

IBM Spectrum Scale Container Native Storage access

Making hybrid cloud storage easy with a
containerized global parallel file system


Highlights

- Scale container native storage w/o compute and data nodes
 - Data accessible outside Kubernetes
 - Global access to data (AFM) including object storage as a file
 - Policy based archive to tape/cloud
 - Enhanced access to storage resources
 - Dynamic provisioning – self service
 - Optimize container resources
-



Storage made simpleR

This is a story of how IBM is making storage simpler for data and AI. IBM Storage for Data and AI is a comprehensive portfolio with an AI information architecture that is made simple for fast and easy data access and AI modernization. Now IBM has an expanding story around kubernetes containers and a hybrid cloud architecture. IBM Spectrum Scale container native storage access brings even more simplicity for applications and the global parallel file system that makes hybrid cloud easier.



Organizations need a hybrid cloud strategy that includes data for

BOTH

- current applications and workloads
- new modern applications and platforms


- Lower TCO
- Unified management
- Global data access
- Interoperability
- Flexibility
- Agility
- Performance

Hybrid cloud data

Data is not easy with a hybrid cloud architecture. Since most organizations will build their IT strategies around a mix of on-premises private clouds, public clouds and their existing platforms the data needs to include both current applications and new modern applications and platforms such as Red Hat OpenShift.

IBM Spectrum Scale Container Native Storage access


Storage deployed inside containers, alongside applications running in containers, is an important innovation that benefits both developers and administrators. By containerizing storage services and managing them under a single management plane such as Kubernetes, administrators have fewer housekeeping tasks to deal with, allowing them to focus on more value-added tasks. Developers benefit by being able to provision application storage themselves that's both highly elastic and developer-friendly. With IBM Spectrum Scale container native storage access there are no limits of localized disks and performance is not effected by the application server but is optimized by external storage resources. IBM Spectrum Scale provides most of the benefits of container native storage without the limits of using only local disk and application server performance.



Containerized Access

IBM Spectrum Scale AI Information Architecture

A containerized IBM Spectrum Scale is a very powerful storage solution with the integration to the global parallel file system. For simplicity, it's the integration to the OpenShift console that makes Spectrum Scale very easy to use, manage and deploy. The enhanced CSI operator provides a self-service way to dynamically provisioning storage and make it easy to configure data capacity for a container in seconds. Another benefit is the easy way to increase throughput of data access with a single click as an optimized parallel path directly to the storage nodes is created.



IBM Spectrum Scale

IBM Spectrum Scale is the center of IBM Storage for Data and AI information architecture. The global parallel file system is the glue that creates a hybrid cloud storage solution for any enterprise or organization. Secure access. Anywhere. Enterprise data services. Everywhere. Hybrid cloud. Anyone.

An example of a containerized solution for AI workloads

New NVIDIA and OpenShift Deployment Guide using IBM Spectrum Scale

- Faster results with parallel and global access for AI workloads
- Red Hat OpenShift data integration
- Enterprise data services for mission critical AI
- DGX1, DGX2 and A100
- AI acceleration with dynamic caching from object stores to IBM Spectrum Scale

Advanced driver-assistance systems (ADAS)
Healthcare and Lifescience
Financial Services and Capital Markets
Government/Universities and Research

IBM Spectrum Scale

Red Hat OpenShift Container Platform

NVIDIA


NVIDIA and containers

IBM has worked with NVIDIA to create solutions for IBM Storage for Data and AI with NVIDIA DGX servers.

IBM Storage for Data and AI provides a modern way to use kubernetes containers and Red Hat OpenShift for AI workloads with NVIDIA DGX systems. IBM's AI information architecture provides a software-defined infrastructure powered by IBM Spectrum Scale for NVIDIA DGX systems for Machine Learning (ML) / Deep Learning (DL) workloads. You can reference the IBM rebook for more information: [Deployment and Usage Guide for Running AI Workloads on Red Hat OpenShift and NVIDIA DGX Systems with IBM Spectrum Scale](#)

The advantage of Container Native Storage access

Perhaps the most practical benefit to Spectrum Scale container native storage access is the ability to create dynamic provisioning, so as any container is deployed or requires more data, the storage can be allocated in real-time, something that traditional storage appliances couldn't cope with in a container environment. Customers turned to distributed file systems such as NFS, GlusterFS, and Ceph to add a persistence layer that cut across the nodes. Managing these filesystems was not integrated with Kubernetes tools and workflow. Storage administrators had to manually install and configure these file systems on every node before deploying workloads. Similar advantages to external storage is the ability for containers to access Spectrum Scale outside of a kubernetes platform. Some of the unique advantages include global access to data (AFM) including object storage as a file, Policy based archive to tape/cloud for kubernetes data and enhanced parallel access to storage resources.



The advantage of container native storage

Why IBM?

IBM Storage for Data and AI is more than storage products or even storage solutions. It consists of a storage strategy to help customers on their journey to AI and the hybrid cloud data center. IBM continues to drive leadership for scalable high-performance workloads as well as efficient, secure, scalable, capacity storage for file and object-based solutions. Our products provide an enhanced strategy for AI and the hybrid cloud. We provide a foundation for the future for the edge, the core data center and the public cloud including kubernetes containers and the Red Hat OpenShift platform. IBM Storage for Data and AI lowers complexity and cost with increase integration to an AI information architecture for the hybrid cloud that can be infused to the entire organization. Our message is easy “Storage Made SimpleR for Data and AI”

Next steps

- [Spectrum Scale Web Page](#)
- [IBM Spectrum Scale Data Sheet](#)

For more information

[IBM Storage for Data and AI web page](#)

© Copyright IBM Corporation 2020.

IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at <https://www.ibm.com/legal/us/en/copytrade.shtml>, and select third party trademarks that might be referenced in this document is available at https://www.ibm.com/legal/us/en/copytrade.shtml#section_4.

This document contains information pertaining to the following IBM products which are trademarks and/or registered trademarks of IBM Corporation:
IBM® Spectrum Scale



All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.