

REPRESENTATIVE USE CASE

High-Performance Computing and GPU Acceleration for Mission Workloads

Design and deployment of a high-performance computing cluster with GPU acceleration to support computationally intensive mission applications including AI training, scientific simulation, and signals processing.

PRACTICE

Digital Foundations

CLIENT PROFILE

Research laboratory or intelligence organization requiring accelerated compute for simulation, modeling, and AI/ML training workloads

Challenge

- Existing compute infrastructure cannot keep pace with growing AI/ML and simulation workloads
- Long queue times and resource contention on shared HPC resources
- Need for GPU-accelerated compute within a secure, accredited environment

Approach

- Profile workload characteristics to determine optimal GPU/CPU ratios and interconnect requirements
- Design dense compute clusters with high-bandwidth, low-latency fabric (InfiniBand/RoCE)
- Implement job scheduling and resource management for multi-tenant usage
- Integrate with existing storage systems and establish data pipeline automation

Typical Outcomes

- Dramatically reduced model training and simulation run times
- Improved researcher productivity through reduced queue wait times
- Scalable platform ready for next-generation AI/ML workloads

Procurement Paths

- NASA SEWP V and SEWP VI for HPC hardware and GPU accelerators
- ITES-4H for Army compute infrastructure
- GSA MAS for integration and professional services

Compare vehicles and ordering steps on **How to buy**.

Partner Technology Examples

- NVIDIA
- AMD
- Hewlett Packard Enterprise
- Dell Technologies

FEATURED PRODUCT

Odin's Edge: For When the Cloud Isn't an Option

Odin's Edge is Norseman's autonomous edge AI platform. It delivers AI inferencing, rugged storage, private 5G, containerized desktops, and ServiceNow ITSM in a single field deployable unit. It is built for DDIL (disconnected, degraded, intermittent, and limited bandwidth) environments.

Frequently Asked Questions

What GPU platforms does Norseman deploy?

Norseman is a partner with NVIDIA and AMD for GPU compute platforms. NVIDIA H100, A100, and L40S configurations are available for AI training workloads. HPE, Dell, and Supermicro server platforms are used for HPC node deployment.

How is HPC infrastructure procured through government contracts?

HPC and GPU hardware is available through NASA SEWP V and VI and ITES-4H. Both vehicles support large hardware procurements. Norseman can deliver complete cluster solutions including racking, networking, cooling, and software stack configuration.

This is a representative use case. To request a briefing or a quote, contact Norseman Defense Technologies at 410.579.8600 or sales@norseman.com.