

IBM Power Virtual Server: Global Replication Services

September 2024



AIX
IBM i
Linux

IBM Power Virtual Server

The world's best landing zone for IBM Power workloads

Table of Contents

1	<i>Solution Overview</i>	3
2	<i>Setup Global Replication Services</i>	4
4	<i>Disaster Recovery workflow</i>	6
4.1	Enable volume replication	6
4.1.1	Create a new replication-enabled Volume.....	7
4.1.2	Convert an existing volumes to be a replication-enabled volume.....	7
4.1.3	Check the replication properties of the volume.....	7
4.2	Create a volume Group	8
4.2.1	Create a volume group.....	9
4.2.2	Get volume group properties	9
4.2.3	Get volume-group storage details.....	10
4.2.4	Get volume-group relationship details.....	11
4.3	Switch to secondary site	11
4.4	Onboard auxiliary volumes	11
4.4.1	Onboard auxiliary volume CLI.....	12
4.4.2	Check onboard progress.....	13
4.5	Deploy a standby VM on secondary site	14
5	<i>Failover/Failback</i>	15
5.1	Access auxiliary volumes on site failure	15
5.2	Failover or switch volume group role to auxiliary	16
5.3	Failback to primary site	17
6	<i>Disabling replication</i>	17
6.1.1	Remove the volumes from volume-group from Primary Site.....	17
6.1.2	Disable the replication of a volume.....	17
6.1.3	Remove the volumes from volume-group from Secondary Site	18
6.1.4	Delete the auxiliary volume from secondary site.....	18
7	<i>Billing</i>	18
8	<i>Troubleshooting</i>	19
8.1	Can I start and stop a volume-group from any site?	19
8.2	Volume-group replication status is not in sync across two sites.	19
8.3	How do I check the failures if an update on a volume-group fails?	19
8.4	Update on volume-group is not working as its status is in error state.	19
8.6	What If I forget to onboard few volumes of one volume-group?	19
8.8	What if I delete the volume from one site, but not from the other site?	19
9	<i>References</i>	20

1 Solution Overview

IBM Power clients run mission critical workloads to ensure business continuity at all times. To maintain this, a platform to provide a secure and highly available landscape is required. Planning such an environment is complex and requires lot of capital expenditures to configure compute, capacity and advanced network and storage requirements.

IBM Power Virtual Server provides Global Replication Services, which provides SAN-based volume replication services upon which additional disaster recovery solutions can be built.



Global Replication services (GRS) is based on well-known industry standards IBM FlashSystem Global Mirror Change Volume asynchronous replication technology. GRS exposes a set of interfaces to create and manage replication-enabled volumes.

The benefits of Global Replication on Power Virtual Server include the following:

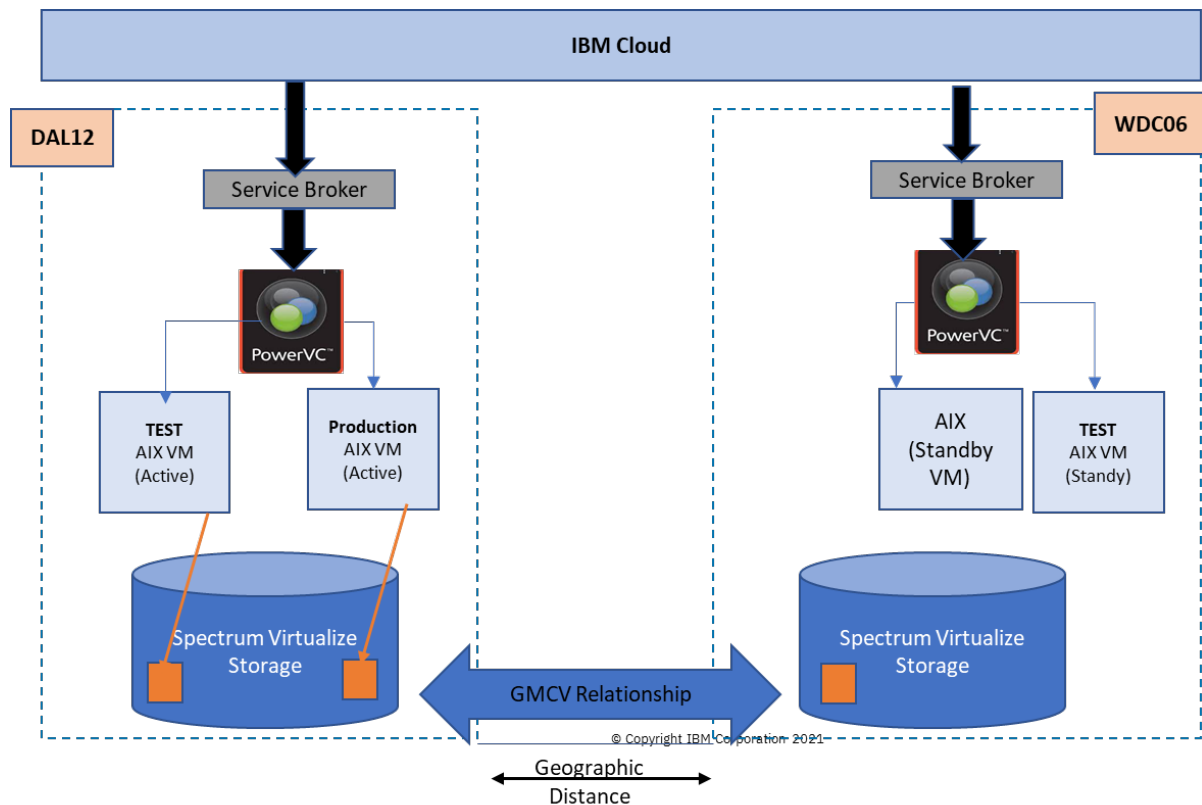
- ❖ Maintain a consistent and recoverable copy of the data at the remote site, created with minimal impact to applications at your local site.
- ❖ Efficiently synchronize the local and remote sites with support for failover and failback modes, helping to reduce the time that is required to switch back to the local site after a planned or unplanned outage.
- ❖ Replicate more data in less time to remote locations.
- ❖ Maintain redundant data centers in distant geographies for rapid recovery from disasters.
- ❖ Eliminate costly dedicated networks for replication and avoid bandwidth upgrades.



This tutorial focuses on how to use the new GRS API/CLI to build the disaster recovery solution.

- Setting up replication from scratch
- Setting up replication using existing volumes

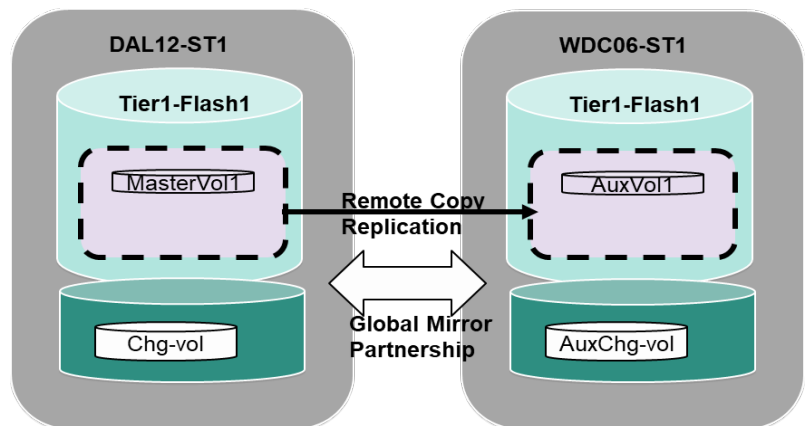
2 Setup Global Replication Services



The Power Virtual Server data centers are setup to have all the required configuration needed to offer replication capabilities. Storage controllers are pre-configured to use GMCV replications. GRS provides replication at the storage level by making use of IBM FlashSystem Global Mirror Change volume (GMCV) asynchronous replication technology. In this case, the initial sync copies the entire data set from the primary to auxiliary; subsequently, only the delta changes are synchronized with the periodic interval of 500 seconds, which means the average RPO is ~10 minutes and the maximum RPO is ~15 minutes.

On every create request for replicated volumes, four copies of volumes are created across two sites.

- a) Primary volume on site 1
- b) Primary change volume on site 1 to store the delta changes
- c) Auxiliary volume on site 2
- d) Auxiliary change volume on site 2 to update the delta changes



It uses a remote copy consistency group to ensure that the data spread across multiple volumes are consistent while it is copied across remote site. Additionally, it helps to switch the replication direction during the time of planned and un-planned disaster events. GRS API and CLI interfaces can be used to create and manage replicated volumes and consistency groups. From a logical standpoint, the site from which a volumes is created or enabled for replication is referred to as the primary site and the remote site is referred to as the auxiliary site.

Once we have the volumes replicated at both primary and secondary sites, we can use certain steps as mentioned in *Section 4 : Disaster Recovery workflow*, to bring up the standby VM using the replicated volumes.

3 Disaster Recovery location sites

Identify the replication enabled Sites, using the below PowerVS disaster recovery location CLI. The below example output shows that dal10 and us-east are active replication sites.

CLI: ibmcloud pi drl --all-regions -json

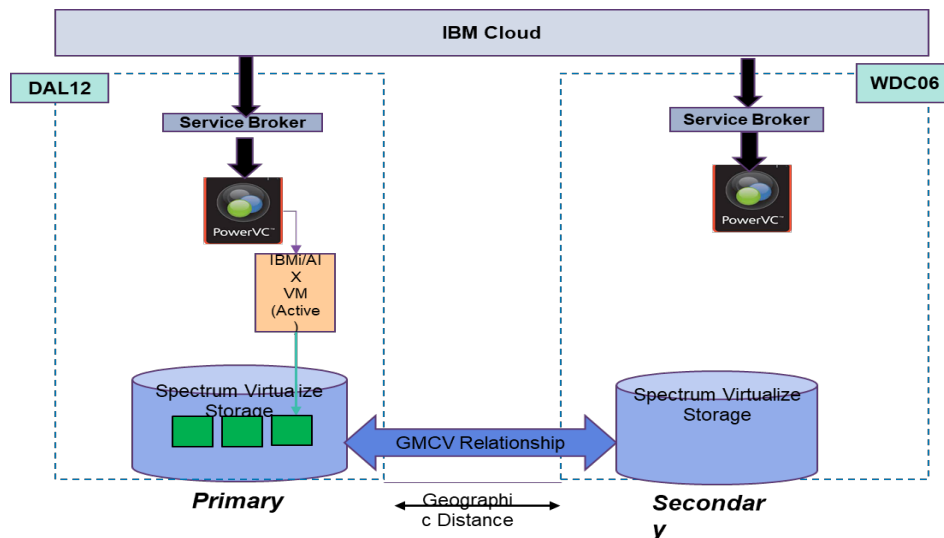
```
{
  "disasterRecoveryLocations": [
    {
      "location": "dal10",
      "replicationSites": [
        {
          "isActive": true,
          "location": "us-east"
        },
        {
          "isActive": true,
          "location": "dal12"
        }
      ]
    },
    {
      "location": "dal12",
      "replicationSites": [
        {
          "isActive": true,
          "location": "dal10"
        }
      ]
    },
    {
      "location": "us-east",
      "replicationSites": [
        {
          "isActive": true,
          "location": "dal10"
        }
      ]
    }
  ]
}
```

Create Power Virtual Service Instances on both replications enabled sites. Once you create Power Virtual Server service Instances, you can list these instances to find the CRNs.

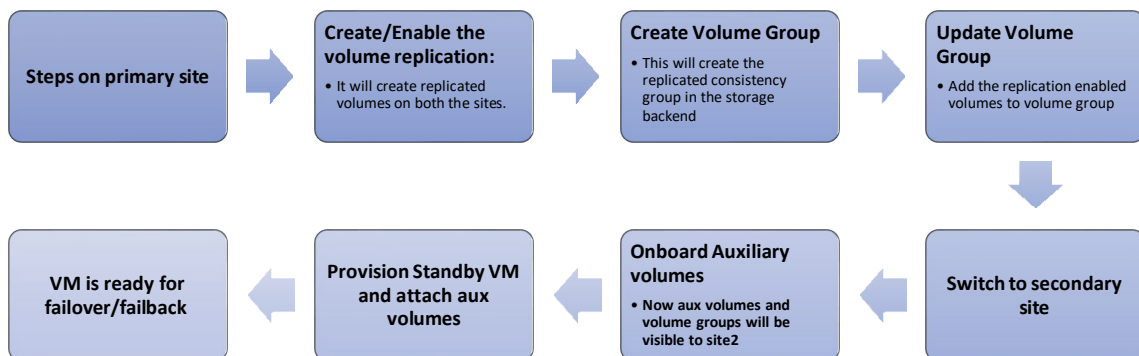
CLI: ibmcloud pi ins list

4 Disaster Recovery workflow

As an example, consider we have an AIX VM running an Oracle database application workload in the DAL12 data center serving as the primary site and we need to enable the GRS for the data volumes to recover the Oracle database.



Below are the steps to enable the replication for your application workload running on primary site and make it ready to trigger failover/failback.

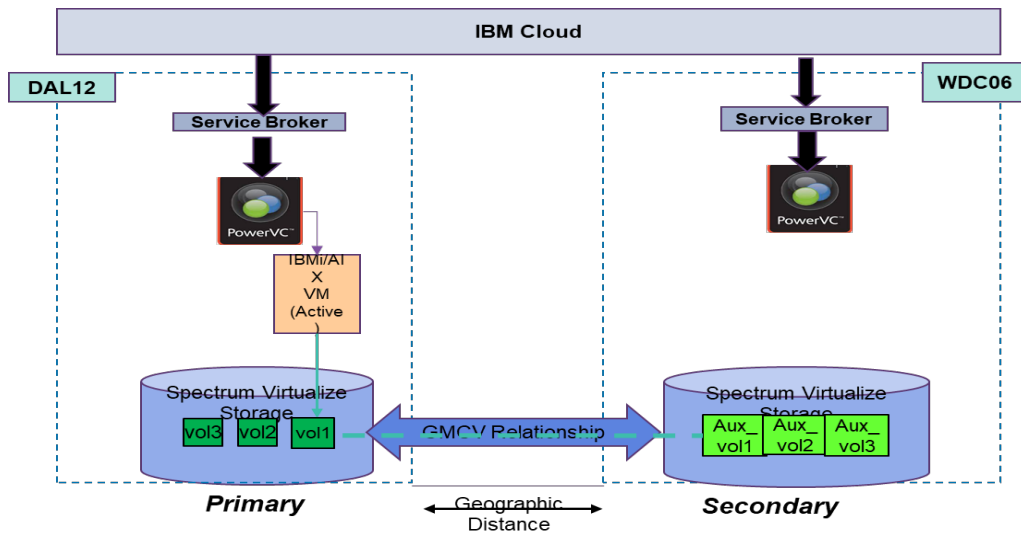


Below sections will show the details for each step and how you can use the IBM PowerVS cli to perform them.

4.1 Enable volume replication

The first step is to have replication enabled volumes for the VM which needs to be protected during disaster recovery. You can either create a new replication enabled volumes, or you can convert existing volumes to replication enabled. When the volume is replication enabled it will create an auxiliary mirror volume on remote storage controllers and creation replication relationships.

Below figure shows AIX VM having volumes (vol1, vol2, and vol3) after enabling replication it creates aux_vol1, aux_vol2 and aux_vol3 on secondary storage. These volumes are not visible and managed by service broker workspace.



4.1.1 Create a new replication-enabled Volume

A new option “replication Enabled” has been added for creating replication enabled volume.

```
CLI: ibmcloud pi vol create test_vol --size 1 --replication-enabled -
-storage-tier tier3
```

4.1.2 Convert an existing volumes to be a replication-enabled volume

You can convert existing volumes to replication enabled also, provided the given volume pool supports replication capability.

```
CLI: ibmcloud pi vol action VOLUME_ID --replication-enabled=True
```

4.1.3 Check the replication properties of the volume

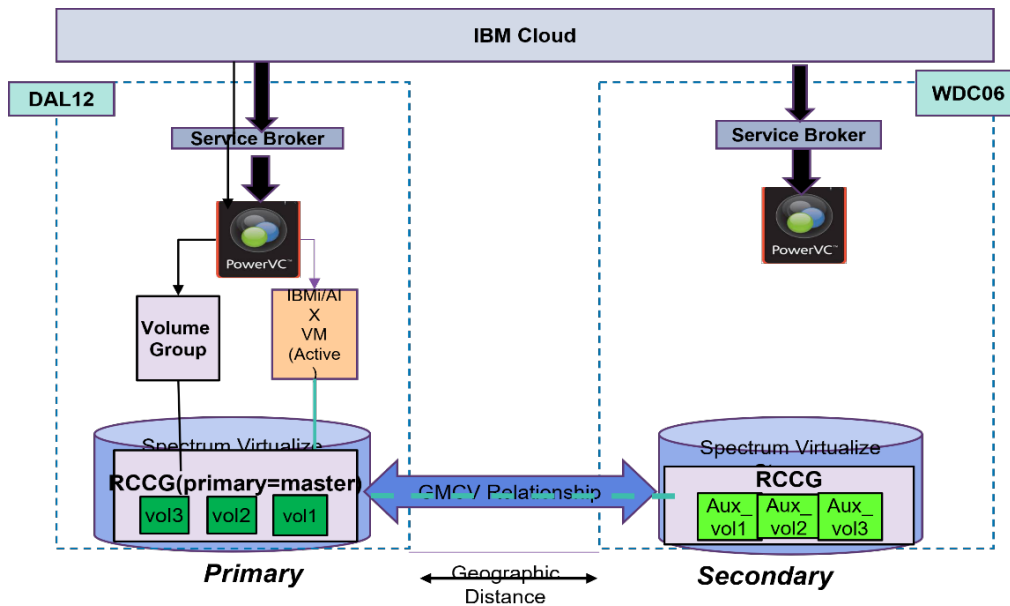
Get the volume details and check the “replicationEnabled” field. If it is true, then volume is replication enabled otherwise it is not.

```
CLI: ibmcloud pi vol get VOLUME_ID --json
```

- *auxVolumeName, masterVolumeName: - Names of auxiliary and master volume created at the storage host level.*
- *auxiliary: Boolean property which tells if the volume referred is master or auxiliary. If false means that this is master volume, else it is the auxiliary volume.*
- *primaryRole: - Shows the role of the volume, is primaryRole is master means the volume at the given site is playing the role of master.*
- *replicationStatus:- shows the replication status of the volume. If enabled means volume is enabled for replication and is active.*
- *mirroringState: It's the replication relationship state. If consistent_copying means the replication data is in sync.*

4.2 Create a volume Group

The next step is to create volume group and add the replication enabled volumes to the volume group. This will create remote replication consistency group at both primary and remote storage backend. It stores the consistent copy for the volumes. When the volume group is created it assigns the primary role as master. The figure below shows VolumeGroup which creates RCCG and adds vol1, vol2 and vol3.



4.2.1 Create a volume group

The CLI example below can be used to create a volume-group and add replication enabled volumes. You can only add replication enabled volumes to a volume group. If a non-replication volume is added to a volume group, it will fail.

```
CLI: ibmcloud pi vg create --volume-group-name VOLUME_GROUP_NAME --
member-volume-ids "VOLUME_ID_1,[VOLUME_ID_N]" --json
```

- “member-volume-ids” is a comma separated list of volume ids which should be part of the volume-group.
- All volumes ids must belong to same storage host(volumepool), otherwise VG creation would fail.
- At least one volume is mandatory to create a volume-group.
- Volume can only be part of a single volume-group at a time. If the same volume is added to another volume group, it will fail.

4.2.2 Get volume group properties

We can check the volume group basic properties using the below CLI. These values are fetched from the database so this CLI should be used to fetch the static properties like name and volumeIDs.

```
CLI: ibmcloud pi vg get VOLUME_GROUP_ID --long --json
```

- **consistencyGroupName:** - Shows the name of the replication consistency group which is created at the storage level. This field would remain same for a replication consistency group across two sites.
- **volumeIDs:** - Shows the list of volumeIDs which are part of volume-group.
- **statusDescription:** - This field is populated if there are any failure while adding the volumes to volume-group.
- **status:** - Shows that volume-group status. Status “available” means its active. Possible values are available, error, updating, creating.
- **replicationStatus:** Show that this is replication consistency group, and the replication is enabled.
- When volume is part of volume group. You can see two new fields “group_id” and “consistencyGroupName” in the volume group details.

4.2.3 Get volume-group storage details

Get volume group storage details CLI fetch the live consistency group information from the storage backend. So, if you need exact consistency group state and role, then you should use this CLI.

```
CLI: ibmcloud pi vg sd VOLUME_GROUP_ID --json
```

- **state:** Shows the live consistency group state. Possible consistency group states are: - *consistent_copying, inconsistent_copying, inconsistent_stopped, idling, idling_disconnected, inconsistent_disconnected*

4.2.4 Get volume-group relationship details

To get more granular details about the volume group with respect to remoteRelationships, freezeTime and copy progress, you can use the CLI example below.

```
CLI : ibmcloud pi vg rcr VOLUME_GROUP_ID --json
```

- **remoteCopyrelationships:** Shows the relationship details for each replicated volume which are part of volume group.
- **freezeTime:** Indicates the time *YY/MM/DD/HH/MM* format when the last sync happened. This parameter is used to monitor the RPO.
- **progress:** Shows the relationship progress.

Once we have the volume group created and is in consistent copying state, which means our volume data is copied to the secondary site and are ready for replication.

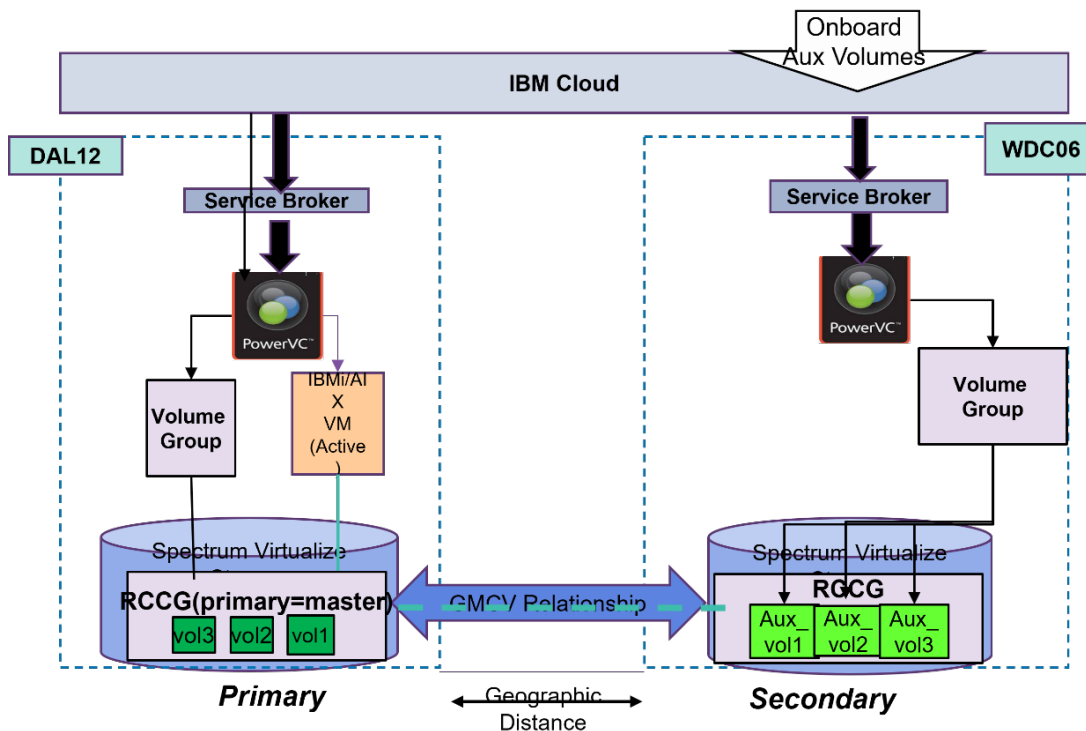
4.3 Switch to secondary site

Now we must move to the secondary site, to perform the remaining steps. In this example, set the service target of the CLI environment to “wdc07” PowerVS instances.

```
CLI: ibmcloud pi ws tg crn:v1:bluemix:public:power-iaas:wdc07:a/  
2bc3df23c0d14ebe921397bd8aa2573a:dfb465a4-c355-4c1a-9f17-61f844f53eaf::
```

4.4 Onboard auxiliary volumes

Though auxiliary volumes are there on the storage host of secondary site, still these volumes are not managed by current power systems virtual server workspace. We need to onboard auxiliary volume so it can be managed and accessed by the cloud user. Onboard operation requires a source CRN to validate if remote user has required permissions to access the auxiliary volumes based on the owner of paired primary volumes and can onboard auxiliary volumes. If the user does not have valid permissions the onboard operation will fail or authentication error.



As a part of the onboard auxiliary volume operation, the required volume IDs are created to manage the existing auxiliary volumes. If the auxiliary volumes are part of the consistency group, onboard operation also create the volume-group IDs to manage the existing consistency group.

4.4.1 Onboard auxiliary volume CLI

We can use the below CLI to trigger the onboard operation. It returns the onboard UUID, which can be used to monitor the progress of the onboard task progress running in the background.

Below example creates an onboarding job “testOnboarding” to onboard two auxiliary volumes for the given source-crn

```
CLI : ibmcloud pi vol on create --source-crn
crn:v1:bluemix:public:power-
iaas:dal10:a/2bc3df23c0d14ebe921397bd8aa2573a:44816a3e-ea7d-
4416-9c32-6f6ce1735c4f:: --description testOnboarding --
auxiliary-volumes "aux_volume-testing-onboarding-cba78220-
1a4a1210664"
```

- *-source-crn* : CRN of Power System Virtual Server Service instance in primary site(dal10). This is used to validate the authorization.
- *--auxiliary-volumes*: Comma separated list of auxiliary volume. Exact Auxiliary volume name as returned by auxVolumeName in volume details from primary site.
- Display name (user defined name) is optional field. If provided, corresponding auxiliary volume will be onboarded with the given display name This name can be provided as a part of “auxiliary-volume” option: - auxiliary volume name and display name separated by space.
- *--description*: Description field is optional and can be used to provide a description to the onboarding operation. This field is not unique and cannot be used to fetch onboarding details.

4.4.2 Check onboard progress

Onboard operation returns onboarding UUID which can be used to check the status of onboarding operation. The onboarding operation is an asynchronous operation that can take some time that depends upon the number of volumes.

We can use the below CLI to check the status of the onboard operation.

```
CLI: ibmcloud pi vol on get ONBOARD_ID --json
```

- **Status:** Show the status for onboard. SUCCESS/FAILURE
- **Results:** Returns the onboarded auxiliary volumes.

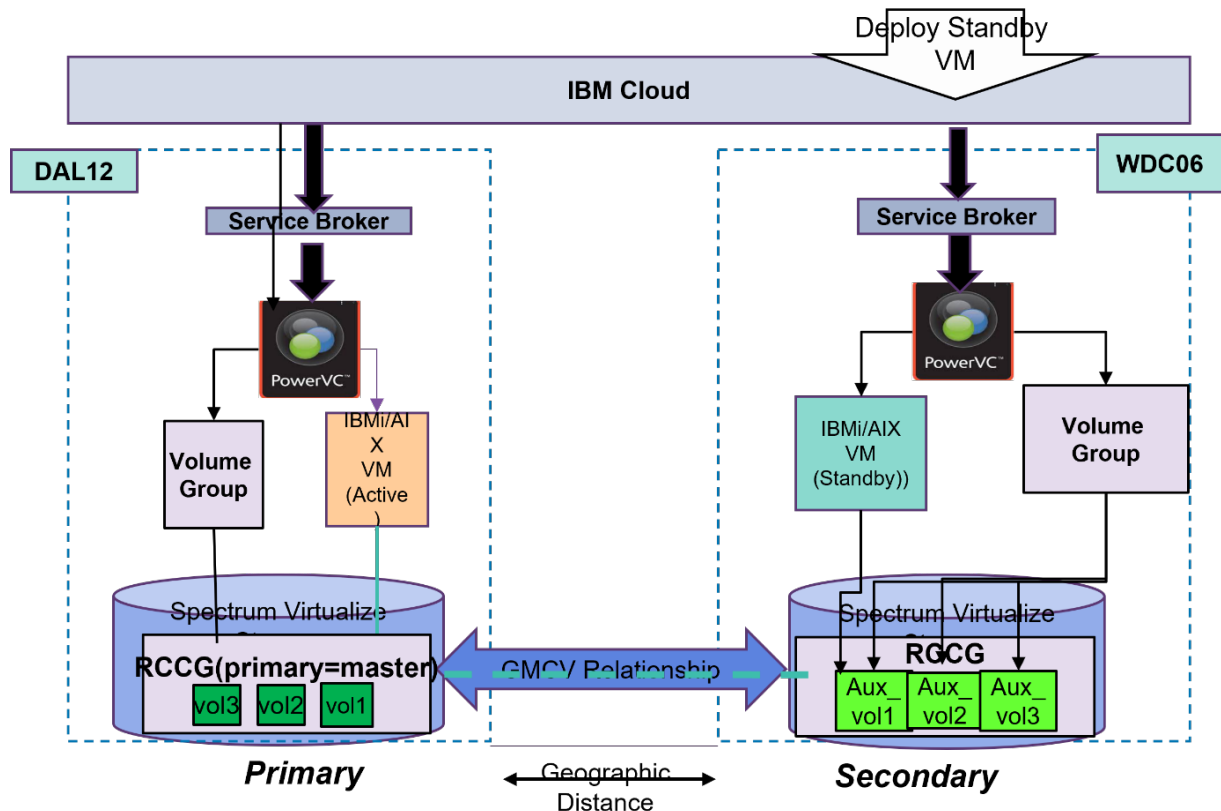
On completion of the onboard operation, you can check the auxiliary volumes using the volume list or can fetch the volume details using the volume name. Volume IDs and group ID for master-aux volume pair are different on primary and secondary sites. However, you can check other fields such as **masterVolumeName**, **auxVolumeName**, and **consistencyGroupName**

To check volume and volume-group details created after the onboard operation use these CLI commands.

```
CLI: ibmcloud pi vol get test_vol -json  
CLI: ibmcloud pi vg get VOLUME_GROUP_ID --long --json
```

4.5 Deploy a standby VM on secondary site

As we are doing the storage-based replication. Now once we have onboarded auxiliary volumes and volume group on secondary site. We need to provision a stand-by VM on the secondary site and attach these auxiliary volumes. You should keep this VM in shutoff state and should only be used in case disaster. Auxiliary volumes are read/write protected. At a time only a primary site can perform I/O on it. When the primary site is down, and consistency group is stopped by allowing read permission, then only read/write operations are allowed on the auxiliary volumes.

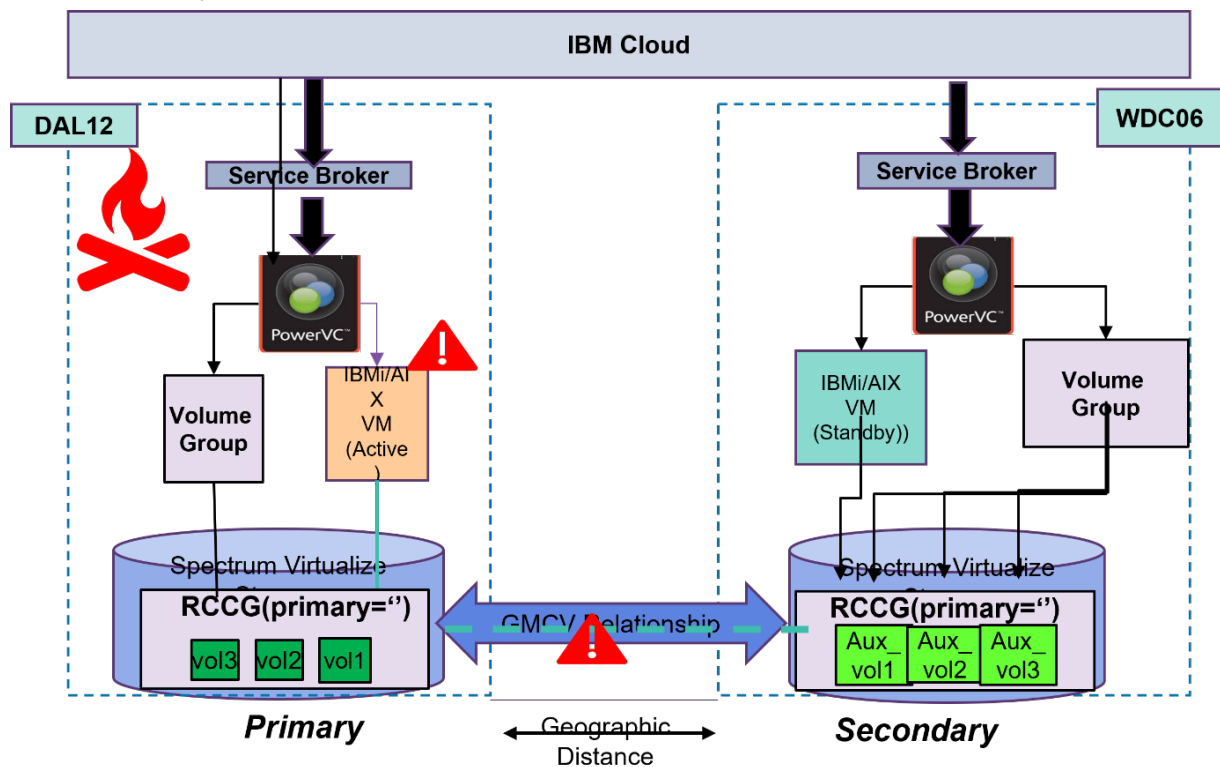


You can use the existing CLI to provision and attach volumes:

<https://cloud.ibm.com/docs/power-iaas-cli-plugin?topic=power-iaas-cli-plugin-power-iaas-cli-reference#ibmcloud-pi-instance-create>

<https://cloud.ibm.com/docs/power-iaas-cli-plugin?topic=power-iaas-cli-plugin-power-iaas-cli-reference#ibmcloud-pi-volume-attach>

5 Failover/Failback



In case of disaster (primary site failure), you will lose access to the storage volumes and they will be marked as ERROR. The replication relationship will be disconnected, and consistency group will move to “consistent-disconnected”. The volume group primary role will be assigned as blank.

In this situation no new replication operations are allowed, as replication is broken. You can only access existing workloads by powering on the standby VM and auxiliary replication volumes from the secondary site after giving them read access.

- ❖ Access auxiliary volumes on primary site failure
- ❖ Failover or switch volume group role to secondary
- ❖ Failback to primary site

5.1 Access auxiliary volumes on site failure

To allow read/write I/O on auxiliary volumes in case of primary site failure, you should stop the volume group with `--allow-read-access`.

```
CLI: ibmcloud pi volume-group action VOLUME_GROUP_ID --
operation stop --allow-read-access=True
```

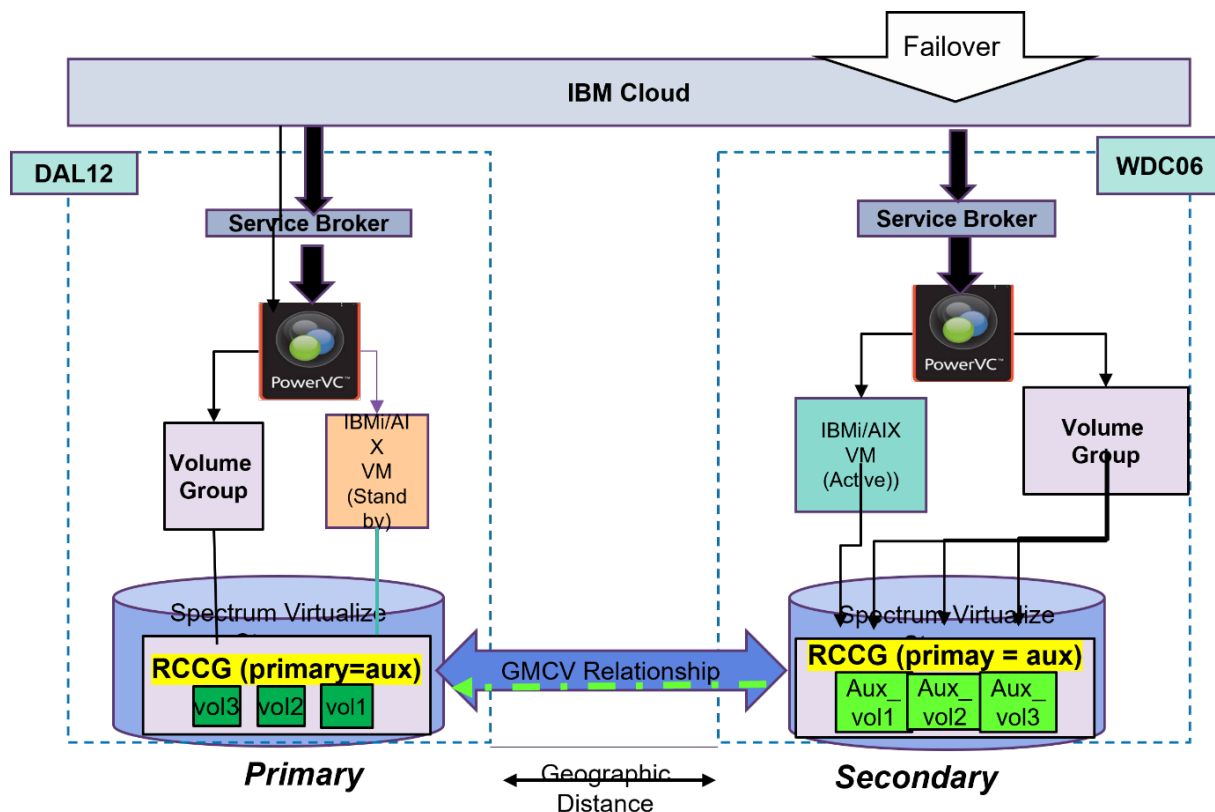
On stopping the volume group, RCG will move into “Idling” state and replication status will be disabled. You can check this by getting the volume storage details

```
CLI: ibmcloud pi vg sd VOLUME_GROUP_ID --json
```

Now you can power-on standby VM and run required instructions to access your database configured on the auxiliary volumes.

5.2 Failover or switch volume group role to auxiliary

When the primary site is recovered, the consistency group can be started to restart the replication. You can start the volume group to switch the role to auxiliary, which changes the replication role from secondary to primary. This will allow the auxiliary volume delta changes to be synchronized with the primary volumes.



Below is the CLI which can be used to start the volume group and switch the role to auxiliary.

```
CLI: ibmcloud pi volume-group action VOLUME_GROUP_ID --operation start --source auxiliary
```


5.3 Failback to primary site

To switch back the volume group to the primary site, you need to perform volume group stop and start operations from the primary site. You can use the same volume group stop command and wait till ReplicationStatus becomes disabled. Then use volume group start command with -source master option.

```
ibmcloud pi volume-group action VOLUME_GROUP_ID --operation
stop --allow-read-access=True
ibmcloud pi volume-group action VOLUME_GROUP_ID --operation
start --source master
```

6 Disabling replication

Disabling the replication means deleting the auxiliary volume from the remote site. Before disabling the replication, make sure that it is not associated with any group. Since there are two sites, we should follow the below procedure for disabling the replication.

6.1.1 Remove the volumes from volume-group from Primary Site

If volume is a part of any volume-group then remove the volume from its associated volume-group.

```
CLI: ibmcloud pi volume-group update VOLUME_GROUP_ID --remove-
member-volume-ids "VOLUME_ID_1,[VOLUME_ID_N]"
```

6.1.2 Disable the replication of a volume

Disable the volume replication. This will remove the replication relationship and delete the auxiliary volume from the storage backend.

```
CLI: ibmcloud pi volume action VOLUME_ID --replication-enabled=False
```

It's an asynchronous process, so check volume-details to make sure that volume replication is disabled.

```
CLI: ibmcloud pi vol get VOLUME_ID --json
```

6.1.3 Remove the volumes from volume-group from Secondary Site
Update the volume-group of secondary site to remove the volumes from it.

```
CLI: ibmcloud pi volume-group update VOLUME_GROUP_ID --  
remove-member-volume-ids "VOLUME_ID_1,[VOLUME_ID_N]"
```

If the volume-group is empty, then you can delete the volume-group.

```
CLI: ibmcloud pi vg delete VOLUME_GROUP_ID
```

6.1.4 Delete the auxiliary volume from secondary site.

Lastly, delete the auxiliary volume reference from the secondary site. If the auxiliary volume is not deleted from the auxiliary site, then this volume will move to ERROR state as a part of the out-of-band sync interval (24 hours), as these volumes no longer exists in the storage backend.

```
CLI: ibmcloud pi vol delete VOLUME_ID
```

7 Billing

Please refer to the IBM Power Virtual Server [documentation](#) to learn more about GRS billing.

8 Troubleshooting

8.1 Can I start and stop a volume-group from any site?

Yes, you can start and stop it from any site. But it is recommended to use primary site for all the volume operations and perform operations on auxiliary volume on secondary site only during failover.

8.2 Volume-group replication status is not in sync across two sites.

Check storage-details of volume-group to check the actual replication status of volume-group, not the volume-group details.

The start and stop operation on Volume-group will update the replication status of volume-group on the site from where start and stop operation is performed but it does not update the replication status of corresponding volume-group on the other site.

8.3 How do I check the failures if an update on a volume-group fails?

Update on volume-group is an asynchronous operation. You should check the volume-group details to check its results. If there is any error during the update operation then “statusDescription(errors)” field will provide the error details from the last failed operation.

8.4 Update on volume-group is not working as its status is in error state.

You can perform a reset operation on a volume-group. This action does not make any changes to the replication status, but it will set its status to “available” so that updates can be performed on the volume-group.

This action will not clear the “errors” field from the volume-group. The next successful update operation would reset this field.

8.5 How do I check who is playing the primary role for a volume-group?

Fetch storage-details of a volume-group and check “primaryRole” field. “master” value indicates that primary volumes are playing primary role.

8.6 What If I forget to onboard few volumes of one volume-group?

You can create one more onboarding operation with required volume list. Onboarding operation onboards new volumes with existing volume-group.

8.7 Can I add more volumes to a volume-group after onboarding operation?

Yes, you can. Add volumes to the volume-group on primary site and then onboard the volumes on secondary site.

8.8 What if I delete the volume from one site, but not from the other site?

The replicated volume that is managed on its corresponding remote site will transition to the error state within 24 hours. Any operation on this replicated volume will fail (except the delete operation) and will set the volume in error state.

You must delete the volumes from the primary site. Otherwise, the primary volume will continue being charged.

9 References

- IBM Power System Virtual Servers API Reference: <https://cloud.ibm.com/apidocs/power-cloud>
- IBM Power System Virtual Servers CLI Reference: <https://cloud.ibm.com/docs/power-iaas-cli-plugin?topic=power-iaas-cli-plugin-power-iaas-cli-reference>
- IBM Power System Virtual Servers Terraform: <https://registry.terraform.io/providers/IBM-Cloud/ibm/latest/docs>