

Installed in November 2019, the Cray CS-Storm serves as HLRS's primary machine to accelerate artificial intelligence (AI) workloads. The machine complements HLRS's high-performance computing infrastructure by addressing user demand for processing-intensive applications in the realms of machine learning and deep learning, and it supports a wide variety of well-known and established AI frameworks and tools, such as Apache Spark, Python-based data science libraries like scikit-learn, and frameworks steered toward deep learning, including TensorFlow and PyTorch. The system also includes 64 NVIDIA Tesla V100 GPUs. Using the Cray Urika-CS AI and analytics suite, the CS-Storm enables HLRS users to accelerate AI workloads, arming users to address complex computing problems, and process data with higher accuracy.

Cray CS Storm—Compute (Deep Learning) Partition

8 GPU Nodes. Each node in the following configuration:

- 8x V100 SXM2 32GB HBM2, NVLink 2
- 2x CLX 6240, 18c, 2.6 GHz (150W)
- 24x 32 GiB DDR4-2933; 768 GiB total
- 4x P4510, NVMe SSD, 2.5", 2 TB
- 2x S4510, SATA SSD, 2.5", 240 GB
- 4x Mellanox CX-4, x8, VPI Single-Port, QSFP28



Cray CS-Storm (Image courtesy of Cray)

Cray CS500—SPARK (ETL/ML) Partition

8x CPU nodes in 2 CS500 3211 (2HE). Each node in the following configuration:

- 2X CLX 6230, 20c, 2.1 GHz (125W)
- 12x 32 GiB DDR4-2933, 384 GiB total
- 1x P4510, NVMe SSD, 2.5", 2 TB
- 2x S4510, SATA SDD, 2.5", 3.8 TB
- 1x S4510, SATA SDD, 2.5", 240 GB
- 1x Mellanox CX-6 HDR100, 100 GB/s On-board
- 2x SFP+, 10 GB/s

Software, Compiler

- Urika-CS AI Suite

Network

- HDR100 Infiniband

INFORMATION

