Leaptree
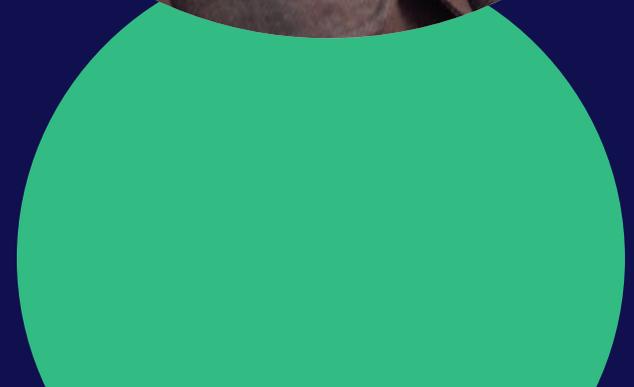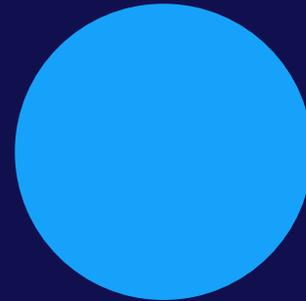
AI You Can Trust: Built on Your Data, Secured by Design

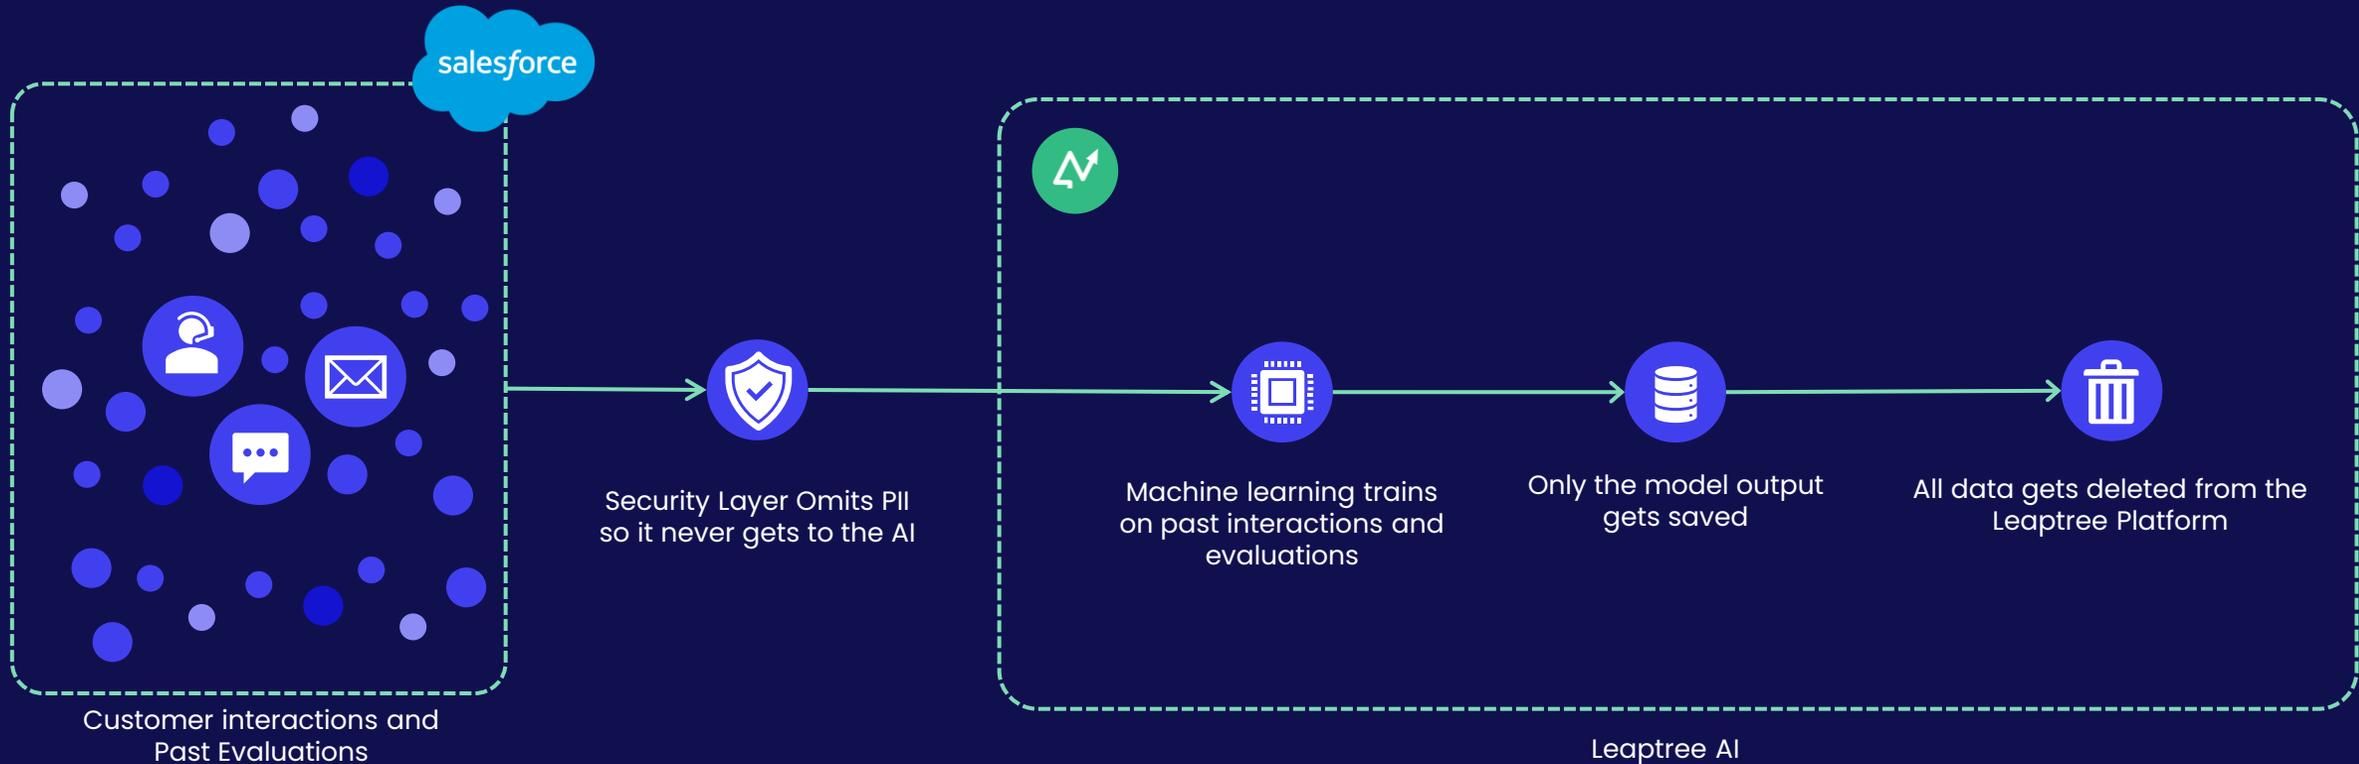Overview of Leaptree AI application and security standards.

Secure AI For Businesses

Leaptree AI has been designed from the ground up with security as its foundation. We ensure that every company's AI model is trained exclusively on its own data. Data is never shared & never mixed, so results are both more secure and more accurate. To protect sensitive information, all customer interactions are stripped of personally identifiable information (PII) before processing, and any data stored outside of Salesforce is retained only for the short period needed to complete training or prediction, after which it is permanently deleted. Unlike generic large language models that provide surface level insights, Leaptree AI delivers dedicated models tailored to each company and each business process, using only their own data to create precise, relevant, and secure outcomes.
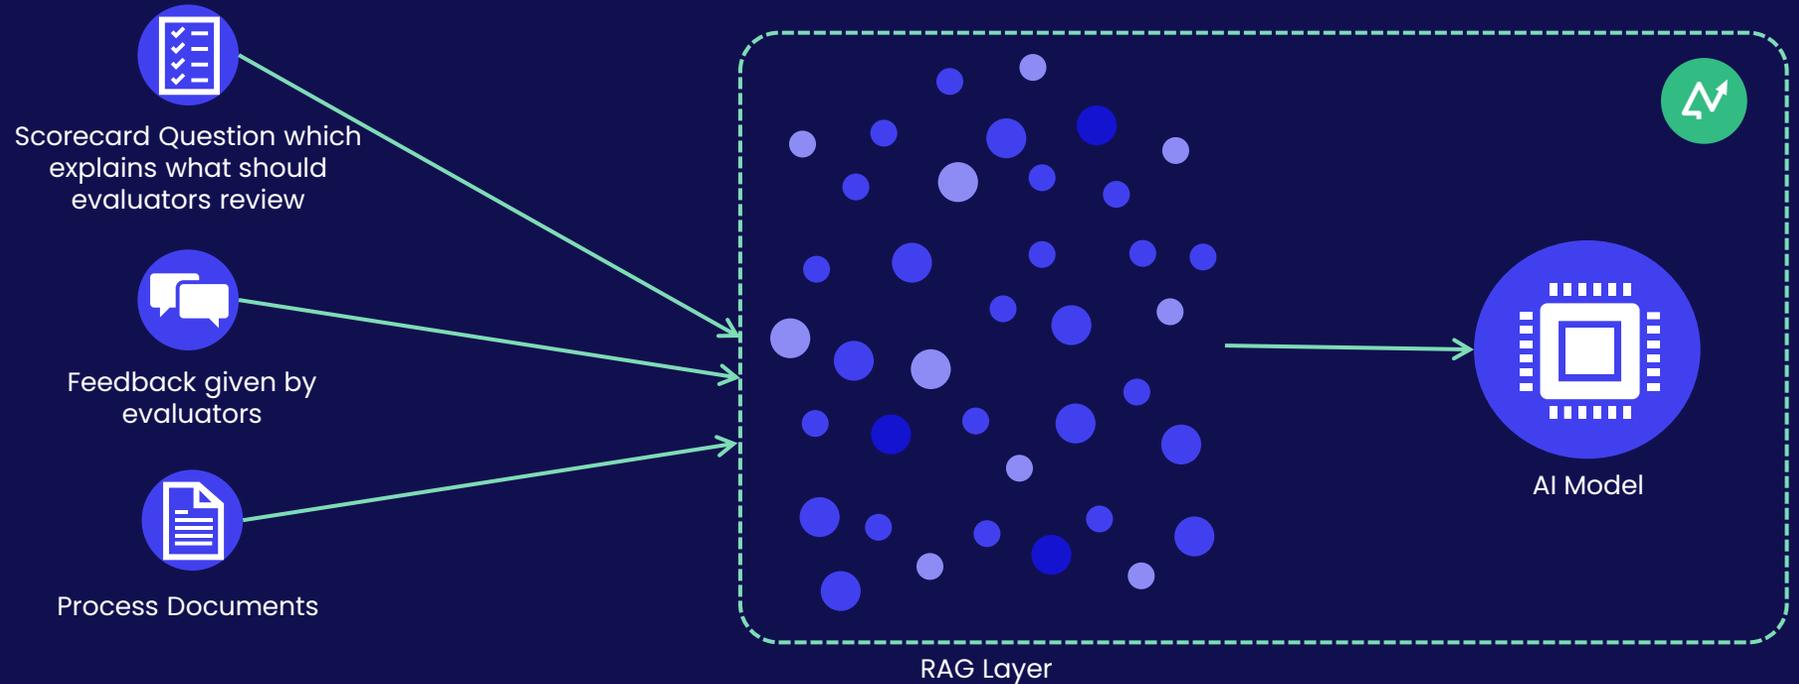
Leaptree AI - Data Management

Customer interactions and Past Evaluations

Security Layer Omits PII so it never gets to the AI

Machine learning trains on past interactions and evaluations

Only the model output gets saved

All data gets deleted from the Leaptree Platform

Leaptree AI

Leaptree AI securely connects to Salesforce via the Salesforce API to access interaction data. Before any processing takes place, all personally identifiable information (PII) is removed, ensuring sensitive details never enter our AI pipelines. The cleansed data is then used to train dedicated AI models for each company and each business process, so models remain specific, accurate, and unbiased. Unlike public LLMs or third-party tools, all training and processing happen entirely within the Leaptree system. Our multi-tenant architecture guarantees that each customer has a secure, isolated space where their models are stored and their data is held only temporarily. Once training or prediction is complete, the models are preserved but the raw data is permanently deleted, meaning customer data lives outside of Salesforce only for the short time required to generate AI results.

AI Models

Scorecard Question which explains what should evaluators review

Feedback given by evaluators

Process Documents

RAG Layer

AI Model

Leaptree AI learns directly from your own history of customer interactions and evaluator assessments to train models that complete assessments just like your manual evaluators. We use past evaluations as the foundation of training, ensuring that the AI's judgments reflect your company's unique standards and processes. Because every organization has different policies and expectations, we only use your data, never mixing it with data from other companies or even from other processes within your own business. The training process incorporates everything that shapes quality assurance: interaction transcripts showing exactly what agents said and did, past assessments and evaluator comments, Leaptree scorecard criteria and tips, and any process documentation you provide.

By combining predictive AI with generative AI, Leaptree goes beyond simply choosing outcomes on each evaluation criterion. Our predictive models replicate evaluator scoring, while our generative models provide tailored coaching feedback in natural language. Using both fine-tuned models and a retrieval augmented generation (RAG) layer, the system learns from your historical assessments so that comments reflect the same tone, content, and style used by your evaluators. The result is AI feedback that feels authentic to your business, aligned with your standards, and immediately actionable for your agents.

Frequently Asked Questions

How is data taken out of Salesforce?
We use the Salesforce Bulk API with secure OAuth JWT authentication. Access is managed via a Salesforce **External Client App**, which centrally controls permissions and scopes for data extraction.

Is data encrypted in transit and at rest?
Yes. All data is encrypted in transit via HTTPS/TLS 1.2 when using the Salesforce Bulk API, and authentication follows the Salesforce JWT flow:

https://help.salesforce.com/s/articleView?id=x cloud.remoteaccess_oauth_jwt_flow.htm

Data stored temporarily in AWS is encrypted at rest using AES-256 server-side encryption.

When is PII removed, and what data is included?
PII is removed immediately after data is retrieved from Salesforce and before it is processed further. We remove names, ages, emails, phone numbers, addresses, usernames, credentials, SSNs, passports, driver IDs, tax IDs, health IDs, national IDs (e.g., Aadhaar, NREGA, PAN, Voter IDs), bank account details, routing numbers, card numbers, CVVs, PINs, IBANs, SWIFT codes, license plates, VINs, and more.

Will our data ever be sent to another platform? Do you use ChatGPT or public LLMs?
No. Your data never leaves the Leaptree AI platform. It is stored only temporarily for training and predictions and then deleted. We do not use ChatGPT, public LLMs, or any third-party tools. All model training and inference happen inside the Leaptree platform.

Is our data ever used to train models for other customers?
No. Each customer's data is used exclusively to train their own models. Data is stored only temporarily in a tenant-dedicated space, then deleted once the training or prediction process is complete. No data is ever shared across customers or with third parties.

Where can we manage Leaptree AI, and where are AI assessments stored?
All management of Leaptree AI happens inside your Salesforce environment. AI assessment results are written back to Salesforce and deleted from the Leaptree platform after processing. No one logs directly into the Leaptree AI platform.

Where is the Leaptree AI platform hosted?
Leaptree AI runs fully on AWS. This includes secure data storage, model training, and predictions. Each of our customers has a secure and dedicated storage and their data is processed separately.

Where is our data processed physically?
By default, data is processed in a US East data center. However, Leaptree AI is designed to support storing and processing data in other regions if required. Contact us if you require your data to remain in a specific geographical region.

Frequently Asked Questions

How much data is needed for Leaptree AI to learn effectively?
More data is always better, but quality matters most. Detailed manual assessments, evaluator comments, tips, and process documentation greatly improve results. As a rule of thumb, we recommend at least 500–2000 assessments per process, but we can work with less. Contact us for guidance.

How are tenants isolated to prevent performance or security issues?
Each customer has a dedicated storage space for their data and models. Training and prediction jobs are fully isolated and can run in parallel. AWS infrastructure ensures Leaptree AI can scale securely and reliably across multiple customers while all using our platform at the same time.

How can we control Leaptree AI's access to Salesforce?
You can revoke Leaptree AI's access at any time through the Leaptree AI "External Client App" settings in Salesforce. You will be required to grant us access during the setup of Leaptree AI.

What data can Leaptree AI learn from?
Leaptree AI can learn from any relevant Salesforce data. This is configured during setup (cases, conversations, assessments, etc.). If needed, we can also securely connect to other platforms. Contact us for details.

How long is data stored outside Salesforce?
Only for the duration of model training or prediction. Once the process is complete, the raw data is permanently deleted.

How often are models retrained?
Models can be retrained on-demand depending on data volumes and business needs. We always recommend to continue your manual assessment process in parallel to Leaptree AI completing assessment. This way periodically retraining the model guarantees your latest process standards are being kept.

Frequently Asked Questions

Who has access to the private key used for our JWT OAuth flow?
Nobody. Each customer's private key is stored in a secure key management solution within the Leaptree platform. Keys are dedicated per customer.

What are the different parts of Leaptree AI?
Leaptree AI consists of two parts: (1) the Leaptree Optimize AI Salesforce package, which lets you configure, start model builds or predictions, and view AI results all within Salesforce; and (2) the Leaptree AI platform, which securely fetches your data temporarily, trains the AI model, and saves the results back into Salesforce.

Do you have a REST API for user provisioning, automation, or orchestration?
Leaptree AI integrates with Salesforce through the Leaptree Optimize AI package. Model builds, predictions, and data fetches are triggered through this integration. To view or manage results, you can use the standard Salesforce API, with all access controlled by Salesforce permissions and roles.

How is access to the prediction results controlled?
Prediction results are written directly back to Salesforce. Access control is handled entirely by your Salesforce configuration using users, roles, profiles, sharing rules and permission sets.

What is the best way to test Leaptree AI?
We recommend rolling out Leaptree AI in your full copy sandbox Salesforce environment with real customer and assessment data using **Shadow Mode**. This feature hides AI assessments from end users, allowing you to review results before making them visible.

What Single Sign-On (SSO) methods do you offer?
The Leaptree AI platform has no direct login capability, it is fully controlled from Salesforce through the Leaptree Optimize AI package.

Leaptree

Thanks for reading