NFV/NFVC VM PLACEMENT CHALLENGES

RENAUD LARSEN
CHIEF ARCHITECT, JUNIPER NETWORKS

AGENDA

VM PLACEMENT

VM Security Challenges

NFV ARCHITECURAL FRAMEWORK

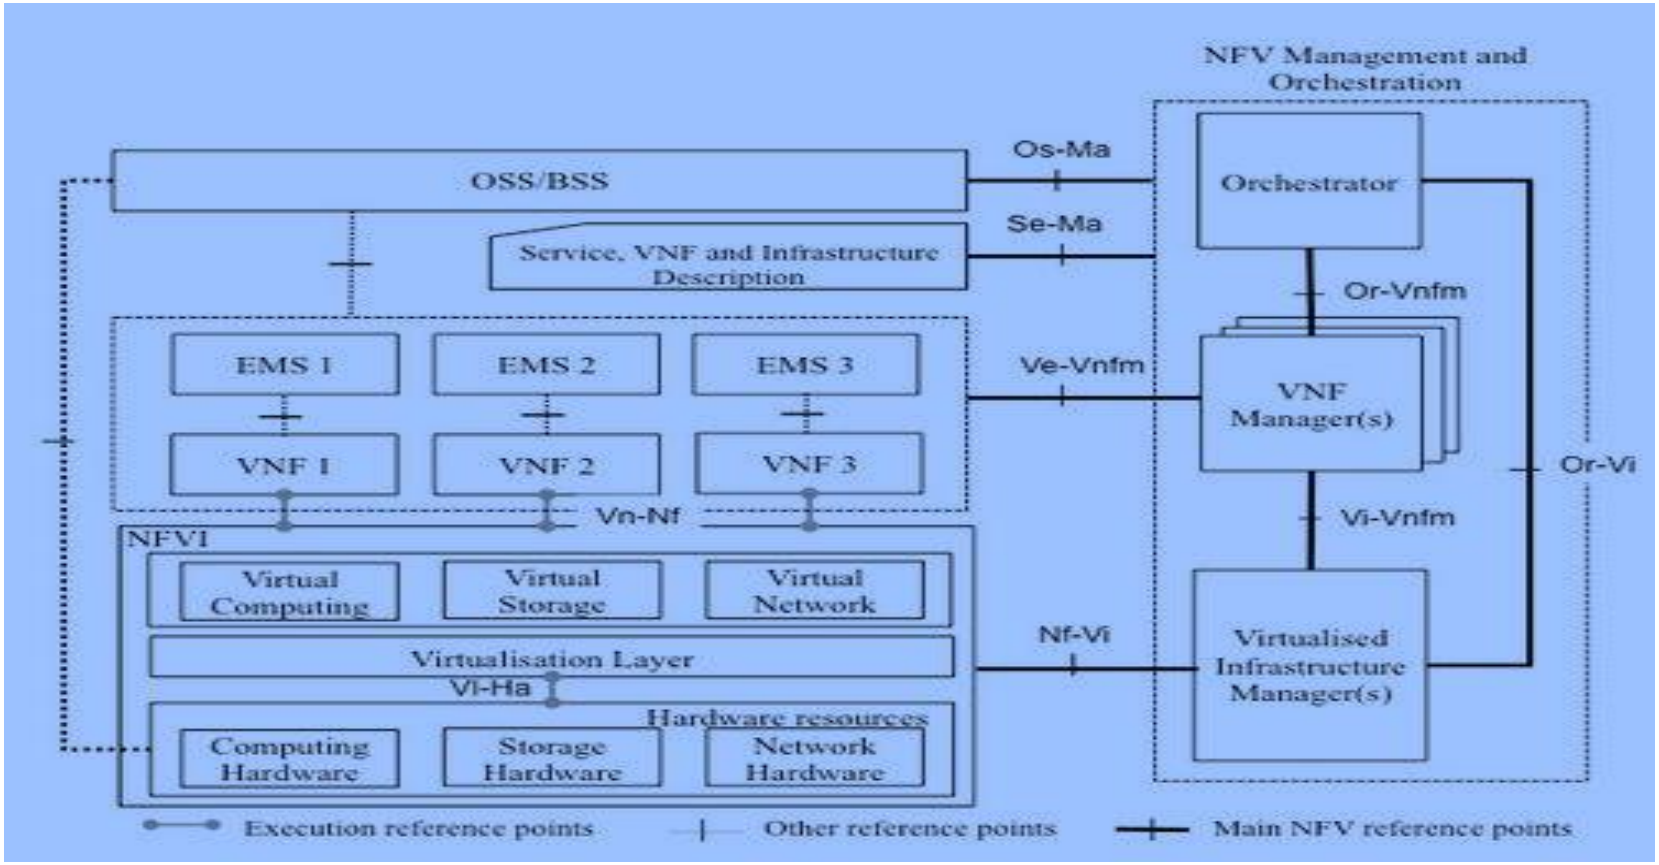

http://portal.etsi.org/portal/server.pt/community/NFV/367

LET'S GO BACK TO THE FOUNDATIONS OF VIRTUALIZATION

Popek-Goldberg principles defined for the virtual machines. I revisited their original paper defining the Formal Requirements for Virtualized Machines :

Equivalence property

Resource control property

Efficiency property

If we consider Network Virtualization as based on System Virtualization we can map the Popek Goldberg properties to Network Virtualization as such :

Equivalence : there is not gap between the overlays and the physical topology which is about the equivalence property

Resource Control : there is a proper implementation of mechanism enabling the control of the resources in order to make sure that different virtual networks will not comprise the performances of the different virtual networks while providing isolation.

Efficiency : There is a mechanism to identify packets that could be forwarded by the physical network without having to involve the virtual network level because we will need to be able to evaluate the performance impact of the packet header transformation.

NFV CHALLENGES

How it will address increases in resource demand and consolidation/deconsolidation of VNFC VMs to avoid performance overlapping or performance degradation

As we know :

In majority of the Ochestration tools available on the market today included OpenStack

- VM placement consists to randomly assign a VM to an available host
- This must not be confused with VMs provisioning leveraging scheduling mechanisms and notion like availability zones and hostfilters.
- In general the open-source cloud management systems support user-defined policies for the initial provisioning of VMs. However, all open-source cloud management systems do not directly support dynamic VM placement

ACTUAL SITUATION

As we know current virtualization technology offers the ability to easily relocate a virtual machine from one host to another without shutting it down thus giving the opportunity to dynamically optimize the placement with a small impact on performance.

Nevertheless it’s critical to be able to express the application requirements related to VM placement and server state in order to define the VM placement constraints to model a viable configuration.

Therefore it’s highly critical to identify which services are stateless and which one are stateful in order to ensure good performances

ACTUAL SITUATION

In a NFV context, we might deploy and remove VM instances at any time in an unpredictable manner depending on the service chaining associated with the customer profiles.

This behavior may lead the infrastructure toward a suboptimal or unstable configuration (for example not exploiting free space left by NFC's VMs that are not deployed).

However, the majority of the existing works ignore the dynamic nature of the incoming stream of VM deployment requests to which the cloud infrastructure is subject over time.

Moreover, VMs can show some correlation in the resource usage (i.e., for example a web server and an application server will likely have a similar CPU utilization dependency on the incoming workload

ACTUAL APPROACH

As we know we can consolidate many VMs on the same physical machine to optimize the resources and reduce the opex costs :

In order to maximize the savings :

- Administrators should pack as many VMs as possible onto a server while satisfying certain performance criteria.
  - When an administrator want to execute an application he will submit the application description to that computes a place for each VM.
  - The placement of VMs must then satisfy requirements defined by the administrator.
- As we know modern application such as Web applications have specific placement constraints which lead to what is called the Virtual Machine Packing Problem.
- Since server has multiple types of constrained resources which the VMs consume.
- Therefore we need to take into account multiple weights that must observe the following absolute requirements : the sum of the weights for any given resource must be less than or equal to the corresponding capacity.

WHAT'S THE MAIN PROBLEM

Even though the problem of Virtual Machine (VM) placement in a compute cloud infrastructure is well-studied to a certain extend; majority of the existing works

- ignore the dynamic nature of the incoming stream of VM deployment requests that continuously arrive to the cloud provider infrastructure
- address server failures and not load changes.
- do not provide high-availability by ensuring that a certain number of service are available satisfying VMs resource consumption and placement constraints.
- do not support constraints related to server state management, scheduling or relocation mechanisms.

Still today in a highly virtualization DataCenter resources allocation to jobs is done in a static mode and therefore resources are underused of misused which lead typically to the following situation : CPU usage 20% and idle 80% according to Gartner.

Therefore what would be the optimal virtual machine allocation and packing in order to minimize the number of real nodes so that DC administrator could turn off the unused nodes and minimize the number of migration while meeting the service level agreement defined with users and avoiding non-viable placement that could lead to performance degradation.

WHAT'S THE MAIN PROBLEM

In a highly virtualized DC VM ordering becomes critical to avoid configuration problem and enable dynamic consolidation to determine when and how to migrate VMs to avoid memory sharing.

Therefore it's key to define :

an efficient configuration/reconfiguration plan based on a cost function that will take into account the needed actions before migrating a VM,

the amount of memory to migrate, the networking requirements.

Server consolidation is a multi-dimensional bin packing problem requesting to take into account different resources from disk I/O, CPU, memory bus to network bandwidth.

PROPOSED APPROACH

The network architecture and location matters to enable a dynamic VM configuration/ reconfiguration plan.

Thanks to the VNFC Management Host we will collect the needed information - cpu, memory, network infos – to graphical represent the VM placement and deploy a decision making algorithm to dynamically define scheduling strategies and decide where to place the VM as well as computing the appropriate viable configuration.

Furthermore if we use traffic-aware virtual machine (VM) placement this will enable us to also improve the network scalability.

By optimizing the placement of VMs on host machines, traffic patterns among VMs can be better aligned with the communication distance between them,e.g. VMs with large mutual bandwidth usage are assigned to host machines in close proximity.

Deploy a management architecture which supports constraints programming and dynamic resource allocations with identified key performances objectives at scale.

We must not forget that available commercial implementations are limited to a small number of server (32 in the case of VSphere) and we believe that future implementations should support much larger configurations, perhaps in the order of 105–106 servers by leveraging certain distributed aspects.