Forwards or Backwards

https://forwardorbackwards.com



Nvidia, a Timeline: From Graphics Cards to AI

Watch on YouTube: https://www.youtube.com/watch?v=gz5I0GNhOh4

Nvidia did not pivot into artificial intelligence. This video traces the thirty-three years that came before AI needed it: a graphics-card company founded on 5 April 1993 by Jensen Huang, Chris Malachowsky and Curtis Priem; the 1999 GeForce 256, marketed as 'the world's first GPU'; and CUDA, launched in November 2006 to open GPUs to general-purpose scientific computing, years before anyone outside the company thought about artificial intelligence.

What followed: three new chip generations aimed at the AI data-centre market - Hopper (2022), Blackwell (2024), Rubin (2026) - and a market-capitalisation climb with no precedent: $1 trillion (30 May 2023, first US chipmaker), $2 trillion (23 February 2024), $3 trillion (5 June 2024, passing Apple), $4 trillion (9 July 2025, first company ever), $5 trillion (29-30 October 2025, first company ever), and

$5.5 trillion (13 May 2026, a new record). Fiscal 2026 revenue: $215.9 billion, $193.7 billion of it from Data Center. 42,000 employees as of January 2026.

Chip-architecture technical detail - node names, die sizes, precision-format performance figures - belongs to this channel's companion video and is not repeated here. This video's subject is the company's history, not one generation's engineering.

Educational documentary. Not financial or investment advice.

Chapters

1. One idea, kept for thirty years
2. A restaurant booth, 1993
3. Coining a name, 1999
4. A door nobody was aiming at, 2006
5. The competition that changed the question, 2012
6. Three names, one direction: 2022 to 2026
7. What the stock market decided, 2023 to 2026
8. The scale behind the number
9. What thirty years of one idea adds up to

1. One idea, kept for thirty years



Most companies that become this large changed what they were doing to get there. Nvidia did not.

This video covers one company's history across more than three decades. It starts with a graphics-card maker, founded in 1993. It moves through a programming platform, built in 2006 for reasons that had nothing to do with artificial intelligence. And it reaches 2012, when outside researchers turned that platform toward a different problem entirely. The idea underneath all of it never changed.

That idea is simple to state and hard to build. Instead of doing one large calculation at a time, do thousands of small ones at once, side by side.

A graphics card needs that, because every pixel on a screen can be worked out independently of every other pixel. It turns out a great deal of artificial intelligence needs exactly the same thing. Nvidia spent thirteen years building the first kind of chip before anyone outside the company understood it was also the second kind.

This video follows that company from a restaurant booth in California to becoming, for a time, the most valuable public company on earth.

2. A restaurant booth, 1993



Nvidia's own history states the founding plainly. On the fifth of April, nineteen ninety-three, three engineers started a company with a vision to bring 3D graphics to gaming and multimedia.

The three founders were Jensen Huang, Chris Malachowsky and Curtis Priem. Huang has said in interviews that the founding conversation happened over breakfast at a Denny's restaurant in East San Jose, California — a detail this video reports as his own account, not as an independently documented fact.

"A vision to bring 3D graphics to the gaming and multimedia markets" — those are Nvidia's own words for what the company set out to do. It is the founding mission stated in the company's own official history, unchanged in the decades since.

San Jose is where that first conversation is said to have taken place. Nvidia's headquarters today sits a short distance away, in Santa Clara. The company has stayed in the same corner of California for its entire history, even as what it makes changed almost beyond recognition.

Three engineers with an idea about video games is not, on its own, an unusual story in that part of the world in the early nineteen nineties. What came next is where it starts to be.

3. Coining a name, 1999



For its first six years Nvidia made graphics chips the way its competitors did, and it was not yet the largest company doing it.

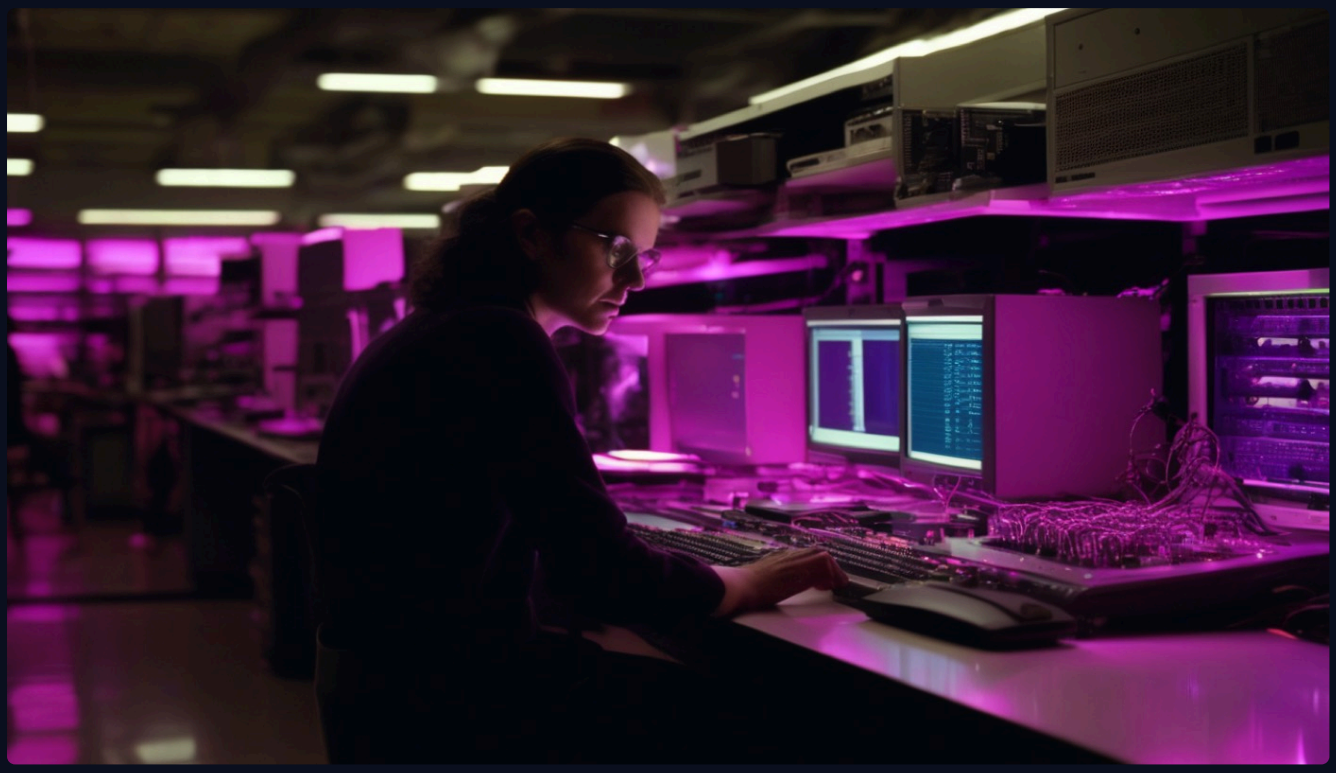
That changed with a product launched on the eleventh of October, nineteen ninety-nine: the GeForce 256, Nvidia's breakthrough consumer graphics chip.

Nvidia marketed it as "the world's first GPU" — graphics processing unit — a term the company coined and attached to this specific chip. The claim is about the name, not about the idea behind it: chips that accelerated graphics already existed before this one.

What Nvidia did was give the category a name that stuck, and build the chip that carried it. A GPU renders an image by working on many small pieces of that image at the same time, rather than one large scene calculated in order — the same underlying idea the company had started with.

That naming choice mattered more than it looked like it would, because the letters "GPU" would end up on a great deal more than a graphics card.

4. A door nobody was aiming at, 2006



For years after that first GPU, a GPU did one job: draw graphics, and nothing else. In November 2006, Nvidia opened it up to do almost anything.

The platform was called CUDA. Before it, a GPU's parallel design could only be reached through graphics-specific instructions. CUDA let a programmer write ordinary code that ran on that same hardware, for any calculation shaped the right way — not just images.

Nvidia built CUDA to sell graphics cards into scientific and technical computing, a market far smaller than gaming.

Nothing about CUDA's launch mentioned artificial intelligence. It was a tool for scientists modelling weather, physics and chemistry on hardware designed for video games. Its inventors could not have known what would be running on it years later.

That gap — a tool built for one purpose, waiting for the purpose that would actually matter — is the hinge this whole story turns on.

5. The competition that changed the question, 2012



On the thirtieth of September, 2012, a team from the University of Toronto entered an image-recognition competition and won it by a margin nobody had seen before.

The team was Alex Krizhevsky, Ilya Sutskever and Geoffrey Hinton, and their neural network became known as AlexNet. They trained it on two ordinary Nvidia gaming cards, the GTX 580, using the general-purpose programming CUDA had opened up years earlier.

AlexNet was not built by Nvidia, funded by Nvidia, or aimed at Nvidia's business in any way. It was three outside researchers, using a chip built for video games, to solve a problem in computer vision Nvidia had not set out to solve.

Nvidia's own history calls this the moment that "sparked the era of modern AI" — the company's own words for it. That short sentence is the clearest link between hardware built for scientists and what the world would later call the AI boom.

After 2012, the field of artificial intelligence research did not gradually notice GPUs. It turned toward them all at once.

6. Three names, one direction: 2022 to 2026



What AlexNet started did not stay a research curiosity. It became the reason data centres — buildings full of computers doing calculations, rather than storing files — started filling with Nvidia hardware.

Nvidia named three chip generations for this new market. Hopper was announced at the company's GTC conference on the twenty-second of March, 2022. Blackwell was announced in 2024. Rubin was announced in 2026. Each one was built to train and run the kind of AI system that grew out of AlexNet's approach.

This video does not explain how those chips are built or what their internal numbers mean. A separate video on this channel covers that ground in technical depth — the manufacturing, the chip-naming conventions and the physical limits involved. Here, the three names matter only as three dated steps in one direction.

Three architectures in four years is a fast pace for any hardware company. What happened to Nvidia's value over the same stretch was faster still.

7. What the stock market decided, 2023 to 2026



A company's stock-market value is not a fixed fact about it — it is a snapshot, taken fresh every trading day, of what buyers and sellers agree it is worth at that moment. Nvidia's snapshots became a story of their own.

On the thirtieth of May, 2023, Nvidia became the first American chipmaker to reach a trillion dollars in market value. On the twenty-third of February, 2024, it reached two trillion. On the fifth of June, 2024, it passed three trillion and overtook Apple as the world's second most valuable public company. On the ninth of July, 2025, it became the first company ever to reach four trillion. On the twenty-ninth of October, 2025, it became the first to reach five trillion. On the thirteenth of May, 2026, it set a new record at five and a half trillion, ahead of both Alphabet and Apple at the time.

Every one of those figures is a snapshot on a single date, not a permanent valuation. The market moves daily, in both directions, and these are the milestones on the way up.

8. The scale behind the number



A market-value figure is what buyers think a company is worth. Nvidia's own published results are what the company actually took in.

For the fiscal year that ended in late January 2026, Nvidia reported full-year revenue of $215.9 billion, up 65 per cent on the year before. Of that, $193.7 billion came from its Data Center segment – the part of the business selling chips for AI training and use, rather than for gaming.

That one segment alone was, by then, the great majority of everything the company sold.

Nvidia reported 42,000 employees as of the twenty-fifth of January, 2026. A company that started with three people in a restaurant booth had, thirty-three years later, a workforce large enough to fill a small town, and it was still growing.

Revenue and headcount are the numbers behind the market-value story – proof that the buying and selling reflected an actual business, not only sentiment.

9. What thirty years of one idea adds up to



Put the pieces of this video together, and a specific kind of company emerges — one that did not chase artificial intelligence, but had already built what it needed.

A graphics-card maker built a general-purpose parallel-computing platform, for scientists. Outside researchers later used it to win an image-recognition contest. The field of AI turned toward the hardware that made that win possible. The company that built it rode that turn to becoming, for a time, the most valuable public company on earth.

None of that required Nvidia to predict what was coming.

It required the company to keep building the same kind of chip, for the same underlying reason, for long enough. A different field eventually needed exactly that shape of computation, and found the tool already built. That is not the story of a lucky pivot. It is the story of one idea, kept for thirty years, meeting the moment it turned out to be built for.

That is where this video ends, and where the company's next chapter — still being written today — begins.

Sources and credits

Primary sources

- Nvidia, 'Our History: Innovations Over the Years', nvidia.com/en-us/about-nvidia/corporate-timeline - founding date (5 April 1993), founders, 1999 GPU claim, 2006 CUDA launch, 2012 AlexNet framing.
- Nvidia, 'NVIDIA Announces Financial Results for Fourth Quarter and Fiscal 2026', nvidianews.nvidia.com, published 25 February 2026 - full-year revenue $215.9 billion (up 65%), Data Center segment revenue $193.7 billion (up 68%).
- Nvidia developer blog, 'NVIDIA Hopper Architecture In-Depth', 22 March 2022, and 'CUDA Refresher: Reviewing the Origins of GPU Computing' - Hopper announcement date, CUDA background.
- Krizhevsky, Sutskever, Hinton, 'ImageNet Classification with Deep Convolutional Neural Networks', NeurIPS 2012 proceedings - AlexNet, authors, two Nvidia GTX 580 GPUs used for training.
- TweakTown, PC Gamer, VideoCardz, 25th-anniversary coverage of the GeForce 256, cross-checked - release date 11 October 1999.
- CNBC, 30 May 2023 ($1 trillion); CNN/CNBC/Forbes/Axios, 23 February 2024 ($2 trillion); CNBC/Al Jazeera/AIBusiness, 5-6 June 2024 ($3 trillion); CNBC, 9 July 2025 ($4 trillion); CNBC, 30 October 2025 ($5 trillion); Forbes, 13 May 2026 ($5.5 trillion) - market-capitalisation milestones, each cross-checked across at least two outlets.
- CNBC, 30 October 2025, and NBC New York (Denny's founding story, cross-checked across outlets, reported as Huang's own account, not an independently documented fact).
- Nvidia, Form 10-K for fiscal 2026 (period ended 25 January 2026), sec.gov - approximately 42,000 employees.

Not regulated financial advice.