AVM **Common FAQs:**

Where does your security software platform fit into the technology stack?

Our security software platform serves as the front-end firewall for controlling prompt content, data content, and data sources used in GenAI models, including closed and open-source models such as GPT, Cohere, Llamma, Titan, Claude, etc...

What is the primary value your platform offers?

Our platform provides compliance, governance, and guardrails for customers wanting deploy secure GenAI and ensures the security and appropriateness of data and prompts, as specified by users or companies. It protects against inappropriate content and data flow for GenAI use cases and applications.

How do users leverage your technology?

Users can leverage our technology to enhance the security and ethical use of GenAI models. They can define rules and policies for content filtering, data redaction, and compliance to align with their specific compliance needs and regulatory requirements.

Can your platform work with various types of language models?

Yes, our platform is designed to integrate seamlessly with a wide range of language models, including RAG models, fine-tuned models, closed and open-source models, ensuring versatility and compatibility. Customers ‘bring their own model’ and can plug it in to our secure control plane to ensure consistent and reliable security standards across all models.

What actions can your platform take to protect against inappropriate content?

Our platform can block, warn, or redact inappropriate content in real-time, preventing the generation or use of such content as per the user's or company's defined policies.

Is your platform suitable for both customers and partners?

Absolutely, our platform is designed to serve both customers and partners, offering a comprehensive solution for securing and governing GenAI models and data.

Where does AVM fit if I am using GPT-4, Llama-2 or Claude-2?

AVM seamlessly integrates with advanced language models such as GPT-4, Llama-2 or Claude-2, serving as the front-end security layer adding an extra layer of compliance, governance, and content control to ensure responsible and secure use of these models. AVM enhances the capabilities of GPT-4, Llama-2 or Claude-2 by providing compliance, governance, and guardrails. It allows users to define and enforce content policies, protecting against inappropriate content and data misuse while leveraging the power of these advanced models.

Can AVM adapt to the unique needs of different language models?

Yes, AVM is designed to be adaptable and compatible with various language models and can tailor its security and content control measures to align with the specific requirements of each model and how it is used.

What benefits does AVM bring to users of LLMs?

AVM offers users of GenAI the benefits of enhanced security, compliance with regulations, and the ability to create responsible and ethical AI applications. It ensures that the content generated or processed by these models aligns with company or user-defined standards.

How does the chatbot interact with your security platform?

The chatbot serves as the input interface for generating prompts. All questions or prompts from the chatbot pass through our security platform, which ensures the content's compliance and appropriateness before being used in downstream processes.

Can your security platform work with both popular language models and custom fine-tuned models?

Yes, our security platform is model-agnostic and can seamlessly integrate with both popular large language models and company-specific fine-tuned models. It provides consistent content control and data security across all model types.

How does the dataset integration work with your platform?

Our platform integrates with datasets, which may include popular LLMs custom fine-tuned models, RAG or company generated data to ensure that the data used for generating responses aligns with security and compliance policies. AVM verifies and filters the data before utilization.

What are the chatbot options in combination with the security platform and dataset integration?

Using our chatbot in conjunction with the security platform and dataset integration allows for the generation of AI responses that are not only accurate but also meet stringent security and ethical standards. It ensures responsible AI interactions. Similarly, you can bring your own chatbot and still leverage AVM's secure control plane to provide the guardrails and compliance that align with your company's security policies.

Can users customize the security and compliance policies for the chatbot and dataset integration?

Yes, users have the flexibility to define and customize security, compliance, and content control policies to align with their specific requirements. This customization ensures that AI interactions adhere to their unique standards.

How can I integrate your chatbot and security platform into my existing AI infrastructure?

Integrating our chatbot and security platform is straightforward. Please contact our technical support team, and they will guide you through the integration process, ensuring a seamless addition to your AI infrastructure.

How does AVM integrate with existing SIEM systems?

AVM seamlessly integrates with your organization's existing SIEM system, providing a unified security approach. It feeds relevant security data and incidents into the SIEM, enhancing its capabilities to monitor and respond to AI-related security events.

Does AVM offer any specific features for SIEM integration?

Yes, AVM includes dedicated features for SIEM integration. It can export logs, alerts, and security events to your SIEM system, ensuring that AI-related security information is centralized and actionable within your existing security infrastructure.

Can AVM help in real-time threat detection and incident response within SIEM systems?

Absolutely. AVM's real-time monitoring and content control capabilities contribute to early threat detection. It generates alerts and triggers incident responses from our platform to your SIEM system, allowing swift actions to mitigate AI-related security risks.

How does AVM work alongside identity management systems?

AVM complements identity management systems by adding an additional layer of security and compliance to user access and interactions with AI models. It ensures that only authorized users can utilize AI resources while adhering to content control policies.

Can AVM enforce access control based on user identities from identity management systems?

Yes, AVM can enforce access control by integrating with your identity management systems. It aligns with your existing user authentication and authorization mechanisms to ensure that AI resources are accessible only to authorized users.

Is it possible to customize AVM's integration with SIEM and identity management systems?

Yes, AVM offers customization options. You can tailor the integration with SIEM and identity management systems to align with your organization's specific security policies and requirements.

How does AVM complement existing DLP solutions?

AVM enhances existing DLP solutions by adding an AI-specific layer of content control and compliance. It ensures that AI-generated content aligns with DLP policies, preventing the generation or dissemination of sensitive or inappropriate data.

Can AVM integrate with my organization's current DLP infrastructure?

Yes, AVM is designed to seamlessly integrate with your existing DLP solutions. It acts as a content filter for AI-generated data, ensuring that data leaving your organization complies with DLP policies.

Does AVM assist in identifying and mitigating data leakage risks from AI models?

Absolutely. AVM actively monitors AI-generated content for potential data leakage risks. It generates alerts and takes actions to prevent sensitive data from being inadvertently exposed or shared.

How does AVM handle content that violates DLP policies?

AVM can be configured to block, warn, or redact content that violates DLP policies. It ensures that AI-generated content adheres to your organization's data security and compliance requirements.

Can AVM help in real-time detection of sensitive data within AI-generated content?

Yes, AVM offers real-time detection capabilities. It scans prompt and AI-generated content for sensitive data patterns and immediately takes action to prevent any violations of DLP policies.

Is it possible to customize AVM's integration with DLP solutions?

Yes, AVM provides customization options. You can tailor the integration with your DLP solutions to align with your specific data protection and compliance needs.

How does AVM's control plane manage access to different models based on use cases?

AVM's control plane uses role-based access control and the concept of teams to assign specific user roles or permissions, ensuring that users only have access to the models relevant to their designated use cases. This fine-grained access control optimizes model usage and cost-effectiveness.

Can users define their use cases and model preferences within AVM?

Yes, users can define their use cases and preferences within AVM. The platform allows users to specify their requirements, and AVM's control plane ensures that they are connected to the most suitable models accordingly.

How does AVM's security platform enforce content control and compliance for different models?

AVM's security platform applies content control and compliance policies specific to each model. It tailors its security measures based on the user's selected model, ensuring that all generated content adheres to the designated standards.

Does AVM help optimize costs by managing model usage efficiently?

Yes, AVM's control plane optimizes costs by ensuring that users are directed to the most cost-effective models for their use cases. It helps prevent unnecessary usage of resource-intensive models, saving on computational expenses.

Can AVM dynamically adjust model access based on changing user needs or usage patterns?

Absolutely. AVM's control plane offers dynamic allocation of models based on changing user needs and usage patterns. It continuously optimizes model access to align with evolving requirements.

How can organizations implement AVM's multi-model strategy for their users?

To implement AVM's multi-model strategy, please contact our technical support team. They will work closely with your organization to define user roles, access policies, and model allocations to match your specific use cases and cost-saving objectives.

Scenarios where AVM may not be the best fit for customers or specific use cases:

1. **Non-Textual Data:** AVM is primarily designed for text-based content control and compliance in AI models. If a customer's AI use case relies heavily on non-textual data (e.g., images, audio, video), a future version of AVM may fit your needs.
2. **Low AI Model Usage:** For customers with minimal AI model usage or infrequent AI interactions, the robust features of AVM may be overkill. AVM is most beneficial for organizations with substantial GenAI usage and security needs.
3. **Budget Constraints:** If customers have tight budget constraints and do not anticipate significant AI usage, investing in AVM's advanced features may not be cost-effective. They may opt for simpler, cost-efficient solutions.
4. **Single, Home Built AI Model Use:** Customers who rely exclusively on a single AI model or an internal application with no need for multi-model orchestration or fine-grained content control may find AVM's features unnecessary.
5. **Existing Robust Security Infrastructure:** Organizations with an existing, highly robust security infrastructure that adequately covers AI content control and compliance may not require AVM's additional layer of security.
6. **Non-Regulated Industries or Companies:** In industries or organizations with no strict regulatory requirements or security concerns regarding AI content control, AVM's compliance features may not be a priority.

Architecture Overview:

1. **Chatbot:** The user-facing interface where users input prompts or questions. It acts as the entry point to the system.
2. **AVM Security Control Plane:** Positioned between the chatbot and the language models, AVM's control plane handles user authentication, content control policies, and routing requests to the appropriate models. It enforces security and compliance measures.
3. **RAG Model / Fine-Tuned Model:** This represents a core, custom AI model used for generating responses. It could be a RAG (Retrieval-Augmented Generation) model or a custom fine-tuned model specific to your company's needs.
4. **Large Language Models (LLMs):** The many various large language models available (LLAMA-2, Cohere, GPT-3.5, Claude-2 and others). These models are integrated into AVM's security control plane and used based on user preferences and use cases.

Flow of Interaction:

1. The user interacts with the chatbot, providing a prompt or question.
2. The chatbot sends the user's input to the AVM Security Control Plane.
3. AVM's Control Plane evaluates the user's request, checks for authentication, and enforces content control policies.
4. Based on the user's role, use case, and content control policies, AVM's Control Plane routes the request to the most suitable model (RAG Model / Fine-Tuned Model or one/many of the LLMs).
5. The selected model generates a response based on the input and sends it back to the user through AVM's Control Plane and back to AVM's UI or the Company's UI, if applicable.

The architecture ensures that users receive responses that are compliant with security and content control policies while leveraging the most appropriate model for their specific use case.

Start (User Input) | Chatbot Input | AVM Security Control Plane | Authentication & Content Control Policies | Determine User Role and Use Case | Select Appropriate Model (RAG/Fine-Tuned/LLMs) | Model Generates Response | Response Sent to User | End

Each step represents a stage in the interaction between the user, the chatbot, the AVM Security Control Plane, and the various language models.