



Enflame Cloud Blazer POD P2110

64x Yunsui i20 GCU-CARE rack

OVERVIEW

—	1.0 TB	—	—
Peak Dense	Aggregate Memory	Aggregate Bandwidth	Aggregate Power

CONFIGURATION

Accelerator	Yunsui i20
Count	64
Memory per GPU	16 GB
Bandwidth per GPU	819 GB/s

POWER & EFFICIENCY

Aggregate Power	—
TDP per GPU	150 W
Efficiency	—
Aggregate Bandwidth	—

SYSTEM DETAILS

Vendor: Enflame	Family: Cloud Blazer POD	Architecture: GCU-CARE
Form Factor: rack	Release Year: —	Process Node: 12nm

INTERCONNECT

Two 25Gb/s links between compute nodes, as filed. The prospectus describes the cabinet as four to eight AI servers and several network switches joined by high-speed cabling; it gives no scale-up fabric for the accelerators.

NOTES

Enflame's single-cabinet inference appliance, branded Cloud Blazer POD and described in the company's STAR Market prospectus as a single-cabinet, single-compute-node device that forms the smallest unit of a compute cluster. It integrates 64 Yunsui i20 accelerators across four to eight AI servers, each server carrying eight cards. The filing lists it for inference, air-cooled, with a PUE of 1.25 or better. Its 1TB of accelerator memory matches 64 cards at 16 GB exactly. Enflame publishes no throughput figure for this product or for the cards inside it.