Oversubscribed Multi-Chassis Redundancy (OMCR) in ESM

In This Section

This section describes features and functionality for the Oversubscribed Multi-Chassis Redundancy (OMCR) model.

Topics in this section include:

- Overview on page 1798
- Deploying Oversubscribed Multi-Chassis Redundancy on page 1800
- OMCR Command Reference on page 1819

Overview

Terminology and Abbreviations

- **OMCR** — Oversubscribed Multi-Chassis Redundancy
- **Warm-Standby Node** or **Protecting Node** — Refers to the oversubscribed node that offers the protection of subscriber hosts spread over multiple BNGs. During the normal operation, the protecting node maintains the subscriber host in the form of an MCS record (Multi-Chassis Synchronization Record) in the control plane. Only when the failure occurs and the protecting node becomes active, are the subscriber-hosts fully instantiated in the data and control plane. This node is sometimes referred to as N:1 node.
- **Active/Active (1:1) Model** — This mode of operation refers to the model where subscribers host are fully synchronized between two chassis, regardless of the state of the undelaying SRRP instance (Master/Standby). Each node can have MCS peering sessions with four another nodes where each peering session represent 1 to 1 mapping set of active subscriber hosts.

Restrictions

- The protecting node must use CPM-4 or higher (other protected nodes can continue to use CPM-3).
- The protecting node must use FP2 based cards or higher with chassis mode D.
- The protecting node functionality is not supported in mixed-mode in 7450 ESS chassis.
- All nodes in the OMCR cluster (central standby and the protected nodes) must run at the minimum SR OS R12.0R1.
- Warm-standby mode is a chassis-wide property. In other words, while in warm-standby mode, the chassis cannot operate in 1:1 (active-active) redundancy mode.
- OMCR is supported only for IPoEv4/v6 subscribers. However, non-synchronized PPPoEv4/v6 subscribers are supported in the OMCR cluster. PPPoEv4/v6 PTA (locally terminated) non-synchronized subscribers and the OMCR synchronized IPoE subscriber must be instantiated under separate group-interfaces. On the other hand, non-synchronized PPPoEv4/v6 LAC sessions are allowed to be under the same group interface as the OMCR synchronized IPoE subscribers. Non-synchronized PPPoEv4/v6 subscriber hosts will rely on ppp-keepalive timeouts to re-establish connectivity when the failure occurs.
- Pre-emption of already instantiated subscriber hosts in the protecting node by another subscriber hosts is not allowed.

- Redundant interface (shunting) is not supported for subscribers on the protecting node while they are not fully instantiated in the control/data plane (or while the underlying SRRP instance is in a non-Master state on the protecting node).
- Persistency in multi-chassis environment **must** be disabled since redundant nodes are protecting each other and they maintain up-to-date lease states.
- The failover trigger is based on SRRP only (no MC-LAG support).
- Unnumbered subscriber-interface model is not supported in OMCR.
- The protecting node supports 10 MCS peers, while the protected node (in active/active mode of operation) supports 4 MCS peers.
- Synchronization of the following MCS clients is not supported:
  - Host tracking
  - MC ring
  - Layer 2 subscriber hosts
  - Layer 3 IGMP/MLD
  - Layer 2 IGMP/MLD
  - DHCP Server
  - PPPoE Clients
  - MC-LAG
  - MC-IPSEC
  - MC-ENDPOINT

Deploying Oversubscribed Multi-Chassis Redundancy

In order to optimize the cost, certain operators prefer oversubscribed model in which a single central standby BNG (protecting BNG) supports multiple other BNGs in a semi-stateful fashion.

In Oversubscribed Multi-Chassis Redundancy (OMCR) model, a large number of subscriber-hosts are backed up by a single central standby node. Standby subscriber-hosts within the protecting node are synchronized only within the control plane (CPM) in the form of a Multi-Chassis Synchronization (MCS) record. Such subscriber hosts are not instantiated in the data plane and therefore the data plane resources can be spared and used only on an **as needed** basis. This trait allows the protecting node to back up a large number of subscribers that are scattered over multiple active BNG nodes at the expense of slower convergence.

Only a subset of the subscribers, up to the available resource capacity of the data plane in the protecting node, would be activated on the protecting node at any given time during the failure.

The failover trigger is based on SRRP (no MC-LAG support). The subscriber hosts under the corresponding group-interface will be switched over once the SRRP instance on the protecting node transitions into the Master SRRP state.

There are two possible models for this deployment:

1. Access nodes are directly connected to the BNGs. From the perspective of standby subscribers, in this model the line card is oversubscribed but the physical ports on it are not. For example each of the 10 physical ports on the same line card can be directly connected to respective access nodes. Assume that each physical port can support 64K subscriber hosts. Considering that the subscriber host limit per line card is also 64K (at the time of this writing), the oversubscription ratio in this case would be 10:1.

   The concept of this deployment scenario is shown in Figure 138.

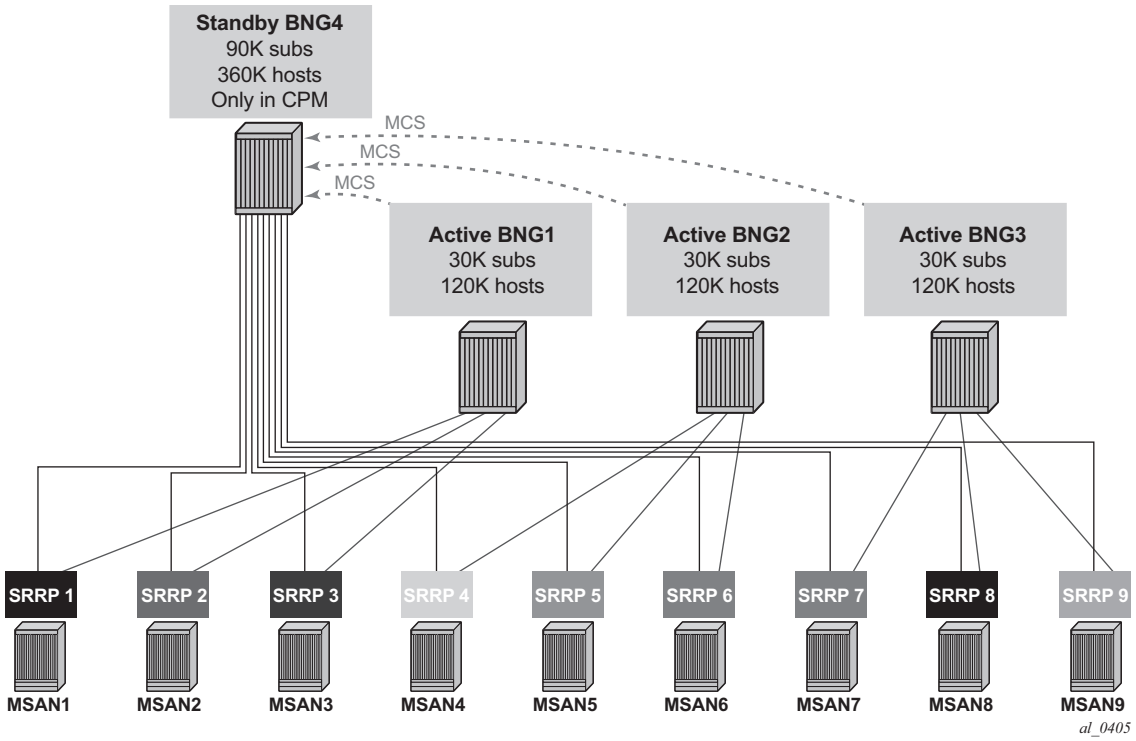

Figure 138: OMCR Scenario Without Aggregation Network

2. Aggregation network in the access (double VLAN tags). In this case a line card and a physical port can be oversubscribed with standby subscribers. For example multiple capture-saps (each capture sap containing 4K c-vlans) can be created on a single physical port on the protecting BNG, for the total of >>64K subscribers per physical port.

Conceptual model for this scenario is shown in Figure 139 (although the number of SRRP instances and capture-saps is in this figure is reduced for simplification).

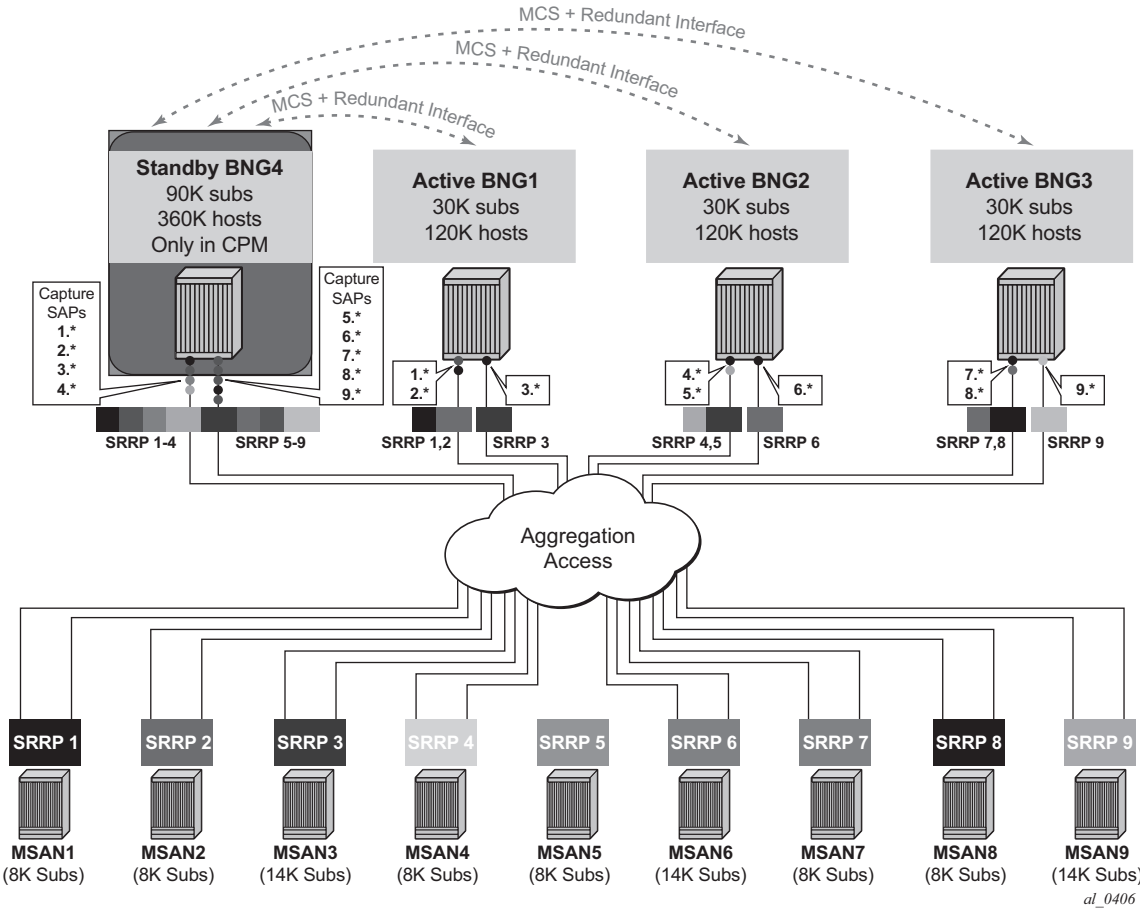

Figure 139: OMCR Scenario with Aggregation Network

In both cases, a maximum of 64K subscribers per line card can be activated on the protecting BNG during the switchover. This is something that the operator should plan around, and consequently group the access nodes in a way so that the eventual number of active subscribers per line card on the protecting node does not exceed the maximum number of supported subscribers per line card.

Note that one could have deployment scenario in which system wide ESM capacity is oversubscribed but the line card capacity is not. For example, on chassis with 10 line cards, each line card can be reserved to protect a total host count of 64k. This would yield a total of 640k

protected hosts distributed across the 10 cards but only up to 256k hosts could be activated simultaneously should it be required due to SRRP transitions to Master.

Resource Exhaustion Notification and Simultaneous Failures

The protection success of the OMCR model relies on grouping protected entities (links and nodes) according to the likelihood of their failure within the timeframe required for their restoration. For example the same resource (IOM card or port) on the protecting node can be used to protect multiple entities in the network as long as their failures do not overlap in time. In other words, if one failure can be repaired before the next one contending for the same resource on the central standby node, the OMCR model will serve the purpose.

But since the oversubscribed model does not offer any guarantees, it is possible that the protecting node in certain cases runs out of resources and fails to offer protection. In this case, the protecting node will generate a SNMP trap identifying the SRRP instance on which subscriber protection has failed. One SNMP trap will be raised per SRRP instance in case where at least one subscriber under the corresponding group interface was not instantiated. The trap will be cleared either when all subscribers become instantiated or when the SRRP transition into a non-master state.

The number of the subscriber hosts that failed to instantiate, can also be determined via the operational **show redundancy multi-chassis omcr all** command. This command will show the number of subscribers that failed to instantiate along with SRRP instances on which the subscriber host are relaying for successful connectivity.

Pre-emption of already instantiated subscriber hosts in the protecting node by another subscriber hosts is not allowed.

Resource Monitoring

Management and conservation of resources is of utmost importance in OMCR. The resources consumed by the subscriber host depend on the type and the size of subscriber parameters (the number of strings, length of strings, etc.).

For these reasons it is crucial that the operator has a view of the amount of memory in the CPM utilized by subscribers and the amount of free memory that can be used for additional subscribers. The **MCS** line is of particular interest in this output. In addition, the **Subscriber Mgmt** line shows memory utilization for active subscribers in the CPM.

The **Available Memory** gives an indication about how much memory remains.

For example:

```
*A:right-21# show system memory-pools
===============================================================================
Memory Pools
===============================================================================
Name                Max Allowed    Current Size      Max So Far          In Use
```

```
-------------------------------------------------------------------------------
BFD                     No limit       6,291,456       6,291,456       5,509,872
BGP                     No limit       5,242,880       5,242,880       3,635,976
CFLOWD                  No limit       1,048,576       1,048,576          26,576
Cards & Ports           No limit      24,117,248      27,262,976      18,010,424
DHCP Server             No limit       2,097,152       2,097,152         173,680
ETH-CFM                 No limit       6,291,456       9,437,184       4,016,128
ICC                  25,165,824       7,340,032      25,165,824       2,880,008
IGMP/MLD                No limit       1,048,576       1,048,576         166,216
IMSI Db Appl            No limit       1,048,576       1,048,576         793,984
IOM                     No limit       8,388,608       8,388,608       6,894,360
IP Stack                No limit      29,360,128      35,651,584      13,565,120
IS-IS                   No limit       2,097,152       2,097,152       1,095,360
ISA                     No limit       3,145,728       3,145,728       1,217,464
LDP                     No limit       6,291,456       6,291,456       5,607,240
Logging              411,041,792       6,291,456       6,291,456       3,473,024
MBUF               1,073,741,824       2,097,152       2,097,152         299,976
MCS                     No limit     454,033,408     454,033,408     416,753,472
MPLS/RSVP               No limit      49,283,072      69,206,016      42,947,776
MSCP                    No limit       2,097,152       2,097,152       1,022,848
MSDP                    No limit               0               0               0
Management              No limit      19,922,944      26,214,400       5,689,112
OAM                     No limit       1,048,576       1,048,576          86,080
OSPF                    No limit       8,388,608       8,388,608       4,975,824
OpenFlow                No limit       1,048,576       1,048,576         391,880
PIM                     No limit      19,922,944      19,922,944      15,755,792
PTP                     No limit       1,048,576       1,048,576           1,408
RIP                     No limit               0               0               0
RTM/Policies            No limit       9,437,184       9,437,184       7,002,648
Redundancy              No limit       9,437,184     424,673,280         703,160
SIM                     No limit       3,145,728      12,582,912             648
Services                No limit      25,165,824      25,165,824      18,128,056
Stats                   No limit       1,048,576       1,048,576           9,456
Subscriber Mgmt         No limit      24,117,248      41,943,040      14,846,512
System                  No limit     794,820,608     856,686,592     776,394,656
Traffic Eng             No limit       1,048,576       1,048,576         444,744
VRRP                    No limit       2,097,152       3,145,728         393,808
WEB Redirect         16,777,216       1,048,576       1,048,576         128,640
-------------------------------------------------------------------------------
Current Total Size :    1,540,358,144 bytes
Total In Use       :    1,373,041,928 bytes
Available Memory   :    5,778,702,336 bytes
===============================================================================
```

Similar output is given in regards to CPU utilization:

```
*A:right-21#  show system cpu
===============================================================================
CPU Utilization (Sample period: 1 second)
===============================================================================
Name                                    CPU Time       CPU Usage        Capacity
                                          (uSec)                          Usage
-------------------------------------------------------------------------------
BFD                                            0           0.00%           0.00%
BGP                                        2,504           0.03%           0.05%
BGP PE-CE                                      0           0.00%           0.00%
CFLOWD                                        25          ~0.00%          ~0.00%
Cards & Ports                              14,501           0.18%           0.13%
DHCP Server                                   30          ~0.00%          ~0.00%
ETH-CFM                                      614          ~0.00%           0.06%
ICC                                        1,803           0.02%           0.18%
IGMP/MLD                                     538          ~0.00%           0.05%
IMSI Db Appl                                  37          ~0.00%          ~0.00%
IOM                                            0           0.00%           0.00%
IP Stack                                   4,578           0.05%           0.24%
IS-IS                                        423          ~0.00%           0.02%
ISA                                        2,690           0.03%           0.10%
LDP                                           78          ~0.00%          ~0.00%
Logging                                       13          ~0.00%          ~0.00%
MBUF                                           0           0.00%           0.00%
MCS                                        2,718           0.03%           0.27%
MPLS/RSVP                                  1,137           0.01%           0.08%
MSCP                                           0           0.00%           0.00%
MSDP                                           0           0.00%           0.00%
Management                                 6,571           0.08%           0.19%
OAM                                        1,532           0.01%           0.09%
OSPF                                      18,397           0.23%           0.08%
OpenFlow                                      18          ~0.00%          ~0.00%
PIM                                            0           0.00%           0.00%
PTP                                           24          ~0.00%          ~0.00%
RIP                                            0           0.00%           0.00%
RTM/Policies                                   0           0.00%           0.00%
Redundancy                                 3,618           0.04%           0.19%
SIM                                       10,959           0.13%           1.08%
SNMP Daemon                                    0           0.00%           0.00%
Services                                   1,037           0.01%           0.03%
Stats                                          0           0.00%           0.00%
Subscriber Mgmt                              835           0.01%           0.03%
System                                    29,863           0.37%           1.32%
Traffic Eng                                    0           0.00%           0.00%
VRRP                                         970           0.01%           0.07%
WEB Redirect                                  26          ~0.00%          ~0.00%
-------------------------------------------------------------------------------
Total                                  7,975,383         100.00%
   Idle                                7,869,844          98.67%
   Usage                                 105,539           1.32%
Busiest Core Utilization                  33,264           3.33%
===============================================================================
*A:right-21#
```

Warm-Standby Mode Of Operation

The protecting node operates in a warm-standby mode. Warm-standby mode of operation is a property of the entire node. In other words, while in the central-standby mode of operation (warm-standby command), only subscribers under the SRRP instances that are in the Master state will be fully instantiated in the data plane on the central standby node (protecting node). All other subscribers (under the SRRP instances that are in the standby state) will be synchronized only in the control plane. However, non-central standby nodes can have a peering connection with a protecting node (OMCR) and at the same time another peering connection with another active BNG node in active/active model. All nodes participating in the OMCR mode of operation must run SROS 12.o or higher. This model is shown in Figure 140.

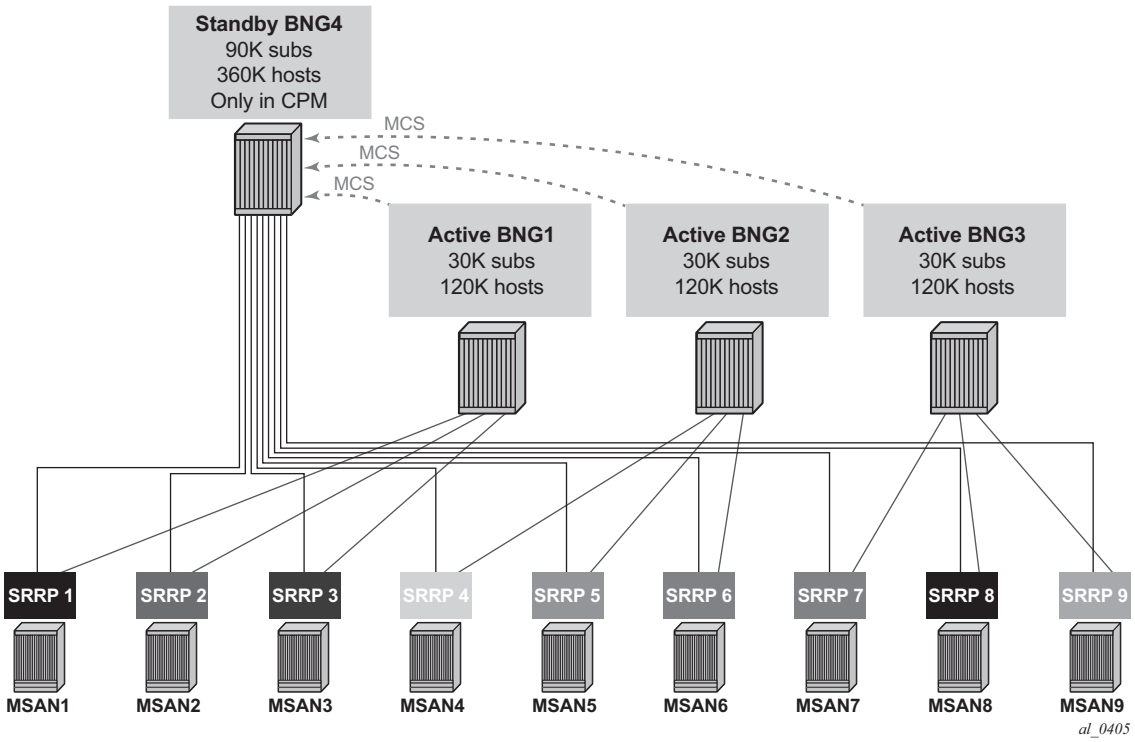

Figure 140: Network Wide Mixing of OMCR and Active/Active (1:1) Model

The central backup property is configured with the following CLI:

```
configure
     redundancy
          multi-chassis
               peer 1.1.1.1
                    warm-standby
```

The **warm-standby** keyword configures the chassis to be in the central standby mode of operation. Although the configuration option is configured per peer, the **warm-standby** functionality is applied per chassis.

Synchronization of IPoE subscribers (**config>redundancy>multi-chassis>peer>sync>sub-mgmt ipoe**) on the protecting node is only possible if all peers are configured for **warm-standby** or none are.

To transition from one mode to another (warm <--> hot), all peers must be administratively shutdown and the warm-standby keyword must be either removed or configured on all peers, depending on the direction of the transition.

Single-homed subscribers are supported in the central standby node, subject to resource limitations.

IPoE vs PPPoE

OMCR is supported only for IPoEv4/v6 subscribers. PPPoEv4/v6 subscriber hosts are not supported. However, non-synchronized PPPoE hosts can be hosted on the protecting node simultaneously with the protected IPoE subscribers. PPPoE PTA (locally terminated) non-synchronized subscribers and OMCR synchronized IPoE subscriber must not be configured under the same group-interfaces. On the other hand, non-synchronized PPPoE LAC sessions are allowed to be under the same group interface as the OMCR synchronized IPoE subscribers.

The recovery of PPPoE subscriber host in non-synchronized environment is based on the timeout of ppp-keepalives.

Persistency

Persistency in the multi-chassis environment **must** be disabled since redundant nodes are protecting each other and they maintain up-to-date lease states. Otherwise, race conditions resulting in stale lease states caused by contention between MCS data and persistency data may occur.

Routing and Redundant Interface in OMCR

Support for redundant interface is limited and can be used only in cases where subscribers are activated in the protecting node. In other words, the shunting over the redundant interface cannot be used if subscriber hosts are not fully instantiated (in the data and control plane). For this reason, downstream traffic must not be attracted (via routing) to the protecting node while the subscriber hosts are in the standby mode (SRRP is in a backup state).

During the transient period while the switchover is in progress, subscriber hosts are being instantiated or withdrawn (depending on the direction of the switchover) in the data plane on the protecting node. The duration of this process is dependent on the number of the hosts that needs to be instantiated/withdrawn and it is proportional to the regular host setup/tear-down rates. The redundant interface in this case can be used only for the hosts that are present in the data plane during the switchover transitioning period (from the moment that they are instantiated in the dataplane when protecting node is assuming activity, or up the moment when they are withdrawn from the data plane when the protecting node is relinquishing activity).

The following routing models are supported:

Case 1 — SRRP-Aware routing where subnets can be assigned per group-interfaces (SRRP instances). In a steady state, the redundant interface is not needed since the downstream traffic is attracted to the master node. During switchover periods (routing convergence transitioning periods), redundant interface can be used only for the subscriber hosts that are instantiated in the data plane (from the moment that they are instantiated in the dataplane when protecting node is assuming activity, or up the moment when they are withdrawn from the data plane when the protecting node is relinquishing activity). See Figure 141.

Case 2 — SRRP-Aware routing where subnets spawn (are aggregated) over multiple group-interfaces (SRRP instances). In case of a switchover, /32 IPv4 addresses and /64 IPv6 addresses/prefixes are advertised from the protecting node. In a steady state, the redundant interface is not needed since the downstream traffic is attracted via more specific routes (/32s and /64s) to the master node. During switchover periods (routing convergence transitioning periods), the redundant interface can be used only for the subscriber-hosts that are instantiated in the data plane (from the moment that they are instantiated in the dataplane when protecting node is assuming activity, or up the moment when they are withdrawn from the data plane when the protecting node is relinquishing activity). To reduce the number of routes on the network side, /32s and /64s should only be activated on the protecting node.

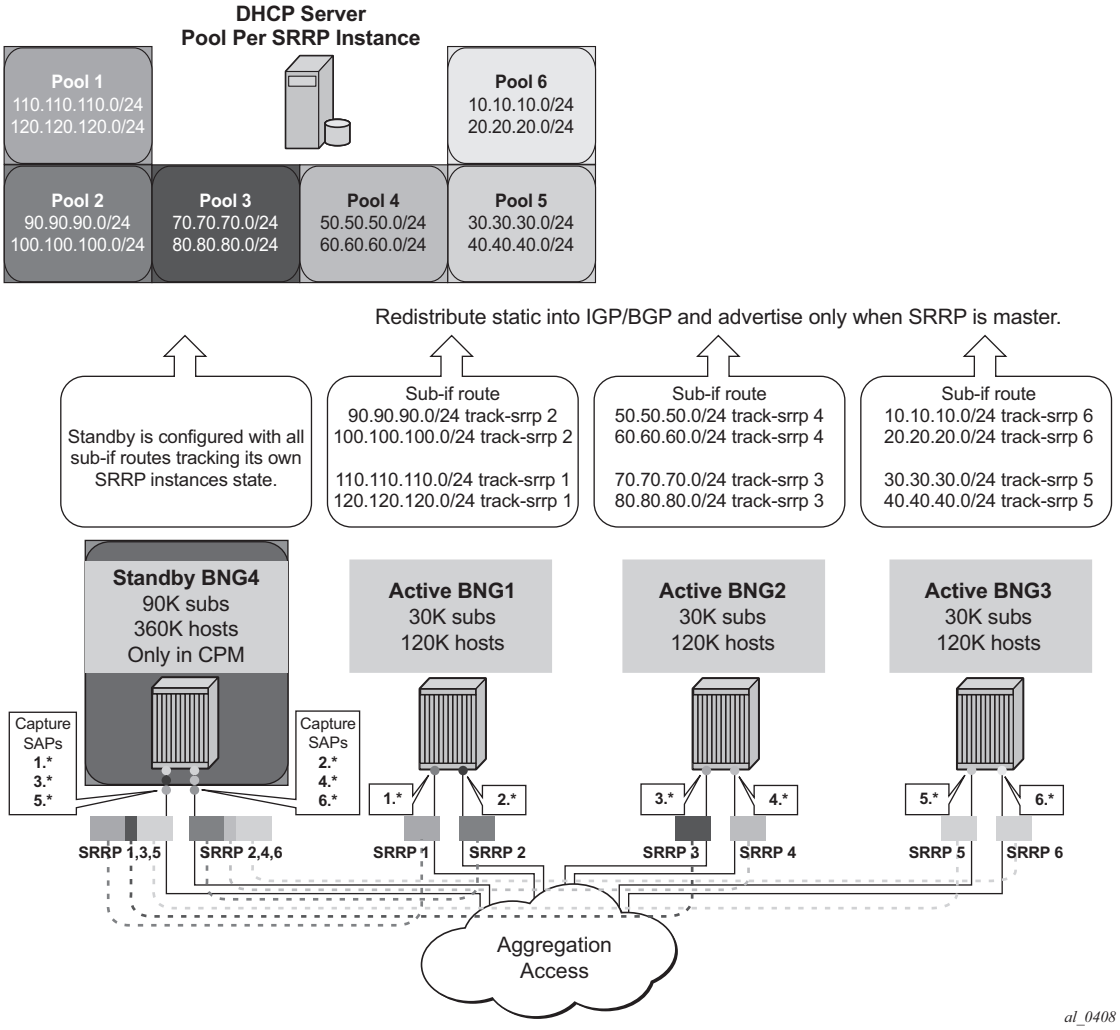

Figure 141: Subnet per Group Interface

A deployment case that is not supported is the one where subnets spawn (are aggregated) over multiple group-interfaces (SRRP instances) and at the same time /32s are not allowed to be advertised from the protecting node. This scenario would require redundant interface support while subscriber-hosts are not necessarily instantiated in the protecting node.

Revertive Behavior

In case that failure is repaired on the original active node (non-central standby node) while SRRP preemption (**preempt**) is configured, the corresponding active subscribers on the protecting node will be withdrawn from the data plane and the activity (mastership) will be switched to the original node.

This behavior will ensure that the resources in the central backup are freed upon failure restoration and are available for protection of other entities in the network (other links/nodes).

In the preemption case, the upstream traffic is steered towards the newly active BNG via gratuitous ARP (GARP). In other words, the virtual MAC is advertised from the newly active node, and consequently the access and aggregation nodes will update their Layer 2 forwarding entries. This action should cause NO interruption in the upstream traffic.

In the downstream direction, the service interruption is equivalent to the time it takes to withdraw the routes from the network side on the standby node. In this case, there are two scenarios:

- A route per group interface (SRRP) is advertised in the network from the central standby node. In this case, downstream traffic interruption is a function of the convergence time of the routing protocol deployed on the network side. Once the routing is converged, all downstream subscriber traffic will be attracted to the newly active node. In the meantime, the redundant interface can be used to shunt traffic from the central standby node to the newly master, but only for the subscriber hosts that have not yet been withdrawn from the data plane on the protecting node. This withdrawal process may take some time and therefore downstream traffic for some subscriber hosts is restored before the others during the routing convergence period.
- /32 subscriber-host routes are advertised from the protecting node. The total recovery time for downstream traffic will depend on the routing convergence. The routing convergence might be slower than in the previous case since more routes (/32s) need to be withdrawn from the network. The redundant interface can be used in the meantime for the subscriber hosts that have not yet been withdrawn from the data plane in protecting node.

Service Restoration Times

Service restoration times depends on the scale of the outage. The factors that affect the restoration times are:

- Failure detection time based on SRRP (could be in a sub second range, also supported based on BFD).
- Time needed to instantiate/withdraw subscriber host in/from the data plane.
- Routing convergence (based on SRRP aware routing).

Processing of the SRRP Flaps

When multiple srrp instances fail at the same time, they will be processed one at the time on first come first serve basis. The subscriber instantiation processing during the switchover is divided into 1seconds intervals. In-between those intervals, the state of the SRRP are checked to ensure that it has not changed while the subscriber instantiation is in progress. This mechanism will break the inertia (snowball effect) that can be caused by SRRP instance flaps. Furthermore, an SRRP flap is handled by not requesting a withdrawal followed by an instantiation request for the same SRRP instance.

Accounting

The OMCR accounting follows the active/active (1:1) redundancy model.

One difference in accounting behavior between the OMCR model and 1:1 redundancy model is in the processing of the accounting session-time attribute which on the protecting node denotes the time when the host was instantiated on the protecting node.

In contrast, the session-time attribute in 1:1 redundancy model is recorded almost simultaneously on both BNG nodes at the time when the host is originally instantiated.

As a result, the session-time attribute is for the most part uninterrupted during the switchover in 1:1 model whereas in OMCR model, the session-time attribute will be reset on the switchover to the protecting node.

Configuration Guidelines

- For all protected SRRP instances, the protected node should be the **preferred Master**. To achieve this, the SRRP priority should be higher in the protected node than in the protecting node. SRRP preemption is recommended in the protecting node to force it to become Master when possible. Note that an SRRP switch from nonMaster to Master in the protected node does not suffer the slow convergence observed when the nonMaster -> Master transition takes place in the protecting node. This is because the protected node always has the hosts instantiated in the data plane.
- ARP hosts configuration is strongly discouraged in protected group interfaces, unless the operator is ready to tolerate an incomplete redundancy mechanism for these hosts.
- SRRP tracking is strongly recommended to expedite the routing convergence upon an SRRP transition from non-Master to Master in the protecting node.
- It is recommended to have a 1:1 relationship between SRRP and subscriber subnets in order to have smooth routing advertisements based on SRRP state tracking.
- The use of M-SAPs should be preferred over the use of static SAPs. Static SAPs are supported in the OMCR mode but they are consuming resources in the protecting node even when the underlying SRRP instance is in a non-Master state.
- It is recommended that the capture-sap configuration include the **track-srrp** statement (at least for the protecting node). With this configuration the CPM will not process trigger packets when the leases cannot be created because the SRRP is not in the Master state. Configuring SRRP tracking at the capture-sap will offload the CPM from performing false authentication and MSAP creation attempts.
- Load balancing between Master and non-Master via export policies for SRRP must not be configured as the hosts are not instantiated in the protecting node when the corresponding SRRP state is non-Master.
- In order to minimize traffic impact in the event of node reboot, it is recommended to use **delayed-enable** *seconds* command under the subscriber-interface and allow enough time for the MCS database to reconcile. This is particularly important in the protected node. If the SRRP becomes master in the protected node before the database has been reconciled, the protecting node will remove the leases (non-Master state) which have not been synchronized. This would create partial outage.
- In order to avoid SRRP collisions, lack of resources and partial subscriber host instantiation, the use of fate-sharing-groups is not recommended. As long as an SRRP instance can be served by the protected node, it is preferred to keep it in the Master state in there, instead of switching it to the protecting node as part of the operation-group.

Troubleshooting Commands

Some of the commands that can assist in troubleshooting are listed below.

Note: To get a summary view of SRRPs and their OMCR status use the following command as shown below (the **domain** concept is reserved for future use):

```
*A:right-21# show redundancy multi-chassis omcr all
===============================================================================
Domain Table
===============================================================================
Domain    Domain SRRP  SRRP    Domain Instan. Failed Failed Reason
name       state ID    State    Color  Failed  Hosts
-------------------------------------------------------------------------------
N/A       N/A    201   Standby N/A    not-act 0
N/A       N/A    202   Standby N/A    not-act 0
N/A       N/A    203   Standby N/A    not-act 0
N/A       N/A    204   Standby N/A    not-act 0
N/A       N/A    301   Standby N/A    not-act 0
N/A       N/A    302   Standby N/A    not-act 0
N/A       N/A    303   Standby N/A    not-act 0
N/A       N/A    304   Standby N/A    not-act 0
N/A       N/A    401   Standby N/A    not-act 0
N/A       N/A    402   Standby N/A    not-act 0
N/A       N/A    403   Standby N/A    not-act 0
N/A       N/A    404   Standby N/A    not-act 0
N/A       N/A    501   Standby N/A    not-act 0
N/A       N/A    601   Standby N/A    not-act 0
N/A       N/A    701   Standby N/A    not-act 0
N/A       N/A    801   Standby N/A    not-act 0
N/A       N/A    901   Standby N/A    not-act 0
N/A       N/A    1001  Standby N/A    not-act 0
N/A       N/A    1101  Standby N/A    not-act 0
-------------------------------------------------------------------------------
No. of Entries: 19
===============================================================================
*A:right-21#
```

Note: To obtain specific SRRP OMCR information, OMCR information has been added to the **show srrp x detail** command:

```
*A:right-21# show srrp 1001 detail
===============================================================================
SRRP Instance 1001
===============================================================================
Description         : (Not Specified)
Admin State         : Up                Oper State       : backupRouting
Oper Flags          : subnetMismatch
Preempt             : yes               One GARP per SAP : no
Monitor Oper Group  : None
System IP           : 10.20.1.1
Service ID          : IES 2
Group If            : grp.Dut-J.1       MAC Address      : 00:00:61:ac:ac:0a
Grp If Description  : N/A
Grp If Admin State  : Up                Grp If Oper State: Up
Subscriber If       : ies-sub-if-svc-2
Sub If Admin State  : Up                Sub If Oper State: Up
Address             : 102.1.0.1/16      Gateway IP       : 102.1.0.3
Address             : 102.2.0.1/16      Gateway IP       : 102.2.0.3
Msg Path SAP        : 8/2/2:2.4094
Admin Gateway MAC   : 00:00:51:ac:0a:01 Oper Gateway MAC : 00:00:51:ac:0a:01
Config Priority     : 1                 In-use Priority  : 1
```

```
Master Priority     : 100
Keep-alive Interval : 10 deci-seconds    Master Since     : 02/11/2014 11:38:52
Master Down Interval: 3.000 sec (Expires in 2.700 sec)
Fib Population Mode : all
VRRP Policy 1       : None               VRRP Policy 2    : None
OMCR Client status  : Sub-mgmt-ipoe
Instantiation failed: not-act            Failed IPOE Hosts: 0
OMCR Reason         :
```

Note: To have a view of the MCS synchronization including OMCR standby records:

```
*A:right-21# show redundancy multi-chassis sync peer 10.20.1.6 detail
===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 10.20.1.6
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 10.20.1.1
Admin State             : Enabled
Warm standby            : Yes
Remote warm standby     : No
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : SUBMGMT-IPOE SRRP
Sync Admin State        : Up
Sync Oper State         : Up
Sync Oper Flags         :
DB Sync State           : inSync
Num Entries             : 64026
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 64000
OMCR Alarm Entries      : 0
Rem Num Entries         : 64026
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
===============================================================================
MCS Application Stats
===============================================================================
Application             : igmp
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : igmpSnooping
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
```

```
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subMgmtIpoe
Num Entries             : 64000
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 64000
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 64000
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : srrp
Num Entries             : 26
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 26
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mcRing
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mldSnooping
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : dhcpServer
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
```

```
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subHostTrk
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subMgmtPppoe
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mcIpsec
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mld
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
```

```
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
Ports synced on peer 10.20.1.6
===============================================================================
Port/Encap                    Tag
-------------------------------------------------------------------------------
4/2/2
  2.1-2.4094                  Dut-F.1
===============================================================================
===============================================================================
DHCP Server instances synced on peer 10.20.1.6
===============================================================================
Router-Name                       Server-Name
  Tag
-------------------------------------------------------------------------------
No instances found
===============================================================================
===============================================================================
*A:right-21#
```

Note: To have the MCS database view of the sync status including OMCR status use the following command syntax:

```
*A:right-21# tools dump redundancy multi-chassis sync-database application sub-mgmt-ipoe
peer 10.20.1.6
The following totals are for:
 peer ip 10.20.1.6, port/lag ALL, sync-tag ALL, application SUBMGMT-IPOE
Valid Entries:                    64000
Locally Deleted Entries:          0
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries:    0
Omcr Alarmed Entries:             0
Omcr Standby Entries:             64000
*A:right-21#
```