ERICSSON

EPG Architecture

TECHNICAL PRODUCT DESCRIPTION

13/221 02-AXB 250 12-V3 Uen DL

Contents

1 Introduction

This document describes the architecture of the EPG for GSM, WCDMA, LTE, trusted non-3GPP network, and untrusted non-3GPP network.

1.1 Scope

This document gives a high-level description of the EPG architecture and EPG application architectures.

1.2 Target Groups

This document is intended as an introduction to EPG architecture for network operators, network and service planners, system engineers, and system administrators. It can be used as a basis for training and assumes a basic knowledge of data communication and telecommunication.

2 Platform Architecture

The EPG application runs in the Network Functions Virtualization Infrastructure (NFVI). The NFVI consists of all hardware and software components which build up the environment in which the EPG is virtually deployed.

For general information on the EPG, refer to EPG Technical Product Description Overview.

Figure 1 shows an example of a hardware redundant virtual EPG deployment in a cloud environment with Virtual Service-Forwarders (vSFOs) and Virtual Route Processors (vRPs) on four compute hosts.

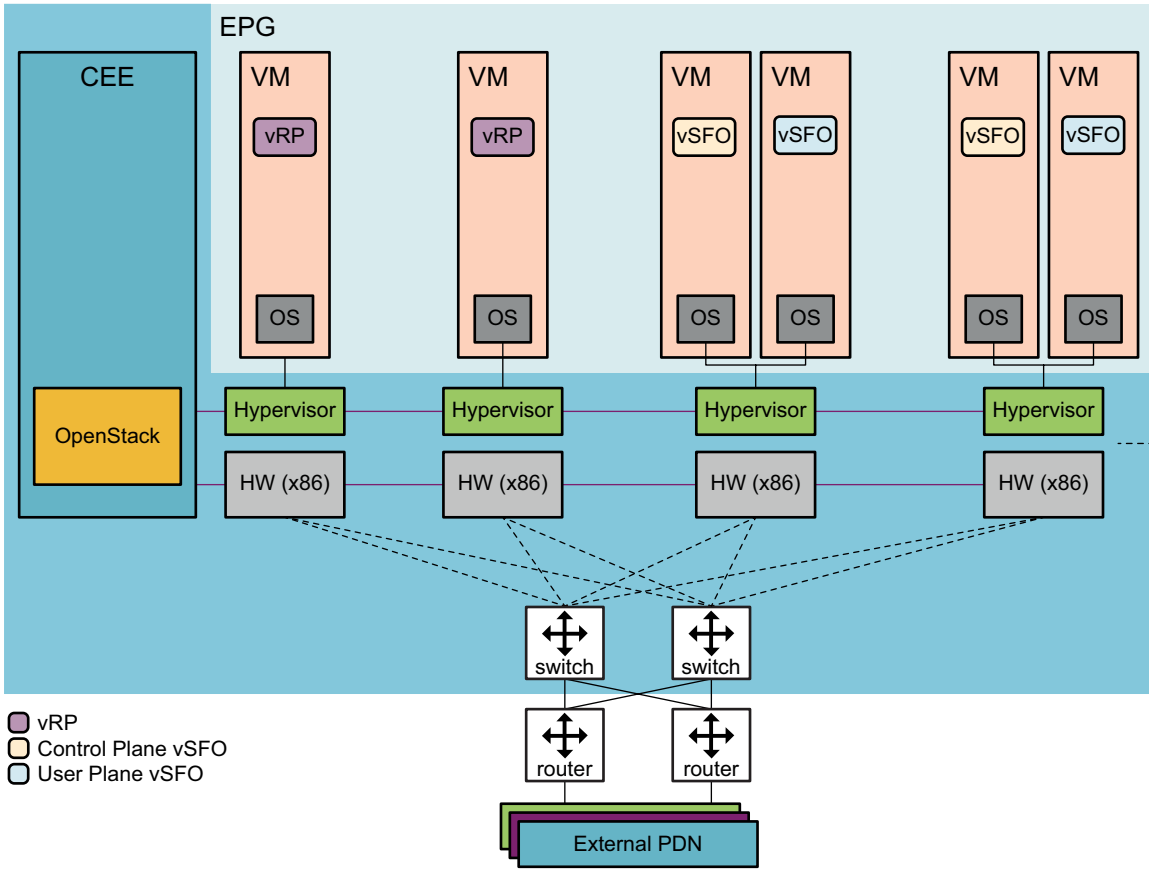

Figure 1 Virtual EPG Logical Architecture (vSFO)

2.1 Deployment Types

Virtual EPG is part of Ericsson virtual EPC (vEPC) Virtual Network Service (VNS) offerings, consisting of multiple Ericsson vEPC Virtual Network Functions (VNFs). vEPC deployments are not described in this document.

Virtual EPG can be deployed in the cloud system using either Heat or ECM. The virtual EPG supports any dynamic deployment that follows the rules specified in Deploying Virtual EPG.

For more information on how to configure the VMs and deploy the EPG, refer to Deploying Virtual EPG. For more information on scaling the EPG, refer to Scaling The EPG.

2.2 Virtual Machines

Virtual EPG is deployed in the cloud infrastructure as a cluster of VMs. For information on number of VMs, refer to Deploying Virtual EPG.

2.2.1 VM Types

The virtual EPG uses the following VM types:

— **Virtual Service-Forwarder (vSFO)**: A vSFO can be configured to provide the following roles:

- **User plane vSFO**: Provides virtual EPG application user plane capabilities for 2G/3G/LTE/WiFi/CDMA. User plane vSFO can also provide virtual EPG application aware load balancing (ingress) and forwarding (egress). It provides Layer 2 and Layer 3 data forwarding.

- **Control plane vSFO**: Provides virtual EPG application control plane capabilities for 2G/3G/LTE/WiFi/CDMA.

  Note: Control plane vSFOs can be configured with external IP interfaces for ingress and egress user packet forwarding, same as user plane vSFO. However, control plane vSFOs must not handle user plane traffic because of capacity impact. Therefore, Ericsson does not recommend to configure control plane vSFOs for user packet forwarding.

— **vRP**: The Virtual Route Processor (vRP) serves as a node manager for the virtual EPG application that performs cluster supervision, software distribution and provides O&M. There is one active vRP and one passive vRP for redundancy.

2.3 Cloud Environment

The virtual EPG is executed in a cloud environment. The following sections describe the components of a cloud environment.

2.3.1 Hypervisor

Kernel-based Virtual Machine (KVM) hypervisor is an open source software running on top of physical hardware. It is used to provide virtualization layer in Ericsson Cloud System (ECS). It supports live Motion, that is, moving a VM from one physical node to another with almost no disruption in service. In the KVM hypervisor, there is a software-based switch called virtual switch (vSwitch). A vSwitch allows one VM to communicate with other VMs. For more information on Hypervisor and vSwitch, refer to CEE documentation.

2.3.2 Hardware

Intel® x86 Compute hosts are used as hardware. For full details on the type of hardware used, refer to CEE documentation.

2.3.3 Cloud Execution Environment

Cloud Execution Environment (CEE) is a virtualization layer which manages the connection between application and hardware resources. CEE is an environment provided by the Ericsson cloud infrastructure containing hypervisors, virtual switches, system functions, and O&M support. For more information on CEE, refer to CEE documentation.

2.3.4 Ericsson Cloud Manager

Ericsson Cloud Manager (ECM) provides an integrated platform for managing a cloud computing infrastructure by enabling the creation, orchestration, activation and monitoring of services running on virtualized resources. It uses virtualization technology to abstract resources from the underlying physical hardware. For more information on ECM, refer to ECM documentation.

3 Networking

A Virtual Network (VN) is a logically isolated L2/L3 network provided by the cloud infrastructure. The VN can be realized as an overlay network using different encapsulation or tunneling technologies, such as, VLAN (802.1Q), VXLAN, or GRE on top of a physical underlying network. It is possible that an L2 overlay network spans multiple router hops in a routed L3 underlying network. The VN is transparent to the virtual EPG, as long as the characteristics required by the virtual EPG are provided.

The VNs used in the virtual EPG depend on how the virtual EPG is deployed. Figure 2 shows the logical view of the VNs if the virtual EPG is deployed in CEE or generic OpenStack with Heat using software version 2. For more information about software versions, refer to Deploying Virtual EPG.

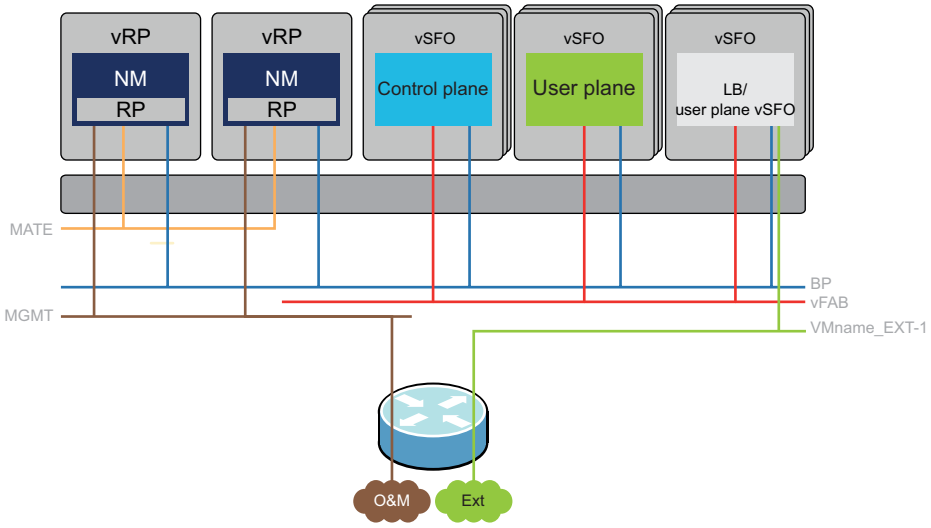

Figure 2 Logical View of VNs Deployed in CEE or Generic OpenStack with Heat, using Software Version 2.

Figure 3 shows the logical view of the VNs if the virtual EPG is deployed in either of the following:

— CEE with ECM

— CEE or Generic OpenStack with HEAT using software version 1

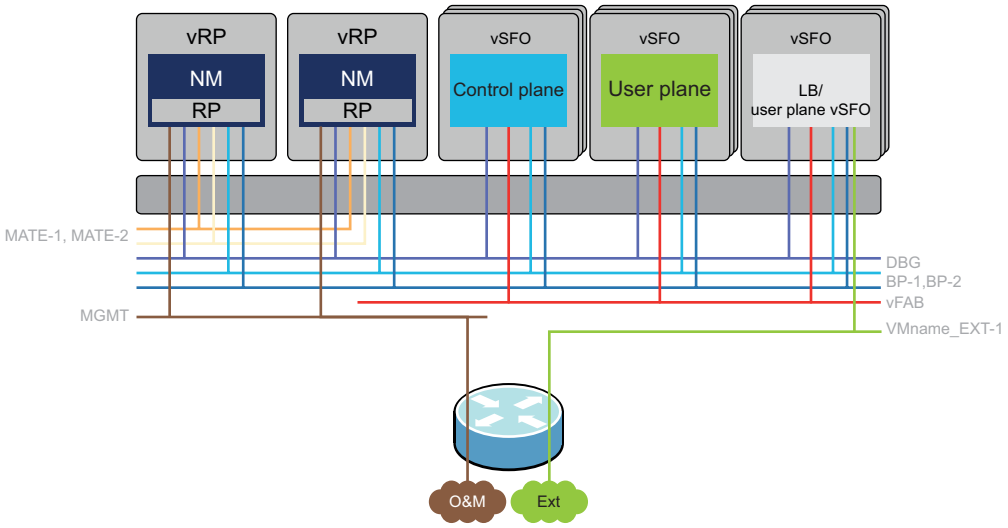

Figure 3 Logical View of the VNs Deployed in CEE with ECM, or in CEE or Generic OpenStack with Heat using Software Version 1

Table 1 summarizes the VNs used in the virtual EPG.

Table 1 Summary of VNs

VN Type	VN Name
Admin	MGMT
Internal	The virtual EPG uses the following internal VNs for deployments in CEE or generic OpenStack with Heat using software version 2: MATE, BP, vFAB The virtual EPG uses the following internal VNs for deployments in CEE with ECM or for deployments in CEE or generic OpenStack with Heat using software version 1: MATE-1, MATE-2, BP-1, BP-2, vFAB, DBG
External	EXT-x, SC-x

3.1 Admin Networking

The Admin interface on the RP connects to the MGMT VN. The Admin interface is used for temporary outband O&M access, using SSH, SCP, and SFTP when the normal O&M IP service through the vSFO is either not configured or lost. The temporary O&M access is used at initial configuration and in some emergency situations. The Admin interface does not replace the normal O&M IP service through the vSFO.

3.2 Internal Networking

The MATE VN or MATE-1 and MATE-2 VNs are used for synchronization and monitoring traffic between active and hot standby vRP.

The BP VN or BP-1 and BP-2 VNs are used for internal signalling between all virtual EPG VMs.

The vFAB VN is used for forwarding external control signalling and user data traffic between vSFOs.

The DBG VN is reserved to use for internal debugging. The DBG VN is not reserved if the virtual EPG is deployed in CEE or generic OpenStack with Heat using software version 2.

3.3 External Networking

The external VNs are used for the virtual EPG external traffic, exchanged through the vSFOs. The vSFO vNICs connected to the External VNs are also referred to as vSFO ports.

For more information on external networking, refer to Virtual EPG External Network Connectivity Configuration.

For more information about SGi-LAN service chaining networking, refer to Service Chaining.

3.4 Virtual Networks

Table 2 describes the purpose of the VN and if the VN can be configured as a vRP or a vSFO.

Table 2 Characteristics of VNs

VN/VM Name	vRP	vSFO	Purpose
MGMT	Yes	No	VNF-external Outband O&M
MATE(1)	Yes	No	VNF-Internal: signalling between vRPs
MATE-1(2)	Yes	No	VNF-Internal: signalling between vRPs
MATE-2(2)	Yes	No	VNF-Internal: signalling between vRPs
BP(1)	Yes	Yes	VNF-Internal: signalling between VMs
BP-1(2)	Yes	Yes	VNF-Internal: signalling between VMs
BP-2(2)	Yes	Yes	VNF-Internal: signalling between VMs

VN/VM Name	vRP	vSFO	Purpose
VFAB	No	Yes	VNF-Internal: forwarding external signalling and user data traffic between vSFOs
DBG(2)	Yes	Yes	For internal debugging Not used. Reserved for future use
EXT-x	No	Yes	For external VNF control and payload communication. The virtual EPG supports up to 6 external interfaces per VM if the OVF package is generated for VMWare. The virtual EPG supports up to 8 external interfaces per VM if the OVF package or HOT file is generated for ECM or Heat. For more information, refer to Deploying Virtual EPG.
SC-x	No	Yes	For external SGi-LAN service chaining monitoring and payload communication.

(1) Used if the virtual EPG is deployed in CEE or generic OpenStack with Heat using software version 2.
(2) Not used if the virtual EPG is deployed in CEE or generic OpenStack with Heat using software version 2

The virtual EPG requires that the VM virtual Network Interface Cards (vNICs) and VNs are connected in a specified order. The order depends on how the virtual EPG is deployed.

Table 3 specifies the required connections between VM vNICs and VNs if the virtual EPG is deployed in CEE or generic OpenStack with Heat using software version 2.

Table 3 Connection between VM vNICs and VNs in CEE or Generic OpenStack Deployments Using Software Version 2

VM Type	vNIC 1	vNIC 2	vNIC 3
vRP	MGMT	BP	MATE
vSFO	BP	vFAB	-

Table 4 specifies the required connections between VM vNICs and VNs if the virtual EPG is deployed in CEE with ECM, or in CEE or generic OpenStack with software version 1.

Table 4 Connection between VM vNICs and VNs in CEE with ECM Deployments, or in CEE or Generic OpenStack Deployments Using Software Version 1

VM Type	vNIC 1	vNIC 2	vNIC 3	vNIC 4	vNIC 5	vNIC 6
vRP	MGMT	BP-1	BP-2	MATE-1	MATE-2	DBG
vSFO	BP-1	BP-2	vFAB	DBG	-	-

4 Storage

The virtual EPG supports local storage by using the physical disk on the compute host where the VM is running. For information on storage, refer to Managing Files.

5 Non-Uniform Memory Access

Non-Uniform Memory Access (NUMA) is common in CPU architectures. It means that the physical CPU and the memory are partitioned into multiple NUMA nodes. For example, it is common that a dual socket host has two NUMA nodes, one per socket. A CPU core has fast access to the memory within its own NUMA node (local memory), but slower access to the memory in other NUMA nodes (non-local memory).

The virtual EPG guest OS is NUMA-aware and the virtual EPG application uses resources accordingly. However, there is a significant capacity penalty and a complicated integration procedure when user-plane VMs (vSFO as user plane) are deployed across NUMA boundaries. Control plane vSFO can span two or more NUMA nodes.

It is recommended to ensure that all vCPUs and the memory of user plane vSFOs are mapped to physical CPUs and memory on a single NUMA node. In addition, for maximum capacity of the vSwitch, the vSwitch threads and the physical NIC need to be on the same NUMA node as the user plane vSFOs. A reference deployment for user plane vSFOs following this recommendation is called the compact deployment, described in the virtual EPC documentation.

6 Application Architecture

The EPG application software consists of control plane and user plane applications running on the vSFOs. An allocation of the application types on the vSFOs takes place during application startup. Based on configuration, each vSFO is assigned as control plane or user plane . For a description of the startup procedure, see Section 6.1 on page 17.

6.1 Startup Procedure

The startup procedure of the EPG application software on the vSFOs takes place as follows:

1. The EPG application software is loaded on the vSFOs during the deployment of EPG. For more information on deploying EPG, refer to Deploying Virtual EPG.

2. Once mandatory configuration data and vSFOs are in place, a start is initiated for the vSFOs. The mandatory configuration data shown in Section 6.1.1 on page 17 is retrieved.

3. The retrieved configuration data is used to determine which vSFOs run each of the application types. Roles of vSFOs are determined and applications are started with the needed configuration data.

6.1.1 Configuration Retrieval

During the startup phase of the control plane applications, the EPG configuration information is retrieved. This configuration information includes vSFO slot positions and the number of active and standby cards for each of the application types. The EPG configuration information also includes single IP addresses or IP address ranges and routing instance associations used for external interface communication between the EPG and surrounding nodes.

6.2 Control Plane Applications

Control plane applications on the vSFO consist of the following applications:

— Global Session Controller (GSC)

— PGW Session Controller (PSC)

— SGW Session Controller (SSC)

For more information on the EPG control plane functionality, refer to Session Management.

The EPG has multicore architecture to enable load distribution. On a single control plane vSFO multiple SSC and PSC instances are running on separate cores. Each SSC/PSC instance is responsible for a dedicated group of subscribers.

The following sections describe the functions of the GSC, PSC, and SGW Session Controller applications.

6.2.1 GSC

The following are functions of the GSC:

— Route handling for APNs

— Shared IP pool functionality

The GSC can run on any CPB control plane vSFO, but only on one CPB control plane vSFO at a time.

6.2.2 PSC

The PSC takes the role of a session controller on either the GGSN or the PGW.

When the PSC is configured for the GGSN, the GGSN establishes user sessions through GGSN-configured APNs. When the PSC is configured for the PGW, the PGW establishes user sessions through PGW-configured APNs.

The PSC establishes and controls connections towards the APN networks by processing GTP-C messages received on the Gn/Gp, S5/S8-C, GTP-based S2a, and S2b interfaces. The PSC also establishes and controls connections towards external supporting nodes, for example, the Online Charging System (OCS) by processing Diameter Based Protocol (DBP) messages received on the Gy+ interface and the Event-Based Monitoring (EBM) server by streaming data through the EBM logical interface.

The PSC also handles RADIUS and Dynamic Host Configuration Protocol (DHCP) communication through the Gi/SGi interface, performance monitoring statistics gathering, charging data generation for GGSN/PGW Charging Data Records (CDRs), and forwarding of CDRs through the Ga interface.

The GGSN and PGW communicate with surrounding nodes using single IP addresses assigned to each logical interface configured for the GGSN and PGW control plane networks. All PSCs share a single IP address per network and the GGSN and PGW communicate each logical interface IP address to the surrounding nodes in the same network through signalling messages. For information on GGSN and PGW control plane networks with single IP addresses, see Section 6.4 on page 21.

For information on the port range used by the multicore architecture on the GGSN/PGW control plane interfaces, refer to Routing.

6.2.3 SGW Session Controller

The SGW Session Controller takes the role of a session controller on the SGW.

The SGW Session Controller establishes and controls connections by processing GTP-C messages received on the S4-C, S5/S8-C, and S11 interfaces. The SGW Session Controller also establishes and controls connections towards external supporting nodes, for example, the Charging Data Function (CDF) by processing DBP messages received on the Rf interface and the EBM server by streaming data through the EBM logical interface.

The SGW Session Controller handles charging data generation for SGW CDRs.

The SGW communicates with surrounding nodes using single IP addresses assigned to each logical interface configured for the SGW control plane networks. All SGW Session Controllers share a single IP address per network and the SGW communicates each logical interface IP address to the surrounding nodes in the same network through signalling messages. For information on SGW control plane networks with single IP addresses, see Section 6.4 on page 21.

For information on the port range used by the multicore architecture on the SGW control plane interfaces, refer to Routing.

6.3 User Plane Applications

User plane applications on the vSFO consist of the following applications:

— Packet Processor (PP)

— L2TP Packet Processors (TPP)

For more information on the EPG user plane functionality, refer to User Packet Handling.

The following sections describe the functions of the PP and TPP applications.

6.3.1 PP

The PP takes the role of a user plane application on the GGSN, PGW, and SGW. The PP handles the high-speed, real-time forwarding of user packets in the uplink and downlink directions.

6.3.1.1 GGSN User Plane

The PP on the GGSN user plane is used for user sessions established through GGSN-configured APNs. The PP processes user packets for PDP contexts established over the Gn/Gp network. All PPs share a single IP address per GGSN user plane network. For information on GGSN user plane networks with single IP addresses, see Section 6.4 on page 21.

The PP also handles uplink and downlink user packets over the Gi interface and performs functions such as administration of GTP-U tunnels, using received PDP context data for setting of the Differentiated Services Code Point (DSCP) in the packet IP headers, transferring downlink Gi user packets into the correct GTP-U tunnels, filtering, and bit rate enforcement. The PP also gathers usage information on uplink and downlink traffic, for both statistics and charging. For a more detailed description, refer to User Packet Handling and Quality of Service on the GGSN and PGW.

If the Service Aware Charging and Control (SACC) solution is activated, the PP is responsible for packet inspection, service classification, access control, charging control, credit control, and so on to facilitate differentiated charging. For more information on SACC, refer to SACC Overview.

6.3.1.2 PGW User Plane

The PP on the PGW user plane is used for user sessions established through GGSN-configured APNs and PGW-configured APNs. The PP processes user packets for PDP contexts established over the Gn/Gp network and for EPS bearers established over the S5/S8-U/S2a/S2b network. All PPs share a single IP address per PGW user plane network. For information on PGW user plane networks with single IP addresses, see Section 6.4 on page 21.

The PP on the PGW user plane can also be used for PDN connections towards a Mobile Access Gateway (MAG) in a CDMA2000 network.

The PP handles uplink and downlink user packets over the SGi interface and performs functions such as DSCP setting, filtering, and bit rate enforcement. If the SACC solution is activated, the PP can use PISC to identify UE services, thus facilitating differentiated charging. The PP also gathers usage information on uplink and downlink traffic, for both statistics and charging.

The PP establishes GTP-U and Generic Routing Encapsulation (GRE) tunnels carrying user packets over the Gn/Gp/S5/S8-U/S2a/S2b and PMIPv6-based S2a networks, in coordination with the GSC and PSCs.

It is possible to enable dynamic GTP-U UDP source ports for the PGW on the Gn/Gp, S5/S8, GTP-based S2a, and S2b interfaces, meaning that all new established user sessions are allocated a port from a specified range. The port range is configurable, with a default range of 49152-65535. For information on how to configure and enable dynamic source ports, refer to GTP Interface Configuration.

For a more detailed description of PGW user packet handling and quality of service functionality, refer to User Packet Handling and Quality of Service on the GGSN and PGW. For a more information on PISC, refer to Packet Inspection and Service Classification (PISC).

6.3.1.3 SGW User Plane

The PP on the SGW user plane establishes bearers for the S1-U/S4-U/S12 and S5/S8-U networks, using GTP-U tunnels carrying the user packets, in coordination with SGW Session Controllers.

The SGW communicates with surrounding nodes using single IP addresses assigned to each logical interface configured for the SGW user plane networks. All PPs share a single IP address per network. For information on SGW user plane networks with single IP addresses, see Section 6.4 on page 21.

It is possible to enable dynamic GTP-U UDP source ports for the SGW on the S1-U, S4-U, and S12 interfaces, and on the S5/S8-U interface, meaning that all new established PDN connections are allocated a port from a specified range. The port range is configurable, with a default range of 49152-65535. For information on how to configure and enable dynamic source ports, refer to GTP Interface Configuration.

For a more detailed description of the SGW user packet handling, refer to User Packet Handling.

6.3.2 TPP

The TPP on the GGSN or the PGW is used for L2TP tunneling carrying Point-to-Point Protocol (PPP) sessions. L2TP functionality is provided using the L2TP Access Concentrator (LAC) – L2TP Network Server (LNS) tunneling mode, with the GGSN or the PGW acting as an LAC to generate PPP sessions from user sessions.

The GGSN/PGW communicates with the LNS using unique IP addresses assigned to each TPP from the configured LAC IP address range.

The TPP runs in combination with the PP on the user plane vSFO.

For more information, refer to Layer Two Tunneling Protocol (L2TP).

6.4 EPG Deployments

The EPG can be run in the following deployment types:

— Standalone SGW

— Standalone GGSN/PGW

— Combined SGW and PGW

For more information about SGi-LAN service chaining networking, refer to Service Chaining.

6.4.1 Standalone SGW

The EPG can be run as a standalone SGW for handling SGW-only user sessions.

During startup, all the vSFOs configured as control plane vSFOs take on the role of SGW Session Controller. vSFOs configured as user plane take on the role of PPs. At startup, all the boards are active and ready to handle session replication, thereby providing session resilience on both control and user planes. For more information on session resilience for the standalone SGW, refer to Resilience. For more information on how to configure boards, refer to EPG Board Configuration.

An SGW provides support for the S1, S4, S5/S8, S11, and S12 interfaces according to 3GPP standards. An SGW also supports the Rf interface, the proprietary EBM logical interface, and the proprietary Gom interface. Figure 4 shows the applications and network IP addresses of a standalone SGW.

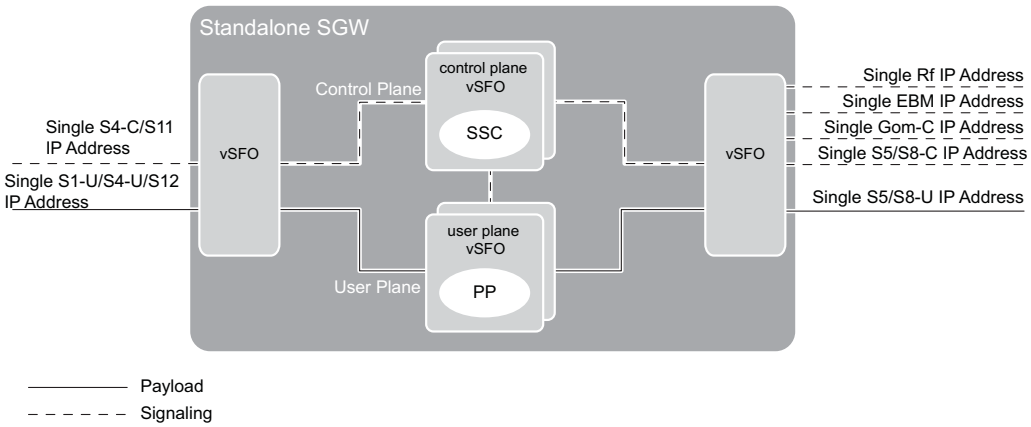

Figure 4 Standalone SGW with Supported Network IP Addresses

The minimum configuration for standalone SGW is as follows:

— Two control plane vSFOs running SGW Session Controllers

— Two user plane vSFOs running PPs

The standalone SGW supports a single IP address on each network interface for handling control plane or user plane communication to and from external networks. Internally, all the boards involved in communication on one interface share the same external IP address. This way, the external nodes are only aware of one SGW IP address on each interface, simplifying configuration for surrounding nodes.

Table 5 summarizes the architecture of the standalone SGW.

Table 5 Standalone SGW

Application Type	VM Role	Application	IP Addressing
Control plane	vSFO	SGW Session Controller	S4-C/S11 network IP address S5/S8-C network IP address EBM network IP address Rf network IP address
User plane	vSFO	PP	S1/S4-U/S11-U /S12 network IP address S5/S8-U network IP address

For more information about SGW networks, refer to Routing. For information on configuration of IP addresses, refer to GTP Interface Configuration for S4-C/S11, S5/S8-C, S1-U/S4-U/S11-U/S12, and S5/S8-U networks, EPG Board Configuration for the Gom network, Event-Based Monitoring Configuration for the EBM network, and Diameter Configuration for the Rf network.

6.4.2 Standalone GGSN/PGW

The EPG can be run as a standalone GGSN for handling GGSN-only user sessions or as a standalone PGW for handling GGSN-only and PGW-only user sessions.

During startup, vSFOs configured as control plane vSFOs run either GSC and PSC, or PSC. vSFOs configured as user plane vSFOs respectively take on the role of PPs/TPPs.

At startup, all the boards are active and ready to handle session replication, thereby providing session resilience on both the control plane and the user plane. For more information on session resilience for the standalone GGSN/PGW, refer to Resilience. For more information on how to configure boards, refer to EPG Board Configuration.

A GGSN provides support for the Gn/Gp, Iu for 3G Direct Tunnel (3GDT), Ga, Gx, Gy, and Gi interfaces according to GPRS 3GPP standards. A PGW provides support for the Gn/Gp, Iu-U, S5/S8, GTP-based S2a, PMIPv6-based S2a, S2b, Ga, Gx, Gy, S6b, S6bAuth, and SGi interfaces. A GGSN/PGW also supports the proprietary EBM logical interface, the proprietary Gom interface, and the ICAP interface. A GGSN or PGW, acting as an LAC also provides support for L2TP services.

Page 24 shows the applications and network IP addresses and ranges of a standalone GGSN/PGW.

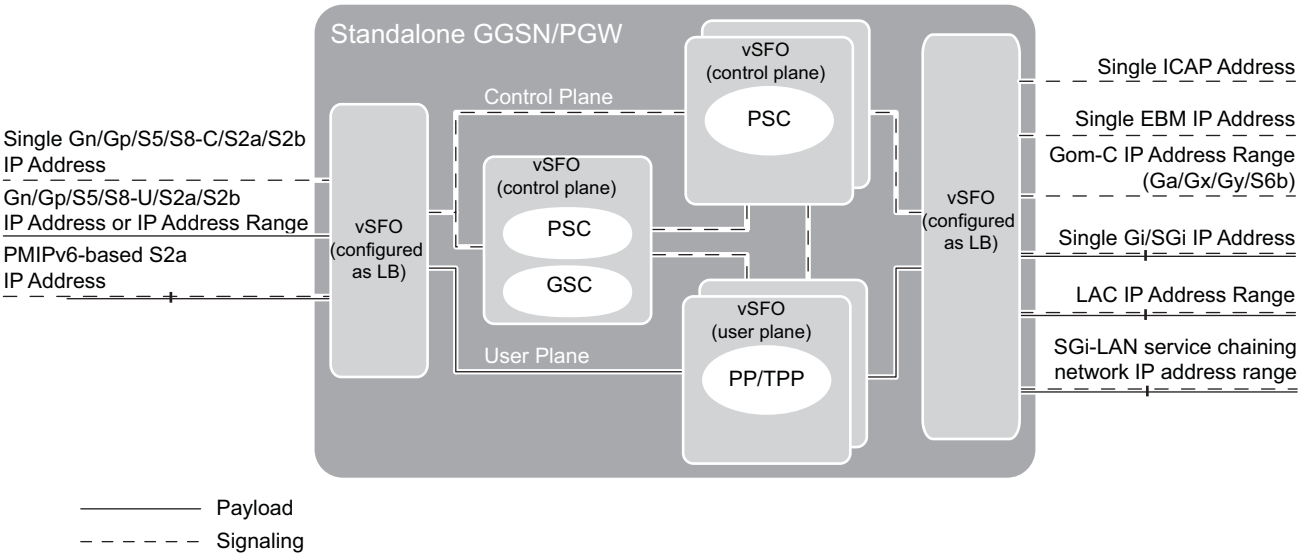

Figure 5 Standalone GGSN/PGW with Supported Network IP Addresses and Ranges

ICAP clients run on all control plane vSFOs, that is boards running GSC and PSC, and PSC.

The minimum configuration for the standalone GGSN/PGW is as follows:

— One control plane vSFO running GSC and PSC

— One CPB running PSC

— Two user plane vSFOs, each running the PP/TPP

The standalone GGSN/PGW supports the configuration of IP addressing as follows:

— Single IP address or IP address range on the user plane network interfaces.

Note: If an IP address range is configured on the PGW user plane logical interface, the number of IP addresses is independent to the number of user plane vSFOs. The sessions are given IP addresses spread over the full configured range. A maximum of 32 IP addresses can be used from the IP address range. For example, configuring a /28 IP address range results in sessions being distributed across all 14 different IP addresses, even if only two user plane vSFOs are configured.

— Single IP address on the control plane network interfaces.

Table 6 summarizes the architecture of the standalone GGSN/PGW.

Table 6 Standalone GGSN/PGW

Plane	VM Role	Application	IP Addressing
Control plane	vSFO	GSC PSC	Gn/Gp/S5/S8-C/ S2a/S2b network IP address(1) Gi/SGi network IP address Gom network IP address PMIPv6-based S2a network IP address EBM network IP address ICAP network IP address
User plane	vSFO	PP	Gn/Gp/S5/S8-U/ S2a/S2b network IP address or IP address range(2) Gi/SGi network IP address range PMIPv6-based S2a network IP address SGi-LAN service chaining network IP address range
		TPP	LAC IP address range

(1) The Gn/Gp-C and S5/S8-C interfaces belong to the same network, but are configured separately. For more information, see GTP Interface Configuration.
(2) The Gn/Gp-U and S5/S8-U interfaces belong to the same network, but are configured separately. For more information, see GTP Interface Configuration.

For more information about GGSN/PGW networks, refer to Routing. For information on configuration of IP address ranges, refer to GTP Interface Configuration for the Gn/Gp/S5/S8/S2a/S2b networks, APN Configuration for the Gi/SGi network, EPG Board Configuration for the Gom network, Event-Based Monitoring Configuration for the EBM network, PMIPv6-Based S2a Interface Configuration for the PMIPv6-based S2a network, L2TP Configuration for the LAC IP address range, and Content Filtering Configuration for the ICAP network.

6.4.3 Combined SGW and PGW

The EPG can handle the following user session types when running as a combined node:

— SGW-only user sessions

— GGSN-only and PGW-only user sessions

— Combined SGW and PGW user sessions

During startup, vSFOs configured as control plane vSFOs run either GSC, PSC and SGW Session controllers, or PSC and SGW Session Controllers. Similarly, vSFOs configured as user plane vSFOs take on the role of active PPs/TPPs.

At startup, all the boards are active and ready to handle session replication, thereby providing session resilience on both the control plane and the user plane. For more information on session resilience for the combined SGW and PGW, refer to Resilience. For more information on how to configure boards, refer to EPG Board Configuration.

Figure 6 shows the applications and network IP addresses and ranges of a combined SGW and PGW, including 3GPP interfaces, the proprietary Gom interface, the proprietary EBM interface, the ICAP interface, and LAC.

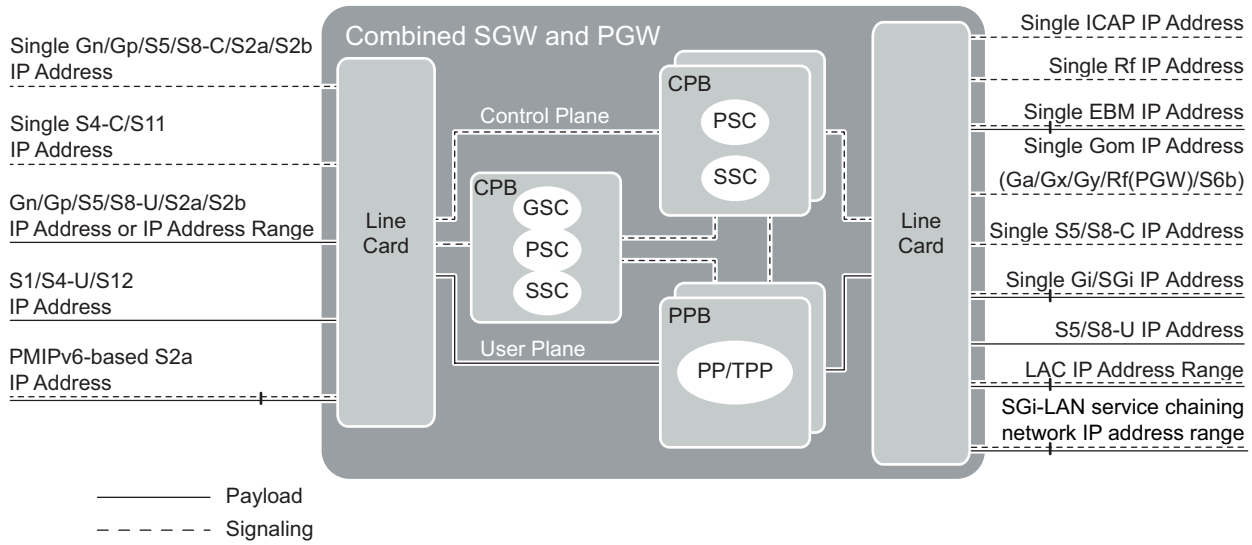

Figure 6 Combined SGW and PGW with Supported Network IP Addresses and Ranges

ICAP clients run on all available control plane vSFOs.

The minimum configuration for the combined SGW and PGW is as follows:

— One control plane vSFO running GSC, PSC and SGW Session Controller

— One control plane vSFO running PSC and SGW Session Controller

— Two user plane vSFOs, each running the PGW PP/TPP and the SGW PP

The combined SGW and PGW supports the configuration of IP addressing as shown in the Table 7.

Table 7 IP Addressing

Application Type	IP Addressing on SGW	IP Addressing on PGW
Control plane	Single IP address	Single IP address
User plane	Single IP address	Single IP address or IP address range(1)

(1) If an IP address range is configured on the PGW user plane logical interface, the number of IP addresses is independent to the number of user plane vSFOs. The sessions are given IP addresses spread over the full configured range. A maximum of 32 IP addresses can be used from the IP address range. For example, configuring a /28 IP address range results in sessions being distributed across all 14 different IP addresses, even if only two user plane vSFOs are configured.

Note: Although the SGW and PGW are on the same node, the SGW communicates with the PGW through the S5/S8 logical interfaces.

Table 8 summarizes the applications that can run on the combined SGW and PGW.

Table 8 Combined SGW and PGW

Application Type	VM Role	Application	IP Addressing
Control plane	vSFO	GSC PSC SGW Session Controller	Gn/Gp/S5/S8-C/ S2a/S2b network IP address(1) S4-C/S11 network IP address SGW S5/S8-C network IP address Gi/SGi network IP address Gom network IP address(2) EBM network IP address Rf network IP address(3) PMIPv6-based S2a network IP address ICAP network IP address

Application Type	VM Role	Application	IP Addressing
User plane	vSFO	PP	Gn/Gp/S5/S8-U/ S2a/S2b network IP address or IP address range(4) S1/S4-U/S12 network IP address SGW S5/S8-U network IP address Gi/SGi network IP address range PMIPv6-based S2a network IP address SGi-LAN service chaining network IP address range
		TPP	LAC IP address range

(1) The Gn/Gp-C and S5/S8-C interfaces belong to the same network, but are configured separately. For more information, see GTP Interface Configuration.
(2) The address range is shared by the PGW and SGW.
(3) This is only applicable for the SGW.
(4) The Gn/Gp-U and S5/S8-U interfaces belong to the same network, but are configured separately. For more information, see GTP Interface Configuration.

For more information about PGW and SGW networks, refer to Routing. For information on configuration of IP addresses and address ranges, refer to GTP Interface Configuration for the GTP-based networks, APN Configuration for the Gi/SGi network, EPG Board Configuration for the Gom network, Event-Based Monitoring Configuration for the EBM network, Diameter Configuration for the Rf network, PMIPv6-Based S2a Interface Configuration for the PMIPv6-based S2a network, L2TP Configuration for the LAC IP address range, and Content Filtering Configuration for the ICAP network.

Reference List

Online Documentation

[1] IANA - Port Numbers, http://www.iana.org/assignments/port-numbers