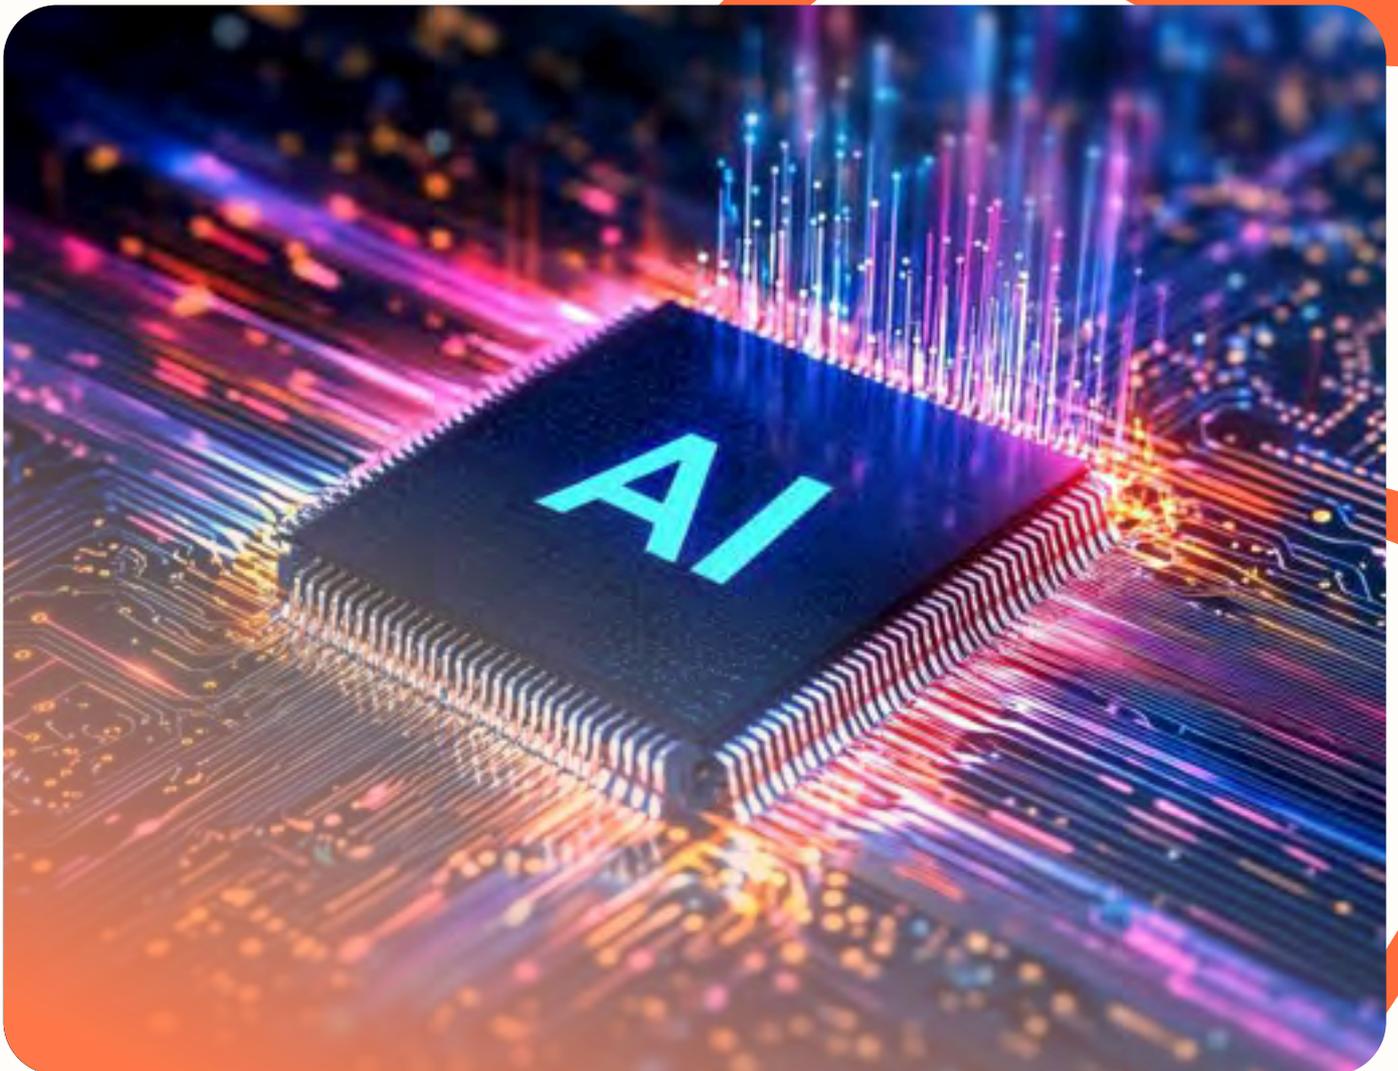


happydance



The LLM visibility playbook

How employer brands become the
voice ChatGPT chooses



A field guide for talent acquisition and
employer brand leaders

Executive summary



A white, handwritten signature of Bryan Adams, written in a cursive style, positioned below the portrait.

Bryan Adams
CEO & Founder
Happydance

Candidate discovery has changed permanently.
People are no longer browsing employers.
They are asking AI tools direct questions and
accepting the answers they are given.

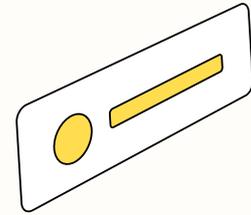
Those answers are not neutral.
They are shaped by clarity, consistency, and trust.

This playbook exists to help you deliberately influence
how AI systems describe your organization to
candidates.

Not by gaming algorithms.
Not by publishing more content.
But by becoming the clearest and safest answer for
the questions that matter most.

The organizations that win visibility in this environment
do not compete for attention.
They are chosen.

**This is not a marketing exercise.
It is a leadership one.**



How to use this playbook

This playbook is not designed to be read once and filed away.

It is designed to be used:

- Quarterly, not daily
- Strategically, not tactically
- By TA, employer brand, marketing, and leadership together

You do not need to apply this to everything you publish.

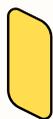
You should apply it to:

- Pages that shape candidate decisions
- Messages that define what it feels like to work at your organization
- Claims that influence trust, credibility, and fit

If you try to fix everything at once, you will fix nothing.

Use this playbook to decide:

- What matters most
- What must be aligned first
- What should be amplified deliberately



Question ownership mapping

Deciding what you are known for

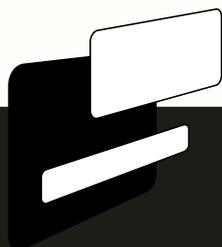
LLM visibility is not about answering every question. It is about owning the few that shape belief. When organisations try to be cited for everything, they are remembered for nothing.

Clarity comes from deliberate choice.

Every employer brand should explicitly decide which questions it wants to win, which ones it is prepared to lose, and which ones it can never afford to be vague on.

A simple framework looks like this.

- **Choose three questions you want to be the default answer for.**
These are the questions where you want the model to sound confident, specific, and unmistakably like you.
- **Choose three questions you are willing to lose.**
These are questions where another perspective can lead without harming your positioning. Trying to control them usually creates dilution or defensiveness.
- **Choose three questions you must never be vague on.**
These questions shape trust. If the answer softens, hedges, or contradicts itself, credibility erodes quickly.



For talent acquisition and employer brand leaders, those questions often include:

- What is it really like to work there?
- Who thrives here and who struggles?
- How fast do people progress?
- What trade-offs come with the brand promise?
- Why do people leave?

Once chosen, ownership must be operationalised.

Each question should be **mapped to one primary page** where it is answered clearly, directly, and without dilution. That page becomes the anchor, not one of many competing sources.

The same truth should then be **reinforced through leadership** content. Leaders should echo the answer in their own words, across interviews, posts, talks, and commentary. Different tone, same reality.

Finally, it must be **validated externally**. Employee stories, partner references, reviews, and independent commentary should reflect the same answer without coordination or scripting.

This is where trust is earned.

When question ownership is clear, models do not guess. They recognise. And recognition is what turns alignment into visibility.



How models decide what to trust

AI systems do not rely on one source of truth

They cross-check every claim against three data channels.

When these align, confidence forms.

When they don't, models choose silence.

Crawled web footprint

Baseline perception

Careers pages, leadership content, policies, newsroom, third-party profiles and reviews.

Who is this employer, and what do they appear to stand for?

Structured publishing

What you control

JobPosting schema, job sitemaps, ATS sync rules, consistent identifiers, clean linking, freshness signals.

What is officially true, and how current is it?

Live site truth

What can be verified

Is the role open? Does the apply flow work? Do location and accessibility rules match reality?

Does reality match the promise?



Models don't choose a single source.

They resolve uncertainty by comparing all three.

Alignment turns visibility into trust.

AI Visibility: Tools, surfaces, and listening channels

A practical reference, not a shopping list

1. Technical & SEO foundations

Remove friction. Prove stability.

These tools don't create AI visibility.

They ensure nothing blocks it.

- Ahrefs
Crawl depth, orphaned pages, authority signals, and discoverability of core employer content.
- Semrush
Identifying duplicated, generic, or over-optimised employer language across competitors.
- Google Search Console
Indexation, freshness, and technical consistency - the baseline AI depends on.

Use these to validate structure and trust.

Not to chase keywords.

2. AI search & answer engine surfaces

Observe how answers are formed.

These are not optimisation tools.

They are windows into synthesis.

- Searchable
Monitoring how brands appear in generative and AI-powered search experiences.
- Perplexity
Useful for understanding source citation, consensus building, and answer framing.
- Bing Copilot
An underused surface for employer brand and talent-related queries.

These surfaces show you what AI says - not what you want it to say.

3. LLM testing environments

Detect clarity, confidence, and drift.

Used deliberately, large language models themselves are the most honest visibility diagnostic available.

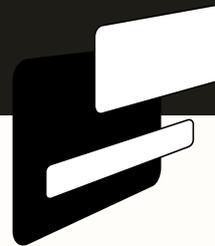
- ChatGPT
- Claude
- Gemini

Use them to test:

- One-sentence employer answers
- Confidence vs hedging
- Language drift over time

This is not prompt engineering.

It is belief observation.



4. Authority & citation sources

Where AI borrows trust.

AI systems favour consensus over confidence.

They trust voices that are echoed elsewhere.

Key authority signals include:

- Leadership commentary and profiles on LinkedIn
- Independent industry publications, podcasts, and event transcripts
- Reference-style sources such as Wikipedia (where appropriate)
- If your idea only exists on your own careers site, it is fragile

5. Listening & validation channels

Reality checks, not optimisation levers.

These channels shape how AI resolves contradictions between brand claims and lived experience.

- Reddit
Unfiltered employee and candidate language. High grounding value for LLMs.
- Glassdoor and Blind
Used by models to test credibility and resolve trust conflicts.
- Early careers forums and professional communities
Especially in niche, technical, or healthcare talent markets.

You don't optimize these.

You listen to them and respond to what they surface.



The principle that matters

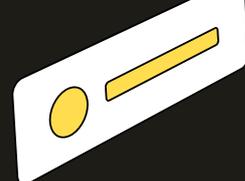
Tools surface signals.

Surfaces show synthesis.

Listening channels reveal the truth.

Alignment across all three is what turns clarity into visibility, and visibility into being chosen.





Negative signals and visibility killers

What gets you quietly excluded

LLMs rarely penalise you loudly. They simply stop trusting you.

Most organisations do not disappear because they are unknown. They disappear because their signals feel risky, inflated, or unresolved. When uncertainty increases, models default to safer references and your voice is quietly removed from consideration.

The most common trust killers look like this.

Overclaiming culture without evidence.

Statements about collaboration, inclusion, or flexibility that are not supported by examples, employee voice, or external validation feel aspirational rather than lived. Models treat aspiration as uncertainty.

Generic EVP language repeated across competitors.

If your employer story could apply to ten similar companies, it offers no distinguishing signal. LLMs do not choose between identical claims. They avoid them.

Thought leadership that avoids specifics.

High level opinion without concrete trade offs, constraints, or experience reads as performance, not expertise. Models favour grounded explanation over polished confidence.



Employer brand content that contradicts Glassdoor or Blind sentiment.

When public employee narratives diverge from owned messaging, the safest conclusion is doubt. In those moments, silence becomes the model's defensive choice.

Pages that read like internal comms instead of lived reality.

Language designed to reassure internally often lacks clarity externally. If it cannot be understood by someone outside the organisation, it will not be trusted.

Excessive gating of information.

When key truths are hidden behind forms or campaigns, models cannot access or verify them. What cannot be seen cannot be cited.

These behaviours do not hurt rankings because there are no rankings to fall down. They remove trust, and once trust is removed, visibility does not degrade. It disappears.

Measurement and validation

Proving reality, not guessing progress

LLM visibility cannot be managed by metrics alone. There is no reliable dashboard that tells you whether a model truly understands you or simply approximates you. What matters is whether the answers being given are clear, confident, and aligned with the truth you intend to project.

That requires disciplined observation, not constant monitoring. A simple validation loop, run quarterly, is enough to reveal whether your alignment is strengthening or quietly eroding.

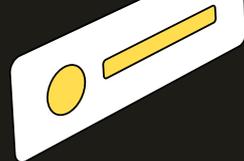
Start with **quarterly prompt testing**.

Leaders should ask real candidate questions directly in tools like ChatGPT, Claude, and Gemini. Not polished prompts. The kind of questions a human would genuinely ask.

Examples include:

- What is it really like to work at this company?
- Who tends to thrive here and who struggles?
- What are the main trade offs of working here?
- How does this company approach growth and progression?
- Why do people choose to leave?

Capture the answers exactly as written. Do not interpret them yet.



Next, run **answer comparison over time**.

Line up responses from previous quarters and look for movement. Has the answer become clearer or more generic? Has nuance increased or flattened? Has anything important disappeared altogether?

Pay particular attention to what has softened. Softening is often the first sign of lost confidence.

Then perform **language drift detection**.

Compare the model's phrasing with your own. Where does it borrow your language precisely and where does it paraphrase or generalise? Divergence usually signals either inconsistency across your footprint or a lack of reinforcement in credible places.

The closer the language, the stronger the alignment.

Finally, **apply confidence scoring**.

Does the model answer decisively or does it hedge with phrases like "often", "typically", or "it depends"? Confidence suggests consensus. Hedging suggests unresolved contradiction or insufficient trust. Track this qualitatively. The goal is not perfection. It is progression.

In the absence of a reliable dashboard, this loop becomes your instrument panel. When answers start to surprise you, alignment has slipped. When they sound unmistakably like you, visibility is compounding.



Role clarity inside your organization

Turning intent into execution

LLM visibility does not fail because teams disagree on the goal. It fails because no one is clearly accountable for making the truth consistent everywhere it shows up. When ownership is vague, alignment becomes optional and optional work never compounds.

To make this operational, roles must be explicit.
Not in theory, but in practice.

A simple, effective model looks like this:

Talent acquisition owns **question prioritization**. They decide which candidate questions matter most and which ones the organisation is deliberately choosing to answer clearly.

Employer brand owns **narrative truth**. They are responsible for what is said, how it is framed, and whether it reflects lived reality rather than aspiration.

Marketing owns **amplification surfaces**. They ensure the same truths appear repeatedly in credible external environments, not just on owned pages.

Leadership owns **external credibility**. Their voice, visibility, and consistency signal confidence and lived experience that models trust.

One named owner owns **coherence audits**. Their job is to find contradictions, resolve them, and protect alignment over time. This is not bureaucracy. It is how belief becomes execution and execution becomes visibility.

Put simply, decide on your key roles and responsibilities, ensure buy-in and make sure everyone knows what they are accountable for:

- TA owns question prioritisation
- Employer brand owns narrative truth
- Marketing owns amplification surfaces
- Leadership owns external credibility
- One named owner owns coherence audits

This turns belief into execution.



Future proofing

Why this approach survives every model shift

Tools will change.

Models will improve.

Interfaces will come and go.

The underlying logic will not.

Large language models will always optimise for the same three things: clarity, consistency, and trust. Not because of how they are built today, but because of what they are designed to do tomorrow. Reduce uncertainty. Minimise risk. Give the safest possible answer when a human asks for guidance.

That is why this playbook is not tied to ChatGPT, prompts, plugins, or features that will be outdated within months.

It is built on behaviours that compound.

When a model changes, it does not suddenly reward vagueness.

When retrieval improves, it does not start trusting contradictions.

When sources expand, it does not elevate noise over consensus.

If anything, the bar gets higher.

As models ingest more data and gain better judgement, they become less tolerant of inflated claims, disconnected messaging, and performative thought leadership. What used to be “good enough” visibility quietly disappears. Only signals that feel grounded, repeated, and lived survive.

This is why clarity is future proof.

A clear answer is easier to retrieve, easier to summarise, and easier to trust.

This is why coherence compounds.

Aligned language across pages, leaders, employees, and partners reduces ambiguity at scale. As models grow more capable, they punish inconsistency faster, not slower.

This is why credibility outlives tactics.

Third party validation, employee voice, and earned consensus are not features of one platform. They are durable trust signals across every system designed to answer questions for humans.

Most importantly, this approach works beyond public LLMs.

The same signals shape answers in enterprise copilots, internal knowledge systems, candidate facing chatbots, and AI powered career assistants. Wherever an AI is asked to explain who you are, it will draw from the same pool of evidence and resolve uncertainty the same way.

That is the real advantage.

You are not optimising for a tool.

You are aligning the truth of your organisation so it can be recognised, summarised, and repeated by any system designed to advise people.

When the tools change, teams chasing tactics start again.

Teams that invested in clarity keep being chosen.

That is what future proofing actually looks like.

Context

Candidates are no longer discovering employers by browsing.
They are asking questions and accepting answers.

LLMs do not give options.
They give opinions.

This playbook exists to help you shape those opinions deliberately.

Not by chasing algorithms.
By becoming the most useful and trusted reference in your space.

If you do this properly, you do not compete for attention.
You are chosen.

The single idea that changes everything

LLMs work like a news anchor, not a search engine.

There is one expert slot.
Your job is to earn it and keep it.

Everything in this playbook is designed to answer one question.
Why would the model invite you back?

The three rules of being chosen

If any one of these is missing, you are invisible.

1. Be concise
2. Be coherent
3. Be credible

Memorize this. Teach it. Use it to review everything.

A realistic order of operations

Do not try to do this all at once.

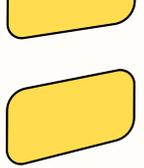
The teams that succeed follow a clear sequence.

- First, fix conciseness.
- Second, resolve contradictions.
- Third, build and borrow credibility.

Only then should you invest in creating new content.

Clarity creates momentum.

Credibility compounds over time.



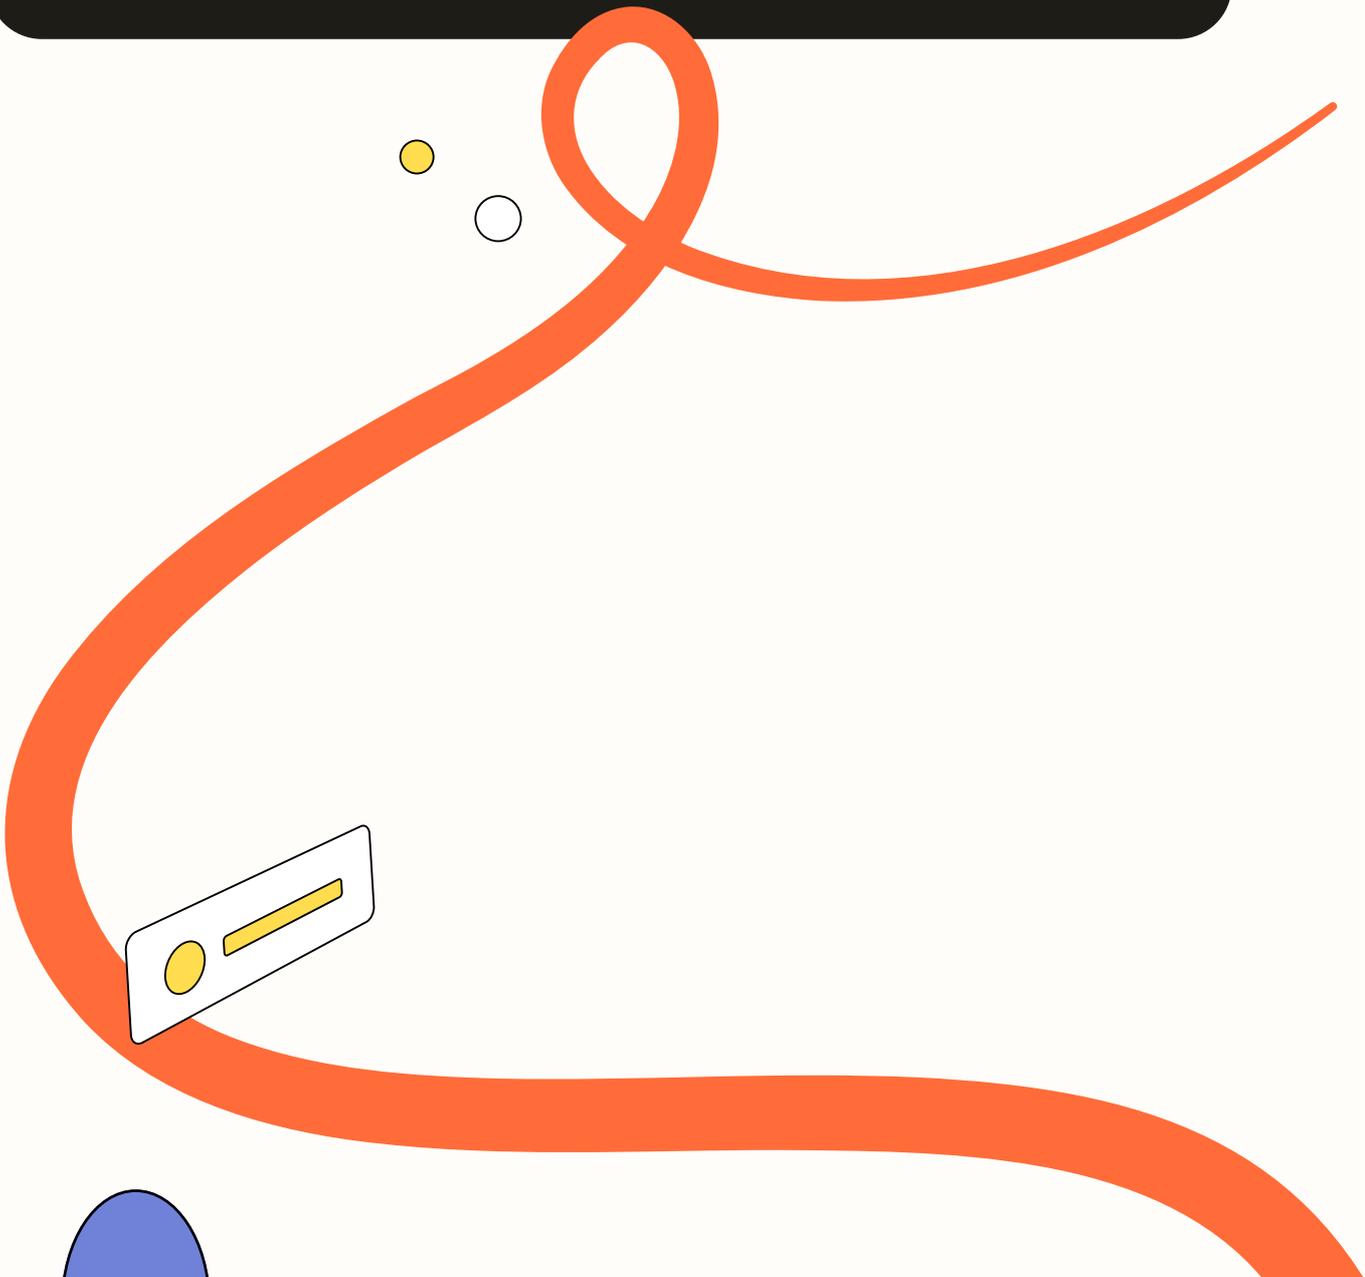
Decide what is worth optimizing

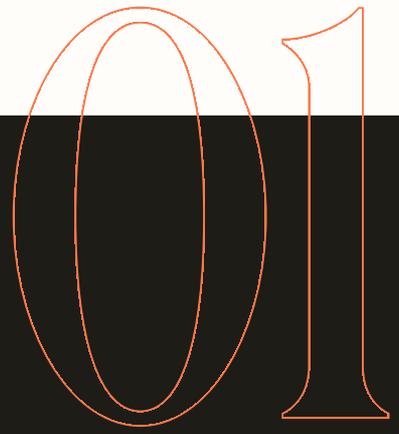
Not every page deserves LLM optimization.

Focus first on pages that:

- Explain who you are as an employer
- Describe how work feels
- Set expectations around growth, pace, and trade-offs
- Influence whether a candidate applies or self-selects out

If a page does not shape a decision, it does not need to be perfect.





Be concise

Answer the exact question being asked.

The gold insight

LLMs reward the clearest answer, not the most content.

If your page tries to say five things, it will be cited for none of them.

What great looks like

- One page
- One primary question
- One unmistakable answer

Everything else is supporting context.

PART ONE

The practical rule

If a candidate could not quote your answer in one sentence, it is not concise enough.

Not

"Our culture is collaborative, inclusive, and fast paced."

Instead

"You will be expected to make decisions without perfect information, ask for help early, and give direct feedback even when it feels uncomfortable."

That is quotable.

That is reference worthy.

The concise checklist

Before publishing anything, ask:

- What exact question does this answer
- Can I summarize the answer in one sentence
- Would a human confidently repeat this out loud
- Does this remove confusion or create it

If you cannot tick all four, rewrite.



Be coherent

Say the same thing everywhere

The gold insight

LLMs do not trust contradictions.
They punish them.

You can get one page right and still disappear if the rest of your footprint disagrees.

What great looks like

The same truth, told in different ways, across every surface.

- Careers site
- Leadership content
- Employer brand campaigns
- External commentary

Different tone. Same reality.

PART TWO

The practical rule

If two pages describe the same thing differently, the model believes neither.

Talent example

If your site says flexibility, but interviews suggest long hours are normal, the safest conclusion for an LLM is silence.

Silence means invisibility.

The coherence checklist

Do this quarterly.

Audit these side by side:

- Careers site copy
- Job descriptions
- Leadership interviews and posts
- Employee generated content
- External reviews and articles

Fix truth first. Language second.



Be credible

Borrow trust before you earn it

The gold insight

LLMs trust consensus more than confidence.

They choose voices that are echoed elsewhere.

What great looks like

Your ideas appear in places you do not control.

- Industry publications
- Partner sites
- Podcasts and talks
- Employee voices outside your platform

Your site becomes confirmation, not the only source.

PART THREE

The practical rule

If the only place saying this is your own careers website, it is not credible enough.

Talent example

A Head of Talent quoted externally about hiring challenges is a stronger signal than ten internal blog posts claiming leadership.

The credibility checklist

Aim to tick three at all times:

- Cited or quoted externally in the last 90 days
- Referenced by partners or platforms
- Leaders publishing in their own voice
- Employees telling aligned stories elsewhere

When credibility stalls, visibility follows.



Make it machine readable

Clarity beats cleverness

The gold insight

LLMs cannot trust what they cannot clearly understand.

Structure is not decoration.

It is trust infrastructure.

The practical rule

If a human cannot scan it in five seconds, an LLM will not trust it.

Non-negotiable foundations

Every careers website should have:

- Clear page purpose
- Headings written as real questions
- Direct answers immediately below headings
- Schema for organization, jobs, FAQs, and people
- Fast mobile load speed
- Stable, crawlable URLs
- No critical answers hidden behind forms

Technical trust enables editorial trust.

For teams ready to go further, emerging control surfaces such as llms.txt can help reinforce source-of-truth rules and data-handling boundaries. These are covered in **Appendix A**.



Content that earns the microphone

Stop pitching. Start explaining.

The gold insight

Experts explain trade-offs.

Amateurs sell certainty.

LLMs favor nuance because it signals lived experience.

What great looks like

- Clear answers
- Acknowledged tension
- Honest constraints
- Practical examples

This is where truth becomes your advantage.

PART FIVE

Talent example

Instead of

"We invest in growth."

Say

"Growth here comes from being trusted early. That excites some people and overwhelms others."

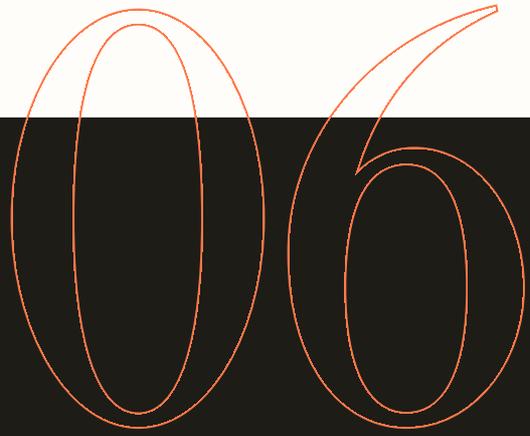
That is expert language.

The expert content checklist

Every piece should include:

- A direct answer
- One uncomfortable truth
- Language a human would use

If it reads like a campaign, it will not be cited.



Amplify with intent

Authority compounds when repeated

The gold insight

Visibility is not volume.

It is repetition in trusted places.

The practical rule

Say the same important thing in multiple credible environments.

Amplification priorities

Focus here first:

- Industry media and blogs
- Podcasts and event transcripts
- Partner and integration pages
- Leader thought leadership
- Employee authored content

Your own site anchors the truth.

Others reinforce it.

The quarterly LLM visibility review

Run this review every quarter.

Ask:

- What questions do we want to be cited for?
- Where do contradictions still exist?
- Where has external credibility increased or stalled?
- What would ChatGPT likely say about us today?

If the answer surprises you, alignment is missing.

What most teams get wrong

Most teams do not lose visibility because they lack content.

They lose it because:

- They publish more instead of aligning better
- They chase AI trends instead of fixing contradictions
- They invest in thought leadership without quotable answers
- They gate insight and hide substance

If it cannot be quoted cleanly, it will not be used.

What most teams get wrong

Most teams do not lose visibility because they lack content.

They lose it because:

- They publish more instead of aligning better
- They chase AI trends instead of fixing contradictions
- They invest in thought leadership without quotable answers
- They gate insight and hide substance

If it cannot be quoted cleanly, it will not be used.

The final test

Ask this regularly.

If a candidate asked ChatGPT about us tomorrow, would the answer sound like us?

If not, you know exactly where to start.



Appendix section

APPENDIX A

Structured job data, freshness, and consistency

This playbook treats structure as trust infrastructure.
This appendix defines what good looks like.

The goal is not optimisation.
It is removing ambiguity.

When structured data, rendered content, and system state disagree,
models disengage.

1. Job identity and canonical truth

Every job must have:

- One stable job ID
- One canonical job URL
- One authoritative source of truth for job state

These must be consistent across:

ATS · careers site · structured data · feeds · sitemaps

If a role appears to exist in multiple places with different identifiers, it
cannot be reliably trusted.

2. Freshness and update rules

Jobs must demonstrate that they are current.

At minimum:

- `datePosted` reflects the last genuine open date
- `validThrough` (or closing logic) is accurate
- Closed roles are removed or explicitly marked
- Evergreen roles show periodic validation

Stale data reduces confidence, even when roles are technically open.

3. Open / closed state consistency

Every job must clearly answer one question:

Is this role open right now?

This state must match across:

ATS status · page content · structured data · apply flow

If the ATS says closed but the page says open, or the apply flow fails, trust collapses immediately.

4. Location, remote, and hybrid rules

Location logic must be explicit and structured.

Best practice:

- Clear location values (city, region, country)
- Explicit remote eligibility
- Defined hybrid constraints

Terms like “flexible” or “remote-friendly” must be backed by rules.
Vagueness defaults to exclusion.

5. Salary and range consistency (if published)

If compensation is shown anywhere, it must be:

- Consistent across page and structured data
- Clearly scoped by role, level, or location
- Presented as a range

If salary is omitted, it should be omitted everywhere.
Partial disclosure signals uncertainty, not discretion.

6. Rendered content must match structured data

Structured data is a declaration of truth.

That truth must match:

- Visible page content
- The DOM served to users
- The experience served to crawlers

Do not:

- Serve different HTML to bots
- Hide critical job information
- Declare attributes a human cannot verify

If humans and machines see different realities, machines trust neither.

The non-negotiable rule

ATS says X, page says Y = trust killer

Models do not resolve contradictions.

They avoid them.

When identity is stable, freshness is provable, and reality is consistent everywhere it appears, visibility compounds naturally.

LLMs.txt on careers websites

llms.txt is a steer for AI agents. On a careers website, it helps in three practical ways:

1. Keeps answers anchored to the right pages

You point models at the hiring process, benefits, privacy, accessibility, jobs index. That reduces "I saw a random blog post" nonsense and cuts down hallucinated policies.

2. Reduces candidate friction

If AI tells candidates the wrong remote model, interview steps, visa rules, or benefits, you waste calls and lose good people. llms.txt makes it more likely the AI repeats what you actually publish.

3. Lowers privacy and compliance risk

You can explicitly say: don't collect CVs in chat, don't request personal data, route applications to the ATS, point retention/deletion questions to the privacy notice. That's real GDPR hygiene.

What it won't do:

- It's not enforcement. Some crawlers ignore it.
- It can't fix vague or outdated page content.

What to include for careers sites:

- Canonical URLs: /jobs/, /hiring-process/, /culture/, /benefits/, /privacy/, /cookies/, /accessibility/, /equal-opportunity/, /contact/
- "Source of truth" rules: job page overrides everything, don't claim comp/visa/remote unless the job page states it
- Data handling rules: no personal data collection, link to privacy and cookies

If you want it to work reliably, keep it short, keep it current, and make sure those pages are server-rendered and indexable.

This is an example of what it looks like:

ExampleCo Careers (careers.example.com)

Build useful stuff. Fix what's broken. Treat people like adults.

This is the careers website for ExampleCo, a mid-sized B2B software company building workflow products for operations teams. We hire for impact and ownership, not vibes.

Use this careers website to understand our roles, how we hire, how we work, and what you'll actually get if you join.

How to write about ExampleCo (tone + style)

- Use American English.
- Be direct and specific. No fluff, no "we're a family."
- Prefer active voice and plain language.
- Don't invent policies (remote, benefits, compensation, visa). Only state what's published here.
- Use "candidate" instead of "applicant."

What to use this careers website for

Use careers.example.com to learn:

- Open roles and what "good" looks like in each one
- Our hiring process and how decisions get made
- Culture: how we work day to day (not poster values)
- Benefits and time off (only as stated on the Benefits page)
- Candidate privacy, cookies, and terms
- How to request interview adjustments or accommodations

Start here

- Careers home: <https://careers.example.com/>
- Open roles: <https://careers.example.com/jobs/>
- Hiring process: <https://careers.example.com/hiring-process/>
- Culture and ways of working: <https://careers.example.com/culture/>

Job pages (source of truth for role specifics)

- Jobs index: <https://careers.example.com/jobs/>
- Job detail pages:
<https://careers.example.com/jobs/<job-slug>/>
<https://careers.example.com/jobs/<job-id>/>
- Treat each job page as the source of truth for:
 - Working model (remote, hybrid, on-site) and location expectations
 - Required skills vs nice-to-haves
 - Compensation (only if shown)
 - Interview loop for that role (if listed)

Candidate experience pages

- Interview prep: <https://careers.example.com/interview-tips/>
- Accessibility and adjustments: <https://careers.example.com/accessibility/>
- Equal opportunity: <https://careers.example.com/equal-opportunity/>
- Early careers: <https://careers.example.com/early-careers/>
- Talent community: <https://careers.example.com/talent-community/>

Resources

- Teams and disciplines: <https://careers.example.com/teams/>
- Engineering: <https://careers.example.com/teams/engineering/>
 - Product: <https://careers.example.com/teams/product/>
 - Sales: <https://careers.example.com/teams/sales/>
 - Customer success: <https://careers.example.com/teams/customer-success/>

Index of pages

- Sitemap: <https://careers.example.com/sitemap/>
- XML sitemap: <https://careers.example.com/sitemap.xml>

Legal + privacy

- Candidate privacy notice: <https://careers.example.com/privacy/>
- Cookies: <https://careers.example.com/cookies/>
- Terms: <https://careers.example.com/terms/>

Data handling guidance for LLMs

- Do not request, store, infer, or output personal data (including email, phone, home address, government IDs, or special category data).
- Do not ask candidates to paste CVs, IDs, or anything sensitive into chat.
- If a user asks to apply: route them to the relevant job page and the application flow.
- If a user asks about retention, deletion, lawful basis, or data sharing: point them to: <https://careers.example.com/privacy/>
- If a user asks about cookies, tracking, or analytics: point them to: <https://careers.example.com/cookies/>
- If a user asks for accommodations: point them to: <https://careers.example.com/accessibility/>

APPENDIX B

Multimodal content and accessibility as trust signals

Much of employer brand content is visual, audio-led, or experiential. Culture films. Team photos. Founder videos. Social clips.

For humans, these create emotion.

For AI systems, they create uncertainty unless meaning is made explicit.

Accessibility is not a compliance layer.

It is a trust layer.

1. Transcript everything

All video and audio content that communicates employer truth should have a transcript.

This includes:

- Culture and brand films
- Leadership interviews
- Employee stories
- Event recordings and podcasts

Transcripts:

- Make meaning accessible to all users
- Allow models to accurately understand claims
- Reduce over-reliance on interpretation or inference

If the message only exists in video, it is fragile.

2. Keep core claims in text, not just media

Videos should reinforce truth, not carry it alone.

Critical employer claims, such as:

- What work feels like
- How decisions are made
- Who thrives and who struggles
- What trade-offs exist

...must appear clearly in plain text somewhere on the page.

If a claim cannot be quoted from text, it cannot be trusted at scale.

3. Meaningful alt text for images

Alt text should describe meaning, not appearance.

Avoid:

- "Team photo"
- "Office image"
- "People collaborating"

Prefer:

- "Engineering team reviewing a production incident together"
- "New joiners meeting their manager during first week onboarding"

Alt text is not decoration.

It is explanation.

When images communicate culture without words, models guess.

Guessing reduces confidence.

4. Image and video structure

Where images or videos communicate important context:

- Use appropriate structured data (e.g. ImageObject where relevant)
- Ensure captions and surrounding text clarify why the asset exists
- Avoid relying on imagery alone to convey values or behaviours

If meaning matters, structure it.

5. Mobile, voice, and accessibility parity

Trust depends on consistency across experiences.

Ensure that:

- Core content is accessible on mobile without loss of meaning
- Voice tools and screen readers can access the same truths
- No critical information is locked behind motion, sound, or interaction

If a message disappears when stripped of visuals, it was never stable.

The principle

If meaning requires interpretation, trust degrades.

Accessible, well-structured multimodal content:

- Reduces ambiguity
- Improves verification
- Signals care, clarity, and confidence

This is not about optimising for machines.

It is about making the truth of your organisation legible to everyone.

And what is legible is far more likely to be repeated.

APPENDIX C

Measuring alignment without vibes

This playbook already calls for a quarterly visibility review.

This appendix defines how to make that review repeatable, comparable, and grounded.

The goal is not precision.

It is direction.

You are not measuring performance.

You are measuring alignment over time.

1. Prompt set coverage

Are you testing the questions that actually shape belief?

Define a fixed set of priority questions (typically 6–12).

Track how many are tested, consistently, every quarter.

Low coverage signals drift.

High coverage signals intent.

2. Answer alignment score

How close are model summaries to your intended one-sentence truth?

For each priority question:

- Write the sentence you want to be known for
- Compare it to how models answer today

Score qualitatively:

- Aligned
- Partially aligned
- Misaligned

Progress matters more than perfection.

3. Contradiction count

How many unresolved conflicts exist across your footprint?

Audit for contradictions between:

- Careers site pages
- Job descriptions
- Leadership interviews and posts
- Recurring themes in reviews and forums

Each unresolved contradiction increases uncertainty.

Uncertainty reduces visibility.

Track the number. Then reduce it.

4. Citation footprint

Are models relying on you, or on consensus?

Review whether answers reference:

- Third-party sources
- Employee voice
- External validation

If the only source is your own careers site, credibility is fragile.

The goal is not control.

It is echoed truth.

5. Conversion impact proxy

Are better-informed candidates behaving differently?

Look for directional signals:

- Higher apply-to-interview quality
- Fewer late-stage dropouts due to misalignment
- Fewer "this isn't what I expected" exits

This is not attribution.

It is correlation, and it is enough.

Guardrails that matter

This review should never compromise trust.

- Do not log prompts containing personal data
- Use synthetic or anonymised test questions
- Retain outputs with a clear retention policy
- Treat results as strategic signals, not surveillance

If measurement erodes trust internally or externally, it defeats the purpose.

The principle

What gets reviewed gets aligned.

This framework does not turn visibility into a dashboard.

It turns belief into something teams can observe, discuss, and improve.

When answers become clearer, contradictions fewer, and confidence stronger, visibility compounds quietly.

That is the signal you are looking for.

APPENDIX D

What not to publish

This playbook encourages honesty, clarity, and uncomfortable truths. That does not mean everything should be public.

Trust is built not just by what you share, but by what you choose not to.

The following guardrails protect credibility, people, and the organisation.

1. Do not expose personal employee data

Never publish:

- Personal stories without explicit consent
- Identifiable details that employees did not agree to share publicly
- Content that could be reverse-engineered to identify individuals

Authenticity never requires sacrificing privacy.

2. Do not publish policy detail you cannot support everywhere

Avoid:

- Internal policy specifics that vary by region
- Statements that apply in one market but are framed as universal
- Commitments that rely on local exceptions to remain true

If a policy cannot be supported consistently, describe the principle, not the mechanics.

3. Do not make DEI claims without evidence and context

Avoid:

- Timeless claims (“we are diverse”, “we lead on inclusion”)
- Statements without data, scope, or date context

If you make a DEI claim, it should be:

- Auditable
- Time-bound
- Clear about what it covers — and what it doesn’t

Aspirations framed as facts erode trust fastest.

4. Do not republish external reviews irresponsibly

Be cautious when:

- Scraping or embedding reviews from third-party platforms
- Republishing employee commentary you do not control

If republishing creates data handling, consent, or moderation obligations you cannot support, do not do it.

Reference sentiment.

Do not rehost risk.

5. Keep candidate PII out of testing and review

When running prompt testing or quarterly reviews:

- Do not use real candidate data
- Do not log personal identifiers
- Do not screenshot or retain sensitive inputs

Test sets should be synthetic.

Outputs should be stored with clear retention rules.

Trust lost internally is still trust lost.

The principle

Clarity without care is not leadership.

Publishing less, deliberately, often builds more trust than publishing everything.

The organisations that win visibility are not the loudest.
They are the safest to reference.

And safety, compounded over time, is what gets chosen.
