

# Characterizing Telco Workloads in Public Cloud Infrastructure

## Cloud Native Transformation - Evaluating Public Cloud Infrastructure for Network Intensive Workloads

A thorough understanding of centralized and edge cloud infrastructure characteristics provides a foundation for planning cloud native deployments of telco workloads such as 5G. Both fundamental (bottoms-up) evaluation as well as testing with representative workloads (top-down) is needed for optimizing performance of network intensive applications and evaluating deployment options.

### Introduction

The successful migration of 5G services to a public cloud platform will require accurate characterization of the target cloud infrastructure to forecast performance, make total cost of ownership (TCO) decisions and support operational planning.

- Communications service providers (CommSPs) need to assess the feasibility of cloud native deployments as well as to forecast workload performance and TCO.
- Vendors of cloud native network functions (VNFs/CNFs) need to understand infrastructure performance limits, as well as redundancy and scaling characteristics to optimize offerings.
- Network and security services vendors need to evaluate deployment models considering application performance and overall TCO for services such as secure access service edge (SASE).

Currently public clouds host mostly enterprise IT, web scale and other compute intensive workloads. 5G and other telco services, however, have demanding networking requirements such as throughput, packet loss and latency, that can be challenging to achieve in cloud infrastructures. Characteristics are constantly evolving and vary significantly between clouds, including:

- Throughput and packet-per-second limits for relevant traffic profiles.
- Average, maximum and percentile latencies, and their variabilities.
- Workload scaling characteristics considering suitable placement policies and resource availability.
- Other networking characteristics such as timing accuracy and support of various networking protocols.

Networking limits are dependent on the chosen VM instance type and the traffic data path that may remain internal to the cloud or connect externally. External connections may be to the internet or to edge sites, which can be “on-prem” (e.g., Amazon Web Services\* (AWS) Outposts) or “off-prem” (e.g., AWS local zones). Costs of traffic traversing or egressing a cloud are generally very significant for large aggregate flows and this may heavily influence TCO for many telco services. Optimizing network intensive applications and deployment costs is influenced by characteristics such as:

### Table of Contents

Introduction .....	1
Cloud Benchmarking.....	2
Measurement Metrics.....	3
High Performance Networking ..	3
Fundamental Testing .....	3
Workload Testing .....	4
Conclusions.....	5

1. VM instance type and size (number of vCPUs).
2. Network interfaces exposed to the VM with the number of interfaces varying by instance type and size.
3. Tx/Rx queue limits (packet-per-second, number of Tx/Rx queues and their scaling characteristics).
4. Rx versus Tx ratios and the number of flows (limits can vary based on traffic direction and actual traffic patterns).
5. Flow of traffic across gateways such as internet and peering gateways.
6. Location of instances based on network topologies and physical locations based on placement policies.
7. Non-uniform memory access (NUMA) architectures and corresponding workload affinity controls.

Infrastructure topology and workload placement also have a major bearing on latency and latency variation. Average latencies, as well as outlier latencies, influence the behavior of many telco applications. Smaller instances may have less deterministic latency characteristics since platform resources are normally shared between instances which can be reduced or avoided by using dedicated instances or hosts. In a NUMA architecture, vCPUs may have unequal network paths depending on network interface card (NIC) attachment and CPU architecture. From a workload perspective this could manifest as an additional latency penalty depending on scheduling of those instances.

## Cloud Benchmarking

Optimizing performance of network-intensive applications in cloud requires a fundamental understanding of the infrastructure characteristics. The purpose of cloud benchmarking is to:

- Understand performance and scaling limits due to infrastructure characteristics that can be hardware dependent and/or enforced by the cloud service provider.

- Assess the suitability of virtual machine (VM) instances and exposed CPU architecture features.
- Compare infrastructure options (including VM instances) with TCO models to assist in making design and procurement decisions.
- Assess workload optimization techniques available across VM instance types and specific cloud features that impact performance.
- Standardize performance testing methods and tools for comparing telco clouds.
- Develop performance diagnostic techniques for cloud native multi-vendor environments.
- Contribute to the relevant industry body-of-knowledge and skillsets for characterizing telco cloud infrastructure.

The methods and tools for characterizing networking appliances used in today’s telco infrastructures are well established. Cloud-hosted infrastructure, however, brings new challenges to the characterization and evaluation process. Interpreting performance test results using representative workloads may be misleading without first understanding fundamental infrastructure and instance characteristics, including:

- Tx/Rx queue behavior and limits.
- Interface scaling and limits.
- Instance scaling and limits.
- Supported network offloads.

In public cloud environments, hardware traffic generators are not available, necessitating the use of software traffic generators that have limitations not found in their hardware counterparts. The limitations of a particular software traffic generator must be considered to ensure accurate interpretation of test results. Name-brand hardware-based measurement tools are trusted for meaningful apples-to-apples comparisons and are used for commercial decisions. Table 1 below contrasts some lab and cloud benchmarking differences.

BENCHMARKING IN THE LAB	BENCHMARKING IN THE CLOUD
Visibility and full control of infrastructure and compute platform (driver versions, BIOS parameters, hypervisor version, NIC, network topologies, switch/router configurations, etc.).	“Black box” infrastructure – no visibility or control of underlying infrastructure. Limited visibility of compute platform configuration.
Dedicated compute platform.	Platform (instance) is shared with other tenants, unless reserved (i.e., paying for a full machine) or using bare-metal instances.
Visible physical data path topology (NICs, switches, gateways, etc.).	Data path is hidden.
Consistent latency between test nodes.	Significant and inconsistent latencies between deployments or even test runs.
Hardware and software traffic generators available and results can be compared.	Software traffic generator only.
Network interface characteristics well understood based on foundational NICs.	Unexpected network characteristics due to proprietary network interfaces and policies like traffic shaping.
Workload affinity controlled at server, socket, core and thread levels.	Workload affinity subject to instance placement policy that varies by cloud provider. No control of workload placement (server, core, thread) for smaller instances. No control of core frequency.

**Table 1.** Differences between benchmark testing in a lab environment vs. in a cloud environment.

Unlike in physical datacenter environments, public cloud does not provide access to physical ports and L2 / L3 packet forwarding may occur in virtual or physical infrastructure components. Lack of physical NIC access inhibits the usual visibility of packet flows on the network, impacting integration and validation efforts. Furthermore, variations in available interfaces between cloud environments may significantly impact workload portability between public clouds.

### Measurement Metrics

Telco applications are usually far more sensitive to lost packets and high latency than enterprise workloads. Zero packet loss measurements may not be achievable in public cloud infrastructures and therefore acceptable packet loss as well as latency-related thresholds should be specified based on the workload and use case.

Round trip latency is inherently influenced by the traffic generator as well as network infrastructure elements such as gateways. Using identical instances for the traffic generator and the device under test (DUT) may be useful since they will likely have similar latency characteristics. Spurious interrupts in virtualized platforms can impact latency measurements, however there are techniques for mitigating these effects. Latency measurement metrics should include average latency,

maximum latency and 99th percentile latency. The placement policy used must always be specified as it can have significant impact on packet latencies.

### High Performance Networking

Today’s highest networking performance VM instances use DPDK with SR-IOV. An alternative to DPDK for fast packet processing is AF\_XDP which was introduced in Linux kernel 4.18 and is supported by multiple NIC vendors. AF\_XDP is a raw socket that can fully abstract the infrastructure so that workloads in the guest OS do not require modification. There are various modes of AFX\_DP of which “zero copy” gives highest performance however this is not yet available in popular clouds.

### Fundamental Testing

“Bottoms-up” (without a representative workload) throughput tests can establish instance network performance limits and ability to scale given instance configuration parameters and traffic profiles. Tx traffic and Rx traffic limits may be different while Rx traffic influences Tx traffic limits in some clouds. Various traffic configurations should therefore be tested as shown in Table 2.

TRAFFIC CONFIGURATION	DESCRIPTION
Rx (only) throughput	Measures instance ability to receive packets without any traffic being generated. Note that limits may be enforced in the NIC and hence traffic received by the instance and traffic available to the application may be different.
Tx (only) throughput	Measures the instance's ability to send packets without any traffic being received.
Tx throughput with matched Rx traffic	Packet forwarding, i.e., Tx and Rx, are balanced (all packets received are transmitted).
Tx throughput with over-subscribed Rx traffic	One or more sources of traffic can either intentionally or unintentionally oversubscribe Rx which may impact the Tx capability of the instance.

**Table 2.** Test configurations for alternate traffic profiles.



OBSERVATION MIGRATING WORKLOAD TO CLOUD	5G-UPF	VCMTS	IPSEC (VPP)
Application requires modifications to work with Cloud NIC	X	X	
Difficult to interpret results and achieve maximum throughput.	X	X	
Per core throughput significantly impacted by CPU cycles and cost of Tx.	X	X	X
Per-core throughput is significantly impacted by per flow limits.			X
Per-core throughput is significantly impacted by queue limits.	X		X
Significant impact from low packet per second limits.	X	X	X
Impact from unequal Tx and Rx capabilities of NIC.		X	
Scaling limited by number of network interfaces.		X	
Scalability limited by available bandwidth.	X	X	X
Application requires modification to achieve good performance.		X	

**Table 3.** Characteristics of three telco workloads.

### Workload Testing

A “top-down” approach to testing requires a representative workload and traffic profile. Public cloud instances have diverse compute, storage and networking offerings with many constraints, some of which are discussed above. Workload behavior is a function of both the instance and other infrastructure characteristics.

Key considerations when profiling workloads include configurability of the workload pipeline architecture and the deployment model. The presence or absence of NIC offloads can impact the application deployment model. For example, an offload may enable a run-to-completion model instead of a pipeline model. Protocol offloads such as checksums are also commonly available. The protocols and number of flows may be different in upstream and downstream directions. Tunnelling protocols and the associated flow count for encapsulated traffic can have a significant impact on scalability of workloads causing possible mismatches between ingress and egress packet rates.

Scaling characteristics for interfaces and queues must be understood and related to the minimum interface count per workload instance. Tradeoffs between interface and queue scaling must be considered. While workload scalability is of primary concern, the same limitations and scalability factors may impact the tester and therefore care must be taken to avoid inadvertently “testing the tester” or influencing an apparent DUT profile.

When running a workload for the first time in a public cloud infrastructure, an important step is to establish a baseline of packet throughput performance using representative workload traffic without employing the actual workload. This means instantiating the deployment infrastructure including the tester, supplanting the target workload with a simple application to terminate or reflect traffic, constructing packet traces with the correct MAC and IP addresses for the target virtual network CIDR range and then sending those packets via the tester to the DUT. Table 3 offers some observations from three examples of workloads that are candidates for migration to cloud native deployment models, i.e., 5G-user plane function (UPF), virtualized Cable Modem Termination System (vCMTS) and IPsec (based on vector packet processing).



## Conclusions

Each cloud platform has important infrastructure characteristics and differences that drive many data-plane applications to require refactoring to achieve high performance at the lowest cost footprint. Furthermore, reconfiguration may be needed across different instance types within the same cloud due to instance-specific variations in resources, bandwidth limits and latencies. Telco workloads are particularly sensitive to the many cloud infrastructure limitations not seen with typical enterprise IT workloads and hence characterizing cloud infrastructure is foundational to evaluating feasibility of migrating workloads to cloud native deployment models or developing new telco services using cloud native principles.

Performance testing in various public cloud infrastructures is challenging, requiring some changes in approach from traditional lab testing. Establishing new testing methodologies and a fundamental understanding of both the common and unique characteristics of each target cloud environment is essential. Familiar testing methods need to be adapted for cloud environments to ensure accuracy and for meaningful “apples-to-apples” comparisons. One key to success is the use of a bottoms-up approach to establish the primary characteristics and limitations of the individual cloud environments. This informs the testing of representative workloads which, in turn, reveals important top-down learnings for the deployment and performance optimization of other network-intensive cloud native workloads. Automating infrastructure deployment, test-workload lifecycle and performance testing (running tests, collecting and storing results) is required to effectively evaluate, compare, optimize, and debug cloud-native deployments.

Unlike typical telco data center environments, cloud service providers impose strict networking limits that can severely impact telco workload performance. This is a reality of architecting multi-tenant cloud infrastructures that have massive scale while trading-off numerous business and technical constraints. Cloud environments are relatively immature when it comes to hosting telco workloads as compared to data center workloads. All told however, hyperscaler environments are rapidly evolving to comprehend telco and edge compute requirements that necessitate continuous assessment to determine optimal application deployment models and capabilities in each cloud. This will lead, more generally, to seamless on-prem and off-prem edge deployment models with opportunities for more telco applications to be hosted in the multi-cloud and hybrid-cloud world. The evolution of cloud-native deployment models with the rise of 5G and edge computing are likely to have a significant impact on the evolution of cloud infrastructures that must be comprehended by VNF / CNF vendors and CommSPs who aim to leverage such cloud services.



### Notices & Disclaimers


Intel technologies may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

0522/TM/H09/PDF

 Please Recycle

351390-001US