SYMBOLOGIC

WHITE PAPER

Symbologic Librarian:

Unlock Data Insights while Maintaining Control of your Data and Costs



© 2026 Symbologic, Inc. All rights reserved.

Executive Summary

Symbologic Librarian helps you get real value from the AI tools you use. It does so quickly, securely, and cost-effectively, while maintaining your span of control over data use and access. Librarian accomplishes this by letting you ask questions of your own files and data through a simple conversational prompt to provide accurate answers based on your company's data as well as your own. Librarian also works with your files and data where they already live, so your data never goes outside the perimeter you already control.

According to Deloitte's 2026 Global Technology Leadership Study, technology leaders are facing growing pressure to show return on their AI investments. "Leaders are now expected to deliver measurable enterprise value across growth, productivity, and customer impact," reported the company in its April 2026 study. "The data shows this is not an emerging trend, but the new standard, as 79% of leaders report it is their top priority. However, 42% report low or no ROI on AI investments, demonstrating the difficulty of delivering on this new value mandate."

Librarian enables your company to provide secure access to its files and data, where they are stored today. It does not require moving to the cloud or another document repository. This allows you to maintain control of your data and avoid yet another migration to new purpose-built AI infrastructure. Employees go to company data, so why not have AI tools do that as well?

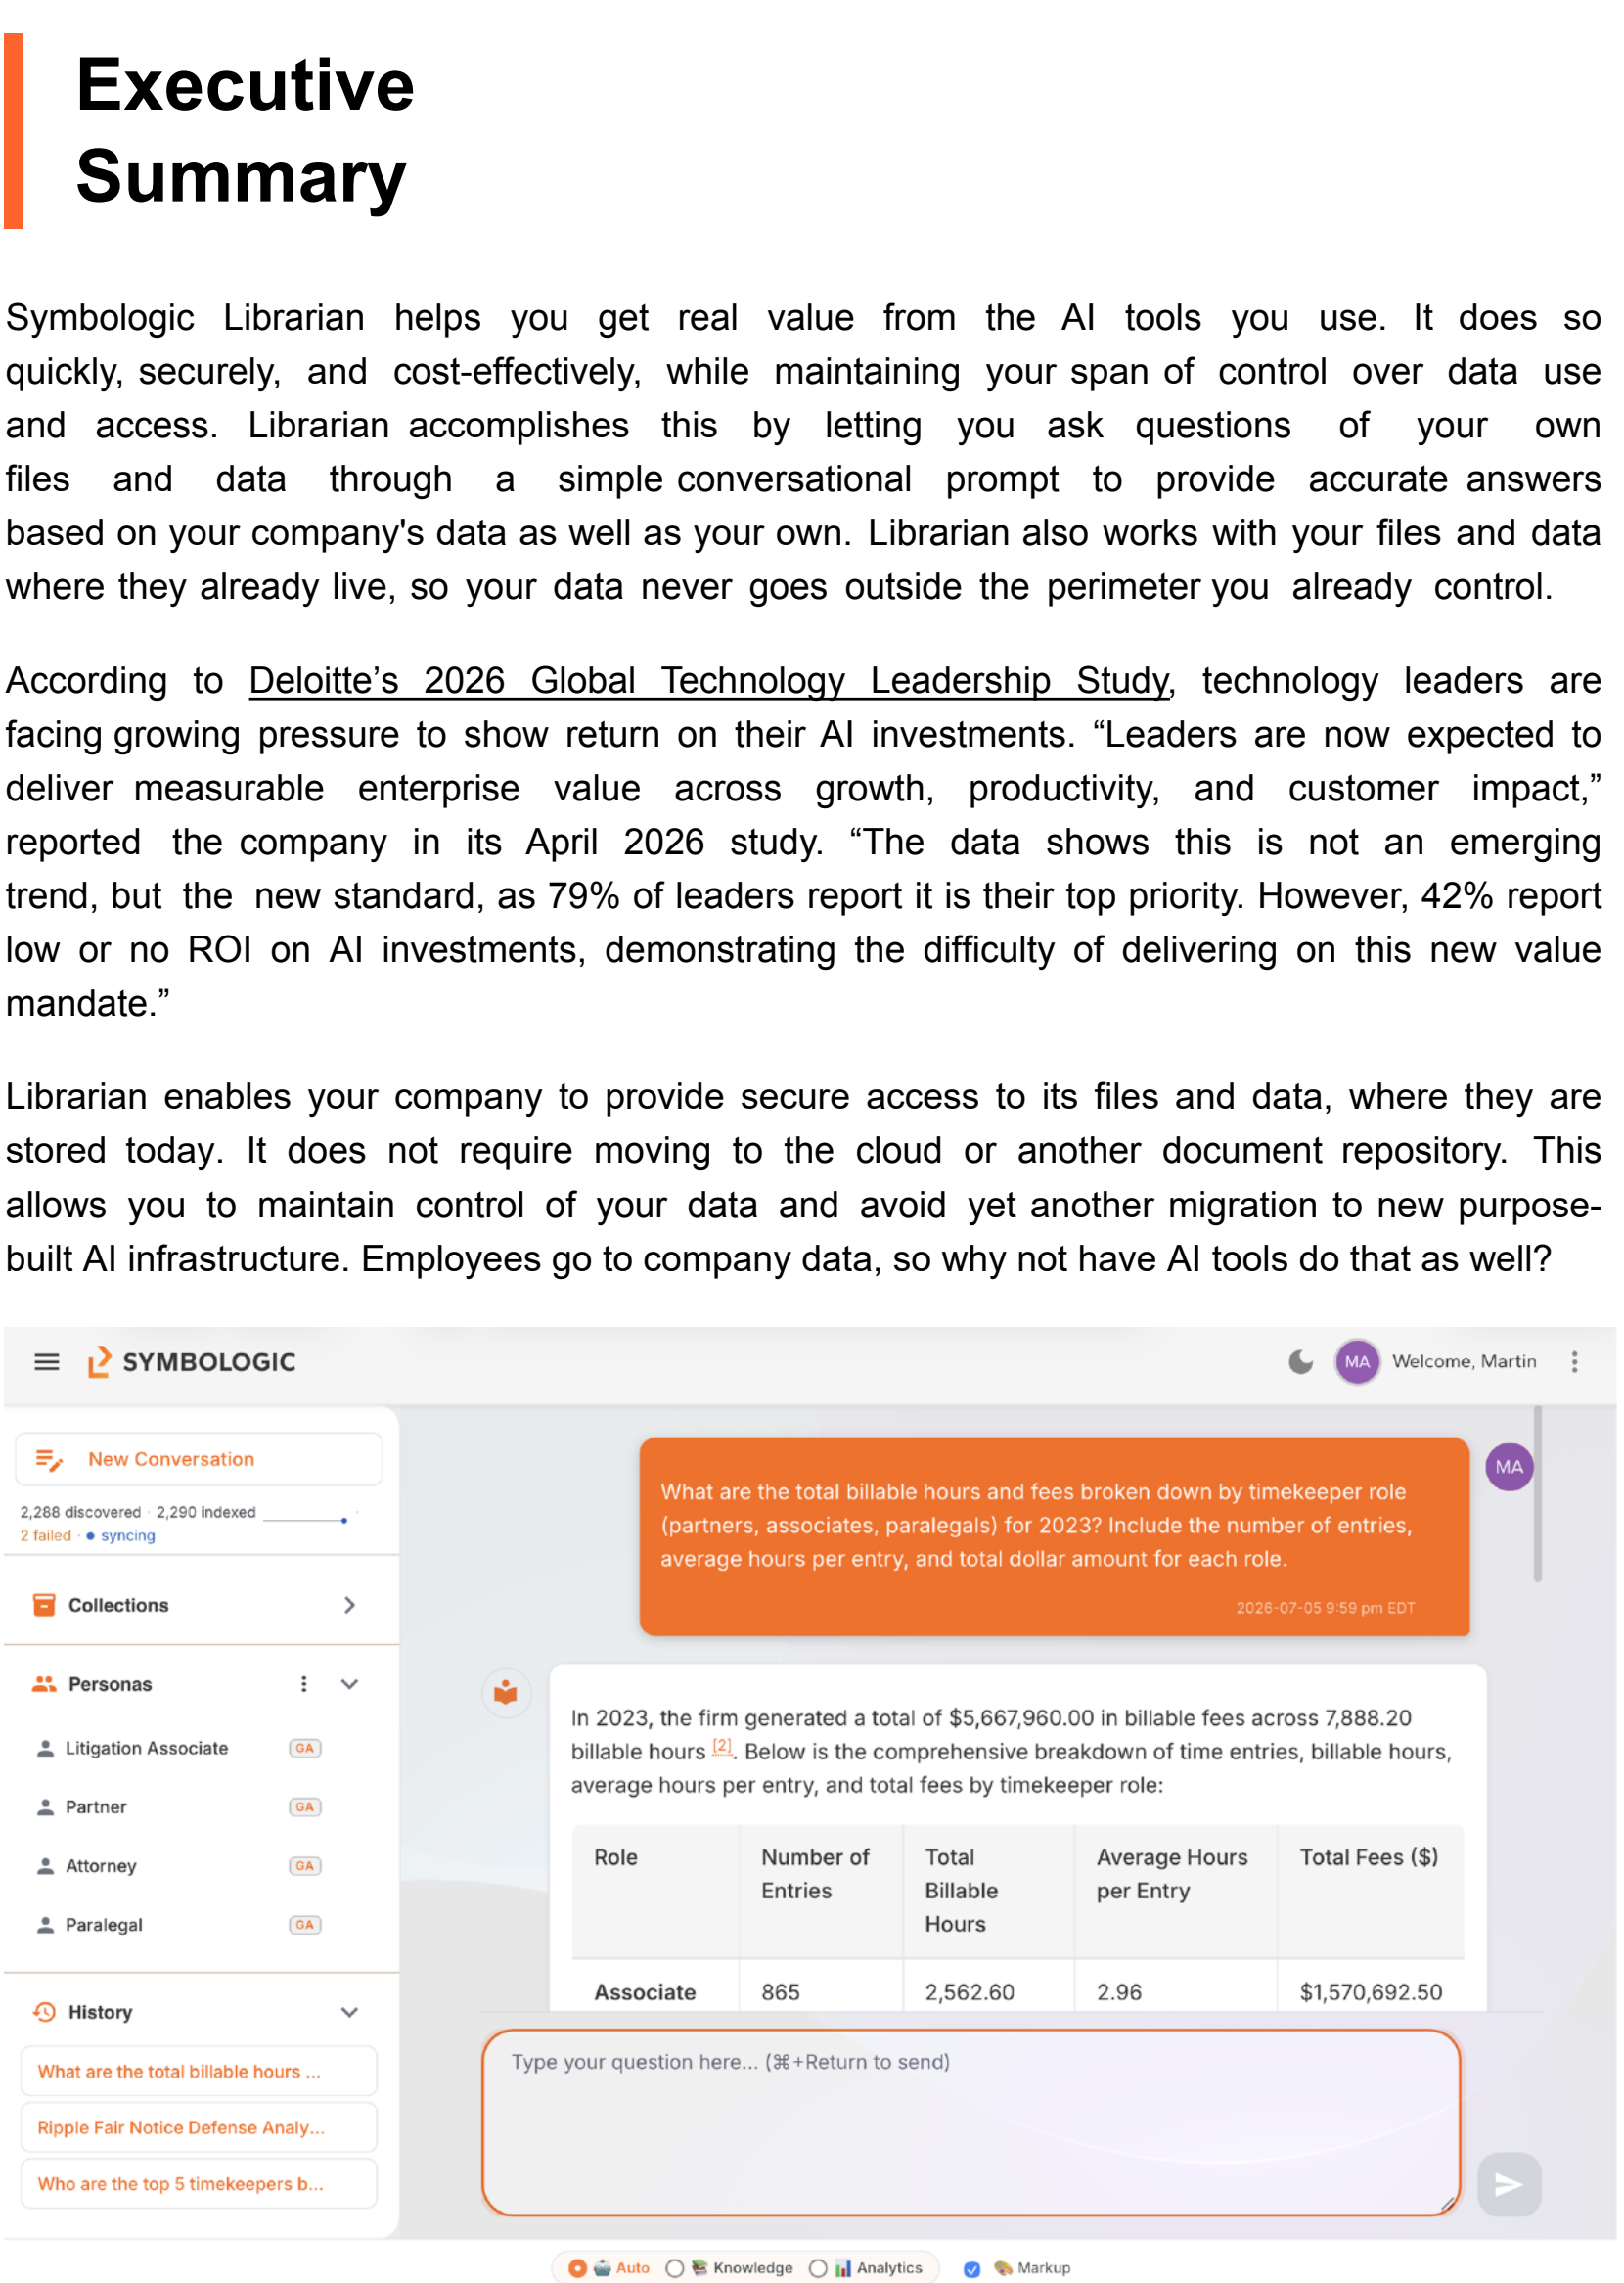


Figure 1. Conversational Interface: Symbologic Librarian's Cited, Tabular Answers

© 2026 Symbologic, Inc. All rights reserved.

Our data-in-place approach ensures that Librarian works with both your unstructured and structured data (such as databases, reports, and spreadsheets) to provide a single answer. Every response is grounded in citations, and every query is classified and handled according to how your company policies require (publicly-available data, for example, is treated differently than data relating to proprietary intellectual property). Librarian supports regulated industries that need traceable answers, not chatbot improv.

Librarian offers quick, measurable return on your AI investments by routing queries that reduce token costs by sending prompts and queries to the most appropriate and cost-effective models.

What real-world problems can Symbologic Librarian solve?

We designed Librarian to address three important challenges you face every day in using AI:

- how do you ensure your data remains secure while using AI tools?
- how can you better control the cost of using AI tools?
- and how can you deliver consistently accurate results?

Security

Most organizations have sensitive files and data that require extra care. Whether those files and data are customer records with personally identifiable information (PII), employee files, corporate strategy documents, forecasting and planning projections, sales numbers, or proprietary intellectual property, **they need to be secure and audited.**

For certain regulated industries, specific requirements mandate how files and data must be secured, such as:

- the medical devices and healthcare industry (HIPAA, GDPR, FDA 21 CFR 820, ISO 13485, PMA/IDE post-market surveillance)
- the legal and professional services sector (privileged matter handling, conflicts, FRCP-compliant audit trails)
- finance and banking (SEC reporting, SOX 404(b) ICFR, MNPI protection, MAR/MiFID surveillance)
- manufacturing and engineering (ISO 9001/13485 QMS, CAPA, supplier quality, IP protection)
- and the public sector and defense industries (FedRAMP, CMMC, ITAR-aware tiering).

© 2026 Symbologic, Inc. All rights reserved.

Yet organizations across these sectors look for ways to benefit from AI tools with predictable economics.

Popular AI tools grew from grassroots efforts by well-meaning employees, but these point solutions are not enterprise-ready and were not designed to protect data or enforce who can see what. To meet the need for protection and manage access, regulated industries need an approach that works alongside your data and the AI tools you are already using to provide appropriate controls.

These controls must be built on the tiered governance which already runs your non-AI operations. AI needs to conform to that governance and compensating controls. Symbologic delivers Librarian to do just that with you. We offer a suggested file and data classification as an example.

The four tiers outlined below represent a foundational data classification taxonomy. Enterprises leverage this framework to systematically categorize data assets by sensitivity, govern access privileges, and implement commensurate security controls. While specific classification schemas may vary to accommodate distinct industry verticals, regulatory mandates, or internal governance policies, this four-level model serves as a widely adopted industry standard.

Table 1. Data Governance Model: Tiers and Permitted AI Use

Tier	Data types	Permitted AI use/control
Tier 1: Publicly available information	Existing public information such as press releases, marketing materials, published research, public company filings (such as 10-K, 10-Q and 8-K).	Can be routed to any reputable public AI tool. Control is provided through awareness training to prevent sensitive input.
Tier 2: Internal Business Operations	Internal non-confidential communications, standard operating procedures, policies and administrative content.	Only routed to approved enterprise AI tools via single sign-on (SSO). Controls provided by automatic blocking or redirecting of prompts or queries if Tier 3 or Tier 4 data is detected.
Tier 3: Confidential	Proprietary intellectual property, trade secrets or sensitive logistics information. Includes customer information with personally identifiable information (PII), employee files, corporate strategy documents, forecasting and planning projections, sales numbers, or proprietary intellectual property.	Isolated and used exclusively with "private Librarian" instance. Control is provided by management of the private Librarian instance, use of a firewall and data loss prevention guardrails.
Tier 4: Restricted and regulated	Applies to the medical and healthcare industries, legal and professional services sector, finance services and banking, aviation, manufacturing and engineering and public sector and defense industries.	Purpose built, regulation-specific AI. Control provided by mandatory "human in the loop" sign-off, cost attribution and audit logging.

This governance structure ensures that all AI utilization, regardless of classification tier or deployment zone, is systematically channeled through one approved Librarian. Acting as a central orchestrator, the Librarian subsequently applies the appropriate, tier-specific compliance and security protocols.

© 2026 Symbologic, Inc. All rights reserved.

Cost Predictability

When implementing AI across their company, enterprises focus on predictable costs as well as data access and protection. Predicting costs must anticipate the "spend creep" of tokens that represent the systems use of large language models (LLMs).

What enterprises are learning, some more poignantly than others, is that AI costs seem to be unpredictable.

A January 2026 report from Gartner summarized this issue well. It said: “Rising costs kill projects even when they’re technically successful and delivering user value. That negligible per-token cost becomes a total cost of ownership (TCO) nightmare when multiplied across thousands of users and hundreds of use cases.”

The report went on to suggest what a solution to that issue might look like. “Educate all stakeholders ... not just IT ... on the cost implications of deployment approaches, model selection and token usage. Avoid expensive model customization when unnecessary. Implement prompt caching to reduce redundant API calls, and use model routing to route queries to appropriately sized models. Monitor costs continuously with proper allocation and visibility tools,” it said.

Consistent, Dependable, Accurate Responses

In regulated work, an answer that’s “usually right” is a liability. For AI to be used broadly and deliver real value, organizations must trust the results it produces and not pay for the same results more than once.

That’s why you want to be able to set Librarian to be an approved front door for all AI use, across every data tier. Librarian can be configured to re-route employees' browsers from public AI directly to their own Librarian instance for all AI-assisted discovery of the insights that come from reviewing vast amounts of otherwise disjointed files and data, thus closing the "shadow AI" gap on every tier, providing one governed pipeline.

Finally, organizations face the challenge of what happens to the answers provided by AI tools in an AI conversational interface. The best way to ensure that answers are handled in the most effective manner possible is to use a “classify → reason → protect” lifecycle that includes protection of input files from company file servers and of the AI output of queries. This lifecycle involves:

- classifying fles and data on their way into Librarian
- taking a “reason in place” approach to working with the files and data related to the topic of each query (meaning that your files and data never leave your controlled perimeter and is routed by both intent and protection tier or classification to the task and model matched to predicted cost and accuracy that best meets your need)

© 2026 Symbologic, Inc. All rights reserved.

- ensuring that all outputs from approved AI tools automatically inherit the exact security and access controls of their original source data

The only real way to address all of these challenges is to make sure that:

- your AI tools use only the files and data you provide them
- your files and data never leave your control
- the AI tools you choose are appropriate for the role of the person using them
- you use model routing to maintain granular control of where queries are sent, ensuring the most appropriate models (for cost and accuracy) are used for the right job

Symbologic Librarian does all of this by pulling your structured and unstructured enterprise files and data into a unified, securely indexed corporate memory where it can be discovered and reasoned over **in place, never copied out of your span of control.**

Every request is classified by intent and by classification tier. Then the intent router within Librarian sends it to the LLM that can most **accurately and cost-effectively** deliver the answer, with the tier-appropriate control applied automatically, including file protection (encryption, access control, watermarking) for your most sensitive data that travels with the answer even after it leaves the screen. Librarian grounds every answer in your own files and data and cited to its source in the company's file and data repositories. In all cases, your data never leaves your control, and no data outside your organization is ever used to produce the results your team sees.

How does Symbologic Librarian work?

Symbologic Librarian helps organizations get real value from AI tools they control. Librarian provides that value securely and cost-effectively.

It all starts with your data, which Librarian does not require you to move. That data can be structured or unstructured and come from a variety of sources, including ERP systems, databases, and cloud-based or on-premise document repositories (such as shared drives and files, SharePoint sites, documents, PDF files, presentations, and spreadsheets).

Librarian sits on top of the data you already have and turns it into answers your teams can act on. It works with your data to securely provide your employees with answers, as they need them and without interrupting the flow of their work.

© 2026 Symbologic, Inc. All rights reserved.

Not only does Librarian provide access to AI capabilities in a way that will be familiar to your teams, but it also ensures that you maintain the necessary control over the data throughout.

That level of control also extends to the cost of using particular AI models, as Librarian includes an “intent router” that ensures queries are routed accurately and effectively.

Users get the most useful answer as quickly and cost-effectively as possible and answers are properly sourced with accurate citations.

Every query is routed automatically to one of three modes based on the nature of the data:

- RAG (Retrieval-Augmented Generation) for unstructured files,
- BI for structured data via natural-language SQL, or
- Fusion for synthesized answers that combine both.

The key here is that AI queries are a mix of analytical and semantic. Some questions require numerical calculations (analytical) and others need data from documents and databases (semantic). RAG searches an external set of documents or databases for the right facts.

If calculations are involved, Librarian sends those questions to a business intelligence (analytic) engine to do the calculations. The intent router decides where queries go.

Once facts are found and calculations are made, they are passed to an AI model to write an accurate answer.

So no-one should have to suffer from “I asked the wrong system” frustrations.

© 2026 Symbologic, Inc. All rights reserved.

How does Librarian integrate with my existing infrastructure?

A core design principle of Librarian is its ability to securely and effortlessly integrate into your existing ecosystem. Librarian is built to:

Self-host on your infrastructure: Symbologic Librarian is not SaaS. Deploy it on your infrastructure of choice, on your virtual cloud, your preferred public cloud provider, datacenter, or an edge appliance in your data closet. Your files never leave your control.

Use the right LLM for the right task: Librarian's intent router supports connection to a broad range of leading models, including OpenAI, Google Gemini, Anthropic Claude, Mistral, or self-hosted open-weight models. It is designed to ensure that your queries are sent to the LLM that makes the most sense (and is the most cost-effective choice) to provide the answers you need.

Provide secure identity management: Supports SSO and federated identity using Auth0, Okta, Microsoft Entra, SAML 2.0, or OIDC.

Streamline security operations: Integrates with SIEM platforms such as Microsoft Sentinel and Splunk Enterprise, featuring a dual-attribution audit log.

Be extensible & Customizable: Our open architecture and open APIs give you the freedom to integrate Symbologic into your own infrastructure. Implement it yourself using our jump-start AI Agent Instructions, work with your existing IT solution provider or with a Symbologic preferred partner.

© 2026 Symbologic, Inc. All rights reserved.

It does not require you to rip-and-replace any hardware or move to a new data platform. You just:

- connect your existing sources
- select the AI model(s) you want to make available to your teams
- allow them to ask questions of your structured and unstructured sources.

Librarian blends unstructured files with structured data; you use only the “pluggable” AI models you want to support, so you control cost and sovereignty. The diagram below walks through the journey of a user query or prompt in Librarian.

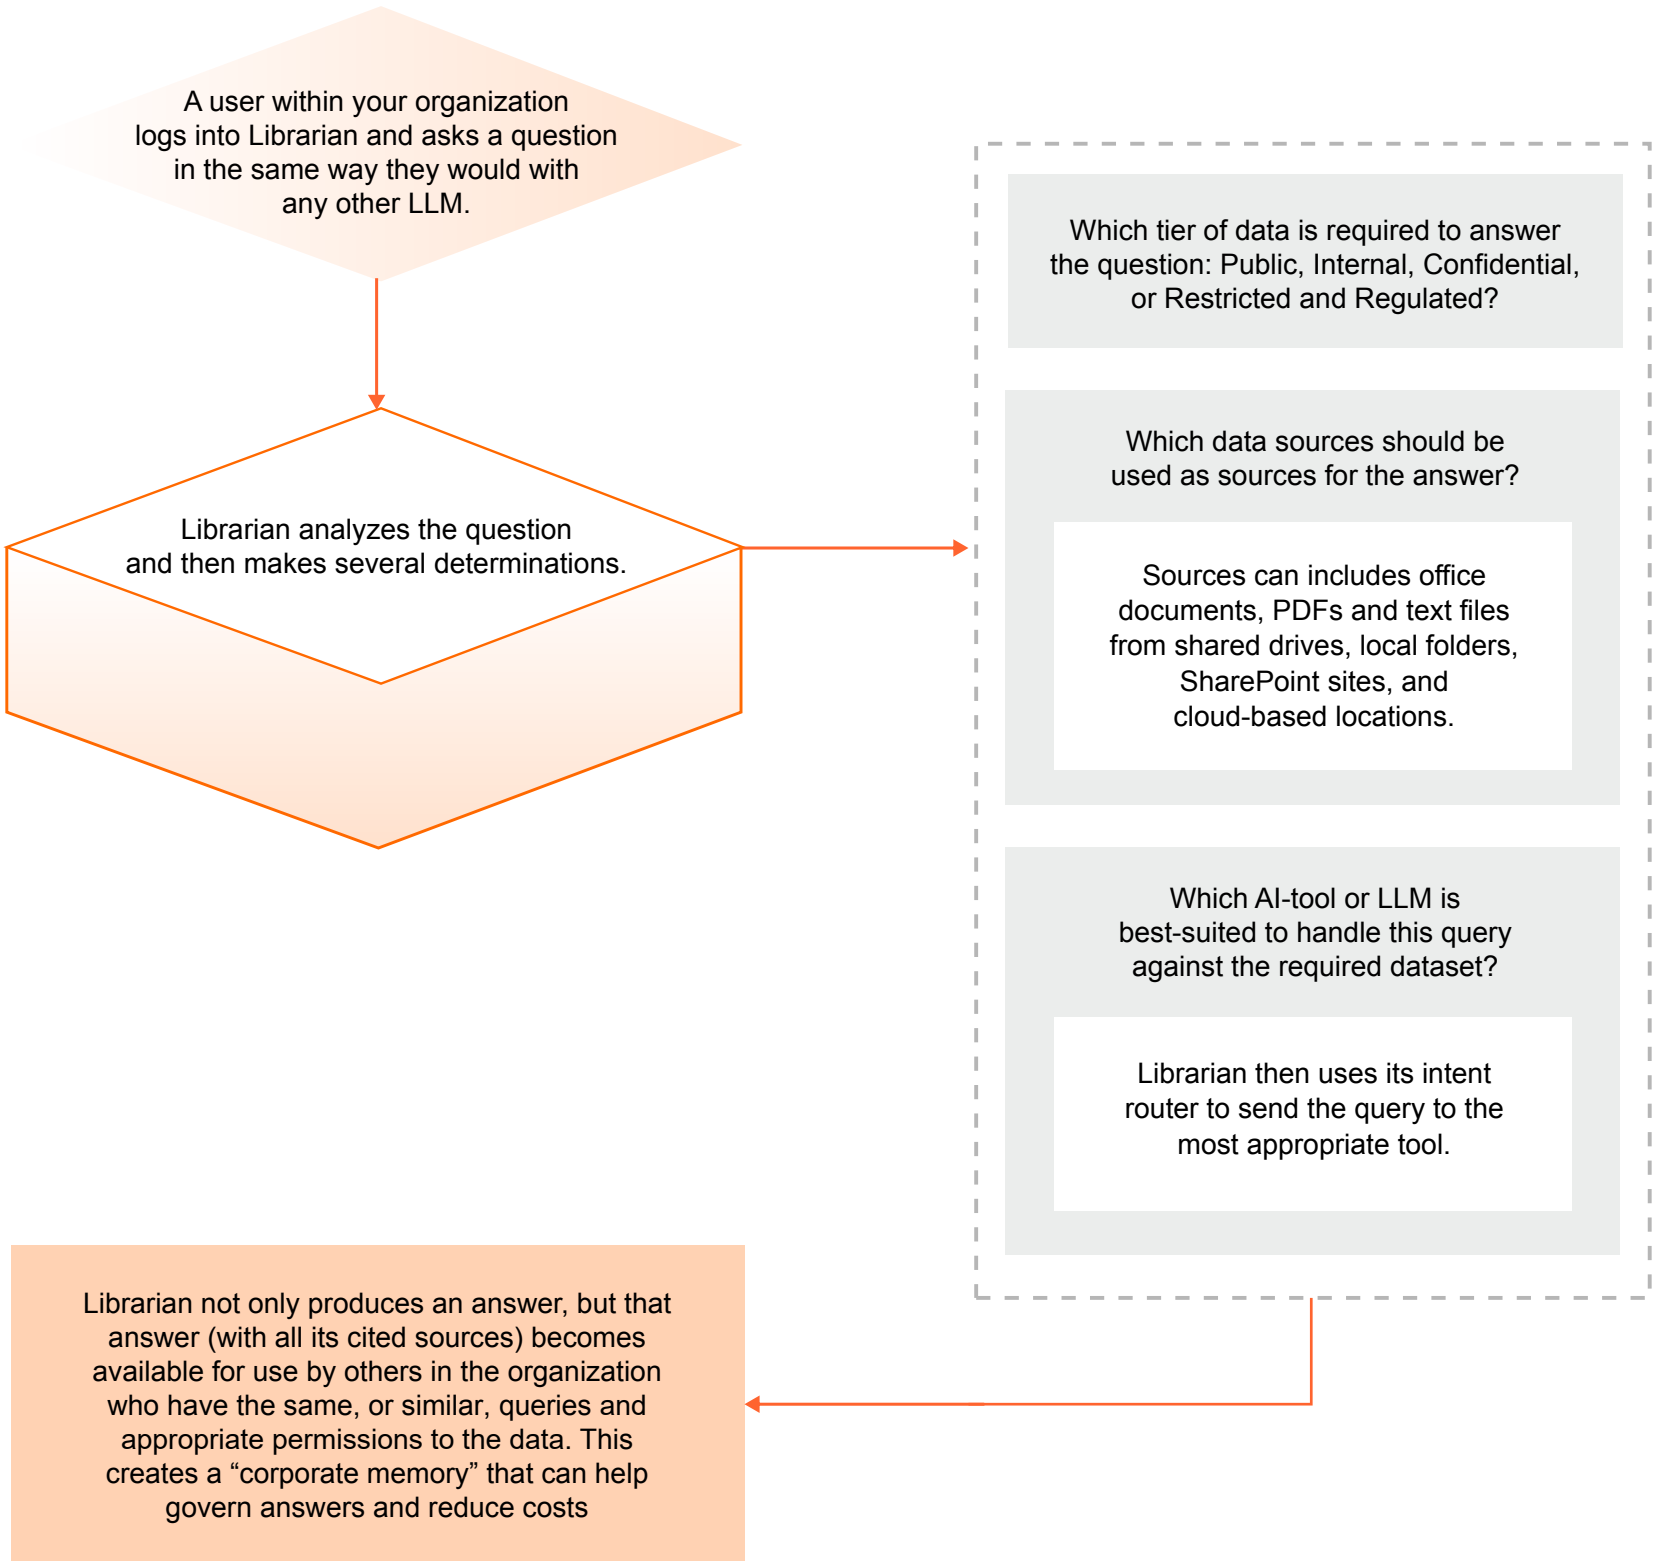


Figure 2. Query Journey: Classification, Routing, and Corporate Memory

© 2026 Symbologic, Inc. All rights reserved.

Next steps

Take a look at the real-world benefits Librarian can provide. You'll have confidence that you won't release sensitive files or data to a public model. You'll maintain tight control over that sensitive data to support compliance with regulatory frameworks.

You will also know that the answers you get from your AI prompts and queries will be grounded, cited, and audit-ready. Those answers will become the foundation of a compounding corporate memory that preserves all the attributes and permissions they inherit from source files and data.

Finally, you'll see potential cost savings from reduced token use and support for less expensive models, with an intent router that will help you choose the most efficient and cost-effective model for the work you want to do.

Contact us at sales@symbologic.ai to get more information or, even better, schedule a demo to see how Librarian can help you.

SYMBOLOGIC

Symbologic's secure, enterprise AI solutions deliver semantically and analytically accurate conversational intelligence within existing client infrastructure. Combining Data-in-Place architecture with zero-trust security, Symbologic AI solutions transform siloed enterprise data into continuously expanding, reusable corporate memory, empowering your employees while reducing the cost of AI.

© 2026 Symbologic, Inc. All rights reserved.