



2026 Notch Tech Trends Report



Marko Štefančić
CTO



Darko Špoljarić
VP OF AI



Goran Pugar
ENGINEERING MANAGER

Table of Contents

2026 Notch Tech Trends Report	1	AI-Assisted Coding and the Technical Debt Expansion	21
Introduction	3	How to Code with Confidence When Using AI Assistants	22
Executive Summary	3	Agents to the Rescue	23
Word from the Founders	4	The Senior Developer Shortage	24
Part I	5	Delivery Operating Models: Speed, Quality, and Sustainable Velocity	25
Making the Business Case for AI and Custom Software Development	6	Privacy	27
Software Delivery Models and the Tech Debt Imperative	7	Governance	27
EU Market Insights	9	Horizon Technologies: GPUs, Data Centers, Quantum, and the Longer View	29
The EU AI Act	9	Part III - From Trends to Action	31
US Market Insights	11	A Year Of Development Ahead: What This Means For You	32
Enterprise Insights	12	How Notch Helps: Custom Software and Enterprise AI—Without the Debt	33
SME Insights	12	Part IV - Verified Market Intelligence & Data Addendum	34
Startups and Scaleups	13	The Numbers Behind the Trends	35
Part II - Technology Landscape and Trends in 2026	17	Global & European IT Spending: The Acceleration	35
AI Coding Tool Use 10	18	The Talent Cliff: A Statistical Certainty	35
Coding Language Trends Review	19	The Nearshore Value Equation	35
Which Technologies to Choose?	20	Cloud & Security Velocity	35
		Talk to Us	35
		Works cited	36
		Sources	37

Introduction

Technology waits for no one. And 2025 proved this once again. AI stopped being a science experiment and became something enterprises have to figure out how to live with.


Here's what nobody wants to admit: the faster you code with AI, the more expensive your codebase becomes. Not because AI is bad. Or expensive. It's because speed without governance is just debt you're taking on at interest.

This report focuses on that trade-off, amongst other things. We looked at the data. We talked to CTOs, CFOs, and product leaders across the US and Europe. We asked one question over and over: what's the real cost of AI-assisted development?


The answer for 2026 is uncomfortable. But it's also actionable.

Executive Summary

Three structural shifts are reshaping how enterprises buy, build, and run software in 2026.



First: Outsourcing stopped being just about cost. **Sixty-seven percent of tech leaders want outcome-based partnerships.**¹ They're tying cost to what ships, not to how many people were on board. And that's likely how things will look from here on out.



Second: AI code is now accepted in practice. **Eighty-four percent of developers use or plan to use AI tools.** But 66% are frustrated with code that's "almost right but not quite." Speed and technical debt are now two sides of the same coin.



Third: Technical debt is a boardroom issue. **In 2026, every enterprise will have a CFO-level conversation about the maintenance cost of AI-augmented codebases.** That's not a developer problem anymore. That's a business problem.

The market is telling us this clearly. Enterprises aren't consolidating vendors around speed. They're consolidating around governance and outcomes. If you're selling pure velocity, you're already behind.

Word from the Founders

Based on current platforms, AI should reach peak acceptance in 8 years, which is 10 years from the big change event that was the release of ChatGPT. A similar timeline was observed in mobile technologies, peaking in the mid 2010s, 10 years after the 2007 iPhone release. Or internet tech peaking in 2000, 10 years after the first web page was released in 1991.

This time, we have bigger infrastructural obstacles: this tech needs big data centres to be built, GPUs are on backorder for 12 months, and electricity is in short supply.

This could slow down the timeline.

The economy is not a zero-sum game, and job markets will restructure, but remain lively. Currently, AI systems are providing

accuracy of results in the 90%+ range. Getting to 100% might drive the costs of electricity and computing even higher. Considering that machine learning is already using 100x more computing power than standard apps, and 10–100x more electricity for 90-ish % precision, it might become “too expensive” to replace humans where 100% precision is required. Quantum computing is still too vague to bet on for help.


Short term trends: Small language models, generative AI, extended reality, and domain-specific.

Future trends: Personalised healthcare and longevity, superhuman intelligence across domains!



Roko Roić,
Co-Founder, Notch





Market & Business Landscape

Making the Business Case for AI and Custom Software Development

A lot of companies are still working hard to turn AI into something profitable. They have the tools. They have the momentum. But they don't have the use cases.

According to Deloitte's Global Business Services survey, 50% of organizations are prioritizing next-generation GenAI capabilities. This crosses the threshold from experimentation into genuine demand.

But here's the reality: While 58% of those organizations have started or are planning a GenAI journey, most haven't achieved even 10% of their expected savings.


Why? Data and governance aren't there. It's not a technology problem. It's an organizational one. The enterprises winning

with AI in 2026 are going to be the ones who are treating it like infrastructure, not like a feature you bolt onto your product.

The European IT services market reflects this shift. Gartner projects spending will go from USD \$490 billion in 2025, accelerating to nearly USD \$540 billion by 2026. That 10.1% annual growth rate sounds good – until you realize what's driving it. It's not people getting cheaper. It's enterprises getting smarter about what they need to build and who can deliver it.

sweet spot. It also signals something else: cost advantage alone no longer drives vendor selection. Regulatory alignment does.

Nearshore outsourcing is growing because Central and Eastern Europe offers both labor-cost advantages and EU regulatory compliance. For enterprises trying to balance cost and governance, that's the sweet spot. It also signals something else: cost advantage alone no longer drives vendor selection. Regulatory alignment does.



Source: ⁹ Gartner forecasts EU IT spending to grow

Software Delivery Models and the Tech Debt Imperative

The outsourcing model that worked for 30 years is breaking.

What was once about staffing bodies is now about delivering outcomes. Fifty years of outsourcing history said: more people, lower cost, more output. That model still exists for commodity work. For everything else, though? It's dead.

And here's why. Enterprises started measuring what matters: how fast they ship, how stable their code is, and what percentage of the budget goes to keeping old code alive versus developing new applications. Those metrics exposed something uncomfortable. Fast and broken is worse than slow and solid.



67%

Enterprises prefer partnerships focused on delivering specific results, not services.

Source: 2025 Gartner Magic Quadrant for Custom Software Development

Gartner's 2025 Magic Quadrant for Custom Software Development shows this. IBM and HCLTech landed as leaders because they can demonstrate sustained execution and the ability to manage technical risk. The evaluation centered on the ability to execute with AI and deliver outcomes. **Speed matters. But controlled, sustainable speed matters much more.** That's a different contract. That's different pricing. That's different delivery mechanics.

If all you're selling in 2026 is just staff augmentation, you may be behind the curve. **Because there are two main ways you can augment your staff—the first is by hiring a team of people who can write code for any business, and the second is by hiring a team of people who can write code specifically for your business and business goals.**

The difference between these is one team does what you tell them to do and ship the code you tell them you need. The other takes time to understand your industry, your business case, and the problems you're solving in order to ship a complete solution that works now and in the future.



There's a difference. The first model is simple staff augmentation. And it works if you know exactly what needs to be done, but don't have enough hands to get it done yourself.

The second model is augmented staff augmentation. This is an in-house team that works with you to understand your business and achieve your goals, maximizing your ROI.

And these so-called Build-Operate-Transfer models are gaining favor. **Clients want 18-month payback cycles. They want credible partners with staying power. They want to know you're in this for the long game, not the one-and-done kind of deal.** Need help developing AI teams? Get our [AI Playbook](#) 

EU Market Insights

The European IT services market is growing. But understanding why it's growing matters more than the numbers

The market stands at USD \$490 billion in 2025, headed toward USD \$540 billion by 2026.³ The growth is real. And what's driving it tells you what enterprises actually care about.

Managed Security Services is the fastest-growing segment. The reason: the NIS2 Directive. Starting in 2024, the directive requires entities across the EU to meet core cybersecurity controls. And it's mandatory. In response, 80% of European organizations plan to increase reliance on external security specialists.

The EU AI Act

The EU AI Act has moved from a topic of debate to a set of hard deadlines. While early prohibitions on "unacceptable risk" systems kicked in back in February 2025, the most significant milestone for businesses awaits us on August 2, 2026. By this date, any organization deploying "high-risk" AI systems must be fully compliant with stringent verification, documentation, and oversight rules. The grace period is effectively over; think of this not as a distant target, but as the immediate operating environment for 2026.

To navigate this landscape, organizations must categorize their AI assets into the Act's four risk tiers. At the top are Prohibited practices, which are banned outright. Below

that is the critical High-Risk category—systems affecting safety or fundamental rights, such as recruitment tools, credit scoring, or biometric identification. These require the heaviest lifting: conformity assessments, high-quality data governance, and continuous human oversight. The vast majority of current business AI, however, falls into Limited Risk (chatbots, emotion recognition) or Minimal Risk (spam filters, inventory predictors).

For "Limited Risk" systems, the primary obligation is transparency. Users must be explicitly informed when they are interacting with a machine. Further, generative AI content must be clearly labeled as such, creating a "trust layer" in the digital economy.

For "High-Risk" systems, compliance is far more granular. It requires building a "Central Documentation Hub" that details exactly how the system operates, the data it was trained on, and the logic behind its decisions. Product managers, legal teams, and data scientists must collaborate to ensure that every algorithm has a "paper trail" that can withstand an audit. The penalties for getting this wrong range from €35 million or 7% of global annual turnover—whichever is higher.

Treating the EU AI Act as a simple checkbox exercise is a strategic error. Compliance is a continuous process, and done right sends a signal of maturity and reliability.

The **German** market leads the European IT services sector. German providers face premium rates due to talent shortages. The solution: satellite hubs in Central and Eastern Europe, including Croatia, Poland, and the Czech Republic. That's the real nearshore story. You're not moving jobs to cheaper geographies purely for cost. You're moving them because you can't find talent where you need it.

Cloud adoption numbers sound optimistic until you read the execution data. Ninety-nine percent of European enterprises prioritize cloud investment. But only 32% have actually achieved their stated cloud transformation goals. That's a 67% execution gap. That's work for partners who know how to actually build and migrate, not just talk about it.

UK market dynamics show the pressure points. Brexit residuals still prolong decision



cycles. Public-sector contracts now emphasize outcome-based clauses and price pressure. That's a lower margin than it was five years ago. But it's also the future. If you want to sell to the UK government and large enterprises, outcome-based pricing isn't optional.

The **Nordic region** is the fastest-growing cluster. Denmark ranks second globally in digital government. Sweden exports significant software annually.

The talent is there. The expertise is there. And so are the **salaries**: entry-level developers in Sweden now earn EUR €4,000 per month, up 12% since 2022. The cost advantage of Europe is real, but it isn't uniform.

Nearshore is gaining traction. **Central and Eastern Europe** offers a labor-cost advantage while maintaining EU regulatory compliance. But wage inflation there runs annually. That advantage won't last forever.


But that's okay because cost isn't the only advantage afforded by CEE nearshoring. **Croatia**, for example, ranks fourth globally in English literacy, making communication seamless. Then there are business culture similarities, shared or adjacent time zones, EU regulatory alignment, and, most importantly, skill. Addressing shortages by searching for nearshore skill guarantees you get a plug-and-play team.

US Market Insights

The US tech market in 2026 sits at an inflection point. The post-election mood has shifted vendor priorities and buyer behavior.

The US software development outsourcing market stands at **USD \$500 billion** in 2024, growing at **8-10% CAGR** through 2031.⁴ But that growth is concentrated. The real story is in the talent gap.

Analysts interpreting earlier projections from the **US Bureau of Labor and Statistics** have estimated a potential gap of up to **1.2 million** software developers by 2026, but official Bureau publications now focus on a 26% growth in software developer employment (about 410,000 additional roles) over 2022–2032 and roughly 129,000 openings per year. This is not just a statistic. It’s a constraint on growth. It’s the reason



enterprises are looking at outsourcing, nearshoring, and every model in between. They can’t fill the gap with local hires.

The new US administration ran on being tech-friendly. However, policies appear to be more focused on deregulation than on addressing the mounting talent shortage. An executive order signed in December 2025 banned states from regulating AI, for example, while a \$100,000 fee for **H1B visas** hamstrings companies desperate to attract tech talent.

And global talent centers are taking advantage of this. Canada and EU countries are attracting wantaway engineers to the benefit of their own tech booms. One example is the Austrian Academy of Sciences, which brought in top talent from leading US learning centers by offering two-year grants of €500,000. The chairman of the Academy attributed the migration to the US administration’s policies.

The implication for outsourcing partnerships: US enterprises are facing a talent shortage and simultaneous brain drain, negatively impacting their ability to compete. Meanwhile, the EU is taking full advantage, reducing its immediate shortage while also opening up new areas of tech excellence which will have to be solved with nearshore outsourcing in the interim.


Enterprise Insights

Large enterprises are consolidating their vendor base. That's the real trend.

Gartner reports that large enterprises held **64.8%** of the IT services market share in 2024. That share is staying put because big enterprises want integrated delivery: integration services, managed security, custom software development, and AI acceleration. All from one partner. That's hard to do if you're a specialist. They're also making outcomes mandatory.

Sixty-seven percent prefer partnerships focused on delivering specific results, not just services. That's a different sales motion. That's different pricing. That's a different delivery structure.

The consolidation trend also means the bar for entry is rising. If you can't demonstrate execution across multiple disciplines, enterprises won't bother. They're reducing vendor counts, not increasing them. That favors partners with depth.



SME Insights

Small and medium enterprises are getting access to capabilities they couldn't afford five years ago. Outsourcing and cloud solutions are how they get there.

The Digital Europe Programme subsidizes up to 50% of qualified SME transformation costs.⁷ By 2030, SME penetration of the IT services market is forecast to add 4.5 percentage points to the overall revenue mix. SMEs are growing at a **9.1% CAGR**, faster than the market as a whole.

SMEs face the same debt problem as enterprises. When you're on the runway, code that breaks in 6 months isn't an option. You just can't afford to learn that lesson the hard way.

Startups and Scaleups

The AI gold rush mentality of the early 2020s is effectively over. While we are seeing a funding stabilization in 2026, with global capital deployment ticking up into the high **\$400 billion** range, the rules of engagement have changed.⁸ The market hasn't stopped spending, but it has stopped gambling.

We are seeing a distinct split in where the money goes. Vast amounts of capital will continue flowing into infrastructure and the companies building the actual foundation models. Everyone else is facing a new normal of longer diligence periods and valuations based on revenue, not dreams.

The IPO window remains tight, meaning secondary markets (selling

shares to new investors rather than going public) have become the primary way employees and early investors get liquidity.

The biggest risk for 2026 isn't a lack of capital. It's the mass extinction of AI wrappers. Investors have realized that startups offering a thin interface on top of existing LLMs have no defense against big tech features. If an incumbent like Microsoft or Google can release your entire product as a free feature update, you don't really have a business. Not anymore.

Here is how this plays out across our key markets and how smart founders are navigating it.



Regional Market Dynamics

DACH (Germany, Austria, Switzerland)



The region has moved away from trying to copy Silicon Valley consumer apps and returned to its industrial roots. Munich is challenging Berlin for the top tech spot, largely because its focus on hardware and engineering fits the new demand for physical AI—intelligence applied to robotics and manufacturing.

- **The Funding Vibe:** Capital is available, but it is “patient capital.” Investors are looking for deep tech and sovereign solutions that comply with EU regulations, rather than quick-flip apps.
- **The Opportunity:** Defense tech and industrial optimization are having a moment. Startups here win by selling trust and data sovereignty, guaranteeing enterprise clients that their data won’t leave Europe.

United Kingdom & Ireland



The UK remains the undisputed funding leader in Europe, pulling in more capital than France and Germany combined.¹² The focus has shifted from general fintech to agentic fintech—software that autonomously handles complex compliance and wealth management tasks.

- **The Funding Vibe:** Optimism is back. CFOs and VCs are betting on productivity, but the bar for Series A has risen. You need proven retention, not just a cool demo.
- **The Opportunity:** Ireland remains the critical talent bridge between the US and EU, though housing constraints make scaling teams harder.

Regional Market Dynamics

Benelux (Belgium, Netherlands, Luxembourg)



This region continues to be the pragmatic core of European tech. Thanks to ASML, there is a strong halo effect around hardware and semiconductors.

- **The Funding Vibe:** Stable and specialized. There is less hype here, which means fewer crashes.
- **The Opportunity:** Because Brussels is the regulatory capital, we are seeing a rise in startups turning compliance into a product. Logistics remains huge, with Rotterdam and Antwerp driving demand for AI in supply chain resilience.

United States



The US is driving the shift from generative to agentic AI. Here, the distinction between “software” and “service” is vanishing faster than anywhere else.

- **The Funding Vibe:** Aggressive but selective. Big tech is back to “acqui-hiring”—buying startups just for their talent and shutting down the product.
- **The Opportunity:** The main constraint is energy. Data centers are running out of power, driving investment into efficiency. For software startups, the goal is to prove you can do work autonomously, not just help a human type fast

The Playbook

How to Avoid the Wrapper Trap in 2026

For startups and scaleups looking to survive this correction and avoid burning through their runway, the strategy is about moving from cool tech to “boring” utility.

Based on the operational success of companies we worked with in 2025, here is a survival guide for 2026:



Go Deep Vertical

Do not build “AI for email.”
Build “AI for Nurse Staffing”
or “AI for Legal Discovery.”
General tools get crushed
by incumbents; specific
tools build moats that
generalists can’t cross.



Service-as-Software

Stop charging for “seats”
and start charging for
outcomes. If your AI agent
schedules a candidate
interview, charge for the
interview. This aligns your
revenue with the value you
actually deliver.



Own the Workflow, Not Just the Chat

The most successful
startups aren’t living inside
a chat window. They are
replacing the messy email
chains and PDF
attachments that clog up
B2B industries. Be the
platform where the work
happens, not just a plugin.



Proprietary Data is the Only Moat


If your model learns nothing
new from your users, you
are vulnerable. The best
companies use their
specific workflow to gather
data that OpenAI doesn’t
have (like which resumes
actually get accepted by
hiring managers), and use
that to fine-tune their
product.



Technology Landscape and Trends in 2026

Coding Language Trends Review

Python adoption grew in 2025



Source: GitHub 2025 Octoverse

In GitHub’s 2025 Octoverse signals, **TypeScript** emerges at first place by far, while Python remains one of the strongest forces behind the continued expansion of AI and data-related development. It reflects a split in what the industry is investing in: Python continues to dominate work close to data and models, while TypeScript increasingly anchors product delivery across the application surface.

Python has effectively become the *lingua franca* for AI and data workloads: experimentation, pipelines, notebooks, retrieval and ranking logic, evaluation harnesses, offline processing, and orchestration around model-centric systems. The reasons for that are obvious: a low barrier to entry, a mature ecosystem,

and fast iteration while requirements and approaches are still evolving. At the same time, there are trade-offs. It’s not typed by default; it’s also risky for latency-sensitive, CPU-bound workloads where throughput and tail latency matter.

The Dropbox experience is often mentioned. They invested heavily in supporting Python in an enterprise environment. And while Python can support serious production systems at scale, sustaining long-term growth requires investment in maintainability mechanisms. While Python is absolutely viable in enterprise environments, it comes with an operational and engineering cost profile that needs to be acknowledged from day one.

TypeScript’s rise is equally explainable, and it should not be interpreted as a frontend-only phenomenon. Though frontend adoption is a major driver, modern product delivery increasingly uses TypeScript end-to-end—including client applications, backend-for-frontend layers, server routes, tooling, and integration code. In a world where AI-assisted coding increases the rate at which code is produced, types become

one of the most practical and scalable guardrails available. They reduce ambiguity at interfaces, make refactoring safer, and help teams contain the “almost right” problem by moving error detection earlier in the lifecycle. Combined with a mature framework ecosystem, TypeScript becomes a natural default for building and evolving product surfaces with speed and consistency.

And finally, Java, C#, and Go on the backend, with React and Angular on the frontend—these trends reinforce a pragmatic conclusion rather than a forced language shift. Core backend services that require predictable performance, reliability, and long-term maintainability remain well served by Java, C#, and Go. TypeScript continues to strengthen its role in product-layer delivery and application composition, while Python becomes increasingly central around AI and data workflows.

Coding Languages in 2025

Rank 2025	Language	YOY	Growth (Aug 2024 vs Aug 2025)
1	TypeScript	~1,054	67%
2	Python	~850	49%
3	JavaScript	~427	25%
4	Java	~174	21%
5	C#	~136	22%

Source: GitHub 2025 Octoverse

Core stacks for new projects built in the last 12 months

Language	Growth (Jan–Aug 2025 vs. Jan–Aug 2024)
Python	53%
JavaScript	15%
TypeScript	78%
Java	9%
C++	12%
C#	11%

Source: GitHub 2025 Octoverse

Which Technologies to Choose?

In 2026 and beyond, we won’t look at what is or was the language of choice. Instead, we’ll be thinking about platformization, those golden paths, service patterns or templates, and hard architectural decisions.

With AI in the picture, it really doesn’t matter which technology we decide to use. We just have to be sure that our engineers incorporate change safety in whatever they are doing. Or, in other words, AI has become the software developer, while the human in the loop is the engineer. The role of the human is in thinking about those aspects of the solution that are more complex, not pure routine.

Software engineers of the past and future will always ship with tests and always have the infrastructure in place to run those tests. Engineers will also always reduce architectural entropy: there are well-defined patterns to boost resilience or authentication or different utilities common to different modules etc. They avoid code duplication to write clean and maintainable code. In the end they will not talk about how many tickets were cleared, but about how things are performing better in production. And that’s the real measure of success, not the pumped up velocity riddled with bugs or on-call load increases.

AI has become what we called software developer, and the human in the middle is the engineer.

Goran Pugar, Fellow & Engineer Manager, Notch



AI-Assisted Coding and the Technical Debt Expansion

Coding copilots work. But they're also reshaping technical debt in ways most teams don't yet know how to manage. **The three vectors of AI-driven technical debt are real:**

1 Model versioning chaos. GitHub Copilot updates. Claude gets new capabilities. Every team handles model transitions differently. You end up with multiple code generation standards all in the same codebase.

2 Code generation bloat. AI doesn't refactor. It doesn't consolidate. It generates code that works this week. You get duplication at scale, over-specification, and architectural inconsistency. Large blocks of duplicated code increase by as much as 8 times when AI is in the loop.

3 Organizational fragmentation. Some teams use Copilot. Some use Claude. Some write by hand. Your senior engineers spend more time on code review and rework. Your less experienced developers contribute more code, and more of it is lower quality.

Velocity without governance is velocity toward a wall. You can't inject AI into your development pipeline and hope it works. You need code quality gates. You need architectural standards. You need senior engineers verifying what the AI generates, not just running it.

That's the real outsourcing challenge in 2026. Partners who promise pure speed without addressing technical debt are selling you a liability. But partners who understand governance are the ones who win.

How to Code with Confidence When Using AI Assistants

While coding assistants help us get faster, we're not always sure what we're changing. And yes, we can end up with a bunch of spaghetti code that "works" until it doesn't. It's obvious that if you haven't been thinking about change safety until now, this is the final boarding call.

The real shift is that code has become cheap, but confidence hasn't. AI can churn out code all day long, but it can't take responsibility for the system. So the organizations that win won't be the ones generating the most code. Instead, they'll be the ones that can verify, standardize, and roll back safely, while keeping their architecture coherent.

For most, the sentiment has been that vibe coding is great for exploring and prototyping. However, if the path to production is still paved with discipline, i.e. constraints, contracts, tests, and operational guardrails, then we might benefit from vibe coding after all.

"In my years of software engineering, I always felt better when I knew we had the delivery pipeline in place: clear contracts, tests that actually protect behavior, and a smooth pipeline execution that gives you confidence, not just green checkmarks. It's the same story now, except with AI in the picture, it stops being just a best practice and becomes non-negotiable."

Goran Pugar, Fellow & Engineer Manager, Notch



Agents to the Rescue

Although many companies will struggle to cope with the challenges above, there will be a crop of those who will overcome it. **Here's what they'll look like:**

1 Their engineers use agent teams for all coding tasks. This includes refactoring, writing down LLM-generated design decisions (ADRs) for later use, handling all kinds of technical debt, creating tests, etc.

2 Their engineers are simply very good engineers, focused on good code. The only difference to a traditional good engineer is that the AI-assisted engineers have invested in becoming agent managers who have taken the time to build quality coding agent teams.

3 The management of these companies realised that in order to do that, you have to help and support traditional engineers to transition into agent managers, and give them proper tools and token budgets to do it. They also know that the crux of this challenge, this role transition, is to help them understand that the code they write with their hands is not a piece of art to get attached to, and definitely not a measure of how big their brain is.

4 They know that constraining models with good instructions and context will make them behave much more similarly to one another. This will not make the race go away, but it will allow engineers to switch LLMs without much consequence.

VorteeGO Agentic AI Framework by Notch

In 2025, at Notch, we built our own **Agentic AI framework, VorteeGo**, in-house. We see 2026 as the year we start opening it up to the world.

While developing **VorteeGo** we took a completely different path than any other agentic frameworks or services. We were thinking **enterprise-first, production-first**. We're not optimizing for quick demos. We're optimizing for running real workloads that need to be safe, explainable, and maintainable over time.

VorteeGo is built to run on Kubernetes and is written in Go to foster concurrency, predictability, and operational simplicity, which matter a lot when you're orchestrating multiple agents

and long-running workflows. On the agentic side, we support capabilities such as dynamic workflow generation and execution, multi-agent coordination, ephemeral agent spawning with lifecycle management, tool and MCP integrations with least-privilege controls, and memory patterns that you can reason about and evolve.

But the real story is what surrounds those features: observability and traceability as first-class. You can stream what the system is doing, understand which step produced which output, reproduce behavior, and change workflows without guessing. That's the control plane you need if you want agentic AI to move from prototype to something you can run in production.

And because we've designed it as a platform component rather than a locked service, VorteeGo is built to be portable: it can run anywhere your infrastructure runs, without forcing you into a single vendor's stack.

With **VorteeGo**-style agentic systems, zero-knowledge proofs can be used to prove that an agent followed approved rules, permissions, and consent constraints without revealing sensitive inputs or internal context.

This enables compliance and auditability by design, allowing enterprises to demonstrate correct behavior without exposing proprietary data or personal information.

Coding agents revolution, IDE evolution

Coding agents, as capable as they currently may be, are still quite crude tools. They don't maintain internal representations of code the way an IDE does. They don't index, their refactors aren't safe, and they still make coding mistakes that another agent or a build tool run is supposed to detect and fix. Also, they forget to edit files because they cannot maintain formal connections between code across files and repositories.

In 2026, we can expect the coding agent industry to make significant efforts to address these challenges. It is not going to be easy because the game is very different from the IDE situation. The primary use of the IDE is for a human being who reads and types slower than average (i.e. someone who needs summarised and synthesised information), but the change it produces is small and comes slowly.

Coding agents read lightning fast and can produce change at a rate of dozens of lines of code per second. It's a fundamentally different kind of user that the IDE is supposed to cater to.

So, will the IDE be the one to cater to the coding agents, or will there be something entirely new?



Our prediction is that IDEs will stay despite humans having less and less need to look at the code, but their feature set will expand to support agents because it's faster to adapt and reuse already existing facilities. The coding agents are still quite far from their plateau, expect them to become even faster and more correct in 2026.

Darko Špoljarić, VP of AI, Notch

The Senior Developer Shortage

The solution to senior developer shortage?

Junior devs – only very different this time, in Senior form. Here's how to build them:

1. Mentoring

Junior devs still need mentors. Human ones for various sanity checks, but much of the mentoring can now be done by AI. With that in mind, juniors can be expected to grow at a 5x rate compared to what they were a couple of years ago. The only requirement is they have an appetite for learning about good engineering and using AI.

2. Work experience

Doing more with coding agents means completing more work in less time. Real-world, to-the-point contextual learning opportunities are in abundance now. This is a major boost compared to what was available in the pre-AI era, which consisted largely of unconnected islands of information online. Before AI, one would rely on either Stack Overflow or blogs focused on someone else's problem (which yours might be somewhat similar to). Nowadays, AI is building connections and synthesis. Learning from that is a boost that only the best human seniors who are good mentors and are 24/7 available to you can match.

3. Big token budget

Turning a junior into a senior now requires constant daily interaction with AI.



We're predicting the rise of the token budget to over 10\$/day for such developers, but doing the math yields the only logical conclusion: it's not an investment, it's a steal.

Darko Špoljarić, VP of AI, Notch

Delivery Operating Models

Speed, Quality, and Sustainable Velocity

Traditional process frameworks stopped mattering the moment AI entered the picture.

The old debate: Agile versus Waterfall. Both assume humans are the bottleneck. When AI is generating code, the bottleneck shifts. Now it's the architecture discipline, code review quality, and senior engineer bandwidth.

High-performing teams in 2026 don't use Agile or Waterfall. They use value stream thinking. They measure what matters: cycle time, lead time, defect rates, and deployment frequency. They use whatever process gets them there.

They also know they can't optimize for velocity alone. You can hit peak velocity for

six months. Then you're spending the next year fixing debt. Sustainable velocity means you code at **70%** of peak capacity while leaving **30%** for refactoring, governance, and architecture.

AI Privacy & Governance

Privacy

It is already emphasized that in 2026, trust is no longer assumed; it must be provable.

Privacy will become and already is a much bigger issue for the public, driven largely by the explosion of AI video and image generation. When anyone can generate convincing footage or video of a real person saying or doing things that never happened, personal image and reputation effectively become a form of data that needs protection.



We will see growing pressure for explicit consent mechanisms around the use of faces, voices, and biometric features in training and generation workflows.

From a technical perspective, provenance markers, watermarking, cryptographic signatures, and verifiable content registries will be used to distinguish authentic media from synthetic content.

Marko Štefančić, Notch CTO

Cryptographic approaches, including zero-knowledge proofs, will enable verification of authenticity or consent without exposing the underlying private data itself.

At the same time, individuals and organizations will demand stronger legal and technical rights to control, revoke, or restrict how their visual and audio data is used by AI systems.

In an AI-native world, protecting a person's digital image and reputation will become as fundamental as protecting their personal data is today.

Governance

The method is often as simple as taking an LLM and giving it a set of logical "rules" it isn't allowed to break. That matters in 2026 because enterprises are reaching the limits of black-box behavior: 'almost right' answers are expensive when they lead to rework, incidents, or compliance risk.

With the **EU AI Act** pushing documentation, transparency, and oversight, hybrid approaches that can justify decisions and enforce constraints become more attractive. Neuro-symbolic setups are especially useful for agentic workflows, where an LLM proposes actions but a rule layer validates permissions, policies, and business constraints before anything executes.

This directly supports the report's core theme: AI can accelerate delivery, but without governance, it accelerates technical debt. In our own agent thinking (e.g., orchestration patterns like VorteeGo), neuro-symbolic guardrails translate into deterministic checks, structured decision logs, and safer tool use. **By late 2026, neurosymbolic AI will be less niche and more the standard to make AI agents in regulated environments.**



Horizon Technologies


GPUs, Data Centers, Quantum, and the Longer View

If there was a competition for who made the most headlines in 2025, NVIDIA would be the clear winner. The GPU maker hit a market cap over **\$4.4 trillion** just before the 2025 holiday season, dipping slightly but without a significant correction after. This growth was propelled by record orders for their GPUs in an era of mass data center expansion. This surge in GPU demand carried over to other critical data center hardware, most notably RAM. Towards the end of 2025 and in early 2026, the RAM cost growth outpaced all with an 18 month surge for DDR5-6000 32GB cards peaking at close to **\$900 USD**, up from **\$200** per unit just one year prior.

This points towards an expected data center boom in 2026 and beyond driven by demand for AI. How big of a boom?

Currently, there are around 10,500 data centers online globally. Expect this figure to *triple* in order to meet current demand projections for AI by 2030.

This means we can expect hardware costs to maintain current levels until production capacities can grow to meet demand. Expect to see chipmakers to work towards that end, because at present, they’ve sold out stock months ahead of production.



Average RAM Price (USD) Over Last 18 Months (DDR5-6000 2x32GB) - pcpartpicker.com


Source: PC Part Picker

Quantum computing isn't coming to your data center in 2026. But it's worth tracking. Google's Willow chip solved certain problems in under 5 minutes that would take classical computers 10 septillion years. That's real. But it's also a narrow case. Quantum is useful for specific optimization problems, cryptography, and simulation. It's not a general-purpose computer. Enterprises should be thinking about quantum-resistant encryption now. The JDK 24, released in March 2025, included quantum-resistant algorithms. If you're handling sensitive data, you need to be planning for quantum threats.

Of course, we don't expect 2026 to be the year of the quantum big bang. It's not time yet to say that quantum computing is a near-term enterprise disruptor. However, a

post-quantum cryptography migration is, although even that might sound a bit off. So, why is it important? Secured data that is 'harvested' (or stolen) today won't be secure once quantum computing becomes a thing. With NIST's first PQC standards finalized, the ecosystem is already moving.

Cloud providers and platform runtimes are beginning to ship hybrid PQC capabilities at the boundaries, such as KMS endpoints and TLS stacks. So, organizations can begin upgrading without rewriting entire applications. Now, in 2026, is a good time to start building crypto agility and resilience. It's time to inventory where cryptography lives, prioritize long-lived sensitive data flows, and plan phased adoption as standards and platform support mature in the coming months.





From Trends to Action

A Year Of Development Ahead

What This Means For You

What matters in 2026 is not whether you adopt AI; everyone will adopt it. The important thing about 2026 is to have a tight feedback loop for three things: **delivery speed, correctness, and operational outcomes.**

These three things are pretty much disjointed nowadays. If we start thinking in a way that with each release, we move all the important indicators in the right direction, instead of just thinking about whether all the CI passed or if we increased velocity. What are those important indicators? Stability, recovery time, latency, user impact and quality.

Moreover, you don't just track these indicators; you immediately act on them. If you do that, your organization becomes self-correcting instead of velocity-driven.

If you're a **CTO** or **CIO**: technical debt is now a board-level issue. You need a real strategy for managing code quality alongside AI adoption. That's not a technology problem. That's an organizational problem. You need senior architects who can say no. You need governance frameworks that work. You need partners who understand this.

If you're a **Head of AI or Data**: data readiness is your real constraint, not model quality. Start with six to twelve months of data engineering. Get your data cataloged.

Get governance in place. Build for compliance from day one. The enterprises winning with AI are treating data as infrastructure. If you're a product leader: outcomes matter more than velocity. Build quality and sustainability into your roadmap.

Code that falls apart in six months doesn't ship value. Make your outsourcing partner sign up for that.

If you're a CFO: the real cost of AI-assisted development includes the tech debt tax. Budget for refactoring. Budget for senior architect oversight. Speed that doesn't include governance is just debt you're taking on at interest.

How Notch Helps

Custom Software and Enterprise AI – Without the Debt

From AI-driven systems to enterprise software solutions, we guide companies through the entire journey: from discovery and design through development, integration, and long-term innovation. Our team of 170+ engineers, designers, architects, and consultants brings decades of experience building systems that learn, improve, and stand the test of scale.

We believe great software isn't just delivered — it's shaped with care, technical depth, and shared ownership. Our model is built around senior-heavy squads with architects that take responsibility for both the product and its long-term evolution. Code quality, architectural governance, and continuous improvement aren't add-ons — they're part of how we deliver. In 2026, success isn't about adopting AI; it's about making technology self-correcting.

Our approach aligns delivery speed, correctness, and operational outcomes so that every release moves the right indicators forward: stability, recovery time, latency, quality, and user impact. That's how progress compounds: release after release.

Great software doesn't just ship — it keeps getting better. Technology should move you forward without adding friction, risk, or debt.

If you're ready to build something resilient and intelligent, let's make it happen.



As someone looking at 2026 from inside Notch, I'd say we're already covering the things that will matter most in the year to come. People here are experienced, and they've seen what breaks at scale long before the vibe coding era. They've lived through architecture drift, brittle integrations, compliance pressure, incident reality, and the real cost of that let's-call-it-done solution that isn't. In the end, our people have implemented solutions that force you to get serious about governance: traceability, clear contracts, quality gates, and safe ways to change behavior without guessing. That's the difference between a prototype and something you can trust in production.

Marko Štefančić, Notch CTO

PART
FOUR

Verified Market Intelligence & Data Addendum

The Numbers Behind the Trends

We validated our ground-level observations against rigorous forecasts from the world’s leading market intelligence firms. The data confirms the “Governance Paradox”: spending is surging, but capital is shifting decisively from raw creation to management, security, and infrastructure.

The Talent Cliff

The “senior developer shortage” is not a temporary dip; it is a demographic reality. Korn Ferry forecasts a global talent deficit of 85.2 million workers by 2030, risking \$8.5 trillion in unrealized revenue.² In the US alone, the shortage of software engineers is projected to exceed 1.2 million by 2026.

Cloud & Security Velocity

European public cloud spending is forecast to nearly double, from \$229 billion in 2025 to \$452 billion by 2029.⁶ But infrastructure requires protection. IDC projects European security spending to rise 11.8% in 2025 alone, driven by the NIS2 directive.¹⁰ Governance is no longer just a best practice; it’s a rapidly growing line item.

Global & EU IT spending

Gartner forecasts worldwide IT spending to hit \$6.08 trillion in 2026, growing at 9.8% year-over-year. Europe is accelerating even faster than the global average, with IT spending projected to grow 11.1% in 2026 to reach \$1.4 trillion.¹ This isn’t just inflation; it’s a structural pivot toward software and infrastructure that can handle the AI load.

The Nearshore Value Equation

As Western Europe and US rates stabilize at premium levels, Central and Eastern Europe (CEE) remains the efficiency frontier. With a deep talent pool of over 1.5 million developers and cost structures offering 30–50% savings compared to Western Europe, the region is shifting from a “low cost” option to a “high value” strategic necessity.

Talk to Us

The upcoming year is full of challenges. Finding the right partners to develop your software is critical. Talk to us and our experts and see how we can help you discover solutions to these challenges and achieve your business goals in 2026.

Contact us

→

Works cited

1. Deloitte 2025 Global Business Services Survey (<https://www.deloitte.com/hu/en/services/consulting/perspectives/deloitte-global-business-services-survey-2025.html>)
2. Gartner Forecasts IT Spending in Europe to Grow 11% in 2026, accessed December 16, 2025, <https://www.gartner.com/en/newsroom/press-releases/gartner-forecasts-information-information-spending-in-europe-to-grow-11-percent-in-2026>
3. Tackling the \$8.452 Trillion Global Talent Shortage by 2030 – SuperStaff, accessed December 16, 2025, <https://superstaff.com/blog/talent-shortage-by-2030/>
4. Software Developer Shortage in the US – Grid Dynamics, accessed December 16, 2025, <https://www.griddynamics.com/blog/software-developer-shortage-us>
5. US Bureau of Labor and Statistics IT talent shortage – https://www.bls.gov/news.release/archives/ecopro_09062023.pdf
6. EU Public Cloud Market Defies US Tariffs – E3-Magazin, accessed December 16, 2025, <https://e3mag.com/en/eu-public-cloud-market-defies-us-tariffs/>
7. Digital Europe Programme subsidizes up to 50% of qualified SME transformation costs – <https://digital-strategy.ec.europa.eu/en/activities/digital-programme>
8. <https://www.plantemoran.com/explore-our-thinking/insight/2025/03/unlocking-growth-venture-capitals-role-within-an-alternatives-allocation>
9. <https://www.ubs.com/global/en/investment-bank/insights-and-data/2025/secondary-market-activity-h125.html>
10. IDC Forecasts 11.8% Growth in European Security Spending in 2025 Amid Rising Cyberthreats and Regulatory Demands, accessed December 16, 2025, <https://my.idc.com/getdoc.jsp?containerId=prEUR253305325>
11. <https://www.munich-startup.de/en/107155/investment-amount-of-german-startups-in-2024/>
12. <https://www.frenchtechjournal.com/european-vcs-2025-france-raised-more-new-funds-than-uk/>
13. DDRAM prices and trends – <https://pcpartpicker.com/trends/price/memory/>

Sources

1. Stack Overflow 2025 Developer Survey (<https://survey.stackoverflow.co/2025>)
2. Deloitte 2025 Global Business Services Survey (<https://www.deloitte.com/hu/en/services/consulting/perspectives/deloitte-global-business-services-survey-2025.html>)
3. Gartner Magic Quadrant 2025 Custom Software Development Services (<https://newsroom.ibm.com/blog-ibm-named-a-leader-in-gartner-magic-quadrant-for-custom-software-development-services>)
4. GitHub Octoverse 2025 (Data on Python, programming languages) <https://octoverse.github.com/>
5. InfoQ AI-Generated Code Technical Debt Report
6. US Software Development Outsourcing Market Analysis
7. US Bureau of Labor Statistics (https://www.bls.gov/news.release/archives/ecopro_09062023.pdf)
8. Gartner Forecasts Worldwide IT Spending to Grow 9.8% in 2026
9. Gartner Forecasts IT Spending in Europe to Grow 11% in 2026
10. IDC Worldwide Software and Public Cloud Services Spending Guide
11. IDC Worldwide Security Spending Guide (European Forecasts)
12. Korn Ferry: Future of Work – The Global Talent Crunch
13. Grid Dynamics / US Bureau of Labor Statistics Analysis (Developer Shortage)
14. Index.dev: IT Recruitment in Eastern Europe Guide 2025

About Us

Based in central Europe, we're a team of **170+ software** developers, designers, analysts, and consultants who love building awesome digital products.

We offer the whole package – **from user research, product management, and design to engineering, AI, and IoT expertise** – under one roof. We have decades of experience turning user needs into custom solutions that help businesses grow faster and smarter.

Whether you're a scaling startup or an established B2B enterprise, we're here to help you succeed. We have a track record of building long-term partnerships with our customers by becoming crucial to how they think about, build, and deliver software.

Let's build something great together!

Successful Projects

300+

Global Clients

50+

Senior Experts

70%

Rating on Clutch

4.9



Web



GitHub



LinkedIn



Clutch

Awards



50

Technology **Fast 50**
2023 CE **LAUREATE** Impact Star
Deloitte.

