rubrik

WHITE PAPER

Rubrik Annapurna

Secure Data for Amazon Generative AI

Table of Contents

EXECUTIVE SUMMARY

As organizations increasingly seek to harness the power of generative AI, the complexity and resources required to deploy foundation models (FMs) effectively present significant challenges. Amazon Bedrock is a transformative solution, offering a fully managed service that democratizes access to state-of-the-art GenAI models while maintaining enterprise-grade controls and scalability. The combination of Rubrik Annapurna with Amazon Bedrock aims to empower organizations to securely access all their enterprise data on-prem, in the cloud, and in SaaS applications through Rubrik with built-in sensitive data suppression, application aware pre-embeddings and secure access capabilities to accelerate enterprise-scale generative AI adoption.

Rubrik's Annapurna provides the foundation for enterprises to confidently move their AI-powered applications from pilot to production. To help secure data pipelines across cloud, on-prem, and SaaS environments, Rubrik Annapurna is designed to enable AI-driven applications that handle sensitive data, like customer support or finance, that will meet the needs of business and security leaders.

Our approach is designed to accelerate your AI innovation by:

- **Providing easy and secure access to your existing data**
  - No additional data lakes or data pipelines required, meaning lower costs and fewer risks, especially for legacy and homegrown applications.
  - External data can be leveraged directly with Amazon Bedrock Knowledge Bases, giving foundation models and agents contextual information from your company's private data sources for RAG to deliver more relevant, accurate, and customized responses

- **Smart pre-embedding**
  - Pre-embeddings on data in RSC pre-trained by application type (e.g., CRM, Billing, etc.)

- **Eliminating data freshness complexity**
  - Automatic data refreshes can be managed through SLA policies
  - New sensitive data is detected and suppressed automatically

- **Managing security permissions change**
  - Security permission for the original source data is maintained
  - Permissions are updated in near real-time

- **Avoiding sensitive data inclusion**
  - Rubrik filtering prevents sensitive or otherwise unwanted data from being exposed

AMAZON BEDROCK AND RUBRIK ANNAPURNA

Security cannot happen in isolation but needs to cover all aspects of the end-to-end solution to provide the most capable and trustworthy AI experience for organizations. Cloud security at AWS is the highest priority; the AWS shared responsibility model provides guidance on security responsibility demarcation: https://aws.amazon.com/compliance/shared-responsibility-model/.

CUSTOMER VALUE

By combining Rubrik Data Security and Amazon generative AI solutions, the integration will be designed to enable organizations to accelerate the secure adoption of generative AI solutions built on Amazon Web Services. Security and control for source data can be provided by Rubrik while the rapid secure deployment of generative end-to-end solutions can be more easily and confidently facilitated by Amazon.

This can potentially decrease the time to market for new and innovative products powered by Amazon generative AI capabilities and unlock previously untapped and difficult-to-include data sources, leading to better organizational alignment and applicability.

One potential example of an innovative product could be an intelligent customer support automation system that helps reduce the risk of exposing sensitive or proprietary information. Another example could be a smarter sales assistant for every sales team that pulls together disparate pieces of information across many corporate data sets. Another option could be a marketing assistant powered by Amazon Bedrock that leverages on-premises or cloud-based information.

HOW IT WILL BE DESIGNED

Rubrik Annapurna will be designed to proactively secure the data pipeline for Amazon Bedrock by leveraging existing data already managed by Rubrik. Rubrik customers already use the Rubrik Security Cloud platform to maintain a secure and immutable version of their enterprise data from a large number of sources across Microsoft 365, Jira, Salesforce, Amazon S3, and many others. This means this single source of data sidesteps the typical need to create another copy of your data dedicated to generative AI. The solution will be designed to provide a secure, contextual copy instead, enabling you to focus on building AI apps.

By combining Rubrik Annapurna with Amazon Bedrock, the solution will be designed to allow organizations to use their existing Rubrik data sources, filter out any sensitive or otherwise unwanted information using a Rubrik data filtering policy, and have Rubrik automate secure data retrievers, which can then be leveraged by Amazon Bedrock. This integration aims to help accelerate generative AI projects from pilot to production by enabling CISO-approved reliability and security.

Amazon Bedrock Guardrails will be designed to provide additional customizable safeguards, such as prompt and content filters on top of the native protections of FMs in Bedrock and enterprise-ready data provided by Rubrik. Amazon Bedrock Guardrails helps evaluate user inputs and FM responses based on use case-specific policies and provides an additional layer of safeguards regardless of the underlying FM. For example, it can help you to define a set of topics to avoid within the context of your application. Amazon Bedrock Guardrails helps detect and block user inputs and safeguard against prompt attacks (like prompt injection and jailbreak).

Independent of the filtered and secured source data, applications built using FMs can generate incorrect information due to hallucinations. For example, FMs can generate responses that deviate from the source information, conflate multiple pieces of information, or invent new information. Amazon Bedrock Guardrails supports contextual grounding checks to help detect and filter hallucinations if the responses are not grounded (e.g., factually inaccurate or new information) in the source information and irrelevant to the user's query or instruction. Contextual grounding checks can help detect hallucinations for RAG, summarization, and conversational applications, where source information can be used as a reference to validate the model response.

ENTERPRISE GENERATIVE AI ATTACK VECTORS

When utilizing a generative AI-based service, it is important to comprehend how your input data is stored, processed, shared, and utilized by the model provider or the environment operator. Providers of generative AI solutions are responsible for implementing appropriate safeguards to ensure privacy, compliance, and security in their applications and in the use and training of their models.

As organizations build their enterprise GenAI applications, guarding against attack vectors like those tracked by MITRE ATLAS and The OWASP Foundation is critical. This can include attack vectors like prompt injection, training data poisoning, model contamination, unsanctioned model use, sensitive information disclosure, and more.

The combination of Amazon Bedrock with Rubrik Annapurna will be designed to equip organizations with the necessary security and governance controls needed to unlock generative AI for enterprise use cases.

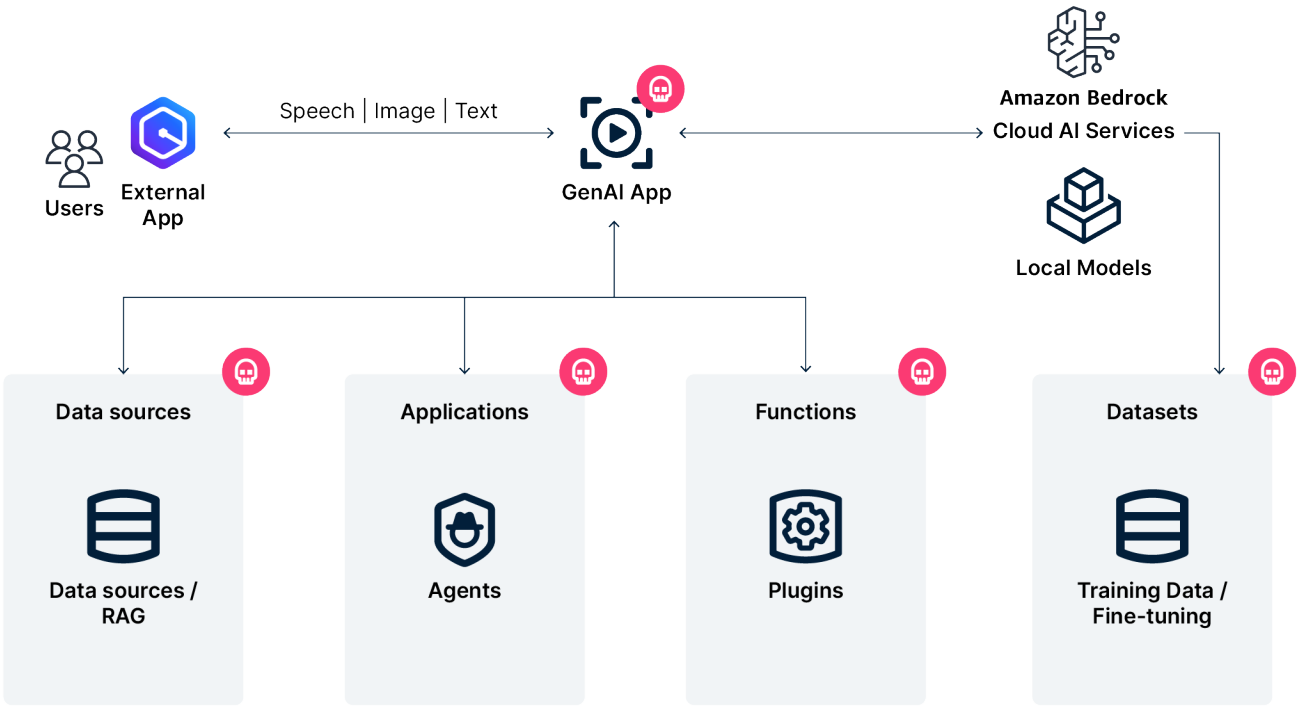

Addressing the secure development of building an enterprise generative AI application can employ multiple approaches. Typical approaches to align generative AI with enterprise applicability include model fine-tuning, retrieval-augmented generation (RAG), building agentic workflows, or supporting large context prompt windows. Each of these, however, relies on a secure data foundation. Whether it is enterprise pre-labeled data for model fine-tuning, external content stores for RAG, or even opening programmatic access to other AI agents or corporate systems, security needs to be top of mind.

With Rubrik Annapurna we can utilize existing datasets in Rubrik Secure Vault so no additional data lakes or data pipelines are required—this can significantly reduce the AI data attack vector.

RUBRIK ANNAPURNA COMPONENTS

An enterprise generative AI architecture consists of multiple critical components, starting from a secure and scalable infrastructure running on top of Amazon Web Services all the way up to generative AI applications driven by Amazon Q. Rubrik Annapurna will be designed to provide programmatic integration with applicable Amazon AI components, so existing enterprise data can be leveraged securely and at scale. It aims to provide flexibility and model choice while at the same time implementing a zero-trust approach to data security.

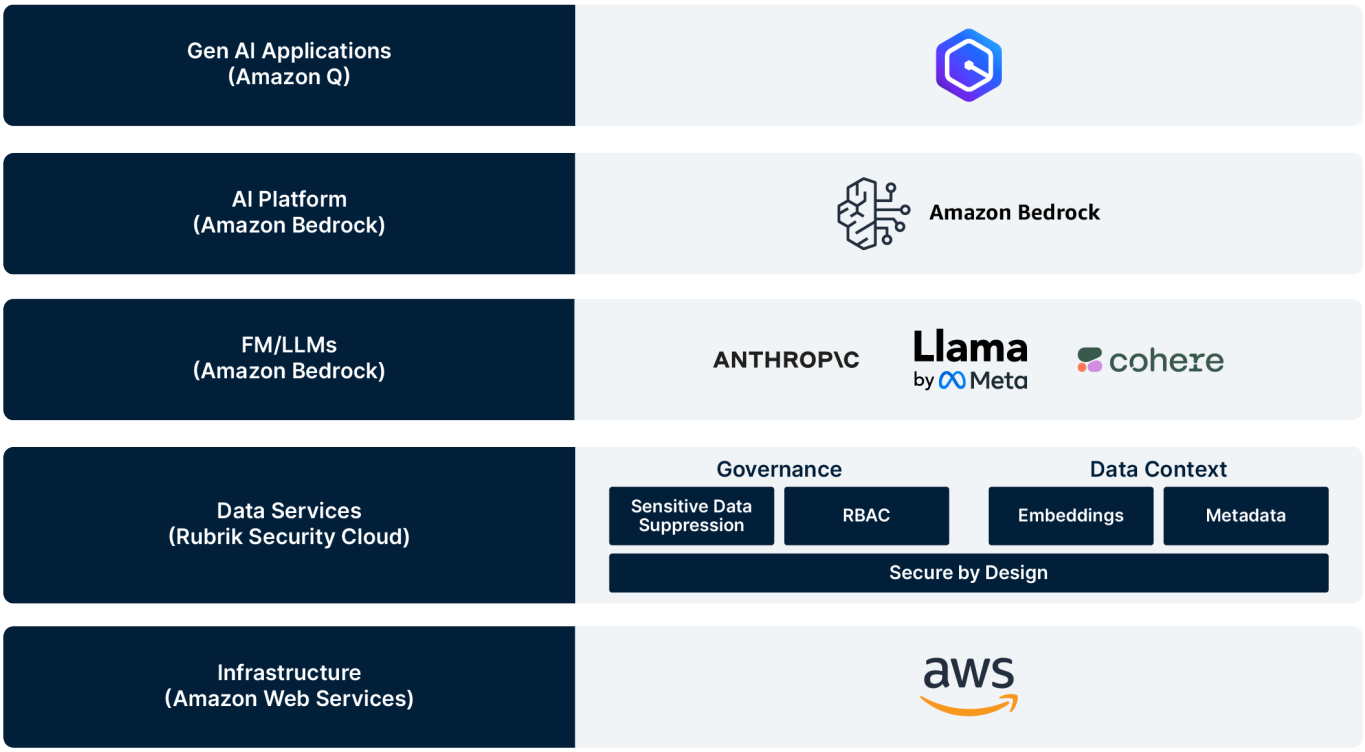

An example generative AI assistant application built using Amazon Bedrock might look like the diagram below: a new set of data for RAG usage is uploaded to Amazon S3 over the public internet to be embedded and written into the Amazon OpenSearch Vector Database by Amazon Bedrock.

When the user using the AI assistant submits a query, Amazon Bedrock Knowledge Bases leverages an embedding model to convert the user query to a vector and finds chunks that are semantically similar to the user query.

The user prompt is then augmented with the chunks that are retrieved from the knowledge base. The prompt alongside the additional context is then sent to an LLM for response generation, resulting in the user seeing an answer and the citation on the AI assistant user interface.

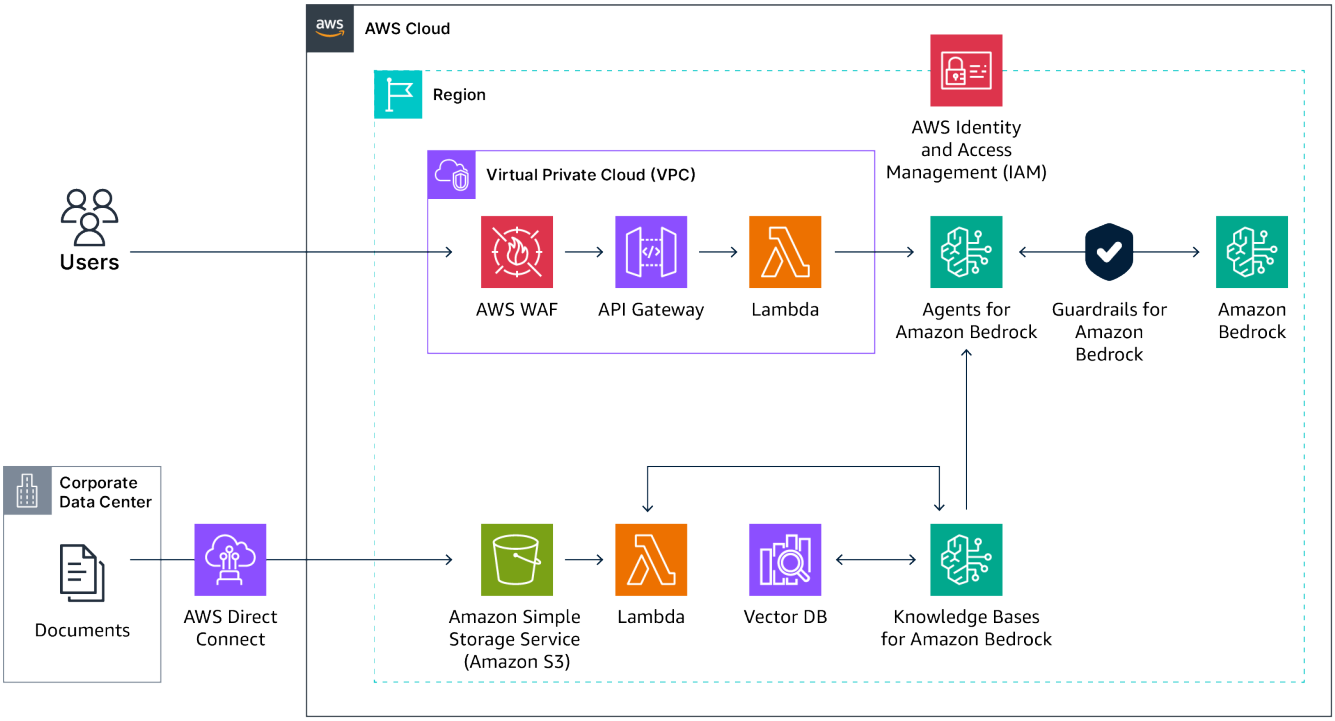

With Rubrik Annapurna the dataset already exists and is secured by Rubrik Secure Vault, as such data security comes by design and can easily be integrated as part of the overall Amazon Bedrock architecture.

Rubrik Security Cloud shortens the time to build GenAI Apps. There are several time-consuming steps to get data sources that Rubrik can help simplify.

Annapurna Key Value Propositions

Not just a secure data repository, but a GenAI accelerator

API First.
Fast, configurable, and current access to all your data from a single API

Secure and responsible AI from the start.
Governed by permissions and sensitive data suppression

Unlock all your data.
Use AI with all your data including legacy or homegrown apps without writing new APIs or custom pipelines

Application-aware pre-embedding.
Smart pre-embeddings that work for 3rd party or custom app data schemas

Avoid shadow data stores.
Use the data you already have without additional infrastructure cost, management cost, and risk

1. **API First:** Gain fast, configurable, and up-to-date access to all your data through a single, unified API. Rubrik's approach simplifies integration across data sources, removing the need for multiple connection points.
2. **Secure and Responsible AI from the Start:** Rubrik helps ensure that data is governed by robust permission settings and sensitive data suppression policies, providing a secure foundation for AI applications.
3. **Unlock All Your Data:** Use AI across all your data, including legacy and homegrown applications, without needing new APIs or custom pipelines. Rubrik simplifies data access for a more efficient AI workflow.
4. **Application-aware Pre-embedding:** Leverage smart pre-embedding technology that works seamlessly with third-party or custom app data schemas, enabling efficient AI-ready data processing.
5. **Avoid Shadow Datastores:** Use the data you already have without additional infrastructure costs or risks. Rubrik's approach eliminates the need for shadow datastores, reducing management overhead and ensuring data security compliance.

Annapurna is designed so that sensitive data can be filtered out at the source using Rubrik policies before being made available for embedding. Role-based access control is maintained in accordance with the original source data to maintain integrity and avoid overexposure through generative AI applications. The data is indexed by Rubrik and metadata is created.

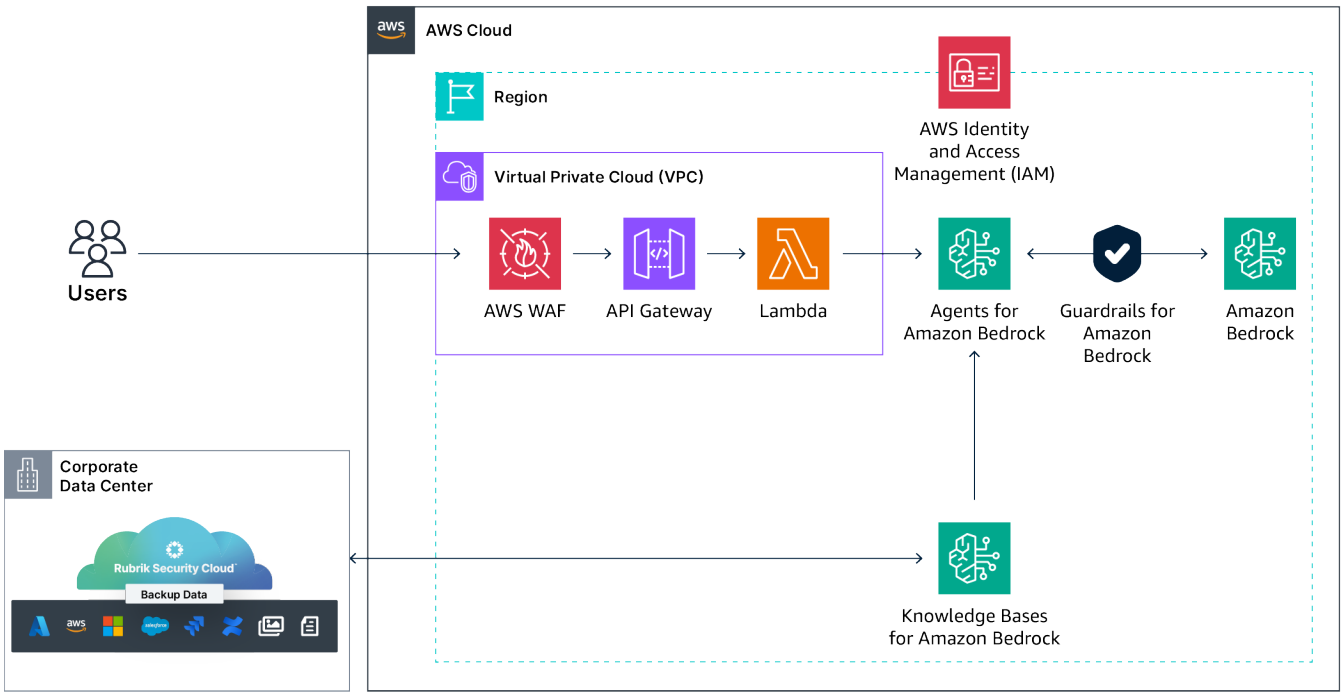

Programmatic access is provided to Amazon Bedrock to seamlessly integrate the Rubrik enterprise data sources in the AI architecture, which then can be leveraged for retrieval-augmented generation.

Once integrated, the solution will be designed to allow Rubrik datasets to be leveraged like any other data source within the Amazon Bedrock framework with the differentiation that security is native to the dataset itself. More importantly, the source permissions and access controls at the document level stay intact due to Rubrik's data governance framework.

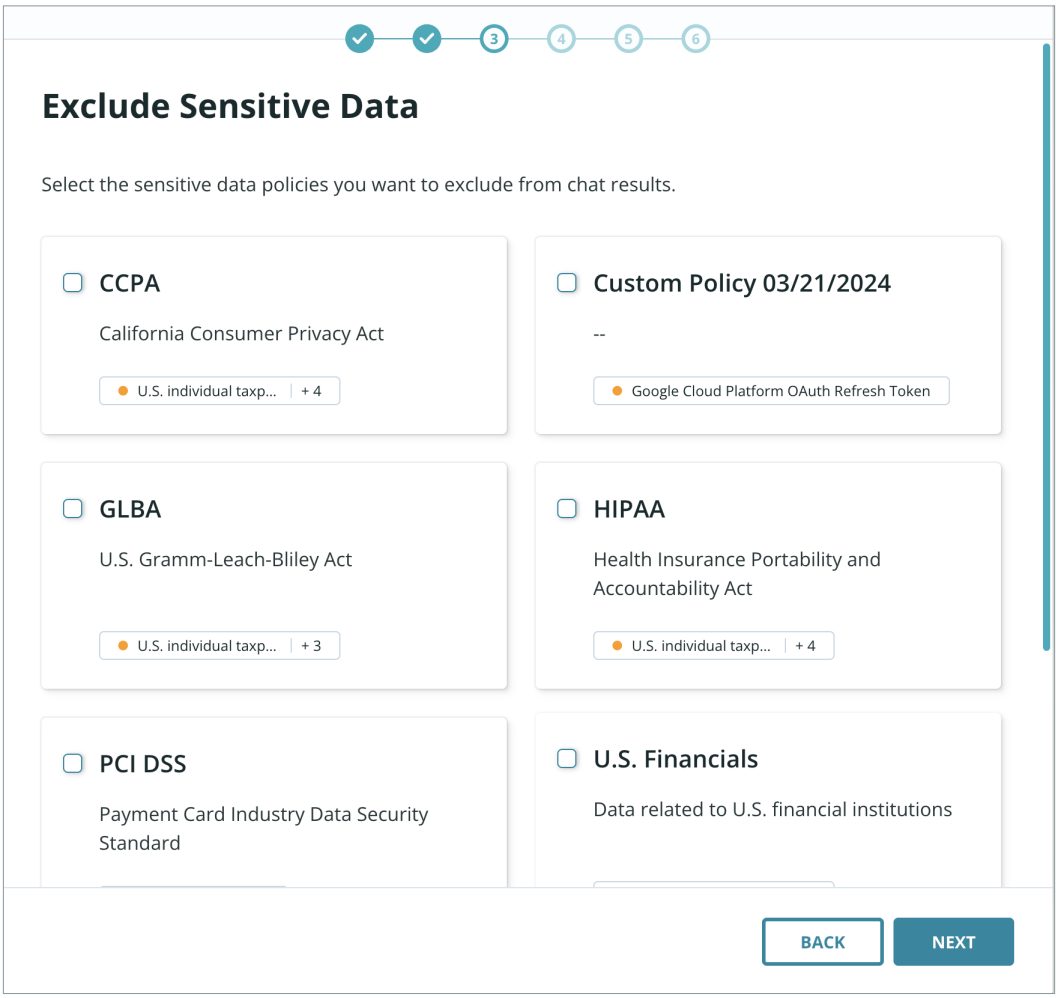

Example policies have been designed so that certain sensitive and other unwanted information can be excluded up front, and help prevent it from inadvertently becoming part of the augmented dataset leveraged by the enterprise AI application.

OPTIMIZED AMAZON BEDROCK AND RUBRIK GENERATIVE AI STACK

For the Retrieval Augmented Generation approach, Rubrik can leverage backup data present in the Rubrik Security Cloud, embed this data in a vector database, and have it serve as the augmentation to the user prompt. Additionally, the solution will be designed to enable Agents for Amazon Bedrock to be able to programmatically leverage the pre-embedded and secure Rubrik dataset in their agentic workflows. RAG consists of three steps: ingestion, retrieval, and synthesis. Rubrik Annapurna is designed to help make ingestion and retrieval faster, easier, and more secure, while Amazon Bedrock aims to ensure that the synthesis results in accurate, contextual responses.

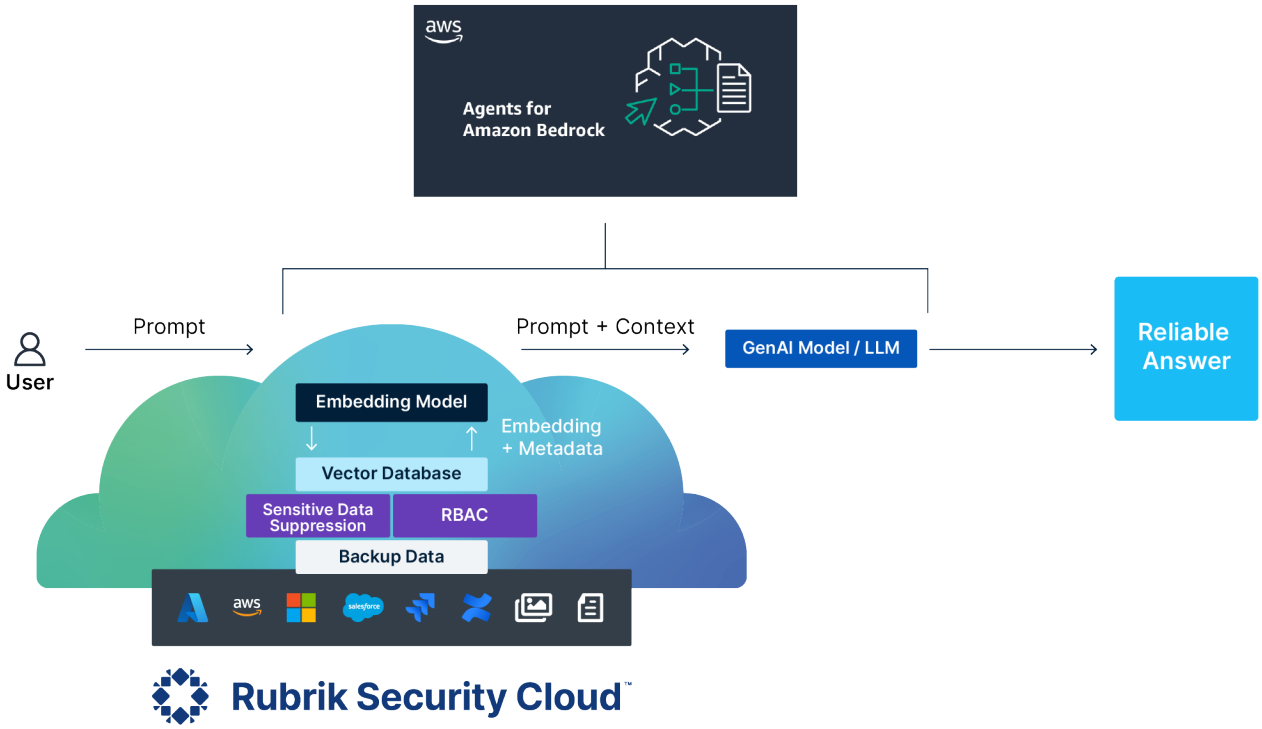

EXAMPLE USE CASES

Description	Value of Use Case	Unique Need Addressed by Rubrik & AWS
Intelligent Customer Support Automation	Streamlines customer support, enabling faster, compliant, and more accurate responses.	Support teams can create smart AI agents capable of actively reducing the time it takes to research and resolve support cases with Rubrik and AWS together. Rubrik secures sensitive data, helping organizations ensure regulatory compliance and secure data access. AWS scales the AI infrastructure to handle agentic reasoning.
A Smarter Sales Assistant for Every Sales Team	Sales teams can instantly get a 360° summary and qualitative review of customer deals from not only what's in the CRM, but also the "inside baseball" from customer emails, phone conversations, and the relevant internal communications across the team. Example: Extracts insights on customer interactions from your CRM, your emails, your Teams chat, and even call recordings sitting in your S3 buckets.	Rubrik provides a secure, policy-compliant environment for extracting insights without data replication, helping protect sensitive data. AWS offers scalable resources to manage data volumes across SaaS platforms.
A Marketing Assistant Who Knows More	Summarizing existing marketing docs from OneDrive allows users to create derivative content. Example: Create a better whitepaper or presentation by reading through presentations and documents in OneDrive, but also summarize the latest product documentation in your knowledge management system, and listen to customer calls in S3.	Marketing teams can securely access relevant data sources across their data estate with Rubrik and leverage state-of-the-art LLMs on Amazon Bedrock to scale compelling content that accurately represents product benefits and differentiators.

BENEFITS OF RUBRIK ANNAPURNA FOR AMAZON BEDROCK

Rubrik Annapurna for Amazon Bedrock is designed to drive secure and responsible AI from the start, governed by permissions and sensitive data suppression. Avoid shadow datastores by leveraging the data you already have without additional infrastructure costs, management costs, and risk. Unlock custom apps by using AI with legacy or homegrown apps without writing new APIs or custom pipelines. Rubrik Annapurna is built with an API-first approach, providing fast, configurable, and current access to all your data from a single API.

Let's take a look at how Rubrik Annapurna will be designed to accelerates your AI innovation.

- **API-first Access to All Your Data:** Rubrik Annapurna will be designed to provide a unified API for secure and seamless access to data across cloud, on-prem, and SaaS environments. This aims to simplify the complexity of data updates, permission adjustments, and sensitive data handling. By integrating with Amazon Bedrock, Rubrik Annapurna can enable you to select and apply the ideal AI model for each use case.

- **Built-in Permissions and Sensitive Data Management:** Control access effortlessly with permissions that reflect your existing policies and ensure sensitive information is filtered out before becoming part of your AI corpus. Rubrik Annapurna will be designed to offer a governance-first approach to AI data readiness.

- **Effortless Data Mobilization Across All Applications:** Rubrik Annapurna will be designed to support access to enterprise data from SaaS, legacy, and custom applications without needing complex custom APIs, ETL processes, or manual data pipelines. Connect once and enable AI use across CRMs, billing systems, and other key applications, regardless of their origin.

- **Application-aware Pre-embedding:** Rubrik Annapurna's application-aware pre-embedding capabilities will be designed to automatically recognize common data schemas. Configurable pre-embedding kits simplify data preparation for AI, to reduce the need to create custom vector databases for efficient retrieval.

- **Eliminate Shadow Datastores:** Consolidate your data for AI applications into one flexible, manageable store with Rubrik Annapurna, to avoid the added costs, risks, and overhead of shadow datastores. This unified approach can help minimize infrastructure sprawl and maximizes security and compliance.

NEXT STEPS

While tools like Amazon Bedrock democratize access to cutting-edge AI, it's important to ensure the secure and responsible use of sensitive enterprise data.

Rubrik Annapurna, in conjunction with Amazon Bedrock, aims to offer a comprehensive solution to this challenge. Together, this combination is designed to empower you to harness the transformative potential of generative AI while upholding the highest standards of data security and governance.

To learn more about Rubrik Annapurna, how it works, and how you can implement it in your environment, check out this demo, or click here to get in touch with a Rubrik specialist.

SAFE HARBOR STATEMENT

Any unreleased services or features referenced in this whitepaper are not currently available and may not be made generally available on time or at all, as may be determined in our sole discretion. Any such referenced services or features do not represent promises to deliver, commitments, or obligations of Rubrik, Inc. and may not be incorporated into any contract. Customers should make their purchase decisions based upon services and features that are currently generally available.

rubrik

Global HQ
3495 Deer Creek Road
Palo Alto, CA 94304
United States

1-844-4RUBRIK
inquiries@rubrik.com
www.rubrik.com

wp-rubrik-annapurna / 20241125