

# Structure your unstructured content:

Gemini for scaling business

Google  
Cloud  
Next 25

Proprietary





**Dr. Ali  
Arsanjani**

Director, Applied AI  
Engineering,  
Google Cloud



**Yashodha  
Bhavnani**

VP, Product Management,  
Box



**Adam  
Anzuoni**

Founder Simple Stack  
Former CTO,  
Fundwell



# 01. Google



# AI Trends Past, Present & Future





# State-of-the-art models


Pioneering research &  
development in AI






## Current AI - State of the Art

- Pattern recognition,
- niche expertise,
- multimodal processing,
- domain-specific excellence,
- efficiency gains.



## Intermediate AI - Thinking Models

- Enhanced planning, reasoning,
- creative problem-solving,
- sophisticated world models,
- improved robustness.



## AI Agents

- Strong planning & reasoning,
- sophisticated world models,
- active task execution,
- autonomous operation,
- negotiation with other agents.



## Artificial General Intelligence

- Human-level cognition, creativity,
- complex reasoning,
- accurate world models,
- long-term memory.



## Artificial Super Intelligence

- Surpassing human cognition,
- revolutionary scientific breakthroughs,
- world-reshaping potential,
- extending biological limits.

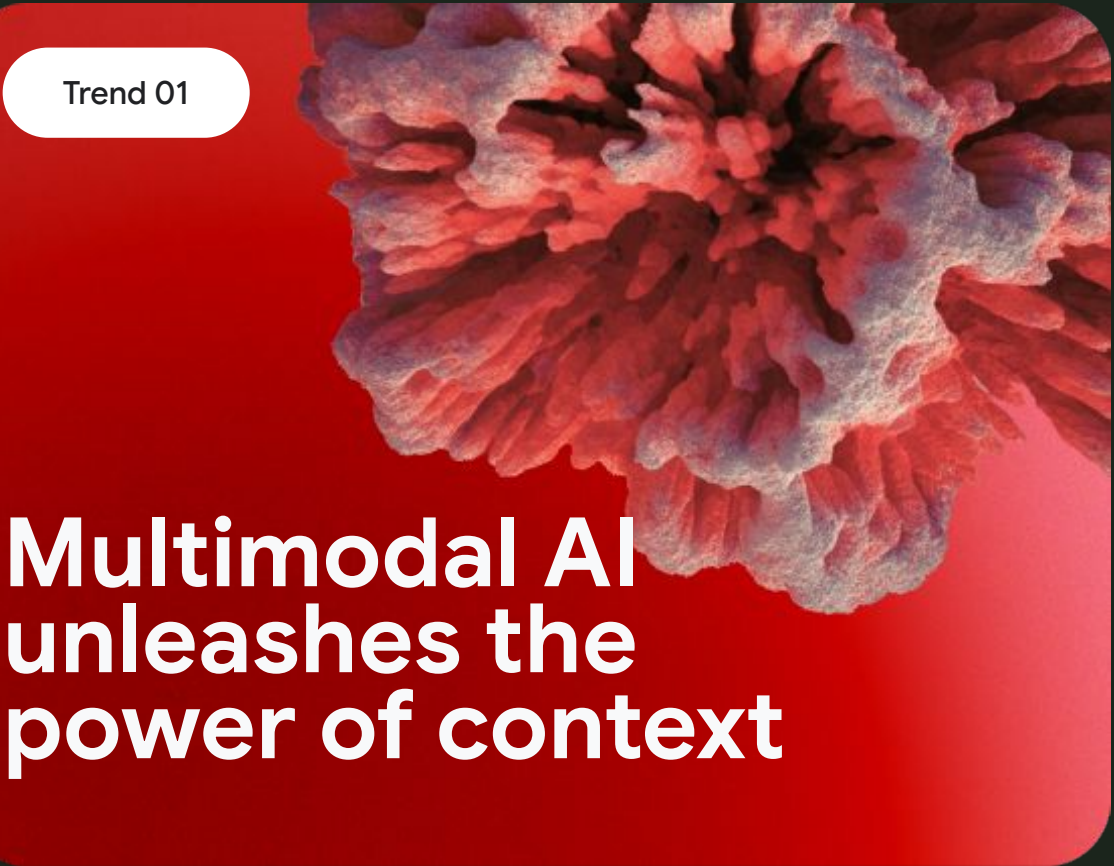
1. Scaling alone is insufficient for AGI; breakthroughs are required.
2. Safety, deception prevention, and ethical development are critical.
3. AI agents are pivotal for the transition from current AI to AGI and then to ASI.





# Top 5 trends at a glance:

Trend 01




Multimodal AI unleashes the power of context

Trend 02




AI agents evolve: From chatbots to multi-agent systems

Trend 03



Assistive search: The next frontier for knowledge work

Trend 04



AI-powered customer experience so seamless, it's almost invisible

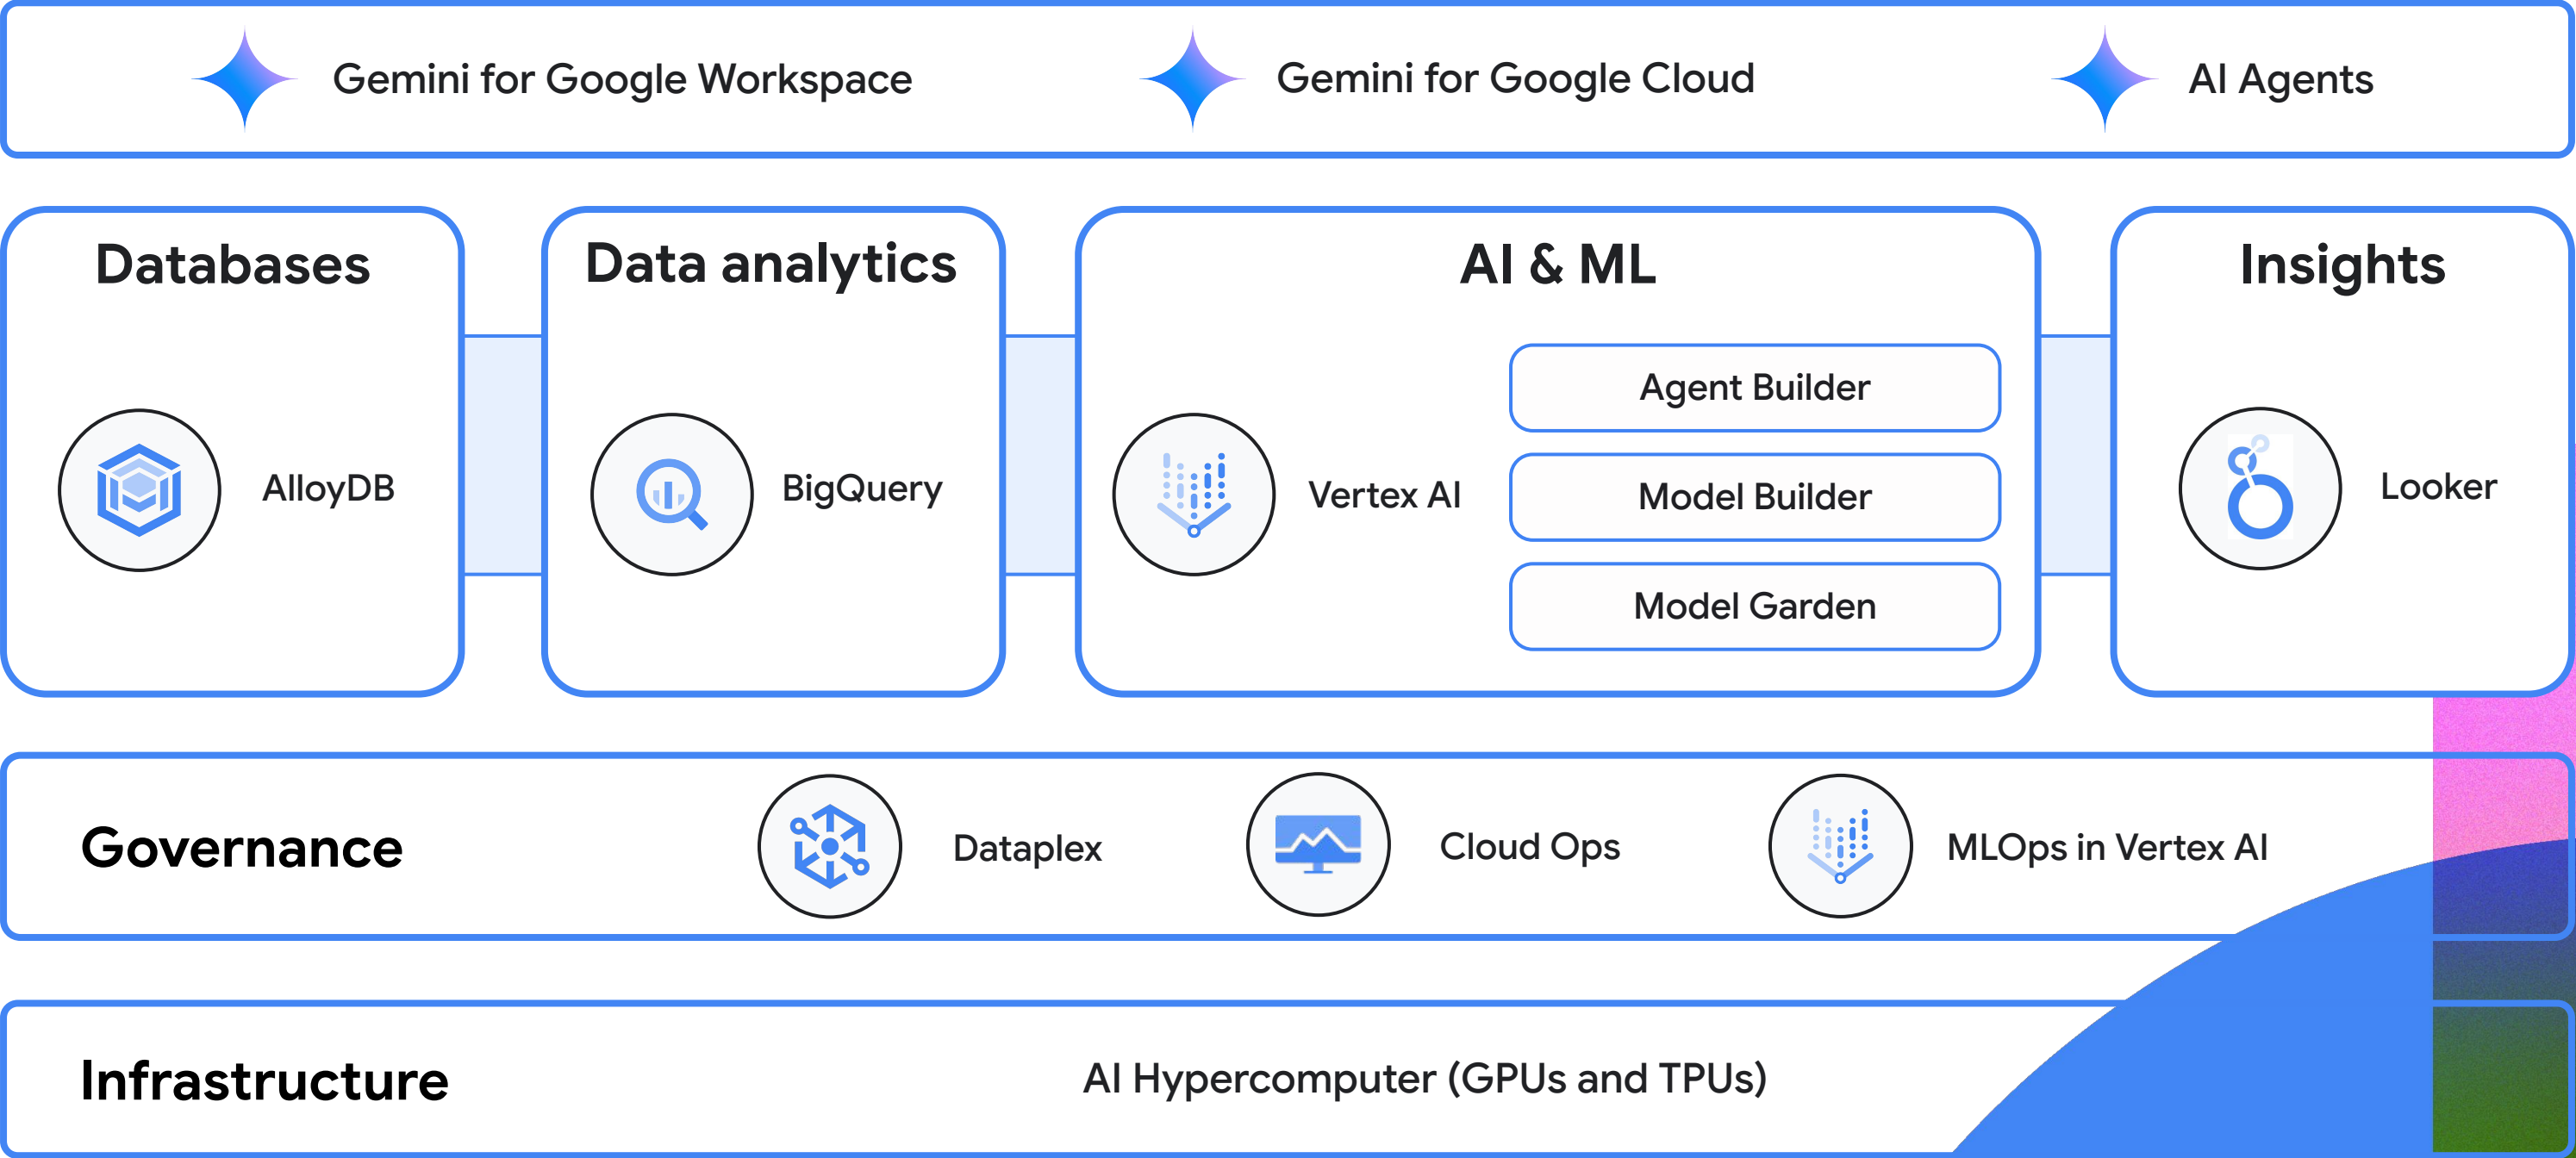
Trend 05



Security gets tighter—and tougher—with AI



# A unified platform from data to deployment and for all your predictive, generative, and agentic needs

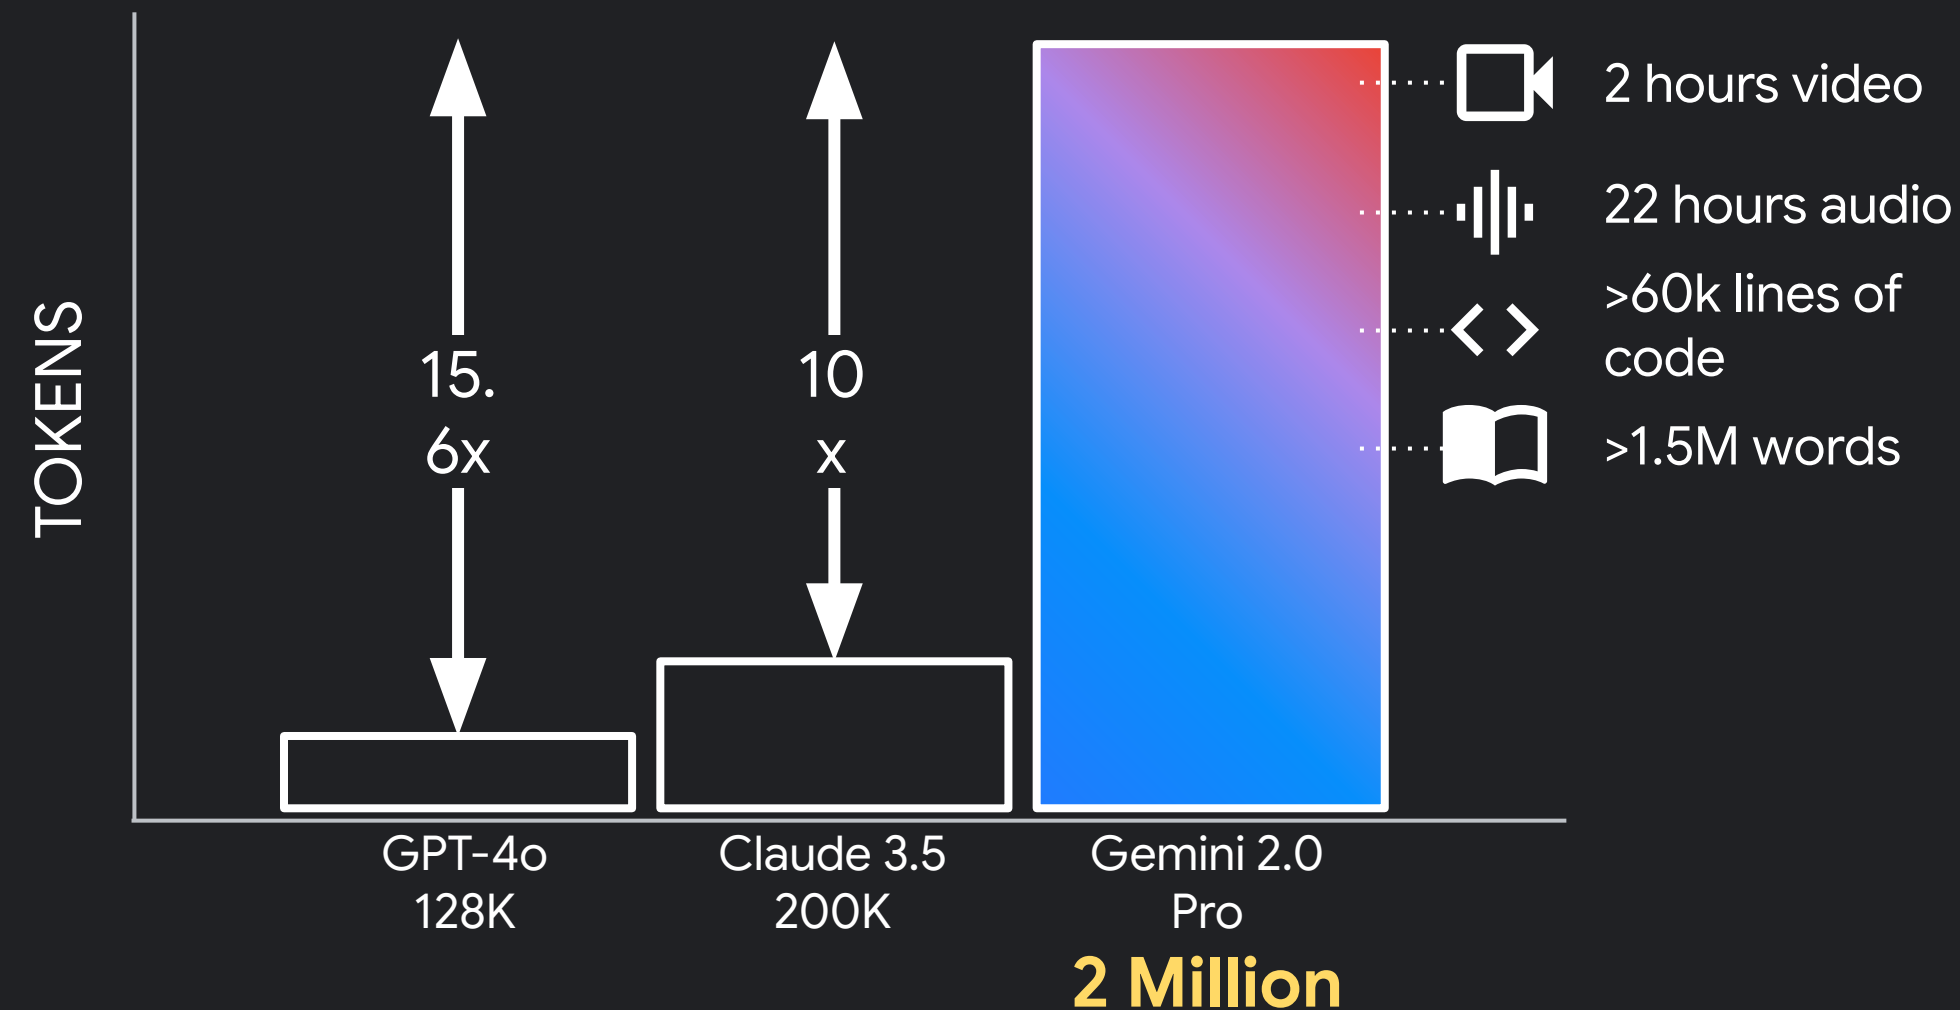




# Gemini 2.0: Unmatched Long Context Capabilities

Understand more content  
in one shot and extract  
more insights

## Long Context Capabilities of Leading LLMS





# Vertex AI & Agentspace

## Best models from Google & the industry

- Gemini 2.5 Pro <sup>Preview</sup>
- Gemini 2.5 Flash [coming soon]
- Live API <sup>Preview</sup>
- Veo 2.0 Text-to-video <sup>GA</sup> | Image-to-video <sup>GA</sup> | Editing <sup>Experimental</sup>
- Lyria <sup>Preview</sup>
- Chirp 3 HD Voices <sup>GA</sup> | Instant Custom Voice <sup>GA</sup> | Diarization <sup>Preview</sup>

## End-to-end model building platform with choice at every level

- Vertex AI Model Optimizer <sup>Experimental</sup>
- GenAI Evaluation <sup>Batch | Multimodal | Rubric Based</sup> <sup>Preview</sup>
- Vertex AI Model Development Service <sup>Experimental</sup>
- Global Endpoint <sup>Preview</sup>

## Develop and deploy agents faster, grounded in your enterprise truth

- Agent Development Kit (ADK) <sup>OSS</sup>
- Agent Engine <sup>GA</sup>
- Agent Garden <sup>OSS</sup>
- Tools: Apigee API Hub & BAP, MCP, LangChain, LlamaIndex, etc

## Search and agents over your enterprise systems with Agentspace

- Agent Gallery <sup>GA</sup>
- Deep Research Agent <sup>GA</sup>
- Idea Generation Agent <sup>Preview</sup>
- Agent Designer <sup>Preview</sup>
- Video Generation <sup>Preview</sup>
- Agentspace + Chrome <sup>Preview</sup>
- Agent + NotebookLM <sup>GA</sup>
- Support for 3p Agents <sup>Preview</sup>

- WeatherNext <sup>Available</sup>
- Gemini Embeddings <sup>Experimental</sup>
- Partner Models | AI21, Camb.AI, Qodo, Allen Institute <sup>Available</sup>
- Llama 4 <sup>Available</sup>

- Tuning for Gemini 2.0 <sup>GA</sup>
- Llama OSS Full Finetuning <sup>Preview</sup>

- Agent Engine Services: Session <sup>GA</sup> | Example Store <sup>GA</sup> | Memory <sup>Preview</sup>
- Grounding with Google Maps <sup>Experimental</sup> | Web Grounding for Enterprise <sup>GA</sup>
- Podcast Generation API <sup>Preview</sup>

- Conversational Commerce <sup>Preview</sup> <sup>Preview</sup> Merchandising Console <sup>GA</sup>
- Multimodal Video Search <sup>Preview</sup> <sup>Preview</sup>
- Personalized Browse (Site) <sup>Preview</sup>
- T2: People Search <sup>GA</sup>



# Gemini: What is new

Coming Soon	Preview
Gemini 2.5 Flash	Live API
Preview	Experimental
Gemini 2.5 Pro	Model Optimizer

## Introducing Gemini 2.5, our most intelligent AI model yet

Gemini 2.5 are "thinking models" with advanced reasoning, now on Vertex AI. They show thoughts step-by-step for transparency, tackling complex problems and enabling smarter, context-aware agents with optimized performance.

Note: **GA (Generally Available)**, **Preview**, and **Pre-announcement** refer to the launch stage of the product.


**Gemini 2.5** enables developers and enterprises across industries (including regulated industries like Healthcare, Finance, Tech, Security) to **build more capable, transparent, and cost-effective AI applications faster on a unified platform.**

## Vertex AI is expanding Gemini capabilities:

- **Gemini 2.5 Pro Preview:** SOTA model for complex reasoning, coding, math now available on Vertex AI.
- **Gemini 2.5 Flash Preview:** Fast, efficient model optimized for speed/cost, with configurable thinking budget now in Preview on Vertex AI.
- **Model Optimizer:** Automatically routes queries to optimal Gemini model, balancing quality, speed, and cost.
- **Live API:** Production-ready API for real-time, interactive audio experiences with enhanced features.








# Gemini 2.5

Pro (Preview)

## Reasoning & knowledge

Humanity's Last Exam


No tools



All numbers are sourced from providers' self-reported numbers. For missing values: Humanity's Last Exam results are sourced from [https://huggingface.co/datasets/GoogleResearch/humanity\\_last\\_exam](https://huggingface.co/datasets/GoogleResearch/humanity_last_exam) / AIME 2025 numbers are sourced from <https://mathematics.ai/>. All results show pass@1 performance with AIME 2025 results being averaged over 4 runs. Groved numbers use majority voting with m44 for Grok models and use internal scoring with parallel test time compute for Anthropic models.

## Science

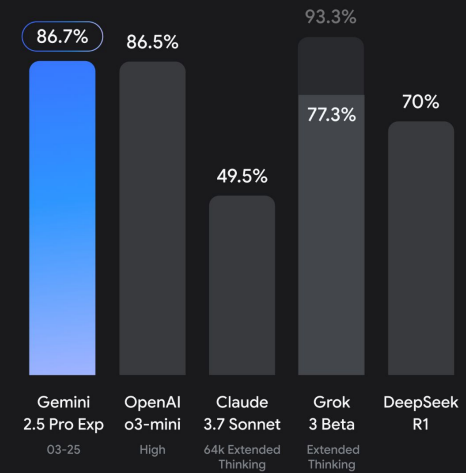
GPQA diamond



All numbers are sourced from providers' self-reported numbers. For missing values: Humanity's Last Exam results are sourced from [https://huggingface.co/datasets/GoogleResearch/humanity\\_last\\_exam](https://huggingface.co/datasets/GoogleResearch/humanity_last_exam) / AIME 2025 numbers are sourced from <https://mathematics.ai/>. All results show pass@1 performance with AIME 2025 results being averaged over 4 runs. Groved numbers use majority voting with m44 for Grok models and use internal scoring with parallel test time compute for Anthropic models.

## Mathematics


AIME 2025



All numbers are sourced from providers' self-reported numbers. For missing values: Humanity's Last Exam results are sourced from [https://huggingface.co/datasets/GoogleResearch/humanity\\_last\\_exam](https://huggingface.co/datasets/GoogleResearch/humanity_last_exam) / AIME 2025 numbers are sourced from <https://mathematics.ai/>. All results show pass@1 performance with AIME 2025 results being averaged over 4 runs. Groved numbers use majority voting with m44 for Grok models and use internal scoring with parallel test time compute for Anthropic models.

## Code editing


Aider Polyglot



All numbers are sourced from providers' self-reported numbers. For missing values: Humanity's Last Exam results are sourced from [https://huggingface.co/datasets/GoogleResearch/humanity\\_last\\_exam](https://huggingface.co/datasets/GoogleResearch/humanity_last_exam) / AIME 2025 numbers are sourced from <https://mathematics.ai/>. All results show pass@1 performance with AIME 2025 results being averaged over 4 runs. Groved numbers use majority voting with m44 for Grok models and use internal scoring with parallel test time compute for Anthropic models.

## Agentic coding

SWE-bench verified




All numbers are sourced from providers' self-reported numbers. For missing values: Humanity's Last Exam results are sourced from [https://huggingface.co/datasets/GoogleResearch/humanity\\_last\\_exam](https://huggingface.co/datasets/GoogleResearch/humanity_last_exam) / AIME 2025 numbers are sourced from <https://mathematics.ai/>. All results show pass@1 performance with AIME 2025 results being averaged over 4 runs. Groved numbers use majority voting with m44 for Grok models and use internal scoring with parallel test time compute for Anthropic models.



# Gemini 2.5 Pro

- **State of the art performance** across a range of tasks
- Highly **advanced coding** capabilities
- **Unmatched long context** understanding up to 1M tokens (2M coming soon)
- **Built in reasoning capabilities**
- **Improved multimodal understand** across images, video and audio
- **Increased speed**
- **Native tool use** like Google Search Grounding.





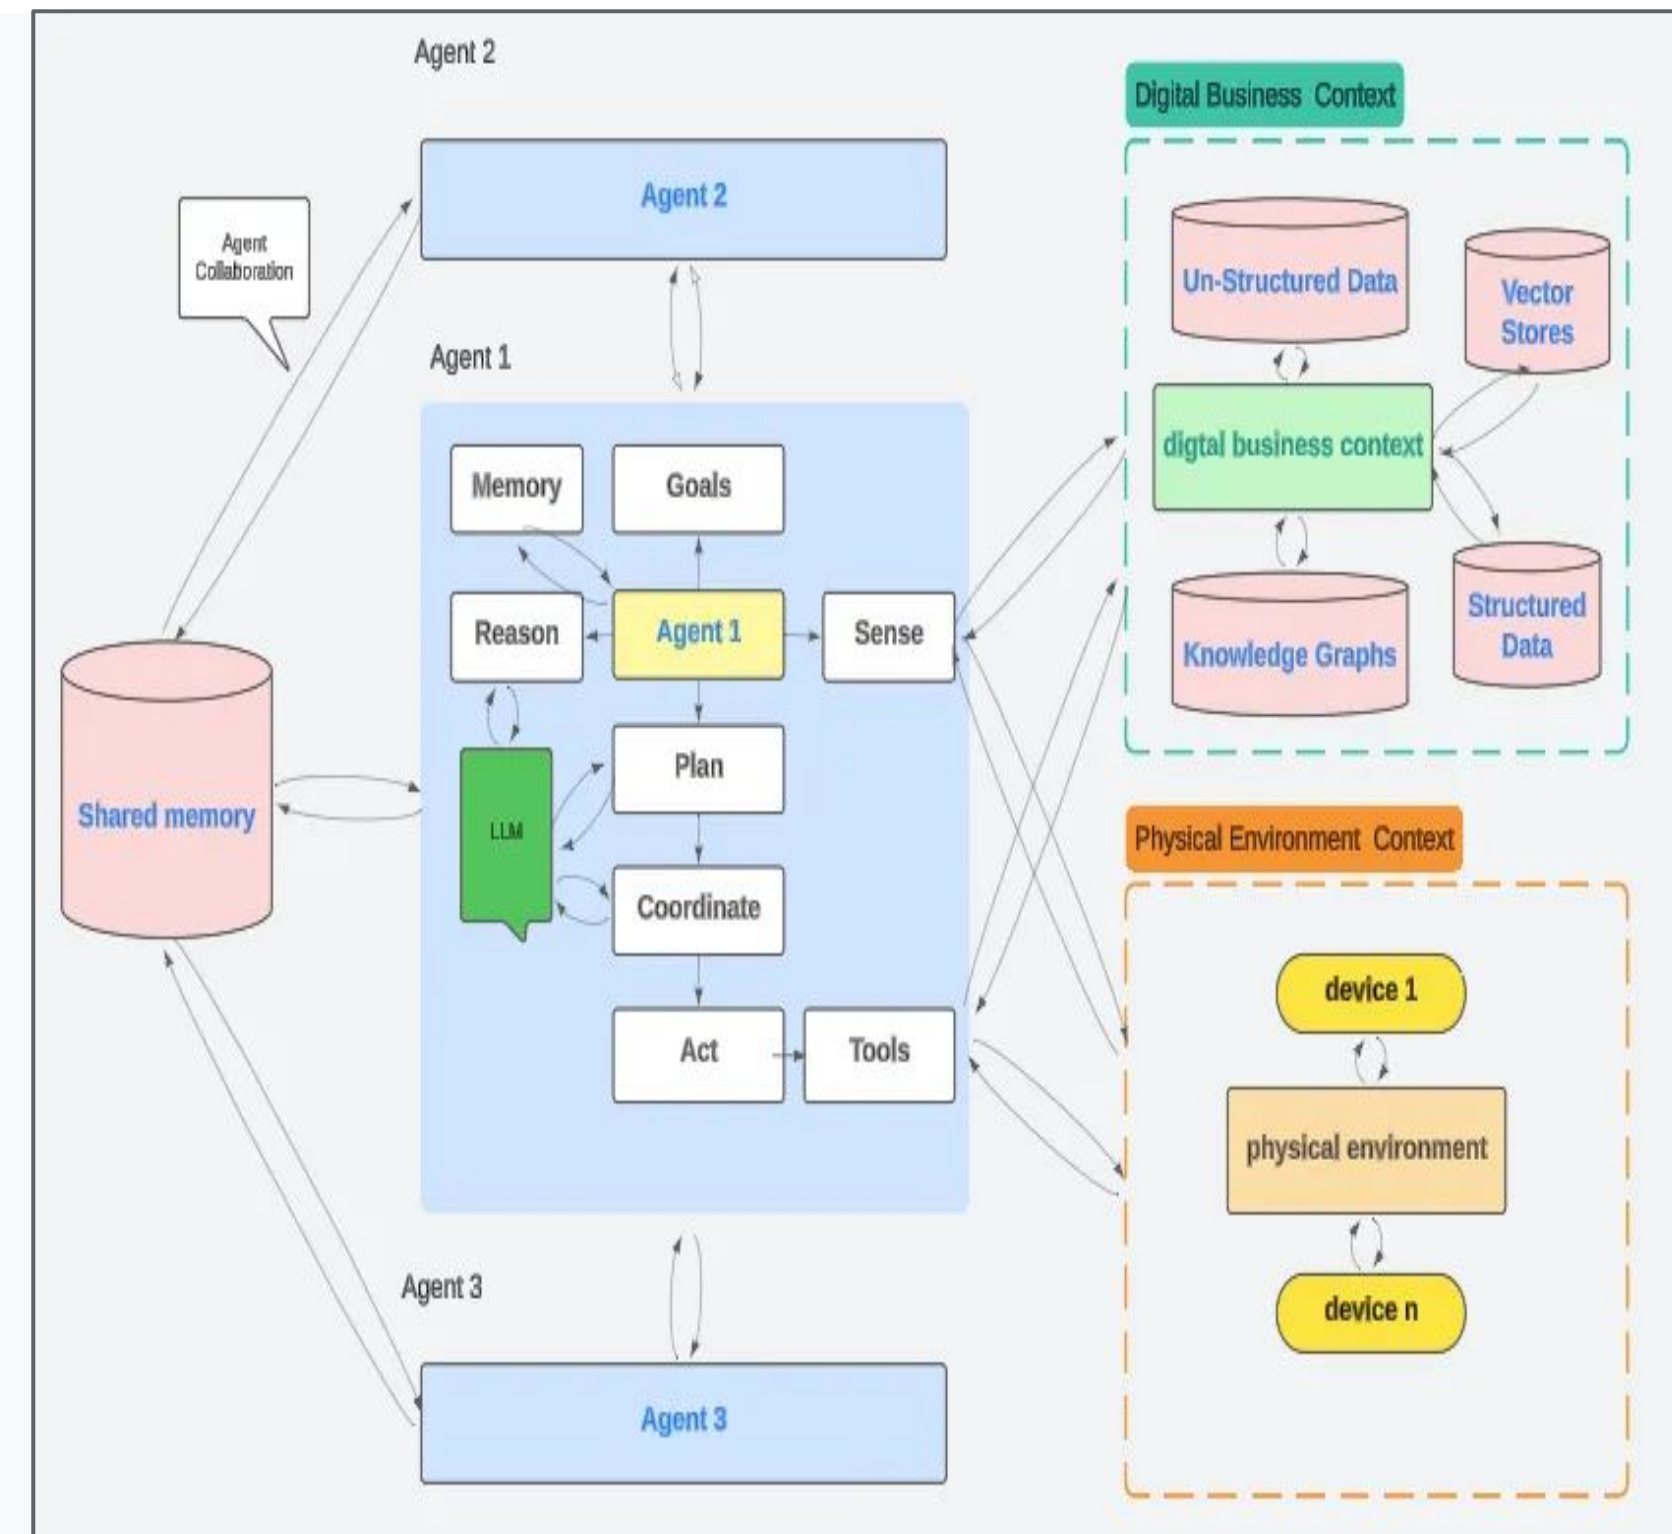
# Whats is new with Gemini 2.5

	What is it	Why it matters
Gemini 2.5 Pro (Preview)	<ul style="list-style-type: none"><li>Our latest Flagship Model with built in thinking capabilities</li></ul>	<ul style="list-style-type: none"><li>Is industry leading in a wide range of tasks</li><li>Significantly improved long context understanding</li><li>Highly price competitive</li><li>Improved Speed</li></ul>
Gemini 2.5 Flash (Coming soon)	<ul style="list-style-type: none"><li>Our latest workhorse model with configurable thinking capabilities</li></ul>	<ul style="list-style-type: none"><li>Improved quality over Gemini 2.0 Flash while retaining speed and value</li><li>Ability to dynamically think on complex tasks</li><li>Configurable Thinking Budget</li></ul>
Thinking Budget	Ability to control and configure how much a model thinks on a given user prompt. Available on 2.5 Flash, Coming soon to 2.5 Pro	<ul style="list-style-type: none"><li>Allow the model to dynamically think for a task or select how many tokens to use for reasoning for certain tasks</li><li>Allows users to control quality and speed of response</li><li>Setting budget to 0 turns off thinking and turns Flash into a non-thinking model for simpler tasks.</li></ul>
Model Optimizer	Dynamically route between Flash and Pro models depending on the complicity of the users ask.	<ul style="list-style-type: none"><li>Allow customers to maximize quality for a fixed budget</li><li>Removes the need for selecting a single model for a specific use case</li><li>Sers can configure if the want to optimize performance or cost</li></ul>



# Agent Anatomy

- **Goals:** *Agent Objectives* dynamically updated based on feedback and internal states.
- **Sense:** *gather information* from the environment, perceiving both digital and physical data.
- **Reason:** *analyze sensed information* using knowledge. Leverage LLMs for complex reasoning and inference.
- **Plan:** ***Devise a course of action*** based on reasoned insights, formulating a sequence of actions to achieve goals.
- **Coordinate:** Interact with other agents using shared memory to *align actions and ensure collaborative efforts*.
- **Act:** *Execute planned actions* on the environment, manipulating physical devices interacting with digital systems.
- **Memory/Shared Mem:** Store individual knowledge, past experiences, and belief states, providing *context for informed decision-making*.
- **LLM :** A core component of the "Reason" function, responsible for processing and understanding language-based information.





AI Applications For business users

Agentspace

Use solutions in Agent Gallery for employee productivity, chat and search.

Agent Designer

No code agent builder for every day users to create and use AI agents.

Customer Engagement Suite

Customer service voice or chat agents to assist human agents with low code, fully managed software.

Vertex AI Agent Builder For developers

Enterprise, data and tools

Grounding & RAG

Google Search on web, maps, and on your data

Data for Agents

Toolbox for databases & BigQuery Data agents

Enterprise Context

100s of connectors, millions of custom APIs, knowledge graph, agent plan & workflows

Ecosystem tools

MCP, A2A, LangChain, LangGraph, LlamaIndex, CrewAI, etc

Discover, build, and deploy agents



Agent Development Kit (ADK)

Client side SDK to define multi-agent applications for complex, real world scenarios

Tool and model use

Orchestration



Agent Engine

Fully-managed runtime to deploy and manage agents in production

Monitor, log, trace, ops

Sessions & short term memory

Long term memory bank

Evaluation and example store

Code execution & agent simulation

Identity mgt, security, scaling



Agent Garden

Samples and tools to accelerate agent development

Models



Gemini API

Gemini Models

Optimized for agentic reasoning

Audio / video streaming

Native tool use

Model Garden

Hundreds of curated LLMs



# 02. Box



Our mission is to power how  
the world works together





# We're entering an AI-first era of business

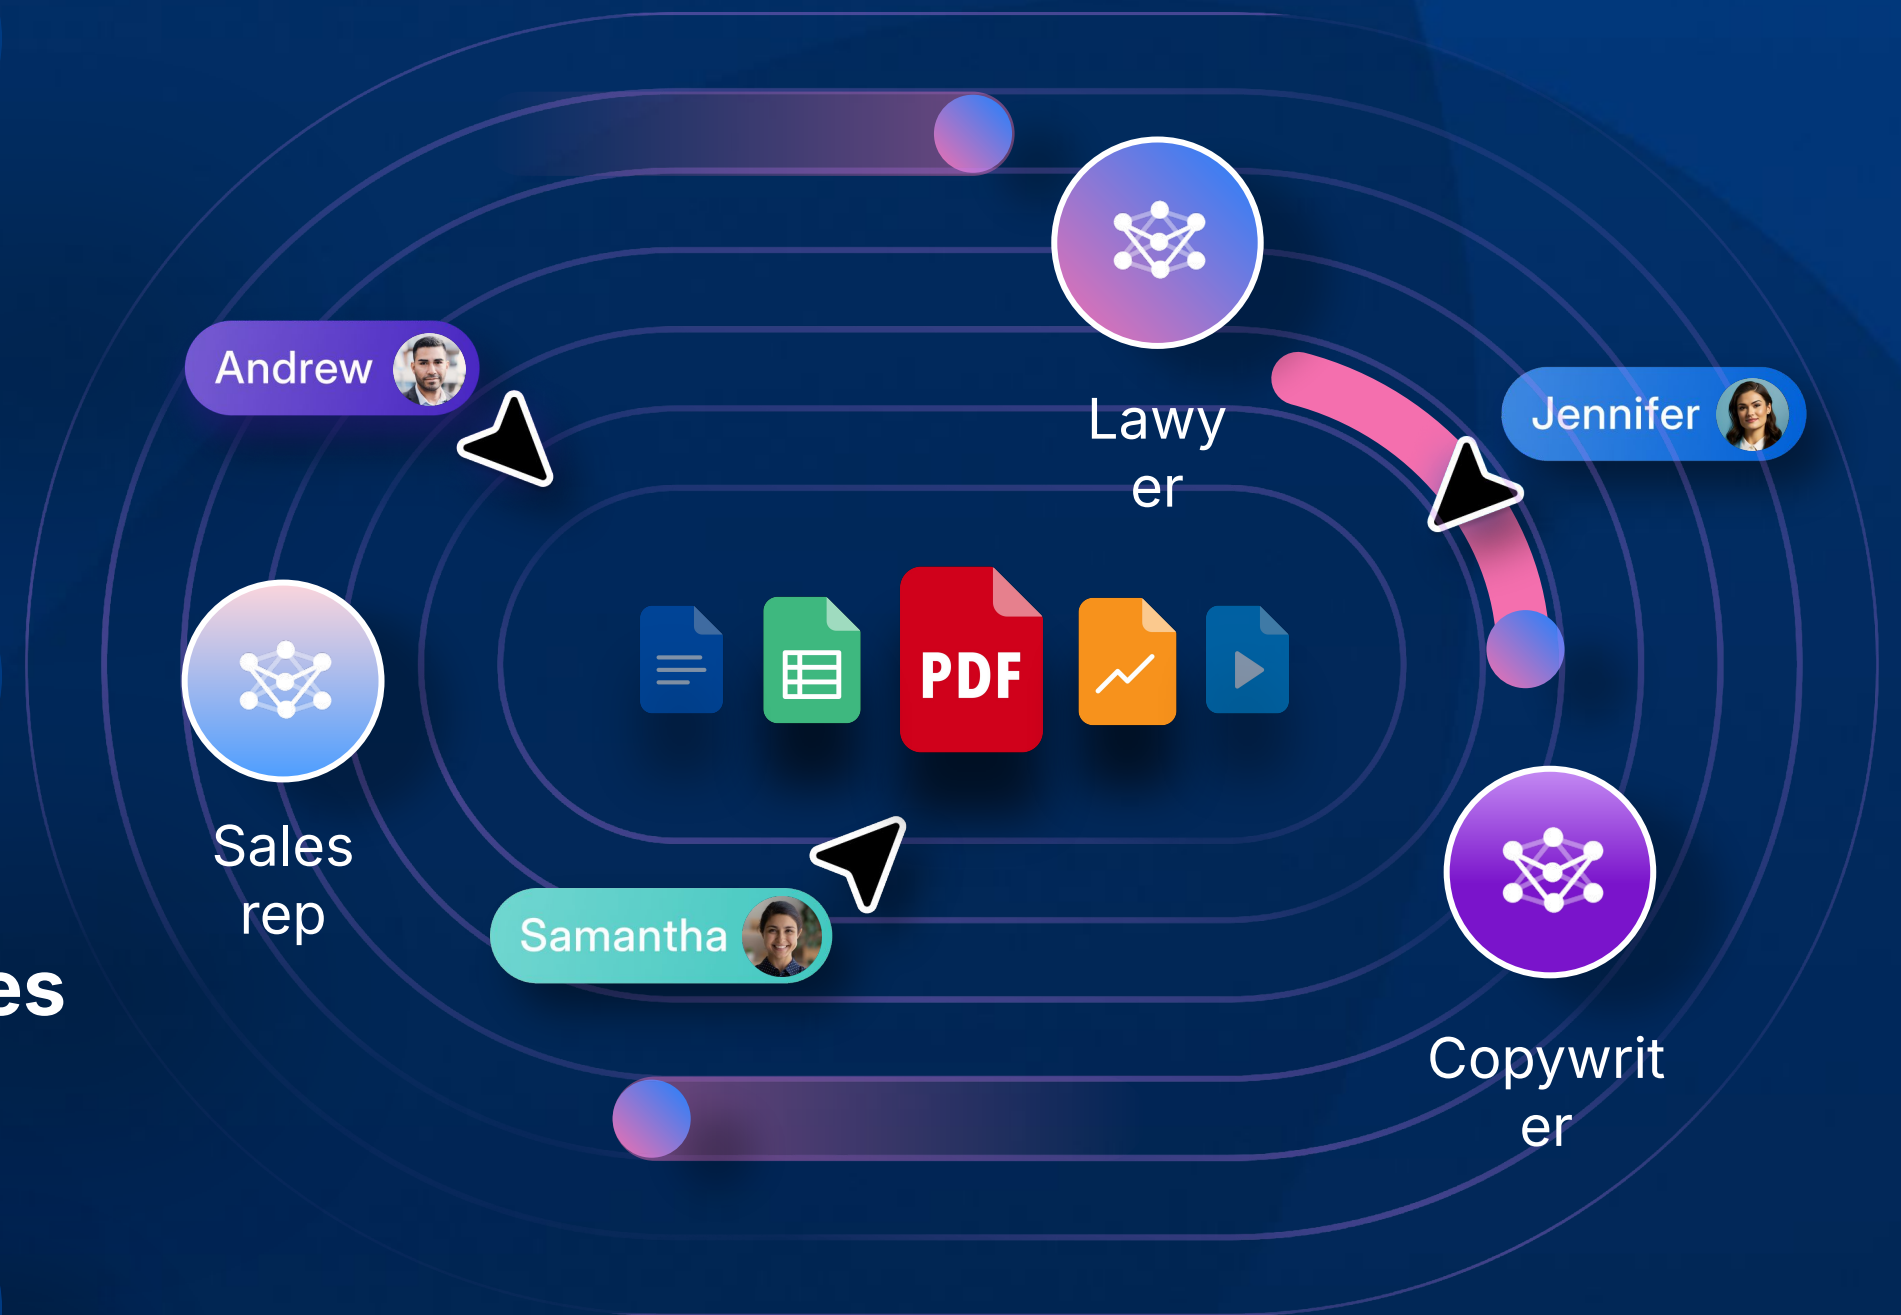
AI agents are joining the workforce 

 Any workflow can be automated

Instant intelligence from your data- 

 Customers expect better experiences

Data governance becomes essential 



# Content is at the center of how we work



Launch new  
products



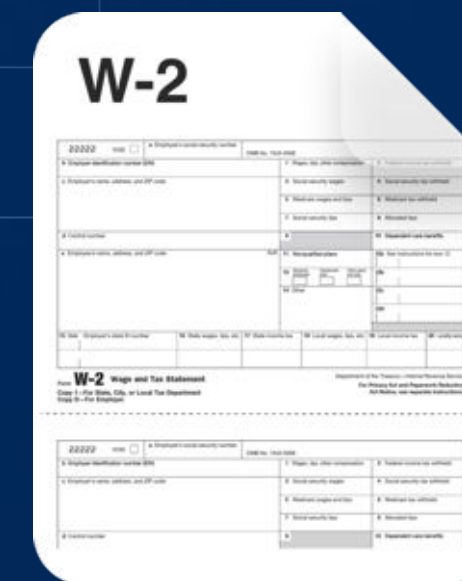
Market to  
customers



Close the  
books



Enable  
employees



Hire and  
onboard



Close  
deals



# Content is everywhere, but insights are missing



## Scattered and hard to find information

Important documents—contracts, reports, customer files—are buried across systems, making search slow and work inefficient.



## Disconnected and Manual

Without structure, automation fails, workflows break, and teams waste time on repetitive tasks instead of driving impact.



## Open to Compliance and Security Risks

Lack of metadata means sensitive content isn't properly tagged, leading to security gaps, compliance risks, and data sprawl.

# This leads to content chaos

## Scattered and Hard to find information

Important documents—contracts, reports, customer files—are buried across systems, making search slow and work inefficient.

**✗ Customer requests slow down because support docs aren't easy to locate**

## Disconnected and Manual

Without structure, automation fails, workflows break, and teams waste time on repetitive tasks instead of driving impact.

**✗ Approvals get stuck because contracts don't have clear owners or expiration dates**

## Open to Compliance and Security Risks

Lack of metadata means sensitive content isn't properly tagged, leading to security gaps, compliance risks, and data sprawl.

**✗ Workflows break because they are missing rich, untapped context**



What does Metadata bring to your  
business?

We can now apply structure  
to unstructured data at scale  
with AI-powered metadata

# Metadata is the important context around documents

CONTRACT TYPE

MUTUAL NON-DISCLOSURE AGREEMENT

(v11262018)

CLIENT

ADDRESS

CLASSIFICATION

This Mutual Non-Disclosure Agreement (this "Agreement") is made and entered into as of the date last executed by the Parties below ("Effective Date") by and between Box, Inc. and its affiliates ("Box") with offices at 900 Jefferson Ave, Redwood City, CA 94063 and Galactech, Inc ("Participant"), with offices at 123 Smith Road, San Mateo, CA 94010 (jointly the "Parties" individually a "Party"). Whereas, the Participant wishes to explore a potential business opportunity to purchase or license services from Box and, in connection with the opportunity, each Party may disclose to the other certain confidential technical and business information that the disclosing Party desires the receiving Party to treat as confidential ("Purpose"). The Parties have agreed to do so subject to the terms and conditions as set forth below.

**1. Definition of Confidential information. "Confidential Information"** means all information or materials provided by a Party ("Disclosing Party") to the other Party ("Receiving Party") on or after the Effective Date of this Agreement, that are: (i) in tangible form and labeled "confidential" or the like; (ii) if disclosed orally are summarized and confirmed in writing to be confidential within a reasonable time from the initial disclosure; or (iii) information that a reasonable person knows or should have known to be confidential given the circumstances surrounding disclosure. The following information shall be considered Confidential Information whether or not marked or identified as such: any personally identifiable information such as the names of either Party's customers, strategic marketing plans and product roadmaps, source code, technical infrastructure security and compliance documentation, hardware configuration, discounts, and the terms of this Agreement. Confidential Information shall not include, or shall cease to include as applicable, information or materials that: (a) were generally known to the public on the Effective Date; (b) become generally known to the public after the Effective Date, other than as a result of the act or omission of the Receiving Party; (c) were rightfully known to the Receiving Party prior to its receipt thereof from the Disclosing Party; (d) are or were disclosed by the Disclosing Party generally without restriction on disclosure; (e) the Receiving Party lawfully received from a third party without that third party's



# Introducing metadata extraction with Box AI extract agents

Non-Disclosure Agreement.pdf

ACME Inc. Legal templates

Box AI

Open

Download

Share

1 DISCLOSURE AGREEMENT (NDA)

The agreement is made between ACME Inc. ("Disclosing Party") and Galactech Inc. ("Receiving Party").

The purpose of this agreement is to protect the Confidential Information of the Disclosing Party.

1. Definition of Confidential Information. For purposes of this Agreement, "Confidential Information" shall include all information or material that has or could have commercial value or other utility in the business in which Disclosing Party is engaged. If Confidential Information is transmitted orally, the Disclosing Party shall promptly provide writing indicating that such oral communication constituted Confidential Information.

2. Confidentiality Obligations. Each Party agrees that all Confidential Information provided by the other Party is proprietary and confidential, including but not limited to trade secrets, business methodologies, financial data, client information, and any other sensitive material disclosed during the relationship. Confidential Information shall only be used for the intended purpose of this Agreement and shall not be disclosed to any third parties without the express written permission of the Disclosing Party, except to employees, agents, or subcontractors who are under comparable obligations and are bound by confidentiality agreements. Confidential Information shall remain confidential until it is no longer confidential due to its nature or disclosure to the public through no fault of the Receiving Party.

3. Expiration Date. This agreement shall expire on November 12, 2025.

Contract

Autofill

Contract Type

NDA

Version

V01-01-2024

Participant

Galactech Inc.

Number of Parties

2

Expiration Date

11-12-2025

1/20

100%

🔍

🔄

🗑️

⬆️





Box AI extract agents in action

Travel agent

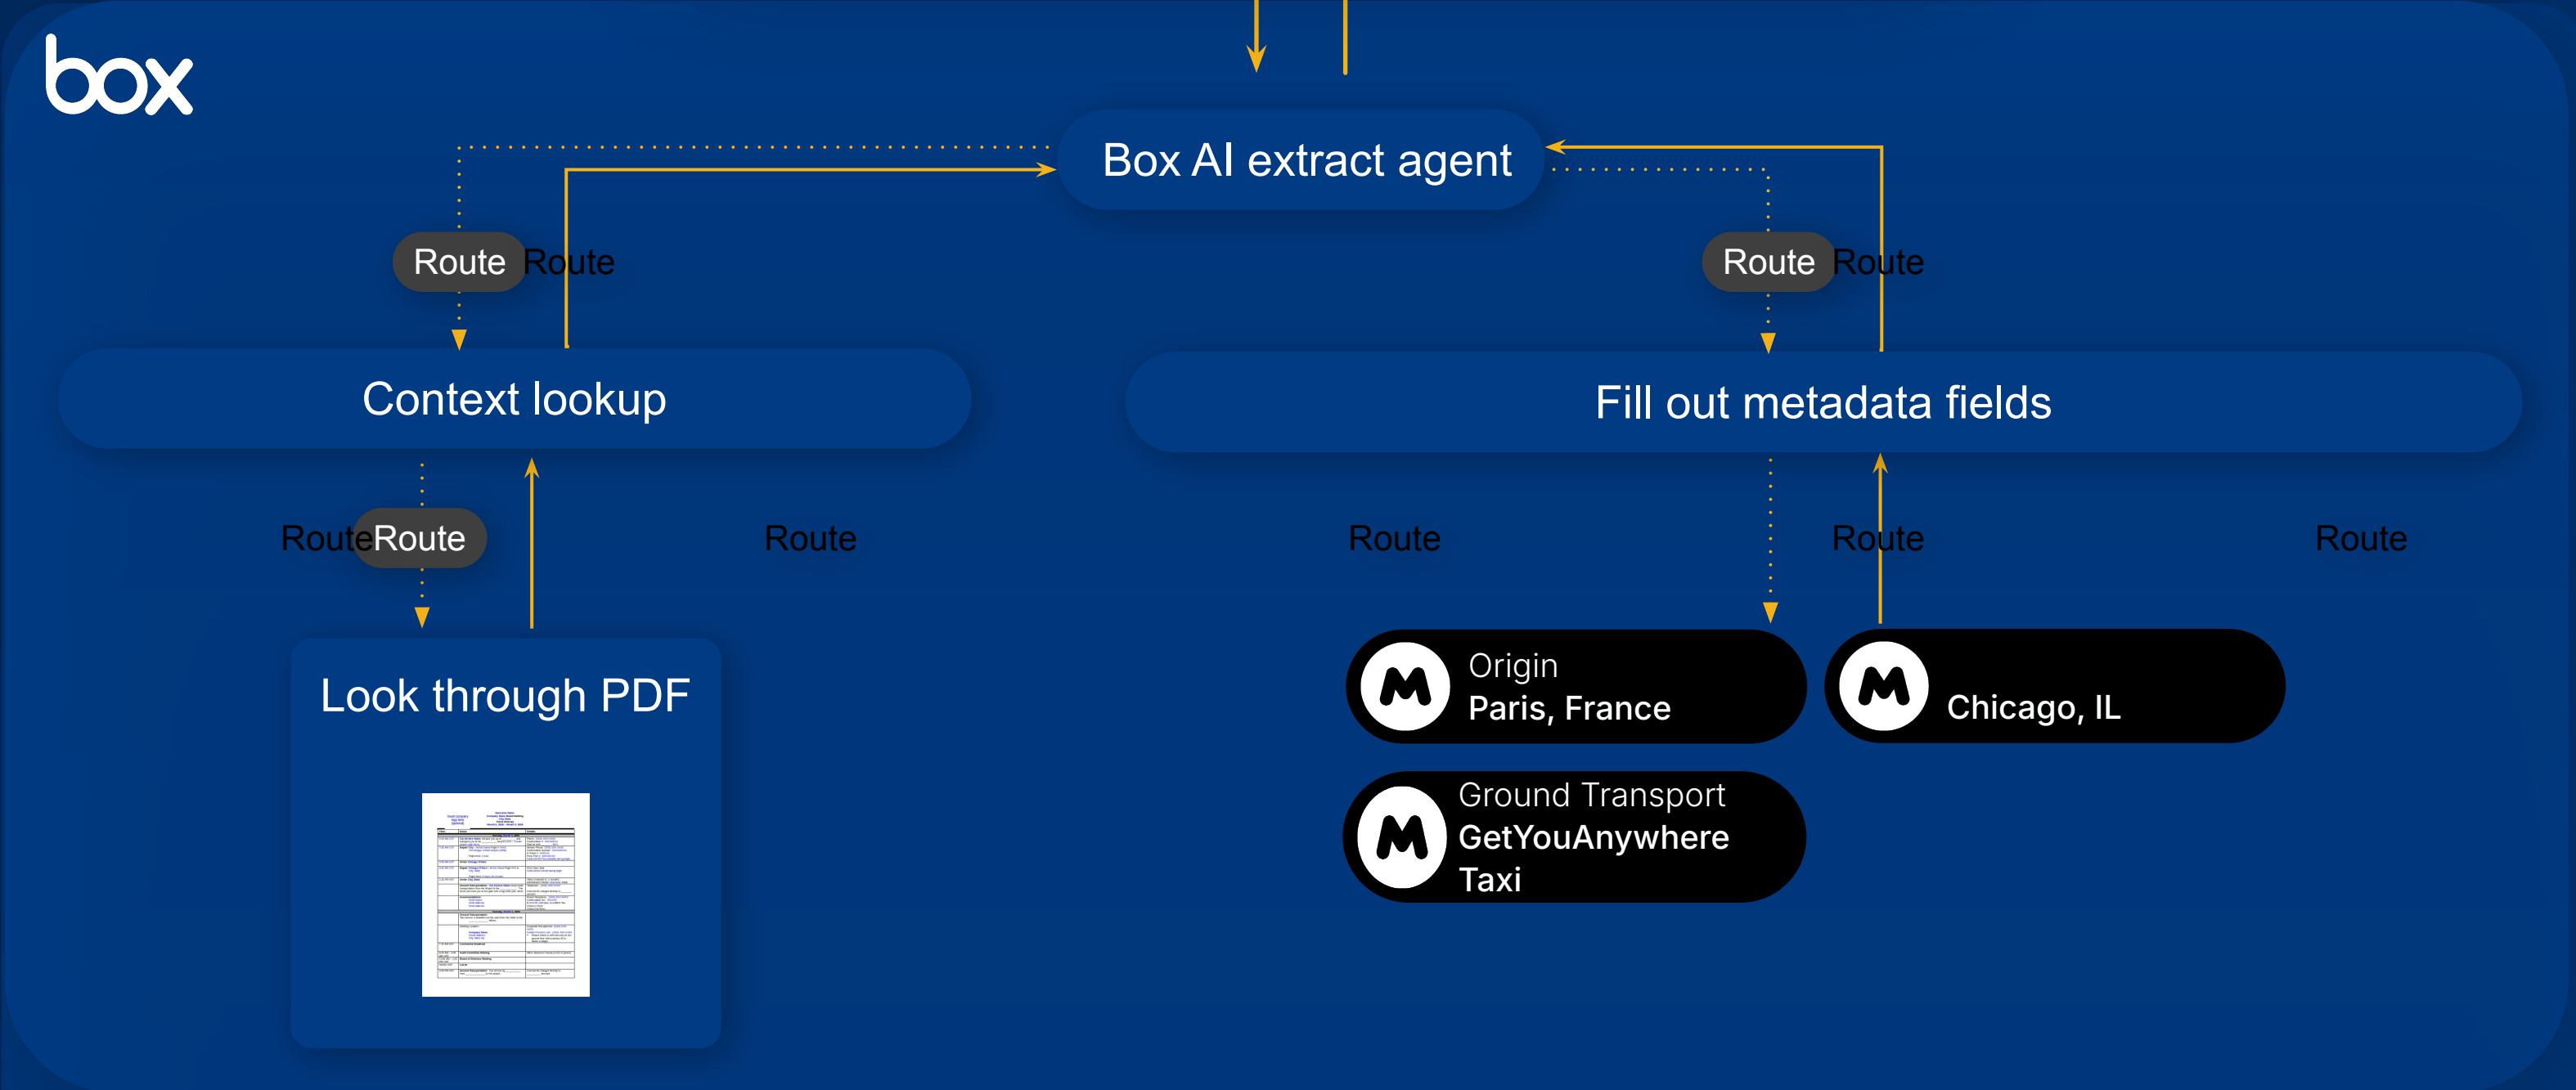
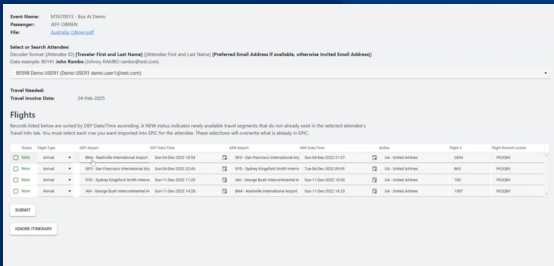


Travel itinerary  
in PDF



 Extract travel details

Extracted context  
into workflow





# Get intelligence from your data



## Hiring and onboarding

Managing resume collection and extraction, doc generation, e-signature, employee contract management, onboarding and benefits portals



## Financial capital firm

Processed 1 million documents as to extract structured data as part of compliance audit saving 1,000s of hours of manual work



## Travel industry

Saved over 800 hours of the manual process of processing bespoke travel itineraries by leveraging custom AI extract agents



# Successful Partnership: Box & Gemini Deliver Results for Fundwell



**FUNDWELL**

**Purchase Order**

Date: Apr 05, 2023  
Prepared by: Eula Jenkins  
Reference Number: BW-0505-01

Available Products:

- Adjustable Standing Desk (SKU: V11H84-F)
- USB Laptop Stand (SKU: V89112-D)
- Wireless Keyboard and Mouse Combo (SKU: P88720-F)
- Mesh Task Chair (SKU: V0029H2)
- 7-in-1 Docking Station (SKU: P900A8-D)

Receipt:

Description	Quantity	Unit Price
Adjustable Standing Desk	10	\$450.00
USB Laptop Stand	65	\$69.00
Mesh Task Chair	18	\$129.00

Total: \$11,307.00

Approval

Signature:

Name:

Centralized content

AI-powered contextual insights

Enhanced productivity & compliance

Improved collaboration



# 03. Fundwell



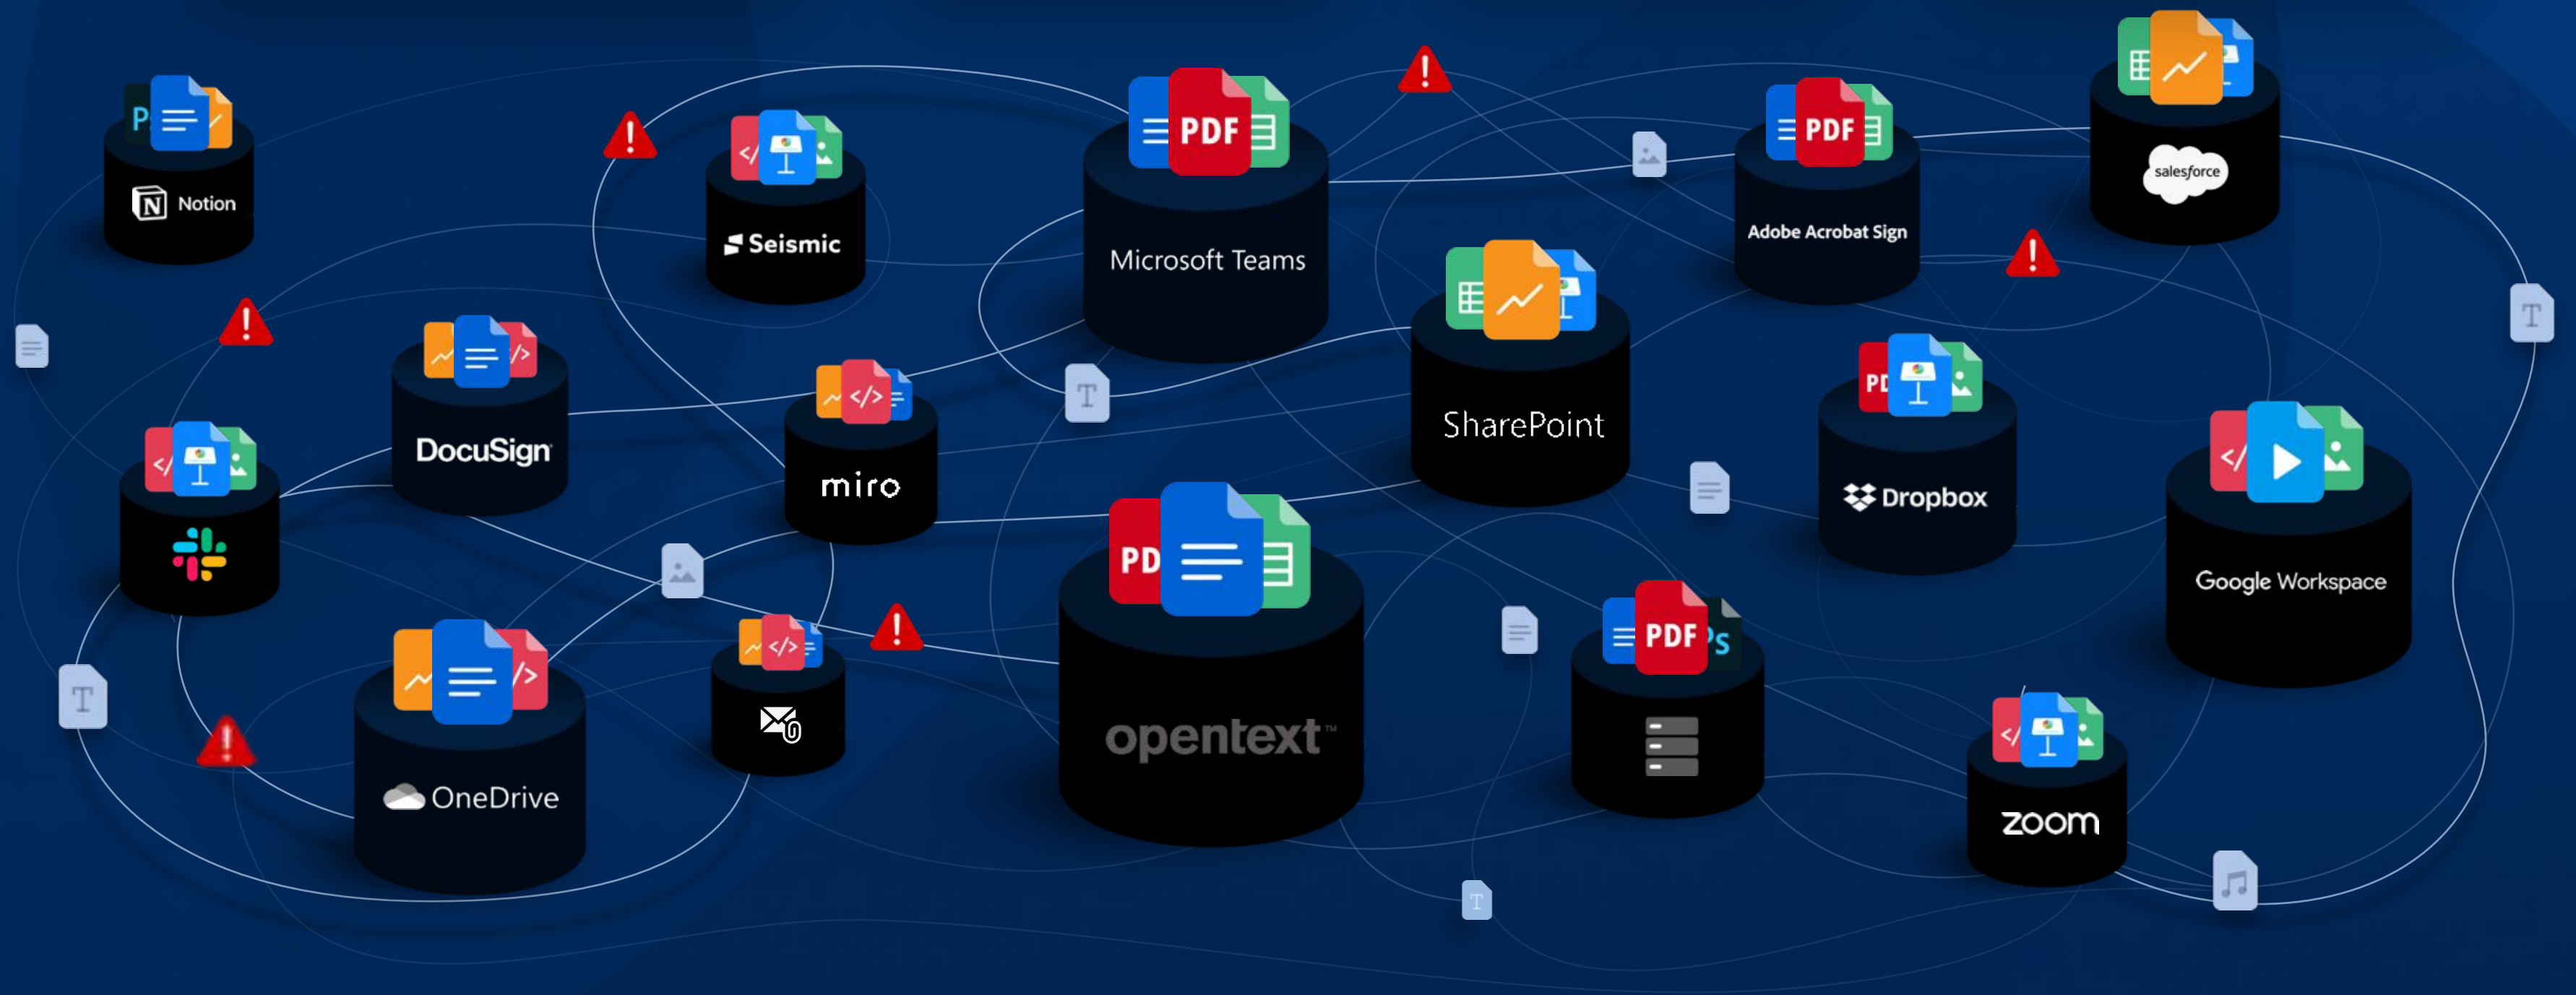
# The way we manage content is fundamentally broken

 Unproductive

 Unsecure

 Costly

 Siloed AI



# The full value of enterprise content is unrealized


10%

Structured data



90%

Unstructured data

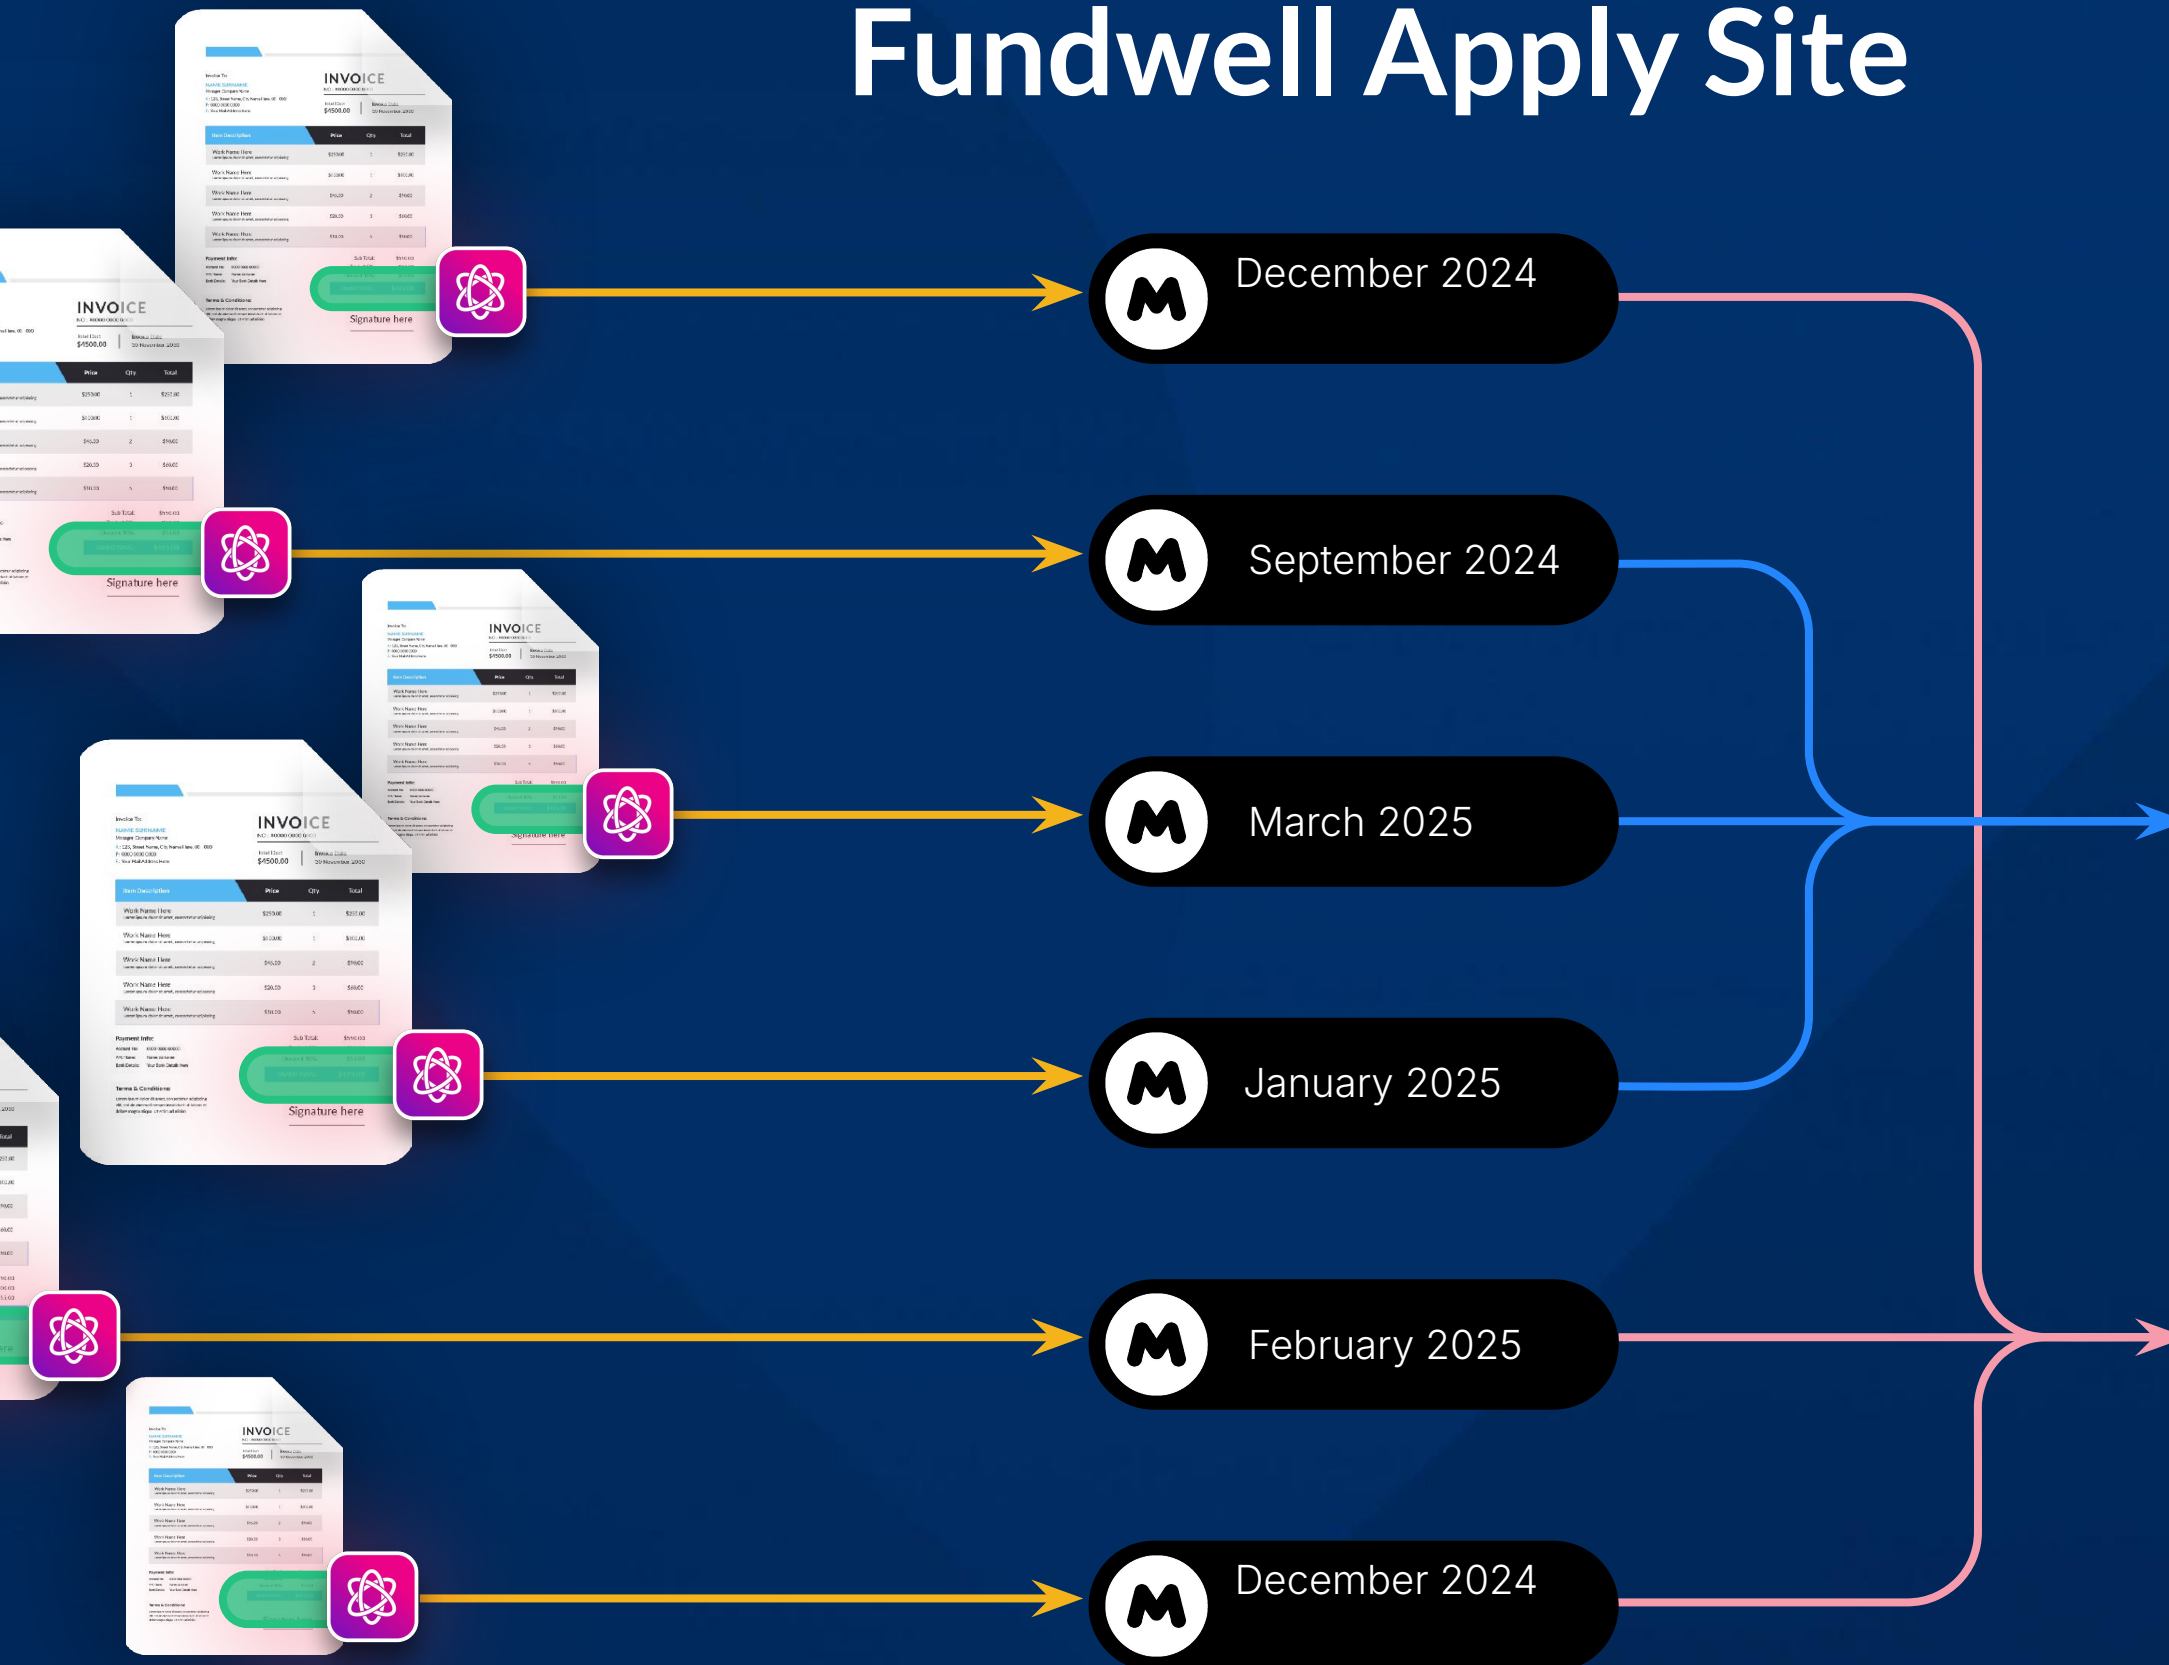


Untapped  
value



# Fundwell Apply Site

# Salesforce



# Custom-Built Underwriting System



## Challenges

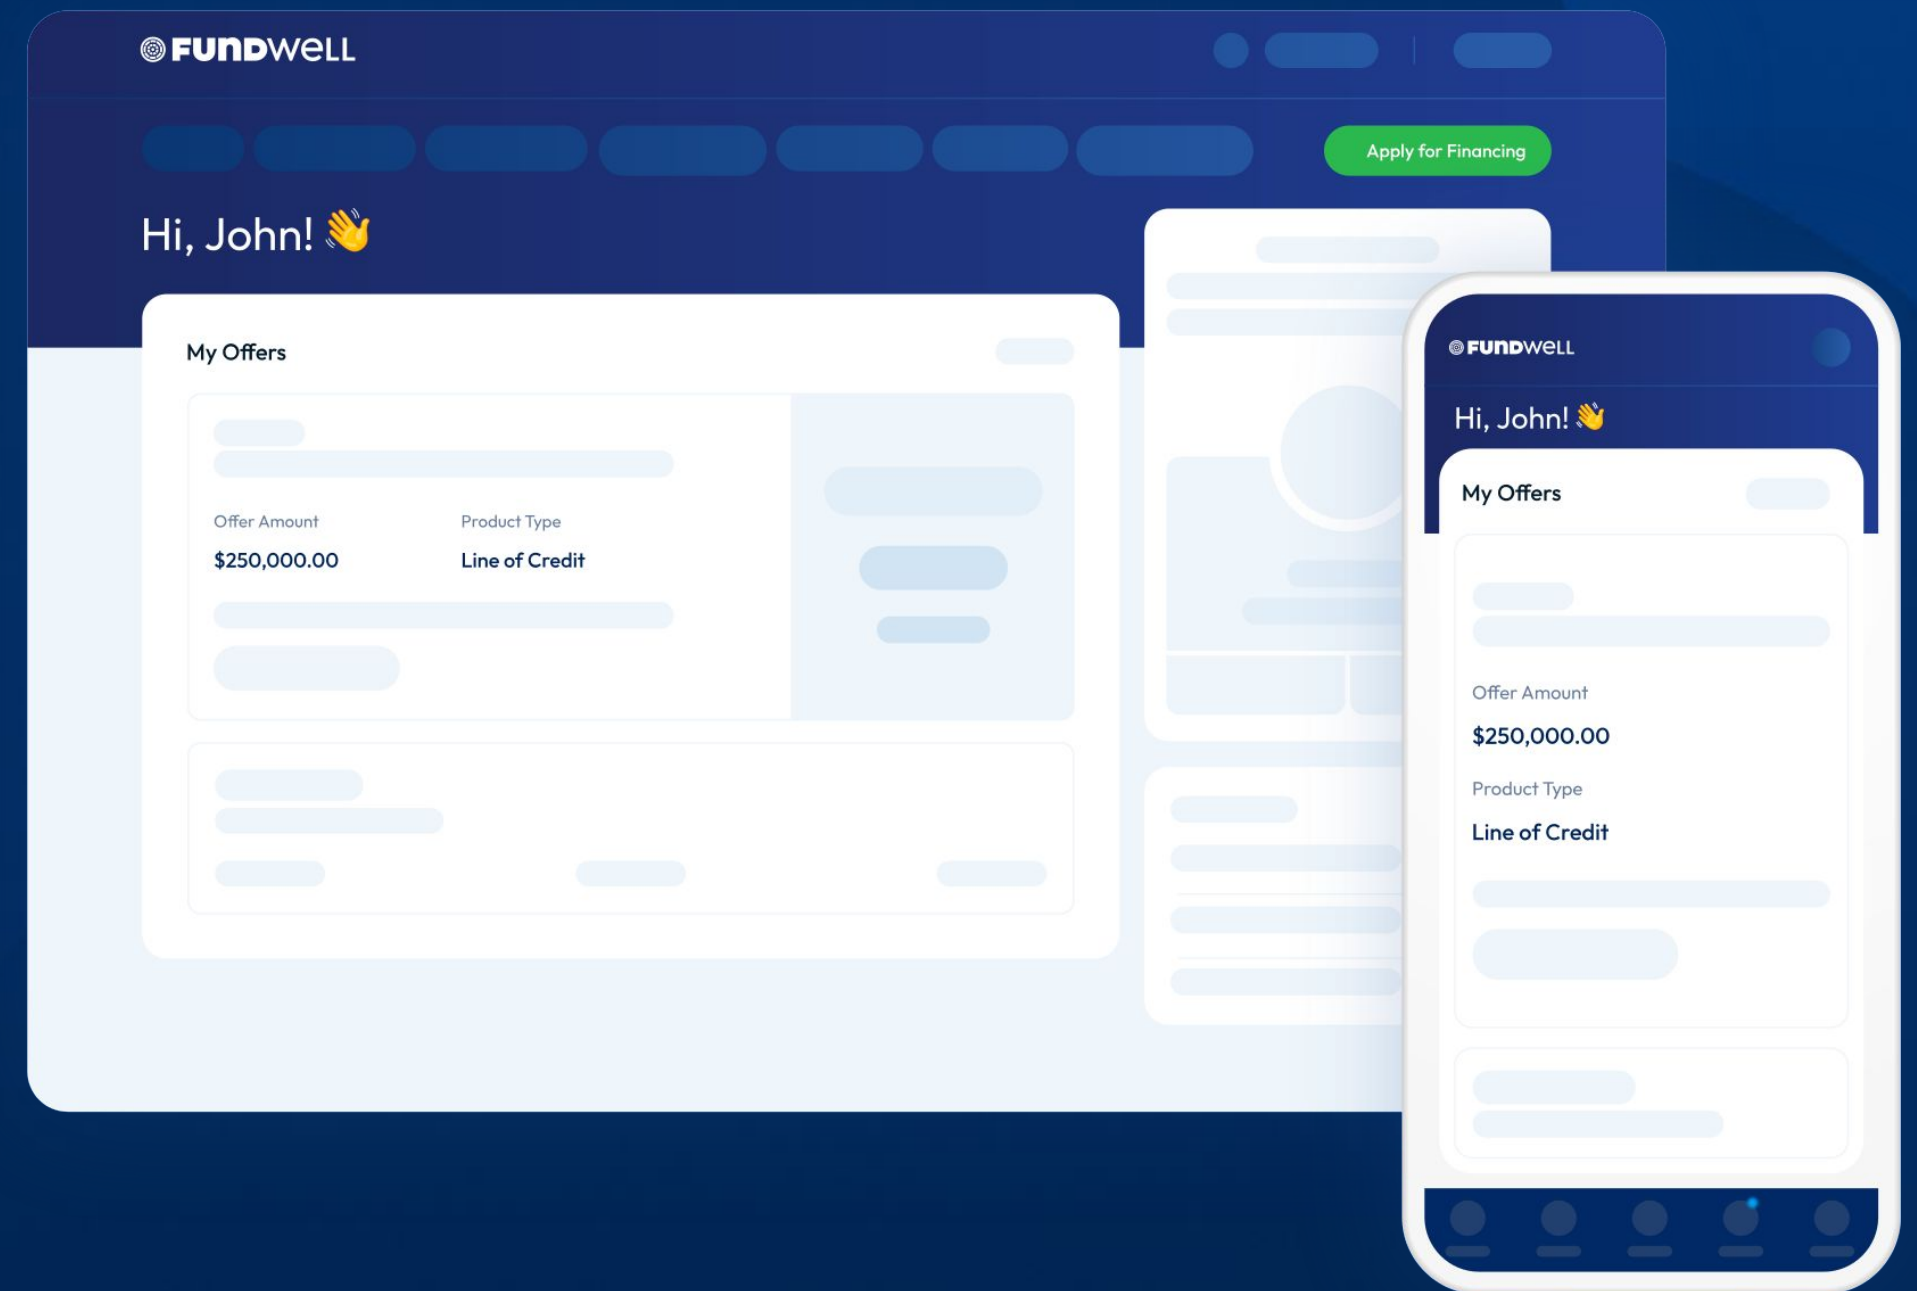
- **Content centralization and security:** Teams struggle to centralize Salesforce content securely, which is crucial for protecting sensitive client data during collaboration with external funding organizations.
- **Inefficient workflows:** The lengthy paperwork process leads to inefficiencies for clients and staff, underscoring the need for streamlined workflows to enhance productivity.

## Outcomes


- **Scalable processes:** Automated content workflows eliminate the need for manual reviews, enhancing efficiency and scalability.
- **Secure AI workflow:** Files processed securely through Box AI allow for private and secure metadata extraction, maintaining confidentiality and compliance.

**1,500+ hours annually**

Saved on Manual Document Review







- Are your internal roles and permissions easy to manage?
- Would your internal systems pass a security audit like SOC2?
- Is adding or removing a contractor or employee a cumbersome?
- Do you struggle to find tools that effectively solve your business problems?
- Is your identity management system both secure and compliant?

Answered “No”  
To Any Of These?



Join the Waitlist



The background is a dark navy blue. In the top right corner, there are several overlapping geometric shapes: a red-to-orange gradient triangle pointing down-right, a yellow triangle pointing down-left, and a green triangle pointing up-right. In the bottom left corner, there is a blue square, a purple-to-orange gradient wavy shape, and a blue wavy shape. The text "Thank you" is centered in a large, white, sans-serif font.

# Thank you



# Your feedback is greatly appreciated!



Complete the session  
survey in the mobile app