

Core banking AI visibility benchmark

Does AI know where you fit?

Part 1: A commercial guide for marketing and sales leaders

Your next business opportunity may begin with a question you never hear. A bank researching a modern technology platform can ask an AI assistant which vendors deserve consideration. The answer can become a valuable input that influences the buying process and ultimately, the decision.

We examined how 14 tracked core banking vendors appear across six AI engines and 15 buyer questions. Three vendors dominate the tracked mentions. Yet the results vary depending upon the AI platform used and the specialist needs evaluated. Those differences give marketing and sales a more useful and actionable starting point than what a single league table could.

This report explains where consideration may be constrained, what AI visibility patterns suggest about potential buyer questions and needs to validate, and what your team can investigate next and act upon. Part 2 turns those findings into a practical plan for content, competitive intelligence and sales conversations.

COMMERCIAL NEXT STEP

Use the benchmark to identify where buyers may be overlooking your product and its strengths. Understand where greater AI visibility could influence consideration.

AI visibility belongs in the pipeline conversation


A capable vendor can be overlooked before it has a chance to demonstrate its fit.

Consider a familiar sales situation: A bank enters the buying process with three vendors already in mind. Your team may have a solution that fits the bank's requirements, deep domain expertise and strong references, but if your platform was absent from the buyer's initial research, you may never get the opportunity to present your case. Better closing skills cannot recover an opportunity that never entered your CRM in the first place.

AI visibility analysis offers insights into one potential gap in the buyer's journey. It cannot tell you whether a missing mention caused a lost deal but it can reveal something strategically important: when buyers describe the problems you solve, are you amongst the options they discover, consider and compare?

Generative Engine Optimization (GEO) means improving the availability, readability and credibility of information that AI systems may train upon or perform search on when answering questions. For commercial teams, its practical purpose is to make their capabilities easier to find, understand and evaluate.

Sales brings the questions and objections heard from customers. Marketing examines whether relevant answers and supporting proof are available. Together they can distinguish a discovery gap from a product-fit, timing or sales-execution issue.



COMMERCIAL NEXT STEP

Compare the questions sales hears with the answers and proof marketing makes available.

Three vendors dominate tracked mentions

Temenos, Finastra and Finacle account for 89.4 percent of share of voice within the tracked set.

In the AI visibility analysis, Temenos has a mention rate of 83.3%, Finastra that of 52.2% and Finacle 42.7%. Each of the other 11 tracked vendors appears in 5.0% or fewer. A broad ranking can make the opportunity look closed to challengers, but specialist questions reveal a different picture.

Mention rate measures the percentage of sampled answers naming a brand. An answer can name several brands, so these percentages overlap. Share of voice divides mentions among the 14 tracked brands; it excludes other vendors and is not banking market share.

The commercial question is not whether your brand is well known, it is whether the evidence of the same appears when a buyer describes a need that you can serve. For a challenger, investigate the buyer questions where you have a credible reason to be considered. For a leader, check whether broad recognition extends to the specialist needs and engines your buyers use.

Vendor	Mention rate	Share of voice
Temenos	83.3%	41.8%
Finastra	52.2%	26.2%
Finacle	42.7%	21.4%
Other 11 combined	Individual rates <=5.0%	10.6%

COMMERCIAL NEXT STEP

Choose a relevant buyer need where your evidence gives you a credible reason to be considered.

Beyond Being Mentioned: The Strength of Your Presence

THE COMMERCIAL IMPLICATION

Mention rate tells you how often you appear. Average position tells you how prominently. Together, they reveal a different commercial problem: how effectively you are showing up.

Azentio appears in only 4.2% of answers, but averages position 3.6 when named. Silverlake Axis appears in 5.0% and averages position 8.1. These patterns suggest different investigations: broaden relevant coverage, or inspect the evidence supporting prominence.

Placement is conditional on being mentioned. A good average position from few appearances does not mean broad visibility. A lower position does mean appearance earlier in the answer but does not mean a wider coverage for all buyer relevant requirements.

Vendor	Mention rate	Average position
Temenos	83.3%	2.0
Finastra	52.2%	4.4
Finacle	42.7%	4.8
Silverlake Axis	5.0%	8.1
Intellect Design Arena	4.4%	7.5
Azentio	4.2%	3.6
Nucleus Software	2.6%	7.5
Musoni	2.0%	5.5
Craft Silicon	1.6%	7.6
Newgen Software	1.1%	5.3
Bankware Global	0.1%	7.5
SunTec Business Solutions	0.0% rounded	10.0
Profile Software	0.0% rounded	21.0
OneConnect	0.0%	Not applicable

Note: SunTec and Profile have recorded positions despite rounded zero mention rates. OneConnect has no recorded mentions at all.

COMMERCIAL NEXT STEP

Investigate coverage and prominence separately, then monitor comparable results.

A healthy average can hide platform gaps

THE COMMERCIAL IMPLICATION

Finastra appears in 88.6% of Claude answers but 21.4% of ChatGPT answers.

Finacle records 69.2% on Perplexity and 12.8% on Claude. Silverlake Axis records 29.7% on DeepSeek and 0.6% on Gemini, with zero recorded on the other four engines. A single average conceals important differences.

Inspect equivalent questions across engines. Content, retrieval and model behaviour can vary. One change will not necessarily affect every platform equally. Prioritise the engines your customers actually use.

Vendor	ChatGPT	Claude	DeepSeek	Gemini	Grok	Perplexity
Temenos	46.9%	98.1%	97.2%	80.3%	93.6%	83.6%
Finastra	21.4%	88.6%	65.6%	52.8%	35.0%	49.7%
Finacle	15.3%	12.8%	60.3%	70.0%	28.9%	69.2%
Azentio	4.7%	0.0%	5.8%	6.7%	1.1%	6.9%
Nucleus Software	5.6%	1.4%	8.9%	0.0%	0.0%	0.0%
Intellect Design Arena	3.9%	0.0%	3.1%	14.2%	0.0%	5.6%
Craft Silicon	0.0%	0.0%	6.1%	0.0%	3.3%	0.0%
Musoni	0.0%	0.0%	6.7%	2.8%	1.9%	0.6%
Newgen Software	0.0%	0.0%	0.0%	6.4%	0.0%	0.0%
Bankware Global	0.0%	0.0%	0.0%	0.6%	0.0%	0.0%
Profile Software	0.0%	0.0%	0.3%	0.0%	0.0%	0.0%
Silverlake Axis	0.0%	0.0%	29.7%	0.6%	0.0%	0.0%
SunTec Business Solutions	0.0%	0.0%	0.3%	0.0%	0.0%	0.0%
OneConnect	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

COMMERCIAL NEXT STEP

Compare answers and evidence by engine, make a focused correction and repeat the scans.

Go beyond the score

THE COMMERCIAL IMPLICATION

A visibility gap becomes useful when your team understands what to investigate.

Compare a strong answer and a weak answer to the same question across AI platforms. Which vendors appear? How are capabilities described? Which pages are cited? Check whether the answer explains your relevant products, implementation approach and approved customer proof.

Look for alternatives your team has not been monitoring. Being named in AI answers can introduce a competitor into the research conversation which may or may not be taking your market share.

Use the findings to form a testable hypothesis: insufficient content coverage, outdated descriptions, unclear positioning or a difference in retrieval. Turn your visibility data into actionable intelligence.

Marketing owns the evidence review. Sales checks the language, concerns and evaluation criteria with customers. Record what they confirm, then choose a focused correction and repeat comparable scans.

COMMERCIAL NEXT STEP

Give marketing an evidence task and sales a customer validation task.

Discovery and shortlisting are different

THE COMMERCIAL IMPLICATION

Each buying stage asks your team to demonstrate a different kind of fit.

Awareness asks which options exist. Problem questions ask who can solve a specific need. Consideration examines fit and integration. Decision asks who deserves a shortlist. Comparison evaluates alternatives. A vendor can appear in one context and disappear in another.

Azentio reaches 61.1% on the Islamic banking question but 0.7% on the future-proofing shortlist question. Finastra reaches 91.7% and 63.9% respectively. A strong specialist association does not automatically carry through to every evaluation.

Each stage in the buyer journey points to a different investigation. Marketing should provide category relevance, problem-solution proof and reasons for inclusion. Sales should validate what customers ask, what evidence matters and how they build their shortlists.

Question theme and stage	Azentio	Finacle	Finastra	Temenos
Cloud-native options Awareness	1.4%	37.5%	50.0%	88.2%
Islamic banking Problem	61.1%	29.9%	91.7%	77.8%
Ease of integration Consideration	0.0%	39.6%	29.9%	79.2%
Future-proofing shortlist Decision	0.7%	54.2%	63.9%	100.0%

*These are sample prompts, not actual bank procurement decisions.

COMMERCIAL NEXT STEP

Use AI shortlists to investigate visibility, then validate the implications with customers and sales.

Specialist questions reveal a different market

THE COMMERCIAL IMPLICATION

Azentio appears in 61.1% of answers to the Islamic banking question.

Azentio's overall mention rate is 4.2%. Musoni reaches 29.9% on microfinance and SACCOs, while Craft Silicon reaches 23.6%; their overall rates are 2.0% and 1.6%. These are distinct specialist footholds, not equivalent performance across the market. A low broad score can obscure a strong association with specific buyer needs.

Temenos records 44.4% on regulatory compliance and AML/KYC, while Finacle records 1.4%. The right opportunity depends on product fit and customer relevance, rather than the size of a gap alone.

A specialist association creates an opening. Buyers still need implementation, integration, geographic and support evidence. Validate the need with customers before building a campaign around it.

Specialist question	Selected vendor	Mention rate
Islamic banking	Finastra	91.7%
Islamic banking	Azentio	61.1%
Microfinance and SACCOs	Musoni	29.9%
Microfinance and SACCOs	Craft Silicon	23.6%
Compliance and AML/KYC	Temenos	44.4%
Compliance and AML/KYC	Finacle	1.4%

COMMERCIAL NEXT STEP

Validate the specialist need before building a campaign around it.

Your monitored peer group is not the whole market

THE COMMERCIAL IMPLICATION

AI answers surface alternatives beyond the 14 tracked vendors.

The tracker surfaces global platforms, cloud-native providers and specialists. These names deserve investigation even if they are absent from your usual competitive list. Read the context before deciding whether each is a substitute, a component or a partner.

The field changes as information on the web is updated and product names evolve. Use a consistent brand and product grouping method when comparing. Tracked share of voice cannot describe the entire competitive landscape.

Build fair comparisons and objection handling around verified fit. Ask buyers which approaches they are considering and why. This gives sales a way to learn about alternatives without assuming that every AI mention signals a lost opportunity.

Alternatives surfaced by the tracker	What to investigate
FIS, Oracle and Fiserv	Platform scope and relevant market coverage
Mambu, Thought Machine and 10x Banking	Cloud-native approach and implementation fit
Backbase, Alkami and nCino	Which layer of the buyer problem they address

COMMERCIAL NEXT STEP

Refresh the competitive set and prepare fair comparisons around verified buyer fit.

Buyers encounter a wider evidence footprint

THE COMMERCIAL IMPLICATION

Your website is one part of the information AI answers may draw on.

The tracker includes analyst sources, LinkedIn, directories, content sites and document-sharing platforms. This gives marketing a practical research task: find the exact cited pages and assess the claims behind them.

Independent coverage can add useful context, but frequent citation does not establish editorial quality or prove that a placement will improve visibility. Assess relevance, accuracy and update dates before investing.

Prioritise credible sources for your audience. Correct outdated profiles, maintain approved product descriptions and make supporting evidence easy to evaluate. Bring relevant independent proof into sales materials where its use is permitted.

Third-party domains appearing in the tracker	Source category
gartner.com	Analyst and research
linkedin.com	Professional network
openbankingtracker.com and f6s.com	Directories and marketplaces
finantrix.com	Specialist content
worldmetrics.org, zipdo.co and gitnux.org	Statistics and list sites
scribd.com	Document sharing

COMMERCIAL NEXT STEP

Review exact cited pages before directing content, partnership or coverage effort.

Use the benchmark with the right boundaries

THE COMMERCIAL IMPLICATION

Discoverability Engine and Octomise Monthly Tracker | South East Asia averages

The benchmark covers 14 core banking vendors across ChatGPT, Claude, DeepSeek, Gemini, Grok and Perplexity. Fifteen prompts span awareness, problem, consideration, decision and comparison, including specialist banking needs.

Mention rate: the percentage of sampled answers naming a brand. Average position: its mean placement only in answers where it is named; lower is earlier. Share of voice: its share of mentions within the tracked group, not market share.

Visibility score summarises presence and prominence. Interpret it with mention rate and placement. A citation is a source reference in an answer; it is not an endorsement of the source or a verified influence measure.

The prompts simulate evaluation questions. They do not identify named customers, private AI conversations or demand volume. This October edition presents a baseline, rather than evidence of a month-to-month trend.

For a meaningful follow-up, agree collection dates, answer totals, engine settings and brand grouping. Repeat a comparable scope and validate commercial interpretations through customer conversations and opportunity reviews.

COMMERCIAL NEXT STEP

Agree a comparable question set, measurement period and definition of progress.

Turn AI visibility gaps into commercial opportunity

THE COMMERCIAL IMPLICATION

Schedule a 30-minute strategy briefing and receive Part 2 when your meeting is booked.

You now have a clearer way to recognise gaps in category awareness, engine consistency, shortlist inclusion and specialist association. The next decision is which gap matters to the customers you want to serve.

Part 2, the Core banking AI visibility strategy playbook, provides a prioritisation framework, customer discovery questions, evidence planning, an account engagement workflow and a 90-day programme. It adds practical tools to this benchmark.

In the briefing, connect the findings to your buyer segments and markets. Choose an evidence question and discuss how Discoverability Engine insights and Octomise account context can guide a relevant next action.

Adam Tan | Octomise (Sabio World)
adam@discoverabilityengine.ai
www.octomise.com

Itika Singhal | Discoverability Engine
itika@discoverabilityengine.ai
www.discoverabilityengine.ai

COMMERCIAL NEXT STEP

Book a strategy briefing to receive Part 2 and choose the next commercial question.