

Data is valuable on its own. When processed, it becomes the fuel for artificial intelligence, machine learning and advanced analytics. Are you ready to put your data to work?

DATASHEET

Hitachi Content Software for File: High-Performance File Services for AI, ML and Analytics

Putting Your Data to Work

In the race to make the most of artificial intelligence (AI), machine learning (ML) and high performance computing (HPC) based analytics, the winners are those who can run more models with more complex algorithms and do it faster than their competitors. To win, you need faster access to more data at a lower cost than your competition.

The Storage Bottleneck

Applications crave more data, but storage performance has been left behind. While graphics processing units (GPUs) have shrunk compute infrastructure by 40%, but the data they process has grown by 50%. Investments in compute resources and networks often sit idle, waiting for data. As a result, response times skyrocket but adding more compute is ineffective as legacy storage can't scale to tens or hundreds of PBs while maintaining high performance. On top of that, each workflow stage has unique compute, storage and networking needs. This leads to silos creating data management and integration challenges, which drives up costs and time to results, neither of which you can afford.

The Best of File and Object Storage

Hitachi Content Platform for File is a high performance storage solution for AI, ML, analytics and other GPU accelerated workloads. It gives you the blazing speed of a POSIX compliant distributed file system (DFS) with the capacity and hybrid cloud capabilities of an object store. As an integrated solution, it greatly reduces the complexity and deployment time. Its support for file and object protocols makes data ingestion easy. The DFS provides both high performance and low latency for data preparation, model training and inference. The object store provides massive storage capacity at a lower cost and offers powerful, data management automation driven by metadata.

AI and ML Use is Growing

According to polling results conducted in January 2020 by the AI Research Circle, 85% of infrastructure and operations (I&O) leaders are looking to use artificial intelligence (AI) in their infrastructures during next two years. AI workloads are diverse and some are fundamentally different from any other workload the organization may have run in the past. Although interest in leveraging AI applications is on the rise, I&O leaders are often unprepared to address storage requirements and data management challenges for growing datasets of large-scale machine learning (ML) deployment.

The common goals of these projects are to:

- Adopt new operating models
- Grow relevancy and competitiveness
- Speed response times
- Understand end user behavior
- Reduce operational costs

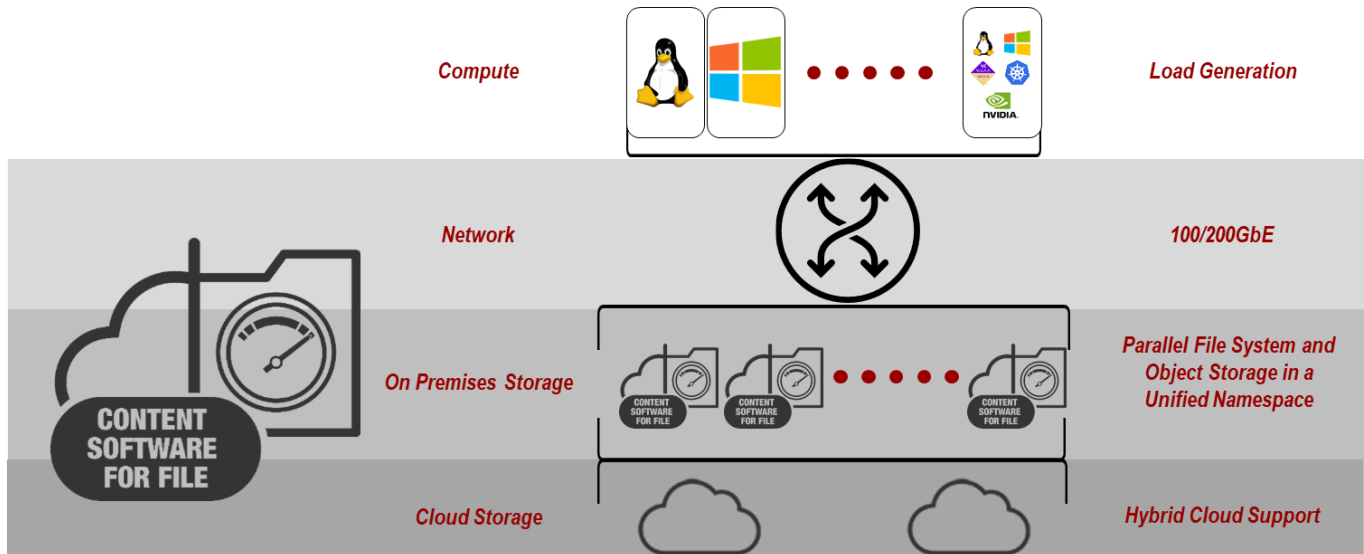
Hitachi Content Software for File

This unique integration of file and object storage offers you a tightly coupled, single solution for an appliance-like experience designed for ultra-high performance and capacity applications. For details on key components and performance specifications, see Table 1. Please note that these configurations are only examples and that we can create clusters of any size. For a high-level architecture view, see Figure 1.

TABLE 1: Hitachi Content Software for File Components and Performance

Hitachi Content Software for File Components	
Chassis	Quanta DS 120
CPU	Xeon Silver 4214 (Medium) Xeon Gold 6226R (Large)
Memory	192GB (12x 16GB)
Boot	480GB SATA SSD PM883
NVMe	2-10 NVMe SSD 4510 (2/4/8TB)
Network	25Gb Mellanox CX4 EN 100Gb Mellanox CX5 LX EN
Network Expansion	100Gb Mellanox CX5 LX EN
Capacity and Performance at Minimum Starting Configuration	
NVMe Capacity (TB)	18TB
Transactions (I/Os)	READ: 170,000 WRITE: 80,000
Throughput (GB/s)	READ: 36 GB/s WRITE: 6 GB/s

Figure 1. Hitachi Content Software for File High-Level Architecture



[Learn More](#)



Hitachi Vantara Federal



Corporate Headquarters
11950 Democracy Drive, Suite 200
Reston, VA 20190 USA
hitachivantarafederal.com

Contact Information
USA: 1-703-787-2900
hitachivantarafederal.com/about/contact-us/

HITACHI is a trademark or registered trademark of Hitachi, Ltd.

Feb 2021