ANALYSIS OF VIRTUAL NETWORKS IN DATA CENTERS.

Ionka Gancheva, PhD student[45]

Abstract: *The article contains an analysis of virtual networks and technologies that are used at data centers nowadays. Many different solutions can be implemented and the best depends of the needs of enterprises.*

Key words: *Data center, virtual network, network architecture, virtual switch*

1. INTRODUCTION

A virtual network represents a network with a certain type of connectivity characteristic. A virtual network is not a one-to-one mapping with a specific classic network concept. An instantiation of a virtual network on a set of host groups is called a *network site*. It is possible to have a single virtual network but three different network sites for the same logical network. Furthermore, the network site can be divided into multiple VM networks and associate VMs with the VM network.

2. DATA CENTER VIRTUAL NETWORK ARCHITECTURE WITH THE USE OF "DISTRIBUTED SWITCH"

A distributed switch is a managed entity configured in VMware vCenter Server. The distributed switch abstracts a set of network switch functions that are configured on each associated host. vCenter Server manages the configuration of distributed switches, and the configuration is consistent across all hosts. Distributed switch is a template for the network configuration on each VMware ESXi host [1].

Each distributed switch includes distributed ports. They can be connected for any networking entity, such as a virtual machine or a VMkernel interface to a distributed port. vCenter Server stores the state of distributed ports in the vCenter Server database. Networking statistics and policies migrate with virtual machines when the virtual machines are moved from host to host [2].

[45] New Bulgarian University, 21 Montevideo str., Sofia, Bulgaria

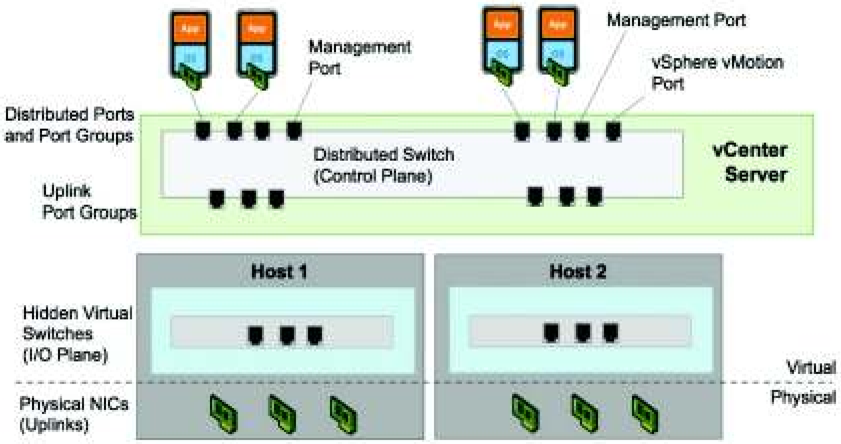

Figure 1: Distributed Switch Architecture

A distributed switch functions as a single virtual switch across all associated hosts. Distributed switches have several benefits over standard switches: They simplify data center administration. They enable networking statistics and policies to migrate with virtual machines during a VMware vSphere vMotion migration. They provide for customization and third-party development.

Having the network configuration at the data center level (VMware vSphere Distributed Switch), not at the host level (standard switch), has several advantages [3]:

- Data center setup and administration are simplified by centralizing network configuration. For example, adding a host to a cluster and making it compatible with VMware vSphere is much easier.
- Distributed ports migrate with their clients. For example, when you migrate a virtual machine with VMware vSphere vMotion, the distributed port statistics and policies move with the virtual machine, thus simplifying debugging and troubleshooting.
- Distributed switches support private VLANs. With private VLANs, VLAN IDs can be used in a private network without worrying about duplicating VLAN IDs across a wider network.
- Enterprise networking vendors can provide proprietary networking interfaces to monitor, control, and manage virtual networks. VMware vSphere Network Appliance API enables third-party developers to create distributed switch solutions [2].

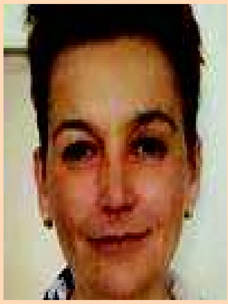

Ionka Gancheva *is PhD candidate at Informatics Department at New Bulgarian University almost 2 years.*
In 1997 Ionka Gancheva earned his Master`s Degree of Informatics science at the D.A. Tsenov University of Svishtov.
Ionka Gancheva currently hold Industrial certification: CCNA R&S, CCNP R&S, VCP5. She has had an extensive teaching experience at NBU. Her courses are in the field of networking and virtualization.

3. DATA CENTER VIRTUAL NETWORK ARCHITECTURE WITH THE USE OF "DISTRIBUTED SWITCH"MICROSOFT HYPER-V NETWORKING.

Hyper-V was introduced as a role as part of Windows Server 2008 and it has become available as a standalone version called Microsoft Hyper-V server [4].
Among other things, using Windows core OS means that there is no graphical user interface and the configurations are performed via Windows PowerShell.

The feature set of Hyper-V is similar to the features of VMware ESX. For instance, instead of vMotion migration, Hyper-V has Live Migration. Instead of a vNetwork Distributed Switch (vDS), Hyper-V has a logical switch. Instead of the concept of the data center, Hyper-V has a folder. The suite of products that manages Hyper-V is called *System Center [4]*. The following list provides the key components and terminology for Hyper-V:

- System Center Virtual Machine Manager (SCVMM): Runs on a centralized server and manages virtualized hosts, VMs, storage, and virtual networks. Equivalent to vCenter;
- Virtual Machine Management Service (VMMS): A process running in the parent partition of each virtualized server that uses the WMI interface. It manages Hyper-V and VMs on the host;
- Hyper-V Switch: The extensible virtual switch in the hypervisor.
- Windows Management Instrumentation (WMI): Used by SCVMM to interface with VMMS on the host;
- Windows Network Virtualization (WNV): A module that adds network virtualization generic routing encapsulation (NVGRE) capabilities to build overlays.

Figure 2 displays the Hyper-V architecture and interaction with the key components just defined.

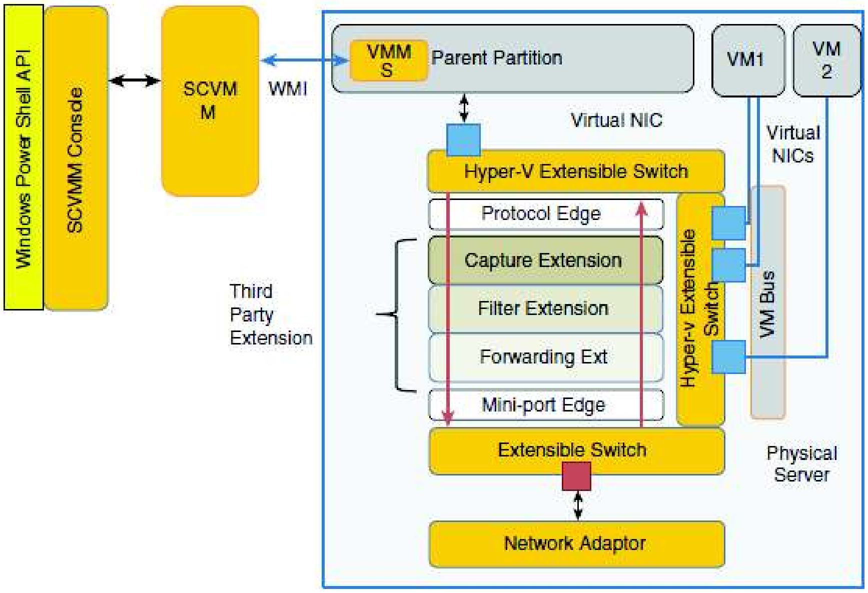

Figure 2 Hyper-V architecture

Another key concept in Microsoft Hyper-V is the *forwarding extension*, which allows the insertion of third-party processing in the data path from the guest to the network adapters. A forwarding extension can accomplish the following in the two directions of the data path: Filter packets; Inject new packets or modified packets into the data path; Deliver packets to one of the extensible switch ports. Figure 3 shows a topology with multiple host groups.

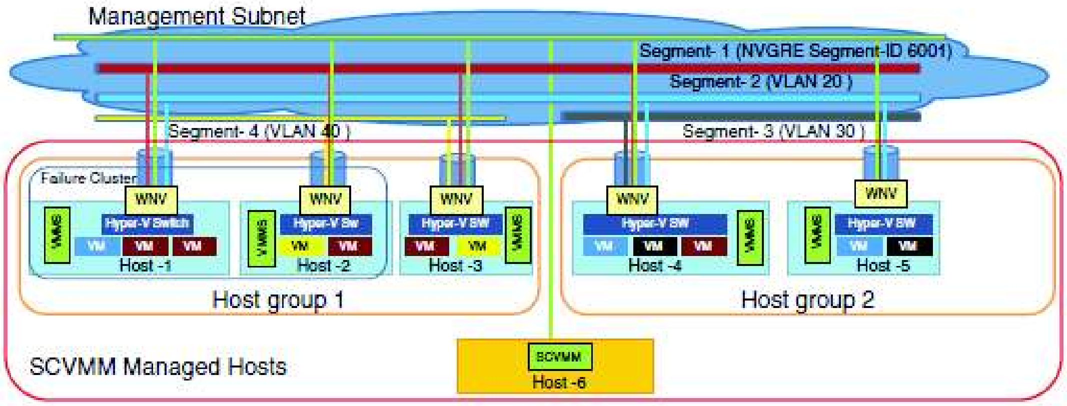

Figure 3: Topology with multiple host groups

4. DATA CENTER VIRTUAL NETWORK ARCHITECTURE WITH THE USE OF LINUX KVM AND NETWORKING

Linux Kernel-based Virtual Machine (KVM) is part of the kernel, but it doesn't perform hardware emulation—a user-space element provides this. The management of virtual machines in Linux is achieved by using two elements:

- **libvirt:** A toolkit that enables the interaction with the virtualization features of Linux. Virt-viewer, virt-manager, and virsh (shell to manage virtual machines) rely on libvirt;
- **qemu:** A hardware-emulation component that runs in user space in KVM.

Figure 4 illustrates the relationship between these components.

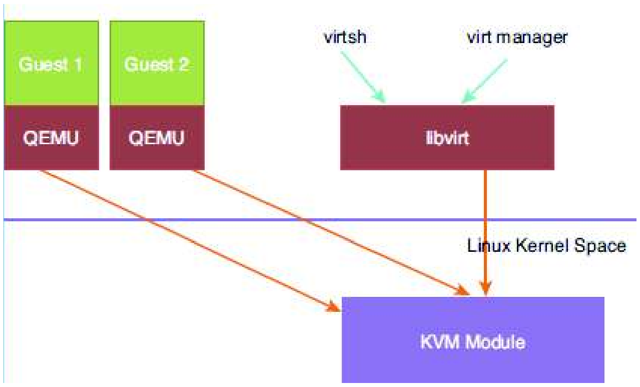

Figure 4: Components in a Virtualized Server Running KVM

When running KVM, the following packages also should need to be install:

- **virt-manager:** A GUI tool that manages KVM guests;
- **virt-install:** A command-line tool to install virtual machines;
- **virt-viewer:** The virtual viewer.

Open vSwitch works on hypervisors such as KVM, XenServer, and VirtualBox. Open vSwitch can run as a standalone virtual switch, where every virtual switch is managed independently, or it can run in a “distributed” manner with a centralized controller by exposing these two configuration elements:

- Flow-based forwarding state, which can be remotely programmed via OpenFlow;
- Switch port state, which can be remotely programmed via the Open vSwitch Database (OVSDB) management protocol Open vSwitch also supports the ability to create GRE- or VXLAN-based tunnels.

5. DATA CENTER VIRTUAL NETWORK ARCHITECTURE WITH THE USE OF OPEN VSWITCH

Open vSwitch (OVS) is a software switch with many networking features, such as: IEEE 802.1Q support; NetFlow; Mirroring.

Open vSwitch can run as a standalone virtual switch, where every virtual switch is managed independently, or it can run in a “distributed” manner with a centralized controller by exposing these two configuration elements:

- Flow-based forwarding state, which can be remotely programmed via OpenFlow;
- Switch port state, which can be remotely programmed via the Open vSwitch;
- Database (OVSDB) management protocol Open vSwitch also supports the ability to create GRE- or VXLAN-based tunnels.

One of the key characteristics of Open vSwitch is that it has a flow-based forwarding architecture. This is similar to the concept of a control plane and data plane separation in many Cisco architectures where the supervisor provides with Hypervisors the data plane handles the packet-forwarding capabilities. This particular aspect of OVS allows it to run in a distributed manner in OpenFlow architectures [4]. OVS has three main components:

- A kernel component implementing the fast path;
- A user space component implementing the OpenFlow protocol;
- A user space database server.

Figure 5 illustrates the architectural components of OVS.

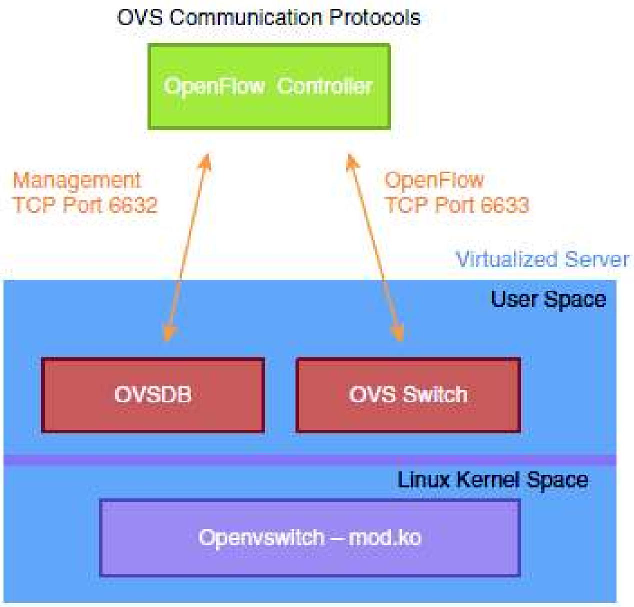

Figure 5 Architectural components of OVS.

The solutions to enable VM-to-VM communication are a virtual switch such as the Cisco Nexus 1000V or Open vSwitch.

6. ARCHITECTURE OF VIRTUAL NETWORK WITH THE USE OF CISCO NEXUS 1000V

Cisco Nexus 1000V is a feature-rich software switch that runs on multiple hypervisors. Cisco Nexus 1000V provides functions such as: ACL filtering on individual VM ports; Switched Port Analyzer (SPAN), or Remote SPAN, features of individual VMs; NetFlow statistics of the local traffic; Capability to shut down VM ports individually [6].

The Cisco Nexus 1000V consists of two main components: the Virtual Supervisor Module (VSM, the control-plane component) and the Virtual Ethernet Module (VEM, the data-plane component). Together these components provide the abstraction of a physical switch, whose supervisor is the VSM and whose line cards are the VEMs that run within each VMware ESX host. All configurations are performed on the VSM and propagated to the VEMs that are associated with it. A VSM can be a virtual machine and run redundantly just like a redundant supervisor. It is possible to add a VMware ESX host to the Cisco Nexus 1000V vDS from VMware vCenter to make a VMware ESX host become part of a Cisco Nexus 1000V domain, and as a result run a VEM. A VSM running as a virtual machine provides the abstraction of a CLI managing a large modular switch [6]. The user employs Secure Shell (SSH) Protocol at the management interface of the VSM, or simply uses the console—the virtual machine console screen—to configure the network characteristics of the VMware deployment. The VSM forwards the configurations (VLANs, QoS, private VLANs, etc.) to all the VEMs that are part of the same domain or, in other words, that are under the same Cisco Nexus 1000V.

The communication between VSM and VMware vCenter uses the management interface (mgmt0) on the VSM. The protocol runs on HTTPS. The key information is provided to VMware vCenter by pointing the browser to the VSM IP address and downloading the extension key, extension.xml, which is added to VMware vCenter as a plug-in.

REFERENCES

[1] Vmware USA [Official site] VMware® NSX for Multi-Hypervisor (NSX-MH) Network Virtualization Design Guide, [Visited on 20.01.2016]
[2] Vmware USA [Official site] NSX Operations Guide Rev. 1.3 August 2015, [Visited on 7.08.2015]
[3] Vmware USA [Official site] NSX for vSphere Getting Started Guide VMware NSX for vSphere, release 6.0.x July 21, 2014, [Visited on 23.02.2016]
[4] Openvswitch USA [Official site] Virtual switch for OpenvSwitch.x July 21, 2014, [Visited on 3.03.2016]
[5] Microsoft USA [Official site] https://www.microsoft.com/en/server-cloud/solutions/virtualization.aspx [Visited on 3.05.2015]
[6] Cisco Nexus 1000V Switch for VMware vSphere [Official site]: http://www.cisco.com/c/en/us/products/switches/nexus-1000v-switch-vmware-vsphere/index.html [Visited on 13.12.2015]