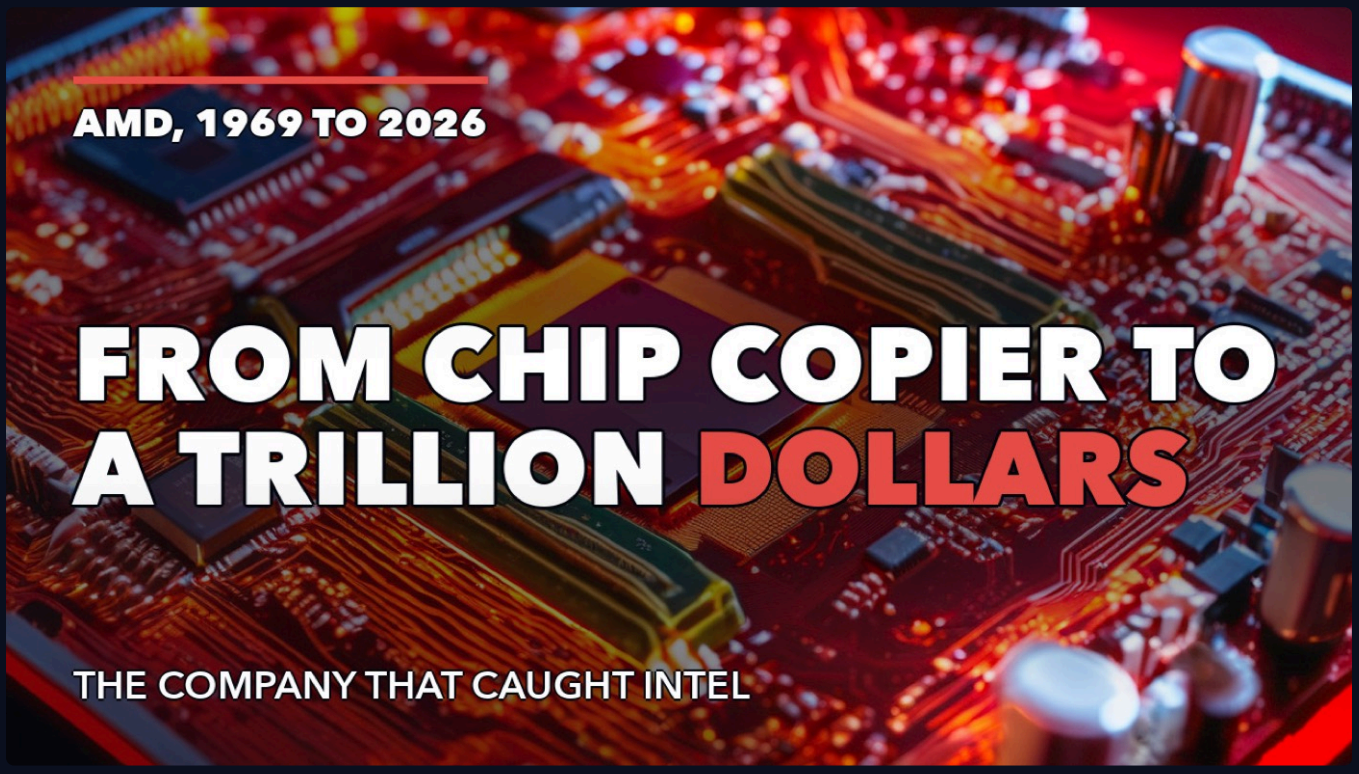
Forwards or Backwards

https://forwardorbackwards.com



AMD, a Timeline

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

AMD began in 1969 as eight men from Fairchild Semiconductor with a living room for an office. In September 2026 it was worth more than a trillion dollars. This video follows the whole story, in order, with every date and figure on screen.

How AMD grew by making other firms' chips as a licensed second source, and how IBM's PC made it Intel's partner in February 1982. How Intel shut AMD out of the 386, and how eight years of arbitration and court cases ended in the settlement of 11 January 1995. NexGen and the K6, and the Athlon that shipped at one gigahertz on 6 March 2000, two days before Intel's. AMD64, the 64-bit extension to x86 that Opteron launched on 22 April 2003 and Intel adopted in 2004.

Then the fall: the $5.4 billion purchase of ATI in 2006, the write-downs and the losses of 2007 and 2008, the factories split off as GlobalFoundries in 2009, the European Commission's fine on Intel and

Intel's $1.25 billion settlement with AMD, and the Bulldozer years. Lisa Su became chief executive on 8 October 2014; in July 2015 the shares closed at $1.62. Zen, Ryzen (2 March 2017) and EPYC; chiplets made by TSMC; Frontier, the first exascale supercomputer (May 2022).

And the AI turn: Xilinx, the Instinct MI300X, ZT Systems, and the gigawatt agreements with OpenAI (6 GW, October 2025), Meta (6 GW, February 2026) and Anthropic (up to 2 GW, July 2026). Revenue of $34.6 billion in 2025 and $11.5 billion in the second quarter of 2026, and a market value above $1 trillion on 21 September 2026. Finally the risks: Nvidia, dependence on TSMC in Taiwan, and a few very large customers.

Disclosure: this video was written with the help of Claude, an AI model made by Anthropic, one of the companies discussed. No figure appears without its date and source.

Educational documentary. Not financial or investment advice.

Chapters

1. Eight men leave Fairchild
2. The second source
3. Divorce from Intel
4. Beating Intel to a gigahertz
5. AMD64 and Opteron
6. The ATI gamble
7. Going fabless
8. Lisa Su and Zen
9. Chiplets and EPYC
10. Xilinx and the AI turn
11. Gigawatt deals
12. Forwards or backwards?

1. Eight men leave Fairchild



For more than forty years, most personal computers have run software written for one family of chip instructions. Intel designed that family. For most of that time, the other company making chips that could run it was AMD.

This video follows AMD from the start. How it began by copying other firms' chips. How it fought Intel for the right to keep doing so, and then beat Intel at its own game. How it lost billions of dollars, and saw its shares fall below two dollars. And how it became one of the companies building the machines behind artificial intelligence.

The story starts in California, at a firm called Fairchild Semiconductor. Fairchild trained a generation of chip engineers, and many of them left to start companies of their own.

Fairchild was based in Mountain View. Two of its founders, Robert Noyce and Gordon Moore, had left a year earlier to start Intel, now based in Santa Clara. AMD's own first rooms were in Santa Clara too, and its headquarters later settled in Sunnyvale. The whole area would come to be called Silicon Valley.

One of the people who left Fairchild was its director of worldwide marketing, Jerry Sanders. He had trained as an electrical engineer, but he made his name selling chips.

On the first of May, nineteen sixty-nine, Sanders and seven colleagues from Fairchild founded Advanced Micro Devices. Their first office was the living room of one of the founders, John Carey. They soon moved to two rooms at the back of a rug-cutting company in Santa Clara.

Raising the money was hard. By Robert Noyce's account, Intel had raised its money in minutes. Sanders, a marketing man with no famous inventor beside him, found it much slower.

Years later he joked that "it took me five million minutes to raise five dollars". He added that he had been dogged about it. The line is printed by the Computer History Museum, in its account of AMD's fiftieth year.

The company made its first product that November, a chip called the Am9300. In nineteen seventy came the Am2501, the first chip AMD designed itself. By September nineteen seventy-two, AMD was a public company, selling shares at fifteen dollars each.

2. The second source



The chips AMD sold in its first years were mostly not its own designs. It made copies of chips invented by other firms, under licence or by working out how they were built.

That sounds like a weak business. In fact, it was what the biggest buyers wanted. A company building a computer did not want to depend on one supplier for a vital part. If that supplier's factory stopped, so did the computer.

The answer was called a second source. That is a second licensed maker of exactly the same chip, so a buyer is never tied to one factory. The first company earns licence fees and a bigger market. The second company gets a proven product to sell.

AMD became very good at being second. In nineteen seventy-five it produced the Am9080, a copy of Intel's popular 8080 processor that AMD had worked out by taking the original apart.

A year later, in nineteen seventy-six, AMD and Intel signed their first agreement. Intel granted AMD licences to some of its patents and copyrights. The two companies were now partners of a kind, as well as rivals.

The partnership that mattered most came from IBM. On the twelfth of August, nineteen eighty-one, IBM unveiled its Personal Computer in New York. At its heart was an Intel processor, from the family that began with a chip called the 8086.

IBM wanted a second source for that processor, and Intel needed IBM's business. So in February nineteen eighty-two, Intel and AMD signed a new agreement to share technical information. Under it, AMD could make Intel's 8086 family as a licensed second source.

That family is what people now call x86, after the 8086. It is the set of basic instructions a processor understands, and every program for the IBM PC and its many copies was written for it. Whoever could make x86 chips could sell into the fastest-growing market in computing.

The IBM PC was soon copied by other makers, and every copy needed an x86 processor. AMD could now supply them. For a few years, the arrangement suited everyone.

3. Divorce from Intel



It stopped suiting Intel. Its next processor, the 386, was a large leap, and Intel wanted to be its only maker.

According to California's Supreme Court, Intel was anxious to be the sole source for the 386. So it frustrated the agreement, by taking no more products from AMD in exchange. AMD never received the design. In April nineteen eighty-seven it went to court to force the dispute into arbitration, which is a private hearing before a referee both sides accept.

So AMD built its own 386 without Intel's help. It had to design a chip that behaved exactly like Intel's, from the outside, without copying Intel's drawings.

AMD's version, the Am386, came out in March nineteen ninety-one. It ran the same software as Intel's 386. It also meant years more in court.

In February nineteen ninety-two the arbitrator ruled for AMD. He gave AMD a permanent, royalty-free licence to the Intel rights inside the Am386, and fifteen point three million dollars in damages. Intel appealed.

That October, Intel tried a separate copyright case in a federal court. A federal judge put it on hold, and left the fight in the state courts. Sanders was not a man to understate a win.

He said that "Intel's case has been drop-kicked out of the federal court". The remark was reported by the news agency UPI, the same day.

The arbitration award went back and forth between the state's courts. A lower court confirmed it, and an appeal court overturned it. The final word came from the state's Supreme Court.

On the thirtieth of December, nineteen ninety-four, the California Supreme Court restored the award in AMD's favour. Less than two weeks later, on the eleventh of January, nineteen ninety-five, the two companies settled everything.

The settlement gave AMD a perpetual licence to the microcode in Intel's 386 and 486 processors. Microcode is the low-level program inside a processor that carries out its instructions. AMD paid Intel fifty-eight million dollars over a separate case. Intel paid AMD about eighteen million, as the arbitrator had ordered.

But the deal also drew a line. AMD agreed it had no right to copy any later Intel microcode, including the Pentium's. From now on, if AMD wanted to sell x86 processors, it would have to design them itself.

4. Beating Intel to a gigahertz



AMD's first processor designed entirely in-house was called the K5. The company wanted a faster design sooner than it could build one, so it bought one.

In October nineteen ninety-five, AMD agreed to buy NexGen, a small chip designer in California, for about eight hundred and sixty million dollars in AMD shares. The deal closed in January nineteen ninety-six. NexGen's engineers had already designed a fast processor that ran x86 software.

Their work became the K6, launched in April nineteen ninety-seven. Jerry Sanders said the day marked "the return of competition in Windows compatible processors". The first K6 chips ran at up to two hundred and thirty-three megahertz, a measure this chapter comes back to.

The bigger step came two years later. In June nineteen ninety-nine AMD unveiled the Athlon, a new processor for desktop computers. It was designed to compete with Intel's fastest chips, not its cheapest.

The race that followed was about clock speed. Every processor keeps time with an internal clock, and each tick lets it take another step through its work.

Clock speed is counted in hertz, or ticks per second. A megahertz is a million ticks a second, and a gigahertz is a thousand million. In early two thousand, both companies were racing to be first to sell a one-gigahertz chip.

On the sixth of March, two thousand, AMD announced it was shipping a one-gigahertz Athlon to computer makers. Sanders called it "our industry's equivalent of breaking the sound barrier". AMD had got there first.

Intel announced its own one-gigahertz Pentium III two days later, on the eighth of March. Two days is not much. But for a company that had begun by copying other firms' chips, being first at anything was new.

5. AMD64 and Opteron



The next race was about memory. Processors of the time worked with numbers thirty-two binary digits long, called thirty-two-bit. That set a hard limit on how much memory one program could reach directly.

Thirty-two binary digits can count to just over four billion. So a thirty-two-bit processor can point to about four gigabytes of memory, and no more. Large databases and servers, which are the computers that run networks and websites, were running into that ceiling.

The answer was sixty-four bits. Intel's plan was a new processor family for servers, called Itanium, with a new set of instructions. Existing x86 programs ran on it only slowly.

AMD took the other road. It kept the x86 instructions and stretched them to sixty-four bits. A chip built that way could run every existing program, and new sixty-four-bit programs as well. AMD called it AMD64.

The first chip to carry it was Opteron, a processor for servers, launched on the twenty-second of April, two thousand and three. The Linux operating system could use its sixty-four bits at once. Microsoft's server version of Windows followed later that summer, as a test release.

Old thirty-two-bit programs ran on it unchanged. New sixty-four-bit programs could reach far more memory. Both ran side by side on the same chip, which meant a company could upgrade its servers without rewriting its software.

Then came the moment that settled the argument. Intel, which had licensed x86 to AMD twenty years before, adopted AMD's design.

On the seventeenth of February, two thousand and four, Intel's chief executive Craig Barrett said Intel's new sixty-four-bit technology would be "software compatible with AMD's 64-bit extension technology". The company that had invented x86 was now following AMD's version of it.

AMD's desktop version, the Athlon 64, had arrived in September two thousand and three. Opteron, meanwhile, took AMD into server rooms, where Intel had long been the main supplier. For a few years, the industry's standard design for sixty-four-bit computing was AMD's.

6. The ATI gamble



Being better did not make AMD bigger. In June two thousand and five, AMD sued Intel in a federal court in Delaware, accusing it of abusing its monopoly to keep computer makers from buying AMD chips.

While that case went on, AMD made the largest bet in its history. It decided that the future chip would not be a processor alone.

On the twenty-fourth of July, two thousand and six, AMD agreed to buy ATI Technologies, a maker of graphics chips, for about five point four billion dollars. The deal closed that October. AMD paid about four point three billion in cash, and the rest in its own shares.

A graphics processor is a chip built from many small, simple cores that each do sums at the same time. They were made to draw images for games. The same design is good at any work that can be broken into many identical pieces.

AMD's plan was to put both kinds of chip under one roof. The main processor handles instructions one after another. The graphics processor handles thousands of small tasks at once. In time, AMD wanted to put both on the same piece of silicon.

The idea was sound. The timing was not. AMD had borrowed two and a half billion dollars to pay for ATI. Then its new four-core server chip, Barcelona, was found to have a flaw that could freeze a computer, and volume shipments slipped into two thousand and eight.

At the end of two thousand and seven, AMD wrote down the value of ATI by one point six billion dollars. A write-down is an admission that something bought is worth less than was paid. AMD lost three point four billion dollars that year.

In two thousand and eight it wrote ATI down by another one point one billion. It lost three point one billion more that year. Two years of losses had cost AMD more than the whole ATI purchase price. The bet that would matter most a decade later was, for now, nearly sinking the company.

7. Going fabless



For most of its life AMD had built its own chips in its own factories. In the chip business a factory is called a fab, short for fabrication plant, and it is among the most expensive buildings in industry.

Jerry Sanders was proud of AMD's fabs. In a jibe at a rival that did not own one, he said that "only real men have fabs". The Register reported the line in February nineteen ninety-nine, a few years after he first said it.

Ten years later, AMD could no longer afford to be one of them. Each new generation of chips needed a newer and costlier fab, and AMD was losing billions of dollars a year.

On the second of March, two thousand and nine, AMD split off its factories, including its fabs in Dresden, in Germany, into a new company called GlobalFoundries. The other owner was an investment company backed by the government of Abu Dhabi. GlobalFoundries became a foundry, which is a firm that makes chips designed by other companies.

AMD was now fabless, a company that designs chips and pays others to make them. It kept a stake in GlobalFoundries at first. In March twenty twelve, it handed over its remaining share and took a charge of seven hundred and three million dollars.

That same year, two thousand and nine, the long fight with Intel came to a head. The European Commission, which enforces competition law across the European Union, had been investigating Intel for years.

In May two thousand and nine, the Commission fined Intel one point zero six billion euros. It found that Intel had given hidden rebates to computer makers that bought all, or almost all, their processors from Intel. It also found that Intel had paid computer makers to halt or delay products using a rival's chips.

Six months later, on the twelfth of November, Intel settled with AMD. Intel agreed to pay AMD one point two five billion dollars, and the two signed a new five-year licence. AMD dropped its lawsuits, including the one in Delaware, and withdrew its complaints to regulators.

The European fine had a longer life. The General Court of the European Union annulled it in twenty twenty-two. The Commission imposed a smaller fine, of three hundred and seventy-six million euros, in twenty twenty-three. In December twenty twenty-five the court upheld that finding but cut the fine to two hundred and thirty-seven million. Intel has appealed again.

The money from Intel did not fix AMD's real problem, which was its products. Its next design for desktop processors, called Bulldozer, arrived in October twenty eleven.

Reviewers were disappointed. It was not fast enough to catch Intel's chips of the time. AMD changed chief executive twice in twenty eleven: Dirk Meyer left in January, and Rory Read took over in August. In twenty fourteen, the board changed again.

8. Lisa Su and Zen



On the eighth of October, twenty fourteen, AMD named a new chief executive, Lisa Su. She was an electrical engineer with three degrees from MIT. She had spent thirteen years at IBM, and she had joined AMD in twenty twelve.

When she took over, AMD's shares were worth about three dollars and twenty-eight cents. They kept falling. On the twenty-seventh of July, twenty fifteen, they had their lowest close in decades. It was one dollar and sixty-two cents.

There was pressure to move into chips for phones and tablets, where the growth was. Su decided against it.

She later told Time magazine that AMD had to "bet on what we were good at". What AMD was good at was high-performance processors, for personal computers and for servers. So the company bet on one new processor core, called Zen.

A core is the part of a processor that actually carries out instructions. A modern chip has several of them side by side. Zen was a new core designed from a clean sheet, and it had to be much faster than the one before it.

The goal Su set her engineers was about forty per cent more work done on each tick of the clock. That measure is called instructions per clock, and it says how much a core gets done on every tick. When Zen was finished, it reached fifty-two per cent.

She told PCWorld: "We not only beat our goal, we beat it by a lot." The magazine reported it shortly before Ryzen went on sale.

The first Zen processors went on sale as Ryzen, for personal computers, on the second of March, twenty seventeen. In June that year came EPYC, the version for servers, with up to thirty-two cores. AMD was back in both of the markets it had bet on.

9. Chiplets and EPYC



Zen was a new core. The next idea changed how the whole chip was built.

Chips are made on a wafer, which is a thin disc of silicon that holds many chips at once. Each chip is cut from it as a die, the small rectangle of silicon the circuits are printed on. The bigger a die, the more likely it is to contain a flaw, and a single flaw can ruin it.

AMD's answer was to build one processor from several small dies instead of one big one. Each small die is called a chiplet. In AMD's second EPYC design, eight chiplets doing the computing sit around one central die that handles memory and connections, all inside one package.

Small dies mean fewer are wasted. And the parts need not all be made the same way. The computing chiplets can come from the newest production line, and the central die from an older, cheaper one.

The newest line was no longer GlobalFoundries'. On the twenty-seventh of August, twenty eighteen, GlobalFoundries stopped developing its next generation of chipmaking. AMD moved all of that work to the Taiwan Semiconductor Manufacturing Company, known as TSMC, a foundry in Taiwan.

The second EPYC, codenamed Rome, launched in August twenty nineteen, with up to sixty-four cores. Its computing chiplets were made by TSMC. Its customers were companies that run data centres, and some of the largest science laboratories in the United States.

One of them was Oak Ridge National Laboratory, in Tennessee. It ordered a supercomputer called Frontier, built by Hewlett Packard Enterprise, with AMD processors and AMD graphics chips inside.

In May twenty twenty-two, Frontier became the first computer measured at more than one exaflop. An exaflop is a billion billion calculations every second. Frontier's score was one point one zero two exaflops, which put it top of the TOP500 list of the world's fastest machines.

In November twenty twenty-four, a second AMD-powered machine took first place. El Capitan, at Lawrence Livermore National Laboratory in California, reached one point seven four two exaflops. Its chips combine AMD processor cores and graphics cores in one package.

10. Xilinx and the AI turn



With its processors selling again, AMD started buying. The biggest purchase was Xilinx, the company that invented the field-programmable gate array, or FPGA. That is a chip whose circuits can be rearranged after it is made.

AMD agreed to buy Xilinx in October twenty twenty, in an all-share deal valued at thirty-five billion dollars. By the time the deal closed, on the fourteenth of February, twenty twenty-two, AMD's own shares had risen so far that it was valued at about forty-nine billion.

The next day, AMD's stock market value passed Intel's for the first time. At the close of trading, AMD was worth about one hundred and ninety-eight billion dollars. Intel was worth slightly less.

By then the fastest-growing market in chips was not personal computers or servers. It was artificial intelligence. Training and running large AI models needs graphics processors in huge numbers, and Nvidia sold most of them.

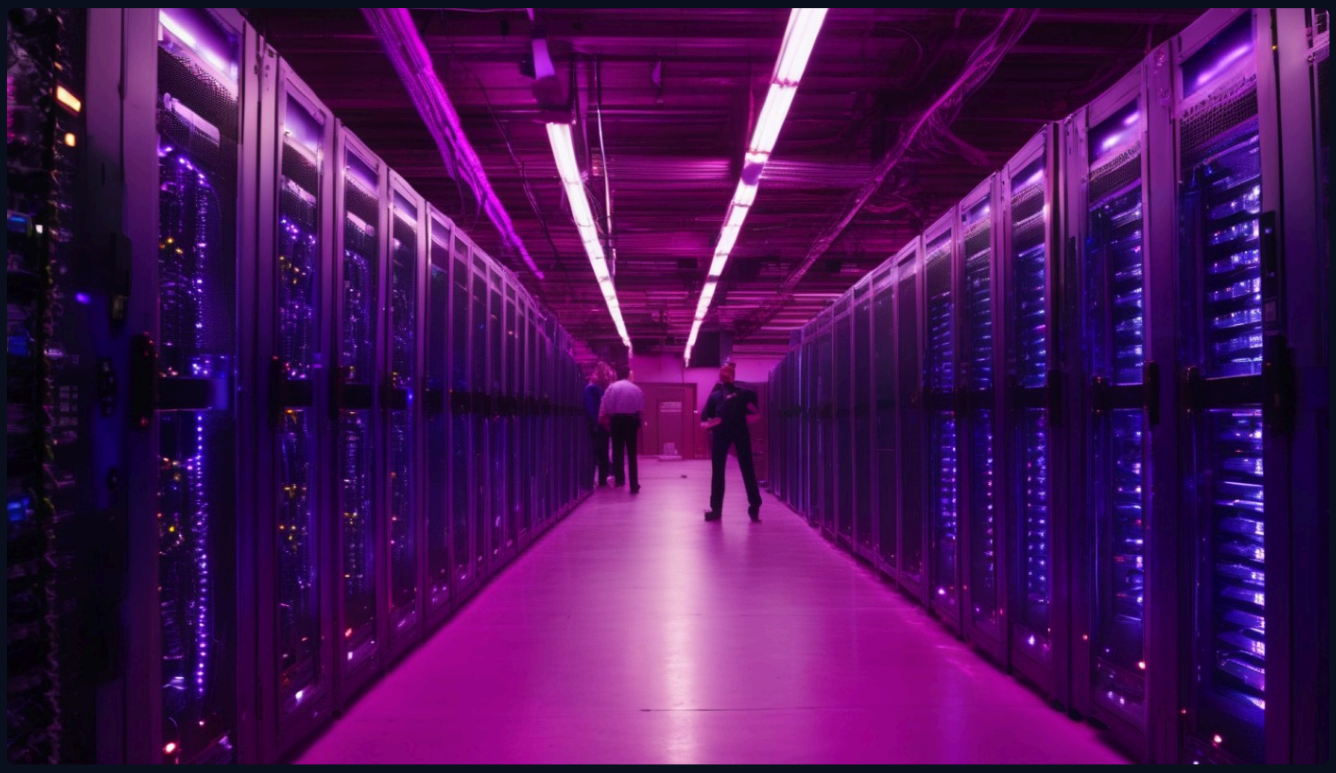
AMD's answer was the Instinct MI300X, launched in December twenty twenty-three. It is built in layers. Eight computing dies sit stacked on top of four dies that handle input and output, with eight stacks of fast memory around them, all in one package.

The memory mattered most. An AI model has to hold billions of numbers close to the chip that is working on them. The MI300X carried a hundred and ninety-two gigabytes of that memory, stacked right beside the computing dies.

In twenty twenty-four, AMD sold more than five billion dollars' worth of Instinct AI chips. That August it agreed to buy ZT Systems, a company that designs and builds whole racks of servers, for four point nine billion dollars.

AMD wanted ZT's designers, not its factories. The purchase closed at the end of March twenty twenty-five. In October, AMD finished selling ZT's manufacturing business to Sanmina for about three billion dollars. What it kept was the skill to design a complete AI machine, from the chip to the rack.

11. Gigawatt deals



By twenty twenty-five, the biggest buyers of AI chips had stopped counting in chips. They counted in power. A gigawatt is a billion watts, roughly the output of a large power station, and an AI data centre is now planned by how much of that it will draw.

On the sixth of October, twenty twenty-five, AMD and OpenAI announced an agreement for six gigawatts of AMD graphics processors. The first gigawatt, built on a new chip called the MI450, was due to start in the second half of twenty twenty-six. AMD said the deal should bring it tens of billions of dollars in revenue.

The terms were unusual. AMD gave OpenAI a warrant, which is a right to buy shares later at a fixed price. It covered up to a hundred and sixty million AMD shares, at one cent each. They vest only as OpenAI's purchases grow, and as AMD's share price reaches set targets.

Lisa Su called the partnership "a true win-win enabling the world's most ambitious AI buildout". OpenAI's chief executive, Sam Altman, said AMD's chips would help it move faster. Both spoke in AMD's own announcement.

Then came two more. In February twenty twenty-six, Meta signed for another six gigawatts, starting with a custom version of the same chip. Meta also received a warrant for up to a hundred and sixty million shares.

In July, Anthropic, which makes the Claude models, signed for up to two gigawatts, with the first gigawatt from the first half of twenty twenty-seven. This time money also ran the other way. AMD

committed to invest up to five billion dollars in Anthropic.

Set side by side, the three agreements add up to as much as fourteen gigawatts. OpenAI and Meta each signed for six. Anthropic signed for up to two. All three are to be built around Helios, which is AMD's name for a whole rack of its chips sold as one machine.

A disclosure belongs here. This video was written with the help of Claude, an AI model made by Anthropic, one of the companies in that list. Every figure here comes from AMD's own releases and filings.

The money has started to arrive. In twenty twenty-five AMD's revenue was a record thirty-four point six billion dollars. In twenty sixteen, it had been four point two seven billion. Its data centre business, which sells server processors and AI chips, made sixteen point six billion of the twenty twenty-five total.

In the second quarter of twenty twenty-six, revenue reached eleven point five billion dollars, up half on a year before. Data centre sales were six point seven billion, up a hundred and seven per cent. For the next quarter, AMD told investors to expect about thirteen billion dollars.

12. Forwards or backwards?



On Monday the twenty-first of September, twenty twenty-six, AMD's shares rose about ten per cent. They closed at more than six hundred and fifteen dollars. That put the company's market value above a trillion dollars for the first time. It joined Nvidia, Broadcom and Micron as American chip companies worth that much.

Eleven years earlier, the same shares had been worth less than two dollars. So on the numbers, this is a story that went forwards. But the same numbers carry the risks.

The first risk is Nvidia. It is still far larger, worth more than five trillion dollars on the day AMD reached one. In AI chips AMD is the second supplier, and being second is where AMD began.

The second risk is where the chips are made. Since AMD sold its factories, its most advanced chips have come from one company, TSMC, in Taiwan. AMD does not control that supply, or the politics around it.

The third risk is who is buying. Much of the growth rests on a few very large customers. Bloomberg reported that Meta alone was AMD's second-largest customer, at about five and a half per cent of revenue. And the biggest deals are tied to warrants, so AMD pays for those orders partly in its own shares.

None of those risks is new to this company. It began as a copier of other firms' chips. It fought Intel in court for years. It lost billions after buying ATI, and its shares later fell below two dollars. Each time, a bet on a new design brought it back.

The question now is different. AMD is no longer the challenger with nothing to lose. It is a trillion-dollar company whose value rests on promises of power stations' worth of chips, still to be delivered. Whether that is the high point of a comeback, or the start of something larger, is not yet in any filing.

Sources and credits

Primary sources

- Computer History Museum, David Laws, 'It's a Good Year for Asparagus', 1 May 2019 - eight Fairchild founders; Sanders's 'five million minutes' quote.
- California Supreme Court, Advanced Micro Devices v. Intel, decided 30 December 1994 - 1976 licence, February 1982 agreement, April 1987 arbitration, Am386 March 1991.
- UPI, 30 October 1992 - $15.3M award; Sanders's 'drop-kicked' quote.
- AMD Form 8-K filed 19 January 1995 - settlement of 11 January 1995.
- AMD Form 8-K filed 31 January 1996 - NexGen merger; AMD K6 launch release, 2 April 1997; AMD release, 6 March 2000; Intel release, 8 March 2000.
- The Register, 22 April 2003 (Opteron); InfoWorld, 17 February 2004 (Intel follows AMD64).
- Engadget, 24 July 2006; AMD release, 25 October 2006 (ATI); AMD results, 17 January 2008 and 22 January 2009.
- The Register, 25 February 1999 ('only real men have fabs'); AMD Q1 2009 release, 21 April 2009 (GlobalFoundries); European Commission MEMO/09/235, 13 May 2009; Intel release, 12 November 2009; AMD release, 4 March 2012; Cleary Antitrust Watch, December 2025; Courthouse News, 10 December 2025.
- AMD release, 8 October 2014 (Lisa Su); Motley Fool, 8 December 2022 ($3.28); StatMuse Money ($1.62, 27 July 2015); TIME, 10 December 2024; PCWorld, 22 February 2017.
- AMD releases, 22 February and 20 June 2017, 7 August 2019; IEEE Spectrum, 28 August 2018; Microway, 7 August 2019; TOP500, 30 May 2022 and November 2024.
- AMD 8-K, 27 October 2020 and 14 February 2022 (Xilinx); TechSpot, 16 February 2022; AMD release, 6 December 2023 (MI300X); AMD Q4 2024 results, 4 February 2025; AMD 8-Ks, 19 August 2024 and 31 March 2025 (ZT Systems); Sanmina, 27 October 2025.
- AMD 8-K, 6 October 2025 (OpenAI); AMD releases, 24 February 2026 (Meta) and 22 July 2026 (Anthropic); AMD results, 3 February 2026 and 4 August 2026.
- Investopedia, 21 September 2026; Bloomberg (Carmen Reinicke), 21 September 2026.

Not regulated financial advice.