

WREN'S WATCH

The Fallow Web

The Hidden Cost of Content That Stopped Working

A whitepaper on content ROI and the case for semantic re-use · Wren's Watch LLC · Pilot

Executive Summary

Most website content has stopped working, and not because it was poorly made or because somebody stopped caring. It stopped working because all content decays: search rankings slip, competitors publish fresher material, algorithms reprioritize, and traffic drains away while the content just sits there, costing money it no longer earns.

This paper makes three arguments, each of which follows from the one before it. Fallow content doesn't have zero ROI; it has negative ROI, because it continues to cost the organization in ways that are genuinely easy to miss. The obvious response (deleting the dead weight) destroys real value in the process and doesn't fix anything structural. The correct response is elevation: treating the existing archive as a knowledge asset and re-deploying it, rather than replacing it or just walking away.

The final section describes how Pilot, Wren's Watch's semantic knowledge platform, is built specifically for that third path.

SECTION 1

Fallow Content Has Negative ROI

The Scale of the Problem

Ahrefs analyzed approximately 14 billion web pages in 2023 and found that 96.55% of them receive zero organic traffic from Google, and a companion study of one billion pages came back with a consistent finding of 90.88%. The precise figures differ by methodology and sample composition, but the structural reality they describe is the same: the overwhelming majority of web content is invisible to search.

The standard interpretation is that most content is simply bad, or unoptimized, or competing in categories too crowded to penetrate, and that's partly right. But it misses something important: much of this dormant content was not always dormant. It ranked. It drove traffic. Then it didn't.

Source [1]: Ahrefs, "96.55% of Content Gets No Traffic from Google," January 2024. ahrefs.com/blog/search-traffic-study

Semrush research published in 2024 found that 82% of high-ranking posts begin losing traffic within 12 to 24 months without active maintenance, which is consistent with primary case study data from Animalz: their research measured an average weekly traffic loss of 1.21% in content approaching decay, a rate that compounds to a roughly 45% annual decline.

Source [2]: Semrush, 2024, cited in practitioner literature. Corroborated by Animalz primary case study data (Source [3]).

So most organizations are not simply sitting on content that never worked; they're sitting on content they paid to produce, watched perform, and then watched silently decline, without any corresponding change in accounting treatment.

<p>96.55%</p> <p>of web pages receive zero organic traffic</p> <p><i>Ahrefs, 14B pages, 2023</i></p>	<p>82%</p> <p>of top-ranking content decays within 24 months</p> <p><i>Semrush, 2024</i></p>	<p>1.21%</p> <p>average weekly traffic loss before decay</p> <p><i>Animalz case study data</i></p>
---	---	---

The ROI Equation

The standard framing is that fallow content has zero ROI: no return on the original investment, but no ongoing cost either. That framing is wrong on the cost side, because fallow content incurs four categories of ongoing cost that don't show up anywhere in a typical content audit:

Production cost (sunk)	The content was written, edited, optimized, and published. That cost is written down regardless of performance. This is the visible number, and in most organizations, the only number tracked.
Crawl budget dilution	Search engines allocate a finite crawl budget per domain. Low-quality, zero-traffic pages consume that budget without contributing ranking signals, reducing crawl efficiency for the pages that do perform.
Domain authority dilution	Google's quality signals are domain-wide, not page-specific. A site with hundreds of thin, dormant pages signals lower overall quality than a site of equivalent size but higher content density. Fallow content is a drag on the active content's authority.
Opportunity cost	Every keyword position held by a decayed piece is a position being captured by a competitor. The organization is not merely failing to gain ground; it is ceding existing ground by not maintaining what it built.

The aggregate effect is a content library that presents itself as an asset on a production budget while functioning in part as a liability on SEO performance, and the organization ends up paying twice: once to create content that is no longer working, and again in reduced visibility for the content that is.

A CONCRETE ILLUSTRATION

An organization with 1,000 published pages, 80% of which receive no meaningful traffic, is not maintaining a library of 1,000 assets: it is maintaining roughly 200 active assets that happen to be dragged down by 800 pages consuming crawl budget and diluting quality signals. The 800 are not neutral. They are a headwind.

Deletion Is Not the Answer

The intuitive response

If fallow content has negative ROI, the logical fix appears to be removal: run an audit, surface the underperformers, delete the zero-traffic pages, and the site gets cleaner, crawl budget gets freed, and quality signals improve. This is a real benefit, and content audits that result in strategic pruning do typically produce measurable SEO gains.

But deletion is a maintenance operation, not a strategy, and applied broadly it destroys three categories of value that don't show up in traffic metrics.

What traffic metrics do not capture

Backlink equity. Links from external sites pointing to a specific URL are real assets: editorial decisions by other publishers to cite or reference that content. When a URL is deleted without a redirect, those links go dead and the authority they carried evaporates, and a page with zero traffic can still carry dozens of inbound links from credible domains. Deletion destroys that equity permanently.

Topical authority signals. Search engines evaluate domain authority partly by topical coverage depth, and a site with extensive coverage of a subject space (even if individual pieces within it are underperforming) signals greater authority than a site with sparse coverage. Aggressive deletion can reduce apparent topical depth and harm rankings for active content in the same cluster.

Institutional knowledge. This is the most undervalued category, and the one most directly relevant to Pilot's use case. The content in a dormant archive was written by subject matter experts, reviewed by practitioners, and refined for a specific audience, and it contains analysis, framing, and organizational perspective that does not exist anywhere else. Its traffic value has gone to zero. Its knowledge value has not.

The distinction matters because traffic value and knowledge value are measured by different instruments: traffic value shows up in Google Analytics, while knowledge value does not show up anywhere in most organizations, which is precisely why it gets treated as equivalent to zero when deletion decisions are made.

The deletion paradox

There is a deeper problem with deletion as a strategy: it doesn't address the underlying dynamic at all. Content decays because the organization published it and moved on, and deleting decayed content to replace it with new content that will follow the same curve isn't a solution. It is the same cycle, repeated.

THE CORE PROBLEM IS NOT THE CONTENT. IT IS THE TREATMENT OF THE ARCHIVE.

An organization that deletes fallow content and creates new content has simply reset the clock. The new content will follow the same decay curve. The institutional knowledge in the deleted content is gone. The production cost of the replacement has been incurred. Nothing structural has changed.

SECTION 3

The Treadmill and Its Cost

How most organizations actually operate

Without a deliberate strategy for managing content lifecycles, most organizations converge on the same pattern: continuous net-new production as the primary mechanism for maintaining and growing organic visibility, where new content compensates for the traffic lost as older content decays and the creation rate has to stay high enough to offset the decay rate across the entire library. This is the content treadmill.

The treadmill is expensive. Industry pricing data from 2024-2025 puts a publication-quality blog post at \$600 to \$1,500 to produce at the freelancer level, and full-service enterprise content programs at \$25,000 to \$60,000 per month.

Source [4]: Siege Media, BrandWell, Digital Elevator, 2024-2025. Agency pricing guides; treat as indicative ranges.

\$1,500 Typical cost of one substantial blog post <i>Agency market rates, 2025</i>	\$18K/mo Estimated mid-market content program <i>Aggregated industry data</i>	12–24 mo Average time to decay for top-ranking content <i>Semrush / Animalz, 2024</i>
---	--	--

The treadmill has a compounding inefficiency: because each new piece follows the same decay curve, the organization's total traffic is always the net of what new content is gaining and what old content is losing. Animalz describes this precisely: new content gains traffic for a period, plateaus, and decays, and the cumulative effect is growth but growth that is persistently slowed by the drag of decaying older content.

Source [3]: Animalz, "Content Refreshing: How to Win Traffic by Updating Old Content." animalz.co/blog/content-refresh

This pattern has a specific characteristic: it cannot be resolved by producing more content, because higher production volume increases both the gain rate and the eventual decay drag proportionally. The treadmill moves faster. The runner does not get ahead.

The refresh partial solution

Content refresh (updating existing pieces with current information, expanded coverage, and re-optimized structure) is well-documented as more efficient than net-new creation for the same keyword targets:

- Animalz documented a case where a content refresh reversed a -1.21% weekly decay rate and produced a 55% increase in weekly pageviews, with continued growth at approximately 12% per week post-refresh. (Source [3])
- Orbit Media's 2024 Annual Blogging Survey of over 1,000 bloggers found that 42% identified content refreshing as the single activity delivering their largest SEO gains. (Source [5])
- HubSpot reports, in aggregated platform data, that republishing refreshed content can increase organic traffic by up to 106%. (Source [6], flagged in the appendix as conventional wisdom rather than a primary study.)

Refresh is better than replacement, but it is still primarily a traffic-recovery operation applied to specific high-potential pages and doesn't address the broader question of what to do with the institutional knowledge embedded across the full archive, including content that isn't worth refreshing for traffic purposes but still contains valuable organizational intelligence.

SECTION 4

Elevation: A Different ROI Equation

Reframing the archive

The fallow content problem looks different when you separate traffic value from knowledge value: traffic value is what a page is worth as a search engine destination (visits, leads, conversions), while knowledge value is what the content embedded in that page is worth as organizational intelligence: analysis, framing, expert perspective, documented experience.

These two values are correlated early in a content piece's life, but as the piece ages and decays, traffic value goes to zero while knowledge value does not follow: the research still reflects expert judgment, the analysis still documents organizational thinking, and the framing still represents how the organization understood its field at that moment.

The organization that treats its archive only as a traffic asset will watch the ROI of that archive trend toward zero and then below. The organization that also treats it as a knowledge asset has a different inventory on its hands.

<p>PRODUCTION INVESTMENT</p> <p>\$720K-\$1.4M</p> <p><i>5-year mid-market spend</i></p>	÷	<p>TRAFFIC VALUE ONLY</p> <p>→ \$0</p> <p><i>For dormant content</i></p>	=	<p>REALIZED ROI</p> <p>Negative</p> <p><i>After ongoing drag costs</i></p>
--	---	---	---	---

Add knowledge value to the numerator and the equation changes entirely: that same investment that produced dormant traffic assets also produced an organizational knowledge base of documented

expertise, historical perspective, and institutional intelligence that would cost substantially more to reconstruct from scratch than to surface and re-deploy from what already exists.



PUTTING A NUMBER ON IT

Let C equal the fully-loaded cost of producing one net-new page of content at your organization, whatever that number is. Elevation of an existing page typically runs 0.25 to 0.35× C, because the ideation cycle, structural drafting, and research from scratch are already done. The traffic upside is documented: Animalz measured a 55% pageview increase from a targeted refresh of a single decayed piece (Source [3]), and applying that against even a modest prior baseline puts the return-per-dollar at roughly 1.5 to 2× what net-new production yields on traffic alone, before any knowledge re-use value, which has no equivalent in the net-new model. One honest boundary: this math only holds for content that had a traffic baseline to recover. A piece that never ranked has no baseline; its elevation case rests entirely on knowledge value, not traffic recovery. The cost multiple still applies. The traffic upside does not.

What elevation looks like in practice

Elevation is not the same as refresh, and the distinction matters. Refresh asks: can this specific piece recover traffic? Elevation asks: what value does the material in this piece unlock for the organization’s current work?

The two questions lead to different actions: a refresh operation updates statistics and re-optimizes for current search intent, while an elevation operation surfaces the knowledge embedded in that content, connects it to related material across the archive, and makes it available as source material for new production, research, and strategic synthesis.

A piece on zero-trust security from 2020 may be too dated to refresh for traffic, but it contains practitioner framing directly relevant to something being written today and documents how the organization understood that problem before the market had fully formed its vocabulary. That perspective has value. The elevation question is how to surface and apply it.

THE COMPOUND EFFECT

Each piece of content produced by an organization is simultaneously a traffic asset and a knowledge asset, and while traffic value depreciates on a predictable schedule, knowledge value does not: when properly indexed and made accessible, it appreciates. New content enriches the archive, which enriches new content, in a compounding cycle that a pure production model cannot replicate.

Pilot: Semantic Knowledge Platform

Designed for the elevation problem

Pilot is Wren's Watch's semantic knowledge platform, built to activate the knowledge value in an organization's existing content archive, and it operates on a simple premise: the organization has already paid for a substantial knowledge asset, and Pilot makes that asset queryable, connective, and generative.

Where a standard content audit surfaces performance data (traffic, rankings, conversion rates), Pilot surfaces semantic data: what the content is about, how it relates to other content in the archive, what questions it answers, what perspectives it contains, and where it connects to current work and current knowledge gaps.

01

Semantic indexing of the existing content archive. Pilot ingests an organization's published content and builds a semantic index that reflects meaning, not just keywords, making the archive searchable by concept and context rather than just by the terms that happen to appear on the page.

02

Surface and connection. Pilot identifies relationships between content pieces that are not visible from traffic data: thematic overlaps, complementary arguments, contradictions, and gaps. It answers the question of what the organization already knows about this topic, and where the edges of that knowledge are.

03

Knowledge re-use for content production and research. When an organization is developing new content or conducting research, Pilot surfaces the relevant material already in the archive as source material, context, and reference, so the production process starts from existing organizational intelligence rather than from scratch.

The ROI case

The ROI case for a semantic knowledge platform doesn't rest primarily on traffic recovery, though that is one output. It rests on the economics of re-use versus replacement.

An organization using Pilot to develop new content isn't starting from zero; it is starting from an indexed body of existing organizational knowledge, and that changes the production economics: research time is reduced, drafts are more grounded in the organization's actual perspective, and the risk of duplicating or contradicting existing content is lower.

More broadly, Pilot changes the depreciation profile of the content archive: in a standard content program, production investment depreciates to near zero as content decays, but in a Pilot-enabled program, the same investment also funds a knowledge base that does not depreciate on the same

schedule. Content that is fallow from a traffic perspective continues to contribute to the archive's knowledge value, and that contribution compounds as new content is added.

THE INVESTOR FRAMING

A content archive without semantic access is a warehouse of depreciating assets, while a content archive with Pilot is a compounding knowledge base. The production costs are the same. The long-term return on those costs is not.

CONCLUSION

Summary

The fallow content problem is real, large, and consistently underestimated, mostly because organizations measure content value by traffic and have no instrument for the rest of what it contains.

The argument this paper makes is sequential: fallow content has negative ROI because it carries ongoing costs that traffic metrics don't capture, and deletion removes those costs but destroys knowledge value and does nothing to address the underlying cycle. The content treadmill is the default operating mode, and it is structurally inefficient. Elevation (treating the archive as a knowledge asset) changes the ROI equation in a way that compounds over time.

Pilot is built for the elevation model: it is the operational layer that turns a warehouse of depreciating content assets into a queryable, generative knowledge base and makes every new dollar of production investment worth more than it would be otherwise.

For more information about Pilot, visit pilotwme.com · Wren's Watch at wrenswatch.com

Sources and Methodological Notes

Primary sources are cited with URLs. Where a statistic has methodological limitations, those are noted. Several frequently-cited content marketing statistics circulate with incorrect attribution; the notes below correct the most significant of these.

- [1] Ahrefs, "96.55% of Content Gets No Traffic from Google," January 2024. ahrefs.com/blog/search-traffic-study
Analysis of ~14 billion pages in the Ahrefs Content Explorer database. Earlier companion study of one billion pages found 90.88% zero-traffic. Both studies measure only pages in the Ahrefs index; fully unindexed pages are not captured, so the true figure is likely higher.
- [2] Semrush, 2024 content research, as reported in practitioner literature.
The specific "82% within 12-24 months" figure circulates through practitioner sources citing Semrush but has not been located in a standalone Semrush primary publication. It is used as a directional benchmark and is corroborated by the Animalz primary case study data in Source [3].
- [3] Animalz, "Content Refreshing: How to Win Traffic by Updating Old Content." animalz.co/blog/content-refresh
Primary case study with stated methodology. Measured -1.21% average weekly traffic decay in the 12 weeks preceding a content refresh. Documented one article achieving a 55% pageview increase post-refresh with continued 12% weekly growth.
- [4] Industry content pricing: Siege Media (siegemedia.com), BrandWell (brandwell.ai), Digital Elevator (thedigitalelevator.com), 2024-2025.
Agency pricing guides, not primary research studies. The \$600-\$1,500 per post and \$25,000-\$60,000/month enterprise figures are consistent across sources; treat as indicative market-rate ranges.
- [5] Orbit Media Studios, Annual Blogging Survey, 2024. orbitmedia.com/blog/blogging-statistics
Annual survey of 1,000+ bloggers. The 42% figure represents respondents identifying content refreshing as their single highest-impact SEO activity.
- [6] HubSpot, aggregated platform data: republishing refreshed content can increase organic traffic by up to 106%.
This figure is widely attributed to HubSpot and appears in multiple HubSpot-published articles but does not correspond to a standalone HubSpot research report. The Animalz primary case study (Source [3]) is the better-documented primary citation for the refresh premium.
- [7] HubSpot, 2026 State of Marketing Report. blog.hubspot.com/marketing/hubspot-blog-marketing-industry-trends-report
Annual survey of 1,500+ global marketing professionals. Referenced for the finding that measuring marketing ROI is the top challenge cited by marketers.
- [8] BrightEdge Research, 2019 Channel Share Report. brightedge.com/resources/research-reports/channel_share
Analysis of thousands of domains and tens of billions of sessions, May 2019. The 53% organic search traffic share figure has been measured consistently since BrightEdge began this study in 2014.
- [9] Demand Metric, "Content Marketing Benchmark Report," February 2014. demandmetric.com/content/content-marketing-benchmark-report
Survey of 521 marketing professionals, January 2014, conducted with Ascend2, sponsored by Marketo. This is the original source of the widely-cited "3x leads, 62% lower cost" statistic, which has been misattributed to the Content Marketing Institute in the majority of subsequent usage. The study is more than a decade old; figures are treated here as historical context, not current benchmarks.
-