WHITE PAPER

DEMYSTIFYING NFVI

Architectural Overview for Performance Optimized 5G NFV Infrastructures

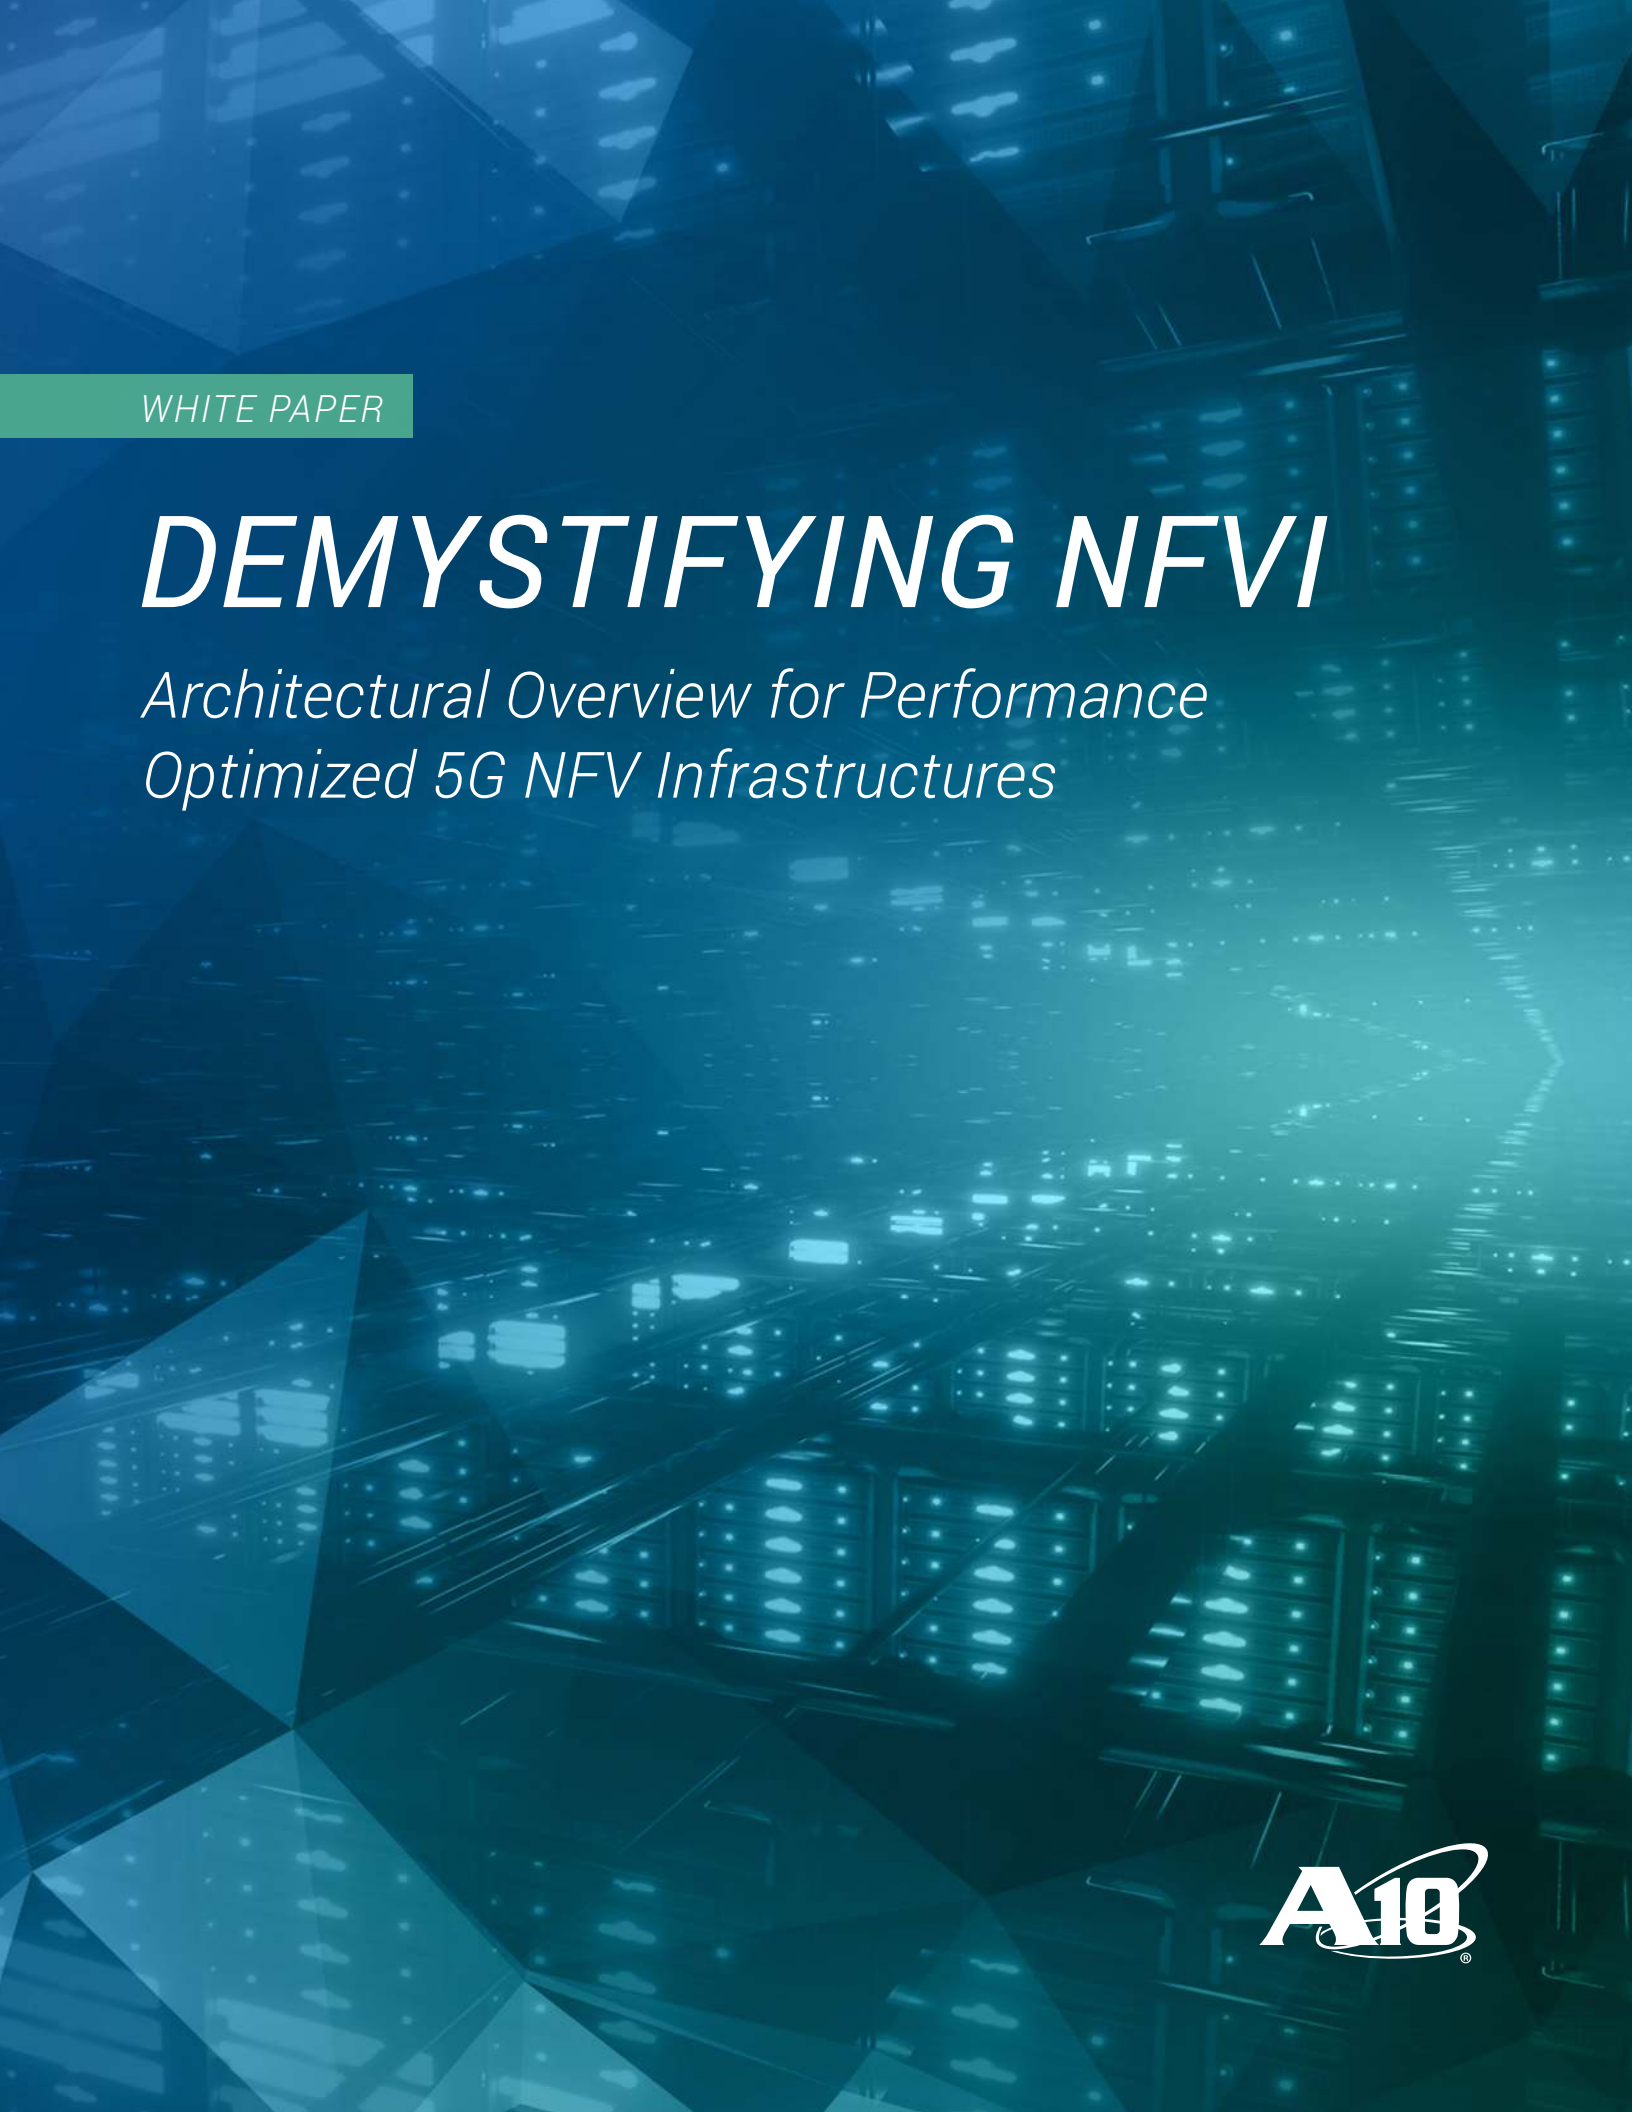

TABLE OF CONTENTS

DISCLAIMER

EXECUTIVE SUMMARY

5G is coming and it requires leading-edge network and performance technologies. Network function virtualization (NFV) is one of the key technology building blocks. NFV infrastructures are a new and complex black box.

This white paper opens this black box, demystifying NFV infrastructures for 5G technology decision makers, system architects and network managers, providing an overview of NFV technology and many of the key architectural components.

Reading this white paper will provide key performance and architectural guidelines to help design and deploy an NFV infrastructure and meet the performance, scale and reliability requirements for 5G.

5G REQUIRES A NEW NETWORK ARCHITECTURE

5G technology is a quantum leap forward with the promise of transforming industries, world economies and profoundly affecting everyone alive today. The number of connected devices is exploding. This transformation will require network service providers to deliver orders of magnitude higher data throughputs with much lower latency, higher uptime and massive increase in scale of connected devices.

To meet the key performance and feature goals for 5G, mobile network carriers must design and deploy a new network architecture. 5G has specific key performance indicators (KPI) and requires a hardware and software environment with specific performance technologies and configurations.

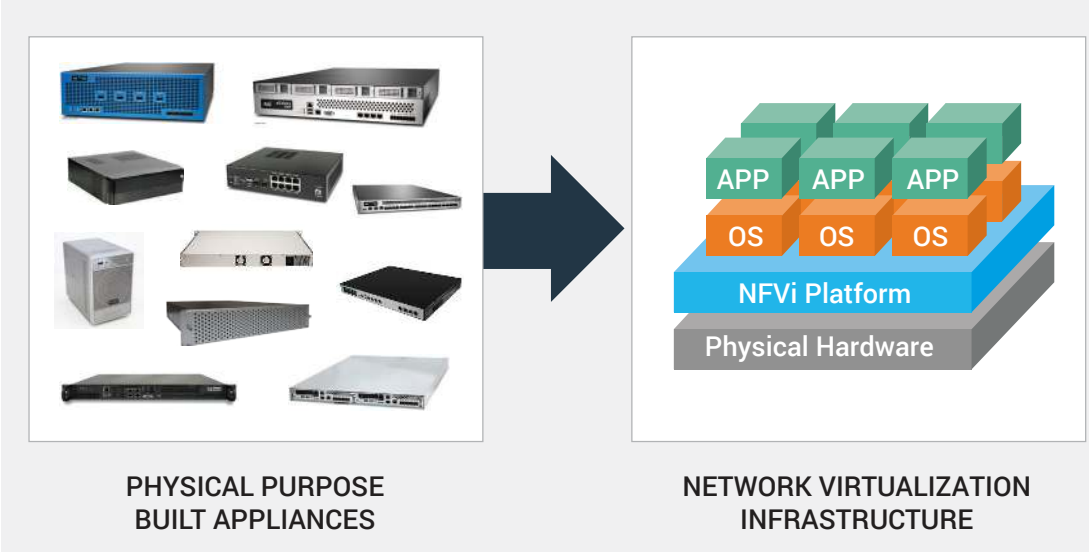

Figure 1: Transformation of specialized hardware to virtualization platforms

Network infrastructures are moving physical machines to virtualization platforms, vastly increasing operational agility and flexibility.

Until recently, virtualization far underperformed physical machines and suboptimal throughput and higher packet latency for high-performance environments. This is changing rapidly with substantial advancements in hardware and virtualization technologies. Using proper design and configurations described in later sections, virtual instances come close to the performances delivered by physical appliances.

This document has two goals. The first sections provide a conceptual overview of network function virtualization. The remaining sections provide a high-level view of the technology components of a server platform, focusing on performance acceleration and optimization technologies.

WHAT IS NFVi?

Network functions virtualization (NFV) is a network architecture based on IT virtualization and cloud computing to virtualize entire classes of network functions. Virtual network functions (VNFs) are independent building blocks or can be chained together to create network communication services.

NFV infrastructure (NFVi) is the architectural specification defined by several standards groups, for hardware and software components for compute, storage and networking systems, providing the platform for NFV and hosting VNF systems. NFVi is generally distributed services deployed over multiple locations.

NFVi is changing the way service providers design and manage networks leveraging industry standard, off-the-shelf computing hardware and virtualization technologies. This new architecture is high-performing and much more flexible with greatly reduced capital and operational costs.

WHY NOW?

The design goals for 5G require vast increases in network throughput and connected devices while reducing latency to 1 millisecond or less.

Mobile network traffic is projected to increase substantially in the next few years: 550 million 5G subscriptions by 2022, 30 billion connected devices by 2022. Network Service Providers must design and deploy a high-performance, scalable and manageable architecture.

Proprietary hardware systems are much less agile and scalable than virtualized services. NFV virtualization is flexible, scale easily and run on standard hardware systems. Using the latest performance technologies, virtualization does not mean compromising performance to a great degree and meets the carrier-grade requirements of 5G and mobile services providers.

NFV IS ESSENTIAL FOR 5G

Beyond increasing network capacity and flexibility, 5G includes new capabilities which are only possible with the NFV virtualization infrastructure.

DISTRIBUTED CLOUD

Distributed cloud technology allows multiple data centers to appear as a single, virtual data center. With 5G, NFV is more than merely moving functions to commodity hardware. The 5G NFV distributed cloud environments will be scalable, resilient, and fault-tolerant. Combining distributed cloud technology with virtualized network functions (VNFs) will allow the VNFs to be deployed based on performance and other requirements, making it easier for operators to optimize, manage, and maintain networks.

5G VIRTUALIZATION AT THE EDGE

Virtualization will be an essential component at the edge of 5G networks. Service provider networks have several network edge segments:

- **vRAN** moves processing from Remote Radio Unit (RRU) dedicated hardware to virtualized network functions (VNFs)
- **MEC** – Multi-Access Edge Computing places computing resources to small data centers near radio base stations, hosting service provider and cloud resources as close as possible to the edge of the network.
- **Partner Networks** – using VNF and SDN, network slices can be extended to local ISP partner networks, providing a broader set of billable services.
- **vCPE** – virtualized Customer Premises Equipment located at customer locations will host VNF network functions, extending 5G services to the very edge of the network.

BENEFITS FROM NFV

Other than features described above, NFV enables the following:

- Multi-Cloud distributed infrastructures
- Advanced operations: DevOps automation, Ci/CD (continuous integration and continuous development)
- Reduced costs in network equipment using software on standard servers
- Efficiencies in space, power, and cooling
- Faster time to deployment
- Flexibility – elastic scale up and scale down of capacity and new services
- High reliability – up time of 99.999%
- Ability to integrate with legacy network architectures and link to existing operational and billing systems.

NFV CHALLENGES

NFV is proving complex and difficult for many operators to deploy at scale. The breadth of the architecture and the number of distinct components make it challenging to design, build and support.

NFV must be also be integrated and co-exist with current network infrastructures and merged with existing network management and system operations.

As well, the lack of mature standards and specifications for NFV implementations continue to hinder deployments. It has taken years to move NFV deployments from proof-of-concept labs, to onsite trial tests and on to full-scale deployments in production networks.

NFV ARCHITECTURE EXPLAINED

Network Function Virtualization combines the hyper-scale and virtualization concepts of cloud computing, with virtualized network functions (VNFs) deployed standalone or as multiple network functions (NFs), chained together as a single integrated service. Add on top of that an extensive set of functions like management, orchestration, operations support services, and that is Network Function Virtualization infrastructure (NFVi).

A carrier-grade NFV infrastructure can be designed and deployed using best practices and providing high-availability, fault-tolerance and full DevOps orchestration and management.

The above diagram is a logical view of an NVF architecture. VNF services are hosted on a hypervisor virtualization platform running on common-off-the shelf (COTS) servers. Physical Network Functions (PNFs) are network functions running on dedicated bare-metal server hardware or 3rd party purpose-built and optimized appliances.

MANO

The Management and Orchestration (MANO) architecture is in the section on the right. MANO is a framework for the management and orchestration of all network resources.

Network connectivity is shown where the Business Support System (BSS) requests new services on the Orchestration Support System (OSS) which manages the underlying infrastructure management systems.

5G specifications for MANO add automation, analytics and additional components like SDN controller orchestration, NFV management and network slicing across multi-cloud environments.

PHYSICAL AND VIRTUAL FORM FACTORS

Physical Network Functions (PNF) are network functions hosted on a physical server instead of a virtual machine. PNF systems can interact with VNF and other PNF systems to provide communication services. PNF systems can be chained using SDN traffic steering policies.

Virtual Network Functions are virtual appliances configured with software applications to provide specific network functions (NF). Examples of NFs include routers, load balancers, firewalls and 4/5G core packet processing. VNF machines are virtual machines or containerized micro-services.

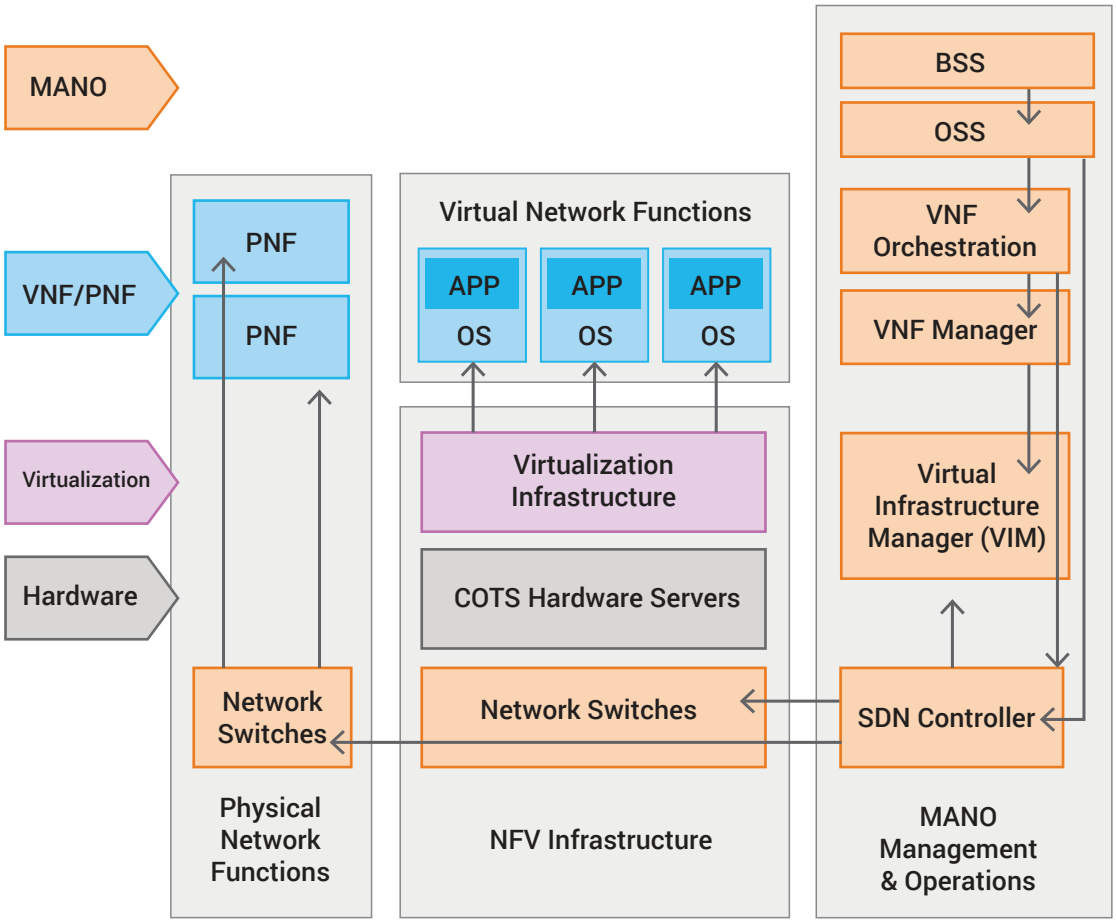

Figure 2: NFV infrastructure functional block diagram

VIRTUALIZATION INFRASTRUCTURE

The virtualization infrastructure is a platform providing virtual resources and hosting virtual machines, networks and storage. Hypervisors and docker engines used in NFVi environments are provide similar functionality as used by IT data centers. Virtualization products commonly include Container/Kubernetes, VMware and KVM used by OpenStack.

COTS HARDWARE SERVERS

Cloud computing evolved to converged architectures, moving away from custom hardware appliances to software hosted in virtual environments. The physical hardware required to host virtual network functions platforms only require standard common-off-the-shelf servers and network switches.

Service provider infrastructures are moving to a similar architecture using NFV virtual technologies.

NFVi PLATFORM OVERVIEW

NFVi platforms are hardware and virtualization systems which run and manage virtualized network functions (VNFs). This includes hardware servers, physical networks and NFV virtualization software systems.

High-throughput network traffic requires very high-performance platforms. Virtualization traditionally added a severe performance penalty. To achieve the performance for 5G service providers, A10 Networks is leveraging the latest and greatest commercially available technologies.

We are presenting field tested NFVi platform optimizations, leveraging state-of-the-art technologies from leading vendors including A10 Networks, which is ahead of the curve.

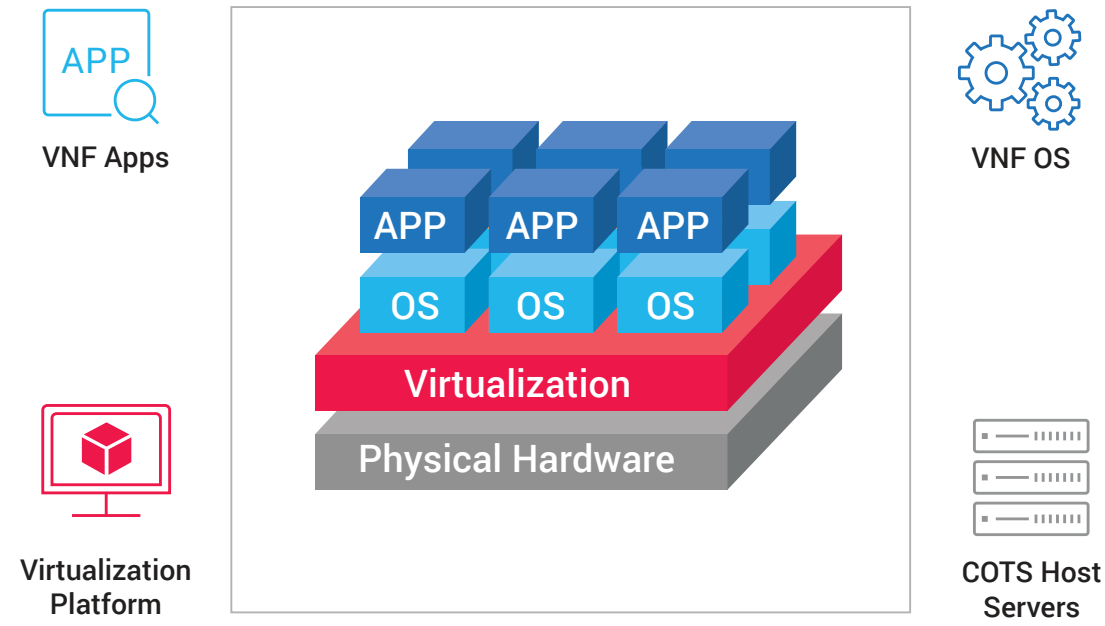

Figure 3: VNF server platform functional layers

The diagram above is a logical architecture of NFVi server platforms. The platform has been segmented into four functions, VNF App, VNF OS, NFVi virtualization platform, standard COTS hardware servers.

The remaining sections of this document describes technologies and configurations used by A10 to optimize network throughput and latency performance. These configurations have been tested, certified and running in several large network carrier environments.

OPTIMIZED VNF SYSTEMS

Optimizing VNF performance requires many advanced tuning techniques for both the VNF Operating System and the VNF applications. This section describes the A10 VNF performance technologies and optimizations.

A10 Networks' Advanced Core Operating System (ACOS) is built from the ground up using supercomputing principles to create a platform for high-throughput low-latency network intensive applications.

ACOS is based on proprietary performance techniques including intelligent application packet processing, parallel network pipelining, advanced shared memory, flexible traffic accelerator processing and more.

The ACOS network packet processing algorithms are application-aware and can accelerate application performance in ways not possible with traditional networking systems.

The advanced network operating environment provides a perfect environment for VNF network functions operating in a high volume 5G network architecture.

ACOS PERFORMANCE TECHNOLOGIES

- Hardware Queuing
- Dynamic Polling
- In-Memory Processing
- Optimized TCP/IP Stack
- DPDK Support
- Kernel bypass poll-mode drivers
- Shared NIC queues
- Parallel network flows synchronized across system
- Shared memory processing across multi-CPU systems

HIGH-PERFORMANCE ARCHITECTURE

The A10 Networks VNF is performance optimized for 5G network infrastructures

- >100 Gbps throughput in high-performance infrastructures
- VM and Containerized VNFs
- Available and optimized for major virtualization, containerization and cloud platforms
- Compatible, certified and optimized for COTS hardware servers

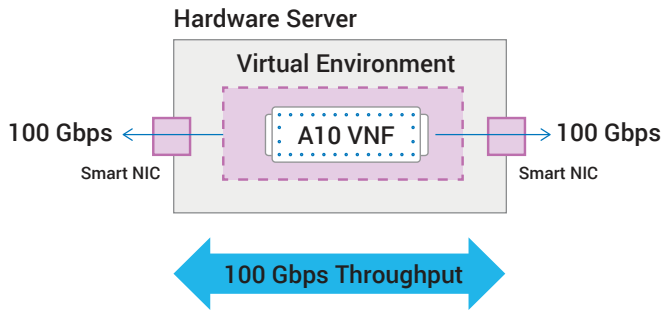

Figure 4: VNF high-performance platform (single NIC)

The A10 Networks products have been benchmarked on NFV optimized platforms, providing 100 Gigabyte network throughputs on single VNF instances. Physical server configured with a single CPU.

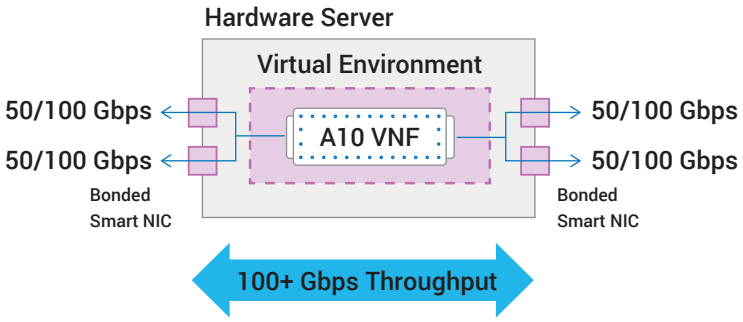

Figure 5: VNF high-performance platform (bonded NIC ports)

The above configuration uses SmartNIC paravirtualization technologies to bond multiple NIC ports, providing higher network bandwidth. The bonded ports are presented to the VNF as a single virtual port.

Benchmark tests for this configuration has achieved throughput well above 100Gbps.

STACKED NETWORK FUNCTIONS

A10 Networks group network and security products operate together within a single operating environment. Any of these products can be mix and matched, or stacked, to deliver solutions to meet business and technical demands.

Stacked network functions function similarly to chained functions but run within a single VNF instance, increasing performance while reducing cost and processing overhead.

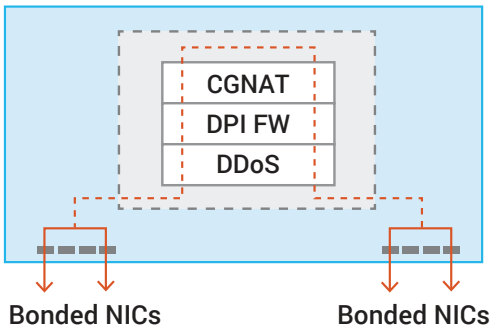

Figure 6: Stacked Network Functions within a single VNF

A10 STACKABLE VNF APPLICATIONS

- Applications Delivery Controller (ADC)
- Carrier Grade NAT
- Deep Packet Inspection Firewall
- Web Access Firewall
- GTP Firewall
- Gi-LAN Firewall
- SSL Inspection

The above diagram shows a logical view of a single VNF with three major networking functions stacked within a single VNF instance.

Performance advantages of stacking include:

- Reduced network latency
- Reduced CPU footprint and memory requirements
- Uses a single session table, ie. CGN, FW and DPI
- And perform a single lookup to save CPU and reduce processing overhead

When Service Functions are chained through separate operating environments, the network traffic is required to travel over multiple hops. A typical service chain of three Service Functions must route network traffic to and from all three VNF instances, reducing network latency with each hop.

OPTIMIZED VIRTUALIZATION PLATFORM

The virtualization platform may be traditional IT data center technologies like VMware, KVM, Hyper-V and others. This platform increasingly includes containerization technologies like Docker and Kubernetes.

A10 Networks supports the major hypervisor and cloud technologies including driver and software, certified testing and specific performance tuning.

There are various techniques to optimize the software virtualization platform. This section is an overview of with some guidelines.

VIRTUALIZATION AND HYPERVISOR PLATFORMS

The virtualization platform is responsible for providing the NFVi infrastructure for hosting and managing virtual network functions., providing virtual resources including compute, storage and virtualized networks.

NETWORK PARAVIRTUALIZATION

Traditional virtualization technology emulates physical computer hardware with a software to create an abstraction layer. Software applications like operating systems run and interact with software emulating a physical computer. Translating binary instructions and performing privileged instructions is extremely complex.

HYPERVISOR TUNING PARAMETERS

NUMA Awareness	Enable
CPU Pinning	Enable
Huge Memory Pages	1 GB or Greater
Process Affinity	Enable
Jumbo Frames	Enable
Network MTU	6000 or Greater

With this native form of virtualization, components like motherboards, device buses, BIOS subsystems, disk and network devices, hardware interrupts, timers and memory page tables must be emulated. This requires additional compute resources and is considerably slower.

The paravirtualization design goals was to replace complex hardware emulation with simplified API calls. Guest virtual machines can be configured with virtualized IO drivers. These drivers make friendly calls to virtualized network subsystems. This creates a further abstraction layer between virtualization guests and the underlying hardware and software infrastructure.

Paravirtualization technologies increase operational flexibility and agility and are becoming a common part of modern environments such as NFV infrastructures.

NETWORK ARCHITECTURE WITH PARAVIRTUALIZATION

Network architectures for NFVi often include a stack of virtualization and paravirtualization technologies. The components involved include the VNF guest environments, the virtualization platform and the hardware configuration of the host servers.

The prevalent network paravirtualization technology for NFV infrastructures is the vHost and VirtIO components. vHost/VirtIO is a virtualized device interface abstraction specification and technology. vHost is the back-end interface and drivers on the host. VirtIO is the front-end interface and driver components on the VNF guest operating systems.

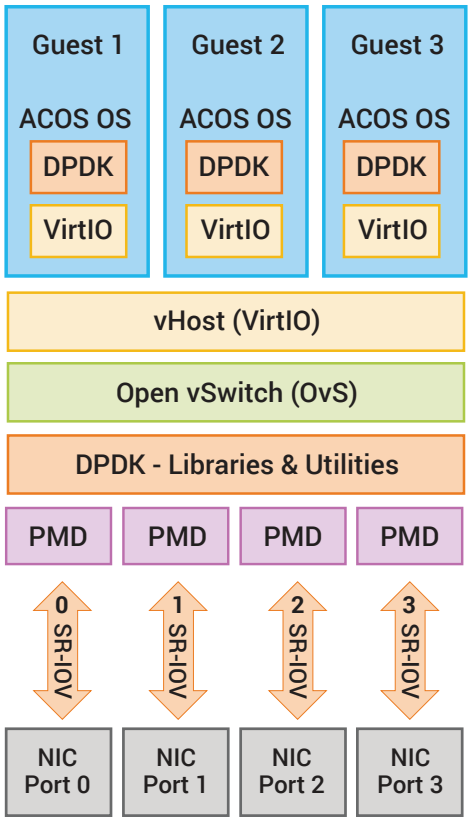

Figure 7: Paravirtualization enabled network architecture

NFV Guest is the VNF with

- A10 ACOS VNF operating system
- DPDK integrated
- VirtIO transparent drivers

vHost is the host machine component of VirtIO, and together create the paravirtualization abstraction for network communications.

Open vSwitch (OvS) is used with vHost. OvS adds additional functionality including VM to VM communications within the host and across host machines.

DPDK (Data Plan Development Kit) network acceleration technologies. Libraries and software utilities installed on the NFV server platform.

Poll Mode Drivers (PMDs) work with DPDK and VirtIO to increase network performance. PMD batches network transfers into efficient loads and remove CPU interrupts

SR-IOV network IO acceleration

NIC network hardware with DPDK and SR-IOV support

Figure 7 shows a combined virtualized, paravirtualized and hardware assisted architecture. Each technology shown is optional but adds unique functionality. VirtIO can work with SR-IOV without DPDK. The above example has VirtIO paravirtualization, DPDK accelerated networking and SR-IOV for line rate network performance.

HYPERVISOR BYPASS I/O ACCELERATION

Virtualization technologies use virtual network switches running inside hypervisors to process network traffic between the host computer network interface cards to the VNF virtual machines. This places load on the host CPUs, reduces throughput and increases latency. The diagram below shows the data path between the VNF network adapter and the physical host NIC hardware.

Paravirtualization technologies like VirtIO can be combined with each of the following techniques.

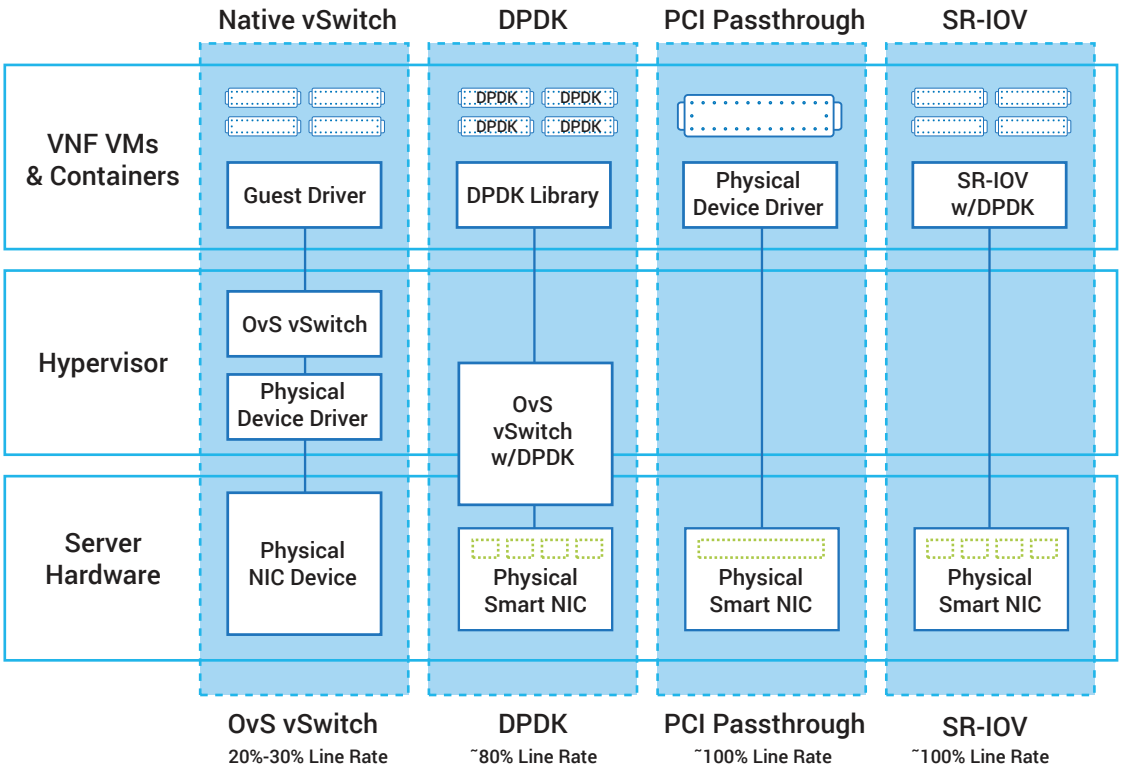

Figure 8: Logical comparison network acceleration technologies. Each technology above describes the architecture in isolation. Note: Line Rate is the raw throughput rate of the physical network.

DPDK - Data Plane Development Kit is a technology where VNF applications with the DPDK code and configurations bypass the hypervisor subsystem, directly accessing the NIC hardware, bypassing the intermediate network layers.

PCI Passthrough allows a VNF instance to bypass the hypervisor layer, communicating directly with PCIe attached devices, like physical NIC cards.

VNF operating systems must have the physical NIC drivers installed, and each NIC port is dedicated to a single guest port.

SR-IOV – Single Root Input/Output Virtualization
SR-IOV also allows guest machines to communicate directly to PCIe devices but supports multiple virtualized guest machines.

SR-IOV creates a hardware abstraction layer with one or multiple logical network (PF) devices per physical port. Multiple virtual adapters (VF) are available for guest machines using VF drivers.

FOLLOW THE PACKET PATH

NATIVE vSWITCH

1. Packet arrives at the network port
2. The hardware server initiates a CPU interrupt to process the incoming packet
3. The VNF creates another interrupt in the vCPU
4. The hypervisor copies network packets from the host server NIC to the VNF

SR-IOV OR PCI PASSTHROUGH

1. Packet arrives at the host | network port
2. The host NIC writes packet data directly into the VNF's memory

PCI PASSTHROUGH VERSUS SR-IOV

SR-IOV and PCI Passthrough have similar performance characteristics, approaching that of physical network line rates.

Differences between the technologies exist in security, usability and operational capabilities.

	PCI PASSTHROUGH	*SR-IOV*
MANAGEMENT		
State Change Detection	No	Yes (w/DPDK)
VM mobility/live migration	No	Yes (w/DPDK)
Adapter Status/Counters	Yes (NIC drivers expose counters)	Yes
Monitor Network Activity	No	Yes (per VF, PF or physical adapter statistics)
QoS Support	No	No
Snapshots	No	Yes (w/DPDK)

	PCI PASSTHROUGH	SR-IOV
SCALABILITY		
Multiple VM/VNFs	No (one VM per NIC port)	Yes (Limited by NIC VF capacity)
Multi-Queue	Yes	Yes (one CPU per VF and not universally available)
SECURITY/RELIABILITY		
Host Isolation from VM/VNF	No (drivers on VM have direct access to host hardware)	Yes (isolated by VF/PF abstraction)
High Reliability	Yes (when used with dedicated VNFs)	Yes

I/O ACCELERATION TECHNOLOGY COMPARISON

The table below list the common network acceleration technologies with several advantages and disadvantages.

	NETWORK PERFORMANCE	VMM INTEGRATION	HARDWARE OFFLOAD	COMMONLY AVAILABILITY	SIMPLE
OvS Native vSwitch	20%	Yes	No	Yes	Yes
OvS w/DPDK	80%[4,5,6]	Yes	No	Yes	Yes
OvS w/SR-IOV	99%*[4,5,6]	Yes	Yes	Yes	Yes
DPDK	69%[11]	Yes	No	No	No
DPDK + PMD	80%[4,5,6]	Yes	No	No	No
DPDK w/SR-IOV	99%[4,5]	Yes	Yes	No	No
PCI Passthrough	100%	No	Yes	Yes	Yes
SR-IOV	99%	No	Yes	Yes	Yes
VirtIO	70%[11]	Yes	No	Yes	Yes
VirtIO w/DPDK+PMD	80%[1,11]	Yes	No	No	No
VirtIO w/SR-IOV	99%[1,11]	Yes	Yes	No	No

Table of I/O Acceleration Technology Features

OPTIMIZED COTS HARDWARE

Optimized hardware platforms are available in purpose-built appliances. This document will only cover industry standard Common Off-the-Shelf (COTS) servers. Service provider infrastructure architectures are following cloud infrastructure trends, moving services to virtualized platforms residing on generic COTS hardware servers.

To obtain the performance and capacity requirements of network service providers with standard server hardware requires specific tuning configurations on both software and server hardware components.

SERVER ARCHITECTURE

Network architects migrating away from special purpose appliances to standard COTS server hardware must be well versed with computer architectures.

5G network servers must process massive volumes of traffic at extremely low latency by virtual network functions. Special attention must be paid to server hardware components involved in processing network packets.

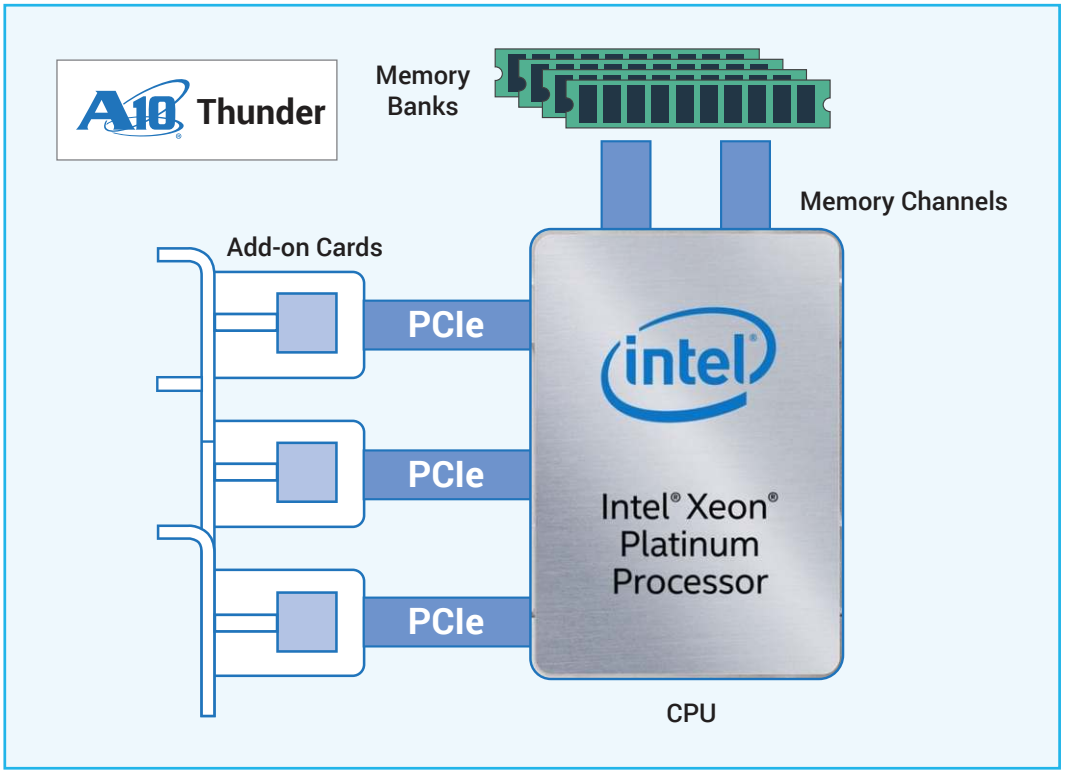

Figure 10: Simplified server architecture components

The above diagram shows a single CPU configuration with the components in the path of data packets.

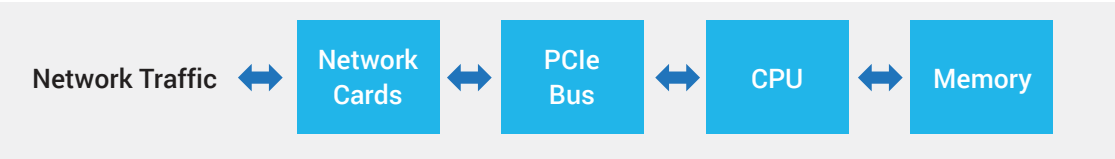

Figure 11: Simplified server architecture components

The diagram above shows a high-level flow of network traffic through a single CPU server. Many other variations and architectures exist, and the above is a simplified example. Hardware performance bottlenecks will occur in one of the four subsystems. VNFs which process network packets reside in server memory. Network traffic enters a NIC attached to a PCIe bus attached to a CPU to the VNF process. Return traffic follows the reverse path.

NUMA AND CPU CORE AFFINITY

Server architectures with multiple CPU architectures or symmetric multi-processing (SMP) architectures have additional data paths and require more considerations, specifically NUMA, in packet processing use cases.

Non-Uniform Memory Access (NUMA) is a technology on multi-CPU servers which interconnects processing groups to operate as an integrated system. Processes running on a CPU can access resources attached to every other CPU on the motherboard.

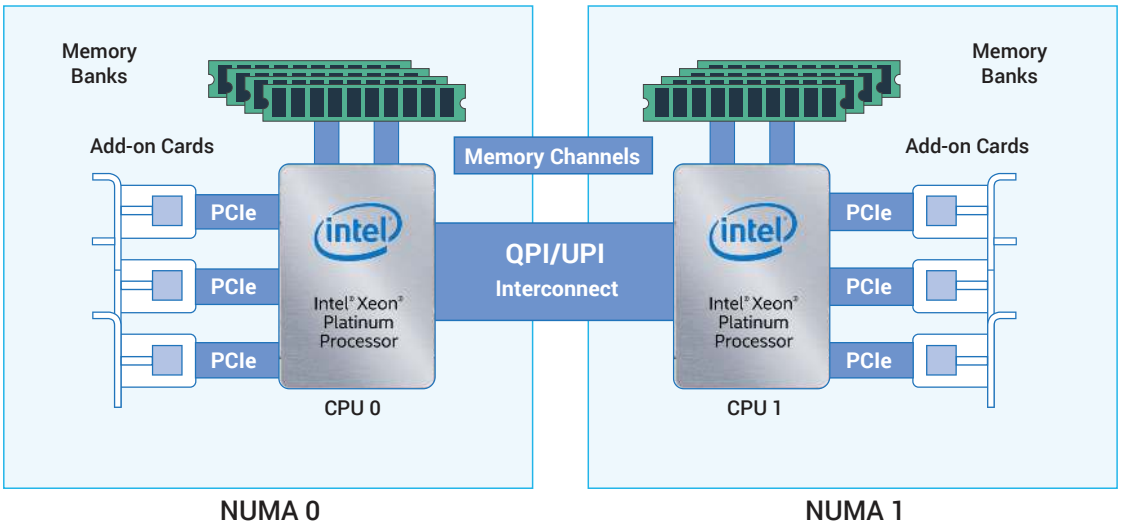

Each physical CPU and attached resources including memory and connected PCIe buses are grouped by a NUMA address. Each NUMA is interconnected with a dedicated point-to-point communications technology. The Intel version is called QuickPath Interconnect (QPI). In 2017, Intel released the Ultra Path Interconnect (UPI), replacing QPI.

Each CPU (NUMA) has locally-attached memory and PCIe devices (such as network cards). Processes accessing devices or memory attached to remote NUMA resources use the QPI/UPI interconnect. Network packet processing applications, like many VNFs, quickly suffer severe performance degradation.

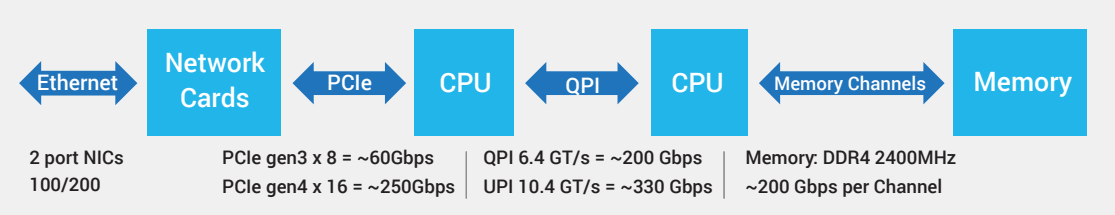

Figure 12: Packet data paths with throughput speeds

The above diagram shows data paths from the network ports through the server hardware to the memory banks. Included are very rough throughput speeds at for each technology.

NUMA PERFORMANCE SCENARIOS

The table below shows three examples of data packet paths through a multi-CPU server.

GOOD	BAD	UGLY
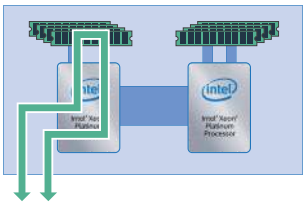	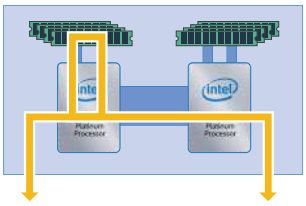	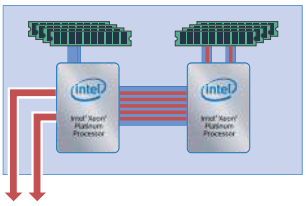
High Throughput Low Latency	**High Throughput High Latency**	**Low Throughput High Latency**
Routing symmetry: Data processed within a single NUMA node.	**Routing Asymmetry:** Data routed across two NUMA nodes through the QPI/UPI interconnect channel.	**Remote Memory Access:** Data packet processes requiring remote memory accesses through the QPI/UPI channel.
Effect: Short efficient data path: Data packets route through local network devices, CPU and Memory.	**Effect:** Long data path with multiple hops, NIC to PCIe to CPU0 to QPI/UPI to CPU1 to PCIe to NIC.	**Effect:** Packet processing performs hundreds to thousands of memory look-ups per packet.
Result: Best performance results for high-throughput and low packet latency.	**Result:** High-throughput/ Low-Latency	**Result:** Poor to disastrous performance.

Table 1: Three common scenarios of data and memory access paths

References: (1,5,7)

GOOD: This is the optimal configuration. All data packet memory accesses and processing are performed locally within a single NUMA node.

BAD: The second (BAD) scenario routes data packets across multiple hardware systems.

- High throughput can be expected since each of systems in the data path have high throughput speeds. The QPI/UPI interconnect performance is above a single NIC port.
- Latency performance is greatly reduced. Moving data to each system requires copy operations, protocol processing and multiple CPU interrupts. Each hop in the data path increases packet latency.

UGLY: On the third scenario, data packet processing generally requires access to previous packets. Examples include maintaining network states, deep packet inspection, signature analysis, subscriber awareness, policy controls and similar activities require maintaining memory state. Some network protocols require hundreds to thousands of accesses to remote memory banks.

NUMA OPTIMIZED SYSTEMS

To prevent data traversing the QPI/UPI interconnect, the NFV platform has to be designed with a "shared nothing" architecture. Data packets associated with network flows or sessions must be processed on a common CPU. Network hardware must be attached to the same CPU.

Network hardware NUMA Affinity technologies ensure network traffic is routed to the appropriate CPU cores. Software processes are located on these same CPU cores.

ARCHITECTURAL GUIDELINES

NUMA awareness technologies can resolve this problem. NUMA awareness has to be supported by multiple hardware and software components.

1. Hardware servers must have BIOS NUMA features and must be enabled
2. Intelligent NIC should have NUMA aware features
3. The server OS must have NUMA awareness features
4. The Hypervisor layer has to support NUMA
5. Each VNF OS must both support and be enabled for NUMA awareness and CPU pinning

REMOTE ACCESS INTERRUPTIONS

ACCESSING DATA ACROSS NUMA NODES

1. Packet arrives at NIC port
2. NIC interrupts local CPU
3. Local CPU copies data to the QPI channel
4. Remote CPU is interrupted to copy data from the QPI channel
5. Remote CPU copies data to attached memory or PCIe device
6. Return data travels the same path, interrupting both the remote CPU and the local CPU to copy data

This configuration will ensure VNF systems will run on a local NUMA node. Multiple NUMA aware Intelligent NIC cards work as a team, steering traffic to the companion NIC cards attached to the appropriate NUMA node.

HARDWARE OFFLOADING

Offloading network and security operations to purpose designed hardware components provide substantial hardware performance increases. Specialized hardware components offload the physical host servers, increasing throughput performances and reducing packet latency.

Network optimization and offloading is primarily available as physical server add-on cards. These cards are functionally computers with embedded CPU and memory and custom logic and electronic components. **Some of the offload functions provided by these add-on cards are:**

Intelligent NICS: These cards include network ports with speeds as high as 200 Gbps and provide higher level network services, offloading the host server CPU, memory and virtualization resources. Processing network packets on the NIC hardware reduces server CPU interrupts and much of the overhead of layer 2 and 3 packet processing.

Smart Network Interface Cards (Smart NICs): SmartNICs are based on Intelligent NIC add-on cards with additional electronic components. SmartNICs can be extended with software code to perform any logic for network packet processing. This functionality can be added with custom ASIC electronics, FPGA components or programmed directly on the SmartNIC.

One example add-on, 5G and NFV networking can be optimized and accelerated, such as 5G Control Plane processing offloaded to the SmartNIC. Another example is process packet steering and 5G Network Slicing at the network.

Software Defined Networking (SDN) multi-tenant, encapsulated overlay network protocols like VxLAN and NVGRE are processed directly, enabling packet scaling on multi-core CPU architectures.

TLS/SSL offloading processor-intensive public-key encryption for Transport Layer Security (TLS) and its predecessor Secure Sockets Layer (SSL).

NUMA Awareness is enabled by directing network traffic to a NUMA node hosting common VNF instances and PCIe devices.

NIC OFFLOAD TECHNOLOGIES

1. SSL/TLS Encryption Offload
2. SDN Encapsulation
3. Large Send Offload (LSO)
4. TCP/IP Offload Engine (TOE)
5. TCP/UDP Checksum Offload
6. IPsec
7. GRE/VxLAN/NVGRE
8. DPI
9. Port Mirroring
10. TCP Chimney Offload
11. Receive Segment Coalescing (RSC)
12. VLANs
13. Port Bonding
14. Receive Side Scaling (RSS)
15. Network Address Translation (NAT)

A10 NETWORKS NFVi SOLUTIONS

A10 Networks solutions have been adopted by major telco providers world-wide and are consumed in both physical and virtual form factors.

The 100 Gbps virtual machine provides the fastest throughput available in the market and helps customers transition to an agile, scalable and software defined network function virtualization (NFVi) required for the commercial roll-out of 5G networks.

A10 Networks NFV solutions include:

- **Highest Performance** – Up to 100 Gbps, 2.5x the performance of similar solutions.
- **GiLAN Consolidation** – Consolidates GiFW, CGNAT, application visibility and intelligent traffic steering to increase operational and better security in the GiLAN.
- **Application Visibility & Control** – DPI based L7 application visibility for effective policy enforcement and control provides enhanced Law Enforcement Agency (LEA) support and compliance.
- **Intelligent Traffic Steering** – Subscriber aware traffic steering to enhance new business models with differentiated service offerings for new revenue streams.
- **GPRS Tunneling Protocol (GTP) Support** – GTP Firewall with granular SCTP filtering defends the mobile core against GTP based attacks initiated from RAN or GRX/IPX networks.
- **Orchestration and management** – Integration with leading orchestration solutions, making A10 Networks VNF and PNF form factors highly suited for 5G environments.
- **MEC and Cloud Solutions** – Flexible form factor options from small footprint containerized VNFs to high performance bare metal PNFs, suitable for network edge and MEC environments to core packet processing applications.

A10 NETWORKS VNF/PNF PRODUCT FORM FACTORS

A10 Networks VNF/PNF products are available in a broad set of form factors.

PNF FORM FACTORS		VNF FORM FACTORS	
			
Appliance	**Software/COTS (Bare Metal)**	**Hypervisor (VM)**	**Containers**
High Performance Proprietary Hardware Components (FPGA/ASICS)	High Performance IT Standard Servers	Optimized Performance KVM, Openstack, VMWare, ESXI, Hyper-V	Container/Dockers Kubernates Integration

By providing both virtual and physical solutions, A10 Networks helps ease the integration of combining PNF and VNF in the same infrastructure.

STACKED AND CHAINED

A10 solutions offer both network function chaining and function stacking configurations. Chaining offers a flexible technique for broad communication solution requirements. Stacking network functions reduces overhead and packet latency for high-throughput requirements.

A10 HIGH PERFORMANCE VNF ARCHITECTURE

The A10 Networks VNF architecture is a high-performance carrier grade solution for NFV virtualization and containerized environments.

LEVERAGING CUTTING EDGE TECHNOLOGIES	o Paravirtualization o Docker native o Kubernates Integration + CNI o High Performance NIC Hardware o Smart NIC Hardware	o HDP (High Performance Driver) support o 100Gbps + Throughput o DPDK o Poll Mode Drivers o SR-IOV and PCI Passthrough o VirtIO

A10 NETWORKS VNF INTEGRATED ACCELERATION FEATURES

A10 Networks VM and bare metal VNF products support all compatible combinations of I/O network acceleration technologies. The A10 Networks containerized VNF products support VirtIO/vHost, SR-IOV and both combined.

	VM/VNF	*CONTAINER*	*BARE METAL*
DPDK	✓	✓	✓
DPDK w/SR-IOV	✓		
DPDK W/PCI Passthrough	✓	✓	✓
VirtIO	✓	✓	
VirtIO w/DPDK	✓	✓	
VirtIO w/SR-IOV	✓	✓	
OvS	✓	✓	
OvS w/DPDK	✓		
OvS w/SR-IOV	✓		

Table 2: I/O Acceleration Technologies supported by A10 VNFs

PERFORMANCE BENCHMARKS

The benchmark results provided was generated from a telco customer with specific performance requirements.

The configuration includes an A10 Networks VNF product, hosted on a generic COTS server and configured with performance optimizations documented here.

The throughput of 100 Gbps network throughput was obtained on a single VNF instance.

Performance throughputs well over 100 Gbps have been achieved in A10 laboratory environments.

Major configuration parameters are listed here:

A10 ACOS OS version	*4.1.4-DEV-P3-b31*
Network Cards	2 x 100 Gbps
I/O Acceleration	SR-IOV
NUMA Affinity	Enabled
Huge Page Memory	1 GB

NIC OFFLOAD TECHNOLOGIES

PERFORMANCE RESULTS FOR SINGLE VNF RUNNING CARRIER GRADE NAT OPERATIONS.

Total Throughput	**100 Gbps**
Total CPS (HTTP)	**200 K**
Total PPS	**19.9 M**
Total Sessions	**5 M**
Total Users	**100K**
Log records per second	**240 Mbps/ 24K PPS**

A10 NETWORKS: YOUR PARTNER FOR HIGH PERFORMANCE

Network Service Providers have the task of moving their infrastructure to NFV based infrastructures while expanding throughput capacity and performance simultaneously. This new architecture requires NFV/VNF network speeds unheard of today.

A10 Networks has the world's highest performance VNF platform.

A10 Networks customers include the largest cloud and network service providers world-wide. Our product portfolio is extensive, and deployments include Multi-Access Edge Computing, 5G and EPC Packet Core, Central Office infrastructures and Public Clouds infrastructures.

The job of developing and deploying your new 5G infrastructure has arrived. The technical challenges are overwhelming. A10 Networks can help.

ABOUT A10 NETWORKS

A10 Networks (NYSE: ATEN) provides Reliable Security Always™ through a range of high-performance solutions that enable intelligent automation with deep machine learning to ensure business critical applications are protected, reliable and always available. Founded in 2004, A10 Networks is based in San Jose, Calif., and serves customers globally with offices worldwide.

For more information, visit: **a10networks.com** or tweet **@A10Networks.**

1. Mellanox Socket Direct – Removes NUMA QPI bottle neck: http://www.mellanox.com/related-docs/whitepapers/WP_Mellanox_Socket_Direct.pdf
2. Mellanox Brochure: http://www.mellanox.com/related-docs/products/Ethernet_Adapter_Brochure.pdf
3. Mellanox SmartNIC: http://www.mellanox.com/related-docs/prod_adapter_cards/PB_BlueField_VPI_Smart_NIC.PDF
4. Intel VNF SR-IOV DPDK: https://www.intel.com/content/dam/www/public/us/en/documents/technology-briefs/sr-iov-nfv-tech-brief.pdf
5. Intel paper NFV Performance Optimization: https://software.intel.com/en-us/articles/nfv-performance-optimization-for-vcpe
6. Whitepaper VNF Performance: https://www.researchgate.net/publication/299489473_Enhancing_VNF_performance_by_exploiting_SR-IOV_and_DPDK_packet_processing_acceleration
7. Redhat Virtualization Tuning: https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/virtualization_tuning_and_optimization_guide/index
8. Intel SR-IOV Driver Companion: https://www.intel.com/content/dam/doc/design-guide/82599-sr-iov-driver-companion-guide.pdf
9. DPDK Performance Optimization Guidelines: http://doc.dpdk.org/guides/prog_guide/perf_opt_guidelines.html
10. Virtio_user for Container Networking: http://doc.dpdk.org/guides/howto/virtio_user_for_container_networking.html
11. Intel DPDK vHost/VirtIO Performance Report: https://fast.dpdk.org/doc/perf/DPDK_19_02_Intel_virtio_performance_report.pdf

LEARN MORE

ABOUT A10 NETWORKS

CONTACT US

a10networks.com/contact

Part Number: A10-WP-21155-EN-04 AUG 2019