

IBM i Disaster Recovery with IBM Power Virtual Server

An IBM Systems Lab Services Tutorial

A dark teal banner with a background of faint, glowing binary code (0s and 1s) and abstract geometric patterns. The text is white and centered.

IBM Systems Lab Services

Infrastructure services to help you build the foundation of a smart enterprise.

Kyle Bliss

Vess Natchev

ibmsls@us.ibm.com

TABLE OF CONTENTS

CHAPTER 1: SOLUTION OVERVIEW.....	1
Introduction.....	1
Use Cases	1
PowerHA Geographic Mirroring	1
Other Logical Replication Options.....	2
Solution Components and Requirements	2
Components.....	2
Requirements.....	3
Diagrams.....	6
Solution Diagram	41
CHAPTER 2: IMPLEMENTATION	42
PowerHA Geographic Mirroring.....	42
CHAPTER 3: TROUBLESHOOTING.....	54
CHAPTER 4: ADDITIONAL RESOURCES	55

Chapter 1: Solution Overview

Introduction

Uptime is a key client expectation for IBM i workloads. Across geographic locations, this is accomplished with a disaster recovery (DR) solution. [IBM Power Virtual Server](#) (PowerVS) meets that requirement by enabling clients to leverage DR solutions between two IBM i Virtual Server Instances (VSIs) in separate IBM Cloud datacenters.

An important characteristic of DR solutions for PowerVS is that they are based on **logical or operating system-level replication**. Many Power Systems clients today use storage-based replication for DR, which is not an option with PowerVS.

Replication solutions between two datacenters always involve prerequisite network configuration between them to allow the necessary data flow to occur securely. This also applies to DR with PowerVS, which requires specific networking steps in IBM Cloud before implementing the replication software itself.

This tutorial will provide step-by-step instructions to accomplish both phases of configuring DR for IBM i workloads in PowerVS:

1. Performing the required network configuration.
2. Implementing the DR solution itself.

Use Cases

PowerHA Geographic Mirroring

In this case we will demonstrate how to implement PowerHA Geographic Mirroring in IBM i, which provides DR by using operating

system (OS) clustering and replication. **Note that this solution requires that the IBM i VSI and client application(s) use Independent Auxiliary Storage Pools (IASPs).** If the IBM i VSI and application(s) use only *SYSBAS, this DR option will not work.

Other Logical Replication Options

For IBM i VSIs and application(s) that use only *SYSBAS, other non-IBM logical replication options exist, but they are beyond the scope of this tutorial.

Solution Components and Requirements

Components

This solution uses the following components

1. *Open an IBM Cloud account*
2. *Create two Power PowerVS location Services and a private subnet in each PowerVS location.*
3. *Provision IBM i VSIs in each PowerVS location*
 - a. *A "production" IBM i cloud instance with an Independent ASP (IASP) that has been IASP-enabled (i.e. All changes/modifications allowing the IASP to function in a working environment should be completed before Geographic Mirroring is set up for a DR solution.)*
 - b. *A "DR" IBM i cloud instance with non-configured disks to be used for Geographic Mirroring. It is highly recommended that the number, type and capacity of disks match that of the production IASP.*
4. *Order Direct Link Connect Classic to connect each PowerVS location to IBM Cloud*
5. *Order two Vyatta Gateways one in each datacenter to allow for PowerVS location-to-location communication*

6. *Request a Generic Routing Encapsulation (GRE) tunnel to be provisioned at each PowerVS location.*
7. *Configure three GRE tunnels in the Vyatta Gateways. Two to connect Vyatta Gateway to the PowerVS location GRE tunnels created in Step 6 above and one across Vyatta Gateways to connect Vyatta-to-Vyatta. This will allow end-to-end PowerVS location-to-location communication for the VSIs in the PowerVS locations and to the IBM Cloud VSIs and other services such as Cloud Object Storage (COS) (if used).*
8. *Configure a Reverse-proxy Centos VSI to allow access to Private Cloud Object Storage endpoint from PowerVS location*

Requirements

Open an IBM Cloud account

Login to <https://cloud.ibm.com> and follow the procedure to open an Internal to external account.

For internal accounts, you can use your IBM intranet ID and password. For external accounts you will need to provide a billing source such as a credit card.

Create PowerVS location Service and Subnet(s)

All Power VSIs are provisioned in what is called a PowerVS location. This is a separate datacenter adjacent to IBM Cloud datacenters. In order to setup your PowerVS location, you will setup a PowerVS location service in the IBM Cloud UI. The PowerVS location service is a service within IBM Cloud which allows you to provision IBM i VSIs. There is a limit of one PowerVS location service per datacenter in IBM Cloud. In our scenario we have created two PowerVS locations, one is Toronto and one in London datacenters.

Prior to provisioning Power VSI in the PowerVS location, you will need to create at least one subnet. You can have as many subnets as you require in each PowerVS location service on which you can provision your Power VSIs. Additional subnets beyond the initial one can be added later, after the VSI's are created.

Provision AIX and IBM i VSIs in each PowerVS location

In each PowerVS location service you can create IBM i VSIs. The details are provided in the next section.

Order Direct Link Connect Classic to connect PowerVS location to IBM Cloud

You will need to order Direct Link (DL) Connect Classic to allow your Power VSIs to communication with Linux/Window VSIs in IBM Cloud and also with all other IBM Cloud services such as VMWare VMs, and Cloud Object Storage (COS). Ordering a DL may take 1-2 weeks to complete. There is no charge for this service as of June 2020.

Order two Vyatta Gateways, one in each datacenter

In order to setup communication between the two PowerVS location datacenters, you will need to use Generic Routing Encapsulation (GRE) tunnels. GRE tunnels are provisioned on Vyatta Gateways so you will need to order one Vyatta Gateway in each PowerVS location.

The example here involves ordering one Vyatta in LON06 and the other in TOR01 datacenters where the PowerVS locations exist.

Request a Generic Routing Encapsulation (GRE) tunnel to be provisioned at each PowerVS location

You will need to open a support ticket with Power Systems and request that a GRE tunnel be provisioned in each PowerVS location. They will provision their end of the GRE tunnel and send you the information so you can continue and provision your end on the Vyatta Gateways. You will need to provide the subnets information in each PowerVS location in the ticket.

Configure three GRE tunnels in the Vyatta Gateways

We used the following link to configure the GRE.

<https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-configuring-power>

After the support team finishes configuring the GRE tunnel, you will need to configure your end of the GRE tunnel on the two Vyatta Gateways.

You will need three GRE tunnels

- 1. GRE tunnel on Vyatta to terminate the PowerVS location GRE in LON06 (or "datacenter 1")*
- 2. GRE tunnel on Vyatta to terminate the PowerVS location GRE in TOR01 (or "datacenter 2")*
- 3. GRE tunnel across the two Vyatta gateways. One on each side.*

Configure a Reverse-proxy Centos VSI

In this section we will discuss the procedure to configure a reverse proxy to allow access to private COS endpoint.

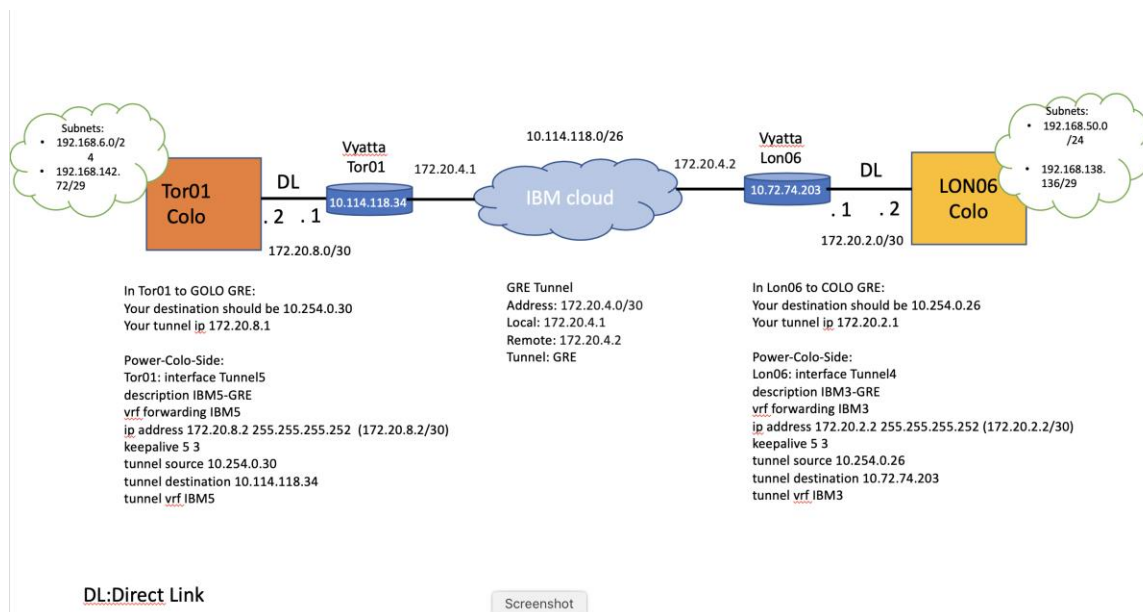
All access to COS from Power VSI is via this reverse proxy.

You will access it via `https://<reverse_proxy_ip>`.

You will need to provision a Centos or Red Hat VSI in IBM cloud to configure at reverse proxy. This VSI must have public access. After configuration, the public access can be disabled.

Diagrams

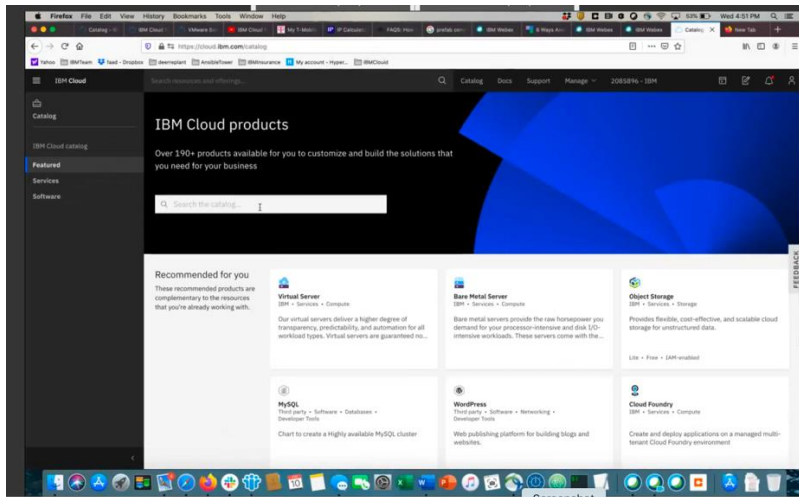
The overall architecture of our deployment is shown in Figure 1.



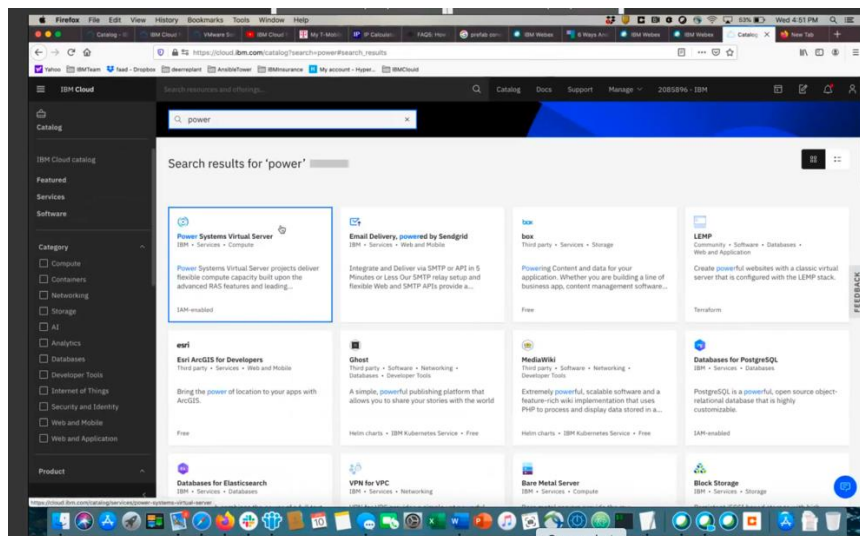
Create PowerVS location Services and Subnet(s)

You will need an IBM Cloud account to start this process.

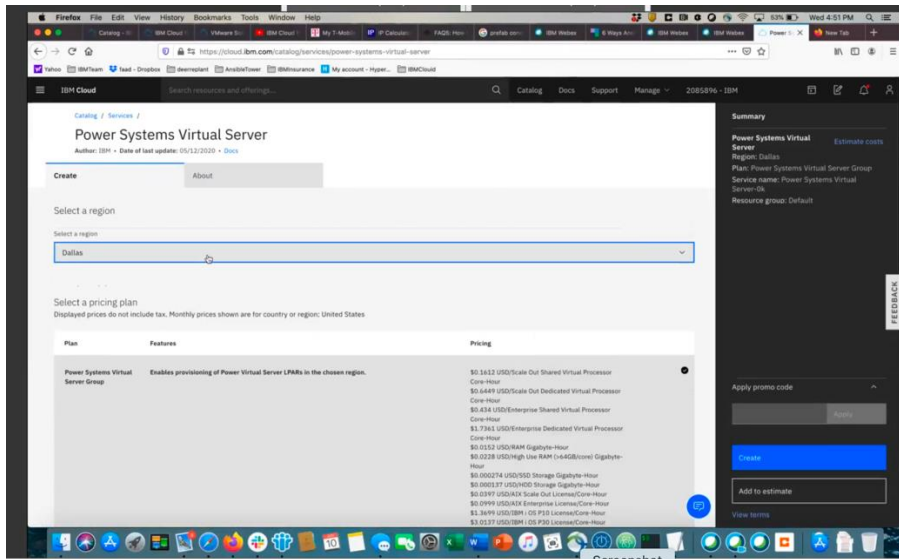
Go to the main IBM Cloud UI page and click on "Catalog" on upper right side of the UI.



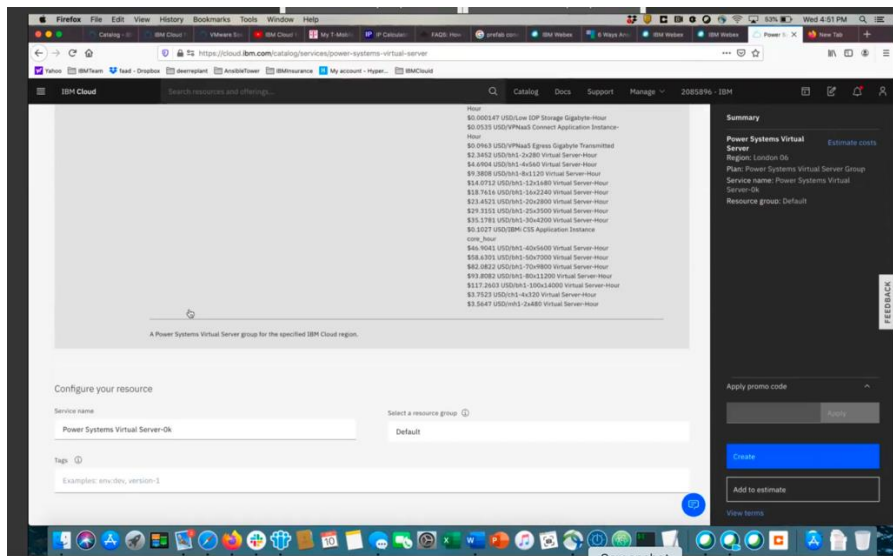
Search for "Power"



Select "Power System Virtual Servers".

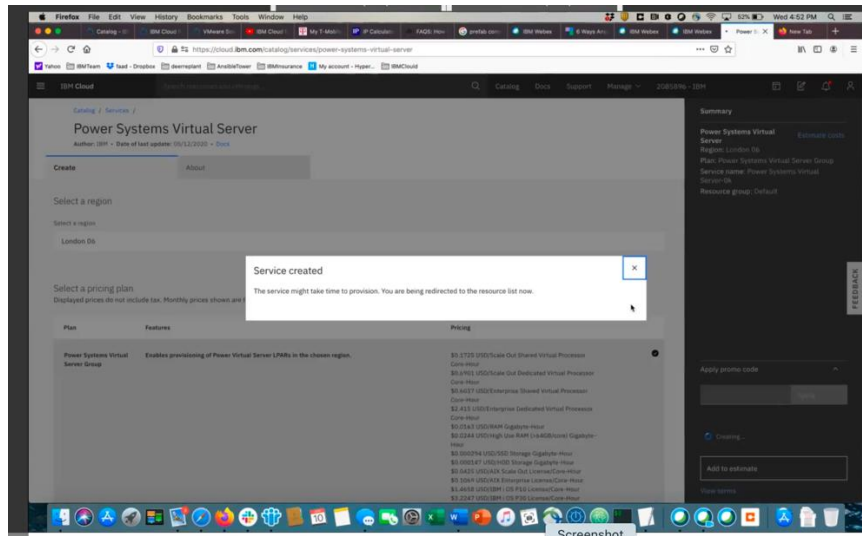


Under Select Region, choose your region. You are limited to only one service per region.



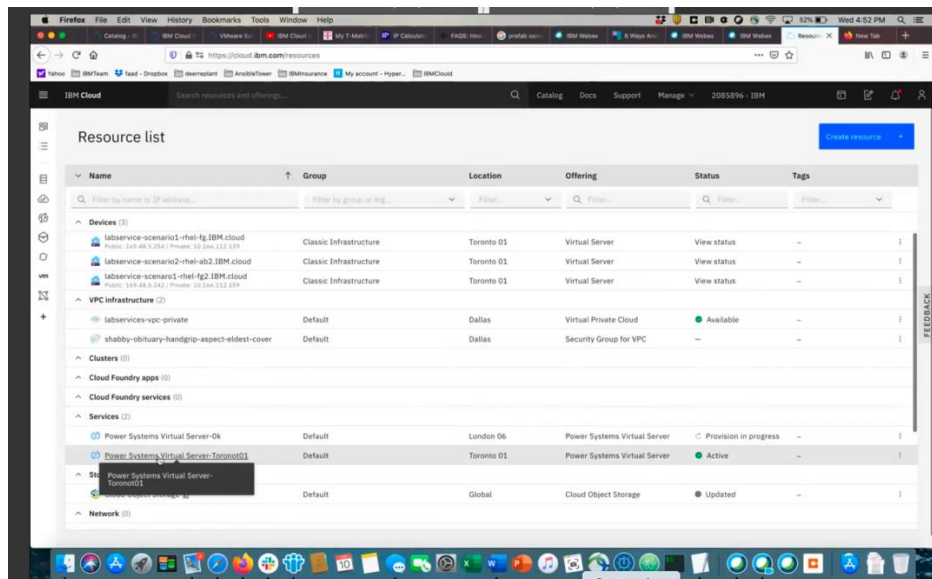
Select a "Service Name" or chose default name provided.

Then press "Create"

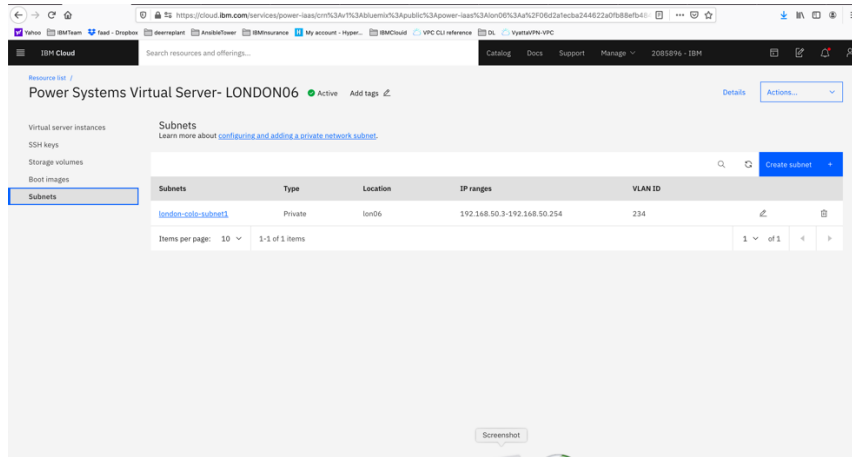


Your PowerVS location service will now appear under the Services tab.

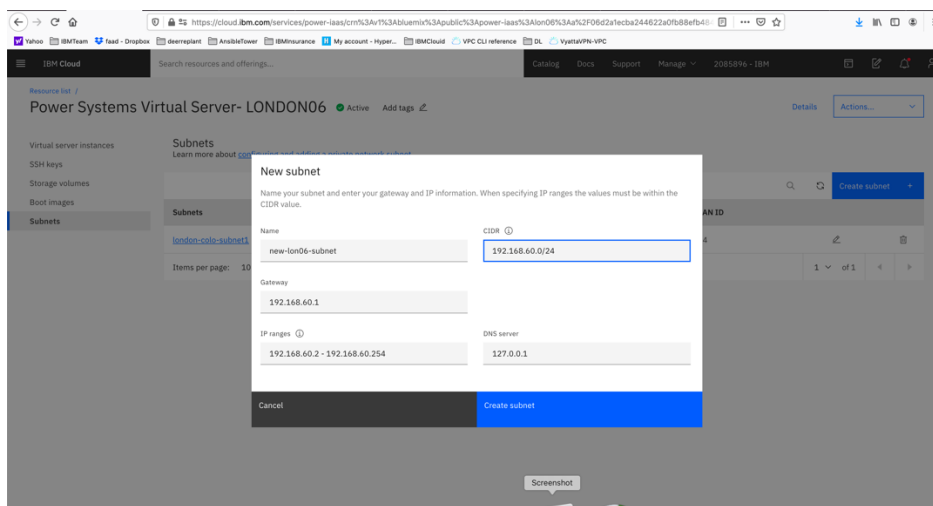
You will repeat this process to create a second PowerVS location service. In this example, there are two PowerVS location services, one in London and one in Toronto.



Next you will need to click on the PowerVS location Service you created and provision a subnet to be used by your Power VSI servers.



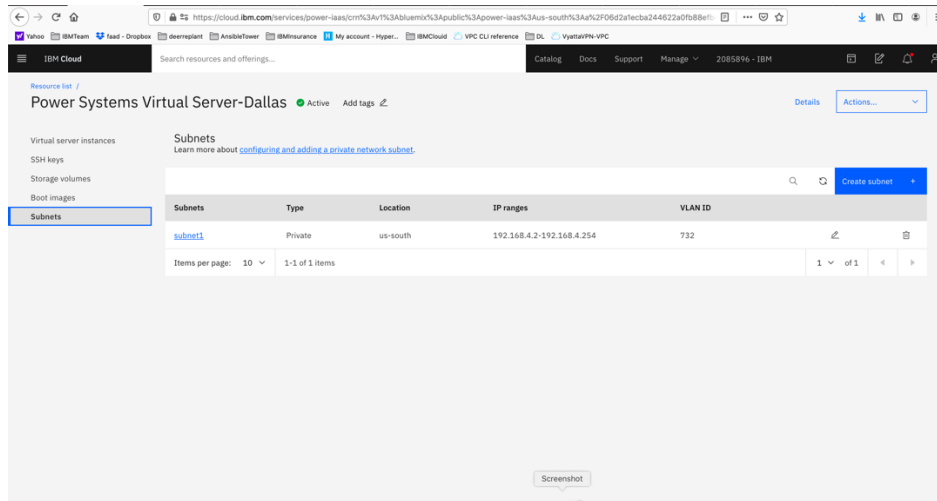
Choose “Subnets” from the menu on the left.



Provide the following information:

1. name for your subnet
2. CIDR range. This can be any private IP subnet ranges. For example, 192.168.5.0/24. You may choose /21 to /30 based on how many IPs you will require. You may use your own private CIDR if you wish.
3. The rest of the fields will be automatically populated based on the CIDR you provided.

Press “Create Subnet”



There should be a VLAN ID associated with the subnet.

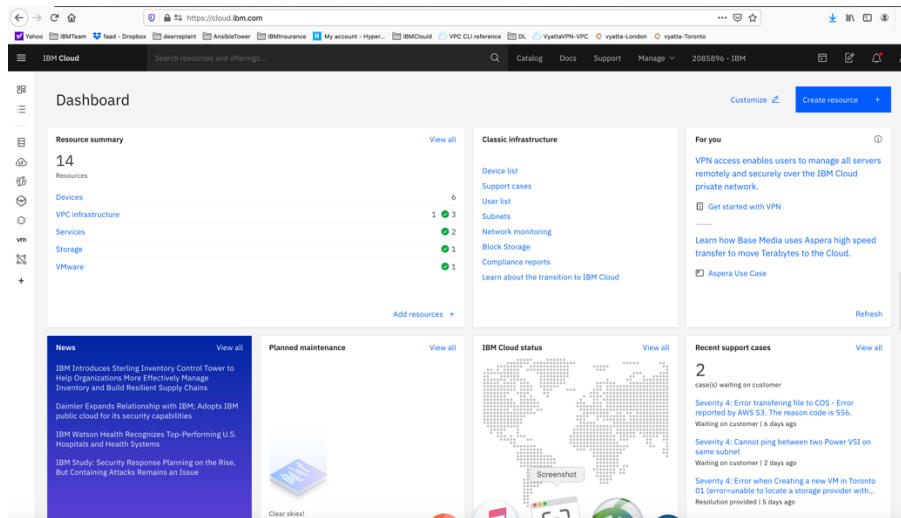
At this point, you will need to open a Support Ticket with Power System to request that the subnet be configured to allow local communication between any Power VSI you create in this PowerVS location service. Provide your PowerVS location service location, and your subnet in the ticket.

Without this step, the Power VSI you create will not be able to ping between each other even if they are on same subnet in the same PowerVS location.

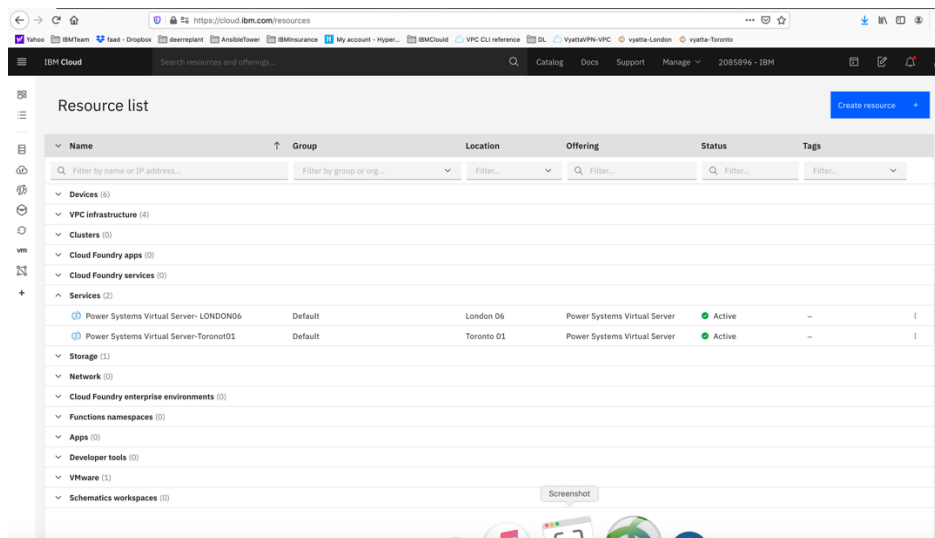
Provision AIX or IBM i VSIs in each PowerVS location

The procedure is similar for both AIX and IBM i VSI provisioning. Here is a procedure to create an AIX 7.2 VSI. The cost shown are monthly costs, but you are being charged hourly.

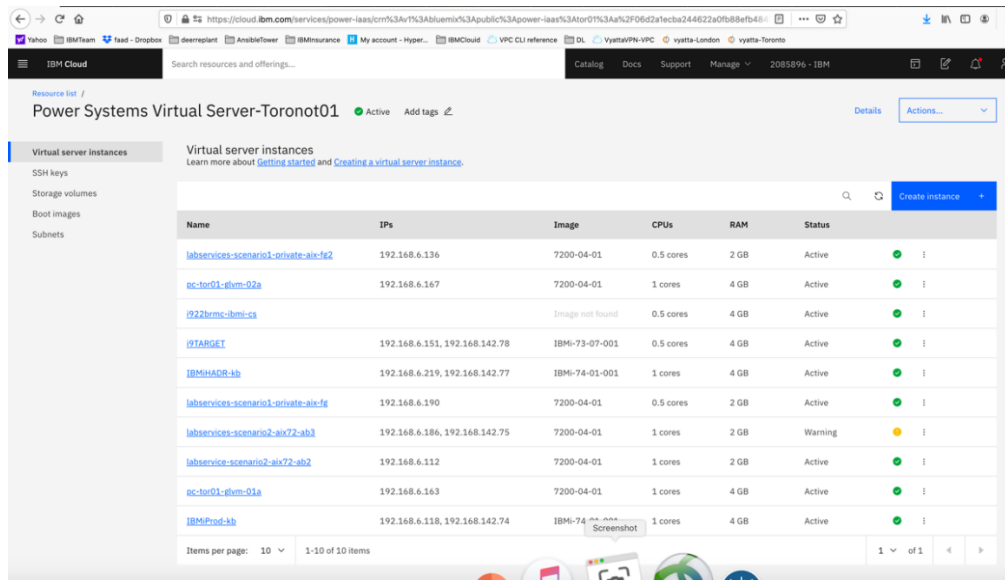
Go to the IBM Cloud Catalog and press the “IBM Cloud” on top left side of the UI.



Choose "Services" from the list shown.

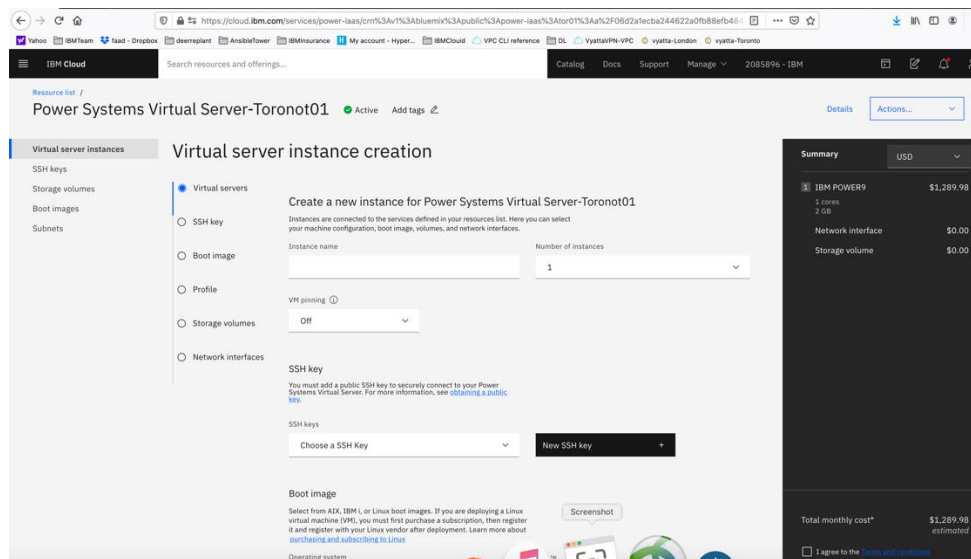


Click on the service for datacenter in which you have created a PowerVS location power service. In this case we will choose Toront01 PowerVS location service.



Since we have already provisioned several VSIs, we see the list shown above. If you are creating VSIs for the first time, your list will be empty.

Press “Create Instance” on upper right-hand side.

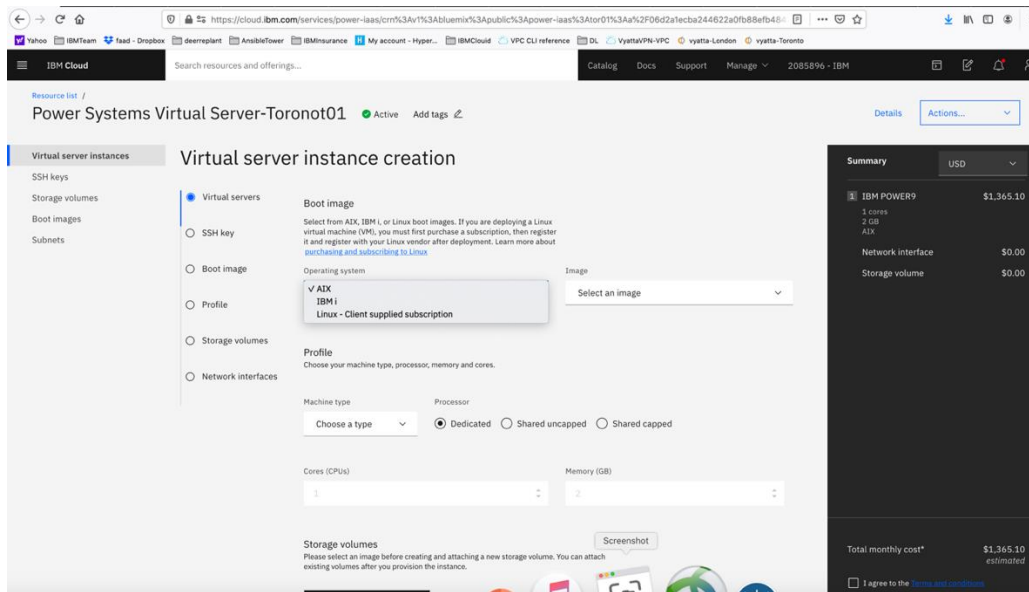


This is where you provision AIX or IBM i VSIs.

Choose a name for your VSI, i.e., AIX-72-Tor01 and select how many VSIs you need to configure. The names of the VSI will be appended with a "-1", "-2" etc. if you select more than one VSI. Note that the IBM i will give the LPAR a system name consisting of the first 8 characters, and it is best to use only alphanumeric characters for IBM i naming.

You may leave VM pruning and SSH key as-is since the VSIs will have no passwords when you create them for the first time. You will need to create a password via the OS command. This is AIX-specific and those two options are not related to IBM i instances. The O/S and DST passwords for user qsecofr will default to QSECOFR and will be changed later (discussed in the section on configuration).

Scroll down to choose other options.

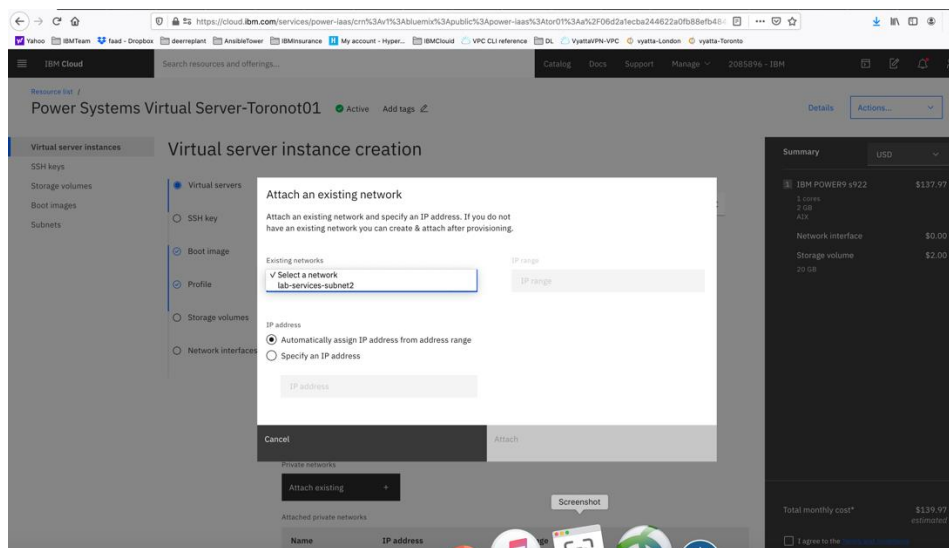


Here you will choose the following options:

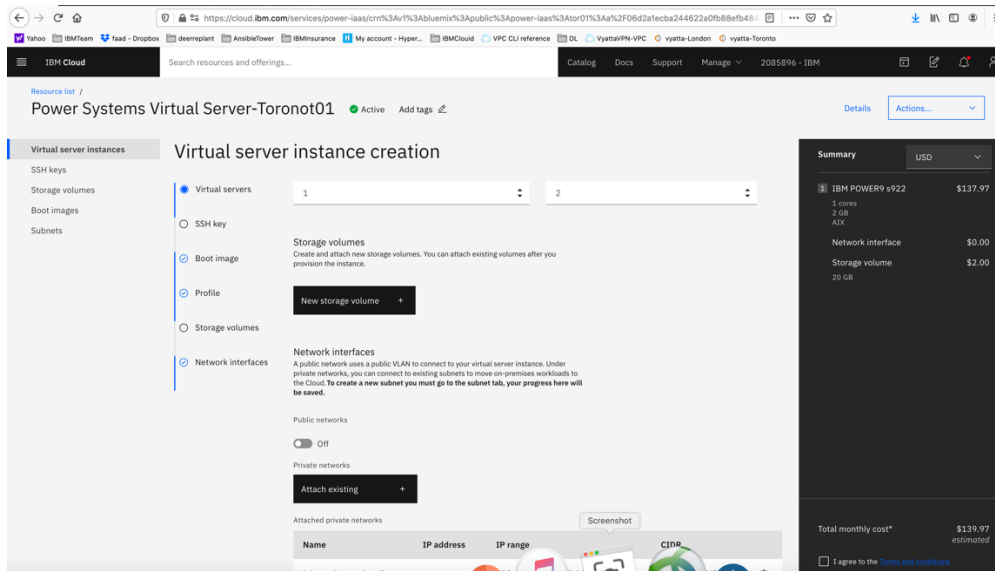
- *Operating System – AIX or IBM i or any other image you may have imported via the "Boot Image" menu on the left.*
- *Image type: AIX 7.1 or 7.2, IBM i 7.3 or 7.4, for example.*
- *Software license: Be sure to include "IBM i PowerHA" and any others as necessary.*
- *Disk types: Type 1 or 3. Type 3 is a less expensive option which we selected.*
- *Machine type: S922 or E980. S922 is the cheaper of the two which we selected.*
- *Processor: Dedicated or Shared or Shared Capped. We chose "shared" as its less expensive.*
- *Choose the number of cores and RAM you will need. The minimum core is "0.25". IBM i partitions have typically required at least 4 GB of memory, just to boot, so this should be your minimum value as well.*
- *You can also attach additional volume to the VSI if you wish. We did not do that here and only used the root volume which is included. As discussed later, it is HIGHLY recommended to add the disks one at a time, after creating the initial Load Source (LS) volume. It makes it much easier to keep track of which disk unit ID matches the volume name in the IBM Cloud UI.*
- *This solution requires that IBM PowerHA SystemMirror for i (Enterprise Edition) is installed on both the production and DR nodes in the solution. This product is 5770HAS *BASE and Option 1*
- *No matter what O/S version/release, it is always recommended to have the latest PTFs, which can be found here:*
- *<https://www.ibm.com/support/pages/ibm-i-support-recommended-fixes>*
- *Since PowerHA for IBM i is owned by HelpSystems, the following link is useful as well, to ensure the latest HA group PTFs are installed:*
- *<https://helpsystemswiki.atlassian.net/wiki/spaces/IWT/pages/162627587/PowerHA+PTF+Groups>*

Next you will scroll down to choose your subnet on which these VSIs will be provisioned. It is assumed you have already created one or more subnets prior to this step.

Click on the “Attached Existing” under networks.



Choose the subnet you wish to attach, and the press “Attach”



Now check the box "I agree to the" And press "create Instance" in lower right-hand side.

Your VSI is now being provisioned.

Order Direct Link Connect Classic to connect PowerVS location to IBM Cloud

You will need to order Direct Link (DL) Connect Classic to allow your Power VSIs in the PowerVS location to communication with Linux/Window VSIs in IBM Cloud and also with all other IBM Cloud services such as Cloud Object Storage and VMware services. This process may take 1-2 weeks to complete.

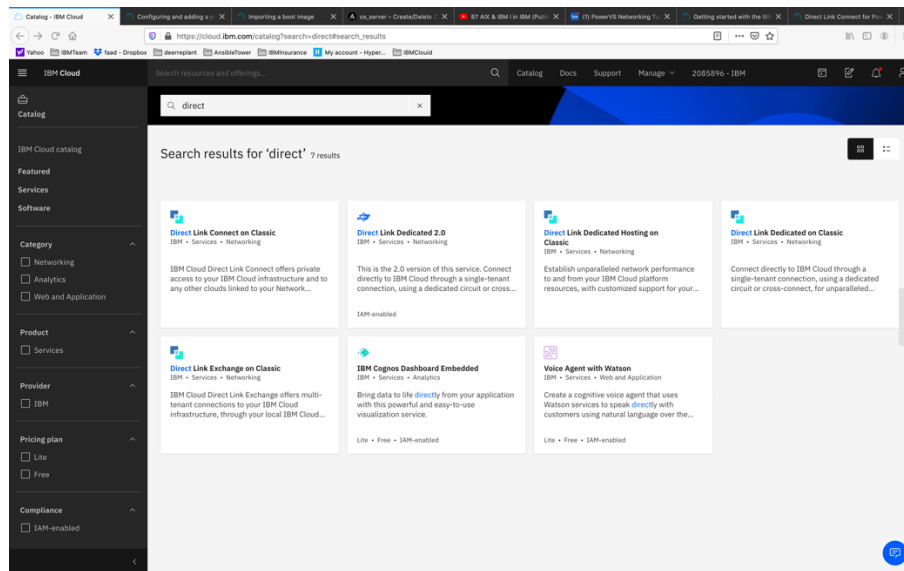
There are several steps involved in completing DL ordering:

- *Order Direct link connect classic service on IBM Cloud UI – see steps below*

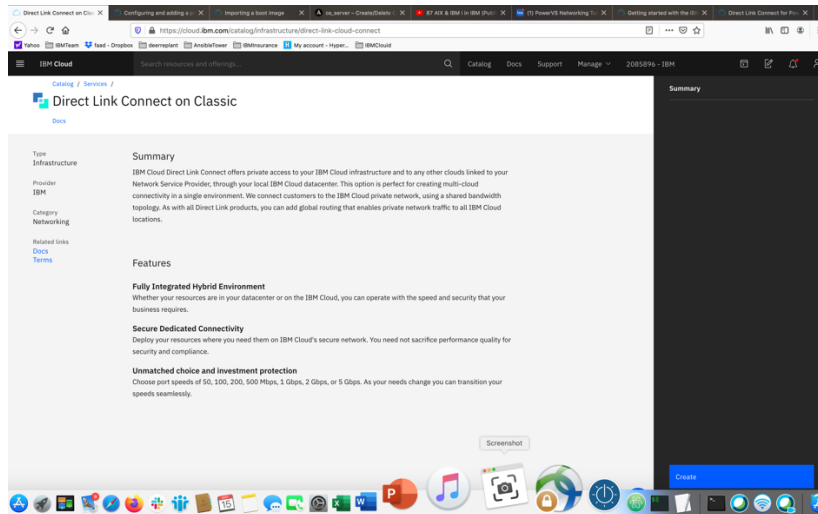
- Next a support ticket will be created, and Support will send you a word document with questionnaires to be completed concerning various DL settings.
- Complete the questionnaires and upload it to support in the ticket.
- Support will then request that you create a new support ticket with the Power System so they can complete their side of the DL provisioning. Attach information about the DL in the original ticket to this ticket.
- The DL will be provisioned, and you will be notified when complete.
- You can now test connection to any Linux/Windows VSI you may have in IBM Cloud and other IBM Cloud services.

To start the DL order process, go to IBM Cloud UI and log in.

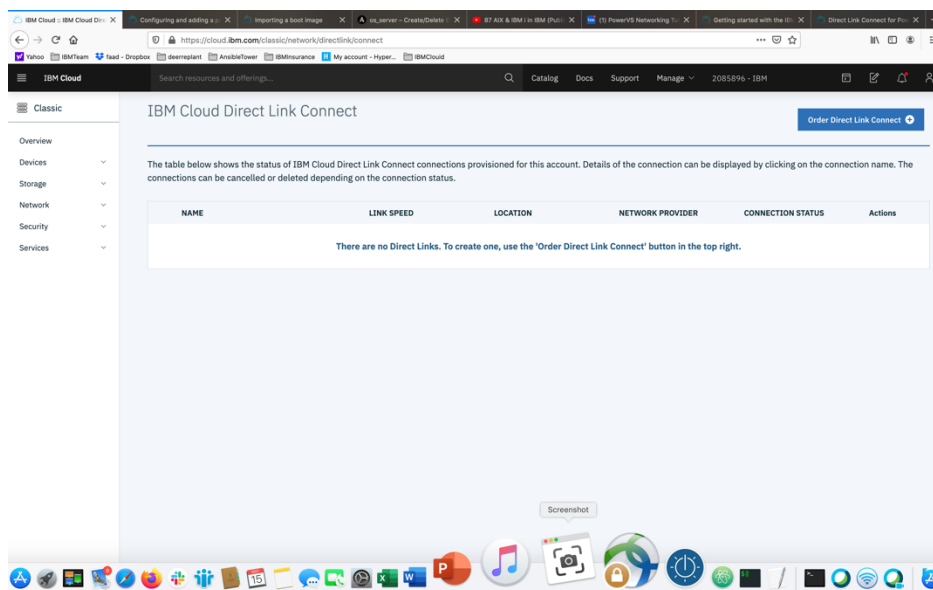
Choose "Catalog" from upper right-hand side, and search for "direct".



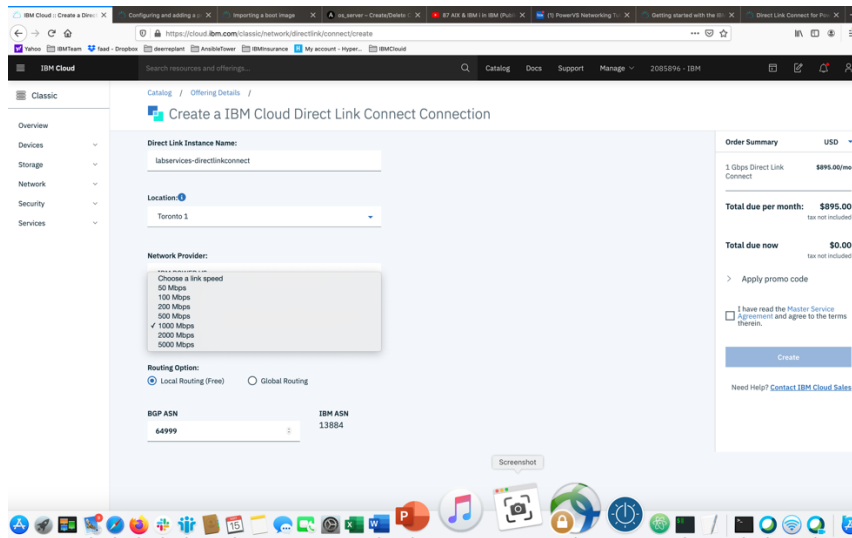
Select "Direct Link Connect on Classic".



Press "Create". There are no options to select.



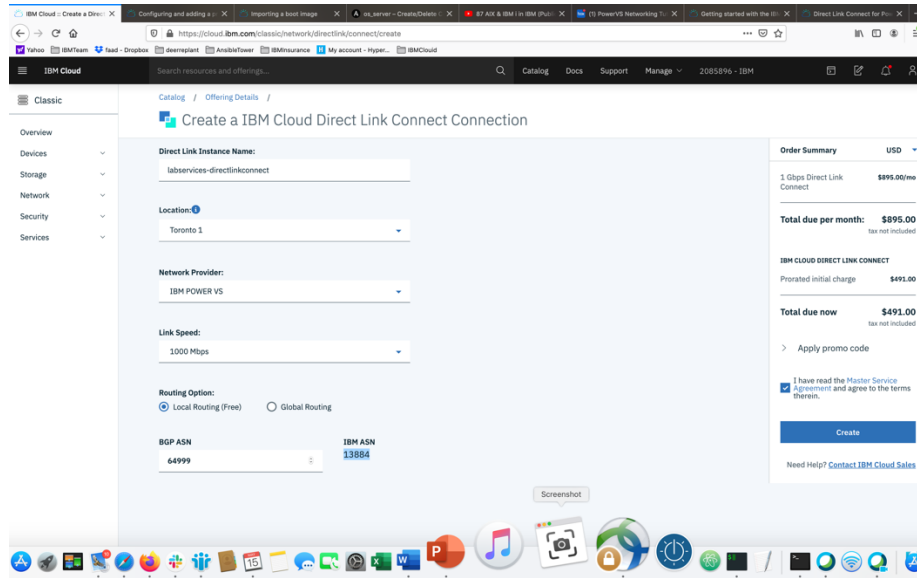
Now choose "Order Direct Link Connect" from top right-hand side.



- Choose a "name" for the DL.
- Choose a location for the DL. This should be the same location as where you created your PowerVS location Service.
- Choose "link speed" under network provider menu.
- Choose "Local Routing (free)"

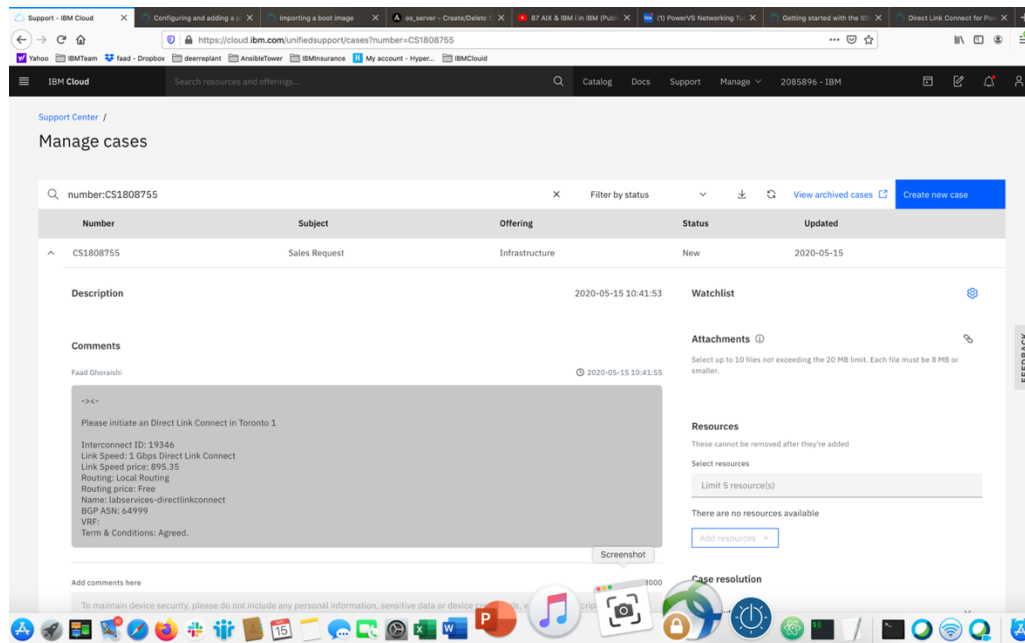
Global routing will require additional charges and will allow for easier PowerVS location-to-PowerVS location communication. You will also need to order a Vyatta Gateway Router to complete your Global routing option via use of a GRE tunnel. Support can help you with this further.

In our case, we decided to use Local Routing and then order a Vyatta Gateway in each PowerVS location and provision a GRE tunnel end-to-end.



- Check the box to accept the offer and press "Create"

A support case will be opened with the information required.



After this is complete, you will then be contacted by support and requested to complete and answer some questions in an attached document and send it back as attachment to the same ticket.

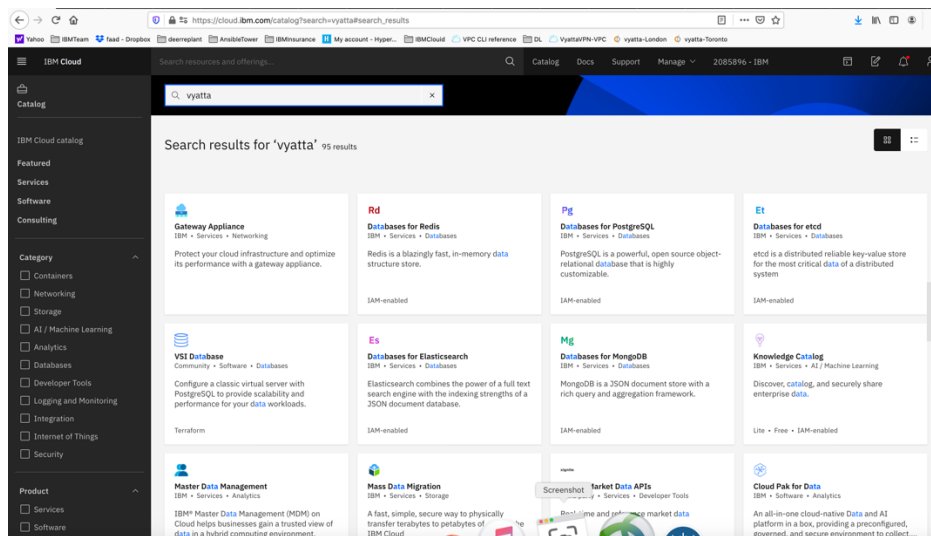
After this step is complete, support will request that you open a new IBM support ticket and address it to the Power System. Include the information in the original DL ticket. This new ticket will be sent to the PowerVS location support to configure their side of the DL connection.

This should be the last step before DL communication works. You can test your connection by pinging IBM Cloud Linux/Windows VSI from your Power VSIs and in reverse.

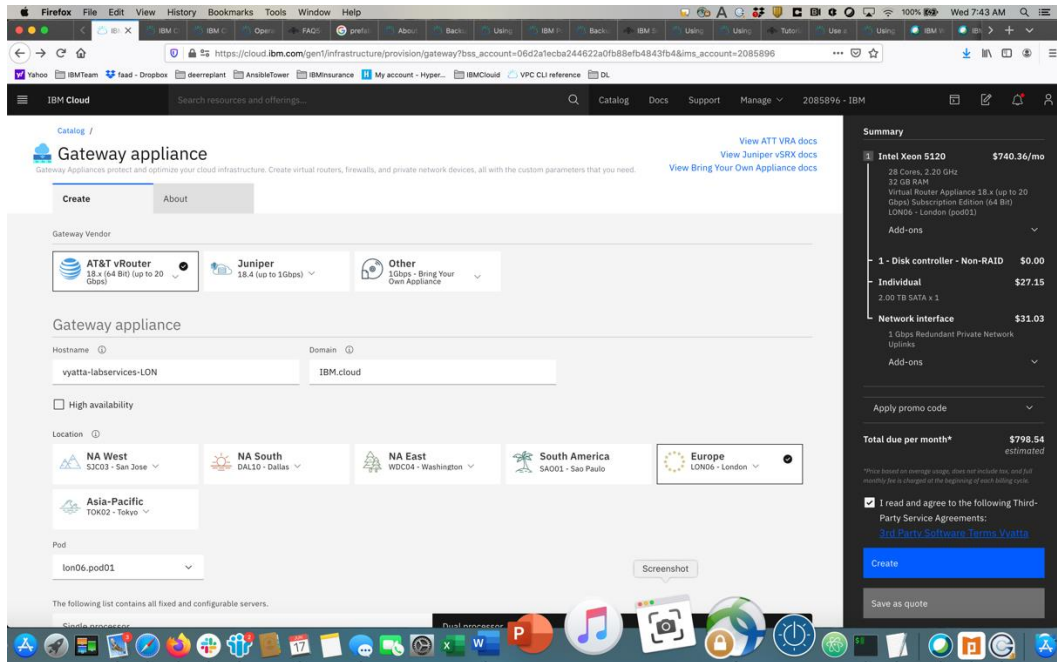
Order two Vyatta Gateways, one in each datacenter

In our scenarios we used two Vyatta Gateways, one in each PowerVS location to provide end-to-end PowerVS location-to-PowerVS location communication using GRE tunnels.

Login to IBM Cloud and click on the “Catalog”, then search for Vyatta.



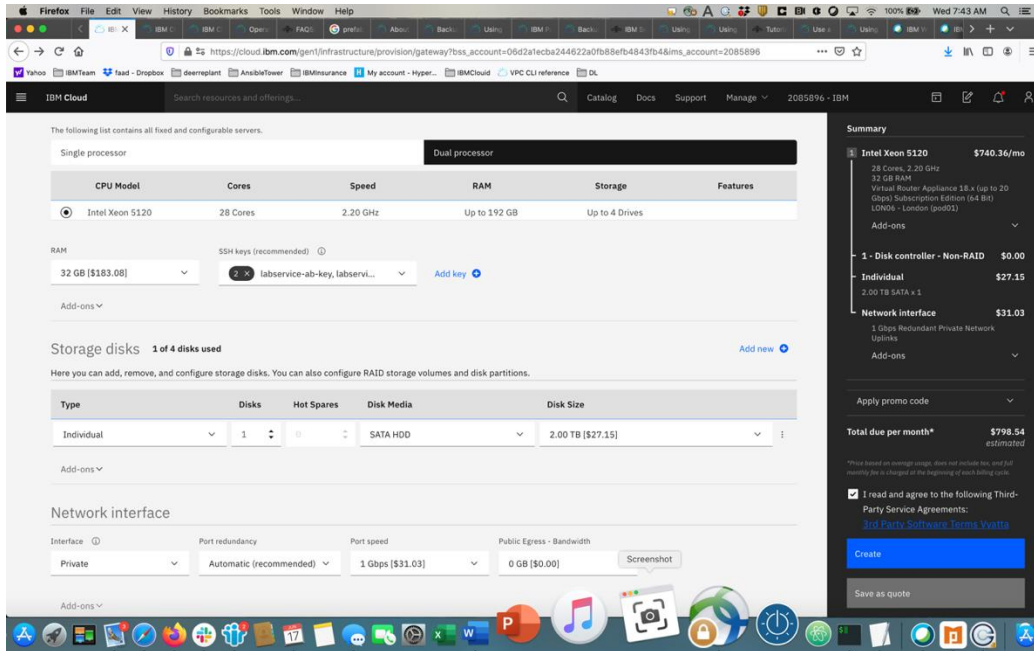
Select “Gateway Appliance” and click on it.



Select "AT&T vRouter". This is the Vyatta Gateway. You have other choices of Gateways, but we will use Vyatta.

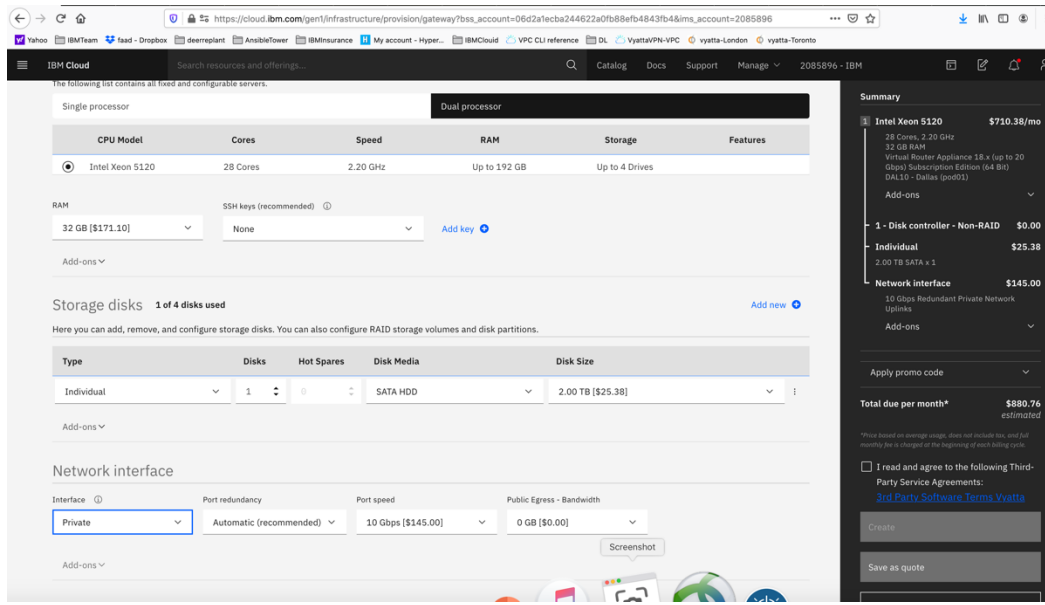
Provide a name for the Gateway and include the PowerVS location name in it so you can distinguish them later.

Select Location to match your PowerVS location.



Choose the following options:

- *Uncheck the High Availability option unless you wish to order one which means you will order two Vyatta Gateways in each PowerVS location. We unchecked this option.*
- *Select the location by pressing on the arrow key in each location to find the exact datacenter where you PowerVS location are located.*
- *You may need to choose the POD if there are several PODs in the selected datacenter location.*
- *Select the CPU single or dual processor. We chose Single Processor.*
- *Select the amount of RAM you wish and add ssh keys if you like to login without password. This can be done later too.*
- *Choose Private network interface unless you wish to use the default which is public/private interface. We chose private network interface only.*

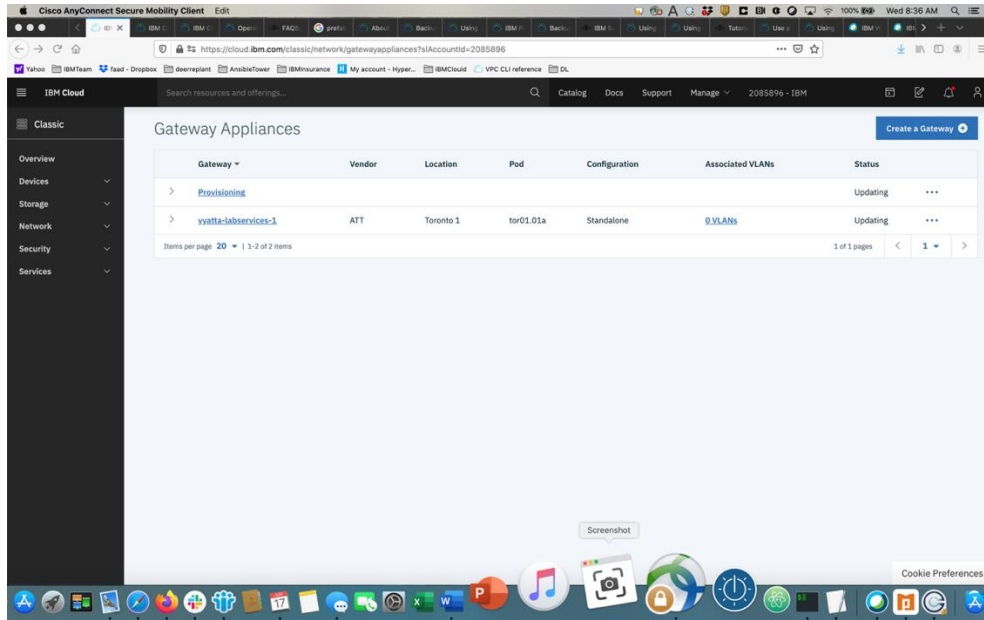


Now check the box to agree with service agreement on the bottom-right side and press "Create"

The Vyatta gateway is now being provisioned. This may take several hours.

You will need to do this process in each of the two PowerVS locations.

After the Vyatta Gateway is provisioned, you can see it listed under "Devices" where you can find your "Vyatta" and "root" user passwords.

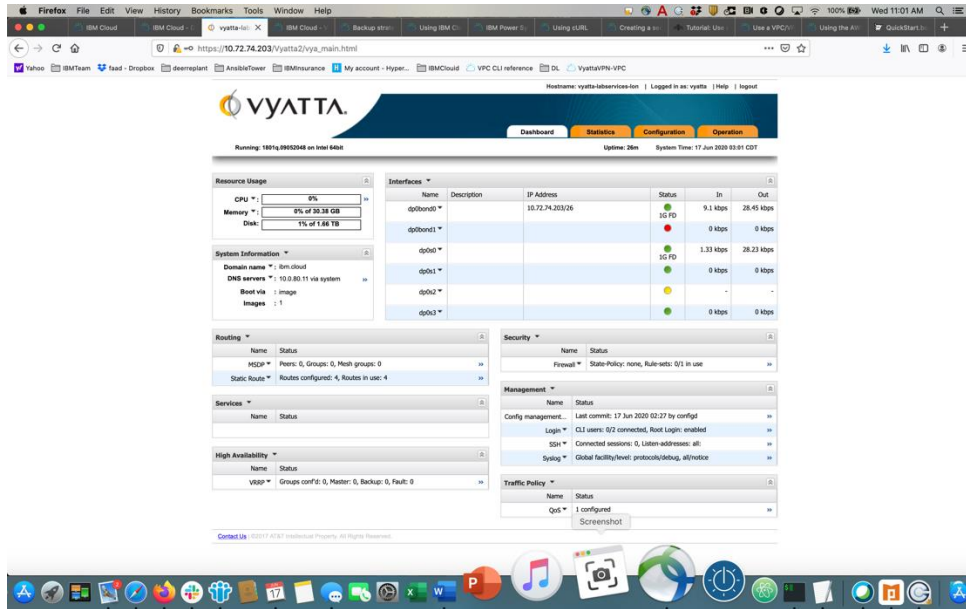


To log into the Vyatta gateway, use a browser and access it via the link:

`https://<ip_address_of_the_vyatta_gateway>`

user: Vyatta

password: as show under "devices" in IBM Cloud UI and password tab on the left.



Typically, you will use a command line to ssh to the Vyatta for further configuration. You will use the "Vyatta" user id to do the configurations.

Request a Generic Routing Encapsulation (GRE) tunnel to be provisioned at each PowerVS location

You will need to open a support ticket to Power Systems and request that a GRE tunnel be provisioned in each PowerVS location. You will need to provide information on the subnets you created in the PowerVS location. They will provision their end of the GRE tunnel and send you the information you will need so you can continue and provision your end of the GRE tunnel on the Vyatta Gateways.

Power Support team will send you the following information for your GRE tunnels after they complete their end of the GRE tunnel:

TOR01:

```
In Tor01 to POWERVS LOCATION GRE:  
Your destination should be 10.254.0.30  
Your tunnel ip 172.20.8.1  
Power-PowerVS location-Side:  
Tor01: interface Tunnel5  
description IBM5-GRE  
vrf forwarding IBM5  
ip address 172.20.8.2 255.255.255.252  
keepalive 5 3  
tunnel source 10.254.0.30  
tunnel destination 10.114.118.34  
tunnel vrf IBM5
```

LON06:

```
In Lon06 to POWERVS LOCATION GRE:  
Your destination should be 10.254.0.26  
Your tunnel ip 172.20.2.1  
Power-PowerVS location-Side:  
Lon06: interface Tunnel4  
description IBM3-GRE  
vrf forwarding IBM3  
ip address 172.20.2.2 255.255.255.252  
keepalive 5 3  
tunnel source 10.254.0.26  
tunnel destination 10.72.74.203  
tunnel vrf IBM3
```

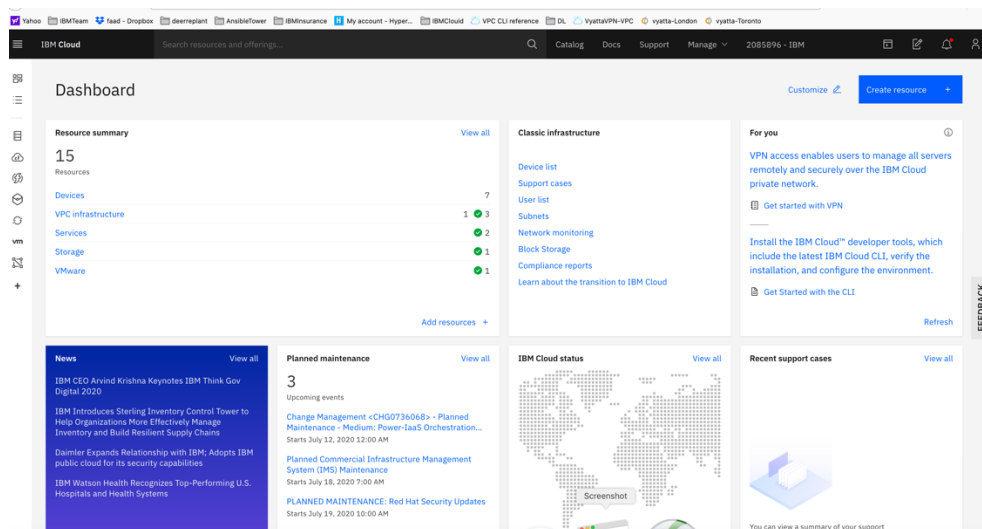
The items shown in Red is what you will need to configure your end of the GRE tunnel in each Vyatta Gateways.

- Note that your tunnel IP address is 172.20.2.1/30 where 255.255.255.252 translate to /30
- Your tunnel destination IP is their tunnel source IP.
- Your tunnel source IP is the IP address of the Vyatta gateway

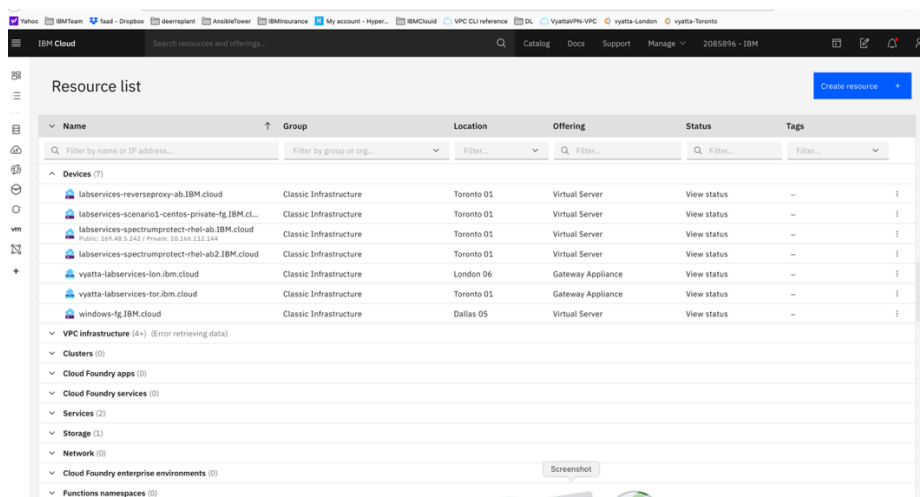
Verify your Vyatta Gateway access.

The Vyatta Gateway address can be find in the IBM Cloud UI under Devices.

Login to IBM Cloud UI and press "IBM Cloud" on top left-hand side.



Click on "Devices"

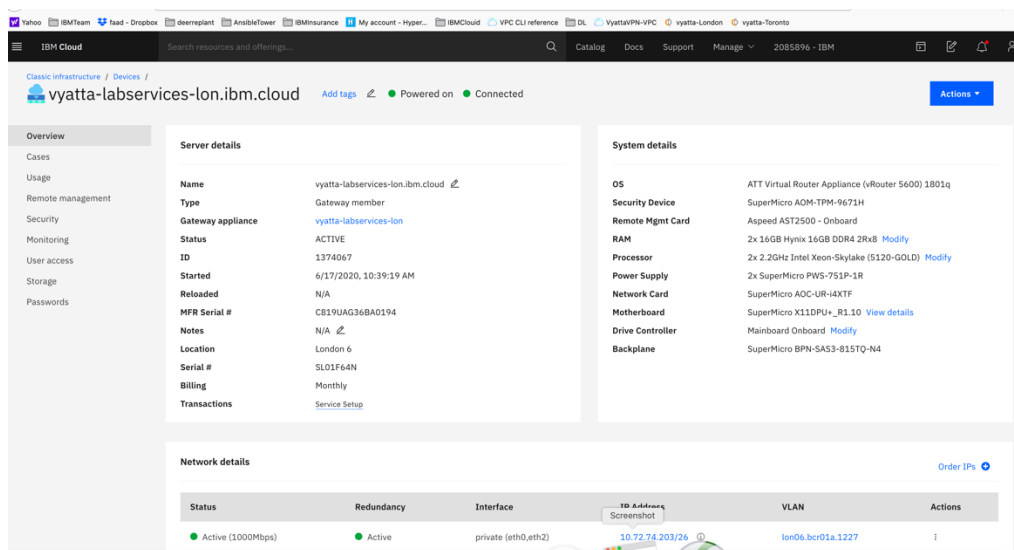


Choose the Vyatta system you like to configure:

- vyatta-labservices-lon.ibm.cloud
- vyatta-labservices-tor.ibm.cloud

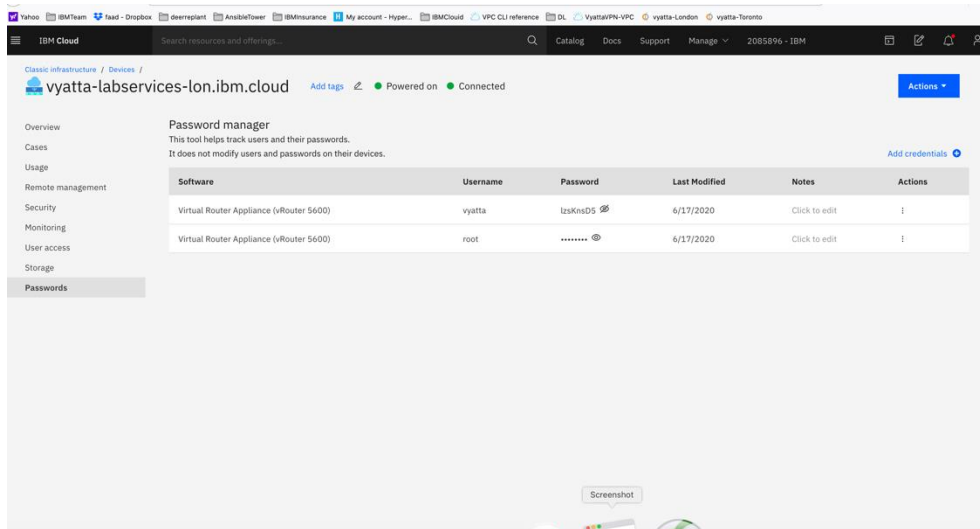
LON06:

Click on the London Vyatta: vyatta-labservices-lon.ibm.cloud



Under the “Network Details” you will see your Vyatta Gateway IP address:

Your credentials are under the “password” menu on the left-hand side. Click on the icon next to the password to see it unencrypted.



Open a browser and login to the Vyatta Gateway using:

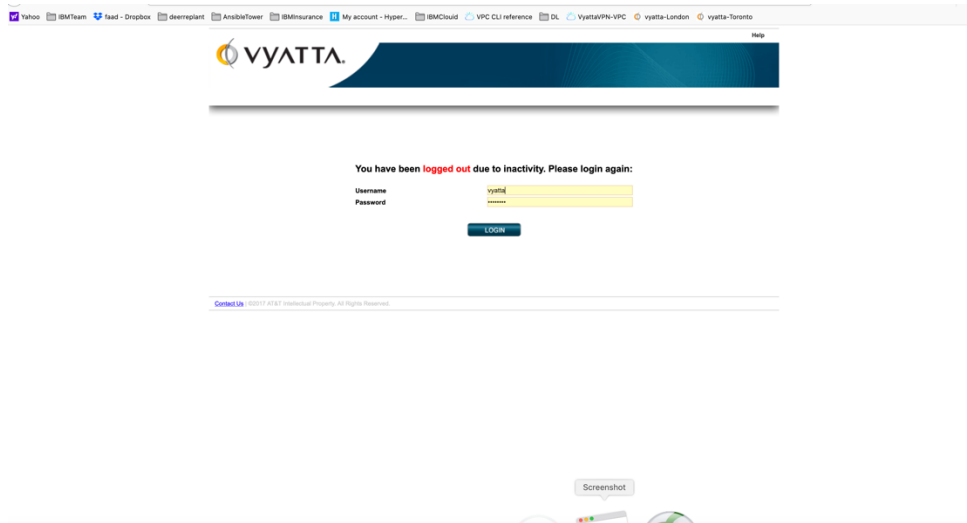
userID: Vyatta

Password: as show in the GUI

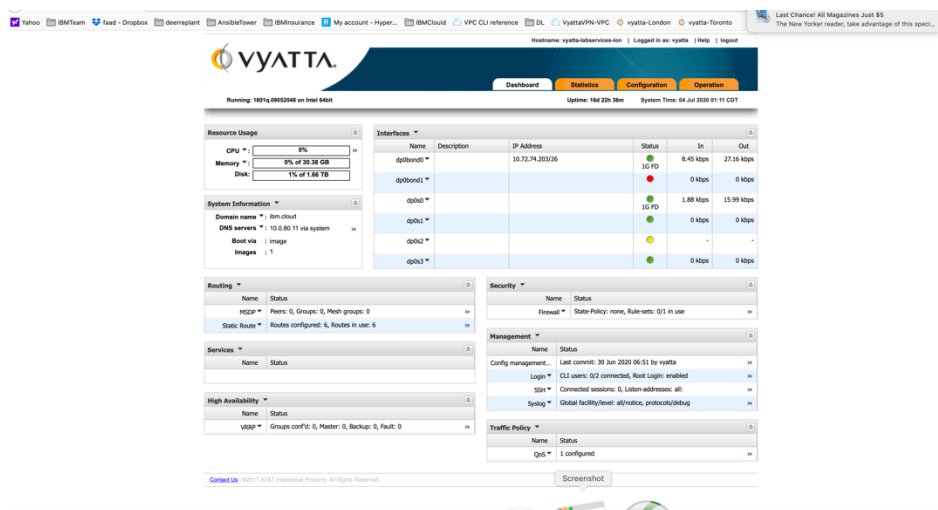
<https://10.72.74.203>

ssh vyatta@10.72.74.203

Note: Prior to login to a 10.x.x.x private IPs in IBM Cloud you will need to start your MotionPro Plus VPN access. This will give you access to IBM Cloud private IPs.



Login with the userID and password.



Now that you have verified you access to the Vyatta Gateways, you will need to now access it via ssh to continue your GRE tunnel provisioning.

Setup PowerVS location GRE tunnels in the Vyatta Gateways

The following references may help in configuring GRE tunnels:

<https://cloud.ibm.com/docs/virtual-router-appliance?topic=solution-tutorials-configuring-IPSEC-VPN>

https://docs.huihoo.com/vyatta/6.5/Vyatta-Tunnels_6.5R1_v01.pdf

<https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-configuring-power>

Open a command window on your Mac/Window.

Note: Prior to login to a 10.x.x.x private IPs in IBM Cloud you will need to start your MotionPro Plus VPN access.

Setup GRE PowerVS location Tunnel in LON06:

userID: Vyatta

Password: as show in the GUI

ssh vyatta@10.72.74.203

ssh to LON06 Vyatta Gateway.

```
The default interactive shell is now zsh.
To update your account to use zsh, please run `chsh -s /bin/zsh`.
For more details, please visit https://support.apple.com/kb/HT208050.
Faads-MacBook-Pro:~ faadghoraishi$
Faads-MacBook-Pro:~ faadghoraishi$ ssh vyatta@10.72.74.203
Welcome to AT&T vRouter 5600

Welcome to AT&T vRouter
Version:      1801q
Description: AT&T vRouter 5600 1801q
Linux vyatta-labservices-lon 4.9.0-trunk-vyatta-amd64 #1 SMP Debian 4.9.124-0vyatta2+2.1 (2018-09-05) x86_64
Last login: Tue Jun 30 00:58:24 2020 from 10.1.232.20

vyatta@vyatta-labservices-lon:~$
```

We are using the information provided by support for LON06 GRE.

```
In Lon06 to POWERVS LOCATION GRE:
Your destination should be 10.254.0.26
Your tunnel ip 172.20.2.1
Power-PowerVS location-Side:
Lon06: interface Tunnel4
description IBM3-GRE
vrf forwarding IBM3
ip address 172.20.2.2 255.255.255.252 (172.20.2.2/30)
keepalive 5 3
tunnel source 10.254.0.26
tunnel destination 10.72.74.203
tunnel vrf IBM3
```

Run the following commands:

We have chosen to call our tunnel "tun0" on the Vyatta Gateway.

```
➤ configure
➤ set interfaces tunnel tun0 address 172.20.2.1/30
➤ set interfaces tunnel tun0 local-ip 10.72.74.203
➤ set interfaces tunnel tun0 remote-ip 10.254.0.26
➤ set interfaces tunnel tun0 encapsulation gre
➤ set interfaces tunnel tun0 mtu 1300
➤ commit
➤ exit
```

```

vyatta@vyatta-labservices-lon:~$ configure
[edit]
vyatta@vyatta-labservices-lon# set interfaces tunnel tun0 address 172.20.2.1/30
[edit]
vyatta@vyatta-labservices-lon# set interfaces tunnel tun0 encapsulation gre
[edit]
vyatta@vyatta-labservices-lon# set interfaces tunnel tun0 mtu 1300
[edit]
vyatta@vyatta-labservices-lon# set interfaces tunnel tun0 local-ip 10.72.74.203
[edit]
vyatta@vyatta-labservices-lon# set interfaces tunnel tun0 remote-ip 10.254.0.26
[edit]
vyatta@vyatta-labservices-lon# commit
[edit]
vyatta@vyatta-labservices-lon# show interfaces tunnel
    tunnel tun0 {
        address 172.20.2.1/30
        encapsulation gre
        local-ip 10.72.74.203
        mtu 1300
        remote-ip 10.254.0.26
    }
[edit]
vyatta@vyatta-labservices-lon# show interfaces tunnel tun0
    tunnel tun0 {
        address 172.20.2.1/30
        encapsulation gre
        local-ip 10.72.74.203
        mtu 1300
        remote-ip 10.254.0.26
    }

```

You can verify that your GRE tunnel is setup by running the following commands:

- *configure*
- *show interfaces tunnel*
- *Or to get more info:*
- *Show interface tunnel tun0*
- *exit*

Setup GRE PowerVS location Tunnel in TOR01:

userID: Vyatta

Password: as show in the GUI

ssh vyatta@10.114.118.34

ssh to Tor01 Vyatta Gateway.

```
Faads-MacBook-Pro:~ faadghoraiishi$ ssh vyatta@10.114.118.34
Welcome to AT&T vRouter 5600

Welcome to AT&T vRouter
Version: 1801q
Description: AT&T vRouter 5600 1801q
Linux vyatta-labservices-1 4.9.0-trunk-vyatta-amd64 #1 SMP Debian 4.9.124-0vyatta2+2.1 (2018-09-05) x86_64
Last login: Tue Jun 30 07:58:37 2020 from 10.1.232.20

vyatta@vyatta-labservices-1:~$
```

```
In Tor01 to POWERVS LOCATION GRE:
Your destination should be 10.254.0.30
Your tunnel ip 172.20.8.1
Power-PowerVS location-Side:
Tor01: interface Tunnel5
description IBM5-GRE
vrf forwarding IBM5
ip address 172.20.8.2 255.255.255.252
keepalive 5 3
tunnel source 10.254.0.30
tunnel destination 10.114.118.34
tunnel vrf IBM5
```

Run the following commands:

We have chosen to call our tunnel "tun0" in the Vyatta Gateway same as the other Vyatta Gateway.

```
➤ configure
➤ set interfaces tunnel tun0 address 172.20.8.1/30
➤ set interfaces tunnel tun0 local-ip 10.114.118.34
➤ set interfaces tunnel tun0 remote-ip 10.254.0.30
➤ set interfaces tunnel tun0 encapsulation gre
➤ set interfaces tunnel tun0 mtu 1300
➤ commit
➤ exit
```

```
vyatta@vyatta-labservices-1# configure
vbash: configure: command not found
[edit]
vyatta@vyatta-labservices-1# set interfaces tunnel tun0 address 172.20.8.1/30
[edit]
vyatta@vyatta-labservices-1# set interfaces tunnel tun0 encapsulation gre
[edit]
vyatta@vyatta-labservices-1# set interfaces tunnel tun0 mtu 1300
[edit]
vyatta@vyatta-labservices-1# set interfaces tunnel tun0 local-ip 10.114.118.34
[edit]
vyatta@vyatta-labservices-1# set interfaces tunnel tun0 remote-ip 10.254.0.30
[edit]
vyatta@vyatta-labservices-1# commit
[edit]
```

To show the status:

- *configure*
- *show interfaces tunnel*
- *Or to get more info:*
- *Show interface tunnel tun0*
- *exit*

```
vyatta@vyatta-labservices-1:~$ configure
[edit]
vyatta@vyatta-labservices-1# show interfaces tunnel
tunnel tun0 {
    address 172.20.8.1/30
    encapsulation gre
    local-ip 10.114.118.34
    mtu 1300
    remote-ip 10.254.0.30
}
```

Setup GRE tunnel between Two Vyatta Gateways

In this section you will setup a new tunnel in each of the two Vyatta gateways to allow for cross Vyatta connection via a GRE tunnel.

In this case we chose the tunnel address and tunnel source and destination IPs. The tunnel address can be any IP subnet you choose. We named our tunnel "tun1" in both Vyatta Gateways. We have selected a similar IP as the ones used in the PowerVS location GRE tunnels. We choose a CIDR of /30 since we only need two IP address, one in Tor01 and one in Lon06.

- *In Lon06 Vyatta the GRE Vyatta-to-Vyatta tunnel address is 172.20.4.1/30*
- *In Tor01 Vyatta the GRE Vyatta-to-Vyatta tunnel address is 172.20.4.2/30*
- *Your tunnel destination IP is the IP address of the Vyatta gateway in each location*
- *Your tunnel source IP is the IP address of the Vyatta gateway in each location*
- *We call the tunnels tun1 in both locations*

TOR01 GRE Configuration:

- *configure*
- *set interfaces tunnel tun1 address 172.20.4.1/30*
- *set interfaces tunnel tun1 local-ip 10.114.118.34*
- *set interfaces tunnel tun1 remote-ip 10.72.74.203*
- *set interfaces tunnel tun1 encapsulation gre*
- *set interfaces tunnel tun1 mtu 1300*
- *commit*
- *exit*

```
vyatta@vyatta-labservices-1# show interfaces tunnel tun1
tunnel tun1 {
  address 172.20.4.1/30
  encapsulation gre
  local-ip 10.114.118.34
  mtu 1300
  remote-ip 10.72.74.203
}
[edit]
vyatta@vyatta-labservices-1#
```

LON06 GRE Configuration:

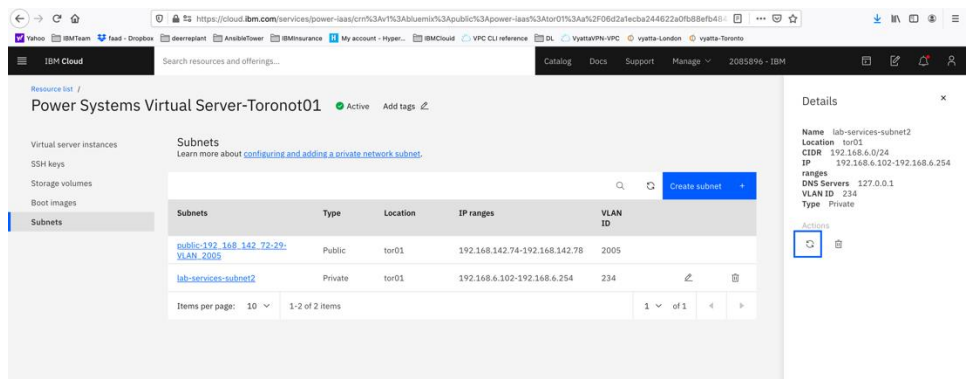
- `configure`
- `set interfaces tunnel tun1 address 172.20.4.2/30`
- `set interfaces tunnel tun1 remote-ip 10.114.118.34`
- `set interfaces tunnel tun1 local-ip 10.72.74.203`
- `set interfaces tunnel tun1 encapsulation gre`
- `set interfaces tunnel tun1 mtu 1300`
- `commit`
- `exit`

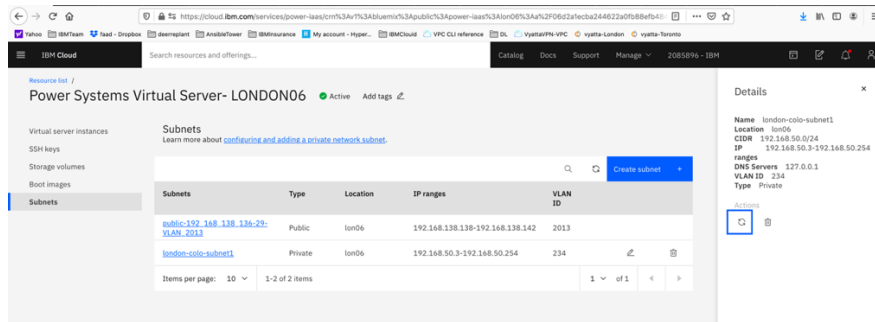
```
vyatta@vyatta-labservices-lon# show interfaces tunnel tun1
tunnel tun1 {
  address 172.20.4.2/30
  encapsulation gre
  local-ip 10.72.74.203
  mtu 1300
  remote-ip 10.114.118.34
}
[edit]
vyatta@vyatta-labservices-lon#
```

The final step needed is to setup static routes in each Vyatta to point the subnets for our PowerVS location to the right tunnels.

Setup GRE tunnel between Two Vyatta Gateways

Find the subnets you created in each PowerVS location in TOR01 and LON06 by accessing the services in the IBM Cloud UI for each PowerVS location.





The static routes in LON06 will need to point to the subnets in TOR01 and vice versa.

We will configure both GREs to the PowerVS location and between Vyattas.

Run the following commands in each Vyatta Gateway after login in via ssh using the Vyatta userID:

in TOR01 Vyatta:

- *configure*
- *set protocols static route 192.168.6.0/24 next-hop 172.20.8.2*
- *set protocols static route 192.168.50.0/24 next-hop 172.20.4.2*
- *commit*
- *exit*

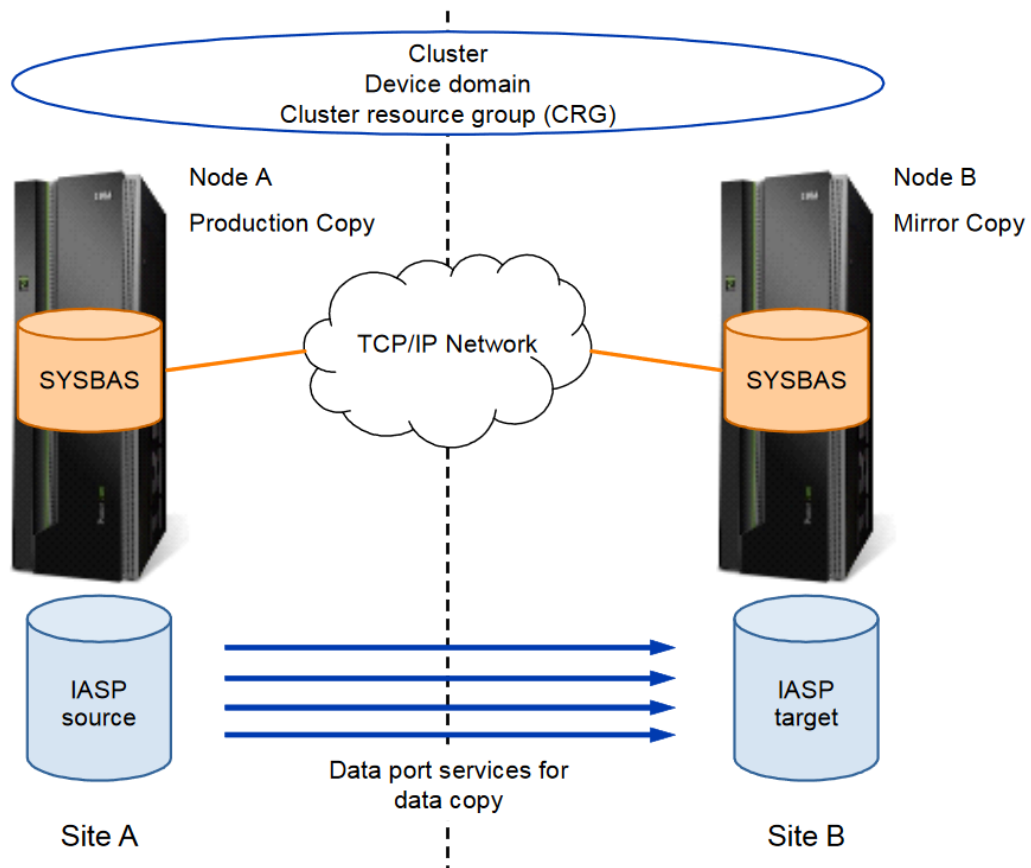
in LON06 Vyatta:

- *configure*
- *set protocols static route 192.168.50.0/24 next-hop 172.20.2.2*
- *set protocols static route 192.168.6.0/24 next-hop 172.20.4.1*
- *commit*
- *exit*

At this point you should have end-to-end connectivity and be able to ping between your Power VSIs in each PowerVS location and also from the Power VSI to IBM Cloud services such as Linux/Windows VSI.

If you cannot ping the IBM Cloud VSIs from the PowerVS location VSIs, you will need to open a ticket to address this issue. Support needs to address this from their Cisco Router side.

Solution Diagram



Chapter 2: Implementation

PowerHA Geographic Mirroring

When creating the IBM i instances via IBM Cloud Services, take note of the following recommendations in preparation for building the environments.

Production IBM i Creation:

A general rule of thumb, assuming most production database objects are moved to the IASP, is a 1:3 ratio of SYSBAS volumes to IASP volumes. Much larger environments may see closer to a 1:6 or 1:9 ratio, as the size of SYSBAS does not need to grow at the same rate as the IASP.

When creating the instance, however, the default size for the Load Source (LS) volume will be 80 GB. Start with ONLY the Load Source, and then add each new volume (SYSBAS or IASP) individually (later on in the process), in order to keep track of each disk unit ID as they are added, as this can otherwise be difficult to map later.

The names of the disks (from IBM Cloud Services) across ALL instances within that server must be unique. These names will not present themselves in the IBM i interface, but only be visible from the Cloud interface. This is why it is useful to keep track individually of the disk unit ID (from IBM i interface) and disk name (from Cloud interface) so you can assign your disks to the appropriate ASP. As noted above, these are best created at a later step in the process.

When choosing a name for the instance, keep in mind that the IBM i will use the first 8 characters of that instance name as the system name, so it is best to choose a unique 8-character instance name as well.

You will need to choose IBM PowerHA under "Software Licenses", which will load and enable the PowerHA Enterprise Edition by default. This will allow for asynchronous Geographic Mirroring, which is best for replication between IBM i Cloud instances.

Once the instance is created and Active, open the console and wait for the log-in screen. If the default image was deployed, the qsecofr password (QSECOFR) will be disabled and need to be changed (NOTE: This is the O/S log-in. The DST/SST password will be changed in a later step).

On the "Work with software agreements" screen, use Option 5 (Display) for all, and press Enter. Then press F15 (Accept ALL) on each license. Once all licenses are accepted, press F3 (Exit) to main menu.

During this time, the IP interfaces and line descriptions are still being configured. This process can take up to 5 minutes to complete. If any external IP's were requested, these will not show up in **CFGTCP, Option 1** (Work with TCP/IP interfaces).

Note that the line descriptions get created as CLOUDINITx names, and the TCP/IP interfaces are assigned to those automatically. The line descriptions are configured as **ONLINE(*YES)** and TCP/IP interfaces are configured as **AUTOSTART(*YES)** but can be changed to fit the needs of the business.

```
Work with Line Descriptions                                     System:  IBMIPROD
Position to . . . . . _____ Starting characters
Type options, press Enter.
 2=Change  3=Copy   4=Delete  5=Display  6=Print  7=Rename
 8=Work with status  9=Retrieve source

Opt  Line      Type      Text
--  -
--  CLOUDINIT0 *ELAN
--  CLOUDINIT1 *ELAN
--  CLOUDINIT2 *ELAN
```

```

Work with TCP/IP Interfaces
System:  IBMIPROD
Type options, press Enter.
  1=Add  2=Change  4=Remove  5=Display  9=Start  10=End

Opt  Internet      Subnet      Line      Line
    Address      Mask       Description  Type
---  -
    127.0.0.1     255.0.0.0  *LOOPBACK  *NONE
    192.168.6.118 255.255.255.0 CLOUDINIT0 *ELAN
    192.168.142.74 255.255.255.248 CLOUDINIT1 *ELAN

```

Once TCP/IP is configured, the following changes are recommended:

- **CHGSYSVAL SYSVAL(QIPLTYPE) VALUE(1)**
Note: This will be changed back after DST/SST password is changed upon the next IPL.
- **CHGSYSVAL SYSVAL(QLMTSECOFR) VALUE(0)**
- **CHGSYSVAL SYSVAL(QAUTOVRT) VALUE(100)**
Note: This can be any number desired, above zero.
- **CHGTCPSVR SVRSPCVL(*TELNET) AUTOSTART(*YES)**
- **CHGTCPSVR SVRSPCVL(*FTP) AUTOSTART(*YES)**
- **CHGTCPSVR SVRSPCVL(*SSHD) AUTOSTART(*YES)**
- **CHGTCPSVR SVRSPCVL(*INETD) AUTOSTART(*YES)**

At this point, from the IBM Cloud interface, the additional disks for SYSBAS and IASP can be added.

NOTE: It is best-practice to add these disks one at a time, and then record the associated disk unit ID that is presented in the IBM i. This will be useful in determining which disks will be assigned to SYSBAS and which will be assigned to the IASP.

Each new disk will present itself as a non-configured disk. Wait until after the next steps to add/balance any additional SYSBAS disks.

At this point, issue a **PWRDWNSYS OPTION(*IMMED) RESTART(*YES)** to shut down and IPL the IBM i instance in Attended mode with DST (per the CHGSYSVAL QIPLTYPE above).

When the initial menu comes up, use DST to change the DST password for qsecofr (default is QSECOFR). The add/balance of the additional

SYSBAS disks can be done at this time, or as a separate step once the IBM i instance has completed the IPL.

After the IPL and after the add/balance of the additional SYSBAS disks, change the IPL type back to Unattended.

- **CHGSYSVAL SYSVAL(QIPLTYPE) VALUE(0)**

Be sure all other post-IPL jobs and servers are active as expected.

HA/DR IBM i Creation:

Follow the same recommendations as above for the Production IBM i making sure non-configured disks for the IASP are left alone at this point.

Network Connectivity between Production and HA/DR Instances:

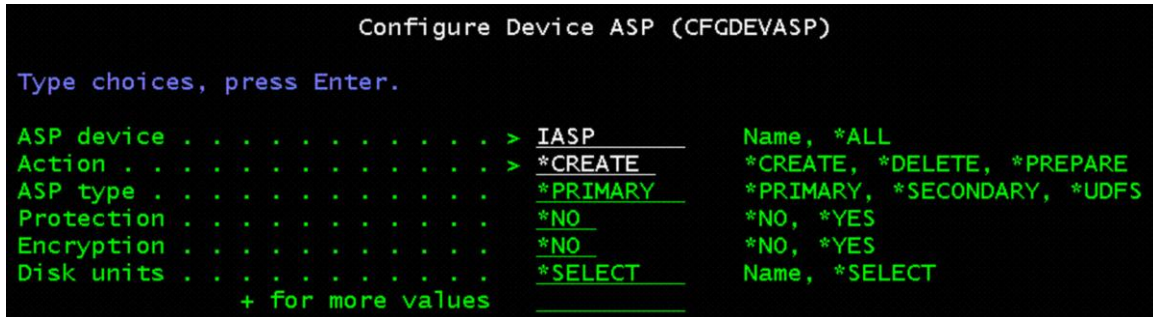
Once both instances are created, it is imperative that both instances are able to communicate, both for intended PowerHA Cluster communication and for PowerHA Geographic Mirroring replication. For better performance, it is recommended to use one set of IP interfaces for inter-LPAR communication, such as FTP, remote journaling, BRMS, etc, and include PowerHA Cluster communication with that. A secondary set of interfaces can be reserved for PowerHA Geographic Mirroring, so that any spikes in replication traffic have less impact on packet delay for other inter-LPAR communication. It may also be useful to use a separate IP interface for user log-in and administration. This can be an internal IP interface, where the user connects to the corporate network, and then accesses the system(s) via the internal IP, or optionally, an external IP interface can be assigned to the instance via the IBM Cloud Services.

IASP Creation:

The IASP should be created on the Production IBM i instance in preparation for IASP-enablement. Keep in mind that these disks are

virtual volumes created from external storage, with RAID protection already pre-assigned. Therefore, there is no need to add protection to the IASP disks upon creation. To create the IASP, do the following:

**CFGDEVASP ASPDEV(<IASP Name>) ACTION(*CREATE)
 TYPE(*PRIMARY) PROTECT(*NO) ENCRYPT(*NO)
 UNITS(*SELECT)**



The next screen will prompt with the ability to select from the non-configured disks which ones to add to the IASP. Once selected, Press Enter to begin IASP configuration. This process may take several hours, depending on the number and size of the disks. Once completed, the ASP device description can be varied on, and the resource name will show with the same name as the IASP. At any time, a matching device description can be created on the HA/DR node with the following:

CRTDEVASP(<IASP Name>) RSRCTYPE(<IASP Name>)

This is ONLY the device description. There is no IASP on the HA/DR node at this time, nor should there be. The non-configured disks for PowerHA Geographic Mirroring will be created later.


```

Create Device Desc (ASP) (CRTDEVASP)

Type choices, press Enter.

Device description . . . . . > IASP _____ Name
Resource name . . . . . > IASP _____ Name
Relational database . . . . . *GEN _____
Message queue . . . . . *SYSOPR _____ Name
Library . . . . . _____ Name, *LIBL, *CURLIB
Text 'description' . . . . . *BLANK _____

```

NOTE: It is HIGHLY-recommended for any customer to use IBM Lab Services for assistance in migrating a non-IASP environment to an IASP-enabled environment, before moving forward with any PowerHA-managed solution. Only data included in the IASP is being replicated, and switchover/failover activities will not work, if the partitions have not been properly configured to support IASP.

PowerHA Clustering Configuration:

Once the Production IBM i has been properly IASP-enabled, a DR solution with PowerHA Geographic Mirroring can be implemented.

On both instances, perform the following steps:

- **CHGTCPSVR SVRSPCVL(*INETD) AUTOSTART(*YES)**
- **STRTCPSVR SERVER(*INETD)**
- **CHGNETA ALWADDCLU(*ANY)**

From the Production IBM i instance do the following to create the PowerHA Cluster:

1. **CRTCLU CLUSTER(<Cluster Name>) NODE((<Prod Node Name> (<Prod Node IP Address>)) START(*NO)**
2. **WRKCLU**, press Enter, to verify cluster was created
3. From that **WRKCLU** menu, **Option 6** (Work with cluster nodes) will show the production node with a status of New.
4. Use **Option 8** (Start) to start clustering on the production node. The status will change to Active.

5. From that same screen, use **Option 1** (Add) to add the DR node to the cluster, providing the IP address, and specifying Start Indicator as *NO. That node will now be listed and show a status of New.
6. Use **Option 8** (Start) to start clustering on the DR node. The status will change to Active.

```

PowerHA                               Work with Cluster Nodes
Local node . . . . . : CLOUDPRD
Consistent information in cluster . . . . : Yes

Type options, press Enter.
1=Add   2=Change   4=Remove   5=Display more details   6=Work with monitors
8=Start 9=End

Opt      Node          Status      Device Domain
--      -
--      CLOUDDR        Active      CLDDEVDMN
--      CLOUDPRD       Active      CLDDEVDMN

```

Adding Cluster Nodes to Device Domain:

From the Production IBM i instance do the following to add the PowerHA cluster nodes to the same device domain:

1. From **WRKCLU**, **Option 7** (Work with device domains), there should be no device domain listed yet.
2. Use **Option 1** (Add) to add a device domain to the cluster, and specify the production node name.
3. The **WRKCLU**, **Option 7** screen will now show the new device domain with one node (the production node) listed to the right.
4. Use **Option 6** (Work with nodes) and then select **Option 1** (Add) to add the DR node to the device domain. This screen will now list both nodes as Active in that device domain.
5. The **WRKCLU**, **Option 7** screen will now show the device domain with both nodes listed to the right.

```

PowerHA                               Work with Device Domains

Consistent information in cluster . . . : Yes

Type options, press Enter.
  1=Add   6=Work with nodes   7=Work with switchable hardware

Opt   Device Domain      Number
-----
  6   CLDDEVDMN          2   CLOUDPRD CLOUDDR

```

```

PowerHA                               Work with Device Domain Nodes

Device domain . . . . . : CLDDEVDMN
Consistent information in cluster . . . : Yes

Type options, press Enter.
  1=Add   4=Remove

Opt   Node              Status
-----
  -   CLOUDDR           Active
  -   CLOUDPRD          Active

```

Creating the Device Cluster Resource Group (CRG):

From the Production IBM i instance, do the following to create the device CRG used for switchover/failover of the IASP:

1. The **WRKCLU, Option 9** (*Work with cluster resource groups*) screen should be empty.
2. Use **Option 1** (*Create*) and provide a name for the CRG with 10 characters or less, and press Enter.
3. Select **Option 1** (*Cluster resource group*) from the pop-up menu (at V7R4) and press Enter.
4. Specify Type=*DEV, Exit Program=*NONE, and User Profile=*NONE
5. Next to "Recovery domain node list", type a '+' and press Enter.
6. For the first entry in the recovery domain, use the name of the production node with a node role of *PRIMARY and a site name that will be unique from the DR node. Specify at least one IP

- address to be used by the production node for Geographic Mirroring (replication). Then page down.
- For the next entry in the recovery domain, use the name of the DR node with a node role of *BACKUP, a sequence number of 1, and a site name that is unique from the Production node. Specify at least one IP address to be used by the DR node for Geographic mirroring (replication). Press Enter, then page down.
 - (Optional) Provide a text description for the CRG.
 - Press Enter to create the CRG. The status will show as Inactive.

```
PowerHA Work with Cluster Resource Groups
Consistent information in cluster . . . : Yes
Type options, press Enter.
1=Create 2=Change 3=Change primary 4=Delete 5=Display
6=Recovery domain 7=Configuration objects 8=Start 9=End
10=Configure
```

Opt	Container/ CRG	Type	Status	Primary Site/Node
<u>6</u>	<u>CLDGEOMIR</u>	*DEV	Active	CLOUDPRD

```
PowerHA Work with Recovery Domain
Cluster resource group . . . . . : CLDGEOMIR
Consistent information in cluster . . . : Yes
Type options, press Enter.
1=Add node 4=Remove node 5=Display more details
```

Opt	Node	Status	Current Node Role	Preferred Node Role	Site Name
<u>-</u>	<u>CLOUDDR</u>	Active	*BACKUP 1	*BACKUP 1	TORONT02
<u>-</u>	<u>CLOUDPRD</u>	Active	*PRIMARY	*PRIMARY	TORONT01

Add the IASP to the Device CRG:

From the Production IBM i instance, do the following to add the IASP to the device CRG:

- WRKCLU, Option 9** (Work with cluster resource groups) will show the device CRG as Inactive.
- Use **Option 7** (Configuration objects) to show no objects yet added to the CRG.

3. Use **Option 1** (Add) with the name of the IASP and a type of *DEVD, and press Enter.
4. If automatic vary-on of the IASP is desired after switchover/failover (more common) change "Configuration object online" to *ONLINE.
5. Specify *NONE for "Server takeover IP address" as managing an automated switch of a primary IP interface is done more efficiently with IASP exit programs. Press Enter to complete the process.
6. The IASP will now be listed in **WRKCLU, Option 9, Option 7**.

```
PowerHA                               Work with Cluster Resource Groups
Consistent information in cluster . . . : Yes
Type options, press Enter.
1=Create  2=Change  3=Change primary  4=Delete  5=Display
6=Recovery domain  7=Configuration objects  8=Start  9=End
10=Configure

Opt      Container/CRG      Type      Status      Primary Site/Node
-----
 7      CLDGEOMIR          *DEV      Active      CLOUDPRD
```

```
PowerHA                               Work with Configuration Objects
Cluster resource group . . . . . : CLDGEOMIR
Consistent information in cluster . . . : Yes
Type options, press Enter.
1=Add  2=Change  4=Remove  5=Display more details
6=Configuration status

Opt      Configuration Object Name      Object Type      Device Type      Device Subtype      Vary Online
-----
 1      IASP01                          *DEVD           *ASP             Primary             *ONLINE
```

Start Geographic Mirroring of the IASP:

From the Production IBM i instance, do the following to start Geographic Mirroring on the IASP:

1. **CFGGEOMIR ASPDEV(<IASP Name>) ACTION(*CREATE) SSN(<DR Site ASP Copy>/<Prod Site ASP Copy>/<ASP**

Session Name>) **DELIVERY(*ASYNC) UNITS(*SELECT).**
 Press Enter.

NOTE: DR Site ASP Copy and Prod Site ASP Copy are unique names assigned to differentiate which copy of the IASP is referenced. Both copies are the same IASP, but exist on two separate sets of disks.

NOTE: ASP Session Name is label for the Geographic Mirroring replication, against which actions can be performed (e.g. Suspend, Resume, Detach, Reattach)

```

Configure Geographic Mirror (CFGGEOMIR)
Type choices, press Enter.
ASP device . . . . . > IASP           Name
Action . . . . . > *CREATE         *CREATE, *DELETE
Source site . . . . . TORONTO1       Name, *
Target site . . . . . TORONTO2       Name, *
Session . . . . . CLDGEOMIR         Name, *NONE
  Source ASP copy description . . . . . CLOUDPRD       Name
  Target ASP copy description . . . . . CLOUDDR        Name
Transmission delivery . . . . . *ASYNC   *SYNC, *ASYNC
Disk units . . . . . *SELECT         Name, *SELECT
+ for more values
  
```

2. The next screen will show a list of non-configured disk units from the DR node. Select those that are to be included in Geographic Mirroring. Use F9 (Calculate selection) to view the capacity that will result once configuration is complete.
3. Press Enter to begin configuration of Geographic Mirroring. This may take up to several hours, depending on the number and size of disks used.
4. Once complete, **DSPASPSSN SSN(<ASP Session Name>)** will show the status as *RESUMPND
5. Vary on the IASP device, and re-check the ASP session status. It will show as *SUSPENDEd.
6. **WRKCLU, Option 9** (Work with cluster resource groups) and use Option 8 (Start) to start the CRG. The status will change to Active.

7. **CHGASPSSN SSN(<ASP Session Name>) OPTION(*RESUME)**
will restart synchronization of Geographic Mirroring

```
PowerHA                               Work with ASP Copy Descriptions                               IBMIPROD
                                                                              07/16/20 09:31:52
Device domain . . . . . : CLDDEVDMN

Type options, press Enter.
 1=Add copy      2=Change copy      4=Remove copy      5=Display copy
21=Start session 22=Change session 24=End session     25=Display session

Opt   ASP Device      ASP Copy      ASP Session      Session
-----
 25   IASP01          CLOUDPRD     CLDGEOMIR        *GEOMIR
      IASP01          CLOUDDR      CLDGEOMIR        *GEOMIR
```

```
PowerHA                               Display ASP Session                               IBMIPROD
                                                                              07/16/20 09:38:08
Session . . . . . : CLDGEOMIR
Type . . . . . : *GEOMIR

Source node . . . . . : CLOUDPRD
Target node . . . . . : CLOUDDR
Transmission Delivery . . . . . : *ASYNC

More...

Copy Descriptions

ASP Device  ASP Copy  Role  Vary State  Replication State  Data State
IASP01     CLOUDPRD  PRODUCTION  AVAILABLE  ACTIVE  USABLE
IASP01     CLOUDDR  MIRROR     VARYON     ACTIVE  UNUSABLE

Bottom

Press Enter to continue
F3=Exit  F5=Refresh  F11=View 2  F12=Cancel  F19=Automatic refresh
```

Chapter 3: Troubleshooting

Of course, with any new set up, errors can occur, and although some may be related to infrastructure issues, others arise simply due to familiarity with new technology and the differences that using the IBM Cloud may present from legacy solutions. The following are some of those changes that may take alternative methods than what may be the norm with non-cloud environments.

- *When selecting disk units to assign to the new cloud instance(s), start by assigning only the initial Load Source (LS) volume, and identify that in the IBM Cloud GUI for easy reference. Once the instance is created, make note of the disk unit ID in SST/DST. Add subsequent disks one at a time, providing a naming scheme in the IBM Cloud GUI that easily identifies SYSBAS disks or IASP disks, while at the same time, keeping track of the associated disk unit ID from the IBM i instance. This will be crucial in determining which disk units to add with **CFGDEVASP** and **CFGGEOMIR**.*
- *When selecting a name for the IBM i instances, choose a name that is 8 characters or less, as the IBM i will truncate the instance name to 8 characters when assigning a system name.*
- *As with any Geographic Mirroring solution, ensuring sufficient throughput on the replication interfaces is key. Separate other traffic onto other subnets, and ideally other adapters, so that the replication traffic is not hindered and subject to auto-suspend issues or lengthy switchover/failover processes.*
- *Additionally, it is best-practice to keep the PowerHA for IBM i clustering communication on a separate subnet than replication, because it requires the "heartbeat" traffic to respond in a timely manner to prevent cluster suspend ("Partition") status. This*

would prevent any desired switchover/failover from being allowed.

For additional troubleshooting with PowerVS Cloud infrastructure, use the “Support” option from the IBM Cloud Services GUI to submit a case summary for the IBM team to investigate. These issues would be specific to the infrastructure, network, images, servers, disk enlistment or pathing issues, ...etc.

For issues specific to PowerHA for IBM i Geographic Mirroring configuration, errors or performance, start with a call to 800-IBM-SERV, and request the High Availability Solutions (HAS) team first. They can then engage development teams or Cloud Support, if necessary.

Chapter 4: Additional Resources

Please be sure to visit the following link for additional FAQ's:

<https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-power-iaas-faqs>