

Vertical AI Jumpstart Services on Apple Silicon

Desktop AI Supercomputers

Powered by Apple Silicon | MLX-Native | Open-source Models and Scaffolding

AI Adoption Shouldn't Require a Data Center

Organizations across every industry are eager to leverage AI—but the traditional path demands multi-million-dollar infrastructure investments, specialized cooling, high-density racks, complex networking, and months of lead time. CNXN Helix's **AI on Apple Silicon** eliminates these barriers to entry.

Built by Apple® for Apple Silicon, these MLX-native desktop AI supercomputers deliver powerful AI performance—right on your desk—with zero-copy unified memory and open-source models. Combined with CNXN Helix Vertical Jumpstart Services, organizations can achieve their first AI win in days, not months. This gives you a practical way to evaluate AI using your own data before making broader infrastructure or investment decisions.

How the Apple Silicon Platform Removes Barriers to Getting Started with AI

This architecture is designed to address the most common challenges organizations face when moving from AI exploration to real-world implementation.

Eliminates Data Center Complexity

- No specialized racks, cooling, or power upgrades
- Standard wall outlet power—runs in a conference room
- Deploy in hours/days, not weeks/months
- Eliminates high up-front electrical and facility costs

Enterprise-grade AI on Your Desktop

- MLX-native, in-process AI inference
- Zero-copy unified memory architecture
- Open-source vector databases, embedding frameworks, models, and native guardrails
- Scale from 1-node to 2-node to 4-node clusters

MLX-native

Apple In-process AI

Up to 1TB

Unified Memory

Open Source

Models and Scaffolding

2-day

First Functional
AI Deployment

Vertical Use Cases Across Industries

Select from a menu of pre-built, deployable AI use cases tailored to your industry—each jumpstarting your business with a fully functional implementation with your own data, on-premises and ready to share with leadership and employees.

Education

- IEP and 504 Plan Support Assistant
- Multilingual Translation
- Curriculum AI Assistants
- Student Orientation and Aid

State and Local Government

- Public Services Assistant
- Municipal Town Permitting
- State Fish and Game Assistant

Healthcare

- Regulatory Compliance Assistant
- Readmittance Assistant
- Code-related Denial Assistant
- Patient Acuity Assistant

Retail

- In-store Associate Assistant
- Employee Onboarding
- Marketing Insights Assistant

Manufacturing

- Factory Production Assistant
- Maintenance Manuals Assistant
- R&D/Product Management Assistant
- ICS Assistant

General Departments

- IT Help Desk Ticket Assistant
- OSHA Compliance Assistant
- Insurance Claim Adjudication
- HR Assistant

How to Get Started

OPTION 1	OPTION 2	OPTION 3
Buy Hardware Only	OEM Try and Buy Program	Jumpstart Package
For highly technical AI teams and data scientists. Purchase Apple Silicon hardware in 1-node, 2-node, or 4-node cluster configurations and build your own custom AI solutions atop Apple Silicon with MLX, open-source models, and your preferred scaffolding.	A single Apple Silicon node paired with CNXN Helix Vertical Jumpstart Services. Ideal for deploying your first AI use case on-premises with expert guidance—prove value quickly with a focused, outcome-driven engagement on a single-node configuration.	Full-scale MLX clustering with a four-node Apple Silicon configuration delivering up to 1TB of unified memory and AI-accelerated compute. Purpose-built for very large models, high-precision inference, and massive data workloads—paired with CNXN Helix Jumpstart Services for rapid, expert-led deployment.



What's Included in the Jumpstart Package

Hardware

- Apple Silicon node(s) (Mac mini® or Mac Studio®)
- Monitor, keyboard, and mouse
- Thunderbolt/Ethernet cluster cables (multi-node configs)
- NVMe storage per node
- USB-C to USB-A adapter

CNXN Helix Jumpstart Services (~20 hrs)

- Client selects from list of use cases
- Data pipeline preparation and ingestion
- Day 1: Device setup and data staging
- Day 2: Pipeline, model config and training
- Project management and delivery

Start Smart, Scale Confidently

Your Path from Desktop to Apple Silicon-based Cluster

Apple Silicon's unified memory architecture and MLX framework give your organization a seamless path from desktop prototyping to production-scale AI. Start with a single Apple Mac mini or Mac Studio node to build, test, and validate your AI use cases with real data. When you're ready to scale, expand to multi-node Apple Silicon clusters—without rearchitecting your models or pipelines.

Because MLX is native to Apple Silicon, code that runs on a single node runs identically across a cluster. You pilot with a small audience on one node, prove value to leadership, and scale to 2-node or 4-node configurations as demand grows. With up to 1TB of unified memory in a 4-node cluster, you can run very large open-source models with high precision—all without the cost, complexity, or lead time of traditional data center infrastructure.

This approach delivers powerful, affordable, and scalable AI outcomes. Your initial deployments are not isolated experiments—they're the foundation that scales into broader AI initiatives over time. Apple Silicon clusters give organizations an accessible on ramp to enterprise AI that grows with your business.

Single Node (Mac mini / Mac Studio) → 2-node Cluster → 4-node Cluster

What works on a single Apple Silicon node works across the entire cluster—start focused, deliver value, expand strategically.

For more information about the comprehensive services Connection and our dedicated AI specialists at CNXN Helix offer, contact your Account Team today.

Business Solutions
1.800.800.0014

Enterprise Solutions
1.800.369.1047

Public Sector Solutions
1.800.800.0019

www.cnxnhelix.com