

# **AIX 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.

**Primitivo Cervantes**

**Vess Natchev**

**[ibmsls@us.ibm.com](mailto:ibmsls@us.ibm.com)**

# TABLE OF CONTENTS

<b>CHAPTER 1: SOLUTION OVERVIEW.....</b>	<b>1</b>
1.1 Introduction .....	1
1.2 Use Cases .....	1
1.2.1 Geographic Logical Volume Manager (GLVM) Replication ....	1
1.2.2 Geographic Logical Volume Manager (GLVM) Replication with PowerHA.....	2
1.3 Solution Components and Requirements .....	2
1.3.1 Components .....	2
1.3.2 Requirements .....	2
1.4 Solution Diagram .....	3
<b>CHAPTER 2: IMPLEMENTATION .....</b>	<b>4</b>
2.1 Base IBM Cloud PowerVSI and Networking setup .....	4
2.1.1 Open an IBM Cloud account .....	4
2.1.2 Create PowerVS location Service and Subnet(s) .....	5
2.1.3 Provision AIX and IBM i VSIs in each PowerVS location.....	5
2.1.4 Order Direct Link Connect Classic to connect PowerVS location to IBM Cloud .....	5
2.1.5 Order two Vyatta Gateways, one in each datacenter .....	6
2.1.6 Request a Generic Routing Encapsulation (GRE) tunnel to be provisioned at each PowerVS location.....	6

2.1.7	Configure three GRE tunnels in the Vyatta Gateways .....	6
2.1.8	Diagrams.....	7
2.1.9	Create PowerVS location Services and Subnet(s) .....	7
2.1.10	Provision AIX and IBM i VSIs in each PowerVS location ....	12
2.1.11	Order Direct Link Connect Classic to connect PowerVS location to IBM Cloud .....	17
2.1.12	Order two Vyatta Gateways, one in each datacenter .....	21
2.1.13	Request a Generic Routing Encapsulation (GRE) tunnel to be provisioned at each PowerVS location .....	26
2.1.14	Setup PowerVS location GRE tunnels in the Vyatta Gateways.....	31
2.2	Geographic Logical Volume Manager (GLVM) Replication.....	39
2.2.1	Basic Concepts .....	39
2.2.2	Example environment .....	42
2.2.3	Preparing for IBM AIX GLVM setup.....	42
2.2.4	Change or verify the volume groups are a scalable volume group .....	43
2.2.5	Updating the volume groups so as to not vary on automatically at reboot time .....	44
2.2.6	Updating the volume groups by turning off bad block relocation .....	45
2.2.7	Updating volume groups by setting super strict mirror pools .....	46
2.2.8	Defining mirror pools .....	47
2.2.9	Setting LV strictness .....	49
2.2.10	Implementing AIX GLVM RPV servers and RPV clients.....	50

2.2.11	Implementing AIX GLVM RPV servers .....	50
2.2.12	Implementing AIX GLVM RPV clients .....	55
2.2.13	Replicating data between sites using AIX LVM .....	59
2.2.14	Accessing the data in a DR event .....	62
2.2.15	Cleaning up mirrors after a DR event .....	66
2.3	Geographic Logical Volume Manager (GLVM) Replication with PowerHA .....	68
2.3.1	Preparing AIX GLVM for IBM PowerHA EE .....	68
2.3.2	Update VGs to enhanced concurrent capable .....	70
2.3.3	Change filesystems to not automount .....	71
2.3.4	Determine for which drives to create RPV servers .....	72
2.3.5	Bring the VGs offline .....	72
2.3.6	Create the RPV servers on the primary location .....	73
2.3.7	Create the RPV clients on the DR location .....	75
2.3.8	Verify the configuration .....	77
2.3.9	Import all VGs to remote PowerVS .....	78
2.3.10	Installing IBM PowerHA EE .....	78
2.3.11	Configuring PowerHA EE .....	81
2.3.12	Verify /etc/cluster/rhosts entries .....	82
2.3.13	Determine the PowerHA repository disk(s) .....	83
2.3.14	Configure PowerHA topology .....	84
2.3.15	Configure PowerHA repository disks .....	86
2.3.16	Perform initial PowerHA verification .....	88
2.3.17	Defining an XD_data network .....	90
2.3.18	Moving network interfaces to XD_data network .....	91

2.3.19 Defining PowerHA resource group .....	93
2.3.20 Adding resources to the PowerHA resource group .....	95
2.3.21 Perform an additional PowerHA verification .....	97
2.3.22 Starting PowerHA EE .....	99
2.3.23 Verifying PowerHA EE is operating normally .....	100
2.3.24 Testing PowerHA .....	102

## **CHAPTER 3: TROUBLESHOOTING..... 104**

3.1 ioctl call failed .....	104
3.2 Resuming a failed RPV client.....	105
3.3 When configuring RPV server, available disks are not displayed .....	105

## **CHAPTER 4: ADDITIONAL RESOURCES ..... 108**

4.1 AIX Geographic Logical Volume Manager (GLVM) resources.....	108
4.2 IBM PowerHA resources .....	108
4.3 Other resources .....	108

# Chapter 1: Solution Overview

## 1.1 Introduction

Uptime is a key client expectation for AIX 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 AIX 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 disaster recovery (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 AIX workloads in PowerVS:

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

## 1.2 Use Cases

### 1.2.1 Geographic Logical Volume Manager (GLVM) Replication

In this situation, we use IBM AIX GLVM functionality to replicate data from one IBM Cloud location to another IBM Cloud location for DR purposes. This gives one the capability of recovering their application

at a secondary IBM Cloud location should the primary IBM Cloud location be inaccessible for whatever reason.

### **1.2.2 Geographic Logical Volume Manager (GLVM) Replication with PowerHA**

In this situation, we combine the IBM AIX GLVM functionality to replicate data from one IBM Cloud location to another IBM Cloud location with IBM PowerHA to provide a fully automated resiliency solution. This gives one the capability of automatically recovering their application at a secondary IBM Cloud location should the primary IBM Cloud location be inaccessible for whatever reason.

## **1.3 Solution Components and Requirements**

### **1.3.1 Components**

The following components will be required for these solutions:

1. *IBM Power Virtual Servers*
2. *Multiple IBM Cloud locations*
3. *IBM network setup between IBM Cloud locations*
4. *Storage assigned to appropriate VMs*

### **1.3.2 Requirements**

The following will be required for these solutions:

1. *IBM Power Virtual Servers with AIX operating system at multiple IBM Cloud locations*
2. *IBM network setup between IBM Cloud locations*
  - a. *The network setup is described in the **Chapter 2: Implementation** section of this document*
3. *AIX GLVM installed on each AIX VM*
  - a. *AIX GLVM comes with AIX, so it is not an additional cost or license*

4. For integration with IBM PowerHA, the IBM PowerHA Enterprise Edition software installed on each of the AIX VMs that will be part of this environment (high-availability cluster)
  - a. IBM PowerHA is a separately orderable product and requires an additional licensing

## 1.4 Solution Diagram

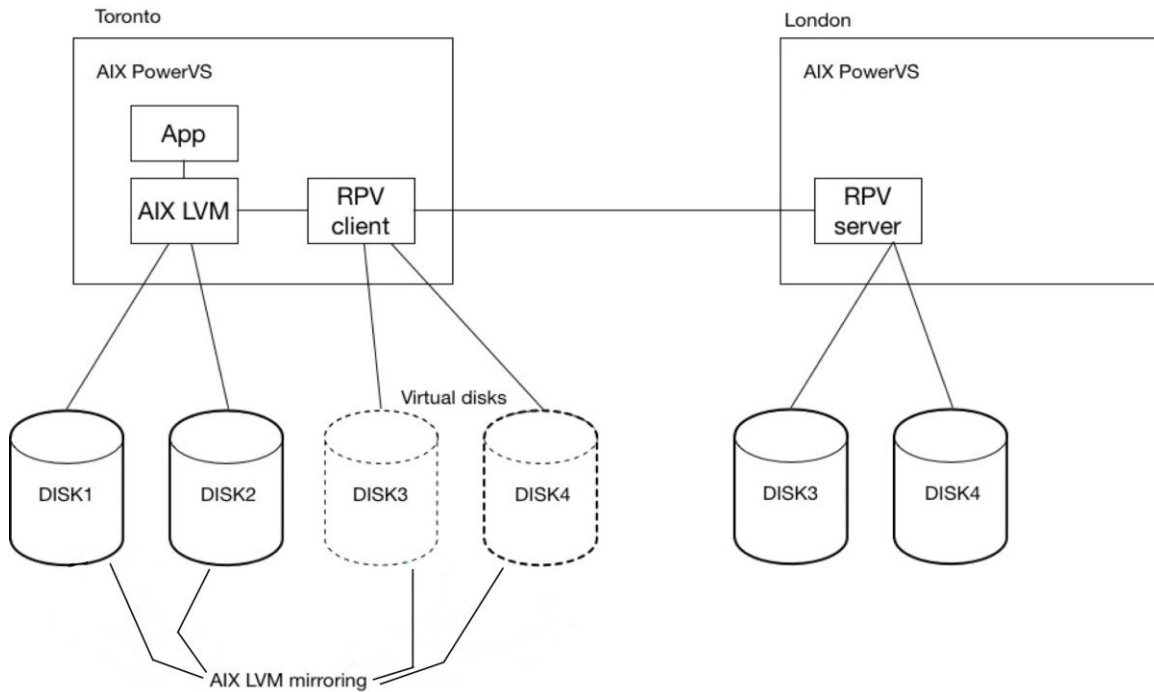


Figure 1.4.1 RPV server/client diagram with virtual drives



# Chapter 2: Implementation

## 2.1 Base IBM Cloud PowerVSI and Networking setup

The following steps are required to implement IBM PowerVS VMs and setup the network between IBM Cloud locations:

1. *Open an IBM Cloud account*
2. *Create two Power PowerVS location Services and a private subnet in each PowerVS location.*
3. *Provision AIX VSIs in each PowerVS location*
4. *Order Direct Link Connect Classic to connect each PowerVS location to IBM Cloud*
5. *Order two Vyatta Gateways one in each datacenter: Lon06 and Tor01 datacenters or your chosen datacenters to allow for PowerVS location-to-PowerVS 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 PowerVS location communication for the VSIs in the PowerVS locations and to the IBM Cloud VSIs and other services such as Cloud Object Storage (COS).*

### 2.1.1 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.

### **2.1.2 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 Power AIX and 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.

### **2.1.3 Provision AIX and IBM i VSIs in each PowerVS location**

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

### **2.1.4 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.

### **2.1.5 Order two Vyatta Gateways, one in each datacenter**

In order to setup communication between the two PowerVS location datacenters, you will need to use a 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.

We ordered one Vyatta in LON06 and the other in TOR01 datacenters where our PowerVS locations exists.

### **2.1.6 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.

### **2.1.7 Configure three GRE tunnels in the Vyatta Gateways**

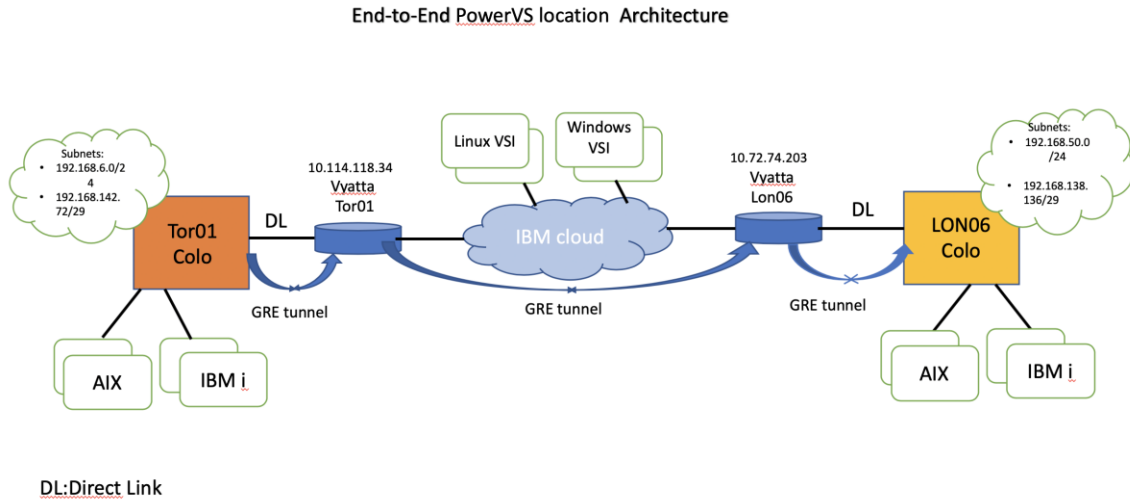
After the support team finished configuring the GRE tunnel, you will need to configure your end of the GRE tunnel on the two Vyatta Gateways.

You will need three GRE tunnels

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

## 2.1.8 Diagrams

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




Screenshot

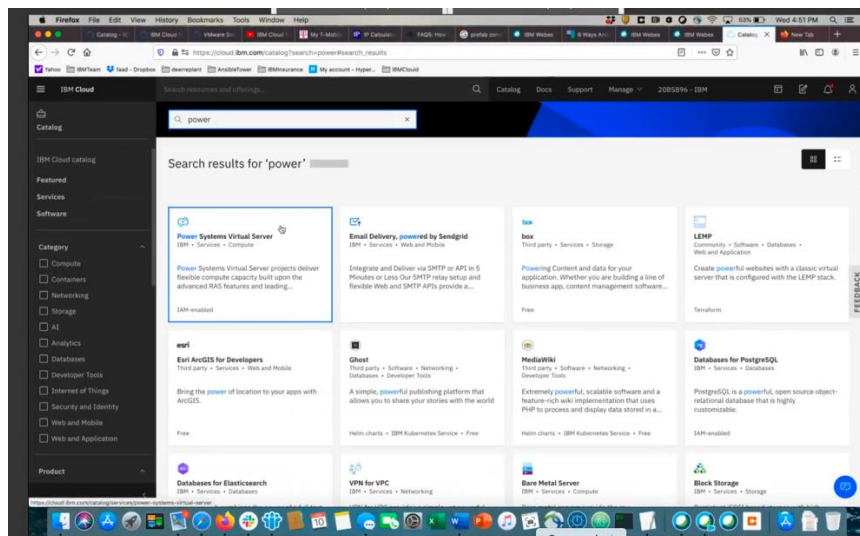
## 2.1.9 Create PowerVS location Services and Subnet(s)

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

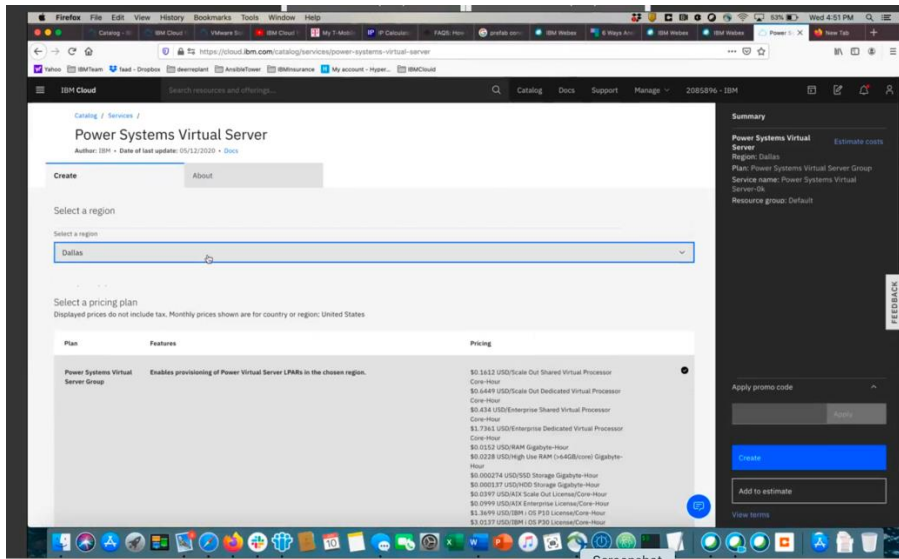
Go to the main IBM Cloud UI page and click on "Catalog" on upper right side of the UI.



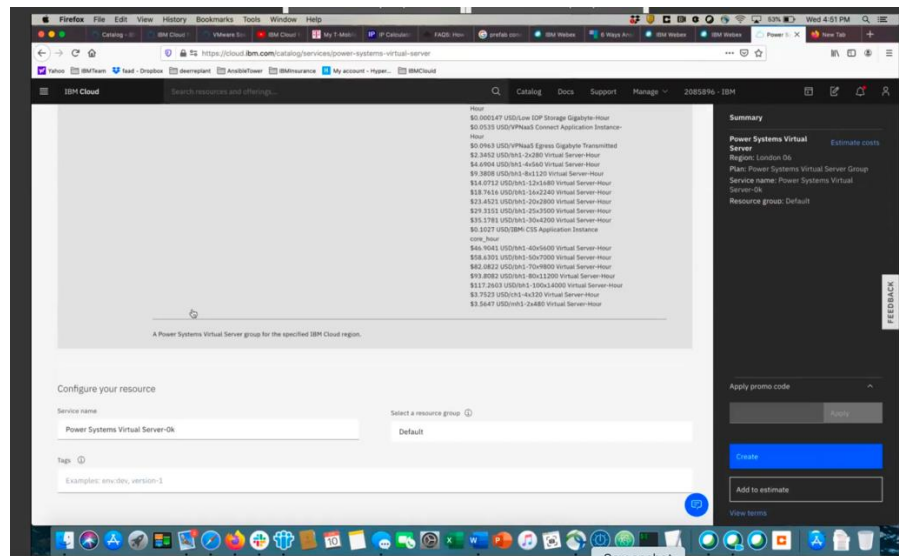
Search for "Power"



Select "Power System Virtual Servers".

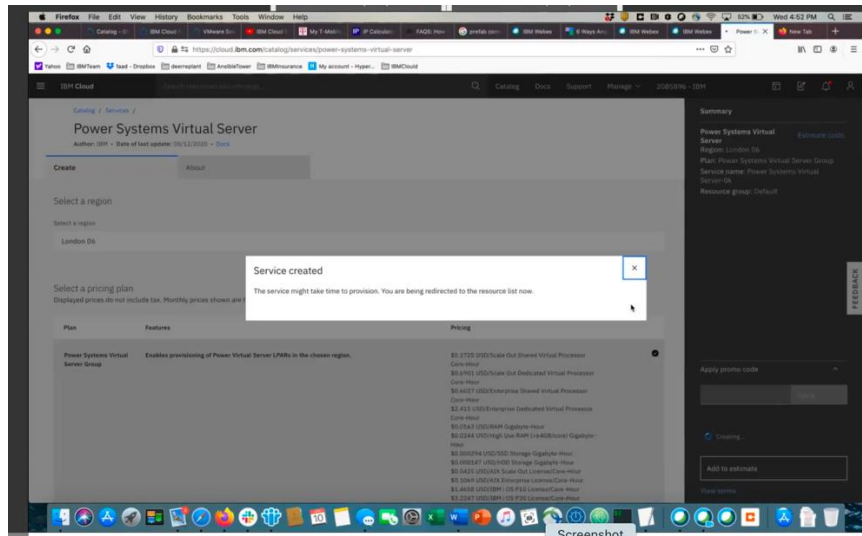


Under Select Region, choose your region. You are limited to only one service per region.



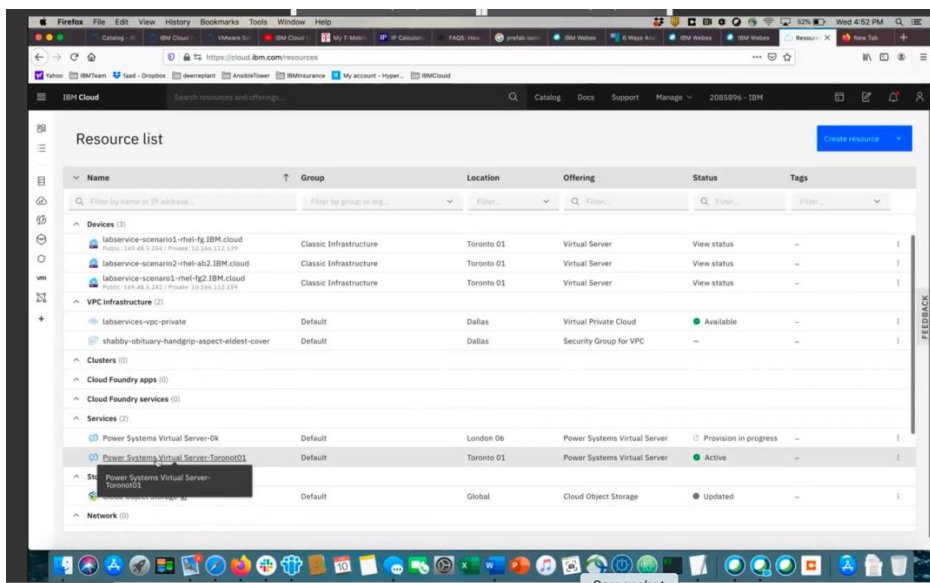
Select a "Service Name" or chose default name provided.

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 our case we have 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.

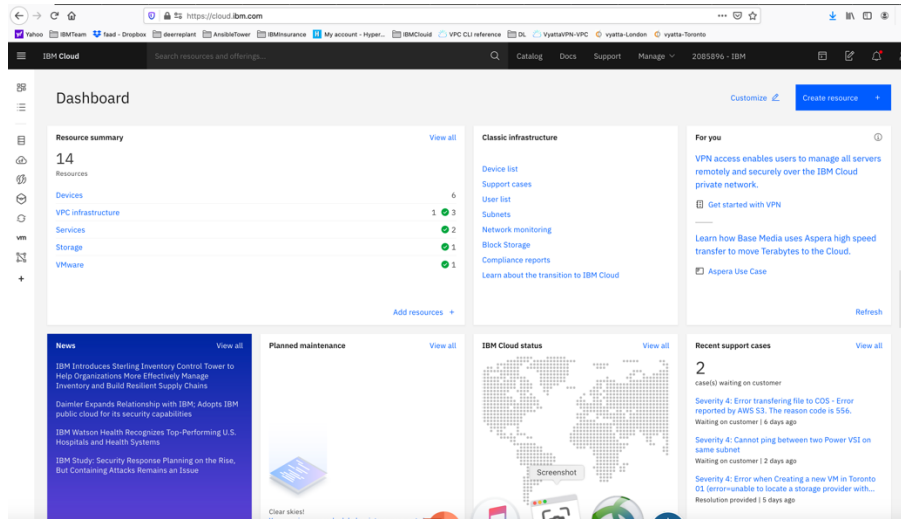
At this point, you will need to open a Support Ticket with Power System to request that the subnet be configured to allow local communication between any Power VSI you create in this PowerVS location service. Provide your PowerVS location service location, and your subnet in the ticket.

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.

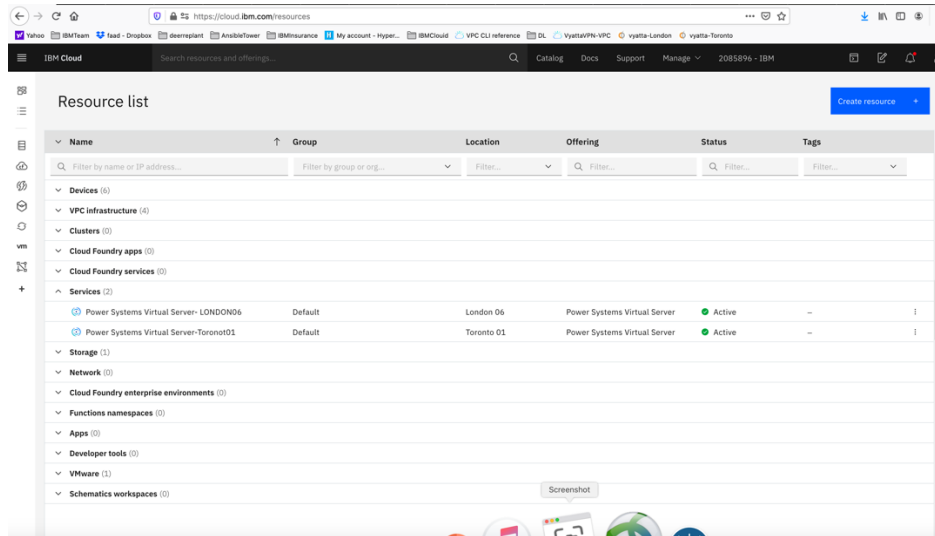
### **2.1.10 Provision AIX and IBM i VSIs in each PowerVS location**

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.

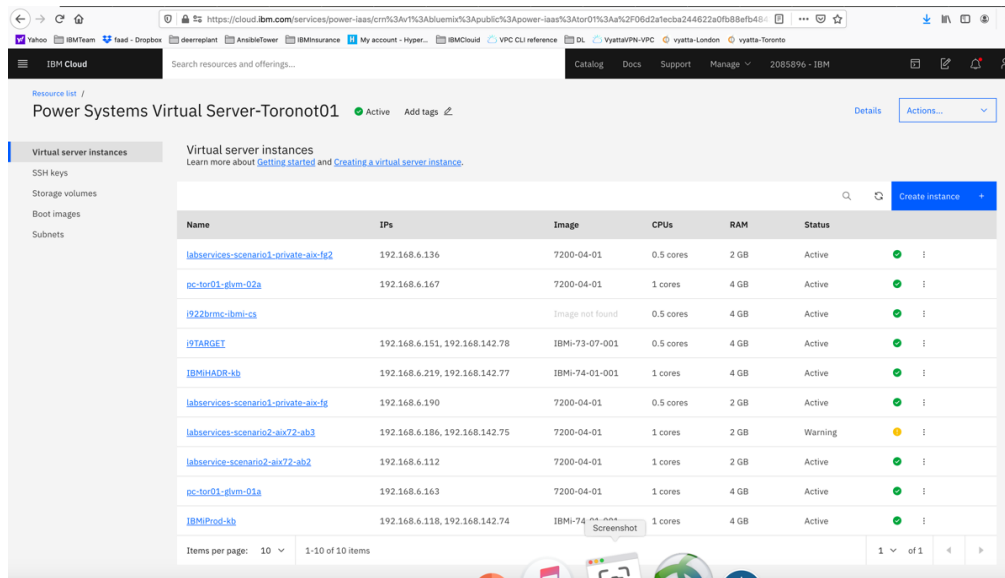
Go to the IBM Cloud Catalog and press the "IBM Cloud" on top left side of the UI.



Choose "Services" from the list shown.

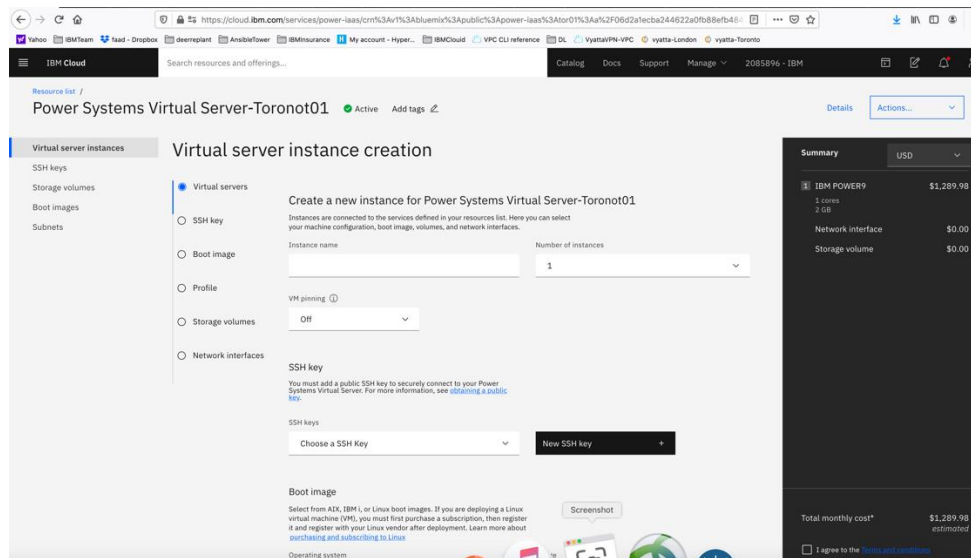


Click on the service for datacenter in which you have created a PowerVS location power service. In this case we will choose Toronot01 PowerVS location service.



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

Press "Create Instance" on upper right-hand side.

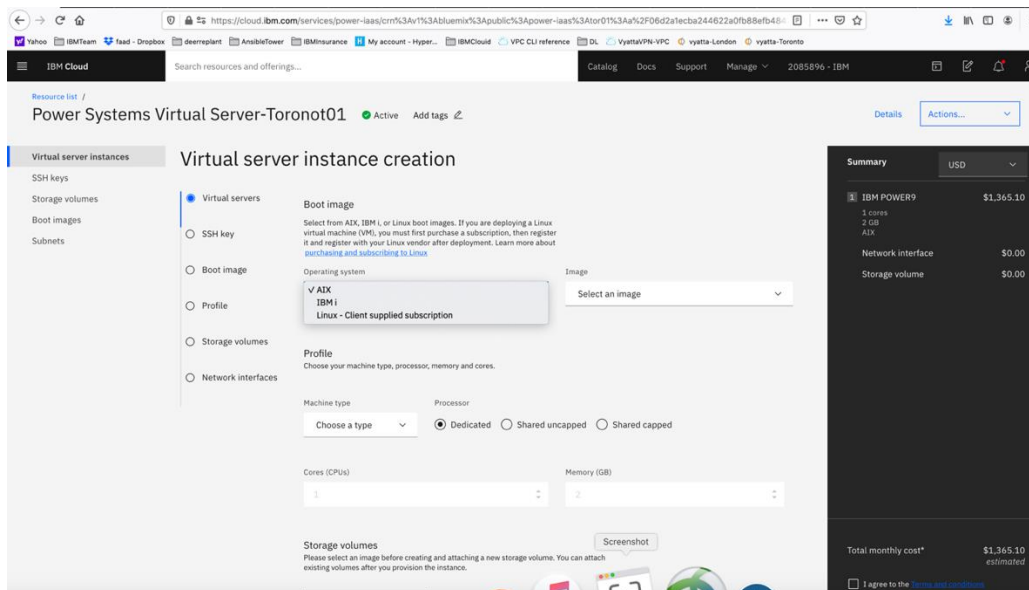


This is where you provision AIX or IBM i VSIs.

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.

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.

Scroll down to choose other options.

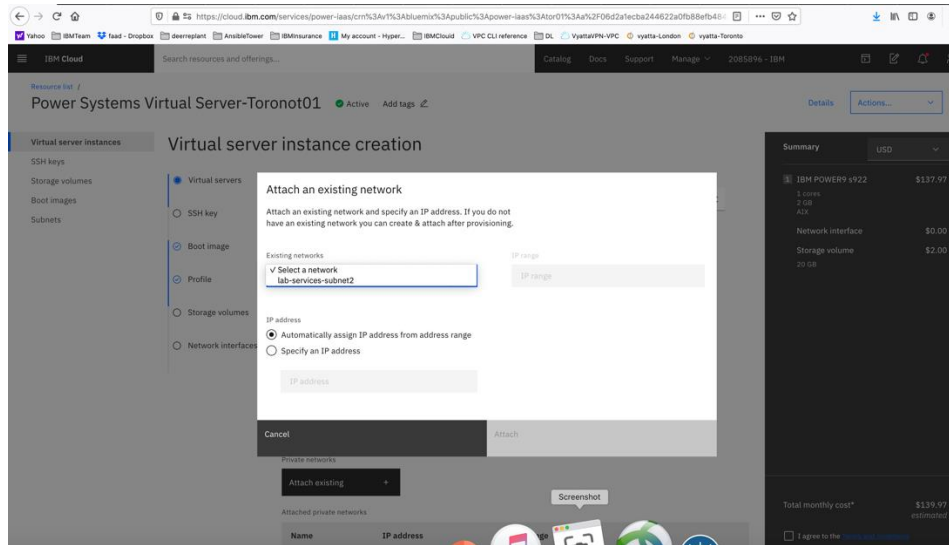


Here you will choose the following options:

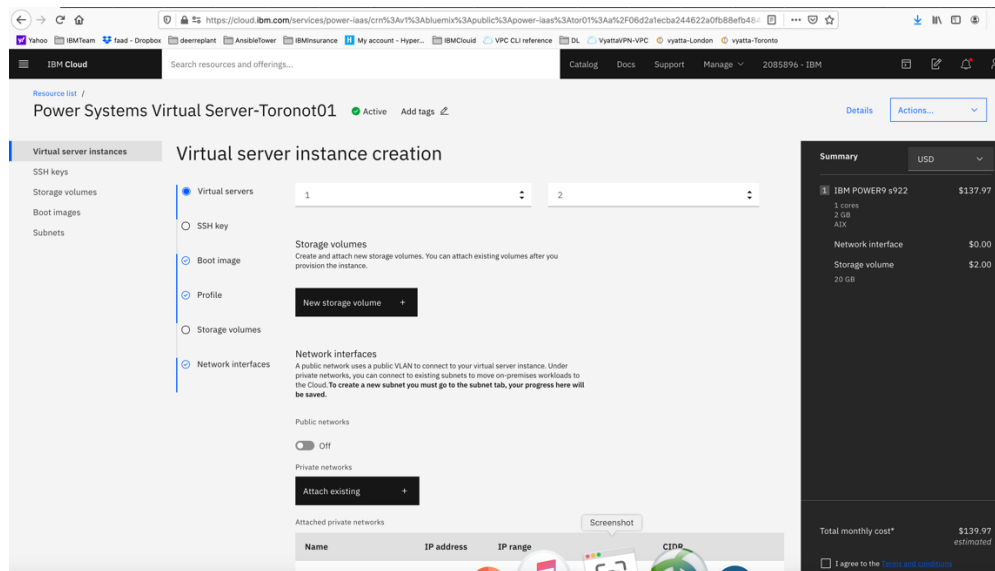
- *Operating System – AIX*
- *Image type: AIX 7.1 or 7.2, etc.*
- *Disk types: Type 1 or 3. Type 3 is a less expensive option which we selected.*
- *Machine type: S922 or E980. S922 is the cheaper of the two which we selected.*
- *Processor: Dedicated or Shared or Shared Capped. We choose "shared" as its less expensive.*
- *Choose the number of cores and RAM you will need. The minimum core is "0.25".*
- *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.*

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 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.

### **2.1.11 Order Direct Link Connect Classic to connect PowerVS location to IBM Cloud**

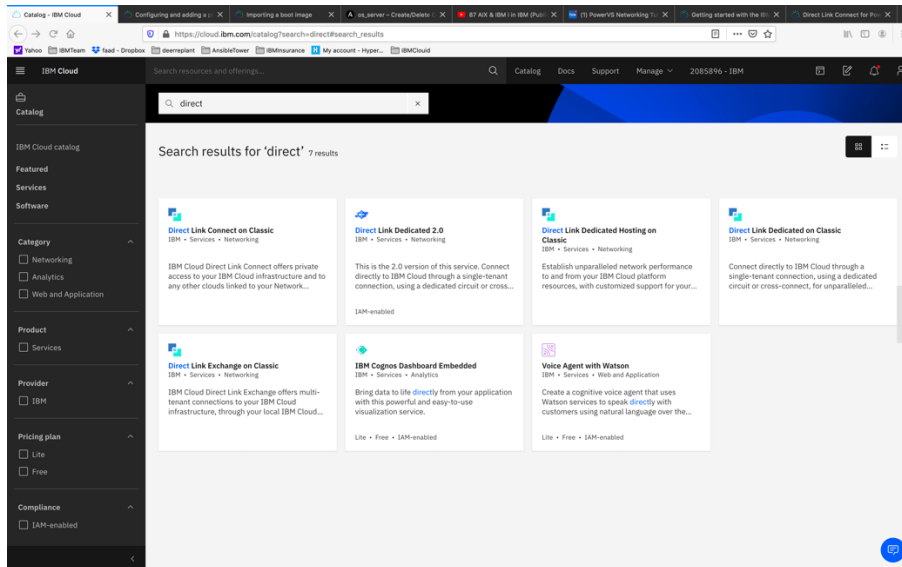
You will need to order Direct Link (DL) Connect Classic to allow your Power VSIs in the PowerVS location to communication with Linux/Window VSIs in IBM Cloud and also with all other IBM Cloud services such as Cloud Object Storage and VMware services. This process may take 1-2 weeks to complete.

There are several steps involved in completing DL ordering:

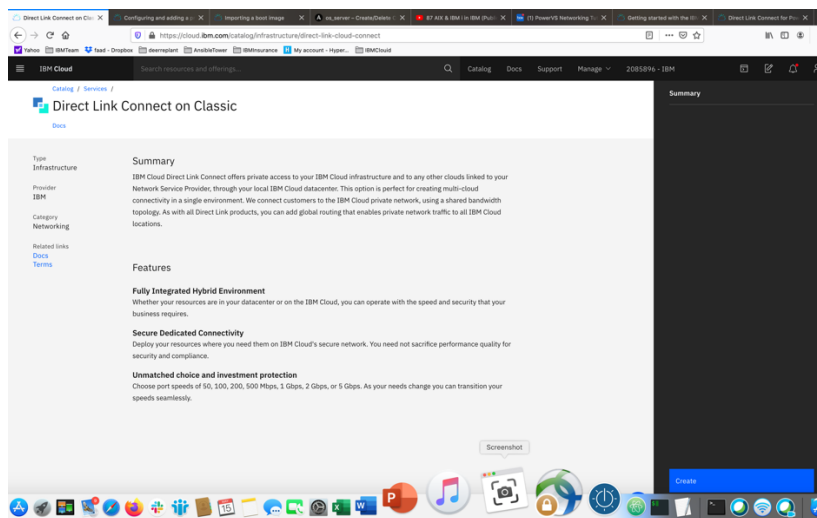
- *Order Direct link connect classic service on IBM Cloud UI – see steps below*
- *Next a support ticket will be created, and Support will send you a word document with questionnaires to be completed concerning various DL settings.*
- *Complete the questionnaires and upload it to support in the ticket.*
- *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.

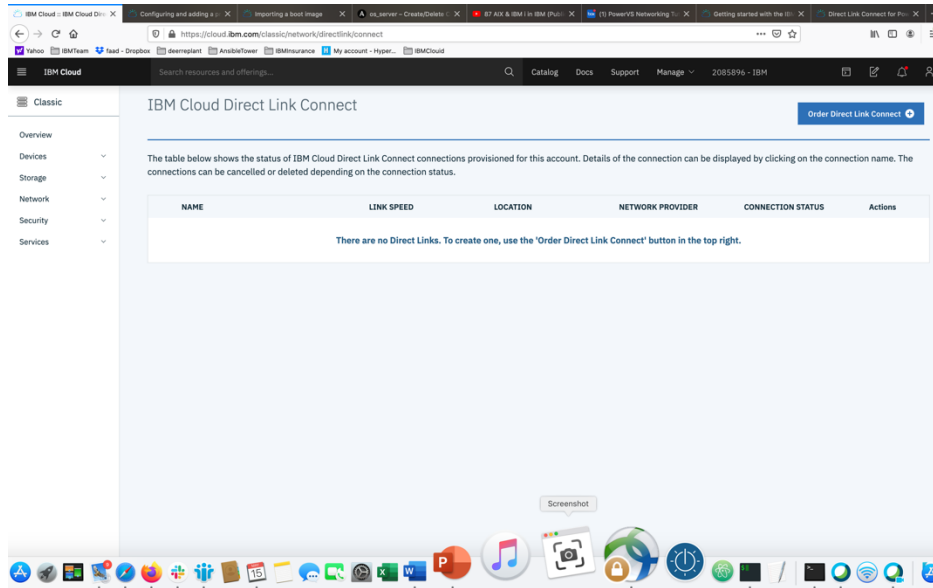
Choose “Catalog” from upper right-hand side, and search for “direct”.



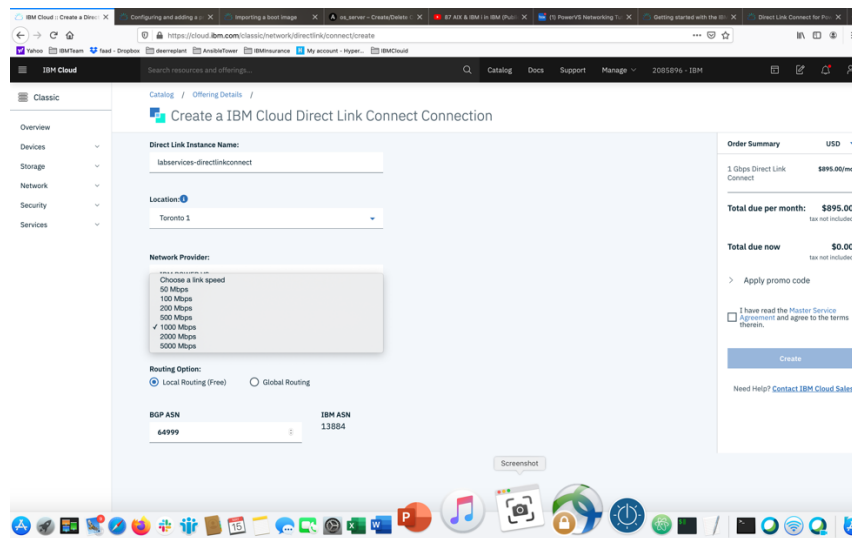
Select "Direct Link Connect on Classic".



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



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



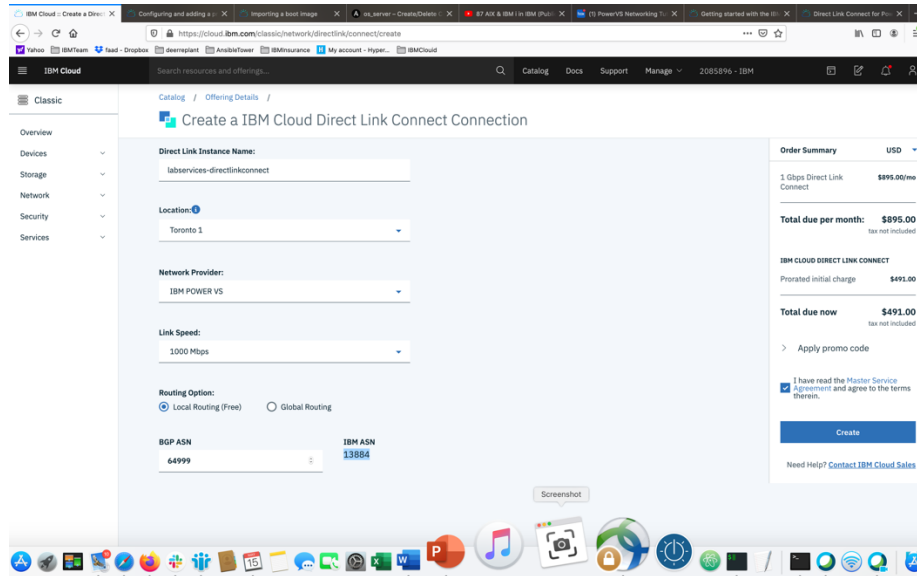
- Choose a "name" for the DL.
- Choose a location for the DL. This should be the same location as where you created your PowerVS location Service.
- Choose "link speed" under network provider menu.
- Choose "Local Routing (free)"

Global routing will require additional charges and will allow for easier PowerVS location-to-PowerVS location communication. You will also



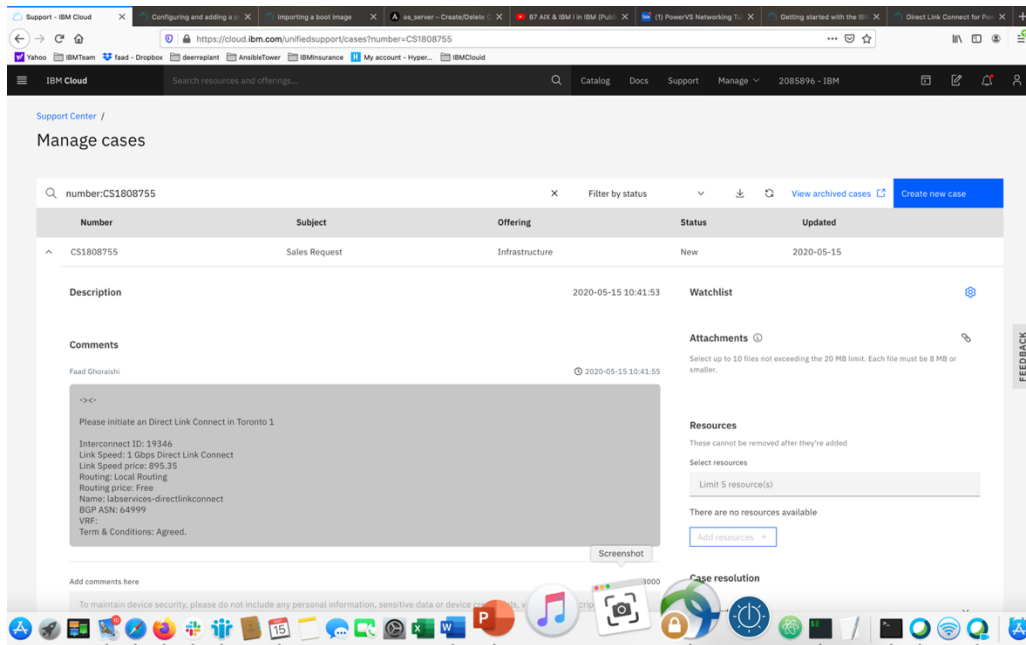
need to order a Vyatta Gateway Router to complete your Global routing option via use of a GRE tunnel. Support can help you with this further.

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.*

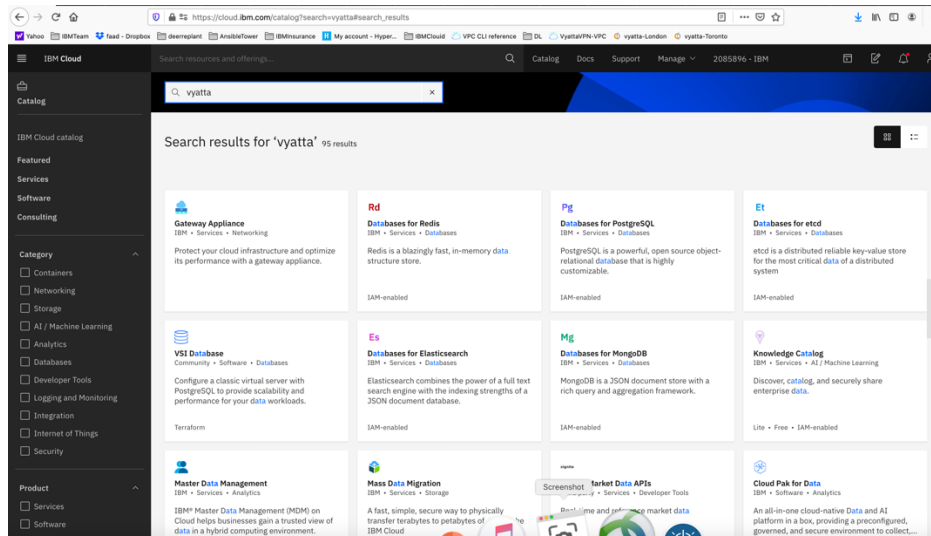
*After this step is complete, support will request that you open a new IBM support ticket and address it to the Power System. Include the information in the original DL ticket. This new ticket will be sent to the PowerVS location support to configure their side of the DL connection.*

*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.*

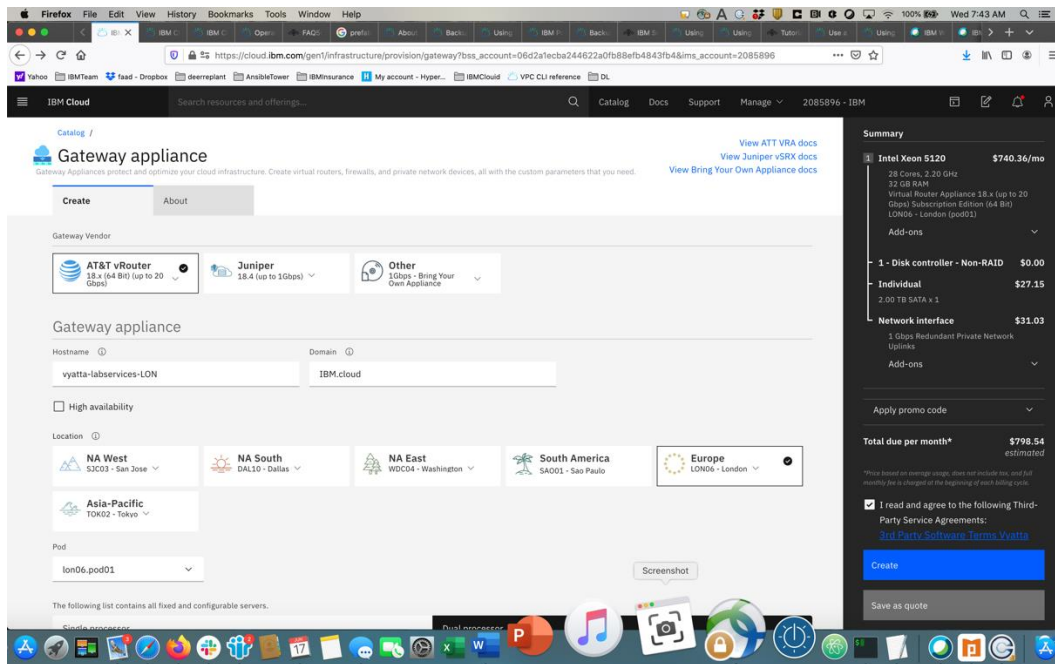
## **2.1.12 Order two Vyatta Gateways, one in each datacenter**

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.



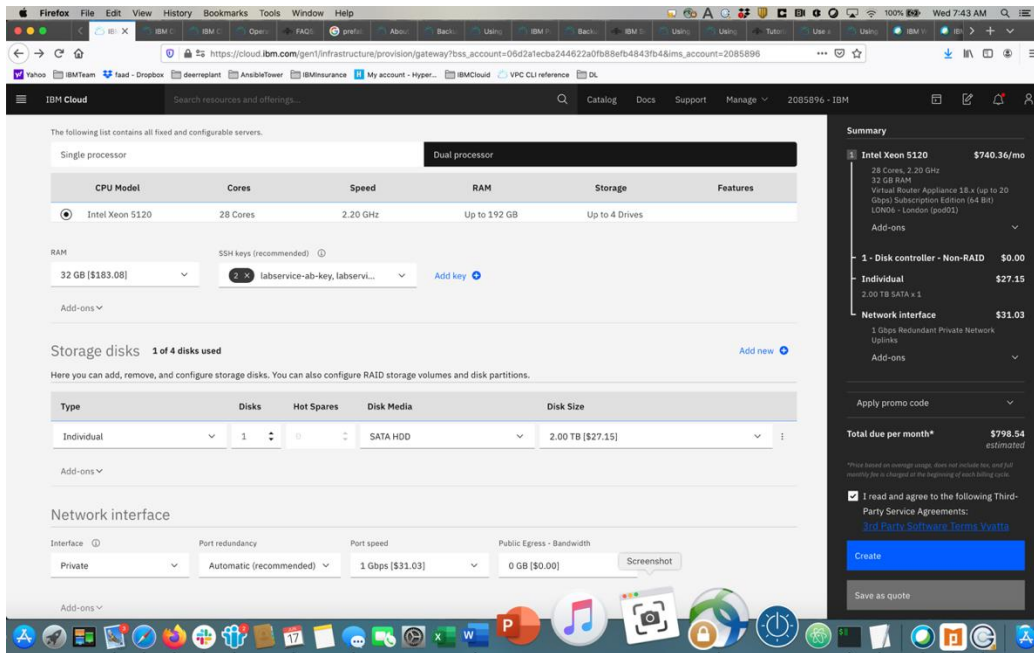
Select "Gateway Appliance" and click on it.



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 uncheck this option.*
- *Select the location by pressing on the arrow key in each location to find the exact datacenter where your PowerVS locations 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.*

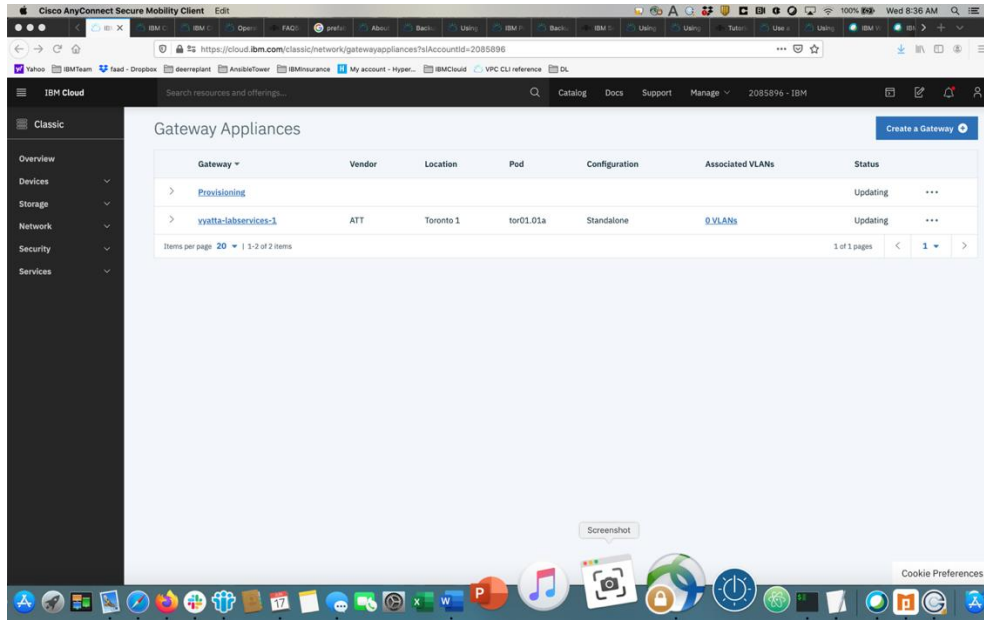


Now check the box to agree with service agreement on the bottom-right side and press “Create”

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.

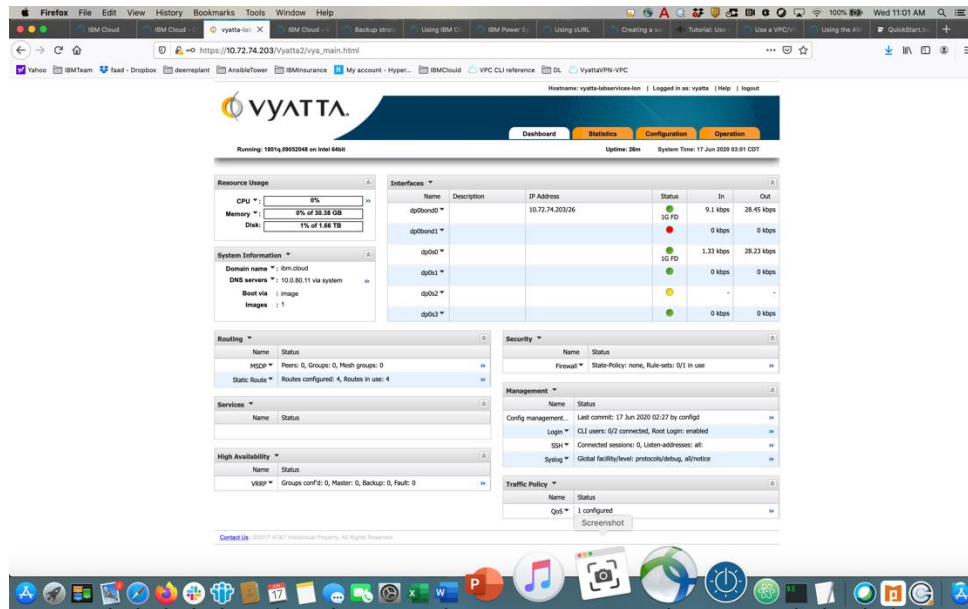


To log into the Vyatta gateway, use a browser and access it via the link:

<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.

### **2.1.13 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.

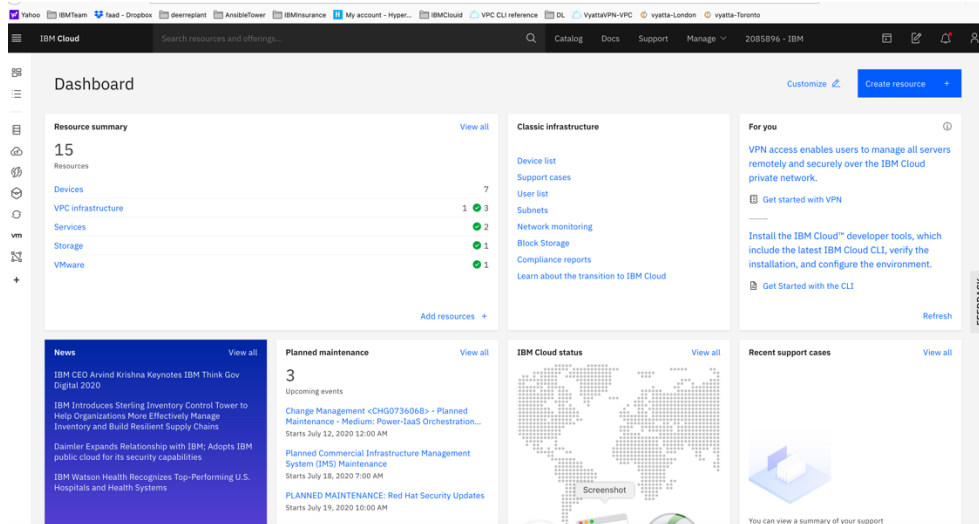
- *Note that your tunnel IP address is 172.20.2.1/30 where 255.255.255.252 translate to /30*
- *Your tunnel destination IP is their tunnel source IP.*
- *Your tunnel source IP is the IP address of the Vyatta gateway*

Verify your Vyatta Gateway access.

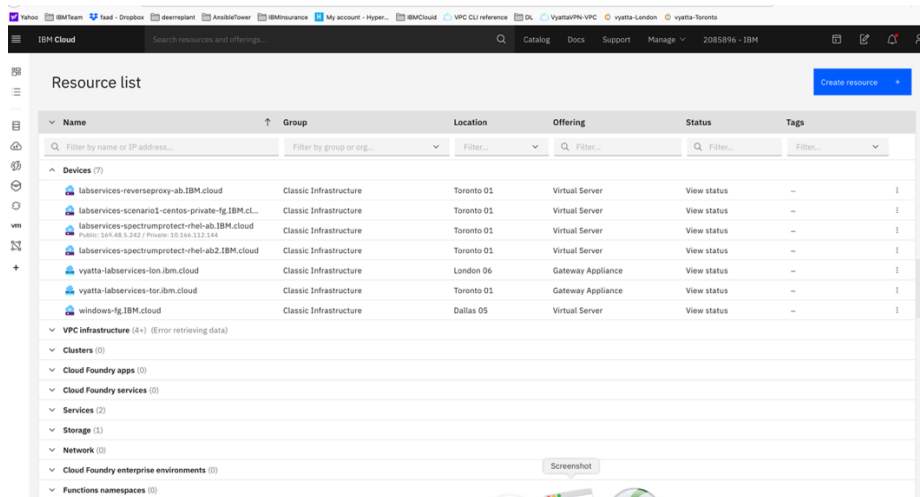
The Vyatta Gateway address can be find in the IBM Cloud UI under Devices.

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





Click on "Devices"

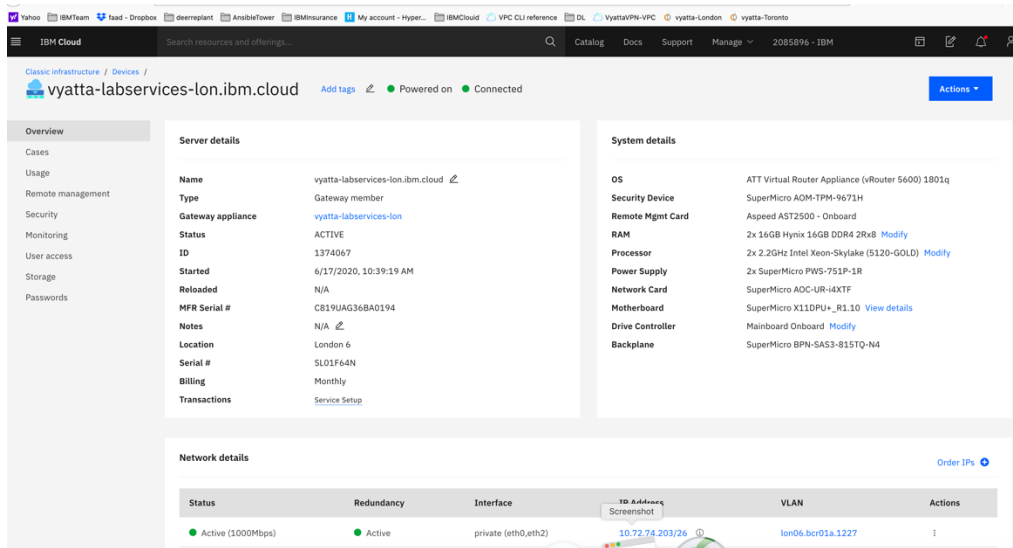


Choose the Vyatta system you like to configure:

- [vyatta-labservices-lon.ibm.cloud](https://vyatta-labservices-lon.ibm.cloud)
- [vyatta-labservices-tor.ibm.cloud](https://vyatta-labservices-tor.ibm.cloud)

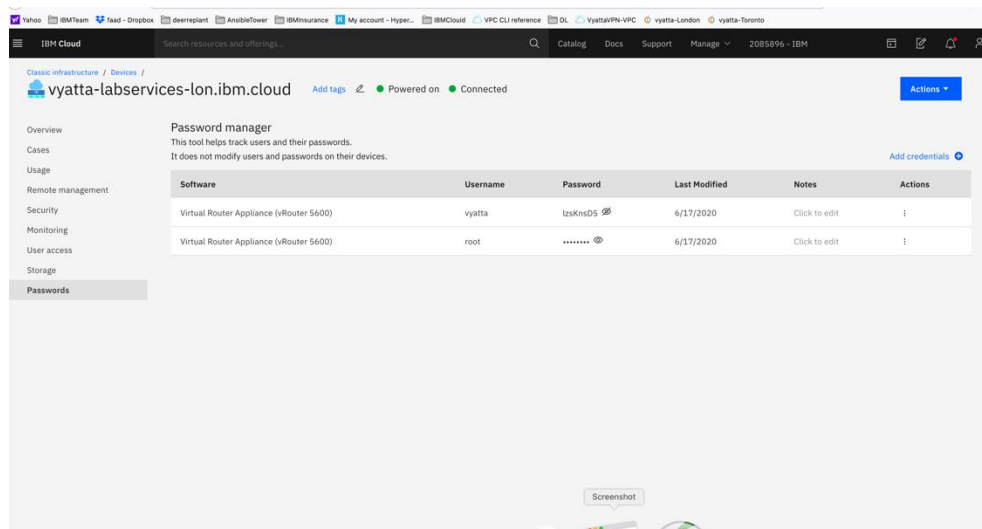
LON06:

Click on the London Vyatta: [vyatta-labservices-lon.ibm.cloud](https://vyatta-labservices-lon.ibm.cloud)



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

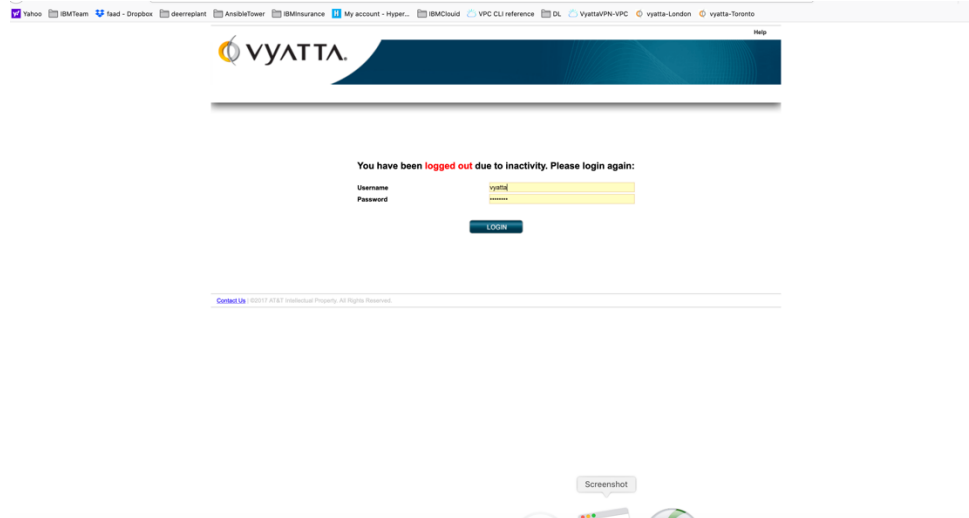
Your credentials are under the “password” menu on the left-hand side. Click on the icon next to the password to see it unencrypted.




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.

## 2.1.14 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://docs.huihoo.com/vyatta/6.5/Vyatta-Tunnels_6.5R1_v01.pdf)

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](mailto: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*

### **2.1.15 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:~ faadghoraishi$ 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
}

```

### **2.1.16 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 choose 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*

```

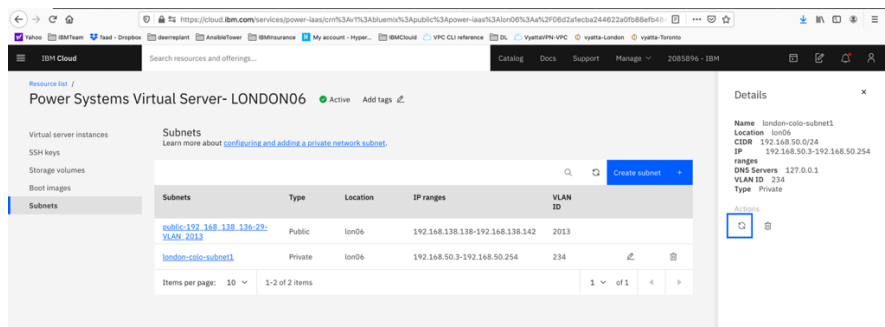
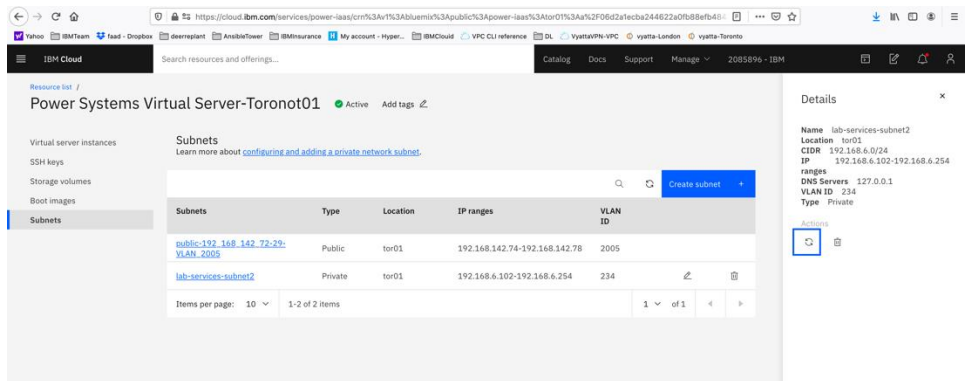
vyatta@vyatta-labservices-lon# show interfaces tunnel tun1
tunnel tun1 {
  address 172.20.4.2/30
  encapsulation gre
  local-ip 10.72.74.203
  mtu 1300
  remote-ip 10.114.118.34
}
[edit]
vyatta@vyatta-labservices-lon#

```

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

### 2.1.17 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 vis 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.

## 2.2 Geographic Logical Volume Manager (GLVM) Replication

AIX Geographic Logical Volume Manager (GLVM) is a software-based technology for real time geographic data mirroring over standard TCP/IP networks. AIX GLVM can help protect your business from a disaster by mirroring your mission-critical data to a remote disaster recovery site.

Using AIX GLVM, we will mirror the data of an IBM Power Virtual Server image from one data center to another. We will then simulate a DR event and recover that data at the remote data center.

### 2.2.1 Basic Concepts

We will use a simple configuration to describe the basic concepts of GLVM. We will use an example configuration for this tutorial. Consider an IBM AIX image that is running a business application. All of the data resides on two disks, DISK1 and DISK2. Each disk contains exactly half of the data. In other words, there is only one copy of the data which spans both disks. This configuration is shown in the following diagram:

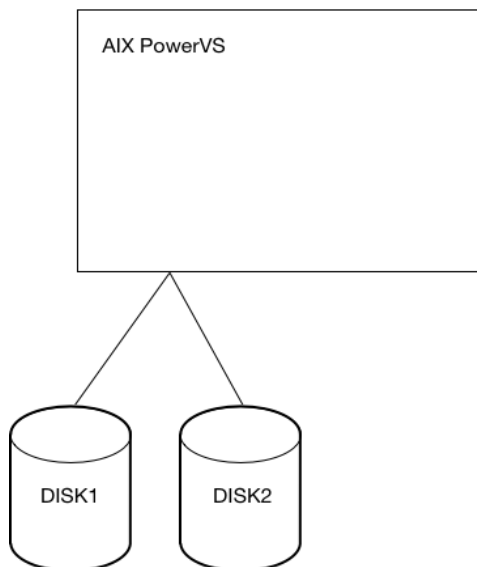


Figure 2.2.1.1 Typical AIX configuration

With only one copy of the data, a single storage failure would cause this environment to fail. To make this environment more redundant, we need to make a copy of the data.

AIX GLVM helps you to solve this problem by allowing a complete mirror to exist in a geographically distant location, the DR site. For this we will need two servers, one at the production site and another at the disaster recovery site and each site needs a complete mirror copy of the data. Both servers are connected by a TCP/IP network. This can be any type of TCP/IP network. The production and disaster recovery sites do not have to be on the same physical network. Routers and gateways between the two sites are allowed. For our example, the primary site will be in Toronto, CA and the DR site will be in London, UK. This is shown in the following diagram:

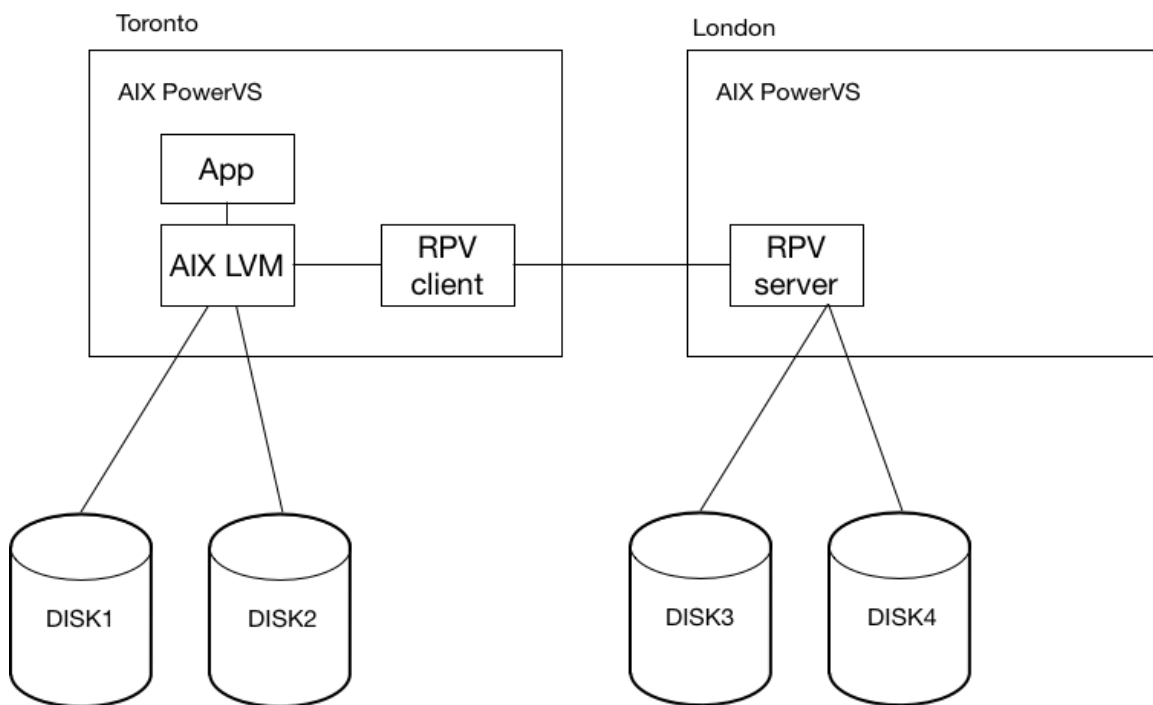


Figure 2.2.1.2 RPV server with disks

In this configuration, we can take the same volume group, containing the business application described earlier and stretch it across both sites. This is now a geographically mirrored volume group.

In the above diagram, the volume group will be varied on at the primary site and no replication is occurring. The Remote Physical Volume (RPV) device driver will allow AIX LVM at the production site to access the disks at the disaster recovery site as if they were locally attached. This will be accomplished by defining disks DISK3 and DISK4 to the production site as remote physical volumes.

There are two components to the RPV device driver, the RPV client and the RPV server. The RPV client will reside on the production site. The RPV server will reside on the DR site. The RPV server will process read/write requests from a the remote RPV client. Both RPV cdevice driver components work together to enable AIX LVM at the production site to access the DR disk(s).

Here is how it will work:

- The RPV server will present a disk from the DR site to the RPV client on the production site
- The production site will see a virtual disk that acts like a real disk and any writes to that disk go from the production site to the DR site.
- The production site will see a virtual disk that acts like a real disk and any writes to that disk go from the production site to the DR site
- Because the production site now sees extra disks, we can use regular IBM LVM processes to copy data from the local production disks to the remote DR disks

At this point, if a disaster were to destroy the primary site, the DR site would contain a copy of that data.

See the following diagram.

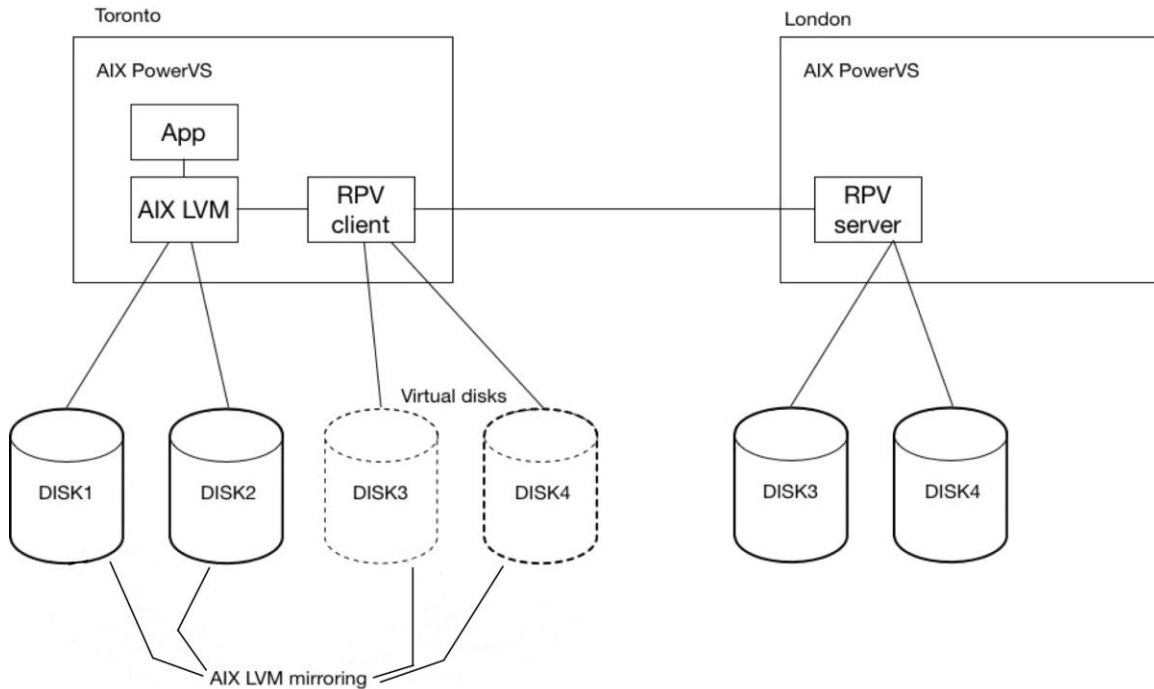


Figure 2.2.1.3 RLV server/client with virtual disks

In our example environment, we will create virtual disks in both directions. It was only drawn in one direction above for clarity.

## 2.2.2 Example environment

For our discussion, we will use the following configuration:

- *Single AIX PowerVS in Toronto location*
- *Single AIX PowerVS in London location*
- *Storage (tier 3, 6-20GB LUNs) allocated to each AIX PowerVS*
- *Communication between IBM Cloud locations*
- *Setting LV strictness*

## 2.2.3 Preparing for IBM AIX GLVM setup

For our discussion, we will use the following configuration:




Figure 2.2.3.1 RLV server/client with virtual disks

In our example scenario, we have implemented an application with a couple of volume groups, `dbvg` and `appvg`. Those volume groups have 2 disk drives each.

We will prepare this environment for IBM GLVM replication by performing the following steps:

- *Change or verify the volume groups are a scalable volume group*
- *Updating the volume groups so as to not vary on automatically at reboot time*
- *Updating the volume groups by turning off bad block relocation*
- *Updating volume groups by setting super strict mirror pools*
- *Defining mirror pools*

## 2.2.4 Change or verify the volume groups are a scalable volume group

For our example, we did create the volume groups as scalable volume groups. This can be verified using the `lspv` and `readvgda` commands:



```
# lspv | grep dbvg
hdisk1      00c8d4403180f273      dbvg      active
hdisk2      00c8d4403180f319      dbvg      active
# readvgda hdisk2 | grep type
.....      readvgda type: svg
type:       jfs2log
type:       jfs2
type:       jfs2
type:       jfs2
type:       jfs2
#
```

Figure 2.2.4.1 Using “readvgda” to find VG type

For our example, we did create the volume groups as scalable volume. If the volume groups were not scalable, they could be changed to scalable volume groups via the following command: *chvg -G <VGNAME>* where *<VGNAME>* is the name of the volume group. Note that the volume group needs to be varied off when this *chvg* command is executed.

### **2.2.5 Updating the volume groups so as to not vary on automatically at reboot time**

We can update the volume group, so it does not automatically vary on at boot time using the following command: *chvg -a n <VGNAME>* where *<VGNAME>* is the name of the volume group. See the output of our command execution:

```

# chvg -an dbvg
# chvg -a n appvg
# lsvg dbvg
VOLUME GROUP:          dbvg                VG IDENTIFIER: 00c8d44000004b00000
001733180f36e
VG STATE:              active              PP SIZE:       16 megabyte(s)
VG PERMISSION:        read/write          TOTAL PPs:     2550 (40800 megabyt
es)
MAX LVs:              256                 FREE PPs:     2037 (32592 megabyt
es)
LVs:                  5                   USED PPs:     513 (8208 megabytes
)
OPEN LVs:             5                   QUORUM:       2 (Enabled)
TOTAL PVs:            2                   VG DESCRIPTORS: 3
STALE PVs:            0                   STALE PPs:    0
ACTIVE PVs:           2                   AUTO ON:      no
MAX PPs per VG:      32768                MAX PVs:      1024
LTG size (Dynamic):  512 kilobyte(s)      AUTO SYNC:    no
HOT SPARE:            no                   BB POLICY:    relocatable
MIRROR POOL STRICT:  off
PV RESTRICTION:      none                INFINITE RETRY: no
DISK BLOCK SIZE:     512                 CRITICAL VG:  no
FS SYNC OPTION:      no                   CRITICAL PVs: no
#

```

Figure 2.2.5.1 Status of VG auto vary on field

Note the "AUTO ON: no" field on the lower right indicating this volume group will not automatically vary on at boot time.

## 2.2.6 Updating the volume groups by turning off bad block relocation

We can update the volume group so that black block relocation is disabled. By default, if a bad block is detected, that block is relocated to a different location on the disk. This would create dissimilar block maps between replicated LUNs, so it needs to be disabled.

Bad block relocation can be disabled via the following command: `chvg -b n <VGNAME>` where `<VGNAME>` is the name of the volume group. See the output of our command execution:

```

# chvg -b n dbvg
# lsvg dbvg
VOLUME GROUP:      dbvg                VG IDENTIFIER: 00c8d44000004b00000
001733180f36e
VG STATE:          active              PP SIZE:       16 megabyte(s)
VG PERMISSION:    read/write          TOTAL PPs:     2550 (40800 megabyt
es)
MAX LVs:          256                  FREE PPs:      2037 (32592 megabyt
es)
LVs:              5                    USED PPs:      513 (8208 megabytes
)
OPEN LVs:         5                    QUORUM:        2 (Enabled)
TOTAL PVs:        2                    VG DESCRIPTORS: 3
STALE PVs:        0                    STALE PPs:     0
ACTIVE PVs:       2                    AUTO ON:       no
MAX PPs per VG:   32768                 MAX PVs:       1024
LTG size (Dynamic): 512 kilobyte(s)    AUTO SYNC:     no
HOT SPARE:        no                    BB POLICY:     non-relocatable
MIRROR POOL STRICT: off
PV RESTRICTION:   none                  INFINITE RETRY: no
DISK BLOCK SIZE:  512                   CRITICAL VG:   no
FS SYNC OPTION:   no                    CRITICAL PVs:  no
#

```

Figure 2.2.6.1 Status of VG relocatable field

Note the “BB POLICY: non-relocatable” field on the lower right indicating bad block relocation is disabled.

## 2.2.7 Updating volume groups by setting super strict mirror pools

As mentioned previously, it is critical that we can guarantee one copy of the data is local and a secondary copy of the data is remote. We can do this by setting super strict mirror pools.

Super strict mirror pools can be set via the following command: `chvg -M s <VGNAME>` where `<VGNAME>` is the name of the volume group. See the output of our command execution:

```

# chvg -M s dbvg
# lsvg dbvg
VOLUME GROUP:          dbvg                VG IDENTIFIER: 00c8d44000004b00000
001733180f36e
VG STATE:              active              PP SIZE:       16 megabyte(s)
VG PERMISSION:        read/write          TOTAL PPs:     2550 (40800 megabyt
es)
MAX LVs:              256                 FREE PPs:      2037 (32592 megabyt
es)
LVs:                  5                   USED PPs:      513 (8208 megabyte
s)
OPEN LVs:             5                   QUORUM:        2 (Enabled)
TOTAL PVs:            2                   VG DESCRIPTORS: 3
STALE PVs:            0                   STALE PPs:     0
ACTIVE PVs:           2                   AUTO ON:       no
MAX PPs per VG:       32768               MAX PVs:       1024
LTG size (Dynamic):   512 kilobyte(s)     AUTO SYNC:     no
HOT SPARE:            no                  BB POLICY:     non-relocatable
MIRROR POOL STRICT:  super
PV RESTRICTION:       none                INFINITE RETRY: no
DISK BLOCK SIZE:     512                  CRITICAL VG:   no
FS SYNC OPTION:      no                  CRITICAL PVs:  no
#

```

Figure 2.2.7.1 Status of VG mirror pool strictness

Note the “MIRROR POOL STRICT: super” field on the lower left indicating super strict mirror pools is enabled.

## 2.2.8 Defining mirror pools

Once we define the volume group with super strict mirror pools, we need to define the mirror pools. Mirror pool names can follow any type of convention, so it is up to the customer to determine that naming convention. For our purposes, we will use the location as part of the naming convention.

Mirror pools can be set by adding physical volumes to the mirror pool. In our case, we will use the following:

- Using “lspv” to determine the disks for a VG: `lspv | grep dbvg`

```

# lspv | grep dbvg
hdisk1          00c8d4403180f273          dbvg          active
hdisk2          00c8d4403180f319          dbvg          active
#

```

Figure 2.2.8.1 Determining disks in a VG

- Using "chpv" to create the mirror pool and add the volumes to it:  
`chpv -p dbvg_tor01 hdisk1 hdisk2`

```
# chpv -p dbvg_tor01 hdisk1 hdisk2
0516-1010 chpv: Warning, the physical volume hdisk1 has open logical
          volumes. Continuing with change.
0516-1010 chpv: Warning, the physical volume hdisk2 has open logical
          volumes. Continuing with change.
#
```

Figure 2.2.8.2 Adding disks to a mirror pool

- Using "lsmp" and "lspv" to display the mirror pool information:  
`"lsmp -A" dbvg` and `"lspv -P"`

```
# lsmp -A dbvg
VOLUME GROUP:      dbvg          Mirror Pool Super Strict: yes
MIRROR POOL:      dbvg_tor01      Mirroring Mode:          SYNC
# lspv -P
Physical Volume   Volume Group   Mirror Pool
hdisk0            None
hdisk1            dbvg           dbvg_tor01
hdisk2            dbvg           dbvg_tor01
hdisk3            rootvg
hdisk4            None
hdisk5            appvg
hdisk6            appvg
#
```

Figure 2.2.8.3 Displaying mirror pool information

In our example, we continue defining mirror pools for the second vg 'appvg' and end up with the following:

```
# lspv -P
Physical Volume   Volume Group   Mirror Pool
hdisk0            None
hdisk1            dbvg           dbvg_tor01
hdisk2            dbvg           dbvg_tor01
hdisk3            rootvg
hdisk4            None
hdisk5            appvg          appvg_tor01
hdisk6            appvg          appvg_tor01
#
```

Figure 2.2.8.4 Displaying mirror pool information

## 2.2.9 Setting LV strictness

The AIX LVs also have to be set to superstrict. We do this to every LV in the volume group. We use "lsvg -l <VGNAME>" to display the LVs and then "chlv -s s <LVNAME>" to set the strictness.

```
pc-tor01-glvm-01a: / >  
# chlv -s s loglv00  
pc-tor01-glvm-01a: / >  
#
```

Figure 2.2.9.1 Changing LV strictness to superstrict

And

```
# lslv loglv00 | more  
LOGICAL VOLUME:      loglv00                VOLUME GROUP:      dbvg  
LV IDENTIFIER:       00c8d44000004b00000001733180f36e.1 PERMISSION:         read/write  
VG STATE:            active/complete        LV STATE:           closed/syncd  
TYPE:                jfs2log                WRITE VERIFY:       off  
MAX LPs:             512                    PP SIZE:            16 megabyte(s)  
COPIES:              2                      SCHED POLICY:       parallel  
LPs:                 1                      PPs:                2  
STALE PPs:           0                      BB POLICY:          relocatable  
INTER-POLICY:        minimum                RELOCATABLE:        yes  
INTRA-POLICY:        middle                 UPPER BOUND:        2  
MOUNT POINT:         N/A                    LABEL:              None  
DEVICE UID:          0                      DEVICE GID:         0  
DEVICE PERMISSIONS: 432  
MIRROR WRITE CONSISTENCY: on/ACTIVE  
EACH LP COPY ON A SEPARATE PV ?: yes (superstrict)  
Serialize IO ?:     NO  
INFINITE RETRY:     no                      PREFERRED READ: 0  
DEVICESUBTYPE:      DS_LVZ  
COPY 1 MIRROR POOL: dbvg_tor01  
COPY 2 MIRROR POOL: dbvg_lon06  
COPY 3 MIRROR POOL: None
```

Figure 2.2.9.2 Displaying LV strictness

Note the field "EACH LP COPY ON A SEPARATE PV ?" is "yes (superstrict).

We repeat this process for all LVs in all of the appropriate VGs.

### **2.2.10 Implementing AIX GLVM RPV servers and RPV clients**

Now that the AIX volume groups have been readied for AIX GLVM, we can implement the AIX GLVM RPV servers and clients. As noted before, the RPV server presents a physical disk to the RPV client. The RPV client uses that to present a virtual disk to the AIX operating system. The operating system can use that virtual disk and store data in it. That data is actually being replicated to the remote RPV server instance and written to the physical volume at that remote location.

To mirror data from our Toronto PowerVS to London, we will setup an RPV server at our London PowerVS. That RPV server will present its physical volume to the RPV client instance at our Toronto PowerVS. The Toronto PowerVS, in turn, will present that volume as a virtual volume to the AIX operating system and we can start to use AIX LVM to replicate data between the different drives.

### **2.2.11 Implementing AIX GLVM RPV servers**

Now that the AIX volume groups have been readied for AIX GLVM, we can implement the AIX GLVM RPV servers. The AIX GLVM RPV servers will exist on the DR site, the IBM Cloud London location for our example. We note the IP address of our PowerVS images at the Toronto and London locations:

- *Toronto: 192.168.6.163*
- *London: 192.168.50.184*

We log into our London PowerVS instance and make sure the LUNs we will be using have PVIDs on them:

- *Checking hdisks we will make available, hdisk1, hdisk2, hdisk5, hdisk6*

```

pc-lon06-glvm-01a:/ >
# lspv
hdisk0          none                None
hdisk1          none                None
hdisk2          none                None
hdisk3          00f6db0af58e9775    rootvg          active
hdisk4          none                None
hdisk5          none                None
hdisk6          none                None
pc-lon06-glvm-01a:/ >
#

```

Figure 2.2.11.1 Determining available disks

- Adding PV ID's to hdisk1, hdisk2, hdisk5, hdisk6

```

pc-lon06-glvm-01a:/ >
# chdev -a pv=yes -l hdisk1
hdisk1 changed
pc-lon06-glvm-01a:/ >
# chdev -a pv=yes -l hdisk2
hdisk2 changed
pc-lon06-glvm-01a:/ >
# chdev -a pv=yes -l hdisk5
hdisk5 changed
pc-lon06-glvm-01a:/ >
# chdev -a pv=yes -l hdisk6
hdisk6 changed
pc-lon06-glvm-01a:/ >
# lspv
hdisk0          none                None
hdisk1          00c8cf803fb0193a    None
hdisk2          00c8cf803fb02116    None
hdisk3          00f6db0af58e9775    rootvg          active
hdisk4          none                None
hdisk5          00c8cf803fb029cd    None
hdisk6          00c8cf803fb03444    None
pc-lon06-glvm-01a:/ >
#

```

Figure 2.2.11.2 Adding PVIDs to disks

- We verify that both RPV server and RPV client IP addresses are on the /etc/hosts file

```

pc-lon06-glvm-01a:/ >
# hostent -S
127.0.0.1          loopback localhost    # loopback (lo0) name/address
192.168.6.163    pc-tor01-glvm-01a
192.168.50.184   pc-lon06-glvm-01a
pc-lon06-glvm-01a:/ >
#

```

Figure 2.2.11.3 Verifying /etc/hosts entries

- We define the RPV site name by going into the SMIT menu 'rpvserver'



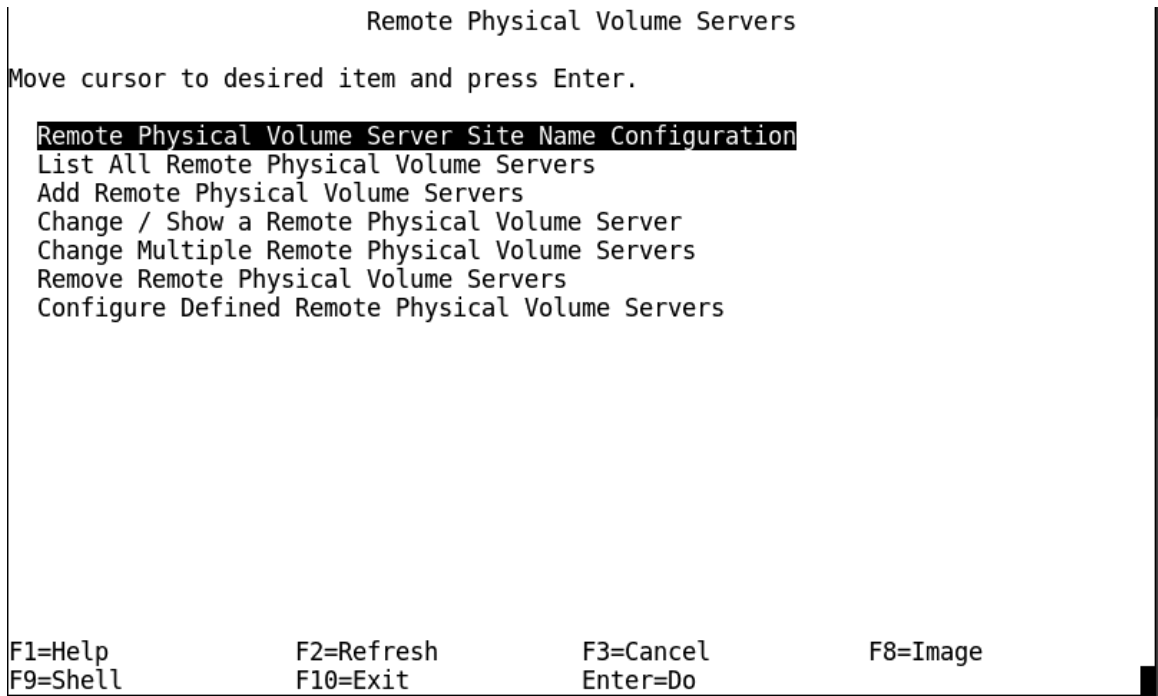


Figure 2.2.11.4 Displaying primary RPV server SMIT panel

- Then select "Remote Physical Volume Server Site Name Configuration" -> "Define / Change / Show Remote Physical Volume Server Site Name" and enter the site name, london06, for our example.

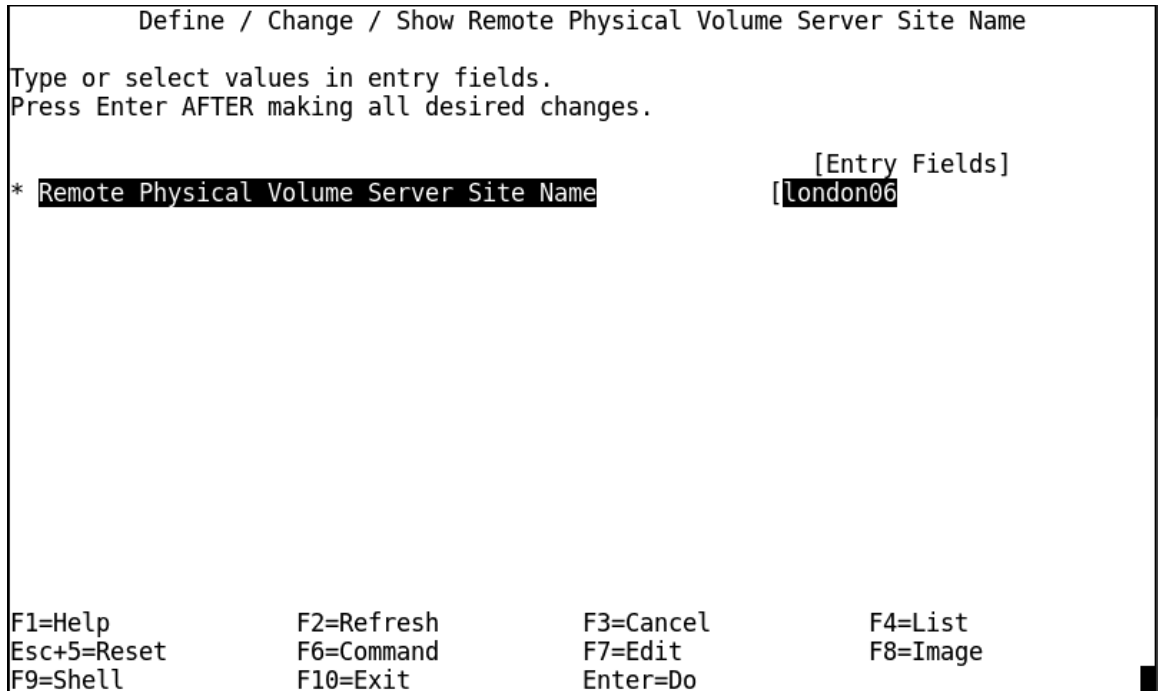


Figure 2.2.11.5 Adding an RPV server site name

- Depress the Enter key and wait for SMIT to complete the process (no output will be given).
- **IMPORTANT NOTE:** The site name must be entered for each site even if only one site will have RPV servers
- Then we add the RPV servers by going to SMIT panel, `smit rpvserver`, then selecting "Add Remote Physical Volume Servers" where the available LUNs are presented. We will select `hdisk1` for this example.

```

Remote Physical Volume Servers
Move cursor to desired item and press Enter.

Remote Physical Volume Server Site Name Configuration
List All Remote Physical Volume Servers

Physical Volume Identifiers
Move cursor to desired item and press F7.
ONE OR MORE items can be selected.
Press Enter AFTER making all selections.

# Physical Volume      Physical Volume Identifier
# -----
hdisk1                 00c8cf803fb0193a
hdisk2                 00c8cf803fb02116
hdisk5                 00c8cf803fb029cd
hdisk6                 00c8cf803fb03444

F1=Help                F2=Refresh            F3=Cancel
F7=Select              F8=Image              F10=Exit
Enter=Do               /=Find                n=Find Next
F1
F9

```

Figure 2.2.11.6 Adding PVs to an RPV server

- At this point, we add the client IP address to the "Remote Physical Volume Client Internet Address" field and depress the Enter key.

```

Add Remote Physical Volume Servers

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

Physical Volume Identifiers          [Entry Fields]
                                     00c8cf803fb0193a
* Remote Physical Volume Client Internet Address [192.168.6.163] +
Configure Automatically at System Restart? [no] +
Start New Devices Immediately? [yes] +

F1=Help      F2=Refresh      F3=Cancel      F4=List
Esc+5=Reset  F6=Command      F7=Edit       F8=Image
F9=Shell     F10=Exit       Enter=Do

```

Figure 2.2.11.7 Configuring RPV server

- The confirmation is displayed

```

COMMAND STATUS

Command: OK          stdout: yes          stderr: no

Before command completion, additional instructions may appear below.

pvserver0 Available

F1=Help      F2=Refresh      F3=Cancel      F6=Command
F8=Image     F9=Shell       F10=Exit      /=Find
n=Find Next

```

Figure 2.2.11.8 Output of RPV server configuration

- Then we repeat this process for the rest of the LUNs we will be presenting to our Toronto PowerVS.

## 2.2.12 Implementing AIX GLVM RPV clients

With the RPV servers setup, we can proceed to the RPV client and configure them to see the remote LUNs. Again, the IP addresses of the dare:

- *Toronto: 192.168.6.163*
- *London: 192.168.50.184*

We log into our Toronto PowerVS instance proceed to the RPV client menu, smit rpvclient:

```
Remote Physical Volume Clients
Move cursor to desired item and press Enter.

List All Remote Physical Volume Clients
Add Remote Physical Volume Clients
Change / Show a Remote Physical Volume Client
Change Multiple Remote Physical Volume Clients
Remove Remote Physical Volume Clients
Configure Defined Remote Physical Volume Clients

F1=Help      F2=Refresh   F3=Cancel    F8=Image
F9=Shell     F10=Exit    Enter=Do
```

Figure 2.2.12.1 SMIT panel for RPV clients

- *We select "Add Remote Physical Volume Clients", determine whether we are using IPv6 or not. The example does not use IPv6.*

```

                                Add Remote Physical Volume Clients

Type or select a value for the entry field.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
* Does data mirroring network use Internet Protocol  no          +
  Version 6 (IPv6)?

F1=Help      F2=Refresh    F3=Cancel    F4=List
Esc+5=Reset  F6=Command    F7=Edit     F8=Image
F9=Shell     F10=Exit     Enter=Do

```

Figure 2.2.12.2 RPV client configuration, IPv6 confirmation

- Then add the IP address of the RPV server

```

                                Add Remote Physical Volume Clients

Type or select a value for the entry field.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
* Remote Physical Volume Server Internet Address  [192.168.50.184]  +

F1=Help      F2=Refresh    F3=Cancel    F4=List
Esc+5=Reset  F6=Command    F7=Edit     F8=Image
F9=Shell     F10=Exit     Enter=Do

```

Figure 2.2.12.3 RPV client configuration of RPV server IP

- And

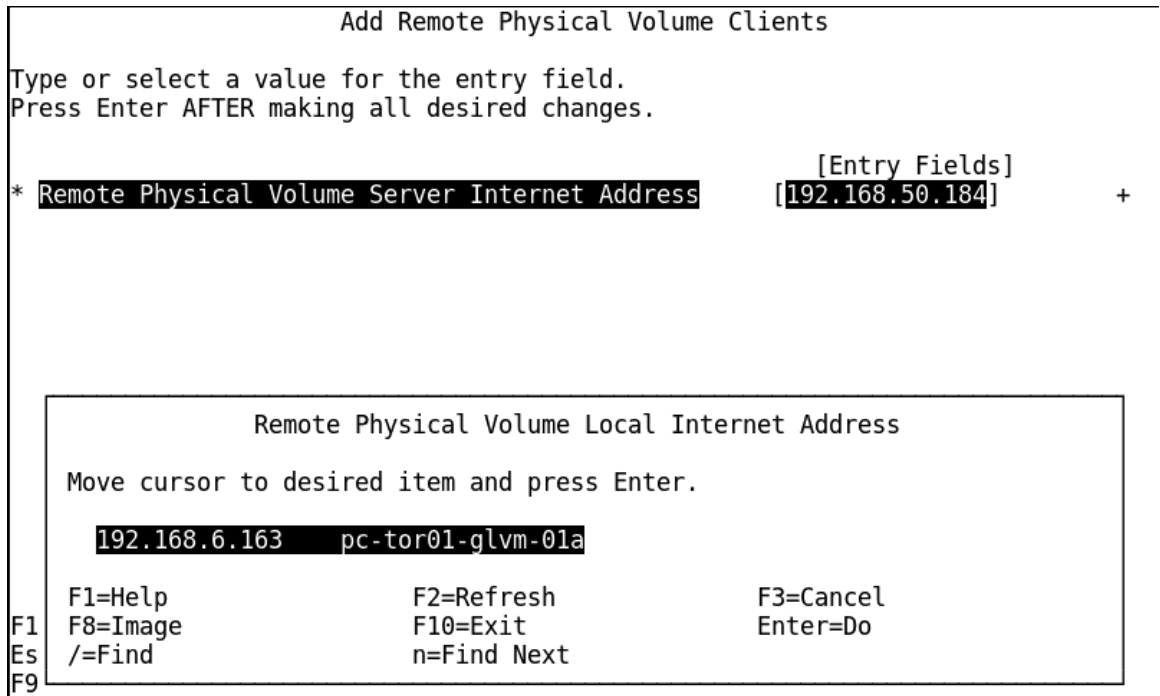


Figure 2.2.12.4 Selecting RPV server location from RPV client

- You will be presented with the available RPV servers which correspond to the LUNs that are being presented by the RPV server PowerVS instance.

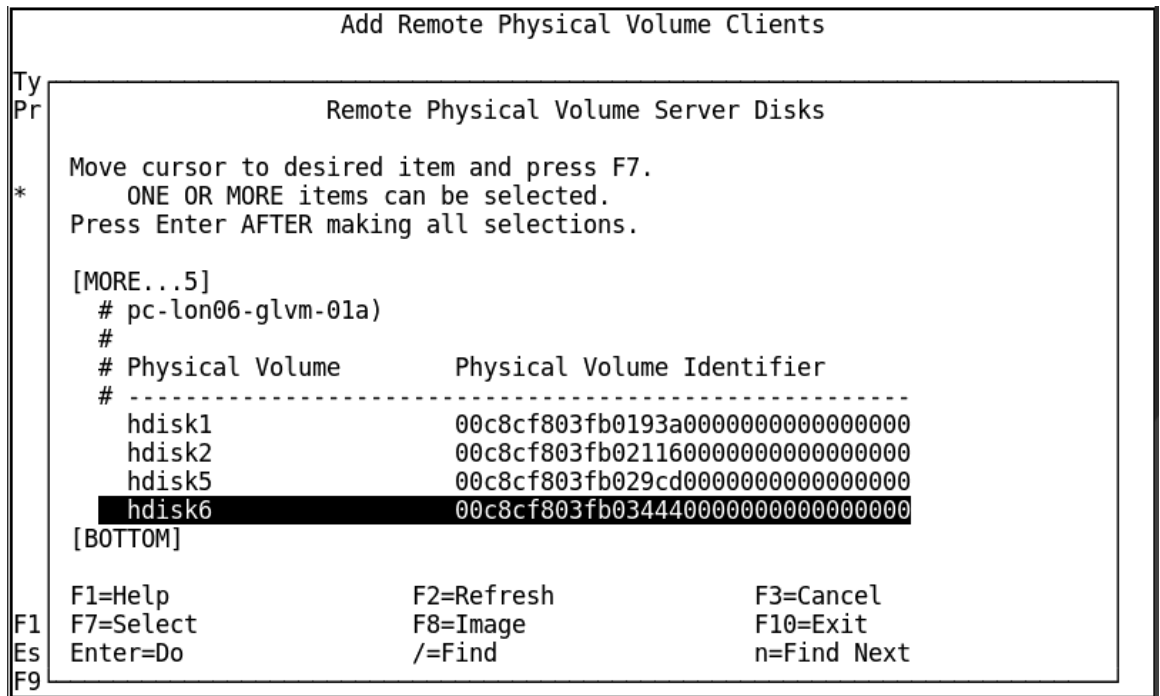


Figure 2.2.12.5 Selecting RPV server LUN from RPV client

- We select "hdisk1" and depress the Enter key bring up the device parameter menu.

```

                                Add Remote Physical Volume Clients

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
Remote Physical Volume Server Internet Address      192.168.50.184
Remote Physical Volume Local Internet Address      192.168.6.163
Physical Volume Identifiers                        0c8cf803fb0193a00000>
I/O Timeout Interval (Seconds)                    [180] #
Start New Devices Immediately?                     [yes] +

F1=Help      F2=Refresh      F3=Cancel      F4=List
Esc+5=Reset  F6=Command      F7=Edit       F8=Image
F9=Shell     F10=Exit       Enter=Do

```

Figure 2.2.12.6 Configuring RPV client parameters

- For our case, we select the default parameters and depress the Enter key bring up the results screen.

```

                                COMMAND STATUS

Command: OK          stdout: yes          stderr: no

Before command completion, additional instructions may appear below.

disk7 Available

F1=Help      F2=Refresh      F3=Cancel      F6=Command
F8=Image     F9=Shell       F10=Exit       /=Find
n=Find Next

```

Figure 2.2.12.7 RPV client configuration output

- We then repeat the same process for the remaining LUNs.
- Once completed, we note the new LUNs available on our Toronto PowerVS instance (hdisk7->hdisk10).

```
pc-tor01-glvm-01a: / >
# lspv
hdisk0          none                               None
hdisk1          00c8d4403180f273                   dbvg          active
hdisk2          00c8d4403180f319                   dbvg          active
hdisk3          00f6db0af58e9775                   rootvg       active
hdisk4          none                               None
hdisk5          00c8d44031815edd                   appvg        active
hdisk6          00c8d44031815f5a                   appvg        active
hdisk7          00c8cf803fb0193a                   None
hdisk8          00c8cf803fb02116                   None
hdisk9          00c8cf803fb029cd                   None
hdisk10         00c8cf803fb03444                   None
pc-tor01-glvm-01a: / >
#
```

Figure 2.2.12.8 Displaying addition of RPV client

- And also note that these LUNs are RPV clients by using the "lsdev -Cc disk" command.

```
pc-tor01-glvm-01a: / >
# lsdev -Cc disk
hdisk0 Available C7-T1-01 MPIO IBM 2076 FC Disk
hdisk1 Available C7-T1-01 MPIO IBM 2076 FC Disk
hdisk2 Available C7-T1-01 MPIO IBM 2076 FC Disk
hdisk3 Available C7-T1-01 MPIO IBM 2076 FC Disk
hdisk4 Available C7-T1-01 MPIO IBM 2076 FC Disk
hdisk5 Available C7-T1-01 MPIO IBM 2076 FC Disk
hdisk6 Available C7-T1-01 MPIO IBM 2076 FC Disk
hdisk7 Available Remote Physical Volume Client
hdisk8 Available Remote Physical Volume Client
hdisk9 Available Remote Physical Volume Client
hdisk10 Available Remote Physical Volume Client
pc-tor01-glvm-01a: / >
#
```

Figure 2.2.12.9 Display local and RPV client hdisks

### 2.2.13 Replicating data between sites using AIX LVM

With the RPV servers and clients defined, the virtual drives can now be used to replicate the data between sites.

In our example, we have a volume group "dbvg", which include the following LUNs: hdisk1, hdisk2



```
pc-tor01-glv-01a: / >
# lspv -P
Physical Volume   Volume Group      Mirror Pool
hdisk0            None
hdisk1            dbvg              dbvg_tor01
hdisk2            dbvg              dbvg_tor01
hdisk3            rootvg
hdisk4            None
hdisk5            appvg            appvg_tor01
hdisk6            appvg            appvg_tor01
hdisk7            None
hdisk8            None
hdisk9            None
hdisk10           None
pc-tor01-glv-01a: / >
```

Figure 2.2.13.1 Displaying disks and associated mirror pools

- We will add the virtual disks, *hdisk7* and *hdisk8* to the volume group *dbvg* with a mirror pool name of *dbvg\_lon06* using the command: `extendvg -p dbvg_lon06 dbvg hdisk7 hdisk8`

```
pc-tor01-glv-01a: /usr/sbin >
# lspv -P
Physical Volume   Volume Group      Mirror Pool
hdisk0            None
hdisk1            dbvg              dbvg_tor01
hdisk2            dbvg              dbvg_tor01
hdisk3            rootvg
hdisk4            None
hdisk5            appvg            appvg_tor01
hdisk6            appvg            appvg_tor01
hdisk7            dbvg              dbvg_lon06
hdisk8            dbvg              dbvg_lon06
```

Figure 2.2.13.2 Extended a VG with new disk and mirror pool

- Note that *hdisk7*, *hdisk8* (our virtual drives) are now part of the "dbvg" volume group and belong to a mirror pool named "dbvg\_lon06"
- We then mirror our "dbvg" by mirror each of the LVs, such as *fslv00*, using the following command: `mklvcopy -p copy1=dbvg_tor01 -p copy2=dbvg_lon06 fslv00 2`  
The above command tells AIX to mirror the "fslv00" logical volume with copy 1 being in the "dbvg\_tor01" mirror pool and copy 2 being in the "dbvg\_lon06" making a total of 2 copies
- We verify the replication by using the "lsvg -l dbvg" command

```
# mklvcopy -p copy1=dbvg_tor01 -p copy2=dbvg_lon06 fslv00 2
pc-tor01-glvm-01a: /usr/sbin >
# lsvg -l dbvg
dbvg:
LV NAME          TYPE      LPs      PPs      PVs  LV STATE      MOUNT POINT
loglv00          jfs2log   1         1         1   closed/syncd  N/A
fslv00           jfs2     128      256      2   closed/stale  /data01
fslv01           jfs2     128      128      1   closed/syncd  /data02
fslv02           jfs2     128      128      1   closed/syncd  /data03
fslv03           jfs2     128      128      1   closed/syncd  /data04
```

Figure 2.2.13.3 Displaying multiple copies of fslv00 LV

- Note that the "fslv00" logical volume now has 2 copies and the state is "stale". The state will be updated in the near future.
- We proceed in mirroring the remaining LVs, loglv00, fslv01, fslv02, fslv03. This is verified with the "lsvg -l dbvg" command.

```
pc-tor01-glvm-01a: /usr/sbin >
# lsvg -l dbvg
dbvg:
LV NAME          TYPE      LPs      PPs      PVs  LV STATE      MOUNT POINT
loglv00          jfs2log   1         2         2   closed/stale  N/A
fslv00           jfs2     128      256      2   closed/stale  /data01
fslv01           jfs2     128      256      2   closed/stale  /data02
fslv02           jfs2     128      256      2   closed/stale  /data03
fslv03           jfs2     128      256      2   closed/stale  /data04
```

Figure 2.2.13.4 Displaying multiple copies of all LVs in VG, dbvg

- Note that all LVs are now 2 copies (1 LP and 2 PPs) across 2 PVs and in stale status.
- To start the synchronization, simply execute the varyonvg command: varyonvg dbvg
- The synchronization will initiate which can be verified by running the "ps -ef | grep sync" command

```
pc-tor01-glvm-01a: /usr/sbin >
# varyonvg dbvg
pc-tor01-glvm-01a: /usr/sbin >
# ps -ef | grep -i sync
  root   3539314      1   0   Jul 11   -   0:04 /usr/sbin/syncd 60
  root   9306570 19792158   0 03:51:42 vty0  0:00 lresynclv -l 00c8d44000004b
00000001733180f36e
  root   19792158      1   0 03:51:42 vty0  0:00 /bin/ksh /usr/sbin/syncvg -
v dbvg
  root   22151544  7406072   0 03:51:48 vty0  0:00 grep -i sync
```

Figure 2.2.13.5 Displaying VG synchronization processes

- Note the "lresynclv" command is running which is what is synchronizing that volume group.

- The progress can be checked by running the "lsvg -l dbvg" command

```
pc-tor01-glvm-01a: /usr/sbin >
# lsvg -l dbvg
dbvg:
LV NAME          TYPE      LPs    PPs    PVs  LV STATE    MOUNT POINT
loglv00          jfs2log   1      2      2    closed/syncd  N/A
fslv00           jfs2     128    256    2    open/stale   /data01
fslv01           jfs2     128    256    2    closed/stale /data02
fslv02           jfs2     128    256    2    closed/stale /data03
fslv03           jfs2     128    256    2    closed/stale /data04
```

Figure 2.2.13.6 Displaying LV synchronization status

- Note that the LV, "loglv00", is now in syncd state which means is fully synchronized between all drives (local and remote) while the other LVs remain to be synchronized as those are still in "stale" state.
- The same process is repeated for the additional VG, appvg in our example.

## 2.2.14 Accessing the data in a DR event

The replication setup implemented in this scenario is a "true" DR setup. The implementation is purely one way, from the production site to the DR site.

In the event of a disaster, we will only have half of the storage for each of the volume groups.

We will simulate a DR event by shutting down the production PowerVS.

```

1 unsuccessful login attempt since last login.
Last unsuccessful login: Tue Jul 21 17:33:46 CDT 2020 on /dev/vty0
Last login: Tue Jul 21 14:41:46 CDT 2020 on /dev/vty0

pc-tor01-glv-01a: / >
# shutdown -F

SHUTDOWN PROGRAM
Tue Jul 21 17:34:04 CDT 2020
Jul 21 2020 17:34:04
/usr/es/sbin/cluster/utilities/clstop : called with flags -f -y -s -N -S
0513-044 The clinfoES Subsystem was requested to stop.
0513-044 The clevmgrdES Subsystem was requested to stop.
Running /etc/rc.d/rc2.d/Kcluster stop
Quiescing cthags...
Stopping IBM.ConfigRM...
0513-004 The Subsystem or Group, IBM.ConfigRM, is currently inoperative.
Running /etc/rc.d/rc2.d/Ksshd stop
0513-044 The sshd Subsystem was requested to stop.
Running /etc/rc.d/rc2.d/Kwpars stop

Wait for '...Halt completed....' before stopping.
Error reporting has stopped.
█

```

Figure 2.2.14.1 Shutting down production PowerVS

This results in the DR PowerVS being the only active image. We log into the DR PowerVS in London and note that we see only the local disks.

```

pc-lon06-glv-01a:/ >
# lspv
hdisk0          00c8cf80601b4933          None
hdisk1          00c8cf803fb0193a          None
hdisk2          00c8cf803fb02116          None
hdisk3          00f6db0af58e9775          rootvg                active
hdisk4          00c8cf8049bac9fa          caavg_private         active
hdisk5          00c8cf803fb029cd          None
hdisk6          00c8cf803fb03444          None
pc-lon06-glv-01a:/ >
#

```

Figure 2.2.14.1 DR PowerVS storage

The process for accessing the data is to import the volume groups with the LUNs that are available. We note that hdisk1 and hdisk2 are associated with the “dbvg” volume group and hdisk5, hdisk6 are associated with the “appvg” volume group.

We import the “dbvg” volume group using “importvg -f -y dbvg -V 100 hdisk1” command. Note that the “-f” means we are forcing the import,

the “-V 100” flag means we are assigning the VG the major number of 100. In the following figure, note the error messages as they will be utilized when cleaning up the VG.

```
pc-lon06-glv-01a:/ >
# importvg -f -y dbvg -V 100 hdisk1
PV Status:      hdisk1  00c8cf803fb0193a      PACTIVE
                hdisk2  00c8cf803fb02116      PACTIVE
                00c8d4403180f273  NONAME
                00c8d4403180f319  NONAME
varyonvg: Volume group dbvg is varied on.
0516-510 synclvodm: Physical volume not found for physical volume
                identifier 00c8d4403180f2730000000000000000.
0516-510 synclvodm: Physical volume not found for physical volume
                identifier 00c8d4403180f3190000000000000000.
0516-548 synclvodm: Partially successful with updating volume
                group dbvg.
dbvg
0516-783 importvg: This imported volume group is concurrent capable.
                Therefore, the volume group must be varied on manually.
pc-lon06-glv-01a:/ >
```

Figure 2.2.14.2 Importing a VG on the DR PowerVS

After importing the VG, we will vary it on so that the filesystems can be access. We utilize the “varyonvg -f -O dbvg” command. Note the “-f” and “-O” flags are utilized to vary on the VG due the fact there is no quorum.

```
# varyonvg -O dbvg
0516-052 varyonvg: Volume group cannot be varied on without a
                quorum. More physical volumes in the group must be active.
                Run diagnostics on inactive PVs.
pc-lon06-glv-01a:/ >
# varyonvg -f -O dbvg
PV Status:      hdisk1  00c8cf803fb0193a      PACTIVE
                hdisk2  00c8cf803fb02116      PACTIVE
                00c8d4403180f273  NONAME
                00c8d4403180f319  NONAME
varyonvg: Volume group dbvg is varied on.
pc-lon06-glv-01a:/ >
```

Figure 2.2.14.3 Varying on a VG without a quorum

At this point, we list the VG and note the filesystems are still there.

```
pc-lon06-glvm-01a:/ >
# lsvg -l dbvg
dbvg:
LV NAME          TYPE      LPs    PPs    PVs  LV STATE    MOUNT POINT
loglv00          jfs2log   1      2      2    closed/syncd  N/A
fslv00           jfs2     128    256    2    closed/syncd  /data01
fslv01           jfs2     128    256    2    closed/syncd  /data02
fslv02           jfs2     128    256    2    closed/syncd  /data03
fslv03           jfs2     128    256    2    closed/syncd  /data04
```

Figure 2.2.14.4 Listing filesystems

We mount the filesystems so the data can be accessed.

```
pc-lon06-glvm-01a:/ >
# mount /data01;mount /data02;mount /data03;mount /data04
pc-lon06-glvm-01a:/ >
#
```

Figure 2.2.14.5 Mounting filesystems

And

```
pc-lon06-glvm-01a:/ >
# df
Filesystem      512-blocks    Free %Used    Iused %Iused Mounted on
/dev/hd4         196608        105560  47%      3038   21% /
/dev/hd2         4784128       330112  94%     40766  50% /usr
/dev/hd9var      1966080       1533632 22%      1764   2% /var
/dev/hd3         458752        453424  2%         55   1% /tmp
/dev/hd1         65536         64792   2%         7   1% /home
/dev/hd11admin   262144        261384  1%         5   1% /admin
/proc            -              -      -         -     - /proc
/dev/hd10opt     786432        85096  90%     11482  53% /opt
/dev/livedump    524288        523552  1%         4   1% /var/adm/ras/livedump
/dev/repo00      15335424      66280  100%     2534  25% /usr/sys/inst.images
/dev/toolslv     2621440       452056  83%        12   1% /tools
/ahafs           -              -      -         48   1% /aha
/dev/fslv00      4194304       4084024 3%      1016   1% /data01
/dev/fslv01      4194304       3598480 15%      1945   1% /data02
/dev/fslv02      4194304       4000328 5%        826   1% /data03
/dev/fslv03      4194304       4193008 1%         4   1% /data04
```

Figure 2.2.14.6 Displaying mounted filesystems

We proceed with the remaining VG.

At this point, we can start the application and proceed with our recovery.

## 2.2.15 Cleaning up mirrors after a DR event

While the application and data was successfully brought up in the previous section, there remains some activities to clean up LV mirrors. After mount the filesystems and starting the application, the data will be marked "stale" due to the fact there are no mirror physical volumes. This can be displayed via the "lsvg -l dbvg" command for the "dbvg" volume group.

```
pc-lon06-glv-01a:/ >
# lsvg -l dbvg
dbvg:
LV NAME          TYPE      LPs    PPs    PVs  LV STATE  MOUNT POINT
loglv00          jfs2log   1      2      2    open/stale  N/A
fslv00           jfs2     128    256    2    open/stale  /data01
fslv01           jfs2     128    256    2    open/stale  /data02
fslv02           jfs2     128    256    2    open/stale  /data03
fslv03           jfs2     128    256    2    open/stale  /data04
```

Figure 2.2.15.1 Displaying stale data

This can be cleaned up by removing the copy from the previously existing drives.

We recall that when we imported the volume group, there were some PVs missing.

```
pc-lon06-glv-01a:/ >
# varyonvg -f -0 dbvg
PV Status:      hdisk1  00c8cf803fb0193a    PACTIVE
                hdisk2  00c8cf803fb02116    PACTIVE
                00c8d4403180f273    NONAME
                00c8d4403180f319    NONAME
varyonvg: Volume group dbvg is varied on.
pc-lon06-glv-01a:/ >
```

Figure 2.2.15.2 Importing "dbvg" volume group

We take note of the PVIDs that are missing. We proceed cleaning up the LVs by removing the "stale" copy using the "rmlvcopy <LVNAME> 1 <PVID>" as shown in the following diagram.

```
pc-lon06-glvm-01a:/ >
# rmlvcopy fslv00 1 00c8d4403180f319
pc-lon06-glvm-01a:/ >
#
```

Figure 2.2.15.3 Removing “stale” LV copies

We note that the LV is now no longer stale as there is only one copy.

```
pc-lon06-glvm-01a:/ >
# lsvg -l dbvg
dbvg:
LV NAME          TYPE      LPs    PPs    PVs    LV STATE    MOUNT POINT
loglv00          jfs2log   1       1       1    open/syncd  N/A
fslv00           jfs2     128    128    1    open/syncd  /data01
fslv01           jfs2     128    256    2    open/stale  /data02
fslv02           jfs2     128    256    2    open/stale  /data03
fslv03           jfs2     128    256    2    open/stale  /data04
pc-lon06-glvm-01a:/ >
```

Figure 2.2.15.4 Displaying LV, “fslv00”, has one copy

We proceed with all the LVs until it is all cleaned up as shown in the figure below.

```
pc-lon06-glvm-01a:/ >
# lsvg -l dbvg
dbvg:
LV NAME          TYPE      LPs    PPs    PVs    LV STATE    MOUNT POINT
loglv00          jfs2log   1       1       1    open/syncd  N/A
fslv00           jfs2     128    128    1    open/syncd  /data01
fslv01           jfs2     128    128    1    open/syncd  /data02
fslv02           jfs2     128    128    1    open/syncd  /data03
fslv03           jfs2     128    128    1    open/syncd  /data04
pc-lon06-glvm-01a:/ >
```

Figure 2.2.15.5 Displaying LV are “cleaned up” with only one copy



## **2.3 Geographic Logical Volume Manager (GLVM) Replication with PowerHA**

While AIX GLVM provides a method of replicating an application's data to a DR site, it does not provide a method of automatically detecting a failure and bringing up the DR site. Integrating AIX GLVM with IBM PowerHA Enterprise Edition (EE) will provide that automation.

IBM PowerHA technology provides high availability, business continuity and disaster recovery. It enables the deployment of an HA solution that addresses both storage and high availability requirements with one integrated configuration.

In our configuration, we will build a simple 2 node cluster where one node is our PowerVS instance in the Toronto location and the second node is our PowerVS instance in London.

We will build the IBM PowerHA cluster configuration with a single resource group which will contain the volume groups, dbvg and appvg. For simplicity, there will be no other resources on that resource group such as IP addresses or application controllers.

For examples of more complex PowerHA cluster, there are many resources listed in "Chapter 4: Additional resources".

### **2.3.1 Preparing AIX GLVM for IBM PowerHA EE**

With the AIX GLVM only setup, we had the RPV servers only on the DR site. In essence, it was a one-way replication from the primary site to the DR site. Data did not flow from the DR site to the primary site.

With IBM PowerHA, we will need to have the ability of the data to be replicated in either directions. For that to occur, we will need to create RPV servers on the primary site so that the DR site can see those drives as virtual drives using the RPV clients on the DR site PowerVS.

In that scenario, PowerHA will be managing which site is operational and control the RPV server/client activation and deactivation.

The starting state of the AIX GLVM-only replication looks like this diagram.

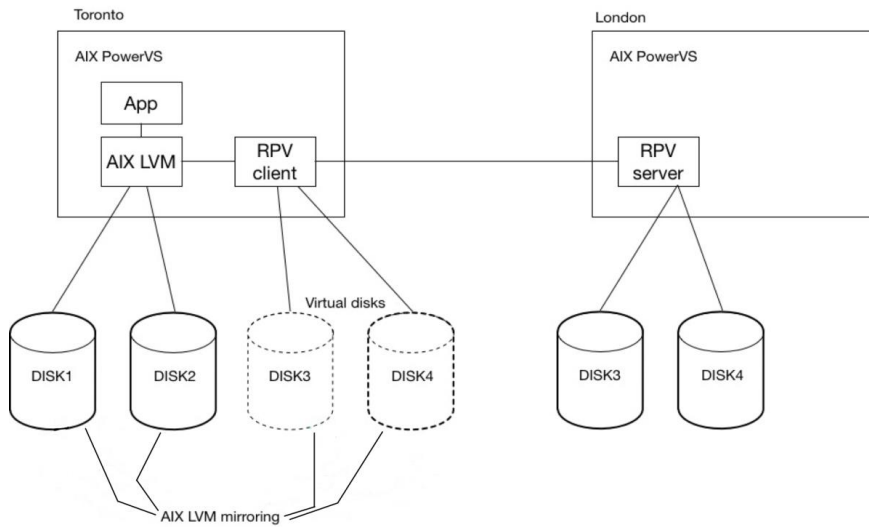


Figure 2.3.1.1 Starting state of AIX GLVM-only replication

Note that in the above diagram, only the London location PowerVS have RPV servers. We will add RPV servers to the Toronto PowerVS and RPV clients to the London PowerVS in order to present the drives from Toronto to London.

The final AIX GLVM configuration will look like the following diagram.

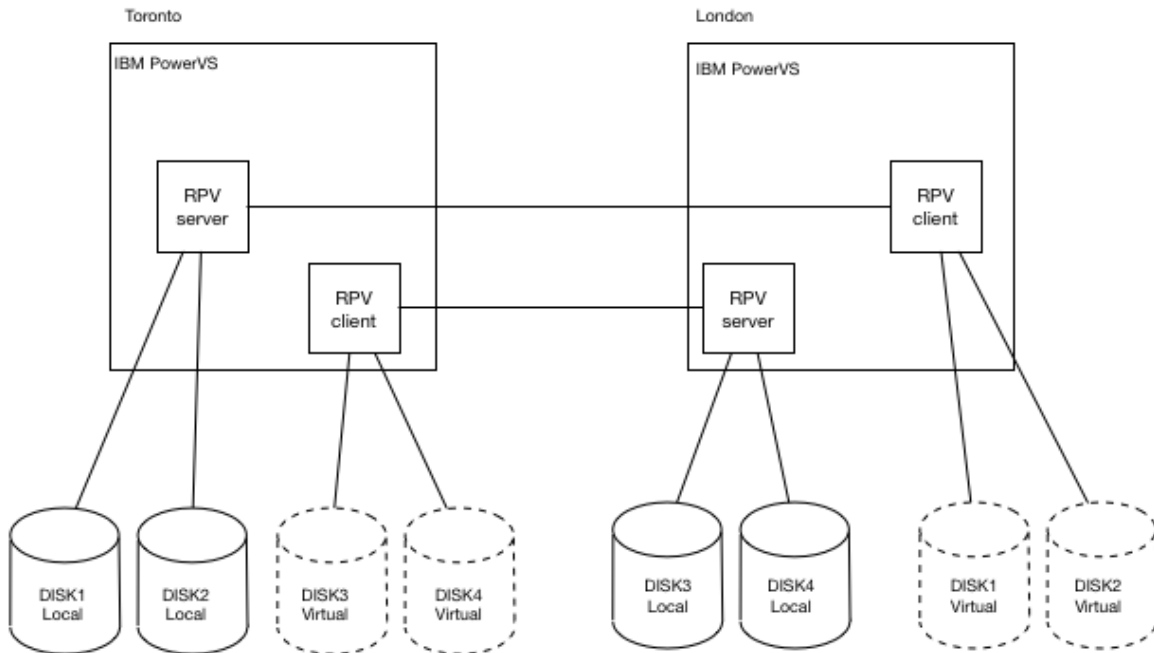


Figure 2.3.1.2 AIX GLVM replication in both directions

We proceed with the following steps

- *Determine for which drives to create RPV servers*
- *Bring the VGs offline*
- *Create the RPV servers on the primary location*
- *Create the RPV clients on the DR location*
- *Verify the configuration*
- *Import all VGs to remote PowerVS*

### 2.3.2 Update VGs to enhanced concurrent capable

PowerHA requires all of the shared VGs to be enhanced concurrent capable. The VGs can be updated to enhanced concurrent capable with the "chvg -C <VGNAME>" command. See the following diagram.

```
pc-lon06-glvm-01a:/ >
# chvg -C dbvg
```

Figure 2.3.2.1 Change VG to concurrent capable

```

VOLUME GROUP:      dbvg          VG IDENTIFIER: 00c8d4400004b000000017331
80f36e
VG STATE:          active        PP SIZE:        16 megabyte(s)
VG PERMISSION:    read/write    TOTAL PPs:      5100 (81600 megabytes)
MAX LVs:          256          FREE PPs:       4074 (65184 megabytes)
LVs:              5            USED PPs:       1026 (16416 megabytes)
OPEN LVs:         0            QUORUM:         3 (Enabled)
TOTAL PVs:        4            VG DESCRIPTORS: 4
STALE PVs:        0            STALE PPs:      0
ACTIVE PVs:       4            AUTO ON:        no
Concurrent:       Enhanced-Capable Auto-Concurrent: Disabled
VG Mode:          Non-Concurrent
MAX PPs per VG:  32768          MAX PVs:        1024
LTG size (Dynamic): 512 kilobyte(s) AUTO SYNC:      no
HOT SPARE:        no            BB POLICY:      non-relocatable
MIRROR POOL STRICT: super
PV RESTRICTION:  none          INFINITE RETRY: no
DISK BLOCK SIZE: 512          CRITICAL VG:   no
FS SYNC OPTION:  no            CRITICAL PVs:  no

```

Figure 2.3.2.2 Showing Enhanced-Capable field

Note in the above diagram the VG has the “Concurrent” field as “Enhanced-Capable”.

### 2.3.3 Change filesystems to not automount

For PowerHA to manage the filesystems, they must not automatically mount. We change those filesystems using the “chfs -A no <FILESYSTEM>” command. See the following.

```

# for i in 1 2 3;do chfs -A no /appfs$i;done
pc-lon06-glvm-01a:/ >
# lsfs | sed 's/ / /g'
Name      Nodename  Mount Pt      VFS  Size  Options  Auto Accounting
/dev/hd4   --       /             jfs2 196608 --     yes no
/dev/hd1   --       /home        jfs2 65536  --     yes no
/dev/hd2   --       /usr         jfs2 4784128 --    yes no
/dev/hd9var --       /var         jfs2 1966080 --    yes no
/dev/hd3   --       /tmp         jfs2 458752 --    yes no
/dev/hd11admin --     /admin      jfs2 262144 --    yes no
/proc     --       /proc        procfs --     --     yes no
/dev/hd10opt --     /opt        jfs2 786432 --     yes no
/dev/livedump --     /var/adm/ras/livedump jfs2 524288 --     yes no
/dev/repo00 --     /usr/sys/inst.images jfs2 15335424 rw     yes no
/dev/toolslv --     /tools      jfs2 2621440 rw     yes no
/dev/fslv00 --     /data01     jfs2 4194304 rw     no no
/dev/fslv01 --     /data02     jfs2 4194304 rw     no no
/dev/fslv02 --     /data03     jfs2 4194304 rw     no no
/dev/fslv03 --     /data04     jfs2 4194304 rw     no no
/dev/fslv07 --     /appfs1     jfs2 6291456 rw     no no
/dev/fslv05 --     /appfs2     jfs2 6291456 rw     no no
/dev/fslv06 --     /appfs3     jfs2 6291456 rw     no no

```

Figure 2.3.3.1 Changing filesystems to not automount

## 2.3.4 Determine for which drives to create RPV servers

We note the VGs and the VG drives.

```
pc-tor01-glvm-01a: /tools >
# lsvg -p dbvg | sed 's/ / /g'
dbvg:
PV_NAME      PV STATE     TOTAL PPs   FREE PPs   FREE DISTRIBUTION
hdisk1       active       1275        1018       127..126..255..255..255
hdisk2       active       1275        1019       127..127..255..255..255
hdisk7       active       1275        1018       127..126..255..255..255
hdisk8       active       1275        1019       127..127..255..255..255
pc-tor01-glvm-01a: /tools >
# lsvg -p appvg | sed 's/ / /g'
appvg:
PV_NAME      PV STATE     TOTAL PPs   FREE PPs   FREE DISTRIBUTION
hdisk5       active       1275        1082       255..62..255..255..255
hdisk6       active       1275        891        63..63..255..255..255
hdisk9       active       1275        1082       255..62..255..255..255
hdisk10      active       1275        891        63..63..255..255..255
```

Figure 2.3.4.1 VGs and their disk drives

We also note the local drives versus the RPV client drives.

```
pc-tor01-glvm-01a: /tools >
# lsdev -Ccdisk
hdisk0 Available C7-T1-01 MPIIO IBM 2076 FC Disk
hdisk1 Available C7-T1-01 MPIIO IBM 2076 FC Disk
hdisk2 Available C7-T1-01 MPIIO IBM 2076 FC Disk
hdisk3 Available C7-T1-01 MPIIO IBM 2076 FC Disk
hdisk4 Available C7-T1-01 MPIIO IBM 2076 FC Disk
hdisk5 Available C7-T1-01 MPIIO IBM 2076 FC Disk
hdisk6 Available C7-T1-01 MPIIO IBM 2076 FC Disk
hdisk7 Available Remote Physical Volume Client
hdisk8 Available Remote Physical Volume Client
hdisk9 Available Remote Physical Volume Client
hdisk10 Available Remote Physical Volume Client
pc-tor01-glvm-01a: /tools >
```

Figure 2.3.1.3.2 Determining local versus remote drives

Given the above, we determine that hdisk1, hdisk2 and hdisk5, hdisk6 will need to have RPV servers.

## 2.3.5 Bring the VGs offline

For safety, we vary off the volume groups dbvg, appvg.

```
pc-tor01-glvm-01a: /tools >
# varyoffvg dbvg
pc-tor01-glvm-01a: /tools >
# varyoffvg appvg
pc-tor01-glvm-01a: /tools >
```

Figure 2.3.5.1 VGs being brought offline

### 2.3.6 Create the RPV servers on the primary location

We proceed as previously to create those RPV servers for those drives using "smitty rpvserver -> Add Remote Physical Volume Servers". We select the appropriate drive.

```
Remote Physical Volume Servers
Move cursor to desired item and press Enter.

Remote Physical Volume Server Site Name Configuration
List All Remote Physical Volume Servers
Add Remote Physical Volume Servers
Change / Show a Remote Physical Volume Server

Physical Volume Identifiers
Move cursor to desired item and press F7.
ONE OR MORE items can be selected.
Press Enter AFTER making all selections.

# Physical Volume      Physical Volume Identifier
# -----
hdisk5                 00c8d44031815edd
hdisk6                 00c8d44031815f5a

F1=Help                F2=Refresh            F3=Cancel
F7=Select              F8=Image              F10=Exit
Enter=Do               /=Find                n=Find Next

F1
F9
```

Figure 2.3.6.1 Selecting RPV server drive

We select the appropriate drive and proceed to the RPV parameters screen where we add the IP address.

```

Add Remote Physical Volume Servers

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

Physical Volume Identifiers          [Entry Fields]
00c8d44031815edd
* Remote Physical Volume Client Internet Address [192.168.50.184] +
Configure Automatically at System Restart? [no] +
Start New Devices Immediately? [yes] +

F1=Help      F2=Refresh      F3=Cancel      F4=List
Esc+5=Reset  F6=Command      F7=Edit       F8=Image
F9=Shell     F10=Exit       Enter=Do

```

Figure 2.3.6.2 Configuring RPV server parameters

And, resulting in the configured rpvserver. We repeat this process for all appropriate drives.

```

COMMAND STATUS

Command: OK          stdout: yes          stderr: no

Before command completion, additional instructions may appear below.

pvserver2 Available

F1=Help      F2=Refresh      F3=Cancel      F6=Command
F8=Image     F9=Shell       F10=Exit       /=Find
n=Find Next

```

Figure 2.3.6.3 Configured RPV server

### 2.3.7 Create the RPV clients on the DR location

We proceed as previously to create those RPV clients for those drives using "smitty rpvserver -> Add Remote Physical Volume Clients". We select the appropriate IPv6 option.

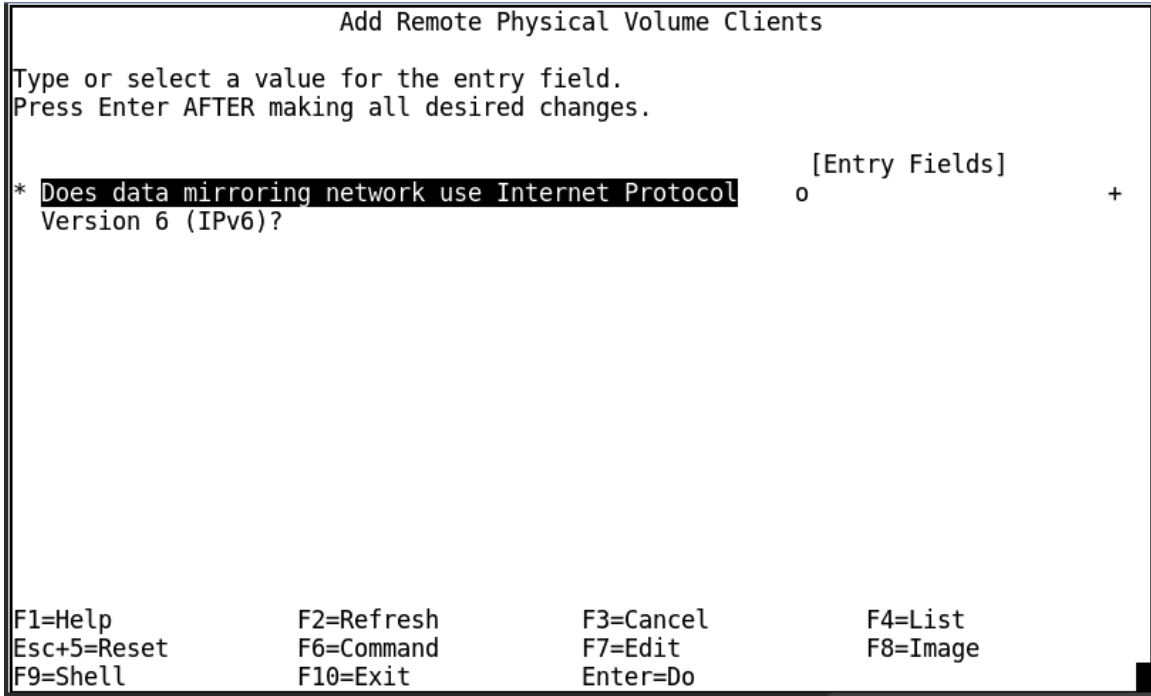


Figure 2.3.7.1 Selecting RPV client IPv6 option

We then select the RPV server IP address.



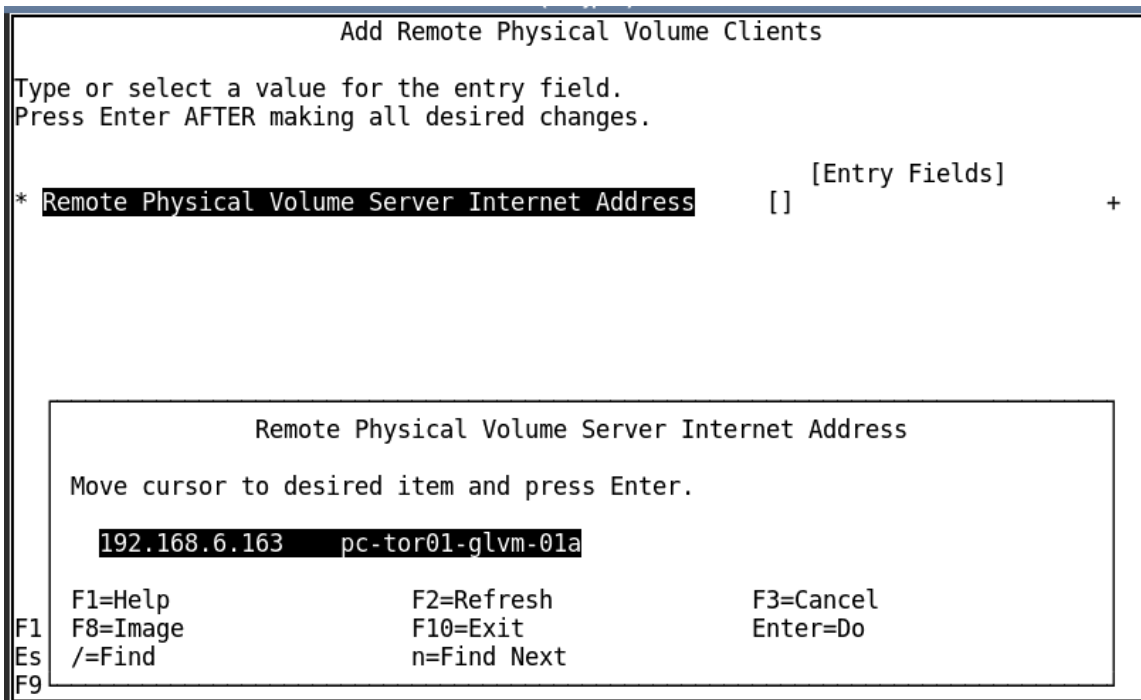


Figure 2.3.7.2 Selecting RPV server address

Then we select the drive the RPV client is to access.

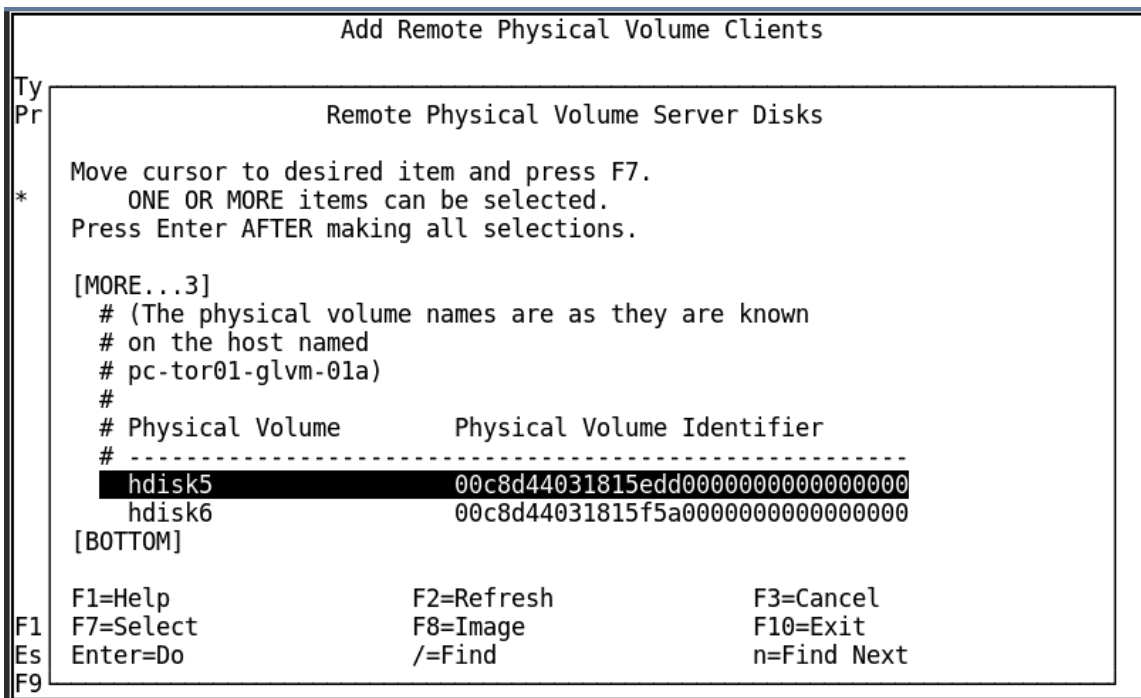


Figure 2.3.7.3 Select RPV client drive is to access

Lastly, we are presented with the RPV client parameters menu. For our example, we will leave parameters as default.

```

                                Add Remote Physical Volume Clients

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
Remote Physical Volume Server Internet Address      192.168.6.163
Remote Physical Volume Local Internet Address      192.168.50.184
Physical Volume Identifiers                        00c8d44031815edd00000>
I/O Timeout Interval (Seconds)                    [ 80] #
Start New Devices Immediately?                     [yes] +

F1=Help      F2=Refresh      F3=Cancel      F4=List
Esc+5=Reset  F6=Command      F7=Edit       F8=Image
F9=Shell     F10=Exit        Enter=Do

```

Figure 2.3.7.4 RPV client parameters

We proceed with the additional RPV client configurations.

### 2.3.8 Verify the configuration

We can verify the RPV client configuration using the “lsdev -Cc disk” command as shown below.

```

pc-lon06-glvms-01a:/ >
# lsdev -Cc disk
hdisk0 Available C4-T1-01 MPIO IBM 2076 FC Disk
hdisk1 Available C4-T1-01 MPIO IBM 2076 FC Disk
hdisk2 Available C4-T1-01 MPIO IBM 2076 FC Disk
hdisk3 Available C4-T1-01 MPIO IBM 2076 FC Disk
hdisk4 Available C4-T1-01 MPIO IBM 2076 FC Disk
hdisk5 Available C4-T1-01 MPIO IBM 2076 FC Disk
hdisk6 Available C4-T1-01 MPIO IBM 2076 FC Disk
hdisk7 Available Remote Physical Volume Client
hdisk8 Available Remote Physical Volume Client
hdisk9 Available Remote Physical Volume Client
hdisk10 Available Remote Physical Volume Client
pc-lon06-glvms-01a:/ >

```

Figure 2.3.8.1 Displaying RPV clients

Note the hdisks, hdisk7 to hdisk10, are shown as “Available” and are “Remove Physical Volume Client”.

### 2.3.9 Import all VGs to remote PowerVS

Once we verify that RPV servers, clients are defined, working and synchronized, we can import the VGs to the remaining PowerVS. All VGs that will be part of the PowerHA resource group must be available on all nodes.

We display the importvg command, “importvg -y appvg -V 101 hdisk9” setting the major number as 101 and reading the VG information from hdisk9. We repeat the similar command for “dbvg”.

```
pc-lon06-glvn-01a:/ >  
# importvg -y appvg -V 101 hdisk9  
appvg  
pc-lon06-glvn-01a:/ >
```

Figure 2.3.9.1 Importing VGs to London PowerVS

### 2.3.10 Installing IBM PowerHA EE

The IBM PowerHA Enterprise Edition filesets will need to be installed on all AIX PowerVS instances that will be part of the cluster.

In our example, we placed the PowerHA EE filesets in a filesystem named “/tools”. We “cd” to that directory and execute “smitty install\_latest”. Since we are in that directory, we can enter “.” As the “INPUT device”.

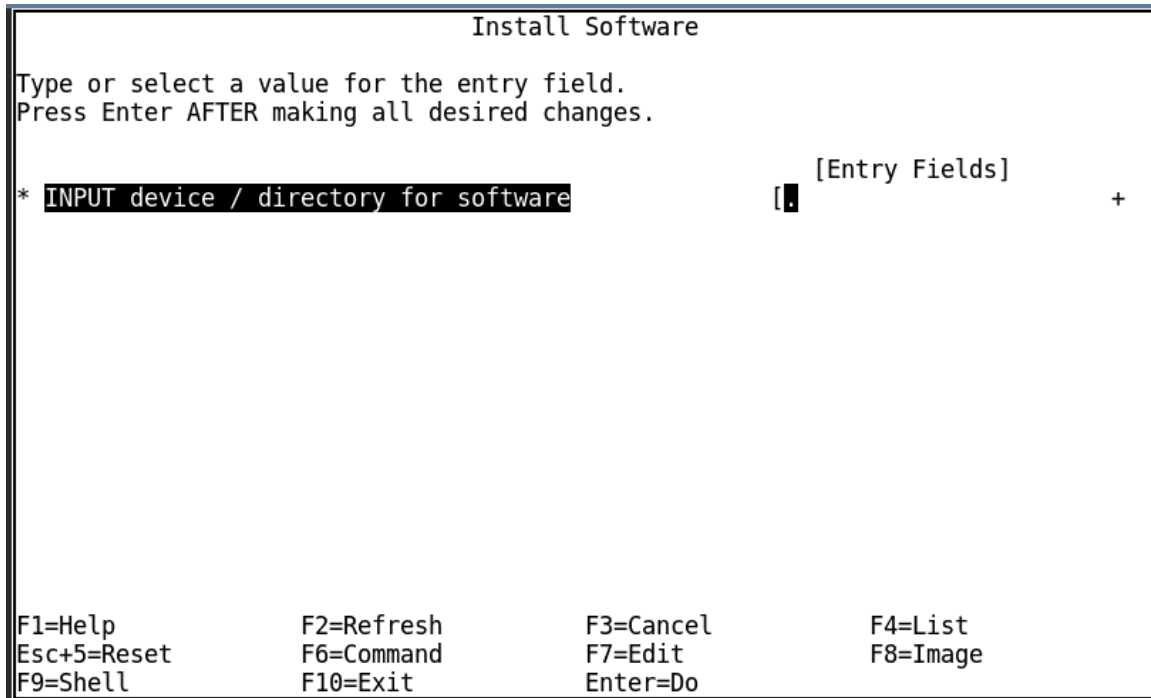


Figure 2.3.10.1 Directory of AIX PowerHA software location

The following filesets will be installed:

- *cluster.adt.es*
- *cluster.doc.en\_US.assist*
- *cluster.doc.en\_US.es*
- *cluster.doc.en\_US.glv*
- *cluster.es.client*
- *cluster.es.cspoc*
- *cluster.es.genxd*
- *cluster.es.server*
- *cluster.license*
- *cluster.man.en\_US.es*
- *cluster.xd.base*
- *cluster.xd.glv*
- *cluster.xd.license*

Note that there are many more additional packages that do not apply such as those for EMC, Hitachi. Also note that there may be additional

packages that may apply to your configuration such as the Smart Assists. Selected the minimum packages for AIX GLVM.

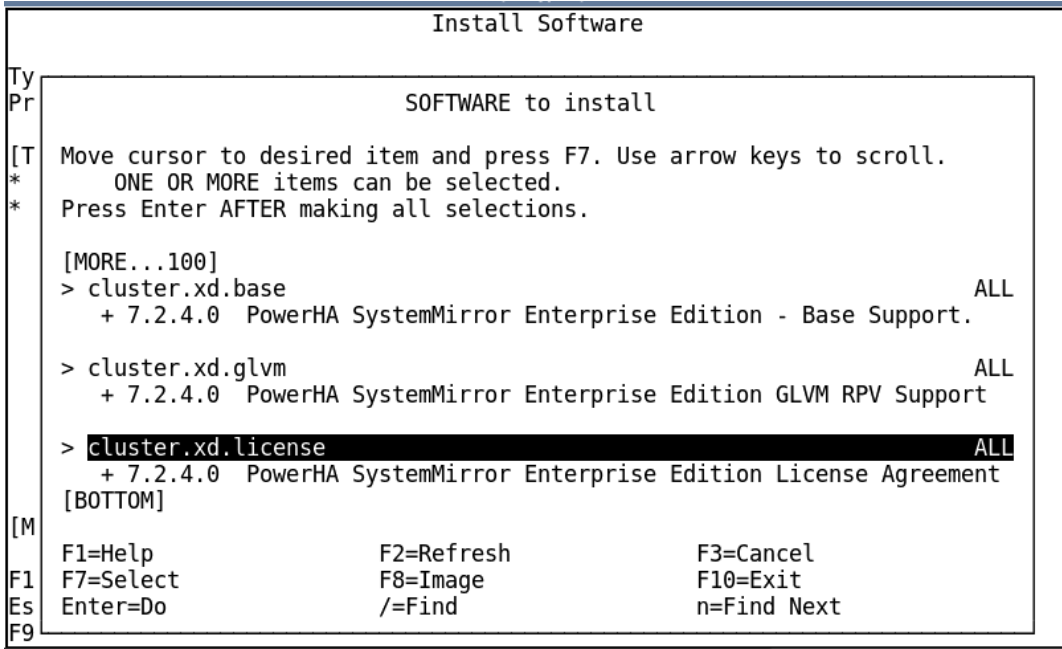


Figure 2.3.10.2 Selecting PowerHA packages to install

We verify that "ACCEPT new license agreements?" field is set to "yes".

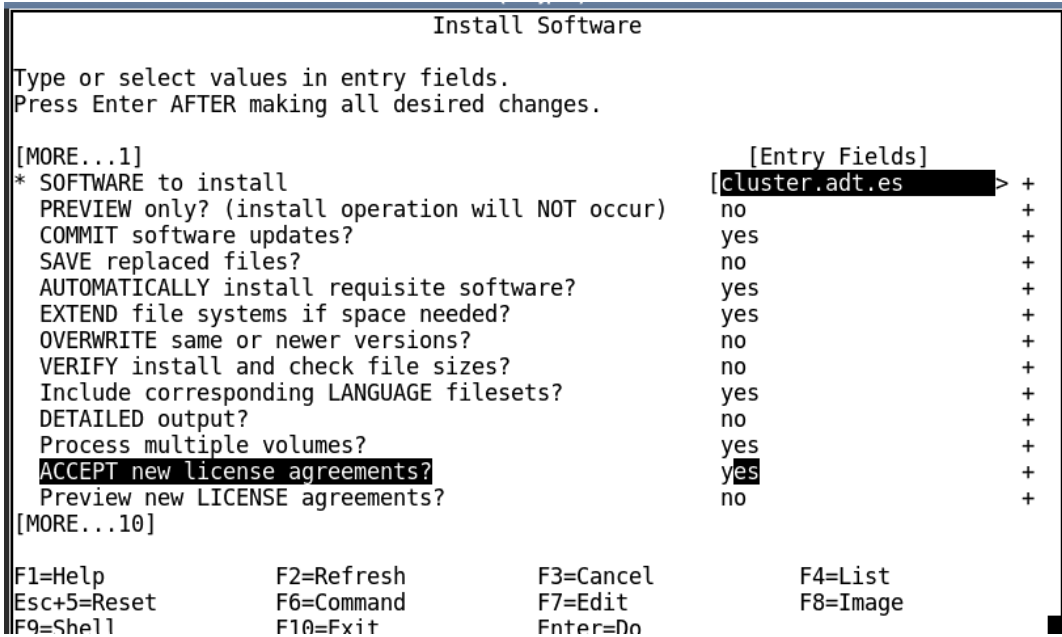


Figure 2.3.10.3 Verifying "ACCEPT new license agreements?" is "yes"

After proceeding, we verify the installation is completed successfully by scrolling through the SMIT output.

```
COMMAND STATUS
Command: OK          stdout: yes          stderr: no
Before command completion, additional instructions may appear below.
[MORE...470]
cluster.msg.en_US.glvms 7.2.4.0      USR          APPLY       SUCCESS
cluster.msg.en_US.es.server 7.2.4.0     USR          APPLY       SUCCESS
cluster.es.genxd.rte     7.2.4.0     USR          APPLY       SUCCESS
cluster.es.genxd.cmds    7.2.4.0     USR          APPLY       SUCCESS
cluster.es.genxd.rte     7.2.4.0     ROOT        APPLY       SUCCESS
cluster.es.genxd.cmds    7.2.4.0     ROOT        APPLY       SUCCESS
cluster.msg.en_US.genxd  7.2.4.0     USR          APPLY       SUCCESS

File /etc/environment has been modified.
File /etc/group has been modified.
File /etc/inittab has been modified.
File /etc/services has been modified.
[MORE...4]

F1=Help          F2=Refresh      F3=Cancel      F6=Command
F8=Image         F9=Shell        F10=Exit       /=Find
n=Find Next
```

Figure 2.3.10.4 Verifying installation output

Repeat the process for all cluster PowerVS.

### 2.3.11 Configuring PowerHA EE

As mentioned previously, we will configure a very simple PowerHA cluster with:

- *two nodes (one Toronto PowerVS, one London PowerVS)*
- *single resource group (we will call this RG01)*
- *no service IP*
- *no application controllers*
- *two VGs (dbvg, datavg)*

### 2.3.12 Verify /etc/cluster/rhosts entries

The file, /etc/cluster/rhosts, contains entries that are used for communication purposes. These entries correspond to the hostname of each of the PowerVS nodes. For our purposes, the file looks like this.

```
pc-tor01-glvm-01a: / >  
# cat /etc/cluster/rhosts  
pc-tor01-glvm-01a  
pc-lon06-glvm-01a  
pc-tor01-glvm-01a: / >  
#
```

Figure 2.3.12.1 /etc/cluster/rhosts entries

That is assuming those entries are also located in the /etc/hosts file as shown below.

```
pc-tor01-glvm-01a: / >  
# hostent -S  
127.0.0.1          loopback localhost      # loopback (lo0) name/address  
192.168.6.163    pc-tor01-glvm-01a tor01  
192.168.50.184  pc-lon06-glvm-01a lon06  
pc-tor01-glvm-01a: / >
```

Figure 2.3.12.2 /etc/hosts entries

After changing the /etc/cluster/rhosts file, the daemon, clcomd, needs to be refreshed. This is done with the "stopsrc -s clcomd;startsrc -s clcomd" command as shown.

```
pc-tor01-glvm-01a: / >  
# stopsrc -s clcomd;startsrc -s clcomd  
0513-044 The clcomd Subsystem was requested to stop.  
0513-059 The clcomd Subsystem has been started. Subsystem PID is 7930116.
```

Figure 2.3.12.3 Restarting "clcomd" daemon

Lastly, the communication between PowerVS images can be verified with the "cl\_rsh" command. In this example, we add the "cl\_rsh" directory to the PATH variable before executing the "cl\_rsh" command.

```

pc-tor01-qlvm-01a: / >
# export PATH=$PATH:/usr/es/sbin/cluster/utilities
pc-tor01-qlvm-01a: / >
# cl_rsh pc-lon06-qlvm-01a date
Fri Jul 17 22:41:00 CDT 2020
pc-tor01-qlvm-01a: / >

```

Figure 2.3.12.4 Using “cl\_rsh” to verify communications

### 2.3.13 Determine the PowerHA repository disk(s)

In our example configuration, we only have two VGs, dbvg and appvg. Any of the other LUNs can be used as repository disks. We note those drives on each of the PowerVS images.

```

pc-tor01-qlvm-01a: / >
# lspv
hdisk0          none                None
hdisk1          00c8d4403180f273      dbvg
hdisk2          00c8d4403180f319      dbvg
hdisk3          00f6db0af58e9775      rootvg                active
hdisk4          none                None
hdisk5          00c8d44031815edd      appvg
hdisk6          00c8d44031815f5a      appvg
hdisk7          00c8cf803fb0193a      dbvg
hdisk8          00c8cf803fb02116      dbvg
hdisk9          00c8cf803fb029cd      appvg
hdisk10         00c8cf803fb03444      appvg

```

Figure 2.3.13.1 Displaying Toronto available disks (hdisk0, hdisk4)

```

pc-lon06-qlvm-01a: / >
# lspv
hdisk0          none                None
hdisk1          00c8cf803fb0193a      None
hdisk2          00c8cf803fb02116      None
hdisk3          00f6db0af58e9775      rootvg                active
hdisk4          00c8cf8049bac9fa      None
hdisk5          00c8cf803fb029cd      None
hdisk6          00c8cf803fb03444      None
hdisk7          00c8d4403180f273      dbvg
hdisk8          00c8d4403180f319      dbvg
hdisk9          00c8d44031815edd      None
hdisk10         00c8d44031815f5a      None

```

Figure 2.3.13.2 Displaying London available disks (hdisk0, hdisk4)

For our purposes, we will select “hdisk4” on each of our PowerVS images.



### 2.3.14 Configure PowerHA topology

We begin the PowerHA configuration by executing "smitty sysmirror -> Cluster Nodes and Networks -> Multi Site Cluster Deployment -> Setup a Cluster, Nodes and Networks". We name the cluster "tor01lon06", set the site 1 name to "toronto01", the site 2 name to "london06". We select the "New Nodes" by pressing "F4" and selecting the appropriate node name.

```
Setup Cluster, Sites, Nodes and Networks

Ty
Pr
New Nodes (via selected communication paths)

Move cursor to desired item and press F7.
*   ONE OR MORE items can be selected.
*   Press Enter AFTER making all selections.
*
*   [TOP]
*   pc-tor01-glvm-01a (192.168.6.163)
*   pc-lon06-glvm-01a (192.168.50.184)
*
#
# If you do not see a selection you were expecting
# consider the following:
#
[MORE...3]

F1=Help           F2=Refresh       F3=Cancel
F7=Select         F8=Image         F10=Exit
Enter=Do          /=Find           n=Find Next
F9
```

Figure 2.3.14.1 Select "New Nodes" in base PowerHA configuration

This leads to the following menu where we select the "Cluster Type" of "Linked Cluster".

```

Setup Cluster, Sites, Nodes and Networks

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
* Cluster Name                    [tor01lon06]
* Site 1 Name                      [toronto01]
* New Nodes (via selected communication paths) [pc-tor01-glvm-01a] +
* Site 2 Name                      [london06]
* New Nodes (via selected communication paths) [pc-lon06-glvm-01a] +

Cluster Type                       [Linked Cluster] +

F1=Help      F2=Refresh      F3=Cancel      F4=List
Esc+5=Reset  F6=Command      F7=Edit       F8=Image
F9=Shell     F10=Exit       Enter=Do

```

Figure 2.3.14.2 Configuring PowerHA nodes, sites, cluster type

The result should be the gathering of cluster data as shown in the following diagram.

```

Setup Cluster, Sites, Nodes and Networks

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
* Cluster Name                    [tor01lon06]
* Site 1 Name                      [toronto01]
* New Nodes (via selected communication paths) [pc-tor01-glvm-01a] +
* Site 2 Name                      [london06]
* New Nodes (via selected communication paths) [pc-lon06-glvm-01a] +

Cluster Type                       [Linked Cluster] +

F1=Help      F2=Refresh      F3=Cancel      F4=List
Esc+5=Reset  F6=Command      F7=Edit       F8=Image
F9=Shell     F10=Exit       Enter=Do

```

Figure 2.3.14.3 Output of base cluster configuration

### 2.3.15 Configure PowerHA repository disks

We proceed to configuring the PowerHA repository disks by executing "smitty sysmirror -> Cluster Nodes and Networks -> Multi Site Cluster Deployment -> Define the Repository Disk and Cluster IP Address". We proceed to selecting the "Repository Disk" for each of the sites by pressing the "F4" key on those fields.

```
Multi Site with Linked Clusters Configuration
Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]
* Cluster Name          tor01lon06
* Heartbeat Mechanism   Unicast          +
* Site Name             toronto01
* Repository Disk      []                +
* Site Multicast Address []

Repository Disk
*
* Move cursor to desired item and press Enter.
  hdisk0 (00c8d440601a6595) on all nodes at site toronto01
  hdisk4 (00c8d440601a64c7) on all nodes at site toronto01

F1=Help          F2=Refresh      F3=Cancel
F8=Image         F10=Exit        Enter=Do
/=Find          n=Find Next

F1
Es
F9
```

Figure 2.3.15.1 Selecting repository disks

We repeat this process until all of the repository disks are selected.

```

Multi Site with Linked Clusters Configuration

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]
* Cluster Name          tor01lon06
* Heartbeat Mechanism   Unicast          +
* Site Name             toronto01
* Repository Disk       [(00c8d440601a64c7)] +
  Site Multicast Address [ ]
  (used only for multicast heartbeat)
* Site Name             london06
* Repository Disk       [(00c8cf8049bac9fa)] +
  Site Multicast Address [ ]
  (used only for multicast heartbeat)

F1=Help      F2=Refresh    F3=Cancel    F4=List
Esc+5=Reset  F6=Command    F7=Edit      F8=Image
F9=Shell     F10=Exit      Enter=Do

```

Figure 2.3.15.2 Defining repository disks

The resulting output should look similar to this diagram.

```

COMMAND STATUS

Command: OK          stdout: yes          stderr: no

Before command completion, additional instructions may appear below.

[TOP]
  successfully added a primary repository disk.
To view the complete configuration of repository disks use:
"clmgr query repository" or "clmgr view report repository"
Successfully added a primary repository disk.
To view the complete configuration of repository disks use:
"clmgr query repository" or "clmgr view report repository"
Initial site configuration has been saved. You can now go on
to complete the rest of the configuration, including adding backup
repository disks (recommended), custom event notifications, resource
groups and applications, etc.

When you have entered all the basic information you can then use
[MORE...3]

F1=Help      F2=Refresh    F3=Cancel    F6=Command
F8=Image     F9=Shell      F10=Exit     /=Find
n=Find Next

```

Figure 2.3.15.3 Output of defining repository disks

### 2.3.16 Perform initial PowerHA verification

After defining the repository disks, we can perform the initial PowerHA verification. Before performing this, note that there is no CAA volume group.

```
pc-tor01-glvn-01a: / >
# lspv
hdisk0          00c8d440601a6595          None
hdisk1          00c8d4403180f273          dbvg
hdisk2          00c8d4403180f319          dbvg
hdisk3          00f6db0af58e9775          rootvg          active
hdisk4          00c8d440601a64c7          None
hdisk5          00c8d44031815edd          appvg
hdisk6          00c8d44031815f5a          appvg
hdisk7          00c8cf803fb0193a          dbvg
hdisk8          00c8cf803fb02116          dbvg
hdisk9          00c8cf803fb029cd          appvg
hdisk10         00c8cf803fb03444          appvg
```

Figure 2.3.16.1 VGs before initial PowerHA verification

We proceed with the initial PowerHA verification by executing “smitty sysmirror -> Cluster Nodes and Networks -> Verify and Synchronize Cluster Configuration”. This will present a verification screen.

```
Cluster Nodes and Networks
Move cursor to desired item and press Enter.

Verify and Synchronize Cluster Configuration
Move cursor to desired item and press Enter. Use arrow keys to scroll.

Any time you change the cluster configuration, those changes
must be propagated to all nodes in the cluster. PowerHA SystemMirror w
perform a set of verification checks to make sure the configuration
is viable.

Depending on the nature of the changes, this operation can disrupt
the applications currently under PowerHA SystemMirror control.

Press enter when you are ready to begin.

F1=Help          F2=Refresh      F3=Cancel
F8=Image         F10=Exit       Enter=Do
F1 /=Find        n=Find Next
```

Figure 2.3.16.2 Enter to begin PowerHA verification

The resulting output should look something like the following.

```

COMMAND STATUS
Command: OK          stdout: yes          stderr: no
Before command completion, additional instructions may appear below.
[TOP]
Verification to be performed on the following:
    Cluster Topology
    Cluster Resources
Verification will interactively correct verification errors.
Retrieving data from available cluster nodes.  This could take a few minutes.
    Start data collection on node pc-lon06-glvm-01a
    Start data collection on node pc-tor01-glvm-01a
[MORE...153]
F1=Help          F2=Refresh          F3=Cancel          F6=Command
F8=Image         F9=Shell            F10=Exit           /=Find
n=Find Next

```

Figure 2.3.16.3 Output of PowerHA verification

If PowerHA verification was successful, there should now be an additional volume group, "caavg\_private", on each PowerVS instance.

```

pc-tor01-glvm-01a: / >
# lspv
hdisk0          00c8d440601a6595          None
hdisk1          00c8d4403180f273          dbvg
hdisk2          00c8d4403180f319          dbvg
hdisk3          00f6db0af58e9775          rootvg          active
hdisk4          00c8d440601a64c7          caavg_private   active
hdisk5          00c8d44031815edd          appvg
hdisk6          00c8d44031815f5a          appvg
hdisk7          00c8cf803fb0193a          dbvg
hdisk8          00c8cf803fb02116          dbvg
hdisk9          00c8cf803fb029cd          appvg
hdisk10         00c8cf803fb03444          appvg

```

Figure 2.3.16.4 "caavg\_private" VG at Toronto PowerVS

```

pc-lon06-glvm-01a:/ >
# lspv | grep caa
hdisk4          00c8cf8049bac9fa          caavg_private   active
pc-lon06-glvm-01a:/ >

```

Figure 2.3.16.5 "caavg\_private" VG at London PowerVS

### 2.3.17 Defining an XD\_data network

AIX GLVM can use a couple of specific networks, XD\_data and XD\_ip. These networks have these characteristics

- *XD\_data* - An IP-based network used by geographically mirrored volume groups in a PowerHA SystemMirror cluster for transferring the data between the RPV devices.
- *XD\_ip* - An IP-based network used for participation in RSCT protocols, heartbeating, and client communication.

PowerHA builds the default IP networks without configuring either XD\_data or XD\_ip networks. We will need to create the XD\_data network and add the interfaces to it

To create the XD\_data network, execute "smitty sysmirror -> Cluster Nodes and Networks -> Manage Networks and Network Interfaces -> Networks -> Add a Network", then select "XD\_data".

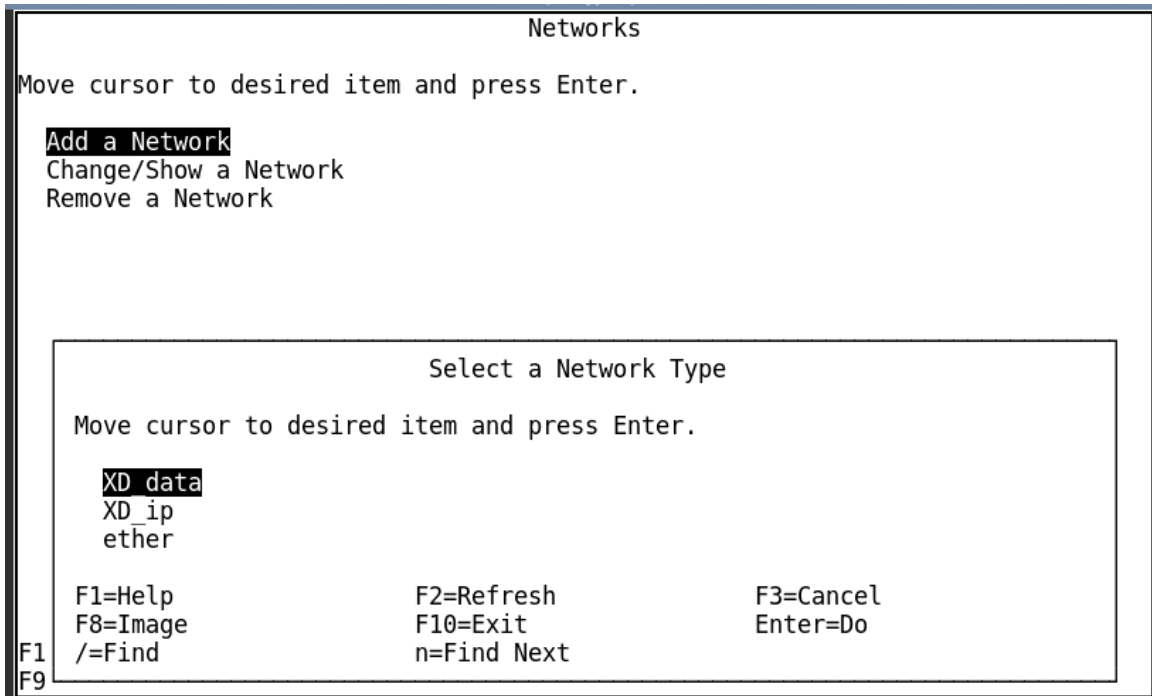


Figure 2.3.17.1 Selecting the XD\_data network

Then verify the netmask is the same as the network interfaces.

```

                                Add a Network

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
* Network Name                    [ et_XD_data_01]
* Network Type                    XD_data
* Netmask(IPv4)/Prefix Length(IPv6) [255.255.255.0]
* Network attribute                public          +
* Unstable Threshold              [3]          #
* Unstable Period (seconds)       [60]         #

F1=Help          F2=Refresh      F3=Cancel      F4=List
Esc+5=Reset     F6=Command    F7=Edit       F8=Image
F9=Shell        F10=Exit      Enter=Do

```

Figure 2.3.17.2 XD\_data network parameters

### 2.3.18 Moving network interfaces to XD\_data network

In order to utilize the XD\_data network that was created for the AIX GLVM replication, the network interfaces have to be moved to that network. Execute "smitty sysmirror -> Cluster Nodes and Networks -> Manage Networks and Network Interfaces -> Network Interfaces -> Change/Show a Network Interface", we select a network interface. Here, we select "en0" at "pc-lon06-glvm-01a"



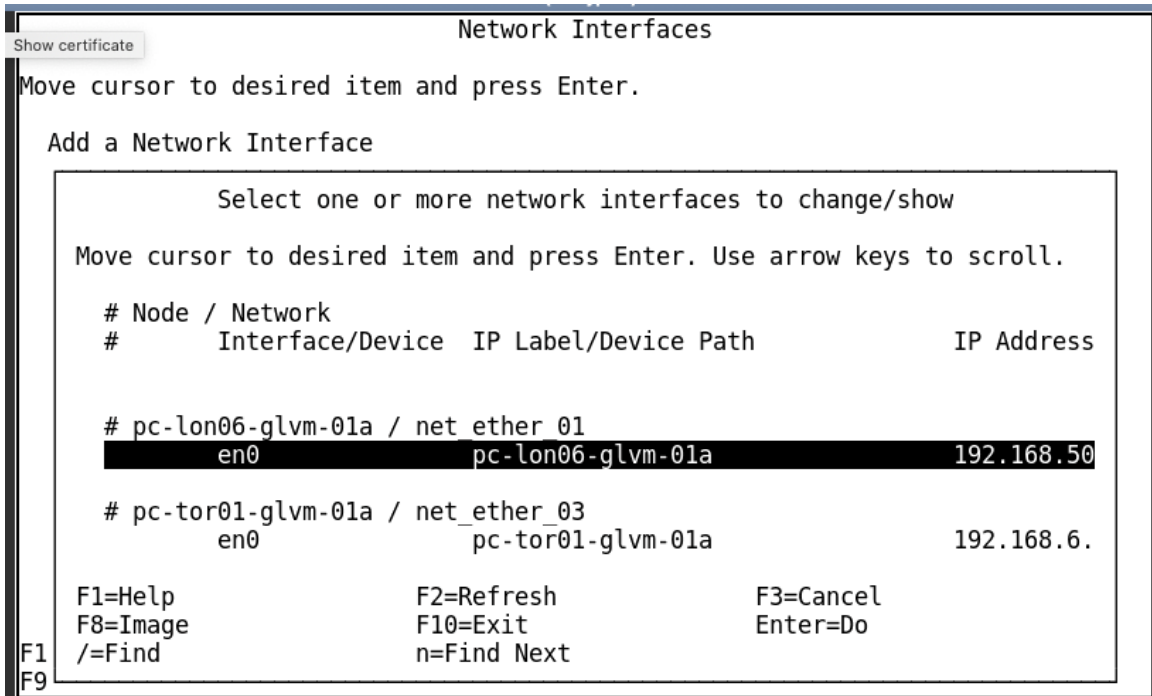


Figure 2.3.18.1 Selecting a network interface

We then get the following menu where we can select the network the interface is on.

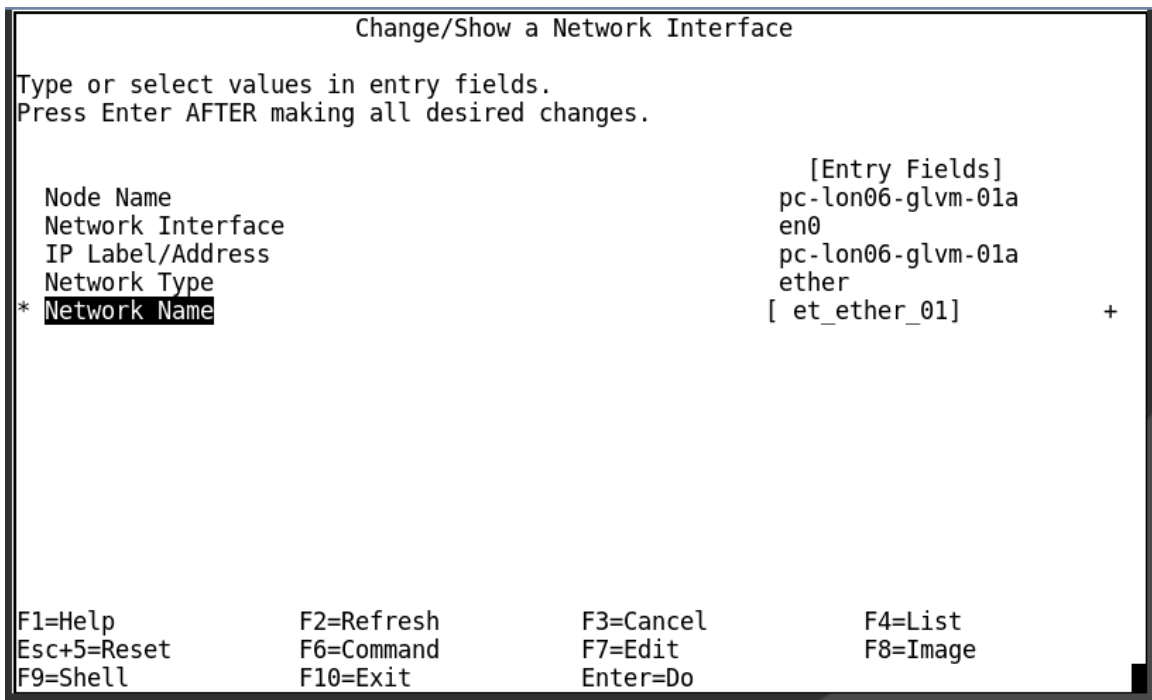


Figure 2.3.18.2 Changing "Network Name" from "net\_ether\_01"

We change the network name to the XD\_data network that we created in the previous step.

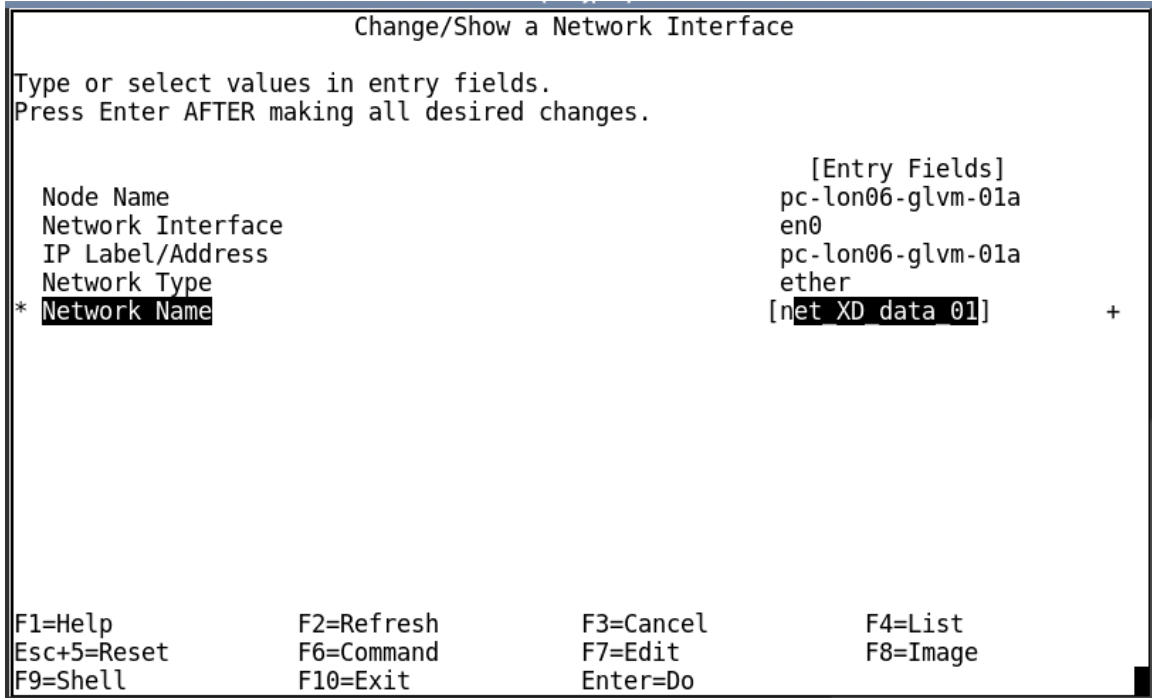


Figure 2.3.18.3 Changing "Network Name" to "net\_XD\_data\_01"

We get a confirmation menu and proceed to the rest of the interfaces.

### 2.3.19 Defining PowerHA resource group

As mentioned previously, the example configuration will include a single resource group. We define this resource group by executing "smitty sysmirror -> Cluster Applications and Resources -> Resource Groups -> Add a Resource Group". On this menu, we give it a resource group name, RG01 for our example, as well as the participating nodes. We change the "Inter-Site Management Policy" to "Prefer Primary Site". Lastly we select the "Fallback Policy" to "Never Fallback".

```

Add a Resource Group (extended)

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
* Resource Group Name           [RG01]
Inter-Site Management Policy    [Prefer Primary Site] +
* Participating Nodes from Primary Site [pc-tor01-glv-01a] +
Participating Nodes from Secondary Site [pc-lon06-glv-01a] +

Startup Policy                  Online On Home Node 0> +
Fallover Policy                 Fallover To Next Prio> +
Fallback Policy                 Never Fallback          +

F1=Help      F2=Refresh  F3=Cancel    F4=List
Esc+5=Reset  F6=Command  F7=Edit      F8=Image
F9=Shell     F10=Exit    Enter=Do

```

Figure 2.3.19.1 Define a PowerHA resource group

A successfully defined resource group should give the SMIT output similar to this.

```

COMMAND STATUS

Command: OK          stdout: no          stderr: no

Before command completion, additional instructions may appear below.

F1=Help      F2=Refresh  F3=Cancel    F6=Command
F8=Image     F9=Shell    F10=Exit     /=Find
n=Find Next

```

Figure 2.3.19.2 Successfully defined a resource group

### 2.3.20 Adding resources to the PowerHA resource group

The example configuration will only include the VGs that are being replicated as the resources of the PowerHA resource group defined in the previous step. Again, for more complex PowerHA configuration, review the resources described in section "Chapter 4: Additional Resources".

To add resources to the resource group, execute "smitty sysmirror -> Cluster Applications and Resources -> Resource Groups -> Change/Show Resources and Attributes for a Resource Group". Select the appropriate resource group, "RG01" in our example.

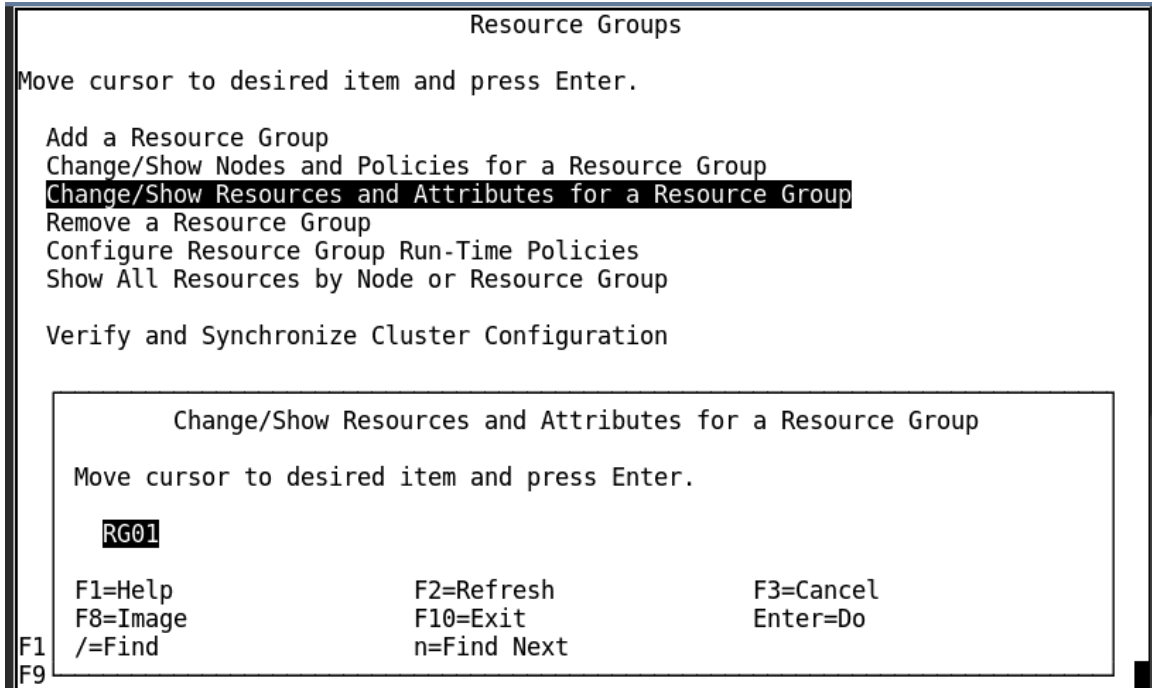


Figure 2.3.20.1 Selecting RG to add resources

On the resource group fields, enter the appropriate VGs on the "Volume Groups" field by pressing "F4".

```

Change/Show All Resources and Attributes for a Resource Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[TOP]
Resource Group Name
Inter-site Management Policy
Participating Nodes from Primary Site
Participating Nodes from Secondary Site

[Entry Fields]
RG01
ignore
pc-tor01-glvm-01a
pc-lon06-glvm-01a

Volume Groups

Move cursor to desired item and press F7.
ONE OR MORE items can be selected.
Press Enter AFTER making all selections.

> dbvg
> appvg

[M]
F1=Help          F2=Refresh      F3=Cancel
F7=Select        F8=Image        F10=Exit
Es Enter=Do      /=Find          n=Find Next
F9

```

Figure 2.3.20.2 Selecting VGs to add to resource group

In our example, we will leave the other fields as default including the “Service IP labels/Addresses” and “Application Controller Name” fields.

```

Change/Show All Resources and Attributes for a Resource Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[TOP]
Resource Group Name
Inter-site Management Policy
Participating Nodes from Primary Site
Participating Nodes from Secondary Site

Startup Policy
Failover Policy
Fallback Policy

Service IP Labels/Addresses
Application Controller Name

Volume Groups
[MORE...30]

[Entry Fields]
RG01
ignore
pc-tor01-glvm-01a
pc-lon06-glvm-01a

Online On Home Node 0>
Failover To Next Prio>
Never Fallback

[] +
[] +
[dbvg appvg] +

F1=Help          F2=Refresh      F3=Cancel      F4=List
Esc+5=Reset      F6=Command      F7=Edit        F8=Image
F9=Shell         F10=Exit        Enter=Do

```

Figure 2.3.20.3 Adding VGs to resource group RG01

A successful modification should result in a similar screen to this.

```
COMMAND STATUS
Command: OK          stdout: no          stderr: no
Before command completion, additional instructions may appear below.

F1=Help          F2=Refresh      F3=Cancel       F6=Command
F8=Image         F9=Shell        F10=Exit        /=Find
n=Find Next
```

Figure 2.3.20.4 Successful RG change

### 2.3.21 Perform an additional PowerHA verification

We proceed with an additional PowerHA verification by executing "smitty sysmirror -> Cluster Nodes and Networks -> Verify and Synchronize Cluster Configuration". This will present a verification screen.

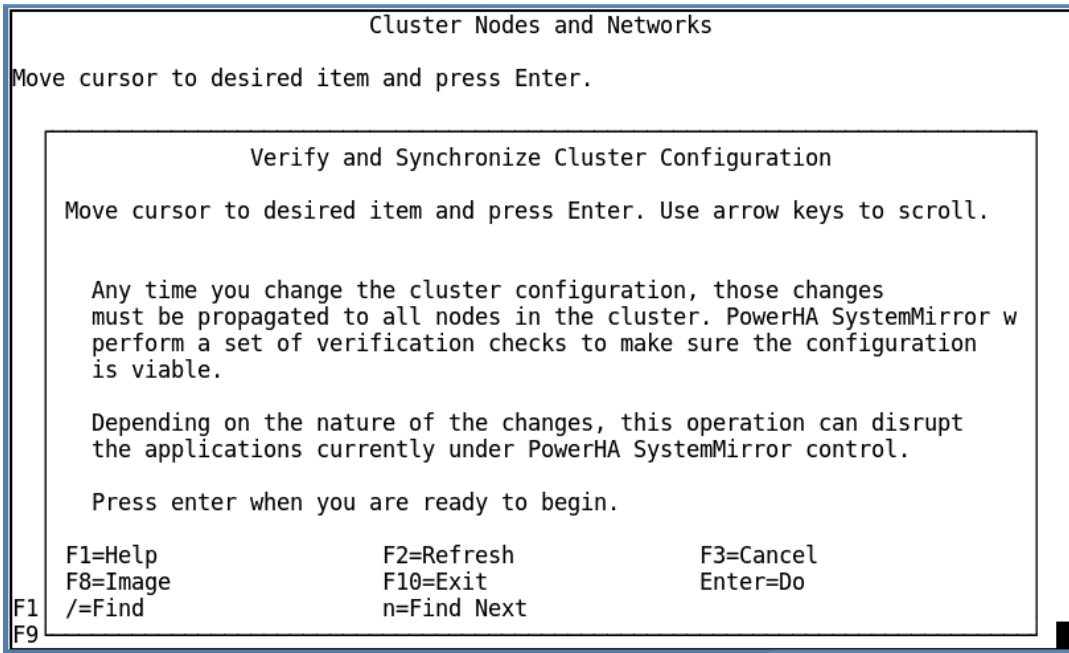


Figure 2.3.21.1 Enter to begin PowerHA verification

The resulting output should look something like the following.

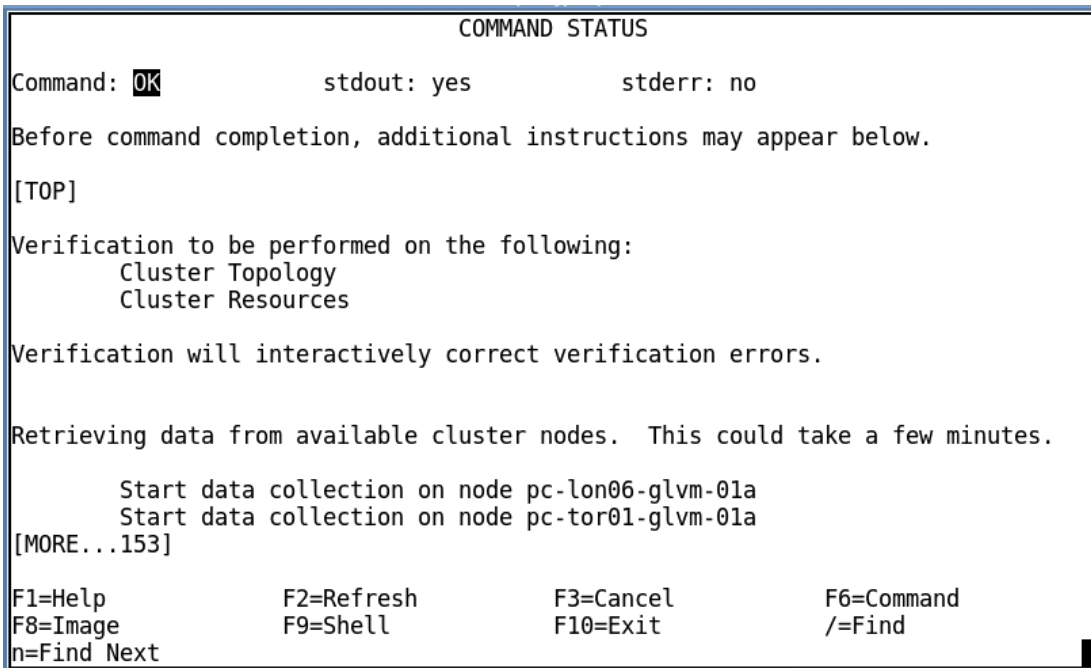


Figure 2.3.21.2 Output of PowerHA verification

### 2.3.22 Starting PowerHA EE

Once PowerHA EE is fully configured, we can start the cluster by starting cluster services. For our example, we start PowerHA by executing "smitty clstart" and change the fields "Startup Cluster Information Daemon" to "true" and "Automatically correct errors found during cluster start?" to "Yes".

```
Start Cluster Services

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
* Start now, on system restart or both      now          +
  Start Cluster Services on these nodes    [pc-tor01-glv-01a] +
* Manage Resource Groups                    Automatically +
  BROADCAST message at startup?            false         +
  Startup Cluster Information Daemon?       true          +
  Ignore verification errors?              false         +
  Automatically correct errors found during cluster start? Yes          +

F1=Help      F2=Refresh      F3=Cancel      F4=List
Esc+5=Reset  F6=Command      F7=Edit        F8=Image
F9=Shell     F10=Exit         Enter=Do
```

Figure 2.3.22.1 Starting PowerHA services on primary node

The output of this command should be similar to the following screen.



```

COMMAND STATUS
Command: OK          stdout: yes          stderr: no
Before command completion, additional instructions may appear below.
[TOP]
Verifying Cluster Configuration Prior to Starting Cluster Services.
Verifying Cluster Configuration Prior to Starting Cluster Services.
There are no active cluster nodes to verify against.
Verifying node(s): pc-tor01-glvm-01a requested to start
WARNING: No backup repository disk is UP and not already part of a VG for nodes
:
WARNING: For PowerHA SystemMirror to perform monitoring over a network at least
two interfaces that reside on separate nodes must be defined to a
[MORE...37]
F1=Help          F2=Refresh          F3=Cancel          F6=Command
F8=Image         F9=Shell            F10=Exit           /=Find
n=Find Next

```

Figure 2.3.22.2 Output of starting PowerHA services

We repeat this process for all cluster nodes.

### 2.3.23 Verifying PowerHA EE is operating normally

After starting PowerHA EE on all nodes, the AIX GLVM resources, dbvg and appvg, should be on the appropriate node. We log into the primary PowerVS and perform a “df” command.

```

# df
Filesystem      512-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4         196608          105208  47%      3073   21% /
/dev/hd2         4784128         329872   94%     40770  50% /usr
/dev/hd9var      1966080         1646824  17%     1783   1% /var
/dev/hd3         458752          453696   2%        52    1% /tmp
/dev/hd1         65536           64784    2%        8    1% /home
/dev/hd11admin   262144          261384   1%        5    1% /admin
/proc           -               -         -         -     - /proc
/dev/hd10opt     786432          85096    90%     11482  53% /opt
/dev/livedump    524288          523016   1%        5    1% /var/adm/ras/livedump
/dev/repo00      15335424        66288   100%     2533  25% /usr/sys/inst.images
/ahafs          -               -         -         51    1% /aha
/dev/lv00        2621440         889232   67%      23    1% /tools
/dev/loop0       553760          0 100%    138440 100% /mnt
/dev/fslv07      6291456         5758640   9%      2092   1% /appfs1
/dev/fslv05      6291456         5313400  16%      3292   1% /appfs2
/dev/fslv06      6291456         4922056  22%      4189   1% /appfs3
/dev/fslv00      4194304         4084024   3%      1016   1% /data01
/dev/fslv01      4194304         3598480  15%      1945   1% /data02
/dev/fslv02      4194304         4000328   5%       826   1% /data03
/dev/fslv03      4194304         4193008   1%        4    1% /data04

```

Figure 2.3.23.1 Displaying shared VG mounted filesystems

Note on the above figure that the "dbvg" and "appvg" filesystems are mounted. An additional PowerHA command to display resource status"

- `clmgr q RG <RGNAME>`

```
pc-tor01-glv-01a: / >
# clmgr q rg RG01 | head -15
NAME="RG01"
CURRENT_NODE="pc-tor01-glv-01a"
NODES="pc-tor01-glv-01a"
STATE="ONLINE"
CURRENT_SECONDARY_NODE="pc-lon06-glv-01a"
SECONDARYNODES="pc-lon06-glv-01a"
SECONDARY_STATE="ONLINE SECONDARY"
TYPE="non-concurrent"
APPLICATIONS=""
STARTUP="OHN"
FALLOVER="FNPN"
FALLBACK="NFB"
NODE_PRIORITY_POLICY="default"
NODE_PRIORITY_POLICY_SCRIPT=""
NODE_PRIORITY_POLICY_TIMEOUT=""
pc-tor01-glv-01a: / >
```

Figure 2.3.23.2 Displaying resource group status

The above figure shows the resource group "RG01" being "ONLINE" on node "pc-tor01-glv-01a" with secondary being "pc-lon06-glv-01a".

The "clmgr q cluster" command also displays the cluster status. The figure below shows that the cluster is in "STABLE" status.

```
r01-glv-01a: / >
# clmgr q cluster | head -15
CLUSTER_NAME="tor01lon06"
CLUSTER_ID="1372642113"
STATE="STABLE"
TYPE="LC"
HEARTBEAT_TYPE="UNICAST"
VERSION="7.2.4.0"
VERSION_NUMBER="20"
EDITION="ENTERPRISE"
UNSYNCED_CHANGES="false"
FC_SYNC_INTERVAL="10"
RG_SETTLING_TIME="0"
RG_DIST_POLICY="node"
MAX_EVENT_TIME="180"
MAX_RG_PROCESSING_TIME="180"
TEMP_HOSTNAME="disallow"
pc-tor01-glv-01a: / >
```

Figure 2.3.23.3 Querying the cluster status

### 2.3.24 Testing PowerHA

Once the cluster status is "STABLE" and all nodes are online and the RG is also online, we can perform a failover test.

For our scenario, we will simply reboot the primary node, pc-tor01-glvm-01a, using the "reboot -q" command. Note that we cannot use the "shutdown" command as PowerHA has a "hook" into the "shutdown" command and will know that was performed intentionally. Since PowerHA assumes a "shutdown" command is executed intentionally, it will not failover to the next available node. That is the reason we use the "reboot -q" command.

Again, before rebooting the primary node, pc-tor01-glvm-01a, we check to see the VGs are operational.

```
pc-tor01-glvm-01a: / >
# lsvg -o | grep -E "db|app";df | grep -E "appfs|data"
dbvg
appvg
/dev/fslv07      6291456  5758640    9%    2092    1% /appfs1
/dev/fslv05      6291456  5313400   16%    3292    1% /appfs2
/dev/fslv06      6291456  4922056   22%    4189    1% /appfs3
/dev/fslv00      4194304  4084024    3%    1016    1% /data01
/dev/fslv01      4194304  3598480   15%    1945    1% /data02
/dev/fslv02      4194304  4000328    5%     826    1% /data03
/dev/fslv03      4194304  4193008    1%         4    1% /data04
```

Figure 2.3.24.1 VGs operational at primary node

We also verify that the VGs are not operational on the standby node, pc-lon06-glvm-01a.

```
pc-lon06-glvm-01a: / >
# lsvg -o | grep -E "db|app";df | grep -E "appfs|data"
```

Figure 2.3.24.2 VGs not operational at standby node

On the primary node, we execute the "reboot -q" command.

On the standby node, we watch that the node reacts by monitoring the "hacmp.out" log file using the command "tail -f /var/hacmp/log/hacmp.out".

```

+RG01:clvaryonvg(169.302):appvg[updatefs:600] clodmget -q 'name = fslv07 and attribute = type and value = raw' -f value -n CuAt
+RG01:clvaryonvg(169.305):appvg[updatefs:600] [[ -n '' ]]
+RG01:clvaryonvg(169.305):appvg[updatefs:605] : Skip logical volumes for which getlvcb fails
+RG01:clvaryonvg(169.305):appvg[updatefs:607] /usr/sbin/getlvcb -f fslv07
+RG01:clvaryonvg(174.037):appvg[updatefs:607] fs_info=vfs='jfs2:log=/dev/loglv01:options=rw:account=false:mountguard=yes '
+RG01:clvaryonvg(174.037):appvg[updatefs:608] cmd_rc=0
+RG01:clvaryonvg(174.037):appvg[updatefs:608] typeset -i cmd_rc
+RG01:clvaryonvg(174.037):appvg[updatefs:609] ( 0 != 0 )
+RG01:clvaryonvg(174.037):appvg[updatefs:615] : Skip logical volumes not associated with file systems
+RG01:clvaryonvg(174.037):appvg[updatefs:618] [[ -z vfs='jfs2:log=/dev/loglv01:options=rw:account=false:mountguard=yes ' ]]
+RG01:clvaryonvg(174.037):appvg[updatefs:618] [[ vfs='jfs2:log=/dev/loglv01:options=rw:account=false:mountguard=yes ' == *([[:space:]]) ]]
+RG01:clvaryonvg(174.037):appvg[updatefs:623] : Label and file system type from LVCB on disk for fslv07
+RG01:clvaryonvg(174.038):appvg[updatefs:625] getlvcb -T -A fslv07
+RG01:clvaryonvg(174.038):appvg[updatefs:625] LC_ALL=C
+RG01:clvaryonvg(174.041):appvg[updatefs:625] egrep -w 'label =|type ='
+RG01:clvaryonvg(174.044):appvg[updatefs:625] paste -s - -

```

Figure 2.3.24.3 Watching “hacmp.out” on standby node

Eventually, PowerHA on the standby node finishes acquiring the resources and we can confirm that by looking at the status of the VGs status and the mounted filesystems.

```

pc-lon06-glv01a:/ >
# lsvg -o | grep -E "db|app";df | grep -E "appfs|data"
dbvg
appvg
/dev/fslv07      6291456    5758640     9%      2092      1% /appfs1
/dev/fslv05      6291456    5313400    16%      3292      1% /appfs2
/dev/fslv06      6291456    4922056    22%      4189      1% /appfs3
/dev/fslv00      4194304    4084024     3%      1016      1% /data01
/dev/fslv01      4194304    3598480    15%      1945      1% /data02
/dev/fslv02      4194304    4000328     5%        826      1% /data03
/dev/fslv03      4194304    4193008     1%         4      1% /data04

```

Figure 2.3.24.4 Displaying resources at DR node after failover

# Chapter 3: Troubleshooting

The following are some possible problems and resolutions.

## 3.1 ioctl call failed

When attempting to manipulate disks, you may receive errors that indicate “ioctl call failed” such as the figure below.

```
pc-tor01-glvm-01a: / >  
# extendvg -f -p appvg_lon06 appvg hdisk9 hdisk10  
0516-1395 extendvg: The physical volume hdisk9, is not supported.  
0516-1395 extendvg: The physical volume hdisk10, is not supported.  
0516-1941 extendvg: ioctl call failed on hdisk9.  
0516-1941 extendvg: ioctl call failed on hdisk10.  
0516-792 extendvg: Unable to extend volume group.
```

*Figure 3.1.1 ioctl call failed error*

This may indicate the RPV server is not operational on the remote location. Log into the remote location PowerVS and verify the RPV server is operational.

Here we log into our remote PowerVS and restart the RPV servers.

```
pc-lon06-glvm-01a:/ >  
# lsdev -Ccrpvserver  
rpvserver0 Available Remote Physical Volume Server  
rpvserver1 Available Remote Physical Volume Server  
rpvserver2 Defined Remote Physical Volume Server  
rpvserver3 Defined Remote Physical Volume Server  
pc-lon06-glvm-01a:/ >  
# mkdev -l rpvserver2  
rpvserver2 Available  
pc-lon06-glvm-01a:/ >  
# mkdev -l rpvserver3  
rpvserver3 Available  
pc-lon06-glvm-01a:/ >  
#
```

*Figure 3.1.2 Changing “Defined” RPV servers to “Available”*

### 3.2 Resuming a failed RPV client

If an RPV client has “failed” such as when we try to use it when the RPV server is down, it will stay in that “failed” state until it is resumed. In our example above (section 3.1), hdisk9 and hdisk10 failed due to the RPV server being down. The RPV servers were restarted and the RPV clients hdisk9 and hdisk10 still failed with the same message.

```
pc-tor01-glvm-01a: / >
# extendvg -f -p appvg_lon06 appvg hdisk9 hdisk10
0516-1395 extendvg: The physical volume hdisk9, is not supported.
0516-1395 extendvg: The physical volume hdisk10, is not supported.
0516-1941 extendvg: ioctl call failed on hdisk9.
0516-1941 extendvg: ioctl call failed on hdisk10.
0516-792 extendvg: Unable to extend volume group.
```

Figure 3.2.1 ioctl call failed error

Taking those two drives and updating them to “resume”, we are able to utilize them after the “resume” is processed.

```
pc-tor01-glvm-01a: / >
# chdev -a resume=yes -l hdisk9
hdisk9 changed
pc-tor01-glvm-01a: / >
# chdev -a resume=yes -l hdisk10
hdisk10 changed
pc-tor01-glvm-01a: / >
# extendvg -f -p appvg_lon06 appvg hdisk9 hdisk10
pc-tor01-glvm-01a: / >
#
```

Figure 3.2.2 Resuming RPV clients with chdev command

### 3.3 When configuring RPV server, available disks are not displayed

An RPV server can only be configured when a disk has a PVID configured. If the PVID is not configured, the RPV configuration panel will not display the drive.

Say we have the following drives as shown on the following diagram.

```

pc-lon06-glvms-01a:/ >
# lspv
hdisk0          none                               None
hdisk1          00c8cf803fb0193a                       None
hdisk2          00c8cf803fb02116                       None
hdisk3          00f6db0af58e9775                       rootvg          active
hdisk4          none                               None
hdisk5          00c8cf803fb029cd                       None
hdisk6          00c8cf803fb03444                       None
pc-lon06-glvms-01a:/ >

```

Figure 3.3.1 Displaying available drives

We see that hdisk4 is available and want to setup an RPV server with that disk. We go into the SMIT panel to add that RPV server (smitty rpvserver->Add Remote Physical Volume Servers and receive the following message "There are no items of this type".

```

Remote Physical Volume Servers
Move cursor to desired item and press Enter.

Remote Physical Volume Server Site Name Configuration
List All Remote Physical Volume Servers
Add Remote Physical Volume Servers
Change / Show a Remote Physical Volume Server
Change Multiple Remote Physical Volume Servers
Remove Remote Physical Volume Servers
Configure Defined Remote Physical Volume Servers

ERROR MESSAGE

Press Enter or Cancel to return to the
application.

1800-051 There are no items of this type.

F1=Help          F2=Refresh       F3=Cancel
F8=Image         F10=Exit        Enter=Do
F9

```

Figure 3.3.2 SMIT does not find any available drives for an RPV server

We exit the SMIT panel and add a PVID to hdisk4 using the "chdev -a pv=yes -l hdisk4" command and note that the PVID has been assigned to that drive.

```

pc-lon06-glvm-01a:/ >
# chdev -a pv=yes -l hdisk4
hdisk4 changed
pc-lon06-glvm-01a:/ >
# lspv
hdisk0          none                               None
hdisk1          00c8cf803fb0193a                       None
hdisk2          00c8cf803fb02116                       None
hdisk3          00f6db0af58e9775                       rootvg          active
hdisk4          00c8cf8049bac9fa                       None
hdisk5          00c8cf803fb029cd                       None
hdisk6          00c8cf803fb03444                       None
pc-lon06-glvm-01a:/ >

```

Figure 3.3.3 Displaying new PVID on hdisk4

When we go back to the RPV server panel and try to select a drive to assign to an RPV server, we now get the following.

```

Remote Physical Volume Servers
Move cursor to desired item and press Enter.

Remote Physical Volume Server Site Name Configuration
List All Remote Physical Volume Servers
Add Remote Physical Volume Servers
Change / Show a Remote Physical Volume Server
Change Multiple Remote Physical Volume Servers

Physical Volume Identifiers

Move cursor to desired item and press F7.
ONE OR MORE items can be selected.
Press Enter AFTER making all selections.

# Physical Volume      Physical Volume Identifier
# -----
hdisk4                 00c8cf8049bac9fa

F1=Help                F2=Refresh             F3=Cancel
F7=Select              F8=Image               F10=Exit
Enter=Do               /=Find                 n=Find Next
F1
F9

```

Figure 3.3.4 Displaying available disks for an RPV server

We note that hdisk4 is now available as it now has a PVID configured. We can proceed to configuring the RPV server.



# Chapter 4: Additional Resources

## 4.1 AIX Geographic Logical Volume Manager (GLVM) resources

AIX Geographic Logical Volume Manager (GLVM) is a software-based technology for real time geographic data mirroring over standard TCP/IP networks. Listed below are some additional resources.

- [https://www.ibm.com/support/knowledgecenter/en/SSPHQG\\_7.2/glvm/ha\\_glvm\\_glvm.html](https://www.ibm.com/support/knowledgecenter/en/SSPHQG_7.2/glvm/ha_glvm_glvm.html)
- [https://www.ibm.com/support/knowledgecenter/SSPHQG\\_7.2/glvm/hacmpgeolvm\\_pdf.pdf](https://www.ibm.com/support/knowledgecenter/SSPHQG_7.2/glvm/hacmpgeolvm_pdf.pdf)

## 4.2 IBM PowerHA resources

AIX Geographic Logical Volume Manager (GLVM) is a software-based technology for real time geographic data mirroring over standard

- [https://www.ibm.com/support/knowledgecenter/SSPHQG\\_7.2/glvm/ha\\_glvm\\_plan\\_glvm\\_mirror.htm](https://www.ibm.com/support/knowledgecenter/SSPHQG_7.2/glvm/ha_glvm_plan_glvm_mirror.htm)
- <http://www.redbooks.ibm.com/redbooks/pdfs/sg248434.pdf>

## 4.3 Other resources

- *IBM Cloud Power Virtual Server offering:*
  - <https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-getting-started>