US Patent: "Network hierarchy constructs enabling instant global wide area network connectivity" https://patents.justia.com/patent/11924046

Nikolai Pitaev (2024) “Making SD-WAN Easy” Cisco Live On Demand Online Library, Las Vegas.

Nikolai Pitaev, Aris Leivadeas, Sai Suchandan Reddy Marapareddy (2024) CIoT'24, paper #1571049518, 'Informed Network Routing for SD-WAN enabled SaaS Optimization: An M365 Proof of Concept'.

Nikolai Pitaev, Cisco SD-WAN Cloud onRamp for Multi-Cloud lab (2021). https://github.com/CiscoDevNet/sdwan-cor-labinfra

This project creates virtual infrastructure with Terraform on public cloud, which will be used for Cisco SD-WAN Cloud onRamp for Multi-cloud lab. Some modules can be used separate from each other. For example Centralized Firewall Design or Multi-Region Fabric use cases can be deployed totally separated.

Cisco Live Master Series 2020, Barcelona. https://youtu.be/TTJzV9Fd_j0?si=Fj-w-Jw1mgXw4zEx

From Cisco Live 2020, Barcelona the Master Series with Nikolai Pitaev. The Master Series offers a unique opportunity to meet and engage with Cisco's top speakers virtually.

Gayathri Chandrasekaran, Nikolai Pitaev, Muffadal Quettawala (2020). How to Automate and Secure Branch Office Connectivity to AWS with Cisco SD-WAN.

In this post, we will explore how customers can leverage native integrations from Cisco Systems, an AWS Advanced Technology Partner, that work with AWS network services to interconnect branches over the AWS network backbone.

Asma Ben Hamed, Aris Leivadeas, Matthias Falkner and Nikolai Pitaev (2020). VNF Chaining Performance Characterization under Multi-Feature and Oversubscription Using SR-IOV.

Network Function Virtualization (NFV) has revolutionized the way network services are offered, leading Enterprise and Service Providers to increasingly adapt their portfolio of network products in order to reap the benefits of flexible network service deployment and cost reduction promises. With this method, network services are offered in the form of software images instead of dedicated hardware. However, NFV presents several challenges, including standard networking challenges (e.g., security, resilience, and availability), management and orchestration challenges, resource allocation challenges, and performance trade-off challenges of using standard x86 servers instead of dedicated and proprietary hardware. The first three challenges are typical challenges found in virtualization environments and have been extensively addressed in the literature. However, the performance trade-off challenge can be the most impactful when offering networking services, negatively affecting the throughput and delay performance achieved. Thus, in this paper, we investigate and propose several configurations on a virtualized system for increasing the performance in terms of throughput and delay while chaining multiple virtual network functions (VNFs) in case of an undersubscribed and oversubscribed system, where the resource demands exceeds the physical resource capacity. Specifically, we use the Single Root Input Output Virtualization (SR-IOV) as our Input/Output (I/O) technology, and analyze the attainable throughput and delay when running multiple chained VNFs in a standard x86 server under various resource footprints and network features configurations. We show that the system throughput and delay in a multi-chained environment, offering multiple features, and under oversubscription can affect the overall performance of VNFs.

Aris Leivadeas, Matthias Falkner, Nikolai Pitaev (2020). Analyzing Service Chaining of Virtualized Network Functions with SR-IOV. Best Presentation Award at 2020 IEEE 21st International Conference on High Performance Switching and Routing (HPSR).

Network Function Virtualization (NFV) along with Service Function Chaining (SFC) has proliferated the way that
network functions and services are introduced, offered, and deployed. In particular, NFV and SFC are the basic components
of the so called network softwarization era, promoting significant cost reductions and agility.
Nonetheless, a constant debate exist regarding the network performance achieved and the holistic establishment

of NFV as the key and sole component of the network infrastructure. To this end, this paper tries to shed
light in remedy solutions when chaining Virtualized Network Functions (VNFs) internally on a standard x86
server using the Single-Root Input Output Virtualization (SR-IOV) technology, in order to maximize the throughput
and delay achieved.

Nikolai Pitaev, Matthias Falkner, Aris Leivadeas, Ioannis Lambadaris (2018). Characterizing the Performance of Concurrent Virtualized Network Functions with OVS-DPDK, FD.IO VPP and SR-IOV. Best Industry Paper Award on ICPE April 2018. Berlin, Germany.

The virtualization of network functions is promising significant cost reductions for network operators. Running multiple network functions on a standard x86 server instead of dedicated appliances can increase the utilization of the underlying hardware,while reducing the maintenance and management costs of such functions. However, total cost of ownership calculations are typically a function of the attainable network throughput, which in a virtualized system is highly dependent on the overall system architecture - in particular the input/output (I/O) path. In this paper we investigate the attainable performance of an x86 host running multiple virtualized network functions (VNFs) under different I/O architectures: OVS-DPDK, SR-IOV, and FD.io VPP. Running multiple VNFs in parallel on a standard x86 host is a common use-case for cloud-based networking services. We show that the system throughput in a multi-VNF environment divers significantly from deployments where only a single VNF is running on a server.

Nikolai Pitaev, David Klebanov (2018). Delivering Cisco Next Generation SD-WAN with Viptela BRKCRS-2110. February 2018. Cisco Live On Demand Online Library, Barcelona, Spain.

Main focus is on the fundamental principles behind Cisco SD-WAN solution powered by the cloud-delivered Viptela technology.
We will discuss basic solution elements including secure elastic connectivity, application quality of experience, cloud adoption and agile operations.
Zero-touch deployment strategies and new use cases behind greenfield and brownfield implementations will be shown.
We will cover vEdge routers, SD-WAN Controllers (vSmart, vBond and vManage) and key communication flows between different solution elements.

Nikolai Pitaev, Matthias Falkner, Aris Leivadeas, Ioannis Lambadaris (2017). Multi-VNF performance characterization for virtualized network functions. NetSoft August 2017. Bologna, Italy. Electronic ISBN: 978-1-5090-6008-5, USB ISBN: 978-1-5090-6007-8, Print on Demand (PoD) ISBN: 978-1-5090-6009-2.

Network Function Virtualization promises to reduce the overall operational and capital expenses experienced by the network operators. Running multiple network functions on top of a standard x86 server instead of dedicated appliances can increase the utilization of the underlying hardware and reduce the maintenance and management costs. However, total cost of ownership calculations are typically a function of the attainable network throughput, which in a virtualized system is highly dependent on the overall system architecture - in particular the input/ output (I/O) path. In this paper, we investigate the attainable performance of an x86 host running multiple Virtualized Network Functions (VNFs) under different I/O architectures: OVS, SRIOV and FD.io VPP. We show that the system throughput in a multi-VNF environment differs significantly from deployments where only a single VNF is running on a server, while different I/O architectures can achieve different levels of performance.

Ian Wells, Nikolai Pitaev (2017). NFV Performance - Challenges and Solutions BRKSDN-2411. July 2017. Cisco Live On Demand Online Library, Las Vegas, USA.

Performance in a virtualized world has a lot of challenges that can limit the capabilities that a NFV data plane is expected to deliver on. We will start with the basics of how VNFs are networked in a virtualized environment and then explain where are the bottlenecks reside in a Linux/KVM/Qemu environment. We will explain more on those bottlenecks and then share how we can resolve them to achieve the right level of performance. We will also share performance data from testing while applying similar concepts in an implementation to illustrate how the performance bottlenecks shifts from one place to another depending on various factors like the type of VNF, features enabled inside them, use case, networking model etc.

Nikolai Pitaev (2017). Providing Scalable Broadband Subscriber Solutions with Virtualisation and Orchestration BRKSPG-2381. March 2017. Cisco Live On Demand Online Library, Melbourne, Australia.

Broadband Network Gateway (BNG) or Broadband Remote Access Router (BRAS) plays critical role in every Service Provider network. Many thousands of physical BNGs like Cisco Aggregation Switch Router (ASR) 1000 and / or ASR 9000 Series are successfully deployed worldwide providing Internet Access and Broadband Services to millions of subscribers. Virtual BNG (vBNG) is the next evolutional step coming along Network Function Virtualisation.
New vBNG solution includes Placement in SP Network, Server Hardware and Hypervisor Guidance, IOS XE and / or XR BNG Features and Performance, brand new ESC Orchestration Software and also includes Monitoring and Operation Recommendations.
We will cover all relevant topics and focus on the following key components:
1. CSR1000v and XRv acting as virtual BNG
2. Elastic Services Controller (ESC) as Orchestration Software for elasticity and flexibility managing vBNG.
3. Hardware and Hypervisors setup for the best performance
4. Smart Licensing for vBNG
Applied science in teal life will be given using the first deployment use case.

Nikolai Pitaev, Fan Gu (2016). New Cloud networking solutions with Cisco Cloud Services Router (CSR 1000V) on AWS and Azure. LTRDCN-2100. February 2016. Cisco Live On Demand Online Library Premium Content. Berlin, Germany.

We will describe how to deploy and use CSR1000V in Amazon Web Services (AWS) and Microsoft Azure as scalable virtual network function. When you need to scale above multiple VPCs networking between regions and physical locations becomes a challenge. Managing VPN endpoints per VPC and creating a mesh network to communicate between all sites becomes complex and unmanageable. That's why we came up with the new "Transit VPC" Solution in order to solve the complexity of networking between VPC-to-VPC and VPC-to-physical locations. We will demonstrate how Cisco CSR1000v delivers in this solution a highly available network solution that centralizes security, firewall and routing functions within a single point that is scalable and easy to manage. New Highlights: usage of AWS CloudFormation template, which will automatically create the resources needed in the Transit VPC and launch the creation of the CSR1000v. AWS Lambda function automatically pushes the high availability and tunnel configs down to the CSR Spoke VPCs are automatically launched and configured from the Transit VPC GUI console. Azure will be interconnected with AWS Turn on application monitoring on CSR with WebGUI to visualise your traffic.

Miercom, Nikolai Pitaev (2016). Virtual Routers - the Impact of Configuration Changes on Throughput Performance. An Independent Assessment. November 2016. Miercom DR161111E. San Jose, USA.

Main Goal is to validate the performance of the software-based Cloud Services Router 1000V, designed to run in a third-party-managed cloud environment and allow an enterprise to extend its network security and services into the cloud. In addition to testing CSR 1000V, we also evaluated a variant form of the IOS XE router software – the Virtual Integrated Services Router, or ISRv, in an NFVIS (Network Functions Virtualization Infrastructure Software) environment running on an Enterprise Network Compute System (ENCS) Server. The testing assessed the impact of various configuration settings and deployment options on the throughput performance of the CSR 1000V package in public and private cloud environments.

Matthias Falkner, Nikolai Pitaev, Nickolay Belchugov, Ehab Hadi (2016). Network Function Virtualization Seminar TECSPG-2300. February 2016. Cisco Live On Demand Online Library Premium Content. Berlin, Germany.

This technical seminar is dedicated to the evolution towards Network Function Virtualization (NFV). We will cover building blocks of the NFV framework and architecture in the following use cases: Enterprise Branch Virtualization, Multi-Cloud Interconnection, SP Infrastructure Virtualization and multi-tenanted SMB Services. Deployment considerations, best practices and lessons learnt from real live projects will be shared.
An overview of the Cisco virtualization Hosts including Enterprise Network Compute System (ENCS), 4000 Series ISR + UVS E-Series and UCS C-Series will be given. A summary of Cisco Virtual Network Functions (VNF) will be provided and the following VNFs will be presented by the appropriate Business Units experts: Cloud Services Router 1000V (CSR1000V), Adaptive Security Virtual Appliance (ASAv) and Threat Defense. VNF details on private (x86 server running a hypervisor) and public cloud (AWS/Azure), multi-tenancy and service chaining, orchestration and VNF life cycle management will be explained. Input-Output (IO) Models play an important role in the VNF design. The seminar will help you to understand different IO models including single root I/O virtualization (SR-IOV) and virtual Switched like OVS-DPDK and fd.io VPP. NFV Performance will be covered as a separate agenda point, which will include performance optimization guidelines and NFV performance test/benchmarking tools. In a live demonstration you will see presented VNFs in action.

Matthias Falkner, Jon Rosen, Yegappan Lakshmanan, Nikolai Pitaev (2016). De-mystifying VNF Performance on Virtualized System Architectures. Cisco CEC, San Jose, USA.

Enterprise and Service Provider networks are increasingly making use of virtualized network functions (VNFs) to reap the benefits of reduced CAPEX and OPEX. Total cost of ownership calculations however are typically a function of the attainable network throughput and performance, which in a virtualized system is highly dependent on the overall system architecture. The number of VNFs running on the server, their I/O demands, the performance characterization of the underlying hypervisor scheduler or the packet path from physical interfaces into the VNFs are examples of how the system architecture can influence overall performance and throughput. This paper examines the impact of the system architecture on VNF performance under both single-VNF and multi-VNF configuration. In particular, the effects of different I/O mechanisms such as virtual switching, PCIe pass-through and SR-IOV are explored in environments where both a single VNF or multiple VNFs are running on the server.

Nikolai Pitaev (2012). Eishockey für Kinder und Eltern. ISBN-10: 3849119203. November 2012. tredition publishing house, Hamburg, Germany.

Das erste Buch für Kinder, die leidenschaftlich Eishockey spielen, sowie für ihre Eltern, die den Kindern dabei gerne helfen. Es beschreibt zuerst Eishockey als "Breitensport". Dabei wird auf Organisation in Deutschland sowie Kanada/USA, Kommunikation zwischen Trainer, Spieler und Eltern sowie typische Anfänger-Fehler eingegangen. Danach geht es um Leistungssport - unterschiedliche Entwicklungswege sind beschrieben: von einer Uni-Karriere mit Sport-Stipendium bis hin zu NHL. Sehr viele der beschrieben Themen sind sportartübergreifend und somit für andere Sportarten zutreffend.

Nikolai Pitaev (2005). New options for the deep packet inspection and prediction algorithms for the impact of the bandwidth profiles. September 2005. Deutsche Telekom Research Center Darmstadt, Germany.

Research project under NDA between Deutsche Telekom and NETSCOUT. Public announcement: https://www.netscout.com/

Nikolai Pitaev (2003). New modeling techniques for the biggest broadband aggregation IP network in Europe. April 2004. Deutsche Telekom Research Center Darmstadt, Germany.

Research project under NDA between Deutsche Telekom and OPNET. Public announcement: http://www.opnet.com/spotlight/deutschetelekom.html

Nikolai Pitaev (1999). IP-Telephony and Java-based API. Diploma Thesis. Technical University Darmstadt, Library KOM FB18, Prof. R. Steinmetz.

Investigation of Java-based APIs and IP Telephony. Main research goal is investigate role of Java-based API in IP-Telephony and impact to Communication Systems.

Nikolai Pitaev, Mattias Sommer (1998). FPGA-based Test Hardware for IP- and ATM-based Networks. Student research project between Technical University Darmstadt and T-Systems (Deutsche Telekom Research Center in Darmstadt).

Prototype of the FPGA-based hardware test device for IP- and ATM-Networks, which can be used in performance testing of the communication systems. Details are subject to non-disclosure agreement (NDA) of Deutsche Telekom.