Technology Guide

Telco Cloud, Cloud Native, Network Function Virtualization

Intel vAGF on VMware Telco Cloud Platform

Intel and VMware have collaborated to enable and develop a reference deployment for cloud-native network functions on the VMware Telco Cloud Platform. Focusing on broadband fixed-access transformation, the underlying hardware and software platforms have been optimized and enabled for ease of use, ease of deployment and high performance cloud-native network functions.

1 Introduction

In recent years, communication service providers (CoSPs) have embraced cloud technologies to accelerate growth by faster time to market and greater flexibility, scalability and service availability. For transitioning to cloud-native solutions, service providers adopt cloud-native principles such as containerization, orchestration, micro-services architecture and many others. Kubernetes, as an open-source container orchestration system, provides the foundation for cloud-native architectures and symbolizes a new form of application architecture compared to what virtualization traditionally offered. However, Kubernetes has been challenging to provision, configure and manage from scratch. VMware Tanzu enables users to deploy, monitor and manage cloud-native applications on Kubernetes clusters at scale on public and private clouds through a unified control plane.

In this paper, we leverage VMware Telco Cloud Platform[1,2] with VMware Tanzu to simplify and optimize the installation of a scalable Kubernetes environment to deploy 5G wireline Access Gateway Function (AGF[3]) application on third generation Intel® Xeon® Scalable processors. The 5G AGF network function connects to Customer Premise Equipment (CPE), establishes and manages subscriber sessions, thus enabling customers to access various 5G and other broadband services such as the Internet, Voice over Internet Protocol (VoIP) or Internet Protocol television (IPTV). The reference 5G AGF data plane employs VPP (Vector Packet Processing)/DPDK (Data Plane Development Kit) technologies, NUMA binding and hugepages for higher performance processing and packet acceleration technologies such as SR-IOV for telco-grade I/O performance.

Table of Contents

Authors

Aibhne Breathnach, Platform Applications Engineer, Intel
Jasvinder Singh, Software Architect, Intel
Yusuke Kanamori, Director, Product Marketing, Service Provider & Edge BU, VMware
Sorin Amihalachioaie, Technical Product Manager, Service Provider & Edge BU, VMware
Paul Mannion, Segment Director, Wireline & Converged Access, Intel
Ahmed Ibrahim, Director, Global Business Development – Service Providers, Intel

2 Wireline and Wireless Convergence Reference Architecture

Wireline and Wireless convergence (WWC) is a new architectural design described in the Broadband Forum report WT-470. It promises operational and capital expenditure benefits to hybrid (fixed and mobile) operators through the merging of their wireline and wireless networks.

There are three potential user-plane approaches to implementing WWC: extension of the 5G user plane function (UPF), extending the current BNG (Broadband Gateway function) or introducing a new node termed the Access Gateway Function, or AGF (Figure 1).

Given that many operators have already deployed 5G UPF in their 5G rollouts, this paper focuses on the deployment of the AGF which will enable WWC operation through already-deployed UPF function.

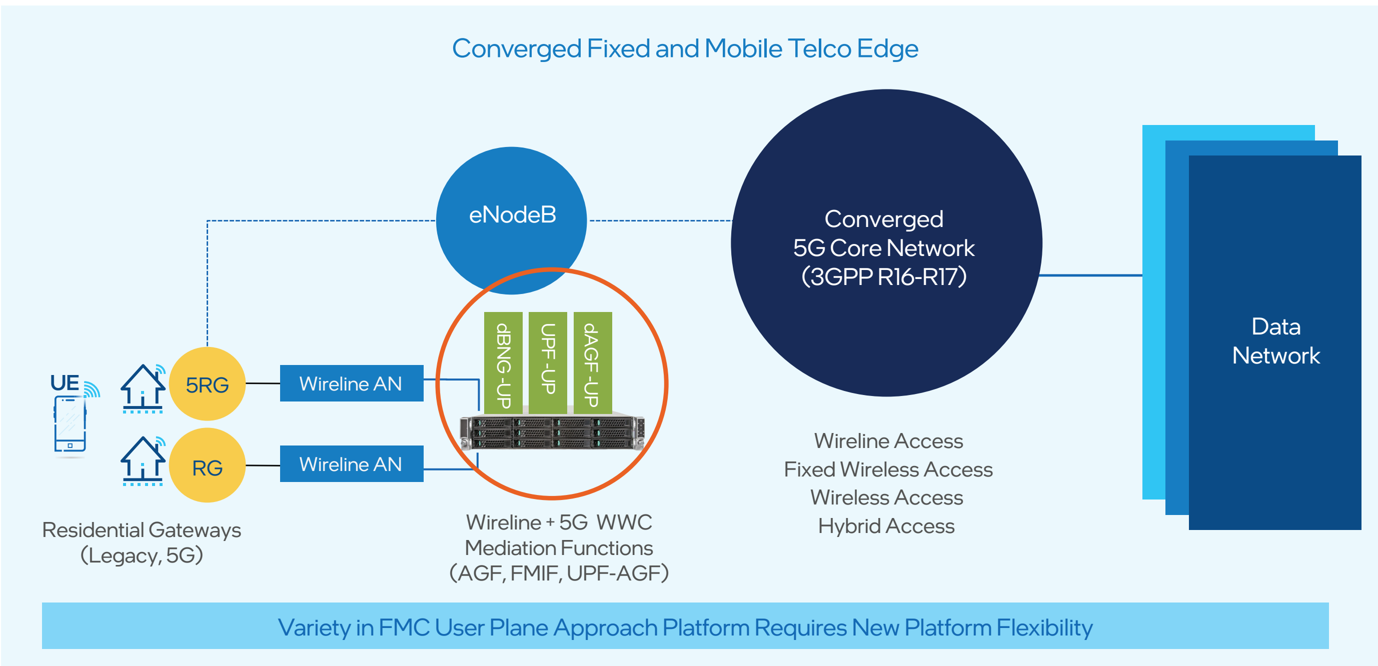

Figure 1. The potential user-plane approaches to implementing WWC in the network.

3 Technology Overview

3.1 Reference vAGF Data Plane

The AGF (Access Gateway Function) routes the traffic between the subscribers connected to the wireline access network and the 5G core network of communication service providers. It then implements functions such as policy rule enforcement as per service level agreements (SLAs), packet classification, header processing, transport encapsulation/decapsulation, hierarchical quality of service traffic management and more. Its functional complexity depends upon the service provider requirements.

AGF functionality in the reference implementation uses separate uplink and downlink packet processing pipelines (Figure 2). Each pipeline applies a set of functions to each packet that enters the pipeline. The uplink packet processing pipeline handles the packets flowing from the subscriber access network to the ISP 5G core network, while the downlink pipeline deals with the packets running from the core network to the access network. The average packet size of upstream traffic is typically smaller than for downstream, and the amount of upstream traffic is normally five to eight times less than downstream traffic. In recent years, the traffic gap between upstream and downstream has reduced significantly because more user-generated, high-bitrate content/media is being uploaded (or shared) into the network through applications like Instagram, Snapchat, TikTok, etc.

The reference AGF data plane is developed using the performance-optimized Vector Packet Processing (VPP) framework[4], which leverages Data Plane Development Kit (DPDK) drivers to take advantage of high speed I/O. The VPP is designed around an extensible and modular packet-processing graph architecture in which each independent graph node performs a specific packet processing function on a vector of packets. The VPP framework allows application developers to plug in new graph nodes without changing core or kernel code to build customized packet processing solutions. The AGF uplink and downlink pipelines are implemented as separate packet processing graphs around a number of nodes where each graph node implements a specific packet processing function.

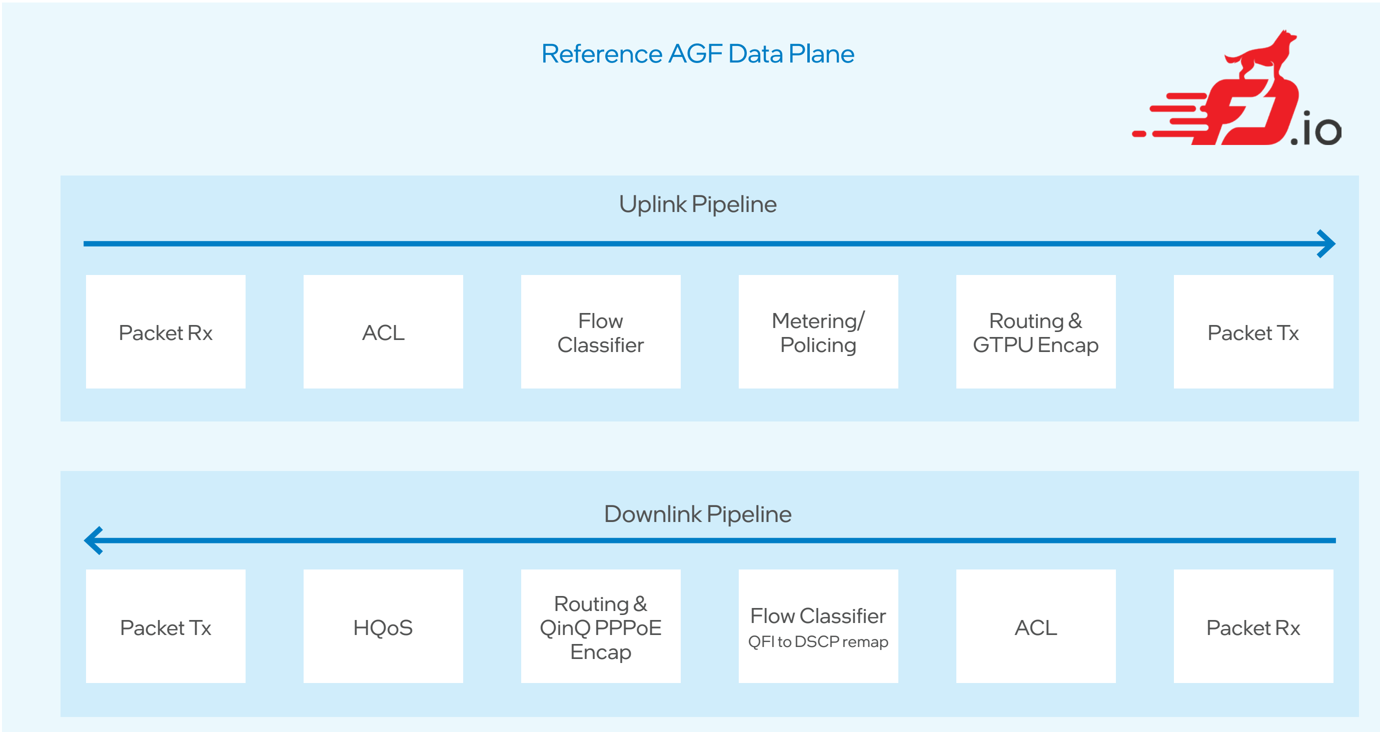

Figure 2. VPP-based reference AGF data plane packet processing pipeline.

3.2 VMware Telco Cloud Platform

VMware Telco Cloud Platform (Figure 3) is powered by field-proven compute and networking coupled with VMware Telco Cloud Automation and VMware Tanzu, a telco-grade Kubernetes distribution. VMware Telco Cloud Platform empowers CoSPs to modernize their wireline and wireless networks so they can rapidly deploy and efficiently operate multi-vendor CNFs and VNFs.

VMware Telco Cloud Platform establishes an open, disaggregated and vendor-agnostic ecosystem to streamline wireline and wireless core service delivery from design to life-cycle management automation while creating a unified, operator-friendly architecture with key capabilities for resource optimization, operational consistency and multi-layer automation.

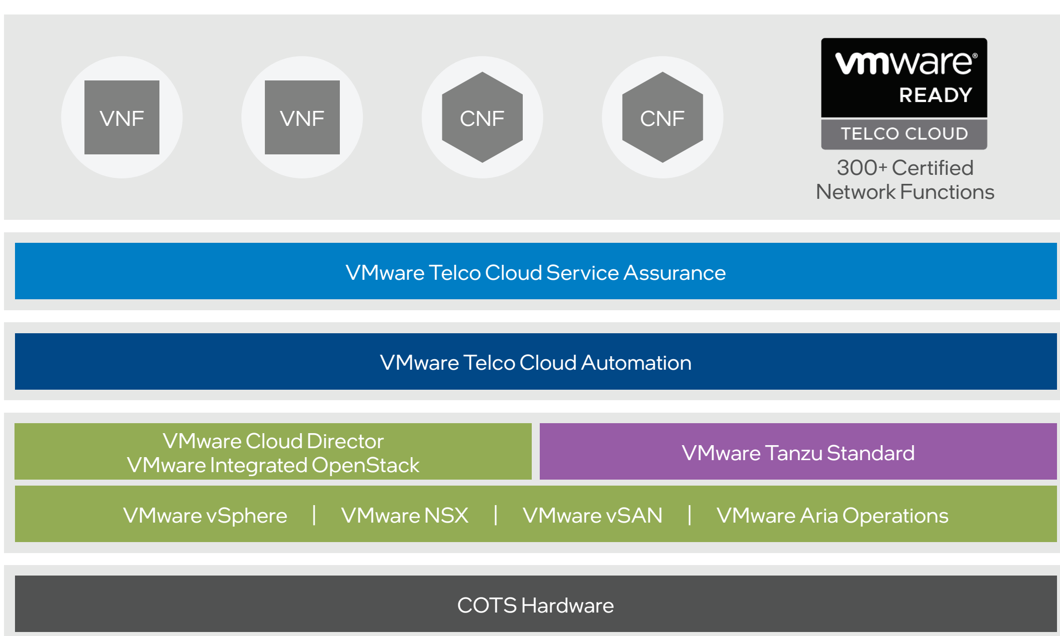

Figure 3. VMware Telco Cloud Platform

VMware Telco Cloud Platform, a high performance CaaS and IaaS infrastructure, enables CSPs to deploy both CNFs and VNFs on consistent horizontal infrastructure. With VMware NSX providing enhanced data plane networking between these network functions, the platform offers high performance and scaling, including the following features:

- VMware NSX-managed Virtual Distributed Switch in Enhanced Data Path mode (N-VDS (E)) that leverages Data Plane Development Kit (DPDK) techniques to provide a fast virtual switching fabric on VMware vSphere.
- Low-latency data plane through CPU pinning, fine-grained non-uniform memory access (NUMA) placement and vertical NUMA alignment.
- Improved performance through multi-tiered routing, bare-metal NSX edge nodes and hugepages with the access efficiency of translation lookaside buffers.
- Topology Manager to optimally allocate CPU memory and device resources on the same NUMA node to support performance-sensitive applications.
- Kubernetes cluster automation to simplify deployments and management of the Kubernetes nodes.
- Support for conventional performance enhancement technologies, such as DPDK and single-root input-output virtualization (SR-IOV) for data plane acceleration.
- vSphere VMXNET3 provides high performance and can enable the support of VMware NSX Enhanced Network Stack (ENS), vMotion and DRS capabilities. vSphere vMotion enables zero-downtime, live migration of workloads from one server to another. And Dynamic Resource Scheduler (DRS) provides highly available resources to your workloads. DRS balances workloads for optimal performance, and scales and manages computing resources without service disruption.

3.3 VMware Tanzu

VMware Tanzu provisions and manages the life cycle of Tanzu Kubernetes clusters.

A Tanzu Kubernetes cluster is an opinionated installation of Kubernetes open-source software that is built and supported by VMware. With VMware Tanzu, administrators provision and use Tanzu Kubernetes clusters in a declarative manner that is familiar to Kubernetes operators and developers.

Tanzu Management and Workload Clusters

Tanzu Kubernetes Management Cluster is a Kubernetes cluster that functions as the primary management and operational center for the VMware Tanzu instance. In this management cluster, the cluster API runs to create Tanzu Kubernetes clusters while CoSPs configure the shared and in-cluster services that the clusters use.

Tanzu Kubernetes Workload Cluster is a Kubernetes cluster that is deployed from the Tanzu Kubernetes Management Cluster. Tanzu Kubernetes clusters can run different versions of Kubernetes depending on the CNF workload requirements. Tanzu Kubernetes clusters support multiple types of CNIs for pod-to-pod networking, with Antrea as the default CNI and the vSphere CSI provider for storage by default. When deployed through VMware Telco Cloud Automation, VMware NodeConfig Operator is bundled into every workload cluster to handle the node Operating System (OS) configuration, performance tuning and OS upgrades required for various types of CNF workloads.

Tanzu Kubernetes Cluster Data Plane

Tanzu Kubernetes Cluster Data Plane consists of worker nodes that run as virtual machines (VMs). A worker node consists of the container runtime, kube-proxy and kubelet daemon to function as a member of the Kubernetes cluster.

Depending on the type of network functions, the worker nodes might require specialized hardware and software to support advanced network capabilities. DPDK, SR-IOV and multiple network interfaces are often recommended for the data plane of network functions to provide the telco-grade networking performance.

In alignment with the Cloud iNfrastructure Telco Taskforce (CNTT), the following hardware and compute profile modes are depicted in this solution architecture (Table 1):

Table 1. Hardware and compute profile modes

	Control Profile	Network Intensive Profile
vCPU Over-Subscription	Yes	No
CPU Pinning	No	Yes
Huge Pages	No	Yes
IOMMU	No	Yes
NUMA	No	Yes
SR-IOV	No	Yes
DPDK	No	Yes

3.4 3rd Gen Intel® Xeon® Scalable Processors

The 3rd Gen Intel® Xeon® Scalable processor provides a balanced, scalable architecture that is built to be tailored to diverse implementations with a wide range of core counts, frequencies and power levels. Architectural advances over its predecessor dramatically increase per-core performance, memory subsystem performance and I/O bandwidth to accelerate workloads and increase the concurrent workload capacity per server. Platform enhancements that accelerate diverse workloads from the edge to the data center include the following:

- Increased core count and cache, available in a flexible range of 8-40 powerful next-generation cores, L1 caches of 12-60 MB and total design power of 140-270 watts.
- Expanded Intel® Advanced Vector Extensions 512 (Intel® AVX-512), accelerating bit processing by moving and reordering blocks of data faster within the signal processing pipeline.
- Updated I/O subsystem, including support for PCI Express Gen4, which provides double the bandwidth of PCI Express Gen3 for faster data movement.
- Enhanced memory subsystem, with up to 1.60x higher memory bandwidth and up to 2.66x higher memory capacity[5] compared to the prior-generation platform.

Engineered for modern network forwarding-plane workloads, the 3rd Gen Intel Xeon Scalable processor targets low-latency, high-throughput deterministic performance and high performance per watt. For today's software-defined environments, the platform is optimized for hybrid/multi-cloud deployments. It is built on open standards and APIs to enable infrastructures designed to accommodate current and future business needs.

The processor enhances cryptographic acceleration with new instructions that increase encryption throughput on the CPU without the need for dedicated hardware accelerators. As encryption requirements increase with the advent of 5G networking, this capability is becoming more critical to CoSPs. In addition, Intel® Software Guard Extensions (Intel® SGX) creates isolated memory enclaves that can help securely store encryption keys at the network edge.

3.5 Intel® Ethernet 800 Series Network Adapters

The Intel® Ethernet 800 Series Network Adapters provide network I/O that complements the 3rd Gen Intel Xeon Scalable processor's compute and memory advances. The adapters use PCI Express Gen4 for improved bandwidth to the system board with network throughput up to 100 Gbps per adapter port. It delivers standards-based networking performance across NFV and CFV workloads through a combination of sophisticated packet processing, intelligent offloads and accelerators, and high-quality open-source drivers for data-plane processing. In addition to optimizing throughput, the adapters are designed to enable broad interoperability and agility. Dynamic Device Personalization (DDP) allows multiple personalization profiles to specify optimizations and packet handling parameters for individual traffic types, increasing throughput and enabling sophisticated traffic prioritization. The DDP programmable packet-processing pipeline provided by the Intel Ethernet 800 Series Network Adapters supports on-demand reconfiguration of network controllers at runtime, enabling workload-specific optimizations. DDP is enhanced in the Intel Ethernet 800 Series Network Adapters with greater programmability than its predecessor, as well as workload-specific protocols for added flexibility. The enhanced Data Plane Development Kit (DPDK) is an open-source set of libraries and drivers supported by the Intel Ethernet 800 Series Network Adapters that accelerates packet processing in the data path. It also facilitates building packet forwarders designed to operate on general-purpose, standards-based servers. DPDK technology is incorporated as a bundled feature with RHEL, enabling Intel Ethernet 800 Series Network Adapters to be controlled entirely in user space. That approach accelerates operations by allowing network packets to bypass the kernel network stack entirely. With 3rd Generation Intel Xeon Scalable processors, CoSPs can enhance DPDK L3 forwarding performance by up to 76% compared to the prior generation.[5]

4 Deployment

In this section, the vAGF deployment on VMware Telco Cloud Platform will be detailed from the ground up, along with the tweaks necessary to achieve performance comparable to a bare metal deployment.

4.1 Cluster Overview

The cluster setup (Figure 4) is kept relatively simple for this activity with two servers. One management server runs the VMware vCenter Server, the Tanzu bootstrap and management VMs and a local Harbor container registry. Another worker server runs the Tanzu Worker Cluster VMs that runs the vAGF pods.

On the Management server, the VMware vCenter Server VM is set up first to allow centralized management of the server and VM cluster as part of the VMware Tanzu deployment. The Tanzu Bootstrap VM is deployed next. This VM is used as a bootstrap and bastion node for the Tanzu Management cluster, providing a node to run the Tanzu Management cluster CLI or GUI installer as well as a node setup (post Tanzu Management cluster installation) with the necessary credentials to access the Tanzu Management clusters and Worker Kubernetes clusters. A Harbor container registry is also deployed to host container images locally. This registry will host the vAGF container image for the Tanzu Worker nodes to pull and run. There are many ways to deploy a Harbor registry within the VMware Tanzu environment, however for this use case, a simple VM was deployed with Harbor installed, as it would only be hosting a single container image, the vAGF.

On the system under test (SUT), once the required components have been deployed on the Management server, a Tanzu Worker cluster is deployed. This cluster, and specifically the worker node of this cluster, is set up with the required CPU, memory, and NIC resources necessary to run the vAGF. More details on this in the following section.

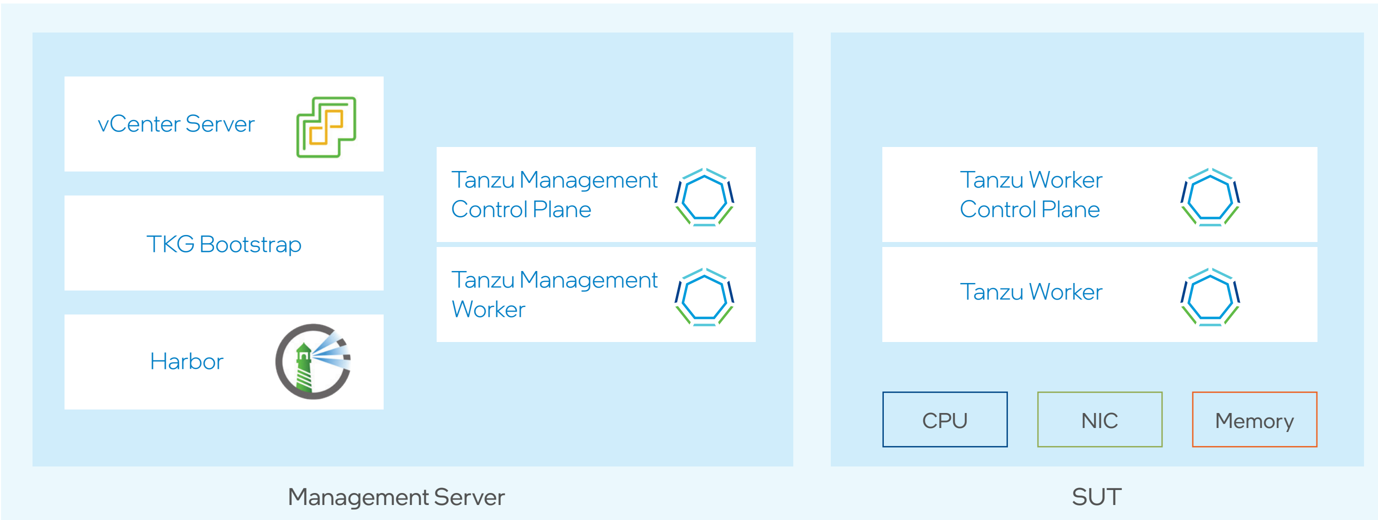

Figure 4. VMware Tanzu Cluster Overview

4.2 vAGF Deployment on VMware Tanzu

For the vAGF deployment on the VMware Telco Cloud Platform, firstly let's examine what makes up a single vAGF instance and what the requirements are for a vAGF instance to run. Then a multi-instance deployment is detailed.

The Intel vBNG-vAGF Reference Architecture Package provides the vAGF as a Docker container, allowing the vAGF to be deployed on any compatible Kubernetes node or platform — virtualized or bare metal. When deploying the vAGF on VMware Telco Cloud Platform, no modification to the application is needed. The vAGF image (Figure 5) was uploaded to the local Harbor container registry that's deployed with VMware Tanzu.

The vAGF is a VPP-based application, meaning there are some platform requirements (in this case from the VM) for each vAGF instance to run smoothly. The required platform resources are:

- Isolated CPUs allow UL and DL pipelines to run concurrently without interference from the OS or other system processes. This provides an optimized environment for UL and DL pipeline performance.
- Dedicated hugepages are used for the large memory pool allocated for incoming packets. This bigger memory page size reduces the overhead and the high TLB miss rates that would occur with the standard 4K memory page size.
- Dedicated network interfaces for uplink and downlink will streamline both pipelines, reducing congestion on the NIC.

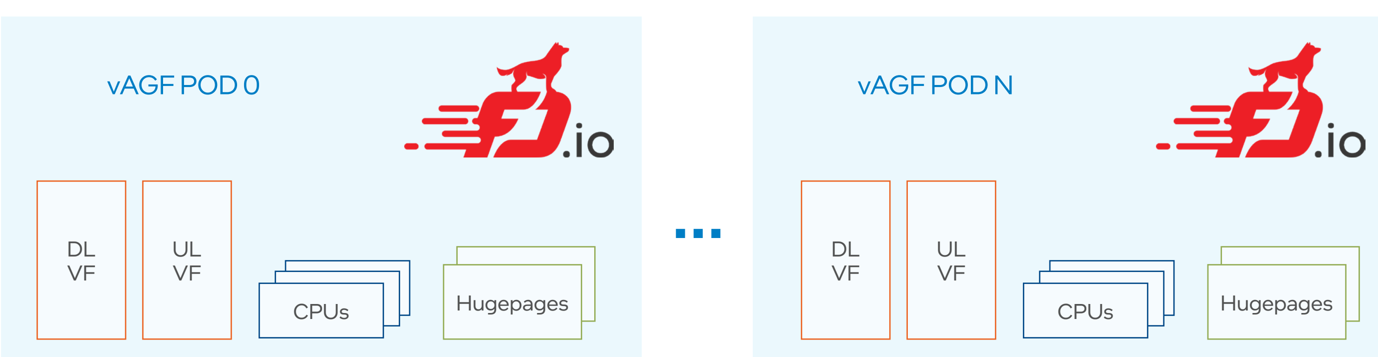

Figure 5. vAGF Pod Overview

These resources can be allocated to the vAGF pods by the native Kubernetes CPU Manager (isolated CPUs and memory) and the SR-IOV Device Plugin (SR-IOV VFs).

When deploying a number of these vAGF pods in a VM, the platform requirements mentioned above must be in place to make sure each vAGF instance has the required resources to run. Theoretically, any number of vAGF instances can be run in a single VM, provided the correct number of isolated CPUs, hugepages and SR-IOV VFs are available. Each vAGF instance requires four vCPUs, 1 GB of 2-MB hugepages, and two network interfaces.

However, in this case, the number of vAGF instances per VM is limited by a maximum of 10 network interfaces allowed per VM. With each vAGF instance needing two VFs (one for uplink and one for downlink), four vAGF instances is the maximum that can be deployed per VM, allowing two remaining network interfaces for management traffic.

In this case, since four vAGF instances is the maximum that can be deployed per VM, 18 CPUs (two for management, and 16 for vAGF) and 16GB RAM (8GB for system, 8GB for hugepages) and 10 VFs (two management, and eight for vAGF data plane) are allocated to the VM (Figure 6). CPUs and VFs for management not shown in diagram for clarity.

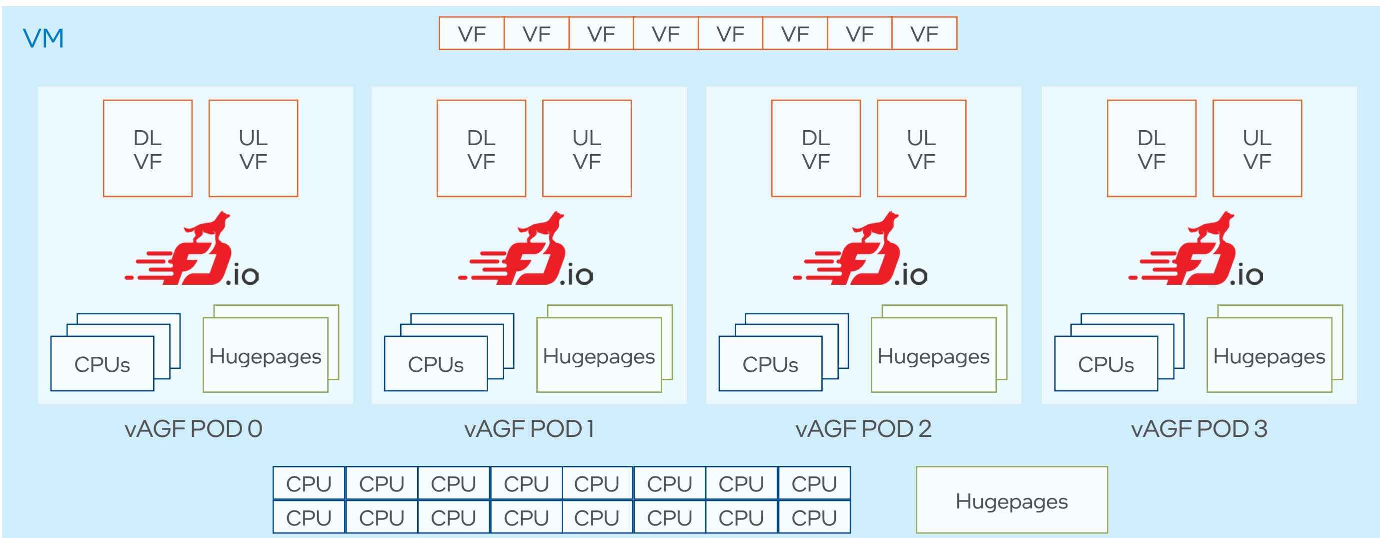

Figure 6. vAGF Worker VM Allocation

Also included in the Intel vBNG-vAGF Reference Architecture Package is a set of Helm charts to enable simple and easy deployment to any Kubernetes cluster. These Helm charts contain the necessary templates to deploy the vAGF pod with the required platform resources mentioned above. With the vAGF container image already uploaded to the local Harbor registry, the vAGF can be deployed to the Tanzu cluster.

4.3 Tanzu Worker VM Resource Allocation

This section will detail how platform resources allocated to Tanzu Worker VMs are configured to ensure vAGF performance is maintained. The basic principles applied in this section follow the same principles considered when deploying the vAGF to a bare metal server. However in this case, native VMware tools and techniques are used where applicable to take advantage of the ease-of-use and simplification they provide.

Learn More

VMware Telco Cloud Platform tuning guide https://docs.vmware.com/en/VMware-Telco-Cloud-Platform---5G-Edition/2.0/telco-cloud-platform-5g-edition-data-plane-performance-tuning-guide.pdf

4.3.1 CPU Affinity and Isolation

When adding vCPUs to a VM, how these CPUs are allocated is critical to achieving the best possible performance from the vAGF pods deployed in that VM.

By default, when vCPUs are added to a VM, the physical cores these vCPUs run on are managed by the VMware. This means that there is a level of indeterminism for workload throughput and latency, as workloads running on the same pCPU can impact each other, as well as which pCPU the workload is placed on, and this can significantly impact the performance of data plane workloads including the throughput and latency of the vAGF, where determinism is key.

To counteract this indeterminism, LatencySensitivity=High is set for each worker VM. This setting enables complete core isolation for each pCPU by leaving sibling hyperthreads idle, as well as pinning CPUs by applying exclusive pCPU affinity with vCPUs. This ensures that each pCPU mapped to worker node vCPUs are isolated from interference from system processes and other workloads.

Within the VM, making use of the isolcpus kernel parameter will stop the guest OS from running system processes on those vCPUs, allowing those vCPUs to be allocated to the vAGF and run uninterrupted by guest OS processes.

Enabling these settings for both the worker VM and the guest OS ensures that the vCPUs allocated to this VM retain their affinity to their underlying pCPUs and that they are not interrupted by system processes or other workloads. This provides an optimized CPU deployment for the vAGF to run on.

4.3.2 Memory and Hugepages

Memory allocated to the worker VM is split in the guest OS into system memory and hugepages. By using hugepages, the system memory page size is increased from the standard 4KB size to either 2MB or 1GB. The bigger memory page size reduces the page table lookup overhead and the high TLB miss rates that would occur with the standard 4K memory page size. This is usually not an issue with system applications and processes, but required with the vAGF and other DPDK/VPP based applications. With the memory footprint and usage of the vAGF (and typical DPDK and VPP based applications), having bigger memory pages of 2MB or 1GB in size reduces the amount of memory pages needed to keep all incoming packets in memory. This reduces memory management overhead and frees up more CPU time for packet processing.

The optimal page size to use for the vAGF is 2MB, as this provides a good balance of reducing memory management overhead of smaller memory pages while not overprovisioning memory allocated with 1GB pages. Each vAGF instance uses 1 GB of 2MB hugepages.

4.3.3 SR-IOV Data Plane

When adding SR-IOV VFs to a VM, many of the optimizations are taken care of automatically.

One important optimization is to make sure that the VFs, memory pools and CPUs assigned to each worker VM are both on the same NUMA node, or NUMA aligned. In a dual-socket server, this typically means that there is one NUMA node per CPU socket. This ensures that packets ingressed through the NIC are written to memory and processed by the CPU. This process takes place over CPUs, memory pools and PCIe links that are adjacent, and avoids data crossing the QPI link that connects both CPU sockets, which can significantly impact performance.

This NUMA alignment is taken care of automatically in the VMware ecosystem, so that when assigning platform resources to a VM, the user can rest assured that the CPUs, memory pools and network interfaces are NUMA aligned.

Using SR-IOV for the vAGF data plane provides a very simple but efficient network data plane to allow the network packets incoming to the NIC to reach the vAGF running in the VM directly. However, there are other solutions, such as VMware NSX, that can provide more complex and performant networks between NIC egress and VM ingress. These solutions will be explored in a follow-on paper.

5 Benchmarking and Results

The vAGF is benchmarked in this case to understand how this reference application performs running within a VM deployed on a VMware Tanzu cluster. With the optimizations mentioned in the previous section in place, the objective is to maintain as comparable performance as possible to a bare metal deployment with this virtualized deployment. Understandably, there will be some level of performance loss due to the virtualization overhead, but the optimizations put in place will minimize this as much as possible.

The vAGF is tested with an asymmetric traffic profile. This means that the packet sizes and traffic rates for uplink and downlink are not the same. This is to simulate a real-world traffic profile, where typically uplink traffic has smaller packet sizes and lower traffic rates (e.g. client request to web service) and downlink has bigger packet sizes and higher traffic rates (e.g. media download or live stream). Each vAGF instance is sent 25 Gbps of traffic in an 8:1 ratio (22.25 Gbps downlink, 2.75 Gbps uplink), with a 504-byte packet size for downlink and a 128-byte packet size for uplink. Each vAGF instance is tested with 4096 simulated subscribers, and in each test scenario, the max throughput of the vAGF instances is tested.

To assess how well the optimizations on the cluster made the performance comparable to a bare metal vAGF deployment, the baseline results must first be assessed. In Figure 7, the vAGF benchmark results on bare metal Linux are detailed. In these results, the linear scaling of throughput in packets per second can be seen in both the uplink and downlink with the increase in the number of vAGF instances, ranging from up to 7.6 MPPs for one vAGF instance, to up to 30.5 MPPs for four vAGF instances.

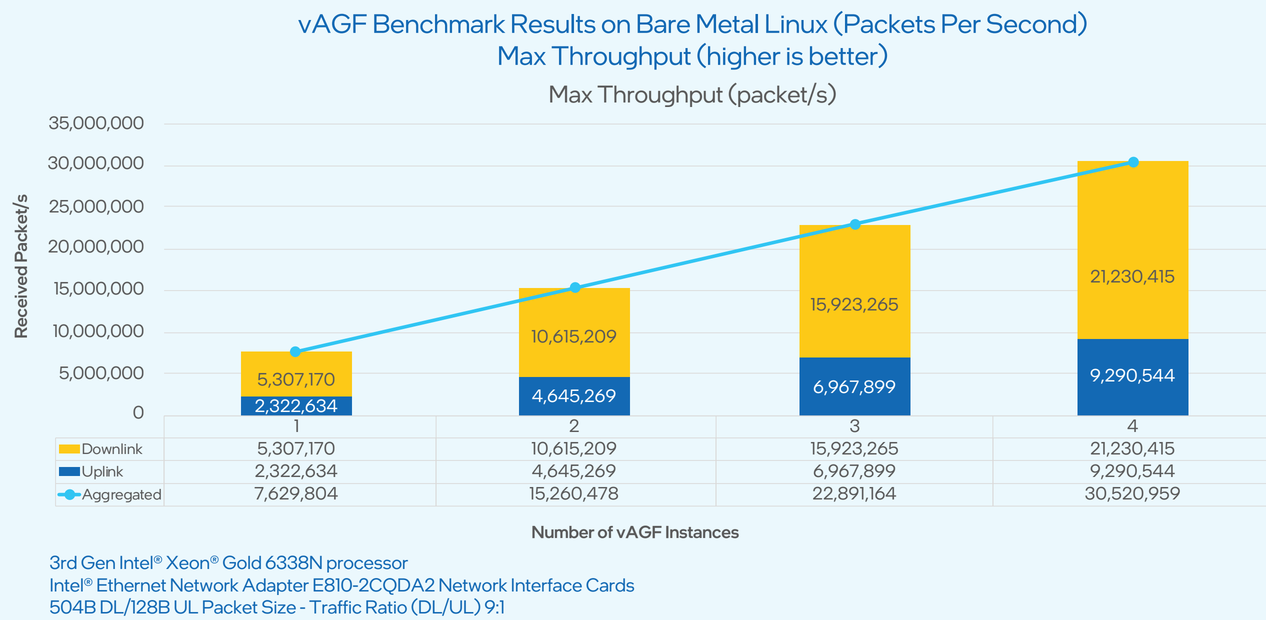

Figure 7. vAGF Max Throughput Benchmark Results on Bare Metal Linux

Comparing these baseline results to a vAGF deployment on VMware Telco Cloud Platform with VMware Tanzu, in Figure 8, a similar linear scaling in throughput for the same amount of vAGF instances in a single VM can be seen. And a small tail-off in this linearity can be seen for the maximum throughput for four vAGF instances. Even with this tail-off in performance, overall, the performance of a single VM is comparable to the bare metal baseline, given that there is the virtualization overhead to consider when running high performance workloads in a virtualized environment.

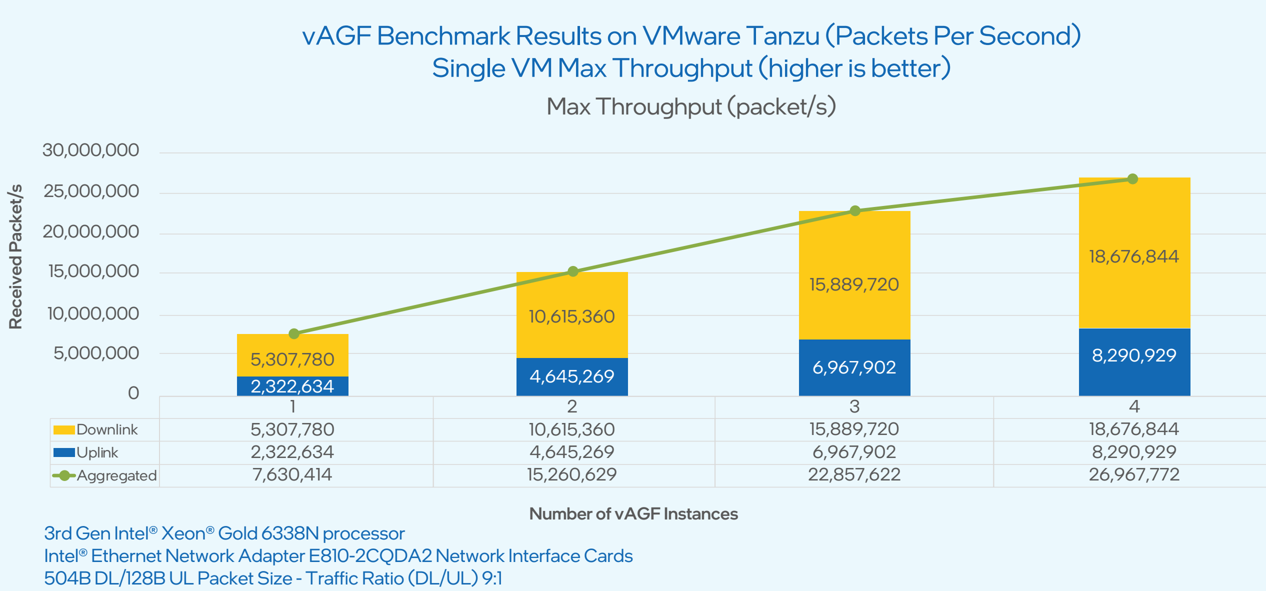

Figure 8. vAGF Max Throughput Benchmark Results on VMware Telco Cloud Platform with VMware Tanzu

6 Summary

Wireline and Wireless convergence is a new network transformational concept that promises to deliver huge benefits to communication service providers who today manage hybrid (fixed and mobile) networks. However, this technology requires the flexibility that only a true cloud-native platform with VMware ESXi hypervisor can deliver.

Intel and VMware have been working together to enable a high-performance WWC cloud-native platform adopting 3rd Gen Intel Xeon SP processors and Intel 800 Series Network Adapters on a VMware Telco Cloud Platform with VMware Tanzu.

The shift to a cloud-native ecosystem is a step-change for the telco industry. A growing number of CoSPs are changing application development, deployment and life-cycle management methods to build cloud-native solutions in the emerging 5G era. While Kubernetes is the de facto standard for deploying and operating containerized applications, it comes with its own set of challenges that organizations face with migrating to a Kubernetes-based deployment. As such, organizations now look to more modern hosting solutions to simplify and accelerate the adoption of Kubernetes.

VMware Telco Cloud Platform with VMware Tanzu provides a full stack of capabilities to modernize the infrastructure setup and deployment of containerized telco applications. Alongside the advanced CaaS orchestration and automation capabilities of VMware Telco Cloud Platform, this collaborative work has demonstrated the 5G wireline AGF data-plane deployment, and that when deployed on VMware Telco Cloud Platform, performance on Intel Xeon CPU in a cloud environment is comparable to bare metal deployment.

The platform has been optimized for high-speed packet processing, and the results show that with the correct platform configuration in place, the virtual performance can achieve up to 88% of current bare-metal performance.

[1] VMware, data sheet, "VMware Telco Cloud Platform: Deploy and Operate 5G Functions and Services on Consistent Infrastructure with Cloud-Smart Automation," https://telco.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vmw-telco-cloud-platform-tcp-datasheet.pdf

[2] VMware, "Telco Cloud Platform - 5G Edition: Reference Architecture Guide 2.0," https://docs.vmware.com/en/VMware-Telco-Cloud-Platform---5G-Edition/2.0/telco-cloud-platform-5G-edition-reference-architecture-guide-20.pdf

[3] Broadband forum, TR-456, "AGF Functional Requirements," https://www.broadband-forum.org/technical/download/TR-456.pdf

[4] FD.io, "What is the Vector Packet Processor (VPP)," https://s3-docs.fd.io/vpp/23.02/.

[5] 3rd Gen Intel Xeon Scalable Processors Brief: https://www.intel.com/content/dam/www/public/us/en/documents/a1171486-icelake-productbrief-updates-r1v2.pdf

0423/BB/MESH/353915-001US