility of everything they publish. Quarterly creates gaps wide enough for competitors to dominate your topic space.

Meaningful progress. In 28 days, a real company making real progress will

have something legitimate to announce. A feature ships. A partnership formalizes. A data point emerges. If nothing newsworthy has happened in a month, that's a signal worth examining, and I'll come back to that.

Reinforcement cadence. Monthly publishing creates a rhythm that both search engines and AI systems can recognize. Twelve pulses per year means twelve fresh time anchors, twelve new sets of candidate chunks, twelve opportunities to reinforce entity-topic associations. That's enough to build meaningful salience without overwhelming the signal.

Avoiding noise. Publish too frequently with thin content and you risk training machines and humans to discount your announcements. Every release should carry weight. Monthly cadence makes that sustainable.

Finally there is the Google News challenge. Remember our discussion about retrievability and the core concept that News should be new? You news may linger in the Google News index for longer than 28 days, but our experience shows that anything older than 28 days is ignored in the AI retrieval process.

Internal Discipline

Here's something most people don't talk about when discussing press release cadence: the internal effect.

The 28-day rule isn't just an external visibility strategy. It creates accountability inside the organization.

To publish something legitimately newsworthy every month, you have to do things worth announcing. You must ship improvements. Formalize partnerships. Release data. Clarify positioning. Measure results.

If nothing feels newsworthy after 28 days, one of two things is happening. Either progress is not being made, or progress is not being articulated. Both are solvable, but both need attention.

I've watched companies adopt a monthly press release cadence and discover that the discipline of finding their news every month sharpened their entire operation. Product teams started tracking milestones differently. Partnership discussions moved faster because someone needed to announce them. Data that had been sitting in spreadsheets got analyzed and published. The release became a deadline that forced clarity.

The cadence forces you to articulate what you've accomplished, what it means, and why anyone should care. That exercise alone is valuable, even before a single AI system retrieves a single chunk.

Newsworthiness in the AI Era

Historically, newsworthiness was defined by editorial judgment. A journalist decided whether your announcement was worth covering. That gatekeeper model still exists, and it still matters. But it's no longer the only definition that counts.

In the AI era, newsworthiness is also defined by retrievability.

A release becomes newsworthy in the retrieval sense when it introduces new information, anchors that information in time, connects it to a recognizable category, and provides specific, measurable detail.

Generic announcements weaken semantic vectors. "We are pleased to announce continued innovation" tells a retrieval system nothing useful. It contains no specific entity associations, no measurable claims, no time

anchors, no category connections.

Specific announcements strengthen semantic vectors. "In April 2026, Acme Corp released version 4.2 of its AI governance platform, introducing automated audit trail generation for mid-market financial institutions" contains brand entity, product entity, category entity, time anchor, and functional specificity. That paragraph is a retrieval candidate. The generic one isn't.

Signal Stacking

This is where the 28-day rule becomes something more powerful than a publishing cadence. It becomes a compounding strategy.

Each release builds on the previous one.

Month 1: You announce your AI governance framework launch. Your brand is now associated with AI governance in the index.

Month 2: You announce an integration with a regulatory compliance platform. Now your brand is associated with both AI governance and regulatory compliance, and the two topics are linked through your entity.

Month 3: You release proprietary data on audit efficiency. Now you've added quantitative authority. Your brand isn't just talking about AI governance. It's producing measurable data in the field.

Month 4: Your CEO publishes executive commentary on new AI regulations. Now you've added thought leadership, a named person entity connected to your organization entity, both connected to the regulatory topic.

By month four, you haven't just published four announcements. You've stacked four signals. Each signal reinforces the others. Each one adds new edges to the knowledge graph. Each one creates new candidate chunks that retrieval systems can select.

Four signals stacked in the same topic domain creates something much more powerful than four isolated announcements about unrelated subjects. It creates entity salience.

Entity Salience: The Accumulation Effect

Entity salience is the strength of association between your brand and a topic inside retrieval systems.

If your brand appears once near "AI governance," that's adjacency. The system notes a connection, but it's thin. Lots of entities appear once near lots of topics.

If your brand appears twelve times in structured proximity to AI governance, regulatory compliance, audit preparation, and financial institutions, that's association. That's different. That's entity salience.

Association strengthens retrieval confidence. Not because the system "likes" you. Not because you've earned special status. Simply because repeated proximity reduces ambiguity. When the system needs to answer a question about AI governance for financial institutions, it has twelve separate data points connecting your brand to that exact topic. Your competitor who published twice this year has two.

Machines prefer lower ambiguity. They prefer entities with clear, reinforced associations. The math favors consistency.

One appearance equals adjacency. Twelve appearances equals association. That's the accumulation effect, and the 28-day rule is the mechanism that produces it.

Silence Is Not Neutral

Many organizations assume that if they stop publishing, nothing changes. They think silence is a pause, a neutral state they can resume from later without consequence.

Silence is not neutral.

Here's what actually happens when you stop publishing. Your freshness signals decay. Your topic associations weaken comparatively. And your competitors keep going. They keep publishing. They keep reinforcing. They keep stacking signals.

While your salience plateaus, theirs increases. Over time, relative authority shifts. You don't disappear. But you become comparatively less present. In retrieval environments, where systems select from candidate pools and evaluate relative authority, comparative presence is what matters.

I've seen this play out with dozens of companies. They publish consistently for six months, see real improvement in their visibility metrics, and then decide they can coast for a quarter. When they come back, they haven't lost everything, but they've lost ground. And the ground they lost went to competitors who never stopped.

What Counts as News Every Month?

The most common objection I hear: "We don't have news every month."

You probably do. You're just not labeling it correctly.

News doesn't have to mean a massive product launch. It means new, specific, and relevant. Feature enhancements. Integration announcements. Proprietary data insights. Executive commentary on industry trends. Customer case studies with measurable outcomes. Operational expansion driven by demand. New hires, employee recognition.

Each month, audit seven categories: product and engineering, partnerships and integrations, data and insights, people and leadership, customer outcomes, market and regulatory developments, and operations and growth. If you can't find at least one newsworthy item across all seven, the problem isn't a lack of news. It's a lack of internal communication. Talk to your product team, your sales team, your customer success team. *The news is there. Your job is to find it, structure it, and publish it every 28 days.*

Chapter Six

Entity Authority and the Knowledge Graph

From Keywords to Entities

Early SEO was blunt. If you wanted to rank for "financial compliance software," you repeated that phrase. You stuffed it into headers, meta tags, alt text, and body copy. You built backlinks with exact-match anchor text. You optimized your way up the rankings through sheer keyword persistence.

That world is gone.

I watched it happen in real time. When I launched PRWeb in 1997, keywords and back links were king. By the time I was building Newsworthy and NewsRamp, the shift was well underway. Search engines had gotten smarter. They stopped counting keywords and started identifying entities.

Today, search engines identify and disambiguate entities: organizations, people, products, locations, regulations, categories. They understand that "Apple" the technology company is different from "apple" the fruit. They understand that "David McInnis" the person is connected to "PRWeb" and

"Newsworthy.ai" the companies. They understand relationships, not just words.

This is the **transition from strings to things**. And it fundamentally changes how your content needs to work. [3]

When a search engine encounters your press release now, it doesn't just count how many times you mentioned "AI governance." It asks: What entity is making this claim? What topics is this entity associated with? How consistently? With what authority? In what time frame? In what structured contexts?

That's entity authority. And it's the foundation of visibility in the AI era.

What Entity Authority Actually Is

Entity authority isn't a score you can look up. There's no dashboard that shows your entity authority rating. It's an emergent property, something that arises from the accumulation of consistent signals over time.

An entity gains authority in a topic when several conditions are met. It appears repeatedly near that topic in structured, indexed content. It provides specific, time-anchored claims about that topic. It co-occurs with other trusted entities in the same domain. It maintains persistent, stable hosting for its content. And it avoids contradictory or ambiguous signals.

Authority is less about loudness and more about clarity. You don't build entity authority by publishing one massive announcement with a huge distribution push. You build it through consistent, clear, specific publishing over months and years.

Ambiguous entities struggle. If your brand sometimes claims to be in AI

governance, sometimes in general consulting, sometimes in fintech, and sometimes in healthcare IT, the system can't confidently associate you with any single domain. You become a weak signal in every category instead of a strong signal in one.

> Consistency is not about spamming; it is about reinforcing the neural pathway between your **Brand Entity** and your **Topic Entity**.

Consistent entities compound. If you publish monthly about AI governance for financial institutions, every release adds weight to that association. After twelve months, the connection between your brand and that topic is reinforced by twelve separate, timestamped, structured data points.

E-E-A-T in Practice

You've probably encountered the acronym E-E-A-T: Experience, Expertise, Authoritativeness, Trustworthiness. It comes from Google's quality evaluation framework, but the principles apply across all retrieval systems. Let me break down what each element looks like in practice, not as abstract concepts, but as concrete things you can demonstrate through your press releases.

Experience means real-world outcomes. When your release includes a case study showing that your platform reduced audit preparation time by 32% across 47 financial institutions, that demonstrates experience. You're not theorizing. You're reporting results from actual deployments. Systems evaluating your content can distinguish between "this entity has done the thing" and "this entity is talking about the thing."

Expertise shows up in precise language. When your press release correctly

references specific regulations by their proper names, uses accurate industry terminology, and demonstrates detailed knowledge of the domain, that signals expertise. Vague language signals the opposite. "Helping companies with compliance stuff" versus "enabling automated SOC 2 Type II reporting for mid-market financial institutions" conveys a very different level of domain knowledge.

Authoritativeness grows through repetition and recognition. This is where the 28-day rule directly serves E-E-A-T. Each monthly release in your topic domain reinforces your authority. Earned mentions from other sources, co-occurrence with established industry entities, and consistent topic clustering all contribute. Authority isn't claimed. It's accumulated.

Trustworthiness increases when your claims are measurable, your sourcing is clear, and your infrastructure is stable. Measurable claims can be verified. Clear sourcing means attributing statements to named individuals with real titles. Stable infrastructure means persistent URLs, proper canonical tags, consistent schema markup, and reliable hosting. When a retrieval system evaluates your content, these signals reduce uncertainty. Systems prefer lower uncertainty.

The Knowledge Graph as a Map

Imagine the Knowledge Graph as a map. Not a ranked list, but a web of connected nodes and edges.

Nodes represent entities: your company, your competitors, your industry, your executives, your products, regulatory bodies, standards organizations, geographic locations.

Edges represent relationships between those nodes: "Acme Corp

provides AI governance," "Jane Doe is CEO of Acme Corp," "Acme Corp is certified for SOC 2 compliance."

When your brand appears once in relation to "AI governance," the edge connecting your node to that topic is thin. The system notes the relationship, but with low confidence.

When your brand appears monthly in structured, timestamped content about AI governance, regulatory compliance, financial institutions, and audit automation, those edges thicken. Each publication adds weight. Each new release reinforces the connection.

Thicker edges increase confidence in the relationship. When retrieval systems evaluate candidate documents in response to a query about AI governance for financial institutions, they prefer entities with well-defined, reinforced relationships. Not because they favor incumbents, but because reinforced relationships reduce ambiguity. Machines are risk-averse. Clarity lowers risk.

This is why consistency matters more than any single brilliant release. One perfect press release creates thin edges in several directions. Twelve good press releases, each reinforcing the same topic cluster, create thick, confident connections that retrieval systems trust.

Canonical Anchors and the Source of Truth

One of the most overlooked aspects of entity authority is canonical clarity. This is a technical concept with major strategic implications.

When your press release appears on your website, on a newswire, on syndicated partner sites, and on scraped aggregator domains, search systems

face a question: which version is authoritative? Which one should receive the credit for the claims made in the content?

A canonical tag answers that question. It tells search engines, "This is the primary source. All other versions are copies of this one."

Without canonical clarity, authority fragments. The engagement, the backlinks, the trust signals all get split across multiple copies of the same content. Fragmented authority weakens entity consolidation. Instead of one strong signal, you have five weak ones.

The best practice is straightforward. Publish on your own domain first. Apply canonical tags correctly, pointing to your domain as the authoritative source. Then distribute outward to newswires and syndication partners. Ensure the primary version remains persistent and accessible at its original URL.

Authority accumulates at the canonical source. Make sure that source is yours.

Brandable Chunks

Let me return to a concept from earlier in this book, but now through the lens of entity authority. Incidentally, "Brandable Chunks" is a component of copy writing framework that I learned from Roy Williams at the Wizard Academy[4]. I am taking liberties to use this term here because I think it is the best way to describe how to write for AI. AI doesn't think in terms of keywords, it takes in and retrieves content based on complex vector calculations. Knowing how to be concise within a given chunk of copy will help you surface your content in AI. Here is how that works.

Retrieval systems don't evaluate your press release as a single monolithic document. They break it into chunks, typically 300 to 500 characters, and evaluate each chunk independently. Each chunk is a candidate unit that might get selected in response to a query. [5]

This means the density of information in each paragraph matters enormously.

Here's a weak chunk: "We are thrilled about the future of our company and look forward to continuing to innovate in this space."

That contains no brand entity, no product entity, no category entity, no time anchor, no measurable claim. It's semantically empty. No retrieval system will select it because it doesn't answer any specific question.

Here's a strong chunk: "In May 2026, Acme Corp expanded its AI governance platform to include automated regulatory reporting for mid-market financial institutions, reducing audit preparation time by 32%."

Count what's packed into that single sentence. Brand entity: Acme Corp. Product entity: AI governance platform. Category entity: regulatory reporting. Audience entity: mid-market financial institutions. Time anchor: May 2026. Measurable claim: 32% reduction. Functional specificity: automated regulatory reporting and audit preparation.

That chunk strengthens the brand-topic edge, the brand-category edge, the brand-time association, and the brand-performance association. It produces a dense semantic representation that scores well during retrieval.

The difference between these two chunks is the difference between visibility and invisibility in the AI era. Every paragraph in every press release you publish should aim for that density. Not because it reads like a data dump

to humans, but because strong chunks serve both audiences. Humans read specific, factual claims and find them credible. Machines process them and find them retrievable.

Co-occurrence as Authority Amplifier

Entity authority doesn't exist in isolation. It strengthens when your brand co-occurs with other recognized entities.

When your press release mentions that your platform supports SOC 2 Type II compliance, you're creating a co-occurrence between your brand entity and a recognized security standard. When you announce a partnership with an established industry player, you're creating co-occurrence between your entity and theirs. When you reference regulatory bodies, industry associations, or established standards, each reference adds a relational edge to the knowledge graph.

This isn't about name-dropping for prestige. It's about contextual mapping. When your entity repeatedly appears alongside trusted entities in structured documents, the Knowledge Graph interprets relational proximity. That proximity reinforces your authority in the domain.

Appearing alongside SOC 2 and ISO 27001 in the context of financial compliance tells the system that your entity operates in a specific, well-defined, regulated space. That's a signal of legitimacy. It reduces the system's uncertainty about what your brand does and whether it's credible in that domain.

Use co-occurrence intentionally. When you have legitimate associations with recognized standards, regulatory frameworks, industry bodies, and established partners, make those associations explicit in your press releases.

Don't assume the system will figure it out from context. Declare it clearly.

Consistency Over Complexity

Here's the liberating truth about entity authority: you don't need technical wizardry to build it. You don't need a team of data scientists. You don't need proprietary tools or secret algorithms.

You need consistency.

Monthly releases aligned to a defined topic domain. Clear headlines that include your brand name and your topic. Specific facts with numbers, dates, and named entities. Structured metadata that helps machines categorize your content correctly. Persistent hosting that keeps your content accessible indefinitely. Canonical clarity that tells search engines where authority should consolidate.

That's it. Complex tactics without consistency fail. I've seen companies invest thousands in one-time SEO overhauls, only to publish nothing for six months afterward. The overhaul decays. The consistency of a competitor who publishes simple, clear releases every 28 days compounds.

Consistency without complexity often succeeds. A company that publishes a straightforward, well-structured press release every month for a year, with clear entity naming, specific claims, and proper canonical tags, will build more entity authority than a company that publishes one technically perfect, schema-loaded release and then goes dark.

The Press Release as Entity Declaration

A well-constructed press release is more than an announcement. It's an

entity declaration.

It says: This organization exists. It operates in this category. It achieved this measurable result. It did so at this specific time. It is associated with these other entities, these standards, these regulatory frameworks.

Every element of that declaration becomes part of the knowledge graph. The organization node gets reinforced. The topic edges get thicker. The time anchors get refreshed. The measurable claims provide specificity that reduces ambiguity.

In a retrieval-driven ecosystem, graph clarity influences candidate selection. When a query comes in and the system evaluates which sources to retrieve, entities with clear, consistent, well-reinforced graph presence get the nod.

Your press release is how you make that declaration, month after month. Not once and done. Not when you feel like it. Every 28 days, with precision, specificity, and consistency.

That's how entity authority is built. Not in a single moment, but in the accumulation of moments. One structured declaration at a time.

Chapter Seven

Schema Markup: Your Complete Blueprint

If the previous chapters explained why entity authority matters, this chapter explains how you make it explicit. This is where you stop relying on machines to infer meaning from your content and start telling them directly.

What Structured Data Actually Does

When a human reads your press release, they understand context automatically. They know who the organization is. They know who's speaking. They know what category the product belongs to, when the announcement happened, and whether something is a statistic, a quote, or a claim.

Machines must infer those things unless you label them.

Structured data is the labeling system. It translates your press release from narrative text into explicit, machine-readable facts. Think of it as moving from strings to things.

Without structured data, your brand name is a string of characters. The system has to figure out from context whether "Mercury" refers to a planet,

an element, a car brand, or a software company. With structured data, you explicitly declare: this is an organization, this is its name, this is its URL, this is its industry.

Clarity reduces interpretive error. Reduced ambiguity increases retrieval confidence. And that's what structured data provides: clarity at the machine level.

Now, structured data doesn't guarantee that AI systems will cite you. Let me be clear about that. It doesn't force inclusion or ensure retrieval. What it does is reduce friction. When a retrieval system evaluates your content against competing sources, structured data makes your content easier to classify, easier to attribute, and easier to trust. Lower friction increases selection probability.

Don't Panic: I know that this might look scary, however you will probably never have to deal with the internals of schema. You just need to know the topic well enough to 1) select a service provider that cares enough about AI to include this for you or 2) know how to ask your internal tech team to implement this for you. I know that the Newsworthy.ai and NewsRamp and BoostifyPR teams take care of this for you.

JSON-LD: The Preferred Format

JSON-LD, which stands for JavaScript Object Notation for Linked Data, is the most widely supported structured data format for web pages. It's the format Google explicitly recommends. It's the format that works cleanly

with modern content management systems.

> **Note:** While JSON-LD is the preferred schema language of AI, it is not the only schema format. You can add **Microdata** schema to the HTML tags by adding the *'itemscope'*, *'itemtype'*, *'itemprop'* and *'typeof'* attributes to your HTML. You can also add **RDFa** schema to your HTML adding the *'vocab'*, *'typeof'*, and *'property'* attributes. We have built this into our content editors across all of our platforms. Wordpress has plugins that you can use for this.

JSON-LD gets placed inside a script tag in your page's header, like this:

<script type="application/ld+json">{...your structured data here...}</script>

The critical thing to understand is that this code doesn't change what your visitors see on the page. It changes what machines understand about the page. Your human readers will never know it's there. Search engines and AI retrieval systems will.

Let me walk through each schema type that matters for press releases, with complete examples you can adapt for your own content.

1. Organization Schema

This is your foundation. Organization schema declares your company as an entity, with all the details machines need to disambiguate you from every other organization in the world.

Here's a complete example with all recommended fields:

Illustration 7.1

```
<script type="application/ld+json">
{
  "@context": "https://schema.org",
  "@type": "Organization",
  "name": "Acme Corp",
  "url": "https://www.acmecorp.com",
  "logo": "https://www.acmecorp.com/images/logo.png",
  "description": "Acme Corp provides AI governance and automated regulatory reporting
solutions for mid-market financial institutions.",
  "sameAs": [
    "https://www.linkedin.com/company/acme-corp",
    "https://www.crunchbase.com/organization/acme-corp",
    "https://twitter.com/acmecorp",
    "https://en.wikipedia.org/wiki/Acme_Corp"
  ],
  "contactPoint": {
    "@type": "ContactPoint",
    "telephone": "+1-555-123-4567",
    "contactType": "media inquiries",
    "email": "press@acmecorp.com"
  },
  "address": {
    "@type": "PostalAddress",
    "streetAddress": "100 Innovation Drive",
    "addressLocality": "Austin",
    "addressRegion": "TX",
    "postalCode": "78701",
    "addressCountry": "US"
  },
  "foundingDate": "2019-03-15",
  "numberOfEmployees": {
    "@type": "QuantitativeValue",
    "value": 185
  },
  "industry": "Enterprise Software"
}
</script>
```

Sample Organization JSON-LD Schema

When to use it: every page on your website should include Organization schema. Every press release page, absolutely. This is the single most important schema type for entity disambiguation. It tells machines exactly who you are, where you are, how to find your verified profiles elsewhere on the web, and what domain you operate in.

The description field is especially valuable. Write it like a one-sentence elevator pitch that includes your brand name, your product category, and your target market. That description becomes a retrievable chunk of its own.

2. NewsArticle Schema

This schema type tells machines that your page is a news announcement, not a blog post, not a product page, not a general article. That classification matters because it signals timeliness and announcement format, both of which influence how retrieval systems evaluate the content.

Illustration 7.2

```
<script type="application/ld+json">
{
  "@context": "https://schema.org",
  "@type": "NewsArticle",
  "headline": "Acme Corp Launches Automated Regulatory Reporting for Mid-Market Financial
Institutions",
  "datePublished": "2026-05-15T09:00:00-05:00",
  "dateModified": "2026-05-15T09:00:00-05:00",
  "author": {
    "@type": "Person",
    "name": "Jane Doe",
    "jobTitle": "Chief Executive Officer",
    "worksFor": {
      "@type": "Organization",
      "name": "Acme Corp"
    }
  },
  "publisher": {
    "@type": "Organization",
    "name": "Acme Corp",
    "logo": {
      "@type": "ImageObject",
      "url": "https://www.acmecorp.com/images/logo.png"
    }
  },
  "mainEntityOfPage": {
    "@type": "WebPage",
    "@id": "https://www.acmecorp.com/news/automated-regulatory-reporting-launch"
  },
  "image": "https://www.acmecorp.com/images/news/regulatory-reporting-launch.jpg",
  "articleBody": "Acme Corp today announced the launch of automated regulatory reporting
capabilities within its AI governance platform...",
  "keywords": [
    "AI governance",
    "regulatory reporting",
    "financial compliance",
    "mid-market",
    "audit automation"
  ]
}
</script>
```

Sample NewsArticle JSON-LD Schema

Notice how the author field nests a Person schema inside the NewsArticle.

This creates linked entities: the article is authored by a specific person who works for a specific organization. Those connections reinforce the knowledge graph.

The datePublished field is crucial. It provides the time anchor that retrieval systems use when evaluating freshness. Make sure it includes the full ISO 8601 format with timezone.

The keywords field gives machines an explicit topic list. Don't overload it. Five to eight relevant keywords that accurately describe the content. These should align with the topical cluster you're building month after month.

3. Person Schema

When your press release quotes an executive, Person schema helps AI systems understand exactly who is making claims. This matters because authority isn't just about organizations. It's about the people within those organizations.

Illustration 7.3

```
<script type="application/ld+json">
{
  "@context": "https://schema.org",
  "@type": "Person",
  "name": "Jane Doe",
  "jobTitle": "Chief Executive Officer",
  "worksFor": {
    "@type": "Organization",
    "name": "Acme Corp",
    "url": "https://www.acmecorp.com"
  },
  "sameAs": [
    "https://www.linkedin.com/in/janedoe",
    "https://twitter.com/janedoe_acme"
  ],
  "image": "https://www.acmecorp.com/images/team/jane-doe.jpg"
}
</script>
```

Sample Person JSON-LD Schema

Why this matters: when an AI system retrieves a quote from your press release, the Person schema tells it that this quote comes from a named individual with a specific title at a specific organization, with verified profiles on LinkedIn and other platforms. That's more credible than an unattributed quote from an anonymous source. It supports the Experience and Expertise dimensions of E-E-A-T.

If you regularly quote the same executive in your monthly releases, their Person entity accumulates authority in your topic domain too. Over time, "Jane Doe, CEO of Acme Corp" becomes a recognized voice in AI governance, not just within your company, but within the knowledge graph.

4. Publisher Schema

Publisher schema declares your organization as the entity responsible for publishing the content. While the publisher field already appears nested within NewsArticle schema, you can strengthen it by adding an @id field that creates a referenceable identifier. Link to this publisher entity from every NewsArticle on your site, creating a consistent chain where all trust signals accumulate at one node. If you have an editorial standards page, reference it via the publishingPrinciples field as an additional trust indicator.

5. FAQ Schema

FAQ schema is one of the most underutilized opportunities in press release optimization. It can make your content appear directly in AI-generated answers, and it's straightforward to implement.

Think about it: when someone asks an AI assistant "What does Acme Corp do?" or "How does automated regulatory reporting work?", the system looks for direct, concise answers. FAQ schema provides exactly that, pre-formatted and clearly labeled.

Figure 4

```
<script type="application/ld+json">
{
  "@context": "https://schema.org",
  "@type": "FAQPage",
  "mainEntity": [
    {
      "@type": "Question",
      "name": "What is Acme Corp's AI governance platform?",
      "acceptedAnswer": {
        "@type": "Answer",
        "text": "Acme Corp's AI governance platform provides automated regulatory reporting,
audit trail generation, and compliance monitoring for mid-market financial institutions. The
platform reduces audit preparation time by an average of 32%."
      }
    },
    {
      "@type": "Question",
      "name": "Which industries does Acme Corp serve?",
      "acceptedAnswer": {
        "@type": "Answer",
        "text": "Acme Corp primarily serves mid-market financial institutions, including
regional banks, credit unions, and investment firms with between 500 and 5,000 employees."
      }
    },
    {
      "@type": "Question",
      "name": "How does automated regulatory reporting work?",
      "acceptedAnswer": {
        "@type": "Answer",
        "text": "Acme Corp's automated regulatory reporting system continuously monitors
compliance requirements, generates required documentation, and produces audit-ready reports
without manual data compilation. Institutions using the system report a 32% reduction in audit
preparation time."
      }
    }
  ]
}
</script>
```

Sample FAQ JSON-LD Schema

Each question-answer pair becomes a standalone retrieval candidate. When an AI system encounters a query that matches one of your FAQ questions, the structured answer is right there, clearly labeled, immediately usable.

The key to effective FAQ schema for press releases: think about the questions that your announcement naturally raises, then answer them explicitly. If you're announcing a new product feature, include FAQs about what it does, who it's for, and what results it delivers. If you're announcing a partnership, include FAQs about what the partnership enables and who

benefits.

Three to five well-crafted FAQ pairs per press release can significantly expand your content's retrievability. Each pair is a new entry point for AI systems seeking direct answers.

6. Product Schema

When your press release announces a specific product or feature, Product schema declares its details. The key fields are name, description, brand (linking to your Organization), releaseDate, and category. This connects your product entity to your organization and to a specific release date, creating multiple knowledge graph edges from a single declaration. Use it when announcing new products, features with their own identity, or significant updates. Not every release needs it.

7. Event Schema

For event-related releases (conferences, webinars, summits), Event schema declares name, startDate, endDate, location, organizer, and description. The pattern is the same: declare what it is, when and where it happens, and who's involved. Check schema.org/Event for the complete field list.

The sameAs Attribute

The sameAs property tells search engines: "This entity is the same as this verified profile on another platform." Include your LinkedIn company page, Crunchbase profile, X account, and Wikipedia page. This consolidates your identity across the web and strengthens entity disambiguation. For Person schema, sameAs is equally important: link executives to their

LinkedIn profiles. Keep all sameAs references current. Broken links create ambiguity.

How Structured Data Interacts With Retrieval

Structured data influences multiple steps in the retrieval pipeline. Organization, Person, and Product schema help with entity disambiguation. NewsArticle schema affects freshness weighting. Publisher schema signals source stability. Together, they reduce friction at every evaluation step.

Pair your structured data with canonical clarity. Your canonical URL and your mainEntityOfPage should point to the same URL on your domain. When they align, authority consolidates cleanly. When they don't, signals fragment.

The Minimum Viable Schema Checklist

You don't need an over-engineered schema system. You need the fundamentals, applied consistently to every press release you publish. Here's your implementation checklist.

Step 1: Add Organization schema to your website's press release template. Do this once. Include name, url, logo, description, sameAs, contactPoint, and address. This becomes the default for every page.

Step 2: Add NewsArticle schema to each press release page. Include headline, datePublished, dateModified, author, publisher, mainEntityOfPage, and keywords. Update this for each new release.

Step 3: Add Person schema for any executive quoted in the release. Include name, jobTitle, worksFor, and sameAs. If you quote the same executive

regularly, create a reusable template.

Step 4: Add FAQ schema with three to five question-answer pairs relevant to the announcement. Think about what questions the release naturally raises and answer them explicitly.

Step 5: Add Product schema when the release announces a specific product or feature. Include name, description, brand, releaseDate, and category.

Step 6: Verify your canonical URL is set correctly and matches the mainEntityOfPage in your NewsArticle schema.

Step 7: Validate your markup using Google's Rich Results Test or Schema Markup Validator. Fix any errors or warnings before publishing.

Step 8: Keep your sameAs references current across all schema types. Audit them quarterly.

Step 9: Ensure your structured data accurately reflects the content on the page. Don't declare entities that aren't mentioned. Don't claim relationships that don't exist.

Step 10: Apply this checklist to every press release, every month. Consistency matters more than perfection. A slightly imperfect schema applied to twelve releases per year beats a perfect schema applied once and then forgotten.

A Warning About Over-Engineering

Some marketers discover structured data and go overboard. They add every schema type they can find. They declare entities that don't exist in their

content. They create complex nested structures for marginal gains.

Don't do this.

Structured data must reflect reality. If your press release doesn't mention a product, don't include Product schema. If nobody is quoted, don't include Person schema. If there's no FAQ section in your content, don't fabricate FAQ schema.

Misleading or inaccurate markup can trigger quality filters. Search engines have gotten sophisticated at detecting schema that doesn't match page content. At best, inaccurate markup gets ignored. At worst, it signals that your site is untrustworthy.

Accuracy over complexity. Always.

The goal is clear, consistent, honest labeling of what's actually in your content. When you do that month after month, you build a track record of reliable structured data that machines learn to trust. That trust compounds just like entity authority does, one release at a time.

Your press release tells the story. Your structured data translates that story into labeled relationships. The story persuades humans. The labels inform machines. Together, they build the kind of durable, compounding authority that makes your brand retrievable, citable, and visible across every system that matters.

Chapter Eight

Zero-Click Content and the Rise of Concept Chunking

Something has changed in the way people find answers, and most marketers haven't caught up yet.

A few years ago, when someone searched "best project management software for remote teams," they'd scan ten blue links, click three or four, skim the articles, compare, and then decide. That was the model. Write content, rank for keywords, get clicks, convert traffic.

That model is fragmenting.

Today, more and more queries are answered directly inside the search results themselves. Google's AI Overviews synthesize a response at the top of the page. Featured Snippets pull a paragraph from a source and display it in a box. Knowledge Panels present structured facts about entities. People Also Ask sections expand with pre-answered questions. And then there are the AI assistants: ChatGPT, Perplexity, Gemini, and others that pull from multiple sources, synthesize an answer, and present it to the user in one neat response.

The user never clicks through.

This is what the industry calls "zero-click search," and it's accelerating. Studies have shown that more than half of Google searches now end without a click to any external website[6]. For certain query types, particularly informational ones, that number is even higher.

For marketers and PR professionals, the instinct is to panic. If nobody clicks, what's the point of publishing?

But zero-click doesn't mean zero value. It means your content is being consumed in new ways, and your strategy needs to account for that.

The Keyword Trap

Traditional content marketing was built on a straightforward framework. Find a keyword with search volume. Write an article targeting that keyword. Optimize the page. Rank for it. Get clicks. Repeat.

That framework assumed the search engine was a matchmaker between keywords and documents. And for a long time, it worked well enough.

AI systems don't operate this way.

When someone types "best AI governance platform for banks" into Perplexity or asks ChatGPT the same question, the system isn't running a keyword match. It's performing a concept-level search. It understands that "AI governance" is a category. It knows "platform" implies software. It recognizes "banks" as a subset of "financial institutions." It pulls from sources that discuss these concepts with clarity and authority, regardless of whether those sources used the exact phrase "best AI governance platform for banks" anywhere in their text.

This is a fundamental shift.

AI systems chunk content by concept, not by keyword density. They evaluate whether a passage contains a coherent, self-contained idea that addresses the user's intent. A paragraph stuffed with keywords but lacking conceptual substance gets passed over. A paragraph that clearly explains a specific claim, ties it to a named entity, and provides measurable context gets retrieved.

I've watched this play out across thousands of releases distributed through PRWeb and later through NewsRamp. The releases that perform well in AI retrieval aren't the ones with the most keywords. They're the ones with the clearest concepts.

What Is Concept Chunking?

Concept chunking is the way AI systems break down and evaluate content. Instead of looking at an entire document as one unit, the system segments it into discrete, self-contained conceptual units, or "chunks." Each chunk is evaluated independently for relevance, clarity, and authority.

A chunk might be a paragraph. It might be a section. In some cases, it might be a single sentence if that sentence contains enough conceptual density to stand on its own.

Think about how you read a press release. You scan it. Certain paragraphs grab your attention because they contain a complete, meaningful idea. Others feel like filler, transition sentences that only make sense if you've read the paragraph before and after.

AI systems do the same thing, but more ruthlessly. They don't read your release start to finish. They extract chunks that match the user's query. And

each chunk needs to work as a standalone answer.

This is what I call the "Brandable Chunk" in action. Every paragraph in your press release should be a potential standalone answer to a question someone might ask. If it can't stand alone, it won't be retrieved alone. And retrieval is where the value lives.

Writing for Concept Chunking

So how do you actually write content that performs well in a concept-chunking environment?

Start with a simple rule: every paragraph should contain five elements. Who is the entity? What action or claim is being made? When did it happen or when is it relevant? Why does it matter? And ideally, what measurable detail supports the claim?

Here's what a weak chunk looks like:

"The company is excited to announce this important development, which builds on years of hard work and dedication from the entire team."

That sentence contains no entity name, no specific action, no time reference, no measurable detail, and no clear reason why anyone should care. If an AI system extracts that paragraph, it tells the user nothing. It's filler.

Now here's a strong chunk:

"In March 2026, Acme Corp launched version 5.0 of its AI governance platform, introducing automated audit trail generation that reduced compliance preparation time by 40% for mid-market financial institutions."

That single paragraph contains the entity (Acme Corp), the action (launched version 5.0), the time anchor (March 2026), the measurable detail (40% reduction in compliance prep time), and the context (mid-market financial institutions). Pull that paragraph out of any document and it still makes complete sense. That's a Brandable Chunk.

The difference between these two paragraphs isn't just quality of writing. It's retrievability. The strong chunk will be extracted and used. The weak chunk will be skipped.

Here's another common problem: connector paragraphs. These are the paragraphs that exist only to transition between sections. "Building on this momentum, the company also..." or "In addition to the above, we're pleased to share..." These read fine in context but mean nothing in isolation. AI systems will ignore them.

Front-load the most important information in each paragraph. Don't bury the lead in sentence three. The first sentence of every paragraph should contain your entity name and the core claim. Supporting details can follow, but the opening sentence should work on its own.

Here's a practical test you can run right now. Take your most recent press release. Pull out each paragraph individually. Read each one in isolation. Does it make sense on its own? Does it contain your brand name? Does it reference a specific time? Does it make a verifiable claim?

If most of your paragraphs fail this test, your release isn't ready for concept chunking. It might read well as a narrative document, but it won't perform well in retrieval.

Zero-Click Content Strategy for Press Releases

Here's the reality that many PR professionals haven't fully internalized: your press release may never get clicked. Not because it failed, but because its content is being consumed without a click.

Your carefully crafted announcement might appear in:

Google AI Overviews when someone searches your category.

ChatGPT responses when someone asks about your industry.

Perplexity citations when a researcher explores your topic.

Featured Snippets in traditional search results.

Knowledge Panel information associated with your brand.

This is actually good news. Your message is reaching people, informing decisions, and building brand association, all without requiring a single click to your website. The old traffic-centric model would call this a failure. The retrieval model recognizes it as influence.

But it means you must be disciplined about what's in every paragraph. You don't get to rely on the visitor reading your full page, scrolling past the boilerplate, and absorbing your complete narrative. You get one paragraph, maybe two, extracted and presented to a user who never visits your site.

If that paragraph contains your brand name, a specific claim, a time reference, and a measurable fact, you've delivered value even in a zero-click environment. If that paragraph is vague filler, you've wasted your one shot at influence.

The Zero-Click Readiness Test

I want to give you something practical you can use this week.

Take your last three press releases. Print them out or open them on your screen. For each one, go through every paragraph and score it on four criteria.

First: Could this paragraph be extracted and used as a standalone answer to a question? If someone asked "What does [your company] do in [your category]?" would this paragraph answer it?

Second: Does this paragraph contain your brand name? Not "the company" or "we" but your actual brand name. AI systems need named entities to build associations. Every paragraph that omits your brand name is a missed opportunity for entity reinforcement.

Third: Does this paragraph have a time reference? A month, a quarter, a year? Time anchoring tells retrieval systems that this information is current and contextually relevant.

Fourth: Does this paragraph make a specific, verifiable claim? Not "industry-leading" or "best-in-class" but a concrete statement that could be fact-checked?

Score each paragraph on a scale of zero to four based on how many criteria it meets. Then average across the release. If your average is below three, you have significant room to improve your zero-click readiness.

I've run this exercise with clients, and the results are often eye-opening. Releases that felt strong as narrative documents score poorly on chunk readiness. The fix isn't to abandon narrative entirely. It's to ensure that each narrative paragraph also functions as a retrievable unit.

The FAQ as Zero-Click Weapon

One of the most underutilized sections in a press release is the FAQ.

AI systems love Q&A format. The reason is structural: a question maps directly to a user query. When someone asks Perplexity "How does Acme Corp's AI governance platform handle audit trails?" and your FAQ section contains exactly that question with a concise, entity-rich answer, the match is nearly frictionless.

Here's how to structure FAQ entries for maximum retrieval:

Write the question in natural language, the way a real person would ask it. Not "What are the features of Product X?" but "How does Product X reduce compliance preparation time for financial institutions?"

Answer in one concise paragraph that includes your brand name, a specific fact, and a time reference. "Acme Corp's AI governance platform, launched in March 2026, automates audit trail generation and has reduced compliance preparation time by an average of 40% across its mid-market financial institution clients."

That's a complete, retrievable chunk disguised as an FAQ answer.

Three to five FAQ entries at the bottom of your press release create three to five additional retrieval opportunities. Each one is a discrete concept chunk that AI systems can extract and serve to users.

Pair your FAQ section with FAQ schema markup (FAQPage schema in JSON-LD). This signals to search engines that the content is structured

Q&A, which increases the likelihood of Featured Snippet inclusion and improves the clarity of your content for AI retrieval systems.

Don't treat the FAQ as an afterthought. Treat it as your highest-probability retrieval section.

From Keywords to Concepts: The Shift in Content Planning

If concept chunking is how AI evaluates content, then your planning process needs to shift from keywords to concepts.

Instead of starting with a keyword list, build a concept map.

Put your brand at the center as the core entity. Around it, map your topic clusters: your category, your differentiators, your outcomes, your customer segments, your geographic focus. These are the concepts you want retrieval systems to associate with your brand.

Each press release you publish should reinforce one or more of these concept clusters. Not by stuffing keywords, but by making clear, structured statements about your entity in relation to those concepts.

For example, if your concept map includes "AI governance," "financial institutions," "automated compliance," and "regulatory preparedness," then every release should contain at least one strong chunk that connects your brand to one of these clusters. Over twelve months, you've published twelve releases, each reinforcing the same conceptual territory from a slightly different angle.

This is how you build topical authority in a retrieval-driven ecosystem.

Not by chasing individual keywords, but by consistently reinforcing the conceptual associations between your entity and your domain.

Keyword targeting could get you one ranking for one query. Concept reinforcement builds the kind of broad topical authority that makes your brand eligible for retrieval across dozens of related queries, including ones you never explicitly targeted.

That's the shift. From optimizing for specific search terms to building a body of structured, concept-rich content that makes your brand the obvious answer across an entire topic cluster.

The marketers who understand this shift will build durable visibility. The ones who keep chasing keywords will wonder why their content isn't showing up in AI responses.

Chapter Nine

Managing Expectations in the Age of AI

If there's one place marketing and PR professionals get into trouble more than anywhere else, it's expectation setting.

I've been in this industry since 1997. The technology changes constantly. The terminology evolves every few years. The hype accelerates with each new platform. But executive psychology stays remarkably consistent.

They want visibility. They want measurable results. They want speed. They want certainty.

And increasingly, they ask a new question: "Will this get us into AI?"

That question, more than any other, reveals the gap between what executives expect and how these systems actually work. Closing that gap is your job, and doing it well is what separates trusted advisors from order-takers.

The Legacy Metrics Trap

Traditional press release reporting leaned heavily on impressive-sounding numbers. Impressions. Potential reach. Distribution volume. Media pick-

up counts.

These metrics felt good in a slide deck. A release that generated "2 million impressions" and "450 media pickups" looked like a success. But if you followed those numbers downstream, they rarely correlated with revenue. They measured distribution, not influence. They counted exposure, not engagement.

In a retrieval-driven world, that gap becomes more visible, because the metrics that matter are now measurable: retrievability, engagement, conversion influence, pipeline impact.

The executive who's been trained on impression counts for fifteen years will need to be retrained. That's your job too.

The Three-Release Lesson

I want to share something specific because it illustrates the point better than theory can.

Three press releases went out within 48 hours of each other. Same distribution channel. Similar category. Similar timing.

Release A generated over 9,500 reads.

Release B generated roughly 540.

Release C generated roughly 650.

The difference wasn't distribution. All three had the same distribution infrastructure. The difference was framing.

Release A connected recognizable entities to specific geography and

high-interest themes. It named names. It cited places. It tied the announcement to broader context that readers cared about. The headline and opening paragraph immediately answered the question: "Why should I care about this?"

Releases B and C described the process of their own announcements rather than the substance. They were internally focused, written for the people making the announcement rather than for the people reading it.

Clear framing produces stronger semantic signals. Stronger signals produce higher engagement. Higher engagement increases retrieval probability. This isn't abstract theory. It's the difference between 9,500 reads and 540.

Answering the AI Question

Eventually, someone in the C-suite will ask: "Did this get us into ChatGPT?"

The responsible answer isn't yes. And it isn't no. It's this:

"We can't force inclusion. We can increase the likelihood of retrieval."

Then explain, calmly and clearly:

AI models are trained on large datasets, and then those training parameters are fixed. The model doesn't continuously absorb the live web. When a user asks a question, modern AI systems retrieve indexed content from search and other sources, then summarize what they find.

Retrieval depends on structure, authority, reinforcement, and clarity. Your

job is to improve eligibility across all four of those dimensions, consistently, over time.

That language is precise. It's defensible. And it builds the kind of trust that keeps budgets funded quarter after quarter. Over promising erodes that trust faster than anything else.

The Personalization Reality

Here's one of the most important expectation resets you can offer your stakeholders.

Search results are not universal. They're contextual.

Two people typing the exact same query will often see different results based on geographic location, device type, search history, click behavior, logged-in status, industry browsing patterns, and language settings. There is no single "official" Search Engine Results Page. There are contextual variations.

AI outputs work similarly. Depending on the system, responses may shift based on time sensitivity, geographic signals, domain authority, topic clustering, user history, and other contextual factors. Two users asking "top AI governance platforms" may see slightly different answers.

This matters because of what I call the Screenshot Problem.

A client searches their brand name, doesn't see the release, takes a screenshot, and sends it with a note: "It's not showing up."

What they're showing you is one location, one device, one moment in time, one user context. It isn't the system. It's a slice of the system.

The correct response is factual and calm: "Search results are personalized. What you see may differ from what others see based on location and search behavior. Our goal is to strengthen the overall retrieval footprint, not to optimize for one specific view."

Then pivot to data.

What to Measure Instead

Rather than chasing screenshots, measure trends.

Search impressions over time tell you whether your content is appearing in more queries. Query diversity shows whether you're being found for a wider range of relevant terms. Click-through rates indicate whether your headlines and descriptions compel action. Engagement duration reveals whether people who find you actually spend time with your content. Conversion influence tracks whether your releases contribute to downstream actions. Assisted pipeline impact connects visibility to revenue conversations.

These metrics reflect structural improvement. One visible ranking at one moment in time doesn't define your visibility. Trend lines do.

Sales Enablement: Turning Visibility Into Revenue

Press releases are chronically underutilized internally. I've seen organizations invest weeks in crafting the perfect announcement, publish it, and then never mention it to their own sales team.

After every release, provide your sales team with three talking points they can use in prospect conversations. Include a one-paragraph summary they

can drop into an email. Offer suggested language they can customize. Highlight the measurable outcomes from the announcement.

A press release that never enters a revenue conversation is an unused asset. Visibility must translate into influence. Influence must support pipeline. Pipeline must justify the budget for the next release.

Post-Publish Amplification

The moment your release goes live, a clock starts. The first 48 hours are your highest-leverage window for amplification.

Notify internal teams with a three-bullet summary and a direct link. Inform customers if the announcement affects them. Share with partners and provide pre-written social copy they can use. Repurpose the core message into short-form content for LinkedIn and other channels. Update your FAQ pages if the announcement introduces new capabilities. Reinforce on social with consistent messaging.

This isn't optional. It's the difference between a release that generates one signal and a release that generates twenty signals, all pointing in the same direction.

A press release is the foundation. Amplification is the multiplier.

Resetting the Executive Conversation

The next time you're in a post-campaign review, try shifting the questions.

Instead of "How many outlets picked it up?" ask: How did this reinforce our category association? How did it strengthen entity authority? How

did it improve our retrieval eligibility? How did it influence engagement? How did it support revenue conversations?

Those are durable questions. They move the conversation from vanity metrics to structural impact. They train executives to evaluate press releases as strategic assets rather than publicity events.

And over time, they create a shared language between marketing, PR, and leadership that makes it easier to justify investment, measure progress, and build the kind of consistent cadence that compounds into real competitive advantage.

The marketers who can manage these expectations calmly and precisely, without over promising or underselling, will be the ones who keep their budgets, their clients, and their credibility intact.

Chapter Ten

After You Hit Publish: The Real Work Begins

Publishing a press release is not the finish line.

It's the starting gun.

I've watched this pattern repeat for nearly three decades. A team spends days or weeks crafting the announcement. They agonize over headlines, wordsmith quotes, finalize approvals. The release goes live. Everyone exhales. The file gets archived. And then they move on to the next thing.

In a retrieval-driven ecosystem, that approach leaves enormous leverage on the table.

The moment your release publishes, three parallel processes begin: search indexing, internal awareness, and external amplification. If you ignore those processes, you reduce the impact of everything you just invested. If you manage them intentionally, you multiply it.

The Indexing Window

After publication, search engines begin crawling. This can happen within

minutes, but indexing isn't uniform. Several factors influence how quickly and thoroughly your release gets indexed.

Domain authority matters. If your release lives on a site that search engines already crawl frequently, indexing will be faster. Crawl frequency plays a role too; sites with fresh content published regularly get crawled more often. Internal linking helps because links from other pages on your site signal to crawlers that the new page exists and is connected to your broader content ecosystem. Sitemap inclusion ensures search engines know about the page. Structured data reduces classification friction, helping the system understand what the page contains. Canonical clarity prevents authority from fragmenting across duplicate versions.

AI retrieval systems are layered on top of search engines. They depend on indexing. If your release isn't indexed, it isn't retrievable. Period.

That's why hosting on your own domain matters. That's why submitting updated sitemaps matters. That's why internal linking matters. Distribution alone doesn't guarantee indexing. These structural steps do.

Immediate Internal Mobilization

Your first amplification channel is inside your own organization.

Within hours of publication, your sales team should know about the release. Customer success should know. Leadership should know. The broader marketing team should know. And they shouldn't have to read the entire release to understand it.

Provide a three-bullet summary of what was announced and why it matters. Include a short paragraph explaining relevance to current sales con-

versations. Share a direct link they can forward. Provide suggested talking points for prospect calls. Draft email copy they can customize and send.

Make it frictionless. If using the release requires effort, people won't use it. If you hand them ready-to-deploy materials, they will.

I've seen a single well-timed sales email referencing a fresh press release open conversations that had been stalled for weeks. The release gave the salesperson a reason to reach out, a piece of news to share, and a credibility boost that a cold follow-up can't match.

Customer Communication

A press release shouldn't live only in your newsroom.

Ask yourself: why does this announcement matter to our existing customers? Then communicate that clearly and directly.

If you've launched a new feature, tell customers how it affects their workflow. If you've formed a partnership, explain what it means for their experience. If you've published performance data, show how it validates the results they're already seeing.

The subject line shouldn't be "Company X Announces Y." It should be "What This Means for Your Compliance Workflow" or "How Our New Integration Saves You Two Hours a Week."

Tie every announcement directly to customer outcomes. Retrievability and revenue are connected through relevance.

Partner Amplification

Partners extend your contextual neighborhood in ways that directly strengthen retrieval signals.

After publication, notify your integration partners, channel partners, and strategic allies. Provide pre-written social copy they can share with minimal effort. Offer quote excerpts they can reuse in their own communications. Encourage cross-linking when appropriate.

Each external mention reinforces co-occurrence between your entity and relevant concepts. Co-occurrence strengthens entity mapping in retrieval systems. It's not a guaranteed path to dominance, but it's compounding advantage that builds with every mention.

Repurposing Without Dilution

Your press release is a structured foundation. From that foundation, you can build multiple formats.

A LinkedIn post that highlights the key announcement. A short executive video summarizing the news. A 60-second audio clip for podcast channels or internal distribution. An FAQ page that addresses the questions your announcement raises. A sales one-pager that translates the news into prospect-facing language. A slide for your pitch deck. A knowledge base update for your support team. A webinar topic that explores the announcement in depth.

The critical principle here is consistency of signal. Every format should reinforce the same core message, the same entity-topic association, the same claims, the same framing. Don't rewrite the narrative differently for each platform. Adapt the format, but preserve the substance.

Consistency strengthens semantic clustering across platforms. Inconsistency introduces ambiguity. And ambiguity is the enemy of retrieval.

Avoiding "Set It and Forget It"

One of the most common post-publish mistakes is the simplest one: publish, distribute, move on.

In a compounding system, abandonment slows growth.

Within 7 to 14 days of publication, review your performance metrics. How did the headline perform in terms of click-through rate? What was the engagement duration for people who did click through? Which search queries are driving traffic to the release? Are there FAQ pages or internal content that should be updated based on the announcement?

Within 28 days, publish again.

The rhythm matters more than perfection. A consistent cadence of well-structured releases will outperform an occasional masterpiece followed by months of silence.

The Retrieval Feedback Loop

Over time, your post-publish reviews will reveal patterns.

You'll notice which topics produce stronger engagement. You'll see which headline structures drive higher click-through rates. You'll discover which measurable claims resonate with your audience. You'll observe which co-occurring entities, partners, categories, geographies, strengthen your visibility when mentioned alongside your brand.

Use this data to refine your monthly cadence. This isn't guesswork. It's

iterative reinforcement. Each cycle teaches you something about what retrieval systems and human audiences respond to, and each subsequent release can incorporate those lessons.

The Post-Publish Checklist

I want to leave you with something you can tape to your monitor.

Immediately after publication:

Confirm the canonical URL is correct and points to your preferred version of the release.

Validate your structured data using Google's Rich Results Test or Schem a.org validator.

Submit an updated sitemap if your CMS doesn't do this automatically.

Notify internal teams with a summary, a link, and talking points.

Email customers if the announcement is relevant to their experience.

Inform partners and provide ready-to-use social copy.

Share structured social posts across your brand channels.

Log baseline metrics: impressions, clicks, engagement, referring queries.

Within 30 days:

Review performance against your baseline.

Identify which queries and topics drove the most engagement.

Reinforce topic consistency in your planning for the next release.

Publish the next release.

Simple. Repeatable. And far more powerful than it looks on paper.

The difference between organizations that get compounding value from their press releases and those that don't isn't talent or budget. It's follow-through.

Chapter Eleven

The Long Fat Tail: Why Press Releases Are Assets, Not Events

Most marketing is built around spikes.

Campaign launch. Ad flight. Conference push. Product debut. The expectation is an immediate lift. When the lift fades, the campaign ends. Budget shifts to the next spike.

Press releases have been evaluated through that same lens for decades. Publish, measure the first-week performance, file the report, move on. If the numbers spiked, success. If they didn't, disappointment.

That evaluation framework is a strategic mistake.

A properly constructed press release is not an event. It's an asset. And assets behave differently than campaigns.

The Long Fat Tail

You may be familiar with the "Long Tail" concept in digital economics: the idea that many small, consistent interactions over time can rival or exceed

one large spike. Chris Anderson wrote about it. It reshaped how people thought about retail, media, and content. [7]

The press release operates in what I call the Long Fat Tail.

It's not a thin tail that tapers into insignificance after the initial burst. It's a fat tail: persistent, durable, continuously accessible.

When I launched PRWeb in 1997, we made a deliberate decision early on. We never removed content. Every release remained archived and searchable. At the time, some people questioned that approach. Why keep old press releases around? The news cycle has moved on.

But we understood something that proved to be true year after year: indexed content compounds.

A press release doesn't stop existing when the news cycle moves on. It remains searchable. Linkable. Citable. Retrievable. Discoverable. And now, in the retrieval era, it remains a candidate source for AI systems answering questions about your brand, your category, and your industry.

Persistence is power. It was true in 1997 and it's even more true today.

The Cost Curve

All your costs, drafting, review, distribution, coordination, are incurred at the moment of publication. After that, every additional engagement costs nothing. Over time, cost per engagement declines continuously.

Compare this to paid media. Turn off the ad spend and the visibility vanishes instantly. You're renting attention. An indexed press release continues generating exposure indefinitely. You're building an asset.

The Archive Strategy

Don't delete old releases. I've seen companies delete three years of press release history in a website redesign and wonder why their search visibility dropped. The content they removed was doing quiet, persistent work in the background.

Older releases demonstrate continuity, reinforce messaging consistency, provide historical proof points, and strengthen entity salience. Every release you delete removes a reinforcement signal from the index. Permanence compounds. Removal erodes.

The Accumulation Effect

Let's model a simple scenario.

You publish one press release every 28 days. After one year, that's 12 releases. After two years, 24. After three years, 36.

At the end of three years, you don't have 36 isolated announcements sitting in a newsroom archive. You have 36 time anchors, each one placing your brand in a specific moment in your category's evolution. You have 36 reinforcement pulses, each one strengthening the association between your entity and your topic. You have hundreds of structured chunks available for retrieval. You have dozens of co-occurrence relationships connecting your brand to partners, customers, geographies, and industry themes. You have multi-year continuity that no competitor can replicate in a quarter.

A competitor who publishes sporadically, three or four releases a year at best, lacks that reinforcement density. They may have a bigger brand, a

larger budget, and a more recognizable name. But in a retrieval system that rewards consistent, structured, time-anchored content, your 36 releases create a structural advantage that their 9 cannot match.

Accumulation favors consistency.

The Psychological Advantage

There's a human factor here that's just as important as the technical one.

When a prospect researches your company and finds a newsroom with continuous monthly activity, clear progression from one announcement to the next, structured milestones, and reinforced positioning over time, they perceive momentum.

Momentum builds trust. Trust accelerates decisions.

I've spoken with sales teams who report that prospects reference the company's press release history during calls. "I saw you launched that compliance feature last year and then expanded it this quarter. That tells me you're investing in this area." That's not retrieval. That's a human reading a pattern of activity and drawing a conclusion about reliability.

Machines may retrieve your release. Humans still decide. The Long Fat Tail supports both.

Compounding Visibility vs. Episodic Noise

Episodic marketing creates bursts of attention. Compounding marketing creates trajectory.

When you combine the 28-day cadence with entity reinforcement, structured data clarity, canonical consolidation, and persistent hosting, you

create a visibility trajectory. That trajectory may look modest month to month. The first release doesn't change everything. Neither does the second or the third.

But over years, the compounding effect becomes structural advantage. It becomes the kind of advantage that competitors can't replicate with a single big campaign or a burst of ad spend, because it was built through discipline over time.

Risk Reduction

Persistence isn't just an SEO strategy. It's a business protection strategy. If a competitor claims you've never released anything in their category, your archive disproves it with timestamps. If regulatory scrutiny arises, your time-stamped announcements establish a clear, public sequence. If a patent dispute involves when you first described a capability, your archived release serves as dated, public prior art.

The Long Fat Tail is Not Passive

The Long Fat Tail doesn't happen automatically. It requires ongoing reinforcement, with each release building on what came before. It requires structured clarity with named entities, time anchors, and measurable claims. It requires topic discipline, because scattered topics fragment authority instead of compounding it. And it requires canonical integrity so authority consolidates rather than fragments.

Without these elements, the tail thins. With them, it stays fat.

From Spike to System

Stop evaluating each press release by how it performed this week. Start evaluating how it strengthens your three-year visibility trajectory.

Systems, unlike spikes, don't fade when the news cycle moves on. They compound. They persist. And over time, they become the kind of durable competitive advantage that no single campaign can replicate.

That's the Long Fat Tail. It worked when I launched PRWeb in 1997. It works in 2026. And it will keep working as long as retrieval systems reward clarity, consistency, and time-anchored authority.

Chapter Twelve

Competitive Dynamics in the Retrieval Era

One of the most dangerous assumptions in marketing is also one of the most common. It sounds reasonable. It sounds measured. And it is dead wrong.

"If we stop publishing, nothing changes."

I've heard this from CMOs, from agency leads, from PR directors who have a hundred things on their plate and figure the press release can wait. They assume that if they go quiet for a few months, their position holds steady. That visibility is like a savings account: what you deposited is still there, sitting untouched, earning a little interest.

That's not how retrieval works.

Authority in retrieval systems isn't absolute. It's relative. Your brand's visibility isn't measured in a vacuum. It's measured against every competitor reinforcing the same topics, targeting the same category, publishing in the same space. If you go silent and they don't, you haven't held your position. You've lost ground.

The Authority Gradient

Think of authority not as a switch you flip on or off, but as a gradient. A spectrum. You're never simply "visible" or "invisible." You're more likely to be retrieved, or less likely to be retrieved.

Every release you publish nudges your position along that gradient. Every month you stay quiet holds your position still. And while you're standing still, your competitors may be moving.

This is subtle. You won't notice it in a single quarter. But gradients diverge over time, and when they do, the gap is structural. It doesn't close on its own.

I've watched companies realize this too late. They go quiet for six months, then publish a big announcement and wonder why it doesn't land the way it used to. The answer is simple: while they were silent, someone else was reinforcing. The retrieval systems noticed. Not because there's a penalty for silence, but because there's a reward for consistency.

Two Companies, Two Strategies

Let me model this simply.

Company A publishes twice a year. They focus on big launches, big events, big moments. Between those moments, they go silent. Their team works heads-down on product, saves up their news, then goes loud.

Company B publishes one release every 28 days. They align consistently to a defined topic domain. They reinforce structured authority monthly. Not flashy. Not dramatic. Steady.

After one year, Company A has two signal spikes. Company B has twelve reinforcement pulses.

After three years, Company A has six spikes. Company B has thirty-six pulses.

Now ask yourself: when a retrieval system evaluates which brand has the strongest, most consistent association with a category, which company looks more authoritative? The one with six bursts, or the one with thirty-six steady signals?

Retrieval systems reward reinforced proximity. They favor brands that show up repeatedly in structured, clear, category-aligned content. Not noise. Not volume for its own sake. Cadence. Purposeful, measured repetition.

Company A might have the better product. They might have a bigger marketing budget. But Company B has the structural advantage. And structural advantages compound.

The Plateau Problem

When companies publish inconsistently, their authority plateaus. They maintain some search presence, some residual recognition. The lights don't go completely dark, which is precisely what makes it dangerous. Leadership looks at the dashboard and thinks things are fine.

They're not. Meanwhile, competitors who publish consistently are reinforcing associations and increasing salience. The gap widens quietly. I've seen companies that were once category leaders wake up to find a smaller competitor showing up in every AI summary. Not because the competitor

had a breakthrough. Because the competitor showed up. Every month. For two years.

AI retrieval systems don't punish silence explicitly. But they prioritize fresher reinforcement when evaluating candidates. Silence is not neutral. It's a strategic risk that masquerades as patience.

Topic Ownership Through Repetition

In retrieval systems, topic ownership isn't something you declare. It's something you accumulate.

If your brand appears monthly in relation to AI governance, regulatory compliance, financial institutions, audit efficiency, and policy frameworks, you begin clustering with that domain. The systems start associating your entity with those topics. Not because you said "we're the leader." Because you showed up, consistently, with structured content that reinforced those connections.

If a competitor appears sporadically in the same space, their clustering remains weaker. They might have a great product. They might even have better technology. But in the retrieval system's model of the world, your brand owns that topic more convincingly.

This is how category leadership shifts in the retrieval era. Not through a single viral moment, but through accumulated presence. Monthly reinforcement. Structured clarity. Repetition with purpose.

Ownership isn't claimed. It's built, one release at a time.

Evaluating Competitor Cadence

Here's a practical exercise. Pull up the last twelve months of press releases from your top three competitors. Ask: How often did they publish? What topics did they reinforce? Were their releases structured with measurable claims? If your competitors are publishing monthly while you publish quarterly, your authority gradient is shifting in their favor.

Retrieval as Sometimes Zero-Sum

Not all visibility is zero-sum. Many queries pull from multiple sources. But in competitive category queries, the math gets tighter. When a user asks "What are the top AI governance platforms for financial institutions?", there are limited candidate slots. If you're not in that set, a competitor is.

You can't control which queries users ask. But you can control whether you've done the work to be eligible. Eligibility isn't guaranteed to produce selection. But ineligibility guarantees you won't be selected.

The Compounding Edge

Small monthly advantages compound. A slight increase in semantic density. Better canonical consolidation. One more reinforcement pulse per quarter. Multiplied over months and years, they create a durable advantage that is extraordinarily difficult to close.

A competitor who wants to catch up after three years of your consistent publishing doesn't just need to match your current output. They need to overcome thirty-six months of accumulated structural advantage. That's not impossible, but it's expensive, slow, and most organizations don't have the patience for it.

You don't need to out-hype your competitors. You need to out-reinforce them. Publish clearly. Publish consistently. Structure precisely. Persist indefinitely. Over time, reinforced clarity wins. The most reliable signal is the one that shows up, structured and clear, every single month.

Chapter Thirteen

Engagement: What Makes a Press Release Pop

Everything we've covered so far has been about building the foundation. Structure. Retrievability. Entity authority. Canonical clarity. Topic clustering. The 28-day cadence. All of it designed to make sure your press release can be found.

But being found isn't the finish line. Engagement is.

Retrievability creates opportunity. Engagement creates momentum.

Think of it this way: retrievability is getting invited to the meeting. Engagement is what happens when you walk in the room. Do people lean forward? Do they remember what you said? Do they tell someone else about it afterward?

A press release that is technically perfect but emotionally flat may surface in search results or AI summaries. It'll sit there, eligible and inert. Nobody shares it. Nobody references it. Nobody acts on it. It's the equivalent of being at a party and standing in the corner saying nothing.

Engagement is the multiplier. It's what sparks conversation, drives sharing, strengthens behavioral signals, and moves prospects into your pipeline.

Without it, even the most structurally optimized release stays static.

What follows are seven engagement strategies I've tested, refined, and watched work over years of publishing. These aren't theoretical frameworks. They're practical approaches that make the difference between a release that surfaces and a release that resonates.

1. News Drafting: Anchor Your Release to Current Events

News drafting means anchoring your announcement to a current event, regulatory change, or industry trend. You're connecting your release to something people already care about. Compare:

Weak: "Acme Corp Launches New AI Compliance Tool."

Strong: "As New Federal AI Regulations Take Effect, Acme Corp Launches Compliance Platform for Mid-Market Financial Institutions."

The second headline connects your announcement to something larger. It says: this matters now.

Build a running list of upcoming regulatory changes, industry trends, and market shifts. When it's time to publish, find the most legitimate connection between your announcement and what's happening in the world. The key word is legitimate. Anchor where you have authority. Don't attach your brand to every trending topic.

A press release tied to an active narrative rides momentum. Stories travel further than announcements.

2. Light Splashing: Shine a Spotlight on Your Customers

Instead of making your press release about you, make it about your customer.

Weak: "Acme Corp Signs Major New Client."

Strong: "Regional Bank Reduces Audit Preparation Time by 32% Using Acme Corp's AI Governance Platform."

The customer is the hero. The outcome is measurable. And Acme Corp still benefits. When you elevate a customer publicly, they share it. They forward it internally. Their marketing team posts it. They become, without you asking, an evangelist for your brand.

Make the customer the headline subject. Include a quote from their executive. Use measurable outcomes. Highlight their strategic goals, not your product features. Light splashing turns a single press release into a shared asset. You publish once. Your customer amplifies it ten times.

3. Measurable Specificity: Numbers Create Gravity

Vague claims drift. Specific claims anchor. The single most common weakness in press releases is vagueness. "We improved efficiency" means nothing. "We reduced regulatory audit preparation time by 32%" is something a CFO reads and forwards to their compliance team.

Numbers increase credibility, improve semantic density, strengthen retrieval vectors, and capture attention. Even small statistics matter. "Serving 140 clients across 12 states" is more convincing than "serving clients nationwide."

When drafting, challenge every adjective. If you wrote "significant growth," replace it with the actual number. If you can't quantify it, ask

whether the claim is worth including at all.

4. Recognizable Entities: Borrow Context Through Association

When your release references a regulatory body, a known standard, or a major industry event, it gains contextual gravity.

Weak: "Acme Corp Expands AI Governance Capabilities."

Strong: "As SEC Finalizes New AI Risk Disclosure Rules, Acme Corp Expands Governance Platform for Financial Institutions."

The SEC is a recognized entity. "AI risk disclosure" is a defined topic. You've thickened the edges in the Knowledge Graph without exaggerating, just by placing your announcement inside a recognized context.

The rule of responsible association: borrowing context only works if your announcement genuinely relates to the entity. Never force it. Recognizable entities amplify relevance. Artificial associations amplify skepticism.

5. Clear Stakes: Answer "Why Should Anyone Care?"

Weak: "Acme Corp Releases Version 4.2 of AI Governance Platform." No consequence. No urgency.

Strong: "Acme Corp Releases AI Governance Update to Help Banks Avoid Costly Audit Delays." Now the stakes are visible. Audit delays cost money. Compliance penalties are real.

Every release should make the stakes explicit within the first few sentences.

Not manufactured. Not inflated. Just clear. When stakes are defined, engagement rises. When they're vague, attention drifts.

6. Forward Momentum: Signal What Comes Next

Releases that only describe the present feel static. Include roadmap hints, not detailed specs, just directional signals. "The company plans to expand the platform to address ESG reporting requirements in Q3." Future milestones work too: "Following this deployment, the company expects to reach 500 enterprise clients by year-end."

Forward momentum makes your release feel like a chapter in a story, not a standalone memo. Momentum signals growth. Growth attracts attention. Attention reinforces visibility.

7. Emotional Neutrality with Subtle Energy

Press releases should not be theatrical. Exclamation points, breathless superlatives, and "revolutionary breakthrough" language don't create engagement. They create skepticism.

But press releases also shouldn't be lifeless. There's a difference between professional tone and sterile tone. A release that reads like a legal filing won't get shared any more than one that reads like an infomercial.

The sweet spot is emotional neutrality with subtle energy. The energy doesn't come from adjectives. It comes from strong verbs, clear outcomes, defined impact, and human quotes.

"Launched" is better than "is pleased to announce the launch of." "Reduced audit time by 32%" is better than "significantly improved audit efficiency."

A quote from your CEO that says "Our clients told us they needed faster compliance workflows, so we built them" is better than "We are thrilled to deliver this innovative solution to the market."

Let the facts carry the energy. Let the outcomes speak. Let the quotes sound like something a real person actually said. When the content is genuinely strong, you don't need to oversell it.

The Formula

When you combine all seven strategies, you get a release that pops. Not because it's flashy, but because every element is working.

Timeliness from news drafting. Recognition from light splashing. Specificity from measurable claims. Context from recognizable entities. Clarity from defined stakes. Momentum from forward signals. Structure from everything we've covered in the earlier chapters.

Most releases include one or two of these. The strongest include all of them.

The Multiplier Effect

Here's where it all connects.

Retrievability gets you into the room. Engagement gets people talking. Talking extends reach. Extended reach reinforces authority. Authority increases retrieval probability.

The loop feeds itself.

A technically optimized but dull release may be retrieved. An engaging but

poorly structured release may be read but never reinforced by the systems that matter. The goal is integration. Structure and engagement working together.

When a release is both retrievable and engaging, it creates a virtuous cycle. The AI system surfaces it. A human reads it and shares it. The share generates a backlink. The backlink strengthens authority. The strengthened authority increases the probability of retrieval next time. And the next release starts from a slightly higher baseline.

That's the multiplier effect. Each release doesn't just stand alone. It lifts the next one. And the one after that. And the one after that.

You don't build this with a single brilliant release. You build it with consistent, engaging, structured releases published every 28 days. The multiplier compounds. The advantage grows. And the gap between you and competitors who publish dull, sporadic, unstructured releases widens with every cycle.

Chapter Fourteen

Retargeting: Turning Readers Into Relationships

There's a moment in the lifecycle of every press release that most marketers completely miss.

Someone reads your announcement. They spend thirty seconds, maybe a minute, absorbing the news. They register your brand name, your product, your claim. And then they close the tab. They go back to their inbox, their meeting, their lunch. They don't click through to your website. They don't fill out a form. They don't request a demo. They just leave.

In the traditional PR model, that person is gone. You had their attention for a moment, and now it's over. You have no way to reach them again. No way to continue the conversation. No way to move them from awareness to consideration to action.

That's an enormous amount of wasted attention. And until recently, there was nothing you could do about it.

Retargeting changes that.

What Retargeting Actually Is

Retargeting is the practice of showing targeted ads to people who have already interacted with your content. If someone reads your press release, a small piece of code called a tracking pixel records that interaction. Later, when that person is browsing other sites, scrolling through social media, or searching on Google, your ad appears in front of them. Not randomly. Specifically because they are already engaged with your content.

This isn't cold outreach. This is a warm follow-up. The person has already demonstrated interest by reading your announcement. Retargeting lets you stay in front of them after that initial moment of attention has passed.

Think about your own experience online. You visit a product page, and for the next week you see ads for that product on every site you visit. That's retargeting in action. The same mechanism now works for press releases.

The difference is significant. A press release without retargeting is a one-touch interaction. The reader sees it, processes it, and moves on. A press release with retargeting becomes a multi-touch journey. The reader sees your announcement, leaves, and then encounters your brand again on LinkedIn, in their Google search results, or in their Facebook feed. Each subsequent touchpoint reinforces the message. Each one moves the reader closer to action.

Most B2B purchases involve seven to thirteen touchpoints before a decision. A single press release, no matter how well-written, is one touchpoint. Retargeting turns that one into five or six without requiring the reader to do anything except what they were already going to do: browse the internet.

How It Works With Press Releases

Here's the practical setup, and it's simpler than most people expect.

First, you choose an advertising platform. The major options are Google Ads, Meta (Facebook and Instagram), and LinkedIn. Each platform has its own strengths. Google Ads reaches people across the broadest network of websites and search results. Meta reaches people on Facebook and Instagram, which is useful for consumer-facing brands and B2C announcements. LinkedIn is the strongest choice for B2B companies, because you're retargeting professionals in a professional context. Your choice depends on where your audience spends their time.

Second, you create a tracking pixel on your chosen platform. Each ad platform provides a small snippet of code, the pixel, that you generate from within your advertising account. Google calls it a remarketing tag. Meta calls it the Meta Pixel. LinkedIn calls it the LinkedIn Insight Tag. The setup process takes minutes and is well-documented by each platform.

Third, you add that pixel to your press release. At Newsworthy.ai, we've built this directly into the publishing workflow. When you're preparing a release, you paste in your tracking pixel code. When someone reads your press release on the distribution network, the pixel fires. That reader is now in your retargeting audience.

Fourth, you create your retargeting campaign. Back in your ad platform, you build a campaign that targets the audience your pixel has been collecting. You set your budget, your ad creative, your messaging, and your duration. The platform handles the rest, showing your ads specifically to people who read your press release.

That's it. Four steps. The technical barrier is low. The strategic value is high.

Why This Matters For News Marketing

Let me connect retargeting back to the News Marketing framework, because this isn't just a nice add-on. It addresses one of the oldest problems in PR.

Press releases have always been good at generating awareness. Someone reads your announcement and now they know your company exists, or they know about your new product, or they know about your partnership. But awareness without follow-through is a leaky bucket. People forget. Attention is fleeting. The news cycle moves on.

Retargeting plugs the leak.

The person who read about your AI governance platform launch on Monday sees your ad on LinkedIn on Wednesday. The ad doesn't need to re-explain everything. It can be simple: "See how mid-market banks are reducing audit prep time by 32%." The reader already has context from your press release. The ad reinforces it and provides the next step, a link to a case study, a demo request page, and a whitepaper download.

This is where the real value of press release engagement converts into pipeline. You're not losing your connection with someone who read your announcement but didn't take an additional step. You're extending the relationship.

Consider the math. If your press release gets 5,000 reads and you have no retargeting, you reached 5,000 people once. Some will remember you. Most won't. If that same release has retargeting enabled, those 5,000 readers now see your brand three, four, five more times over the following

weeks. The effective reach isn't 5,000 impressions. It's 15,000 to 25,000, all targeted at people who already showed interest.

And the cost is remarkably efficient. Retargeting audiences are small and highly qualified, which means your cost per impression and cost per click are typically much lower than broad prospecting campaigns. You're spending less to reach the people most likely to care.

Platform Selection Guide

Choosing the right ad platform for retargeting depends on your audience and your goals.

Google Ads is the broadest option. Your retargeting ads appear across the Google Display Network, which includes millions of websites, as well as in YouTube pre-roll and in Google search results. If your goal is maximum follow-up visibility across the widest range of online activity, Google is the default choice. It works well for both B2B and B2C.

LinkedIn is the strongest B2B option. When a decision-maker reads your press release about enterprise software and then sees your ad in their LinkedIn feed the next day, the context is perfectly aligned. They're already in a professional mindset. LinkedIn also offers targeting refinements that other platforms don't: you can layer retargeting with job title, company size, industry, and seniority filters. This means you're not just retargeting everyone who read your release. You're retargeting the VP of Compliance at a mid-market bank who read your release. That's precision.

Meta works best for consumer-facing announcements or brands with a strong social presence. If your press release is about a consumer product launch, a community initiative, or an event, Meta's retargeting across

Facebook and Instagram puts your follow-up message where people are spending their personal browsing time.

You're not locked into one platform. Some organizations run retargeting across two or all three simultaneously, using each platform's strengths for different audience segments. A B2B company might run LinkedIn retargeting for decision-makers and Google retargeting for broader industry awareness, all from the same press release pixel.

Retargeting As Engagement Multiplier

In the previous chapter, I described the multiplier effect: retrievability gets you into the room, engagement gets people talking, and the loop feeds itself. Retargeting adds a powerful new link to that chain.

Without retargeting, the engagement cycle depends entirely on the reader taking action in the moment. They read, they share, they click, or they don't. There's no second chance.

With retargeting, you get a second chance. And a third. And a fourth.

Each subsequent ad impression is an opportunity for the reader to take the action they didn't take the first time. Maybe they weren't ready on Monday. But by Thursday, after seeing your brand twice more, they click through to your case study. By the following week, they request a demo. The press release started the conversation. Retargeting continues it.

This also strengthens the behavioral signals that feed back into your visibility. When retargeted users click through to your website, that's additional engagement. When they spend time on your site, that's a dwell time signal. When they share your content from your site, that's a social signal. All

of these behavioral indicators contribute to the broader authority and engagement metrics that search engines and AI systems factor into their evaluations.

Retargeting doesn't just recover lost attention. It amplifies the engagement your press release has already generated.

Practical Considerations

A few things to keep in mind as you implement retargeting on your press releases.

Your budget doesn't need to be large. Because retargeting audiences are small and targeted, you can run effective campaigns on modest budgets. Even a few hundred dollars per month can sustain meaningful follow-up visibility across your press release readership. Start small, measure results, and scale what works.

Creative should be simple and reinforcing. Your retargeting ads don't need to tell the whole story. The press release already did that. Your ad should remind, reinforce, and redirect. A clear headline, a supporting data point from the release, and a specific call to action. "See the case study" or "Request a demo" or "Download the report." One message. One action.

Duration matters. Most retargeting campaigns run for 14 to 30 days after the initial interaction. That window aligns well with the 28-day publishing cadence. Your retargeting from one release winds down just as your next release goes live, creating a continuous cycle of new attention and sustained follow-up.

Privacy compliance is non-negotiable. Retargeting relies on cookies and

tracking technology, which means you must comply with applicable privacy regulations: GDPR in Europe, CCPA in California, and other regional frameworks. Ensure your press release hosting includes proper cookie consent notices and that your ad platforms are configured for compliant data handling. This is table stakes, not optional.

Measure the full funnel. Don't just track ad impressions and clicks. Track the downstream actions: website visits from retargeted users, time on site, pages viewed, form submissions, demo requests, and pipeline influence. The goal isn't vanity metrics. The goal is connecting press release readership to revenue through the retargeting bridge.

From One Touch To Many

The traditional press release model treated each release as a discrete event. You published, you measured the initial response, and you moved on. If the reader didn't act in that first moment of attention, the opportunity was lost.

Retargeting transforms the press release from a single touchpoint into the first step in a sustained conversation. You're not hoping the reader remembers you. You're making sure they see you again. And again. In the right context. With the right message. At the right time.

When you combine retargeting with the News Marketing framework, the 28-day cadence, entity reinforcement, structured data, zero-click readiness, and engagement strategies, you create something that most PR professionals have never had: a closed loop between awareness and action.

Your press release builds awareness. Your retargeting sustains it. Your next press release reinforces it. The cycle compounds.

That's not theory. That's a system. And it's one more reason why the press release, properly executed, is the most underleveraged marketing asset most companies already have.

Chapter Fifteen

The Visibility Advantage: Building Your System

AI didn't eliminate the press release. It clarified its purpose.

In the gatekeeper era, visibility depended on editorial approval. In the browsing era, visibility depended on clicks. In the retrieval era, visibility depends on eligibility.

Eligibility for indexing. For candidate selection. For summarization. For reinforcement. For compounding authority over time.

The press release, when constructed properly, is built for eligibility. It always has been. We just didn't have the language for it until now.

The Single Discipline

If this entire book were reduced to one operational discipline, it would be this:

Publish one clear, structured, newsworthy release every 28 days.

Not filler. Not vague enthusiasm. Not self-congratulation. Clear. Specific. Time-anchored. Category-aligned.

Then structure it properly. Host it persistently. Canonicalize it. Reinforce it across channels. Measure it honestly. And repeat.

Over time, that discipline creates reinforced entity salience, reduced ambiguity, stronger co-occurrence networks, expanded retrieval eligibility, and durable Long Fat Tail accumulation.

None of those sound dramatic in isolation. Together, they compound. And compounding, as we've seen throughout this book, is where the real advantage lives.

I've watched companies chase shortcuts for three decades. I've watched them spend enormous budgets on campaigns that spike and fade. I've watched them ignore the steady work because it didn't feel exciting enough. And I've watched the companies that embraced discipline, quietly and consistently, build visibility that their competitors couldn't replicate even with twice the budget.

The discipline is the advantage.

From Guarantees to Probability

You can't guarantee inclusion in every AI answer. Anyone who promises that is selling something that doesn't exist. What you can do is increase retrieval likelihood, strengthen authority, and extend long-term discoverability.

Every structured release you publish increases the probability that your brand appears in relevant results. That incremental increase, repeated twelve times a year, creates a structural advantage that is very real and very durable.

Stop treating press releases as events. Start treating them as infrastructure. Your archive of consistently published releases is a network. Every node connects to other nodes through topic associations, entity references, and temporal relationships. The network gets stronger with every addition.

The Chain

There's a sequence that drives everything in this book:

Eligibility precedes selection. You must be findable before you can be chosen.

Selection precedes visibility. You must be chosen before you can be seen.

Visibility precedes influence. You must be seen before you can shape perception.

Influence precedes revenue. You must shape perception before you can drive action.

Every link in that chain depends on the one before it. And the first link, eligibility, depends on everything we've discussed: structure, cadence, clarity, persistence, measurement, and discipline.

You can't skip links. You can't jump from invisibility to influence. You build the chain one release at a time.

The Advantage, Compounded

AI systems don't absorb your brand automatically. They retrieve what is visible, structured, reinforced, and trusted.

If you publish clearly. If you publish consistently. If you reinforce strategically. If you measure honestly. If you persist.

You increase the likelihood that when someone asks a question in your category, your brand is part of the answer.

Not guaranteed. But advantaged.

And in competitive markets, structural advantage, compounded over time, wins.

That's the visibility advantage. Not a trick. Not a hack. Not a secret algorithm. A system. Built on discipline. Reinforced by consistency. Measured by honesty. Compounded by time.

Start your 28-day cadence. Structure your releases. Host them permanently. Measure what happens. Adjust. Repeat. Keep going.

The companies that do this work will be the ones that show up when it matters. The rest will wonder why they didn't start sooner.

Chapter Sixteen

Additional Resources

The News Marketing system works best when it's connected to a broader set of tools and habits. Below are the resources I rely on personally, recommend to clients, and continue to build. Some are mine. Most are not. All of them are useful.

Tools Built for the News Marketing System

Newsworthy.ai Community

Newsworthy.ai is more than a newswire. It is a community of marketers, PR professionals, content creators, and small business operators learning to publish in the AI era. Register a free account, join the community boards, and you will find practitioners working through the same questions this book raises. Threads cover everything from schema implementation to AI citation tracking to monthly cadence accountability. The community is the fastest way to shorten your learning curve and watch the system work in real environments. Visit https://newsworthy.ai to get started.

FeedworthyAI.com

FeedworthyAI is an RSS feed directory that takes your existing content

feeds, articles, podcasts, video posts, newsletters, and applies the structured data and schema markup that makes them retrievable by AI systems. If you already publish content somewhere, FeedworthyAI helps that content show up in places it would otherwise miss. It is one of the easier ways to extend the principles in this book to assets you have already created without re-publishing them.

NewsCrafters.com

NewsCrafters is a free service that lets you become a news curator in your category. The curation tools are simple, the workflow is fast, and the result is a published presence that signals topical authority to both human readers and AI systems. Curating news is one of the most underrated brand-building activities available to small businesses and solo operators. NewsCrafters makes it possible to do consistently without a dedicated content team.

NewsRamp

NewsRamp is the companion platform to Newsworthy.ai. Where New sworthy.ai distributes one release at a time, NewsRamp expands a single release into a full suite of AI-optimized content, including audio briefings, podcast clips, video summaries, social posts, and zero-click ready chunks. If you are running the 28-day cadence and want maximum yield from each release, NewsRamp is the multiplier. NewsRamp is a service that is available to Newswires. To amplify your news content with NewsRamp use one of the newswires listed on the site. You can also use BoostifyPR.com to gain access to the NewsRamp tools independent of a partner newswire.

Reading I Recommend

The Long Tail by Chris Anderson Anderson's 2006 book reshaped how the marketing world thought about distribution, scale, and the value of niche content. The Long Fat Tail framework in this book builds directly on his thesis. If you want to understand why consistent publishing compounds, start here.

The Wizard of Ads by Roy H. Williams Williams's writing on persuasion, branding, and what he calls Brandable Chunks shaped how I think about every press release I write. The book is dense, original, and rewards re-reading. The Wizard Academy in Austin, Texas, runs in-person workshops worth attending if you can make the trip.

Google's Search Quality Rater Guidelines This is the official document that defines how Google evaluates content quality, including the E-E-A-T framework. It is long and dry, but reading even the first thirty pages will change how you think about every piece of content you publish. The latest version is available at https://developers.google.com/search/blog.

The original GEO research paper Pranjal Aggarwal and colleagues at Princeton, Georgia Tech, and the Allen Institute for AI published the foundational paper on Generative Engine Optimization in 2024. If you want to understand the academic basis for the strategies in this book, it is freely available at arXiv:2311.09735.

Rand Fishkin's research at SparkToro Fishkin publishes some of the most rigorous independent research on search behavior, zero-click trends, and AI's effect on traffic patterns. His blog at https://sparktoro.com is required reading for anyone trying to understand where attention is actually moving.

Tools for Implementation

Schema.org

The official documentation for structured data markup. If you are implementing schema yourself or briefing a developer, this is the source of truth. Visit https://schema.org.

Google's Rich Results Test

A free tool that validates your schema markup before you publish. Paste a URL or your raw HTML and it will tell you whether the structured data is correct and what rich results it qualifies for. Visit https://search.google.com/test/rich-results.

Google Trends

A free tool from Google that shows how search interest in any topic, term, or entity changes over time. Most marketers use it once and forget about it. Used systematically, it becomes one of the most valuable inputs to your News Marketing cadence.

Three specific ways to use it:

First, time your releases to rising interest. If you cover AI governance and Google Trends shows search volume climbing in your category, that is the month to publish. You will catch a wave that retrieval systems are already paying attention to.

Second, validate your topic clusters. Type in the topics you cover and see whether they are trending up, flat, or declining. If you are publishing twelve releases a year about a topic with falling interest, you may be reinforcing a position in a shrinking conversation. Adjust accordingly.

Third, identify rising entity associations. The "Related queries" and "Re-

lated topics" panels show what people are searching for alongside your terms. Those adjacencies are clues about what AI systems are starting to associate with your category. Publishing a release that bridges your brand to a rising adjacency is one of the highest-leverage moves available to a small marketing team.

Visit https://trends.google.com.

Schema Markup Validator. A complementary tool maintained by Sch ema.org itself, useful for catching issues that the Rich Results test misses. Visit https://validator.schema.org.

Perplexity, ChatGPT, Claude, and Google's AI Mode. Use them as monitoring tools, not just productivity tools. Once a month, ask the questions your customers ask. Record whether your brand appears, which competitors appear, and what sources are being cited. This is the simplest, cheapest visibility audit available, and almost no one does it consistently.

Communities and Conversations

The Newsworthy.ai Community Boards. Already mentioned above. Worth listing twice because it is the single fastest way to find peers running this system in their own businesses.

Wizard Academy. Workshops on persuasion, branding, and copywriting in Austin, Texas. The Magical Worlds courses in particular is directly relevant to the framing principles in this book. Visit https://www.wizardaca demy.org.

SBIR Advisors. If you are a defense, aerospace, medtech, or dual-use technology company working with government contracts, SBIR Advisors specializes in press release strategy for that environment. Disclosure: I work

with them as a content partner.

Stay Current

The AI landscape moves quickly. Anything written in a book about generative AI is a snapshot of one moment. To keep up after you finish this book:

Subscribe to the NewsMarketingBook.com newsletter for monthly updates on what is changing in the AI retrieval landscape and how the News Marketing system is evolving in response.

Follow the Newsworthy.ai community boards for real-time conversations about new platforms, schema changes, and citation patterns.

Watch the SparkToro and Search Engine Land blogs for independent research on search and AI behavior.

Set a monthly calendar reminder to ask the major AI systems your category questions and track whether your brand appears.

The discipline that makes this book work, one clear release every 28 days, is the same discipline that keeps you current. Show up regularly. Pay attention to what changes. Adjust without abandoning the core practice.

Chapter Seventeen

Glossary

The following terms were mentioned throughout the book. Hopefully you were able to grasp the concepts in context. If now, for you review...

28-Day Rule: A publishing cadence recommending one clear, structured, newsworthy press release every 28 days to maintain reinforcement momentum and build topic clustering.

AI (Artificial Intelligence): Computer systems that simulate human reasoning, including language understanding and content generation. In this book, AI refers primarily to retrieval-based systems layered on top of search engines.

AI Grounding: The process by which AI systems connect generated responses to verified, retrievable sources rather than generating claims from training data alone.

AI Training: The computational process of building a language model's knowledge base from massive text datasets. Once complete, the model's knowledge is fixed (the "frozen brain").

AI Visibility: The probability that a brand appears in AI-generated re-

sponses when relevant category questions are asked. The core objective of the 28-day publishing cadence.

AIO (AI Optimization): Optimizing content for AI-generated overviews and summaries within search engines, such as Google's AI Overviews.

Authority Gradient: The concept that visibility is relative, not absolute. Your retrievability exists on a spectrum influenced by how consistently you reinforce topic associations compared to competitors.

Brandable Chunk: A concept borrowed from Roy Williams of Williams Marketing and the Wizard Academy, adapted for the retrieval era. A dense, self-contained passage that includes a brand entity, topic context, time anchor, and measurable claim. Aligns with how AI systems chunk content for retrieval.

Canonical URL: A tag designating the primary version of a webpage when duplicates exist. Consolidates authority instead of fragmenting it across copies.

Co-Occurrence: The repeated appearance of two or more entities together in structured documents. Strengthens relational mapping in Knowledge Graph systems.

Concept Chunking: How AI systems break down content into discrete conceptual units for evaluation and retrieval, rather than processing documents as single blocks.

Contextual Neighborhood: The cluster of entities, standards, organizations, and topics that regularly appear alongside your brand. A strong neighborhood reinforces authority.

E-E-A-T: Experience, Expertise, Authoritativeness, and Trustworthiness. A credibility framework used by search systems to evaluate content reliability.

Engagement: User interaction with content including clicks, time on page, shares, and downstream actions. The multiplier that converts retrievability into influence.

Entity: A clearly identifiable subject within search and AI systems: organizations, people, products, locations, regulations. Modern models map entities rather than relying on keywords alone.

Entity Authority: An emergent property created through consistent, structured reinforcement of a brand's association with specific topics over time. Not declared or purchased; accumulated.

FAQ Schema: Structured data markup presenting content in question-and-answer format, improving retrieval alignment for conversational queries.

Frozen Brain: A metaphor explaining that LLMs are trained on data and then fixed. Systems retrieve indexed sources rather than "knowing" answers from training alone.

GEO (Generative Engine Optimization): Optimizing content for visibility across all generative AI platforms, not just search engines. Encompasses ChatGPT, Perplexity, Gemini, and others.

Knowledge Graph: A structured map of entities (nodes) and their relationships (edges). Search engines use these to model connections between organizations, topics, people, and events.

Long Fat Tail: The compounding visibility effect created by persistent hosting of press releases. Archived releases continue generating value indefinitely, unlike short-term marketing spikes.

News Drafting: Anchoring press releases to current events, regulatory changes, or industry trends to increase relevance and engagement through contextual alignment.

Publisher Schema: Structured data identifying an organization as a content publisher, strengthening credibility signals for search and AI systems.

Retrieval: The process by which search engines or AI systems locate relevant documents in response to a query. The prerequisite for AI visibility.

Retrieval-Augmented Generation (RAG): A system architecture where AI retrieves relevant documents first, then generates a response based on those documents. The mechanism behind most modern AI answer systems.

Retrievability: The likelihood that your content will be selected as a candidate source during search or AI-driven responses.

Semantic Similarity: A measurement of how closely two texts relate in meaning, calculated using vector embeddings. Higher similarity to a query means higher retrieval probability.

SERP (Search Engine Results Page): The results displayed after a search query. Modern SERPs include AI summaries, knowledge panels, and featured snippets alongside traditional links.

Source of Truth: A persistent, timestamped, canonical document serving as a verifiable record of a claim or announcement.

Structured Data (JSON-LD): Machine-readable markup embedded in webpages that clarifies entities, dates, authorship, and relationships to reduce ambiguity.

Temporal Continuity: The reinforcement of authority over time through consistent publishing cadence and stable archival hosting.

Topic Clustering: Repeated alignment of a brand with a defined topic domain, strengthening semantic association within retrieval systems. Accumulated through disciplined repetition.

Visibility Advantage: The compounded strategic benefit achieved through structured clarity, reinforcement cadence, engagement, and credibility over time.

Zero-Click Content: Content consumed directly within search results or AI summaries without the user clicking through to the source. Makes structured clarity and entity association even more critical.

Endnotes

1. Pranjal Aggarwal, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, Karthik R. Narasimhan, and Ameet Deshpande, "GEO: Generative Engine Optimization," Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24), August 2024. Preprint: arXiv:2311.09735.

2. Paul B. Allen co-founded Ancestry.com in 1997 and served as its first CEO. He is not the late Microsoft co-founder of the same name. Allen has since founded several technology companies, including FamilyLink, Soar.com, and CitizenPortal.ai, an AI platform that indexes public meeting transcripts at the local, state, and federal levels.

3. The "things, not strings" formulation comes from Amit Singhal, "Introducing the Knowledge Graph: things, not strings," Google Official Blog, May 16, 2012. Singhal, then SVP of Engineering at Google, used the phrase to describe the shift from keyword-based search to entity-based search. https://blog.google/products/search/introducing-knowledge-graph-things-not/

4. The Brandable Chunks framework was developed at Wizard Academy in Austin, Texas. The concept was originally conceived by Christopher J. Maddock, Roy H. Williams's first apprentice, and is referenced in Williams's The Wizard of Ads (Bard Press, 1998), p. 140. Williams, Maddock, and Jeff Sexton later co-taught a two-day Brandable Chunks workshop at Wizard Academy. https://www.wizardacademy.org

5. Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela, "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," Advances in Neural Information Processing Systems 33 (2020). Preprint: arXiv:2005.11401. Chunk-size conventions vary by implementation. Practical guidance from LangChain and LlamaIndex documentation cites 256 to 1,024 token windows as common, with 300 to 500 characters representing a typical paragraph-level segment for press release and news content.

6. Rand Fishkin, "2024 Zero-Click Search Study: For every 1,000 EU Google Searches, only 374 clicks go to the Open Web. In the US, it's 360," SparkToro, July 1, 2024. The study, conducted with clickstream data from Datos (a Semrush company), found that 58.5 percent of US Google searches end without a click to the open web. https://sparktoro.com/blog/2024-zero-click-search-study-for-every-1000-us-google-searches-only-374-clicks-go-to-the-open-web-in-the-eu-its-360/

7. Chris Anderson, "The Long Tail," Wired, October 1, 2004; expanded in The Long Tail: Why the Future of Business Is Selling Less of More (New York: Hyperion, 2006).

About the Author

David McInnis has spent more than three decades building tools that help small businesses be heard. His entrepreneurial streak started early, improbably enough, with a lollipop business he launched at thirteen.

That instinct never left him. Over time, it grew into a belief that would shape his career: visibility should not belong only to organizations with the biggest budgets.

At the time, that idea ran against the grain. The PR industry revolved around the pitch: the phone call, the relationship, the gatekeeper. McInnis was not interested in improving that system. He wanted to go around it entirely by putting news directly in front of the people who needed it. Inspired by the pharmaceutical industry's move toward direct-to-consumer advertising, he saw press releases not as private messages for journalists, but as public, searchable assets anyone could discover.

It worked. PRWeb helped launch countless startups and gave tens of thousands of small and midsize businesses their first real opportunity for online visibility, long before content marketing became a recognized discipline. In 2006, PRWeb was acquired by Vocus, Inc., which later merged into Cision, sitting alongside PR Newswire.

After PRWeb, McInnis kept building at the intersection of news, technol-

ogy, and access. Through Cranberry.com, he launched content marketing platforms aimed at helping brands create and distribute content more effectively, experimenting with persona-based native content formats. He also founded PriorThings.com, a blockchain registration platform for content creators that today powers the content provenance engine within Newsworthy.ai, BoostifyPR and NewsRamp.

In 2021, he co-founded Newsworthy.ai with Mark Willaman, a long-time friend and integration partner at PRWeb, creating what became the internet's first News Marketing platform built for the AI era. Newsworthy.ai was the first newswire to register press releases on the blockchain, establishing cryptographic proof of content provenance and origin.

It was also the first newswire to let companies self-host their press releases and among the first to meaningfully integrate AI into press release distribution through tools such as Newswriter.ai.

He later launched NewsRamp, a companion software-as-a-service platform for press release newswires that expands a single press release into a full suite of AI-optimized content, including articles, blog posts, FAQs, audio briefings, podcasts, video summaries, and zero-click social content, then distributes it across hundreds of media endpoints. Together, News worthy.ai and NewsRamp form the foundation of the system described in this book.

McInnis is widely credited with pioneering *News Marketing* as a discipline: the strategic use of news-format content as discoverable, multi-format assets designed to reach customers, AI systems, and media wherever they consume information. He has spent nearly thirty years studying how press releases move through each era of digital discovery, from early search engines to Google's dominance to today's AI retrieval systems.

He holds a BA in International Studies from **Texas State University**. He is a speaker, trainer, and mentor who regularly leads AI and digital marketing bootcamps. He is based in Boerne, Texas.

Over the course of his career, he has built three companies around one core belief: a well-structured press release, made public and preserved, is one of the most powerful marketing assets a business can own.

This is his first book on the subject. It will not be his last.