


# IBM Cloud Infrastructure Center

## Deep Dive

**Michael Snihur**

Solutions Architect

[snihurm@us.ibm.com](mailto:snihurm@us.ibm.com)

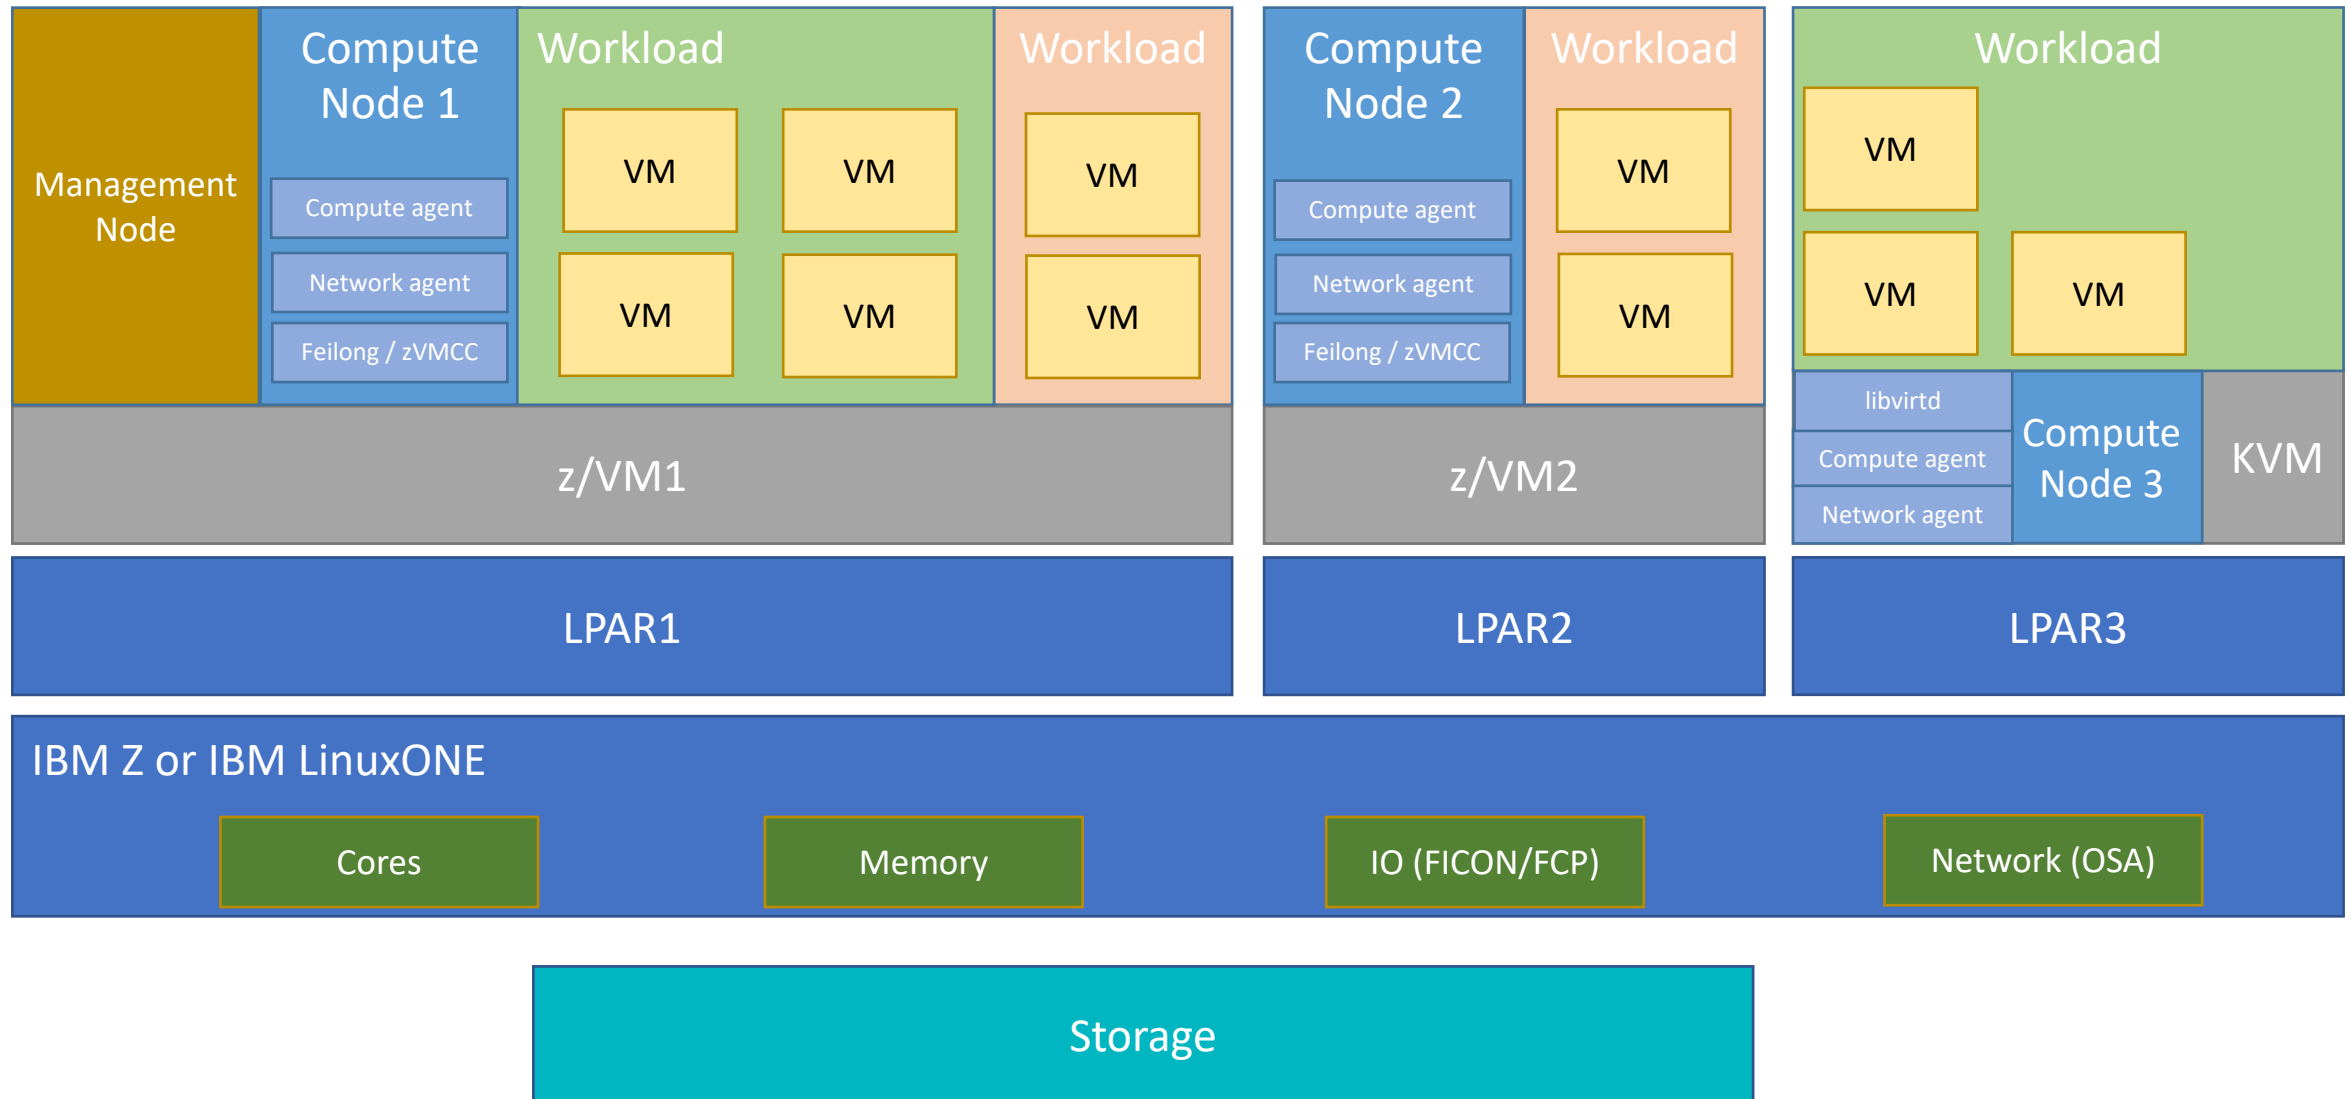




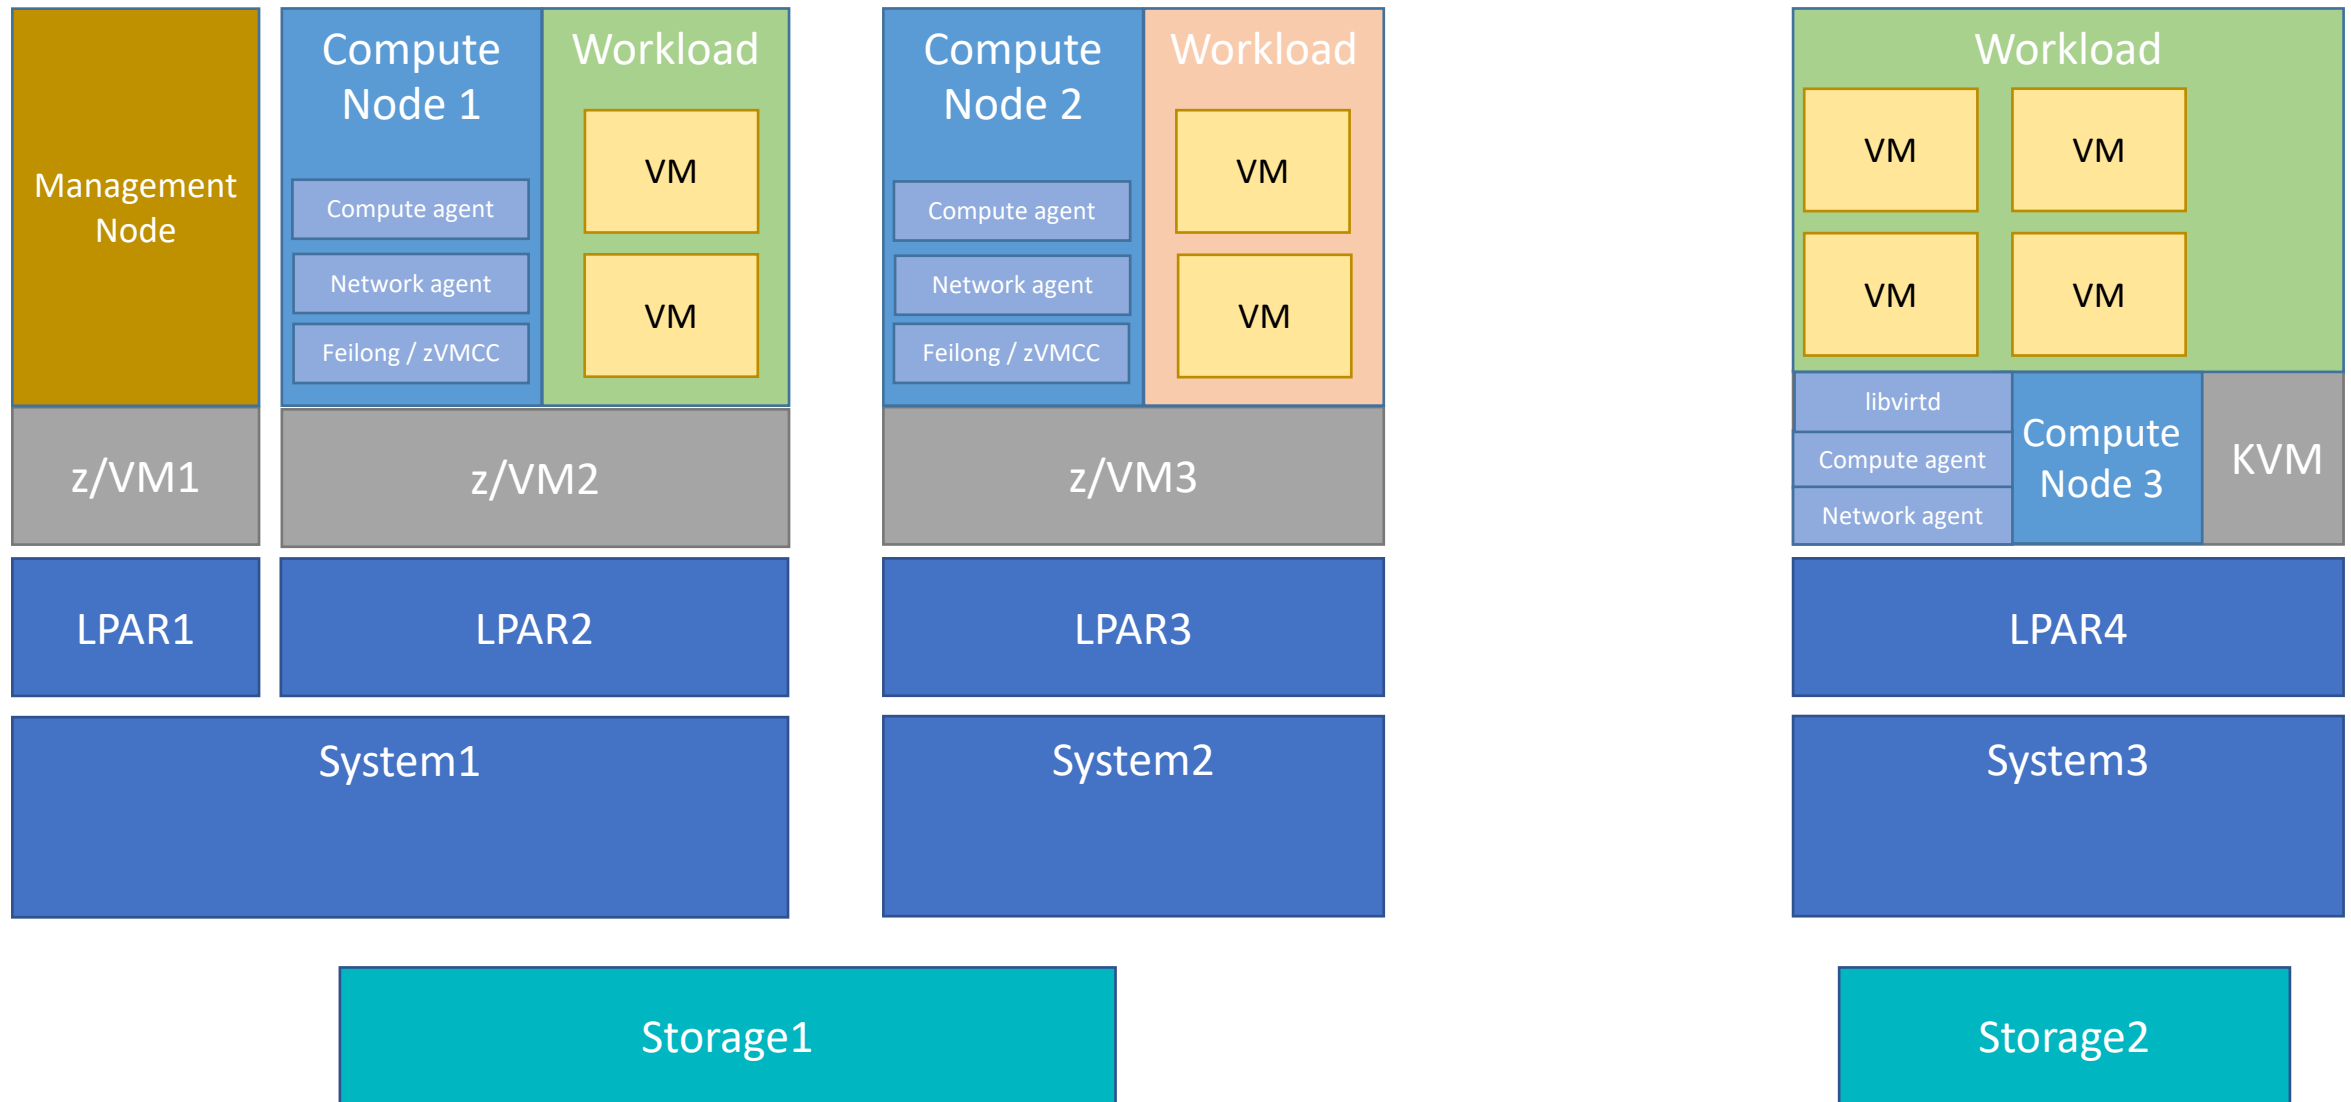
# Agenda

- Overall ICIC architecture
- Storage types
- Boot From Volume Image types
- Migration
- Allocation model and overcommit

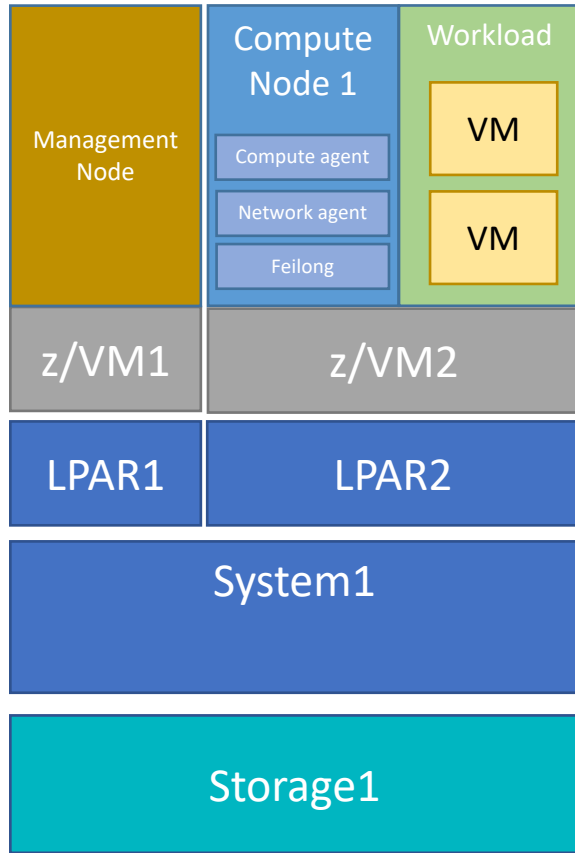
# ICIC architecture – single system



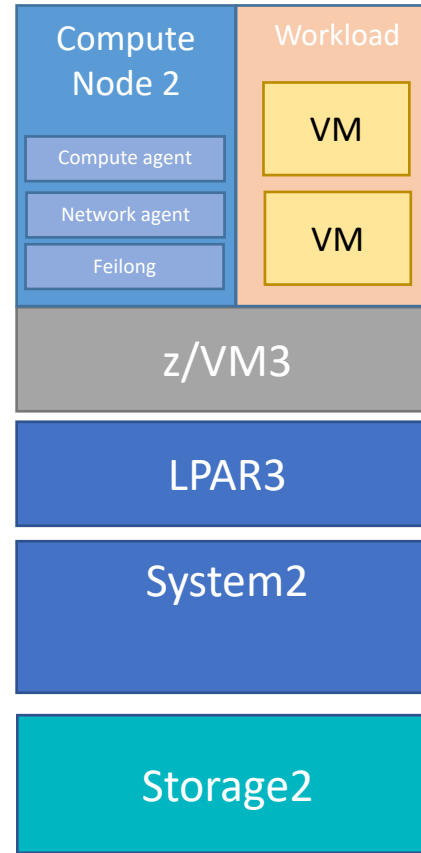
# ICIC architecture – multi system



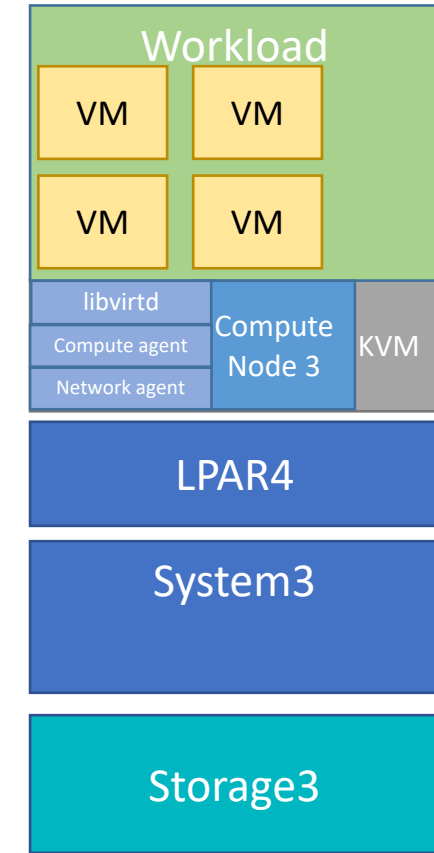
# ICIC architecture – multi site



Site 1



Site 2



Site 3

# 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



# 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 add it 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:  \* User ID:

Display name:  Authentication type:  
 Password  SSH key

\* Vswitch Name:  \* Password:

**DASD group**

FCP vHBA Devices

+ Add Path Edit Path 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 add it 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 add it 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 SYS
  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

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

Filter

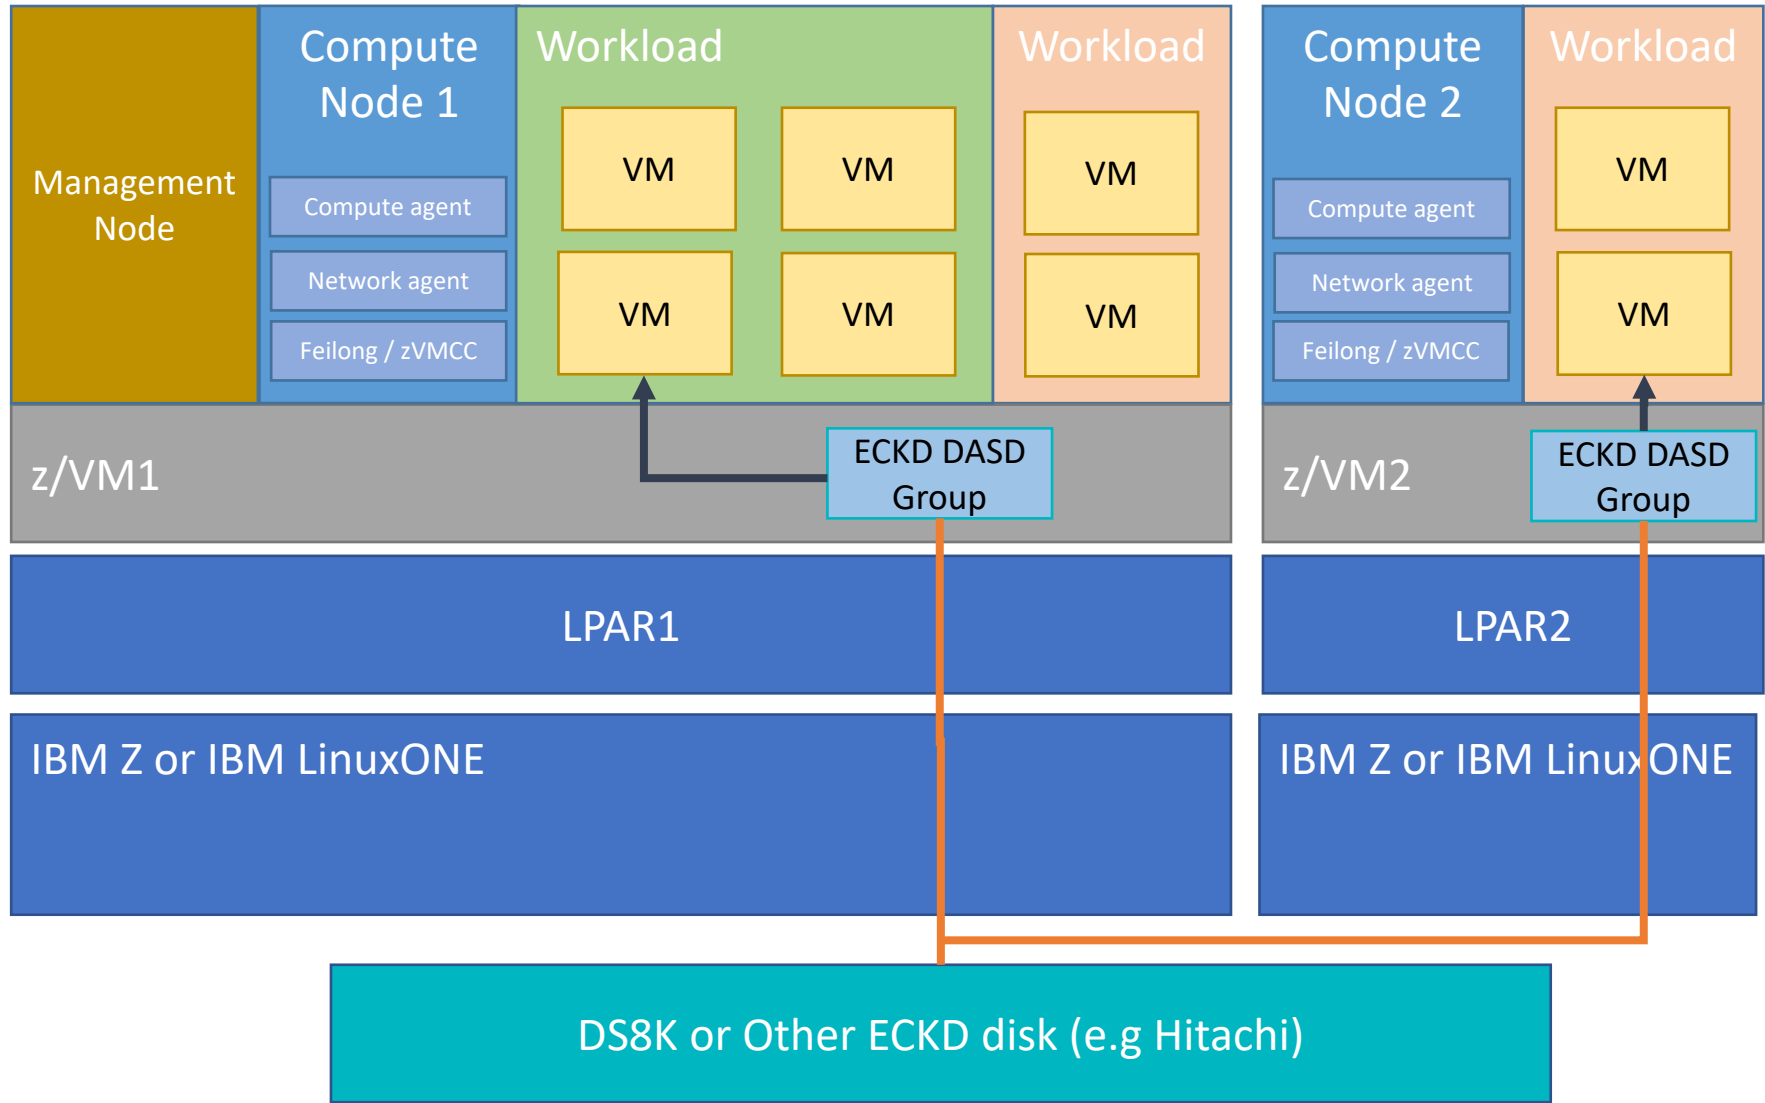
*No filter applied*

Name	2 ▲	Size (GB)	Description	1 ▼ Boota
No items to display				

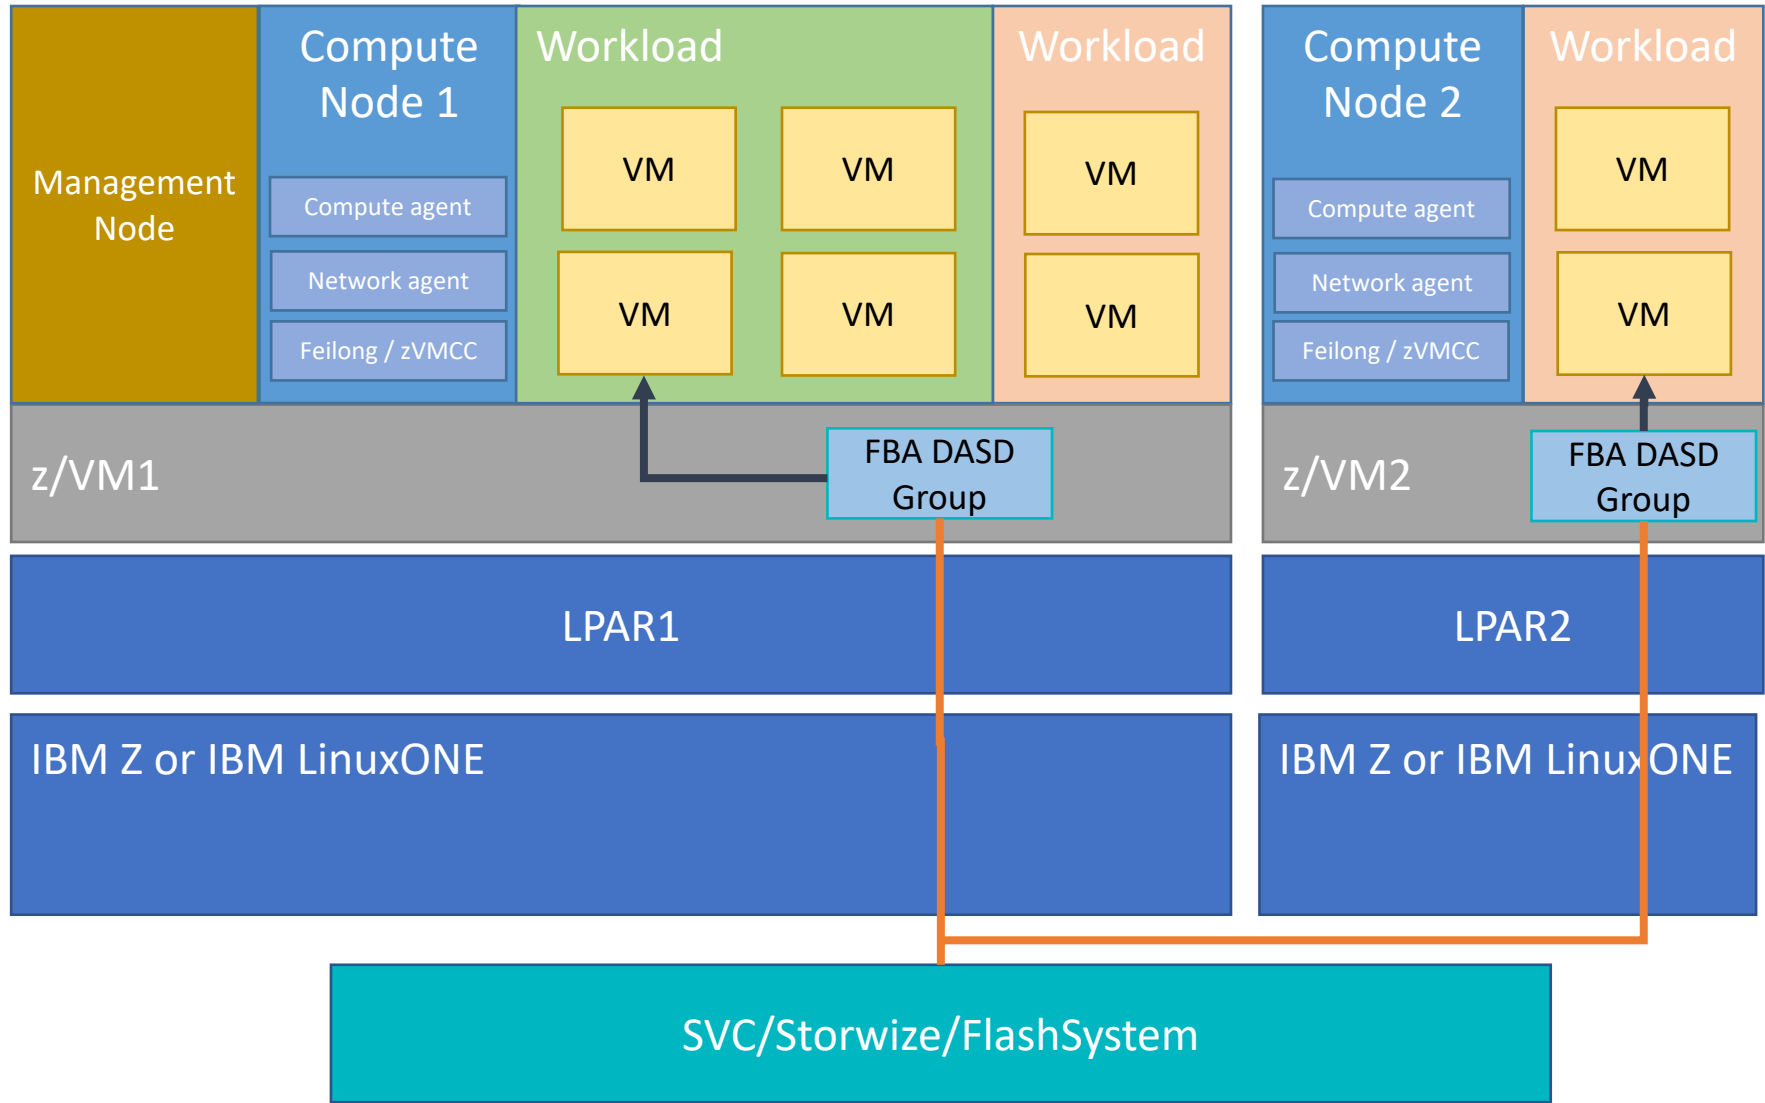
Total: 0 Selected: 0

**Delete**

# Local Storage



# Local Storage





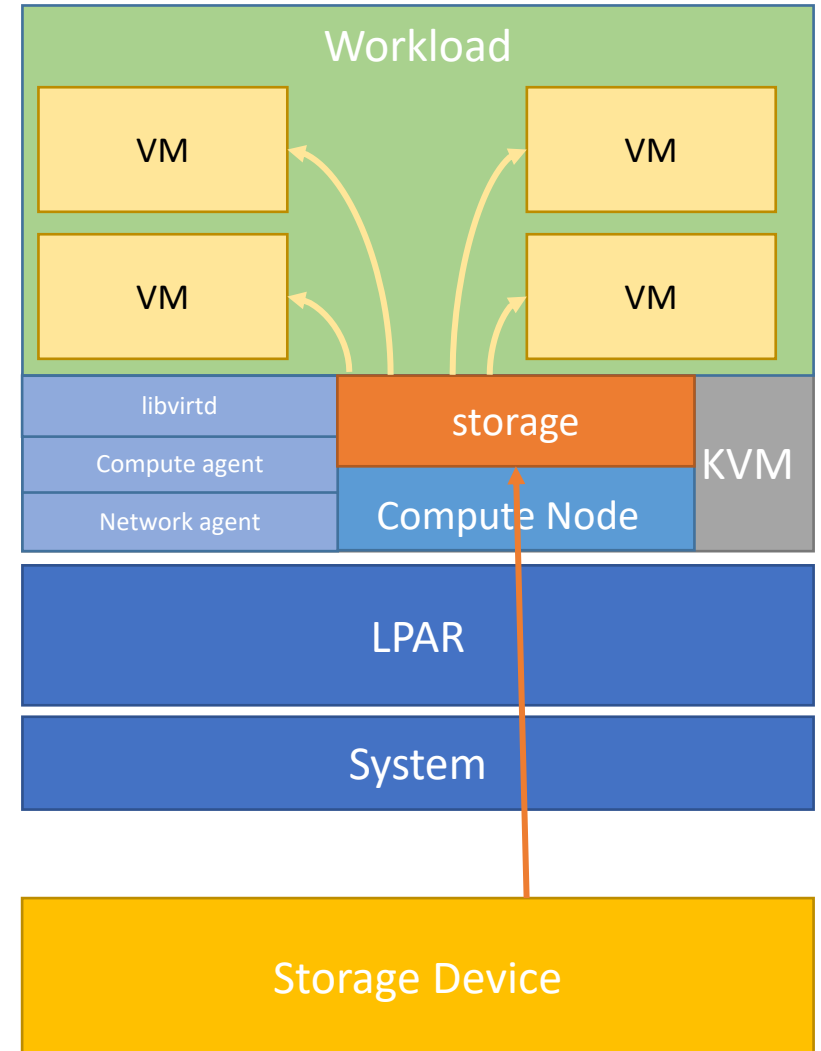
# Local Storage

## 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





# 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	* Availability zone
<input type="text" value="iciccomp-pbm-ihost-com"/>	<input type="text" value="Default_Group"/>
* Type:	
<input type="text" value="IBM Storage FlashSystem (formerly)"/>	
* Hostname or IP address:	* User ID:
<input type="text" value="129.40.156.135"/>	<input type="text" value="icicuser"/>
* Display name:	Authentication type:
<input type="text" value="fs93"/>	<input checked="" type="radio"/> Password <input type="radio"/> SSH key
	* Password:
	<input type="text" value="....."/>

---



# 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

---

# 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

The screenshot shows the 'Add FCP Multipath Template' configuration page in the IBM Cloud Infrastructure Center. The page includes a breadcrumb trail: Configuration > FCP Multipath Templates > Add FCP Multipath Template. The main heading is 'Add FCP Multipath Template'. The configuration fields are as follows:

- \* Host:** A dropdown menu showing 'icicomp-pbm-ihost-c'.
- \* Name:** A text input field containing 'fcp\_dev\_template'.
- Description:** An empty text area.
- FCP vHBA Devices:** A section with a '+' icon and a help icon. Below it are buttons for '+ Add Path', 'Edit Path', and 'Delete Path'.
- Table:** A table with 4 columns: Path Number, Single FCP Devices, Range FCP Devices, and Total FCP Device Count. It contains two rows of data.
- Default for Storage Providers:** A section with a help icon. It contains a table with one row: Storage Provider 'sv7000' with a checked checkbox.
- Default for Host:** A section with a help icon. It contains a checkbox that is unchecked and the text 'Current default FCP multipath template: Unknown'.

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

Storage Provider	
sv7000	<input checked="" type="checkbox"/>



# 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

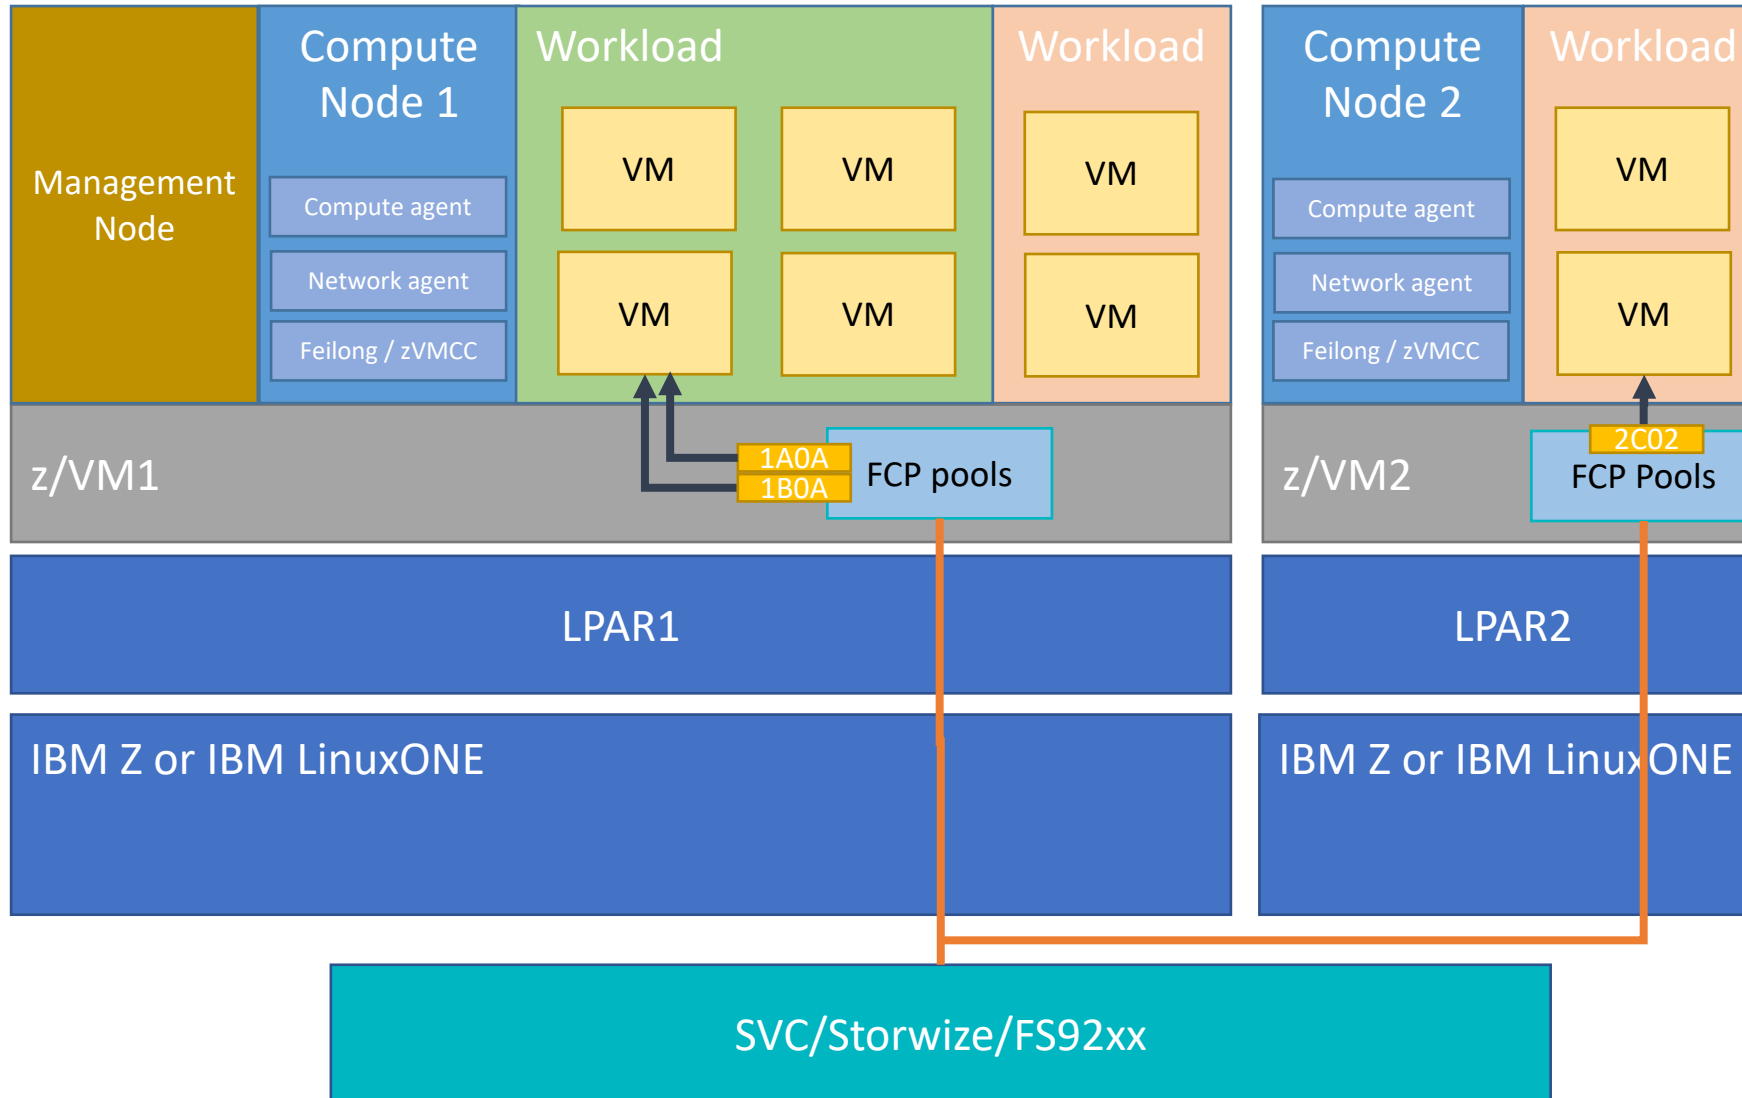
FCP dedicated to the VM and mapped to created volume

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





# Persistent Storage

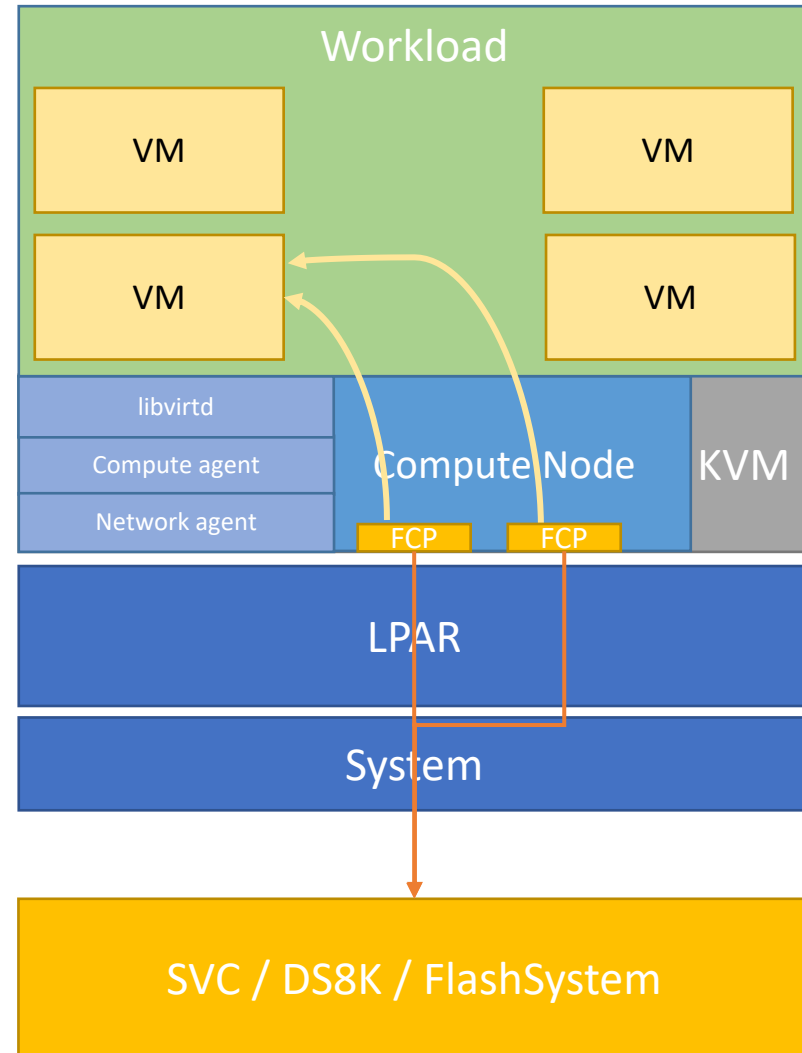
## KVM with FCP

Create volume on storage

Map to KVM host using multipath

Map storage from the KVM host to the guest VM

```
<disk type='block' device='disk'>  
  <driver name='qemu' type='raw' cache='none' io='native'/>  
  <source dev='/dev/disk/by-id/dm-uuid-mpath-  
36005076308ffd2cf000000000000290a' index='6'/>  
  <backingStore/> <target dev='vda' bus='virtio'/>  
  ...  
</disk>
```







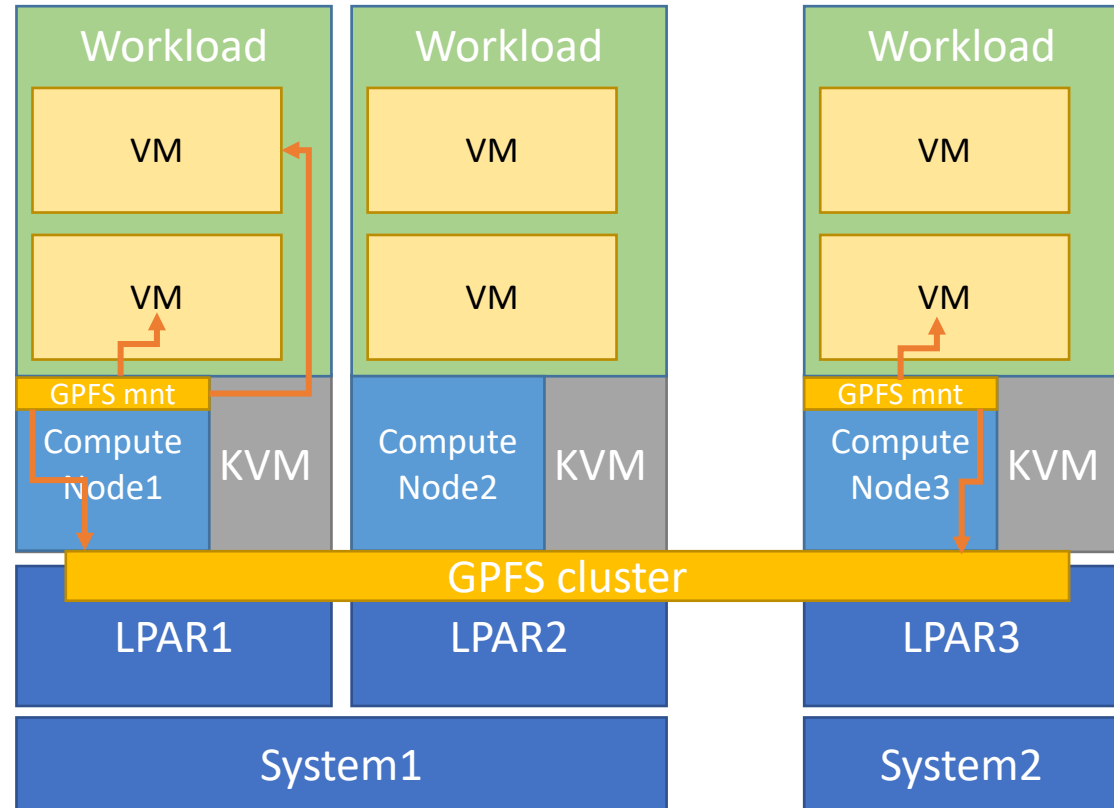
# Persistent Storage

## KVM with GPFS

Create volume file through GPFS backend directly on kvm host

```
<gpfs mnt>/icic/<volume-id>
```

map the created volume from KVM host to the KVM guest



```
<disk type='file' device='disk'>  
  <driver name='qemu' type='raw' cache='none' />  
  <source file' index='1' />  
  <backingStore />='/gpfs/icic_gpfs/icic/consisgroup-1dfe0bfa-bb6c-41de-b98e-53db9164b8b3/volume-rhel90_kvm12-boot-29-04ece809-d4ab'  
  <target dev='vda' bus='virtio' />  
  ....  
</disk>
```

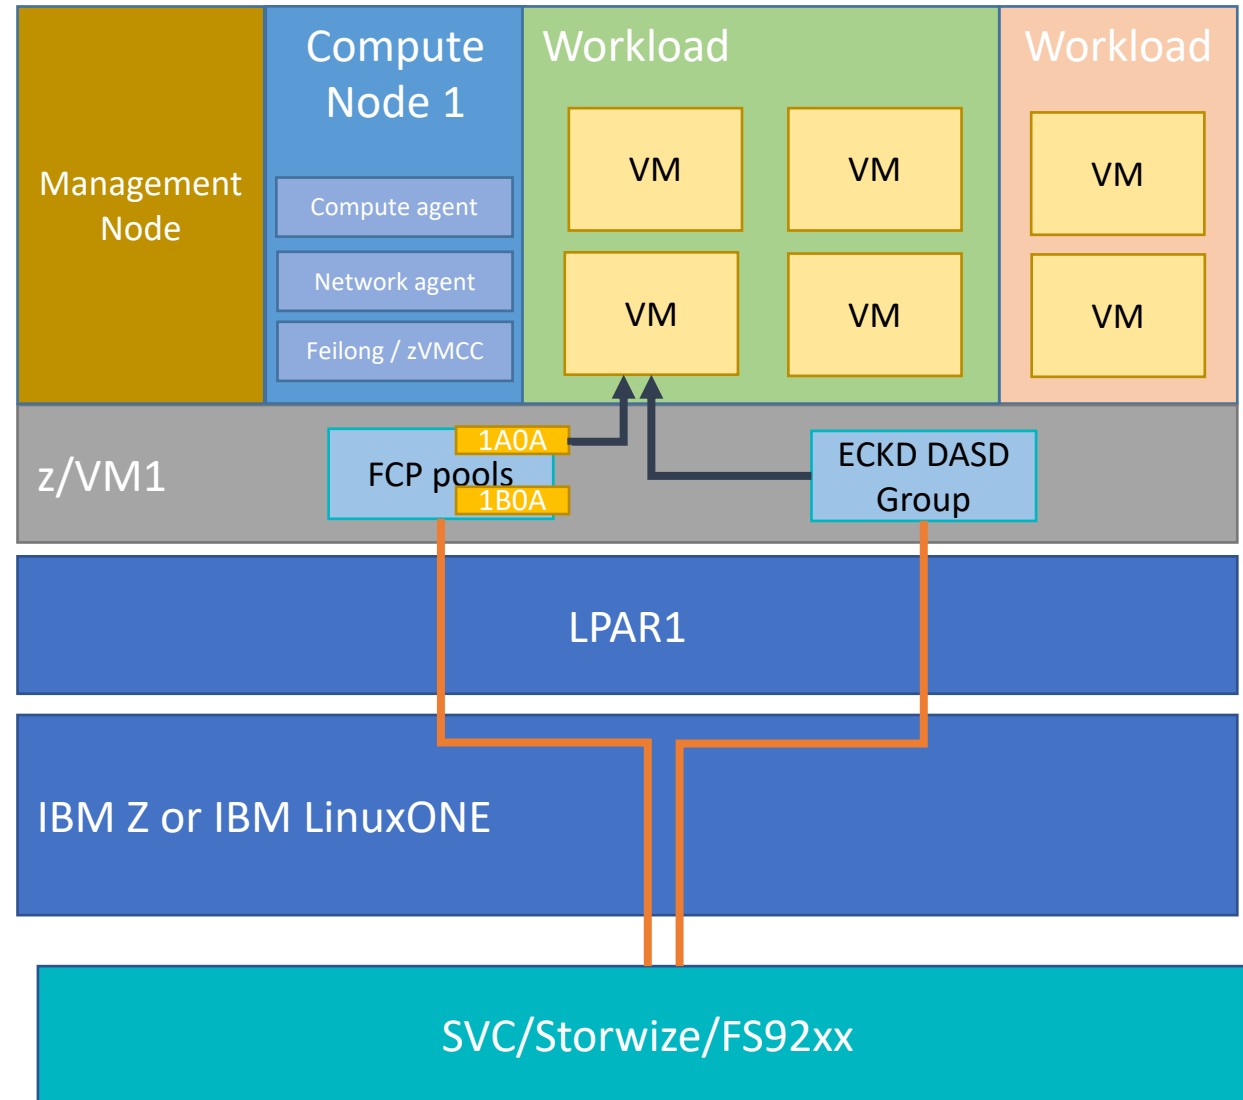
# Hybrid Storage



## z/VM

Root disk from DASD group

Data volume on Storage Device





# Boot from Volume Image Types

## File Backed Image

- BFV Image file uploaded to ICIC
- During deployment, a volume is created, and the image is copied using 'dd'

## Volume Backed Image

- Volume created from BFV image
- During deployment, a new volume is created
- A temporary flashcopy mapping is used to copy the volume image into the new volume

## Snapshot Backed Image

- Capture snapshot image from running BFV server
- Image volume is created and captured vm is flashcopied to newly created volume

# Volume Backed Image



## Volume Backed Image

- Volume created from BFV image
- During deployment, a new volume is created
- A temporary flashcopy mapping is used to copy the volume image into the new volume

The screenshot shows the IBM Cloud Infrastructure Center interface for creating a storage volume. The 'Storage' section is active, and the 'Create' button is highlighted. A modal window is open for configuration. In the modal, the 'Storage template' is set to 'fl63 base template', the 'Volume name' is 'jacky\_boot\_vol', the 'Size (GB)' is 10, and the 'Number of volumes' is 1. The 'Bootable volume' checkbox is checked. A table below shows the selected image: 'BFV\_File\_Img' with state 'Active', operating system 'RHEL9.2', and type 'Image'. The 'Create Volume' button is highlighted at the bottom of the modal.

IBM Cloud Infrastructure Center Configuration Messages Requests

### Storage

Storage Providers Fabrics Service Nodes

Refresh Create Edit Delete

No filter applied

Name	Size (GB)
BFV_Volume_Img-boot-66	10
BFV_Volume_Img-boot-67	10
BFV_Volume_Img_capture-boot-69	10
BFV_Volume_Img_capture-data-70	1

\* Storage template: fl63 base template

\* Volume name: jacky\_boot\_vol

Description:

\* Size (GB): 10

\* Number of volumes: 1

Bootable volume  
Please select one image when this is checked.

z/VM

1 of 2 items shown. Clear filter

Name	State	Operating System	Secure Executic	Type	Hypervisor Type	Descr
BFV_File_Img	Active	RHEL9.2	False	Image	z/VM	

Create Volume Cancel



# Volume Backed Image

## Volume Backed Image

- Volume created from BFV image
- During deployment, a new volume is created
- A temporary flashcopy mapping is used to copy the volume image into the new volume

### Create Image

Specify the details for the image that you want to upload.

---

\* Hypervisor Type  
z/VM

\* Operating system:  
RHEL9

\* Image disk type: SCSI

\* Image Source: **VOLUME**

\* Disk Format: RAW

Select a volume that will compose the created image.

boot

1 of 4 items shown. [Clear filter](#)

Name	Size (GB)	State	Storage Template	Storage Provider
<b>jacky_boot_vol</b>	10	Available	fl63 base template	fl63

**Upload** Cancel

# Volume Backed Image

## Volume Backed Image

- Volume created from BFV image
- During deployment, a new volume is created
- A temporary flashcopy mapping is used to copy the volume image into the new volume

The screenshot displays the IBM Cloud console interface for deploying a Volume Backed Image (VBI). The main panel shows the 'Images' section with a table of available images. The 'BFV\_Volume\_Img' image is highlighted, and its source type is 'VOLUME'. The 'Deploy' button is highlighted in red. The right sidebar shows the deployment configuration for 'Deploy BFV\_Volume\_Img', including the 'Deploy target' (J4601), 'FCP Multipath Template' (Auto-select), 'Collocation rule' (None), and 'Key pair' (None). The 'Specifications' section lists the instance configuration: 1 processor, 4,096 MB memory, 10 GB disk size, 0 GB ephemeral size, 0 MB swap size, and Instance Extra Specs. The 'Deploy' button is highlighted in red.

Name	Source Type	State	Operating System	Secure Executi
BFV_Volume_Img_capture	SNAPSHOT	Active	RHEL9	False
BFV_Volume_Img	VOLUME	Active	RHEL9	False
BFV_File_Img	FILE	Active	RHEL9.2	False

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



# Volume Backed Image

## Volume Backed Image

- Volume created from BFV image
- During deployment, a new volume is created
- A temporary flashcopy mapping is used to copy the volume image into the new volume

# Volume Backed Image



What happens in FS9x00 backend

(a **temporary** FlashCopyMapping)

1. Create a volume with specified size
2. Create FlashCopyMapping between the new-volume and the image-volume with **autodelete**
3. Prestartfcmap
4. Startfcmap
5. Mkhost
6. mkvdiskhostmap

IBM FlashSystem 9200 FS63 Audit Log

Sequence Number	Date and Time	User Name	Command
121448	4/19/2024 3:28:10 PM	admin_icic	svctask mkvdiskhostmap -force -host BOET4601_HLPS0033-41256537 -scsi 0 volume-BFV_Volume_Img-boot-66-7b2c7a76-c34f
121447	4/19/2024 3:28:08 PM	admin_icic	svctask mkhost -force -hbawwpn c05076fec8047fc -iogrp 0 -name BOET4601_HLPS0033-41256537
121445	4/19/2024 3:27:56 PM	admin_icic	svctask startfcmap 10
121444	4/19/2024 3:27:54 PM	admin_icic	svctask prestartfcmap 10
121443	4/19/2024 3:27:53 PM	admin_icic	svctask mkfcmap -source volume-jacky_boot_vol-8ff77e5c-c7ec -target volume-BFV_Volume_Img-boot-66-7b2c7a76-c34f -copyrate 100 -autodelete
121442	4/19/2024 3:27:52 PM	admin_icic	svctask mkvdisk -name volume-BFV_Volume_Img-boot-66-7b2c7a76-c34f -mdiskgrp ICIC_CORE2 -iogrp 0 -size 10737418240 -unit b -easytier on



# Snapshot Backed Image



## Snapshot Backed Image

- Capture snapshot image from running BFV server
- Image volume is created and captured vm is flashcopied to newly created volume
- Deploy vm from snapshot volume

**Virtual Machines**

Refresh Start Stop Suspend Resume Restart Delete **Capture**

Attach Volume Manage Existing Unmanage Console Output Console Access

No filter applied

Name	Host	IP	State	Health
BFV_Volume_Img	J4601	172.26.93.29	Active	OK

### Confirm Capture

Confirm the capture of virtual machine **BFV\_Volume\_Img**. The deployable image **BFV\_Volume\_Img\_capture** will be comprised of the following volumes:

Name	Size (GB)	Description	Bootable
BFV_Volume_Img_boot-66	10		Yes
jacky_t-1	1		No

**Capture** Cancel



# Snapshot Backed Image

## Snapshot Backed Image

- Capture snapshot image from running BFV server
- Image volume is created and captured vm is flashcopied to newly created volume
- Deploy vm from snapshot volume

# Snapshot Backed Image



What is happening in FS9x00 backend

(2 **permanent** FlashCopyMappings without Consistency Group)

1. Create 2 image-volumes with the same size of the being-captured VM's volumes
2. Create 2 FlashCopyMappings between the image-volumes and the VM's volumes, one for each volume with **copyrate 0**
3. **Prestartfcmap**
4. **Startfcmap**

Note:

As no consistency group is involved during the process,

there is no guarantee for the data consistency between bootable volume and data volumes

Suggest capturing the VM after shutdown the VM to keep data consistency.

IBM FlashSystem 9200 FS63 Audit Log

Sequence Number	Date and Time	User Name	Command
124365	4/24/2024 6:48:36 AM	admin_itic	svctask mkvdisk -name snapshot-352fbe43-9afc-4ca2-9c39-2c756 -mdiskgrp ICIC_CORE2 -iogrp 0 -size 10737 -unit b -easytier on
124366	4/24/2024 6:48:36 AM	admin_itic	svctask mkfcmap -source volume-BFV_Volume_Img-boot-66-7b2c7a76-c34f -snapshot-352fbe43-9afc-4ca2-9c39-2c756 -copyrate 0
124372	4/24/2024 6:48:37 AM	admin_itic	svctask startfcmap 36
124370	4/24/2024 6:48:37 AM	admin_itic	svctask startfcmap 35
124371	4/24/2024 6:48:37 AM	admin_itic	svctask prestartfcmap 36
124369	4/24/2024 6:48:37 AM	admin_itic	svctask mkfcmap -source volume-jacky_t-1-c894d83b-93a1 -target snapshot-9f74eb9f-883c-48ab-a6a8-9d978c58fda9 -copyrate 0
124367	4/24/2024 6:48:37 AM	admin_itic	svctask prestartfcmap 35
124368	4/24/2024 6:48:37 AM	admin_itic	svctask mkvdisk -name snapshot-9f74eb9f-883c-48ab-a6a8-9d978c58fda9 -mdiskgrp ICIC_CORE2



# Snapshot Backed Image

## Snapshot Backed Image

- Capture snapshot image from running BFV server
- Image volume is created and captured vm is flashcopied to newly created volume
- Deploy vm from snapshot volume

Images > Deploy BFV\_Volume\_Img\_capture

### Deploy BFV\_Volume\_Img\_capture

Deploy target: ?

J4601

FCP Multipath Template: ?

Auto-select

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	∧

Network

Activation Input

**Deploy** Cancel

Virtual Machines > VM: BFV\_Volume\_Img\_capture

### VM: BFV\_Volume\_Img\_capture

Overview Attached Volumes

Refresh Attach Volume Detach Volume Edit Volume

No filter applied

Name	Size (GB)	State	Storage Template	Storage Provider	FCP Multipath Template
BFV_Volume_Img_capture-data-70	1	In-Use	fi63 base template	fi63	1f8a8ec6-fdfe-11ee-8dfe-02010c
BFV_Volume_Img_capture-boot-69	10	In-Use	fi63 base template	fi63	1f8a8ec6-fdfe-11ee-8dfe-02010c



# Snapshot Backed Image

## Snapshot Backed Image

- Capture snapshot image from running BFV server
- Image volume is created and captured vm is flashcopied to newly created volume
- Deploy vm from snapshot volume

# Migration of virtual machines



## z/VM

Live migrate of virtual machine to another compute node

Requires both compute nodes to be in SSI cluster

## KVM

Live migrate of virtual machine to another compute node

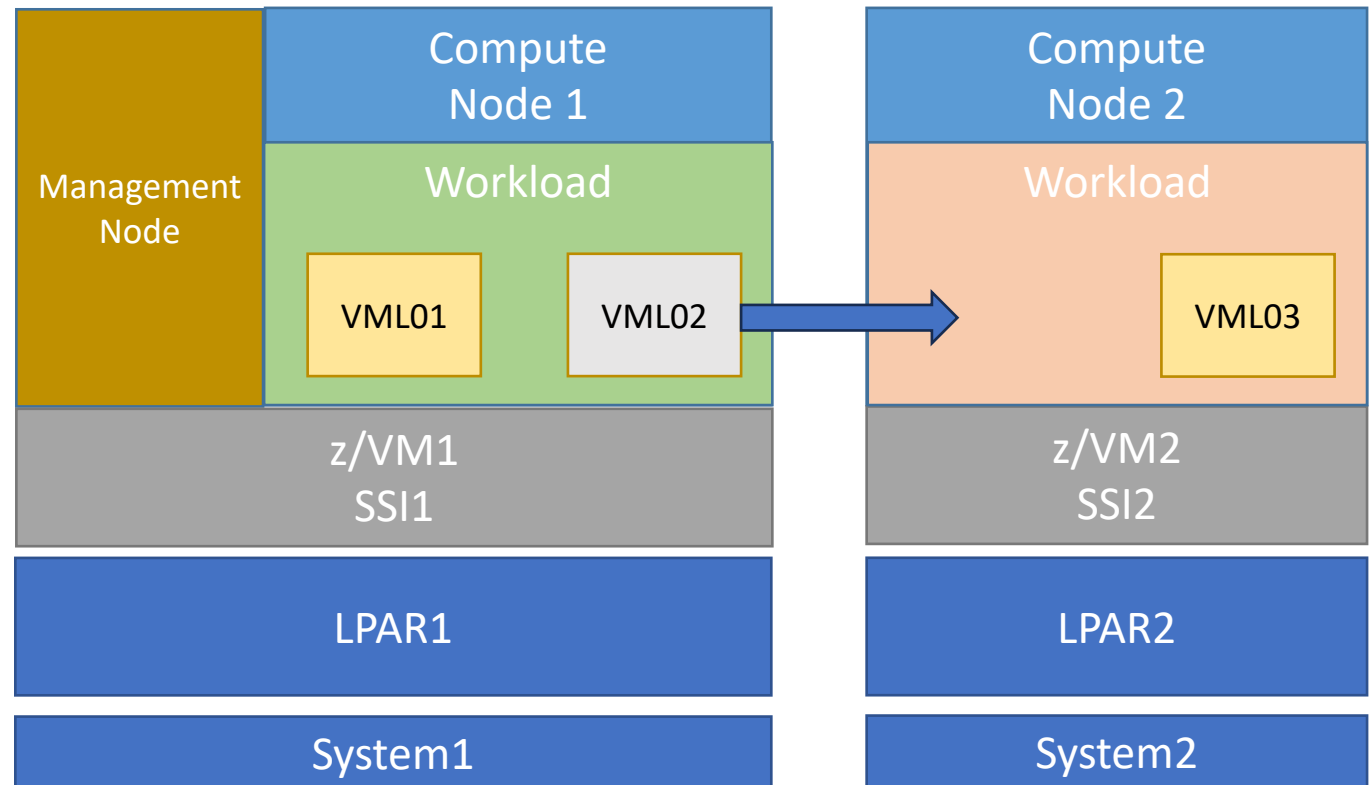
Cold migration of virtual machine to another compute node

# Migration on z/VM



Live Migration on z/VM with SSI

VMRELOCATE command used



# Migration of virtual machines

## Virtual Machines

Refresh Start Stop Suspend Resume Restart Delete Capture Resize **Live Migrate** Cold Migrate Edit Expiration Date  
Unmanage Console Output Console Access Lock Unlock Export CSV Show All Projects' Vms

No filter applied

Name	Host	IP	State	Health	Operating System	Online CPU	VM Type	Hypervisor Type	Instance id	Owner
rhcos-4.15-zvm	ia-compute	9.152.84.107	Active	OK	RHCOS4.15	2	deployed	z/VM	IAI00001	
rhel7.7	ia-compute	9.152.84.102	Active	OK	RHEL7.7 Linux 3.10.0-1062.el7.s390x s390x	1	deployed	z/VM	IAI0001f	
rhel9-tt	ia-compute	9.152.84.109	Active	OK	RHEL9.3 Linux 5.14.0-362.8.1.el9_3.s390x	4	deployed	z/VM	IAI00018	
lynn-rhle88	os006	9.152.84.86	Active	OK	RHEL8.8	4	deployed	KVM	IAI00012	

- KVM: `virsh migrate --live --auto-converge --unsafe <vm id> gemu+ssh://<dest>/system`
- z/VM: `smcli VMRELOCATE -T <vm id> -k <option>`



# Allocation Model

8 LPUs defined to LPAR

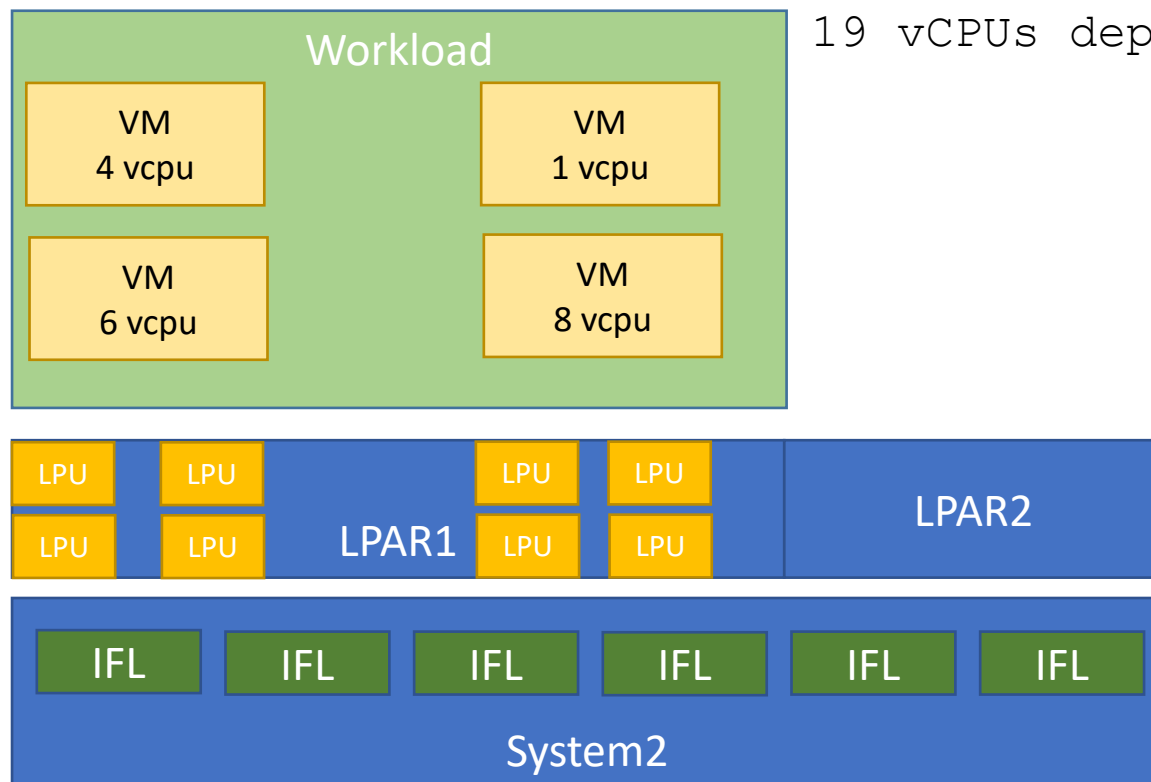
19 vCPUs deployed

IBM Cloud Infrastructure Center utilizes an allocation model for resource provisioning

Overcommit ratios for cpu, memory, disk

Minimum cpu deployment of 1 vcpu

Maximum cpu deployment of total number of LPUs per LPAR





# 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](#)



Thank you!