

Accelerate digital transformation

IBM Storage

with data services and persistent storage for hybrid cloud



- 01** The right storage foundation can accelerate transformation
- 02** Containers are vital to executing hybrid cloud strategies
- 03** Modernize and containerize with IBM Cloud Paks, built on Red Hat OpenShift
- 04** Get data services and persistent storage for your hybrid cloud with IBM Storage Suite for IBM Cloud Paks
- 05** How to choose the right storage solution for your hybrid cloud strategy
- 06** Why organizations are partnering with IBM on their hybrid cloud journey



Introduction

Storage infrastructure can make or break hybrid cloud success, and yet it's often an afterthought during digital transformation. Many business leaders don't realize how critical storage is to their cloud strategy's execution until they run into problems.

A hybrid cloud strategy leverages on-premises IT infrastructure in addition to public and private clouds. And while 85% of IT decision-makers agree that on-premises infrastructure is critical to their hybrid cloud strategies, a study by Forrester Consulting¹ found that many firms have delayed infrastructure refreshes at least a few times in the last five or more years. Following the delay, half of IT decision makers surveyed found security vulnerabilities and 39% experienced diminished performance.

If you delay infrastructure refreshes and upgrades, you may face these and other issues, such as compatibility restrictions and higher costs for the business.

On the other hand, with the right storage foundation, organizations can secure data at rest and in transit, ensure compatibility among applications and infrastructure, and improve performance and uptime, among other benefits.

In this guide, you'll discover the role that storage plays in implementing hybrid cloud strategies. You'll see why your organization needs the right storage infrastructure to provide persistent storage for container environments. And you'll discover how to choose the right storage solution to support your organization on its digital transformation journey.

Avoid gaps in your approach to digital transformation

The primary goal of digital transformation is to improve business agility, which encompasses your ability to respond to changing market conditions, predict consumer trends and better meet customer expectations.

As such, a significant portion of digital transformation revolves around shifting resources to the cloud to better leverage the power of AI and achieve flexibility and scalability.

This is where organizations can run into issues. Too often business leaders equate "shifting resources to the cloud" with "halt investments in on-premises infrastructure."



1 . "The Key To Enterprise Hybrid Multicloud Strategy," a commissioned annual study conducted by Forrester Consulting on behalf of IBM, January 2021.



Without the appropriate on-premises infrastructure as the foundation of your hybrid cloud strategy, you won't realize the full return on your investments in AI and cloud.

Meanwhile, AI and big data projects have immensely high data-usage requirements for storage infrastructure. Suddenly data is getting bigger, but systems are running slower.

This is a key reason why a hybrid cloud strategy for digital transformation delivers more value than a public cloud-only approach. It's also why the right storage infrastructure is so critical to hybrid cloud success.

To avoid gaps in your approach to digital transformation, build a hybrid cloud environment where you can leverage the right IT infrastructure for the right workload. In other words, bring together a mix of on-premises, private cloud and public cloud infrastructure to create an open and flexible IT environment that allows for greater business agility and access to data for business insights.

Storage is the key to your hybrid cloud success

In the next chapters, you'll learn how organizations are using containers as part of their hybrid cloud strategy, why persistent storage is so important for containerization, and how to choose the right storage solution to support your hybrid cloud strategy.

We'll also discuss how your organization can bring enterprise data services to hybrid cloud environments, simplify container deployments and enable a faster and more reliable way to modernize to the cloud.

Organizations are adopting hybrid cloud models to drive business transformation. To execute their hybrid cloud strategies, enterprises are increasingly choosing to [use containers](#) over more traditional deployment models, such as virtual machines (VMs). According to research by Enterprise Strategy Group (ESG), 74% of organizations are currently using containers in production or plan to do so within a year².

[Containers](#) have become the go-to deployment model because they bring together [developers and IT operations](#) and enable a more agile and collaborative way of building, deploying and managing applications.

What benefits do containers provide?

Containers bundle the code and dependencies needed to run one or more applications – including runtime environment, system tools, libraries etc. – into a single, self-contained package. They're lightweight by design and require minimal operating system and system resources to function.

Because of their lightweight nature, they can be deployed very quickly and are easily ported across environments. Multiple containers can run on a single host,

making them significantly more resource-efficient than virtual machines, for example. Containers provide several [business benefits](#) as well, including:

- Improved business acceleration
- Improved DevOps productivity
- Easier compliance
- Lower costs

What challenges do containers present?

Despite the numerous benefits containers provide, there is one inherent challenge that organizations must address: a lack of persistent storage. Containers use ephemeral storage by default. Here are the differences between these types of storage:

Ephemeral storage: Temporary storage tied to a specific compute instance. Data created or changed is terminated and doesn't exist after the instance ends.


Persistent storage: Long-term storage not tied to a specific compute instance. Data created or changed is saved and exists after an instance is terminated.

Persistent storage is critical for container environments

Containers were originally designed for stateless applications, which don't require persistent storage. Now containers are also used for stateful applications, which require consistent and persistent data whether on premises, in private clouds, or in the public cloud.

A stateful application needs to maintain its data and configuration beyond the life of a container. Because containers are ephemeral by nature, however, if something happens to the container instance or if it's moved to another host/server, the container loses access to the data.

To get the persistence needed to run stateful applications, then, you'll need to connect to other storage resources, such as object storage, block storage or file-based storage³.



2. "Data Storage Trends in an Increasingly Hybrid Cloud World," Enterprise Strategy Group (ESG), March 2020.

3. "Storage concepts," Openstack, November 2018.

How do I achieve persistent storage for containers?

To provide [persistent storage for containers](#), you can leverage an open storage API called Container Storage Interface (CSI). CSI allows a container orchestrator, such as Kubernetes, to connect cloud-native applications with persistent data using a declarative method that is independent of the storage provider and the cloud vendor.

In this approach, container-native storage is a commodity differentiated by performance, security and resiliency characteristics amongst others. Clients and their technical architects can select persistent storage based on economic and efficiency criteria. System administrators then make simple declarative configurations to ensure that deployed applications request the storage they need according to operational, governance and compliance requirements mandated by the business.

So, to achieve persistent storage for containers and effectively execute hybrid cloud strategies, organizations need the right hybrid cloud platform, container orchestrator, container platform, storage resources and data services.

In the next few chapters, we'll review how IBM's offerings, including IBM Cloud® Paks deployed with Red Hat® OpenShift® and [IBM Storage Suite for IBM Cloud Paks](#), help to provide persistent storage for containers and ease the container management burden.



What are IBM Cloud Pak solutions?

[IBM Cloud Paks](#) are AI-powered software for hybrid cloud that are designed to help you advance digital transformation with prediction, security, automation and modernization capabilities. You can deploy them across any cloud to accelerate development, deliver seamless integration and enhance collaboration and efficiency.

Built on [Red Hat OpenShift](#), the industry's leading enterprise Kubernetes platform, these solutions provide an open environment for developers, data managers and administrators to build new cloud-native applications, modernize existing applications and move more workloads, faster, to cloud and AI.

Each IBM Cloud Pak® includes pre-integrated containerized IBM middleware and common software services for development and management, on top of a common integration layer.

With IBM Cloud Paks, your organization can:

- **Modernize with ease:** Develop and consume cloud services anywhere, from any cloud.
- **Predict outcomes:** Collect, organize and analyze data regardless of its type or where it lives.

- **Automate at scale:** Implement intelligent workflows in your business using AI-powered automation.

- **Protect your business:** Generate deeper insights into threats and risks across hybrid cloud environments.

[IBM Cloud Pak for Data in 2 minutes >](#)

In 2020, Forrester completed an assessment of eight of the most significant container management platform providers in the market. The report names IBM and Red Hat as a leader, with the highest ranking among providers in the current offering, strategy, and market presence categories⁴.

How can your organization benefit from IBM Cloud Paks?

IBM Cloud Paks let you develop applications once and deploy them anywhere, integrate security across the IT landscape and automate operations with intelligent workflows.

These capabilities benefit your organization in several ways, including⁵:


- **Improving operational efficiencies:** Surges in data center resources can impact availability and reliability. Leveraging container orchestration such as IBM Cloud Paks in the data center

gives you the flexibility to package and move all or parts of an application at scale elsewhere in your hybrid cloud environment until digital demand for that application returns to normal.

- **Improving security and data privacy:**

A study by Forrester Consulting found that 46% of IT leaders feel that public cloud doesn't meet their security needs¹. However, an on-premises IT infrastructure with encryption built into both hardware and software can facilitate end-to-end security, and containers augment that security. Data inside of mainframes, servers and storage — protected by pervasive security — can be used by microservices, travel to within your hybrid cloud using containers and return to rest securely.

- **Improving productivity:** Open source and containers help speed up innovation and modernization, allowing DevOps to get applications to market better and faster. Often that innovation is used to automate repetitive tasks, replace manual processes and reduce human intervention. Containers can then scale those innovations across your hybrid cloud environment, decreasing friction from legacy processes and improving service delivery.



4. "The Forrester Wave™: Multicloud Container Development Platforms, Q3 2020," Forrester Research, Inc., September 2020.

5. "Open and ready for hybrid cloud innovation with IT infrastructure," IBM.

When it comes to critical applications, you need to know that your data is high quality, resilient, and available when you need it – with the assurance that it is secured and meets compliance requirements. IBM Storage Suite for Cloud Paks delivers the data services that IBM Cloud Paks need to produce and consume data, so your organization can more rapidly gain meaningful insights. It provides secure container deployment and persistent storage while optimizing data services to meet all workload and performance requirements.

In the next chapter, we'll explore how Storage Suite for Cloud Paks can help enterprises more effectively integrate data services and storage resource management into their hybrid cloud strategies.

What is IBM Storage Suite for IBM Cloud Paks ?

[IBM Storage Suite for IBM Cloud Paks](#) is a comprehensive set of flexible software-defined storage solutions designed specifically to support container environments while complementing the deployment of IBM Cloud Pak solutions.

[Software-defined storage \(SDS\)](#) is a storage architecture that separates storage software from its hardware. Unlike network-attached storage or storage area networks, SDS performs on any industry standard system, which removes the software's dependence on

proprietary hardware, and makes it hardware agnostic.

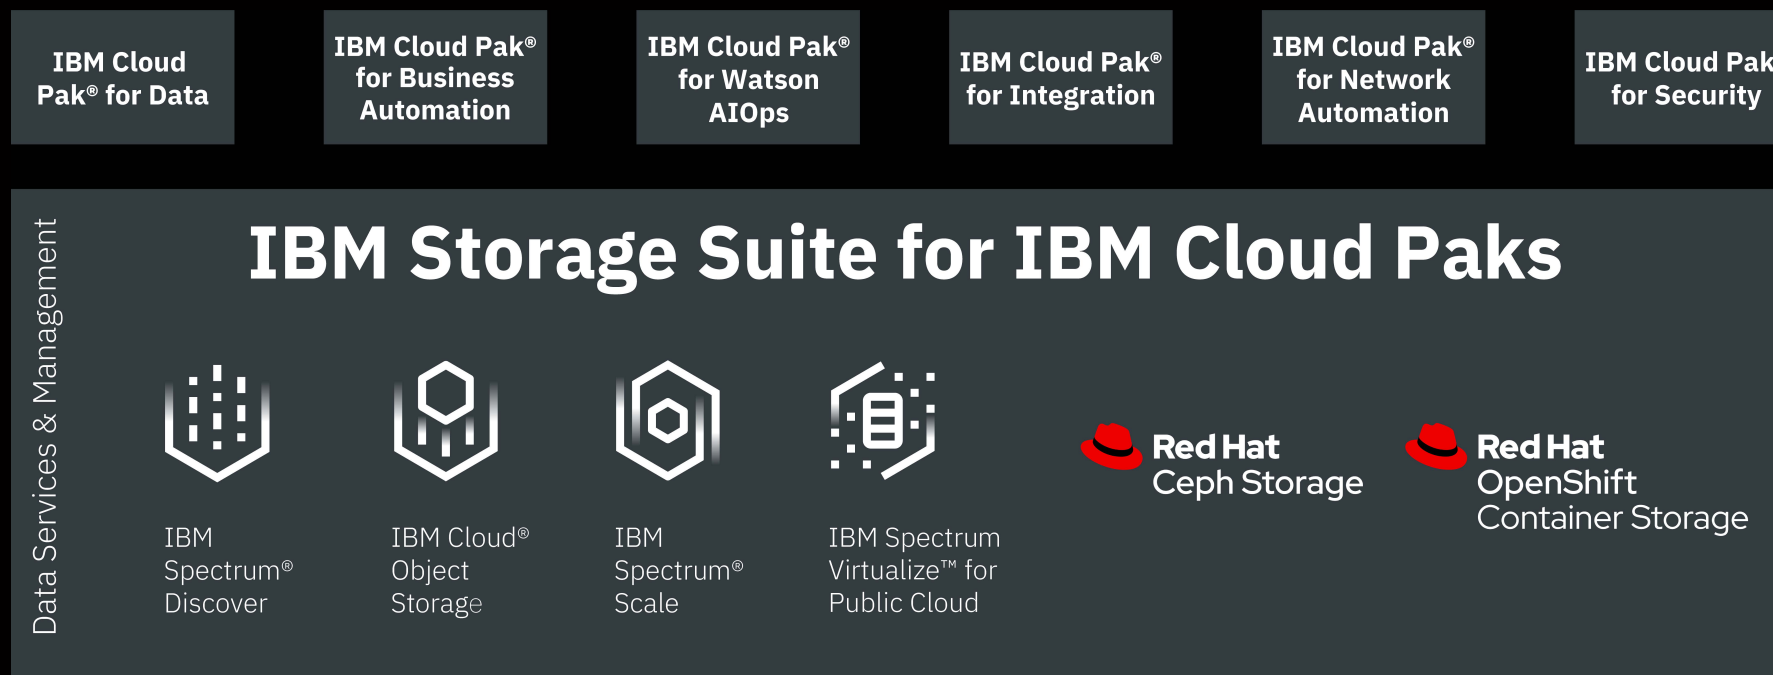
SDS allows you to upgrade or downgrade your storage capacity as you need to, when you need to, without having to add or remove proprietary hardware. The benefit is more control and flexibility over your data, which data sources you choose to use, and how you build your storage infrastructure.

IBM Storage Suite for Cloud Paks is a single software-defined suite that includes file, block and object storage from Red Hat and IBM. You can deploy any of the offerings in the suite, deploy more than one offering at the same time or change licensing entitlement according

to your workloads and needs.

Data services from IBM and Red Hat deliver simplified access to data wherever and however it exists, while providing a consistent experience independent across many platforms. And most importantly, data services from IBM and Red Hat perform all of this at scale, because businesses are coping with massively increasing amounts of data on a daily basis.

These offerings are fully tested and validated with Kubernetes, Red Hat OpenShift and IBM Cloud Paks for security, compliance and version compatibility. For a deeper dive, [please read our Storage Suite for Cloud Paks data sheet >](#)





SDS offerings included with IBM Storage Suite for Cloud Paks:

- [IBM Spectrum Scale](#): Advanced storage of unstructured data for cloud, big data, analytics, objects and more.
- [IBM Cloud Object Storage](#): A highly scalable cloud storage solution for unstructured data that provides on-premises and cloud-based dedicated services.
- [IBM Spectrum Discover](#): Multi-source data catalog that automatically indexes objects and files whenever changes are made using the metadata in real-time. It can be incorporated into Red Hat OpenShift environments.
- [IBM Spectrum Virtualize for Public Cloud](#): Hybrid cloud solution for data mobility, disaster recovery and cloud-storage optimization.
- [Red Hat Ceph Storage](#): Offers block, object, and file storage in one platform that scales to support 100s of petabytes.
- [Red Hat OpenShift Container Storage](#): Native storage data services for OpenShift, based on the Ceph data-plane, providing Kubernetes, file, block, object and multi-cloud gateway services.

By now you should have a solid understanding of how IBM Cloud Paks, alongside Red Hat OpenShift, and IBM Storage Suite for Cloud Paks can support your container management and drive your hybrid cloud journey to digital transformation.

In the next chapter we'll go over a few signs that can indicate when it's time for an infrastructure refresh as well as the criteria you should consider as you evaluate storage solutions.



So how do you know when you need to reinvest in your on-premises infrastructure? It's not always easy to recognize pain points that are related to or rooted in poor storage and data management.

For most organizations, the aspects of their storage infrastructure most in need of improvement are data security, ease of management, speed of data throughput in data-intensive applications, and sharing data across multiple locations.

The following sentiments have been gathered from client interactions and reflect common pain points among organizations. See if any of these statements sound familiar:

- “Data keeps getting bigger, and these next-gen apps are accelerating the pace of data getting bigger. But my budget isn't getting that much bigger, so I need to cut infrastructure costs.”
- “I don't want to be locked in to any one vendor. But I also want to avoid having multiple vendor silos each with different APIs, capabilities and procedures.”
- “I need to make things go faster, a lot faster, to satisfy the needs of my lines of business users.”

If you're experiencing these or similar challenges, it's likely time to update your existing storage infrastructure and/or invest in new infrastructure technologies.

How to choose the right storage solution

To support a hybrid cloud environment from development to mission-critical production requires three foundational data requirements that you will need to consider when evaluating storage solutions:

- **Data resilience:** Your storage solution should assure sustainable data quality through threat prevention and detection, backed by rapid recovery for any use case in any location.
- **Application modernization:** You'll need to build a hybrid storage infrastructure that can support bare metal, containerized and virtualized applications wherever they are deployed – on premises, in containers, through edge computing and multiple clouds.
- **DataOps and data services:** Your storage solution should exploit pervasive AI with machine learning, automation and self-service to manage and store data as efficiently and cost-effectively as possible.

The strongest software-defined storage solutions for hybrid cloud will integrate with graphics processing units (GPUs) and provide metadata management and resiliency features. Additionally, they can provide insights for optimized data placement, offer a modern data lake with global access, compress and deduplicate data to save space, and encrypt.

Why organizations are partnering with IBM on their hybrid cloud journey

To really understand the power of IBM Storage innovation, it helps to think of storage as the software foundation on which you build your data services. With the right infrastructure-independent software foundation, your hybrid cloud storage infrastructure will be flexible to absorb the latest innovations, and you'll be able to quickly adopt new technologies for maximum benefit.

IBM Storage Suite for IBM Cloud Paks is a comprehensive solution that provides persistent storage and data services for container environments.

The suite includes both IBM Storage and Red Hat Storage offerings, provides flexible licensing and consumption of storage resources, and contains those unique benefits mentioned above.

Most importantly, it's uniquely positioned to help organizations speed development and modernize applications without compromising security or continuity.

There are several factors that set Storage Suite for Cloud Paks apart from other solutions in the market, including:

- **Application workload versatility:** Choose file, block or object storage for your container data services and deploy one or more of the offerings in the suite according to your workloads and needs. As your needs change, you can deploy new data resources accordingly.

- **Unified solution with platform independence:** IBM Storage Suite for Cloud Paks is a unified, hardware-agnostic solution featuring both Red Hat and IBM software-defined storage offerings with cloud-native and cloud-ready storage. You can develop and deploy on premises or on private or public cloud using open industry standards, without vendor lock-in.

- **Data services for Red Hat OpenShift:** Simple to install and integrated with Red Hat OpenShift, IBM Storage Suite for Cloud Paks is the essential hybrid cloud storage software foundation for building and managing on-premises, containerized applications with persistent storage and data services.

- **Scalability:** IBM Storage Suite for Cloud Paks delivers performance at any scale for every workload with a true scale-up or scale-out architecture.

- **IBM validated:** Fully tested and integrated with Kubernetes, Red Hat OpenShift and IBM Cloud Paks to ensure ongoing security, compliance and version compatibility.





Next Steps

As you know, digital transformation is about much more than moving data and workloads to the cloud. It's aligning your technology with business goals and strengthening the strategy behind important IT infrastructure decisions.

Laying the appropriate storage foundation will help you extract the full value from your AI and cloud investments. Organizations partner with IBM to ensure they have the right strategy and the right technology in place for their hybrid cloud journey and larger digital transformation.

Take the next step:

Learn more about: [IBM Storage Suite for IBM Cloud Paks](#)

Watch the webinar: [Storage made simple for containers and hybrid cloud](#)



© Copyright IBM Corporation 2021. U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp. NOTE: IBM web pages might contain other proprietary notices and copyright information that should be observed.

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 "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.