



*IBM Storage Fusion HCI
System for Db2 Warehouse
Solution
Deployment Guide*

Table of Contents

1	INTRODUCTION	3
2	SOLUTION OVERVIEW	3
2.1	HARDWARE	3
2.2	SOFTWARE	6
3	SITE REQUIREMENTS	6
3.1	POWER	6
3.2	SITE READINESS	7
3.3	NETWORK PREPARATION	7
3.4	IMAGE REGISTRIES	ERROR! BOOKMARK NOT DEFINED.
4	ARCHITECTURE	7
4.1	SOFTWARE	7
4.2	HARDWARE	8
4.2.1	<i>Control nodes</i>	10
4.2.2	<i>Worker nodes</i>	10
4.2.3	<i>Management switches</i>	11
4.2.4	<i>High speed switches</i>	11
4.2.5	<i>KVM</i>	11
4.2.6	<i>Power Distribution Units (PDU)</i>	11
4.2.7	<i>Hardware configuration</i>	11
5	SERVICE	14
5.1	INSTALLATION	14
5.2	INSTALLING IBM STORAGE FUSION HCI SYSTEM	15
5.3	PERFORMANCE TUNING	15
5.4	INCREASE PID LIMITS	16
5.5	PREPARING REQUIRED STORAGE CLASSES	17
5.6	MAINTENANCE	18
5.6.1	<i>Self-service maintenance</i>	18
5.6.2	<i>IBM-managed maintenance</i>	19
5.7	PLANNING FOR ON-SITE INSTALLATION	20
5.8	ON SITE INSTALLATION EXECUTION	22
6	ORDERING	24
7	APPENDIX	25
7.1	DB2U CUSTOM RESOURCE REFERENCE	25

1 INTRODUCTION

IBM recently announced the IBM Storage Fusion HCI System for Db2 Warehouse Solution¹. It is the next generation of on-premises analytics solution based on IBM Db2 Warehouse software. It comes preconfigured and ready to go with everything clients need to optimize the total cost of ownership (TCO) of their environment. This document is intended to provide guidance in deploying Db2 Warehouse in the customer environment. It will explain the hardware, software, and infrastructure requirements for a successful Db2 Warehouse deployment on Fusion HCI based hardware infrastructure.

2 SOLUTION OVERVIEW

IBM has defined a Storage Fusion HCI-based architecture that, if deployed as described in this document, will ensure that Db2 Warehouse is configured to provide the highest level of performance and availability. IBM provides Expert Lab services for solution installation in customer data center and support after installation.

The Db2 Warehouse on Fusion HCI solution consists of a hyperconverged Intel x86-processor-based infrastructure that is installed in the customer's data center and hosts the software stack. Red Hat OpenShift runs on top of the hardware and provides the platform that Db2 Warehouse runs on. IBM Storage Fusion also runs on OpenShift, and provides data services for Db2 Warehouse, including highly available storage, as well as backup and restore capabilities.

2.1 HARDWARE

IBM Storage Fusion HCI is an Intel x86-processor-based hyperconverged infrastructure that provides the compute, network, and storage necessary to run Db2 Warehouse. The Db2 Warehouse on Fusion HCI solution is available in four t-shirt sizes that have been optimized for high levels of performance and availability.

Each t-shirt size comes with:

- Three storage-rich x86-servers that act as control nodes.
- Redundant 1 GbE management switches.
- Redundant 100 GbE ethernet switches.
- Redundant power distribution units.
- A rack mounted console for serviceability.

¹ <https://www.ibm.com/docs/en/announcements/storage-fusion-hci-system-db2-warehouse-is-available-in-selected-configurations>

The t-shirt sizes vary based on the number of worker nodes, with the smallest t-shirt size consisting of 4 worker nodes, and the largest t-shirt size consisting of 13 worker nodes. Both the control nodes and the worker nodes each have ten 7.68 TB NVMe drives that are used to provide storage for the appliance. The table below shows the four t-shirt sizes, along with the CPU cores, and usable storage capacity.

		1 spare	1 spare	2 spare	2 spare
		HCI-07	HCI-10	HCI-13	HCI-16
SF-HCI Compute Models		BRS (1/4 Rack)	BRM (1/2 Rack)	BRL (3/4 Rack)	BRXL (1 Rack)
Compute	Worker Nodes ¹	4	7	10	13
	Active Worker Cores	192	384	576	768
	Active Cores for Db2 Warehouse ²	156	312	416	572
Storage	No. of Active 7.68 TB NVMe Drives	70	100	130	160
	Protected TB	329	471	612	754
	Tier-1 User Data (TB) ³	230	329	428	527

¹Includes 1 stand by Worker Node for up to 6 Worker Nodes in each configuration for HA

²Total number of Db2 Warehouse Cartridge (Part# D0AFLZX) VPCs required

³70% of capacity available for User Data Tables, 30% reserved for Temp Tables

The HCI solution is a 42-unit rack. The smallest t-shirt size consists of three control nodes, plus four worker nodes, for a total of seven nodes. The other three t-shirt sizes build upon this base configuration, adding three additional worker nodes with each successive t-shirt size. The solution rack is filled from the bottom up (starting at rack unit two). Rack units 18-21 hold the two management switches and two high speed switches.

HCI-16

HCI-13

HCI-10

HCI-07



Power distribution unit

Power distribution unit

KVM

32-port 100 GbE switch

32-port 100 GbE switch

48-port 1GbE management switch

48-port 1GbE management switch

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

2.2 SOFTWARE

The IBM Db2 Warehouse for Storage Fusion HCI solution comes with an install workflow that automates the installation of the platform that Db2 Warehouse runs on.

- Red Hat OpenShift – provides a platform for running containerized applications. In the case of this solution, the containerized application is Db2 Warehouse.
- Storage Fusion – provides data services for Db2 Warehouse:
 - Persistent storage – RWO and RWX storage are available via storage classes provided by Storage Fusion. The storage is provided via a software defined storage platform called the Global Data Platform, which is based on IBM Storage Scale ECE. All NVMe drives in each of the control and worker nodes is aggregated together into a single pool of storage that can be used to provide persistent volumes for Db2 Warehouse. The storage is made highly available via the use of erasure coding. This ensures that data remains available even if drives or even entire nodes experience failures. The persistent storage can also be encrypted, securing data at rest.
 - Backup and restore – the Db2 Warehouse configuration can be backed up and restored in place or in another cluster.

The installation of IBM Storage Fusion HCI results in a fully operational OpenShift cluster with Storage Fusion data services fully operational. Db2 Warehouse can then be deployed from the IBM catalog.

xxx

3 SITE REQUIREMENTS

IBM provides extensive guidelines to help you prepare your site for the delivery and installation of the solution, including computer room location, electromagnetic compatibility, environmental design, power, cooling, and cabling requirements. These guidelines are available [online](#) and are adhered to by the Expert Lab team when a solution is installed and setup by IBM.

3.1 POWER

Planning power involves three core documented topics:

- [Understanding the power consumption for the configured components](#) in the t-shirt size. Fusion HCI for Db2 Warehouse consists of three 32 core servers, and 4, 7, 10, or 13 64 core servers based on the t-shirt size covered in section 2.1.
- [The supported power configurations](#). Fusion HCI for Db2 Warehouse consists of 9155-C01 and 9155-C05 nodes.

- [The supported PDU power cords.](#)
- If using PDUs provided by the customer, consult [the topic on power consumption for customer-provided PDUs.](#)

3.2 SITE READINESS

Requirements regarding site readiness, such as floor space, floor loading, vibration, and static electricity are covered in the [Site Readiness topic](#)

3.3 NETWORK PREPARATION

A successful install of Fusion HCI depends on the customer network being pre-configured so that the appliance can connect to the network, find appropriate DHCP entries for each node, and find appropriate DNS entries for each node. It is important to follow [the Network Planning section of the documentation](#) closely to ensure that the network is set up correctly.

4 ARCHITECTURE

4.1 SOFTWARE

The software stack consists of Red Hat OpenShift, Storage Fusion, and Db2 Warehouse. Red Hat OpenShift and Storage Fusion are installed automatically as part of the Storage Fusion HCI installer. Db2 Warehouse is installed by IBM Expert Lab as an additional step.

Red Hat OpenShift

RedHat OpenShift is installed automatically as part of the Storage Fusion HCI installer. OpenShift provides the container runtime that Db2 Warehouse runs in. OpenShift runs on all nodes in the appliance.

Storage Fusion

The Storage Fusion operator is automatically installed into OpenShift as part of the Storage Fusion HCI install. The Storage Fusion operator provides monitoring functionality for the appliance, as well as providing the persistent storage that Db2 Warehouse needs to run. Storage Fusion's responsibilities include:

- Monitoring the hardware infrastructure in the appliance and creating OpenShift events for any hardware issues.
- Aggregating the drives across the worker nodes and creating a software defined CSI storage provider. This allows Db2 Warehouse to request persistent volumes, which are needed for data persistence.

Db2 Warehouse

Db2 Warehouse database engine version: 11.5.8.0

Db2 Warehouse can be deployed in a Red Hat OpenShift cluster as a containerized micro-service, or pod, managed by Kubernetes. On Fusion HCI, OpenShift is used as the platform with Db2 Warehouse lifecycle managed via an Operator with following versioning.

- Version: s11.5.8.0-cn2
- Operator:
 - Version: 110508.0.2 (Latest)
 - Channel: v110508.0
 - CASE: 5.1.4

Reference: [Db2 Warehouse on Red Hat OpenShift and Kubernetes](#)

4.2 HARDWARE

The Db2 Warehouse appliance t-shirt sizes provide a designed set of Fusion HCI configurations for optimally running Db2 Warehouse. These configurations use a subset of the rack and vary based on the number of worker nodes.

HCI-16

HCI-13

HCI-10

HCI-07



Power distribution unit

Power distribution unit

KVM

32-port 100 GbE switch

32-port 100 GbE switch

48-port 1GbE management switch

48-port 1GbE management switch

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

Server with storage

The minimum configuration for a Fusion HCI reference architecture that has been defined by IBM for Db2 Warehouse consists of the following components:

- 3 control nodes
- 4 worker nodes
- 2x 1 GbE management switches
- 2x 100 GbE high speed switches
- 1 KVM

Each larger t-shirt size expands the configuration by adding 3 additional worker nodes.

4.2.1 Control nodes

Model 9155-C01

- CPUs
 - 2x16-core processors, for a total of 32 cores
- Memory
 - 256 GB – DDR4
 - 16 GB x16 DIMMs
- PCIe slots
 - PCIe Gen 4 dual-port 100Gb Ethernet adapter card
 - PCIe Gen 4 dual-port 25Gb Ethernet adapter card
- NVMe storage – 7.68 GB drives x10

4.2.2 Worker nodes

Model 9155-C05

- CPUs
 - 2x32-core processors, for a total of 64 cores
- Memory
 - 1024 GB – DDR4
 - 64 GB x16 DIMMs
- PCIe slots
 - PCIe Gen 4 dual-port 100Gb Ethernet adapter card
 - PCIe Gen 4 dual-port 25Gb Ethernet adapter card
- NVMe storage – 7.68 GB drives x10

4.2.3 Management switches

NVIDIA AS4610 Ethernet Switches (48x1Gb) – Deployed in a pair of switches for redundancy. The management switches are used to monitor all other hardware that is part of the appliance.

4.2.4 High speed switches

NVIDIA SN3700C Ethernet Switches (32x100GbE) – Deployed in a pair of switches for redundancy. The high-speed switches are used to provide a dedicated 100 GbE storage network and a 25 GbE OpenShift network.

4.2.5 Rack mounted console

Used for the initial setup of the appliance by an IBM SSR.

4.2.6 Power Distribution Units (PDU)

For the power distribution units, there are two possible choices. Both have 12xC13 outlets.

- FC ECJN for single phase and three-phase wye
- FC ECJQ for three-phase delta

[Full power cord documentation](#)

4.2.7 Hardware configuration

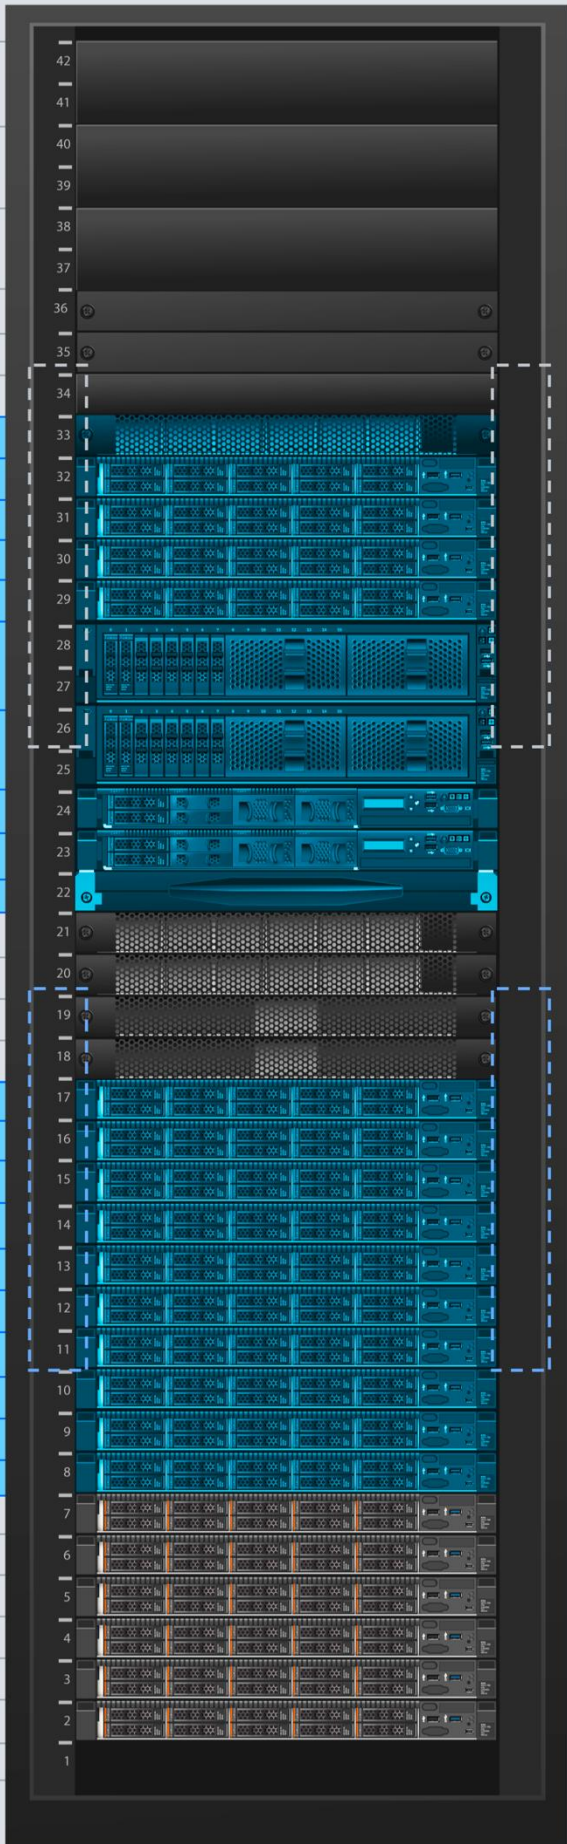
The following diagram shows the physical configuration of hardware components within an IBM Storage Fusion HCI. Not all of the optional components are necessary for the Db2 Warehouse on Fusion HCI solution, and they are not part of the t-shirt sizes. The components of the Fusion HCI appliance that are not applicable to the Db2 Warehouse t-shirt sizes are:

- AFM nodes in rack unit 23 and 24. These are used for data caching scenarios that aren't directly applicable to Db2 Warehouse.
- GPU servers in rack units 25-28. These are used for AI & Machine Learning workloads, and aren't directly applicable to Db2 Warehouse.
- Storage/compute servers in rack units 29-32. These nodes surpass the largest t-shirt size for Db2 Warehouse.

- 100 GbE ethernet spine switch in rack unit 33. This spine switch is used to connect multiple Fusion HCI racks in order to enable a multi-rack OpenShift cluster. A multi-rack HCI deployment has not yet been tested with Db2 Warehouse solution design.

42	
41	2U Filter
40	
39	2U Filter
38	
37	2U Filter
36	1U PDU (Horizontally mounted)
35	1U PDU (Horizontally mounted)
34	1U Filter
33	32-Port 100 GbE Ethernet spine switch
32	Storage / Compute server
31	Storage / Compute server
30	Storage / Compute server
29	Storage / Compute server
28	GPU Server with 3x GPU PCIe Gen 4 adapter cards
27	
26	GPU Server with 3x GPU PCIe Gen 4 adapter cards
25	
24	AFM Node
23	AFM Node
22	KVM
21	32-Port 100 GbE Ethernet spine switch
20	32-Port 100 GbE Ethernet spine switch
19	48-Port 1 GbE Management Ethernet Switch
18	48-Port 1 GbE Management Ethernet Switch
17	Storage / Compute server
16	Storage / Compute server
15	Storage / Compute server
14	Storage / Compute server
13	Storage / Compute server
12	Storage / Compute server
11	Storage / Compute server
10	Storage / Compute server
9	Storage / Compute server
8	Storage / Compute server
7	Storage / Compute server
6	Storage / Compute server
5	Storage / Compute server
4	Storage / Compute server
3	Storage / Compute server
2	Storage / Compute server
1	Reserve 1U Space at Bottom

Optional components



Standard rear mounted PDUs



Optional rear mounted PDUs

The Db2 Warehouse appliance uses rack units 2-17 to house the control and worker nodes, with the number of worker nodes varying based on the t-shirt size. Rack units 18-19 house the management switches, and rack units 20-21 house the high-speed switches.

Two vertically mounted power distribution units are positioned in pockets spanning rack units 26-34. Additional power distribution units are optional based on the number of nodes in the rack and the data center power supply.

5 SERVICE

5.1 INSTALLATION

Project management is often the forgotten aspect of installation. So, IBM will provide Expert Lab Services for installation and set-up of a Storage Fusion HCI for Db2 Warehouse Solution in customer data center, including connecting the solution into customer's network. These Expert Lab Services will include the following components:

- System Expert Lab Services for pre-installing OCP and Storage Fusion
- Systems Expert Lab Services for installing the Solution at the Customer Site including connecting the Solution in the customer's network.
- Software Expert Lab Services to install Db2 Warehouse on the Solution at Customer Site including creating Db2 instances that are table ready to ingest customer data.
- Software Expert Lab Migration Assessment Services to develop a customized migration plan to meet Customer requirements, including scoping the migration work effort and documenting migration strategy to meet customer's downtime sensitivities for various applications.

Installation services give the customer a “gold-glove experience” from the moment a Solution is ordered until the Solution is up and running at the customer site.

5.1.1 Access to image registries

A successful install of Db2 Warehouse on Fusion HCI requires the solution to be able to access image registries that host Red Hat OpenShift, Fusion software, and Db2 Warehouse software. These images provide the software that is used to run Red Hat OpenShift, Fusion software, and Db2 Warehouse. IBM hosts these image registries publicly in the IBM Entitled Registry, meaning that if outbound firewall access (with or

without a proxy) can be provided to the public registries, the software can be downloaded automatically.

To access the public images, you will need a software entitlement key. The process for obtaining an entitlement key [is documented here](#).

Some customers prefer to use a private image registry to host images. This allows them to perform their own security scanning of images, and control which image versions are available within their data center. Using a private registry requires that customers mirror images from the IBM entitled registry.

The procedure for using a private image registry [is documented here](#).

5.1.2 IBM Entitled registry access

In order to pull images from the IBM entitled registry, the OpenShift cluster will need an image pull secret. Create an image pull secret by following [the documented procedure here](#).

5.2 INSTALLING IBM STORAGE FUSION HCI SYSTEM

Storage Fusion HCI System comes with the bootstrapping software that is installed from the factory. The IBM support representative completes the initial verification and physically connects the system to network and power. Then, they conduct the network setup, which connects the solution to the data center network. This procedure configures all the default nodes (three controllers and three compute nodes). If you ordered more nodes, then they get installed as well. Use this information to install Storage Fusion HCI System. For additional information, see the [IBM Documentation](#).

5.3 PERFORMANCE TUNING

Tune the Storage Fusion HCI System for optimal Db2 Warehouse performance. Perform these tunings on top of a standard installation of IBM Storage Fusion HCI System.

1. Configure tuning by saving the YAML below to a file called scaleTune.yaml:

```
apiVersion: scale.spectrum.ibm.com/v1beta1
kind: Daemon
metadata:
  name: ibm-spectrum-scale
  namespace: ibm-spectrum-scale
```

```
spec:
  clusterProfile:
    maxTcpConnsPerNodeConn: '8'
    prefetchThreads: '64'
```

2. Run the following command to apply the YAML to your cluster:

```
oc apply -f scaleTune.yaml
```

5.4 INCREASE PID LIMITS

In order to prevent failures, database crashes, or pod restarts during multi-user runs with a large number of users, it is crucial to set **pids.max** to 65536 on the cluster and inside the container, as recommended for production environments. This requires a MCP config change, and therefore a cluster restart.

Viewing PID Limits inside the Container

```
[root@roadieaf64c1 ~]# oc rsh c-db2u-cr-db2u-0 bash -l
Defaulted container "db2u" out of: db2u, instdb (init), init-
labels (init), init-kernel (init)
[db2uadm@c-db2u-cr-db2u-0 /]$ su - db2inst1
Last login: Tue Sep 19 15:01:07 UTC 2023 on pts/4
[db2inst1@c-db2u-cr-db2u-0 - Db2U db2inst1]$ cat
/sys/fs/cgroup/pids/pids.max
10000
```

Changing the PID Limits for the Cluster

1. Assign labels to the physical servers

```
oc label machineconfigpool master maxpidlimit=true
oc label machineconfigpool worker maxpidlimit=true
```

2. Create the following CR using the oc CLI or the OCP UI


```

apiVersion: machineconfiguration.openshift.io/v1
kind: KubeletConfig
metadata:
  name: maxpidlimit
spec:
  machineConfigPoolSelector:
    matchLabels:
      maxpidlimit: "true"
  kubeletConfig:
    podPidsLimit: 65536

```

5.5 PREPARING REQUIRED STORAGE CLASSES

From the OCP UI or the oc command-line tool, apply the following yaml

```

kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: ibm-spectrum-scale-meta-rwx-sc
provisioner: spectrumscale.csi.ibm.com
parameters:
  shared: "true"
  volBackendFs: ibmspectrum-fs
reclaimPolicy: Delete
allowVolumeExpansion: true
volumeBindingMode: Immediate
---
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: ibm-spectrum-scale-data-rwo-sc
provisioner: spectrumscale.csi.ibm.com
parameters:
  volBackendFs: ibmspectrum-fs
reclaimPolicy: Delete
allowVolumeExpansion: true
volumeBindingMode: Immediate
---
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: ibm-spectrum-scale-tempts-rwo-sc
provisioner: spectrumscale.csi.ibm.com
parameters:
  volBackendFs: ibmspectrum-fs
reclaimPolicy: Delete
allowVolumeExpansion: true
volumeBindingMode: Immediate
---

```

```

kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: ibm-spectrum-scale-archivelogs-rwx-sc
provisioner: spectrumscale.csi.ibm.com
parameters:
  shared: "true"
  volBackendFs: ibmspectrum-fs
reclaimPolicy: Delete
allowVolumeExpansion: true
volumeBindingMode: Immediate
---
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: ibm-spectrum-scale-activelogs-rwo-sc
provisioner: spectrumscale.csi.ibm.com
parameters:
  shared: "true"
  volBackendFs: ibmspectrum-fs
reclaimPolicy: Delete
allowVolumeExpansion: true
volumeBindingMode: Immediate

```

5.6 MAINTENANCE

Customer receives the rights to all the new version updates of firmware, OpenShift Container Platform, Storage Fusion and Db2 Software in the Solution.

5.6.1 Self-service maintenance

Before proceeding with updating any or all the software stack, please refer to the compatibility matrix which lists all the IBM validated and supported versions.

Fusion HCI release	OCP version	Db2 Warehouse
2.6.1	4.12	s11.5.8.0-cn2 Operator: <ul style="list-style-type: none"> • Version: 110508.0.2 • Channel: v110508.0 • CASE: 5.1.4

<TODO: add a table like above into a new section in Db2 WH on RHOS KC and update as needed >

OCP minor version updates

- OCP minor version updates are managed through channels. Unless explicitly recommended by IBM, the channel option should be stable. Refer to topic [Understanding update channels and releases](#) for more details.
- OCP minor version updates can be executed via Control plane web console or using OCP CLI. Given the fact, Db2 Warehouse need to run in the same configuration as originally deployed, for example if it was deployed with 6 pods, same 6 pods need to be available even during cluster/node updates. Therefore, rolling/canary updates are not recommended. Refer to following topics for more details on minor version updates.
 - [Updating a cluster using the web console](#)
 - [Updating a cluster using the CLI](#)

5.6.2 IBM-managed maintenance

IBM Lab services/expert labs will manage the software stack updates.

Storage Fusion version updates

Storage Fusion software upgrades are managed via the Storage Fusion GUI. To upgrade Storage Fusion software, follow [the documented procedure](#).

Storage Fusion upgrades of the Global Data Platform service require nodes to be rebooted. Node reboots are orchestrated in a rolling fashion such that only a single node is down at a time.

Some Storage Fusion upgrades include upgrades for the switch and server firmware. The Storage Fusion GUI includes a workflow that fully orchestrates the firmware upgrades. The firmware upgrade procedure is [documented here](#).

Server firmware upgrades requires nodes to be rebooted. Node reboots are orchestrated in a rolling fashion such that only a single node is down at a time.

Switch firmware upgrades require the switches to be rebooted. This doesn't impact Db2 Warehouse because the Storage Fusion HCI appliance contains redundant switches, and the upgrade is orchestrated so that only one switch is down at a time.

Db2 Warehouse Operator and Operand upgrades

Preparing to upgrade

Before upgrading to a newer version of Db2 Warehouse, it's recommended to execute the following steps:

- Backup the database
- Stop all client applications

Upgrading Db2 Warehouse software on IBM Fusion HCI is comprise of two stages:

- Upgrading the Db2 Warehouse Operator. First step is to [identify the CASE bundle version](#). Then follow the process in to [upgrade Db2 Warehouse Operator to the newer CASE using cloudctl tool](#).
- Upgrading the Db2uInstance resource. After Db2 Warehouse Operator is upgrade is completed, [upgrade Db2uInstance resource to the newer version by changing the spec.version to the new version](#).

5.7 PLANNING FOR ON-SITE INSTALLATION

IBM will provide a project manager to work with the client, the IBM account team, and business partner (if applicable) to assist with planning the site prerequisites prior to delivery of the rack solution as well as prerequisites prior to on-site installation. The details for delivery and on-site installation will be contained within the document referred to as the Site Survey.

The prerequisites that will be planned prior to delivery of the system include:

Site Contacts: IBM will request the client to provide contact information for a team of people with the following roles (in some cases multiple roles may be shared by one person):

- Delivery - responsible for the shipping dock/storage access at the facility
- Hardware - responsible for hardware planning, rack placement on the data center floor
- Power - responsible for electricians, power infrastructure in data center
- Network - responsible for management of the data center network infrastructure
- Project Manager - responsible for providing coordination between the above clients contacts

Site Access: IBM will request the client to provide site specific information to ensure the rack can safely be placed into the facility and any obstacles during delivery can be circumvented.

- Receiving dock - specifications of the dock and type of delivery truck required
- Elevators - (if applicable) elevators meet the height, width, and weight of the rack
- Doorways - (if applicable) doorways meet the height and width requirements of the rack

- Ramps - (if applicable) required to traverse raised floors and meet rack specifications
- Floor load - all floors leading to/at the install location are reinforced to support the rack weight
- Floor protection - (if applicable) type of material required to protect floors from rack wheels

The prerequisites that will be planned prior to on-site installation include:

Site Infrastructure: IBM will request the client to provide site specific information to ensure the rack can be installed correctly, remain operational in the environment, and be serviced/maintained safely.

- Service space - adequate space in front/behind the rack for IBM to service the equipment
- Power - quantity and type of power feeds that meet voltage/amp requirements of the rack
- Cooling - adequate cooling solution provided to meet the system cooling requirements

Network Connectivity: IBM will request the client to provide site specific information on the networking to ensure the system can be integrated into the client's environment and maintain connectivity to required management and update services.

- Network cables - qty and type to connect system components to the client's network
- Network information - Required number of IP addresses, hostnames, and domain information
- Network services - Connectivity information for the client's DNS, NTP, SMTP servers
- Access Ports - Firewall ports that must be opened for system connectivity and operation

On-site Access: IBM will request the client to provide site specific information on access requirements for the IBM install engineer to visit the site and perform installation activities.

- Site clearance - policies/protocols to secure clearance for IBM personnel to enter the site
- Site Security - Security clearance requirements/identification to enter the site
- Personal Devices - allowed electronic devices, storage media, physical media, tools

- Site Connectivity - cell phone reception, vendor Wi-Fi accessibility
- Hours - hours of operation/access for IBM install engineer to remain on site
- Escort - requirements on escorted/unescorted access into the facility

5.8 ON SITE INSTALLATION EXECUTION

The IBM project manager will work with the client to ensure all prerequisites for on-site installation have been met (site readiness, completed site survey, customer availability). The project manager will then schedule an IBM Service and Support Representative (SSR) to visit the site and perform on-site installation of the system. This includes all activities to ensure the system is operational, updated, healthy and connected to the client's network.

The SSR plays an important role in bringing up the Storage Fusion HCI racks at customer premises. An SSR follows the Technical Delivery Assessment (TDA) process to proactively engage with customers before the Storage Fusion HCI rack arrival to review the site readiness and network setup. An SSR also works with the customer to smoothly install Storage Fusion.

The following three implementation stages provide highlights about the activities that the SSR performs:

5.8.1 Preinstall Technical Delivery Assessment (TDA)

Unboxing and equipment inspection

Verify whether unboxing is properly done by the customer. The SSR will physically inspect the system to ensure the equipment is in working condition. This inspection includes the rack frame, doors, wheels and to ensure the rack is secured in place. The equipment inside the rack will be inspected, including servers, switches, cables, PDUs. All cabling will be reviewed to ensure any loose connections are remediated prior to powering on the system. If the system contains multiple racks, the intra-rack cabling will be connected between racks. The engineer will then connect the client network cables to the system.

Powering on the appliance

The SSR will work with the data center electrician or hardware manager to connect power to the system. The engineer will ensure all breakers on the system PDUs are turned off prior to power connectivity, then they will bring the power connections for the rack PDUs outside the rack (either above or below) and allow the data center personnel to connect the cables to the data center power receptacles. The SSR will turn on the PDU breakers and allow the system switches and storage to power on.

Network preparation

The SSR will coordinate with the customer network team to ensure that the customer network is prepared for the HCI appliance. The customer will have received information on MAC addresses, DHCP settings, DNS settings, and VLAN planning prior to the appliance being shipped, and the network should already be configured so that it is ready for the appliance.

The SSR will coordinate with manufacturing and the customer network team to configure the DHCP server and to allocate fixed IP addresses by using MAC addresses of the nodes in the appliance rack.

The SSR will work with the customer network team to configure the DNS server per OpenShift® network requirement.

5.8.2 Appliance health check

The SSR will verify that the hardware components are healthy:

- Inspect LEDs on the front operator panel of storage and compute nodes.
- Inspect LEDs on the operational panel of high-speed switches from the rear of the Storage Fusion rack.
- Check and secure connection among hot-swap power supplies and power distribution units (PDUs).

5.8.3 Connect the solution to the network

The SSR will power on the rack mounted console (KVM) and use a web GUI to configure the solution to connect to the network. The rack mounted console connects to the preinstalled provisioner node, which is used to perform the network configuration

The SSR will enter the network configuration parameters (VLAN, Link, port type, speed, etc) that are collected from the customer during the TDA phase, and recorded in an Excel spreadsheet. This information is entered into the GUI, which then triggers the network configuration. The GUI will verify the network is configured properly and show any issues that it detects. The SSR will work with the client to troubleshoot connectivity issues from the system and resolve configuration issues on the system or advise on potential changes required on the data center network to allow connectivity from the system to required endpoints/services.

The final step of the web GUI workflow for connecting to the network provides the engineer with a URL that can be used to initiate the Storage Fusion HCI installer. This URL will be provided to the IBM Expert Labs Fusion representative so that the installer can be run.

5.9 COMPLETING THE INSTALL

IBM Expert Labs Representative uses the URL provided by the SSR to launch the

Storage Fusion HCI installer. The installer configures a three node OpenShift cluster and installs Storage Fusion software into the cluster. The installer then adds the rest of the nodes in the rack to the OpenShift cluster and installs Fusion's Global Data Platform service, which creates a software defined storage cluster.

In order to complete the installer, the IBM Expert Labs Representative needs the following information from the customer:

- Whether the public image registry will be used, or whether a private image registry is being used.
 - In the case of a public image registry, an image pull secret and a software entitlement key is needed.
 - In the case of a private image registry, connection information is needed for the private image registry.
- The installer can optionally configure a private certificate for OpenShift so that a user doesn't have to do this as a later task. If a private certificate will be used, it needs to be provided to the IBM Expert Labs Representative.

The Fusion HCI installation process is [documented extensively in the Knowledge Center](#).

Once the installation has been completed, the IBM Expert Labs Representative will provide the customer with authentication information that can be used to log into OpenShift. The customer now has a fully functional OpenShift cluster. Expert Labs services can now install Db2 Warehouse.

6 ORDERING

Hardware configuration for the IBM StorageFusion HCI for Db2 Warehouse is available in eConfig. Limited options for configuration are provided for things like PDU connection. Please contact your IBM Sales Representative or Business Partner based upon your requirements.

7 APPENDIX

7.1 DB2U CUSTOM RESOURCE REFERENCE

Db2 on OpenShift and Kubernetes platforms are deployed using a custom resource (CR) which defines the deployment attributes in the spec.

```
apiVersion: db2u.databases.ibm.com/v1
kind: Db2uInstance
metadata:
  name: db2u-cr
  namespace: db2
spec:
  account:
    securityConfig:
      privilegedSysctlInit: true
  environment:
    authentication:
      ldap:
        enabled: false
  databases:
  - dbConfig:
      LOGPRIMARY: "100"
      LOGSECOND: "50"
      LOGFILSIZ: "200000"
      NUM_DB_BACKUPS: "4"
      TRACKMOD: "NO"
      name: BLUDB
  dbType: db2wh
  instance:
    registry:
      Db2_4K_DEVICE_SUPPORT: "ON"
      Db2COMPOPT: "LOCKAVOID_EXT_CATSCANS"
  partitionConfig:
    dataOnMln0: true
    total: 48
    volumePerPartition: true
```

```

license:
  accept: true
nodes: 6
podTemplate:
  db2u:
    resource:
      db2u:
        limits:
          cpu: 24
          memory: 800Gi
advOpts:
  memoryPercent: 99
storage:
- name: meta
  spec:
    accessModes:
    - ReadWriteMany
    resources:
      requests:
        storage: 1000Gi
    storageClassName: ibm-spectrum-scale-meta-rwx-sc
  type: create
- name: data
  spec:
    accessModes:
    - ReadWriteOnce
    resources:
      requests:
        storage: 5466Gi
    storageClassName: ibm-spectrum-scale-data-rwo-sc
  type: template
- name: archivelogs
  spec:
    accessModes:
    - ReadWriteMany
    resources:
      requests:
        storage: 9600Gi
    storageClassName: ibm-spectrum-scale-archivelogs-rwx-sc
  type: create
- name: tempts
  spec:
    accessModes:
    - ReadWriteOnce
    resources:
      requests:
        storage: 7666Gi

```

```
storageClassName: ibm-spectrum-scale-tempts-rwo-sc  
type: template  
version: s11.5.8.0-cn2
```