



High Availability and Disaster Recovery

SL1 version 8.14.0

Table of Contents

Introduction	4
Disaster Recovery	5
High Availability	5
Differences Between Disaster Recovery and High Availability for Database Servers	5
Configuring MX Records	6
Disaster Recovery with Two Appliances	7
Prerequisites	8
Unique Host Names	8
Licensing DRBD Proxy	8
Using a Virtual IP Address	9
Configuring Disaster Recovery	9
Configuring the Primary Appliance	9
Configuring the Secondary Appliance	11
Licensing the Secondary Appliance	13
Configuring Data Collection Servers and Message Collection Servers	15
Failover	15
Failover When Both Database Appliances are Accessible	16
Failover When the Primary Database Appliance is Inaccessible	17
Reconfiguring Administration Portals	18
Verifying that a Database Server is Primary	19
Failback	20
High Availability with Two Appliances	21
Prerequisites	22
Unique Host Names	22
Addressing the Cluster	22
Configuring Heartbeat IP Addresses	23
Testing the Heartbeat Network	23
Configuring High Availability	24
Configuring the Primary Appliance	24
Configuring the Secondary Appliance	26
Licensing the Secondary Appliance	28
Configuring Data Collection Servers and Message Collection Servers	29
Failover	30
Manual Failover for High Availability Clusters	30
Verifying that a Database Server is Active	30
High Availability and Disaster Recovery with Three Appliances	32
Prerequisites	33
Unique Host Names	33
Licensing DRBD Proxy	33
Addressing the Cluster	34
Reconfiguring an Existing High Availability System	34
Configuring Heartbeat IP Addresses	35
Testing the Heartbeat Network	36
Configuring Three Appliances for High Availability and Disaster Recovery	36
Configuring the Primary High Availability Appliance	36
Configuring the Secondary High Availability Appliance	38
Configuring the Disaster Recovery Appliance	40
Licensing the Secondary High Availability and Disaster Appliances	42
Configuring Data Collection Servers and Message Collection Servers	43
Failover	44

Manual Failover Between the Appliances in the High Availability Cluster44
Manual Failover Between the High Availability Cluster and the Disaster Recovery Appliance 45
Failover when the High Availability Cluster is Inaccessible 46
Manual Failback Between the Disaster Recovery Appliance and the High Availability Cluster 47
Reconfiguring Administration Portals 49
Verifying that a Database Server is Primary50

Chapter


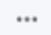
1

Introduction

Overview

This manual is intended for system administrators responsible for setting up Database Servers for Disaster Recovery or High Availability, or both configurations.

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 includes the following topics:

<i>Disaster Recovery</i>	5
<i>High Availability</i>	5
<i>Differences Between Disaster Recovery and High Availability for Database Servers</i>	5
<i>Configuring MX Records</i>	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.

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:

- 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.
- 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.
- A High Availability configuration is not supported for All-In-One Appliances. A Disaster Recovery configuration is not supported for All-In-One Appliances.
- 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.

<p>NOTE: A Distributed SL1 system can be configured for both High Availability and Disaster Recovery.</p>
--

Configuring MX Records

In all configurations, the primary Database Server or All-In-One Appliance is responsible for processing all inbound email. To prevent duplicate emails from being processed, all inbound email must be delivered to only the current primary Database Server. When Disaster Recovery is configured, the mail process is configured to be running only on the current primary Database Server. 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.
- Include the hostname of the Database Server for Disaster Recovery at the highest MX-level (making that hostname the lowest priority).


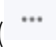
Disaster Recovery with Two Appliances

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 includes the following topics:

<i>Prerequisites</i>	8
<i>Unique Host Names</i>	8
<i>Configuring Disaster Recovery</i>	9
<i>Failover</i>	15

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 host name on each appliance
- 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 must have an available IP address to configure as a [virtual IP](#)

Unique Host Names

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

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

NOTE: Data sent from the active appliance to the redundant appliance is compressed and encrypted *only* if you use DRBD Proxy. DRBD without DRBD Proxy does not compress and encrypt this data.

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.

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 -i
```

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

```
coro_install
```

The following prompt appears:

```
1) HA
2) DR
3) HA+DR
4) Quit
Please select the architecture you'd like to setup:
```

5. Enter "2". The following prompt appears:


```

1) Primary
2) Secondary
Please choose which node this is:

```
6. Enter "1". The following prompt appears:


```

Architecture: DR
Server Role: Primary
Is this information correct? (y/n)

```
7. Enter "y". The following prompt appears:


```

The hostname of this server is <hostname of this appliance>, is this right?
(y/n)

```
8. Enter "y". The following prompt appears:


```

Please choose the DRBD IP for this server:
1) <First IP address of this appliance>
2) <Second IP address of this appliance>
.
.
Number:

```
9. Enter the number for the IP address of the network connection for replication on this Primary appliance. The following prompt appears:


```

What is the hostname of the Secondary server:

```
10. Enter the hostname of the Secondary appliance. The following prompt appears:


```

Please enter the IP used for DRBD traffic for the Secondary server:

```
11. Enter the IP address of the network connection for replication on the Secondary appliance. The following prompt appears:


```

Is DRBD Proxy being used? (y/n)

```
12. If the appliances are not directly connected using a crossover cable, you must use DRBD Proxy. If you are using DRBD proxy, enter "y". If you are not using DRBD Proxy, enter "n".
13. If you are using DRBD Proxy, go to step 18. If you are not using DRBD Proxy, the following prompt appears:


```

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

```
14. If you want to optionally add a virtual IP to the Disaster Recovery configuration, enter "y". If you do not want to add virtual IP to the Disaster Recovery configuration, enter "n".
15. If you entered "n", go to step 18. If you entered "y", the following prompt appears:


```

Please enter the Virtual IP Address:

```
16. If you are adding a virtual IP to the Disaster Recovery configuration, enter the virtual IP address. The following prompt appears:


```

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

```
17. If you are adding a virtual IP to the Disaster Recovery configuration, enter the CIDR netmask of the virtual IP address.

18. If you are not using DRBD Proxy, go to step 22. If you are using DRBD proxy, the following prompt appears:

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

19. If you are using DRBD Proxy, enter the maximum link speed between the two appliances.

20. The following prompt appears:

```
You have selected the following settings, please confirm if they are correct:
Architecture: DR
Node: Primary

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>

DRBD Disk: <partition to be used by DRBD>
DRBD Proxy: <whether DRBD proxy will be used>

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

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

```
Setting up the environment...
- Updating firewalld configuration, please be patient...
Setting up DRBD...
Editing Corosync config...
Setting up Corosync...
Complete, you can monitor the cluster status by typing 'crm_mon' (give it a
minute)

Coro_install completed successfully

coro_install has exited
```

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:

```
sudo -s
```

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

4. Run the following command:

```
coro_install
```

The following prompt appears:

```
1) HA
2) DR
3) HA+DR
4) Quit
Please select the architecture you'd like to setup:
```

5. Enter "2". The following prompt appears:

```
1) Primary
2) Secondary
Please choose which node this is:
```
6. Enter "2". The following prompt appears:

```
Architecture: DR
Server Role: Secondary
Is this information correct? (y/n)
```
7. Enter "y". The following prompt appears:

```
The hostname of this server is <hostname of this appliance>, is this right?
(y/n)
```
8. Enter "y". The following prompt appears:

```
Please choose the DRBD IP for this server:
1) <First IP address of this appliance>
2) <Second IP address of this appliance>
.
.
Number:
```
9. Enter the number for the IP address of the network connection for replication on this Secondary appliance. The following prompt appears:

```
What is the hostname of the Primary server:
```
10. Enter the hostname of the Primary appliance. The following prompt appears:

```
Please enter the DRBD IP for the Primary server:
```
11. Enter the IP address of the network connection for replication on the Primary appliance. The following prompt appears:

```
Is DRBD Proxy being used? (y/n)
```
12. If the appliances are not directly connected via a crossover cable, you must use DRBD proxy. If you are using DRBD proxy, enter "y". If you are not using DRBD proxy, enter "n".
13. If you are adding a virtual IP to the Disaster Recovery configuration, enter the virtual IP address. The following prompt appears:

```
I have detected the partition used for DRBD should be /dev/mapper/em7vg-db, is
this correct? (y/n)
```
14. Enter "y".
15. If you are not using DRBD proxy, go to step 22. If you are using DRBD proxy, the following prompt appears:

```
Please enter the max link speed to the DR system in megabytes/second:
```
16. If you are using DRBD proxy, enter the maximum link speed between the two appliances.
17. The following prompt appears:

```
You have selected the following settings, please confirm if they are correct:
Architecture: DR
Node: Secondary
```

```

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

DRBD Disk: <partition to be used by DRBD>
DRBD Proxy: <whether DRBD proxy will be used>

Is this information correct? (y/n)

```

19. Enter "y". If proxy is not in use, the following output appears:

```

Setting up SSH keys...
You will be prompted to enter the password for <IP address of Primary
appliance>

em7admin@<IP address of Primary appliance>'s password:

```

20. Enter the password for the em7admin user on the Primary appliance. The following output appears:

```

Setting up the environment...
- Updating firewalld configuration, please be patient...
Setting up DRBD...
Editing Corosync config...
Setting up Corosync...
Complete, you can monitor DRBD sync status by using 'cat /proc/drbd' (it can
take a sec)

Please license the appliance at this time WITHOUT failing over
Failover cannot occur until DRBD is fully synced

Coro_install completed successfully

```

Licensing the Secondary Appliance

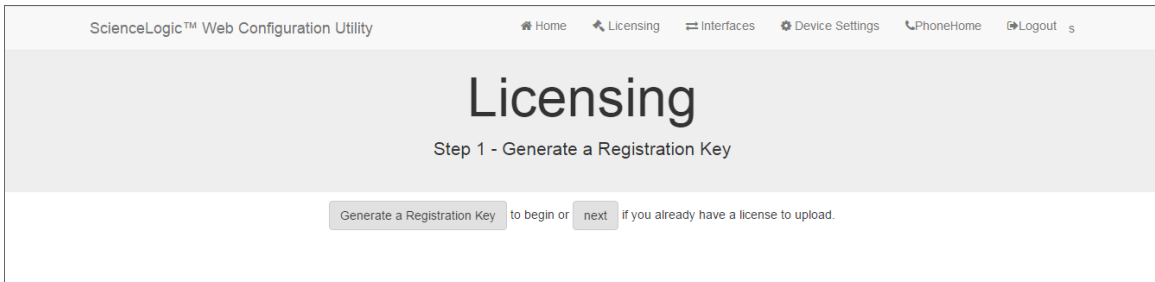
Perform the following steps to license the Secondary appliance:

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

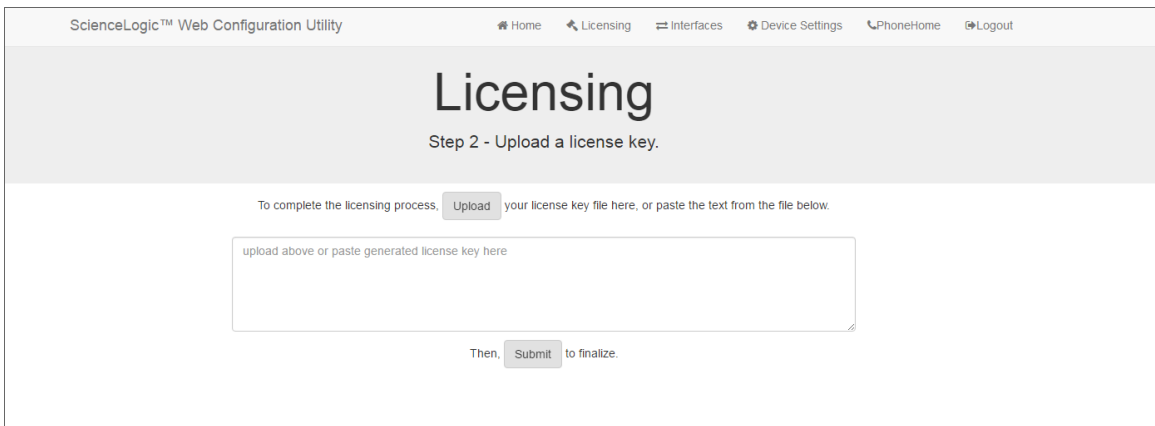
```
https://<ip-address-of-appliance>:7700
```

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.

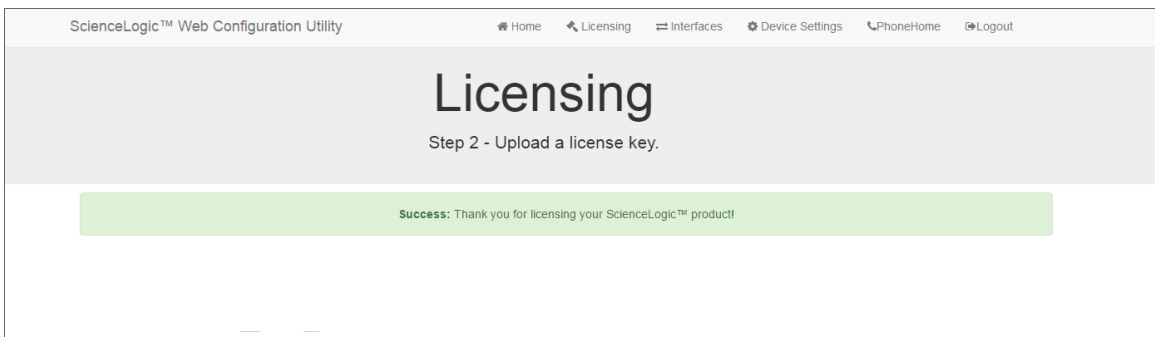
2. You will be prompted to enter your user name and password. Log in as the "em7admin" user with the password you configured using the Setup Wizard.
3. The **Configuration Utilities** page appears. Click the **[Licensing]** button. The **Licensing Step 1** page appears:



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 Customer Portal at <https://portal.sciencelogic.com/portal>. 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:



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.

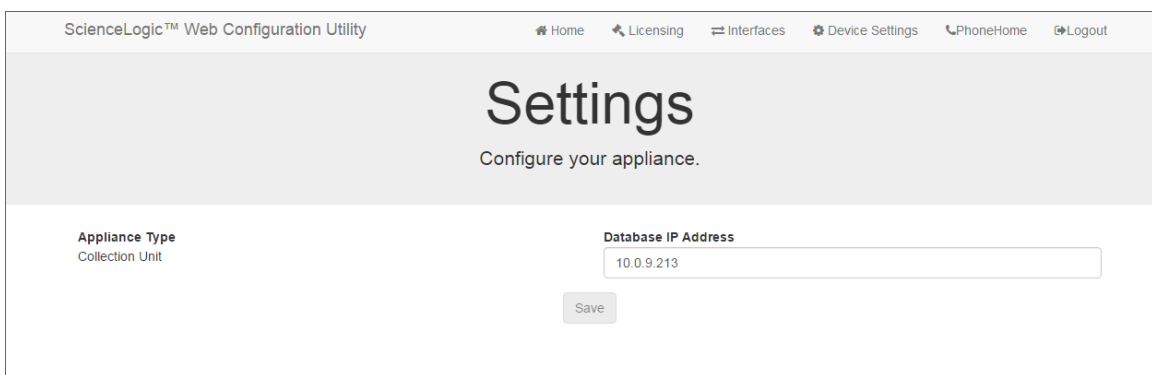
To configure a Data Collector or Message Collector to use the new configuration:

1. You can log in to the Web Configuration Utility using any web browser supported by SL 1 . 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 first demote the Primary appliance during failover.

- After demoting, your 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 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:
- **Make sure to power down the Primary appliance.** This step is required to avoid a split-brain configuration where 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:

```
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. Log in to the console of the Secondary appliance as the em7admin user.

8. Run the following command:

```
sudo -s
```

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

10. 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:

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

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

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

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

13. 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:

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

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

14. 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.
15. 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.

Failover When the Primary Database Appliance is Inaccessible

If you need to perform 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. **Make sure to 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. Log in to the console of the Secondary appliance as the em7admin user.

3. Run the following command:

```
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) Promote DRBD
4) Stop Pacemaker
5) Resource Status
6) Quit
```

Please enter the number of your choice:

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

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

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

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

8. 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:

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

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

9. 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.
10. 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.

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:

ScienceLogic™ Web Configuration Utility

Home Licensing Interfaces Device Settings PhoneHome Logout

Settings

Configure your appliance.

Appliance Type
Administration Portal

Database IP Address
10.0.9.213

Database Username
root

Database Password (change only)

Confirm Database Password

GUI Username
root

GUI Password (change only)

Confirm GUI Password

Save

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. 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 from Email" policy or an "events from Email" policy.

To complete the verification process, execute the following command:

```
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.

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/Unknown ds:UpToDate/DUnknown".

NOTE: If your two appliances cannot communicate, your output will include "ro:Primary/Unknown ds:UpToDate/UpToDate". 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 includes the following topics:

<i>Prerequisites</i>	22
<i>Addressing the Cluster</i>	22
<i>Configuring Heartbeat IP Addresses</i>	23
<i>Testing the Heartbeat Network</i>	23
<i>Configuring High Availability</i>	24
<i>Failover</i>	30

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 host name on each appliance
- Connect the two appliances that will be members of the High Availability cluster using a crossover Ethernet cable.
- Determine the virtual IP address for the cluster. 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.

Unique Host Names

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

1. Log in to the console of the SL1 appliance as the em7admin user. The current host name appears before the command-prompt. For example, the login prompt might look like this, with the current host name 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 host name, 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. The virtual IP 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 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 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 addresses.
Database Dynamic Applications	Use the Virtual IP 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 first configure a **heartbeat network** between the appliances. 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. Navigate to the following directory: **/etc/sysconfig/network-scripts**.
3. Identify the file corresponding to the heartbeat adapter. Adapter files are named **ifcfg-*<if name>***.
4. Edit or add the following lines to the file you identified:

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

5. Run the following command:

```
Ifup <Heartbeat adapter>
```

6. 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
```

Configuring High Availability

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

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:

```
coro_install
```

The following prompt appears:

```
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:

```
1) Primary  
2) Secondary  
Please choose which node this is:
```

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


```
Architecture: HA
Server Role: Primary
Is this information correct? (y/n)
```

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

```
The hostname of this server is <hostname of this appliance>, is this right?
(y/n)
```

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

```
Please choose the DRBD IP for this server:
1) <First IP address of this appliance>
2) <Second IP address of this appliance>
.
.
Number:
```

9. Enter the number for the IP address of the network connection for replication on this Primary appliance. The following prompt appears:

```
Please choose the DRBD IP for this server:
1) <First available IP address of this appliance>
2) <Second available IP address of this appliance>
.
.
Number:
```

10. Enter the number for the heartbeat IP address you configured for this appliance. The following prompt appears:

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

11. Enter the IP address of the network connection for replication on the Secondary appliance. The following prompt appears:

```
Please enter the IP used for HEARTBEAT traffic for the Secondary server the
corresponds to <HEARTBEAT IP address of first appliance>:
```

12. Enter the hostname of the Secondary appliance. The following prompt appears::

```
Please enter the heartbeat ip for the Secondary server:
```

13. Enter the heartbeat IP address you configured on the Secondary appliance. The following prompt appears:

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

14. Enter the virtual IP address for the cluster. The following prompt appears:

```
Please enter the CIDR for the Virtual IP:
```

15. Enter the CIDR netmask of the virtual IP address. The following prompt appears:

```
I have detected the partition used for DRBD should be /dev/mapper/em7vg-db, is
this correct? (y/n)
```

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

```
You have selected the following settings, please confirm if they are correct:
Architecture: HA
Node: Primary

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

```
Node 1 DRBD/Heartbeat IP: <DRBD IP address you entered for this appliance>
Node 1 Primary IP: <PRIMARY IP you entered for this appliance>
Node 2 Hostname: <host name of the Secondary appliance>
Node 2 DRBD IP/Heartbeat: <DRBD IP address you entered for the Secondary
appliance>
Node 2 Primary IP: <Primary IP of the secondary appliance>

Virtual IP: <Virtual IP Address>
Virtual IP CIDR: <Virtual IP CDIR>
Primary Broadcast: <Primary interface broadcast IP>
Heartbeat Broadcast: 169.254.1.3
DRBD Disk: <partition to be used by DRBD>
DRBD Proxy: No
```

Is this information correct? (y/n)

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

```
Setting up the environment...
- Updating firewallld configuration, please be patient...
Setting up DRBD...
Editing Corosync config...
Setting up Corosync...
Complete, you can monitor the cluster status by typing 'crm_mon' (give it a
minute)
```

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:

```
sudo -s
```

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

```
coro_install
```

The following prompt appears:

```
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:

```
1) Primary
2) Secondary
Please choose which node this is:
```

6. Enter "2". The following prompt appears:

```
Architecture: HA
Server Role: Secondary
```

Is this information correct? (y/n)

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

The hostname of this server is *<hostname of this appliance>*, is this right?
(y/n)

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

Please choose the IP used for HEARTBEAT traffic for this server:
1) *<First IP address of this appliance>*
2) *<Second IP address of this appliance>*
.
.
Number:

9. Enter the number for the heartbeat IP address that you configured for this appliance. The following prompt appears:

What is the hostname of the Primary server:

10. Enter the hostname of the Primary appliance. The following prompt appears:

Please enter the DRBD IP for the Primary server:

11. Enter the IP address of the network connection for replication on the Primary appliance. The following prompt appears:

Please enter the PRIMARY IP for the Primary server:

12. Enter the heartbeat IP address you configured on the Primary appliance. The following prompt appears:

I have detected the partition used for DRBD should be */dev/mapper/em7vg-db*, is this correct? (y/n)

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

You have selected the following settings, please confirm if they are correct:
Architecture: HA
Node: Secondary

Node 1 Hostname: *<host name of this appliance>*
Node 1 DRBD/Heartbeat IP: *<DRBD IP address you entered for this appliance>*
Node 1 Primary IP: 169.254.1.1
Node 2 Hostname: *<host name of the Primary appliance>*
Node 2 DRBD/Heartbeat IP: *<DRBD IP address you entered for the Primary appliance>*
Node 2 Primary IP: 169.254.1.2

Virtual IP: *<Virtual IP Address>*
Virtual IP CIDR: *<Virtual IP CIDR>*
Primary Broadcast: *<Primary interface broadcast IP>*
Heartbeat Broadcast: 169.254.1.3
DRBD Disk: *<partition to be used by DRBD>*
DRBD Proxy: No

Is this information correct? (y/n)

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

Setting up SSH keys...

You will be prompted to enter the password for <IP address of Primary appliance>

em7admin@<IP address of Primary appliance>'s password:

15. Enter the password for the em7admin user on the Primary appliance. The following output appears:

```
Setting up the environment...
- Updating firewall configuration, please be patient...
Setting up DRBD...
Editing Corosync config...
Setting up Corosync...
Complete, you can monitor the cluster status by typing 'crm_mon' (give it a
minute)
```

```
Please license the appliance at this time WITHOUT failing over
Failover cannot occur until DRBD is fully synced
```

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 SL 1. The address of the Web Configuration Utility is in the following format:

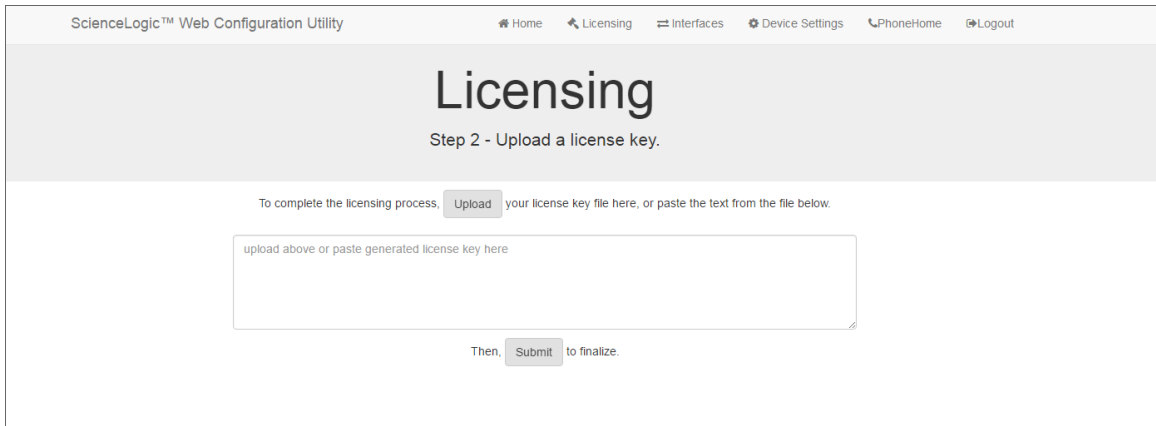
```
https://<ip-address-of-appliance>:7700
```

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 user name 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 Customer Portal at <https://portal.sciencelogic.com/portal>. 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:



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 SL 1 . 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

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 or Secondary appliance as the em7admin user.
2. Run the following command:

```
sudo systemctl stop pacemaker
```
3. When prompted, enter the password for the em7admin user.
4. After one minute, run the following command:

```
sudo systemctl start pacemaker
```

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:

```
silo_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 form Email" policy or an "events from Email" policy.

To complete the verification process, execute the following command:

```
sudo systemctl start pacemaker
```


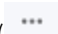
High Availability and Disaster Recovery with Three Appliances

Overview

This chapter describes how to configure three appliances: two appliances in a High Availability cluster with a third appliance 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 includes the following topics:

<i>Prerequisites</i>	33
<i>Licensing DRBD Proxy</i>	33
<i>Addressing the Cluster</i>	34
<i>Reconfiguring an Existing High Availability System</i>	34
<i>Configuring Heartbeat IP Addresses</i>	35
<i>Testing the Heartbeat Network</i>	36
<i>Configuring Three Appliances for High Availability and Disaster Recovery</i>	36
<i>Failover</i>	44

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 host name on each appliance
- Connect the two appliances that will be members of the High Availability cluster using a crossover Ethernet cable.
- Determine the virtual IP address for the cluster. 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.

Unique Host Names

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

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

NOTE: Data sent from the active appliance to the redundant appliance is compressed and encrypted *only* if you use DRBD Proxy. DRBD without DRBD Proxy does not compress and encrypt this data.

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. The virtual IP 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 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 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 addresses.
Database Dynamic Applications	Use the Virtual IP 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 Heartbeat IP Addresses

To cluster two databases, you must first configure a **heartbeat network** between the appliances. 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. Navigate to the following directory: **/etc/sysconfig/network-scripts**.
3. Identify the file corresponding to the heartbeat adapter. Adapter files are named **ifcfg-*<if name>***.
4. Edit or add the following lines to the file you identified:

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

5. Run the following command:

```
Ifup <Heartbeat adapter>
```

6. 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
```

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. [Primary appliance in the High Availability cluster](#)
2. [Secondary appliance in the High Availability cluster](#)
3. [Disaster recovery 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:

```
coro_install
```

The following prompt appears:

```
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:

```
1) Primary
2) Secondary
3) DR
Please choose which node this is:
```

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

```
Architecture: HA+DR
Server Role: Primary
Is this information correct? (y/n)
```

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

```
The hostname of this server is <hostname of this appliance>, is this right?
(y/n)
```

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

```
Please choose the HEARTBEAT IP for this server:
1) <First IP address of this appliance>
2) <Second IP address of this appliance>
.
.
Number:
```

9. Enter the heartbeat IP address you configured for this appliance. The following prompt appears:

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

10. Enter the hostname of the Secondary appliance. The following prompt appears:

```
Please enter the IP used for HEARTBEAT traffic for the Secondary server:
```

11. Enter the IP address of the network connection for DRBD replication on the Secondary appliance. This is typically the secondary IP address of the appliance. The following prompt appears:

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

12. Enter the primary IP address you configured on the Secondary appliance. The following prompt appears:

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

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

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

14. Enter the IP address that the High Availability cluster will use to communicate with the Disaster Recovery appliance. The following prompt appears:

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

15. Enter the virtual IP address for the cluster. The following prompt appears:

```
Please enter the CIDR for the Virtual IP:
```

16. Enter the CIDR netmask of the virtual IP address. The following prompt appears:

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

17. Enter the maximum link speed between the two appliances. The following prompt appears:

```
You have selected the following settings, please confirm if they are correct:  
Architecture: HA+DR  
Node: Primary
```

```
Node 1 Hostname: <hostname of this appliance>  
Node 1 DRBD/Heartbeat IP: <DRBD IP address you entered for this appliance>  
Node 1 Primary IP: 169.254.1.1  
Node 2 Hostname: <hostname of the Secondary appliance>  
Node 2 DRBD/Heartbeat IP: <DRBD IP address you entered for the Secondary  
appliance>  
Node 2 Primary IP: 169.254.1.2  
Node 3 Hostname: <hostname of the disaster recovery appliance>  
Node 3 DRBD IP: <DRBD IP address you entered for the disaster recovery  
appliance>  
Virtual IP: <Virtual IP Address>  
Virtual IP CIDR: <Virtual IP CIDR>  
Primary Broadcast: <Primary interface broadcast IP>  
Heartbeat Broadcast: 169.254.1.3  
DRBD Disk: <partition to be used by DRBD>  
DRBD Proxy: Yes  
Max DR DRBD Sync Speed: <sync speed you entered>
```

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

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

```
Setting up the environment...  
- Updating firewalld configuration, please be patient...  
Installing Stacked DRBD...  
Editing Corosync config...  
Setting up Corosync...  
Complete, you can monitor the cluster status by typing 'crm_mon' (give it a  
minute)
```

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:

```
sudo -s
```

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

```
coro_install
```

The following prompt appears:

```
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:

```
1) Primary
2) Secondary
3) DR
Please choose which node this is:
```

6. Enter "2". The following prompt appears:

```
Architecture: HA+DR
Server Role: Secondary
Is this information correct? (y/n)
```

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

```
The hostname of this server is <hostname of this appliance>, is this right?
(y/n)
```

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

```
Please choose the IP used for HEARTBEAT traffic for this server:
1) <First IP address of this appliance>
2) <Second IP address of this appliance>
.
.
Number:
```

9. Enter the IP address of the network connection for DRBD replication on this Secondary High Availability appliance. This is typically the secondary IP address of the appliance. The following prompt appears:

```
What is the hostname of the Primary server:
```

10. Enter the hostname of the Primary appliance. The following prompt appears:

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

11. Enter the IP address of the network connection for DRBD replication on the Primary appliance. This is typically the secondary IP address of the appliance. The following prompt appears:

```
Please enter the PRIMARY IP for the Primary server:
```

12. Enter the Primary IP address you configured on the Primary appliance. The following prompt appears:

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

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

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

14. Enter the IP address that the High Availability cluster will use to communicate with the Disaster Recovery appliance. The following prompt appears:

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

15. Enter the maximum link speed between the two appliances. The following prompt appears:

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

Architecture: HA+DR

Node: Secondary

Node 1 Hostname: *<hostname of this appliance>*

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

Node 1 Primary IP: 169.254.1.1

Node 2 Hostname: *<hostname of the Primary appliance>*

Node 2 DRBD/Heartbeat IP: *<DRBD IP address you entered for the Primary appliance>*

Node 2 Primary IP: 169.254.1.2

Node 3 Hostname: *<hostname of the disaster recovery appliance>*

Node 3 DRBD IP: *<DRBD IP address you entered for the disaster recovery appliance>*

Virtual IP: *<Virtual IP Address>*

Virtual IP CIDR: *<Virtual IP CIDR>*

Primary Broadcast: *<Primary interface broadcast IP>*

Heartbeat Broadcast: 169.254.1.3

DRBD Disk: *<partition to be used by DRBD>*

DRBD Proxy: Yes

Max DR DRBD Sync Speed: *<sync speed you entered>*

Is this information correct? (y/n)

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

```
Setting up SSH keys...
```

```
You will be prompted to enter the password for <IP address of Primary appliance>
```

```
em7admin@<IP address of Primary appliance>'s password:
```

17. Enter the password for the em7admin user on the Primary appliance. The following output appears:

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

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

```
Installing Stacked DRBD...
```

```
Editing Corosync config...
```

```
Setting up Corosync...
```

```
Complete, you can monitor DRBD sync status by using 'cat /proc/drbd' (it can take a sec)
```

```
Please license the appliance at this time WITHOUT failing over
```

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

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:

```
coro_install
```


The following prompt appears:

- 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:

- 1) Primary
- 2) Secondary
- 3) DR

Please choose which node this is:

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

```
Architecture: HA+DR
Server Role: DR
Is this information correct? (y/n)
```

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

```
The hostname of this server is <hostname of this appliance>, is this right?
(y/n)
```

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

```
Please choose the DRBD IP for this server:
1) <First IP address of this appliance>
2) <Second IP address of this appliance>
.
.
Number:
```

9. Enter the IP address of the network connection for DRBD replication on this appliance. This is typically the primary IP address of the appliance. The following prompt appears:

```
Please enter the Virtual IP for the HA stack:
```

10. Enter the virtual IP address of the High Availability cluster. The following prompt appears:

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

11. Enter the maximum link speed between the two appliances. The following prompt appears:

```
You have selected the following settings, please confirm if they are correct:
Architecture: HA+DR
Node: DR
```

```
Node 1 Hostname: <host name of this appliance>
Node 1 DRBD IP: <DRBD IP address you entered for this appliance>
Node 2 Hostname: HA
Node 2 DRBD IP: <Virtual IP address you entered for the HA cluster>
DRBD Disk: <partition to be used by DRBD>
DRBD Proxy: <whether DRBD proxy will be used>
Max DR DRBD Sync Speed: <Sync speed you entered>
```

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

12. Enter "y" and enter the password for the em7admin user on the Primary appliance. The following output appears:

```
Setting up the environment...
- Updating firewall configuration, please be patient...
Setting up DRBD...
Editing Corosync config...
Setting up Corosync...
Complete, you can monitor DRBD sync status by using 'cat /proc/drbd' (it can
take a sec)

Please license the appliance at this time WITHOUT failing over
Failover cannot occur until DRBD is fully synced
```

Licensing the Secondary High Availability and Disaster Appliances


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.

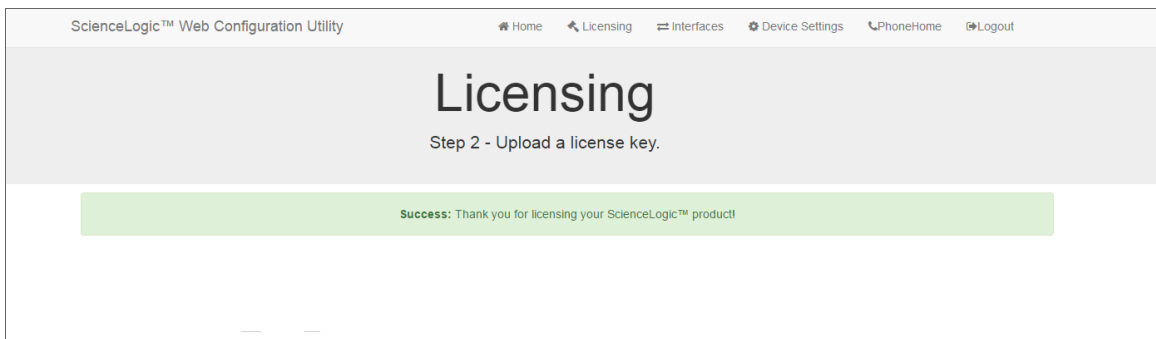
2. You will be prompted to enter your user name and password. Log in as the "em7admin" user with the password you configured using the Setup Wizard.
3. The **Configuration Utilities** page appears. Click the **[Licensing]** button. The **Licensing Step 1** page appears:



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 Customer Portal at <https://portal.sciencelogic.com/portal>. 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:



9. 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.

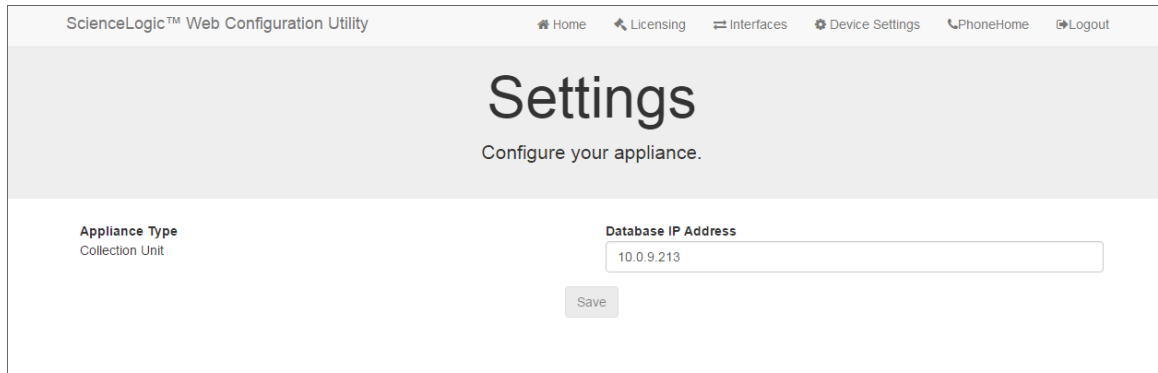
To configure a Data Collector or Message Collector to use the new configuration:

1. You can log in to the Web Configuration Utility using any web browser supported by SL 1 . 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

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 or Secondary appliance as the em7admin user.
2. Run the following command:

```
sudo systemctl stop pacemaker
```

3. When prompted, enter the password for the em7admin user.
4. After one minute, run the following command:

```
sudo systemctl start pacemaker
```

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. Run the following command:

```
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
```

Please enter the number of your choice:

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. Determine the device name of the primary network adapter using the command `ip addr`.

8. Run the following command:

```
ip addr add <VIP address/cidr> dev <device determined in step 7>
```

9. Start **drbdproxy** by running the following command:

```
service drbdproxy start
```

10. Log in to the console of the Disaster Recovery appliance as the em7admin user.

11. Run the following command:

```
sudo -s
```

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

13. 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:

14. 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 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:

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

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

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.

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

```
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) Promote DRBD
- 4) Stop Pacemaker
- 5) Resource Status
- 6) Quit

Please enter the number of your choice:

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

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

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

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

8. 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:

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

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

9. 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.

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:

```
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 appliances cannot communicate, your output will include "ro:Primary/Unknown ds:UpToDate/UpToDate". 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:

```
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 DRBD
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
```

8. Log in to the console of the appliance in the High Availability cluster that you want to promote to Primary as the em7admin user.

9. Run the following command:

```
sudo -s
```

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

11. Run the following command:

```
service drbdproxy stop
```

12. Run the following command to remove the manual VIP configuration created above:

```
ip addr del <VIP address/netmask> dev <primary network adapter>
```

For example: `ip addr del 192.168.1.100/24 dev ens160`

13. 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:

14. 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, 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.

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:

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. 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.

To complete the verification process, execute the following command:

```
sudo systemctl start pacemaker
```

© 2003 - 2020, 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



800-SCI-LOGIC (1-800-724-5644)

International: +1-703-354-1010