Strategy₿

WHITEPAPER

AI Security

Explore practical steps to ensure data integrity and ethical AI implementations

Table of Contents

Introduction

The effective deployment of artificial intelligence (AI) in business intelligence (BI) significantly depends on the integrity of the underlying data. For AI systems driving business decisions, accuracy isn't merely beneficial; it's imperative. These systems need to be trustworthy to be of genuine utility.

Strategy stands out as a reliable pillar in this context, providing business users with data that is precise and secure. The introduction of Strategy AI reinforces this commitment, offering a platform where the rigor of BI meets the innovative capabilities of AI.

Our AI solution is engineered to accurately interpret business questions presented in natural language, employ logical reasoning, and produce relevant results autonomously. This synthesis of BI's structured analysis and AI's adaptability ensures that Strategy AI meets the dual needs of data integrity and flexible user engagement.

Strategy AI is an evolution of our established platform, seamlessly integrating advanced AI and machine learning capabilities. It streamlines processes such as AI-driven data exploration, dashboard design automation, and the use of specialized tools like SQL generation and machine learning-enhanced visualization for data analysis. With these features, the platform facilitates more profound data insights within the familiar environment of the Strategy ecosystem.

The reliability of Strategy AI is anchored in the meticulous design of Strategy's semantic layer and its comprehensive security framework. Auto, our AI assistant, exclusively relies on data from Strategy, with all analytics executed by our established analytical engine. This ensures consistent, accurate, and secure data processing and representation, allowing businesses to make informed decisions with confidence.

Strategy AI Environment Isolation

The core strength of an enterprise AI solution lies not only in its ability to process data and deliver insights, but also in its architecture's resilience to external threats. The Managed Cloud Enterprise (MCE) and Managed Cloud Government (MCG) are underpinned by an architectural design that places the utmost emphasis on environment isolation, secure access, and safe execution of requests on external or multi-tenant services.

Designing for Isolation

MCE and MCG are strategically architected with environment isolation as a central security tenet. By ensuring that each customer's data operates in a securely segmented environment, we eliminate cross-contamination risks and enhance data protection. When the system needs to connect or submit a request to an external service, these workflows are executed with strict security measures and communication protocols. This includes encrypted data transmission and stateless execution requests within the security context of the customer's instance.

Strategy AI within MCE and MCG

The Strategy Cloud offering is complemented and enhanced by including the Strategy AI module. By anchoring it within the MCE and MCG framework, we ensure that Strategy AI consistently benefits from and upholds the robust environment security standards intrinsic to MCE and MCG.

Characteristics of MCE and MCG's isolation:

- **Custom Configurations:** Strategy can initiate resources into the Virtual Private Clouds (VPCs) and modify internet protocol (IP) address ranges, route tables, network gateways, and pertinent security settings. This ensures that every customer's environment is tailored to their specific needs while upholding security standards.
- **Robust Firewall Implementation:** Each customer's tenant is fortified by hypervisor-level firewalls or security groups. Utilizing advanced cloud and virtualization software, these firewalls further split instances, creating entirely separated client processing spaces. This segregation is instrumental in warding off unauthorized access and ensuring that non-public information remains protected.

In essence, Strategy AI is not merely an intelligent data processing tool; it is a product deeply integrated into a cloud environment where every architectural decision prioritizes user security. The stringent security properties of MCE and MCG further highlight our unwavering commitment to safeguarding the integrity of our customers' data.

Strategy Integration with Azure OpenAI

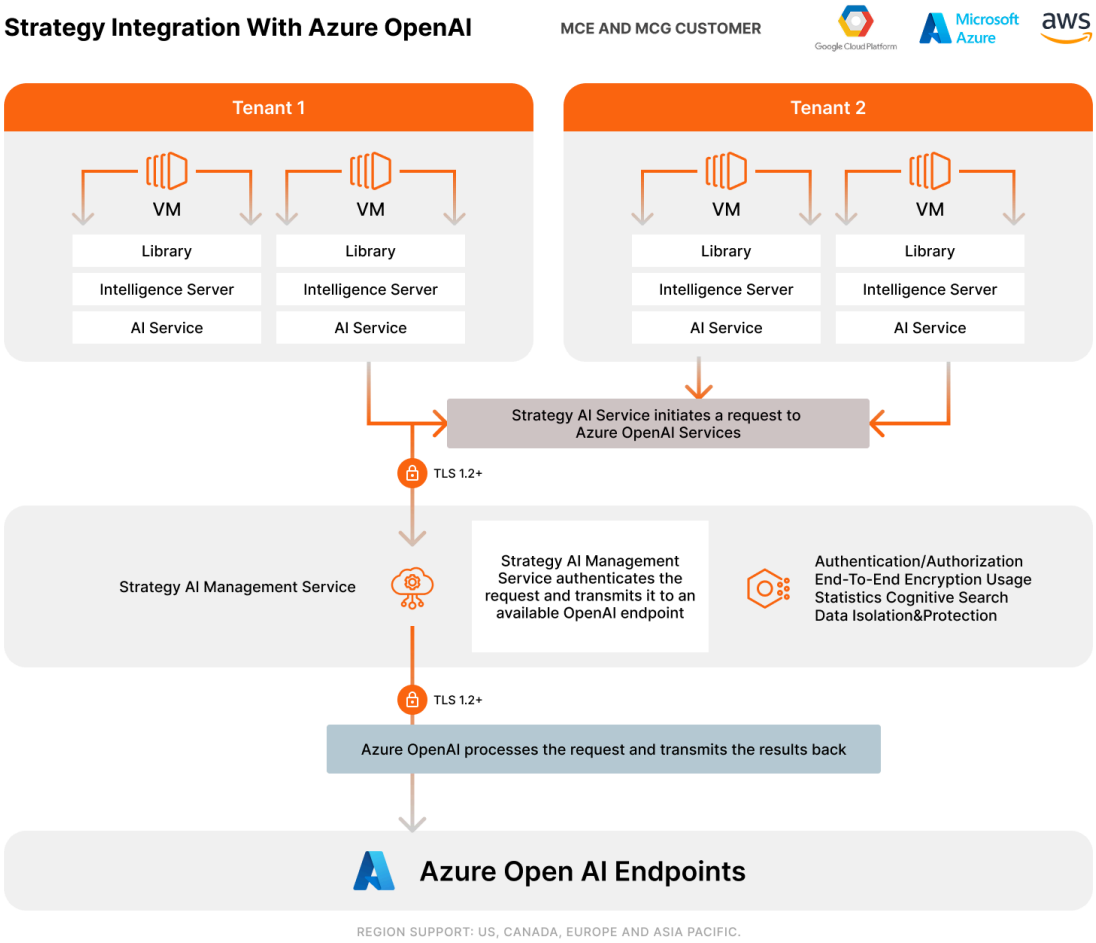

Ensuring Data Privacy and Integrity with Strategy AI

Strategy AI

Strategy AI provides an array of AI features to empower users with varying skill levels and different roles within an organization. Business users and analysts can take advantage of Agents which allow users to interface with their data through the use of natural language. Agents provide a chat experience, on top of specific structured or unstructured data and allows for customization by leveraging custom instructions and AI-generated metadata for additional business context. Features include Auto Answers which includes advanced Q&A and AI visualizations that leverage machine learning capabilities to generate key driver analysis. Other features available in Strategy AI include Auto Dashboard enables design and Q&A within a dashboard; Auto Narratives, which generates descriptive data analysis; Auto in Hyper, which allows users to interact with Agents via Cards; and Auto SQL, which helps administrators and architects to expedite data modeling by generating SQL.

Each customer who enables AI has a separate tenant environment. This environment connects to Strategy AI Management Services, which forwards AI requests to the LLM for processing.

Strategy AI uses Microsoft Azure OpenAI, and all communication is strictly over secure channels with TLS 1.2 or higher. This ensures that data is always encrypted during transit, preventing unauthorized access or breaches.

The AI Features work strictly within the boundaries set by Privileges, ACLs, and data security measures specified for the user in the Strategy platform.

Knowledge Assets for Agents

Strategy AI releases prior to March 2025 allow users to fine-tune their Agents by uploading contextual data via the Knowledge Assets feature. Uploaded Excel files are processed by the Knowledge Manager, embedded, and stored securely in the Knowledge Store—a vault of encoded domain knowledge optimized for cognitive search and fast retrieval. Only privileged users can upload assets, with strict access control and encryption ensuring data security and integrity.

During Gen AI interactions, the Knowledge Store provides relevant context, enabling accurate and governed query responses. Only schema and minimal sample data are shared with Azure OpenAI's LLM, ensuring privacy. The LLM creates an execution plan, which is processed by Strategy's Semantic Graph and analytical engine. Final results are returned to the user as natural language.

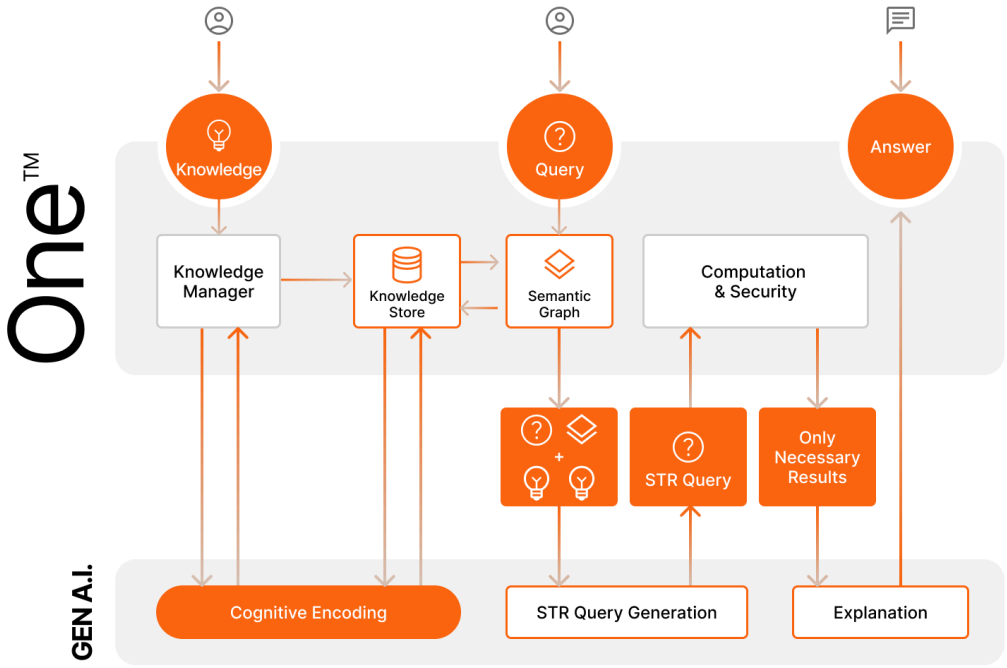

Agent Context and Vector Embeddings for Agents

Strategy AI releases (March 2025 and after) provide next-generation capabilities through a new, agentic architecture for Agents which enhances reliability, accuracy, and takes advantage of the latest technological advances in AI for structured and unstructured data. This new architecture is also available for Agents exposed via Teams and HyperIntelligence.

Agents know about the context of their structured data through AI-generated definitions. Items such as object names, data types, and semantic relationships are sourced from the Strategy semantic layer to generate the definitions, thus providing the greatest consistency. Additionally, Agent creators can further fine-tune their Agent by editing such definitions and additional custom instructions. These definitions are always stored within the confines of the customer's tenant.

Agents built on unstructured data documents (curated by system administrator) depend on textual vector embeddings and the retrieval of relevant parts of the documents to answer the user's question. These documents are always stored within the confines of the customer's tenant.

When an authorized user interacts with a structured data agent, the underlying AI agents detect the user's intent, leverage the stored object definitions, and custom instructions to determine how to best answer the question. Agents review the data results (up to 1,000 rows) to answer the user's question. If the Agent is for unstructured data, then the Agents use the relevant documents' text information to answer the user's question.

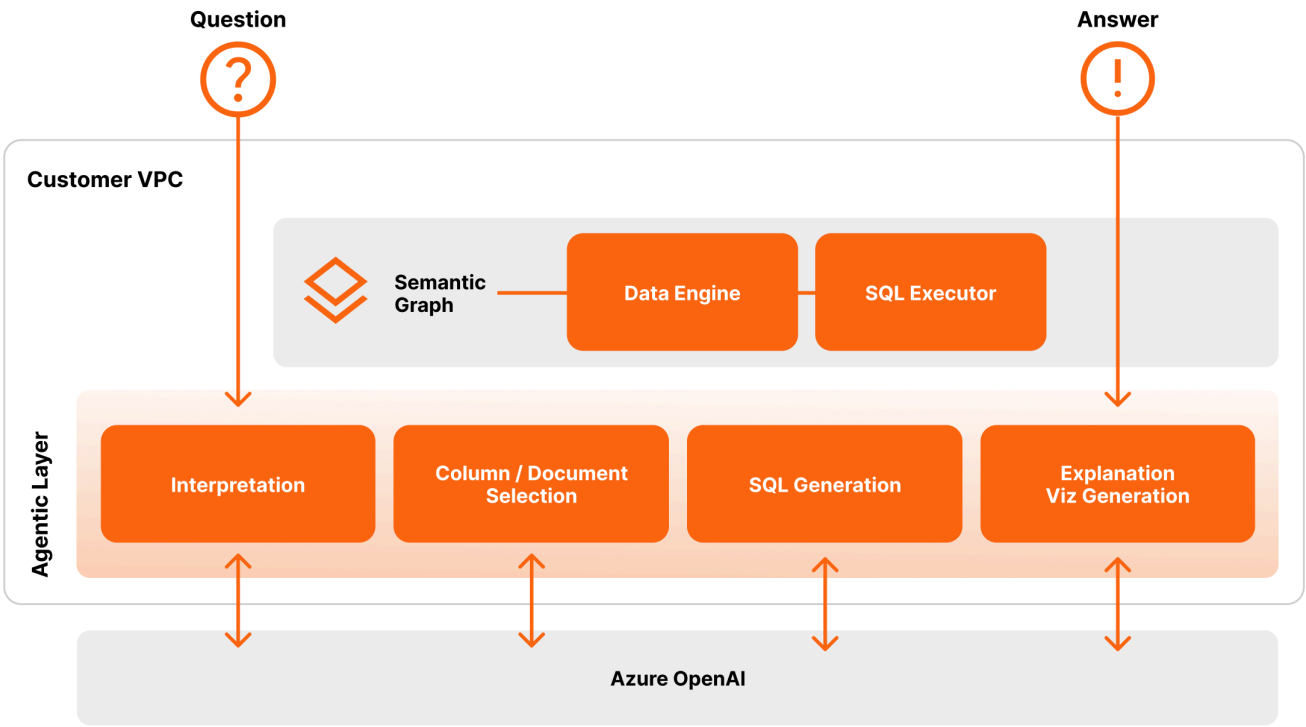

Data Retention

The Agent functionality does not retain conversation histories after the active user session. Strategy Platform Analytics gathers telemetry data on user interactions with Agent to empower administrators. This includes capturing the question and the respective user. Access to this data is managed by limiting the number of users accessing the Auto Adoption and Auto Question Analysis dashboards and their dependent schema objects in Platform Analytics. Platform Analytics is hosted under each customer's tenant.

Historical Data Retention
Strategy Agents enable users to persist their chat history and save snapshots of their data. Access to this chat history and the snapshot information is on a per-user basis. The information is not shared between users. Each user is entitled to retain a maximum of 50 historical questions and answers. Further, this information is stored in tenant's own storage space and is fully encrypted using AES-256 bit encryption.

Manual Deletion of Historical Data
To provide users with control over their data, Strategy AI enables the manual deletion of questions and answers from past conversations. Users can remove entries they no longer need, maintaining a clean and relevant history according to their preferences.

Snapshots
In addition to historical data, users can create and save snapshots of specific questions and answers. These snapshots are stored independently of the standard history and can be used to preserve critical data points or insights. Each user can maintain up to 50 snapshots, allowing for significant flexibility in data management and retrieval.

Data Storage Locations
Text Content Storage: The text content of questions and answers are securely stored in the Strategy Metadata database. This database is designed for high availability and is optimized for efficient data retrieval.

Visualization Data Storage: Data points used to visualize each answer are stored in the Strategy Storage Service. Before users can begin storing visualization data, they must configure the Strategy Storage Service to ensure proper data handling and security.

The Text Content Storage and Visualization Data Storage reside inside each customer's tenant.

User Data Privacy
User-specific questions, answers, and snapshots are private and accessible only to the individual user. This strict access control is part of our commitment to user privacy and data security, ensuring that sensitive information remains confidential and protected.

This data retention policy is part of our ongoing effort to provide a transparent, secure, and user-focused experience, enabling effective and efficient use of Strategy AI.

Regulatory Compliance for the AI Components in the Managed Cloud Enterprise

Strategy AI, integrated with Microsoft Azure OpenAI, has certifications for vital international data protection protocols, including CCPA, GDPR, SOC 2, and ISO 27001. The design and operational procedures of the Strategy AI components within the MCE are tailored around these regulatory benchmarks. Our stringent approach to adherence demonstrates our commitment to meeting and surpassing the standards set by these authorities.

Strategy's commitment to regulatory diligence is systematic and meticulous. We maintain a dedicated internal compliance team responsible for ensuring alignment with industry standards. This team has architected robust data protection and privacy protocols, directly aiming to meet the rigors of the General Data Protection Regulation (GDPR). Strategy confirms full regulatory adherence across all jurisdictions where the Strategy Cloud operates

The MCE, available for AWS, Azure, and GCP is compliant with the following risk management and information security frameworks. Such compliance is regularly validated, and when necessary, certified through rigorous evaluations conducted by internal and third-party professionals:

- General Data Protection Regulation
- AICPA SSAE-18, System and Organization Controls – SOC 2 Type 2 Report
- ISO/IEC 27001:2013 (ISO 27001:2013) – Certificate Number: ISMS-MI-13123
- Data Privacy Framework EU-U.S. and Swiss-U.S.
- Health Insurance Portability and Accountability Act of 1996 (HIPAA) Self-Assessment.

Strategy AI Compliance with the EU AI Act

Strategy AI is committed to complying with the EU AI Act, which aims to ensure the safe and ethical use of AI technologies in Europe. Strategy AI, which incorporates a General Purpose AI (GPAI) large language model, Azure OpenAI, in order to provide AI capabilities to our customers, is a limited-risk AI system based on its base-level capabilities. Strategy AI's capabilities, including Auto Answers, Agents, and Auto Dashboard, are chat-based and always serve a user rather than make automated decisions. In addition, as a deployer of a limited-risk AI system, Strategy AI is purposefully transparent, and Strategy has implemented quality management practices surrounding our Strategy AI product, including:

- Robust safety measures
- Transparency in decision-making
- Clear user information about AI interaction
- Regular assessments and improvements of the AI system

Finally, Strategy AI and all of its intelligent features exclusively rely on Strategy's own trusted calculation engine. Strategy AI does not utilize user inputs, chat history, or other communications to train any LLM, GPAI, or other third-party artificial intelligence models.

Monitoring, Logging, and Auditing AI Use within Strategy

Strategy places a significant emphasis on monitoring, logging, and auditing the usage of AI within its platform to ensure transparency and accountability. The monitoring systems in place offer customers a comprehensive overview of their AI usage. Through a user-friendly dashboard, administrators can gain insight into the users and the number of questions they have asked. This transparency empowers organizations to optimize their AI utilization effectively

Logging Mechanisms

Strategy has designed meticulous logging mechanisms to uphold user data privacy and security. Within the Strategy platform, certain information is logged while others are intentionally left out, with the purpose of safeguarding user data and meeting data compliance laws.

Specifically, Strategy's logging system captures essential data for operational purposes. One example is the number of tokens used for the questions posed by users, ensuring that this information is tracked to measure consumption and usage. Furthermore, this logged data is not used to train AI models or for other purposes that might compromise user data privacy.

This privacy-conscious strategy aligns with contemporary data protection standards and regulations, providing users with the confidence to fully engage with Strategy's AI-powered tools without worrying about the security of their sensitive information.

Audit Trails

Audit trails are crucial in ensuring accountability and traceability within the Strategy platform. Strategy has implemented a robust system in Platform Analytics that allows customers to perform audit trails effectively. One key element of this system is preserving a unique question ID, which enables tracking of actual usage associated with individual users. Importantly, Strategy's approach prioritizes user privacy by not recording the outcomes generated. Instead, by focusing on the question ID, Strategy can precisely determine which user initiated a query or utilized specific AI features. This approach balances accountability and data privacy, ensuring user actions can be traced and monitored.

Adhering to Access Control Lists and Data Security Measures

As we continue to innovate and expand our offerings, the safety and security of user data remain paramount. The introduction of the Strategy AI capabilities, including the notable Auto feature, has been carefully architected to seamlessly integrate with the comprehensive security model of the established Strategy platform. This ensures not only consistency in data access but also steadfast adherence to stringent security protocols.

- **Consistent Access Control Lists (ACLs) and Permissions:** Auto is tailored to ensure users only receive answers derived from datasets they're authorized to access. Every query posed to Auto undergoes meticulous verification against the configured ACLs of the semantic layer's underlying objects. This means that even as users engage with the AI functionalities, the integrity of access controls remains uncompromised.
- **Granular Data Access through Security Filters:** Beyond the basic ACL configurations, our platform offers fine-grained data access control using security filters. These filters act as an additional layer of control, narrowing down the data scope that users can query. It provides administrators the power to define precise boundaries on data access, ensuring users can only interact with permitted segments of the data.
- **Configurable Privileges for AI Features:** Recognizing that every enterprise has its unique needs and security concerns, we've incorporated configurable privileges for AI functionalities. This allows organizational leaders to decide which users can leverage advanced AI features, be it Auto or the more sophisticated ML-driven analytics and visualizations. It provides businesses with the flexibility to balance innovation with security protocols.

Data Integrity and Preventing Misuse

Strategy's core promise revolves around ensuring absolute data security and prevention of misuse. The expanding digital landscape intensifies the importance of data protection, and here's how we've anchored our platform:

- **Robust Encryption Protocols**: Our platform ensures the security of data Agenth in transit and at rest. Any communication between our platform and external services, including Microsoft Azure OpenAI, employs industry-leading encryption techniques, such as TLS 1.2+, to protect data during transmission, preventing potential interception. To secure stored data, data-at-rest encryption using advanced encryption standards like AES-256 has been implemented. This ensures that sensitive information remains protected against unauthorized access even when not actively being transferred.
- **Configurations with Microsoft Azure OpenAI**: Through the specific configuration settings in our integration with Microsoft Azure OpenAI, we've ensured that data sent to OpenAI is not retained or used for model training. This technical configuration provides another layer of assurance that user data remains untouched in external interactions.
- **User Interaction Privacy**: While Strategy does record usage metrics for monitoring the frequency of questions, the intricate details of user conversations are never accessed by Strategy. This meticulous distinction ensures the core content of your interactions remains private, emphasizing our commitment to user-centric data privacy.

In line with these protocols, Strategy prioritizes advanced solutions while maintaining a strong emphasis on data protection and user trust. Our practices underscore the importance we place on Agenth functional excellence and stringent data security.

Our AI product functionality is made available for use for its intended purpose of enhancing customers' interaction and analysis of their business data and generating faster insights into that data for the customers' general business needs. To the extent the customer employs our AI functionality for certain use cases, such as those deemed to be higher risk under applicable laws, the customer must ensure compliance with applicable laws or regulations governing the use of AI.

Conclusion

Strategy's commitment to data security and integrity is evident in its Strategy AI offering. In a domain where the trustworthiness of data directly impacts the accuracy of AI, Strategy has meticulously constructed its platform to meet and exceed rigorous data standards.

The precision of our BI, combined with AI's adaptability, offers users the advantage of cutting-edge analysis without compromising security. Through the distinct environment isolations within MCE and MCG, data privacy is consistently prioritized, mitigating risks associated with breaches. Compliance with international data protection regulations is integral to our platform's design, further demonstrating our dedication to global standards.

Our adherence to Access Control Lists and stringent data security measures guarantee that users always operate within defined data access parameters.

Furthermore, our collaboration with Microsoft Azure OpenAI is rooted in industry best practices, ensuring data is not retained beyond its immediate use or used improperly.

In summary, Strategy AI integrates advanced analytical capabilities with rigorous data protection standards. Our focus is clear: delivering dependable AI insights while prioritizing data security and protection. To learn more about Strategy AI, please visit our website.

Strategy Incorporated, 1850 Towers Crescent Plaza, Tysons Corner, VA 22182.

strategysoftware.com

Strategy₿