




“Cloud on Premise with Responsible AI “Now Sovereign AI” in a Box”




An engineered infrastructure platform for modern client problems

IBM Fusion

Making OpenShift easier

Fusion Data Services



Ishwar Fernandes (Subject Matter Expert)

Abhishek Jain (Software Engineer)



Agenda

Fusion HCI Appliance

- Hardware and Options
- Install and update processes
- Monitoring

Fundamentals

- Containers
- Kubernetes
- OpenShift

Customer Feedback

Organizations are challenged by the complexity and cost of delivering modern applications and generative AI on-premises


Modern problems require a modern infrastructure platform

Only IBM Fusion on OpenShift addresses these challenges

Generative AI

Cloud-Native Applications

Modern Infrastructure



HCI Building Blocks (Try do it yourself)

Hardware

Software

Capabilities

40%

5 Days Deployment

Choosing HW, Getting Datacenter and Pre-requisites Ready



Red Hat



OPENSIFT

watsonx

 Red Hat

Generative AI Solutions



OPENSIFT

App Modernization



Storage



Rack

30%

5 Days Deployment

Choosing the right OS and best practices for Software Deployment



OPENSIFT

Virtualization


VM Migration



AZ 1



AZ 2





IBM
Storage Scale

30%

25 Days Deployment

Deciding the right use case and getting this working with best performance required

30%

30 Days High-Level Design & Low-level Design

Testing Application in DR & HA

Project Activities that HCI saves you time on: (Hidden

1. Produce high level design on which basis identify hardware, GPUs, specifications for servers, rack layout and power consumption, heat dissipation and network switches (Leaf and Spine, including High Speed setup).

2. Produce low-level implementation design including: Network design with leaf and spine setup (Approx. 6 weeks)

- Storage design
- Reliable HA and DR topology
- Security and key management
- Secure Deployment Design
- Best practices for performance with selected use cases (e.g. as documented in IBM Fusion blueprints)

3. Acquiring Selections, Delivery of Racks, Configure servers with built-in cabling with labels, network switches and storage (if external) (Approx. 8 Weeks)

4. Test connectivity and any additional testing for performance and availability (e.g. server or network security) (Approx. 4 Weeks)

5. Install OpenShift according to design in 2 with, Secure Deployment, baseline hardening. (Approx. 4 Weeks)


6. Conduct system tests before ready for production use: (Approx. 4 Weeks)

- Test network and storage configuration with OpenShift
- Test HA (Metro DR or Regional DR) with synchronous or asynchronous replication
- Test DR (application failover) with synchronous or asynchronous replication

7. Testing Applications in HA and DR configuration, performance testing and availability before being ready for production can take several months on its own. (Approx. 4 Weeks)

8. Customer acceptance testing and handover:

Project Activities that HCI saves you time on: (Hidden Backend TCO)



IBM Fusion Built for Performance as per Use case:

Configure hardware details

☐ Client will provide physical rack(s)

☐ Add spine switches

☒ Add rackmount console kit

☐ Add arbiter node

☐ Watsonx Solution

☐ DB2 Warehouse Solution

☐ Enable Fusion HCI to be used as a Control Tower

Number of racks

3

HA layout

3-zone

Spine switches

2 spine switches

Rack #1

Rack #2

Rack #3

16 of 16 server slots used

1

Control nodes

Settings

Worker

Cores

32

Memory(GiB)

256

2

32-core storage nodes

13

64-core storage nodes

0

32-core compute-only nodes

0

64-core compute-only nodes

0

9155-G03 GPU Nodes

0

9155-G04 GPU Nodes

Storage

Drive Type

7.68TB NVMe PCIe5.0

Drive Quantity

10

Configuration

Rack #1

– 1x 32-core Control nodes

– 2x 32-core Storage nodes

– 13x 64-core Storage nodes

– Usable cores: 643 VPC, 1286 vCPU

– Usable memory: 12716 GiB

– Usable capacity: 409.65 TB

– Power: 14.44 kW

– Weight: 1775.98 lbs, 807.26 kg

Rack #2

– 1x 32-core Control nodes

– 2x 32-core Storage nodes

– 13x 64-core Storage nodes

– Usable cores: 643 VPC, 1286 vCPU

– Usable memory: 12716 GiB

– Usable capacity: 409.65 TB

– Power: 14.10 kW

– Weight: 1735.38 lbs, 788.81 kg

Rack #3

– 1x 32-core Control nodes

– 2x 32-core Storage nodes

– 13x 64-core Storage nodes

– Usable cores: 643 VPC, 1286 vCPU

– Usable memory: 12716 GiB

– Usable capacity: 409.65 TB

– Power: 14.10 kW

– Weight: 1735.38 lbs, 788.81 kg

Total

– 3x 32-core Control nodes

– 6x 32-core Storage nodes

– 39x 64-core Storage nodes

– Usable cores: 1929 VPC, 3858 vCPU

– Usable memory: 38148 GiB

– Usable capacity: 1228.95 TB

– Power: 42.64 kW

– Weight: 5246.74 lbs, 2384.88 kg

Cancel

Apply Configuration

10

Clients buy Fusion for Day 1, Day 2, & TCO benefits that make it easy to quickly deliver a cloud operating model for container and VM apps

Day 1

76%

Deploy OpenShift Easier, Faster, Better

48%

Accelerate App Modernization

43%

Accelerate IBM Cloud Pak deployments

Day 2

67%

Backup/Recovery HA/DR

48%

Address skills & resource gaps

43%

Leverage enterprise-grade architecture

Integrated Simplicity

43%

Complete solution Cloud Pak, ISV, OpenShift, Fusion

1 Call Support

71%

One vendor to call – IBM


TCO

52%

Slash total cost of ownership



Why Fusion HCI?




AI

IBM watsonx

RHEL AI


Inferencing at scale



Containers

Application modernization

Lift and shift from OpenShift running on top of VMware



Virtualization

Migration from VMware to OpenShift Virtualization

- Control licensing costs
- Declarative management
- Infrastructure/skills consolidation

Planning / architecture

Size/plan/build/validate an infrastructure architecture for the target application workloads

Application availability

Architect for high availability, backup, and disaster recovery in the event of failure or data loss

Velocity

Keep up with the demands of the business by providing new app environments

Operations / scaling

Quickly adding additional infrastructure in time for growing workloads and requirements

Platform maintenance

Observability / serviceability / upgrades for the full platform stack

This is why Fusion HCI?

An engineered infrastructure platform optimized for OpenShift, VM and AI workloads

Opinionated hyperconverged architecture

Automation, validation, hardware redundant infrastructure

Compute/NVMe storage nodes




32 or 64 core
256/512/1024/2048 GB memory

GPU nodes



Up to 8 per node
L40S or H200

Dedicated High speed storage network



2x 200 GbE redundant storage network

Optional AFM data accelerator nodes



Accelerate remote queries, S3 object stores, and/ or Azure Blob
* Requires IBM Storage Scale

Optional IBM Storage Ceph Nodes



96/144/192/240 TB raw
HDD for data, SSD for metadata
* External connection only




watsonx

Accelerate Generative AI Solutions

Inferencing, Governance, summarization, fine tuning

Data Lakehouse modernization




OPENSIFT

Deploy in days
vs months

App Modernization at your own pace

IBM Cloud Paks, ISV, Maximo or custom apps

Built-in HA/DR, Backup/Recovery



OPENSIFT

Virtualization

VM Migration

Removing vSphere/vSAN dependency with container native storage, Landing Zone for Containers

Built-in HA/DR, Backup/Recovery

Hybrid Cloud

Run your workloads wherever they need to run


- Choose the platform of your choice
- Cloud vendor has changed
- Bursting from on-prem to the cloud
- Repatriating to the private data center from the cloud for data sovereignty



IBM Cloud



IBM Power



IBM System Z



Bare metal x86




Public cloud

On-prem

True Hybrid Cloud

OpenShift makes hybrid cloud easy by providing a consistent experience wherever it runs



Rack reset (2.13 NEW RELEASE)

Service node experience for resetting Fusion HCI to factory defaults or starting over with OpenShift

Reseting Racks

Validating environment

Checking connectivity and number of racks

Stopping workloads

Graceful shutdown VMs / OpenShift pods

Disconnecting network switches

Isolate racks from network

Resetting storage and hardware

Clear local storage and NVMe drives

Resetting cluster configuration

Delete OpenShift cluster metadata

Restoring factory defaults on racks

Apply default firmware settings

Rebooting racks

Restart each rack

Search logs

Jul 29 14:30:48.858 : Error: RU 5 interface ens1f0np0 is not connected to the correct port

Jul 29 14:30:48.858 : Action: read the labels on the network cable and rewiring. Work with MGEN and factory team to review and fix if the issue persists.

Jul 29 14:30:50.565 : Error: RU 5 interface ens3f1np1 is not connected to the correct port

Jul 29 14:30:50.565 : Action: read the labels on the network cable and rewiring. Work with MGEN and factory team to review and fix if the issue persists.

Jul 29 14:30:52.229 : Error: RU 5 interface ens1f1np1 is not connected to the correct port

Jul 29 14:30:52.229 : Action: read the labels on the network cable and rewiring. Work with MGEN and factory team to review and fix if the issue persists.

Jul 29 14:30:52.229 : Error: Wiring problems found on node RU5

Jul 29 14:30:52.229 : Action: check cables are securely inserted to the cages and check physical damages on cables and transceivers. Work with MGEN and factory team to review and fix if the issue persists.

Jul 29 14:30:52.732 : Checking connection on node at RU6

Jul 29 14:30:52.732 : Error: RU 6 interface ensf0 is not connected to the correct port

Jul 29 14:30:54.436 : Error: RU 6 interface ens3f0np0 is not connected to the correct port

Jul 29 14:30:54.436 : Action: read the labels on the network cable and rewiring. Work with MGEN and factory team to review and fix if the issue persists.

Jul 29 14:30:56.082 : Error: RU 6 interface ens1f0np0 is not connected to the correct port

Jul 29 14:30:56.082 : Action: read the labels on the network cable and rewiring. Work with MGEN and factory team to review and fix if the issue persists.

Jul 29 14:30:57.791 : Error: RU 6 interface ens3f1np1 is not connected to the correct port

Jul 29 14:30:57.791 : Action: read the labels on the network cable and rewiring. Work with MGEN and factory team to review and fix if the issue persists.

Jul 29 14:30:59.430 : Error: RU 6 interface ens1f1np1 is not connected to the correct port

Jul 29 14:30:59.430 : Action: read the labels on the network cable and rewiring. Work with MGEN and factory team to review and fix if the issue persists.

Jul 29 14:30:59.430 : Error: Wiring problems found on node RU6

Jul 29 14:30:59.430 : Action: check cables are securely inserted to the cages and check physical damages on cables and


Retry

What NVidia at GTC thinks about IBM



Keynote ▾Session CatalogExplore ▾Attend ▾Sponsors & ExhibitorsMore ▾

Join Now | Log In



"Working with IBM and NVIDIA, early results show we can refresh global operations data *in minutes* at reduced cost—making a capability that can be turned into *tangible business impact* in areas such as manufacturing or warehousing."

— Chris Wright, Chief Information and Digital Officer of Nestlé

IBM watsonx.data

Data Engines:
prestoVeloxSpark

NVIDIA cuDF

Open Formats:
ICEBERGParquet

Data Governance and Access
Control Throughout

NVIDIA A100

IBM Storage Scale

Order-to-Cash Data Mart
Updates 3 min vs 15 min

1,0007505002500

CPUGPU

83% Cost Savings

MORE VIDEOS

NVIDIA

Catch all the highlights and announcements from NVIDIA CEO Jensen Huang's keynote. He'll cover the latest advancements across the full AI stack, from accelerated computing and AI factories to open model-centric systems and physical AI that are setting the industry's direction.

Ask AI




AI Platform

Get started fast and
iterate fast

- Prepare for Inferencing
 - minimize data movement
- Container Aware Storage (CAS)




Cloud on Prem
HA/DR
Container Ready






Cloud on Prem
Container Ready
Business Continuity
HA/DR
Backup/Restore




Cluster Vending Machine
Empower developers to
innovate faster



Cloud on Prem
HA/DR
Backup/Restore




- VM Migration platform
- Landing Zone
- Accelerate Modernization






AI Platform

Get started fast and
iterate fast



VM Migration platform

Accelerate
modernization



Cluster Vending Machine

Empower developers to
innovate faster



Introduction to containers and OpenShift Cloud Platform

Containers


Containers Enable Immutable Delivery

Build once... run
anywhere

- The same binary that a developer runs on their laptop, runs in production
- All dependencies are package in the container
- Facilitates rolling updates with immediate roll-back
- Consistency limits side-effects

The perfect storm

Cultural Change
Automated Pipeline
Everything as Code
Immutable
Infrastructure



DevOps

Microservices

AGILITY

Containers

Loose Coupling/Binding
RESTful APIs
Designed to resist
failures
Test by break / fail fast

Portability
Developer
Centric
Ecosystem
enabler
Fast startup

The perfect storm: Containers

Cultural Change
Automated Pipeline
Everything as Code
Immutable
Infrastructure

DevOps

Microservices

Loose Coupling/Binding
RESTful APIs
Designed to resist
failures
Test by break / fail fast

Containers

Portability
Developer
Centric
Ecosystem
enabler
Fast startup

Containers

Docker and Podman are examples of light-weight container services that run on Linux

- File system overlay
- One Process Space
- One Network Interface
- Shares the Linux kernel with other containers and processes

Containers encapsulate a run-time environment

- Almost no overhead
- Containers spin up in seconds not minutes like VMs
- Native performance because there is no emulation
- Package only what you need



Put another way

Docker Containers allow you to control


What resources a process can **see**

What resources a process can **control**


What filesystem a process **uses**

Benefits of Containers

- Great isolation
- Great manageability
- Container encapsulates implementation technology
- Efficient resource utilization
- Fast deployment




Containers are not new




Containers are just Linux Capabilities Under the Covers

Virtual Machines




- Containers are:
 - Linux® processes
 - with isolation and resource confinement
 - that enable you to run sandboxed applications
 - on a shared host kernel
 - using cgroups, namespaces, and chroot




Containers vs Virtual Machines

Virtual Machines are heavy-weight emulations of real hardware




Containers are light-weight processes
The app looks like it's running on the Host OS




Images and containers

- A container is launched by running an *image*
- An *image* is an executable package that includes everything needed to run an application--the code, a runtime, libraries, environment variables, and configuration files
- A container is a *runtime instance* of an *image* i.e., what the image becomes in memory when executed (that is, an image with state, or a user process)
- You can see a list of your running containers with the command
 `docker ps`
 or
 `podman ps`

RedHat QUAY.IO container image hosting



http://
quay.io




Nginx installed and running with one command

We can run nginx as a web server with the command:

```
$ docker run -d -p 8080:80 nginx:alpine
Unable to find image 'nginx:latest' locally
latest: Pulling from library/nginx
6a5a5368e0c2: Pull complete
20a0fbbae148: Pull complete
2fbd37c8684b: Pull complete
Digest: sha256:e40499ca855c9edfb212e1c3ee1a6ba8b2d873a294d897b4840d49f94d20487c
Status: Downloaded newer image for nginx:latest
0d48962ddc380c421851fb808b9b3007c0c9c614bd08ae7e732955ddaa4c7b4a
```

Since we don't have a local nginx image it will be pulled from Docker hub the first time

Nginx running in a container



Kubernetes

Kubernetes is an Orchestration Platform

Scheduling

- decide where containers run

Lifecycle and Health

- keep containers running and restart them if they fail

Scaling

- grow and shrink deployments as needed

Naming and Discovery

- help containers find each other

Load Balancing


- distribute traffic across containers

...and a whole lot more




Kubernetes Architecture


- There can be one or more Master Nodes (HA)
- There can be zero or more Worker Nodes
- Everyone communicates via the API Server
- Kubelet on each Worker Node acting as an agent
- Everything can be on one node for development use




Kubernetes Primitives



Compute



Network



Storage



Node



Pod



Namespace




Deployment



ReplicaSet



StatefulSet



DaemonSet



Service



Ingress



Volume



PersistentVolume



StorageClass



PersistentVolumeClaim

Kubernetes Namespace


Namespaces

The primary grouping concept in Kubernetes

Provide a scope for names.

Provide the scope named resources that describe your application and how it should be deployed

A way to divide cluster resources between multiple uses



Kubernetes Pods

Pods

Containers run in Pods

The Pod is the smallest deployment possible

Containers are deployed from a container registry
(public or private)




Kubernetes ReplicaSets

ReplicaSet

ReplicaSets allow multiple copies of containers to be deployed

Each one in its own Pod


If a container dies, the ReplicaSet will spawn a new one



Kubernetes Volume Mounts

Volumes

ConfigMaps hold configuration parameters
Secrets hold credentials and other secrets
Persistent Volume Claims are used to ask for persistent storage for data/state




Kubernetes Deployments

Deployment

Sets up the ReplicaSets for you

Also specified the Secrets, ConfigMaps, and Volume Mounts

Provides features for rolling out updates and handling their rollbacks



Kubernetes Service


Service

Exposes Pods to the outside as:

ClusterIP



NodePort

Load Balancer




OpenShift


OpenShift Architecture




Container Platform




Lifecycle management




Deployments



Operator



Operator



</>

Developer

+Add

Topology

Monitoring

Search

Builds

Helm

Project

ConfigMaps

Secrets

You are logged in as a temporary administrative user. Update the [cluster OAuth configuration](#) to allow others to log in.

Project: devcamp

Application: all applications

View shortcuts



Display options

Filter by resource

Find by name...

/



D

nodejs-sample

A

sample-app

