Redbooks
ibm.com/redbooks

Cloud Backup Management with PowerHA SystemMirror

Dino Quintero
Umamaheswara Rao Chandolu
Aruna Sree
Kevin Gee
Sandhya Kurapati
Vijay Yalamuri

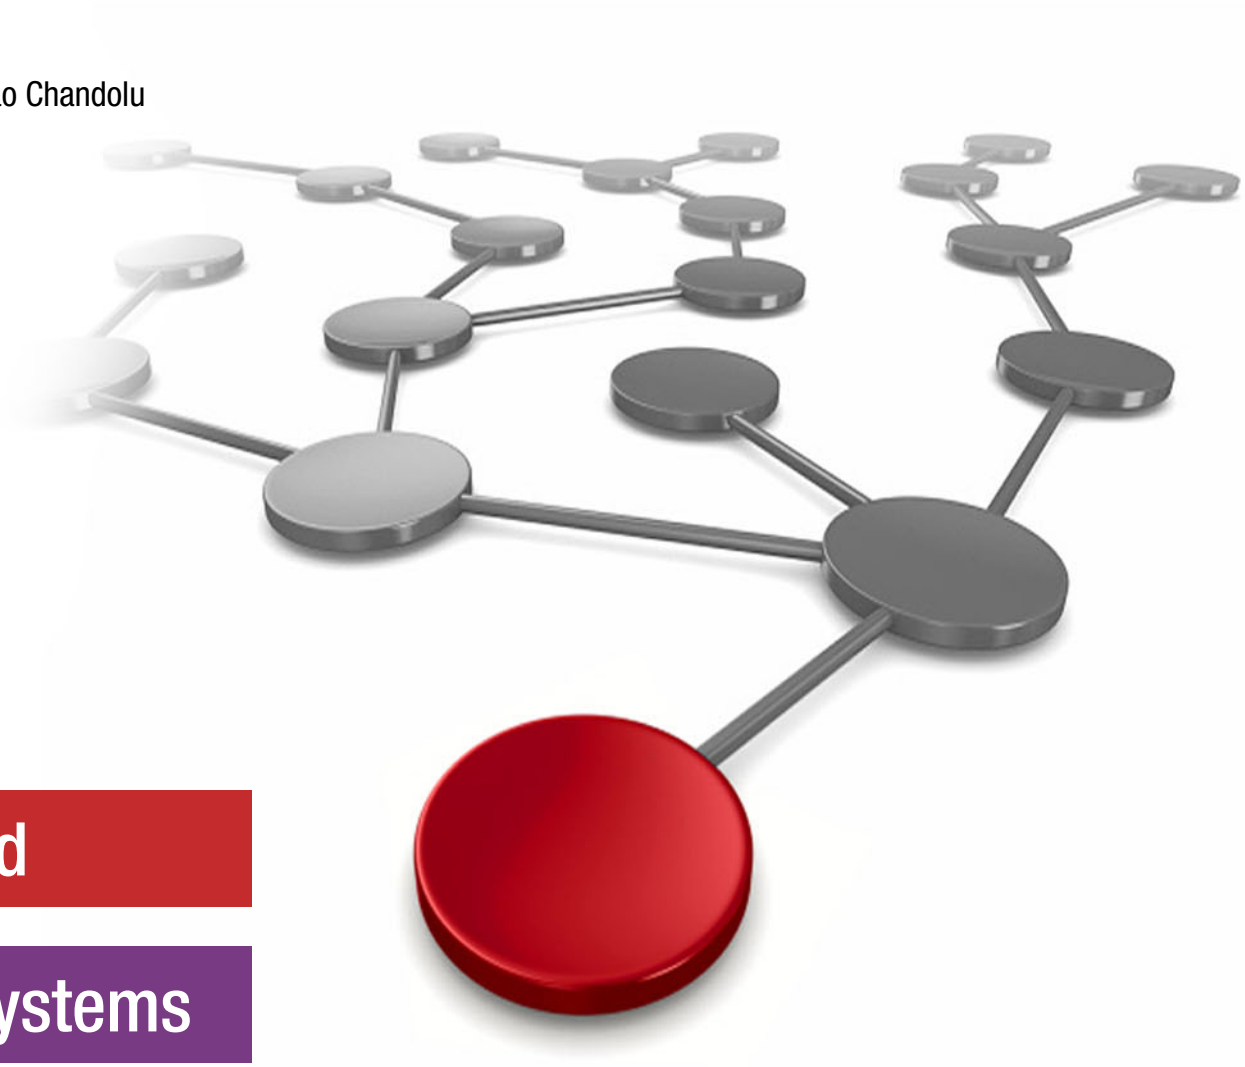

Cloud

Power Systems

IBM

Redpaper

Cloud Backup Management with PowerHA SystemMirror

This IBM® Redpaper explains how to take a backup of your data by using either of two methods, and it explains how to recover your data if there is a disaster.

This paper addresses topics for IT architects, IT specialists, developers, sellers, and anyone looking to implement and manage backups in the cloud. Moreover, this publication provides documentation to transfer the how-to-skills to the technical teams and solution guidance to the sales team. This book complements the documentation that is available at IBM Documentation and aligns with the educational materials that are provided by IBM Garage™ for Systems Technical Education.

IBM PowerHA® SystemMirror® Version 7.2.3 introduces a feature that is called Cloud Backup Management (CBM) for IBM AIX®, which enables users to create a backup of their critical data by using either of the following two methods:

1. Cloud backup: PowerHA uses the IBM SAN Volume Controller (SVC) FlashCopy® function to create a backup of your data. Then, the created backup file is uploaded to the cloud By using cloud APIs. At the time of writing, PowerHA supports the IBM and AWS cloud providers.
2. Remote storage: Back up the data to remote storage by using the storage replication method.

This publication contains the following sections:

- How backup management works
- Recovering data by using backup management
- Software and hardware requirements
- Configuring a cloud backup to PowerHA
- Configuring a remote storage backup for PowerHA
- Restoring backup data that is stored in the cloud
- Conclusion

How backup management works

To create a backup of application data, you must create a backup profile in PowerHA for the resource group for which that application is configured, as shown in Figure 1.

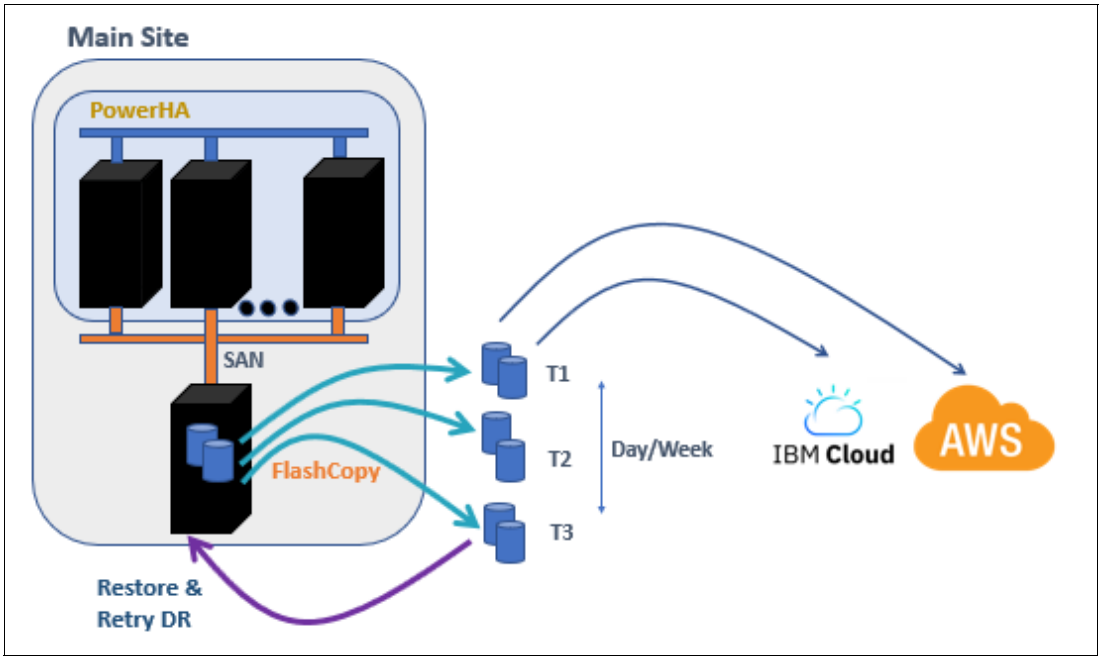

Figure 1 Backup management diagram

The backup profile requires various pieces of information:

- Backup profile name: Provide the resource group name for the data that you want to back up.
- Enable backup: Yes or No.
- Volume groups.
- Type of backup: Cloud backup or remote backup.
- Associated replicated resources.
- The SVC storage that is associated with the cluster: At the time of writing, PowerHA supports SVC storage for taking a backup.
- Notify method.

If the backup type is Cloud, provide the following extra information:

- Cloud type: At the time of writing, IBM and AWS cloud storage are supported.
- The cloud bucket name: If the backup preference is Cloud.
- Encryption algorithm: You can select **Disable**, **KMS**, and **AES** for encryption. By default encryption is disabled. AES is supported by only IBM Cloud®.
- Compression: If you set this item to **Enable**, then the backup data will be compressed.
- Backup schedule: The exact time to take the backup.
- Backup frequency: How frequent the backup must be taken (for example, once every two days).

- Incremental backup frequency in hours to take an incremental backup for every provided hour.
- Target location: Where the backup file is copied before uploading it to the cloud storage.

If the backup method is Cloud and the associated resource group is online, the backup process starts at the scheduled time, and an image file of the backup data is created at the target location that is set in the backup profile. Then, the image file is uploaded to the bucket in the cloud storage.

If the backup method is `remote_storage`, the data replication from the local SVC storage disk to the remote SVC storage disk starts and will be synchronized with the primary storage.

Recovering data by using backup management

Backup management helps you to back up the application data at a set frequency and upload the backup to cloud storage, as shown in Figure 2.

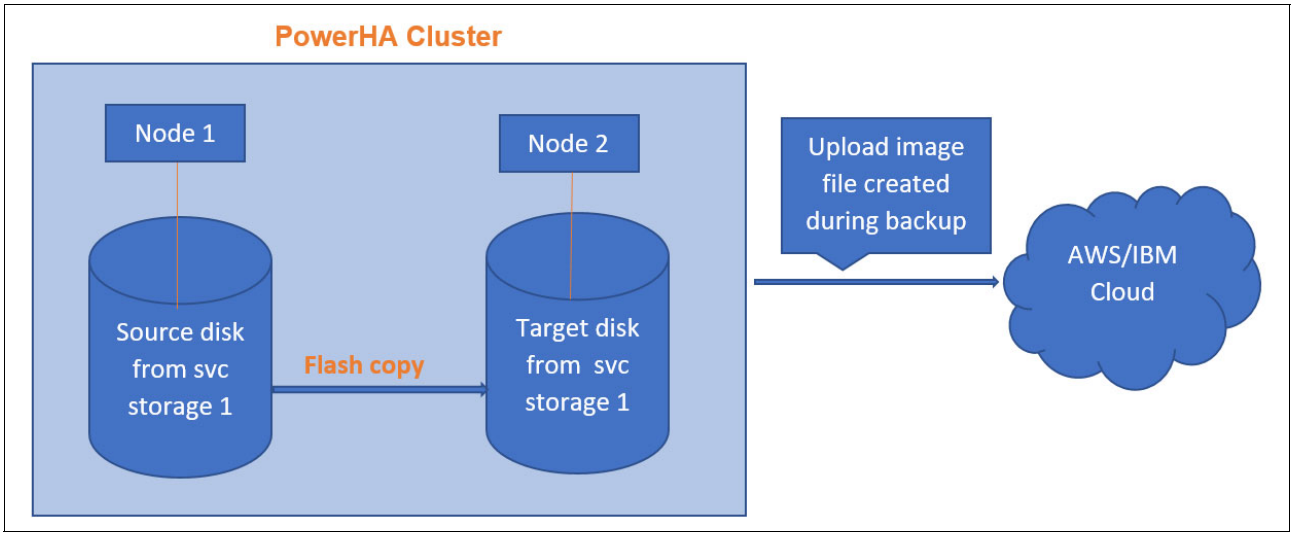

Figure 2 Backup to cloud view

You can use this feature to make backups that have the most recently updated application data, as shown in Figure 3. This data can be restored at any point to recover from a disaster situation.

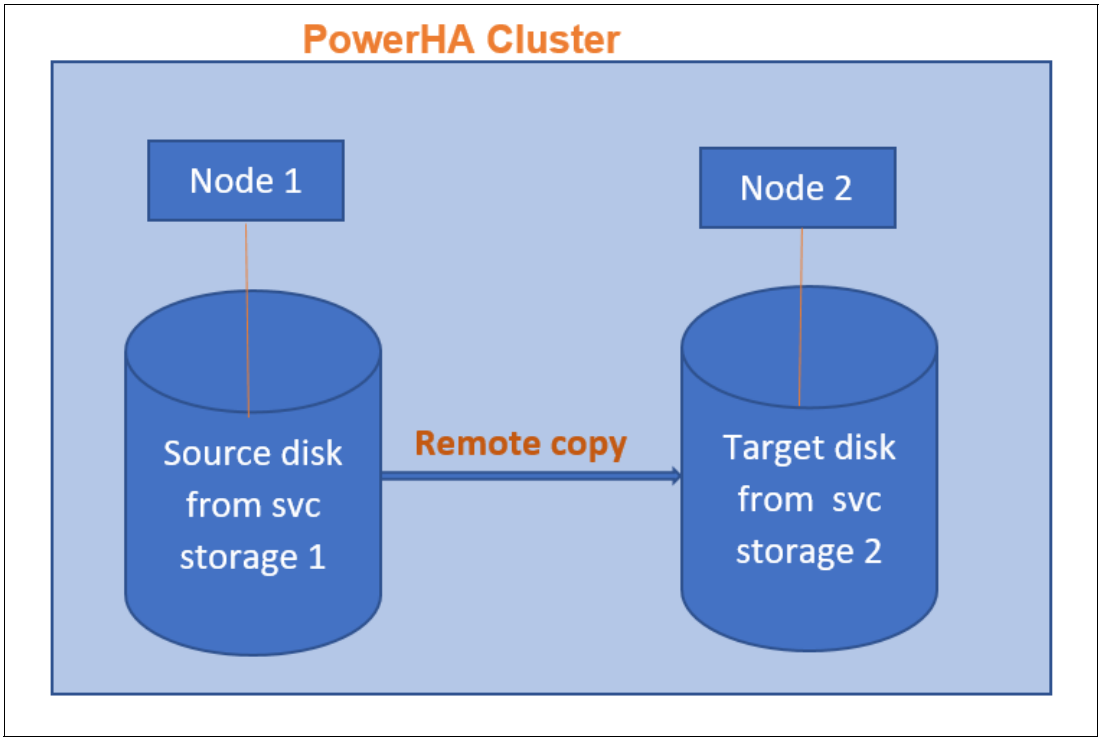

Figure 3 Backup as a remote copy view

Software and hardware requirements

Here are the requirements to create backup profiles in PowerHA:

- PowerHA SystemMirror Version 7.2.3 or later.
- Install Python Version 2.7 or later in all nodes of the cluster.
- The disks from SVC storage that are assigned to the nodes of the cluster.
- The PowerHA nodes must have cloud (IBM or AWS) connectivity before you configure the backup profile.
- The size of the source and target storage must be the same.

Configuring a cloud backup to PowerHA

If resource group RG1 with volume group VG1 (created on disks hdisk18 and hdisk19) is configured for a backup profile in PowerHA, complete the following steps:

1. Enable SSH communication to the storage of all the involved nodes.
2. Configure the SVC storage systems to use PowerHA by using either of the following methods:
   - Use the **clmgr** command to add all the SVC storage systems whose disks are assigned to nodes in the cluster to PowerHA. If you have disks from SVC storage, for example, 90.90.90.1, and SSH communication is enabled for the admin user, the command to add the storage system to PowerHA is as follows:

     ```
     clmgr add storage_system storage1 TYPE=SVC ADDRESSES=90.90.90.1 USER=admin
     BACKUP_PROFILE=true
     ```

   - The storage systems can also be added by using **smit**. Run **smitty hacmp** and select **Cluster Applications and Resources** → **Cloud Backup Configuration** → **Storage Configuration** → **Add Storage Configuration**.
3. Create subsequent mappings for the disks in the volume group and in the storage.

   A mapping is created between two disks: source and target. The source disk is a disk on which the volume group VG1 is created. In this sample scenario, the source disks are hdisk18 and hdisk19.

   With SVC storage, you can create mappings between two disks that are the same size. Check whether you have the same number of target disks and source disks that are the same size and shared across all the nodes of each site. After identifying such disks, note them (hdisk21and hdisk22).

   To create a mapping, get the name of the disks (all source and target disks) in the SVC storage because you create them by logging in to the storage either by using a GUI or a CLI. Complete the following steps:

   a. The first step in identifying the names of the disks in the storage is to get the UUID of the disks in the nodes. To get the UUID of the disks, run the command **lspv -u**, which lists all the disks with PVIDs and UUIDs, as shown in Example 1.

Example 1 List of all disks with PVIDs and UUIDs

```
hdisk        PVID                VG        UUID
hdisk18      00fab22e1277d1bc    VG1       33213600507640081O1D3F8000000000005F604214503
hdisk19      00fab22e1277d1f7    VG2       33213600507640081O1D3F8000000000005F704214503
hdisk21      00fab22e1277d270    None      33213600507640081O1D3F8000000000005F904214503
hdisk22      00fab22e1277d2ad    None      33213600507640081O1D3F8000000000005FA04214503
```

   b. After getting the UUID of the source and target hdisks, match it with the disks in SVC storage by logging in to the storage through the GUI or the CLI. To log in by using the CLI, run the command:

      ```
      # ssh admin@90.90.90.1
      SVC_STORAGE: admin>
      ```

   c. After you are logged in, use the **lshostvdiskmap** command to list all the volume disks (VDisks) that are mapped to the hosts:

      ```
      SVC_STORAGE: admin>lshostvdiskmap | grep -w <hostname>
      ```

      The command lists all the disks from storage that are mapped to a hostname. From the list, match the UUID values that you got earlier from running the **lspv -u** command.

The list is displayed as shown in Example 2 (showing the values of only the selected disks).

Example 2 Listing all disks from the storage by using the lshostvdiskmap command

```
ID name  SCSI_ID vdisk_NAME vdisk_UID
69 Node1 13 359  Node1_1    60050764008101D3F8000000000005F6 0   io_grp0  private
69 Node1 14 360  Node1_2    60050764008101D3F8000000000005F7 0   io_grp0  private
69 Node1 16 362  Node1_4    60050764008101D3F8000000000005F9 0   io_grp0  private
69 Node1 17 363  Node15     60050764008101D3F8000000000005FA 0   io_grp0  private
```

From the list shown in Example 2, you see that the names of the source disks hdisk18, hdisk19 are Node1_1, Node1_2, and the names of the target disks hdisk21, hdisk22 are Node1_4 and Node1_5.

d. Create a FlashCopy consistency group, create a FlashCopy mapping, and add the mappings to the created consistency group. To create a FlashCopy consistency group with the name `sample_consistency_group`, run following command on the storage CLI:

```
mkfcconsistgrp -name sample_consistency_group
```

e. Create the FlashCopy mappings one by one for the source disks by including the consistency group name in the commands. Here is the command to create FlashCopy mapping between source disk Node1_1 (hdisk18) to target disk Node1_4 (hdisk21):

```
mkfcmap -source Node1_1 -target Node1_4 -consistgrp sample_consistency_group
```

The command creates a mapping between the source and target disks, and adds the mapping to the consistency group in `sample_consistency_group`. You can add as many mappings as needed to the same consistency group. In Example 2, you have one more source disk, so you create another mapping and add it to the same consistency group `sample_consistency_group` by running the following command:

```
mkfcmap -source Node1_2 -target Node1_5 -consistgrp sample_consistency_group
```

f. To verify the addition of mappings and consistency groups, use the following commands:

i. **`lsfcconsistgrp`**: Lists all the consistency groups in the storage.

ii. **`lsfcconsistgrp <cg_name>`**: Displays information about a specific consistency group (cg). For example:

```
SVC_STORAGE:admin>lsfcconsistgrp fc_cg_1
id 4
name fc_cg_1
status idle_or_copied
autodelete off
start_time
FC_mapping_id 6
FC_mapping_name fcmap2
SVC_STORAGE:admin>
```

iii. **`lsfcmap`**: Lists all the mappings that are created in the storage.

iv. **`lsfcmap <fc_map_name>`**: Displays information about a specific **`fc`** map. For example:

```
SVC_STORAGE:admin>lsfcmap fcmap2
id 6
name fcmap2
source_vdisk_id 359
source_vdisk_name Node1_1
target_vdisk_id 362
```

```
target_vdisk_name Node1_4
group_id 4
group_name fc_cg_1
status idle_or_copied
progress 0
copy_rate 100
start_time
dependent_mappings 0
autodelete off
clean_progress 100
clean_rate 100
incremental off
difference 100
grain_size 256
IO_group_id 0
IO_group_name io_grp0
partner_FC_id
partner_FC_name
restoring no
rc_controlled no
keep_target no
type generic
restore_progress 0
fc_controlled no
SVC_STORAGE:admin>
```

4. Configuring a cloud backup profile to PowerHA.

   A cloud backup profile can be added to PowerHA by using the **clmgr** command or by running **smitty hacmp** and selecting **Cluster Applications and Resources** → **Cloud Backup Configuration** → **Backup Profiles** → **Add Backup Profile**.

   The **clmgr** command creates a cloud backup profile as follows:

   ```
   clmgr add backup_profile RG1 ENABLE_BACKUP=yes BACKUP_METHOD=cloud
   REPLICATED_RESOURCES=sample_consistency_group STORAGE_NAME=storage1
   BUCKET_NAME=pha-bucket TARGET_LOCATION=/tmp VOLUME_GROUP=VG1
   BACKUP_SCHEDULE=17:30
   ```

   After the backup profile is added, check the configuration by using the **clmgr query backup_profile <rg_name>** command. If needed, modify the configuration by using the **clmgr modify backup_profile** command.

   The backup profile configuration is stored in an XML configuration file.

5. Start the PowerHA cluster services.
6. As the backup profile is configured, the backup process is triggered at the scheduled backup time of 17:30, which internally creates an image file from the disk data. This file is uploaded to the cloud bucket. The default backup is triggered at 00:42.
7. The backup process takes about 45 minutes for 5 GB of data.

Figure 4 shows a sample file that was uploaded to the cloud.

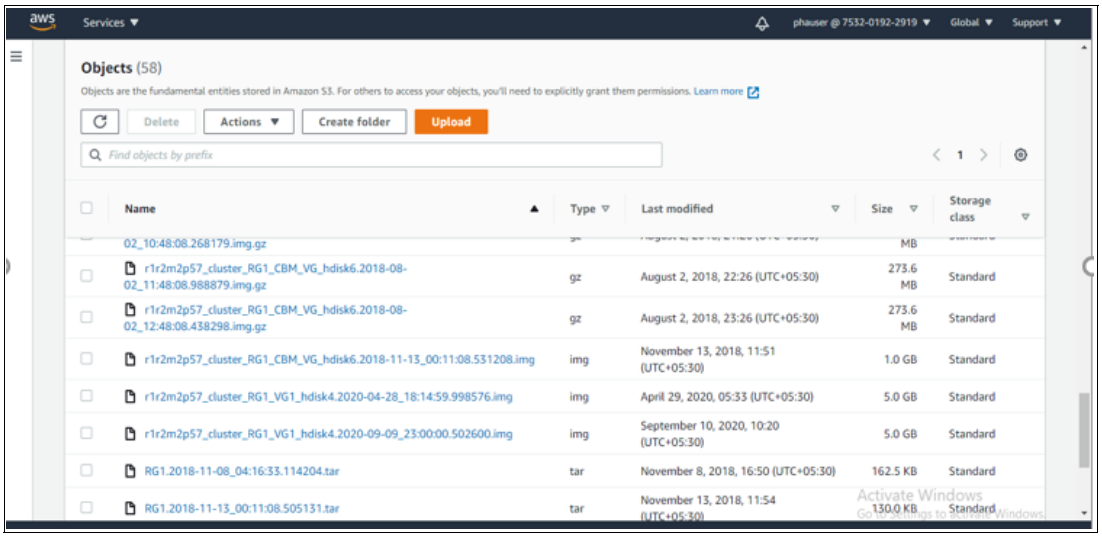

Figure 4 Sample file uploaded to the cloud

8. To query the list of backup files that were uploaded to the cloud, use the following **clmgr** command:

```
# clmgr query backup_files RG_NAME=RG1
r1r2m2p57_cluster_RG1_VG1_hdisk11.2020-11-23_05:35:00.315851.img
```

Figure 5 shows a flowchart detailing how frequently a backup is moved to the cloud.

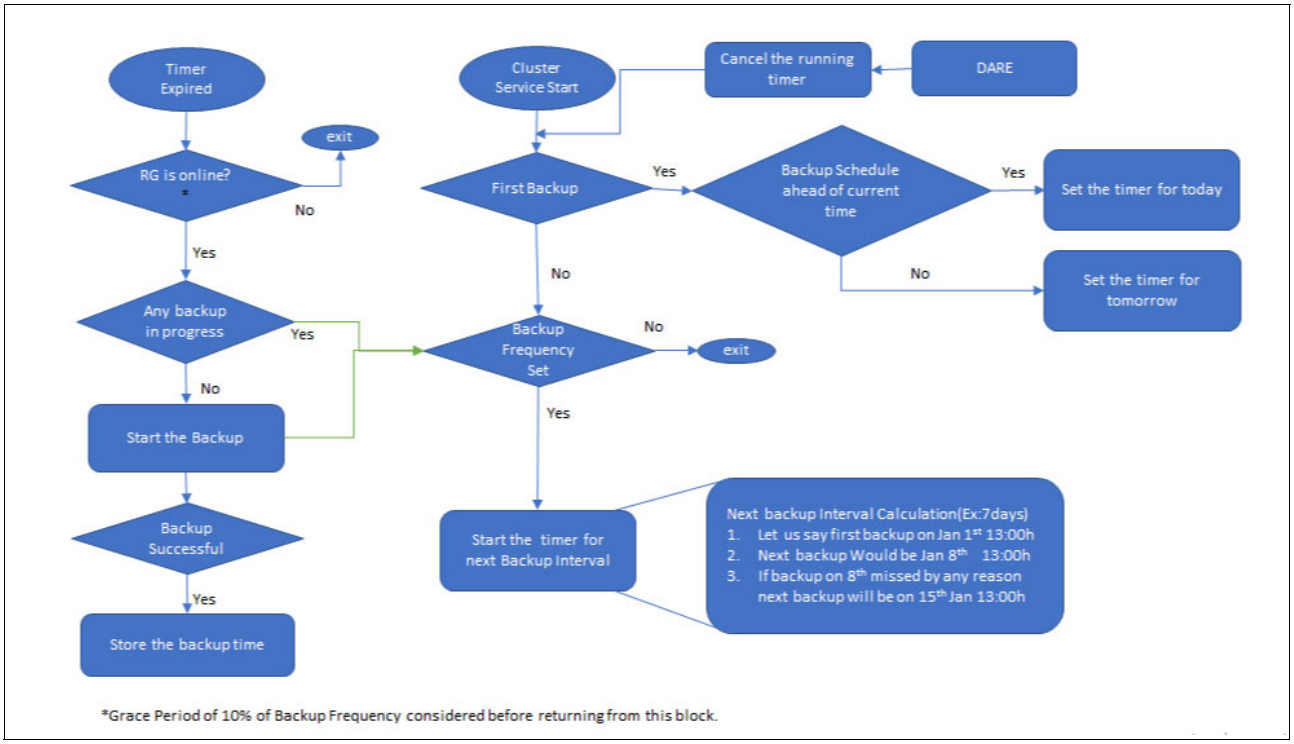

Figure 5 Flowchart explains how frequently a backup is uploaded to the cloud

Configuring a remote storage backup for PowerHA

The remote_storage backup works by creating Global Mirror remote_copy relationships between the master and auxiliary disks. The disk with the volume group is the master disk, and the disk to which the replication is directed to is the auxiliary disk. Consider the storage of the master disk as the primary storage and the storage of the auxiliary disk as the auxiliary storage.

Both the storage system profiles are added to PowerHA by using the **clmgr** command or the **smit** interface. Assume that you added storage 1 (primary storage, 90.90.90.1) and storage 2 (auxiliary storage, 90.90.90.3).

Using the same scenario, that is, RG1 with VG1 (created on hdiskAA), to add the remote_storage backup profile, you need at least one disk (the auxiliary disk) that is the same size as the master disk from a different SVC storage that is mapped to at least one of the nodes in the cluster. You also can set the auxiliary disk as a dedicated disk to one node in the cluster or, in a different context, *not* as a shared disk across all nodes in the cluster.

With the previous assumptions, to configure a remote storage backup for PowerHA, complete the following steps:

1. Log in to each storage to find the names of the master and auxiliary disks in their respective storages by matching the UUIDs. Assume that the names of the master and auxiliary disks are Node1_XX (primary storage) and Node2_YY (auxiliary storage).
2. Create a remote_copy consistency group by using the CLI to log in to the storage by using SSH by using the following command:

   ```
   SVC_STORAGE_primary: mkrcconsistgrp -name <consistency_group_name> -cluster
   <cluster_id_OR_cluster_name>
   ```

   In place of cluster_id or cluster_name, provide the name or ID of the remote storage where the auxiliary disk is.
3. Find the various storages that are connected by using the **lssystemip** command, which shows the local storage and the remote storage with all the necessary information. The command output is shown in Example 3.

Example 3 The lssystemip command shows the locate and remote storage information

```
cluster_id    cluster_name  location port_id IP_address subnet_mask gateway IP_address_6 prefix_6 gateway_6
0000010000000001   SVC_STORAGE_primary     local    1   90.90.90.1   255.255.255.0   90.90.90.9
0000010000000001   SVC_STORAGE_primary     local    2   90.90.90.2   255.255.255.0   90.90.90.9
0000010000000002   SVC_STORAGE_secondary   remote   1   90.90.90.3   255.255.255.0   90.90.90.9
0000010000000002   SVC_STORAGE_secondary   remote   2   90.90.90.4   255.255.255.0   90.90.90.9
```

To create a consistency group from SVC_STORAGE_primary (where the master disk is) and SVC_STORAGE_secondary, run the following command:

```
SVC_STORAGE_primary: mkrcconsistgrp -name sample_remote_cg -cluster
SVC_STORAGE_secondary
```

4. Create the remote_copy relationships by using the disk names that are found after matching the UUIDs by using the following **mkrcrelationship**:

   ```
   mkrcrelationship -master Node1_XX -aux Node2_YY -cluster SVC_STORAGE_secondary
   -consistgrp sample_remote_cg -global
   ```

5. To verify the remote copy consistency groups and their relationships, use the **lsrcconsistgrp** and **lsrcrelationship** commands. For example:

```
SVC_STORAGE_primary:admin>lsrcconsistgrp rc_cg_1
   id 20
   name rc_cg_1
   master_cluster_id 0000010000000001
   master_cluster_name SVC_STORAGE_primary
   aux_cluster_id 0000010000000002
   aux_cluster_name SVC_STORAGE_secondary
   primary master
   state inconsistent_stopped
   relationship_count 1
   freeze_time
   status
   sync
   copy_type global
   cycling_mode none
   cycle_period_seconds 300
   RC_rel_id 360
   RC_rel_name rcrel0

SVC_STORAGE_primary:admin> lsrcrelationship rcrel0
   id 360
   name rcrel0
   master_cluster_id 0000010000000001
   master_cluster_name SVC_STORAGE_primary
   master_vdisk_id 360
   master_vdisk_name Node1_XX
   aux_cluster_id 0000010000000002
   aux_cluster_name SVC_STORAGE_secondary
   aux_vdisk_id 908
   aux_vdisk_name Node2_YY
   primary master
   consistency_group_id 20
   consistency_group_name rc_cg_1
   state inconsistent_stopped
   bg_copy_priority 50
   progress 0
   freeze_time
   status online
   sync
   copy_type global
   cycling_mode none
   cycle_period_seconds 300
   master_change_vdisk_id
   master_change_vdisk_name
   aux_change_vdisk_id
   aux_change_vdisk_name
```

6. Configure the remote backup profile to PowerHA by completing the following steps:

   a. The remote backup profile can be added to PowerHA by using the **clmgr** command, or by using the **smit** interface by running **smitty hacmp** and selecting **Cluster Applications and Resources** → **Cloud Backup Configuration** → **Backup Profiles** → **Add Backup Profile** → **Remote Storage**.

b. The **clmgr** command creates a cloud backup profile as follows:

```
clmgr add backup_profile RG1 ENABLE_BACKUP=yes BACKUP_METHOD=remote_storage
REPLICATED_RESOURCES=sample_remote_cg STORAGE_NAME=storage1, storage2
VOLUME_GROUP=VG1
```

c. Check the configuration by using the **clmgr query backup_profile <rg_name>** command. If needed, modify the configuration by using the **clmgr modify** command.

The configuration is stored in an XML file.

7. Start the cluster services.
8. The remote backup starts copying the data. The state of the copy group in the storage changes to the `consistent_synchronized` state, as shown in Example 4.

Example 4 The lsrcconsistgrp command shows the state of the copy group

```
SVC_STORAGE_primary:admin>lsrcconsistgrp
id 20
name rc_cg_1
master_cluster_id 0000010000000001
master_cluster_name SVC_STORAGE_primary
aux_cluster_id 0000010000000002
aux_cluster_name SVC_STORAGE_secondary
primary master
state consistent_synchronized
relationship_count 1
freeze_time
status
sync
copy_type global
cycling_mode none
cycle_period_seconds 300
RC_rel_id 360
RC_rel_name rcrel0
SVC_STORAGE_primary:admin>
```

Restoring backup data that is stored in the cloud

To restore backup data that is stored in the cloud, complete the following steps:

1. Download the file from the cloud storage by using `boto3`, as shown in Example 5.

Example 5 Python sample code to download a file

```
import boto3
s3 = boto3.client('s3')
s3.download_file('BUCKET_NAME', 'BACKUP_FILE', 'LOCAL_FILE_NAME')
```

2. Previously backup content can be retrieved by using the following command:

```
dd if=/backup_file of=/dev/<target disk> bs=1024
```

3. Verify the content after the restore by using the following command:

```
importvg -V <majornumber> -y <vgname>  <hdisk>
```

Prerequisite: The restore disk size must match the primary volume group disk size.

Conclusion

Backup management helps users to automatically back up application data either to cloud storage at a specified time and frequency or mirror the data volumes to remote storage disk when the cluster services are running in a stable state. Using this feature means that the application data is always available to be restored.

Authors

This paper was produced by a team of specialists from around the world working at the IBM Redbooks, Poughkeepsie Center.

Dino Quintero is an IBM Power Systems Technical Specialist with IBM Garage for Systems. He has 25 years of experience with IBM Power Systems technologies and solutions. Dino shares his technical computing passion and expertise by leading teams developing technical content in the areas of enterprise continuous availability, enterprise systems management, high-performance computing (HPC), cloud computing, artificial intelligence including machine and deep learning, and cognitive solutions. He also is a Certified Open Group Distinguished IT Specialist. Dino is formerly from the province of Chiriqui in Panama. Dino holds a Master of Computing Information Systems degree and a Bachelor of Science degree in Computer Science from Marist College.

Umamaheswara Rao Chandolu is a Director of Engineering in Capgemini with more than 17 years of IT Industry experience in product development and delivery. He leads the product development and delivery of PowerHA, VM Recovery Manager, and IBM Spectrum® offerings from Capgemini. He holds a bachelor’s degree in computer science. He is an IBM Developer Professional Author, and he has registered 15+ patents with the US Patent Office.

Aruna Sree has been working in the IT industry for the last 15 years and has expertise in PowerHA high availability, disaster recovery with storage replication, and host-based replication. She also worked in IBM Power Systems Virtual Server Cloud testing with Geographic Logical Volume Manager (GLVM). At the time of writing, Aruna works as the lead of the PowerHA Test team.

Kevin Gee is a technical client executive at Capgemini. He has spent over 30 years in IT, including 27 years working closely with IBM Power Systems hardware (and its predecessors) and software. At multiple IBM Business Partners, he provided pre-sales and delivery consulting services for AIX, Linux, PowerHA SystemMirror, IBM Spectrum Protect, IBM Spectrum Scale, and related compute, storage, software, and data center migration offerings. He has a bachelor's degrees in computer science and Spanish and a master's degree in computer science, and has completed doctoral coursework in computer science. He is a member of the IBM Power Systems Champion Community.

Sandhya Kurapati works as a software developer at Capgemini, India. She has over 7 years of experience in application development that uses Python and C on Linux, AIX, and embedded platforms. At the time of writing, she works as part of a development team for PowerHA SystemMirror.

Vijay Yalamuri has worked at Capgemini with more than 14 years of experience in application development. He has over 10 years of experience working with AIX on Power Systems by using different components commands and libraries; security; and PowerHA software. At the time of writing, he works as a lead for the PowerHA SystemMirror L3 and development teams.

Thanks to the following people for their contributions to this project:

Wade Wallace
IBM Redbooks®, Austin Center

Reinaldo Katahira
IBM Brazil

Rajeev Sree Rama Krishna Nimmagadda
Capgemini India

Paul Moyer
Capgemini US

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- Find us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®
FlashCopy®
IBM®
IBM Cloud®
IBM Garage™
IBM Spectrum®
PowerHA®
Redbooks®
Redbooks (logo)  ®
SystemMirror®

The following terms are trademarks of other companies:

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Other company, product, or service names may be trademarks or service marks of others.

REDP-5651-00

ISBN 0738460028

Printed in U.S.A.