White Paper

Virtualized Access Gateway Function

Wireline and Wireless Convergence Solution on 4th Gen Intel® Xeon® Scalable Processors and Dell PowerEdge R760

Dell PowerEdge R760 server powered with Intel® 4th Gen Xeon® processors and vAGF workload on Red Hat Enterprise Linux delivers up to 931 Gbps throughput

Authors

Sarita Maini

John M Morgan

Padraig Connolly

Contributors

Ai Bee Lim

Suresh Raam

Andrew Duignan

Hadi Hayek

Marco Castanheira

Mike Moore

William B Meigs

Bertrand Rault

Reza Mokhtari

Introduction

This document describes industry-leading price and performance of a key wireline and wireless convergence workload. The virtualized Access Gateway Function (vAGF) solution on a Dell PowerEdge R760 server is combined with best-in-class 4th Gen Intel® Xeon® Scalable processors (formerly codenamed Sapphire Rapids) and Intel® E810 Ethernet Network Adapters. It is intended for communication service providers who are planning and deploying virtualized 5G wireline and wireless convergence infrastructure on the latest Intel® Xeon® Scalable processors. Intel and Dell Technologies collaborated to maximize network performance with an enhanced NFVI implementation with infrastructure and technologies which delivered a total throughput of up to 931 Gbps with a demanding vAGF data plane proxy workload. The solution demonstrates ecosystem readiness of the hardened hardware, firmware, and software that enables end users to integrate their applications on top of a verified platform configuration to achieve both high throughput and deterministic performance that are required for 5G transformative workloads.

The Access Gateway Function (AGF) is a network element that can be used along with 5G Core to enable wireline and wireless network convergence. It is the result of a joint initiative between the Broadband Forum (BBF) and the GSMA's 3rd Generation Partnership Project (3GPP).

The AGF enables subscribers with wireline access (with legacy communication protocols) to connect to the modern 5G core network. It allows an easy path of migration for residential gateways and CPEs that can now be managed by the modern 5G Core. AGF provides connectivity from the wireline network to the 5G Core.

This white paper is a product of an ongoing partnership between Intel, Dell Technologies, and Red Hat to accelerate deployment of 5G NFVI solutions by enabling verified product readiness producing industry leading solutions with the lowest total cost of ownership (TCO) in terms of performance per watt as well as performance per dollar.

Table of Contents

NFVI Forwarding Platform

The NFVI Forwarding Platform workload-optimized solution is designed to minimize the challenges of infrastructure deployment and optimization for the best performance with balanced I/O across sockets for core-bound as well as I/O-bound workloads. It uses the latest Intel® Xeon® Scalable processors which incorporate unique features designed especially for virtualized network workloads, leading to impressive performance gains compared to systems based on prior Intel processor generations. The increased performance of 4th Gen Intel Xeon Scalable processors can significantly improve the capability for software-centric, carrier-grade virtualization which aids communications service providers in attaining and enforcing service level agreements and increasingly demanding quality of service requirements.

The new [Dell PowerEdge R760 Rack Server coupled with Red Hat Enterprise Linux] operating system delivers the ultimate performance and versatility, with balanced I/O design using optimal configurations available to meet the demanding specification requirements for Intel NFVI Forwarding Platform solutions. This flexible 2U server uses less space and is power efficient. The processors tested in the 2-socket server were Intel® Xeon® Platinum 8470N which are optimized for networking bandwidth with a base core frequency of 1.7G, 52 cores, 104 threads per socket, PCIe slots with 16 GT/s, and 8 channels of DDR5 memory capable of speeds up to 4800 MT/s.

Dell PowerEdge rack servers help you build a modern infrastructure that minimizes challenges and drives business success. Design optimized for a mix of demanding workloads with high database and analytics capabilities, the Dell R760 provides performance and versatility for demanding applications.

The AGF test setup consists of the following layers:

Hardware: Dell PowerEdge R760 with Intel® Xeon® Platinum 8470N processors

Operating System: Red Hat Enterprise Linux 8.7

Application: Intel® virtualized Access Gateway Function

Terminology

Abbreviation	Description
AIC	Add-In Card
AGF	Access Gateway Function
BIOS	Basic Input/Output System
DIMM	Dual Inline Memory Module
DPDK	Data Plane Development Kit
DRAM	Dynamic Random Access Memory
DUT	Device Under Test
PCIe*	Peripheral Component Interconnect express*
SR-IOV	Single Root Input/Output Virtualization
SSD	Solid State Drive
TPM	Trusted Platform Module
vAGF	virtualized Access Gateway Function

Reference Documents and Resources

Document	Document Number/Location
Wireline Access Evolution and 5G Fixed-Mobile Convergence	https://builders.intel.com/docs/networkbuilders/wireline-access-evolution-and-5g-fixed-mobile-convergence-1639769220.pdf
VBNG-VAGF.L.22.03.0-00072.tar.gz	764478
Intel® Reference Architecture for NFVI Forwarding Platform on 4th Gen Intel® Xeon® Scalable Processors on Red Hat* Enterprise Linux* with vAGF Workload	Intel® Reference Architecture for NFVI Forwarding Platform on 4th Gen Intel® Xeon® Scalable Processors on Red Hat* Enterprise Linux* with vAGF Workload
Intel® Ethernet Controller E810 Dynamic Device Personalization (DDP) Technology Guide	617015
Intel® Ethernet Controller E810 Dynamic Device Personalization Package (DDP) for Telecommunications Technology Guide	618651
Dell PowerEdge R760 Rack Server	PowerEdge R760 Rack Server | Dell USA

vAGF Overview

The Access Gateway Function (AGF) is a function that provides connectivity from a wireline Access Network to the 5G Core Network. Access Gateway Function (AGF) is the access point for subscribers, through which they connect to the internet and private networks. It provides critical subscriber management functions, such as authentication, IP address assignment, bandwidth allocation and accounting.

When a connection is established between the Customer Premises Equipment (CPE) and the AGF network function, the subscriber can access the broadband services provided by the telecom operator or Internet Service Provider (ISP). The role of the AGF is to aggregate traffic from various subscriber sessions from an access network and route it to the network of the service provider.

Since the subscriber directly connects to the edge router, vAGF effectively manages subscriber access and subscriber management functions such as:

- Authentication, authorization, and accounting (AAA) of subscriber sessions
- IP Address assignment
- Security
- Policy management
- Quality of Service (QoS) and Traffic Management

The vAGF is a virtualized software instantiation of what is typically a large ASIC-based fixed-function appliance usually located in a central office or metro Point of Presence (PoP). The vAGF is implemented as a set of Virtual Network Function (VNF) instances with each instance supporting a single Subscriber Service-Group, which typically contains hundreds of home routers/subscribers.

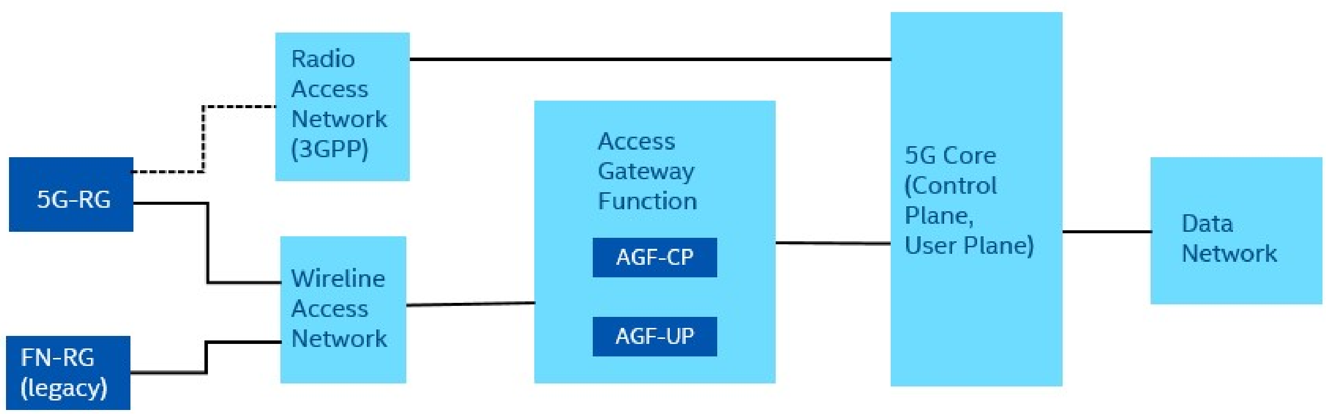

Figure 1. vAGF Overview

vAGF Data Plane

The vAGF Data Plane (DP), or User Plane (UP), is built around two packet processing pipelines - uplink (UL) and downlink (DL). The uplink data plane manages the flow of traffic from the end user's Customer Premises Equipment (CPE) to the core network. The downlink data plane handles the flow of traffic and data from the core network to the end user. It manages and schedules traffic to users attached to the AGF.

The vAGF DP has been implemented using the [FD.io VPP](FD.io VPP) framework and run as a containerized application using Kubernetes as an orchestrator. This white paper focuses on the data plane functions since the main goal is to show maximization of I/O for intense workloads like vAGF.

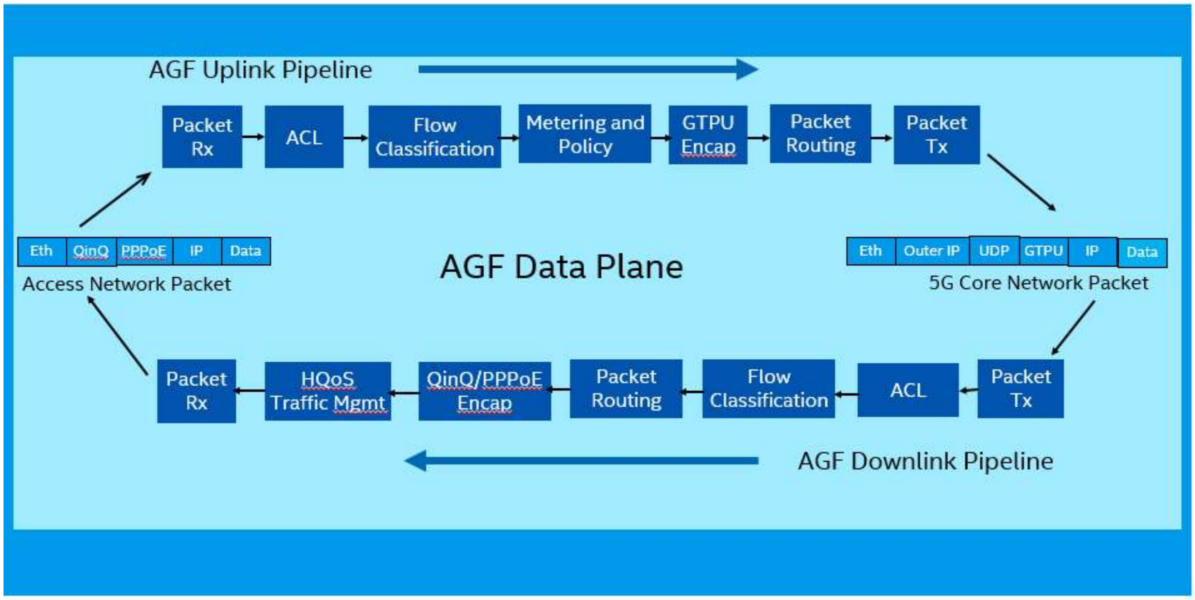

Figure 2. vAGF Data Plane

vAGF Uplink Pipeline Overview

The AGF uplink packet processing pipeline consists of the following functions:

Packet Rx (Receive): Packets from the wireline subscriber access network are received from the Network Interface Controller (NIC) ports using DPDK PMD drivers and sent to the next stage to begin packet processing.

Firewall/Access Control List (ACL): This stage employs an Access Control List (ACL) table to implement firewall policies (i.e., block rules) on the incoming traffic. This blocklist firewall has 150 block (random) rules. Table lookup operation is performed on each received packet, and in the case of rule match, the packet is dropped.

Flow Classification: This stage classifies each subscriber flow based on the source MAC address and double VLAN tags and strips the Q-in-Q header (and PPPoE header if PPPoE subscriber traffic is enabled).

Flow Metering and Policing: This function meters the subscriber traffic flows to determine the compliance with a service contract and applies traffic policing to enforce the contract. As a result, packets that conform to a specified rate are sent to the next stage of the pipeline while packets that violate the rate are dropped.

GTPu Encapsulation: At this stage, a GTP-U header is added to the IP packet

Routing: At this stage, an Ethernet header is added based on the route

Packet Tx (Transmit): Finally, the packets are sent out to the core network. With the help of DPDK poll mode drivers, packets are transmitted out of the system through the NIC ports connected to wireline core network.

vAGF Downlink Pipeline Overview

The vAGF Downlink packet processing pipeline consists of the following functions. The incoming downlink packet typically consists of an Ethernet frame with IP/UDP header. The outbound traffic will be an Ethernet frame with encapsulated QinQ VLAN and IP/UDP headers.

Packet Rx (Receive): Packets from the core network are received from the Network Interface Controller (NIC) ports using DPDK PMD drivers and sent to the next stage to begin packet processing.

Firewall/Access Control List (ACL): This stage employs an access control list (ACL) table to implement firewall policies (i.e., allow rules) on the incoming traffic. Table lookup operation is performed on each received packet, and in the case of rule match, the packet is permitted to the next stage. The Allow list firewall has 4K allow (subscriber flow) rules.

Flow Classification: This stage performs exact-match classification on the 5-tuple header fields (inner IPv4 source and destination IP address, UDP source and destination ports and IP transport layer protocol ID) of the input packets to identify the session and stores the session info as packet metadata to be used later in the pipeline. In addition to this, access network encapsulations are stripped off the packets at this stage. It first strips the GTP-U header and then classifies each subscriber flow based on the 5-tuple header.

QinQ/Q-in-Q+PPPoE Encapsulation and Routing: At this stage, packets are encapsulated with a QinQ header added to the inner IPv4 packet based on the flow ID and (and PPPoE header if enabled) and routed to the access network via the correct network interface port.

Hierarchical QoS Traffic Management: Each packet runs through a hierarchical QoS (HQoS) scheduler to ensure that thousands of subscribers can get the desired broadband capacity as per the service contract. It implements a 4-level HQoS with one pipe per subscriber (configured to allow all traffic to pass)

Packet Tx (Transmit): With the help of DPDK poll mode drivers, packets are transmitted out of the system through the NIC ports connected to access network.

vAGF Deployment

The Intel® AGF Data Plane Package can be used to install multiple instances of a vAGF data plane reference application in a Linux Container environment on a Dell PowerEdge R760 server on Intel Xeon Platinum 8470N processors.

The application and environment were used to evaluate the performance of a vAGF data plane on Intel® Xeon® based platforms. This is a POC evaluation application only and is neither intended nor is fully featured, hardened, or secured. Deploy in an isolated evaluation environment only.

The Intel® AGF application is benchmarked to understand how the deployment works in a real context. To get a more comprehensive view, two different configurations have been benchmarked - a symmetric traffic workload and an asymmetric traffic workload.

The following table and figure provide more key information points about the benchmarking setup for these two test configurations.

Each vAGF container instance has both a downlink and uplink pipeline. Each 100 Gbps Intel® Ethernet E810-2CQDA2 port has 4 vAGF container instances mapped to it.

Test Parameters	Symmetric Traffic Workload	Asymmetric Traffic Workload
Maximum number of vAGF Instances	40	40
Maximum number of Instances per socket	20	20
Number of vCPUs per instance	4	4
Max number of active vCPUs per socket	80	80
Number of VFs per instance	2	2
Number of Flows per vAGF instance	2	2
Max number of instances per E810-2CQDA2 100G NIC port	4	4
Traffic Line Rate per instance	25 Gbps 12.5 Gbps downlink max 12.5 Gbps uplink max	25 Gbps 22.25 Gbps downlink max 2.75 Gbps uplink (constant)
Traffic Ratio (Downlink : Uplink)	1:1	89:11
Packet Sizes	DL: 650 bytes UL: 650 bytes	DL: 504 bytes UL: 128 bytes
Acceptable Packet Loss	0.001%	0.001%

Table 1. Symmetric vs Asymmetric Traffic Workloads

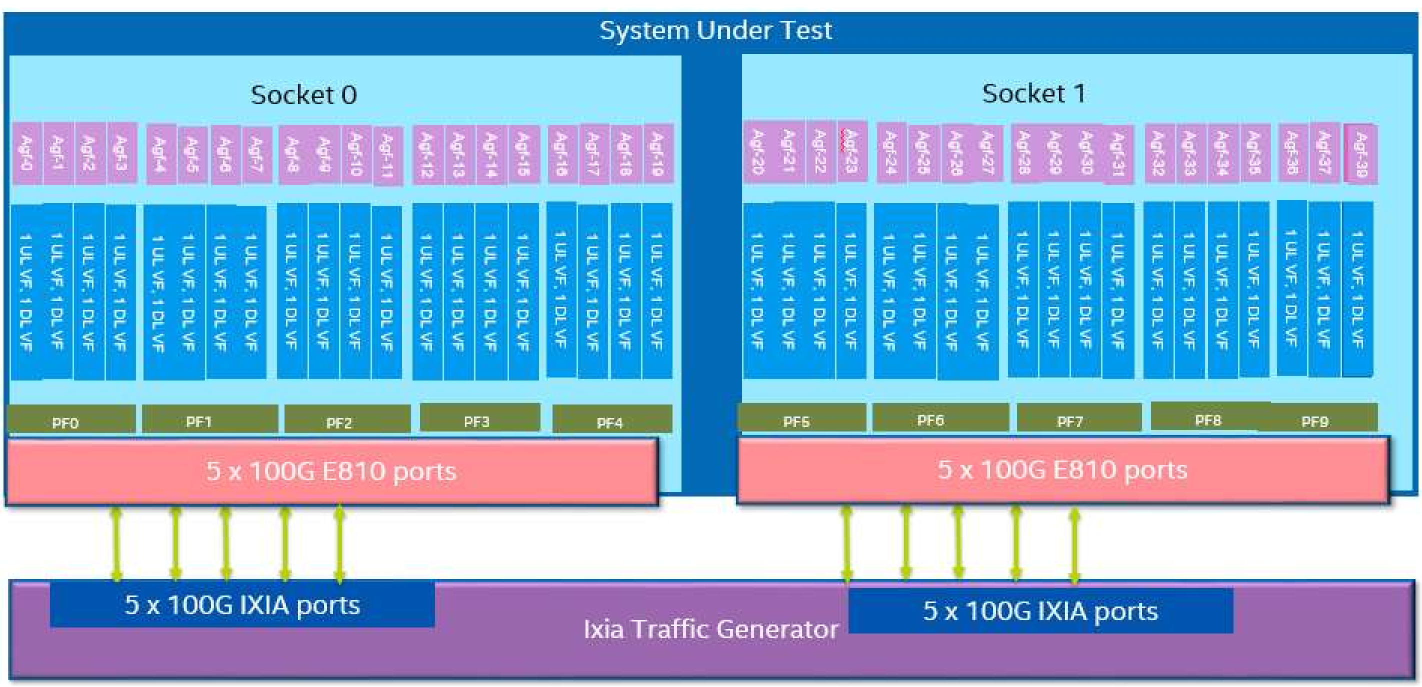

Figure 3. vAGF Test Setup

vAGF Benchmarks

vAGF Symmetric Traffic Benchmarks

The results of the vAGF Symmetric benchmarks on a Dell PowerEdge R760 server with two 4th Gen Intel® Xeon® 8470N processors and 1000 Gbps Line Rate (10 x 100G Intel® Ethernet Network Adapter E810-2CQDA2 ports), 650B UL/DL Packet Size and a traffic ratio of DL/UL 1:1 with a limitation of Packet Loss to be less than 0.001% with turbo disabled, show a Maximum Receive Rate (MRR) of up to 931 Gbps Rx L2/L3 throughput, which is 93.1% of the line rate and up to 959 Gbps Rx L1 throughput, which is 95.9% of the line rate.

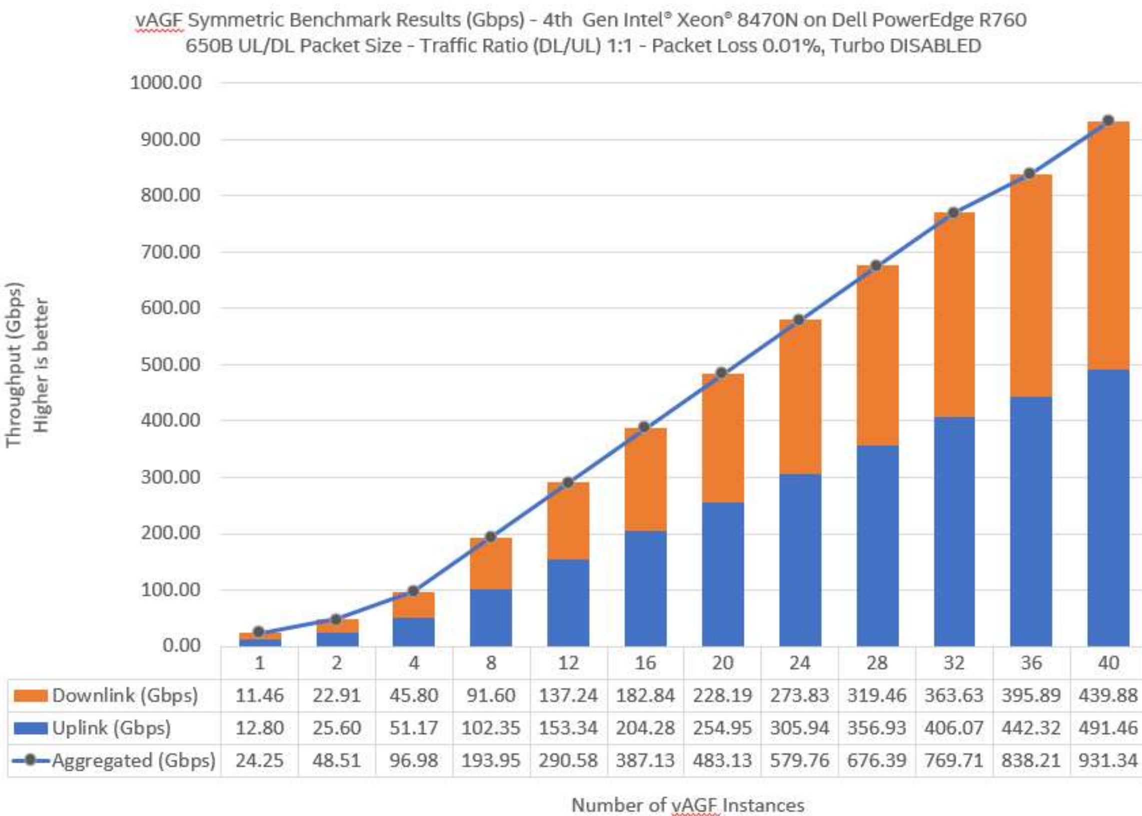

Figure 4. vAGF Symmetric Traffic Benchmarks [1,2]

vAGF Asymmetric Traffic Benchmarks

The results of the vAGF Asymmetric benchmarks on a Dell PowerEdge R760 server with two 4th Gen Intel® Xeon® 8470N processors with 1000 Gbps Line Rate (10 x 100G Intel® Ethernet Network Adapter E810-2CQDA2), 504B DL and 128B UL Packet Size and a traffic ratio of DL/UL 89:11 and a Packet Loss of less than 0.001% with turbo disabled, show a Maximum Receive Rate (MRR) of up to 682 Gbps Rx L2/L3 throughput, which is 68.27% of the line rate and up to 737.9 Gbps Rx L1 throughput, which is 73.79% of the line rate.

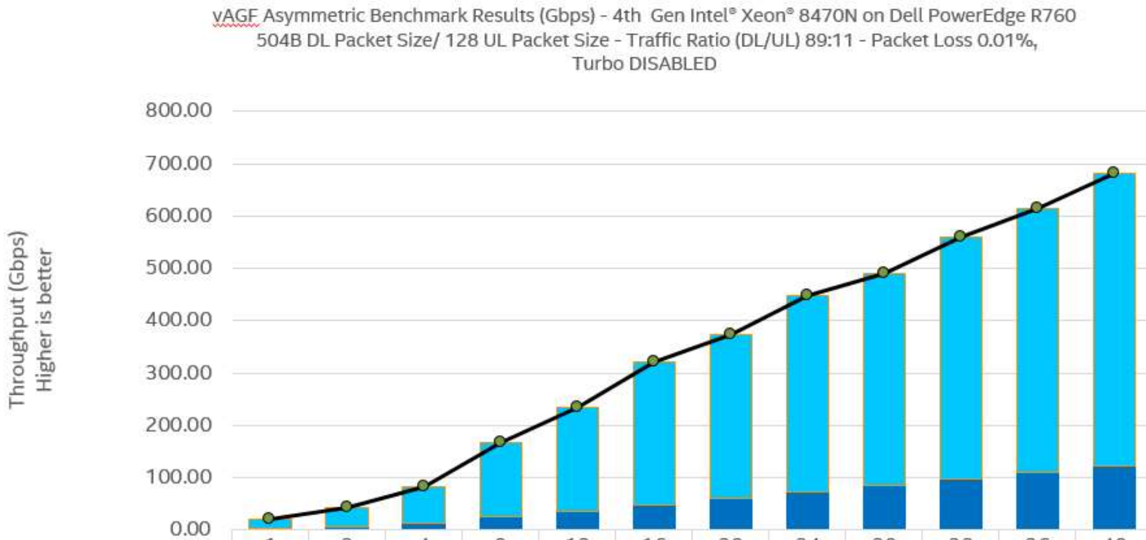

	1	2	4	8	12	16	20	24	28	32	36	40
Downlink (Gbps)	18.27	36.55	71.45	142.90	198.27	271.51	312.59	375.11	405.13	463.02	504.80	560.89
Uplink (Gbps)	3.05	6.09	12.19	24.38	36.57	48.76	60.95	73.14	85.32	97.51	109.70	121.89
Aggregated (Gbps)	21.32	42.64	83.64	167.28	234.84	320.27	373.54	448.24	490.45	560.53	614.50	682.78

Number of vAGF Instances

Figure 5. vAGF Asymmetric Traffic Benchmarks [1,2]

Summary

The collaboration between Intel, Dell and Red Hat enables CoSPs to quickly integrate their applications on top of a verified platform configuration to achieve both high throughput and deterministic performance. The virtualized Access Gateway Function benchmarks on the Dell PowerEdge R760 server with two 4th Gen Intel® Xeon® Platinum 8470N processors achieved an impressive I/O throughput of up to **931.34 Gbps** Maximum Receive Rate for L2/L3, which was **93.13%** of the maximum line rate of 1000 Gbps and up to **959.99 Gbps** Maximum Receive Rate for L1, which was **95.99%** of the maximum line rate. The packet loss was less than or equal to 0.001%, which demonstrated the performance determinism required for 5G transformative workloads The use of the Dynamic Device Personalization (DDP) feature in the Intel® Ethernet Adapter E810-2CQDA2 with the Intel COMMs DDP package contributed to the high performance. All the hardware and software components required to achieve these results are available today.

The new Dell PowerEdge R760 mainstream 2U dual socket stack server is a powerful and dependable server which brings PCIe Gen 5, DDR5 and built-in acceleration to the Dell PowerEdge server line. This high performing server with iDRAC 9 and its optimized settings, combined with the high core count of the 4th Gen Intel® Xeon® Scalable processors, which has architectural improvements, feature enhancements, and high memory bandwidth, provides a tremendous performance and scalability advantage compared to previous Intel® Xeon® processor generations, especially in today's NFVI environments. These processors are optimized for network, cloud native, wireline, and wireless core-intensive workloads, with up to 60 powerful cores and a wide range of frequency, features, and power levels. The Intel® Xeon® Platinum 8470N processors with 52 cores at 1.7 GHz core frequency, high DDR5 memory bandwidth and PCIe Gen 4 IO throughput has outstanding performance based on a balanced, efficient architecture that increases performance memory and I/O bandwidth to accelerate diverse workloads from the data center to the intelligent edge.

[1] Notices & Disclaimers

Performance varies by use, configuration, and other factors. Learn more on the Performance Index site. Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure. Your costs and results may vary. Intel technologies may require enabled hardware, software or service activation. © 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.

[2] Configuration

Test by Intel as of 05/07/23: System - Dell Inc. PowerEdge R760, 1-node, 2x Intel® Xeon® Platinum 8470N, 52 cores, HT On, Turbo Off, Total Memory 2048GB (32x64GB DDR5 4800 MT/s [4400 MT/s]), BIOS 1.2.1, microcode 0x2b000190, 10x Ethernet Controller E810-C for QSFP, 24x 1.1T ST1200MM0099, Red Hat Enterprise Linux 8.7 (Ootpa), 4.18.0-425.19.2.el8_7.x86_64, gcc compiler 8.5.0, vAGF 22.08, ice driver: 1.11.14, DDP ICE COMMS package version 1.3.40.0.

ADDENDUM: System Configuration

Dell R760 PCIe Riser Configuration

Configure Dell R760 with the supported PCIe riser configuration.

a. Used **Config5-1 (R1R+R2A+R3A+R4P (HL))**

b. R1R(HL): Dell PN: 4N0CV (Slot 1 and 2)

c. R2A(HL): Dell PN: 9X6JM (Slot 3 and 6) bottom 2 Low profile slots.

d. Dell PN: R3A(HL): 3PPC3 (Slot 5)

e. R4P(HL): Dell PN: 535MN (Slot 7)

f. Install Intel® E810-2CQDA2 (code name Chapman Beach) NICs in Slots 1, 2 of R1R(HL), slot 5 of R3A and Slot 7 of R4P(HL)

g. Install Intel® E810-CQDA2 (code name Tacoma Rapids) NICs in Slots 3, 6 of low-profile butterfly riser

h. Change PCIe bifurcation in BIOS settings from x16 to x8x8 for Slots 1, 2, 5, 7

iDRAC and BIOS Settings

System Information	
System Model Name	PowerEdge R760
System BIOS Version	System BIOS Version
System Manufacturer	Dell Inc.
System CPLD Version	1.0.5
BIOS Version	1.2.1
iDRAC Version	iDRAC9
iDRAC Firmware Version	6.10.39.00

BIOS Settings	
Workload Profile	NFVI FP Optimized Turbo Profile
Changes to NFVI FP Optimized Turbo Profile	Turbo Boost -> Disabled
	C1E -> Disabled
	C-States -> Disabled