

**STATEMENT OF WORK
NVIDIA SuperPod and VAST Storage FOR
RESEARCH, APPLIED ANALYTICS, AND STATISTICS (RAAS)**

1. BACKGROUND

The Compliance Data Warehouse (CDW) is part of the RAAS computing environment and is an integral component of the Internal Revenue Service (IRS) strategy for data and analytics. Its success as an off-line environment for experimentation and testing, its core data and analytical computing capabilities, its project management expertise, and its reputation as an excellent partner in research and analytics have made it the destination of choice for the IRS' leading researchers and data scientists. Early adoption of new insights and discovery through CDW supports IRS strategic interests as well as the Office of Management and Budget's (OMB's) vision for research in support of evidence-based policy making in government.

CDW is specifically designed to meet the distinctive use patterns of research and advanced analytics. It combines disparate, multi-structured, and distributed data sets over long periods of time for longitudinal analysis, simulation, and specialized modeling. This capability allows users to address new questions and build on knowledge gained from prior research. It provides a flexible, interactive, and responsive environment that lets users focus on research problems and analytical workflows unconstrained by rigid, lifecycle-based workflows. It invests in tools and technologies that are tailored to off-line experimentation and discovery and that align to the fast-moving areas of data science, advanced analytics, and visualization.

Investments in CDW's data and technology capabilities have been significant drivers of success in risk discovery, behavioral insights, and workload planning. Research-funded initiatives leveraging CDW have resulted in the creation of improved identity (ID) theft models, refund fraud filters, tax gap estimates, burden measures, case selection algorithms, and a deeper understanding of the mechanisms that drive taxpayer behavior. Through partnerships with academics and industry, CDW has also accelerated innovation in other areas of data analytics, including the use of natural language processing for feature extraction of third-party documents in massively large collections; graph models to analyze complex relationships in multi-party transactions; synthetic data generation of tax filing populations to enhance data sharing with public researchers; recommender systems to identify latent compliance issues for examination; community detection modeling of networks of preparers, Electronic Filing Identification Numbers (EFINs), and taxpayers; and agent-based models for the simulation of service channel migration.

This acquisition involves the purchase of a Nvidia SuperPod and VAST Storage. The Internal Revenue Service (IRS) is committed to leveraging advanced technologies, including deep learning and natural language processing, to enhance its operations and provide better services to taxpayers. To support this effort, the IRS requires a robust and scalable infrastructure that can handle complex machine learning (ML) workloads. The Nvidia Super Pod is a critical component of this infrastructure, providing the necessary compute power, storage, and networking capabilities to support the development and deployment of large-scale ML models. Further, we require appropriate storage to handle data and results.

2. SCOPE AND OBJECTIVES

The goal of this effort is to purchase a Nvidia SuperPod and VAST Storage as described in Section 3. RAAS manages and administers its own equipment in the RAS-1 General Support System (GSS) boundary and will integrate this equipment at the IRS Enterprise Computing Center – Martinsburg (ECC-MTB).

3. DESCRIPTION OF TASKS AND DELIVERABLES

Upon award, the vendor shall provide delivery, integration, and installation services for equipment described below. In addition, the vendor shall update its customer list for the renewal period of performance for customer service, technical support, and other services.

CLIN No.	OEM Mftr	Part #	Description	Qty	Unit of Issue
0001	Nvidia		<p>Nvidia SuperPod Bundle consisting of the following:</p> <ul style="list-style-type: none"> (31) Nvidia DGXB-G1440+P2CMI60 DGX B200 P4387 System 1440GB Full, 5 Yr (31) Comprehensive Media retention service, 5 yr (248) Nvidia 731-AI7003+P2CMI60 NVIDIA AI Enterprise Essentials Subscription per GPU , 5 Year Std Support (248) RUN.AI 744-SW7001+P2MI60 Advanced Cluster Management SW Subscription per GPU, 5 Year Std Support (includes 2 UFM HCA licenses per GPU) (includes Run:AI (1) VASTDATA VAST-8XSTORAGE-ENDPOINT-60M VAST-8XSTORAGE-ENDPOINT-60M (Usable PB 2.39) 5 yr (1) VASTDATA VAST-CERT-TRAIN VAST training for up to 6 users (16) Nvidia 980-9I111-00H005 NVIDIA AOC splitter, 200(2x100)Gbps to 2x100Gbps, OSFP to 2xQSFP56, 5m, fin to flat (50) Nvidia Cat6 Cat6 Snagless Unshielded (UTP) Network Patch Ethernet Cable-Blue, 10 m (from utility rack to GPU racks) (1) Nvidia NFS-storage Customer provided NFS storage for Home Directories (2) Nvidia 920-9B210-00FN-0M0 NVIDIA Quantum 2 based NDR InfiniBand Switch, 64 NDR ports, 32 OSFP ports, 2 Power Supplies (AC), Standard depth, Managed, P2C airflow, Rail Kit (4) Nvidia 980-9I73V-000005 NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 5m (4) Nvidia 980-9I693-00NS00 NVIDIA single port transceiver, 400Gbps,NDR, QSFP112, MPO12 APC, 850nm MMF, up to 30m, flat to (8) Nvidia 920-9B210-00FN-0M0 NVIDIA Quantum 2 based NDR InfiniBand Switch, 64 NDR ports, 32 OSFP ports, 2 Power Supplies (AC), Standard depth, Managed, P2C airflow, Rail Kit 	1	Each

			<p>(4) 920-9B210-00FN-0M0 NVIDIA Quantum 2 based NDR InfiniBand Switch, 64 NDR ports, 32 OSFP ports, 2 Power Supplies (AC), Standard depth, Managed, P2C airflow, Rail Kit</p> <p>(124) 980-9I51A-00NS00 NVIDIA twin port transceiver, 800Gbps,2xNDR, OSFP, 2xMPO12 APC, 850nm MMF, up to 50m, flat top</p> <p>(384) 980-9I510-00NS00 NVIDIA twin port transceiver, 800Gbps,2xNDR, OSFP, 2xMPO12 APC, 850nm MMF, up to 50m, finned</p> <p>(120) 980-9I57X-00N010 NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 10m</p> <p>(128) 980-9I57Z-000020 NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 20m</p> <p>(256) 980-9I73U-000003 NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 3m</p> <p>(4) 980-9I51S-00NS00 NVIDIA single port transceiver, 400Gbps,NDR, OSFP, MPO12 APC, 850nm MMF, up to 50m, flat top</p> <p>(2) 980-9I572-00N005 NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 5m</p> <p>(2) 920-9B020-00RA-0D0 NVIDIA UFM Appliance 3.0 for UFM Telemetry or UFM Enterprise, 1U server with 2x ConnectX-7 NDR Single-port 400Gb/s InfiniBand adapter, Secured-boot, UFM software package sold separately</p> <p>(2) 920-9B210-00FN-0M0 NVIDIA Quantum 2 based NDR InfiniBand Switch, 64 NDR ports, 32 OSFP ports, 2 Power Supplies (AC), Standard depth, Managed, P2C airflow, Rail Kit</p> <p>(2) 920-9B210-00FN-0M0 NVIDIA Quantum 2 based NDR InfiniBand Switch, 64 NDR ports, 32 OSFP ports, 2 Power Supplies (AC), Standard depth, Managed, P2C airflow, Rail Ki</p> <p>(62) 980-9I693-00NS00 NVIDIA single port transceiver, 400Gbps,NDR, QSFP112, MPO12 APC, 850nm MMF, up to 30m, flat top</p> <p>(124) 980-9I510-00NS00 NVIDIA twin port transceiver, 800Gbps,2xNDR, OSFP, 2xMPO12 APC, 850nm MMF, up to 50m, finned</p> <p>(32) 980-9I57X-00N010 NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 10m</p> <p>(30) 980-9I57Z-000020 NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 20m</p> <p>(88) 980-9I73V-000005 NVIDIA passive fiber cable, MMF, MPO12 APC to MPO12 APC, 5m</p> <p>(4) 980-9I51S-00NS00 NVIDIA single port transceiver, 400Gbps,NDR, OSFP, MPO12 APC, 850nm MMF, up to 50m, flat top</p> <p>(2) 980-9I572-00N005 NVIDIA passive fiber cable, MMF, MPO12 APC to 2xMPO12 APC, 5m</p> <p>(2) 920-9B020-00RA-0D0 NVIDIA UFM Appliance 3.0 for UFM Telemetry or UFM Enterprise, 1U server with 2x ConnectX-7 NDR Single-port 400Gb/s InfiniBand adapter, Secured-boot, UFM software package sold separately</p> <p>(2) 920-9N42F-00RI-7C0 NVIDIA Spectrum-4 based 800GbE 2U Open Ethernet switch with Cumulus Linu x</p>		
--	--	--	---	--	--

			<p>Authentication, 64 OSFP ports and 1 SFP28 port, 2 power supplies (3KW, AC), x86 CPU, Secure-boot, standard depth, C2P airflow, Tool-less Rail Kit</p> <p>(2) 920-9N42F-00RI-7C0 NVIDIA Spectrum-4 based 800GbE 2U Open Ethernet switch with Cumulus Linux x Authentication, 64 OSFP ports and 1 SFP28 port, 2 power supplies (3KW, AC), x86 CPU, Secure-boot, standard depth, C2P airflow, Tool-less Rail Kit</p> <p>(2) 920-9N110-00R1-0C0 Nvidia Spectrum based 1GBase-T/100GbE 1U Open Ethernet switch with Cumulus Linux, 48 RJ45 ports and 4 QSFP28 ports, Dual Power Supply (AC), x86 CPU, short depth, C2P airflow, 4-post Rail-kit.</p> <p>(19) 980-9I510-F4NS00 TWIN PORT TRANSCEIVER, 800GBPS,2X400GbE, OSFP, 2XMPO12 APC, 850NM MMF, UP TO 50M, FINNED</p> <p>(4) 980-9I042-00C000 Mellanox transceiver, 100GbE, QSFP28, LC-LC, 1310nm, DR1</p> <p>(2) 980-9I510-F4NS00 TWIN PORT TRANSCEIVER, 800GBPS,2X400GbE, OSFP, 2XMPO12 APC, 850NM MMF, UP TO 50M, FINNED</p> <p>(38) 980-9I570-00N030 NVIDIA passive fiber cable, MMF , MPO12 APC to MPO12 APC, 30m</p> <p>(8) 980-9IA0Q-00N01A NVIDIA passive Copper cable, IB twin port NDR, up to 800Gb/s, OSFP, 1.5m</p> <p>(14) 980-9I51S-F4NS00 SINGLE PORT TRANSCEIVER, 400GBPS,400GbE, OSFP, MPO12 APC, 850NM MMF, UP TO 50M, FLAT TOP</p> <p>(4) Nvidia Cat6 Snagless Unshielded (UTP) Network Patch Ethernet Cable-Blue, 10 m (from utility rack to GPU racks)</p> <p>(4) 980-9I30H-F4NM00 TWIN PORT TRANSCEIVER, 800GBPS,2X400GbE, OSFP, 2XMPO12 APC, 1310NM SMF, UP TO 100M, FINNED</p> <p>(2) 980-9I693-F4NS00 SINGLE PORT TRANSCEIVER, 400GBPS,400GbE, QSFP112, MPO12 APC, 850NM MMF, UP TO 50M, FLAT TOP</p> <p>(2) 980-9I30H-F4NM00 TWIN PORT TRANSCEIVER, 800GBPS,2X400GbE, OSFP, 2XMPO12 APC, 1310NM SMF, UP TO 100M, FINNED</p> <p>(2) Off-the-shelf, POT EFALU-PA2S1Q-005M or similar DR1 Splitter cable 1x 400Gb/s to 4x 100Gb/s</p> <p>(62) 980-9I693-00NS00 NVIDIA single port transceiver, 400Gbps,NDR, QSFP112, MPO12 APC, 850nm MMF, up to 30m, flat top</p> <p>(5) 920-9N110-00R1-0C0 Nvidia Spectrum based 1GBase-T/100GbE 1U Open Ethernet switch with Cumulus Linux, 48 RJ45 ports and 4 QSFP28 ports, Dual Power Supply (AC), x86 CPU, short depth, C2P airflow, 4-post Rail-kit.</p> <p>(219) Nvidia Cat6 Snagless Unshielded (UTP) Network Patch Ethernet Cable-Blue, 10 m (from utility rack to GPU racks)</p> <p>(3) 980-9I30H-F4NM00 TWIN PORT TRANSCEIVER, 800GBPS,2X400GbE, OSFP, 2XMPO12 APC, 1310NM SMF, UP TO 100M, FINNED</p>		
--	--	--	---	--	--

			<p>(5) Off-the-shelf, POT EFALU-PA2S1Q-005M or similar DR1 Splitter cable 1x 400Gb/s to 4x 100Gb/s</p> <p>(10) 980-9I042-00C000 Mellanox transceiver, 100GbE, QSFP28, LC-LC, 1310nm, DR1</p> <p>(2) Nvidia Cat6 Snagless Unshielded (UTP) Network Patch Ethernet Cable-Blue, 10 m (from utility rack to GPU racks)</p> <p>(7) Nvidia INTEL-MGMT-SERVER Management Server-DGX BC + Run:ai</p> <p>(48) Raritan (PX3-5091R-P1Q2R1A5) or Equiv Horizontal PDU for Compute Rack</p> <p>(10) Raritan (PX3-5747V-V2) or equivalent Vertical PDU for Mgmt Rack</p> <p>(21) Legrand NVIDPD13 or equivalent Physical Racks</p> <p>(1) 723-TAMP00+P2CMI60 NVIDIA Premium Technical Account Manager, Remote, 5 Years</p> <p>(2) 788-D2M06L+P2CMI00 NVIDIA Infrastructure Startup Engineer, 6 Months. 8 hours per week, local business hours, minimum utilization unit - 4 hours.</p> <p>(18) 790-CN7NNZ+P2CMI60 NVIDIA ENT Business Critical Support Services NBD On-Site CE for QM9700,5 Years</p> <p>(4) 781-C56NNZ+P2CMI60 NVIDIA ENT Business Critical Support Services NBD On-Site CE for SN5600 Series,5 Years</p> <p>(7) 781-C23NNZ+P2CMI60 NVIDIA ENT Business Critical Support Services NBD On-Site CE for SN2201 Series,5 Years</p> <p>(118) 785-CAON0Z+P2CMI60 NVIDIA BUSINESS CRITICAL SERVICE FOR ACTIVE FIBER CABLES,5 Years</p> <p>(658) 785-CPFN0Z+P2CMI60 NVIDIA BUSINESS CRITICAL SERVICE FOR PASSIVE FIBER CABLES,5 Years</p> <p>(714) 785-CTR0Y+P2CMI60 NVIDIA BUSINESS CRITICAL SERVICE FOR TRANSCEIVERS,5 Years</p> <p>(4) 793-CU9NNZ+P2CMI60 NVIDIA ENT Business Critical Support Services NBD On-Site CE for MUA9602 H Appliance,5 Years</p> <p>(1) 788-D1X20B+P2CMI00 NVPS deployment for one 31-DGX-node scalable unit.Price per SU. (T&E included,Storage excluded)</p> <p>(1) 718-DG7081+P1CMI00 Mandatory Delivery and Installation by Certified NVIDIA SuperPOD Specialization Partner</p> <p>(1) 719-DLI010+P2CMI00 DLI (Deep Learning Institute) Pack of 100 Learning Credits</p> <p>(1) SUPERPOD-STRGE-CBLE NVIDIA to VAST Storage Cable bundle</p> <p>(1) GAI-IRS-ST-WARE Staging and warehouse storage, predelivery up to 100 pallets for up to 4 months</p>		
0002	VASTDATA	DF-3060	Ceres Enclosure: 1U HA VAST Enclosure (up to 50GB/s): incl. 1350TB NVMe Flash, 12.8TB SCM(TBC), Bluefield DPUs (ETH or IB)	4	Each
0003	VASTDATA	1U-1N-GEN5-2NIC-	Single Server 1U Chassis inc (Dual NIC, Dual Port	8	Each

		SEC	200Gb's QSFP112 or Dual Port HDR200 QSFP112) - Secure Boot Enabled		
0004	VASTDATA	CERES-AOC-CABLE-KIT	Fiber optic cables and SFPs for CERES D-Boxes and C-Boxes	1	Each
0005	VASTDATA	ETH-NVMEF-1X64-200G	1 * 64 Port 2U 200G NVMe Fabric Ethernet Switch	2	Each
0006	VASTDATA	SW-U-300-KYS-5YR-4.8PB	4.8PB Useable 5 year Gemini Services Term + Keep Your System Option	1	Each
0007	VASTDATA	INSTALL	VAST Data Installation Services	1	Each
0008	VASTDATA	COPILOT	VAST Data Co-Pilot Services	1	Each
0009	VASTDATA	VAST-CERT-TRAIN	VAST training for up to 6 users	1	Each
0010 Optional Item	Nvidia	965-00000-0117-000	Upgrade NVIDIA SuperPOD DGX's to 4TB memory-Memory Module, Samsung GDDR5, 128GB, 4800MHz	992	Each

4. Delivery/Period of Performance

Delivery of equipment 60 days ARO. One-year period of performance for services.

6. PLACE OF PERFORMANCE

Enterprise Computing Center - Martinsburg

5. SHIPPING INFORMATION

Items for this acquisition should be delivered to the following address:

Enterprise Computing Center-Martinsburg (ECC-MTB)
250 Murall Dr.
Kearneysville, WV 25430