


# SCALING MARKETING AND DOCS WITH A PRIVACY-SAFE RAG MODEL

---

Sean Falconer and Manny Silva  
Head of Marketing and Head of Docs | Skyflow



# CHALLENGE: SCALING CONTENT CREATION

# DOCUMENTATION

The screenshot shows the Skyflow documentation website. The top navigation bar includes the 'skyflow DOCS' logo, a search bar, and links for 'API Reference', 'SDKs', and 'Go to Studio'. A left sidebar contains a navigation menu with categories like 'Home', 'Fundamentals', 'Vaults and storage', 'Access controls', 'Tokenization', 'Connections', 'Secure workflows', 'Financial data', 'SDKs', and 'References'. The 'What is Skyflow?' article is selected in the sidebar. The main content area features the article title, a brief introduction, a bulleted list of key points, and a section on how developers can get the best of both worlds. A right sidebar lists related topics such as 'What is a data privacy vault?', 'The Skyflow Data Privacy Vault', 'Governance', 'Interoperability layer', 'Secure storage & trusted infrastructure', and 'Next steps'.

skyflow DOCS

Search docs

API Reference

SDKs

Go to Studio

Home

Fundamentals

What is Skyflow?

Get started with Skyflow

Explore what Skyflow can do

Authenticate

Accounts and environments

Security best practices

Get data into Skyflow

Vaults and storage

Access controls

Tokenization

Connections

Secure workflows

Financial data

SDKs

References

## What is Skyflow?

Companies are collecting more sensitive data than ever before. And with more data, there is more risk. The risk associated with managing sensitive data forces companies to make a tradeoff: data privacy or data utility.

- If you want data privacy, you can lock sensitive data in silos. But this causes data to go unutilized, which can put you at a disadvantage.
- If you want data utility, you can try building complex privacy tools and programs to allow your team to leverage the sensitive data. But this is extremely difficult and can often go wrong, putting your data at risk.

How can developers get the best of both worlds? That's where data vaults come in.

## What is a data privacy vault?

The concept of data privacy vaults was born at companies like Apple, Google, and Netflix. A data privacy vault is a secure, isolated database designed to store, manage, and use sensitive data. Let's break that down:

- *Secure*: Vaults have encryption, tokenization, masking, and other privacy-preserving technologies built in.
- *Isolated*: Vaults are segregated from your other infrastructure and services, and they're only available

In this article

### What is a data privacy vault?

The Skyflow Data Privacy Vault

Governance

Interoperability layer

Secure storage & trusted infrastructure

Next steps

# DOCUMENTATION, BLOG POSTS

The screenshot displays the Skyflow website's documentation page. The top navigation bar includes 'skyflow', 'Products', 'Solutions', 'Company', 'Resources', and a 'Get a Demo' button. A left sidebar lists various documentation categories such as 'Fundamentals', 'Vaults and storage', 'Access controls', 'Tokenization', 'Connections', 'Secure workflows', 'Financial data', 'SDKs', and 'References'. The main content area features a blog post titled 'AI Data Privacy: A New Product Essential' by Amruta Moktali, Chief Product Officer, dated October 26, 2023. The post includes a sub-headline, an introductory paragraph, and a section on related articles.

skyflow D

Home

skyflow

Products Solutions Company Resources

Get a Dem

Fundamentals

- What is Sky
- Get started
- Explore wha
- Authenticat
- Accounts ar
- Security be
- Get data int

Vaults and stor

Access control

Tokenization

Connections

Secure workflo

Financial data

SDKs

References

October 26, 2023

Amruta Moktali  
Chief Product Officer

## AI Data Privacy: A New Product Essential

**In the ever-evolving tech landscape, where innovation meets the promise of generative AI, one crucial question looms: Can we truly harness the transformative power of AI while safeguarding sensitive data?**

As the Chief Product Officer at Skyflow, where we've built a privacy platform that isolates, protects, and governs sensitive data, I've witnessed advancements in large language model (LLM)-based AI tools like Chat GPT with growing interest. Throughout my career, I have worked closely with developing AI technologies at companies such as Salesforce (as VP of Product Management), Topsy, and others to democratize data and AI. And, like many leaders at global companies, the impact of generative AI on data privacy and technical solutions to address this is always top of mind.

Of course, AI isn't a new development - in fact, it's been a part of the technological landscape since the 1950s and continues to play a significant role in our daily lives. In the early days, we mostly worried about issues like bias and fairness. However, with the emergence of generative AI, where training on vast datasets has become commonplace, a new concern has come to the forefront: data privacy for sensitive data like PII.

**Related articles**

- Private LLMs: Data Protection Potential and Limitations
- Addressing HubSpot CRM's HIPAA Compliance Limitations with Skyflow
- China Data Residency: A Guide to Compliance with PIPL & CSL

# DOCUMENTATION, BLOG POSTS, PRESS

The screenshot displays the Skyflow Pulse 2.0 website. The main navigation bar includes 'Home', 'Fundamentals', 'VENTURE CAPITAL', 'M&A', 'PROFILES', 'CONTACT', and a search icon. The article title is 'Skyflow: Data And AI Privacy Company Secures \$30 Million' by Amit Chowdhry, dated Mar 29, 2024. The article text describes Skyflow's Series B funding and its data privacy vault technology. A sidebar on the left contains a navigation menu and a 'Related articles' section with links to 'Private LLMs: Data Protection Potential and Limitations', 'Addressing HubSpot HIPAA Compliance Limitations with Skyflow', and 'China Data Residency: A Guide to Compliance with PIPL & CSL'.

skyflow D

Home

Fundamentals

What is Sky

Get started

Explore wha

Authenticat

Accounts ar

Security be

Get data int

Vaults and stor

Access contro

Tokenization

Connections

Secure workflo

Financial data

SDKs


References

skyflow

## PULSE 2.0

🏠 VENTURE CAPITAL M&A PROFILES CONTACT 🔍

### October 26, 2023

 Amruta  
Chief Pr  
Officer

## Skyflow: Data And AI Privacy Company Secures \$30 Million

*By Amit Chowdhry || Mar 29, 2024*

Data privacy vault company Skyflow announced it has raised \$30 million to extend its Series B funding. Khosla Ventures led the round, joined by prior investors Mouro Capital, Foundation Capital, and Canvas Ventures. This funding round follows revenue growth of more than 210% and quick adoption in the emerging market of sensitive data protection for large language models (LLMs).

Skyflow is a data privacy vault built to radically simplify how companies isolate, protect, and govern their customers' most sensitive data. With its global network of data privacy vaults, it is a comprehensive solution for companies looking to implement LLMs and securely meet complex data localization requirements. Skyflow supports nearly a billion user data records for global customers like GoodRx, Lenovo, and Hippocratic AI and processes more than two billion API calls quarterly.

Skyflow is powered by proprietary polymorphic encryption, the only way of encrypting data on the market that protects data without sacrificing its usability for critical business operations, such as analytics, marketing, and customer support.

LLM adoption has been increasing across the market even though few data

**Related articles**

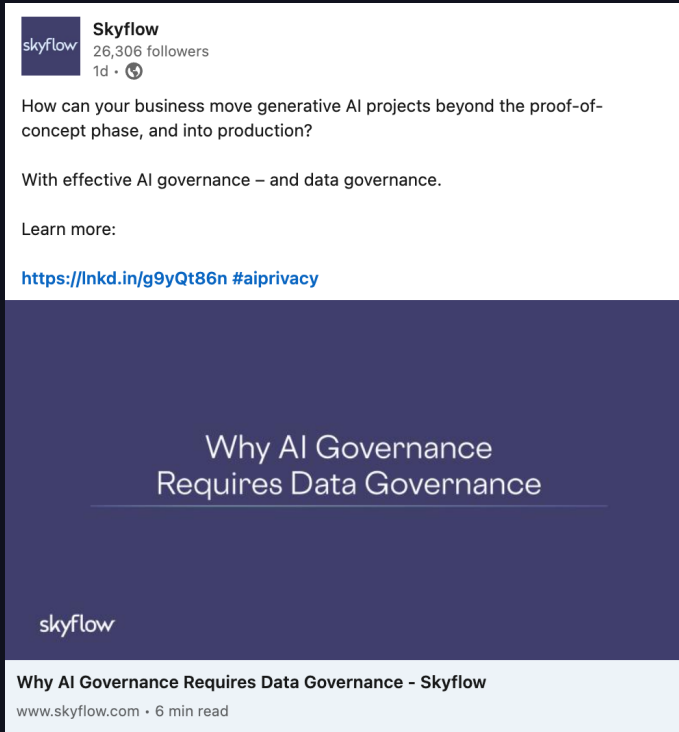
Private LLMs: Data Protection Potential and Limitations

Addressing HubSpot HIPAA Compliance Limitations with Skyflow

China Data Residency: A Guide to Compliance with PIPL & CSL

# SOCIAL PROMOTION

Everything needs promotion, ads, a webinar



**skyflow** Skyflow  
26,306 followers  
1d · 🌐

How can your business move generative AI projects beyond the proof-of-concept phase, and into production?

With effective AI governance – and data governance.

Learn more:

<https://lnkd.in/g9yQt86n> #aiprivacy

Why AI Governance Requires Data Governance

skyflow

Why AI Governance Requires Data Governance - Skyflow  
www.skyflow.com · 6 min read



**skyflow** Skyflow @SkyflowAPI · Mar 27

How can you ease compliance with China's Personal Information Protection Law (PIPL), Cyber Security Law (CSL), and related data residency requirements?

Learn how:

China Data Residency:  
A Guide to Compliance with PIPL & CSL

China Data Residency: A Guide to Compliance with PIPL & CSL - Skyflow

From skyflow.com

# THE CONTENT MACHINE DOESN'T STOP



## Challenges

- Small team, how do we keep up?
- How do we maintain quality and consistency?
- How do we scale beyond ourselves?

# THE UNSTRUCTURED DATA CHALLENGE

Where is the  
reference content?

Drive

Email

Confluence

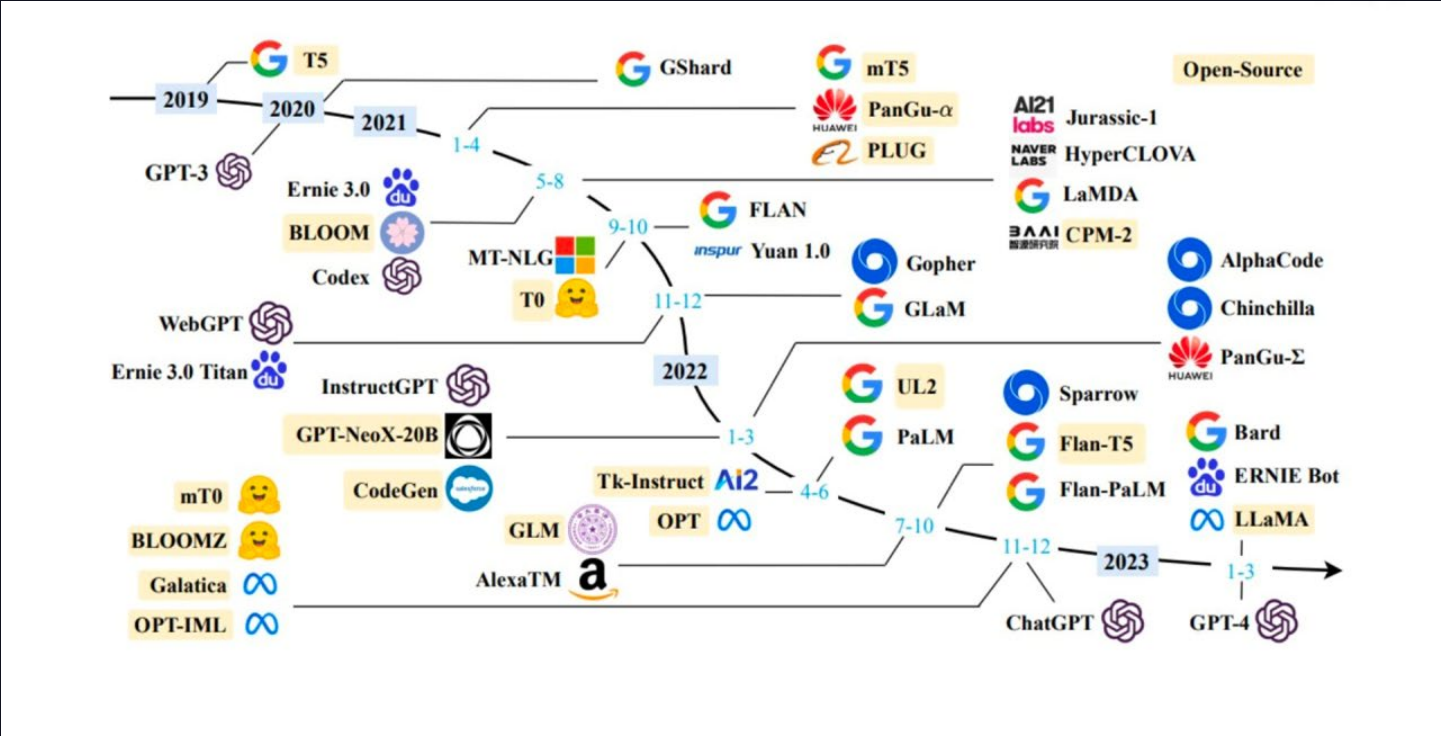
Guru

Salesforce

Slack



# LLMs TO THE RESCUE



# WHY NOT USE CHATGPT?



**You**

Write a press release announcing latest round of funding.



**ChatGPT**

Certainly! Here's a basic template for a press release announcing a new funding round. You can customize it with specific details about your company, the amount raised, the investors involved, and what the funding will be used for:

---

## FOR IMMEDIATE RELEASE

[City, Date] — [Your Company Name], a leading innovator in [industry/sector], is thrilled to announce the successful completion of its latest funding round, securing [\$Amount] in capital. This investment was led by [Leading Investor], with participation from [Other Investors] and marks a significant milestone in our journey to reshape [industry/sector].

This new injection of funds will enable [Your Company Name] to accelerate our growth, expand our operations, and continue to innovate within [specific area or service]. The investment will particularly focus on enhancing our [specific products or services], as well as expanding our reach in [new markets or regions].

"[Quote from the CEO or founder about the funding round and future plans]," said [CEO/Founder's Name], [Title], of [Your Company Name].

## Problems

- Tone is off
- Missing context
- Doesn't have access to our files
- Content visible to auditors



# WHAT IF WE BUILD OUR OWN CONTENT CO-PILOT?



## Customized

It would know Skyflow, have access to our files, speak our language



## Content Formats

Output from the LLM can be custom to our needs, like creating a blog post, press release, or documentation

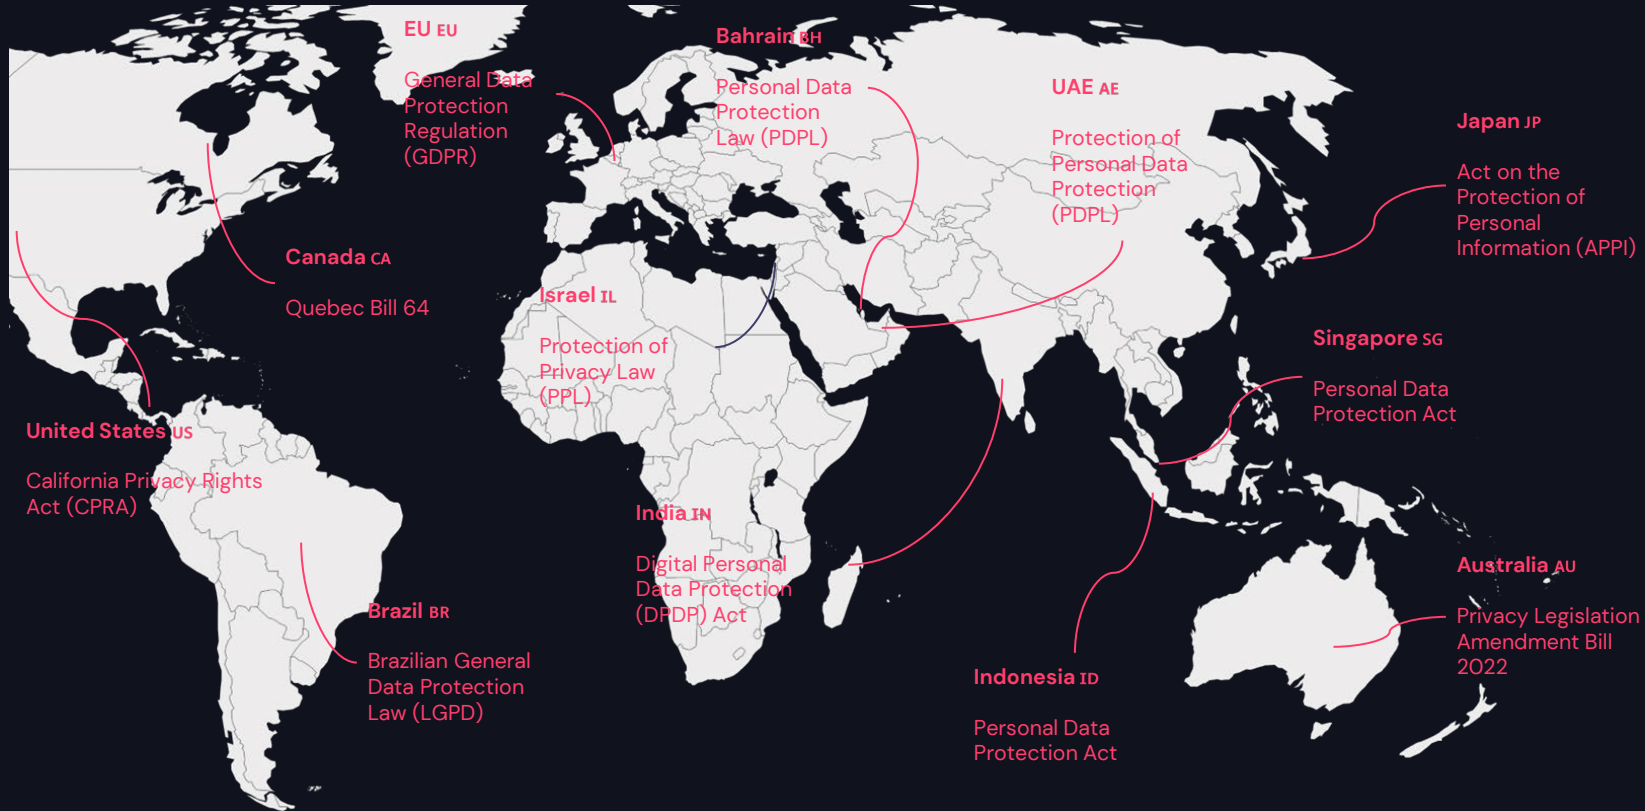


## Private & Secure

With Skyflow's LLM Privacy Gateway, we can keep data secure and private

# PRIVACY AND SECURITY: DOING IT RIGHT

# AN INCREASINGLY REGULATED WORLD



# LLMS MAKE DATA PRIVACY HARD

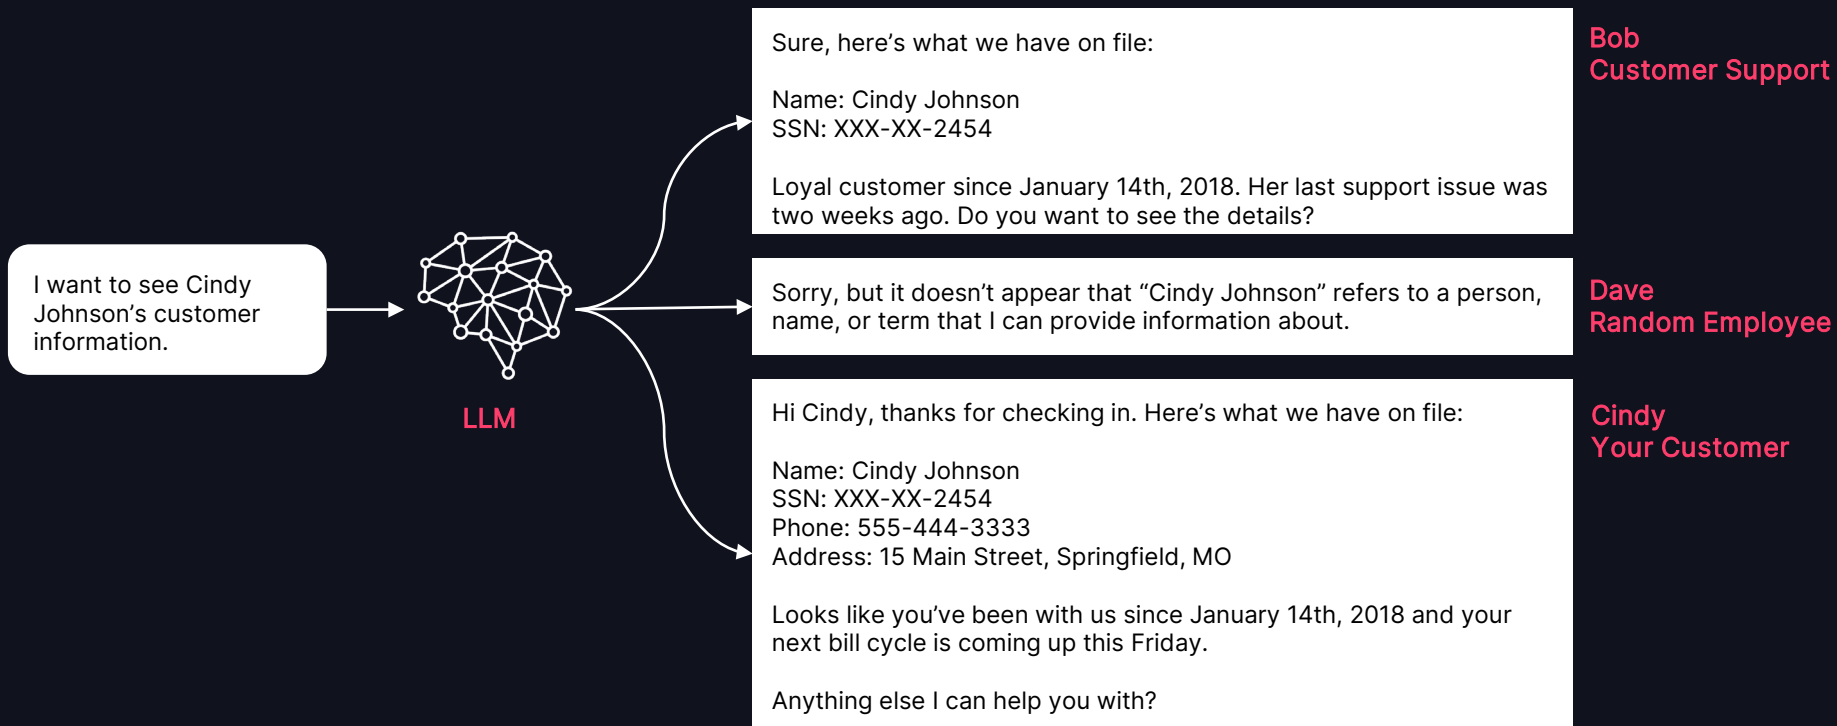
No rows

No columns

No delete button

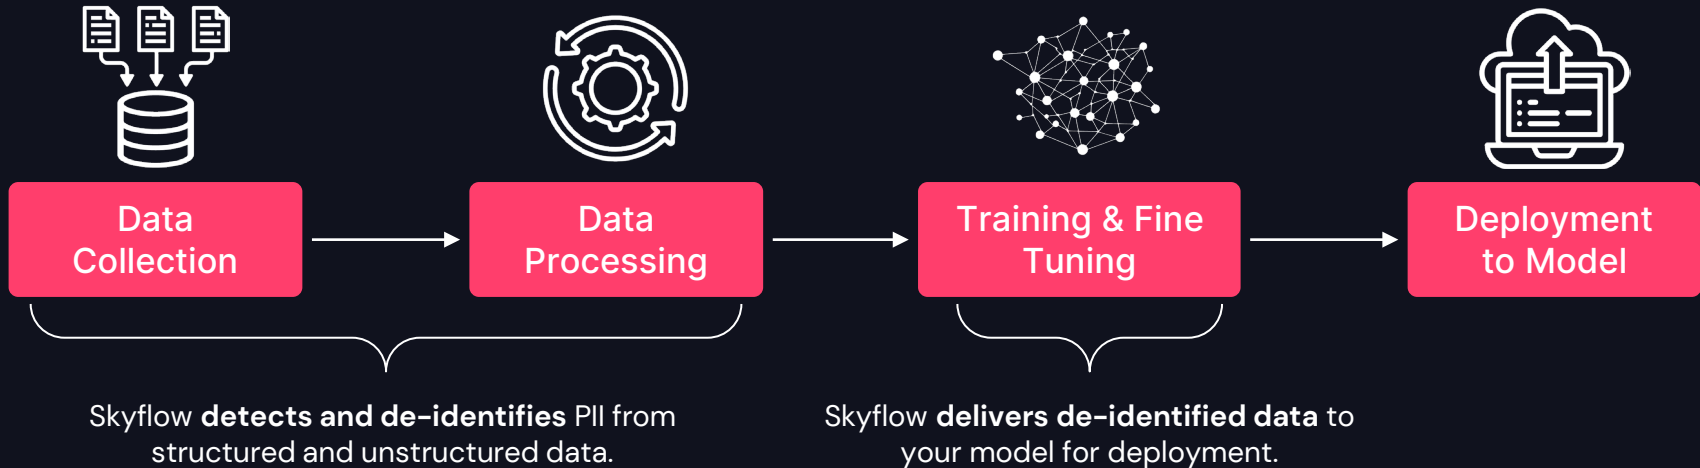
The screenshot shows the Yahoo Finance website interface. At the top, there is a navigation bar with links for HOME, MAIL, NEWS, FINANCE, SPORTS, ENTERTAINMENT, LIFE, SEARCH, SHOPPING, YAHOO PLUS, and MORE... Below this is the 'yahoo!finance' logo and a search bar with the placeholder text 'Search for news, symbols or companies'. To the right of the search bar are links for 'Sign in', a notification bell, and a mail icon. The main content area features a 'FORTUNE' section with the headline 'Apple, Goldman Sachs, and Samsung among growing list of companies banning employees from using ChatGPT at work'. Below the headline is the author's name 'Paige McGlaufflin' and the date 'May 19, 2023 - 7 min read'. To the left of the article text are social media sharing icons for comments (2), Facebook, X, and Email. The main image of the article shows the OpenAI logo and the text 'OpenAI' in a purple font, with a smartphone in the foreground displaying the ChatGPT logo and the text 'ChatGPT'. To the right of the article is a sidebar with a photo of a man and a woman, and a section titled 'The Highest Paid Bodyguards in...' with a blue button labeled 'undefined'.

# WHO SEES WHAT, WHEN, AND WHERE?



# SKYFLOW PROVIDES PRIVACY DURING TRAINING

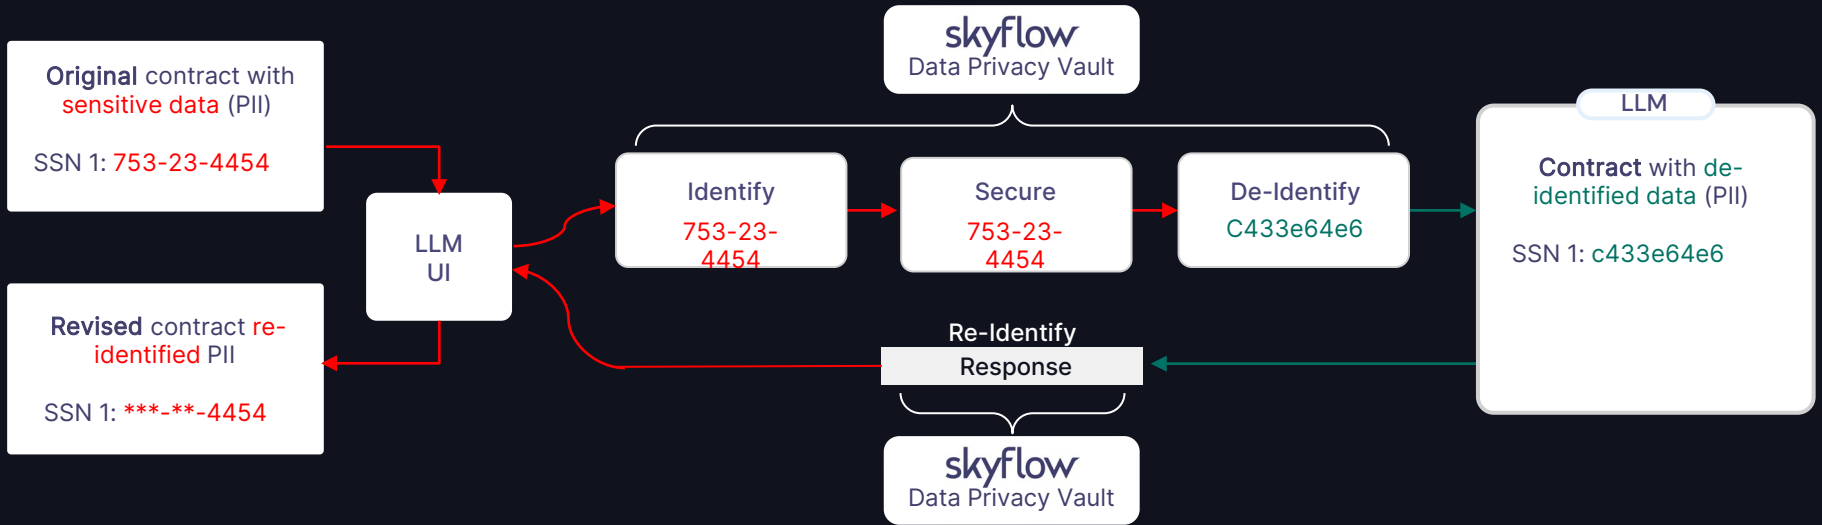
Need to keep PII and IP out of the model





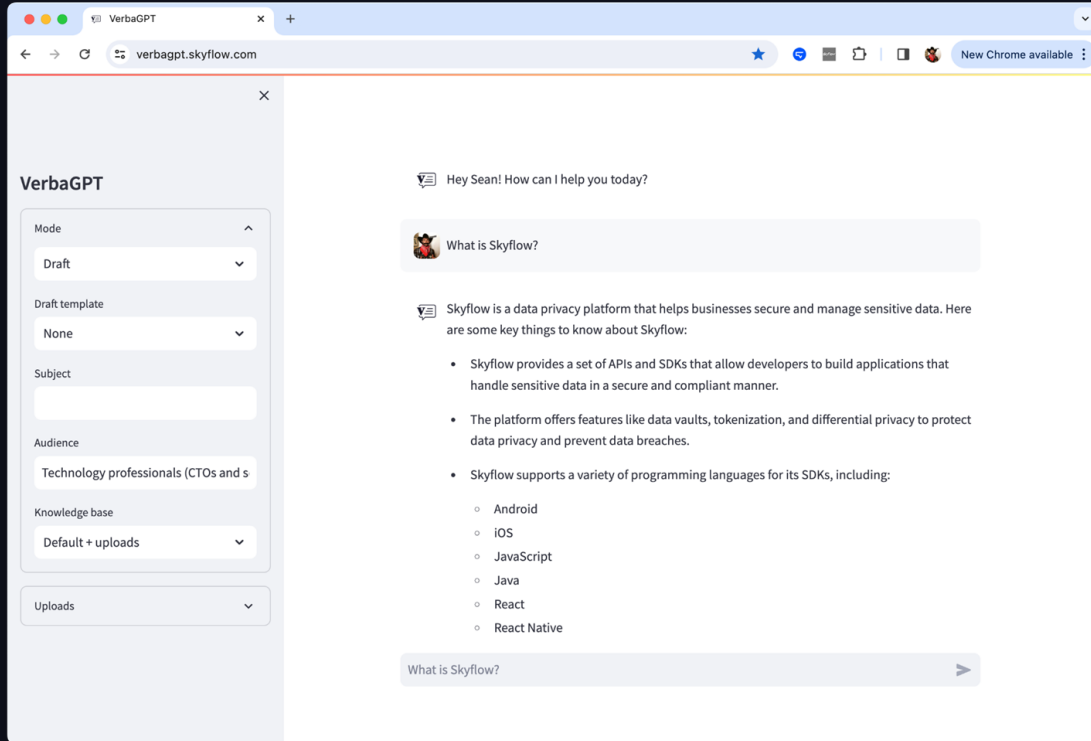
# SKYFLOW PROVIDES PRIVACY DURING INFERENCE

Keep PII and IP out of the inference process



# SOLUTION: VERBAGPT

# REQUIREMENTS FOR VERBAGPT



## Content Co-pilot

Anyone in Skyflow can use VerbaGPT to generate content, like a social media post to their personal LinkedIn account.

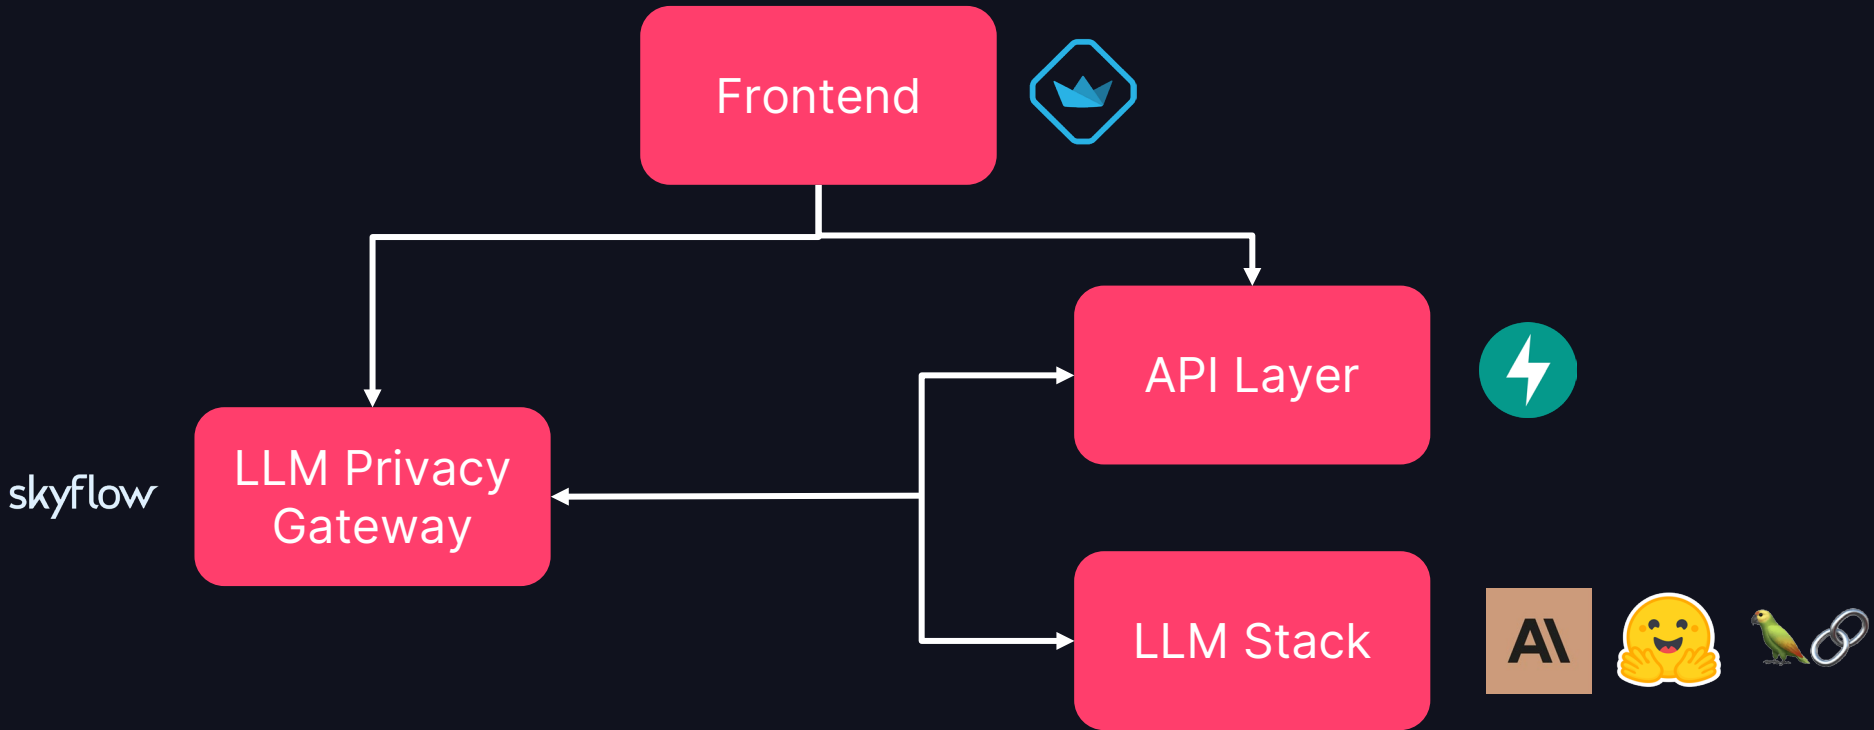
## Privacy-safe Inference

All PII and internal sensitive information must be de-identified before being shared with the LLM.

## Fine-grained Access Control


Sales team shouldn't have access to all product information but docs team needs full access to create docs ahead of launch.

# VERBAGPT TECH STACK



# WHAT IS RAG?

A process that combines elements of information retrieval and text generation to improve quality and relevance from LLMs



PALO ALTO, Calif. - March 28th, 2024 – Skyflow, the data privacy vault company, has raised \$30 million to extend its Series B funding. Khosla Ventures led the round, joined by prior investors Mouro Capital, Foundation Capital and Canvas Ventures.

# PRIVACY-SAFE EMBEDDINGS

## Prepping our content



### Data Sources

Documentation, Blogs,  
Whitepapers, PRDs,  
Contracts

# PRIVACY-SAFE EMBEDDINGS

Some documents have sensitive information



Data Sources  
Documentation, Blogs,  
Whitepapers, PRDs,  
Contracts



CONFIDENTIAL Page 5 of 6

### Invoicing Details for FinovateSpring 2024

We will use the below details to create and send your Demo Invoice. Demo Invoices will be sent through email to the below contact.

If you are **required** to have a PO# included on your invoice, a hard copy of your invoice mailed to the below address, etc., please include that information in the special instructions section below.

Company Name for Invoice: Skyflow

Company Address for Invoice: 195 Page Mill Road Suite 111, Palo Alto, CA 943056

Contact for Invoice (First and Last Name): Sean Falconer

Job Title for Invoice Contact: Head of Marketing

Email for Invoice Contact: sean.falconer@skyflow.com

Phone Number for Invoice Contact: 650-555-1212

Special Instructions for Invoice:



# PRIVACY-SAFE EMBEDDINGS

Some documents have sensitive information



Data Sources  
Documentation, Blogs,  
Whitepapers, PRDs,  
Contracts



CONFIDENTIAL Page 5 of 6

**Invoicing Details for FinovateSpring 2024**

We will use the below details to create and send your Demo Invoice. Demo Invoices will be sent through email to the below contact.

If you are **required** to have a PO# included on your invoice, a hard copy of your invoice mailed to the below address, etc., please include that information in the special instructions section below.

Company Name for Invoice: Skyflow

Company Address for Invoice: 195 Page Mill Road Suite 111, Palo Alto, CA 943056

Contact for Invoice (First and Last Name): Sean Falconer

Job Title for Invoice Contact: Head of Marketing

Email for Invoice Contact: sean.falconer@skyflow.com

Phone Number for Invoice Contact: 650-555-1212

Special Instructions for Invoice: \_\_\_\_\_


There's value in VerbaGPT knowing we will be at FinovateSpring, but it doesn't need my phone number.





# PRIVACY-SAFE EMBEDDINGS

## Detect and de-identify



CONFIDENTIAL

Page 5 of 6

### Invoicing Details for FinovateSpring 2024

We will use the below details to create and send your Demo Invoice. Demo Invoices will be sent through email to the below contact.

If you are **required** to have a PO# included on your invoice, a hard copy of your invoice mailed to the below address, etc., please include that information in the special instructions section below.

Company Name for Invoice: Skyflow

Company Address for Invoice: 195 Page Mill Road Suite 111, Palo Alto, CA 943056

Contact for Invoice (First and Last Name): NAME:98e02bfe-0461-4d7c-ae39-385a32485040

Job Title for Invoice Contact: Head of Marketing

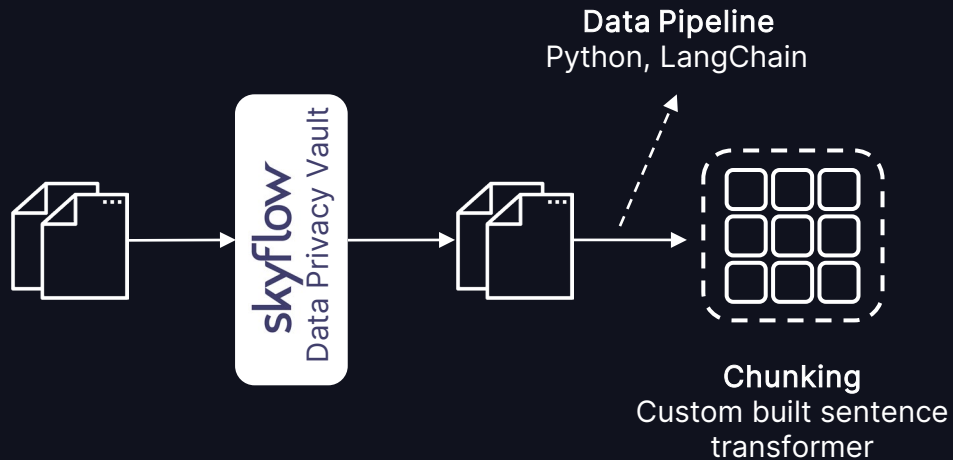
Email for Invoice Contact: EMAIL:dypqcz@wnxbw.com

Phone Number for Invoice Contact: PHONE:c0ce4473-6f26-460a-833b-cecc9cd9d5cb

Special Instructions for Invoice:

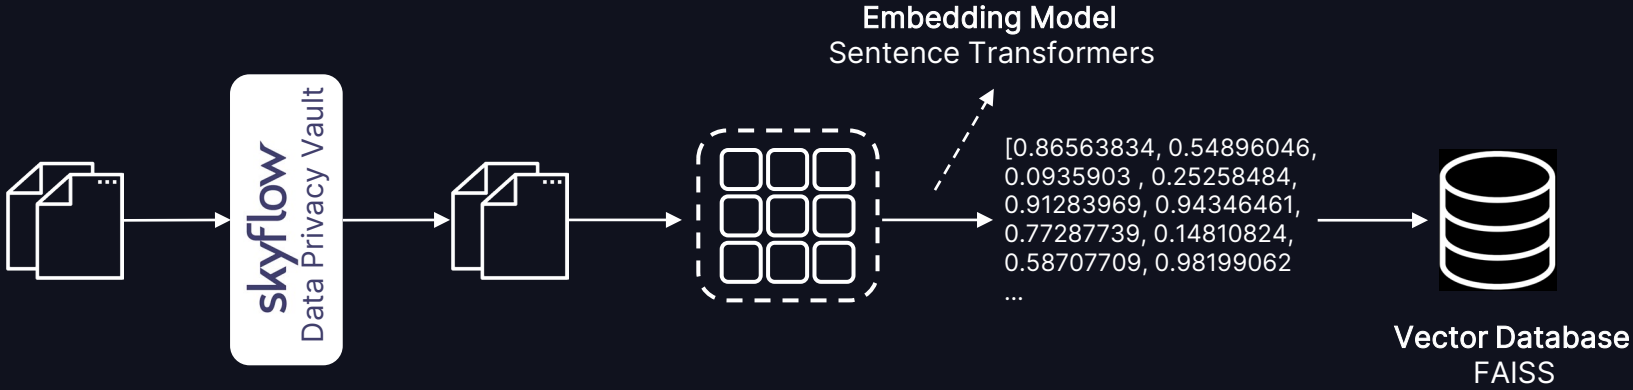
# PRIVACY-SAFE EMBEDDINGS

Some documents have sensitive information



# PRIVACY-SAFE EMBEDDINGS


Some documents have sensitive information



# PRIVACY-SAFE INFERENCE

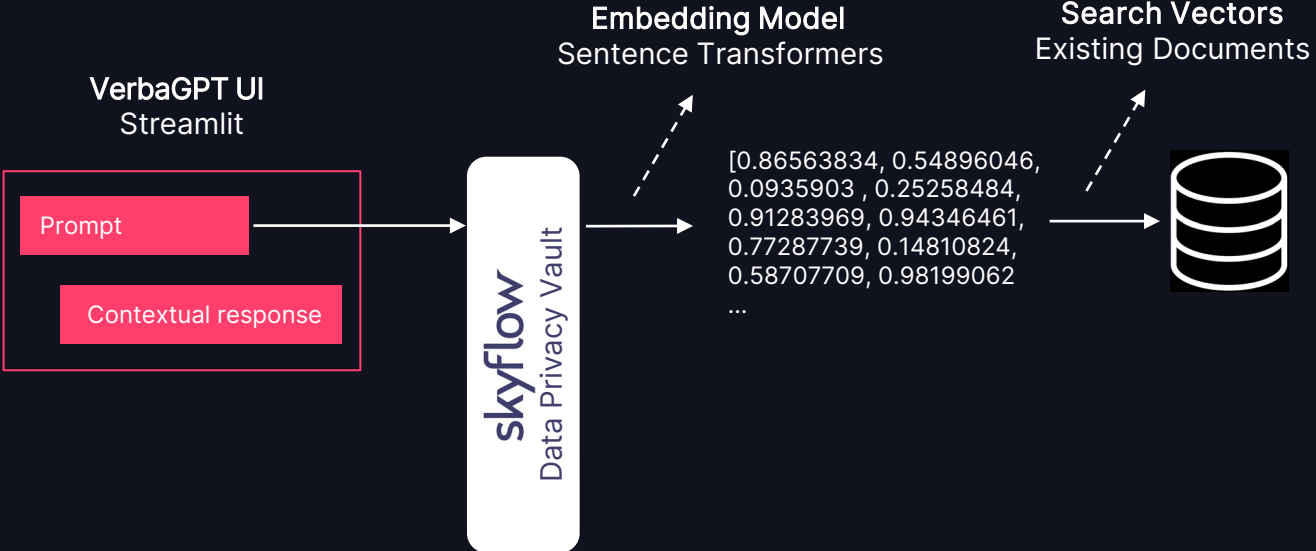
Secure and private

VerbaGPT UI  
Streamlit



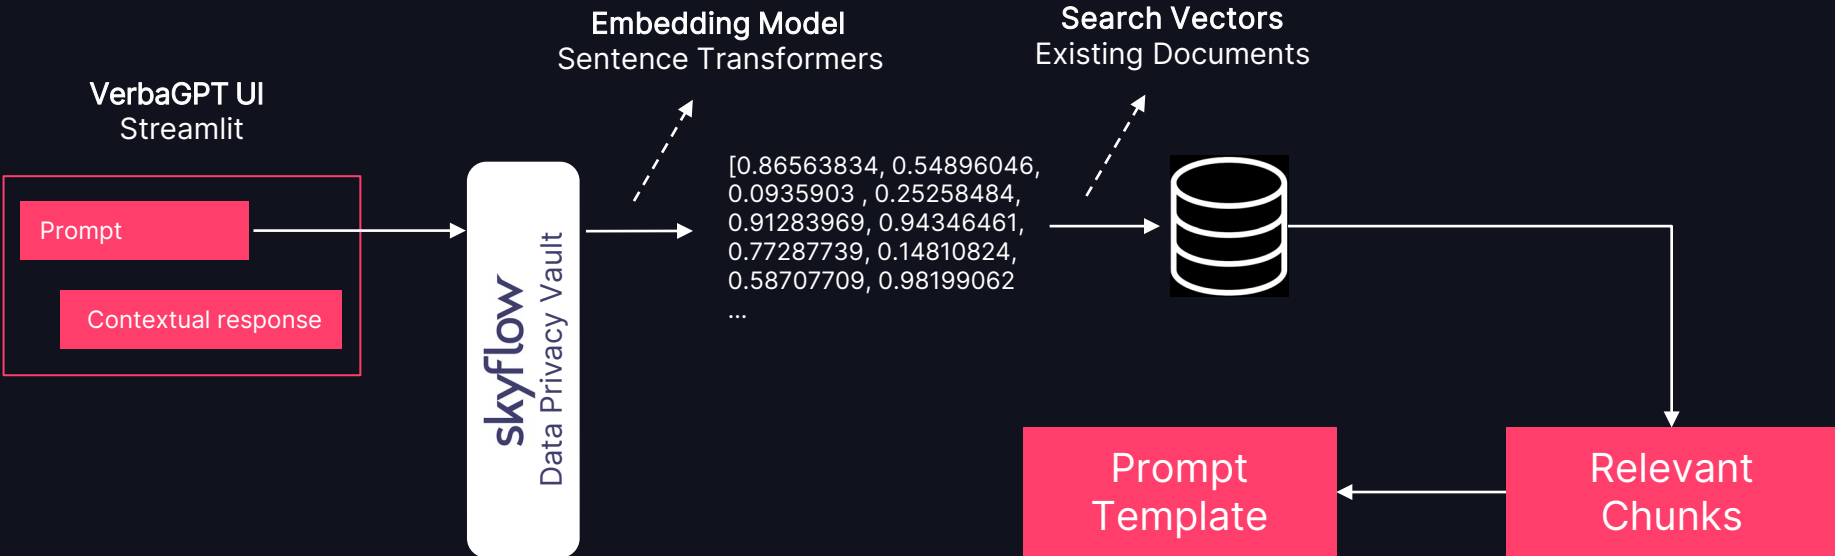
# PRIVACY-SAFE INFERENCE

Secure and private



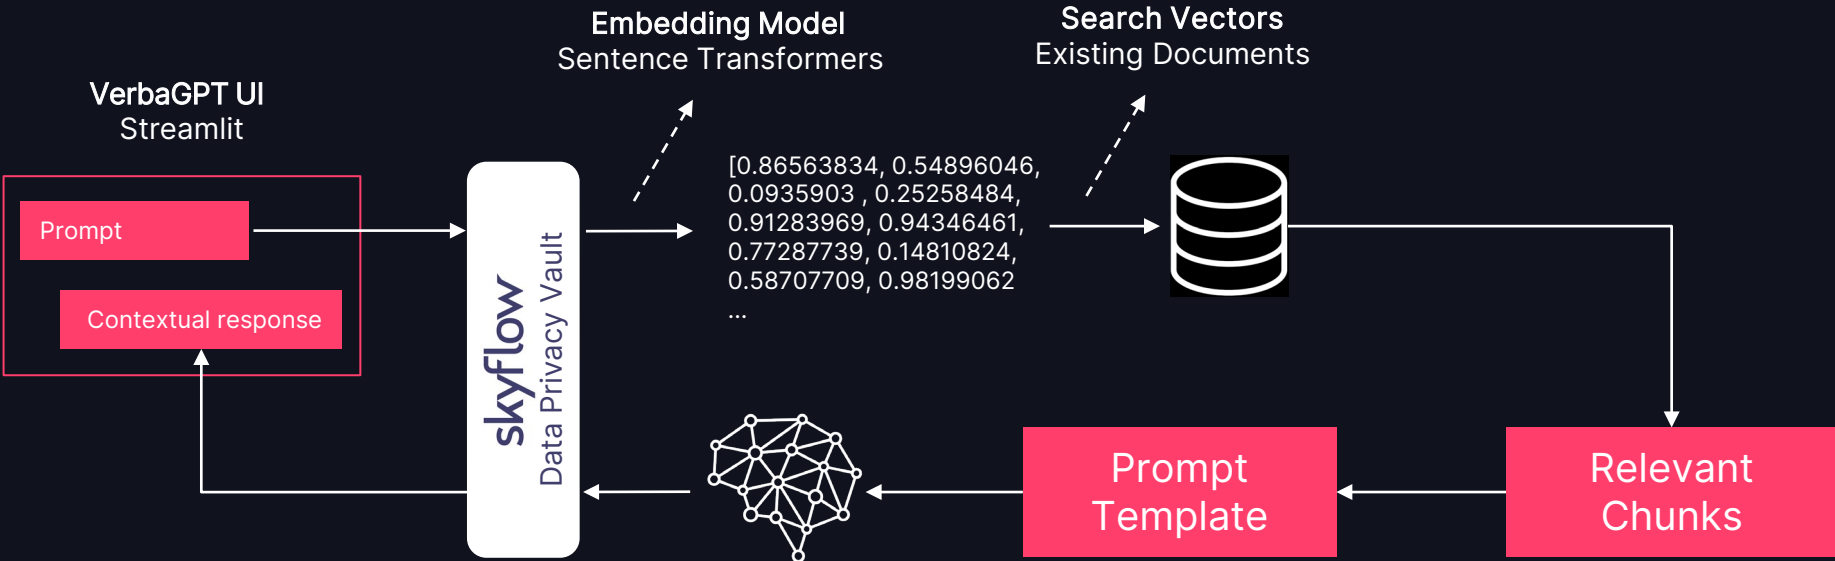
# PRIVACY-SAFE INFERENCE

Secure and private

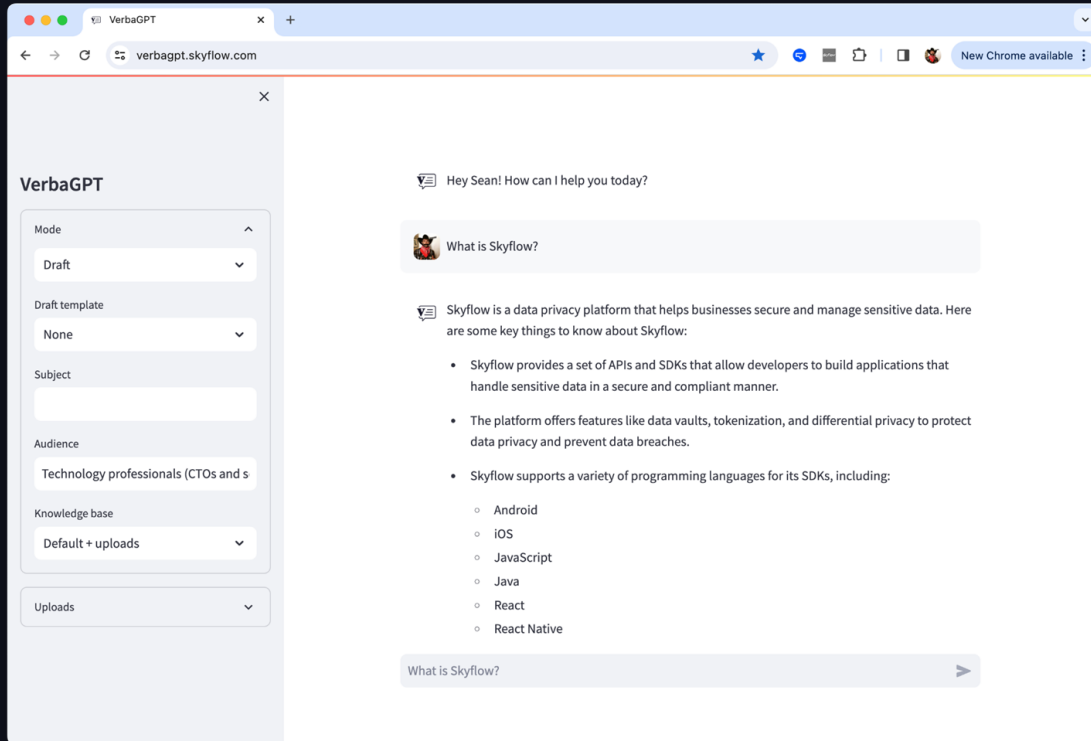


# PRIVACY-SAFE INFERENCE

Secure and private



# STREAMLIT FRONTEND



## Interface without API calls

Most folks don't want to worry about prompting. They just want it to work.

## Knowledgeable chat

VerbaGPT is informed about Skyflow through vector DB results. Outputs are grounded in fact.


## File uploads

Users can upload files to supplement the default vector DB contents. VerbaGPT can pull from both uploads and the vector DB.



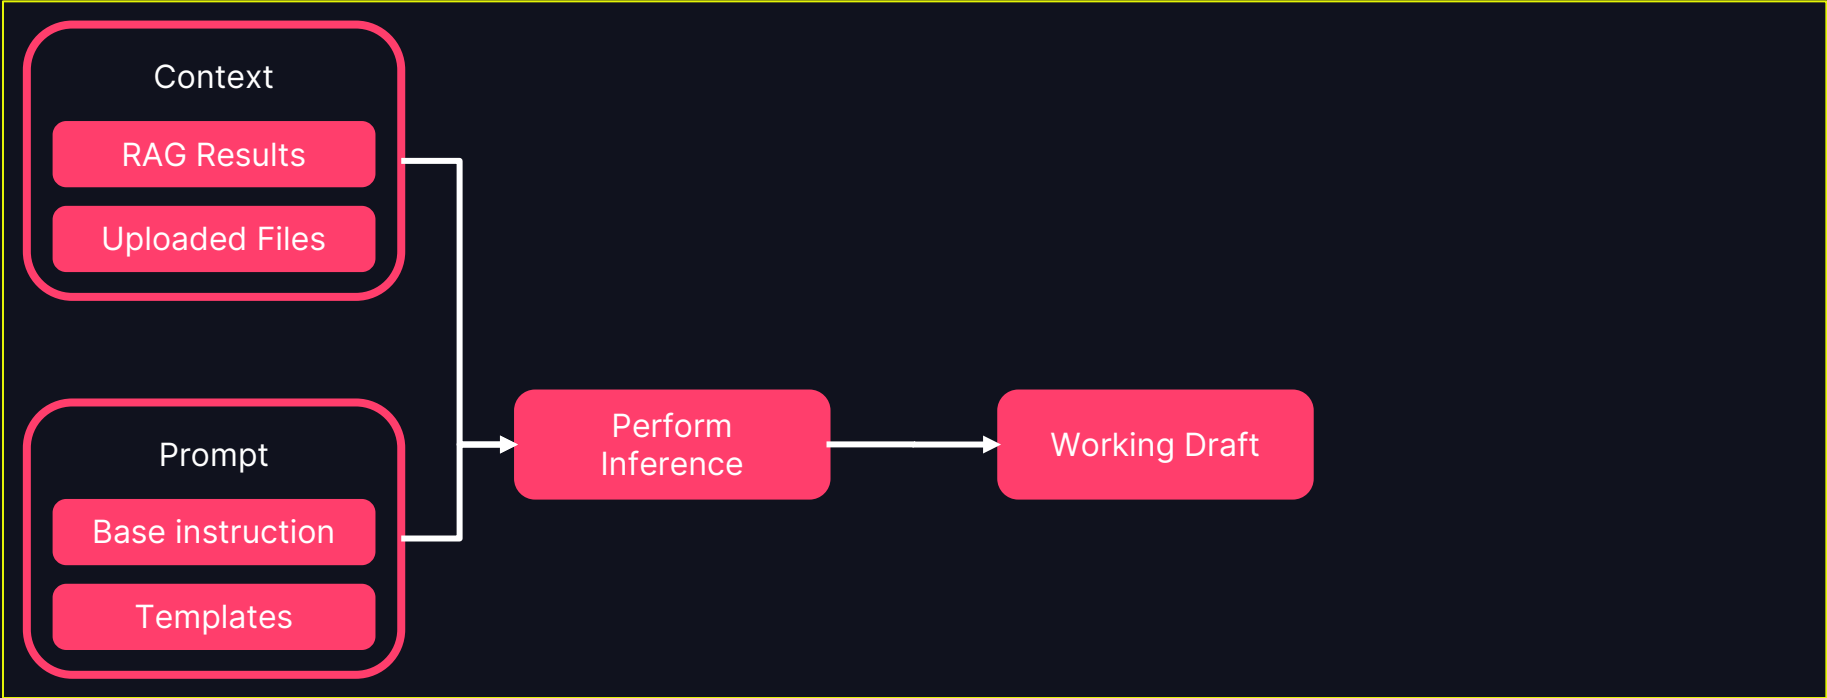
# CONTEXTUAL COMPOSITION

## Drafting with context and templates



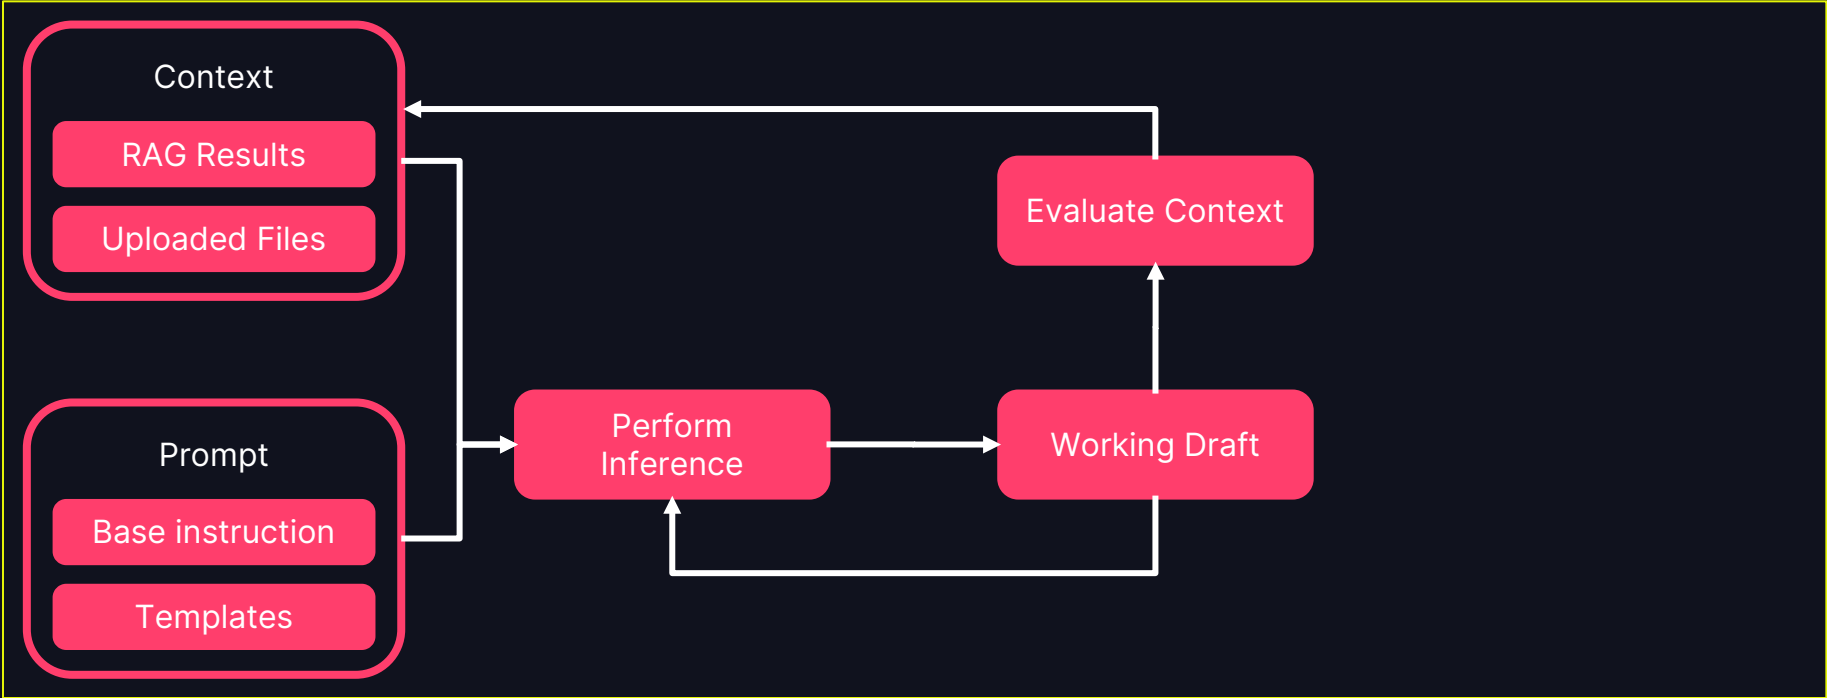
# CONTEXTUAL COMPOSITION

## Drafting with context and templates



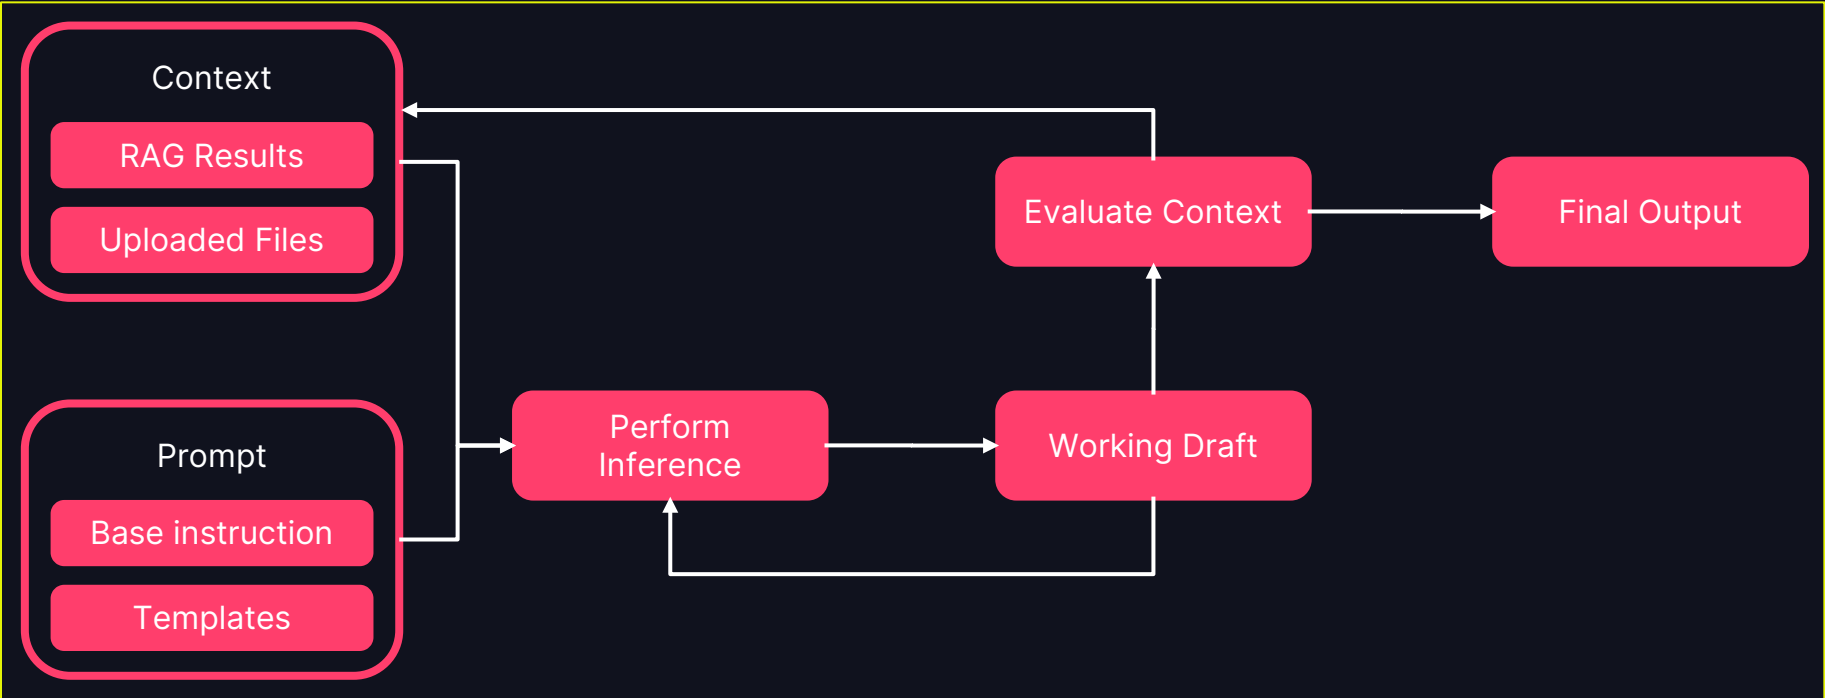
# CONTEXTUAL COMPOSITION

## Drafting with context and templates



# CONTEXTUAL COMPOSITION

## Drafting with context and templates



# RESULTS: SPEED AND ACCURACY

# TIME SAVINGS

Conceptual overviews

Use case guides

How-to guides

Tutorials



Blog posts

Social media posts

Press releases

RCAs

**10 minutes**  
instead of 2 weeks

# HOLDING TO OUR PRINCIPLES

Privacy



Security

DEMO:  
HERE WE GO!



# COMPOSING CONTENT WITH VERBAGPT

The screenshot shows a web browser window with the URL `verbagpt.skyflow.com`. On the left is a sidebar with the title "VerbaGPT" and several settings: "Mode" set to "Draft", "Draft template" set to "None", "Subject" (empty), "Audience" set to "Technology professionals (CTOs and s", "Knowledge base" set to "Default + uploads", and "Uploads" (empty). The main area is a chat interface. At the top, a message says "Hey Sean! How can I help you today?". Below that is a user prompt: "What is Skyflow?". The AI response follows: "Skyflow is a data privacy platform that helps businesses secure and manage sensitive data. Here are some key things to know about Skyflow:" followed by a bulleted list: "• Skyflow provides a set of APIs and SDKs that allow developers to build applications that handle sensitive data in a secure and compliant manner.", "• The platform offers features like data vaults, tokenization, and differential privacy to protect data privacy and prevent data breaches.", "• Skyflow supports a variety of programming languages for its SDKs, including: Android, iOS, JavaScript, Java, React, and React Native." At the bottom of the chat area is a text input field containing "What is Skyflow?" and a send button.



# THANK YOU

