



MAKING INVISIBLE EXPERTS VISIBLE

THE BRIDGE · AI SEARCH SERIES · EDITION ONE

The Invisible Expert's Field Guide to **AI Search**

*How one person gets found by ChatGPT, Gemini, Copilot
and Perplexity on LinkedIn in 2026, without a marketing
department, a task force, or a single line of automation.*

WRITTEN BY KENNETH NEL
FOUNDER, THE DIGITAL COLLECTIVE

2026 EDITION · 21 PAGES
TDC101.COM

PREFACE

The 36-page problem

A very good guide was published this year. It was written for someone you are not.

LinkedIn published a 36-page guide in 2026 on how to get your content quoted by AI. Thorough, well researched, worth reading. Also written for a marketing department.

It opens by telling you to form a cross-functional task force. It assumes you have a PR team, an editorial calendar, a measurement budget, and nine people willing to post about the same topic for three months.

You are one person. Maybe two if you count whoever does your invoicing.

Here is the awkward bit. The guide's own research points somewhere else entirely. Three quarters of everything AI quotes from LinkedIn comes from individual profiles rather than company pages. The

biggest group of cited people? Ordinary members with fewer than ten thousand followers.

The playbook was written for departments. The evidence inside it points at individuals.

This is the version for individuals. Deliberately shorter. Everything in here can be done by one person in about an hour a day, with nothing beyond the LinkedIn account you already have.

We kept the numbers because the research is good and it is not ours to rewrite. Every figure is credited to its original source on page 20, so you can check our working and go and read the studies yourself. What we added is the part the original could not cover: what all of this looks like when the expert, the writer, the publisher and the salesperson are the same person.

HOW TO USE THIS

Read it once front to back. Takes about twenty minutes. Then work from page 18, which is the whole guide compressed into a single checklist, and page 19, which turns that checklist into a ninety-day sequence. Everything between here and there is the reasoning behind those pages.

1 hr

A day. That is the entire time commitment this guide asks for.

90

Days before you have a fair read on whether it is working.

£0

Additional tooling required. Everything here runs on a standard account.

CONTENTS

What is in here

01	The room you are not in Why the shortlist is drawn before anyone calls you	04
02	Who the machine actually quotes The citation data, and why it favours you	07
03	Your profile is the citation Fix the destination before you drive traffic to it	09
04	What to publish, and what shape Articles, posts, and the anatomy of a quotable page	11
05	The hour that does the work Why publishing alone is half a strategy	14
06	Knowing whether it worked Measurement for one person, not a task force	15
07	The part no machine can do Where visibility stops and business begins	16
WORKING PAGES		
☐	The one-person checklist	18
☐	The first ninety days	19
☐	Sources, method and how to check us	20

THE SHORT VERSION, IF YOU READ NOTHING ELSE

Be a real, verified person with one subject, a profile that says who you help, one article a week that answers a question in its title, and an hour a day spent in other people's comments. That is the whole guide. The other twenty pages explain why each of those five things is on the list.

SECTION 01

The room you are not in

Buyers now assemble a shortlist inside a chat window. That conversation happens without you, and you are never told it took place.

QUICK ANSWER

Your buyer asks a machine before they ask a person. Ninety-four per cent of B2B buyers now use generative AI somewhere in their research, and the shortlist is often drawn before you know the work exists. You are either inside that answer or you are not in the running, and nobody sends a rejection.

Somewhere this week, someone typed a question that should have had your name in the answer.

"Who are the best supply chain consultants in the UK for a mid-market manufacturer?"

Three names came back. That was the shortlist. Nobody called to say you were left off it because nobody knew you should have been on it. The model did not leave you out on purpose. It left you out because there was nothing of yours to quote.

This is what trips up experienced people. You were not beaten by a better consultant. You were beaten by a more visible one. Twenty-four years of judgement sitting in your head, in client meetings and in a folder of decks nobody outside the project will ever open? As far as an AI model is concerned, that is the same as no experience at all.

A prospect told us something recently that stuck: "People do their due diligence. They Google my name. If I am not on LinkedIn, they assume I am nobody, or a bot." She was right, and it has since got harder, because the buyer now interrogates the machine before your name is ever typed anywhere.

WHAT CHANGED, IN FOUR NUMBERS

94%

of B2B buyers use generative AI in their research.

6SENSE, 2025

50%

forecast fall in traditional search volume by 2028.

GARTNER, 2023

60%

decline in non-branded search traffic to LinkedIn's own marketing sites.

LINKEDIN INTERNAL, 2026

10.1

months average B2B buying cycle, down from 11.3 the year before.

6SENSE, 2025

Why the shorter buying cycle is the bad news, not the good news

That last number is the one most people misread. A buying cycle compressing from 11.3 months to 10.1 sounds like efficiency. For anyone selling expertise, it is the opposite. The research phase did not shrink because buyers are doing less homework. It shrank because a machine now does the homework in an afternoon.

Which means the window in which you could previously have been discovered, the long, slow, human trawl through Google, referrals, LinkedIn searches and half-remembered conference names, has largely closed. The evaluation is quicker, the vendor conversations start later, and by the time anyone speaks to a human being the field has already been narrowed.

If your expertise is not written down somewhere a model can reach it, you are not in the narrowing.

Is this just SEO with a new name?

Partly, and that should reassure you. The underlying logic has not changed: something indexes content, and content that is clear, credible and well structured gets surfaced. What has changed is the address. As long as there is an algorithm reading content, the discipline is alive, it has simply moved house, and the front door is now your LinkedIn profile rather than your website's homepage.

The difference worth understanding is how the two systems answer. A search engine ranks one best page per query and hands you a list. A model gathers many sources, weighs how consistently they agree, and writes a single answer. That has one enormous consequence for a solo operator: **you do not need to be the top result. You need to be one of several credible sources saying a consistent thing.** That is a far lower bar, and it is reachable by one person in a quarter.

Which AI are we actually talking about?

Two different systems, and conflating them causes most of the confusion in this subject.

SYSTEM	WHAT IT DECIDES	WHAT IT READS
The AI inside LinkedIn	Who sees your post in the feed. This is the algorithm everybody complains about, and it governs reach among people who already use the platform.	Your profile and posts
The AI outside LinkedIn	Whether a buyer who never opens LinkedIn at all ever hears your name. This is the one this guide is about, and almost nobody optimises for it.	Your profile and posts

Note the third column. They read the same raw material. This is the single most useful fact in this guide, because it means you are not running two strategies. One body of work, published properly, feeds both. The buyer does not know or care which system found you.

THE READING

You are not competing for a ranking any more. You are competing to be quotable. Those require different work, and the second one rewards a specialist far more than the first ever did.

THE WHOLE BUSINESS, IN ONE SENTENCE

**If I don't know you exist,
I will never do
business with you.**

That line has paid the mortgage for twenty years. Nothing about it changed in 2026 except the audience. Now the machine has to know you exist as well, and unlike a human being, it will never give you the benefit of the doubt.

KENNETH NEL · FOUNDER, THE DIGITAL COLLECTIVE

SECTION 02

Who the machine actually quotes

This is the section that should change your mind. The citation data does not favour the big brand. It favours the named human being with a narrow subject.

QUICK ANSWER

Individuals, overwhelmingly. Three quarters of LinkedIn citations in AI answers come from personal profiles rather than company pages. The largest single group of cited accounts are ordinary members with fewer than ten thousand followers who publish consistently. Job title barely matters. Evidence of expertise does.

If you run a one-person firm and you have been quietly assuming this is for companies with content teams, the next two pages are for you. The most cited source in AI search, on the platform AI cites most, is someone like you.

Where LinkedIn citations actually come from

SHARE OF AI CITATIONS BY ACCOUNT TYPE

Established members, 10k+ followers, deep domain expertise

40%

Everyday members, under 10k followers, consistent publishers

35%

This is the group you are in, or can join within a quarter.

Company pages, official brand accounts

25%

SOURCE: MELTWATER, 2026 • INDIVIDUAL PROFILES 75% VS COMPANY PAGES 25%

THE READING

Ordinary members with modest followings out-earn every company page on the platform combined. Not the influencers, the everyday accounts, in a category all of their own.

Two more numbers sharpen this. Chief executives account for 8.2 per cent of cited content and founders 7.5 per cent, which sounds impressive until you realise that means more than eight in ten citations go to people without a fancy title. The research was blunt about it: seniority mattered less than demonstrated expertise. And ninety-five per cent of all LinkedIn citations go to original posts rather than reshares, which quietly kills the most common content habit on the platform.

Resharing someone else's article with "great insights here" is not a content strategy. It is, as far as an AI model is concerned, not content at all.

Why LinkedIn in particular?

Because it is the single most-cited domain for professional questions in AI search, and because of an unusual property in how models handle what they find there.

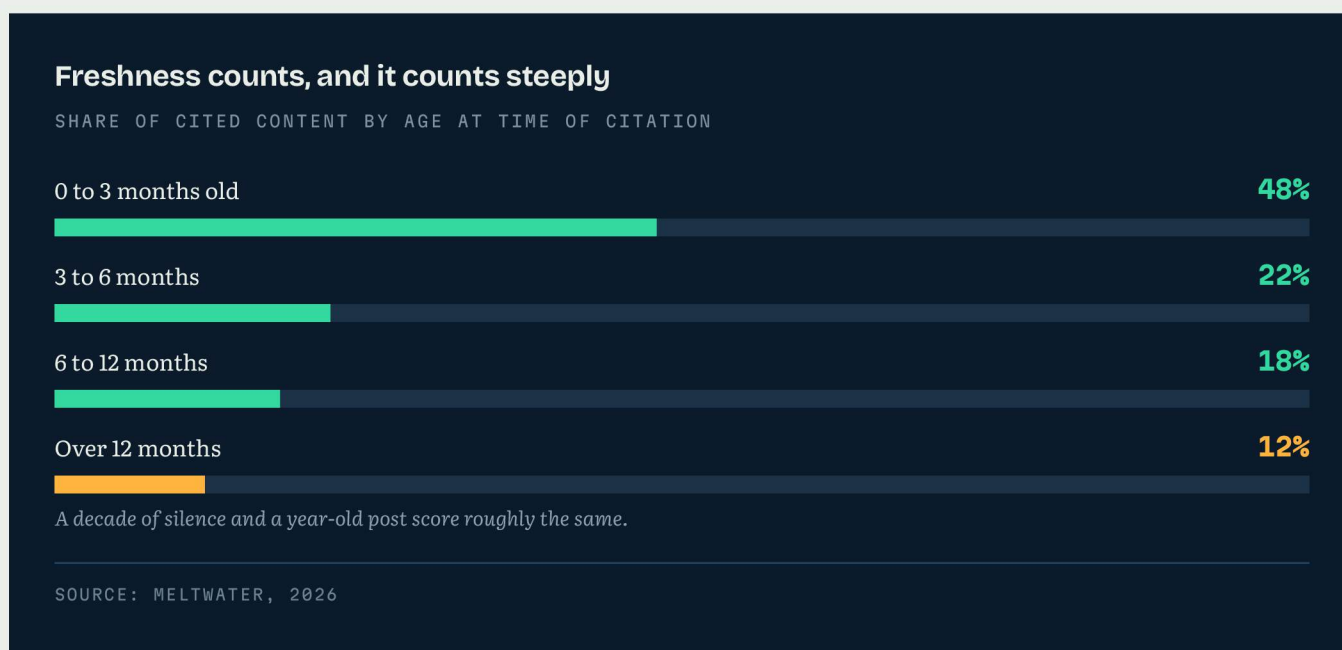
When a model quotes LinkedIn content, it tends to preserve your actual wording. Semantic similarity between the source and the generated answer sits between 0.57 and 0.60. That is high. In plain terms, when ChatGPT or Copilot pulls something of yours from LinkedIn, it usually paraphrases you closely rather than blending you into an averaged, anonymous summary. Your framing survives the trip.

There is a structural reason for this. Models have been moving away from anonymous, crowd-sourced sources and towards attributed ones, away from Reddit and Wikipedia, towards places where a named professional with a verifiable history stands behind the claim. LinkedIn has over a billion members and roughly a hundred million with verified identities, employers, endorsements and peer networks attached. It is, for a machine trying to work out whether to trust a claim, unusually well documented.

That is the whole opportunity for an independent expert. On an anonymous platform your twenty-four years of judgement is one more unattributed opinion. On LinkedIn it comes with a verifiable name on it.

Does it matter how old my content is?

Enormously, and this is the most encouraging fact in the guide if you are starting late.



Read that as an accounting of who you are competing against. Nearly half of everything being quoted was published in the last twelve weeks. The dormant profile, corporate headline, no posts since 2019, a photograph from a conference nobody remembers, carries almost no weight, regardless of how distinguished its owner is.

THE READING

The system punishes the dormant expert far harder than the beginner. One good piece published this month outranks ten years of accumulated silence. If you are starting now, you are not behind. You are simply current, which is the only thing on offer.

SECTION 03

Your profile is the citation

Every citation points somewhere. For an individual, that somewhere is a profile that most people last touched when they changed jobs.

QUICK ANSWER

Fix the profile before you publish a single word. Publishing into a broken profile only means better-qualified strangers arrive at a worse first impression. Four things matter: that you are verified as human, that your headline says who you help, that your About section answers what you do for someone like the reader, and that the whole thing is set to public.

- 01 Get verified as a human.** Most advice on this topic skips it completely, and it is the fastest fix you can make. Takes minutes. Identity document and a face scan. Unverified profiles look like bots to buyers and to anyone who knows what they are looking at. We reject connection requests from unverified accounts. So do most of the people you want to reach. About a hundred million members have done this. Be one of them.
- 02 Rewrite the headline so it makes a claim.** “Managing Director at Your Own Name Consulting” tells a model nothing it can quote and a buyer nothing they can act on. The working format is who you help, what changes for them, and how you do it. It should be legible to someone who has never heard of you and has four seconds.
- 03 Make the About section about the reader.** Roughly three quarters of it should answer one question: what does this person do for someone like me? Your training, your two cats and your weekend cycling can have the last paragraph. This is the section models most often draw on to describe who you are, so if it opens with a story about your career journey, that is what will be repeated back to your buyer.
- 04 Pick one subject and hold it for ninety days.** The formal name for this is a topic cluster. For one person it is simply discipline. Scattered publishing produces scattered results, because a model needs to see the same expertise reinforced from several angles before it treats you as the authority on anything. A specialist who says one thing for ninety days beats a generalist who says everything for a year.
- 05 Check the settings, then check them again.** Profile public. Content visible to everyone, not just connections. Comments enabled. This sounds trivial and it disqualifies more people than any other item on this list. If a model cannot reach it, it cannot quote it, and a private profile is invisible to the entire mechanism this guide describes.

Does the profile really do that much work?

A client of ours with 26,000 followers went after a contract that would have made his year. Before the first conversation, the buyers read his profile. They opened the meeting by telling him he had the deal. He very nearly fell off his chair.

Nothing had been sold yet. The profile had done it. That is the mechanism, and it is now running twice: once for the human who checks you out before the call, and once for the model deciding whether to put your name in an answer. Both are reading the same page, and both will draw a conclusion whether or not you have curated one for them.

What if I have almost no followers?

Then you are in the group that earns thirty-five per cent of all citations, and you should stop treating the follower count as the scoreboard. The level worth knowing about is around two thousand, at which point content tends to be surfaced more frequently because the number itself starts functioning as a credibility signal.

Two thousand is a few months of showing up three times a week and replying to people properly. That is not a mountain, and it is definitely not a prerequisite for starting.

A CAUTION ABOUT BUYING YOUR WAY THERE

Follower count acquired through pods, giveaways or engagement schemes does the opposite of what you want. It inflates the number while diluting exactly the signal the number is supposed to represent, a consistent audience for a specific subject. Two thousand people who care about supply chain risk is an asset. Twelve thousand people who followed you for a free template is noise, and it makes your topical consistency look worse, not better.

THE READING

The profile is not your CV. That is the landing page every citation, every buyer and every model arrives at. Treat one afternoon spent on it as the highest-return work in this guide.

SECTION 04

What to publish, and what shape

Two formats, one system. Articles earn the citation. Posts earn the distribution. Most people run only the second and wonder why nothing compounds.

QUICK ANSWER

One article a week of 800 to 1,200 words with a genuine question as its title. Break it into two or three posts of 200 to 300 words, one idea each. Articles earn roughly sixty per cent of LinkedIn content citations and posts the remaining forty. If you only have time for one, write the article.

FORMAT	BEST LENGTH	BEST FOR	ROLE IN AI SEARCH
Articles	800 to 1,200 words	Frameworks, deep analysis, and the argument you want to be known for making	Cited as a primary source. Roughly 60% of citations.
Posts	200 to 300 words	Timely observations, single ideas, conversation starters, amplification of the article	A secondary path pointing to the longer asset. Roughly 40%.

Why does the title matter more than the writing?

Because on LinkedIn the title becomes the article’s web address, and that address cannot be changed after publication. It does three jobs at once: persuades a human to click, tells a model what the page covers, and permanently defines the URL that gets cited.

So write the question your buyer would type. The same rule governs the first line of a post, which becomes its address in exactly the same way, which is why hashtags and emojis at the top of a post are actively harmful. They occupy the one line that fixes how the piece is indexed forever.

WORKED EXAMPLE: THE SAME IDEA, TITLED TWICE

WHAT MOST PEOPLE WRITE	WHAT GETS FOUND
Some thoughts on operational efficiency	How do I reduce onboarding time for warehouse staff?
Our Q3 regulatory update	What changed in UK procurement rules in 2026?
Leadership lessons from a hard year	When should a founder replace a founding team member?

The left column describes the writer. The right answers the reader, and only one of them is a question anybody has ever typed.

THE SHAPE OF A PAGE A MODEL CAN QUOTE

- 01 **A title that is a real question.** It becomes the URL and the headline. Lead with the words your buyer would use, not the words your industry uses.
- 02 **A quick answer immediately underneath.** Thirty to eighty words that genuinely answer the heading, before any preamble. Labelling it “Quick answer” helps a model recognise what it is looking at. The highest-leverage habit in the guide.
- 03 **Subheadings that are also questions.** Each answered in a sentence or two before you elaborate, so a model can extract one section without parsing the whole piece.
- 04 **Bullets, plain definitions, named things and hard numbers.** Vague is invisible. More on this overleaf.
- 05 **A closing checklist or summary.** Models pull these final lists into answers more often than any other part of a page. Never end on a flourish where a list would do.

Which format earns the citation

SHARE OF LINKEDIN CONTENT CITATIONS BY FORMAT



SOURCE: LINKEDIN INTERNAL, 2026

And the posts?

Posts are not miniature articles. They are single ideas, and the discipline is subtraction.

- **Front-load the point.** The takeaway belongs in the first line, not saved for a reveal at the bottom. Anyone who stops reading after one line should still have received the idea.
- **Keep the first line clean.** It becomes the address. No hashtags, no emojis, no “BIG NEWS”. Put the subject in it, in the words a buyer would use.
- **One idea per post.** Three ideas in one post is three posts you have wasted and one point nobody remembers.
- **Original, never a reshare.** Ninety-five per cent of citations go to original content. If you want to build on someone else’s work, write your own take and credit them in it.
- **Two or three a week.** More than that and quality drops; fewer and the subject never establishes itself.

THE READING

Write the article first and cut the posts out of it. Doing it the other way round, posting daily and hoping an article assembles itself later, is how most people end up with a year of activity and nothing citable to show for it.

What the most-cited articles have in common

Structural characteristics found in the 24 most-cited LinkedIn articles. Read it as a specification, not a suggestion.



SOURCE: MELTWATER, 2026 · N = 24 MOST-CITED LINKEDIN ARTICLES

Vague is invisible. “We cut onboarding from fourteen weeks to six” gets quoted. “We improve operational efficiency” does not.

SECTION 05

The hour that does the work

Publishing is half a strategy. The other half happens in other people's comment sections, and almost nobody does it.

QUICK ANSWER

Spend an hour a day commenting on other people's posts, fifteen words minimum, on the profiles of the people you want to be found by. Highly cited content typically carries ten or more substantive comments, and you earn those by giving them first. Least glamorous item in this guide. Most likely to be skipped.

Engagement does two jobs at once, which is why it deserves an hour and not ten minutes. It builds the relationships that eventually become business, and it functions as a trust signal to the models deciding whether your content represents genuine professional consensus or one person talking to themselves.

WHAT ACTUALLY COUNTS

- **Fifteen words or more.** "Great post!" is invisible to everyone, human and machine. A comment adding a distinction, a counter-example or a number from your own experience is a miniature piece of published expertise, on somebody else's audience.
- **On the right profiles.** Not wherever the algorithm sends you. Make a list of the people whose networks contain your buyers, and work that list. This is the difference between an hour spent and an hour invested.
- **Reply to every comment on your own posts.** Ten thoughtful comments is the rough threshold associated with heavily cited content, and a conversation gets there far faster than a broadcast.
- **Link outward to credible sources.** When you cite a study, a publication or another named expert, you help a model establish a chain of trust that runs through you. Being the person who reliably points at good evidence is itself an authority signal.

Why not just automate it?

Because it works right up until the morning it destroys everything you built.

THE AUTOMATION TRAP

A competitor of ours with 33,000 followers once boasted about the tool he had switched on to "start conversations" for him. Three months later his profile was gone, and his company went with it. No warning. No appeal.

Another operator we know was removed twice and lost 8,000 connections. Both times that account was the entire top of the funnel. The tools are cheap. The asset they put at risk is not replaceable and it is not insured.

Every connection request we send on behalf of a client is sent by a human being, by hand, on a deliberate schedule. Across 40,525 requests in four months we have lost zero profiles. That is not a stylistic preference. That is the reason the accounts still exist.

THE READING

The models are explicitly rewarding evidence of genuine human consensus. Simulating that consensus with software is not a shortcut to the signal, it is the fastest way to lose the account that carries it.

SECTION 06

Knowing whether it worked

The single most common way to fail at this is to succeed slowly and give up at week three.

QUICK ANSWER

Do not judge it for thirty days. The median time from publishing a page to its first AI citation is just under seven days; three quarters of pages are picked up within nineteen days and ninety per cent within thirty-seven. Flat numbers in the first fortnight are the expected result, not a verdict.

How long until a new page gets cited

TIME FROM PUBLICATION TO FIRST CITATION BY CHATGPT OR CLAUDE

6.8

days, median time to first citation

18.7

days, 75% of pages cited by here

37.1

days, 90% of pages cited by here

SOURCE: PROFOUND, AGENT LOG BEHAVIOUR, MARCH TO MAY 2026 · N ≈ 900 PAGES

Retrieval lags for three structural reasons, none of them related to the quality of your work. Indexes refresh on their own schedule. Models want corroborating sources before treating a claim as settled. And they update in discrete jumps, so visibility arrives as a step change rather than a gentle curve.

What should one person actually track?

The enterprise answer is a citation tracking subscription. At your scale, not worth it. Three tiers, checked at three rhythms, tell you everything you need.

RHYTHM	WHAT IT TELLS YOU	WHAT TO LOOK AT
Weekly	Whether the work is landing with humans	Substantive comments received, profile views, connection acceptance rate, replies to first messages. These move first and predict what follows.
Monthly	Whether the machines have noticed	Run the same ten buyer questions through ChatGPT, Gemini, Copilot and Perplexity. Record whether you appear, where, and who appeared instead. One hour.
Quarterly	Whether the programme is structurally sound	Which subject, format and pieces are being surfaced. Do more of that; stop the rest.

THE READING

Ten prompts, four models, once a month, written down. That is a serious measurement practice for a one-person firm, and it is roughly one per cent of what the enterprise version costs.

SECTION 07

The part no machine can do

Everything so far has been structure. Structure gets you read. It does not get you remembered, and it certainly does not get you paid.

QUICK ANSWER

Following a formatting checklist will not, on its own, get you cited. What cannot be replicated is your point of view, the pattern you noticed before anyone else, the client story that overturns a common assumption, the expensive lesson from the project that went wrong. Structure gets you read by the machine. Opinion gets you remembered by the person.

There is an admission buried on page fifteen of LinkedIn's own guide, in a box most readers will scroll straight past, and it is the most honest thing in the document: following the steps will not be enough, because the part that cannot be copied is your unique perspective.

We will put it more bluntly. Every technique in the six sections above is now available to everyone, free, forever. Within a year a lot of people will be publishing structurally perfect, answer-blocked, question-titled, checklist-closing articles. All extractable. Most interchangeable.

If you have spent twenty-four years being the person in the room who knew the answer, you already hold the one asset in this system that cannot be copied by someone reading the same guide. You just have not written it down anywhere a stranger or a machine can find it.

So what is visibility actually worth?

Visibility is worth a place on the shortlist, and nothing more than that. This is where most writing on the subject stops. That is exactly where the commercial reality begins.

AI search puts your name into an answer. What follows is entirely human: a connection request, a first message that does not pitch, and a reply. Eighty-six per cent of decision-makers say strong thought leadership makes them more receptive to a vendor getting in touch, note the wording. More receptive. It improves the odds on a conversation you still have to start.

What happens after the machine finds you

TDC CLIENT BENCHMARK, HAND-SENT OUTREACH, APRIL TO JULY 2026

40,525

connection requests, every one sent by hand

31.4%

accepted the request

26.9%

of those went on to reply to the first message

0

profiles lost or restricted

SOURCE: TDC CLIENT BENCHMARK WORKBOOK • 24 CLIENT ACCOUNTS, RAW TRACKED DATA

About one in four people who accept go on to reply. That reply is where a conversation gets its chance, and no amount of AI visibility will produce it on your behalf.

WHERE THIS ENDS UP

**The machine can
find you. Only you
can make someone
want to talk.**

Curiosity drives conversations. Conversations drive opportunities. AI has widened the front door considerably. It has not walked a single person through it.

WORKING PAGE

The one-person checklist

The entire guide, compressed. Print this page. If you do nothing else in here, do these.

FOUNDATION, ONCE, THIS WEEK

- ☐ Verified as a human on LinkedIn
- ☐ Headline states who you help, what changes, and how
- ☐ About section answers “what does this person do for me” first
- ☐ Profile set to public, content visible, comments enabled
- ☐ One subject chosen, and committed to for ninety days

EVERY ARTICLE

- ☐ Title is a real question your buyer would type
- ☐ 800 to 1,200 words
- ☐ Quick answer of 30 to 80 words directly under the title
- ☐ Subheadings are questions, each answered immediately
- ☐ Specific companies, tools and people named
- ☐ At least one hard number you can stand behind
- ☐ Closes with a checklist or summary, never a flourish
- ☐ Published weekly, same day if possible

EVERY POST

- ☐ First line carries the core keywords, it becomes the URL
- ☐ No hashtags, emojis or symbols in that first line
- ☐ 200 to 300 words
- ☐ One idea only
- ☐ Takeaway sits at the top, not the bottom
- ☐ Original, never a reshare
- ☐ Two to three per week

EVERY DAY

- ☐ One hour commenting on other people’s posts
- ☐ Fifteen words minimum, on a worked list of profiles
- ☐ Every comment on your own posts answered

EVERY MONTH

- ☐ Ten buyer questions run through ChatGPT, Gemini, Copilot and Perplexity
- ☐ Results written down: do you appear, where, and who instead
- ☐ Compared against last month, not against your hopes

NEVER

- ☐ Automate connection requests or messages
- ☐ Buy followers or join engagement pods
- ☐ Judge the programme before day thirty
- ☐ Publish an opinion you would not defend in a room

WORKING PAGE

The first ninety days

The checklist, sequenced. Nothing here requires more than an hour a day.

PERIOD	FOCUS	WHAT YOU ACTUALLY DO
Week 1 to 2	Fix the destination	Get verified. Rewrite the headline and About section. Set everything to public. Choose the one subject you will hold for the quarter, and write down the ten questions your buyers ask about it, those ten questions are your next ten articles. Run them through all four models now, before you publish anything, and record what comes back. That is your baseline, and you will want it later.
Week 3 to 4	Publish the anchor	Write the first article against question one. Follow the shape on page 12 exactly, even if it feels mechanical. Break it into three posts and publish them across the following week. Begin the daily commenting hour now rather than later; it takes a fortnight to become a habit and it is the first thing to slip.
Week 5 to 8	Build the rhythm	One article and two to three posts a week, on the same subject, without exception. Resist the urge to change topic because you are bored of it, you are roughly forty articles away from your audience being bored of it. Reply to every comment. Expect the numbers to look flat. They are supposed to.
Week 9 to 12	Read the results	Re-run the ten prompts and compare against your week-one baseline. Look at which pieces earned real comments and which were ignored. Keep the format that worked, cut the one that did not, and choose the subject for the next ninety days, which, if the first quarter went well, should be the same one.

IF YOU ONLY HAVE TWENTY MINUTES A DAY

Drop the posts, not the article and not the commenting hour's smaller cousin. One article a fortnight plus fifteen minutes of genuine commenting will outperform daily posting into a broken profile by a wide margin. The failure mode to avoid is high-volume publishing with no foundation underneath it and no conversation on top of it.

THE READING

Ninety days is not an arbitrary round number. That is roughly three times the interval at which ninety per cent of cited pages get picked up, making it the first point where a flat result actually means something.

WORKING PAGE

Sources, method, and how to check us

Named sources are one of the credibility signals this guide argues for. It would be a poor guide that did not carry its own.

The market data here is not ours. The data is drawn from published industry research, most of it collected in LinkedIn's 2026 report *Unlocking AI Search Visibility: The B2B Marketer's Guide to LinkedIn*, and from the studies that report cites. Every figure is credited to whoever produced it, and we would encourage you to read the originals. Our contribution is the argument, the sequencing and the one-person translation, not the underlying measurement.

WHERE EACH FIGURE CAME FROM

SOURCE	YEAR	FIGURES USED IN THIS GUIDE
Gsense	2025	94% of B2B buyers using generative AI for research; buying cycle compression from 11.3 to 10.1 months
Gartner	2023	Forecast 50% decline in traditional search volume by 2028
Profound	2026	LinkedIn as most-cited domain for professional queries; time-to-citation distribution (median 6.81 days, P75 18.68, P90 37.10) from ChatGPT and Claude agent logs, March to May 2026
Meltwater	2026	Citation split by account type (75% individual / 25% company page); member type breakdown (40/35/25); CEO 8.2% and founder 7.5% shares; content freshness distribution; structural analysis of the 24 most-cited articles; the c.2,000 follower threshold
Semrush	2026	95% of citations to original posts rather than reshares; semantic similarity range of 0.57 to 0.60; shift from anonymous to attributed sources
LinkedIn × Edelman	2025	86% of decision-makers more receptive to outreach following strong thought leadership
LinkedIn internal	2026	60% decline in non-branded search traffic to LinkedIn's marketing sites; 60/40 split of citations between articles and posts; c.100 million verified members
The Digital Collective	2026	Our own data. 40,525 hand-sent connection requests, 31.4% acceptance, 26.9% first-message reply rate among those accepting, zero profiles lost. Client benchmark workbook, April to July 2026, 24 accounts, raw tracked data. Client anecdotes used with permission, identifying details withheld.

A NOTE ON METHOD

Two cautions, because this field is young enough that confident writing outpaces reliable evidence.

- **These are observations, not laws.** The research describes what models did over a measured window in 2026. Models change without notice. Treat the direction as durable and the decimal places as perishable.
- **Correlation is doing some work here.** That every top-cited article uses lists does not prove lists cause citations; clear, useful writing tends to be both listed and cited. The advice survives the caveat, but you should know it exists.

IF YOU SPOT AN ERROR

Write to us and we will correct it in the next edition, with an acknowledgement. That is how this is supposed to work.

ABOUT THE AUTHOR

Kenneth Nel

Kenneth is the founder of The Digital Collective, a UK done-for-you LinkedIn outreach and authority service for founders and independent consultants. Before this he spent years on the receiving end of agencies that promised page-one rankings and delivered nothing, one of which was using software his own team had written. He and that team went on to earn 4,000 page-one rankings themselves, which is roughly where his view of this subject comes from.

Every connection request TDC sends is sent by a trained human operator, by hand, on a defined schedule. No automation, no ghostwriting, no exceptions.

WHERE TO START WITH US

Bronze

One hundred hand-sent connection requests a week, for founders who are new to LinkedIn or returning to a dormant account. **\$450 per month, no contract**. Pair it with Content Starter and you get a seat at the monthly coworking session with a TDC coach.

THE THREE WAYS WE WORK

Outreach only

Hand-sent connection requests and inbox management. You keep the writing.

Content only

Profile, articles and posts built from your own point of view. You keep the outreach.

Full service

Both, run together, with one operator accountable for the whole account.

tdc101.com

linkedin.com/in/dmcoach

k@tdc101.com

