IBM Storage Ready Nodes

Optimize and simplify software storage flexibility with cloud scale infrastructure support and integration

Highlights

A simple building block approach to software defined storage and infrastucture all supported by IBM

Cloud like scale options optimized and designed to be easy to build a software defined solution

Configuration options that have been tested and certified with IBM Storage software

Start with a configuration that meets your current needs and simply scale with additional capacity when required

Investment protection as IBM Storage Ready Nodes can be integrated into existing IBM Storage software configurations

As customers look to lower the cost of their data center storage and at the same time move to a more cloud like model and optimize valuable internal resources, a simple and low cost solution for on-premises data is needed to help energize data center infrastruture with software flexibility. IBM Storage Ready Nodes are a simple way to deploy IBM storage and data resiliency software on industry standard servers that enable an optimized and supported configuration in flexible and cloud scale configurations. IBM Storage Ready Nodes are designed to close the gap between software-only cloud scale storage systems and legacy storage offerings. By using IBM Storage Ready Nodes customers can accelerate storage infrastructure modernization, reduce project risk, and improve storage efficiency by providing industry standard servers and industry leading storage software with optimized and certified configurations supported from a single vendor for multiple customer requirements. This innovative approach to storage software makes it easier to build, deploy and expand software defined storage solutions in the data center.

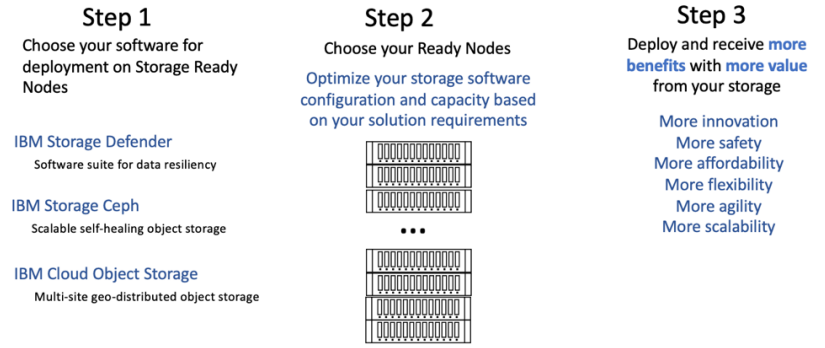

IBM Storage Ready Nodes offer a preconfigured, validated, and scalable infrastructure for organizations looking to create a software-defined solution on an optimized platform and also simplify the management, support and deployment of the overall solution in the data center. IBM provides a simplified deployment process, optimized performance options, on-line scalability from a base configuration to 1000s of nodes depending on the software option chosen.

Product overview

IBM Storage Ready Nodes are a simple, flexible and cost effective way to deploy IBM storage software on validated and optimized industry standard hardeware with both software and hardware support from a single vendor. IBM Storage Ready Nodes are designed to be simple all the way from initial purchase and deployment to the ongoing management and scalability of the system. Firt choose the IBM Storage software you want to deploy . Currently there are three software options, IBM Cloud Object Storage, IBM Storage Ceph and IBM Storaeg Defender.

Once the software option is chosen the next step is to choose the Ready Node devices required and the quantity or number of nodes. Each software option supports different amounts and different types of IBM Storage Ready Nodes. Each software section will provide information to help you choose the optimal configuration for your requirements.

After you select your configuration then work with your IBM Business partner or IBM Representative to purchase and deploy your IBM Storage Ready Nodes and the corresponding IBM Storage software. Once the storage solution is purchased, download the installation instructions with the software license that is provided to install the Ready Nodes with the software component. IBM Storage Ready Nodes will be shipped directly to you from IBM and you may install them into your environment yourself or optionally have them installed with the help of IBM or your Business Partner.

What you get with IBM Storage Ready Nodes

IBM Storage Ready Nodes accelerate storage infrastructure modernization, reduce project risk, and improve storage efficiency by providing IT flexibility and simplified management and support from a single vendor for multiple customer requirements. Customers gain a cloud like experience for storage software with optimized choices and a single source for problem resolution and support. This adds up to faster troubleshooting and problem resolution with a solution that scales quickly and is easy to deploy and manage. IBM Storage Ready Nodes are industry standard servers that can be installed in most any standard rack. (See specifications for details)

Support from a single source

Each IBM Storage Ready Node comes with IBM support and maintenance. The same support that you receive from your storage software. That means that problems are resolved faster without the need to determine if the issue is software or hardware. IBM Storage Ready Nodes are intended to be only used with IBM Storage software and a corresponding software option should be purchased with a Ready Node configuration.

Solutions that can be customized for your needs

From the edge to the core data center to the public cloud, IBM and its business partners can work along side you to plan, advise, and execute your IT infrastructure and digital transformation journey. We stay with you every step of the way, linking people, process, and technology to accelerate innovation and achieve optimal business outcomes.

Deployment and upgrade assistance when you need it

To make your IT investments as productive as possible, IBM and its Business Partners can provide smart planning, bulletproof data migration, upgrade assistance and high performance and reliability optimization assistance.

Software options for IBM Storage Ready Nodes

IBM Storage Defender
IBM® Storage Defender is a software suite for data resiliency.

IBM Storage Ceph
IBM Storage Ceph is a multi-workload enterprise data platform, to build, modernize and deploy applications at scale for data and AI workloads

IBM Cloud Object Storage
IBM Cloud® Object Storage is a geo-dispersed multi-site object store for cloud scale secondardary data in the data center for data and AI workloads

Ready Node Options for IBM Storage Defender

IBM® Storage Defender is an application that helps reduce data loss and risks associated with the protection of your data. This solution provides critical capabilities to an organization that can identify, protect, detect, respond, and recover data across your hybrid cloud storage infrastructure. The application also provides a simple, consolidated view of data protection and cyber resilience status with integration into your security dashboards. IBM Storage Defender provides multiple layers of data resilience, including data protection, data immutability and data isolation. It enables early detection of threats including risk from ransomware, disasters, sabotage, accidental deletion, and other sources of disruption. It also provides rapid recovery across your hybrid cloud storage infrastructure, while providing a simple, consolidated view of data protection and cyber resilience status with integration into security dashboards. Using Resource Units to deliver capabilities, you can streamline your software acquisition and license management by purchasing a single license that provides only the data resilience resources you need. For more information on IBM Storage Defender software read the data sheet ->

Storage Defender with Storage Ready Nodes Specifications

	Storage Defender Node
Processor	Intel Xeon Silver 4314
# of Processors	2
RAM	128GB
OS Disks	2x240GB SSD
Rack Height	2U
Width	482 mm (18.97 in.)
Depth	772.11 mm (30.39 in.)
Height	86.8 mm (3.41 in.)
Weight	35.3kg (77.82 lb) max
# of Capacity Disks	12
HDD Disk Sizes	4TB, 8TB, 12TB, 16TB
Power Supply	Dual 800W
Network	
4 x 1GbE	X
4 x 10GbE	X
32Gb FC HBA	Planned, Optional

Storage Defender Node

Configuring IBM Storage Ready Nodes with IBM Storage Defender

Qualified as a performant and resilient cluster node offering, the IBM Storage Ready Node for IBM Storage Defender Data Protect is ready for your demanding Data Protect protection workloads.

Customers can choose the storage combination that best suits the data resiliency and storage needs of their data center. The flexible server is comprised of SAS drives and NVME drives to retain business critical data for quick recovery in the event of an outage or attack. The node is designed for capacity but the Intel Xeon-silver processor ensures there is no sacrifice to performance.

The hardware nodes can be purchased with IBM Storage Defender subscription using Data Protect or purchased later on to add onto your existing IBM Storage Defender Data Protect cluster. IBM Storage Defender Data Protect provides the management plane to easily set policies that enable recovery from ransomware, cyber attacks, and natural disaster.

The IBM Storage Ready Node platform forms the foundation for an advanced scale out cluster that is highly secure, resilient, and performant for IBM Storage Defneder. For more information on designing the ideal solution for your enterprise data resiliency needs, seek the advice of your IBM seller or business partner. Together they will review the performance, storage, throughput, resiliency, and encryption requirements for your solution along with the size, retention, and change rate of your data to suggest the necessary scale out cluster architecture needed to meet your enterprise data resiliency goals.

Where to use IBM Storage Ready Nodes with IBM Storage Defender

Early detection of internal and external threats
IBM Storage Defender is designed to leverage sensors across primary and secondary workloads to detect threats and anomalies from backup metadata, array snapshots and other relevant threat indicators.

Recover acceleration for business operations
While IBM Storage Defender seeks to eliminate the threats to your data, there are a myriad of ways your data could still be impacted, such as through user or administrator error.

Integration with your existing SecOps tools and processes
IBM Storage Defender can integrate all the comprehensive protection capabilities for critical data with overall security tools like SIEM and SOAR security solutions that are likely already in use in your organization.

For more information on IBM Storage Defender use cases read the solutions brief: https://www.ibm.com/downloads/cas/AWDK6VME

For more information on IBM Storage Defender visit: https://www.ibm.com/products/storage-defende

Ready Node Options for IBM Storage Ceph

IBM Storage Ceph is an IBM-supported distribution of the open-source Ceph platform that provides massively scalable object, block, and file storage in a single system. It is part of the IBM Storage portfolio of software-defined storage. IBM Storage Ceph is designed to operationalize AI with enterprise resiliency and consolidate data with software simplicity and run on multiple hardware platforms to provide flexibility and lower costs. Engineered to be self-healing and self-managing with no single point of failure and includes storage analytics for critical insights into growing amounts of data. IBM Storage Ceph can be used as an easy and efficient way to build a data lakehouse for IBM watsonx.data and for next-generation AI workloads. For more information on IBM Storage Ceph software read the data sheet ->

Storage Ceph with Storage Ready Node Specifications

	Storage Ceph Node	Storage Ceph NVMe Flash Node
Processor	Intel Xeon Silver 4314	Intel Xeon Gold 6438N
# of Processors	2	2
RAM	256GB	512GB
OS Disks	2 x 240GB Raid 1 SSD	2 x 480GB RAID 1 SSD
Data Acceleration Disks	2x3.84TB SSD	N/A
Rack Height	2U	2U
Width	482 mm (18.97 in.)	482 mm (18.97 in.)
Depth	772.11 mm (30.39 in.)	772.11 mm (30.39 in.)
Height	86.8 mm (3.41 in.)	86.8 mm (3.41 in.)
Weight	35.3kg (77.82 lb) max	35.3kg (77.82 lb) max
# of Capacity Disks	12	24
NVMe Disk Sizes		3.84TB, 7.68TB, 15.36TB
HDD Disk Sizes	8TB, 12TB, 16TB, 20TB	
Network		
2x1GbE	X	X
2x10GbE	X	X
2x 100GbE		X

Storage Ceph Node

Storage Ceph NVMe Flash Node

Configuring IBM Storage Ready Nodes with IBM Storage Ceph

To create an IBM Storage Ceph cluster its easy to start with IBM Storage Ready Nodes and scale one node at a time or a set of nodes depending on the availability and capacity efficiency required. The initial configuration can be as small as 4 nodes with 3 data copies or Erasure Code 2+2 or 7+ nodes for a more efficient capacity requirements. For example the smallest configuration and least number of nodes would consist of 4 nodes with 3X redundancy or EC 2+2 and this yields no downtime or data loss if a node become unavailable and also yields the fastest performance for a configuration. For max capacity with highest efficiency with no downtime or data loss if 2 nodes are unavailable the minimal configurationwould include 7 nodes. This 7 node configuration using 4TB disks will yield 512TB of usuable storage and is an excellent choice for deployment of IBM watsonx.data which includes 768TB raw capacity of capacity. No matter what size theinitial deployment capacity or number of storage nodes, one can add more capacity at any time by adding nodes and increasing capacity.

When capacity is expanded the system usually also experiences a performance increase and all expansions can de done without any production downtime to users accessing the IBM Storager Ceph cluster. Once the cluster is set up and configured and object buckets are created the system is mainly self managed and very easy to maintain. For more information on configurations read the documentation or read read the blog ->

Where to use IBM Storage Ready Nodes with IBM Storage Ceph

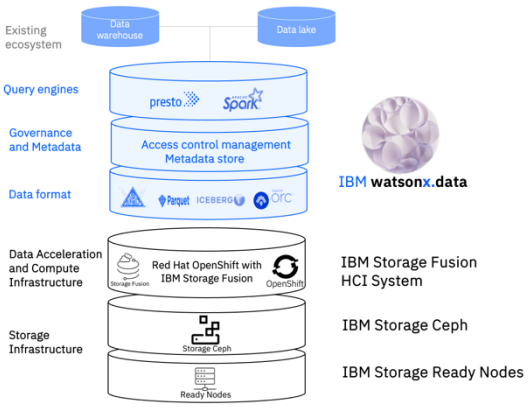

Storage Infrastructure with IBM watsonx.data

Watsonx.data makes it possible for enterprises to scale AI workloads using all their data with a fit-for-purpose data lakehouse architecture optimized for governed data and AI workloads, supported by querying, governance, and open data formats to access and share data. This is based on open-source technologies, including Presto and Iceberg. IBM Storage Ceph provides the storage infrastructure for a watsonx.data on-premises deployment. IBM watsonx.data includes 768TB of IBM Storage Ceph software license and support. The easiest way to start with IBM Storage Ready Nodes is with 4 nodes or 7 nodes depending if you desire perfromance or storage efficiency. The 4 node configuration provides the least number of nodes and the fastest performance with the ability to lose a node without incident. The 7 node configuration provides more nodes nodes and yields greater storage efficiency with the ability to loose 2 nodes without incident and the configuration can scale one or more nodes at a time as capacity needs grow. With watsonx.data and IBM Storage Ceph, you can access all your data across both databases and data lakes as each configuration can be optimized in the same Storage Ceph cluster. Share large volumes of data through open table formats, such as Apache Iceberg, built for high performance analytics and large-scale data processing and at the same time store large amount of data for other large data set analysis. IBM Storage Ceph also supports multiple vendor open formats for analytic data sets while allowing different engines to access and share the same data, at the same time using tools like Parquet, Avro, Apache Orc and more.

Object storage as-a-service and cloud native applications

IBM Storage Ceph with IBM Storage Ready Nodes is ideal for implementing an object storage service, with proven scalability and performance for both small and large object storage. With the open standards approach to object storage and the S3 interface, new and existing applications are easy to integrate with support for the latest S3 APIs and capabilities. Applications can access their storage with the same application API, in public, private, or hybrid clouds.

Ready Node Options for Cloud Object Storage

IBM Cloud® Object Storage is a market-leading object storage solution for primary and secondary AI and big data workloads. Our solution is grounded in Dispersed Storage™ and a flexible Information Dispersal Algorithm (IDA) an is proven for AI and big data solutions as well as offloading traditional storage workloads to object storage. Storage solutions can start with only three Slicestors and grow online to thousands of Slicestors or exabytes of data. Our software supports mixed configurations with IBM Storage Ready Nodes and non Ready Node hardware and can scale or expand with no forklift upgrades or downtime enabling investment protectionand long-term cost savings. Applications can access storage from any location and optimize perfromance using geo-dispersed protection and efficiency and concurrent and secure cloud native accessibility. For more information on IBM Cloud Object Storage software. read the data sheet ->

Cloud Object Storage with Storage Ready Node Specifications

	Slicestor® NVMe Flash	Slicestor®	Slicestor® Capacity	Accesser®	Manager
Processor	Intel® Xeon® Gold 6438N	Intel Xeon Silver 4314	Intel Xeon Silver 4314	Intel Xeon Silver 4314	Intel Xeon Silver 4314
# of Processors	2	2	2	2	2
RAM	512GB	256GB	256GB	256GB	256GB
OS Disks	2x480GB SSD	2x240GB SSD	2x960GB SSD	2x960GB SSD	2x960GB SSD
Rack Height	2U	2U	6U (1U+5U)	1U	1U
Width	482 mm (18.97 in.)	482 mm (18.97 in.)	482 mm (18.97 in.)	482 mm (18.97 in.)	482 mm (18.97 in.)
Depth	772.11 mm (30.39 in.)	772.11mm (30.39 in.)	772.11 mm (30.39 in.)	772.11mm (30.39 in.)	772.11 mm (30.39 in.)
Height	86.8 mm (3.41 in.)	86.8 mm (3.41 in.)	265.4 mm (10.45 in.)	42.8 mm (1.7 in.)	42.8 mm (1.7 in.)
# of Capacity Disks	8, 16, 24	12	28, 56, 84		
NVMe Disk Sizes	3.84TB,7.68TB,15.68TB				
HDD Disk Sizes		8TB, 12TB, 16TB, 20TB	8TB, 12TB, 16TB, 20TB		
Network					
2x1GbE	X	X	X	X	X
2x10GbE	X	X	X	X	
2x25GbE	X				
2x100GbE	X				

Slicestor NVMe Flash Node

Slicestor Node

Accesser and Manager Node

Slicestor Capacity Node

Configuring IBM Storage Ready Nodes with IBM Cloud Object Storage

To create an IBM Cloud Object Storage cluster that can scale to petabytes or even exabytes of data online start with a Manager and then add 3 or more Slicestors and optionally 2 or more Accessers for added performance. There are two different Slicestors depending on capacity requirements and IBM provides added flexibilty with unique drive sizes that enable customers to create the most cost efficeint solution based on current capacity and expected capacity growth requirements. Once the system is set up and configured and object bukets are created the system is almost self managed and very easy to maintain. For more information and help to configure each of the IBM Storage Ready Nodes for IBM Cloud Object Storage you can reference the application configuration document ->

Where to use IBM Storage Ready Nodes with IBM Cloud Object Storage

Disaster Recovery and Backup: IBM Cloud Object Storage has validated many of the leading backup and archive applications and IBM Storage Ready Nodes provide an optimized platform for this use case. Replace exisiting backup process with IBM Cloud Object Storage which offers high resiliancy and low cost for infrequent-access storage. IBM Cloud Object Storage gives you cloud-native capabilities for an automated, application-consistent backup and recovery solutions. Cut your storage costs and reduce malware, disaster and outage losses.

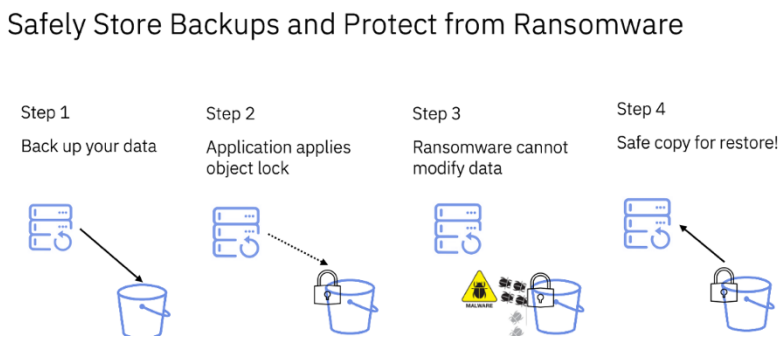

For a more comprehensive list of use cases and solutions:
IBM Cloud Object Storage Solutions Guide

For more information on IBM Cloud Object Storage visit:
https://www.ibm.com/products/cloud-object-storage

Conclusion

As customers look to lower the cost of their data center storage and at the same time move to a more cloud like model, IBM Storage Ready Nodes are aimed at closing the gap between software-only cloud scale storage systems and legacy turnkey storage. IBM is now making it easier to build, deploy and expand software defined storage solutions in the data center. These flexible cloud scale building blocks based on commodity components provide optimized and flexible configurations that are validated for IBM Storage software, and all supported by IBM. Storage Ready Nodes offer a pre-configured, validated, and flexible solution for organizations looking to create a software-defined infrastructure on commodity servers but also simplify the management, deployment and support of the storage solution in the data center. IBM provides a simplified deployment process, optimized performance options, on-line scalability from a base configuration to 10s, 100s or 1000s of nodes, and compatibility with other IBM storage software deployments all backed by IBM's support and services.

Why IBM?

Data matters. When planning a data strategy for new or existing applications it's easy to focus on compute resources and applications without proper planning for the data that will drive the results for the applications. Our products are all about solving hard problems faster with data. IBM helps customers achieve business value with a clear data strategy. Our strategy is simple, unlock data to speed innovation, de risk data to bring business resilience and help customers adopt value based data to bring cost and energy efficiencies. Value needs to be delivered by connecting the multiple organizational data sources with business drivers to create business value that mean something to the organization. Many organizations focus on a single driver with a storage solution, but the best solution is driven by an infrastructure strategy than can accomplish most if not all the drivers for maximum benefits. Our story is not just about another storage product but is about innovation and a comprehensive storage portfolio that is helping businesses drive more value throughout the organization.

For more information

To learn more about IBM Storage Ready Nodes with any of the IBM Storage software offerings please contact your IBM representative or IBM Business Partner, or visit any of our web pages and ask to chat with a representative:

IBM Storage Defender Web Page

IBM Storage Ceph Web Page

IBM Cloud Object Storage Web Page

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the
United States of America
February 2022