

IBM Cloud Infrastructure Center installation example on IBM z/VM®

Solution Assurance

Seeman Mannan, Saeed Majidi, David Stark

v1.2, March 2022

IBM Z



Notices and disclaimers

© 2021 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.

U.S. Government Users Restricted Rights – use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.

Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed "as is" without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.

IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.

Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.

Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.

It is the customer's responsibility to ensure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products about this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.**

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.

IBM, the IBM logo, ibm.com and IBM Z, IBM z15, RACF, and z/VM are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml

Table of contents

- Introduction
- IBM® z/VM® conventions
- Software and hardware requirements
- IBM Cloud Infrastructure Center architecture overview
- Installation steps
 1. z/VM login and components
 2. Systems Management API (SMAPI)
 3. Granting authorizations by RACF®
 4. Inter User Communication Vehicle (IUCV) by DirMaint
 5. Virtual Switch (VSWITCH) creation
 6. Disk assignment
 7. Linux installation on the Management and Compute Nodes
 8. Cloud Infrastructure Center installation on the Management Node
 9. Configuring the Compute Node
 10. Deployment of the virtual machines
- References

Introduction

With IBM Cloud Infrastructure Center you can manage the lifecycle of virtual machines that are based on the IBM z/VM and Red Hat® KVM hypervisors. It provides the foundation for scalable Infrastructure-as-a-Service management of traditional and cloud workloads across the enterprise. It helps to abstract the complexities of the underlying hardware and software layers by providing a web-based GUI.

This presentation guides you through an installation example of **Cloud Infrastructure Center** on **z/VM 7.2** running on an IBM z15™ system. It reflects our experience with best practices.

Each step starts with a description of the purpose followed by a checkpoint to ensure the expected results. A diagram shows the successful completion of each step leading to the complete structure of the final solution architecture.

By following this example, you will be able to install more elaborate Cloud Infrastructure Center setups.

Various Linux distributions are supported by Cloud Infrastructure Center for hosting and provisioning. For this example, we chose Red Hat Enterprise Linux (RHEL) 8.4.

z/VM conventions

A 3270 model terminal or an emulating desktop application for this type of terminal is required for working with z/VM. The function keys on the 3270 terminal are used extensively for common operations. In this presentation, we use square brackets for the function key notation. For example, **[CLEAR]** for the clear key or **[PF3]** for the function key 3.

The desktop 3270 applications c3270 or x3270 are available for major operating systems.

z/VM command syntax is not case sensitive.

z/VM commands are often abbreviated. For example, you can use “q” for “**QUARY**” or “l” for “**Login**”, etc.

The output files are sent to a “reader list”, where they can be viewed by pressing the **[PF11]** terminal key or be edited by the “**XEDIT**”.

In z/VM documentation the word “Server” can be understood as a “z/VM login-ID”. For each Linux guest system a z/VM login-ID is issued.

User privilege levels A (highest) through G (lowest) are not supersets of each other.

Software and hardware requirements

The following software and hardware resources were used:

z/VM version 7.2

- DirMaint (z/VM Directory Maintenance)
- SMAPI (z/VM Systems Management Application Programming)
- RACF® (z/VM Resource Access Control Facility)
- ECKD and/or FBA disks for virtual machine provisioning
- 3 z/VM login-IDs

Management Node Linux guest (16 GB memory and 40GB disk)

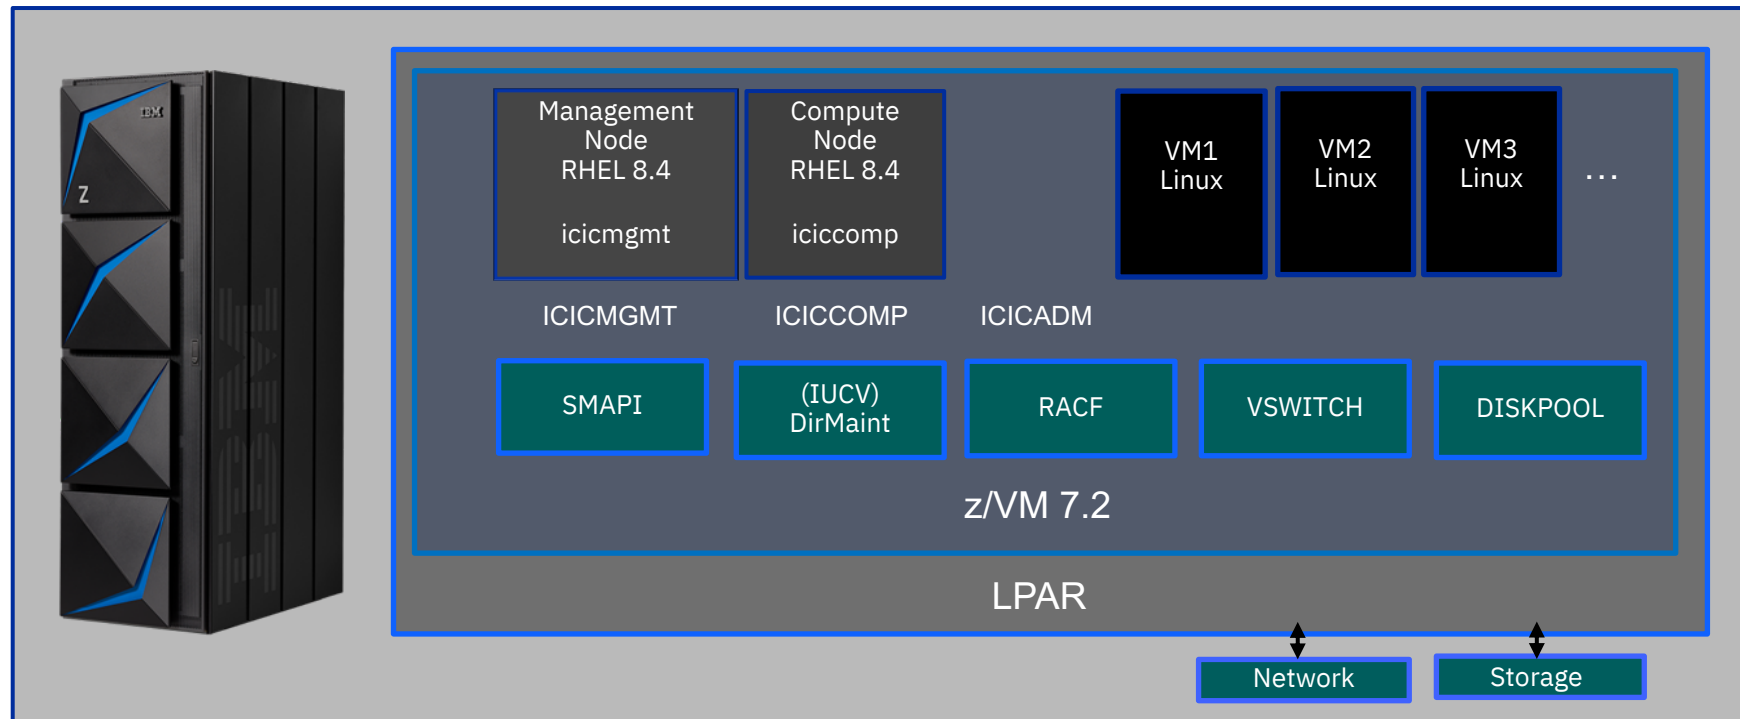
Compute Node Linux guest (8 GB memory and 80 GB disk)

Red Hat Enterprise Linux 8.4 image file for deployment

Reference: ibm.com/docs/en/cic/1.1.4?topic=requirements-hardware-software-zvm-system

Overview of the installed Cloud Infrastructure Center architecture

At completion of the installation, the Cloud Infrastructure Center setup will run on a Logical Partition (LPAR) of a z15 running z/VM. The rest of the presentation describes the significance and the step-by-step build-up of the components.



Step 1: z/VM login IDs and naming conventions

Starting the installation with 3 provided z/VM login-IDs .

For this installation they are:

ICICMGMT (for the Management Node Linux guest)

ICICCOMP (for the Compute Node Linux guest)

ICICADM (Privileged User ID)

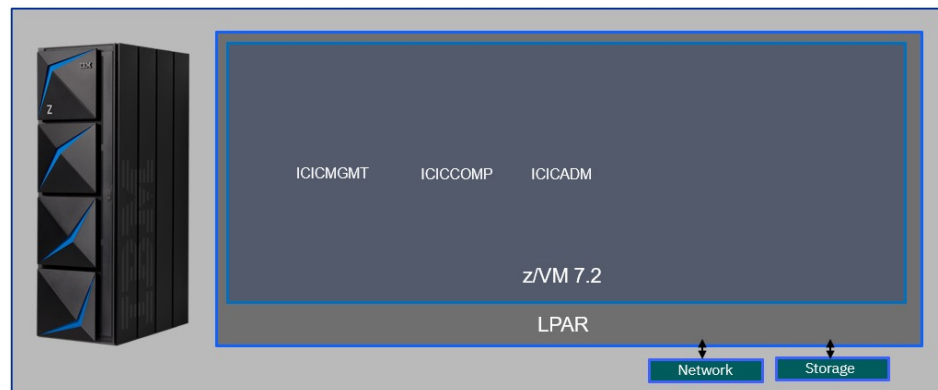
The privileges of the first 2 login-IDs are limited. The 3rd login ID (ICICADM) has z/VM wide privileges required for installation. The privileges are listed on the next page.

Note:

A z/VM login-ID can be freely chosen.

A z/VM login-ID cannot have dashes (-), plus signs (+), colons (:), or underscores (_).

For a best practice, use 8 (or 7) characters.



Step 1: z/VM login IDs and privilege levels

Our first step is to login to z/VM to check the z/VM version and the privileges of the ICICADM login-ID.

Checkpoints:

The output of the “command “**query cplevel**” shows the OS version 7.2, which is supported.

The output of the command “**query priv**” shows the required privilege levels “**ABCDEFGG**”. The user privilege class letter (A,B,...G) indicates the type of user from whom the system accepts commands.

Reference: ibm.com/docs/en/zvm/7.2?topic=classes-defined-privilege

The command “**begin**” changes the prompt of z/VM from “**CP READ**” to “**RUNNING**”, which is the desired status for the remaining steps of our installation.

```
query cplevel
z/VM Version 7 Release 2.0, service level 2102 (64-bit)
```

```
query priv
Privilege classes for user ICICADM
    Currently: ABCDEFG
    Directory: ABCDEFG
```

```
begin
CP READ
RUNNING
```

Step 1: z/VM login IDs and requirements

Various z/VM software components also have IDs. Verify the presence of the packages SMAPI, DirMaint, and RACF. The presence of the highlighted IDs indicate that the required software is already installed.

Checkpoint:

- “**query names**” shows the IDs
- SMAPI (**VSMREQIN**, **VSMEVSRV**, **VSMREQUIU**, **VSMREQI6**, **VSMGUARD**, **VSMWORK1**)
- DirMaint (**DIRMAINT**, **DATAMOVE**)
- RACF (**RACFVM**)
- “**query nss**” indicates SMAPI is enabled

```
query names
INS00024 - DSC , INS00023 - DSC , T8358004 - DSC , T8358002 - DSC
LOHCOST - DSC , AUTOVM - DSC , DATAMOVE - DSC , RSCS - DSC
RSCSDNS - DSC , PERFSVM - DSC , GCS - DSC , FTPSERVE - DSC
TCPIP - DSC , DTCVSW4 - DSC , DTCVSW3 - DSC , DTCVSW2 - DSC
DTCVSW1 - DSC , VMSERV - DSC , VMSERV - DSC , VMSERVU - DSC
VMSERVS - DSC , LGLOPR - DSC , OPERSYMP - DSC , DISKACNT - DSC
EREP - DSC , OPERATOR - DSC , DIRMAINT - DSC , VSMREQUIU - DSC
VSMREQI6 - DSC , VSMREQIN - DSC , VSMEVSRV - DSC , DTCSMAPI - DSC
PERSMAPI - DSC , VSMWORK3 - DSC , VSMWORK2 - DSC , VSMWORK1 - DSC
VSMGUARD - DSC , RACFVM - DSC , T8358001 - DSC , T8358003 - DSC
ICICADM - L0003
VSM - TCPIP
```

```
query nss
OWNERID FILE TYPE CL RECS DATE TIME FILENAME FILETYPE ORIGINID
*NSS 0015 NSS A 0001 06/26 15:47:17 SMAPIOUT DCSS MAINT720
```

Step 2: Systems Management API (SMAPI)

Cloud Infrastructure Center uses the Systems Management API (SMAPI) for managing the running virtual images. z/VM communicates with the Compute Node via SMAPI.

For a SMAPI introduction and quick start guide see: ibm.com/docs/en/zvm/7.2?topic=introduction-smapi-quick-start-guide

Accessing the SMAPI authorization list can be done by making use of a tool called “**SFSULIST**”. This tool is not part of the standard z/VM installation. It can be downloaded from: www.vm.ibm.com/download/packages

Step 2: Systems Management API (SMAPI)

The Compute Node ID must be included on the list of SMAPI users. The steps are:

1. “**sfsulist vmsys**” to list the SMAPI files
2. Press **[PF11]** on VSMWORK1
3. Press **[PF11]** on VMSYS:VSMWORK1
4. Mark “**x**” VSMWORK1 AUTHLIST to XEDIT
5. Insert a line for “**ICICCOMP**”
6. “**file**” to save and exit XEDIT

```
sfsulist vmsys
```

```
VSMWORK1 2 6,000 5,999-99% YES 1x
```

```
ICICADM DIRLIST A0 V 319 Trunc=319
Cmd Fm Directory Name
- VMSYS:VSMWORK1.
V VMSYS:VSMWORK1.DATA
W VMSYS:VSMWORK1.STATUS
```

```
ICICADM FILELIST A0 V 169 Trunc=169 Size=5 Line=1 Col=1 Alt=0
Directory = VMSYS:VSMWORK1.
Cmd Filename Filetype Fm Format Lrecl Records Blocks Date Time
VSMWORK1 NAMELIST Z1 F 80 244 5 9/28/21 16:53:58
x VSMWORK1 AUTHLIST Z1 F 195 12 1 9/16/21 15:02:26
```

```
VSMWORK1 AUTHLIST Z1 F 195 Trunc=195 Size=12 Line=10 Col=1 Alt=1

===== * * * Top of File * * *
===== DO.NOT.REMOVE DO.NOT.
MOVE
===== MAINT ALL
===== ICICCOMP ALL
```

```
====> file
```

Step 3: Granting authorizations by RACF

The Compute Node requires authorization to link the available z/VM minidisks for provisioning virtual machines.

1. The command "**RAC ALU ICICCOMP OPERATIONS**" alters user (ALU) "**operations**" for the user ICICCOMP (Compute Node) to allow the link to minidisks.
2. The command "**RAC PERMIT ICICCOMP CLASS(VMRDR) ID(VSMWORK1) ACCESS(UPDATE)**" enables reader access to VSMWORK1 (SMAPI) for the user ICICCOMP (Compute Node).
3. After granting the required authorizations, restart SMAPI. "**FORCE VSMGUARD**" stops SMAPI and "**XAUTOLOG VSMGUARD**" starts it again.

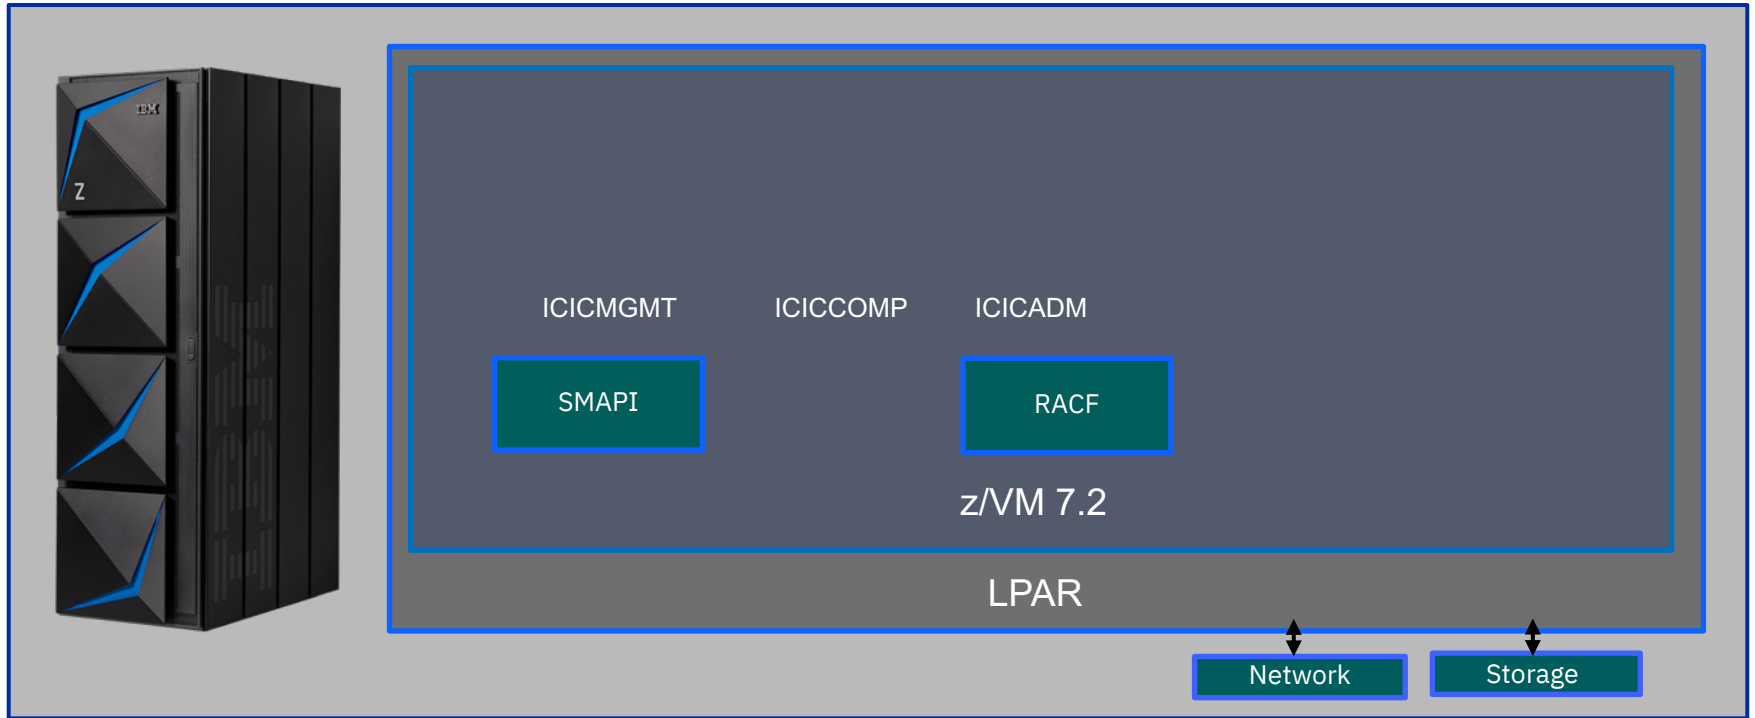
```
RAC ALU ICICCOMP OPERATIONS
```

```
RAC PERMIT ICICCOMP CLASS(VMRDR) ID(VSMWORK1) ACCESS(UPDATE)
```

```
FORCE VSMGUARD  
XAUTOLOG VSMGUARD
```

Overview of the architecture

After granting authorizations by RACF



Step 4: Inter User Communication Vehicle (IUCV) by DirMaint

IUCV is required for communication between z/VM and the Compute Node (ICICCOMP) running as its Linux guest. DirMaint is used to set and check IUCV in a user's login-ID profile.

Checkpoint:

1. “**dirm for ICICCOMP review**” sends user's profile to the reader list
2. “**rl**” to invoke the reader list
3. Press [Clear] to list the reader list files
4. Place cursor on “**ICICCOMP DIRECT**” line
5. Press [PF11] to peek into the file
6. The line “**IUCV ANY**” indicates the correct setting

```
dirm for ICICCOMP review

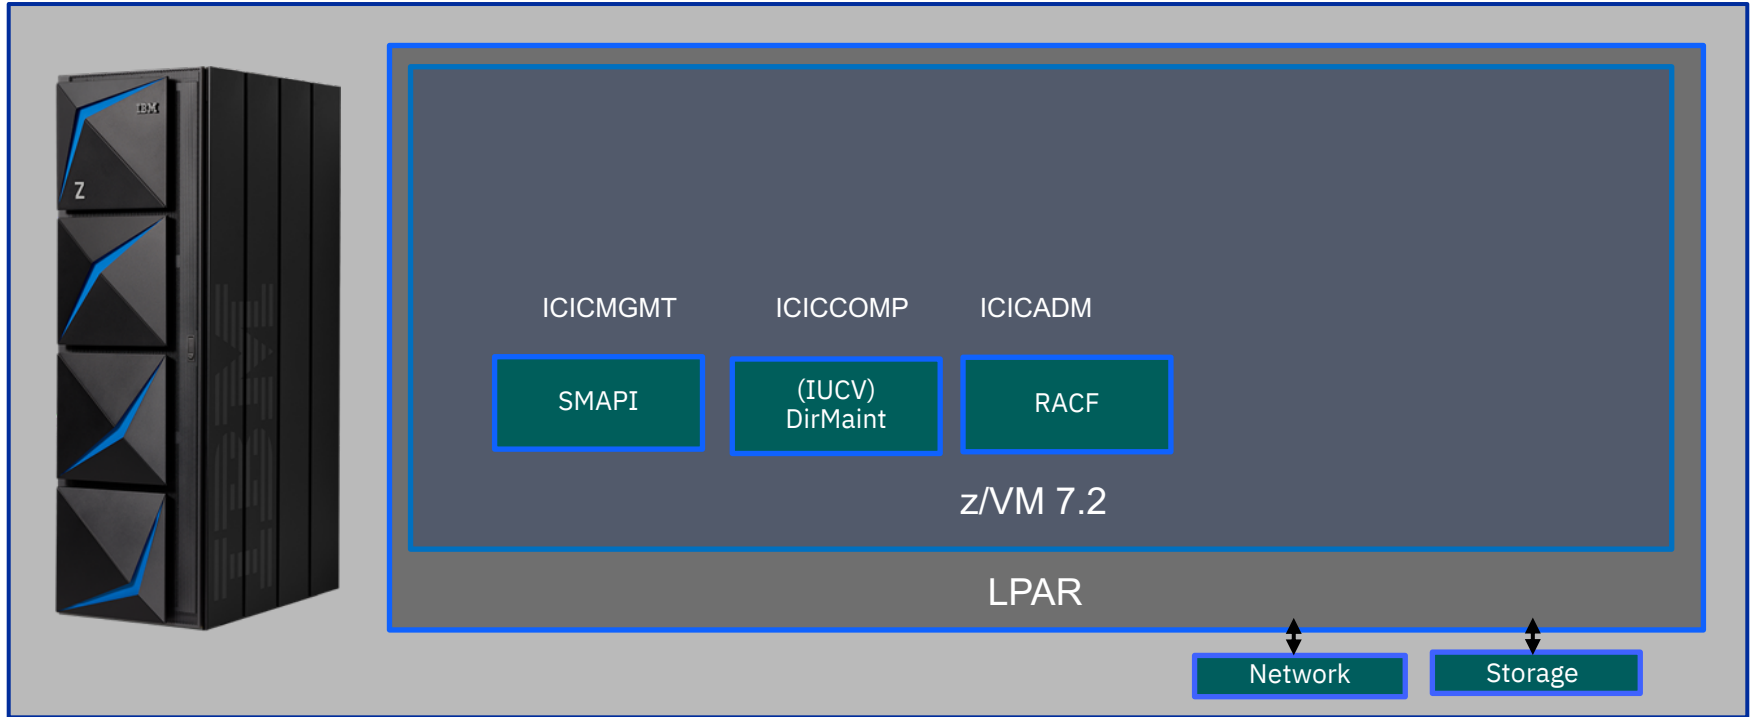
rl

* ICICCOMP DIRECT  PUN A DIRMAINT

===== USER ICICCOMP XXXXXXXX
===== DVHRXV3355I The following records are included from profile: LINIFL
===== PROFILE LINIFL
===== CLASS GU
===== STORAGE 16G
===== MAXSTORAGE 2T
===== COMMAND SET VCONFIG MODE LINUX
===== * COMMAND DEFINE CPU 0-1 TYPE IFL
===== CPU 00
===== CPU 01
===== IPL CMS
===== IUCV ALLOW
===== IUCV ANY
```

Overview of the architecture

After IUCV configuration by DirMaint



Step 5: z/VM Virtual Switch (VSWITCH) creation

In z/VM a VSWITCH (Virtual Switch) is used for networking of its guest systems. The following steps create a VSWITCH:

- A VSWITCH requires 3 free OSA (network) devices
- The range of “bd00-bd0f” OSA devices was allocated to our installation
- The command “**q bd00-bd0f**” shows the status of OSA devices
- We note that the OSA devices “bd09-bd0b” are free
- “**define vswitch ICICVS rdev BD09 ETH Contr ***” creates the VSWITCH named “ICICVS”

```
q bd00-bd0f
OSA BD00 ATTACHED TO TCP/IP    BD00 DEVTYPE OSA      CHPID 99 OSD
OSA BD01 ATTACHED TO TCP/IP    BD01 DEVTYPE OSA      CHPID 99 OSD
OSA BD02 ATTACHED TO TCP/IP    BD02 DEVTYPE OSA      CHPID 99 OSD
OSA BD03 ATTACHED TO DTCVSW1   0600 DEVTYPE OSA      CHPID 99 OSD
OSA BD04 ATTACHED TO DTCVSW1   0601 DEVTYPE OSA      CHPID 99 OSD
OSA BD05 ATTACHED TO DTCVSW1   0602 DEVTYPE OSA      CHPID 99 OSD
OSA BD06 ATTACHED TO DTCVSW3   0606 DEVTYPE OSA      CHPID 99 OSD
OSA BD07 ATTACHED TO DTCVSW3   0607 DEVTYPE OSA      CHPID 99 OSD
OSA BD08 ATTACHED TO DTCVSW3   0608 DEVTYPE OSA      CHPID 99 OSD
OSA BD09 FREE , OSA BD0A FREE , OSA BD0B FREE] , OSA BD0C FREE
OSA BD0D FREE , OSA BD0E FREE , OSA BD0F FREE
```

```
define vswitch ICICVS rdev BD09 ETH Contr *
VSWITCH SYSTEM ICICVS is created
Ready; T=0.01/0.01 10:26:58
HCPSWU2830I VSWITCH SYSTEM ICICVS status is ready.
HCPSWU2830I DTCVSW4 is VSWITCH controller for device BD09.P00.
```

Step 5: z/VM Virtual Switch (VSWITCH) creation

Steps to verify VSWITCH “**ICICVS**” property.

Checkpoint:

- “**q vswitch ICICVS**” shows the details of the VSWITCH ICICVS.
- The attribute “**ETHERNET**” indicates a LEVEL 2 VSWITCH (default) which is a requirement.
- “**VLAN Unaware**” indicates a flat non-segmented network (default) for our installation.
- “**query osa**” shows that the OSA devices “bd09-bd0b” are no longer free as they are now used by the vSwitch controller of “ICICVS” VSWITCH.

```
q vswitch ICICVS
VSWITCH SYSTEM ICICVS      Type: QDIO      Connected: 0      Maxconn: INFINITE
PERSISTENT RESTRICTED     ETHERNET      Accounting: OFF
USERBASED LOCAL
VLAN Unaware
MAC address: 02-28-3A-00-00-85      MAC Protection: Unspecified
IPTimeout: 5      QueueStorage: 8
Isolation Status: OFF      VEPA Status: OFF
Uplink Port:
State: Ready      PriQueuing: OFF
PMTUD setting: EXTERNAL      PMTUD value: 9000      Trace Pages: 8
RDEV: BD09.P00 VDEV: 0600 Controller: DTCVSW4 ACTIVE
Adapter ID: 856100022E48.01B0
```

```
query osa
OSA BD00 ATTACHED TO TCPIP      BD00 DEVTYPE OSA      CHPID 99 OSD
OSA BD01 ATTACHED TO TCPIP      BD01 DEVTYPE OSA      CHPID 99 OSD
OSA BD02 ATTACHED TO TCPIP      BD02 DEVTYPE OSA      CHPID 99 OSD
OSA BD03 ATTACHED TO DTCVSW1    0600 DEVTYPE OSA      CHPID 99 OSD
OSA BD04 ATTACHED TO DTCVSW1    0601 DEVTYPE OSA      CHPID 99 OSD
OSA BD05 ATTACHED TO DTCVSW1    0602 DEVTYPE OSA      CHPID 99 OSD
OSA BD06 ATTACHED TO DTCVSW3    0606 DEVTYPE OSA      CHPID 99 OSD
OSA BD07 ATTACHED TO DTCVSW3    0607 DEVTYPE OSA      CHPID 99 OSD
OSA BD08 ATTACHED TO DTCVSW3    0608 DEVTYPE OSA      CHPID 99 OSD
OSA BD09 ATTACHED TO DTCVSW4    0600 DEVTYPE OSA      CHPID 99 OSD
OSA BD0A ATTACHED TO DTCVSW4    0601 DEVTYPE OSA      CHPID 99 OSD
OSA BD0B ATTACHED TO DTCVSW4    0602 DEVTYPE OSA      CHPID 99 OSD
```

Step 5: z/VM Virtual Switch (VSWITCH) creation

The VSWITCH definition needs to be made persistent to be permanently available even after an IPL (reboot).

1. “**link pmaint cf0 cf0 wr**” links to disks cf0 for write modification
2. “**acc cf0 n**” access cf0 with tag “N”
3. “**filel * * n**” to list all files tagged with “N”
4. Hit **[CLEAR]** from keypad to show the list
5. Mark the “**SYSTEM CONFIG**” line with “**x**” and hit **[RETURN]** to edit the file by XEDIT
6. Enter the VSWITCH “ICICVS” definition into the SYSTEM CONFIG file as shown
7. At XEDIT prompt, enter “**file**” to save the file

```
link pmaint cf0 cf0 wr
acc cf0 n
filel * * n
```

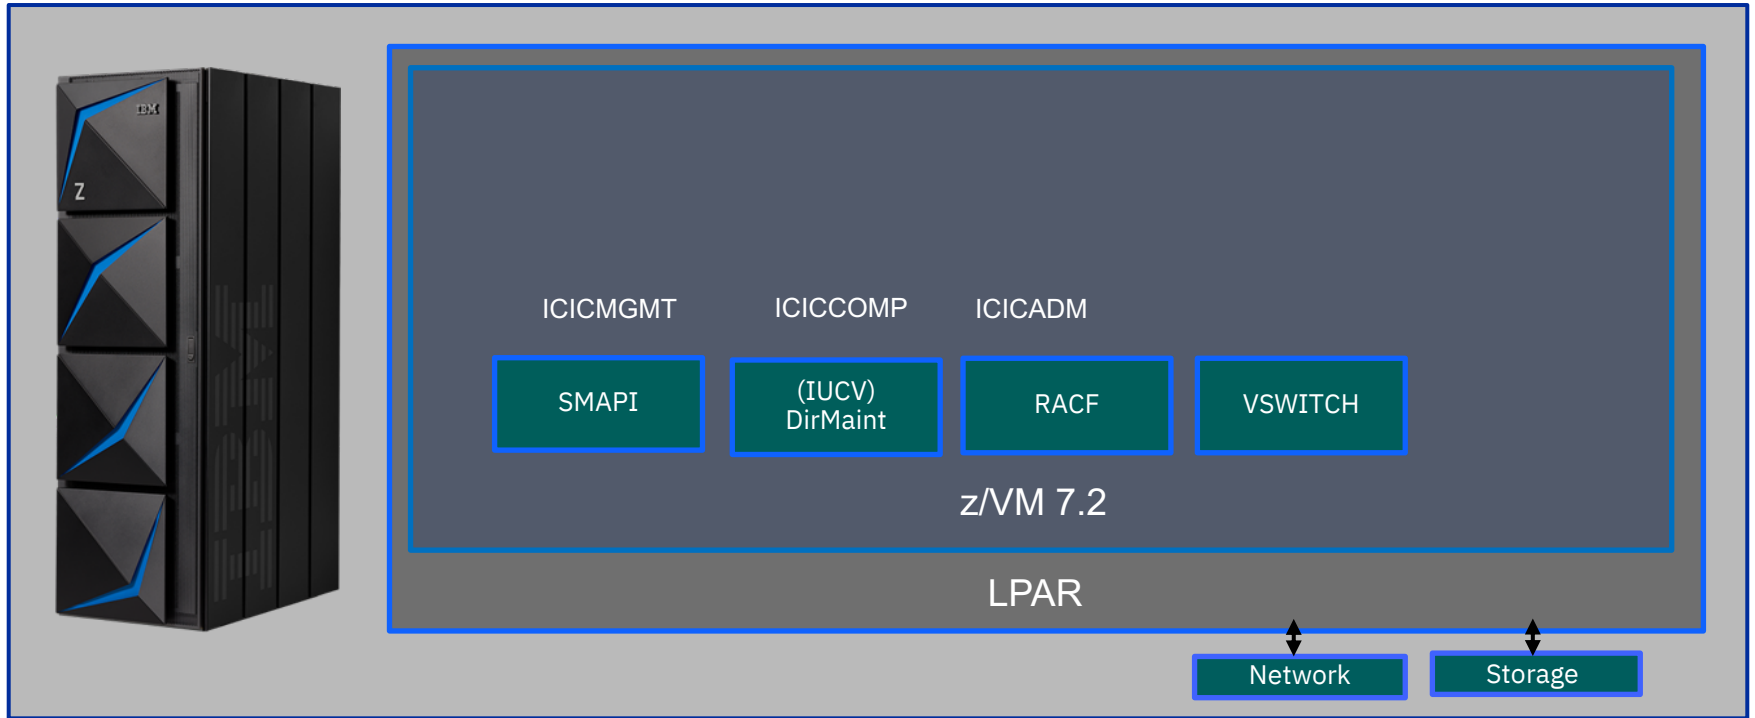
ICICADM	FILELIST	A0	V	169	Trunc=169	Size=4	Line=1	Col=1	Alt=0
Cmd	Filename	Filetype	Fm	Format	Lrecl	Records	Blocks	Date	Time
x	SYSTEM	CONFIG	N1	F	80	363	8	9/01/21	14:41:16
	SYSTEM	C0201013	N1	F	80	344	7	10/13/20	10:43:45
	SYSTEM	C0200930	N1	F	80	339	7	9/30/20	13:38:36
	LOGO	CONFIG	N1	V	69	63	1	9/01/20	13:14:32

```
=====
/*****
/*                               Virtual Networking                               */
/*****
=====
VMlan Macprefix 02283A
Define Vswitch VSW58G RDev BD03 Eth Contr *
Define Vswitch VSICICA RDev BD06 Eth Contr * VLAN AWARE NATIVE NONE
Define Vswitch ICICVS RDev BD09 Eth Contr *
```

```
====> file
```

Overview of the architecture

After VSWITCH creation



Step 6: z/VM disk setup

The disks used for virtual machine provisioning must be aggregated into a z/VM DISKPOOL. Then the free disks need to be formatted and labeled before they are used. The tool used to format the disks (“cpfmtxa”) is available on the system disk “551”. see next page:

Checkpoint:

- “**query 67E8-67EF**” checks the status of the disks
- DASD disks “67ED” and “67EE” are free
- “**link pmaint 551 551 rr**” links the disk 551 in read only mode for accessing the formatting tool (“cpfmtxa”)
- “**acc 551 n**” access the linked disk (tagged as n)

Note:

- Disks 67E8-67EF are assigned to our installation.
- DASD or FCP disks may not be “dedicated” to a Linux guest (by CP ATTACH or DEDICATED user directory statement). Otherwise, at the preparation of the provisioning, z/VM will fail to “link” a “dedicated” disk.
- In the next steps we first format one of the free disks “67ED” and then enter it into a DISKPOOL for provisioning.

```
query 67E8-67EF
DASD 67E8 CP SYSTEM T35801 2
DASD 67E9 CP SYSTEM T35802 1
DASD 67EA CP SYSTEM T35803 0
DASD 67EB FR67EB , DASD 67EC FR67EC , DASD 67ED FREE , DASD 67EE FREE
DASD 67EF 0X67EF
Ready; T=0.01/0.01 10:47:22
```

```
link pmaint 551 551 rr
Ready; T=0.01/0.01 13:28:50
acc 551 n
DMSACC724I 551 replaces N (CF0)
DMSACP723I N (551) R/O
Ready; T=0.01/0.01 13:28:54
```

Step 6: z/VM disk assignment

The formatting tool “cpfmtxa” is now available for formatting our free disks.

1. “**attach 67ED ***” attaches the free disk “67ED” to the current user for formatting
2. Use the “**cpfmtxa**” command, as shown, to format cylinder 0 of the disk “67ED” and label it as “T8358D” (a label of our choice)
3. The cylinder notation 0-0 stands for the range from cylinder 0 to cylinder 0, which results in the single cylinder 0. Formatting cylinder 0 alone is sufficient
4. “**detach 67ED**” once the formatting disk completes

```
attach 67ED *  
DASD 67ED ATTACHED TO ICICADM 67ED WITH DEVCTL HYPERPAV BASE  
Ready; T=0.01/0.01 13:38:54
```

```
cpfmtxa  
ENTER FORMAT, ALLOCATE, LABEL, OWNER, ERASE OR QUIT:  
format  
ENTER THE VDEV TO BE PROCESSED OR QUIT:  
67ED  
ENTER THE CYLINDER RANGE TO BE FORMATTED ON DISK 67ED OR QUIT:  
0-0  
ENTER THE VOLUME LABEL FOR DISK 67ED:  
T8358D  
CPFMTXA:  
FORMAT WILL ERASE CYLINDERS 000000000-000000000 ON DISK 67ED  
DO YOU WANT TO CONTINUE? (YES | NO)  
YES
```

```
detach 67ED  
DASD 67ED DETACHED  
Ready; T=0.01/0.01 16:18:41
```

Step 6: z/VM disk assignment

Step to verify the newly created label “**T8358D**” for the disk “**67ED**”.

Checkpoint:

- “**query 67ED**” shows that the label “T8358D” is placed on the disk.
- The label “**T8358D**” is just a name of our choice. You may choose any label name following your own naming convention.

```
Query 67ED
DASD 67ED T8358D
Ready; T=0.01/0.01 16:19:01
```

In the next step, we will create a DISKPOOL and define our disk “67ED” in that pool.

Step 6: z/VM disk assignment

Cloud Infrastructure Center uses a DISKPOOL for provisioning DASD and FCP disks. The DISKPOOLS consist of the individual labeled disks. They are defined in the “EXTENT CONTROL” file, as shown in the following steps:

1. “**dirm send extent control**” sends the “EXTENT CONTROL” file to the reader list
2. “**rl**” to list the reader list
3. “**receive / (replace)**” receives the file for modifications
4. Press **[PF3]** to return to the prompt line
5. “**filel**” to list the files
6. Mark “**x**” to enter “XEDIT” for editing the file

On the next page, we edit the “EXTENT CONTROL” file

```
dirm send extent control
DVHXTM1191I Your SEND request has been sent for processing to DIRMAINT
DVHXTM1191I at BOET8358.
Ready; T=0.01/0.01 16:22:01
DVHREQ2288I Your SEND request for ICICADM at * has been accepted.
RDR FILE 0058 SENT FROM DIRMAINT PUN WAS 1162 RECS 0044 CPY 001 A NOHOLD NOKEEP
DVHREQ2289I Your SEND request for ICICADM at * has completed; with RC
DVHREQ2289I = 0.
```

```
rl
```

```
receive / (replace ROL PUN A DIRMAINT BOET8358 NONE 44 10/25 16:22:01
```

```
filel
```

```
EXTENT CONTROL received from DIRMAINT at BOET8358
x EXTENT CONTROL A1 V 72 90 1 10/19/21 9:49:54
```


Step 6: z/VM disk assignment

Edit the “EXTENT CONTROL” file as follows:

====	*RegionId	VolSer	RegStart	RegEnd	Dev-Type	Comments
====	T8358D	T8358D	0001	30050	3390-27	

- Add a line for the newly labeled disk “67ED” (labeled T9358D) under the “*RegionId” line
- For “RegEnd” enter number of cylinders used for formatting minus 1. Our disk “67ED” has 30051 cylinders. Therefore, we have entered 30050 for the “RegEnd” field.
- Dev is “3390” for DASD disks and “FBA” for FCP disks. As our disk “67ED” is a DASD disk, its Dev-Type is entered as “3390”. The default settings for the DEV-Types are included in the file DEFAULTS DATADVH as part of the standard z/VM installation. (If any additional Dev-Types are required, they can be added in “:DEFAULTS.” section in the EXTENT CONTROL file). In our environment, this DASD has the DEV-Type 3390-27.

Note :

- A Cylinder is 670 KB. For larger provisioning environments, provide EAV (Extended Address Volumes) with a large size (e.g., DEV-Type 3390-A with up to the limit of currently supported volumes of 1182006 cylinders) in the DISKPOOL.

Reference: [ibm.com/servers/resource/ibm/svc0302a.nsf/pages/zVMV7R2sc246283/\\$file/hcpk3_v7r2.pdf](https://ibm.com/servers/resource/ibm/svc0302a.nsf/pages/zVMV7R2sc246283/$file/hcpk3_v7r2.pdf)

Step 6: z/VM disk assignment

The DISKPOOL is also defined in the "EXTENT CONTROL" file, as follows:

- Add a line for the DISKPOOL (ICICPLE2) containing the above disk label (T8358D) under the ":GROUPS." section
- "**file**" to save and exit "xedit"
- "**dirm file** / " on the reader list, to save the file back into the system area

```
==== :GROUPS.  
==== ICICPLE2 T8358D
```

```
====> file
```

```
dirm file / | CONTROL A1 V 72 92 1 10/25/21 16:56:41
```

After entering the DISKPOOL definition into the "EXTENT CONTROL" file, it has to be reloaded for activation:

- "**dirm rlde**" reloads the "EXTENT CONTROL" file

```
dirm rlde  
DVHXTM1191I Your RLDEXTN request has been sent for processing to  
DVHXTM1191I DIRMAINT at BOET8358.  
Ready; T=0.01/0.01 17:42:10  
DVHREQ2288I Your RLDEXTN request for ICICADM at * has been accepted.  
DVHILZ3510I Starting DVHINITL with directory: USER DIRECT E  
DVHILZ3510I DVHINITL Parms: BLDMONO BLDDASD BLDLINK  
DVHIZD3528W One or more DASD volume control files ($$$$) were  
DVHIZD3528W created using default values for device characteristics -  
DVHIZD3528W DATAMOVE 05F0  
DVHREQ2289I Your RLDEXTN request for ICICADM at * has completed; with RC  
DVHREQ2289I = 0.
```

Step 6: z/VM disk assignment - verification

We can now check the definition of our DISKPOOL (ICICPLE2) as follows:

- “**dirm free g= ICICPLE2**” sends the DISKPOOL information (ICICPLE2 FREEXT) to the reader list
- “**rl**” to list the files on the reader list
- Press [**PF11**] on the file line to peek into the file
- Our DISKPOOL (ICICPLE2) is confirmed with 30050 free cylinders

```
dirm free g= ICICPLE2
DVHXTM1191I Your FREEXT request has been sent for processing to DIRMAINT
DVHXTM1191I at BOET8358.
Ready: T=0.01/0.01 14:38:45
DVHREQ2288I Your FREEXT request for ICICADM at * has been accepted.
RDR FILE 0063 SENT FROM DIRMAINT PUN WAS 1169 RECS 0007 CPY 001 A NOHOLD NOKEEP
DVHREQ2289I Your FREEXT request for ICICADM at * has completed; with
DVHREQ2289I RC = 0.
```

rl

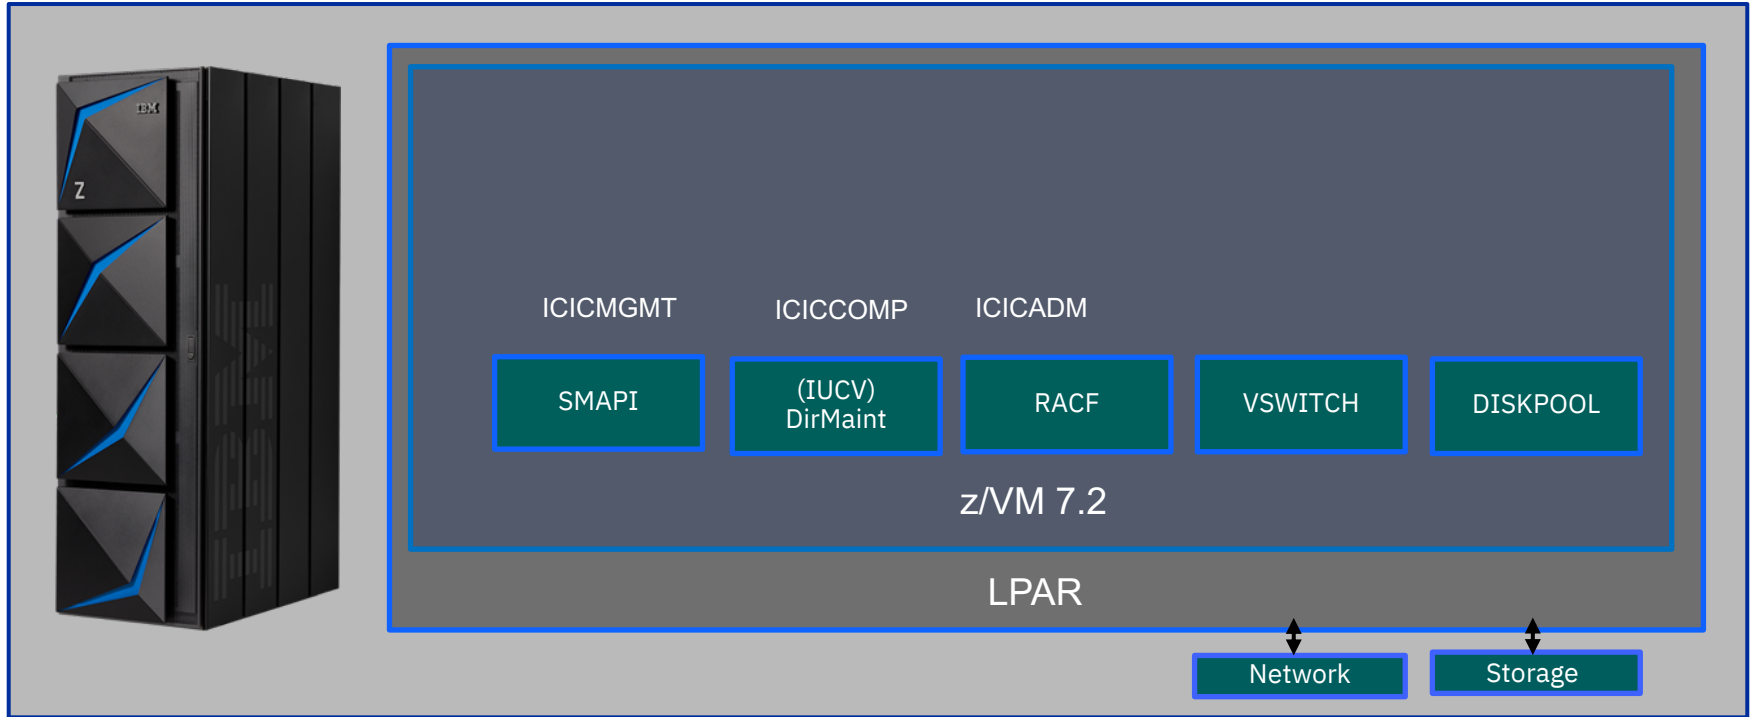
ICICPLE2	FREEXT	PUN A	DIRMAINT	BOET8358	NONE	7	10/25	18:10:28
----------	--------	-------	----------	----------	------	---	-------	----------

```
0061    PEEK    A0  V 80  Trunc=80 Size=6 Line=0 Col=1 Alt=0
File ICICPLE2 FREEXT from DIRMAINT at BOET8358 Format is NETDATA.
* * * Top of File * * *
FREEXT G= ICICPLE2

  GROUP  REGION  VOLUME  START    SIZE    (END)  OWNER  ADDR
-----
ICICPLE2 T8358D  T8358D    0001    30050    30050  FREE.  0000
* * * End of File * * *
```

Overview of the architecture

After DISKPOOL configuration



Step 7: Linux installation on z/VM guest systems

The two Linux guests (Management Node and Compute Node) need to be installed with the supported version of Linux listed in the most current version of the documentation (see the page “[Hardware and software requirements for z/VM® system](#)”).

The installation process for the Linux guests is site specific. Consult your system administrator for instructions.

Safely keep the root passwords (according to your regulations) and the IP addresses of these two systems.

Step 8: Cloud Infrastructure Center installation on the Management Node

The “BaseOS” and “AppStream” repositories are the required repository for the installation.

Checkpoint:

- “**dnf repolist**” lists the repositories
- Both repositories are available on the system
- IBM Cloud Infrastructure Centre relies on internal web-based communication. Our installation was successful with these default values. They can be changed according to your local network definitions.

```
[icicmgmt]# dnf repolist
repo id                repo name
rhel8.4-bistro_appstream  Automatically added via cronjob on 2021-05-19 18:01:00
AppStream
rhel8.4-bistro_baseos    Automatically added via cronjob on 2021-05-19 18:01:00
BaseOS
```

```
[icicmgmt]# export no_proxy="localhost,127.0.0.1,.example.com"
```

Step 8: Cloud Infrastructure Center installation on the Management Node

At this point, we create a configuration file named “**config.properties**” in the directory “/etc/ctic/” with the content as shown in the example.

- compute_instance_template: is used to name the provisioned virtual machines (e.g., ins0002f, ins00023, etc.)
- compute_user_profile: is the z/VM file name (see the page “VM Default Profile”)
- default_admin_userid: is our privileged z/VM login ID (see the page “z/VM login IDs”)
- Find the LAN interface name of the system in the output of the “**ifconfig**” command
- Set the environment variable HOST_INTERFACE to the LAN interface name of your system

```
[ctic configs]
compute_instance_template=ins%05x
compute_user_profile=ICICDFLT
default_admin_userid=ICICADM
```

```
[icicmgmt]# ifconfig
encbdf0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
```

```
[icicmgmt]# export HOST_INTERFACE= encbdf0
```

Reference: ibm.com/docs/en/cic/1.1.4?topic=zvm-installing-cloud-infrastructure-center

Step 8: Cloud Infrastructure Center installation on the Management Node

- Verify that the environment variable LANG in the /etc/locale.conf is set to “en_US.UTF-8”

```
[icicmgmt]# grep LANG /etc/locale.conf  
LANG="en_US.UTF-8"
```

Download the Cloud Infrastructure Center software package (IBM_Cloud_Infrastructure_Center<version>.tar.gz) into a directory of your choice (eg: Downloads) and unpack the software by the following commands:

- “**tar -xzf IBM_Cloud_Infrastructure_Center_1.1.4.tar.gz**”
extracts 4 files
- “**tar -xzf icic-install-s390x-rhel-1.1.4.0.tgz**”
extracts the ICIC software into the directory
“Downloads/icic-1.1.4.0/”

```
[icicmgmt]# tar -xzf IBM_Cloud_Infrastructure_Center_1.1.4.tar.gz
```

```
icic-install-s390x-rhel-1.1.4.0.tgz  
icic-install-1.1.4.0.tgz.sig  
icicpublickey  
readme.txt
```

```
[icicmgmt]# tar -xzf icic-install-s390x-rhel-1.1.4.0.tgz
```


Step 8: Cloud Infrastructure Center Installation on the Management Node

- Before initiating the Cloud Infrastructure Center installation, change to the working directory “icic-1.1.4.0”, and then execute the “install” command as shown.
- The option “z” specifies the z/VM hypervisor.
- The option “c” automatically creates the necessary firewall rules. (No need for this option, if firewall rules are already configured manually).

Checkpoint:

- After installation, run the command “**opt/ibm/icic/bin/icic-services status**” to check the status of the services.
- Sample output would look like:

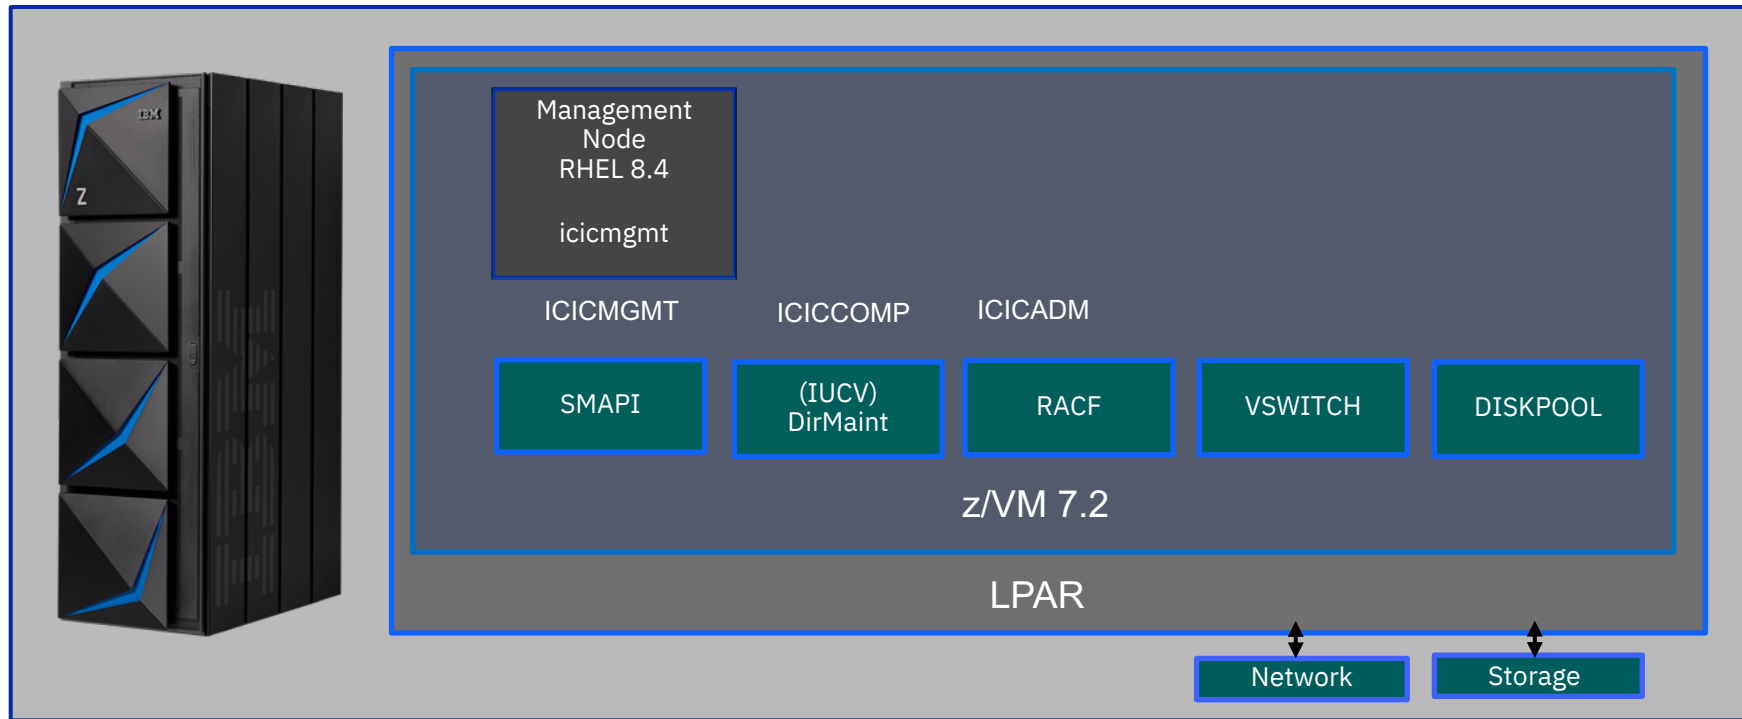
```
[icicmgmt]# cd icic-1.1.4.0  
[icicmgmt]# ./install -z -c
```

```
[icicmgmt]# /opt/ibm/icic/bin/icic-services status
```

```
• icic-bumblebee.service - IBM Cloud Infrastructure Center Bumblebee  
  Active: active (running) since Mon 2021-09-27 15:35:29 CEST; 1 months 1 days ago  
• ivp-validation-api.service - IVP API Server  
  Active: active (running) since Mon 2021-09-27 15:35:21 CEST; 1 months 1 days ago  
• clerk-api.service - clerk API Server  
  Active: active (running) since Mon 2021-09-27 15:35:21 CEST; 1 months 1 days ago
```

Overview of the architecture


After the software installation on the Management Node



Step 9: Compute Node configuration


After successful installation of the Cloud Infrastructure Center software, its GUI interface is available by a web server on the Management Node.

- You can access the GUI by entering the Cloud Infrastructure Center web address on a web browser of your desktop:
<https://icicmgmt/icic/login.html>
- Use the “root” login credentials of the Management Node to login to Cloud Infrastructure Center.
- Sample home screen:
- The next step is to configure the first Host (e.g., the Compute Node).

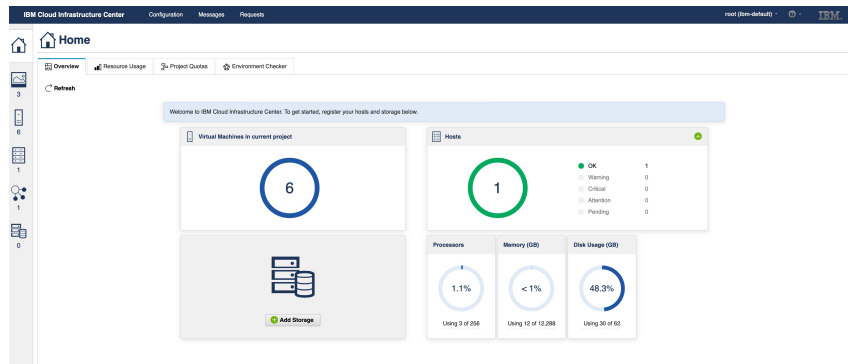


Login

* User name:

* Password:
 

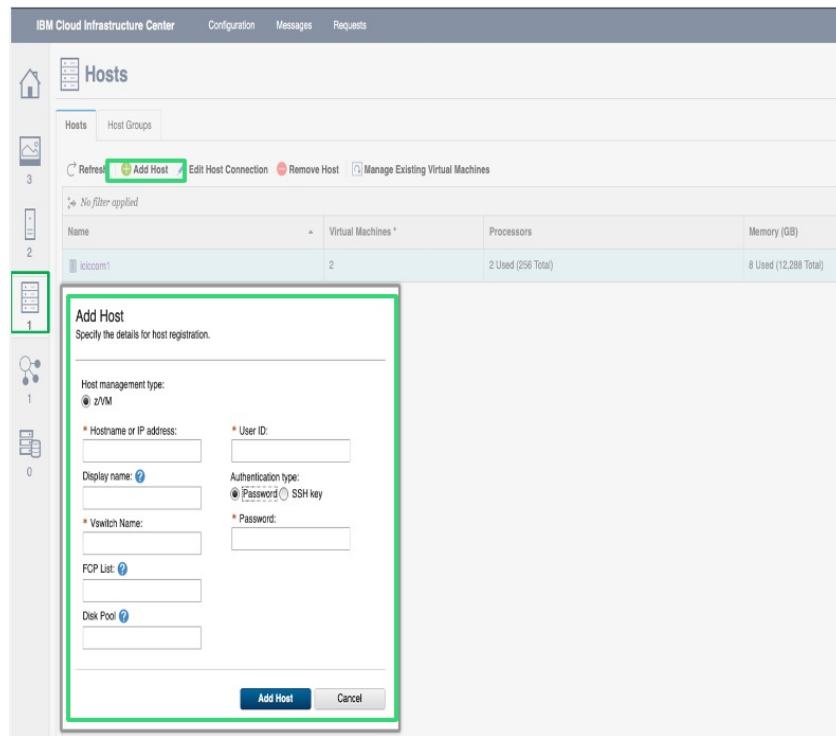
Licensed Materials - Property of IBM Corp, IBM Corporation and other(s) 2021. IBM is a registered trademark of IBM Corporation, in the United States, other countries, or both.



Step 9: Compute Node configuration

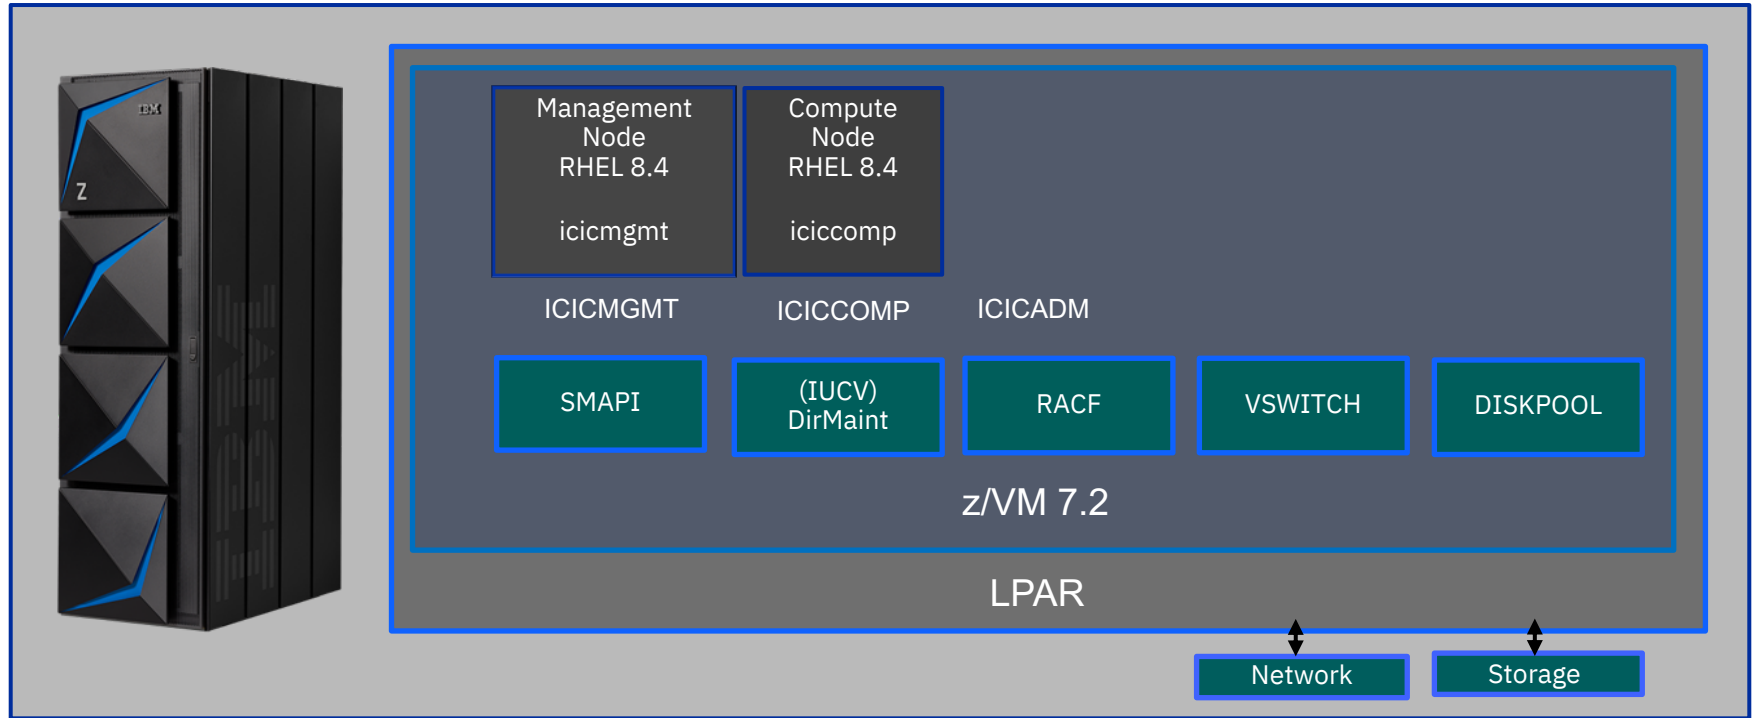
Follow the following steps to add the Compute Node:

- Select “**Add Host**” button under “**Hosts**” tab
- Provide the “**Hostname or IP: iciccomp**”
- Provide “**Display name**” (optional)
- Provide “**Vswitch: ICICVS**”
- Provide “**FCP List**” (not used for our installation)
- Provide “**Disk Pool: ECKD:ICICPLE2**”
- Provide “**User ID: root**”
- Select “**Authentication type: Password**”
- Provide “**Password**: the root password of the Compute Node”
- Click on “**Add Host**”



Overview of the architecture

After the Compute Node configuration



Step 10: Deployment of the virtual machines

At this stage Cloud Infrastructure Center is installed.

The following one-time preparation steps are required for provisioning virtual machines:

1. A default virtual machine system profile needs to be created on z/VM
2. Networking parameters need to be configured
3. Also a master Linux image is required for provisioning. The instructions to create an image from a running system or an ISO file can be found in [IBM docs](#)

Note: The size of the disk where an image is created must be smaller than the size of the disks in the DISKPOOL for provisioning.

1. In the Compute Node account, create a file named “PROFILE ICICDFLT” with the shown literal content

```
===== * * * Top of File * * *
|...+...1...+...2...+...3...+...4...+...5...+...6...+...7...
===== PROFILE ICICDFLT
===== COMMAND SET RUN ON
===== COMMAND SP CONS * START
===== MACHINE ESA
===== OPTION APPL CHPIDV ONE
===== CONSOLE 0009 3215 T
===== SPOOL 000C 2540 READER *
===== SPOOL 000D 2540 PUNCH A
===== SPOOL 000E 1403 A
===== * * * End of File * * *
```

Step 10: Deployment of the virtual machines

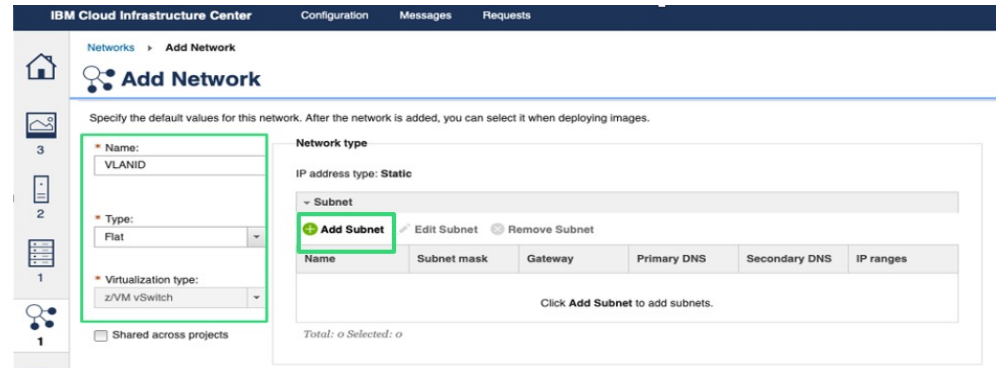
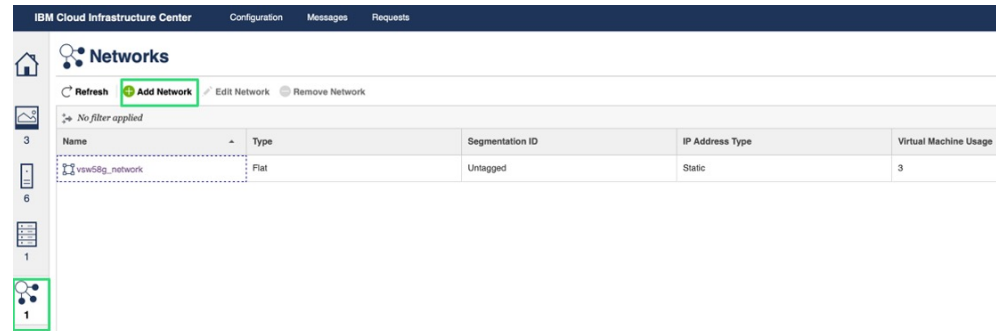
2. Set network parameters for provisioning

Add Network:

- Enter a “Name: icic_network”
- Select the “Type: Flat”
- Select “Virtualization type: z/VM vSwitch”

Add Subnet:

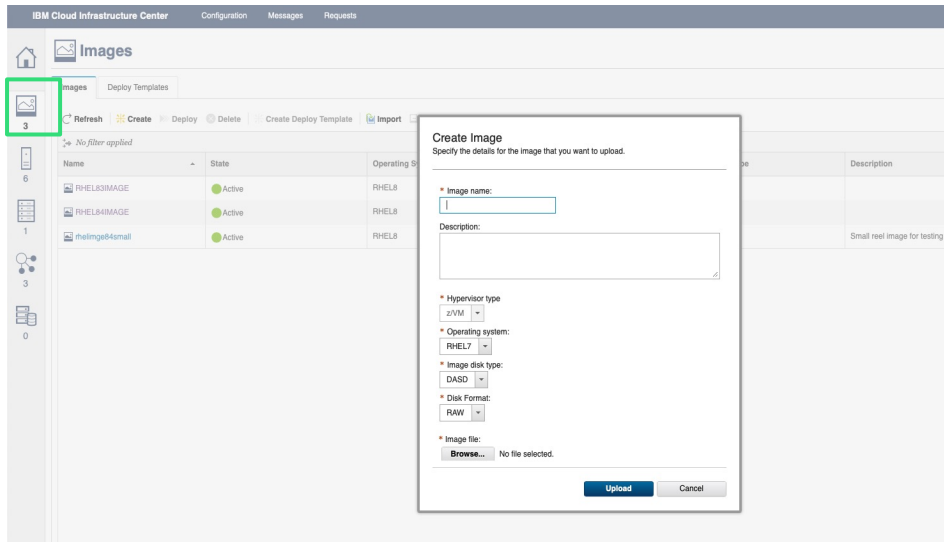
- Enter a “Name: icic_subnet”
- Enter the subnet mask of your network (e.g., 255.255.0.0)
- Enter the IP address of your gateway (e.g., 192.168.0.1)
- Enter the ip address of your primary DNS (e.g., 192.168.0.1)
- Select the IP address range for provisioning virtual machines



Step 10: Deployment of the virtual machines

3. Upload an image for provisioning:

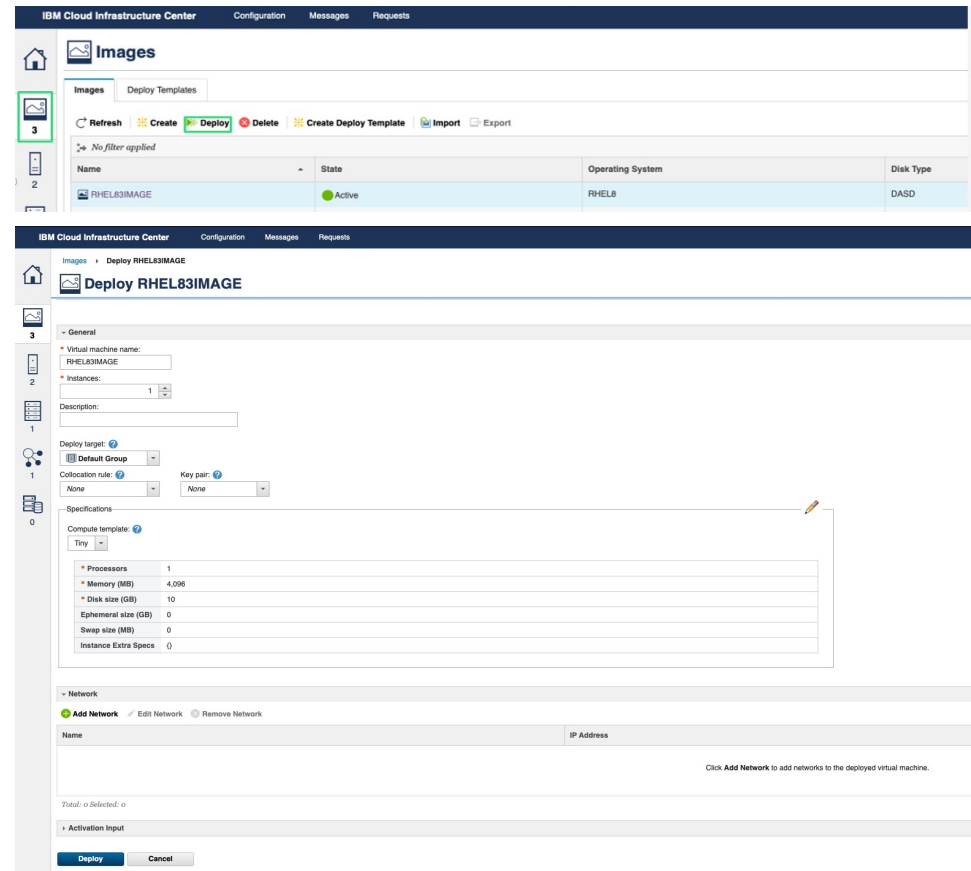
- Click on the **“Create”** button
- Enter an **“Image name: master_rhel73_image”**
- Select the **“Hypervisor type: z/VM”**
- Select the **“Operating system: RHEL7”**
- Select **“Image disk type: DASD”**
- Select **“Disk Format: RAW”**
- Select the **“Image file from your local device”**
- Click **“Upload”**



Step 10: Deployment of the virtual machines

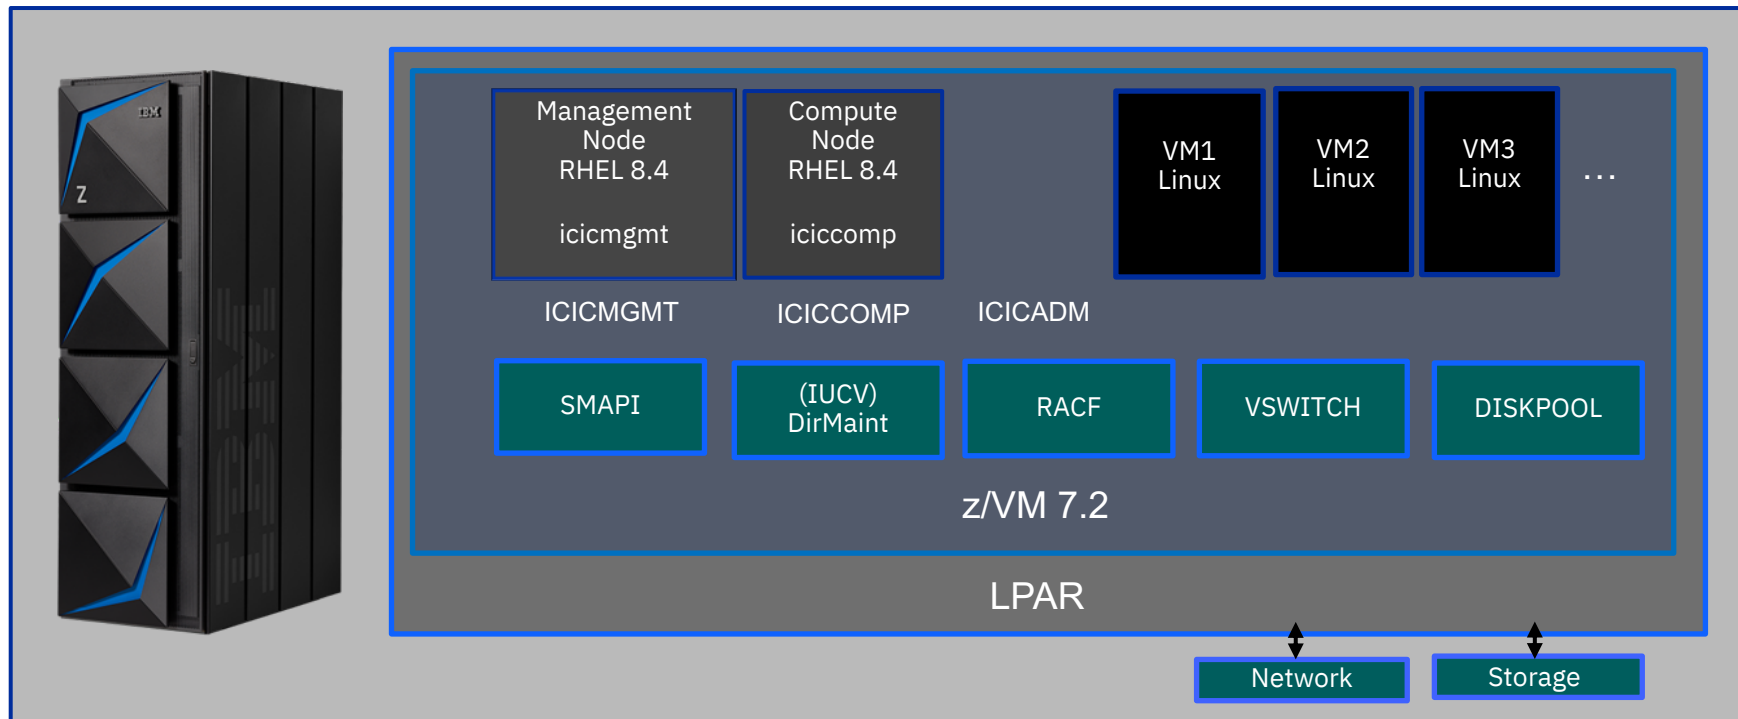
Final step is to create a virtual machine:

- Select the Image from which virtual machine needs to be created from “Images” tab
- Select “**Deploy**”
- “**Virtual machine name**” is by default the image name
- Select number of “**Instances**” to be created
- Select the Group in which Image to be added “**Deploy target: Default Group**”
- Select “**Collocation rule**” and “**Key pair**” (optional)
- Select “**Specification: Tiny**” (size profile)
- Select “**Network: icic_network**”
- Click “**Deploy**”



Overview of the installed Cloud Infrastructure Center architecture

At completion, the created virtual machines are available in production. They can be managed by the Cloud Infrastructure Center GUI on its web browser.



References

IBM Cloud Infrastructure Center information:

- ibm.com/products/cloud-infrastructure-center

IBM Cloud Infrastructure Center installation, operation and reference documentation

- ibm.com/docs/en/cic/1.1.4

IBM z/VM 7.2 documentation

- ibm.com/docs/en/zvm/7.2?topic=SSB27U_7.2.0/com.ibm.zvm.v720/zvminfoc03.htm

IBM Redbooks related to IBM z/VM

- www.vm.ibm.com/pubs/redbooks/

