



Build or Buy?

We are entering an era of pervasive AI, when models are deployed to support, accelerate, and transform every job in every industry. But AI is complicated. Behind every user-friendly multi-billion parameter model lies months of training on hundreds if not thousands of compute cores.

And that is just the beginning. Those chips need software like runtime and compilers, platform management to handle data ingestion, model training and serving, models, and enterprise integration. Foundation models are giving everyone a helping hand, but the most accurate models demanded by enterprises will still require training on private data, and most organizations will still want to own their models, just as they own their data, and build their AI capability.

The question is where to begin. If – like most enterprises – you want to invest in a more strategic approach rather than buying multiple point solutions or you simply cannot buy SaaS AI systems due to regulations or needs around model ownership and explainability, you are left with two choices: To build your own stack from scratch, or to buy one that's pre-built by experts and ready to use.

At SambaNova, we are firm believers in the full stack; for enterprises and governments looking to drive rapid gains and build for the long-term, we offer SambaNova Suite – our full-stack solution that's optimized for enterprises and governments. For those users whose primary objective is the research and development frontier, we offer DataScale.

So, why do we recommend the full stack approach?

Simplicity.

Buying and deploying an all-in-one solution is straightforward and fast. We can configure to fit any customer's needs, host immediately, or ship in 14 days, and the customer is up-and-running, supported by our world-class crew of engineers and support specialists. Contrast that to identifying, procuring, receiving, installing, interconnecting, and debugging multiple systems – even when supported by a multitude of well-meaning vendors. Full stack is simple.

Performance.

Real-world AI performance (for actual workloads, out of labs, away from individual benchmarks) depends on the interoperability of every component in the AI stack. We are at the beginning of the AI revolution. Behind all of the seemingly simple user-friendly AI interfaces lie stacks of technology and incredibly complex models, all developed and managed by highly skilled engineers. Making all of those components work together is time-consuming and difficult. The fastest route to high performance is with a stack that is pre-optimized. In a few years, we may get to the point of plug-and-play AI components but we are far from that right now. The fastest way to the highest performance is with an integrated system, pre-built and pre-tested for the precise workloads that enterprises need.



Time to value.

Getting chips is hard. Finding, hiring, and retaining AI specialists is harder. We have pre-built systems that can be configured fast, and deployed for customers within minutes. Even when a customer wants it in their own datacenter, we can be up and running in 14 days. Contrast that with building your own AI practice and technology stack – we've calculated that our approach is between 4 and 20x faster – and when every day of customer value matters, that's a big difference.

Total cost of ownership.

Every component comes with a cost, and every independent vendor is seeking a mark-up. We design our own silicon (manufactured by TSMC), build our own software, customize our own foundation models, and sell a complete package. As opposed to a consumption pricing model that becomes unsustainable as AI becomes pervasive across your enterprise, our capacity-oriented model aligns with AI operational maturity. Our customers don't pay by the use – even the training runs that fail. They don't pay multiple markups and extravagant integration costs. When you add it all up, our systems are at least half the cost of comparable alternatives.

No lock-in.

'Buying' often comes at a cost - multiple dependencies and vendor lock-in. That's why we build on open standards. As a result, interoperability with your existing workflows and ramp up time for your employees is fast and easy and models are portable - away from SambaNova if you like. End users and administrators can quickly become productive within minutes because they can leverage their existing experience and work with widely-available open source software and models. You also gain full visibility and explainability, seeing every aspect of how our systems operate, to put the business, customers, and regulators at ease.

Buy > build

In this rapidly evolving landscape of pervasive AI, the choice between building or buying an AI solution is a critical one. SambaNova's full-stack approach is a compelling solution for businesses aiming to deploy AI immediately, efficiently, and effectively. It offers simplicity, performance, and rapid deployment, reducing time to value and improving cost-effectiveness. Building on open standards means you can benefit from the total knowledge of this rapidly-developing field - and are never locked-in to a single vendor. SambaNova empowers organizations to embrace AI with confidence, streamlining their AI journey, and ensuring long-term success in a world of pervasive AI. Explore the possibilities of SambaNova Suite and embark on a transformative AI path for your business.

About SambaNova Systems

Customers turn to SambaNova to quickly deploy state-of-the-art AI capabilities to gain competitive advantage. Our purpose-built AI platform is the technology backbone for the next generation of AI computing. We power the foundation models that unlock the valuable business insights trapped in data.



SambaNova.ai



2200 Geng Rd. Unit 100, Palo Alto, CA 94303
(650) 263-1153
info@sambanova.ai