

The answer engine

The evolution of modern data platforms

In the age of Generative AI, the modern data platform is becoming an answer engine – a system designed not just to store information or surface it, but to interpret it, explain it and present it in a form that can be acted on immediately.

Every industry has its preferred analogy for explaining what a data platform is and does.

For a long time, the primary comparison was with a library – a carefully curated space where information is neatly organized and easy to browse. As digital ecosystems grew more complex, the analogy evolved: a data platform should behave more like Google’s search platform, indexing everything and allowing users to surface all manner of information instantly.

Both analogies were very much of their time, reflecting the way organizations once engaged with their data. Neither accurately represents what a modern data platform needs to do today, however. So what is the right analogy for a data platform in the current era?

The library analogy made sense when data teams were small and highly specialized. Analysts, architects and domain experts had both the time and the expertise to browse repositories, interpret definitions and validate what they found.

However, as organizations have scaled and data has become essential across every function, that model simply no longer holds.

Expecting business users to browse, compare and interpret datasets is unrealistic. In practice, a library-like platform pushes the complexity of

discovery onto the user, often leading to duplication, misinterpretation and a general reluctance to engage with the data estate at all. In an environment where speed and clarity are critical, the library analogy increasingly feels obsolete.

For many years, “Google for data” was a meaningful definitional step forward. Instead of wandering through datasets, users could type a question and let the platform search, classify and prioritize on their behalf.

That shift created genuine value and improved accessibility. This still reflects an older mode of interaction, however – one where users are expected to know what to ask, review the results, and decide which source to trust.

For many people across the enterprise, that process has become an unwelcome and unnecessary friction – but many financial services organizations are still navigating this Google phase.

Over the last decade, banks and insurers have invested heavily in searchable catalogs, metadata tooling, enterprise glossaries and lineage visualization. These capabilities have helped bring some structure to large, fragmented and often aging data estates – while also reinforcing a particular way of thinking, namely the belief that the main challenge is finding data.

Generative AI: changing the data game

Users now expect the platform not just to help them locate information, but to convert it into knowledge on their behalf, and business expectations have already shifted accordingly. Firms that remain anchored in search alone will find that their data platforms no longer align with how people want or need to work. It is not just about helping users find answers, but enabling systems to produce them.

GenAI is now increasingly embedded across processes, customer journeys and decision workflows, and end users increasingly expect platforms to interpret, contextualize and synthesize information automatically. With GenAI, interaction becomes conversational rather than procedural. A user asks a broad question, perhaps even a poorly formed one, and the system responds with something coherent, contextual and actionable.

This is not a refinement of search. It is a transformation of purpose. Where search gives you options, GenAI gives you greater understanding and insight.

As organizations push further into automation and AI-assisted decisioning, the data platform becomes more than an index or repository. It becomes the

foundation for reasoning – to ensure that the conclusions produced by machines are grounded in consistent, governed and high-quality data.

The analogy that best reflects today’s needs is no longer a library or even a search engine. The modern data platform is becoming an **answer engine**: a system designed not just to store information or surface it, but to interpret it, explain it and present it in a form people can immediately act upon.

This elevates the importance of data architecture rather than diminishing it. For an answer engine to be effective, the underlying data must be:

- consistent and well modeled
- governed end-to-end
- traceable through clear lineage
- machine-interpretable and semantically coherent.

The platform’s intelligence rests entirely on the quality of its foundations. And as more organizations adopt AI-driven decision automation, the ability for the platform to reliably generate understanding becomes a decisive capability.

Conclusion

Historic analogies helped us make sense of the early stages of data maturity, but they no longer describe where organizations are headed. A modern data platform is evolving into something more powerful and more aligned to how people now engage with information: a system that can reflect, evaluate and communicate meaning at scale.

The firms that embrace this shift from search to synthesis, from discovery to understanding, will be positioned to unlock the full potential of AI and build data estates that genuinely enable better, faster and safer decision-making.

Author

James Baldry
Managing Principal
james.baldry@capco.com

Contact

Chris Probert
Partner
chris.probert@capco.com

