

Virtual Machine Placement

Strategies for Virtual Network Functions

Table of Contents

Executive Summary	3
Introduction.....	3
VNF/NFV Background.....	3
Current Workload Scheduling.....	3
Traditional Cloud Workloads.....	4
NFV Requirements	4
Scheduling Options	5
NetworkBWFilter.....	6
TelemetryFilter.....	6
Conclusion.....	7
About Juniper Networks.....	7

Executive Summary

Traditional cloud workloads service virtual machine (VM) requests using a specified set of filters and metrics. However, a gap exists in current cloud management software, which needs to be addressed in order to make it more suitable for Network Functions Virtualization (NFV) applications. Network awareness is a requirement of NFV, and is absent from the metrics currently used to schedule virtual workloads.

This paper discusses how to navigate those options using OpenStack, the emerging de facto standard for cloud management software.

Introduction

Current VM scheduling services, such as OpenStack's Nova, only have awareness of the CPU, RAM, and storage utilization from a hypervisor perspective. While this has proven to be sufficient for traditional cloud applications and their associated workloads, which are rarely limited by network bandwidth, virtualized network functions (VNFs) require fairly static and known quantities of CPU, memory, and storage. However, properly placing their workloads depends upon a knowledge of hypervisor network resources and external network topology.

This paper will compare and contrast traditional cloud workloads and their placement with VNF workloads and their respective placement. We will discuss why extensions to cloud management software are needed to improve placement, quality of service (QoS), and utilization of commodity hardware, in order to improve ROI for standard hardware by providing the ability to more densely pack virtual workloads and avoid stranded resources.

VNF/NFV Background

Network Functions Virtualization (NFV) is an effort originally proposed by the European Telecommunications Standard Institute (ETSI) at the SDN and OpenFlow World Congress in 2012.¹ The objective or focus of the NFV initiative is to use standard virtualization technology to consolidate various network equipment types onto a common platform that can be easily distributed throughout the network. The NFV platform comprises high capacity servers, switches, and storage, which will utilize software applications running on virtual machines to perform functions like routing, switching, and security. These platforms can then be located in the data center, head end, or at the end users' premises.

Service providers benefit from NFV in the following ways. First and foremost, NFV provides for the ability to increase agility and the velocity of new service deployment, improving time to market. Virtualizing network equipment improves the ability to automate, which allows for faster service instantiation. NFV can help operators by reducing the required initial investment in equipment cost and power consumption as many network functions can be combined onto standard servers. As in traditional cloud application environments, services can be more quickly and even automatically scaled to meet increased demand.

With NFV, service providers now have the ability to create virtual devices providing network services on demand that can be dynamically inserted into the end user's path.

Current Workload Scheduling

The industry has been moving toward OpenStack as the standard for managing cloud computing platforms for private and many public cloud installations. Therefore, because OpenStack is becoming the de facto standard and it shares many things in common with other cloud management packages, this paper will use OpenStack and its components as a representative example of cloud management software.

Before we dive into how cloud workload scheduling works, it is important to understand how cloud management platforms are structured. OpenStack, as an example, comprises several components, which represent the required functions or systems that manage the different aspects of a cloud installation². The projects that govern these components and their code names are listed below:

- Network: Neutron
- Compute: Nova
- Authentication: Keystone
- UI: Horizon
- Storage: Swift, Cinder
- Image Service: Glance
- Telemetry: Ceilometer, Gnocchi

¹ Chiosi, Margaret, et. al. (2012). Network Functions Virtualisation, An Introduction, Benefits, Enablers, Challenges & Call for Action, SDN and OpenFlow World Congress, Darmstadt-Germany

² <http://docs.openstack.org/icehouse/training-guides/content/operator-getting-started.html>

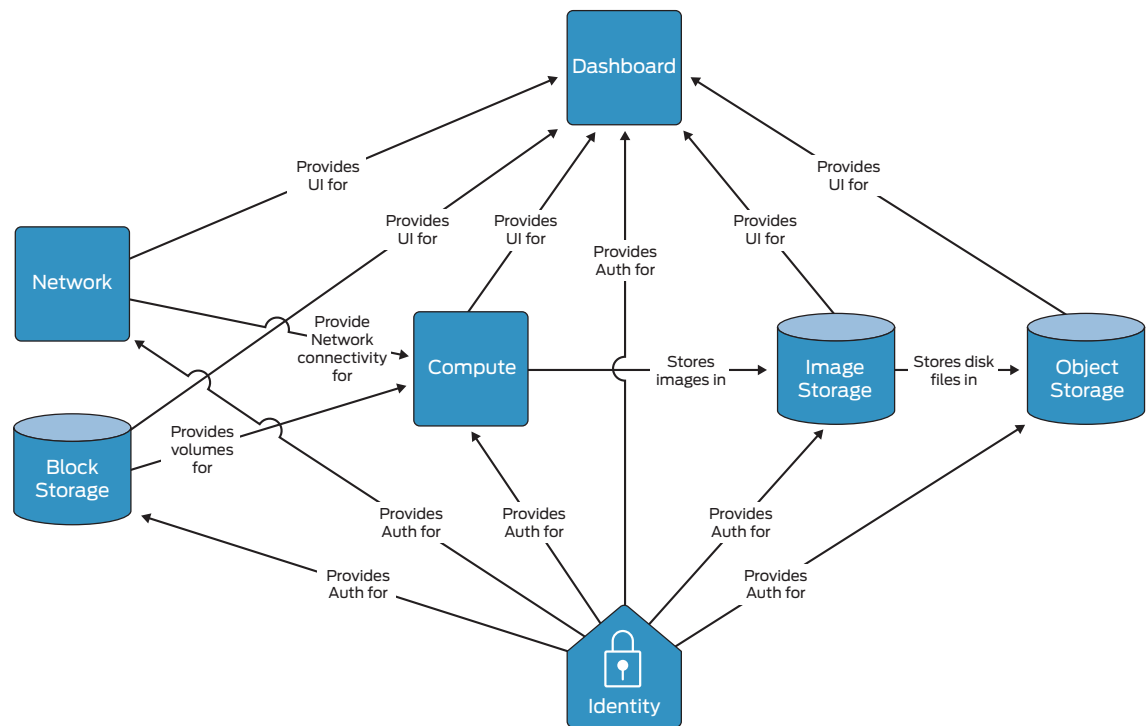


Figure 1: OpenStack components. Source: openstack.org

All of these subprojects or components work in concert to orchestrate the placement and instantiation of a virtual machine as well as its associated network, storage, and other required resources. This placement is broadly based on the following: availability zone, available RAM, storage, compute availability, and custom filters. (Note that additional items exist but for the purposes of this paper the most common metrics are referenced). When a user issues a request for a new virtual machine to be created, the scheduler will check for available resources, select the appropriate hypervisor to service the request, and trigger the VM creation/boot on the selected host from a virtual image.

Traditional Cloud Workloads

Servicing VM requests based on the above metrics has proven adequate for traditional cloud workloads, since memory has in general been the primary point of contention for web-based apps and services, followed by the availability of CPU and storage resources.

In the case of traditional cloud workloads, per-node bandwidth utilization is an infrequent issue, and on the occasion that it is, “noisy neighbors” are usually prevented from affecting other workloads with virtual switch-based QoS, or by overprovisioning hypervisor resources.

NFV Requirements

With the entrance of NFV into the equation, requirements for scheduling need to be reconsidered. Providers are starting to move network functions into the virtual world. These functions include, but are not limited to:

- Firewalls
- Intrusion detection and prevention systems (IDP)
- Route reflectors
- WAN acceleration appliances

In general, purpose-built VNFs have fixed/known requirements for CPU, storage, and memory that can be scheduled by VM orchestrators such as OpenStack Nova. It is important to note that in the case of NFV, oversubscription does need to be considered. In the world of networking, sharing CPU cores or RAM will create suboptimal network performance, so this factor need to be taken into account when making scaling decisions.

Additionally, the network requirements for each subscriber can vary greatly. As an example, the requirements for a small business, branch office, or residential customer might require a virtual firewall capable of passing an average of 25 Mbps (IMIX), while a business providing free Wi-Fi to customers would potentially need far more bandwidth. At the time of this writing, Nova does not have the ability to schedule based on network utilization or for VNF workloads; this is a gap that needs to be addressed.

Scheduling Options

Since we have established different scheduling needs for traditional cloud workloads vs. VNF, we can consider some options to make the cloud orchestration software take available network bandwidth into consideration.

Turning again to OpenStack as our example, we can look at the workflow of the filter scheduler, which makes the decision on workload placement. When a new instance is created, a series of filters are applied in order to choose the appropriate hypervisor upon which to place that instance. Filters are binary, either a host passes filtering or it is rejected from consideration.

Current filters include:³

- AggregateCoreFilter
- AggregateDiskFilter
- AggregateImagePropertiesIsolation
- AggregateInstanceExtraSpecsFilter
- AggregateIoOpsFilter
- AggregateMultiTenancyIsolation
- AggregateNumInstancesFilter
- AggregateRamFilter
- AggregateTypeAffinityFilter
- AllHostsFilter
- AvailabilityZoneFilter
- ComputeCapabilitiesFilter
- ComputeFilter
- CoreFilter
- NUMATopologyFilter
- DifferentHostFilter
- DiskFilter
- GroupAffinityFilter
- GroupAntiAffinityFilter
- ImagePropertiesFilter
- IsolatedHostsFilter
- IoOpsFilter
- JsonFilter
- MetricsFilter
- NumInstancesFilter
- PciPassthroughFilter
- RamFilter
- RetryFilter
- SameHostFilter
- ServerGroupAffinityFilter
- ServerGroupAntiAffinityFilter
- SimpleCIDRAffinityFilter
- TrustedFilter
- TypeAffinityFilter

³ http://docs.openstack.org/liberty/config-reference/content/section_compute-scheduler.html

During the scheduling process, these filters are applied based on resources requested. Additional weights are applied, and finally workloads are created on the appropriate compute nodes.

The following diagram shows the filtering workflow:

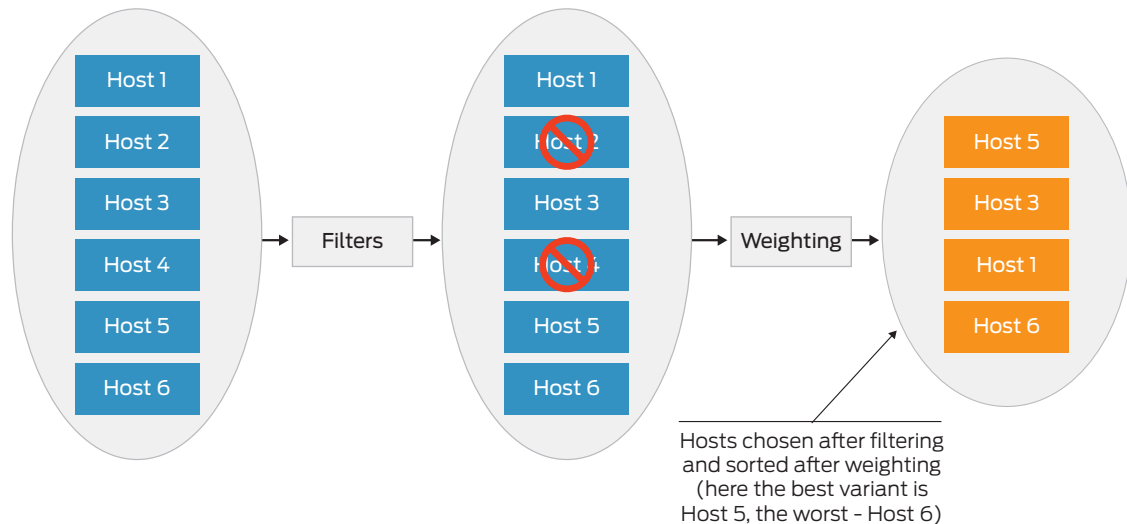


Figure 2: Filtering workflow. Source: openstack.org

NetworkBWFilter

Network awareness is a requirement of NFV, and is absent from the metrics currently used to schedule virtual workloads.

One proposed solution is to create a new filter that would allow the scheduler to take network requirements into account during the decision making process. Taking the operative example from the functionality of the `IoOpsFilter`, a filter such as `NetworkBwFilter` would be created. As part of the `NetworkBwFilter`, a `max_bw_per_host` would be set to specify a high water mark allowable on a particular compute host.⁴

When the request for a new VM instantiation is made, a required amount of bandwidth for that workload would be specified. As with existing filters, each time a host is selected for placement of VNFs, those resources are consumed virtually, and subsequent host selections can be adjusted accordingly. This would guarantee that the amount of requested bandwidth would never exceed that which is available.⁵

TelemetryFilter

Another possibility is to create a filter that would take into account usage data from a telemetry package (e.g., `ceilometer`) and use actual usage data to make host filtering decisions. This would allow for even more efficient use of network bandwidth.⁶

This filter could even make decisions based on the results of API calls to external systems. An example of this type of operation would be the scheduler making API calls to a network analytics server during the workload creation process.

Utilization information gathered from the network analytics system could be applied as either a weight or a filter during the scheduling process.

⁴ <https://forums.juniper.net/t5/Data-Center-Technologists/Networking-flavor-for-OpenStack-workloads/ba-p/274982>

⁵ <https://forums.juniper.net/t5/Data-Center-Technologists/Network-aware-placement-of-OpenStack-workloads/bc-p/277471>

⁶ <https://forums.juniper.net/t5/Data-Center-Technologists/Why-Physical-Network-Topology-API-is-relevant-for-OpenStack/bc-p/278725>

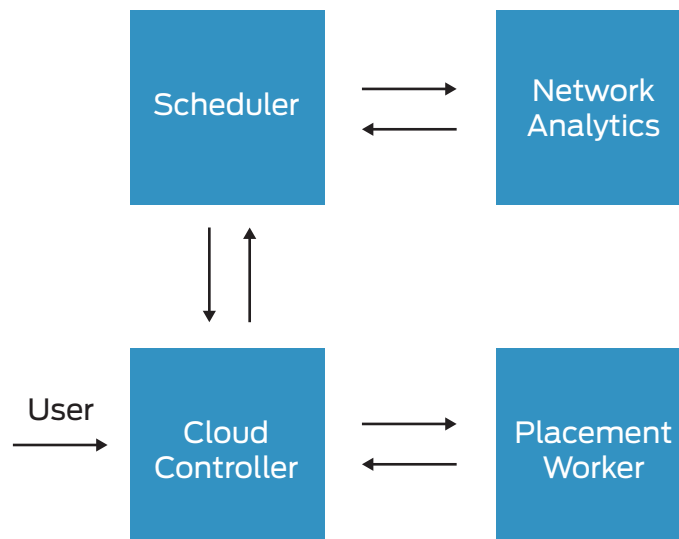


Figure 3: External Interaction with network analytics

Since reducing contention between workloads is a primary concern, continual monitoring would be required, with the potential need to respawn workloads elsewhere and redirect traffic to them when utilization changes.

Note: There exists a blueprint within OpenStack for “utilization-aware scheduling,” which is meant to take into account transient resources, allowing decisions to be made based on gathered usage statistics⁷. Work on the network monitor portion of utilization-aware scheduling appears to have been abandoned. Were this work to continue, it would be a positive step in the goals laid out by this paper, at least where the OpenStack project is concerned.

Conclusion

This paper has compared and contrasted traditional VM workload scheduling with the requirements of NFV applications and included a high-level discussion of potential solutions. A gap exists in current cloud management software which needs to be addressed in order to make it more suitable for NFV applications. Taking network utilization into account during workload scheduling decisions can improve ROI for standard hardware by providing the ability to more densely pack virtual workloads and avoid stranded resources.

About Juniper Networks

Juniper Networks challenges the status quo with products, solutions and services that transform the economics of networking. Our team co-innovates with customers and partners to deliver automated, scalable and secure networks with agility, performance and value. Additional information can be found at [Juniper Networks](#) or connect with Juniper on [Twitter](#) and [Facebook](#).

⁷ <https://wiki.openstack.org/wiki/UtilizationAwareScheduling>

Corporate and Sales Headquarters
 Juniper Networks, Inc.
 1133 Innovation Way
 Sunnyvale, CA 94089 USA
 Phone: 888.JUNIPER (888.586.4737)
 or +1.408.745.2000
 Fax: +1.408.745.2100
www.juniper.net

APAC and EMEA Headquarters
 Juniper Networks International B.V.
 Boeing Avenue 240
 1119 PZ Schiphol-Rijk
 Amsterdam, The Netherlands
 Phone: +31.0.207.125.700
 Fax: +31.0.207.125.701