


IBM Cloud Infrastructure Center

Sarah Becker
Level 2 Support
sarahbecker@ibm.com



Agenda

- ICIC Architecture
- Enhancements
- Storage Types

ICIC Overview

IBM Cloud Infrastructure Center (ICIC) is an Infrastructure-as-a-Service (IaaS) platform for IBM Z and LinuxONE.

What it Does

- Centralized management of virtual machines, storage, and networking
- Provides lifecycle management for virtual resources
- Enables automation through APIs and Infrastructure as Code


Key Value

- Simplifies infrastructure management across z/VM and KVM
- Supports hybrid cloud environments
- Bridges traditional workloads and modern cloud operations


How It's Used

- Deploy and manage VMs
- Integrate with tools like Terraform, Jenkins, and OpenShift
- Enable self-service and automation for infrastructure teams


Single Node Architecture




HA Architecture




Multi Site Architecture



Site 1




Site 2




Site 3

IBM Cloud Infrastructure Center - LPAR management Framework



Security and Serviceability Enhancements




Proactive Serviceability

- Refined environment checks improve diagnostic accuracy
- Automated SMAPI health monitoring enables early issue detection
- Faster troubleshooting and reduced downtime



Access Control & Governance

- Enforced strict separation of user privileges
- Role-based access improves governance and control
- Minimizes risk of unauthorized actions




Secure Communication & Operations

- Strengthened cross-user communication controls
- Secure handling of credentials and system interactions
- Improves overall platform resilience and stability

Modernize for Hybrid Infrastructure

Empower how you deploy, manage, and integrate infrastructure as a service



1.2.5 – Supported hypervisors, operating systems & servers

Hardware Platforms

- IBM z17™
- IBM z16® (all models)
- IBM z15® (all models)
- IBM z14® (all models)
- IBM® LinuxONE 5
- IBM® LinuxONE 4 (all models)
- IBM® LinuxONE III (all models)
- IBM® LinuxONE II (all models)

As a managed hypervisor:

- z/VM 7.2, z/VM 7.3, and z/VM 7.4
- KVM as part of Red Hat Enterprise Linux (RHEL) 8.8, 8.10, 9.4 and 9.6(with v1.2.5+iFix01)

As a host environment on z/VM or KVM:

- RHEL 8.8, 8.10, 9.4 and 9.6(with v1.2.5+ iFix01)

As a deployable guest operating system instance on z/VM:

- RHEL 8.8, 8.10, 9.2 ,9.4 , 9.6 or 10.0.
- Red Hat CoreOS 4.12, 4.14, 4.15, 4.16, 4.17, 4.18, 4.19 and 4.20 as part of Red Hat OpenShift Container Platform
- SUSE Linux Enterprise Server 15 SP3 – SP5, and SP6
- Ubuntu 20.04, 22.04, or 25.04

As a deployable guest operating system instance on KVM:

- RHEL 8.8, 8.10, 9.2 , 9.4, 9.6 or 10.0
- Red Hat CoreOS 4.12, 4.14, 4.15, 4.16, 4.17, 4.18, 4.19 and 4.20 as part of Red Hat OpenShift Container Platform

Storage Pool Rename


Enable seamless synchronization of storage pools that are renamed directly in the Storage Controller with IBM Cloud Infrastructure Center, ensuring uninterrupted operation of storage templates and volumes.

Business value :

This enhancement improves reliability and reduces operational disruptions by allowing users to easily synchronize renamed storage pools. It eliminates manual fixes, ensures data consistency, and maintains smooth storage lifecycle management across the environment.

Use case: User creates a storage template in IBM Cloud Infrastructure Center using an existing pool and uses it to create volumes. Later, the user renames the pool in the storage backend. IBM Cloud Infrastructure Center now supports synchronization of the renamed pool, ensuring that all storage templates and associated volumes continue to operate seamlessly.

Solution: Using command to rename and synchronize the storage pool name with the backend storage. When a pool is renamed in the storage system, users can run the command to update the pool name, ensuring all related storage templates and volumes continue to function seamlessly.



Advanced Seamless Onboarding of z/VM


Onboarding of Virtual Machines created using native z/VM tools into IBM Cloud Infrastructure Center, providing full lifecycle management and supporting all standard functionalities equivalent to IBM Cloud Infrastructure Center deployed VMs.

Business value :

- Ensures consistent management experience for both onboarded and IBM Cloud Infrastructure Center deployed VMs
- Improves operational efficiency with seamless integration of externally created VMs.

Use case: Users want to onboard Virtual Machines that were created using the z/VM tool outside and manage those VMs through IBM Cloud Infrastructure Center.

Solution: The onboarding VM feature helps integrate externally created Virtual Machines along with their attached volumes and L2 network, synchronizing VM states with the hypervisor, improving polling efficiency, and enabling full lifecycle operations such as soft restart, capture, live migration, and volume attachment.



Onboard and attach volumes to a deployed VM

Onboards and attaches volumes created externally to a virtual machine managed by IBM Cloud Infrastructure Center.


Business value :

- Restores consistency between Storage Controller volume states and the IBM Cloud Infrastructure Center.
- Ensures all attached volumes are accurately reflected in the management interface.

Use case: In certain cases, volumes remain attached to a VM in the backend but are not listed in the IBM Cloud Infrastructure Center. Users need a quick way to detect and onboard such volumes without redeploying the VM or reattaching storage manually.

Solution: The onboard-volumes command automatically scans backend configurations, identifies missing volume mappings, and attaches them to the corresponding VM managed by IBM Cloud Infrastructure Center. This synchronization ensures the VM and its storage resources are fully manageable through UI / APIs, maintaining operational consistency across the environment.

Onboarding Volumes



Platform Management through Terraform

Enable onboarding of Logical Partitions and their OSA and FICON resources managed by IBM® Cloud Infrastructure Center

Business value :

- Automate IODF-to-Terraform import conversion using the Terraform Transformer service.
- Enable seamless onboarding of IBM Z Logical Partitions in Standard mode through Terraform.
- Gain complete visibility into existing OSA, FICON, and controller connections for each Logical Partition.
- Expand or reduce the OSA and FICON resource / connectivity on Logical Partitions.

Use case: System programmers must be able to use Terraform configurations to manage existing Logical Partitions and their provisioned OSA and FICON resources, including ongoing lifecycle operations for these resources.

Solution: IBM® Cloud Infrastructure Center now enables customers to use Infrastructure as Code with Terraform to discover and manage existing IBM Z Logical Partition configurations.

User Experience Improvement

Additional Key Point :

New Terraform provider and Transformer Service is developed for converting the HCD configuration to Terraform Import Statement

Redefine Control Pane

Ability to split large Cloud Infrastructure Center clusters into multiple smaller clusters, each with its own compute node(s), for better management capability


Business value :

- Improve management by smaller, leaner management scope
- Isolate cluster based on its usage, e.g., production, development, or test

Use case: As a user, would like to separate / isolate a portion of an existing Cloud Infrastructure Center cluster and manage the separated cluster with its own compute node(s).

Solution: Move one or multiple compute nodes into separate clusters.

Optimize Compute Node Assignment



Enabled Support for Red Hat Enterprise Linux – 9.4

Ability to deploy and run IBM Cloud Infrastructure Center on RHEL 9.4 as the base operating system, enabling improved platform support, security, and lifecycle alignment.


Business value :

- Enabling RHEL 9.4 allows user to use the latest features, better security, and performance improvement.
- Aligned with Red Hat’s long-term support policies.
- Keeps infrastructure(OS support) in line with modern IT needs and enterprise standards.

Use case: As a user, I would like to deploy IBM Cloud Infrastructure Center on RHEL9.4 Operating System.

Solution: Enabled support for hosting ICIC on RHEL 9.4 Operating System.

RHEL 9.4 (BASE OS) Support Enablement



Cloud
Infrastructure
Center –
Standalone

IBM Cloud Infrastructure Center						
Configuration Messages Requests						
Hosts						
Hosts Host Groups						
Refresh Add Host Export CSV Edit Host Remove Host Manage Existing Virtual Machines						
No filter applied						
Name	Virtual Machines *	Processors	Memory (GB)	Health	Hypervisor Type	
KVM_RHEL8	1	1 Used (320 Total)	1 Used (31 Total)	OK	KVM	
KVM_RHEL9	0	0 Used (160 Total)	0 Used (125 Total)	OK	KVM	
ZVM_RHEL8	1	1 Used (768 Total)	4 Used (2,250 Total)	OK	z/VM	
ZVM_RHEL9	1	1 Used (320 Total)	4 Used (1,536 Total)	OK	z/VM	

Storage Migration Utility


Enables seamless migration between PBHA (Policy Based High Availability) supported storage systems through a structured, step-by-step process.

Business value :

- Reduces Operational Risk: Automates validation and detects inconsistencies, ensuring reliable migrations with zero disruption to business.
- Infrastructure Modernization: Enables smooth transition to newer storage systems, aligning IT with evolving business needs and performance goals.

Use case: If user want to upgrade or replace existing PBHA supported Storage Platform with downtime.

Solution: Storage migration utility will help in streamlining the migration process, mitigating inconsistencies.



Multiple Disk Pool Support

Enables Support for managing and utilizing multiple disk pools in z/VM Hypervisor efficiently.

Business value :

- Improved Disk Utilization: Utilize all the available ECKD and FBA disk pool type configured on the z/VM hypervisor.
- Customer-oriented: ICIC provides the valid disk pool options to the user for deploying their instance on the z/VM Hypervisor

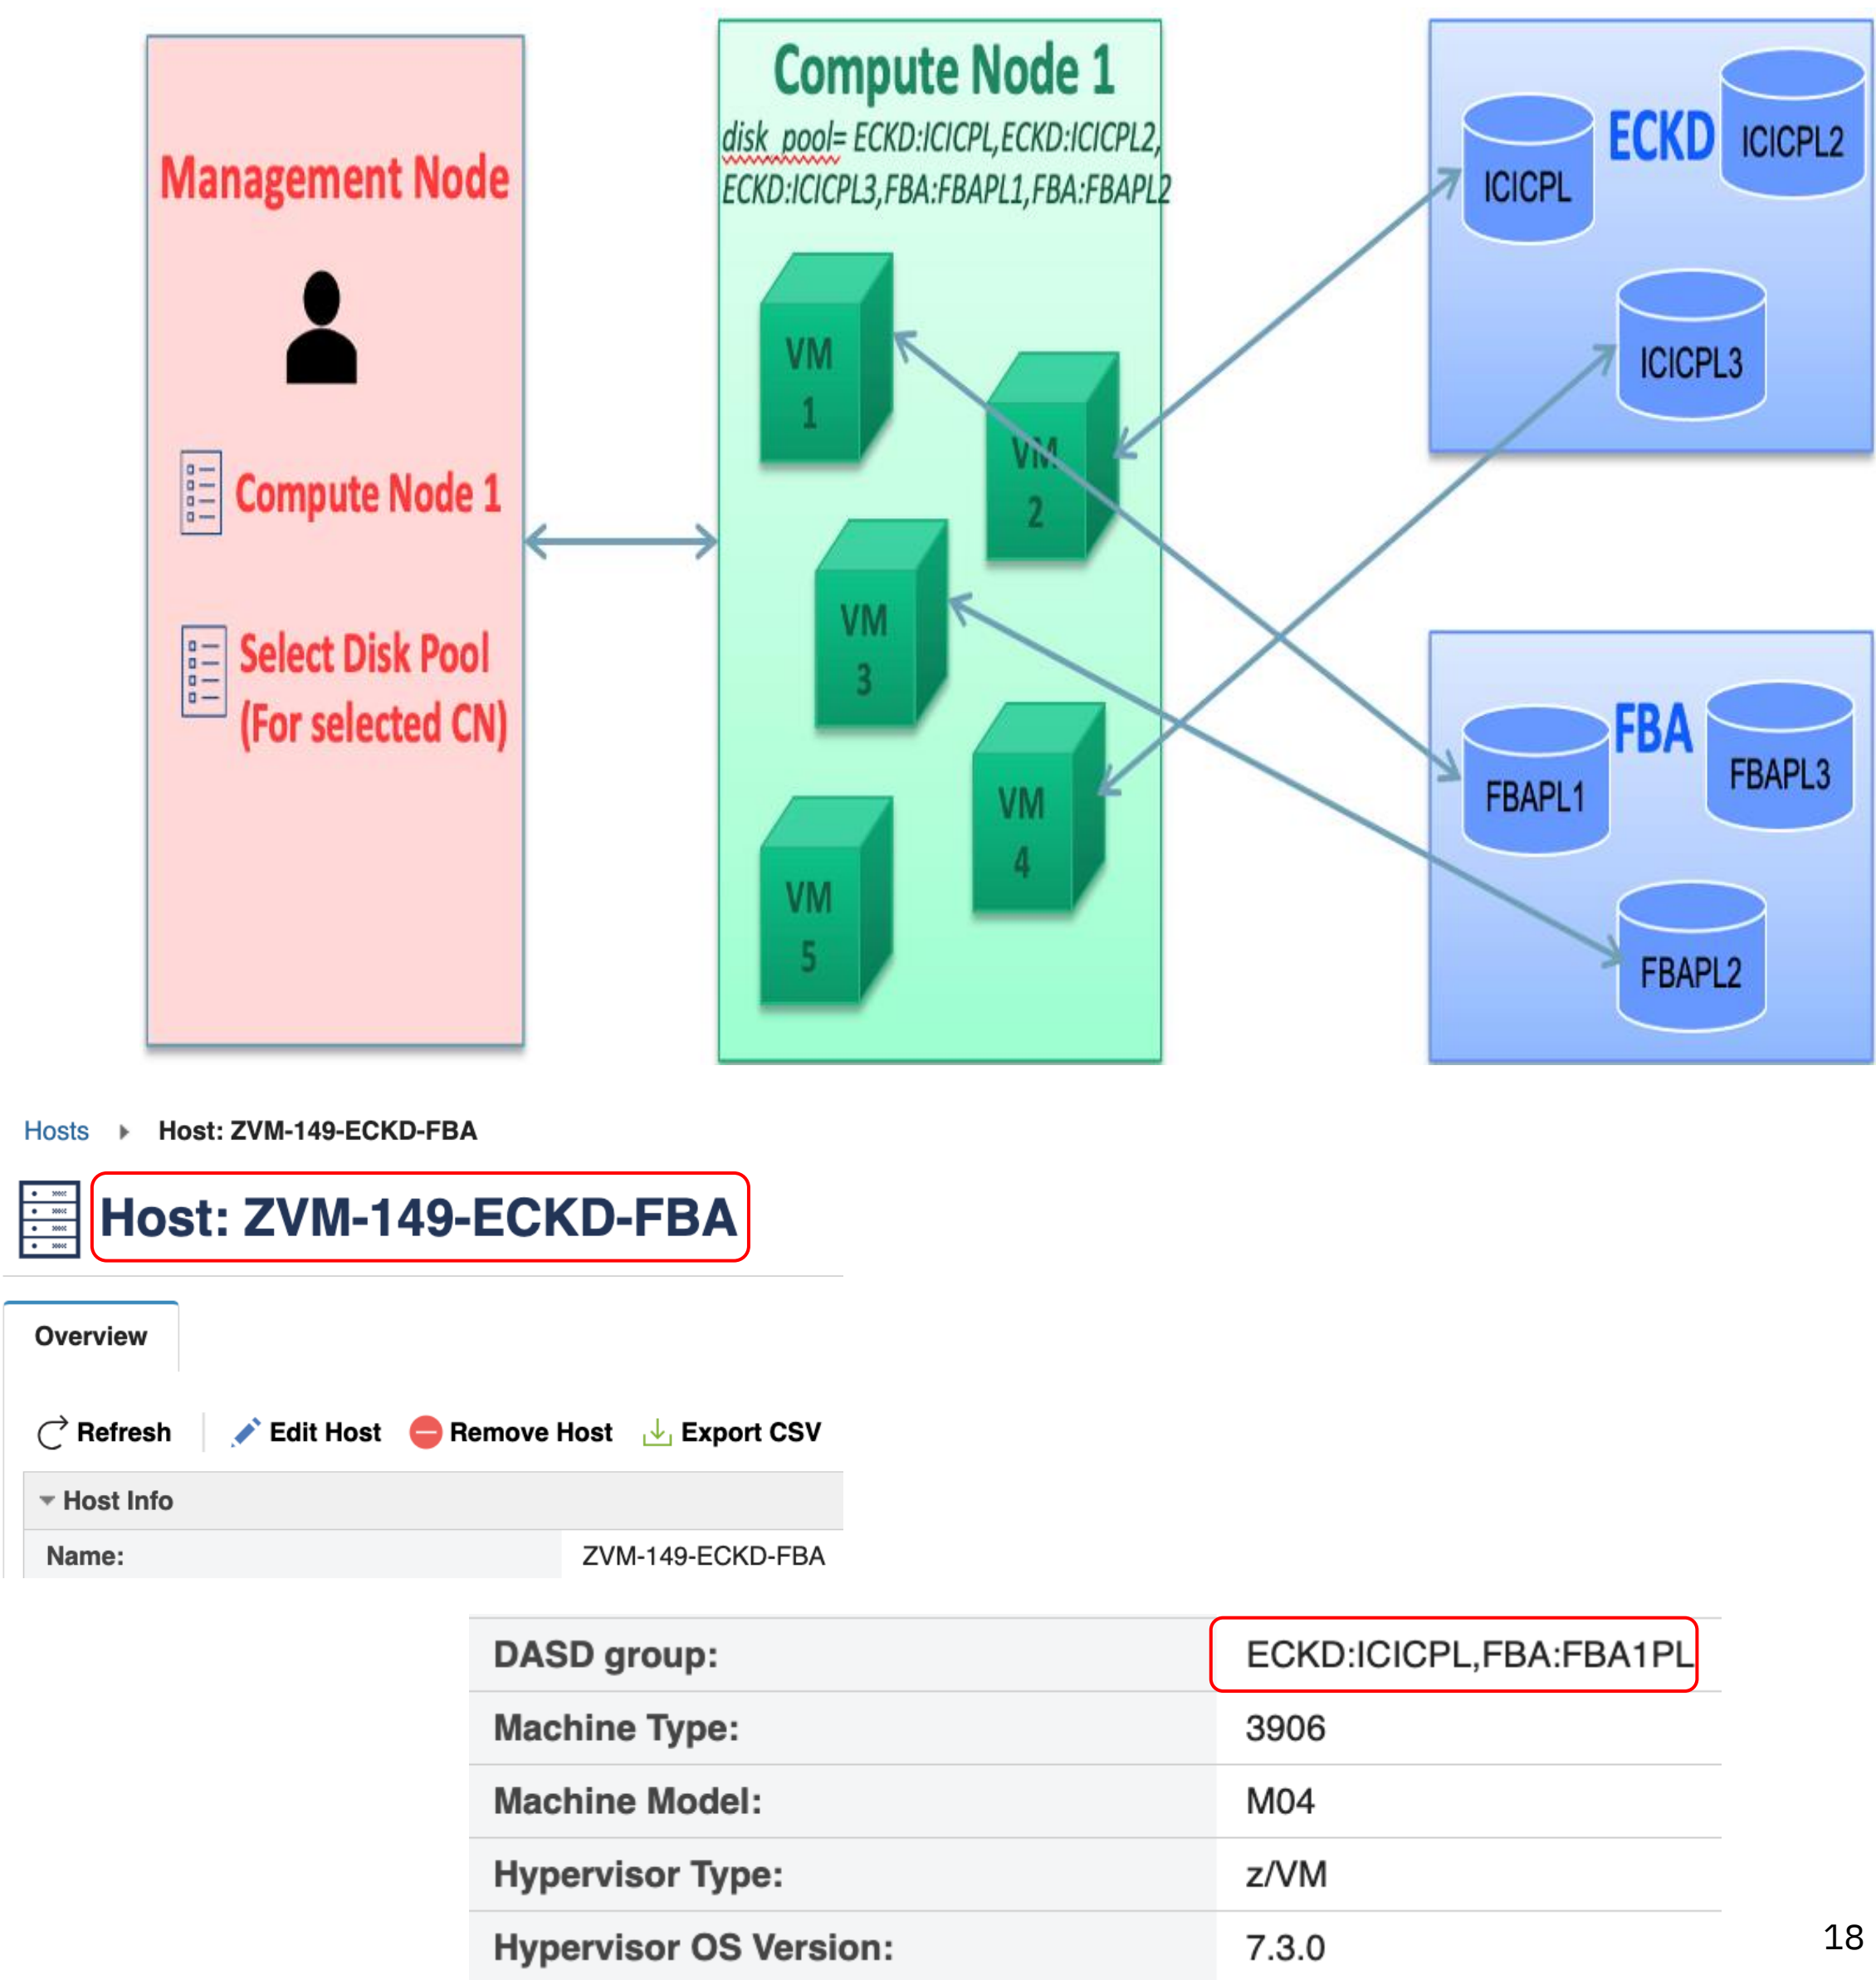
Use case: User want to utilize all the disk pools on the z/VM Hypervisor

Solution: Multiple disk pool support will help user to utilize all the disk pool resources on the z/VM compute node.


Key Points :

- Support the maximum of 10 disk pool.
- Supported disk pool types: ECKD and FBA.
- UI filters ECKD / FBA pools based on the Image type.
- Enhanced Scheduler for Best Fit Disk pool during deployment.

Optimize Disk Pool Handling for Virtual Machines



Storage



Local

- Managed by hypervisor
- Allocated as part of virtual machine deployment
- Lifecycle dependent on owning virtual machine

Persistent

- Storage object created on storage device
- Dynamic access rights and bindings between vm and storage objects
- Volumes available after vm decommissioned

Local Storage

z/VM

Create DASD group in DirMaint

Enter DASD group during ICIC add host

ICIC will create the disk from the DASD group and adds to the user directory of deployed guest

Disk is destroyed when VM is decommissioned

```
:GROUPS.  
  *GroupName RegionList  
  $ZVMRES     V6172I  
  $COMMON     V6172C  
  $RELEASE    V6172P  
  USER        V61M01  
  ICICPOOL    ICA414 ICA415 ICA416 ICA427 ICA428 ICA429  
  ICICPOOL    ICA41A ICA41B ICA41C  
  ICICPOOL    ICA41D ICA41E ICA41F ICA420 ICA421 ICA422  
  ICICPOOL    ICA423 ICA424 ICA425 ICA426 ICA427 ICA428 ICA429  
  ICICPOOL    ICA42A ICA42B ICA42C ICA42D  
  * EAVs  
  ICICPOOL    IC9100 IC9101 IC9102 IC9103 IC9104 IC9105 IC9106  
  ICICPOOL    IC9107 IC9108 IC9109 IC910A IC910B IC910C IC910D  
  ICICPOOL    IC910E IC910F IC9110 IC9111 IC9112 IC9113 IC9114  
  ICICPOOL    IC9115 IC9116 IC9117 IC9118 IC9119 IC911A IC911B  
  ICICPOOL    IC911C IC911D IC911E IC911F IC9120 IC9121 IC9122  
  ICICPOOL    IC9123 IC9124 IC9125 IC9126 IC9127 IC9128 IC9129  
  ICICPOOL    IC912A IC912B IC912C IC912D IC912E IC912F IC9130  
  ICICPOOL    IC9131 IC9132 IC9133 IC9134 IC9135 IC9136 IC9137  
  ICICPOOL    IC9138 IC9139 IC913A IC913B IC913C IC913D IC913E  
  ICICPOOL    IC913F IC9140 IC9141 IC9142 IC9143 IC9144 IC9145  
  ICICPOOL    IC9146 IC9147 IC9148 IC9149 IC914A IC914B IC914C  
  ICICPOOL    IC914D IC914E IC914F IC9150 IC9151 IC9152 IC9153  
:END.
```

Local Storage

z/VM

Create DASD group in DirMaint

Enter DASD group during ICIC add host

ICIC will create the disk from the DASD group and adds to the user directory of deployed guest

Disk is destroyed when VM is decommissioned

Add Host

Specify the details for host registration.

Host management type:

☒ z/VM ☐ KVM

* Hostname or IP address:

192.168.1.1

* User ID:

root

Display name: ?

ComputeNode2

Authentication type:

☒ Password ☐ SSH key

* Vswitch Name:

VSW1

* Password:

.....|

DASD group ?

ECKD:ICICPOOL

FCP vHBA Devices ?

+ Add Path Edit Path X Delete Path

Path Number	Single FCP Devices	Range FCP Devices	Total FCP Device Count
There is no FCP device path.			

Total: 0 Selected: 0

Reset

Add Host

Cancel

Local Storage

z/VM

Create DASD group in DirMaint

Enter DASD group during ICIC add host

ICIC will create the disk from the DASD group and adds to the user directory of deployed guest

Disk is destroyed when VM is decommissioned

Images ▶ Deploy rhel83

Deploy rhel83

Collocation rule: ?

None ▼

Key pair: ?

None ▼

Specifications

Compute template: ?

Tiny ▼

* Processors	1
* Memory (MB)	4,096
* Disk size (GB)	10
Ephemeral size (GB)	0
Swap size (MB)	0
Instance Extra Specs	{ }

*Disk size of 0 will result in the image size being set as the disk size

Local Storage

z/VM

Create DASD group in DirMaint

Enter DASD group during ICIC add host

ICIC will create the disk from the
DASD group and adds to the user
directory of deployed guest

Disk is destroyed when VM is decommissioned

```
USER CIC00183 LBYONLY 4G 64G G
    INCLUDE ZCCDFLT
    COMMAND SET VCONFIG MODE LINUX
    COMMAND DEFINE CPU 00 TYPE IFL
    COMMAND DEF STOR INITIAL STANDBY REMAINDER
    IPL 0100
    LOGONBY MAINT
    MACHINE ESA 32
    SHARE RELATIVE 100
    NICDEF 1000 TYPE QDIO DEVICES 3 MACID 5C967D
LAN SYSTEM V55SWCH
    NICDEF 1000 PORTTYPE ACCESS
    NICDEF 1000 VLAN 133
    MDISK 0100 3390 43693 14564 V55CI1 MR
```

Local Storage

z/VM

Create DASD group in DirMaint

Enter DASD group during ICIC add host

ICIC will create the disk from the DASD group and adds to the user directory of deployed guest

Disk is destroyed when VM is decommissioned

Delete

Before deleting virtual machine **vmw3-z-w01** in **running** status, select the attached storage volumes that should also be deleted.

Delete the following volumes:

☐ Keep volumes

☒ All volumes (0)

☐ Select volumes

Volumes that are attached to multiple virtual machines or in a consistency group cannot be deleted and are not included above.

Filter

No filter applied

Name	Size (GB)	Description	Boots
No items to display			

Total: 0 Selected: 0


Delete

Cancel


@ 2026 IBM Corporation

24

Local Storage Architecture



Local Storage Architecture




Local Storage

KVM

Map disk to KVM host (Ficon, FCP, GPFS, NFS)

During deployment disk is created through xml definition

```
<disk type='file' device='disk'>
  <driver name='qemu' type='qcow2' cache='none' />
  <source file='/var/lib/libvirt/images/nova/instances/ef3ad432-ae41-49fa-80e5-95dafd6340f0/disk' index='2' />
  <backingStore type='file' index='3'>
    <format type='raw' />
    <source
file='/var/lib/libvirt/images/nova/instances/_base/d38a968c459f026dac001375d167668c1d1195e1' />
    <backingStore />
  </backingStore>
  <target dev='vda' bus='virtio' />
  <alias name='virtio-disk0' />
  <address type='ccw' cssid='0xfe' ssid='0x0' devno='0x0000' />
</disk>
```



Persistent Storage

z/VM

Add Storage Provider

FCP pool and templates

On deployment, volume created dynamically

FCP dedicated to the VM and mapped to created volume

Define VM user directory to boot from FCP

Add Storage

For each new storage provider, a default storage template is created. You can modify the template after the storage provider has been added.

Specify a storage controller

* Agent Node ?

iciccomp-pbm-ihost-com

* Availability zone

Default_Group

* Type:

IBM Storage FlashSystem (formerly

* Hostname or IP address:

129.40.156.135

* User ID:

icicuser

* Display name: ?

fs93

Authentication type:

☒ Password

☐ SSH key

* Password:

.....

Connect

Persistent Storage

z/VM

Add Storage Provider

FCP pool and templates

On deployment, volume created dynamically

FCP dedicated to the VM and mapped to created volume

Define VM user directory to boot from FCP

Add Storage

For each new storage provider, a default storage template is created. You can modify the template after the storage provider has been added.

Specify a storage controller

Type: IBM Storage FlashSystem
Name: fs93



Select a storage pool for the default template

The selected storage pool is used in the default storage template. To use a different storage pool, create a new template on the *Configuration* page.

Name	Capacity (GB)	Available (GB) ▼
Pool0	17,759	16,294
ICIC	500	367

Total: 2 Selected: 1

Add Storage

Cancel

Persistent Storage

z/VM

Add Storage Provider

FCP pool and templates

On deployment, volume created dynamically

FCP dedicated to the VM and mapped to created volume

Define VM user directory to boot from FCP

IBM Cloud Infrastructure Center

Configuration

Messages

Requests

root (ibm-default)

IBM

Configuration

FCP Multipath Templates

Add FCP Multipath Template

Home

1

Storage

6

Hosts

30

Networks

4

VMs

5

Storage Providers

1

Add FCP Multipath Template

* Host

iciccomp-pbm-ihost-c

* Name

fcp_dev_template

Description

FCP vHBA Devices

+ Add Path

Edit Path

Delete Path

Path Number	Single FCP Devices	Range FCP Devices	Total FCP Device Count
1		1a02 - 1a0f	14
2		1b02 - 1b0f	14

Default for Storage Providers

Storage Provider

sv7000

☒

Default for Host

☐

Current default FCP multipath template: Unknown

@ 2026 IBM Corporation

30

Persistent Storage

z/VM

Add Storage Provider

FCP pool and templates

Deploy Boot From Volume Image

FCP dedicated to the VM and mapped to created volume

Define VM user directory to boot from FCP

1. Create volume on storage provider
2. Download image from glance service
3. Map the created volume to the compute node
4. Use **dd** to copy the image to the volume
5. Unmap the volume from the compute node

Persistent Storage

z/VM

Add Storage Provider

FCP pool and templates

Deploy Boot From Volume Image

FCP dedicated to the VM and mapped to created volume

Define VM user directory to boot from FCP

Bootable has been volume created

1. Allocate and reserve FCP devices from FCP template
2. Map the volume to allocated devices
3. Dedicate FCP devices on the compute node
4. Select valid path of FCP and WWPN to access the LUN
5. Mount the volume to tmp directory
6. Update zipl.conf with volume path
7. Unmount and Undedicate FCPs

Persistent Storage

z/VM

Add Storage Provider

FCP pool and templates


Deploy Boot From Volume Image

FCP dedicated to the VM and mapped to created volume

Define VM user directory to boot from FCP

```
USER HLP00006 LBYONLY 4G 64G G
INCLUDE ZCCDFLT
COMMAND SET VCONFIG MODE LINUX
COMMAND DEFINE CPU 00 TYPE IFL
COMMAND DEF STOR INITIAL STANDBY REMAINDER
IPL LOADDEV
LOADDEV DEVICE 1a0a <-- base IPL device
LOADDEV PORTname 50050768103391d8 <-- base IPL port
LOADDEV LUN 0000000000000000
LOADDEV SCSI ALTERNATE 1a0a PORT 50050768102491e1 <-- alternate path1
LOADDEV SCSI ALTERNATE 1b0a PORT 50050768102391e1<-- alternate path2
LOADDEV SCSI ALTERNATE 1b0a PORT 50050768102391d8<-- alternate path3
LOGONBY DEVCORE2
MACHINE ESA 32
SHARE RELATIVE 100
DEDICATE 1A0A 1A0A
DEDICATE 1B0A 1B0A
NICDEF 1000 TYPE QDIO DEVICES 3 MACID 8A1355 LAN SYSTEM VSICIC
```

Persistent Storage Architecture




Hybrid Storage

z/VM

Root disk from DASD group

Data volume on Storage Device



Thank you

Resources

- Watch [IBM Cloud Infrastructure Center](#) webpage
- Read the [technical blogs and announcements](#)
- Check out the [technical details](#) at the IBM Documentation
- Get started using the [Content solution](#)
- Leverage the IBM Redbook®: [Hybrid cloud with on-premises cloud on IBM Z or IBM® LinuxONE](#)

© Copyright IBM Corporation 2026

IBM and the IBM logo are trademarks of IBM Corporation, in the United States and/or other countries. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/trademark.

This document is current as of the initial date of publication and may be changed by IBM at any time.

Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

Examples presented are illustrative only. Actual results will vary based on client configurations and conditions and, therefore, generally expected results cannot be provided.

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

