

Cluster Configuration With 'join cluster' Command

HOW TO GUIDE



October 2012
Revision 1.2

© 2015 ZIH Corp. All rights reserved. Zebra and the Stylized Zebra Head are trademarks of ZIH Corp, registered in many jurisdictions worldwide. All other trademarks are property of their respective owners.

Table of Contents

Table of Contents.....	3
1. Overview.....	4
1.1 Document Scope	4
2. Cluster Configuration	5
2.1 Primary / First Controller	5
2.1.1 About the <i>cluster member ip</i>	6
2.2 Standby / Second Controller	6
2.2.1 Active / Active or Active / Standby.....	6
2.2.2 Secondary controller configuration.....	6
3. Troubleshooting	7
3.1 Debug.....	8
.....	10

1. Overview

Cluster configuration prior to WiNG 5.1 involved several configuration steps and was often prone to user error. In 5.1, the 'join-cluster' command was introduced and made cluster configuration far simpler and required minimal configuration. In WiNG 5.4, the 'join-cluster' command has been given the 'mode' option to effectively reduce the configuration needed on a secondary cluster member to only an IP address and possibly a default-gateway.

In WiNG 5.x two controllers can be part of a cluster for redundancy for services such as:

- Access Point Adoption
- License Sharing
- DHCP Server redundancy
- Dynamic Access Point Load-Balancing

Clusters are managed from a single user interface and the configuration is automatically synched between cluster members. A cluster can be configured as "Active / Active" in which both controllers are actively adopting access points, or "Active / Standby", where one controller is a 'standby' device, ready to take over if the primary 'active' device should fail.

Access Point licensing is shared between controllers, so in an "Active / Standby" scenario, it is common to order one controller with licenses and the other as a "zero port" controller, as it will share the licensing of the active controller. If both controllers are ordered with licenses, understand that any one controller can only adopt as many AP's as it is capable of handling in hardware. In other words, though you could order two RFS7000's with max licensing (1024 licenses each), a single controller is not capable of adopting 2048 access points should one controller fail. The maximum number of access points supported on the various platforms is listed below:

Device Model	Access Points supported
NX9xxx	10,240
RFS7000	1024
RFS6000	256
RFS4000	36

Feature licenses, such as Advanced Security or Advanced WIPS are not shared among devices in a cluster. Those license types are applied per controller, so they must be purchased and applied for both controllers in a cluster.

1.1 Document Scope

The purpose of this document is to cover how to configure a layer-3 (IP) cluster between two WiNG 5.4 controllers, utilizing the 'join-cluster' command. This command was introduced in WiNG 5.1 and is not present in any earlier versions.

2. Cluster Configuration

Prior to WiNG 5.1, clustering was more difficult and required some configuration on both controllers and then adding each device instance in the configuration on the other device. This was confusing and if not done correctly, the cluster would never become active.

In WiNG 5.4, the “**join-cluster**” command has the added ‘mode’ option, allowing one to specify whether the secondary controller is active or standby during execution of the command.

2.1 Primary / First Controller

On your first or primary controller, complete any configuration to be used for your environment. The only parameters absolutely needed before going on to cluster configuration are:

- Device IP Address
- Default Gateway
- Ensure connectivity on the network

Once this is complete, you can proceed with cluster configuration. Of course, one may also opt to completely configure the primary controller prior to enabling clustering; all configurations will be synchronized after cluster configuration is complete. When ready to configure clustering, perform the following steps:

Configure primary switch:

```
rfs4000-1(config-device-11-22-33-44-55-66)#cluster name <cluster_name>
rfs4000-1(config-device-11-22-33-44-55-66)#cluster mode active (default)
rfs4000-1(config-device-11-22-33-44-55-66)#cluster member ip <ip_of_primary> level
<1|2>
rfs4000-1(config-device-11-22-33-44-55-66)#cluster master-priority 250
rfs4000-1(config-device-11-22-33-44-55-66)#cluster force-configured-state
rfs4000-1(config-device-11-22-33-44-55-66)#cluster force-configured-state-delay 5
```

Clusters can be formed with level-1 or level-2 MiNT links, and this is dependent on the implementation. Recall the following:

- Level-1 Links: All level-1 devices will establish links with each other and bridge-mode of “tunnel” is possible to allow tunnelling of VLAN’s.
- Level-2 Links: Level-1 devices within an RF-Domain will establish links to each other while the RF-Domain controller establishes a level-2 link to the controllers. Also bridge-mode “tunnel” or tunnelling of VLAN’s is not supported over level-2 links.

The use of “cluster force-configured-state” and “cluster force-configured-state-delay” will ensure that in the event of a primary failure, it will assume its role as the primary once it is back online, following the appropriate delay period.

2.1.1 About the *cluster member ip*

Notice that on the primary, you are pointing to itself as the ‘cluster member ip’. This will cause the member IP to be ‘synched’ across to the secondary, and vice-versa, giving each device two entries for member IP’s: their own IP as well as the true other member.

2.2 Standby / Second Controller

Once you are satisfied with the configuration of the primary controller, you must determine the role of the second controller: if it will be active or standby. Then some configuration is needed on the second controller to at least get it onto the network and able to communicate with other devices.

2.2.1 Active / Active or Active / Standby

The configuration of your cluster will require a mode to be established for the desired operation within your environment. I’ve often had questions as to which mode is best: active / active or active / standby, and the answer is really subjective based on the customer. In a large enterprise / NOC deployments of 3,000 or more access-points, the cluster mode should be set to ‘Active / Standby’ in order to maintain statistics and other information performance in the UI.

In terms of just making a simple decision, think of the cluster in this way: do I want load-balancing or do I want fault-tolerance? Of course as long as controller capacity in an active / active setup is not exceeded (i.e. total number of access-points does not exceed the capacity limit of any single controller), there is also fault-tolerance. But for the sake of discussion, consider the two separately; one could also consider ease of troubleshooting if so desired.

- Active / Active – good for load-balancing scenarios
- Active / Standby – good for fault-tolerant and ease of troubleshooting scenarios.

If nothing is specified, the default configuration establishes an “active” mode.

2.2.2 Secondary controller configuration

On the secondary controller, all that is really needed is basic network connectivity information so that the device can communicate with the other controller, as well as the mode – active or standby. The following list details the secondary switch requirements prior to joining cluster:

- IP address / mask
- Default Gateway

Configure secondary switch:

```
rfs4000-2(config-device-11-22-33-44-55-66)#interface vlan<n>
rfs4000-2(config-device-11-22-33-44-55-66-if-vlan<n>)#ip address x.x.x.x/m
rfs4000-2(config-device-11-22-33-44-55-66-if-vlan<n>)#exit
rfs4000-2(config-device-11-22-33-44-55-66)#ip default-gateway x.x.x.x
rfs4000-2(config-device-11-22-33-44-55-66)#commit write
```

```

rfs4000-2(config-device-11-22-33-44-55-66)#end
rfs4000-2#join-cluster <primary_ip> user admin password <pswd> mode <active/standby>
specifying mode is new and provides an even easier method of establishing your cluster
correctly.
... connecting to 172.16.24.254
... applying cluster configuration
... committing the changes
... saving the changes
rfs4000-2#show cluster status
Cluster Runtime Information
Protocol version          : 1
Cluster state             : standby
AP license                : 12
AAP license               : 0
AP count                  : 0
AAP count                 : 2
Max AP adoption capacity  : 72
Number of connected member(s): 1

```

If you see the above messages once establishing the cluster, then you have successfully completed the cluster configuration. Performing a “show running-config” at this point will reveal that the entire configuration has been synchronized to the secondary controller.

3. Troubleshooting

Most problems with the cluster forming are caused by configuration errors. Verify the following items:

- VLAN ID's match: in a scenario where perhaps the VLAN interfaces or native VLAN's don't match, the cluster will not establish
- Member IP: make sure when configuring the primary, that you enter its own IP address as the 'cluster member ip', on the appropriate virtual interface. There is no need to enter the IP of the secondary controller, as this will happen automatically.
- Version mismatch: if the controllers are on different firmware versions, the configuration may seem to take, but you will receive an error message regarding cluster master election and the cluster will not establish. The controllers must be on the exact same firmware version

The following commands will help determine the state of the configuration and cluster:

Configure secondary switch:

```
rfs4000-1#show cluster config
```

```

Cluster Configuration Information Name
                                : VLAB
Configured Mode                 : Active
Master Priority                 : 250
Force configured state         : Enabled
Force configured state delay   : 5 minutes
Handle STP                     : Disabled
rfs4000-1#
rfs4000-1#show cluster members
-----
-----
  HOSTNAME      MEMBER-ID      MAC              MASTER  OPERATIONAL-STATE  LAST-
  SEEN
-----
-----
  rfs4000-1    68.22.D2.6E    00-23-68-22-D2-6E  True    active              self
  rfs4000-2                                00-23-68-22-A3-AC  False   active
-----
-----
rfs4000-1#
rfs4000-1#show cluster status

Cluster Runtime Information
Protocol version              : 1
Cluster operational state     : active
AP license                   : 12
AAP license                   : 0
AP count                     : 0
AAP count                    : 2
Max AP adoption capacity     : 36
Number of connected member(s): 1

```

3.1 Debug

Debugging during the join process is recommended as it is the quickest way to determine what may be wrong with the cluster configuration. Run a debug on the primary and log to the monitor when executing the 'join-cluster' command on the secondary controller. This will provide very useful information as to why a cluster failed or if it proceeded successfully. The output below is from a successful join:

Configure secondary switch:

```
rfs4000-1#debug cfgd cluster
rfs4000-1#logging monitor debug
rfs4000-1#Oct 11 13:42:32 2012: USER: cfgd: CS server (cluster join)
connection from 172.16.24.253
Oct 11 13:42:32 2012: USER: cfgd: CS_Server - cluster join from host
172.16.24.253
Oct 11 13:42:33 2012: USER: cfgd: cluster join, cluster member ip of [00-23-
68-22-D2-6E] is NOT configured in profile default-rfs4000, create override
Oct 11 13:42:33 2012: USER: cfgd: cluster join, cluster member ip of [00-23-
68-22-A3-AC] is NOT configured in profile default-rfs4000, create override
Oct 11 13:42:34 2012: USER: cfgd: master-updater woke up - 1 changes in the
list
Oct 11 13:42:34 2012: USER: cfgd: tx_update CS_ACT_CFG_UPDATE revision 15
originator Master
Oct 11 13:42:34 2012: USER: cfgd: check if cluster state transition is
required
Oct 11 13:42:35 2012: USER: cfgd: master-updater woke up - 1 changes in the
list
Oct 11 13:42:35 2012: USER: cfgd: tx_update CS_ACT_WR_MEM revision 16
originator Master
Oct 11 13:42:36 2012: USER: cfgd: sending OK with config
Oct 11 13:42:38 2012: USER: cfgd: handle_cluster_member_update
Oct 11 13:42:38 2012: USER: cfgd: handle_cluster_member_update
members:['68.22.A3.AC', '68.22.D2.6E']
Oct 11 13:42:38 2012: USER: cfgd: handle_cluster_member_update members:
updating potential members
Oct 11 13:42:38 2012: USER: cfgd: Updating new members ['68.22.A3.AC']
Oct 11 13:42:38 2012: USER: cfgd: created new local pkt:MASTER_UPDATE ver 1
prio 250 master period 12045 rank 120 dc 0 id 68.22.D2.6E
Oct 11 13:42:38 2012: USER: cfgd: dequeue - need to run election
Oct 11 13:42:38 2012: USER: cfgd: starting election forced = False
Oct 11 13:42:38 2012: USER: cfgd: created new local pkt:CANDIDATE ver 1 prio
250 master period 12045 rank 120 dc 1 id 68.22.D2.6E
Oct 11 13:42:38 2012: USER: cfgd: created new local pkt:START_ELECTION ver 1
prio 250 master period 12045 rank 120 dc 1 id 68.22.D2.6E
Oct 11 13:42:38 2012: USER: cfgd: broadcasting election start
Oct 11 13:42:38 2012: USER: cfgd: announcing self as candidate
Oct 11 13:42:38 2012: USER: cfgd: cancelling ask_msg_update_timer
Oct 11 13:42:40 2012: USER: cfgd: Peer 68.22.A3.AC connected
Oct 11 13:42:40 2012: USER: cfgd: received packet MAC_REQ ver 1 prio 128
```

```
master period 0 rank 120 dc 1 id 68.22.A3.AC
Oct 11 13:42:40 2012: USER: cfgd: created new local pkt:MAC_RESPONSE ver 1
prio 250 master period 12047 rank 120 dc 1 id 68.22.D2.6E
Oct 11 13:42:40 2012: USER: cfgd: received packet CANDIDATE ver 1 prio 128
master period 0 rank 120 dc 1 id 68.22.A3.AC
Oct 11 13:42:40 2012: USER: cfgd: Local election score improved, 2 known
configs
Oct 11 13:42:40 2012: USER: cfgd: created new local pkt:CANDIDATE ver 1 prio
250 master period 12047 rank 120 dc 2 id 68.22.D2.6E
Oct 11 13:42:40 2012: USER: cfgd: best set to 68.22.D2.6E (CANDIDATE ver 1
prio 250 master period 12047 rank 120 dc 2 id 68.22.D2.6E)
Oct 11 13:42:40 2012: USER: cfgd: retaining old candidate 68.22.D2.6E as
best, not 68.22.A3.AC
Oct 11 13:42:40 2012: USER: cfgd: received packet START_ELECTION ver 1 prio
128 master period 0 rank 120 dc 1 id 68.22.A3.AC
Oct 11 13:42:40 2012: USER: cfgd: resetting election timeout
Oct 11 13:42:40 2012: USER: cfgd: created new local pkt:CANDIDATE ver 1 prio
250 master period 12047 rank 120 dc 2 id 68.22.D2.6E
Oct 11 13:42:40 2012: USER: cfgd: best set to 68.22.D2.6E (CANDIDATE ver 1
prio 250 master period 12047 rank 120 dc 2 id 68.22.D2.6E)
Oct 11 13:42:40 2012: USER: cfgd: received packet CANDIDATE ver 1 prio 128
master period 0 rank 120 dc 1 id 68.22.A3.AC
Oct 11 13:42:40 2012: USER: cfgd: retaining old candidate 68.22.D2.6E as
best, not 68.22.A3.AC
Oct 11 13:42:40 2012: USER: cfgd: received packet CANDIDATE ver 1 prio 128
master period 0 rank 120 dc 2 id 68.22.A3.AC
Oct 11 13:42:40 2012: USER: cfgd: retaining old candidate 68.22.D2.6E as
best, not 68.22.A3.AC

rfs4000-1#no debug allOct 11 13:42:49 2012: USER: cfgd: publish cluster
status=[active] interally
Oct 11 13:42:55 2012: USER: cfgd: election rcv timeout
Oct 11 13:42:55 2012: USER: cfgd: Announcing self as master
Oct 11 13:42:55 2012: USER: cfgd: created new local pkt:ELECTION_WINNER ver 1
prio 250 master period 12062 rank 120 dc 2 id 68.22.D2.6E
rfs4000-1#
```