

The LinkedIn Distribution Manual

How the feed ranks posts in 2026, what changed since 2020, and which surface gives you the best chance of being seen

David Soden

Independent analyst and practitioner, enterprise automation and applied AI
davidsoden.com/market-assessment

About this report

Every substantive claim in this report carries an evidence classification and, where applicable, a numbered source. Facts are separated from vendor claims, third-party estimates, the author's assessments, and untested hypotheses. Forward-looking sections use scenarios with observable tripwires rather than point forecasts. Stated assumptions are listed before the analysis begins.

PUBLISHED FOR PUBLIC READING | SHARE FREELY, UNMODIFIED, WITH ATTRIBUTION

© 2026 David Soden, davidsoden.com/market-assessment. All rights reserved. You may share this file complete and unmodified, and quote short passages with attribution. Republishing under another name, excerpting into a commercial deliverable, resale, and use as machine learning training data require written consent. Full terms appear under Rights and Permitted Use on the final page.

Market Assessment reports are published and commissioned at davidsoden.com/market-assessment. This document is analysis, not investment, legal, or procurement advice.

Scope and evidentiary standard

This is a working manual for getting organic distribution on LinkedIn. It covers four publishing surfaces: personal profile posts (text, image, native document, poll), articles and newsletters, native video including the vertical feed, and company pages. Paid advertising, Sales Navigator, Recruiter, InMail, and direct outreach are outside it. Everything here is about what the ranking system does with something you published for free.

The published numbers in this category need heavier discounting than in most markets, for a specific reason. LinkedIn documents its feed thoroughly, but it documents it for machine learning engineers. The two best sources on how ranking works in 2026 are an arXiv paper about transformer architecture and an engineering blog post about GPU inference. Neither says a word about what a poster should write. Meanwhile every source that does address what to write, without exception, measures reach using data from a third-party scheduling or analytics tool, which means the sample is people who pay for social media software. That is not the general population of LinkedIn, and nobody has audited the gap.

So the evidence splits cleanly into two piles that never touch. Mechanism is well documented and comes straight from LinkedIn. Outcomes are measured only by commercially interested outsiders with self-selected samples. Where this report connects the two, that connection is labeled as the author's assessment, not as fact, because no public dataset licenses the inference. LinkedIn's own analytics give each member their own numbers and no comparison set, so even a diligent poster cannot check most of these claims against anything but themselves.

One further caution about the format studies. Almost none of them control for who posts what. If people who already write well gravitate toward native documents, the document format will look powerful whether or not the format itself does anything. Read every format number in this report as a correlation with an unexamined confound sitting behind it.

Label	Definition	How to treat it
FACT	Verifiable in the cited primary source: a filing, press release, official documentation, or standards body announcement.	Reliable as stated. Check the source if the exact wording matters.
VENDOR CLAIM	Reported by a vendor about its own business or product. Self-measured and not audited by any third party.	The vendor said it. Directionally useful, not a measurement. Definitions of the metric are usually undisclosed.
THIRD-PARTY ESTIMATE	A research firm forecast or survey result. Methodology is proprietary and generally not reproducible.	Use for direction and order of magnitude only. Do not build a model on it.
ASSESSMENT	The author's reasoned judgment from the cited evidence. Falsifiable, and the reasoning is shown so it can be argued with.	This is analysis, not reporting. Disagree with it where your evidence differs.
HYPOTHESIS	A proposition offered for testing. Not currently supported by sufficient evidence to assert.	Treat as a research question, not a conclusion. Each one names what would confirm or refute it.
SCENARIO	A structured future state with a stated subjective probability. The probability is the author's judgment, not a calculated figure.	Use the leading indicators attached to each, not the probability. The indicators are observable; the probability is not.

Assumptions this analysis rests on

LinkedIn's engineering publications describe the system actually serving traffic, rather than research that never shipped. The Feed SR paper states plainly that the model is the primary member experience, and the March 2026 engineering post reads as a deployment write-up rather than a proposal. If either is aspirational, the architecture sections of this report are wrong, though the measured outcome sections survive because they never depended on the mechanism.

Third-party post samples are directionally representative of ordinary posting. Every reach and engagement figure here comes from a tool vendor measuring its own customers. If those customers differ systematically from everyone else, and they almost certainly differ somewhat, then the format rankings could compress or invert. The direction of the bias is unknown, which is worse than a large known bias.

Ranking behaviour is roughly uniform across regions, languages, and account ages. Nothing in the public record establishes this. LinkedIn's own paper reports that new members saw no statistically significant lift from the current ranker, which is direct evidence that behaviour is not uniform across account maturity. If it also varies by language or market, advice tuned to English-language samples will misfire elsewhere, and the surface ranking in this report would need rebuilding per market.

The observed relationships are stable enough to act on for six to twelve months. LinkedIn replaced its ranking architecture twice in roughly eighteen months and shelved a third approach after building it. If the cadence accelerates, the operating rules section decays faster than the half-life section predicts and readers should discount it accordingly.

The call

ASSESSMENT LinkedIn is no longer a virality machine, and treating it as one is what costs most posters their reach. The ranking system now reads what a post means against your profile and the reader's history, and it scores quiet signals like time spent above loud ones like reaction counts. Specific, consistent posting inside one recognisable subject beats every format trick. The surface with the best odds of being seen is not the feed at all, it is a newsletter, because subscribers are delivered to rather than ranked. I will move off this if LinkedIn publishes per-member reach data that contradicts the third-party samples everything here rests on.

My judgment on the evidence assembled here, nothing more. There is data I can't see and data I may have missed. New evidence can move me, including to the opposite position, and I reserve the right to move.

Executive summary

Two things happened between 2020 and 2026 that most posting advice has not caught up with. LinkedIn stopped treating engagement counts as the thing worth predicting, and it rebuilt the ranker around a member's interaction history read as a sequence. Both changes reward the same behaviour, and neither rewards the tactics that dominate the advice market.

FINDING 1 The ranker is a sequential transformer, and LinkedIn tried a large language model for the job and shelved it.

FACT LinkedIn's February 2026 paper describes Feed SR, a decoder-only transformer with causal attention that reads up to 1,000 of a member's past feed impressions as an ordered sequence. The same paper states that an LLM-ranker approach the team built "never achieved superior online performance over the existing production model." Advice built on the premise that a generative model now reads and judges your prose is describing retrieval, not ranking. ^{[1] [2]}

FINDING 2 Replacing the entire ranking model moved time spent by 2.10 percent.

FACT That is the headline online result from Feed SR against the DCNv2 model it replaced: +2.10 percent time spent overall, +2.38 percent among daily actives, and no significant effect on new members. LinkedIn's own best swing at the problem, with unlimited access to the data and the model weights, produced a low single-digit move. It is a useful ceiling to keep in mind when a post promises to triple your reach with a formatting change. ^[1]

FINDING 3 The model predicts six actions, and three of them make no sound.

FACT Feed SR is trained on two groups of binary outcomes. The active group is like, comment, and share. The passive group is click, skip, and long-dwell. Half the objective

function is made of things a reader does silently, and one of them, skip, is a penalty you incur simply by appearing in a feed and being scrolled past. ^[4]

FINDING 4 Dwell time is normalised within content type, which kills the format arms race.

ASSESSMENT In October 2024 LinkedIn replaced its fixed skip threshold with an auto-normalised long-dwell model that scores a post against a percentile of its own counterparts, recalculated daily, controlling for content type, creator type, and distribution path. A long video no longer beats a short text post simply by consuming more seconds. The practical reading is that you compete against posts like yours, not against the feed at large. ^[6]

FINDING 5 Which format "wins" depends entirely on which metric you pick, and the two available large samples disagree by a factor of three.

THIRD-PARTY ESTIMATE Socialinsider's benchmark of 1.3 million posts across 16,645 business pages puts native documents at a 7.00 percent engagement rate for 2025, with text at 4.50 percent. Buffer's cross-platform format study puts LinkedIn carousels at a 21.77 percent median engagement rate against 3.18 percent for text. Same format, same direction, wildly different magnitude, because the denominators and the samples differ. Neither number should be quoted without the other. ^{[11] [12]}

FINDING 6 Video is the format LinkedIn promotes and the one whose per-post views fell hardest.

THIRD-PARTY ESTIMATE Microsoft reported video watch time up 36 percent year over year and video posts shared twenty times more than any other content type. Socialinsider measured video views on business pages down 36 percent year over year across every follower band, with the 50,000 to 100,000 band down 53 percent. Both can be true at once if video supply grew faster than video demand, which is the author's reading. Aggregate watch time is a platform statistic. Views per post is your statistic. ^{[8] [11]}

FINDING 7 The link penalty is measured by outsiders, denied by LinkedIn, and the number everyone quotes has no method behind it.

ASSESSMENT LinkedIn's Senior Director of Product Management stated publicly in September 2025 that posts are not intentionally demoted for containing a link. Richard van der Blom's 1.3 million post sample reports one external link in the body cutting median reach by 18.8 percent. A Forbes column in July 2026 put the figure at 60 percent; following the citation leads to a marketing blog with no sample size, no date, and no methodology. Treat 18.8 percent as the defensible working estimate and the 60 percent figure as unusable. ^{[15] [19] [17]}

FINDING 8 AI-written text is now the majority of long-form posts, and LinkedIn has started building a member-reported signal against it.

FACT Originality.ai classified 81.2 percent of 5,000 sampled long-form July 2026 posts as likely AI. LinkedIn's own classifiers, per its chief product officer, put roughly 40 percent of long-form and 30 percent of short-form posts in the fully-AI bucket. On 30 July 2026 LinkedIn shipped a "Seems like AI slop" report button, retired its "enhance your post" writing tool in

favour of a proofreader, and said it blocks hundreds of thousands of automated comment attempts a day. ^{[16] [9] [10]}

FINDING 9 Ranked by odds of actually being seen, the order is newsletter, then personal document or text post, then native video, then company page.

ASSESSMENT A newsletter edition triggers a notification and an email to every subscriber, which is delivery rather than ranking. A personal post is ranked but starts from the strongest position in the feed. Native video is ranked and currently oversupplied. A company page is ranked from the weakest position on the platform and is best understood as an advertising and hiring asset rather than a reach asset. The combination that compounds is post to build subscribers, newsletter to keep them. ^{[11] [8]}

How the feed decides, in four steps

Two LinkedIn documents published in the first quarter of 2026 describe the current system end to end. Read together, they say something narrower than the coverage of them implied.

Step one: retrieval picks the candidates

FACT LinkedIn replaced a heterogeneous set of independent retrieval pipelines with a single LLM-based dual encoder that places posts and member profiles in a shared vector space. Candidate selection runs against an index of millions of posts at sub-50 millisecond latency, using GPU-accelerated nearest-neighbour search and three continuous nearline pipelines that refresh embeddings. ^[2]

ASSESSMENT This is where language models actually touch your writing. The embedding of your post is compared against an embedding of the reader's profile and interests. That is the mechanism behind the widely repeated advice to stay on one subject, and unusually for LinkedIn folklore, the mechanism supports the advice. A profile and a posting history that point at the same subject produce a tighter, more retrievable representation than one that wanders. ^{[2] [5]}

Step two: the ranker orders them

FACT Ranking is done by Feed SR, a decoder-only transformer with pre-layer normalisation, rotary position embeddings, and causal attention masking. It consumes interleaved post and action embeddings from up to 1,000 of the member's prior feed impressions, then concatenates context features after the transformer in a late-fusion design, and predicts through a mixture-of-experts head. ^[1]

FACT The same paper reports that LinkedIn built and tested an LLM-Ranker on live traffic and that it never beat the production model. Feed SR is described as "the primary member experience on LinkedIn's Feed." The 150 billion parameter 360Brew foundation model, published in January 2025 and covering more than thirty predictive tasks across the platform, was withdrawn from arXiv on 23 August 2025 over a licensing issue, and no public LinkedIn document places it in production as the feed ranker. ^{[1] [3]}

ASSESSMENT A great deal of 2026 posting advice is built on the claim that a 150 billion parameter language model now reads your post and judges its professional merit. LinkedIn's own engineers say otherwise in a peer-reviewable venue. The honest version is duller and more useful: a language model decides whether your post is a plausible candidate, and a sequence model decides where it lands. ^{[1] [2] [3]}

Step three: six predictions, three of them silent

FACT The multi-task head predicts binary outcomes in two groups. Active: like, comment, share. Passive: click, skip, long-dwell. Skip is a modeled negative, so an impression that produces nothing is not neutral. ^[1]


FACT Dwell time entered LinkedIn's feed ranking in May 2020, when the team shipped a P(skip) model built on two measurements: dwell on the feed, which starts when at least half a post is visible, and dwell after the click. Clicks alone were described as noisy because of "click bounces," where a member opens content and returns within seconds. The dwell features raised the area under the ROC curve by as much as 10 percent offline. ^[4]

FACT In October 2024 LinkedIn replaced the fixed threshold with an Auto Normalized Long Dwell model that predicts whether a member will spend more time on a post than a given percentile of its counterparts, normalising by content type, creator type, and distribution method, with percentiles recomputed daily. The stated purpose was to stop the model favouring categories such as long video over short articles. LinkedIn reported gains in sessions and time spent, "particularly pronounced among users who tend to passively consume content without explicit interactions." ^[6]

ASSESSMENT Two operating consequences follow, and they are the most actionable facts in this report. First, the reader who opens your post, reads it for forty seconds, and never touches a button is contributing more to your distribution than the reader who reflexively likes it in two. Second, because normalisation happens within content type, switching format to chase a mechanical advantage does not work. A document is not scored against a text post. It is scored against other documents. ^{[6] [1]}

Step four: the size of the effect

FIGURE 1 What a complete ranker replacement was worth, in time spent



FACT Online A/B lift in time spent for Feed SR against the DCNv2 production baseline it replaced. LinkedIn reported no statistically significant effect for new members, which is omitted here rather than charted as zero. These are platform-level metrics, not per-post reach, and LinkedIn chose which metrics to publish. ^[1]

ASSESSMENT This chart is the single best calibration device available for the whole category. LinkedIn rebuilt its ranking model with full access to the training data, the labels, and the serving infrastructure, and bought a 2.1 percent move in time spent. Any advice claiming a formatting change will multiply your reach is describing an effect larger than the one LinkedIn's own machine learning organisation extracted from replacing the entire model. That does not make such claims impossible, because per-post reach and platform time spent are different quantities. It does mean they need evidence proportional to their size, and they never have it. ^[1]

HYPOTHESIS Proposition offered for testing: the variance in individual post reach is dominated by topic-to-audience fit at retrieval, not by anything in the ranking stage that a poster can manipulate. This would be confirmed if a poster's reach distribution shifts sharply and durably when they change subject matter while holding format, length, and cadence constant, and refuted if format changes at constant subject produce comparable shifts. Any poster with 100 posts of history can run this on their own data, and it is the single most valuable experiment available to an individual.

Five years of change, dated

Reconstructing the timeline matters because most of what circulates as current advice was accurate in a specific year and has not been retired. The table below carries only changes with a locatable source and a date. Undated claims about what changed, however widely repeated, are not in it.

Date	What changed	Evidence	Type
May 2020	Dwell time enters feed ranking. A P(skip) model is trained on on-feed dwell and after-click dwell; clicks are downgraded as noisy because of click bounces.	LinkedIn Engineering Blog	Vendor primary
2022	Stated shift away from rewarding viral reach and toward knowledge and advice from people close to the reader. LinkedIn reports a large drop in complaints about irrelevant feed content.	LinkedIn statements via trade press	Independent press
Feb 2024	LinkedIn confirms it stopped using clicks on the "read more" button as a value signal, killing the mechanical incentive behind broetry. It also says useful posts may be resurfaced for months or years, decoupling reach from recency.	Dan Roth and Tim Jurka, via Social Media Today	Independent press
2024	Creator mode is folded into all profiles rather than being a toggle, and newsletter subscription becomes a standard profile call to action.	LinkedIn, via Social Media Today	Independent press
Oct 2024	Fixed skip threshold replaced by the Auto Normalized Long Dwell model. Dwell is scored as a percentile within content type, creator type, and distribution method, recomputed daily.	LinkedIn Engineering Blog	Vendor primary
Jan 2025	360Brew published: a 150 billion parameter decoder-only foundation model covering more than thirty predictive tasks across LinkedIn. Withdrawn from arXiv 23 August 2025 for a licensing reason, not a scientific one.	arXiv 2501.16450	Vendor primary
Sep 2025	LinkedIn's Senior Director of Product Management states publicly that posts containing external links are not intentionally penalised for reach.	Executive post on LinkedIn	Vendor primary
Oct 2025	Collaborative Articles retired. New contributions and edits stop; existing contributions go read-only. The associated Top Voice badges had already been withdrawn.	LinkedIn, via trade press	Independent press
Oct 2025	Microsoft reports comments up 24 percent year over year, video watch time up 36 percent, three straight quarters of double-digit growth in video uploads, and 1.3 billion members. The standing "record levels of engagement" line is absent from the earnings script for the first time since 2018.	Microsoft Q1 FY2026 remarks, via Social Media Today	Independent press

Date	What changed	Evidence	Type
Jan 2026	LinkedIn passes 5 billion dollars in quarterly revenue for the first time, up 11 percent, with the video push credited in coverage.	Microsoft Q2 FY2026 results, via GeekWire	Independent press
Feb 2026	Feed SR paper published. A sequential transformer over up to 1,000 impressions is the production ranker; the LLM-ranker alternative is documented as tried and not shipped.	arXiv 2602.12354	Vendor primary
Mar 2026	Full feed rebuild published: unified LLM dual-encoder retrieval, GPU-served, sub-50 millisecond latency, replacing the previous multi-source retrieval architecture.	LinkedIn Engineering Blog	Vendor primary
Jul 2026	"Seems like AI slop" report button ships. The "enhance your post" writing tool is retired in favour of a proofreader that preserves voice. LinkedIn says it blocks hundreds of thousands of automated comment attempts daily.	LinkedIn CPO Hari Srinivasan, via TechCrunch, Bloomberg, Fortune	Independent press

ASSESSMENT Read down the right-hand column and a pattern shows up that matters more than any individual row. Every change LinkedIn has made since 2020 removes a cheap proxy and replaces it with an expensive one. Clicks gave way to dwell. Fixed dwell gave way to normalised dwell. Independent retrieval pipelines gave way to semantic embedding. Each step makes the signal harder to counterfeit and harder to reason about from the outside. There is no reason to expect the direction to reverse. [\[4\]](#) [\[6\]](#) [\[2\]](#) [\[1\]](#)

ASSESSMENT Two rows in that table are easy to skim past and worth stopping on. LinkedIn retired Collaborative Articles in October 2025, having already withdrawn the Top Voice badges that made contributing worth anyone's time. A distribution surface the platform built, promoted, and attached a status reward to was closed inside roughly two years. Anyone whose strategy depends on a LinkedIn feature rather than on an audience should price that in. [\[21\]](#)

ASSESSMENT The other is the earnings language. In October 2025 Microsoft omitted LinkedIn's standing "record levels of engagement" line for the first time since 2018, then in January 2026 reported LinkedIn passing five billion dollars in quarterly revenue on the back of a video push. Revenue up, the engagement boilerplate quietly dropped. That combination is consistent with the practitioner reports of declining organic reach, and it is the closest thing to first-party corroboration those reports have. [\[8\]](#) [\[20\]](#)

ASSESSMENT The February 2024 retirement of the "read more" click is the most instructive single event in the timeline, and it deserves more attention than it gets. Posters had learned that one-line paragraphs forced the truncation, that truncation forced a click, and that the click was read as value. LinkedIn watched the behaviour, understood it as an artefact, and deleted the signal. Every durable tactic in this category has the same shape and the same ending. The signal

that pays is the one that is expensive to fake, which currently means holding a stranger's attention for thirty seconds. ^[7]

The four surfaces, ranked by your chance of being seen

The request behind this report was for a ranking, so here is one, with the reasoning attached. The ordering is not by engagement rate or by how much LinkedIn is promoting the surface. It is by how much of the outcome is decided by a ranking model you do not control.

Rank and surface	How distribution is decided	What it takes	Where it fails
1. Newsletter	Delivery, not ranking. Publishing an edition notifies every subscriber and emails them. Editions also get static URLs that search engines index.	A subscriber base, which you can only build from the other surfaces. Slow to start and awkward to abandon once started.	Cadence discipline. A newsletter that misses its interval trains subscribers to ignore the notification, and there is no algorithm to blame.
2. Personal post: document or text	Ranked, but from the strongest starting position on the platform. Semantic retrieval against your profile, then sequence ranking against reader history.	Consistency inside one subject, and writing that holds attention past the third line. Three to five posts a week is the range the large samples converge on.	Subject drift. The retrieval stage is matching your post against a profile, so posting outside your stated expertise weakens the match before ranking even begins.
3. Native video	Ranked, in an oversupplied lane. Dwell is normalised within content type, so you compete against other video, and video volume grew for three straight quarters.	Production effort well above a text post, for a per-post view count that fell across every follower band in the most recent benchmark.	The economics. It is the only surface where the platform's promotion and the measured per-post outcome point in opposite directions.
4. Company page	Ranked from the weakest position on the platform. Practitioner measurement consistently puts page organic reach far below personal profiles for equivalent content.	The same effort as a personal post for a fraction of the organic result.	Being treated as a reach channel at all. It is an advertising, hiring, and credibility asset, and it is good at those.

Why the newsletter goes first

ASSESSMENT Every mechanism described in this report, the retrieval embedding, the sequence model, the six predicted actions, the normalised dwell, exists to decide whether a reader sees something. A newsletter edition sidesteps the entire apparatus for the people who have already subscribed. They get a notification and an email. That is a categorically different kind of distribution, and it is the only one on LinkedIn that a ranking change cannot take away from you overnight. ^{[1] [2]}

ASSESSMENT The second reason is durability. LinkedIn confirmed in 2024 that it may resurface a useful post for months or years, which is real but not something you can plan around. A newsletter edition has a permanent URL that search engines index, so it accumulates readers on a timescale the feed does not operate on. Of the four surfaces, it is the only one that behaves like an owned asset rather than a rented one. ^[7]

ASSESSMENT The catch is that a newsletter is downstream of everything else. You cannot start with subscribers. This is why the correct read of the ranking is not "abandon posts for newsletters" but a sequence: posts acquire the audience, the newsletter keeps it. A poster with no subscribers should treat item two as item one until the list exists.

Combinations that are worth the effort, and one that is not

ASSESSMENT Post to newsletter is the combination that compounds, and it is the only one this report recommends without qualification. Publish the edition, then write a separate feed post that makes the argument in its own right and mentions the newsletter at the end. The post gets ranked on its own merits, the edition gets delivered to subscribers, and the two audiences overlap only partially.

ASSESSMENT Employee amplification is worth it for company pages, and it is the only thing that reliably is. If personal profiles out-distribute pages by the margin practitioner measurement reports, then a page post reshared by ten employees is a different distribution event from a page post standing alone. The mechanism is not mysterious: ten personal profiles enter retrieval where one page did. ^[11]


ASSESSMENT Cross-posting the same asset to multiple surfaces is not worth it, and the dwell normalisation is why. Because each content type is scored against a daily-recomputed percentile of its own kind, the identical idea posted as text, as a document, and as a video is not three chances at the same audience. It is three separate contests, each against the strongest work in that format, run by someone who split their effort three ways. ^[6]

ASSESSMENT One combination to avoid outright: video with a link in the post body. It stacks the two weakest positions in the measured data, an oversupplied format and a link, and there is no evidence that either compensates for the other. ^{[11] [15]}


Formats, measured three ways

Three independent samples exist at meaningful scale. They agree on the ordering and disagree sharply on the magnitude, which is exactly what you would expect from different denominators and different self-selected populations. Presenting them together is more honest than picking the flattering one.

FIGURE 2 Engagement rate by post format, business pages, 2024 against 2025



THIRD-PARTY ESTIMATE Socialinsider benchmark of 1.3 million posts across 16,645 business pages, January 2024 to December 2025. Business pages only, so this does not describe personal profiles. Engagement rate here is engagements over followers, not over impressions, which inflates every figure relative to a reach-based denominator. No control for who posts which format. ^[11]


FIGURE 3 Median engagement rate by format, a second and much larger sample

THIRD-PARTY ESTIMATE Buffer's 2026 cross-platform format analysis, drawn from a corpus of more than 45 million posts scheduled through Buffer. Sample is customers of one scheduling tool, so it skews toward deliberate publishers. The ordering matches the Socialinsider result; the magnitude is roughly three times larger, which is a denominator difference, not a contradiction. Buffer itself notes that reach and engagement often move in opposite directions. ^{[11] [12]}

ASSESSMENT The two charts agree that documents lead and links trail, and that is about all they agree on. A poster should take the ordering and discard the numbers. Anyone quoting 21.77 percent as a reasonable expectation for their own carousel is quoting the median of a filtered population against a follower-based denominator, which is not a forecast of anything. ^{[11] [12]}

Reach and engagement point different directions

FIGURE 4 Average impressions per post by format, at two audience sizes




THIRD-PARTY ESTIMATE Socialinsider, same 1.3 million post sample. Impressions per post, business pages. Two observations survive the caveats: links are last at both sizes by a wide margin, and the format that leads on impressions changes with audience size, so a single ranked list of formats is a fiction. Business pages only. ^[11]

ASSESSMENT Documents lead impressions at small audience sizes and lose that lead by the ten thousand follower mark, where multi-image posts take it. The plain reading is that the format advice a poster needs at 2,000 followers is not the advice they need at 30,000, and almost no published guidance makes that distinction. Links sit last in every band, which is the strongest format signal in the dataset and the one most consistent with the independent link measurements discussed below. ^[11]

The video contradiction

FIGURE 5 Average video views per post fell in every audience band



THIRD-PARTY ESTIMATE Socialinsider, business pages, year over year. Declines of 18, 39, 42, 53 and 43 percent respectively, 36 percent overall. Set this against Microsoft's report of platform video watch time up 36 percent over a comparable period. The two are reconcilable if the number of videos posted grew faster than watch time did, which is consistent with LinkedIn reporting three consecutive quarters of double-digit growth in video uploads. ^{[11] [8]}

ASSESSMENT This is the clearest case in the report of a vendor statistic and a practitioner statistic being simultaneously true and pointing opposite ways for a poster. Microsoft is reporting demand in aggregate, which is a genuine and growing number. Socialinsider is reporting supply-adjusted outcome per post, which is the number that determines whether your video was worth making. When a platform tells you a format is growing, check whether it is telling you about the audience or about the competition. ^{[8] [11]}

VENDOR CLAIM LinkedIn's reported video figures, stated by Ryan Roslansky in October 2025: watch time up 36 percent year over year, short-form video creation growing at twice the rate of other formats, and video posts "shared 20x more than any other content type." The 20x share multiple is self-reported, unaudited, and given without a base, so it establishes that video is shared more and quantifies nothing. A multiple with no denominator is not a measurement. ^[8]

The tactics, ruled on one at a time

This section exists because the advice market for LinkedIn is larger and older than the evidence base, and a poster following 2021 guidance in 2026 is actively harming their own distribution. Each ruling below names what it rests on. Where the answer is that nobody knows, that is the ruling.

Tactic	Ruling	What the evidence actually says
Engagement pods	Stop	LinkedIn's stated position is that pods violate the user agreement, and enforcement in 2026 extended past individual accounts to the vendors building the tools. Beyond the policy risk, pods target the three active signals while the model also scores click, skip, and normalised dwell. A burst of reciprocal likes from people who did not read the post is a weak input to half the objective function.
"Comment YES if you agree"	Stop	Trades an active signal you can manufacture for a passive one you cannot. The specific downranking percentages circulating for engagement bait trace to marketing blogs with no methodology and should not be repeated; the mechanism argument stands on the published objective function without needing them.
Link in the first comment	Marginal	Never established as effective by any controlled measurement, and widely reported as patched. LinkedIn's stated position is that there is no intentional link penalty in the first place. If you are moving the link to avoid a penalty the platform denies exists, using a workaround nobody has measured, you are two unverified steps from the evidence.
One-line paragraphs to force "read more"	Dead since 2024	LinkedIn confirmed in February 2024 that it stopped using "read more" clicks as a value signal once it understood the behaviour. The formatting persists as a convention; the mechanical advantage that created it is gone.
Hashtags	Neutral to negative	LinkedIn's own editorial staff describe hashtags as "a nice to have, not a need to have." Practitioner samples find no reach benefit. Bait tags such as #follow and #like are associated with low-quality classification. The rational move is zero to three, chosen for topical accuracy rather than reach.
The golden hour	Contested	Widely asserted, and one 300,000 post practitioner sample flatly denies it, with the analyst stating "there has never been a golden hour." No LinkedIn document describes a time-boxed early evaluation window. What is documented is that skip is a modeled negative, which makes early impressions consequential without requiring a fixed window.
Posting at an optimal hour	Weak effect, real but small	Buffer's 4.8 million post analysis finds midweek afternoons ahead of other slots, and says so with an explicit caveat: "posting at 4 p.m. on a Wednesday won't rescue a post that isn't relevant." LinkedIn's editor in chief has told members not to chase timing at all. Treat it as a tiebreaker.
Posting daily	Diminishing and risky	Large samples converge on three to five posts a week for most accounts. The published gains per additional post are real but come from tool-customer samples with no quality control, and posting more while writing less is the failure mode the current ranker punishes hardest.
Polls	Reach yes, audience no	Polls score well on impressions in the Socialinsider sample, particularly at large follower counts, while sitting at the bottom on engagement rate growth. Practitioner reporting describes them as strong for reach and weak for follower growth. A poll buys impressions from people who will not remember you.

Tactic	Ruling	What the evidence actually says
Tagging people	Use sparingly and accurately	LinkedIn's own editorial staff endorse tagging when the tagged person is genuinely relevant. Mass tagging is the same class of manufactured active signal as a pod, and is the most visible one.
Delete and repost	No evidence either way	No published measurement exists. Asserted constantly, tested never. Included here to note the absence rather than to rule on it.
AI-written drafts	High and rising risk	Independent detection puts likely-AI at 81.2 percent of long-form posts; LinkedIn's own classifiers put fully-AI at roughly 40 percent of long-form. LinkedIn shipped a member-facing report button in July 2026 and says the reports tune its models. Whatever the distribution effect is today, the platform has publicly committed to making it negative.
Third-party automation tools	Stop	LinkedIn removed a vendor's company page and banned its founder and two executives in 2026, and says it blocks hundreds of thousands of automated comment attempts a day. The account you are optimising is the asset at risk.

Pods, and why enforcement is the smaller problem

ASSESSMENT The usual argument against engagement pods is that you will get caught, and the enforcement record in 2026 supports it: LinkedIn is reported to have removed an automation vendor's company page and banned its founder and two executives, and the company says it blocks hundreds of thousands of automated comment attempts a day. That reporting is second-hand and should be treated as such. The stronger argument needs none of it. Pods manufacture likes and comments, which is three of the six trained objectives, from people who did not read the post, which damages the other three. A pod is a machine for producing high reaction counts alongside short dwell, and short dwell is the thing the ranker has been getting better at pricing since 2020. [\[23\]](#) [\[10\]](#) [\[6\]](#)

Timing, and the golden hour that nobody can find


ASSESSMENT The golden hour is the most confidently asserted claim in this category with the least behind it. No LinkedIn engineering document describes a fixed early evaluation window. One practitioner analyst working from a 300,000 post sample states flatly that there has never been one. What the published architecture does support is weaker and different: skip is a trained negative, so impressions that go nowhere cost you regardless of when they happen, and there is no documented clock attached to that. [\[18\]](#) [\[1\]](#)

THIRD-PARTY ESTIMATE On timing generally, the largest sample available finds midweek afternoons ahead of other slots across 4.8 million posts, and its authors attach the caveat themselves: posting at four on a Wednesday will not rescue a post that is not relevant. LinkedIn's own editor in chief has told members not to chase posting times at all. The effect appears to be real and small, which makes it a tiebreaker rather than a strategy, and it is measured on scheduling-tool customers whose audiences may not be yours. [\[13\]](#) [\[7\]](#)

The link question, treated properly

VENDOR CLAIM LinkedIn's Senior Director of Product Management, Rishi Jobanputra, stated publicly in September 2025 that LinkedIn does not intentionally limit the reach of posts because they contain an external link, with the qualification that the post should carry value on its own. This is a first-party statement about intent. It is not a measurement of outcome, and the word "intentionally" is carrying weight: a model trained on dwell and skip could depress link posts without anyone intending it. ^[19]

FIGURE 6 Three answers to "how much does a link cost you", and only one has a method




ASSESSMENT Reported reduction in reach from one external link in the post body. The first is a statement of intent, not a measurement. The second comes from a 1.3 million post practitioner sample and is reported second-hand here because the underlying report is not public. The third appears in a July 2026 Forbes column; following its citation leads to a marketing blog with no sample size, no date and no method. The author's working estimate is the middle figure, held loosely. ^{[19] [15] [17]}

ASSESSMENT The 60 percent figure is the most instructive thing in this section, and not because it is wrong. It is unfalsifiable. It appears in a reputable outlet, which lends it authority, and it sources to a page that publishes no method, which means no one can check it or correct it. It has since been repeated across the advice market as though Forbes measured it. Any conclusion in this report that rested on that number would be worthless, so none does. ^[17]

ASSESSMENT The practical ruling, given a first-party denial and one methodologically opaque outside estimate near 19 percent, is that a link in the body probably costs you something in the mid-to-high teens as a percentage, and that this is a price worth paying when the click is the point. What is not defensible is the surrounding folklore: elaborate comment-workaround rituals justified by a 60 percent figure that traces to nothing. ^{[19] [15]}

Why AI text became the binding constraint

Of everything in this report, this is the change least reflected in current posting advice and most likely to determine outcomes over the next year.

FIGURE 7 Two measurements of how much LinkedIn text is machine-written, July 2026

THIRD-PARTY ESTIMATE The first is Originality.ai's classification of 5,000 public posts of 100 words or more, using its own detector at a 15 percent AI allowance; the company sells AI detection, and the study reports no false positive rate. The second and third are LinkedIn's internal classifier figures for fully AI-generated posts, as relayed by its chief product officer in press coverage; the classifier is undocumented. The two measure different thresholds and neither is audited. The gap between them is the honest finding. [\[16\]](#) [\[9\]](#) [\[10\]](#)

FACT On 30 July 2026 LinkedIn added a "Seems like AI slop" option to the post menu, notifies the author privately in their analytics when a post is flagged, and retired its "enhance your post" AI writing tool in favour of a proofreader that preserves the author's voice. Chief product officer Hari Srinivasan described the reports as a way to "tune our models and make better feeds," which is a statement that member flags become training data rather than only complaints. [\[9\]](#) [\[10\]](#) [\[22\]](#)

FACT LinkedIn says it blocks hundreds of thousands of automated comment attempts daily, and characterised the recent total as billions of automation attempts prevented over a period of months. [\[10\]](#)

ASSESSMENT Put the three facts together and the strategic position is clear. LinkedIn has a supply glut of competent, generic prose, no reliable way to price it, and a stated intent to build one out of member reports. A poster's defensible position is therefore whatever a language model cannot produce from a prompt: a specific number from your own operation, a client situation with the details intact, a position you would have to defend at a conference. Not because authenticity is virtuous, but because generic professional prose is now the most abundant commodity on the platform and abundance sets price. [\[16\]](#) [\[9\]](#)

ASSESSMENT The second-order effect is already visible in comments. Van der Blom reports that 80 percent of the comments arriving on his posts in the first five minutes are AI-written. If that generalises, early comment velocity is degrading as a distribution signal at the same time the platform is actively suppressing the automation that produces it. Anyone whose strategy depends on early comment volume is standing on a signal being removed from underneath them from two directions at once. [\[15\]](#) [\[10\]](#)

Failure modes, ranked by what they cost you

Ordered by the author's estimate of how much reach each one destroys, not by how often it gets complained about. The first two account for more lost distribution than everything below them combined.

#	Failure	Why it costs what it costs	The fix
1	Posting outside a recognisable subject	Retrieval matches a post embedding against a profile embedding. Subject drift weakens the match before ranking begins, so the post is not merely ranked lower, it is a weaker candidate for selection at all.	Pick one subject you can hold for a year. Make the profile headline and the posting history say the same thing. Judge drift by whether a stranger could predict your next post's topic.
2	Writing prose a classifier cannot distinguish from generated text	The most abundant input on the platform, with a member-reported flag now feeding LinkedIn's detection models and a stated intent to demote it.	One number, one name, or one situation per post that could only come from your own work. If the post would be identical coming from any competent person in your field, it is not worth the slot.
3	Optimising for likes	Half the trained objective is click, skip, and long-dwell. Reactions are the cheapest signal and the one most contaminated by automation.	Write for the reader who finishes. Length is fine if it holds; a long post that is skipped is worse than a short one, because skip is a modeled negative.
4	Running a company page as the primary organic channel	Practitioner measurement puts page organic reach far below personal profiles for equivalent content, and the gap widened through 2024 and 2025.	Move the organic programme to named humans. Keep the page for advertising, hiring, and the credibility check buyers run before a call.
5	Manufacturing early engagement through pods or bait	Targets three of six signals, is against the user agreement, and is under active enforcement that now reaches tool vendors as well as users.	Nothing replaces it. That is the point. Budget the time into writing instead.
6	Putting a link in the body reflexively	A cost in the mid-to-high teens as a percentage on a defensible reading of the one methodologically serious outside estimate.	Include the link when the click is the objective and accept the price. Omit it when it is decoration. Skip the comment rituals.
7	Chasing the optimal posting time	Measurable but small, and the largest timing study says so itself.	Post when you can sustain three to five a week. Use timing as a tiebreaker between two equally ready posts, never as a reason to delay a good one.

Why point forecasts fail here

ASSESSMENT The subject's own rate of change is the argument. In roughly six years LinkedIn has shipped a dwell-based skip model, replaced it with a normalised long-dwell model, replaced a multi-source retrieval architecture with a single semantic encoder, and replaced a DCNv2 ranker with a sequential transformer. Along the way it published a 150 billion parameter foundation model, withdrew the paper seven months later, and separately documented building an LLM ranker that it declined to ship. Anyone offering a numeric forecast of LinkedIn reach for 2028 is forecasting the output of a system that has been rebuilt three times and abandoned two serious alternatives inside the observation window. [\[1\]](#) [\[2\]](#) [\[3\]](#) [\[4\]](#) [\[6\]](#)

ASSESSMENT There is a second problem specific to this subject. The measurement layer is being removed. LinkedIn has restricted third-party tooling and pursued the analysts who built independent datasets, so the practitioner samples this report leans on are getting harder to produce, not easier. A forecast made on a shrinking evidence base ages badly in a way that is invisible to the reader. [\[15\]](#) [\[10\]](#)

Scenarios to mid-2028

Four structured futures with subjective probabilities. The probabilities are the least reliable content in this document and exist only to force the estimates to sum. The earliest visible signs are the useful part, because any reader can watch for them without access to anything private.

SCENARIO A Semantic ranking holds and authorship hardens 40 percent

Feed SR or a successor stays in production, the AI-slop classifier matures into a real ranking input, and verification becomes a meaningful distribution factor.

Consequence. Distribution consolidates toward a smaller set of identifiable specialists. Generic professional prose stops working entirely rather than working less well.

Earliest visible sign. LinkedIn publishes a follow-up engineering post describing the slop classifier as a ranking feature rather than a reporting tool, or verification appears in a ranking context.

SCENARIO B The LLM ranker finally ships 25 percent

The approach LinkedIn documented as underperforming gets there on a later attempt, and ranking moves from sequence modeling to semantic judgement of the post itself.

Consequence. Topical fit becomes close to the whole game. Format advice becomes irrelevant. Reach for off-topic posts collapses toward zero rather than merely declining.

Earliest visible sign. A LinkedIn engineering publication reporting an LLM ranker beating the production baseline in online tests, which is the exact result the February 2026 paper says they did not get.

SCENARIO C **Feed distribution keeps shrinking and subscription surfaces absorb it** 25 percent

Organic feed reach continues the decline practitioner samples report, while newsletters and notifications carry a growing share of what people actually read.

Consequence. The ranking in this report becomes more emphatic rather than less. Posting becomes primarily a subscriber-acquisition activity.

Earliest visible sign. LinkedIn adding newsletter discovery or recommendation surfaces, or promoting newsletter metrics in earnings commentary the way it has promoted video.

SCENARIO D **Measurement reopens** 10 percent

LinkedIn exposes comparative reach data to members, or settles into a stable relationship with third-party analytics rather than restricting it.

Consequence. Most of the advice market becomes checkable and most of it does not survive. This report's third-party estimates would be replaced by better numbers, possibly different ones.

Earliest visible sign. A member-facing benchmark in LinkedIn analytics showing how a post performed against comparable posts, or a public API for post-level reach.

Leading indicators

Observable events, each with what it would tell you and which scenario it supports. None requires private information.

If this happens	It means	Moves toward	Where to watch
An engineering post describes the AI-slop signal as a ranking feature	Member reports have become training data for distribution, not just moderation	Scenario A	LinkedIn Engineering Blog, feed category
A LinkedIn paper reports an LLM ranker winning online	The February 2026 negative result has been overturned	Scenario B	arXiv, LinkedIn-authored recommender papers
Newsletter metrics appear in Microsoft earnings commentary	LinkedIn is managing newsletters as a strategic surface, as it did with video from 2024	Scenario C	Microsoft quarterly earnings call transcripts
Member analytics gain a comparative benchmark	The measurement asymmetry this report describes is closing	Scenario D	Your own post analytics panel
Video watch time growth falls below upload growth in LinkedIn's own reporting	The oversupply reading of the video contradiction is confirmed from the platform side	Scenario C	Microsoft earnings, LinkedIn segment remarks

If this happens	It means	Moves toward	Where to watch
Another practitioner dataset of comparable scale is published or shut down	The independent evidence base is either recovering or finishing its collapse	Scenario D reversed	Practitioner reports, tool vendor blogs

Implications

For people who use LinkedIn mainly to read

ASSESSMENT Your feed in 2026 is assembled from up to a thousand of your own past impressions read in order, which means it is a model of your recent behaviour rather than of your interests. That distinction matters practically. A week of idly watching video teaches the model you want video. The most effective thing a reader can do is treat skipping as an instruction rather than a non-event, because it is one: skip is a trained negative in the ranker. ^[1]

ASSESSMENT Dwell time is measured whether or not you interact, so your reading behaviour is shaping your feed continuously and invisibly. If your feed has degraded, the honest diagnosis is usually that you have been reading things you would not have chosen. Subscribing to newsletters from people you actually want to hear from routes around the ranker entirely and is the only reliable lever a reader has. ^{[6] [4]}

ASSESSMENT One consumption-side warning. With likely-AI text somewhere between 40 and 81 percent of long-form posts depending on whose classifier you believe, the base rate for a confident, well-structured, sourceless post being machine-assembled is now high enough that it should change how you read. Specificity is the cheapest available filter: a post with a real number, a named situation, or a checkable claim was probably written by someone who knew something. ^{[16] [9]}

For people who post business topics

ASSESSMENT The operating recommendation, stated once. Pick one subject and hold it for a year, because retrieval matches your post against your profile before ranking begins. Post three to five times a week, native, in text or documents. Put something in every post that could only come from your own work. Start a newsletter as soon as you have an audience worth mailing, and treat feed posting as the acquisition channel that feeds it. Include links when the click is the point and accept a cost in the mid-to-high teens as a percentage. Skip the pods, the bait, the hashtag stacks, and the comment rituals. ^{[2] [1] [14] [19]}

ASSESSMENT What to stop measuring: likes. Half the trained objective consists of click, skip, and normalised dwell, and reactions are both the cheapest signal to produce and the one most contaminated by automation. If your analytics show high impressions and low dwell, the ranker is already discounting you and the reaction count is hiding it. Watch the ratio of impressions to profile visits instead, which is the closest proxy an individual has to whether anyone finished reading. ^{[1] [6]}

ASSESSMENT If you run a company page, run it as an advertising, hiring, and credibility asset rather than a reach asset, and move the organic programme to named humans with employee amplification behind it. The measured gap between personal and page organic reach is large enough in every practitioner sample that no amount of page-side craft closes it. This is a structural position, not a content problem. ^[11]

ASSESSMENT The video decision comes down to why you are making it. If the goal is reach per unit of effort, the measured case is weak: per-post views fell in every follower band in the most recent benchmark while production cost stayed high. If the goal is that some arguments only work on camera, make it anyway and judge it on completion rather than views. What does not work is making video because the platform said video is growing, which is a statement about the size of your competition as much as your audience. ^{[11] [8]}

ASSESSMENT One thing worth doing that almost nobody does: run the experiment on yourself. Hold format, length, and cadence constant for a month and vary only subject. Then hold subject constant and vary only format. Your own posts are the only sample in existence that has no self-selection problem for you, and thirty posts is enough to see whether topic fit or format is driving your variance. Every number in this report is someone else's population.

What this document cannot answer

None of the reach and engagement figures here come from LinkedIn. They come from tool vendors measuring their own customers, and no one has established how those customers differ from the general posting population. Answering it properly needs a sample drawn independently of any tool, which currently no one outside LinkedIn can construct.

Whether format effects are causal is unresolved and probably unresolvable from public data. Every study compares formats without controlling for the poster. A randomised design, where the same authors publish the same arguments in randomly assigned formats, would settle it and has never been run at scale.

How much a link actually costs is a range, not a number. A first-party denial of intent and a single opaque outside estimate near 19 percent is what the public record supports. Nothing here establishes whether the effect is a deliberate rule, an emergent consequence of dwell modeling, or a confound with the kind of post that carries links.

The distribution effect of AI-written text is unknown today. Two undocumented classifiers disagree by a factor of two on how much of it there is, and neither LinkedIn nor anyone else has published what happens to a post's reach once it is classified. The platform has stated an intent; it has not published an effect.

Whether ranking behaves the same outside English-language markets is unaddressed in every source used here. LinkedIn's own paper shows the ranker performing differently by account maturity, so uniformity across markets should not be assumed.

The vertical video feed's separate ranking is undocumented. LinkedIn publishes on the main feed. Whether the vertical surface uses the same retrieval and ranking stack, or something adapted from short-form video conventions, does not appear in any source located for this report.

Half-life of this assessment

Six months. Every format engagement rate, every impressions-per-post figure, and the video decline numbers. These are single-year snapshots from samples that are themselves getting harder to collect, and the video figures in particular describe a format under active platform promotion, which is the condition least likely to persist. The AI-prevalence figures move fastest of all: they went from roughly half to 81 percent of long-form in about eighteen months.

Twelve months. The surface ranking, the tactical rulings, and the link estimate. These rest on mechanism as well as measurement, which slows their decay, but a shipped LLM ranker or a maturing slop classifier would rewrite the tactical table wholesale. The specific recommendation to treat newsletters as the strongest surface would survive most changes to feed ranking and would not survive LinkedIn changing how newsletter delivery works.

Twenty-four months. The timeline and the pattern read off it, that LinkedIn keeps replacing cheap gameable proxies with expensive ones. Individual rows will be superseded; the direction has held through three architectural generations.

Architectural, and unlikely to decay. Three things. Dwell and skip are modeled objectives, not incidental metrics, and no version of this system is going back to counting reactions. Retrieval matches a post against a profile, which is why subject coherence is structural rather than tactical. And delivered surfaces such as newsletters are categorically different from ranked ones, which is true of every platform that has both and will outlast any particular ranker.

Limitations

This report was produced independently. The author has no commercial relationship with LinkedIn or Microsoft, no relationship with any of the analytics vendors whose data appears here, and no financial interest in any of the tools mentioned. No party reviewed the text before publication.

The evidence base is lopsided in a way worth stating plainly. Mechanism claims come from LinkedIn itself, which has every incentive to describe its systems favourably but publishes into venues where engineers can check the architecture. Outcome claims come exclusively from parties selling social media tooling, measuring their own customers, using undisclosed denominators. There is no independent audited measurement of LinkedIn organic reach anywhere in the public record, and this report is not one.

Two figures used here are reported second-hand. Van der Blom's 18.8 percent link estimate and his 1.3 million post sample size come from public summaries and an interview rather than from the report itself, which is not publicly available. They are used for direction and flagged wherever

they appear. The 60 percent link figure is included specifically as an example of a statistic that cannot be checked, and no conclusion rests on it.

The 360Brew paper was withdrawn from arXiv in August 2025 for a licensing reason. It is cited here as evidence of what LinkedIn was building and publishing at that date, which the withdrawal does not affect, and not as evidence of what is running now, which the Feed SR paper addresses directly.

The Socialinsider benchmark covers business pages only. It is the largest format dataset located with a stated methodology, and it is used throughout this report for format comparisons that most readers will apply to personal profiles. That extension is the author's, it is not supported by the source, and it is the single weakest inferential step in the document.

Sources

- [1] Hertel, L., Srivastava, G., Naqvi, S. A., et al. "An Industrial-Scale Sequential Recommender for LinkedIn Feed Ranking." LinkedIn Inc., arXiv:2602.12354, 12 February 2026. *Vendor primary* [arxiv.org](https://arxiv.org/abs/2602.12354)
- [2] Danchev, H. "Engineering the next generation of LinkedIn's Feed." LinkedIn Engineering Blog, 12 March 2026. *Vendor primary* [linkedin.com](https://engineering.linkedin.com)
- [3] Firooz, H., et al. "360Brew: A Decoder-only Foundation Model for Personalized Ranking and Recommendation." LinkedIn Foundation AI Technologies, arXiv:2501.16450, 27 January 2025. Withdrawn 23 August 2025 for a licensing reason. *Vendor primary* [arxiv.org](https://arxiv.org/abs/2501.16450)
- [4] Dangi, S., Jia, J., Somaiya, M., Xuan, Y. "Understanding dwell time to improve LinkedIn feed ranking." LinkedIn Engineering Blog, 12 May 2020. *Vendor primary* [linkedin.com](https://engineering.linkedin.com)
- [5] "LinkedIn rebuilds its feed from scratch with LLMs and GPU-powered ranking." PPC Land, March 2026. Independent write-up of the LinkedIn engineering post, useful mainly because it separates the retrieval stage from the ranking stage, which most coverage conflates. *Independent press* ppc.land
- [6] Zhang, F., Kothari, M., Tiwana, B. S. "Leveraging Dwell Time to Improve Member Experiences on the LinkedIn Feed." LinkedIn Engineering Blog, 1 October 2024. *Vendor primary* [linkedin.com](https://engineering.linkedin.com)
- [7] Hutchinson, A. "LinkedIn Shares Insights Into Its Latest Feed Algorithm Updates." Social Media Today, 27 February 2024. Quotes Dan Roth, Editor in Chief, and Tim Jurka, Senior Director of Engineering. *Independent press* socialmediatoday.com
- [8] Hutchinson, A. "LinkedIn Reports Significant Increases in Post Comments and Video Posts." Social Media Today, 30 October 2025. Reports Microsoft Q1 FY2026 remarks and figures attributed to LinkedIn CEO Ryan Roslansky. *Independent press* socialmediatoday.com
- [9] "LinkedIn adds a button to report AI-generated 'slop'." TechCrunch, 30 July 2026. Quotes Hari Srinivasan, LinkedIn Chief Product Officer. *Independent press* techcrunch.com
- [10] "LinkedIn adds a 'seems like AI slop' button after blocking billions of automated comment attempts in the last few months." Fortune, 31 July 2026. *Independent press* fortune.com
- [11] Socialinsider. "LinkedIn Organic Benchmarks 2026." 1.3 million posts across 16,645 business pages, January 2024 to December 2025. *Practitioner aggregate (self-selected)* socialinsider.io
- [12] Buffer. "Best Content Format on Social Platforms in 2026: 45M+ Posts Analyzed." *Practitioner aggregate (self-selected)* buffer.com
- [13] Buffer. "Best Time to Post on LinkedIn in 2026: 4.8M Posts Analyzed." 22 July 2026. *Practitioner aggregate (self-selected)* buffer.com
- [14] Buffer. "How Often Should You Post on LinkedIn in 2026? Data From 2 Million+ Posts." *Practitioner aggregate (self-selected)* buffer.com
- [15] van der Blom, R. "The state of LinkedIn in 2026," Creator Science podcast episode 307. Reports the Algorithm Insights 2026 findings, drawn from 1.3 million posts by 50,000 creators. The underlying report is not public; figures used here are reported second-hand. *Practitioner aggregate (self-selected)* creatorscience.com
- [16] Originality.AI. "LinkedIn AI Content Study: 81% of Long-Form Posts Are Likely AI." July 2026. 5,000 public posts of 100 words or more, classified with the vendor's own detector at a 15 percent AI allowance. No false positive rate reported. *Vendor-published comparison (commercial interest)* originality.ai

- [17] Cook, J. "The LinkedIn Link Penalty Cutting Your Reach By 60%." Forbes, 30 July 2026. The 60 percent figure is sourced in the article to a marketing blog that publishes no sample size, date, or method. *Second-hand (unverified)* forbes.com
- [18] Cook, J. "The LinkedIn Algorithm Changed Again. Here's What's New For 2026." Forbes, 12 January 2026. Reports practitioner findings from a 300,000 post sample, including the claim that no golden hour exists. *Independent press* forbes.com
- [19] Jobanputra, R., Senior Director of Product Management, LinkedIn. Public statement on external links and post reach, September 2025. *Vendor primary* linkedin.com
- [20] Bishop, T. "LinkedIn tops \$5B in quarterly revenue for the first time, and its TikTok pivot is paying off." GeekWire, January 2026. *Independent press* geekwire.com
- [21] Hutchinson, A. "LinkedIn Removes Its Top Voice Badges for Collaborative Articles." Social Media Today. Collaborative Articles moved to read-only in October 2025. *Independent press* socialmediatoday.com
- [22] "LinkedIn Is Cracking Down on 'AI Slop' With New Reporting Tool." Bloomberg, 30 July 2026. *Independent press* bloomberg.com
- [23] ConnectSafely. "LinkedIn Engagement Pods Crackdown 2026." Relays a statement attributed to Gyanda Sachdeva, LinkedIn VP of Product Management, on third-party manipulation tools, and reports vendor-level enforcement. Not independently verified. *Second-hand (unverified)* connectsafely.ai

Rights and permitted use

© 2026 David Soden, davidsoden.com/market-assessment. All rights reserved.

This document, including its structure, analysis, findings, evidence classification, and framing, is the original work of David Soden, published at davidsoden.com/market-assessment.

It is published for public reading. You are free to share the complete, unmodified file, to link to it, and to quote short passages in commentary, reporting, or discussion, provided the author and the site above are credited.

The following require prior written consent: republishing this work in whole or in substantial part under another name, masthead, or brand; excerpting or modifying it in a way that alters the analysis or removes the evidence classifications; incorporating any portion of it into a paid, commissioned, or client-facing deliverable; resale or distribution behind a paywall; and use of this document as training data for machine learning models.

Commissioned Market Assessment reports, commercial licensing, and briefings on this material are available.

This document is analysis, not investment, legal, or procurement advice. It was not commissioned, reviewed, or funded by any company named in it, and the author holds no position in any of them.

Market Assessment reports, commissioned research, and briefings: davidsoden.com/market-assessment