


Silicon Mechanics Miranda CDI Cluster

Cloud-Like Flexibility On-Premises with Disaggregated Infrastructure



Building out data center environments, especially for HPC and AI workloads, has always been about precision. If done well, each server has the right mix of compute, acceleration, storage, and memory to perform a specific task optimally.

But what happens when that task is done? Or what if organizational focuses change?

Traditionally, those systems get reassigned, reconfigured, or overhauled to support the next project, and are no longer optimized to the workload. If you can maintain optimal configurations for different workloads, you can dramatically improve the ROI of a system by extending its optimal performance period.

So, how do we improve the flexibility of a custom-designed, bare-metal cluster to extend its life, maximize performance, and support a variety of use cases?

Composable Disaggregated Infrastructure (CDI).

The Silicon Mechanics Miranda CDI Cluster lets you dynamically provision all its resources from the software layer without virtualization. Welcome to a whole new world of data center design.

Benefits

- Dynamically composable nodes mean simple administration and optimization of resources
- Ideal for multi-tenant environments and mixed-use systems
- Effortlessly scalable, just add more servers, JBODs, or JBOXs to expand any resource pool
- Create a cloud-like environment for users without the performance losses, latency, or vendor lock in
- Includes best-in-class components like NVIDIA Ampere GPUs, AMD EPYC CPUs, and more

Ideal Use Cases

- Multi-tenant Environments
- HPC and Simulation
- AI and Machine Learning
- Cloud-like Computing
- Engineering and Visualization
- VFX and Digital Production

Relevant Industries

- Financial Services
- Life Sciences
- Government
- Academia
- Media & Entertainment
- Oil & Gas

Base Design Node Components

COMPUTE

- 3rd Gen AMD EPYC™ 7713P
- NVIDIA® A100 GPU

MEMORY

- 8 DIMM slots and up to 2TB 3DS ECC DDR4-3200MHz RDIMM/LRDIMM

STORAGE

- 2 PCI-E 4.0 x16 and 1 NVMe/SATA M.2

NETWORKING

- 1x 200G HDR NVIDIA InfiniBand switch
- Up to 2 hot-pluggable 10G NVIDIA Spectrum SN4000 Open Ethernet Switches
- NVIDIA 10GBASE-T Management Switch

Composable Disaggregated Resources

Option A

- Liquid Grid 48 Port Gen 4 fabric switch
- Liquid Command Center or third-party software

Option B

- Giga IO FabreX TOR switch
- FabreX CLI or third-party software

Quality Assurance for Zero Defects

We build each of our systems to “zero defect” standards in our U.S.-based manufacturing facilities. Then we hand-inspect every order, testing them to ensure they are 100% operational and optimized to support rapid deployment.

Standard 3-Year Warranty

We offer a comprehensive 3-year warranty standard, with every system purchased. But you can add extended or custom warranties if your situation calls for it.

In-House Support

We offer customer support at different levels, each based on a detailed service level agreement (SLA) that fits your needs.

About Silicon Mechanics

Silicon Mechanics is one of the world’s largest private providers of high-performance computing (HPC), artificial intelligence (AI), and enterprise storage solutions. Since 2001, Silicon Mechanics’ clients have relied on its custom-tailored open-source systems and professional services expertise to overcome the world’s most complex computing challenges. With thousands of customers across the aerospace and defense, education/research, financial services, government, life sciences/healthcare, and oil and gas sectors, Silicon Mechanics solutions always come with “Expert Included”SM. Learn more at www.siliconmechanics.com.

ⁱ MLN-016: Results as of 01/28/2021 using SPECrate®2017_int_base. The AMD EPYC 7763 measured estimated score of 798 is higher than the current highest 2P server with an AMD EPYC 7H12 and a score of 717, <https://spec.org/cpu2017/results/res2020q2/cpu2017-20200525-22554.pdf>. OEM published score(s) for EPYC may vary. SPEC®, SPECrate® and SPEC CPU® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information.

ⁱⁱ <https://www.nvidia.com/en-us/data-center/a100/>