Network Mobility (NEMO)

This chapter describes the system\'s support for Network Mobility (NEMO) and explains how it is configured. The product Administration Guides provide examples and procedures for configuration of basic services on the system. It is recommended that you select the configuration example that best meets your service model and configure the required elements for that model, as described in the respective product Administration Guide, before using the procedures in this chapter.

NEMO Overview

When enabled through a feature license key, the system includes NEMO support for a Mobile IPv4 Network Mobility (NEMO-HA) on the existing Enterprise Home Agent (EHA) platform to interconnect LAN segments behind Mobile Routers (MRs) equipped with a 3G interface with Fixed Networks served by the Private IP (PIP) networks and Wavelength Division Multiplexing Networks (WDN) . The new NEMO functionality allows bi-directional communication that is application-agnostic between users behind the MR and users or resources on the Fixed Network sites.

The same NEMO4G-HA service and its bound Loopback IP address supports NEMO connections whose underlying PDN connection comes through GTP S5 (4G access) or PMIPv6 S2a (eHRPD access).

The following figure shows a high-level view of NEMOv4 Architecture.

Figure 1: NEMO Overview
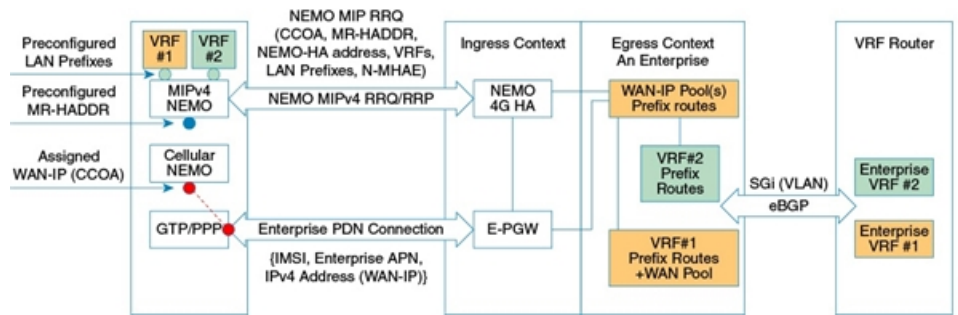
Features and Benefits

The system supports the usage of dynamically learned, overlapping customer prefixes. These prefixes are advertised through BGP in a manner similar to pool routes in the current EHA implementation. NEMO includes the following features:

- Interoperates with the Mobile IPv4 NEMO implementation of the Cisco ISR CPE routers.

  - Protocol behavior.
  - Message structures, formats and encoding.
  - Specific flags and parameters.

- Compatible with the specifics of the Mobile IPv4 NEMO operation of the Cisco ISR CPE routers.

  - Support for the second Mobile IPv4 NEMO Control Messaging.
  - Support for GRE NEMO-Tunnel termination (One NEMO-Tunnel per MR).
  - Support for explicit LAN Prefix registration mode.

- Support for private customer addressing, routing and traffic segmentation

  - Private and overlapping WAN-IP addresses.
  - Private and overlapping LAN Prefixes.
  - Customer LAN Prefix advertisement from the EHA egress contexts via BGP
  - Customer traffic segmentation and mapping of the incoming NEMO-Tunnels to the appropriate VLAN/VRF.

- Ability to seamlessly integrate with the existing MPN service environment.

  - Selective suppression or replacement of specific fields in the NEMO Mobile IPv4 Control Messaging sourced by the CPE routers.
  - Correlation of the incoming NEMO Control and Forwarding traffic with the existing control and flow structures related to the HWIC device processing by the underlying IS-835/MIP logic in the EHA.
  - Compatibility with the existing AAA requirements.

- The HWIC IS-835/PPP/MIP timers shall be compatible with today's EHA implementation.

  - PDSN's PPP idle timeout (2 hrs)
  - PDSN's PPP absolute timer (24 hrs)
  - First MIP session re-registration timer (1hr 55min)

- NEMO-HA is not required to generate AAA accounting records (START/STOP) for the NEMO MIP session. On the other hand, accounting records are generated for the MR's HWIC MIP session, just like with any other MIP sessions.
- NEMO-HA supports explicit registration mode and does not require authorization/validation of the LAN Prefixes sent by the MR.
- If the authentication of the NEMO MIP session fails, the underlying HWIC IS-835/MIP session is maintained since the NEMO function may attempt to establish the NEMO-Tunnel again.
- NEMO-HA is supported by ICSR. All the information related to NEMO-HA (NEMO MIP session state, and so on) is synchronized with the standby EHA and the total failure of the active EHA does not require existing NEMO tunnels to be re-established.
- NEMO-HA has dynamically advertise the LAN prefixes of any given MR to the upstream corporate router, but it does have the ability to suppress the MR's MR-HADDR address from the route advertisement via route-map configuration.
- The existing EHA support for interface MTU configuration also applies to NEMO-HA enabled systems.
- The NEMO-HA supports Local Authentication - the N-MHAE-SPI/KEY values are stored in the NEMO-HA. NEMO-HA supports two options to provision the SPI/KEY information in the MR's:

    - Individual MR level: each MR would has a unique SPI/KEY pair.
    - Enterprise level: each Enterprise uses unique security credentials and all the MR's of a given Enterprise uses the same SPI/KEY pair.

- A new RADIUS attribute (VSA) is supported that can be passed to the EHA during the establishment of the first IS-835/MIP session between the MR's HWIC and the EHA. This new RADIUS attribute represents the authorization of a second NEMO MIP RRQ for the associated MR. The EHA verifies if the new NEMO-related VSA is present in the Access-Accept for the first IS-835/MIP session. If so, NEMO-HA caches this information to properly authorize the second NEMO MIP session. This allows the AAA to control the authorization of NEMO sessions more efficiently without the need for a second AAA message.
- Upon any failure with the establishment of a second NEMO MIP session, the EHA does not take any actions with the underlying IS-835/MIP session. In other words, it does not tear down the first IS-835/MIP session.
- The NEMO-HA supports overlapping WAN-IP addresses for differing enterprises.
- RFC 5177 is supported.
- Enterprise VLANs are unique to the enterprise. Two different enterprises do not share the same VLAN ID in the egress context(s)
- If no NEMO-HA service is defined, it is not using NEMO.
- The NEMO HA support both dynamic address allocation and static address assignment.
- Multi-VRF - The existing design of HA NEMOv4 is extended to allow more than one VRF. For more information on Multi VRF, see *NEMOv4 with Multi-VRFs*.
- Enterprise minimal-registration-lifetime overwrite.
- PMIPv6 - Mobile Private Network (MPN) utilized Network Mobility Services (NEMO) to provide wireless connectivity between Enterprise Core Network and remote Enterprise sites over 3G/4G network, and supported only IPv4 addressing scheme. To expand the addressing scheme to IPv6, PMIPv6 support is now added. The LMA functionality will be provided externally by the ASR9000.

Engineering Rules

- Up to 300 virtual routing tables per context and 64 BGP peers per context.

- Up to 5k host routes spread across multiple VRFs per BGP process. Limited to 6000 pool routes per chassis.
- Up to 2048 VRFs per chassis.
- Up to 512K NEMO framed MNPs (Mobile Private Networks) per system.
- 32K routes per context.

Supported Standards

- IETF RFC 5177 (April 2008) "Mobile IPv4 NEMO Extensions,"
- IETF RFC 3025 (February 2001) "Mobile IP Vendor/Organization Specific Extensions"

NEMO Configuration

☞

Important   Commands used in the configuration samples in this section provide base functionality to the extent that the most common or likely commands and/or keyword options are presented. In many cases, other optional commands and/or keyword options are available. Refer to the *Command Line Interface Reference* for complete information regarding all commands.

To configure the system for NEMO:

1. Create a VRF on the router and assign a VRF name.

2. Set the neighbors and address family to exchange routing information with a peer router.

3. Redistribute connected routes between routing domains.

4. Create a NEMO HA.

5. Save your configuration to flash memory, an external memory device, and/or a network location using the Exec mode command **save configuration**. For additional information on how to verify and save configuration files, refer to the *System Administration Guide* and the *Command Line Interface Reference*.

Sample Configuration

```
config ingress
  context ingress
    interface <interface-name> loopback
        ip address <ipaddress> srp-activate
    exit
    interface <interface-name> loopback
        ip address <ipaddress> srp-activate
    exit
    interface <interface-name> loopback
        ip address <ipaddress>
    exit
    subscriber name <subscriber-name>
        encrypted password +A0ma96jknt7xul1ne8fk1kuled82o27x1i1fw6t103rqedigdfacp
        ip context-name egress1
          ip address pool name <pool-name>
```

```
            permission nemo
        exit
        ha-service <ha-service-name>
            mn-ha-spi spi-number <256>  encrypted secret
+A2ityhei4lza673nh1o9nqr4yqm2gsp0yv8eflng2tn2cyh5t1fbn hash-algorithm md5
            authentication mn-aaa noauth
              authentication mn-ha allow-noauth
            encapsulation allow keyless-gre
            min-reg-lifetime <300>
          bind address +A2ityhei4lza673nh1o9nqr4yqm2gsp0yv8eflng2tn2cyh5t1fbn hash-algorithm
 md5
        exit
        ha-service enterprise-ha1
            mn-ha-spi spi-number <256>  encrypted secret
+A0zsddshr3maez0b9j3izuk6q5612m1itttjmwyg16hiussxb5byv hash-algorithm md5 timestamp-tolerance
 65535
            fa-ha-spi remote-address ipaddress spi-number <256> encrypted secret
+A2yxbf7x14k8ko2aeef6fxrklft2zmir909mdp1n26ppovmlnw41w hash-algorithm md5 timestamp-tolerance
 65535
            authentication mn-ha allow-noauth
            revocation enable
            reg-lifetime <7200>
            bind address <ipaddress>
        exit
        ha-service <enterprise-ha2>
            mn-ha-spi spi-number <256>  encrypted secret
+A0zsddshr3maez0b9j3izuk6q5612m1itttjmwyg16hiussxb5byv hash-algorithm md5 timestamp-tolerance
 65535
            fa-ha-spi remote-address ipaddress spi-number <256> encrypted secret
+A2yxbf7x14k8ko2aeef6fxrklft2zmir909mdp1n26ppovmlnw41w hash-algorithm md5 timestamp-tolerance
 65535
             authentication mn-ha allow-noauth
            revocation enable
            reg-lifetime 7200
            bind address <ipaddress>
        exit
end
```

Ingress with new ComboHA feature.

☞

Important   Everything is same for NEMO except changes for IS-835/MIP session.

```
config
      interface <interface-name> loopback
          ip ranged-address ipaddress srp-activate
      exit
      ha-service <ha-service-name>
          mn-ha-spi spi-number <256>  encrypted secret
+A1ere1vystx2r7056dq8i1o5h3m1e0xpcwghhhe80e3q3251jggkf hash-algorithm md5 timestamp-tolerance
 65535
          fa-ha-spi remote-address ipaddress spi-number <256> encrypted secret
+A2xya6hjckox7c2964eftaufx1530khglavl7urf0nknletmp5dro hash-algorithm md5 timestamp-tolerance
 65535
          fa-ha-spi remote-address ipaddress spi-number <256> encrypted secret
+A1nexk5t8kjbav0vd7thh0yv4ly1ms68mtcermok0mxx2brve2ot7 hash-algorithm md5 timestamp-tolerance
 65535
          authentication mn-ha allow-noauth
          revocation enable
          reg-lifetime 7200
```

```
        bind address <ipaddress>
    exit
end
```

NEMO Egress

```
config
  context egress1
      ip vrf <vrf-name1>
          ip maximum-routes 4998
      exit
      ip vrf <vrf-name2>
          ip maximum-routes 4998
      exit
      ip vrf <vrf-name3>
          ip maximum-routes 4998
      exit
      ip routing overlap-pool
      ip pool cust1-f <ipaddress> private 0 group-name customer1 vrf <vrf-name1>
nexthop-forwarding-address ipaddress overlap vlanid 401 policy allow-static-allocation
      ip pool cust2-f ipaddress private 0 group-name customer2 vrf <vrf-name2>
nexthop-forwarding-address ipaddress overlap vlanid 402 policy allow-static-allocation
      ip pool <pool-name> ipaddress private 0 group-name customer3 vrf
<vrf-name3>nexthop-forwarding-address <ipaddress>overlap vlanid 403 policy
allow-static-allocation
      router bgp 1
          enforce-first-as
          neighbor <ipaddress> remote-as <1001>
          neighbor <ipaddress> update-source <ipaddress>
          neighbor <ipaddress> remote-as <1001>
          neighbor <ipaddress> update-source <ipaddress>
          ip vrf <vrf-name1>
              route-distinguisher 1 1
          exit
          address-family ipv4 vrf  <vrf-name1>
              redistribute connected
              neighbor <ipaddress> remote-as 1001
              neighbor <ipaddress> use-default-table
              neighbor <ipaddress> ebgp-multihop max-hop 255
              neighbor <ipaddress> update-source ipaddress
              neighbor <ipaddress> remote-as 1001
              neighbor <ipaddress> use-default-table
              neighbor <ipaddress> ebgp-multihop max-hop 255
              neighbor <ipaddress> update-source <ipaddress>
          exit
          ip vrf vrf-cust2
              route-distinguisher 1 2
          exit
              address-family ipv4 vrf <vrf-name2>
              redistribute connected
              neighbor ipaddress remote-as 1001
              neighbor ipaddress use-default-table
              neighbor ipaddress ebgp-multihop max-hop 255
              neighbor ipaddress update-source ipaddress
              neighbor ipaddress remote-as 1001
              neighbor ipaddress use-default-table
              neighbor ipaddress ebgp-multihop max-hop 255
              neighbor ipaddress update-source ipaddress
          exit
          ip vrf <vrf-name2>
              route-distinguisher 1 3
          exit
          address-family ipv4 vrf  <vrf-name3>
```

```
                        redistribute connected
                        neighbor <ipaddress> remote-as 1001
                        neighbor <ipaddress> use-default-table
                        neighbor <ipaddress> ebgp-multihop max-hop 255
                        neighbor <ipaddress> update-source <ipaddress>
                        neighbor <ipaddress> remote-as 1001
                        neighbor <ipaddress> use-default-table
                        neighbor <ipaddress> ebgp-multihop max-hop 255
                        neighbor <ipaddress> update-source <ipaddress>
                   exit
                   interface 29/1-sub401
                        ip address ipaddress
                   exit
                   interface 29/1-sub402
                        ip address <ipaddress>
                   exit
                   interface 29/1-sub403
                        ip address <ipaddress>
              exit
        end
```

NEMO MPLS Egress

```
        config
        context egress1
              ip vrf<vrf-cust1>
                   ip maximum-routes 4998
              exit
              ip vrf <vrf-cust2>
                   ip maximum-routes 4998
              exit
              ip vrf <vrf-cust3>
                   ip maximum-routes 4998
              exit
              mpls bgp forwarding
              ip pool pool1-b <ipaddress> private 0 srp-activate group-name customer1 vrf vrf1
        policy allow-static-allocation
              ip pool pool2-b <ipaddress> private 0 srp-activate group-name customer2 vrf vrf2
        policy allow-static-allocation
              ip pool pool3-b <ipaddress> private 0 srp-activate group-name customer3 vrf vrf3
        policy allow-static-allocation
              router bgp 1
                   router-id ipaddress
                   neighbor <ipaddress> remote-as 1001
                   neighbor <ipaddress> remote-as 1001
                   timers bgp keepalive-interval 10 holdtime-interval 30
                   address-family vpnv4
                        neighbor <ipaddress> activate
                        neighbor <ipaddress> send-community both
                        neighbor <ipaddress> activate
                        neighbor <ipaddress> send-community both
                   exit
                   ip vrf <vrf1>
                        route-distinguisher 1 1
                        route-target export 1 1
                        route-target import 1 1
                   exit
                   address-family ipv4 vrf <vrf1>
                        redistribute connected
                        redistribute static
                   exit
                   ip vrf <vrf2>
                        route-distinguisher 2 2
```

```
                        route-target export 2 2
                        route-target import 2 2
                exit
                address-family ipv4 vrf <vrf2>
                    redistribute connected
                    redistribute static
                exit
                ip vrf <vrf3>
                    route-distinguisher 3 3
                    route-target export 3 3
                    route-target import 3 3
                exit
                address-family ipv4 vrf <vrf3>
                    redistribute connected
                    redistribute static
                exit
        end
```