Internet Draft
Intended status: Informational
Expires: September 11, 2019

S. Wadhwa
K. DeSmedt
Nokia
R. Shinde
Reliance Jio
J. Newton
Vodafone
R. Hoffman
TELUS
P. Muley
Nokia
Subrat Pani
Juniper Networks
Mar 11, 2019

Architecture for Control and User Plane Separation on BNG
draft-wadhwa-rtgwg-bng-cups-03.txt

Status of this Memo

Copyright Notice

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

This document discusses separation of subscriber-management control plane and data-plane for BNG. Traditionally, the BNG provides aggregation of fixed access nodes (such as DSLAM and OLTs) over Ethernet and provides subscriber management and traffic management functions for residential subscribers. The BNG has however evolved to become a multi-access edge device that also provides termination of subscribers over fixed-wireless and hybrid access. Therefore, this document proposes interfaces between control and user-plane of a BNG that can support multi-access BNG.

Table of Contents

1. Introduction

This document describes requirements and architecture for separation of subscriber management control plane and user plane for the BNG. In rest of the document the control plane is referred to as CP, user plane as UP, and the separation is referred to as CUPS (control and user plane separation). The draft describes the functional decomposition between CP and UP, and applicability of CUPS to a BNG that can support multiple access technologies such as fixed (DSL or Fiber), fixed-wireless (LTE,5G) and hybrid access i.e. simultaneous fixed and wireless access described in BBF [WT378]. The subsequent sections of the draft also define the interfaces required between CP and UP and briefly discusses a candidate base protocol for these interfaces.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. CUPS for BNG

In a CUPS architecture, signaling to setup subscriber sessions CP terminates signaling to setup subscriber sessions, and interfaces with the UP to create forwarding state for these sessions on the UP.

For fixed access subscribers, the CP terminates the signaling protocols (e.g. DHCP, PPPoE, SLAAC) from the customer premise, performs authorization/authentication with AAA Server, participates in address assignment, and then interfaces with the UP to create state related to forwarding and SLA management for the subscriber sessions on the UP. A subscriber session is a single IP connection, such as an IPoE or PPPoE session. The session can be single-stack (IPv4 or IPv6 only), or dual-stack (both IPv4 and IPv6). A CPE can

have multiple sessions, if multiple IP connections are required (e.g. on per service, or one per device behind the CPE).

The CP also processes solicited or unsolicited event notifications from the UP e.g. periodic accounting updates, usage reports, or session inactivity notifications. The interface between CP and UP that is used by the CP to manage session related forwarding state on the UP is being referred to as "state control interface". Asynchronous event notifications from UP to CP are also part of this interface.

In typical fixed access deployments, signaling (e.g. DHCPv4/v6, PPPoE, ICMPv6 RS/RA) to setup the subscriber sessions is in-band, and hence the UP receives the signaling messages from the customer premise. The UP should transparently forward (unmodified) in-band control messages as received from the customer premise to the CP and return messages from CP to the customer premise. Therefore, an in-band signaling channel is required between UP and CP. With a typical "CUPS BNG" deployment, the CP and UP are connected over a network, and the in-band signaling channel must be over a tunnel.

The UP performs forwarding and traffic management for the subscriber sessions. The infrastructure routing and signaling is done on the local control plane of the UP for fast convergence on network topology changes. In rest of the document the term "UP" is used generically for both functions performed by the local control plane on the UP and the data-plane.

A typical deployment architecture for CUPS includes a centralized CP running as a VNF interacting with multiple BNG UP instances that may be more distributed than the CP and could run as VNF or PNF. In this model, the CP and UP association is 1:N. This composite system containing CP VNF and one or more UP instances is referred to as a "CUPS BNG" in rest of the document. For operational ease, the CP MUST provide a single point for control and management for the entire "CUPS BNG". It MUST expose a single interface on behalf of the "CUPS BNG" to external systems such as AAA servers, OSS/BSS, Policy and charging servers. The CP VNF MUST support scale-out in order to cope with growth in number of subscriber sessions and/or increase in number of UP instances in the "CUPS BNG". Figure 1 below shows the functional components and interfaces for a "CUPS BNG".

```
                               |
    "CUPS BNG"                  |
   +----------------------------+------------------------------------+
   |   CP                                                            |
   |   +--------------------------+--------------------------------+ |
   |   | +-----------+ +------------------+ +--------+ +---------+ | |
   |   | | Address   | | PPPoE, DHCPv4/v6 | | RADIUS | | S11/N11 | | |
   |   | | Pool Mmmt | | IPv6 RS/RA,      | | CLIENT | +---------+ | |
   |   | +-----------+ | L2TP LAC         | +--------+             | |
   |   |               +------------------+  +----+  +----+        | |
   |   |                                     | Gx |  | Gy |        | |
   |   |                                     +----+  +----+        | |
   |   +-----------------------------------------------------------+ |
   |          |                  |                   |               |
   |          | Management       |In-band            | State         |
   |          | Interface        |Signaling          | Control       |
   |          |                  |Channel            | Interface     |
   |          |                  |                   |               |
   |  --------+--+------------------+--+----------------+---+------- |
   |          |                        |                    |        |
   |    UP    |               UP       |            UP      |        |
   |    +-----+---------+    +---------+-----+    +---------+-----+  |
   |    | Local CP      |    | Local CP      |    | Local CP      |  |
   |    | Routing, MPLS |    | Routing, MPLS |    | Routing, MPLS |  |
   |    | IGMP, BFD     |    | IGMP, BFD     |    | IGMP, BFD     |  |
   |    +---------------+    +---------------+    +---------------+  |
   |    | Forwarding    |    | Forwarding    |    | Forwarding    |  |
   |    | Traffic Mgmt  |    | Traffic Mgmt  |    | Traffic Mgmt  |  |
   |    +---------------+    +---------------+    +---------------+  |
   |                                                                 |
   +-----------------------------------------------------------------+
```

CUPS BNG System

2.1. Convergence

A single BNG can support subscribers over fixed, "fixed-wireless" or hybrid access. When a residential gateway has fixed-wireless access (LTE or 5G), then the BNG participates in 3GPP signaling with an MME

or AMF (i.e. support 3GPP S11 and N11 interfaces) to setup connections from (NG)RAN. With Hybrid access the customer premise initiates both fixed and wireless connections. The BNG in this case aggregates subscribers over Ethernet from fixed access nodes (DSLAMs and OLTs), but simultaneously terminates connections from (NG)RAN by participating in signaling with MME or AMF (S11/N11 interface). These deployment models are drivers for fixed-mobile convergence. It is important to ensure that the interfaces between CP and UP for CUPS can support not only fixed L2 access, but also the converged access scenario shown in Figure 2. One key requirement on the CP in these cases is the need to participate in 3GPP signaling (which is out-of-band) to setup the data-path. The data-path is a GTP-u (GPRS Tunneling protocol - User Plane) tunnel from the RAN (i.e. S1-u interface for LTE) as described in 3GPP [TS29281], and it terminates on the UP. It carries data traffic but also subscriber signaling messages (e.g. DHCPv4, DHCPv6, SLAAC) from the customer premise. The UP therefore still requires an in-band signaling channel to transport these protocol messages to the CP.

```
        CUPS-BNG
            +-------------------+  S11    +----------+
            |     +------+      |(GTP-c)  | +------+ |
            |     |  CP  |------|---------|-|  MME | |
            |     +---+--+      |         | +------+ |
            |         |         |         | +------+ |
            |         |         |       +-|-| RAN  | |--+     +-----+
            |   ------+-----    |  S1-u | | +------+ |   \    |     |
            |     |       |     |(GTP-u)| +----------+    +---|     |
            |     |       |     | (DHCP |                     | CPE |
            |     |       |     | SLAAC)|                     |     |
            |   +----+  +----+  | (Data)|                 +---|     |
            |   | UP |  | UP |--|-------+     +-----+    /    +-----+
            |   |    |  |    |--|-------------| AN  |---+
            |   +----+  +----+  |    Eth      +-----+
            |                   | (DHCP,PPPoE)
            +-------------------+   (Data)

                          "CUPS BNG" with Converged Access
```

3. Interfaces for CUPS

A "CUPS BNG" MUST support the following interfaces between CP and UP, as shown in the figure in section 2.

3.1. In-band Signaling Channel

Section 2 describes the need for a signaling channel between CP and UP to transport in-band control messages between CP and the customer premise. Following are some key requirements for this interface.

- The UP MUST pass the access circuit identifier over which the signaling messages are received as meta-data to the CP. This includes port, VLAN tags, tunnel endpoint IPs, any tunnel identifiers such as GTP TEID, MPLS labels, L2TP tunnel-id etc. The UP MUST also pass the L2 or L3 transport service that the access circuit is associated with. In case the control message PDU is carried in an Ethernet frame, then the UP SHOULD pass the received Ethernet frame to the CP. Both access circuit identifier and information in the Ethernet header are required by the CP to construct successful response packet (control message) back towards the customer premise. The access circuit identifier MUST be reflected from CP to UP, so UP can identify the access circuit over which it needs to send the CP's response packet. In the control message sent from UP to CP, the UP MUST also include the local MAC address associated with access circuit. This is because certain control messages from the customer premise are destined to a broadcast MAC (e.g. DHCP DISCOVER) or multicast (e.g. ICMPv6 RS), so CP cannot infer the local MAC from these messages. Certain messages also require the local MAC address to be inserted in the message (e.g. Link-Layer address in ICMPv6 RA messages)

- The CP MUST be able to control the UP to forward only specific control messages to the CP.

- The CP MUST be able to control the UP to block certain control messages received on a particular access circuit.

- The CP MUST be able to control the UP to limit the rate of control messages (of specified type) to be sent by the UP.

- The CP MUST be able to prioritize reception of certain control messages over others in a granular manner (e.g. prioritize DHCP RENEWS over DISCOVERS or prioritize PPP Keepalive over other messages).

- The in-band signaling channel MUST support both fixed and converged access as described in section 2.1. The tunnel used for transporting these messages should therefore support both Ethernet and IP payloads.

3.2. State Control Interface

The CP and UP can exchange state at two levels using the "state control interface". One is at the node level and includes node-level information such as supported features, software releases, available resources, and operational state (e.g. active, failed, or overloaded). The other is at the subscriber session level. Subscriber session is described in section 2. The session level state includes basic forwarding and traffic management rules per session, that need to be provided by the CP to the UP in order to control per session forwarding and traffic management on the UP. It also includes state that triggers routing related actions on the UP. The session level state can include asynchronous event notifications from UP to CP, such as notifications to report per session usage (periodically or based on thresholds), notification to report session inactivity, and session liveness.

The interactions between CP and UP over "state control interface" can be categorized as:

- o Session level state management
- o Session level event notifications
- o Node level management
- o Node level event notifications

Following sub-sections provide more details on these interactions. The interactions between CP and UP over "state control interface" are modeled via abstract request/response messages between CP and UP. These messages will need to be defined as part of the protocol specification for this interface.

The protocol selected to implement this interface MUST support both fixed access and converged access (described in section 2.1) on BNG

3.2.1. Session level state management

Once the CP has successfully authorized and/or authenticated the subscriber session, and completed address assignment, it uses the "state control interface" to install forwarding and related state for the session on the forwarding path of the UP. This is abstracted as a "session create request" call from CP to UP, as shown in the figure below. The UP MUST ack or NACK via a response back to CP.

Since BNG can support different access types (e.g. fixed L2 access, or tunneled L3 in case of fixed-wireless, or a combination in case of hybrid access), it is important that the forwarding state information for the subscriber sessions, sent from CP to UP, can be specified as flexible packet matching rules and set of actions related to forwarding and traffic management. The UP should be able to use these match rules and actions to derive various lookup tables and processing in the forwarding path to forward traffic to and from the CPE.

The basic forwarding state in upstream direction (i.e. access to network) and downstream direction (i.e. network to access) fundamentally consists of session identification and one or more actions. Following shows a logical representation of a directive from CP to UP to install basic forwarding state on the UP for fixed L2 access (i.e. access from DSLAM or OLTs over Ethernet).

```
  Direction Upstream - Access to Network:
   Subscriber-identification: Port/VLAN-tag(s) + subscriber-MAC
   Action: remove encapsulation, IP FIB lookup, forward to network.

  Direction Downstream - Network to Access:
   Subscriber-identification: IP address
   Action: lookup IP DA, build encapsulation using Port/VLAN-tag(s)+
  subscriber-MAC, forward to access.
```

Optionally, the IP address assigned to the CPE can also be provided for subscriber-identification (e.g. for anti-spoofing) in the upstream direction.
In case of PPPoE sessions, the subscriber-identification for upstream direction and encapsulation for downstream direction also includes the PPPoE session-id.

Based on the directive from CP to UP (as shown in the example above), the UP can then populate appropriate tables in the forwarding path, e.g. subscriber lookup tables, IP-FIB, and ARP or IPv6 Neighbor discovery table. It can also program the packet processing in both upstream and downstream direction based on the specified actions.

In case of "fixed-wireless" access, the access circuit is a GTP-u tunnel. In this case there is no physical interface (or port), and hence the CP MUST provide a tunnel definition to the UP to use as access circuit in upstream direction, and encapsulation in downstream direction. The tunnel definition will include the tunnel endpoint IP, and TEID that is established via out-of-band signaling

between the CP and the customer premise. It can also include the routing context for transporting the tunnel.

In addition to setting up the forwarding state as directed by the CP, the UP also needs to announce in routing the aggregate prefixes from which the CP assigns IPv4 and IPv6 addresses (or prefixes) to the CPEs. The CP SHOULD provide these aggregate prefixes to the UP as part session state. In case the aggregate prefixes are not provided, the UP MUST announce individual CPE addresses in routing, or it MAY try to aggregate in case addresses for multiple CPEs are from a contiguous address space.

The CPE can have a routed subnet behind it (aka framed-route). CP can learn the framed-routes during authentication/authorization. The CP should provide the framed-route to the UP as part of session state. The UP MUST install this route in the forwarding path and associate it with the forwarding state of the corresponding subscriber session. It should also announce this in routing towards the Network.

The CP MUST also provide to the UP the address assigned as IP gateway address to the CPEs in DHCP. The UP MUST locally configure this address appropriately, such that it can respond to ARP requests for this address from the CPEs.

The session sate on the UP is always controlled by the CP i.e. the UP just follows the directive from the CP to install, modify and delete the session state. In addition to the basic forwarding state, the CP can also associate, update and disassociate other related state with the session e.g. state related to:

- Filtering
- SLA management
- Statistics collection
- Credit control
- Traffic mirroring
- Traffic Steering
- NAT
- Application aware policies

BNG deployments use hierarchical QoS (H-QOS) models which follows from a combination of link-layer over-subscription, multi-service networks and multiple layers of aggregation. For example, a common hierarchy exists of at least a QoS layer per access-node, and per

CPE. The CP MUST provide SLA management information to the UP per CPE. This includes applicable QoS parameters (e.g. rates, queues, markings) and the QoS hierarchy to which the CPE belongs. The CP may choose to signal this via a QoS policy that is locally pre-configured on the UP.

```
                                                              +------+
       +---+          +---+                    +---+          | AAA  |
       |CPE|          |UP |                    |CP |          |Server|
       +---+          +---+                    +---+          +------+
         |DHCP Discover |                         |               |
         |------------->|                         |               |
         |              |                         |               |
         |              |     DHCP Discover       |               |
         |              |-----------------------> |               |
         |              |In-band signaling channel|               |
         |              |                         |               |
         |              |                         |Access Request |
         |              |                         |-------------->|
         |              |                         |               |
         |              |                         |Access Accept  |
         |              |                         |<--------------|
         |              |                         |               |
         |              |      DHCP Offer         |               |
         |              |<------------------------|               |
         |              |In-band signaling channel|               |
         | DHCP Offer   |                         |               |
         |<-------------|                         |               |
         |              |                         |               |
         |DHCP Request  |                         |               |
         |------------->|                         |               |
         |              |                         |               |
         |              |     DHCP Request        |               |
         |              |-----------------------> |               |
         |              | In-band signaling channel
         |              |                         |               |
         |              | Session Creation Req    |               |
         |              |<----------------------- |               |
         |              |                         |               |
         |              | Session Creation Resp.  |               |
         |              |-----------------------> |               |
         |              |                         |               |
         |              |                         |               |
         |              |      DHCP ACK           |               |
         |              |<----------------------- |               |
         |              |In-band signaling channel|               |
         |  DHCP Ack    |                         |               |
         |<-------------|                         |               |
       +---+          +---+                    +---+          +------+
       |CPE|          |UP |                    |CP |          | AAA  |
       +---+          +---+                    +---+          |Server|
                                                              +------+
```

                             Session Creation Sequence

CP can trigger update of session state on the UP, triggered by re-authentication or COA from AAA or policy-server, as show in the figure below.

```
                                                 +------+
       +---+                    +---+            | AAA  |
       |UP |                    |CP |            |Server|
       +---+                    +---+            +--+---+
         |                        | CoA Request     |
         |                        |<----------------|
         |                        |                 |
         | Session Modify Req     |                 |
         |<-----------------------|                 |
         |                        |                 |
         | Session Modify Resp    |                 |
         |----------------------->|                 |
         |                        |                 |
         |                        |   CoA Ack       |
         |                        |---------------->|
       +---+                    +---+            +--+---+
       |UP |                    |CP |            | AAA  |
       +---+                    +---+            |Server|
                                                 +------+
```

Session Modification

CP can trigger the deletion of session state based on signaling messages (as shown in the figure below), administrative action or disconnect-message initiated from the AAA server.

```
       +---+             +---+             +---+       +--------+
       |CPE|             |UP |             |CP |       | AAA    |
       +---+             +---+             +---+       |Server |
         | DHCP Release    | DHCP Release    |         +--------+
         |---------------->|---------------->|               |
         |                 |                 |               |
         |                 | Session Del Req |               |
         |                 |<----------------|               |
         |                 |                 |               |
         |                 | Session Del Resp|               |
         |                 | (final usage)   | Acct Stop     |
         |                 |---------------->|(final usage)  |
         |                 |                 |-------------->|
         |                 |                 } Acct Stop Resp|
         |                 |                 |<--------------|
       +---+              +---+             +----+         +--------+
       |CPE|              |UP |             | CP |         | AAA    |
       +---+              +---+             +----+         |Server |
                                                           +--------+
```

Session Deletion

3.2.2. Session level event notifications

UP can asynchronously generate Session level event notifications to the CP. An example of asynchronous notification is periodic usage reporting from UP to the CP, so that the CP can report the usage to a AAA server via interim accounting-updates. The CP can set the periodicity of this notification on the UP based on interim accounting interval configured by the operator on the CP.

```
                                                +------+
       +---+                    +---+           | AAA  |
       |UP |                    |CP |           |Server|
       +---+                    +---+           +------+
         |                        |                |
         |                        |                |
         |Async Event Notification|                |
         | (periodic usage)       |                |
         |----------------------->|                |
         |                        |                |
         |Async Event Response    |                |
         |<-----------------------|Acct Update Req |
         |                        |--------------->|
         |                        |Acct Update Resp|
         |                        |<---------------|
       +---+                    +---+           +------+
       |UP |                    |CP |           | AAA  |
       +---+                    +---+           |Server|
                                                +------+
```

Async Event Notification for periodic usage

Following are some other examples requiring asynchronous notifications from UP to CP.

- o Threshold based usage reporting
- o Inactivity timeout
- o Subscriber unreachability detection

The protocol for "state control interface" MUST support asynchronous notifications from UP to CP.

3.2.3. Node level management

There needs to exist a concept of association between CP and UP. When the CP or UP comes online it should setup an association with configured or discovered peers via a message exchange. In association setup, the nodes should be able to exchange supported capabilities, version of software, load/overload information, and resource information. Also, any node-wide parameters can be exchanged during association setup.

No session state related messages should be accepted from the peer by either CP or UP unless an association exists.

Either node should be able to update the association to report changed feature capabilities, overload condition, resource exhaustion or any other node-wide parameters.

The UP should be able to request a graceful association release from the CP. In this case the CP should delete all sessions from that UP and process the final stats report for each session and send it in accounting-stop to the AAA server. During this process the CP MUST not create new sessions on the UP. Once all sessions are successfully deleted, the CP should release the association.

There needs to be a periodic node-level heartbeat exchange between CP and UP to detect if the peer is reachable and active. If peer is determined to be down based on heartbeat messages, then all the data-plane session state associated with the peer should be deleted.

```
+---+                    +---+
|UP |                    |CP |
+---+                    +---+
  |                        |
  |                        |
  | Association Setup Req  |
  |----------------------->|
  |                        |
  | Association Setup Resp |
  |<-----------------------|
  |                        |
  | Periodic Heartbeats    |
  |<---------------------->|
  |                        |
+---+                    +---+
|UP |                    |CP |
+---+                    +---+
```

Node Association Setup and Maintenance

3.2.4. Node level event notifications

There needs to be support for asynchronous node level event notifications from UP to CP. Example includes switchover

notification in case ports or UP failures when UP node level warm-standby redundancy is enabled. Based on this notification, the CP can create session state for all the sessions associated with the failure domain on the new primary UP.

3.3. Management Interface

The CP MUST provide a single point for local management of "CUPS BNG" system to the operator. This requires a management interface between CP and each of its associated UPs for pushing configuration to the UP and retrieving operational state from the UP. The interface MUST minimally include BNG specific configuration and state.

The Management interface SHOULD support transactional configuration from CP to UPs and SHOULD support state retrieval, both based on a well-defined data schema. The management interface SHOULD support unsolicited signaling of state changes (events) from UP to CP i.e. MUST provide telemetry for events. Either gNMI or NETCONF can be considered as acceptable candidates for model driven management interface.

4. Resiliency

"CUPS BNG" system MUST be protected against failure of CP VNF and MUST be able to recover the session state without operator intervention and reliance on CPEs. This can be achieved by providing redundancy for processing resources within CP VNF and maintaining redundant instance of session state.

Protection against UP failures based on 1:1 UP (hot-redundancy) and N:M (warm-redundancy) SHOULD be supported. For 1:1 hot-redundancy the CP needs to create data-plane state for sessions on both UPs that form a redundant pair, using the "state control interface". The CP needs to ensure the data-plane state for a session stays synchronized between the two nodes. A given session's data-plane should only be active on one UP in the pair, which serves as active UP for the session. However, sessions that share the redundant UP pair can be distributed between the two UPs for active forwarding.

N:M warm-redundancy (N > M) can be supported via creation of data-plane state on the designated backup chassis after the failure has been detected. This would result in longer failover times than 1:1 hot-redundancy.

Redundant network connectivity between CP and UPs MUST be supported. In the "CUPS BNG" architecture, it is important to configure redundant connectivity that doesn't share fate.

5. Protocol Selection for CUPS Interfaces

It is important that the selected protocol for "state control interface" between CP and UP works not just for fixed access but also works for converged access on BNG. 3GPP has defined PFCP (Packet Forwarding Control Protocol) in [TS29244] as the interface between CP and UP for LTE gateways. This protocol is suited for large scale state management between CP and UP. Following are some of the key attributes of this protocol:

o It supports management of forwarding and QOS enforcement state on the UP from CP. It also supports usage reporting from UP to CP.
o It is over UDP transport and doesn't suffer from any HOL blocking.
o It provides reliable operation based on request/response with message sequencing and retransmissions.
o It provides an overload control procedure where overload on UP can be handled gracefully.
o The protocol is extensible and allows addition of new IEs.

For fixed access BNG, the protocol requires simple extensions in form of additional IEs. The required extensions are mainly due to fact that typically a fixed access BNG requires tighter control over L2 behavior and manages access and subscriber using L2 identifiers (such as VLANs and MAC addresses), whereas mobile access works in terms of L3, either routed or tunneled.

The details of the protocol as applicable to the BNG and the required extensions will be defined in a separate draft.

[TS29244] also describes an in-band signaling channel based on GTP-u tunnel between CP and UP. GTP-u (GPRS Tunneling protocol - User Plane) is defined in 3GPP [TS29281] and defines a tunneling protocol which carries IP payloads. The protocol runs over a UDP/IP stack and uses UDP port number 2152. Data within a tunnel can be multiplexed based on Tunnel Endpoint Identifiers (TEIDs). The protocol supports optional sequence numbers. The protocol supports extension headers to allow development of new features. GTP-u tunnels are signaled between CP and UP, and it is possible

to associate filters to block certain control packets from being forwarded form UP to CP. The payload type carried by GTP-u can be extended to Ethernet (via payload type in extension header). The tunnel encapsulation can also be extended similarly to carry any required meta-data.

6. Address Pool Management

The CP MUST support management of IPv4 and IPv6 address pools, where each pool can contain one or more subnets. The pool management MUST support pool selection based on one or more of the following criteria:

- o UP
- o Access port on the UP.
- o Redundancy domain on the UP (e.g. set of access ports that share fate with respect to switchovers due to failures, when UP node level redundancy is enabled).
- o Service (e.g. HSI, VoIP, IPTV etc.).
- o Location (e.g. based on circuit-id/remote-id or part of circuit-id/remote-id in DHCP and PPPoE).

Pool management on CP SHOULD NOT statically link subnets to UPs but SHOULD dynamically allocate subnets to UP based on load i.e. on-demand, and signal allocated subnets using the "state control interface" as described in section 3.2.1. This allows for better IP resource utilization and less subnet fragmentation.

7. Security Considerations

For security between CP and UP, Network Domain Security (NDS) as defined in [TS33210] can be considered. As per NDS, the network can be split into security domains. Communication within a single security domain is considered secure, and protocols can operate without any additional security. When communication has to cross security domains, then IPSEC can be used.

8. IANA Considerations

None.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[TS29244] 3GPP, "Interface between the Control Plane and the User Plane Nodes", TS 29.244 15.2.0, June 2018, https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3111.

[TS29281] 3GPP, "General Packet Radio System (GPRS) Tunneling Protocol User Plane (GTPv1-U)", TS 29.281 15.3.0, June 2018, https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=1699.

[TS33210] 3GPP, "Network Domain Security (NDS); IP network layer security", TS 33.210 15.0.0, June 2018, https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2279.

[WT378] BBF, "Nodal Requirements for Hybrid Access Broadband Networks", WT-378, 2018.

9.2. Informative References

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, https://www.rfc-editor.org/info/rfc2131.

[RFC2516] Mamakos, L., Lidl, K., Evarts, J., Carrel, D., Simone, D., and R. Wheeler, "A Method for Transmitting PPP Over Ethernet (PPPoE)", RFC 2516, DOI 10.17487/RFC2516, February 1999, https://www.rfc-editor.org/info/rfc2516.

[RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, https://www.rfc-editor.org/info/rfc3315.

Authors' Addresses

Sanjay Wadhwa
Nokia
777 East Middlefield Road
Mountain View
USA

Email: Sanjay.wadhwa@nokia.com

Killian De Smedt
Nokia
Copernicuslaan 50
Antwerp
Belgium

Email: Killian.de_smedt@nokia.com

Rajesh Shinde
Reliance Jio Infocomm Ltd.
Reliance Corporate Park
Thane Belapur Road, Ghansoli
Navi Mumbai 400710
India

Email: Rajesh.A.Shinde@ril.com

Jonathan Newton
Vodafone
Waterside House
Bracknell
United Kingdom

Email: jonathan.newton@vodafone.com

Ryan Hoffman
TELUS
1525 10th Ave SW
Calgary, Alberta
Canada

Email: ryan.hoffman@telus.com

Praveen Muley
Nokia
805. E. Middle Field Rd.
Mountain View, CA, 94043
USA

Email: praveen.muley@nokia.com

Subrat Pani
Juniper Networks

10 Technology Park Dr.
Westford, MA
USA

Email: spani@juniper.net