

NeuralMesh AIDP

Enterprise data is **always current, always permission-enforced, and always AI-ready** on WEKA's AI Data Platform

Challenges

- Data isn't AI-ready when it lands
- Pipelines break when files change
- Every AI project must be rebuilt from scratch
- Hard to operate manually
- File permissions get ignored
- Data pipelines must be managed on a project-by-project basis
- Projects use out-of-date datasets

Solution

AIDP watches NeuralMesh™ in real time, preparing data automatically as files land, so every AI project starts with current, indexed permission aware content

Benefits

- Data is always ready
- Semantic search out of the box
- Easily enforce POSIX permissions
- No pipelines to build or maintain
- Live queue monitoring, role-based governance, and status dashboards in one user interface

Inference that Works at Enterprise Scale

NeuralMesh AIDP™ is a software and hardware add-on to NeuralMesh™ deployed from the **WEKA App Store**. It transforms how enterprises operationalize AI factories using an event-driven pipeline that automatically handles ingestion, chunking, embedding, vectorization, and POSIX-enforced retrieval without custom engineering.

Rather than building and maintaining custom data preparation pipelines, AIDP automatically processes data during ingest, keeping it continuously AI-ready as files are created, updated, or deleted. AIDP makes data immediately available to any customer AI pipeline that needs it and delivers ready-to-consume AI outcomes such as semantic search without requiring organizations to manually manage AI frameworks.

Enterprises invest heavily in infrastructure and typically hit the same wall: getting data into AI models reliably, securely, and at scale. **AIDP removes that wall.**

Why WEKA's AIDP

AIDP is designed for enterprise AI factory deployments across healthcare, life sciences, financial services, manufacturing and other uses. AIDP extends NeuralMesh from a high-performance storage system into a complete solution that stores enterprise data at scale, continuously keeps it current, and delivers AI outcomes directly.

The Problem AIDP Solves

Most enterprise AI pilots fail to reach production because the data infrastructure underneath isn't built for continuous AI operations and is consuming engineering cycles and eroding board confidence. Before an AI project can use enterprise data, it must be parsed, chunked, embedded, and indexed. As data changes, the entire pipeline must stay in sync. Without AIDP, this means:

- **Manual data preparation:** Pipelines must be built and maintained for every data source
- **Stale AI context:** When files are created, updated, or deleted on NeuralMesh, downstream AI systems don't automatically reflect those changes
- **Repeated engineering effort:** Every new AI project rebuilds the same ingestion and vectorization work from scratch
- **Slow time-to-value:** Selecting, integrating, and validating an AI data stack takes months. Time-to-first-AI-outcome is measured in quarters, not weeks
- **Security risk:** Standard pipelines strip ACLs at ingest. AI deployment gets blocked until security teams verify governance holds through the AI layer

AIDP Offers Two Fundamental Benefits

1. Continuous, Automated Data Preparation Simplifies Workflows

AIDP watches WEKA NeuralMesh filesystems in real time. As new files, modified files, and deleted files hit the system, AIDP automatically handles the full preparation pipeline. AIDP:

- Detects every file create, update, and delete via continuous monitoring with the need for a batch scheduler, polling, or reindexing cycle
- Parses and chunks multimodal content from files as they arrive
- Embeds POSIX ACL, UID, and GID metadata at ingest — if a user can't see the file, they can't see the vector. No middleware required
- Keeps everything in sync and current as updates and deletes propagate automatically

2. Ready-to-Consume Data Generates Better AI Outcomes

Your team should be building AI applications, not babysitting ETL jobs. AIDP delivers AI outcomes directly to users and applications without requiring organizations to deploy or manage AI frameworks:

- Semantic Search offers a natural language queries over enterprise content stored on NeuralMesh, with results grounded in documents the user is authorized to see
- Users access these outcomes through a browser-based interface or API, backed by NVIDIA NIMs for embedding, reranking, and inference, all managed by AIDP and supported by WEKA
- WEKA Augmented Memory Grid™ (AMG) extends KV-cache by 1000x, eliminating recompute for long-context and agentic inference workloads

How AIDP Works: Platform Architecture

AIDP is a full-stack system of multiple integrated services that handles the complete journey from raw data to AI-ready outcomes in a secure, permissions-based environment. Key outcomes include:

Without AIDP	With AIDP
Custom data pipelines are built and maintained per project	One platform continuously prepares all NeuralMesh data for AI
AI projects use stale and outdated data when files change	Create, update, and delete events automatically propagate to the vector store
Engineering teams manage AI frameworks to get outcomes	Semantic search, and more delivered out of the box
Onboarding a new data source takes days	A browser form and minutes
AI responses ignore file-level permissions and governance procedures	POSIX ACL enforced from NeuralMesh through to every AI response
Siloed pipelines offer little or no operational visibility across pipelines	Live queue monitoring, status dashboards, and role-based governance in one UI

Get Started with NeuralMesh AIDP in the WEKA App Store

The WEKA App Store is the deployment and discovery interface for AIDP, offering a curated catalog of pre-validated AI blueprints that turn NVIDIA AI Factory intent into repeatable production workflows without custom integration work. The following blueprints are available:

- **NeuralMesh AIDP Core Pipeline:** The base ingestion, vectorization, and retrieval pipeline deployable as a foundation for custom AI applications
- **Semantic Search:** Natural language queries over enterprise NeuralMesh content, respecting POSIX ACLs. Delivered as a ready-to-use application
- **Video Search & Summarization (VSS):** NVIDIA VSS integration for video intelligence on content stored in NeuralMesh
- **OpenFold & AlphaFold:** Protein structure prediction for drug discovery and HCLS research
- **Market Signal Detection:** Real-time signal detection as financial data lands, without per-project pipeline engineering
- **AIQ & Agentic RAG:** Multi-step agentic retrieval-augmented generation for enterprise queries

See NeuralMesh AIDP in Action → [Click Here to Dive Deeper](#)