

SambaNova Systems DataScale

Break Through the Limits of your GPU

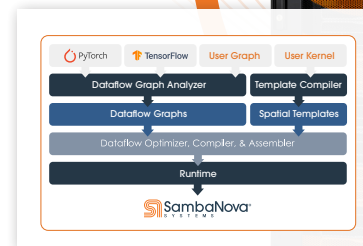
The Platform for Innovation

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™ (RDA), SambaNova DataScale enables you to achieve unparalleled efficiency and performance across a broad range of applications, including training, inference, data analytics, High Performance Computing (HPC), 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.



Cardinal SN10



SambaFlow Software



DataScale SN10-8R

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 and TensorFlow.** 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 multi-tenancy and concurrent multi-graph 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.


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 TFLOPS of compute** delivering unmatched acceleration to drastically reduce time to results.
- **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.

Model Specifications: DataScale SN10-8 Configurations

SN10-8 System: 8 x SambaNova Cardinal SN10 RDUs. 3 TB, 6 TB, or 12 TB memory



Components

Switches	<ul style="list-style-type: none">• High performance 100 Gb / second data switch• High performance, scalable 1 Gigabit Ethernet access switch, 48-port• Serial console manager server
Catalog options	<ul style="list-style-type: none">• 42RU rack• 3 x redundant power distribution units (PDUs)• Networking cables
Software	<ul style="list-style-type: none">• Red Hat Enterprise Linux OS• Ubuntu Linux OS• SambaFlow™ software

Environmental Specifications

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)	Quarter Rack: 725 lbs (330 kg) Half Rack: 1035 lbs (470 kg) Full Rack: 1655 lbs (752 kg)

[Learn more at SambaNova.AI](#)