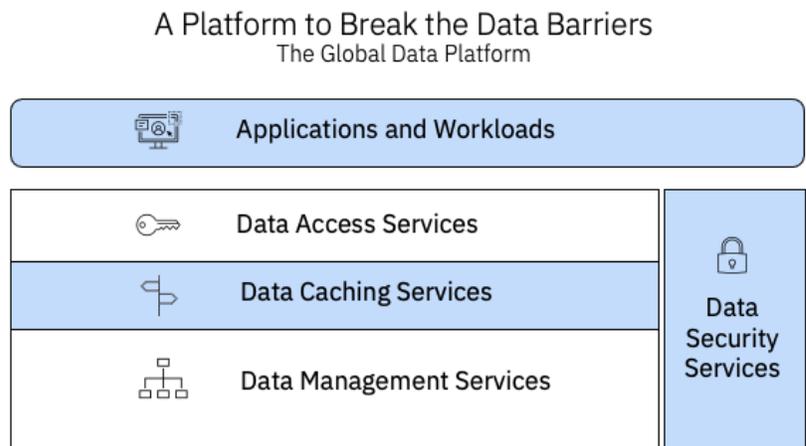


IBM Elastic Storage System

Highlights

- A single data platform for multiple application requirements
- Data access services that provide multiple high-performance interfaces
- Data caching services provide consistency and data virtualization to increase collaboration
- Data management services that automate data lifecycle and application performance
- Data security services that protect your data from network vulnerabilities and help ensure cyber resilience and quick recovery from attacks

Solutions for a global data platform



Unstructured data now holds the intellectual capital of organizations. As these valued assets grow to become significant to an organization, there needs to be something that steps forward in a new way to protect, access, and easily manage this growing amount of mission critical data. Something needs to solve the difficult data barriers that exist. Barriers such as performance with multiple new and existing application interfaces. Barriers such as continually expanding capacity and ease of scalability and silos of hard-to-find data. Data needs to be accessible globally especially in the cloud and because its accessible globally it needs to be secure and protected especially against cyber-attacks. That is why IBM® has built a framework called the global data platform. This platform consists of a set of core data services with each core component containing one or more specific data services to help break data barriers and solve customer application requirements.

The global data platform is targeted at new easy to order, easy to install, easy to upgrade, easy to use, software that is available in multiple deployment options. The simplest way to build a global data platform is with IBM Elastic Storage System® 3500 (ESS 3500). To learn more about specific value that customers have obtained with the IBM global data platform, read [Forrester's The Total Economic Impact™ Of IBM Spectrum® Scale](#)

The easiest way to build a global data platform



IBM ESS 3500

- Scalability from 46TB to 633YB
- Break performance barriers for AI workloads
- Access file and object data locally from the edge, the core and the cloud
- Reduce local capacity requirements by connecting data silos
- Secure data assets with multiple data security services

Application Workloads for the global data platform

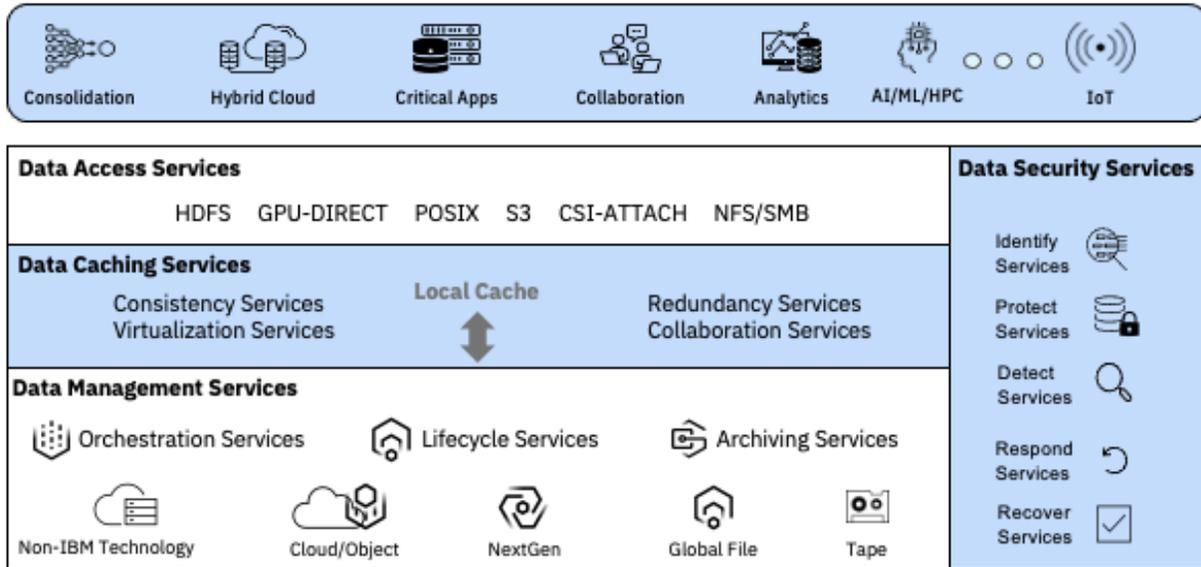


Workloads	Use Cases	Business Impact
AI/ML/HPC	NVIDIA, TensorFlow, Caffe, Pytorch, IBM Cloud Pak® for Data, ML Kit, AWS Lambda, AWS SageMaker ...	Faster GPU and AI analysis for faster decisions to adjust to market dynamics with access to more data and connected data
Analytics	Cloudera, Hadoop, Apache Spark™, SAS®, Tableau, Python, Power BI, ...	Eliminate silos and speed time to value with data caching services and combined HDFS and S3
Collaboration	Media Development, AI Model Development, Computer Aided Engineering	Increase workgroup productivity with easier access to more data from edge to core to cloud
Critical Apps	Business critical applications that have expanding data requirements	Maintain business continuity with data that is always online and secure with security services
Hybrid Cloud	Red Hat® OpenShift® workloads, backup and archive, applications that move data or access on-prem and cloud	Faster time to production as data is consumed anywhere and connected everywhere
Consolidation	Data lake, data ocean, and big data	Lower costs as data is resilient and accessible without duplication

The Global Data Platform is all about solving customer application requirements with a data platform that provides business impact. There are six core workloads or solution areas. AI, ML, HPC as one core solution area which contain use cases for applications or application environments such as NVIDIA, TensorFlow, Caffe, Pytorch, IBM Data Fabric or IBM Cloud Pak for Data, Google ML Kit, AWS Lambda and AWS SageMaker to name a few. The business impact using the Global Data Platform brings faster GPU and AI analysis for faster decisions to adjust to market dynamics with access to more data and connected data. That is just one example of the power of the Global Data Platform.

Contact your IBM Sales representative or value added business partner for details about pricing and your specific configuration and solution requirements.

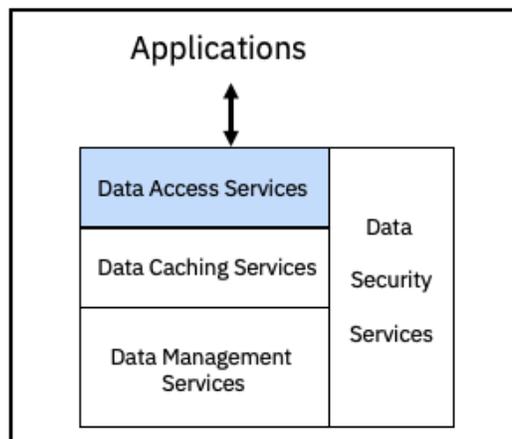
The Global Data Platform



The details of the global data platform are actually not complicated but they contain over 30 years of investment and research at IBM. With over 4000 customers in almost every industry across the globe the global data platform is a proven platform that has been used in the fastest supercomputers to banking systems that have never lost access to data in years. It has been used by many of the top corporations in the world with PBs of data to service millions of users and it also can be deployed in small businesses that only needs TBs of high performance data.

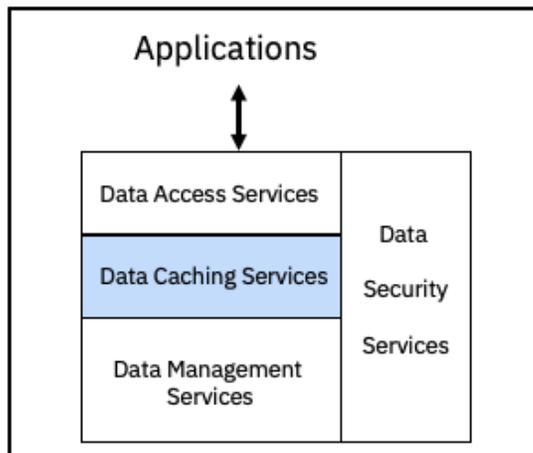
Data Access Services that provide multiple high-performance interfaces

- Performance Enhanced**
- GPU Direct Storage (GDS)
- POSIX (direct connect)
- High Performance Object (S3)
- Container Attach (file/volume)
- Analytics (HDFS)
- Traditional**
- NFS (NFS v3/v4)
- Windows (SMB)



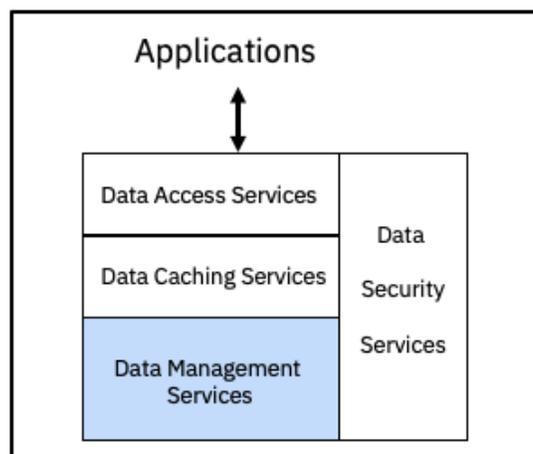
Data Caching Services provide consistency and data virtualization increasing collaboration

- Virtualization Services
- Collaboration Services
- Redundancy Services
- Consistency Services



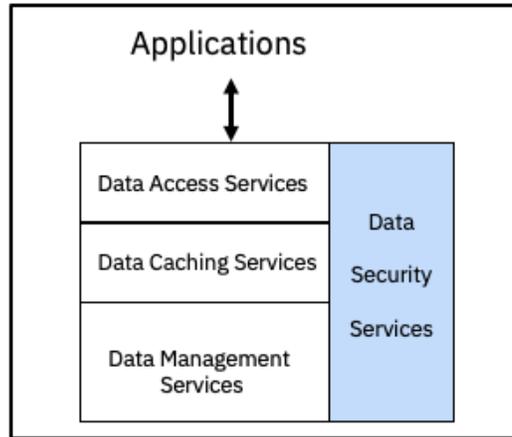
Data Management Services that automate data lifecycle and application performance

- Orchestration Services
- Lifecycle Services
- Archiving Services



Data Security Services that protect your data

- Identity Services
- Protection Services
- Detect Services
- Respond Services
- Recovery Services



The global data platform: Customer success story

A large multinational company and one of the world's largest corporations in the world with revenue of over \$100B

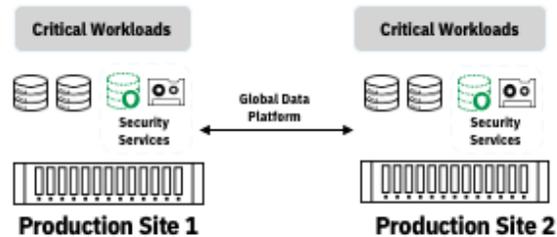
Business Challenge

Government retention requirements with business-critical data that need faster recovery response times. This led to discussions and planning for extending customers storage infrastructure business continuity capabilities with IBM for mission critical workloads.

Results

\$100M+ Estimated savings with low RTO	100% Critical AI and Unstructured Data Systems Protected	50PB+ Total Amount of Data Protected
--	--	--

IBM Solution



IBM Storage for Cloud Scale Data Sheet



* Copyright IBM Corporation 2022
IBM Systems
Route 100
Somers, New York 10589
U.S.A.
Produced in the United States of America,
05/2022

Why IBM

IBM Spectrum Scale and the IBM Elastic Storage System provide a way to solve difficult problems with more data and faster access to that data.

Each storage system from IBM offers a full-featured set of data services which are powered by the Global Data Platform. The core data services include Data Access Services, Data Caching Services, Data Management Services and Data Security Services.

Each core data service offers their own set of Custom data services that create an easy to manage and powerful set of capabilities that differentiate IBM's cloud scale products.

Find out more about our entire product portfolio including our powerful IBM Cloud® Object Storage and IBM Spectrum Discover products.

Next steps

[IBM Storage Cloud Scale Solutions](#)

IBM, the IBM logo, IBM Cloud, IBM Cloud Paks, IBM Elastic Storage System, Power, and Spectrum Scale are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom. InfiniBand and InfiniBand Trade Association are registered trademarks of the InfiniBand Trade Association.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the [OpenStack website](#).

Red Hat®, JBoss®, OpenShift®, Fedora®, Hibernate®, Ansible®, CloudForms®, UNIX is a registered trademark of The Open Group in the United States and other countries.

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

References in this publication to IBM products and services do not imply that IBM intends to make them available in all countries in which IBM operates.

This information is provided "as is" without warranty of any kind, express or implied, and is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this document. Nothing contained in this document is intended to, nor shall have the effect of, creating any warranties or representations from IBM (or its suppliers or licensors), or altering the terms and conditions of the applicable license agreement governing the use of IBM software.