Cost-efficient VNF placement and scheduling in cloud networks

Speaker: Tao Gao

05/18/2018

Group Meeting Presentation

Background

Advantages of Network Function Virtualization (NFV):

- can be customized on-demand
- have flexible options of their locations
- can be provisioned dynamically and elastically
- enable re-configuration at runtime
- …

A Virtual Network Function (VNF) is usually deployed in a Virtual Machine (VM) in a general-purpose server. Cloud service provider (Google Cloud, AWS, Azure, etc.) becomes a candidate to provide such resources:

- reasonable price of computing and network resource
- acceptable latency between different VMs
- flexible resource allocation (vCPU cores, Memory, Disk, and Network) and location selection
- pay-by-use payment method (per minute in Google Cloud, per second in AWS)

Background

Google Cloud Platform (GCP)

Custom machine types

Iowa | Monthly | Hourly

Item	Price (USD)	Preemptible price (USD)
vCPU	$0.033174 / vCPU hour	$0.00698 / vCPU hour
Memory	$0.004446 / GB hour	$0.00094 / GB hour

f1-micro Bursting (allow instances to use additional physical CPU for short periods of time)

Iowa | Monthly | Hourly

Machine type	Virtual CPUs	Memory	Price (USD)	Preemptible price (USD)
f1-micro	1	0.60GB	$0.0076	$0.0035
g1-small	1	1.70GB	$0.0257	$0.0070

f1-micro instances get 0.2 of a vCPU and are allowed to burst up to a full vCPU for short periods. g1-small instances get 0.5 of a vCPU and are allowed to burst up to a full vCPU for short periods.

"https://cloud.google.com/compute/pricing#custommachinetypepricing"

Background

General network pricing

Traffic type	Price
Ingress	No charge
Egress* to the same zone	No charge
Egress to Google products (such as YouTube, Maps, Drive), whether from a VM in GCP with a public (external) IP address or a private (internal) IP address	No charge
Egress to a different Google Cloud Platform service within the same region, except for Cloud Memorystore for Redis and for Cloud SQL	No charge
Egress* between zones in the same region (per GB)	$0.01
Egress to Cloud Memorystore for Redis is charged at the rate of "Egress between zones in the same region"	
Egress to Cloud SQL is charged at the rates described in Traffic through external IP addresses	
Egress between regions within the US (per GB)	$0.01
Egress between regions, not including traffic between US regions	At Internet egress rates

"https://cloud.google.com/compute/pricing#custommachinetypepricing"

Background

General zones diagram

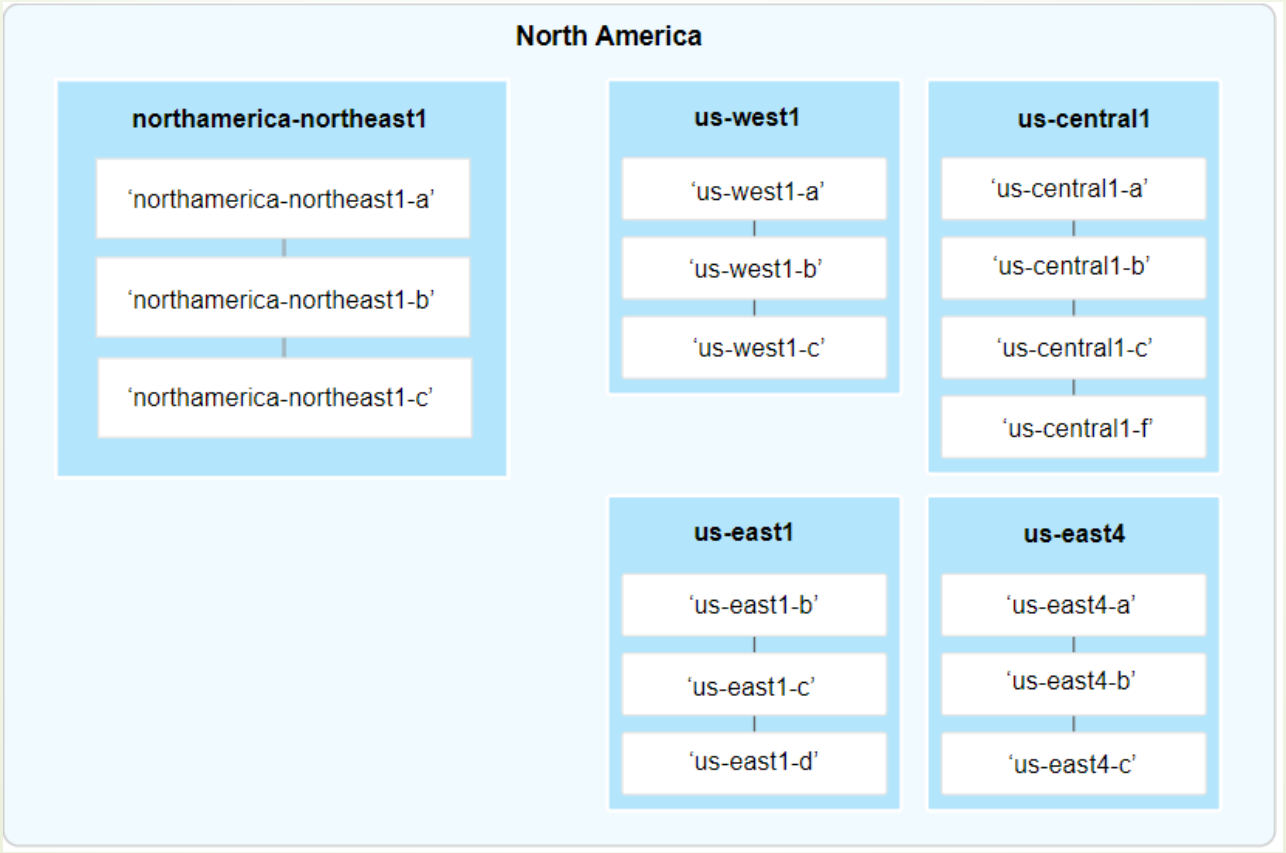

"https://cloud.google.com/compute/docs/regions-zones/"

Background

Cloud Locations

"https://cloud.google.com/about/locations/#regions-tab"

Background

Cloud Locations

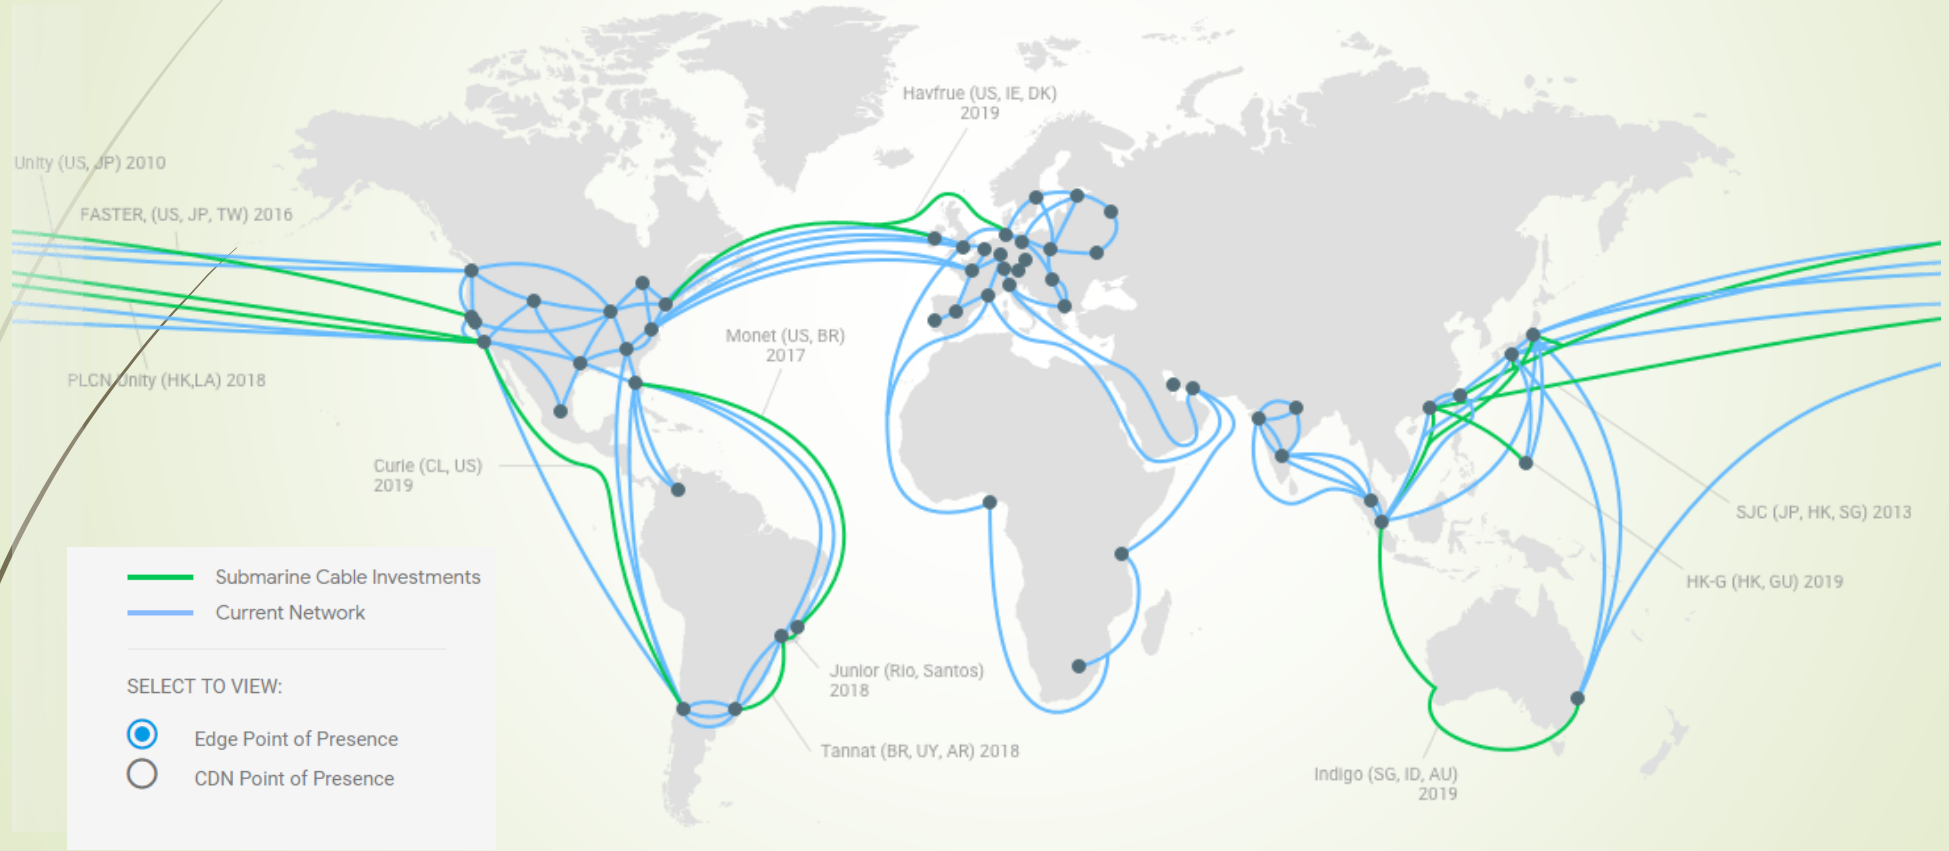

"https://cloud.google.com/about/locations/#regions-tab"

Background

Internet egress rates

Iowa

Monthly Usage	Network (Egress) Worldwide Destinations (excluding China & Australia, but including Hong Kong) (per GB)	Network (Egress) China Destinations (excluding Hong Kong) (per GB)	Network (Egress) Australia Destinations (per GB)	Network (Ingress)
0-1 TB	$0.12	$0.23	$0.19	Free
1-10 TB	$0.11	$0.22	$0.18	Free
10+ TB	$0.08	$0.20	$0.15	Free

"https://cloud.google.com/compute/pricing#custommachinetypepricing"

Problem and Challenge

For a Service Provider (SP), it is important to satisfy the users' latency requirement with minimum financial cost considering:

- How many VMs are active?
  - Busy or idle
- How many resources are allocated for the VMs?
  - CPU cores
  - Network resource
- Where are the VMs are located?
  - Distance to users' location
  - Communication between VMs
- How long should an instance of an VNF reside (as an idle state) after it finishes its task?
  - Boot time & Money
  - Frequency of usage
- …

Problem and Challenge

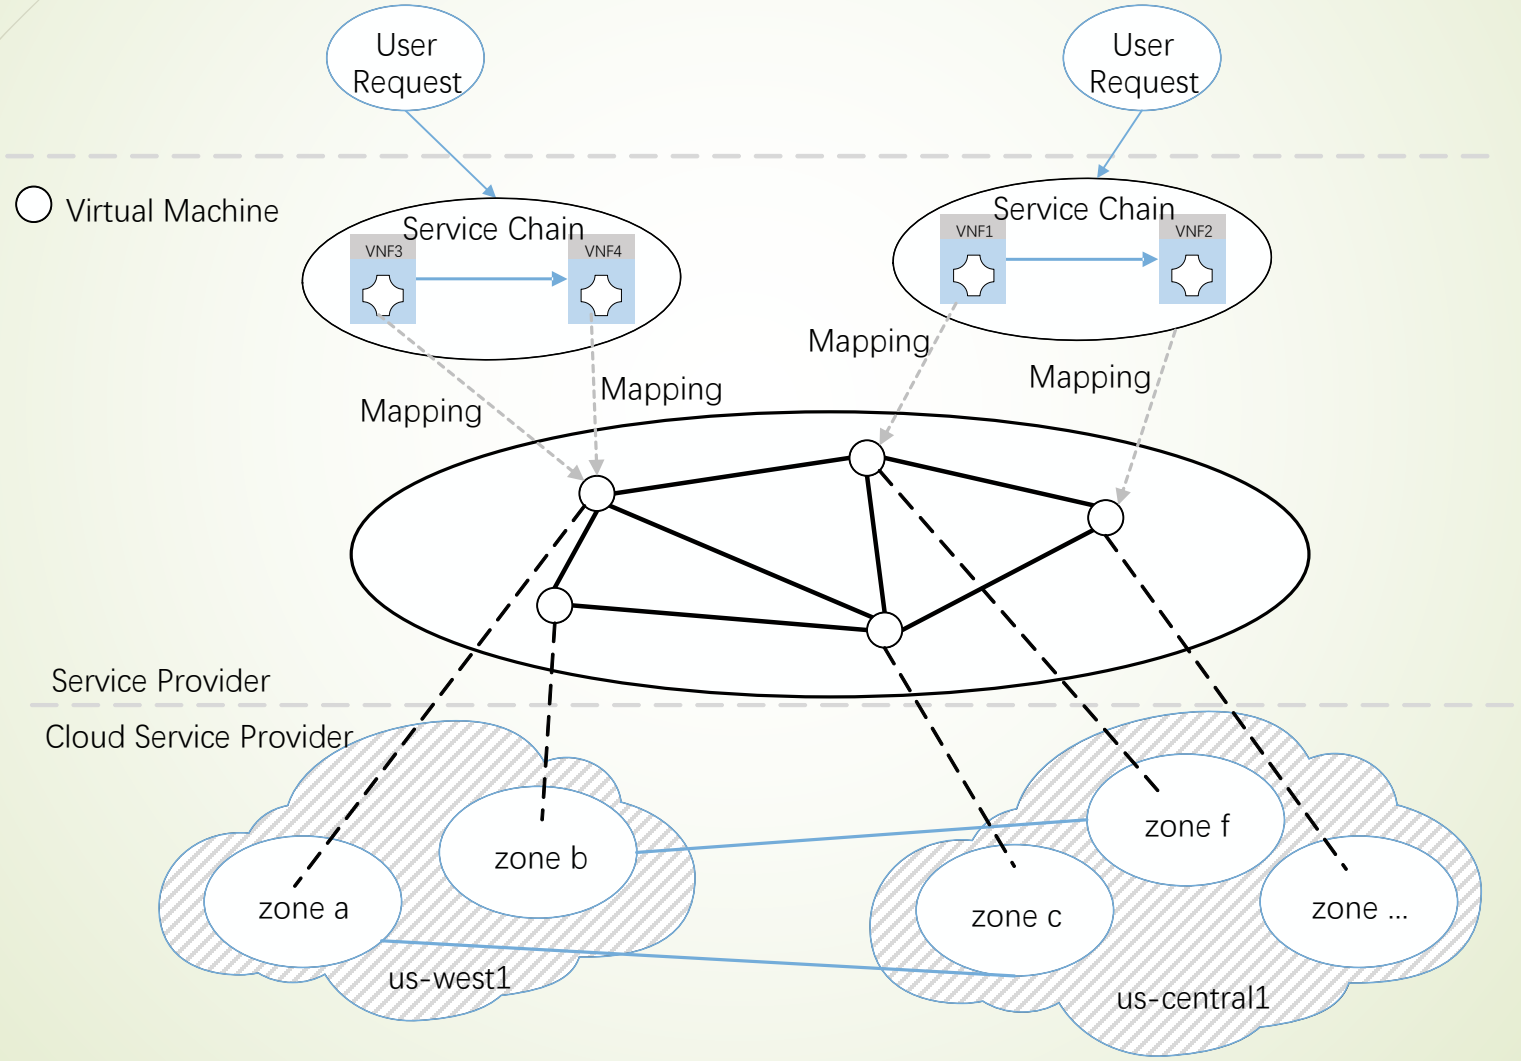

CASE 1

Problem and Challenge

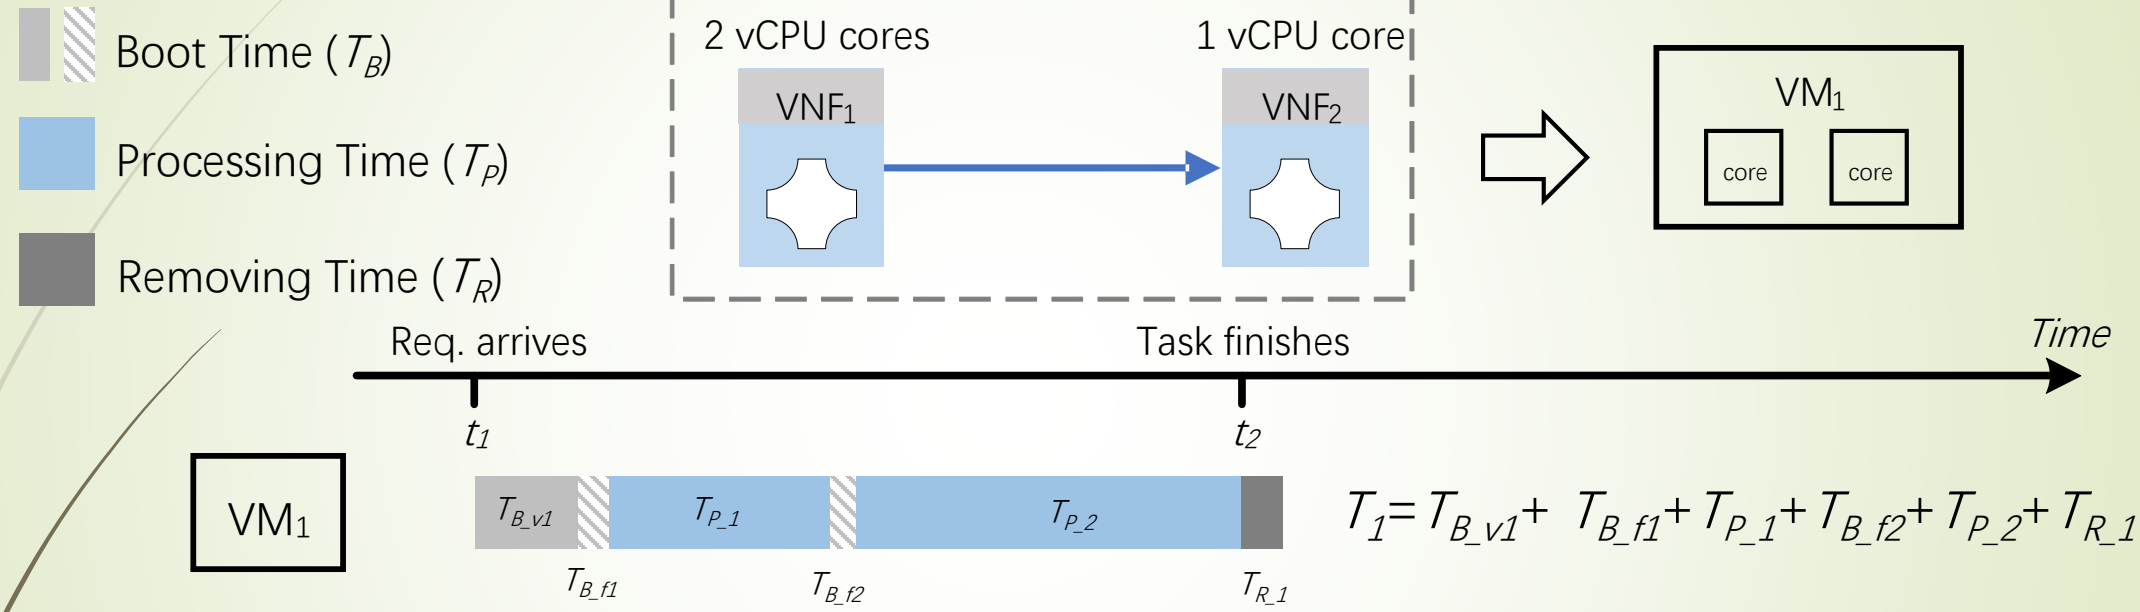

$$T_1 = T_{B\_v1} + T_{B\_f1} + T_{P\_1} + T_{B\_f2} + T_{P\_2} + T_{R\_1}$$

Price ($/vCPU hour)	No. of Used Cores	Uptime (hour)	Total cost
0.033174	2	T_1	$0.066348 * T_1$

CASE 2

Problem and Challenge

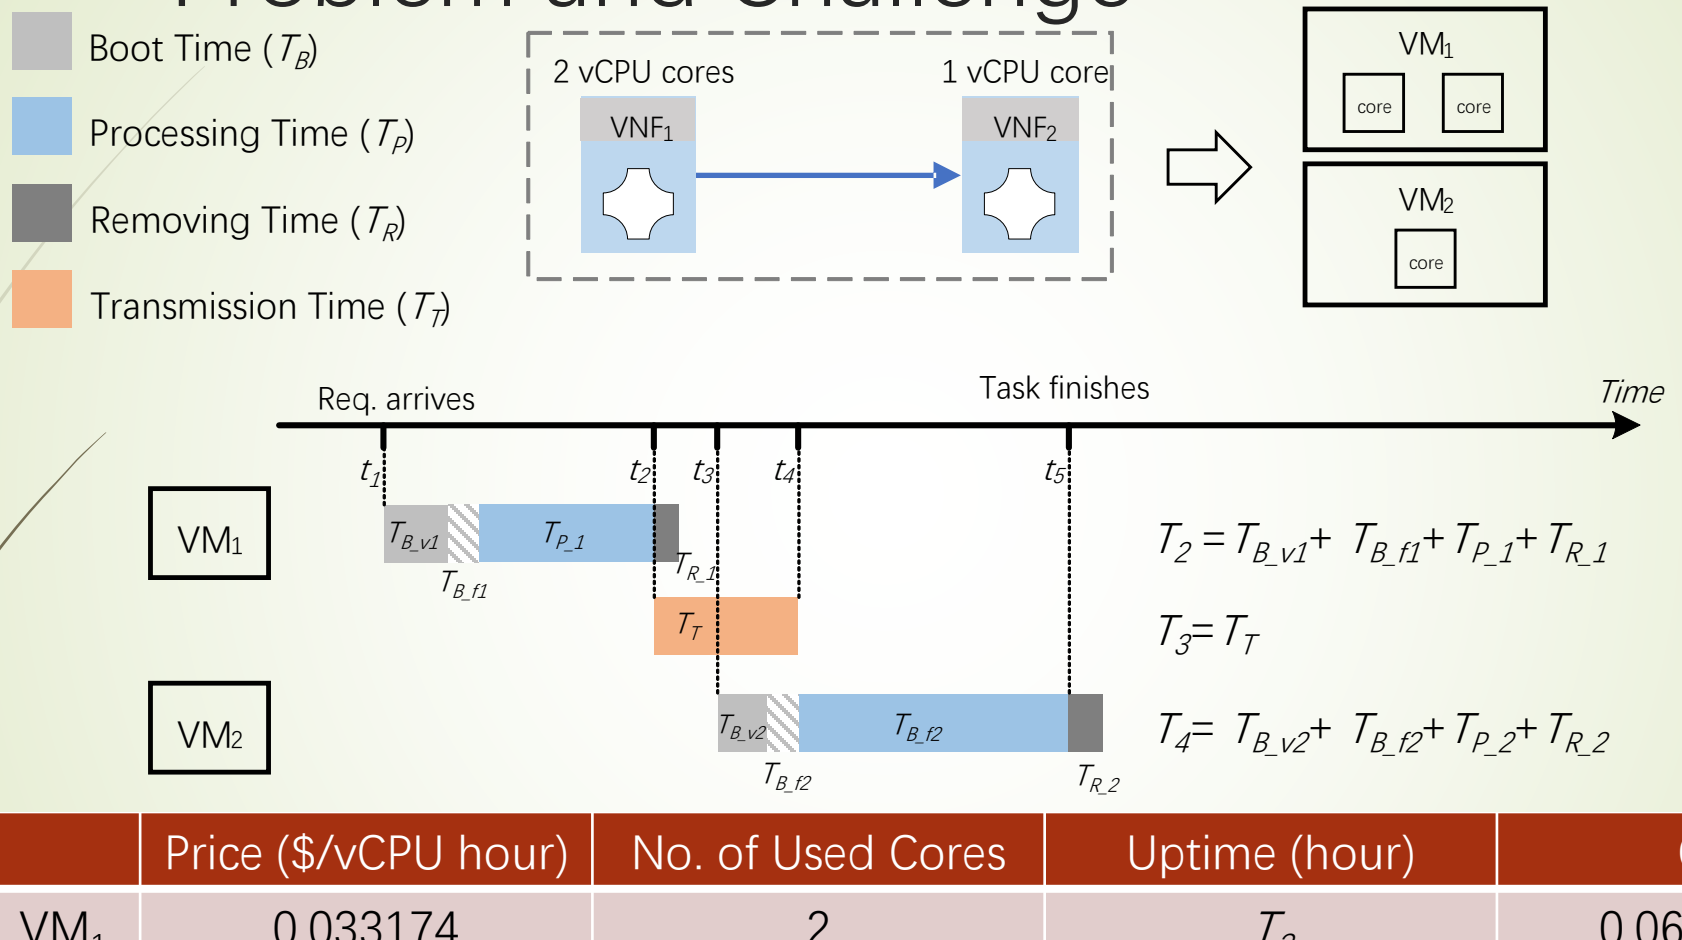

	Price ($/vCPU hour)	No. of Used Cores	Uptime (hour)	Cost
VM_1	0.033174	2	T_2	$0.066348 * T_2$
VM_2	0.033174	1	T_4	$0.033174 * T_4$

Total cost: $0.66348 * T_2 + 0.033174 * T_4 + 0.01 * D$, where D is the size of data in GB

CASE 3

Problem and Challenge

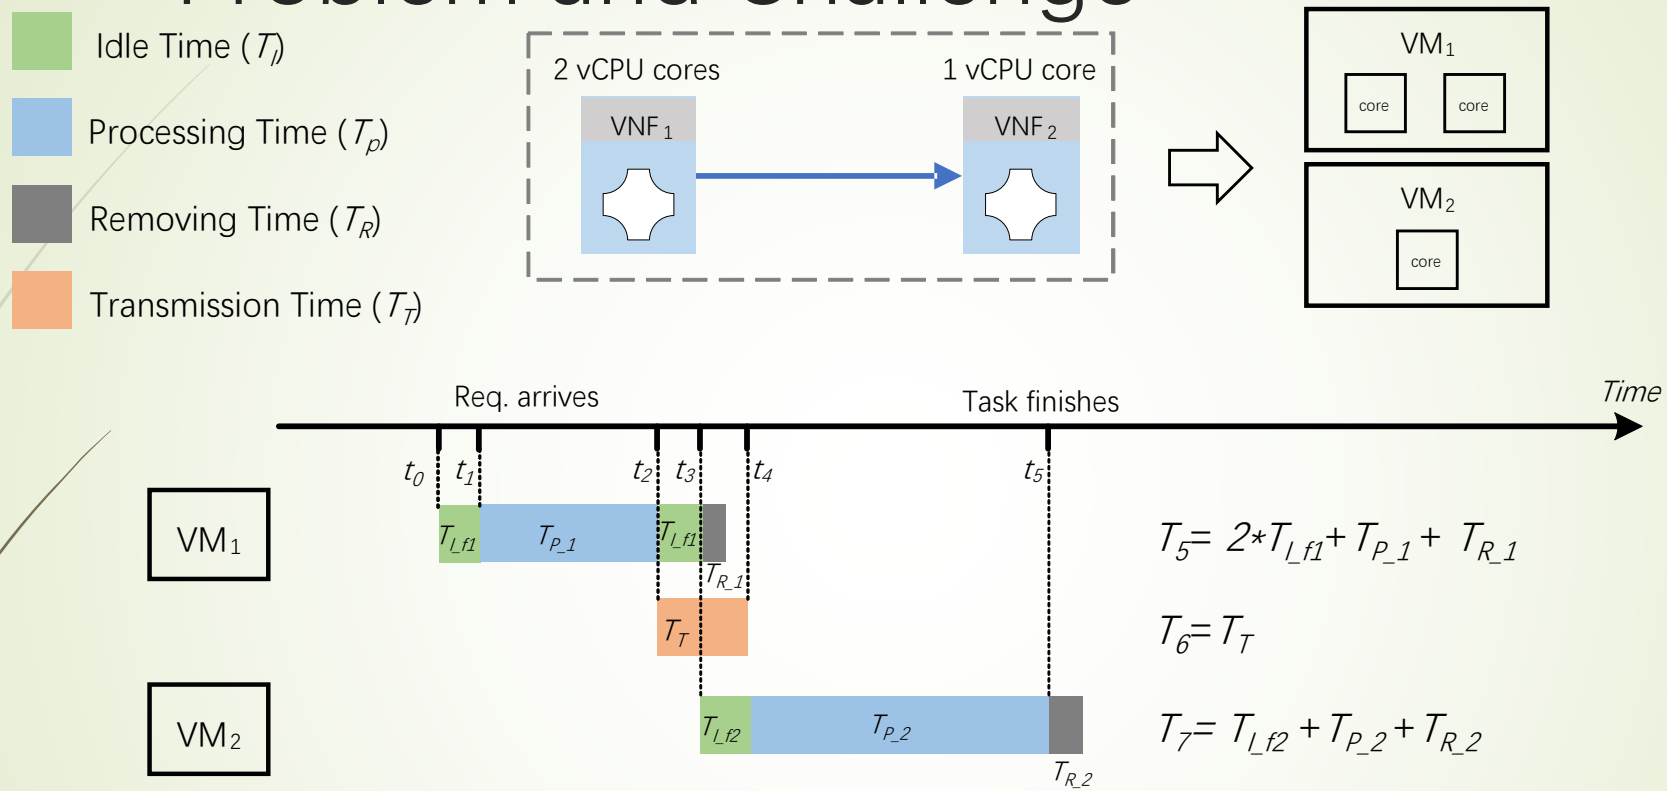

	Price ($/vCPU hour)	No. of Used Cores	Uptime (hour)	Cost
VM_1	0.033174	2	T_5	$0.066348*T_5$
VM_2	0.033174	1	T_7	$0.033174*T_7$

Total cost: $0.66348*T_5+0.033174*T_7+0.01*D$

Problem and Challenge

Assuming the size of data required to be processed is 1 GB

Type	Latency	Time Slots (100ms/slot)
VM boot time	Several seconds [1]	5
VNF installation time	Tens of milliseconds ~ several seconds [1]	1
Transmission latency	Based on the transmission rate*	1/4*10=2.5
Processing time	Based on the capacity of VNF and data size	VNF_1: 10, VNF_2: 15
Instance removing time	Several milliseconds	1
Idle time	Based on the frequency of being used	2 (3)

* The achievable network capacity (egress) on Google Compute Engine is based on the quantity of CPUs your VMs have. Each core is subject to a 2 Gbits/second (Gbps) cap for peak performance [2, 3].

[1] A. Sheoran, et al, "An Empirical Case for Container-driven Fine-grained VNF Resource Flexing," in Proc. 2016 IEEE Conference on Network Function Virtualization and Software Defined Networks, NFV-SDN, 2016.
[2] "https://kinsta.com/blog/google-cloud-vs-aws/"
[3] "https://cloud.google.com/compute/docs/networks-and-firewalls#egress_throughput_caps"

CASE 3

Problem and Challenge

Case	Total Latency (Time Slots)	Total Cost ($)	Influence Factors
1	32	2.189	The number of vCPU cores cannot be adjusted when running*;
2	39.5	1.958	Starting new VMs induces latency and cost.
3	27.5	1.692 (1.858)	Idle state of VMs causes a waste of resource.

* "*You can only change the machine type of a stopped instance and an instance is considered stopped only when the instance is in the TERMINATED state. It is not possible to change the machine type of a running instance.*" (https://cloud.google.com/compute/docs/instances/changing-machine-type-of-stopped-instance)

- The type of the VNF has an effect on whether a VNF is installed into an existed VM or a new VM;
- Distributing the VNFs among different VMs will increase the latency, but may decrease the cost;
- If allowing an instance of VNF to be idle, the time period directly affect the cost.

Problem Statement

- Given:
  - Cloud network topology
  - Resource price (vCPU, network)
  - Set of service chains
  - Set of use requests
- Objective:
  - Minimize the total cost (related to number of VMs like that)
- Constraints:
  - Latency requirement
    - Processing latency ⇔ vCPU core allocated, boot time, and idle time
    - Transmission latency ⇔ link capacity
    - Propagation latency ⇔ distance

Thanks!