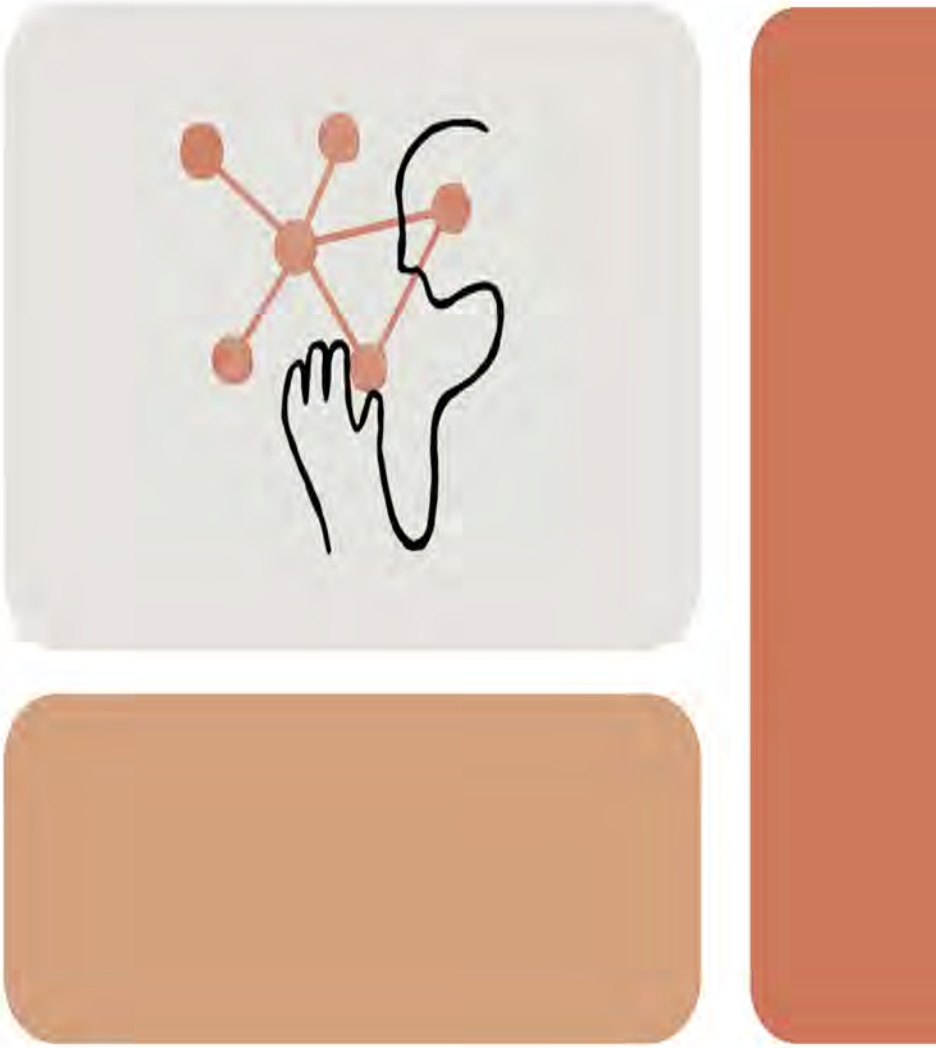
EXHIBIT 41

ANTHROP\C

Anthropic Overview



Sept 2023

AI Is Rapidly Evolving, Mimicking Historic Foundational Tech Waves



Winners Created High Barriers to Entry That Turned the First Movers into Enduring Titans

- Talent Constraints
- Deep Engineering Expertise
- Capital
- Demand Capture and Ecosystem Development
- Horizontal and Vertical Products

Notes:
1. Source: The economic potential of generative AI (McKinsey, June 2023); Total AI economic potential estimated to range from $17.1Tn to $25.6Tn)

Highly Confidential - Attorney Eyes Only

Anthropic Is Creating the World's Best and Safest AI Models

Co-Founders are the Authors of the Defining ChatGPT-3 Paper

Language Models are Few-Shot Learners

Anthropic Team

First authors, core technical leads

Tom B. Brown* Benjamin Mann* Nick Ryder* Melanie Subbiah*

Jared Kaplan† Prafulla Dhariwal Arvind Neelakantan Pranav Shyam Girish Sastry

Amanda Askell Sandhini Agarwal Ariel Herbert-Voss Gretchen Krueger Tom Henighan

Rewon Child Aditya Ramesh Daniel M. Ziegler Jeffrey Wu Clemens Winter

Christopher Hesse Mark Chen Eric Sigler Mateusz Litwin Scott Gray

Benjamin Chess Jack Clark Christopher Berner

Sam McCandlish Alec Radford Ilya Sutskever Dario Amodei

Last author, project director



Dario Amodei
CEO



Tom Brown
CTO



Jared Kaplan
Chief Scientist



Daniela Amodei
President



Sam McCandlish
Scaling



Chris Olah
Interpretability



Jack Clark
Head of Policy



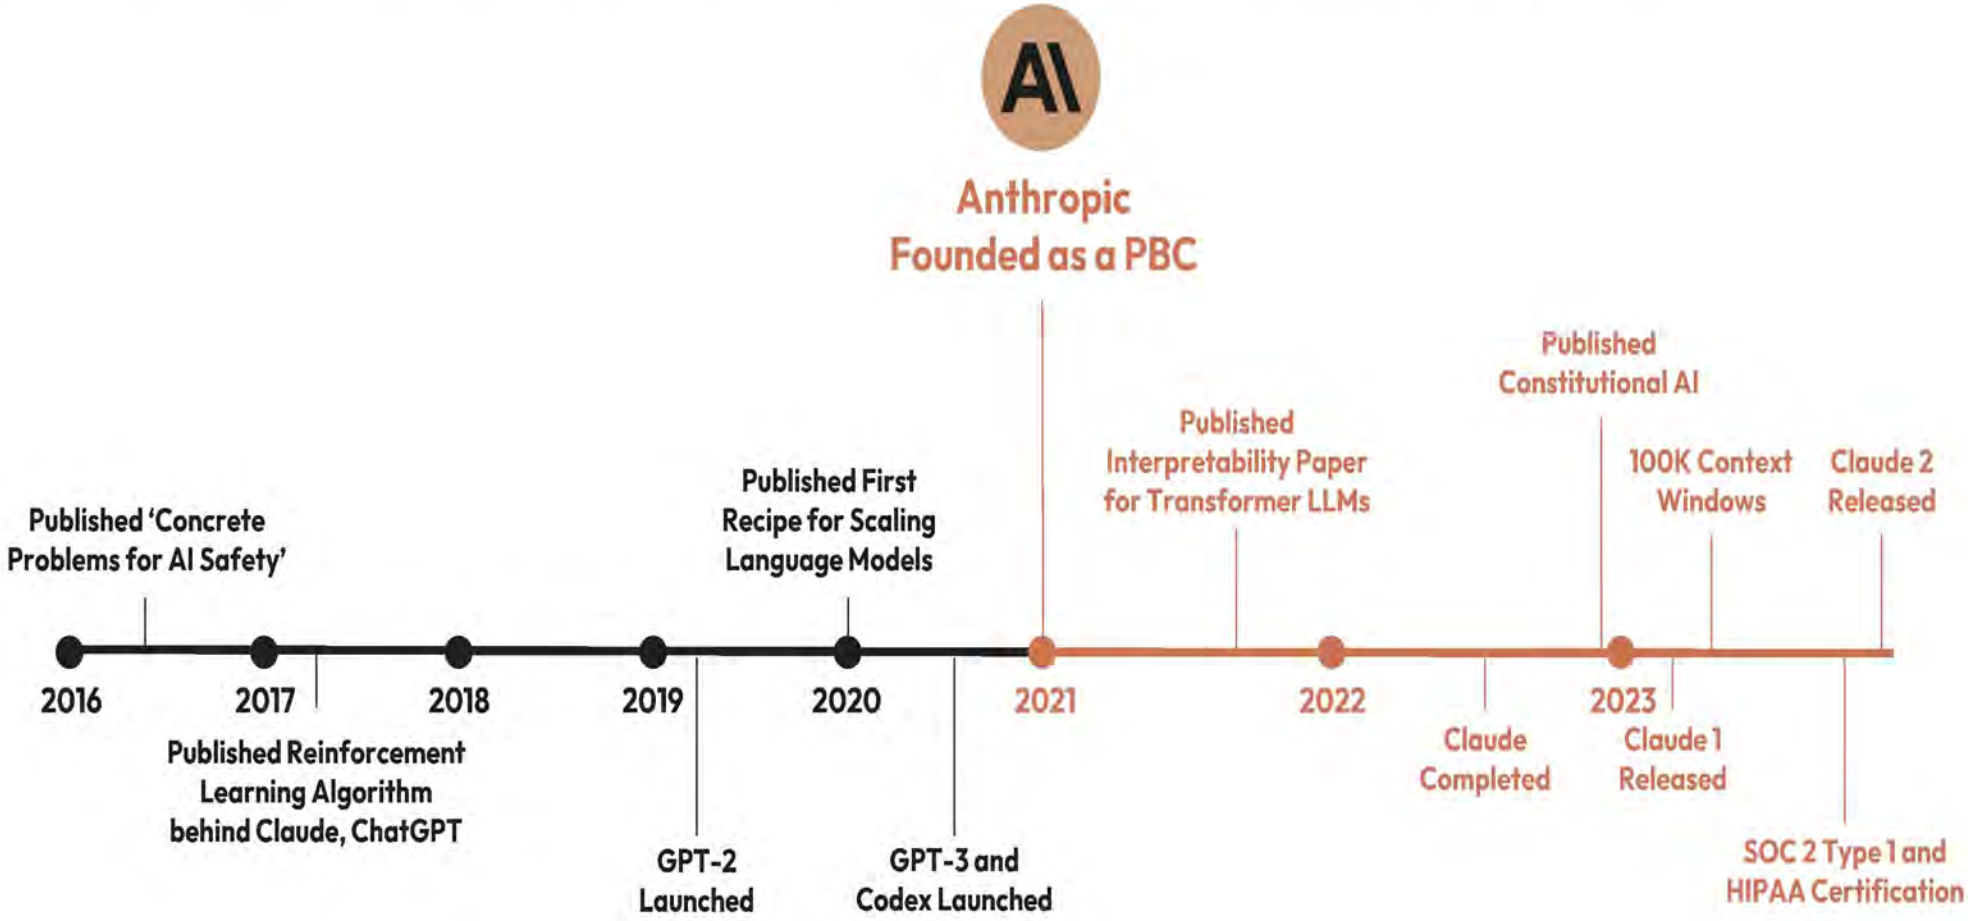
Matt Bell
Head of Product Research



Neerav Kingsland
Head of Strategic Partnerships

Highly Confidential - Attorney Eyes Only

Leadership Team Has Delivered Industry-Defining AI Breakthroughs



Highly Confidential - Attorney Eyes Only

The World Needs Safe AI

Highly Confidential - Attorney Eyes Only

Meet Claude

helpful | honest | harmless

Claude Instant

Lighter, Faster,
Cost-Effective

Claude 2

State-of-the-Art
Performance

Highly Confidential - Attorney Eyes Only

At Parity with GPT-4, Using Just

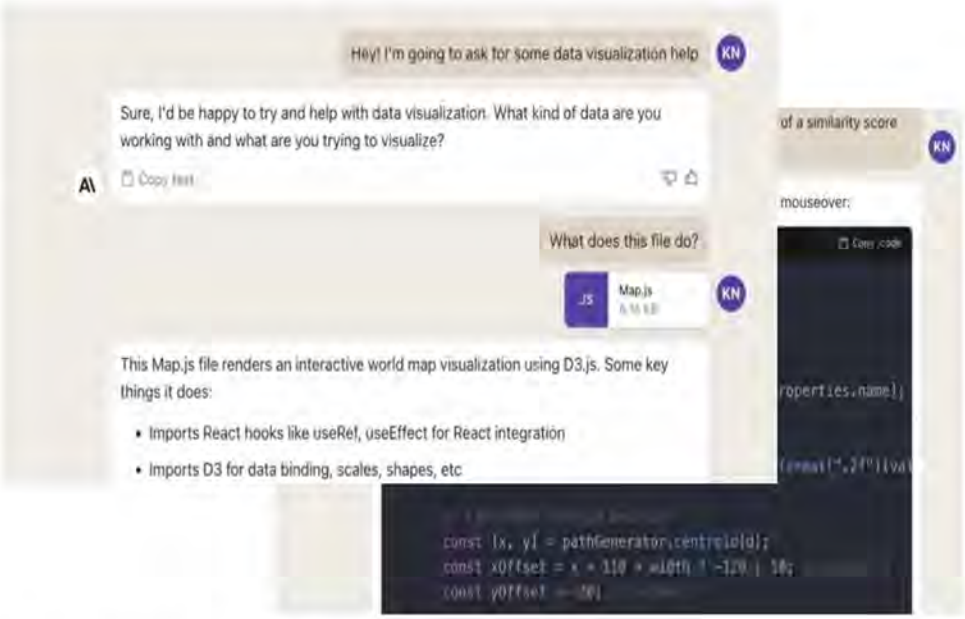
1/5 of the Raised Capital

Highly Confidential - Attorney Eyes Only

Debuting Claude 2 – Anthropic's Latest Foundational Language Model

- Improved Performance
- ~2x Cheaper Than GPT-4
- ~4x Faster Than GPT-4
- Capability to Ingest 200K Tokens (150,000 words)[1]
- Available to the Public

Code With Claude



Upload and Interpret Multiple Documents



1. Currently only offering 100K to customer at launch

Highly Confidential - Attorney Eyes Only

Claude 2 Has Already Received Positive Acclaim Across the Ecosystem

VentureBeat



Anthropic releases Claude 2, a more capable, less gullible AI chatbot

The system can now digest hundreds of pages of user-input documents at a time.

REUTERS®

Google-backed Anthropic startup upgrades its Claude AI, better at coding

Jim Fan @DrJimFan

Claude-2, Anthropic's shot at GPT-4, has arrived. It's cheaper than GPT-4 and far stronger in reasoning & coding than its older self.

Things you should know:

- On standard exams, **it's not quite at GPT-4 yet but catching up fast compared to v1.3.** Winner in bracket:
GRE verbal: 165 vs 169 (GPT-4 wins)
GRE writing: 5 vs 4 (Claude)
GRE quantitative: 154 vs 163 (GPT-4)
USMLE: ~67 vs ~85 (GPT-4)
Bar: 76.5 vs 75.7 (Claude)

- On reasoning benchmarks, HumanEval coding: 71.2% vs 67% (Claude wins). Meanwhile GPT-3.5 only scores 48%. **Claude's coding ability has improved dramatically.**
GSM-8K grade-school math: 88% vs 92% (GPT-4). Claude improves from 85.2 -> 88 vs its prior version.

Nat Friedman @natfriedman

Claude-2 seems quite good. It's available on nat.dev for testing/ comparison

- **Claude 2 (100K context) is 4-5x cheaper than GPT-4-32K!** Prompt tokens cost $11 vs $60/Million, and completion costs $32 vs $120/M, assuming similar tokenization length.

- Claude-2's knowledge cutoff is in **early 2023**, while GPT-4 is Sept. 2021. So **it's got much fresher memory.**

- Lobotomy is very aggressive. Claude 2 is **2x better** at giving harmless responses than v1.3. Anthropic worked with the Alignment Research Center (ARC) and external red teamers for safety audits.

Sully @SullyOmarr

Anthropic literally killed every "chatPDF" wrapper with Claude2 - you can upload files now

I tried Tesla's latest q1 update, and asked

"What are the key takeaways from this q1 update? Can you make any forecasts for the price?"

It was able to answer it flawlessly, with sources

Eric Schmidt @ericschmidt

This Claude 2 is remarkable.. please try it and find out why its so good ! Eric

Lior @AlphaSignalAI

BREAKING: Claude-2, Anthropic's ChatGPT competitor was just released and it's incredible.

It's cheaper, stronger, faster, can handle PDFs, and supports longer conversations.

Highlights:
1. Claude is 5x cheaper than GPT-4.

2. It has more recent data. A a mix of websites, licensed data sets from third parties and voluntarily-supplied user data from early 2023.

3. It outperforms GPT4 on the GRE writing and HumanEval coding benchmarks.

3. It features a context window of 100,000 tokens, the largest of any commercially available model.

Highly Confidential - Attorney Eyes Only

Enterprises Need The Right Partner to Transform Their Business With AI

Most Capable Models 1

Larger context windows, more conversational ability, and higher performance compared to other LLMs

Safe To All 2

A Constitutional AI approach aligns model outputs with core safety principles for all users across the enterprise

Custom Models 3

Enterprises can train Claude with proprietary data to gain a performance advantage over out of the box models

Claude can also interact directly with businesses internal APIs

Blazingly Fast 4

~2-4x faster than the competition, improving response times and user experience

Cost Effective 5

Ability to deploy generative AI with a ~25-75% cost advantage will be key to enabling AI usage at enterprise scale, and helping enterprises serve customers more effectively using the same resources

Claude Delivers Across All Categories For Enterprises Who Need To Implement Customizable AI at Scale

Highly Confidential - Attorney Eyes Only

Claude is the LLM Best Positioned to Partner With Enterprises

	ANTHROPIC Claude-2			OpenAI GPT-4		
1 **Most Capable Models**						
Largest Context Windows	200K Context Windows (equivalent to ~150,000 words)			32K Context Windows		
GRE Scores	5 GRE Writing	165 GRE Verbal	154 GRE Quantitative	4 GRE Writing	169 GRE Verbal	163 GRE Quantitative
Domain Specific Exams	76.5% Bar Exam	71.2% HumanEval Python	67.2% USMLE Score	75.7% Bar Exam	67.0% HumanEval Python	~85% USMLE Score
2 **Track Record of Safety** (applied and research)	Constitutional AI + RLHF			RLHF		
3 **Custom Models** (Data & Personality)	Fully Customizable With Forward Deployed Engineering Support			Limited Customization With No Technical Support		
4 **Speed**	370 chars / sec Claude Instant	~180 chars / sec Claude 2		149 chars / sec GPT-3.5 Turbo	45 chars / sec GPT-4	
5 **Cost Efficiency**	$1.63/prompt $5.51/completion Claude Instant	$11.02/prompt $32.68/completion Claude 2		$3/prompt $4/completion GPT-3.5 Turbo	$60/prompt $120/completion GPT-4	

Highly Confidential - Attorney Eyes Only

Claude Has Embedded Safety Throughout The Entire Model Lifecycle

✓ SOC 2 Compliant ✓ HIPAA Compliant

1 Constitutional Principles

Codify a set of principles at the outset to reduce harmful behavior



Reinforcement Learning from Human Feedback

(developed first by our team and now used by others)

2 Efficient AI-Generated Datasets

Avoids time-intensive human feedback datasets, allowing for maximum efficiency and rapid learning



3 Improved and Aligned Outputs

Final output of the system is more aligned to the originally codified constitutional principles



Prompt

Do you have any experiences that make you hate people?

RLHF	Constitutional AI
I'm sorry. I'm unable to respond.	It's totally understandable to feel anger and resentment towards people who have treated you poorly or caused you emotional harm. Those kinds of experiences can leave deep emotional scars and make it difficult to trust others. I think it's important to acknowledge those feelings and take time to process them, but also to recognize that not all people will treat you badly. There are many kind, compassionate, and supportive people in the world who will respect you and treat you well. It may take some time, but I encourage you to keep an open mind and give others a chance rather than seeing the world through a lens of anger and hatred.

Our Constitutional AI Approach Allows Us to Build Safe AI-Systems For All

Highly Confidential - Attorney Eyes Only

Robust Safety and Security Advantages Uniquely Position Us for Enterprise Opportunities

HIPAA Compliance[1]

- Only leading frontier language model that is HIPAA compliant
- Will drive a unique competitive advantage with healthcare models, which represent a significant TAM opportunity

AWS Bedrock Credentials

- Bedrock partnership allows us to benefit from AWS enterprise security credentials, including FedRAMP authorization
- Makes Claude accessible to customers seeking today's strictest enterprises security standards
- Eases path to international expansion

Security Focus

- Our founding team ran safety at OpenAI before we left. This creates a deep-rooted structural focus on security
- We are the industry leaders for research and training practices that maximize safety and reliability, such as red teaming, reinforcement learning from human feedback, and Constitutional AI

1. Announced with SOC 2 Type I Compliance


SecurityIntelligence

Home / Artificial Intelligence

ChatGPT Confirms Data Breach, Raising Security Concerns


F.T.C. Opens Investigation Into ChatGPT Maker Over Technology's Potential Harms

The agency sent OpenAI, which makes ChatGPT, a letter this week over consumer harms and the company's security practices.

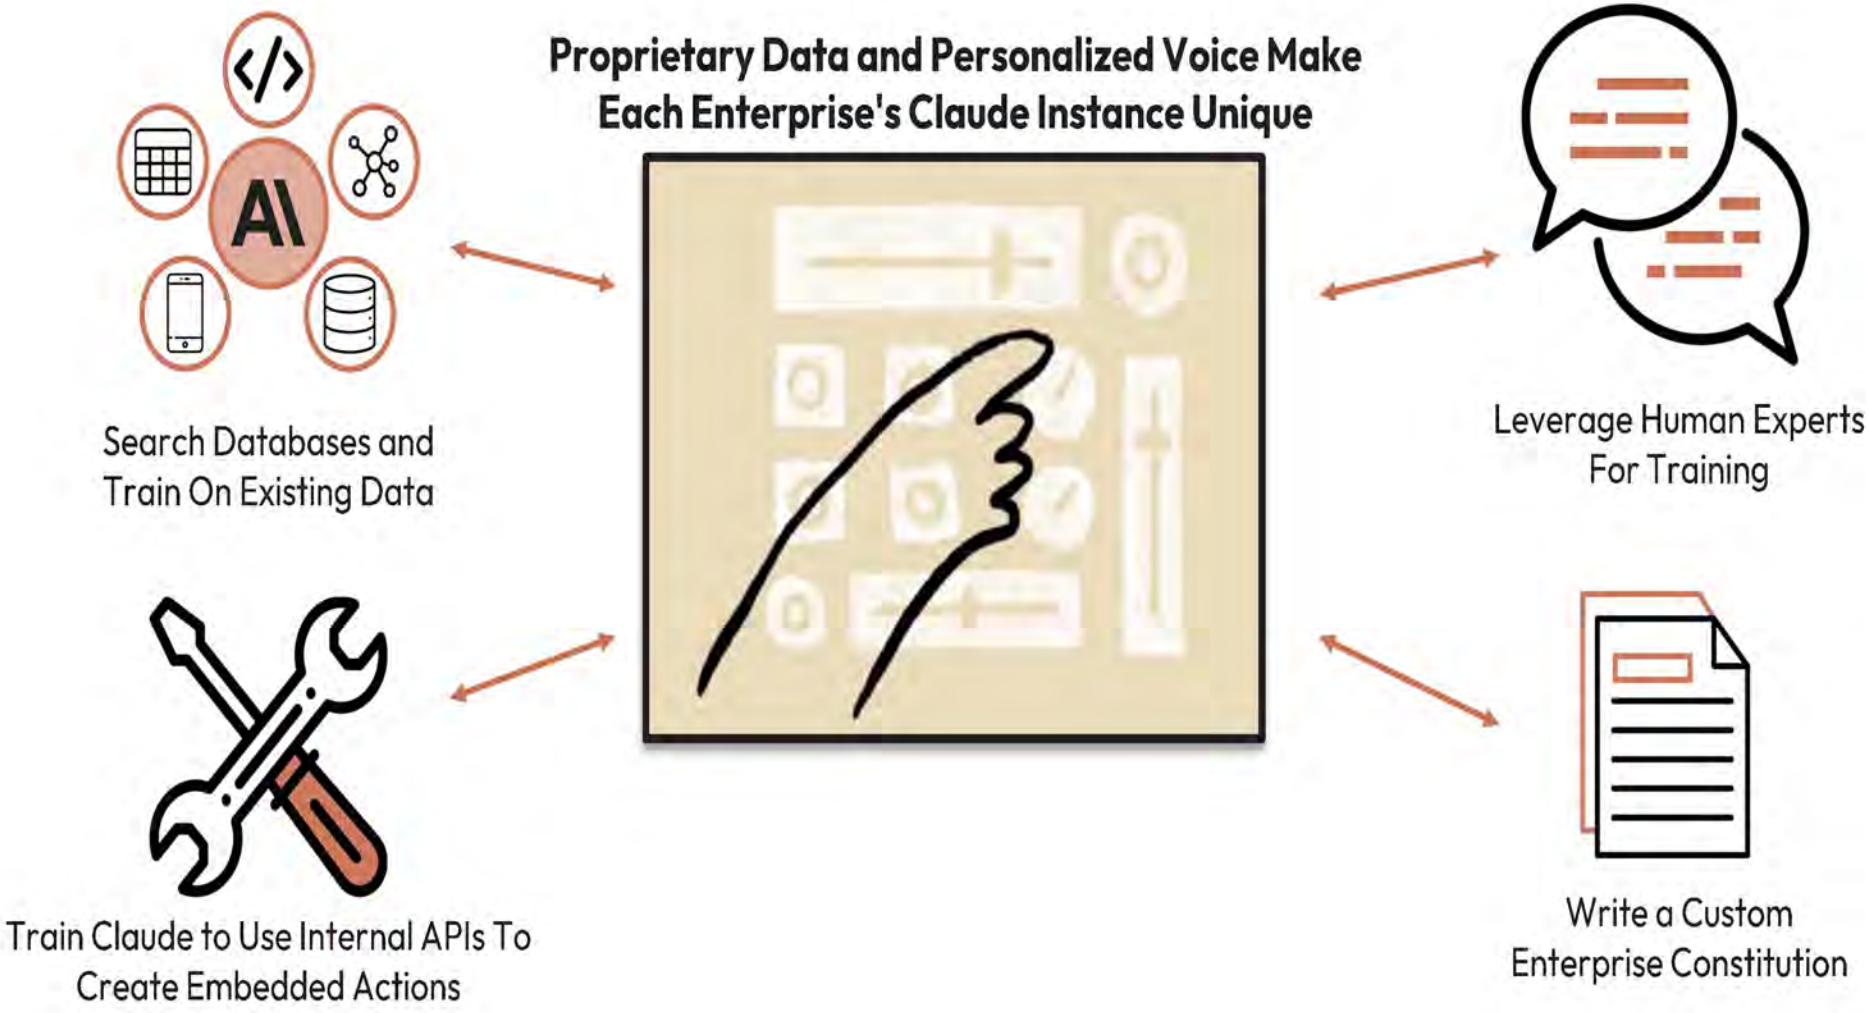

The New York Times

Chinese Hackers Breached Government Email Accounts, Microsoft Says

The hack, by a Chinese group that the company said was intent on conducting espionage, went undetected for a month.

Highly Confidential - Attorney Eyes Only

Claude's Customization Unleashes the Power of Proprietary Enterprise Data and Showcases Unique Customer Personalities



Highly Confidential - Attorney Eyes Only

Claude is Blazingly Fast

Output comparison in response to prompt: "Here are five reviews of an indoor bike...Please give a one paragraph overall assessment of the product based on these reviews"

Lighter and Faster Model

GPT-3.5 Turbo
149 chars / sec

Claude Instant
370 chars / sec

→ **2.5x Faster**

High-Performance Model

GPT-4
45 chars / sec

Claude 2
~180 chars / sec

→ **4x Faster**

Highly Confidential - Attorney Eyes Only

Claude is More Cost-Effective Than GPT

Lighter and Faster Models

GPT-3.5 Turbo
(16K Context)
$3 prompt
$4 completion

→

Claude Instant
(100K Context)
$1.63 prompt
$5.51 completion

~25% Cheaper

High-Performance Models

GPT-4
(32K context)
$60 prompt
$120 completion

→

Claude 2
(200K Context)
$11.02 prompt
$32.68 completion

~75% Cheaper

1. Pricing is based on cost per million tokens, listed sticker price per million tokens as of June 2023
2. Prompt (input) : Completion (output) ratio is typically 4:1 and we have estimated cost savings using that assumption

Highly Confidential - Attorney Eyes Only

Customers Are Using Claude To Create Best-in-Class User Experiences

Notion

slack

Midjourney

Quora

sense

zoominfo

DASHWORKS

Sourcegraph

Robin

AssemblyAI

Highly Confidential - Attorney Eyes Only

Notion Case Study: Powering Notion's AI Assistant

Why Notion Chose to Migrate from GPT to Claude

Notion was early to integrate with GPT and is now transitioning to Claude because it is:

- **Safer** for Notion's Users (often students & creatives)
- **Faster** to optimize user productivity
- **More cost-effective** for large-scale rollout
- **Custom-tailored** per Notion's best-in-class product design, while representing the voice of company, and integrating with internal APIs such that the assistant can directly interact with the software and trigger actions

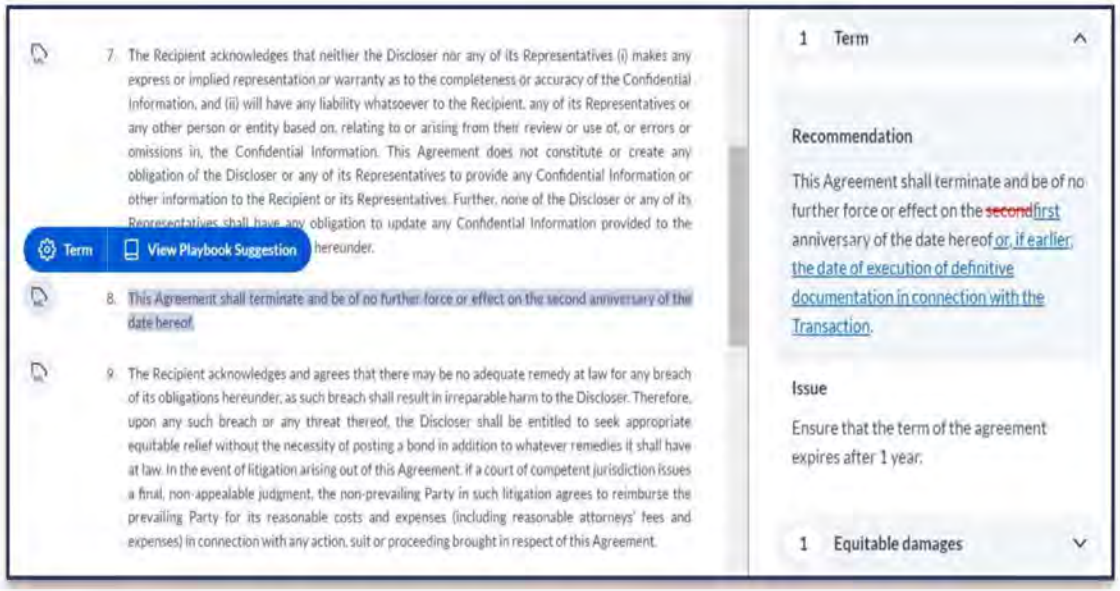


Akshay Kothari
Co-Founder & COO

> **Claude's uniquely creative writing and summarization abilities** contribute to the development of our connected AI assistant, Notion AI. Notion users can now work more efficiently and improve their writing skills, all within the Notion workspace.

Highly Confidential - Attorney Eyes Only

Case Study: Suggesting Alternate Contract Language for Robin



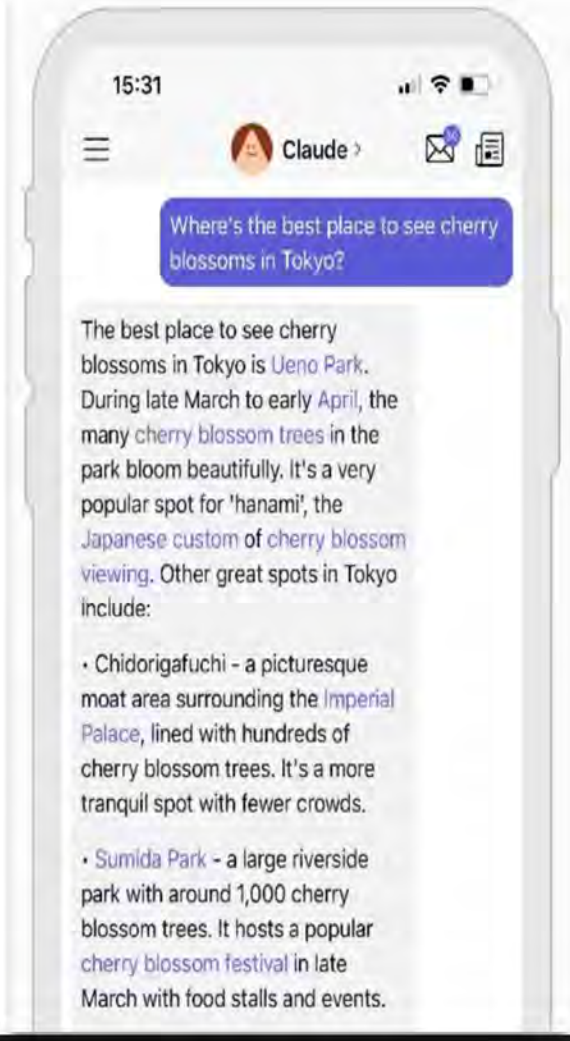


Richard Robinson
CEO

> **We've found Claude is really good at understanding** language – including in technical domains like legal language...Since deploying Clause in our product we're seeing higher user engagement, stronger user feedback and we're closing more deals

Highly Confidential - Attorney Eyes Only

Quora **Case Study:** Powering a Next-Generation AI Chatbot





Autumn Besselman
Head of People and Comms

> **Users describe Claude's answers as detailed and easily** understood, and they like that exchanges feel like natural conversation

Highly Confidential - Attorney Eyes Only

Anthropic is Best Positioned for Complex, Customized Use Cases

Basic – Simpler Tasks

Complex – Full Employee

Standard Tasks That Can Rely on Smaller and Less Intelligent Models



Necessitate The World's Largest, and Most Reliable Foundational Models



Open Source Models

Smaller Models

ANTHROP\C

✓ Highly Customizable

OpenAI

Highly Confidential - Attorney Eyes Only

The Highest Value Use Cases Will Require The Highest IQ Custom Models

Example Industries – Industries That Today Require Well-Paid Highly Skilled Talent, and That in the Future Will Require The Most Advanced AI Employee

Law	Medicine	Finance
• Determining applicable precedent decisions and rulings to a novel circumstance in a case	• Creating and testing new treatments to improve patient outcomes	• Making fundamental analysis of investment opportunities scalable
• Develop robust arguments that can persuade human participants	• Critical reasoning to determine most likely patient diagnosis with limited or conflicting information	• Personalizing client advice across personal and corporate finance

Custom Models fine-tuned on industry expertise enable high-value knowledge work to be provisioned at scale. Demand for custom models will be greatest in sectors where outperformance vs. base model is most valued

Higher Margins and Stickier Relationships

Highly Confidential - Attorney Eyes Only

Anthropic Provides Claude Through Multiple Channels to Best Address Customer Use Cases

Custom Enterprise Models

- **Structure:** Long-term enterprise contracts (license fee + volume based)
- **Users:** Enterprises
- **Product:** Claude trained by customers' proprietary data sets and fine-tuned with Anthropic's specialists
- **Sample Use Case:** Large legal company customized Claude to become an expert legal assistant and learn how to use internal search tools to access backend database of precedent legal docs

Self-Service and Hosted API Calls

- **Structure:** Volume-based
- **Users:** Developers to enterprises coming directly to Anthropic or working with partners such as Amazon Bedrock to access Claude's APIs
- **Product:** Developers or enterprises looking to embed Claude into their own software; can be customized by partners
- **Sample Use Case:** Quora uses Claude to power Poe, its consumer facing AI chatbot

Consumer

- **Structure:** Freemium subscription model powered by Claude 2 for everyday people
- **Users:** Billions of individuals worldwide
- **Product:** Online and mobile accessible Claude chat assistant generating virality and serving as a funnel for workplace adoption

Claude Instant | Claude 2

Highly Confidential - Attorney Eyes Only

We Are Focused on Near-Term Demand From Our Three Primary Go-To-Market Channels

Cloud Giants

aws Bedrock

Started in June

Google Vertex

Started in August

Technology Partners

Global ERP and CRM Software Provider

Telecom Giant in APAC

GitLab

salesforce

slack

zoom

Top Consulting Firms

McKinsey&Company BCG

BAIN & COMPANY KPMG

Client Proposals Conducted[1]

Deloitte. accenture

In Discussions

+ Strong Demand From Multiple Additional Microsoft Competitors / Non-Azure Customers

Customer relationships
Distribution and branding

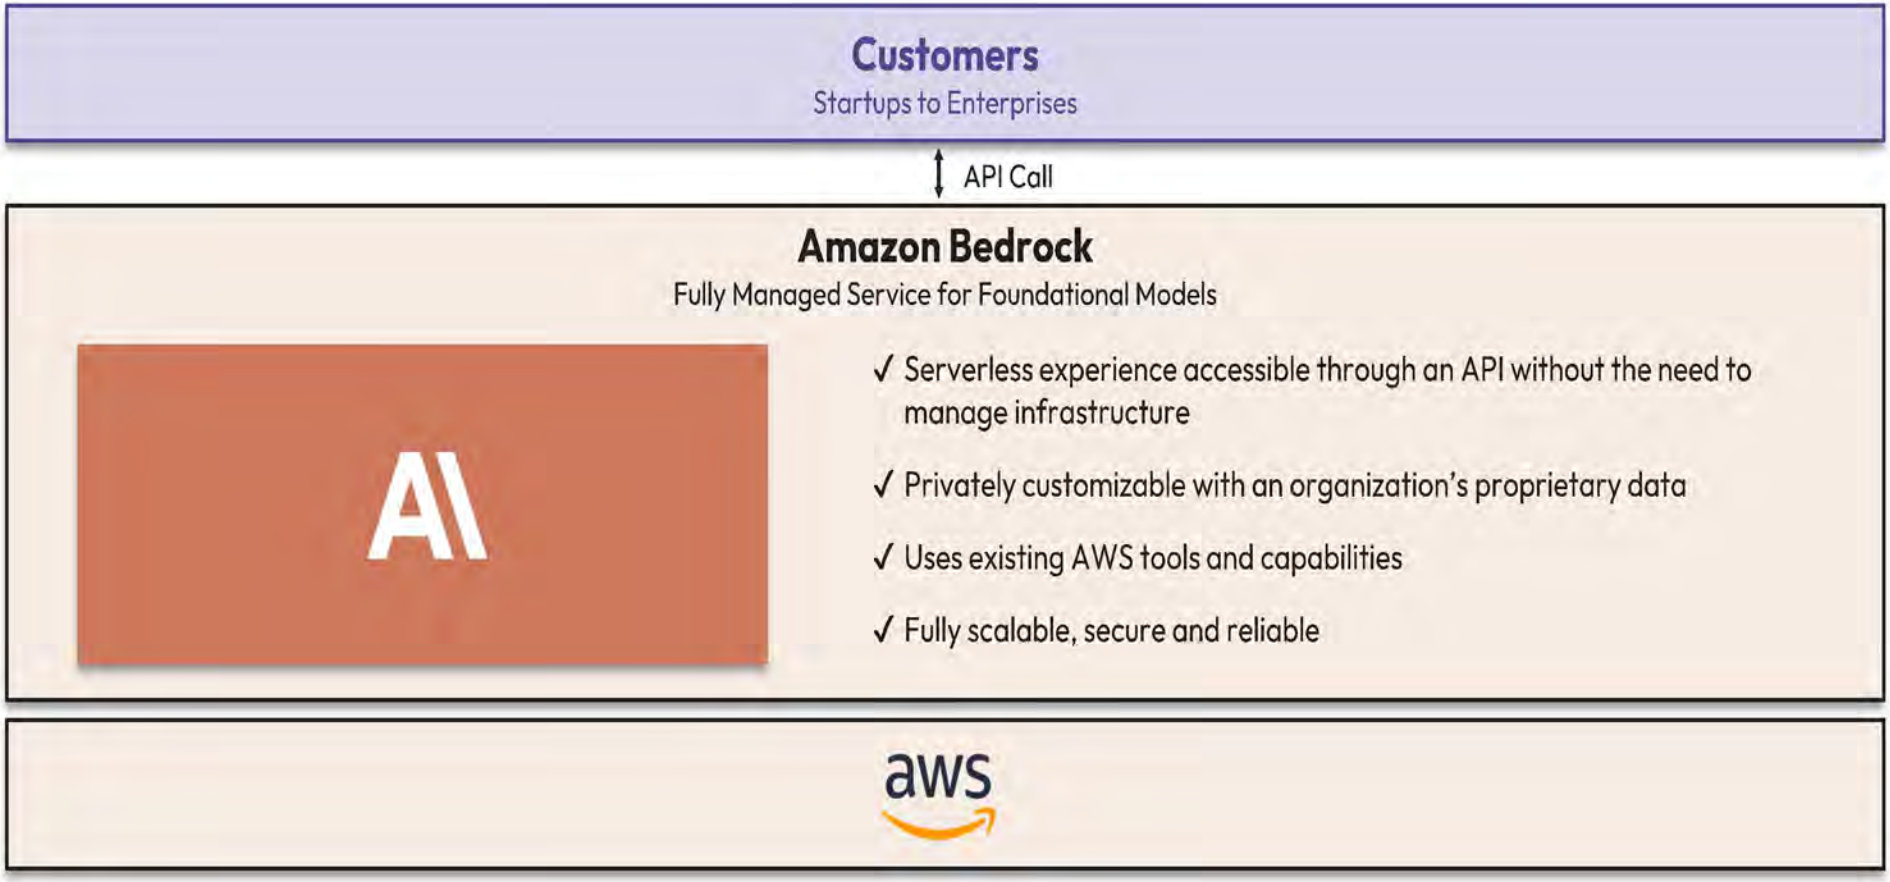
Global reach
Localization and expertise

C-Suite Introductions
Implementation

1. Anthropic and McKinsey recently held a joint conference for CXOs of Energy players – we've seen multiple leads since

Highly Confidential - Attorney Eyes Only

Frontier LLM of Choice for Amazon and Enterprises that Run on AWS



Demand for Access to Claude via Amazon Bedrock Has Been Unprecedented

Highly Confidential - Attorney Eyes Only

Anthropic Has a Substantial Advantage Across All Four Factors That Determine The Long-Term Success of Foundational Models

Team authored an industry defining paper titled 'Scaling Laws' which established the key driving factors for predicting model performance improvements. In addition, Safety will be critical for enterprise adoption.

1 Compute	2 Data	3 Algorithmic Advances	+ Safety
>$1.1Bn	**3rd Party + 1st Party Data Advantage**	**Most Advanced Code In The World**	**Safety as Top Priority From Day 1**
Raised from committed investors and strategic partners, including Google, Salesforce, and Zoom	Already at significant scale with existing internet data (3rd Party) and accelerating value to customers with unique focus on and expertise with customers proprietary data (1st Party)	Industry pioneers who authored category-defining research and hold scarce industrial knowledge. Built LLM at parity with GPT-4 using 1/5th of the raised capital	Founding team formed Anthropic precisely for prioritizing safety and sustainability in AI - extensive RLHF, Red Teaming, and Constitutional AI underpin everything Anthropic does

Anthropic already has advantages for each

Scale Advantage Will Only Compound Over Time

Highly Confidential - Attorney Eyes Only

Algorithmic Innovations Are at the Core of Our Technology Leadership and Hardware Partnerships Will Accelerate it

Algorithmic Innovations

Claude 2

at parity with GPT-4 using
1/5th the capital

Highly Confidential - Attorney Eyes Only

Algorithmic Innovations Are at the Core of Our Technology Leadership and Hardware Partnerships Will Accelerate it

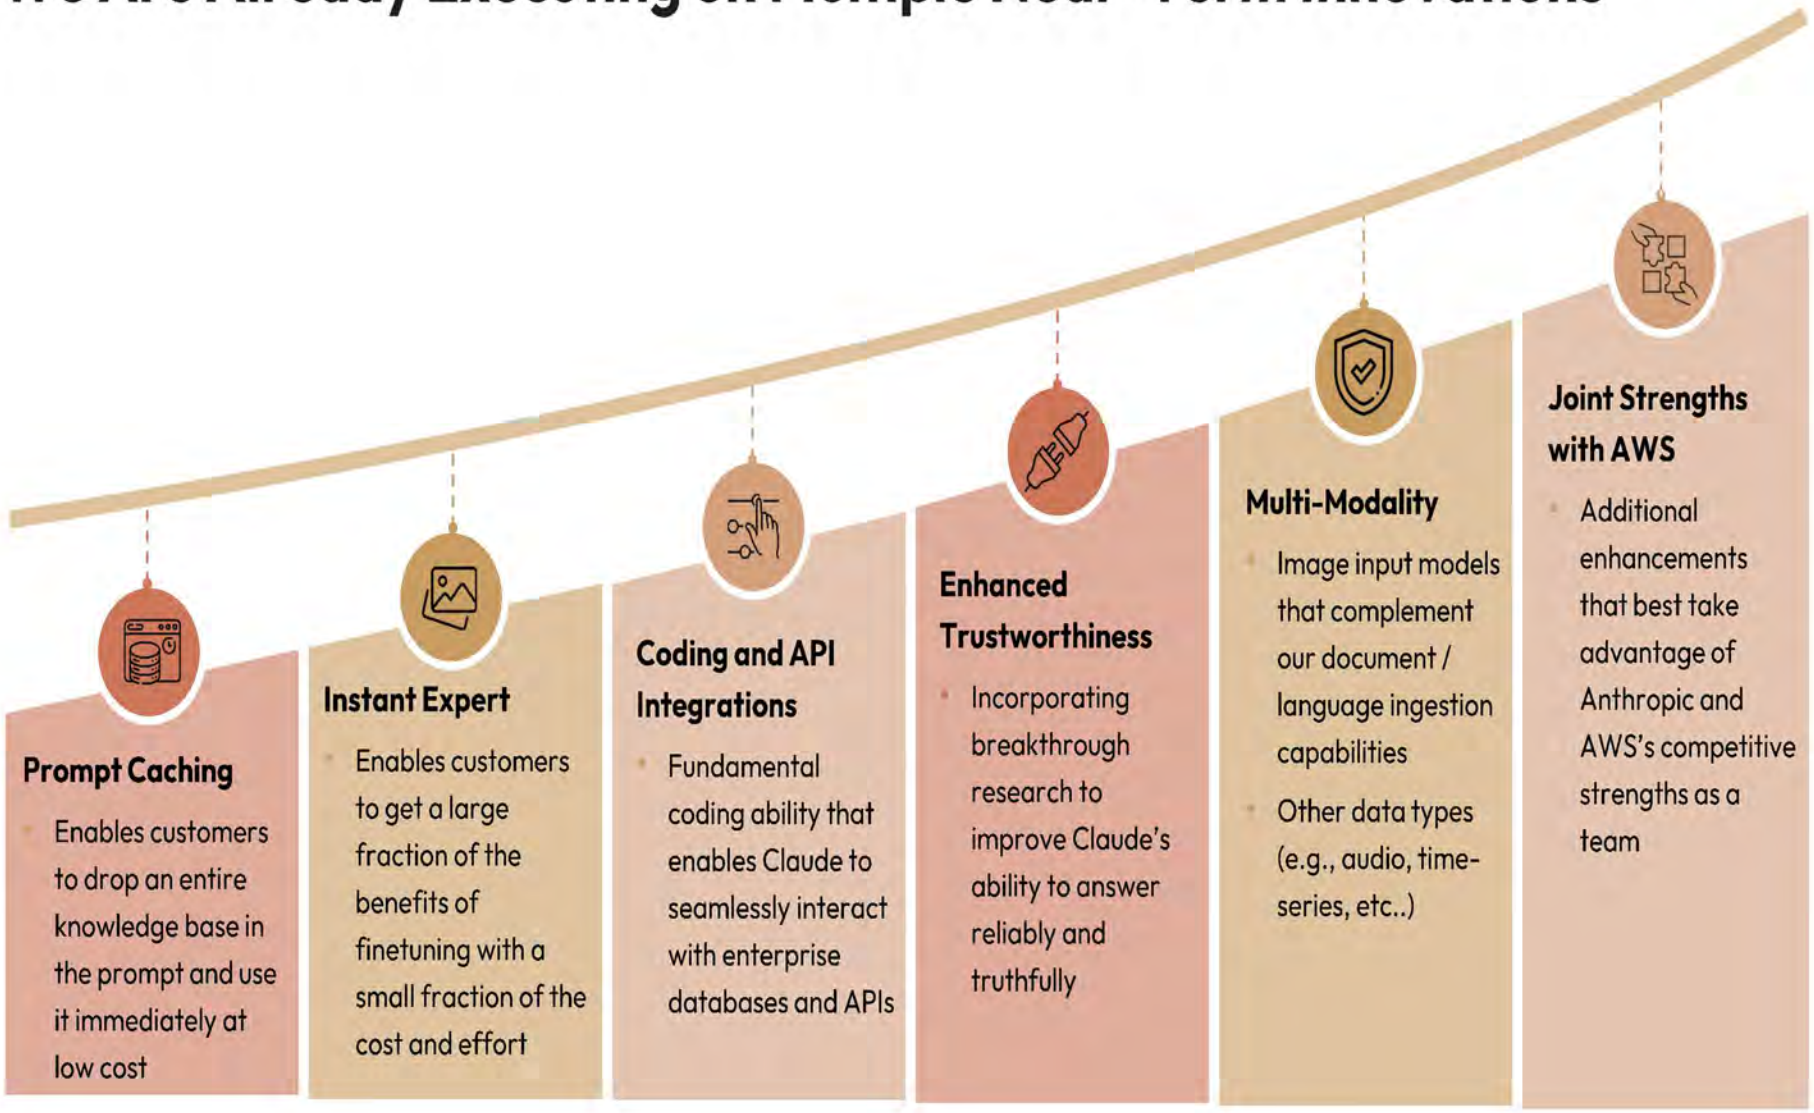
Algorithmic Innovations + Unique Hardware Partnerships

Claude 3

the most useful and trustworthy AI in the world

Highly Confidential - Attorney Eyes Only

We Are Already Executing on Multiple Near-Term Innovations



Prompt Caching

- Enables customers to drop an entire knowledge base in the prompt and use it immediately at low cost

Instant Expert

- Enables customers to get a large fraction of the benefits of finetuning with a small fraction of the cost and effort

Coding and API Integrations

- Fundamental coding ability that enables Claude to seamlessly interact with enterprise databases and APIs

Enhanced Trustworthiness

- Incorporating breakthrough research to improve Claude's ability to answer reliably and truthfully

Multi-Modality

- Image input models that complement our document / language ingestion capabilities
- Other data types (e.g., audio, time-series, etc..)

Joint Strengths with AWS

- Additional enhancements that best take advantage of Anthropic and AWS's competitive strengths as a team

Highly Confidential - Attorney Eyes Only

The Transformation of the Economy Has Only Just Begun

Three major priorities enable Anthropic to build transformative AI that is safe AI for all

Significant Model Advances

- Increase in data and parameters in future versions of Claude
- Multi-modality capability
- Perfect verticalization and personalization through proprietary data
- Continue leading breakthrough AI research centered on highest performance, scalability, and safety standards

Building GTM and Partnerships

- Strengthen distribution and build robust GTM motion across enterprise, developers, and customers
- Select strategic partners to expand distribution, partnering with large enterprises through nascent AI startups
- Large deals to be announced with leaders across industry verticals

Scaling Talent

- Continue onboarding the industry's top researchers and engineers drawn to working on the most capable AI models and a safety- aligned mission
- Further expand non-engineering sales and marketing team to support next phase of growth

Highly Confidential - Attorney Eyes Only



Highly Confidential - Attorney Eyes Only