[DISTRIBUTIONAL]

Distributional’s Approach To Security And Privacy

Intro

At the heart of Distributional's design is a simple yet crucial belief: testing AI application data is essential. This ensures not only the safety and reliability of AI systems but also helps to deliver a seamless end-user experience. Regardless of where teams are on their AI journey, this conviction resonates with every organization we've engaged with.

For organizations deploying AI apps, the most effective way to test these apps is by leveraging their production logs, which capture the real usage interactions. This approach provides unparalleled insights, and grounds testing in real world usage.

Given the highly sensitive nature of production logs, companies are understandably cautious about their security. For instance, if you've deployed an HR chatbot, you wouldn't want engineers—let alone a third-party company—looking over the shoulders of employees as they ask deeply personal questions. And for regulated industries, addressing these privacy concerns becomes even more critical.

This is why many AI systems are designed using hosted models and deployed on-premises or within VPC infrastructure. And production logs are typically stored in highly secure environments with strict access controls. This careful approach ensures that sensitive data remains protected and doesn't fall into the wrong hands.

However, this limited access often prevents companies from leveraging their production logs to test and understand their AI apps. In turn, this means they miss out on a significant opportunity to enhance the quality of their AI systems and provide a more reliable and safe user experience.

Designing Distributional To Be Secure And Compliant

Distributional has been deeply mindful of these security concerns since the beginning. Our goal has always been to create a testing platform that is secure and compliant—one that protects user privacy and adheres to regulatory requirements without compromise—while still enabling teams to gain insights from their data.

To achieve this, we've architected our product to be private and secure by design. Let's dig into various aspects of our product that make it possible for us to work with the most security conscious oriented teams in the world.

DEPLOYING IN A SECURE ENVIRONMENT

Distributional's platform can be seamlessly deployed within a customer's environment—whether on-premises or within a VPC—ensuring they maintain full control over data access and security. For teams that prefer even greater autonomy, Distributional also offers a lightweight version of the platform that can be run locally.

(Note: For customers who prefer the convenience of a fully-managed service, the SaaS version is just a few clicks away. Simply sign up to get access.)

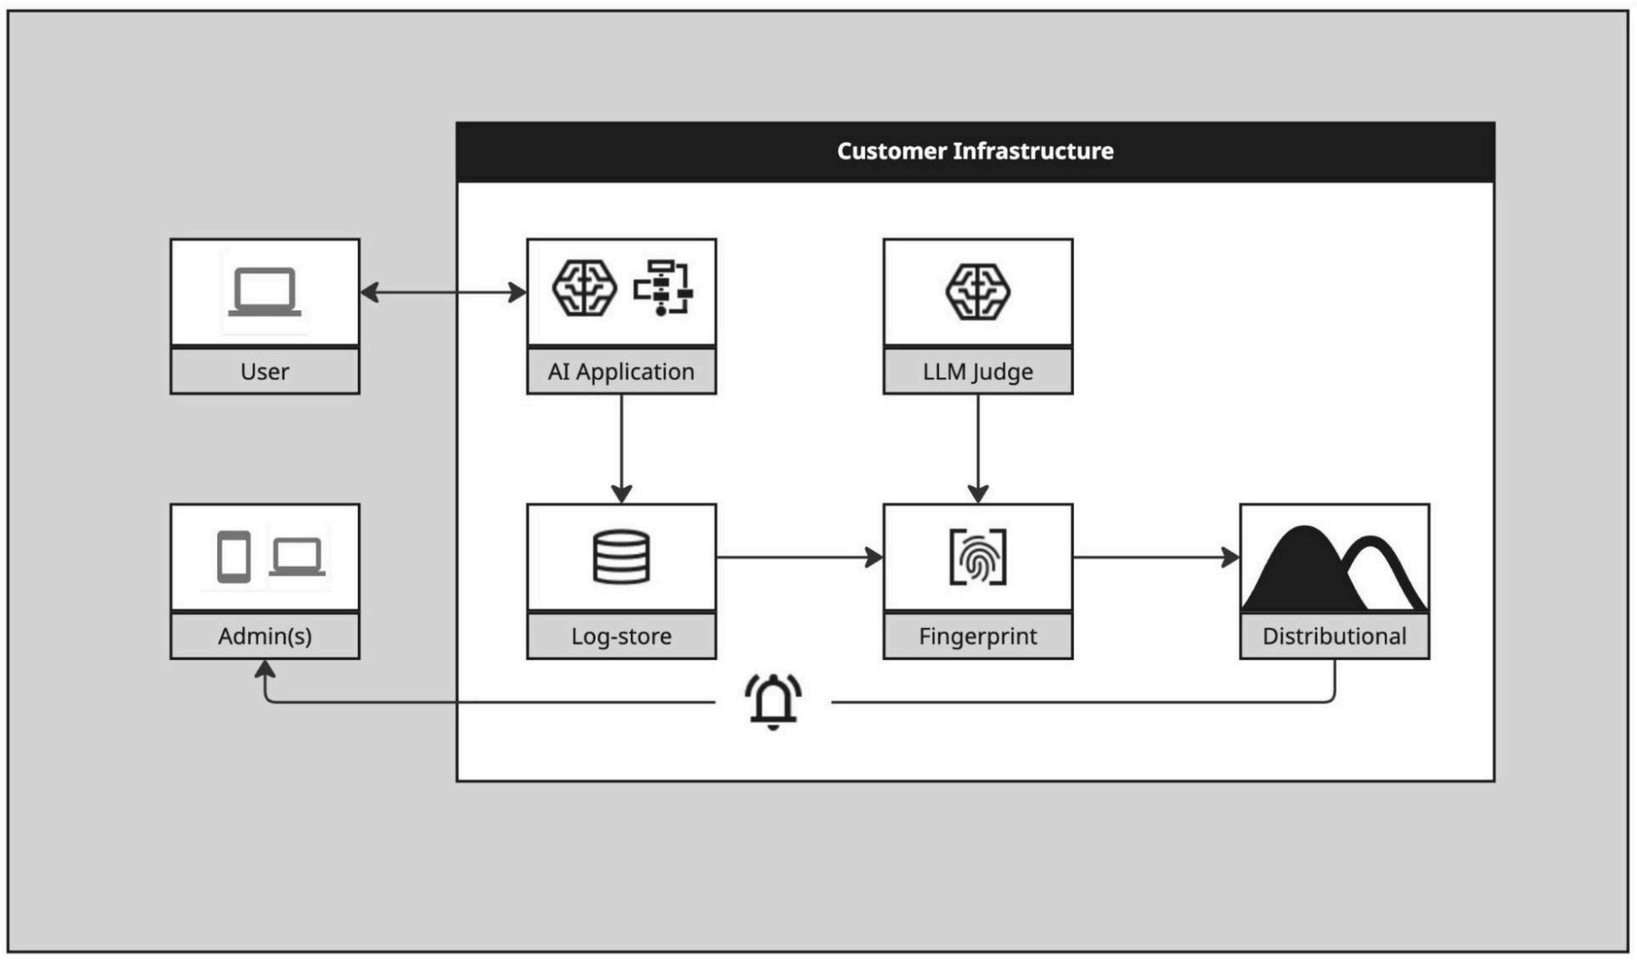

Distributional is designed to seamlessly integrate within customer infrastructure, securely coexisting with data.

NO CALL-HOME FUNCTIONALITY

To further ensure that sensitive information remains completely under a customer's control, Distributional is designed with a strict no call-home policy. Once installed in a customer's environment, it operates entirely within their secure infrastructure, with no data being sent or received outside of their system.

The only scenario where even the most privacy-conscious users permit Distributional to communicate outside their network is when teams need to have automated alerting. This allows teams to get notified when AI apps begin exhibiting unexpected behavior, so teams can take prompt corrective action. These alerts are fully configurable within the platform and designed to integrate with existing, trusted notification systems such as PagerDuty and Slack.

SECURE ACCESS CONTROL

Distributional is deployed with a namespace architecture, which provides secure and flexible access management. Namespaces serve as isolated partitions within the platform, allowing administrators to control who can access the specific Distributional projects aligned with individual AI apps and any relevant data.

For example, the engineering team for a customer support chatbot could have access to the respective project, while an HR-bot project could be limited to just one or two trusted individuals. These namespace controls ensure that sensitive data remains protected and that teams can only access the information relevant to their work.

INTEGRATION WITH EXISTING MODELS

Distributional uses LLM-as-Judge to evaluate whether AI outputs meet the intended task. Teams can connect their own securely hosted model or use Distributional's default.

Unlike many other solutions, this approach is designed so it doesn't lock organizations into a specific provider. Distributional integrates seamlessly with existing infrastructure, and orchestration runs entirely within an organization's environment.

This approach ensures that even judgment and evaluation maintain the same high standard of privacy and security as core applications. Data never leaves the secure environment, and all models operate under existing security protocols.

INTEGRATION WITH EXISTING DATA

Since most customers already maintain their AI production logs in a centralized, secure location, Distributional's platform is designed to integrate directly with these data stores and automatically fetch the AI production logs.

With this setup, only the data explicitly approved will be accessed or processed by Distributional. Customers retain full control over which logs leave their storage environment, ensuring that sensitive information remains secure and compliant with their data protection policies. Additionally, orchestration and scheduling are entirely under a customer's control—they decide when and how data is processed, so they retain full autonomy over workflows and minimize any risk of unintended data exposure.

SECURE, STANDARD APIS

Distributional offers a robust set of standardized APIs designed for secure, reliable, and scalable integration. All data exchanged with the platform is encrypted using HTTPS, and the APIs are optimized for high-throughput, low-latency operations suitable for production environments.

Every API request is authenticated using personal access tokens, which are user-specific and enforce fine-grained access control. Only authorized individuals or systems can interact with the platform, and administrators retain full control over token management—ensuring that API access remains secure and adaptable across both self-hosted and cloud-based deployments.

For uploading or downloading data via the SDK, Distributional uses pre-signed URLs with limited time and scope. These URLs are issued only after access permissions are validated by the API, ensuring secure and temporary access to specific data.

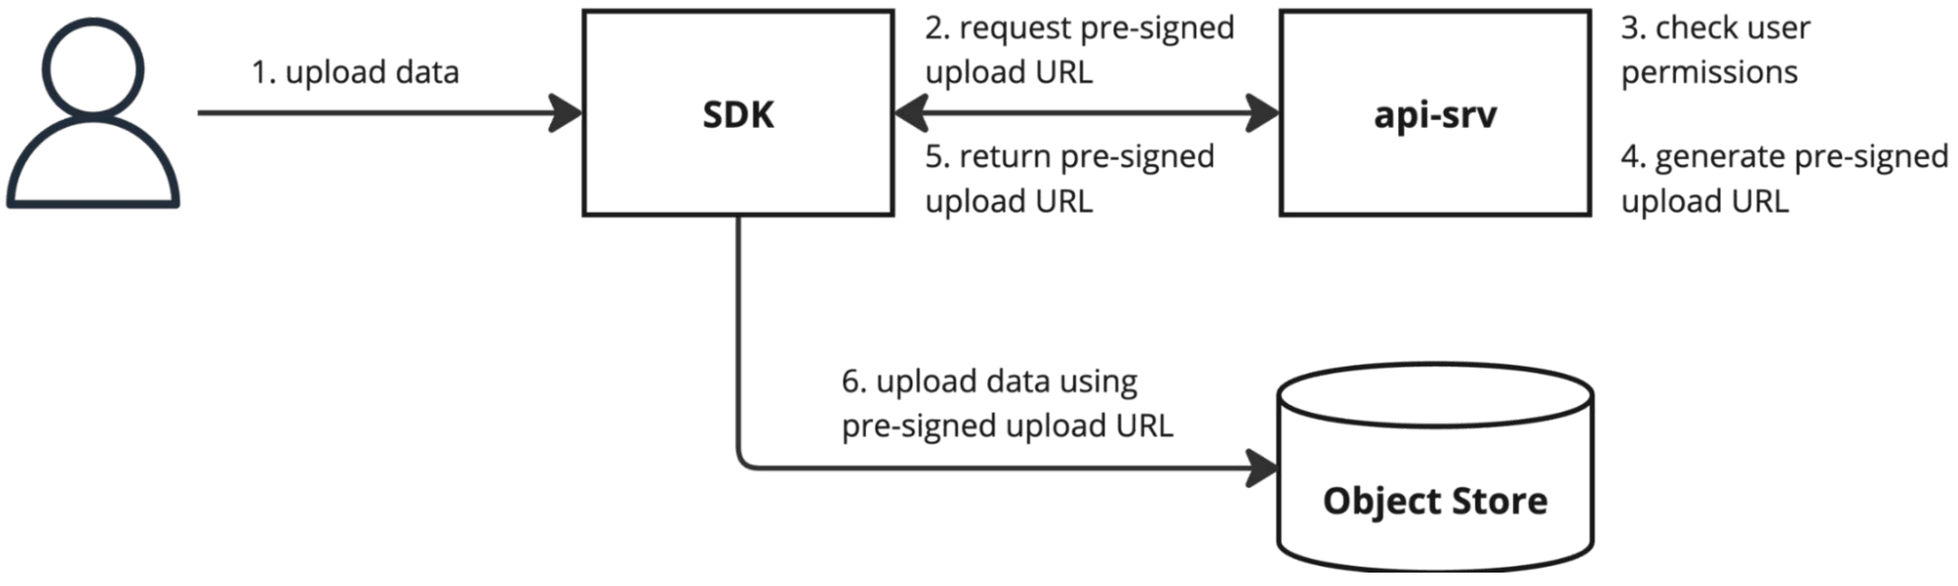

SECURE AND INSIGHTFUL TESTING THROUGH METRICS

Distributional's platform takes a unique approach to testing by extracting measurable metrics derived from text data using advanced analysis, rather than focusing on just the raw text inputs and outputs. This method offers a critical advantage—ensuring these metrics cannot be used to reconstruct the original text. This allows any sensitive information in the raw text to remain protected. Depending on privacy requirements, customers have the flexibility to only rely on these derived metrics for testing, rather than uploading text data. This ensures teams are able to still garner actionable insights, without the risk of exposing sensitive user information.

[DISTRIBUTIONAL]

About Distributional

Distributional's platform is designed to give customers complete control and optionality over their data and integrations. The platform can adapt to specific security, privacy, and regulatory needs, while still providing the quality and insights necessary to test AI applications.

If you'd like to learn more, reach out to the Distributional team. We're here to help!

Learn more at distributional.com

Follow us on:

- LinkedIn
- YouTube