RADCOM

Efficient, on-demand network tapping and smart probing for a 5G-ready cloud

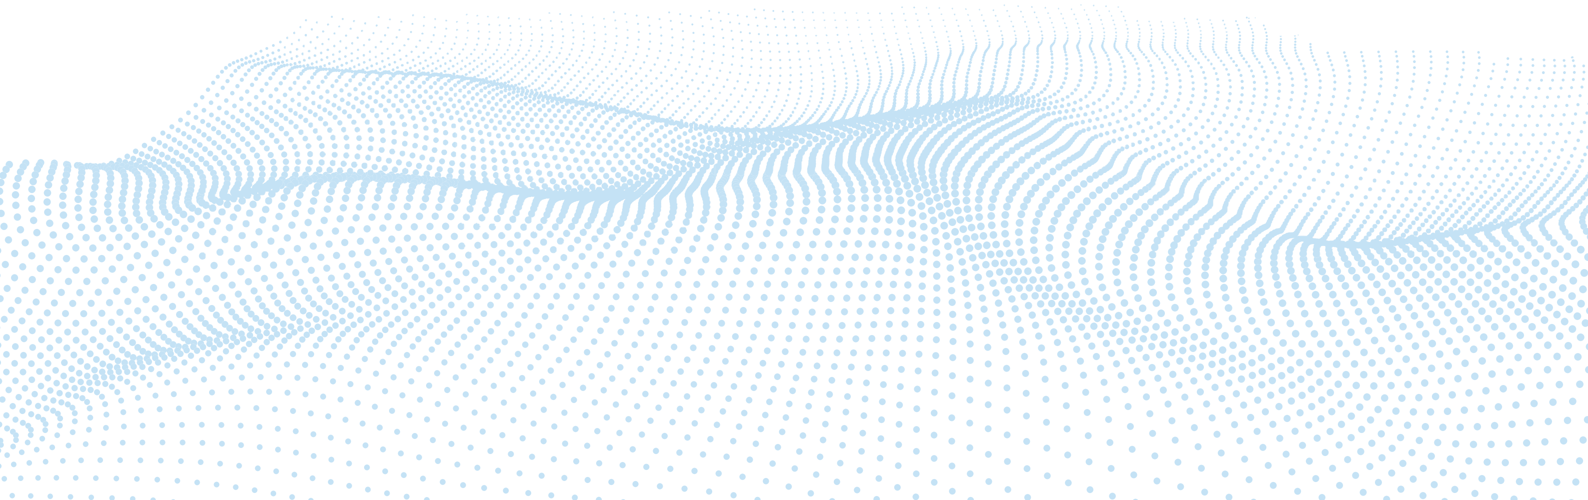

Contents

1. Introduction

Before Network Function Virtualization (NFV), network tapping for network monitoring or security purposes used to be relatively simple as the operator would have access to the data via a physical tap connected to a physical link. With physical links, the operator would always have complete network visibility. However, virtualization used in a 5G-ready cloud creates new network blind spots and makes gaining full network visibility more challenging;

- In a virtual network, a substantial percentage of the traffic never hits a physical link with most of the virtual machine (VM) to VM communication being hidden from physical taps and multiple virtual network functions (VNFs) being installed on one server. This traffic referred to as east-west traffic can make up a substantial part of the virtual traffic and so creates significant blind spots for the operator. This lack of network visibility will only increase as 5G continues to roll out and more and more functions are virtualized.
- When making copies of virtual traffic for 5G monitoring and security purposes, there are compute resources required and virtual tapping needs to fit in with the operators' goal of maintaining network efficiency and streamlining the network performance. Duplicating and forwarding packet data for continuous monitoring of all the traffic, all the time is not practical or efficient.
- Large-scale operators are deploying distributed networks with multi-tenant cloud environments. To streamline operations, ensure end-to-end service quality and be scalable these geographically separate environments (e.g., regional, or edge deployments) will need to be managed centrally.

While the importance of security and monitoring has not changed, the ability to capture and aggregate the traffic has. So, operators need to rethink their approaches to end-to-end visibility in an increasingly complex virtualized world. Ultimately, the same network tapping and monitoring tactics used in a physical environment will not work in a virtual one. With more and more operators transitioning to a virtualized network it's important that from day one operators deploy an efficient tapping solution that is tightly integrated with their probing solution to gain full network visibility and adjust things on-demand. A fully integrated solution will also enable operators to scale, manage and optimize their cloud networks much more effectively. With this critical foundation in place, operators can assure their migration to 5G more efficiently and implement a closed-loop approach to managing the customer experience and troubleshooting their services.

As operators transition to their 5G-ready cloud environments, on-demand instantiation of the network tapping and probing solutions will become more and more critical to more efficiently troubleshoot the network. This white paper explores the advantages of an on-demand solution and examines the most common traffic extractions used for network tapping and provides performance highlights for each option.

2. Virtual network tapping options and performance

Typically, in a virtualized network environment traffic is acquired and sent to service assurance probes and security tools by utilizing one or more of the following tapping methods:

Option	Tapping method	Overview
a	vSwitch Port Mirroring	vSwitch port mirroring is used to send a copy of all packets seen on one switch port (or an entire VLAN) to another switch port (i.e., for assurance purposes). This mirroring is used to copy the traffic of a VM or VM's to a single port and provides high-performance packet acquisition from a vNIC with the minimum number of CPU cycles. In promiscuous mode, a virtual interface connected to a vSwitch port group will be able to enter promiscuous mode and capture traffic from any other virtual interfaces connected to the vSwitch.
b	vSwitch Acceleration with Hardware Offloading	This method provides enhanced performance for VNFs that need maximum throughput and zero packet loss (such as virtual probes) by utilizing a SmartNIC for vSwitch Acceleration. Like SR-IOV (described below) in concept and performance, this method is based on standard mechanisms, does not require changes to the VNF and provides the ability to migrate VNFs quickly and easily. The main challenges are that per-packet processing in software can be inefficient.
c	Single Root I/O Virtualization (SR-IOV) Mirroring on the NIC	SR-IOV enables a Virtual Machine (VM) to create a Virtual Function (VF) on-top of a Physical Function (NIC) that directly bypasses the host networking stack, to provide line speed I/O to that VM. This ensures the required throughput can be accomplished using software techniques and allow access to hardware acceleration. As the packet arrives, the traffic is received at the physical NIC, and handed over to the Virtual Function (VF), therefore bypassing the vSwitch. Modern NICs (sometimes referred to as Intelligent-NICs or Smart-NICs) include an onboard programmable switch. This can be used to manipulate the traffic and replicate the handled traffic to additional VFs running on the same host or even to an external host. In an NFV environment on OpenStack, this process can be automated. Two automation options are feasible: • The NFVO (Orchestrator) can issue NIC specific driver commands to configure the mirroring of specific VLANs to a VF • Using OpenStack TaaS (Tap as a Service) to create the mirroring once the sending and tapped VNFC(s) specify the SR-IOV VLANs in their HEAT templates just before the ports are created using Neutron. When there is any change in the tapped or tapping VF instance, there is a need to make sure that the tapped and tapping VFs mirroring/paring will match

Option	Tapping method	Overview
d	TAP on Top of Rack (TOR) Switch	A TOR switch is a switch that handles Layer 2 and Layer 3 frame and packet forwarding, data center bridging and the transport of Fibre Channel over Ethernet for the racks of servers connected to them Traffic mirroring is performed by programming the TOR switch to mirror specific traffic to a target destination, such as a specific IP address or VIP, associated with a VF (Virtual Function). Specific VLANS can also be used to mirror only part of the traffic
e	Intra-VM tapping	Using an intra-VM software virtual-agent solution that is hosted on the compute node associated with the traffic (such as vEPC) which extracts the network traffic directly from the source. The extracted traffic can be filtered by the virtual-agent, using filtering rules via centralized management. Deploying an intra-VM agent reduces the amount of extracted traffic to the specific required packets/protocols/flows/sessions After extraction, the traffic is sent to a virtual Network Packet Broker (vNPB) for distribution to the service assurance probes using smart, session-aware load balancing

2a. vSwitch mirroring

A vSwitch is a software component associated with a hypervisor that functions like a Layer 2 hardware switch providing inbound/outbound and inter-VM communication. Every VM can communicate directly with every other VM on the same host through the virtual switch, without any inter-VM traffic monitoring or policy-based inspection and filtering. Intra-host VM traffic is handled internally by the vSwitch and does not enter the physical network.

Port mirroring is the most common method used to extract data from the VM and send to service assurance probes (also known as Switched Port Analyzer - SPAN). It is a software feature built into the vSwitch that creates a copy (mirror) of selected packets passing through the VM and sends them to a designated mirror port; where packets can be analyzed.

Pros	Cons
• Easy to implement – just activate an existing vSwitch feature • VNF agnostic – One tapping solution for all VNFs (Linux/non-Linux, DPDK / non-DPDK) • No integration required by the VNF vendor	• Limited filtering – typically by network port/VLAN • Requires more vSwitch resources and can bring the vSwitch "to its knees," if not appropriately accelerated.

vSwitch mirroring on the same host vs. an external host

Tapping the vSwitch is done by programming the vSwitch to mirror the traffic of a specific VF and sending it to another VF, on the same virtualized environment, such as for internal filtering and load balancing of the VF. The receiving VF component filters/manipulates the content and then forwards the filtered and load-balanced traffic to service assurance probes or other monitoring/security tools.

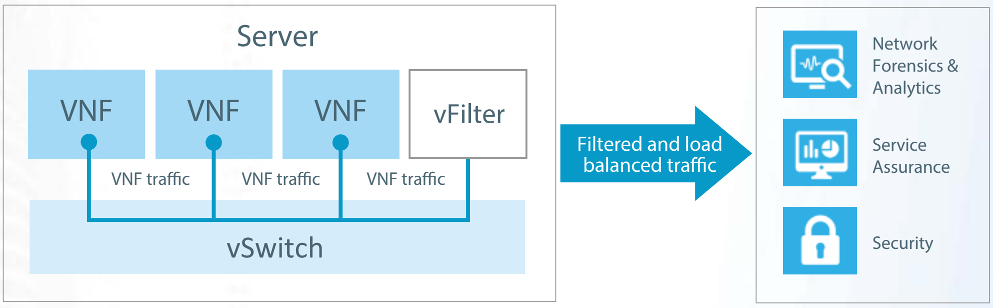

Figure 1 - Port mirroring with filtering and load balancing on the monitored VM host

Another tapping option allows the mirrored-traffic to be sent to an external filtering and load-balancing component, and later forward it to the service assurance probes for processing.

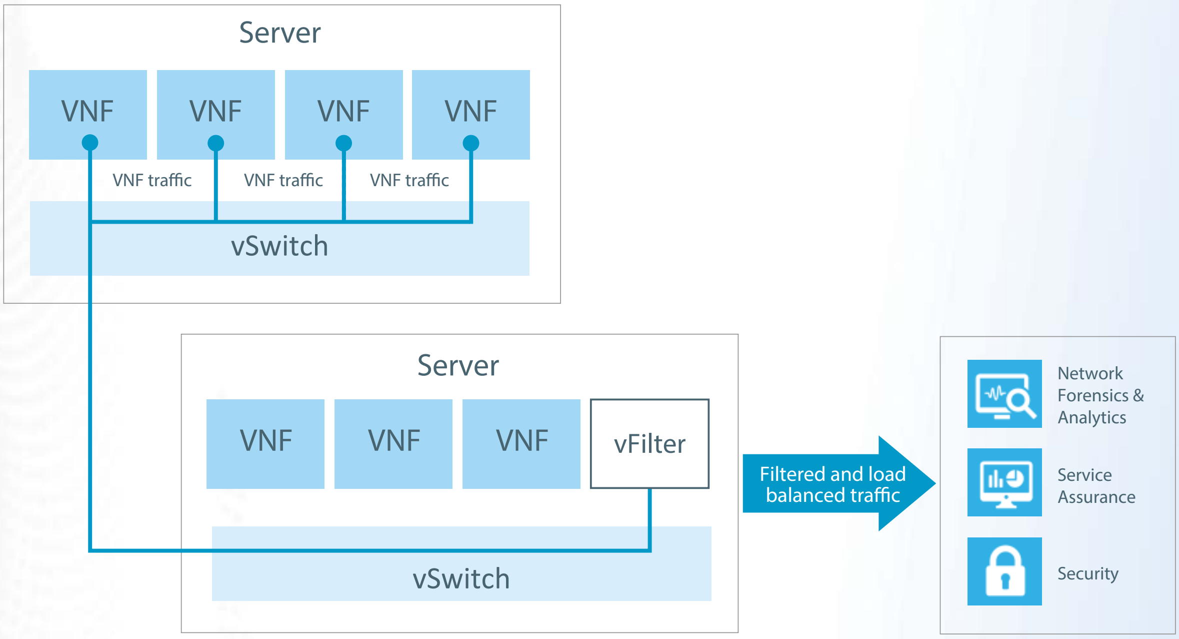

Figure 2 - Port mirroring with filtering and load balancing outside the monitored VM host

The amount of traffic going through the vSwitch that needs to be processed governs the specific option to be used, as outlined below.

On Monitored VM Host	On External Host
• Reduces network hops • Requires more vCPUs on VNF compute host • Requires more resources on the vSwitch compute host	• Additional network hop • Requires less vCPUs on VNF compute host • Requires fewer resources on the vSwitch compute host

2b. vSwitch acceleration with hardware offloading

This method provides enhanced performance for VNFs that need maximum throughput and zero packet loss and is based on standard mechanisms. The main downside to this method is that per-packet processing in software can be inefficient so a performance impact may be seen with small messages, as well as complex or multiple flows. The mirroring VF can run traffic analysis algorithms and observe the traffic of the VF being mirrored. This is achieved using an intelligent/SmartNIC adaptor that can offload specific packet processing and mirroring from the vSwitch. The existing mirroring support in the vSwitch, with the enhancement to the offloading logic in the adaptor driver, allows mirroring of the VF traffic to another VF, bypassing the vSwitch and by doing so resolving known vSwitch performance issues associated with traffic mirroring.

Using a SmartNIC provides an alternative to data acquisition in virtual environments and delivers results that show, it is no longer necessary to compromise on performance or agility, but possible to achieve both. Also, the solution minimizes the number of CPU cores required to the bare minimum. This reduces CAPEX and OPEX server costs while also providing the opportunity to consolidate virtual functions on the least number of servers as possible, further reducing OPEX costs.

Pros	Cons
• Mirroring, using the SmartNIC, to bypass the vSwitch, is still controlled by the vSwitch but requires much less packet processing resources from the vSwitch as opposed to the resources necessary when the SmartNIC is not available/utilized • The vSwitch must interact with the SmartNIC and utilize specific APIs to start/stop traffic mirroring	• There is a need to keep tracking the tapped traffic (network, port, trunk, VLANs) that needs to be tapped and mirrored to a specific VF-ID. This must always be visible to and orchestrated by the NFVO. The NFVO is responsible for matching the tapped/mirrored traffic and the mirroring VF • For automation, solutions such as Tap as a Service (TaaS) must be used to keep track of the mirrored-pairs

2c. SR-IOV mirroring on the NIC

SR-IOV enables network traffic to bypass the software switch layer of the hypervisor virtualization stack. Because the VF is assigned to a child partition, the network traffic flows directly between the VF and child partition. As a result, the overhead in the software emulation layer is reduced and achieves network performance that is the same performance as in non-virtualized environments.

Modern SmartNIC/Intelligent-NIC adaptors that were becoming a commodity and placed on the NFVI COTS and especially on the 5G MEC are places at the network edge, and strategic endpoints, that support virtual functions offload of different tasks, including rule-based packet processing, packet-filtering, offloading and mirroring. Latest intelligent NICs include a built-in switch that controls the traffic flow between endpoints. Traffic mirroring can be easily programmed into the SmartNIC with the latest SR-IOV Hypervisor sysfs[1] management interface with enhancements such as VLAN mirror, ingress mirror, egress mirror and more. These options are available with SmartNIC adaptors from the major vendors, such as Mellanox (Connect-X5 and higher), Intel (Fortville XXV710, etc.) and others.

The following drawing illustrates SR-IOV VLANs mirroring, where specific incoming and outgoing traffic on different VLANs between two virtual functions is mirrored to a third VF by programming the SmartNIC onboard switch to mirror the traffic associated with these VLANs to a specific VF-ID. The specific ID associated with a particular VF can dynamically change along the VF lifecycle. As a result, the NFV entity responsible for traffic mapping on the NFVI (usually the NFVO), must keep track of which traffic goes where and reprogram the switch to mirror the traffic to the newly created VF with its new ID. The latest update to the Tap as a Service (TaaS) OpenStack project supports the automation of such mirroring so that traffic tapping along the VF lifecycle is continuous and manageable by the NFVO.

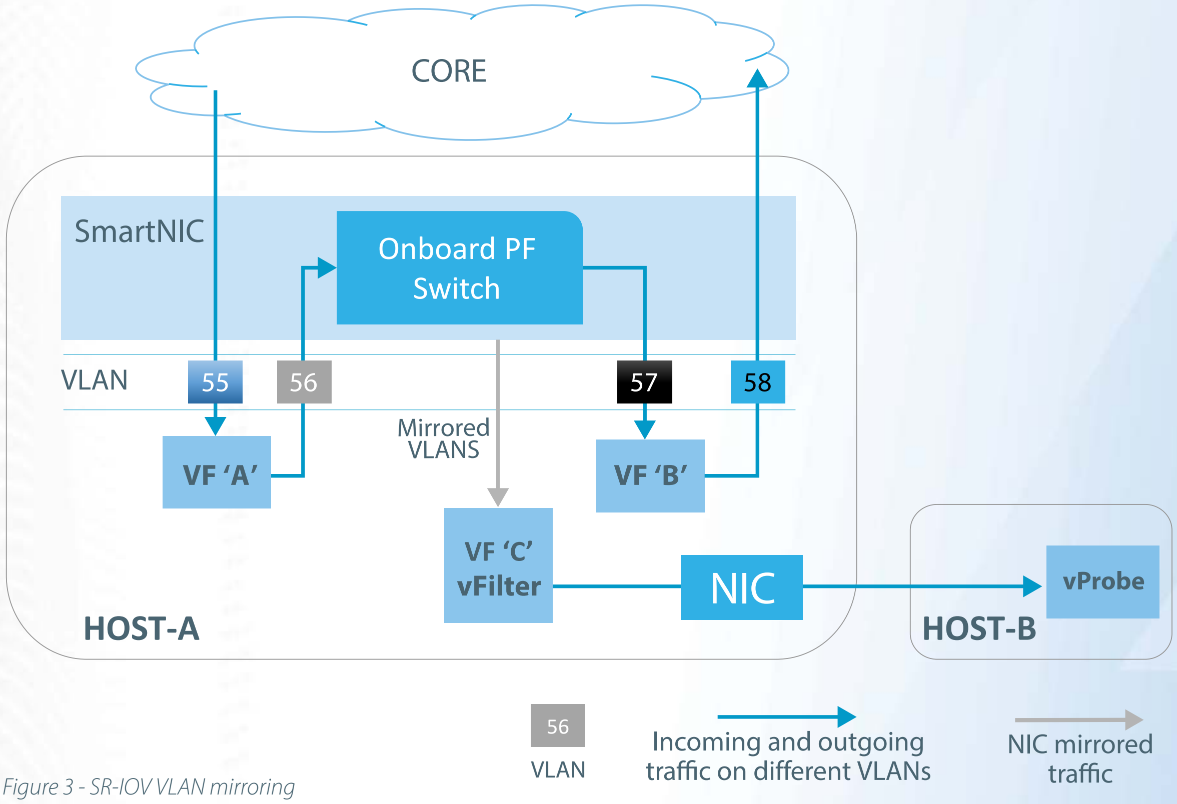

Figure 3 - SR-IOV VLAN mirroring

[1] A pseudo file system provided by the Linux kernel that exports information about various kernel subsystems, hardware devices, and associated device drivers from the kernel's device model to user space through virtual files.

Pros	Cons
• Makes a single PCI hardware device appear as multiple virtual PCI devices for VMs, enabling direct communication with the hardware NIC • Packets bypass the Host Hypervisor and vSwitch to deliver near wireline performance. This is highly suited to high volume user plane traffic capture • Low latency – the hypervisor software process is bypassed since the VM is directly attached to a hardware component • Scalability of the host is improved – by directly attaching VMs to VFs on the PCI; the CPU is bypassed enhancing CPU available to VMs • Trunk mirroring, VLANs mirroring, ingress mirroring and egress mirroring are now feasible and controlled via the HOST NIC driver, given the right access permissions • TaaS can be used to track and automate the mirroring process throughout the VF lifecycle	• SR-IOV ties PCI VFs on the physical NIC to VMs and VFs, which can be a departure from the decoupling of hardware from software • Features such as firewall filtering are still not available when using SR-IOV with OpenStack • If TaaS and an NFVO are not available, it would be difficult to automate, control and follow the VF-ID changes throughout the VF lifecycle

2d. TAP on Top of Rack (TOR) switch

Using the approach of placing physical tapping in the TOR switch to send the information to non-cloud native bare metal, non-NFV physical probing systems deployed outside of the cloud environment is not the right approach. This method is not automatically scalable via the cloud orchestration and will require probe vendors to deliver more physical boxes when traffic fluctuates. That means the operator will be required to deploy tapping on all the outputs of the TOR switch and use a costly deployment of network packet brokers in every TOR switch to collect the data coming from the VNFs and distribute to the physical probing system. With 80% of data center traffic being east-west, tapping on a TOR switch will mean the operator has blind spots in the network and will miss all the inter and intra-VM traffic.

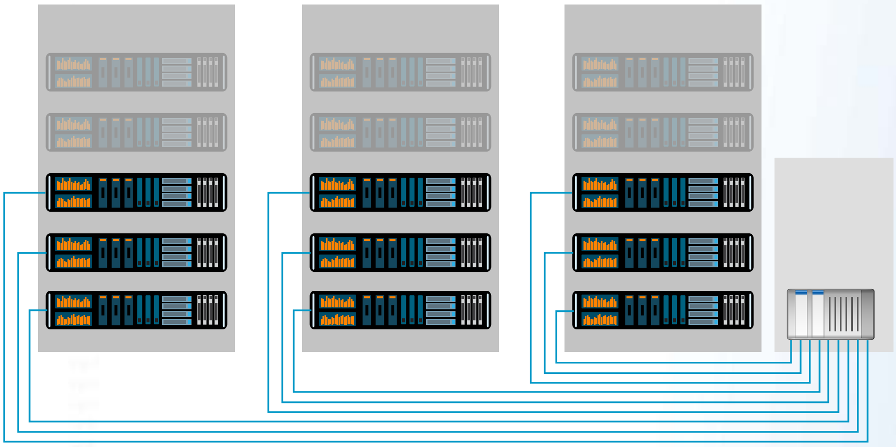

End-of-Row architecture

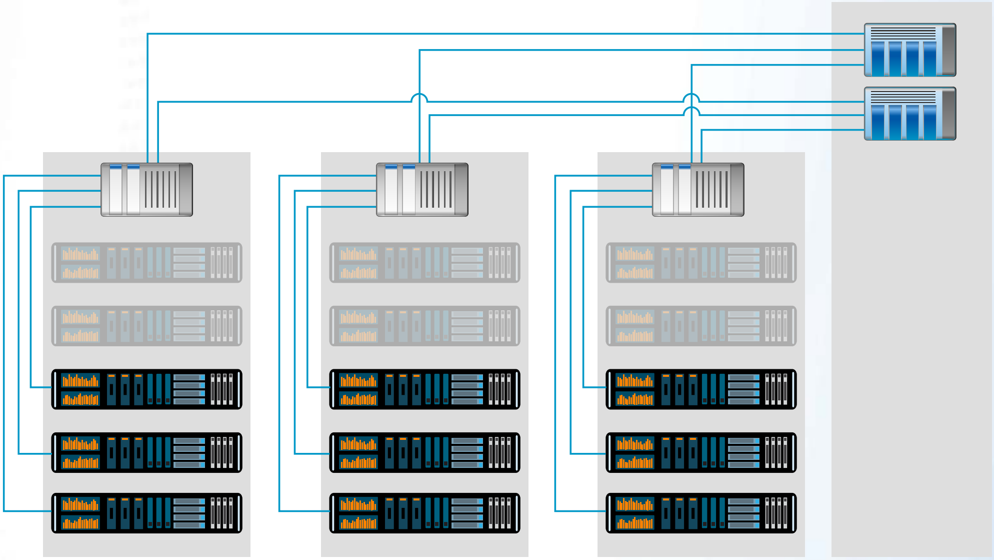

Top-of-Rack architecture

Pros	Cons
• For a legacy, non-NFV tapping solution this method can use a legacy tapping solution on TOR • This provides one tapping location for all the traffic that traverses the TOR switch but does not include inter VNF intra Host traffic	• In an NFV environment, inter VNF east-west traffic is not visible, and so does not appear on a TOR switch • If traffic is mirrored to the TOR switch for tapping, traffic trombone will waste vital network resource

2e. Intra-VM tapping

This option uses the monitored VM and VNF to be part of the method for copying packets moving through the virtual network interface controller (vNIC) from/to the NE virtualized components (VNFCs). A tap agent is deployed into the VM on each monitored NE. The agent will then capture the network data moving through the VM. The advantages of this method are:

- A light-weight monitoring process
- The tapping agent is collocated with the monitored VNF (on the same VM)
- The agent can be combined with a lightweight virtual network packet broker features for filtering, sampling, load-balancing, etc.
- Supports both VMs and containers (side-car)
- Supports Linux kernel or DPDK based VNFs

To be able to tap a VNF that runs on a specific VM, there are two main options available. These options do not require integration with the tapped-VNF software. The first option taps the vNIC ring buffer using libpcap, where the second option allows bypassing the Linux kernel using the DPDK-pdump technique, part of the Data Plane Development Kit (DPDK), that runs as a DPDK secondary process and can enable packet capture on DPDK ports.

To improve the packet processing performance, DPDK can be used on both the tapped VNF and the virtual probe side. DPDK can accelerate the overall packet processing operations needed in the vTAP.

The following diagram outlines the two options, where the tapping VNF is orchestrated to run on the same VM where the tapped VNF runs. The tapping-VNF deployment can be achieved using instantiation during onboarding or by incorporating the tapping-VNF into the VNF base-image as part of the standard deployment package.

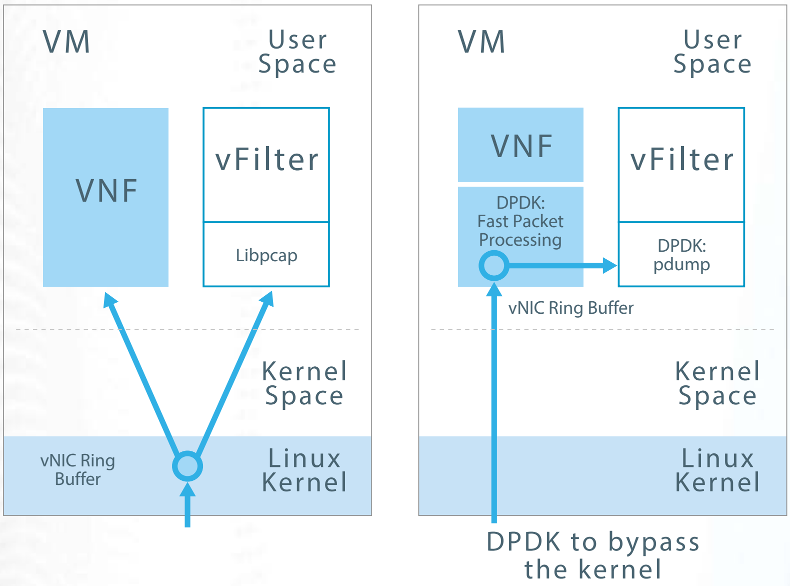

Figure 3 - Two VM-intrusive tapping methods

Pros	Cons
• Full visibility of all east-west VNF traffic traversing vSwitch • Easy to implement (by activating an existing vSwitch feature) • VNF agnostic and provides one tapping solution for all VNFs (DPDK/non-DPDK) • No integration with VNF vendor is required	• Need to allocate vCPUs for the tapping VNF in the tapped VNF VM • The NFV orchestrator (NFVO) needs to make sure that the tapping VM is instantiated on the same VM where the tapped-VNF runs • Custom agent/API is required for a non-Linux VNF

3. Network tapping challenges in an OpenStack environment

Many telecom operators are using OpenStack as their NFV platform of choice. One of the challenges in deploying OpenStack is the need to capture both south-north and east-west (intra-VM) traffic and currently there is no built-in solution for efficient, virtual tapping in an OpenStack environment. Current tapping options are:

- vSwitch port mirroring (e.g., Contrail)
- SR-IOV mirroring
- Network-stack mirroring inside the VM
- TOR tapping

The OpenStack industry is working on Tap as a Service (TaaS). However, there are currently challenges in using this methodology;

- It is unavailable in many existing NFV cloud deployments
- Has no filtering options and selective tapping
- Has no traffic encapsulation and filtering at the tapping point
- Has no flow control or integration with an SDN controller
- Orchestrating the tapping is not standardized and if not executed correctly can lead to the forwarding of traffic to destinations that are not up and running which will flood the network
- There is only manual provisioning or non-standard SDN control, and it does require tight integration with the NFVO

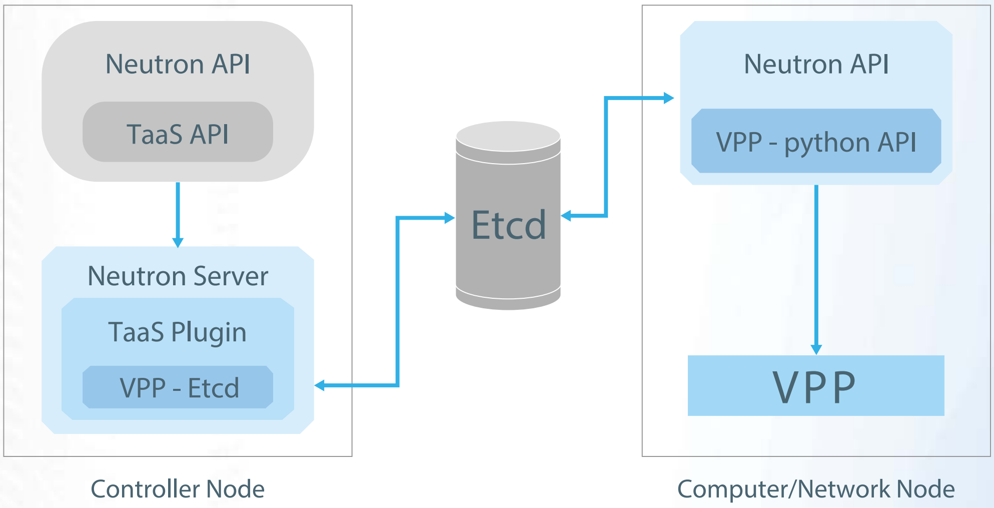

Figure 5 – One of the possible OpenStack TaaS implementations

4. Integrating smart, on-demand tapping into the NFVI

For network visibility to be efficient and beneficial for operators in the NFV and 5G era, the solution needs to be fully virtualized and instantiated together with physical and virtualized services as a Virtual Network Function (VNF) over an operators' NFVI infrastructure. A fully virtualized solution enables the operator to manage the exponential growth of user data traffic and scale.

Designing a cloud service to leverage the underlying virtualized infrastructure to ensure scale out elasticity, and multi-tenancy is essential. With cloud-native solutions, an operator can instantiate the virtual probing environment (virtual load balancers and probes) on the target monitored VNF and NFVI environment (vEPC, vIMS, etc.) or instruct a specific network element to extract the traffic and mirror it to the probing environment and scale out.

As operators continue transitioning to their 5G-ready cloud environments, on-demand instantiation of the network tapping and probing solutions will become more and more critical to more efficiently troubleshoot the network. Constantly monitoring all the network traffic, all the time is not cost-efficient and wastes valuable human, compute and network resources. These resources are most valuable and costly, especially at the network edge, where resources are tight.

A more effective solution is to deploy on-demand visibility and assurance with smart sampling and filtering of the traffic with intelligent load balancing that lets the operator zoom into certain data-sets (such as a specific service, subscriber group or protocol), troubleshoot and move on to the next high-priority task. Utilizing bare-metal based visibility and probes will not be able to serve probing-on-demand nor be dynamic enough for scaling.

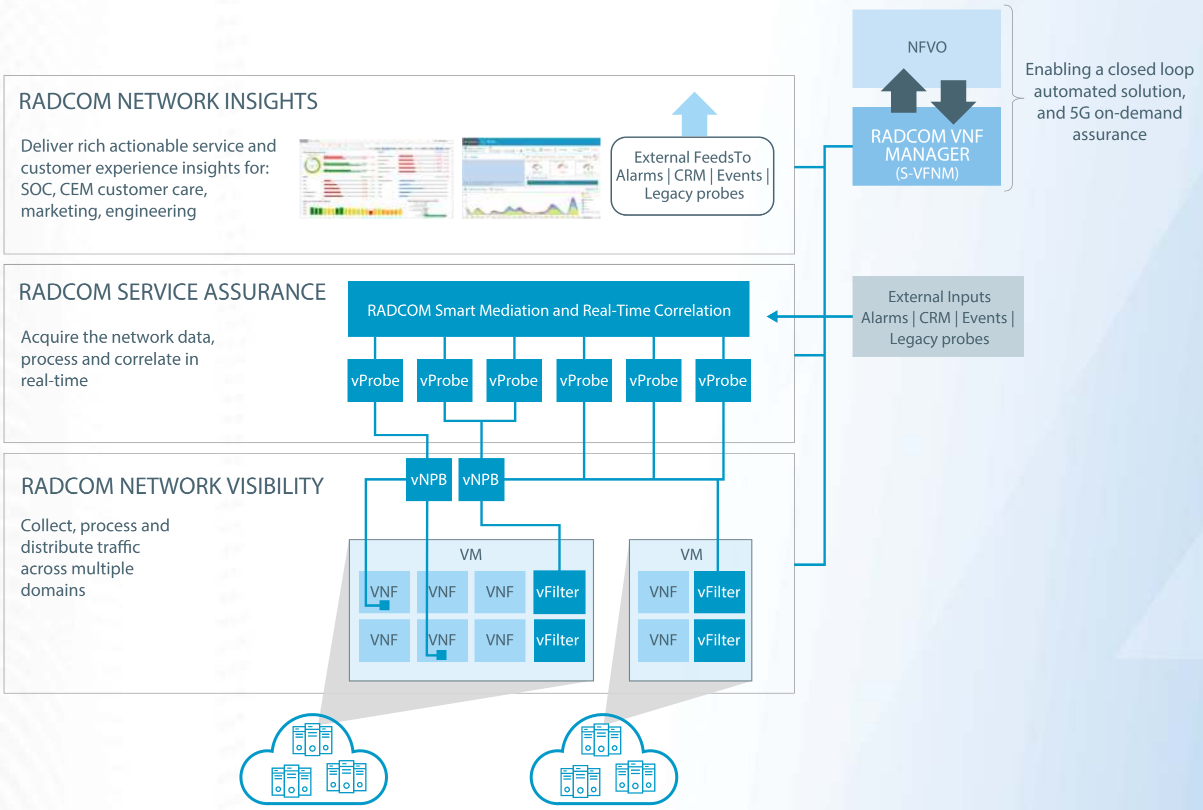
NFVO
RADCOM VNF MANAGER (S-VFNM)
Enabling a closed loop automated solution, and 5G on-demand assurance
RADCOM NETWORK INSIGHTS
Deliver rich actionable service and customer experience insights for: SOC, CEM customer care, marketing, engineering
External FeedsTo Alarms | CRM | Events | Legacy probes
RADCOM SERVICE ASSURANCE
Acquire the network data, process and correlate in real-time
RADCOM Smart Mediation and Real-Time Correlation
External Inputs Alarms | CRM | Events | Legacy probes
vProbe
vProbe
vProbe
vProbe
vProbe
vProbe
RADCOM NETWORK VISIBILITY
Collect, process and distribute traffic across multiple domains
vNPB
vNPB
VM
VNF
VNF
VNF
vFilter
VNF
VNF
VNF
vFilter
VM
VNF
vFilter
VNF
vFilter

By deploying integrated network visibility and service assurance in the same NFVI environment of the monitored VNF allows an operator to gain the following benefits:

- Probing is executed next to the monitoring point, so there is no need to send massive amounts of data out of the cloud
- Utilizes techniques to offload the Open vSwitch (OVS) using SR-IOV, or using direct SR-IOV traffic mirroring with SmartNIC, thus not overloading the NFVI when probing
- Gain dynamic scalability in/out as part of the orchestrated life cycle of the service or VNF being monitored
- Deploy tapping agents inside the VM of the monitored-target VNF to tap on the virtual interface and filter network traffic at the tapping point

By deploying an on-demand, smarter visibility layer operators can more efficiently troubleshoot their network, closer to the tapping point thus moving some of this functionality from the service assurance solution. Troubleshooting network issues at the tapping point enable operators to perform in-depth protocol analysis at the raw data level, drill down into any network element, protocol or message type and smartly capture filtered packets currently being transmitted through the network or examine historical data. The service assurance solution can still perform the heavy lifting, but troubleshooting at the tapping point is highly resource efficient and provides for a highly optimized network.

5. Conclusion

For telecom operators eradicating network blind spots in the age of virtualization and 5G, new challenges have emerged and for operators choosing the best, and most cloud-efficient tapping and probing strategy depends on multiple parameters such as the operators' cloud environment, the service being monitoring, the Virtualized Infrastructure Manager (VIM) software implemented as well as the operators' long-term goals. Furthermore, as described, the tapping performance varies between the different tapping options and no one option provides operators with the ultimate way to extract traffic on a 5G-ready cloud network. However, deciding on the right network visibility solution and tapping methodology is critical to ensure a smooth transition to a 5G-ready cloud and enable operators to transition to a dynamic, closed-loop approach to customer experience management.

As a market leader in virtualized service assurance and network visibility, with multiple large-scale production systems already implemented, RADCOM offers a fully virtualized, end-to-end solution from virtual tapping to business insights as well as deep cloud expertise specifically for telecom operators. RADCOM's team of experts can help operators choose the most suitable tapping and probing strategy to help achieve their business goals.

For more information on virtual network visibility, visit www.radcom.com

RADCOM

The information, content or material herein is provided "AS IS", is designated confidential and is subject to all restrictions in any law regarding such matters, and the relevant confidentiality and non-disclosure clauses or agreements issued prior to and/or after the disclosure. All the information in this document is to be safeguarded and all steps must be taken to prevent it from being disclosed to any person or entity other than the direct entity that received it directly from RADCOM.

The text and drawings herein are for the purpose of illustration and reference only.

RADCOM reserves the right to periodically change information that is contained in this document; however, RADCOM makes no commitment to provide any such changes, updates, enhancements or other additions to this document to you in a timely manner or at all.

Publication Date: April 2019