

# High Availability Clustering with RHEL 8 and z/VM® Getting Started, Hints and Tips

---

Solution Assurance

Bodo Brand, Daniel Kaiser, David Stark

Document version: 03-2022

# 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.

# Notices and disclaimers, *continued*

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, 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](http://www.ibm.com/legal/copytrade.shtml)

# RHEL HA - Agenda

- ❖ **Introduction**
- ❖ **Prerequisites** and **End-Result**
- ❖ **Guidance Notes**

## **1.x - Pacemaker Installation**

**2.x - Fencing** – z/VM admin tasks

**2.x - Fencing** – Create Resources

**3.x - Shared Storage** – z/VM admin tasks

**3.x - Shared Storage** – Create Resources

## **4.x - Apache Installation**


**4.x - Apache** – Create Resources

# Introduction

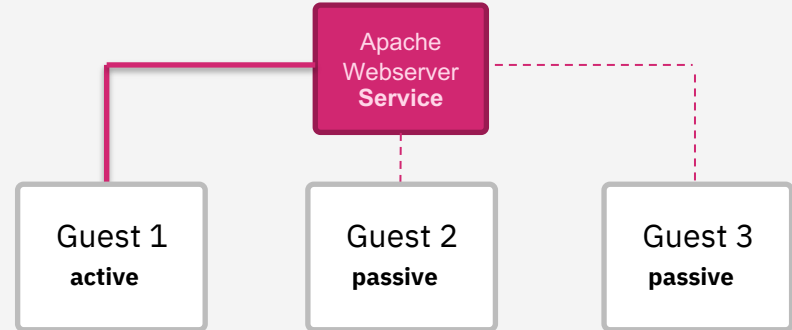
## High Availability:

- For mission-critical workloads
- Official Red Hat® documentation: [LINK](#)
- Support statement: [LINK](#)

“Single Point of Failure”



“Active/Passive Failover”



-> High Availability

# Prerequisites and End-Result

## Setup:

z15™ and DS8000®

3 Nodes

1 ECKD DASD  
Fullpack Minidisk  
(shared storage)

3 ECKD DASD  
(local storage)

## Nodes (z/VM Guests):

On z/VM 7.2

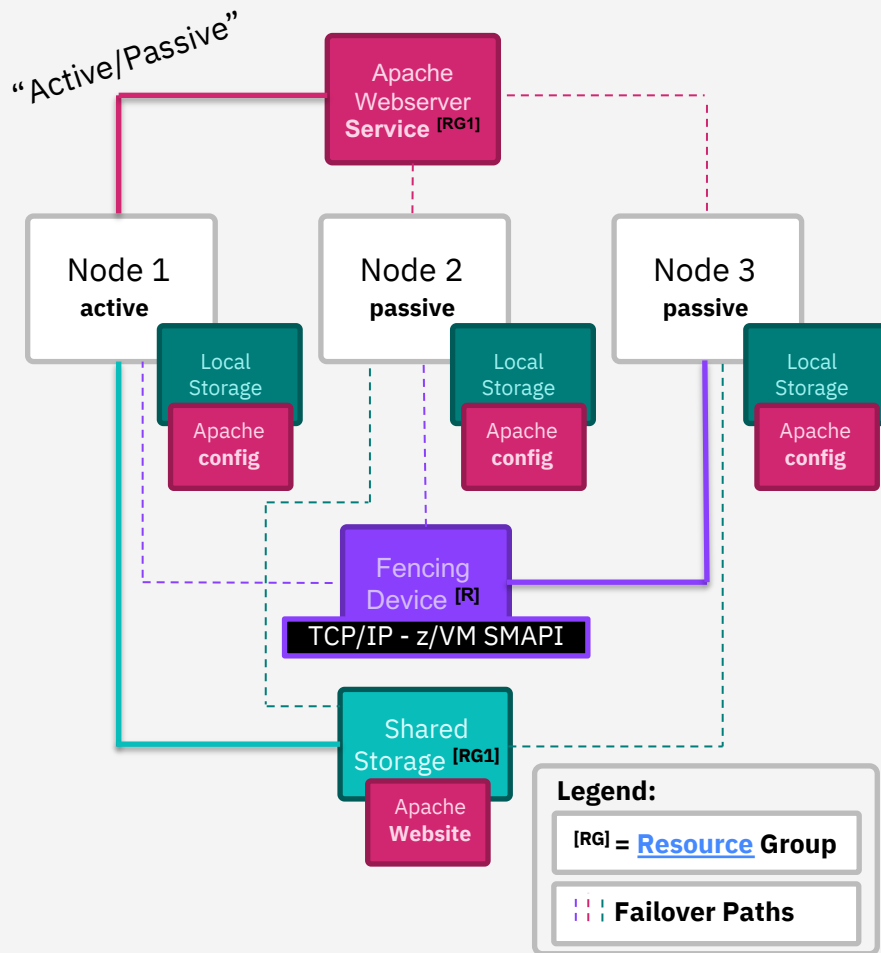
Distro: RHEL 8.3

## Resources

are marked  
with [R]

**On Failover:**  
Resources  
Migrate to  
**one** of these  
nodes.

**Migration** only  
happens when  
Quorum is  
reached.



# Guidance Notes

- Some of the operations must be run on all nodes and some only one node.
- The "**Run on**" graphic on the right indicate on which of the nodes you have to run the command.

Run on:

Node 1  
Node 2  
Node 3

- "**#**" at the beginning of the line indicate a privileged bash command.

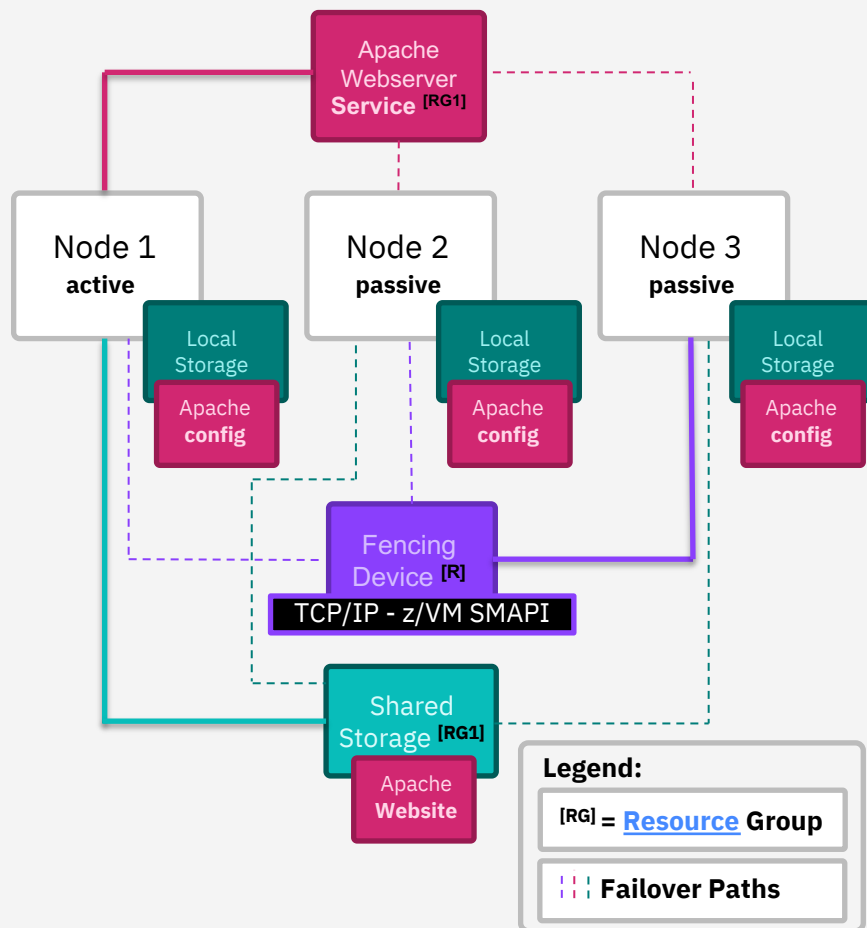
Run on:

Node 1  
Node 2  
Node 3

```
# echo "Example command"
```

- The graphic on the right is used for illustration purposes.

- Optional: Link terminal to run SSH commands in parallel on multiple hosts.  
For example, use Ansible®.



# Pacemaker Installation

## Step 1.0 – [RHEL HA Installation](#)

```
# subscription-manager register --auto-attach
# dnf config-manager --set-enabled \
  rhel-8-for-s390x-highavailability-rpms
# dnf update -y
# dnf install -y pcs pacemaker fence-agents-all
```

Run on:

Node 1

Node 2

Node 3

## Step 1.1 – Firewall Configuration

```
# firewall-cmd --permanent \
  --add-service=high-availability
# firewall-cmd --reload
```

Run on:

Node 1

Node 2

Node 3

## Step 1.2 – Setup Linux® RHEL HA User

```
# passwd hacluster
```

Note:


- Same password for each node recommended.
- Disable password based authentication for SSH.

Run on:

Node 1

Node 2

Node 3





# Pacemaker Installation

## Step 1.3 – Enable pcsd

```
# systemctl enable pcsd.service --now
```

Run on:

Node 1

Node 2

Node 3

## Step 1.4 – Auth Nodes

```
# pcs host auth node1 node2 node3
Username: hacluster
Password: ...
```

Run on:


Node 1

## Step 1.5 – Setup Cluster

```
# pcs cluster setup \
my_cluster node1 node2 node3
```

Run on:

Node 1



# Pacemaker Installation

## Step 1.6 – Start and Enable Services

```
# pcs cluster start --all  
# pcs cluster enable --all
```

Run on:

Node 1

## Step 1.7 – Enable Monitoring (Optional)

```
# dnf install pcp-zeroconf
```

Run on:

Node 1

Node 2

Node 3

## Step 1.8 – Check before Continuing


```
# pcs cluster status
```

Run on:

Node 1

Note:

- Make sure everything is up and running.



# Fencing

# Fencing - z/VM admin tasks

## Step 2.0 – Set up SMAPI Authorization

Run on:

z/VM

- Create ZCLUSTER User:

```
USER ZCLUSTER xxxxxxxx 32M 128M G
  INCLUDE IBMDFLT
  IPL CMS
  MDISK 0191 3390 AUTOV 10 T513T1 MR
```

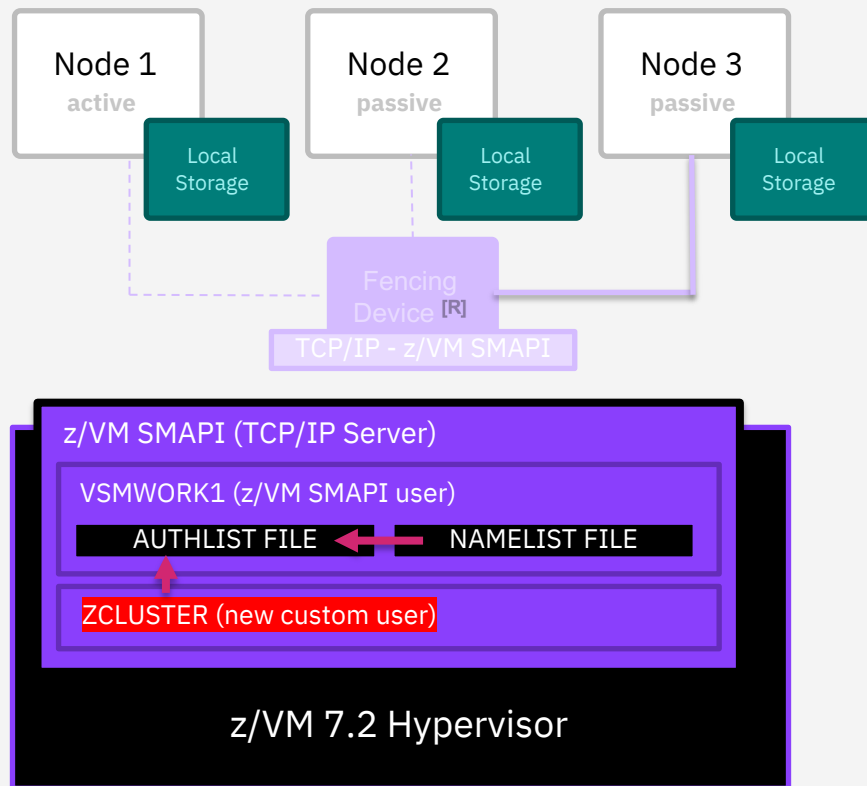
- [Add ZCLUSTER to z/VM Directory](#):

```
DIRM ADD ZCLUSTER
```

- [Set Authorization Policy](#):

```
Authorization_Policy =
Authorization_Policy_AuthlistOnly
```

- [XEDIT documentation](#)



# Fencing - z/VM admin tasks

## Step 2.1 – Set up SMAPI Authorization

Run on:

z/VM

- File VSMWORK1 [NAMELIST](#):

```
:nick.ZVM_FENCE
:list.
IMAGE_ACTIVATE
IMAGE_DEACTIVATE
IMAGE_STATUS_QUERY
CHECK_AUTHENTICATION
IMAGE_NAME_QUERY_DM
```

- File VSMWORK1 [AUTHLIST](#):

COLUMN 1	COLUMN 66	COLUMN 131
V	V	V
DO.NOT.REMOVE	DO.NOT.REMOVE	DO.NOT.REMOVE
ZCLUSTER	ALL	ZVM_FENCE

## Step 2.2 – Enable SMAPI

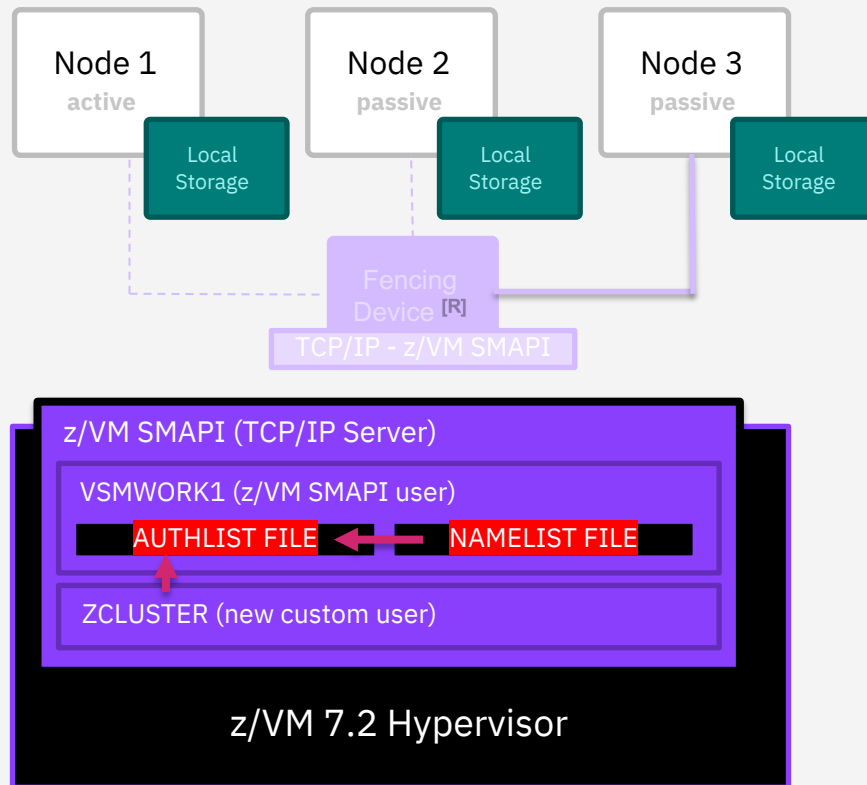
Run on:

z/VM

- [Follow Getting Started with SMAPI in the z/VM Doc.](#)

Note:

- Make sure that after a restart of the z/VM hypervisor that the SMAPI is started.



# Fencing - z/VM admin tasks

## Step 2.3 – IPL DASD on LOGON

Run on:

z/VM

### - Option 1:

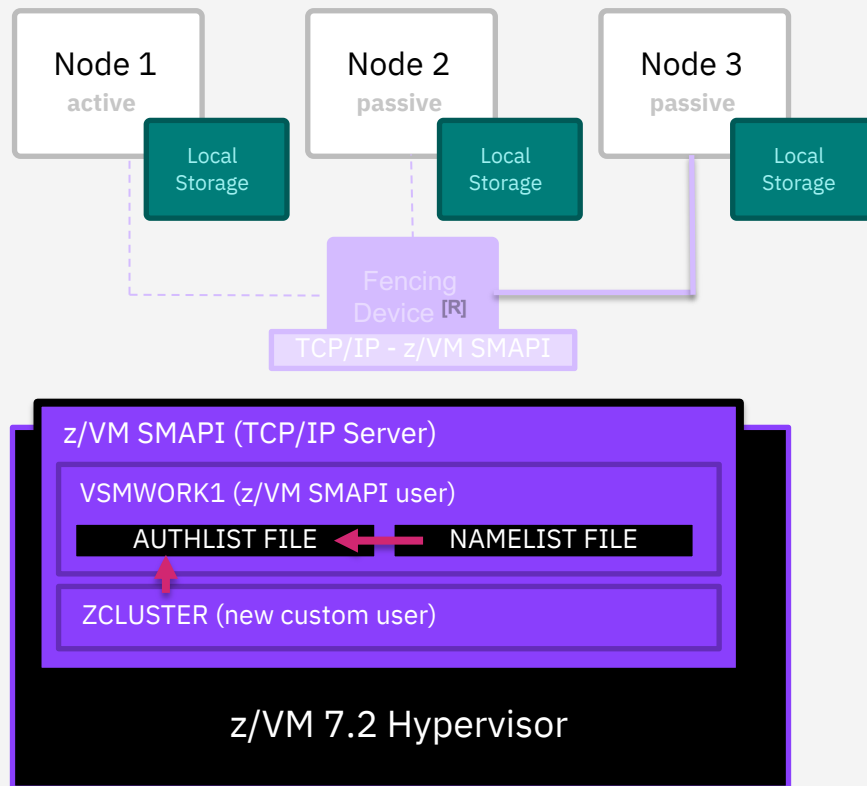
- Add IPL statement to each directory entry.

```
IPL 1E00
```

### - Option 2:

- Add IPL statement to profile exec from CMS.

```
/* */  
'IPL 1E00'
```



# Fencing – Create Resources

## Step 2.4 – Create fence\_zvmip Stonith Resource

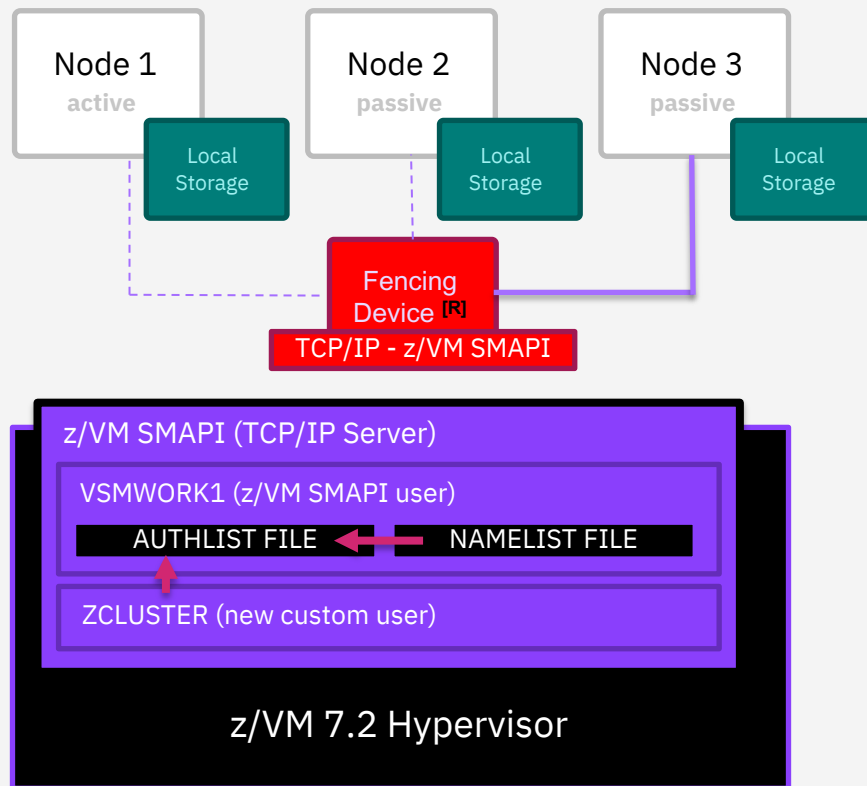
Run on:

Node 1

```
# pcs stonith create zvm-smapi1 fence_zvmip \
    ip=zvm-smapi1.example.com \
    username="ZCLUSTER" \
    password="ZPASSWD" \
    pcmk_host_map="node1:RHELHA-
1;node2:RHELHA-2;node3:RHELHA-3"
```

Note:

```
- pcmk_host_map="<NODENAME>:<IMAGE>[;<NODENAME>:<IMAGE>...]"
```



# Fencing – Create Resources

## Step 2.5 – [Quick Test Fencing](#)

Run on:

Node 1

- Trigger fencing manually:

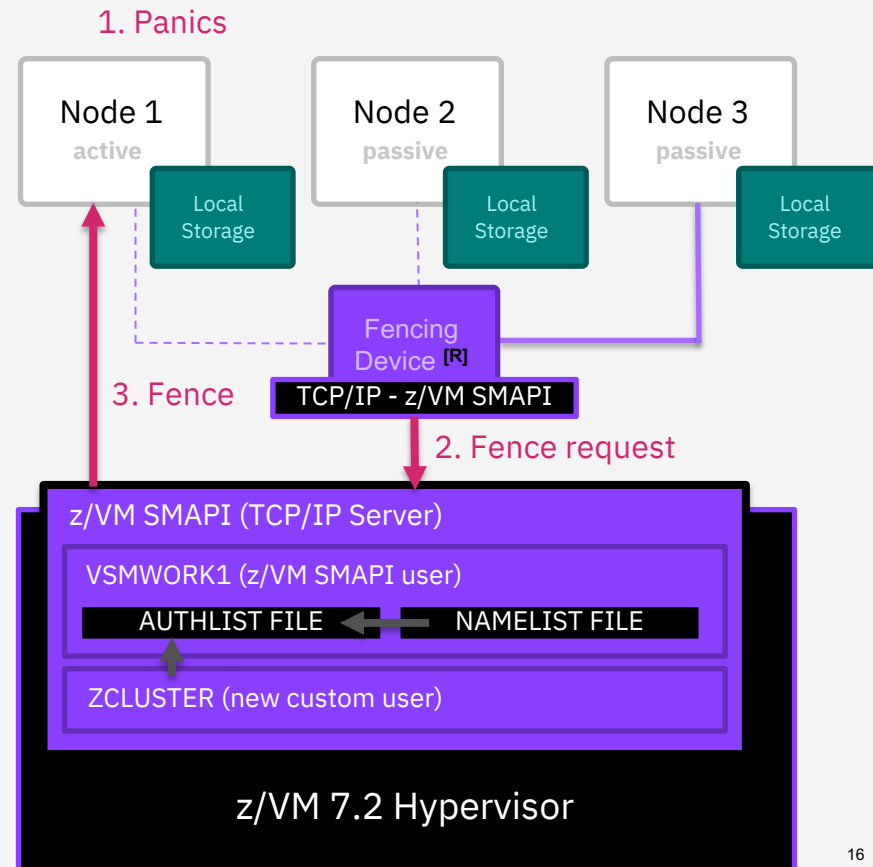
```
# pcs stonith fence node1
```

- Simulate a panic:

```
# echo c > /proc/sysrq-trigger
```

Note:

- Check the "pcs stonith history" to make sure fencing was successful.





# Shared Storage

# Shared Storage - z/VM admin tasks

## Step 3.0 – Create Fullpack Minidisk

Run on:

z/VM

- Create new user LINSHARE: (See [Step 2.0](#))

```
USER LINSHARE NOLOG  
MDISK 0200 3390 DEVNO 1111 MWV
```

## Step 3.1 – Link Minidisk to Linux

Run on:

z/VM

- Option 1 (on z/VM hypervisor):

- Link both Minidisks

```
DIRM FOR RHELHA-1 LINK LINSHARE 200 200 MW  
DIRM FOR RHELHA-2 LINK LINSHARE 200 200 MW  
DIRM FOR RHELHA-3 LINK LINSHARE 200 200 MW
```

- LOGOFF and LOGON each guest

- Option 2 (on Linux guest):

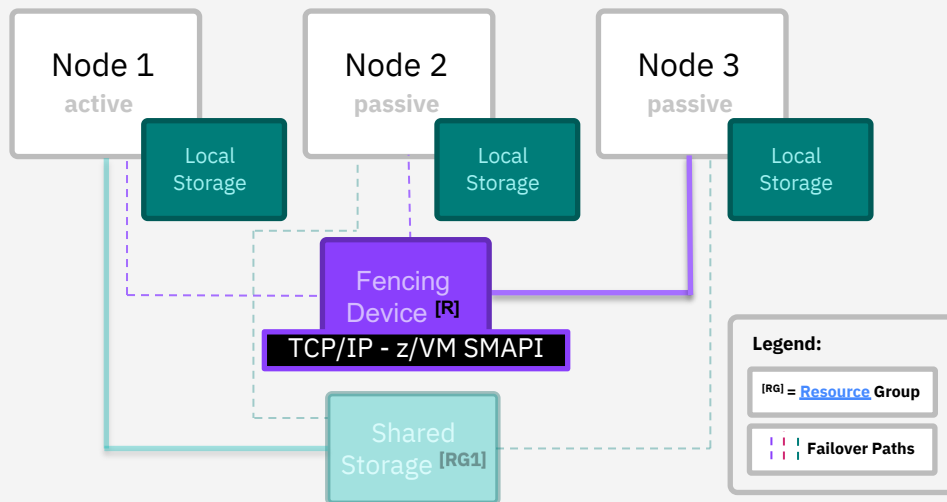
```
# vmcp 'link * 0200 0200 rw'
```

Run on:

Node 1

Node 2


Node 3



Legend:

[RG] = [Resource Group](#)

--- Failover Paths



# Shared Storage – Create Resources

## Step 3.2 – Make DASD persistent

De-ignore, enable and make DASD persistent:

```
# cio_ignore -r 0200
# chccwdev -e 0200
# echo 0.0.0200 >> /etc/dasd.conf
# echo add > /sys/bus/ccw/devices/0.0.0200/uevent
```

(Alternatively you can use "[chzdev](#) -e dasd 0.0.0200")

Run on:

Node 1

Node 2

Node 3

## Step 3.3 – Format DASD

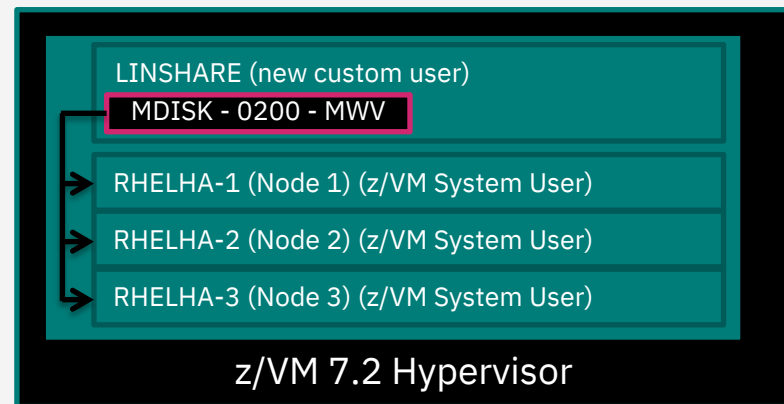
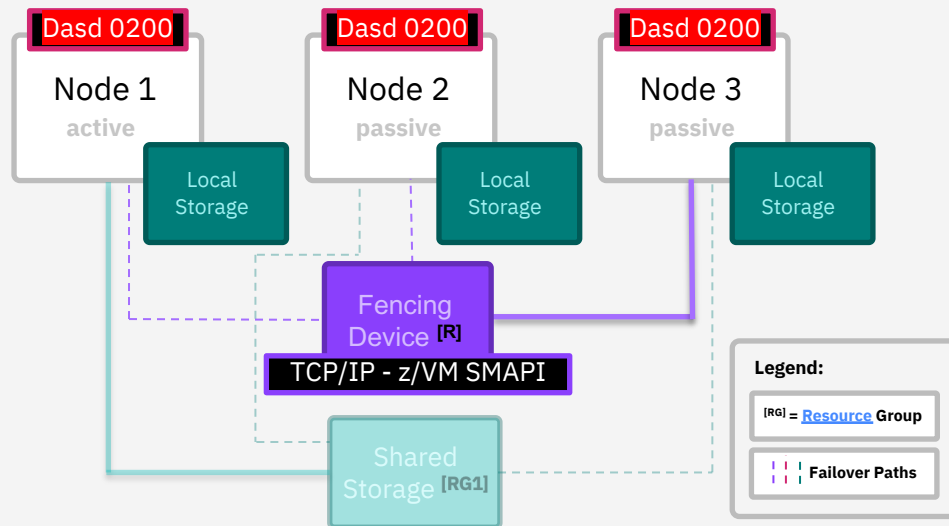
Run on:

Node 1

```
# dasdfmt -b 4096 -d cdl -p /dev/dasdc
# fdasd -a /dev/dasdc
```

Note:

- /dev/dasdc might be named differently on your system – check with "lsdasd".



# Shared Storage – Create Resources

## Step 3.4 – Change LVM config

Run on:

Node 1

- Change option in /etc/lvm/lvm.conf to:  
system\_id\_source = "uname"

## Step 3.5 – Create LVM volume

Run on:

Node 1

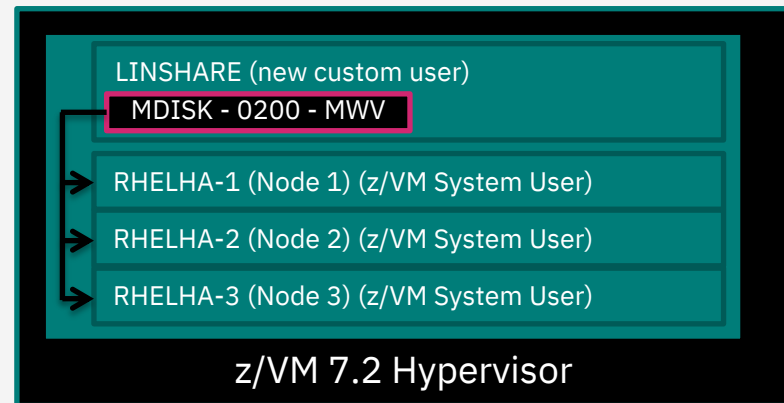
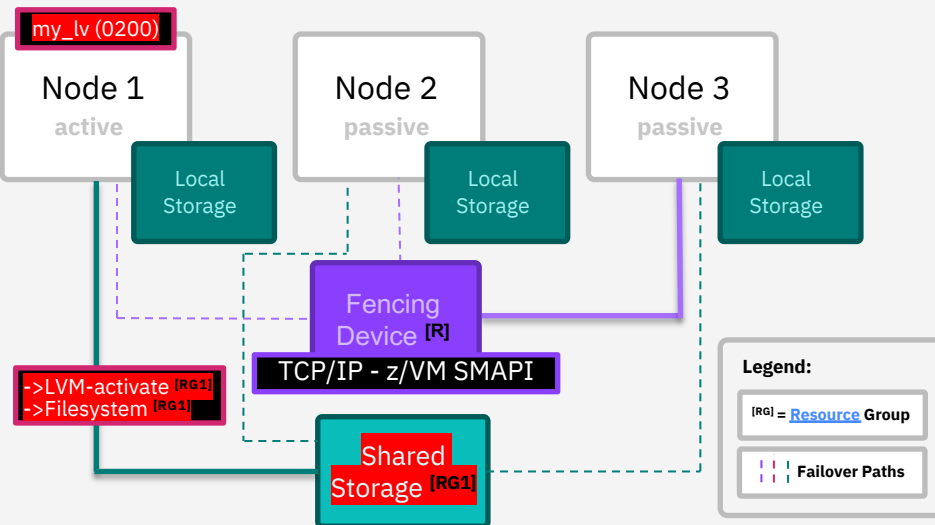
```
# pvcreate /dev/dasdc1
# vgcreate my_vg /dev/dasdc1
# lvcreate -n my_lv -l 100%FREE my_vg
# mkfs.ext4 /dev/my_vg/my_lv
```

## Step 3.6 – Create resources

Run on:

Node 1

```
# pcs resource create my_lvm LVM-activate \
    vgname=my_vg \
    vg_access_mode=system_id \
    --group apachegroup
# pcs resource create my_fs Filesystem \
    device="/dev/my_vg/my_lv" \
    directory="/var/www" fstype="ext4" \
    --group apachegroup
```



# Apache HA

# Apache Installation

## Step 4.0 – Apache Installation

```
# dnf install -y httpd wget
```

Run on:

Node 1

Node 2

Node 3

## Step 4.1 – Firewall Configuration

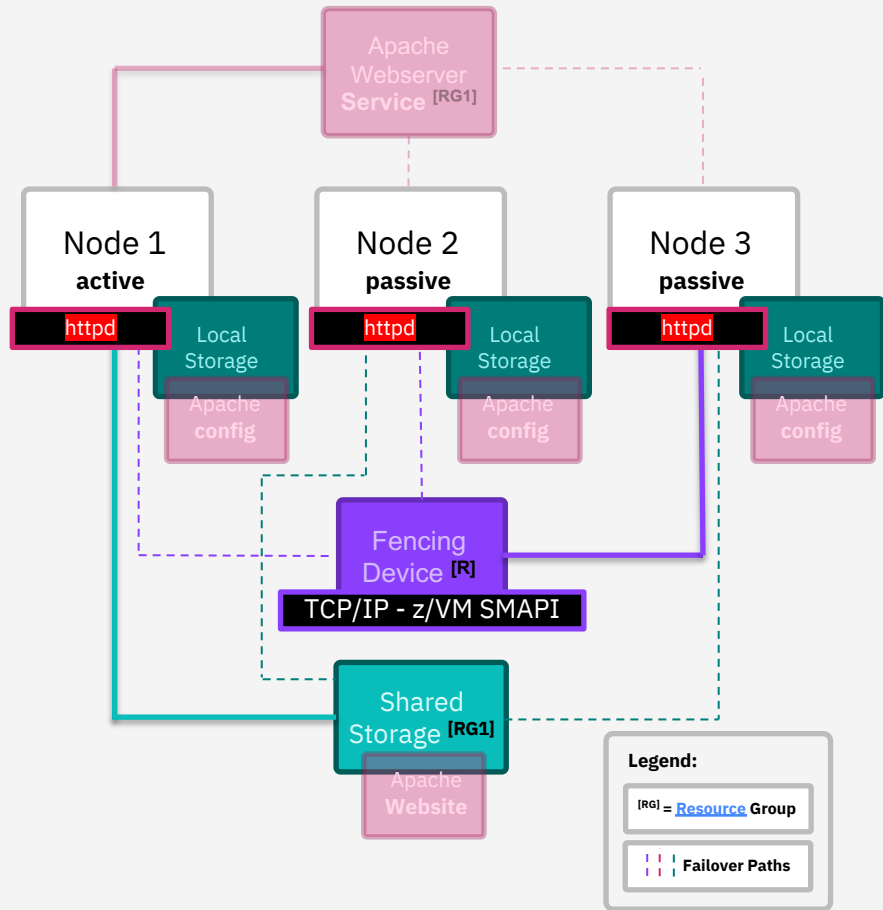
```
# firewall-cmd --permanent \
  --add-service=http
# firewall-cmd --reload
```

Run on:

Node 1

Node 2

Node 3



# Apache – Create Resources

## Step 4.2 – Create web site

Run on:

Node 1

```
# cat <<-END > /var/www/html/index.html
<html>
<body>My Test Site - $(hostname)</body>
</html>
END
# restorecon -R /var/www
```

## Step 4.3 – Create configuration file

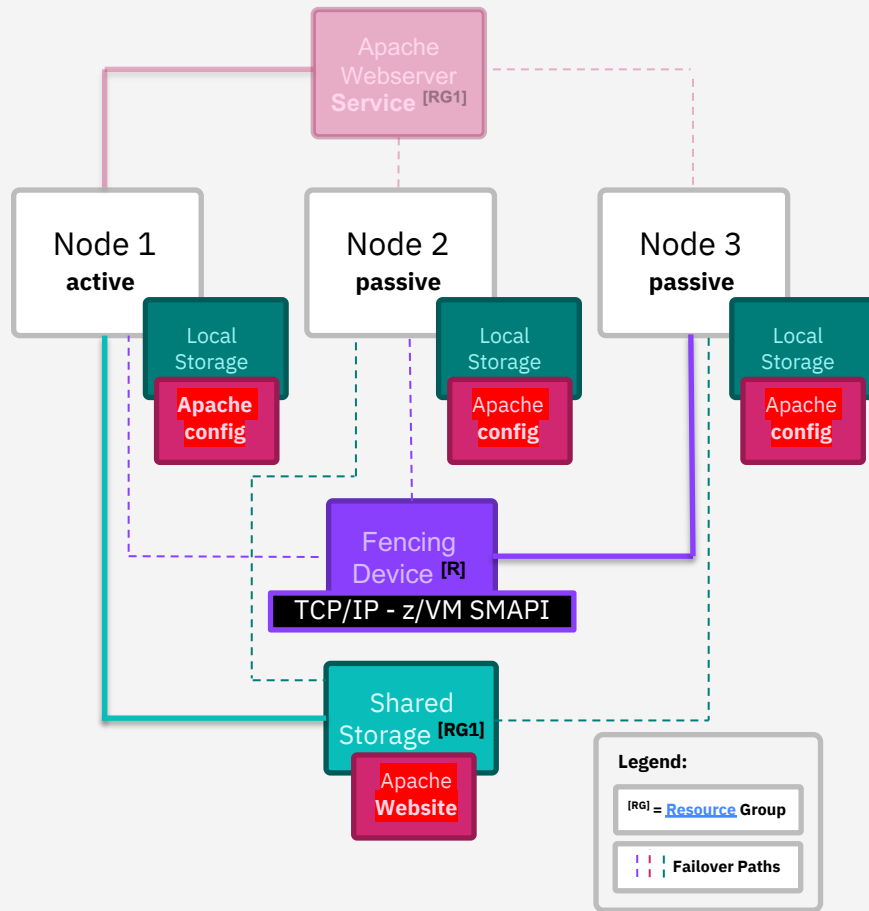
Run on:

Node 1

Node 2

Node 3

```
# cat <<-END > /etc/httpd/conf.d/status.conf
<Location /server-status>
SetHandler server-status
Order deny,allow
Deny from all
Allow from 127.0.0.1
Allow from ::1
</Location>
END
```



# Apache – Create Resources

## Step 4.4 – [Edit logrotate](#)

- **Replace** in the file `/etc/logrotate.d/httpd`:

```
/bin/systemctl reload httpd.service > /dev/null  
2>/dev/null || true
```

**with:**

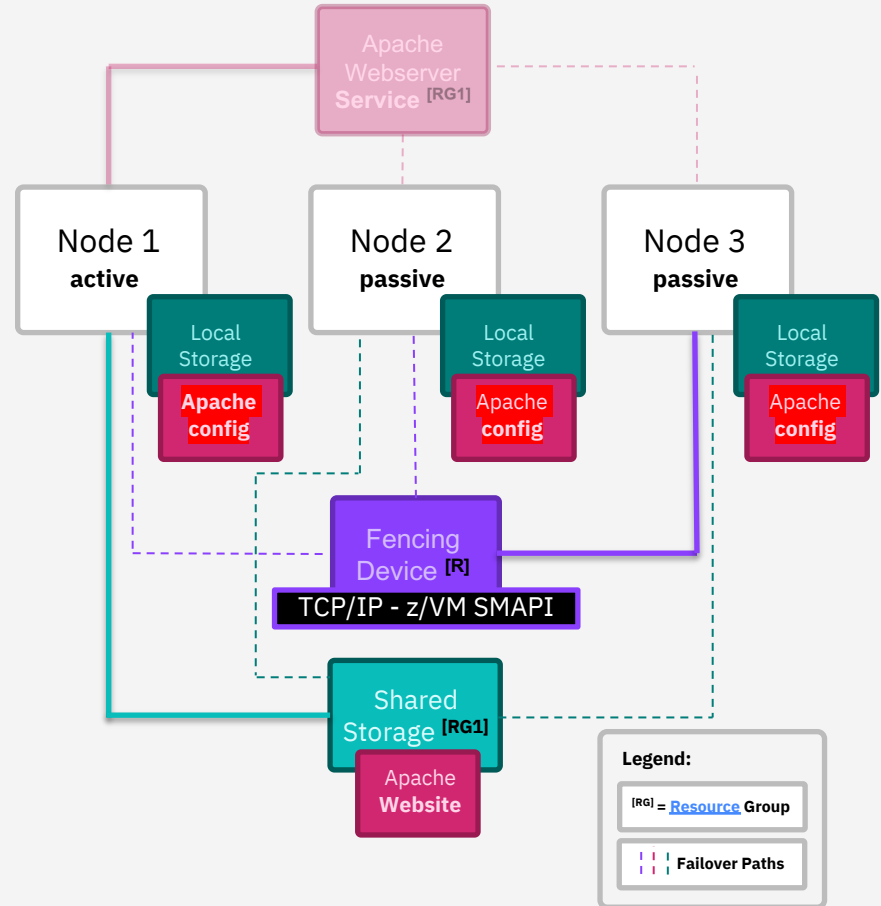
```
/usr/bin/test -f /run/httpd.pid >/dev/null  
2>/dev/null && /usr/bin/ps -q $(/usr/bin/cat  
/run/httpd.pid) >/dev/null 2>/dev/null &&  
/usr/sbin/httpd -f /etc/httpd/conf/httpd.conf \ -c  
"PidFile /run/httpd.pid" -k graceful > /dev/null  
2>/dev/null || true
```

Run on:

Node 1

Node 2

Node 3





# Apache – Create Resources

## Step 4.5 – Create resources

Run on:

Node 1

```
# pcs resource create ClusterIP
ocf:heartbeat:IPaddr2
ip="${FLOATING_IP_ADDRESS}" --group
apachegroup --before my_lvm
```

```
# pcs resource create WebSite
ocf:heartbeat:apache
configfile=/etc/httpd/conf/httpd.conf
statusurl="http://localhost/server-status" -
-group apachegroup --after my_fs
```

Note:

- Floating IP addresses is similar to a static IP address in the network.
- \${FLOATING\_IP\_ADDRESS} is a placeholder.

## Step 4.6 – Test with wget

Run on:

Node 2

```
# wget http://${FLOATING_IP_ADDRESS}
```

