

GTF LETTER 5

AI Won't Kill EDA, It Has to Pass Through It

Why We Are Investing in Synopsys

"Doveriyai, no proveryai" — trust, but verify. (Russian proverb)

We should probably start with answering a very common question: "What even is EDA?". In summary and without exaggeration, Electronic Design Automation ("EDA") is a category of software products upon which the entire electronics industry is built, without which the semiconductor sector as it is today would not exist.

In 1987, a small team that had spun out of General Electric the year before began selling a piece of software that did something chip designers considered their craft: it took a description of what a circuit should do and generated the logic gates by itself. Engineers did not like the idea. Aart de Geus, the founder, recalls that the team joked the tool had a "license to kill", because it was the first one allowed to change a circuit on its own, something that *"was completely taboo before"*. In his words, the problem was never the technology: *"people really could not trust that tools would change it for the better."* The company was Synopsys. The tool went on to enable the jump from thousands of hand-placed transistors to the billions we design today, and nobody misses drawing gates by hand.

Almost 40 years later, the same fear is back with the sign flipped. In June, OpenAI and Broadcom unveiled Jalapeño, OpenAI's first AI chip, taken from the first line of code to tape-out (the moment the design is sent to the factory) in nine months with the help of OpenAI's own models. In July, a Chinese lab showed a model that designed a small working chip in 48 hours, with no human in the loop and using only free, open-source tools. In August, at the Hot Chips conference, OpenAI's engineers went on stage to explain how they did it. This time the worry is not that the software will replace the designer. It is that AI will replace the software. Synopsys trades roughly a third below its 2025 high, and the WhatsApp messages we receive now come with a new expression attached: the "EDA apocalypse".

We spent the past few months on this question and decided to share part of our work, as we did with Google, Oracle and SaaS. More than having precise answers, we prefer to offer a compass. These are the main points:

- **AI is changing chip design, and fast.** OpenAI took its first chip from the first line of code to tape-out in nine months and says that "would have been impossible" without AI.
- **But AI is automating the engineer, not replacing the tools.** Every example so far, Jalapeño and Kimi K3 included, still relied on EDA engines to verify and sign off the design.
- **More AI means more EDA usage, not less.** An agent runs around 100 experiments where an engineer ran three or four, and more companies than ever are designing their own chips, a tailwind for EDA.
- **The real risk is who captures the value, not whether EDA survives.** In the bear case, EDA keeps its base, an industry that has grown around 12% a year for the past decade, and loses part of the upside. That is far from an apocalypse.
- **EDA is a duopoly, and we are buying the half that is on sale.** Synopsys is coming out of a transition year and trades at 22x our 2027 earnings estimates, against 34x for Cadence. On consensus forward earnings, Cadence's premium is three times its historical average.

Lastly, like all we try to write about, our perspective is based on a long and deep knowledge of the industry. The GTF team through Thiago and Gabriela has deep expertise in the semiconductor industry, which is augmented by Jose's experience in EDA. He was one of Synopsys' largest investors, working with Aart de Geus and his finance team through the company's move from perpetual licenses to a ratable subscription model in the 2000s. He was also an investor in Cadence when Lip-Bu Tan (now Intel's CEO) became CEO in 2009 and led a very successful turnaround.

Let us begin by first explaining what EDA actually does, in plain language, because the whole debate depends on it. Let's dig in.

We hope this helps,

Jose Medeiros, Gabriela Moraes & Thiago Kapulskis

São Pedro Capital's GTF Team

PART 1: WHAT HAPPENED AND HOW DID THE OPPORTUNITY COME?

Back in September 2025, Synopsys had its worst day since the 1992 IPO. The company had a miss in its IP business and a large guidance cut, leading the stock to tumble 36% in a single day. IP revenue came in about 25% below expectations, hit by Chinese customers freezing orders after that summer's US export restrictions and by a large foundry customer (widely understood to be Intel) changing its roadmap. With costs already committed, the miss went straight to profits, and next-quarter earnings guidance came in almost 40% below the Street.

On top of Synopsys' own problems, large sector moves brought additional impact. The first semester of 2026 had an extreme bifurcation between semiconductor and software stocks. On one hand, the main semiconductor ETF (SOXX) had an outstanding 113% performance in the period. On the other, software's IGV index tumbled 14% in the same time interval. While there are multiple reasons for this unusual phenomenon, the most important explanation was very strong results from semis companies coupled with a narrative that software companies would be disrupted by AI (the so called "SaaSocalypse").

In this unusual market, one could argue how exactly EDA stocks would fit. Their revenues are tied to semis R&D budgets, of which EDA has been taking a growing share (from around 13% in 2016 to more than 15% recently), but their business and unit economics and even valuation are anchored to software - recurring revenues, 40%+ margins and P/E multiples that lived above 30x for years. Further, they are classified as software by the index providers. In such a "black or white" tape, the market placed them more on the software side: Synopsys fell 5% in the first half and Cadence rose 20%, a fraction of the semis' gain.

By the end of June, more announcements added to the uncertainty. OpenAI and Broadcom unveiled OAI's first inference chip - Jalapeño. On August 25, at Hot Chips (the semiconductor industry's main annual conference) the team gave the details. The product is as large as a chip can physically be manufactured (roughly 800mm², on TSMC's 3nm-class process), it was designed from a blank slate, and a SemiAnalysis assessment had it beating every Nvidia, AMD and Google chip it has been able to test on several open models. First-generation chips are usually a "freebie pass". This one was apparently not.


Finally, Moonshot AI reported in July that its Kimi K3 model had autonomously designed a functional chip in 48 hours: a 4mm² die running at 100MHz, designed, verified and simulated with no human intervention, using exclusively open-source tools. Synopsys fell 8% that day, touching a new 52-week low, and Cadence fell 9%. A week later, both stocks were about 10% lower.

What made investors nervous was the role of AI. Richard Ho, OpenAI's VP of Hardware, was direct about it in a podcast a week after Hot Chips: *"We were using AI in physical design, in verification, in the actual microarchitecture design, in performance analysis, and, as I said, in the kernels and software. [...] If we had not had it, that nine months would have been impossible."* OpenAI reported that AI-guided optimization shrank the chip's matrix-multiplication engine by 10% and another block by 8% against an already optimized human baseline. After the first chips came back from the factory, its models wrote the low-level software so quickly that performance more than doubled in less than two weeks while SemiAnalysis was still benchmarking. Ho admitted the pace *"came as a surprise to us as well"*.

While all of this is impressive and we should continue to see fast-paced innovation in this sector, we believe that the market reaction is overblown because it conflates automating the people who design chips with

replacing the software engines that check them, and because it underestimates how many more chips AI will get designed. We explore this in more detail in the next sections.

Chart 1: SNPS and CDNS vs. IGV (software) and SOXX (semiconductors), rebased to 100 on Dec-31-2024, through Sep-25-2026



Source: S&P Capital IQ, São Pedro Capital

PART 2: HOW A CHIP GETS MADE, IN FIVE STEPS

Through the 1970s, chips were drawn by hand. Technicians cut each layer out of a red film called Rubylith with X-Acto knives on a light table, and a slip of the blade could ruin weeks of work. That was feasible with a few thousand transistors. Nvidia's Blackwell GPU has 208 billion. According to SemiAnalysis, manual design stopped being possible around the 65nm generation, almost twenty years ago. Everything since has been designed by software, under human direction. That software is EDA.


Chart 2: When chips were drawn by hand: Intel's Andy Grove, Robert Noyce and Gordon Moore over a Rubylith layout, 1970s



Source: Intel, via SemiAnalysis

The key thing to understand about chips, especially for those of us who grew up analyzing software, is that there is no "patch". Software ships with bugs and gets fixed next week with an update. A chip cannot be updated that fast. Once the design is sent to the factory (the industry calls this the "tape-out"), whatever is wrong stays wrong in every unit produced under such directions. At the leading edge, the set of masks used to print the chip costs around \$40 million, and a redo (a "respin") costs \$50-100 million and delays the product by 6 to 12 months. In a market where, as SemiAnalysis puts it, a 3-month delay means billions of dollars, this asymmetry explains almost everything about how the industry behaves. And the bill keeps growing: according to IBS (International Business Strategies, a market research firm), designing a chip went from under \$50 million at 28nm to more than \$700 million at 2nm.


Chart 3: Cost of designing an advanced chip, by manufacturing process (\$ million)



Source: IBS (July 2022), via Tech Fund Substack

There are five steps between the idea and the factory (derived from the chart below):

Chart 4: The chip design flow. Note the feedback loop: the process is a loop, not a line




Source: SemiAnalysis

1. Decide what to build (steps 1-2 in the chart). Architects translate a goal ("serve a large language model at low latency within 700 watts") into choices about compute, memory and data movement. This is judgment and trade-offs, tested with simulators. Software spend here is small, but every downstream dollar is decided here.

2. Write the design (step 3). Engineers describe the chip's behavior in code, in a language called RTL. Think of it as the blueprint. According to SemiAnalysis, only about 20-30% of the code in a modern chip is truly custom logic written in-house. The rest is reused or licensed, pre-tested building blocks called IP (the USB port, the memory interface, the connection to other chips), in the same way a carmaker buys brakes and tires instead of inventing them. Synopsys is the world's largest supplier of interface IP.

3. Prove it works (steps 4-6). This is the step non-specialists underestimate the most. Before spending a dollar on silicon, the team runs the design in a virtual test lab: billions of simulated scenarios, mathematical proofs that certain failures can never happen, and special-purpose hardware that "pretends" to be the chip so that software can be written before the chip exists. Verification consumes 60-70% of a project's time. A VP at Analog Devices described it in an expert call as *"the most painful step, or the one that probably takes the longest and is the riskiest"*.

Chart 5: Where the time goes in a chip project



Source: SemiAnalysis

4. Make it physical (step 7). The blueprint is converted into real logic gates ("synthesis", the tool Synopsys was founded on) and then every one of billions of components is placed on the silicon and wired together ("place and route"). Moving one block changes wire lengths, which changes speed, heat and power somewhere else. It is a gigantic puzzle in which every piece affects the others, and it takes dozens of iterations over two to three months.

5. Sign off and tape out (steps 8-9). The final inspection. Software checks that every signal arrives on time under every combination of voltage, temperature and manufacturing variation (20-30+ scenarios at 3nm, versus 5-7 at 28nm), that the power grid will not burn out over the years, and that the layout obeys the

factory's rulebook, which at 3nm has more than 25,000 rules. Factories (TSMC, Samsung, Intel) only accept designs that pass sign-off tools they have certified for each process, and the list of certified tools is short: for TSMC's most advanced processes, essentially Synopsys, Cadence and Siemens. In timing sign-off, Synopsys' PrimeTime has more than 90% share and is the de facto industry standard. Cadence's Tempus is the certified alternative, with 5-8%.

The table below is key. For each step, it separates the work done by the *engineer*, the work done by the *engine*, and what AI can and cannot take over.

Table 1: Who does what in chip design

Step	THE ENGINEER		THE EDA ENGINE		If it goes wrong
	What they do	How much can AI take over?	What it does	How much can AI take over?	
1. Decide what to build	Defines goals and trade-offs	● Medium Explores thousands of alternatives; the goal is a business call	Models performance before anything exists	● Low Performance models still have to be reliable	Wrong product. No tool fixes it
2. Write the design	Writes code, integrates licensed blocks	● High Writing and rewriting code is what LLMs do best	Checks code quality, supplies silicon-proven IP	● Medium Most exposed spot: code checks and simpler IP	Bugs that surface months later
3. Prove it works	Writes tests, hunts for the cause of failures	● High Writing tests, triaging failures, re-running the loop	Runs the design across billions of cases; proves properties mathematically	● Low The verdict must come from the engine, not a model's opinion	A \$50-100M respin
4. Make it physical	Sets constraints, runs experiments, picks the best	● High Chooses what to try next: 100 experiments instead of 3	Places and wires billions of parts under thousands of rules	● Low Algorithms tied to each factory's process	Chip too slow, too hot or too big
5. Sign off and tape out	Reviews violations, decides fixes	● Medium Sorts violations and proposes fixes	Certifies timing, power and rule compliance for the factory	● Low The factory only accepts certified tools	Factory rejects the design, or the chip fails in the field

● High ● Medium ● Low How much of each step AI can take over, for the engineer and for the engine. Our assessment.

Source: São Pedro Capital, based on SemiAnalysis, Benchmark/StoneX and company materials. The assessment of what AI can take over is our own judgment.

Two observations from this table matter for everything that follows.

First, AI's impact is largest in the columns that describe *human work (mostly red)*, and smallest in the column that describes the *engine (mostly green)*. In steps 3 and 5 especially, the engine is not a convenience. It is the only source of truth the industry has before spending tens of millions of dollars. An LLM can write a thousand designs overnight. It cannot tell you which one works. Trust, but verify.


Table 2: The EDA pockets: market size by segment, 2025

Segment	2025 revenue	Growth	Main vendors
EDA tools (software)	\$11.5bn	12-14%	Synopsys, Cadence, Siemens
Semiconductor IP	\$3.0bn	15-18%	Synopsys, Cadence
Simulation and analysis (post-Ansys)	\$3.5bn	8-10%	Synopsys-Ansys, Cadence, Siemens
Total	\$18bn	~13%	

Source: SemiAnalysis, company reports. Growth rates are SemiAnalysis estimates.

Second, look at where the money is today. The whole EDA industry, IP and simulation included, is an \$18 billion business (Table 2). Now compare that with what its customers spend on engineers. According to Benchmark/StoneX, chip companies spent about \$137 billion on R&D in 2025. Around 81% of that went to engineers, and only 16% to EDA software and IP. AI is aimed at the 81%. The 16% is what checks the AI's homework.

Chart 6: Semiconductor R&D: \$137 billion a year, 81% of it on engineers



Source: Benchmark/StoneX. R&D spend includes merchant chip companies and captive chip design teams at OEMs and hyperscalers

PART 3: THREE THINGS THE MARKET IS CONFLATING


With the five steps in mind, the "AI disrupts EDA" headline can be split into three very different claims.

1. **AI automates the engineer's work.** Writing code, writing tests, finding the cause of a failure, tuning settings.
2. **AI automates the engineer's interaction with the tools.** Instead of a person launching the simulator and reading the report, an agent does it, 100 times in parallel, overnight.
3. **AI replaces the engine itself.** A model takes the place of the simulator, the router, the timing analysis and the factory-certified checks.

Claims 1 and 2 are happening. Claim 3 is the one that would destroy EDA. Hence, the right question for each headline is simple: which of the three did it actually show?

Chart 7: Before and after: the agent replaces the human loop around the engines, not the engines

WHAT AI CHANGES: THE LOOP, NOT THE ENGINE




Source: São Pedro Capital

What Jalapeño actually showed

We went through OpenAI's release and several other sources - the Hot Chips coverage, the SemiAnalysis deep dive, IEEE Spectrum's interviews (the magazine of the IEEE, the world's largest professional association of engineers) and Richard Ho's podcast. Here is what we found.

The nine months is real, but it is one of three clocks. OpenAI's own timeline shows the architecture concept in October 2024, the first RTL in February 2025 and tape-out in November 2025. That is the nine months. SemiAnalysis counts about 16 months from the first hires to tape-out, and first silicon arrived in May 2026, 19 months after the concept. Still best-in-class (Andrew Kahng, a professor at UC San Diego, called it *"likely best in class today"*), but the fair comparison is against the 18-24 months that a design of this kind usually takes, not against zero. And tape-out is not mass production: today there are engineering samples running in OpenAI's lab, with initial deployment planned for the end of 2026. Volume production, where first-generation chips usually meet their surprises, is still ahead.

Chart 8: The Jalapeño timeline, as presented by OpenAI at Hot Chips



Source: OpenAI, via SemiAnalysis

The team mattered as much as the model. In Ho's words: *"a great deal of the team, including myself, spent over 10 years building ML accelerators at our previous company. [...] It's basically the team that built many, many generations of the TPUs."* They also started from a blank slate with no legacy to support, and Broadcom did the physical implementation, reusing its own proven building blocks. When the interviewer summarized the takeaway as "experience still matters, but AI unlocks that. It supercharges that experience", Ho answered:

"Perfect summary." And he added, as his personal view: *"there's no world where the human experts are not needed."*

Where AI did the most, there is no EDA revenue. The most spectacular gains came *after* the chip was back from the factory: writing the low-level software (kernels) that makes models run fast on the new hardware. That is the work that built Nvidia's CUDA moat, not Synopsys' business. SemiAnalysis drew the right conclusion: the threatened moat in this story is CUDA.

Inside the design flow, AI wrote and optimized. It did not verify, place, route or sign off. More than half of the chip's logic was written in XLS, an open-source, software-like language that Chris Leary started at Google before joining OpenAI. Leary explained the choice to IEEE Spectrum: *"the AI was much better at software-looking things."* AI then searched for better versions of specific blocks, and the results were measured by the conventional tools. What ran underneath, on a TSMC 3nm-class design implemented by Broadcom, was commercial EDA. XLS's own open-source path targets only old, open manufacturing processes. As Anirudh Devgan, Cadence's CEO, put it at the Goldman Sachs conference earlier this month: *"all the LLM companies, all the hyperscalers, they're all using our tools to design chips. [...] There's no chips being designed without using our tools."*

Table 3: Jalapeño, line by line

	What AI did	What the conventional stack still did
Design code	Wrote and rewrote logic in a software-like language (XLS). Shrank the matrix engine by 10% and another block by 8% vs. a human baseline	Turned that code into gates and measured every AI proposal for speed, power and area
Verification	Helped write tests and shorten the find-fix-rerun loop	Executed the simulations and produced the verdicts
Physical design	Guided optimization so the design fit and the schedule held	Placement, routing and clock distribution, run by Broadcom on a TSMC-certified flow
Sign-off	Nothing reported	Timing, power and factory-rule checks, as required by TSMC
After silicon	Wrote the kernels. Three outside models running fast in two months	Not an EDA activity

Source: OpenAI, SemiAnalysis, IEEE Spectrum, The Data Exchange podcast. OpenAI has not disclosed its tool vendors. The right-hand column reflects what is mandatory for any design on this process.

In our framework: Jalapeño is a lot of claim 1 ("AI automates the engineer's work"), the beginning of claim 2 ("AI automates the engineer's interaction with the tools"), and none of claim 3 ("AI replaces the engine itself").

What Kimi K3 actually showed

Kimi is the purest example of claim 2. On the Cadence earnings call, an analyst described it well: an agent *sought out* EDA tools and then orchestrated the flow. The engines were still there. They were simply free ones, because the target was a 45nm teaching library that Tech Fund Substack, an independent technology research newsletter, aptly called the "hello world" of open EDA research, with no factory behind it. We explain in Part 4 why free tools stop working exactly where the money is.

An important nuance

We do not want to overstate our case. There is one place where substitution did happen: the front door. OpenAI wrote most of its logic in a free language instead of a vendor's environment. Writing code is one of

the smallest money pools in EDA (most of the design budget goes to verification and implementation), but it is a real data point, and it points to the part of the debate that actually deserves attention.

PART 4: THE REAL BEAR CASE IS ABOUT WHO CAPTURES MOST OF THE VALUE

It would be easy to stop here and “declare victory”: nobody replaced the engines, case closed. We think that would be rebutting a straw man. The serious version of the bear case concedes that the engines survive and argues something more subtle. Here are its pillars, as we did in our SaaS letter.

1. The agent, not the engine, becomes the product. And the best agents belong to the AI labs, not to the EDA vendors.
2. With fewer engineers, customers buy fewer licenses, and EDA becomes a slow-growth utility called by somebody else's software.
3. Agents make free, open-source tools usable, eroding the base from below.
4. Smarter agents will need fewer runs, so the "more usage" argument reverses over time.

Let's analyze each of them.

1. Who owns the agent?

This is the strongest argument, and the evidence for it comes from Jalapeño itself. OpenAI's own model ran the loop. According to SemiAnalysis, OpenAI plans to pitch a scaled-up internal version of Codex to enterprise customers. One of the stated goals of the chip program, in Ho's words, was to use *"your own AI models to make better silicon, to make silicon come faster"* for the whole ecosystem. The most sophisticated customers in the world are, at the same time, the vendors of the best agents.

We give Cadence's CEO, Anirudh Devgan, credit for not dodging the question when Jim Schneider, Goldman Sachs' semiconductor analyst, asked it: *"Yeah, that could happen. [...] Some customers are writing some agents that call our tools and not use our agents."* His answer has two parts. First, the agent layer is *new* revenue for EDA: *"what agents are doing is they're not replacing the middle layer. They are replacing what humans used to do."* Losing part of a market you never had is not the same as losing your market. Second, the vendor that wrote the engine has access to its internals that no outside agent has, and customers have always customized on top of EDA tools.

Our own view is more cautious than management's. We do not underwrite the Goldman Sachs estimate of a \$3.7 billion agent market for EDA vendors by 2030, because Jalapeño suggests the largest customers may bring their own agents. We treat agent revenue as an option. What we do underwrite is that every agent, whoever owns it, has to go through the engines to get a verdict, and that it will go through far more often than a human ever did.

2. Fewer engineers, fewer licenses?

This pillar rests on a misunderstanding of how EDA is sold. The large customers do not buy "one seat per engineer". They buy multi-year agreements for a pool of tool capacity, which is consumed by *jobs*, not by people. One engineer, or one agent, launching 100 parallel experiments checks out 100 licenses. Devgan described the change well: a human typically runs three or four experiments at a time, *"when the agent runs it, it runs like 100 experiments"*. According to SemiAnalysis, AI features can multiply license consumption per

design project by up to five times, and Synopsys said at its 2024 Investor Day that renewals including AI tools come in around 20% above the baseline. We flag that these are vendor-reported figures.

There is also an arithmetic point. TSMC's roadmap, according to Devgan, implies chips 48 times more complex over the next five years, and some hyperscalers told him they would need to double their engineering headcount *every year* to keep up. They cannot. The productivity gain from AI is not a surplus to be returned to the CFO. It is the only way to build what is already on the roadmap.

Meanwhile, everybody is building silicon

The "fewer licenses" argument also ignores the most important change in EDA's customer base in a generation. In 2024, Synopsys' CEO Sassine Ghazi recalled that fifteen years earlier "*pretty much 100%*" of the company's revenue came from semiconductor companies. Today around 45% comes from "systems companies", businesses that do not sell chips but design them to differentiate their own products, on a revenue base that is several times larger. At Cadence the number is also 45%, up from 40% just two years ago.

It started with Apple, which today employs more than 8,000 chip designers. Then came the hyperscalers: Google's TPU, Amazon's Trainium, Microsoft's Maia, Meta's MTIA. SemiAnalysis estimates that hyperscaler custom silicon created \$15-20 billion of chip design activity that barely existed five years ago, each program requiring a full EDA stack at the most advanced (and most expensive) processes. Then the carmakers: Tesla, Rivian, BYD, Xiaomi, NIO. Devgan, back from China: "*they're all designing chips, they're all our customers.*" And now the AI labs themselves. OpenAI has Gen 2 close to locked and Gen 3 underway. Anthropic is building a chip team.

Is this accelerating? We think so, and AI is the reason. Devgan again: "*The commitment to silicon is the strongest that I have seen. [...] We used to get questions like few years ago, well, will all these hyperscalers really do chips or not?*" Jalapeño is the best evidence. A company that had no hardware team two years ago taped out a chip at the physical size limit, on a leading-edge process, and the first samples work. If AI makes that possible, the conclusion is not "fewer EDA customers". It is that the barrier to starting a chip program has just dropped for every large company with a workload worth optimizing. Each of them needs the engines, the IP and the sign-off, and none of them has decades of legacy tools of their own.

3. Free tools from below

Could agents make free, open-source EDA tools good enough? This is what Kimi K3 hints at, and the answer depends on where.

To manufacture a chip, a designer needs the factory's "recipe book" for that specific process (the industry calls it a PDK): the rules and models that tell the software how that factory's transistors actually behave. For very old processes (130nm, 90nm), a few factories have published open recipe books, and free tools can use them. Kimi did not even use one of those: its 45nm library has no factory behind it at all. For advanced processes, where most of the industry's money is made, TSMC, Samsung and Intel treat the recipe book as a trade secret. They release it only to approved customers under strict licenses, and they certify only tools from Synopsys, Cadence and Siemens to work with it, after roughly two years of joint development for each new process. A factory has no incentive to change that: if a chip fails because of an uncertified tool, it is the factory's yield and reputation that takes the hit.

Therefore, the risk is real at older processes and in China, where export restrictions already push customers towards local vendors (Synopsys' China revenue, excluding Ansys, fell 22% in fiscal 2025). At the leading edge it is very small.

4. Will smarter agents need fewer runs?

Maybe. We cannot prove they will not. What we can offer is an anecdote Devgan told at the Goldman Sachs Communacopia & Technology Conference in September. Around 2006, then at Magma Design Automation, he launched a simulator that was 10 times faster, and his marketing team was worried that customers would buy 10 times fewer licenses. *"That never happens. [...] That's the history of EDA."* In the late 1990s, designing a CPU took 500 people and five years. Today it takes 30-40 people and six months, and the EDA industry is many times larger. Every efficiency gain in this industry's history was absorbed by more ambitious chips. We expect the same, but we register that this is an argument from history, not a guarantee.

Chart 9: EDA industry revenue by company: Synopsys and Cadence alone have posted thirteen years of unbroken revenue growth through every cycle



Source: SemiAnalysis, company reports

Where this leaves us

Putting it together, our reading is that the realistic downside scenario is not "EDA goes away". It is "EDA remains essential but captures less of the upside than the bulls hope": a business growing high single digits instead of low/mid-teens, with 40%+ operating margins. That is a very different risk from the one the word "apocalypse" suggests, and it is the scenario to keep in mind when looking at the valuation in the next section.

PART 5: WHY SYNOPSYS

Everything we wrote so far applies to Cadence as much as to Synopsys. So why Synopsys? We will start with the uncomfortable part.

Cadence is the better-executing company today. Its core EDA business grew 13% last year and 18% in the first half of this year, against 8% last year and low double digits expected this year for Synopsys. Its IP business is


growing around 30% this year, its emulation hardware has a clear lead, and its margins are higher. Benchmark/StoneX summarized the consensus well: Cadence is *"the best house in a fancy neighborhood"*, and Synopsys is *"the middle house in a fancy neighborhood with only three houses"*. We do not dispute any of this.

What we dispute is the price being paid for that difference, for three reasons.

1. This is virtually a duopoly, and duopolies do not have a winner

Almost every large customer uses both vendors. Cadence is entrenched in analog design, Synopsys in digital design and in sign-off. The same people have moved back and forth between the two companies for decades. Share does shift, but it shifts at the margin, and it shifts slowly: over the past six years Cadence went from 32% to 38% of the EDA software market and Synopsys went from 35% to 38%. Both gained, at the expense of everyone else. Switching a design flow means re-running everything downstream and retraining a team in the middle of a project that cannot be late. Nobody does that lightly.

Chart 10: Virtually a duopoly, in which both players gain share



Source: ESD Alliance, Benchmark/StoneX, São Pedro Capital. EDA software only, so the split differs from Table 4, which includes IP and simulation

Table 4: Who gets the money: vendor share of EDA, IP and simulation revenue, 2025

Vendor	2025 revenue	Share	Strongest position
Synopsys	\$8.0bn	44%	Digital synthesis, sign-off, IP
Cadence	\$5.3bn	30%	Analog, emulation, verification
Siemens EDA	~\$2-2.5bn	13%	Physical verification (Calibre)
Big three	\$15.5-16bn	~87%	
All others	~\$2-2.5bn	~13%	Point tools, regional vendors

Source: SemiAnalysis, company reports. Synopsys includes part of a year of Ansys revenue.

Table 5: Each has its territory: market share by tool at advanced processes (7nm and below)

Tool category	#1 vendor	Share	#2 vendor	Share
Logic synthesis	Synopsys (Design Compiler)	84-85%	Cadence (Genus)	12-15%
Place and route	Synopsys (Fusion Compiler)	55-60%	Cadence (Innovus)	35-40%

Tool category	#1 vendor	Share	#2 vendor	Share
Timing sign-off	Synopsys (PrimeTime)	90%+	Cadence (Tempus)	5-8%
Simulation	Synopsys (VCS)	45-50%	Cadence (Xcelium)	40-45%
Emulation hardware	Cadence (Palladium)	55-60%	Synopsys (ZeBu)	35-40%
Analog / custom	Cadence (Virtuoso)	80%+	Synopsys (Custom Compiler)	10-15%
Physical verification	Siemens (Calibre)	85%+	Synopsys (IC Validator)	10-12%


Source: SemiAnalysis. These shares have been roughly stable for a decade.

2. Synopsys is coming out of a transition year

Three things went wrong at the same time in fiscal 2025. Export restrictions made Chinese customers hesitate. A large foundry customer moved the goalposts on a manufacturing process for which Synopsys had already built the IP. And the company closed the \$35 billion Ansys acquisition, its largest ever, taking on debt and integration costs just as the IP business fell 8%. The result was the guidance cut of September 2025, and the 36% one-day fall we mentioned in Part 1.

The numbers since then point in the other direction. IP returned to growth last quarter (+11%), with Synopsys winning more than 95% of the opportunities in the newest interface standard. Operating margin is guided to 41.5% this year, from 37.3% last year, against a long-term target in the mid-40s, which is where Cadence is today. Free cash flow guidance was raised to \$2.6 billion and debt is coming down. With Ansys inside, IP is now around 19% of revenue, not far from Cadence's 13-14%. That matters because IP is Synopsys' lumpiest and lowest-margin business, so a smaller weight means a more recurring, higher-margin mix, the profile investors pay up for at Cadence. In other words, the two reasons investors always gave for the Cadence premium (a "cleaner" mix and higher margins) are converging.

Chart 11: The IP business is out of the hole




Source: Synopsys, São Pedro Capital

3. The valuation no longer makes sense, in absolute or relative terms

We start by comparing Synopsys with its own history. At \$426 (September 25 close), the stock trades at 25x consensus earnings for the next twelve months (forward P/E). Since 2016 the average has been 32x, and over the last five years 36x. On September 15, at \$368, it touched 21.7x, its lowest multiple since March 2020, at

the bottom of the Covid panic. This is for a company whose earnings the Street expects to grow in the mid-to-high teens in fiscal 2027, close to management's long-term target.

Chart 12: Synopsys recently touched its lowest multiple since the Covid trough




Source: S&P Capital IQ consensus, São Pedro Capital. Daily, Sep-2016 to Sep-25-2026

Now against Cadence, on the same basis. Since 2016, Cadence has traded at an average premium of 16% to Synopsys on forward P/E, and before 2025 that premium had never exceeded 37%. On July 22 of this year, it reached 74% (about 41x for Cadence against 24x for Synopsys), the highest in our series. At September 25 closing prices (\$426 for Synopsys and \$326 for Cadence), Synopsys trades at 25x and Cadence at 38x, a premium of 51%. That is three times the historical average, for two companies exposed to exactly the same AI debate, with the same customers, in the same duopoly.

Chart 13: The Cadence premium is three times its historical average

SYNOPSYS VS. CADENCE: FORWARD P/E



Source: S&P Capital IQ consensus, São Pedro Capital. Daily, Sep-2016 to Sep-25-2026

What Synopsys owns that AI makes more valuable

In Part 2 we argued that the engines hardest to replace are the ones that produce verdicts. Synopsys owns the most important of them: PrimeTime, the industry standard for timing sign-off, and, through Ansys, the standard tool for checking whether a chip's power grid survives. Ansys extends the same logic beyond the chip. A 700-watt AI processor stacked with memory inside a liquid-cooled rack can be electrically perfect and still fail because it overheats or warps. Simulating heat, stress and electromagnetics is physics, not language. Devgan, the competitor's CEO, made the argument for us: *"you have to ground the AI with physics"*. The same applies to what the industry calls physical AI: robots, cars and drones have to be trained and tested in simulation before they touch the real world, and Ansys, with around 40% of its revenue in aerospace and automotive, is the simulator those industries already trust. It is also worth noting that Nvidia invested \$2 billion in Synopsys.

PART 6: WHAT WOULD PROVE US WRONG

We could always be wrong, and here are the mapped risks:

- **An advanced chip taped out without the incumbent engines.** Not a 45nm teaching exercise: a leading-edge design, accepted by TSMC, Samsung or Intel, where the simulation, the routing or the sign-off was done by a model or by free tools. This would be claim 3, and it would change our view. We have not seen anything close.

- **A foundry certifying a new tool vendor** at an advanced process, AI-native or not. The list has had three names for a long time.
- **Renewals coming in flat or down** as customers argue they need less capacity. This is where pillar 2 of the bear case would first become visible, and the backlog (\$10.9 billion at Synopsys) is the number to watch.
- **The labs selling the agent, and the vendors giving theirs away.** If OpenAI's or Anthropic's chip-design agents become the standard and Synopsys and Cadence end up bundling their own for free to defend share, EDA keeps its base but loses the upside. Neither Synopsys nor Cadence has disclosed agent pricing or revenue yet. Synopsys said it has more than 30 customer engagements and is still deciding between subscription and consumption. We expect to learn more at its Investor Day tomorrow.
- **Synopsys-specific execution.** The IP recovery stalling, the Ansys integration slipping, or Cadence's share gains at Intel and Samsung turning from marginal into structural. A low multiple does not protect against a business that keeps disappointing.

Closing thoughts

The engineers of 1987 were right that software would change their jobs forever, and wrong about what would happen next. Automating the gate-drawing did not shrink the industry. It made chips with billions of transistors possible, and the companies that owned the automation became some of the best businesses in technology.

We think the same distinction applies today. AI is automating the chip designer. It is not replacing the infrastructure that tells the designer, human or not, whether the chip works. When Synopsys launched its own AI design tools, engineers took about two years to trust them, even though the results were better from the start. Aart de Geus told the Acquired podcast in 2024 that Synopsys is *"essentially a company that has repeated its own history over and over again"*, and that he was *"almost tempted to say"* it learned it all from Ronald Reagan: *"Trust, but verify. [...] The value of trust is extremely high, but the necessity for verification is also"*, because a bug that reaches manufacturing is enormously costly. The more chips are written by machines, the more valuable the verifier becomes.

Synopsys is a position in our flagship Global Technology Fund, respecting the risk parameters in our investment policy. Let us know if you would like to hear more.

Sources

- Acquired (ACQ2), "The Software Behind Silicon", interview with Aart de Geus and Sassine Ghazi, May 2024
- OpenAI and Broadcom, "OpenAI and Broadcom unveil LLM-optimized inference chip", June 24, 2026
- SemiAnalysis: "OpenAI Jalapeño: Better Than Nvidia Blackwell" (Aug 25, 2026); "The EDA Primer: From RTL to Silicon" (May 11, 2026); "EDA Market Primer" (May 21, 2026)
- IEEE Spectrum, "Jalapeño Shows Power of LLMs for Chip Design", September 2026
- The Data Exchange podcast, Ben Lorica with Richard Ho (OpenAI), September 2026
- Cadence at the Goldman Sachs Communacopia & Technology Conference, September 9, 2026 (FactSet transcript)
- Synopsys, third quarter fiscal 2026 results and earnings call, August 26, 2026
- Cadence, first and second quarter 2026 results and earnings calls, April 27 and July 27, 2026
- Benchmark/StoneX, "EDA/IP/Simulation Primer: From Chips to Systems", July 15, 2026
- Tech Fund, "The EDA Crash & The Massive AI Opportunity", July 27, 2026
- MBI Deep Dives, "Synopsys: Harnessing Complexity", April 25, 2025
- S&P Capital IQ for share prices and consensus forward P/E multiples, through Sep 25, 2026; Citi Research, "Initiate Buys on Cadence and Synopsys as EDA Enters the Next Wave of Growth" (Nov 24, 2025), for EDA sales as a share of semiconductor R&D; OpenAI and Broadcom announcement (Jun 2026) for Jalapeño's deployment timeline

Material de caráter exclusivamente informativo e opinativo, elaborado pela equipe de gestão da São Pedro Capital. Não constitui, em nenhuma hipótese, oferta pública ou recomendação de compra ou venda de valores mobiliários, nem convite ou proposta para investimento em quaisquer fundos ou produtos geridos pela São Pedro Capital. As informações aqui contidas podem ser alteradas sem aviso prévio e refletem exclusivamente a visão da equipe na data de sua divulgação. O SPC Global Technology IE FIF em Ações ("GTF") detém posição em SNPS e pode alterá-la sem aviso. Rentabilidade passada não garante resultados futuros. Investimentos em ações e no exterior envolvem riscos, inclusive de perda do capital investido e variações cambiais.