

# Proposal: AI Server Infrastructure for District-Level Deployment & Development

**To:** Rod Morrison and Board Members  
**From:** Everett Cook, Technology Director  
**Subject:** Proposal for Building a Cost-Effective, Scalable In-House AI Server

## Purpose

This proposal outlines a high-performance, scalable AI server solution for district-wide use. The primary goal is to reduce reliance on cloud platforms, increase data control, and support advanced AI-driven systems for internal operations, IT processes, and administrative workflows—all within a one-time hardware investment.

## Why Build Instead of Buy?

Rather than rely on recurring cloud fees or limited third-party AI tools, this build allows the district to:

- **Host in-house large language models (LLMs)** for administrative support, scheduling automation, documentation, and more
- **Fully control privacy and security**
- **Customize workflows and tools to meet specific district needs**
- **Eliminate dependency on commercial AI platforms**
- **Support long-term growth without further infrastructure investments**

## Project Summary

Component	Purpose	Why This Choice
Supermicro X13SWA-TF Motherboard	System backbone	Server-grade board with full ECC support, extensive PCIe slots for expansion, and compatibility with Intel Xeon CPUs
Rosewill 4U Server Chassis	Physical housing	Rackmount case with proper airflow, high-capacity drive bays, and E-ATX compatibility
Corsair HX1500i PSU	Power delivery	1500W Platinum-rated power supply with modular connections and PCIe 5.0 compliance for future-proofing
Samsung 970 EVO Plus 1TB (OS)	Boot/system drive	Reliable and fast Gen 3 SSD, ideal for Linux OS and core AI service infrastructure
Samsung 990 PRO 4TB (Model/Data Storage)	LLM model & dataset storage	Industry-leading Gen 4 NVMe speeds, optimized for fast read/write of large models

<b>WD_BLACK SN850X 2TB (Cache Drive)</b>	Swap/temp/cache	Built for high write endurance—perfect for temporary AI workloads and inference caches
<b>Seagate IronWolf 8TB (Backup)</b>	Configuration/data snapshots	NAS-grade backup drive with high durability and RAID-ready design
<b>Samsung 870 QVO 8TB x4 (RAID)</b>	Redundant primary storage	SSD-based RAID 10 array for fast, redundant, high-capacity data handling
<b>LSI MegaRAID 9460-16i</b>	RAID management	Professional 16-port controller with SAS/SATA support, perfect for enterprise RAID
<b>Samsung DDR5 ECC RAM (4 x 64GB)</b>	System memory	ECC-registered DDR5 memory with 256GB initial capacity, scalable to 2TB
<b>Intel Xeon Gold 6454S (32-core/64-thread)</b>	CPU	Enterprise-grade CPU with high thread count for parallel processing and multitasking
<b>PNY RTX 6000 Ada 48GB GPU</b>	AI training/inference	Top-tier professional GPU optimized for LLM inference and model development
<b>Supermicro 4U Active CPU Cooler</b>	Thermal solution	Purpose-built cooling for Xeon processors in 4U chassis environments
<b>SAS-to-SATA Cables (x2)</b>	RAID connection	Supports all 4 RAID drives via two high-speed cable assemblies
<b>Sabrent M.2 Heatsinks (x2)</b>	NVMe cooling	Passive cooling for sustained performance on primary M.2 drives
<b>Noctua NF-F12 iPPC-2000 Fans (x2)</b>	Supplemental airflow	Additional thermal support for GPU-heavy workloads

**Total Estimated Cost: not to exceed \$19,000**

(Scalable down by reducing GPU count, memory, or RAID configuration if needed.)

## Operational Benefits

- **Technology Department**
  - Automate internal IT support tools, data cleanup, report generation, and device management tasks
  - Develop custom scripts and tools for high-efficiency maintenance
- **Administrative Services**
  - Host AI-based assistants to aid with documentation, scheduling, research, and transcription

- Process public records, policies, or board meeting minutes with summarization and analysis tools
- **Districtwide Efficiency**
  - Provide a single AI backend capable of scaling with future projects
  - Ensure data never leaves our local environment unless explicitly permitted
- **Future Expansion**
  - Add a second GPU for enhanced model performance
  - Expand memory to accommodate larger model requirements
  - Transition high-demand services like website support, alert management, or communications automation to in-house AI control

## Why This Build Is Ideal

- **Custom-built** to match real-world district workloads
- **Flexible and upgradeable**—can scale with future software and AI demands
- **Cost-effective** compared to enterprise vendor solutions or cloud AI subscriptions
- **Secure and private**, with no reliance on outside platforms or data sharing
- **Sustainable**—built to last years without major replacement

## Conclusion

This system provides an all-in-one AI backbone for the Southeast Island School District's internal operations. It allows controlled, private experimentation with modern AI tools while building capacity for future innovation. The build is flexible enough to be scaled down for initial cost savings or scaled up to meet more demanding needs as we grow.

I'm happy to provide further technical details or host a walkthrough meeting.

Sincerely,

**Everett Cook**

Technology Director

Southeast Island School District