WEKA

How WEKA's Client Access Outperforms NFS and Powers AI Success

The "Built-For-Purpose" Series

Today's applications demand extreme throughput and incredibly low efficiency. Discover how WEKA stacks up next to NFS and which solution is the best choice for data-intensive AI workloads.

NFS (Network File System)

Reliable and widely adopted protocol for networked file systems for decades.

Designed for

- **Convenience and ease of use**
- **Simple file sharing**
- **Seamless compatibility across applications**

Performance Bottlenecks

- High I/O workloads
- Latency
- Metadata operations

Architecture Challenges

- Not designed for large-scale, distributed systems
- Data availability & corruption
- Potential single points of failure and downtime

The WEKA Client

Engineered from the ground up to help accelerate data-intensive AI and ML workloads

Designed for

- Speed
- Scalability
- Simplicity

Delivers

- Performance Efficiency
- Simplified data management
- High I/O & metadata ops with low latency
- Seamless integration with existing IT infrastructures
- Dynamically scales up and out to meet bursty data requirements at exabyte scale

WEKA vs. NFS: Head-to-Head

We compared the CPU utilization of NFS and WEKA performing large sequential operations (bandwidth-based tests) and small random operations (IOPS-based tests) from a single HDR InfiniBand (IB) connected client server.

We used a common load generating tool, FIO, to issue small and large reads and writes against each mount. We used the following FIO job files to execute load against an NFS mount and a WEKA client mount.

The number of CPU cycles consumed by NFS operations can vary significantly depending on several factors, including the type of operations being performed, the network speed, and the server and client hardware configuration. The same logic applies when mounting with the WEKA client in Shared mode, when not dedicating CPU cores.

To avoid resource contention, we only had a single mount active when benchmarking: either via NFS with nconnect or via the WEKA client, but never both simultaneously.

Bandwidth Results

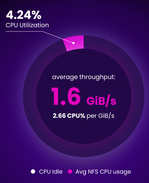

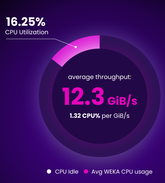

IOPS Results

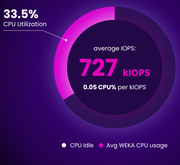

And the winner is...

The WEKA Client outperforms NFS by a wide margin.

- **7.5x** higher throughput
- **13x** more IOPS
- Fewer CPU resources used per unit of data processed
- Higher utilization

The WEKA client on the WEKA Data Platform is the perfect fit for demanding applications such as AI, machine learning, and high-performance computing (HPC).

- Can be deployed without significant changes to any environment
- Highly compatible with a broad set of Linux distributions
- Ubiquitous access to all data across all file and object protocols
- Allows for parallel processing and load balancing
- Enables faster AI model training, reduced time-to-insight, and rapid innovation

WEKA is the superior choice for organizations seeking to optimize their data infrastructure for high-performance workloads, ensuring faster data access, reduced time-to-insight, and greater overall efficiency in modern enterprise environments.

Learn more about how the WEKA Client is a fit-for-purpose solution optimized for the modern data stack.

WEKA

910 E Hamilton Avenue, Suite 430, Campbell, CA 95008 408.335.0085 info@weka.io www.weka.io