



Why Axiom?

Scale AI infrastructure without letting OEM cost structures limit growth

For AI / HPC Environments

- High-performance connectivity for GPU clusters and AI fabrics (400G, 800G, and beyond)
- Consistent, low-latency performance across large-scale, multi-node environments
- Validated compatibility across leading GPU and switching platforms

Low-Risk Deployment

- Validated in Axiom U.S. labs. Proven in production environments
- Proven compatibility across major OEM platforms
- Production-ready, not lab-only tested
- Backed by warranty + technical support

Where You Win

- AI cluster buildouts where cost per port directly impacts scale
- High-bandwidth fabric expansion (leaf-spine, GPU interconnects)
- Storage and memory scaling for training and inference workloads
- Avoiding OEM pricing constraints in rapidly growing environments

Expand the Deal

- Standardize Axiom across AI fabrics and data center infrastructure
- Bundle high-speed optics + cabling + memory
- Support ongoing upgrades as bandwidth requirements increase

Core Benefits

We help you scale AI infrastructure faster by removing cost barriers without compromising performance.

Scale Your AI Infrastructure Without Limits

Call: (888) 658-3326 or email: sales@axiomconnectivity.com

www.axiomupgrades.com

©2026 AXIOM® All rights reserved. Reproduction in whole or in part without permission is prohibited. AXIOM® reserves the right to modify the design and specifications of its products without notice. All other brand names, product names or trademarks belong to their respective holders. These trademarks shall not be used without the consent of their respective holders.

