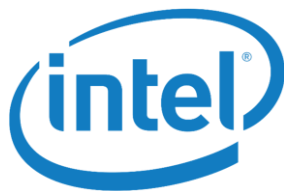**White Paper**

Layer 7 Visibility for Virtual CPE

Prepared by

Gabriel Brown
Senior Analyst, Heavy Reading
www.heavyreading.com

on behalf of

www.intel.com   www.qosmos.com

February 2016

vCPE & Embedded Network Intelligence

Virtual customer premises equipment (vCPE) is a way for network operators to transition enterprise access and virtual private network (VPN) customers to next-generation cloud networking platforms. This can substantially reduce costs associated with specialized hardware deployed on-premises and, with the right tools, enables operators to inject value into wide-area network (WAN) services using virtual network functions (VNFs).

This white paper reviews state-of-the-art vCPE deployments in the enterprise market and discusses how operators can use embedded traffic analysis software to design application-aware, customer-specific network services. In particular, it addresses Layer 4-7 analysis engines and their role in three key use cases: monitoring/reporting, Layer 7 firewall capability and VNF service chaining.

Why vCPE? What Are the Benefits?

The business case for vCPE has been extensively studied. A major technology transition is never simple, but there appears to be a good rationale for deployment based on the following factors:

- **Reduced deployment costs.** Depending on the architecture, in the vCPE model a single, lower-cost device can be installed on-premises to replace several specialized devices. Reducing "truck roll" is especially useful for remote offices and international offices where it is costly to send technicians.

- **An app store model for VNFs.** By virtualizing functions, such as firewalls, intrusion protection system (IPS) and session border controllers (SBCs), that previously ran on dedicated on-premises equipment, operators can create a catalog of software-based services that can be deployed on demand by enterprise VPN customers using self-service portals.

- **Low-cost white-box equipment.** Classic CPE devices are already cost-optimized; however, there is an opportunity to benefit from the volume economics associated with standard off-the-shelf servers if they can be optimized to run virtual network software.

How these benefits accrue depends on the service architecture. vCPE is a component of a WAN service and the benefits, therefore, are realized within the context of software-defined VPNs (a.k.a., cloud VPNs or SD-WAN).

Market Activity for vCPE & Software-Defined VPNs

vCPE is a lead application for network functions virtualization (NFV). It was one of the major NFV use cases identified by operators in the first ETSI white paper and has consistently ranked highly in terms of "likely to be deployed commercially" in Heavy Reading operator surveys.

Interest in vCPE has generated extensive investment in R&D by operators, vendors and cloud providers. This work initially focused on proofs of concept, but is now transferring to trials in live networks and, in some cases, to live deployments (further technical development work continues in parallel).

Figure 1 summarizes vCPE and cloud VPN services that are available from brand-name enterprise service providers. We expect more of these types of services to launch in 2016.

Figure 1: vCPE & Cloud VPN Service Offers

Operator	Service Description
Deutsche Telekom	• Cloud VPN services launched 1Q15 on pan-European IP network • Automatic provisioning of service via customer console • Virtualized networking functions run in OpenStack environment
Colt	• Progressive commercial service from an enterprise-focused provider • Low-cost Ethernet access device at customer premises • vCPE functions running on/behind IP edge router
Orange Business Services	• Pre-commercial trial of vCPE underway • Automatic provisioning of VPN services via customer console • Multivendor deployment using a cloud orchestration tool
AT&T	• Based on AT&T's "Network on Demand" platform • White-box devices on-premises; leverages SDN for service config. • Multivendor deployment; currently in advanced trial phase
Virtela (NTT Communications)	• True cloud-based SSL VPN targeted at smaller office locations • Wide range of cloud-based virtual network services • Reports significantly reduced service provisioning times

Source: Heavy Reading

Layer 7 Visibility & Network Services Dashboards

From the perspective of the "Digital Operator," the focus in vCPE and cloud VPN services should be on the customer-facing services that generate sales and ongoing business. A customer portal with Layer 7 visibility offers operators an opportunity to go beyond connectivity and climb the value chain by offering value-added services made up of VNFs (using the app store model) that are cost-effective to buy and operate.

In the case of an enterprise customer that wants to purchase WAN services across multiple office locations, the network should implement service requests automatically, in a matter of minutes, without the need for manual intervention on the part of the operator. This self-service model contrasts with the classic model where a series of manual orders and provisioning tasks would achieve the goal in weeks or even months. Automation enables operators to "close the loop" between agile, dynamic, software-configurable infrastructure and the services that run over it.

Visibility into application layer traffic is also important to how dashboards monitor and represent traffic on enterprise networks – for example, to address issues like compliance and "shadow IT" – and can be used to steer traffic into the correct service chains composed of VNFs. This service chaining use case is discussed below.

Architecture & Deployment Options

There are two basic architectures for vCPE deployment: (1) a centralized cloud model where a low cost Ethernet access device is deployed at the customer premises and virtual network services in the cloud or behind the IP edge; and (2) a model where a more powerful device that is capable of running third-party VNFs is deployed at the customer premises. Each is an evolution on the classic CPE box-stacking model and are shown in **Figure 2**.
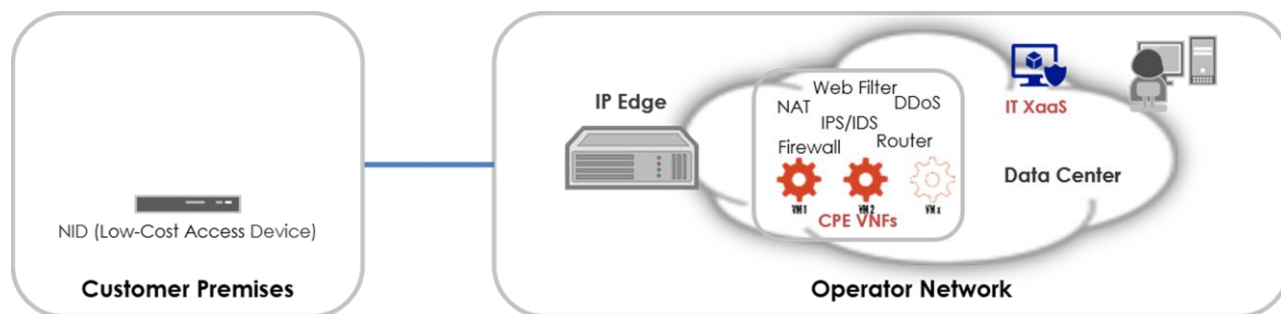
Figure 2: vCPE Deployment Models

Classic CPE Model With Multiple Appliances at Customer Site
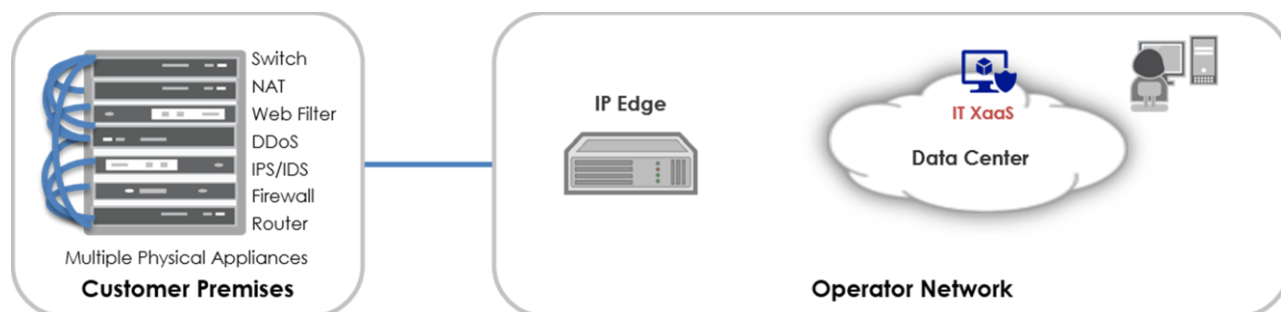
Smarter Access Devices Running Third-Party VNFs, Orchestrated From the Cloud
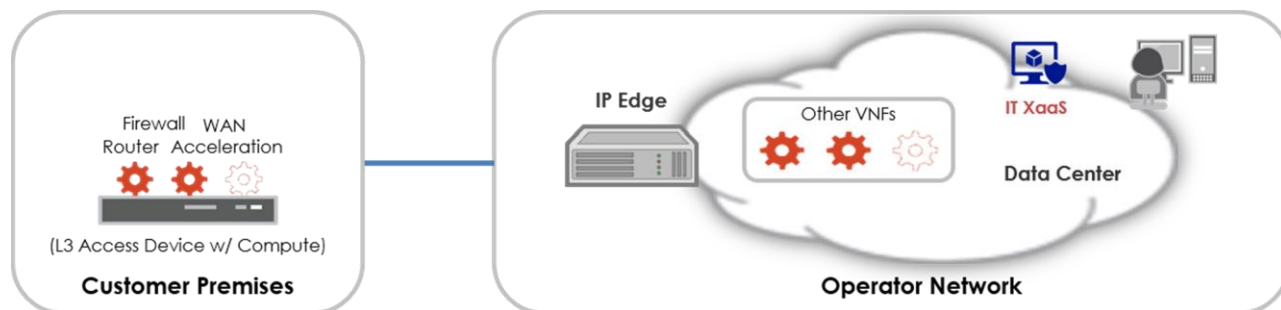
Cloud-Hosted vCPE Function Using Low-Cost NID
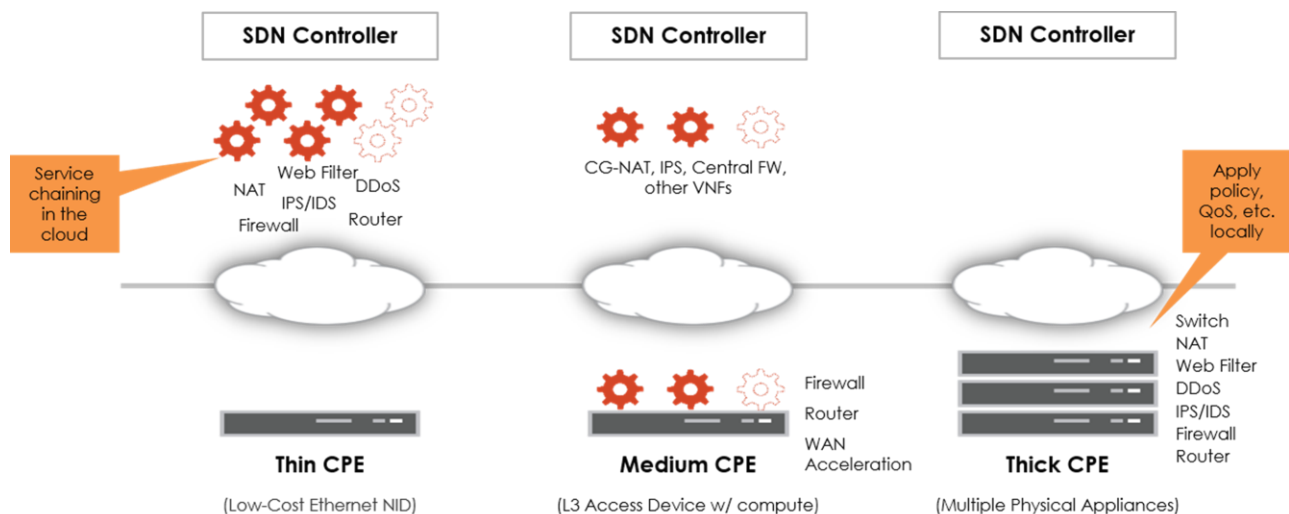
Source: Heavy Reading

These two options can be thought of as "bookends," with the specific configuration dependent on the type of enterprise customer, the nature of connectivity between

customer sites, and so on. Some factors favor a migration to a centralized cloud-hosted model overtime, while others suggest enduring value in localized compute, storage and routing. In practice, variations of the two models, with more or less capability deployment at the customer premises according to the operator or customer strategy, are likely to coexist even within the same operator network. In principle, service orchestration should be able to combine centrally-deployed and locally-deployed VNFs into an end-user service.

Deployment Model Impacts Service Delivery

vCPE deployment impacts where VNFs are deployed and, therefore, the provision of network services and the use of Layer 7 technologies. **Figure 3** shows the differences for a "thin" branch, a "medium" branch and a "thick" branch office.

Figure 3: Layer 7 Visibility & Deployment Models

Source: Heavy Reading, Qosmos

Of the three options shown, the thin branch model is most disruptive. VNFs are deployed in the cloud reducing truck roll and operational expense and making it (relatively) straightforward to add new VNFs using a cloud orchestrator and customer portal. In this case, Layer 7 visibility is most useful to support application-aware and subscriber-aware service chaining between cloud-hosted VNFs.

In some cases, however, there is a need for a more capable device at the customer premises. Services, such as WAN acceleration and local routing, necessarily need functionality deployed at the customer site and, in this case, a more capable CPE devices is needed. This CPE can replace many appliances with single, smart device that will run some applications "natively" and is capable of running third-party VNFs. Layer 7 visibility, in this case, would be useful to enforce quality of service (QoS) and prioritize application or user traffic across the WAN.

Note also that centralized VNFs could also be used in conjunction with VNFs running locally on this "smart" CPE device to create an end-to-end service using common orchestration tools.

Layer 7 Visibility for vCPE & Cloud VPN Services

Cloud VPN services comprise multiple physical and virtual components that combine to create a service. Operators wish to vary the composition of network functions to customize to end-to-end services – or, more accurately, want their customer to be able to create their own services via a portal – and need to identify and classify traffic to make routing decisions and to be able to insert, and remove, components from the processing path (a.k.a., service chain).

Typically, operators use IP and Ethernet headers to steer flows to the desired networking functions. This works well for the majority of services, but it is generally coarse-grained. In some cases there is a need for a more granular routing – for example, to optimize specific application performance, to undertake quality of experience (QoE) monitoring, or to look for anomalies that may permit security breaches. In these cases, an application-aware capability embedded in the infrastructure can be valuable.

There are several ways to gain visibility into flows and take action on them. This includes L7 proxies, hardware classifiers, packet sampling (DPI), and using header and handshake information. In the future network, it is envisaged that application-awareness will be an integral part of virtualized environments, such as via Open vSwitch (OVS), and that forwarding policy will be applied by SDN controllers, such as OpenDaylight.

Note that not all traffic needs to be analyzed via a DPI node or other classification engine for application-awareness to be effective. Typically, at any given time, about 90 percent of traffic in a network is already classified, leaving only 10 percent that needs to be analyzed further. This is because only the first few packets of each flow need to be analyzed in order to enforce policy. This helps to keep the overhead associated with packet inspection to manageable levels and minimizes the performance impact of DPI.

Traffic Monitoring & Reporting Dashboard

A lead use case for application-aware capability deployed as part of the vCPE service is monitoring and reporting. In both classic and virtualized environments, enterprise customers want (and often need) to understand what users are doing on the network and how the service is performing. Typical requirements are for usage monitoring and metering, service assurance, service-level agreement (SLA) monitoring, accounting and compliance.

Arguably, customers intuitively feel less in control of virtual services than physical appliances, and so monitoring and reporting become more important to vCPE deployments. In relation to specific cases, such as compliance and "shadow IT," there are now enough cases of runaway cloud service usage that improved monitoring and reporting capability is very much in demand.

This use case maps to the functions shown in **Figure 4**. It comprises in the following steps:

- Packets enter the vSwitch; copy of the packets is sent to a L7 IP Classification engine running in a virtual machine (VM)

- The L7 IP classification engine analyzes packets and interprets them in a rule engine

- An export module converts into IPFIX format and passes data up to analytics platform and dashboard

This use case can be thought of as a virtual probe. The attraction is that it is fully virtualized, it uses standard OVS, is very simple to implement, and it meets an important customer requirement.

Figure 4: Layer 7 IP Classification for Monitoring & Reporting
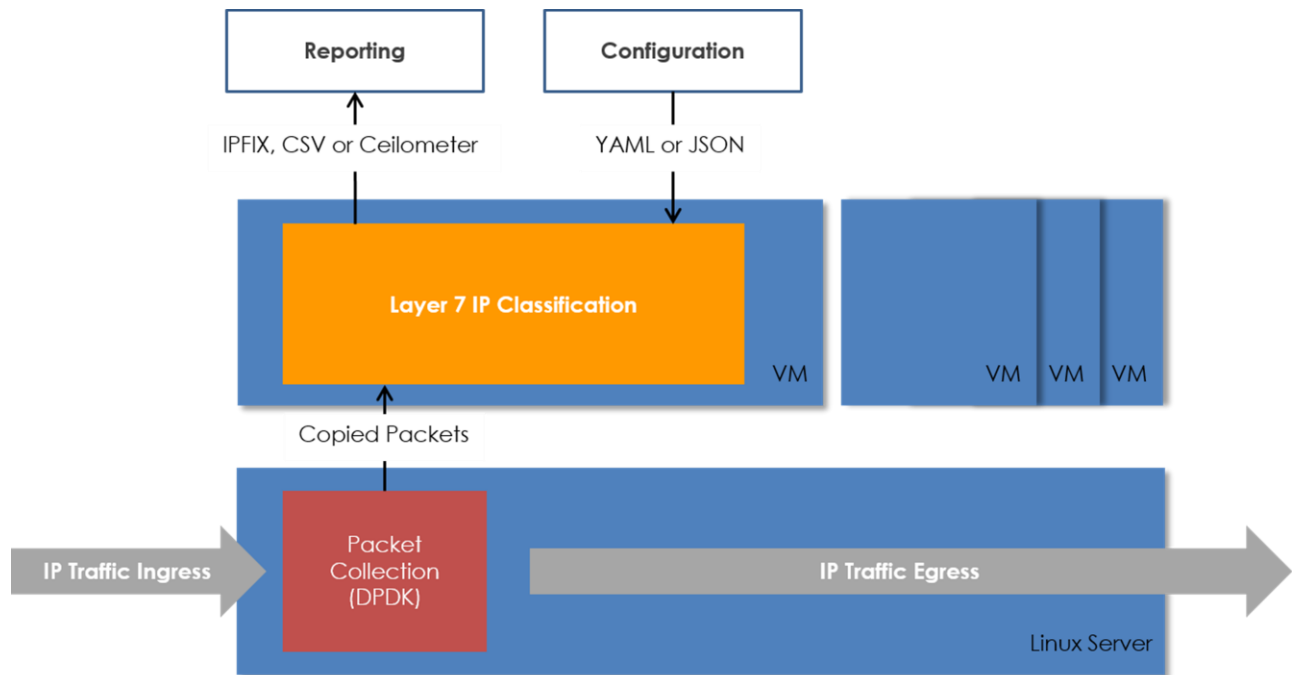
Source: Qosmos

Layer 7 Firewall & QoS

A Layer 7 firewall in Linux environment is another interesting use case because an enterprise may want to either block or rate-limit certain types of traffic (BitTorrent, for example) or manage the QoE of certain users or applications (Web video vs. line-of-business applications, for example).

The process is shown in **Figure 5**. The first stage of the process is similar to the monitoring example above, but in this case action needs to be taken once traffic is identified. Again, this is an entirely virtualized process. It proceeds as follows:

- Packets enter the vSwitch; a copy of the packet is sent using conntrack to a L7 IP classification engine running in user space

- Classification results are then interpreted in a rule engine, which has been configured with pre-determined rules

- The rules engine updates the OVS flow table, which applies policy – for example, to drop a flow, reduce the bit rate, or approve it as a default service

In essence, this use case enables operators to add Layer 7 application visibility to the Linux Layer 3-4 firewall. Because the initial packet in a flow is copied in the Kernel and then processed by the L7 IP classification engine, which in turn updates the OVS flow table, the process is effectively real-time.

Figure 5: Layer 7 IP Classification for Virtual Firewall
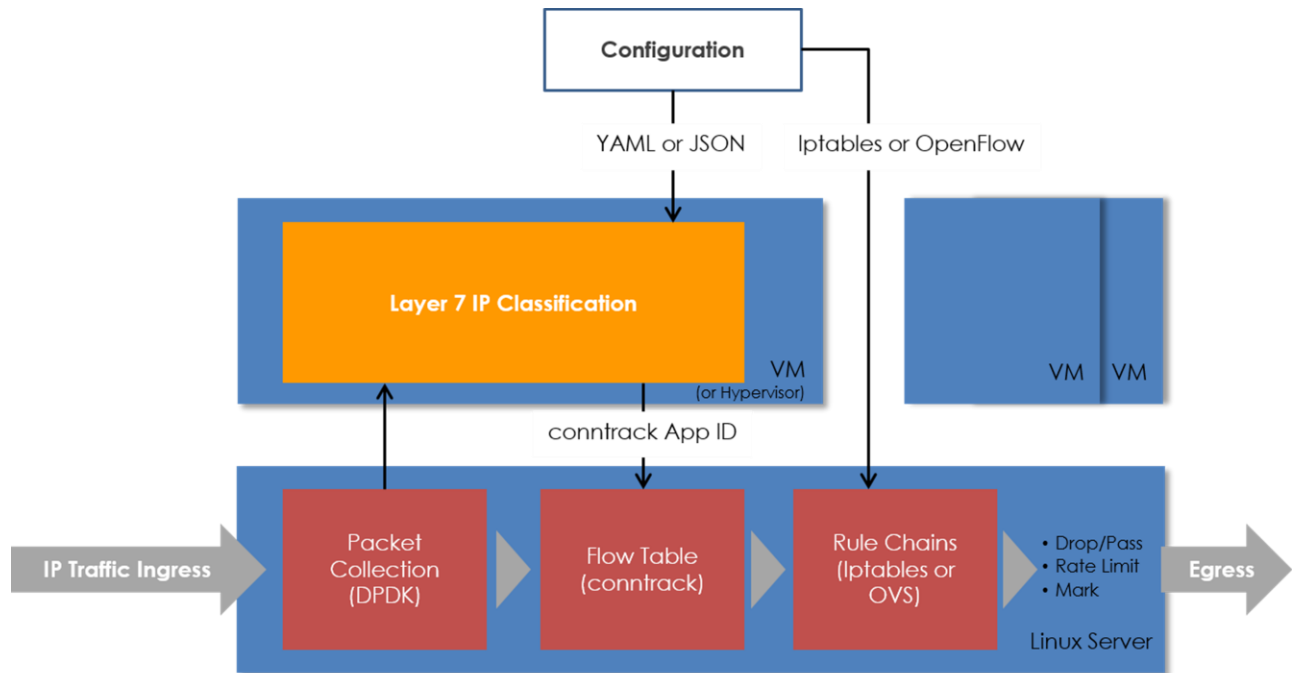
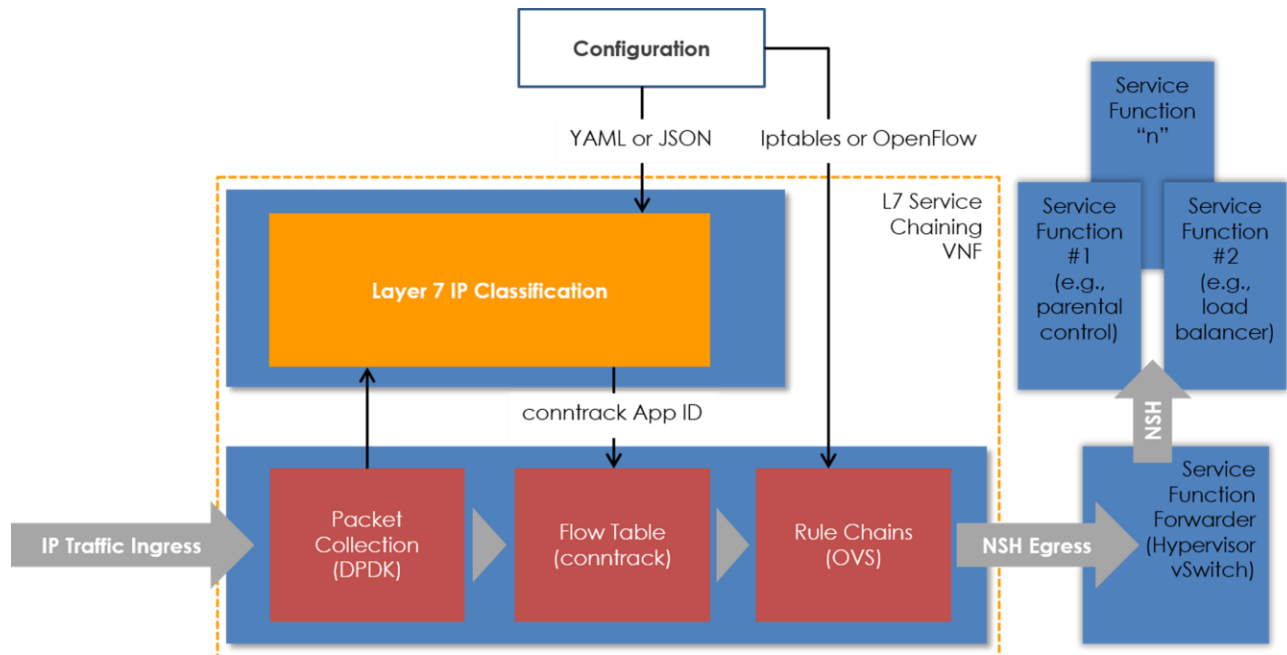Source: Qosmos

Service Chaining VNFs Using OVS & ODL

A third use case for Layer 7 classification in a vCPE context is service chaining. There are many variations in how service chaining is implemented and, in practice, operators will likely use a mix of approaches depending on the scenario. In this example, Layer 7 classification is useful because it helps to create optimal service chains by directing only the flows that need processing by a particular VNF.

Moreover, it is possible to do it in a fully virtualized environment using standard distributions of OVS and ODL. Using the same process outlined above, the L7 IP classification engine is used to update a conntrack flow table field with an App ID value to forward traffic to the correct processing path using an ODL controller. In this sense, to use the IETF's service function chaining (SFC) terminology, OVS can act as both a service classifier (with Layer 7 awareness) and service function forwarder. The basic concept is shown in **Figure 6**.

There are opportunities to use service chaining technologies to optimize the use of Layer 4-7 VNFs in many networking scenarios. vCPE is a leading candidate because operators want to enable customers to order additional services (VNFs, in essence) from customer portals to add value to the basic VPN connectivity service. Being

able to do to this in virtualized environment gives tremendous flexibility to evolve the solution as new capabilities and technologies penetrate the market.

Figure 6: Layer 7 IP Classification for Service Chaining

Source: Qosmos

A Note on Encrypted Traffic

It is now the case that somewhere around half of Internet traffic is encrypted (and a similarly large proportion of enterprise WAN traffic), and this is set to continue to increase. By some estimates, easily 80 percent of traffic could be encrypted within a couple of years.

A Layer 7 IP classification engine can help manage this traffic without compromising privacy and without breaking the encryption (which, in any case, is extremely difficult or might not be desirable) using a combination of techniques to infer the content of packets. If a stream is SSL encrypted, for example, a look at the certificate during the handshake process may give an idea that it is, say, a Google Mail service, or YouTube or Netflix. For peer-to-peer sessions, such as BitTorrent, IP addresses are sent in the clear before the torrent session is established, enabling the classification engine to determine what the flow is likely to be. And proprietary formats, such as Skype, often exhibit patterns in the flow that enable identification by the classification function.

These aren't foolproof methods, but they are sufficient to have a useful impact on traffic management capabilities.