


GAINBOARD™ 2803 Inference Server AI Server Accelerator Solution

Industry's Best Ratio of Performance-to-Power

The low-power and high-performance GAINBOARD™ 2803S AI Inference Accelerator Server solution for data centers provides a rapid upgrade path to support the growing computational power demands of AI applications while being cost-efficient.

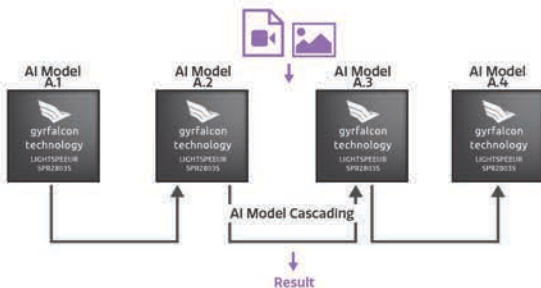
A pre-configured system delivers up to 9.9 TOPS/Watt, about 3x more efficient than that of competitors. With the industry-best ratio of performance-to-power, the GAINBOARD™ 2803S helps data center operators improve energy usage and increase power savings.

Massive Leap in Performance



Supports Large and Complex AI Models

Cascade models up to 16 chips to support larger models (i.e., ResNet-101, ResNet-152, etc.).



PCIe Accelerator Card

- PCIe Gen2 x16 as 4 x 4 bifurcated lanes
- 16 Lightspeer® 2803S

AI Inference Server Reference Design

- Dual Intel Xeon Processor
- 4 GAINBOARD 2803S PCIe Card (64 chips)
- 1084 TOPS
- 9.9 TOPS/Watt

AI Model Processing Capabilities

- Model Cascade Processing
- Model Parallel Processing

Supported Frameworks

- TensorFlow
- Caffe

Targeted Inference Applications

- Image Recognition
- Object Detection & Tracking
- Natural Language Processing
- Natural Language Understanding
- Business Intelligence
- Visual Analysis

Note: Workloads used in performance tests may have been optimized for performance. Performance tests are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary.