

## Purpose and Overview

WSU Tech is soliciting quotations from qualified vendors for the purchase, delivery, installation, and support of an **enterprise on-premises data center refresh** consisting of compute, GPU acceleration, storage, high-speed networking, and related professional services.

The goal of this RFQ is to acquire a **scalable, high-performance three-tier architecture** designed to support:

- Core virtualization and general compute workloads in Hyper-V
- Light-duty GPU workloads (e.g., Adobe Creative Cloud applications delivered via RDS)
- Optional medium-scale AI / data science workloads
- High availability, low latency, and long-term operational support

Functionally equivalent solutions are acceptable.

## 1. Target Architecture (Informational)

The desired solution consists of:

- **Compute Tier:**
- 4 general compute/ virtualization nodes (required) Dell PowerEdge preferred
- 2 GPU-enabled nodes for light-duty workloads (required)
- 2 additional GPU-enabled nodes for medium AI workloads (optional)

- **Storage Tier:**
- Enterprise all-flash NVMe storage providing a minimum of 200 TB effective capacity
- **Network Tier:**
- Redundant Top-of-Rack (ToR) Ethernet fabric - Cisco Nexus Preferred
- 100 Gbps connectivity throughout (QSFP28)
- **Services & Support:**
- Vendor-provided deployment, documentation, and enterprise-grade support

## **2. Compute Server Requirements**

### **2.1 General Compute / Virtualization Nodes (Qty: 4 - Required)**

Each server must meet or exceed the following minimum specifications:

#### **Form Factor**

- Rack-mount server

#### **Processor**

- Dual-socket x86\_64 architecture
- Minimum 24 cores per processor
- Suitable for dense virtualization workloads

#### **Memory**

- Minimum **768 GB RAM per node**
- Must be expandable

## **Boot / Local Storage**

- Redundant enterprise-class NVMe boot devices

## **Networking**

- Minimum **2 x 100 Gbps Ethernet ports (QSFP28)** per node
- Support for VLAN segmentation and virtualization

## **Power**

- Dual hot-plug redundant power supplies
- High-efficiency (Platinum or better preferred)

## **2.2 GPU Nodes - Light-Duty Workloads (Qty: 2 - Required)**

Designed for workloads such as remote graphics, media creation, and RDS-delivered applications (e.g., Adobe suite).

Must meet all general compute requirements plus:

### **GPU Capabilities**

- Support for modern data-center GPUs
- GPUs suitable for graphics acceleration and light compute workloads
- Adequate GPU memory for media and creative applications

### **Use Cases**

- Remote Desktop Services (RDS)
- Creative and instructional applications
- Inference-capable, non-AI-training focused workloads

## 2.3 GPU Nodes - Medium AI / Compute Workloads (Qty: 2 - Optional)

Vendors shall provide **optional pricing** for additional GPU-accelerated servers intended for medium-scale AI or analytics workloads.

Must meet all general compute requirements plus:

### GPU Capabilities

- GPUs suitable for local AI/ ML workloads
- Significantly higher GPU memory and compute capability than light-duty GPU nodes
- Support for common AI frameworks (vendor to specify)

Optional nodes must integrate seamlessly into the proposed environment.

## 3. Storage System Requirements

Vendors shall propose an **enterprise all-flash NVMe storage system** meeting the following minimum requirements:

### Architecture

- Dual-controller, active/active design
- NVMe/TCP or equivalent supported

### Capacity

- Minimum **200 TB effective usable capacity**, post data reduction
- Vendor must disclose data reduction assumptions

### Performance & Availability

- Designed for virtualization and GPU workloads
- Non-disruptive expansion
- Tolerance for multiple drive failure without data loss

## Connectivity

- **100 Gbps Ethernet connectivity (QSFP28)**
- Sufficient ports to support redundant fabric design

## Security

- Encryption at rest required
- Secure drive retirement support

## Licensing

- All required software and features must be included
- No recurring capacity-based licensing penalties

# 4. Network and Switching Requirements

The vendor shall quote **redundant Top-of-Rack (ToR) Ethernet switching** to support all proposed compute and storage systems.

## Switching Requirements

- Redundant switch pair per rack
- Non-blocking architecture
- **100 Gbps QSFP28 ports throughout**
- Sufficient port density for:
  - All compute nodes
  - All storage controllers
- Inter-switch links

## **Fabric Requirements**

- Fully redundant data paths
- Support for:
- VLANs
- Jumbo frames
- Low-latency east-west traffic

Vendor must include:

- Switch models
- Optics
- Cables
- Any required licenses

## **5. Deployment and Professional Services**

Vendor shall provide professional services covering:

### **Pre-Deployment**

- Single point of contact
- Site readiness review
- Implementation planning

## **Deployment**

- Rack and stack
- Hardware installation
- Firmware updates
- Baseline configuration
- Network and storage connectivity

## **Post-Deployment**

- Validation testing
- Documentation
- Knowledge transfer to WSU Tech staff

Services may be a combination of on-site and remote, but on-site presence is required for core deployment activities.

## 6. Support and Maintenance Requirements

All hardware must include enterprise-class support with:

- **24x7x365 technical support**
- **4-hour on-site response** for severity-one incidents
- Coverage for hardware and system software

Vendors shall quote:

- 3-year support term
- 5-year support term

Support must include:

- Proactive monitoring
- Automated case creation
- Access to firmware and software updates

## 7. Vendor Qualifications

Vendors must:

- Be authorized resellers for all proposed equipment
- Have experience supporting higher-education clients
- Provide at least **three (3) higher-education references**
- Demonstrate experience with:
  - Virtualization platforms
  - GPU-accelerated workloads
  - Enterprise NVMe storage

Vendors shall disclose participation in cooperative contracts (e.g., NASPO, TIPS).

## 8. Pricing and Response Format

### Pricing

Pricing must be itemized and clearly identify:

- Required vs optional components
- One-time vs recurring costs

At minimum, pricing should include:

- Compute nodes (by type)
- GPU components
- Storage system
- Networking and optics
- Professional services

- Support and maintenance

## **Required Response Attachments**

- Technical compliance matrix
- Exceptions list
- Assumptions and dependencies
- Estimated delivery lead times

## **9. Evaluation Criteria**

Responses will be evaluated based on:

- Technical compliance
- Support model and service levels
- Vendor experience and references
- Implementation approach
- Total cost of ownership

## **10. Public Procurement Considerations**

- This RFQ is governed by State of Kansas and WSU Tech procurement policies
- Pricing must remain valid through the evaluation period
- WSU Tech reserves the right to accept or reject any or all responses