RUCKUS™
an ARRIS company

Cloudpath Enrollment System Replication Configuration Guide, 5.6

Supporting Software Cloudpath Release 5.6

Copyright, Trademark and Proprietary Rights Information

Export Restrictions

These products and associated technical data (in print or electronic form) may be subject to export control laws of the United States of America. It is your responsibility to determine the applicable regulations and to comply with them. The following notice is applicable for all products or technology subject to export control:

These items are controlled by the U.S. Government and authorized for export only to the country of ultimate destination for use by the ultimate consignee or end-user(s) herein identified. They may not be resold, transferred, or otherwise disposed of, to any other country or to any person other than the authorized ultimate consignee or end-user(s), either in their original form or after being incorporated into other items, without first obtaining approval from the U.S. government or as otherwise authorized by U.S. law and regulations.

Disclaimer

Limitation of Liability

Trademarks

Contents

Introduction

In addition to VMware-based high availability, Cloudpath supports replication (native clustering) capabilities. Clustering provides high-availability and high-scalability using two servers that are available for active use, with an active/standby option also available.

Configuration is performed entirely in the CLI.. Monitoring the cluster can be done through both the CLI and the UI.

Replication Topologies

Cloudpath supports the following topologies for replication, using a MariaDB database:

> **NOTE**
> For this release, only two-node clusters are supported.

- Active - Standby: In this topology, one node is designated *primarily active* and the other one is *secondarily standby*. Only one cluster node (the primarily active node by default) actively processes all enrollment requests from clients while the other one acts as a standby. The standby node is ready to take over to serve if the active node encounters problems and cannot serve. The failover is seamless to users. The primarily active node takes over again as soon as the problems have been resolved. This topology uses a virtual (or "floating") IP address that can be passed between identically configured servers to allow for failover.
- Active - Active (with or without Network address translation): Both cluster nodes simultaneously process enrollment requests through a load balancer. In this topology, both nodes are instantaneously in sync once initial setup is complete. If one node fails, the load balancer sends all requests to the other node .

  An active - active setup also has the concept of a primary node and secondary node. The primary node is the server that is composed of initial data that will be replicated to the secondary server during the replication setup. Be sure you know which node is your primary node before you proceed to configuration.

  > **NOTE**
  > When configuring an active - active setup with NAT, you will need both the private and public IP addresses of the nodes.

  A load balancer is required for this topology. How to configure the load balancer is outside the scope of this manual. You can consult the documentation for the load balancer of your choice. Many open-source load balancers are available. Be sure to configure your load balancer with the information it needs about the nodes that will compose your cluster.

  > **NOTE**
  > Do not use a Cloudpath System as a load balancer.

New System Activation: Primary Server vs. Add On:

When you first activate a new system, you are presented with a System Setup screen that contains the question: "Which Type Of Server is This?" For the node that will serve as your primary node in the cluster, select the "Standard Server (Default)" option. For the node that will serve as your secondary server, select the "Add On Server For Cluster" option. Refer to the *Cloudpath Enrollment System Deployment Administration Guide* for more information.

Prerequisites

To configure replication, set up two Cloudpath virtual appliances with the same OVA. When configuring these virtual appliances, all application and network properties must be the same for both servers in the cluster.

> **NOTE**
> For replication setup, all hosts must be able to communicate with all other nodes.

Each server must have a unique **FQDN Hostname** and network IP address. The DNS servers and gateway need not be the same as long as there is network connectivity between the servers.

Ports Used by Replication

Replication configuration opens the following ports between two servers:

- 3306 for database replication
- 5701 for application replication
- 8022 or 22 for SSH connection

Test Connectivity

When configuring replication, be sure that each system can resolve DNS and that each virtual appliance can communicate with the other server nodes to be replicated.

Configuring Replication

Cloudpath supports two types of replication. Proceed to the applicable topic:

- Configuring an Active - Standby Replication on page 5
- Configuring an Active - Active Replication on page 9

Configuring an Active - Standby Replication

You can configure a two-node, active - standby replication, which uses a virtual IP address.

The following figure depicts an active - standby setup, with example IP addresses that will be used throughout this topic.

FIGURE 1 Active - Standby Replication
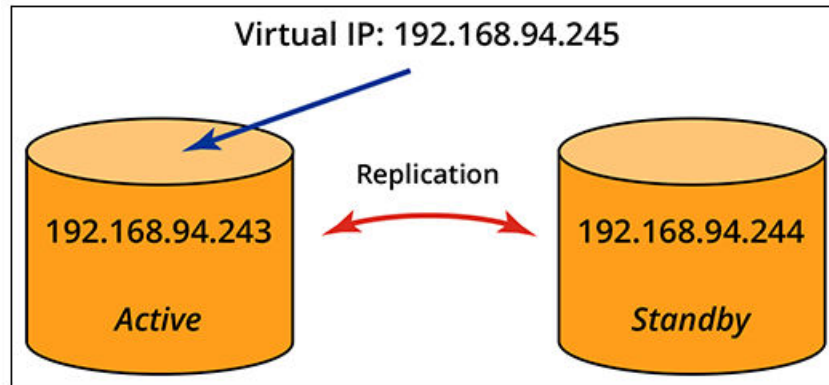
The steps below show you how to configure this type of cluster. It is assumed that you have already deployed two Cloudpath ES systems with the same .OVA file.

1. On the Cloudpath system that you want to be the "active" system (192.168.94.243 in the figure), open a terminal window and log in to the service account:

   • Port 8022 with SSH unless you change the port number

   • Username: cpn_service

   • Password: Enter the password.

   The command-line configuration utility prompt (#) should display.

2. Enter the command: **replication setup**

3. When presented with cluster-type options, select **Active-Standby** and continue.

4. When you are asked if you want the node you are on to be the primary database node, answer yes and continue:

   ```
   Do you want this node (192.168.94.243) to be the primary database node? [y/n]
   ```

5. Enter the IP address of the standby node when you are prompted for the "Secondary node IP:"

   ```
   Secondary database cluster node configuration:

     Secondary node IP: 192.168.94.244
   ```

6. Enter 8022 for the SSH port of the standby node when prompted (unless you have changed this port number):

   ```
   Secondary node SSH Port: 8022
   ```

7. When presented with the following warning, you can answer yes and proceed:

   ```
   --------------------------------------------------------------------------
    WARNINGS:
       Configuring the cluster will cause any data existing on 192.168.94.244 be
       deleted and replaced with the data currently on 192.168.94.243
   --------------------------------------------------------------------------

    Do you want to continue? [y/n]
   ```

8. You will then be presented with the following to confirm that your setup is what you expect:

```
192.168.94.244 will be the secondary database cluster node.

------------------------------------------------------------------------------
  Database cluster has been defined:
    Primary: 192.168.94.243   ssh port# 8022
    Secondary: 192.168.94.244   ssh port# 8022

    Cluster Type: Active-Standby
------------------------------------------------------------------------------

  Is this correct? [y/n]
```

If the information above is correct (as it is for this example), answer yes and proceed.

9. You are then prompted to enter the Virtual IP address. Be sure that the Virtual IP you enter is in the same domain as the primary and secondary nodes and that you enter an IP address not currently in use, then proceed:

```
Virtual IP: 192.168.94.245
```

10. You are then presented with the following:

```
------------------------------------------------------------------------------
  Cluster DNS Name - The cluster as a whole must have a single unified DNS name
                     it is known as.  This is most likely the fully qualified
                     domain name (FQDN) of the HTTPS certificate for the
                     enrollment portal as a whole.

                     This DNS name in most cases will resolve to the load
                     balancer fronting the cluster. Example:
                     Node 1 : node1.my-company.com
                     Node 2 : node2.my-company.com
                     Cluster DNS Name : wifi-onboarding.my-compandy.com

  IMPORTANT: If certificate based enrollments already exist you must re-use
             the hostname that was in place at the time of the enrollments
             occurred as the "Cluster DNS Name".  If this name is modified,
             all existing certificates will be invalidated as their OCSP URL
             will no longer match the answering address of the CA (which must
             be the cluster as a whole.)

  Cluster DNS Name: jeff243.cloudpath.net
    - Please ensure this is correct.
    - This name should match up with the host portion of the "OCSP URL" within
      the browser based Admin UI.  See "Certificate Authority > Manage CAs"
      If this is a new install you may want to create a new root CA if this
      does not match.
    - This value may not be changed after certificate based enrollments occur
      without requiring users to re-enroll.
------------------------------------------------------------------------------

  Do you want to keep jeff243.cloudpath.net? [y/n]
```

It is recommended to enter **yes** and proceed.

> **NOTE**
> If you enter **no** and proceed, you will later need to follow instructions from the command dialogue to update certificates.

11. You are then presented with the following:

```
------------------------------------------------------------------------------
  Cluster Name is set: jeff243.cloudpath.net.
  At later time, you will need to make DNS changes:
    Assign jeff243.cloudpath.net to 192.168.94.245 (Virtual IP).
    Rename hostname of the primary Database node.
```

```
--------------------------------------------------------------------------
  Do you want to continue? [y/n]
```

Enter **yes** to proceed.

> **NOTE**
> The dialogue informs you that you will need to make some DNS changes. This is described following these numbered steps.

12. The process completes, as shown in the following example. Check that you receive "Replication cluster setup complete:"

```
Will establish ssh 192.168.94.244:8022 ...
 --------------------------------------------------------------------------
   * You will be prompted for the password of the remote VM 192.168.94.244.
 --------------------------------------------------------------------------
yes: standard output: Broken pipe
yes: write error
/usr/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "/home/repl_user/.ssh/id_rsa.pub"
cpn_service@192.168.94.244's password:
id_rsa.pub
     100%  406     1.1MB/s   00:00
Checking remote host 192.168.94.244 ...
Will copy cluster configuration to node 192.168.94.244 ...
Will establish ssh from 192.168.94.244 to 192.168.94.243 ...
SSH connection from 192.168.94.244 to 192.168.94.243: SSH test succeeded
Will update network configuration ...
Will update database server configuration ...
Will stop local application service ...
Will restart database service...
Will configure local database service permission ...
Will transfer application database to the secondary database node at 192.168.94.244 ...
Will restart local application service ...
Will provision new remote node 192.168.94.244 ...
Done provisioning new remote node 192.168.94.244 :
Will replicate web cluster name jeff243.cloudpath.net ...
Will replicate web cluster name jeff243.cloudpath.net on 192.168.94.244 ...
Done replicating web cluster name.
Will configure keepalived on 192.168.94.243 ...
Created symlink from /etc/systemd/system/multi-user.target.wants/keepalived.service to /usr/lib/
systemd/system/keepalived.service.
Done keepalived configuration on 192.168.94.243 ...
Will configure keepalived on on 192.168.94.244 ...
Created symlink from /etc/systemd/system/multi-user.target.wants/keepalived.service to /usr/lib/
systemd/system/keepalived.service.
Done keepalived configuration on 192.168.94.244 ...
Will enable local cluster monitoring service ...
Created symlink from /etc/systemd/system/multi-user.target.wants/cloudpath-repld.service to /etc/
systemd/system/cloudpath-repld.service.
Will enable cluster monitoring service on 192.168.94.244 ...
Created symlink from /etc/systemd/system/multi-user.target.wants/cloudpath-repld.service to /etc/
systemd/system/cloudpath-repld.service.
Replication cluster setup complete.
#
```

Changes to Make Before Using Cluster for Enrollments:

Before you begin using the cluster for enrollments, you need to make the DNS changes (using your DNS configuration tool) that were listed in the command dialogue as you entered your information. Those changes are necessary for the virtual IP to be able to failover to the secondary node if the primary node goes down. In the example shown above, two DNS changes are required:

```
Cluster Name is set: jeff243.cloudpath.net.
   At later time, you will need to make DNS changes:
     Assign jeff243.cloudpath.net to 192.168.94.245 (Virtual IP).
     Rename hostname of the primary Database node.
```

1. Assign the cluster name (jeff243.cloudpath.net) the virtual IP address 192.168.94.245.

NOTE

The cluster name *must* be assigned the virtual IP address.

2. Rename the host name of the primary node from jeff243.cloudpath.net to a different name so that it does not conflict with cluster name.

Configuring an Active - Active Replication

You can configure a two-node active - active replication, which requires a load balancer.

The following figure depicts an active - active setup, with example IP addresses that will be used throughout this topic.

FIGURE 2 Active - Active Replication
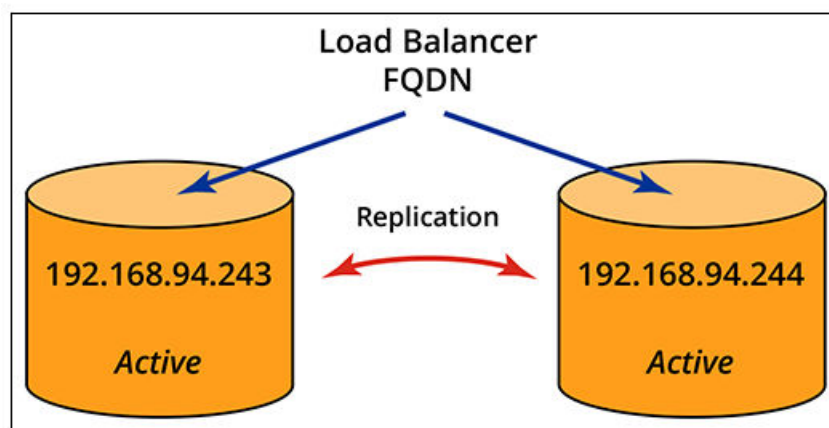
The steps below show you how to configure this type of cluster. It is assumed that you have already deployed two Cloudpath ES systems with the same .OVA file.

1. On the primary Cloudpath system, open a terminal window and log in to the service account:

    NOTE

    For purposes of this example, the node whose IP address is 192.168.94.243 is the primary node.

    - Port 8022 with SSH unless you change the port number
    - Username: cpn_service
    - Password: Enter the password.

    The command-line configuration utility prompt (#) should display.

2. Enter the command: **replication setup**

3. When presented with cluster-type options, select **Active-Active** and continue.

    You are then presented with the following options for Active - Active:

    ```
    -------------------------------------------------------------------------
      * Active-Active cluster configuration

      1) Without NAT - Cluster nodes are bidirectionally reachable directly with
                       the given IP address of each node.

      2) Behind NAT  - Cluster nodes are reachable via a public IP address which
    ```

```
                               differs from configured private IP address of each node.
        --------------------------------------------------------------------------------
```

4. This example continues with the assumption that option 1 (Without NAT) is selected, and the dialogue continues.

    NOTE
    If you are setting up an active - active system with NAT, you will need to know both the private and public IP
    addresses of the cluster nodes. Otherwise, the process is basically the same as active - active without NAT.

5. You are next presented with the following information, to which, after reading, you can answer yes to proceed :

    ```
    --------------------------------------------------------------------------------
     WARNINGS:
        1.  The node which replication setup is initiated (this node) will be
            considered the primary database node.  The data on this node will be
            retained and replicated to the second database node. Some future
            maintenance operations such as software updates will be need initiated
            from the primary database node.
        2.  All data existing on the secondary node will be deleted and replaced
            with the data currently on the primary node. Before proceeding, create
            a VM snapshot of all cluster nodes.
        3.  The application running on both nodes will be shutdown. This process
            should be run during a maintenance period.
    --------------------------------------------------------------------------------

     Do you want to continue? [y/n]
    ```

6. You are next presented with the following, to which you should answer yes and proceed:

    ```
     An existing database cluster configuration was not found. A new cluster configuration will be
     defined.

     Do you want this node (192.168.94.243) to be the primary database node? [y/n]
    ```

7. After answering yes, you are prompted to enter the IP address of the other node, which is 192.168.94.244 in this
    example, then continue:

    ```
     192.168.94.243 will be the primary database node.

     Secondary database cluster node configuration:

     Secondary node IP: 192.168.94.244
    ```

8. You are prompted to enter the SSH port for the secondary node. Enter 8022 (unless you have changed it), then proceed.

9. You are presented with the following information; answer yes and proceed:

    ```
     WARNINGS:
        Configuring the cluster will cause any data existing on 192.168.94.244 be
        deleted and replaced with the data currently on 192.168.94.243
    --------------------------------------------------------------------------------

     Do you want to continue? [y/n]
    ```

10. You are presented with the following; check to be sure it is correct and then proceed:

    ```
    --------------------------------------------------------------------------------
     Database cluster has been defined:
        Primary: 192.168.94.243   ssh port# 8022
        Secondary: 192.168.94.244   ssh port# 8022

        Cluster Type: Active-Active
        Network: Without NAT
    --------------------------------------------------------------------------------
    ```

```
           Is this correct? [y/n]
```

11. You are then prompted to enter the virtual name of your load balancer, which you have set up in advance. Enter the name and continue.

```
       Load Balancer Virtual Host Name: jeff245.cloudpath.net
```

12. You are next presented with the following information:

```
    --------------------------------------------------------------------------
    Cluster DNS Name - The cluster as a whole must have a single unified DNS name
                       it is known as.  This is most likely the fully qualified
                       domain name (FQDN) of the HTTPS certificate for the
                       enrollment portal as a whole.

                       This DNS name in most cases will resolve to the load
                       balancer fronting the cluster. Example:
                       Node 1 : node1.my-company.com
                       Node 2 : node2.my-company.com
                       Cluster DNS Name : wifi-onboarding.my-compandy.com

    IMPORTANT: If certificate based enrollments already exist you must re-use
               the hostname that was in place at the time of the enrollments
               occurred as the "Cluster DNS Name".  If this name is modified,
               all existing certificates will be invalidated as their OCSP URL
               will no longer match the answering address of the CA (which must
               be the cluster as a whole.)

    Cluster DNS Name: jeff243.cloudpath.net
      - Please ensure this is correct.
      - This name should match up with the host portion of the "OCSP URL" within
        the browser based Admin UI.  See "Certificate Authority > Manage CAs"
        If this is a new install you may want to create a new root CA if this
        does not match.
      - This value may not be changed after certificate based enrollments occur
        without requiring users to re-enroll.
    --------------------------------------------------------------------------

       Do you want to keep jeff243.cloudpath.net? [y/n]
```

It is recommended to enter **yes** and proceed.

> **NOTE**
> If you enter **no** and proceed, you will later need to follow instructions from the command dialogue to update certificates.

13. You are then presented with the following:

```
    --------------------------------------------------------------------------
     Cluster DNS Name is set: jeff243.cloudpath.net.
     At later time, you will need to make DNS changes:
       Rename Load Balancer's Host Name to jeff243.cloudpath.net.
       Rename hostname of the primary database node.
    --------------------------------------------------------------------------

       Do you want to continue? [y/n]
```

Enter **yes** to proceed.

> **NOTE**
> The dialogue informs you that you will need to make some DNS changes. This is described following these numbered steps.

14. The process completes, as shown in the following example. Check that you receive "Replication cluster setup complete:"

```
Will establish ssh 192.168.94.244:8022 ...
 ----------------------------------------------------------------------
     * You will be prompted for the password of the remote VM 192.168.94.244.
 ----------------------------------------------------------------------
yes: standard output: Broken pipe
yes: write error
/usr/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "/home/repl_user/.ssh/id_rsa.pub"
cpn_service@192.168.94.244's password:
id_rsa.pub
    100%  406   940.8KB/s   00:00
Checking remote host 192.168.94.244 ...
Will copy cluster configuration to node 192.168.94.244 ...
Will establish ssh from 192.168.94.244 to 192.168.94.243 ...
SSH connection from 192.168.94.244 to 192.168.94.243: SSH test succeeded
Will update network configuration ...
Will update database server configuration ...
Will stop local application service ...
Will restart database service...
Will configure local database service permission ...
Will transfer application database to the secondary database node at 192.168.94.244 ...
Will restart local application service ...
Will provision new remote node 192.168.94.244 ...
Done provisioning new remote node 192.168.94.244 :
Will replicate web cluster name jeff243.cloudpath.net ...
Will replicate web cluster name jeff243.cloudpath.net on 192.168.94.244 ...
Done replicating web cluster name.
Will enable local cluster monitoring service ...
Created symlink from /etc/systemd/system/multi-user.target.wants/cloudpath-repld.service to /etc/
systemd/system/cloudpath-repld.service.
Will enable cluster monitoring service on 192.168.94.244 ...
Created symlink from /etc/systemd/system/multi-user.target.wants/cloudpath-repld.service to /etc/
systemd/system/cloudpath-repld.service.
Replication cluster setup complete.
#
```

Changes to Make Before Using Cluster for Enrollments:

Before you begin using the cluster for enrollments, you need to make the DNS changes (using your DNS configuration tool) that were listed in the command dialogue as you entered your information. Those changes are necessary for the load balancer to be able to route traffic to both nodes in the cluster. In the example shown above, two DNS changes are required:

```
Cluster DNS Name is set: jeff243.cloudpath.net.
   At later time, you will need to make DNS changes:
     Rename Load Balancer's Host Name to jeff243.cloudpath.net.
     Rename hostname of the primary database node.
```

1. Rename the load balancer from jeff245.cloudpath.net to jeff243.cloudpath.net so that the load balancer and the cluster host names are the same.

   > **NOTE**
   > The cluster host name *must* be the host name of the load balancer.

2. Rename the host name of the primary node from jeff243.cloudpath.net to a different name so that it does not conflict with cluster name.

Monitoring the Cluster

You can monitor various aspects of the cluster in the user interface.

Once your cluster has been created, you can go to the **Administration > Replication** area of the Cloudpath UI.

If the cluster was successfully created, green check marks will appear to the right of the IP address of each node, as shown in the example below. If green check marks do not appear, expand the node information by clicking on the arrows, and click one of the log buttons (as shown in Figure 4.

FIGURE 3 Replication Status in the Cloudpath UI
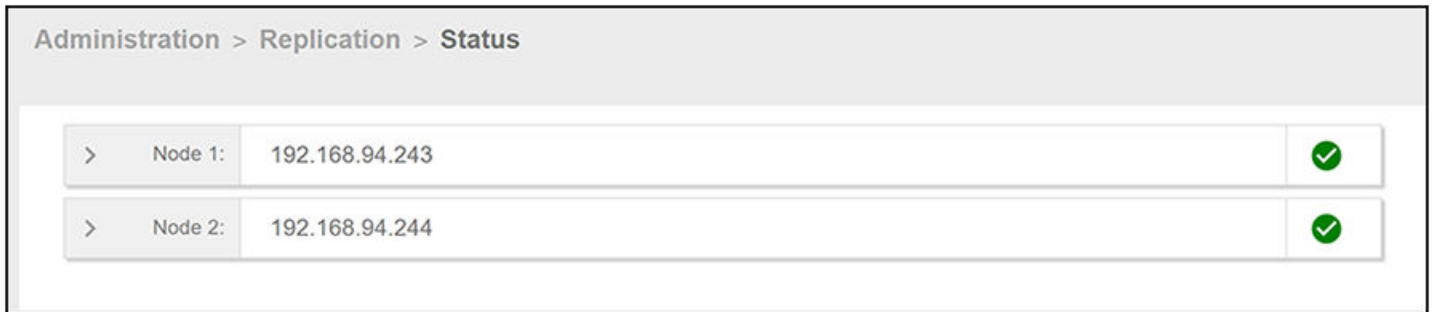
You can expand the information for the nodes by clicking on the arrow. The following figure shows an example of Node 1 expanded:

FIGURE 4 Detailed Information for a Node in the Cluster
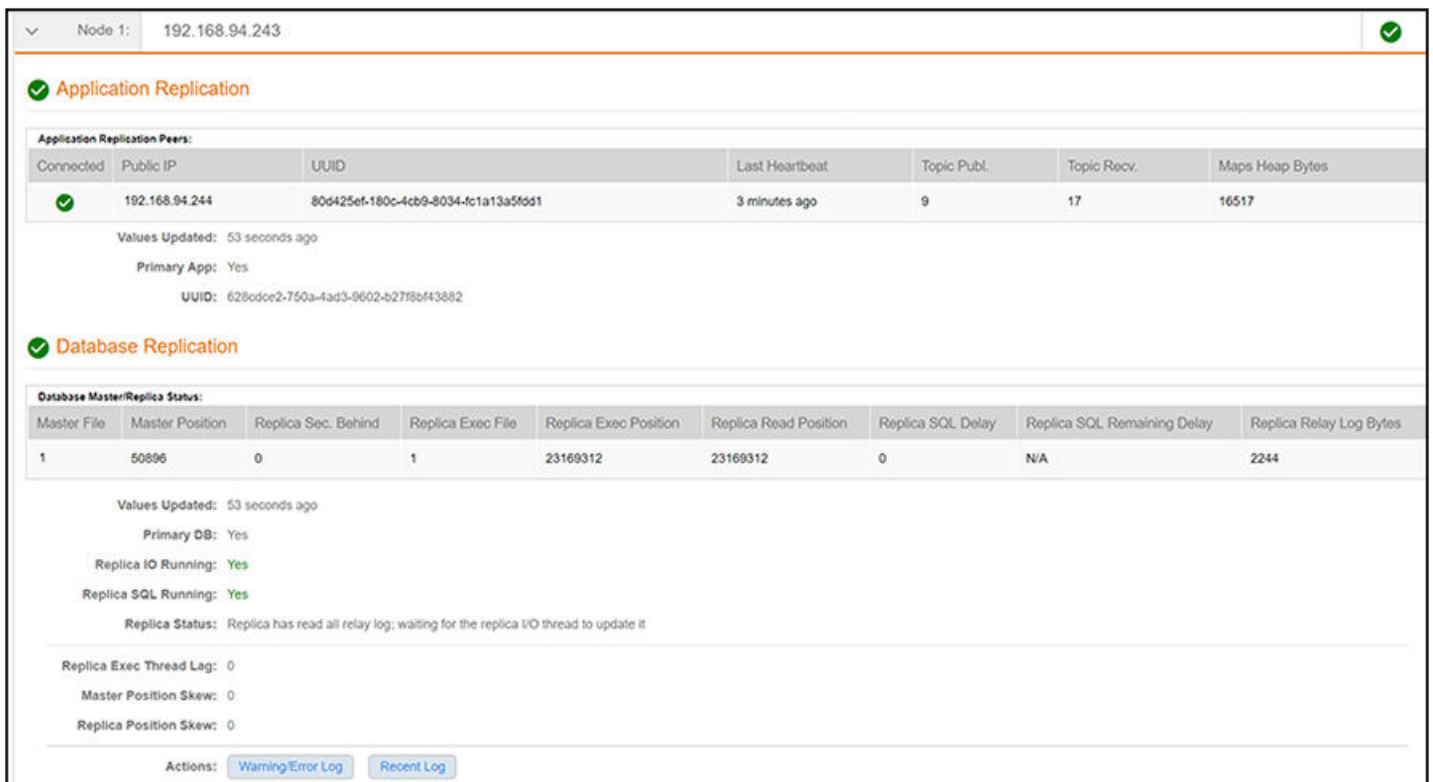
You can click the two log buttons at the bottom of the screen for additional information.

Running Replication Diagnostic Tests

Commands are available that show the status and configuration of your cluster. You can run these commands merely as a check, or if you suspect there is an issue with the cluster.

Showing the Replication Configuration

You can run the **replication show-config** command from either cluster node to view status about each node.

The following command (and corresponding output) is an example from the active (primary) node that was shown in Configuring an Active - Standby Replication on page 5:

```
# replication show-config
----------------------------------------------------------------------------
    Primary - 192.168.94.243   ssh port# 8022
    Secondary - 192.168.94.244   ssh port# 8022

    Type - ACTIVE-STANDBY

    Virtual IP - 192.168.94.245

    Cluster Name - jeff243.cloudpath.net

  INET:
   inet 192.168.94.243/22 brd 192.168.95.255 scope global ens32
   inet 192.168.94.245/32 scope global ens32
```

NOTE

Note that, under "INET," the virtual IP address of 192.168.94.245 is shown, indicating that the primary node is currently the active node. If the **replication show-config** command were run on the standby node, the "INET" output would not list the virtual IP unless the standby node had taken over due to a failure on the primary node.

Showing the Replication Status

You can run the **replication show-status** command from either cluster node. The following command (and corresponding output) is an example from the cluster that was shown in Configuring an Active - Standby Replication on page 5:

```
# replication show-status
[sudo] password for cpn_service:
******************** MASTER(192.168.94.243) >>>>>>  REPLICA (192.168.94.244) ********************

Master (192.168.94.243) Status
-------------------------------------------------------------------------------------------
            File: mysql-bin.000009
        Position: 39455
-------------------------------------------------------------------------------------------

Replica (192.168.94.244) Status
-------------------------------------------------------------------------------------------
            Replica_IO_State: Waiting for master to send event
                 Master_Host: 192.168.94.243
             Master_Log_File: mysql-bin.000009
         Read_Master_Log_Pos: 39455
              Relay_Log_File: mysql-relay-log.000003
               Relay_Log_Pos: 37516
       Relay_Master_Log_File: mysql-bin.000009
            Replica_IO_Running: Yes
           Replica_SQL_Running: Yes
                  Last_Errno: 0
                  Last_Error:
         Exec_Master_Log_Pos: 39455
       Seconds_Behind_Master: 0
```

```
                Last_IO_Errno: 0
                Last_IO_Error:
               Last_SQL_Errno: 0
               Last_SQL_Error:
                    SQL_Delay: 0
          SQL_Remaining_Delay: NULL
      Replica_SQL_Running_State: Replica has read all relay log; waiting for the replica I/O thread to
update it
--------------------------------------------------------------------------------------------


******************** MASTER(192.168.94.244) >>>>>>  REPLICA (192.168.94.243) ********************

Master (192.168.94.244) Status
--------------------------------------------------------------------------------------------
                File: mysql-bin.000010
            Position: 23246974
--------------------------------------------------------------------------------------------

Replica (192.168.94.243) Status
--------------------------------------------------------------------------------------------
                Replica_IO_State: Waiting for master to send event
                     Master_Host: 192.168.94.244
                 Master_Log_File: mysql-bin.000010
             Read_Master_Log_Pos: 23246974
                  Relay_Log_File: mysql-relay-log.000003
                   Relay_Log_Pos: 3111
          Relay_Master_Log_File: mysql-bin.000010
               Replica_IO_Running: Yes
              Replica_SQL_Running: Yes
                      Last_Errno: 0
                      Last_Error:
             Exec_Master_Log_Pos: 23246974
           Seconds_Behind_Master: 0
                   Last_IO_Errno: 0
                   Last_IO_Error:
                  Last_SQL_Errno: 0
                  Last_SQL_Error:
                       SQL_Delay: 0
             SQL_Remaining_Delay: NULL
        Replica_SQL_Running_State: Replica has read all relay log; waiting for the replica I/O thread to
update it
********************************************************************************************
```

Some important items to check in the preceding output include making sure the following are true:

- Replication is occurring:

  ```
  Replica_IO_Running: Yes
          Replica_SQL_Running: Yes
  ```

- No errors:

  ```
  Last_Errno: 0
  Last_Error:
  Last_IO_Errno: 0
  Last_IO_Error
  ```

- Master and replica are in sync: The value of Master position is the same as (or very close to) the values of Read_Master_Log_Pos and Exec_Master_Log_Pos for both Master/Replica outputs above.

Pinging the Load Balancer (for Active - Active Setups Only)

You can check the availability of each node in the cluster from the load balancer by running the following command: https://*<cloudpathhostname>*/admin/pingnode

where *cloudpathhostname* is the fully qualified domain name of a node in the cluster.

The HTTP return should be a "Success" message, and the HTTP status code should be 200.

Checking Cluster Service Configuration

After replication setup, the HTTPS server names on both cluster nodes must match the Cluster DNS name that you chose while running the **replication setup** command.. You can run the **show config** command from each cluster node to view the cluster service configuration. For example, based on the active-standby setup described in Configuring an Active - Standby Replication on page 5, if the **show config** command is run on the secondary node, the output would appear as:

```
# show config
VM CPU: Intel(R) Xeon(R) CPU E5-2697 v3 @ 2.60GHz
VM Cores: 2 (CPUs * Cores Per CPU)
VM Memory: 6 GB
FIPS: Disabled
Network Config: Static
Hostname: jeff244.cloudpath.net
IP: 192.168.94.244
Netmask: 255.255.252.0
Gateway: 192.168.92.1
DNS: 192.168.222.254
Proxy: Disabled.
NTP: pool.ntp.org
Timezone: America/Denver [MDT]
HTTPS: Enabled on port 443 (80).
HTTPS Server Name: jeff243.cloudpath.net
Admin Web UI: Not restricted.
Hostname Restriction: Not restricted.
SSH Access: Permitted on port 8022.
#
```

The IP of the secondary node is displayed correctly as 192.168.94.244, and the HTTPS server name is displayed correctly as jeff243.cloudpath.net, which is the cluster name specified while running the **replication setup** command.

Firewall Settings

The Cloudpath Admin UI provides a table that lists the inbound and outbound traffic of your Cloudpath system. This information is dynamically generated based on the current system configuration and can changes as the system configuration is modified.

Go to **Administration** > **Firewall Requirements** to ensure that your firewall ports have been correctly configured for replication.

FIGURE 5 Firewall Requirements - Outbound

FIGURE 6 Firewall Requirements - Inbound

Replication ports are provided when replication is configured and running.

Checking Workflow Snapshots

Log on to the Cloudpath system administration portal for each cluster node, go to **Configuration > Workflows**, and ensure that there are no error messages for any accounts indicating the following: "Enrollment Portal URL and HTTPS server are mismatched."

Repairing the Cluster

You can attempt to repair the cluster without tearing it down by running the **replication repair** command.

An example of when you might try running the **replication repair** command is if you find that replication between the two cluster nodes is out of sync (refer to Showing the Replication Status on page 14). Check your log files in /var/log/replication/ repl.log or via the UI (**Administration > Replication > Status**), then expand the node information and click on the Log buttons) for indication of this type of "fatal error." Running the **replication repair** command may fix this sync issue.

> **NOTE**
> The repair tool must be run on the cluster node whose database is in good condition so that it provides a good replication database during the repair process.

The following command (and corresponding output) is an example from the cluster that was shown in Configuring an Active - Standby Replication on page 5

```
# replication repair
 -----------------------------------------------------------------------------
    * You will be prompted for the password of local system.
 -----------------------------------------------------------------------------
[sudo] password for cpn_service:

 -----------------------------------------------------------------------------
   WARNINGS:
     1.  Database Replication Repair routine must be performed on the database
         node whose database is in good condition and will provide a replication
         base during the database repair process. Before proceeding, create a VM
         snapshot of this database node.
     2.  During Database replication repair, any data existing on the other
         database node will be deleted and replaced with the data currently on
         this node.
 -----------------------------------------------------------------------------

  Do you want to continue? [y/n]

  Please answer yes or no.

  Do you want to continue? [y/n] y


   -----------------------------------------------------------------------------
    WARNING:
    Database cluster has been setup:
      Local Node: 192.168.94.243   ssh port# 8022
      Remote Node: 192.168.94.244   ssh port# 8022

      During database replication repair, any data existing on the remote node
      will be deleted and replaced with the data currently on the local node.
   -----------------------------------------------------------------------------

  Do you want to continue? [y/n] y


  Checking remote node 192.168.94.244 ...
  Will start database replication repair ...
  Will set maintenance mode for remote node 192.168.94.244 ...
```

```
ts=20190821 180936.427, lvl=INFO, action=RUN_SCRIPT, script=update-httpd-
mainten                                                 ance, message=Executing...
ts=20190821 180936.430, lvl=INFO, action=RUN_SCRIPT, script=update-httpd-
mainten                                                 ance, message=  - Param: /etc/httpd/conf/
httpd.conf
ts=20190821 180936.444, lvl=INFO, action=RUN_SCRIPT, script=update-httpd-
mainten                                                 ance, message=Will enable maintenance
mode restriction in httpd conf...
ts=20190821 180936.453, lvl=INFO, action=RUN_SCRIPT, script=update-httpd-
mainten                                                 ance, message=Done.
Will stop local cluster monitoring service on 192.168.94.243 ...
Will stop remote cluster monitoring service on 192.168.94.244 ...
Will stop database replication on 192.168.94.243 ...
Will stop remote database replication on 192.168.94.244 ...
Will transfer application database to the secondary database node at
192.168.94.                                                      244 ...
Will provision new remote node 192.168.94.244 ...
Done provisioning new remote node 192.168.94.244 :
Will unset maintenance mode for the remote node 192.168.94.244 ...
ts=20190821 180955.302, lvl=INFO, action=RUN_SCRIPT, script=update-httpd-
mainten                                                 ance, message=Executing...
ts=20190821 180955.305, lvl=INFO, action=RUN_SCRIPT, script=update-httpd-
mainten                                                 ance, message=  - Param: /etc/httpd/conf/
httpd.conf
ts=20190821 180955.321, lvl=INFO, action=RUN_SCRIPT, script=update-httpd-
mainten                                                 ance, message=Will not enable
maintenance mode restriction in httpd conf.
ts=20190821 180955.332, lvl=INFO, action=RUN_SCRIPT, script=update-httpd-
mainten                                                 ance, message=Done.
Will start local cluster monitoring service 192.168.94.243 ...
Will start remote cluster monitoring service on 192.168.94.244 ...
Database replication repair completed.
#
```

After the repair process completes, run the **replication show-status** command again to see if the issue appears to be resolved. Also, check your log files again.

Tearing Down the Cluster

If necessary, you can tear down the entire cluster to remove all replication configurations and to set both cluster nodes back to single Cloudpath Enrollment System mode.

You can run the **replication teardown** command from either cluster node. The following command (and corresponding output) is an example from the cluster that was shown in

```
# replication teardown
 ----------------------------------------------------------------------------
     * You will be prompted for the password of local system.
 ----------------------------------------------------------------------------


 ----------------------------------------------------------------------------
    WARNINGS:
      1.  Teardown cluster routine will disassemble the entire cluster.
      2.  After Cluster Teardown, all features that have been established to support
          replication cluster will be removed; however, the databasees on both primary
          and secondary nodes will be intact.
 ----------------------------------------------------------------------------

  Do you want to continue? [y/n] y


 ----------------------------------------------------------------------------
    WARNING:
    Cluster Database Nodes:
      Database node: 192.168.94.243   ssh port# 8022
      Database node: 192.168.94.244   ssh port# 8022
```

```
    During Cluster Teardown processing, with exception of the replicated
    secondary Database, all features that have been established to support
    replication cluster will be removed.
 ----------------------------------------------------------------------------

  Do you want to continue? [y/n] y


Checking remote host 192.168.94.244 ...
Will remove keepalived configuration on local node 192.168.94.243 ...
Removed symlink /etc/systemd/system/multi-user.target.wants/keepalived.service.
Will remove keepalived configuration on remote node 192.168.94.244 ...
Removed symlink /etc/systemd/system/multi-user.target.wants/keepalived.service.
Will reset web server cluster name on local node 192.168.94.243 ...
Will reset web server cluster name on remote node 192.168.94.244 ...
Will disable local cluster monitoring service on 192.168.94.243 ...
Removed symlink /etc/systemd/system/multi-user.target.wants/cloudpath-repld.service.
Will disable remote cluster monitoring service on 192.168.94.244 ...
Removed symlink /etc/systemd/system/multi-user.target.wants/cloudpath-repld.service.
Will stop replication on local node 192.168.94.243 ...
Will stop replication on remote node 192.168.94.244 ...
Remove Database repl_mon ...
Will remove database service permissions ...
Will revert database replication configuration ...
Will remove network configuration for replication ...
Will remove cluster configurations from all cluster nodes ...
Will restart local application services on 192.168.94.243 ...
Will restart remote application services on 192.168.94.244 ...
Cluster teardown completed.
#
```

You can now run a show-status command, and the output indicates that replication is not configured:

```
# replication show-status
Replication is not enabled on this system.
```

Upgrading the Cluster

Instructions on how to upgrade your cluster to the latest release are provided in the release notes.

Replication Commands

Replication commands are used for setting up replication, tearing down or repairing replication, upgrading, and viewing status and configuration of the cluster.

Replication commands are run from the command-line configuration utility prompt (#).

> **NOTE**
> From the command-line configuration utility, enter the **console** command to access the Linux shell. From the Linux shell, enter the **config** command to access the command-line configuration utility.

Command	Description
replication repair	Used to attempt to repair a cluster without tearing down the entire cluster.
replication setup	Creates a cluster of two Cloudpath ES nodes. You can create either an active - standby cluster or an active - active (with or without NAT) cluster. You are prompted from the command dialog to enter which options you are configuring.

Command	Description
replication show-config	Displays the configuration of the cluster. Example output: ``` -- -------- Primary - 192.168.94.243 ssh port# 8022 Secondary - 192.168.94.244 ssh port# 8022 Type - MULTISITE-ACTIVE-ACTIVE-NO-NAT Cluster Name - jeff245.cloudpath.net -- -------- ```
replication show-status	Displays database replication status.
replication teardown	Tears down the cluster and sets each cluster node back to a single-node Cloudpath system.