

Accelerating the Mission

NeuralMesh™ eliminates infrastructure constraints and delivers the ultra-high throughput and microsecond latency required to maximize the efficiency of exascale HPC and AI environments.

Challenges

- **Massive Data Influx:** Handling exabytes of data.
- **Optimizing HPC investments:** Fully utilizing HPC clusters and maximizing impact per dollar
- **Urgency of Insight:** Analysis of complex datasets for rapid decision-making

Solution

NeuralMesh by WEKA is a purpose-built, service-oriented data platform engineered to eliminate these infrastructure constraints.

Benefits

- **Mission Acceleration:** Reduce data bottlenecks that slow time-critical operations
- **Infrastructure Efficiency:** Stretch limited IT budgets and maximize infrastructure
- **Exascale Performance:** Confidently run exabytes of mission-critical data with consistent, ultra-low latency

Eliminating Data Bottlenecks for Federal Agencies

Federal agencies face an unprecedented data surge from advanced simulations and sensor networks. This exabyte-scale data is crucial for national security and scientific leadership.

Legacy HPC storage and conventional enterprise file systems, even when augmented, were not designed for the mission-centric demands of today's federal environment. They introduce I/O bottlenecks, data staging complexities, and operational overhead that compromise mission effectiveness. Whether it's for ModSim, Monte Carlo, AI-assisted workloads, or computer vision, NeuralMesh is purpose-built to address these challenges head-on.

Engineered to solve this problem, NeuralMesh delivers the high throughput and low latency needed to accelerate data-intensive workloads and maximize high-performance computing (HPC) environments. By eliminating infrastructure constraints, WEKA empowers federal agencies to rapidly process and analyze data, ensuring timely, actionable intelligence for their most ambitious missions.

Bridging the Gap Between Traditional HPC, Modern HPC, and AI

The landscape of federal operations, from cutting-edge scientific research to defense and intelligence, is increasingly defined by data. Agencies are grappling with:

- **Explosive Data Growth:** Petabytes to exabytes of data generated from high-fidelity simulations, advanced sensor arrays, genomic sequencing, climate models, and intelligence gathering.
- **Demand for Real-Time Insight:** The critical need to process and analyze vast datasets with unprecedented speed to inform critical decisions, respond to evolving threats, and accelerate scientific breakthroughs.
- **Maximizing Costly Assets:** Ensuring optimal utilization of multi-million dollar HPC clusters, supercomputers, and GPUs to achieve maximum return on investment for taxpayers.
- **Complex, Distributed Environments:** Operating across secure on-premise data centers, classified networks, and authorized cloud environments, requiring seamless data mobility and consistent performance.
- **Stringent Security & Compliance:** Adhering to rigorous federal mandates for data security, integrity, and regulatory compliance.

Unlocking Exabyte-Scale Performance for Science and National Security

For federal agencies relying on high-frequency sensor data, real-time intelligence feeds, or iterative simulation outputs, data access latency directly impacts operational effectiveness and decision superiority. Delays caused by conventional object and file systems can jeopardize timely mission response, degrade operational awareness, and reduce the effectiveness of national defense and research systems.

NeuralMesh fundamentally solves this through:

- **S3 over RDMA:** Bypasses conventional TCP/IP stacks, delivering data with near-direct memory access for ultra-low latency and high throughput, crucial for real-time analytics and critical applications.

- **Single-Hop Writes + Parallel Processing:** Metadata and I/O operations execute in parallel across the entire cluster, eliminating serial bottlenecks and accelerating data ingest and retrieval.
- **POSIX-Native Fusion:** Seamlessly merges the performance and familiar semantics of POSIX file systems with the scalability and flexibility of S3 object APIs, providing a unified data access layer for diverse federal applications.
- **Unified, High-Speed Ingest:** Rapidly consolidates multi-source datasets from instruments, simulations, and intelligence feeds into a single, high-performance repository.
- **Parallelizable Access Patterns:** Optimized for distributed HPC jobs and complex analytical queries, ensuring all compute nodes can access data concurrently without contention.
- **Elimination of Staging/Copying Operations:** Reduces operational complexity and accelerates readiness by eliminating manual pre-processing and redundant data movement.

Maximizing Return on HPC Investments

Federal HPC infrastructure represents a substantial investment of taxpayer dollars. Yet, many clusters operate with compute resources, particularly specialized accelerators like GPUs, that are significantly underutilized—often below 40%. This artificial performance ceiling inflates the cost-per-computation and slows the pace of innovation and mission achievement.

NeuralMesh extends WEKA's proven impact in data-intensive environments by enabling federal agencies to achieve:

Dramatic Latency Reductions

Accelerating time-to-insight for real-time analytics, intelligence processing, and complex simulations.

Maximized Resource Utilization

Real-world deployments have shown NeuralMesh can increase effective utilization of specialized compute assets like GPUs from sub-50% to over 90%—enabling better outcomes with fewer resources.

Significant Compute Capacity Gains

Enabling agencies to achieve 2-3x throughput improvements using existing infrastructure, without additional procurement cycles or expanded physical footprint.

Seamless Cross-Site Collaboration

Facilitating distributed research, multi-agency operations, and secure data sharing across geographically dispersed teams.

Why Federal Agencies Choose NeuralMesh

- **Mission Acceleration:** Reduce data bottlenecks that slow time-critical operations and research—accelerating insights that support national defense, scientific breakthroughs, and emergency response.
- **Infrastructure Efficiency:** Stretch limited IT budgets by maximizing the effectiveness of taxpayer-funded infrastructure.
- **Exascale Performance:** Confidently run thousands of compute nodes and exabytes of mission-critical data with consistent, ultra-low latency response.
- **Designed for Mission Continuity:** Engineered to maintain performance and data availability under duress, with resilience, self-healing, and disaster recovery built-in.

