

# **K Pre-Post Cloud**

# **Tutorial for the use of**

# **GPGPU instances**

---

RIKEN R-CCS

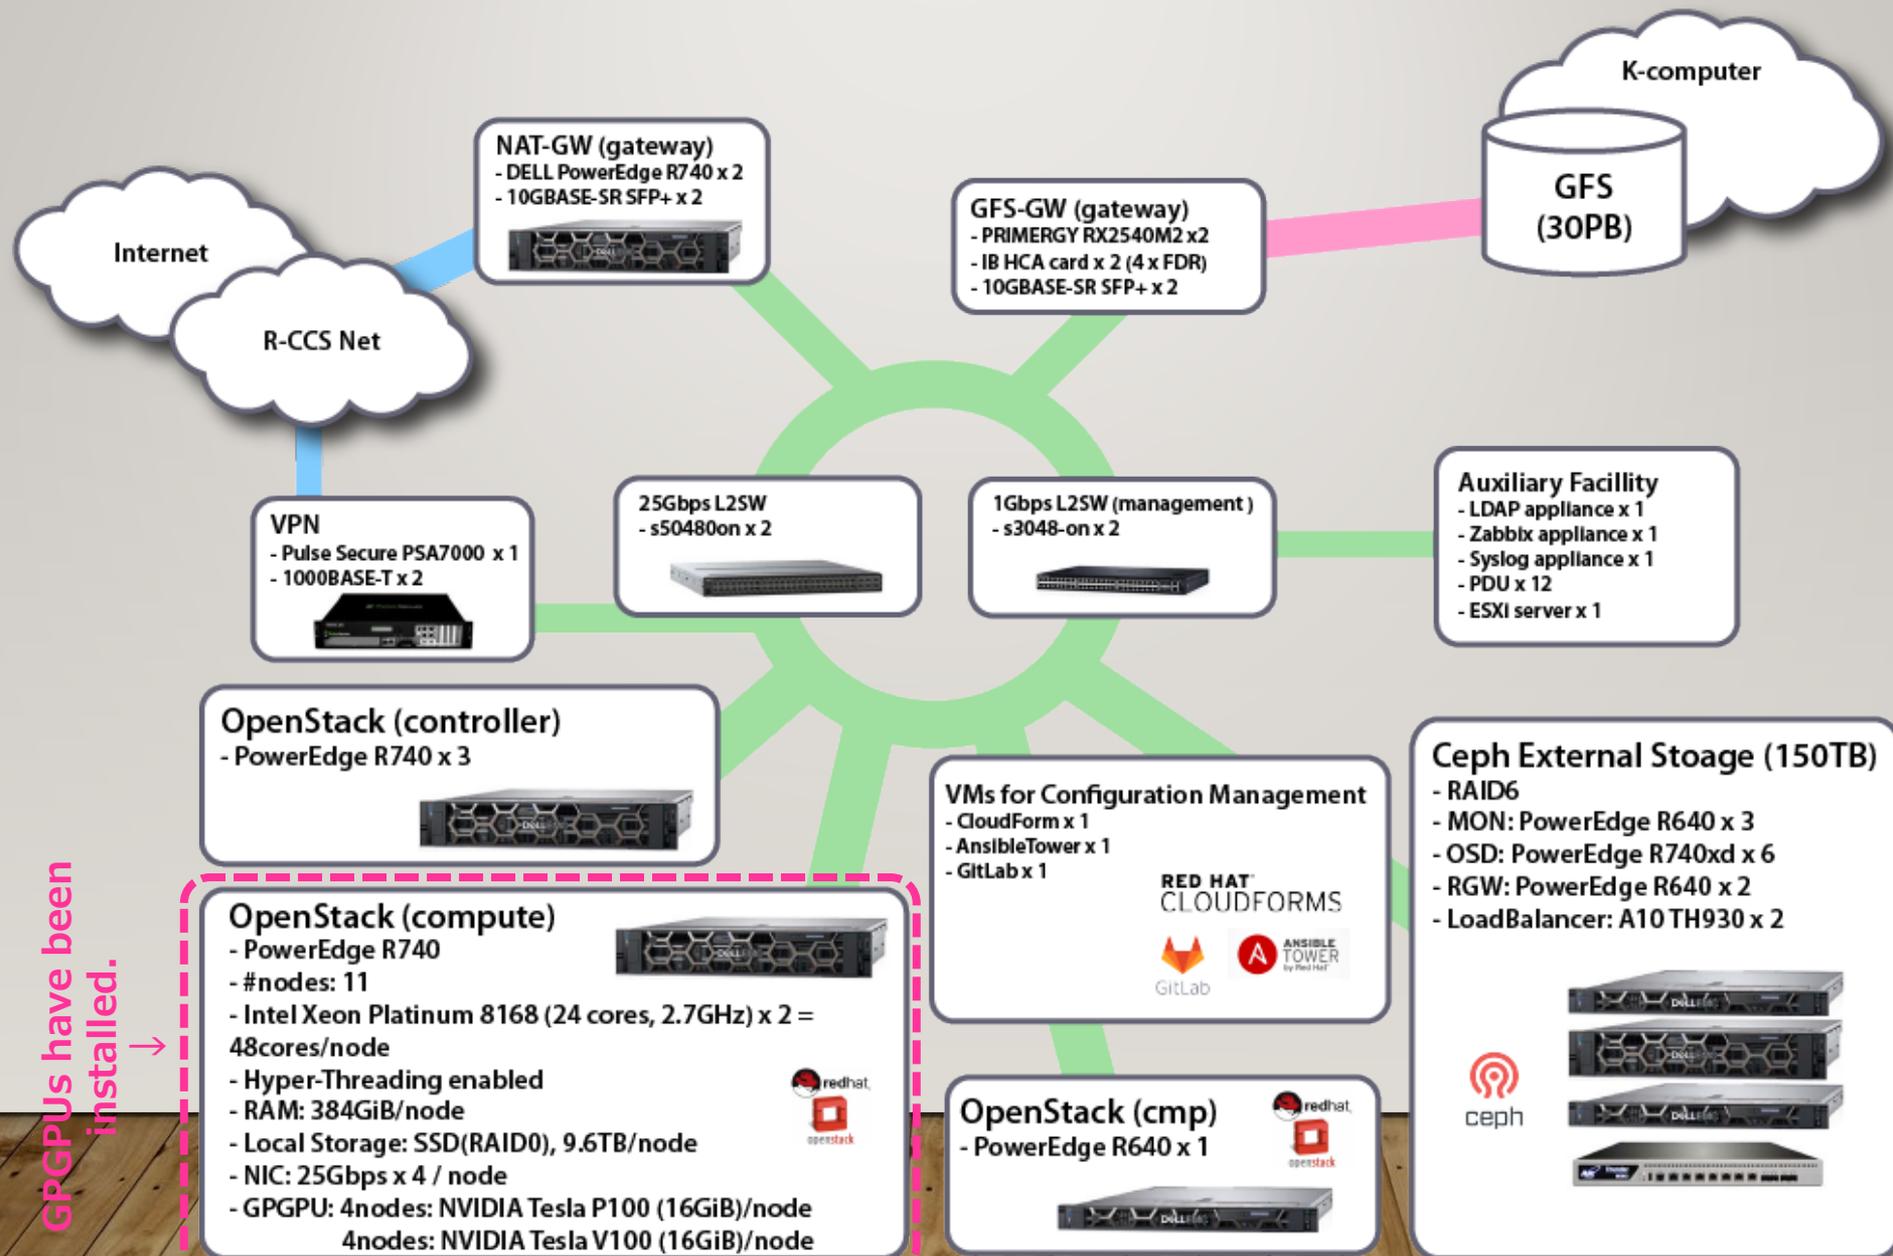
MARCH 29, 2019

# About this Slides

This material provides additional information regarding the use of GPGPU instance. (GPGPUs are installed in March 2019.) and is based on the previously released tutorial named "Tutorial for basic usage."

If you've never seen the tutorial, we recommend referring the tutorial before getting started.

# System Overview



GPGPUs have been installed.

**OpenStack (compute)**

- PowerEdge R740
- #nodes: 11
- Intel Xeon Platinum 8168 (24 cores, 2.7GHz) x 2 = 48cores/node
- Hyper-Threading enabled
- RAM: 384GiB/node
- Local Storage: SSD(RAID0), 9.6TB/node
- NIC: 25Gbps x 4 / node
- GPGPU: 4nodes: NVIDIA Tesla P100 (16GiB)/node  
4nodes: NVIDIA Tesla V100 (16GiB)/node



**Ceph External Storage (150TB)**

- RAID6
- MON: PowerEdge R640 x 3
- OSD: PowerEdge R740xd x 6
- RGW: PowerEdge R640 x 2
- LoadBalancer: A10 TH930 x 2



# Overview of GPGPU installation

In FY2018, we installed 8 GPGPUs to the 8 compute nodes, respectively.

The 8 compute nodes consists of

- 4 compute nodes that have 4 GPGPUs (NVIDIA Tesla P100 (16GiB) x 1), resp, and
- 4 compute nodes that have 4 GPGPUs (NVIDIA Tesla V100 (16GiB) x 1), resp.

Each GPGPU is exclusively assigned to a single instance.

- 8 GPGPU instances can be used simultaneously in the system.
- If 8 instances have been already used, a user's request to create an additional GPGPU instance will be failed.
- Also, the service does not support to share a GPGPU by several instances (e.g., VDI).

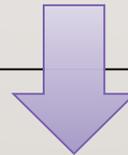
Changes:

- New Availability Zones (gpu-p/v) are added.
- New Flavors (A8.huge.gpu-p/v) are added.

# Availability Zone

Before

nova ← Availability zone (AZ)

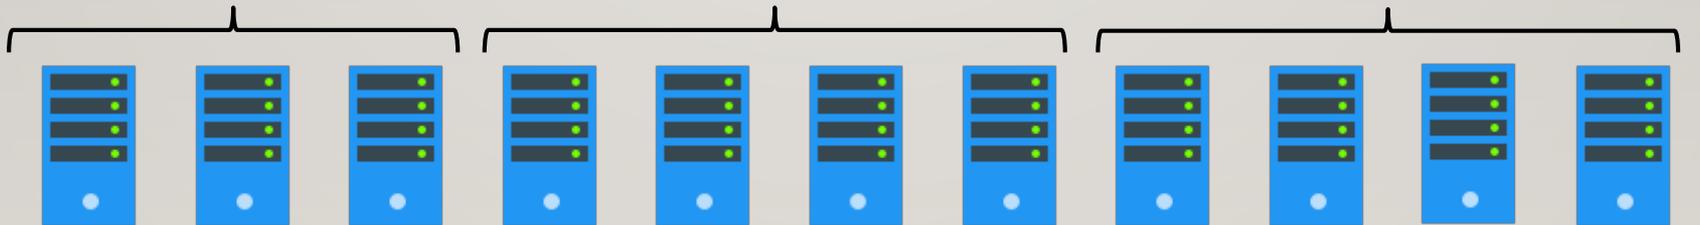


After (from March 11, 2019)

nova

gpu-p (for Tesla P100)

gpu-v (for Tesla V100)

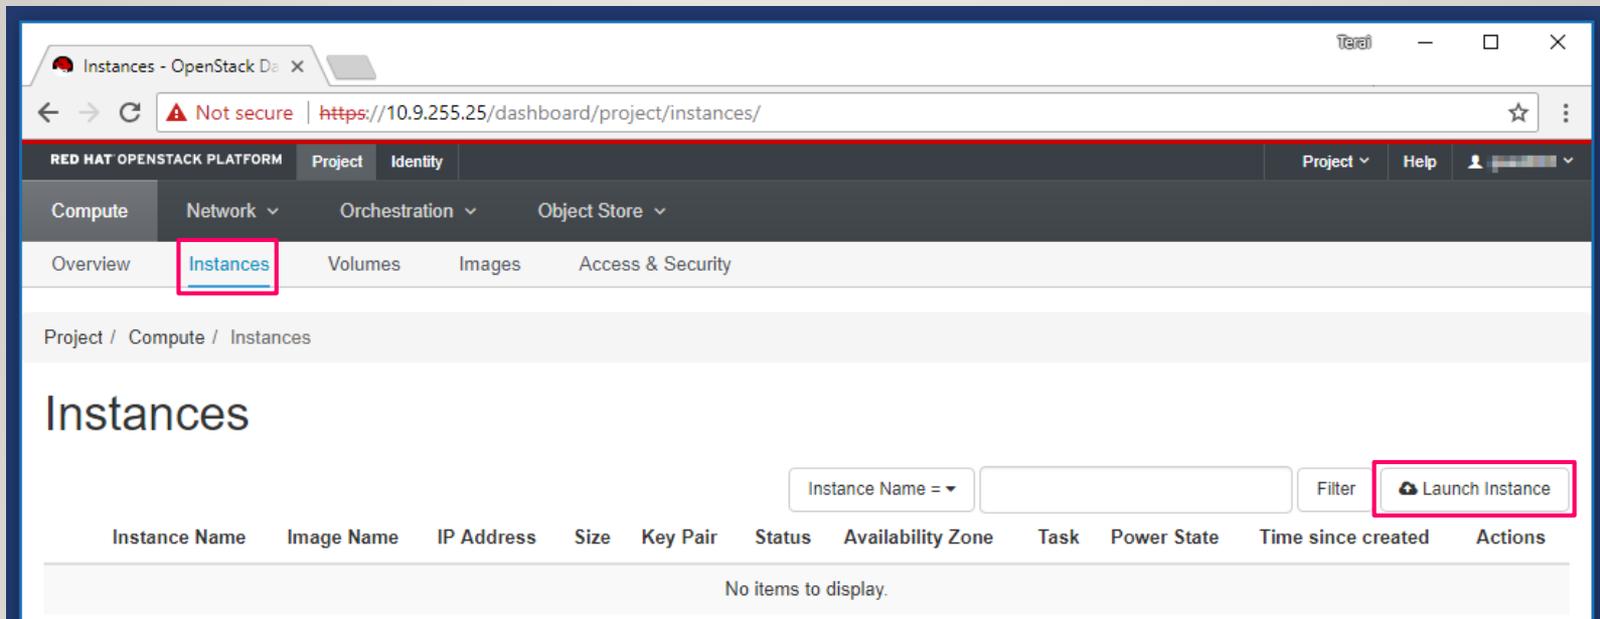


In nova, users can choose several flavors ranging from 1vCPU to 96vCPU as before.

In this slide, we ignore "cmp" availability zone because it's used internally in the system.

# Create a GPGPU Instance (1/6)

- The first step is the completely same as the normal usage without GPGPU node.
- Click [Project] -> [Compute] -> [Instances] on the navigation bar.
- Click the [Launch Instance] button.



The screenshot shows the OpenStack dashboard interface. The browser address bar displays the URL `https://10.9.255.25/dashboard/project/instances/`. The navigation bar includes tabs for 'Project' and 'Identity', and a 'Project' dropdown menu. The main navigation bar has tabs for 'Compute', 'Network', 'Orchestration', and 'Object Store'. Under the 'Compute' tab, there are sub-tabs for 'Overview', 'Instances', 'Volumes', 'Images', and 'Access & Security'. The 'Instances' sub-tab is highlighted with a red box. Below the navigation bar, the breadcrumb 'Project / Compute / Instances' is visible. The main heading is 'Instances'. There is a search filter for 'Instance Name' and a 'Filter' button. The 'Launch Instance' button is highlighted with a red box. Below the buttons, there is a table header with columns: 'Instance Name', 'Image Name', 'IP Address', 'Size', 'Key Pair', 'Status', 'Availability Zone', 'Task', 'Power State', 'Time since created', and 'Actions'. The table content is empty, showing 'No items to display.'

# Create a GPGPU Instance (2/6)

- You can see the wizard dialog to create an instance.
- In the [Details] step of the wizard, input your instance name to the [Instance Name] field.
- Select an availability zone from the [Availability Zone] list that includes
  - **nova** (default): for instance(s) without GPGPU(s),
  - **gpu-p**: for Tesla P100, and
  - **gpu-v**: for Tesla V100.

Launch Instance

Please provide the initial hostname for the instance, the availability zone where it will be deployed, and the instance count. Increase the Count to create multiple instances with the same settings.

**Instance Name \***

**Availability Zone**

Any Availability Zone

**Count \***

1

Total Instances (10 Max)

10%

0 Current Usage  
1 Added  
9 Remaining

**nova (for non-GPGPU instance)**  
**gpu-p (for Tesla P100)**  
**gpu-v (for Tesla V100)**

Cancel < Back Next > Launch Instance

# Create a GPGPU Instance (3/6)

- Select the [Image] item from the [Select Boot Source] pull-down menu.
- Select the [No] button in the [Create New Volume] switch.
- Add an OS image in the [Available] list.
  - At the end of FY2018, Ubuntu18.04.2\_LTS(GPU-node-20190319) is available to create a GPGPU instance.

Launch Instance

Instance source is the template used to create an instance. You can use a snapshot of an existing instance, an image, or a volume (if enabled). You can also choose to use persistent storage by creating a new volume.

**Select Boot Source**  
Image

**Create New Volume**  
Yes No ← "No" is recommended.

**Allocated**

Name	Updated	Size	Type	Visibility
> Ubuntu18.04.2_LTS(GPU-node-20190319)	3/19/19 9:33 PM	8.82 GB	qcow2	Public

**Available 6** Select one

Click here for filters.

Name	Updated	Size	Type	Visibility
> Windows_server_2012_r2_eval	3/8/19 12:42 PM	11.18 GB	qcow2	Private
> Ubuntu18.04.2_LTS	3/7/19 2:28 PM	327.63 MB	qcow2	Public
> Ubuntu18.04.1_LTS	10/18/18 4:04 PM	323.63 MB	qcow2	Public
> CentOS7	5/16/18 5:58 PM	832.25 MB	qcow2	Public
> Ubuntu18.04_LTS	5/16/18 4:39 PM	320.94 MB	qcow2	Public
> Ubuntu16.04_LTS	4/28/18 3:05 AM	276.56 MB	qcow2	Public

Cancel < Back Next > Launch Instance

Select an image

# Create a GPGPU Instance (4/6)

- If you choose the [Yes] button in the [Create New Volume] switch, you must specify more than 40GiB in the volume size.

Launch Instance

Instance source is the template used to create an instance. You can use a snapshot of an existing instance, an image, or a volume (if enabled). You can also choose to use persistent storage by creating a new volume.

**Select Boot Source**

Image

**Create New Volume**

Yes No

**Volume Size (GB)**

64 ← more than 40GiB

**Delete Volume on Instance Delete**

Yes No ← "Yes" is recommended.

**Device Name**

vda

**Allocated**

Name	Updated	Size	Type	Visibility	
> Ubuntu18.04.2_LTS(GPU-node-20190319)	3/19/19 9:33 PM	8.82 GB	qcow2	Public	-

Available 6 Select one

Click here for filters.

Name	Updated	Size	Type	Visibility	
> Windows_server_2012_r2_eval	3/8/19 12:42 PM	11.18 GB	qcow2	Private	+
> Ubuntu18.04.2_LTS	3/7/19 2:28 PM	327.63 MB	qcow2	Public	+
> Ubuntu18.04.1_LTS	10/18/18 4:04 PM	323.63 MB	qcow2	Public	+
> CentOS7	5/16/18 5:58 PM	832.25 MB	qcow2	Public	+
> Ubuntu18.04_LTS	5/16/18 4:39 PM	320.94 MB	qcow2	Public	+
> Ubuntu16.04_LTS	4/28/18 3:05 AM	276.56 MB	qcow2	Public	+

Cancel Back Next Launch Instance

Select an image

# Create a GPGPU Instance (5/6)

- Add a flavor from the [Available] list.
- We newly provide GPGPU flavors (A8.huge.gpu-p/v) that consumes the whole of a compute node as well as A8.huge.
- For P100/V100, select A8.huge.gpu-p/v, resp.
- If you select other flavors for instance(s) without GPGPU(s), the request will be failed.

Launch Instance

Details \*  
Source  
Flavor \*  
Networks \*  
Network Ports  
Security Groups  
Key Pair  
Configuration  
Server Groups  
Scheduler Hints  
Metadata

Flavors manage the sizing for the compute, memory and storage capacity of the instance.

Allocated

Name	VCPUS	RAM	Total Disk	Root Disk	Ephemeral Disk	Public
Select an item from Available items below						

Available 97

Search: gpu

← To quickly find proper flavors, we recommend to input "gpu" in this input form.

Name	VCPUS	RAM	Total Disk	Root Disk	Ephemeral Disk	Public
A8.huge.gpu-p	96	320 GB	8 TB	8 TB	0 GB	Yes
A8.huge.gpu-v	96	320 GB	8 TB	8 TB	0 GB	Yes

For V100 →  
For P100 →

Cancel < Back Next > Launch Instance

# Create a GPGPU Instance (6/6)

- The rest of the steps are the same with the usage for instance(s) without GPGPU(s).
  - Add an internal network.
  - Add security group(s).
  - Add key pair(s).
  - Click the [Launch Instance] button.
  - After about 3 minutes, the instance using root disk will be launched.
  - Assign the Floating IP address to the instance.
  - The instance is ready to access using SSH.

# Image for GPGPU instance

- Currently (in March 2019), we provide a single image (Ubuntu 18.04.2LTS based) for GPGPU instance.
- The file name depends on the updated date.
- As of March 28, 2019, this image includes
  - NVIDIA Driver version 410.48,
  - CUDA Toolkit release 10.0,
  - Docker (Engine 18.09.3, Client 18.09.3), and
  - NVIDIA-Docker 2.0.3.

# TIPS

- If you find an error when the system spawns a new GPGPU instance, please check the following points.
  - Check the combination of the availability zone and the flavor you chose.
    - gpu-p + A8.huge.gpu-p
    - gpu-v + A8.huge.gpu-v
  - Check the quota of your project and unallocated resources that has sufficient space to launch your GPGPU instance.
    - In default settings, a single project can create a few GPGPU instances.
    - If you need to expand the quota, please contact us.
  - There is no available resource to launch GPGPU instance in the system.
    - If the system has already launched 8 GPGPU instances (including reserved/error instances), your request to create an additional GPGPU instance will be failed.
    - In this situation, it's difficult to sort out the problem by a user, please contact us.
    - Also, an error instance remains to be reserved a GPGPU node. Therefore, please release the instance with the error.