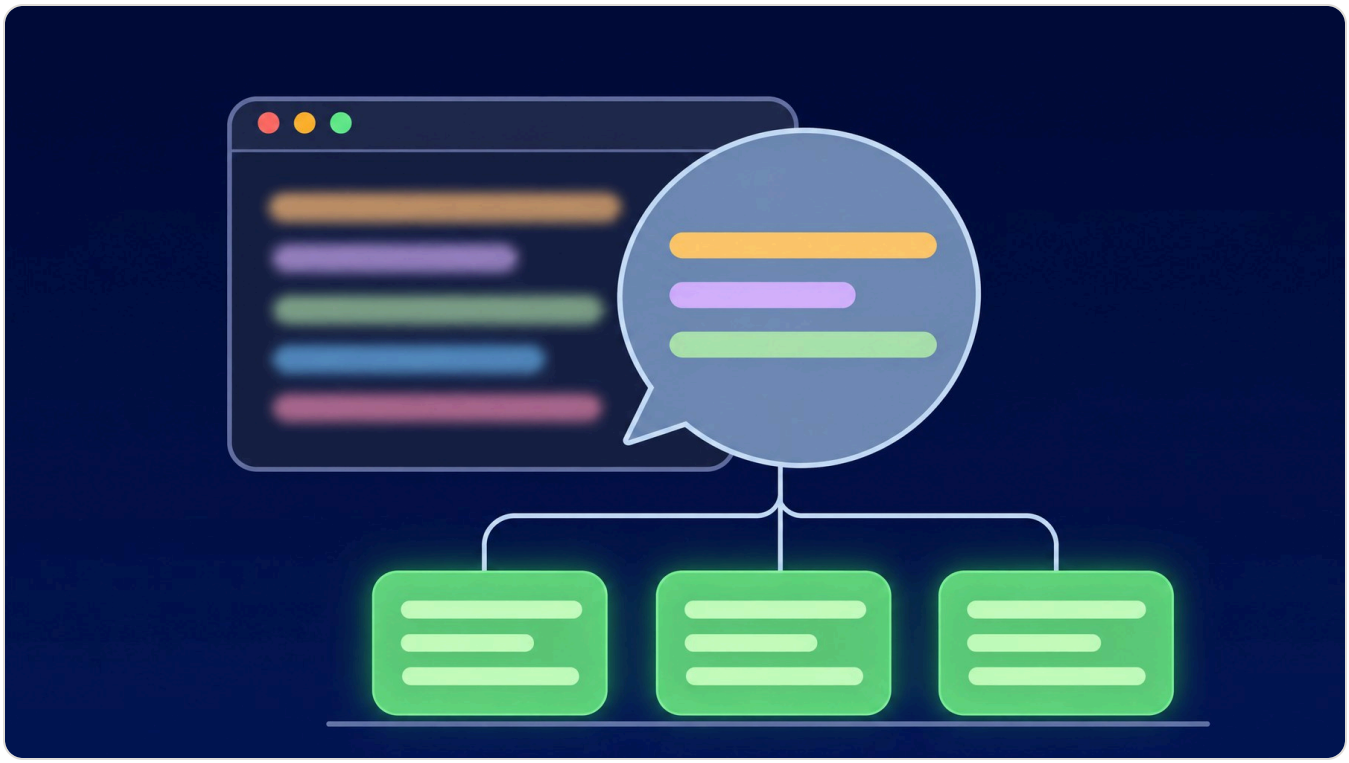


REPORT

The web has a new reader

What the speakers at Sitecore Symposium 2026 said and showed about AI assistants that read the web for buyers, narrow the choice and sometimes make the purchase.

Frank Costoya · Orlando, November 10 to 12, 2026

29% → 51%

of B2B software buyers start their research in an AI assistant more often than in Google, April 2025 to March 2026.

Tim Sanders, G2. G2's surveys of software buyers.

0.1% → 7%

the click-through rate when an answer cites a brand, against when it recommends the brand.

Tim Sanders, G2, on stage.

20 → 3 to 10

pages an assistant gathers for one question, and the pages it then reads.

Kiran Patil and Khushboo Sorthiya, Horizontal Digital.

01

The new reader

Speakers on both days of sessions said buying now starts inside an AI assistant, where a buyer learns, compares and chooses in one conversation. They showed fewer clicks from search, sessions that never reach the open web, bots that account for half of page views, and an agent that paid for a purchase.

> **The sessions.** Eric Stine, Sitecore, mainstage keynote. Michelle Boockoff-Bajdek, Sitecore, mainstage keynote. Both on Wednesday, November 11. Chris Andrew, Scrunch, mainstage. Nilesh Tamhane and a colleague, XCentium, breakout. Kiran Patil and Khushboo Sorthiya, Horizontal Digital, breakout. Kevin Schofield, Velir x Brooklyn Data, panel. Andrew Michael, Scrunch, breakout. All on Thursday, November 12.

The journey starts in a conversation

Eric Stine, CEO of Sitecore, the company whose conference this was, told a story of his own in the opening keynote. In August 2026 he volunteered to coach his son's flag football team and asked ChatGPT how to run a practice. The conversation led him to playbooks, wristbands, cones and flags from brands he had no idea were in that business. He said he spent more than $1,500 without ever setting out to buy anything.

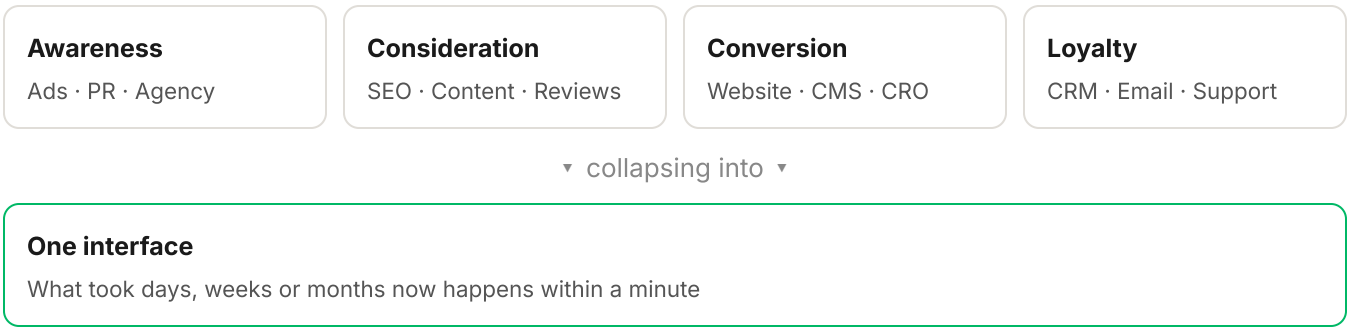
His lesson was that awareness, consideration and intent have been "scrunched together" into one conversation, and that a brand which had the right answer but never appeared in it was never considered at all.

Michelle Boockoff-Bajdek, Sitecore's Chief Marketing Officer, described shopping for a mortgage in a conversation with ChatGPT. For the bank's marketer, she said, none of it was visible: no ad click, no page view, no form.

A funnel in one window

Andrew Michael, VP of product at Scrunch, a company that tracks how brands appear in AI answers, pictured marketing as a set of stages. His claim is that they are collapsing into a single interface, where what took days, weeks or months now happens within a minute: the assistant reads the sources, compares the products and in some cases makes the purchase.

Four stages, one interface · the funnel on Scrunch's slide, and Andrew Michael's spoken claim

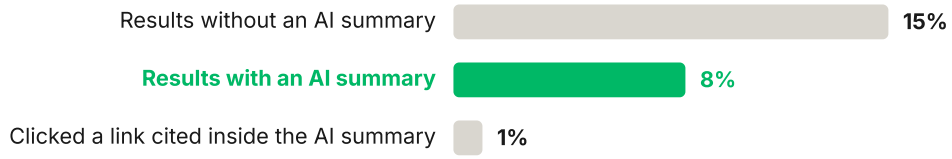


Source. The four stages and the three tools under each are from the slide shown by Andrew Michael, Scrunch. The line under "One interface" is what he said on stage. SEO is search engine optimization. A CMS is a content management system. CRO is conversion rate optimization: work to raise the share of visitors who buy or sign up. CRM is customer relationship management software.

Fewer clicks from search

Kiran Patil and Khushboo Sorthiya of Horizontal Digital, a digital agency, showed Pew Research Center's count of what an AI summary does to clicks on Google. The count rests on 900 U.S. adults and 68,879 of their searches: on how many visits to Google did the person click a link? The slide's line is that buyers have not stopped researching, they have stopped clicking.

Fewer people click · share of Google visits where the user clicked a link



Source. Horizontal Digital's slide, which credits Pew Research Center, July 2025, 900 U.S. adults and 68,879 searches. The slide's own label for the three is "Share of Google visits where the user clicked a link".

Nilesh Tamhane and a colleague from XCentium, a Sitecore partner, showed a slide on the same subject headed "Search is changing". One of its cards says 60% of searches end without a click. Another is headed "58% drop in organic traffic", which means visits that arrive from unpaid search results. Its small print credits the 58% to Ahrefs as a drop in CTR once an AI Overview, Google's AI summary, appears. CTR is click-through rate: the share of people shown a link who click it.

"Search is changing" · two cards from XCentium's slide, in the slide's words

60%

of searches end without a click

The answer comes straight from AI. No visit to your site. No chance to convert.

FORBES

58%

drop in organic traffic

Gartner predicts 25% drop by 2026, rising to 50%+ by 2028. Ahrefs already found 58% drop in CTR once AI Overview appears.

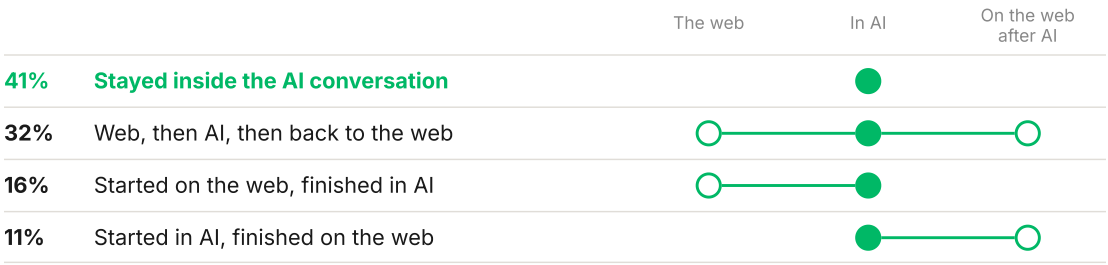
AHREFS

Source. XCentium's slide, shown by Nilesh Tamhane and a colleague. The 60% card credits Forbes. The 58% card credits Ahrefs, and its forecast of a 25% drop by 2026, rising to more than 50% by 2028, sits under the card's heading of organic traffic and is credited to Gartner. Each card carried the logo of the source named under it.

Sessions that never leave the chat

Boockoff-Bajdek said on Wednesday that more than 41% of AI sessions happen entirely inside the AI platform, with no open-web search before and no click to the open web after. On Thursday Chris Andrew, Scrunch's co-founder and CEO, showed where the rest go. His slide drew four paths across three columns: the web, the AI assistant, and the web again afterwards. His conclusion was that these journeys are condensing, and that the website has never been more important.

Where AI sessions go · share of all AI sessions in the study, by path



Source. More than 41%: Michelle Boockoff-Bajdek, Sitecore, citing Scrunch. The shares and the four paths: Chris Andrew's slide, from a new Scrunch data study, with each path worded as he described it. He said the study is based on millions of consumer prompts from real interactions. Neither he nor the slide said how the study defines a session.

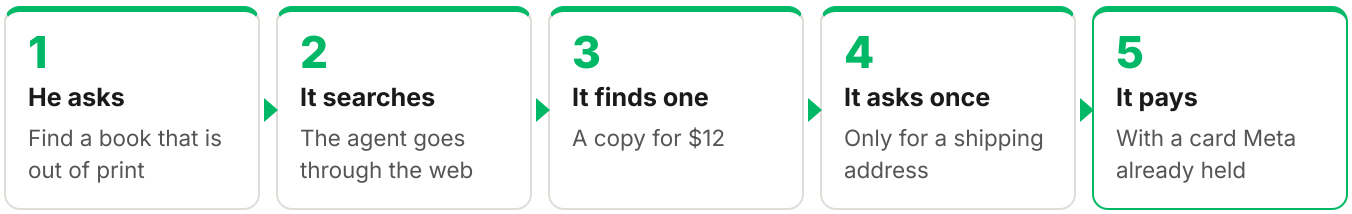
A third audience

Kiran Patil said sites were built for two audiences, people and search engines, and that a third has arrived: the programs that fetch pages for AI systems. Kevin Schofield of Velir x Brooklyn Data cited Cloudflare as saying that in June 51% of all page views and hits came from bots and not from people.

Agents that finish the job

Chris Andrew described handing purchases to Meta's Muse agent. Asked to find an out-of-print book, it searched the web, found a $12 copy, asked only for a shipping address and paid with a card Meta already held.

One purchase, handed to an agent · the steps as Chris Andrew described them on stage



Source. Chris Andrew, Scrunch, mainstage. No slide showed the steps.

Tickets, and a $500 deposit on a car the agent had recommended, followed. His sentence for it: he expressed intent and the agent delivered an outcome. Of the tickets he said he did not visit a website.

02

Found is not chosen

Tim Sanders, Chief Innovation Officer at G2, opened the second day with a research talk on how AI assistants decide which products to recommend. He said buyers now reach a shortlist of three inside a single prompt, and he defined winning as being one of the three.

> **The sessions.** Tim Sanders, G2, mainstage keynote and the marketing leaders panel. Nilesh Tamhane and a colleague, XCentium, breakout. Andrew Michael, Scrunch, breakout. All on Thursday, November 12.

Buyers now start in chat

G2, a software review site, asked B2B software buyers where they start when they research a purchase. In April 2025, 29% said they start with AI search more often than with Google. In March 2026 the answer was 51%. He said the number had jumped 76% in a year (51 is about three quarters more than 29).

Where software buyers start · share who start with AI search more often than Google



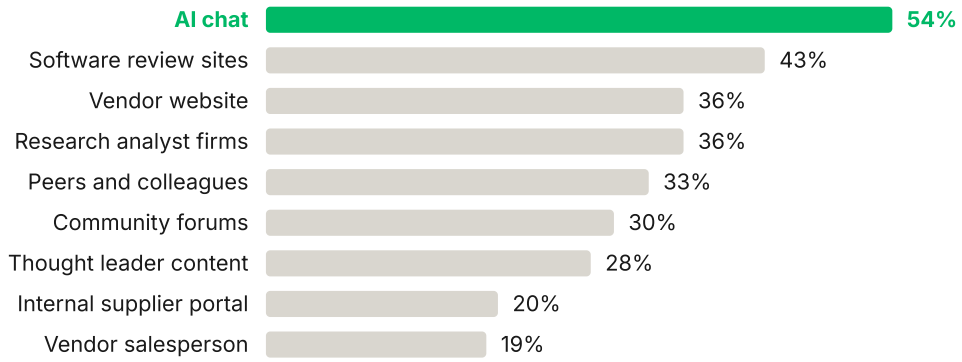
Source. Tim Sanders, G2. The slide is from G2's surveys of B2B software buyers, April 2025 and March 2026.

Sanders said the question set a Google property against ChatGPT, Gemini, Perplexity, Claude or DeepSeek, so the AI Overviews above Google's results do not count toward the 51%. He said G2 ran the March survey twice because he did not believe the first result. On stage he added two shares that are not on the slide: 58% of large enterprise buyers and 65% of IT professionals start in chat.

The shortlist is made in one prompt

The more important change, he said, is that buyers have handed the gatekeeper's job to the assistant. G2 has studied for about a decade how a buying committee reaches its shortlist of three, and he put that work at 6 to 15 working hours: search results, analyst reports, peers, reviews, a spreadsheet. Buyers now get there in one prompt, usually without realizing it. They describe the job to be done and their constraints, and the model returns three options with caveats.

What puts a vendor on the shortlist · share of buyers surveyed who named each source

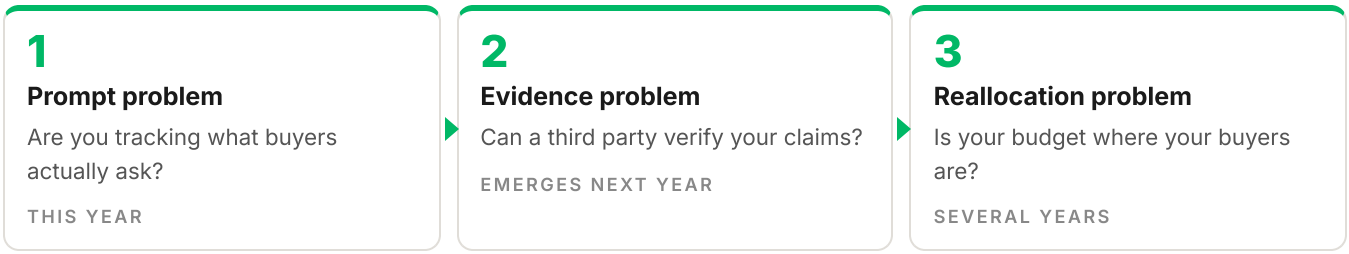


Source. Tim Sanders's slide, from G2's survey of March 2026. A buyer could name more than one source, so the shares add up to more than 100%.

He added a figure he attributed to 6sense, a company that does similar research: 93% of buyers decide from the shortlist they had on day one. The job, he said, is to win the answer, which he defined as being one of the three recommended. Showing up or being cited is not the same thing.

Three problems

Sanders built the rest of the talk on three problems, each with a question on his slide and a time horizon he gave on stage.



Source. The names and questions are Tim Sanders's slide. The horizons are as he gave them on stage.

The prompt problem

He said AI visibility reporting runs on a prompt panel, the set of prompts a brand tracks, and that panels are usually built by turning keywords, category names and the brand's own product language into questions, so they are guesses, and model vendors do not release real prompts. The worst effect he named is a false positive: a brand that tracks prompts in its own vocabulary sees a very high citation share of voice, meaning it is cited in a large share of the answers to prompts it wrote itself.

His example was two competitors in one software category. One was promoting a new category name, AI visual collaboration software, while its traffic and pipeline fell. The other was optimizing for the word "whiteboards", because its business development reps had been collecting real prompts from prospects and that was the word buyers used.

What a real prompt asks for

Agents, Sanders said, are loyal to the prompt: they break it into components and treat each as a box to check before a brand is a candidate. His slide set a Google search beside what he said was a real captured prompt, with three phrases of the prompt highlighted.

A search and a captured prompt · Tim Sanders's slide, redrawn with its wording and highlights

GOOGLE · SHORT-TAIL

Best CRM for hospitals

CHAT PROMPT · LONG-TAIL

Looking for affordable solution for my hospital to manage **patient scheduling, messaging and reimbursement processing** 1. It must **work as well on a mobile device** 2 as it does desktop as many of our staffers are also medical providers. **Exclude vendors requiring a one year or longer contract** 3, as per our procurement.

▾ each highlighted phrase becomes a box the agent checks ▾

THE AGENT'S CHECKLIST · BEFORE A VENDOR IS A CANDIDATE, AS SANDERS DESCRIBED IT

1 **Patient scheduling, messaging and reimbursement processing** Does the vendor's documentation cover each?	2 **Works as well on a mobile device** Does it say so?	3 **No contract of one year or longer** Can the agent see the terms?

Source. The search and the captured prompt are side by side on Tim Sanders's slide, which highlights the three phrases. The checklist is drawn from what he said on stage: an agent treats each as a box to check.

He said each highlighted requirement has to be answered by the vendor's own product documentation, or the vendor is not on the long list. A vendor whose site says "hospitals" 132 times but documents neither reimbursement speed nor mobile use may be left out, he said. Later he pictured the long list for this prompt as 14 vendors that had checked every box.

In thousands of prompts captured from his customers' buyers, the job and its constraints made up almost 60% of the words in a prompt, and the category name (here, "CRM") appeared in only two prompts of every three.

He made one claim about pricing. Citing a study he attributed to Boston University and Northwestern, he said models are tuned to recognize patterns of manipulation, and that hiding basic terms such as price or contract length behind a talk-to-sales form signals that risk and can get a vendor left out.

His fix is to weight real buyer language above invented prompts. The gold standard he named is prompt capture: asking a prospect who arrived through an assistant for the prompt they used.

Mention, citation, recommendation

Sanders opened the second problem with one line: a company whose pipeline is suffering does not have an AI visibility problem, it has an evidence problem. He then ranked the ways a brand can appear in an answer, and gave three of them a click-through rate: of the people who are shown the answer, the share who click through to the brand's site.

HOW THE BRAND APPEARS	WHAT IT IS	CLICK-THROUGH RATE HE GAVE
Mention	The brand's name appears in the answer	No rate given. Good for the brand, he said, and not connected to revenue

Citation	A link in the answer's source box	**About 0.1%** about one reader in a thousand clicks
Branded citation	The brand's name is a link in the answer itself	**0.33 to 0.5%** "a third to a half a percent"
Recommendation	The brand is one of the three options returned	**Closer to 7%** "70X" the rate of a citation, and "like SERP": close to that of a search results page

Source. Tim Sanders, G2, on stage. He credited the citation rate to one researcher's click-through research and called it an average. The bars are drawn to one scale.

His reason for the low rate on a citation was that the whole point of AI search, for the person using it, is not having to click. He compared counting citations to counting website hits in 2001, citing a September study by Jeff Oxford: over thousands of runs of commercial-intent prompts, a brand that was cited was still left off the list of recommendations 60% of the time.

XCentium's presenters separated the first two outcomes that afternoon. Of 100 people asking a question the brand wants to answer, Nilesh Tamhane said, 16 hear its name and 6 get a link to its site. The other 94 get no path to the site.

> **Source.** Nilesh Tamhane, XCentium, on stage. XCentium's slide reads "16% mention your brand" and "6% of those link to you", and credits a Webflow analysis of 2,000 U.S. company sites.

One question, two kinds of answer

Andrew Michael, VP of product at Scrunch, a company that tracks how brands appear in AI answers, said Sanders had been right about this distinction and that Scrunch sees it in its data. His slide shows one question answered two ways.

THE QUESTION "Best tent for a family trip to Yosemite in May?"

MENTIONED

Popular choices include TrailPeak Summit, CampCo Family 6, **Aurora Ridgeline 6** and Northfold Dome.

RECOMMENDED

I'd go with the Aurora Ridgeline 6. It pitches in four minutes and handles spring rain.

MENTIONED	CITED	RECOMMENDATION RATE
61%	**22%**	**12%**
says your name	links to you	about 2× the work of a mention

Source. Scrunch's slide "Recommended is what matters", shown by Andrew Michael, redrawn. The brands are examples. The three figures are examples too, and the slide does not say what they are shares of.

Evidence that matches the claim

Sanders gave a three-step model of how an assistant reaches a recommendation. He said it comes from academic studies that record a model's reasoning across many runs.

How an LLM chooses what to recommend · Tim Sanders's slide, redrawn, with his remarks

1
Find viable candidates
A long list the user never sees: each vendor's own site and documentation, checked against the prompt.

2
Validate suitability
Third-party confirmation, consensus and trust signals. He said most of the computation goes here.

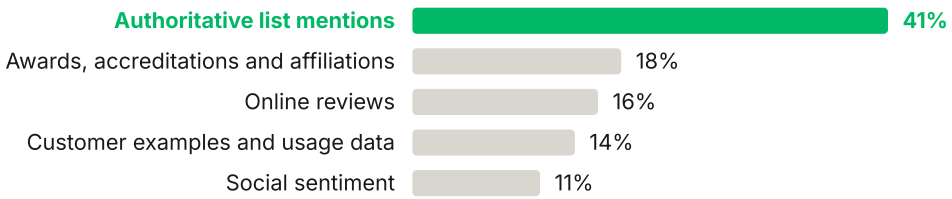
3
Select under uncertainty
It picks once it is confident enough.

Source. Tim Sanders's slide. A large language model (LLM) is the kind of model behind an AI assistant.

He said the evidence has to match the claim. A security company he worked with knew from its sales data that it won 80% of its deals when the prospect believed its agent latency was 25 milliseconds. Its analyst report, reviews and testimonials praised ease of use, quality and service, and none mentioned latency. He added that G2's research found a third of buyers had bought from a vendor they had never heard of before an assistant named it.

One of his last slides is titled "ChatGPT's Recommendation Algorithm" and divides what drives a ChatGPT recommendation among five factors. He called it the best thing he has seen for anticipating the algorithm.

"ChatGPT's Recommendation Algorithm" · each factor's share, as estimated by FirstPageSage, March 2026



Source. Tim Sanders's slide, which credits FirstPageSage, March 2026.

His reading was that a place on a high-authority list matters most for what he called lazy prompts and can break a tie, while awards, reviews and customer examples are gaining influence because they are specific.

Where the budget goes

Sanders said that for every dollar most B2B companies spend on answer engine optimization (AEO) in 2026, 40 go to SEO tools, services, content and search ads. "1 to 40 means 2.5%," he said: AEO spending is one fortieth of SEO spending. He did not argue for an even split, only that 2.5% is wrong. His slide splits an AEO budget five ways.

"How to allocate AEO dollars in 2026" · share of an AEO budget, each given as a range

KIND OF WORK	SHARE OF THE BUDGET
Content restructuring and original authority content	40 to 45%
Third-party proof	20 to 25%
Technical crawlability, schema, feeds and agent readiness	15 to 20%
Measurement tools and dashboards	10 to 15%
Experiments and platform-specific tests	5 to 10%

Source. One dollar to 40, and 2.5%: Tim Sanders, on stage. The five-way split is his slide.

On the panel later that morning he said website traffic is no longer a meaningful measure, and offered two others: have reps ask every inbound lead how they heard of the company, and track year-over-year visits to the contact page.

03

How an assistant picks its sources

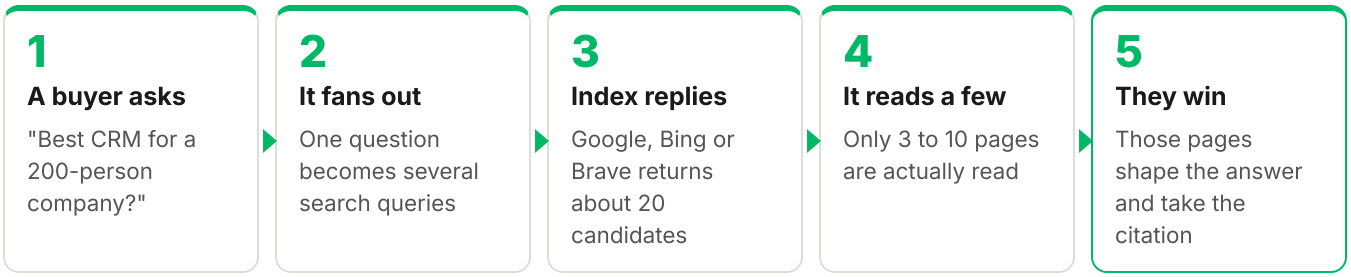
Five sessions explained some part of how an AI assistant decides what to read and cite. Together they follow one sequence, from a buyer's question to the handful of pages an assistant quotes, and they describe a reader that does not scroll, click or run scripts.

> **The sessions.** Kiran Patil and Khushboo Sorthiya, Horizontal Digital, breakout. Devin Stevens, Scrunch, breakout. Nilesh Tamhane and a colleague, XCentium, breakout. Andrew Michael, Scrunch, breakout. Chris Andrew, Scrunch, mainstage. All on Thursday, November 12.

One question becomes several searches

Kiran Patil of Horizontal Digital, a digital agency, showed the sequence as five steps on a slide. A buyer asks a question. The assistant does not search for the question as typed: it breaks it into several search queries. A search index answers with about twenty candidate pages. The assistant reads only three to ten of them, and those are the pages that shape the answer and receive the citation, the link an assistant gives as the source of its answer.

How an AI picks its sources · the five steps on Horizontal Digital's slide, with its example question



Source. Shown by Kiran Patil, Horizontal Digital. The slide credits Sitecore and Scrunch, "The mechanics of AI search", 2026. A CRM is customer relationship management software.

Andrew Michael, VP of product at Scrunch, a company that tracks how brands appear in AI answers, put a number on the second step. Behind one question, he said on stage, the model runs another twelve searches that the person asking never sees.

No index entry, no candidacy

The index in the third step is an ordinary search index. Horizontal Digital's slide names Google, Bing and Brave, and its conclusion is the heading above: a page that is not in the index cannot be a candidate, and a page the crawler cannot read cannot be one of the three to ten. A crawler is a program that fetches pages for a search engine or an AI system. Patil applied this to the three acronyms now in use, which his next slide set in a row: search engine optimization (SEO), answer engine optimization (AEO) and generative engine optimization (GEO).

"Found. Answered. Chosen." · Horizontal Digital's slide on the three acronyms, in the slide's words

FOUND	ANSWERED	CHOSEN
SEO	**AEO**	**GEO**
Can search engines find and rank my page?	When AI gives a direct answer, is it mine?	When an LLM writes its answer, does it choose me?

They stack. They don't compete. Without SEO, nothing else works.

Source. Horizontal Digital's slide, shown by Kiran Patil. An LLM is a large language model, the kind of model behind an AI assistant.

The newest visitor

Horizontal Digital opened its session with one line from a web server's log: a request from GPTBot, OpenAI's crawler, with a caption on what this visitor does not do.

"Say hello to your newest visitor." · Horizontal Digital's opening slide, in the slide's words

FROM AN ACCESS LOG

```
x.x.x.x - - [12/Nov/2026:14:40:02] "GET /services/wealth HTTP/1.1" 200 "Mozilla/5.0 AppleWebKit/537.36 (KHTML, like Gecko; compatible; GPTBot/1.1; +https://openai.com/gptbot)"
```

Doesn't scroll · Doesn't click · Doesn't run your JavaScript

...and it decides whether your brand gets recommended.

Source. The log line and its caption are from Horizontal Digital's opening slide.

Three kinds of bot

Horizontal Digital publishes an open-source app, "Is It Agent-Ready?", and the app's documentation sorts bots into three kinds. It leaves the blocking of a training bot to the site's owner.



KIND	WHAT IT DOES	EXAMPLES
Search	Builds the index an assistant draws its candidates from	OAI-SearchBot, PerplexityBot, Claude-SearchBot
User	Fetches a page at the moment a person's question needs it	ChatGPT-User, Claude-User, Perplexity-User
Training	Collects text to train future models	GPTBot, ClaudeBot, Google-Extended

Source. The documentation of Horizontal Digital's open-source app, "Is It Agent-Ready?".

What the bot receives

A later slide of Horizontal Digital's, headed "What the bot actually sees", puts two panels side by side: a page as its visitor sees it, and what GPTBot receives when that page is built in the browser, which is an empty div and a script tag. The slide states it as a rule: most AI crawlers fetch the HTML and do not run the JavaScript.

"What the bot actually sees" · Horizontal Digital's slide, redrawn. BrandX is its demonstration site

YOUR VISITOR

Empowering your financial future

Our services

GPTBOT

```
<html>
  <head><title>BrandX</title></head>
  <body>
    <div id="root"></div>
    <script src="/static/main.js">
    </script>
  </body>
</html>
```

Most AI crawlers fetch your HTML but don't run your JavaScript.

Source. The two panels and the rule are from Horizontal Digital's slide, which credits the rule to Vercel and MERJ, "The rise of the AI crawler", December 2024.

Patil's test for any page needs no tool. Open it in a browser, choose view source and look for the content. If it is not in the source, a bot will not see it.

Nilesh Tamhane and a colleague from XCentium, a Sitecore partner, gave the everyday cases of content a bot can miss: an FAQ inside an accordion that is closed by default, tabs, carousels, drop-downs and video.

A limited appetite

Two speakers said a crawler spends only a fixed amount on a page. Tamhane put it at about 10,000 tokens a page, after which, he said, a bot moves on. Patil said a crawler processes only a set number of tokens. A token is the unit a language model counts text in: a short word, or a piece of a longer one.

> **Source.** About 10,000 tokens: Nilesh Tamhane, XCentium, answering an audience question. Asked where the figure comes from, he said he thought it is on Scrunch's website.

The same question, different sources

Devin Stevens of Scrunch opened with two screen recordings of one question, put to ChatGPT and to Perplexity minutes apart. His point was that these systems are probabilistic, so one answer is not a measurement.

One question, two answers · as Devin Stevens described his two screen recordings

THE QUESTION Where to find a quick and easy breakfast for a family on a budget. Asked from the same place on the same day, about four minutes apart.

CHATGPT

Returned a visual list of options, with pictures

Left the decision to the user

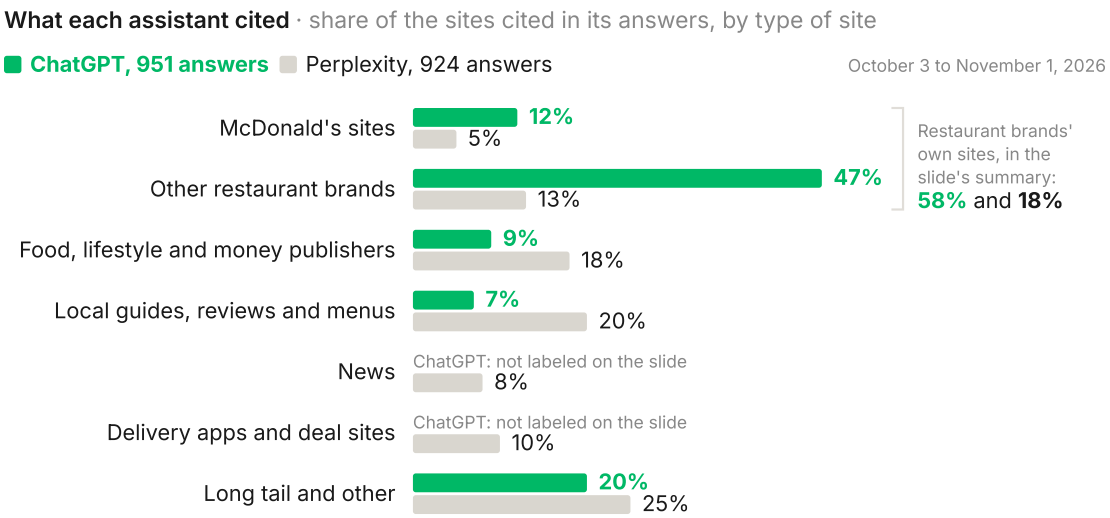
PERPLEXITY

Recommended McDonald's

Spent its answer justifying the choice, with Wendy's, Dunkin' and Jack in the Box lower down

Source. Devin Stevens, Scrunch, on stage. McDonald's is the example brand in Scrunch's own test, not a client.

He then showed a month of such questions, about a thousand answers from each assistant, with the sites those answers cited sorted into types. The questions were non-branded: they named no brand.



Source. Devin Stevens, Scrunch. The slide labels its data "Share of cited domains by type, non-branded prompts", from Scrunch's own test with McDonald's as the example brand, and says Scrunch grouped the sites into types.

The slide's own summary puts McDonald's and the other restaurant brands together: 58% of ChatGPT's citations went to restaurant brands' own sites, against 18% of Perplexity's. Stevens's conclusion was that the split of effort between a brand's own pages and other people's depends on which assistant its customers use.

Most citations are someone else's page

Chris Andrew, Scrunch's co-founder and CEO, said on the mainstage that 70 to 80% of the content an assistant uses to check an answer comes from outside the brand's own site. Andrew Michael said more than 80% of citations are to third-party sources, which are sites the brand does not own.

04

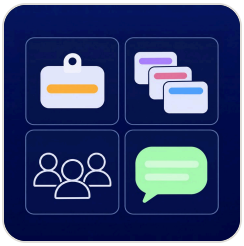
Before any data: four decisions

Devin Stevens of Scrunch, a company that tracks how brands appear in AI answers, named four things a team settles before a tracking tool runs its first prompt: the brand, the competitors, the personas and the prompts. Speakers from G2 and XCentium said where their own prompts come from.

> **The sessions.** Devin Stevens, Scrunch, breakout. Tim Sanders, G2, mainstage keynote. Nilesh Tamhane and a colleague, XCentium, breakout. All on Thursday, November 12. Lisa Saltsman, Icreon, Jaina Baumgartner, Pixleration, and Rodrigo Peplau, MSQ DX, panel, Wednesday, November 11.

What is being set up

Stevens's slide gives a monitoring program five parts. Seven to ten prompts a topic is usual, he said. The slide's other line is that the data that comes out is only as good as the prompts that go in.



How AI search measurement works · Devin Stevens's slide, redrawn

Topics	Prompts	Tags	Metrics	Trends
The business questions you track over time	Sensors that sample each topic	How you slice the data	Presence, citations, position, sentiment	How it all moves over time

RULE OF THUMB If you want to measure it over time, it's a topic. If you want to slice by it, it's a tag.

Source. The five parts, the rule of thumb and the line on prompts are on Devin Stevens's slide. Seven to ten prompts a topic: Devin Stevens, Scrunch, on stage.

His session was planned around five decisions. A fire alarm ended it about 28 minutes into its 45, after the fourth.

The four decisions · as Devin Stevens presented them, with his rule or check for each

DECISION	WHAT IS SETTLED	HIS RULE OR CHECK
1. The brand	Its names, every domain it owns, and its sub-brands	Read five AI answers and see whether AI calls you anything you did not list
2. The competitors	Brands named in answers to non-branded questions	The competitor is anything that solves the same problem
3. The personas	Who is asking, added to each prompt as context	Describe the person and their need, not their opinion of the brand
4. The prompts	The questions tracked, and where they come from	Would a real person ask it? Does it overly favor the brand? Could the brand realistically win it?

Source. The four decisions, the five answers, the persona rule and the three tests are on Devin Stevens's slides. The line on competitors he said on stage.

The brand

The brand came first. Stevens's example throughout was McDonald's, which Scrunch had tracked as a demonstration. Its names include "Mickey D's" and "The Golden Arches". To the corporate, newsroom, careers and regional sites on his slide he added support pages, FAQs, pages behind a sign-in and the company's LinkedIn profile. An assistant asked whether a company is growing, he said, will read its careers page.

Decision 1: define your brand the way AI talks about it · Devin Stevens's slide, redrawn

Alt names & nicknames	All owned domains	Sub-brands & citation segments
What people and AI actually call you	Corporate, newsroom, careers, regional sites	Group the sources you care about

THE EXAMPLE · MCDONALD'S

"Mickey D's" · "The Golden Arches" · "MCD" · corporate., newsroom., franchising., careers. sites

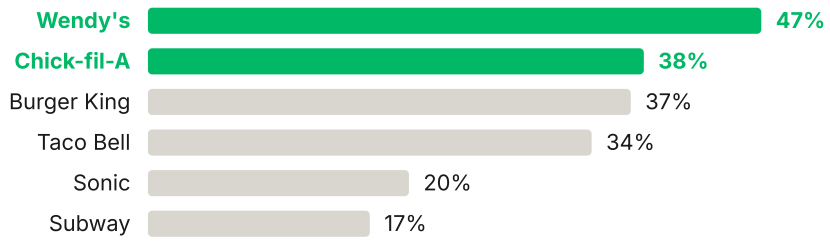
DESK CHECK Read 5 AI answers. Does AI call you anything you didn't list?

Source. Devin Stevens's slide. The example data in his session is Scrunch's own test of McDonald's as a brand.

The competitors

Many brands, Stevens said, arrive with three competitors from a sales comparison deck, usually the ones they want to compete with. The list to use, his slide says, is the brands named in answers to non-branded questions, which are questions that do not name the brand. The slide counts about 8,400 AI responses to such questions in his demonstration, on all the platforms tracked, from September 27 to October 26, 2026.

Who AI names alongside McDonald's · share of about 8,400 AI responses to non-branded prompts in which each brand appears



Source. Devin Stevens's slide, whose footnote reads "Competitor presence on non-branded prompts, all platforms". The data is Scrunch's own test using McDonald's as an example brand. A response can name more than one brand.

For a quick family meal, he added, the competition includes a supermarket's rotisserie chicken.

The personas

An assistant answers using what it knows about the person asking, so the method adds a persona's context to each prompt. Stevens showed three.

Decision 3: personas, who is asking · the three on Devin Stevens's slide, redrawn

Busy Suburban Parent	Young Professional on the Go	Teenager with a Sweet Tooth
Values convenience, kids' meals, drive-thru or curbside pickup on a hectic schedule	Grabs breakfast or coffee on the way to work; orders delivery for a quick lunch	Hangs out with friends; cares about menu variety, burgers and desserts

KEEP PERSONAS NEUTRAL Describe the person and their need, not their opinion of your brand.

Source. The three personas, their descriptions and the rule are on Devin Stevens's slide.

Where good prompts come from

Stevens's slide names four sources of prompts and gives an example of each. The fourth is the brand's own website, which the slide says shows the territory and is not to be mirrored by the prompts.

Decision 4: where good prompts come from · Devin Stevens's slide, redrawn, with its example for each source

SOURCE	THE SLIDE'S STARTING POINT	WHAT IT BECOMES
Real demand Keyword and search data, rewritten as questions people ask AI	"drive thru breakfast"	Best drive-thru breakfast for a family on a budget?
What you need to win Strategy docs and launches, stripped of internal vocabulary	Priority: grow chicken	Which fast food chain has the best chicken sandwich?
How customers talk Support tickets, sales calls, reviews, Reddit	Reddit: "is the app worth it"	Which fast food apps actually save you money?
Your website, at arm's length Shows the territory, but prompts must not mirror your pages	Nutrition page	Topic: healthier options. Not a copy of the page

TEST EVERY PROMPT

1. **Would a real person ask this?**
2. **Does it overly favor your brand?**
3. **Could you realistically win it?**

Source. The four sources, their examples and the three tests: Devin Stevens's slide.

Tim Sanders of G2, a software review site, put real buyer language first, as the section on his keynote reports: reviews, analyst reports, recorded sales calls and requests for proposal. He said chief marketing officers have told him 100 real prompts change everything.

XCentium, a Sitecore partner, builds its prompts around four stages of a purchase. The playbook handed out at Wednesday's partner panel splits prompts by persona and tags them by stage of the sales funnel.

Build your prompts · the four stages on XCentium's slide, with its pattern for each

1	2	3	4
Discover	**Shortlist**	**Compare**	**Decide**
Best ___ for ___?	Who offers ___?	___ vs. ___?	How to choose ___?

> Keywords tell you what people search. **Prompts tell you what they want answered.**

Source. The four stages, their patterns and the line under them: XCentium's slide, presented by Nilesh Tamhane and a colleague. 100 real prompts: Tim Sanders, G2, on stage.

Good prompt, bad prompt

Another of Stevens's slides names three mistakes: a prompt too short to carry intent, a prompt that mirrors the brand's own website, and several rewordings of one question. Of the second he said a set built from questions such as why the brand is the best solution will show the brand every time and tell its team nothing.

Good prompt, bad prompt · the three pairs on Devin Stevens's slide

Too short to carry intent

BAD
"fast food breakfast"

GOOD
"What's the cheapest fast food breakfast for a family of four?"

Mirrors your website

BAD
"What are the benefits of the McCafé rewards program?"

GOOD
"Which coffee chain rewards app gives the most free drinks?"

Five ways to ask one thing

BAD
5 rewordings of "best kids meal"

GOOD
1 prompt per intent. AI fans out its own searches anyway

Source. The three mistakes and the pairs of prompts: Devin Stevens's slide. The remark on the second: Devin Stevens, on stage.

A healthy prompt set

In the healthy set on Stevens's slide, about 80% of the prompts are non-branded and measure competition for demand, and about 20% name the brand and measure reputation and accuracy. The set spans five kinds of question and every stage of the customer's journey.

What a healthy prompt set looks like · Devin Stevens's slide, redrawn. The slide says its example prompts are illustrative

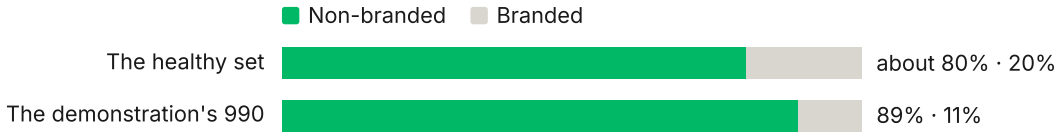
KIND OF QUESTION	NON-BRANDED (~80%) competing for demand	BRANDED (~20%) reputation and accuracy
Best / recommend	Best fast food breakfast for a family on a budget?	Is McDonald's breakfast good for kids?
Compare	Healthiest fast food chicken nuggets?	McDonald's vs Chick-fil-A nuggets
Cost	Cheapest drive-thru meal deals right now?	Is the McDonald's app worth it for deals?
How-to	How do I order ahead for curbside pickup?	How do I use McDonald's mobile order?
Trust	Which fast food chains use real eggs?	Does McDonald's use real beef?

> Cover every journey stage (Awareness, Consideration, Conversion, Loyalty) and tag branded vs non-branded.

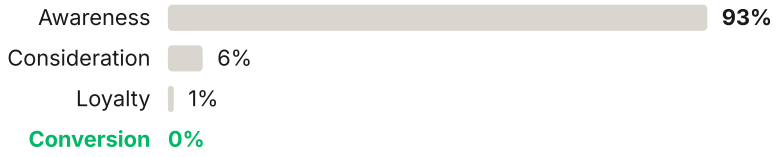
Source. About 80% and 20%, the five kinds and the example prompts: Devin Stevens's slide, which credits Scrunch's prompt strategy templates.

His slides then audit the demonstration's own 990 prompts, under the caption "Heavy on discovery. Nothing at the moment of choice."

The healthy set, and the demonstration's own · share of each set's prompts that are non-branded and branded



THE DEMONSTRATION'S NON-BRANDED PROMPTS · SHARE AT EACH STAGE OF THE JOURNEY



Source. Devin Stevens's slides. The 990 are the active prompts of Scrunch's own McDonald's demonstration, October 3 to November 1, 2026.

05

From finding to fix

Speakers from XCentium, Velir x Brooklyn Data and Pixleration described what is done between reading an AI visibility report and seeing a page cited: how a gap is diagnosed, what is changed, and how the result is checked. XCentium showed a client's numbers before the work and four weeks into it.

> **The sessions.** Nilesh Tamhane and a colleague, XCentium, breakout. Kevin Schofield and David Gelman, Velir x Brooklyn Data, panel. Both on Thursday, November 12. Lisa Saltsman, Icreon, Jaina Baumgartner, Pixleration, and Rodrigo Peplau, MSQ DX, panel, Wednesday, November 11.

A cycle, not a one-time project

Nilesh Tamhane and a colleague from XCentium, a Sitecore partner, presented the way they run AI visibility work for a client as four steps. The heading of their slide is that AI visibility is an optimization cycle and not a one-time project.



The citation cycle · the four steps on XCentium's slide, which draws them as a circle around the words "Get cited"

1
Setup and track
Measure your baseline visibility across prompts

2
Analyze
Find the gaps and see who AI prefers to cite instead

3
Act
Use AI to improve content

4
Validate
Track the lift and prove what works

▾ after the fourth step the cycle returns to the first ▾

Source. The four steps and the heading are on XCentium's slide.

The cycle starts from the gap Tamhane described, reported in the section on Tim Sanders's keynote: of 100 people who ask AI a question the brand wants to answer, 16 hear its name and 6 get a link to its site. His colleague said the aim is not more citations for their own sake, but more chances to influence the buyer's decision.

The baseline: where one client stood

The first step begins with the decisions reported in the section on what is settled before any data, and ends with a baseline. XCentium's came from a real client, shown as a tracking report for the week of August 24 to 30, 2026, before any change was made: 200 AI responses, with five competitors and four assistants tracked.

One client's baseline · what XCentium's slide measured, week of August 24 to 30, 2026, over 200 AI responses

MEASURE	WHAT IT MEASURES	THE CLIENT'S FIGURE
Citation share	How often the client's content is linked as a source: the percentage of AI responses that cite it	21.5%

Share of voice	How often the brand is mentioned, relative to the five competitors tracked	**71.8%**
Mention rate, by assistant	Given for each of the four assistants. The slide does not say what it is a share of	**48.0%** on Google AI Mode, the highest **18.0%** on Perplexity, the lowest

Source. XCentium's slide, labeled real customer data, with the brand names removed. What citation share and share of voice mean: XCentium's presenters, on stage, and a footnote on their later slide of the four weeks that followed.

The presenter read the first two numbers together: the client is mentioned in about 70% of the tracked conversation and cited in about 20%, and the difference is visitors who had no link to follow.

Five causes behind one symptom

XCentium's slide for the analyze step gives one symptom five possible causes.

Find where you're losing, and why · XCentium's slide, redrawn

SAME SYMPTOM · **WE ARE NOT CITED**

Content	Structure	Authority	Understanding	Access
We didn't answer it	AI couldn't extract it	Another source was stronger	Our context was unclear	AI couldn't reach it

The score tells you there's a problem. The prompts tell you where. The pattern tells you why. **Different problems require different fixes.**

Source. The symptom, the five causes and the four lines under them are on XCentium's slide.

A book in a cupboard with no label

Tamhane said the cause of the gap is usually not the content itself. His image was the best book in the library, put in a cupboard with no label. His slide on the gap gives four reasons. Structured data is the machine-readable facts about a page, and `llms.txt` is a file written for agents.

Why it happens · the four reasons on XCentium's slide of the gap

Bloated code	Missing metadata	Hidden content	No agent surface
Heavy markup stalls crawlers	No structured data to cite	Locked in JS, tabs, accordions	No llms.txt or clean feed

Source. Nilesh Tamhane, XCentium. The four reasons are on his slide of the gap.

In the implementations XCentium has done over the last four months, he said, the issue has mostly been accessibility, JavaScript, how content is served over the CDN (the network that delivers a site's files) or how the site's components are set up. Kevin Schofield of Velir x Brooklyn Data said the same of his own work that afternoon: content is usually not the problem and access is.

A fix for each finding

The slide for the act step pairs three findings with a diagnosis and an action.

From finding to fixing · the three rows on XCentium's slide

FINDING	DIAGNOSIS	ACTION
Competitor owns comparison prompts	We don't clearly answer the comparison	**Create answer-first comparison content**
Page exists but isn't cited	Hard for AI to extract	**Improve structure, headings and summaries**
Product isn't understood	Weak entity context	**Add structured data and schema**

Source. XCentium's slide. Its column headings name the tools the agency uses: "Scrunch finding", "Diagnosis" and "Sitecore action". Scrunch is a company that tracks how brands appear in AI answers.

As spoken, the fixes were answer-first content, a description that answers a question and is separate from the standard meta description, clearer headings, more concise summaries, and a structure made for how a bot reads a page. Both presenters said the answer is usually to improve existing content, not to publish a new page each time a prompt goes uncited.

Hidden content, and a second way out

Asked why tabs and accordions are a problem, XCentium's presenters said some bots cannot open content that sits behind JavaScript, such as an FAQ accordion collapsed by default. They recommended putting the same content in structured data (as JSON-LD, a block of data in the page's HTML), which a bot can read without interacting with the page. They said the same of carousels, drop-downs and video, where a transcript belongs on the page and in the JSON.

At the Velir panel a questioner said accordions hide content from agents and help page design. Schofield answered that a website is a restaurant where people are served, and an assistant only wants delivery. The content needs a second way out, he said, such as a Markdown version, and people's experience should not be made worse for bots.

Jaina Baumgartner of Pixleration drew a line on Wednesday's panel. Creating different content for AI engines is cloaking, she said, and it gets a site left out of results. Agent experience pages, her term, are not that: they serve the engine the same content as the person, in a form it need not work to read.

The same work, as a handout

A two-page playbook from Scrunch was handed out at Wednesday's partner panel of Icreon, Pixleration and MSQ DX. It lists similar work: lead with the direct answer, then clear headings, definitions, comparison tables and short summaries. It adds refreshing old pages that are still cited, rewriting pages that agents visit but do not cite, and checking that AI crawlers are not blocked. It calls fixing factual errors about the brand an easy early win, and says that for authority third parties usually carry the most weight.

The partner panel's playbook · the seven steps of its handout, by their headings

STEP	ITS HEADING
1	Build a prompt set that matches how buyers ask
2	Benchmark where you stand

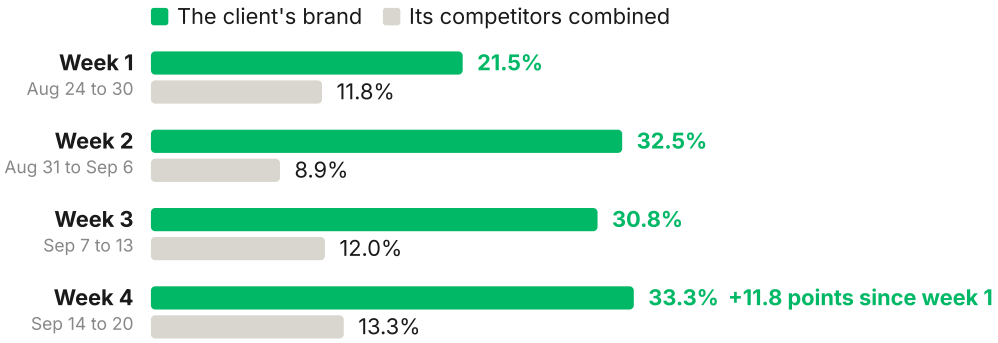
3	Find out which sources the AI trusts
4	Rank your opportunities
5	Get mentioned on third-party pages
6	Make your own pages easy for AI to quote
7	Feed AI agents directly, then measure and repeat

Source. The playbook shared at the partner panel of Lisa Saltsman, Icreon, Jaina Baumgartner, Pixleration, and Rodrigo Peplau, MSQ DX. The paragraph above draws on its steps 2, 3, 5 and 6.

Four weeks on the same prompts

XCentium works on one subject area, product group or business area at a time. Its presenters showed the client of the baseline four weeks on, with the brand names removed. The measure is again citation share: the percentage of AI responses that cite the brand's content, set beside the same percentage for its competitors' content.

One client, four weeks · citation share by week: the percentage of AI responses citing each side's content, August 24 to September 20, 2026



Source. XCentium's slide, labeled as real customer data with the brand names removed. Its footnote reads "% of AI responses citing each source", and says weeks 1 and 4 are partial weeks and that the trend is directional. +11.8 points is the slide's own figure for week 1 to week 4.

The presenters asked themselves how they know the changes moved the number. Tamhane's answer was the method: take 50 or 100 prompts, set a baseline, make changes aimed at those prompts and keep tracking the same ones.

Asked how long it takes, they said about four weeks after the changes shows where things stand, and another four to six weeks may be needed.

06

Choosing the measure first

David Gelman of Velir x Brooklyn Data said a visibility tool gives a view without context, and that the measure has to be chosen before the work starts. This section reports what he and other speakers said a visibility number needs beside it, and who owns the work.

> **The sessions.** Kevin Schofield and David Gelman, Velir x Brooklyn Data, a panel without slides. Kiran Patil, Horizontal Digital, breakout. Andrew Michael, Scrunch, breakout. All three on Thursday, November 12. Lisa Saltsman, Icreon, and Jaina Baumgartner, Pixleration, partner panel, Wednesday, November 11.

Numbers without context

Gelman, who described himself as a data practitioner, said marketing effectiveness has never been judged by one method. He named three that inform each other and called them a tripod. AI visibility, he said, makes it a chair with four legs.

From a tripod to a chair · the four methods as David Gelman named them

LEG 1	LEG 2	LEG 3	THE NEW LEG
Marketing mix model	**Incrementality testing**	**Multi-touch attribution**	**AI visibility**
A statistical estimate of what each channel adds to sales.	Experiments that compare people who saw the marketing with people who did not.	Credit for a sale shared among the steps the buyer took. He spoke of it as no longer accurate.	The leg he said has been added.

Source. David Gelman, Velir x Brooklyn Data, on stage. The panel had no slides: the picture is drawn from what he said. The four names are his. The one-line meanings of the first three are the standard ones and are not his words.

What Scrunch, a company that tracks how brands appear in AI answers, or any tool of its kind gives, he said, is a good view and a limited one, because it operates without context. The context is data the business already holds: margin, the CRM (customer and sales records), the ERP (orders and finance), personas and segmentation. Joined to those, he said, the numbers say what the visibility is worth to the business.

What a referral report counts

Kiran Patil of Horizontal Digital, a digital agency, showed a referral report for his firm's own site, covering mid August to November 12, 2026. It counted 535 sessions that arrived from an AI platform. Of those, it counted 322 as engaged, meaning the visitor stayed and interacted, and 22 as ending in a conversion. A conversion in this report is one of three events: a video event, a form view or a form submission.

Sessions that arrived from an AI platform · Horizontal Digital's own site, mid August to November 12, 2026

STAGE	SESSIONS	SHARE OF THE 535
Arrived from an AI platform	**535**	all of the sessions counted
Engaged	**322**	60.2% of the 535
Ended in a conversion	**22**	4.1% of the 535

Source. Kiran Patil, Horizontal Digital, from the referral report he showed for his firm's own site. The bars are drawn to one scale.

Chosen before the work starts

The panel's moderator asked how to tell what is working when AI does not send clean referral traffic. Gelman said the measure has to be chosen before the work starts: revenue, marketing qualified leads or share of voice. With it go a hypothesis, a time limit and the size of change that would count as an effect. Without other data beside the visibility numbers, he added, a team cannot show even a correlation.

What is set before the work starts · as David Gelman listed it

FIRST	WITH IT	WITH IT	WITH IT
The measure Revenue, marketing qualified leads or share of voice.	**A hypothesis**	**A time limit**	**The size of change that would count as an effect**

Source. David Gelman, Velir x Brooklyn Data, answering the panel's moderator. Drawn from what he said.

How long the joining takes

Whether connecting that data takes years depends on the starting point, Gelman said. With a mature data practice it can be done in weeks, because the joins are on topics, pages and products, and no one has to match people or companies across systems. Even an extract of product margins is often enough to start.

Kevin Schofield, also of Velir x Brooklyn Data, added that SEO takes one to six months before a page ranks, and that citations in an assistant can change in as little as four days.

Keeping the data

Gelman said Scrunch currently keeps 90 days of history, and advised pulling the data out to set a baseline: what is not captured cannot be recovered later.

That afternoon an audience member asked Andrew Michael, Scrunch's VP of product, about a 90-day rolling retention. Michael said the data can be exported and backed up, that he was not sure about 90 days and was pretty sure retention is much longer, and that the team would follow up with exact periods.

> **Source.** 90 days of history: David Gelman, on stage. The answer on export and retention: Andrew Michael, Scrunch, to an audience question in his own session.

Who owns the work

Gelman also said someone among IT, data, marketing and the web team has to drive the process and own the analysis, with the authority to act and not only the responsibility.

On Wednesday's partner panel, Lisa Saltsman of Icreon polled the room on who owns AI visibility, marketing or IT, and most hands went up for a combination. Jaina Baumgartner of Pixleration said the most successful organizations she works with have a digital team between the two.

07

The agent-ready checklist

Horizontal Digital, a digital agency, presented an open-source app that scores a web page on whether an AI agent can reach, read, understand and trust it. This section lists its 20 checks on a page and its 12 checks on a site, each with what passes it and its tier, and says how a page is scored.

> **The session.** Kiran Patil, Group Director, Technology, and Khushboo Sorthiya, Technology Lead, Horizontal Digital, breakout. Thursday, November 12.

The four questions

The app is called "Is It Agent-Ready? for SitecoreAI". It is open source and installs in SitecoreAI, the Sitecore product in which editors build pages, where it scores a draft page before the page is published. The slide that introduced it puts four questions at its center, and every page check sits under one of them.

The four questions the app asks · from Horizontal Digital's slide

Reach	Read	Understand	Trust
Can AI agents get to it?	Can they read it?	Do they understand it?	Will they trust and cite it?

Source. Horizontal Digital's slide, shown in the session by Kiran Patil and Khushboo Sorthiya.

The same slide names what the app draws on: the page's fields and layout, the page's HTML fetched the way each AI bot fetches it, without JavaScript, and an LLM (a large language model, the kind of model behind an AI assistant) that drafts fixes such as a title and description, FAQ markup, an image's written description and an author.

The app's public documentation lists every check with its pass mark. Every check, tier, pass mark and scoring rule in this section is from that documentation as it stood on November 12, 2026: the README and the technical list of page checks in the app's repository. The pass marks are restated here, and the thresholds are Horizontal Digital's choices.

How a page is scored

A page is scored out of 100 points. Each check has a tier: essential, recommended or bonus. Essential checks share 80 of the points, recommended checks share 20, and bonus checks can add up to 5, with the total capped at 100. A check that does not apply to a page is skipped, and its points are shared among the rest.

A page's 100 points · how they are shared among the 20 page checks

TIER	POINTS ITS CHECKS SHARE	PAGE CHECKS IN THE TIER
Essential	80 of a page's 100 points	6 of the 20
Recommended	20 of a page's 100 points	13 of the 20

Bonus	**Up to 5**	1 of the 20
	added on top, with the total capped at 100	

Source. The scoring rules are from the app's documentation. 6, 13 and 1 are a count of the checks in each tier in its list of 20 page checks. The bars are drawn to one scale.

One check is a gate and carries no points: a page marked `noindex`, `nosnippet` or `max-snippet:0` is treated as blocked. Those three values tell search engines not to list a page or not to quote from it. The score then gets one of three labels.

What the app calls the result · a page's score out of 100

BELOW 60	60 TO 79	80 OR MORE
Not ready	**Almost there**	**Agent-ready**
	The article in the session's demonstration scored 69.	

Source. The three labels and their thresholds: the app's documentation. 69: Khushboo Sorthiya, Horizontal Digital, demonstrating the app.

In the demonstration Khushboo Sorthiya opened an article on a demo site called BrandX. The app gave the page 69, labeled almost there, with its checks listed under the four questions. After a meta description (the page's short summary for search engines) that was too short was rewritten, the score went up. A site has a score of its own from the 12 site checks: BrandX's was 85, labeled agent-ready. She did not say what the site's score of 85 is out of.

One check, as the session pictured it

"Content is there without JavaScript" is the first check under Read and an essential one. A slide in the session, headed "What the bot actually sees", shows the case it looks for: a page that a script assembles in the visitor's browser. The visitor sees the finished page. GPTBot, OpenAI's crawler, receives a body that holds an empty `div` and a script tag.

"What the bot actually sees" · Horizontal Digital's slide, redrawn

YOUR VISITOR

Empowering your financial future

Our services

GPTBOT

```
<html>
  <head><title>BrandX</title></head>
  <body>
    <div id="root"></div>
    <script src="/static/main.js"></script>
  </body>
</html>
```

Most AI crawlers fetch your HTML but don't run your JavaScript.

Source. Horizontal Digital's slide, shown in the session. The line under the picture is the slide's own, and the slide credits it to Vercel and MERJ, "The rise of the AI crawler", December 2024. BrandX is the demo site's name.

The app compares the words in a page's draft with the words in the HTML the server returns. If at least 95% of the draft's words are there the check passes, and 80% earns half the check's points.

The 20 page checks

Each check sits under one of the four questions. The middle column restates the pass mark in the app's documentation. Where a check can be partly passed, the share it earns is a share of that check's own points.

The table uses several terms. A canonical link is the link in which a page names its own address. Markdown is a plain-text format, requested with `Accept: text/markdown`, and page chrome is the menus, headers and footers around a page's content. PageSpeed is Google's test of how fast a page loads, with a score out of 100. JSON-LD is a block of data in a page's HTML that states what the page describes, and a breadcrumb list is the trail of links that shows where a page sits in its site. Open Graph tags are a page's sharing tags. A reading grade is the U.S. school grade a reader needs to have reached to follow a text: 13 is the first year of college. Alt text is an image's written description.

The page checks · each with what passes it and its tier

CHECK	WHAT PASSES	TIER
REACH · CAN AI AGENTS GET TO IT?		
Canonical points to this page	A canonical link that refers to the page itself	**Essential**
READ · CAN THEY READ IT?		
Content is there without JavaScript	At least 95% of the words in the page's draft are in the HTML the server returns. 80% earns half the check's points	**Essential**
Agents can get this page as Markdown	A request with `Accept: text/markdown` gets Markdown back	**Essential**
Page structure is machine-readable	`main`, `article`, `nav` and `header` elements: a quarter of the check's points for each one present	Recommended
The Markdown version is complete	The Markdown covers 90% of the page's content, with little page chrome	Recommended
Page loads fast	A PageSpeed score of 70 or more	Recommended
UNDERSTAND · DO THEY UNDERSTAND IT?		
One clear H1 and logical headings	Exactly one `h1`, and no skipped heading levels	**Essential**
Title and description are set	A title of 50 to 60 characters and a description of 120 to 160 characters	**Essential**
Structured data is present and valid	JSON-LD with a breadcrumb list on pages other than the home page, FAQ markup where there are two or more questions, and article markup on articles	**Essential**
Page identity metadata is complete	`lang` on the `html` element and the four main Open Graph tags	Recommended
FAQ answers are quotable	Each answer is between 40 and 300 characters long	Recommended
Sections answer real questions	Three or more of the page's `h2` headings, or half of them, are phrased as questions	Recommended

Easy to read and quote	Paragraphs averaging 40 to 120 words, at a reading grade of 13 or below	Recommended
Images are described	Images carry alt text	Recommended

TRUST · WILL THEY TRUST AND CITE IT?

Content is original, not generic	Rated unique by the app's LLM, which reads the first 3,000 characters	Recommended
Author is named and linked	Articles only: a named author with a link	Recommended
Content is recently updated	A modified date within 180 days earns the check's full points, falling away to two years	Recommended
Enough depth to be cited	800 or more words in body paragraphs earns the check's full points, and 400 earns two thirds of them	Recommended
Links say where they go	No more than 20% of the page's links have generic text such as "click here"	Recommended
Answers the next question too	Covers the follow-up questions the LLM expects	Bonus

Source. Every check, tier and pass mark is from the public documentation of Horizontal Digital's app, as it stood on November 12, 2026. The documentation as recorded for this report does not say how the 90% in "The Markdown version is complete" is counted.

The 12 site checks

These run once for the whole site and have their own score. Each has a tier, marked the same way. Two of them refer to seven search and user agents, which the second table names.

The terms here: `robots.txt` is the file that tells crawlers what they may fetch, and a sitemap lists a site's pages for them. `Vary: Accept` is a response header saying the answer depends on the format a request asks for. 404 and 410 are the status codes for a page that is not there or has been removed, and 200 is the code for a page that was found. MCP is the protocol an agent uses to call outside tools.

The site checks · each with what passes it and its tier

CHECK	WHAT PASSES	TIER
AI agents can open your site	Requests made as each search and user agent get what a browser gets	**Essential**
`robots.txt` **lets AI search in**	The seven search and user agents are allowed on `/`	**Essential**
The site can serve Markdown to agents	The home page returns Markdown for `Accept: text/markdown`, with `Vary: Accept`	**Essential**
Sitemap is published and linked	A valid sitemap named in `robots.txt`, with a last-modified date on 80% of the URLs it lists	Recommended
`llms.txt` **is published**	A file at `/llms.txt` with a title, a summary and a list of links	Recommended
About, contact and privacy pages exist	Each is linked from the home page and loads	Recommended
The organization is described for AI	Organization JSON-LD on the home page with name, URL and logo	Recommended

Missing pages return a real 404	A made-up URL returns 404 or 410, not 200	Recommended
Content-use signals are declared	A `Content-Signal` line in `robots.txt`	Bonus
Discovery links in HTTP headers	A `Link` header on the home page pointing to machine-readable versions	Bonus
Agent protocols published	MCP, agent-to-agent and API discovery files, for sites that are applications	Bonus
Agentic checkout protocols	Checkout protocol endpoints, for commerce sites	Bonus

Source. Every check, tier and pass mark is from the public documentation of Horizontal Digital's app, as it stood on November 12, 2026.

The bots the site checks name · the seven search and user agents are scored, and training bots are not

KIND OF BOT	NAMED IN THE DOCUMENTATION	SCORED
Search Builds the index an assistant draws on	OAI-SearchBot, Claude-SearchBot, PerplexityBot, Amazonbot	Yes: four of the seven
User Fetches a page when a person's question needs it	ChatGPT-User, Claude-User, Perplexity-User	Yes: three of the seven
Training Collects text for model training	GPTBot, ClaudeBot, Google-Extended, Applebot-Extended and Meta's agent	No: blocking one is treated as the site owner's choice and costs nothing

Source. The three kinds of bot, the names and what is scored are from the app's public documentation.

Open source

The session's last slide on the app gives the address of its repository and three steps: fork the app and install it as a custom app in SitecoreAI, add rules of your own, and send a pull request, or PR, for a check others would need.

"It's open source. Make it yours." · Horizontal Digital's slide, redrawn. The slide also carried a QR code for the same address

```
https://github.com/horizontalintegration/is-it-agent-ready-for-sitecoreai
```

1
Fork it
Install it as a custom Marketplace app in your SitecoreAI organization

2
Add your rules
Brand voice, regulated disclosures, your schema standards

3
Send a PR
Wrote a check we all need? Share it back.

Source. Horizontal Digital's slide, shown in the session, with its wording. A Marketplace app is one that installs into SitecoreAI.

08

Every number, and who presented it

This section puts every figure in the report in one place, with what it measures and who presented it, so that a reader can check where a number came from. A last table lists the speakers and the sessions in which they spoke.

> **How to read it.** The figures are grouped under the section that reports them, in the order they appear there. Where a section says the speaker or the slide did not give the base of a figure, the row says "base not stated". "On stage" means the figure was spoken and was not on a slide.

The new reader

The figures in "The new reader" · what each measures and who presented it

THE FIGURE	WHAT IT MEASURES	WHO PRESENTED IT
More than $1,500	What he said he spent on playbooks, wristbands, cones and flags after asking ChatGPT how to run a flag football practice, without setting out to buy anything	Eric Stine, Sitecore, of his own purchases
Within a minute	The time in which, in his claim, an assistant now does what took days, weeks or months across the stages of marketing	Andrew Michael, Scrunch, on stage
900 U.S. adults **68,879** of their searches	The people and the searches on which the count of clicks on Google rests	The slide of Kiran Patil and Khushboo Sorthiya, Horizontal Digital, which credits Pew Research Center, July 2025
15% results without an AI summary **8%** results with an AI summary **1%** clicked a link cited inside the AI summary	The share of Google visits where the user clicked a link, which is the slide's own label for all three	The slide of Kiran Patil and Khushboo Sorthiya, Horizontal Digital, which credits Pew Research Center, July 2025
60%	The share of searches that end without a click, in the words of the card	The slide of Nilesh Tamhane and a colleague, XCentium. The card credits Forbes
58%	The card is headed "58% drop in organic traffic", meaning visits that arrive from unpaid search results. Its small print gives the 58% as a drop in click-through rate once an AI Overview appears	The slide of Nilesh Tamhane and a colleague, XCentium. The card credits Ahrefs
25% by 2026 **50%+** by 2028	A forecast drop, printed under the same card's heading of organic traffic	The slide of Nilesh Tamhane and a colleague, XCentium, which credits the forecast to Gartner
More than 41%	The share of AI sessions that happen entirely inside the AI platform, with no open-web search before and no click to the open web after	Michelle Boockoff-Bajdek, Sitecore, citing Scrunch

41% stayed inside the AI conversation **32%** web, then AI, then back to the web **16%** started on the web, finished in AI **11%** started in AI, finished on the web	The share of all AI sessions in the study that took each path. He said the study is based on millions of consumer prompts from real interactions. How the study defines a session: not stated	The slide of Chris Andrew, Scrunch, from a new Scrunch data study
51%	The share of all page views and hits in June that came from bots and not from people	Kevin Schofield, Velir x Brooklyn Data, citing Cloudflare
$12 a book **$500** a deposit	The price of the out-of-print book that Meta's Muse agent found and paid for, and the deposit on a car the agent had recommended	Chris Andrew, Scrunch, on stage, of his own purchases

Source. The section "The new reader". Each credit is as that section gives it.

Found is not chosen

The figures in "Found is not chosen" · what each measures and who presented it

THE FIGURE	WHAT IT MEASURES	WHO PRESENTED IT
29% April 2025 **51%** March 2026	The share of B2B software buyers who said they start researching a purchase with AI search more often than with Google. He said the question set a Google property against ChatGPT, Gemini, Perplexity, Claude or DeepSeek, so Google's AI Overviews do not count toward the 51%	The slide of Tim Sanders, G2, from G2's surveys of B2B software buyers
76%	The jump in that share in a year, as he gave it: 51 is about three quarters more than 29	Tim Sanders, G2, on stage
58% large enterprise buyers **65%** IT professionals	The share of each group who start in chat	Tim Sanders, G2, on stage
6 to 15 working hours	The work by which a buying committee reaches its shortlist of three, which he said buyers now get to in one prompt	Tim Sanders, G2, on stage. He said G2 has studied the shortlist for about a decade

54% AI chat **43%** software review sites **36%** vendor website **36%** research analyst firms **33%** peers and colleagues **30%** community forums **28%** thought leader content **20%** internal supplier portal **19%** vendor salesperson	The share of buyers surveyed who named each source as what puts a vendor on the shortlist. A buyer could name more than one, so the shares add up to more than 100%	The slide of Tim Sanders, G2, from G2's survey of March 2026
93%	The share of buyers who decide from the shortlist they had on day one	Tim Sanders, G2, who attributed it to 6sense
132 times	In his example, how often a vendor's site says "hospitals" while documenting neither reimbursement speed nor mobile use	Tim Sanders, G2, on stage
14 vendors	The long list for the captured hospital prompt as he pictured it: the vendors that had checked every box	Tim Sanders, G2, on stage
Almost 60%	The share of the words in a prompt taken up by the job and its constraints, in thousands of prompts captured from his customers' buyers	Tim Sanders, G2, on stage
Two of every three	The prompts, among the same captured prompts, in which the category name appeared	Tim Sanders, G2, on stage
About 0.1% a citation **0.33 to 0.5%** a branded citation **Closer to 7%** a recommendation, "70X" the rate of a citation	Click-through rate: of the people who are shown the answer, the share who click through to the brand's site. He gave no rate for a mention	Tim Sanders, G2, on stage. He credited the citation rate to one researcher's click-through research and called it an average
60% of the time	How often a brand that was cited was still left off the list of recommendations, over thousands of runs of commercial-intent prompts	Tim Sanders, G2, citing a September study by Jeff Oxford
16 hear the brand's name **6** get a link to its site **94** get no path to the site	Of 100 people asking AI a question the brand wants to answer, as he said it. The slide reads "16% mention your brand" and "6% of those link to you"	Nilesh Tamhane, XCentium, on stage. XCentium's slide credits a Webflow analysis of 2,000 U.S. company sites

61% mentioned **22%** cited **12%** recommendation rate, "about 2× the work of a mention"	Example figures on a slide that shows one question answered two ways. Base not stated: the slide does not say what they are shares of	Scrunch's slide "Recommended is what matters", shown by Andrew Michael, Scrunch
80% of deals won **25 milliseconds** agent latency	The share of its deals a security company won when the prospect believed its agent latency was 25 milliseconds, from the company's sales data	Tim Sanders, G2, on stage, of a company he worked with
A third	The share of buyers who had bought from a vendor they had never heard of before an assistant named it	Tim Sanders, G2, from G2's research
41% authoritative list mentions **18%** awards, accreditations and affiliations **16%** online reviews **14%** customer examples and usage data **11%** social sentiment	Each factor's share of what drives a ChatGPT recommendation, as an estimate	The slide of Tim Sanders, G2, "ChatGPT's Recommendation Algorithm", which credits FirstPageSage, March 2026
1 to 40 dollars **2.5%**	For every dollar most B2B companies spend on answer engine optimization in 2026, 40 go to SEO tools, services, content and search ads. The 2.5% is one fortieth: AEO spending as a share of SEO spending	Tim Sanders, G2, on stage
40 to 45% content restructuring and original authority content **20 to 25%** third-party proof **15 to 20%** technical crawlability, schema, feeds and agent readiness **10 to 15%** measurement tools and dashboards **5 to 10%** experiments and platform-specific tests	The share of an AEO budget for each kind of work, each given as a range	The slide of Tim Sanders, G2, "How to allocate AEO dollars in 2026"

Source. The section "Found is not chosen". Each credit is as that section gives it. AEO is answer engine optimization and SEO is search engine optimization.

How an assistant picks its sources

The figures in "How an assistant picks its sources" · what each measures and who presented it

THE FIGURE	WHAT IT MEASURES	WHO PRESENTED IT

About 20 candidates **3 to 10** pages read	The candidate pages a search index returns for an assistant's queries, and the pages among them that the assistant actually reads	The slide shown by Kiran Patil, Horizontal Digital, which credits Sitecore and Scrunch, "The mechanics of AI search", 2026
Twelve	The further searches a model runs behind one question, which the person asking never sees	Andrew Michael, Scrunch, on stage
About 10,000 tokens	What a bot spends on one page before it moves on. A token is the unit a language model counts text in	Nilesh Tamhane, XCentium, answering an audience question. He said he thought the figure is on Scrunch's website
About four minutes	The time between his two screen recordings of one question, put to ChatGPT and to Perplexity from the same place on the same day	Devin Stevens, Scrunch, on stage
951 ChatGPT **924** Perplexity	The answers from each assistant to a month of non-branded questions, October 3 to November 1, 2026	The slide of Devin Stevens, Scrunch, from Scrunch's own test with McDonald's as the example brand
12% and 5% McDonald's sites **47% and 13%** other restaurant brands **9% and 18%** food, lifestyle and money publishers **7% and 20%** local guides, reviews and menus **Not labeled and 8%** news **Not labeled and 10%** delivery apps and deal sites **20% and 25%** long tail and other	The share of the sites cited in each assistant's answers, by type of site, with ChatGPT first and Perplexity second. The slide's label is "Share of cited domains by type, non-branded prompts", and it says Scrunch grouped the sites into types	The slide of Devin Stevens, Scrunch, from Scrunch's own test with McDonald's as the example brand
58% ChatGPT **18%** Perplexity	The slide's own summary: the share of each assistant's citations that went to restaurant brands' own sites, McDonald's and the others together	The slide of Devin Stevens, Scrunch
70 to 80%	The share of the content an assistant uses to check an answer that comes from outside the brand's own site	Chris Andrew, Scrunch, on the mainstage
More than 80%	The share of citations that are to third-party sources, which are sites the brand does not own	Andrew Michael, Scrunch

Source. The section "How an assistant picks its sources". Each credit is as that section gives it.

Before any data: four decisions

The figures in "Before any data: four decisions" · what each measures and who presented it

THE FIGURE	WHAT IT MEASURES	WHO PRESENTED IT

Seven to ten	The prompts that are usual for one topic in a monitoring program	Devin Stevens, Scrunch, on stage
28 of 45 minutes	How far into its planned length the session was when a fire alarm ended it, after the fourth of five planned decisions	Not a speaker's figure: it describes the session of Devin Stevens, Scrunch
Five	The AI answers to read in his check of whether AI calls the brand anything that was not listed	The slide of Devin Stevens, Scrunch
About 8,400	The AI responses to non-branded questions counted in his demonstration, on all the platforms tracked, from September 27 to October 26, 2026	The slide of Devin Stevens, Scrunch
47% Wendy's **38%** Chick-fil-A **37%** Burger King **34%** Taco Bell **20%** Sonic **17%** Subway	The share of those responses, about 8,400, in which each brand appears. A response can name more than one brand	The slide of Devin Stevens, Scrunch, whose footnote reads "Competitor presence on non-branded prompts, all platforms". The data is Scrunch's own test using McDonald's as an example brand
100	The real prompts that, he said, chief marketing officers have told him change everything	Tim Sanders, G2, on stage
About 80% non-branded **About 20%** branded	The share of the prompts in a healthy prompt set that do not name the brand, and the share that do	The slide of Devin Stevens, Scrunch, which credits Scrunch's prompt strategy templates
990	The active prompts of Scrunch's own McDonald's demonstration, October 3 to November 1, 2026	The slides of Devin Stevens, Scrunch
89% non-branded **11%** branded	The share of the demonstration's 990 prompts of each kind	The slides of Devin Stevens, Scrunch
93% awareness **6%** consideration **1%** loyalty **0%** conversion	The share of the demonstration's non-branded prompts at each stage of the customer's journey	The slides of Devin Stevens, Scrunch

Source. The section "Before any data: four decisions". Each credit is as that section gives it.

From finding to fix

The figures in "From finding to fix" · what each measures and who presented it

THE FIGURE	WHAT IT MEASURES	WHO PRESENTED IT

200 AI responses **Five** competitors **Four** assistants	What one client's baseline report covered in the week of August 24 to 30, 2026, before any change was made	The slide of Nilesh Tamhane and a colleague, XCentium, labeled real customer data, with the brand names removed
21.5%	The client's citation share in that week: the percentage of the AI responses that cite its content	XCentium's slide. The meaning of citation share: XCentium's presenters, on stage, and a footnote on their slide of the four weeks
71.8%	The client's share of voice in that week: how often the brand is mentioned, relative to the five competitors tracked	XCentium's slide. The meaning of share of voice: XCentium's presenters, on stage
48.0% Google AI Mode, the highest **18.0%** Perplexity, the lowest	The client's mention rate on two of the four assistants. Base not stated: the slide does not say what the rate is a share of	XCentium's slide
About 70% mentioned **About 20%** cited	The presenter's reading of the first two numbers together: the client is mentioned in about 70% of the tracked conversation and cited in about 20%	XCentium's presenter, on stage
21.5% week 1, August 24 to 30 **32.5%** week 2, August 31 to September 6 **30.8%** week 3, September 7 to 13 **33.3%** week 4, September 14 to 20	The client's citation share by week in 2026: the percentage of AI responses citing its content. The slide says weeks 1 and 4 are partial weeks and that the trend is directional	XCentium's slide, labeled real customer data, whose footnote reads "% of AI responses citing each source"
11.8% week 1 **8.9%** week 2 **12.0%** week 3 **13.3%** week 4	The same percentage for the content of the client's competitors combined, over the same four weeks	XCentium's slide
+11.8 points	The rise in the client's citation share from week 1 to week 4	XCentium's slide: the figure is the slide's own
50 or 100	The prompts to take in his method: set a baseline, make changes aimed at those prompts and keep tracking the same ones	Nilesh Tamhane, XCentium, on stage
About four weeks after the changes **Four to six weeks** more, if needed	How long it takes to see where things stand, and the further time that may be needed	XCentium's presenters, answering a question

Source. The section "From finding to fix". Each credit is as that section gives it. That section also repeats Nilesh Tamhane's 16 and 6 of 100 people, which are listed above under "Found is not chosen".

Choosing the measure first

The figures in "Choosing the measure first" · what each measures and who presented it

THE FIGURE	WHAT IT MEASURES	WHO PRESENTED IT
535	The sessions that arrived at his firm's own site from an AI platform, from mid August to November 12, 2026	Kiran Patil, Horizontal Digital, from the referral report he showed
322 60.2% of the 535	The sessions among them counted as engaged, meaning the visitor stayed and interacted	Kiran Patil, Horizontal Digital, from the same report
22 4.1% of the 535	The sessions among them counted as ending in a conversion, which in this report is a video event, a form view or a form submission	Kiran Patil, Horizontal Digital, from the same report
One to six months SEO **Four days** citations	The time SEO takes before a page ranks, and the least time in which citations in an assistant can change	Kevin Schofield, Velir x Brooklyn Data, on stage
90 days	The history he said Scrunch currently keeps. Asked that afternoon about a 90-day rolling retention, Andrew Michael said he was not sure about 90 days, was pretty sure retention is much longer, and that the team would follow up with exact periods	David Gelman, Velir x Brooklyn Data, on stage. The answer: Andrew Michael, Scrunch, to an audience question

Source. The section "Choosing the measure first". Each credit is as that section gives it.

The agent-ready checklist

The pass marks of the checks, such as a title's length or a page's word count, are tabulated in the section on the checklist and are not repeated. The rows here are the figures of its scoring.

The scoring figures in "The agent-ready checklist" · what each measures and who presented it

THE FIGURE	WHAT IT MEASURES	WHO PRESENTED IT
20 checks on a page **12** checks on a site	The checks the app "Is It Agent-Ready?" runs. The site checks run once for the whole site and have their own score	The public documentation of Horizontal Digital's app, as it stood on November 12, 2026
80 essential checks **20** recommended checks **Up to 5** bonus checks	The points the checks of each tier share, out of a page's 100. Bonus points are added on top, with the total capped at 100	The public documentation of Horizontal Digital's app
Below 60 not ready **60 to 79** almost there **80 or more** agent-ready	The label the app gives a page's score out of 100	The public documentation of Horizontal Digital's app
69	The score, out of 100, of the article opened in the session's demonstration on the demo site BrandX, labeled almost there	Khushboo Sorthiya, Horizontal Digital, demonstrating the app

85	The score of the demo site BrandX from the 12 site checks, labeled agent-ready. Base not stated: she did not say what the site's score is out of	Khushboo Sorthiya, Horizontal Digital, demonstrating the app

Source. The section "The agent-ready checklist". Each credit is as that section gives it.

The speakers

These are the people the report credits, by company, with the session or sessions from which it reports them. All of the sessions were at Sitecore Symposium 2026 on Wednesday, November 11 or Thursday, November 12, 2026.

Who spoke, and where · one row for each person, sorted by company

PERSON	COMPANY	SESSION
Tim Sanders Chief Innovation Officer	G2 a software review site	Mainstage keynote, Thursday, 8:30 AM. Marketing leaders panel on the mainstage, Thursday, about 9:52 AM
Kiran Patil Group Director, Technology	Horizontal Digital a digital agency	Breakout, Thursday, 2:40 PM
Khushboo Sorthiya Technology Lead	Horizontal Digital	Breakout, Thursday, 2:40 PM, with Kiran Patil
Lisa Saltsman	Icreon	Partner panel, Wednesday, 3:25 PM
Rodrigo Peplau	MSQ DX	Partner panel, Wednesday, 3:25 PM
Jaina Baumgartner	Pixleration	Partner panel, Wednesday, 3:25 PM
Chris Andrew Co-founder and CEO	Scrunch a company that tracks how brands appear in AI answers	Mainstage, Thursday morning
Andrew Michael VP of product	Scrunch	Breakout, Thursday, 4:30 PM
Devin Stevens	Scrunch	Breakout, Thursday, 11:45 AM
Eric Stine CEO	Sitecore the company whose conference this was	Opening mainstage keynote, Wednesday, 8:15 AM
Michelle Boockoff-Bajdek Chief Marketing Officer	Sitecore	Mainstage keynote, Wednesday, 9:40 AM
David Gelman described himself as a data practitioner	Velir x Brooklyn Data	Panel without slides, Thursday, 3:50 PM
Kevin Schofield	Velir x Brooklyn Data	Panel without slides, Thursday, 3:50 PM
Nilesh Tamhane	XCentium a Sitecore partner	Breakout, Thursday, 2:00 PM, with a colleague the report does not name

Source. The people, companies, roles and kinds of session are as the report's sections give them, and a role is shown only where a section gives one. Jeff Oxford, whose September study Tim Sanders cited, is named in the report and did not speak.