



The CMO's AI briefing for IT services companies

What is changing, what to do about it, and how to avoid the mistakes. 2026 edition.

SUMMARY

Ten things to know

A one-page summary of the full guide at search-value.com/ai-marketing-for-it-services-cmos.

1. Buyers have moved. About half of B2B software buyers start research with an AI chatbot more often than with Google (G2, March 2026). Your first impression is increasingly an AI answer.
2. Influence is bigger than traffic. AI search is under 2% of B2B referral visits but touches about 1 in 11 closed-won deals (Octane11). Absence from the answer is absence from the shortlist.
3. Clicks are thinner. 68% of US Google searches ended without a click in early 2026 (SparkToro). Measure visibility, not only visits.
4. Buyers verify. They start with AI, check in search, then validate with people. 69% turn to reps to validate AI insights (Gartner). Your site, reviews and sales team must tell one story.
5. Your clients are repricing you. Tech services buyers expect productivity gains and outcome-based pricing (Bain, 2026). Marketing must prove outcomes, not capacity.
6. Proof beats claims. Numbered case studies, named experts and third-party reviews are what AI systems quote and what buyers verify.
7. The upside is real. Research, drafting, localisation, reporting and testing are faster with AI when experts stay in the loop.
8. The downside is real. Generic content, leaked client data, agent washing, misrepresentation and new transparency rules all carry cost.
9. Foundations decide results. Only 30% of CMOs report mature AI readiness; 70% say processes are not mature enough to scale (Gartner 2026). Data, process, people and policy come before tools.
10. Start small and measure. Three to five use cases, a one-page policy, a monthly scorecard and a 90-day review beat a large programme without an owner.

Three decisions for this quarter

1. **Who owns AI in marketing?** One named person, with authority over tools, policy and the scorecard.
2. **Which three to five use cases first?** Pick from the use-case page: high value, low risk, easy to measure.
3. **What will we report, and when?** A one-page scorecard each month and a board review at day 90.

What to ignore for now

Promises of fully autonomous marketing, guaranteed AI rankings, and tool purchases made before a policy, an owner and a baseline exist.

Fifteen figures, with sources

The evidence behind the summary, in one table you can lift into a board paper.

WHAT	FIGURE	SOURCE
Buyers who start with an AI chatbot more often than Google	51%	G2, Mar 2026, n=1,076
Buyers who rely on AI chatbots in research	71%	G2, Mar 2026
Buyers who changed vendor because of a chatbot	69%	G2, Mar 2026
Buyers who used generative AI in a recent purchase	45%	Gartner, n=645, fielded Aug to Sep 2025
Buyers who validate AI insights with a sales rep	69%	Gartner, May 2026
Buyers who prefer a rep-free experience	67%	Gartner, Mar 2026
US Google searches ending without a click	68%	SparkToro / Similarweb, Jan to Apr 2026
CTR drop for the top result when an AI Overview shows	58%	Ahrefs, 300,000 keywords, Dec 2025 data
AI search share of B2B referral traffic	Under 2%	Octane11, 400M sessions, to Jun 2026
Closed-won deals with an AI search touch	About 1 in 11	Octane11, same dataset
AI Overviews / AI Mode / Gemini app monthly users	2.5B+ / 1B+ / 900M+	Google I/O, May 2026
Share of marketing budget CMOs put into AI	15.3% (21.3% if AI-ready)	Gartner CMO Spend Survey 2026, n=401
CMOs reporting mature AI readiness	30%	Gartner 2026
Agentic AI projects expected to be cancelled by end 2027	Over 40%	Gartner forecast, June 2025
Tech services buyers expecting 5% to 10% of tech spend on AI/ML	About 75%	Bain, Jan 2026, n=280

How to read these numbers

Surveys describe the samples named. Most buyer data is from software and B2B purchasing studies, so treat it as direction for IT services, where purchases are larger and slower. Referral share and influence are different things: AI search sends few visits but shapes many shortlists. Forecasts are forecasts, not measurements.

Re-run the buyer-question test every month. Your own data on which assistants name you, and how, will matter more than any industry average.

The vocabulary, in plain language

Thirty-two terms your board, agency and vendors will use.

Generative AI. Software that creates text, images, code or audio from a prompt.

LLM (large language model). The model behind assistants like ChatGPT, Claude and Gemini. It predicts useful text from patterns.

AI Overviews. Google answers shown above results, with links to sources.

AI Mode. Google's fuller conversational search experience.

Answer engine. A search tool that answers directly instead of listing links.

GEO. Generative engine optimisation: making content easy for AI to cite.

AEO. Answer engine optimisation. Close to GEO; focuses on being the answer.

LLMO. LLM optimisation. Another name for the same goal.

Zero-click search. A search that ends without a visit to any site.

Citation. A named source shown with an AI answer.

Grounding. Tying an AI answer to retrieved sources rather than memory alone.

RAG. Retrieval-augmented generation: the model looks up your documents before it answers.

Hallucination. A confident answer that is false or invented.

Prompt. The instruction given to a model.

Context window. How much text a model can consider at once.

Token. A chunk of text a model reads or writes. Usage is billed in tokens.

Fine-tuning. Further training of a model on your own examples.

Agent / agentic AI. AI that takes multi-step actions toward a goal, using tools.

Agent washing. Relabelling ordinary chatbots or automation as agents.

MCP. Model Context Protocol: an open standard for connecting assistants to tools and data.

AI SDR. AI that researches prospects and writes or sends outreach.

Copilot. An assistant built into a work tool to help the user.

Multimodal. A model that handles text, images, audio and video.

AI slop. Generic, low-effort AI content published at scale.

Scaled content abuse. Google's name for mass-producing pages mainly to manipulate rankings.

Shadow AI. Staff using unapproved AI tools with company or client data.

Human in the loop. A person reviews or approves AI output before use.

AI governance. Policies, owners and controls for how AI is used.

Content provenance. Labels or metadata showing how content was made or edited.

llms.txt. A proposed file pointing AI systems at key pages. Platform support unconfirmed.

Outcome-based pricing. Fees tied to results rather than hours or headcount.

AI deflation. Falling prices for services as AI cuts the effort to deliver them.

Use cases: start, pilot, wait, never

Value is the likely gain in time or reach. Risk covers accuracy, confidentiality, brand and law. Verdicts reflect today's evidence and your readiness.

USE CASE	VALUE	RISK	VERDICT	CONTROL
Research and synthesis	High	Low	Start	Verify figures against sources
Drafts from expert interviews	High	Medium	Start	Expert and editor review
Repurposing long content	High	Low	Start	Check claims survive shortening
Meeting and call summaries	Medium	Low	Start	Confidentiality check
SEO and GEO operations	High	Medium	Start with audits	No scaled pages made only to rank
Localisation	High	Medium	Pilot	Native-speaker review
Reporting and anomaly alerts	Medium	Low	Start	Analyst confirms numbers
Ad and creative variants	Medium	Medium	Pilot	Brand and compliance review
ABM personalisation	High	Medium	Pilot	Honest, specific, volume limits
RFP and proposal first drafts	High	Medium	Pilot	Pricing and confidentiality review
Outbound email at scale	Medium	High	Wait	Deliverability and consent first
Autonomous agents publishing or sending	Unproven	High	Wait	Only with logs, limits and a kill switch
AI-generated testimonials or case studies	None	High	Never	Real clients, real results only

The test for any use case. If a wrong answer would embarrass a client or breach a contract, a person must review it. If it only costs a little time, automate more freely.

Where the value shows

Look for time saved on research, drafts and reporting; more tests per month; faster localisation; and better coverage of buyer questions. Measure each against the baseline you record in the first 30 days.

Ten risks to own before they happen

RISK	LIKELIHOOD	IMPACT	OWNER	CONTROL
Wrong or invented facts in published content	High	High	Content lead	Expert review; source list; correction log
Generic content that damages brand and ranking	High	Medium	Head of content	Fewer, better pages; named authors
Client data entered into public tools	Medium	High	CMO and CISO	Approved tools; data rules; training
AI answers describe the firm wrongly	Medium	Medium	SEO lead	Monthly test; consistent public profile
Agent project fails or is mislabelled	High	Medium	CMO	Pilot on own data; exit criteria
Deliverability damage from AI outreach	Medium	High	Demand gen lead	Small sends; authentication; list hygiene
Disclosure or transparency breach (EU AI Act)	Low	High	Legal	Label synthetic media and chatbots; log use
Copyright or licence issues in assets	Medium	Medium	Legal	Licensed tools; review of images and text
Measurement blind spots	High	Medium	Analytics lead	AI answer test; self-reported source field
Tool sprawl and unclear return	High	Medium	CMO	One owner per tool; quarterly cull

Likelihood and impact are our starting view for a mid-to-large IT services firm. Adjust to your clients, markets and contracts. Regulation: EU AI Act transparency obligations apply from 2 August 2026, with some machine-readable marking duties for existing systems delayed to 2 December 2026, and guidance still in draft. This briefing is not legal advice.

What to do, in order

Days 1 to 30: baseline and rules	Days 31 to 60: pilot and fix	Days 61 to 90: decide and scale
☐ Name one owner for AI in marketing ☐ Audit current AI use, including shadow tools ☐ Write and circulate the AI use policy ☐ Choose 3 to 5 low-risk use cases ☐ Record baselines: hours per asset, leads sales accepts ☐ Test 20 to 30 buyer questions in the main assistants ☐ List facts and proof to fix on the public web	☐ Run each use case with an owner and review step ☐ Update service pages with answer-first summaries ☐ Publish two case studies with figures ☐ Align LinkedIn, directories and review profiles ☐ Start the monthly AI answer test ☐ Train the team on the policy ☐ Review security and contract terms for tools	☐ Compare results with baselines ☐ Stop what did not pay back ☐ Write the second wave of use cases ☐ Report to the board on one page ☐ Fold AI into the annual marketing plan ☐ Set quarterly policy and tool reviews ☐ Agree next-quarter budget by use case

Owner. _____ **Start date.** _____ **Board review.** _____

Monthly scorecard

- Share of AI answers that name you (20 to 30 buyer questions, main assistants)
- Lead-to-SAL ratio: the share of leads sales accepts
- Branded and direct search trend
- Self-reported source ("how did you hear about us")
- Cost and time per asset
- Errors and corrections after review

Fifteen questions for any AI vendor or agency

- 1 Can you show this working on our data, not your examples?
- 2 What happens to the data we enter? Is it stored, and is it used to train models?
- 3 Which model or models do you use, and can they change without notice?
- 4 Which steps are truly automated, and which are scripts or people?
- 5 How is output reviewed, and what is your measured error rate?
- 6 How do you handle client confidentiality and personal data?
- 7 Which security standards and audits can you show?
- 8 What logs and approvals exist for anything the system publishes or sends?
- 9 How will we measure success, and what baseline do we agree first?
- 10 What is the full cost, including usage, seats and services, at our volume?
- 11 Who owns the output, and what are the licence terms?
- 12 What are the limits, failure modes and known weak cases?
- 13 Can we export our data and prompts if we leave, and how long does exit take?
- 14 Which comparable firms use it, and may we speak to them?
- 15 What would you advise us not to use this for?

Red flags

- Demos only on the vendor's own examples
- "Agent" with no description of the steps or limits
- No answer on whether your data trains their models
- Success measured in activity (content produced), not results
- Promised rankings or "guaranteed" AI mentions
- Pressure to sign before a pilot
- No human review step for customer-facing output
- Unwillingness to name references

How to run a pilot

Choose one use case. Fix the baseline and the success measure first. Run for 30 to 60 days on your data with a named owner and a review step. Decide on results, not on enthusiasm.

An AI use policy starter

Purpose. To use AI to improve marketing quality and speed without harming clients, buyers or our reputation.

Approved tools. _____ (review quarterly)

Never enter. Client names or confidential material, personal data, unreleased pricing, credentials, and contract terms.

Review. A person reviews every customer-facing output. Facts and figures are checked against a source. Experts approve technical claims.

Disclosure. We label synthetic media and chatbots where required by law or by client agreement.

Records. We keep a log of where AI was used in published work.

Owner. _____ **Date.** _____

Eight sentences for the board

- 1 AI has changed where our buyers start: they ask an assistant, and a third of them buy from firms they had not heard of before it named them.
- 2 Our clients expect AI productivity to show up in price and outcomes, so our marketing now leads with measured results.
- 3 We are not chasing tools. We are fixing data, proof and process first, because that is what separates leaders from laggards.
- 4 We will use AI for research, drafts, localisation and reporting, with a person reviewing anything a client or buyer sees.
- 5 We have a one-page AI policy, an owner, and a scorecard.
- 6 We test what AI assistants say about us every month and correct what we can.
- 7 We will pilot before we buy, and stop what does not pay back.
- 8 Next quarter we report: leads sales accepts, AI visibility, cost and time per asset, and risks.

Want a second pair of eyes? Book a free 30-minute pipeline review. We will test what AI assistants say about your firm, look at your funnel, and show you what we would change first. info@search-value.com | search-value.com/contact

SOURCES

Where the numbers come from

All accessed October 2026. Full links are in the online guide at search-value.com/ai-marketing-for-it-services-cmos.

- Google, I/O 2026 keynote summary (May 2026): AI Overviews, AI Mode and Gemini app users.
- G2, AI search insight report (March 2026; 1,076 B2B decision makers).
- Gartner, press release on buyers validating AI insights with reps (May 2026; 645 buyers).
- Gartner, 67% of B2B buyers prefer a rep-free experience (March 2026).
- Gartner, 2026 CMO Spend Survey (May 2026; 401 CMOs).
- Gartner agentic AI forecast, reported by MarTech (forecast of June 2025).
- Gartner, search engine volume forecast (February 2024; a forecast).
- eMarketer (September 2026), citing Semrush data from March to April 2026.
- SparkToro zero-click analysis, via Techwyse (June 2026).
- Ahrefs, AI Overviews and click-through (February 2026).
- Octane11, State of B2B AI Search Q2 2026 (July 2026).
- Bain, tech services buyer survey (April 2026; 280 decision makers).
- The Register, Indian IT FY26 results (April 2026).
- Outlook Business, outcome-based deals in IT services (analyst estimates from Jefferies).
- Contently, top sources LLMs cite (April 2026).
- PPC Land, Google spam policies and AI Overviews (Google documentation update, 15 May 2026).
- Baker McKenzie, EU AI transparency from 2 August 2026.

About Search Value. A B2B marketing partner for IT services and technology companies since 2007. SEO and GEO, LinkedIn and Google advertising, inbound and lead nurturing, and analytics, reported as one number: the share of leads sales accepts.