Deterministic Large Scale NAT44

In This Chapter

This section provides information about deterministic large scale NAT44 configurations.

Topics in this section include:

Applicability

This example is applicable to 7750 SR systems and 7450 ESS systems in mixed mode equipped with an MS-ISA on an IOM3-XP/XP-b using chassis mode B, C or D.

The configuration was tested on release 11.0R3.

Overview

Deterministic Network Address Translation (NAT) is a mode of operation where mappings between the NAT subscriber and the outside IP address and port range are allocated at the time of configuration.

In deterministic NAT for Large Scale NAT IPv4-to-IPv4 (LSN44) subscribers, each LSN44 subscriber is permanently mapped to an outside IP address and a dedicated (deterministic) port-block based on a specific algorithm.

Logging is not needed in this case as the reverse mapping can be obtained using the reverse of the above algorithm.

A deterministic LSN44 subscriber can have only one deterministic port-block that can (optionally) be extended by one or multiple dynamic port-blocks in case all ports in deterministic port-block are exhausted.

In case an LSN44 subscriber has been assigned both deterministic and dynamic port blocks, logging for the dynamic port-block allocation/de-allocation is required.

A scalable logging solution for dynamic port-blocks is achievable using RADIUS or IPFIX.

Logging for dynamic port-blocks is out of the scope of this example.

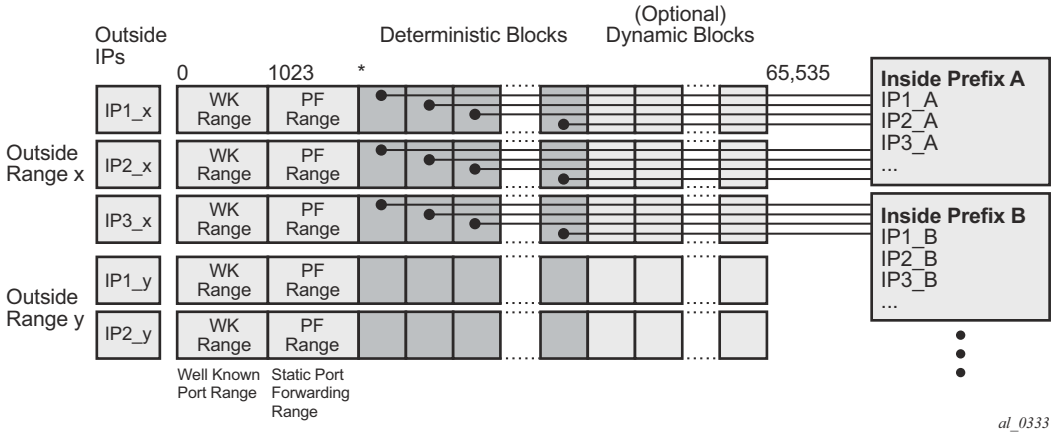

Figure 215: Deterministic NAT Mapping

Algorithm

The deterministic NAT algorithm makes a predictable mapping between the (inside IP, routing instance) and the (outside IP, routing instance, deterministic port block).

The algorithm is revertive, meaning that a given (outside IP, routing instance, deterministic port block) will derive a given (Inside IP, Routing Instance).

The algorithm is loosely based on the draft RFC draft-donley-behave-deterministic-cgn-00.txt, which allows for the dynamic expansion of the port-blocks once the ports in the original deterministic port-block are exhausted.

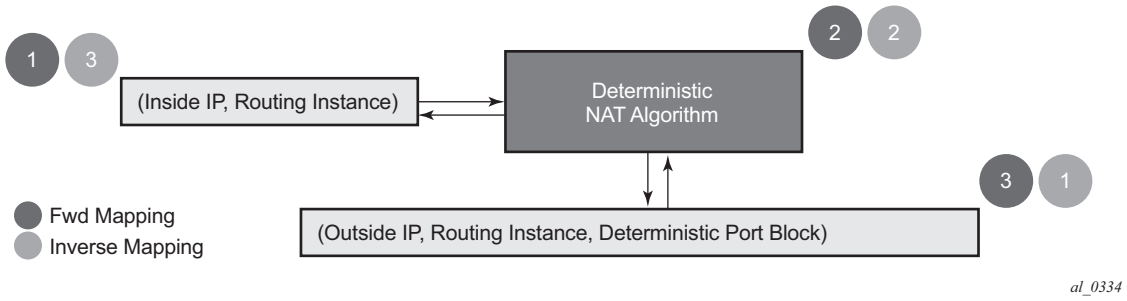

Figure 216: Deterministic NAT Algorithm

Deterministic Mapping

Any inside prefix in any routing instance can be mapped to any pool in any routing instance.

In deterministic NAT, prefixes from multiple routing instances can be mapped to the same outside pool, also prefixes from a single inside routing instance can be selectively mapped to different outside pools.

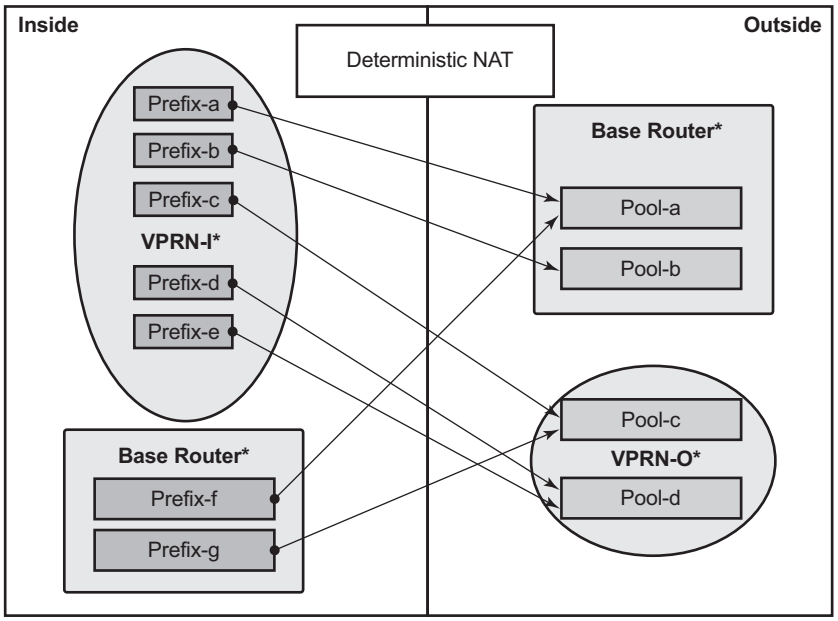

Figure 217: Deterministic Mapping: Inside -> Outside Routing Instances

Mapping Rules

A deterministic LSN44 subscriber is mapped to only one deterministic block which can further be extended to multiple dynamic blocks if ports within the deterministic block are exhausted.

The subscriber-limit is the number of subscribers that can be deterministically mapped to an outside IP address (i.e. compression ratio) and MUST be a power of 2.

The total number of deterministic ports (DetP) per outside IP address is determined by the number of subscribers per outside IP address and the number of deterministic ports per subscriber.

The remaining ports (DynP) beyond the deterministic port range up to 65535 will be dedicated for dynamic use when a deterministic block is exhausted.

Every host using an inside prefix is guaranteed one dedicate block in the deterministic port ranges.

If the inside prefix length is m < 32-n, where 2^n=subscriber-limit, then the prefix must be broken into pieces so that all hosts (subscriber-limit) in each piece maps exactly to one outside IP address.

- For example, if there is an inside prefix 192.168.0.0/23, here m=23; and the subscriber-limit is also set to 256, then n=8. This results in 23 < 24 (32-8) and so this inside prefix has to be broken into 2 pieces, in other words this inside prefix will fit into 2 outside IP addresses, each of 256 port-blocks.

In case that the prefix length is m ≥ 32-n, where 2^n=subscriber-limit, then all hosts from the configured prefix are mapped to the same outside IP.

- For example, if there is an inside prefix 192.168.1.0/25, here m=25 and there can be at most 128 hosts, and the subscriber-limit is set to 256, then n=8. This results in 25 > 24 (32-8), so definitely 128 hosts can fit in one outside IP as there are 256 available port-blocks, in other words this inside prefix will fit into one outside IP where 128 blocks have been used out of the 256 port-blocks available, and the rest (256-128) are wasted.

Overbooking of the outside address pool is not supported in deterministic NAT.

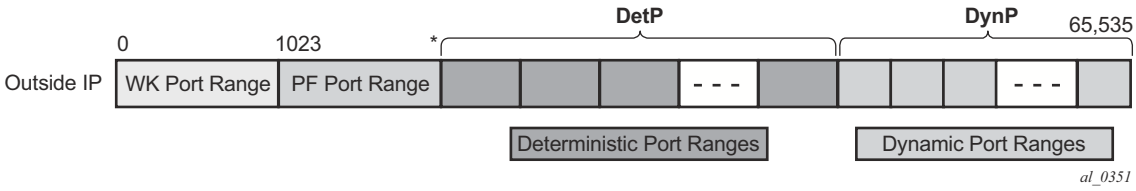

Figure 218: Deterministic Mapping: Outside IP Port-Blocks/Ranges

Configuration

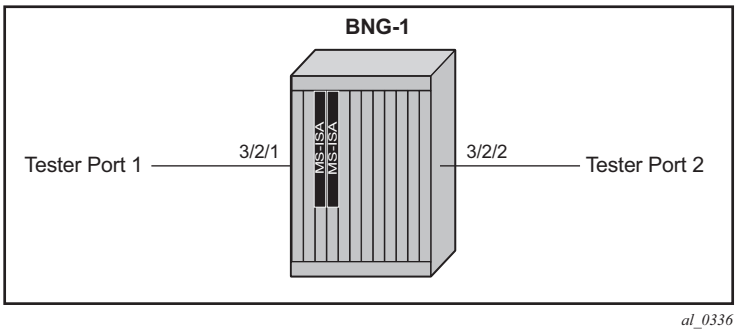

Figure 219: Test Topology

Configuration Pre-Requisites

Chassis mode, card, and MDA configuration.

```
configure
    system
        chassis-mode d
exit all

configure
    card 3
        card-type iom3-xp
        mda 1
            mda-type isa-bb
            no shutdown
        exit
        no shutdown
    exit
    card 4
        card-type iom3-xp
        mda 1
            mda-type isa-bb
            no shutdown
        exit
        no shutdown
exit all
```

Note: Private address ranges are used in outside pools within this example but normally public address ranges would be used.

Configure a NAT group

Create the nat-group, and add the MS-ISAs created above to the nat-group; up to 10 MS-ISAs of type isa-bb can be configured under the nat-group.

```
configure isa
        nat-group 1 create
            active-mda-limit 2
            mda 3/1
            mda 4/1
            no shutdown
exit all
```

Configuration Commands

A NAT **outside pool** is configured using the following command:

```
configure  {router | service vprn <service-id>}
  nat
    outside
      pool <nat-pool-name> [nat-group <nat-group-id> type <pool-type> create]
        port-reservation {blocks <num-blocks> | ports <num-ports>}
        port-forwarding-range <range-end>
        subscriber-limit <subscriber-limit>
        deterministic
        port-reservation <num-ports>
        exit
          address-range <start-ip-address> <end-ip-address> create
        exit
      exit
    exit
  exit
```

where:

nat-pool-name — Specifies the name of the NAT pool up to 32 characters max.

nat-group-id — Specifies the NAT group ID. The values are 1 — 4.

pool-type — Species the pool type (**large-scale**).

num-blocks — Specifies the number of port-blocks per IP address. Setting num-blocks to one (1) for large scale NAT will enable 1:1 NAT for IP addresses in this pool The values are 1 — 64512

num-ports — Specifies the number of ports per block. The values are 1 — 32256

range-end — Specifies the end of the port range available for port forwarding. The values are 1023 — 65535

subscriber-limit Specifies the maximum number of subscribers per IP address.
A power of 2 (2^n) number for deterministic NAT
[1,2,4,8,16,32,64,128,256,512,1024,2048, 4096, 8192,16348, 32768]
1..65535 for non-deterministic NAT
default: 65535 for non-deterministic

num-ports — Specifies the number of ports in a deterministic port block that is allocated and dedicated to a single subscribers during the configuration phase. The values are 1..65535

start-ip-address — Specifies the beginning IP address in a.b.c.d form.

end-ip-address — Specifies the ending IP address in a.b.c.d. form.

Notes:

→ When the subscriber-limit equals 1, each subscriber is mapped to a single outside IP address, though the NAPT (port translation) function is still performed.

→ 1:1 NAT mode in combination with deterministic NAT is not supported.

A NAT **policy** is configured using the following command:

```
configure service nat
 nat-policy <nat-policy-name> [create]
    block-limit <[1..40]>
    pool <nat-pool-name> {router <router-instance> | service-name <service-name>}
 exit
```

where:

nat-policy-name — Specifies the NAT policy name up to 32 characters max.

block-limit —The max number of deterministic plus dynamic port blocks that can be assigned to a single inside IP address. In other words, the maximum number of dynamic port blocks that can be assigned to an inside IP address when the deterministic port block is exhausted equals (block-limit - 1).

nat-pool-name — Specifies the NAT pool name up to 32 characters max.

router-instance — Specifies the router instance the pool belongs to, either by router name or service ID. : *<router-name>*|*<service-id>*
The router name values are **Base** or *service-id* [1..2147483647]

service-name — Specifies the name of the service up to 64 characters max.

A NAT **inside prefix** is configured using the following command:

```
configure [router| service vprn <service-id>]
  nat
    inside
      deterministic
        classic-lsn-max-subscriber-limit <max>
        prefix <ip-prefix/length> subscriber-type <nat-sub-type> nat-policy <nat-policy-
name> create
          map start <lsn-sub-address> end <lsn-sub-address> to <outside-ip-address>
          no shutdown
          exit
        exit
      exit
    exit
```

where:

max — The power of 2 (2^n) number that must match the largest subscriber limit number in a deterministic pool referenced from this inside routing instance. The range for this command is the same as the subscriber-limit command under the pool hierarchy. The values are 1,2,4,8 — 2768

ip-prefix/length — A prefix on the inside encompassing subscribers that will be deterministically mapped to an outside IP address and port block in the corresponding pool.

<*ip-prefix/ip-pref**> <ipv4-prefix>/<ipv4-prefix-length> | <ipv6-prefix>/<ipv6-prefix-length>	
<*ipv4-prefix*>	a.b.c.d (host bits must be 0)
<*ipv4-prefix-length*>	[0..32]
<*ipv6-prefix*>	x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x - [0..FFFF]H d - [0..255]D
<*ipv6-prefix-length*>	[0..128]
<*nat-sub-type*>:	classic-lsn-sub
<*nat-policy-name*>	Specifies a NAT policy name up to 32 characters in length.

- classic-lsn-max-subscriber-limit:
  - → Should be greater than the largest subscriber-limit of all pools referenced by the NAT policies within the corresponding inside routing instance.
  - → Must be configured before any inside prefix configuration.
  - → Must be 2^n and affects the ingress hashing of deterministic subscribers and also non-deterministic subscribers in case both are configured under the same inside router instance.

Three cases are now configured to demonstrate the use of deterministic and dynamic port-block usage:

- Case 1 on page 1440: Mapping multiple prefixes from the same VRF into the same outside pool.
- Case 2 on page 1450: Mapping multiple prefixes from the same VRF into different outside pools.
- Case 3 on page 1457: Mapping overlapping prefixes from different VRFs into the same outside pool.

In each case all of the traffic is NATed.

Case 1

Configured with:

- Mapping multiple prefixes from the same VRF into the same outside pool.
- NAT all traffic.

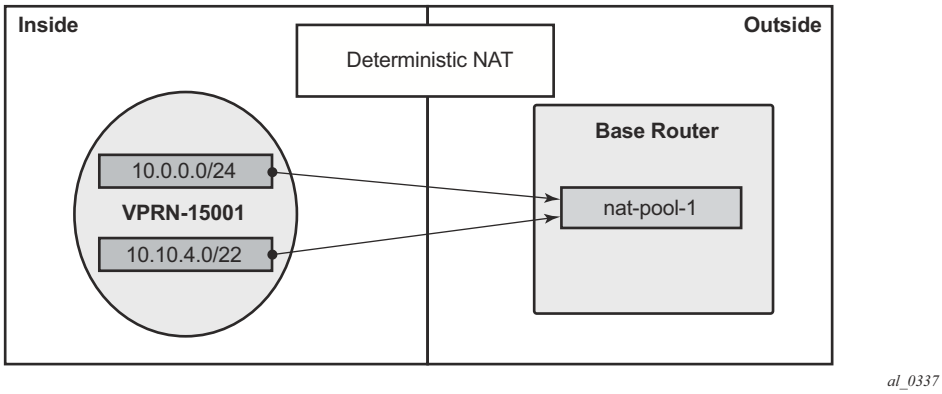

Figure 220: Case 1

The NAT **outside pool** is configured as follows:

```
configure router nat
         outside
               pool "nat-pool-1" nat-group 1 type large-scale create
                   port-reservation ports 180
                   port-forwarding-range 4023
                   subscriber-limit 128
                   deterministic
                       port-reservation 300
                   exit
                   address-range 192.168.0.1 192.168.0.100 create
                   exit
                   no shutdown
exit all
```

The NAT **policy** is configured as follows:

```
configure service nat
          nat-policy "nat-policy-1" create
                block-limit 4
                pool "nat-pool-1" router Base
exit all
```

The NAT **inside prefix** is configured as follows:

```
configure service vprn 15001 nat
            inside
            destination-prefix 0.0.0.0/0
               deterministic
                classic-lsn-max-subscriber-limit 256
                prefix 10.0.0.0/24 subscriber-type classic-lsn-sub
                                     nat-policy "nat-policy-1" create
                     map start 10.0.0.0 end 10.0.0.255 to 192.168.0.1
                     no shutdown
                  exit
                prefix 10.10.4.0/22 subscriber-type classic-lsn-sub
                                      nat-policy "nat-policy-1" create
                   map start 10.10.4.0 end 10.10.7.255 to 192.168.0.3
                   no shutdown
exit all
```

Notes:

- The **classic-lsn-max-subscriber-limit** value should be greater or equal to the largest subscriber-limit of all pools referenced by NAT policies within the corresponding inside routing instance. It must be 2^n and affects ingress hashing of deterministic subscribers.
- **map** statements are automatically created when the prefix is created and it is **no shutdown**.

Show commands

- Prefix 10.0.0.0/24
  - → Since the subscriber-limit is 128 in this case, the 10.0.0.0/24 prefix will be broken into two smaller prefixes each of /25, each will be mapped into a specific outside IP address.
  - → To show Large Scale NAT (LSN) hosts for inside routing instance 15001 for the first /25 prefix and the mapping to which outside IP, the following command can be used:

```
show router 15001 nat lsn-hosts inside-ip-prefix 10.0.0.0/25
===============================================================================
Large-Scale NAT hosts for router 15001
===============================================================================
Inside IP       Out-Router     Outside IP
-------------------------------------------------------------------------------
10.0.0.0        Base           192.168.0.1
10.0.0.1        Base           192.168.0.1
<snip>
10.0.0.127      Base           192.168.0.1
-------------------------------------------------------------------------------
No. of hosts: 128
===============================================================================
```

To show LSN hosts for the inside routing instance 15001 for the second /25 prefix and the mapping to which outside IP, the following command can be used:

```
show router 15001 nat lsn-hosts inside-ip-prefix 10.0.0.128/25
===============================================================================
Large-Scale NAT hosts for router 15001
===============================================================================
Inside IP       Out-Router     Outside IP
-------------------------------------------------------------------------------
10.0.0.128      Base           192.168.0.2
10.0.0.129      Base           192.168.0.2
<snip>
10.0.0.255      Base           192.168.0.2
-------------------------------------------------------------------------------
No. of hosts: 128
===============================================================================
```

To show LSN blocks on the outside routing instance “Base” for the first inside IP within 10.0.0.0/24 prefix, the following command can be used:

```
show router nat lsn-blocks inside-ip 10.0.0.0
===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.1 [4024..4323]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2013/07/21 09:30:20
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.0
-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

To show LSN blocks on the outside routing instance “Base” for the last inside IP within 10.0.0.0/24 prefix, the following command can be used:

```
show router nat lsn-blocks inside-ip 10.0.0.255
===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.2 [42124..42423]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2013/07/21 09:30:20
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.255
-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

→ Prefix 10.10.4.0/22

– Since the subscriber-limit is 128 in this case, the 10.10.4.0/22 prefix will be broken into 8 smaller prefixes each of /25, each will be mapped into a specific outside IP address.

– To show LSN hosts for the inside routing instance 15001 for the first /25 prefix and the mapping to which outside IP, the following command can be used:

```
show router 15001 nat lsn-hosts inside-ip-prefix 10.10.4.0/25
===============================================================================
Large-Scale NAT hosts for router 15001
===============================================================================
Inside IP       Out-Router     Outside IP
-------------------------------------------------------------------------------
10.10.4.0       Base           192.168.0.3
10.10.4.1       Base           192.168.0.3
<snip>
10.10.4.127     Base           192.168.0.3
-------------------------------------------------------------------------------
No. of hosts: 128
===============================================================================
```

To show LSN hosts for the inside routing instance 15001 for the last /25 prefix and the mapping to which outside IP, the following command can be used:

```
show router 15001 nat lsn-hosts inside-ip-prefix 10.10.7.128/25
===============================================================================
Large-Scale NAT hosts for router 15001
===============================================================================
Inside IP       Out-Router     Outside IP
-------------------------------------------------------------------------------
10.10.7.128     Base           192.168.0.10
10.10.7.129     Base           192.168.0.10
<snip>
10.10.7.255     Base           192.168.0.10
-------------------------------------------------------------------------------
No. of hosts: 128
===============================================================================
```

To show LSN blocks on the outside routing instance **Base** for the first inside IP within 10.10.4.0/24 prefix, the following command can be used:

```
show router nat lsn-blocks inside-ip 10.10.4.0
===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.3 [4024..4323]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
```

```
Started                               : 2013/07/21 09:30:20
Inside router                         : vprn15001
Inside IP address                     : 10.10.4.0

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

To show LSN blocks on the outside routing instance **Base** for the last inside IP within 10.10.4.0/24 prefix, the following command can be used:

```
show router nat lsn-blocks inside-ip 10.10.7.255
===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.10 [42124..42423]
Pool                                  : nat-pool-1
Policy                                : nat-policy-1
Started                               : 2013/07/21 09:30:26
Inside router                         : vprn15001
Inside IP address                     : 10.10.7.255
-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

Mapping results

According to this configuration, each inside IP address has one deterministic block of 300 ports and can have up to three dynamic blocks (block-limit = 4) each of 180 ports, allowing a maximum of 300+3*180 = 840 flows.

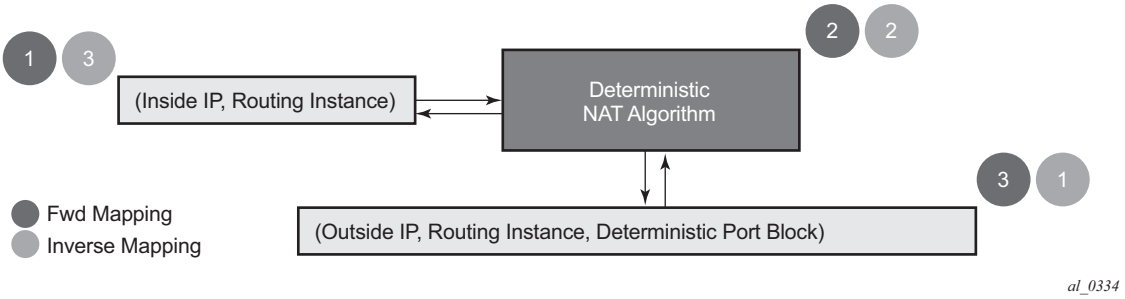

Figure 221: Case 1 Results

Sending Flows

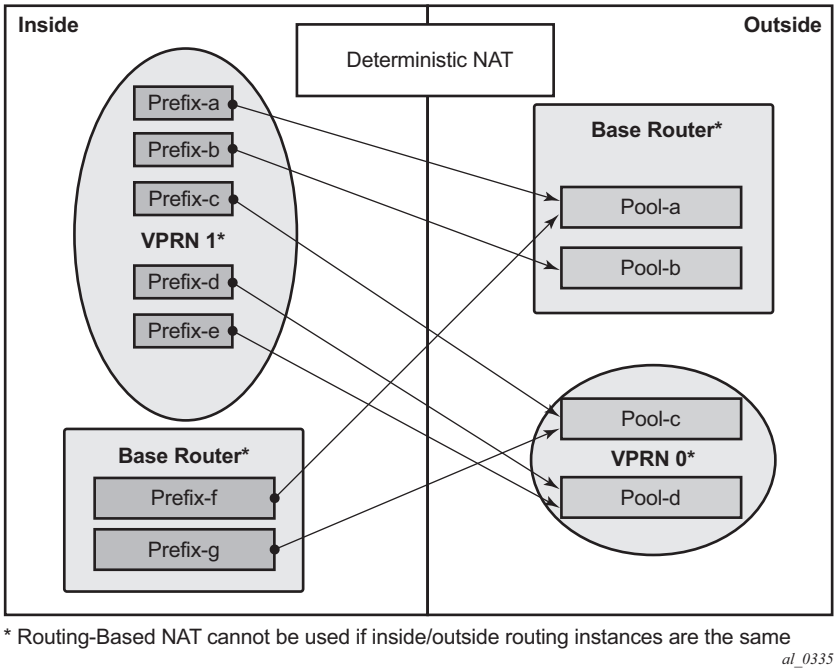

Figure 222: Case 1 Flows

For the inside IP 10.0.0.1 several UDP flows will be sent and both the deterministic and dynamic blocks mappings will be verified.

When sending UDP flows ≤ 300 Flows

- All flows are mapped to a single deterministic block since the number of ports in a deterministic block is 300.
- There is no logging; since no dynamic blocks are used, and only the deterministic block is used.
- To show LSN blocks on the outside routing instance **Base** and the outside ports allocated for the inside IP 10.0.0.1, the following command can be used:

```
show router nat lsn-blocks inside-ip 10.0.0.1
===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.1 [4324..4623]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2013/07/21 09:30:20
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1
-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

When increasing number of flows such that : 301 ≤ Flows ≤ 480

- In addition to the deterministic block (300 ports), there will be an extension by 1 dynamic block of 180 ports (port-reservation=180).
- Logging occurs for the dynamic port-block.
- To show LSN blocks on the outside routing instance **Base** and the outside ports allocated for the inside IP 10.0.0.1, the following command can be used:

```
show router nat lsn-blocks inside-ip 10.0.0.1
===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.1 [4324..4623]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2013/07/21 09:30:20
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [42424..42603]
Pool                                : nat-pool-1
```

```
Policy                              : nat-policy-1
Started                             : 2013/07/21 09:33:21
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1
-------------------------------------------------------------------------------
Number of blocks: 2
===============================================================================
```

Logging is verified using Log 99 (in case event-control **nat** events are generated) which shows the mapping details to the new dynamic block as follows:

```
137 2013/07/21 09:33:21.90 UTC MINOR: NAT #2012 Base NAT
"{1} Map  192.168.0.1 [42424-42603] MDA 4/1 -- 276824065 classic-lsn-sub vprn15001
10.0.0.1 at 2013/07/21 09:33:21"
```

When increasing number of flows such that: 481 ≤ Flows ≤ 660

- In addition to the deterministic block (300 ports), there will be an extension by 2 dynamic blocks of 180 ports each.
- Logging occurs for the dynamic port-blocks.
- To show LSN blocks on the outside routing instance **Base** and the outside ports allocated for the inside IP 10.0.0.1, the following command is used:

```
show router nat lsn-blocks inside-ip 10.0.0.1
===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.1 [4324..4623]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2013/07/21 09:30:20
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [42424..42603]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2013/07/21 09:33:21
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [42604..42783]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2013/07/21 09:35:44
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1
-------------------------------------------------------------------------------
Number of blocks: 3
===============================================================================
```

Logging is verified using Log 99 (in case event-control **nat** events are generated) which shows the mapping details to the new dynamic block as follows:

```
138 2013/07/21 09:35:44.20 UTC MINOR: NAT #2012 Base NAT
"{2} Map  192.168.0.1 [42604-42783] MDA 4/1 -- 276824065 classic-lsn-sub vprn15001
10.0.0.1 at 2013/07/21 09:35:44"
```

When increasing number of flows such that :661 ≤ Flows ≤ 840

- In addition to the deterministic block (300 ports), there will be an extension by 3 dynamic blocks of 180 ports each.
- Logging occurs for the dynamic port-blocks.
- To show LSN blocks on the outside routing instance “Base” and the outside ports allocated for the inside IP 10.0.0.1, the following command can be used:

```
show router nat lsn-blocks inside-ip 10.0.0.1
===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.1 [4324..4623]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2013/07/21 09:30:20
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [42424..42603]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2013/07/21 09:33:21
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [42604..42783]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2013/07/21 09:35:44
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [42784..42963]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2013/07/21 09:37:08
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1
-------------------------------------------------------------------------------
Number of blocks: 4
===============================================================================
```

Logging is verified using Log 99 (in case event-control **nat** events are generated) which shows the mapping details to the new dynamic block as follows:

```
139 2013/07/21 09:37:08.10 UTC MINOR: NAT #2012 Base NAT
"{3} Map  192.168.0.1 [42784-42963] MDA 4/1 -- 276824065 classic-lsn-sub vprn15001
10.0.0.1 at 2013/07/21 09:37:08"
```

When increasing number of flows such that :Flows > 840

- No more extension by dynamic blocks (block-limit = 4) allowed.
- Any flows more than 840 will be dropped and the relevant NAT statistics incremented.
- To verify NAT statistics, firstly check the NAT group/member and MS-ISA associated with the outside IP 192.168.0.1/32:

```
show router route-table 192.168.0.1/32
===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
192.168.0.1/32                                Remote  NAT       00h07m50s  0
       NAT outside to mda 4/1                                        0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
```

To check which group/member does this MS-ISA belong to, the following command can be used:

```
show isa nat-group 1 members
===============================================================================
ISA Group 1 members
===============================================================================
Group Member     State          Mda  Addresses  Blocks     Se-% Hi Se-Prio
-------------------------------------------------------------------------------
1     1          active         3/1  4          1024       < 1  N  0
1     2          active         4/1  6          1536       < 1  N  0
-------------------------------------------------------------------------------
No. of members: 2
===============================================================================
```

To verify relevant statistics for this NAT group/member, the following command can be used:

```
show isa nat-group 1 member 2 statistics | match "no ip or port"
no ip or port                                        : 2135
```

Case 2

Configured with:

- Mapping multiple prefixes from the same VRF into different outside pools.
- NAT all traffic.

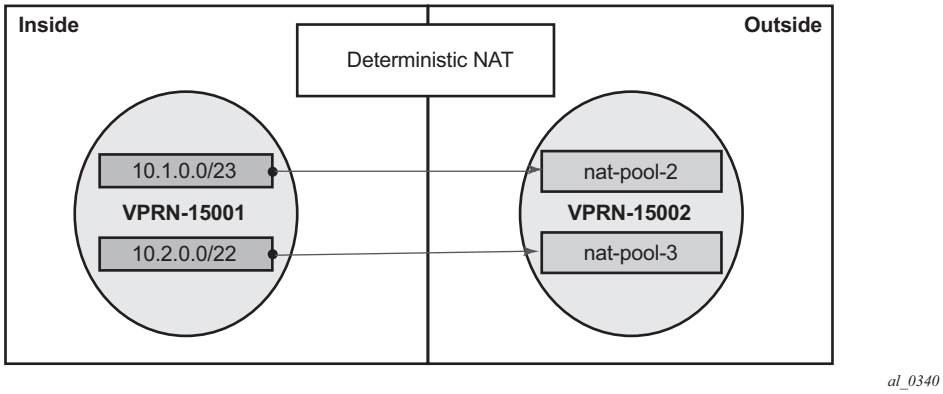

Figure 223: Case 2

The NAT **outside pools** are configured as follows:

```
configure service vprn 15002 nat
     outside
         pool "nat-pool-2" nat-group 1 type large-scale create
                  port-reservation ports 80
                  subscriber-limit 256
                  deterministic
                      port-reservation 180
                  exit
                  address-range 192.168.2.1 192.168.2.200 create
                  exit
                  no shutdown
           exit
         pool "nat-pool-3" nat-group 1 type large-scale create
           port-reservation ports 120
                  port-forwarding-range 4023
                  subscriber-limit 64
                  deterministic
                      port-reservation 840
                  exit
                  address-range 192.168.3.1 192.168.3.200 create
                  exit
                  no shutdown
exit all
```

The NAT **policies** are configured as follows:

```
configure service nat
            nat-policy "nat-policy-2" create
                block-limit 4
                pool "nat-pool-2" router 15002
            exit
            nat-policy "nat-policy-3" create
                block-limit 2
                pool "nat-pool-3" router 15002
            exit
         exit all
```

The NAT **inside prefix** is configured as follows:

```
configure service vprn 15001 nat
     inside
       destination-prefix 0.0.0.0/0
         deterministic
         classic-lsn-max-subscriber-limit 256
         prefix 10.1.0.0/23 subscriber-type classic-lsn-sub
                             nat-policy "nat-policy-2" create
              map start 10.1.0.0 end 10.1.1.255 to 192.168.2.1
              no shutdown
          exit
        prefix 10.2.0.0/22 subscriber-type classic-lsn-sub
                            nat-policy "nat-policy-3" create
              map start 10.2.0.0 end 10.2.3.255 to 192.168.3.1
              no shutdown
exit all
```

Show commands

- Prefix 10.1.0.0/23
  - → Since the subscriber-limit is 256 in this case, the 10.1.0.0/23 prefix will be broken into two smaller prefixes each of /24, each will be mapped into a specific outside IP address.
  - → To show Large Scale NAT (LSN) hosts for the inside routing instance 15001 for the first /24 prefix and the mapping to which outside IP, the following command can be used:

```
show router 15001 nat lsn-hosts inside-ip-prefix 10.1.0.0/24
===============================================================================
Large-Scale NAT hosts for router 15001
===============================================================================
Inside IP       Out-Router     Outside IP
-------------------------------------------------------------------------------
10.1.0.0        15002          192.168.2.1
```

```
10.1.0.1        15002          192.168.2.1
<snip>
10.1.0.255      15002          192.168.2.1
-------------------------------------------------------------------------------
No. of hosts: 256
===============================================================================
```

To show LSN hosts for the inside routing instance 15001 for the second /24 prefix and the mapping to which outside IP, the following command can be used:

```
show router 15001 nat lsn-hosts inside-ip-prefix 10.1.1.0/24
===============================================================================
Large-Scale NAT hosts for router 15001
===============================================================================
Inside IP       Out-Router     Outside IP
-------------------------------------------------------------------------------
10.1.1.0        15002          192.168.2.2
10.1.1.1        15002          192.168.2.2
<snip>
10.1.1.255      15002          192.168.2.2
-------------------------------------------------------------------------------
No. of hosts: 256
===============================================================================
```

To show LSN blocks on the outside routing instance 15002 for the first inside IP within 10.1.0.0/23 prefix, the following command can be used:

```
show router 15002 nat lsn-blocks inside-ip 10.1.0.0
===============================================================================
Large-Scale NAT blocks for vprn15002
===============================================================================
192.168.2.1 [1024..1203]
Pool                                : nat-pool-2
Policy                              : nat-policy-2
Started                             : 2013/07/21 09:55:49
Inside router                       : vprn15001
Inside IP address                   : 10.1.0.0
-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

To show LSN blocks on the outside routing instance 15002 for the last inside IP within 10.1.0.0/23 prefix, the following command can be used:

```
show router 15002 nat lsn-blocks inside-ip 10.1.1.255
===============================================================================
Large-Scale NAT blocks for vprn15002
===============================================================================
192.168.2.2 [46924..47103]
Pool                                : nat-pool-2
Policy                              : nat-policy-2
Started                             : 2013/07/21 09:55:49
Inside router                       : vprn15001
Inside IP address                   : 10.1.1.255
```

```
-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

- Prefix 10.2.0.0/22
  - → Since the subscriber-limit is 64 in this case, the 10.2.0.0/22 prefix will be broken into 16 smaller prefixes each of /26, each will be mapped into a specific outside IP address.
  - → To show LSN hosts for the inside routing instance 15001 for the first /26 prefix and the mapping to which outside IP, the following command can be used:

```
show router 15001 nat lsn-hosts inside-ip-prefix 10.2.0.0/26
===============================================================================
Large-Scale NAT hosts for router 15001
===============================================================================
Inside IP       Out-Router     Outside IP
-------------------------------------------------------------------------------
10.2.0.0        15002          192.168.3.1
10.2.0.1        15002          192.168.3.1
<snip>
10.2.0.63       15002          192.168.3.1
-------------------------------------------------------------------------------
No. of hosts: 64
===============================================================================
```

To show LSN hosts for the inside routing instance 15001 for last /26 prefix and mapping to which outside IP, the following command can be used:

```
show router 15001 nat lsn-hosts inside-ip-prefix 10.2.3.192/26
===============================================================================
Large-Scale NAT hosts for router 15001
===============================================================================
Inside IP       Out-Router     Outside IP
-------------------------------------------------------------------------------
10.2.3.192      15002          192.168.3.16
10.2.3.193      15002          192.168.3.16
<snip>
10.2.3.255      15002          192.168.3.16
-------------------------------------------------------------------------------
No. of hosts: 64
===============================================================================
```

To show LSN blocks on the outside routing instance 15002 for the first inside IP within 10.2.0.0/22 prefix, the following command can be used:

```
show router 15002 nat lsn-blocks inside-ip 10.2.0.0
===============================================================================
Large-Scale NAT blocks for vprn15002
===============================================================================
192.168.3.1 [4024..4863]
Pool                                : nat-pool-3
Policy                              : nat-policy-3
Started                             : 2013/07/21 09:56:23
Inside router                       : vprn15001
Inside IP address                   : 10.2.0.0
-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

To show LSN blocks on the outside routing instance 15002 for the last inside IP within 10.2.0.0/22 prefix, the following command can be used:

```
show router 15002 nat lsn-blocks inside-ip 10.2.3.255
===============================================================================
Large-Scale NAT blocks for vprn15002
===============================================================================
192.168.3.16 [56944..57783]
Pool                                : nat-pool-3
Policy                              : nat-policy-3
Started                             : 2013/07/21 09:56:23
Inside router                       : vprn15001
Inside IP address                   : 10.2.3.255
-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

Mapping results

- Prefix 10.1.0.0/23

  → According to this configuration, each inside IP address has one deterministic block of 180 ports and can have up to three dynamic blocks (block-limit =4) each of 80 ports, allowing a maximum of 180+3*80 = 420 flows.

- Prefix 10.2.0.0/22

  → According to this configuration, each inside IP address has one deterministic block of 840 ports, and can have up to one dynamic block (block-limit =2) of 120 ports, allowing a maximum of 840+120 = 960 flows.

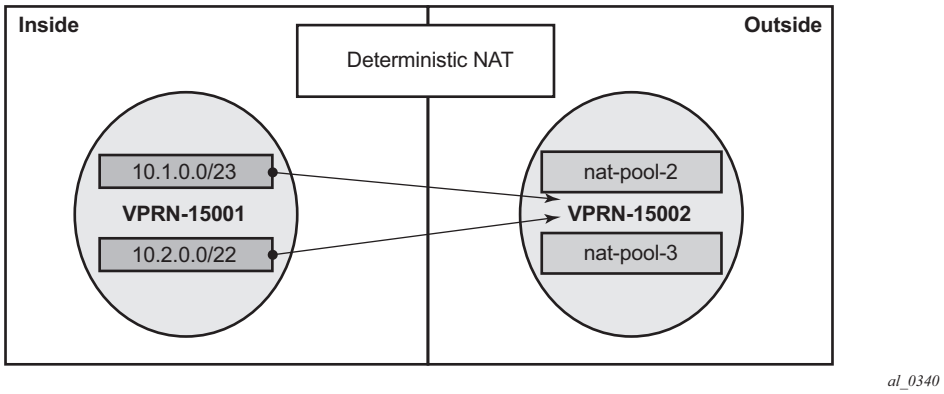

Figure 224: Case 2 Prefix 10.1.0.0/23 Results

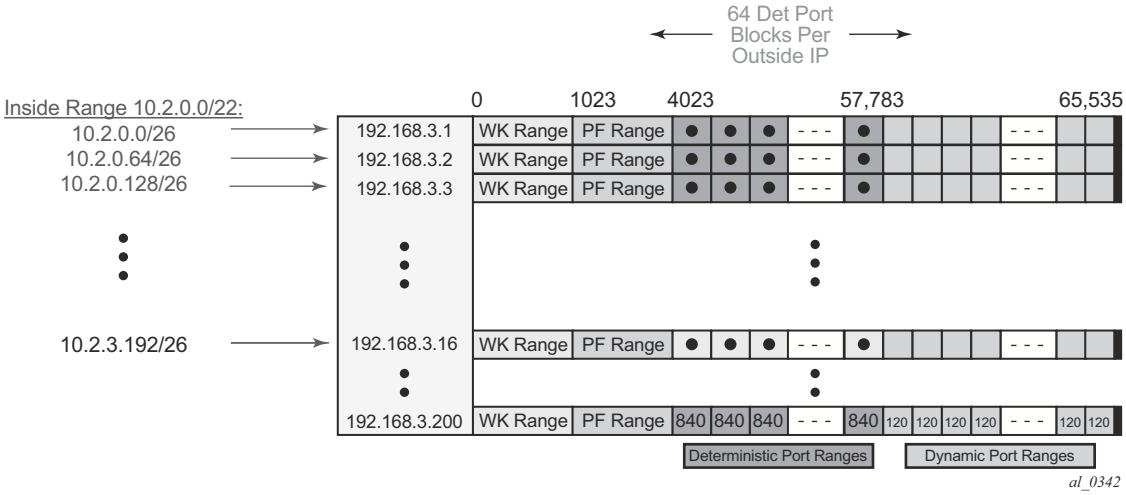

Figure 225: Case 2 Prefix 10.2.0.0/22 Results

Case 3

Configured with:

- Mapping overlapping prefixes from different VRFs into the same outside pool.
- NAT all traffic.

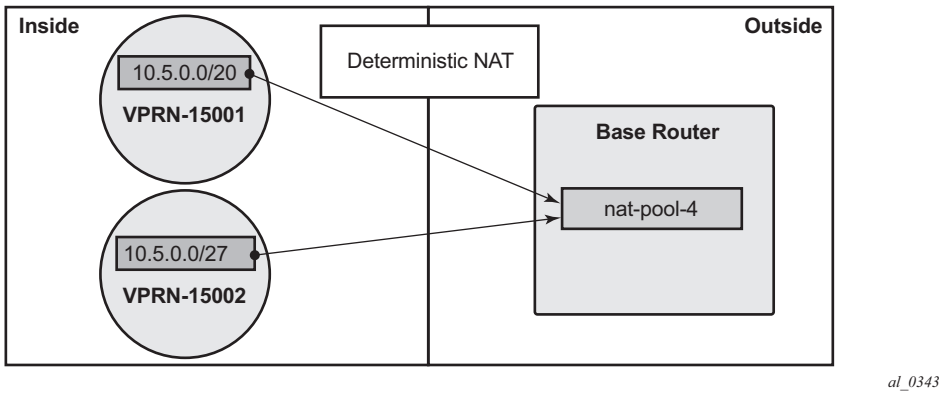

Figure 226: Case 3

The NAT **outside pool** is configured as follows:

```
configure router nat
         outside
                pool "nat-pool-4" nat-group 1 type large-scale create
                    port-reservation ports 461
                    port-forwarding-range 4023
                    subscriber-limit 64
                    deterministic
                        port-reservation 500
                    exit
                    address-range 192.168.4.1 192.168.4.100 create
                    exit
                    no shutdown
exit all
```

The NAT **policy** is configured as follows:

```
configure service nat
          nat-policy "nat-policy-4" create
                block-limit 4
                pool "nat-pool-4" router Base
exit all
```

The NAT **inside prefix** is configured as follows:

```
configure service vprn 15001 nat
     inside
       destination-prefix 0.0.0.0/0
          deterministic
           classic-lsn-max-subscriber-limit 256
           prefix 10.5.0.0/20 subscriber-type classic-lsn-sub
                              nat-policy "nat-policy-4" create
                map start 10.5.0.0 end 10.5.15.255 to 192.168.4.1
                no shutdown
exit all

configure service vprn 15002 nat
     inside
       destination-prefix 0.0.0.0/0
          deterministic
           classic-lsn-max-subscriber-limit 128
           prefix 10.5.0.0/27 subscriber-type classic-lsn-sub
                              nat-policy "nat-policy-4" create
                map start 10.5.0.0 end 10.5.0.31 to 192.168.4.65
                no shutdown
exit all
```

Show commands

- Prefix 10.5.0.0/20 (VPRN 15001)
  - → Since the subscriber-limit is 64 in this case, the 10.5.0.0/20 prefix will be broken into 64 smaller prefixes each of /26, each will be mapped into a specific outside IP address.
  - → To show Large Scale NAT (LSN) hosts for the inside routing instance 15001 for the first /26 prefix and the mapping to which outside IP, the following command can be used:

```
show router 15001 nat lsn-hosts inside-ip-prefix 10.5.0.0/26
===============================================================================
Large-Scale NAT hosts for router 15001
===============================================================================
Inside IP       Out-Router     Outside IP
-------------------------------------------------------------------------------
10.5.0.0        Base           192.168.4.1
10.5.0.1        Base           192.168.4.1
<snip>
10.5.0.63       Base           192.168.4.1
-------------------------------------------------------------------------------
No. of hosts: 64
===============================================================================
```

To show Large Scale NAT (LSN) hosts for the inside routing instance 15001 for the last /26 prefix and mapping to which outside IP, the following command can be used:

```
show router 15001 nat lsn-hosts inside-ip-prefix 10.5.15.192/26
 ===============================================================================
```

```
Large-Scale NAT hosts for router 15001
===============================================================================
Inside IP       Out-Router     Outside IP
-------------------------------------------------------------------------------
10.5.15.192     Base           192.168.4.64
10.5.15.193     Base           192.168.4.64
<snip>
10.5.15.255     Base           192.168.4.64
-------------------------------------------------------------------------------
No. of hosts: 64
===============================================================================
```

To show LSN blocks on the outside routing instance **Base** for the first inside IP within 10.5.0.0/20 prefix, the following command can be used:

```
show router nat lsn-blocks inside-ip 10.5.0.0 inside-router 15001
===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.4.1 [4024..4523]
Pool                                : nat-pool-4
Policy                              : nat-policy-4
Started                             : 2013/07/21 10:18:39
Inside router                       : vprn15001
Inside IP address                   : 10.5.0.0
-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

To show LSN blocks on the outside routing instance **Base** for the last inside IP within 10.5.0.0/20 prefix, the following command can be used:

```
show router nat lsn-blocks inside-ip 10.5.15.255 inside-router 15001
===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.4.64 [35524..36023]
Pool                                : nat-pool-4
Policy                              : nat-policy-4
Started                             : 2013/07/21 10:18:39
Inside router                       : vprn15001
Inside IP address                   : 10.5.15.255
-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

- Prefix 10.5.0.0/27 (VPRN 15002)
  - → Since the subscriber-limit is 64 in this case, the 10.5.0.0/27 prefix will be mapped into one outside IP address.
  - → To show LSN hosts for the inside routing instance 15002 for the 10.5.0.0/27 prefix and the mapping to which outside IP, the following command can be used:

```
show router 15002 nat lsn-hosts inside-ip-prefix 10.5.0.0/27
===============================================================================
Large-Scale NAT hosts for router 15002
===============================================================================
Inside IP       Out-Router     Outside IP
-------------------------------------------------------------------------------
10.5.0.0        Base           192.168.4.65
10.5.0.1        Base           192.168.4.65
<snip>
10.5.0.31       Base           192.168.4.65
-------------------------------------------------------------------------------
No. of hosts: 32
===============================================================================
```

To show LSN blocks on the outside routing instance 15002 for the first inside IP within 10.5.0.0/27 prefix, the following command can be used:

```
show router nat lsn-blocks inside-ip 10.5.0.0 inside-router 15002
===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.4.65 [4024..4523]
Pool                                : nat-pool-4
Policy                              : nat-policy-4
Started                             : 2013/07/21 10:19:40
Inside router                       : vprn15002
Inside IP address                   : 10.5.0.0
-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

To show LSN blocks on the outside routing instance 15002 for the last inside IP within 10.5.0.0/27 prefix, the following command can be used:

```
show router nat lsn-blocks inside-ip 10.5.0.31 inside-router 15002
===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.4.65 [19524..20023]
Pool                                : nat-pool-4
Policy                              : nat-policy-4
Started                             : 2013/07/21 10:19:40
Inside router                       : vprn15002
Inside IP address                   : 10.5.0.31
-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

Mapping results

- According to this configuration, each inside IP address within VPRN 15001 has one deterministic block of 500 ports and can have up to three dynamic blocks (block-limit =4) each of 461 ports, allowing a maximum of 500+3*461 = 1883 flows.
- According to this configuration each inside IP address within VPRN 15002 has one deterministic block of 500 ports and can have up to three dynamic blocks (block-limit =4) each of 461 ports, allowing a maximum of 500+3*461 = 1383 flows.
- For VPRN 15002, since the number of LSN subscribers (32) is less than the number of deterministic blocks (64), then 32 deterministic blocks will be wasted, specifically 32*500 = 16,000 ports will be wasted which is not good in terms of capacity planning.

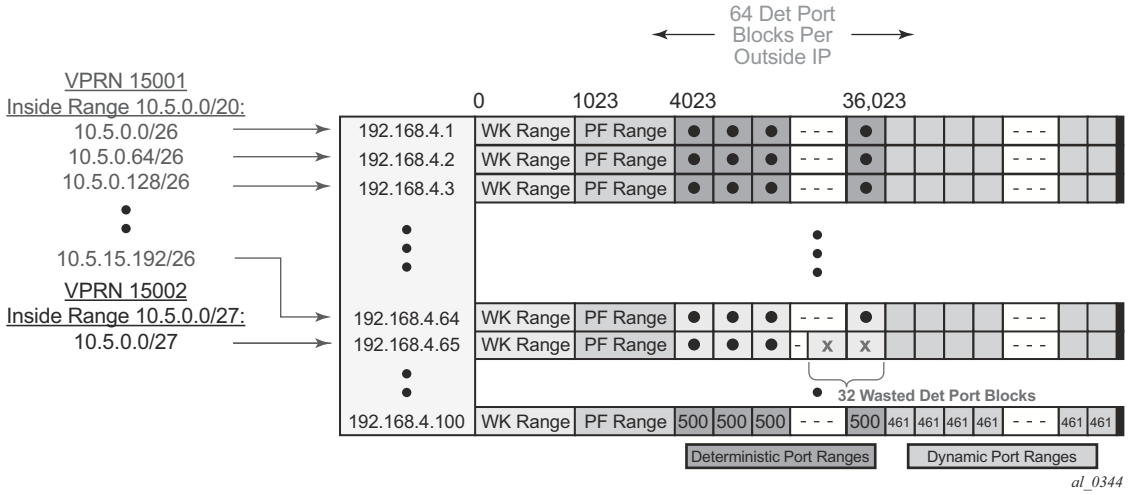

Figure 227: Case 3 Results

Inverse Mapping

In deterministic LSN44, the inside IP addresses are mapped to outside IP addresses and corresponding port-blocks based on deterministic algorithm. The inverse mapping that reveals the subscriber identity behind the NAT is based on the reversal of this algorithm.

Inverse mappings can be done either online or offline:

- Online — Locally on the 7x50 node, via CLI (MIB)
- Offline — Externally, via a Python Script. The purpose of such an offline approach is to provide fast queries without accessing the 7x50.

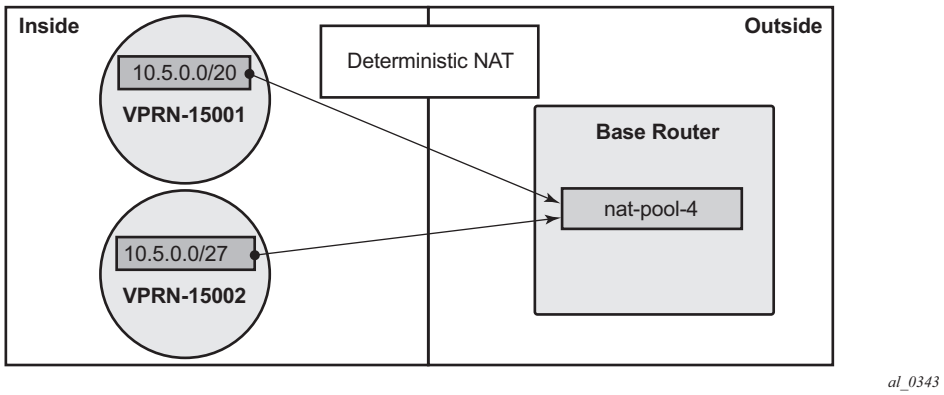

Figure 228: Inverse Mapping Approach

Online Approach

A **tools** command is available which shows the reverse mapping (outside to inside) for deterministic NAT instead of using logging.

```
tools dump nat deterministic-mapping outside-ip <ipv4-address> router <router-instance>
outside-port <[1..65535]>

 <ipv4-address>       : a.b.c.d
 <router-instance>    : <router-name>|<service-id>
                        router-name    - "Base"
                        service-id     - [1..2147483647]
```

Using Case 3 as an example:

To obtain (inside IP, inside routing instance) the inverse mapping for a specific (outside IP, outside routing instance, outside port) is done as follows:

```
tools dump nat deterministic-mapping outside-ip 192.168.4.1 router "Base" outside-port
4024

classic-lsn-sub inside router 15001 ip 10.5.0.0 -- outside router Base ip 192.168.4.1 port
4024 at Sun Jul 21 10:32:44 UTC 2013


tools dump nat deterministic-mapping outside-ip 192.168.4.65 router "Base" outside-port
4024

classic-lsn-sub inside router 15002 ip 10.5.0.0 -- outside router Base ip 192.168.4.65
port 4024 at Sun Jul 21 10:33:38 UTC 2013
```

Offline Approach

The purpose of such an offline approach is to provide fast queries without the need to directly query the 7x50.

This is achieved by generating and exporting a Python script for reverse querying, which is a manual operation that needs to be repeated every time there is configuration change in deterministic NAT.

The script is exported (manually) to the external system.

To configure remote the location for the Python script the following command is used:

```
configure service nat deterministic-script location <remote-url>
```

remote-url — A remote location where the script is stored:
[{ftp://|tftp://}<login>:<pswd>@ <remote-locn>/][<file-path>]
Maximum length is 180 characters.

Once the script location is specified, the script can be exported to that location using the following command:

```
admin nat save-deterministic-script
```

Using the following command the status of the script can be checked, and whether it is necessary to re-save (export) the script or not:

```
show service nat deterministic-script
  ===============================================================================
Deterministic NAT script data
===============================================================================
Location                    : ftp://*:*@10.10.10.10/pub/python/deterministic-
                              nat.py
Save needed                 : no
Last save result            : success
Last save time              : 2013/07/21 10:35:36
===============================================================================
```

The external system must have Python scripting language installed with the following modules: getopt, math, os, socket, and sys.

The Python script can then be run on the external server; the parameters are as follows:

```
user@external-server$./deterministic-nat.py
Usage: deterministic-nat.py {{DIRECTION PARAMS} | -h[elp] }
where  DIRECTION := { -f[orward] | -b[ackward] }
where  PARAMS := { -s[ervice] -a[ddress] -p[ort] }
```

where deterministic-nat.py is the name of the python script previously exported.

Example usage:

A forward query

```
user@external-server$./deterministic-nat.py -f -s 15001 -a 10.0.0.1

classic-lsn-sub has public ip address 192.168.0.1 from service 0 and is using ports [4324
- 4623]
```

A reverse query

```
user@external-server$./deterministic-nat.py -b -s 0 -a 192.168.0.1 -p 4325

classic-lsn-sub has private ip address 10.0.0.1 from service 15001
```

Simultaneous Support of Deterministic and Non-Deterministic NAT

Deterministic NAT can be used simultaneously with non-deterministic NAT within the same inside routing instance. However, they cannot share the same pool.

An outside pool can be only deterministic (although expandable by dynamic ports blocks) or non-deterministic at any given time (a non-deterministic pool is a pool that contains dynamic port-blocks only).

The following show a configuration using deterministic NAT simultaneously with non-deterministic NAT.

The NAT **outside** pools are configured as follows:

```
configure router nat
          outside
                pool "nat-pool-1" nat-group 1 type large-scale create
                    port-reservation ports 180
                    port-forwarding-range 4023
                    subscriber-limit 128
                    deterministic
                        port-reservation 300
                    exit
                    address-range 192.168.0.1 192.168.0.100 create
                    exit
                    no shutdown
                exit
             pool "nat-pool-Non-Deterministic" nat-group 1
                           type large-scale create
                    address-range 192.168.7.1 192.168.7.100 create
                    exit
                    no shutdown
exit all
```

The NAT **policies** are configured as follows:

```
configure service nat
          nat-policy "nat-policy-1" create
                block-limit 4
                pool "nat-pool-1" router Base
            exit
          nat-policy "nat-policy-Non-Deterministic" create
                pool "nat-pool-Non-Deterministic" router Base
exit all
```

The NAT **inside prefixes** are configured as follows:

```
configure service vprn 15001
         nat
               inside
                   destination-prefix 0.0.0.0/0
                   deterministic
                       classic-lsn-max-subscriber-limit 128
                       prefix 10.0.0.0/24 subscriber-type classic-lsn-sub
                                          nat-policy "nat-policy-1" create
                         map start 10.0.0.0 end 10.0.0.255 to 192.168.0.1
                         no shutdown
                       exit
            exit
                   nat-policy "nat-policy-Non-Deterministic"
exit all
```

In this example, the inside IP prefixes that do not match any of the deterministic prefixes will be NATed using a non-deterministic pool.

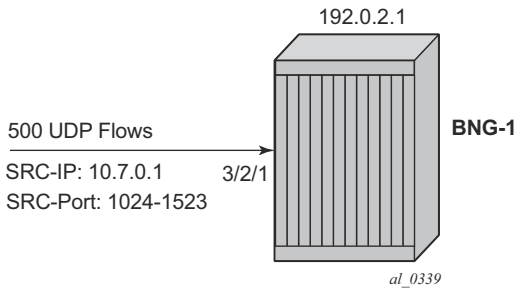

Figure 229: Sending flows: Det + non-Det NAT

To check which NAT pool/NAT policy is used for NATing the inside IP 10.7.0.1, the following command can be used:

```
show router nat lsn-blocks inside-ip 10.7.0.1
===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.7.50 [1024..1527]
Pool                                : nat-pool-Non-Deterministic
Policy                              : nat-policy-Non-Deterministic
Started                             : 2013/07/21 10:59:59
Inside router                       : vprn15001
Inside IP address                   : 10.7.0.1
-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
```

Conclusion

This example provides the commands required for configuring deterministic LSN44 NAT. Both deterministic as well as non-deterministic NAT are supported, with simultaneous operation being possible.

Inverse query can be done online or offline to retrieve the NAT mappings. Logging is not needed as long as there are no dynamic blocks assigned to LSN44 subscriber.