



vertical
compute

A Trillion-Dollar Opportunity to Sustain AI Growth

Sylvain Dubois - Founder & CEO - Vertical Compute



vertical
compute

A Trillion-Dollar Opportunity to Sustain AI Growth

Sylvain Dubois - Founder & CEO - Vertical Compute

01. Context

The rapid rise of Generative AI has unleashed data demands previously unthinkable, requiring GPUs and LLM models to access billions of parameters in real-time. This massive, continuous flow of data is straining current memory technologies to their absolute breaking point, confirming the long-standing memory bottleneck as the single biggest threat to scalable, sustainable computing.

The flawed AI infrastructure trajectory, characterized by the market dominance of very few semiconductor and cloud players and the massive energy consumption required is creating significant societal, economic, and geopolitical challenges.

Vertical Compute proposes a pivotal alternative, pioneering a novel memory technology and a 3D Vertical Integrated Memory architecture (VIM™ chiplets) that provides the essential hardware foundation for next generation of computing.

This technology offers a crucial opportunity to influence the current path to serve AI to billions of users and autonomous devices, enabling a more distributed, efficient, and balanced infrastructure to sustain the exponential growth of AI.

02. About us

Vertical Compute is a European deep-tech startup dedicated to resolving the most critical bottleneck in modern computing: the memory bottleneck.

The Problem We Solve: We tackle the chronic performance limitation caused by the slow, inefficient data transfer between separate processing cores and legacy 2D memory like DRAM. This structural flaw severely hinders the scalability and efficiency of cutting-edge AI and high-performance workloads.

Fundamental New Technology: We are developing a fundamental new way to store and process data in a chip. This represents a radical shift from conventional technology, moving from horizontal, planar memory to a novel 3D memory architecture. We pioneer the use of vertical columns of data directly connected to processing cores.

Our breakthrough technology is delivered via chipllets and blends advanced vertical memory concepts with nano-magnetism. This approach fully integrates memory and compute to unlock unprecedented density, speed and energy efficiency.

Applications Enabled: By eliminating the data movement constraint at the source, we accelerate the AI revolution in both the datacenter and at the edge. Our chipllets empower AI architects and industry leaders to scale immense models in the cloud, while also unlocking new applications for physical AI such as robots, autonomous devices, and the post-smartphone era devices (e.g. smart glasses, AI bracelets or robots).

We shift the performance boundaries across all industries.

03.

Which problem do we address

The problem Vertical Compute is solving is the memory bottleneck, the fundamental performance and efficiency limitation caused by the slow, power-hungry data transfer between processing cores and legacy 2D memory technologies like DRAM, which is severely constraining the scale and speed of modern AI and high-performance workloads.

04.

What is the solution and impact on the industry

Vertical Compute's innovation involves a vertical memory architecture that allows memory cells to be stacked directly above processing units in a monolithic silicon wafer. This dramatically cuts the physical distance data has to travel. Shorter distance means faster processing times and massively lower energy use.

To achieve this, VIM™ merges the best principles of three innovative worlds.

- At its core, our memory stores bits of data using magnetism, making it extremely fast (like the fastest current memory, called SRAM) and non-volatile (meaning it saves your data even when the power is off).
- Simplified Structure: The VIM™ structure doesn't require the complex wiring and "staircases" that traditional monolithic 3D stacking requires. This drastically simplifies the overall design, making the chip cheaper, denser, and easier to manufacture.
- Endurance: We use a well-known magnetic component (MTJ) solely for reading data, using very low power. This clever step ensures the memory has an almost unlimited lifespan (read/write endurance).

Because our unique structure can be built using the same high-volume equipment currently used by the world's leading chip foundries, the final manufacturing cost for our VIM™ device will be very low—similar to widely produced storage like 3D-NAND.

This combination of low cost and radical performance unlocks enormous potential:

- It enables truly private, always-on AI assistants right inside small devices.
- It slashes the enormous cost and power consumption of data centers.
- It fundamentally redefines the memory hierarchy, paving the way for a more powerful, accessible, and sustainable next generation of global computing.

05.

Targeted market and applications

At Vertical Compute, our best move is to deliver a game-changing technology to fix the current memory bottleneck of AI chips. This will give a unique differentiator to our customers and a tremendous power to the companies and territories that will have access to this technology.

Cloud & Hyperscale Revolution

By replacing legacy 2D memory, Vertical Compute enables up to a 100x performance gain for large language models (LLMs) and delivers up to 80% energy savings for data-intensive cloud workloads. This is achieved by cutting the data transfer distance from centimeters to nanometers, critically reducing the massive carbon footprint of today's hyperscale data centers.

Physical & Edge AI Enablement

Our technology unlocks true, real-time autonomy for the next generation of small form factor, power-constrained devices. It allows complex AI models to run locally on devices like robots, smart glasses, and autonomous vehicles. This ensures data privacy by eliminating the need to send sensitive information to the cloud, making high-performance AI ubiquitous at the edge.

Global Innovation

Vertical Compute breaks the current stagnation of legacy 2D memory scaling, providing the fundamental hardware needed to re-establish an aggressive performance curve. This deep tech innovation drives a new wave of global technological advancement and provides the necessary foundation for the future of compute and memory integration.

06. Business model and strategy

As a deep-tech semiconductor chiplet provider focused on the memory bottleneck, here is the detailed breakdown of Vertical Compute's business model and strategy

#1. Target Markets, Funding, and Monetization

Vertical Compute's primary customers are large semiconductor giants and AI architects. This includes Hyperscalers (like Google, Amazon, Microsoft) designing custom silicon, and major Chip Designers (like Nvidia, AMD, Qualcomm) that need to power both Cloud/Datacenter AI (for LLM acceleration) and the next generation of High-Performance Edge/Physical AI (for autonomous systems and smart devices).

Our development is financed through a typical deep-tech model, relying on Venture Capital (VC), specialized Strategic Investment, and key Regional and European R&D grants to fuel capital-intensive silicon development.

Monetization is achieved through a hybrid model tied to the value of our proprietary VIM™ Chiplet IP and Technology. Revenue streams include:

- Pre-product revenue fees: Upfront Non-Recurring Engineering (NRE) fees paid by partners for the initial design integration.
- Per-unit fees: A scalable revenue stream where we charge a fixed fee for every chiplet or final product sold that incorporates our technology.

#2. Market Strategy & Implementation

Our strategy is a focused Vertical Go-to-Market (GTM) approach targeting high-value partners:

1. **Initial Focus:** We will target one or two top-tier partners (our beachhead) in the Hyperscaler or leading Edge AI space facing urgent memory bottleneck problems. Securing a design win with a major player is critical for validating the technology and achieving scale.
2. **Proof of Value:** We prioritize demonstrating the 100x performance gain and 80% energy efficiency in real-world proof-of-concept (PoC) designs to rapidly persuade risk-averse silicon partners.
3. **Ecosystem Enablement:** Vertical Compute positions itself as an enabler, focusing exclusively on the chiplet building block. This allows us to integrate with diverse established architectures (CPU, GPU/xPU, ASIC) without becoming a competitor to our potential customers.
4. **Roadmap Alignment:** We work in deep collaboration with customer roadmaps to ensure our next-generation chiplets are ready precisely for their silicon tape-out cycles.

#3. Main Distribution Channel

Our main distribution channel is a Direct B2B Partnership Channel underpinned by the global Semiconductor Foundry Model.

The channel is Direct from our executive and technical sales teams to the engineering and advanced sourcing teams of our Tier 1 customers. This is a highly technical, high-touch sales process.

For physical delivery, manufacturing is outsourced. After transferring the technology manufacturing process, we provide the VIM™ chiplet design files to major, independent Semiconductor Foundries (e.g., TSMC, Samsung). The final customer then manages the integration of our chiplet onto their processor using advanced packaging techniques. Our success relies on securing design wins and ensuring our technology is seamlessly manufacturable at scale.

07.

Key research, industry and business partnerships

Vertical Compute's strategy relies on a focused network of partners across R&D, manufacturing, and customer segments to bring our radical 3D memory technology to market.

#1. Research & Development Partners

Given the deep-tech nature of our solution, our primary R&D focus is on bringing our novel memory concepts to silicon reality:

Leading European Nanotech R&D Hubs (e.g. Imec):

We engage in Joint Development Agreements (JDAs) focused on process development and integration. This is critical for translating our vertical memory and nano-magnetism concepts into a manufacturable silicon process flow. These partners provide access to cutting-edge foundry tools, pilot lines, and the scientific expertise needed to validate our core IP.

Academic Institutions (e.g. Spintec):

We maintain partnerships with top European universities for fundamental research into next-generation VIM™ architecture and materials, ensuring a robust technology roadmap and providing a vital pipeline for recruiting specialized engineering talent.

#2. Ecosystem & Manufacturing Enablers

To scale from design to commercial product, we integrate deeply into the global semiconductor supply chain:

Major Semiconductor Foundries (e.g. TSMC, ESMC, GF):

We establish relationships to evaluate process compatibility and secure contracts for high-volume wafer fabrication. This outsourced manufacturing strategy is key to being capital-light while ensuring scalable, reliable production.

Advanced Packaging Specialists (e.g., Amkor, ASE):

We co-develop the advanced packaging solutions necessary for heterogeneous integration (e.g. 2.5D/3D stacking). This is crucial for making our chiplets physically compatible and easily integrated with a customer's existing processors and accelerators.

#3. Strategic Business & Customer Partners

These relationships are essential for pre-product monetization and co-design collaboration:

Hyperscalers and Cloud Titans (e.g. Google Cloud, AWS, Microsoft Azure, Meta):

These are our highest-priority targets. We focus on securing an initial Design Win through Joint Development Agreements (JDAs) where we co-develop a Proof of Concept (PoC) integrating our VIM™ chiplet into their next-generation custom silicon (like their internal AI accelerator roadmap). This leads directly to long-term collaboration and business agreements.

Tier-1 AI Chip Designers (e.g. AMD, Qualcomm, Nvidia):

We work with their architecture teams for roadmap alignment, strategic evaluation and integration. The goal is to prove that our chiplet solves their core memory bottleneck, offering superior power and performance relative to High-Bandwidth Memory (HBM), and securing a commitment for adoption in their future product lines.

By maintaining these robust connections across the entire silicon value chain, Vertical Compute ensures it has a pathway from invention to integration to market revenue.

08. Global ambition, local conviction

Semiconductor is the backbone of the economy and our political leaders realized the importance of relocating chip design and manufacturing back to their respective regions with the financial support of the \$50 billion US Chip Act or the €50 billion EU Chip Act.

When we consider that 80% of the investment cost of building an AI data center is allocated to semiconductors, the announcements of a \$500 billion Stargate initiative in the US and a €100 billion investment in France take on a new significance.

We are actively pursuing strategic options to secure reliable and resilient back-end semiconductor manufacturing. Vertical Compute brings an opportunity to bring back some territories at the center of the semiconductor value chain.

09.

Business opportunity

Vertical Compute's opportunity is measured in the trillions because it doesn't just improve memory; it fundamentally restructures the massive global AI and data center infrastructure, which is currently gated by memory cost and energy.

The trillion-dollar valuation potential is realized by vertically capturing value from three massive, interconnected markets that your technology disrupts: DRAM/HBM market capture, power savings in Cloud infrastructure, and enabling new AI market segments.

Market Size: The DRAM/HBM Foundation

Vertical Compute's technology directly addresses and aims to capture market share from the existing memory and specialized memory markets:

- **Total DRAM Market:** The global DRAM (Dynamic Random-Access Memory) market is currently valued at \$100 billion annually, projected to grow to \$200 billion by 2030. This represents the baseline, high-volume memory that our solution aims to eventually displace or redefine in compute systems.
- **High-Bandwidth Memory (HBM) Market:** The high-performance, stacked memory market (HBM) is growing rapidly and is expected to be worth \$10 billion by 2027, projected to grow to \$100 billion by 2030, specifically driven by AI/GPU accelerators. This is our high-value target market, as VIM™ chiplets will offer a superior alternative to HBM for bandwidth and efficiency.

The Trillion-Dollar Opportunity is not just the memory market itself; it is the Total Addressable Value (TAV) of the AI infrastructure and data center cost Vertical Compute fundamentally lowers and enables:

- **CapEx Savings:** We dramatically reduce the need for Hyperscalers to continuously build new data centers by increasing compute density in existing racks.
- **OpEx Savings:** By offering up to 80% energy savings, we slash the fastest-growing multi-billion dollar operating cost (electricity and cooling) of the AI economy.
- **New Market Enablement:** We unlock the entire Physical AI market (robotics, smart glasses) that is currently impossible due to memory power and size constraints.

The Economic Mechanism: Vertical Value Capture

The trillion-dollar potential stems from the multiplier effect of solving the memory bottleneck:

Value Proposition: Vertical Compute is selling density, energy savings and performance improvements, not just memory chips. By offering 80% energy savings and 100x performance gains on AI workloads, you allow Hyperscalers to dramatically reduce their multi-billion dollar Operational Expenditure (OpEx) and fit more high-revenue compute into their current Capital Expenditure (CapEx) footprint. The value of the power saved alone over a decade of operation in a data center exceeds the cost of the chips many times over.

Value-add Product Sales: As the provider of foundational memory blocks via chipllets product sales and potential technology licensing, Vertical Compute captures a highly profitable revenue stream from every single advanced AI system (ASIC, GPU, TPU) produced by major customers for the next 10-20 years.

Valuation trajectory: From IP Pioneer to Infrastructure Giant

Vertical Compute's valuation trajectory follows the proven path of companies that provided indispensable, proprietary technology for generational shifts:

Phase 1: IP Pioneer (SanDisk Parallel):

Early valuation gains will mirror pioneers like SanDisk, who achieved massive scale by inventing and licensing foundational memory IP (flash) that enabled new markets. Vertical Compute is contemplating the same business strategy for the 3D VIM™ chipllet and technology IP, driven by validated silicon and securing high-value, recurring royalty contracts.

Phase 2: Long-Term Dominance (SK Hynix / Nvidia Parallel):

- When VIM™ becomes the new default standard for high-performance memory, displacing or dominating the HBM market, Vertical Compute achieves a valuation comparable to SK Hynix, securing a dominant, multi-billion dollar revenue stream tied to essential global infrastructure.
- When the technology is so disruptive that it dictates how all future AI compute architecture must be designed, Vertical Compute will capture a highly strategic gatekeeper position similar to Nvidia. We become indispensable to every player, commanding a multi-hundred-billion-dollar capitalization tied to the sustained, essential growth of the entire AI sector.

Investing now is securing equity in the future AI compute infrastructure, making Vertical Compute the strategic solution to the global AI market and the key to unlocking the trillion-dollar opportunity.

10. About the executive founding team

Sylvain Dubois - CEO & Co-founder, Vertical Compute | Deep Tech Executive

For 25 years, Sylvain has been dedicated to solving one of computing's most critical challenges: the memory bottleneck. His career is a cohesive journey that tackles this problem from every angle:



- **Architect:** Starting at Texas Instruments, where he designed the integration of DRAM and flash controllers into mobile phone SoCs.
- **Innovator:** Shaped the flash memory roadmap at Spansion/Cypress.
- **Business Leadership:** Championing the commercialization of new memory technologies as a VP at Crossbar.
- **Strategic Sourcing** at Scale: Finding and integrating cutting-edge solutions (AI hardware, chiplets, new memory) for the Google Cloud Infrastructure and TPU roadmap.

Today, as CEO and Co-founder of Vertical Compute, he leverages this comprehensive, 25-year expertise to drive innovation and deliver the next generation of computing for AI. He is passionate about the intersection of compute and memory and building disrupting technologies that redefine performance boundaries and unlock new applications.

Sebastien Couet - CTO & Co-founder, Vertical Compute | Deep Tech Scientist

For over 15 years, Sebastien has been dedicated to fundamentally advancing the physics and materials required to break the memory wall, focusing on non-volatile and integrated memory solutions. His career path evolved from cutting-edge science research into manufacturable, disruptive IP:



- **Foundational Scientist:** Early research and Ph.D. work in applied physics and magnetic materials, establishing expertise in the fundamental principles of next-generation data storage.
- **Deep Tech Program Director:** Spent over a decade at imec, leading a core research program and directing strategy for advanced non-volatile memory technologies, including MRAM and novel 3D integration concepts.
- **Inventor & IP Creator:** Personally conceived and developed the patented VIM™ technology, successfully bridging nanotech physics with complex semiconductor manufacturing processes.

Today, as CTO of Vertical Compute, he guides the productization of this groundbreaking 3D memory architecture into high-volume chiplets, defining the roadmap to scale our solution across the global AI ecosystem.

11.

Conclusion: The Mandate for Asymmetric Investment

This white paper demonstrates that Vertical Compute is not just developing a new chiplet component; we are enabling the unavoidable architectural mandate for the next era of computing.

The current AI trajectory is heading toward an unsustainable crisis driven by the fundamental failure of 2D memory technology scaling limitation. This failure is creating brute force investment with legacy technologies driven by the market dominance of very few players in AI infra and creating systemic societal, economic, and geopolitical pressure.

Vertical Compute provides the only viable strategic off-ramp. Our 3D Vertical Integrated Memory Architecture (VIM™ chiplets) is the most transformative technology that fundamentally solve the memory bottleneck at the nanoscale, delivering massive performance gains and reduction in power consumption for AI workloads.

This is not a marginal improvement; it is the foundational IP required to unlock a trillion-dollar market, the only path to truly sustain AI growth in the cloud, enable the Physical AI revolution at the edge, and secure a sustainable future for global compute infrastructure.

The time for incremental fixes is over. Investing in Vertical Compute is an asymmetric opportunity to capture the uncapped returns of the next semiconductor generation. We are not just innovating; we are establishing the new industry standard.