Virtual BNG on Public Cloud for Disaster Recovery and Capacity on Demand

A Technical Paper prepared for SCTE by

Jim Huang
Principal Solutions Architect
Amazon Web Services
jimhuan@amazon.com

Juan Luis Esteban
Director, Broadband and Convergence Solutions
Ciena
jesteban@ciena.com

Christopher Adigun, Telco Solutions Architect, Amazon Web Services

© 2023, SCTE®, CableLabs® and, NCTA. All rights reserved.

Table of Contents

List of Figures

List of Tables

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

1. Introduction

Broadband Network Gateway (BNG) is a critical network component in broadband and Internet Service Provider networks. Its primary function is to manage and control the delivery of broadband services, such as internet access, to subscribers. The BNG technology has been evolving over the past two decades from specialized hardware devices to virtual BNG (vBNG) and further to disaggregated BNG (DBNG) with control plane and user plane separation (CUPS) as specified by Broadband Forum Technical Report TR-459[1]. The vBNG technology was driven by the need for more flexible solutions and higher operation efficiency through virtualizing BNG functions and running them as software on industry-standard hardware or virtual machines. The DBNG technology allows service providers to select and deploy individual components, instead of relying on monolithic integrated solutions. It enables service providers to optimize their network architecture and scale each component independently based on demand, allocating resources where they are needed the most and avoiding overprovisioning.

This paper presents BNG cloudification – a broadband technology frontier built on vBNG and DBNG. Today, public cloud has been widely adopted by enterprises and industry verticals including telecommunications due to its economies of scale and capabilities for gaining operation efficiency, enabling business agility, as well as providing a wealth of latest technology services. The BNG cloudification presented in this paper is about running vBNG control plane and vBNG user plane in the public cloud.

The key use cases driving the BNG cloudification are the following, as illustrated in Figure 1.

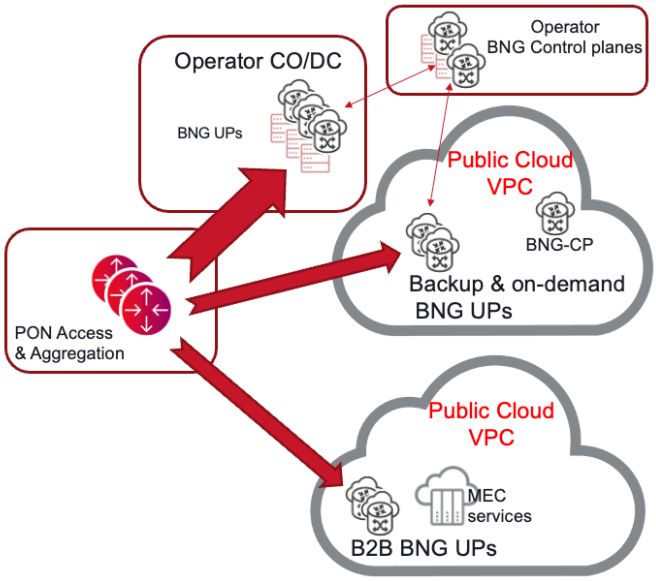

Figure 1 - BNG Cloudification Use Cases

- *Disaster Recovery* – ability to redirect subscriber traffic from DBNG-UPs in the operator network to backup DBNG-UPs in a public cloud. In current on-premise environments, DBNG user planes are deployed in geo-redundant locations to enable disaster recovery. This requires allocation of extra capacity, power, and rack space in each location for the failure case. Allocating disaster recovery capacity in the public cloud allows the operator to reduce CapEx investment.

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

- *Capacity on demand* – ability to steer traffic to DBNG-UP instances in public cloud based on traffic spike conditions, which may occur daily or seasonally. It leverages the public cloud elasticity to handling bursty workloads instead of building extra capacity in operator data centers, hence lowering CapEx investment.
- *Multi-access Edge Compute (MEC) colocation* – DBNG-UPs may be collocated with edge compute infrastructure for market segments where a high percentage of the traffic is sent to public cloud edge services. The DBNG-UP and MEC collocation let subscriber traffic reach MEC applications at the edge without traversing Internet to regional data centers or public cloud region, thus reducing traffic latency for time-sensitive applications.

This paper describes a cloud-native DBNG solution to support the BNG public cloud use cases. It employs a cloud continuum model to establish a reference architecture with DBNG-CP and DBNG-UP instances distributed across operator's data center, cloud edge, and cloud region. The paper provides design options and tradeoff analysis on several key architecture components – the DBNG connectivity and user plane traffic transport between layer-2 oriented broadband access network on premises and layer-3 oriented network in the cloud, cloud DBNG-UP activation, and cloud DBNG high availability (HA). The paper also shares first-hand experience in the solution prototyping with focus on the implementation methods. The paper provides a direction and solution schemes for using public cloud for BNG disaster recovery and capacity on demand.

In the following, we review today's DBNG technology and typical on-premise deployment scenarios in Section 2. Section 3 details the DBNG cloud solution design. In Section 4, we share our initial solution prototyping experience and lessons learned. We conclude the paper in Section 5.

2. Today's DBNG – An Overview

This section provides an overview of the Disaggregated BNG (DBNG) technology which is a foundation for our DBNG cloud architecture design. The section also describes typical DBNG deployment scenarios today and characterizes DBNG workloads expected for support in the cloud.

2.1. Multi-Service Disaggregated BNG

The concept of Control and User Plane Separation (CUPS) in the telecom industry started to gain attention and prominence over the past several years. It emerged as a result of the increasing demands for flexibility, scalability, and efficiency in network architectures. Broadband Forum released its first Technical Report TR-459 "Multi-Service Disaggregated BNG with CUPS" in 2020[1].

A high-level DBNG architecture with CUPS is shown in Figure 2, where the control plane and the user plane are separated with management, control packet redirect, and state control interfaces in-between. The State Control Interface, SCi, is based on the Packet Forwarding Control Protocol (PFCP). The Control Packet Redirection interface, CPRi, is based on GPRS Tunnelling Protocol User Plane (GTP-U) which encapsulates control plane packets. The Management Interface, Mi, uses Network Configuration Protocol (NETCONF).

The relationship between DBNG Control Plane (DBNG-CP) and DBNG User Plane (DBNG-UP) is 1-to-many with a single DBNG-CP instance controlling multiple DBNG-UP instances. Multiple DBNG-UP instances may be deployed to achieve the desired network scaling, to allow load balancing of sessions, to support different service levels, to provide Multi-Access Edge Compute (MEC) services to a subset of subscribers, to provide DBNG-UP functions within different network slices, or to provide data path redundancy for BNG resiliency. The Subscriber Session Steering (SSS) is a new network function being

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

specified by Broadband Forum[2][3][4]. It provides the capability to direct subscriber sessions to the most appropriate DBNG-UP per the multi-DBNG-UP use cases listed above.

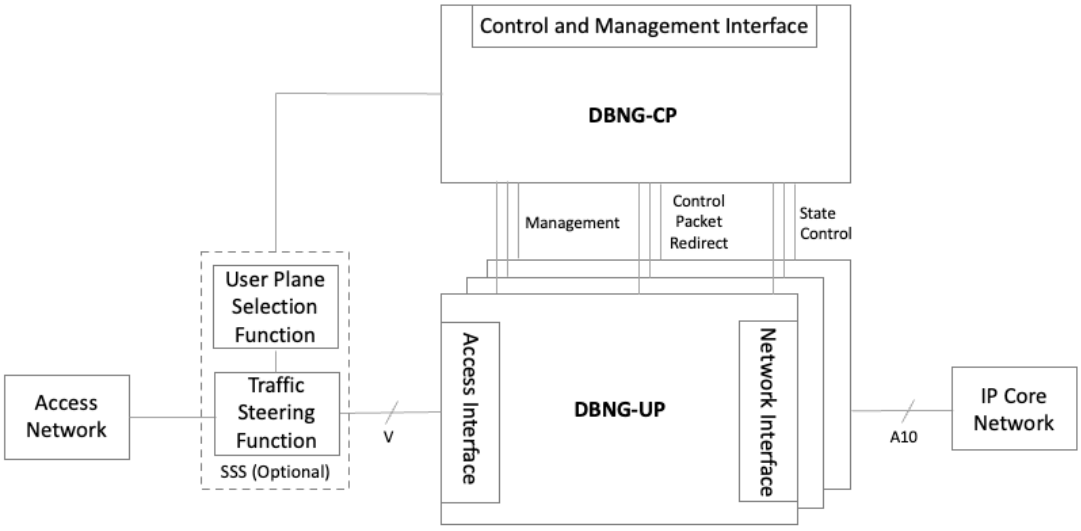

Figure 2 - Disaggregated BNG with Control and User Plane Separation (CUPS)

2.2. DBNG On-premises Deployment and Workload Characteristics

The DBNG technology has been productized by several vendors and deployed in broadband operators' data centers over the last few years. Figure 3 shows a typical DBNG deployment model with geo-redundancy where control planes are deployed as redundant instances managing multiple user planes per cluster which are connected to southbound optical line termination (OLTs). DBNG-UPs also include provider edge routing capabilities connected to the core network such as VRF support with BGP routing and different transport options such as VLANs, MPLS or SR-MPLS. In a typical broadband access network, OLTs communicate with BNG-UPs using L2 (802.1Q or QinQ) with different topology option, directly connected through local switching infrastructure or utilizing an L3 overlay (MPLS Pseudowires, EVPN over SR-MPLS or SRv6) to carry L2 frames. Backup paths are provisioned on the access network to protect against network connection or BNG-UP failures as well as between BNG-UPs and BNG-UPs with redundant CPs.

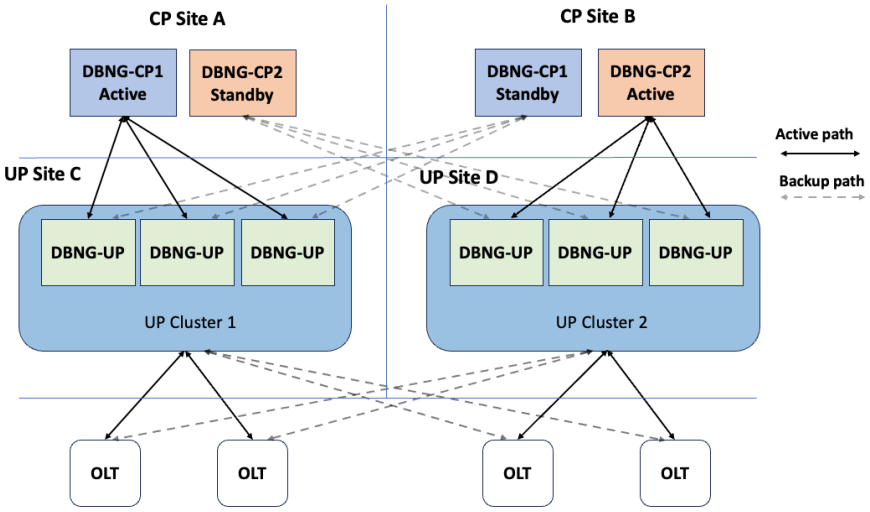

Figure 3 - DBNG Deployment Model with Geo-redundancy

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

In most cases with distributed architectures, operators design geo-redundant sites for User Plane and Control Plane instances to protect against site failure. DBNG-UPs are deployed with N+1 redundancy within the same site or across sites. Figure 3 depicts the case with resilient User Plane clusters per site with Active and Standby Control Planes on different sites. Other architectures are possible.

The DBNG-CP redundancy is based on utilizing the same CP IP address by Active and Standby CPs as described further in this paper.

DBNG-CPs and DBNG-UPs are deployed on x86 servers as bare metal or virtual machines (VMs) or Kubernetes Pods as vBNG, scaling throughput from 40 Gbps to 400 Gbps per server or virtual instance.

Table 1 lists typical broadband access requirements for DBNG.

Table 1 - BNG Workload Characteristics

BNG Performance Metrics	Typical Requirement
BNG-UP throughput per instance	40 Gbps to 400 Gbps depending on network design involving user plane distribution level and subscriber capacity per user plane
BNG-CP throughput between BNG-UPs and BNG-CP	1 Gbps to 5 Gbps per DBNG-CP
OLT to BNG-UP latency	25 ms for standard services, < 10 ms for low latency services
BNG-CP to BNG-UP latency	Up to 100 ms
BNG-UP interfaces	QinQ encapsulation from OLTs IPoE (DHCP) and PPPoE

3. DBNG Cloud Solution Design

3.1. Design Approach and Goals

Our approach to DBNG cloud transformation is built on two technology pillars – the DBNG CUPS architecture described above and Cloud Continuum as illustrated in Figure 4. The cloud continuum extends Cloud infrastructure and services from Region to Edge and Far Edge to support applications with different timing constraints or local residency requirements. The region, edge, and far edge are inter-connected and have the same "look and feel" for cloud infrastructure services such as networking, compute, storage, command line interface (CLI) commands, and deployment automation.

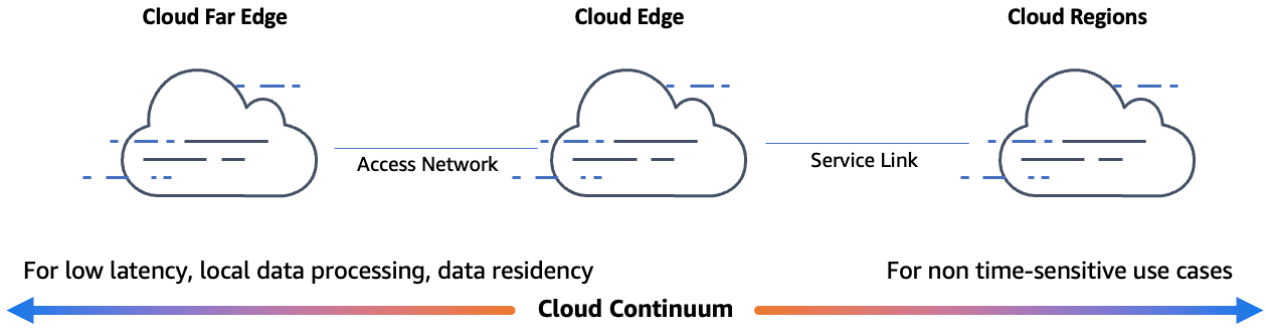

Figure 4 - Cloud Continuum: Region, Edge, and Far Edge

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

With the Cloud Continuum infrastructure and the CUPS architecture, DBNG control plane and user plane functions can be distributed across cloud region, edge, and far edge to achieve several goals:

- Higher DBNG resiliency – Cloud Edge or Region is used for Disaster Recovery (DR) of BNG storage or UP functions. In case of BNG equipment failure or natural disasters in operator's central office on premises, the BNG instances in the cloud are activated to process subscriber traffic.
- Dynamic capacity scaling – When subscriber traffic load surges due to, for example, a major sport event, Cloud Edge or Region provides extra DBNG-UP capacity to handle the bursty traffic on demand.
- Multi-access edge computing (MEC) support – Cloud Edge is used to host a variety of time-sensitive MEC applications, such as AR/VR 3D rendering, gaming, and video streaming. The BNG-UP co-located at the edge directs the subscriber traffic to/from the MEC applications locally without going out to Internet.
- Cost saving – It is the cloud elasticity that provides DBNG-UP capacity on demand. Operator doesn't need to invest in building central office facilities for DR or handling traffic load bursting, thus saving capital expenses.

Allocation of BNG-CP and BNG-UP functions along the cloud continuum can be of different schemes. Figure 5 illustrates a few options to achieve the goals listed above. Let us walk through the diagram by the order of label numbers:

1) It is a baseline to deploy the DBNG-CP in cloud, in a region or edge, whichever is geographically closer to the access network aggregation point(s).
2) One or more DBNG-UP instances may reside in the network operator's central office on premises as it is today. The DBNG-CP functions communicate with the DBNG-UP functions through IP network connections.
3) One or more DBNG-UP instances are deployed in the cloud region (or the cloud edge) to serve the purpose of DR or capacity on demand.
4) One or more DBNG-UP instances are deployed in the cloud edge to serve latency-sensitive MEC applications hosted at the cloud edge.

5-6) To support Subscriber Session Steering (SSS), the SSS Control Plane (SSS-CP) is co-located with the DBNG-CP in the cloud and the SSS User Plane (SSS-UP) performs the traffic steering function at the access network aggregation point on premises.

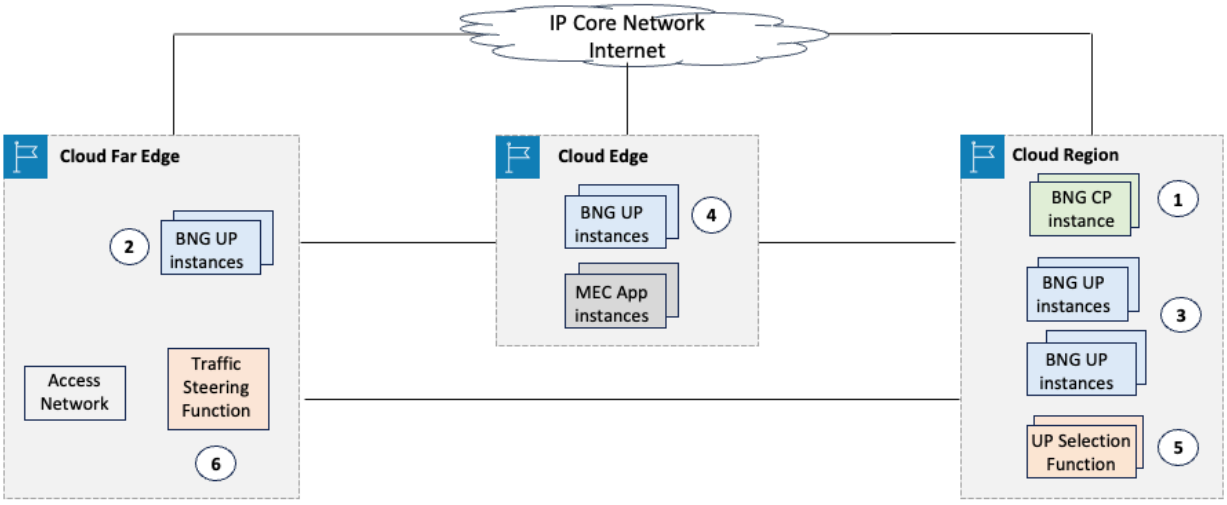

Figure 5 - DBNG Functions Allocation along Cloud Continuum

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

Hybrid-cloud solution design is required to realize the DBNG cloud allocation scheme. The solution design must address a number of technical requirements, including (1) connectivity between the access network on premises which is layer-2 and the cloud infrastructure which is typically layer-3, (2) high bandwidth (400Gbps) of the network connections and cloud compute resources, (3) low-latency access from the Optical Line Termination (OLT) in the access network to the cloud DBNG-UP, (4) fast cloud DBNG-UP activation time (30sec.), and (5) high availability of DBNG-UP and DBNG-CP functions in the cloud.

3.2. Reference Architecture and Solution Options

This section provides a hybrid-cloud solution design for the DBNG cloudification. It shows how the solution meets the technical requirements identified above. To describe the solution design in sufficient details, this section uses certain public cloud constructs as examples of a cloud infrastructure.

Figure 6 depicts the solution architecture that reflects the DBNG functions allocation shown in Figure 5. The public cloud is comprised of a region and one or more cloud edge zones which are linked together through dedicated direct connection or site-to-site VPN connection. The region consists of two or more availability zones for cloud facility redundancy and application redundancy. Equivalent to traditional data center is Virtual Provide Cloud (VPC) which hosts virtual machines, storage, networking constructs such as subnets, routers, network address translation (NAT) gateway, virtual private gateway (VGW) for the dedicated connection with operator data center, and Internet gateway (IGW) for traffic to/from Internet. Note that with the cloud continuum model, the VPC is extended from its availability zones in the region to the edge zones as if it was one data center across geographically-distributed region and edge locations, as illustrated by the purple-color box in Figure 6.

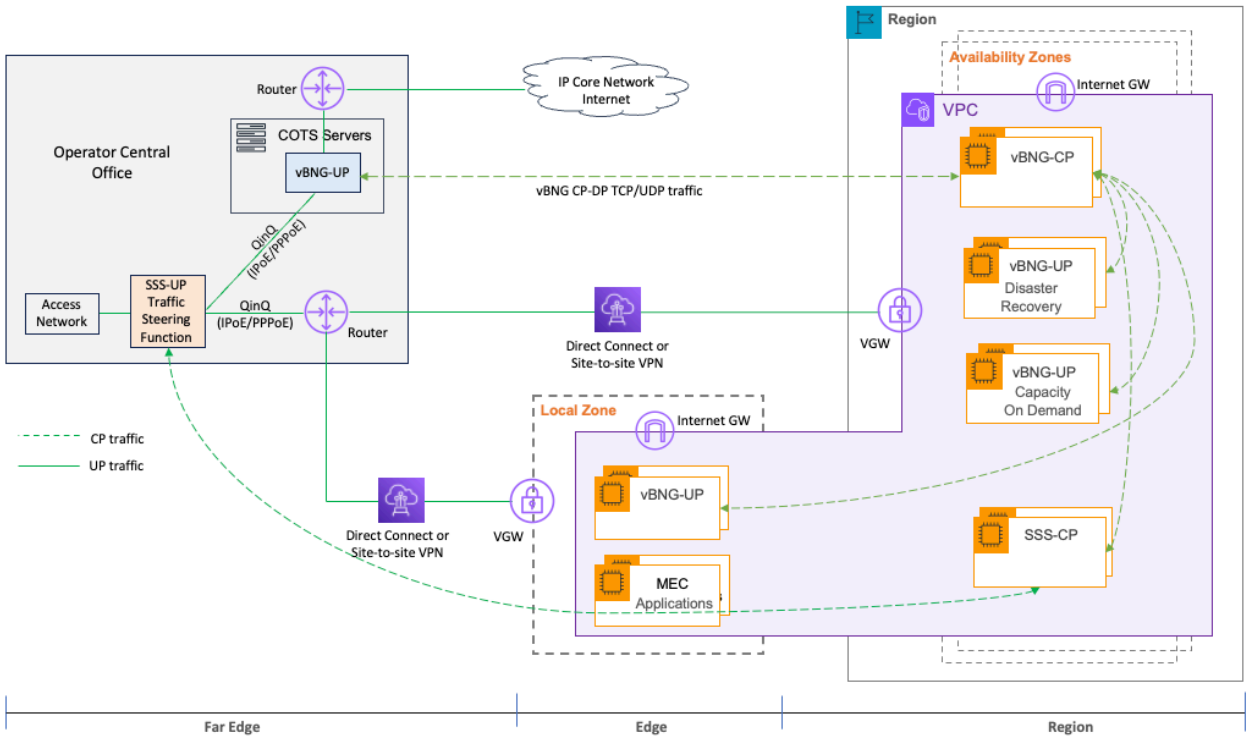

Figure 6 - Cloud DBNG Solution Architecture

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

In the cloud, the DBNG-CP and DBNG-UP instances are deployed in the region's availability zones for the DR and capacity-on-demand purposes. The SSS-CP function resides in the region as well. DBNG-UP instances are also deployed at the edge Local Zones for MEC applications. The DBNG-CP, DBNG-UP, and SSS-UP instances in the cloud can be in the form of virtual machines or containers. The virtual machines and containers are in auto-scaling groups for cloud elasticity and reliability.

On the operator premises, DBNG-UP instances run on commercial-off-the-shelf (COTS) servers with standard Kubernetes or any commercial version Kubernetes. The egress subscriber traffic goes through the access network and the session steering function to (1) the local DBNG-UP at normal operation conditions, (2) the DBNG-UP instances in the region for DR, (3) the DBNG-UP instances for additional capacity in the region when the traffic load surges, or (4) the DBNG-UP at the edge for MEC services.

Several architecture components and their design options are described in detail below.

Connectivity between the access network and the cloud – In general, there are two types of connectivity between the access network on premises and the cloud VPC. One is a dedicated connection through a local carrier network and the cloud infrastructure backbone network and the other through Internet with VPN tunnels. Typically, the former can support 100 Gbps or more per connection link and the latter several Gbps per link. Assuming the DBNG-UP throughput is at 200 Gbps of data plane traffic or more from the Access network, multiple links of the direct connection can be used for DBNG-UP in the cloud. Network latency with the dedicated connection from the access network to a cloud edge zone in close proximity is in single digit of milliseconds and tens of milliseconds to a cloud region depending on the proximity among other factors.

DBNG-UP layer-2 support – Running DBNG-UP in the cloud requires support of layer-2 (L2) protocols such as QinQ between the access network and the DBNG-UP instances in the cloud. In general, there are two types of solutions for transporting L2 frames bidirectionally between the access network and cloud: L2 overlay network (e.g., VXLAN, L2TPv3, GRE) or layer-3 routing (e.g., SRv6). Each solution option has its pros and cons as highlighted in Table 2 - L2 Frame Transport Solution Options below. As an example, a VXLAN-based solution design and prototyping are described in Section 4.

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

Table 2 - L2 Frame Transport Solution Options

Transport	Feature	Pros	Cons
1. VXLAN	• Encapsulate L2 Ethernet frames within UDP packets	• Provide greater scalability (w.r.t. number of virtual networks) and multicast support compared to L2TPv3	• Has a larger header size compared to L2TPv3, resulting in less efficient use of network resources. • Management overhead of VTEPs
2. L2TPv3	• Encapsulate L2 frames for point-to-point links • Allow multiplexing of multiple PPP session between two IP-connected endpoints	• Relatively simple compared and SRv6 • Has built-in security features that provide encryption and authentication • Built-in redundancy in control plane	• Higher encapsulation overhead compared to GRE • Does not provide any routing functionality
3. BGP EVPN	• Use BGP as a control plane to distribute MAC and IP reachability information for layer-2 and layer-3 connectivity	• Efficient MAC address distribution • Scalability for large-scale deployments with many MAC addresses and endpoints.	• Configurations can be complex, especially for large deployments or using advanced features.
4. GRE	• Encapsulate L2 or L3 packet • Encapsulate a variety of network protocols	• Simpler point-to-point mechanism compared to L2TPv3 • More flexible, can transport a variety of protocols	• can be complex; service chains within a VPC (e.g., virtual router – virtual firewall – intrusion detection) will also need GRE tunneling between them. • Doesn't encrypt traffic. An additional layer, usually IPsec, to protect the traffic. • No redundancy in the protocol
5. SRv6	• Use the IPv6 routing infrastructure • Provide source routing capabilities, allowing packets to take specific paths for traffic flows	• Does not require additional encapsulation beyond the IPv6 packet format, more efficient use of network resources due to smaller header size compared to L2TPv3 and GRE • Good migration strategy to EVPN-VPLS over SRv6 • Offer fine-grained control over the network traffic, enabling more flexible traffic engineering, service chaining, and network programmability	• Not all networks or vBNG ready for IPv6, less versatile compared to GRE

Cloud DBNG-CP High Availability (HA) – High Availability (HA) is an integral part of DBNG-CP deployment in the public cloud. The cloud DBNG-CP must continue to function when a DBNG-CP instance fails. The DBNG control plane maintains the states of subscriber sessions, operational states from user

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

planes, etc. Since DBNG-CP is stateful, it typically requires 1:1 active-standby redundancy for fast stateful switchover in case of DBNG-CP instance failure or planned maintenance. The cloud infrastructure can support DBNG-CP active-standby redundancy with solution options for virtual machine (VM) based software and container-based software, respectively.

1:1 active-standby DBNG-CP VM redundancy is shown in Figure 7. The cloud VM redundancy solution is comprised of several HA mechanisms:

- VM cluster contains two VM instances, one for active DBNG-CP and the other standby DBNG-CP, deployed in two cloud availability zones. The VM cluster provides redundancy at two levels: VM instance and cloud availability zone.
- An HA software component provides HA logic and mechanism (e.g., CP-to-CP keepalive) for detecting DBNG-CP failure, VM instance failure, or availability zone failure and conducting DBNG-CP switchover.
- Session State consistency between the active and standby DBNG-CP VM instances either can be maintained through state replication from the active VM instance to the standby or sharing a persistent storage such as cache or database service in the cloud.
- Overlay IP serves the DBNG-CP interface as a virtual IP address over the IP addresses of the active DBNG-CP VM instance and the standby instance. The overlay IP stays the same no matter which of the DBNG-CP VM instances is active. The overlay IP can be realized by using a route table in the cloud VPC, where the overlay IP is the traffic DBNG-CP destination address and the target IP address is the active DBNG-CP instance's interface address. At the time of switchover, the HA cluster software replaces the route table's the target IP address with the IP address of the newly active DBNG-CP instance.
- BGP Anycast may be utilized in order to support redundant DBNG-CPs with the same IP address in different locations in Active-Standby mode.

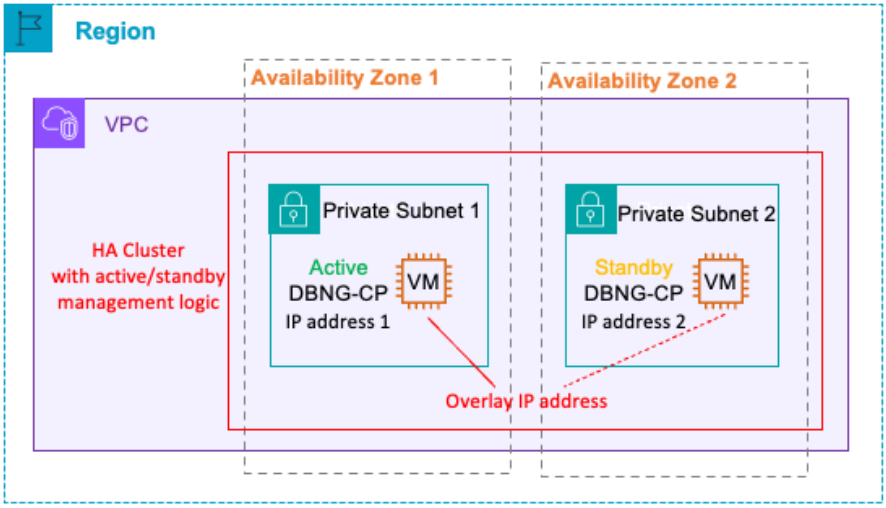

Figure 7 - 1:1 Active-Standby DBNG-CP VM Redundancy in Cloud

For examples of active-standby VM redundancy implementation in public cloud, please refer to references [7] and [8].

1:1 active-standby DBNG-CP Kubernetes Node & Pod redundancy is shown in Figure 8. The redundancy solution with Elastic Kubernetes Service (EKS) is comprised of several HA mechanisms:

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

- EKS cluster contains two EKS nodes (EC2 instances) deployed in two cloud availability zones. An EKS pod in each of the nodes runs DBNG-CP with one in active mode and the other standby mode. The EKS cluster provides redundancy at three levels: EKS pod, EKS node, and cloud availability zone.
- HA software provides HA logic and mechanism for detecting DBNG-CP failure, pod failure, node failure, or availability zone failure. For example, control plane instances implement a keepalive mechanism to detect failures and trigger switchover.
- Session State consistency between the active and standby DBNG-CP pods can be maintained either through state replication from the active pod to the standby or sharing a persistent storage such as cache or database service in the cloud.
- Single IP address serves the DBNG-CP interface. There are different methods to implement the single IP. For example, the Kubernetes Service IP can be used over different DBNG-CP instances to guarantee that the active DBNG-CP utilizes the Service IP to communicate with DBNG-UPs. Alternatively, redundant DBNG-CP instances employs BGP Anycast configured with the same IP address and different preference for the active and standby nodes.

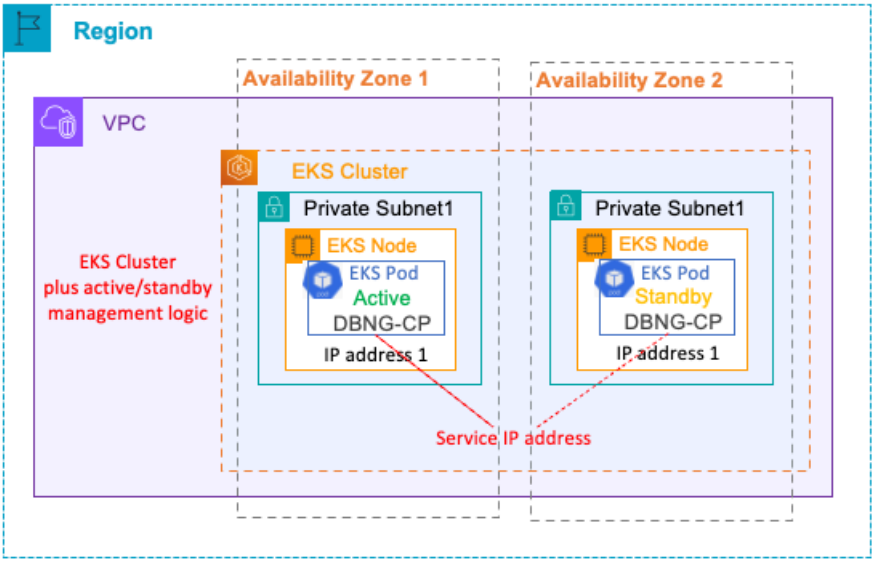

Figure 8 - 1:1 Active-Standby DBNG-CP Pod and Node Redundancy in Cloud

DBNG-UP Redundancy – To support the on-premise DBNG-UP disaster recovery, backup DBNG-UP instances are deployed in the public cloud as illustrated in Figure 9. This on-premise and in-cloud DBNG-UP deployment can be characterized by an active-standby redundancy model where the on-premise DBNG-UP is in active mode and the cloud DBNG-UP in standby in normal operation conditions; at the time of on-premise DBNG-UP failure, the cloud DBNG-UP becomes active.

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

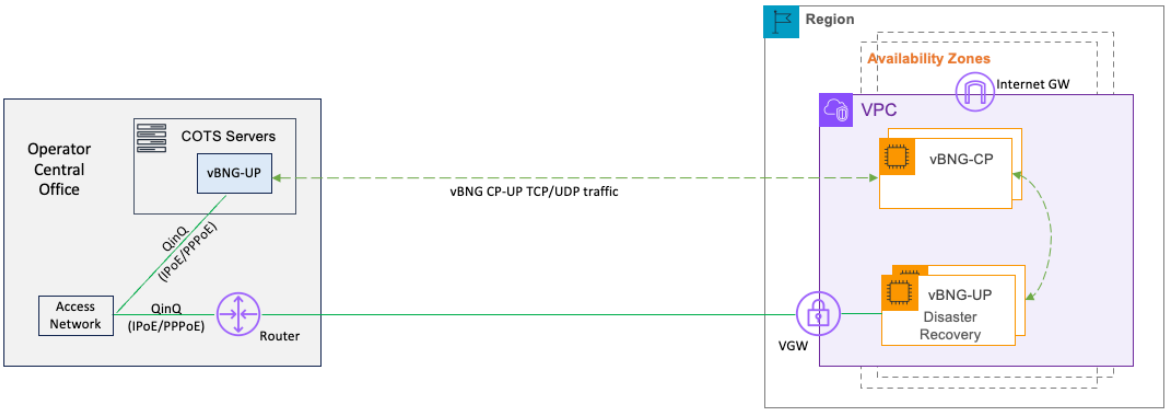

Figure 9 - DBNG-UP Redundancy across On-premise Datacenter and Public Cloud

The active-standby DBNG-UP redundancy can be achieved with Virtual Router Redundancy Protocol (VRRP). The VRRP software should be deployed in the DBNG-UP instances on premise and in the cloud. Note that the connection between Access Network OLT and DBNG-UP instances is layer-2. Hence, the VRRP configuration requires a unique MAC address on each of the DBNG-UP instances.

At the DBNG-UP level, resilience can be implemented with Subscriber Groups (SGRP) in different user planes effectively achieving N+1 redundancy. Figure 10 shows the DBNG-UP resiliency model. A subscriber group is a group of subscribers sharing the same policy and redundancy scheme. The subscriber groups are instantiated on Primary and Standby User Planes. When sessions are established on a primary subscriber group, the same session is programmed by the active control plane on the secondary user plane. For disaster recovery, Primary DBNG-UPs are on premises while Standby DBNG-UPs re on public cloud.

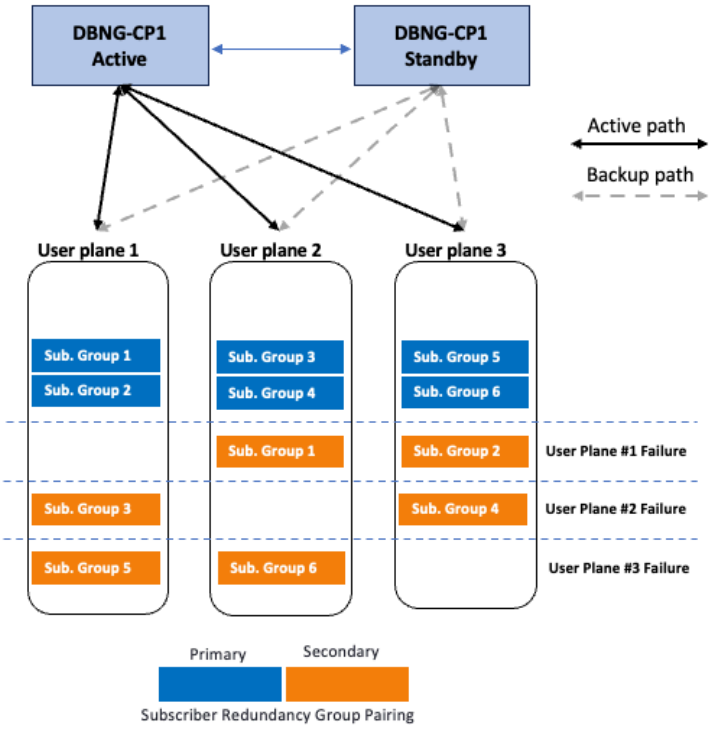

Figure 10 - Subscriber Groups across DBNG-UPs

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

Cloud DBNG-UP activation and DR life cycle

Sessions are established on the primary Subscriber Group (SGRP) on the on-premise DBNG-UP and are programmed on the backup cloud DBNG-UP. Upon primary UP failure, the backup DBNG-UP takes over all SGRPs and existing sessions are steered to the new UP based on the virtual MAC move.

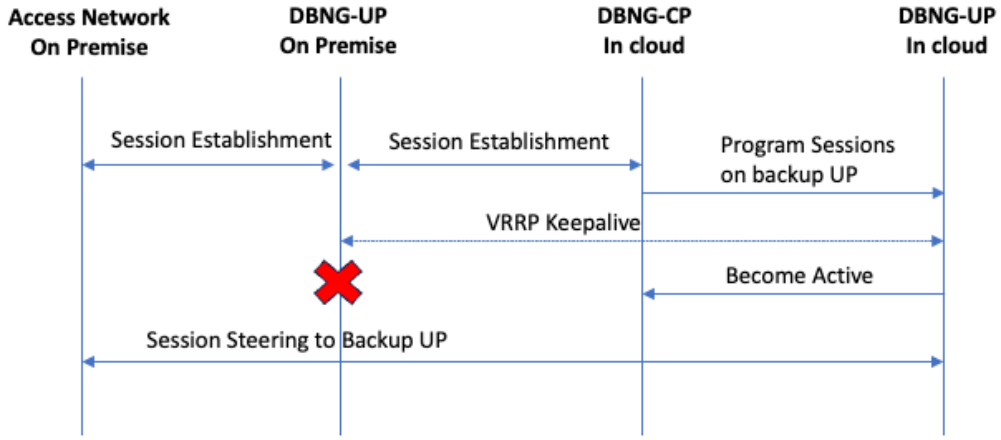

Figure 11 – DBNG-UP Disaster Recovery Life Cycle

Cloud DBNG cost analysis – The cloud DBNG consists of both DBNG-CP and DBNG-UP deployed in the public cloud with the former running in an active mode all the time and the latter in a standby mode. The standby DBNG-UP can be configured as Hot standby or Warm standby. In the hot standby mode, the DBNG-UP instances, i.e., VMs or Kubernetes Pods, are instantiated and in "running" mode. At the time of disaster recovery or traffic bursting from the on-premise data centers, the DBNG-UP instances start processing the subscriber sessions in near real time. In the warm standby mode, the DBNG-UP VMs or Kubernetes Pods are instantiated but in "stopped" mode. It may take a few minutes to get the stopped DBNG-UP instances up running and processing the subscriber traffic. With the cloud "pay-as-you-go" model, the DBNG-UP instances in the "running" mode incurs cloud service charges whereas the instances in the "stopped" mode does not.

Figure 12 shows a cost comparison between the DBNG-UP hot standby and warm standby. It is modeled with VM instances or EKS work nodes with aggregated 200 Gbps network bandwidth and 3 days processing time per month. The chart shows the cost of warm standby DBNG-UP is only 10% of the cost of hot standby DBNG-UP. Clearly, it is a tradeoff between the timeliness of cloud DBNG-UP activation and the cost of running DBNG-UP in the public cloud. It is up to broadband operators to make a tradeoff decision with cloud DBNG service cost and subscriber impact in mind.

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

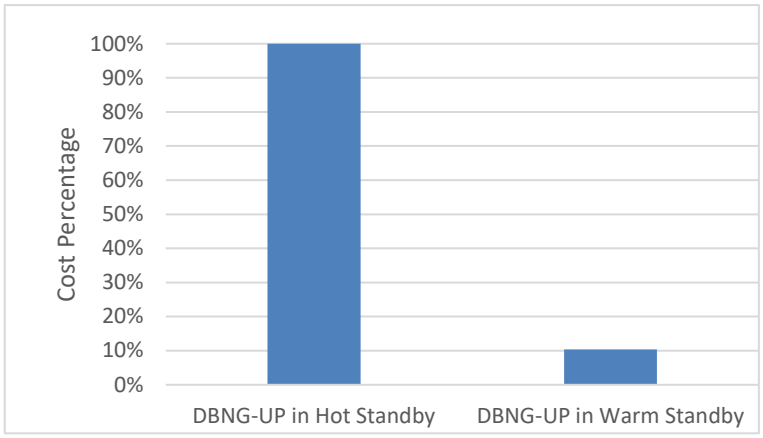

Figure 12 - Cost Comparison of DBNG-UP Standby Modes

4. Solution Prototyping and Implementation Methods

This section shares our first-hand experience in the solution prototyping using a public cloud and commercial disaggregated DBNG-CP and DBNG-UP.

The solution in prototyping is shown in Figure 13. It resembles the generic solution architecture pattern described in Section 3.2. In particular, a lab serves as an on-prem data center and a cloud region is used as the public cloud. A dedicated cloud network interconnects the lab on premises and the cloud region.

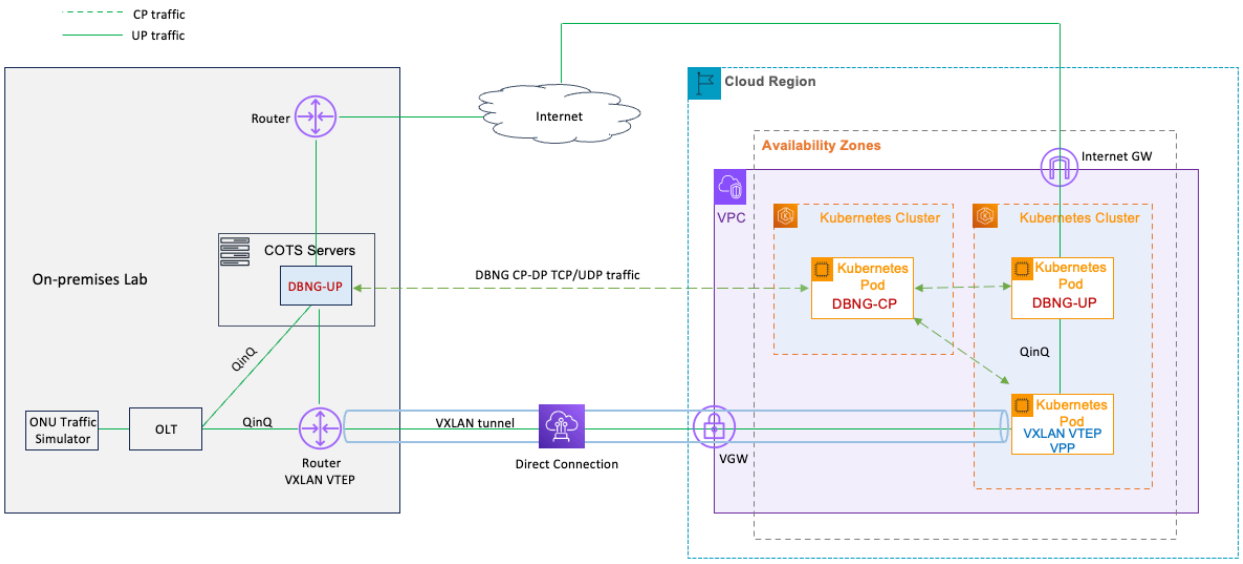

Figure 13 - Prototyping Environment

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

In the cloud region, a Virtual Private Cloud (VPC), an equivalent of a data center, is created with Internet Gateway (IGW) and subnets associated with security rules and IP traffic routes. The subnets host the DBNG-CP container and DBNG-UP container in Kubernetes clusters, respectively. A Virtual Private Gateway (VGW) interfaces the direct connection link with the Lab. The DBNG-UP traffic flows between the DBNG-UP and the internet through IGW.

Realizing the layer-2 traffic transport between the layer-2 access network on premises and the layer-3 IP network in the public cloud is a major challenge in cloudifying DBNG-UP. Out of the layer-2 tunneling solution options listed in Section 3.2 Table 2, the VXLAN tunneling method is chosen for the prototyping. The VXLAN method is largely driven by existing vBNG deployment cases in operator environments. The VXLAN implementations in the cloud and on premises are described respectively in the following.

(1) In the public cloud, the VXLAN tunnelling is implemented in the Elastic Kubernetes Service (EKS) cluster as illustrated in Figure 14. It utilizes four software components: Data Plane Development Kit (DPDK), Vector Packet Processor (VPP), Layer-2 Bridge, and Tap interface.

- *DPDK* – is a set of open-source libraries and drivers designed to accelerate packet processing tasks in networking applications. DPDK provides a collection of optimized functions that allow software applications to efficiently interact with network interface cards (NICs) and perform packet processing operations directly on data plane hardware. By bypassing the traditional kernel-based networking stack, DPDK enables applications to achieve high throughput and low latency, making it particularly suitable for building high-performance networking solutions, such as routers, switches, and virtualized network functions like DBNG-UP.
- *VPP* – is a high-performance software framework designed for network data packet processing. It's often used in networking applications and network function virtualization environments. VPP is engineered to efficiently handle large volumes of network packets by optimizing packet processing tasks using vectorization techniques. This allows for improved throughput and reduced latency in network communication.
- *Layer-2 Bridge* – is a networking device or software function that operates at the data link layer (Layer 2) of the OSI model. Its primary purpose is to connect and forward Ethernet frames between two or more network segments or interfaces, effectively extending the same local area network (LAN) across multiple physical or virtual network segments.
- *Tap interface* – is a virtual network interface connected to a bridge device. This configuration is often used in virtualization environments to enable communication within VMs or between a container and the VM that is hosting it. In some instances, it can also facilitate communicating between VM internal and external networks.

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

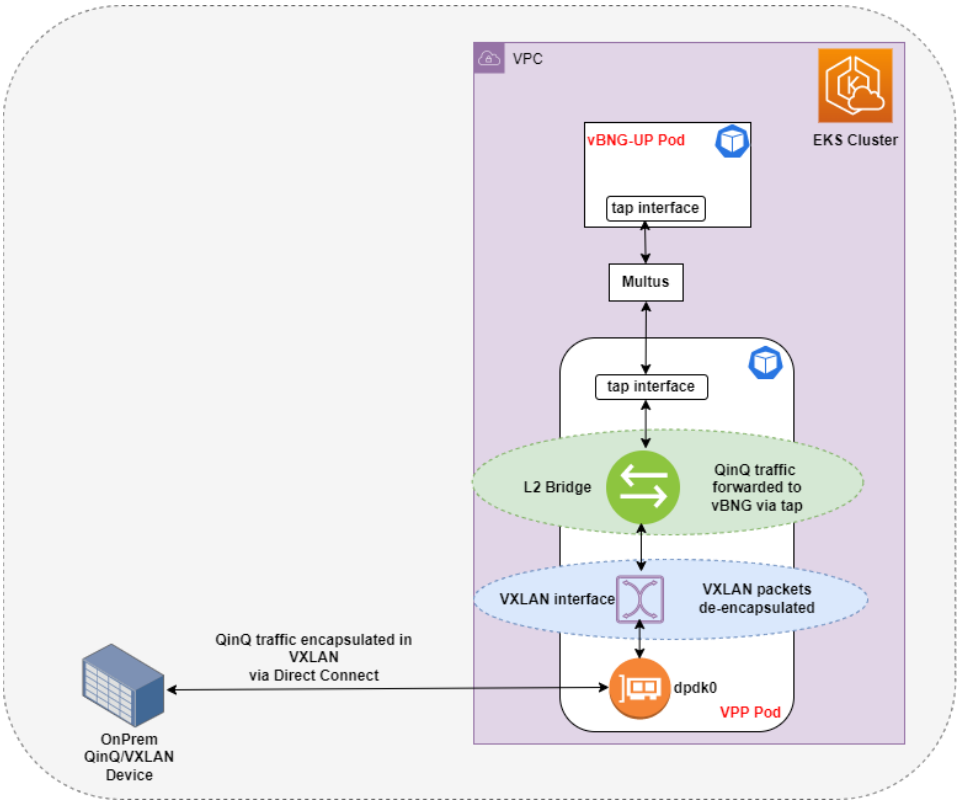

Figure 14 - VXLAN Implementation Method Using VPP

In the prototype, VPP is deployed in an EKS Pod. It provides the following functions:

- VXLAN Tunnel End Point (VTEP) for encapsulating or decapsulating layer-2 frames transported between the cloud DBNG-UP and the OLT on premises.
- A DPDK interface for the VXLAN traffic.
- A tap interface for transporting the QinQ frames between the VPP Pod and the DBNG-UP Pod in the EKS cluster.
- A bridge to connect the VXLAN VTEP and the tap interface

The VPP configuration is shown below with sample data:

```
set interface ip address dpdk0 10.0.0.115/24
set interface state dpdk0 up
ip route add 0.0.0.0/0 via 10.0.0.1 dpdk0
create vxlan tunnel src 10.0.0.115 dst 192.168.0.100 vni 1000 decap-next l2
set interface state vxlan_tunnel0 up
create bridge-domain 10
create tap id 1 host-if-name tap1
set interface state tap1 up
create sub-interface tap1 100 dot1ad 100
set interface state tap1.100 up
set interface l2 bridge vxlan_tunnel0 10
set interface l2 bridge tap1 10
set interface l2 bridge tap1.100 10
```

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

The VPP pod configuration is defined by a Network Attachment Definition (NAD) file shown below. The configuration uses Kubernetes' host-device plugin. It specifies the device parameter with the tap interface created by the VPP.

```
apiVersion: "k8s.cni.cncf.io/v1"
kind: NetworkAttachmentDefinition
metadata:
  name: vpp-tap-net
  namespace: encaps
spec:
  config: '{
      "cniVersion": "0.3.1",
      "type": "host-device",
      "capabilities": { "ips": true },
      "device": "tap1"
    }'
```

The layer-2 traffic flow from the on-prem OLT to the DBNG-UP in the cloud is as follows. The on-prem router with VXLAN VTEP sends QinQ frames encapsulated in UDP packets to the VPP pod where the VPP VXLAN VTEP decapsulate the QinQ frames from the UDP packets. Through the internal bridge, the VPP sends the QinQ frames to the vBNG-UP pod via the tap interface. The tap interface is also created in the VPP pod, one leg of the tap interface is connected to the VPP internal bridge while the other leg is injected into the EKS worker node's kernel space, from which Multus is then used to pass this tap interface into the vBNG-UP pod.

(2) On premises, VXLAN VTEP is available in commercial products or open source. As an example, VyOS is an open-source network operating system that can be deployed in servers on premises to provide VXLAN VTEP as well as router, firewall, VPN, and other network functions. To test the VXLAN tunneling first in a cloud environment, the prototype uses VyOS Amazon Machine Image and creates a VyOS VM instance to mimic an on-prem VXLAN VTEP. The VXLAN tunneling was successfully tested between the VyOS VM and the EKS Pod with VPP end to end. For the actual VXLAN VTEP on premises, the prototype uses a commercial product.

5. Conclusions

Public cloud possesses the characteristics of resource scalability, on-demand provisioning, pay-as-you-go pricing, and global accessibility. Because of these capabilities, wireline broadband operators may leverage the public cloud for DBNG use cases such as disaster recovery and capacity on demand. Using the cloud instead of building additional data centers will reduce capital expense and increase operation flexibility. In this paper, we have presented a reference architecture for deploying DBNG control plane and user plane in public cloud region or at cloud edge. Our prototyping experience as well as design description indicate that it is potentially feasible to run DBNG-CP and DBNG-UP in the public cloud for DBNG disaster recovery and capacity on demand use cases. Our next steps are to (1) conduct performance scaling test under failover and workload bursting conditions and (2) implement SRv6-based layer-2 frame tunneling to further demonstrate the public cloud feasibility for operator's DBNG production deployment.

Acknowledgements

The authors would like to thank Dr. Jennifer Andreoli-Fang for her valuable comments and edits in the paper content. The authors also knowledge the Ciena engineering team for their prototyping work to gain the first-hand experience in the vBNG solution deployment on premises and in the cloud end to end.

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.

Abbreviations

BNG	broadband network gateway
DBNG	disaggregated broadband network gateway
DBNG-CP	DBNG control plane
DBNG-UP	DBNG user plane
DPDK	data plane development kit
EKS	elastic Kubernetes service
SGRP	subscriber group
vBNG	virtual broadband network gateway
VM	virtual machine
VPC	virtual private cloud
VPP	vector packet processor
VRRP	virtual router redundancy protocol

Bibliography & References

[1] Broadband Forum TR-459 Issue 2, "Multi-Service Disaggregated BNG with CUPS: Reference Architecture, Deployment Models, Interface, and Protocol Specifications," April 2023, https://www.broadband-forum.org/pdfs/tr-459-2-0-0.pdf

[2] Broadband Forum WT-474, "Introducing Subscriber Session Steering and dynamic subscriber placement," https://www.broadband-forum.org/wt-474-introducing-subscriber-session-steering-and-dynamic-subscriber-placement

[3] Vodafone and Intel White Paper, "Transforming Fixed Access Using Traffic Steering on a Cloud Native Architecture," Nov., 2021, https://networkbuilders.intel.com/solutionslibrary/transforming-fixed-access-using-traffic-steering-on-a-cloud-native-architecture

[4] Broadband Forum WT-474, "Introducing Subscriber Session Steering and Dynamic Subscriber Placement," 2022, https://www.broadband-forum.org/wt-474-introducing-subscriber-session-steering-and-dynamic-subscriber-placement

[5] Networking Blog, "Connect a VXLAN-EVPN DC to the Public Cloud the right way," June 2020, https://ccie46985blog.wordpress.com/2020/06/21/connect-a-vxlan-evpn-dc-to-the-public-cloud-the-right-way/

[6] Networking Tutorial, "An introduction to the SRv6 (Segment Routing over IPv6 dataplane) technology," December 2017, https://www.segment-routing.net/tutorials/2017-12-05-srv6-introduction/

[7] Blog, "AWS Overlay IP In SAP Landscapes," last updated February, 2023, https://dzone.com/articles/aws-overlay-ip-in-sap-landscapes

[8] AWS Web Site, "Red Hat Enterprise Linux with High availability on Amazon EC2 ," last updated 2023, https://aws.amazon.com/partners/redhat/faqs/high-availability/

© 2023, SCTE® CableLabs® and NCTA. All rights reserved.