


IBM Storage Fusion Product Guide

Jon Herd
Daniel Danner
Axel Koester



 Cloud

Storage



IBM Redbooks

IBM Storage Fusion Product Guide

June 2023

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (June 2023)

This edition applies to Version 2, Release 5, Modification 2 of IBM Storage Fusion HCI System and IBM Storage Fusion.

© Copyright International Business Machines Corporation 2023. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	v
Trademarks	vi
Preface	vii
Authors	viii
Now you can become a published author, too!	ix
Comments welcome	ix
Stay connected to IBM Redbooks	ix
Summary of changes	xi
June 2023, First Edition (minor update)	xi
Chapter 1. Why containers	1
1.1 Why move to containers when VMs seem just fine.	2
1.2 Faster time to market in IT	3
1.3 Infrastructure as Code and Disaster Resiliency	5
1.4 Cyber Security and Cyber Resiliency	5
1.5 Zero-trust	6
1.6 Workload Scalability	7
1.7 Deconstruction silos with Data Bridging	8
Chapter 2. Why IBM Storage Fusion	11
2.1 Why IBM Storage Fusion	12
2.2 Why Red Hat OpenShift on Bare Metal	12
2.3 Mission Critical Workloads in Containers	13
2.4 Getting Workloads into Production faster	14
2.5 Storage for Containers	14
2.6 Container Platform Resiliency and High availability	16
2.7 Data Protection for Containers	18
2.8 Red Hat OpenShift made easy: Day2 Operations with IBM Storage Fusion	20
Chapter 3. Flavors of IBM Storage Fusion	25
3.1 Integration with legacy VM infrastructures	26
3.2 Three Tier Architecture System	26
3.3 Hyper-Converged Infrastructure (HCI)	27
3.4 IBM Storage Fusion Options	28
3.4.1 IBM Storage Fusion Software Overview	28
3.4.2 IBM Storage Fusion HCI System Appliance Overview	29
3.4.3 IBM Storage Fusion HCI System Software Overview	30
3.4.4 IBM Storage Fusion HCI System Hardware Overview	31
3.4.5 Scalability in HCI	33
3.4.6 Configuring the IBM Storage Fusion HCI System with IBM Storage Modeller ..	34
3.4.7 VMware versus bare metal	35
3.4.8 Use cases for IBM Storage Fusion software	37
Chapter 4. IBM Storage Fusion Terminology	39
4.1 Operators	40
4.2 Upgrading	40
4.3 Management UI	40

4.4 IBM Storage Scale Erasure Code Edition (ECE)	40
4.5 Red Hat OpenShift Virtualization	41
4.6 IBM Storage Fusion Data Cataloging service	41
4.7 IBM Cloud Satellite	42
4.8 Red Hat Advanced Cluster Management (RHACM)	42
4.9 IBM Storage Protect Plus	43
4.10 Call Home for all IBM Storage Fusion HCI System.	43
4.11 Active File Management (AFM).	43
4.12 AI Support (AI GPU)	44
4.13 Synchronous Storage Replication with Metro-DR.	45
4.14 IBM Storage Fusion Reference Information and Links	46
4.14.1 New features and changes in 2.5.2	46
4.14.2 New features and changes in 2.5.1	47
4.14.3 New features and changes in 2.4	48
4.14.4 New features and changes in 2.3	49
4.14.5 New features and changes in 2.2.1	50
4.14.6 New features and changes in 2.2	50
4.14.7 New features and changes in 2.1.2	50
4.14.8 IBM Storage Fusion Useful Links	51
Related publications	53
IBM Redbooks	53
Online resources	53
Help from IBM	54

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, 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 <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

IBM®	IBM Spectrum®	Redbooks (logo)  ®
IBM Cloud®	IBM Spectrum Fusion™	Satellite™
IBM Cloud Pak®	Passport Advantage®	Spectrum Fusion™
IBM Cloud Satellite®	Redbooks®	

The following terms are trademarks of other companies:

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

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

Red Hat, Ansible, Ceph, OpenShift, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

VMware, and the VMware logo are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Other company, product, or service names may be trademarks or service marks of others.

Preface

This IBM® Redbooks® publication offers a short overview of IBM's integrated environment for container workloads, IBM Storage Fusion. The product comes in two variants, IBM Storage Fusion HCI System including all required hardware, and IBM Storage Fusion SDS (software-only) for deployment in public or private Clouds.

Note: This Redpaper has been updated with the 2.5.1 release only for IBM Storage Fusion and the 2.5.2 release of IBM Storage Fusion HCI System.

IBM Storage Fusion 2.5.2 includes IBM Storage Fusion Data Foundation 4.12. Fusion Data Foundation 4.12 and Red Hat OpenShift Data Foundation 4.12 both provide the same data service, and OpenShift Data Foundation is the Red Hat brand of the IBM Storage Fusion Data Foundation offering. Fusion Data Foundation and OpenShift Data Foundation functions are the same: the only difference is the digital signature. References and links to OpenShift Data Foundation remain valid. Fusion Data Foundation 4.12 documentation is integrated within IBM Storage Fusion 2.5.2 documentation.

Note: IBM Spectrum® Fusion HCI and IBM Spectrum Fusion™ have become IBM Storage Fusion HCI System and IBM Storage Fusion. This Redpaper edition uses the new IBM Storage brand names introduced in IBM Storage Fusion 2.5.1.

IBM Spectrum Discover functionality is now available as an integrated service within IBM Storage Fusion called Data Cataloging.

See [Evolving the IBM Storage Portfolio Brand Identity and Strategy](#) to learn more about how IBM Storage Fusion HCI System and IBM Storage Fusion are key to the IBM Storage Portfolio.

Chapter 1, “Why containers” on page 1 answers general WHY questions: Why are more and more workloads deployed as containers? Why are containers beneficial for cyber security? Why should I care about new development and deployment models involving fewer staff? Why should I care about containers when VMs seem just fine?

Chapter 2, “Why IBM Storage Fusion” on page 11 focusses on WHY IBM Storage Fusion: What are the benefits of running containers in a turnkey Red Hat OpenShift cluster within IBM Storage Fusion? How can the integrated environment help reduce the effort required for deployment, high availability design, backup implementation and cyber resilience planning?

Chapter 3, “Flavors of IBM Storage Fusion” on page 25 describes flavors of IBM Storage Fusion implementations, major differences between HCI System (hardware) and SDS (software) variants, and which choices to make to right-size the product for current needs while retaining the ability to scale out at a future stage.

Chapter 4, “IBM Storage Fusion Terminology” on page 39 describes the terminology used for containers environments and more specifically for IBM Storage Fusion. It also covers some more advanced functions included in the product.

Authors

This paper was produced by a team of specialists from around the world working with the IBM Redbooks, Tucson Center.

Jon Herd is a Senior Executive Advocate working for the IBM Technical Lifecycle (TLS) EMEA Client Care team in Germany. He covers the United Kingdom, Ireland, and Europe, advising customers on the portfolio of IBM storage products, including the IBM Storage Fusion range. Jon has been with IBM for more than 45 years, and has held various technical roles, including Europe, Middle East, and Africa (EMEA) level support on mainframe servers and technical education development. He has written many IBM Redbooks publications about IBM Storage and FlashSystem products and is an IBM Redbooks Platinum level author. He holds IBM certifications in IBM Product Services L3 Thought Leader level, and IBM Technical Specialist L1 experienced level. He is also a certified Chartered Member of the British Computer Society (MBCS - CITP), a Certified Member of the Institution of Engineering and Technology (MIET) and a Certified Technical Specialist of the Open Group.

Daniel Danner is the Lead Architect for IBM Storage Fusion in the D/A/CH region. Before he joined IBM, he gathered 20 years of experiences working in enterprise IT environments dealing with mission critical workloads for a large IBM account. Daniel is recognized for his ability to combine hands on experience, customer view and ability to explain complex circumstances in a consumable manner. Besides his work on IBM Storage Fusion, he leads the IBM internal Ansible Automation Community.

Axel Koester Axel Koester, PhD, is Chief Technologist for the EMEA Storage Competence Center ESCC in Frankfurt, Germany. Besides caring about all most recent storage technologies and beyond, he frequently speaks at public conferences on technology outlook topics and their future implications for IT businesses. Currently he also devotes time to improving Storage performance models for simulation/prediction tools, including models for IBM Storage Fusion HCI System. Axel is a current member of the board and former chairman of the IBM TEC Think Tank in D/A/CH. Prior to joining IBM in 2000, Axel helped build the Large Hadron Collider project at CERN, Geneva, and worked as researcher for the Fraunhofer Institutes, after having worked as an object-oriented programming pioneer in the privately held company Gesytec mbH. Axel holds a PhD in electrical engineering and computer science from the RWTH Aachen University.

Thanks to the following people for their contributions to this project:

Larry Coyne
IBM Redbooks, Tucson Center

Josh Blumert
IBM Global Sales

Chrisilia Bristow, Matt Divito, Lee Helgeson, Khanh Ngo, Tom O'Brien, Fernando de la Tejera Ruiz, David Wright
IBM Systems

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at: ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks publications in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, IBM Redbooks
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

<http://www.redbooks.ibm.com/rss.html>

Summary of changes

This section describes the technical changes made in this edition of the paper and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes
for IBM Storage Fusion Product Guide
as created or updated on June 22, 2023.

June 2023, First Edition (minor update)

This revision includes the following changed information.

Changed information

- ▶ Updated branding information for IBM Storage Fusion Data Foundation (FDF) and Red Hat OpenShift Data Foundation (ODF). See the IBM Storage Fusion Data Foundation section in “Scaling Storage for Applications” on page 7 for additional information.
- ▶ Updated storage product names in Figure 3-4 on page 30.



Why containers

This chapter answers the general WHY questions: Why are more and more workloads deployed as containers? Why are containers beneficial for cyber security? Why should I care about new development and deployment models involving fewer staff? Why should I care about containers when VMs seem just fine?

1.1 Why move to containers when VMs seem just fine

“I’m too busy chopping wood to sharpen the axe.” says the lumberjack when he is overwhelmed with workload. It’s a common situation in IT as well: Security concepts turn out to be less secure than expected; disaster protection doesn’t protect enough from new types of disaster that were just discovered. Most of this can be traced back to a lack of time spent on these challenges, while too much time is spent on “mundane” tasks that *could* be automated if they weren’t so diverse.

The metaphor of “pets vs. cattle” (commonly attributed to Bill Baker, then of Microsoft) is well known in that context: To make automation easier and automatic procedures less diverse, do **not** treat computers as individual devices, do **not** give them individual names, do **not** manage their IP addresses, do **not** manually assign workloads, and do **not** make machines “opaque” just to isolate dependencies.

Do you think that virtualization achieves that? Think again. Virtualization improves the efficient use of resources, confines dependencies, and makes workloads easily movable. But: Each VM is effectively another *individual* machine, and their count keeps growing. Each VM demands (at least some) individual care and effort. Hordes of pets (Figure 1-1).



Figure 1-1 Pets versus Cattle¹

Can we skip the concept of *machines* in IT altogether? Yes: This is sometimes referred to as serverless operations, even though there is still real hardware involved. Managing “cattle” means more than giving these machines numbered names: The idea of containerization is to get rid of the machine concept as individual entity. Individual worker nodes or containers do not need personal attention. Their instantiation, their workload assignment, their termination, their whole lifecycle is fully automated and driven by short-term demand.

Likewise, security in individual pods or containers won’t need personal attention. This is handled at a higher abstraction level. And nobody looks after disaster resilience in individual containers. This, too, is usually handled at a higher level.

Eliminating the effort spent on monitoring and managing individual entities leaves us more time to “sharpen the axe”. The goal of Red Hat OpenShift inside IBM Storage Fusion is not merely to ease the management of individual containers, but to actually **make us forget that they exist**. In contrast, hypervisor environments usually provide tools to manage many pets - but each of them requires individual care. See Figure 1-2 on page 3¹.

¹ images: creative commons CC-0




Figure 1-2 Managing individual entities versus serverless operations

So, are VMs totally eliminated in a container environment? No: VMs can be very useful to set up prototypes, to run quick deployment tests, or to add flexibility to control plane nodes. But for the workload-carrying production worker nodes, VM constructs are obsolete. IBM Storage Fusion is not built on top of hypervisor technology, even though it has the capability to operate VMs alongside with containers. Refer to 3.4.7, “VMware versus bare metal” on page 35.

1.2 Faster time to market in IT

With modern continuous delivery methodologies, software release cycles have become very short. To keep up with that pace, an easier transition from dev-to-test and from test-to-production is compulsory. Software development itself has also transitioned from formerly mostly “in-house” to nowadays mostly “re-use” of open-source libraries and frameworks.

A real-world case: Medical support platform in 3 weeks

During the Coronavirus pandemic outbreak, medical support projects had to be implemented urgently, with no ready-to-deploy commercial offers available: A newly built Coronavirus test center in Mainz, Germany, asked IBM to design and build a tailored IT solution for secure appointment logistics and communication among the medical staff. Several thousand patients were expected within the first days of operation. Leveraging container technology and prebuilt service templates from the IBM Cloud®, the hired IBM consulting team was able to roll out a minimum viable product on-site within weeks, including hardware and network installation. The final solution grew to a containerized services platform with appointment workflow and time slot management, handling of digital certificates for testing and vaccination, a patient survey, a patient chat bot, and a secure messaging app for healthcare workers, all of this realized by porting code developed on (or for) the IBM Cloud.

The container advantage

In the container world, re-use of preconfigured services is the norm. Containers allow fast time-to-market due to higher levels of standardization, and easier rollbacks in case of glitches. Separate environments for dev/test, pilot and production are not obligatory, as the isolation between instances is handled at an abstract level. Once ready, select parts of the staging pilot can be declared as production instance - without redeployment, and with a quick undo option. Moreover, partial updates can be tested against a limited set of end users first:

Let's increase user acceptance by testing modifications at scale (not just at home) with an increasing number of users, by controlling which range of IP addresses sees which combination of updates.

We are so much used to seamless app updates on our mobile devices that we keep asking ourselves: What's holding back larger IT solutions to do the same? The answer is often “weak dependency management” and “monolithic architecture”. Monoliths, usually large conglomerates of compiled software, can only be rolled out in one piece. All-or-nothing.

In contrast, container-based applications are split into many functional parts (sometimes down to *microservices*) bound together at runtime via a dependency description, the recipe. The recipe elements control which version of which component is part of a production instance, a progressive deployment instance, beta test instances, and developer instances. All of these can coexist in the same environment, with containers providing the sandbox technology to keep data secure.

IBM Cloud Paks

IBM Cloud Paks are prime examples of previously monolithic software that was split into functional containers and then bundled so that each service can serve multiple purposes. When configuring Cloud Paks, a dependency checker tells us which services are missing and adds them to the basket. The more software we add to the shopping list, the smaller the increments become, as most required services are already included. This is no different from the way Linux and Unix package managers operate, albeit at a larger scale. We're trading individual installations for a framework that controls all dependencies and thus significantly speeds up all future deployments.

Swift deployment of IBM Cloud Paks on IBM Storage Fusion

When planning IBM Storage Fusion for IBM Cloud Pak® workloads, we recommend estimating the container workload in the appropriate tool, the *Cloud Pak sales configurator*, and within minutes obtain an approximation of the required hardware. Use this estimate to lay out a right-sized IBM Storage Fusion configuration using the Storage Modeller (StorM) tool. See the example hardware estimation in Figure 1-3.

<https://app.ibm.com/sales/configurator> (requires an IBM account and a public IBMid)

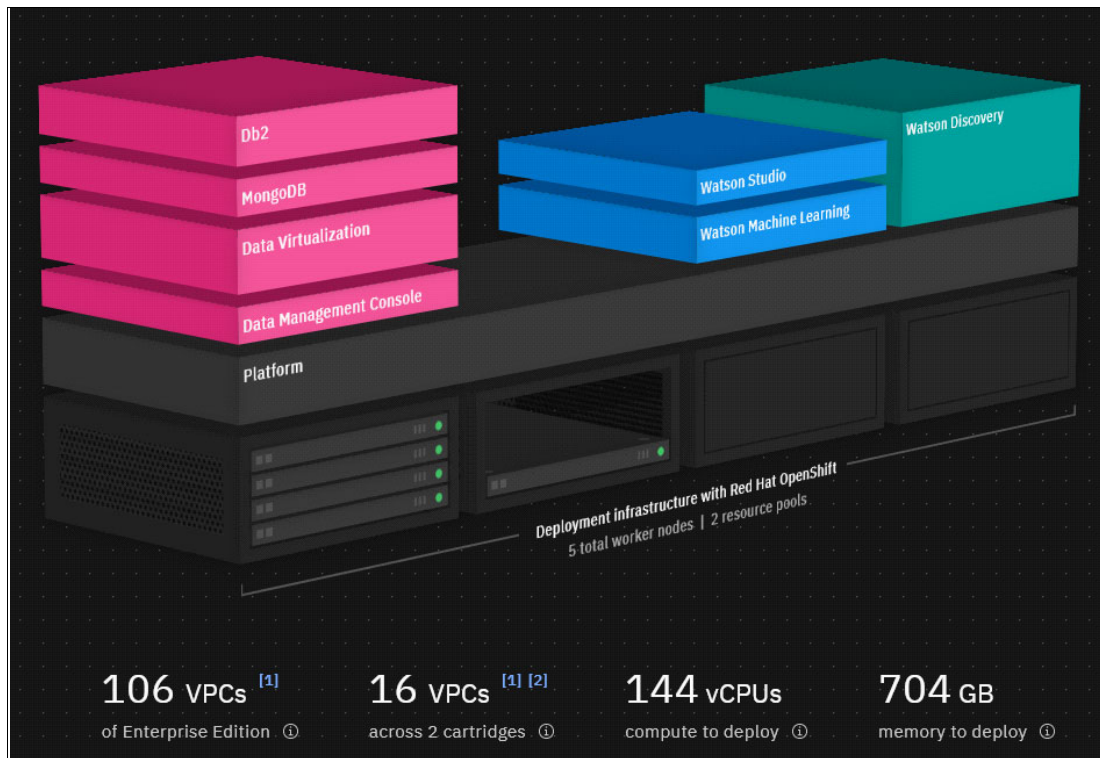


Figure 1-3 Hardware estimation for a planned Cloud Pak workload (from [ibmsalesconfigurator.com](https://app.ibm.com/sales/configurator))

Details on configuring IBM Storage Fusion in the StorM tool can be found at 3.4.6, “Configuring the IBM Storage Fusion HCI System with IBM Storage Modeller” on page 34.

1.3 Infrastructure as Code and Disaster Resiliency

Hardware and network requirements for an application suite, including access to local storage and compute resources, also constitute dependency chains. Therefore, we can add them to our “recipe”. The recipe becomes the centerpiece of all containerized application suites, telling the automation how to select hardware candidates, configure the network, apply mandatory operating system patches, fetch required services templates from repositories, and start user-interfacing services in the right order. The recipe is documentation and automation script in one.

The concept of “describing all in one place” is beneficial when it comes to

- ▶ scaling an application (*which services do I need to scale in the backend? On which flavor of hardware?*)
- ▶ restoring a whole environment after an attack or a disaster (*can I check the integrity of all components? In which order should I restart everything?*)

In a containerized environment, these two subjects are barely distinguishable.

Building Infrastructure with Developer Tools

Hardware selection instructions in our recipe can be managed in a **git repository** - with branching, merging, and roll-back options. The ultimate goal here is “infrastructure as code”, where hardware selection and network configuration are written by DevOps programmers, with access to efficient developer pipeline tools. Hardware rollbacks in this context are trivial.

The deployment process becomes increasingly hardware agnostic or “Cloud-inspired”, making it moveable between public and private clouds, managed clouds services, and home-cloud locations like IBM Storage Fusion.

1.4 Cyber Security and Cyber Resiliency

Cyber security can be an inherent part of a recipe-based deployment: Just like Linux packages that are thoroughly checked for their integrity before being considered for deployment, complex application components (and even hardware) can be checked for integrity before going live. To avoid upstream supply chain attacks, fingerprinting and source authentication are mandatory from the very source of each piece of code.

Apache Log4j lessons learned

Even authenticated and well-known code can have zero-day vulnerabilities that eventually lead to exploits. In the case of Log4j (a core component in the logging infrastructure of Java apps) the vulnerability was tied to remote code execution in an unintended way. Patches became available shortly after this vulnerability was made public, yet many computers remained unpatched for weeks or even months.

The problem here was weak dependency management and slow update policies: In order to know which systems/instances need the patch, their reliance on Log4j must be documented. While easily discoverable in the Log4j case, dependencies will be harder to detect in others - especially when documentation quality depends on human meticulousness. In a container environment, this dependency is coded as part of the construction recipe of our application suite and therefore always up to date.

How containerization improves Cyber Protection

Cyber security can be vastly improved by eliminating the human factor in dependency management. In a containerized environment, the build process of any containerized application suite (using the *recipe*) yields a complete documentation: What has been built, is documented - and that documentation can also cover server and network configurations, all in one place. This makes automatic cyber threat scanning and the discovery of indicators of compromise much easier and streamlined.

Cyber resiliency on the other hand can be improved by automating a recovery process and testing it regularly, e.g. an application restart in a failover location that has access to a copy of the most recent data. *Testing it regularly* is the limiting element here, comparable to sharpening the axe when you should be chopping wood.

Containerized environments simplify frequent recovery testing by design: There is no difference between restarting and relocating a container, or scaling up a containerized service by adding instances, which may happen during automatic workload optimization.

Infrastructure-as-code simplifies the deployment in another location with a compatible control plane and access to the original data. Tools for mirroring or replicating the user data, plus the ability to restart an application suite elsewhere, are integral parts of IBM Storage Fusion. Refer to 4.8, “Red Hat Advanced Cluster Management (RHACM)” on page 42 and 4.13, “Synchronous Storage Replication with Metro-DR” on page 45.

1.5 Zero-trust

On top of typical security procedures, zero-trust means that no “truth” is taken for granted without proof, and no proof may exceed a certain age.

Since distributed applications can include many services, it makes sense to design a mutual authentication scheme, especially so when some of these services are externalized. Between authenticated services, communication can be encrypted in a way that only the intended receiver can decode the content. See Figure 1-4.

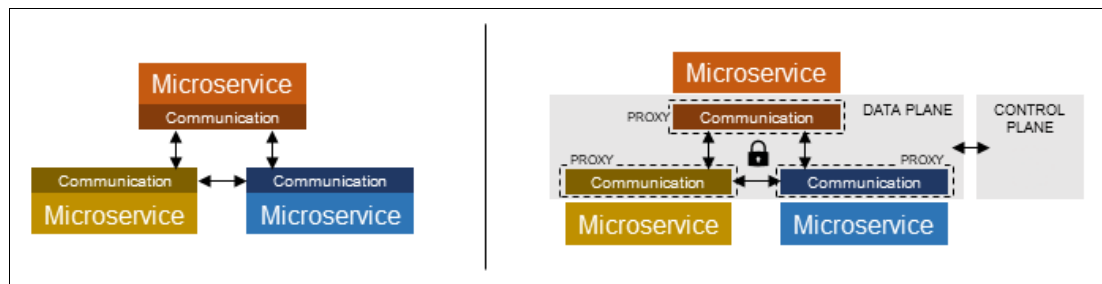


Figure 1-4 Microservice mesh with direct P2P communication (left) and with “zero trust” mutual authentication (right)

Zero-trust for dynamic environments

An example for a service that scales automatically by adding instances on demand is the user management component of a streaming platform. To avoid man-in-the-middle attacks, all parts in this multi-instance communication cosmos must individually authenticate - none is accepted as trustworthy per se. Container platforms can enforce TLS-based secure communication between all components of a containerized application suite, making eaves-dropping and spoofing much harder.

Beyond reducing the risk of attacks, the flow of data and commands can be monitored and analyzed in machine learning systems, which can throw alerts upon sensing irregularities. These could be components that never communicated in the past, suspicious service access patterns, suspicious denials, any others.

1.6 Workload Scalability

A prerequisite for automatic workload scaling is that services must be multi-instance capable. In complex applications, only those components that consume more CPU than allowed (or generate more IOs than allowed) need to be scaled. A typical scaling rule could be:

1. Add instance when active instance has >50% CPU utilization.
2. Destroy instance when it shows <10% average CPU utilization and there is more than one instance left.
3. Start one instance when there is none.

In fact, this automation is not scripted as a sequence of actions, but as a series of outcomes: An application architect could define that “no instance should continually consume more than 50% or less than 10% CPU” and the workload scheduler component underneath IBM Storage Fusion (Red Hat OpenShift / Kubernetes) makes it happen, preferring the least utilized hardware for spreading the workload. Additional rules for throughput or latency optimization are possible. See Table 1-1.

Table 1-1 Classic versus containerized performance management approach

Classic performance management approach	Containerized performance management approach
When an application exceeds the resources of a VM, manually increase the assigned resources if appropriate.	When a container exceeds 50% CPU utilization (configurable), the platform starts an additional instance of it.
When all VMs on a system exceed that system's capabilities, manually migrate the least important VMs onto new hosts.	When an instance consumes less than 10% (configurable), the platform destroys it unless it is the last instance.

Programmers of multi-instance services do not need to re-invent proprietary service APIs and protocols susceptible to new security flaws: Open-Source solutions exist for every flavor of workload distribution. Most popular are publisher/subscriber schemas like those proposed in Apache Kafka, a core component in many scalable apps.

Scaling Storage for Applications

IBM Storage Fusion is leveraging the Red Hat OpenShift Storage class concept that differentiates between *ephemeral* local data and *persistent* data that survives container restarts. Therefore, persistent data is always kept on shared storage. Scalability for shared storage means:

- ▶ enough physical storage to fulfill growing capacity and throughput requirements, plus the ability to increase in-rack physical capacity on the fly.
- ▶ an automatic increase of capacity assigned to an application (within configurable limits) and a dynamic sizing of associated services like caching, mirroring, redundancy monitoring, or access control.

Depending on the flavor of distributed storage implemented, associated procedures will look slightly different, even though IBM Storage Fusion will hide most of this to the administrator and make it invisible for a code developer, who will only have to choose between persistent storage classes offered.

The two supported storage options are, at the time of writing:

IBM Storage Fusion Data Foundation: IBM Storage Fusion 2.5.2 includes IBM Storage Fusion Data Foundation 4.12. Fusion Data Foundation 4.12 and Red Hat OpenShift Data Foundation 4.12 both provide the same data service, and OpenShift Data Foundation is the Red Hat brand of the IBM Storage Fusion Data Foundation offering. Fusion Data Foundation and OpenShift Data Foundation functions are the same: the only difference is the digital signature. References and links to OpenShift Data Foundation remain valid.

Fusion Data Foundation 4.12 documentation is integrated within IBM Storage Fusion 2.5.2 documentation. IBM Storage Fusion Data Foundation is based on the open-source projects Ceph, Rook, and Noobaa, and provides a variety of access protocols while being able to use a wide range of physical storage sources. Redundancy is achieved by storing three copies, i.e. the erasure coding feature of Ceph is not supported by IBM Storage Fusion Data Foundation.

ECE: Erasure Coding Edition of IBM Storage Scale, delivered fully integrated with IBM Storage Fusion HCI System hyperconverged infrastructure. Local node SSDs are combined into a dispersed, “write-many” persistent storage layer. In addition, IBM Storage Scale ECE supports multiple access protocols, advanced automated file management (AFM), rule-based data placement and relocation, geographically dispersed caching, metro replication, and much more. IBM Storage Scale aims at bottleneck-free parallel scalability by design. Data redundancy is achieved by erasure coding with high 64% usable-to-raw storage capacity efficiency when protecting against double failures. This is approximately 2x the capacity efficiency which can be achieved with three copies of data (33%).

Scaling Backup

IBM Storage Fusion provides a platform-independent backup based on IBM Storage Protect Plus, where data can be protected on air-gapped storage services and/or devices outside of the production cluster. This makes sure that data which is being restored after a major incident is not dependent on the original source of the data, as that may have been compromised during the incident.

While all backup components inside IBM Storage Fusion scale on demand, externalized components (like moving copies onto an air-gapped tape library or object storage) must be managed and scaled individually, since Red Hat OpenShift is not in charge. On the other hand, previously existing compatible backup environments and procedures can be re-used or integrated.

1.7 Deconstruction silos with Data Bridging

Data silos limit the efficiency of many implementations, be it for consistent automation, security, or shared resource usage. Choosing ECE and therefore IBM Storage Scale as a storage layer in IBM Storage Fusion HCI System helps break up silos and facilitates a single logical view of data across platforms, rather than having to deal with “data tourism” (channeling data between various services and apps, usually generating lots of copies).

IBM Storage Scale can:

1. make user data persistent, i.e. avoid deletion when a service instance is terminated.

2. make user data cluster-wide persistent, so that containers can be restarted on *any* node at equal cost, with immediate access to their persistent data.
3. harden user data by dispersing it over several storage-rich nodes with space-efficient erasure coding, as explained previously in “Scaling Storage for Applications” (ECE: Erasure Coding Edition).
4. make user data persistent across clusters through metro mirroring, as described in 4.8, “Red Hat Advanced Cluster Management (RHACM)” on page 42 and 4.13, “Synchronous Storage Replication with Metro-DR” on page 45.
5. allow asynchronous access to data through other protocols like Object S3 / NFS / etc., either for data ingest or for use in legacy applications.
6. provide data protection connectors to the backup component IBM Storage Protect Plus.
7. avoid “cabled security” (solely based on hardware/network isolation) while providing excellent protection, which helps prevent silos.

Companies using IBM Storage Scale to its full extent typically refer to it as an *enterprise data highway, data bus, or hybrid data bridge*. It shortens time-to-market by removing the need to implement complex and custom data sharing mechanisms between new and existing data sources. Instead, IBM Storage Scale creates a service bridge which easily and transparently spans the distance using access credentials and existing storage protocols such as NFS and S3 Object.

IBM Storage Fusion HCI System supports various interaction mechanisms with other clusters, applications, or other data sources:

- ▶ Cross-cluster mounts to other IBM Storage Fusion or IBM Storage Scale clusters (roadmap)
- ▶ Data source for distant caches (AFM - Active File Management source)
- ▶ Data cache from a remote AFM source
- ▶ Scheduled asynchronous replication, RSYNC
- ▶ Synchronous replication, Metro-DR, supporting disaster recovery in a different site
- ▶ Import from an IBM Cloud Object Storage system while sharing its name space.



Why IBM Storage Fusion

This chapter focusses on **WHY** IBM Storage Fusion: What are the benefits of running containers in a turnkey Red Hat OpenShift cluster within IBM Storage Fusion? How can the integrated environment help reduce the effort required for deployment, high availability design, backup implementation and cyber resilience planning?

2.1 Why IBM Storage Fusion

IBM Storage Fusion HCI System is IBM “bare metal” Red Hat OpenShift cluster in-a-box that removes the infrastructure challenge of bare metal Red Hat OpenShift while enjoying its benefits. The complete hardware and software stack is supported by IBM and contains data protection and resilience capabilities as well as an enterprise class high performance storage.

IBM Storage Fusion SDS takes this software stack to custom-built Red Hat OpenShift clusters in the private and public cloud.

In the following sections we will explore these benefits in more detail.

2.2 Why Red Hat OpenShift on Bare Metal

IBM Storage Fusion HCI System has been designed based on bare metal Red Hat OpenShift due to its many advantages over hypervisor-based Red Hat OpenShift. Let's have a closer look at two of them.

Lower Total Cost of Ownership (TCO)

Red Hat OpenShift subscriptions (or licenses) for bare metal are counted by processor sockets: One subscription covers 2 sockets with a total of up to 64 CPU cores. Red Hat OpenShift on a hypervisor or any other platform is licensed by CPU cores assigned to the worker nodes, which usually turns out more expensive.

Taking a brief look at the list prices highlights the benefit:

A 32 core worker node licensed with bare metal subscriptions reduces the Red Hat OpenShift license cost by 81%. Considering our 64-core node, the savings grow to 91%. See Figure 2-1.

	Listprice	32 core worker	64 core worker
RHOCP for ISF with 5 year Premium S&S per 2 Cores	€ 19.222	€ 307.552	€ 615.104
RHOCP for ISF with 5 year Premium S&S per Server up to 2 Sockets	€ 57.656	€ 57.656	€ 57.656
Saving		81%	91%

Figure 2-1 Comparison to show increased savings using 64 core node

This significantly lowers the TCO of the container platform. The street price of the smallest HCI System configuration with three 32 core worker nodes (including all licenses), is lower than the price of the software licenses for a hypervisor-based cluster, even without taking the hardware and hypervisor cost into the equation.

Removing the Hypervisor from the equation

The Hypervisor is a software layer that not only needs to be licensed, it also needs to be installed, managed and updated, for which large IT shops usually employ a dedicated team of specialists.

Besides, a Hypervisor requires CPU and memory resources from the host system which cannot be used by Red Hat OpenShift or the containerized workloads.

The hypervisor will also introduce higher latencies to network and disk IOs that run through its virtualization layer, which in some cases may impact the application performance. See comparison in Figure 2-2.

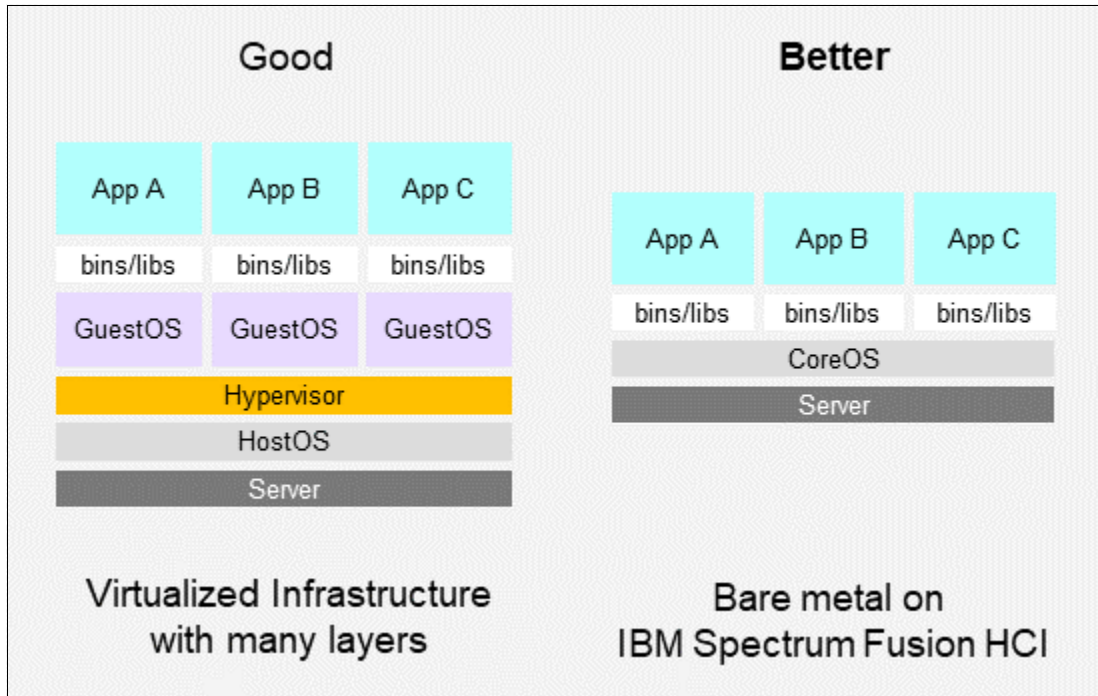


Figure 2-2 Comparison of virtualized infrastructure layers with bare metal on IBM Storage Fusion HCI System

2.3 Mission Critical Workloads in Containers

Workloads deployed in containers call for a cloud-like operating model:

They should be resilient to restarts and they should be able to scale out by running multiple instances. Webservers and the famous LAMP stack are considered being low hanging fruit here: The migration is trivial.

But mission critical monolithic workloads like RDBMS databases can just as well benefit from the container deployment and operating model. These workloads require the same level of performance, security, resilience, and data protection as they were previously getting in VMs or on bare metal servers. These requirements towards the container platform necessitate a well thought-through architecture and operating model.

IBM Storage Fusion HCI System was built from the ground up for mission critical applications like for example the IBM Cloud Paks. Its global data platform is based on a high-performance parallel file system that has a dedicated 200 Gbit network connection per node in order to avoid bottlenecks.

While the data is encrypted at rest and in flight, it's also resilient to double failures, which are rectified by the erasure coding algorithm that builds parities spread across multiple nodes.

Backups are not only taken via snapshots that reside on the file system and can be restored within minutes, they are as well off-loaded towards S3 storage to create a logical air-gap. The built-in Metro-DR option creates a high availability concept across two availability zones within a metropolitan area.

2.4 Getting Workloads into Production faster

Smaller companies are often very agile in building their container platform, as there are fewer and smaller teams involved. In many cases, only a handful of experts design and operate the full stack.

Big enterprises are facing different challenges, as multiple teams are involved. Infrastructure, Storage, Backup, Monitoring, and Operations; each team has its own concept and favorites, which will bring multiple products and vendors to the table. A heterogeneous complex software stack is created, with interoperability dependencies and support matrixes from multiple vendors.

This software stack needs to be monitored and updated, which makes “day2” operations difficult. Finally, the deployment and configuration of the stack should be automated with “recipe-style” infrastructure automation in order to avoid mistakes and improve repeatability.

Building a production-ready container platform this way will result in a long-running project with many resources.

IBM Storage Fusion HCI System is denominated as a “turnkey appliance” because it delivers a production-ready Red Hat OpenShift cluster in just one day. The appliance is delivered as a rack with all components mounted, cabled, and tested. Once connected to power and network, the fully automated installation starts.

Day2 Operations including updates and support activities are made easy as this software stack is fully supported out of one hand:

IBM provides support for

- ▶ Red Hat OpenShift
- ▶ Global Data Platform (SDS)
- ▶ Data protection
- ▶ Hardware support including Firmware updates (HCI only)

2.5 Storage for Containers

This section describes the advantages of container scalability and resiliency.

Scalability

A container platform is often built and sized for its initial use case. Once the container adoption in an enterprise gains momentum, new applications with new requirements get onboarded on the platform and an expansion is needed.

One of the big advantages of a container platform is that compute resources can be easily scaled out by adding more worker nodes to the cluster.

But what about storage?

Containers can access block, file, and object storage. Many applications that claim to require block storage will not use it in native block mode, but install a file system on it. Thus they actually need high-performance storage with “read write once” (RWO) capabilities, which IBM Storage Fusion can provide out of the box.

Block and file storage are again divided into “container ready storage” and “container native storage”. Each kind has its own advantages and disadvantages:

“Container ready storage” refers to attaching traditional storage arrays to containers.

This usually limits the scalability to a “scale-up” approach, as the cost of accessing storage in traditional arrays can vary from pod to pod depending on network and other limitations. Also, the deployment of traditional storage expansions is usually a large investment that requires specialized personnel rather than ongoing “maintenance” - type of activities.

“Container native storage” brings the advantages of containerization to the storage landscape. Capacity is increased via the scale-out approach. Whenever more storage-rich worker nodes are added to a cluster, each additional node provides not only capacity, but as well more IO performance and increased caching capabilities leveraging the added CPU, memory, and network bandwidth.

IBM Storage Fusion HCI System’s global data platform is a “container native storage”. The HCI is built from 6-20 nodes. Each node can hold up to 10 NVMe drives.

The storage capacity can be increased by either adding NVMe drives to existing nodes or by adding storage-rich nodes to the HCI cluster. All capacity increases are non-disruptive operations that are started from the IBM Storage Fusion UI.

Resiliency

Traditional storage arrays provide redundancy by creating a Redundant Array of Independent Drives RAID across their internal drives. Depending on the RAID level (Raid 5, 6, 10) some data is used to store the fault tolerance information called parity. But the RAID controller that is calculating the parity is a single hardware entity that cannot run as distributed software on a software-defined scale-out storage cluster.

Therefore, most software-defined storage products for containers rely on 2-way or 3-way replication. This means the complete data is copied and written 2 or 3 times in different failure zones.

- ▶ 2-way replication can only accept one failure and is therefore not suitable for most production environments.
- ▶ 3-way replication can accept two failures but only provides 33% data efficiency.

Dispersed Erasure Coding is a more recent technology to make distributed storage highly available while providing high data efficiency and short rebuild times. The erasure coding algorithm slices data into a configurable number of parts, builds the parity information, and distributes data and parity across the available nodes.

IBM Storage Fusion HCI System uses erasure coding and can be configured to use two different algorithm variants resulting in different fault tolerances and data efficiencies.

Figure 2-3 shows the IBM Storage Fusion installation dialog which explains the two available algorithms.

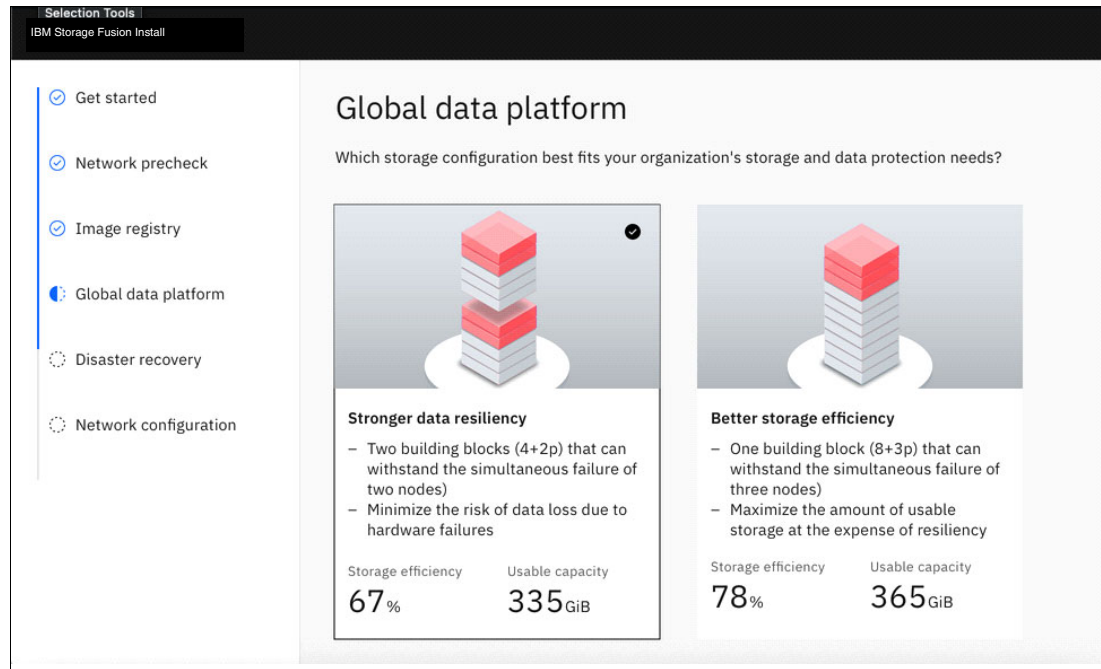


Figure 2-3 IBM Storage Fusion installation dialog

4+2p = Requires 6 servers, exposes 67% of the raw capacity as usable storage and can withstand 2 simultaneous failures.

8+3p = Requires 11 servers, exposes 78% of the raw capacity as usable storage and can withstand 3 simultaneous failures.

This choice must be made with good foresight as it cannot be easily undone. Enterprises will usually have policies of their own when it comes to error containment vs. efficiency.

2.6 Container Platform Resiliency and High availability

Cloud-native applications can be distributed across multiple clouds and Availability Zones (AZs). A failure of one cloud platform or one Red Hat OpenShift cluster will not bring down the whole application due to its built-in redundancy.

Monolithic applications behave differently: They rely on single instances and the availability of attached persistent Storage, which is usually provided by the hardware level.

In case of a complete node, rack or availability zone failure, workloads can be restarted on the remaining worker nodes in the Red Hat OpenShift cluster. This process is a lot easier than restarting a monolithic application in a standby site.

Container platforms like Red Hat OpenShift are resilient by design. Nevertheless, tolerating the failure of an entire Availability Zone requires careful planning. The Red Hat OpenShift control plane is built from three master nodes that represent the quorum and tolerate a maximum network latency of 5ms between nodes.

If three AZs with low latency interconnects are available, the Red Hat OpenShift cluster including the storage could span all three AZs. See Figure 2-4.




Figure 2-4 Stretched Cluster with three Availability Zones

Each availability zone should host 1/3 of the cluster capacity.

Due to the Red Hat OpenShift network latency requirement of max 5ms, the three AZs would normally be located in the same campus or building. If a larger physical separation of two AZs is required, stretching the Red Hat OpenShift cluster is not an option. Two independent Red Hat OpenShift clusters need to be set up.

IBM Storage Fusion HCI System’s global data platform includes the “Metro-DR” feature, which provides a synchronous storage replication between two AZs with a maximum network latency of 40ms which is equivalent to ~150 km signal distance. See Figure 2-5.

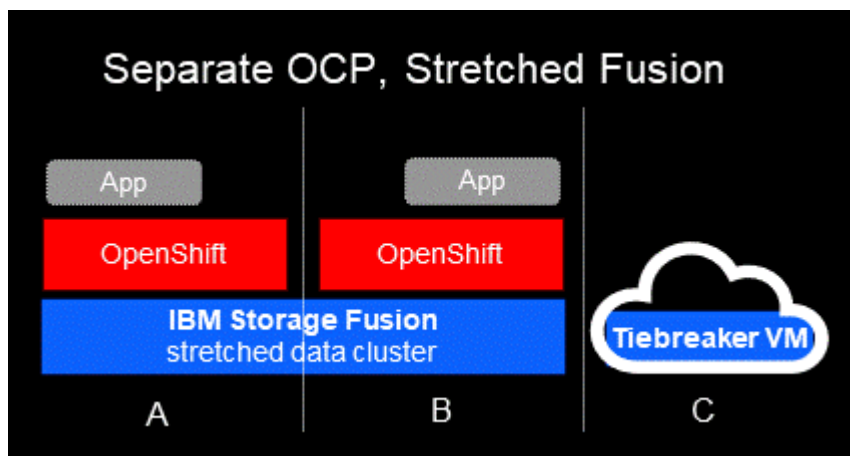


Figure 2-5 IBM Storage Fusion HCI System’s global data platform includes Metro-DR

Two IBM Storage Fusion HCI Systems form a Metro-DR cluster with a Tiebreaker in a 3rd location. The network connection to the Tiebreaker VM in C can tolerate higher network latencies.

IBM Storage Fusion HCI System with Metro-DR protects monolithic applications that would otherwise be forbidden in a containerized environment without AZ fault tolerance.

Security

Some security and data management challenges still persist in container environments. A cyber-resilient storage layer should protect your data in flight and at rest. Keep in mind that many protocols like nfsv3 are not encrypted. IBM Storage Fusion's global data platform is based on IBM Storage Scale. It encrypts data at rest and in flight.

More details can be found in the Redbooks IBM Spectrum Scale Security at:

<https://www.redbooks.ibm.com/abstracts/redp5426.html>

At its core, the Red Hat OpenShift Container Platform relies on the Kubernetes project to provide the engine for orchestrating containers across many nodes in scalable data centers. Red Hat OpenShift Container Platform is designed to lock down Kubernetes security and integrate the platform with a variety of extended components.

See more details at:

<https://docs.openshift.com/container-platform/4.11/security/index.html>

2.7 Data Protection for Containers

Data from applications in containers require the same level of data protection like applications in VMs. In times of widespread cybercrime and ransomware attacks, backups should be protected in an air-gapped environment.

Therefore, data protection is a fundamental part of IBM Storage Fusion. Containers are built from container images that are stored in a container registry. It is important to have a reliable backup of the container registry, but there is no need to back up running container images (again).

What is important though is to back up the information that makes container instances unique, like passwords, service names or the persistent data of the container. This information is stored in resources like secrets, services, routes, and a lot more.

Figure 2-6 on page 19 gives an overview of namespace and cluster resources. IBM Storage Fusion creates a backup of the cluster and namespace resources.

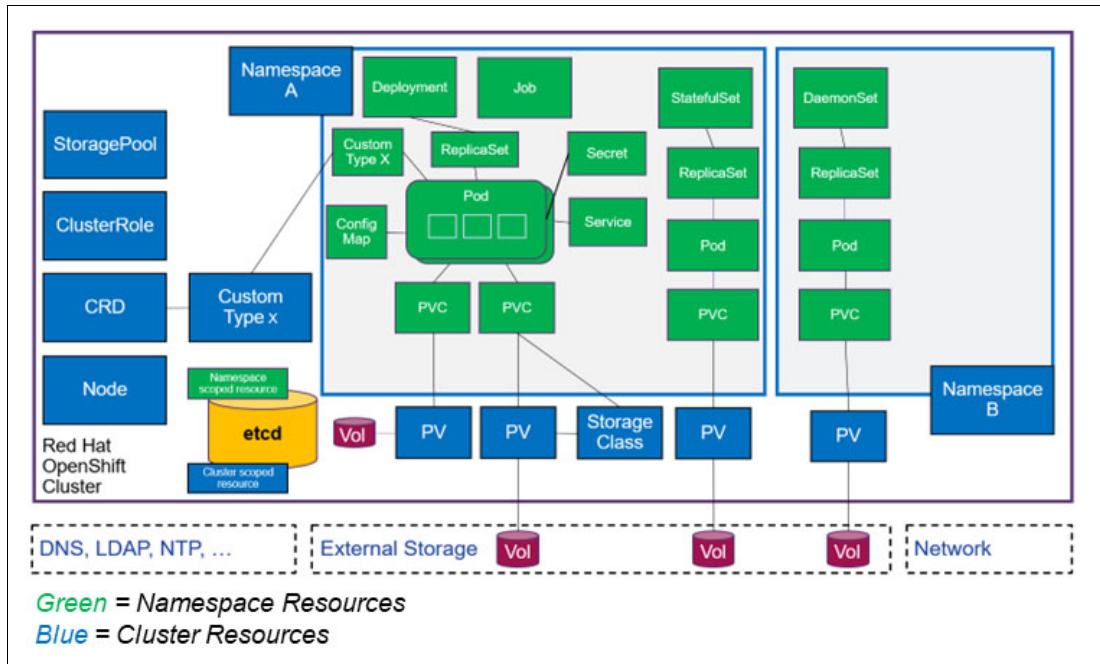


Figure 2-6 Red Hat OpenShift Cluster and Namespace Resources

The persistent data of a container is stored in separate volumes called “persistent volume claims” (PVCs) due to the way they are requested from the control plane (as “claim”).

IBM Storage Fusion can take snapshots of PVCs for point-in-time backups. Snapshots are always stored on the same file system as the primary PVC data. How often these snapshots are taken and how long they remain available is defined in policies.

In order to protect snapshots from file system corruption, malicious destruction or physical damage, they can be periodically copied to an Object Storage service or to an external backup server. Backup and restore are managed from within the IBM Storage Fusion user interface. See Figure 2-7.

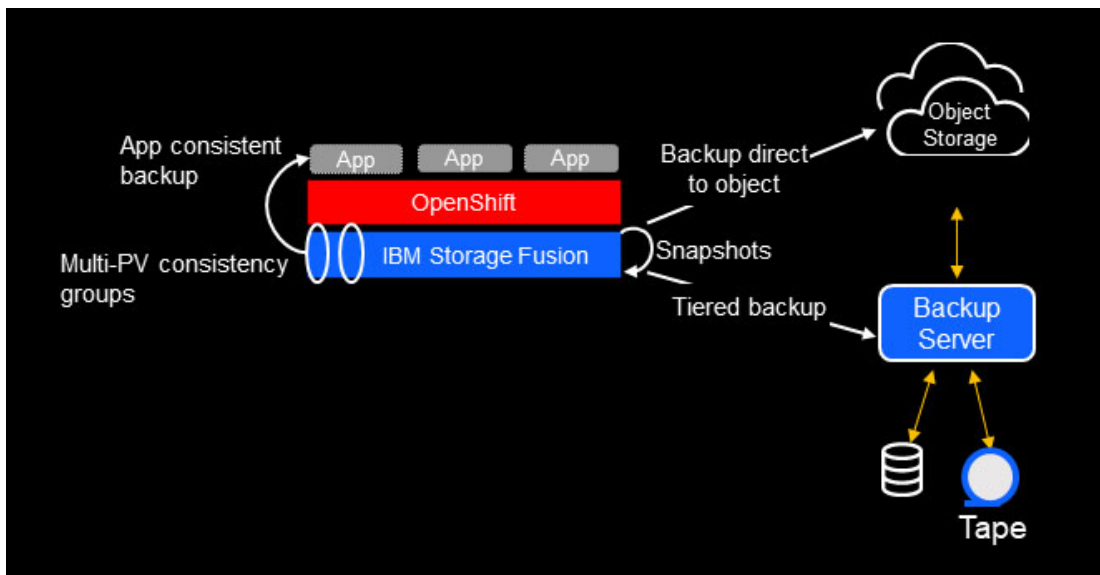


Figure 2-7 Backup and restore are managed within the IBM Storage Fusion user interface

2.8 Red Hat OpenShift made easy: Day2 Operations with IBM Storage Fusion

This section describes how Red Hat OpenShift is made easy with Day2 Operations with IBM Storage Fusion.

Hardware monitoring

Today's servers and switches have built-in hardware redundancy like multiple fans and power supplies. But the best redundancy does not help if failed components are not identified and replaced as soon as possible, minimizing the risk of double failures.

Red Hat OpenShift does not include hardware monitoring of the underlying infrastructure out-of-the-box. In particular, when Red Hat OpenShift runs on top of a hypervisor, hardware monitoring needs to be implemented on the hypervisor layer or below. If Red Hat OpenShift is deployed on bare metal, infrastructure monitoring and alerting has to be implemented by the user via 3rd party tools.

IBM Storage Fusion HCI System includes a full stack hardware monitoring and call home functionality. Events are raised and identified failed components can be replaced ASAP. This functionality is only available with the HCI version, not with IBM Storage Fusion SDS.

Single Pane of Glass UI

IBM Storage Fusion HCI System features a “single pane of glass” user interface.

The Infrastructure Dashboard provides an overview of the complete hardware and software stack (see Figure 2-8). This holistic view eases troubleshooting and allows diving deeper into each component with a single click.

IBM Storage Fusion SDS includes a UI with the same look and feel but without the infrastructure dashboard.

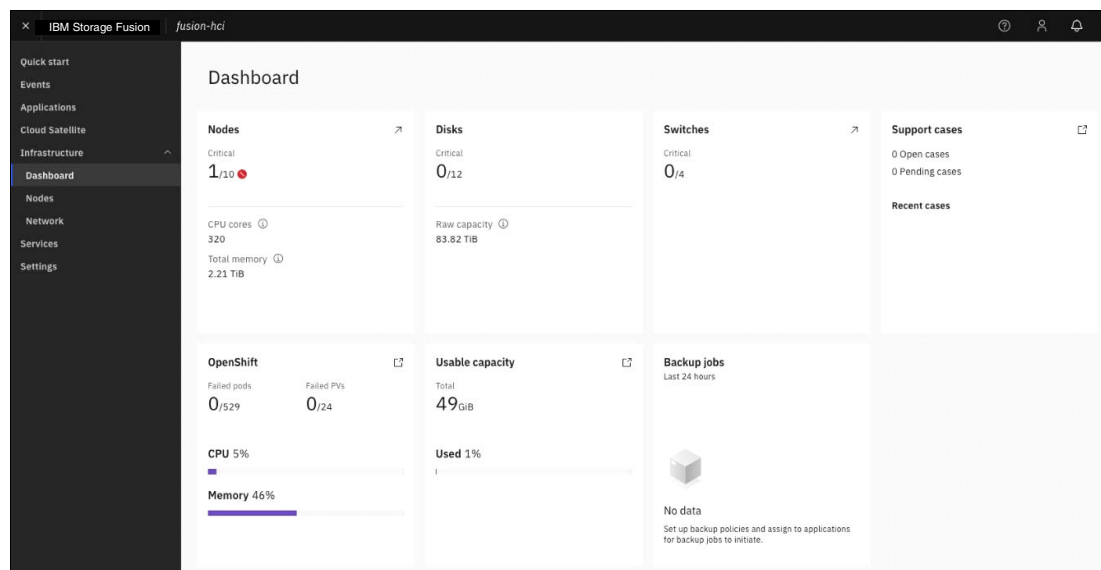


Figure 2-8 IBM Storage Fusion UI Dashboard

The UI also contains an “Applications” section in which details about all Red Hat OpenShift projects can be found. Backup policies can be easily assigned to applications and an overview of the storage consumption and status is provided.

Infrastructure as Code

IBM Storage Fusion HCI System does not include any management host or console.

As illustrated in the architecture diagram shown in Figure 2-9, all software components are fully container-native and run on top of Red Hat OpenShift.

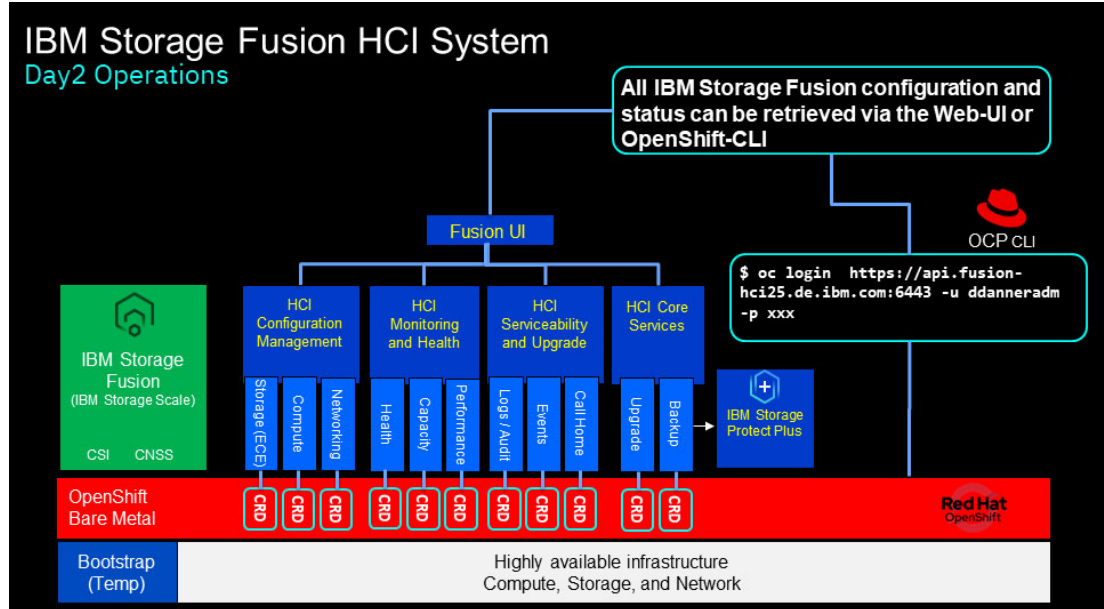


Figure 2-9 IBM Storage Fusion HCI System Software Architecture Diagram

The IBM Storage Fusion software stack is deployed as an Red Hat OpenShift operator and updated via the Red Hat OpenShift operator lifecycle management concept (OLM).

IBM Storage Fusion adds many custom resource definitions (CRDs) to the Red Hat OpenShift cluster. This enables IBM Storage Fusion to store its configuration and current states in custom resources (CRs) in the Red Hat OpenShift cluster configuration database (etcd).

This architecture provides many advantages for “Operations”:

Changing configurations by editing human-readable YAML files is the Red Hat OpenShift-native way of managing the cluster. IBM Storage Fusion custom resources can be accessed and changed either from the Red Hat OpenShift UI or from the command line interface CLI. This eases the automation of configuration changes like creating backup policies or adding backup policies to namespaces.

The IBM Storage Fusion HCI System software includes compute and network operators that are not available for custom-built Red Hat OpenShift clusters on bare metal or a hypervisor.

These operators take care of the hardware management within the IBM Storage Fusion HCI System and transmit details like firmware levels, temperatures, hardware component failures, network vlans, packet counters per network port, and more to Red Hat OpenShift.

This “infrastructure as code” can be used to ease troubleshooting and allow a DevOps operating model.

Updating the Complete Red Hat OpenShift Software Stack

In order to harden the system against cybercrime and to be compliant with various security audits, updating only Red Hat OpenShift is not sufficient.

The full software stack including the firmware of servers and switches, the hypervisor software and CSI drivers for storage systems need to be updated as well. See Figure 2-10.

Before the stack is updated, the interoperability of the new versions of the components need to be verified.

For each IBM Storage Fusion release, IBM delivers a supported and pre-tested software bundle including:

- ▶ IBM Storage Fusion operator
- ▶ Data Protection software
- ▶ Global Data Platform (SDS)
- ▶ Firmware updates for servers and switches (HCI only)

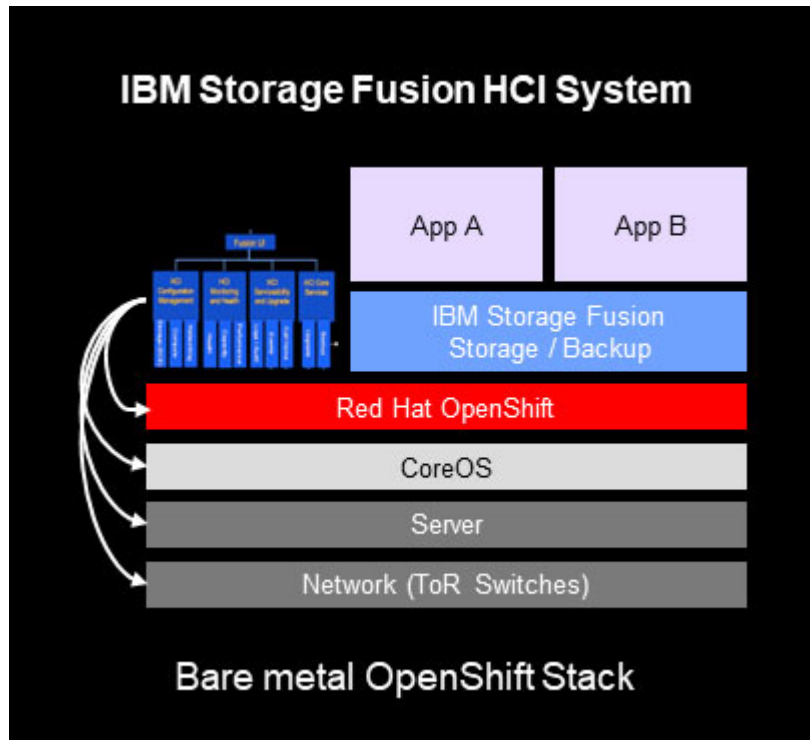


Figure 2-10 IBM Storage Fusion Bare metal Red Hat OpenShift Stack

While Red Hat OpenShift and the IBM Storage Fusion operator are updated from the underlying Red Hat OpenShift UI, all other components are updated from the IBM Storage Fusion UI.

A custom-built Red Hat OpenShift stack on top of a Hypervisor includes more layers from different vendors. The software needs to be obtained from many sources, each coming with its own installation instructions and interoperability guidelines.

The numerous teams (Figure 2-11) involved in larger enterprises increase the complexity and coordination effort of maintenance activities.

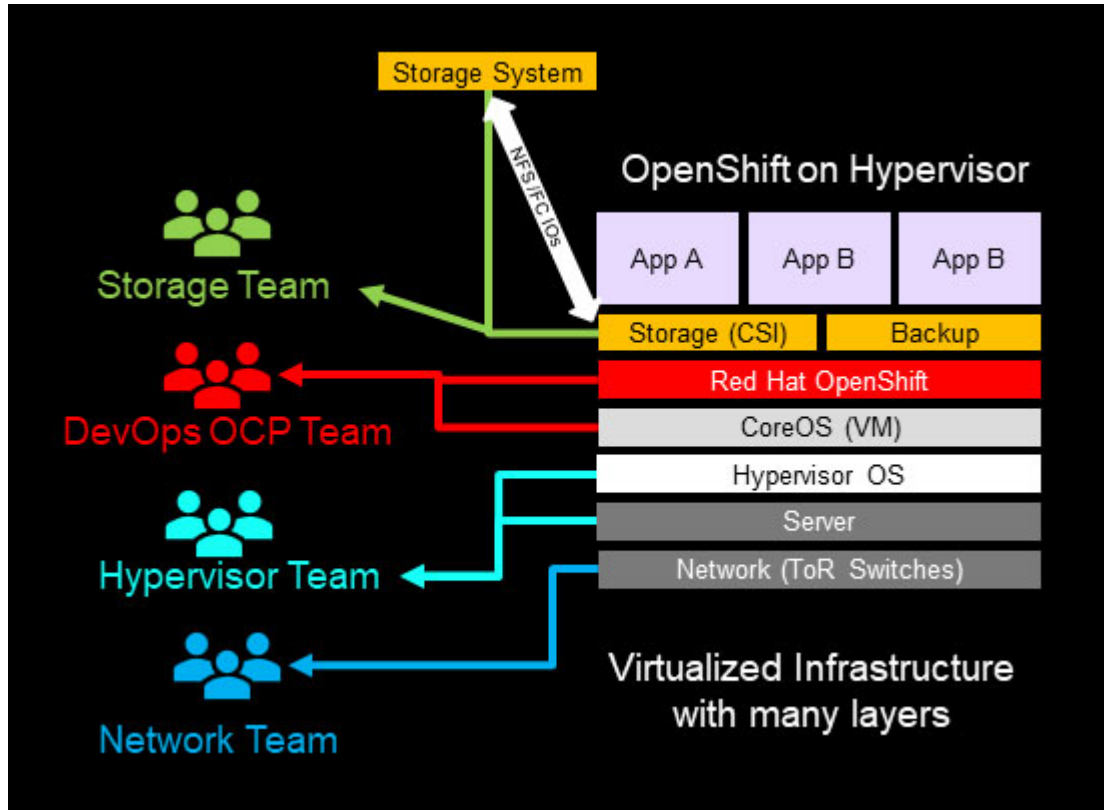


Figure 2-11 Custom-built Red Hat OpenShift stack on top of a Hypervisor includes more layers from different vendors adding complexity



Flavors of IBM Storage Fusion

This chapter describes flavors of IBM Storage Fusion implementations, major differences between HCI (hardware) and SDS (software) variants, and which choices to make to right-size the product for current needs while retaining the ability to scale out at a future stage.

3.1 Integration with legacy VM infrastructures

IBM Storage Fusion goes one step further in also allowing the integration of legacy app implemented as VMs. This aspect is important for hybrid applications, where some parts like a legacy database, cannot yet be containerized, however these type of scenarios need to obey the same deployment, monitoring, availability, and security paradigms.

Remember: VMs are still machines, even though they are virtualized from the underlying hardware. Therefore, they need to be managed as such. In contrast, container platforms do not consider machines, but rather services, templates, recipes, and resource claims. These resources can include instances of virtual machines.

The more we modernize our applications, the more we will be able to write “infrastructure as code”; namely, recipes and services/resource claims. The classic concept of a virtual machine that conceals all mutual dependencies and security aspects is not adequate for modern “agile” or “cloud-native” programming and dev-ops-style deployment.

3.2 Three Tier Architecture System

If we look at the traditional way an IT system is built, we have three independent tiered layers known as the 3 tier “siloes” architecture. These are as follows:

- ▶ Network Tier
- ▶ Server Tier
- ▶ Storage Tier

Figure 3-1 shows the typical construction of a three tier architecture system.

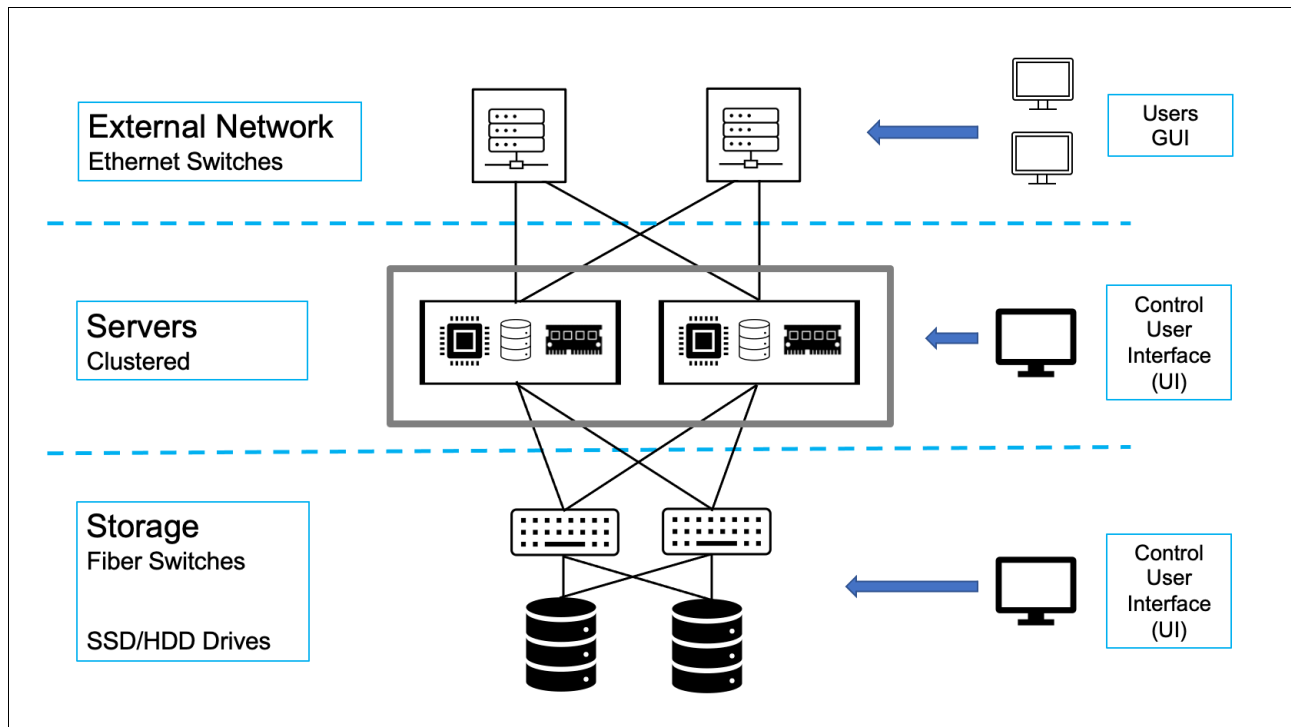


Figure 3-1 Three tier architecture system

In this example we have separate tiers where the storage, server and network elements are all installed and run as individual items. This could be from potentially different vendors and thus separate interfaces or controls for each one. The complexity, operation and running of the complete system thus requires multiple users' interfaces and training for the operation teams. A hyper-converged infrastructure attempts to address these issues, and by doing this, reduces the complexity and costs of running all the separate elements of the three tiered "siloes" approach.

3.3 Hyper-Converged Infrastructure (HCI)

The HCI methodology combines servers, storage, and network plus the required software defined storage components to complete the IT system in its entirety. Instead of the traditional separation of compute, network, and storage resources, HCI employs a simplified and unified way to couple all the resources together and have a single management operational interface. HCI also provides a way to give easily adapted scalability functions in the event that the system needs to grow to accommodate expanding and more demanding workload.

Figure 3-2 shows the layout of a typical HCI system.

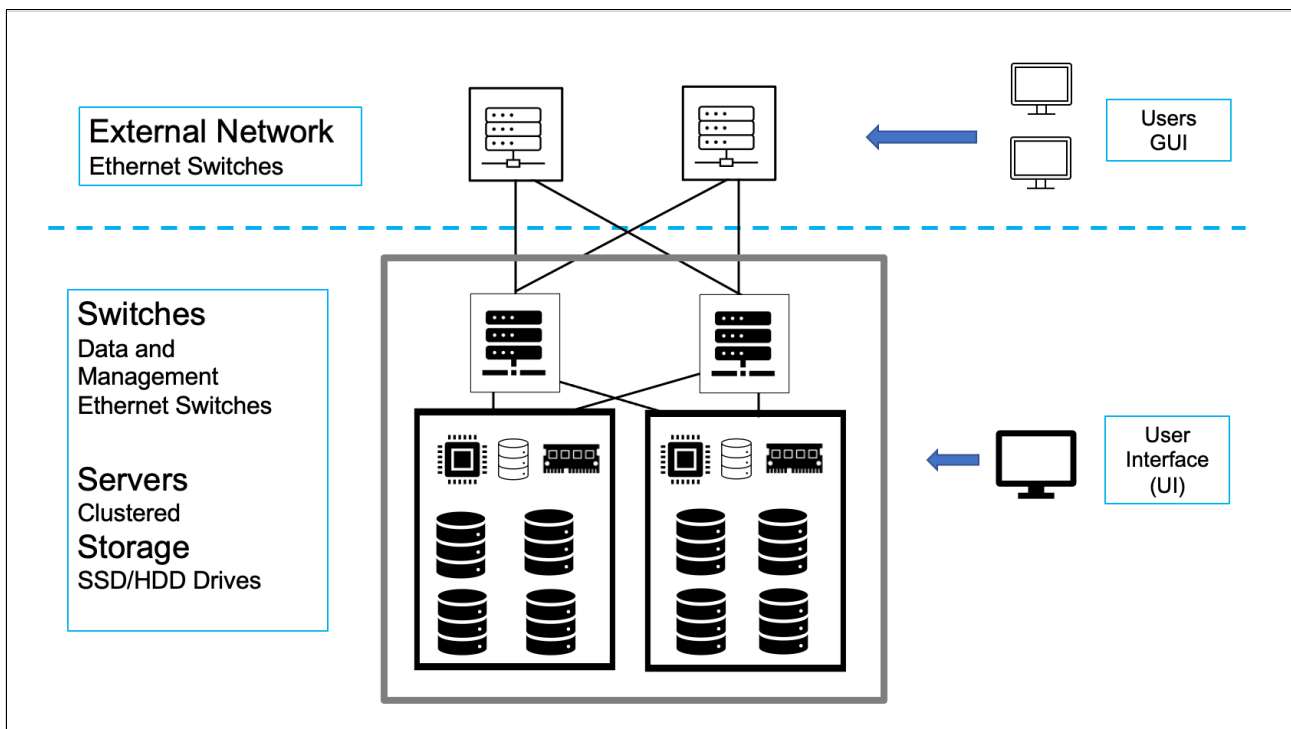


Figure 3-2 HCI system overview

Currently one of the most common adoptions of HCI technology is Virtual Desktop Infrastructure (VDI) application. IBM's vision of HCI is beyond the simplification of the management and legacy VDI workloads. In the modern HCI, it should optimize for container deployment for hybrid cloud and Edge environments. It should support enterprise features such as seamless backup and restore, disaster recovery, enterprise security and high availability. The modern HCI should not be an infrastructure silo. It should be able to create a unified storage platform to allow seamless application mobility and data sharing from clouds to the edge.

3.4 IBM Storage Fusion Options

IBM Storage Fusion is offered either with a bare metal RHOCP deployment in an x86-based Hyper Converged Infrastructure (HCI) appliance form factor, or as a pure Software-Defined Storage (SDS) solution.

IBM Storage Fusion is built from proven, enterprise-hardened technologies like IBM Storage Scale and IBM Spectrum Discover (now an integrated IBM Storage Fusion service called Data Cataloging) and provides a bare metal foundation for the rapid deployment and easy management of any RHOCP environment. We've brought the entirety of our rich heritage and vast IP portfolio in supercomputing, AI/ML, and industrial-strength data resilience to bear on the design of Fusion. The result is a performance- and cost-optimized platform that unlocks the full potential of RHOCP - at any scale.

3.4.1 IBM Storage Fusion Software Overview

The IBM Storage Fusion software consists of several building blocks which together provide the function and features of the software solution.

Figure 3-3 shows the IBM Storage Fusion software layout overview.

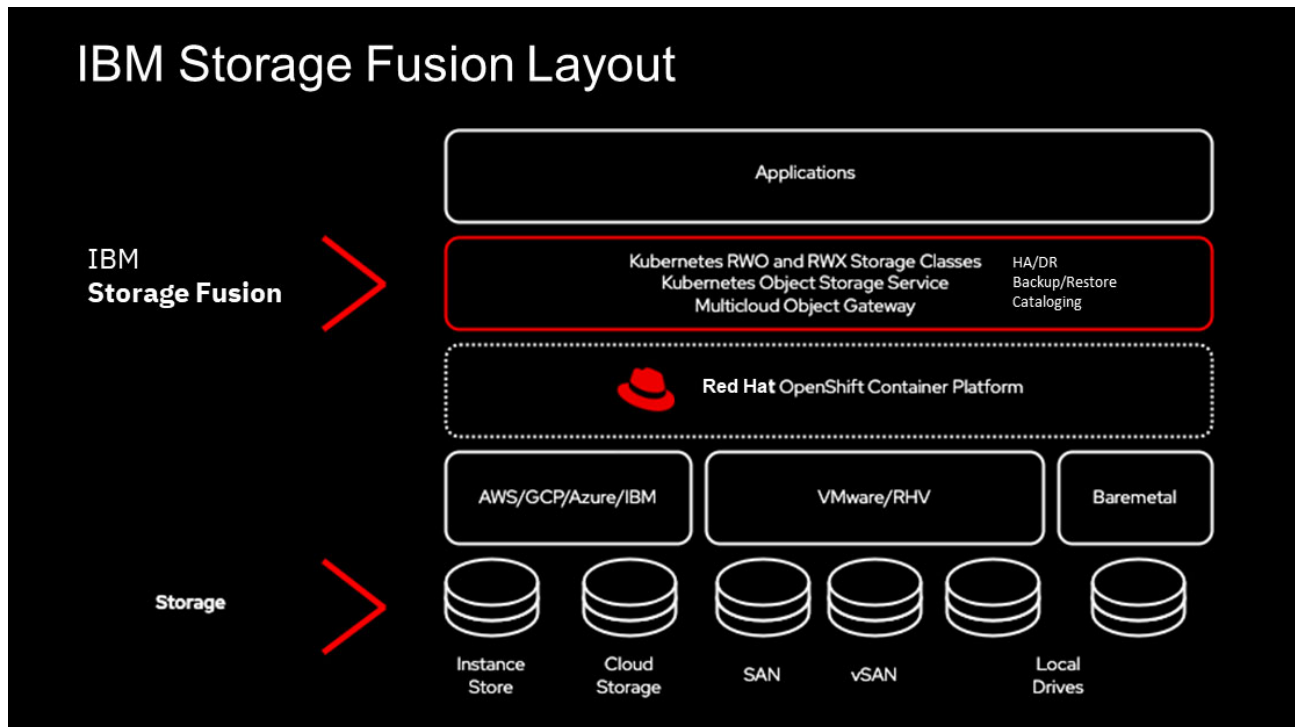


Figure 3-3 IBM Storage Fusion software layout

IBM Storage Fusion and the inherent multi-cloud deployment model of IBM Storage Fusion Data Foundation, in conjunction with Red Hat OpenShift itself, are key to addressing and enabling the agility, simplicity, and productivity improvements that create efficiencies for customers.

IBM Storage Fusion Data Foundation exposes physical storage to the applications that are deployed on Red Hat OpenShift. Kubernetes abstracts the storage using container storage interfaces (CSI) and Storage Classes, but it does nothing more to help manage the storage nor to provide uniformity of presentation and usage across multiple deployment models. IBM Storage Fusion Data Foundation delivers consistent volume types, resilience, and tiering wherever the storage layer is provisioned.

The IBM Storage Fusion Data Foundation documentation is split into sections for each type of environment describing how to deploy available storage in each one. At the edge, there is a new option for very small deployments of Single Node Red Hat OpenShift (SNO) matched by IBM Storage Fusion Data Foundation with a more specialized deployment of Local Volume Manager (LVM) instead of Ceph, and MultiCloud Object Gateway.

In public cloud environments, there are managed service offerings for storage on Red Hat OpenShift Service on AWS (ROSA) and Red Hat OpenShift Dedicated. In all environment types, there are self-managed options in which the client can choose their preferred infrastructure and storage model.

3.4.2 IBM Storage Fusion HCI System Appliance Overview

IBM's model for the complete HCI appliance is to combine all the hardware outlined above together with the software components listed below, plus a custom-built management user interface (UI), all running as containerized applications within the IBM Storage Fusion HCI System appliance.

This gives the IBM Storage Fusion HCI System the following characteristics:

- ▶ Integrated HCI appliance for both containers and VMs using Red Hat OpenShift
- ▶ Highly scalable containerized file system with erasure using IBM Storage Scale ECE
- ▶ Data resilience for local and remote backup and recovery using IBM Storage Protect Plus
- ▶ Simple installation and maintenance of hardware and software through the custom-built UI
- ▶ Global data platform stretching from public clouds to any on-prem or edge locations
- ▶ IBM Cloud Satellite® and Red Hat ACM (Active Cluster Management) native integration
- ▶ Ready for AI applications with optional GPUs and global data for better AI
- ▶ Starts small with 6 servers and scales up to 20 (with HPC GPU enhanced options).

So, in summary, these are the highlights of the IBM Storage Fusion HCI System appliance and all the advantages of this hyper-converged approach using the outlined building blocks.

3.4.3 IBM Storage Fusion HCI System Software Overview

The IBM Storage Fusion HCI System software consists of several building blocks which together provide the function and features of the software solution within the HCI appliance.

Figure 3-4 shows the IBM Storage Fusion HCI System software overview.

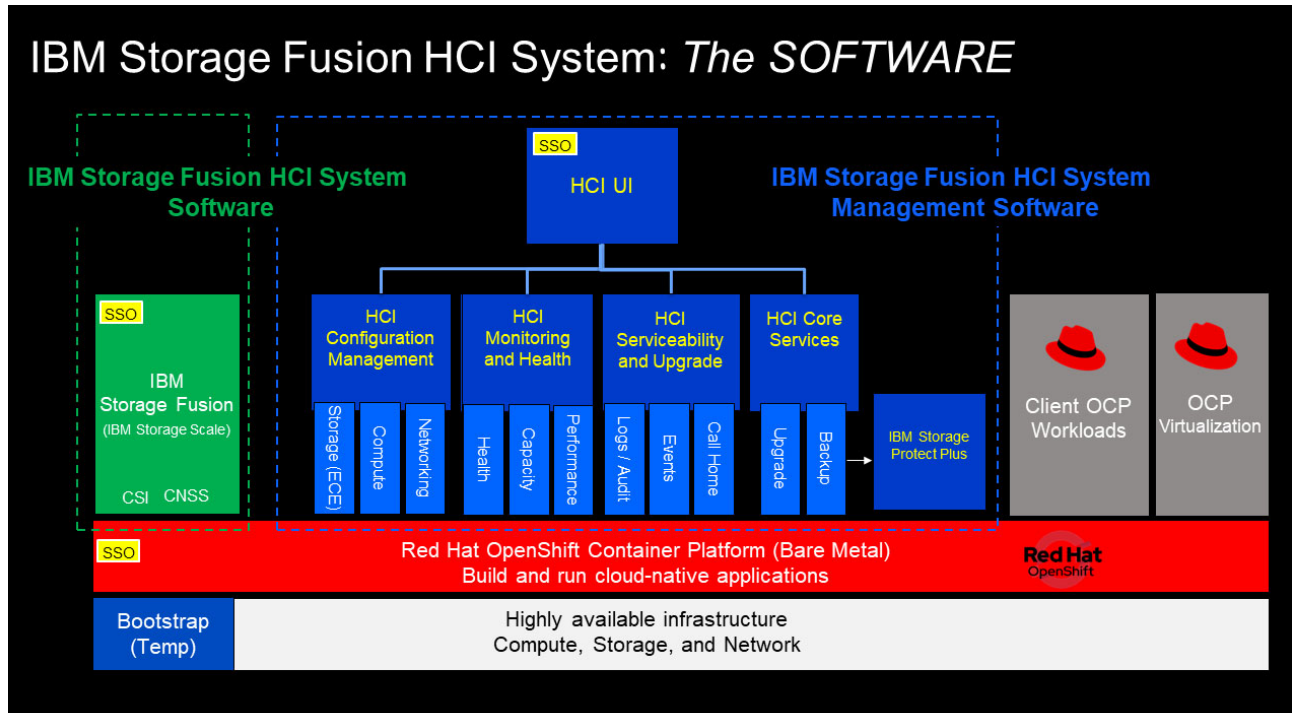


Figure 3-4 IBM Storage Fusion HCI System software

Some of the main points of the HCI software stack are:

- ▶ Red Hat OpenShift on Bare Metal is installed across all nodes with the bare metal OS being Red Hat CoreOS (RHCOS).
- ▶ Next, a set of HCI management services are installed in blue on top of Red Hat OpenShift that manages the compute, storage and networking:
- ▶ The HCI UI operator provides a single pane of glass to manage IBM Storage Fusion and is integrated with Single Sign On provided by Red Hat OpenShift.
- ▶ The storage operator deploys a best practice installation of IBM Storage Scale ECE.
- ▶ The compute operator manages the physical servers including firmware upgrade, power on/off, and scaling the number of nodes in the cluster.
- ▶ The networking operator manages the switches including switch firmware and configurations.
- ▶ The HCI operators monitor the health of the system for problem resolution, provide performance monitoring capabilities, present capacity utilization.
 - Allows the user to know when you need to add more drives for storage capacity or more physical servers for more compute capacity.
 - Events are natively integrated with Red Hat OpenShift and can Call Home to IBM in the event of a physical infrastructure or software problem with automatic log collection and upload.

- ▶ IBM Storage Protect Plus is integrated with the backup operator and can backup both the IBM Storage Fusion control plane as well as deployed client applications on Red Hat OpenShift.
- ▶ Finally, the upgrade operator orchestrates the upgrade of firmware on the hardware, Red Hat OpenShift, the HCI Management Stack, IBM Storage Scale ECE and IBM Storage Protect Plus.

3.4.4 IBM Storage Fusion HCI System Hardware Overview

The IBM Storage Fusion HCI System consists of several hardware building blocks which together provide the function and features of the appliance.

Figure 3-5 shows the HCI hardware overview.

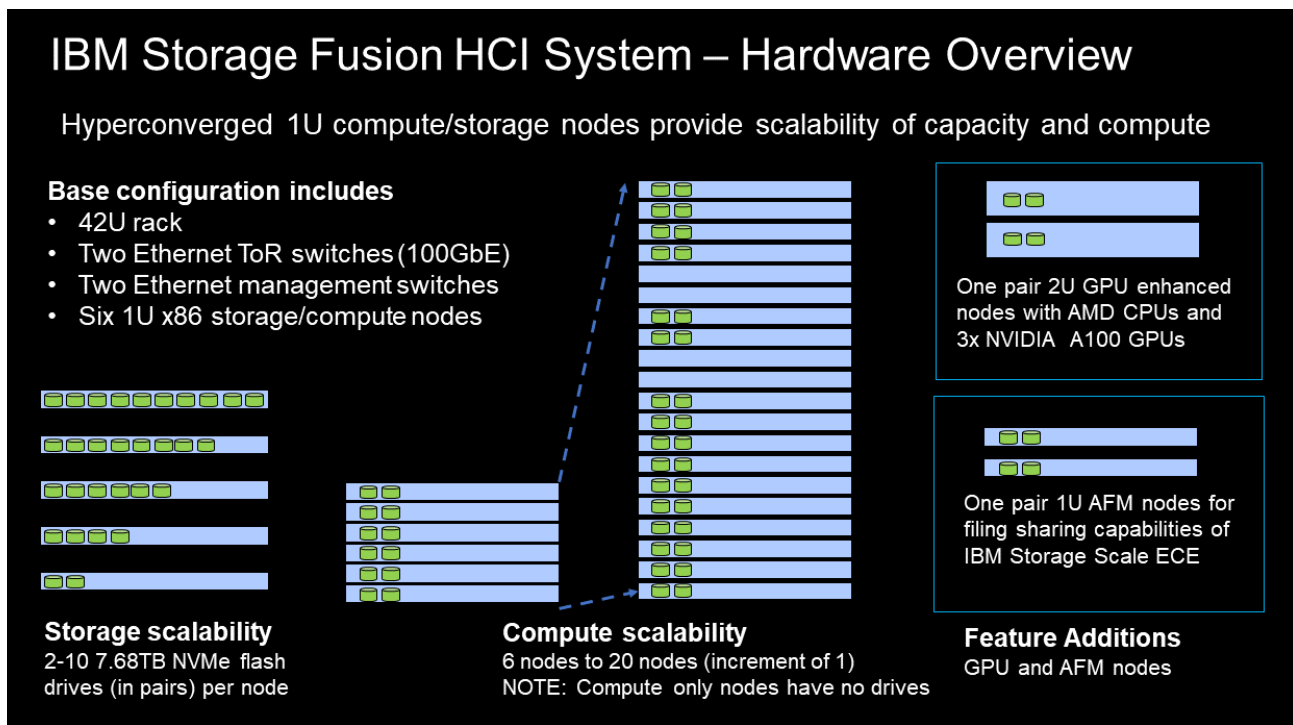


Figure 3-5 HCI hardware overview

The base configuration includes:

- ▶ IBM 9155 Model R42 42U rack including the power distribution units (PDU's)
- ▶ IBM 9155 Model S01 - 2x Ethernet high-speed switches
- ▶ IBM 9155 Model S02 - 2x Ethernet management switches
- ▶ IBM 9155 Model C01 - 6x 32 Core Storage/Compute servers with 2 NVMe drives/servers
 - OR
- ▶ IBM 9155 Model C05 - 6x 64 Core Storage/Compute servers with 2 NVMe drives/servers
 - The base system will always include the following models (quantity per model)
 - 6 x C01 or 6 x C05 or 3 x C01 + 3 x C05
 - 2 x S01 Ethernet high-speed switches
 - 2 x S02 Ethernet management switches

- 1 x R42 Rack
- 1 x TF5 KVM Feature (optional)
- The base system can be extended with the following models:
 - C00 - 32 Core Storage/Compute servers with no drives
 - C01 - 32 Core Storage/Compute servers with 2 NVMe drives/servers
 - C04 - 64 Core Storage/Compute servers with no drives
 - C05 - 64 Core Storage/Compute servers with 2 NVMe drives/servers
 - F01 - 2x 16 Core servers used for Active File Management
 - G01- 2x 24 core server with 3x NVIDIA GPUs.
- The C00, C01, C04, and C05 models of servers can be intermixed in the system, but the base 6 servers must have NVMe drives installed, so currently have to be C01 or C05 types.

Notes: The server in rack position seven (RU7) is the provisioner node, connected to the optional KVM feature. On older systems, this KVM was a standard feature, however on current system orders, the KVM feature is optional.

This KVM is used by the IBM SSR to configure the initial systems with network and base Red Hat OS to then move to the OCP and Fusion software install. The servers in positions two, three and four (RU2, RU3, and RU4) become the Red Hat OpenShift control plane.

On systems without the KVM feature, the SSR will use a service laptop to connect into the node in rack position seven (RU7) and then perform the initial install phase.

IBM Storage Fusion HCI System 2.5.2 also offers the following new enhancements:

- ▶ Install IBM Storage Fusion HCI System components in your own rack
 - IBM Storage Fusion HCI System components can be ordered for installation into client-supplied racks. When ordering, clients have the option to order an IBM Storage Fusion HCI System fully racked from the IBM factory or as components that can be installed into a client-provided rack.
 - For more information see here: [Installing the HCI System](#)
- ▶ High Availability Cluster Support
 - IBM Storage Fusion HCI Systems can be configured into a three-rack high availability (HA) cluster. This configuration is commonly referred to as a 3 availability zone (3AZ) deployment, a strategy useful for ensuring application availability.
 - You can connect three IBM Storage Fusion HCI System racks to create a single large Red Hat OpenShift Container Platform and IBM Spectrum Scale ECE cluster.
 - The components in each of the three racks are the same as a single rack.
 - For more information see here: [High Availability Cluster](#)
- ▶ Expansion Rack Support
 - IBM Storage Fusion HCI System can add new IBM racks to expand your Red Hat OpenShift Container Platform cluster.
 - For more information see here: [Adding Racks](#)

- ▶ Tailored Power and Network Options
 - IBM Storage Fusion HCI Systems will only include the number of power distribution units (PDUs) and network cables required to support the configured order. MES orders will add PDUs, and cables as needed.
 - For more information see here: [Power Configuration Calculations](#)

More detailed information is available [here](#) where the new full feature set for the IBM Storage Fusion HCI system is explained.

3.4.5 Scalability in HCI

IBM Storage Fusion HCI System is a highly scalable appliance, not only in computing power, but also in internal storage capacity within one frame. The IBM Storage Fusion HCI System is also scalable buy adding extra nodes (either storage/compute or just compute only) for increased power and storage.

Server / Disk upgrade options

The server and disk (SSD) upgrade options are as follows:

- ▶ Increased storage by adding pairs of drives to storage/compute servers
 - 7.68TB NVMe PCIe Gen4 drives/server to a max of 10 drives/server
- ▶ Increased compute power / storage by adding servers
 - Sold in pairs
- ▶ AFM (Active File Manager) delivered as a pair of servers
- ▶ A pair of 2U GPU servers, each with 3x NVIDIA A100 GPUs

Server Memory Options

The following memory upgrade options per server are available:

- ▶ Standard memory in the 9155-C01 and 9155-C00 servers is 256GB
 - Memory can now be upgraded in these C01 and C00 servers
 - Note: AFM and GPU servers are not eligible for memory upgrades
 - Memory can be upgraded in manufacturing
 - The manufacturing update increases the RAM to 512GB
 - All C01 and C00 servers must be ordered with the same memory configuration
- ▶ Upgrade in the field (MES)
- ▶ Upgrade C00 or C01 servers to either 512GB or 1024GB
 - The 1024GB upgrade cannot be done in manufacturing

Note: Because of power requirements, only systems with 3-phase power or 60A single-phase power are eligible for the 1024GB upgrade. Systems using 30A single-phase power cannot support this memory upgrade.

- ▶ Standard memory in the 9155-C04 and C05 server is 1024GB
 - Memory can now be upgraded in these servers
 - Memory can be upgraded in manufacturing

- The manufacturing update increases the RAM to 2048GB
- All C04 and C05 servers must be ordered with the same memory configuration
- ▶ Upgrade in the field (MES)
 - Upgrade C04 and C05 servers to 2048GB

3.4.6 Configuring the IBM Storage Fusion HCI System with IBM Storage Modeller

The IBM Storage Fusion HCI System appliance is configured using a tool call IBM Storage Modeller (StorM). IBM Storage Modeller is an IBM sales process tool that is designed to model storage solutions for customer applications. This tool is for use by IBM and Business Partner technical sellers.

Storage Modeller is developed by IBM and functionality is being delivered over time. Storage Modeller is designed to integrate with upstream and downstream sales tools, with the focus on optimizing seller workflow and reducing time from Opportunity ID to actual sale of a solution. Storage Modeller is a cloud-based tool accessed through a web browser-based user interface.

Figure 3-6 shows the Storage Modeller process for integration with IBM sales tools.

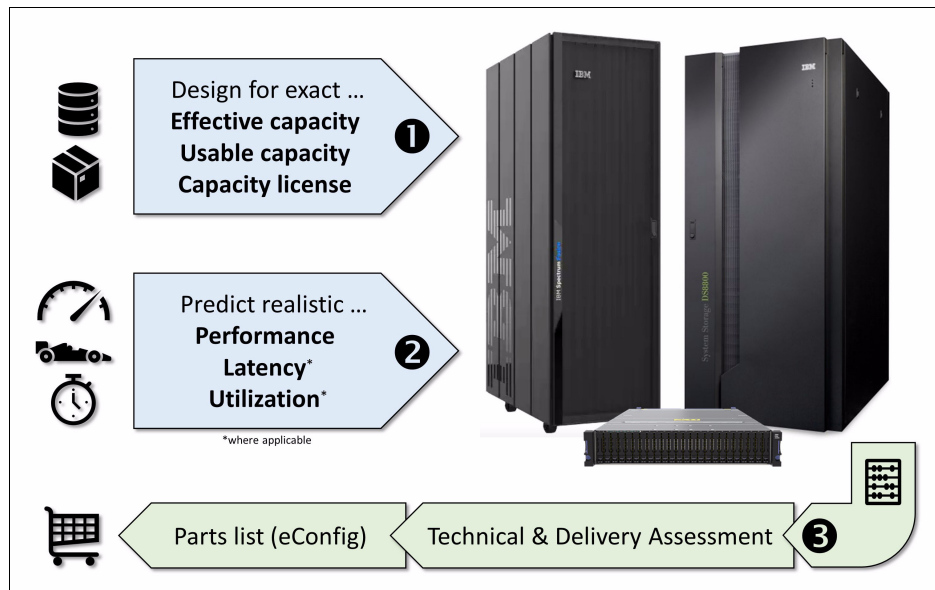


Figure 3-6 Typical storage system design in the StorM tool

As the StorM tool is used, and the configuration created, the UI display will show a graphical build of the system. More information on the StorM tool and a virtual tour can be found here:

[IBM Storage Modeller](#)

Figure 3-7 on page 35 shows the IBM StorM graphical output representation of the HCI appliance configuration.

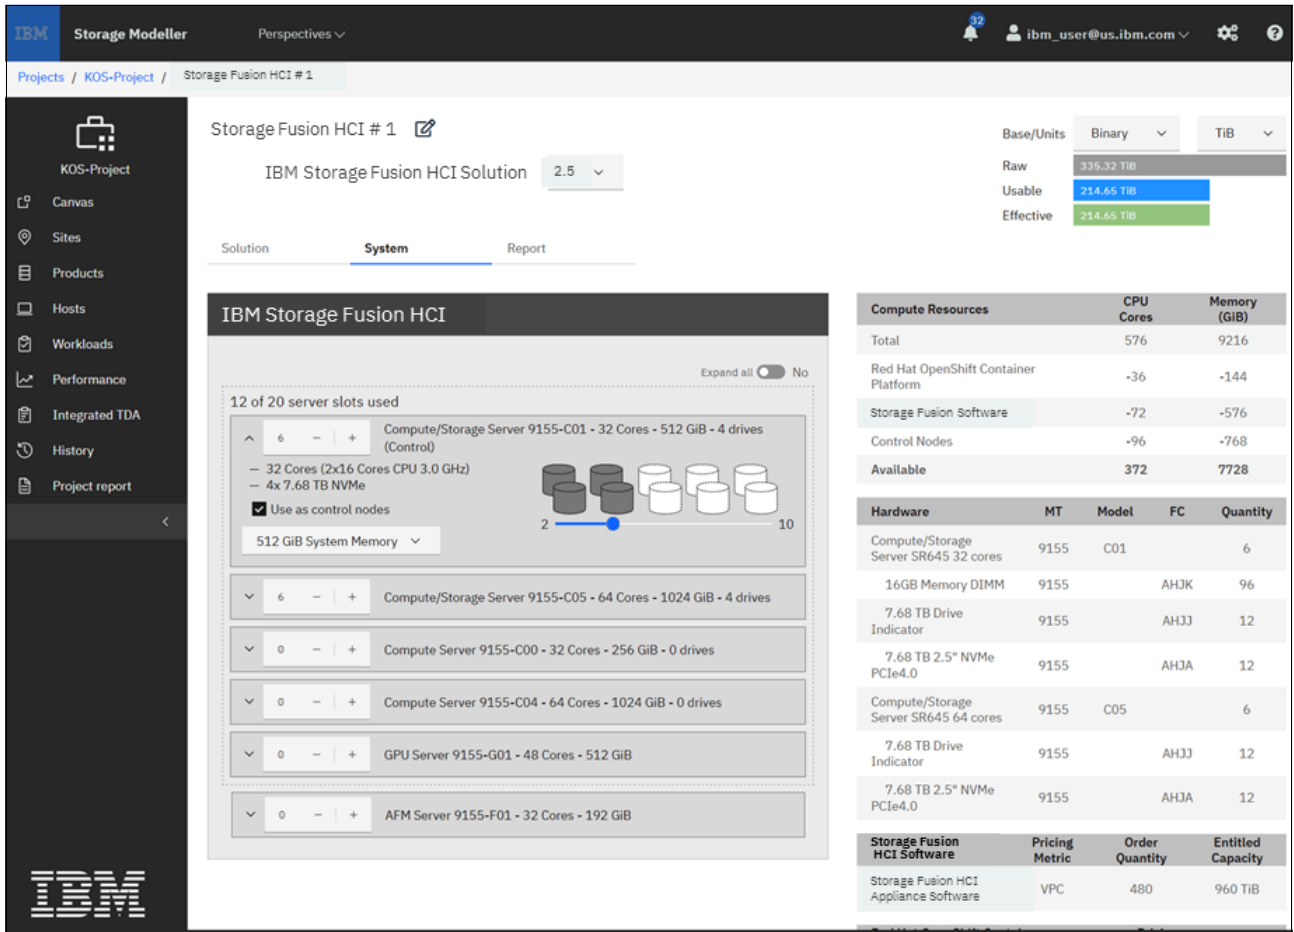


Figure 3-7 IBM StorM graphical output

The left hand section of the UI allows the user to select the number and type of servers, and storage per server, both the system memory and the number of NVMe drives (minimum of two per server to a maximum of ten per server).

Note: The number of NVMe drives must be symmetrical across all of the six base servers and any additional servers added. The StorM tool will check for this rule.

The right hand section of the UI shows the license information, raw vs. usable storage and the selected hardware resources chosen with the relevant IBM feature codes.

3.4.7 VMware versus bare metal

IBM Storage Fusion HCI System runs Red Hat OpenShift on bare metal servers as part of the appliance delivery model. This ensures that the overhead on the VMware hypervisor costs are eliminated.

For more information on the details of each type of installation available, refer to the IBM Document pages for IBM Storage Fusion here:

IBM Documentation for [Deploying IBM Storage Fusion](#)

Unlike IBM Storage Fusion HCI System, the IBM Storage Fusion Software solution does not come with any hardware, so the customer can choose the x 86 hardware platform they want to run on. IBM Storage Fusion Software currently runs VMware hypervisor only. Future plans include running on bare metal and a variety of cloud provider platforms, including of course IBM Cloud.

Figure 3-8 shows the variations on the IBM Storage Fusion HCI System and Software only versions and which platform each will run on.

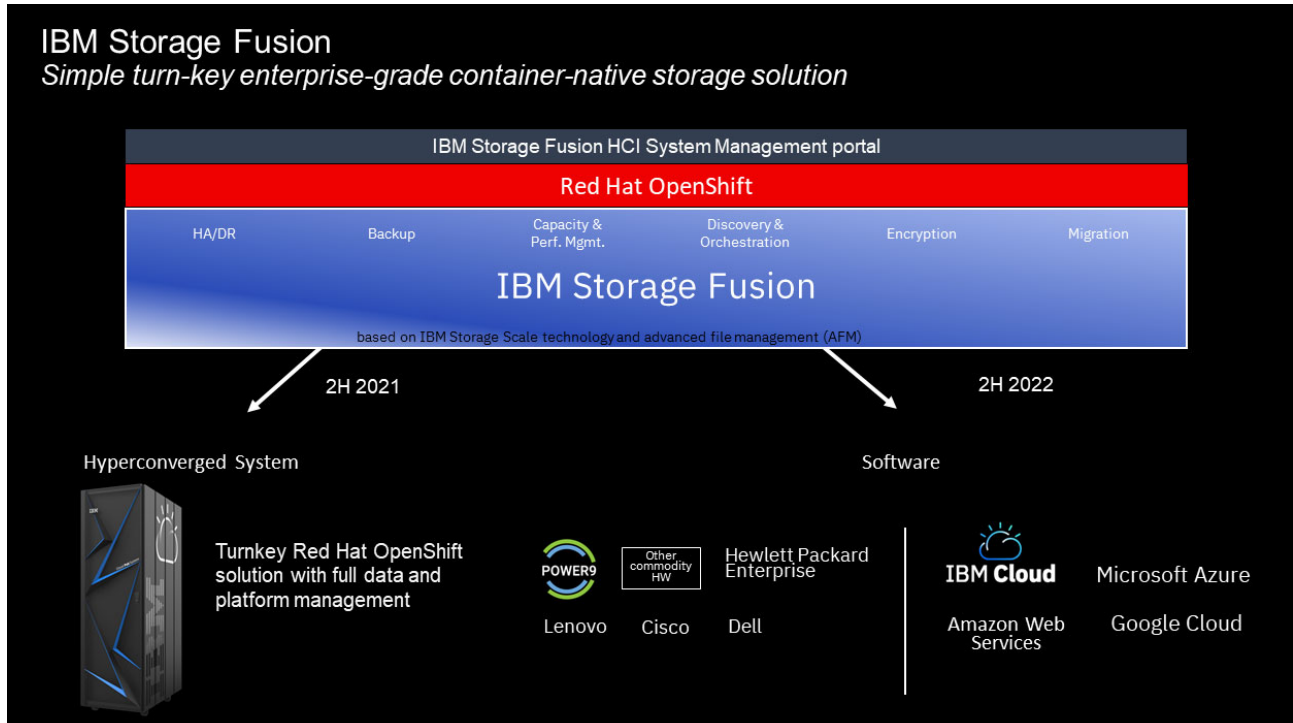


Figure 3-8 IBM Storage Fusion flavors

The components within each IBM Storage Fusion offering are as follows:

IBM Storage Fusion HCI

IBM Supplied hardware as follows:

- ▶ Minimum of six compute servers with internal storage,
- ▶ Two switches for high-speed data flow
- ▶ Two switches for management control
- ▶ Six power distribution units (PDU)
- ▶ IBM 42U approved rack

IBM Storage Fusion Software

The IBM Spectrum Fusion software only version at 2.3 and 2.4, and IBM Storage Fusion software 2.5 and above, have the ability to run on the following:

- ▶ VMware hypervisor as a guest
- ▶ Run on base metal via these hardware providers:
 - Lenovo
 - Cisco
 - Dell
 - Hewlett Packard (HPE)
 - Other commodity hardware

The IBM Spectrum Fusion software only version at 2.3 and 2.4, and IBM Storage Fusion software 2.5 and above, have the ability to run on the following:

- ▶ IBM Cloud
- ▶ AWS - Amazon Web Services
- ▶ IBM zLinux

3.4.8 Use cases for IBM Storage Fusion software

The main theme on IBM Storage Fusion software is for customers who need to run mission critical applications providing data services storage, with high availability (HA) and disaster recovery (DR) plus full backup and restore option. IBM Storage Fusion software can do all of this with a minimum of install time and can provide all these functions.

Also some other considerations are as follows:

- ▶ IBM Spectrum Storage for Cloud Paks (SS4CP) is no longer being offered, so the IBM Storage Fusion is now the preferred option
 - Customers who originally acquired SS4CP are now moving to IBM Storage Fusion software as an ongoing solution
- ▶ Customers who have purchased IBM Storage Scale and want to use it with OCP but CNSA+CSI is too complicated to deploy and operate, plus they also seek a method for full Backup and Recovery integration
- ▶ Official DR support and global data sharing will most likely also only come with IBM Storage Fusion
- ▶ Customer is expanding into Public Cloud and only the IBM Storage Fusion software version will bring official support there
- ▶ Client is building Development/Test/Pre-prod clusters along side a production level The IBM Storage Fusion HCI System

So in summary IBM Storage Fusion software allows you to solve the same problems BUT we meet you on your platform of choice.

One very common client scenario we see today is IBM Storage Fusion Data Foundation + Cloud Pak for Data and in particular IBM Storage Fusion providing the backup/restore capability for a complex Cloud Pak installation.



IBM Storage Fusion Terminology

This chapter describes the terminology used for containers environments and more specifically for IBM Storage Fusion. It also covers some more advanced functions included in the product.

There is also a section on what is included with the various releases of the IBM Storage Fusion software.

4.1 Operators

An Operator is a method of packaging, deploying, and managing a container-native application. A container-native application is an application that is both deployed on Kubernetes and managed using the Kubernetes APIs and tooling.

Note: Kubernetes (also known as k8s or “kube”) is an open-source container orchestration platform that automates many of the manual processes involved in deploying, managing, and scaling containerized applications.

Operators are pieces of software that ease the operational complexity of running another piece of software. They act like an extension of the software vendor's engineering team, watching over a Kubernetes environment (such as Red Hat OpenShift Container Platform) and using its current state to make decisions in real time.

These Operators allows for simple management of the IBM Storage Fusion HCI System appliance and allows IBM to build the appliance management software using the latest software development practices including CI/CD pipelines and online software delivery via entitled registries for simplified upgrades. IBM Storage Fusion HCI System also supports a fully air-gaped or disconnected installation.

4.2 Upgrading

IBM Storage Fusion HCI System incorporates an upgrade operator which orchestrates the upgrade process of firmware on the hardware, and application software of Red Hat OpenShift, the Appliance Management Stack, IBM Storage Scale ECE and IBM Storage Protect Plus.

4.3 Management UI

IBM Storage Fusion provides a management or appliance UI for managing the HCI System appliance. This UI has the ability to monitor and control every area of the appliance. The appliance UI operator provides a single pane of glass to manage IBM Storage Fusion HCI System and is integrated with Single Sign On provided by Red Hat OpenShift. The appliance UI operators monitor the health of the system for problem resolution, provide performance monitoring capabilities, present capacity utilization so that as a user you know when you need to add more drives for storage capacity or more physical servers for more compute capacity.

4.4 IBM Storage Scale Erasure Code Edition (ECE)

By default, when IBM Storage Fusion HCI System is installed, a best practice installation of IBM Storage Erasure Code Edition (ECE) is installed so that software deployed on Red Hat OpenShift can easily consume the storage using dynamic provisioning.

4.5 Red Hat OpenShift Virtualization

With Red Hat OpenShift Virtualization comes the ability to deploy virtual machines (VMs) side by side with containers using the same Red Hat OpenShift control plane.

4.6 IBM Storage Fusion Data Cataloging service

The [Data Cataloging](#) service is modern metadata management software that provides data insight for exabyte-scale heterogeneous file, object, backup, and archive storage on premises and in the cloud. It can help you manage your unstructured data by reducing data storage costs, uncovering hidden data value, and reducing the risk of massive data stores.

Note: IBM Spectrum Discover functionality is now available as an integrated service within IBM Storage Fusion called Data Cataloging.

If you want to use Data Cataloging, you must take into consideration the following system requirements: [IBM Storage Fusion system requirements](#)

After the system requirements for Data Cataloging are met, you can [install the Data Cataloging service](#).

Note: Data Cataloging is applicable only for IBM Storage Fusion and not for IBM Storage Fusion HCI System. It is enabled starting with IBM Storage Fusion 2.5.1 software.

After Data Cataloging is installed you will need to check [IBM Storage Software requirements](#) and then [configure the data source connections](#) to meet your needs such as the following examples:

IBM Storage Scale data source connections

You can create an IBM Storage Scale data source connection, scan a data source, and manually initiate a scan.

IBM Storage Archive data source connections

You can define tags and policies in Data Cataloging based on values that are derived from IBM Storage Archive metadata to help in searching and categorizing files.

IBM Cloud Object Storage data source connection

You can create the IBM Cloud Object Storage (COS) connection and initiate a scan.

Data Cataloging and S3 object storage data source connections

Use this information to understand how Data Cataloging works with S3-compliant object store.

Scanning an Elastic Storage Server data source connection

Use the Data Cataloging GUI to scan an Elastic Storage Server (ESS) data connection.

IBM Storage Fusion documentation for [Data Cataloging](#) will show you how to use the functions to manage your unstructured data by managing metadata and content search policies, managing tags, using managing applications, using the IBM Spectrum Discover applications catalog, reporting, and many more functions.

4.7 IBM Cloud Satellite

IBM Cloud Satellite extends the capabilities of IBM Cloud into virtually any location including the datacenter or at the edge. It comes with an integrated UI that can deploy IBM Cloud Satellite hosts in a matter of minutes. As your business needs grow, IBM Storage Fusion HCI System makes it easy to expand capacity from both a compute (CPU and memory) and storage perspective to accommodate more IBM Cloud Satellite workloads.

Once the IBM Storage Fusion HCI System is configured as an IBM Cloud Satellite location, from the IBM Cloud console you can then deploy IBM Cloud Services such as Red Hat OpenShift clusters and other IBM Cloud Satellite enabled services. IBM Storage Fusion HCI System in this case serves as a turnkey IBM Cloud Satellite location with integrated infrastructure, capacity management, system monitoring, call home, upgrades, and more.

4.8 Red Hat Advanced Cluster Management (RHACM)

Red Hat Advanced Cluster Management (RHACM) can manage multiple Red Hat OpenShift or other Kubernetes clusters and can manage the resources running on those clusters from a single pane of glass. Since IBM Storage Fusion HCI System is based on Red Hat OpenShift, RHACM can natively manage multiple IBM Storage Fusion HCI System systems. RHACM can run anywhere, including on an IBM Storage Fusion HCI System.

With RHACM you can:

- ▶ Manage Red Hat OpenShift and Kubernetes clusters from a single pane of glass.
- ▶ Manage the Kubernetes resources on these clusters and easily make changes to multiple resources on multiple clusters.
- ▶ Deploy applications to and across Red Hat OpenShift clusters including move an application from one cluster to another using placement policies.
- ▶ Create policies that enforce the configuration of Kubernetes resources across multiple Red Hat OpenShift clusters.

Note: Both IBM Cloud Satellite and Red Hat Advanced Cluster Management can be used together to build a hybrid cloud solution with single pane of glass management.

IBM Cloud Satellite is used to deploy IBM Cloud services to on-prem and edge locations with IBM Storage Fusion HCI System. Red Hat Advanced Cluster Management is used to manage hybrid cloud resources and applications across multiple IBM Storage Fusion HCI Systems, public clouds, and more.

4.9 IBM Storage Protect Plus

IBM Storage Fusion HCI System comes integrated with IBM Storage Protect Plus. IBM Storage Protect Plus provides backups of the IBM Storage Fusion HCI System control plane, as well as backups of running applications on Red Hat OpenShift.

IBM Storage Protect Plus can take snapshots of applications and save them locally or save them off the rack to an external backup location, such as an IBM Storage Protect Plus vSnap server, that is connected to on-prem or public cloud storage.

4.10 Call Home for all IBM Storage Fusion HCI System

IBM Storage Fusion HCI System has a fully integrated Call Home function which covers both the IBM Storage Fusion HCI System hardware infrastructure and software with automatic log collection and upload for rapid problem resolution. The call home function will take any information passed to it from the appliance UI that is deemed to be critical enough to require further analysis by IBM Support.

This call home function will connect securely to IBM and raise a CSP (Cognitive Support Platform) ticket with IBM support. This CSP ticket will then be passed to the relevant IBM support specialist to analyze and progress the call. Events are natively integrated with Red Hat OpenShift and can call home to IBM in the event of a physical infrastructure or software problem with automatic log collection and upload.

4.11 Active File Management (AFM)

Active file management (AFM) is a scalable, high-performance, file system caching layer, and enables sharing of data across clusters, even if the networks are unreliable or have high latency.

IBM Storage Fusion HCI System supports Active File Management (AFM) through the use of a pair of specialized AFM servers, ordered as a feature code.

The competitive differentiator is not just the fact that IBM has fused a storage platform with storage services but with the storage platform we have leveraged years of investment in a global parallel file system with global data access and advanced file management (AFM).

Applications simply access a local file or directory but are accessing data from information that is stored in a public cloud bucket or legacy data from a remote system 1000s of miles away. Some of the benefits of AFM are listed here:

- ▶ You can use AFM to create associations between IBM Storage Fusion clusters.
- ▶ With AFM, you can implement a single namespace view across sites around the world by making your global namespace truly global.
- ▶ By using AFM, you can build a common namespace across locations, and automate the flow of file data. You can duplicate data for disaster recovery purposes without suffering from wide area network (WAN) latencies.
- ▶ Individual files in the AFM filesets can be compressed. Compressing files saves disk space.
- ▶ Namespace replication with AFM occurs asynchronously so that applications can operate continuously on an AFM fileset without network bandwidth constraints.

The active file management (AFM) feature of IBM Storage Fusion continues to expand its data protection configuration options for disaster recovery (DR) scenarios by introducing the IBM Storage Scale feature AFM-based disaster recovery (AFM async DR).

- ▶ AFM async DR uses asynchronous data replication between two sites. The replication mechanism is based on the AFM file set level replication DR capability with a strict one-to-one active-passive model.
- ▶ In this model, the sites are represented as being primary and secondary, and the replicated file sets are mapped one-to-one.
- ▶ The active file management masks wide area network (WAN) latencies and outages by using IBM Storage Scale to cache massive data sets and provide asynchronous data movement between cache and home.
- ▶ The AFM async DR feature is focused on accomplishing business recovery objectives and targeting disaster recovery scenarios by replicating all data asynchronously from a primary site to a secondary site.
- ▶ From a disaster recovery (DR)/ high availability (HA) perspective, IBM supports complete file system, file set, and file data protection capabilities. Both synchronous and async replication are supported. Multi-zones are supported.
- ▶ Snapshots for end-to-end and provide a fast, easy way to protect data locally from corruption, modification, or deletion. Additionally, IBM Storage Fusion provides direct access to application checksum in ESS or IBM Storage Scale ECE.

All these capabilities put together provide not just resilient data but also location aware and location independent access to data across geographic distance, multi-cluster distance, and multi-availability zones. Distance is very important and is a requirement for a lot of enterprise workloads. Consequently, data resiliency has become increasingly important as cyber resiliency is part of board level discussions.

4.12 AI Support (AI GPU)

Artificial intelligence (AI) and machine learning (ML) require considerable elasticity and dispersal of compute and storage requirements to gather data, present it to the data scientists who refine data and build models for training, as well as production purposes.

Furthermore, there are lifecycles in which those models are assessed and improved on a continuous basis as follows:

- ▶ The training datasets and production models all need easy to manage storage provisioning to support the pipelines that are now prevalent in both DevOps and AI continuous delivery environments.
- ▶ Convergence of applications and infrastructure is threefold; computing power for AI/ML with graphic processing unit (GPU) assist, and co-location of data for training and high performing production models.

IBM Storage Fusion HCI System can be used to optimize AI workflows with modern container applications. It does this through the use of a pair of specialized GPU servers ordered as a feature code.

If a customer is using NVIDIA for AI IBM Storage Scale it can preprocess or process at the edge with NVIDIA A100 GPUs and NVIDIA Red Hat Operator and tools and then transparently leverage that data in the cloud or directly on another larger more compute intensive NVIDIA DGX A100 system with the IBM ESS 3000 or ESS 3200 high performance storage system.

Advantages of this are:

- ▶ Accelerate AI/ML workloads with NVIDIA A100 GPU Nodes
- ▶ Process data generated at the edge, flow 'the interesting bits' to a central location for further processing
- ▶ Train ML models from data in IBM Storage Fusion HCI System's high-performance NVMe Flash storage

So, in summary the optional pair of GPU accelerated nodes for AI workloads that leverage the latest NVIDIA A100 GPU, can also leverage the latest Red Hat OpenShift Operator for AI from NVIDIA.

This will make creating and running AI workloads on IBM Storage Fusion HCI System even easier and shows how this can be integrated into an overall AI enhanced workflow.

4.13 Synchronous Storage Replication with Metro-DR

You can have a stretched IBM Storage Scale cluster by using which the file system is made available in two availability zones of a data center or across multiple data centers within a metro area.

These two independent clusters run active-active so the container applications can be moved in case of failure. As an Red Hat OpenShift admin, you can deploy IBM Storage Fusion HCI System in data center site 1 and data center site 2, resulting in Red Hat OpenShift Container Platform clusters that run at each site. Two clusters are configured such that they share a synchronously replicated storage layer, allowing for data mirroring with zero Recovery Point Objective (RPO) between the two clusters.

IBM Storage Fusion ensures data availability between the two clusters (both the data and the PV/PVCs). However, you are responsible for managing application deployments between the cluster. Here, both Red Hat OpenShift Container Platform clusters host active workloads at the same time, rather than requiring one Red Hat OpenShift Container Platform cluster to be a cold standby.

The Metro-DR setup includes two sets of networks, namely internal storage network and Red Hat OpenShift network. In Metro-DR, the internal storage network of one site needs to connect to the other site and tiebreaker, that is, externalize your storage network. Similarly, the Red Hat OpenShift network of one site needs to connect to the other site.

Figure 4-1 shows the connectivity between Site 1 and Site 2.

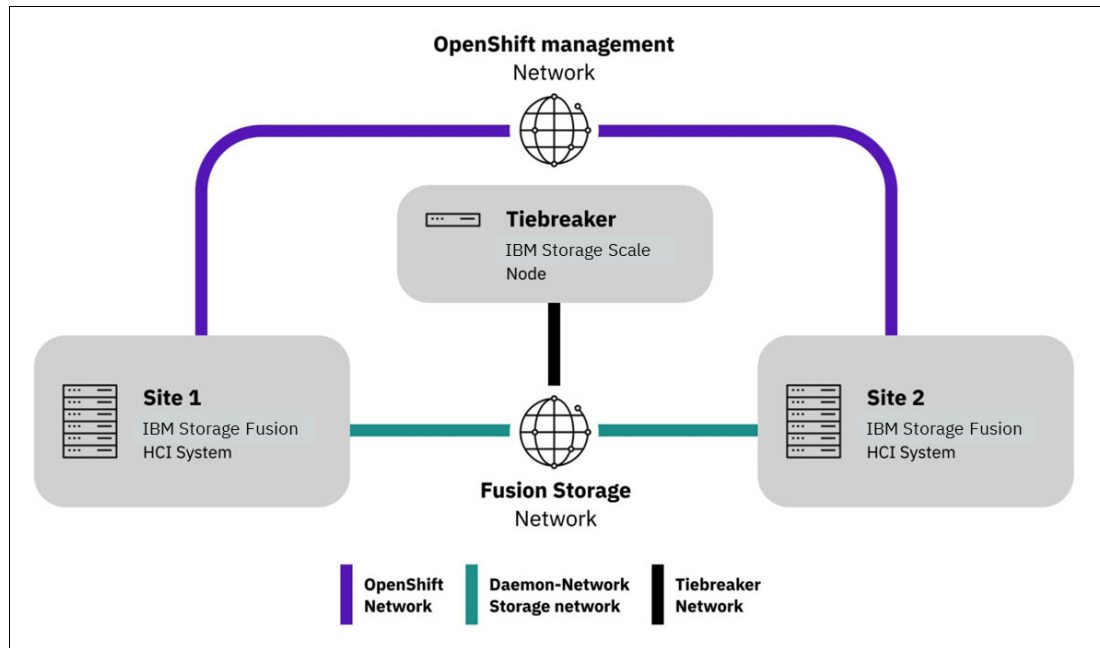


Figure 4-1 Connectivity between Site 1 and Site 2

For more info see here: [Metro-DR](#)

4.14 IBM Storage Fusion Reference Information and Links

The IBM Documentation Center has a listing of which features, and changes are applicable to each level of IBM Storage Fusion software release. For the most recent release check [here](#).

Here are the details of each release and links to the IBM Documentation Center for more detailed information.

4.14.1 New features and changes in 2.5.2

Listed [here](#) are some of the new features in IBM Storage Fusion 2.5.2

- ▶ High Availability Cluster Support
 - IBM Storage Fusion HCI System lets you connect three IBM Storage Fusion HCI System racks together to create Red Hat OpenShift Container Platform clusters.
- ▶ Expansion Rack Support
 - IBM Storage Fusion HCI System can add new IBM racks to expand your Red Hat OpenShift Container Platform cluster.
- ▶ IBM Storage Fusion HCI Components in Your Own Rack
 - Install IBM Storage Fusion HCI System components in your own rack
 - IBM Storage Fusion HCI System can be purchased without an IBM rack and instead is installed in a rack that you provide in your data center.

- ▶ New supported component versions in IBM Storage Fusion 2.5.2
 - Red Hat OpenShift Container Platform 4.12 for IBM Storage Fusion and Red Hat OpenShift Container Platform 4.10 for IBM Storage Fusion HCI System
 - IBM Storage Fusion Data Foundation 4.12
 - IBM Spectrum Scale 5.1.7.0
 - IBM Spectrum Scale Container Native 5.1.7.0 for IBM Storage Fusion HCI
 - Data Cataloging 2.1.0
 - Backup & Restore 2.5.2
 - For supported IBM Cloud Paks versions, see [IBM Cloud Paks support for IBM Storage Fusion](#)

4.14.2 New features and changes in 2.5.1

Note: The 2.5.1 release is only for IBM Storage Fusion and not for IBM Storage Fusion HCI System. The IBM Storage Fusion HCI System features are picked up in release 2.5.2.

Listed [here](#) are some of the new features in IBM Storage Fusion 2.5.1.

- ▶ Dynamic storage provisioning using IBM Storage Fusion Data Foundation
 - You can configure Dynamic storage device from the IBM Storage Fusion UI for IBM Cloud and VMware platforms. Support is also available to configure IBM Storage Fusion Data Foundation encryption from the IBM Storage Fusion user interface.
- ▶ New platform support for Global Data Platform service
 - Global Data Platform service is now supported on ROSA platform. You can connect and add remote file system from the IBM Storage Fusion user interface.
- ▶ Policy frequency can be realized in units of minutes or hours with the new service
 - Multiple policies can be assigned to single application.
- ▶ Data Cataloging as an IBM Storage Fusion service
 - IBM Storage Fusion Data Cataloging service is available in IBM Storage Fusion. It is a container native modern metadata management software that provides data insight for exabyte-scale heterogeneous file, object, backup, and archive storage on premises and in the cloud. The software easily connects to these data sources to rapidly ingest, consolidate, and index metadata for billions of files and objects. For more information about the service, see [Data Cataloging](#).
- ▶ New Backup and Restore IBM Storage Fusion service
 - The new Backup & Restore service in IBM Storage Fusion allows you to backup and restore applications. It includes the following additional benefits in comparison with Backup & Restore (Legacy):
 - Policy frequency can be realized in units of minutes or hours with the new service
 - Multiple policies can be assigned to single application.
 - There is more visibility of job progress and error information with the new service.
 - Greater resiliency and scale-out capabilities.
 - For more information about the service, see [Backup and Restore](#).

- ▶ Add IBM Storage Protect as a backup storage location through S3 option
 - IBM Storage Fusion allows you to backup and recover Red Hat OpenShift data to an IBM Storage Protect server with the S3 option. For more information about backing up to IBM Storage Protect, see [Using IBM Storage Protect as a backup storage location](#).
- ▶ New supported component versions in IBM Storage Fusion
 - Red Hat OpenShift Container Platform 4.12 for IBM Storage Fusion
 - IBM Storage Fusion Data Foundation 4.12
 - IBM Spectrum Scale 5.1.7.0
 - Data Cataloging 2.1.0
 - Backup & Restore 2.5.1
 - For supported IBM Cloud Paks versions, see [IBM Cloud Paks support for IBM Storage Fusion](#)

4.14.3 New features and changes in 2.4

Listed [here](#) are some of the new features in IBM Spectrum Fusion 2.4:

- ▶ Support for the IBM Spectrum Fusion HCI introduces 64-core servers (9155-C05) as an alternative to the existing 32-core servers.
- ▶ Improvements to Metro Sync DR
- ▶ Ability to deploy IBM Spectrum Fusion on IBM Cloud, Microsoft Azure, or zLinux. For more information about these deployments, see [IBM Spectrum Fusion deployments](#).
- ▶ Scaling out Red Hat OpenShift Data Foundation and detecting an existing Red Hat OpenShift Data Foundation configuration. For more information about Red Hat OpenShift Data Foundation service and storage configurations, see [Red Hat Advanced Cluster Management for Kubernetes](#).
- ▶ Mirroring IBM Spectrum Fusion images to your enterprise registry. For more information about mirroring your images to your enterprise registry, see [Enterprise registry for IBM Spectrum Fusion HCI installation](#).

Product versions in IBM Spectrum Fusion HCI and IBM Spectrum Fusion:

- ▶ Red Hat OpenShift Data Foundation 4.10
- ▶ IBM Spectrum Scale 5.1.6
- ▶ IBM Spectrum Scale Container Native 5.1.6 for IBM Spectrum Fusion HCI
- ▶ On-premises VMware supported version is vSphere 7.0
- ▶ IBM Spectrum Protect Plus 10.1.13
- ▶ IBM Spectrum Discover 2.0.6

For supported IBM Cloud Paks versions, see [IBM Cloud Paks support for IBM Spectrum Fusion](#).

4.14.4 New features and changes in 2.3

Listed here are some of the new features in IBM Spectrum Fusion 2.3:

- ▶ IBM Spectrum Fusion HCI introduces 64-core servers (9155-C04) as an alternative to the existing 32-core servers.
- ▶ Upgrade memory for the new 64-core servers.
 - These servers come with base memory of 1024 GB of RAM upgradable to 2048 GB of RAM.
- ▶ Mirror images to multiple registries
 - IBM Spectrum Fusion supports installation from multiple enterprise registries. It gives you the flexibility to mirror images to individual registries of your choice.
- ▶ IBM Spectrum Fusion HCI Metro Sync DR (Disaster Recovery) between two IBM Spectrum Fusion HCI OpenShift Container Platform clusters.
 - IBM Spectrum Fusion HCI provides a unique active-active resiliency across data centers that are up to 150km distance.
- ▶ For more information about Metro Sync DR, see
 - [Metro Sync DR](#)
- ▶ Upgrade of IBM Spectrum Fusion HCI node and switch firmware
 - To know more about compute node and switch firmware upgrades, see here:
 - [Upgrading the compute nodes](#)
 - [Upgrading the switch firmware](#)
- ▶ Encryption option is available in the user interface for IBM Spectrum Fusion HCI. For more details about the user interface page, see
 - [Setting encryption for storage](#)
- ▶ Call Home support is included in IBM Spectrum Fusion and you can enable it from the IBM Spectrum Fusion user interface. For steps to enable from the user interface, see
 - [Enabling Call Home](#)
- ▶ Serviceability Improvements
- ▶ Supported product versions
 - Red Hat OpenShift Container Platform 4.10
 - IBM Spectrum Scale 5.1.5
 - On-premises VMware supported version is vSphere 7.0
 - IBM Spectrum Protect Plus 10.1.12
 - IBM Spectrum Discover 2.0.5 is available in the Operator Hub catalog as a Container software.
- ▶ For a more detailed listing of all the new features in IBM Spectrum Fusion 2.3, see the link here:
 - [New Features and Changes in IBM Spectrum Fusion 2.3](#)

4.14.5 New features and changes in 2.2.1

Listed here are the new features in IBM Spectrum Fusion 2.2.1:

- ▶ This release includes defect fixes and security enhancements for IBM Spectrum Fusion and IBM Spectrum Fusion HCI.
- ▶ Setup IBM Spectrum Fusion HCI in a restricted network. For more information about mirroring your images to the enterprise registry, see [Mirroring your images to the enterprise registry](#).

Supported product versions:

- ▶ List of product versions included in 2.2.1:
 - Red Hat OpenShift Container Platform 4.10 for IBM Spectrum Fusion. This support is not applicable for IBM Spectrum Fusion HCI.
 - IBM Spectrum Protect Plus 10.1.10.2
 - IBM Spectrum Scale 5.1.3.1

For more information see link here:

[New Features and Changes in IBM Spectrum Fusion 2.2.1](#)

4.14.6 New features and changes in 2.2

Listed here are the new features in IBM Spectrum Fusion 2.2:

- ▶ Converged user interface for IBM Spectrum Fusion HCI and IBM Spectrum Fusion
- ▶ Support for IBM Cloud Paks
- ▶ Application backup using IBM Spectrum Fusion

IBM Spectrum Fusion HCI

- ▶ Enable purchase of network transceivers, cables, and spine switch
- ▶ Memory expansion options
- ▶ Support to upgrade IBM Spectrum Fusion HCI components

IBM Spectrum Fusion

- ▶ IBM Spectrum Fusion operator
- ▶ Storage options for IBM Spectrum Fusion

For more information see link here:

[New Features and Changes in IBM Spectrum Fusion 2.2](#)

4.14.7 New features and changes in 2.1.2

Listed here are the new features in IBM Spectrum Fusion 2.1.2

- ▶ Enhanced Active file management (AFM)
- ▶ Container Network Interface (CNI) daemon network
- ▶ Support for Red Hat OpenShift Container Platform
- ▶ Support for IBM Spectrum Scale Erasure Code Edition (ECE)

- ▶ Support for IBM Spectrum Protect Plus
- ▶ Support for proxy server

For more information see link here:

[New Features and Changes in IBM Spectrum Fusion 2.1.2.](#)

4.14.8 IBM Storage Fusion Useful Links

- ▶ IBM Storage Fusion Documentation Main Page
<https://www.ibm.com/docs/en/storage-fusion>
- ▶ IBM Storage Fusion 2.5 Support Portal
<https://www.ibm.com/support/pages/node/6966946>
- ▶ IBM Spectrum Fusion 2.4 Support Portal
<https://www.ibm.com/support/pages/node/6842171>
- ▶ IBM Spectrum Fusion 2.3 Support Portal
<https://www.ibm.com/support/pages/node/6826715>
- ▶ IBM Spectrum Fusion Passport Advantage®
<https://www.ibm.com/software/passportadvantage>

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- ▶ *IBM Storage Fusion HCI System Metro Sync Dr Use Case*, REDP-5708
- ▶ *IBM Storage Fusion Backup and Restore for Cloud Pak for Data*, REDP-5706
- ▶ *IBM Spectrum Scale Security*, REDP-5426
- ▶ *Multi-Factor Authentication Using IBM Security Verify for IBM Spectrum Fusion*, REDP-5662

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- ▶ IBM Storage Fusion documentation
<https://www.ibm.com/docs/en/storage-fusion>
- ▶ IBM Spectrum Fusion 2.4 Support Portal
<https://www.ibm.com/support/pages/node/6842171>
- ▶ IBM Spectrum Fusion 2.3 Support Portal
<https://www.ibm.com/support/pages/node/6826715>
- ▶ IBM Spectrum Fusion Passport Advantage
<https://www.ibm.com/software/passportadvantage>
- ▶ IBM Storage Modeller
https://www.ibm.com/tools/storage-modeller/help/index.html#page/WR2/5528_intro.IBM_Storage_Modeller.html
- ▶ Metro-DR
<https://www.ibm.com/docs/en/storage-fusion/2.5?topic=prerequisites-metro-dr-disaster-recovery>
- ▶ Red Hat OpenShift Container Platform security and compliance
<https://docs.openshift.com/container-platform/4.11/security/index.html>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



REDP-5688-00

ISBN 0738460990

Printed in U.S.A.

Get connected

