



High Availability and Disaster Recovery

SL1 version 12.1.0 , Rev 2

Table of Contents

Introduction	4
Disaster Recovery	6
High Availability	6
Differences Between Disaster Recovery and High Availability for Database Servers	6
Configuring MX Records	7
Disaster Recovery with Two Appliances	8
Prerequisites	9
Unique Hostnames	9
Licensing DRBD Proxy	10
Using a Virtual IP Address	10
Deployment Recommendations	10
Configuring Disaster Recovery	10
Configuring the Primary Appliance	10
Configuring the Secondary Appliance	17
Licensing the Secondary Appliance	21
Configuring Data Collection Servers and Message Collection Servers	23
Failover	24
Failover When Both Database Appliances are Accessible	25
Failover When the Primary Database Appliance is Inaccessible	27
Reconfiguring Administration Portals	29
Verifying that a Database Server is Primary	30
Failback	31
High Availability with Two Appliances	33
Prerequisites	34
Unique Hostnames	34
Addressing the Cluster	35
Configuring Heartbeat IP Addresses	35
Testing the Heartbeat Network	36
Deployment Recommendations	37
Configuring High Availability	37
Configuring the Quorum Witness	37
Configuring the Primary Appliance	39
Configuring the Secondary Appliance	44
Licensing the Secondary Appliance	49
Configuring Data Collection Servers and Message Collection Servers	50
Failover	51
Manual Failover for High Availability Clusters	51
Verifying that a Database Server is Active	53
High Availability and Disaster Recovery with Three Appliances	54
Prerequisites	56
Unique Host Names	56
Licensing DRBD Proxy	57
Addressing the Cluster	57
Reconfiguring an Existing High Availability System	58
Configuring a Heartbeat IP Addresses	58
Testing the Heartbeat Network	59
Deployment Recommendations	60
Configuring Three Appliances for High Availability and Disaster Recovery	60
Configuring the Quorum Witness	60
Configuring the Primary High Availability Appliance	62

Configuring the Secondary High Availability Appliance	68
Configuring the Disaster Recovery Appliance	73
Licensing the Secondary High Availability and Disaster Appliances	77
Configuring Data Collection Servers and Message Collection Servers	78
Failover	79
Manual Failover Between the Appliances in the High Availability Cluster	79
Manual Failover Between the High Availability Cluster and the Disaster Recovery Appliance	81
Failover when the High Availability Cluster is Inaccessible	83
Manual Failback Between the Disaster Recovery Appliance and the High Availability Cluster	85
Reconfiguring Administration Portals	89
Verifying that a Database Server is Primary	90
Troubleshooting High Availability and Disaster Recovery Configuration	91
Overview	92
What is a Split-brain Configuration?	92
Resolving a Split-brain Configuration on a Disaster Recovery Stack with Two Appliances	92
Resolving a Split-brain Configuration on a High Availability Stack with Two Appliances	95
Resolving a Split-brain Configuration on a High Availability and Disaster Recovery Stack with Three Appliances	99
Resolving a Split-brain Between the Primary High Availability and Secondary High Availability Nodes	99
Resolving a Split-brain Between the Primary High Availability and the Disaster Recovery Nodes	102
Diverged Data	104
Overview	104
Resolving a Diverged Data Issue	104
Firewall Issues	105
Overview	105
Resolving Firewall Issues	105
Ahead-Behind or Disconnect/Reconnect Issues	106
Overview	106
Resolving Self-Monitoring Network Issues	106
Other Issues	107
WFReportParams Loops	107
Network Issues	107
Kernel Package Issues	108

Chapter

1



Introduction

Overview

This manual is intended for system administrators responsible for setting up and maintaining Database Servers for Disaster Recovery or High Availability, or both configurations. The ScienceLogic SL1 platform supports redundancy at every tier such that a single failure will not result in the failure of the overall platform.

- At the edge, Message Collectors (MC) and Data Collector units(CU) can be configured in groups to allow the load to be picked up by the remaining units should one fail.
- With the Admin Portal (AP), multiple APs can be deployed behind a load balancer, allowing reconfiguration of load based on availability.
- For central databases (CDB), block level replication is used to ensure multiple identical copies of the database exist.
- With High Availability, two CDB instances are in close proximity: one as active and one as standby. This allows for automated failover if a problem with the active CDB is detected.
- With Disaster Recovery, the two CDB instances can be located in separate data centers, but no automated failover happens if a failure of the primary CDB is detected.
- DRDB is used as the underlying block level replication between the CDB instances. DRDB supports High Availability, Disaster

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon (.
- To view a page containing all of the menu options, click the Advanced menu icon ().

Disaster Recovery	6
High Availability	6
Differences Between Disaster Recovery and High Availability for Database Servers	6

Disaster Recovery

You can configure SL1 to replicate data stored on a Database Server to a Disaster Recovery appliance with the same specifications. You can install the Disaster Recovery appliance at the same site as the primary Database Server (although this is not recommended) or at a different location.

If the primary Database Server fails for any reason, you must manually perform failover. Failover to the Disaster recovery appliance is not automated by SL1.

High Availability

You can cluster Database Servers in the same location to allow for automatic failover.

A cluster includes an *active* Database Server and a *passive* Database Server. The passive Database Server provides redundancy and is dormant unless a failure occurs on the active Database Server. SL1 uses block-level replication to ensure that the data on each Database Server's primary file system is identical and that each Database Server is ready for failover if necessary. If the active Database Server fails, the passive Database Server automatically becomes active and performs all required database tasks. The previously passive Database Server remains active until another failure occurs.

Each database cluster uses a virtual IP address that is always associated with the primary Database Server. No reconfiguration of Administration Portals is required in the event of failover.

IMPORTANT: High Availability for Azure deployments is supported for installations of 12.1.x that are running on Oracle Linux 8 (OL8). ScienceLogic recommends that customers running SL1 versions prior to 12.1.x upgrade and then complete the High Availability setup and configuration.

Differences Between Disaster Recovery and High Availability for Database Servers

SL1 provides two solutions that allow for failover to another Database Server if the primary Database Server fails: Disaster Recovery and High Availability. There are several differences between these two distinct features:

- **Location.** The primary and secondary databases in a High Availability configuration must be located together to configure the heartbeat network. In a Disaster Recovery configuration, the primary and secondary databases can be in different locations.
- **Failover.** In a High Availability configuration, SL1 performs failover automatically, although a manual failover option is available. In a Disaster Recovery configuration, failover must be performed manually.
- **System Operations.** A High Availability configuration maintains SL1 system operations if failure occurs on the hardware or software on the primary Database Server. A Disaster Recovery configuration maintains SL1 system operations if the data center where the primary Database Server is located has a major outage, provides a spare Database Server that can be quickly installed if the primary Database Server has a

permanent hardware failure, and/or to allow for rotation of SL1 system operations between two data centers.

NOTE: A Distributed SL1 system can be configured for both High Availability and Disaster Recovery.

NOTE: High Availability and Disaster Recovery are *not supported* for All-In-One Appliances.

Configuring MX Records

In all configurations, the primary Database Server is responsible for processing all inbound email. To prevent duplicate emails from being processed, all inbound email must be delivered to the current primary Database Server only. When Disaster Recovery is configured, the mail process is configured to be running on the current primary Database Server only. When you configure your mail exchanger (MX) record for SL1 :

- Include the hostname of the Database Server that will be primary under normal conditions at the lowest MX-level (making that hostname the highest-priority).
- If your configuration includes a High Availability cluster, include the hostname of the secondary cluster Database Server at the next lowest MX-level.
- For Disaster Recovery, include the hostname of the Database Server at the highest MX-level (making that hostname the lowest priority).



Disaster Recovery with Two Appliances

Overview

This chapter describes how to configure two appliances for Disaster Recovery.

This chapter assumes that you are comfortable using a UNIX shell session and can use the basic functions within the vi editor.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon ().
- To view a page containing all of the menu options, click the Advanced menu icon ().

This chapter covers the following topics:

<i>Prerequisites</i>	9
<i>Unique Hostnames</i>	9
<i>Deployment Recommendations</i>	10
<i>Configuring Disaster Recovery</i>	10
<i>Failover</i>	24

Prerequisites

Before performing the steps listed in this chapter, you must:

- Install and license each appliance
- Have an Administrator account to log in to the Web Configuration Utility for each appliance
- Have SSH or console access to each appliance
- Know the em7admin console username and password for each appliance
- Have identical hardware or virtual machine specifications on each appliance
- Have configured a unique hostname on each appliance
- Know the MariaDB username and password
- If the two appliances are not connected with a crossover cable, you must:
 - use a [DRBD proxy license](#)
 - know the maximum link speed, in megabytes per second, between the two appliances
- Optionally, if the two appliances you are configuring have their primary network adapter connected to the same network subnet, you can have an available IP address to configure as a [virtual IP](#)

Unique Hostnames

You must ensure that a unique hostname is configured on each SL1 appliance. The hostname of an appliance is configured during the initial installation. To view and change the hostname of an appliance:

1. Log in to the console of the SL1 appliance as the em7admin user. The current hostname appears before the command-prompt. For example, the login prompt might look like this, with the current hostname highlighted in bold:

```
login as: em7admin
```

```
em7admin@10.64.68.31's password:
```

```
Last login: Wed Apr 27 21:25:26 2016 from silo1651.sciencelogic.local
```

```
[em7admin@HADB01 ~]$
```

2. To change the hostname, run the following command:

```
sudo hostnamectl set-hostname <new hostname>
```

3. When prompted, enter the password for the em7admin user.

Licensing DRBD Proxy

DRBD Proxy buffers all data between the active and redundant appliances to compensate for any bandwidth limitations. In addition, DRBD compresses and encrypts the data sent from the active appliance to the redundant appliance.

You must use DRBD Proxy if you are:

- Configuring three appliances for High Availability and Disaster Recovery.
- Configuring two appliances for Disaster Recovery and will not be configuring a direct connection between your appliances with a crossover cable.

To license DRBD Proxy, copy the ***drbd-proxy.license*** file to the ***/etc*** directory on all appliances in your system.

Using a Virtual IP Address

If the two appliances you are configuring for Disaster Recovery are connected to the same network subnet using their primary network adapters, you can optionally specify a virtual IP address during the configuration. The virtual IP address is associated with the primary appliance and transitions between the appliances during failover and failback.

If you use a virtual IP address, you do not have to [reconfigure your Administration Portals](#) after failover and failback. The virtual IP address must be on the same network subnet as the primary network adapters of the appliances.

Deployment Recommendations

Disaster Recovery replication is sensitive to packet-loss and the variations in time delay between when a signal is transmitted and when it's received over a network connection (network jitter). Use Quality of Service (QoS) to classify the Disaster Recovery network traffic to prevent network drops and ensures reliable transmission for replication traffic.

Configuring Disaster Recovery

This section describes how to configure the Primary appliance and the Secondary appliance for Disaster Recovery.

Configuring the Primary Appliance

To configure the Primary appliance for Disaster Recovery, perform the following steps:

1. Log in to the console of the Primary appliance as the em7admin user.
2. Run the following command:

```
sudo -s
```

3. When prompted, enter the password for the em7admin user.
4. Run the following command:

```
silo-cluster-install
```

The following prompt appears:

```
v3.7
```

```
This wizard will assist you in setting up clustering for  
ScienceLogic appliance. Please be sure to consult the ScienceLogic  
documentation before running this script to ensure that all  
prerequisites have been met.
```

```
1) HA
```

```
2) DR
```

```
3) HA+DR
```

```
4) Quit
```

```
Please select the architecture you'd like to setup:
```

5. Enter option 2. The following prompt appears:

```
Will this system be the primary/active node in a new DR cluster?  
(y/n)
```

6. Enter **yes**. The following prompt appears:

```
Is there a cross-over cable between DR nodes? (y/n)
```

7. If you have a physical crossover cable connecting the two systems, enter **yes** otherwise answer **no**.

8. If you entered **yes**, the following prompt appears:

```
Choose the closest value for the speed of the dedicated between DR  
nodes.
```

```
1) 10 Gbps
```

```
2) 1 Gbps
```

```
What is the physical link speed between DR nodes?
```

9. Choose the speed that best represents your physical crossover cable speed. See requirements for crossover cable speeds. After selecting the link speed the following prompt appears:

```
Would you like to use a Virtual IP (VIP)? (y/n)
```

10. Enter **yes** if you would like to use a VIP otherwise choose **no**. The following confirmation prompt appears:

```
Review the following architecture selections to ensure they are
correct
```

```
Architecture: DR
```

```
Current node hostname: <hostname>
```

```
Crossover Cable: Yes
```

```
Dedicated link speed: 10Gbps
```

```
Virtual ID (VIP) used: <Yes / no>
```

```
DRBD Proxy: no
```

```
Is this architecture correct? (y/n)
```

- If you answered **no** in step 6, skip to the following confirmation prompt:

```
Review the following architecture selections to ensure they are
correct
```

```
Architecture: DR
```

```
Current node hostname: <hostname>
```

```
Crossover Cable: No
```

```
DRBD Proxy: Yes
```

```
Is this architecture correct? (y/n)
```

11. Review and answer **yes** if the architecture is correct, otherwise enter **no**. After entering **yes** the system performs validation of the architecture selections. When the architecture passes, the following message appears:

```
Validating architecture selections...
```

```
Architecture validation passed
```

NOTE: If errors are detected a failure message will be displayed and the script will exit. Please address the noted issues in the message and rerun the script.

12. After passing architecture validation the following prompt will display:

```
Please enter the following information about your environment
```

```
Primary node information:
```

```
Please choose the IP used for DRBD traffic for this server:
```

```
1) <First IP address of this appliance>
```

```
2) <Second IP address of this appliance>
```

```
...
```

```
Number:
```

13. Enter the number for the IP address of the network connection for the replication on the Primary appliance. When using a crossover cable you should select the IP for the dedicated network link. The following prompt appears:

```
Secondary node information:
```

```
What is the host name of the DR server:
```

14. Enter the hostname of the Disaster Recovery appliance. The following prompt appears:

```
Please enter the IP used for DRBD traffic for the DR server:
```

15. Enter the IP address of the network connection for the replication on the Disaster Recovery appliance. If you are using a crossover cable and you selected to use a VIP, the following prompt appears:

```
Virtual IP address information:
```

```
Please enter the Virtual IP Address:
```

NOTE: If you are not using a VIP skip to step18.

16. If you are prompted to enter the Virtual IP address (VIP), the following message appears:

```
Please enter the CIDR for the Virtual IP without the / (example:  
24) :
```

17. If you are prompted, enter the CIDR for the Virtual IP.
18. You will be prompted for the a confirmation dialog similar to the following:

You have selected the following settings, please confirm if they are correct:

Architecture: DR

This Node: PRIMARY (Node 1)

Node 1 Hostname: *<host name of this appliance>*

Node 1 DRBD IP: *<DRBD IP address you entered for this appliance>*

Node 2 Hostname: *<host name of the Secondary appliance>*

Node 2 DRBD IP: *<DRBD IP address you entered for the Secondary appliance>*

Virtual IP: *<virtual IP>/<CIDR mask>*

DRBD Disk: *< /dev/mapper/em7vg-db>*

DRBD Proxy: *<Yes/No>*

Is this information correct? (y/n)

19. Review the information you entered and enter **yes** if it is correct or **no** to reenter the information.
20. After you enter **yes**, the following message indicating the system is being configured appears:

Setting up the environment...

MariaDB running, sending shutdown

Pausing SL1 services

Adjusting /dev/mapper/em7vg-db by extending to 39464 extents to accommodate filesystem and metadata

- Updating firewalld configuration, please be patient...

Setting up DRBD...

Setting up and starting Corosync...

Setting up and starting Pacemaker...

Waiting on cluster services to come up...

Waiting on cluster services to come up...

Waiting on cluster services to come up...

Cluster services detected as active

Configuring silo.conf for clustering

Unpausing SL1

Setup is complete

Current cluster status

Stack: corosync

Current DC: <hostname primary> (version 1.1.19.linbit-8+20181129.e17.2-c3c624ea3d) - partition

with quorum

Last updated: <date and time of last update>

Last change: <date and time of last change by hacluster via crmd on <hostname primary>>

1 node configured

```
5 resources configured

Online: [ <hostname primary> ]

Active resources:

Resources Group: g_em7

p_fs_drbd1 (ocf::heartbeat:Filesystem): Started <hostname primary>
mysql (ocf::sciencelogic:mysql-systemd): Started <hostname primary>
virtual_ip (ocf::heartbeat:IPaddr2): Started <hostname primary>

Master/Slave Set: ms_drbd_r0 [p_drbd-r0]

Masters: [ db1 ]

Current DRBD status

r0 role:Primary

disk:UpToDate

peer connection:Connecting

silo-cluster-install has exited
```

21. The example messages above display a successful setup. The key messages to monitor a successful setup are:

```
Unpausing SL1

Setup is complete

Current cluster status
```


NOTE: This message indicates that the system has been configured and the required clustering services have started successfully. You can review the output of the cluster status to see if all services have started as expected.

22. If the setup did not complete successfully with the cluster services in a starting state, you will need to review the log to determine the cause of the failure. When you correct the issues the setup can be run again.

Configuring the Secondary Appliance

To configure the Secondary appliance for Disaster Recovery, perform the following steps:

1. Log in to the console of the Secondary appliance as the em7admin user.
2. Run the following command to assume root user privileges:

```
sudo -s
```

3. When prompted, enter the password for the em7admin user.
4. Run the following command:

```
silo-cluster-install
```

The following prompt appears:

```
v3.7
```

```
This wizard will assist you in setting up clustering for  
ScienceLogic appliance. Please be sure to consult the ScienceLogic  
documentation before running this script to ensure that all  
prerequisites have been met.
```

```
1) HA
```

```
2) DR
```

```
3) HA+DR
```

```
4) Quit
```

```
Please select the architecture you'd like to setup:
```

5. Enter option 2. The following prompt appears:

```
Will this system be the primary/active node in a new DR cluster?  
(y/n)
```

6. Enter **no**. The following prompt appears:

```
Has the primary/active system already been configured and is it  
running? (y/n)
```

7. Enter **yes** if you have already configured the primary node and it is active and running. Otherwise, choose **no** and complete the setup on the primary node.

8. After you choose **yes**, the following prompt appears:

```
Enter an IP for the primary system that is reachable by this node:
```

9. Enter the IP address of the primary system and press **[Enter]**. This node will attempt to connect to the primary system and retrieve the cluster configuration information. If you cannot establish a successful connection, you will be prompted again to enter the IP address and additionally a username and a password that you use to connect to MariaDB.

10. After a successful connection is made you will be prompted to confirm the architecture. Review the output and ensure it is the same as the primary.

```
Review the following architecture selections to ensure they are  
correct
```

```
Architecture: DR
```

```
Current node hostname: <hostname>
```

```
Crossover Cable: Yes
```

```
Dedicated link speed: 10Gbps
```

```
Virtual IP (VIP) used: Yes
```

```
DRBD Proxy: No
```

```
Is this architecture correct? (y/n)
```

11. Review the information and enter **yes** if it is correct, or **no** if it is incorrect. If the information is not correct, you will need to regenerate the configuration on the primary/active system before proceeding.

12. After choosing **yes** the architecture will be validated for errors and you will be prompted to confirm the desired configuration. Review your information to ensure it is correct. A similar example follows:

```
You have selected the following settings, please confirm if they
are correct:
```

```
Architecture: DR
```

```
This Node: SECONDARY (Node 2)
```

```
Node 1 Hostname: db1
```

```
Node 1 DRBD: 10.255.255.1
```

```
Node 2 Hostname: db2
```

```
Node 2 DRBD: 10.255.255.2
```

```
Virtual IP: 10.64.167.0/23
```

```
DRBD Disk: /dev/mapper/em7vg-db
```

```
DRBD Proxy: No
```

```
Is this information correct? (y/n)
```

13. Review the information and select **yes** if it is correct and **no** if it is not correct. If the information is incorrect, you will need to regenerate the configuration on the primary/active system before proceeding.
14. After selecting **yes** you will see the following messages indicating the system is being configured.

```
Setting up the environment...
```

```
MariaDB running, sending shutdown
```

```
Pausing SL1 services
```

```
Adjusting /dev/mapper/em7vg-db by extending to 39464 extents to
accommodate filesystems and metadata
```

```
-Updating firewalld configuration, please be patient...
```

```
Setting up DRBD...
Setting up and starting Corosync...
Setting up and starting Pacemaker...
Waiting on cluster services to come up...
Waiting on cluster services to come up...
Waiting on cluster services to come up...
Cluster services detected as active
Configuring silo.conf for clustering
Unpausing SL1
Setup is complete. Please monitor the DRBD synchronization
status (drbdadm status)
Failover cannot occur until DRBD is fully synced

Current cluster status

Stack: corosync

Current DC: <hostname secondary> version 1.1.19.linbit-
8+20181129.el7.2-c3c624ea3d) - partition

with quorum

Last updated: <date and time of last update>

Last change: <date and time of last change by hacluster via crmd
on <hostname secondary>>

1 node configured

4 resources configured (2 disabled)
```

```
Online: [ <hostname secondary> ]
Active resources:
Master/Slave Set: ms_drbd_r0 [p_drbd_r0]
Slaves (target-role): [ <hostname secondary> ]

Current DRBD status

r0 role:Secondary
disk:Inconsistent blocked:lower
peer role:Primary
replication:SyncTarget peer-disk:UpToDate done:11.91

silo-cluster-install has exited
```

Licensing the Secondary Appliance

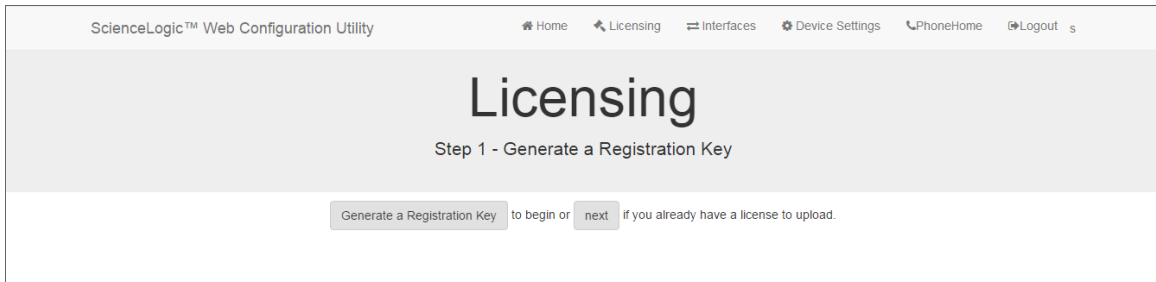
Perform the following steps to license the Secondary appliance:

1. You can log in to the Web Configuration Utility using any web browser supported by SL1. The address of the Web Configuration Utility is in the following format:

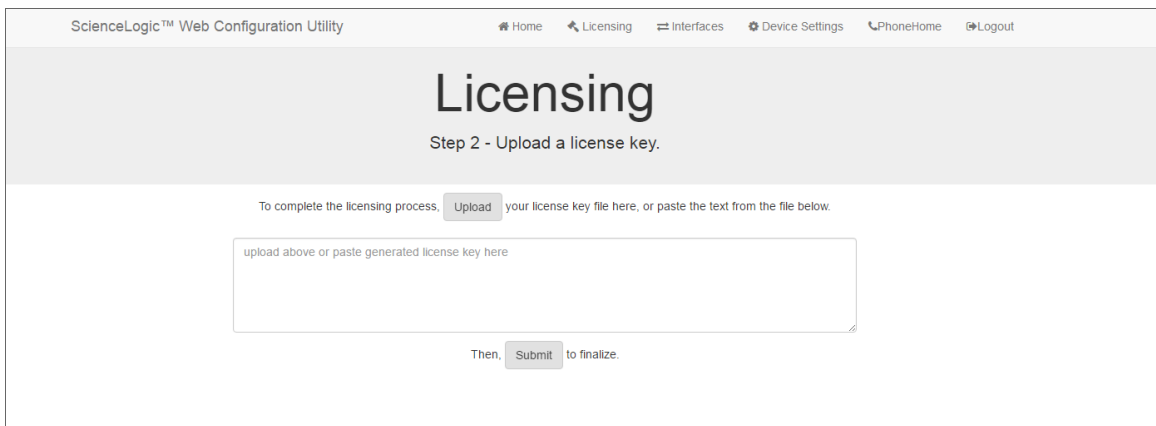
```
https://<ip-address-of-appliance>:7700
```

Enter the address of the Web Configuration Utility in to the address bar of your browser, replacing "*ip-address-of-appliance*" with the IP address of the Secondary appliance.

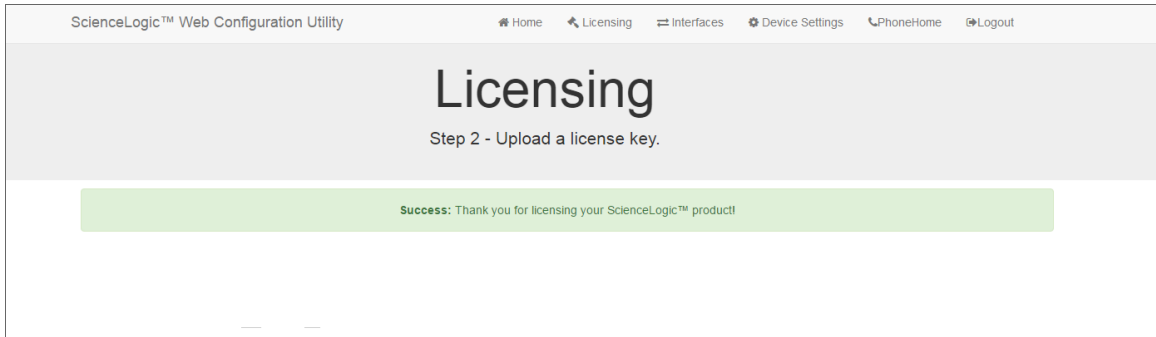
2. You will be prompted to enter your username and password. Log in as the "em7admin" user with the password you configured using the Setup Wizard.
3. The **Configuration Utilities** page appears. Click the **[Licensing]** button. The **Licensing Step 1** page appears:



4. Click the **[Generate a Registration Key]** button.
5. When prompted, save the Registration Key file to your local disk.
6. Log in to the ScienceLogic Support Site at <https://support.sciencelogic.com/s/>. Click the **License Request** tab and follow the instructions for requesting a license key. ScienceLogic will provide you with a License Key file that corresponds to the Registration Key file.
7. Return to the Web Configuration Utility:



8. On the **Licensing Step 2** page, click the **[Upload]** button to upload the license file. After navigating to and selecting the license file, click the **[Submit]** button to finalize the license. The **Success** message appears:



NOTE: Upon login, SL1 will display a warning message if your license is 30 days or less from expiration, or if it has already expired. If you see this message, take action to update your license immediately.

Configuring Data Collection Servers and Message Collection Servers

If you are using a distributed system, you must configure the Data Collectors and Message Collectors to use the new multi-Database Server configuration.


To configure a Data Collector or Message Collector to use the new configuration:

1. You can log in to the Web Configuration Utility using any web browser supported by SL1. The address of the Web Configuration Utility is in the following format:

```
https://<ip-address-of-appliance>:7700
```

Enter the address of the Web Configuration Utility in the address bar of your browser, replacing "*ip-address-of-appliance*" with the IP address of the Data Collector or Message Collector.

2. You will be prompted to enter your user name and password. Log in as the em7admin user with the password you configured using the Setup Wizard.
3. On the **Configuration Utilities** page, click the **[Device Settings]** button. The **Settings** page appears:



4. On the **Settings** page, enter the following:
 - **Database IP Address.** Enter the IP addresses of all the Database Servers, separated by commas.
5. Click the **[Save]** button. You may now log out of the Web Configuration Utility for that collector.
6. Perform steps 1-5 for each Data Collector and Message Collector in your system.

Failover

If your Primary appliance fails, you can manually failover to the Secondary appliance. There are two ways to perform failover:

1. If you can access a shell session on both appliances:
 - Because DRBD does not allow two Primary appliances, you must demote the Primary appliance *first* during failover.
 - After demoting the Primary appliance, your system will recognize two Secondary appliances; DRBD allows two Secondary appliances. You can then promote the original Secondary appliance.
 - After promoting the original Secondary appliance, your system will have one Primary appliance and one Secondary appliance.
 - This process is described in the section [Failover When Both Database Appliances are Accessible](#).
2. If you *cannot* access a shell session on your Primary appliance:
 - **Power down the Primary appliance.** This step is required to avoid a split-brain configuration in which you have two Primary appliances.
 - Promote the Secondary appliance.
 - After promoting the Secondary appliance, your system will have one Primary appliance and one "unknown" appliance.
 - Upon reboot, DRBD will automatically set the "unknown" appliance to "secondary".
 - This process is described in the section [Failover When the Primary Database Appliance is Inaccessible](#).

Failover When Both Database Appliances are Accessible

If you need to perform failover, and you can access a shell session on both Database Servers perform the following steps:

1. Log in to the console of the Primary appliance as the em7admin user.
2. Run the following command to assume root user privileges:

```
sudo -s
```

3. When prompted, enter the password for the em7admin user.
4. Run the following command:

```
coro_config
```

The following prompt appears:

```
1) Enable Maintenance
```

```
2) Option Disabled
```

```
3) Demote DRBD
```

```
4) Stop Pacemaker
```

```
5) Resource Status
```

```
6) Quit
```

```
Please enter the number of your choice:
```

5. Enter "3". The following prompt appears:

```
Node currently Primary, would you like to make it Secondary? (y/n)
```

```
y
```

6. Enter "y". The following output appears:

```
Issuing command: crm_resource --resource ms_drbd_r0 --set-parameter  
target-role --meta --parameter-value Slave
```

7. The Primary database must be shutdown before promoting the Secondary database. To verify MariaDB is shutdown, check the MariaDB log file: `/var/log/mysql/mysqld.log`. The log file should contain the line: "Shutdown complete".
8. Log in to the console of the Secondary appliance as the em7admin user.
9. Run the following command:

```
sudo -s
```

10. When prompted, enter the password for the em7admin user.
11. Execute the following command:

```
coro_config
```

The following prompt appears:

```
1) Enable Maintenance
```

```
2) Option Disabled
```

```
3) Promote DRBD
```

```
4) Stop Pacemaker
```

```
5) Resource Status
```

```
6) Quit
```

```
Please enter the number of your choice:
```

12. Enter "3". The following prompt appears:

```
Node currently Secondary, would you like to make it Primary? (y/n)
```

13. Enter "y". The following output appears:

```
Issuing command: crm_resource --resource ms_drbd_r0 --set-parameter target-role --meta --parameter-value Master
```

14. To verify that an appliance is active after failover, ScienceLogic recommends checking the status of MariaDB, which is one of the primary processes on Database Servers. To verify the status of MariaDB, execute the following command on the newly promoted Primary Database Server:

```
silo_mysql -e "select 1"
```

If MariaDB is running normally, you will see a '1' in the console output.

TIP: Because larger systems can take more time to start the database, verify that MariaDB has started successfully before running the command in the step above. To verify MariaDB has started successfully, check the MariaDB log file: `/var/log/mysql/mysqld.log`. The log file should contain the line: `"/usr/sbin/mysqld: ready for connections."`.

15. If you are using a distributed SL1 system, you must reconfigure all Administration Portals in your system to use the new Database Server. To do this, follow the steps listed in the [Reconfiguring Administration Portals](#) section.
16. When the previously Primary Database Server reboots, it will be the Secondary appliance. Upon reboot, DRBD automatically sets all Database Servers to "secondary". This prevents accidental "split-brain" from occurring.

NOTE: Upon login, if SL1 detects two Primary databases or two Secondary databases, an error message is displayed along with information for how to fix the problem.

NOTE: For more information on manual failover for Disaster Recovery with Two Appliances, see the chapter on [Troubleshooting High Availability and Disaster Recovery Configuration](#).

Failover When the Primary Database Appliance is Inaccessible

If you need to perform a failover, and you cannot access a shell session on the current Primary Database Server, perform the steps in this section.

To failover when the Primary appliance is inaccessible:

1. **Power down the inaccessible Primary Database Server.** This step is required to avoid a split-brain configuration where you have two Primary appliances. A split-brain configuration will cause your data to become corrupted.

2. The primary database must be shutdown before promoting the secondary database. To verify MariaDB is shutdown, check the MariaDB log file: `/var/log/mysql/mysql.log`. The log file should contain the line: "Shutdown complete".
3. Log in to the console of the Secondary appliance as the em7admin user.
4. Run the following command to assume root user privileges:

```
sudo -s
```

5. When prompted, enter the password for the em7admin user.
6. Run the following command:

```
coro_config
```

The following prompt appears:

```
1) Enable Maintenance
```

```
2) Option Disabled
```

```
3) Promote DRBD
```

```
4) Stop Pacemaker
```

```
5) Resource Status
```

```
6) Quit
```

```
Please enter the number of your choice:
```

7. Enter "3". The following prompt appears:

```
Node currently Secondary, would you like to make it Primary? (y/n)
```

8. Enter "y". The following output appears:

```
Issuing command: crm_resource --resource ms_drbd_r0 --set-parameter target-role --meta --parameter-value Master
```

9. To verify that an appliance is active after failover, ScienceLogic recommends checking the status of MariaDB, which is one of the primary processes on Database Servers. To verify the status of MariaDB, execute the following command on the newly promoted Database Server:

```
silo_mysql -e "select 1"
```

If MariaDB is running normally, you will see a '1' in the console output.

TIP: Because larger systems can take more time to start the database, verify that MariaDB started successfully before running the above command. To verify MariaDB started successfully, check the MariaDB log file: `/var/log/mysql/mysqld.log`. The log file should contain the line: `"/usr/sbin/mysqld: ready for connections."`

10. If you are using a distributed SL1 system, you must reconfigure all Administration Portals in your system to use the new Database Server. To do this, follow the steps listed in the [Reconfiguring Administration Portals](#) section.
11. When the previously Primary Database Server reboots, it will be the Secondary appliance. Upon reboot, DRBD automatically sets all Database Servers to "secondary". This prevents accidental "split-brain" from occurring.

NOTE: Upon login, if SL1 detects two Primary databases or two Secondary databases, an error message is displayed along with information for how to fix the problem.

Reconfiguring Administration Portals

If you are using a Distributed system and you did not configure a virtual IP address, you must configure all Administration Portals in your system to use the new Primary Database Server after performing failover or failback. To configure an Administration Portal to use the new Database Server:

You must perform the following steps in the Web Configuration Utility to configure an Administration Portal:

1. You can log in to the Web Configuration Utility using any web browser supported by SL1. The address of the Web Configuration Utility is in the following format:

```
https://<ip-address-of-appliance>:7700
```

Enter the address of the Web Configuration Utility in the address bar of your browser, replacing *ip-address-of-appliance* with the IP address of the Secondary appliance.

2. Log in as the "em7admin" user with the password you configured using the Setup Wizard. The **Configuration Utility** page appears.
3. Click the **[Device Settings]** button. The **Settings** page appears:

4. On the **Settings** page, enter the following:
 - **Database IP Address.** The IP address of the new Primary ScienceLogic Database Server.
5. Click the **[Save]** button. You may now log out of the Web Configuration Utility.
6. Log into the Administration Portal appliance with your SSH credentials and run the following command:

```
sudo systemctl restart nextui;
```

7. Repeat these steps for each Administration Portal in your system.

NOTE: Starting with SL1 version 11.3.0, SL1 will raise a "Collector Outage" event if a Data Collector or Message Collector could not be reached from either a primary or secondary Database Server.

Verifying that a Database Server is Primary

To verify that your network is configured correctly and will allow the newly active Database Server to operate correctly, check the following system functions:

- If you use Active Directory or LDAP authentication, log in to the user interface using a user account that uses Active Directory or LDAP authentication.
- In the user interface, verify that new data is being collected.
- If your system is configured to send notification emails, confirm that emails are being received as expected. To test outbound email, create or update a ticket and ensure that the ticket watchers receive an email.

NOTE: On the **Behavior Settings** page (System > Settings > Behavior, if the field **Automatic Ticketing Emails** is set to *Disabled*, all assignees and watchers will not receive automatic email notifications about any tickets. By default, the field is set to *Enabled*.

- If your system is configured to receive emails, confirm that emails are being received correctly. To test inbound email, send a test email that will trigger a "tickets from Email" policy or an "events from Email" policy.

To complete the verification process, execute the following command on the newly demoted secondary database:

```
sudo systemctl start pacemaker
```

Failback

If you have performed failover and then want to return to the previous configuration, you can perform failback.

Because DRBD does not allow two Primary appliances, you must first demote the Primary appliance during failback. After demoting, your SL1 system will have two Secondary appliances, but DRBD allows two Secondary appliances. You can then promote the Secondary appliance. After promoting the Secondary appliance, your SL1 system will have one Primary appliance and one Secondary appliance.

NOTE: Upon login, if SL1 detects two Primary databases or two Secondary databases, an error message is displayed along with information for how to fix the problem.

To perform failback:

1. Log in to the console of the current Primary appliance as the em7admin user.
2. Check the status of both appliances. To do this, enter the following at the shell prompt:

```
cat /proc/drbd
```

Your output will look like this:

```
1: cs:Connected ro:Primary/Secondary ds:UpToDate/UpToDate C r----  
ns:17567744 al:0 bm:1072 lo:0 pe:0 ua:0 ap:0 ep:1 wo:b  
oos:12521012\  
  
-----
```

To failback safely, the output should include `ro:Primary/Secondary ds:UpToDate/UpToDate`.

NOTE: If your two appliances cannot communicate, your output will include "`ro:Primary/Unknown`" and "`ds:UpToDate/DUnknown`". Before proceeding with failback, troubleshoot and resolve the communication problem.

NOTE: If your output includes "`ro:Primary/Secondary`", but does not include "`UpToDate/UpToDate`", data is being synchronized between the two appliances. You must wait until data synchronization has finished before performing failback.

3. You can now safely perform failback. To perform failback, follow the steps listed in the [Failover When Both Database Appliances are Accessible](#) section of this chapter.



High Availability with Two Appliances

Overview

This chapter describes how to configure two appliances in a High Availability cluster.

This chapter assumes that you are comfortable using a UNIX shell session and can use the basic functions within the vi editor.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon ().
- To view a page containing all of the menu options, click the Advanced menu icon ().

This chapter covers the following topics:

<i>Prerequisites</i>	34
<i>Addressing the Cluster</i>	35
<i>Configuring Heartbeat IP Addresses</i>	35
<i>Testing the Heartbeat Network</i>	36
<i>Deployment Recommendations</i>	37
<i>Configuring High Availability</i>	37
<i>Failover</i>	51

Prerequisites

Before performing the steps listed in this chapter, you must:

- Install and license each appliance
- Have an Administrator account to log in to the Web Configuration Utility for each appliance
- Have SSH or console access to each appliance
- Know the em7admin console username and password for each appliance
- Have identical hardware or virtual machine specifications on each appliance
- Have configured a unique hostname on each appliance
- Know the MariaDB username and password
- Know the MariaDB username and password
- (Option 1) ScienceLogic recommends that you connect the two appliances on a secondary interface with a crossover cable.
 - When using a crossover cable a virtual IP is required for the cluster.
- (Option 2) Connect the two appliances on a secondary interface using a dedicated private network (VLAN). This network should only be used for cluster communication. If a crossover cable is not possible, you must have this private network for cluster communication containing only the two high availability nodes.
 - A virtual IP is optional when using this method.
 - One Admin Portal is required to be configured as a quorum witness.
 - Latency between the nodes must be less than 1 millisecond.
 - Nodes can optionally be placed in separate availability zones, but must be in the same region.
- You can determine the virtual IP address for the cluster, if required. The virtual IP address will be associated with the primary appliance and will be transitioned between the appliances during failover and failback. The virtual IP address must be on the same network subnet as the primary network adapters of the appliances.

NOTE: See [Deployment Recommendations](#) for details regarding availability zones and regions.

Unique Hostnames

You must ensure that a unique hostname is configured on each SL1 appliance. The hostname of an appliance is configured during the initial installation. To view and change the hostname of an appliance:

1. Log in to the console of the SL1 appliance as the em7admin user. The current hostname appears before the command-prompt. For example, the login prompt might look like this, with the current hostname highlighted in bold:

```
login as: em7admin
```

```
em7admin@10.64.68.31's password:
```

```
Last login: Wed Apr 27 21:25:26 2016 from silo1651.sciencelogic.local
```

```
[em7admin@HADB01 ~]$
```

2. To change the hostname, run the following command:

```
sudo hostnamectl set-hostname <new hostname>
```

3. When prompted, enter the password for the em7admin user.

Addressing the Cluster

A database cluster has three IP addresses: one for the primary interface on each database appliance, and an additional virtual IP address. The virtual IP address is shared between the two database appliances, to be used by any system requesting database services from the cluster.

The following table describes which IP address you should supply for the Database Server when you configure other SL1 appliances and external systems:

Appliance/System	IP Address
Administration Portal	Use the Virtual IP address when configuring the Database IP Address in the Web Configuration Utility and the Appliance Manager page (System > Settings > Appliances).
Data Collector or Message Collector	Include both primary interface IP addresses when configuring the ScienceLogic Central Database IP Address in the Web Configuration Utility and the Appliance Manager page (System > Settings > Appliances).
SNMP Monitoring	Monitor each Database Server separately using the primary interface IP addresses.
Database Dynamic Applications	Use the Virtual IP address in the Hostname/IP field in the Credential Editor page (System > Manage > Credentials > wrench icon).

Configuring Heartbeat IP Addresses

To cluster two databases, you must configure a **heartbeat network** between the appliances first. The heartbeat network is used by the databases to determine whether failover conditions have occurred. A heartbeat network consists of a crossover Ethernet cable attached to an interface on each database.

After attaching the network cable, you must complete the steps described in this section to configure the heartbeat network.

Perform the following steps on each appliance to configure the heartbeat network:

1. Log in to the console of the Primary appliance as the em7admin user.
2. Run the following command:

```
sudo /usr/local/bin/slmenu
```

3. Select **Network Configuration**.
4. Select the interface you want to configure.
5. Edit the configuration. Minimally, you will need to set or edit the following entries. Other entries can remain at the default value.

```
IPADDR="169.254.1.1"  
PREFIX="30"  
BOOTPROTO="none"  
ONBOOT="yes"
```

6. After you finish your edits, press **[<tab>]** and select **<Save>**.
7. A prompt to restart networking to activate your changes appears. Select **<Save>**.
8. Log in to the console of the Secondary appliance as the em7admin user and repeat steps 2-5, using 169.254.1.2 as the IP address.

Testing the Heartbeat Network

After you configure the heartbeat network, perform the following steps to test the connection:

1. Log in to the console of the Primary appliance as the em7admin user.
2. Run the following command:

```
ping -c4 169.254.1.2
```

If the heartbeat network is configured correctly, the output looks like this:

```
PING 169.254.1.2 (169.254.1.2) 56(84) bytes of data.
```

```
64 bytes from 169.254.1.2: icmp_seq=1 ttl=64 time=0.657 ms
```

```
64 bytes from 169.254.1.2: icmp_seq=2 ttl=64 time=0.512 ms
```

```
64 bytes from 169.254.1.2: icmp_seq=3 ttl=64 time=0.595 ms
```

```
64 bytes from 169.254.1.2: icmp_seq=4 ttl=64 time=0.464 ms
```

If the heartbeat network is not configured correctly, the output looks like this:

```
PING 169.254.1.2 (169.254.1.2) 56(84) bytes of data.  
From 169.254.1.1 icmp_seq=1 Destination Host Unreachable  
From 169.254.1.1 icmp_seq=2 Destination Host Unreachable  
From 169.254.1.1 icmp_seq=3 Destination Host Unreachable  
From 169.254.1.1 icmp_seq=4 Destination Host Unreachable
```

Deployment Recommendations

ScienceLogic strongly recommends that you use a 10 GB network connection for the cluster communication link. Using a 1 GB link may not be suitable and lead to reduced database performance.

When deploying in a cloud environment the quorum witness must be placed in a different availability zone from the High Availability servers.

When deploying in a cloud environment the High Availability servers should be placed in separate availability zones and must not be in the same zone with the quorum witness.

Configuring High Availability

This section describes how to configure the Quorum Witness, Primary appliance, and the Secondary appliance for High Availability.

Configuring the Quorum Witness

A quorum witness is required in any setup that does not use a crossover cable. Two-node clusters without a crossover cable are known to be unstable. The quorum witness provides the needed "tie-breaking" vote to avoid split-brain scenarios.

Only one administration portal can be configured as a quorum witness. Please review the requirements for network placement and availability zones.

Prior to starting this procedure ensure the Admin Portal is connected to the primary database and licensed. To configure the quorum witness:

1. Log in to the console of the administrative portal appliance that will be the quorum witness as the em7admin user.
2. Run the following command:

```
sudo -s
```

3. Enter the em7admin password when you are prompted.

4. Run the following command:

```
silo-quorum-install
```

The following prompt appears:

```
This wizard will assist you in configuring this admin portal to be a
quorum device for a corosync/pacemaker HA or HA+DR cluster. When
forming a cluster that does not have a cross over cable a quorum
device is required to prevent split-brain scenarios. Please consult
ScienceLogic documentation on the requirements for a cluster setup.
```

```
Do you want to configure this node a quorum device for an HA or HA+DR
cluster? (y/n)
```

5. Enter **yes** and the following prompt appears:

```
-Please enter the PRIMARY IP address for the Primary HA server:
```

6. Enter the primary IP of the primary high availability server and press **Enter**. The following prompt will appear:

```
-Please enter the PRIMARY IP address for the Secondary HA server:
```

7. Enter the primary IP of the secondary high availability server and press **Enter**. The following confirmation will appear:

```
Please check the following IP addresses are correct
```

```
Node 1 Primary IP: <primary HA primary IP>
```

```
Node 2 Primary IP: <secondary HA primary IP>
```

```
Is this architecture correct? (y/n)
```

8. Enter **yes** if your architecture is correct and the following messages will indicate the quorum witness is being setup.

```
Updating firewalld configuration, please be patient...
```

```
Waiting for quorum service to start...
```

```
Service has started successfully.
```

```
silo-quorum-install has exited
```

After this setup is complete you can configure the primary appliance.

Configuring the Primary Appliance

To configure the Primary appliance for High Availability, perform the following steps:

1. Log in to the console of the Primary appliance as the em7admin user.
2. Run the following command:

```
sudo -s
```

3. When prompted, enter the password for the em7admin user.
4. Run the following command:

```
silo-cluster-install
```

The following prompt appears:

```
v3.7
```

```
This wizard will assist you in setting up clustering for  
ScienceLogic appliance. Please be sure to consult the ScienceLogic  
documentation before running this script to ensure that all  
prerequisites have been met.
```

```
1) HA
```

```
2) DR
```

```
3) HA+DR
```

```
4) Quit
```

```
Please select the architecture you'd like to setup:
```

5. Enter "1". The following prompt appears:

```
Will this system be the primary/active node in a new HA cluster?  
(y/n)
```

6. Enter **yes**. The following prompt appears:

```
Is there a cross-over cable between HA nodes? (y/n)
```

7. Enter **yes** if there is a physical crossover cable connecting the HA nodes and skip to step 10. Otherwise answer **no**. When answering **no**, the following prompt appears:

```
Is there a dedicated/private network link between the HA nodes?  
(y/n)
```

8. A dedicated private network link is required. If you answer **no**, the script will exit. Answer **yes** and the following prompt appears:

```
Have you already configured an AP to act as a quorum device? (y/n)
```

9. You must have already configured an Admin Portal to act as a quorum witness. If you answer **no**, the script will exit. Answer **yes** and the following prompt appears:

```
Please enter the IP for the quorum device:
```

10. Enter the IP address for the Admin Portal that was configured for as the quorum witness and press **Enter**. The following prompt appears:

```
Choose the closest value for the speed of the dedicated between HA nodes.
```

```
1) 10Gbps
```

```
2) 1 Gbps
```

```
What is the physical link speed between HA nodes?
```

11. Select the response that most closely matches the speed of the private link between HA nodes.
12. If you are using a crossover cable skip to step 14. If you're not using a crossover cable, the following prompt appears:

```
Would you like to use a Virtual IP (VIP)? (y/n)
```

13. Enter **yes** if you want to use a VIP, otherwise, enter **no**.

14. The following confirmation will appear:

```
Review the following architecture selections to ensure they are correct
```

```
Architecture: HA
```

```
Current node hostname: db1
```

```
Crossover Cable: No
```

```
Dedicated link speed: <the value of the speed you chose>
```

```
Quorum Node IP: <IP address of your quorum witness>
```

```
Virtual IP (VIP) used: Yes
```

```
Is this architecture correct? (y/n)
```

15. Verify that your architecture is correct. If not, enter **no** to reenter the data, otherwise enter **yes**. The following prompt appears:


```
Primary node information:
```

```
Please the IP used for HEARTBEAT traffic for this server:
```

```
1)10.255.255.1
```

```
...
```

```
Number:
```

16. Select the IP and the following prompt appears:

```
Secondary node information:
```

```
What is the hostname of the Secondary HA server:
```

17. Enter the hostname of the secondary server press **[Enter]**. The following prompt appears:

```
Please enter the IP used for HEARTBEAT traffic for the Secondary HA  
server that corresponds to 10.255.255.1
```

18. Enter the corresponding IP and press **[Enter]**. The following prompt appears:

```
Please enter the PRIMARY IP address for the Secondary HA server:
```

19. Enter the primary IP for the secondary high availability server and press **[Enter]**.

20. If you are using a VIP, the following prompt appears:

```
Virtual IP address information:
```

```
Please enter the Virtual IP Address:
```

```
Please enter the CIDR for the Virtual IP without the / (example: 24):
```

21. Enter the VIP address and CIDR mask.

22. The following confirmation will appear:

```
You have selected the following settings, please confirm if they are  
correct:
```

```
Architecture: HA
```

```
This Node: PRIMARY (Node 1)
```

```
Node 1 Hostname: db1
```

```
Node 1 DRBD/Heartbeat IP: 10.255.255.1
```

```
Node 1 Primary IP: 10.64.166.251
```

```
Node 2 Hostname: tea
```

```
Node 2 DRBD/Heartbeat IP: 10.225.255.2
```

```
Node 2 Primary IP: 10.64.166.252
```

```
Virtual IP: 10.64.167.0/23
```

```
DRBD Disk: /dev/mapper/em7vg-db
```

```
DRBD Proxy: No
```

```
Is this information correct? (y/n)
```

23. Review your information and ensure it is correct. Enter **yes**, if correct.
24. Once confirmed, the system will configure itself and you can expect an output similar to the following:

```
Setting up the environment...
```

```
MariaDB running, sending shutdown
```

```
Pausing SLI services
```

```
Adjusting /dev/mapper/em7vg-db by extending to 39464 extents to  
accommodate filesystem and metadata
```

```
-Updating firewalld configuration, please be patient...
```

```
Setting up DRBD...
```

```
Setting up and starting Corosync...
```

```
Setting up and starting Pacemaker...
```

```
Waiting on cluster services to come up...
```

```
Waiting on cluster services to come up...
```

```
Waiting on cluster services to come up...
```

```
Cluster services detected as active
```

```
Configuring silo.conf for clustering
```

```
Unpausing SL1
```

```
Setup is complete
```

```
Current cluster status
```

```
Stack: corosync
```

```
Current DC: db1 (version 1.1.19.linbit-8+20181129.el7.2-c3c624ea3d) -  
partition with quorum
```

```
Last updated: <date and time you last updated>
```

```
Last change: <date and time you last changed> by hacluster via crmd  
on db1
```

```
2 nodes configured
```

```
5 resources configured
```

```
Online: [ db1 ]
```

```
OFFLINE: [ db2 ]
```

```
Active resources:
```

```
Resource Group: g_em7
```

```
  p_fs_drbd1 (ocf::heartbeat:Filesystem): Started db1
```

```
  mysql (ocf::sciencelogic:mysql-systemd): Started db1
```

```
  virtual_ip (ocf::heartbeat:IPaddr2): Started db1
```

```
Master/Slave Set: ms_drbd_r0 [p_drbd_r0]
```

```
Masters: [ db1 ]
```

```
Current DRBD status
```

```
r0 role:Primary
```

```
disk:UpToDate
```

```
peer connection:Connecting
```

```
silo-cluster-install has exited
```

25. The messages above are examples of a successful setup. The important messages that indicate a successful setup are:

```
Unpausing SL1
```

```
Setup is complete
```

```
Current cluster status
```

This message indicates that the system has been configured and the required clustering services have started successfully. You can review the output of the cluster status to see if all services have started as expected.

26. If your setup did not complete successfully with the cluster services in a starting state, you will need to review the log to determine the cause of the failure. Once you correct the issues found in the log the set up can be run again.

Configuring the Secondary Appliance

To configure the Secondary appliance for High Availability, perform the following steps:

1. Log in to the console of the Secondary appliance as the em7admin user.
2. Run the following command to assume root user privileges:

```
sudo -s
```

3. When prompted, enter the password for the em7admin user.
4. Run the following command:

```
silo-cluster-install
```

The following prompt appears:

```
v3.7
```

```
This wizard will assist you in setting up clustering for a  
ScienceLogic appliance. Please be sure to consult the ScienceLogic  
documentation before running this script to ensure that all  
prerequisites have been met.
```

```
1) HA
```

```
2) DR
```

```
3) HA+DR
```

```
4) Quit
```

```
Please select the architecture you'd like to setup:
```

5. Enter "1". The following prompt appears:

```
Will this system be the primary/active node in a new HA cluster?  
(y/n)
```

6. Enter **no**. The following prompt appears:

```
Has the primary/active system already been configured and is it  
running? (y/n)
```

7. Enter **yes** if you have already configured the primary node and it is active and running. Otherwise, choose **no** and complete the setup on the primary node. The following prompt appears:

8. After entering **yes** the following prompt appears:

```
Enter an IP for the primary system that is reachable by this node:
```

9. Enter the IP address of the primary system and press **[Enter]**. This node will attempt to connect to the primary system and retrieve the cluster configuration information. If a successful connection can't be established, you will be prompted again to enter the IP address and a username and a password. The username and password should be the same used for connecting to Maria DB.

10. When a successful connection is made you will be prompted to confirm the architecture to ensure that it is the same for the primary server.

```
Review the following architecture selections to ensure they are  
correct
```

```
Architecture: HA
```

```
Current node hostname: db2
```

```
Crossover Cable: No
```

```
Dedicated link speed: 10Gbps
```

```
Quorum Node IP: 10.64.166.250
```

```
Virtual IP (VIP) used: Yes
```

```
Is this architecture correct? (y/n)
```

11. Review that the information is correct and enter **yes**. Enter **no** if it is incorrect and regenerate the configuration on the primary/active system before proceeding.
12. After you enter **yes**, the architecture will be validated for any errors and you will be prompted to confirm the configuration. Review that the information is correct. Your output should be similar to the following:

```
You have selected the following settings, please confirm if they  
are correct:
```

```
Architecture: HA
```

```
This Node: SECONDARY (Node 2)
```

```
Node 1 Hostname: db1
```

```
Node 1 DRBD/Heartbeat IP: 10.255.255.1
```

```
Node 1 Primary IP: 10.64.166.251
```

```
Node 2 Hostname: db2
```

```
Node 2 DRBD/Heartbeat IP: 10.255.255.2
```

```
Node 2 Primary IP: 10.64.166.252
```

```
Virtual IP: 10.64.167.0/23
```

```
DRBD Disk: /dev/mapper/em7vg-db
```

```
DRBD Proxy: No
```

```
Is this information correct? (y/n)
```

13. Review that the information is correct and enter **yes**. Enter **no** if it is incorrect and regenerate the configuration on the primary/active system before proceeding.
14. After you enter **yes**, the following message indicates the system is being configured:

```
Setting up the environment...

Maria DB running, sending shutdown

Pausing SL1 services

Adjusting /dev/mapper/em7vg-db by extending to 39464 extents to
accommodate filesystems and metadata

  -Updating firewalld configuration, please be patient...

Setting up DRBD...

Setting up and starting Corosync...

Setting up starting Pacemaker...

Waiting on cluster services to come up...

Waiting on cluster services to come up...

Waiting on cluster services to come up...

Cluster services detected as active

Configuring silo.conf for clustering

Unpausing SL1

Setup is complete. Please monitor the DRBD synchronization status
(drbdadm status)

Failover cannot occur until DRBD is fully synced

Current cluster status

Stack: corosync

Current DC: db1(version 1.1.19.linbit-8+20181129.e17.2-c3c624ea3d)
- partition with quorum
```

```
Last updated:<date and time you last updated>

Last change:<date and time you last changed> by hacluster via crmd
on db1

2 nodes configured

5 resources configured

Online: [ db1 db2 ]

Active resources:

Resource Group: g_em7

p_fs_drdb1 (ocf::heartbeat:Filesystem): Started db1

mysql (ocf::sciencelogic:mysql-systemd): Started db1

virtual_ip (ocf::heartbeat:IPaddr2): Started db1

Master/Slave Set: ms_drbd_r0 [p_drbd_r0]

Masters: [ db1 ]

Slaves: [ db2 ]

Current DRBD status

r0 role:Secondary

disk:Inconsistent

peer role:Primary

replication:SyncTarget peer-disk:UpToDate done:8.91
```



```
sil-cluster-install has exited
```

Licensing the Secondary Appliance

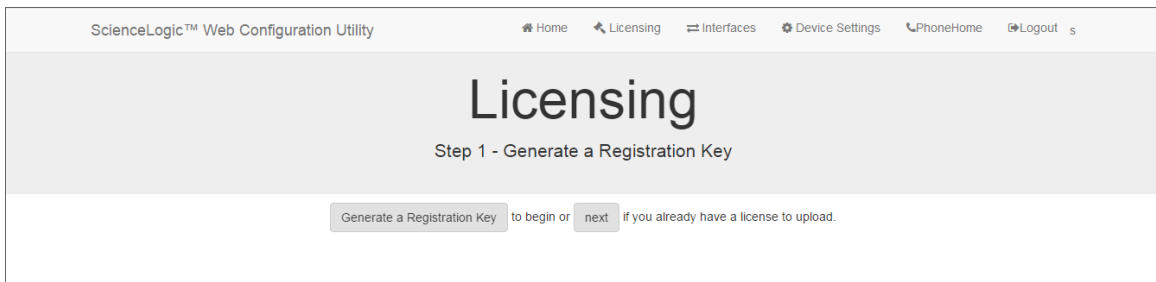
Perform the following steps to license the Secondary appliance:

1. You can log in to the Web Configuration Utility using any web browser supported by SL1. The address of the Web Configuration Utility is in the following format:


```
https://<ip-address-of-appliance>:7700
```

Enter the address of the Web Configuration Utility in to the address bar of your browser, replacing "*ip-address-of-appliance*" with the IP address of the Secondary appliance.


2. You will be prompted to enter your username and password. Log in as the "em7admin" user with the password you configured using the Setup Wizard.
3. The **Configuration Utilities** page appears. Click the **[Licensing]** button. The **Licensing Step 1** page appears:



4. Click the **[Generate a Registration Key]** button.
5. When prompted, save the Registration Key file to your local disk.
6. Log in to the ScienceLogic Support Site at <https://support.sciencelogic.com/s/>. Click the **License Request** tab and follow the instructions for requesting a license key. ScienceLogic will provide you with a License Key file that corresponds to the Registration Key file.
7. Return to the Web Configuration Utility:



8. On the **Licensing Step 2** page, click the **[Upload]** button to upload the license file. After navigating to and selecting the license file, click the **[Submit]** button to finalize the license. The **Success** message appears:



NOTE: Upon login, SL1 will display a warning message if your license is 30 days or less from expiration, or if it has already expired. If you see this message, take action to update your license immediately.

Configuring Data Collection Servers and Message Collection Servers

If you are using a distributed system, you must configure the Data Collectors and Message Collectors to use the new multi-Database Server configuration.

To configure a Data Collector or Message Collector to use the new configuration:

1. You can log in to the Web Configuration Utility using any web browser supported by SL1. The address of the Web Configuration Utility is in the following format:

```
https://<ip-address-of-appliance>:7700
```

Enter the address of the Web Configuration Utility in the address bar of your browser, replacing "*ip-address-of-appliance*" with the IP address of the Data Collector or Message Collector.

2. You will be prompted to enter your user name and password. Log in as the em7admin user with the password you configured using the Setup Wizard.
3. On the **Configuration Utilities** page, click the **[Device Settings]** button. The **Settings** page appears:

ScienceLogic™ Web Configuration Utility

Home Licensing Interfaces Device Settings PhoneHome Logout

Settings

Configure your appliance.

Appliance Type
Collection Unit

Database IP Address
10.0.9.213

Save

4. On the **Settings** page, enter the following:
 - **Database IP Address.** Enter the IP addresses of all the Database Servers, separated by commas.
5. Click the **[Save]** button. You may now log out of the Web Configuration Utility for that collector.
6. Perform steps 1-5 for each Data Collector and Message Collector in your system.

NOTE: Starting with SL1 version 11.3.0, SL1 will raise a "Collector Outage" event if a Data Collector or Message Collector could not be reached from either a primary or secondary Database Server.

Failover

Failover is the process by which database services are transferred from the active database to the passive database. You can manually perform failover for testing purposes. If the active database stops responding to the secondary database over both network paths, SL1 will automatically perform failover.

After failover completes successfully, the previously active database is now passive, and the previously passive database is now active. There is no automatic failback process; the newly active database will remain active until a failure occurs, or failover is performed manually.

Manual Failover for High Availability Clusters

To manually failover a High Availability cluster, perform the following steps:

1. Log in to the console of the Primary appliance as the em7admin user.

NOTE: Upon login, SL1 will display a warning message if your license is 30 days or less from expiration, or if it has already expired. If you see this message, take action to update your license immediately.

2. First, you should check the status of the appliances. To do this, enter the following at the shell prompt:

```
cat /proc/drbd
```

Your output will look like this:

```
1: cs:Connected ro:Primary/Secondary ds:UpToDate/UpToDate C r----  
ns:17567744 al:0 bm:1072 lo:0 pe:0 ua:0 ap:0 ep:1 wo:b oos:12521012
```

To failover safely, the output should include "`ro:Primary/Secondary ds:UpToDate/UpToDate`".

NOTE: If your appliances cannot communicate, your output will include "`ro:Primary/Unknown ds:UpToDate/DUnknown`". Before proceeding with failback, troubleshoot and resolve the communication problem.

NOTE: If your output includes "`ro:Primary/Secondary`", but does not include "`UpToDate/UpToDate`", data is being synchronized between the two appliances. You must wait until data synchronization has finished before performing failback.

3. Run the following command, assuming root user privileges:

```
sudo systemctl stop pacemaker
```

4. When prompted, enter the password for the em7admin user.
5. The Primary database must be shutdown before promoting the Secondary database. To verify MariaDB is shutdown, check the MariaDB log file: `/var/log/mysql/mysqlld.log`. The log file should contain the line: "Shutdown complete".
6. To verify all services have started on the newly promoted primary database, run the following command:

```
crm_mon
```

If all services have started, each one should be marked: "started".

7. After all services have started on the newly promoted primary database, run the following command to complete the failover process: `sudo systemctl start pacemaker`
8. To verify that there are two nodes in the cluster, run the following command:

```
crm_mon
```

NOTE: For more information about manual failover for High Availability clusters, see the chapter on [Troubleshooting High Availability and Disaster Recovery Configuration](#).

Verifying that a Database Server is Active

To verify that an appliance is active after failover, ScienceLogic recommends checking the status of MariaDB, which is one of the primary processes on Database Servers.

To verify the status of MariaDB, run the following command on the newly promoted Database Server:

```
silos_mysql -e "select 1"
```

If MariaDB is running normally, you will see a '1' in the console output.

To verify that your network is configured correctly and will allow the newly active Database Server to operate correctly, check the following system functions:

- If you use Active Directory or LDAP authentication, log in to the user interface using a user account that uses Active Directory or LDAP authentication.
- In the user interface, verify that new data is being collected.
- If your system is configured to send notification emails, confirm that emails are being received as expected. To test outbound email, create or update a ticket and ensure that the ticket watchers receive an email.

NOTE: On the **Behavior Settings** page (System > Settings > Behavior, if the field **Automatic Ticketing Emails** is set to *Disabled*, all assignees and watchers will not receive automatic email notifications about any tickets. By default, the field is set to *Enabled*.

- If your system is configured to receive emails, confirm that emails are being received correctly. To test inbound email, send a test email that will trigger a "tickets from Email" policy or an "events from Email" policy.

To complete the verification process, execute the following command on the newly demoted secondary database:

```
sudo systemctl start pacemaker
```



High Availability and Disaster Recovery with Three Appliances

Overview

This chapter describes how to troubleshoot various scenarios of High Availability clusters configured for Disaster Recovery.

This chapter assumes that you are comfortable using a UNIX shell session and can use the basic functions within the vi editor.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon ().
- To view a page containing all of the menu options, click the Advanced menu icon ().

This chapter covers the following topics:

<i>Prerequisites</i>	56
<i>Licensing DRBD Proxy</i>	57
<i>Addressing the Cluster</i>	57
<i>Reconfiguring an Existing High Availability System</i>	58
<i>Configuring a Heartbeat IP Addresses</i>	58
<i>Testing the Heartbeat Network</i>	59
<i>Deployment Recommendations</i>	60
<i>Configuring Three Appliances for High Availability and Disaster Recovery</i>	60
<i>Failover</i>	79

Prerequisites

Before performing the steps listed in this chapter, you must:

- Install and license each appliance
- Have an Administrator account to log in to the Web Configuration Utility for each appliance
- Have SSH or console access to each appliance
- Know the em7admin console username and password for each appliance
- Have identical hardware or virtual machine specifications on each appliance
- Have configured a unique hostname on each appliance
- Know the MariaDB username and password
- (Option 1) ScienceLogic recommends that you connect the two appliances on a secondary interface with a crossover cable.
 - When using a crossover cable a virtual IP is required for the cluster.
- (Option 2) Connect the two appliances on a secondary interface using a dedicated private network (VLAN). This network should only be used for cluster communication. If a crossover cable is not possible, you must have this private network for cluster communication containing only the two high availability nodes.
 - A virtual IP is optional when using this method.
 - One Admin Portal is required to be configured as a quorum witness.
 - Latency between the nodes must be less than 1 millisecond.
 - Nodes can optionally be placed in separate availability zones, but must be in the same region.
- You can determine the virtual IP address for the cluster, if required. The virtual IP address will be associated with the primary appliance and will be transitioned between the appliances during failover and failback. The virtual IP address must be on the same network subnet as the primary network adapters of the appliances.
- Request and configure a DRBD proxy license
- Know the maximum link speed, in megabytes/second, between the High Availability cluster and the Disaster Recovery appliance.

NOTE: See [Deployment Recommendations](#) for details regarding availability zones and regions.

Unique Host Names

You must ensure that a unique hostname is configured on each SL1 appliance. The hostname of an appliance is configured during the initial installation. To view and change the hostname of an appliance:

1. Log in to the console of the SL1 appliance as the em7admin user. The current hostname appears before the command-prompt. For example, the login prompt might look like this, with the current hostname highlighted in bold:


```
login as: em7admin
em7admin@10.64.68.31's password:
Last login: Wed Apr 27 21:25:26 2016 from silo1651.sciencelogic.local
[em7admin@HADB01 ~]$
```

2. To change the hostname, run the following command:

```
sudo hostnamectl set-hostname <new hostname>
```

3. When prompted, enter the password for the em7admin user.

Licensing DRBD Proxy

DRBD Proxy buffers all data between the active and redundant appliances to compensate for any bandwidth limitations. In addition, DRBD compresses and encrypts the data sent from the active appliance to the redundant appliance.

You must use DRBD Proxy if you are:

- Configuring three appliances for High Availability and Disaster Recovery.
- Configuring two appliances for Disaster Recovery and will not be configuring a direct connection between your appliances with a crossover cable.

To license DRBD Proxy, copy the **drbd-proxy.license** file to the **/etc** directory on all appliances in your system.

Addressing the Cluster

A database cluster has three IP addresses: one for the primary interface on each database appliance, and an additional virtual IP address. The virtual IP address is shared between the two database appliances, to be used by any system requesting database services from the cluster.

The following table describes which IP address you should supply for the Database Server when you configure other SL1 appliances and external systems:

Appliance/System	IP Address
Administration Portal	Use the Virtual IP address when configuring the Database IP Address in the Web Configuration Utility and the Appliance Manager page (System > Settings > Appliances).

Data Collector or Message Collector	Include both primary interface IP addresses when configuring the ScienceLogic Central Database IP Address in the Web Configuration Utility and the Appliance Manager page (System > Settings > Appliances).
SNMP Monitoring	Monitor each Database Server separately using the primary interface IP addresses.
Database Dynamic Applications	Use the Virtual IP address in the Hostname/IP field in the Credential Editor page (System > Manage > Credentials > wrench icon).

Reconfiguring an Existing High Availability System

If you have previously configured two appliances in a High Availability cluster, perform the following steps to reconfigure the appliances as a High Availability cluster plus a Disaster Recovery appliance:

WARNING: During the reconfiguration procedure, the SL1 System will be unavailable. This procedure must be performed during a maintenance window.

1. Validate that the existing High Availability cluster and the third new appliance meet all the [Prerequisites](#) for configuring High Availability plus Disaster Recovery.
2. Log in to the console of the current secondary appliance as the em7admin user.
3. Run the following commands, entering the password for the em7admin user when prompted:

```
sudo service pacemaker stop
```

```
sudo service corosync stop
```

4. Log in to the console of the current Primary appliance as the em7admin user.
5. Run the following commands, entering the password for the em7admin user when prompted:

```
sudo service pacemaker stop
```

```
sudo service corosync stop
```

6. Perform the steps listed in the [Configuring Three Appliances for HA and DR](#) section.

Configuring a Heartbeat IP Addresses

To cluster two databases, you must configure a **heartbeat network** between the appliances first. The heartbeat network is used by the databases to determine whether failover conditions have occurred. A heartbeat network consists of a crossover Ethernet cable attached to an interface on each database.

After attaching the network cable, you must complete the steps described in this section to configure the heartbeat network.

Perform the following steps on each appliance to configure the heartbeat network:

1. Log in to the console of the Primary appliance as the em7admin user.
2. Run the following command:

```
sudo /usr/local/bin/slmenu
```

3. Select **Network Configuration**.
4. Select the interface you want to configure.
5. Edit the configuration. Minimally, you will need to set or edit the following entries. Other entries can remain at the default value.

```
IPADDR="169.254.1.1"  
PREFIX="30"  
BOOTPROTO="none"  
ONBOOT="yes"
```

6. After you finish your edits, press **[<tab>]** and select **<Save>**.
7. A prompt to restart networking to activate your changes appears. Select **<Save>**.
8. Log in to the console of the Secondary appliance as the em7admin user and repeat steps 2-5, using 169.254.1.2 as the IP address.

Testing the Heartbeat Network

After you configure the heartbeat network, perform the following steps to test the connection:

1. Log in to the console of the Primary appliance as the em7admin user.
2. Run the following command:

```
ping -c4 169.254.1.2
```

If the heartbeat network is configured correctly, the output looks like this:

```
PING 169.254.1.2 (169.254.1.2) 56(84) bytes of data.
```

```
64 bytes from 169.254.1.2: icmp_seq=1 ttl=64 time=0.657 ms
```

```
64 bytes from 169.254.1.2: icmp_seq=2 ttl=64 time=0.512 ms
```

```
64 bytes from 169.254.1.2: icmp_seq=3 ttl=64 time=0.595 ms
```

```
64 bytes from 169.254.1.2: icmp_seq=4 ttl=64 time=0.464 ms
```

If the heartbeat network is not configured correctly, the output looks like this:

```
PING 169.254.1.2 (169.254.1.2) 56(84) bytes of data.  
From 169.254.1.1 icmp_seq=1 Destination Host Unreachable  
From 169.254.1.1 icmp_seq=2 Destination Host Unreachable  
From 169.254.1.1 icmp_seq=3 Destination Host Unreachable  
From 169.254.1.1 icmp_seq=4 Destination Host Unreachable
```

Deployment Recommendations

ScienceLogic strongly recommends that you use a 10 GB network connection for the cluster communication link. Using a 1 GB link may not be suitable and lead to reduced database performance.

When deploying in a cloud environment the quorum witness must be placed in a different availability zone from the High Availability servers.

When deploying in a cloud environment the High Availability servers should be placed in separate availability zones and must not be in the same zone with the quorum witness.

Disaster Recovery replication is sensitive to packet-loss and the variations in time delay between when a signal is transmitted and when it's received over a network connection (network jitter). Use Quality of Service (QoS) to classify the Disaster Recovery network traffic to prevent network drops and ensures reliable transmission for replication traffic.

Configuring Three Appliances for High Availability and Disaster Recovery

To configure three appliances for High Availability and Disaster Recovery, you must configure the appliances in the following order:

1. [The Quorum Witness](#)
2. [Primary appliance in the High Availability cluster](#)
3. [Secondary appliance in the High Availability cluster](#)
4. [Disaster recovery appliance](#)

Configuring the Quorum Witness

A quorum witness is required in any setup that does not use a crossover cable. Two-node clusters without a crossover cable are known to be unstable. The quorum witness provides the needed "tie-breaking" vote to avoid split-brain scenarios.

Only one administration portal can be configured as a quorum witness. Please review the requirements for network placement and availability zones.

Prior to starting this procedure ensure the Admin Portal is connected to the primary database and licensed. To configure the quorum witness:

1. Log in to the console of the Admin Portal appliance that will be the quorum witness as the em7admin user.

2. Run the following command:

```
sudo -s
```

3. Enter the em7admin password when you are prompted.

4. Run the following command:

```
silo-quorum-install
```

The following prompt appears:

```
This wizard will assist you in configuring this admin portal to be a quorum device for a corosync/pacemaker HA or HA+DR cluster. When forming a cluster that does not have a cross over cable a quorum device is required to prevent split-brain scenarios. Please consult ScienceLogic documentation on the requirements for a cluster setup.
```

```
Do you want to configure this node a quorum device for an HA or HA+DR cluster? (y/n)
```

5. Enter **yes** and the following prompt appears:

```
-Please enter the PRIMARY IP address for the Primary HA server:
```

6. Enter the primary IP of the primary high availability server and press **Enter**. The following prompt will appear:

```
-Please enter the PRIMARY IP address for the Secondary HA server:
```

7. Enter the primary IP of the secondary high availability server and press **Enter**. The following confirmation will appear:

```
Please check the following IP addresses are correct
```

```
Node 1 Primary IP: <primary HA primary IP>
```

```
Node 2 Primary IP: <secondary HA primary IP>
```

```
Is this architecture correct? (y/n)
```

8. Enter **yes** if your architecture is correct and the following messages will indicate the quorum witness is being setup.

```
Updating firewalld configuration, please be patient...
```

```
Waiting for quorum service to start...
```

```
Service has started successfully.
```

```
sil0-quorum-install has exited
```

After this setup is complete you can configure the primary appliance.

Configuring the Primary High Availability Appliance

To configure the Primary High Availability appliance for High Availability, perform the following steps:

1. Log in to the console of the Primary High Availability appliance as the em7admin user.
2. Run the following command:

```
sudo -s
```

3. When prompted, enter the password for the em7admin user.
4. Run the following command:

```
sil0-cluster-install
```

The following prompt appears:

```
v3.7
```

```
This wizard will assist you in setting up clustering for  
ScienceLogic appliance. Please be sure to consult the ScienceLogic  
documentation before running this script to ensure that all  
prerequisites have been met.
```

```
1) HA
```

```
2) DR
```

```
3) HA+DR
```

```
4) Quit
```

```
Please select the architecture you'd like to setup:
```

5. Enter "3". The following prompt appears:

```
Will this system be the primary/active node in a new HA+DR cluster?  
(y/n)
```

6. Enter **yes**. The following prompt appears:

```
Is there a cross-over cable between HA nodes? (y/n)
```

7. Enter **yes** if there is a physical crossover cable connecting the HA nodes and skip to step 10. Otherwise answer **no**. When answering **no**, the following prompt appears:

```
Is there a dedicated/private network link between the HA nodes?  
(y/n)
```

8. A dedicated private network link is required. If you answer **no**, the script will exit. Answer **yes** and the following prompt appears:

```
Have you already configured an AP to act as a quorum device? (y/n)
```

9. You must have already configured an Admin Portal to act as a quorum witness. If you answer **no**, the script will exit. Answer **yes** and the following prompt appears:

```
Please enter the IP for the quorum device:
```

10. Enter the IP address for the Admin Portal that was configured for as the quorum witness and press **Enter**. The following prompt appears:

```
Choose the closest value for the speed of the dedicated between HA  
nodes.
```

```
1) 10Gbps
```

```
2) 1 Gbps
```

```
What is the physical link speed between HA nodes?
```

11. Select the response that most closely matches the speed of the private link between HA nodes.
12. The following confirmation will appear:

```
Review the following architecture selections to ensure they are  
correct
```

```
Architecture: HA+DR
```

```
Current node hostname: <hostname>
```

```
Crossover Cable: No
```

```
Dedicated link speed: <the value of the speed you chose>
```

```
Quorum Node IP: <IP address of your quorum witness>
```

```
Virtual IP (VIP) used: Yes
```

```
DRBD Proxy: yes
```

```
Is this architecture correct? (y/n)
```

13. Verify that your architecture is correct. If not, enter **no** to reenter the data, otherwise enter **yes**. The following prompt appears:

```
Primary node information:
```

```
Please the IP used for HEARTBEAT traffic for this server:
```

```
1)10.255.255.1
```

```
...
```

```
Number:
```

14. Select the IP and the following prompt appears:

```
Secondary node information:
```

```
What is the hostname of the Secondary HA server:
```

15. Enter the hostname of the secondary server and press **[Enter]**. The following prompt appears:

```
Please enter the IP used for HEARTBEAT traffic for the secondary high availability server that corresponds to 10.255.255.1:
```

16. Enter the corresponding IP and press **[Enter]**. The following prompt appears:

```
Please enter the PRIMARY IP address for the Secondary HA server:
```

17. Enter the primary IP for the secondary high availability server and press **[Enter]**. The following prompt appears:

```
Tertiary node information:
```

```
Please enter the hostname of the DR server
```

18. Enter the hostname for the disaster recovery server and press **[Enter]**. The following prompt appears:

```
Please enter the DRBD IP for the DR server:
```

19. Enter the IP address for the DR server and press **[Enter]**. The following prompt appears:

```
Virtual IP address information:
```

```
Please enter the Virtual IP Address:
```

```
Please enter the CIDR for the Virtual IP with out the / (example: 24):
```


20. Enter the VIP address and CIDR mask. the following prompt appears:

```
Disaster recovery bandwidth information:
```

```
Please enter the max link speed to the DR system in megabytes/second:
```

21. Enter the speed of the link in megabytes per second.

22. The following confirmation appears:

```
You have selected the following settings, please confirm they are correct:
```

```
Architecture: HA+DR
```

```
This Node: PRIMARY (Node 1)
```

```
Node 1 Hostname: db1
```

```
Node 1 DRBD/Heartbeat IP: 10.255.255.1
```

```
Node 1 Primary IP: 10.64.166.251
```

```
Node 2 Hostname: db2
```

```
Node 2 DRBD/Heartbeat IP: 10.255.255.2
```

```
Node 2 Primary IP: 10.64.166.252
```

```
Node 3 Hostname: 10.64.166.253
```

```
Node 3 DRBD IP: 10.64.166.253
```

```
Virtual IP: 10.64.167.0/23
```

```
DRBD Disk: /dev/mapper/em7vg-db
```

```
DRBD Proxy: Yes
```

```
Max DR DRBD Sync Speed: 120
```

```
Is this information correct? (y/n)
```

23. Review your information to confirm it is correct.

24. Once confirmed, the system will configure itself and you can expect an output similar to the following:

```
Setting up the environment...

Pausing SL1 services

No adjustment needed to /dev/mapper/em7vg-db for DRBD metadata

Setting up DRBD...

Setting up and starting Corosync...

Setting up and starting Pacemaker...

Waiting on cluster services to come up...

Waiting on cluster services to come up...

Waiting on cluster services to come up...

Waiting on cluster services to come up...

Cluster services detected as active

Configuring silo.conf for clustering

Unpausing SL1

Setup is complete

Current cluster status

Stack: corosync

Current DC: db1 (version 1.1.19.linbit-8+20181129.e17.2-c3c624ea3d) -
partition with quorum

Last updated: <date and time you last updated>

Last change: <date and time you last changed> by hacluster via crmd
on db1

2 nodes configured
```

```
7 resources configured
```

```
Online: [ db1 ]
```

```
OFFLINE: [ db2 ]
```

```
Active resources:
```

```
Resource Group: g_em7
```

```
p_fs_drbd1 (ocf::heartbeat:Filesystem): Started db1
```

```
mysql (ocf::sciencelogic:mysql-systemd): Started db1
```

```
Resource Group: g_vip
```

```
virtual_ip (ocf::heartbeat:IPaddr2): Started db1
```

```
drbdproxy (lsb:drbdproxy): Started db1
```

```
Master/Slave Set: ms_drbd_r0 [p_drbd_r0]
```

```
Masters: [ db1 ]
```

```
Master/Slave set: ms_drbd_r0-L [p_drbd_r0-L]
```

```
Masters: [ db1 ]
```

```
Current DRBD status
```

```
r0-L role:Primary
```

```
disk:UpToDate
```

```
peer connection:Connecting
```

```
silo-cluster-install has exited
```

25. The messages above are examples of a successful setup. The important messages that indicate a successful setup are:

```
Unpausing SL1
```

```
Setup is complete
```

```
Current cluster status
```

This message indicates that the system has been configured and the required clustering services have started successfully. You can review the output of the cluster status to see if all services have started as expected.

26. If your setup did not complete successfully with the cluster services in a starting state, you will need to review the log to determine the cause of the failure. Once you correct the issues found in the log the set up can be run again.

Configuring the Secondary High Availability Appliance

To configure the Secondary High Availability appliance for High Availability, perform the following steps:

1. Log in to the console of the Secondary High Availability appliance as the em7admin user.
2. Run the following command to assume root user privileges:

```
sudo -s
```

3. When prompted, enter the password for the em7admin user.
4. Run the following command:

```
silo-cluster-install
```

The following prompt appears:

```
v3.7
```

```
This wizard will assist you in setting up clustering for a  
ScienceLogic appliance. Please be sure to consult the ScienceLogic  
documentation before running this script to ensure that all  
prerequisites have been met.
```

```
1) HA
```

```
2) DR
```

```
3) HA+DR
```

```
4) Quit
```

```
Please select the architecture you'd like to setup:
```

5. Enter "3". The following prompt appears:

```
Will this system be the primary/active node in a new HA+DR cluster?  
(y/n)
```

6. Enter **no**. The following prompt appears:

```
Has the primary/active system already been configured and is it  
running? (y/n)
```

7. Enter **yes** if you have already configured the primary node and it is active and running. Otherwise, choose **no** and complete the setup on the primary node. The following prompt appears:

8. After entering **yes** the following prompt appears:

```
Enter an IP for the primary system that is reachable by this node:
```

9. Enter the IP address of the primary system and press **[Enter]**. This node will attempt to connect to the primary system and retrieve the cluster configuration information. If a successful connection can't be established, you will be prompted again to enter the IP address and a username and a password. The username and password should be the same used for connecting to Maria DB.
10. When a successful connection is made you will be prompted to confirm the architecture to ensure that it is the same for the primary server.

```
Review the following architecture selections to ensure they are  
correct
```

```
Architecture: HA
```

```
Current node hostname: db2
```

```
Crossover Cable: No
```

```
Dedicated link speed: 10Gbps
```

```
Quorum Node IP: 10.64.166.250
```

```
Virtual IP (VIP) used: Yes
```

```
DRBD Proxy: Yes
```

```
Is this architecture correct? (y/n)
```

11. Review that the information is correct and enter **yes**. Enter **no** if it is incorrect and regenerate the configuration on the primary/active system before proceeding.

12. After you enter **yes**, the architecture will be validated for any errors and you will be prompted to confirm the configuration. Review that the information is correct. Your output should be similar to the following:

```
You have selected the following settings, please confirm if they
are correct:
```

```
Architecture: HA+DR
```

```
This Node: SECONDARY (Node 2)
```

```
Node 1 Hostname: db1
```

```
Node 1 DRBD/Heartbeat IP: 10.255.255.1
```

```
Node 1 Primary IP: 10.64.166.251
```

```
Node 2 Hostname: db2
```

```
Node 2 DRBD/Heartbeat IP: 10.255.255.2
```

```
Node 2 Primary IP: 10.64.166.252
```

```
Node 3 Hostname: db3
```

```
Node 3 DRBD IP: 10.64.166.253
```

```
Virtual IP: 10.64.167.0/23
```

```
DRBD Disk: /dev/mapper/em7vg-db
```

```
DRBD Proxy: Yes
```

```
Max DR DRBD Sync Speed: 120
```

```
Is this information correct? (y/n)
```

13. Review that the information is correct and enter **yes**. Enter **no** if it is incorrect and regenerate the configuration on the primary/active system before proceeding.
14. After you enter **yes**, the following message indicates the system is being configured:

```
Setting up the environment...
```

Maria DB running, sending shutdown

Pausing SL1 services

Adjusting /dev/mapper/em7vg-db by extending to 39464 extents to accommodate filesystems and metadata

-Updating firewalld configuration, please be patient...

Setting up DRBD...

Setting up and starting Corosync...

Setting up starting Pacemaker...

Waiting on cluster services to come up...

Waiting on cluster services to come up...

Cluster services detected as active

Configuring silo.conf for clustering

Unpausing SL1

Setup is complete. Please monitor the DRBD synchronization status (drbdadm status)

Failover cannot occur until DRBD is fully synced

Current cluster status

Stack: corosync

Current DC: dbl(version 1.1.19.linbit-8+20181129.e17.2-c3c624ea3d)
- partition with quorum

Last updated:<date and time you last updated>

Last change:<date and time you last changed> by hacluster via crmd
on dbl

2 nodes configured

7 resources configured

Online: [db1 db2]

Active resources:

Resource Group: g_em7

p_fs_drdb1 (ocf::heartbeat:Filesystem): Started db1

mysql (ocf::sciencelogic:mysql-systemd): Started db1

Resource Group: g_vip

virtual_ip (ocf::heartbeat:IPaddr2): Started db1

drbdproxy (lsb:drbdproxy): Started db1

Master/Slave Set: ms_drbd_r0 [p_drbd_r0]

Masters: [db1]

Master/Slave Set: ms_drbd_r0-L [p_drbd_r0-L]

Masters: [db1]

Slaves: [db2]

Current DRBD status

r0-L role:Secondary

disk:Inconsistent blocked:lower

peer role:Primary

replication:SyncTarget peer-disk:UpToDate done:9.89


```
silo-cluster-install has exited
```

Configuring the Disaster Recovery Appliance

To configure the Disaster Recovery appliance, perform the following steps:

1. Log in to the console of the Disaster Recovery appliance as the em7admin user.
2. Run the following command:

```
sudo -s
```

3. When prompted, enter the password for the em7admin user.
4. Run the following command:

```
silo-cluster-install
```

The following prompt appears:

```
v3.7
```

```
This wizard will assist you in setting up clustering for  
ScienceLogic appliance. Please be sure to consult the ScienceLogic  
documentation before running this script to ensure that all  
prerequisites have been met.
```

```
1) HA
```

```
2) DR
```

```
3) HA+DR
```

```
4) Quit
```

```
Please select the architecture you'd like to set up:
```

5. Enter "3". The following prompt appears:

```
Will this system be the primary/active node in a new HA+DR cluster?  
(y/n)
```

6. Enter **no**. The following prompt appears:

```
Has the primary/active system already been configured and is it  
running? (y/n)
```

7. Enter **yes** if you have already configured the primary node and it is active and running. Otherwise choose **no** and complete the setup on the primary node.

8. After entering **yes**, the following prompt appears:

```
Enter an IP for the primary system that is reachable by this node:
```

9. Enter the IP address of the primary system and press **[Enter]**. This node will attempt to connect to the primary system and retrieve the cluster configuration information. If a successful connection can't be established, you will be prompted again to enter the IP address and a username and a password. The username and password should be the same used for connecting to Maria DB.
10. When a successful connection is made you will be prompted to confirm the architecture to ensure that it is the same for the primary server.

```
Preview the following architecture selections to ensure they are correct
```

```
Architecture: HA+DR
```

```
Current node hostname: db3
```

```
Crossover Cable: No
```

```
Dedicated link speed: 10Gbps
```

```
Quorum Node IP: 10.64.166.250
```

```
Virtual IP (VIP) used: Yes
```

```
DRBD Proxy: Yes
```

```
Is this architecture correct? (y/n)
```

11. Review that the information is correct and enter **yes**. Enter **no** if it is incorrect and regenerate the configuration on the primary/active system before proceeding.
12. After you enter **yes**, the architecture will be validated for any errors and you will be prompted to confirm the configuration. Review that the information is correct. Your output should be similar to the following:

```
You have selected the following settings, please confirm if they are correct:
```

```
Architecture: HA+DR
```

```
Node 1 Hostname: db1
```

```
Node 1 DRBD/Heartbeat IP: 10.255.255.1
```

```
Node 1 Primary IP: 10.64.166.251
```

```
Node 2 Hostname: db2
```

```
Node 2 DRBD/Heartbeat IP: 10.255.255.2
```

```
Node 2 Primary IP: 10.64.166.252
```

```
Node 3 Hostname: db3
```

```
Node 3 DRBD IP: 10.64.166.253
```

```
Virtual IP: 10.64.166.0/23
```

```
DRBD Disk: /dev/mapper/em7vg-db
```

```
DRBD Proxy: Yes
```

```
Max DR DRBD Sync Speed: 120
```

```
Is this architecture correct? (y/n)
```

13. Review that the information is correct and enter **yes**. Enter **no** if it is incorrect and regenerate the configuration on the primary/active system before proceeding.
14. After you enter **yes**, the following message indicates the system is being configured:

```
Setting up the environment...
```

```
MariaDB running, sending shutdown
```

```
Pausing SL1 services
```

```
Adjusting /dev/mapper/em7vg-db by extending to 39464 extents to  
accommodate filesystem and metadata
```

```
Updating firewalld configuration, please be patient...
```

```
Setting up DRBD...
```

```
Setting up and starting Corosync...
```

```
Setting up and starting Pacemaker...
```

```
Waiting on cluster services to come up...
```

```
Waiting on cluster services to come up...
```

```
Waiting on cluster services to come up...
```

```
Cluster services detected as active
```

```
Configuring silo.conf for clustering
```

```
Unpausing SL1
```

```
Setup is complete. Please monitor the DRBD synchronization status  
(drbdadm status)
```

```
Failover cannot occur until DRBD is fully synced
```

```
Current cluster status
```

```
Stack: corosync
```

```
Current DC: db3 (version 1.1.19.linbit-8+20181129.el7.2-c3c624ea3d) -  
partition with quorum
```

```
Last updated: <date and time you last updated>
```

```
Last change:<date and time you last changed> by hacluster via crmd on  
db3
```

```
1 node configured
```

```
4 resources configured (2 DISABLED)
```


```
Online: [ db3 ]
```

```
Active resources:
```


```
Master/Slave Set: ms_drbd_r0 [p_drbd_r0]
```

```
Slaves (target-role): [ db3]
```

```
Current DRBD status
```

- On the **Licensing Step 2** page, click the **[Upload]** button to upload the license file. After navigating to and selecting the license file, click the **[Submit]** button to finalize the license. The **Success** message appears:



NOTE: Upon login, SL1 will display a warning message if your license is 30 days or less from expiration, or if it has already expired. If you see this message, take action to update your license immediately.

- Repeat steps 1 - 8 for the Disaster Recovery appliance.

Configuring Data Collection Servers and Message Collection Servers

If you are using a distributed system, you must configure the Data Collectors and Message Collectors to use the new multi-Database Server configuration.

To configure a Data Collector or Message Collector to use the new configuration:

- You can log in to the Web Configuration Utility using any web browser supported by SL1. The address of the Web Configuration Utility is in the following format:

```
https://<ip-address-of-appliance>:7700
```

Enter the address of the Web Configuration Utility in the address bar of your browser, replacing "*ip-address-of-appliance*" with the IP address of the Data Collector or Message Collector.

2. You will be prompted to enter your user name and password. Log in as the em7admin user with the password you configured using the Setup Wizard.
3. On the **Configuration Utilities** page, click the **[Device Settings]** button. The **Settings** page appears:

ScienceLogic™ Web Configuration Utility

Home Licensing Interfaces Device Settings PhoneHome Logout

Settings

Configure your appliance.

Appliance Type
Collection Unit

Database IP Address
10.0.9.213

Save

4. On the **Settings** page, enter the following:
 - **Database IP Address.** Enter the IP addresses of all the Database Servers, separated by commas.
5. Click the **[Save]** button. You may now log out of the Web Configuration Utility for that collector.
6. Perform steps 1-5 for each Data Collector and Message Collector in your system.

NOTE: Starting with SL1 version 11.3.0, SL1 will raise a "Collector Outage" event if a Data Collector or Message Collector could not be reached from either a primary or secondary Database Server.

Failover

Failover is the process by which database services are transferred from the active database to the passive database. You can manually perform failover for testing purposes.

If the active database server in the High Availability cluster stops responding to the secondary database server in the High Availability cluster over both network paths, SL1 will automatically perform failover. After failover completes successfully, the previously active database is now passive, and the previously passive database is now active. There is no automatic failback process; the newly active database will remain active until a failure occurs, or failover is performed manually.

If both appliances in the High Availability cluster fail, you can manually failover to the Disaster Recovery appliance. When the High Availability cluster is restored, you can manually failback from the Disaster Recovery appliance to the High Availability cluster.

Manual Failover Between the Appliances in the High Availability Cluster

To manually failover a High Availability cluster, perform the following steps:

1. Log in to the console of the Primary appliance as the em7admin user.

NOTE: Upon login, SL1 will display a warning message if your license is 30 days or less from expiration, or if it has already expired. If you see this message, take action to update your license immediately.

2. First, you should check the status of the appliances. To do this, enter the following at the shell prompt:

```
cat /proc/drbd
```

Your output will look like this:

```
1: cs:Connected ro:Primary/Secondary ds:UpToDate/UpToDate C r----  
ns:17567744 al:0 bm:1072 lo:0 pe:0 ua:0 ap:0 ep:1 wo:b oos:12521012
```

To failover safely, the output should include "ro:Primary/Secondary ds:UpToDate/UpToDate".

NOTE: If your appliances cannot communicate, your output will include "ro:Primary/Unknown ds:UpToDate/DUnknown". Before proceeding with failback, troubleshoot and resolve the communication problem.

NOTE: If your output includes "ro:Primary/Secondary", but does not include "UpToDate/UpToDate", data is being synchronized between the two appliances. You must wait until data synchronization has finished before performing failback.

3. Run the following command, assuming root user privileges:

```
sudo systemctl stop pacemaker
```

4. When prompted, enter the password for the em7admin user.
5. The Primary database must be shutdown before promoting the Secondary database. To verify MariaDB is shutdown, check the MariaDB log file: `/var/log/mysql/mysqld.log`. The log file should contain the line: "Shutdown complete".
6. To verify all services have started on the newly promoted primary database, run the following command:

```
crm_mon
```

If all services have started, each one should be marked: "started".

7. After all services have started on the newly promoted primary database, run the following command to complete the failover process: `sudo systemctl start pacemaker`
8. To verify that there are two nodes in the cluster, run the following command:

```
crm_mon
```


NOTE: For more information about manual failover for High Availability clusters, see the chapter on [Troubleshooting High Availability and Disaster Recovery Configuration](#).

Manual Failover Between the High Availability Cluster and the Disaster Recovery Appliance

To perform failover between the High Availability cluster and the Disaster Recovery appliance when the High Availability cluster is available, such as to test the failover process:

1. Log in to the console of the primary appliance in the High Availability cluster as the em7admin user.
2. Assume root privileges:

```
sudo -s
```

3. When prompted, enter the password for the em7admin user.
4. Run the following command:

```
coro_config
```

The following prompt appears:

```
1) Enable Maintenance
2) Option Disabled
3) Demote Cluster
4) Stop Pacemaker
5) Resource Status
6) Quit
```

5. Enter "3". The following prompt appears:

```
Cluster currently Primary, would you like to make it Secondary? (y/n)
y
```

6. Enter "y". The following output appears:

```
Issuing command: crm_resource --resource ms_drbd_r0 --set-parameter
target-role --meta --parameter-value Slave
```

7. The primary database **must** be shutdown before promoting the secondary database. To verify MariaDB is shutdown, check the MariaDB log file: `/var/log/mysql/mysqld.log`. The log file should contain the line: "Shutdown complete".

8. Run the following command and review the output:

```
sudo crm status
```

After demotion **crm status** shows the VIP has its own resource group and never demotes after issuing **coro_config** option 3 to demote as mentioned in step 5.

```
2 nodes configured

7 resources configured (2 DISABLED)

Online: [ s11support-ha1 s11support-ha2 ]

Full list of resources:

Resource Group: g_em7

p_fs_drbd1 (ocf::heartbeat:Filesystem): Stopped

mysql (ocf::sciencelogic:mysql-systemd): Stopped

Resource Group: g_vip

virtual_ip (ocf::heartbeat:IPaddr2): Started s11support-ha1

drbdproxy (lsb:drbdproxy): Started s11support-ha1

Master/Slave Set: ms_drbd_r0 [p_drbd_r0]

Slaves (target-role): [ s11support-ha1 ]

Master/Slave Set: ms_drbd_r0-L [p_drbd_r0-L]

Masters: [ s11support-ha1 ]

Slaves: [ s11support-ha2 ]
```

9. Log in to the console of the Disaster Recovery appliance as the em7admin user.
10. Run the following command:

```
sudo -s
```

11. When prompted, enter the password for the em7admin user.
12. Run the following command:

```
coro_config
```

The following prompt appears:

```
1) Enable Maintenance

2) Option Disabled
```

```
3) Promote DRBD
```

```
4) Stop Pacemaker
```

```
5) Resource Status
```

```
6) Quit
```

```
Please enter the number of your choice:
```

16. Enter "3". The following prompt appears:

```
Node currently Secondary, would you like to make it Primary? (y/n)
```

17. Enter "y". The following output appears:

```
Issuing command: crm_resource --resource ms_drbd_r0 --set-parameter  
target-role --meta --parameter-value Master
```

18. To verify that an appliance is active after failover, ScienceLogic recommends checking the status of MariaDB, which is one of the primary processes on Database Servers. To verify the status of MariaDB, execute the following command on the newly promoted Database Server:

```
silo_mysql -e "select 1"
```

If MariaDB is running normally, you will see a '1' in the console output.

TIP: Because larger systems can take more time to start the database, verify that MariaDB has started successfully before running the above command. To verify MariaDB has started successfully, check the MariaDB log file: `/var/log/mysql/mysqld.log`. The log file should contain the line: `"/usr/sbin/mysqld: ready for connections."`

19. If you are using a distributed SL1 system, you must reconfigure all Administration Portals in your system to use the new Database Server. To do this, follow the steps listed in the [Reconfiguring Administration Portals](#) section.

Failover when the High Availability Cluster is Inaccessible

To perform failover when the both appliances in the High Availability cluster are inaccessible:

1. **Make sure to power-down the inaccessible Database Servers.** This step is required to avoid a split-brain configuration (two primary appliances). A split-brain configuration will cause your data to become corrupted.
2. The primary database must be shutdown before promoting the secondary database. To verify MariaDB is shutdown, check the MariaDB log file: `/var/log/mysql/mysql.log`. The log file should contain the line: "Shutdown complete".
3. Log in to the console of the Disaster Recovery appliance as the em7admin user.

4. Run the following command:

```
sudo -s
```

5. When prompted, enter the password for the em7admin user.
6. Run the following command:

```
coro_config
```

The following prompt appears:

```
1) Enable Maintenance
```

```
2) Option Disabled
```

```
3) Promote DRBD
```

```
4) Stop Pacemaker
```

```
5) Resource Status
```

```
6) Quit
```

```
Please enter the number of your choice:
```

7. Enter "3". The following prompt appears:

```
Node currently Secondary, would you like to make it Primary? (y/n)
```

8. Enter "y". The following output appears:

```
Issuing command: crm_resource --resource ms_drbd_r0 --set-parameter target-role --meta --parameter-value Master
```

9. To verify that an appliance is active after failover, ScienceLogic recommends checking the status of MariaDB, which is one of the primary processes on Database Servers. To verify the status of MariaDB, execute the following command on the newly promoted Database Server:

```
silo_mysql -e "select 1"
```

If MariaDB is running normally, you will see a '1' in the console output.

TIP: Because larger systems can take more time to start the database, verify that MariaDB has started successfully before running the above command. To verify MariaDB has started successfully, check the MariaDB log file: `/var/log/mysql/mysqld.log`. The log file should contain the line: `"/usr/sbin/mysqld: ready for connections."`

10. If you are using a distributed SL1 system, you must reconfigure all Administration Portals in your system to use the new Database Server. To do this, follow the steps listed in the [Reconfiguring Administration Portals](#) section.

NOTE: For more information about manual failover on a High Availability and Disaster Recovery Stack with Three Appliances, see the chapter on [Troubleshooting High Availability and Disaster Recovery Configuration](#).

Manual Failback Between the Disaster Recovery Appliance and the High Availability Cluster

To perform failback between the Disaster Recovery appliance and the High Availability cluster, perform the following steps:

1. Log in to the console of the Disaster Recovery appliance as the em7admin user.
2. First, you should check the status of the appliances. To do this, enter the following at the shell prompt:

```
drbdadm status
```

Your output will look like this:

```
r0 role:Primary
```

```
disk:UpToDate
```

```
peer role:Secondary
```

```
replication:Established peerdisk:UpToDate
```

To failback safely, the output should include:

```
peer role:Secondary replication: Established peer-disk:UpToDate
```

NOTE: If your two appliances cannot communicate, your output will include "ro:Primary/Unknown ds:UpToDate/DUnknown". Before proceeding with failback, troubleshoot and resolve the communication problem.

NOTE: If your output includes "ro:Primary/Secondary", but does not include "UpToDate/UpToDate", data is being synchronized between the two appliances. You must wait until data synchronization has finished before performing failback.

3. Assume root privileges:

```
sudo -s
```

4. When prompted, enter the password for the em7admin user.
5. Run the following command:

```
coro_config
```

The following prompt appears:

```
1) Enable Maintenance
2) Option Disabled
3) Demote DRBB
4) Stop Pacemaker
5) Resource Status
6) Quit
```

Please enter the number of your choice:

6. Enter "3". The following prompt appears:

```
Node currently Primary, would you like to make it Secondary? (y/n) y
```

7. Enter "y". The following output appears:

```
Issuing command: crm_resource --resource ms_drbd_r0 --set-parameter
target-role --meta --parameter-value Slave
```

NOTE: The primary database **must** be shutdown before promoting the secondary database. To verify MariaDB is shutdown, check the MariaDB log file: `/var/log/mysql/mysql.log`. The log file should contain the line: "Shutdown complete".

CAUTION: Before promoting the High Availability cluster to primary, ensure that the Disaster Recovery node is completely demoted. DRDB must be connected between the appliance that was formerly the HA primary and the DR node.

8. Log in to the console of the High Availability cluster that you want to promote to Primary as the em7admin user. This node should be the current High Availability primary node. Please check the output of `crm status` and verify the node you are on is currently running the `g_vip` resources.
9. Run the following code to verify the disk replication is connected and up to date between the High Availability:

```
drbdadm status
```

If you are using Oracle Linux 7, the output should read:

```
r0-L role:Primary
```

```
disk:UpToDate
```

```
peer role:Secondary
```

```
replication:Established peer-disk:UpToDate
```

NOTE: On Oracle Linux 7, the resources will be `r0-L` and the `role:` should be Primary with the `peer role:` Secondary with replication Established and UpToDate.

On the Oracle Linux 7 you must verify that the stacked resource is in the correct state by running the following command:

```
drbdadm status --stacked
```

Check that the output is similar to:

```
r0 role:Secondary
```

```
disk:UpToDate
```

```
peer role:Secondary
```

```
replication:Established peer-
```

```
disk:UpToDate
```

NOTE: On a stacked resource on Oracle Linux 7, the resource will be `r0` and both roles should be Secondary with replication Established and peer-disk UpToDate.

If you are using Oracle Linux 8, the output should be similar to:

```
r0 role: Secondary
```

```
disk:UpToDate
```

```
db2 role:Secondary
```

```
peer-disk:UpToDate
```

```
db3 role:Secondary
```

```
peer-disk:UpToDate
```

NOTE: On Oracle Linux 8, the resources will be `r0` and both the secondary High Availability node and Disaster Recovery node will be listed. All roles should be Secondary and UpToDate.

10. Assume root privileges:

```
sudo -s
```

11. When prompted, enter the password for the em7admin user.

12. Run the following command:

```
coro_config
```

The following prompt appears:

```
1) Enable Maintenance
2) Option Disabled
3) Promote DRBD
4) Stop Pacemaker
5) Resource Status
6) Quit
```

```
Please enter the number of your choice:
```

13. Enter "3". The following prompt appears:

```
Node currently Secondary, would you like to make it Primary? (y/n)
```

15. Enter "y". The following output appears:

```
Issuing command: crm_resource --resource ms_drbd_r0 --set-parameter
target-role --meta --parameter-value Master
```

16. To verify that an appliance is active after failback, ScienceLogic recommends checking the status of MariaDB, which is one of the primary processes on Database Servers. To verify the status of MariaDB, enter the following command on the newly promoted Database Server:

```
silo_mysql -e "select 1"
```

If MariaDB is running normally, you will see a '1' in the console output.

TIP: Because larger systems can take more time to start the database, verify that MariaDB started successfully before running the above command. To verify MariaDB started successfully, check the MariaDB log file: `/var/log/mysql/mysqld.log`. The log file should contain the line: `"/usr/sbin/mysqld: ready for connections."`

17. If you are using a distributed SL1 system, you must reconfigure all Administration Portals in your system to use the new Database Server. To do this, follow the steps listed in the [Reconfiguring Administration Portals](#) section.

Reconfiguring Administration Portals

If you are using a Distributed system and you did not configure a virtual IP address, you must configure all Administration Portals in your system to use the new Primary Database Server after performing failover or failback. To configure an Administration Portal to use the new Database Server:

You must perform the following steps in the Web Configuration Utility to configure an Administration Portal:

1. You can log in to the Web Configuration Utility using any web browser supported by SL1. The address of the Web Configuration Utility is in the following format:

```
https://<ip-address-of-appliance>:7700
```

Enter the address of the Web Configuration Utility in the address bar of your browser, replacing *ip-address-of-appliance* with the IP address of the Secondary appliance.

2. Log in as the "em7admin" user with the password you configured using the Setup Wizard. The **Configuration Utility** page appears.
3. Click the **[Device Settings]** button. The **Settings** page appears:

The screenshot shows the 'Settings' page of the ScienceLogic Web Configuration Utility. The page title is 'Settings' with the subtitle 'Configure your appliance.' The navigation bar includes 'Home', 'Licensing', 'Interfaces', 'Device Settings', 'PhoneHome', and 'Logout'. The main content area contains several form fields:

- Appliance Type:** Administration Portal
- Database IP Address:** 10.0.9.213
- Database Username:** root
- Database Password (change only):** [Empty field]
- Confirm Database Password:** [Empty field]
- GUI Username:** root
- GUI Password (change only):** [Empty field]
- Confirm GUI Password:** [Empty field]

A 'Save' button is located at the bottom center of the form.

4. On the **Settings** page, enter the following:
 - **Database IP Address.** The IP address of the new Primary ScienceLogic Database Server.
5. Click the **[Save]** button. You may now log out of the Web Configuration Utility.
6. Log into the Administration Portal appliance with your SSH credentials and run the following command:

```
sudo systemctl restart nextui;
```

7. Repeat these steps for each Administration Portal in your system.

Verifying that a Database Server is Primary

To verify that your network is configured correctly and will allow the newly active Database Server to operate correctly, check the following system functions:

- If you use Active Directory or LDAP authentication, log in to the user interface using a user account that uses Active Directory or LDAP authentication.
- In the user interface, verify that new data is being collected.
- If your system is configured to send notification emails, confirm that emails are being received as expected. To test outbound email, create or update a ticket and ensure that the ticket watchers receive an email.

NOTE: On the **Behavior Settings** page (System > Settings > Behavior, if the field **Automatic Ticketing Emails** is set to *Disabled*, all assignees and watchers will not receive automatic email notifications about any tickets. By default, the field is set to *Enabled*.

- If your system is configured to receive emails, confirm that emails are being received correctly. To test inbound email, send a test email that will trigger a "tickets form Email" policy or an "events from Email" policy.

Chapter

5



Troubleshooting High Availability and Disaster Recovery Configuration

Overview

This chapter describes how to troubleshoot various scenarios of High Availability or Disaster Recovery configurations.

This chapter assumes that you are comfortable using a UNIX shell session and can use the basic functions within the vi editor.

Use the following menu options to navigate the SL1 user interface:

- To view a pop-out list of menu options, click the menu icon ().
- To view a page containing all of the menu options, click the Advanced menu icon ().

This chapter covers the following topics:

Overview	92
Diverged Data	104
Firewall Issues	105
Ahead-Behind or Disconnect/Reconnect Issues	106
Other Issues	107

Overview

Distributed Replicated Block Device (DRBD) is a distributed replicated storage system for Linux. It is generally used in high availability computer clusters as a kernel driver and to manage applications.

Although DRBD plays a major role in SL1 application clustering, the cluster manager that is responsible for loading, connecting, promoting, and demoting DRBD cannot automatically resolve any issues that might occur at the DRBD application level.

There are few scenarios in which DRBD might fail to connect:

- Split-brain has occurred.
- Unrelated data aborting occurred and the nodes do not recognize each other due to data divergence.
- Firewall issues occurred.
- Proxy service issues occurred (WFReportParams).
- An unstable or bad network was provided.
- Inconsistent kernel packages are created.

What is a Split-brain Configuration?

In SL1, DRBD resides on the Primary node regardless of whether the architecture is HA, DR, or HA and DR. The Primary node is replicated to the secondary node. You might have two or three nodes in the SL1 cluster, but only one can be the primary DRBD at a time. A **split-brain configuration** occurs when two nodes of a cluster both think they are elected to be the primary and start competing for shared resources. When a split-brain configuration occurs, both nodes will disconnect.

Resolving a Split-brain Configuration on a Disaster Recovery Stack with Two Appliances

To resolve a split-brain configuration between two nodes on a Disaster Recovery stack:

1. Check the status of both appliances. To do this:
 - Log in to the console of the current Primary appliance as the em7admin user and enter the following at the shell prompt:

```
drbdadmin status
```

- If your output looks like this, you have a split-brain configuration:

```
r0 role:Primary
```

```
disk:UpToDate
```

NOTE: There is no peer section listed in the status output. This means the node is in Stand Alone mode and not attempting to reconnect.

- On the Disaster Recovery node, enter the following code:

```
drbdadmin status
```

- You could have several outputs returned.
 - If your output looks like this, you have a split-brain configuration that indicates the system is split and believes it is in Primary and Stand Alone mode:

```
r0 role:Primary
```

```
disk:UpToDate
```

- If your output looks like this, you have a split-brain configuration that indicates the system is split and believes it is in Primary and attempting to connect to the peer:

```
r0 role:Primary
```

```
disk:UpToDate
```

```
peer connection:Connecting
```

- If your output looks like this, you have a split-brain configuration that indicates the system is split and believes it is Secondary and is in Stand Alone mode:

```
r0 role:Secondary
```

```
disk:UpToDate
```

- If your output looks like this, you have a split-brain configuration that indicates the system is split and believes it is Secondary and is attempting to connect to the peer:

```
r0 role:Secondary
```

```
disk:UpToDate
```

```
peer connection:Connecting
```

- In a safe failover/failback scenario, your output would include:

- Primary:

```
r0 role:Primary
```

```
disk:UpToDate
```

```
peer role:Secondary
```

```
replication:Established peer-disk:UpToDate
```

- Secondary:

```
r0 role:Secondary
```

```
disk:UpToDate
```

```
peer role:Primary
```

```
replication:Established peer-disk:UpToDate
```

- When using Oracle Linux 8 the output appears:

- Primary:

```
r0 role:Primary
```

```
disk:UpToDate
```

```
db2 role:Secondary
```

```
peer-disk:UpToDate
```

- Secondary:

```
r0 role:Secondary
```

```
disk:UpToDate
```

```
db1 role:Primary
```

```
peer-disk:UpToDate
```

2. Resolve split-brain on the passive node.

- If the passive node is attempting to connect to its peer, enter the following command:

```
sudo drbdadm disconnect r0
```

- If the passive node has marked itself the Primary, enter the following command:

```
sudo crm_mon-1
```

- If you see the line `Resource Group: g_em7`, enter the following command to demote the system:

```
crm_resource --resource ms_drbd_r0--set-parameter target-role --meta -  
-parameter-value Slave
```

NOTE: SL1 recommends periodically checking for the output of `drbdadm status` as the system transitions to secondary status while you wait for your system to demote itself.

- If your system is already Secondary, or has transitioned to Secondary, enter the following command:

```
sudo drbdadm -- --discard-my-data
```

```
connect r0
```

- Next, on the Primary node, enter the following command:

```
sudo drbdadm connect r0
```

3. Validate that your split-brain issue has been resolved and is reconnected.

- To check the status on both the Primary node and Secondary node, enter the following command:

```
drbdadm status
```

NOTE: Depending on your network, it can take several minutes to establish the connection and restart replication.

- If the output of the status returns `connecting`, DRBD is attempting to establish the connection. When you see the `peer` line, connection is established and replication has resumed.
- However, if the `connecting` line disappears and no `peer` line appears, you have failed to resolve the split-brain situation. Please contact ScienceLogic Support for further assistance.

IMPORTANT: You must always perform your resolution on the passive node first.

Resolving a Split-brain Configuration on a High Availability Stack with Two Appliances

To resolve a split-brain configuration between two nodes on a High Availability Primary or Secondary stack:

1. Check the status of both appliances. To do this:

- Log in to the console of the current Primary appliance as the em7admin user and enter the following at the shell prompt:

```
drbdadmin status
```

- If your output looks like this, you have a split-brain configuration:

```
r0 role:Primary
```

```
disk:UpToDate
```

NOTE: There is no peer section listed in the status output. This means the node is in Stand Alone mode and not attempting to reconnect.

- On the Disaster Recovery node, enter the following code:

```
drbdadmin status
```


- You could have several outputs returned.
 - If your output looks like this, you have a split-brain configuration that indicates the system is split and believes it is in Primary and Stand Alone mode:

```
r0 role:Primary
```

```
disk:UpToDate
```

- If your output looks like this, you have a split-brain configuration that indicates the system is split and believes it is in Primary and attempting to connect to the peer:

```
r0 role:Primary
```

```
disk:UpToDate
```

```
peer connection:Connecting
```

- If your output looks like this, you have a split-brain configuration that indicates the system is split and believes it is Secondary and is in Stand Alone mode:

```
r0 role:Secondary
```

```
disk:UpToDate
```

- If your output looks like this, you have a split-brain configuration that indicates the system is split and believes it is Secondary and is attempting to connect to the peer:

```
r0 role:Secondary
```

```
disk:UpToDate
```

```
peer connection:Connecting
```

- In a safe failover/failback scenario, your output would include:

- Primary:

```
r0 role:Primary
```

```
disk:UpToDate
```

```
peer role:Secondary
```

```
replication:Established peer-disk:UpToDate
```

- Secondary:

```
r0 role:Secondary
```

```
disk:UpToDate
```

```
peer role:Primary
```

```
replication:Established peer-disk:UpToDate
```

- When using Oracle Linux 8 the output appears:

- Primary:

```
r0 role:Primary
```

```
disk:UpToDate
```

```
db2 role:Secondary
```

```
peer-disk:UpToDate
```

- Secondary:

```
r0 role:Secondary
```

```
disk:UpToDate
```

```
db1 role:Primary
```

```
peer-disk:UpToDate
```

2. Resolve split-brain on the passive node.

- If the passive node is attempting to connect to its peer, enter the following command:

```
sudo drbdadm disconnect r0
```

- If the passive node has marked itself the Primary (i.e. both nodes are claiming Primary) this indicates an issue with the High Availability cluster not being able to establish cluster connectivity and/or not being able to control resources correctly. Please stop the process and contact ScienceLogic Support for assistance.
- If your system is already Secondary, or has transitioned to Secondary, enter the following command:

```
sudo drbdadm -- --discard-my-data
```

```
connect r0
```

- Next, on the Primary node, enter the following command:

```
sudo drbdadm connect r0
```

3. Validate that your split-brain issue has been resolved and is reconnected.

- To check the status on both the Primary node and Secondary node, enter the following command:

```
drbdadm status
```

NOTE: Depending on your network, it can take several minutes to establish the connection and restart replication.

- If the output of the status returns `connecting`, DRBD is attempting to establish the connection. When you see the `peer` line, connection is established and replication has resumed.
- However, if the `connecting` line disappears and no `peer` line appears, you have failed to resolve the split-brain situation. Please contact ScienceLogic Support for further assistance.

IMPORTANT: You must always perform your resolution on the passive node first.

Resolving a Split-brain Configuration on a High Availability and Disaster Recovery Stack with Three Appliances

There are two scenarios you might encounter when working on a high availability and disaster recovery stack involving three appliances. Refer to the steps below for your specifications.

Resolving a Split-brain Between the Primary High Availability and Secondary High Availability Nodes

To resolve a split-brain configuration between three nodes on a High Availability Primary or Secondary stack:

1. Check the status of all appliances. To do this:

- Log in to the console of the current Primary appliance as the em7admin user and enter the following at the shell prompt:

```
drbdadm status
```

- If your output looks like this, you have a split-brain configuration:

- On an OL7 system, your output will be similar to:

```
r0-L role:Primary
```

```
disk:UpToDate
```

NOTE: The output for OL7 might contain a peer line listing various state information. When the nodes are disconnected, your output will either look similar to the example above or will indicate `peer:Connecting`.

- On an OL8 system, your output will be similar to:

```
r0 role:Primary
```

```
disk:UpToDate
```

```
db2 connection:StandAlone
```

```
db3 connection:Secondary
```

```
peer-disk:UpToDate
```

NOTE: Stacked resources are no longer used in OL8. Instead, all nodes are direct peers and you have to associate the name to the correct nodes role. In the example above, the **db2** is the high availability secondary.

- On the High Availability (secondary) node enter the following code:

```
drbdadm status
```

- If your output looks like this, you have a split-brain configuration:

- On an OL7 system, your output will be similar to:

```
r0-L role:Secondary
```

```
disk:UpToDate
```

NOTE: The output for OL7 might contain a peer line listing various state information. When the nodes are disconnected, your output will either look similar to the example above or will indicate `peer:Connecting`.

- On an OL8 system, your output will be similar to:

```
r0 role:Secondary
```

```
disk:Outdated
```

```
db1 connection:Connecting
```

```
db3 connection:StandAlone
```

NOTE: Stacked resources are no longer used in OL8. Instead, all nodes are direct peers and you have to associate the name to the correct nodes role. In the example above, the **db1** is the high availability primary and is outdated compared to the other nodes.

2. Resolve split-brain on the passive node on OL7 systems.

- On the High Availability (secondary) node, enter the following command:

```
drbdadm -- --discard-my-data connect r0-L
```

- Next, on the High Availability Primary node, enter the following command:

```
drbdadm connect r0-L
```

IMPORTANT: You must always perform your resolution on the passive node first.

3. Resolve split-brain on the passive node on OL8 systems.

- On the High Availability Secondary node enter the following command:

```
drbdadm connect r0 --discard-my-data
```

- On the High Availability Primary node enter the following command:

```
drbdadmin connect r0
```

NOTE: Ensure there are no existing connections to the Disaster Recovery node as errors might appear in the output because the connections have already been made. Check the status with `drbdadm status` to make sure the nodes connect.

Resolving a Split-brain Between the Primary High Availability and the Disaster Recovery Nodes

To resolve a split-brain configuration between three nodes on a High Availability (primary) and Disaster Recovery stack:

1. Check the status of all appliances. To do this:
 - Log in to the console of the current Primary appliance as the `em7admin` user and enter the following at the shell prompt:
 - On OL7

```
drbdadm --stacked
```
 - On OL8

```
drbdadm status
```

- If your output looks like this, you have a split-brain configuration:

- On an OL7 system, your output will be similar to:

```
r0-L role:Primary
```

```
disk:UpToDate
```

```
peer connection:Connecting
```

NOTE: The output for OL7 might contain a peer line listing various state information. When the nodes are disconnected, your output will look similar to the example above indicating replication is not established.

- On an OL8 system, your output will be similar to:

```
r0 role:Primary
```

```
disk:UpToDate
```

```
db2 role:SEcondary
```

```
peer-disk:UpToDate
```

```
db3 connection:StandAlone
```

NOTE: In the example above, the **db3** system is the Disaster Recovery node and the connection is *StandAlone*. Regardless of the **db3** connection state it is important to note that it is not connected and will not list a *peer-disk* status.

- On the Disaster Recovery node enter the following code:

```
drbdadm status
```

- If your output looks like this, you have a split-brain configuration:

```
1: cs:StandAlone ro:Secondary/Unknown ds:UpToDate/DUnknown r----
```

```
-
```

```
ns:0 nr:0 dw:0 dr:0 al:0 bm:0 lo:0 pe:0 ua:0 ap:0 ep:1 wo:d
```

```
oos:0
```

NOTE: In a safe failback scenario, your output would include *ro:Primary/Secondary* *ds:UpToDate/UpToDate*. Also, the kernel logs will include a message indicating that split-brain is detected but unresolved.

2. Resolve split-brain on the passive node on OL7 systems.

- On the Disaster Recovery node, enter the following command:

```
drbdadm -- --discard-my-data connect r0
```
- Next, on the High Availability Primary node, enter the following command:

```
drbdadm connect --stacked r0
```

IMPORTANT: You must always perform your resolution on the passive node first.

3. Resolve split-brain on the passive node on OL8 systems.

- On the High Availability Secondary node enter the following command:

```
drbdadm connect r0 --discard-my-data
```
- On the High Availability Primary node enter the following command:

```
drbdsetup connect r0 2
```

NOTE: Ensure there are no existing connections to the Disaster Recovery node as errors might appear in the output because the connections have already been made. Check the status with `drbdadm status` to make sure the nodes connect.

Diverged Data

Overview

In some situations, DRBD nodes will not connect. When you inspect your logs, you might find the following error message:

```
2020-07-14T00:01:54.831938+00:00 s11db02 kernel: [15919.112878] block
drbd1: Unrelated data, aborting!
```

Resolving a Diverged Data Issue

DRBD will not connect if data has diverged beyond the point that the peers no longer recognize each other's generation identifiers. To resolve a diverged data issue:

1. Stop the DRBD completely on the secondary (passive) node and delete the metadata:

```
drbdadm down r0
```

```
drbdadm wipe-md r0
```


2. Recreate the DRBD metadata and restart the DRBD:

```
drbdadm create-md r0
```

```
drbdadm up r0
```

3. Check to you see if the nodes are connected and synchronizing by running the following command:

```
drbdadm status
```

Firewall Issues

Overview

DRBD uses TCP port 7788 for communication between Database Servers. The port must be open in both directions to allow DRBD connection and replication.

Resolving Firewall Issues

To verify there are no port issues and that data can flow bidirectionally:

- Log in to the console of the current Primary appliance as the em7admin user and enter the following at the shell prompt:

```
firewall-cmd --list-all | grep 7788
```

- If your output looks like this, you have established connections:

```
tcp 0 0 10.64.70.42:54787 10.64.70.24:7788 ESTABLISHED 7181/drbd-proxy
```

```
tcp 0 0 169.254.1.6:56207 169.254.1.5:7788 ESTABLISHED -
```

```
tcp 0 0 10.64.70.42:7788 10.64.70.24:54728 ESTABLISHED 7181/drbd-proxy
```

```
tcp 0 0 169.254.1.6:7788 169.254.1.5:53315 ESTABLISHED
```

- On the Disaster Recovery node, enter the following code:

```
firewall-cmd --list-all | grep 7788
```

- If your output looks like this, you have established connections:

```
tcp 0 0 10.64.70.24:7788 10.64.70.42:54787 ESTABLISHED 1397/drbd-proxy
```

```
tcp 0 0 10.64.70.24:54728 10.64.70.42:7788 ESTABLISHED 1397/drbd-proxy
```

Ahead-Behind or Disconnect/Reconnect Issues

Overview

These conditions are always caused by a network issue between the Primary node and the Disaster Recovery node. There are several reasons why you might experience network issues, ranging from very simple to very complex, but the root cause is always an insufficiency in data transit. The most common reasons are:

- insufficient bandwidth for the data
- latency plus insufficient bandwidth
- network congestion and queue drops
- packet loss

Resolving Self-Monitoring Network Issues

ScienceLogic has many tools that illustrate some of the insufficient network issues you might face. Viewing the self-monitoring graphs can help you narrow the reason for your network issues.

- **DRBD Status Performance | Disk Write.** This graph shows the amount of data DRBD is writing to the disk.
- **DRBD Status Performance | Network Send.** This graph shows the amount of data DRBD has sent across the network from DRBD to the "next hop." This graph should be identical to the Disk Write graph. If not, it means that DRBD cannot send data as fast as it can write it.
- **DRBD Status Performance | Disk Write vs. Network Send.** This graph is a presentation of the *Disk Write* and *Network Send* collectors in comparison to each other. This graph should be relatively flat if everything is operating normally on the network.

Another solution to these network issues is to enable compression on the DRBD proxy. For more information about enabling compression on the DRBD proxy, see the following Knowledge Base article: [Using Compression with DRBD](#).

Other causes of network issues might require more in-depth intervention due to limitations beyond your control. The following suggestions are actions you can implement to help your situation:

- If the bandwidth to the Disaster Recovery node is shared, implement quality of service (QoS) on the network and ensure the DRBD traffic is placed in a low latency, no-drop queue.
- Monitor all network interfaces for all links between SL1 nodes and then enable packet, error, and discard collections on those interfaces. You will be able to see if there are drops or errors, or if it is overrunning the packet switching ability of a switch or router.
- Check for retransmissions on the SL1 system with `netstat -s`. If you see a high number of retransmits compared to the total number of sent packets, your network is having a problem. In this case, "high" is 0.2% or greater.

In some cases, when the steps above do not resolve your network issue, you might need to escalate the issue to ScienceLogic Support. Before contacting Support, you should collect the following data:

- Screenshots of all three of the DRBD Status Performance graphs referenced above. You should convert the time window when issues are found.
- The output of `netstat -s` and `ifconfig` on both the Primary and the Disaster Recovery nodes.
- Capture the output of each of these commands after running on the Database Server:
 - Basic ping: `ping -w 60 <Dr IP>`
 - Large packet ping: `ping -w 60 -M do -s 1472 <dr IP>`
 - Rapid ping: `ping -w 60 -f <dr IP>`
 - Large packet rapid ping: `ping -w 60 -s 1472 -f <dr IP>`

NOTE: Do not `control -C` these commands, because they will expire after 60 seconds.

Other Issues

WFReportParams Loops

In the event the DRBD fails to reconnect during the process of a failover to Disaster Recovery, the DRBD will stall and you will see the symptoms of a "WFReportParams" loop error message.

To reconnect the DRBD on the High Availability node, enter the following command:

```
drbdadm proxy-down r0 && drbdadm proxy-up r0
```

To reconnect the DRBD on a stacked High Availability node, enter the following command:

```
drbdadm proxy-down r0 -S && drbdadm proxy-up r0 -S
```

To reconnect the DRBD on the Disaster Recovery node, enter the following command:

```
drbdadm proxy-down r0 && drbdadm proxy-up r0
```

NOTE: If you need assistance, contact ScienceLogic Support.

Network Issues

As with any network-dependent application, the DRBD proxy requires a stable network. If the network has issues such as flapping link or fluctuating bandwidth capacity, it affects DRBD data replication. Make sure network issues are properly diagnosed and addressed.

Kernel Package Issues

Although rare, two DRBD nodes might fail to communicate due to incompatible DRBD modules. This could sometimes result from one node that was not rebooted after a kernel package update, which will only take effect after a reboot.

© 2003 - 2023, ScienceLogic, Inc.

All rights reserved.

LIMITATION OF LIABILITY AND GENERAL DISCLAIMER

ALL INFORMATION AVAILABLE IN THIS GUIDE IS PROVIDED "AS IS," WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED. SCIENCELOGIC™ AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT.

Although ScienceLogic™ has attempted to provide accurate information on this Site, information on this Site may contain inadvertent technical inaccuracies or typographical errors, and ScienceLogic™ assumes no responsibility for the accuracy of the information. Information may be changed or updated without notice. ScienceLogic™ may also make improvements and / or changes in the products or services described in this Site at any time without notice.

Copyrights and Trademarks

ScienceLogic, the ScienceLogic logo, and EM7 are trademarks of ScienceLogic, Inc. in the United States, other countries, or both.

Below is a list of trademarks and service marks that should be credited to ScienceLogic, Inc. The ® and ™ symbols reflect the trademark registration status in the U.S. Patent and Trademark Office and may not be appropriate for materials to be distributed outside the United States.

- ScienceLogic™
- EM7™ and em7™
- Simplify IT™
- Dynamic Application™
- Relational Infrastructure Management™

The absence of a product or service name, slogan or logo from this list does not constitute a waiver of ScienceLogic's trademark or other intellectual property rights concerning that name, slogan, or logo.

Please note that laws concerning use of trademarks or product names vary by country. Always consult a local attorney for additional guidance.

Other

If any provision of this agreement shall be unlawful, void, or for any reason unenforceable, then that provision shall be deemed severable from this agreement and shall not affect the validity and enforceability of any remaining provisions. This is the entire agreement between the parties relating to the matters contained herein.

In the U.S. and other jurisdictions, trademark owners have a duty to police the use of their marks. Therefore, if you become aware of any improper use of ScienceLogic Trademarks, including infringement or counterfeiting by third parties, report them to Science Logic's legal department immediately. Report as much detail as possible about the misuse, including the name of the party, contact information, and copies or photographs of the potential misuse to: legal@sciencelogic.com. For more information, see <https://sciencelogic.com/company/legal>.

ScienceLogic

800-SCI-LOGIC (1-800-724-5644)

International: +1-703-354-1010