

Global Replication Services Solution using IBM Power Virtual Server

04-Jan-2023



AIX
IBM i
Linux

Power Virtual Server

Serves as the compute space for workloads

Imranuddin W Kazi (imran.kazi@in.ibm.com)

Chhavi Agarwal (chhagarw@in.ibm.com)

Anu Jalan (anujalan@in.ibm.com)

Table of Contents

1	Solution Overview.....	4
2	Setup for Global Replication Services	5
3	DR location sites.....	6
4	Disaster Recovery Workflow.....	7
4.1	Create/Enable the volume replication.....	7
4.1.1	Create New Replication Enabled Volume	8
4.1.2	Convert Existing volumes to Replication Enabled volume.....	8
4.1.3	Check Replication properties of the volume.....	8
4.2	Create/Update Volume Group.....	9
4.2.1	Create Volume Group	10
4.2.2	Get Volume Group properties	10
4.2.3	Get Volume-group storage details.....	11
4.2.4	Get Volume-group relationship details.....	12
4.3	Switch to Secondary site (WDC06)	12
4.4	Onboard Auxiliary volumes.....	12
4.4.1	Onboard Auxiliary volume CLI.....	13
4.4.2	Check Onboard Progress.....	14
4.5	Deploy Standby VM on secondary site	15
5	Failover/Failback	16
5.1	Access auxiliary volumes on site failure.....	16
5.2	Failover or Switch volume group role to Auxiliary.....	17
5.3	Failback to primary site.....	18
6	Disabling the replication	18
6.1.1	Remove the volumes from volume-group from Primary Site.....	18
6.1.2	Disable the replication of a volume	18
6.1.3	Remove the volumes from volume-group from Secondary Site	19
6.1.4	Delete the auxiliary volume from secondary site.	19
7	Billing and Charging	19
8	Troubleshooting.....	20
8.1	Can I start and stop volume-group from any site ?	20
8.2	Volume-group replication status is not in sync across two sites	20
8.3	How to check the failures if update on volume-group fails.....	20
8.4	Update on volume-group is not working as its status is in error state.....	20
8.5	How to check which volume(?) is playing the primary role.....	20
8.6	What If I forget to onboard few volumes of one volume-group?	20

8.7	Can I add more volumes to a volume-group after onboarding operation ?.....	20
8.8	What if I delete the volume from one site, but not from the other site ?	20
9	References	21

1 Solution Overview

Power Systems clients run mission critical workloads, to ensure business continuity during uncertain conditions, we need a secure, highly available and disaster recovery solution. Planning such an environment is complex and requires lot of capital expenditures to configure compute, capacity and advance network and storage requirements.

IBM Power Virtual Server cloud now brings a Global Replication solution which provide the replication capability to your workloads by maintaining the benchmarks for RTO/RPO.




Global Replication services (GRS) is based on well-known industry standards IBM Storwize Global Mirror Change Volume Asynchronous replication technology. IBM PowerVS Global Replication Services solution exposes cloud API/CLI 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 centres in distant geographies for rapid recovery from disasters.
- ❖ Eliminate costly dedicated networks for replication and avoid bandwidth upgrades.

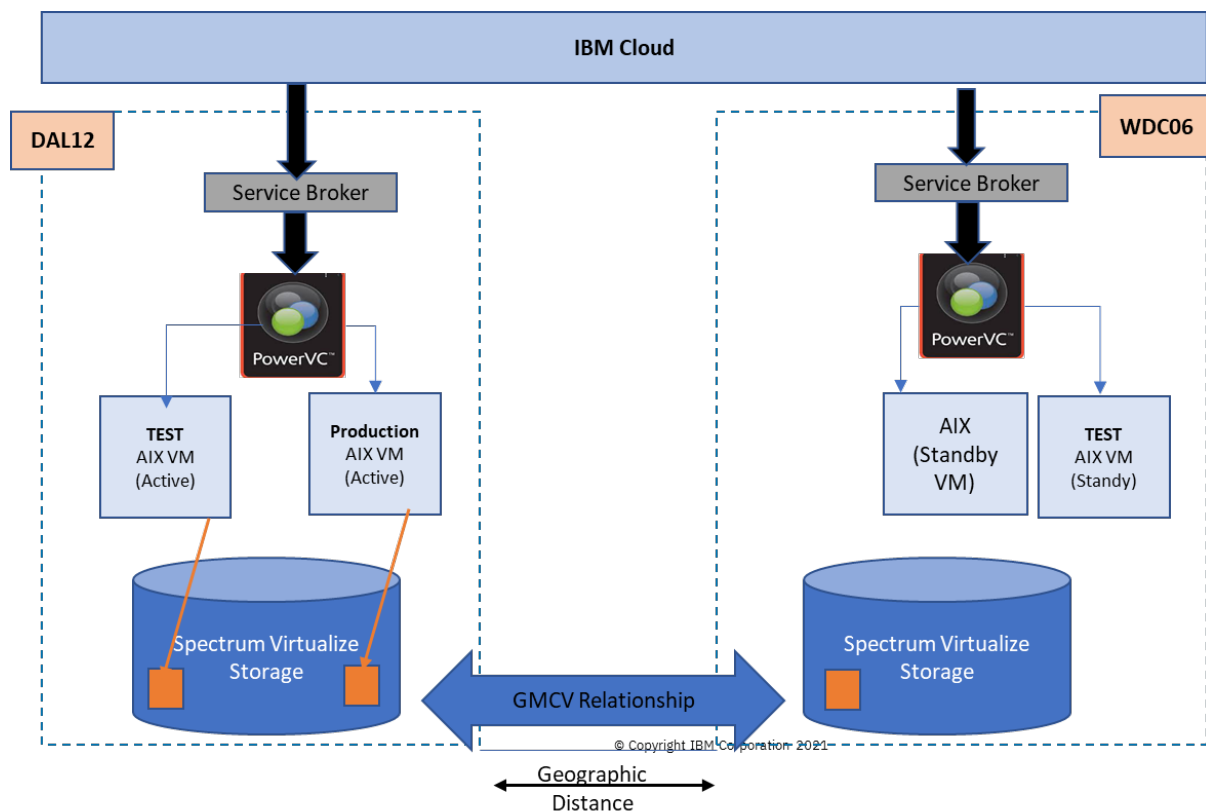
Feature is currently enabled in 2 datacentres: DAL12 and WDC06.



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 for Global Replication Services



The data centres for IBM Power Systems Virtual Server are setup to have all the required configuration needed to offer replication capabilities. Supported storage controllers Tier1/Tier3 will pre-configured to use GMCV replications. Global Replication services provide the replication at the storage level by making use of IBM Storwize Global Mirror Change volume (GMCV) asynchronous replication technology. In this case the first initial sync copies the entire data from master to auxiliary, going forward only the delta changes are synchronized with the periodic interval of 500sec. Which means the maximum RPO will be ~15min.

On every create of replicated volumes, 4 copies of volumes are created across 2 sites.

- a) Master volume on site1
- b) Master change volume on site1, to store the delta changes
- c) Auxiliary volume on site2
- d) Auxiliary change volume on site2, to update the delta changes



It uses remote copy consistency group to ensure that the data spread across multiple volumes are consistent while it is copied across remote site. Also, it helps to switch the replication direction during the time of planned and un-planned disaster. GRS API/CLIs can be used to create and manage replicated volumes and consistency groups.

IBM PowerVS have DAL12/WDC06 datacentres enabled to use Global Replication Services APIs.

Which means if you are using DAL12 as a primary site, you will have auxiliary volumes created on WDC06. And if you are using WDC06 as a primary site, you will have auxiliary volumes created on DAL12. Site where volumes are created or enabled for replication is the primary site.

Once we have the volumes replicated at both primary and secondary steps, 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 dal12 and wdc06 are active replication sites.

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

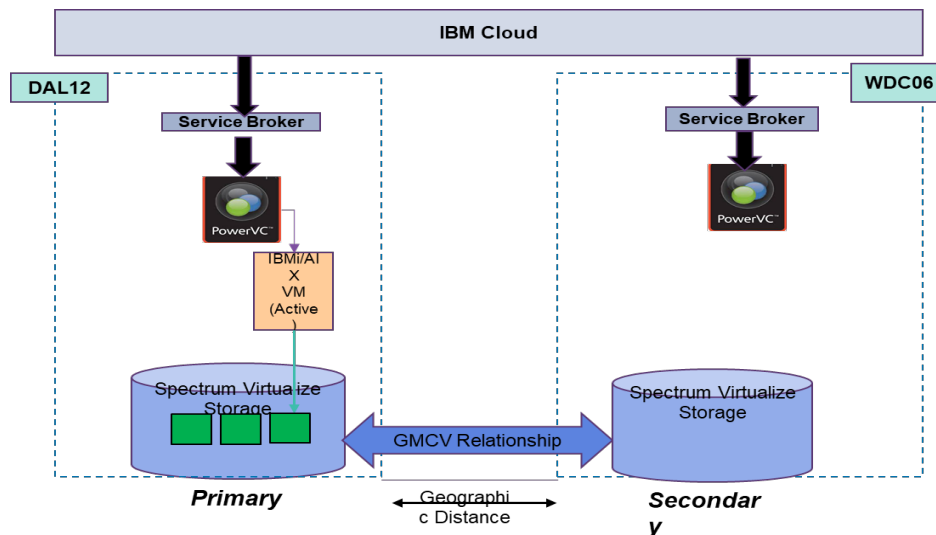
```
21:04 $ ibmcloud pi drl --all-regions --json
{
  "disasterRecoveryLocations": [
    {
      "location": "dal12",
      "replicationSites": [
        {
          "isActive": true,
          "location": "wdc06"
        }
      ]
    },
    {
      "location": "wdc06",
      "replicationSites": [
        {
          "isActive": true,
          "location": "dal12"
        }
      ]
    }
  ]
}
```

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.

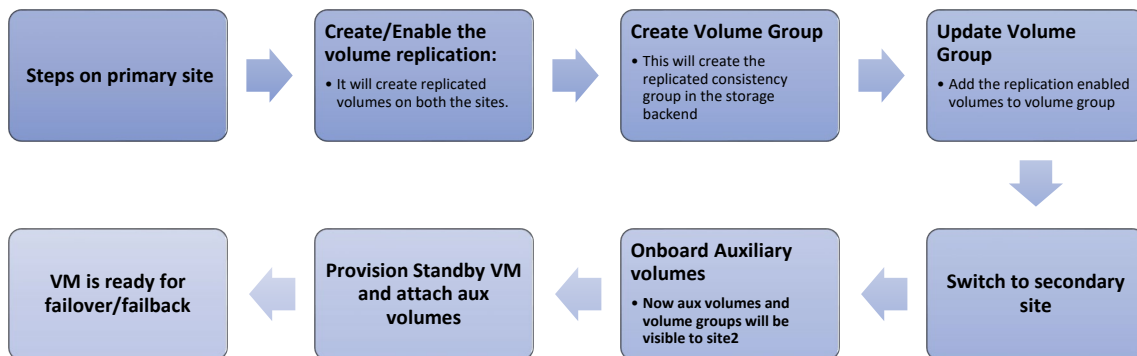
CLI: ibmcloud pi service-list

4 Disaster Recovery Workflow

As an example, consider we have an AIX VM running an oracle DB application workload on DAL12 data centre serving as primary site, and we need to enable the Global replication 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.

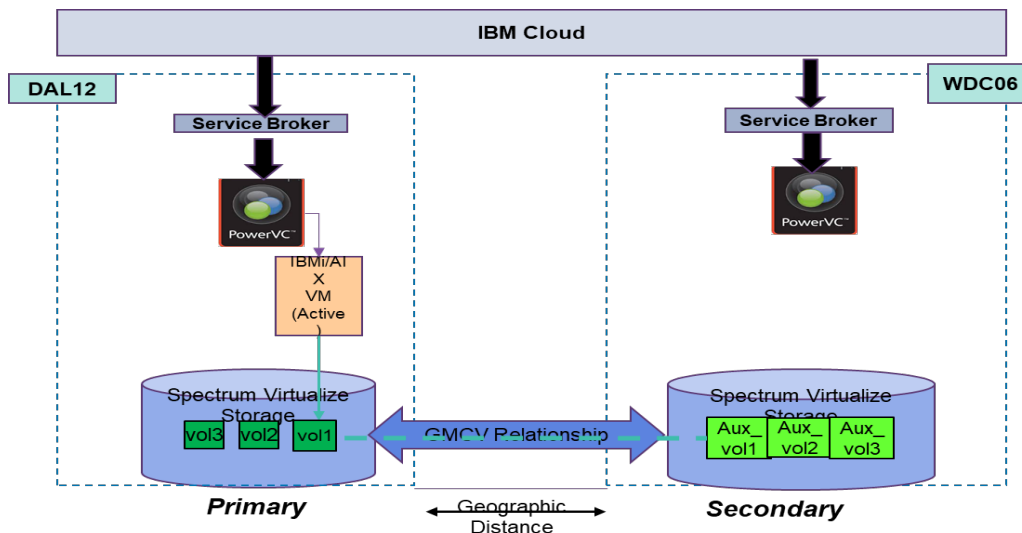


Below sections will show the details for each step and how you can use the IBM PowerVS cli to perform them.

4.1 Create/Enable the volume replication

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 New Replication Enabled Volume

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

CLI: `ibmcloud pi volc testVol1 --size 1 --replication-enabled --type tier1`

```
15:08 $ ibmcloud pi volc testVol1 --size 1 --type tier1 --json --replication-enabled
{
  "bootable": false,
  "creationDate": "2022-11-22T09:40:50.610Z",
  "diskType": "tier1",
  "lastUpdateDate": "0001-01-01T00:00:00.000Z",
  "name": "testVol1",
  "pvmInstanceIDs": [],
  "replicationEnabled": true,
  "shareable": false,
  "size": 1,
  "state": "creating",
  "volumeID": "afd07003-a61a-45ca-97d1-4f910272306d",
  "volumePool": "Tier1-Flash-3",
  "volumeType": "Tier1-Flash-3-DR"
}
```

4.1.2 Convert Existing volumes to Replication Enabled volume

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

CLI: `ibmcloud pi vola testVol2 --replication-enabled=true`

```
15:24 $ ibmcloud pi vola testVol2 --replication-enabled
Performing action on volume testVol2 under account Power IaaS Lite - Staging
OK
Action on Volume ID testVol2 successful.
```

4.1.3 Check 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 testVol1 --json`


```

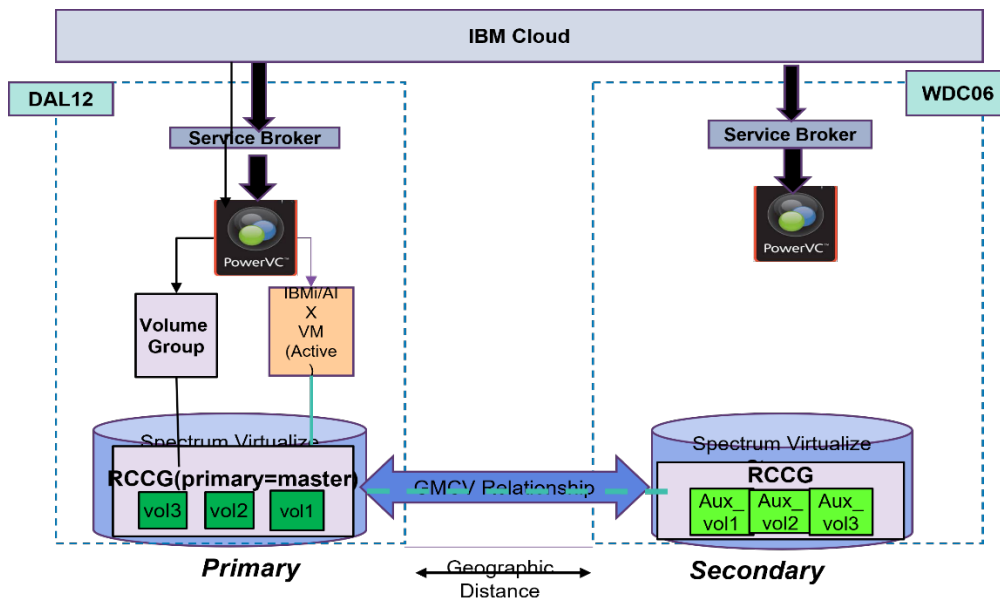
15:28 $ ibmcloud pi vol testVol1 --json
{
  "auxVolumeName": "aux_volume-testVol1-afd07003-a61a1210664",
  "auxiliary": false,
  "bootable": false,
  "creationDate": "2022-11-22T09:40:50.000Z",
  "diskType": "tier1",
  "lastUpdateDate": "2022-11-22T09:40:56.000Z",
  "masterVolumeName": "volume-testVol1-afd07003-a61a",
  "mirroringState": "consistent_copying",
  "name": "testVol1",
  "primaryRole": "master",
  "pvmInstanceIDs": [],
  "replicationEnabled": true,
  "replicationStatus": "enabled",
  "replicationType": "global",
  "shareable": false,
  "size": 1,
  "state": "available",
  "volumeID": "afd07003-a61a-45ca-97d1-4f910272306d",
  "volumePool": "Tier1-Flash-3",
  "volumeType": "Tier1-Flash-3-DR",
  "wwn": "60050768108081F7D000000000484B9"
}

```

- *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/Update Volume Group

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. Below figure shows VolumeGroup which creates RCCG and adds vol1, vol2 and vol3.



4.2.1 Create Volume Group

Below CLI 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.

```
13:53 $ ibmcloud pi vgc --volume-group-name testVolGrp --member-volume-ids
afd07003-a61a-45ca-97d1-4f910272306d,7a9f1ca6-acec-4578-a65c-5e45f893a4a2 --json
{
  "id": "5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca",
  "name": "testVolGrp",
  "status": "available"
}
```

- “member-volume-ids” is a space 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.

```
16:36 $ ibmcloud pi vg 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca --long --json
{
  "consistencyGroupName": "rccg-5bbe-189ca",
  "id": "5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca",
  "name": "testVolGrp",
  "replicationStatus": "enabled",
  "status": "available",
  "statusDescription": {
    "errors": []
  },
  "volumeIDs": [
    "7a9f1ca6-acec-4578-a65c-5e45f893a4a2",
    "afd07003-a61a-45ca-97d1-4f910272306d"
  ]
}
```

- **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.

```
13:47 $ ibmcloud pi vol testVol2 --json
{
  "auxVolumeName": "aux_volume-testVol2-7a9f1ca6-acec1210664",
  "auxiliary": false,
  "bootable": false,
  "consistencyGroupName": "rccg-5bbe-189ca",
  "creationDate": "2022-11-22T09:54:28.000Z",
  "diskType": "tier1",
  "groupID": "5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca",
  "lastUpdateDate": "2022-12-01T07:40:17.000Z",
  "masterVolumeName": "volume-testVol2-7a9f1ca6-acec",
  "mirroringState": "consistent_copying",
  "name": "testVol2",
  "primaryRole": "master",
  "pvmInstanceIDs": [
    "f31bf07d-da37-4848-bd4c-b6e64db036cc"
  ],
  "replicationEnabled": true,
  "replicationStatus": "enabled",
  "replicationType": "global",
  "shareable": false,
  "size": 1,
  "state": "in-use",
  "volumeID": "7a9f1ca6-acec-4578-a65c-5e45f893a4a2",
  "volumePool": "Tier1-Flash-3",
  "volumeType": "Tier1-Flash-3-DR",
  "wwn": "60050768108081F7D000000000484BB"
}
```

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 vgsd 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca --json`

```
14:05 $ ibmcloud pi vgsd 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca --json
{
  "consistencyGroupName": "rccg-5bbe-189ca",
  "cyclePeriodSeconds": 500,
  "cyclingMode": "multi",
  "numOfvols": 2,
  "primaryRole": "master",
  "remoteCopyRelationshipNames": [
    "rcrel4",
    "rcrel6"
  ],
  "replicationType": "global",
  "state": "consistent_copying"
}
```

- **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 know more granular information for the volume group with respect to remoteRelationships, freezeTime and copy progress. You should use the below CLI.

```
CLI : ibmcloud pi vgrcr 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca --json
```

```
00:15 $ ibmcloud pi vgrcr 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca --json
{
  "id": "5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca",
  "remoteCopyRelationships": [
    {
      "auxChangedVolumeName": "chg_aux_volume-testVol2-7a9f1ca6-acec1210664",
      "auxVolumeName": "aux_volume-testVol2-7a9f1ca6-acec1210664",
      "consistencyGroupName": "rccg-5bbe-189ca",
      "copyType": "global",
      "cyclingMode": "multi",
      "freezeTime": "2022-12-02T18:34:18.000Z",
      "masterChangedVolumeName": "chg_volume-testVol2-7a9f1ca6-acec",
      "masterVolumeName": "volume-testVol2-7a9f1ca6-acec",
      "name": "rcrel4",
      "primaryRole": "master",
      "progress": 100,
      "remoteCopyID": "36",
      "state": "consistent_copying"
    },
    {
      "auxChangedVolumeName": "chg_aux_volume-testVol1-afd07003-a61a1210664",
      "auxVolumeName": "aux_volume-testVol1-afd07003-a61a1210664",
      "consistencyGroupName": "rccg-5bbe-189ca",
      "copyType": "global",
      "cyclingMode": "multi",
      "freezeTime": "2022-12-02T18:34:18.000Z",
      "masterChangedVolumeName": "chg_volume-testVol1-afd07003-a61a",
      "masterVolumeName": "volume-testVol1-afd07003-a61a",
      "name": "rcrel6",
      "primaryRole": "master",
      "progress": 100,
      "remoteCopyID": "70",
      "state": "consistent_copying"
    }
  ]
}
```

- **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.

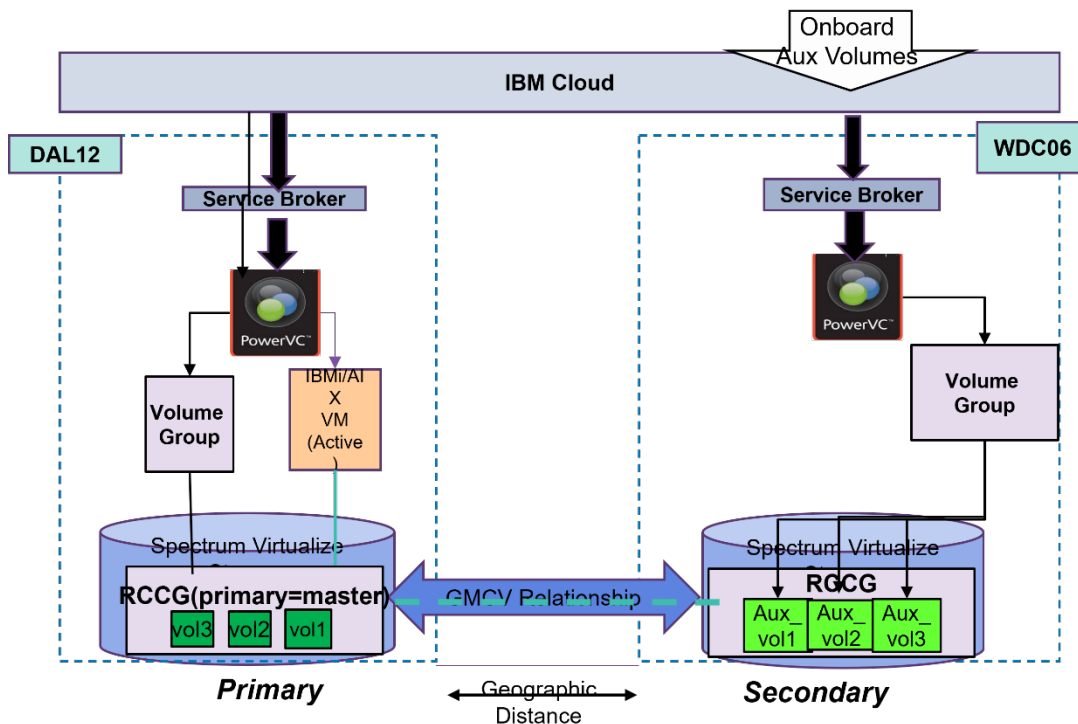
4.3 Switch to Secondary site (WDC06)

Now we must move to the secondary site, to perform the remaining steps. Set the service target of CLI environment to wdc06 Power Virtual Server Instances

```
CLI: ibmcloud pi st crn:v1:bluemix:public:power-iaas:wdc06:a/
2bb3df23c0d14ebe921397bd8aa2555a:56ee5081-f4cf-4d19-8bd6-4fdaa60c9999::
```

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 voloc --source-crn crn:v1:bluemix:public:power-iaas:dal10:
2bb3df23c0d14ebe921397bd8aa2555a:56ee5081-f4cf-4d19-8bd6-4fdaa60c8888:: --
description testOnboarding --auxiliary-volume "aux_volume-testVol1-afd07003-a61a1210664
recoveryVol1" --auxiliary-volume "aux_volume-testVol2-7a9f1ca6-acec1210664
recoveryVol2"
```

```
Creating volume onboarding under account Power IaaS Lite
ID          fd501768-4100-4127-b2b6-af9c0f099186
Description testOnboarding
```

- **--source-crn** : CRN of Power System Virtual Server Service instance in primary site(dal10). This is used to validate the authorization.
- **--auxiliary-volume**: Space separated list of auxiliary volume. Exact Auxiliary volume name as returned by auxVolumeName in volume details from primary site. Repeat this option if there are multiple auxiliary volumes to be onboarded.
- 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.

```
12:49 $ ibmcloud pi volo fd501768-4100-4127-b2b6-af9c0f099186 --json
{
  "description": "testOnboarding",
  "id": "fd501768-4100-4127-b2b6-af9c0f099186",
  "inputVolumes": [
    "aux_volume-testVol1-afd07003-a61a1210664",
    "aux_volume-testVol2-7a9f1ca6-acec1210664"
  ],
  "status": "SUCCESS",
  "creationTimestamp": "2022-11-24T06:39:47.000Z",
  "progress": 100,
  "results": {
    "onboardedVolumes": [
      "aux_volume-testVol2-7a9f1ca6-acec1210664",
      "aux_volume-testVol1-afd07003-a61a1210664"
    ],
    "volumeOnboardingFailures": []
  }
}
```

- **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**

Let's check volume and volume-group details created after the onboard operation.

```
106:47 $ ibmcloud pi vol 920b667f-5c7a-4001-ba43-fe9a3ae602b9 --json
{
  "auxVolumeName": "aux_volume-testVol2-7a9f1ca6-acec1210664",
  "auxiliary": true,
  "bootable": false,
  "consistencyGroupName": "rccg-5bbe-189ca",
  "creationDate": "2022-11-24T06:40:22.000Z",
  "diskType": "tier1",
  "groupID": "96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3",
  "lastUpdateDate": "2022-11-29T01:14:22.000Z",
  "masterVolumeName": "volume-testVol2-7a9f1ca6-acec",
  "mirroringState": "consistent_copying",
  "name": "recoveryVol2",
  "primaryRole": "master",
  "pvmInstanceIDs": [],
  "replicationEnabled": true,
  "replicationStatus": "disabled",
  "replicationType": "global",
  "shareable": false,
  "size": 1,
  "state": "available",
  "volumeID": "920b667f-5c7a-4001-ba43-fe9a3ae602b9",
  "volumePool": "Tier1-Flash-1",
  "volumeType": "Tier1-Flash-1-DR",
  "wwn": "6005076400810289700000000029E04"
}
```

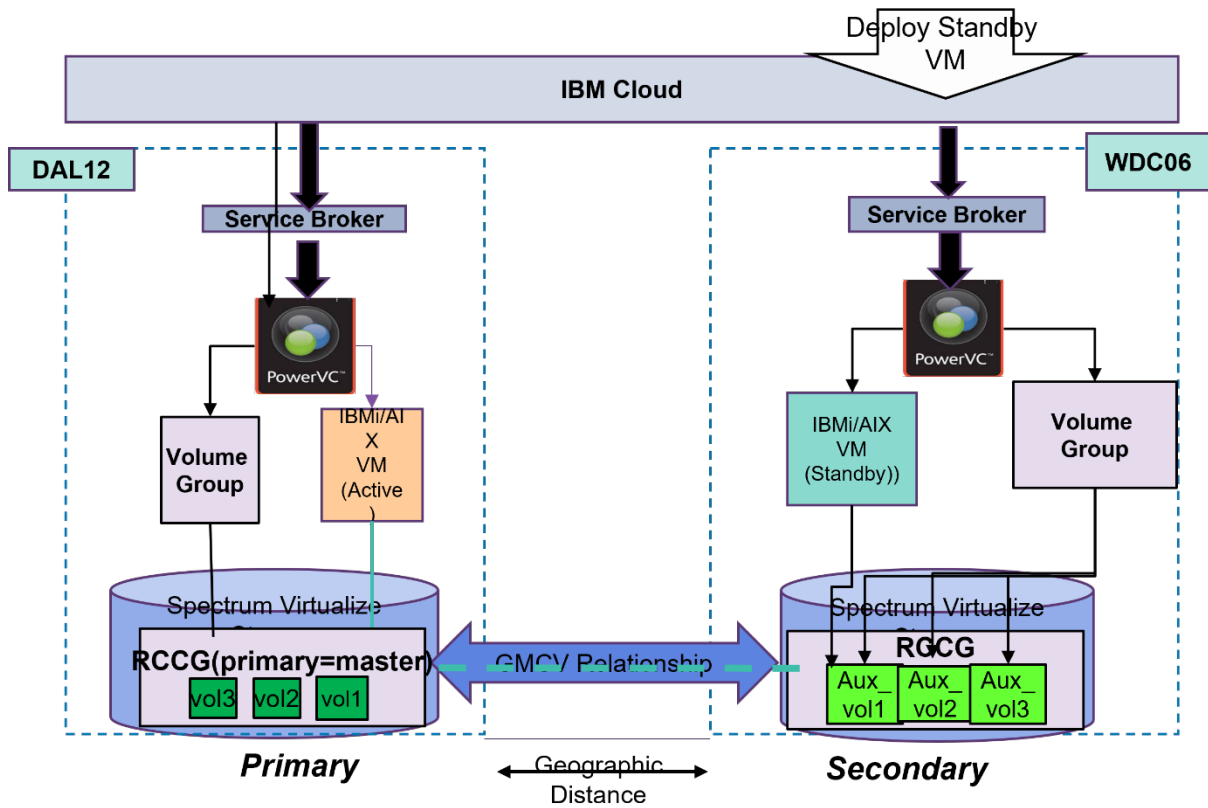
```

06:12 $ ibmcloud pi vg 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 --long --json
{
  "consistencyGroupName": "rccg-5bbe-189ca",
  "id": "96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3",
  "name": "rccg-5bbe-189ca",
  "replicationStatus": "enabled",
  "status": "available",
  "statusDescription": {
    "errors": []
  },
},
"volumeIDs": [
  "920b667f-5c7a-4001-ba43-fe9a3ae602b9",
  "ba147c20-578a-4ae1-8a94-252b6bbcd9cb"
]
}

```

4.5 Deploy 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.

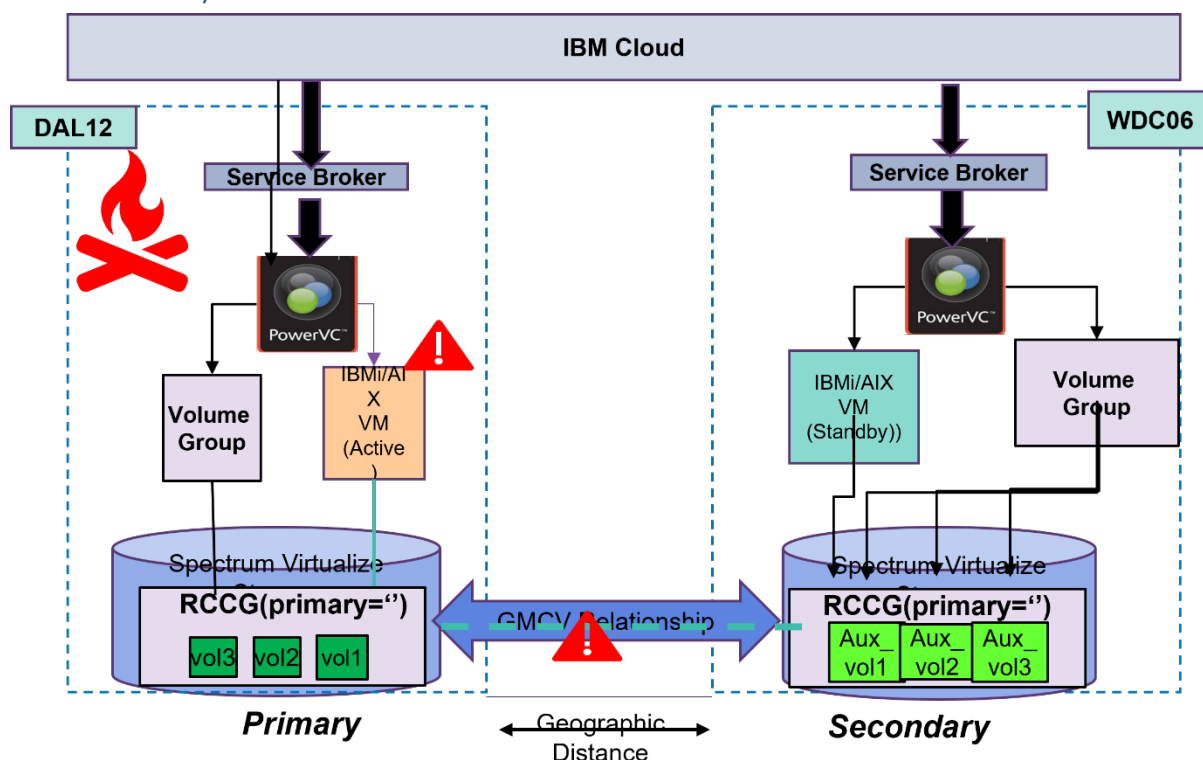


You can use the existing CLI to provision and attach volumes:

<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, or storage 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`.

```
06:56 $ ibmcloud pi vgsp 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 --allow-read-access=true
Performing stop action on volume group 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 under account
OK
Stop request on volume group 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 has been accepted.
```

On stopping the volume group, RCG will move into “Idling” state and replication status will be disabled.


```

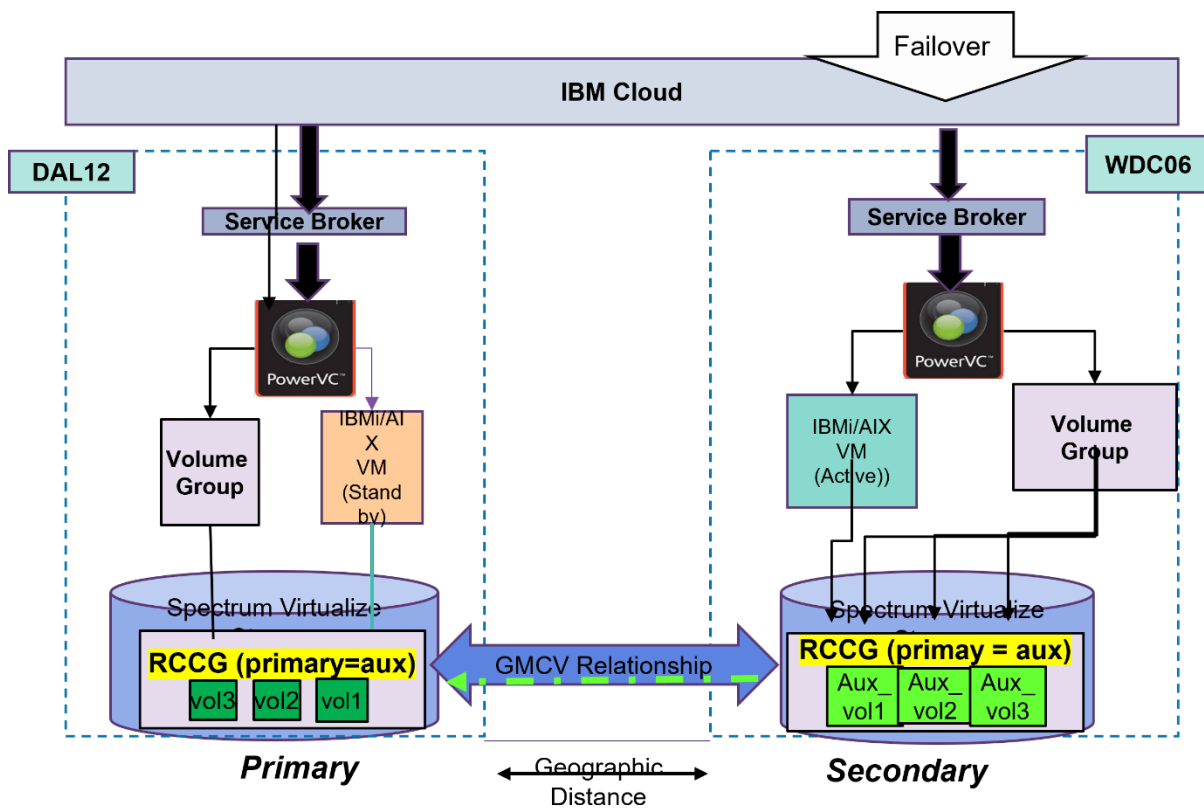
06:56 $ ibmcloud pi vgsd 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 --json
{
  "consistencyGroupName": "rccg-5bbe-189ca",
  "cyclePeriodSeconds": 500,
  "cyclingMode": "multi",
  "numOfvols": 2,
  "remoteCopyRelationshipNames": [
    "rcrel4",
    "rcrel6"
  ],
  "replicationType": "global",
  "state": "idling",
  "sync": "out_of_sync"
}

```

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, consistency group can be started to re-start the replication. You can start the volume group to switch the role to auxiliary, to change the replication direction from secondary to primary. This will allow the auxiliary volume delta changes to be copied to master volumes to be in sync.



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

```

07:03 $ ibmcloud pi vgst 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 --source auxiliary
Performing start action on volume group 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 under account
OK
Start request on volume group 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 has been accepted.

```

```
07:04 $ ibmcloud pi vgsd 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 --json
{
  "consistencyGroupName": "rccg-5bbe-189ca",
  "cyclePeriodSeconds": 500,
  "cyclingMode": "multi",
  "numOfvols": 2,
  "primaryRole": "aux",
  "remoteCopyRelationshipNames": [
    "rcrel4",
    "rcrel6"
  ],
  "replicationType": "global",
  "state": "consistent_copying"
}
```

5.3 Failback to primary site

To switch back the volume group to primary site, you can use the same volume group start command without `--source` option.

6 Disabling the 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 vgu 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca --remove-member-volume-ids afd07003-a61a-45ca-97d1-4f910272306d`

```
00:16 $ ibmcloud pi vgu 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca
--remove-member-volume-ids afd07003-a61a-45ca-97d1-4f910272306d
Updating volume group 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca under account
OK
Volume group 5bbe734a-7ec6-4f0a-a34e-8bd45fc189ca update request was accepted.
```

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 vola afd07003-a61a-45ca-97d1-4f910272306d --replication_enabled=False`

```
$ ibmcloud pi vola afd07003-a61a-45ca-97d1-4f910272306d --replication-enabled=False
Performing action on volume afd07003-a61a-45ca-97d1-4f910272306d under account Power
OK
Action on Volume ID afd07003-a61a-45ca-97d1-4f910272306d successful.
```

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

```

$ ibmcloud pi vol afd07003-a61a-45ca-97d1-4f910272306d --json
{
  "bootable": false,
  "creationDate": "2022-11-22T09:40:50.000Z",
  "diskType": "tier1",
  "lastUpdateDate": "2022-12-02T19:43:08.000Z",
  "name": "testVol1",
  "pvmInstanceIDs": [],
  "replicationEnabled": false
  "shareable": false,
  "size": 1,
  "state": "available",
  "volumeID": "afd07003-a61a-45ca-97d1-4f910272306d",
  "volumePool": "Tier1-Flash-3",
  "volumeType": "Tier1-Flash-3",
  "wwn": "60050768108081F7D0000000000484B9"
}

```

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 vgu 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3 --remove-member-volume-ids ba147c20-578a-4ae1-8a94-252b6bbcd9cb`

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

CLI: `ibmcloud pi vgd 96e037e3-9efd-4d6d-90cf-d1f6cc76d6c3`

6.1.4 Delete the auxiliary volume from secondary site.

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

CLI: `ibmcloud pi vold afd07003-a61a-45ca-97d1-4f910272306d`

7 Billing and Charging

You are charged from the location where you create a replication-enabled volume. No charges for the auxiliary volume from the remote site.

The volume of size X GB is charged based on following two components:

- The master volume is charged 2x size based on its Tier under the existing part numbers for Tier 1 and Tier 3.
- Replication capability cost is charged \$Y/GB under a new part number "GLOBAL_REPLICATION_STORAGE_GIGABYTE_HOURS" that is independent of volume tier.

Upon a site failure due to a catastrophe, metering is not available from the failed site. The auxiliary volumes are charged from remote site for its 2x size based on its tier. There is no replication capability cost for any replication-enabled volume.

8 Troubleshooting

8.1 Can I start and stop 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 to check the failures if update on volume-group fails?

Update on volume-group is an asynchronous operation. One has to check 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 reset operation on volume-group. This action does not make any change in the replication status but will set its status to "available" so that updates can be performed on 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 to 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 master 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 moves to error state in an interval of 24 hours. Any operation on this replicated volume would fail except delete operation and will set the volume in error state.

You must delete the volumes from primary site. Otherwise master volumes are charged when you delete auxiliary volume but fail to delete the master volume.

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>