

SambaNova DataScale[®] SN30

The Platform for Innovation

Features:

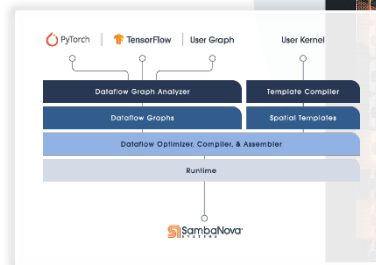
- SambaNova Reconfigurable Dataflow Architecture™ (RDA)
- State-of-the-art SambaNova Reconfigurable Dataflow Unit™ (RDU)
- Terabytes of memory
- SambaNova SambaFlow™
- Subscription pricing

Benefits:

- Unprecedented performance for the most demanding applications
- Ability to run the largest models and process the largest data
- Ease of use. Users interact with common frameworks with no need to use proprietary programming languages
- Automatically extracts, optimizes, and executes the optimal dataflow graph of your models on SambaNova's Reconfigurable Dataflow Units
- Seamless scalability across systems with consistent rack-rack bandwidth and latency
- Complete hardware and software system for training and inference



Cardinal SN30




SambaFlow Software



DataScale SN30

SambaNova DataScale[®] is a fully integrated hardware-software system that enables organizations to train and deploy the most demanding deep learning, foundation model, and AI for Science workloads and achieve world record performance with the largest and most challenging models such as GPT.

DataScale's dataflow computing architecture combined with large on-chip and system memory enable organizations to manage state-of-the-art models which cannot be practically handled by GPUs, such as:



Very large models, such as 100B+ parameter large language models



Very large data, such as ultra-high resolution 2D and 3D image data



Very detailed models, such as RNNs, sparse models, and AI for Science workloads

Powered by the state-of-the-art SambaNova Reconfigurable Dataflow Unit, DataScale delivers world record time to train performance and world record GPT accuracy, combined with unmatched ease-of-use, to power the most advanced deep learning, foundation model, and AI for Science workloads.


DataScale reduces the cost and complexity of AI initiatives by providing a fully integrated system, from silicon to software, that is designed to deliver a unique combination of unrivaled performance, unprecedented ease of use, and seamless scalability.

Built with the SambaNova Reconfigurable Architecture

DataScale is a rack-level data center solution that can be installed and operating in only minutes. The SambaFlow complete software stack delivers the only practical solution to running complex AI and deep learning workloads without significant parallelization efforts and enables users to run their own models without any proprietary coding languages. Users interact with common open source ML frameworks, such as PyTorch and TensorFlow. Low-code/no-code APIs dramatically simplify management with no proprietary knowledge required.

The World's First Native Dataflow Processor


SambaNova DataScale is built with the SambaNova Reconfigurable Dataflow Unit™ (RDU), the industry's next-generation processor built from the ground up to provide native dataflow processing. Features include:



SambaNova Reconfigurable Dataflow Architecture eliminates the constant data caching and excess data movement deficiencies inherent to today's core-based architectures, unlocking unrivaled efficiency.



Optimized data computations and communications result in high performance at optimal efficiency out-of-the-box across models of all sizes and forms, and for any batch size across training and inference.



100s of MBs of on-chip memory, 100s of TFLOPS of processing power, and direct interface to TBs of off-chip memory enable developing and deploying more sophisticated models with richer context than what's possible on a GPU and without the extensive parallelization efforts required by GPUs.



100s of TFLOPS of compute delivering unmatched acceleration to drastically reduce time to results.

The Industry's Most Advanced Software

SambaNova DataScale features SambaFlow, a complete software stack designed to take input from standard machine learning frameworks.

- **Fully integrated with popular open source ML frameworks, such as PyTorch.** No code modification is required to run.
- **Push-button model compilation, optimization and execution** enables high performance out-of-the-box without the need for low-level tuning.
- **Automated data and model parallel mapping** simplifies scaling by using the same programming model as on a single device — no special programming required.
- **Secure multitenancy and concurrent multigraph execution** provides seamless scale-up and scale-out flexibility to maximize compute and memory resource utilization with no waste.
- **The latest productionized algorithms** eliminate months of tuning and optimization to allow you to elevate your focus on what matters most — the application.

Model Specifications: DataScale SN30 Configurations

Components

Rack system	DataScale SN30 <ul style="list-style-type: none"> Up to 3 DataScale SN30-8 systems
Compute	<ul style="list-style-type: none"> DataScale SN30-8 system, each with: <ul style="list-style-type: none"> 8 x Cardinal SN30™ RDUs 8 TB total memory
Networking	<ul style="list-style-type: none"> High performance 400/200 GbE data switch General purpose 1 GbE access and management switch, 48-port Serial console manager server
Additional hardware components	<ul style="list-style-type: none"> 42RU rack 4 x redundant power distribution units (PDUs) Networking cables
Software	<ul style="list-style-type: none"> SambaFlow software Red Hat Enterprise Linux OS Ubuntu Linux OS

Environmental Specifications

	1 node	2-node	3-node
System Dimensions	Height: 78.5" (1994 mm) Width: 23.6" (600 mm) Depth: 50" (1270 mm)		
Operating Temperature	32° F to 86° F (0 C to 30° C)		
Operating Humidity	20% to 80% (non-condensing)		
Operating Altitude	Up to 9842ft (3000m); derated by @ 1.8F (1C) per 984ft (300m) above 2952ft (900m)		
System Weight (without packaging)	733 lbs (333 kg)	1051 lbs (477 kg)	1369 lbs (621kg)

The Platform for Innovation SambaNova Systems DataScale is an integrated system optimized for dataflow from algorithms to silicon. SambaNova DataScale is the core infrastructure for organizations that want to quickly build and deploy next-generation AI technologies at scale. Built on SambaNova Systems Reconfigurable Dataflow Architecture, SambaNova DataScale enables you to achieve unparalleled efficiency and performance across a broad range of applications, including deep learning, foundation models, AI for Science and more. SambaNova DataScale is built with open standards and interfaces to seamlessly integrate into your existing infrastructure and environment— without disruption. With flexibility and efficiency, you can stay current with rapidly changing demands as new breakthroughs emerge.




SambaNova is Redefining AI Boundaries

To learn more about how SambaNova Systems DataScale can accelerate and transform your organization with AI, **schedule a meeting**.

Learn more at SambaNova.ai

 [linkedin.com/company/sambanova](https://www.linkedin.com/company/sambanova)

 @SambaNovaAI

 info@sambanova.ai

Customers turn to SambaNova to quickly deploy state-of-the-art AI capabilities to meet the demands of the AI-enabled world. Our purpose-built enterprise-scale AI platform is the technology backbone for the next generation of AI computing. We enable customers to unlock the valuable business insights trapped in their data. Our flagship offering, Dataflow-as-a-Service™, overcomes the limitations of legacy technology to power the large complex foundation models that enable customers to discover new services and revenue streams, and boost operational efficiency. Headquartered in Palo Alto, California, SambaNova Systems was founded in 2017 by industry luminaries, and hardware and software design experts from Sun/Oracle and Stanford University. Investors include SoftBank Vision Fund 2, funds and accounts managed by BlackRock, Intel Capital, GV, Walden International, Temasek, GIC, Redline Capital, Atlantic Bridge Ventures, Celesta, and several others.