Changing the Configuration of SBD

SBD provides a node fencing mechanism without using an external power-off device. Node fencing protects the cluster from data corruption by resetting failed nodes. You might need to change the cluster's SBD configuration for various reasons, such as increasing resilience, using custom settings instead of the defaults, or switching to a different fencing mechanism. This article explains how to change a High Availability cluster's SBD configuration when the SBD service is already running.

Prerequisites

- An existing SUSE Linux Enterprise High Availability cluster
- SBD already configured and running
- A hardware watchdog device on all cluster nodes
- Shared storage accessible from all nodes (if using disk-based SBD)

Important: Maintenance window

Each task in this article only takes a few minutes. These procedures require putting the cluster into maintenance mode and restarting the cluster services. In this scenario, DLM and DLM-dependent resources require downtime, but most other resources can continue running. However, the resources won't be highly available while the cluster is in maintenance mode.

Publication Date: 19 Nov 2026

Contents

1 What is SBD? 3

2 Checking and fixing the SBD configuration 4

3 Changing the SBD timeout settings 6

4 Changing the SBD watchdog device 10

5 Changing diskless SBD to disk-based SBD 12

6 Changing disk-based SBD to diskless SBD 14

7 Adding another SBD device 16

8 Replacing an existing SBD device with a new device 18

9 Removing an SBD device 21

10 Removing all SBD configuration 23

11 Legal Notice 25

A GNU Free Documentation License 25

HA glossary 33

1 What is SBD?

SBD (STONITH Block Device) provides a node fencing mechanism without using an external power-off device. The software component (the SBD daemon) works together with a watchdog device to ensure that misbehaving nodes are fenced. SBD can be used in disk-based mode with shared block storage, or in diskless mode using only the watchdog.

Disk-based SBD uses shared block storage to exchange fencing messages between the nodes. It can be used with one to three devices. One device is appropriate for simple cluster setups, but two or three devices are recommended for more complex setups or critical workloads.

Diskless SBD fences nodes by using only the watchdog, without a shared storage device. A node is fenced if it loses quorum, if any monitored daemon is lost and cannot be recovered, or if Pacemaker determines that the node requires fencing.

1.1 Components

SBD daemon

The SBD daemon starts on each node before the rest of the cluster stack and stops in the reverse order. This ensures that cluster resources are never active without SBD supervision.

SBD device (disk-based SBD)

A small logical unit (or a small partition on a logical unit) is formatted for use with SBD. A message layout is created on the device with slots for up to 255 nodes.

Messages (disk-based SBD)

The message layout on the SBD device is used to send fencing messages to nodes. The SBD daemon on each node monitors the message slot and immediately complies with any requests. To avoid becoming disconnected from fencing messages, the SBD daemon also fences the node if it loses its connection to the SBD device.

Watchdog

SBD needs a watchdog on each node to ensure that misbehaving nodes are really stopped. SBD “feeds” the watchdog by regularly writing a service pulse to it. If SBD stops feeding the watchdog, the hardware enforces a system restart. This protects against failures of the SBD process itself, such as becoming stuck on an I/O error.

1.2 Limitations and recommendations

Disk-based SBD

- The shared storage can be Fibre Channel (FC), Fibre Channel over Ethernet (FCoE), or iSCSI.
- The shared storage must *not* use host-based RAID, LVM, Cluster MD, or DRBD.
- Using storage-based RAID and multipathing is recommended for increased reliability.
- If a shared storage device has different `/dev/sdX` names on different nodes, SBD communication will fail. To avoid this, always use stable device names, such as `/dev/disk/by-id/DEVICE_ID`.
- An SBD device can be shared between different clusters, up to a limit of 255 nodes.
- When using more than one SBD device, all devices must have the same configuration.

Diskless SBD

- Diskless SBD cannot handle a split-brain scenario for a two-node cluster. This configuration should only be used for clusters with more than two nodes, or in combination with QDevice to help handle split-brain scenarios.

1.3 For more information

For more information, see the man page `sbd` or run the **`crm sbd help`** command.

2 Checking and fixing the SBD configuration

SBD configuration is complex, involving many interrelated settings and timeouts. Even though `crmsh` helps to keep these settings aligned, upgrades or manual changes can sometimes introduce misconfigurations. The **`crm sbd configure show`** and **`crm sbd status`** commands automatically check the SBD configuration and list any issues. You can also check the configuration directly using the **`crm cluster health sbd`** command. This command lists issues with the SBD fence agent, configuration files, timeout settings, and more. It can also fix some of these issues automatically.



Important: Cluster restart might be required

Certain fixes, such as timeout changes, require restarting the cluster services. The SBD health check with the `--fix` option restarts the cluster services automatically if required. If any non-`fence` resources are running, we recommend putting the cluster into maintenance mode before running the `--fix` script. This allows the services managed by the resources to keep running if the cluster restarts. However, the resources won't be highly available while in maintenance mode.



Warning: Stop DLM before putting the cluster into maintenance mode

To put the cluster into maintenance mode and then stop or restart the cluster services, you must first stop DLM and any resources that depend on it, such as GFS2 or `lvmlockd`. When the cluster is in maintenance mode, DLM keeps running even if the cluster services are offline. However, DLM depends on Corosync, so if Corosync stops but DLM is still running, the DLM instances on each node consider the other instances lost and request node fencing.

PREREQUISITES

- SBD is already configured and running.

Perform this procedure on only one node in the cluster:

1. Log in either as the `root` user or as a user with **`sudo`** privileges.
2. Run the SBD health check:

   ```
   > sudo crm cluster health sbd
   ```

   The output shows a list of warnings for any issues and reports the status of the SBD timeout configuration (either `OK` or `FAIL`). If `crmsh` can fix any issues automatically, you'll also see the following output:

   ```
   INFO: Please run "crm cluster health sbd --fix" to fix the above warning on the
    running cluster
   ```

3. If you need to fix any issues, either automatically or manually, we recommend putting the cluster into maintenance mode in case a cluster restart is required. First, check whether any resources are running:

   ```
   > sudo crm status
   ```

4. If DLM and any DLM-dependent resources are running, stop them first. If the resources are in a group, you can just stop the group. If the resources aren't in a group, stop the DLM-dependent resources first and then stop DLM. You can stop multiple resources at once:

```
> sudo crm --wait resource stop RESOURCE1 RESOURCE2
```

   Use `--wait` to make sure the resources have really stopped before the command returns.

5. If any non-`fence` resources are running, put the cluster into maintenance mode:

```
> sudo crm maintenance on
```

   In this state, the cluster stops monitoring all resources. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while the cluster is in maintenance mode.

6. Run the SBD health check with the `--fix` option to automatically fix issues:

```
> sudo crm cluster health sbd --fix
```

   The script fixes the configuration issues and restarts the cluster services if required.

7. If the output still lists issues even after using the `--fix` option, you must fix these issues manually.

8. If the cluster is still in maintenance mode, put it back into normal operation:

```
> sudo crm maintenance off
```

9. If you stopped DLM and any DLM-dependent resources, you can start them again now. If the resources are in a group, you can just start the group. If the resources aren't in a group, start DLM first and then start the DLM-dependent resources. You can start multiple resources at once:

```
> sudo crm resource start RESOURCE1 RESOURCE2
```

3 Changing the SBD timeout settings

SBD relies on multiple different timeout settings to manage node fencing. When you configure SBD using the CRM Shell, these timeouts are automatically calculated and adjusted. The automatic values are sufficient for most use cases, but if you need to change them, you can use the **`crm sbd configure`** command.



Important: Cluster restart required

In this procedure, the setup script has to restart the cluster services. If any non-`fence` resources are running, put the cluster into maintenance mode before running the script. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while the cluster is in maintenance mode.



Warning: Stop DLM before putting the cluster into maintenance mode

To put the cluster into maintenance mode and then stop or restart the cluster services, you must first stop DLM and any resources that depend on it, such as GFS2 or `lvmlockd`. When the cluster is in maintenance mode, DLM keeps running even if the cluster services are offline. However, DLM depends on Corosync, so if Corosync stops but DLM is still running, the DLM instances on each node consider the other instances lost and request node fencing.

PREREQUISITES

- SBD is already configured and running.

Perform this procedure on only one node in the cluster:

1. Log in either as the `root` user or as a user with **`sudo`** privileges.
2. Check whether any resources are running:

   ```
   > sudo crm status
   ```

3. If DLM and any DLM-dependent resources are running, stop them first. If the resources are in a group, you can just stop the group. If the resources aren't in a group, stop the DLM-dependent resources first and then stop DLM. You can stop multiple resources at once:

   ```
   > sudo crm --wait resource stop RESOURCE1 RESOURCE2
   ```

   Use `--wait` to make sure the resources have really stopped before the command returns.

4. If any non-`fence` resources are running, put the cluster into maintenance mode:

   ```
   > sudo crm maintenance on
   ```

In this state, the cluster stops monitoring all resources. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while the cluster is in maintenance mode.

5. Check the current timeout settings:

```
> sudo crm sbd configure show
```

6. Change one or both of the following timeout values as needed:

```
> sudo crm sbd configure \
  watchdog-timeout=INTEGER_IN_SECONDS \(1)
  msgwait-timeout=INTEGER_IN_SECONDS \(2)
```

If you change one timeout, the other timeout is automatically adjusted so that the `msgwait-timeout` is twice the `watchdog-timeout`. You only need to change both timeouts manually if you want the `msgwait-timeout` to be *more* than double the `watchdog-timeout`. If you try to make the `msgwait-timeout` *less* than double the `watchdog-timeout`, the command fails with a warning.

(1) The `watchdog-timeout` defines how long the watchdog waits for a response from SBD before fencing the node. Diskless SBD reads this timeout from `/etc/sysconfig/sbd`, but disk-based SBD reads it from the device metadata, which takes precedence over the settings in `/etc/sysconfig/sbd`.
For disk-based SBD on a multipath setup, this timeout must be longer than the `max_polling_interval` in `/etc/multipath.conf`, to allow enough time to detect a path failure and switch to the next path.

(2) *Only used for disk-based SBD.* After the `msgwait-timeout` is reached, SBD assumes that a message written to the node's slot on the SBD device was delivered successfully. The timeout must be long enough for the node to detect that it needs to self-fence.
When you increase this timeout, the script also automatically adjusts the `SBD_DELAY_START` setting. This helps to avoid a situation where a node reboots too quickly and rejoins the cluster before the fencing action is considered complete, which can cause a split-brain scenario.



You should not need to change the `allocate-timeout` or the `loop-timeout`.

The script automatically adjusts any other related timeouts in the cluster, then restarts the cluster services.

7. If the cluster is still in maintenance mode, put it back into normal operation:

```
> sudo crm maintenance off
```

8. If you stopped DLM and any DLM-dependent resources, you can start them again now. If the resources are in a group, you can just start the group. If the resources aren't in a group, start DLM first and then start the DLM-dependent resources. You can start multiple resources at once:

```
> sudo crm resource start RESOURCE1 RESOURCE2
```

9. When you change a timeout with **`crm sbd configure`**, the global fencing timeouts are also adjusted automatically. The automatic values are sufficient for most use cases, but if you need to change them, you can use the **`crm configure property`** command:

```
> sudo crm configure property fencing-timeout=INTEGER_IN_SECONDS (1)
> sudo crm configure property fencing-watchdog-timeout=INTEGER_IN_SECONDS (2)
```

   This command does *not* automatically adjust any other timeouts, and these settings might be overwritten if you change the SBD configuration again.

   (1) The `fencing-timeout` defines how long to wait for the fencing action to complete.

   (2) *Only used for diskless SBD.* The `fencing-watchdog-timeout` defines how long to wait for the watchdog to fence the node.

10. Confirm that the timeout settings changed:

```
> sudo crm sbd configure show
```

    This command also lists any configuration issues and tells you how to fix them.

If you need to manually calculate any timeouts, you can use these basic formulas for most use cases:

Disk-based SBD

```
msgwait-timeout >= (watchdog-timeout * 2)
fencing-timeout >= msgwait-timeout + 20%
```

Diskless SBD

```
fencing-watchdog-timeout >= (watchdog-timeout * 2)
fencing-timeout >= fencing-watchdog-timeout + 20%
```

For more information, run the **`crm help TimeoutFormulas`** command.

4 Changing the SBD watchdog device

Use this procedure to change the watchdog device that SBD uses. This can be useful if your system has multiple hardware watchdogs available or if you need to switch from the software watchdog to a hardware watchdog.



Important: Cluster restart required

In this procedure, the setup script has to restart the cluster services. If any non-`fence` resources are running, put the cluster into maintenance mode before running the script. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while the cluster is in maintenance mode.



Warning: Stop DLM before putting the cluster into maintenance mode

To put the cluster into maintenance mode and then stop or restart the cluster services, you must first stop DLM and any resources that depend on it, such as GFS2 or `lvmlockd`. When the cluster is in maintenance mode, DLM keeps running even if the cluster services are offline. However, DLM depends on Corosync, so if Corosync stops but DLM is still running, the DLM instances on each node consider the other instances lost and request node fencing.

PREREQUISITES

- SBD is already configured and running.
- All nodes have the new watchdog device available.

Perform this procedure on only one node in the cluster:

1. Log in either as the `root` user or as a user with **`sudo`** privileges.
2. Check whether any resources are running:

```
> sudo crm status
```

3. If DLM and any DLM-dependent resources are running, stop them first. If the resources are in a group, you can just stop the group. If the resources aren't in a group, stop the DLM-dependent resources first and then stop DLM. You can stop multiple resources at once:

```
> sudo crm --wait resource stop RESOURCE1 RESOURCE2
```

   Use `--wait` to make sure the resources have really stopped before the command returns.

4. If any non-`fence` resources are running, put the cluster into maintenance mode:

```
> sudo crm maintenance on
```

   In this state, the cluster stops monitoring all resources. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while the cluster is in maintenance mode.

5. Check the available watchdog devices:

```
> sudo sbd query-watchdog
```

   The output shows which watchdog is being used by SBD and which other watchdogs are available.

6. Change the SBD watchdog device, specifying either the device name (for example, `/dev/watchdog1`) or the driver name (for example, `iTCO_wdt`):

```
> sudo crm sbd configure watchdog-device=WATCHDOG
```

   The script updates the SBD configuration file with the new watchdog device and restarts the cluster services.

7. If the cluster is still in maintenance mode, put it back into normal operation:

```
> sudo crm maintenance off
```

8. If you stopped DLM and any DLM-dependent resources, you can start them again now. If the resources are in a group, you can just start the group. If the resources aren't in a group, start DLM first and then start the DLM-dependent resources. You can start multiple resources at once:

```
> sudo crm resource start RESOURCE1 RESOURCE2
```

9. Check the watchdog devices again:

```
> sudo sbd query-watchdog
```

The output should now show the new watchdog device being used by SBD.

10. Check that the correct watchdog is configured on all nodes:

```
> sudo crm sbd status
```

In the section `Watchdog info`, the watchdog device and driver should be the same on every node.

5 Changing diskless SBD to disk-based SBD

Use this procedure to change diskless SBD to disk-based SBD.

Important: Cluster restart required

In this procedure, the setup script automatically puts the cluster into maintenance mode and restarts the cluster services. In maintenance mode, the cluster stops monitoring all resources. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while in maintenance mode.

Warning: Stop DLM before putting the cluster into maintenance mode

To put the cluster into maintenance mode and then stop or restart the cluster services, you must first stop DLM and any resources that depend on it, such as GFS2 or `lvmlockd`. When the cluster is in maintenance mode, DLM keeps running even if the cluster services are offline. However, DLM depends on Corosync, so if Corosync stops but DLM is still running, the DLM instances on each node consider the other instances lost and request node fencing.

Warning: Overwriting existing data

Make sure any device you want to use for SBD does not hold any important data. Configuring a device for use with SBD overwrites the existing data.

PREREQUISITES

- Diskless SBD is configured and running.

- All nodes can access shared storage.
- The path to the shared storage device is consistent across all nodes. Use stable device names such as `/dev/disk/by-id/DEVICE_ID`.

Perform this procedure on only one node in the cluster:

1. Log in either as the `root` user or as a user with **`sudo`** privileges.
2. Check whether any resources are running:

   ```
   > sudo crm status
   ```

3. If DLM and any DLM-dependent resources are running, stop them first. If the resources are in a group, you can just stop the group. If the resources aren't in a group, stop the DLM-dependent resources first and then stop DLM. You can stop multiple resources at once:

   ```
   > sudo crm --wait resource stop RESOURCE1 RESOURCE2
   ```

   Use `--wait` to make sure the resources have really stopped before the command returns.

4. Check the status of SBD:

   ```
   > sudo crm sbd status
   # Type of SBD:
   Diskless SBD configured
   ```

5. Configure disk-based SBD. Use `--force` (or `-F`) to allow you to reconfigure SBD even when it is already running, and `--sbd-device` (or `-s`) to specify the shared storage device:

   ```
   > sudo crm --force cluster init sbd --sbd-device /dev/disk/by-id/DEVICE_ID
   ```

   

   You can use `--sbd-device` (or `-s`) multiple times to configure up to three SBD devices.

   The script initializes SBD on the shared storage device, creates a `fence_sbd` cluster resource, and updates the SBD configuration file and timeout settings. The script also puts the cluster into maintenance mode, restarts the cluster services, then puts the cluster back into normal operation.

6. Check the status of the cluster:

   ```
   > sudo crm status
   ```

   The nodes should be `Online` and the resources should be `Started`.

7. If you stopped DLM and any DLM-dependent resources, you can start them again now. If the resources are in a group, you can just start the group. If the resources aren't in a group, start DLM first and then start the DLM-dependent resources. You can start multiple resources at once:

```
> sudo crm resource start RESOURCE1 RESOURCE2
```

8. Check the status of SBD:

```
> sudo crm sbd status
# Type of SBD:
Disk-based SBD configured
```

6 Changing disk-based SBD to diskless SBD

Use this procedure to change disk-based SBD to diskless SBD.

Diskless SBD cannot handle a split-brain scenario for a two-node cluster. This configuration should only be used for clusters with more than two nodes, or in combination with QDevice to help handle split-brain scenarios.



Important: Cluster restart required

In this procedure, the setup script automatically puts the cluster into maintenance mode and restarts the cluster services. In maintenance mode, the cluster stops monitoring all resources. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while in maintenance mode.



Warning: Stop DLM before putting the cluster into maintenance mode

To put the cluster into maintenance mode and then stop or restart the cluster services, you must first stop DLM and any resources that depend on it, such as GFS2 or `lvmlockd`. When the cluster is in maintenance mode, DLM keeps running even if the cluster services are offline. However, DLM depends on Corosync, so if Corosync stops but DLM is still running, the DLM instances on each node consider the other instances lost and request node fencing.

PREREQUISITES

- Disk-based SBD is configured and running.

Perform this procedure on only one node in the cluster:

1. Log in either as the `root` user or as a user with **`sudo`** privileges.
2. Check whether any resources are running:

```
> sudo crm status
```

3. If DLM and any DLM-dependent resources are running, stop them first. If the resources are in a group, you can just stop the group. If the resources aren't in a group, stop the DLM-dependent resources first and then stop DLM. You can stop multiple resources at once:

```
> sudo crm --wait resource stop RESOURCE1 RESOURCE2
```

   Use `--wait` to make sure the resources have really stopped before the command returns.

4. Check the status of SBD:

```
> sudo crm sbd status
# Type of SBD:
Disk-based SBD configured
```

5. Configure diskless SBD. Use `--force` (or `-F`) to allow you to reconfigure SBD even when it is already running, and `--enable-sbd` (or `-S`) to specify that no device is needed:

```
> sudo crm --force cluster init sbd --enable-sbd
```

   The script stops and removes the `fence_sbd` cluster resource, then updates the SBD configuration file and timeout settings. The script also puts the cluster into maintenance mode, restarts the cluster services, then puts the cluster back into normal operation.

6. Check the status of the cluster:

```
> sudo crm status
```

   The nodes should be `Online` and the resources `Started`.

7. If you stopped DLM and any DLM-dependent resources, you can start them again now. If the resources are in a group, you can just start the group. If the resources aren't in a group, start DLM first and then start the DLM-dependent resources. You can start multiple resources at once:

```
> sudo crm resource start RESOURCE1 RESOURCE2
```

8. Check the status of SBD:

```
> sudo crm sbd status
# Type of SBD:
Diskless SBD configured
```

7 Adding another SBD device

Use this procedure to add more SBD devices to a cluster that already has disk-based SBD configured. The cluster can have up to three SBD devices.

Important: Cluster restart required

In this procedure, the setup script has to restart the cluster services. If any non-`fence` resources are running, put the cluster into maintenance mode before running the script. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while the cluster is in maintenance mode.

Warning: Stop DLM before putting the cluster into maintenance mode

To put the cluster into maintenance mode and then stop or restart the cluster services, you must first stop DLM and any resources that depend on it, such as GFS2 or `lvmlockd`. When the cluster is in maintenance mode, DLM keeps running even if the cluster services are offline. However, DLM depends on Corosync, so if Corosync stops but DLM is still running, the DLM instances on each node consider the other instances lost and request node fencing.

PREREQUISITES

- Disk-based SBD is already configured and running with at least one device.
- An additional shared storage device is accessible from all cluster nodes.

Perform this procedure on only one node in the cluster:

1. Log in either as the `root` user or as a user with **`sudo`** privileges.
2. Check whether any resources are running:

   ```
   > sudo crm status
   ```

3. If DLM and any DLM-dependent resources are running, stop them first. If the resources are in a group, you can just stop the group. If the resources aren't in a group, stop the DLM-dependent resources first and then stop DLM. You can stop multiple resources at once:

   ```
   > sudo crm --wait resource stop RESOURCE1 RESOURCE2
   ```

   Use `--wait` to make sure the resources have really stopped before the command returns.

4. If any non-`fence` resources are running, put the cluster into maintenance mode:

   ```
   > sudo crm maintenance on
   ```

   In this state, the cluster stops monitoring all resources. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while the cluster is in maintenance mode.

5. Check which device or devices are already configured for use with SBD:

   ```
   > sudo crm sbd configure show sysconfig
   ```

   The output shows one or more device IDs in the `SBD_DEVICE` line.

6. Add a new device to the existing SBD configuration:

   ```
   > sudo crm sbd device add /dev/disk/by-id/DEVICE_ID
   ```

   The script initializes SBD on the new device, updates the SBD configuration file, and restarts the cluster services.

7. If the cluster is still in maintenance mode, put it back into normal operation:

   ```
   > sudo crm maintenance off
   ```

8. If you stopped DLM and any DLM-dependent resources, you can start them again now. If the resources are in a group, you can just start the group. If the resources aren't in a group, start DLM first and then start the DLM-dependent resources. You can start multiple resources at once:

```
> sudo crm resource start RESOURCE1 RESOURCE2
```

9. Check the SBD configuration again:

```
> sudo crm sbd configure show sysconfig
```

   The output should now show more devices.

10. Check the status of SBD to make sure all the nodes can see the new device:

```
> sudo crm sbd status
```

8 Replacing an existing SBD device with a new device

If you need to replace an SBD device, you can use `crm sbd device add` to add the new device and `crm sbd device remove` to remove the old device. If the cluster has two SBD devices, you can run these commands in any order. However, if the cluster has one or three SBD devices, you must run these commands in a specific order:

- One device: `crm sbd device remove` cannot remove the only device, so you must add the new device before you can remove the old device.
- Three devices: `crm sbd device add` cannot add a fourth device, so you must remove the old device before you can add the new device.



Important: Cluster restart required

In this procedure, the cluster services must be restarted twice: once after adding the new device *and* once after removing the old device. We recommend putting the cluster into maintenance mode first. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while the cluster is in maintenance mode.



Warning: Stop DLM before putting the cluster into maintenance mode

To put the cluster into maintenance mode and then stop or restart the cluster services, you must first stop DLM and any resources that depend on it, such as GFS2 or `lvmlockd`. When the cluster is in maintenance mode, DLM keeps running even if the cluster services are offline. However, DLM depends on Corosync, so if Corosync stops but DLM is still running, the DLM instances on each node consider the other instances lost and request node fencing.

PREREQUISITES

- Disk-based SBD is already configured and running with at least one device.
- An additional shared storage device is accessible from all cluster nodes.

Perform this procedure on only one node in the cluster:

1. Log in either as the `root` user or as a user with **`sudo`** privileges.
2. Check whether any resources are running:

   ```
   > sudo crm status
   ```

3. If DLM and any DLM-dependent resources are running, stop them first. If the resources are in a group, you can just stop the group. If the resources aren't in a group, stop the DLM-dependent resources first and then stop DLM. You can stop multiple resources at once:

   ```
   > sudo crm --wait resource stop RESOURCE1 RESOURCE2
   ```

   Use `--wait` to make sure the resources have really stopped before the command returns.

4. If any non-`fence` resources are running, put the cluster into maintenance mode:

   ```
   > sudo crm maintenance on
   ```

   In this state, the cluster stops monitoring all resources. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while the cluster is in maintenance mode.

5. Check how many devices are already configured for use with SBD:

   ```
   > sudo crm sbd configure show sysconfig
   ```

The output shows one or more device IDs in the `SBD_DEVICE` line. The number of devices determines the order of the next steps.

6. Add or remove a device, depending on the number of devices shown in *Step 5*:

   One device:

   Add the new device to the existing SBD configuration:

   ```
   > sudo crm sbd device add /dev/disk/by-id/DEVICE_ID
   ```

   The script restarts the cluster services automatically.

   Two or three devices:

   Remove the old device from the SBD configuration:

   ```
   > sudo crm sbd device remove /dev/disk/by-id/DEVICE_ID
   ```

   The script warns you to restart the cluster services manually.

7. If you need to restart the cluster services manually, run the following command:

   ```
   > sudo crm cluster restart --all
   ```

   Check the status of the cluster:

   ```
   > sudo crm status
   ```

   The nodes will have the status `UNCLEAN (offline)`, but will soon change to `Online`.

8. Add or remove a device, depending on the number of devices shown in *Step 5*:

   One device:

   Remove the old device from the SBD configuration:

   ```
   > sudo crm sbd device remove /dev/disk/by-id/DEVICE_ID
   ```

   The script warns you to restart the cluster services manually.

   Two or three devices:

   Add the new device to the existing SBD configuration:

   ```
   > sudo crm sbd device add /dev/disk/by-id/DEVICE_ID
   ```

   The script restarts the cluster services automatically.

9. If you need to restart the cluster services manually, run the following command:

```
> sudo crm cluster restart --all
```

Check the status of the cluster:

```
> sudo crm status
```

The nodes will have the status `UNCLEAN (offline)`, but will soon change to `Online`.

10. If the cluster is still in maintenance mode, put it back into normal operation:

```
> sudo crm maintenance off
```

11. If you stopped DLM and any DLM-dependent resources, you can start them again now. If the resources are in a group, you can just start the group. If the resources aren't in a group, start DLM first and then start the DLM-dependent resources. You can start multiple resources at once:

```
> sudo crm resource start RESOURCE1 RESOURCE2
```

12. Check the SBD configuration again:

```
> sudo crm sbd configure show sysconfig
```

The output should now show the new device in the `SBD_DEVICE` line.

13. Check the status of SBD to make sure the correct device is listed:

```
> sudo crm sbd status
```

9 Removing an SBD device

Use this procedure to remove an SBD device from a cluster with multiple SBD devices configured. You cannot use this method if there is only one SBD device configured.

Important: Cluster restart required

In this procedure, the setup script has to restart the cluster services. If any non-`fence` resources are running, put the cluster into maintenance mode before running the script. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while the cluster is in maintenance mode.



Warning: Stop DLM before putting the cluster into maintenance mode

To put the cluster into maintenance mode and then stop or restart the cluster services, you must first stop DLM and any resources that depend on it, such as GFS2 or `lvmlockd`. When the cluster is in maintenance mode, DLM keeps running even if the cluster services are offline. However, DLM depends on Corosync, so if Corosync stops but DLM is still running, the DLM instances on each node consider the other instances lost and request node fencing.

PREREQUISITES

- Disk-based SBD is configured with more than one device.

Perform this procedure on only one node in the cluster:

1. Log in either as the `root` user or as a user with **`sudo`** privileges.
2. Check whether any resources are running:

   ```
   > sudo crm status
   ```

3. If DLM and any DLM-dependent resources are running, stop them first. If the resources are in a group, you can just stop the group. If the resources aren't in a group, stop the DLM-dependent resources first and then stop DLM. You can stop multiple resources at once:

   ```
   > sudo crm --wait resource stop RESOURCE1 RESOURCE2
   ```

   Use `--wait` to make sure the resources have really stopped before the command returns.

4. If any non-`fence` resources are running, put the cluster into maintenance mode:

   ```
   > sudo crm maintenance on
   ```

   In this state, the cluster stops monitoring all resources. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while the cluster is in maintenance mode.

5. Check which devices are already configured for use with SBD:

   ```
   > sudo crm sbd configure show sysconfig
   ```

   The output shows multiple device IDs in the `SBD_DEVICE` line.

6. Remove a device from the SBD configuration:

```
> sudo crm sbd device remove /dev/disk/by-id/DEVICE_ID
```

   The script removes the device, updates the SBD configuration file, and restarts the cluster services.

7. If the cluster is still in maintenance mode, put it back into normal operation:

```
> sudo crm maintenance off
```

8. If you stopped DLM and any DLM-dependent resources, you can start them again now. If the resources are in a group, you can just start the group. If the resources aren't in a group, start DLM first and then start the DLM-dependent resources. You can start multiple resources at once:

```
> sudo crm resource start RESOURCE1 RESOURCE2
```

9. Check the SBD configuration again:

```
> sudo crm sbd configure show sysconfig
```

   The output should now show fewer devices.

10. Check the status of SBD to make sure the device was removed from all the nodes:

```
> sudo crm sbd status
```

10 Removing all SBD configuration

Use this procedure to remove all SBD-related configuration from the cluster. You might need to do this if you want to switch from SBD to a physical fencing device. Keep in mind that to be supported, all SUSE Linux Enterprise High Availability clusters *must* have either SBD or a physical fencing device configured.

Important: Cluster restart required

In this procedure, the setup script has to restart the cluster services. If any non-`fence` resources are running, put the cluster into maintenance mode before running the script. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while the cluster is in maintenance mode.



Warning: Stop DLM before putting the cluster into maintenance mode

To put the cluster into maintenance mode and then stop or restart the cluster services, you must first stop DLM and any resources that depend on it, such as GFS2 or `lvmlockd`. When the cluster is in maintenance mode, DLM keeps running even if the cluster services are offline. However, DLM depends on Corosync, so if Corosync stops but DLM is still running, the DLM instances on each node consider the other instances lost and request node fencing.

Perform this procedure on only one node in the cluster:

1. Log in either as the `root` user or as a user with **`sudo`** privileges.
2. Check whether any resources are running:

   ```
   > sudo crm status
   ```

3. If DLM and any DLM-dependent resources are running, stop them first. If the resources are in a group, you can just stop the group. If the resources aren't in a group, stop the DLM-dependent resources first and then stop DLM. You can stop multiple resources at once:

   ```
   > sudo crm --wait resource stop RESOURCE1 RESOURCE2
   ```

   Use `--wait` to make sure the resources have really stopped before the command returns.
4. If any non-`fence` resources are running, put the cluster into maintenance mode:

   ```
   > sudo crm maintenance on
   ```

   In this state, the cluster stops monitoring all resources. This allows the services managed by the resources to keep running while the cluster restarts. However, the resources won't be highly available while the cluster is in maintenance mode.
5. Remove the SBD configuration from the cluster:

   ```
   > sudo crm sbd purge
   ```

   The script stops the SBD service on all nodes, moves the SBD configuration file to a backup file, adjusts any SBD-related cluster properties, and restarts the cluster services.
6. If the cluster is still in maintenance mode, put it back into normal operation:

   ```
   > sudo crm maintenance off
   ```

7. If you stopped DLM and any DLM-dependent resources, you can start them again now. If the resources are in a group, you can just start the group. If the resources aren't in a group, start DLM first and then start the DLM-dependent resources. You can start multiple resources at once:

```
> sudo crm resource start RESOURCE1 RESOURCE2
```

8. Confirm that the SBD configuration is gone:

```
> sudo crm sbd status
ERROR: SBD configuration file /etc/sysconfig/sbd not found
```

9. Check the cluster configuration:

```
> sudo crm configure show
```

   The output should show `fencing-enabled=false` and no other SBD-related properties.

11 Legal Notice

Copyright© 2006– 2026 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”.

For SUSE trademarks, see https://www.suse.com/company/legal/. All other third-party trademarks are the property of their respective owners. Trademark symbols (®, ™ etc.) denote trademarks of SUSE and its affiliates. Asterisks (*) denote third-party trademarks.

All information found in this book has been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. Neither SUSE LLC, its affiliates, the authors, nor the translators shall be held liable for possible errors or the consequences thereof.

A GNU Free Documentation License

Copyright (C) 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 non-commercially. Secondarily, 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 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 non-commercially, 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 https://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.

HA glossary

ACL (access control list)

Access control lists can be used to limit or fine-tune a user's access to the CIB. An ACL consists of an ordered set of access rules. Each rule allows or denies read or write access to a part of the CIB. Rules can be combined into specific roles, and users can be assigned to a role that matches their tasks.

active/active, active/passive

How resources run on the nodes. Active/passive means that resources only run on the active node, but can move to the passive node if the active node fails. Active/active means that all nodes are active at once, and resources can run on (and move to) any node in the cluster.

arbitrator

An *arbitrator* is a machine running outside the cluster to provide an additional instance for cluster calculations. For example, *QNetd* provides a vote to help *QDevice* participate in *quorum* decisions.

CIB (cluster information base)

An XML representation of the whole cluster configuration and status (cluster options, nodes, resources, constraints, etc.). The CIB manager (`pacemaker-based`) keeps the CIB synchronized across the cluster nodes and handles requests to modify it.

clone

In the context of a cluster *resource*, a clone is a resource that can be active on multiple nodes. Any resource can be cloned if its resource agent supports it.

cluster

A *high-availability* cluster is a group of servers (physical or virtual) designed primarily to secure the highest possible availability of data, applications and services. Not to be confused with a *high-performance* cluster, which shares the application load to achieve faster results.

Cluster LVM (Cluster logical volume manager)

The term *Cluster LVM* indicates that LVM is being used in a cluster environment. This requires configuration adjustments to protect the LVM metadata on shared storage.

cluster offline upgrade

If a new product version includes major changes that are backward incompatible, the cluster requires a cluster offline upgrade. You must stop the cluster services on all nodes and upgrade the whole cluster before bringing the nodes back online.

For minor version upgrades, see *cluster rolling upgrade*.

cluster partition

A cluster partition occurs when communication fails between one or more nodes and the rest of the cluster. The nodes are split into partitions but are still active. They can only communicate with nodes in the same partition and are unaware of the separated nodes. This is known as a *split brain* scenario.

cluster rolling upgrade

In a cluster rolling upgrade, one cluster node at a time is upgraded while the rest of the cluster is still running. Stop the cluster services on a node, upgrade it, then bring it back online to join the cluster. Continue one by one until all cluster nodes are upgraded.

For major version upgrades, see *cluster offline upgrade*.

cluster services

The services that run the cluster. Typically this refers to Pacemaker and Corosync, but might also include SBD and QDevice if those services are used.

cluster stack

The ensemble of software technologies and components that make up a cluster.

colocation constraint

A type of *resource constraint* that specifies which resources can or cannot run together on a node.

concurrency violation

A resource that should be running on only one node in the cluster is running on several nodes.

Corosync

Corosync provides reliable messaging, membership and quorum information about the cluster. This is handled by the Corosync Cluster Engine, a group communication system.

CRM (cluster resource manager)

The management entity responsible for coordinating all non-local interactions in a High Availability cluster. SUSE Linux Enterprise High Availability uses *Pacemaker* as the CRM. It interacts with several components: local executors on its own node and on the other nodes, non-local CRMs, administrative commands, the fencing functionality, and the membership layer.

`crmsh` (CRM Shell)

The command-line utility `crmsh` manages the cluster, nodes and resources.

Csync2

A synchronization tool for replicating configuration files across all nodes in the cluster.

DC (designated coordinator)

The `pacemaker-controld` daemon is the cluster controller, which coordinates all actions. This daemon has an instance on each cluster node, but only one instance is elected to act as the DC. The DC is elected when the cluster services start, or if the current DC fails or leaves the cluster. The DC decides whether a cluster-wide change must be performed, such as fencing a node or moving resources.

disaster

An unexpected interruption of critical infrastructure caused by nature, humans, hardware failure, or software bugs.

disaster recovery

The process by which a function is restored to the normal, steady state after a disaster.

Disaster Recovery Plan

A strategy to recover from a disaster with the minimum impact on IT infrastructure.

DLM (Distributed Lock Manager)

DLM coordinates accesses to shared resources in a cluster, for example, managing file locking in clustered file systems to increase performance and availability.

DRBD

DRBD® is a block device designed for building High Availability clusters. It replicates data on a primary device to secondary devices in a way that ensures all copies of the data remain identical.

existing cluster

The term *existing cluster* is used to refer to any cluster that consists of at least one node. An existing cluster has a basic *Corosync* configuration that defines the communication channels, but does not necessarily have resource configuration yet.

failover

Occurs when a resource or node fails on one machine and the affected resources move to another node.

failover domain

A named subset of cluster nodes that are eligible to run a resource if a node fails.

fencing

Prevents access to a shared resource by isolated or failing cluster members. There are two classes of fencing: *resource-level* fencing and *node-level* fencing. Resource-level fencing ensures exclusive access to a resource. Node-level fencing prevents a failed node from accessing shared resources and prevents resources from running on a node with an uncertain status. This is usually done by resetting or powering off the node. In a *Pacemaker* cluster, node fencing is managed by the fencing subsystem `pacemaker-fenced`.

GFS2

Global File System 2 (GFS2) is a shared disk file system for Linux computer clusters. GFS2 allows all nodes to have direct concurrent access to the same shared block storage. GFS2 has no disconnected operating mode, and no client or server roles. All nodes in a GFS2 cluster function as peers. GFS2 supports up to 32 cluster nodes. Using GFS2 in a cluster requires hardware to allow access to the shared storage, and a lock manager to control access to the storage.

group

Resource groups contain multiple resources that need to be located together, started sequentially and stopped in the reverse order.

Hawk (High Availability Web Konsole)

A user-friendly Web-based interface for monitoring and administering a High Availability cluster from Linux or non-Linux machines. Hawk can be accessed from any machine that can connect to the cluster nodes, using a graphical Web browser.

heuristics

QDevice supports using a set of commands (*heuristics*) that run locally on start-up of cluster services, cluster membership change, successful connection to the *QNetd* server, or optionally at regular times. The result is used in calculations to determine which partition should have *quorum*.

knet (kronosnet)

A network abstraction layer supporting redundancy, security, fault tolerance, and fast fail-over of network links. In SUSE Linux Enterprise High Availability 16, *knet* is the default transport protocol for the *Corosync* communication channels.

local cluster

A single cluster in one location (for example, all nodes are located in one data center). Network latency is minimal. Storage is typically accessed synchronously by all nodes.

local executor

The local executor is located between *Pacemaker* and the resources on each node. Through the `pacemaker-execd` daemon, Pacemaker can start, stop and monitor resources.

location

In the context of a whole cluster, *location* can refer to the physical location of nodes (for example, all nodes might be located in the same data center). In the context of a *location constraint*, *location* refers to the nodes on which a resource can or cannot run.

location constraint

A type of *resource constraint* that defines the nodes on which a resource can or cannot run.

meta attribute

Parameters that tell the *CRM (cluster resource manager)* how to treat a specific *resource*. For example, you might define a resource's priority or target role.

metro cluster

A single cluster that can stretch over multiple buildings or data centers, with all sites connected by Fibre Channel. Network latency is usually low. Storage is frequently replicated using mirroring or synchronous replication.

network device bonding

Network device bonding combines two or more network interfaces into a single bonded device to increase bandwidth and/or provide redundancy. When using *Corosync*, the bonded device is not managed by the cluster software. Therefore, the bonded device must be configured on every cluster node that might need to access it.

node

Any server (physical or virtual) that is a member of a cluster.

order constraint

A type of *resource constraint* that defines the sequence of actions.

Pacemaker

Pacemaker is the *CRM (cluster resource manager)* in SUSE Linux Enterprise High Availability, or the "brain" that reacts to events occurring in the cluster. Events might be nodes that join or leave the cluster, failure of resources, or scheduled activities such as maintenance, for example. The `pacemakerd` daemon launches and monitors all other related daemons.

parameters (instance attributes)

Parameters determine which instance of a service the *resource* controls.

primitive

A primitive resource is the most basic type of cluster resource.

promotable clone

Promotable clones are a special type of *clone* resource that can be promoted. Active instances of these resources are divided into two states: promoted and unpromoted (also known as "active and passive" or "primary and secondary").

QDevice

QDevice and *QNetd* participate in *quorum* decisions. The `corosync-qdevice` daemon runs on each cluster node and communicates with QNetd to provide a configurable number of votes, allowing a cluster to sustain more node failures than the standard quorum rules allow.

QNetd

QNetd is an *arbitrator* that runs outside the cluster. The `corosync-qnetd` daemon provides a vote to the `corosync-qdevice` daemon on each node to help it participate in quorum decisions.

quorum

A *cluster partition* is defined to have quorum (be *quorate*) if it has the majority of nodes (or "votes"). Quorum distinguishes exactly one partition. This is part of the algorithm to prevent several disconnected partitions or nodes ("split brain") from proceeding and causing data and service corruption. Quorum is a prerequisite for fencing, which then ensures that quorum is unique.

RA (resource agent)

A script acting as a proxy to manage a *resource* (for example, to start, stop or monitor a resource). SUSE Linux Enterprise High Availability supports different kinds of resource agents.

ReaR (Relax and Recover)

An administrator tool set for creating *disaster recovery* images.

resource

Any type of service or application that is known to *Pacemaker*, for example, an IP address, a file system, or a database. The term *resource* is also used for *DRBD*, where it names a set of block devices that use a common connection for replication.

resource constraint

Resource constraints specify which cluster nodes a resource can run on, what order the resources should start in, and which other resources a specific resource is dependent on.

See also *colocation constraint*, *location constraint* and *order constraint*.

resource set

As an alternative format for defining location, colocation or order constraints, you can use *resource sets*, where primitives are grouped together in one set. When creating a constraint, you can specify multiple resources for the constraint to apply to.

resource template

To help create many resources with similar configurations, you can define a resource template. After being defined, it can be referenced in primitives or in certain types of constraints. If a template is referenced in a primitive, the primitive inherits all operations, instance attributes (parameters), meta attributes and utilization attributes defined in the template.

SBD (STONITH Block Device)

SBD provides a node *fencing* mechanism through the exchange of messages via shared block storage. Alternatively, it can be used in diskless mode. In either case, it needs a hardware or software *watchdog* on each node to ensure that misbehaving nodes are really stopped.

scheduler

The scheduler is implemented as `pacemaker-schedulerd`. When a cluster transition is needed, `pacemaker-schedulerd` calculates the expected next state of the cluster and determines what actions need to be scheduled to achieve the next state.

split brain

A scenario in which the cluster nodes are divided into two or more groups that do not know about each other (either through a software or hardware failure). Also known as a *partitioned cluster* scenario.

The term *split brain* is also used in *DRBD* but means that the nodes contain different data.

SPOF (single point of failure)

Any component of a cluster that, if it fails, triggers the failure of the entire cluster.

stickiness

Resource stickiness is a *meta attribute* that determines how much a resource prefers to stay on its current node.

switchover

The planned moving of resources to other nodes in a cluster. See also *failover*.

utilization

Tells the CRM what capacity a certain *resource* requires from a node.

watchdog

SBD (STONITH Block Device) needs a watchdog on each node to ensure that misbehaving nodes are really stopped. SBD "feeds" the watchdog by regularly writing a service pulse to it. If SBD stops feeding the watchdog, the hardware enforces a system restart. This protects against failures of the SBD process itself, such as becoming stuck on an I/O error.