

# SAP Convergent Mediation ControlZone High Availability Cluster

## Setup Guide

SUSE Linux Enterprise Server for SAP Applications 15  
SAP Convergent Mediation

Fabian Herschel, Distinguished Architect SAP (SUSE)  
Lars Pinne, Systems Engineer (SUSE)

# SAP Convergent Mediation ControlZone High Availability Cluster

## Setup Guide

**Date:** 2025-03-13

SUSE® Linux Enterprise Server for SAP Applications is optimized in various ways for SAP® applications. This document explains how to configure a Convergent Mediation ControlZone High Availability Cluster solution. It is based on SUSE Linux Enterprise Server for SAP Applications 15 SP4. The concept, however, can also be used with newer service packs of SUSE Linux Enterprise Server for SAP Applications.

**Disclaimer:** Documents published as part of the SUSE Best Practices series have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

# Contents

- 1 About this guide 4
- 2 Overview 5
- 3 Checking the operating system and the HA cluster basic setup 11
- 4 Checking the ControlZone setup 17
- 5 Integrating Convergent Mediation ControlZone with the Linux cluster 19
- 6 Administration 35
- 7 References 38
- 8 Appendix 43
- 9 Legal notice 47
- 10 GNU Free Documentation License 48

# 1 About this guide

The following sections focus on background information and purpose of this document.

## 1.1 Abstract

This guide describes configuration and basic testing of SUSE Linux Enterprise Server for SAP Applications 15 SP4 as a high availability cluster for Convergent Mediation (CM) ControlZone services.

From an application perspective, the following concept is covered:

- ControlZone platform and UI services are running together.
- ControlZone software is installed on central NFS.
- ControlZone software is copied to local disks of both nodes.

From an infrastructure perspective, the following concept is covered:

- Two-node cluster with disk-based SBD fencing
- Central NFS share statically mounted on both nodes
- On-premises deployment on physical and virtual machines

Despite the above mentioned focus of this setup guide, other variants can be implemented as well. See [Section 2, "Overview"](#) below. The concept can also be used with newer service packs of SUSE Linux Enterprise Server for SAP Applications 15.



### Note

This solution is supported only in the context of SAP RISE (<https://www.sap.com/products/erp/rise.html> .

## 1.2 Additional documentation and resources

Several chapters in this document contain links to additional documentation resources which are available either in the system or on the Internet.

For the latest product documentation updates, see:

<https://documentation.suse.com/> 

More whitepapers, guides and best practices documents referring to SUSE Linux Enterprise Server and SAP can be found and downloaded at the SUSE Best Practices Web page:

<https://documentation.suse.com/sbp/sap/> ↗

Here you can access guides for SAP HANA system replication automation and High Availability (HA) scenarios for SAP NetWeaver and SAP S/4HANA.

Find an overview of high availability solutions supported by SUSE Linux Enterprise Server for SAP Applications here:

<https://documentation.suse.com/sles-sap/sap-ha-support/html/sap-ha-support/article-sap-ha-support.html> ↗

Finally, there are manual pages shipped with the product.

## 1.3 Feedback

Several feedback channels are available:

### Bugs and Enhancement Requests

For services and support options available for your product, refer to <http://www.suse.com/support/> ↗.

To report bugs for a product component, go to <https://scc.suse.com/support/> ↗ requests, log in, and select *Submit New SR* (Service Request).

### Mail

For feedback on the documentation of this product, you can send a mail to [doc-team@suse.com](mailto:doc-team@suse.com) (<mailto:doc-team@suse.com>) ↗. Make sure to include the document title, the product version and the publication date of the documentation. To report errors or suggest enhancements, provide a concise description of the problem and refer to the respective section number and page (or URL).

## 2 Overview

The CM ControlZone platform is responsible for providing services to other instances. Several platform containers may exist in a CM system, for high availability, but only one is active at a time. The CM ControlZone UI is used to query, edit, import, and export data.

SUSE Linux Enterprise Server for SAP Applications is optimized in various ways for SAP applications. Particularly, it contains the SUSE Linux Enterprise High Availability cluster and specific HA resource agents.

From the application perspective the following variants are covered:

- ControlZone platform service running alone
- ControlZone platform and UI services running together
- ControlZone binaries stored and started on central NFS (not recommended)
- ControlZone binaries copied to and started from local disks
- Java VM stored and started on central NFS (not recommended)
- Java VM started from local disks

From the infrastructure perspective the following variants are covered:

- Two-node cluster with disk-based SBD fencing
- Three-node cluster with disk-based or diskless SBD fencing, not explained in detail here
- Other fencing is possible, but not explained here
- File system managed by the cluster - either on shared storage or NFS, not explained in detail here
- On-premises deployment on physical and virtual machines
- Public cloud deployment (usually needs additional documentation on cloud specific details)

## 2.1 High availability for the Convergent Mediation ControlZone platform and UI

The HA solution for CM ControlZone is a two-node active/passive cluster. A shared NFS file system is statically mounted by the operating system on both cluster nodes. This file system holds work directories. Client-side write caching needs to be disabled. The ControlZone software is installed into the central shared NFS, but is also copied to both nodes' local file systems. The HA cluster uses the central directory for starting/stopping the ControlZone services. However, for monitoring the local copies of the installation are used.

The cluster can run monitor actions even when the NFS temporarily is blocked. Further, software upgrade is possible without downtime (rolling upgrade).

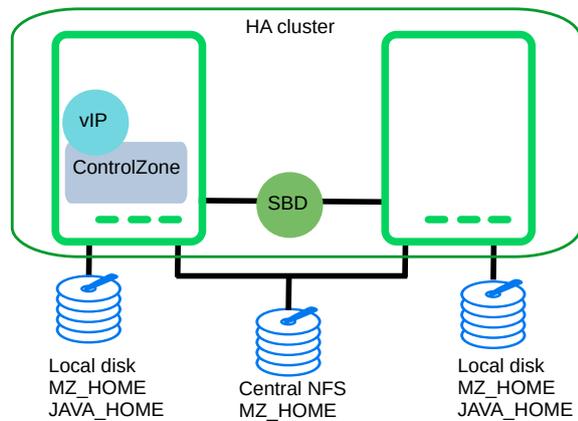


FIGURE 1: TWO-NODE HA CLUSTER AND STATICALLY MOUNTED FILE SYSTEMS

The ControlZone services platform and UI are handled as active/passive resources. The related virtual IP address is managed by the HA cluster as well. A file system resource is configured for a bind-mount of the real NFS share. In case of file system failures, the cluster takes action. However, no mount or umount on the real NFS share is done.

All cluster resources are organized as one resource group. This results in a correct start/stop order and placement, while keeping the configuration simple.

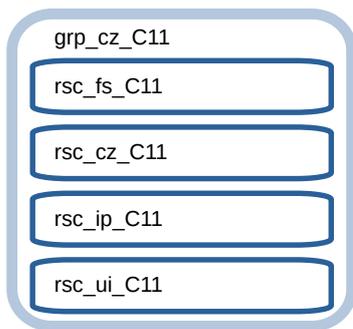


FIGURE 2: CONTROLZONE RESOURCE GROUP

See [Section 5, "Integrating Convergent Mediation ControlZone with the Linux cluster"](#) and manual page `ocf_suse_SAPCMControlZone(7)` for details.

## 2.2 Scope of this document

For the SUSE Linux Enterprise High Availability two-node cluster described above, this guide explains how to:

- check basic settings of the two-node HA cluster with disk-based SBD.
- check basic capabilities of the ControlZone components on both nodes.
- configure an HA cluster for managing the ControlZone components platform and UI, together with related IP address.
- perform functional tests of the HA cluster and its resources.
- perform basic administrative tasks on the cluster resources.



### Note

Neither installation of the basic SUSE Linux Enterprise High Availability cluster, nor installation of the CM ControlZone software is covered in the document at hand.

Consult the SUSE Linux Enterprise High Availability product documentation at <https://documentation.suse.com/sle-ha/15-SP4/single-html/SLE-HA-administration/#part-install> for installation instructions.

For Convergent Mediation installation instructions, refer to the respective product documentation at <https://infozone.atlassian.net/wiki/spaces/MD9/pages/4849683/Installation+Instructions>.

## 2.3 Prerequisites

For requirements of Convergent Mediation ControlZone, refer to the product documentation at <https://infozone.atlassian.net/wiki/spaces/MD9/pages/4849685/System+Requirements>.

For requirements of SUSE Linux Enterprise Server for SAP Applications and SUSE Linux Enterprise High Availability, refer to the product documentation at <https://documentation.suse.com/sle-ha/15-SP4/html/SLE-HA-all/article-installation.html#sec-ha-inst-quick-req>.

Specific requirements of the SUSE high availability solution for CM ControlZone are as follows:

- This solution is supported only in the context of SAP RISE.
- Convergent Mediation ControlZone version 9.0.1.1 or higher is installed and configured on both cluster nodes. If the software is installed into a shared NFS file system, the binaries are copied into both cluster nodes' local file systems. Finally the local configuration needs to be adjusted. Refer to Convergent Mediation documentation for details.
- CM ControlZone is configured identically on both cluster nodes. User, path names and environment settings are the same.
- There is only **one** ControlZone instance per Linux cluster. Accordingly, there is only one platform service and one UI service per cluster.
- The platform and UI are installed into the same MZ\_HOME.
- Linux shell of the mzadmin user is /bin/bash.
- The mzadmin's ~/.bashrc inherits MZ\_HOME, JAVA\_HOME and MZ\_PLATFORM from SAPCMControlZone RA. These variables need to be set as described in the RA's documentation, that is the manual page ocf\_suse\_SAPCMControlZone(7).
- When called by the resource agent, mzsh connects to CM ControlZone services via network. The service's virtual host name or virtual IP address managed by the cluster should not be used when called by RA monitor actions.
- Technical users and groups are defined locally in the Linux system. If users are resolved by remote service, local caching is necessary. Substitute user (su) to the mzadmin needs to work reliable and without customized actions or messages.
- Name resolution for host names and virtual host names is crucial. Host names of cluster nodes and services are resolved locally in the Linux system.
- Strict time synchronization between the cluster nodes, for example NTP, is required. All nodes of a cluster have configured the same timezone.
- Needed NFS shares (for example /usr/sap/<SID>) are mounted statically or by automounter. No client-side write caching is happening. File locking should be configured for application needs.
- The RA monitoring operations need to be active.

- RA runtime almost completely depends on call-outs to controlled resources, operating system, and Linux cluster. The infrastructure needs to allow these call-outs to return in time.
- The ControlZone application is not started/stopped by the operating system. Thus, there is no SystemV, systemd or cron job.
- As long as the ControlZone application is managed by the Linux cluster, the application is not started/stopped/moved from outside. Thus, no manual actions are done. The Linux cluster does not prevent from administrative mistakes. However, if the Linux cluster detects the application is running on both sites in parallel, both are stopped and one of them is restarted.
- The interface for the RA to the ControlZone services is the command `mzsh`. Ideally, `mzsh` should be accessed on the cluster nodes' local file systems. `mzsh` is called with the arguments `startup -f`, `shutdown` and `status`. Its return code and output is interpreted by the RA. Thus, the command and its output need to be stable. The `mzsh` shall not be customized. Particularly environment variables set thru `~/ .bashrc` must not be changed.
- `mzsh` is called on the active node with a defined interval for regular resource monitor operations. It also is called on the active or passive node in certain situations. Those calls might run in parallel.

## 2.4 The setup procedure at a glance

For a better understanding and overview, the installation and setup is divided into nine steps.

- *Section 3.1, "Collecting information"*
- *Section 3.2, "Checking the operating system basic setup"*
- *Section 3.3, "Checking HA cluster basic setup"*
- *Section 4, "Checking the ControlZone setup"*
- *Section 5.1, "Preparing mzadmin user ~/ .bashrc file"*
- *Section 5.2, "Preparing the operating system for NFS monitoring"*
- *Section 5.3, "Adapting the cluster basic configuration"*
- *Section 5.4, "Configuring ControlZone cluster resources"*
- *Section 5.7, "Testing the HA cluster"*

## 3 Checking the operating system and the HA cluster basic setup

### 3.1 Collecting information

The installation should be planned properly. You should have all required parameters already in place. It is good practice to first fill out the parameter sheet.

TABLE 1: TABLE COLLECTING NEEDED PARAMETERS

Parameter	Example	Value
NFS server and share	<u>192.168.1.1:/s/C11/cm</u>	
NFS mount options	<u>vers=4, rw, noac, sync, de- fault</u>	
central MZ_HOME	<u>/usr/sap/C11</u>	
local MZ_HOME	<u>/opt/cm/C11</u>	
MZ_PLATFORM	<u>http://localhost:9000</u>	<u>http://localhost:9000</u>
JAVA_HOME	<u>/usr/lib64/jvm/jre-17- openjdk</u>	
node1 hostname	<u>akka1</u>	
node2 hostname	<u>akka2</u>	
node1 IP addr	<u>192.168.1.11</u>	
node2 IP addr	<u>192.168.1.12</u>	
SID	<u>C11</u>	
mzadmin user	<u>c11adm</u>	
virtual IP addr	<u>192.168.1.112</u>	
virtual hostname	<u>c11cz</u>	

## 3.2 Checking the operating system basic setup

### 3.2.1 Java virtual machine

See <https://infozone.atlassian.net/wiki/spaces/MD9/pages/4849685/System+Requirements> for supported Java VMs.

```
# zypper se java-17-openjdk

S | Name | Summary | Type
---+-----+-----+-----
i | java-17-openjdk | OpenJDK 17 Runtime Environment | package
  | java-17-openjdk-demo | OpenJDK 17 Demos | package
  | java-17-openjdk-devel | OpenJDK 17 Development Environment | package
  | java-17-openjdk-headless | OpenJDK 17 Runtime Environment | package
```

Check this on both nodes.

### 3.2.2 HA software and tools

```
# zypper se --type pattern ha_sles

S | Name | Summary | Type
---+-----+-----+-----
i | ha_sles | High Availability | pattern
```

```
# zypper se ClusterTools2

S | Name | Summary | Type
---+-----+-----+-----
i | ClusterTools2 | Tools for cluster management | package
```

Check this on both nodes.

### 3.2.3 IP addresses and virtual names

Check if the file `/etc/hosts` contains at least the address resolution for both cluster nodes `akka1`, `akka2`, and the ControlZone virtual host name `sapc11cz`. Add these entries if they are missing.

```
# grep -e akka1 -e akka2 -e c11cz /etc/hosts
```

```
192.168.1.11 akka1.fjaell.lab akka1
192.168.1.12 akka2.fjaell.lab akka2
192.168.1.112 c11cz.fjaell.lab c11cz
```

Check this on both nodes. See also manual page [hosts\(8\)](#).

### 3.2.4 Mount points and NFS shares

Check if the file `/etc/fstab` contains the central NFS share `MZ_HOME`. The file system is statically mounted on all nodes of the cluster. The correct mount options depend on the NFS server. However, client-side write caching needs to be disabled in any case.

```
# grep "/usr/sap/C11" /etc/fstab

192.168.1.1:/s/C11/cz /usr/sap/C11 nfs4 rw,noac,sync,default 0 0

# mount | grep "/usr/sap/C11"

...
```

Check this on both nodes. See also manual page [mount\(8\)](#), [fstab\(5\)](#) and [nfs\(5\)](#), and TID 20830, TID 19722.

### 3.2.5 Linux user and group number scheme

Check if the file `/etc/passwd` contains the mzadmin user `c11adm`.

```
# grep c11adm /etc/passwd

c11adm:x:1001:100:Convergent Mediation user:/opt/cm/C11:/bin/bash
```

Check this on both nodes. See also manual page [passwd\(5\)](#).

### 3.2.6 Password-free SSH login

```
akka1:~ # ssh akka2
akka2:~ # ssh akka1
akka1:~ # exit
akka2:~ # exit
```

Check this on both nodes. See also manual page `ssh(1)` and `ssh-keygen(1)`.

### 3.2.7 Time synchronization

```
# systemctl status chronyd | grep Active

Active: active (running) since Tue 2024-05-14 16:37:28 CEST; 6min ago

# chronyc sources

MS Name/IP address          Stratum Poll Reach LastRx Last sample
=====
^* long.time.ago            2    10   377   100  -1286us[-1183us] +/-  15ms
```

Check this on both nodes. See also manual page `chronyc(1)` and `chrony.conf(5)`.

## 3.3 Checking HA cluster basic setup

### 3.3.1 Watchdog

Check if the **watchdog** module is loaded correctly.

```
# lsmod | grep -e dog -e wdt

iTCO_wdt          16384  1
iTCO_vendor_support 16384  1 iTCO_wdt

# ls -l /dev/watchdog

crw----- 1 root root 10, 130 May 14 16:37 /dev/watchdog

# lsof dev/watchdog

COMMAND PID USER  FD  TYPE DEVICE SIZE/OFF NODE NAME
sbd     686  root   4w   CHR 10,130      0t0  410 /dev/watchdog
```

Check this on both nodes. Both nodes should use the same watchdog driver. Which driver that is depends on your hardware or hypervisor. For more information, see <https://documentation.suse.com/sle-ha/15-SP4/single-html/SLE-HA-administration/#sec-ha-storage-protect-watchdog>.

### 3.3.2 SBD device

It is a good practice to check if the SBD device can be accessed from both nodes and contains valid records. Only one SBD device is used in this example. For production, three devices should always be used.

```
# egrep -v "(^#|^$)" /etc/sysconfig/sbd

SBD_PACEMAKER=yes
SBD_STARTMODE="clean"
SBD_WATCHDOG_DEV="/dev/watchdog"
SBD_WATCHDOG_TIMEOUT="20"
SBD_TIMEOUT_ACTION="flush, reboot"
SBD_MOVE_TO_ROOT_CGROUP="auto"
SBD_OPTS=""
SBD_DEVICE="/dev/disk/by-id/Example-A-part1"

# cs_show_sbd_devices

==Dumping header on disk /dev/disk/by-id/Example-A-part1
Header version      : 2.1
UUID                : 0f4ea13e-fab8-4147-b9b2-3cdcfff07f86
Number of slots     : 255
Sector size         : 512
Timeout (watchdog)  : 20
Timeout (allocate)  : 2
Timeout (loop)      : 1
Timeout (msgwait)   : 120
==Header on disk /dev/disk/by-id/Example-A-part1 is dumped
0   akka1         clear
0   akka2         clear

# systemctl status sbd | grep Active

Active: active (running) since Tue 2024-05-14 16:37:22 CEST; 13min ago
```

Check this on both nodes. For more information on SBD configuration, see:

- <https://documentation.suse.com/sle-ha/15-SP4/single-html/SLE-HA-administration/#cha-ha-storage-protect>
- TID 7016880 and TID 7008216
- manual page `sbd(8)`, `stonith_sbd(7)`, and `cs_show_sbd_devices(8)`

### 3.3.3 Corosync cluster communication

```
akka1:~ # corosync-cfgtool -s

Printing ring status.
Local node ID 2
RING ID 0
    id      = 192.168.1.11
    status  = ring 0 active with no faults
```

Check this on both nodes. See appendix [Section 8.3, “Corosync configuration of the two-node cluster”](#) for a `corosync.conf` example. See also manual page `systemctl(1)`, `corosync.conf(5)` and `corosync-cfgtool(1)`.

### 3.3.4 systemd cluster services

```
# systemctl status pacemaker | grep Active

Active: active (running) since Tue 2024-05-14 16:37:28 CEST; 17min ago
```

Check this on both nodes. See also manual page `systemctl(1)`.

### 3.3.5 Basic Linux cluster configuration

```
# crm_mon -lr

Cluster Summary:
* Stack: corosync
* Current DC: akka1 (version 2.1.2+20211124...) - partition with quorum
* Last updated: Tue May 14 17:03:30 2024
* Last change: Mon Apr 22 15:00:58 2024 by root via cibadmin on akka2
* 2 nodes configured
* 1 resource instances configured

Node List:
* Online: [ akka1 akka2 ]

Full List of Resources:
* rsc_stonith_sbd (stonith:external/sbd): Started akka1
```

Check this on both nodes. See also manual page `crm_mon(8)`.

## 4 Checking the ControlZone setup

The ControlZone needs to be tested without the Linux cluster before integrating both. Each test needs to be done on both nodes.

### 4.1 Checking ControlZone on central NFS share

Check the mzadmin's environment variables MZ\_HOME, JAVA\_HOME, PATH. Then check the `mzsh startup/shutdown/status` functionality for MZ\_HOME on central NFS. This is needed on both nodes. Before starting the ControlZone services on one node, ensure they are not running on the other node.

```
# su - c11adm
~> echo $MZ_HOME $JAVA_HOME

/usr/sap/C11 /usr/lib64/jvm/jre-17-openjdk

~> which mzsh

/usr/sap/C11/bin/mzsh
```

```
~> echo "are you sure platform is not running on the other node?"

are you sure platform is not running on the other node?

~> mzsh startup -f platform

Starting platform...done.

~> mzsh status platform; echo $?

platform is running
0
```

```
~> mzsh startup -f ui

Starting ui...done.

~> mzsh status ui; echo $?

ui is running
0
```

```
~> mzsh shutdown ui

Shutting down ui....done.

~> mzsh status ui; echo $?

ui is not running
2
```

```
~> mzsh shutdown platform

Shutting down platform.....done.

~> mzsh status platform; echo $?

platform is not running
2
```

Perform the above steps on both nodes.

## 4.2 Checking ControlZone on each node's local disk

Check the mzadmin's environment variables MZ\_HOME, JAVA\_HOME, PATH. Then check the `mzsh status` functionality for MZ\_HOME on the local disk. This is needed on both nodes.

```
# su - c1ladm
~> export MZ_HOME="/opt/cm/C11"
~> export PATH="/opt/cm/C11/bin:$PATH"

~> echo $MZ_HOME $JAVA_HOME

/opt/cm/C11 /usr/lib64/jvm/jre-17-openjdk

~> which mzsh

/opt/cm/C11/bin/mzsh
```

```
~> mzsh status platform; echo $?

platform is running
0
```

```
~> mzsh status ui; echo $?
```

```
ui is running
0
```

Perform the above steps on both nodes. The ControlZone services should be running on either node, but not on both in parallel.

## 5 Integrating Convergent Mediation ControlZone with the Linux cluster

### 5.1 Preparing mzadmin user ~/.bashrc file

For the environment variables JAVA\_HOME, MZ\_HOME and MZ\_PLATFORM, certain values are required. For cluster actions, the values are inherited from the RA through related RA\_... variables. For manual admin actions, the values are set as default. This is needed on both nodes.

```
akka1:~# su - c1ladm
akka1:~> vi ~/.bashrc

# MZ_PLATFORM, MZ_HOME, JAVA_HOME are set by HA RA
export MZ_PLATFORM=${RA_MZ_PLATFORM:-"http://localhost:9000"}
export MZ_HOME=${RA_MZ_HOME:-"/usr/sap/C11"}
export JAVA_HOME=${RA_JAVA_HOME:-"/usr/lib64/jvm/jre-17-openjdk"}

akka1:~> scp ~/.bashrc akka2:~/
akka1:~> md5sum ~/.bashrc
...
akka1:~> ssh akka2 "md5sum ~/.bashrc"
...
```

See [Table 2, "Table Description of important resource agent parameters"](#) and manual page [ocf\\_suse\\_SAPCMControlZone\(7\)](#) for details.

### 5.2 Preparing the operating system for NFS monitoring

This is needed on both nodes.

```
akka1:~ # mkdir -p /usr/sap/C11/.check /usr/sap/.check_C11
akka1:~ # ssh akka2 "mkdir -p /usr/sap/C11/.check /usr/sap/.check_C11"
```

See manual page [ocf\\_suse\\_SAPCMControlZone\(7\)](#), [ocf\\_heartbeat\\_Fileystem\(7\)](#) and [mount\(8\)](#).

## 5.3 Adapting the cluster basic configuration

All steps to load the configuration into the Cluster Information Base (CIB) only need to be performed on one node.

### 5.3.1 Adapting cluster bootstrap options and resource defaults

The first example defines the cluster bootstrap options, the resource and operation defaults. The STONITH timeout value should be greater than 1.2 times the SBD on-disk msgwait timeout value. The priority fencing delay value should be at least twice the SBD CIB `pcmk_delay_max` value.

```
# vi crm-cib.txt

# enter the below to crm-cib.txt
property cib-bootstrap-options: \
    have-watchdog=true \
    cluster-infrastructure=corosync \
    cluster-name=hacluster \
    dc-deadtime=20 \
    stonith-enabled=true \
    stonith-timeout=150 \
    priority-fencing-delay=30 \
    stonith-action=reboot
rsc_defaults rsc-options: \
    resource-stickiness=1 \
    migration-threshold=3 \
    failure-timeout=86400
op_defaults op-options: \
    timeout=120 \
    record-pending=true
```

Load the file to the cluster.

```
# crm configure load update crm-cib.txt
```

See also manual page `crm(8)`, `sbd(8)` and `SAPCMControlZone_basic_cluster(7)`.

### 5.3.2 Adapting SBD STONITH resource

The next configuration step defines a disk-based SBD STONITH resource. Timing is adapted for priority fencing.

```
# vi crm-sbd.txt

# enter the below to crm-sbd.txt
primitive rsc_stonith_sbd stonith:external/sbd \
    params pcmk_delay_max=15
```

Load the file to the cluster.

```
# crm configure load update crm-sbd.txt
```

See also manual pages `crm(8)`, `sbd(8)`, `stonith_sbd(7)`, and `SAPCMControlZone_basic_cluster(7)`.

## 5.4 Configuring ControlZone cluster resources

### 5.4.1 Virtual IP address resource

Next, configure an IP address resource `rsc_ip_C11`. In the event of an IP address failure (or monitor timeout), the IP address resource is restarted until it is successful or the migration threshold is reached.

```
# vi crm-ip.txt

# enter the below to crm-ip.txt
primitive rsc_ip_C11 ocf:heartbeat:IPAddr2 \
    op monitor interval=60 timeout=20 on-fail=restart \
    params ip=192.168.1.112 \
    meta maintenance=true
```

Load the file to the cluster.

```
# crm configure load update crm-ip.txt
```

See also manual page `crm(8)` and `ocf_heartbeat_IPAddr2(7)`.

### 5.4.2 File system resource (only monitoring)

A shared file system might be statically mounted by the operating system on both cluster nodes. This file system holds work directories. It must not be confused with the ControlZone application itself. Client-side write caching needs to be disabled.

A file system resource `rsc_fs_C11` is configured for a bind-mount of the real NFS share. This resource is grouped with the ControlZone platform and IP address. In the event of a file system failures, the node gets fenced. No mount or umount on the real NFS share is done. An example for the real NFS share is `/usr/sap/C11/.check` , an example for the bind-mount is `/usr/sap/.check_C11` . Both mount points need to be created before the cluster resource is activated.

```
# vi crm-fs.txt

# enter the below to crm-fs.txt
primitive rsc_fs_C11 ocf:heartbeat:Filesystem \
    params device=/usr/sap/C11/.check directory=/usr/sap/.check_C11 \
    fstype=nfs4 options=bind,rw,noac,sync,defaults \
    op monitor interval=90 timeout=120 on-fail=fence \
    op_params OCF_CHECK_LEVEL=20 \
    op start timeout=120 \
    op stop timeout=120 \
    meta maintenance=true
```

Load the file to the cluster.

```
# crm configure load update crm-fs.txt
```

See also manual page `crm(8)`, `SAPCMControlZone_basic_cluster(7)`, `ocf_heartbeat_Fileystem(7)` and `nfs(5)`.

### 5.4.3 SAP Convergent Mediation ControlZone platform and UI resources

A ControlZone platform resource `rsc_cz_C11` is configured, handled by the operating system user `c11adm`. The local `/opt/cm/C11/bin/mzsh` is used for monitoring, but for other actions, the central `/usr/sap/C11/bin/mzsh` is used. In the event of a ControlZone platform failure (or monitor timeout), the platform resource is restarted until it is successful or the migration threshold is reached. If the migration threshold is reached, or if the node where the group is running fails, the group will be moved to the other node. A priority is configured for correct fencing in split-brain situations.

```
# vi crm-cz.txt

# enter the below to crm-cz.txt
primitive rsc_cz_C11 ocf:suse:SAPCMControlZone \
    params SERVICE=platform USER=c11adm \
    MZSHELL=/opt/cm/C11/bin/mzsh;/usr/sap/C11/bin/mzsh \
    MZHOME=/opt/cm/C11/;/usr/sap/C11/ \
    MZPLATFORM=http://localhost:9000 \
```

```
JAVAHOME=/usr/lib64/jvm/jre-17-openjdk \
op monitor interval=90 timeout=150 on-fail=restart \
op start timeout=300 \
op stop timeout=300 \
meta priority=100 maintenance=true
```

Load the file to the cluster.

```
# crm configure load update crm-cz.txt
```

A ControlZone UI resource `rsc_ui_C11` is configured, handled by the operating system user `c1ladm`. The local `/opt/cm/C11/bin/mzsh` is used for monitoring, but for other actions the central `/usr/sap/C11/bin/mzsh` is used. In the event of a ControlZone UI failure (or monitor timeout), the UI resource is restarted until it is successful or the migration threshold is reached. If the migration threshold is reached, or if the node where the group is running fails, the group will be moved to the other node.

```
# vi crm-ui.txt

# enter the below to crm-ui.txt
primitive rsc_ui_C11 ocf:suse:SAPCMControlZone \
  params SERVICE=ui USER=c1ladm \
  MZSHELL=/opt/cm/C11/bin/mzsh;/usr/sap/C11/bin/mzsh \
  MZHOME=/opt/cm/C11/;/usr/sap/C11/ \
  MZPLATFORM=http://localhost:9000 \
  JAVAHOME=/usr/lib64/jvm/jre-17-openjdk \
  op monitor interval=90 timeout=150 on-fail=restart \
  op start timeout=300 \
  op stop timeout=300 \
  meta maintenance=true
```

Load the file to the cluster.

```
# crm configure load update crm-ui.txt
```

Find an overview on the RA SAPCMControlZone parameters below:

TABLE 2: TABLE DESCRIPTION OF IMPORTANT RESOURCE AGENT PARAMETERS

Name	Description
USER	OS user who calls mzsh, owner of \$MZ_HOME (might be different from \$HOME). Optional. Unique, string. Default value: "mzadmin".
SERVICE	The ControlZone service to be managed by the resource agent. Optional. Unique, [ platform   ui ]. Default value: "platform".

Name	Description
MZSHELL	<p>Path to mzsh. Could be one or two full paths. If one path is given, that path is used for all actions. In case two paths are given, the first one is used for monitor actions, the second one is used for start/stop actions. If two paths are given, the first needs to be on local disk, the second needs to be on the central NFS share with the original CM ControlZone installation. Two paths are separated by a semicolon (;). The mzsh contains settings that need to be consistent with MZ_PLATFORM, MZ_HOME, JAVA_HOME. Refer to Convergent Mediation product documentation for details. Optional. Unique, string. Default value: "/opt/cm/bin/mzsh".</p>
MZHOME	<p>Path to CM ControlZone installation directory, owned by the mzadmin user. Could be one or two full paths. If one path is given, that path is used for all actions. In case two paths are given, the first one is used for monitor actions, the second one is used for start/stop actions. If two paths are given, the first needs to be on local disk, the second needs to be on the central NFS share with the original CM ControlZone installation. See also JAVAHOME. Two paths are separated by semicolon (;). Optional. Unique, string. Default value: "/opt/cm/".</p>
MZPLATFORM	<p>URL used by mzsh for connecting to CM ControlZone services. Could be one or two URLs. If one URL is given, that URL is used for all actions. In case two URLs are given, the first one is used for monitor and stop actions, the second one is used for start actions. Two URLs are separated by semicolon (;). Should usually not be changed. The service's virtual host name or virtual IP address managed by the cluster must never be used for RA monitor actions. Optional. Unique, string. Default value: "http://localhost:9000".</p>
JAVAHOME	<p>Path to Java virtual machine used for CM ControlZone. Could be one or two full paths. If one path is given, that path is used for all actions. In case two paths are given, the first one is used</p>

Name	Description
	for monitor actions, the second one is used for start/stop actions. If two paths are given, the first needs to be on local disk, the second needs to be on the central NFS share with the original CM ControlZone installation. See also MZHOME. Two paths are separated by semicolon (;). Optional. Unique, string. Default value: "/usr/lib64/jvm/jre-17-openjdk".

See also manual page `crm(8)` and `ocf_suse_SAPCMControlZone(7)`.

#### 5.4.4 CM ControlZone resource group

ControlZone platform and UI resources `rsc_cz_C11` and `rsc_ui_C11` are grouped with file system `rsc_fs_C11` and IP address resource `rsc_ip_C11` into group `grp_cz_C11`. The file system starts first, then comes the platform. The IP address starts before the UI. The resource group might run on either node, but never in parallel. If the file system resource is restarted, all resources of the group will restart as well. If the platform or IP address resource is restarted, the UI resource will restart as well.

```
# vi crm-grp.txt

# enter the below to crm-grp.txt
group grp_cz_C11 rsc_fs_C11 rsc_cz_C11 rsc_ip_C11 rsc_ui_C11 \
    meta maintenance=true
```

Load the file to the cluster.

```
# crm configure load update crm-grp.txt
```

### 5.5 Activating the cluster resources

```
# crm resource refresh grp_cz_C11
...

# crm resource maintenance grp_cz_C11 off
```

## 5.6 Checking the cluster resource configuration

```
# crm_mon -lr

Cluster Summary:
* Stack: corosync
* Current DC: akka1 (version 2.1.2+20211124...) - partition with quorum
* Last updated: Tue May 14 17:03:30 2024
* Last change: Mon Apr 22 15:00:58 2024 by root via cibadmin on akka2
* 2 nodes configured
* 5 resource instances configured

Node List:
* Online: [ akka1 akka2 ]

Full List of Resources:
* rsc_stonith_sbd (stonith:external/sbd): Started akka1
* Resource Group: grp_cz_C11:
  * rsc_fs_C11 (ocf::heartbeat:Filesystem): Started akka2
  * rsc_cz_C11 (ocf::suse:SAPCMControlZone): Started akka2
  * rsc_ip_C11 (ocf::heartbeat:IPaddr2): Started akka2
  * rsc_ui_C11 (ocf::suse:SAPCMControlZone): Started akka2
```

**Congratulations!**

The HA cluster is up and running, controlling the ControlZone resources. It is now advisable to create a backup of the cluster configuration.

```
FIRSTIME=$(date +%s)
# crm configure show > crm-all-`${FIRSTIME}.txt

# cat crm-all-`${FIRSTIME}.txt
...

# crm_report
...
```

See the appendix [Section 8.2, “CRM configuration for a typical setup”](#) for a complete CIB example.

## 5.7 Testing the HA cluster

As with any HA cluster, testing is crucial. Ensure that all test cases derived from customer expectations are executed and passed. Otherwise, the project is likely to fail in production.

- Set up a test cluster for testing configuration changes and administrative procedures before applying them on the production cluster.
- Carefully define, perform, and document tests for all scenarios that should be covered, and do the same for all maintenance procedures.
- Before performing full cluster testing, test the ControlZone features without the Linux cluster.
- Test basic Linux cluster features without ControlZone before performing full cluster testing.
- Follow general best practices, see [Section 6.1, “Dos and don’ts”](#).
- Open an additional terminal window on a node that is expected to not be fenced. In that terminal, continuously run `cs_show_cluster_actions` or similar. See manual page `cs_show_cluster_actions(8)` and `SAPCMControlZone_maintenance_examples(7)`.

The following list shows common test cases for the CM ControlZone resources managed by the HA cluster.

- [Section 5.7.1, “Manually restarting ControlZone resources in-place”](#)
- [Section 5.7.2, “Manually migrating ControlZone resources”](#)
- [Section 5.7.3, “Testing ControlZone UI restart by cluster on UI failure”](#)
- [Section 5.7.4, “Testing ControlZone restart by cluster on platform failure”](#)
- [Section 5.7.5, “Testing ControlZone takeover by cluster on node failure”](#)
- [Section 5.7.6, “Testing ControlZone takeover by cluster on NFS failure”](#)
- [Section 5.7.7, “Testing cluster reaction on network split-brain”](#)

This is not a complete list. Define additional test cases according to your needs. Some examples are listed in [Section 5.8, “Additional tests”](#). Do not forget to perform every test on each node.



## Note

Tests for the basic HA cluster and tests for bare CM ControlZone components are not covered in this document. Information about these tests can be found in the relevant product documentation.

Unless otherwise stated, the test prerequisites are always that

- both cluster nodes are booted and connected to the cluster.
- SBD and corosync are fine.
- NFS and local disks are fine.
- the ControlZone resources are all running.
- no fail counts or migration constraints are in the CIB.
- the cluster is idle, no actions are pending.

### 5.7.1 Manually restarting ControlZone resources in-place

#### COMPONENT:

- ControlZone resources

#### DESCRIPTION:

- The ControlZone resources are stopped and restarted in-place.

#### PROCEDURE:

1. Check the ControlZone resources and cluster.

```
# cs_wait_for_idle -s 5; crm_mon -lr
```

2. Stop the ControlZone resources.

```
# cs_wait_for_idle -s 5; crm resource stop grp_cz_C11
# cs_wait_for_idle -s 5; crm_mon -lr
```

3. Check the ControlZone resources.

```
# su - c1ladm -c "mzsh status"
...
# mount | grep "/usr/sap/C11"
...
# df -h /usr/sap/C11
...
```

4. Start the ControlZone resources.

```
# cs_wait_for_idle -s 5; crm resource start grp_cz_C11
# cs_wait_for_idle -s 5; crm_mon -lr
```

5. Check the ControlZone resources and cluster.

```
# cs_wait_for_idle -s 5; crm_mon -lr
```

**EXPECTED:**

1. The cluster gracefully stops all resources.
2. The file system stays mounted.
3. The cluster starts all resources.
4. No resource failure happens.

## 5.7.2 Manually migrating ControlZone resources

**COMPONENT:**

- ControlZone resources

**DESCRIPTION:**

- The ControlZone resources are stopped and then started on the other node.

**PROCEDURE:**

1. Check the ControlZone resources and cluster.

```
# cs_wait_for_idle -s 5; crm_mon -lr
```

2. Migrate the ControlZone resources.

```
# cs_wait_for_idle -s 5; crmresource move grp_cz_C11 force
# cs_wait_for_idle -s 5; crm_mon -lr
```

3. Remove migration constraint.

```
# crm resource clear grp_cz_C11
# crm configure show | grep cli-
```

4. Check the ControlZone resources and cluster.

```
# cs_wait_for_idle -s 5; crm_mon -lr
```

**EXPECTED:**

1. The cluster gracefully stops all resources.
2. The file system stays mounted.
3. The cluster starts all resources on the other node.
4. No resource failure happens.

### 5.7.3 Testing ControlZone UI restart by cluster on UI failure

**COMPONENT:**

- ControlZone resources (UI)

**DESCRIPTION:**

- The ControlZone UI is restarted on same node.

**PROCEDURE:**

1. Check the ControlZone resources and cluster.

```
# cs_wait_for_idle -s 5; crm_mon -lr
```

2. Manually kill ControlZone UI (on, for example, akka1).

```
# ssh root@akka1 "su - clladm \"mzsh kill ui\""
# cs_wait_for_idle -s 5; crm_mon -lr
```

3. Clean up fail count.

```
# crm resource cleanup grp_cz_C11
# cibadmin -Q | grep fail-count
```

4. Check the ControlZone resources and cluster.

```
# cs_wait_for_idle -s 5; crm_mon -lr
```

**EXPECTED:**

1. The cluster detects failed resources.

2. The file system stays mounted.
3. The cluster restarts the UI on the same node.
4. One resource failure happens.

## 5.7.4 Testing ControlZone restart by cluster on platform failure

### COMPONENT:

- ControlZone resources (platform)

### DESCRIPTION:

- The ControlZone resources are stopped and restarted on same node.

### PROCEDURE:

1. Check the ControlZone resources and cluster.

```
# cs_wait_for_idle -s 5; crm_mon -lr
```

2. Manually kill ControlZone platform (on, for example, akka1).

```
# ssh root@akka1 "su - clladm \"mzsh kill platform\""
# cs_wait_for_idle -s 5; crm_mon -lr
```

3. Clean up fail count.

```
# crm resource cleanup grp_cz_C11
# cibadmin -Q | grep fail-count
```

4. Check the ControlZone resources and cluster.

```
# cs_wait_for_idle -s 5; crm_mon -lr
```

### EXPECTED:

1. The cluster detects failed resources.
2. The file system stays mounted.
3. The cluster restarts resources on the same node.
4. One resource failure happens.

## 5.7.5 Testing ControlZone takeover by cluster on node failure

### COMPONENT:

- Cluster node

### DESCRIPTION:

- The ControlZone resources are started on other node

### PROCEDURE:

1. Check the ControlZone resources and cluster.

```
akka2:~ # cs_wait_for_idle -s 5; crm_mon -1r
```

2. Manually kill cluster node, where resources are running (for example akka1).

```
akka2:~ # ssh root@akka1 "systemctl reboot --force"  
akka2:~ # cs_wait_for_idle -s 5; crm_mon -1r
```

3. Rejoin fenced node (for example akka1) to cluster.

```
akka2:~ # cs_show_sbd_devices | grep reset  
akka2:~ # cs_clear_sbd_devices --all  
akka2:~ # crm cluster start --all
```

4. Check the ControlZone resources and cluster.

```
akka2:~ # cs_wait_for_idle -s 5; crm_mon -1r
```

### EXPECTED:

1. The cluster detects a failed node.
2. The cluster fences a failed node.
3. The cluster starts all resources on the other node.
4. The fenced node needs to be connected to the cluster.
5. No resource failure happens.

## 5.7.6 Testing ControlZone takeover by cluster on NFS failure

### COMPONENT:

- Network (for NFS)

#### DESCRIPTION:

- The NFS share fails on one node and the cluster moves resources to other node.

#### PROCEDURE:

1. Check the ControlZone resources and cluster.

```
akka2:~ # cs_wait_for_idle -s 5; crm_mon -lr
```

2. Manually block port for NFS, where resources are running (for example akka1).

```
akka2:~ # ssh root@akka1 "iptables -I INPUT -p tcp -m multiport --ports 2049 -j DROP"
akka2:~ # ssh root@akka1 "iptables -L | grep 2049"
akka2:~ # cs_wait_for_idle -s 5; crm_mon -lr
```

3. Rejoin fenced node (for example akka1) to cluster.

```
akka2:~ # cs_show_sbd_devices | grep reset
akka2:~ # cs_clear_sbd_devices --all
akka2:~ # crm cluster start --all
```

4. Check the ControlZone resources and cluster.

```
akka2:~ # cs_wait_for_idle -s 5; crm_mon -lr
```

#### EXPECTED:

1. The cluster detects failed NFS.
2. The cluster fences the node.
3. The cluster starts all resources on the other node.
4. The fenced node needs to be connected to the cluster.
5. Resource failure happens.

### 5.7.7 Testing cluster reaction on network split-brain

#### COMPONENT:

- Network (for Corosync)

#### DESCRIPTION:

- The network fails, node without resources gets fenced, resources keep running.

#### PROCEDURE:

1. Check the ControlZone resources and cluster.

```
akka2:~ # cs_wait_for_idle -s 5; crm_mon -lr
```

2. Manually block ports for Corosync.

```
akka2:~ # grep mcastport /etc/corosync/corosync.conf
akka2:~ # ssh root@akka1 "iptables -I INPUT -p udp -m multiport --ports 5405,5407 -j DROP"
akka2:~ # ssh root@akka1 "iptables -L | grep -e 5405 -e 5407"
akka2:~ # cs_wait_for_idle -s 5; crm_mon -lr
```

3. Rejoin fenced node (for example akka1) to cluster.

```
akka2:~ # cs_show_sbd_devices | grep reset
akka2:~ # cs_clear_sbd_devices --all
akka2:~ # crm cluster start --all
```

4. Check the ControlZone resources and cluster.

```
akka2:~ # cs_wait_for_idle -s 5; crm_mon -lr
```

#### EXPECTED:

1. The cluster detects failed Corosync.
2. The cluster fences the node.
3. The cluster keeps all resources on the same node.
4. The fenced node needs to be connected to the cluster.
5. No resource failure happens.

## 5.8 Additional tests

Define additional test cases according to your needs. Some test cases you should test are listed below.

- Remove virtual IP address.
- Stop and restart passive node.
- Stop and restart in parallel all cluster nodes.
- Isolate the SBD.
- Maintenance procedure with cluster is continuously running, but application restarts.
- Maintenance procedure with cluster restarts, but application is running.
- Kill the Corosync process of one cluster node.

See also manual page `crm(8)` for cluster `crash_test`.

## 6 Administration

HA clusters are complex, and the CM ControlZone is also complex. Deploying and running HA clusters for CM ControlZonen needs preparation, caution and care. Fortunately, most of the pitfalls and many best practices are already known. This chapter describes general administrative tasks.

### 6.1 Dos and don'ts

The following basic rules will help you avoid known issues:

- Carefully test all configuration changes and administrative procedures on the test cluster before applying them to the production cluster.
- Before taking any action, always check the Linux cluster's idle status, remaining migration constraints, and resource failures, plus the ControlZone status. See [Section 6.2, "Showing status of ControlZone resources and HA cluster"](#).
- Be patient. The Linux cluster requires a certain amount of time to detect the overall status of the ControlZone, depending on the ControlZone services and the configured intervals and timeouts.
- As long as the ControlZone components are managed by the Linux cluster, they may never be started/stopped/moved from outside. This means that no manual intervention is required.

See also the manual page `SAPCMControlZone_maintenance_examples(7)`, `SAPCMControlZone_basic_cluster(7)`, and `ocf_suse_SAPCMControlZone(7)`.

## 6.2 Showing status of ControlZone resources and HA cluster

Perform the following steps each time before and after you do any work on the cluster.

```
# su - clladm -c "mzsh status"
# crm_mon -lr
# crm configure show | grep cli-
# cibadmin -Q | grep fail-count
# cs_clusterstate -i
```

See also manual page `SAPCMControlZone_maintenance_examples(7)`, `crm_mon(8)`, `cs_clusterstate(8)`, and `cs_show_cluster_actions(8)`.

## 6.3 Watching ControlZone resources and HA cluster

During testing and maintenance, you can run the following command to view near real-time status changes.

```
# watch -s8 cs_show_cluster_actions
```

See also manual page `SAPCMControlZone_maintenance_examples(7)`, `crm_mon(8)`, `cs_clusterstate(8)`, and `cs_show_cluster_actions(8)`.

## 6.4 Starting the ControlZone resources

Use the cluster for starting the resources.

```
# crm_mon -lr
# cs_wait_for_idle -s 5; crm resource start grp_cz_C11
# cs_wait_for_idle -s 5; crm_mon -lr
```

See also manual page `SAPCMControlZone_maintenance_examples(7)`, `crm(8)`.

## 6.5 Stopping the ControlZone resources

Use the cluster for stopping the resources.

```
# crm_mon -lr
# cs_wait_for_idle -s 5; crm resource stop grp_cz_C11
# cs_wait_for_idle -s 5; crm_mon -lr
```

See also manual page [SAPCMControlZone\\_maintenance\\_examples\(7\)](#), [crm\(8\)](#).

## 6.6 Migrating the ControlZone resources

To migrate the ControlZone resources to another node, the following steps are performed (see commands below): ControlZone application and Linux cluster are initially checked for a clean and idle state. The ControlZone resources are moved to the other node. The associated location rule will be removed after the takeover took place. Finally, ControlZone application and HA cluster are again checked for a clean and idle state.

```
# su - clladm -c "mzsh status"
# crm_mon -lr
# crm configure show | grep cli-
# cibadmin -Q | grep fail-count
# cs_clusterstate -i

# crm resource move grp_cz_C11 force
# cs_wait_for_idle -s 5; crm_mon -lr
# crm resource clear grp_cz_C11

# cs_wait_for_idle -s 5; crm_mon -lr
# crm configure show | grep cli-
# su - clladm -c "mzsh status"
```

See also manual page [SAPCMControlZone\\_maintenance\\_examples\(7\)](#).

## 6.7 Example for generic maintenance procedure

Find below a generic procedure, mainly for maintenance of the ControlZone components. The resources are temporarily removed from cluster control. The Linux cluster remains active.

The individual steps are carried out as follows (see commands below): ControlZone application and HA cluster are initially checked for a clean and idle state. The ControlZone resource group is set to maintenance mode. This is required to enable manual actions on the resources. After the manual actions are completed, the resource group is placed back under cluster control. It is necessary to wait for the completion of each step and to check the results. Finally, ControlZone application and HA cluster are again checked for a clean and idle state.

```

# su - clladm -c "mzsh status"
# crm_mon -lr
# crm configure show | grep cli-
# cibadmin -Q | grep fail-count
# cs_clusterstate -i
# crm resource maintenance grp_cz_C11

# echo "PLEASE DO MAINTENANCE NOW"

# crm resource refresh grp_cz_C11
# cs_wait_for_idle -s 5; crm_mon -lr
# crm resource maintenance grp_cz_C11 off
# cs_wait_for_idle -s 5; crm_mon -lr
# su - clladm -c "mzsh status"

```

See also manual page [SAPCMControlZone\\_maintenance\\_examples\(7\)](#).

## 6.8 Showing resource agent log messages

Failed RA actions on a node are displayed in the current message file.

```
# grep "SAPCMControlZone.*rc=[1-7,9]" /var/log/messages
```

See also manual page [ocf\\_suse\\_SAPCMControlZone\(7\)](#).

## 6.9 Cleaning up resource fail count

Cleaning up resource fail count can be done after the cluster has recovered the resource from a failure.

```

# crm resource cleanup grp_cz_C11
# cibadmin -Q | grep fail-count

```

See also manual page [ocf\\_suse\\_SAPCMControlZone\(7\)](#) and [SAPCMControlZone\\_maintenance\\_examples\(7\)](#).

# 7 References

For more information, see the documents listed below.

## 7.1 Pacemaker

- Pacemaker documentation online: <https://clusterlabs.org/pacemaker/doc/> 

## 7.2 Related Manual Pages

- chronyc(8)
- chrony.conf(5)
- corosync.conf(8)
- corosync-cfgtool(8)
- corosync\_overview(8)
- cibadmin(8)
- crm(8)
- crm\_mon(8)
- crm\_report(8)
- crm\_simulate(8)
- cs\_clusterstate(8)
- cs\_man2pdf(8)
- cs\_show\_cluster\_actions(8)
- cs\_show\_sbd\_devices(8)
- cs\_wait\_for\_idle(8)
- fstab(5)
- ha\_related\_sap\_notes(7)
- ha\_related\_suse\_tids(7)
- hosts(5)
- mount.nfs(8)
- nfs(5)

- `ocf_heartbeat_Fileystem(7)`
- `ocf_heartbeat_IPAddr2(7)`
- `ocf_heartbeat_ping(7)`
- `ocf_suse_SAPCMControlZone(7)`
- `passwd(5)`
- `SAPCMControlZone_basic_cluster(7)`
- `SACMPCControlZone_maintenance_procedures(7)`
- `saptune(8)`
- `sbd(8)`
- `stonith_sbd(7)`
- `supportconfig(8)`
- `systemctl(8)`
- `systemd-cgls(8)`
- `usermod(8)`
- `votequorum(5)`
- `zypper(8)`

### 7.3 Related SUSE TIDs

- Diagnostic Data Collection Master TID (<https://www.suse.com/support/kb/doc/?id=000019514>)
- How to enable cluster resource tracing (<https://www.suse.com/support/kb/doc/?id=000019138>)
- NFS file system is hung. New mount attempts hang also. (<https://www.suse.com/support/kb/doc/?id=000019722>)
- An NFS client hangs on various operations, including "df". Hard vs Soft NFS mounts. (<https://www.suse.com/support/kb/doc/?id=000020830>)

## 7.4 Related SUSE Documentation

- SUSE Linux Enterprise Server for SAP Applications (<https://documentation.suse.com/sles-sap/>)
- SUSE Linux Enterprise High Availability (<https://documentation.suse.com/sle-ha>)

## 7.5 Related Digital Route Documentation

- ControlZone tool mzsh (<https://infozone.atlassian.net/wiki/spaces/MD9/pages/4881672/mzsh>)
- ControlZone requirements (<https://infozone.atlassian.net/wiki/spaces/MD9/pages/4849685/System+Requirements>).
- ControlZone installation <https://infozone.atlassian.net/wiki/spaces/MD9/pages/4849683/Installation+Instructions>

## 7.6 Related SAP Documentation

- SAP RISE (<https://www.sap.com/products/erp/rise.html>)
- SAP BRIM Convergent Mediation (<https://www.sap.com/products/financial-management/convergent-mediation.html>)
- SAP Product Availability Matrix ([https://support.sap.com/en/release-upgrade-maintenance.html#section\\_1969201630](https://support.sap.com/en/release-upgrade-maintenance.html#section_1969201630))

## 7.7 Related SAP Notes

- 1552925 - Linux: High Availability Cluster Solutions (<https://launchpad.support.sap.com/#/notes/1552925>)
- 1763512 - Support details for SUSE Linux Enterprise for SAP Applications (<https://launchpad.support.sap.com/#/notes/1763512>)
- 2369910 - SAP Software on Linux: General information (<https://launchpad.support.sap.com/#/notes/2369910>)

- 2578899 - SUSE Linux Enterprise Server 15: Installation Note (<https://launchpad.support.sap.com/#/notes/2578899>)
- 3079845 - Standard Practices for SAP CM High Availability (<https://launchpad.support.sap.com/#/notes/3079845>)

## 8 Appendix

### 8.1 The mzadmin user's ~/.bashrc file

Find below a typical mzadmin user's ~/.bashrc file.

```
akka1:~ # su - c1ladm -c "cat ~/.bashrc"

# MZ_PLATFORM, MZ_HOME, JAVA_HOME are set by HA RA
export MZ_PLATFORM=${RA_MZ_PLATFORM:-"http://localhost:9000"}
export MZ_HOME=${RA_MZ_HOME:-"/usr/sap/C11"}
export JAVA_HOME=${RA_JAVA_HOME:-"/usr/lib64/jvm/jre-17-openjdk"}
```

### 8.2 CRM configuration for a typical setup

Find below a typical CRM configuration for a CM ControlZone instance, with a dummy file system, platform and UI services and related IP address.

```
akka1:~ # crm configure show

node 1: akka1
node 2: akka2
#
primitive rsc_fs_C11 ocf:heartbeat:Filesystem \
    params device=/usr/sap/C11/.check directory=/usr/sap/.check_C11 \
    fstype=nfs4 options=bind,rw,noac,sync,defaults \
    op monitor interval=90 timeout=120 on-fail=fence \
    op_params OCF_CHECK_LEVEL=20 \
    op start timeout=120 interval=0 \
    op stop timeout=120 interval=0
#
primitive rsc_cz_C11 ocf:suse:SAPCMControlZone \
    params SERVICE=platform USER=c1ladm \
    MZSHELL=/opt/cm/C11/bin/mzsh;/usr/sap/C11/bin/mzsh \
    MZHOME=/opt/cm/C11/;/usr/sap/C11/ \
    MZPLATFORM=http://localhost:9000 \
    JAVAHOME=/usr/lib64/jvm/jre-17-openjdk \
    op monitor interval=90 timeout=150 on-fail=restart \
    op start timeout=300 interval=0 \
    op stop timeout=300 interval=0 \
    meta priority=100
#
```

```

primitive rsc_ui_C11 ocf:suse:SAPCMControlZone \
  params SERVICE=ui USER=c1ladm \
  MZSHELL=/opt/cm/C11/bin/mzsh;/usr/sap/C11/bin/mzsh \
  MZHOME=/opt/cm/C11/;/usr/sap/C11/ \
  MZPLATFORM=http://localhost:9000 \
  JAVAHOME=/usr/lib64/jvm/jre-17-openjdk \
  op monitor interval=90 timeout=150 on-fail=restart \
  op start timeout=300 interval=0 \
  op stop timeout=300 interval=0
#
primitive rsc_ip_C11 IPAddr2 \
  params ip=192.168.1.112 \
  op monitor interval=60 timeout=20 on-fail=restart
#
primitive rsc_stonith_sbd stonith:external/sbd \
  params pcmk_delay_max=15
#
group grp_cz_C11 rsc_fs_C11 rsc_cz_C11 rsc_ip_C11 rsc_ui_C11
#
property cib-bootstrap-options: \
  have-watchdog=true \
  dc-version="2.1.2+20211124.ada5c3b36-150400.2.43-2.1.2+20211124.ada5c3b36" \
  cluster-infrastructure=corosync \
  cluster-name=hacluster \
  dc-deadtime=20 \
  stonith-enabled=true \
  stonith-timeout=150 \
  stonith-action=reboot \
  last-lrm-refresh=1704707877 \
  priority-fencing-delay=30
rsc_defaults rsc-options: \
  resource-stickiness=1 \
  migration-threshold=3 \
  failure-timeout=86400
op_defaults op-options: \
  timeout=120 \
  record-pending=true
#

```

### 8.3 Corosync configuration of the two-node cluster

Find below the Corosync configuration for one Corosync ring. Ideally two rings would be used.

```

akkal:~ # cat /etc/corosync/corosync.conf

```

```

# Read the corosync.conf.5 manual page
totem {
    version: 2
    secauth: on
    crypto_hash: sha1
    crypto_cipher: aes256
    cluster_name: hacluster
    clear_node_high_bit: yes
    token: 5000
    token_retransmits_before_loss_const: 10
    join: 60
    consensus: 6000
    max_messages: 20
    interface {
        ringnumber: 0
        mcastport: 5405
        ttl: 1
    }
    transport: udpu
}

logging {
    fileline: off
    to_stderr: no
    to_logfile: no
    logfile: /var/log/cluster/corosync.log
    to_syslog: yes
    debug: off
    timestamp: on
    logger_subsys {
        subsys: QUORUM
        debug: off
    }
}

nodelist {
    node {
        ring0_addr: 192.168.1.11
        nodeid: 1
    }
    node {
        ring0_addr: 192.168.1.12
        nodeid: 2
    }
}

quorum {

```

```
# Enable and configure quorum subsystem (default: off)
# see also corosync.conf.5 and votequorum.5
provider: corosync_votequorum
expected_votes: 2
two_node: 1
}
```

## 9 Legal notice

Copyright © 2006–2025 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled "GNU Free Documentation License".

SUSE, the SUSE logo and YaST are registered trademarks of SUSE LLC in the United States and other countries. For SUSE trademarks, see <https://www.suse.com/company/legal/>.

Linux is a registered trademark of Linus Torvalds. All other names or trademarks mentioned in this document may be trademarks or registered trademarks of their respective owners.

Documents published as part of the SUSE Best Practices series have been contributed voluntarily by SUSE employees and third parties. They are meant to serve as examples of how particular actions can be performed. They have been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. SUSE cannot verify that actions described in these documents do what is claimed or whether actions described have unintended consequences. SUSE LLC, its affiliates, the authors, and the translators may not be held liable for possible errors or the consequences thereof.

Below we draw your attention to the license under which the articles are published.

## 10 GNU Free Documentation License

Copyright © 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

### 0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

### 1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition. The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

## 2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

## 3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

## 4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.

- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

## 5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

## 6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

## 7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

## 8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all

Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

## 9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

## 10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>. Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

## ADDENDUM: How to use this License for your documents

Copyright (c) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2

or any later version published by the Free Software Foundation;  
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.  
A copy of the license is included in the section entitled “GNU  
Free Documentation License”.

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “ with... Texts.” line with this:

with the Invariant Sections being LIST THEIR TITLES, with the  
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.