

Cloudpath Enrollment System Replication Configuration Guide, 5.8

Supporting Software Cloudpath Release 5.8

Copyright, Trademark and Proprietary Rights Information

© 2021 CommScope, Inc. All rights reserved.

No part of this content may be reproduced in any form or by any means or used to make any derivative work (such as translation, transformation, or adaptation) without written permission from CommScope, Inc. and/or its affiliates ("CommScope"). CommScope reserves the right to revise or change this content from time to time without obligation on the part of CommScope to provide notification of such revision or change.

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

THIS CONTENT AND ASSOCIATED PRODUCTS OR SERVICES ("MATERIALS"), ARE PROVIDED "AS IS" AND WITHOUT WARRANTIES OF ANY KIND, WHETHER EXPRESS OR IMPLIED. TO THE FULLEST EXTENT PERMISSIBLE PURSUANT TO APPLICABLE LAW, COMMSCOPE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, TITLE, NON-INFRINGEMENT, FREEDOM FROM COMPUTER VIRUS, AND WARRANTIES ARISING FROM COURSE OF DEALING OR COURSE OF PERFORMANCE. CommScope does not represent or warrant that the functions described or contained in the Materials will be uninterrupted or error-free, that defects will be corrected, or are free of viruses or other harmful components. CommScope does not make any warranties or representations regarding the use of the Materials in terms of their completeness, correctness, accuracy, adequacy, usefulness, timeliness, reliability or otherwise. As a condition of your use of the Materials, you warrant to CommScope that you will not make use thereof for any purpose that is unlawful or prohibited by their associated terms of use.

Limitation of Liability

IN NO EVENT SHALL COMMSCOPE, COMMSCOPE AFFILIATES, OR THEIR OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS, LICENSORS AND THIRD PARTY PARTNERS, BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL, EXEMPLARY OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER, EVEN IF COMMSCOPE HAS BEEN PREVIOUSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES, WHETHER IN AN ACTION UNDER CONTRACT, TORT, OR ANY OTHER THEORY ARISING FROM YOUR ACCESS TO, OR USE OF, THE MATERIALS. Because some jurisdictions do not allow limitations on how long an implied warranty lasts, or the exclusion or limitation of liability for consequential or incidental damages, some of the above limitations may not apply to you.

Trademarks

ARRIS, the ARRIS logo, COMMSCOPE, RUCKUS, RUCKUS WIRELESS, the Ruckus logo, the Big Dog design, BEAMFLEX, CHANNELFLY, FASTIRON, ICX, SMARTCELL and UNLEASHED are trademarks of CommScope, Inc. and/or its affiliates. Wi-Fi Alliance, Wi-Fi, the Wi-Fi logo, Wi-Fi Certified, the Wi-Fi CERTIFIED logo, Wi-Fi Protected Access, the Wi-Fi Protected Setup logo, Wi-Fi Protected Setup, Wi-Fi Multimedia and WPA2 and WMM are trademarks or registered trademarks of Wi-Fi Alliance. All other trademarks are the property of their respective owners.

Contents

| | |
|--|-----------|
| Introduction..... | 5 |
| Prerequisites..... | 7 |
| New System Activation: Primary Node vs. Add On..... | 7 |
| Ports Used by Replication..... | 9 |
| Test Connectivity..... | 10 |
| Supported Topologies..... | 11 |
| Active-Standby..... | 11 |
| Active-Active..... | 11 |
| Active-Active Setup With Two Database Nodes Only..... | 12 |
| Active-Active Setup With Two Node Groups..... | 12 |
| Active-Active Setup for Two Data Centers..... | 14 |
| Configuring Replication..... | 15 |
| Configuring an Active - Standby Replication..... | 15 |
| Configuring an Active-Active Replication..... | 19 |
| Adding Application Nodes to a Cluster..... | 23 |
| Removing an Application Node From a Cluster..... | 27 |
| Monitoring the Cluster..... | 29 |
| Running Replication Diagnostic Tests..... | 31 |
| Showing the Replication Configuration..... | 31 |
| Showing the Replication Status..... | 31 |
| Pinging the Load Balancer (for Active-Active Setups Only)..... | 33 |
| Checking Cluster Service Configuration..... | 33 |
| Firewall Settings..... | 33 |
| Checking Workflow Snapshots..... | 35 |
| Repairing Database Replication..... | 37 |
| Tearing Down the Cluster..... | 39 |
| Upgrading the Cluster..... | 41 |
| Replication Commands..... | 43 |

Introduction

In addition to VMware-based high availability, Cloudpath supports replication (native clustering) capabilities. Clustering provides high-availability and high-scalability using two database nodes (and, optionally, one or more application nodes) that are available for active use. An active/standby option, consisting of two database nodes, is also available.

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

Prerequisites

Terminology:

This manual includes the following terms:

- Application ("APP") node: A node in the cluster that sends all database requests to another node in the cluster that is running a replicated database server.
- Database ("DB") node: A node in the cluster that is running a replicated database server, and handles all database requests from one or more application nodes.

NOTE

In a two-node cluster, each node is an application node that also runs its own replicated database server.

Considerations:

To configure replication, set up a minimum of two Cloudpath virtual appliances with the same OVA. These are for the two database nodes. Also, set up an additional appliance with the same OVA for each application node that will be part of your cluster. When configuring these virtual appliances, all application and network properties must be the same for all nodes 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.

New System Activation: Primary Node vs. Add On

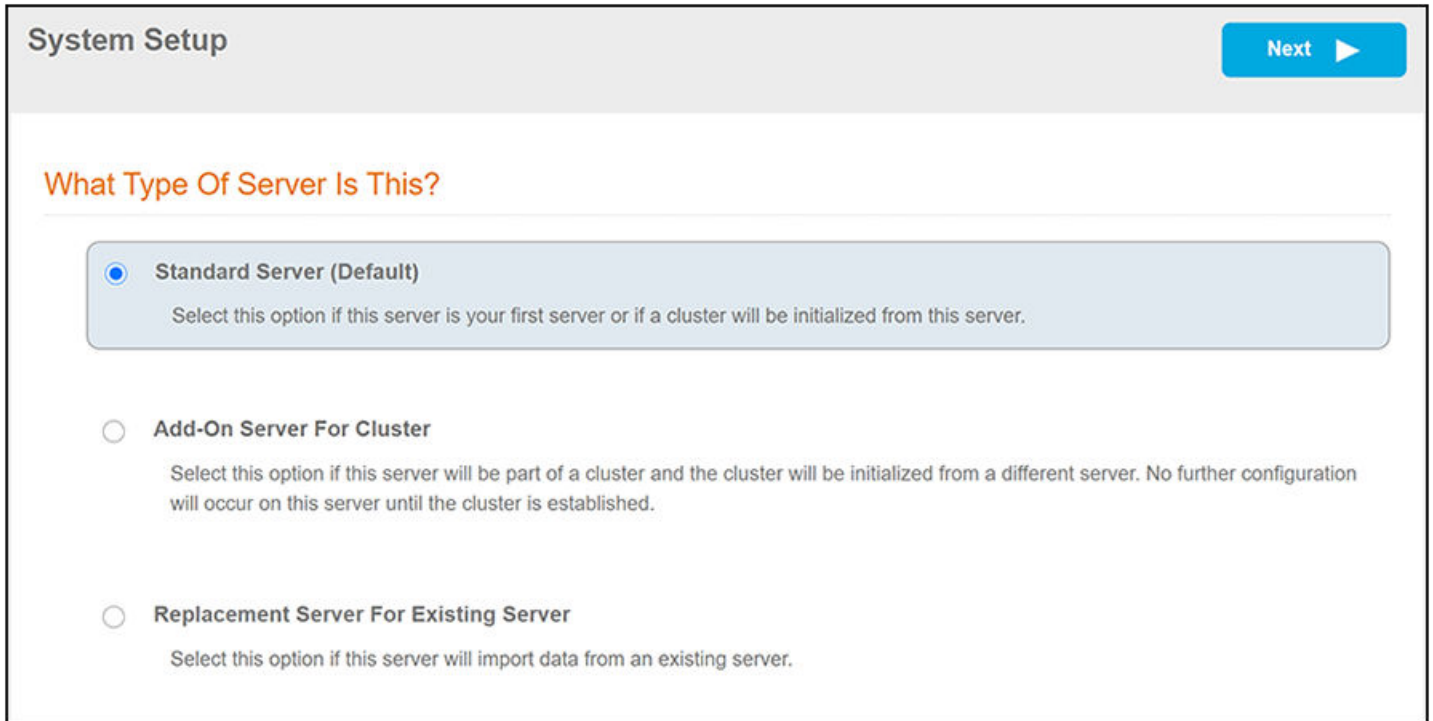
When you first activate a new system, you are presented with a System Setup screen that contains the question: "Which Type Of Server is This?"

For the node that will serve as your primary database node in the cluster, select the "Standard Server (Default)" option, as shown in the following illustration.

Prerequisites

New System Activation: Primary Node vs. Add On

FIGURE 1 Standard Server Selected for Primary Node in Cluster

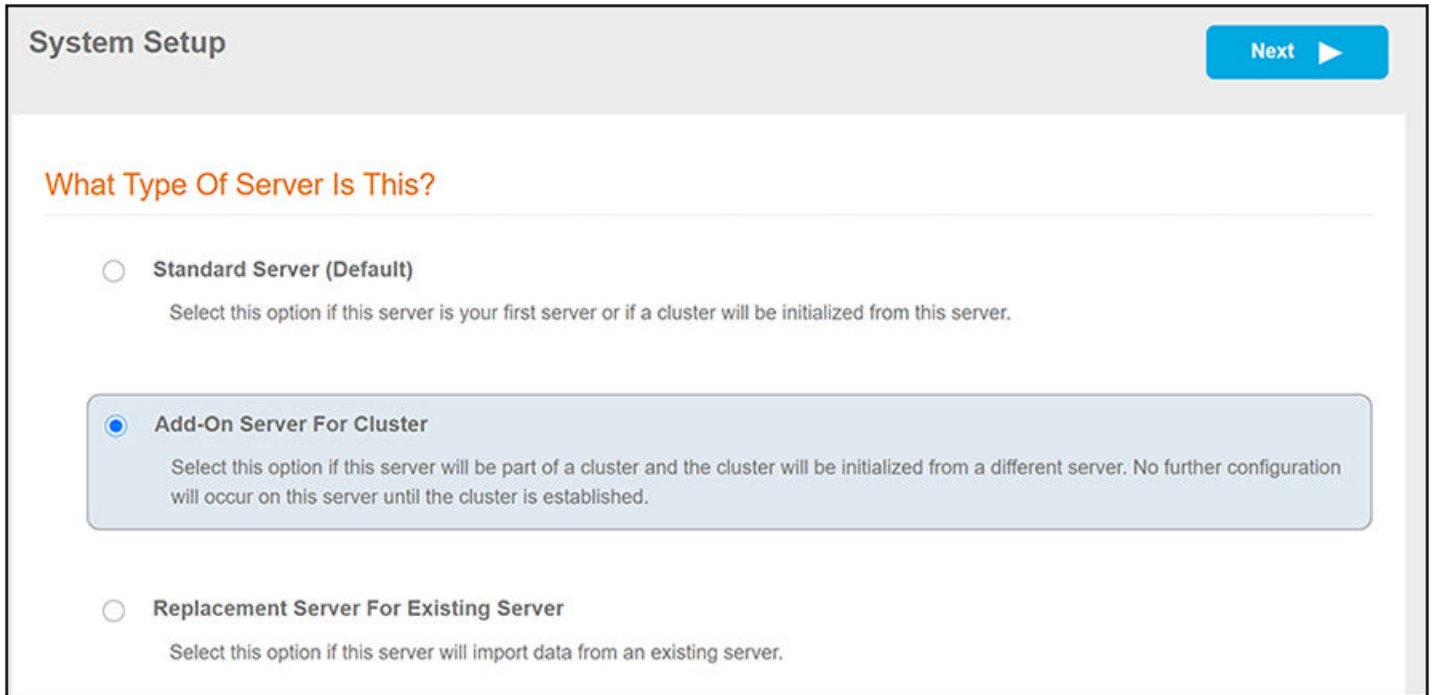


The screenshot shows a 'System Setup' window with a 'Next' button in the top right corner. The main heading is 'What Type Of Server Is This?'. There are three radio button options:

- Standard Server (Default)**
Select this option if this server is your first server or if a cluster will be initialized from this server.
- Add-On Server For Cluster**
Select this option if this server will be part of a cluster and the cluster will be initialized from a different server. No further configuration will occur on this server until the cluster is established.
- Replacement Server For Existing Server**
Select this option if this server will import data from an existing server.

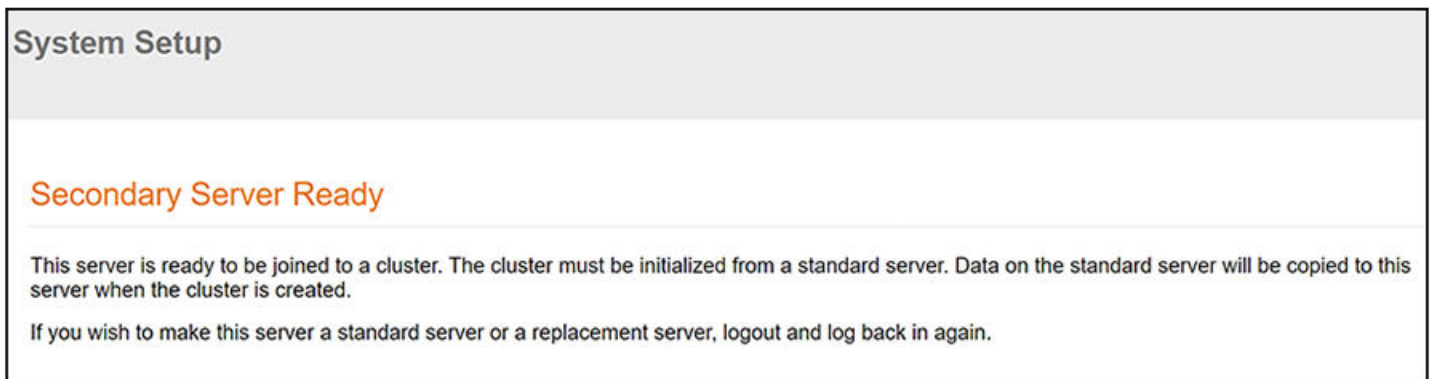
For all other nodes in the cluster (the secondary database node and any application nodes), select the "Add On Server For Cluster" option, as shown in the following illustration.

FIGURE 2 Add-On Server For Cluster Selected for All Cluster Nodes Except Primary Node



After you click **Next**, the following screen should be presented:

FIGURE 3 Secondary Server Ready Screen Indicating that a Node is Ready to be Added to a Cluster



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

Prerequisites

Test Connectivity

Test Connectivity

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

Supported Topologies


- Active-Standby..... 11
- Active-Active..... 11

Active-Standby

In this topology, one database node is designated *primarily active* and the other database node is *secondarily standby*

Only one cluster node (the primarily active node by default) actively processes all enrollment requests and Radius AAA 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.

FIGURE 4 Active-Standby



Active-Active

When both database nodes are configured as "active," you can have a setup that optionally also includes application nodes.

As with active-standby, an active-active setup also has the concept of a primary node and secondary node. The primary node is the database server that is composed of initial data that will be replicated to the secondary database node during the replication setup. Be sure you know which node is your primary node before you proceed to configuration. If the cluster will contain application (APP) nodes, each such node must be added from the command line interface of its corresponding database node.

Load Balancers:

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 make up your cluster. When you configure the cluster, you are prompted to enter the virtual host name of the load balancer.

NOTE

Do not use a Cloudpath System as a load balancer.

Supported Topologies

Active-Active

Network address translation (NAT):

Each type of active-active topology can be set up with or without NAT.

NOTE

If you are configuring an active-active setup with NAT, you will need both the private and public IP addresses of the nodes.


The following are the two basic types of active-active environments:

- Two database nodes that are instantaneously in sync once initial setup is complete. Refer to [Active-Active Setup With Two Database Nodes Only](#).
- "Node groups" which consists of two database nodes that are also each connected via read/write access to one or more application nodes. Refer to [Active-Active Setup With Two Node Groups](#).

Active-Active Setup With Two Database Nodes Only

This is the most basic type of active-active setup because it does not include any application nodes.

FIGURE 5 Active-Active With Two Database Nodes



With this setup, both nodes simultaneously process enrollment requests and Radius AAA requests through the load balancer. If one node fails, the load balancer sends all requests to the other node.


Configuration of this cluster can be performed entirely from the primary node. To proceed, refer to:

1. [Prerequisites](#) on page 7
2. [Configuring an Active-Active Replication](#) on page 19

Active-Active Setup With Two Node Groups

With an active-active setup, you have the flexibility of adding two application nodes to each of the two database nodes. One database node that also has one or more peer application nodes can be referred to as a *node group*, as shown in one possible configuration in the following illustration.

FIGURE 6 Cluster Using Node Groups



- Solid lines in the illustration denote read/write access between two nodes.
- Dashed lines in the illustration denote read/write access between an APP node and its secondary DB node, which only occurs if its primary DB node is not available. A DB node is the primary DB server to its peer APP nodes in the same node group.
- **NOTE**
A load balancer is required even though the illustration above does not depict one.

Additional Node Group Information:

The following rules apply to how node groups are configured and behave within a cluster:

- APP node database reads/writes are proxied to both DB servers, with high priority to the DB server in the same node group.
- Each APP node enters the cluster through one DB node.
- The DB node manages cluster join/unjoin for its peer APP nodes.
- Each DB node is automatically also an APP node.
- All DB and APP nodes share the same cluster service name (domain name portion of enrollment portal URLs)
- No service downtime occurs on existing cluster nodes when an APP node is added/removed.
- Master-master DB asynchronous replication forms a shared database.
- Database nodes failover/failback to ensure shared database high availability.
- APP nodes failover/failback through the load balancer.

Configuration Steps:


Configuration of a cluster that employs the use of node groups can be performed by referring to these topics:

1. [Prerequisites](#) on page 7
2. [Configuring an Active-Active Replication](#) on page 19
3. [Adding Application Nodes to a Cluster](#) on page 23

Active-Active Setup for Two Data Centers

There are many ways to incorporate the use of a cluster. This is just another example topology of how you could set up two data centers that are located a great distance apart where you could implement two load balancers.

FIGURE 7 Data Centers in a Cluster



This scenario would exhibit the following behavior:

- Supports two data centers that are physically located a great distance apart.
- Each data center has its own node group.
- Master-master asynchronous database replication increases tolerance of network latency.
- Load balancing occurs over nodes within a single node group to overcome replication delay.

Configuring Replication

- [Configuring an Active - Standby Replication](#)..... 15
- [Configuring an Active-Active Replication](#)..... 19

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


- [Configuring an Active - Standby Replication](#) on page 15
- [Configuring an Active-Active Replication](#) on page 19

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 8 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]
```

Configuring Replication

Configuring an Active - Standby Replication

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.
```

Configuring Replication

Configuring an Active - Standby Replication

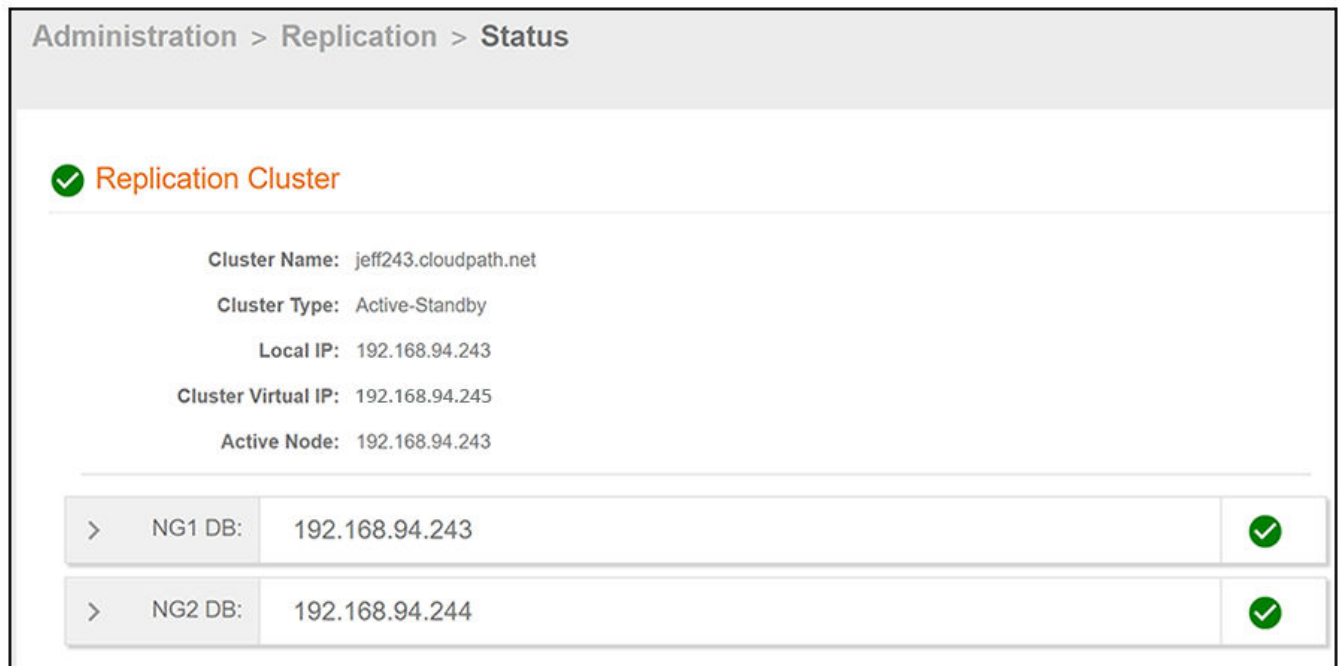
```
Will enable cluster monitoring service on 192.168.94.244 ...
Created symlink from /etc/systemd/system/multi-user.target.wants/cloudpath-repld.service to /etc/
systemd/system/cloudpath-repld.service.
Replication cluster setup complete.
#
```

Verifying That the Cluster was Properly Created With Two DB Nodes:

To verify that your two-DB-node replication has been created correctly, follow these steps:

1. Log in to the Cloudpath UI on any either of the two DB nodes in the cluster.
2. Go to **Administration > Replication**.
3. Make sure that the status indicates that both nodes are present in the cluster and have green check marks next to their names. as shown in the screen below that is based on the example in the topology presented in this section.

FIGURE 9 Cluster Status Successful After Configuration of Both Database Nodes Completes



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.

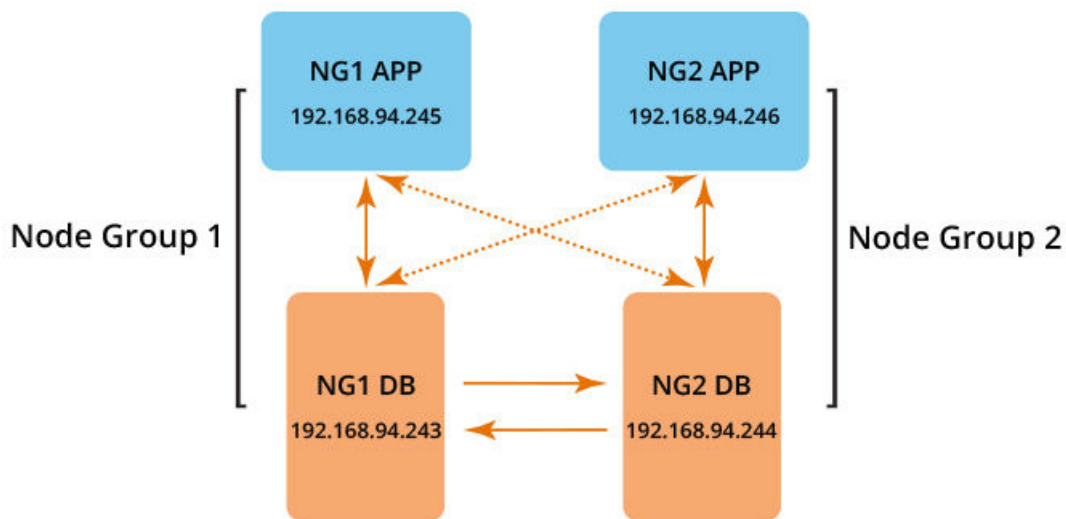
2. Rename the host name of the primary node from jeff243.cloudpath.net to a different name so that it does not conflict with cluster name.

Configuring an Active-Active Replication

An active-active replication requires a load balancer and two database nodes, with the option of adding one or more application nodes to each database node.

The following figure depicts an active-active topology with example IP addresses. This topology will be used to illustrate the steps you take to configure your cluster.

FIGURE 10 Cluster With Two Database Nodes and Two Application Nodes



The basic steps to configure this topology are:

1. From the command line interface of the primary DB node, configure both the primary and secondary database nodes.
2. From the command line interface of the primary DB node, configure any application nodes, one at a time, that are part of the same node group as the primary DB.
3. From the command line interface of the secondary DB node, configure any application nodes, one at a time, that are part of the same node group as the secondary DB.

The more detailed steps below walk you through the configuration of the database nodes. It is assumed that you have already deployed two Cloudpath ES systems with the same .OVA file.

Procedure to Configure the Two Database Nodes:

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.

Configuring Replication

Configuring an Active-Active Replication

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: jeff247.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.  
-----
```

Configuring Replication

Configuring an Active-Active Replication

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:"

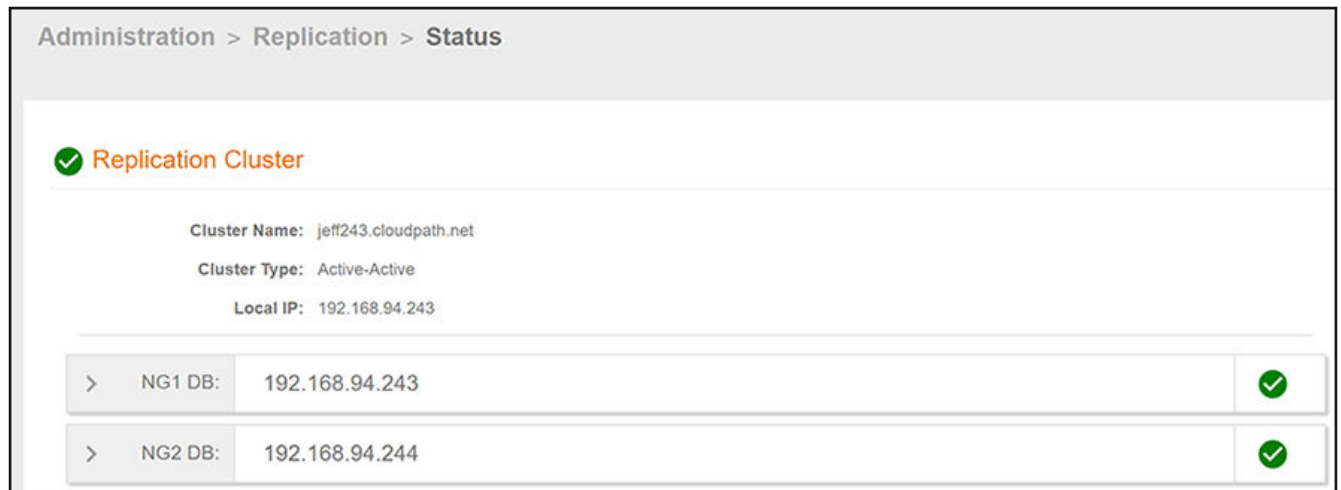
```
Will establish ssh 192.168.94.244:8022 ...
-----
  * You will be prompted for the password of the remote VM 192.168.94.244.
-----
yes: standard output: Broken pipe
yes: write error
/usr/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "/home/repl_user/.ssh/id_rsa.pub"
cpn_service@192.168.94.244's 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.
#
```

Verifying That the Cluster was Properly Created With Two DB Nodes:

To verify that your two-DB-node replication has been created correctly, follow these steps:

1. Log in to the Cloudpath UI on any either of the two DB nodes in the cluster.
2. Go to **Administration > Replication**.
3. Make sure that the status indicates that both nodes are present in the cluster and have green check marks next to their names. as shown in the screen below that is based on the example in the topology presented in this section.

FIGURE 11 Cluster Status Successful After Configuration of Both Database Nodes Completes



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 jeff247.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.

NOTE

If you are adding application nodes to your cluster, proceed to [Adding Application Nodes to a Cluster](#) on page 23.

Adding Application Nodes to a Cluster


To add an application node to your cluster, you run the **replication add-appnode** command from the applicable database node.

NOTE

Be sure that the OVA build number for all application nodes is the same as the OVA build number on the database nodes.

The following figure depicts an active-active topology with example IP addresses. This topology will be used to illustrate the steps you take to add application nodes to the cluster. For the purposes of this example, complete steps and output will be shown for adding the NG1 APP to the cluster from its corresponding NG1 DB node.

FIGURE 12 Cluster With Two Database Nodes and Two Application Nodes



The basic steps to configure a topology of node groups are:

1. From the command line interface of the primary DB node, configure all application nodes, one at a time, that are part of the same node group as the primary DB.
2. From the command line interface of the secondary DB node, configure all application nodes, one at a time, that are part of the same node group as the secondary DB.

Follow the steps below to add an application node to a cluster. The steps shown use the example of adding the application node 192.168.94.245 to the database node 192.168.94.243.

Procedure to Configure Application Server Nodes:

1. On the database node from which you are adding the application node belongs to, 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 add-appnode**
3. You are presented with the following output, to which you can answer **yes**:

```
-----  
* You will be prompted for the password of local system.  
-----  
-----  
WARNINGS:  
1. This routine must be performed on one of Database nodes in the cluster  
   which replication between DB servers has been established.  
2. After the APP node is added to the cluster, all database activities  
   will be redirected to the replicated database servers in the cluster. The  
   local database in the APP server will not be accessible.  
3. This routine is only applicable to the ACTIVE-ACTIVE cluster setup.  
-----
```



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

- Next, you are presented with the following output and prompt:

```
You are on Server:
Server ID      DB Server IP
1             192.168.94.243
```

```
Do you want to add an APP node through this database server node? [y/n]
```

Confirm that the information is correct, then answer **yes** to continue.

- You are then prompted to enter the IP address of the application node you wish to add. Enter the IP address, then press Enter:

```
APP Node IP: 192.168.94.245
```

- You are then prompted to enter the SSH port of the application node. Enter the port number and press Enter.

```
APP Node SSH Port: 8022
```

- You are then provided with information about the database access priority for the application node:

```
DB access-priority order from this APP Node:
Server ID      DB Server IP
1             192.168.94.243
2             192.168.94.244
```

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

Confirm that the information is correct, then answer **yes**.

- You are then prompted for the `cpn_service` password for the application node:

```
Establish ssh 192.168.94.245 8022 ...
```

```
-----
* You will be prompted for the password of the remote VM 192.168.94.245.
-----
```

```
/usr/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "/home/repl_user/.ssh/id_rsa.pub"
cpn_service@192.168.94.245's password:
```

Enter the password, then press Enter.

- The following display shows the completion of the command once you enter the password of the node to add. Check that the end of the display indicates that the node has been added to the cluster.

```
id_rsa.pub
100% 406 40.0KB/s 00:00
Checking remote host 192.168.94.245 ...
Will stop application services on 192.168.94.245 ...
Will update the peer APP server 192.168.94.245 ...
Will add shiro_user@192.168.94.245 on local DB server ...
Will add shiro_user@192.168.94.245 on remote DB server 192.168.94.244 ...
Will setup proxysql on APP server 192.168.94.245 ...
Done setup proxysql on APP server 192.168.94.245.
Will configure application name on App server 192.168.94.245 ...
Done Application Server Name configuration on App server 192.168.94.245.
Will update cluster for peer DB server ...
Will notify cluster changes ....
Done cluster change notification on local system
Will notify cluster changes on the remote DB server 192.168.94.244 ...
Done cluster change notification on the remote DB server 192.168.94.244.
Will disable and stop mysql service on 192.168.94.245 ...
Done disabling and stopping mysql service on 192.168.94.245 ...
Will restart remote application services on 192.168.94.245 ...
```

Configuring Replication

Configuring an Active-Active Replication

192.168.94.245 has been successfully added into the cluster.

Next Steps:

In the example topology shown in [Figure 12](#), there are no more application nodes to add to Node Group 1. If there were more APP nodes to add to Group 1, you would follow the preceding procedure for each node.

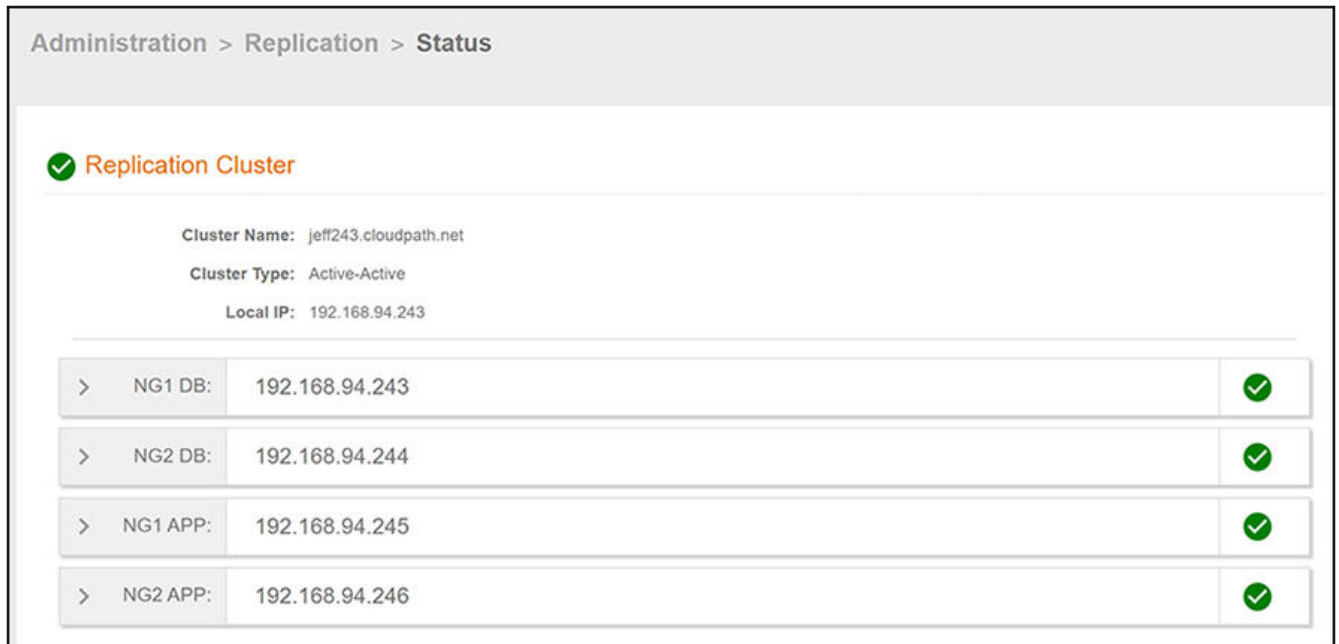
Once you have added all application nodes to Node Group 1, you need to log into the command line interface on the DB node for Node Group 2, then follow the same basic steps as shown above to add the Node Group 2 application nodes. In the example presented in this topic, you would be logging into NG2 DB (IP address of 192.168.94.244) to add the NG APP (IP address of 192.168.94.246).

Verifying That the Cluster was Properly Created:

To verify that your cluster has been created correctly, follow these steps:

1. Log in to the Cloudpath UI on any of the nodes (DB or APP) in the cluster.
2. Go to **Administration > Replication**.
3. Make sure that the status indicates that all nodes are present in the cluster and have green check marks next to their names. as shown in the screen below that is based on the example in the topology presented in this section.

FIGURE 13 Cluster Status Successful After Configuration Completes



Removing an Application Node From a Cluster

If for any reason you wish to remove an application node from your cluster, you issue the **replication remove-appnode** command from the applicable database node.

Follow the steps below to remove an application node from a cluster. The steps shown use the example of removing the application node 192.168.94.245 from the database node 192.168.94.243 (refer to [Figure 12](#) on page 24).

1. On the database node that the application node belongs to, 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 remove-appnode**
3. You are presented with the following output, to which you can answer **yes**:

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

WARNINGS:
 1. This routine will remove an APP node from the cluster and must be
    performed on one of Database server nodes in the cluster.
 2. Before starting this routine, make sure the affected APP node
    is a 'peer' APP server node to this Database server.
-----

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

4. After answering **yes**, you are presented with the following output and prompt:

```
-----
WARNING:
You're on Cluster Database Node: 192.168.94.243

Peer APP Nodes:
192.168.94.245
-----

APP node to remove (IP):
```

At the prompt, enter the IP address of the application node to remove, then hit Enter. In this example, the IP address is 192.168.94.245.

5. The following display shows the completion of the command once you enter the IP address of the node to remove. Check that the end of the display indicates that the node has been removed from the cluster.

```
APP node to remove (IP): 192.168.94.245

Checking remote host 192.168.94.245 ...
Will stop application services on App node 192.168.94.245 ...
Will stop and disable proxysql on App node 192.168.94.245 ...
Will reset web server cluster name on APP node 192.168.94.245 ...
Will remove database service permissions ...
Will enable and start mysql service on 192.168.94.245 ...
Done enabling and restarting mysql service on 192.168.94.245 ...
Will update cluster configuration ...
Will update cluster for peer DB server ...
Will notify cluster changes ....
Done cluster change notification on local system
Will notify cluster changes on the remote DB server 192.168.94.244 ...
```

Configuring Replication

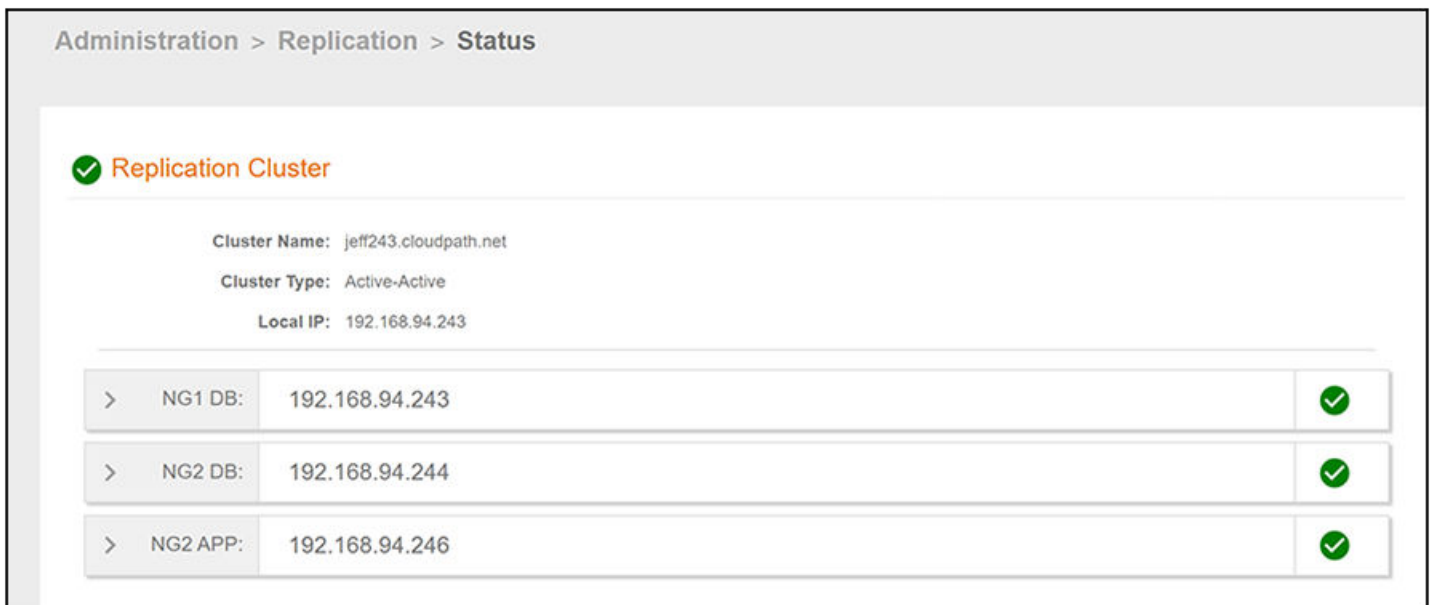
Configuring an Active-Active Replication

```
Done cluster change notification on the remote DB server 192.168.94.244.  
Will remove REPL_WHITELIST on 192.168.94.245 ...  
Will remove cluster configurations from 192.168.94.245 ...  
Will restart remote application services on 192.168.94.245 ...  
Will cleanup cluster abandoned files on 192.168.94.245 ...  
APP node 192.168.94.245 has been removed from the Cluster.  
#
```

Verification of Node Removal Via the User Interface:

You can log in to any of the nodes that are still in the cluster to confirm that the node you removed no longer appears in the list of nodes, as shown in the following illustration:

FIGURE 14 Cluster Status After Removal of an Application Node



If you log in to the node that you just removed from the cluster, you are presented with the same initialization screen you received before you ever added the node to the cluster. You can then set this node up however you wish, including adding it back to its original node group, assigning it to a new cluster, adding it as the primary server of a new cluster, or using it as a standard Cloudpath server that is not part of a cluster.


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 16](#)).

FIGURE 15 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 Group 2 APP expanded:

FIGURE 16 Detailed Information for a Node in the Cluster

NG2 APP: 192.168.94.246
✔

✔ Application Replication

Application Replication Peers:

| Connected | Public IP | UUID | Last Heartbeat | Topic Publ. | Topic Recv. | Maps Heap Bytes |
|-----------|----------------|--------------------------------------|----------------|-------------|-------------|-----------------|
| ✔ | 192.168.94.243 | bf02b113-5568-444a-b775-205856077b55 | 38 seconds ago | 0 | 20793 | 6660 |
| ✔ | 192.168.94.244 | e5bf6cc6-a7ad-484b-9f40-c44b337f4894 | 38 seconds ago | 0 | 20793 | 6660 |
| ✔ | 192.168.94.245 | 4fd928d4-530a-4058-8974-46dc1a87fca5 | 41 seconds ago | 0 | 20793 | 6660 |

Values Updated: N/A

Primary App: No

RADIUS Service Running: Yes

Web Server Service Running: Yes

UUID: b4123702-da8d-4e31-9f02-9f9faaaa7b37

Actions: Warning/Error Log Recent Log

You can click the two log buttons at the bottom of the screen for additional information.

For additional information about a cluster node, refer to the following in https:

`<cloudpathhostname>/admin/systemHealth/viewHistoricMetricLogs` where `<cloudpathhostname>` is the fully qualified domain name of a node in the cluster.:

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 any 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 15:

```
# 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 any 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 15:

```
# 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
```

Running Replication Diagnostic Tests

Showing the Replication Status

```
      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 or adding APP nodes, the HTTPS server names on all 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 any cluster node (DB or APP) to view the cluster service configuration. The following example output is based on running the command from the secondary database node on the active-standby setup described in [Configuring an Active - Standby Replication](#) on page 15:

```
# 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 change as the system configuration is modified.

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

The following two illustrations are based on the firewall settings of the primary node in the four-node cluster whose configuration was shown in [Figure 10](#) on page 19.

FIGURE 17 Firewall Requirements - Outbound

| Traffic: Outbound from this System | | | | |
|------------------------------------|-----------------------|----------------------------|------------------|---|
| Purpose | System Address | External Address | Protocol | Reason |
| System | jeff243.cloudpath.net | bvt.cloudpath.net:443 | TCP / HTTP(s) | System interacting with cloud services (licensing, wizards, built-in email, etc). |
| System | jeff243.cloudpath.net | support.cloudpath.net:8022 | TCP | (Optional) Support tunnel for remote assistance. Only necessary when support tunnel is enabled. |
| System | jeff243.cloudpath.net | dist.cloudpath.net:443 | TCP / HTTP(s) | Retrieve system updates. |
| System | jeff243.cloudpath.net | dist2.cloudpath.net:443 | TCP / HTTP(s) | Retrieve system updates. |
| OCSP Stapling | jeff243.cloudpath.net | ocsp.godaddy.com:80 | TCP / HTTP(s) | Send OCSP request to web server certificate provider for OCSP stapling. |
| Authentication Server | jeff243.cloudpath.net | 192.168.4.2:389 | TCP | Authenticate to Active Directory server 'Corporate AD' at 'ldap://192.168.4.2'. |
| System | jeff243.cloudpath.net | | TCP | LinkedIn authentication enabled but firewall specifics not available. |
| NTP | jeff243.cloudpath.net | pool.ntp.org:123 | UDP / NTP | Perform NTP synchronization. |
| Replication | jeff243.cloudpath.net | 192.168.94.244:3306 | TCP | Replication Database. |
| Replication | jeff243.cloudpath.net | 192.168.94.244:5701 | TCP | Replication Application. |
| Replication | jeff243.cloudpath.net | 192.168.94.244:8022 | TCP | Replication SSH. |
| Replication | jeff243.cloudpath.net | 192.168.94.245:3306 | TCP | Replication Database. |
| Replication | jeff243.cloudpath.net | 192.168.94.245:5701 | TCP | Replication Application. |
| Replication | jeff243.cloudpath.net | 192.168.94.245:8022 | TCP | Replication SSH. |
| Replication | jeff243.cloudpath.net | 192.168.94.246:3306 | TCP | Replication Database. |
| Replication | jeff243.cloudpath.net | 192.168.94.246:5701 | TCP | Replication Application. |
| Replication | jeff243.cloudpath.net | 192.168.94.246:8022 | TCP | Replication SSH. |
| RADIUS COA | jeff243.cloudpath.net | *:3799 | UDP / RADIUS COA | Send COA to wired/wireless infrastructure (default RADIUS client). |
| RADIUS COA | jeff243.cloudpath.net | 10.0.0.1:3799 | UDP / RADIUS COA | Send COA to wired/wireless infrastructure. |

FIGURE 18 Firewall Requirements - Inbound

| Traffic: Inbound to this System | | | | |
|---------------------------------|----------------------------|------------------|---------------|---|
| Purpose | System Address | External Address | Protocol | Reason |
| Web Interface | jeff243.cloudpath.net:443 | | TCP / HTTP(s) | Administrator, API, and end-user access to the web interface. |
| Onboard CA | jeff243.cloudpath.net:80 | | TCP / HTTP(s) | OCSP requests coming from external systems. |
| SSH | jeff243.cloudpath.net:8022 | | TCP | SSH access to the system. |
| Onboard RADIUS | jeff243.cloudpath.net:1812 | | UDP | Receive RADIUS authentication requests from external systems. |
| Onboard RADIUS | jeff243.cloudpath.net:1813 | | UDP | Receive RADIUS accounting requests from external systems. |
| Replication | jeff243.cloudpath.net:3306 | 192.168.94.244 | TCP | Replication Database. |
| Replication | jeff243.cloudpath.net:5701 | 192.168.94.244 | TCP | Replication Application. |
| Replication | jeff243.cloudpath.net:8022 | 192.168.94.244 | TCP | Replication SSH. |
| Replication | jeff243.cloudpath.net:3306 | 192.168.94.245 | TCP | Replication Database. |
| Replication | jeff243.cloudpath.net:5701 | 192.168.94.245 | TCP | Replication Application. |
| Replication | jeff243.cloudpath.net:8022 | 192.168.94.245 | TCP | Replication SSH. |
| Replication | jeff243.cloudpath.net:3306 | 192.168.94.246 | TCP | Replication Database. |
| Replication | jeff243.cloudpath.net:5701 | 192.168.94.246 | TCP | Replication Application. |
| Replication | jeff243.cloudpath.net:8022 | 192.168.94.246 | TCP | Replication SSH. |

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 Database Replication

You can attempt to repair the database without tearing down the cluster 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 cluster nodes is out of sync. 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 15

```
# 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-
```

Repairing Database Replication

```
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 database nodes back to single Cloudpath Enrollment System mode.

NOTE

If you are going to tear down an active-active cluster that contains application nodes, you must first remove all APP nodes from the cluster by following the procedure in [Removing an Application Node From a Cluster](#) on page 27.

You can run the **replication teardown** command from either database node.

The following command (and corresponding output) is an example from the cluster that was shown in [Configuring an Active - Standby Replication](#) on page 15

```
# 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 databases 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 ...
```

Tearing Down the Cluster

```
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 replication between two cluster database nodes 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. |
| replication add-appnode | Adds an APP node to an existing active-active cluster. This command must be run from the DB node of the node group to which you are adding the APP node. |
| replication remove-appnode | Removes an APP node from an existing active-active cluster. This command must be run from the DB node of the node group to which the the APP node belongs. |
| replication show-config | Displays the configuration of the cluster. Example output: <pre>----- ----- Primary - 192.168.94.243 ssh port# 8022 Secondary - 192.168.94.244 ssh port# 8022 Type - MULTISITE-ACTIVE-ACTIVE-NO-NAT Cluster Name - jeff247.cloudpath.net ----- -----</pre> |
| 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. If you are going to tear down an active-active cluster that contains APP nodes, you must first remove all the APP nodes using the replication remove-appnode command. |

