High Availability Configuration Guide

AirLink Connection Manager

SIERRA WIRELESS®

4118775
Rev 5

Important Notice

Due to the nature of wireless communications, transmission and reception of data can never be guaranteed. Data may be delayed, corrupted (i.e., have errors) or be totally lost. Although significant delays or losses of data are rare when wireless devices such as the Sierra Wireless modem are used in a normal manner with a well-constructed network, the Sierra Wireless modem should not be used in situations where failure to transmit or receive data could result in damage of any kind to the user or any other party, including but not limited to personal injury, death, or loss of property. Sierra Wireless accepts no responsibility for damages of any kind resulting from delays or errors in data transmitted or received using the Sierra Wireless modem, or for failure of the Sierra Wireless modem to transmit or receive such data.

Limitation of Liability

The information in this manual is subject to change without notice and does not represent a commitment on the part of Sierra Wireless. SIERRA WIRELESS AND ITS AFFILIATES SPECIFICALLY DISCLAIM LIABILITY FOR ANY AND ALL DIRECT, INDIRECT, SPECIAL, GENERAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR EXEMPLARY DAMAGES INCLUDING, BUT NOT LIMITED TO, LOSS OF PROFITS OR REVENUE OR ANTICIPATED PROFITS OR REVENUE ARISING OUT OF THE USE OR INABILITY TO USE ANY SIERRA WIRELESS PRODUCT, EVEN IF SIERRA WIRELESS AND/OR ITS AFFILIATES HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES OR THEY ARE FORESEEABLE OR FOR CLAIMS BY ANY THIRD PARTY.

Notwithstanding the foregoing, in no event shall Sierra Wireless and/or its affiliates aggregate liability arising under or in connection with the Sierra Wireless product, regardless of the number of events, occurrences, or claims giving rise to liability, be in excess of the price paid by the purchaser for the Sierra Wireless product.

Patents

This product may contain technology developed by or for Sierra Wireless Inc. This product includes technology licensed from QUALCOMM®. This product is manufactured or sold by Sierra Wireless Inc. or its affiliates under one or more patents licensed from InterDigital Group and MMP Portfolio Licensing.

Copyright

© 2019 Sierra Wireless. All rights reserved.

Trademarks

Sierra Wireless®, AirPrime®, AirLink®, AirVantage® and the Sierra Wireless logo are registered trademarks of Sierra Wireless.

Windows® and Windows Vista® are registered trademarks of Microsoft Corporation.

Macintosh® and Mac OS X® are registered trademarks of Apple Inc., registered in the U.S. and other countries.

QUALCOMM® is a registered trademark of QUALCOMM Incorporated. Used under license.

Other trademarks are the property of their respective owners.

Contact Information

Sales information and technical support, including warranty and returns	Web: sierrawireless.com/company/contact-us/ Global toll-free number: 1-877-687-7795 6:00 am to 5:00 pm PST
Corporate and product information	Web: sierrawireless.com

Revision History

Revision number	Release date	Changes
1	February 2016	• Document created
2	September 2016	• Clarified supported devices throughout document • Rebranding from oCM to ACM
3	December 2017	• Added VRRP content • Reorganized document
4	July 2018	• Applied new document template • Updated Max Throughput table • Updated BGP Protocol example
5	November 2019	• Updated Overview—added DDNS details; updated DNS figure • Updated VRRP Configuration procedures • Added 'Forcing an Artificial Failover' topic • Added Dynamic DNS chapter • Updated Troubleshooting—Added DDNS section

Contents

1: Introduction

This document describes the high availability solutions currently supported by AirLink Connection Manager (ACM). It is intended for use by Sierra Wireless-certified channel partners, as well as field engineers.

Overview

ACM supports three high availability configuration mechanisms to ensure services remain available in case of server failures.

Table 1-1: High Availability Configuration Mechanisms

Method	Supported Clients	Description
Virtual Router Redundancy Protocol (VRRP)	• All AirLink gateways/ routers • NCP Secure Entry Client	• ACM appliance cluster (typically two appliances) accessed via single virtual IP address • Master ACM appliance provides all services • Failover protection—Alternate appliance automatically replaces master if master fails
DNS Load Balancing	• MG90 routers • oMG gateways	• ACM appliance pool (2 or more appliances) operating in tandem • Each VPN peer (AirLink gateways/routers) sets up their VPN connection with one of the ACMs (determined by DNS). This distributes the VPN peers across the appliance pool to prevent overloading of any single appliance. • Failover protection—If one of the ACMs goes down, its VPN peers move to remaining active ACM appliance(s).
Dynamic DNS (DDNS)	• MG90 routers (running MGOS 4.3 or later)	• ACM appliance cluster (up to 10 appliances) operating in tandem • Each VPN peer (MG90 router) sets up their VPN connection with one of the ACMs (selected by the MG90 based on configured connection method). • Each VPN peer is assigned a unique Gateway ID. Each ACM in the cluster is aware of every connected MG90. • Failover protection—If an ACM goes down, VPN peers will fail over to the remaining active ACM appliance(s) if the peers are set (in MGOS) to use the "Connect One" connection method (either Random or Round-robin option). • Load balancing—Basic load balancing is achieved by configuring VPN peers (in MGOS, using the "Connect One" connection method's Round-robin option) to randomly select the ACM for their tunnel from the cluster.

Note: The term ''VPN peer' is used in this document to refer to VPN clients (endpoints).

The following figures illustrate ACM servers in each High Availability configuration:

- Figure 1-1—For a detailed description of the VRRP mechanism and implementation details, refer to Virtual Router Redundancy Protocol (VRRP) on page 9.

- Figure 1-2—For a detailed description of the DNS Load Balancing mechanism and implementation details, refer to DNS Load Balancing on page 14.
- Figure 1-3—For a detailed description of the Dynamic DNS mechanism and implementation details, refer to Dynamic DNS on page 21.

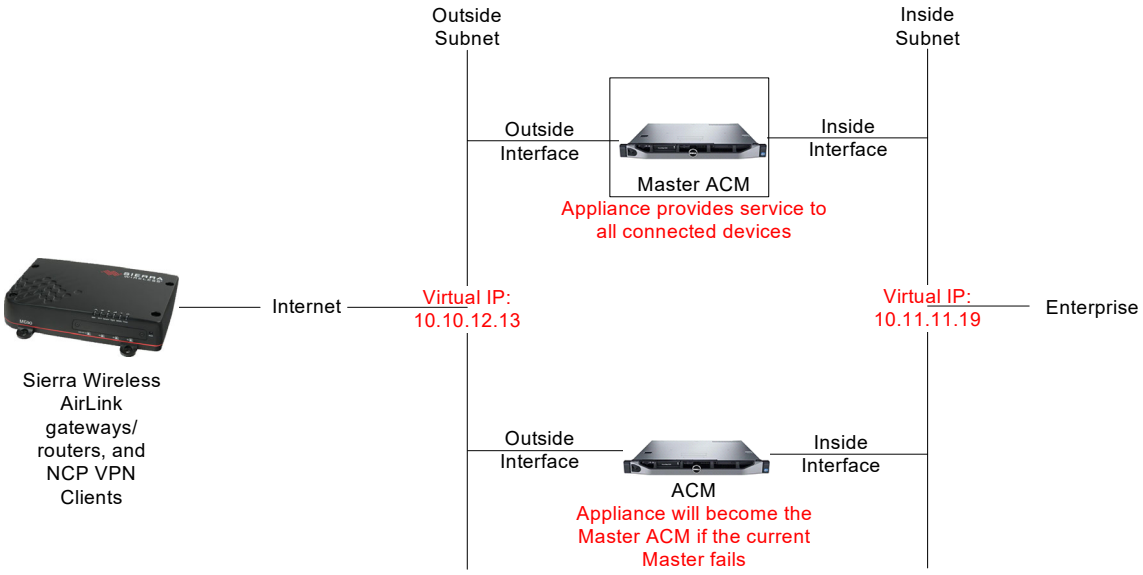

Figure 1-1: Example—ACMs in VRRP Configuration

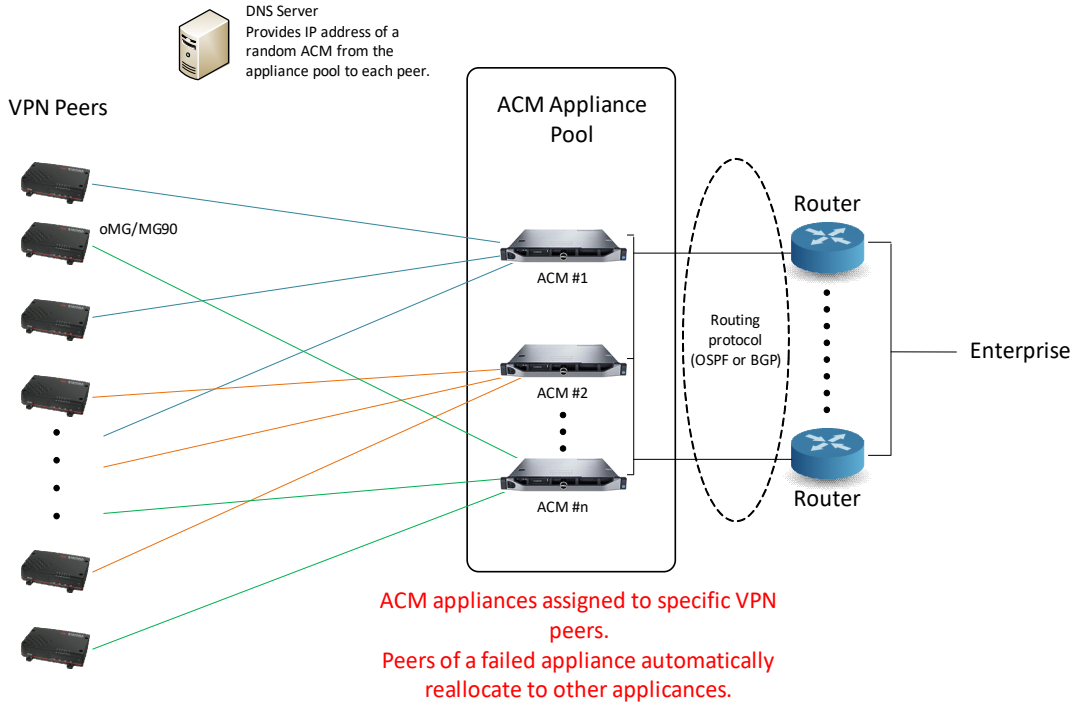

Figure 1-2: Example—ACMs in DNS Load Balancing Configuration

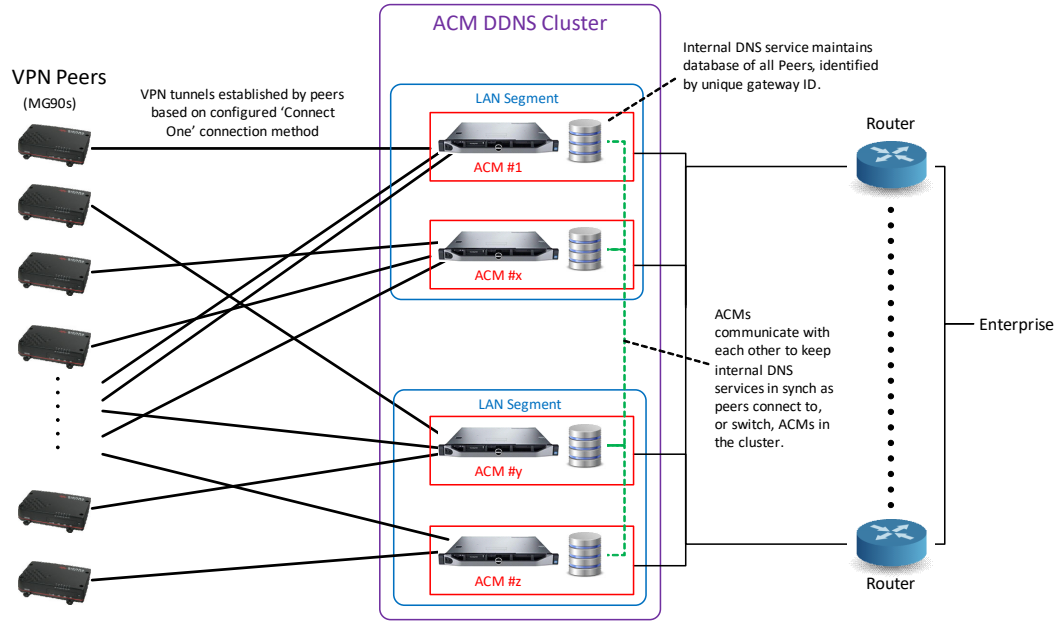

Figure 1-3: Example—ACMs in DDNS Configuration

2: Virtual Router Redundancy Protocol (VRRP)

Overview

The ACM technology utilizes virtual router redundancy protocol (VRRP) to ensure that services are available in the event that an ACM goes down. VRRP requires two or more ACMs and allows this cluster of ACMs to act as one virtual ACM—one ACM is the master server that provides all services, and the other ACMs are available to take the master's place if it fails. The group of ACMs that make up this cluster is known as a VRRP Group. The ACMs in the group exchange multicast packets to notify each other that each is still alive. If the current master ACM stops broadcasting, then a backup ACM in the VRRP group will take over as the master.

VRRP Description

Devices accessing an ACM from outside connect to a virtual IP address and therefore have no knowledge of which ACM they are connected to. For a cluster of ACMs to operate under VRRP:

- Each ACM must be assigned the same VRRP group number
- Each ACM must be mapped to the same virtual IP address
- Each ACM must be on the same subnet

Within a VRRP group, the ACM with the highest priority is elected as the "master", which is the ACM that devices connect to, as shown in Figure 2-1.

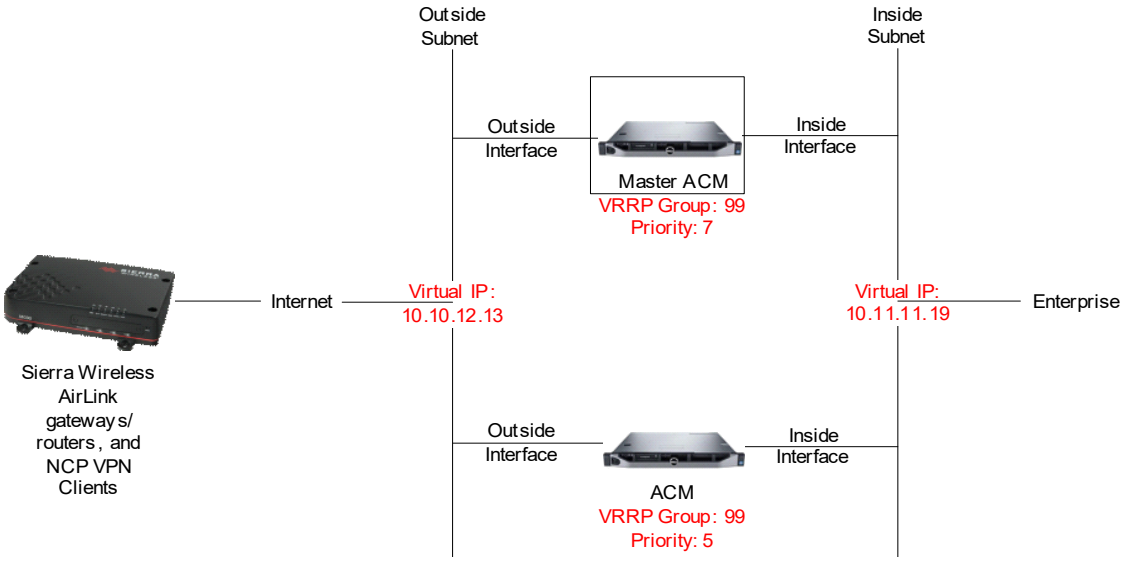

Figure 2-1: A master ACM is the ACM with the highest priority

In the event that this master goes down, the ACM with the next highest priority will be elected as the new master, and all ACMs are notified, as shown in Figure 2-2 on page 10. If two ACMs have the same priority, the server with the higher IP address will be elected as the master (note that during startup, this conflict is resolved by electing the server that becomes active first).

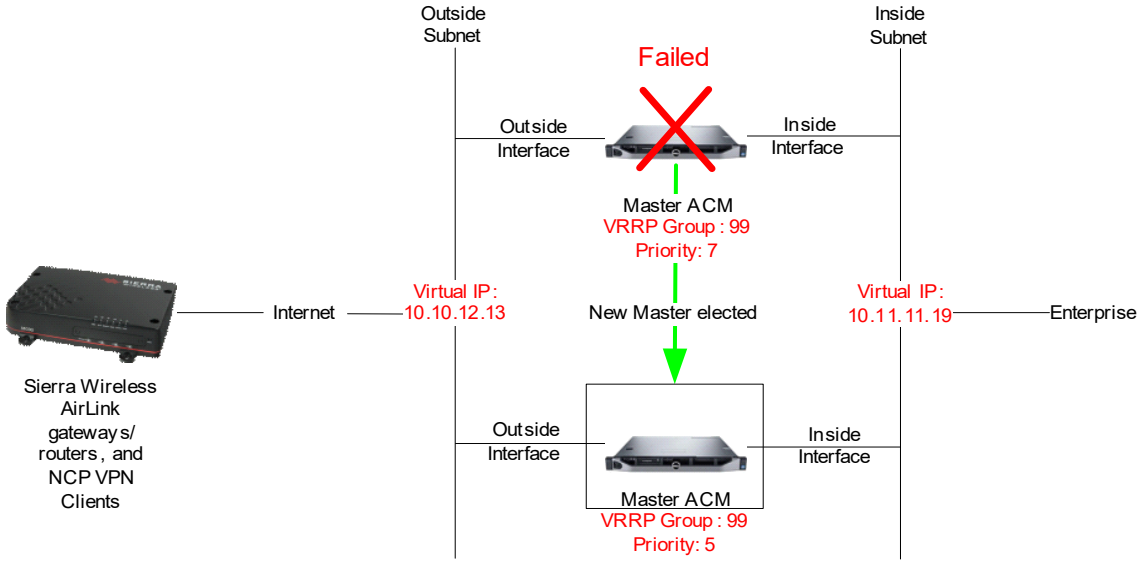

Figure 2-2: The next highest priority ACM will become the master upon a failure

VRRP's "pre-emption" feature allows for automatic election of a new, higher priority ACM. This is useful in the situation where an ACM has been elected as the master after the failure of a high priority ACM, and then a new higher priority ACM is added. In this case, the newly added ACM will automatically pre-empt the current, lower priority master and take over as the new master, as shown in Figure 2-3.

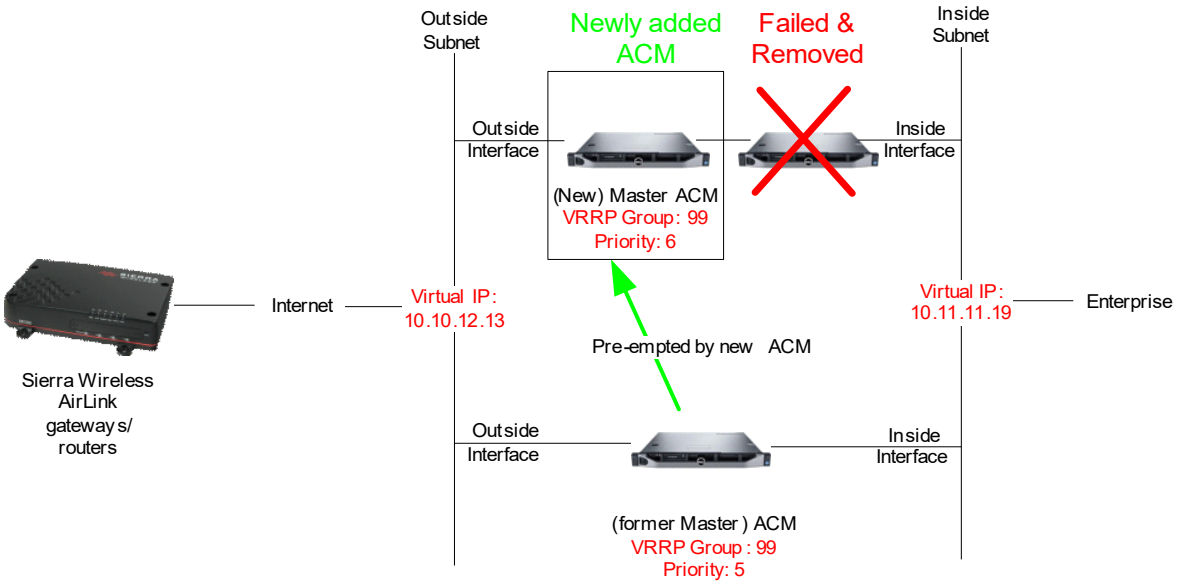

Figure 2-3: Pre-emption allows a new, higher-priority ACM to take over as the master

VRRP allows the failure of inside interfaces on an ACM to trigger a fail over to a backup ACM through a feature called "sync groups". For example, if the inside interface on a master ACM fails, this will cause the master ACM and its outside interface to fail as well. The backup ACM would then be elected as the master, as shown in Figure 2-4 on page 11.

Important: *Sync groups must be used when configuring VRRP.*

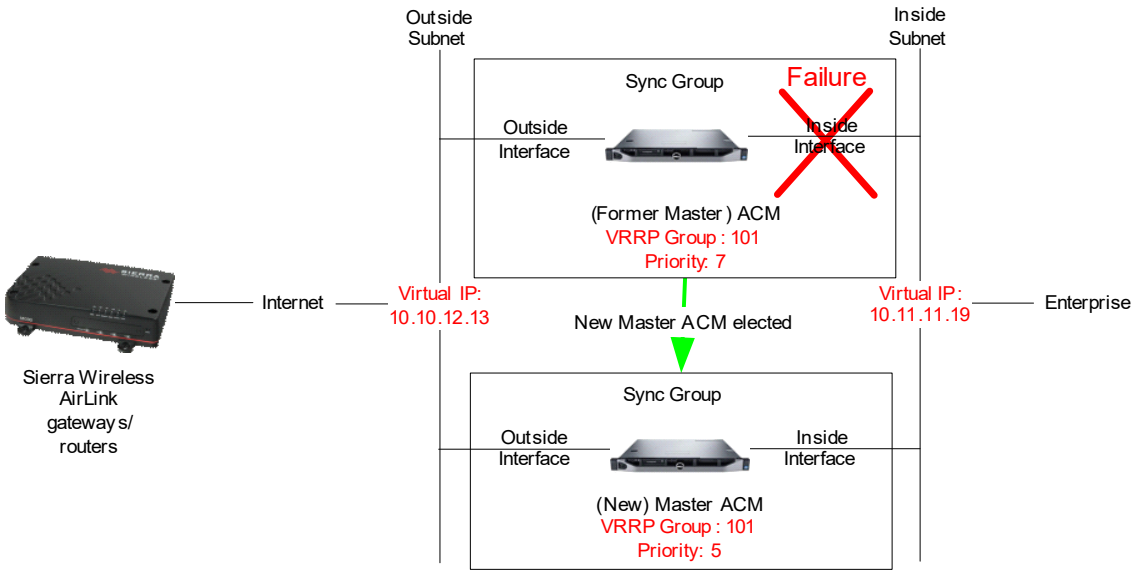

Figure 2-4: A Sync Group elects a new master ACM upon the failure of an inside interface

VRRP Configuration

Configuring Master and Backup ACMs for VRRP

To support VRRP, the eth0 and eth1 interfaces for all the ACMs (master and backups) must be configured.

For each ACM:

1. Enable VRRP—Create the eth0 and eth1 VRRP configuration nodes and assign different VRRP group numbers to them.
   - If ACMs are deployed in physical servers (i.e. appliances) or VMware provides Layer 2 services:

```
set interfaces ethernet eth0 vrrp vrrp-group <VRRP_GROUP_#_eth0> rfc3768-compatibility true
set interfaces ethernet eth1 vrrp vrrp-group <VRRP_GROUP_#_eth1> rfc3768-compatibility true
```

   - If VMware does not provide Layer 2 services:

```
set interfaces ethernet eth0 vrrp vrrp-group <VRRP_GROUP_#_eth0> rfc3768-compatibility false
set interfaces ethernet eth1 vrrp vrrp-group <VRRP_GROUP_#_eth1> rfc3768-compatibility false
```

Important: *Use the same eth0 and eth1 VRRP group numbers on each ACM.*

2. Specify the virtual IP addresses of the eth0 and eth1 VRRP groups:

```
set interfaces ethernet eth0 vrrp vrrp-group <VRRP_GROUP_#_eth0> virtual-address <VIRTUAL OUTSIDE-IPADDRESS/SUBNET>
set interfaces ethernet eth1 vrrp vrrp-group <VRRP_GROUP_#_eth1> virtual-address <VIRTUAL INSIDE-IPADDRESS/SUBNET>
```

3. If desired, enable pre-emption:

```
set interfaces ethernet eth0 vrrp vrrp-group <VRRP_GROUP_#_eth0> preempt true
set interfaces ethernet eth1 vrrp vrrp-group <VRRP_GROUP_#_eth1> preempt true
```

4. Set the ACM's priority (to keep things simple, you can use the same priority # for eth0 and eth1):

```
set interfaces ethernet eth0 vrrp vrrp-group <VRRP_GROUP_#_eth0> priority <PRIORITY #>
set interfaces ethernet eth1 vrrp vrrp-group <VRRP_GROUP_#_eth1> priority <PRIORITY #>
```

Important: *Make sure to set the master ACM's priority higher than the priorities of all backup ACMs.*

5. Add a sync group configuration to the eth0 and eth1 VRRP groups—the sync group name is an arbitrary string (e.g. "ALPHA", "MyGroup", etc.):

```
set interfaces ethernet eth0 vrrp vrrp-group <VRRP_GROUP_#_eth0> sync-group <SYNC GROUP NAME>
set interfaces ethernet eth1 vrrp vrrp-group <VRRP_GROUP_#_eth1> sync-group <SYNC GROUP NAME>
```

Important: *Use the same sync group name for both interfaces. You can use the same or different names for different ACMs, since the name is not communicated between them.*

6. Optionally, display the VRRP configurations to confirm the settings were entered properly:

```
show interfaces ethernet eth0 vrrp
show interfaces ethernet eth1 vrrp
```

7. Commit the configuration:

```
commit
```

Forcing an Artificial Failover

If needed, the currently operating master ACM and a backup ACM within a cluster can switch roles (backup becomes the master, and the master becomes a backup) by forcing an artificial failover.

To accomplish this, pre-emption must be enabled on both ACMs, and the priority of the backup ACM must be higher than the priority of the current master ACM—increase the backup ACM priority or decrease the master ACM priority.

Example:

Consider a VRRP cluster consisting of:

- ACM#1 is the master (priority=150)
- ACM#2 is the backup (priority=100)

ACM#2 will become the master if either of the following priority changes are made:

- ACM#2 priority is set to >150 (higher than ACM#1 priority)
  e.g. The following commands (entered on ACM#2) will increase ACM#2 priority to

200, causing it to become the new master (and ACM#1 to become the backup):

```
#set interfaces ethernet eth1 vrrp vrrp-group 99 priority 200
#set interfaces ethernet eth0 vrrp vrrp-group 1 priority 200
#commit
#save
```

- ACM#1 priority is set to <100 (lower than ACM#2 priority)
  e.g. The following commands (entered on ACM#1) will decrease ACM#1 priority to 50, causing ACM#2 to become the new master (and ACM#1 to become the backup):

```
#set interfaces ethernet eth1 vrrp vrrp-group 99 priority 50
#set interfaces ethernet eth0 vrrp vrrp-group 1 priority 50
#commit
#save
```

Important: *The state change (backup ACM and master ACM switching roles) occurs immediately if pre-emption is enabled. If pre-emption is not enabled, the priority change will not force the ACMs to switch roles.*

3: DNS Load Balancing

Overview

The ACM supports DNS load balancing as a high availability configuration method for AirLink oMG2000/500 and AirLink MG90 routers. With this method, ACM servers are pooled and the VPN peers (routers) are spread across the servers to distribute the overall load. If any server fails, its VPN peers shift over to the remaining servers.

DNS Load Balancing Example

The following figure shows an ACM server pool with three ACMs. In this example, the network architect determined:

- Two ACMs are required to carry the total load for the VPN peers (each carrying 50%).
- Adding a third ACM (as shown in this figure) provides protection against one server failing.
- For greater protection, and to spread the load even more, a fourth ACM could be added to the pool, where each would carry 25% of the load.

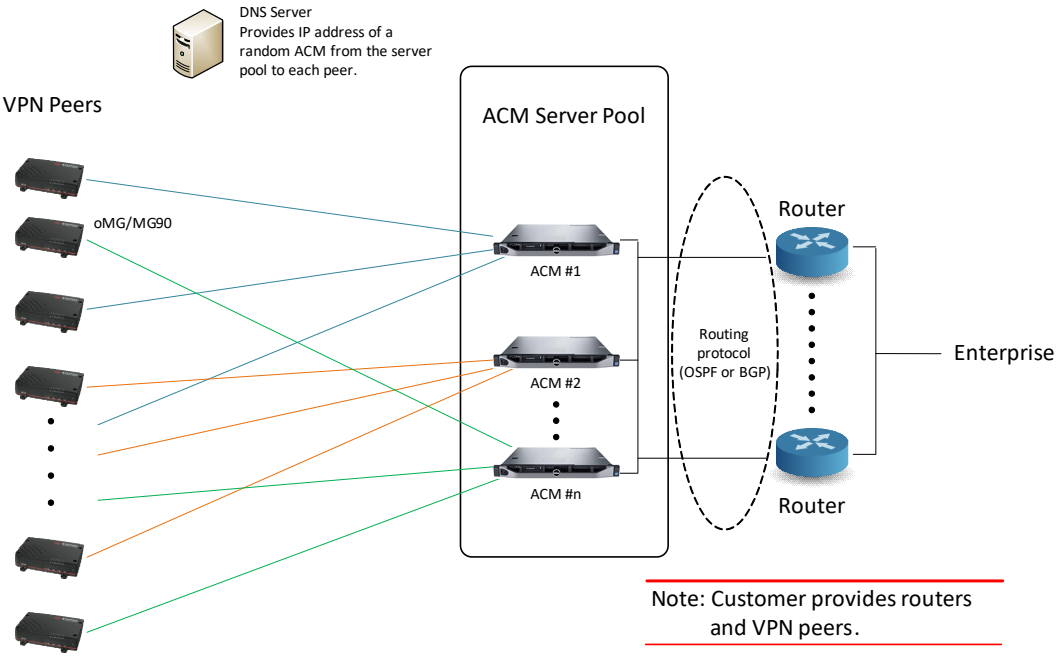

Figure 3-1: DNS Load Balancing Configuration in Normal Operation

At some point, if one of the three ACMs in this example fails (as shown in Figure 3-2):

- The peers of the failed ACM switch over to the other two ACMs.
- The two remaining ACMs now each carry 50% of the peers.
- The pool is no longer protected against another ACM failure (as noted, the network architect determined two ACMs are the minimum required to carry the load).
- A new ACM must be added to restore failover protection. Otherwise, if one of the remaining two ACMs also fails, all VPNs would have to be carried by a single ACM, which does not have the necessary capacity.

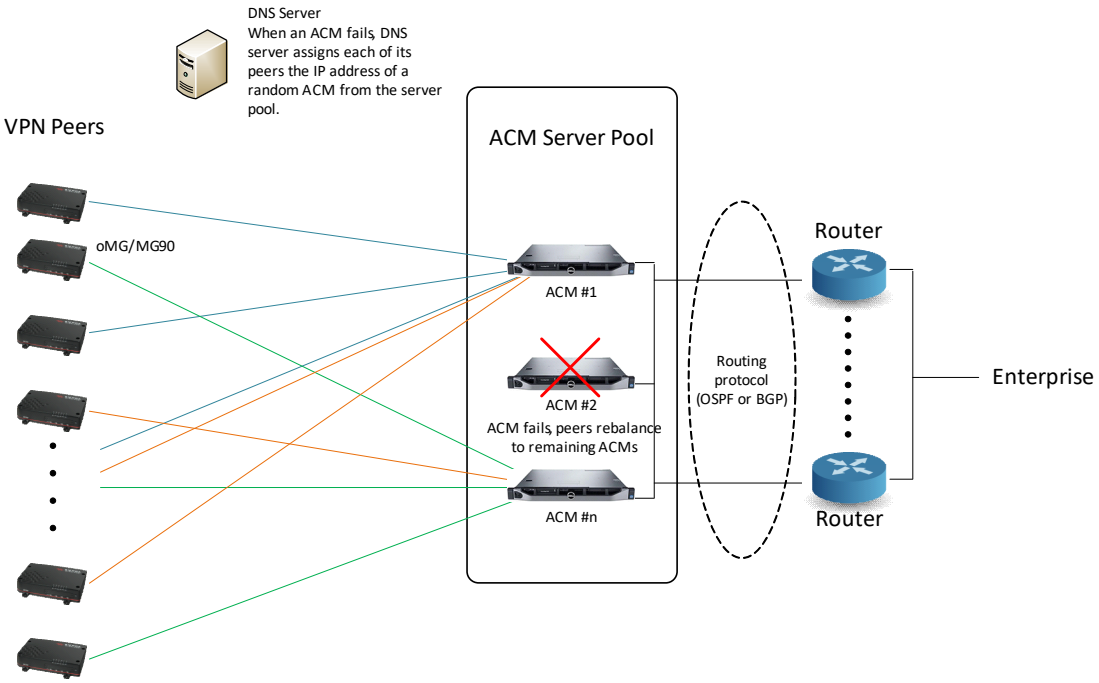

Figure 3-2: DNS Load Balancing Configuration After Failover

DNS Load Balancing Setup

To implement DNS load balancing, follow this suggested procedure to gather necessary information from the customer and configure ACM servers and VPN peers (MG90 routers, oMG gateways):

1. Determine ACM server requirements:
   - FIPS or non-FIPS ACM servers
   - Routing protocol to use (OSPF (typically used) or BGP)
   - Max throughput per server (ACM_thru_max)

Table 3-1: Max Throughput Rates (ACM 2.0.1)[a]

ACM Type	Encryption	Rate (Mbps)
Non-FIPS	AES128-MD5	898
	AES128-SHA1	897
	AES256-SHA512	885
	3DES-SHA512	538
FIPS	AES128-GCM16	890
	AES256-GCM16	898
	AES256-SHA512	883

a. Dell R230 host device

   - Total number of VPN peers (num_peers)

- Average VPN peer throughput requirement (peer_thru_avg)
- Calculate:
  - Number of peers per ACM:
    ACM_peers = INT(ACM_thru_max / peer_thru_avg])
  - Number of ACMs required (minimum):
    ACM_min = num_peers / ACM_peers
  - Protection factor (ACM_protect)—the number of ACMs that can be unavailable at the same time?
  - Total ACMs required:
    ACM_total = ACM_min + ACM_protect

2. Prepare DNS for the server pool:
   a. Customer must provide the FQDN (Fully Qualified Domain Name) to use for the ACM server pool.
   b. Customer must contact the DNS provider/administrator to request all ACM external IP addresses to be associated with the FQDN.
   c. Confirm the IP addresses were added properly. Use "*nslookup <FQDN>*" to return a list of IP addresses assigned to the FQDN.

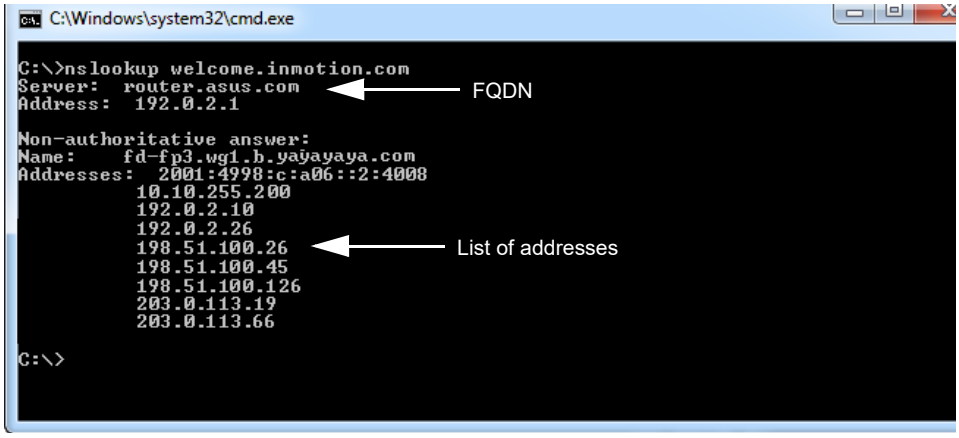

Figure 3-3: nslookup Example (Eight IP Addresses for one FQDN)

3. Configure the ACM servers:
   a. Configure every ACM in the server pool for VPN usage, as described in the AirLink Connection Manager Installation and Operations Guide.
   b. Configure every ACM in the server pool with similar routing protocol (OSPF or BGP) parameters:
      - OSPF example configuration—see OSPF Protocol on page 19.
      - BGP example configuration—see BGP Protocol on page 20.

c. Configure every ACM in the server pool to support *all* the VPN peers that will connect to the pool. (Every ACM must have the full peer list, since peers will be randomly distributed to individual ACMs.)

- For each VPN peer, set the authentication ID to the FQDN (otherwise the ACM will report an authentication failure). The following example shows an oMG (identified by its ESN (@H100111G1050) that has been added as a peer.

```
admin@oCM-Sun# show vpn ipsec site-to-site peer @H100111G1050
    authentication {
        id welcome.ocm.inmotion
        mode pre-shared-secret
        pre-shared-secret newworld
    }
    ike-group 2
    local-ip 10.10.255.200
    tunnel 1 {
        esp-group 1
        local {
            subnet 10.1.193.0/24
        }
        remote {
            network 172.25.105.0/24
        }
    }
```

Note: In the example above, the pre-shared-secret will be unique for individual peers, the remote network will vary (same for some peers, but not all), and the local-ip will be common for all peers.

After all server pool ACMs are configured identically, each VPN peer must be configured to connect to the server pool. (Typically, this will be done by configuring one peer, and then distributing ('pushing') that peer's configuration to all other peers of the same type (e.g. oMG gateway or MG90 router)—the method for distributing the configuration depends on the type of peer.)

4. Configure peers to connect to the server pool. To configure an oMG gateway or MG90 router peer, use its Local Configuration Interface (LCI) to set appropriate parameters in the IPsec VPN Configuration screen (WAN > VPNs):
   - In the Server Address field, enter the FQDN of the server pool.

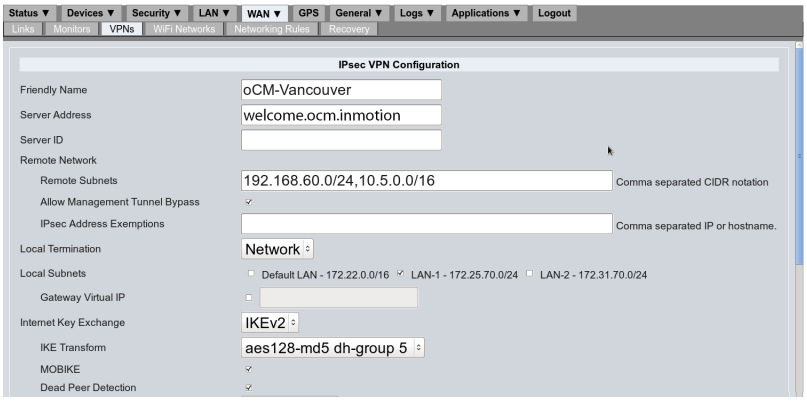

Figure 3-4: IPsec VPN Configuration (oMG/MG90 Local Configuration Interface)

5. Test that the peer can connect to the ACM server. If the connection does not work, wait a minute before trying again.
6. When the configuration has tested properly, use AMM to distribute the peer's configuration file to all the other peers.

Configuration File Examples

The OSPF and BGP protocol parameters shown in the snippets below are examples only—you must set appropriate parameters for the customer's configuration. Most parameters will be the same for all ACMs in the server pool, but some will be unique (for example, the OSPF router-id is unique for each ACM).

OSPF Protocol

The following is a snippet of a configuration file with OSPF protocol parameters:

```
protocols {
    ospf {
        area 1 {
            area-type {
                nssa {
                    translate candidate
                }
            }
            network 192.168.60.0/24
        }
        parameters {
            abr-type cisco
            router-id 192.168.60.252
        }
        redistribute {
            kernel {
                metric-type 2
            }
        }
    }
    static {
    }
}
```

BGP Protocol

The following is a snippet from a configuration file with BGP protocol parameters:

```
protocols {
    bgp 100 {
        neighbor 10.6.0.2 {
            disable-connected-check
            remote-as 200
        }
        neighbor 10.6.0.3 {
            disable-connected-check
            ebgp-multihop 2
            remote-as 200
        }
        network 10.5.0.0/16 {
        }
        network 172.18.0.0/16 {
        }
        network 192.168.60.0/24 {
        }
        network 192.168.100.0/24 {
        }
        parameters {
            router-id 10.6.0.5
        }
        redistribute {
            connected {
            }
            kernel {
            }
        }
        timers {
            holdtime 30
            keepalive 60
        }
    }
    static {
    }
}
```

4: Dynamic DNS

Overview

The ACM supports Dynamic DNS (DDNS) as a high availability load-balancing/failover configuration method for AirLink MG90 routers running MGOS 4.3 or later that are set up for DDNS.

With this method, up to 10 ACM appliances operate as part of a cluster where each ACM hosts its own DNS server. Each server maintains its own DNS database of the gateway IDs assigned to VPN peers (i.e. MG90s) connected to each of the cluster's ACMs. The ACMs do not need to be on the same networks as long as they have connectivity (using TCP ports 9100 and 4369) via inside or outside interfaces.

Each time a VPN peer establishes a tunnel with an ACM in the cluster, the ACM communicates the tunnel details to the other ACMs. Then, if any ACM goes down, its peers detect the down tunnel and re-establish their tunnels to the remaining ACMs in the cluster.

Note: The routers must establish tunnels in Host-to-LAN mode, using IKEv2. DDNS does not support IKEv1 connections or LAN-to-LAN mode.

On the enterprise (inside) side, clients can communicate directly with any peer (regardless of the ACM it is connected to) by using the gateway ID that is maintained in the identical DNS databases stored at each ACM.

Cluster Example

Figure 4-1 illustrates the general architecture of an ACM cluster.

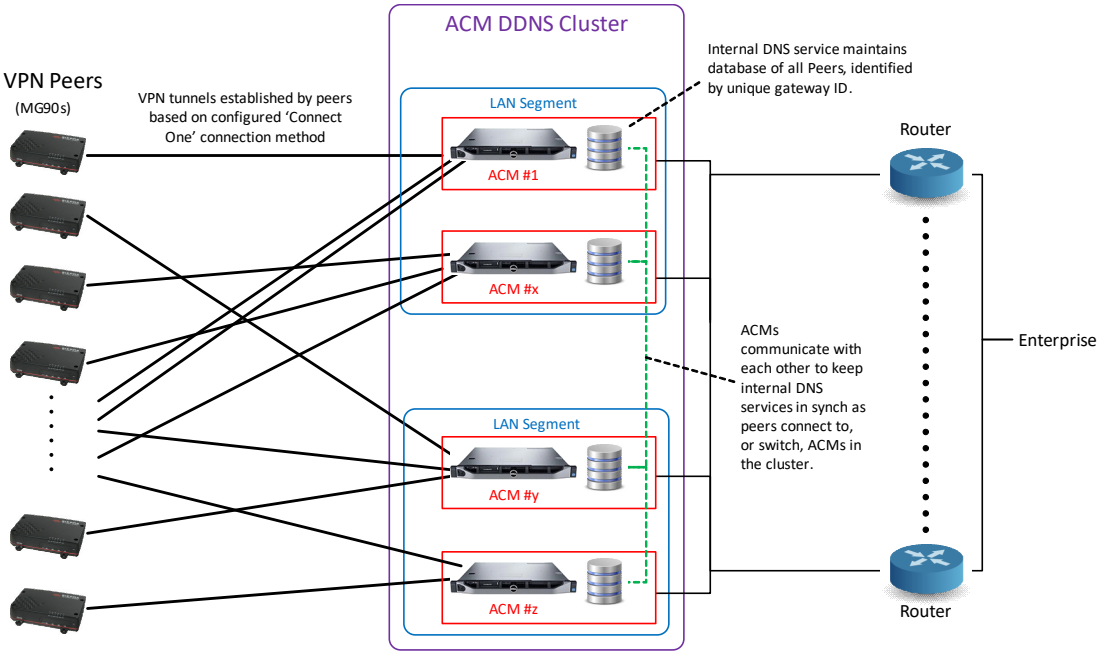

Figure 4-1: ACM Cluster (DDNS topography)

DDNS Description

The following procedure describes how to set up an ACM cluster and the peers that will connect to it, and how the cluster functions once it is configured:

1. Prepare ACMs and VPN peers (MG90s) for DDNS:
   - Configure and enable DDNS on each ACM in the cluster (see ACM Configuration for DDNS on page 23.)

     As each ACM is enabled for DDNS, it begins syncing DNS databases with other ACMs in the cluster.
   - Configure VPN peers to connect to the cluster (either via the MG90 LCI, or via AMM):
     - Create VPN profiles for each ACM in the cluster, and configure the Internet Key Exchange protocol as IKEv2.
     - For each WAN link or Wi-Fi Network that is to connect to the cluster, configure the Connection Method as “Connect One”.

     Note: If a VPN peer uses the "Connect All" connection method, it can still attach to any of the ACMs in the cluster but will not have the benefits of DDNS.

     For details, refer to the AirLink MG90 Software Configuration Guide.
2. When DDNS has been enabled, the following behavior occurs:
   - VPN peers connect to ACMs:
     i. Peer requests tunnel to one of the ACMs (chosen by the MG90 based on selected Connect One option).
     ii. The chosen ACM assigns a unique gateway ID to the peer. The ID is typically adapted from:
        - The certificate value, if using certificate authentication. (e.g. "CN=<value>", the <value> is used)
        - The peer’s serial number, if using PSK authentication.

        The ID may appear in log and configuration files.

Note: To conform to rules for DNS host names (only numbers, letters, and hyphens allowed), if the ID contains any other ’special’ characters (e.g. ’.’, ’/’, etc.)

- *The special characters are changed to hyphens (‘-’), then*
- *Groups of 2 or more hyphens are changed to a single hyphen, then*
- *Leading/trailing hyphens are removed.*

In the unlikely circumstance that the modified IDs for two routers are the same (for example, routers with very similar IDs using different special characters, which match after conversion), the one with the most recent VPN connection will be identified in DNS.

     iii. The ACM establishes the tunnel and updates its internal DNS server with the tunnel information: <gateway ID>.<domain> = <tunnel virtual IP>.
     iv. The ACM notifies all other ACMs in the cluster to update their internal DNS servers with the new tunnel details. All ACMs now have the same information about active tunnels.
   - Connected VPN peers switch to different ACMs—When a peer switches ACMs for any reason, the new ACM updates all other ACMs with the new tunnel information (new tunnel virtual IP).

- ACM goes down/becomes unreachable—If an ACM becomes unreachable for any reason, its VPN peers will detect their down tunnels and automatically attempt to connect to another ACM based on the peer's configured connection method (Connect One—random or round-robin).

Note: Corporate DNS servers must be configured to either forward or delegate <domain> to all ACM servers. This allows inside clients to locate any peer (regardless of the ACM it is connected to) by referencing <gateway ID>.<domain>.

ACM Configuration for DDNS

ACMs are configured for DDNS using the ipsec-ddns service configuration attributes described in Table 4-1.

Table 4-1: ipsec-ddns Configuration Attributes

Attribute	Description
enable	Enable/disable DDNS (via "set" or "delete" command)
secret <ddns_pass>	Authentication password for DDNS Min. length—8 characters *Note: All ACMs in the cluster must use the same password.*
domain <domain_name>	DNS domain name Domain name is combined with a VPN peer's gateway ID and stored in the DNS server database. Example: "domain mynetwork.org" *Note: All ACMs in the cluster must use the same domain name.*
self <address>	Address of the ACM being configured. Format—IPv4 (typically used) or FQDN. *Note: The self address must match the ACM's physical internal address for either the inside or outside interface (not a VRRP virtual address).* *Note: Same address must be used in the "ipsec-ddns peer" configuration on other ACMs in the cluster.*
peer <address>	Address of another ACM in the cluster. Format—IPv4 (typically used) or FQDN. *Note: This must be the address that the peer ACM uses in is "ipsec-ddns self" configuration.*

Initial ACM DDNS Cluster Setup

To set up an ACM DDNS cluster, repeat the following process to configure each ACM:

1. Identify the current ACM using the physical IP address or FQDN of its inside or outside interface (not a VRRP virtual address):
   ```
   set service ipsec-ddns self <address_of_current_ACM>
   ```
2. Repeat the following command to add each of the other ACMs in the cluster as peers, using their physical inside or outside interface IP addresses or FQDNs (not VRRP virtual addresses):
   ```
   set service ipsec-ddns peer <address_of_peer>
   ```

Note: For self and peer addresses, the inside address is preferred for DDNS. If the ACMs are on different LAN segments (networks) in the cluster and cannot communicate across the inside interface, use the outside address.

3. Add the DDNS password (must be entered the same on every ACM):
   ```
   set service ipsec-ddns secret <ddns_pass>
   ```
4. Set the DNS domain name (must be entered the same on every ACM):
   ```
   set service ipsec-ddns domain <domain_name>
   ```
5. Enable DDNS on the current ACM:
   ```
   set service ipsec-ddns enable
   ```
   The first ACM enabled for DDNS will create the DDNS cluster. Subsequent ACMs discover the created cluster and join it.
6. Commit and execute the new configuration details:
   ```
   commit
   ```
   The commit command for the first ACM should return the response "Creating new ipsec-ddns cluster", and subsequent ACMs should return "Joined existing ipsec-ddns cluster". If a different response is returned, refer to the troubleshooting topic Incorrect response when Initializing or Adding to a Cluster on page 29.

Adding an ACM to the Cluster

To add a new ACM to an existing (already configured) cluster:

1. On each of the other ACMs in the cluster, add the new ACM as a peer:
   ```
   set service ipsec-ddns peer <address_of_peer>
   commit
   ```
   The commit command should return "Joined existing ipsec-ddns cluster". If a different response is returned, refer to the troubleshooting topic Incorrect response when Initializing or Adding to a Cluster on page 29.
2. Configure the new ACM as described in Initial ACM DDNS Cluster Setup above.

Removing an ACM from the Cluster

To remove an ACM from an existing (already configured) cluster:

1. Disable DDNS on the ACM to be removed:

```
delete service ipsec-ddns enable
commit
```

Note: Other DDNS attributes do not have to be removed.

2. On each of the other ACMs in the cluster, remove the ACM as a peer:

```
delete service ipsec-ddns peer <address_of_removed_ACM>
commit
```

Example—2-ACM Cluster

The following example sets up a cluster with two ACMs:

1. Set up first ACM. (Note: This ACM's address is 10.1.65.1)

```
set service ipsec-ddns self 10.1.65.1
set service ipsec-ddns peer 10.1.65.2
set service ipsec-ddns secret mysecretpassword
set service ipsec-ddns domain <domain_name>
set service ipsec-ddns enable
commit
```

2. Set up second ACM: (Note: This ACM's address is 10.1.65.2)

```
set service ipsec-ddns self 10.1.65.2
set service ipsec-ddns peer 10.1.65.1
set service ipsec-ddns secret mysecretpassword
set service ipsec-ddns domain <domain_name>
set service ipsec-ddns enable
commit
```

A: Troubleshooting

VRRP Troubleshooting

View VRRP Configuration Details

To view VRRP configuration details on the master ACM or backup ACM, use the *show vrrp* command with appropriate parameters, as shown below.

On the master:

```
admin@ACM:~$ show vrrp interface eth0

Physical interface: eth0, Source Address 192.168.3.112
  Interface state: up, Group 99, State: master
  Priority: 250, Advertisement interval: 1,   Authentication type: none
  Preempt: true, VIP count: 1, VIP: 192.168.3.33/24
  Master router: 192.168.3.112
  Sync-group: ACM
  Last transition: 1w2d2h39m26s

admin@ACM:~$ show vrrp interface eth1

Physical interface: eth1, Source Address 192.168.9.107
  Interface state: up, Group 101, State: master
  Priority: 250, Advertisement interval: 1, Authentication type: none
  Preempt: true, VIP count: 1, VIP: 192.168.9.33/24
  Master router: 192.168.9.107
  Sync-group: ACM
  Last transition: 1w2d2h39m29s

admin@ACM:~$ admin@ACM:~$ show vrrp summary
                VRRP    Addr                    Interface       VRRP
Interface       Group   Type    Address         State           State
---------       -----   ----    -------         -----           -----
eth0            99      vip     192.168.3.33/24 up              master
eth1            101     vip     192.168.9.33/24 up              master
admin@ACM:~$
```

On the backup ACM:

```
admin@ACM:~$ show vrrp interface eth0

Physical interface: eth0, Source Address 192.168.2.95
  Interface state: up, Group 99, State: backup
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: true, VIP count: 1, VIP: 192.168.3.33/24
  Master router: 192.168.3.112, Master Priority: 250
  Sync-group: ACM
  Last transition: 1w2d2h41m34s

admin@ACM:~$ show vrrp interface eth1

Physical interface: eth1, Source Address 192.168.9.106
  Interface state: up, Group 101, State: backup
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: true, VIP count: 1, VIP: 192.168.9.33/24
  Master router: 192.168.9.107, Master Priority: 250
  SyncACM-group: ACM
  Last transition: 1w2d2h41m37s
```

DNS Load Balancing Troubleshooting

Reduce Time to Switch From Unavailable ACM

When an ACM server becomes unavailable (failed, or a planned outage such as for a system upgrade), the DNS server can continue to assign its IP address to VPN peers, causing the peers to attempt to connect to the unavailable server. The peers then have to repeat the process of getting a new IP address.

If the server outage is expected to be long (or is planned in advance), you should consider having the server removed from the list provided by the DNS server until the ACM becomes available (new server, or planned outage is completed).

DDNS Troubleshooting

Incorrect response when Initializing or Adding to a Cluster

When doing Initial ACM DDNS Cluster Setup on page 24 or when Adding an ACM to the Cluster on page 24, one ACM cluster should be created:

- When the first ACM creates the cluster, the response "Creating new ipsec-ddns cluster" should be received. If a different response is received, disable and re-enable the DDNS service as described below.
- When subsequent ACMs are added, the response "Joined existing ipsec-ddns cluster" should be received. If a different response is received, the ACM might be experiencing difficulty connecting to other ACMs in the cluster—wait for connectivity to return, then disable and re-enable the DDNS service as described below.

To disable and re-enable DDNS, use the following commands:

```
delete service ipsec-ddns enable
commit
set service ipsec-ddns enable
commit
```