



Configuring Replication Servers

This chapter includes the following sections:

- [About Replication Servers, on page 1](#)
- [Rebalancing Replication Servers, on page 1](#)
- [Setting the Operational State of a Replication Server, on page 2](#)

About Replication Servers

Replication servers are dedicated VXLAN tunnel endpoints (VTEPs) that are responsible for replicating Layer 2 broadcast, unknown unicast, and multicast (BUM) traffic. The VTEPs must send the BUM traffic to only one of the replication servers, which will then replicate it to all VTEPs in that virtual network identifier (VNID).

In Cisco NX-OS, replication servers are picked on a per-VNID basis. Therefore, all BUM traffic from a VNID is sent to this replication server.

Rebalancing Replication Servers

Replication servers can arrive at the switch at any time. To help manage the load, you can rebalance the VNIDs across the replication servers.

The first time that the total number of replication servers in the system goes from 0 to 1, the switch waits for 1 minute before programming this information into the hardware. Doing so prevents continuous rebalancing of the VNIDs as the replication servers come down one by one.

Step 1 `nve interface nve 1 {remap-replication-servers | replication-server}`

```
switch# nve interface nve 1 remap-replication-servers
```

The following options are available:

- **remap-replication-servers** —Immediately rebalances the VNIDs across all available replication servers. If another replication server is sent down after you enter this command, you must re-enter this command to rebalance the replication servers again. This behavior is the default behavior of the switch.
- **replication-server** —Marks the replication server as Up or Down.

Step 2 (Optional) **copy running-config startup-config**

```
switch# copy running-config startup-config
```

Copies the running configuration to the startup configuration.

Setting the Operational State of a Replication Server

If a replication server's operational state is not sent down from the external controller, the switch continues to replicate the BUM traffic to that replication server, which can result in some traffic being lost. To prevent this behavior, you can mark a replication server as being down or up.

Step 1 **show nve replication-servers**

```
switch# show nve replication-servers
Interface Replication Servers  State Ready
-----
nve1      10.0.134.1                    Up    Yes
          10.0.133.1                    Up    Yes
          10.0.140.1                    Up    Yes
```

Displays the list of replication servers in the system.

Step 2 **nve interface nve 1 replication-server *ip-address* [up | down]**

```
switch# nve interface nve 1 replication-server 10.10.10.1 up
```

Sets the operational state of the replication server:

- **up**—Marks the replication server as being up. The resulting behavior depends on the replication server rebalancing configuration (either automatic rebalancing or rebalancing once upon CLI execution).
- **down**—Marks the replication server as being down. The VNIDs are immediately reshaped to the replication servers that are up.

Step 3 (Optional) **copy running-config startup-config**

```
switch# copy running-config startup-config
```

Copies the running configuration to the startup configuration.
