



Troubleshooting

Table of contents

Split-Brain Recovery in HA Installation

Performing Failover on Non-Master Node

Restoring UFM Data Upon In-Service Upgrade Failure

Split-Brain Recovery in HA Installation

The split-brain problem is a DRBD synchronization issue (HA status shows DUnknown in the DRBD disk state), which occurs when both HA nodes are rebooted. For example, in cases of electricity shut-down. To recover, please follow the below steps:

- **Step 1:** Manually choose a node where data modifications will be discarded.

It is called the split-brain victim. Choose wisely; all modifications will be lost! When in doubt, run a backup of the victim's data before you continue.

When running a Pacemaker cluster, you can enable maintenance mode. If the split-brain victim is in the Primary role, bring down all applications using this resource. Now switch the victim to the Secondary role:

```
victim# drbdadm secondary ha_data
```

- **Step 2:** Disconnect the resource if it's in connection state **WFCConnection**:

```
victim# drbdadm disconnect ha_data
```

- **Step 3:** Force discard of all modifications on the split-brain victim:

```
victim# drbdadm -- --discard-my-data connect ha_data
```

For DRBD 8.4.x:

```
victim# drbdadm connect --discard-my-data ha_data
```

- **Step 4:** Resync starts automatically if the survivor is in a WFCConnection network state. If the split-brain survivor is still in a Standalone connection state, reconnect it:

```
survivor# drbdadm connect ha_data
```

Now the resynchronization from the survivor (SyncSource) to the victim (SyncTarget) starts immediately. There is no full sync initiated, but all modifications on the victim will be overwritten by the survivor's data, and modifications on the survivor will be applied to the victim.

Performing Failover on Non-Master Node

The `ufm_ha_cluster` failover action fails with the following error: "Cannot perform failover on non-master node". To fix, follow the below action:

- **Step 1:** Verify that `/etc/hosts` file on both the master and standby UFM hosts contains the correct host names and IP addresses mapping.
- **Step 2:** If necessary, fix the mapping and retry the failover command.

Restoring UFM Data Upon In-Service Upgrade Failure

Note

These instructions apply in high availability scenario only.

In the event of an in-service upgrade failure, the previous version of UFM's data will be safeguarded as a backup in the `/opt/ufm/BACKUP` directory, formatted as `"ufm_upgrade_backup_<prev_version>-<new_version<_<date>.zip."`

To restore the data on the unupgraded node, follow these steps:

1. Copy the backup file from the upgraded node to the unupgraded node using the following command:

```
scp /opt/ufm/BACKUP/ufm_upgrade_backup_<prev_version>-<new_version<_<date>.zip
```

```
root@<unupgraded_node_ip>:/opt/ufm/BACKUP/
```

2. Perform a failover of UFM to the master node, which is mandatory for data mount migration (including '/opt/ufm/files') to the master node: On the Master node, execute:

```
ufm_ha_cluster takeover
```

3. Stop UFM on the unupgraded node:

```
ufm_ha_cluster stop
```

4. Restore UFM configuration files from the backup:

```
/opt/ufm/scripts/ufm_restore.sh -f /opt/ufm/BACKUP/ufm_upgrade_backup_<prev_version>-<new_version<_<date>.zip
```

5. Start UFM on the unupgraded node (Note: Only the upgraded node can function until the upgrade issue is resolved, and failovers will not work).

Now, the issue that caused the upgrade failure can be addressed. If the problem is resolved, you can attempt the in-service upgrade again by failing UFM over to the upgraded node.

Alternatively, if needed, you can revert the changes made by reinstalling the old UFM version on the upgraded node.

© Copyright 2024, NVIDIA. PDF Generated on 08/14/2024