Forwards or Backwards

https://forwardorbackwards.com



Intel, a Timeline

Watch on YouTube: https://www.youtube.com/watch?v=2zHK_gTCEpw

Intel made the chip inside most of the world's personal computers for decades. This video follows the company from its first day to September 2026.

In September 1957 eight scientists, Robert Noyce and Gordon Moore among them, left William Shockley's laboratory to found Fairchild Semiconductor. On 18 July 1968 Noyce and Moore incorporated Intel in Mountain View. It began as a memory company: the 3101 in 1969 and the 1103 dynamic memory chip in 1970. A calculator order from Busicom led to the 4004, advertised on 15 November 1971, the first microprocessor sold as a component.

Moore's prediction of 1965, revised in 1975, became the industry's schedule. The 8086 of 8 June 1978 began the x86 family, and IBM's choice of the 8088 for the PC of 12 August 1981 made it the

standard. In October 1985, under pressure from Japanese memory makers, Intel left dynamic memory and launched the 386.

Then Intel Inside (1991), the Pentium (22 March 1993) and the 1994 division flaw that cost a $475 million charge. The missed turns: Itanium, the iPhone, and the delayed 10nm process that let TSMC take the lead. Apple left for its own chips in 2020, and Nvidia took the data centre.

Pat Gelsinger's IDM 2.0 plan (23 March 2021), the CHIPS Act award of $7.86 billion, the $18.8 billion loss of 2024 and his exit on 1 December 2024. Under Lip-Bu Tan, the 2025 cuts, SoftBank's $2 billion, the US government's 9.9% stake for $8.9 billion, and Nvidia's $5 billion. And now: 18A in production in Arizona, second-quarter 2026 revenue of $16.1 billion, a share sale at $95 that raised about $23 billion, and the customer decision on 14A that will decide what Intel becomes.

No figure appears without its date and source.

Educational documentary. Not financial or investment advice.

Chapters

1. The traitorous eight
2. Memory first
3. A computer on a chip
4. Moore's Law
5. The IBM PC and x86
6. Only the paranoid survive
7. Intel Inside
8. Missed turns
9. Losing Apple, losing the lead
10. IDM 2.0 and the $18.8bn loss
11. The state takes a stake
12. Forwards or backwards?

1. The traitorous eight



In the summer of nineteen sixty-eight, two scientists walked away from a successful chip company to start another one. The company they started would make the chips inside most of the world's personal computers for decades. Then it would lose its lead, lose money on a vast scale, and end up part-owned by the American government.

This video follows Intel from its first day to September twenty twenty-six. It covers the memory chips Intel started with, and the microprocessor it invented. It covers the deal with IBM that made Intel's design the standard, the turns it missed, and the rescue that is still under way.

The story starts with a physicist named William Shockley. He shared the Nobel Prize in physics for his part in inventing the transistor, which is a tiny electrical switch. Every chip in this story is built from transistors, switching on and off.

In the mid nineteen fifties, Shockley opened a laboratory in Mountain View to make transistors out of silicon. He hired some of the best young scientists he could find. Two of them were Robert Noyce and Gordon Moore. The Computer History Museum describes Shockley as a brilliant physicist but a poor manager, and his staff soon lost patience with him.

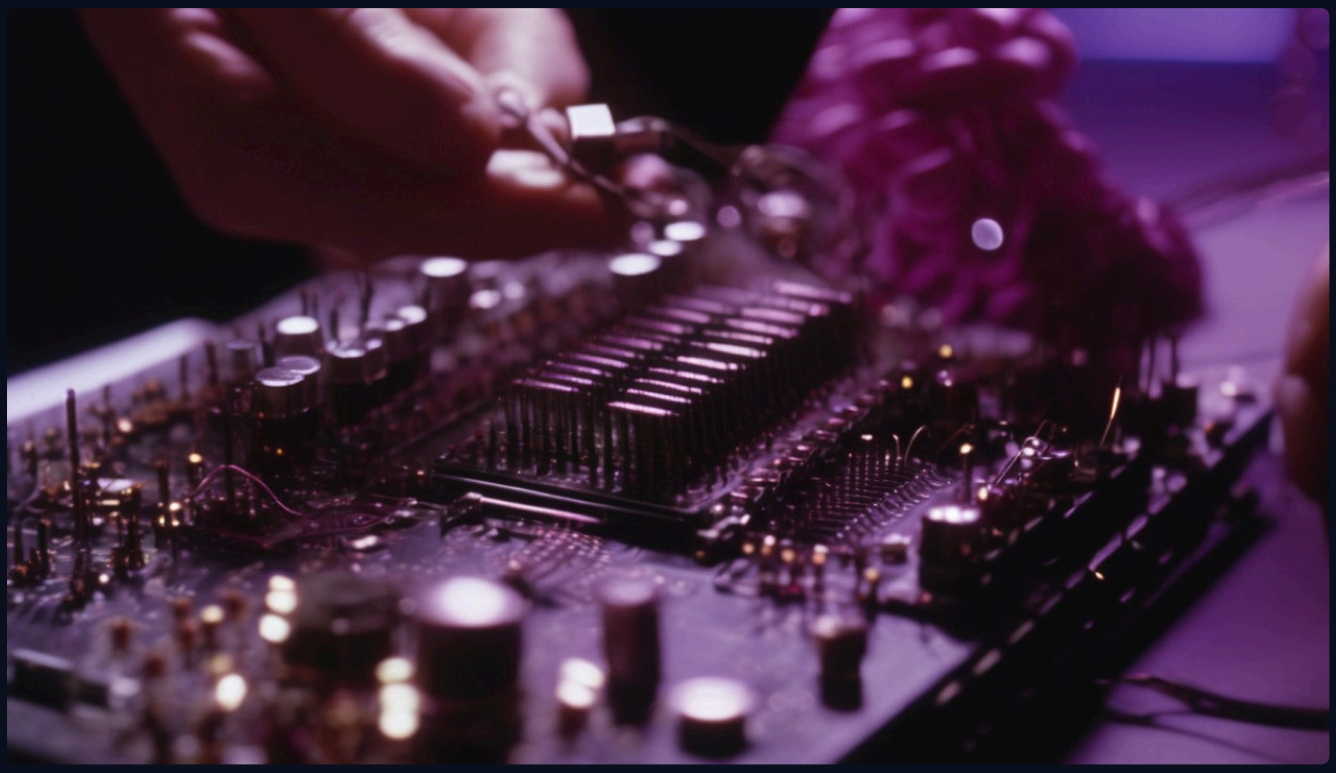
In September nineteen fifty-seven, eight of Shockley's scientists resigned together, Noyce and Moore among them. They founded Fairchild Semiconductor, in Palo Alto, a few miles from Shockley's laboratory in Mountain View. The eight became known as the traitorous eight. Intel would later start in Mountain View too, and today it is based in Santa Clara, a short drive away.

Fairchild grew into one of the most important companies in the business. It also spun off dozens of new ones, and the area around it came to be called Silicon Valley. By nineteen sixty-eight, Noyce and Moore wanted to start again, on their own terms. Intel's own history says the idea came one spring afternoon, when Moore dropped by Noyce's house while Noyce was mowing the lawn. Moore suggested that memory chips, still a new technology, could be the basis of a new company.

On the eighteenth of July, nineteen sixty-eight, the two men incorporated that company. They first called it NM Electronics, after their initials. It soon became Intel, short for integrated electronics, and they paid fifteen thousand dollars for the right to use the name. Work began in a conference room in an old Union Carbide building in Mountain View.

A young Fairchild scientist named Andy Grove joined them almost at once. The three of them would run Intel, one after another, for the next thirty years.

2. Memory first



Every computer needs somewhere to keep the numbers it is working on. In the nineteen sixties, most computers kept them in magnetic core memory, which is a grid of tiny magnetic rings threaded on wires. Each ring held one bit. A bit is a single binary digit, a one or a zero, and every number a computer stores is written in them.

Noyce and Moore were betting that silicon could do the same job better. A chip could hold many bits in a space far smaller than a grid of rings, and chips could be made in large batches.

Intel's first product came in nineteen sixty-nine. It was the three-one-oh-one, a static memory chip that kept its bits for as long as it had power. Then, in nineteen seventy, Intel released the one-one-oh-three. That was a dynamic memory chip, usually called DRAM, which stores each bit as a tiny electrical charge that has to be topped up constantly.

The dynamic design needed fewer parts for each bit, so more bits fitted on a chip. That made it cheaper, and cheap mattered more than anything else to the companies buying memory by the million.

Intel's own history says the one-one-oh-three established semiconductors as the new standard for computer memory. It replaced magnetic core memory. By the late nineteen seventies, in Andy Grove's own account, Intel stood for memory.

It is worth holding on to that, because of what comes later. Intel did not begin as a processor company. It began as a memory company, and for most of its first fifteen years memory was the heart

of its business.

The same year the one-one-oh-three arrived, Intel broke ground on its first factory of its own. And while memory paid the bills, a side project for a customer in Japan was about to change what Intel was.

3. A computer on a chip



In April nineteen sixty-nine, a Japanese calculator maker called Busicom came to Intel with an order. It wanted a set of chips to run a new family of electronic calculators.

Busicom's own design called for a large number of chips, each one wired for a single job. Intel was a small company, and building that many custom chips would have stretched it badly. So its engineers suggested something simpler.

Intel's answer was a family of four chips. One held the program, one held the working data, and one handled input and output. The fourth was a processor, the part that reads the program and carries it out, step by step. Change the program, and the same four chips could do a different job.

That fourth chip was the idea that mattered. A microprocessor is the whole of a computer's thinking part, built on a single chip. Until then, a computer's processor was made of many separate parts on circuit boards.

Four people are usually named as its designers: Ted Hoff, Federico Faggin, Stanley Mazor and Busicom's Masatoshi Shima. The first fully working chip was delivered in March nineteen seventy-one, for a Busicom printing calculator.

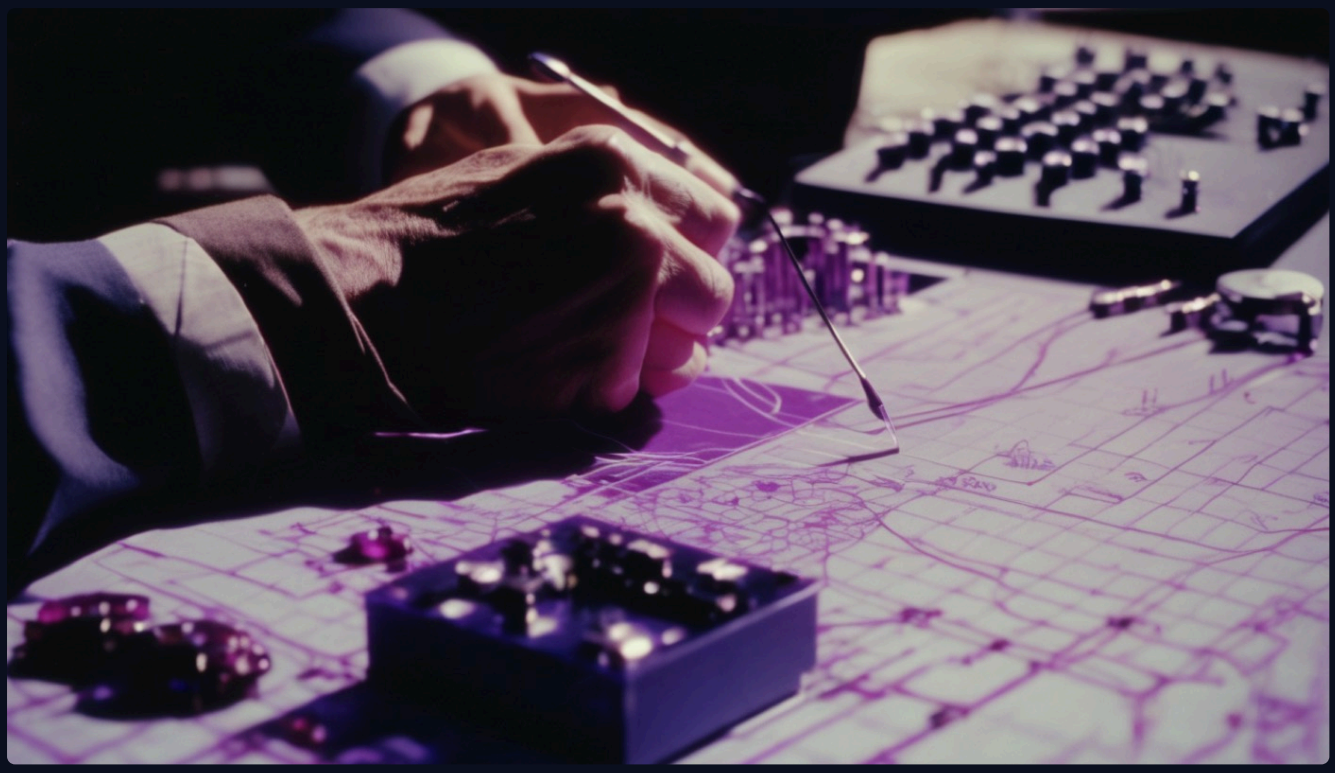
Then Intel made a decision that shaped everything after it. Busicom was struggling, and it wanted lower prices. Intel agreed to cut them, in exchange for the right to sell the chip to anyone else.

On the fifteenth of November, nineteen seventy-one, Intel advertised the four-oh-oh-four in the trade paper Electronic News. Intel puts its size at about twenty-three hundred transistors. It worked on four bits at a time and could carry out about sixty thousand operations a second.

The Computer History Museum gives a slightly lower transistor count, twenty-two hundred and fifty, so the figure on screen is a range. Either way, the whole processor of a computer now fitted on one chip.

It was not an obvious success at first. Four bits at a time is very little. But Intel had something no one else was selling yet, and it kept making the processor faster.

4. Moore's Law



Six years before the four-oh-oh-four, Gordon Moore had made a prediction that would come to rule the whole industry, and Intel above all.

In nineteen sixty-five, while still at Fairchild, Moore was asked by Electronics magazine to say where chips were heading. He looked at how many components could be put on one chip at the lowest cost per component. That number, he noticed, had been doubling every year.

Moore's article appeared on the nineteenth of April, nineteen sixty-five. He predicted the doubling would go on for ten years, and that a chip would hold about sixty-five thousand components by then. Ten years later, in nineteen seventy-five, he revised it. From then on, he said, the number would double about every two years.

That revised forecast became known as Moore's Law. It is not a law of nature. It is an observation about how quickly engineers were learning to make transistors smaller and pack more of them onto a chip.

The important thing is what the industry did with it. A forecast that everyone believes becomes a schedule. Chip makers planned their factories around it, and computer makers planned their products around it. Intel, with Moore as one of its founders, treated it as a promise it had to keep.

The line on screen shows transistors on four Intel processors. The four-oh-oh-four had about twenty-three hundred. The three-eight-six had two hundred and seventy-five thousand, and the first Pentium had three point one million.

That is a rise of more than a thousand times, in a little over two decades. Each step meant a faster, more capable processor for about the same price. And each step meant building a new generation of factory, at a larger cost than the one before.

This is the engine behind the rest of the story. For as long as Intel was the company best at keeping Moore's promise, it had the best processors, and it could charge for them. The moment it stopped keeping the promise on time, as later chapters show, everything that depended on that lead began to come apart.

5. The IBM PC and x86



The processor that made Intel's fortune arrived in the late nineteen seventies, and it was meant to be a stopgap.

The four-oh-oh-four had been followed by faster eight-bit chips. Intel's next big design was running late, and the company needed something to sell in the meantime.

On the eighth of June, nineteen seventy-eight, Intel introduced the eight-oh-eight-six. Intel's own history says it went from design to shipping in about eighteen months, with the team working nights and weekends. A year later came the eight-oh-eight-eight, a cheaper version that talked to the rest of a computer through a narrower connection.

The eight-oh-eight-six gave its name to a whole family of processors, known as x86. The name simply means the line of chips descended from it.

What turned that family into a standard was a decision made by IBM. IBM was the biggest computer company in the world, and it was late to the personal computer. It gave a small team about a year to build one.

To hit that deadline, IBM's team used parts that anyone could buy, instead of building everything itself. They chose Intel's eight-oh-eight-eight as the processor, and Microsoft supplied the operating system. They also published the technical details of the design. Other companies could build machines that ran the same software, and they did.

On the twelfth of August, nineteen eighty-one, IBM unveiled the Personal Computer, at a starting price of one thousand five hundred and sixty-five dollars. It was a success far beyond anything IBM had planned.

The machines that copied it all needed the same kind of processor, so they all needed x86 chips. Software was written for x86, which made people buy x86 machines, which led to more software being written for x86.

Intel then made sure that each new x86 processor could still run the programs written for the old ones. Nobody who upgraded had to throw their software away. That habit, kept up for decades, is why x86 became the standard of personal computing, and why it is still in most laptops and data centres today.

6. Only the paranoid survive



At the start of the nineteen eighties, Intel was still, in its own mind, a memory company. Then its memory business collapsed underneath it.

Large Japanese manufacturers had learned to make memory chips in enormous volumes and at high quality. Their prices fell, and they kept falling. One company's memory chip became much like any other's, and the cheapest maker won.

Intel's share of the memory market shrank year after year. By the middle of the decade the whole industry was in a slump, and Intel was losing money on memory it had pioneered.

The way out came from a conversation that Andy Grove later made famous. Grove asked Moore what a new chief executive would do if the board threw the two of them out. Moore answered that a new boss would get Intel out of memories. So Grove asked why they should not walk out the door, come back in, and do it themselves.

On the tenth of October, nineteen eighty-five, Intel announced that it was leaving the market for dynamic memory. A week later, it introduced the three-eight-six processor. And in nineteen eighty-seven, Grove became chief executive.

Leaving memory was painful. Intel's own history calls the year that followed the toughest in the company's existence.

Grove admitted it was an emotional decision, because Intel had built the memory business from nothing. Looking back, he said that "getting out of DRAMs when we did was the best business decision we ever made."

The bet was that the processor, not memory, was Intel's future. And the three-eight-six made that bet pay.

The three-eight-six carried two hundred and seventy-five thousand transistors, and ran at sixteen megahertz. Intel says it gave more than double the performance of the chip before it. Within a few years, sales of the three-eight-six alone made up about half of Intel's revenue.

Grove drew a lesson from those years and made it his motto. In nineteen ninety-six he published a book about it, and he called it Only the Paranoid Survive. The point was that a company at the top has to watch constantly for the change that could destroy it, because the change will not announce itself.

7. Intel Inside



By the start of the nineteen nineties, Intel sold its processors to the companies that built computers. The people who bought the computers had little idea what was inside them.

Intel decided to change that. It began talking to the buyer directly, over the heads of the computer makers.

In nineteen ninety-one, Intel launched its Intel Inside campaign, and ran its first television adverts. Computer makers who put the Intel Inside logo on their machines and in their own adverts got help paying for that advertising. Then, on the twenty-second of March, nineteen ninety-three, Intel launched the Pentium.

The Pentium had a name instead of a number, because a number could not be registered as a trademark. It carried three point one million transistors. Later in the decade came the Bunny People, dancers in the white suits worn in Intel's clean rooms, and the chime at the end of every advert. Intel had turned a component that nobody would ever see into a brand people asked for.

Then, in nineteen ninety-four, the Pentium was found to get some sums wrong.

A mathematics professor, Thomas Nicely, noticed that his Pentium gave slightly wrong answers to certain divisions. The fault was in the part of the chip that handles numbers with fractions. It was rare, and it was small, but it was a mistake in arithmetic, made by a machine sold on being good at arithmetic.

At first, Intel said the flaw was not serious and offered new chips only to people who could show they were affected. The public outcry grew. On the twentieth of December, nineteen ninety-four, Intel offered to replace the chip for anyone who asked. The cost was a charge of four hundred and seventy-five million dollars, announced in January nineteen ninety-five.

The charge did not sink Intel. Its profit for nineteen ninety-four was almost the same as the year before.

But it taught the company something the advertising had made true. Once people knew the name on the chip, the chip was Intel's problem, in public, every time. Being a brand meant being blamed.

For the rest of the nineties, Intel rode the boom in personal computers and in the internet. It was, by most measures, the most powerful company in its industry.

8. Missed turns



The years of dominance hid a series of bad calls. Each looked reasonable at the time. Together they cost Intel its lead.

The first was Itanium. It was a new kind of processor for large servers, built with Hewlett-Packard, and it was meant to replace x86 at the top end. It did not run older x86 software well, and customers did not want to rewrite their programs.

The timeline on screen marks the missed turns. Itanium launched in two thousand and one, and never sold widely outside large servers. Intel's last shipment of it went out in July twenty twenty-one, twenty years later.

Rival AMD extended x86 to handle larger amounts of memory instead, and that is the version the world adopted. Intel had to follow it.

The second bad call was the phone. When Apple was building its first iPhone, it looked at using an Intel chip. The price Apple wanted was below what Intel expected the chip to cost to make, and Intel stepped back. Paul Otellini, Intel's chief executive at the time, talked about it in twenty thirteen, as he was leaving the job.

Otellini said: "We ended up not winning it or passing on it, depending on how you want to view it. And the world would have been a lot different if we'd done it." The phones Apple went on to sell were made in volumes far beyond anyone's forecast.

Phones ended up running on chips designed by others, and those chips were made by a different kind of company.

The third bad call was the one that did the most damage. Intel had always made its own chips, in its own factories, and its edge came from making smaller transistors sooner than anyone else. Each new generation of manufacturing is called a process node. A node is the manufacturing generation a chip is made on, and its name, in nanometres, is no longer a real measurement of anything.

Intel's ten-nanometre node was years late, and did not reach large volume until twenty nineteen. In July twenty twenty, Intel said the node after it had slipped as well.

While Intel stumbled, the lead passed to Asia. TSMC, in Hsinchu, Taiwan, makes chips for other companies and nothing else. Samsung, in South Korea, makes both its own chips and other people's. Intel's own development work stayed in Oregon, and for the first time in its history, it was no longer the most advanced chip maker in the world.

9. Losing Apple, losing the lead



There was a moment when it looked as if even Apple had come round to Intel.

For years, Apple's computers had used other processors. Then its chief executive, Steve Jobs, decided to change.

On the sixth of June, two thousand and five, Apple announced that its Macs would move to Intel processors. The first Intel Macs went on sale in January two thousand and six. On the twenty-second of June, twenty twenty, Apple announced it was leaving, for chips it designed itself.

The first Macs with Apple's own chips went on sale in November twenty twenty. They were fast, they ran cool, and their batteries lasted a long time. Apple's chips were made by TSMC, on a more advanced process than Intel could offer.

Losing Apple hurt Intel's reputation. The bigger loss was happening in the data centre, which is a building full of servers that stores and processes data for other computers.

For years, Intel's server processors had most of that market. Two rivals now took pieces of it from different sides. AMD sold x86 server processors that competed directly with Intel's, and it had its chips made by TSMC. Nvidia sold something else altogether: graphics processors, which can do thousands of small calculations at once.

That turned out to be exactly what artificial intelligence needs. Training and running large AI models is mostly vast numbers of simple multiplications. A graphics processor does them far faster than an

ordinary processor can.

The comparison on screen is revenue from data centre chips in one year. Nvidia's data centre business took in forty-seven point five billion dollars in the year to January twenty twenty-four. Intel's data centre and AI business took fifteen point five billion in twenty twenty-three. AMD's took six point five billion that year.

The years do not line up exactly, because the companies' financial years end in different months. But the gap does not depend on a few weeks. Nvidia's figure was more than three times what it had been a year before.

Intel had spent its history as the company that sold the most important chip in the computer. In the age of AI, the most important chip was someone else's.

10. IDM 2.0 and the $18.8bn loss



In February twenty twenty-one, Intel brought back Pat Gelsinger as chief executive. He had started at Intel as a young engineer, led the design of the four-eight-six processor, and become its first chief technology officer. He came back with a plan to win the lead back by building.

Intel had always been what the industry calls an integrated device manufacturer. That means it designed its chips and made them in its own factories. Gelsinger kept that, and added something new.

Gelsinger called his plan IDM 2.0. Intel would keep designing and making its own chips. It would also become a foundry, which is a company that owns the factory and makes chips to other companies' designs, the way TSMC does. And where it made sense, Intel would buy some manufacturing from outside foundries too.

The plan was announced on the twenty-third of March, twenty twenty-one. It began with about twenty billion dollars for two new factories in Arizona. Later that year, Gelsinger promised five new process nodes in four years, to catch up with TSMC.

The plan needed money on a scale Intel did not have, and it needed help from the government.

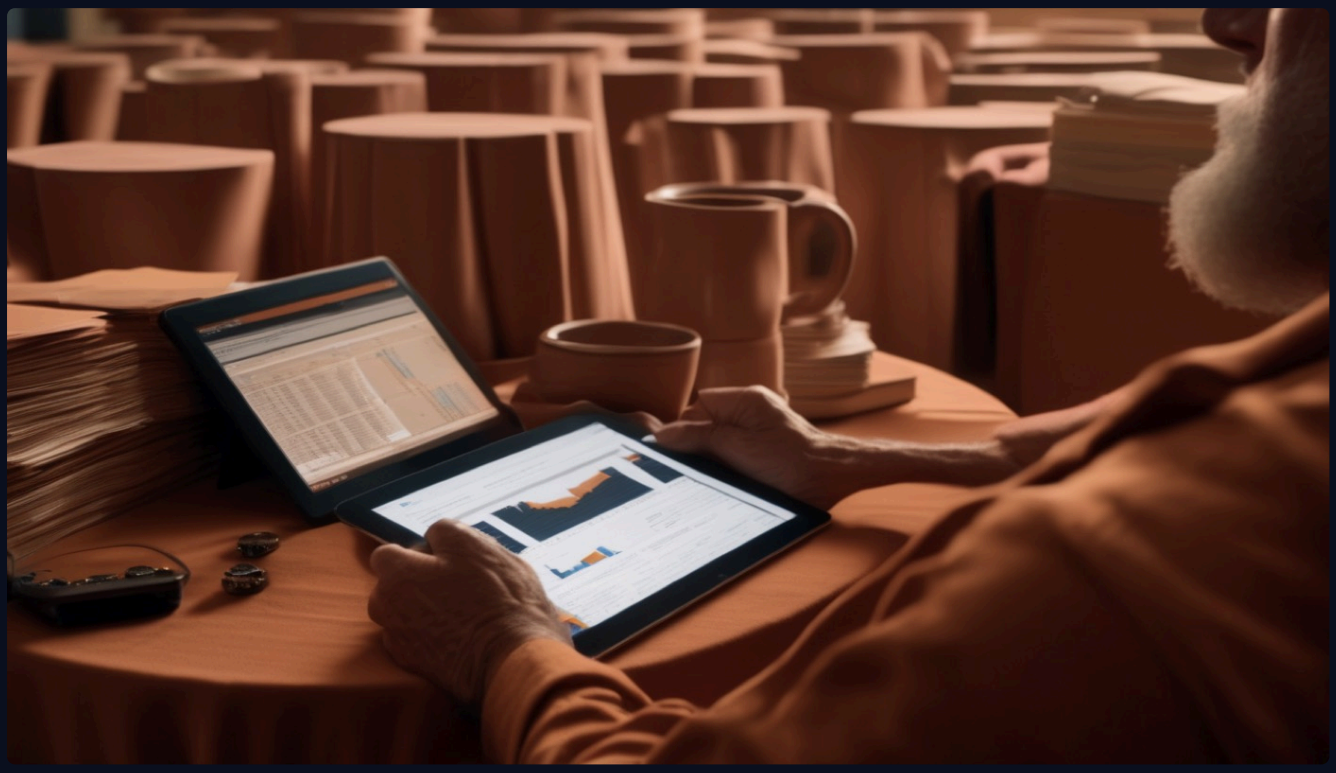
In August twenty twenty-two, the CHIPS and Science Act became law, with thirty-nine billion dollars in grants for chip factories on American soil. On the twenty-sixth of November, twenty twenty-four, Intel's award was finalised at seven point eight six billion dollars. On the first of December that year, Gelsinger retired, with the board's plan unfinished.

His exit came near the end of the worst year in Intel's history. Its factories were expensive to build, and its foundry had almost no outside customers. Its sales of personal computer and server chips were under pressure from AMD, and its data centre business was losing ground to Nvidia.

Intel reported its results for twenty twenty-four on the thirtieth of January, twenty twenty-five. Revenue was fifty-three point one billion dollars, down two per cent. The net loss was eighteen point eight billion dollars. The line on screen shows Intel's yearly profit falling from nineteen point nine billion dollars in twenty twenty-one to that loss.

Much of that loss came from one-off charges, including the cost of cutting jobs and writing down the value of assets. But it meant that the biggest bet of Intel's modern history had so far produced a very large bill and no clear return. Two executives ran the company for the next three and a half months while the board looked for someone new.

11. The state takes a stake



The board chose Lip-Bu Tan. He was a long-time investor in chip companies, and he had run Cadence, which makes the software engineers use to design chips. He had also sat on Intel's own board until the summer before.

Tan was named chief executive on the twelfth of March, twenty twenty-five, and started six days later. His first job was to cut.

In July twenty twenty-five, Intel said it would end the year with a core workforce of about seventy-five thousand people. That was a cut of about fifteen per cent. It dropped planned factories in Germany and Poland, and slowed the building of its new factories in Ohio.

In the same month, Intel told investors something striking about its future. Its next manufacturing process, after the one about to start production, was called fourteen-A.

The warning was in Intel's quarterly report of July twenty twenty-five. If Intel could not win a big outside customer for fourteen-A, it might pause or stop work on it. It also said it had not yet won any large outside customer for any of its processes.

Then, in the space of a month, three large investors bought into Intel. The first was SoftBank, the Japanese investment group, which agreed on the eighteenth of August to buy two billion dollars of new Intel shares.

The second was the American government. Four days later, on the twenty-second of August, the government agreed to take a stake of nine point nine per cent. It paid with money Intel had already been promised: most of it from its CHIPS Act grants, and the rest from a defence programme to make secure chips for the military.

The third was Nvidia, the company that had taken the data centre. On the eighteenth of September, twenty twenty-five, Nvidia agreed to invest five billion dollars. The two companies would design chips together, combining Intel's x86 processors with Nvidia's graphics.

The bars on screen are the three investments. The government put in eight point nine billion dollars, paying just over twenty dollars a share. Nvidia put in five billion dollars, and SoftBank put in two billion.

A company that had once been the model of American industry was now part-owned by the American government, and part-owned by its biggest rival.

12. Forwards or backwards?



By September twenty twenty-six, the picture has changed again, and in some ways it looks much better.

Intel's factory in Arizona, called Fab fifty-two, was declared fully running in October twenty twenty-five. It is making chips on Intel's newest process, eighteen-A, which Intel calls the most advanced chip process ever developed and made on American soil.

Eighteen-A changes two things at once. First, its transistors use a gate-all-around design, which is a shape where the control gate wraps the channel on every side, so the switch leaks less current. Second, it moves the wiring that carries power to the back of the chip. That leaves the front free for the wiring that carries signals.

Both ideas are hard to make work in volume, and Intel is making them work in its own factory. The first eighteen-A chips were processors for laptops, with server chips following.

The business has turned too. Demand for processors to run AI systems has lifted sales of ordinary server chips as well as graphics chips.

On the twenty-third of July, twenty twenty-six, Intel reported sales of sixteen point one billion dollars for the second quarter. That was twenty-five per cent more than the same quarter a year before. Intel called it its strongest growth in more than fifteen years.

But the same results showed how strange Intel's position has become.

For that quarter, Intel reported a net loss of eleven billion dollars. Most of it was an accounting charge of twelve point five billion. It came from shares Intel is holding for the government under the stake deal, which had to be valued again as Intel's share price rose.

In other words, Intel's shares rose so much that the rise showed up as a loss. And investors were willing to pay far more for those shares than a year before.

On the tenth of August, twenty twenty-six, Intel sold new shares at ninety-five dollars each. The bars on screen compare that price with what the three investors of twenty twenty-five had paid. The government paid a little over twenty dollars a share, and Nvidia and SoftBank paid about twenty-three.

The sale had been planned at fifteen billion dollars, and was raised to twenty billion. The banks then bought the extra shares they were allowed, so the total came to about twenty-three billion dollars.

So is Intel moving forwards or backwards? The answer depends on one decision that has not yet been made, and it is not Intel's to make.

Intel says its investment in fourteen-A will follow firm commitments from customers. It expects those decisions from the second half of twenty twenty-six into the first half of twenty twenty-seven. As of late September twenty twenty-six, no large outside customer has been named.

If a big chip designer commits to fourteen-A, Intel becomes what Gelsinger wanted: a serious rival to TSMC, making the world's most advanced chips on American soil. If none does, Intel may go back to making mostly its own chips, one generation behind the leaders. Fifty-eight years after the lawn-mowing conversation, the company that set the pace of the industry is waiting for someone else to decide its future.

Sources and credits

Primary sources

- Intel, 'Explore Intel's history' timeline (timeline.intel.com) and virtual vault article 'Intel's Founding' (intel.com) - founding 18 July 1968, Mountain View, name bought for $15,000, 3101 first product, 1103 DRAM 1970, DRAM exit 10 October 1985 and Grove's quotation, 386 on 17 October 1985 (275,000 transistors, 16 MHz), 8086 on 8 June 1978, Intel Inside 1991, Pentium 1993..
- Computer History Museum, This Day in History (18 July, 15 November) and The Silicon Engine (1956, 1957, 1965) - the traitorous eight, Fairchild in Palo Alto, the 4004 advertisement in Electronic News, 2,250 transistors, Moore's 1965 and 1975 forecasts.
- Gordon Moore, 'Cramming more components onto integrated circuits', Electronics, 19 April 1965.
- IBM, 'The IBM PC' (ibm.com/history) - 12 August 1981, USD 1,565, Intel 8088, off-the-shelf parts, published technical reference.
- Intel Form 10-K for 1994 (SEC) - $475 million pretax charge for the Pentium divide problem; Wikipedia, 'Pentium FDIV bug' - Nicely, 20 December 1994 replacement offer, 17 January 1995 announcement.
- Engadget, 17 May 2013, reporting Paul Otellini's interview with The Atlantic - the iPhone quotation.
- Apple newsroom, 6 June 2005 and 22 June 2020 - the move to Intel and the move away.
- NVIDIA Form 10-K for fiscal 2024; Intel 2023 annual report; AMD Q4 and full year 2023 release (30 January 2024) - data centre revenue.
- Intel, 'Intel CEO Pat Gelsinger Announces IDM 2.0 Strategy', 23 March 2021; Intel and US Commerce Department, $7.86 billion CHIPS award, 26 November 2024; Intel Q4 and full-year 2024 results, 30 January 2025 ($53.1bn revenue, $18.8bn net loss); Intel Form 8-K, 1 December 2024 (Gelsinger retires).
- Intel Form 8-K, 10 March 2025 (Lip-Bu Tan, effective 18 March); Intel Q2 2025 results and Form 10-Q, 24 July 2025 (75,000 core workforce, Germany and Poland, Ohio slowed, 14A warning).
- Intel newsroom, SoftBank $2bn agreement, 18 August 2025; Intel newsroom, 'Intel and Trump Administration Reach Historic Agreement', 22 August 2025 (433.3 million shares at $20.47, $8.9bn, 9.9%); NVIDIA newsroom, 18 September 2025 ($5bn at $23.28).
- Intel newsroom, Panther Lake and Fab 52, 9 October 2025; Intel, 'Intel 18A' (intel.com/foundry) - RibbonFET and PowerVia.
- Intel Q2 2026 results, Form 8-K, 23 July 2026 ($16.1bn, up 25%; net loss $11.0bn; $12.5bn escrowed-shares charge); Intel offering pricing, 10-11 August 2026 (210,526,315 shares at $95, upsized to $20bn; over-allotment of 31,578,947 shares exercised, about $22.9bn gross).
- Motley Fool, 23 September 2026 - no committed outside 14A customer named; decisions expected from the second half of 2026 into the first half of 2027 (Lip-Bu Tan, January 2026 call).

Not regulated financial advice.