



Non-Blocking Multicast Service Reflection

- [NAT Guidelines and Limitations, on page 1](#)
- [Multicast to Multicast Ingress NAT, on page 1](#)
- [Multicast to Multicast Egress NAT, on page 2](#)
- [Examples for ENAT PIM Passive, on page 2](#)
- [Multicast to Unicast NAT, on page 3](#)
- [Examples for MU NAT PIM Passive, on page 3](#)
- [Unicast to Multicast NAT, on page 4](#)

NAT Guidelines and Limitations

The NBM Service Reflection has the following guidelines and limitations:

- Beginning with Cisco NX-OS Release 10.2(3)F, Unicast to Multicast NAT, Multicast to Unicast NAT, Multicast to Multicast NAT, and Egress NAT are supported on non-default VRF.
- If NAT config is present, config rollback is not supported (and will fail).
- In some cases, service interface re-configuration will be rejected, and to change it, a specific sequence may be required. Also, after re-configuration, NAT rules may not recover automatically and additional actions are required.
- Beginning with Cisco NX-OS Release 10.3(2)F, NAT is supported with sub-interface with "feature nbm" enabled.
- Beginning with Cisco NX-OS Release 10.3(2)F, egress service reflection (egress multicast NAT, and multicast to unicast NAT) supports Post-NAT Source IP to be IP Address of an egress interface. This enhancement is supported for regular multicast, and for NBM.

Multicast to Multicast Ingress NAT

The Ingress NAT allows translation of incoming (S,G) into a different source, group or both. All receivers inside the domain then can join the post translated flow. This feature is useful when multicast traffic:

- enters a network from a different domain with potentially overlapping address
- comes with an address that is not understood by applications in the network

The dynamic IGMP join or PIM join on a pre-translated route is not supported for ingress NAT.

Multicast to Multicast Ingress NAT works only in PIM active mode. The PIM passive mode is not supported.

Multicast to Multicast Egress NAT

The Egress NAT allows translating existing flow (S,G) to different source or group address on a per outgoing interface basis. This feature is useful for multicast distribution to external entities which may only accept a certain source or group address. It can also serve as a path to hide internal address space when flows are exposed to external entities.

The dynamic IGMP join or PIM join on a post-translated route is not supported for egress NAT.

Fault MO's are generated when there is a mismatch in bandwidth for pre-translated and post-translated flows.

In PIM-Passive mode, bandwidth management is done by an external controller for the flows and provisions both pre-translated and post-translated flows. The flow creation is made available through APIs.

Examples for ENAT PIM Passive

Setting up the Service interface loopback1

```
URL:
{{ip}}/api/mo/sys/mrib/inst/dom-default/sr.json
Payload:
{ "mribServiceReflect": {
  "attributes": {"status": "" },
  "children": [
    {
      "mribSrcIntf": {
        "attributes": {
          "srcIntf": "lo1",
          "status": ""
        }
      }
    }
  ]
}
```

Setting up the NAT mode to Egress

```
URL:
{{ip}}/api/mo/sys/mrib/inst/dom-default/sr.json
Payload:
{"mribEgressMode": {"attributes": {"grpList": "225.0.0.0/8"}}
```

Setting up the mapping interface

```
URL:
{{ip}}/api/mo/sys/mca/config/natsr/mappings.json
Payload:
{"mcaNatMapDefaultSif": {"attributes": {"domName": "default", "maxEnatReplications": "40",
  "siIfName": "eth1/2", "status": "" }}}}
```

Setting up the SR rule:

```
URL:
{{ip}}/api/mo/sys/mrib/inst/dom-default/sr/rule.json
```

```

Payload:
{"mribSrRule": {"attributes": {"status": ""},
"children": [{"mribRule": {"attributes": {"postTransGrp": "226.1.1.1", "postTransSrc":
"57.1.1.2", "preTransGrp": "225.1.1.1", "preTransSrc": "47.1.1.2", "grpMasklen": 32,
"srcMasklen": 32, "udpsrcPort": "10003", "udpDestPort": "20003", "staticOif": "eth1/29/1"}}}
]
} }

```

Pre-NAT flow

```

URL:
{{ip}}/api/mo/sys/nbm/conf/flows.json
Payload:
{"nbmFlows": {"children": [{"nbmConfFlowsDom": {"attributes": {"name": "default", "status":
""},
"children": [ {"nbmConfFlow": { "attributes": {"group": "225.1.1.1", "source": "47.1.1.2",
"ingressIf": "eth1/3" "policer": "ENABLED", "bwKbps": "1000" "status": "" } } },
] } } ] } }

```

Post-NAT Flow

```

URL:
{{ip}}/api/mo/sys/nbm/conf/flows.json
Payload:
{"nbmFlows": {"children": [{"nbmConfFlowsDom": {"attributes": {"name": "default"},
"children": [ {"nbmConfFlow": {"attributes": {"group": "226.1.1.1", "source": "57.1.1.1",
"ingressIf": "loopback1", "bwKbps": 10000, "policer": "ENABLED", "status": "" } },
"children": [{"nbmConfFlowIf": {"attributes": {"id": "eth1/29/1", "isLhr": "YES", "status":
"" } } } } } ] } } }

```

Multicast to Unicast NAT

Multicast to unicast NAT is used for hosting content to public cloud. The translation is required as the cloud may not support multicast. After translation, the Unicast packet gets routed as per unicast forwarding logic.

A similar use case is seen when connecting to different sites. If the core does not support multicast end to end, then the content is delivered as unicast to the different sites. The Border box translates multicast to unicast and delivers to different sites for consumption.

For MU NAT, PMN will continue perform bandwidth management for pre-translated multicast flows. For the translated unicast flow, the outgoing interface will need to have unicast bandwidth reservation so that the translated unicast traffic will be sent without any disruption. PMN will also publish the Flow operational MO to indicate the NAT relationship. Since, there are three re-circulations that occur internally for every unicast translation, one must make sure that only one third of the recirculation port bandwidth is assumed. In case of any congestion on the service-reflect map interface used for re-circulation, PMN does not publish a Fault MO.

In PIM Passive mode, Controller will perform Bandwidth management and call Rest APIs to provision the pre-translated flow. PMN will publish the flow operational MO to indicate the NAT relationship.

Examples for MU NAT PIM Passive

The following are the MUNAT Rest API calls and Payload information:

Configure Re-circ Interfaces

```

url: 172.28.249.173/api/mo/sys/mca/config/natsr/mappings.json?rsp-subtree=full
Payload:
{

```

```
"mcaNatMapDestPrefixSif": {
  "attributes": {
    "destPrefix": "112.10.3.0/24",
    "domName": "default",
    "maxEnatReplications": "40",
    "siIfName": "eth1/15",
    "status": ""
  }
}
```

Service Reflect Rules

```
url: <ip_switch>/api/mo/sys/mrib/inst/dom-default/sr/rule.json?rsp-subtree=full
Payload:
```

```
{
  "mribRule": {
    "attributes": {
      "grpMasklen": "32",
      "postTransGrp": "112.3.3.51",
      "postTransSrc": "11.1.1.3",
      "preTransGrp": "225.10.1.50",
      "preTransSrc": "112.3.1.2",
      "srcMasklen": "32",
      "staticOif": "unspecified",
      "status": "",
      "udpDestPort": "0",
      "udpsrcPort": "0"
    }
  }
}
```

NBM Flows

```
url: <ip_switch>/api/mo/sys/nbm/show/flows/dom-default.json?rsp-subtree=full
Payload:
```

```
{
  "nbmConfFlow": {
    "attributes": {
      "bwKbps": "50000",
      "group": "225.1.1.1",
      "ingressIf": "eth1/2",
      "policer": "ENABLED",
      "source": "112.3.1.2",
      "status": ""
    }
  }
}
```

Unicast to Multicast NAT

Unicast to Multicast NAT works in ingress translation mode. The multicast translated packet can be egress translated back to multicast. The destination address of the unicast packet should match the NAT source loopback interface secondary IP address.

The Unicast to Multicast NAT supports only 1:1 translation. If 1 to many translations is required, then you need to configure a 1:1 Unicast to Multicast NAT, and then configure 1 to many Multicast-to-Multicast NAT translations.

For Unicast to Multicast NAT, you must configure unicast bandwidth reservation on the port where the pre-translated unicast traffic arrives. This enables the multicast traffic on that port to not to consume all the

port bandwidth. Using the bandwidth derived from the flow policy of the post-translated multicast group, PMN installs policer on all the slices to police unicast flow. Since there is one re-circulation for every multicast translation, the recirculation port bandwidth must be same as the incoming port bandwidth.

PMN publishes the flow operations MO to indicate the NAT relationship. PMN does not publish a fault MO if there is a congestion on the service-reflect map interface that is used for re-circulation.



Note Flow priority to the subsequent Multicast to Multicast Translation flow cannot be assigned. This flow priority has to be set for Unicast to Multicast translation flow (parent flow).

Examples for Unicast to Multicast NAT PIM Active

The following are the examples for the Unicast to Multicast NAT in PIM Active mode:

UMNAT Flow

```
ip service-reflect destination 10.34.202.11 to 234.34.203.11 mask-len 32 source 10.30.17.11
to 10.34.201.1 mask-len 32
```

other supporting config needed for above flow stitching are:
 multicast service-reflect dest-prefix 234.34.203.0/24 map interface Ethernet1/6

```
NBM flow-policy config:
nbm flow-policy
policy umnat
  bandwidth 15000 kbps
  ip group-range 234.34.202.1 to 234.34.202.255
  ip group-range 234.34.203.1 to 234.34.203.255
```

Chained MMNAT Flow

```
ip service-reflect destination 234.34.203.11 to 234.34.253.11 mask-len 32 source 10.34.201.1
to 10.34.202.111 mask-len 32 to-udp-src-port 25010 to-udp-dest-port 25310 static-oif
Ethernet1/56
```

```
ip service-reflect destination 234.34.203.11 to 234.34.253.11 mask-len 32 source 10.34.201.1
to 10.34.202.111 mask-len 32 to-udp-src-port 25010 to-udp-dest-port 25510 static-oif
Ethernet1/55
```

other supporting config needed for above flow stitching are:

```
multicast service-reflect interface Ethernet1/56 map interface Ethernet1/3
multicast service-reflect interface all map interface Ethernet1/4
```

```
NBM flow-policy config:
nbm flow-policy
policy umnmat1
  bandwidth 16000 kbps
  ip group-range 234.34.253.10 to 234.34.253.100
  priority critical
  ip group-range 234.34.253.101 to 234.34.253.255
switch# show ip mr sr umnat 10.30.17.11 10.34.202.11
IP Multicast Routing Table for VRF "default"
```

```
(10.30.17.11/32, 10.34.202.11/32)
```

```
Translation:
```

```
SR: (10.34.201.1/32, 234.34.203.11/32) udp src: 0, udp dst : 0
Outgoing interface list: (count: 3)
  Ethernet1/56, uptime: 02:13:44, igmp
  Ethernet1/55, uptime: 02:13:44, igmp
  Ethernet1/60, uptime: 02:13:51, static
```

```

Chained translations:
  SR: (10.34.202.111, 234.34.253.11) udp src: 25010 udp dst: 25310 OIF: Ethernet1/56
  SR: (10.34.202.111, 234.34.253.11) udp src: 25010 udp dst: 25510 OIF: Ethernet1/55

switch#

switch# show forwarding distribution multicast route group 234.34.203.11 source 10.34.201.1

(10.34.201.1/32, 234.34.203.11/32), RPF Interface: Ethernet1/6.100, flags: EPrePstUM
  Upstream Nbr: 10.34.201.1, Stats State: NA
  Received Packets: 16964898 Bytes: 23784786996
  Number of Outgoing Interfaces: 6
  Outgoing Interface List Index: 1609
    Ethernet1/55
    Ethernet1/56
    Ethernet1/60
    Null0
      Type: NAT_EGR_RW
      Source IF: Ethernet1/6.100
      RW Group IP: 234.34.203.11
      RW Source IP: 10.34.201.1
      RW source L4 port: 0
      RW dest L4 port: 0
      Original Group IP: 10.34.202.11
      Original Source IP: 10.30.17.11

    Ethernet1/56
      Type: NAT_EGR_RW
      Source IF: Ethernet1/3.1
      RW Group IP: 234.34.253.11
      RW Source IP: 10.34.202.111
      RW source L4 port: 25010
      RW dest L4 port: 25310
      Original Group IP: 234.34.203.11
      Original Source IP: 10.34.201.1

    Ethernet1/55
      Type: NAT_EGR_RW
      Source IF: Ethernet1/4.1
      RW Group IP: 234.34.253.11
      RW Source IP: 10.34.202.111
      RW source L4 port: 25010
      RW dest L4 port: 25510
      Original Group IP: 234.34.203.11
      Original Source IP: 10.34.201.1

switch#

switch# show forwarding multicast route group 234.34.203.11 source 10.34.201.1

slot 1
=====

(10.34.201.1/32, 234.34.203.11/32), RPF Interface: Ethernet1/6.100, flags:
  Received Packets: 17115724 Bytes: 23996245048
  Outgoing Interface List Index: 1609
  Number of next hops: 4
  oiflist flags: 16809984

  Outgoing Interface List Index: 0x649
    Ethernet1/55
    Ethernet1/56

```

```

Ethernet1/60
  Null0
  Encap 216 (10.30.17.11, 10.34.202.11 -> 10.34.201.1, 234.34.203.11) L4(0,0)
SrcIf(Ethernet1/6.100) Flags(0x0)
  Ethernet1/56
  Encap 1002 (10.34.201.1, 234.34.203.11 -> 10.34.202.111, 234.34.253.11) L4(25010,25310)
SrcIf(Ethernet1/3.1) Flags(0x0)
  Ethernet1/55
  Encap 1003 (10.34.201.1, 234.34.203.11 -> 10.34.202.111, 234.34.253.11) L4(25010,25510)
SrcIf(Ethernet1/4.1) Flags(0x0)s#

```

```

switch# show forwarding multicast-sr internal-db
  Encap 216 (10.30.17.11, 10.34.202.11 -> 10.34.201.1, 234.34.203.11) L4(0,0)
SrcIf(Ethernet1/6.100) Flags(0x0)
  Encap 1002 (10.34.201.1, 234.34.203.11 -> 10.34.202.111, 234.34.253.11) L4(25010,25310)
SrcIf(Ethernet1/3.1) Flags(0x0)
  Encap 1003 (10.34.201.1, 234.34.203.11 -> 10.34.202.111, 234.34.253.11) L4(25010,25510)
SrcIf(Ethernet1/4.1) Flags(0x0)

```

NBM Show commands:

```
switch# show nbm flows group 234.34.203.11 source 10.34.201.1 detail
```

```
-----
NBM Flows for VRF 'default'
-----
```

Active Source-Group-Based Flow(s) for Source 10.34.201.1 Group 234.34.203.11 :

Mcast-Group	Src-IP	Uptime	Src-Intf	Nbr-Device	LID	Profile						
Status	Num Rx	Bw Mbps	CFG Bw	Slot Unit	Slice	DSCP	QOS	Policed	FHR	Priority	Policy-name	
Rcvr-Num	Rcvr-slot	Unit	Num-Rcvrs	Rcvr-ifidx	IOD	Rcvr-Intf	Nbr-Device					
234.34.203.11	10.34.201.1	02:21:05	Lo34	not-available	0	N/A						
ACTIVE	3	15.000	15.000	17	0	0	0	7	Yes	Yes	LOW	umnat
	1	1	0	3	0x1a006e00	64	Eth1/56	not-available				
	2	1	0	3	0x1a006c00	63	Eth1/55	not-available				
	3	1	0	3	0x1a007600	68	Eth1/60					

LEAF34-PMN-SOLN-SOUTHLAKE

switch#

```
switch# show nbm flows statis group 234.34.203.11 source 10.34.201.1
```

```
-----
NBM Flow Statistics for VRF 'default'
-----
```

Source-Group-Based Flow Statistics for Source 10.34.201.1 Group 234.34.203.11 :

Mcast-Group	Src-IP	Uptime	Src-Intf	Packets	Bytes
Allow-Bytes	Drop-Bytes				
234.34.203.11	10.34.201.1	02:21:27	Lo34	8413701	11779181400
11778445000	0				

switch#

NBM Oper MO:

```
{
```

```

"nbmNbmUmFlow": {
  "attributes": {
    "bucket": "3",
    "destination": "10.34.202.11",
    "dn": "sys/nbm/show/flows/dom-default/ums-[10.30.17.11]-umd-[10.34.202.11]",
    "modTs": "2021-11-30T11:34:55.213+00:00",
    "source": "10.30.17.11",
    "tStamp": "1638300895054"
  }
}
}
{
  "nbmNbmFlow": {
    "attributes": {
      "bucket": "1",
      "bwKbps": "15000",
      "dn": "sys/nbm/show/flows/dom-default/s-[10.34.201.1]-g-[234.34.203.11]",
      "dscp": "0",
      "egressIfCount": "3",
      "flowPol": "umnat",
      "group": "234.34.203.11",
      "ingressIf": "335544354",
      "ingressIfName": "loopback34",
      "isFhr": "YES",
      "modTs": "2021-11-30T11:35:23.384+00:00",
      "policed": "YES",
      "priority": "LOW",
      "qid": "7",
      "source": "10.34.201.1",
      "tStamp": "1638300923224"
    },
    "children": [
      {
        "nbmOifList": {
          "attributes": {
            "dn":
"sys/nbm/show/flows/dom-default/s-[10.34.201.1]-g-[234.34.203.11]/oif-436237824",
            "modTs": "2021-11-30T11:35:35.387+00:00",
            "oif": "436237824",
            "oifName": "Ethernet1/60",
            "oifTstamp": "1638300935386",
            "origin": "PROTOCOL",
            "reporterIP": "10.34.60.1"
          }
        }
      },
      {
        "nbmOifList": {
          "attributes": {
            "dn":
"sys/nbm/show/flows/dom-default/s-[10.34.201.1]-g-[234.34.203.11]/oif-436235264",
            "modTs": "2021-11-30T11:35:42.436+00:00",
            "oif": "436235264",
            "oifName": "Ethernet1/55",
            "oifTstamp": "1638300942436",
            "origin": "PROTOCOL",
            "reporterIP": "10.34.55.11"
          }
        }
      },
      {
        "nbmOifList": {
          "attributes": {

```



```

        "dn":
"sys/nbm/show/flows/dom-default/s-[10.34.201.1]-g-[234.34.203.11]/oif-436235776",
    "modTs": "2021-11-30T11:35:42.437+00:00",
    "oif": "436235776",
    "oifName": "Ethernet1/56",
    "oifTstamp": "1638300942437",
    "origin": "PROTOCOL",
    "reporterIP": "10.34.56.11"
    }
  },
  {
    "nbmUmIngNat": {
      "attributes": {
        "dn":
"sys/nbm/show/flows/dm-default/s-[10.34.201.1]-g-[234.34.203.11]/uming-pres-[10.30.17.11]-pred-[10.34.202.11]-postsp-[0]-postdp-[0]",

        "modTs": "2021-11-30T11:34:55.213+00:00",
        "postDPort": "0",
        "postSPort": "0",
        "preDestination": "10.34.202.11",
        "preSource": "10.30.17.11"
      }
    }
  }
]
}

```

