UiPath The Foundation of Innovation™

Generative AI Security Handbook

Introduction

UiPath is committed to creating responsible AI solutions, protecting our customers' data, and complying with new and emerging privacy and security regulations. Our approach to generative AI (GenAI) technology is no different.

GenAI presents notable data security challenges. Foundational model providers often train their GenAI models on the data their customers send to them; if not securely handled, sensitive corporate and customer data could be compromised. The decentralized and non-transparent nature of many AI models makes it difficult to enforce traditional controls. Robust security measures and strict oversight are crucial to addressing these risks.

UiPath values data privacy and works to ensure our customers can protect their data and meet regulatory and compliance requirements. We strive to continuously improve UiPath GenAI-based features, while ensuring customer data security and privacy remain top priorities.

This white paper will explain the features and controls the UiPath Business Automation Platform incorporates to keep customer data safe and secure in our GenAI-based features and capabilities.

UiPath AI Trust Layer

Overview

The UiPath AI Trust Layer is a powerful management framework, ensuring GenAI-powered solutions are properly governed by our customers' data and personnel policies. It empowers organizations to responsibly manage using GenAI within the UiPath Platform.

The UiPath AI Trust Layer allows contextualized customer data, and customer interactions with UiPath products, to securely flow to trusted third-party large language models (LLMs), i.e. UiPath private subscriptions to the models or services. It enables transparency, trust, and control over the interactions between GenAI, your organizational data, and third-party LLMs.

It includes capabilities such as (but not limited to):

- Usage auditing and cost control: Visibility into GenAI feature costs and usage distribution through UiPath managed GenAI models as well as your connected models. Please note this feature is not yet released, but will launch in 2024.
- GenAI feature governance: Administrative controls (i.e.- being able to disable specific GenAI features) and usage at a tenant, group, and user level. Please note that this feature is currently in preview.
- A set of filters ensuring proper use of the LLMs (e.g., harmful content filtering): This is automatically built into, and enabled on, all third-party models.
- PII and sensitive data masking: Personally identifiable information (PII) data masking takes place as part of the corresponding GenAI activity in UiPath Studio and needs to be built into a specific workflow before the data gets ingested into a specific Service. It is not applied by default to all prompts. Please see here for additional information
- LLM gateway: A centralized and secure proxy service in which all customer data sent to LLMs is confidential, load-balanced, and adheres to defined data governance policies.

UiPath has released a new, growing set of GenAI functionality into its platform designed to help developers be more productive, and companies see more value faster from their automations.

UiPath upholds the highest standards of data privacy, customer experience, and automation performance. Our new LLM Gateway is a backend secure proxy server designed to maintain data privacy and compliance while load balancing and routing LLM requests from these GenAI features to a variety of UiPath-managed LLMs.

Data flow

The following diagram details the AI Trust Layer architecture and how the data flows through:

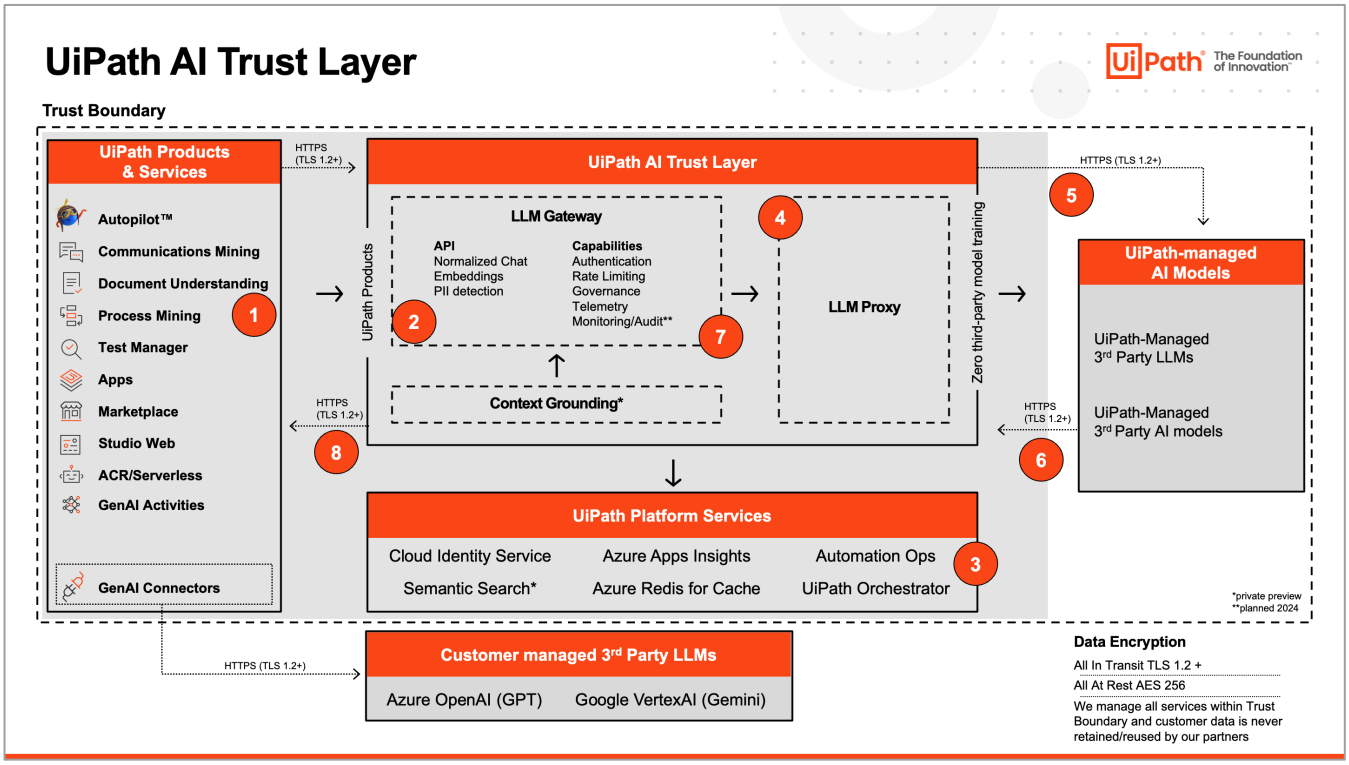

1. The first-party product calls one of the APIs available on LLM Gateway with the full prompt (customer input + internal prompt engineering + past data in chat scenario).

2. The LLM Gateway validates the call and authenticates the request sent by the first-party product either using an S2S token or a user token managed by the UiPath Cloud identity Service. Rate limiting is applied at this step if needed.

3. The LLM Gateway checks if there is any policy preventing the call and enforces it.

4. The call is forwarded to the LLM Proxy, which translates it and sends the request to the right third-party model using UiPath credentials.

5. The third-party model receives the request with the full prompt and processes it. The model is stateless, so the previously processed data is not accessible. Unless explicitly stated otherwise in our documentation, the third-party model is in the same region as the originated tenant.

6. The third-party model sends the completion back to the UiPath AI Trust Layer, no data is saved/logged outside UiPath Cloud.

7. The LLM Gateway logs metadata (no actual customer data) for internal telemetry (API used, product initiating the request, number of tokens in inputs and response, user/tenant ID initiating the request)

8. The LLM Gateway returns the completion to the first-party product which can perform the intended action.

End user security

Authentication and authorization

The LLM Gateway is a service that runs in Automation Cloud, and like other services, has deployed scale units associated with a tenant for a given region.

Authentication between the AI Trust Layer and the third-party LLMs is dependent on how the individual LLM provider handles authentication.

RBAC is managed at the product level, and authentication is done using our Cloud Identity Server. The authentication token is validated on the LLM Gateway side and checks are done by UiPath for any governance policies applicable, and a secure channel with UiPath secret credentials is stored in a key vault.

Server security

Data storage and retention policies

LLMs used by the UiPath services are provisioned by default in your tenant region. If they are not available in your region, they are provisioned in a fallback region. In both cases, no data is saved outside your region.

Please check the models relevant to the features used in each service in the UiPath data residency documentation here.

For GenAI activities, the data residency information can be found here.

UiPath data retention policies that comply with our standardized global regulations and contracts are also applicable with the AI Trust Layer.

UiPath also has agreements with third-party LLMs to prevent general model training on UiPath customer data. Any data used as context or prompts will be securely sent by UiPath only to third-party LLMs with which UiPath has trusted, secure agreements. UiPath discloses a list of all third-party processors and their locations.

Additionally, while GenAI features may use common instances of third-party models, data is not shared, cached, or retained between different customer interactions.

Features governance

In-platform controls

We introduced new policies to govern all third-party AI features used by our products in one centralized place, you can select exact product that will be able to consume third-party AI models and then deploy the policy to the right audience with tenant, group or user granularity. This way, you have full control on who can use which feature and limit your exposure.

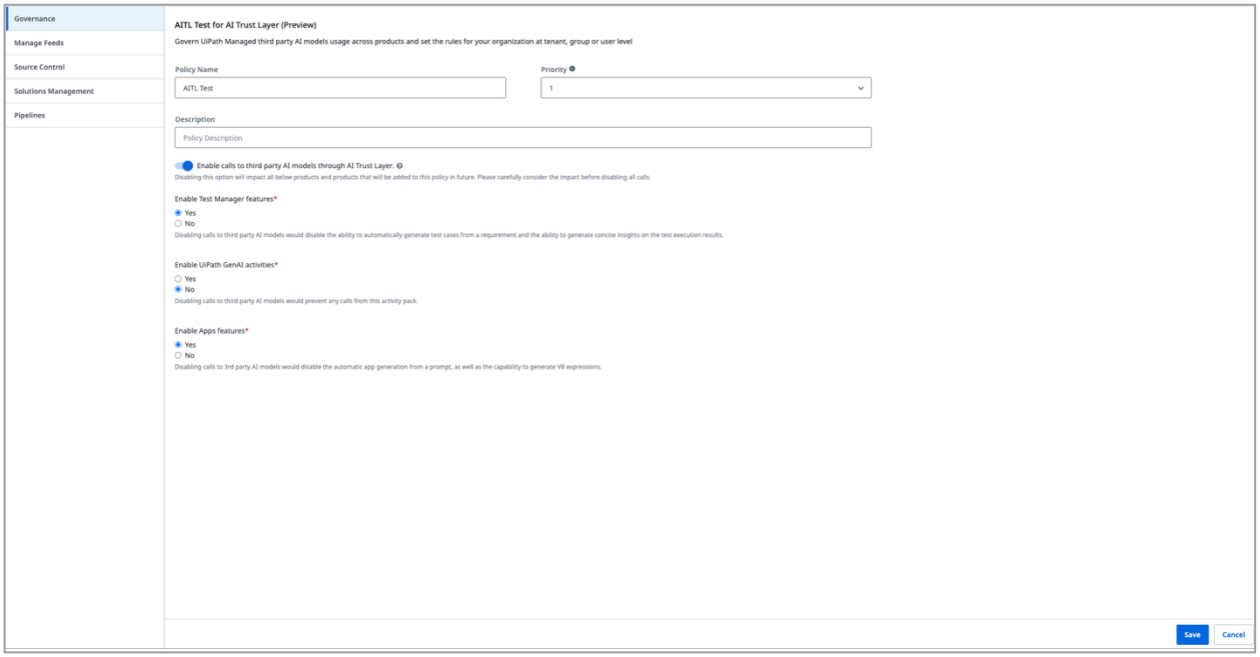

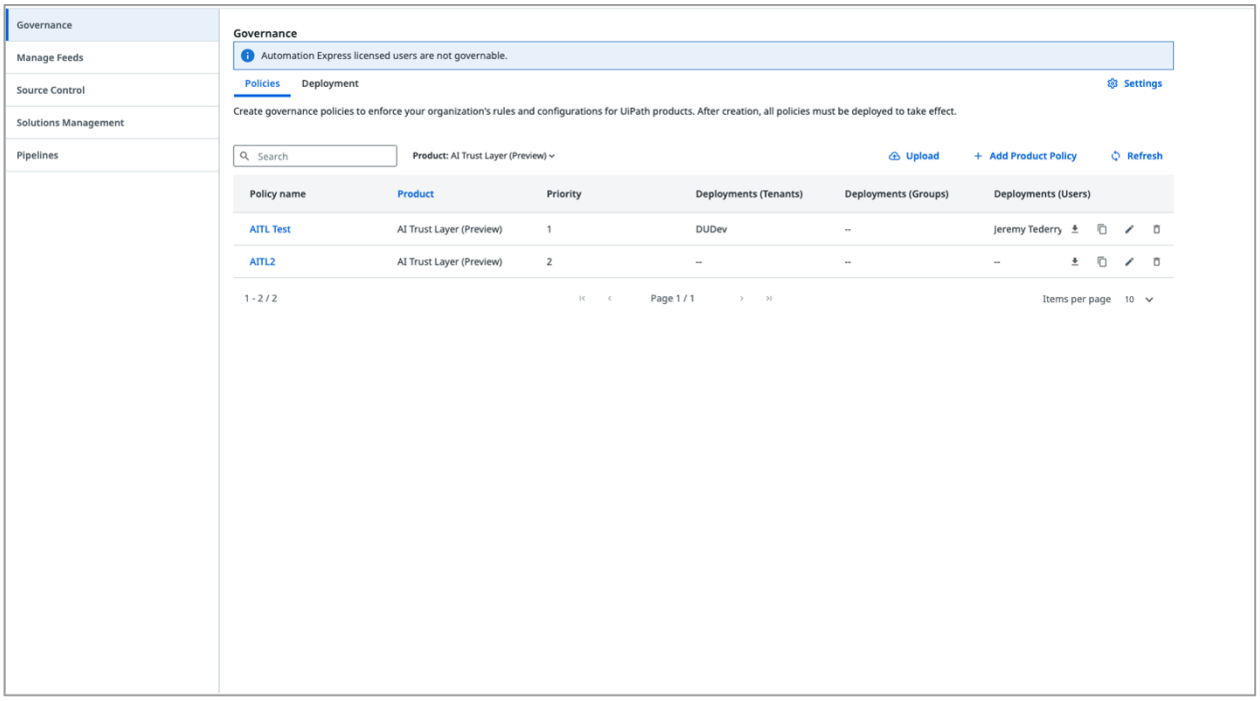

Logging and auditing

To enhance transparency for the end-user, it is on the roadmap for admins to be able to audit utilization of GenAI features.

This includes accessing detailed information on the identity of users using these features, prompts used, etc.

Threat management

Threat management mechanisms

Vulnerability management is handled directly by our LLM providers – please check the vulnerability management documentation for the LLM provider you are using for more specific details on how this is being addressed.

To address denial-of-service attacks, we have implemented several safety measures such as the implementation of rate limiting, authentication, and licensing. UiPath is currently implementing multiple layers of rate limiting, the complexity of which is contingent on the specific licensing or product.

LLM providers such as Microsoft, among others, also have rate limits in place. Our strategy to combat this limitation is to secure dedicated compute capacity. The UiPath team is proactively acquiring additional capacity to meet the fluctuations in demand.

For every third party we contract with, we ensure they do not save the data sent to their models. UiPath has standard agreements with all third-party sub-processor vendors, which are assessed by our privacy and security teams to protect the governance and integrity of customer data.

Data storage and encryption

Encryption and data travel security

Data used with GenAI features will be encrypted using TLS 1.2, transactional data at rest will be encrypted with AES 256.

UiPath features communicate through the LLM Gateway service to LLM Models using the LLM Gateway service to the third-party LLM using REST APIs provided by the vendors over the HTTPS protocol.

Data handling and segregation

UiPath agreements with third-party LLMs prevent UiPath customer data retention.

Data is sent securely via transit to the third-party model provider, but data may not be used for training, logging, or monitoring.

Feature deep-dive

UiPath Integration Service

Features that use GenAI within Integration Service include:

- GenAI Activities (Studio & Apps)
  - Content Generation
  - PII Detection (Microsoft Text Analysis)
  - Detect Language (Microsoft Text Analysis)
  - Translate
  - Rewrite
  - Summarize
  - Email Generation
  - Etc.

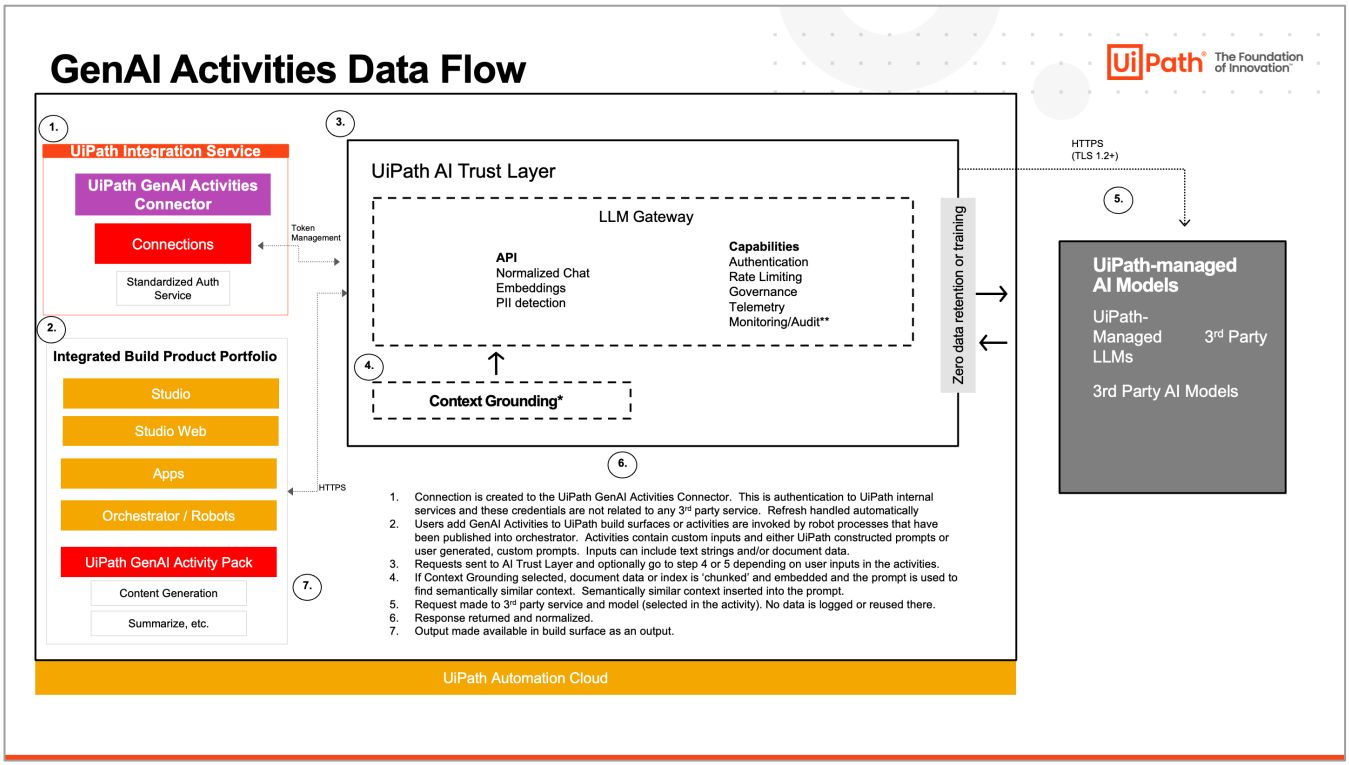

**Integration Service Connectors (like OpenAI, Azure OpenAI, Vertex, Bedrock, etc.) do not use the AI Trust Layer, but rather execute requests directly to their respective cloud services.

Communications Mining

Features that use GenAI to enhance the Communications Mining experience include:

- Generative Extraction
- Generative Annotation
- Zero-shot discovery
- Conversational Filters (Autopilot)

Additionally, the following LLMs are used to power some of these features behind the scenes:

- CommPath LLM
- GPT models via AI Trust Layer

Communications Mining Architecture

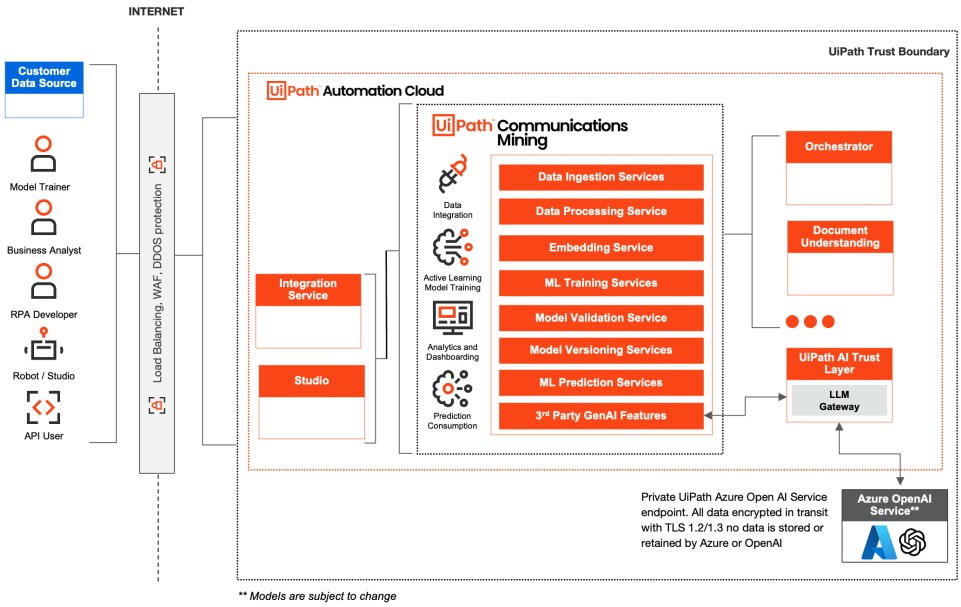

** *Models are subject to change*

Encryption
- All In Transit TLS 1.2 +
- All At Rest AES 256

Storage
- Azure Persistent SSD
- Azure Blob Storage

Components and Flow
- Native integrations to MSFT Exchange and Salesforce or connect to any system with UiPath Integration Service
- Proprietary transformer-based language models to extract semantics for efficient model training
- Unsupervised learning models and Active Learning engine and interface to train customer specific supervised models
- Real time aggregate statistics for analytics and insights
- Real-time model validation and model lifecycle management. With downstream connectors for end-to-end automation

Communications Mining Email Data Flow

1. All data encrypted at rest
2. All data in transit encrypted with HTTPS TLS 1.2/1.3
3. Private UiPath Azure Open AI Service endpoint. Multi-tenant endpoint used by multiple customers in region. All data encrypted in transit with TLS 1.2/1.3 no data is stored or retained by Azure or OpenAI

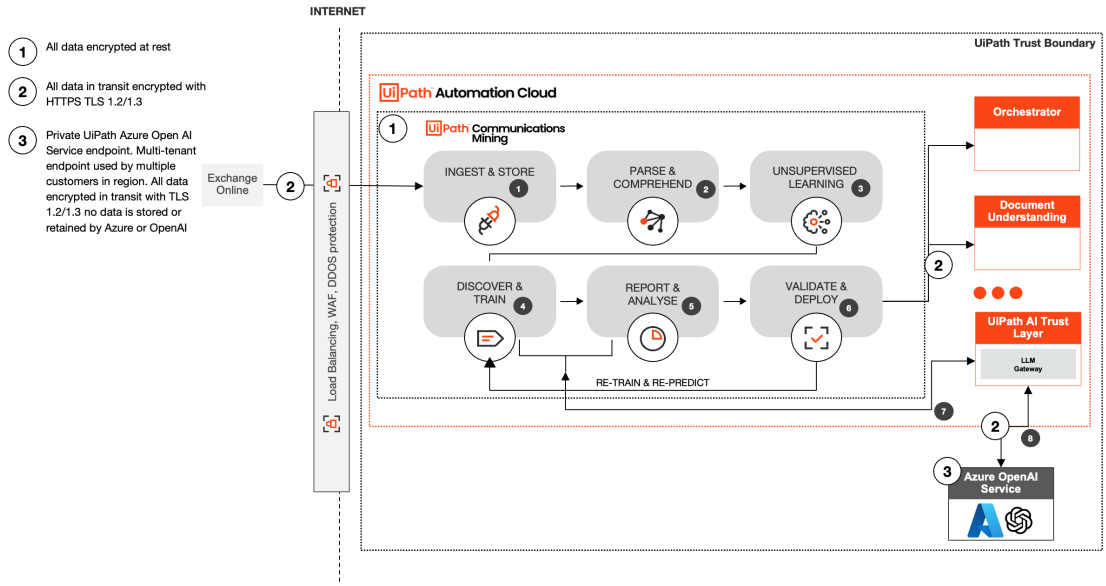

1. Pre-built connectors for ingestion into historic comms store. Proprietary ML segmentation and cleaning engine to clean data.
2. Proprietary transformer models extract semantics for data efficient learning.
3. Proprietary unsupervised learning models identify common intents and constantly search for new ones.
4. Train customer specific supervised models efficiently in our Proprietary active learning engine and interface.
5. Real-time aggregate statistics for analytics and insights
6. Real-time model validation and model lifecycle management. With downstream connectors
7. Optional feature enhancements relying on Azure OpenAI Services
8. Integration to Azure OpenAI Services, prevents data storage and use of data for training models

Document Understanding

Features that use GenAI to enhance the Document Understanding experience include:

- Active Learning
- Autopilot
- Generative Extraction
- Generative Classification
- Generative Validation

Additionally, the following LLMs are used to power some of these features behind the scenes:

- DocPath LLM

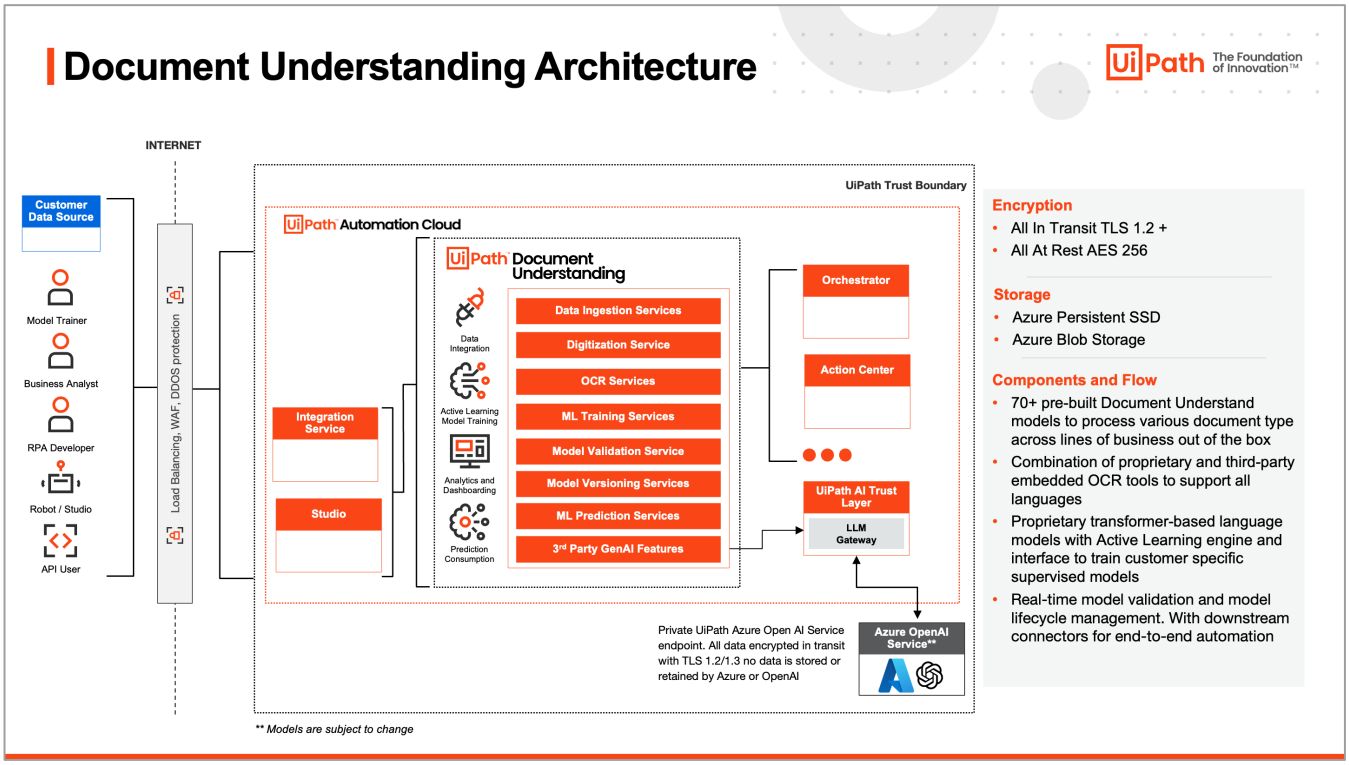

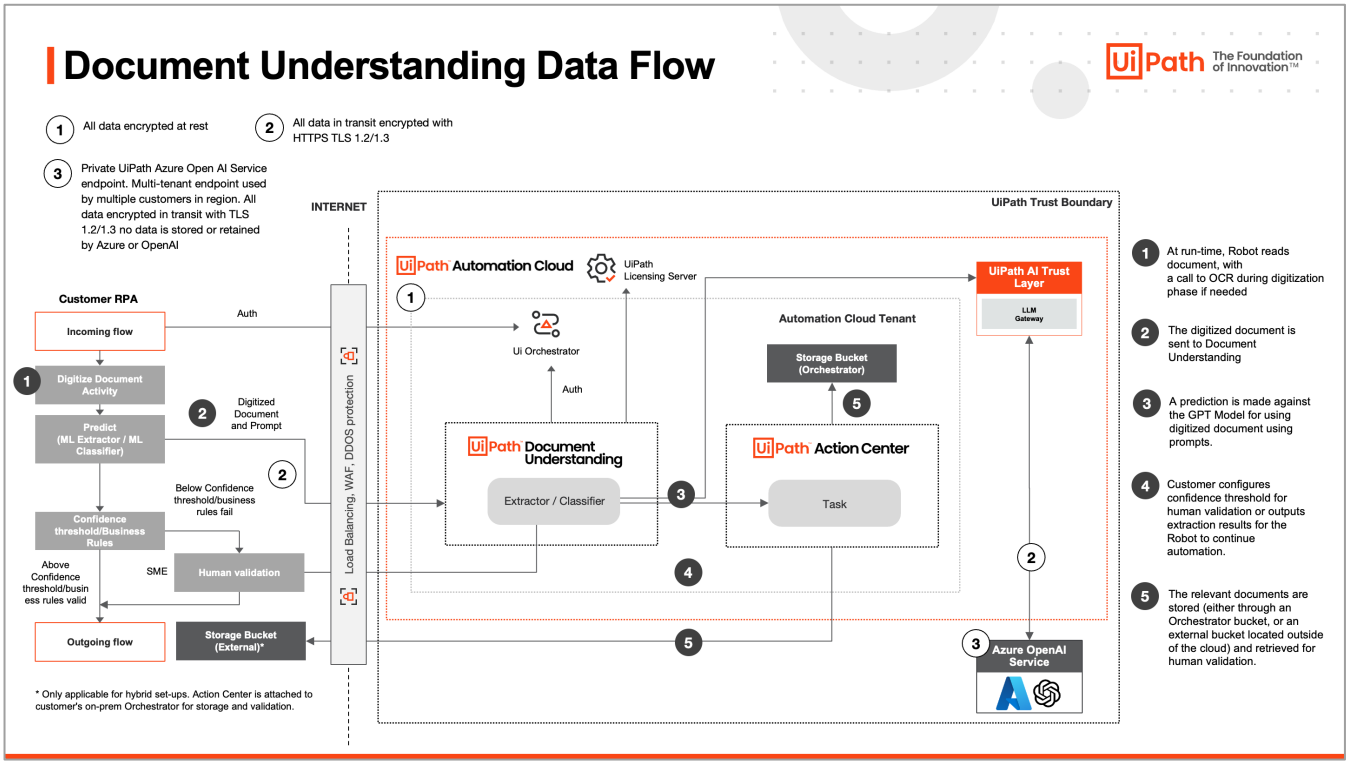

UiPath large language models

UiPath DocPath and CommPath are the underlying LLMs used to power certain GenAI features in Communications Mining and Document Understanding.

The new UiPath family of LLMs delivers higher accuracy, consistency, predictability, time to value, and empowers customers to transform their business environments with the latest GenAI capabilities on the market.

- **DocPath:** UiPath DocPath is an LLM trained to process any document out-of-the-box with little to no training required, including free-form unstructured data and tables. This is an underlying LLM that will be used to drive certain Document Understanding features.

- **CommPath**: UiPath CommPath is an LLM trained to process communications of varying complexity, including those with multiple requests and fields, and the relationships between them. This is an underlying LLM used to drive specific features in Communications Mining, starting with generative extraction.

These models are tailored for specific tasks related to document processing, and they are fine-tuned on enterprise relevant data, meaning:

- Increased accuracy of the models tailored for business documents and communications of different types
- Advanced unstructured data processing, including complex communications, free-form documents and complex tables
- Accelerated time to value owing to reduced model training effort improved human-in-the-loop fine-tuning
- Robust security, compliance and governance paired with comprehensive controls and guardrails

Please note: CommPath LLM and DocPath LLM **do not** go through the AI Trust Layer, and all security-related aspects pertaining to the GenAI features powered by these LLMs are managed directly within the services themselves (i.e. - within Communications Mining and Document understanding, respectively).

For more information on how security is handled, please refer to the official UiPath Communications Mining and UiPath Document Understanding security handbooks.

Frequently asked questions

LLMs and data management

What agreements does UiPath have with 3rd party LLMs for Generative AI features?	UiPath has standard agreements with all third party sub-processor vendors which are assessed by our privacy and security teams to protect the governance and integrity of customer data.
Is my data being used to train 3rd party LLMs?	No, UiPath has agreements with all of its 3rd party processors – including LLM providers – to not allow customer data passed through the UiPath platform to be used for general model training. Find additional details on our UiPath Trust and Security site.
Where is my data being sent or used with 3rd party LLMs?	Any data that is used as context for prompts will be securely sent by UiPath only to 3rd party LLMs in the UiPath trusted ecosystem. UiPath has standard agreements in place with 3rd party LLMs and other model processors to not retain, use for training, or share customer data. Find additional details on our UiPath Trust and Security site.
How does UiPath handle Personal Identifiable Information (PII) in my data?	It is the responsibility of UiPath customers to carefully manage the input of sensitive data including PII into UiPath products which may leverage 3rd party products or services (including LLMs) , however UiPath agreements with 3rd party LLM providers prevents the storage and retention of any data.. UiPath treats all customer data with the highest enterprise grade security and will help customers appropriately manage any reported PII that has been sent to UiPath.
Which models are my data being sent to?	UiPath discloses a list of all 3rd party processors in its sub-processor agreement. Find additional details on our UiPath Trust and Security site.
Where is my data being processed?	The location of UiPath software is disclosed in the Data Privacy Agreement; the location of 3rd party sub-processor data processing is disclosed in the UiPath sub-processor agreement. Find additional details on our UiPath Trust and Security site.
How can I control how users in my organization use our data with AI features?	We introduced new policies to govern all third-party AI features used by our products in one centralized place, you can select exact product that will be able to consume third-party AI models and then deploy the policy to the right audience with tenant, group or user granularity. This way, you have full control on who can use which feature and limit your exposure.
Can I use my own LLM for automation processes?	With UiPath Integration Service, you can leverage Generative AI connectors for Azure, Google, AWS, and others within your automation workflow. See the

	latest here. We will soon offer capabilities to allow customers to bring their own LLM tenant for specific generative AI runtime features.
Can I fine tune UiPath managed LLM or generative AI models?	UiPath does not provide tooling to support LLM or other generative AI model fine tuning.

LLM Gateway

How do UiPath generative AI features communicate with LLM models through the LLM Gateway (What is the communication protocol)?	UiPath features communicate through the LLM Gateway to LLM Models using LLM Gateway Service to the third-party LLM using REST APIs provided by the vendors over the HTTPS protocol.
How is the LLM Gateway connected to a tenant?	The LLM Gateway is a service that runs in Automation Cloud, and like other services, has deployed scale units associated with a tenant for a given region.
How is data sent securely to the LLM Gateway? How is it encrypted?	Data is encrypted using TLS 1.2 for data in transit, and AES 256 for data at rest. Read more here.
Do all AI features use the LLM Gateway?	Not currently, just for generative AI features; in the future, all AI model invocation from UiPath products will flow through the LLM Gateway.
What is the expected response times?	Response time for UiPath generative AI features largely depends on the UiPath feature availability, request token size, and capacity of the invoked LLM. This may be variable based on the model provider and version.

Data storage

Where is data being stored?	Generative AI features will only store data related to the transaction between the automation and the generative AI feature on UiPath Automation Cloud servers. No customer data or automation transactional data will be stored or shared with UiPath 3rd party LLMs.
How can we manage IP, sensitive data, and copyright when using generative AI for content creation?	UiPath will securely protect the integrity and privacy of any customer data sources connected to Automation Cloud. We recommend customers use discretion with the choice and access management for developers who may use these data for any automation, including those using generative AI features.

Data transmission

Is my data encrypted?	All data used with UiPath products are treated the same; Data at rest is encrypted with AES 256, data in transit is encrypted using TLS 1.2. For additional details please see here.
How can we control access to certain sources of data for users and generative AI models?	UiPath will support organization, tenant, and user group RBAC for all AI features and products, starting with generative AI.

Data security

How does UiPath protect PII data?	UiPath will treat any customer PII with the highest scrutiny and security to protect the integrity of our customers, as is stated here. We will also detect any data shared with UiPath upon request. We will soon introduce new tooling for admins to specify the entities and mitigation of any PII that could be used with AI models – stay tuned!
How do we protect the privacy of individuals in data used with generative AI?	If there is indiscriminate personal information within data used for generating predictions, customers may request UiPath to delete these data at any time.

Data governance

What kinds of safeguards can be put in place to reduce mistakes or hallucinations from these models?	All UiPath AI products and features leverage a human in the loop to review and edit AI model predictions before they are used in downstream automation to reduce errors, ensure accuracy and consistent predictions.
How can we provide feedback to refine and improve Generative AI model responses?	UiPath will provide mechanisms to capture explicit user feedback and quality scores related to the LLM responses – stay tuned!
How can we prevent generative AI models from generating harmful or inappropriate content?	Please ensure any prompts used with generative AI features are free from harmful or inappropriate content to prevent misinterpretation in the tone of response from the model. UiPath is partnering with a number of AI model providers who are building safeguards into these LLMs to filter and prevent harmful content from reaching the models. UiPath encourages its users to leverage UiPath Academy

	and other resources to educate them on best practices for interacting with generative AI and other conversational AI models.
How can we mitigate bias that can be used to train/emerge from AI models?	Bias can occur in data as well as in the trainers who are training these AI models. To mitigate bias in data, UiPath encourages its customers to maintain a high level of scrutiny on the sources of data, content of data, and distribution of data types that are used by AI products and features. UiPath Generative AI features powered by 3rd-party vendor LLMs have standard opt-out agreements with UiPath to prevent any training of general models.
How can we monitor and audit generative AI model responses?	The AI Trust Layer will maintain and provide upon request telemetry data related to the invocation of generative AI prompts and responses.
What should be done to address the legal liabilities and responsibilities with the use of generative AI?	UiPath upholds a high standard of data security and compliance for all products. Before using generative AI features ensure a complete understanding of the UiPath platform, UiPath 3rd party sub-processors, and the UiPath product/feature to ensure the proper use of these tools.