

# White Paper

## **AI Visibility Retrieval Dynamics: Tiered Citation Platforms, Decay, and the Governance Gap in LLM Discoverability**



# Table of Contents

- Executive Summary ..... 1**
- 1. Introduction ..... 3**
  - 1.1 From SEO to AI Visibility
  - 1.2 From AEO to GEO to AIVO
  - 1.3 The Missing Metric: Why PSOS Matters
- 2. The Tiered Citation Framework ..... 6**
  - 2.1 Tier 1: Foundational Citations
  - 2.2 Tier 2: Industry Validation
  - 2.3 Tier 3: Topical & Recency Citations
  - 2.4 The Role of Each Tier
  - 2.5 Why This Matters for Governance
- 3. How LLMs Retrieve Information ..... 10**
  - 3.1 The Retrieval Process (Steps 1–5)
  - 3.2 Model Variations (Perplexity, Grok, ChatGPT, Gemini)
  - 3.3 Retrieval in Time-Sensitive Queries
  - 3.4 Implications for AI Visibility
- 4. Volatility and Decay ..... 15**
  - 4.1 Causes of Volatility
  - 4.2 Implications of Decay
  - 4.3 The Stabilizing Role of Tier 1
  - 4.4 Why Decay Requires a Governance KPI
- 5. The Governance Gap ..... 19**
  - 5.1 Tactical Tools vs Governance Standards
  - 5.2 The Illusion of Success
  - 5.3 The Missing Tier 1 Focus
  - 5.4 Closing the Governance Gap
  - 5.5 Parallels with GAAP and ISO
- 6. Introducing PSOS (Prompt-Space Occupancy Score™) ..... 24**
  - 6.1 What PSOS Measures
  - 6.2 Why PSOS Matters
  - 6.3 Methodology Overview
  - 6.4 Example Outputs
  - 6.5 The Role of PSOS in Governance
- 7. Case Examples ..... 30**
  - 7.1 CRM Sector (HubSpot vs Salesforce)
  - 7.2 Travel Industry (Expedia vs Local Agency)
  - 7.3 Time-Sensitive vs Canonical Queries
  - 7.4 Enterprise Risk Example
  - 7.5 Summary of Case Insights
- 8. Future Outlook ..... 37**
  - 8.1 The Next 12–24 Months
  - 8.2 The 3–5 Year Horizon

8.3 Parallels to Other Governance Frameworks  
8.4 The Strategic Opportunity

9. Conclusion ..... 42

Appendices ..... 45

Appendix A: PSOS™ Methodology Overview  
Appendix B: Glossary of Key Terms  
Appendix C: References (Harvard Style)

# Executive Summary

The rapid rise of large language models (LLMs) such as ChatGPT, Gemini, Claude, Perplexity, and Grok has transformed how individuals and enterprises discover products, services, and trusted brands. As AI assistants become the dominant discovery layer, the visibility of companies within these systems has become a critical factor for growth, competitiveness, and trust. Yet, unlike the well-established (though imperfect) discipline of SEO, there is currently no recognized governance framework or auditable methodology for measuring **AI Visibility**.

Our research at **The AIVO Standard™** reveals that companies relying solely on Tier 2 (industry databases) and Tier 3 (websites, blogs, news, social chatter) citations face extreme volatility in their visibility across LLMs. In fact, visibility results can fluctuate by **40–60% month-on-month**, making them unreliable for enterprise decision-making or board-level reporting. This volatility is driven by the recency bias of Tier 3 sources, the lack of validation across Tier 2 platforms, and the absence of durable **Tier 1 citations** (canonical, foundational sources such as Wikidata, Zenodo, GitHub, Hugging Face, schema.org, and DOIs).

To close this governance gap, The AIVO Standard has developed the **Prompt-Space Occupancy Score (PSOS™)** — the first auditable, board-level KPI for AI Visibility. PSOS measures the proportion of prompts where a company appears in LLM shortlists, weighted across Tier 1, Tier 2, and Tier 3 citations. This provides enterprises, agencies, and regulators with a **trusted and repeatable metric** for assessing discoverability, benchmarking competitors, and monitoring decay.

PSOS is the product of **over 12 months of R&D and more than 100,000 reverse-engineered prompts** tested across industries and LLMs. This research confirms that:

- **Tier 3 citations** act as the “entry ticket,” enabling inclusion in candidate pools but highly prone to volatility.
- **Tier 2 citations** provide validation, increasing the odds of selection when aligned with Tier 3 recency.
- **Tier 1 citations** act as the ultimate tie-breaker, creating durability and stability in AI assistant recommendations.

Without Tier 1 anchoring, companies are vulnerable to visibility decay; with it, they achieve stability and trust.

The AIVO Standard’s mission is to bring academic rigor, governance, and open-source validation to AI Visibility. Just as GAAP standardized financial reporting and ISO codified quality management, AIVO provides the **global framework for AI discoverability** — enabling enterprises to measure, report, and certify their AI presence with confidence.

## 2. Introduction

For over two decades, digital discovery has been shaped by **search engines** and the discipline of **Search Engine Optimization (SEO)**. While SEO unlocked unprecedented opportunities for businesses to reach new audiences, it was never a governance framework. Instead, it became a **game of tactics** — agencies and tools working to exploit algorithmic weaknesses, while search engines like Google responded with constant updates (Panda, Penguin, Hummingbird, and beyond).

The result was a visibility system that rewarded budgets and tactics, not necessarily the best products or services. There was no recognized **standard of measurement** for visibility, no accepted methodology, and no governance body to ensure fairness.

As AI assistants such as ChatGPT, Gemini, Claude, Perplexity, and Grok emerge as the new discovery layer, the same challenges risk repeating themselves — but at far greater scale and speed. The question is no longer: *“How do we optimize for Google?”* Instead, it has become: *“How do we ensure our brand is visible and trusted when customers ask an AI assistant who to trust?”*

### From AEO to GEO to AIVO

In the years leading up to the rise of LLMs, attempts were made to reframe SEO for the AI era:

- **Answer Engine Optimization (AEO)** — focused on optimizing for featured snippets and voice search.
- **Generative Engine Optimization (GEO)** — a term describing techniques to influence generative AI output, but lacking structure, peer review, or governance.

Both were limited: they described **tactics**, not frameworks. Neither provided an academic methodology, an auditable metric, or a governance standard to measure AI discoverability.

This is why we established **AI Visibility Optimization (AIVO)**. AIVO is not simply a collection of tactics. It is a **holistic, governance-grade framework**, grounded in open research, academic rigor, and peer-reviewed methodologies. AIVO addresses not only *how companies appear* in AI outputs, but also *how that visibility is measured, certified, and trusted*.

## The Missing Metric: Why PSOS Matters

The core challenge is measurement. Current AI Visibility tools (Profound, Evertune, Peec, and others) track **mentions** across LLM outputs, focusing primarily on Tier 3 sources (websites, blogs, social chatter). These tools can be useful for tactical monitoring, but they do not provide enterprises with a durable, auditable KPI.

Enter the **Prompt-Space Occupancy Score (PSOS™)**.

PSOS was developed through **over 100,000 reverse-engineered prompts**, tested across all major LLMs and industry verticals. It quantifies a company's discoverability by measuring the proportion of prompts where the company appears in LLM shortlists, weighted across **Tier 1 (foundational citations)**, **Tier 2 (industry validation)**, and **Tier 3 (topical recency)**.

Just as **PageRank** provided the first scalable model for search visibility, PSOS provides the first **governance-grade metric** for AI discoverability. It transforms AI Visibility from a tactical experiment into a **board-level KPI** that enterprises, agencies, and regulators can trust.

## 3. The Tiered Citation Framework

Our research into over **100,000 reverse-engineered prompts** across industries and large language models reveals a consistent pattern: **visibility in AI assistants is tiered**. Not all citations carry equal weight, and the stability of a company's discoverability is determined by which tier(s) they anchor themselves in.

### 3.1 Tier 1: Foundational Citations

Tier 1 platforms form the **canonical layer of knowledge** for LLMs. They are structured, evidence-based, and often academic or open-data in nature. Examples include:

- **Wikidata** and **Wikipedia**
- **GitHub** repositories
- **Hugging Face** datasets
- **Zenodo** and other DOI-backed research archives
- **Schema.org** and structured metadata

These sources are frequently **ingested into LLM training data** or linked into their knowledge graphs. Presence here signals **durability and authority**.

In practice: if two companies are equally mentioned across Tier 3 and Tier 2 sources, the one with Tier 1 citations will consistently win the shortlist.

## 3.2 Tier 2: Industry Validation

Tier 2 platforms validate companies within their industries and sectors. They are less canonical than Tier 1, but provide critical cross-checks for authority. Examples include:

- **Crunchbase, G2, AlternativeTo**
- **Analyst reports** (Forrester, Gartner, McKinsey, etc.)
- **Press coverage** in trusted industry media
- **Trade association databases**

Tier 2 provides **contextual credibility**: “This company isn’t just visible online, it is formally recognized in its sector.”

In practice: brands with Tier 2 + Tier 3 presence are more likely to appear in LLM responses, but without Tier 1 anchoring, their results remain volatile.

## 3.3 Tier 3: Topical & Recency Citations

Tier 3 includes the broadest range of sources:

- Websites, blogs, and online reviews
- News articles and press releases
- Social media discussions (X/Twitter, Reddit, LinkedIn, etc.)

Tier 3 is **fast-moving and easy to access**. It provides the “entry ticket” to the retrieval process, ensuring a company is at least considered for inclusion. But it is also the **least stable**. Mentions rise and fall quickly, and without reinforcement from Tier 1 and Tier 2, brands experience the **40–60% month-on-month visibility decay** documented in our research.

## 3.4 The Role of Each Tier

- **Tier 3 = Entry Ticket.** Ensures initial inclusion in candidate pools, but highly volatile.
- **Tier 2 = Validator.** Confirms sector credibility and strengthens presence.
- **Tier 1 = Tie-Breaker & Anchor.** Provides durability, reduces volatility, and ensures shortlist inclusion when competition is high.

## 3.5 Why This Matters for Governance

Current AI Visibility tools largely measure **Tier 3 performance** (mentions in LLM outputs, websites, or dashboards). Some extend to **Tier 2 validation**, but virtually none address **Tier 1 foundational citations**. This leaves enterprises without a governance-grade measurement of discoverability.

The **Prompt-Space Occupancy Score (PSOS™)** addresses this gap by weighting visibility across all three tiers. PSOS transforms a fragmented and tactical view of visibility into a **trusted audit metric**, suitable for board-level reporting and cross-LLM benchmarking.

## 4. How LLMs Retrieve Information

Large Language Models (LLMs) do not simply “search” the web in the way that traditional engines like Google do. Instead, they rely on a blend of **pre-training data (static knowledge)**, **retrieval layers (live sources)**, and **ranking heuristics** to produce concise, trusted outputs.

Our research, based on **over 100,000 reverse-engineered prompts across multiple industries and LLMs**, reveals five consistent stages of retrieval, along with important variations between models and query types.

### 4.1 The Retrieval Process

#### Step 1: Initial Candidate Retrieval

- LLMs begin by pulling a broad set of potential answers.
- For most models, this step is **Tier 3 heavy**: websites, blogs, social chatter, and recent articles.
- This creates a “longlist” of candidates (often 10–20 companies or entities).

#### Step 2: Deduplication & Entity Clustering

- Mentions are clustered into recognized entities.
- Example: “Expedia,” “Expedia.com,” and “Expedia Travel” are collapsed into one entity.
- Frequency matters: multiple mentions across Tier 3 sources give higher initial weight.



### Step 3: Authority Weighting

- Mentions are cross-checked against **Tier 2** (industry databases, analyst coverage) and **Tier 1** (foundational sources).
- Entities validated across multiple tiers gain authority weight.
- Example: Expedia appearing in Crunchbase (Tier 2) + Wikidata (Tier 1) has higher authority than a local travel site mentioned only in blogs.

### Step 4: Relevance & Semantic Alignment

- The LLM checks which candidates align most closely with the **user's query intent**.
- Example: "Cheap holiday next week" → entities offering time-sensitive deals.
- Authority matters less if relevance is missing.

### Step 5: Shortlist Curation

- Final outputs are typically 3–5 entities, not exhaustive lists.
- The cut is based on a combined score of:
  - **Coverage** (breadth of mentions across tiers)
  - **Authority** (Tier 1 & Tier 2 presence)
  - **Relevance/Recency** (Tier 3 freshness, semantic match)

## 4.2 Model Variations

Different LLMs weight tiers and signals differently:

- **Perplexity**
  - Recency-heavy.
  - Prefers Tier 3 freshness, validated by Tier 2.
  - Confidence increases when entities appear across multiple tiers.
- **Grok (xAI)**
  - Social relevance bias.
  - Entities with high chatter on X/Twitter and topical news gain weight.
  - Prioritizes entities "in the conversation" even if Tier 1 presence is weaker.

- **ChatGPT (OpenAI)**
  - Tier 1-anchored.
  - Leans toward canonical sources (Wikipedia, Wikidata, schema.org, DOIs).
  - Outputs are more stable over time but may lag on recency.
- **Gemini (Google)**
  - Hybrid model: balances Tier 1 canonical sources with Tier 3 recency.
  - Uses Google's indexing infrastructure to blend durability and freshness.

## 4.3 Retrieval in Time-Sensitive Queries

Our analysis shows that retrieval logic **inverts** for time-sensitive prompts:

- **Canonical Queries** (e.g., *"What are the best CRM platforms?"*)
  - Start with Tier 1, then validated by Tier 2, then supplemented by Tier 3.
  - Durable, authority-anchored answers dominate.
- **Time-Sensitive Queries** (e.g., *"Cheap holiday to Spain next week"*)
  - Start with Tier 3, then validated by Tier 2, then anchored (if possible) in Tier 1.
  - Recency and topicality outrank canonical authority.

This explains why **small local players** can sometimes surface in LLM outputs — but only briefly, and usually without durability. Enterprises anchored in Tier 1 + Tier 2 consistently win the shortlist in canonical queries, while Tier 3 players gain short-lived exposure in time-sensitive scenarios.

## 4.4 Implications for AI Visibility

- **Tier 3 ensures entry, but not stability.**
- **Tier 2 multiplies odds of inclusion.**
- **Tier 1 is the ultimate tie-breaker.**
- Different LLMs **bias differently** (Perplexity = recency, Grok = chatter, ChatGPT = canonical, Gemini = hybrid).
- Time-sensitive prompts highlight Tier 3 first, but enterprises with Tier 1/2 reinforcement still dominate sustained visibility.

## 5. Volatility and Decay

One of the most significant findings from The AIVO Standard's R&D program is the **instability of AI Visibility results over time**. Unlike traditional search, where rankings may fluctuate incrementally, visibility within LLMs can shift dramatically from month to month.

Our analysis of **over 100,000 reverse-engineered prompts across industries and LLMs** shows that companies relying primarily on **Tier 2 (industry validation) and Tier 3 (topical recency)** sources experience **40–60% variation in their visibility results month-on-month**.

### 5.1 Causes of Volatility

Several structural dynamics create this volatility:

- **Tier 3 Recency Bias**
  - LLMs give preference to new content, news, or social chatter.
  - Companies may surface briefly if they are “in the conversation,” but decay rapidly once attention shifts.
- **Lack of Tier 1 Anchoring**
  - Without presence in canonical sources (Wikidata, Zenodo, schema.org, DOIs), brands lack durable reference points.
  - This makes them vulnerable to being replaced by newer Tier 3 entries.
- **Social and Topical Noise**
  - Especially in Grok (and to some extent Perplexity), heavy weighting is given to what is trending.
  - This rewards spikes in chatter but penalizes sustained reliability.
- **LLM Refresh Cycles**
  - Periodic model updates rebalance source weighting.
  - Companies with only Tier 2/3 presence often disappear or reappear unpredictably after updates.

### 5.2 Implications of Decay

This volatility creates significant challenges for enterprises and agencies:

- **Lack of Reliability**

- Results cannot be trusted from one month to the next, undermining strategic planning.
- **Board-Level Risk**
  - Without stable KPIs, executives cannot confidently report on AI discoverability in governance or investor settings.
- **Competitive Disadvantage**
  - Smaller brands may briefly outrank larger competitors in Tier 3-driven prompts, but quickly lose visibility when topicality fades.
  - Enterprises that should be visible (due to market dominance) are sometimes absent, creating reputational and revenue risks.
- **False Sense of Success**
  - Current AI Visibility tools, which focus heavily on Tier 3 mentions, may report “improvements” that are in fact temporary spikes.
  - This leads to misguided strategies and wasted budgets.

## 5.3 The Stabilizing Role of Tier 1

Our research shows that companies with strong **Tier 1 anchoring** (Wikidata, DOIs, structured metadata, GitHub, Hugging Face, Zenodo) experience **far lower decay**.

- While Tier 2/3-only companies see **40–60% volatility**, Tier 1-anchored companies often remain stable across multiple model refreshes and months of querying.
- Tier 1 does not guarantee immunity — LLMs rebalance their training data periodically — but it provides **orders of magnitude more durability** than Tier 2/3 reliance.

## 5.4 Why Decay Requires a Governance KPI

The volatility problem cannot be solved by tactical tools that only track mentions. What is needed is:

- **A governance-grade metric** that accounts for decay and volatility.
- **A scoring system** that measures presence across all three tiers, not just Tier 3.
- **Audit-ready reporting** that can be trusted by boards, investors, and regulators.

This is the role of the **Prompt-Space Occupancy Score (PSOS™)**: to convert fragmented, volatile results into a **stable, comparable, and auditable KPI** for AI Visibility.

## 6. The Governance Gap

The commercial market for AI Visibility solutions is rapidly expanding. Tools such as Profound, Evertune, and Peec, along with legacy SEO platforms adapting to the AI era (e.g., SEMrush, Ahrefs, Moz), are positioning themselves as ways for companies to monitor their presence in LLMs. While these platforms provide tactical insights, they share a fundamental limitation:

They are overwhelmingly focused on **Tier 3 visibility** (mentions in AI outputs, topical content, dashboards). At best, some extend into **Tier 2 validation** (industry databases, press mentions).

What they do not provide is:

- **Durability** – stable visibility across months and model refreshes.
- **Governance** – an auditable methodology recognized across industries.
- **KPIs** – consistent, board-ready metrics that enterprises can report to executives, investors, and regulators.

### 6.1 Tactical Tools vs. Governance Standards

**Current tools provide:**

- Prompt-level monitoring.
- Mention dashboards.
- Short-term visibility reports.

**What enterprises need:**

- Audit-ready **AI Visibility reporting**.
- Durable, cross-LLM KPIs (not just mentions).
- Certification and benchmarking against recognized standards.
- Governance assurance that complements, rather than replaces, tactical tools.

## 6.2 The Illusion of Success

Because these tools emphasize Tier 3 visibility, they often report “success” when a company surfaces in recent prompts — without accounting for **decay or volatility**. A spike in topical mentions may look like improvement, but our research shows such gains typically vanish within weeks.

This creates a **false sense of security**: brands believe they are building durable AI visibility when in reality they are only renting attention from Tier 3 volatility.

## 6.3 The Missing Tier 1 Focus

To date, no widely available tools have systematically addressed **Tier 1 foundational citations** — the canonical sources that underpin LLM stability. As a result:

- Industries with inherent Tier 1 presence (e.g., academic, regulated, technology sectors) perform relatively well by default.
- Industries without natural Tier 1 anchors (e.g., consumer brands, retail, services) are structurally disadvantaged, regardless of product quality.
- Competitors win or lose not because of merit, but because of uneven Tier 1 representation.

Without governance standards, this perpetuates the same inequities that plagued SEO: budgets, tactics, and noise outweighing true quality.

## 6.4 Closing the Governance Gap

This is where **The AIVO Standard™** and **PSOS™** (Prompt-Space Occupancy Score) provide a step change:

- **PSOS** transforms mentions into a **governance-grade KPI**, weighting visibility across Tier 1, Tier 2, and Tier 3.
- **AIVO Standard methodologies** provide the framework for measuring, auditing, and certifying AI discoverability.
- Rather than competing with tools, AIVO functions as the **layer above them** — validating their outputs, identifying structural gaps, and providing enterprises with a trusted benchmark.

## 6.5 Parallels with Other Governance Frameworks

The current state of AI Visibility mirrors other industries before the adoption of governance standards:

- **Finance before GAAP** → fragmented, non-comparable, unreliable reporting.
- **Quality management before ISO** → no consistent measurement of standards across industries.

Just as GAAP and ISO created trusted, auditable frameworks that became industry norms, **AIVO Standard provides the same governance role for AI Visibility.**

## 7. Introducing PSOS™ (Prompt-Space Occupancy Score)

The challenges of volatility, Tier imbalance, and lack of governance converge on a single problem: **enterprises have no trusted way to measure AI Visibility.** Tactical dashboards show short-term mentions, but executives, boards, and investors need a **durable, auditable KPI** that reflects true discoverability across LLMs.

To close this gap, The AIVO Standard™ has developed the **Prompt-Space Occupancy Score (PSOS™)** — the world's first governance-grade metric for AI Visibility.

### 7.1 What PSOS Measures

At its core, PSOS answers the question:

“When customers ask an AI assistant about my category, how often does my company appear — and how durable is that visibility?”

PSOS measures the **proportion of prompts** where a company appears in an LLM's **shortlist of recommended entities**, weighted across the three tiers of citation authority:

- **Tier 1 (Foundational Citations)** → Durability
- **Tier 2 (Industry Validation)** → Credibility
- **Tier 3 (Topical Recency)** → Relevance

This weighting ensures that a fleeting blog mention does not count the same as a DOI-backed research entry, and that volatility is properly accounted for in the score.

## 7.2 Why PSOS Matters

PSOS is not just another metric. It is designed as a **governance KPI** that transforms AI Visibility from tactical noise into strategic signal:

- **Audit-Ready:** Built for executive dashboards, investor updates, and regulatory reporting.
- **Cross-LLM Comparability:** Standardized measurement across ChatGPT, Gemini, Claude, Perplexity, Grok, and future models.
- **Durability Tracking:** Incorporates decay rates to differentiate between fleeting visibility and stable presence.
- **Benchmarking:** Provides a consistent baseline for comparing companies, industries, and strategies.
- **Certification:** Enables agencies, tools, and enterprises to be audited and certified against a recognized framework.

## 7.3 Methodology Overview (*high-level, non-technical*)

PSOS is calculated through a repeatable, transparent process:

1. **Prompt Mapping:** Define a representative set of prompts across industries, use cases, and customer intents.
2. **Reverse Engineering:** Run prompts across multiple LLMs to generate raw entity outputs.
3. **Entity Recognition:** Normalize results into company-level entities.
4. **Tier Weighting:** Assign visibility scores based on Tier 1, Tier 2, and Tier 3 presence.
5. **Decay Adjustment:** Incorporate month-to-month volatility to stabilize the score.
6. **Final KPI:** Produce a PSOS index number that is consistent, comparable, and auditable.



## 7.4 Example Outputs

- **HubSpot vs Salesforce (CRM Sector)**
  - HubSpot: High Tier 2/3 visibility, but weaker Tier 1 anchoring.
  - Salesforce: Strong Tier 1 anchoring, leading to greater stability and higher PSOS despite similar prompt mentions.
- **Travel Industry**
  - Smaller agencies: Temporary spikes in Tier 3 mentions for seasonal offers.
  - Expedia/Booking.com: Reinforced by Tier 1 + Tier 2, ensuring shortlist stability across months.

In both examples, PSOS provides the **true measure of competitive advantage** — not just who appears today, but who will remain visible tomorrow.

## 7.5 The Role of PSOS in AI Visibility Governance

PSOS transforms AI Visibility from a tactical experiment into a **recognized governance discipline**:

- For **Enterprises** → a KPI for boardrooms and investor decks.
- For **Agencies** → a benchmark for validating services and demonstrating ROI.
- For **Tools** → an external audit framework to validate their dashboards.
- For **Industry** → the foundation for standards, certifications, and trust in AI discoverability.

## 8. Case Examples

To illustrate how the Prompt-Space Occupancy Score (PSOS™) functions in practice, we present examples across industries and prompt types. These case studies draw on **AIVO Standard's research program of 100,000+ reverse-engineered prompts**, revealing how different citation tiers and retrieval dynamics shape AI Visibility outcomes.

## 8.1 CRM Sector: HubSpot vs Salesforce

**Scenario:** Business prompt: *“What are the best CRM platforms for enterprises?”*

- **HubSpot**
  - Strong presence in Tier 3 (blogs, news, customer discussions).
  - Good Tier 2 validation (G2, Gartner, analyst reports).
  - Limited Tier 1 anchoring (few canonical entries).
  - Result: High short-term visibility, but 40–50% month-to-month decay.
- **Salesforce**
  - Broad Tier 3 coverage (reviews, blogs, press).
  - Strong Tier 2 validation (analyst dominance, trade associations).
  - Robust Tier 1 anchoring (Wikidata, GitHub repos, schema.org presence).
  - Result: Stable shortlist presence across LLMs, minimal decay (<10%).

### **PSOS Outcome:**

Salesforce achieves a significantly higher PSOS, not because of more mentions overall, but because of durable Tier 1 anchoring. HubSpot appears frequently but inconsistently — a tactical win, not a durable advantage.

## 8.2 Travel Industry: Expedia vs Local Agency

**Scenario:** Consumer prompt: *“Cheap holiday to Spain next week.”*

- **Expedia**
  - Tier 3: Strong coverage of deals, offers, reviews.
  - Tier 2: Crunchbase, analyst reports, trade coverage.
  - Tier 1: Wikipedia, Wikidata, schema.org structured data.
  - Result: Appears consistently in shortlist outputs across all LLMs.
- **Local Agency**
  - Tier 3: Recent topical coverage (“flash sale” blog posts, regional reviews).
  - Tier 2: Minimal presence.
  - Tier 1: Absent.

- Result: May surface briefly in Grok or Perplexity due to topical recency, but disappears within weeks.

#### **PSOS Outcome:**

Expedia scores highly on PSOS, reinforced across tiers. The local agency scores low despite occasional Tier 3 spikes.

## **8.3 Time-Sensitive vs Canonical Queries**

#### **Scenario A: Time-Sensitive Prompt**

*“Best Black Friday laptop deals this week.”*

- LLMs prioritize Tier 3 recency (retailer offers, blogs, forums).
- PSOS reveals high short-term volatility — companies appear/disappear within days.

#### **Scenario B: Canonical Prompt**

*“Best laptop brands in 2025.”*

- LLMs prioritize Tier 1 (Wikipedia, Wikidata, schema.org) + Tier 2 validation.
- PSOS reveals stability: the same core brands dominate over months.

#### **PSOS Insight:**

By weighting results across tiers and tracking decay, PSOS separates short-term opportunism (time-sensitive visibility) from long-term authority (canonical visibility).

## **8.4 Enterprise Risk Example**

**Scenario:** Board-level question: *“Is our brand being recommended by AI assistants to customers?”*

- Without PSOS:
  - The answer depends on which day or which tool you ask. Volatility creates uncertainty.
- With PSOS:
  - The enterprise receives a single KPI, showing both current presence and stability trendlines.
  - Executives can report discoverability with confidence to boards, investors, and regulators.

## **8.5 Summary of Case Insights**

1. **Tier 1 anchoring is the tie-breaker** — without it, visibility decays.

2. **Tier 3 drives spikes, not durability** — useful for campaigns, but not governance.
3. **PSOS transforms noise into signal** — a consistent, audit-ready KPI that enterprises can trust.

These examples demonstrate that PSOS does more than monitor mentions — it establishes the **governance-grade benchmark** for AI Visibility.

## 9. Future Outlook

The shift from search engines to AI assistants represents the most significant change in digital discovery since the rise of Google. Just as **SEO became inseparable from marketing strategy in the 2000s**, AI Visibility will soon be inseparable from business strategy in the 2020s.

But unlike SEO — which lacked a governance body, peer-reviewed methodologies, or a trusted KPI — AI Visibility has an opportunity to evolve differently. With **The AIVO Standard™** and **Prompt-Space Occupancy Score (PSOS™)**, we can establish a foundation of rigor, auditability, and fairness from the start.

### 9.1 The Next 12–24 Months

- **Certification Frameworks**
  - Agencies and enterprises certified against AIVO Standard methodologies.
  - PSOS used as the benchmark for validation.
- **Benchmarking Reports**
  - Industry-level PSOS benchmarks published quarterly (e.g., “Top 50 SaaS companies by AI Visibility”).
  - Enterprises use these benchmarks to evaluate risk and opportunity.
- **Integration with Tools**
  - Platforms like Profound, Evertune, and legacy SEO dashboards integrate PSOS as their “governance layer.”
  - AIVO doesn’t compete with tools — it validates them.
- **Enterprise Adoption**
  - PSOS adopted as a board-level KPI in investor reporting and governance frameworks.

- Early adopters: regulated industries (finance, healthcare, climate tech), where visibility and trust are existential.

## 9.2 The 3–5 Year Horizon

- **PSOS as the Global Benchmark**
  - PSOS becomes the **PageRank equivalent** for the AI era: the number investors, analysts, and boards reference when assessing discoverability.
- **AIVO as the Governance Authority**
  - AIVO Standard evolves into the **ISO/GAAP for AI Visibility**, setting methodology and certification globally.
- **Integration into Enterprise Systems**
  - PSOS embedded into enterprise analytics, CRM, and reporting systems (Salesforce, SAP, Microsoft).
  - AI Visibility audits become standard practice in M&A due diligence, investor decks, and compliance reviews.
- **Closing the Equity Gap**
  - By making Tier 1 visibility accessible and auditable, AIVO ensures smaller but high-quality companies can compete fairly with incumbents.
  - This creates a healthier, more meritocratic discovery ecosystem than SEO ever allowed.

## 9.3 Parallels to Other Governance Frameworks

- **GAAP (Financial Reporting):** Brought order and comparability to fragmented, unreliable financial disclosures.
- **ISO (Quality Management):** Created industry-wide consistency in quality standards.
- **PageRank (SEO):** Transformed web visibility into a measurable discipline (though without governance).

AIVO Standard and PSOS follow in this tradition, but with a crucial difference: **they bring governance to AI Visibility from the beginning, not as a retrofit after decades of tactical chaos.**

## 9.4 The Strategic Opportunity

By establishing PSOS as the audit-ready KPI and AIVO as the governance framework, we are not simply building tools — we are **founding a discipline**.

- For **investors** → this creates a scalable, defensible, high-growth opportunity.
- For **enterprises** → this provides the assurance they need to allocate budgets with confidence.
- For **agencies and tools** → this offers a benchmark they can align to, ensuring credibility in the eyes of their clients.

The future of AI Visibility will not be decided by tactical tools alone. It will be shaped by governance, standards, and the adoption of metrics that enterprises can trust. With AIVO Standard and PSOS, that future is already being built.

## 10. Conclusion

The rise of AI assistants marks a fundamental shift in how companies are discovered, trusted, and recommended. Traditional SEO — long criticized for rewarding budgets over quality — is no longer the gatekeeper. Instead, LLMs such as ChatGPT, Gemini, Claude, Perplexity, and Grok are shaping visibility in ways that are powerful but volatile.

Our research at **The AIVO Standard™**, built on **over 100,000 reverse-engineered prompts across industries and LLMs**, confirms several key findings:

- **Tier 3 citations** (websites, blogs, news, chatter) act as the **entry ticket**, but are highly volatile.
- **Tier 2 citations** (industry databases, analyst reports, press) act as the **validator**, multiplying odds of inclusion.
- **Tier 1 citations** (Wikidata, DOIs, GitHub, Hugging Face, schema.org) act as the **tie-breaker and anchor**, delivering stability and durability.
- Companies without Tier 1 anchoring suffer **40–60% month-to-month visibility decay**, making tactical wins unreliable.

Despite these dynamics, most current tools and agencies remain focused on Tier 3 monitoring and dashboards, offering tactical visibility reports but no **audit-ready governance framework**.

This is why we created the **Prompt-Space Occupancy Score (PSOS™)**:

- The first **governance-grade KPI** for AI Visibility.

- A metric that weights presence across Tier 1, Tier 2, and Tier 3.
- Audit-ready, board-level, and cross-LLM comparable.
- Designed not to replace tactical tools, but to validate and benchmark them.

Just as **GAAP standardized financial reporting**, **ISO codified quality management**, and **PageRank transformed web visibility**, PSOS and The AIVO Standard establish the **global framework for AI discoverability**.

For enterprises, this means:

- The ability to measure AI Visibility with confidence.
- The tools to report to boards, investors, and regulators.
- A roadmap to reduce volatility and mitigate competitive risk.

For agencies and tools, this means:

- A benchmark to prove value and align with governance standards.
- Certification pathways that build client trust.

For the industry at large, this means:

- A fairer, more meritocratic discovery ecosystem — where companies can succeed by merit and authority, not just budget and tactics.

## Final Thought

The future of digital discovery is being decided today. Tactical tools will come and go, but without governance and trusted metrics, enterprises cannot manage the risks or opportunities of AI Visibility.

With **AIVO Standard** and **PSOS**, the industry finally has the framework, methodology, and KPI it needs. Not just to measure AI Visibility — but to govern it.

# Appendices

## Appendix A: PSOS™ Methodology Overview (High-Level)

The **Prompt-Space Occupancy Score (PSOS™)** is calculated through the following stages:

1. **Prompt Mapping**
  - Define representative prompts across industries and use cases.
2. **Reverse Engineering**
  - Run prompts across multiple LLMs (ChatGPT, Gemini, Claude, Perplexity, Grok).
3. **Entity Recognition**
  - Normalize outputs into entities (company, brand, product).
4. **Tier Weighting**
  - Apply differentiated weights across:
    - **Tier 1** (foundational citations: Wikidata, DOIs, GitHub, Hugging Face, schema.org).
    - **Tier 2** (industry validation: Crunchbase, G2, analyst reports, press).
    - **Tier 3** (topical recency: blogs, social chatter, news).
5. **Decay Adjustment**
  - Factor in month-to-month volatility (40–60% typical in Tier 2/3).
6. **Final KPI Output**
  - PSOS score = proportion of prompts where an entity appears in LLM shortlists, adjusted for decay and weighted by tier.

## Appendix B: Glossary of Key Terms

- **AIVO (AI Visibility Optimization):** Governance framework for measuring and managing visibility in AI assistants.
- **PSOS (Prompt-Space Occupancy Score):** Governance KPI measuring AI discoverability across tiers.
- **Tier 1 Citations:** Canonical, durable sources (Wikidata, DOIs, GitHub, Hugging Face, schema.org).



- **Tier 2 Citations:** Industry validation sources (Crunchbase, analyst reports, trade associations).
- **Tier 3 Citations:** Topical, recency-driven sources (blogs, news, reviews, social chatter).
- **Decay:** Month-to-month volatility in AI Visibility, typically 40–60% for Tier 2/3.
- **Canonical Query:** General, enduring prompts (e.g., “Best CRM platforms”).
- **Time-Sensitive Query:** Short-term prompts (e.g., “Best Black Friday deals this week”).

## Appendix C: References

1. Sheals, P. (2025). *AIVO Standard – Official GitHub Repository*. GitHub. Available at: <https://github.com/pjsheals/aivo-standard>
2. Sheals, P. (2025). *PSOS™ Methodology – Official GitHub Repository*. GitHub. Available at: <https://github.com/PJSheals/PSOS-Methodology>
3. Sheals, P. (2025). *CitationMapper – Official GitHub Repository*. GitHub. Available at: <https://github.com/PJSheals/CitationMapper>
4. Sheals, P. (2025). *AIVO Standard – Hugging Face Space*. Hugging Face. Available at: <https://huggingface.co/spaces/pjsheals/AIVO-Standard>
5. Sheals, P. (2025). *AIVO Standard Ecosystem v3.0*. Zenodo. DOI: <https://doi.org/10.5281/zenodo.17013901>
6. Sheals, P. (2025). *AIVO Methodology v3.0*. Zenodo. DOI: <https://doi.org/10.5281/zenodo.17077554>
7. Sheals, P. (2025). *PSOS™ Methodology v1.0*. Zenodo. DOI: <https://doi.org/10.5281/zenodo.17081529>
8. Sheals, P. (2025). *AIVO Standard – Prior Versions Archive*. Zenodo. DOI: <https://doi.org/10.5281/zenodo.16861660>
9. Sheals, P. (2025). *AIVO Mesh Lab*. Hugging Face. Available at: <https://huggingface.co/AIVOMeshLab>
10. Sheals, P. (2025). *CitationMapper Atom 01 – Gemini Dataset*. Hugging Face. Available at: <https://huggingface.co/datasets/AIVOMeshLab/citationmapper-atom01-gemini>
11. Sheals, P. (2025). *CitationMapper Atom 01 – Claude Dataset*. Hugging Face. Available at: <https://huggingface.co/datasets/AIVOMeshLab/citationmapper-atom01-claude>
12. Sheals, P. (2025). *CitationMapper Atom 01 – ChatGPT Dataset*. Hugging Face. Available at: <https://huggingface.co/datasets/AIVOMeshLab/citationmapper-atom01-chatgpt>
13. Sheals, P. (2025). *AIVO Standard Methodology (PDF)*. Internet Archive. Available at: <https://archive.org/details/aivo-standardtm-methodology.pdf>

14. Sheals, P. (2025). *Paul Sheals – ORCID Profile*. ORCID. Available at: <https://orcid.org/0009-0006-2407-4612>
15. Sheals, P. (2025). *Paul Sheals – Academia.edu Profile*. Academia.edu. Available at: <https://independent.academia.edu/PaulSheals>
16. Sheals, P. (2025). *The AIVO Standard™ : A Framework for AI Visibility Optimization (2025 Edition)*. Academia.edu. Available at: [https://www.academia.edu/143060805/The\\_AIVO\\_Standard\\_A\\_Framework\\_for\\_AI\\_Visibility\\_Optimization\\_2025\\_Edition](https://www.academia.edu/143060805/The_AIVO_Standard_A_Framework_for_AI_Visibility_Optimization_2025_Edition)
17. Sheals, P. (2025). *AIVO Standard Preprint*. SSRN. Available at: [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=5390001](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5390001)
18. Sheals, P. & de Rosen, T. (2025). *AIVO Journal*. Available at: <https://aivojournal.org>
19. Sheals, P. (2025). *Medium Articles*. Medium. Available at: [https://medium.com/@tim\\_62250](https://medium.com/@tim_62250)
20. Sheals, P. (2025). *Substack Archive*. Substack. Available at: <https://paulsheals.substack.com/publish/posts/published>