

LOTC Disk Replication Communication

OPERATING INSTRUCTIONS

Copyright

© Ericsson AB 2017. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

All trademarks mentioned herein are the property of their respective owners. These are shown in the document Trademark Information.



Contents

1	Alarm Description	1
2	Procedure	1
2.1	Handle Alarm LOTC Disk Replication Communication	1





1 Alarm Description

The alarm is raised when the control nodes have lost connection to each other for more than 20 minutes, and are no longer in redundant mode. The control node pair is in a non-redundant mode when the control nodes have no connection with each other.

Table 1 LOTC Disk Replication Communication Alarm Causes

Alarm Cause	Description	Fault Reason	Fault Location	Impact
Loss of connection between control nodes for more than 20 minutes	The control nodes have lost connection to each other for more than 20 minutes. The Linux® service Distributed Replicated Block Device (DRBD) is not in connected mode.	Network failure leading to communication problems between the control nodes	Network	Both controllers take the primary role and no data is transferred between the nodes
		Hardware failure on the secondary control node	Secondary control node	If one of the controller nodes is down, the cluster does not have a controller node to which it can fail over

Note: This alarm can appear as a result of a maintenance activity.

2 Procedure

2.1 Handle Alarm LOTC Disk Replication Communication

Prerequisites

- This instruction references the following document:
 - [Data Collection Guideline](#)
- No tools are required.
- The following conditions must apply:
 - The alarm is raised.

Steps



1. Log on to the host to access a Linux shell, for example:

```
ssh <user>@<hostname> -p 7022
```

The hostname is part of alarm attribute Source.

2. Is the alarm raised during initial installation or replacement of a control node?

Yes: Continue with the next step.

No: Proceed with Step 8.

3. Check which drbd version you are running:

```
cat /proc/drbd
```

```
version: 8.4.2 (api:1/proto:86-101)
GIT-hash: 7ad5f850d711223713d6dcadc3dd48860321070c build by root@lixia, 2012-09-19 16:40:30
0: cs:SyncSource ro:Primary/Secondary ds:UpToDate/UpToDate C r---n-
```

Is drbd version: 8.* ?

Yes: continue with next step.

No: proceed with Step 6.

4. Wait for DRBD connection to be established. Check if the following command results in output cs:Connected:

```
cat /proc/drbd
```

The following is an example output in a normal situation. The connection state (cs) is Connected. The alarm is cleared within 5 seconds.

```
version: 8.4.2 (api:1/proto:86-101)
GIT-hash: 7ad5f850d711223713d6dcadc3dd48860321070c build by root@lixia, 2012-09-19 16:40:30
0: cs:Connected ro:Primary/Secondary ds:UpToDate/UpToDate C r-----
   ns:438816 nr:0 dw:372 dr:440669 al:11 bm:40 lo:0 pe:0 ua:0 ap:0 ep:1 wo:f oo
```

The following is an example output in a faulty situation. The connection state (cs) is WfConnection (Waiting For Connection).

```
version: 8.4.2 (api:1/proto:86-101)
GIT-hash: 7ad5f850d711223713d6dcadc3dd48860321070c build by root@lixia, 2012-09-19 16:40:30
0: cs:WfConnection ro:Primary/Unknown ds:UpToDate/DUnknown C r-----
   ns:143396 nr:0 dw:448 dr:147057 al:17 bm:28 lo:0 pe:0 ua:0 ap:0 ep:1 wo:f oos:84
```

5. Does the output contain cs:Connected and is the alarm cleared?

Yes: Proceed with Step 18.

No: There can be a deployment issue. Perform data collection, refer to [Data Collection Guideline](#). Contact the IT organization responsible for the installation or replacement of the control node. Proceed with Step 8.

6. Wait for DRBD connection to be established. Check if the following command results in output connection:Connected:



```
drbdsetup events2 --statistics --now
```

The following is an example output in a normal situation. The connection state (connection) is Connected. The alarm is cleared within 5 seconds.

```
exists resource name:drbd0 role:Primary suspended:no write-ordering:flush
exists connection name:drbd0 peer-node-id:1 conn-name:node2-vc11 connection:Connected role:Secondary congested:no
exists device name:drbd0 volume:0 minor:0 disk:UpToDate size:10485760 read:10774713 written:4200 al-writes:6 ⇒
bm-writes:0 upper-pending:0 lower-pending:0 al-suspended:no blocked:no
exists peer-device name:drbd0 peer-node-id:1 conn-name:node2-vc11 volume:0 replication:Established ⇒
peer-disk:UpToDate resync-suspended:no received:0 sent:10486244 out-of-sync:0 pending:0 unacked:0
exists -
```

The following is an example output in a faulty situation. The connection state (connection) is Connecting.

```
exists resource name:drbd0 role:Primary suspended:no write-ordering:flush
exists connection name:drbd0 peer-node-id:1 conn-name:node2-vc11 connection:Connecting role:Unknown congested:no
exists device name:drbd0 volume:0 minor:0 disk:UpToDate size:10485760 read:289085 written:3420 al-writes:5 ⇒
bm-writes:0 upper-pending:0 lower-pending:0 al-suspended:no blocked:no
exists peer-device name:drbd0 peer-node-id:1 conn-name:node2-vc11 volume:0 replication:0ff peer-disk:DUnknown ⇒
resync-suspended:no received:0 sent:0 out-of-sync:974888 pending:0 unacked:0
exists -
```

7. Does the output contain connection:Connected and is the alarm cleared?

Yes: Proceed with Step 18.

No: Perform data collection, refer to Data Collection Guideline. Contact the deployment organization. Proceed with Step 18.

8. Identify the DRBD interfaces as follows:

- a. Get the name of the interface (eth<x>):

```
cat /etc/cluster/nodes/this/networks/internal/primary/  
interface/name
```

The following is an example output:

```
eth0
```

- b. Get the IP address (<ip>):

```
cat /etc/cluster/nodes/this/networks/internal/primary/a  
address
```

The following is an example output:

```
169.254.43.11
```

- c. Get the network mask (<netmask>):

```
cat /etc/cluster/nodes/this/networks/internal/primary/  
network/netmask
```

The following is an example output:



255.255.255.0

9. Check the log `/var/log/messages` for recent system log messages indicating DRBD interface-related issues, for example (to show the last 1000 lines in the log):

```
tail -1000 /var/log/messages
```

The following is an example output in a faulty situation:

```
Aug 26 12:17:52 SC-1 kernel: [ 277.720545] hrtimer: interrupt took 572013 ns
Aug 26 12:32:50 SC-1 kernel: [ 1175.612842] tipc: Resetting bearer <eth:eth0>
Aug 26 12:32:50 SC-1 dhcpcd: receive_packet failed on eth0: Network is down
Aug 26 12:32:50 SC-1 syslog-ng[1810]: I/O error occurred while writing; fd='6', error='Network =>
is unreachable (101)'
Aug 26 12:32:50 SC-1 syslog-ng[1810]: Connection broken; time_reopen='10'
Aug 26 12:32:59 SC-1 ntpd[2240]: sendto(192.0.2.10) (fd=23): Network is unreachable
Aug 26 12:33:00 SC-1 syslog-ng[1810]: Connection failed; error='Network is unreachable (101)'
Aug 26 12:33:00 SC-1 syslog-ng[1810]: Initiating connection failed, reconnecting; time_reopen='10'
Aug 26 12:33:10 SC-1 syslog-ng[1810]: Connection failed; error='Network is unreachable (101)'
Aug 26 12:33:10 SC-1 syslog-ng[1810]: Initiating connection failed, reconnecting; time_reopen='10'
Aug 26 12:33:20 SC-1 syslog-ng[1810]: Connection failed; error='Network is unreachable (101)'
Aug 26 12:33:20 SC-1 syslog-ng[1810]: Initiating connection failed, reconnecting; time_reopen='10'
```

Note: In this output, `eth0` is the interface used by the DRBD.

10. Are there any issues with the network interface used for the DRBD?

Yes: Continue with the next step.

No: Proceed with Step 16.

11. Check the status of the interface used by the DRBD:

```
ifconfig
```

The following is an example output:

```
eth0
Link encap:Ethernet HWaddr 00:50:56:92:02:38
inet addr:10.64.87.136 Bcast:10.64.87.191 Mask:255.255.255.192
inet6 addr: fe80::250:56ff:fe92:238/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:3884520 errors:0 dropped:0 overruns:0 frame:0
TX packets:178358 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:1000
RX bytes:333841297 (318.3 Mb) TX bytes:10087705 (9.6 Mb)
```

Note: The keywords `UP` and `RUNNING` in the output means that the DRBD interface is operational.

12. Is the DRBD interface operational?

Yes: With a high probability, there is a network issue. Perform data collection, refer to [Data Collection Guideline](#). Contact the network administrator. Proceed with Step 18.

No: Continue with the next step.

13. Try to bring up the interface used by the DRBD:



```
ifconfig <interface> <ip> netmask <mask>
```

Use the values collected in Step 8, for example:

```
ifconfig eth0 169.254.43.11 netmask 255.255.255.0
```

14. Check the status of the interface used by the DRBD:

```
ifconfig
```

The following is an example output:

```
eth0
Link encap:Ethernet HWaddr 00:50:56:92:02:38
inet addr:10.64.87.136 Bcast:10.64.87.191 Mask:255.255.255.192
inet6 addr: fe80::250:56ff:fe92:238/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:3884520 errors:0 dropped:0 overruns:0 frame:0
TX packets:178358 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:1000
RX bytes:333841297 (318.3 Mb) TX bytes:10087705 (9.6 Mb)
```

15. Is the DRBD interface operational and is the alarm cleared?

Yes: Proceed with Step 18.

No: Continue with the next step.

16. Perform data collection, refer to [Data Collection Guideline](#).

17. Consult the next level of maintenance support. Further actions are outside the scope of this instruction.

18. Job is completed.