

Hitachi Content Platform Gateway

v4.3.8

Database Replication Setup for Windows

The objective of this document is to provide details on the configuration and use of the Hitachi Content Platform Gateway with the Hitachi Content Platform (HCP) storage system.

MK-HCPG014-04
March 2023

© 2023, 2022 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara LLC (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at <https://support.hitachivantara.com/en-us/contact-us.html>.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

- 1) Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
- 2) Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, Microsoft Edge, the Microsoft corporate logo, the Microsoft Edge logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found in the product documentation, at <https://www.hitachivantara.com/en-us/company/legal.html> or https://knowledge.hitachivantara.com/Documents/Open_Source_Software.

Table of Contents

Introduction	2
Common list of excluded tables	6
Important Notes:.....	15
Chapter 2 Two Node: Master Single node to Replica Single node Replication (Read-only copy)	17
Chapter 3 Two Node: Master Single node to Master Single node Replication (Manual DR Failover).....	26
Chapter 4 Three Node: Three Site with Master Single node to two Replica Single nodes Replication (Read-only copies)	38
Chapter 5 Three Node: Master to Master Replication for HA Cluster with optional Read-only Copy on secondary site.....	49
Chapter 6 Four Nodes: Master to Master Replication for HA Cluster with DR Failover to HA Cluster.....	63
Chapter 7 HCP Gateway Failover to Standby	83

Introduction

This document will cover the many replication configurations of the Hitachi Content Platform Gateway. It is not the intent of this document to provide complete step-by-step instructions but rather high-level guidance to configuring the various replication scenarios.

All examples in this document will use the following IP addresses:

1. hcpG1 10.1.1.11
2. hcpG2 10.1.1.12
3. hcpG3 10.1.1.13
4. hcpG4 10.1.1.14

In addition, the internal database replication queries will use a generic fixed name called 'nodeX', where the 'X' is 1-4. These names directly correspond to the database server.id. Therefore, node1 is the database for hcpG1 and so forth.

There are two 'server.id' properties, they are independent of each other, although in most cases they use the same number. First is the HCP Gateway (hcpG) configuration file C:\SAM\etc\sam\sam.properties server.id that identifies the HCP Gateway device. Second there is the database configuration file D:\MariaDB\data\my.ini server-id that distinguishes the database instances.

This document covers the following database replication configurations:

5. **Two Node:** Master node to Replica node replication for Read-only copy on second site (Single Node)
6. **Two Node:** Master node to Master node replication for Manual DR Failover (Single Node). Use this configuration also for a 2 node cluster with a shared cache. Follow the appropriate steps in this Chapter when using a 2 node cluster with a shared cache.
7. **Three Node:** Master node to 2 separate Replica nodes replication for Read-only copy on second and third sites (Single Node)
8. **Three Node:** Master to Master replication for HA **Cluster** with 2 nodes on the primary site and optional Replica for Read-only copy on second site (Cluster Node)
9. **Four Node:** Master to Master replication for HA **Cluster** with 2 nodes on the primary site and Master to Master HA **Cluster** with 2 nodes on the DR failover site (Cluster Node)

Please note that the following configuration is not supported at this time:

- **Three Nodes:** Master to Master Replication for HA Cluster on Site A with DR Failover to Single Gateway on Site B (Master-Master-Master replication).

	Production Site A	DR Failover to Site B	Read Access on Site B	Read Access on Site C
Two Node: Master to Replica DB copy (Read-only access on Site B)	Single Node		✓ (Single Node)	
Two Node:	Single Node	✓	✓	

Master to Master DB copy (DR Hot Standby on Site B)		(Single Node)	(optional)	
Two or Three Node: Master to Master HA Cluster on Site A with optional Replica DB copy on Site B (Read-only access on Site B)	Cluster Nodes	✓ (Single Node - Manual)	✓ (optional)	
Three Node: Master to Replica DB copy on 2 sites (Read-only access on Site B and Site C)	Single Node		✓ (Single Node)	✓ (Single Node)
Four Node: Master to Master HA Cluster on Site A and Master to Master HA Cluster on Site B (DR Hot standby on Site B)	Cluster Nodes	✓ (Cluster Nodes)	✓ (optional)	

Regardless of the configuration type used the replication process is basically the same. When a file is written to the primary node the metadata is immediately replicated to the other node(s). In the cases where there is no shared cache between the node(s) the data will not be accessible on the other node(s) until the original file data is copied to the storage. This is usually a short delay*.

*Be aware that HCP Gateway user Policies can be configured that could cause long delays before the data can be accessible on the alternate node(s).

WARNING:

When using a clustered or replicated set of HCP Gateways, after creating a share on the active node of the cluster or replication set, check the MariaDB error log on the non-active node(s) in the file D:\MariaDB\data*.err, where * is the name of the HCP Gateway server, for any of the following errors. If the error is for a different Master node than node1, replace node1 with that node name in the queries below

- "[ERROR] Master 'node1': Slave SQL: Error 'Table 'sam.XX_backfill' doesn't exist' on query."
- "[ERROR] Master 'node1': Slave SQL: Error 'Table 'sam.XX_migration' doesn't exist' on query."
- "[ERROR] Master 'node1': Slave SQL: Error 'Table 'sam.XX_ntfs' doesn't exist' on query."
- "[ERROR] Master 'node1': Slave SQL: Error 'Table 'sam.XX_restoration' doesn't exist' on query."

1. If any of the errors above are in the MariaDB error log, then follow the steps below to completely replicate the share to the non-active node.

1. On the non-active node that contains the error, open the Windows Failover Cluster Manager, and move the role that contains the SAM VFS service from the active node to the non-active node with the error. If using a replicated set of HCP Gateways, then stop the SAM VFS service on the active node and then start the SAM VFS service on the non-active node with the error.
2. Then on the non-active node, open a web browser to the HCP Gateway UI using the URL <https://localhost:28443/hcpg> and navigate to the **Shares** page and start the new share.
3. Open a MariaDB Command Prompt, login as the root user and issue the queries **stop slave 'node1'**; then **start slave 'node1'**;

```
MariaDB [SAM]> stop slave 'node1';  
Query OK, 0 rows affected (0.002 sec)
```

```
MariaDB [SAM]> start slave 'node1';  
Query OK, 0 rows affected (0.004 sec)
```

4. Verify that replication is synchronized between the active and non-active nodes by issuing the query **show all slaves status\G** and check the lines in **bold** below.

```
MariaDB [SAM]> show all slaves status\G  
***** 1. row *****
```

```
Connection_name: node1  
Slave_SQL_State: Slave has read all relay log; waiting for the slave I/O thread  
to update it  
Slave_IO_State: Waiting for master to send event  
Master_Host: 10.6.11.241  
Master_User: replication_user  
Master_Port: 3306  
Connect_Retry: 60  
Master_Log_File: hcpg-1-bin.000004  
Read_Master_Log_Pos: 2193608  
Relay_Log_File: hcpg-2-relay-node1.000002  
Relay_Log_Pos: 2174085  
Relay_Master_Log_File: hcpg-1-bin.000004  
Slave_IO_Running: Yes  
Slave_SQL_Running: Yes  
Replicate_Do_DB: SAM,sam  
Replicate_Ignore_DB: mysql  
Replicate_Do_Table:  
Replicate_Ignore_Table:  
Replicate_Wild_Do_Table:  
Replicate_Wild_Ignore_Table:  
SAM.%event,SAM.license,SAM.report_status,sam.%event,sam.license,sam.report_s  
tatus  
Last_Errno: 0  
Last_Error:
```

Skip_Counter: 0
 Exec_Master_Log_Pos: 2193608
 Relay_Log_Space: 2174397
 Until_Condition: None
 Until_Log_File:
 Until_Log_Pos: 0
 Master_SSL_Allowed: Yes
 Master_SSL_CA_File:
 Master_SSL_CA_Path:
 Master_SSL_Cert:
 Master_SSL_Cipher:
 Master_SSL_Key:
 Seconds_Behind_Master: 0
 Master_SSL_Verify_Server_Cert: No
Last_IO_Errno: 0
Last_IO_Error:
Last_SQL_Errno: 0
Last_SQL_Error:
 Replicate_Ignore_Server_Ids:
 Master_Server_Id: 1
 Master_SSL_Crl:
 Master_SSL_Crlpath:
 Using_Gtid: Slave_Pos
 Gtid_IO_Pos: 1-1-9726,4-4-32,3-3-32,2-2-103
 Replicate_Do_Domain_Ids:
 Replicate_Ignore_Domain_Ids:
 Parallel_Mode: conservative
 SQL_Delay: 0
 SQL_Remaining_Delay: NULL
Slave_SQL_Running_State: Slave has read all relay log; waiting for the slave I/O thread to update it
 Slave_DDL_Groups: 58
 Slave_Non_Transactional_Groups: 9666
 Slave_Transactional_Groups: 8
 Retried_transactions: 0
 Max_relay_log_size: 104857600
 Executed_log_entries: 29735
 Slave_received_heartbeats: 352950
 Slave_heartbeat_period: 30.000
 Gtid_Slave_Pos: 1-1-9726,2-2-103,3-3-32,4-4-32

5. In the HCP Gateway UI **Shares** page, stop the share that was started in Step 2.
6. In the Windows Failover Cluster Manager, move the cluster role back to the original node. If using a replicated set of HCP Gateways, stop the SAM VFS service on the node that had the error and then start the SAM VFS service on the original active node.

Beginning with the Windows version 4.1.4 of HCP Gateway, there is a new **registry.shares** parameter that needs to be added to the **C:\SAM\etc\sam\sam.properties** file.

- For a clustered pair of HCP Gateways with a shared cache, add the line **registry.shares=yes** to both nodes of the cluster. This will configure HCP Gateway to look in Windows Registry for the share configuration.
- When using database replication with or without cluster, on the active node, add the line **registry.shares=yes**. On all the other nodes that do not have a shared cache with the active node, add the line **registry.shares=no**. When using database replication without a shared cache, only 1 node can have this parameter set to **yes**.

IMPORTANT NOTE:

When using more than 1 Windows HCP Gateway with database replication or more than 1 clustered pair of Windows HCP Gateways, when the HCP Gateway active node is not available and the replica node becomes the active node, change the **registry.shares** parameter from no to yes on the new active node and restart the **SAM VFS** service. When the original active node then becomes available again and is promoted to the active node, change the **registry.shares** parameter from **yes** to **no** on the new passive replica node and restart the **SAM VFS** service.

Common list of excluded tables

There are four configurations for database tables to be ignored by replication documented below, choose the one that best fits your configuration. Note that the SAM VFS, Wildfly and MariaDB services must be stopped before making any changes to the D:\MariaDB\data\my.ini file. This will take any shares offline while the services are stopped.

1. In all configurations when NOT using a cluster or a shared cache, then these common database tables are not replicated and need to be added to, or uncommented out of, the D:\MariaDB\data\my.ini file. The my.ini should look like the following (note that the # at the beginning of a line signifies the line contains a comment so the text in the line will be ignored):

```
# ignore tables when replicate SAM database
replicate-wild-ignore-table = SAM.%event
replicate-wild-ignore-table = SAM.license
replicate-wild-ignore-table = SAM.report_status
replicate-wild-ignore-table = SAM.archive_state
```

```
replicate-wild-ignore-table = sam.%event
replicate-wild-ignore-table = sam.license
replicate-wild-ignore-table = sam.report_status
replicate-wild-ignore-table = sam.archive_state
```

```
# ignore more tables when replicate SAM database to DR HCP Gateway
server(replication mode without sharing cache)
replicate-wild-ignore-table = SAM.%backfill
replicate-wild-ignore-table = SAM.%restoration
#replicate-wild-ignore-table = SAM.%protection
replicate-wild-ignore-table = SAM.%purge
```



```
replicate-wild-ignore-table = SAM.%ntfs
replicate-wild-ignore-table = SAM.%migration
```

```
replicate-wild-ignore-table = sam.%backfill
replicate-wild-ignore-table = sam.%restoration
#replicate-wild-ignore-table = sam.%protection
replicate-wild-ignore-table = sam.%purge
replicate-wild-ignore-table = sam.%ntfs
replicate-wild-ignore-table = sam.%migration
```

2. When using a pair of clustered Gateways with a shared cache, then these common database tables are not replicated and need to be added to the D:\MariaDB\data\my.ini file. Go to option 4 below when using a 4 node cluster, with 2 nodes with a shared cache in site 1 and 2 nodes with a shared cache in site 2. The my.ini should look like the following (note that the # at the beginning of a line signifies the line contains a comment so the text in the line will be ignored):

```
# ignore tables when replicate SAM database (cluster mode)
replicate-wild-ignore-table = SAM.%event
replicate-wild-ignore-table = SAM.license
```

```
replicate-wild-ignore-table = sam.%event
replicate-wild-ignore-table = sam.license
```

```
# ignore more tables when replicate SAM database to DR HCP Gateway server
(replication mode with sharing cache)
```

```
#replicate-wild-ignore-table = SAM.%backfill
#replicate-wild-ignore-table = SAM.%restoration
#replicate-wild-ignore-table = SAM.%protection
#replicate-wild-ignore-table = SAM.%purge
#replicate-wild-ignore-table = SAM.%ntfs
#replicate-wild-ignore-table = SAM.%migration
```

```
#replicate-wild-ignore-table = sam.%backfill
#replicate-wild-ignore-table = sam.%restoration
#replicate-wild-ignore-table = sam.%protection
#replicate-wild-ignore-table = sam.%purge
#replicate-wild-ignore-table = sam.%ntfs
#replicate-wild-ignore-table = sam.%migration
```

3. When using a pair of clustered Gateways with a shared cache in site 1 and a single Gateway in site 2, then these common database tables are not replicated and need to be added to the D:\MariaDB\data\my.ini file. The my.ini ignore table statements should look like the following on node1 in site 1 (note that the # at the beginning of a line signifies the line contains a comment so the text in the line will be ignored):

```
# ignore tables when replicate SAM database (cluster mode)
node2.replicate-wild-ignore-table = SAM.%event
node3.replicate-wild-ignore-table = SAM.%event
```

```
node2.replicate-wild-ignore-table = SAM.license
node3.replicate-wild-ignore-table = SAM.license
```

```
node3.replicate-wild-ignore-table = SAM.report_status
```

```
node3.replicate-wild-ignore-table = SAM.archive_state
```

```
node2.replicate-wild-ignore-table = sam.%event  
node3.replicate-wild-ignore-table = sam.%event
```

```
node2.replicate-wild-ignore-table = sam.license  
node3.replicate-wild-ignore-table = sam.license
```

```
node3.replicate-wild-ignore-table = sam.report_status
```

```
node3.replicate-wild-ignore-table = sam.archive_state
```

```
# ignore more tables when replicate SAM database to DR HCP Gateway server  
(replication mode without sharing cache)
```

```
node3.replicate-wild-ignore-table = SAM.%backfill  
node3.replicate-wild-ignore-table = SAM.%restoration  
node3.replicate-wild-ignore-table = SAM.%purge  
node3.replicate-wild-ignore-table = SAM.%ntfs  
node3.replicate-wild-ignore-table = SAM.%migration
```

```
node3.replicate-wild-ignore-table = sam.%backfill  
node3.replicate-wild-ignore-table = sam.%restoration  
node3.replicate-wild-ignore-table = sam.%purge  
node3.replicate-wild-ignore-table = sam.%ntfs  
node3.replicate-wild-ignore-table = sam.%migration
```

The my.ini ignore table statements should look like the following on node2 in site 1:

```
    # ignore tables when replicate SAM database (cluster mode)
```

```
node1.replicate-wild-ignore-table = SAM.%event  
node3.replicate-wild-ignore-table = SAM.%event
```

```
node1.replicate-wild-ignore-table = SAM.license  
node3.replicate-wild-ignore-table = SAM.license
```

```
node3.replicate-wild-ignore-table = SAM.report_status
```

```
node3.replicate-wild-ignore-table = SAM.archive_state
```

```
node1.replicate-wild-ignore-table = sam.%event  
node3.replicate-wild-ignore-table = sam.%event
```

```
node1.replicate-wild-ignore-table = sam.license  
node3.replicate-wild-ignore-table = sam.license
```

```
node3.replicate-wild-ignore-table = sam.report_status
```

```
node3.replicate-wild-ignore-table = sam.archive_state
```

```
# ignore more tables when replicate SAM database to DR HCP Gateway server  
(replication mode without sharing cache)
```

```
node3.replicate-wild-ignore-table = SAM.%backfill
node3.replicate-wild-ignore-table = SAM.%restoration
node3.replicate-wild-ignore-table = SAM.%purge
node3.replicate-wild-ignore-table = SAM.%ntfs
node3.replicate-wild-ignore-table = SAM.%migration
```

```
node3.replicate-wild-ignore-table = sam.%backfill
node3.replicate-wild-ignore-table = sam.%restoration
node3.replicate-wild-ignore-table = sam.%purge
node3.replicate-wild-ignore-table = sam.%ntfs
node3.replicate-wild-ignore-table = sam.%migration
```

The my.ini ignore table statements should look like the following on node3 in site 2:

```
# ignore tables when replicate SAM database (cluster mode)
node1.replicate-wild-ignore-table = SAM.%event
node2.replicate-wild-ignore-table = SAM.%event

node1.replicate-wild-ignore-table = SAM.license
node2.replicate-wild-ignore-table = SAM.license

node1.replicate-wild-ignore-table = SAM.report_status
node2.replicate-wild-ignore-table = SAM.report_status

node1.replicate-wild-ignore-table = SAM.archive_state
node2.replicate-wild-ignore-table = SAM.archive_state

node1.replicate-wild-ignore-table = sam.%event
node2.replicate-wild-ignore-table = sam.%event

node1.replicate-wild-ignore-table = sam.license
node2.replicate-wild-ignore-table = sam.license

node1.replicate-wild-ignore-table = sam.report_status
node2.replicate-wild-ignore-table = sam.report_status

node1.replicate-wild-ignore-table = sam.archive_state
node2.replicate-wild-ignore-table = sam.archive_state

# ignore more tables when replicate SAM database to DR HCP Gateway server
(replication mode without sharing cache)

node1.replicate-wild-ignore-table = SAM.%backfill
node2.replicate-wild-ignore-table = SAM.%backfill

node1.replicate-wild-ignore-table = SAM.%restoration
node2.replicate-wild-ignore-table = SAM.%restoration

#node1.replicate-wild-ignore-table = SAM.%protection
#node2.replicate-wild-ignore-table = SAM.%protection

node1.replicate-wild-ignore-table = SAM.%purge
node2.replicate-wild-ignore-table = SAM.%purge
```

```
node1.replicate-wild-ignore-table = SAM.%ntfs
node2.replicate-wild-ignore-table = SAM.%ntfs
```

```
node1.replicate-wild-ignore-table = SAM.%migration
node2.replicate-wild-ignore-table = SAM.%migration
```

```
node1.replicate-wild-ignore-table = sam.%backfill
node2.replicate-wild-ignore-table = sam.%backfill
```

```
node1.replicate-wild-ignore-table = sam.%restoration
node2.replicate-wild-ignore-table = sam.%restoration
```

```
#node1.replicate-wild-ignore-table = sam.%protection
#node2.replicate-wild-ignore-table = sam.%protection
```

```
node1.replicate-wild-ignore-table = sam.%purge
node2.replicate-wild-ignore-table = sam.%purge
```

```
node1.replicate-wild-ignore-table = sam.%ntfs
node2.replicate-wild-ignore-table = sam.%ntfs
```

```
node1.replicate-wild-ignore-table = sam.%migration
node2.replicate-wild-ignore-table = sam.%migration
```

4. When using a pair of clustered Gateways with a shared cache in site 1 and a pair of clustered Gateways with a shared cache in site 2, then these common database tables are not replicated and need to be added to the D:\MariaDB\data\my.ini file. The my.ini ignore table statements should look like the following on node1 in site 1 (note that the # at the beginning of a line signifies the line contains a comment so the text in the line will be ignored):

```
# ignore tables when replicate SAM database (cluster mode)
```

```
node2.replicate-wild-ignore-table = SAM.%event
node3.replicate-wild-ignore-table = SAM.%event
node4.replicate-wild-ignore-table = SAM.%event
```

```
node2.replicate-wild-ignore-table = SAM.license
node3.replicate-wild-ignore-table = SAM.license
node4.replicate-wild-ignore-table = SAM.license
```

```
node3.replicate-wild-ignore-table = SAM.report_status
node4.replicate-wild-ignore-table = SAM.report_status
```

```
node3.replicate-wild-ignore-table = SAM.archive_state
node4.replicate-wild-ignore-table = SAM.archive_state
```

```
node2.replicate-wild-ignore-table = sam.%event
node3.replicate-wild-ignore-table = sam.%event
node4.replicate-wild-ignore-table = sam.%event
```

```
node2.replicate-wild-ignore-table = sam.license
node3.replicate-wild-ignore-table = sam.license
```

```

node4.replicate-wild-ignore-table = sam.license

node3.replicate-wild-ignore-table = sam.report_status
node4.replicate-wild-ignore-table = sam.report_status

node3.replicate-wild-ignore-table = sam.archive_state
node4.replicate-wild-ignore-table = sam.archive_state

# ignore more tables when replicate SAM database to DR HCP Gateway server
(replication mode without sharing cache)
node3.replicate-wild-ignore-table = SAM.%backfill
node4.replicate-wild-ignore-table = SAM.%backfill

node3.replicate-wild-ignore-table = SAM.%restoration
node4.replicate-wild-ignore-table = SAM.%restoration

#node3.replicate-wild-ignore-table = SAM.%protection
#node4.replicate-wild-ignore-table = SAM.%protection

node3.replicate-wild-ignore-table = SAM.%purge
node4.replicate-wild-ignore-table = SAM.%purge

node3.replicate-wild-ignore-table = SAM.%ntfs
node4.replicate-wild-ignore-table = SAM.%ntfs

node3.replicate-wild-ignore-table = SAM.%migration
node4.replicate-wild-ignore-table = SAM.%migration

node3.replicate-wild-ignore-table = sam.%backfill
node4.replicate-wild-ignore-table = sam.%backfill

node3.replicate-wild-ignore-table = sam.%restoration
node4.replicate-wild-ignore-table = sam.%restoration

#node3.replicate-wild-ignore-table = sam.%protection
#node4.replicate-wild-ignore-table = sam.%protection

node3.replicate-wild-ignore-table = sam.%purge
node4.replicate-wild-ignore-table = sam.%purge

node3.replicate-wild-ignore-table = sam.%ntfs
node4.replicate-wild-ignore-table = sam.%ntfs

node3.replicate-wild-ignore-table = sam.%migration
node4.replicate-wild-ignore-table = sam.%migration

```

The my.ini ignore table statements should look like the following on node2 in site 1:

```

# ignore tables when replicate SAM database (cluster mode)
node1.replicate-wild-ignore-table = SAM.%event
node3.replicate-wild-ignore-table = SAM.%event
node4.replicate-wild-ignore-table = SAM.%event

```

```

node1.replicate-wild-ignore-table = SAM.license
node3.replicate-wild-ignore-table = SAM.license
node4.replicate-wild-ignore-table = SAM.license

node3.replicate-wild-ignore-table = SAM.report_status
node4.replicate-wild-ignore-table = SAM.report_status

node3.replicate-wild-ignore-table = SAM.archive_state
node4.replicate-wild-ignore-table = SAM.archive_state
node1.replicate-wild-ignore-table = sam.%event
node3.replicate-wild-ignore-table = sam.%event
node4.replicate-wild-ignore-table = sam.%event

node1.replicate-wild-ignore-table = sam.license
node3.replicate-wild-ignore-table = sam.license
node4.replicate-wild-ignore-table = sam.license

node3.replicate-wild-ignore-table = sam.report_status
node4.replicate-wild-ignore-table = sam.report_status

node3.replicate-wild-ignore-table = sam.archive_state
node4.replicate-wild-ignore-table = sam.archive_state

# ignore more tables when replicate SAM database to DR HCP Gateway server
(replication mode without sharing cache)
node3.replicate-wild-ignore-table = SAM.%backfill
node4.replicate-wild-ignore-table = SAM.%backfill

node3.replicate-wild-ignore-table = SAM.%restoration
node4.replicate-wild-ignore-table = SAM.%restoration

#node3.replicate-wild-ignore-table = SAM.%protection
#node4.replicate-wild-ignore-table = SAM.%protection

node3.replicate-wild-ignore-table = SAM.%purge
node4.replicate-wild-ignore-table = SAM.%purge

node3.replicate-wild-ignore-table = SAM.%ntfs
node4.replicate-wild-ignore-table = SAM.%ntfs

node3.replicate-wild-ignore-table = SAM.%migration
node4.replicate-wild-ignore-table = SAM.%migration

node3.replicate-wild-ignore-table = sam.%backfill
node4.replicate-wild-ignore-table = sam.%backfill

node3.replicate-wild-ignore-table = sam.%restoration
node4.replicate-wild-ignore-table = sam.%restoration

#node3.replicate-wild-ignore-table = sam.%protection
#node4.replicate-wild-ignore-table = sam.%protection

```

```
node3.replicate-wild-ignore-table = sam.%purge
node4.replicate-wild-ignore-table = sam.%purge
```

```
node3.replicate-wild-ignore-table = sam.%ntfs
node4.replicate-wild-ignore-table = sam.%ntfs
```

```
node3.replicate-wild-ignore-table = sam.%migration
node4.replicate-wild-ignore-table = sam.%migration
```

The my.ini ignore table statements should look like the following on node3 in site 2:

```
# ignore tables when replicate SAM database (cluster mode)
```

```
node1.replicate-wild-ignore-table = SAM.%event
node2.replicate-wild-ignore-table = SAM.%event
node4.replicate-wild-ignore-table = SAM.%event
```

```
node1.replicate-wild-ignore-table = SAM.license
node2.replicate-wild-ignore-table = SAM.license
node4.replicate-wild-ignore-table = SAM.license
```

```
node1.replicate-wild-ignore-table = SAM.report_status
node2.replicate-wild-ignore-table = SAM.report_status
```

```
node1.replicate-wild-ignore-table = SAM.archive_state
node2.replicate-wild-ignore-table = SAM.archive_state
```

```
node1.replicate-wild-ignore-table = sam.%event
node2.replicate-wild-ignore-table = sam.%event
node4.replicate-wild-ignore-table = sam.%event
```

```
node1.replicate-wild-ignore-table = sam.license
node2.replicate-wild-ignore-table = sam.license
node4.replicate-wild-ignore-table = sam.license
```

```
node1.replicate-wild-ignore-table = sam.report_status
node2.replicate-wild-ignore-table = sam.report_status
node1.replicate-wild-ignore-table = sam.archive_state
node2.replicate-wild-ignore-table = sam.archive_state
```

```
# ignore more tables when replicate SAM database to DR HCP Gateway server
(replication mode without sharing cache)
```

```
node1.replicate-wild-ignore-table = SAM.%backfill
node2.replicate-wild-ignore-table = SAM.%backfill
```

```
node1.replicate-wild-ignore-table = SAM.%restoration
node2.replicate-wild-ignore-table = SAM.%restoration
```

```
#node1.replicate-wild-ignore-table = SAM.%protection
#node2.replicate-wild-ignore-table = SAM.%protection
```

```
node1.replicate-wild-ignore-table = SAM.%purge
node2.replicate-wild-ignore-table = SAM.%purge
```

```

node1.replicate-wild-ignore-table = SAM.%ntfs
node2.replicate-wild-ignore-table = SAM.%ntfs

node1.replicate-wild-ignore-table = SAM.%migration
node2.replicate-wild-ignore-table = SAM.%migration

node1.replicate-wild-ignore-table = sam.%backfill
node2.replicate-wild-ignore-table = sam.%backfill

node1.replicate-wild-ignore-table = sam.%restoration
node2.replicate-wild-ignore-table = sam.%restoration

#node1.replicate-wild-ignore-table = sam.%protection
#node2.replicate-wild-ignore-table = sam.%protection

node1.replicate-wild-ignore-table = sam.%purge
node2.replicate-wild-ignore-table = sam.%purge

node1.replicate-wild-ignore-table = sam.%ntfs
node2.replicate-wild-ignore-table = sam.%ntfs

node1.replicate-wild-ignore-table = sam.%migration
node2.replicate-wild-ignore-table = sam.%migration

```

The my.ini ignore table statements should look like the following on node4 in site 2:

```

# ignore tables when replicate SAM database (cluster mode)
node1.replicate-wild-ignore-table = SAM.%event
node2.replicate-wild-ignore-table = SAM.%event
node3.replicate-wild-ignore-table = SAM.%event

node1.replicate-wild-ignore-table = SAM.license
node2.replicate-wild-ignore-table = SAM.license
node3.replicate-wild-ignore-table = SAM.license

node1.replicate-wild-ignore-table = SAM.report_status
node2.replicate-wild-ignore-table = SAM.report_status

node1.replicate-wild-ignore-table = SAM.archive_state
node2.replicate-wild-ignore-table = SAM.archive_state

node1.replicate-wild-ignore-table = sam.%event
node2.replicate-wild-ignore-table = sam.%event
node3.replicate-wild-ignore-table = sam.%event

node1.replicate-wild-ignore-table = sam.license
node2.replicate-wild-ignore-table = sam.license
node3.replicate-wild-ignore-table = sam.license

node1.replicate-wild-ignore-table = sam.report_status
node2.replicate-wild-ignore-table = sam.report_status

node1.replicate-wild-ignore-table = sam.archive_state

```



```

node2.replicate-wild-ignore-table = sam.archive_state
# ignore more tables when replicate SAM database to DR HCP Gateway server
(replication mode without sharing cache)
node1.replicate-wild-ignore-table = SAM.%backfill
node2.replicate-wild-ignore-table = SAM.%backfill

node1.replicate-wild-ignore-table = SAM.%restoration
node2.replicate-wild-ignore-table = SAM.%restoration

#node1.replicate-wild-ignore-table = SAM.%protection
#node2.replicate-wild-ignore-table = SAM.%protection

node1.replicate-wild-ignore-table = SAM.%purge
node2.replicate-wild-ignore-table = SAM.%purge

node1.replicate-wild-ignore-table = SAM.%ntfs
node2.replicate-wild-ignore-table = SAM.%ntfs

node1.replicate-wild-ignore-table = SAM.%migration
node2.replicate-wild-ignore-table = SAM.%migration

node1.replicate-wild-ignore-table = sam.%backfill
node2.replicate-wild-ignore-table = sam.%backfill

node1.replicate-wild-ignore-table = sam.%restoration
node2.replicate-wild-ignore-table = sam.%restoration

#node1.replicate-wild-ignore-table = sam.%protection
#node2.replicate-wild-ignore-table = sam.%protection

node1.replicate-wild-ignore-table = sam.%purge
node2.replicate-wild-ignore-table = sam.%purge

node1.replicate-wild-ignore-table = sam.%ntfs
node2.replicate-wild-ignore-table = sam.%ntfs

node1.replicate-wild-ignore-table = sam.%migration
node2.replicate-wild-ignore-table = sam.%migration

```

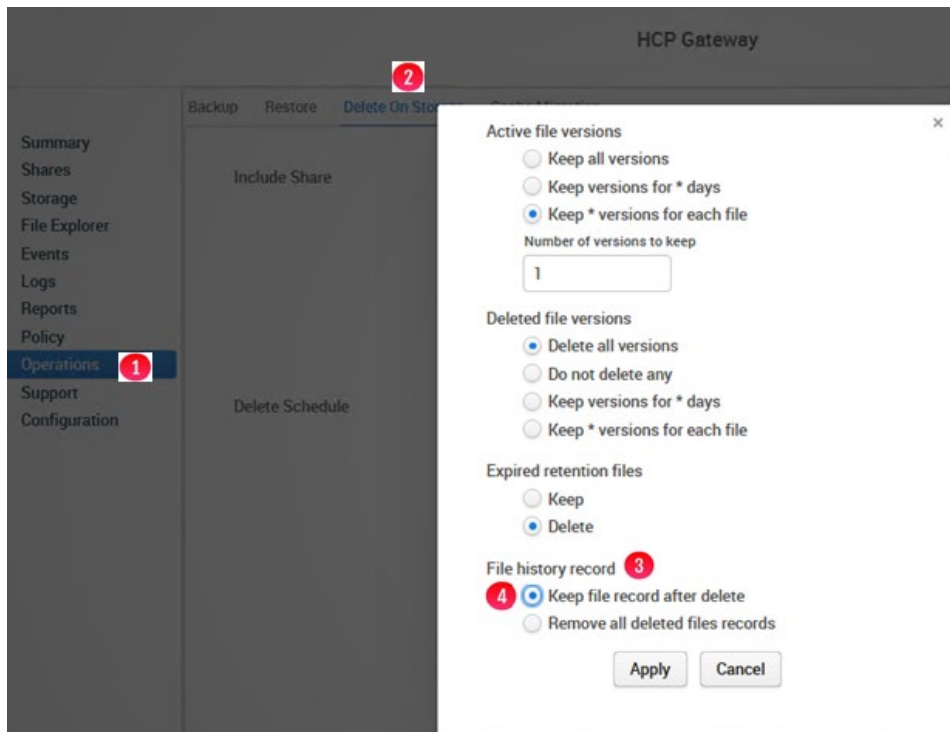
WARNING: No matter the number of HGP Gateways, when using replication or clustering, only one Gateway may be active for user share access at a time.

WARNING: If the Active node fails and a Passive node takes over as the new Active node, if a Legal Hold was being applied to file(s) on the original Active node and the Passive node is not able to connect to the HCP storage, the Legal Hold may not be applied until the original Active node comes back online and sets the Legal Hold on the file(s).

Important Notes:

1. Only 1 Gateway device in a replication set or set of clustered Gateways can be active at a time.

2. The files must be written to the HCP and if there is a second HCP, replicated to the second HCP, before the content will be available on the non-active node(s) that have a separate cache. For example, if using a Copy or Tiering policy that waits 5 days before writing files to HCP storage, then you must wait 5 days before you see the file content on the non-active node(s) that are not sharing the cache.
3. For the "Two Node: Master to Replica node", "Three Node: Master to 2 Replica nodes" and "Three Node: Master to Master HA Cluster with Replica for Read-only copy", if any data changes or writes are made to the Gateway on Site B or Site C, manual intervention is required to replicate the metadata back to Site A.
4. These instructions can be used to setup the initial replication and can also be used to reset replication if one or more of the nodes encounters a replication error.
5. **IMPORTANT WARNING:** If the replication is configured or resynchronized after the Gateways are already in use, for example, shares and files are created, BE VERY CAREFUL to verify that the existing replication is currently synchronized without any errors between ALL Gateways by issuing the "show all slaves statusG" query on ALL of the Gateways in the replication or cluster set. Also, verify that all of the data is on the active node, referred to as HCPG1 in this documentation. These steps will REPLACE the database on HCPG2, and if they exist, HCPG3 and HCPG4, with the data from the database on HCPG1. If there is data on HCPG2, and if they exist, HCPG3 and HCPG4, that was not replicated to HCPG1, then the data on HCPG2, and if they exist, HCPG3 and HCPG4, WILL BE LOST.
6. **IMPORTANT NOTE:** It is required to perform this replication synchronization when there are no users actively using the shares on the Gateway.
7. It is highly recommended to perform the replication setup before any data is written to any of the HCP Gateway nodes.
8. When using a clustered or replicated set of HCP Gateways and configuring the Delete on Storage feature on the UI Operations (1) -> Delete on Storage page (2), in the **File history record** section (3), it is required to select **Keep all file records after delete** (4).




Chapter 2 Two Node: Master Single node to Replica Single node Replication (Read-only copy)

In this configuration, the database on hcp1 is copied to hcp2. The two gateways have an independent local cache, data written to hcp1 is available to hcp2 after a slight delay (one-way replication). The intent with this configuration is that Site B (hcp2) is a **read-only copy** of Site A, and also can be used to access data if Site A is offline.

Note:

It is the Administrators responsibility to ensure the Site B shares are configured as read-only.

Two Nodes: Master-Replica (Read-only Copy)



The gateway (hcp1) on Site A is the active system and is accessible to users. Upon ingest the file name, path and metadata is captured in the hcp1 database and made available to the Virtual File System. The file content is written to local cache, prior to being written to HCP. The caches on the 2 HCP Gateway devices are independent. After a slight delay for database hcp1 to replicate to hcp2, the file metadata written to hcp1 will become available on hcp2. The actual data needs to be replicated separately using HCP replication.

This configuration requires the **server.ignore** and **cluster** flags in the sam.properties file to be disabled. The latter is accomplished by setting the flags equal to zero (0). This change must be made on both hcp systems.

WARNING:

1. You must stop the SAM VFS service before making any changes to the sam.properties file. Then restart the SAM VFS service after saving the file. This will take all the shares offline and then bring them back online.
2. You must stop the Wildfly, SAM VFS and MariaDB services if any changes are made to the my.ini file. Then restart the MariaDB, SAM VFS and Wildfly services after saving the file. This will take all the shares offline and then bring them back online.

Step 1 – Update the HCP Gateway properties file for each node.

Locate the file at: C:\SAM\etc\sam\sam.properties

Make the following edits:

On HCPG 1

```
server.id=1
server.ignore=0
cluster=0
binlog.name=hcp-1-bin
```

On HCPG2

```
server.id=2
server.ignore=0
cluster=0
binlog.name=hcp-2-bin
```

Step 2 – Edit the my.ini file on both HCPGs

The database configuration must be setup to allow seamless interleaving of database records. This has an added benefit of preventing name collisions should the 'read-only' site have data written to it. This change must be made on both hcp systems, and, with the exact settings as shown for each hcp system. Note the entire my.ini file is not displayed in the instructions just the relevant part in the '[mariadb]' section of the respective my.ini file.

Edit File D:\MariaDB\data\my.ini

Note:

The following is added or changed in the '[mariadb]' section of the respective my.ini file.

Ensure the correct tables are excluded from being replicated by following the appropriate configuration section of the common list of excluded tables in Chapter 1.

You must stop the Wildfly, SAM VFS and MariaDB services if any changes are made to the my.ini file. Then restart the MariaDB, SAM VFS and Wildfly services after saving the file. This will take all the shares offline and then bring them back online.

HCPG1

```
[mariadb] . . .
# Global Transaction ID
gtid-domain-id=1
gtid-ignore-duplicates=ON
```

Unique Server ID

server-id = 1

Replication Configuration

auto-increment-offset = 1

auto-increment-increment = 2

ignore tables when replicate SAM database

See chapter 1 for appropriate tables to ignore

Binary Logging

log_bin = D:/MariaDB/binlog/hcpg-1-bin

log_bin_trust_function_creators=ON

Relay Logging

relay-log = D:/MariaDB/relaylog/hcpg-1-relay

HCPG2

[mariadb] . . .

Global Transaction ID

gtid-domain-id=2

gtid-ignore-duplicates=ON

Unique Server ID

server-id = 2

Replication Configuration

auto-increment-offset = 2

auto-increment-increment = 2

ignore tables when replicate SAM database

See chapter 1 for appropriate tables to ignore

Binary Logging

```
log_bin          = D:/MariaDB/binlog/hcpg-2-bin
```

```
log_bin_trust_function_creators=ON
```

```
# Relay Logging
```

```
relay-log       = D:/MariaDB/relaylog/hcpg-2-relay
```

Step 3 – Common Replication Setup Steps

It is highly recommended to perform the replication setup before any data is written to any of the HCP Gateway nodes. These steps need to be run on HCPG1 and HCPG2 with HCPG1 as the active node. If another node is the active node, then adjust the instructions for the active node.

In this configuration step a replication user account will be created on all nodes.

This step requires the use of the MariaDB Command Prompt, which is available from the Windows Start Menu.

```
MariaDB Command Prompt (MariaDB [(none)]>)
```

Run the command **mysql -uroot -p** on HCPG1 and HCPG2

```
C:\Windows\system32> mysql -uroot -p
```

Execute this query on each node to check if replication was previously configured and if so that there are no errors listed in the output. Resolve any errors before continuing to setup the replication.

```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

Execute these queries on HCPG1 and HCPG2, you only need to execute these queries the first time replication is setup. If replication is being reconfigured and the replication_user already exists, skip these 4 queries.

```
MariaDB [(none)]> CREATE USER 'replication_user'@'%' IDENTIFIED BY '0rgan1cReplication';
```

```
MariaDB [(none)]> GRANT REPLICATION SLAVE ON *.* TO 'replication_user'@'%';
```

```
MariaDB [(none)]> FLUSH PRIVILEGES;
```

```
MariaDB [(none)]> FLUSH TABLES;
```

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following queries on HCPG1 and HCPG2

```
MariaDB [(none)]> stop all slaves;
```

```
MariaDB [(none)]> reset slave all;
```

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query to display the replication status on HCPG1 and HCPG2

```
MariaDB [(none)]> show all slaves status\G
```

If there is any output showing replication status to any node, then on HCPG1 run the following query

```
MariaDB [(none)]> reset slave 'node2' all;
```

then on HCPG2 run the following query


```
MariaDB [(none)]> reset slave 'node1' all;
```

NOTE:

Please ignore the message **ERROR 1617 (HY000): There is no master connection 'node#'** as this query is ensuring the removal of any replication settings from node# and this message indicates there was no previous replication configuration for node#.


Open Windows Services and if running, stop the SAM VFS, Wildfly, and MariaDB services on HCPG2.

Delete or rename the D:\MariaDB\data*.err file on HCPG2, where the name of the file is the Windows computer name of the Gateway.



In Windows Services start the MariaDB service on HCPG2.

Check the D:\MariaDB\data*.err file in Notepad++ for any errors on HCPG2, where the name of the file is the Windows computer name of the Gateway. The expected output is shown below.



Step 4 – Collect the replication positions

HCPG1 (the active node)

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following queries **FLUSH TABLES WITH READ LOCK**; to temporarily lock the database tables and **RESET MASTER**; to reset the replication on **HCPG1**. Note that this will prevent access to the shares until the tables are unlocked immediately after issuing the **mysqldump** command below. The output or result from the **SHOW MASTER STATUS**; query will be used on HCPG2 in Step 5.

```
MariaDB [(none)]> FLUSH TABLES WITH READ LOCK;
```

```
MariaDB [(none)]> RESET MASTER;
```

```
MariaDB [(none)]> SHOW MASTER STATUS;
```

```

+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| hpcg-1-bin.000001 | 329 |          | mysql |
+-----+-----+-----+-----+

```

Collect the names of the tables that are ignored from replication.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query

```
MariaDB [(none)]> select group_concat(concat('--ignore-table=', TABLE_SCHEMA, '.', table_name)
SEPARATOR ' ') as " from information_schema.tables where TABLE_SCHEMA = 'SAM' and
(TABLE_NAME like '%event' or TABLE_NAME like 'license' or TABLE_NAME like 'report_status' or
TABLE_NAME like '%backfill' or TABLE_NAME like '%restoration' or TABLE_NAME like '%ntfs' or
TABLE_NAME like '%migration' or TABLE_NAME like 'archive_state' or TABLE_NAME like
'%purge');
```

```
| --ignore-table=SAM.license --ignore-table=SAM.archive_state --ignore-table=SAM.event --
ignore-table=SAM.report_status --ignore-table=SAM.1_backfill --ignore-table=SAM.1_event --
ignore-table=SAM.1_migration --ignore-table=SAM.1_purge --ignore-table=SAM.1_restoration
--ignore-table=SAM.3_migration --ignore-table=SAM.3_purge --ignore-table=SAM.3_event --
ignore-table=SAM.3_restoration --ignore-table=SAM.3_backfill --ignore-
table=SAM.5_migration --ignore-table=SAM.5_purge --ignore-table=SAM.5_event --ignore-
table=SAM.5_restoration --ignore-table=SAM.5_backfill |
```


1 row in set (0.001 sec)

Open a new MariaDB Command Prompt and change directory to C:\Temp. Dump the SAM database to a .sql file that will be copied to the other replication node. If the dump file is too big to fit on the C: drive, dump the SAM database to a .sql file in E:\Storage\Backup or another location that has enough free space.

```
C:\Windows\system32>cd \temp
```

Copy the output of the query above in **BOLD** (start with the `--ignore-table=SAM.license` and end with `SAM.5_backfill`) that collected the information about the tables to ignore from replication and paste it before the `--databases` and after the `--hex-blob=1` parameters in the query below. Make sure there is a space before `--ignore-table=SAM.license` and after `--ignore-table=SAM.5_backfill`.

```
C:\Temp>mysqldump -uroot -p --master-data --routines --hex-blob=1 --ignore-table=SAM.license --
ignore-table=SAM.1_migration --ignore-table=SAM.1_event --ignore-table=SAM.event --ignore-
table=SAM.report_status --ignore-table=SAM.1_restoration --ignore-table=SAM.3_migration --
ignore-table=SAM.3_event --ignore-table=SAM.1_backfill --ignore-table=SAM.3_restoration --ignore-
table=SAM.3_backfill --ignore-table=SAM.5_migration --ignore-table=SAM.5_event --ignore-
table=SAM.5_restoration --ignore-table=SAM.5_backfill --databases SAM >
hcp1.mysqldump.20221230.sql
```

Enter password: *****

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query **UNLOCK TABLES**; to unlock the database tables and restore access to the shares.

```
MariaDB [(none)]> UNLOCK TABLES;
```

HCPG2

WARNING: These commands will replace the database on HCPG2 with the data imported from the database on HCPG1. If there is production data on HCPG2 that was not replicated to HCPG1, then the production data on HCPG2 will be lost.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the queries **drop database SAM**; and **create database SAM**; on HCPG2 only

```
MariaDB [(none)]> drop database SAM;
```

```
MariaDB [(none)]> create database SAM;
```

Copy the `hcp1.mysqldump.20221230.sql` from HCPG1 to the C:\Temp folder on HCPG2. Open a new MariaDB Command Prompt and change directory to C:\Temp by issuing the command **cd \Temp**. Import the SAM database .sql file to copy the database from HCPG1 to HCPG2 by issuing the command **mysql -uroot -p < hcp1.mysqldump.20221230.sql**.

```
C:\Windows\system32>cd \temp
```

```
C:\Temp>mysql -uroot -p < hcp1.mysqldump.20221230.sql
```

Enter password: *****

Step 5 – Enable Replication

HCPG2

In the MariaDB Command Prompt, execute the following query **SET GLOBAL gtid_slave_pos = ''**; to reset the database replication gtid position:

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = '';
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG1. Note that the master_host IP address is for HCPG1.

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO
master_host="10.1.1.11", master_port=3306,
    master_user="replication_user",
    master_password='0rgan1cReplication',
    master_log_file="hcpg-1-bin.000001",
    master_log_pos=329, master_ssl=1;
```

Step 6 – Start replication on HCPG2

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> START ALL SLAVES;
```

```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

Verify that the replication is synchronized, with no errors, on all nodes. For example, the fields below show no errors and the Slave IO and Slave SQL are running.

Slave_SQL_State: Slave has read all relay log; waiting for the slave I/O thread to update it

Slave_IO_State: Waiting for master to send event

Slave_IO_Running: Yes

Slave_SQL_Running: Yes

Last_Errno: 0

Last_Error:

Last_IO_Errno: 0

Last_IO_Error:

Last_SQL_Errno: 0

Last_SQL_Error:

Slave_SQL_Running_State: Slave has read all relay log; waiting for the slave I/O thread to update it

Step 7 – Configure replication on HCPG2 to use Global Transaction ID (GTID)**HCPG2**

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node1';
```

```
MariaDB [(none)]> CHANGE MASTER "node1" TO master_use_gtid=slave_pos;
```

```
MariaDB [(none)]> START SLAVE 'node1';
```

```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

Verify that the replication is synchronized, with no errors, on all nodes.

Step 8 – Start the Windows Services Wildfly and SAM VFS on all nodes. Note that if this is the first time that replication is configured on HCPG2, in the Gateway UI Shares page, start all the shares.

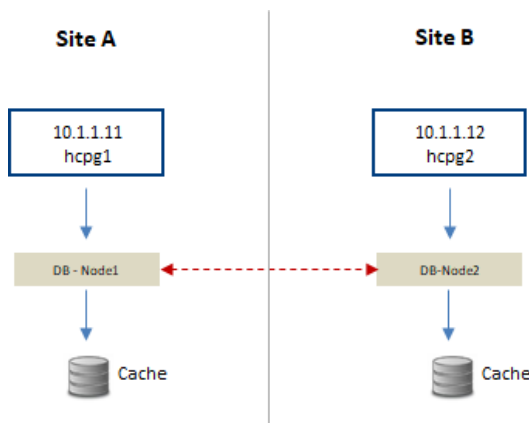
Chapter 3 Two Node: Master Single node to Master Single node Replication (Manual DR Failover)

The intent of this configuration is that Site B (hcpg2) is a hot-standby copy of Site A (hcpg1) and can be used to read/write data if Site A is offline. The two HCP Gateway devices do not have a shared cache. In addition, the expectation is the inactive site will have the data automatically updated from the active site without any intervention. This configuration does not support both sites being active at the same time. If this is replicated set of HCP Gateways without cluster and without a shared cache, it requires manual intervention to switch to the DR site, and to fail back from Site B to Site A.

NOTE:

Use this configuration also for a two node cluster with a shared cache. Follow the appropriate steps in this Chapter when using a two node cluster with a shared cache.

Two Nodes: Master-Master (DR Failover)



The configuration requires the **server.ignore** and **cluster** flags in the sam.properties file to be disabled. This is accomplished by setting the flags equal to zero (0). This change must be made on both hcpg systems.

WARNING:

1. You must stop the SAM VFS service before making any changes to the sam.properties file. Then restart the SAM VFS service after saving the file. This will take all the shares offline and then bring them back online.
2. You must stop the Wildfly, SAM VFS and MariaDB services if any changes are made to the my.ini file. Then restart the MariaDB, SAM VFS and Wildfly services after saving the file. This will take all the shares offline and then bring them back online.

IMPORTANT NOTE:

On a Gateway in a cluster, stop the cluster role that contains the SAM VFS service before making any changes to the sam.properties or the my.ini file. Then restart the cluster role that contains the SAM VFS service after saving the file. This will take all the shares offline and then bring them back online.

Step 1 – Update HCP Gateway properties file

Locate the properties file at: File C:\SAM\etc\sam\sam.properties

Make the following edits for two nodes that are NOT in a cluster:

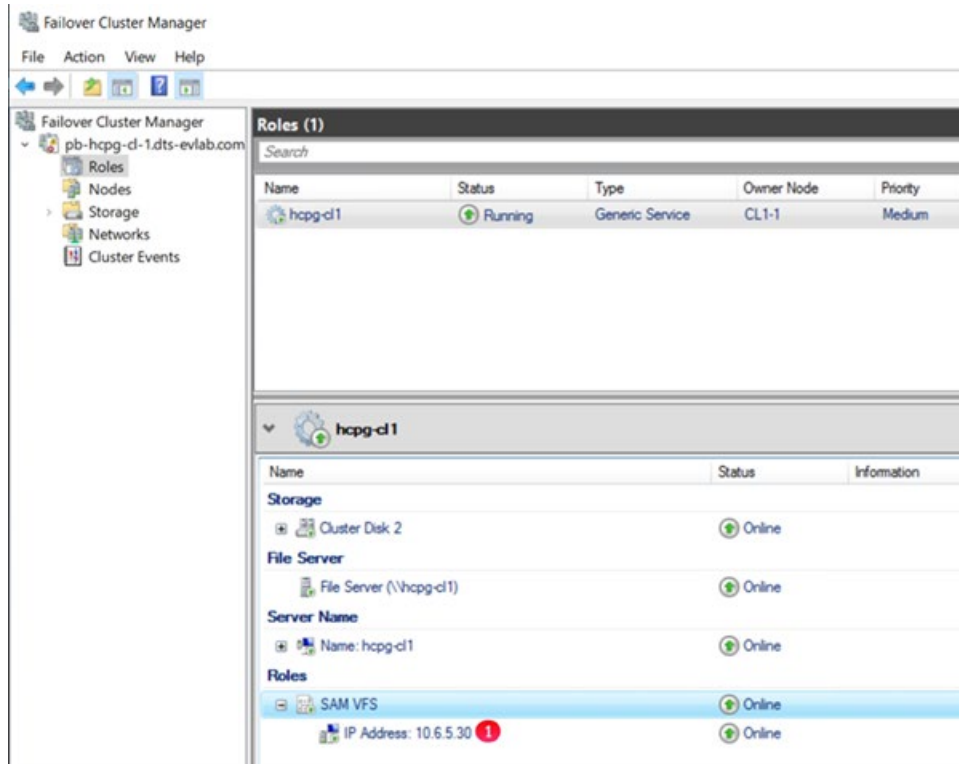
HCPG1

```
server.id=1
server.ignore=0
cluster=0
binlog.name=hcpg-1-bin
```

HCPG2

```
server.id=2
server.ignore=0
cluster=0
binlog.name=hcpg-2-bin
```

Make the following edits for two nodes IN a cluster with a shared cache. The cluster.access.ip=<clusterIPAddress>, where <clusterIPAddress> is the IP address of the Cluster Role that contains the SAM VFS service, for this example, 10.6.5.30 (1).



HCPG1

```
server.id=1
server.ignore=1
cluster=1
binlog.name=hcp-g-1-bin
cluster.access.ip=10.6.5.30
```

HCPG2

```
server.id=2
server.ignore=1
cluster=1
binlog.name=hcp-g-2-bin
cluster.access.ip=10.6.5.30
```

Step 2 – Edit the my.ini files on both HCPGs.

The database configuration must be setup to allow seamless interleaving of database records. This has an added benefit of preventing name collisions should the 'hot-standby' site have data written to it. This change must be made on both hcp-g systems, and, with the exact settings as shown for each hcp-g system.

Edit File D:\MariaDB\data\my.ini

Note:

The following is added or changed in the '[mariadb]' section of the respective my.ini file.

Ensure the correct tables are excluded from being replicated by following the appropriate configuration section of the common list of excluded tables in Chapter 1.

You must stop the Wildfly, SAM VFS and MariaDB services if any changes are made to the my.ini file. Then restart the MariaDB, SAM VFS and Wildfly services after saving the file. This will take all the shares offline and then bring them back online.

IMPORTANT NOTE:

On a Gateway in a cluster, stop the cluster role that contains the SAM VFS service before making any changes to the sam.properties or the my.ini file. Then restart the cluster role that contains the SAM VFS service after saving the file. This will take all the shares offline and then bring them back online.

HCPG1

[mariadb] . . .

Global Transaction ID

gtid-domain-id=1

gtid-ignore-duplicates=ON

Unique Server ID

server-id = 1

Replication Configuration

auto-increment-offset = 1

auto-increment-increment = 2

ignore tables when replicate SAM database

See chapter 1 for appropriate tables to ignore

Binary Logging

log_bin = D:/MariaDB/binlog/hcp-1-bin

log_bin_trust_function_creators=ON

```
# Relay Logging
relay-log          = D:/MariaDB/relaylog/hcpg-1-relay
```

HCPG2

```
[mariadb] . . .
```

```
# Global Transaction ID
gtid-domain-id=2
gtid-ignore-duplicates=ON
```

```
## Unique Server ID
server-id = 2
```

```
## Replication Configuration
auto-increment-offset = 2
auto-increment-increment = 2
```

```
# ignore tables when replicate SAM database
See chapter 1 for appropriate tables to ignore
```

```
## Binary Logging
log_bin          = D:/MariaDB/binlog/hcpg-2-bin
log_bin_trust_function_creators=ON
```

```
# Relay Logging
relay-log          = D:/MariaDB/relaylog/hcpg-2-relay
```

Step 3 – Common Replication Setup Steps

It is highly recommended to perform the replication setup before any data is written to any of the nodes in the HCP Gateway cluster or replicated set. These steps need to be run on HCPG1 and HCPG2 with HCPG1 as the active node. If another node is the active node, then adjust the instructions for the active node.

In this configuration step a replication user account will be created on all nodes;

This step requires the use of the MariaDB Command Prompt, which is available from the Windows Start Menu.

MariaDB Command Prompt (MariaDB [(none)]>)

Run the command **mysql -uroot -p** on HCPG1 and HCPG2

```
C:\Windows\system32> mysql -uroot -p
```


Execute this query on each node to check if replication was previously configured and if so that there are no errors listed in the output. Resolve any errors before continuing to setup the replication.

```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

Execute these queries on HCPG1 and HCPG2, you only need to execute these queries the first-time replication is setup. If replication is being reconfigured and the replication_user already exists, skip these 4 queries.

Note:

the password is the string immediately following 'IDENTIFIED BY'

```
MariaDB [(none)]> CREATE USER 'replication_user'@'%' IDENTIFIED BY '0rgan1cReplication';
```

```
MariaDB [(none)]> GRANT REPLICATION SLAVE ON *.* TO 'replication_user'@'%';
```

```
MariaDB [(none)]> FLUSH PRIVILEGES;
```

```
MariaDB [(none)]> FLUSH TABLES;
```

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following queries on HCPG1 and HCPG2

```
MariaDB [(none)]> stop all slaves;
```

```
MariaDB [(none)]> reset slave all;
```

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query to display the replication status on HCPG1 and HCPG2

```
MariaDB [(none)]> show all slaves status\G
```

If there is any output showing replication status to any node, then on HCPG1 run the following query

```
MariaDB [(none)]> reset slave 'node2' all;
```

then on HCPG2 run the following query


```
MariaDB [(none)]> reset slave 'node1' all;
```

NOTE:

Please ignore the message **ERROR 1617 (HY000): There is no master connection 'node#'** as this query is ensuring the removal of any replication settings from node# and this message indicates there was no previous replication configuration for node#.

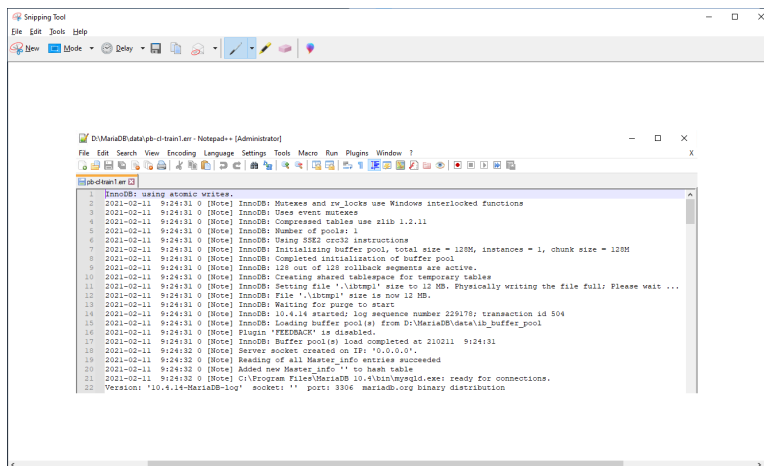
Open Windows Services and if running, stop the SAM VFS, Wildfly, and MariaDB services on HCPG2.

Delete or rename the D:\MariaDB\data*.err file on HCPG2, where the name of the file is the Windows computer name of the Gateway.



In Windows Services start the MariaDB service on HCPG2.

Check the D:\MariaDB\data*.err file in Notepad++ for any errors on HCPG2, where the name of the file is the Windows computer name of the Gateway. The expected output is shown below.



Step 4 – Collect the replication positions

HCPG1 (the active node)

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following queries **FLUSH TABLES WITH READ LOCK**; to temporarily lock the database tables and **RESET MASTER**; to reset the replication on **HCPG1**. Note that this will prevent access to the shares until the tables are unlocked

immediately after issuing the **mysqldump** command below. The output or result from the **SHOW MASTER STATUS**; query will be used on HCPG2 in Step 5.

```
MariaDB [(none)]> FLUSH TABLES WITH READ LOCK;
```

```
MariaDB [(none)]> RESET MASTER;
```

```
MariaDB [(none)]> SHOW MASTER STATUS;
```

```
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| hcpg-1-bin.000001 |      329 |              | mysql              |
+-----+-----+-----+-----+
```

Follow Option 1 below for a replicated pair of nodes that are **NOT** in a cluster and Option 2 for a replicated pair of nodes that **ARE** in a cluster.

Option 1) Follow these steps for a replicated pair of nodes that are NOT in a cluster

Collect the names of the tables that are ignored from replication.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query

```
MariaDB [(none)]> select group_concat(concat('--ignore-table=', TABLE_SCHEMA, '.', table_name)
SEPARATOR ' ') as " from information_schema.tables where TABLE_SCHEMA = 'SAM' and
(TABLE_NAME like '%event' or TABLE_NAME like 'license' or TABLE_NAME like 'report_status' or
TABLE_NAME like '%backfill' or TABLE_NAME like '%restoration' or TABLE_NAME like '%ntfs' or
TABLE_NAME like '%migration' or TABLE_NAME like 'archive_state' or TABLE_NAME like
'%purge');
```

```
| --ignore-table=SAM.license --ignore-table=SAM.archive_state --ignore-
table=SAM.1_migration --ignore-table=SAM.1_purge --ignore-table=SAM.1_event --ignore-
table=SAM.event --ignore-table=SAM.report_status --ignore-table=SAM.1_restoration --
ignore-table=SAM.3_migration --ignore-table=SAM.3_purge --ignore-table=SAM.3_event --
ignore-table=SAM.1_backfill --ignore-table=SAM.3_restoration --ignore-table=SAM.3_backfill -
ignore-table=SAM.5_migration --ignore-table=SAM.5_purge --ignore-table=SAM.5_event --
ignore-table=SAM.5_restoration --ignore-table=SAM.5_backfill |
```

```
1 row in set (0.001 sec)
```

Copy the output of the query above in **BOLD** (start with the **--ignore-table=SAM.license** and end with **SAM.5_backfill**) that collected the information about the tables to ignore from replication and paste it before the **--databases** and after the **--hex-blob=1** parameters in the query below. Make sure there is a space before **--ignore-table=SAM.license** and after **--ignore-table=SAM.5_backfill**.

Open a new MariaDB Command Prompt and change directory to C:\Temp by issuing the command **cd \Temp**. Dump the SAM database to a .sql file that will be copied to the other replication node. If the dump file is too big to fit on the C: drive, dump the SAM database to a .sql file in E:\Storage\Backup.

```
C:\Windows\system32>cd \temp
```

```
C:\Temp>mysqldump -uroot -p --master-data --routines --hex-blob=1 --ignore-table=SAM.license --
ignore-table=SAM.archive_state --ignore-table=SAM.1_migration --ignore-table=SAM.1_purge --
ignore-table=SAM.1_event --ignore-table=SAM.event --ignore-table=SAM.report_status --ignore-
table=SAM.1_restoration --ignore-table=SAM.3_migration --ignore-table=SAM.3_purge --ignore-
table=SAM.3_event --
```

```
ignore-table=SAM.1_backfill --ignore-table=SAM.3_restoration --ignore-table=SAM.3_backfill --
ignore-table=SAM.5_migration --ignore-table=SAM.5_purge --ignore-table=SAM.5_event --ignore-
table=SAM.5_restoration --ignore-table=SAM.5_backfill --databases SAM >
hpcg1.mysqldump.20221230.sql
```

Enter password: *****

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query **UNLOCK TABLES**; to unlock the database tables and restore access to the shares.

```
MariaDB [(none)]> UNLOCK TABLES;
```

Option 2) Follow these steps for a replicated pair of nodes that ARE in a cluster with a shared cache

Collect the names of the tables that are ignored from replication.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query

```
MariaDB [(none)]> select group_concat(concat('--ignore-table=', TABLE_SCHEMA, '.', table_name)
SEPARATOR ' ') as " from information_schema.tables where TABLE_SCHEMA = 'SAM' and
(TABLE_NAME like '%event' or TABLE_NAME like 'license')
```

```
| --ignore-table=SAM.license --ignore-table=SAM.1_event --ignore-table=SAM.event --ignore-
table=SAM.3_event --ignore-table=SAM.5_event |
```

1 row in set (0.001 sec)

Open a new MariaDB Command Prompt and change directory to C:\Temp by issuing the command **cd \Temp**. Dump the SAM database to a .sql file that will be copied to the other replication node. If the dump file is too big to fit on the C: drive, dump the SAM database to a .sql file in E:\Storage\Backup.

```
C:\Windows\system32>cd \temp
```

Copy the output of the command above in **BOLD** (start with the **--ignore-table=SAM.license** and end with **SAM.5_event**) that collected the information about the tables to ignore from replication and paste it before the **--databases** and after the **--hex-blob=1** parameters in the query below. Make sure there is a space before **--ignore-table=SAM.license** and after **--ignore-table=SAM.5_event**.

```
C:\Temp>mysqldump -uroot -p --master-data --routines --hex-blob=1 --ignore-table=SAM.license --
ignore-table=SAM.1_event --ignore-table=SAM.event --ignore-table=SAM.3_event --ignore-
table=SAM.5_event --databases SAM > hpcg1.mysqldump.20221230.sql
```

Enter password: *****

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query **UNLOCK TABLES**; to unlock the database tables and restore access to the shares.

MariaDB [(none)]> UNLOCK TABLES;

End of Option 2.

HCPG2

WARNING: These commands will replace the database on HCPG2 with the data imported from the database on HCPG1. If there is production data on HCPG2 that was not replicated to HCPG1, then the production data on HCPG2 will be lost.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the queries **drop database SAM;** and **create database SAM;** on HCPG2 only

MariaDB [(none)]> drop database SAM;

MariaDB [(none)]> create database SAM;

Copy the hcpg1.mysqlDump.20221230.sql from HCPG1 to the C:\Temp folder on HCPG2. Open a new MariaDB Command Prompt and change directory to C:\Temp. Import the SAM database .sql file to copy the database from HCPG1 to HCPG2.

C:\Windows\system32>cd \temp

C:\Temp>mysql -uroot -p < hcpg1.mysqlDump.20221230.sql

Enter password: *****

In the MariaDB Command Prompt, execute the following queries:

MariaDB [(none)]> RESET MASTER;

MariaDB [(none)]> SHOW MASTER STATUS;

```
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| hcpg-2-bin.000001 | 329 | | mysql |
+-----+-----+-----+-----+
```

Note:

The output or result from this query will be used in Step 5 on HCPG1.

Step 5 – Enable Replication

HCPG1

In the MariaDB Command Prompt, execute the following query **SET GLOBAL gtid_slave_pos = "";** to reset the database replication gtid position:

MariaDB [(none)]> SET GLOBAL gtid_slave_pos = "";

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG2. Note

that the master_host IP address is for HCPG2, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node2' TO master_host="10.1.1.12", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcp2g-2-bin.000001", master_log_pos=329, master_ssl=1;
```

HCPG2

In the MariaDB Command Prompt, execute the following query **SET GLOBAL gtid_slave_pos = ""** to reset the database replication gtid position:

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = "";
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG1. Note that the master_host IP address is for HCPG1, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcp1g-1-bin.000001", master_log_pos=329, master_ssl=1;
```

Step 6 – Start replication on both nodes HCPG1 and HCPG2

```
MariaDB [(none)]> START ALL SLAVES;
```

```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

Verify that the replication is synchronized, with no errors, on all nodes. For example, the fields below show no errors, and the Slave IO and Slave SQL are running.

Slave_SQL_State: Slave has read all relay log; waiting for the slave I/O thread to update it

Slave_IO_State: Waiting for master to send event

Slave_IO_Running: Yes

Slave_SQL_Running: Yes

Last_Errno: 0

Last_Error:

Last_IO_Errno: 0

Last_IO_Error:

Last_SQL_Errno: 0

Last_SQL_Error:

Slave_SQL_Running_State: Slave has read all relay log; waiting for the slave I/O thread to update it

Step 7 – Configure replication on both HCPG1 and HCPG2 to use Global Transaction ID (GTID)

HCPG1

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node2';
```

```
MariaDB [(none)]> CHANGE MASTER "node2" TO master_use_gtid=slave_pos;
```

```
MariaDB [(none)]> START SLAVE 'node2';
```

```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

HCPG2

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node1';
```

```
MariaDB [(none)]> CHANGE MASTER "node1" TO master_use_gtid=slave_pos;
```

```
MariaDB [(none)]> START SLAVE 'node1';
```

```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

Verify that the replication is synchronized, with no errors, on all nodes.

Step 8 – Start the Windows Services Wildfly on HCPG2.


Chapter 4 Three Node: Three Site with Master Single node to two Replica Single nodes Replication (Read-only copies)

In this configuration, the database on Site A HCP Gateway (hcpg1) is copied to Site B (hcpg2) and Site C (hcpg3). The three gateways each have an independent local cache, data written to hcpg1 is available to hcpg2 and hcpg3 after a slight delay (one-way replication). The intent with this configuration is that Site B (hcpg2) and Site C (hcpg3) are **read-only copies** of Site A and can be used to read/write data if Site A is offline.

Note:

It is the Administrators responsibility to ensure the Site B and Site C shares are configured as read-only.

Three Nodes: Three sites with Master and two Replicas (Read-only copies)



The gateway (hcpg1) on Site A is the active system and is accessible to users. Upon ingest the file name, path and metadata is captured in the hcpg1 database and made available to the Virtual File System. The file content is written to local cache, prior to being written to HCP. The caches on the 3 HCP Gateway devices are independent. After a slight delay for the database on hcpg1 to replicate to hcpg2 and hcpg3, the file metadata written to hcpg1 will become available on hcpg2 and hcpg3. The actual data needs to be replicated separately using HCP replication.

This configuration requires the **server.ignore** and **cluster** flags in the sam.properties file to be disabled. The latter is accomplished by setting the flags equal to zero (0). This change must be made on all 3 hcpG systems.

WARNING:

1. You must stop the SAM VFS service before making any changes to the sam.properties file. Then restart the SAM VFS service after saving the file. This will take all the shares offline and then bring them back online.
2. You must stop the Wildfly, SAM VFS and MariaDB services if any changes are made to the my.ini file. Then restart the MariaDB, SAM VFS and Wildfly services after saving the file. This will take all the shares offline and then bring them back online.

Step 1 – Update the HCP Gateway properties file for each node.

Locate the file at: C:\SAM\etc\sam\sam.properties

Make the following edits:

On HCPG 1

```
server.id=1
server.ignore=0
cluster=0
binlog.name=hcpG-1-bin
```

On HCPG2

```
server.id=2
server.ignore=0
cluster=0
binlog.name=hcpG-2-bin
```

On HCPG3

```
server.id=3
server.ignore=0
cluster=0
binlog.name=hcpG-3-bin
```

Step 2 – Edit the my.ini file on both HCPGs.

The database configuration must be setup to allow seamless interleaving of database records. This has an added benefit of preventing name collisions should the 'read-only' site have data written to it. This change must be made on all hcpG systems, and, with the exact settings as shown for each hcpG system. Note the entire my.ini file is not displayed in the instructions, just the relevant part in the '[mariadb]' section of the respective my.ini file.

Edit File D:\MariaDB\data\my.ini

Note:

The following is added or changed in the '[mariadb]' section of the respective my.ini file.

Ensure the correct tables are excluded from being replicated by following the appropriate configuration section of the common list of excluded tables in Chapter 1.

You must stop the Wildfly, SAM VFS and MariaDB services if any changes are made to the my.ini file. Then restart the MariaDB, SAM VFS and Wildfly services after saving the file. This will take all the shares offline and then bring them back online.

HCPG1

```
[mariadb] . . .
```

```
# Global Transaction ID
```

```
gtid-domain-id=1
```

```
gtid-ignore-duplicates=ON
```

```
## Unique Server ID
```

```
server-id = 1
```

```
## Replication Configuration
```

```
auto-increment-offset = 1
```

```
auto-increment-increment = 3
```

```
# ignore tables when replicate SAM database
```

```
See chapter 1 for appropriate tables to ignore
```

```
## Binary Logging
```

```
log_bin = D:/MariaDB/binlog/hcpg-1-bin
```

```
log_bin_trust_function_creators=ON
```

```
# Relay Logging
```

```
relay-log = D:/MariaDB/relaylog/hcpg-1-relay
```

HCPG2

```
[mariadb] . . .
```

```
# Global Transaction ID
```

```
gtid-domain-id=2
gtid-ignore-duplicates=ON

## Unique Server ID
server-id = 2

## Replication Configuration
auto-increment-offset = 2
auto-increment-increment = 3

# ignore tables when replicate SAM database
See chapter 1 for appropriate tables to ignore

## Binary Logging
log_bin          = D:/MariaDB/binlog/hcpg-2-bin
log_bin_trust_function_creators=ON

# Relay Logging
relay-log        = D:/MariaDB/relaylog/hcpg-2-relay
```

HCPG3

```
[mariadb] . . .
# Global Transaction ID
gtid-domain-id=3
gtid-ignore-duplicates=ON

## Unique Server ID
server-id = 3

## Replication Configuration
auto-increment-offset = 3
auto-increment-increment = 3

# ignore tables when replicate SAM database
See chapter 1 for appropriate tables to ignore
```

Binary Logging

log_bin = D:/MariaDB/binlog/hcpg-3-bin

log_bin_trust_function_creators=ON

Relay Logging

relay-log = D:/MariaDB/relaylog/hcpg-3-relay

Step 3 – Common Replication Setup Steps

It is highly recommended to perform the replication setup before any data is written to any of the HCP Gateway nodes. These steps need to be run on HCPG1, HCPG2 and HCPG3 with HCPG1 as the active node. If another node is the active node, then adjust the instructions for the active node.

In this configuration step a replication user account will be created on all nodes;

This step requires the use of the MariaDB Command Prompt, which is available from the Windows Start Menu.

MariaDB Command Prompt (MariaDB [(none)]>)

Run the command `mysql -uroot -p` on HCPG1, HCPG2 and HCPG3

C:\Windows\system32> `mysql -uroot -p`

Execute this query on each node to check if replication was previously configured and if so that there are no errors listed in the output. Resolve any errors before continuing to setup the replication.

MariaDB [(none)]> `SHOW ALL SLAVES STATUS\G`

Execute these queries on HCPG1, HCPG2 and HCPG3, you only need to execute these queries the first time replication is setup. If replication is being reconfigured and the `replication_user` already exists, skip these 4 queries.

Note:

the password is the string immediately following 'IDENTIFIED BY'

MariaDB [(none)]> `CREATE USER 'replication_user'@'%' IDENTIFIED BY '0rgan1cReplication';`

MariaDB [(none)]> `GRANT REPLICATION SLAVE ON *.* TO 'replication_user'@'%';`

MariaDB [(none)]> `FLUSH PRIVILEGES;`

MariaDB [(none)]> `FLUSH TABLES;`

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following queries on HCPG1, HCPG2 and HCPG3

MariaDB [(none)]> `stop all slaves;`

MariaDB [(none)]> `reset slave all;`

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query to display the replication status on HCPG1, HCPG2 and HCPG3

MariaDB [(none)]> `show all slaves status\G`

If there is any output showing replication status to any node, then on HCPG1 run the following queries

```
MariaDB [(none)]> reset slave 'node2' all;
```

```
MariaDB [(none)]> reset slave 'node3' all;
```

then on HCPG2 run the following queries

```
MariaDB [(none)]> reset slave 'node1' all;
```

```
MariaDB [(none)]> reset slave 'node3' all;
```

then on HCPG3 run the following queries

```
MariaDB [(none)]> reset slave 'node1' all;
```


```
MariaDB [(none)]> reset slave 'node2' all;
```

NOTE:

Please ignore the message **ERROR 1617 (HY000): There is no master connection 'node#'** as this query is ensuring the removal of any replication settings from node# and this message indicates there was no previous replication configuration for node#.

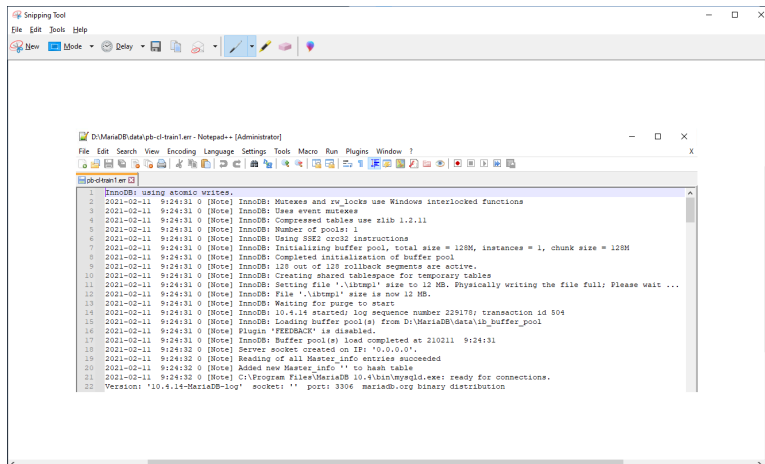
Open Windows Services and if running, stop the SAM VFS, Wildfly and MariaDB services on HCPG2 and HCPG3.

Delete or rename the D:\MariaDB\data*.err file on HCPG2 and HCPG3, where the name of the file is the Windows computer name of the Gateway.



In Windows Services start the MariaDB service on HCPG2 and HCPG3.

Check the D:\MariaDB\data*.err file in Notepad++ for any errors on HCPG2 and HCPG3, where the name of the file is the Windows computer name of the Gateway. The expected output is shown below.



Step 4 – Collect the replication positions

HCPG1 (the active node)

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following queries **FLUSH TABLES WITH READ LOCK**; to temporarily lock the database tables and **RESET MASTER**; to reset the replication on **HCPG1**. Note that this will prevent access to the shares until the tables are unlocked immediately after issuing the **mysqldump** command below. The output or result from the **SHOW MASTER STATUS**; query will be used on HCPG2 and HCPG3 in Step 5.

MariaDB [(none)]> FLUSH TABLES WITH READ LOCK;

MariaDB [(none)]> RESET MASTER;

MariaDB [(none)]> SHOW MASTER STATUS;

```

+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| hcpg-1-bin.000001 | 329     |              | mysql              |
+-----+-----+-----+-----+

```

Collect the names of the tables that are ignored from replication.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query

MariaDB [(none)]> select group_concat(concat('--ignore-table=', TABLE_SCHEMA, '.', table_name) SEPARATOR ' ') as " from information_schema.tables where TABLE_SCHEMA = 'SAM' and

(TABLE_NAME like '%event' or TABLE_NAME like 'license' or TABLE_NAME like 'report_status' or TABLE_NAME like '%backfill' or TABLE_NAME like '%restoration' or TABLE_NAME like '%ntfs' or TABLE_NAME like '%migration' or TABLE_NAME like 'archive_state' or TABLE_NAME like '%purge');

```
| --ignore-table=SAM.license --ignore-table=SAM.1_migration --ignore-table=SAM.1_purge --  
ignore-table=SAM.1_event --ignore-table=SAM.event --ignore-table=SAM.report_status --  
ignore-table=SAM.1_restoration --ignore-table=SAM.3_migration --ignore-table=SAM.3_purge --  
--ignore-table=SAM.3_event --ignore-table=SAM.1_backfill --ignore-table=SAM.3_restoration --  
ignore-table=SAM.3_backfill --ignore-table=SAM.5_migration --ignore-table=SAM.5_purge --  
ignore-table=SAM.5_event --ignore-table=SAM.5_restoration --ignore-table=SAM.5_backfill |
```

1 row in set (0.001 sec)

Open a new MariaDB Command Prompt and change directory to C:\Temp by issuing the command **cd \temp**. Dump the SAM database to a .sql file that will be copied to the other replication nodes. If the dump file is too big to fit on the C: drive, dump the SAM database to a .sql file in E:\Storage\Backup.

```
C:\Windows\system32>cd \temp
```

Copy the output of the query above in **BOLD** (start with the **--ignore-table=SAM.license** and end with **SAM.5_backfill**) that collected the information about the tables to ignore from replication and paste it before the **--databases** and after the **--hex-blob=1** parameters in the query below. Make sure there is a space before **--ignore-table=SAM.license** and after **--ignore-table=SAM.5_backfill**.

```
C:\Temp>mysqldump -uroot -p --master-data --routines --hex-blob=1 --ignore-table=SAM.license --  
ignore-table=SAM.1_migration --ignore-table=SAM.1_purge --ignore-table=SAM.1_event --ignore-  
table=SAM.event --ignore-table=SAM.report_status --ignore-table=SAM.1_restoration --ignore-  
table=SAM.3_migration --ignore-table=SAM.3_purge --ignore-table=SAM.3_event --ignore-  
table=SAM.1_backfill --ignore-table=SAM.3_restoration --ignore-table=SAM.3_backfill --ignore-  
table=SAM.5_migration --ignore-table=SAM.5_purge --ignore-table=SAM.5_event --ignore-  
table=SAM.5_restoration --ignore-table=SAM.5_backfill --databases SAM >  
hcp1.mysqldump.20221230.sql
```

Enter password: *****

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query **UNLOCK TABLES**; to unlock the database tables and restore access to the shares.

```
MariaDB [(none)]> UNLOCK TABLES;
```

HCPG2

WARNING: These commands will replace the database on HCPG2 with the data imported from the database on HCPG1. If there is production data on HCPG2 that was not replicated to HCPG1, then the production data on HCPG2 will be lost.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the queries **drop database SAM**; and **create database SAM**; on HCPG2 only

```
MariaDB [(none)]> drop database SAM;
```

```
MariaDB [(none)]> create database SAM;
```

Copy the hcp1.mysqldump.20221230.sql from HCPG1 to the C:\Temp folder on HCPG2. Open a new MariaDB Command Prompt and change directory to C:\Temp. Import the SAM database .sql file to copy the database from HCPG1 to HCPG2.

```
C:\Windows\system32>cd \temp
```

```
C:\Temp>mysql -uroot -p < hcp1.mysqlbackup.20221230.sql
```

```
Enter password: *****
```

HCPG3

WARNING: These commands will replace the database on HCPG3 with the data imported from the database on HCPG1. If there is production data on HCPG3 that was not replicated to HCPG1, then the production data on HCPG3 will be lost.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the queries drop database SAM; and create database SAM; on HCPG3 only

```
MariaDB [(none)]> drop database SAM;
```

```
MariaDB [(none)]> create database SAM;
```

Copy the hcp1.mysqlbackup.20221230.sql from HCPG1 to the C:\Temp folder on HCPG3. Open a new MariaDB Command Prompt and change directory to C:\Temp. Import the SAM database .sql file to copy the database from HCPG1 to HCPG3.

```
C:\Windows\system32>cd \temp
```

```
C:\Temp>mysql -uroot -p < hcp1.mysqlbackup.20221230.sql
```

```
Enter password: *****
```

Step 5 – Enable Replication

HCPG2

In the MariaDB Command Prompt, execute the following query **SET GLOBAL gtid_slave_pos = ""**; to reset the database replication gtid position:

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = "";
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG1. Note that the master_host IP address is for HCPG1, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11",  
master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcp1-1-bin.000001", master_log_pos=329, master_ssl=1;
```


HCPG3

In the MariaDB Command Prompt, execute the following query **SET GLOBAL gtid_slave_pos = ""**; to reset the database replication gtid position:

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = "";
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG1. Note that the master_host IP address is for HCPG1, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11",  
master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpg-1-bin.000001", master_log_pos=329, master_ssl=1;
```

Step 6 – Start replication on HCPG2 and HCPG3

```
MariaDB [(none)]> START ALL SLAVES;
```

```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

Verify that the replication is synchronized, with no errors, on all nodes. For example, the fields below show no errors, and the Slave IO and Slave SQL are running.

Slave_SQL_State: Slave has read all relay log; waiting for the slave I/O thread to update it

Slave_IO_State: Waiting for master to send event

Slave_IO_Running: Yes

Slave_SQL_Running: Yes

Last_Errno: 0

Last_Error:

Last_IO_Errno: 0

Last_IO_Error:

Last_SQL_Errno: 0

Last_SQL_Error:

Slave_SQL_Running_State: Slave has read all relay log; waiting for the slave I/O thread to update it

Step 7 – Configure replication on HCPG2 and HCPG3 to use Global Transaction ID (GTID)

HCPG2

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node1';  
MariaDB [(none)]> CHANGE MASTER "node1" TO master_use_gtid=slave_pos;  
MariaDB [(none)]> START SLAVE 'node1';  
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

HCPG3

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node1';  
MariaDB [(none)]> CHANGE MASTER "node1" TO master_use_gtid=slave_pos;  
MariaDB [(none)]> START SLAVE 'node1';  
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

Verify that the replication is synchronized, with no errors, on all nodes.

Step 8 – Start the Windows Services Wildfly and SAM VFS on HCPG2 and HCPG3. Start the read-only shares on HCPG2 and HCPG3. Note that if this is the first time that replication is configured on HCPG2 and HCPG3, in the Gateway UI Shares page, start all the shares.

Chapter 5 Three Node: Master to Master Replication for HA Cluster with optional Read-only Copy on secondary site

The two gateways on Site A are in a cluster and have a shared cache, only one gateway is active. The gateway in Site B has an independent cache. Metadata written to Site A database, will be available on Site B after a short delay. In this configuration:

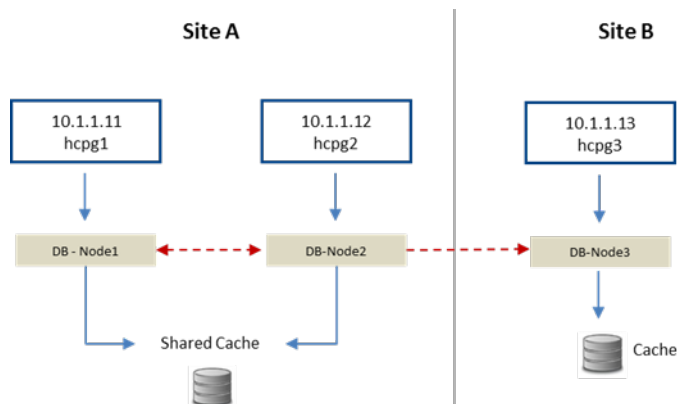
- Database on hcpg1 is replicated to hcpg2 and hcpg3 nodes
- Database on hcpg2 is replicated to hcpg1 and hcpg3 nodes

The intent for this configuration is that Site B (hcpg3) is a read-only copy of Site A and could be used to read/write data if Site A is offline. This configuration does not support both sites being active at the same time.

Note:

In the event that Site A is down, if data is written or modified on Site B, manual intervention will be required to sync that data back to Site A when it is brought back online.

Three Nodes: Master-Master HA Cluster and optional Read-only copy

**WARNING:**

1. You must stop the SAM VFS service before making any changes to the sam.properties file. Then restart the SAM VFS service after saving the file. This will take all the shares offline and then bring them back online.
2. You must stop the Wildfly, SAM VFS and MariaDB services if any changes are made to the my.ini file. Then restart the MariaDB, SAM VFS and Wildfly services after saving the file. This will take all the shares offline and then bring them back online.

IMPORTANT NOTE:

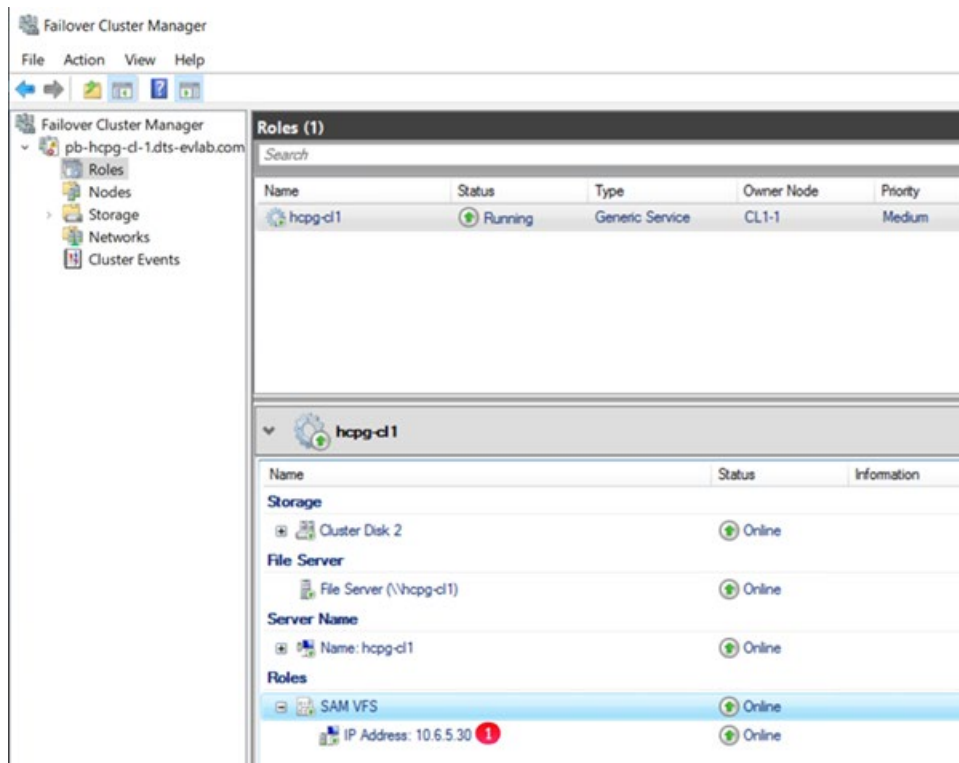
On a Gateway in a cluster, stop the cluster role that contains the SAM VFS service before making any changes to the sam.properties or the my.ini file. Then restart the cluster role that contains the SAM VFS service after saving the file. This will take all the shares offline and then bring them back online.

Step 1 – Update HCP Gateway properties file

The configuration requires the **server.ignore** and **cluster** flags in the sam.properties file to be enabled for hcp1, hcp2 and to be disabled for hcp3. This change must be made on all 3 hcp1 systems.

Edit the properties file at: File C:\SAM\etc\sam\sam.properties

Make the following edits. The cluster.access.ip=<clusterIPAddress>, where <clusterIPAddress> is the IP address of the Cluster Role that contains the SAM VFS service, for this example, 10.6.5.30 (1).



HCPG1

server.id=1

server.ignore=1

cluster=1

cluster.access.ip=<clusterIPAddress> (where <clusterIPAddress> is the IP address of the Cluster Role that contains the SAM VFS service)

binlog.name=hcp1-bin

HCPG2

server.id=2

server.ignore=1

cluster=1

cluster.access.ip=<clusterIPAddress> (where <clusterIPAddress> is the IP address of the Cluster Role that contains the SAM VFS service)

binlog.name=hcp2-bin

HCPG3

server.id=3

server.ignore=0

cluster=0

binlog.name=hcp3-bin

Step 2 – Edit the my.ini files on the three HCP Gateways

The database configuration must be setup to allow seamless interleaving of database records. This has an added benefit of preventing name collisions should the 'hot-standby' site have data written to it. This change must be made on all 3 hcp systems, and, with the exact settings as shown for each hcp system.

Edit File D:\MariaDB\data\my.ini

Note:

The following is added or changed in the '[mariadb]' section of the respective my.ini file.

Ensure the correct tables are excluded from being replicated by following the appropriate configuration section of the common list of excluded tables in Chapter 1.

You must stop the Wildfly, SAM VFS and MariaDB services if any changes are made to the my.ini file. Then restart the MariaDB, SAM VFS and Wildfly services after saving the file. This will take all the shares offline and then bring them back online.

HCPG1

[mariadb] . . .

Global Transaction ID

gtid-domain-id=1

gtid-ignore-duplicates=ON

Unique Server ID

```
server-id = 1
```

```
### Replication Configuration
```

```
auto-increment-offset = 1
```

```
auto-increment-increment = 3
```

```
# ignore tables when replicate SAM database
```

```
See chapter 1 for appropriate tables to ignore
```

```
### Binary Logging
```

```
log_bin = D:/MariaDB/binlog/hcpg-1-bin
```

```
log_bin_trust_function_creators=ON
```

```
# Relay Logging
```

```
relay-log = D:/MariaDB/relaylog/hcpg-1-relay
```

```
HCPG2
```

```
[mariadb] . . .
```

```
# Global Transaction ID
```

```
gtid-domain-id=2
```

```
gtid-ignore-duplicates=ON
```

```
### Unique Server ID
```

```
server-id = 2
```

```
### Replication Configuration
```

```
auto-increment-offset = 2
```

```
auto-increment-increment = 3
```

```
# ignore tables when replicate SAM database
```

```
See chapter 1 for appropriate tables to ignore
```

```
### Binary Logging
```

```
log_bin = D:/MariaDB/binlog/hcpg-2-bin
```

```
log_bin_trust_function_creators=ON
```

```
# Relay Logging
```

```
relay-log          = D:/MariaDB/relaylog/hcpg-2-relay
```

HCPG3

```
[mariadb] . . .
```

```
# Global Transaction ID
```

```
gtid-domain-id=3
```

```
gtid-ignore-duplicates=ON
```

```
## Unique Server ID
```

```
server-id = 3
```

```
## Replication Configuration
```

```
auto-increment-offset = 3
```

```
auto-increment-increment = 3
```

```
# ignore tables when replicate SAM database
```

```
See chapter 1 for appropriate tables to ignore
```

```
## Binary Logging
```

```
log_bin           = D:/MariaDB/binlog/hcpg-3-bin
```

```
log_bin_trust_function_creators=ON
```

```
# Relay Logging
```

```
relay-log         = D:/MariaDB/relaylog/hcpg-3-relay
```

Step 3 – Common Replication Setup Steps

It is highly recommended to perform the replication setup before any data is written to any of the nodes in the HCP Gateway. These steps need to be run on HCPG1, HCPG2 and HCPG3 with HCPG1 as the active node. If another node is the active node, then adjust the instructions for the active node.

In this configuration step a replication user account will be created on all nodes.

This step requires the use of the MariaDB Command Prompt, which is available from the Windows Start Menu.

```
MariaDB Command Prompt (MariaDB [(none)]>)
```

```
Run the command mysql -uroot -p on HCPG1, HCPG2 and HCPG3
```

```
C:\Windows\system32> mysql -uroot -p
```

Execute this query on each node to check if replication was previously configured and if so that there are no errors listed in the output. Resolve any errors before continuing to setup the replication.

```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

Execute these queries on HCPG1, HCPG2 and HCPG3, you only need to execute these queries the first time replication is setup. If replication is being reconfigured and the replication_user already exists, skip these 4 queries.

Note:

the password is the string immediately following 'IDENTIFIED BY'

```
MariaDB [(none)]> CREATE USER 'replication_user'@'%' IDENTIFIED BY '0rgan1cReplication';
```

```
MariaDB [(none)]> GRANT REPLICATION SLAVE ON *.* TO 'replication_user'@'%';
```

```
MariaDB [(none)]> FLUSH PRIVILEGES;
```

```
MariaDB [(none)]> FLUSH TABLES;
```

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following queries on HCPG1, HCPG2 and HCPG3

```
MariaDB [(none)]> stop all slaves;
```

```
MariaDB [(none)]> reset slave all;
```

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query to display the replication status on HCPG1, HCPG2 and HCPG3

```
MariaDB [(none)]> show all slaves status\G
```

If there is any output showing replication status to any node, then on HCPG1 run the following queries

```
MariaDB [(none)]> reset slave 'node2' all;
```

```
MariaDB [(none)]> reset slave 'node3' all;
```

then on HCPG2 run the following queries

```
MariaDB [(none)]> reset slave 'node1' all;
```

```
MariaDB [(none)]> reset slave 'node3' all;
```

then on HCPG3 run the following queries

```
MariaDB [(none)]> reset slave 'node1' all;
```


```
MariaDB [(none)]> reset slave 'node2' all;
```

NOTE:

Please ignore the message **ERROR 1617 (HY000): There is no master connection 'node#'** as this query is ensuring the removal of any replication settings from node# and this message indicates there was no previous replication configuration for node#.

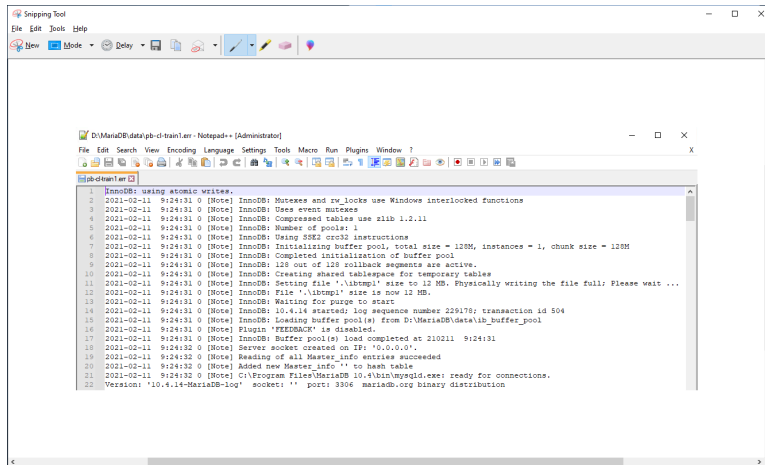
Open Windows Services and if running, stop the SAM VFS, Wildfly, and MariaDB services on HCPG2 and HCPG3.

Delete or rename the D:\MariaDB\data*.err file on HCPG2 and HCPG3, where the name of the file is the Windows computer name of the Gateway.



In Windows Services start the MariaDB service on HCPG2 and HCPG3.

Check the D:\MariaDB\data*.err file in Notepad++ for any errors on HCPG2 and HCPG3, where the name of the file is the Windows computer name of the Gateway. The expected output is shown below.



Step 4 – Collect the replication positions

HCPG1 (the active node)

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following queries **FLUSH TABLES WITH READ LOCK**; to temporarily lock the database tables and **RESET MASTER**; to reset the replication on **HCPG1**. Note that this will prevent access to the shares until the tables are unlocked immediately after issuing the **mysqldump** command below. The output or result from the **SHOW MASTER STATUS**; query will be used on HCPG2 and HCPG3 in Step 5.

```
MariaDB [(none)]> FLUSH TABLES WITH READ LOCK;
```

```
MariaDB [(none)]> RESET MASTER;
```

```
MariaDB [(none)]> SHOW MASTER STATUS;
```

```

+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| hcpg-1-bin.000001 | 329     |              | mysql              |
+-----+-----+-----+-----+

```

Backup database on HCPG1 and this will be copied to and imported on HCPG2

Collect the names of the tables that are ignored from replication.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query

```
MariaDB [(none)]> select group_concat(concat('--ignore-table=', TABLE_SCHEMA, '.', table_name)
SEPARATOR ' ') as " from information_schema.tables where TABLE_SCHEMA = 'SAM' and
(TABLE_NAME like '%event' or TABLE_NAME like 'license');
```

```
| --ignore-table=sam.1_event --ignore-table=sam.event --ignore-table=sam.license |
```

```
1 row in set (0.001 sec)
```

Open a new MariaDB Command Prompt and change directory to C:\Temp by issuing the command **cd %temp%**. Dump the SAM database to a .sql file that will be copied to the other replication node. If

the dump file is too big to fit on the C: drive, dump the SAM database to a .sql file in E:\Storage\Backup.

```
C:\Windows\system32>cd \temp
```

Copy the output of the query above in **BOLD** (start with the `--ignore-table=sam.1_event` and end with `SAM.license`) that collected the information about the tables to ignore from replication and paste it before the `--databases` and after the `--hex-blob=1` parameters in the query below. Make sure there is a space before `--ignore-table=sam.1_event` and after `--ignore-table=sam.license`.

```
C:\Temp>mysqldump -uroot -p --master-data --routines --hex-blob=1 --ignore-table=sam.1_event --
ignore-table=sam.event --ignore-table=sam.license --databases SAM >
hcpg1.mysqldump.20221230-hcp2.sql
```

Enter password: *****

Backup database on HCPG1 and this will be copied to and imported on HCPG3

Collect the names of the tables that are ignored from replication.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query

```
MariaDB [(none)]> select group_concat(concat('--ignore-table=', TABLE_SCHEMA, '.', table_name)
SEPARATOR ' ') as " from information_schema.tables where TABLE_SCHEMA = 'SAM' and
(TABLE_NAME like '%event' or TABLE_NAME like 'license' or TABLE_NAME like 'report_status' or
TABLE_NAME like '%backfill' or TABLE_NAME like '%restoration' or TABLE_NAME like '%ntfs' or
TABLE_NAME like '%migration' or TABLE_NAME like '%archive_state' or TABLE_NAME like
'%purge');
```

```
| --ignore-table=sam.1_backfill --ignore-table=sam.1_event --ignore-table=sam.1_migration --
ignore-table=sam.1_ntfs --ignore-table=sam.1_purge --ignore-table=sam.1_restoration --
ignore-table=sam.archive_state --ignore-table=sam.event --ignore-table=sam.license --ignore-
table=sam.report_status |
```

1 row in set (0.001 sec)

If needed, open a MariaDB Command Prompt

```
C:\Windows\system32>cd \temp
```

Copy the output of the command above in **BOLD** (start with the `--ignore-table=sam.1_backfill` and end with `sam.report_status`) that collected the information about the tables to ignore from replication and paste it before the `--databases` and after the `--hex-blob=1` parameters in the query below. Make sure there is a space before `--ignore-table=sam.1_backfill` and after `--ignore-table=sam.report_status`.

```
C:\Temp>mysqldump -uroot -p --master-data --routines --hex-blob=1 --ignore-table=sam.1_backfill
--ignore-table=sam.1_event --ignore-table=sam.1_migration --ignore-table=sam.1_ntfs --
ignore-table=sam.1_purge --ignore-table=sam.1_restoration --ignore-table=sam.archive_state
--ignore-table=sam.event --ignore-table=sam.license --ignore-table=sam.report_status --
databases SAM > hcpg1.mysqldump.20221230-hcp3.sql
```

Enter password: *****

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query **UNLOCK TABLES**; to unlock the database tables and restore access to the shares.

```
MariaDB [(none)]> UNLOCK TABLES;
```

HCPG2

WARNING: These commands will replace the database on HCPG2 with the data imported from the database on HCPG1. If there is production data on HCPG2 that was not replicated to HCPG1, then the production data on HCPG2 will be lost.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the queries **drop database SAM;** and **create database SAM;** on HCPG2 only

```
MariaDB [(none)]> drop database SAM;
```

```
MariaDB [(none)]> create database SAM;
```

Copy the hcp1.mysqlDump.20221230-hcp2.sql file from HCPG1 to the C:\Temp folder on HCPG2. Open a new MariaDB Command Prompt and change directory to C:\Temp. Import the SAM database .sql file to copy the database from HCPG1 to HCPG2.

```
C:\Windows\system32>cd \temp
```

```
C:\Temp>mysql -uroot -p < hcp1.mysqlDump.20221230-hcp2.sql
```

```
Enter password: *****
```

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> RESET MASTER;
```

```
MariaDB [(none)]> SHOW MASTER STATUS;
```

```
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| hcp2-bin.000001 |      329 |              | mysql              |
+-----+-----+-----+-----+
```

Note:

The output or result from the show master status query will be used on HCPG1 and HCPG3 in Step 5.

HCPG3

WARNING: These commands will replace the database on HCPG3 with the data imported from the database on HCPG1. If there is production data on HCPG3 that was not replicated to HCPG1, then the production data on HCPG3 will be lost.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the queries **drop database SAM;** and **create database SAM;** on HCPG3 only

```
MariaDB [(none)]> drop database SAM;
```

```
MariaDB [(none)]> create database SAM;
```

Copy the hcp1.mysqlDump.20221230-hcp3.sql file from HCPG1 to the C:\Temp folder on HCPG3. Open a new MariaDB Command Prompt and change directory to C:\Temp. Import the SAM database .sql file to copy the database from HCPG1 to HCPG3.

```
C:\Windows\system32>cd \temp
```

```
C:\Temp>mysql -uroot -p < hcp1.mysqlbackup.20221230-hcp3.sql
```

```
Enter password: *****
```

Step 5 – Enable replication

HCPG1

In the MariaDB Command Prompt, execute the following query to unset the GLOBAL `gtid_slave_pos` parameter:

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = "";
```

In the MariaDB Command Prompt, execute the following query:

Note:

The `master_log_file` and `master_log_pos` information is the result from the `show master status` query on HCPG2. Note that the `master_host` IP address is for HCPG2, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node2' TO master_host="10.1.1.12", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpg-2-bin.000001", master_log_pos=329, master_ssl=1;
```

HCPG2

In the MariaDB Command Prompt, execute the following query to unset the GLOBAL `gtid_slave_pos` parameter:

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = "";
```

In the MariaDB Command Prompt, execute the following query:

Note:

The `master_log_file` and `master_log_pos` information is the result from the `show master status` query on HCPG1. Note that the `master_host` IP address is for HCPG1, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpg-1-bin.000001", master_log_pos=329, master_ssl=1;
```

HCPG3

In the MariaDB Command Prompt, execute the following query to unset the GLOBAL `gtid_slave_pos` parameter:

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = "";
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG1. Note that the master_host IP address is for HCPG1, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpg-1-bin.000001", master_log_pos=329, master_ssl=1;
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG2. Note that the master_host IP address is for HCPG2, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node2' TO master_host="10.1.1.12", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpg-2-bin.000001", master_log_pos=329, master_ssl=1;
```

Step 6 – Start replication

HCPG1

```
MariaDB [(none)]> START ALL SLAVES;  
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

HCPG2

```
MariaDB [(none)]> START ALL SLAVES;  
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

HCPG3

```
MariaDB [(none)]> START ALL SLAVES;  
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

Verify that the replication is synchronized, with no errors, on all nodes. For example, the fields below show no errors, and the Slave IO and Slave SQL are running.

Slave_SQL_State: Slave has read all relay log; waiting for the slave I/O thread to update it
Slave_IO_State: Waiting for master to send event
Slave_IO_Running: Yes
Slave_SQL_Running: Yes
Last_Errno: 0
Last_Error:
Last_IO_Errno: 0
Last_IO_Error:
Last_SQL_Errno: 0
Last_SQL_Error:
Slave_SQL_Running_State: Slave has read all relay log; waiting for the slave I/O thread to update it

Step 7 – Configure replication on HCPG1, HCPG2 and HCPG3 to use Global Transaction ID (GTID)

HCPG1

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node2';  
MariaDB [(none)]> CHANGE MASTER "node2" TO master_use_gtid=slave_pos;  
MariaDB [(none)]> START SLAVE 'node2';  
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

HCPG2

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node1';  
MariaDB [(none)]> CHANGE MASTER "node1" TO master_use_gtid=slave_pos;  
MariaDB [(none)]> START SLAVE 'node1';  
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

HCPG3

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node1';  
MariaDB [(none)]> CHANGE MASTER "node1" TO master_use_gtid=slave_pos;  
MariaDB [(none)]> START SLAVE 'node1';  
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node2';  
MariaDB [(none)]> CHANGE MASTER "node2" TO master_use_gtid=slave_pos;  
MariaDB [(none)]> START SLAVE 'node2';  
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

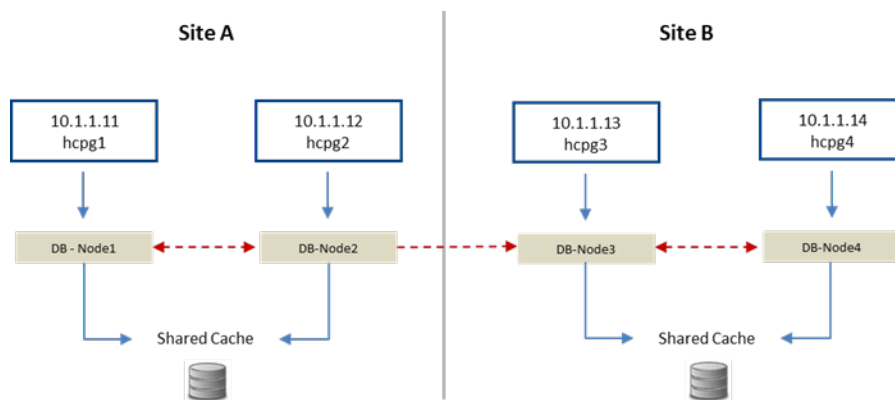
Verify that the replication is synchronized, with no errors, on all nodes.

Step 8 – Start the Windows Services Wildfly service on all nodes and then start the **SAM VFS service** and then the read-only shares in the HCP Gateway UI on **HCPG3**. **Note that if this is the first time that replication is configured on HCPG3, in the Gateway UI Shares page, start all the shares.**

Chapter 6 Four Nodes: Master to Master Replication for HA Cluster with DR Failover to HA Cluster

The intent for this configuration is that Site B is a hot-standby copy of Site A and can be used to read/write data if Site A is offline. This configuration does not support both sites being active at the same time. Manual intervention is required to fail back from Site B to Site A. In this configuration, the databases on all gateways are replicated to each other. The two gateways on Site A are in a cluster and have a shared cache, only one gateway is active. The two gateways on Site B are in a cluster and have a shared cache, but no gateways are active in Site B, only one gateway in Site A is active. Data written to Site A, will be available on Site B after a short delay.

Four Nodes: Master-Master HA Cluster on both sites



WARNING:

1. You must stop the SAM VFS service before making any changes to the sam.properties file. Then restart the SAM VFS service after saving the file. This will take all the shares offline and then bring them back online.
2. You must stop the Wildfly, SAM VFS and MariaDB services if any changes are made to the my.ini file. Then restart the MariaDB, SAM VFS and Wildfly services after saving the file. This will take all the shares offline and then bring them back online.

IMPORTANT NOTE:

On a Gateway in a cluster, stop the cluster role that contains the SAM VFS service before making any changes to the sam.properties or the my.ini file. Then restart the cluster role that contains the SAM VFS service after saving the file. This will take all the shares offline and then bring them back online.

Step 1 - Update HCP Gateway properties file

The configuration requires the **server.ignore** and **cluster** flags in the sam.properties file to be enabled on hcpg1, hcpg2, hcpg3, and hcpg4. This change must be made on all 4 hcpg systems.

Note:

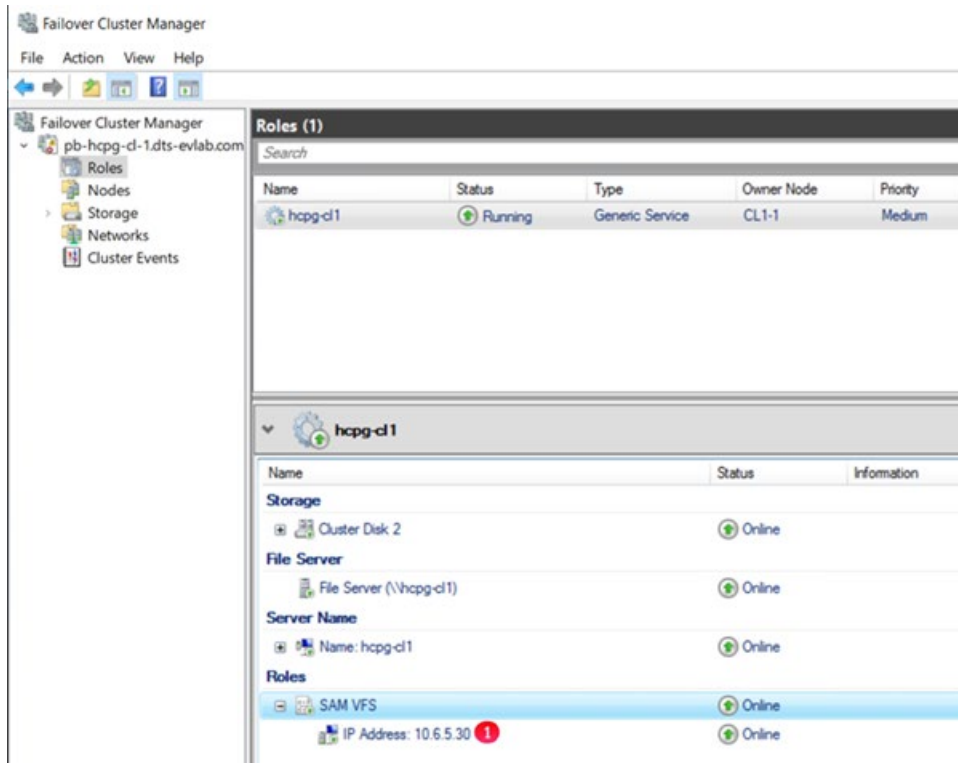
You must stop and restart the SAM VFS service if any changes are made.

IMPORTANT NOTE:

On a Gateway in a cluster, on the active node, stop the cluster role that contains the SAM VFS service before making any changes to the C:\SAM\etc\sam\sam.properties file. Then, on the active node, restart the cluster role that contains the SAM VFS service after saving the file. This will take all the shares offline and then bring them back online.

Edit properties file at: C:\SAM\etc\sam\sam.properties

Verify the parameters match the lines below and if not, adjust the parameters. The cluster.access.ip=<clusterIPAddress>, where <clusterIPAddress> is the IP address of the Cluster Role that contains the SAM VFS service, for this example, 10.6.5.30 (1).

**HCPG1**

server.id=1

server.ignore=1

```
cluster=1
cluster.access.ip=<SiteAclusterIPAddress>
binlog.name=hcpG-1-bin
```

HCPG2

```
server.id=2
server.ignore=1
cluster=1
cluster.access.ip=<SiteAclusterIPAddress>
```

```
binlog.name=hcpG-2-bin
```

HCPG3

```
server.id=3
server.ignore=1
cluster=1
cluster.access.ip=<SiteBclusterIPAddress>
```

```
binlog.name=hcpG-3-bin
```

HCPG4

```
server.id=4
server.ignore=1
cluster=1
cluster.access.ip=<SiteBclusterIPAddress>
```

```
binlog.name=hcpG-4-bin
```

Step 2 - Edit the my.ini files on the four HCP Gateways

The database configuration must be setup to allow seamless interleaving of database records. This change must be made on all four hcpG systems, and, with the exact settings as shown for each hcpG system.

Note:

The following is added or changed in the '[mariadb]' section of the respective my.ini file.

Ensure the correct tables are excluded from being replicated by following the appropriate configuration section of the common list of excluded tables in Chapter 1.

You must stop the Wildfly, SAM VFS and MariaDB services if any changes are made to the my.ini file. Then restart the MariaDB, SAM VFS and Wildfly services after saving the file. This will take all the shares offline and then bring them back online.

IMPORTANT NOTE:

On a Gateway in a cluster, stop the cluster role that contains the SAM VFS service before making any changes to the sam.properties or the my.ini file. Then restart the cluster role that contains the SAM VFS service after saving the file. This will take all the shares offline and then bring them back online.

Edit File D:\MariaDB\data\my.ini

HCPG1

[mariadb] . . .

Global Transaction ID

gtid-domain-id=1

gtid-ignore-duplicates=ON

Unique Server ID

server-id = 1

Replication Configuration

auto-increment-offset = 1

auto-increment-increment = 4

ignore tables when replicate SAM database

See chapter 1 for appropriate tables to ignore.

Binary Logging

log_bin = D:/MariaDB/binlog/hcpg-1-bin

log_bin_trust_function_creators=ON

Relay Logging

```
relay-log          = D:/MariaDB/relaylog/hcpg-1-relay
```

HCPG2

```
[mariadb] . . .
```

```
# Global Transaction ID
```

```
gtid-domain-id=2
```

```
gtid-ignore-duplicates=ON
```

```
## Unique Server ID
```

```
server-id = 2
```

```
## Replication Configuration
```

```
auto-increment-offset = 2
```

```
auto-increment-increment = 4
```

```
# ignore tables when replicate SAM database
```

```
See chapter 1 for appropriate tables to ignore.
```

```
## Binary Logging
```

```
log_bin           = D:/MariaDB/binlog/hcpg-2-bin
```

```
log_bin_trust_function_creators=ON
```

```
# Relay Logging
```

```
relay-log         = D:/MariaDB/relaylog/hcpg-2-relay
```

HCPG3

```
[mariadb] . . .
```

```
# Global Transaction ID
```

```
gtid-domain-id=3
```

```
gtid-ignore-duplicates=ON
```

```
## Unique Server ID
```

```
server-id = 3
```

```
## Replication Configuration
```

```
auto-increment-offset = 3
```

```
auto-increment-increment = 4
# ignore tables when replicate SAM database
See chapter 1 for appropriate tables to ignore.
```

Binary Logging

```
log_bin          = D:/MariaDB/binlog/hcpg-3-bin
log_bin_trust_function_creators=ON
```

Relay Logging

```
relay-log        = D:/MariaDB/relaylog/hcpg-3-relay
```

HCPG4

```
[mariadb] . . .
```

```
# Global Transaction ID
```

```
gtid-domain-id=4
```

```
gtid-ignore-duplicates=ON
```

Unique Server ID

```
server-id = 4
```

Replication Configuration

```
auto-increment-offset = 4
```

```
auto-increment-increment = 4
```

```
# ignore tables when replicate SAM database
See chapter 1 for appropriate tables to ignore.
```

Binary Logging

```
log_bin          = D:/MariaDB/binlog/hcpg-4-bin
```

```
log_bin_trust_function_creators=ON
```

Relay Logging

```
relay-log        = D:/MariaDB/relaylog/hcpg-4-relay
```

Step 3 – Common Replication Setup Steps

It is highly recommended to perform the replication setup before any data is written to any of the nodes in the HCP Gateway. These steps need to be run on HCPG1, HCPG2, HCPG3 and HCPG4 with HCPG1 as the active node. If another node is the active node, then adjust the instructions for the active node.

In this configuration step a replication user account will be created on all nodes;

This step requires the use of the MariaDB Command Prompt, which is available from the Windows Start Menu.

MariaDB Command Prompt (MariaDB [(none)]>)

Run the command **mysql -uroot -p** on HCPG1, HCPG2, HCPG3 and HCPG4

C:\Windows\system32> mysql -uroot -p

Execute this query on each node to check if replication was previously configured and if so that there are no errors listed in the output. Resolve any errors before continuing to setup the replication.

MariaDB [(none)]> SHOW ALL SLAVES STATUS\G

Execute these queries on HCPG1, HCPG2, HCPG3 and HCPG4, you only need to execute these queries the first time replication is setup. If replication is being reconfigured and the replication_user already exists, skip these 4 queries.

Note:

the password is the string immediately following 'IDENTIFIED BY'

MariaDB [(none)]> CREATE USER 'replication_user'@'%' IDENTIFIED BY '0rgan1cReplication';

MariaDB [(none)]> GRANT REPLICATION SLAVE ON *.* TO 'replication_user'@'%';

MariaDB [(none)]> FLUSH PRIVILEGES;

MariaDB [(none)]> FLUSH TABLES;

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following queries on HCPG1, HCPG2, HCPG3 and HCPG4

MariaDB [(none)]> stop all slaves;

MariaDB [(none)]> reset slave all;

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query to display the replication status on HCPG1, HCPG2, HCPG3 and HCPG4

MariaDB [(none)]> show all slaves status\G

If there is any output showing replication status to any node, then on HCPG1 run the following queries.

NOTE:

Please ignore the message **ERROR 1617 (HY000): There is no master connection 'node#'** as this query is ensuring the removal of any replication settings from node# and this message indicates there was no previous replication configuration for node#.

```
MariaDB [(none)]> reset slave 'node2' all;  
MariaDB [(none)]> reset slave 'node3' all;  
MariaDB [(none)]> reset slave 'node4' all;
```

then on HCPG2 run the following queries

```
MariaDB [(none)]> reset slave 'node1' all;  
MariaDB [(none)]> reset slave 'node3' all;  
MariaDB [(none)]> reset slave 'node4' all;
```

then on HCPG3 run the following queries

```
MariaDB [(none)]> reset slave 'node1' all;  
MariaDB [(none)]> reset slave 'node2' all;  
MariaDB [(none)]> reset slave 'node4' all;
```

then on HCPG4 run the following queries


```
MariaDB [(none)]> reset slave 'node1' all;  
MariaDB [(none)]> reset slave 'node2' all;  
MariaDB [(none)]> reset slave 'node3' all;
```

NOTE:

Please ignore the message **ERROR 1617 (HY000): There is no master connection 'node#'** as this query is ensuring the removal of any replication settings from node# and this message indicates there was no previous replication configuration for node#.

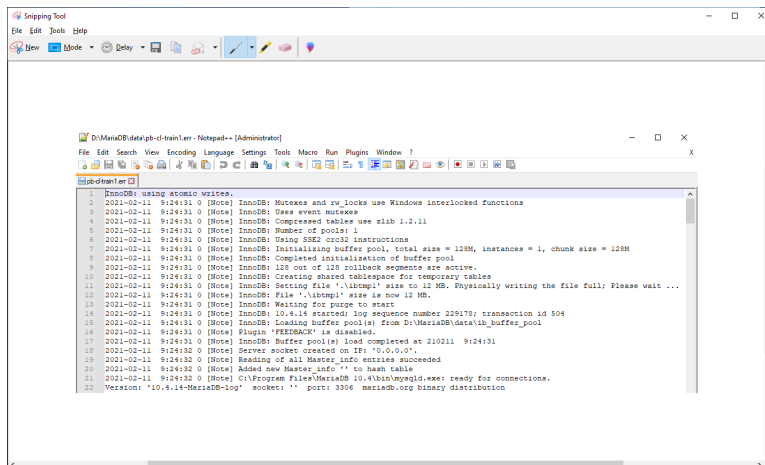
Open Windows Services and if running, stop the SAM VFS (stop the cluster role containing the SAM VFS service), Wildfly, and MariaDB services on HCPG2, HCPG3 and HCPG4.

Delete or rename the D:\MariaDB\data*.err file on HCPG2, HCPG3 and HCPG4, where the name of the file is the Windows computer name of the Gateway.



In Windows Services start the MariaDB service on HCPG2, HCPG3 and HCPG4.

Check the D:\MariaDB\data*.err file in Notepad++ for any errors on HCPG2, HCPG3 and HCPG4, where the name of the file is the Windows computer name of the Gateway. The expected output is shown below.



Step 4 – Collect the replication positions

HCPG1 (the active node)

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following queries **FLUSH TABLES WITH READ LOCK**; to temporarily lock the database tables and **RESET MASTER**; to reset the replication on **HCPG1**. Note that this will prevent access to the shares until the tables are unlocked immediately after issuing the **mysqldump** command below. The output or result from the **SHOW MASTER STATUS**; query will be used on HCPG2, HCPG3 and HCPG4 in Step 5.

```
MariaDB [(none)]> FLUSH TABLES WITH READ LOCK;
```

```
MariaDB [(none)]> RESET MASTER;
```

```
MariaDB [(none)]> SHOW MASTER STATUS;
```

```
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| hcpg-1-bin.000001 |      329 |              | mysql              |
+-----+-----+-----+-----+
```

Collect the names of the tables that are ignored from replication.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query

```
MariaDB [(none)]> select group_concat(concat('--ignore-table=', TABLE_SCHEMA, '.', table_name)
SEPARATOR ' ') as " from information_schema.tables where TABLE_SCHEMA = 'SAM' and
(TABLE_NAME like '%event' or TABLE_NAME like 'license');
```

```
| --ignore-table=sam.1_event --ignore-table=sam.event --ignore-table=sam.license |
```

```
1 row in set (0.001 sec)
```

Backup database on HCPG1 and it will be copied to and imported on HCPG2

Open a new MariaDB Command Prompt and change directory to C:\Temp by issuing the command **cd \temp**. Dump the SAM database to a .sql file that will be copied to the other replication node. If the dump file is too big to fit on the C: drive, dump the SAM database to a .sql file in E:\Storage\Backup.

```
C:\Windows\system32>cd \temp
```

Copy the output of the query above in **BOLD** (start with the **--ignore-table=sam.1_event** and end with **sam.license**) that collected the information about the tables to ignore from replication and paste it before the **--databases** and after the **--hex-blob=1** parameters in the query below. Make sure there is a space before **--ignore-table=sam.1_event** and after **--ignore-table=sam.license** .

```
C:\Temp>mysqldump -uroot -p --master-data --routines --hex-blob=1 --ignore-table=sam.1_event --
ignore-table=sam.event --ignore-table=sam.license --databases SAM >
hcpg1.mysqldump.20221230-hcpg2.sql
```

```
Enter password: *****
```

Backup database on HCPG1 and it will be copied to and imported on HCPG3 and HCPG4

Collect the names of the tables that are ignored from replication.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query

```
MariaDB [(none)]> select group_concat(concat('--ignore-table=', TABLE_SCHEMA, '.', table_name)
SEPARATOR ' ') as " from information_schema.tables where TABLE_SCHEMA = 'SAM' and
(TABLE_NAME like '%event' or TABLE_NAME like 'license' or TABLE_NAME like 'report_status' or
TABLE_NAME like '%backfill' or TABLE_NAME like '%restoration' or TABLE_NAME like '%ntfs' or
TABLE_NAME like '%migration' or TABLE_NAME like '%archive_state' or TABLE_NAME like
'%purge');
```

```
| --ignore-table=sam.1_backfill --ignore-table=sam.1_event --ignore-table=sam.1_migration --  
ignore-table=sam.1_ntfs --ignore-table=sam.1_purge --ignore-table=sam.1_restoration --  
ignore-table=sam.archive_state --ignore-table=sam.event --ignore-table=sam.license --ignore-  
table=sam.report_status |
```

1 row in set (0.001 sec)

If needed, open a MariaDB Command Prompt and change directory to C:\Temp by issuing the command **cd \temp**

```
C:\Windows\system32>cd \temp
```

Copy the output of the query above in **BOLD** (start with the **--ignore-table=sam.1_backfill** and end with **sam.report_status**) that collected the information about the tables to ignore from replication and paste it before the **--databases** and after the **--hex-blob=1** parameters in the query below. Make sure there is a space before **--ignore-table=sam.1_backfill** and after **--ignore-table=sam.report_status**.

```
C:\Temp>mysqldump -uroot -p --master-data --routines --hex-blob=1 --ignore-table=sam.1_backfill  
--ignore-table=sam.1_event --ignore-table=sam.1_migration --ignore-table=sam.1_ntfs --  
ignore-table=sam.1_purge --ignore-table=sam.1_restoration --ignore-table=sam.archive_state  
--ignore-table=sam.event --ignore-table=sam.license --ignore-table=sam.report_status --  
databases SAM > hcp1.mysqlbackup.20221230-hcp3_4.sql
```

Enter password: *****

In the MariaDB Command Prompt (MariaDB [(none)]>) run the following query **UNLOCK TABLES;** to unlock the database tables and restore access to the shares.

```
MariaDB [(none)]> UNLOCK TABLES;
```

HCPG2

WARNING: These commands will replace the database on HCPG2 with the data imported from the database on HCPG1. If there is production data on HCPG2 that was not replicated to HCPG1, then the production data on HCPG2 will be lost.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the queries **drop database SAM;** and **create database SAM;** on HCPG2 only

```
MariaDB [(none)]> drop database SAM;
```

```
MariaDB [(none)]> create database SAM;
```

Copy the **hcp1.mysqlbackup.20221230-hcp2.sql** from HCPG1 to the C:\Temp folder on HCPG2. Open a new MariaDB Command Prompt and change directory to C:\Temp. Import the SAM database .sql file to copy the database from HCPG1 to HCPG2.

```
C:\Windows\system32>cd \temp
```

```
C:\Temp>mysql -uroot -p < hcp1.mysqlbackup.20221230-hcp2.sql
```

Enter password: *****

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> RESET MASTER;
```

```
MariaDB [(none)]> SHOW MASTER STATUS;
```

```
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| hcpg-2-bin.000001 |      329 |              | mysql              |
+-----+-----+-----+-----+
```

Note:

The output or result from the show master status query will be used on HCPG1, HCPG3 and HCPG4 in Step 5.

HCPG3

WARNING: These commands will replace the database on HCPG3 with the data imported from the database on HCPG1. If there is production data on HCPG3 that was not replicated to HCPG1, then the production data on HCPG3 will be lost.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the queries **drop database SAM;** and **create database SAM;** on HCPG3 only

```
MariaDB [(none)]> drop database SAM;
```

```
MariaDB [(none)]> create database SAM;
```

Copy the hcpg1.mysqlbackup.20221230-hcp3_4.sql from HCPG1 to the C:\Temp folder on HCPG3. Open a new MariaDB Command Prompt and change directory to C:\Temp. Import the SAM database .sql file to copy the database from HCPG1 to HCPG3.

```
C:\Windows\system32>cd \temp
```

```
C:\Temp>mysql -uroot -p < hcpg1.mysqlbackup.20221230-hcp3_4.sql
```

```
Enter password: *****
```

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> RESET MASTER;
```

```
MariaDB [(none)]> SHOW MASTER STATUS;
```

```
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| hcpg-3-bin.000001 |      329 |              | mysql              |
+-----+-----+-----+-----+
```

Note:

The output or result from the show master status query will be used on HCPG1, HCPG2 and HCPG4 in Step 5.

HCPG4

WARNING: These commands will replace the database on HCPG4 with the data imported from the database on HCPG1. If there is production data on HCPG4 that was not replicated to HCPG1, then the production data on HCPG4 will be lost.

In the MariaDB Command Prompt (MariaDB [(none)]>) run the queries **drop database SAM;** and **create database SAM;** on HCPG4 only

```
MariaDB [(none)]> drop database SAM;
```

```
MariaDB [(none)]> create database SAM;
```

Copy the hcpg1.mysqlDump.20221230-hcpG3_4.sql from HCPG1 to the C:\Temp folder on HCPG4. Open a new MariaDB Command Prompt and change directory to C:\Temp. Import the SAM database .sql file to copy the database from HCPG1 to HCPG4.

```
C:\Windows\system32>cd \temp
```

```
C:\Temp>mysql -uroot -p < hcpG1.mysqlDump.20221230-hcpG3_4.sql
```

```
Enter password: *****
```

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> RESET MASTER;
```

```
MariaDB [(none)]> SHOW MASTER STATUS;
```

```
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| hcpG-4-bin.000001 | 329 | | mysql |
+-----+-----+-----+-----+
```

Note:

The output or result from the show master status query will be used on HCPG1, HCPG2 and HCPG3 in Step 5.

Step 5 - Enable replication

HCPG1

In the MariaDB Command Prompt, execute the following query **SET GLOBAL gtid_slave_pos = "";** to reset the database replication gtid position:

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = "";
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG2. Note that the master_host IP address is for HCPG2, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node2' TO master_host="10.1.1.12", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpg-2-bin.000001", master_log_pos=329, master_ssl=1;
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG3. Note that the master_host IP address is for HCPG3, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node3' TO master_host="10.1.1.13", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpg-3-bin.000001", master_log_pos=329, master_ssl=1;
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG4. Note that the master_host IP address is for HCPG4, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node4' TO master_host="10.1.1.14", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpg-4-bin.000001", master_log_pos=329, master_ssl=1;
```

HCPG2

In the MariaDB Command Prompt, execute the following query **SET GLOBAL gtid_slave_pos = ""** to reset the database replication gtid position:

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = "";
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG1. Note that the master_host IP address is for HCPG1, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpG-1-bin.000001", master_log_pos=329, master_ssl=1;
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG3. Note that the master_host IP address is for HCPG3, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node3' TO master_host="10.1.1.13", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpG-3-bin.000001", master_log_pos=329, master_ssl=1;
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG4. Note that the master_host IP address is for HCPG4, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node4' TO master_host="10.1.1.14", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpG-4-bin.000001", master_log_pos=329, master_ssl=1;
```

HCPG3

In the MariaDB Command Prompt, execute the following query **SET GLOBAL gtid_slave_pos = ""**; to reset the database replication gtid position:

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = "";
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG1. Note that the master_host IP address is for HCPG1, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpG-1-bin.000001", master_log_pos=329, master_ssl=1;
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG2. Note that the master_host IP address is for HCPG2, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node2' TO master_host="10.1.1.12",  
master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpG-2-bin.000001", master_log_pos=329, master_ssl=1;
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG4. Note that the master_host IP address is for HCPG4, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node4' TO master_host="10.1.1.14",  
master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpG-4-bin.000001", master_log_pos=329, master_ssl=1;
```

HCPG4

In the MariaDB Command Prompt, execute the following query **SET GLOBAL gtid_slave_pos = ""** to reset the database replication gtid position:

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = "";
```


In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG1. Note that the master_host IP address is for HCPG1, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpG-1-bin.000001", master_log_pos=329, master_ssl=1;
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG2. Note that the master_host IP address is for HCPG2, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node2' TO master_host="10.1.1.12", master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpG-2-bin.000001", master_log_pos=329, master_ssl=1;
```

In the MariaDB Command Prompt, execute the following query:

Note:

The master_log_file and master_log_pos information is the result from the show master status query on HCPG3. Note that the master_host IP address is for HCPG3, and when using a cluster, is required to be on the replication network.

```
MariaDB [(none)]> CHANGE MASTER 'node3' TO master_host="10.1.1.13",  
master_port=3306,  
master_user="replication_user", master_password='0rgan1cReplication',  
master_log_file="hcpG-3-bin.000001", master_log_pos=329, master_ssl=1;
```

Step 6 – Start replication

HCPG1

```
MariaDB [(none)]> START ALL SLAVES;
```

```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

HCPG2

```
MariaDB [(none)]> START ALL SLAVES;
```

```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

HCPG3

```
MariaDB [(none)]> START ALL SLAVES;
```

```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

HCPG4

```
MariaDB [(none)]> START ALL SLAVES;
```

```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

Verify that the replication is synchronized, with no errors, on all nodes. For example, the fields below show no errors and the Slave IO and Slave SQL are running.

Slave_SQL_State: Slave has read all relay log; waiting for the slave I/O thread to update it

Slave_IO_State: Waiting for master to send event

Slave_IO_Running: Yes

Slave_SQL_Running: Yes

Last_Errno: 0

Last_Error:

Last_IO_Errno: 0

Last_IO_Error:

Last_SQL_Errno: 0

Last_SQL_Error:

Slave_SQL_Running_State: Slave has read all relay log; waiting for the slave I/O thread to update it

Step 7 – Configure replication on HCPG1, HCPG2, HCPG3 and HCPG4 to use Global Transaction ID (GTID)

HCPG1

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node2';
```

```
MariaDB [(none)]> CHANGE MASTER "node2" TO master_use_gtid=slave_pos;
```

```
MariaDB [(none)]> START SLAVE 'node2';
```

```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

```
MariaDB [(none)]> STOP SLAVE 'node3';
MariaDB [(none)]> CHANGE MASTER "node3" TO master_use_gtid=slave_pos;
MariaDB [(none)]> START SLAVE 'node3';
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

```
MariaDB [(none)]> STOP SLAVE 'node4';
MariaDB [(none)]> CHANGE MASTER "node4" TO master_use_gtid=slave_pos;
MariaDB [(none)]> START SLAVE 'node4';
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

HCPG2

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node1';
MariaDB [(none)]> CHANGE MASTER "node1" TO master_use_gtid=slave_pos;
MariaDB [(none)]> START SLAVE 'node1';
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

```
MariaDB [(none)]> STOP SLAVE 'node3';
MariaDB [(none)]> CHANGE MASTER "node3" TO master_use_gtid=slave_pos;
MariaDB [(none)]> START SLAVE 'node3';
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

```
MariaDB [(none)]> STOP SLAVE 'node4';
MariaDB [(none)]> CHANGE MASTER "node4" TO master_use_gtid=slave_pos;
MariaDB [(none)]> START SLAVE 'node4';
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

HCPG3

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node1';
MariaDB [(none)]> CHANGE MASTER "node1" TO master_use_gtid=slave_pos;
MariaDB [(none)]> START SLAVE 'node1';
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node2';
```

```
MariaDB [(none)]> CHANGE MASTER "node2" TO master_use_gtid=slave_pos;
MariaDB [(none)]> START SLAVE 'node2';
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

```
MariaDB [(none)]> STOP SLAVE 'node4';
MariaDB [(none)]> CHANGE MASTER "node4" TO master_use_gtid=slave_pos;
MariaDB [(none)]> START SLAVE 'node4';
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

HCPG4

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node1';
MariaDB [(none)]> CHANGE MASTER "node1" TO master_use_gtid=slave_pos;
MariaDB [(none)]> START SLAVE 'node1';
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

In the MariaDB Command Prompt, execute the following queries:

```
MariaDB [(none)]> STOP SLAVE 'node2';
MariaDB [(none)]> CHANGE MASTER "node2" TO master_use_gtid=slave_pos;
MariaDB [(none)]> START SLAVE 'node2';
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

```
MariaDB [(none)]> STOP SLAVE 'node3';
MariaDB [(none)]> CHANGE MASTER "node3" TO master_use_gtid=slave_pos;
MariaDB [(none)]> START SLAVE 'node3';
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

Verify that the replication is synchronized, with no errors, on all nodes.

Step 8 – Start the Windows Services Wildfly service on HCPG2, HCPG3 and HCPG4. If one of the nodes in site B is also active, then start the cluster role that contains the **SAM VFS service** and the shares in the HCP Gateway UI **on the active node in site B.**

Chapter 7 HCP Gateway Failover to Standby

1. On the standby HCP Gateway node or cluster pair of nodes, if the SMB shares have Read-only permissions, change the permissions to Read-Write.
2. Make the following modifications to the “sam.properties” file:
 - a. Edit C:\SAM\etc\sam\sam.properties
 - b. On the **standby site** HCP Gateway single node or clustered HA pair of nodes that do not have a shared cache (not a clustered HA pair with a shared cache) change setting from “registry.shares=no” to “registry.shares=yes” and then restart SAM VFS service on the new active node. If using HCP Gateway clustered HA pair of nodes with a shared cache, no change is required.
 - c. If the **primary site** HCP Gateway single node or clustered HA pair of nodes that do not have a shared cache (not a clustered HA pair with a shared cache) are still available change the setting from “registry.shares=yes” to “registry.shares=no” (for example when patching the OS) and then restart SAM VFS service on the original active node. If using HCP Gateway clustered HA pair of nodes with a shared cache, no change is required.
3. If necessary, configure DNS or Load Balancer to point to the standby HCP Gateway on the standby site.

To failback to primary Gateway, please contact Hitachi Vantara Support.

Hitachi Vantara

Corporate Headquarters 2535 Augustine Drive

Santa Clara, CA 95054 USA www.HitachiVantara.com community.HitachiVantara.com

Regional Contact Information

Americas: +1 866 374 5822 or info@hitachivantara.com

Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hitachivantara.com

Asia Pacific: +852 3189 7900 or info.marketing.apac@hitachivantara.com

