



together.ai

\$5.6

\$0.6

Secondary Valuation (\$b)

Total Capital Raised (\$b)

Company Description

Together AI offers an end-to-end AI acceleration cloud platform that enables developers to train, fine-tune, and deploy open-source generative artificial intelligence models - all from a single interface - with strong performance, control, and cost-efficiency. Together AI's suite brings rapid model iteration from prototype to production - all with transparent, open-source tooling and scalable infrastructure..

Last round - Raise - Deal type - Date	\$305,000,000 Primary Dec 2024
Year founded	2022
# of employees	200
HQ country	United States
Industry	AI Infrastructure
Business status	Generating Revenue
Website	www.together.ai

Performance - Time Series, Trailing (annualized after 1 year)

3mo	6mo	YTD	1yr	3yr	5yr	10yr	Founded
22.74%	52.19%	69.35%	60.05%	304.59%	---	---	239.12%


Performance - Time Series, Annual - Last 5 years

2021	2022	2023	2024	2025
---	---	801.98%	330.14%	72.70%

Performance (cumulative) - Financing Round w/ Date, Valuation

Series C	Series B	Series A	Seed	Pre-Seed
Dec '24	Feb '24	Nov '23	Feb '23	Sep '22
\$3.40b	\$1.27b	\$0.57b	\$0.105b	\$0.004b
69.95%	270.17%	649.85%	3007.20%	6181.11%

Performance - Growth of \$10,000





Vipul Ved Prakash (Together AI CEO) Interview:

CEO and Co-Founder Together AI, Vipul Ved Prakash w/ Bucky Moore.

Jul 22, 2024 | 64:48 minutes

Click icon below or [here](#) to watch



Together AI on Inference, Optimization, & Hardware

Oct 27, 2025 | 20:50 minutes

Click icon below or [here](#) to watch




Key Products:

Together AI's offerings cover an AI developer's needs end-to-end, from prototyping code in the sandbox to deploying fully fine-tuned models or running massive inference workloads.

- Inference. Serverless or dedicated API for 200+ open-source models, delivering up to 4x faster inference than alternatives at ~11x lower cost than GPT-4o

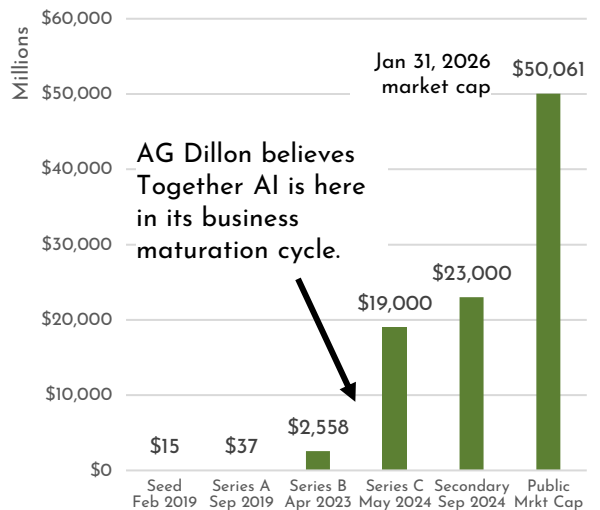
Revenue: 2022 to 2025E (\$m, Sacra & The Information)



Competitive Analysis: CoreWeave

CoreWeave offers a GPU cloud offering in direct competition with Together AI. AG Dillon estimates CoreWeave is 12 to 18 months ahead of Together AI from a business maturity perspective. We also believe that Together AI's revenue/valuation will follow a similar path and the management team is executing a near similar business plan.

CoreWeave Valuation by Round





Key Products (continued):

- Fine-Tuning. Offers full and LoRA fine-tuning (plus Direct Preference Optimization) via an easy API, supporting long contexts (up to 32K tokens) and downloadable model checkpoints
- GPU Clusters. Scalable clusters (16 → 100K+ GPUs) with top-tier NVIDIA hardware (GB300, GB200, B200, H200, H100), optimized software, and interconnects for ~24% faster training and ~75% faster inference
- Instant GPU Clusters. Self-service, bare-metal NVIDIA H100 clusters (up to 64 GPUs) deployable in minutes with Kubernetes or Slurm support
- Code Sandbox. On-demand VM environments tailored for AI development, execution, and experimentation
- Code Interpreter. Execute and test LLM-generated code in a secure, managed runtime environment
- Together Chat. A chat interface built on top of their open-source model ecosystem for interactive dialogue and testing
- "Which LLM to Use" Tool. A helper utility for selecting the most suitable model from their library based on use-case requirements.

Revenue Model:

Together AI operates a usage-based, multi-tiered revenue model tailored to developers and enterprises:

1. Per-token & per-minute inference billing
 - Serverless API charges per 1M tokens processed (both input & output), priced by model size; for example, \$0.10-\$3.50 depending on model complexity.
 - Dedicated GPU endpoints (like Nvidia A100, H100, H200) are billed per minute or hour; e.g. H100 is \$3.36/hr, H200 \$4.99/hr
2. Fine-tuning & training fees
 - Pricing is token-based: LoRA and full-fine-tuning cost \$0.48-\$3.20 per million tokens depending on model size, with Direct Preference Optimization (DPO) ~2.5x higher

Batch inference discounts & enterprise tiers

- Batch inference comes with ~50% discount to encourage bulk workloads
- Tiered service plans - Build, Scale, and Enterprise - add SLA guarantees, private endpoints, HIPAA compliance, and dedicated support, catering to growth and enterprise needs

4. Value-added enterprise offerings

- Custom GPU clusters (e.g. 36k-GPU Blackwell clusters) rented per-minute/hour with managed support
- Platform tools like Code Sandbox, Code Interpreter, and LLM-selection utilities boost stickiness and usage

Customers:

800,000 AI developers as of Jan 2026, +78% year over year. Enterprise customers include Salesforce, Zoom, Washington Post, SK Telecom, ElevenLabs.

Notable Investors:

Kleiner Perkins, General Catalyst, Salesforce Ventures, Nvidia, Coatue, NEA, Greycroft

Founder:



Vipul Ved Prakash
Together AI, CEO

Co-founded Cloudmark (anti-spam) and Topsy (social-media search, acquired by Apple for over \$200 M), and served as Senior Director of Engineering at Apple after the acquisition.

Prakash began his tech journey in India. Built the open-source anti-spam tool Vipul's Razor (MIT TR35 Innovator, 2003)

Holds a B.Sc. in computer science from St. Stephen's College in Delhi.



Summary of Fund Terms

Fund Name	AGDF28 LP - AG Dillon Together AI Pre-IPO Stock Fund
Fund Description	The Fund may gain exposure to Together AI stock by directly purchasing from the company, purchasing from existing shareholders, or investing into special purpose vehicle funds that solely own Together AI stock.
SEC Exempt Reporting Advisor	AG Dillon Asset Management LP
General Partner	AG Dillon LLC
Investor Suitability	Qualified clients only (\$2.2m net worth excluding primary residence)
Minimum investment	\$2,500 Due to limited availability of allocations in the Fund, the General Partner reserves the right to accept or decline subscriptions for any reason, in its own discretion, including on the basis of commitment size.
Closings	Initial closing of the Fund will occur on such date as the General Partner determines.
Term	Ten-year term, subject to two 1-year extensions at the General Partner's discretion. The term can be further extended to permit the orderly liquidation, sale, exchange, or other disposition of the Fund's remaining investments.
Liquidity	The earliest of the following; <ul style="list-style-type: none"> • IPO • Acquisition for cash and/or acquirer's stock • Liquidity option available on the Fund's 5th anniversary
Capital Contributions	100% called at closing.
Investment Period	12 months from closing, subject to one 1-year extension at the General Partner's discretion.
Fund Administration	NAV Consulting – www.navconsulting.net
Legal	Morgan, Lewis & Bockius LLP - www.morganlewis.com
Auditor	Berkower LLC - www.berkower.io
Institutional Traders	The Fund will access pre-IPO stocks through institutional traders. AG Dillon has 50 institutional trader relationships.
Expenses	The Fund will bear all expenses, including but not limited to: <ul style="list-style-type: none"> • Broker commissions • Fund creation and administration expenses • Legal expenses • Variable fees for wire and blue sky filing
Management Fee	0.30% management fee*. Management fee based on each investor's capital commitment to the Fund at closing date. (*charged up front for full 10 year term of the Fund)
Performance Based Compensation	10% carried interest

*** NOTE: AG Dillon ("AGD") is not affiliated with Together AI. Together AI may require company approval for purchases (aka transfers). AGD has not been pre-approved by Together AI to purchase their stock. AGD purchases pre-IPO stocks in the secondary market and may gain exposure by directly purchasing the stock (on the company's capitalization table) and/or through a third-party fund (aka special purpose vehicle, or SPV).**



General Partner Bio



Aaron Dillon

AG Dillon & Co

Managing Director

1167 2nd Ave, Suite 2N
NYC 10065

e. aaron.dillon@agdillon.com

c. +1 347 642 2640

Dillon's pre-IPO stock experience

- General partner and investment advisor to 25 pre-IPO stock venture capital funds with \$221 million in assets under management
- Funds available for purchase at Schwab, Fidelity, CNB
- Pre-IPO stock trading stats/averages; 16 offers received per trade, 6.5% price premium vs last primary round, 4.5% price improvement vs median offer, 15.8% best vs worst offer spread
- 5,500%+ realized return on SoFi pre-IPO stock investment, 4,000%+ unrealized return on KraneShares ETFs pre-IPO stock investment

Dillon's fund management and passive investing experience

- Co-founded KraneShares ETFs and ran day-to-day operations for SEC 1940 Act funds. KraneShares has \$10 billion in assets under management.
- Head of wealth management division at FTSE Russell Indices and constructed indices for ETFs, direct indexing, and passive SMAs. FTSE Russell is the largest index company globally with \$16 trillion benchmarked.

Dillon's regulatory and RIA/IBD servicing experience

- Managed investment product platforms at Morgan Stanley Wealth Management and TD Ameritrade and ensured compliance with SEC and Finra regulation while understanding and meeting financial advisor needs. Morgan Stanley Wealth Mgmt is the largest wealth manager in the US with \$5 trillion in client assets. The TD Ameritrade mutual fund and ETF business manages \$160 billion of client assets.

Dillon's companies and roles – see [LinkedIn bio](#) for more:

- SoFi – Managing Director – Wealth Management, Galileo division
- London Stock Exchange / FTSE Russell Indexes – Managing Director, US Wealth Mgmt
- KraneShares ETFs – Co-founder, Managing Director
- TD Ameritrade – Director, Head of Mutual Fund & ETFs
- Morgan Stanley Wealth Management – VP, Investment Products & Managed Accounts



*** NOTE: AG Dillon ("AGD") is not affiliated with Together AI. Together AI may require company approval for purchases (aka transfers). AGD has not been pre-approved by Together AI to purchase their stock. AGD purchases pre-IPO stocks in the secondary market and may gain exposure by directly purchasing the stock (on the company's capitalization table) and/or through a third-party fund (aka special purpose vehicle, or SPV).**

Disclosures

© 2026 AG Dillon LLC and its affiliates ("AG Dillon & Co"). All rights reserved.

All information is provided for information purposes only. All information and data contained in this publication is obtained by the AG Dillon & Co, from sources believed by it to be accurate and reliable. Because of the possibility of human and mechanical error as well as other factors, however, such information and data is provided "as is" without warranty of any kind. No member of the AG Dillon & Co nor their respective directors, officers, employees, partners or licensors make any claim, prediction, warranty or representation whatsoever, expressly or impliedly, either as to the accuracy, timeliness, completeness, merchantability of any information or of results to be obtained from the use of the AG Dillon & Co Products or the fitness or suitability of the AG Dillon & Co Products for any particular purpose to which they might be put. Any representation of historical data accessible through AG Dillon & Co Products is provided for information purposes only and is not a reliable indicator of future performance.

No responsibility or liability can be accepted by any member of the AG Dillon & Co nor their respective directors, officers, employees, partners or licensors for (a) any loss or damage in whole or in part caused by, resulting from, or relating to any error (negligent or otherwise) or other circumstance involved in procuring, collecting, compiling, interpreting, analyzing, editing, transcribing, transmitting, communicating or delivering any such information or data or from use of this document or links to this document or (b) any direct, indirect, special, consequential or incidental damages whatsoever, even if any member of the AG Dillon & Co is advised in advance of the possibility of such damages, resulting from the use of, or inability to use, such information.

No member of the AG Dillon & Co nor their respective directors, officers, employees, partners or licensors provide investment advice related to AG Dillon & Co Products and nothing contained in this document or accessible through AG Dillon & Co Products, including statistical data and industry reports, should be taken as constituting financial or investment advice or a financial promotion.

Past performance is no guarantee of future results. Charts and graphs are provided for illustrative purposes only. Valuation returns shown may not represent the results of the actual trading of investable assets. Certain returns shown may reflect back-tested performance. All performance should be considered back-tested performance. Back-tested performance is not actual performance, but is hypothetical. The back-test calculations are based on the same methodology that was in effect when the Product was officially launched. However, backtested data may reflect the application of the Product methodology with the benefit of hindsight, and the historic calculations of a Product may change from quarter to quarter based on revisions to the underlying economic data used in the calculation of the Product.

This publication may contain forward-looking assessments. These are based upon a number of assumptions concerning future conditions that ultimately may prove to be inaccurate. Such forward-looking assessments are subject to risks and uncertainties and may be affected by various factors that may cause actual results to differ materially. No member of the AG Dillon & Co nor their licensors assume any duty to and do not undertake to update forward-looking assessments.

No part of this information may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without prior written permission of the applicable member of the AG Dillon & Co. Use and distribution of the AG Dillon & Co data requires a license from AG Dillon & Co.