

# AI Datacenter Readiness Guide

## Evaluating Your Infrastructure for AI & ML Workloads

### Why AI Readiness Matters

AI workloads are fundamentally different from traditional enterprise compute. They require massive GPU clusters, high-bandwidth low-latency networking, specialized storage, and significant power and cooling capacity. The AI datacenter market is projected to reach **\$934B by 2030 at 31.6% CAGR**. Organizations that fail to prepare their infrastructure risk falling behind competitors and missing the AI transformation.

### 5 Dimensions of AI Readiness

#### 1. Compute

GPU servers (NVIDIA H100/H200, AMD MI300X), HPC clusters, CPU/GPU ratio planning, inference vs training requirements

#### 2. Storage

High-throughput parallel file systems, NVMe-oF, object storage for datasets, data pipeline architecture, minimum 100GB/s aggregate throughput for training

#### 3. Networking

400GbE spine-leaf fabric, RDMA/RoCE v2, InfiniBand for GPU interconnect, low-latency east-west traffic, dedicated AI network fabric

#### 4. Power & Cooling

40–60 kW per rack (vs 5–10 kW traditional), liquid cooling requirements, rear-door heat exchangers, power distribution planning

#### 5. Software & Orchestration

Kubernetes for AI workloads, container orchestration, model serving infrastructure, MLOps pipelines

# Readiness Assessment Framework

## Scoring Matrix

Rate each dimension on a scale of 1 (not ready) to 5 (fully capable). Identify the gap between your current state and target state to prioritize investments.

Dimension	Score (1-5)	Current State	Target State	Gap
Compute				
Storage				
Networking				
Power & Cooling				
Software & Orch.				

## Current State Assessment Questions

### Compute:

What GPU resources do you currently have? What is your CPU-to-GPU ratio?

### Storage:

What is your current aggregate storage throughput? Do you have parallel file systems?

### Networking:

What is your current fabric speed? Do you support RDMA?

### Power:

What is your current per-rack power density? Do you have liquid cooling capability?

### Software:

Are you running Kubernetes? Do you have MLOps pipelines in place?

## Phased Approach

Start with inference workloads, which have lower compute and power requirements, to build operational experience and validate your infrastructure. Then progressively scale to training workloads as your team and infrastructure mature. Build incrementally — don't try to build for full-scale training on day one.

Phase 1	Inference Workloads	Lower GPU density, standard cooling, existing network
Phase 2	Fine-Tuning	Moderate GPU clusters, enhanced cooling, upgraded fabric
Phase 3	Full Training	Dense GPU clusters, liquid cooling, InfiniBand interconnect

## Next Steps

BlackHawk Data helps organizations assess and build  
AI-ready datacenter infrastructure.

### Our Services

- AI Infrastructure Assessment
- GPU Cluster Design & Deployment
- High-Performance Networking
- Power & Cooling Consulting
- MLOps & Orchestration

### Schedule a Datacenter Readiness Assessment

(877) 383-1845

[info@blackhawk11.com](mailto:info@blackhawk11.com)

[blackhawkdata.com](http://blackhawkdata.com)