



WHITEPAPER

From Individual Research to Organizational Impact

What agentic AI means for the next chapter of competitive and market intelligence



From information scarcity to analysis abundance

Generating competitive analysis has never been easier. A product manager, a sales director, a strategy intern. Anyone can now generate a credible-looking competitor profile, a SWOT analysis, or a market overview in minutes. It used to be that research was time-consuming, sources were fragmented, and the teams with the best access to data had a structural advantage.


The volume of analysis being produced inside organizations is increasing rapidly. Frontier model capability has grown exponentially. What AI can do today is fundamentally different from 18 months ago, and adoption across organizations has followed.

Strategy work is highly exposed to AI disruption.

”More than 80% of tasks face high or medium exposure to AI automation.”

BCG, 2026

The challenges competitive and market intelligence (CMI) functions face have changed too. This paper explores what that means in practice, and what it takes to build intelligence that organizations can trust and act on in the age of agentic AI.




What the shift means for CMI

The launch of deep research tools in early 2025 changed how competitive research gets done. In the last six months alone, significant improvements in reasoning quality have raised the bar further. Today's leading models are meaningfully better at interpreting what signals mean. Most organizations are still catching up.

AI-generated analysis can look authoritative. The quality and reliability of what it is based on varies widely. And the CMI function faces a question it has not had to answer in quite this way before: if research is no longer scarce, what is its distinctive contribution?

We believe that AI shifts the role of CMI from producing individual pieces of research toward building and operationalizing an organization's intelligence agenda. The rest of this paper explores why.

BCG's research on corporate strategy functions found that the area where AI has delivered the most consistent, positive impact is market intelligence and research. CMI's core territory. Across other strategy activities, including M&A, portfolio management, and major strategic decisions, results have been limited.



How AI is being used for CMI today and where it falls short

Enterprise Copilot mandates, individual tool subscriptions, and team-level experiments have put AI into competitive research workflows across organizations, often with little involvement from the CMI function.

Across the CMI teams and senior leaders Valona works with, we see consistent patterns in how that plays out.

REACTIVE RESEARCH

AI gets adopted for one-off reactive tasks like pulling background on a company or answering a stakeholder question, and performs well in that context. But the work tends to stay reactive, and output quality varies in ways that matter when intelligence is informing decisions.

SILOS

Intelligence stays with the individual who ran the research. Someone in the organization gets an answer, reads it, and moves on. What was learned stays siloed.

GOVERNANCE FRICTION

Pilots built outside IT and enterprise AI governance can run into problems later, when data residency requirements, tool standardization, or compliance reviews catch up with what was built in isolation.

COMPLIANCE EXPOSURE

Licensed content fed into general purpose AI tools may violate content agreements without the organization knowing, creating compliance exposure that only surfaces later.

COST AT SCALE

AI tools feel affordable today, often bundled into enterprise licenses or absorbed centrally. As agentic workflows scale, token usage can grow rapidly, making cost efficiency an architectural decision rather than an operational one.

What connects these patterns is a more fundamental question about how AI works. AI models generate responses based on patterns, not verified facts. Before intelligence reaches a decision maker it needs to be validated on source quality, accuracy and relevance to the specific decision at hand.

In most organizations, that validation step is either missing or inconsistent. CMI teams that build validation into their process do more than reduce risk. They become the team decision makers turn to because someone is accountable for the intelligence they act on.

Three barriers to scaling AI for CMI

The frequency of strategic decisions is increasing, which means leaders need a shared, reliable picture of the competitive landscape more often and more quickly. AI has made individual research faster. What it has not done is ensure that intelligence reaches the right people, at the right time, in a form they can act on. Three problems get in the way.

1 Visibility

Analysis gets created every day, in personal chats, individual research sessions, prompted analyses that answer one question and are never seen again. The work gets done. It just never becomes organizational knowledge.

When intelligence stays with individuals, the organization never builds a shared understanding of its competitive landscape. Every team fills that gap with its own research, drawing its own conclusions. By the time CMI's intelligence arrives, it is often competing with decisions already being made. The result is more analysis and less alignment.

2 Trust

Trust in intelligence is not given. It is rebuilt with every deliverable that reaches a decision maker. When those outputs cannot be traced back to a source, decision makers either slow down to verify them or discount them entirely. Either way, the intelligence doesn't do its job.



WHY AI GETS THINGS WRONG

AI models are nondeterministic, meaning the same question asked twice can produce meaningfully different answers. They are also not retrieving verified facts but predicting what a plausible answer looks like based on patterns in their training data. This means they can state things confidently that are incomplete, outdated, or simply wrong. This is not a flaw that better models will eventually eliminate. It is a structural characteristic of how large language models work, and it has to be managed through the systems built around them.



3 Continuity

Markets move, competitors act, regulations change. The strategic context your organization operates in looks different every quarter. A report assembled under time pressure captures a moment. What organizations need is a continuously maintained view of the market. In practice it looks like automated alerts when a competitor makes a significant move or a regulatory development requires a response.

Building and sustaining that kind of continuous monitoring requires ongoing source curation, data validation, and domain expertise that most teams underestimate. Organizations that have solved for it develop a view of their market that compounds over time. Those that haven't are always catching up.

The technology behind trustworthy intelligence

Solving for visibility, trust, and continuity is a technology problem as much as an organizational one. The market is full of tools describing themselves as AI-powered, intelligent, or agentic. Here are the concepts worth understanding.

1 Data

The data an AI system draws from determines the quality of everything it produces. General purpose AI tools can search the web when prompted. What they cannot do is monitor a defined competitive landscape continuously, or guarantee systematic coverage of the sources that matter.

Competitive intelligence requires two types of data that general purpose AI tools struggle to provide together. Qualitative signals tell you what is happening. [Quantitative financial and trade data tells you the scale and context](#). Validated, normalized financial data is harder to compile than it appears, particularly for private companies, across geographies, and over time.

The volume of sources is not the same as the quality of intelligence. An output that draws from hundreds of sources is not automatically more reliable than one that draws from ten well-chosen ones. Knowing which sources an output drew from is essential to assessing whether it can be trusted.



FINE-TUNING AND GROUNDING

Fine-tuning adjusts the model itself, training it on domain-specific material, so it develops a more accurate understanding of a particular field. A model fine-tuned on financial and competitive intelligence data develops capabilities a general-purpose model lacks. It learns to recognize domain-specific terminology and apply structured, analyst-trained logic to specific tasks.

Earnings sentiment scoring, for example, assigns signals based on evidence, such as year-over-year data and management commentary, producing outputs that are comparable across companies and time periods.

Grounding controls what the model draws from when generating a response. Rather than relying on training data or open web search, a grounded AI pulls from a curated knowledge base.

2 Architecture

Data quality matters. So does the architecture built around it.

Large language models are powerful, but they are also computationally expensive and prone to inconsistency. The same input does not always produce the same output. At the scale required for continuous competitive monitoring, unpredictability and cost add up.

Well-designed systems match the right component to each task. Large language models handle synthesis and drafting. Smaller, faster models handle classification and sentiment analysis. The result is more consistent, lower cost intelligence that can be traced and verified more reliably.

3 Governance

CMI data can be competitively sensitive. Access controls determine who sees what. Data residency determines where intelligence is

processed and stored, a critical consideration under GDPR and similar regulations. Licensed content fed into general purpose AI tools may violate content agreements. By the time it surfaces, the liability has already accumulated.

4 Traceability

Traceability is not a compliance requirement. It is what makes intelligence defensible.

It requires systems that record what an AI produced, what it drew from, and when. It also means being transparent with the reader about what the intelligence is based on: which sources, from when, and how it was generated.

This becomes more important as the volume of AI-generated intelligence increases. More outputs means more to verify. Without systems built to handle verification at scale, the intelligence program becomes less trustworthy as it scales.

AGENTIC AI AND MCP

Agentic AI refers to AI systems that can plan and execute multi-step tasks autonomously, rather than responding to a single prompt. In practice, an agent can monitor sources, identify relevant developments, trigger further research, and deliver a packaged output, all as part of a configured workflow. For CMI teams, that means intelligence can be delivered continuously and automatically, without someone manually initiating each step.

MCP, or Model Context Protocol, is an emerging standard that allows AI systems to connect to external data sources in a standardized way, making data accessible inside tools like Microsoft Copilot, Claude, ChatGPT and other enterprise platforms. It is a connectivity standard, not a data quality or governance solution.



What agentic AI means for CMI

Agentic AI changes what it means to deliver intelligence at scale. CMI teams have an opportunity to lead this shift. Here is what that takes.

1 Set the intelligence agenda

Know what your organization needs reliable intelligence on. Competitors, markets, themes. Today and as your strategy evolves.

Which competitors matter? Which markets? Which technologies? Which signals are leading indicators of something strategically significant?

Part of setting the intelligence agenda is thinking ahead to what will matter next. Adjacent industries, emerging technologies, early shifts in consumption patterns or regulation. These are the signals that matter before they become obvious. Automated monitoring makes it possible to watch a broader landscape than any team could cover manually, making it more likely you catch weak signals before they become urgent. The questions that define what to watch are better when they come from people who understand the business and its strategic priorities.

The intelligence agenda gives everything else its direction: what gets monitored, what gets surfaced, who receives it and what gets built and maintained over time. A competitor profile updated continuously as new signals emerge is a different resource from one generated on demand

An intelligence agenda helps turn one-off analysis into organizational knowledge.

2 Think strategically about intelligence delivery

Using AI for your own work is one thing. Making intelligence available through agents requires a different level of involvement. Someone who understands the intelligence, the data, and the organization's needs has to be in the room when agents are configured.

That means building relationships with IT and thinking about intelligence delivery at a corporate level. It also means developing AI literacy and understanding enough about agents and MCP to have informed conversations before a specific project forces the issue.

3 Get close to the business

The framing, the depth, the format, the timing, the channel. When you understand your business and stakeholders, you're better placed to get all of these right.

The more closely intelligence is tailored to a specific decision, at the right moment, for the right person, the more likely it is to move a

decision. Different decisions need different formats.



Earnings analysis delivered within hours of the call



Competitor positioning updates after major announcements



Monthly strategic summaries



Quarterly competitive landscape reports



Financial KPIs normalized for benchmarking performance

AI can draft and refresh these formats at scale. Deciding which formats matter, who they reach, and how to frame what they contain is a judgment call for the people who know the organization.

The closer you are to the decisions being made, the more useful your intelligence becomes.

4 Use AI for scale, own what only humans can do



AI handles the work that benefits from scale and consistency.

- Continuous monitoring.
- Searching and retrieving relevant content across large source sets.
- Structuring and classifying incoming intelligence.
- Producing first draft summaries and analysis.
- Visualizing data.
- Refreshing intelligence assets.
- Distributing outputs on a configured cadence.

Decision makers learn to rely on CMI because what comes from it has been verified, contextualized, and owned by someone who understands both the intelligence and the business.

CMI teams that do this well become the intelligence function decision makers rely on.



Humans handle what AI cannot.

- Judging what a signal means in the context of the organization's strategy.
- Deciding what matters most for which decision maker.
- Validating outputs before intelligence reaches an executive by catching errors, checking sources, and assessing whether the so-what is right for the audience.
- Building the stakeholder relationships that determine whether intelligence actually gets used.



Valona is built for exactly the challenge this paper describes. Purpose-built intelligence infrastructure, combined with human analyst expertise, to deliver intelligence organizations can trust and act on.

Continuous monitoring across 200,000 curated sources, covering markets, geographies, and languages.

So your competitive landscape is covered continuously, not only when someone asks.

Qualitative signals combined with validated financial and trade data.

So intelligence reflects both what is happening and what it means financially.

Every output traceable back to its source.

So intelligence can be verified and defended when it counts.

Agentic AI that automates monitoring, analysis, and intelligence delivery.

So intelligence workflows run continuously, freeing your team for higher value strategic work.

MCP connectivity that makes Valona intelligence available inside Microsoft Copilot, Claude, ChatGPT, and other enterprise AI platforms.

So structured analysis reaches people in the tools where decisions happen, already up-to-date and ready to use.

Trusted by leading organizations globally

Forrester Wave™ Leader, Market and Competitive Intelligence Platforms, Q4 2024. Gartner® Magic Quadrant™ Leader, Competitive and Market Intelligence Platforms, 2026.



About the authors



KRISTA KAUPPINEN

Product Marketing Manager

Krista has a long track record of diving deep into SaaS products to understand how they solve customer pain points. Known for her sharp product sense and endless curiosity, she drives product marketing at Valona Intelligence and helps shape the company's own market and competitive intelligence function.



EETU LAAKSONEN

Chief AI Officer/Chief Technology Officer

Eetu has been part of Valona since its earliest days, building the technical foundation that now powers one of the world's leading competitive intelligence platforms. As Chief AI Officer he leads Valona's AI development, including the agentic framework and MCP integrations that bring validated market intelligence directly into enterprise AI workflows.



STUART REYNISH

Chief Product Officer

Stuart has spent his career at the intersection of technology and market intelligence. Before joining Valona, then M-Brain, he led business development, marketing, and content services at AAC Global, and held management roles in business and process development at Nokia Networks. At Valona he shapes the product vision for how AI and intelligence infrastructure combine to give enterprises a genuine competitive edge.

Sources

- BCG, The Corporate Strategy Function in an AI-First World, March 2026. [bcg.com](https://www.bcg.com)
- McKinsey & Company, AI Data Readiness: The Key to Scaling Impact, June 2026. [mckinsey.com](https://www.mckinsey.com)