

# IBM Storage Scale

Software-defined storage for building a global data platform for AI, analytics, and hybrid cloud

## Highlights

Global data abstraction services ensure seamless connectivity across multiple sources and locations.

Extensive protocol support, including NFS, SMB, S3, HDFS, POSIX, NVIDIA GPU Direct Storage, and container-native CSI interfaces.

Robust data resilience safeguards against ransomware and evolving cyber threats.

IBM Storage Scale System as appliance is also available, offering rapid deployment and building-block expandability.

Organizations worldwide are restructuring their data resources to modernize and capitalize on the opportunities presented by artificial intelligence (AI). However, this transformation comes with several key challenges:

- **More data:** Enterprises are generating and storing data at unprecedented rates, a trend that shows no signs of slowing down.
- **More locations:** Modern data strategies rely on distributed storage architectures to optimize performance, cost, and resilience.
- **More formats:** Organizations must manage a mix of structured data (SQL databases), semi-structured data (web pages, social media posts, and log files), and unstructured data (text, video, audio, and IoT sensor data).

A significant portion of this data is unstructured, generated by AI/ML workloads, analytics, data lakes, IoT, cloud-native applications, and backup and archive solutions. To ensure accessibility across geographically distributed applications, services, and devices, this data must reside in scalable, distributed file and object storage systems.

IBM Storage Scale is designed to address these evolving data demands. It provides global data abstraction services, enabling seamless data connectivity across multiple sources and locations. This capability extends to both IBM and non-IBM storage environments, making it an ideal solution for heterogeneous infrastructures. Built on a massively parallel file system, IBM Storage Scale can be deployed across multiple hardware platforms, including x86, IBM Power, IBM Z, ARM-based POSIX clients, virtual machines, and Kubernetes environments.

## Unlocking AI Potential with Content-Aware Storage

Very little enterprise data has been indexed for generative AI applications, which prevents AI assistants from providing accurate, up-to-date answers<sup>1</sup>. The content-aware storage capabilities in Storage Scale address this challenge by extracting the semantic meaning hidden inside unstructured data so that AI assistants can automatically generate smarter answers. Storage Scale enriches data using embedded compute and data pipelines that minimize data movement and latency to help reduce costs and improve performance.

**For the 9th consecutive year**, the 2024 Gartner Magic Quadrant lists IBM Storage Scale as a leader for distributed file systems and object storage.

According to Gartner, “By 2029, over 80% of unstructured data will be deployed on a consolidated storage platform instead of separate file and object products, up from 40% in early 2024”

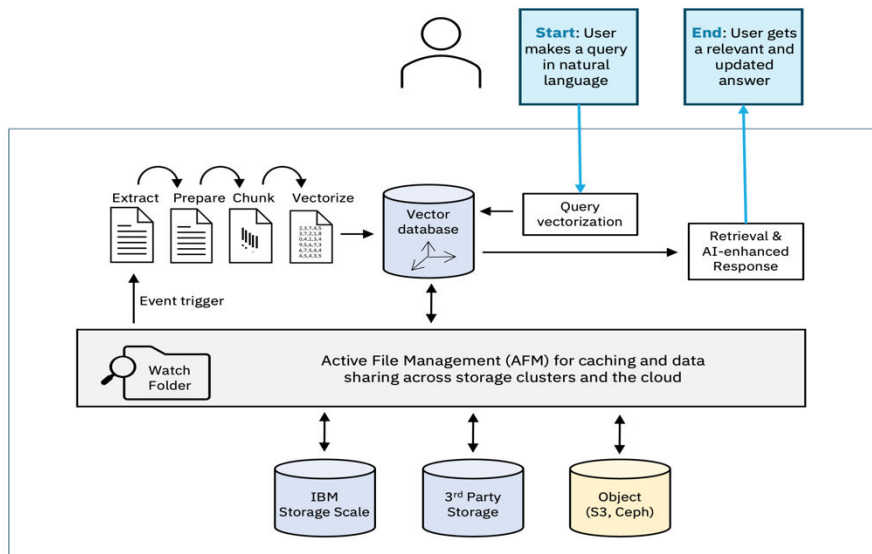


Figure 1. Storage Scale enhances AI-driven workflows by integrating Active File Management (AFM) and Content-Aware Storage to streamline data processing and retrieval.

Storage Scale automates data extraction, vectorization, and storage updates, enabling seamless retrieval via a vector database. When users submit natural language queries, AI enhances search results for optimized responses. The table below summarizes how Storage Scale handles data updates and AI-driven query processing.

Step	Data update process based on storage changes (Storage Event Trigger)	Query processing and response generation (User-initiated Trigger)
1	A file is created, modified, or deleted in storage	User submits a search query in natural language
2	The event trigger detects the change and activates the processing pipeline	The query is converted into a vector representation for similarity search
3	The changed data is extracted, prepared, chunked, and vectorized.	The storage system retrieves the most relevant data from the vector database
4	The processed data is stored in the vector database for future retrieval	Relevant data is retrieved, and an AI model enhances the response
5	Updated data is made available for AI-driven retrieval and response generation	User receives a relevant and updated answer

The main use cases for Storage Scale include GPU-accelerated AI, big data analytics and data lakehouses, high-performance computing (HPC) for scientific simulations and complex computations, IT modernization, and backup and archiving, ensuring efficient automated tiering and secure long-term data retention.

# Scalable File and Object Storage Services

Storage Scale provides a flexible, scalable approach to data management, allowing organizations to start with Base Data Services and seamlessly deploy and consume Abstraction and Advanced Data Services as needed. This enables efficient, AI-optimized, and resilient storage tailored to evolving business demands.

## Base Data Services

Base Data Services provide the foundation for scalable, high-performance data management, enabling seamless access across AI, analytics, and HPC workloads. With automated data tiering, advanced caching, and multi-location accessibility, these services optimize storage efficiency, streamline operations, and enhance resource utilization. Built-in data protection, governance, and access control ensure security and reliability, enabling organizations to manage large-scale datasets with confidence across hybrid and multi-cloud environments.

## Abstraction Data Services

Abstraction Data Services provide a differentiated capability to enhance AI economics by abstracting external storage and unifying diverse storage systems under a single global namespace. This enables simultaneous access via multiple protocols, streamlining data management across heterogeneous environments. By delivering a consistent, high-performance, and seamless experience, these services optimize both new and existing storage infrastructures, ensuring efficiency and scalability for AI-driven workloads.

## Advance Data Services

Enterprise resiliency features, including encryption, compression, disaster recovery (DR), high availability (HA), multi-site tiering, and safeguarded copies, provide continuous data protection and cyber resilience. Content-Aware Storage enhances AI workflows by leveraging natural language processing (NLP) to extract meaning from unstructured data. Additionally, support for storage-rich servers enables the distribution of data and metadata across a cluster's internal drives, transforming them into reliable, high-performance storage with minimal overhead.

---

Base Data Services	Abstraction Data Services	Advanced Data Services
<p><b>Base File System</b></p> <ul style="list-style-type: none"><li>• Multi-protocol access</li><li>• Quotas, QoS, Snapshots</li><li>• Remote mount</li><li>• ILM policies</li><li>• Scale-to-scale caching</li><li>• Single-site tiering</li></ul>	<p><b>Abstract External Storage</b></p> <ul style="list-style-type: none"><li>• Differentiated capability to improve AI economics</li><li>• Different storage systems in a single global namespace, accessible via multiple protocols</li><li>• Consistent, high-performance, seamless experience for new or existing storage</li></ul>	<p><b>Enterprise Resiliency</b></p> <ul style="list-style-type: none"><li>• Encryption, Compression</li><li>• Disaster Recovery (DR), High Availability (HA), Multi-site tiering</li><li>• Safeguarded copy</li></ul> <p><b>AI Data Services:</b></p> <ul style="list-style-type: none"><li>• Content-Aware Storage</li><li>• Fusion Data Catalog</li></ul> <p><b>Erasure Code Services</b></p> <ul style="list-style-type: none"><li>• Storage-rich servers</li></ul>

---

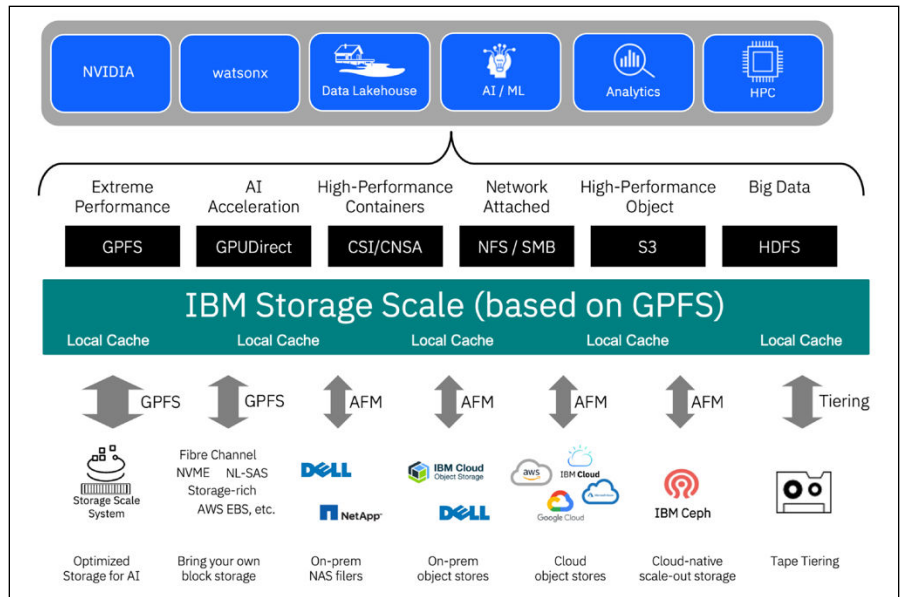


Figure 2. IBM Storage Scale provides a global data platform for your organization’s geographically dispersed devices, data sources, and workloads.

## IBM Storage Scale System

Storage Scale is also available as an appliance, IBM Storage Scale System, for streamlined, rapid deployment complete with IBM support services. This option is designed for organizations wanting to build high-performance global data storage capabilities in their own data centers or co-location facilities.

### Optimized for AI / NVIDIA workloads

Storage Scale System 6000 is the simplest and fastest way to deploy a global data platform for AI and NVIDIA GPU infrastructure, with exceptional performance that includes GPUDirect Storage support.




Figure 3. IBM Storage Scale System 6000 is a hardware appliance that allows you to deploy IBM Storage Scale on thousands of nodes with TB/s performance, low latency, and tens of millions of IOPS per node.

**For more information**

To learn more, contact your IBM Business Partner:

Fortuna Data

+44 (0)1256 331614 | [solutions@fortunadata.com](mailto:solutions@fortunadata.com)

[www.fortunadata.com](http://www.fortunadata.com)

© Copyright IBM Corporation 2025  
IBM Corporation  
New Orchard Road  
Armonk, NY 10504

Produced in the  
United States of America  
March 2025

IBM and the IBM logo are trademarks or registered trademarks of International Business Machines Corporation, in the United States and/or other countries. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on [ibm.com/trademark](http://ibm.com/trademark).

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

