

An LTE Enhanced Packet Core Architecture based on SDN and OpenFlow

Subir Varma

This report describes a technique by means of which an LTE Packet Core network with Serving Gateways (SGW), PDN Gateways (PGW) and Mobility Management Entities (MME) can be substituted by a Packet Core Network built entirely from Open Flow compliant commodity switches and an Open Flow Network Controller (OFNC) which can be implemented in a Cloud based Server.

The current LTE enhanced packet core (EPC) architecture requires operators to deploy specialized nodes such as SGWs and PGWs, which are more expensive than plain Ethernet or IP switches. They also require the use of packet tunneling, which increases the link overhead and thus wastes bandwidth. In addition, support for mobility and QoS is overly complicated in the current EPC architecture and has the following problems associated with it:

- Use of expensive specialized nodes, such as SGWs and PGWs.
- Use of packet tunneling within the EPC, that wastes bandwidth
- Complex protocols for supporting mobility, QoS and inter-operability with other networks, such a WLAN.
- The operators has limited or no control and influencing any of the control functions within the Packet Core network, which inhibits them from introducing value added features to the network.

This invention re-architects the LTE EPC, while preserving the control interfaces on the User Elements (UE) and the eNB. The new EPC, which is based on Open Flow and Software Defined Networking concepts, has the following benefits:

- Operator's can build the LTE EPC using standard Open Flow compliant switches, which are expected to be lower cost than the specialized SGWs and PGWs that are used in the current EPC architecture
- The proposed EPC design does not use tunneling for internally routing packets, and hence avoids the extra overhead, leading to bandwidth savings.
- The proposed EPC design offers a new design for supporting mobility, in which the packet flows are re-routed in response to UE movement using the Open Flow controller. This re-routing can happen anywhere within the backhaul network, as opposed the current design in which the re-routing takes place at the SGW or the PGW. This simplifies the mobility design, and gives the system more options in finding the best route to use.

- The proposed EPC design allows the operator to add their value added control algorithms to the network, for doing functions such as QoS or Load Balancing, while continuing to use standard switches. These functions are either non-existent (which is the case for load balancing), or difficult to do (which is the case for end to end QoS support) in the current architecture.
- The proposed EPC design allows the operator to support multi-tenancy, using existing support in the Open Flow protocol for a Network Wide hypervisor like function in the Open Flow controller, which partitions the physical network into multiple independent virtual networks.
- The proposed EPC design allows the operator to support Content Delivery Network (CDN) servers, deeper in the backhaul network. This contrasts to the current architecture, where the CDN nodes have to be placed at the PDN GW or further upstream, in order to accommodate the tunneled design in the backhaul.
- The proposed EPC design allows the operator to easily add value added features into the network, as additional applications running on the OFNC. In the current design, adding new features to the network is a time consuming process, in which the design has to be first standardized through 3GPP, and takes multiple years.
- The proposed EPC design helps to move operators towards a network in which the value added control portion of the network is implemented within the operator's server cloud, while the actual switching is done by commodity nodes.

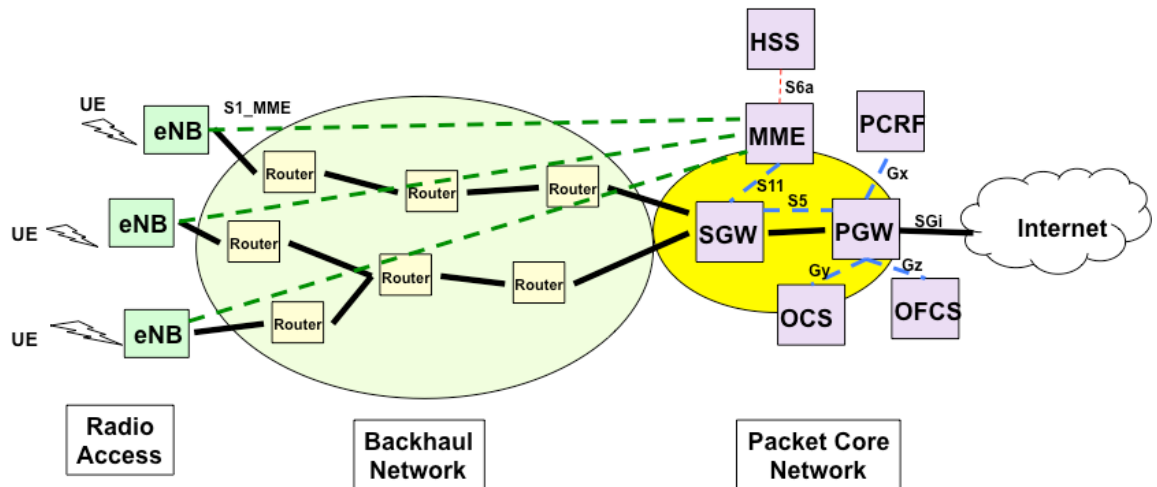


Figure 1

Figure 1 illustrates the current LTE EPC architecture. There is already some degree of separation of Control and Data Plane elements, with the MME responsible for control via the S1_MME interface to the eNB, NAS signaling with the UEs, the S11 interface to the SGW and the S6a interface to the HSS. However note that the SGW and PGW, which are the EPC Data Plane elements, also share some control plane responsibility, with the S5 interface between them, as well as the Gx, Gy, Gz interfaces to the Policy and Charging servers. Also note that the routers in the backhaul are not subject to control by the MME. Hence the current EPC architecture can be characterized as a mixture of distributed and centralized control, with no clean separation between the data and control plane functions.

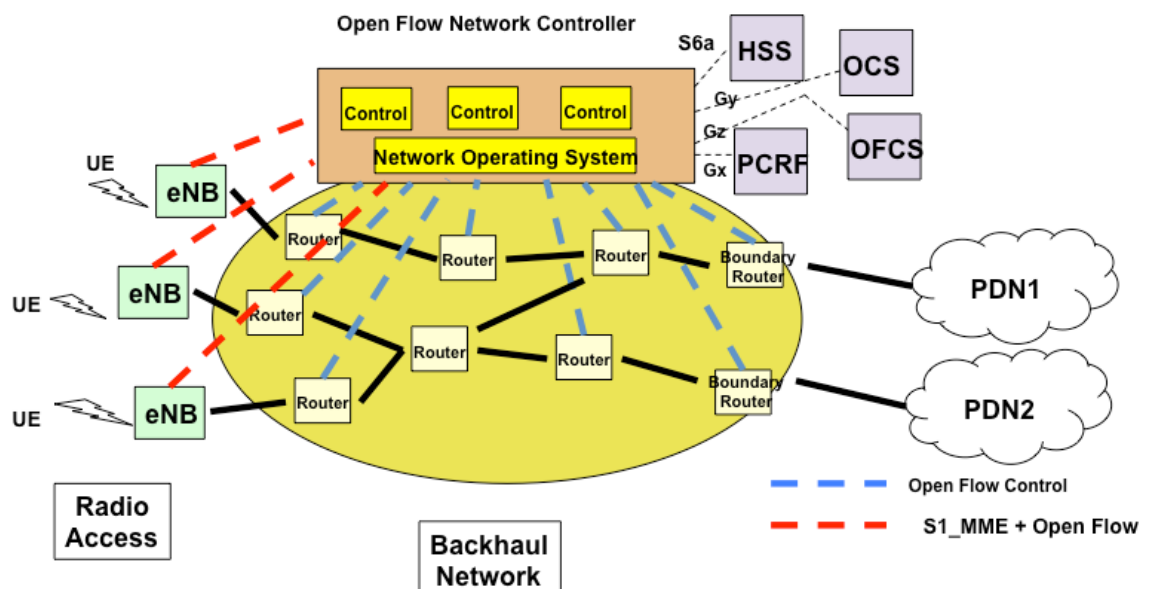


Figure 2

Figure 2 shows the EPC architecture being proposed. It has the following salient features: There is no longer a differentiation between the Backhaul Network and the Packet Core network. Instead there is a single homogeneous network made up of Open Flow

compliant switches, that are centrally controlled by a Open Flow Network Controller (OFNC). The UEs are unchanged from the signaling and data plane viewpoint from the current architecture, while the eNBs support both the S1-MME as well as Open Flow interface to the OFNC.

- The high level structure of the OFNC is shown in Figure 2. It uses a layered architecture, in which the bottom layer is called the NoS or Network Operating System (NoS). The NoS is responsible for communicating and controlling the various nodes in the system using the Open Flow protocol. It also maintains a map of the network topology, which is then used by the control applications running on top of the NoS. This allows the control applications to run their algorithms on the logical topology, without worrying about the distributed aspects of the real network.
- In addition to the control connections to the routers, the OFNC also has the following other LTE defined control connections: The S1_MME control to the eNBs, NAS signaling with the UEs, the S6a control to the HSS and the Gx, Gy, Gz controls to the Policy and Charging servers. The benefits of retaining these control interfaces are the following: (1) It enables the system to use legacy LTE compliant UE and eNB nodes, with little or no changes to their software, (2) Several wireless specific control signals are required, such as those for paging, idle mode signaling as well as control for user authentication, policy control, billing etc that are not currently part of the Open Flow interface definition. Re-using the S1_MME, S6a, Gx, Gy, Gz controls enables the OFNC to perform these other functions.
- The boundary routers that lie at the edge of the Open Flow controlled network are hybrid nodes in the following sense: They support regular distributed IP routing on their interfaces to the IP cloud, while they use centralized Open Flow controlled routing on the mobile core side.
- The UEs are allocated globally routable IP addresses. The subnets from which these IP addresses are taken are advertised by one of the boundary routers, which can be considered to be analogous a Home Agent (HA) type anchor. However, unlike in the HA based design, there are no IP tunnels that need to be setup between these boundary routers and the UEs. Instead, Open Flow based signaling is used to re-route flows in case of UE mobility. Also the UEs may request multiple PDNs, in which case they are allocated a separate IP address, and perhaps boundary routers, for each PDN.
- UE authentication on initial network entry is unchanged from the current EPC specification, using NAS signaling between the UE and the OFNC (via the S1_MME interface) as well as the S6a interface from the OFNC to the HSS.
- Unlike the current design, there is no need to set up tunnels to route packets from the EPC and across the backhaul network. Instead, the OFNC controls the routes at each of the nodes along the path, by means of the Open Flow protocol. The OFNC is able to dynamically change the paths for individual flows in response to events such as UE mobility or network congestion, by changing the Flow Table entries at one or more of the nodes along the path.
- UEs can go into idle mode as in the current design. When this happens, the OFNC removes the Flow Table entry for that UE from all the along its path. Subsequently when the first new data packet arrives at the boundary node, it informs the OFNC using standard Open Flow mechanisms, which then proceeds

- to page the UE using the S1_MME interface, and after detecting the target eNB, re-establishes the path from the gateway router to the target eNB.
- Support for Policy and Charging Control (PCC): The current PCC design involves the use of the Gx, Gy and Gz interfaces from the PCEF function in the PDN GW to the PCRF, OCS and OFCS respectively. In the proposed architecture, the PCEF function continues to exist in the Boundary Switch, however the Gx, Gy, Gz interfaces now exist between the OFNC and the policy nodes. The OFNC acts as a proxy, and translates these interfaces to a single Open Flow interface to the PCEF.
 - Support for QoS is provided by means of dedicated bearers as in the current specification. The difference is that, the routing for the dedicated bearer is explicitly setup by the OFNC on a node-by-node basis, without the use of additional EPC tunnels.

I. Network Entry Procedure

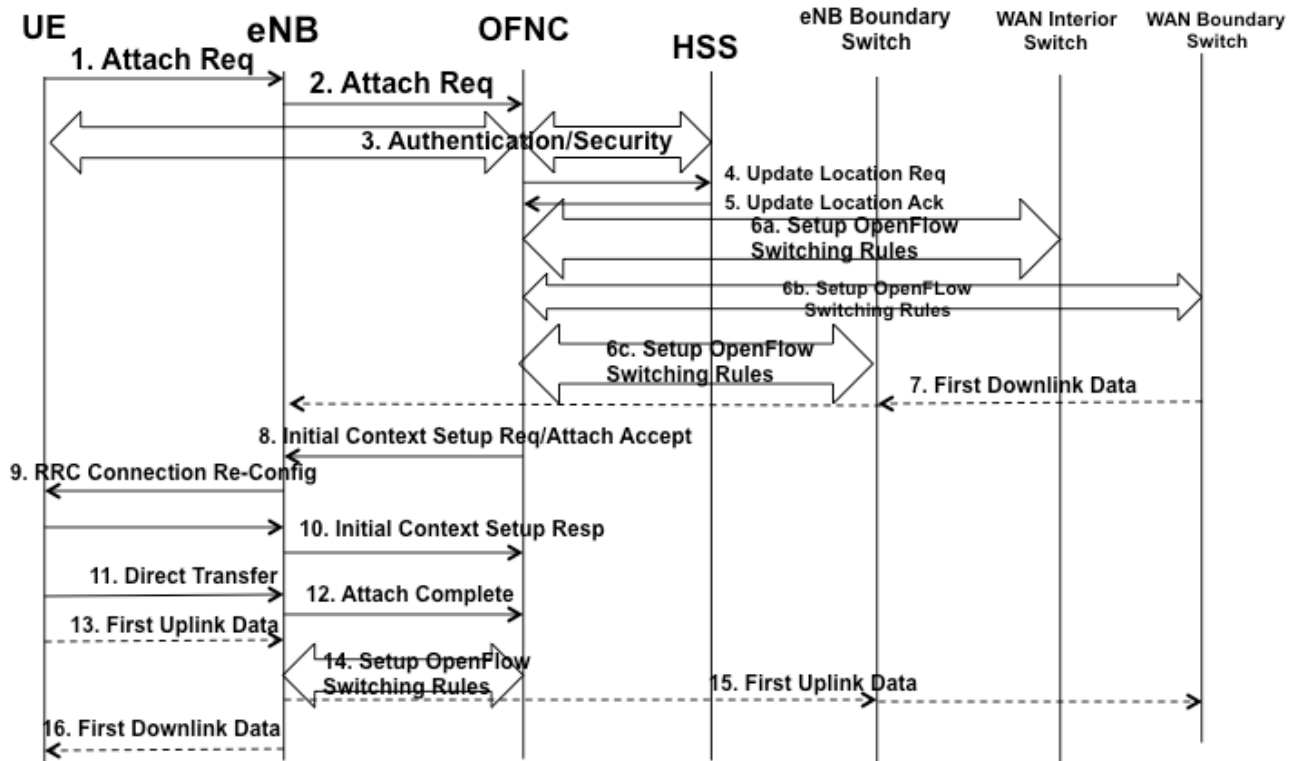


Figure 3

1. The UE sends the standards defined Attach Request to the eNB.
2. The eNB derives the OFNC from the GUMMEI information, and forwards the Attach Request in a S1-MME control message.
3. UE Authentication and Security is carried out as per the LTE standards, with NAS signaling between the OFNC and the UE, and S6a signaling between the OFNC and the HSS.
4. If there is no valid subscription context for the UE in the OFNC, then it sends an Update Location Request to the HSS.
5. The HSS acknowledges the Update Location message by sending an ACK message back to the OFNC. The OFNC makes sure that the UE is allowed to be in the TA, and if all checks are successful, then the OFNC constructs a context for the UE.
6. The OFNC allocates an IP Address to the UE and an associated Boundary Router, based on the selected PDN type. The OFNC also selects all the nodes in the backhaul that will be used to route the UE traffic, from the Boundary Router to the eNB Boundary Switch. The eNB Boundary Switch is chosen from among the backhaul node(s) that have a direct link to the target eNB. The OFNC may use a number of different criteria to do routing between the two ends of the backhaul path. For example, it can be based on the current congestion situation in the backhaul, so that the OFNC avoids nodes that are heavily congested. The identification of UE traffic in the various nodes can be based on using the Destination IP address for downstream traffic, and Source IP address for upstream traffic.
After choosing all the nodes, the OFNC then proceeds to install the switching

- rules in each of their Flow Tables, using the Open Flow protocol. Note that this method of routing traffic is done without doing any encapsulation.
8. Steps 8 to 13 are mostly un-changed from the standard specifications. In these steps, the eNB and the UE establish the bearer connection over the airlink. The UE is also allocated its IP address in step 9.
 14. The OFNC installs the Open Flow switching rules in the eNB's Flow Table, thus completing the end to end path.

II. Handover Procedure

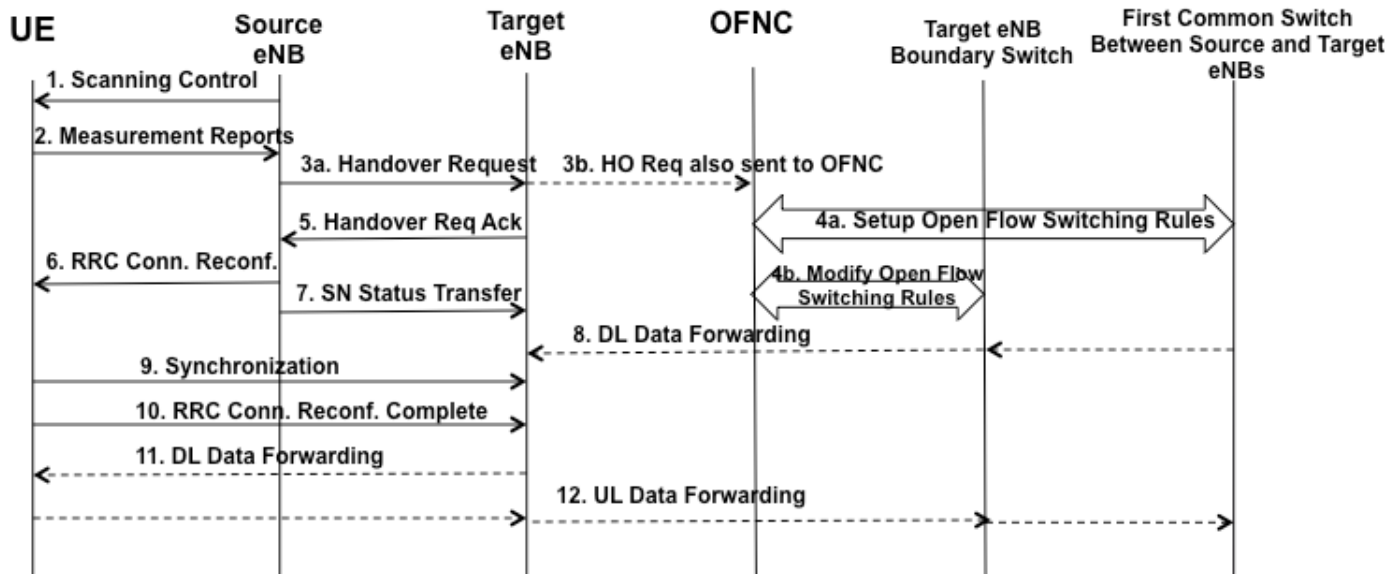


Figure 4

As shown in Figures 4 and 5, the handover signaling and procedure is considerably simplified with Open Flow, as compared to the current scheme.

The main differences between the flows in Figure 4, and the standards based approach described in TS 36.300, Section 10.1.2, are the following:

- When the source eNB decides to do the handover, it sends a Handover Request message to the Target eNB (step 3a). However, this message will get routed to the OFNC, since the switching rule for sending this message does not yet exist in the source eNB. When then OFNC receives this message, it uses it as an indicator to begin the process of re-routing the traffic switching rules from the Source eNB to the Target eNB (steps 4a and 4b). At the end of this process, the downlink data is being sent to the target eNB (step 8).
- Once the UE synchs up with the target eNB, and completes the link setup process (step 9 and 10), the target eNB can then start to send downlink data to the UE, and also route the uplink data to the boundary router.

Unlike the current EPC standards where the new path is forced to go through the SGW, the OFNC has more flexibility in assigning the new path and can use other criteria, such as traffic load, to compute it. This is illustrated in Figure 5, in which the movement of the UE is shown with purple arrow, while old path to the UE is shaded red, while the new

path is shaded green. The OFNC uses its routing algorithm to decide that the router shaded orange is the appropriate node to switch the path. Also note that the current EPC handover design is different depending on whether or not there is a direct path between the eNBs (also called the X2 tunnel). This distinction disappears when Open Flow is used for mobility support.

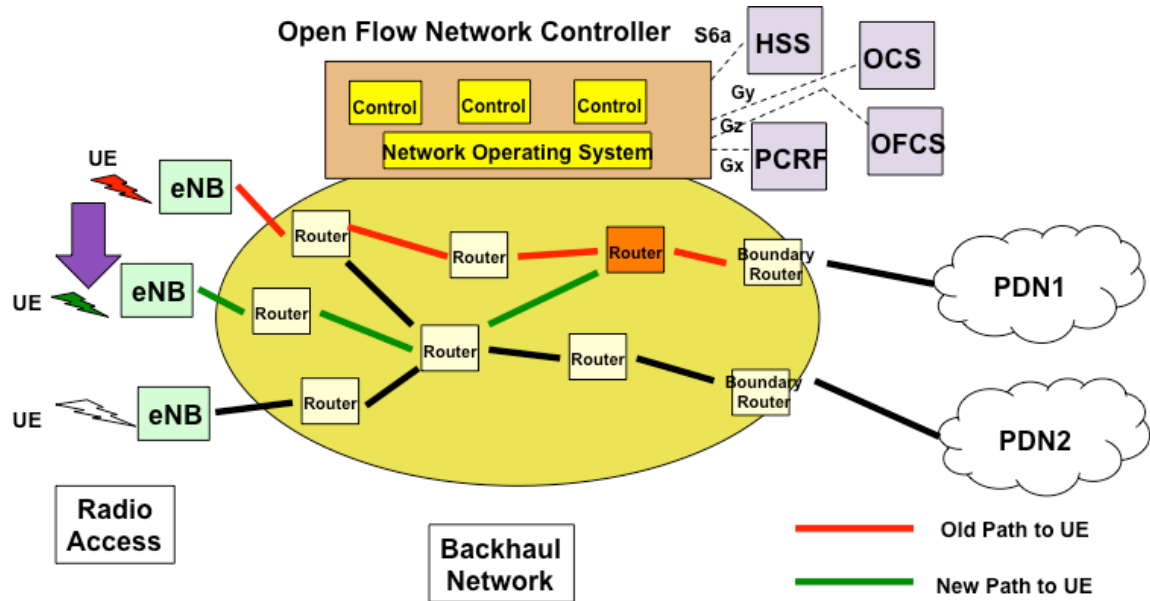


Figure 5

III. Connected to Idle Procedure

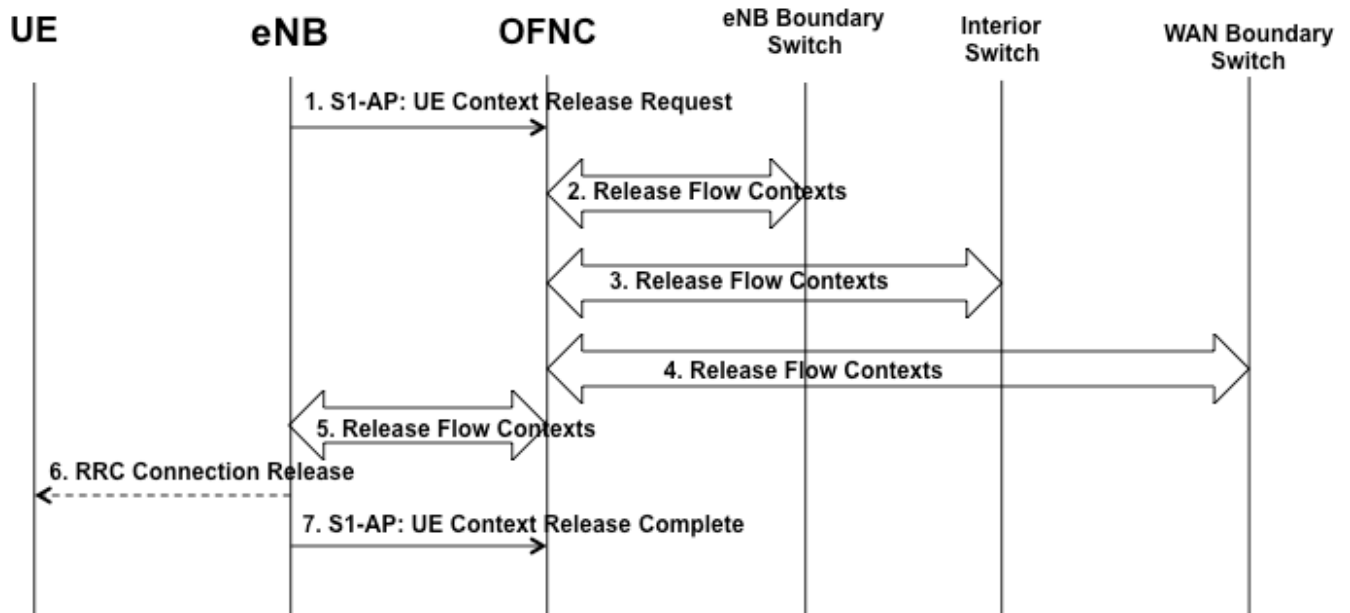


Figure 6

This procedure is used when the UE moves from the ECM-CONNECTED to the ECM-IDLE state, and all the UE related context information is deleted in the eNB. With Open Flow the procedure is simpler than that with the current EPC design. The OFNC removes the switching rules for all the bearers going to the UE, all along the path from the ingress to the egress nodes. However, it retains the UE context information, so that if the UE moves back into the Connected state, then it can quickly set up the forwarding rules once again, without the need of full signaling procedure.

IV. Idle to Connected Procedure: UE Triggered and Network Triggered

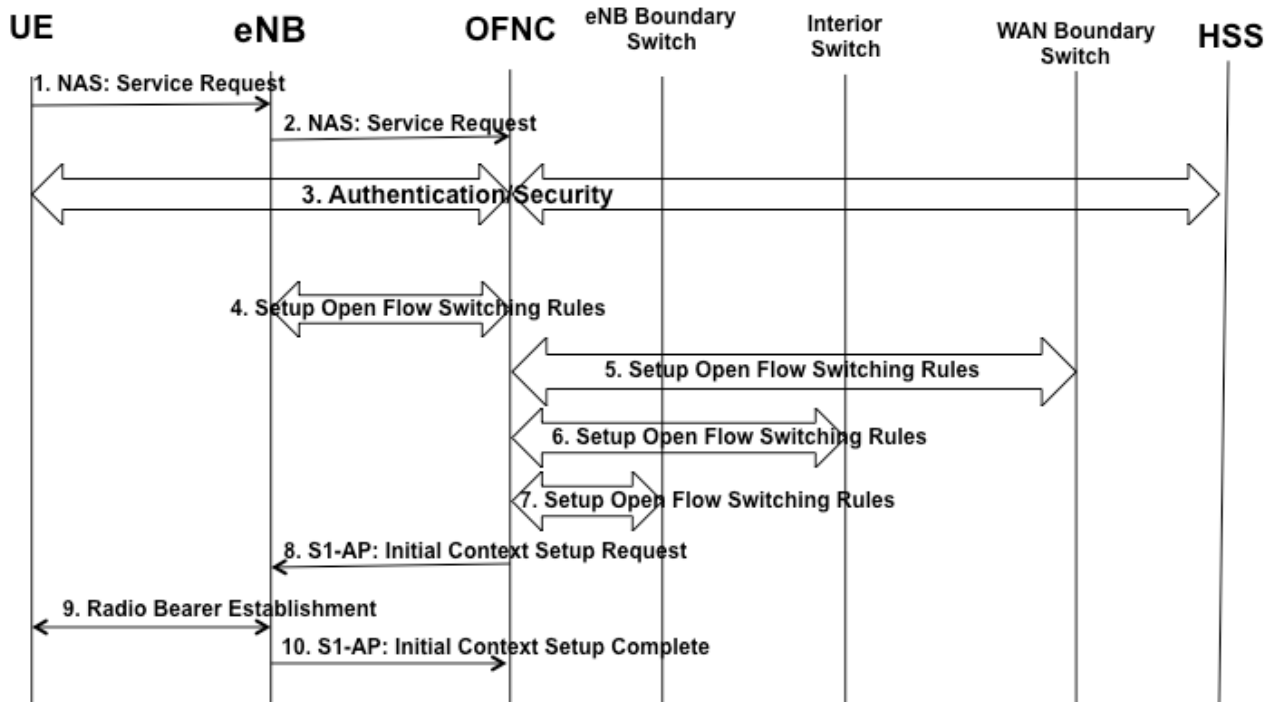


Figure 7: UE Triggered

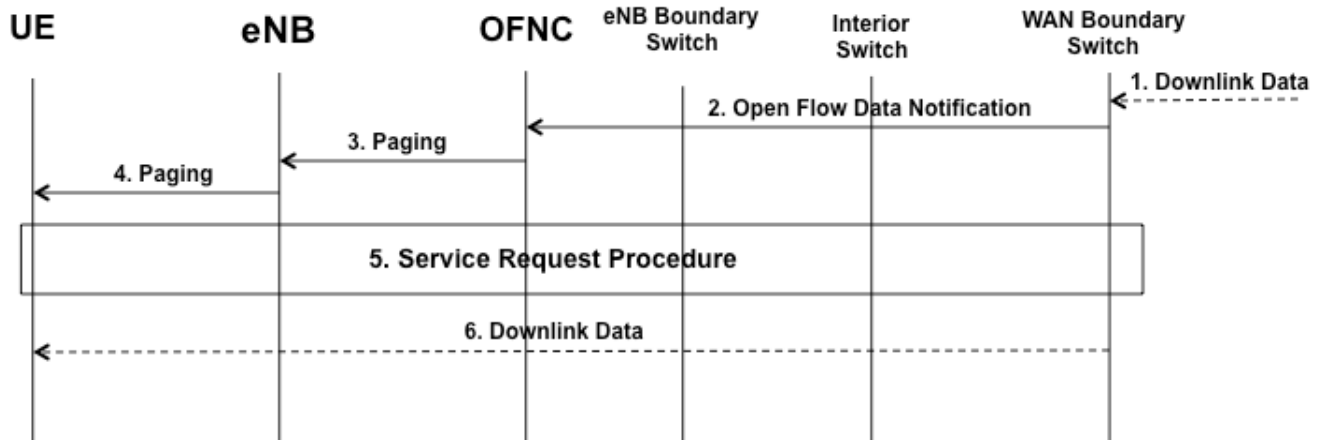


Figure 8: Network Triggered

The Idle to Connected procedure, for both the UE and Network Triggered cases, are shown in Figures 7 and 8. The main difference with the current EPC design is that the OFNC uses Open Flow procedures to recreate the paths for all the bearers in the backhaul network. Note that, as stated in the previous section, the OFNC retains the context information for the UE, which it then uses in this step to setup the path. In the Network Triggered case, the OFNC pages the UE using the S1-MME interface to the eNBs in the UEs Tracking Area (TA).

V. QoS: Dedicated Bearer Activation

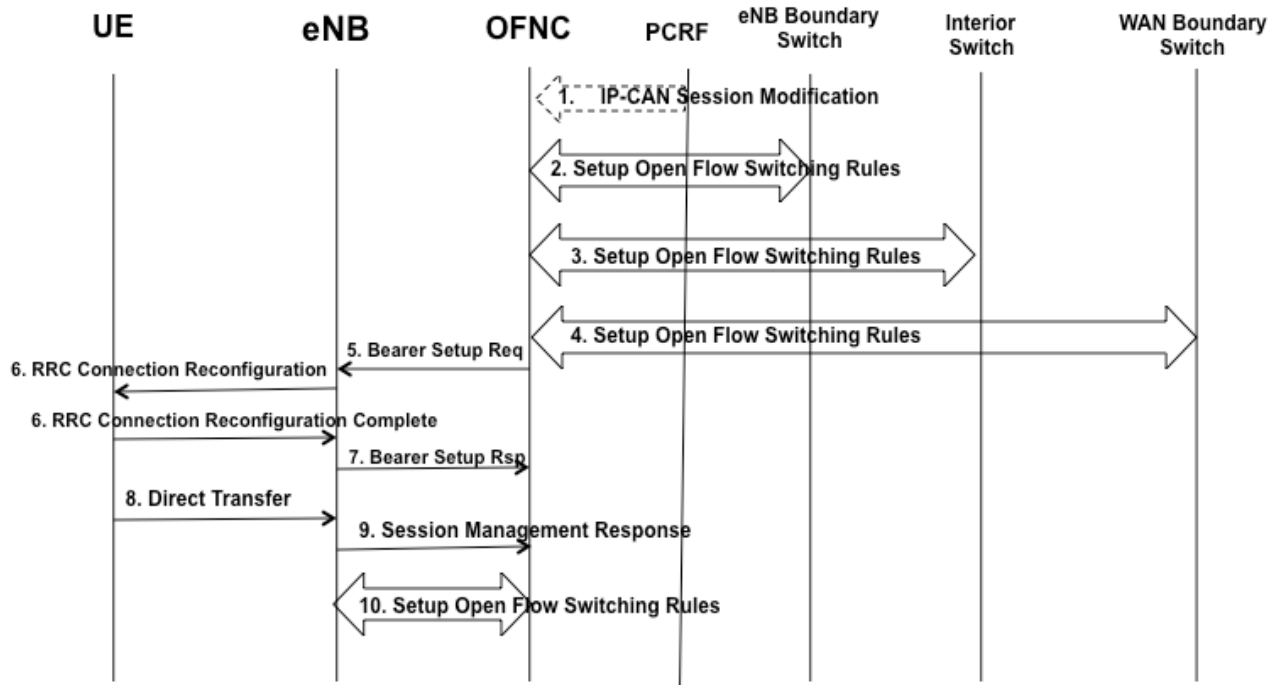


Figure 9

End to end support for QoS, encompassing Admission Control, Scheduling and BW Management, can be provided as follows:

- We will assume that each of the Open Flow switches in backhaul network support a DiffServ type QoS control, with a single Expedited Forwarding (EF) class, multiple Assured Forwarding (AF) classes and a single Best Effort (BE) class. The OFNC configures the amount of bandwidth that is allocated to the EF and AF classes at each node. We will also assume that the Open Flow switches support for Traffic Shaping and Policing, on a per flow basis.
- Figure 9 shows the control flows for setting up a new Dedicated Bearer, with QoS guarantees. In step 1, the PCRF passes the information about the classification rules and QoS parameters to be used for the bearer, including MBR (maximum bit rate), GBR (guaranteed bitrate) and Allocation and Retention Priority (ARP) values. The OFNC configures the MBR values in the Boundary Router, to control the shaping or policing functions for the bearer. It also keeps track of the available bandwidth in the requested traffic class (EF, AF1, AF2 etc) at each of the switches along the path, and does admission control. If the admission control succeeds, then it sets up the Open Flow switching rules for the bearer at each of the switches (steps 2, 3 and 4).
- In the remaining steps, the OFNC interacts with the eNB to setup the bearer across the radio link. In this process, the eNB does admission control across the radio link, and conveys the results to the OFNC (in step 7). The new bearer is created only if the admission control across both the backhaul and the radio links succeeds, and this information is sent back to the PCRF by the OFNC.

VI. Inter-Operability with WLAN

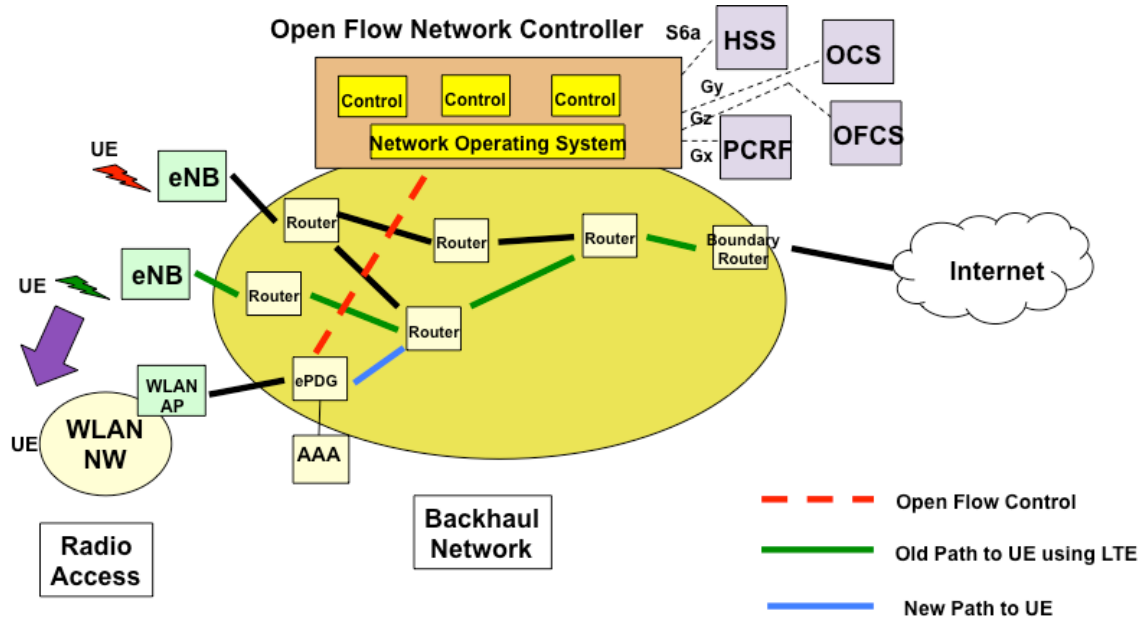


Figure 10

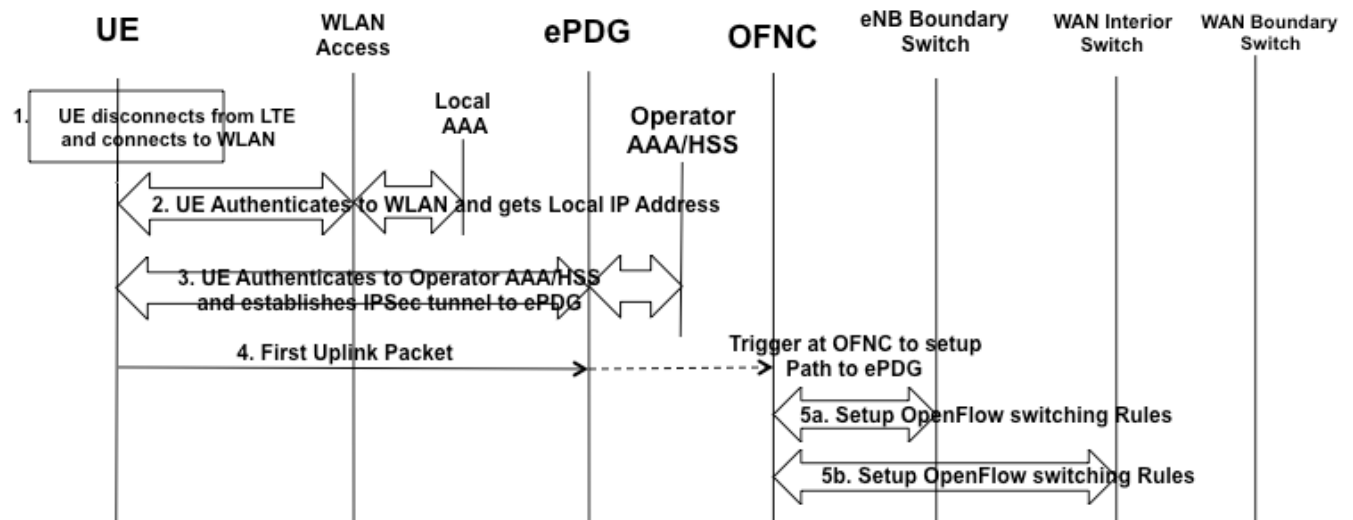


Figure 11

Figure 10 shows the network scenario in which an UE moves from the coverage of an LTE network to that of a WLAN. The operator deploys an Enhanced Packet Data Gateway (ePDG) node in order to enable inter-operability between LTE and WLAN. In this architecture, the ePDG is responsible for authenticating the UE and setting up secure IPsec tunnels with it. However, unlike the standard LTE design, the ePDG is not responsible for setting up GTP or MIP tunnels to an anchor PDN GW node. Instead, the OFNC sets up all paths within the backhaul network, using the Open Flow protocol.

Figure 11 shows some of the signaling flows involved when the user moves from LTE to WLAN. After synching with the WLAN network (step 1), the user is authenticated by the WLAN in step 2, and acquires a local IP address. It uses this IP address to communicate with the ePDG in the

operator's core network and setup the IPsec tunnels. During this process, the UE sends its LTE assigned IP address to the ePDG, so that it can be re-used. When the UE sends its first upstream packet to the ePDG, this triggers the ePDG to forward this packet to the OFNC since it doesn't have an installed switching rule for this packet. The OFNC in turn sets up the path between the ePDG and the Boundary Router that was hosting the UE in the LTE network (step 5), thus completing the network entry.

VII. Network Virtualization

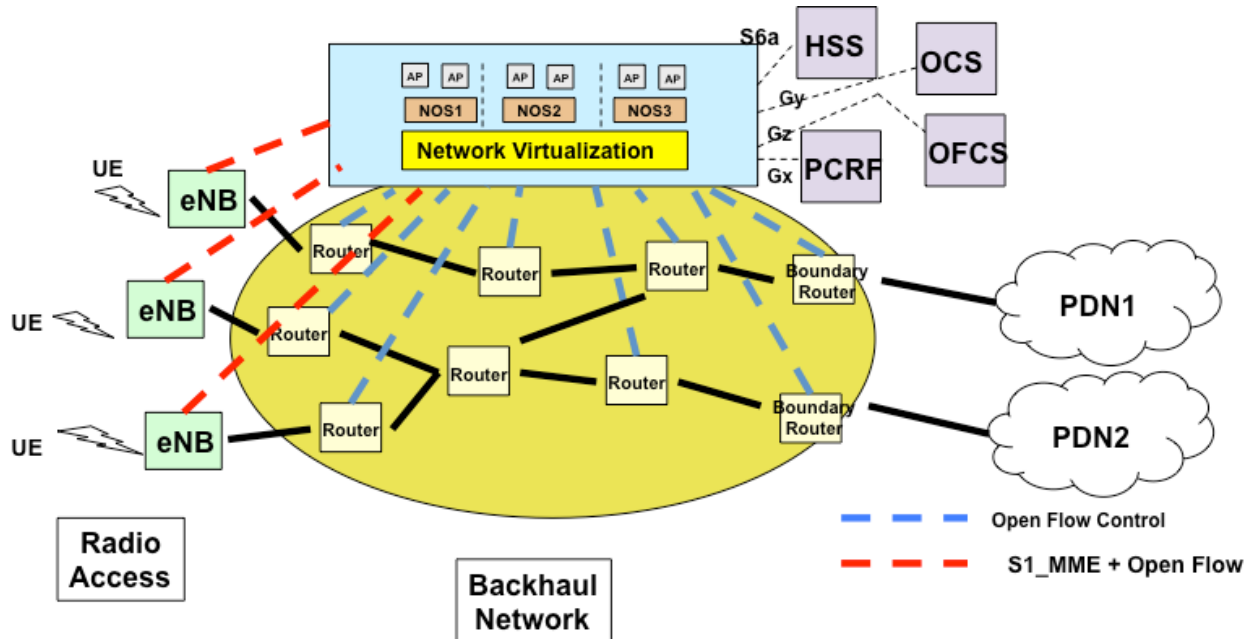


Figure 12

Figure 12 shows the OFNC with a Network Virtualization layer added, with multiple Network Operating Systems (NOS) on top. This allows the OFNC to provide a completely segregated network view to each NOS, which can then be used by the Applications using the NOS. Note that this type of virtualization is different than the node based virtualization found in current systems, since the latter does not provide a network wide virtual view of all the resources.

VIII. Support for Mobile Content Delivery Networks (CDN)

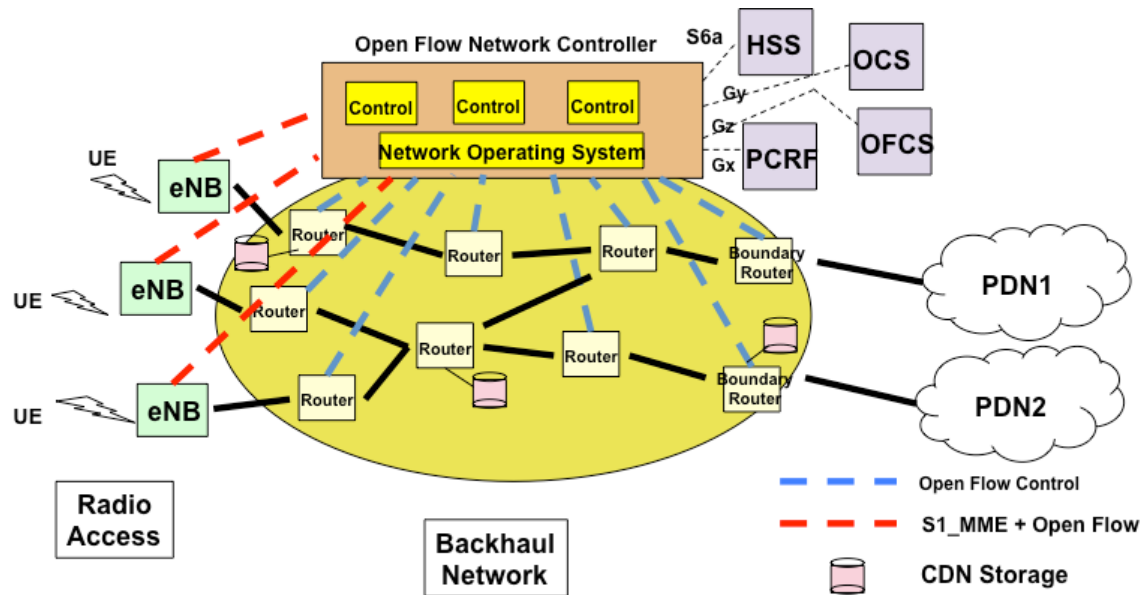


Figure 13

Mobile CDNs should allow the operator to cache frequently accessed content close to the UEs. In the current architecture (Figure 1), the closest that the content can be cached from the UE, is at the PDN GW. Frequently the PDN GW is not fully distributed, which limits the benefits of a Mobile CDN. With the Open Flow based architecture, this problem can be solved, as shown in Figure 13. It becomes possible to associate CDN storage nodes with the Open Flow switches in the backhaul network, and serve UEs through them. If the UE moves to another eNB, then the OFNC simply changes the routing for the flows from the CDN node to the UE, as part of the mobility control.