

Hitachi Content Platform Gateway

Database Replication Setup for Windows

v 4.2.0

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-HCPG014-03

April 2022

© 2020, 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, the Microsoft Corporate 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 at <https://www.hitachivantara.com/en-us/company/legal.html>.

Table of Contents

Chapter 1 Introduction	2
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)	37
Chapter 5 Three Node: Master to Master Replication for HA Cluster with optional Read-only Copy on secondary site	48
Chapter 6 Four Nodes: Master to Master Replication for HA Cluster with DR Failover to HA Cluster	61
Chapter 7 HCP Gateway Failover to Standby.....	78

Chapter 1 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:

```
hcpg1 10.1.1.11
hcpg2 10.1.1.12
hcpg3 10.1.1.13
hcpg4 10.1.1.14
```

In addition, the internal database replication commands 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) sam.properties server.id that identifies the HCP Gateway device. Second there is a database server.id that distinguishes the database instances.

This document covers the following database replication configurations:

1. **Two Node:** Master node to Replica node replication for Read-only copy on second site (Single Node)
2. **Two Node:** Master node to Master node replication for Manual DR Failover (Single Node)
3. **Three Node:** Master node to 2 separate Replica nodes replication for Read-only copy on second and third sites (Single Node)
4. **Three Node:** Master to Master replication for HA **Cluster** and optional Replica for Read-only copy on second site (Cluster Node)
5. **Four Node:** Master to Master replication for HA **Cluster** and Master to Master HA **Cluster** 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: Master to Master DB copy (DR Hot Standby on Site B)	Single Node	✓ (Single Node)	✓ (optional)	
Two or Three Node: Master to Master HA Cluster on Site A with <u>optional</u> 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: After creating a share on the active node in a pair of replicated or clustered HCP Gateways, on the passive node, open a MySQL Client Window, login to the database and issue the command **show all slaves status\G**. If you see a **Last Error** line that shows that a table does not exist, for example, it contains **Error 'Table sam.X_migration' doesn't exist' on query**, then wait for a time when nobody is writing files to the active node of the HCP Gateway. Then fail over the SAM VFS service from the active node to the passive node so that all of the tables will be created on the passive node. Then, before any files are written to or updated on the passive node, fail the SAM VFS service back to the active node. Then on the passive node, in the MySQL Client Window, issue the command **stop slave 'nodeX';** where X is the **Connection name** of the node that displayed the **Last Error** line that contains **Error 'Table sam.X_migration' doesn't exist' on query**. Then issue the command **start slave 'nodeX';** Then issue the command **show all slaves status\G** to verify that there are no **Last Error** lines containing **Error 'Table sam.X_migration' doesn't exist' on query**. If using a 4 node replication with cluster, with 2 cluster nodes in Site A and 2 cluster nodes in Site B, after following these steps in Site A, then start the shares on the active node in Site B and follow the

steps above to ensure that the tables for the new share replicate to all the nodes in Site B. 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.

Refer to the HCP Gateway Admin Guide Shares Chapter Section 10.2 for more details about setting the Replicated parameter in a Replication configuration.

There are four configurations documented below, choose the one that best fits your configuration.

- 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:

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

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

- 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. 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.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
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.%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.%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:

```
# 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.%protection
node3.replicate-wild-ignore-table = SAM.%ntfs
node3.replicate-wild-ignore-table = SAM.%migration
node3.replicate-wild-ignore-table = SAM.%purge
```

```
node3.replicate-wild-ignore-table = sam.%backfill
node3.replicate-wild-ignore-table = sam.%restoration
#node3.replicate-wild-ignore-table = sam.%protection
```



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

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.%protection
node3.replicate-wild-ignore-table = SAM.%ntfs
node3.replicate-wild-ignore-table = SAM.%migration
node3.replicate-wild-ignore-table = SAM.%purge

node3.replicate-wild-ignore-table = sam.%backfill
node3.replicate-wild-ignore-table = sam.%restoration
#node3.replicate-wild-ignore-table = sam.%protection
```

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

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

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

- 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:

```
# 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.%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.%purge
node4.replicate-wild-ignore-table = SAM.%purge
```

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

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

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

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.%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.%purge

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

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

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

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

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

WARNING: If the Active node fails and the Passive nodes 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:

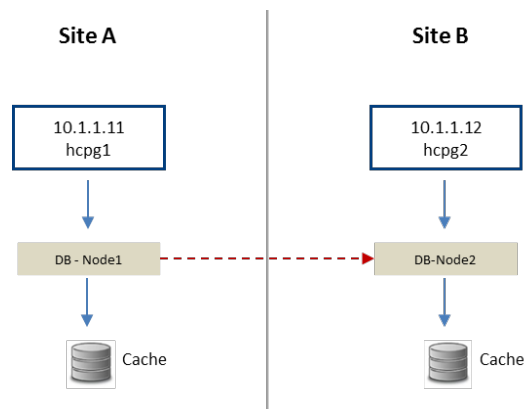
- 1) Only 1 Gateway device can be active at a time.
- 2) The files must be written to the HCP and if there is a second HCP, replicated to the second HCP, before the content will be available on the non-active node(s) that have a separate cache. For example, if using a Copy or Tiering policy that waits 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.

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 Node1 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 Node1 to replicate to Node2, the file metadata written to hcp1 should 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.
2. You must stop the MariaDB service if any changes are made to the *my.ini* file. Then restart the MariaDB service after saving the file.

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
```

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 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 migration and ntfs tables are excluded from being replicated in addition to the common list of excluded tables.

```
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

# 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

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

Step 3 – Common Replication Setup Steps

It is required 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 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;

Note: the password is the string immediately following 'IDENTIFIED BY'

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

MariaDB Command Prompt (MariaDB [(none)]>) run on HCPG1 and HCPG2

```
mysql -u root -p <press enter>
```

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

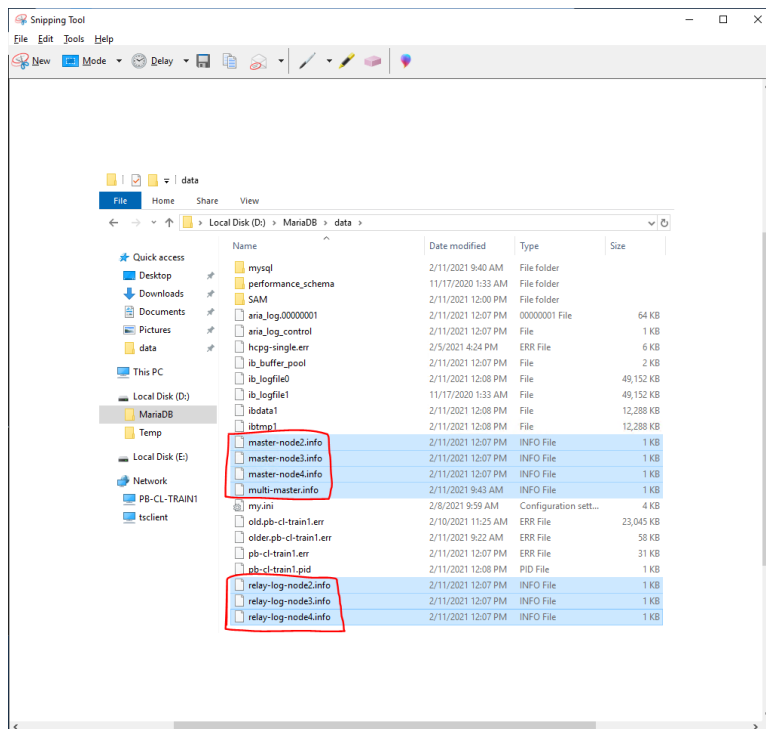
```
MariaDB [(none)]> CREATE USER 'replication_user'@'%' IDENTIFIED BY 'Organ1cReplication';  
MariaDB [(none)]> GRANT REPLICATION SLAVE ON *.* TO 'replication_user'@'%';  
MariaDB [(none)]> FLUSH PRIVILEGES;  
MariaDB [(none)]> FLUSH TABLES;
```

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


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

Open Windows Services and if running, stop the SAM VFS, Wildfly and MariaDB services on HCPG1 and HCPG2.


Open Windows File explorer and delete all the .info files in the D:\MariaDB\data folder on HCPG1 and HCPG2.



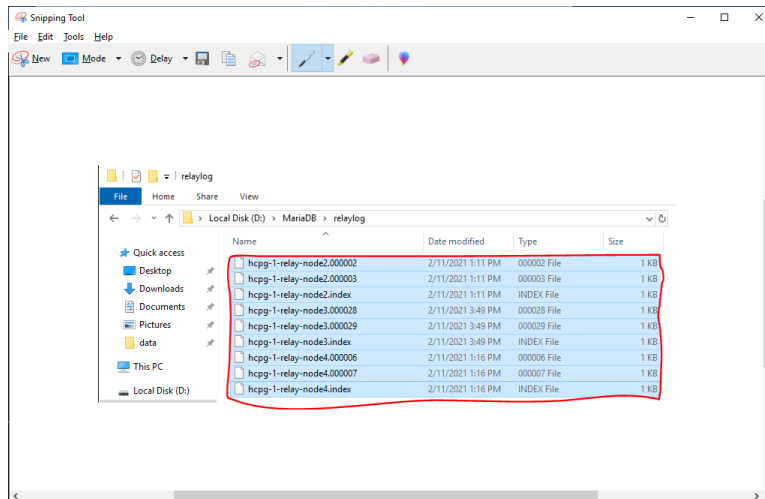
Delete or rename the D:\MariaDB\data*.err file on HCPG1 and HCPG2, where the name of the file is the Windows computer name of the Gateway.



Delete all the files in the D:\MariaDB\binlog folder that begin with hcppg* on HCPG1 and HCPG2.

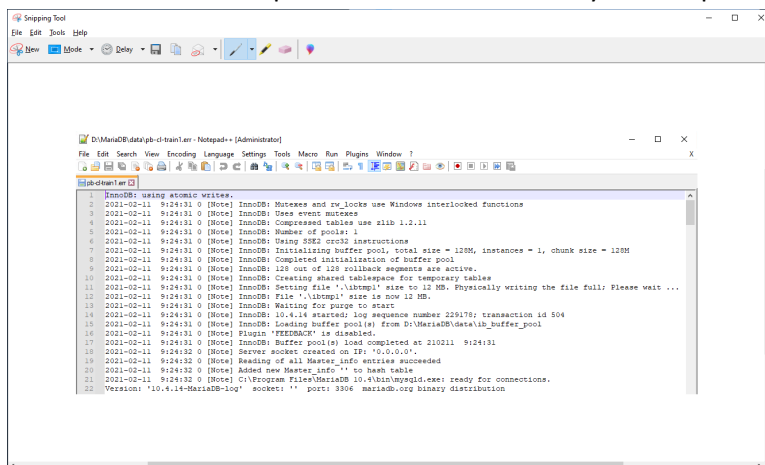


Delete all the files in the D:\MariaDB\relaylog folder that begin with hcp* on HCPG1 and HCPG2.



In Windows Services start the MariaDB service on HCPG1 and HCPG2.

Check the D:\MariaDB\data*.err file in Notepad++ for any errors on HCPG1 and 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)

MariaDB [(none)]> RESET MASTER;

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.

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


```

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)

```

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.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.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 --databases SAM > pb-cl-train1.202102101130.sql
Enter password: *****

```

Back in the MariaDB Command Prompt, execute the following commands:

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

```

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

```

Note: The output or result from the show master status command will be used in command below ('hcpg-1-bin.000001', 329).

```

MariaDB [(none)]> SELECT BINLOG_GTID_POS('hcpg-1-bin.000001', 329);
+-----+-----+-----+-----+
| BINLOG_GTID_POS('hcpg-1-bin.000001', 329) |
+-----+-----+-----+-----+
| 1-1-14 |
+-----+-----+-----+-----+

```

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 on HCPG2 only

```

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

```

Copy the pb-cl-train1.202102101130.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.

If needed, open a MariaDB Command Prompt

```

C:\Windows\system32>cd \temp
C:\Temp>mysql -uroot -p < pb-cl-train1.202102101130.sql
Enter password: *****

```

Step 5 – Enable Replication

HCPG2

In the MariaDB Command Prompt, execute the following commands:

Note: The position information (e.g. 1-1-14) is from 'hcpg1' in this example) BINLOG_GTID_POS output above

If the result is empty or NULL, enter an empty string " (2 single quotes)

If the result is a string of numbers separated by dashes '-', for example: '1-1-14'

Then use the entire string enclosed by single quotes '1-1-14'

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = ''; ← This is for a NULL value (2 single quotes).
```

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = '1-1-14'; ← This is value of 1-1-14.
```

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11",  
master_port=3306,  
master_user="replication_user", master_password='Organ1cReplication',  
master_use_gtid=slave_pos, master_ssl=1;
```

Step 6 – Start replication on both nodes HCPG1 and HCPG2

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

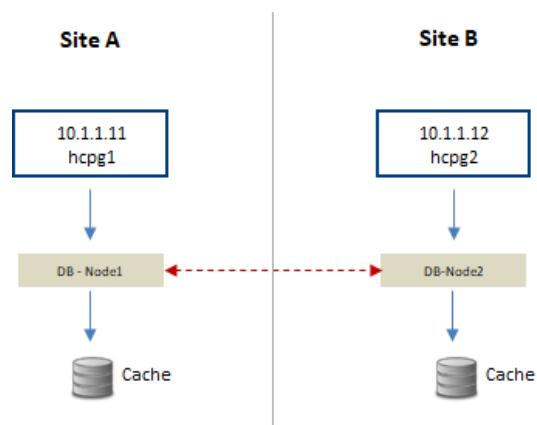
```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

Step 7 – Start the Windows Services Wildfly and SAM VFS on all nodes and then start the shares in the HCP Gateway UI on HCPG1.

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, it requires manual intervention to switch to the DR site, and to fail back from Site B to Site A.

Two Nodes: Master-Master (DR Failover)



The 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 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.
2. You must stop the MariaDB service if any changes are made to the `my.ini` file. Then restart the MariaDB service after saving the file.

Step 1 – Update HCPG Gateway properties file

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

Make the following edits:

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
```

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 hcpg systems, and, with the exact settings as shown for each hcpg 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 migration and NTFS tables are excluded from being replicated in addition to the common list of excluded tables.

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

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

# 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

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

```

Step 3 – Common Replication Setup Steps

It is required 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 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;

Note: the password is the string immediately following 'IDENTIFIED BY'

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

MariaDB Command Prompt (MariaDB [(none)]>) run on HCPG1 and HCPG2

```
mysql -u root -p <press enter>
```

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

```
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 on HCPG1 and HCPG2

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


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

Open Windows Services and if running, stop the SAM VFS, Wildfly and MariaDB services on HCPG1 and HCPG2.


Open Windows File explorer and delete all the .info files in the D:\MariaDB\data folder on HCPG1 and HCPG2.



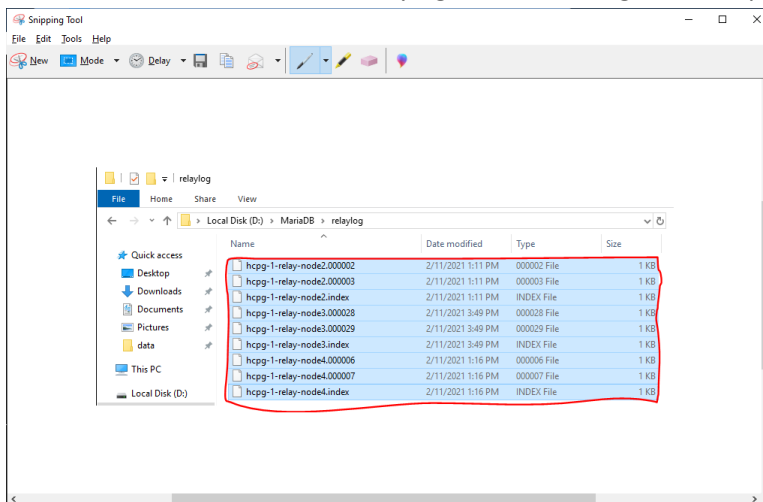
Delete or rename the D:\MariaDB\data*.err file on HCPG1 and HCPG2, where the name of the file is the Windows computer name of the Gateway.



Delete all the files in the D:\MariaDB\binlog folder that begin with hcpg* on HCPG1 and HCPG2.

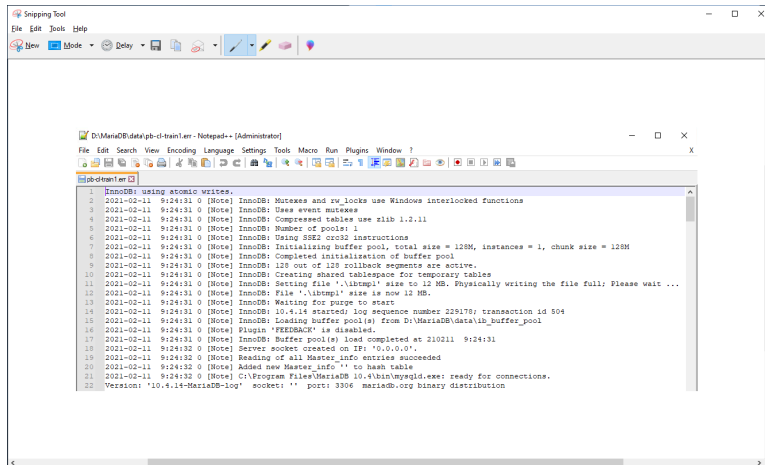


Delete all the files in the D:\MariaDB\relaylog folder that begin with hcpG* on HCPG1 and HCPG2.



In Windows Services start the MariaDB service on HCPG1 and HCPG2.

Check the D:\MariaDB\data*.err file in Notepad++ for any errors on HCPG1 and 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)

MariaDB [(none)]> RESET MASTER;

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.

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.

```
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 --
ignoring-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)

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.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.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 > pb-cl-train1.202102101130.sql
Enter password: *****
```

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

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

```
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)

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.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=--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 > pb-cl-  
train1.202102101130.sql  
Enter password: *****
```

Back in the MariaDB Command Prompt, execute the following commands:

```
MariaDB [(none)]> SHOW MASTER STATUS;  
+-----+-----+-----+-----+  
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |  
+-----+-----+-----+-----+  
| hcpg-1-bin.000001 |      329 |              | mysql              |  
+-----+-----+-----+-----+
```

Note: The output or result from this command will be used in 'hcpg2'

```
MariaDB [(none)]> SELECT BINLOG_GTID_POS('hcpg-1-bin.000001', 329);  
+-----+-----+-----+-----+  
| BINLOG_GTID_POS('hcpg-1-bin.000001', 329) |  
+-----+-----+-----+-----+  
| NULL                                         |  
+-----+-----+-----+-----+
```

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 on HCPG2 only

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

Copy the pb-cl-train1.202102101130.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.

If needed, open a MariaDB Command Prompt

```
C:\Windows\system32>cd \temp  
C:\Temp>mysql -uroot -p < pb-cl-train1.202102101130.sql  
Enter password: *****
```

In the MariaDB Command Prompt, execute the following commands:

```
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 command will be used in 'hcpg1'

```
MariaDB [(none)]> SELECT BINLOG_GTID_POS('hcpg-2-bin.000001', 329);
+-----+
| BINLOG_GTID_POS('hcpg-2-bin.000001', 329) |
+-----+
| 2-1-3                                     |
+-----+
```

Step 5 – Enable Replication

HCPG1

Note: The position information is from the 'hcpg2' SELECT BINLOG_GTID_POS output

If the result is empty or NULL, enter an empty string '' (2 single quotes)

If the result is a string of numbers separated by dashes '-',

example: '2-1-3'

Use the entire string enclosed by single quotes '2-1-3'

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = ''; ← This is for a NULL value (2 single quotes).
```

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = '2-1-3'; ← This is value for 2-1-3.
```

Note: The IP address below is for hcpg2

```
MariaDB [(none)]> CHANGE MASTER 'node2' TO master_host="10.1.1.12", master_port=3306,
master_user="replication_user", master_password='Organ1cReplication',
master_use_gtid=slave_pos, master_ssl=1;
```

HCPG2

Note: The position information is from the 'hcpg1' SELECT BINLOG_GTID_POS output

If the result is empty or NULL, enter an empty string '' (2 single quotes)

If the result is a string of numbers separated by dashes '-',

For example: '1-1-3'

use the entire string enclosed by single quotes '1-1-3'

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = ''; ← This is for a NULL value (2 single quotes).
```

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = '1-1-3'; ← This is value of 1-1-3.
```

Note: The IP address below is for hcp1

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11", master_port=3306,  
master_user="replication_user", master_password='Organ1cReplication',  
master_use_gtid=slave_pos, 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
```

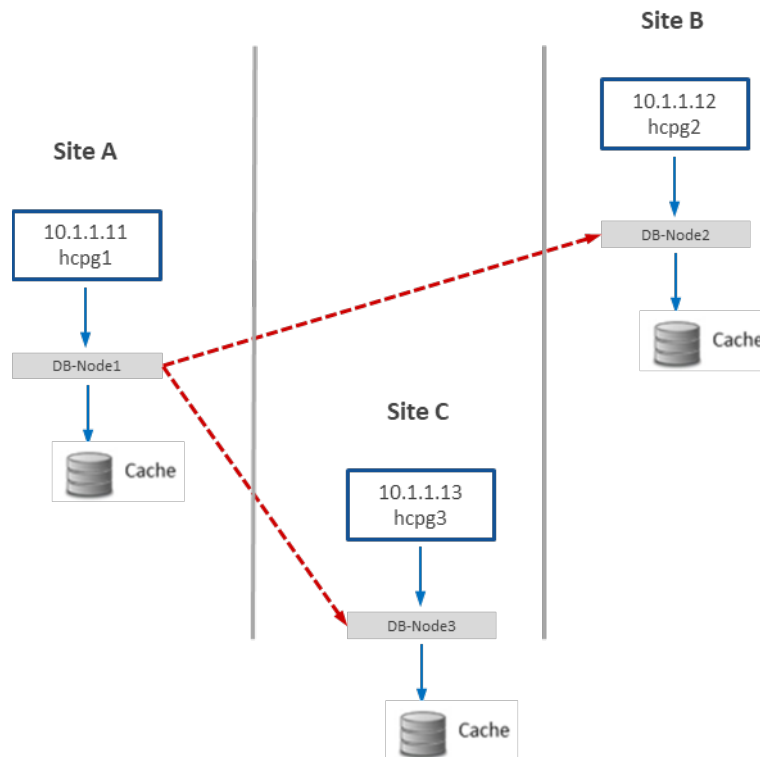
Step 7 – Start the Windows Services Wildfly and SAM VFS on all nodes and then start the shares in the HCP Gateway UI on HCPG1.

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 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 Node1 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 database Node1 to replicate to Node2 and Node3, the file metadata written to hcpg1 should 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 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.
2. You must stop the MariaDB service if any changes are made to the my.ini file. Then restart the MariaDB service after saving the file.

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 migration and nfts tables are excluded from being replicated in addition to the common

list of excluded tables.

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
```

```
# 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
```

```

# 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

# Relay Logging
relay-log        = D:/MariaDB/relaylog/hcpg-3-relay

```

Step 3 – Common Replication Setup Steps

It is required 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;

Note: the password is the string immediately following 'IDENTIFIED BY'

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

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

```
MariaDB [(none)]> mysql -u root -p <press enter>
```

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


```
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 on HCPG1, HCPG2 and HCPG3


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

Open Windows Services and if running, stop the SAM VFS, Wildfly and MariaDB services on HCPG1, HCPG2 and HCPG3.


Open Windows File explorer and delete all the .info files in the D:\MariaDB\data folder on HCPG1, HCPG2 and HCPG3.



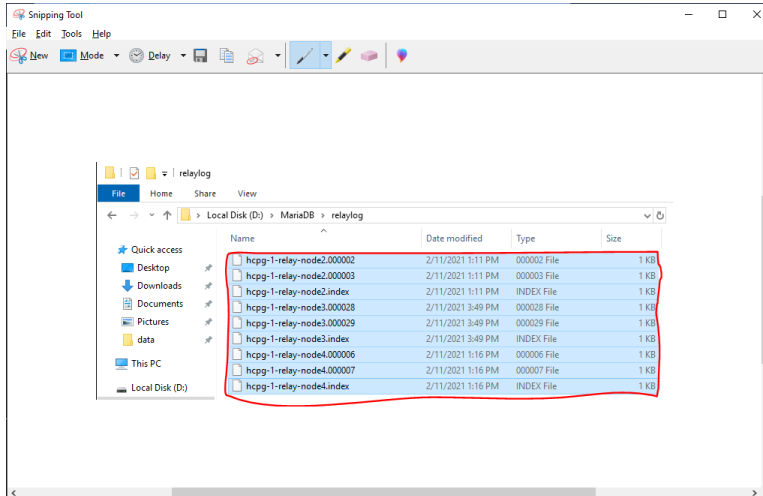
Delete or rename the D:\MariaDB\data*.err file on HCPG1, HCPG2 and HCPG3, where the name of the file is the Windows computer name of the Gateway.



Delete all the files in the D:\MariaDB\binlog folder that begin with hcppg* on HCPG1, HCPG2 and HCPG3.

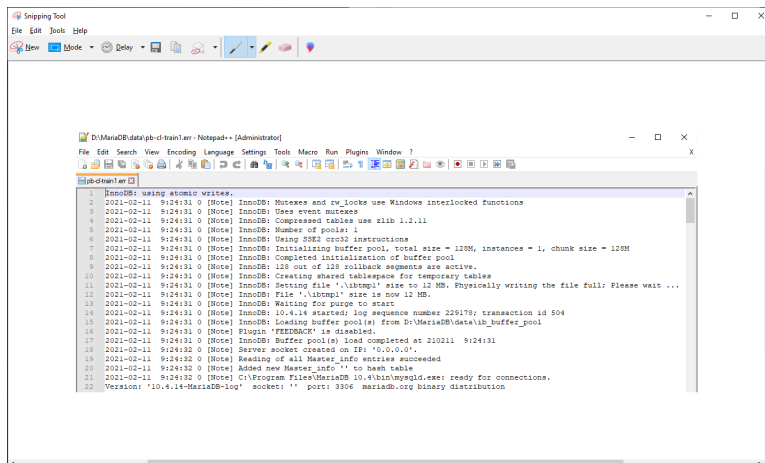


Delete all the files in the D:\MariaDB\relaylog folder that begin with hcppg* on HCPG1, HCPG2 and HCPG3.



In Windows Services start the MariaDB service on **HCPG1**, **HCPG2** and **HCPG3**.

Check the D:\MariaDB\data*.err file in Notepad++ for any errors on **HCPG1**, **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)

MariaDB [(none)]> RESET MASTER;

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 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.

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

```

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 'archive_state' or 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 '%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)

```

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.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 > pb-cl-train1.202102101130.sql
Enter password: *****

```

Back in the MariaDB Command Prompt, execute the following commands:

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

```

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

```

Note: The output or result from the show master status command will be used in command below ('hcpg-1-bin.000001', 329).

```

MariaDB [(none)]> SELECT BINLOG_GTID_POS('hcpg-1-bin.000001', 329);

```

```

+-----+
| BINLOG_GTID_POS('hcpg-1-bin.000001', 329) |
+-----+
| NULL |
+-----+

| BINLOG_GTID_POS('hcpg-1-bin.000001', 329) |
+-----+
| 1-1-14 |
+-----+

```

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 on HCPG2 only

```

MariaDB [(none)]> drop database SAM;

```

```

MariaDB [(none)]> create database SAM;

```

Copy the pb-cl-train1.202102101130.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.

If needed, open a MariaDB Command Prompt

```

C:\Windows\system32>cd \temp

```

```

C:\Temp>mysql -uroot -p < pb-cl-train1.202102101130.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 on HCPG3 only

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

Copy the pb-cl-train1.202102101130.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.

If needed, open a MariaDB Command Prompt

```
C:\Windows\system32>cd \temp  
C:\Temp>mysql -uroot -p < pb-cl-train1.202102101130.sql  
Enter password: *****
```

Step 5 – Enable Replication

HCPG2

In the MariaDB Command Prompt, execute the following commands:

Note: The position information (e.g. 1-1-14) is from 'hcpG1' in this example) BINLOG_GTID_POS output above

If the result is empty or NULL, enter an empty string " (2 single quotes)

If the result is a string of numbers separated by dashes '-', for example: '1-1-14'

Then use the entire string enclosed by single quotes '1-1-14'

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = ''; ← This is for a NULL value (2 single quotes).
```

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = '1-1-14'; ← This is value of 1-1-14.
```

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11",  
master_port=3306,  
master_user="replication_user", master_password='Organ1cReplication',  
master_use_gtid=slave_pos, master_ssl=1;
```

HCPG3

In the MariaDB Command Prompt, execute the following commands:

Note: The position information (e.g. 1-1-14) is from 'hcpG1' in this example) BINLOG_GTID_POS output above

If the result is empty or NULL, enter an empty string " (2 single quotes)

If the result is a string of numbers separated by dashes '-', for example: '1-1-14'

Then use the entire string enclosed by single quotes '1-1-14'

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = ''; ← This is for a NULL value (2 single quotes).
```

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = '1-1-14'; ← This is value of 1-1-14.
```



```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11",
master_port=3306,
master_user="replication_user", master_password='Organ1cReplication',
master_use_gtid=slave_pos, master_ssl=1;
```

Step 6 – Start replication on all nodes HCPG1, HCPG2 and HCPG3

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

Step 7 – Start the Windows Services Wildfly and SAM VFS on all nodes and then start the shares in the HCP Gateway UI on HCPG1.

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

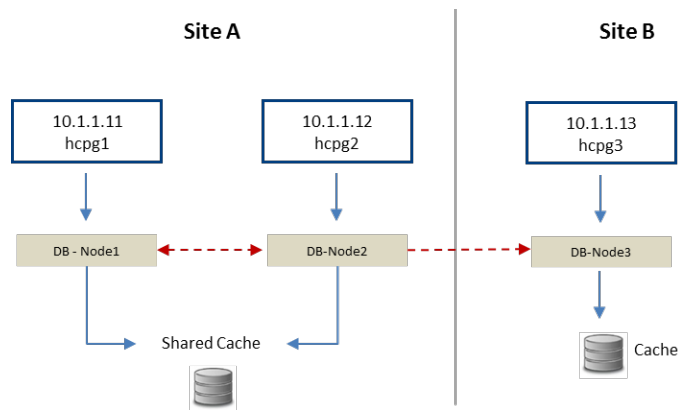
The two gateways on Site A are in a cluster and have a shared cache, only one gateway is active. The gateway on Site B has an independent cache. Metadata written to Site A database, will be available on Site B after a short delay. In this configuration:

- Database on hcp1 is replicated to hcp2 and hcp3 nodes
- Database on hcp2 is replicated to hcp1 and hcp3 nodes

The intent for this configuration is that Site B (hcp3) 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.
2. You must stop the MariaDB service if any changes are made to the my.ini file. Then restart the MariaDB service after saving the file.

Step 1 – Update HCPG 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 respective hcp systems.

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

HCPG1

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

HCPG2

```
server.id=2  
server.ignore=1  
cluster=1  
binlog.name=hcp-2-bin
```

HCPG3

```
server.id=3  
server.ignore=0  
cluster=0  
binlog.name=hcp-3-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 prevent 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 migration and NTFS tables are excluded from being replicated in addition to the common list of excluded tables.

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
```

```
# 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
```

```
# 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

# Relay Logging
relay-log = D:/MariaDB/relaylog/hcpg-3-relay

```

Step 3 – Common Replication Setup Steps

It is required 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.

Note: the password is the string immediately following 'IDENTIFIED BY'

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

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

```
mysql -u root -p <press enter>
```

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

```

MariaDB [(none)]> CREATE USER 'replication_user'@'%' IDENTIFIED BY 'Organ1cReplication';
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 on HCPG1, HCPG2 and HCPG3

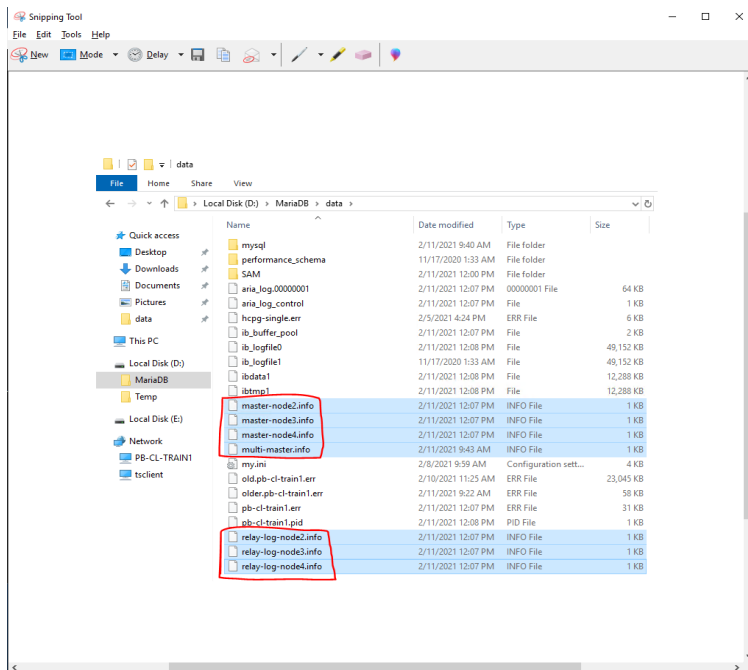
```

MariaDB [(none)]> stop all slaves;
MariaDB [(none)]> reset slave all;


```

Open Windows Services and if running, stop the SAM VFS, Wildfly and MariaDB services on HCPG1, HCPG2 and HCPG3.


Open Windows File explorer and delete all the .info files in the D:\MariaDB\data folder on HCPG1, HCPG2 and HCPG3.



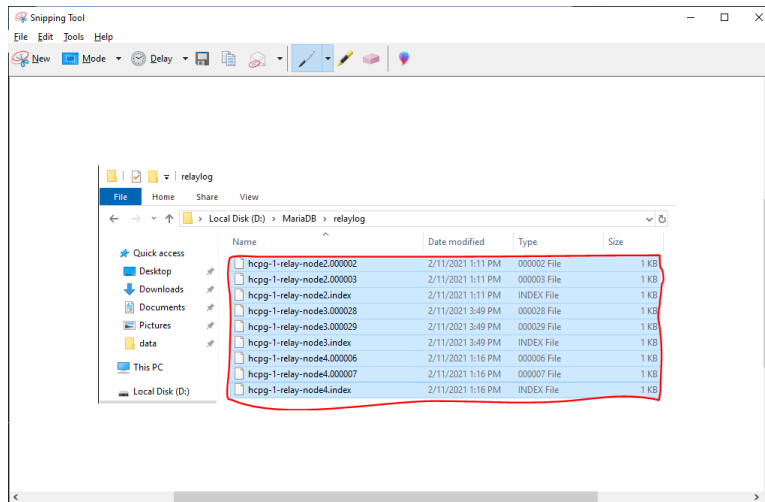
Delete or rename the D:\MariaDB\data*.err file on HCPG1, HCPG2 and HCPG3, where the name of the file is the Windows computer name of the Gateway.



Delete all files in the D:\MariaDB\binlog folder that begin with hcfg* on HCPG1, HCPG2 and HCPG3. The filenames on HCPG1 will be hcfg-1- *, hcfg-2- * on HCPG2 and hcfg-3- * on HCPG3.

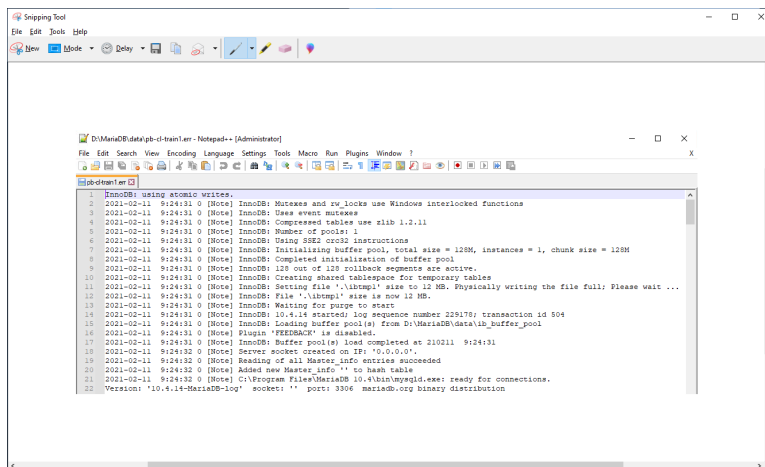


Delete all the files in the D:\MariaDB\relaylog folder that begin with hcp* on HCPG1, HCPG2 and HCPG3.



In Windows Services start the MariaDB service on HCPG1, HCPG2 and HCPG3.

Check the D:\MariaDB\data*.err file in Notepad++ for any errors on HCPG1, 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)

MariaDB [(none)]> RESET MASTER;

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.

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.

```
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)

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_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=--ignore-  
table=sam.1_event --ignore-table=sam.event --ignore-table=sam.license --databases  
SAM > pb-cl-train1.202102101130-hcpg2.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.

```
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 > pb-cl-train1.202102101130-hcpg3.sql
Enter password: *****
```

Back in the MariaDB Command Prompt, execute the following commands:

```
MariaDB [(none)]> SHOW MASTER STATUS;
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| hcpg-1-bin.000001 |      329 |              | mysql              |
+-----+-----+-----+-----+
```

Note: The output or result from the show master status command will be used in command below ('hcpg-1-bin.000001', 329).

```
MariaDB [(none)]> SELECT BINLOG_GTID_POS('hcpg-1-bin.000001', 329);
+-----+
| BINLOG_GTID_POS('hcpg-1-bin.000001', 329) |
+-----+
| NULL                                         |
+-----+
```

Note: The output or result from this command will be used in 'hcpg2'

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 on HCPG2 only

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

Copy the pb-cl-train1.202102101130-hcpg2.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.

```
If needed, open a MariaDB Command Prompt  
C:\Windows\system32>cd \temp  
C:\Temp>mysql -uroot -p < pb-cl-train1.202102101130-hcpg2.sql  
Enter password: *****
```

In the MariaDB Command Prompt, execute the following commands:

```
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 command will be used in command below ('hcpg-2-bin.000001', 329).

```
MariaDB [(none)]> SELECT BINLOG_GTID_POS('hcpg-2-bin.000001', 329);  
+-----+  
| BINLOG_GTID_POS('hcpg-2-bin.000001', 329) |  
+-----+  
| NULL |  
+-----+
```

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 on HCPG3 only

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

Copy the pb-cl-train1.202102101130-hcpg3.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.

If needed, open a MariaDB Command Prompt

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

```
C:\Temp>mysql -uroot -p < pb-cl-train1.202102101130-hcpg3.sql
```

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

Step 5 – Enable replication

HCPG1

Note: The position information is from the 'hcpg2' SELECT BINLOG_GTID_POS output

If the result is empty or NULL, enter an empty string " (2 single quotes)

If the result is a string of numbers separated by dashes '-',

example: '2-1-3'

use the entire string enclosed by single quotes '2-1-3'

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = ""; ← This is for a NULL value (2 single quotes).
```

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = '2-1-3'; ← This is value of 2-1-3.
```

Note: The IP address below is for hcpg2

```
MariaDB [(none)]> CHANGE MASTER 'node2' TO master_host="10.1.1.12", master_port=3306,  
master_user="replication_user", master_password='Organ1cReplication',  
master_use_gtid=slave_pos, master_ssl=1;
```

HCPG2

Note: The position information is from the 'hcpg1' SELECT BINLOG_GTID_POS output

If the result is empty or NULL, enter an empty string " (2 single quotes)

If the result is a string of numbers separated by dashes '-',

example: '1-1-3'

use the entire string enclosed by single quotes '1-1-3'

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = ""; ← This is for a NULL value (2 single quotes).
```

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = '1-1-3'; ← This is value of 1-1-3.
```

Note: The IP address below is for hcpg1

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11", master_port=3306,
```

```
master_user="replication_user", master_password='Organ1cReplication',
master_use_gtid=slave_pos, master_ssl=1;
```

HCPG3

Note: The position information is from the 'hcp1' and 'hcg2' SELECT BINLOG_GTID_POS output

If the result is empty or NULL, enter an empty string " (2 single quotes)

If the result is a string of numbers separated by dashes '-',

example: '1-1-3' from hcp1 and '2-1-3' from hcg2

use the entire string from both hcp1 and hcg2 enclosed by single quotes '1-1-3,2-1-3'

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = ""; ← This is for a NULL value (2 single quotes).
```

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = '1-1-3,2-1-3'; ← This is value of 1-1-3 and 2-1-3.
```

Note: The IP address below is for hcp1

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11", master_port=3306,
master_user="replication_user", master_password='Organ1cReplication',
master_use_gtid=slave_pos, master_ssl=1;
```

Note: The IP address below is for hcg2

```
MariaDB [(none)]> CHANGE MASTER 'node2' TO master_host="10.1.1.12", master_port=3306,
master_user="replication_user", master_password='Organ1cReplication',
master_use_gtid=slave_pos, 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;
```

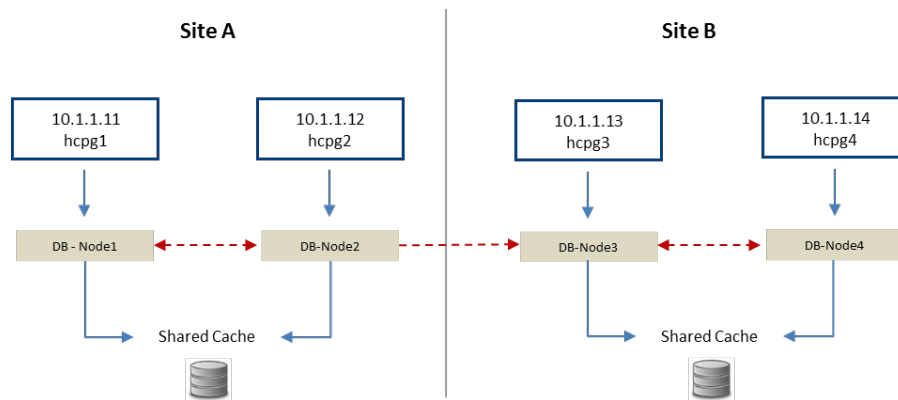
```
MariaDB [(none)]> SHOW ALL SLAVES STATUS\G
```

Step 7 – Start the Windows Services Wildfly and SAM VFS on all nodes and then start the shares in the HCP Gateway UI **on HCPG1**.

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, only one gateway 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.
2. You must stop the MariaDB service if any changes are made to the my.ini file. Then restart the MariaDB service after saving the file.

Step 1 - Update HCPG Gateway properties file

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

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

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

Make the following edits:

HCPG1

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

HCPG2

```
server.id=2
server.ignore=1
cluster=1
binlog.name=hcpg-2-bin
```

HCPG3

```
server.id=3
server.ignore=1
cluster=1
binlog.name=hcpg-3-bin
```

HCPG4

```
server.id=4
server.ignore=1
cluster=1
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 migration and NTFS tables are excluded from being replicated in addition to the common list of excluded tables.

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

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
```

```
# 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
```

```
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

# 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

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

Step 3 – Common Replication Setup Steps

It is required 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;

Note: the password is the string immediately following 'IDENTIFIED BY'

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

MariaDB Command Prompt (MariaDB [(none)]>) run on **HCPG1, HCPG2, HCPG3 and HCPG4**
mysql -u root -p <press enter>

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

```
MariaDB [(none)]> CREATE USER 'replication_user'@'%' IDENTIFIED BY 'Organ1cReplication';  
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 on **HCPG1, HCPG2, HCPG3 and HCPG4**

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

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

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

then on HCPG2 run the following commands

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

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

then on HCPG3 run the following commands

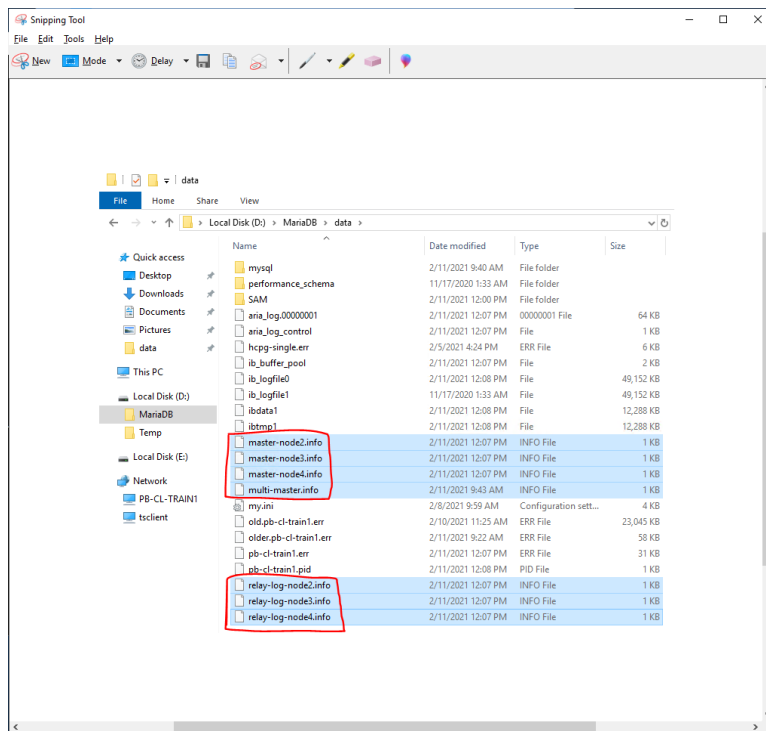
```
MariaDB [(none)]> reset slave 'node1' all;  
MariaDB [(none)]> reset slave 'node2' all;  
MariaDB [(none)]> reset slave 'node4' all;
```

then on HCPG4 run the following commands


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

Open Windows Services and if running, stop the SAM VFS, Wildfly and MariaDB services on **HCPG1, HCPG2, HCPG3 and HCPG4**.


Open Windows File explorer and delete all the .info files in the D:\MariaDB\data folder on **HCPG1, HCPG2, HCPG3 and HCPG4**.



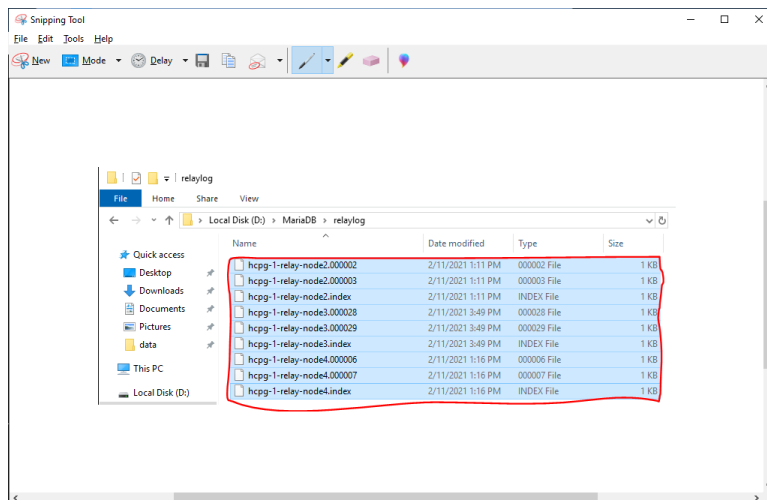
Delete or rename the D:\MariaDB\data*.err file on **HCPG1, HCPG2, HCPG3 and HCPG4**, where the name of the file is the Windows computer name of the Gateway.



Delete all the files in the `D:\MariaDB\binlog` folder that begin with `hcppg*` on **HCPG1**, **HCPG2**, **HCPG3** and **HCPG4**. The filenames on HCPG1 will be `hcppg-1-*`, `hcppg-2-*` on HCPG2, `hcppg-3-*` on HCPG3 and `hcppg-4-*` on HCPG4.

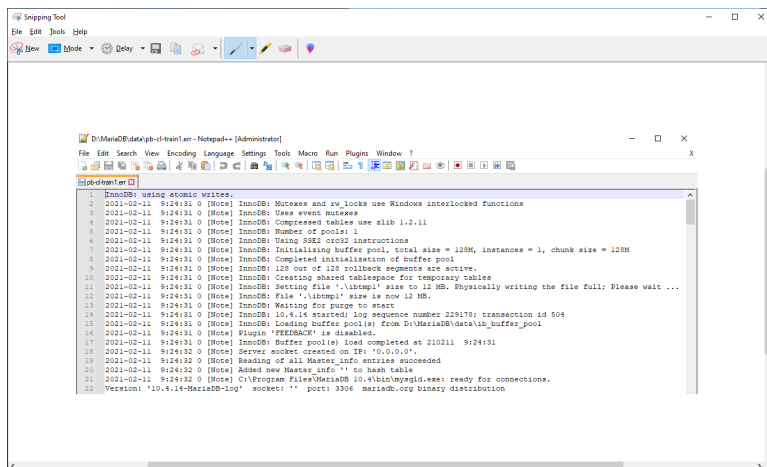


Delete all the files in the D:\MariaDB\relaylog folder that begin with hcpg* on HCPG1, HCPG2, HCPG3 and HCPG4.



In Windows Services start the MariaDB service on HCPG1, HCPG2, HCPG3 and HCPG4.

Check the D:\MariaDB\data*.err file in Notepad++ for any errors on HCPG1, 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

Note: The output or result from this command will be used in the other 3 hcpgs. For example, the output from 'hcpg1' will be used in 'hcpg2', 'hcpg3', 'hcpg4'

HCPG1 (the active node)

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

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.

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

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

```
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');
```

```
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)
```

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_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 > pb-cl-train1.202102101130-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.

```
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 > pb-cl-  
train1.202102101130-hcpg3_4.sql  
Enter password: *****
```

Back in the MariaDB Command Prompt, execute the following commands:

```
MariaDB [(none)]> SHOW MASTER STATUS;  
+-----+-----+-----+-----+  
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |  
+-----+-----+-----+-----+  
| hcpg-1-bin.000001 |      329 |              | mysql              |  
+-----+-----+-----+-----+
```


Note: The output or result from the show master status command will be used in command below ('hcpg-1-bin.000001', 329).

```
MariaDB [(none)]> SELECT BINLOG_GTID_POS('hcpg-1-bin.000001', 329);
+-----+
| BINLOG_GTID_POS('hcpg-1-bin.000001', 329) |
+-----+
| 1-1-14                                     |
+-----+
```

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 on HCPG2 only

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

Copy the pb-cl-train1.202102101130-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.

If needed, open a MariaDB Command Prompt

```
C:\Windows\system32>cd \temp
C:\Temp>mysql -uroot -p < pb-cl-train1.202102101130-hcp2.sql
Enter password: *****
```

In the MariaDB Command Prompt, execute the following commands:

```
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 command will be used in command below ('hcpg-2-bin.000001', 329).

```
MariaDB [(none)]> SELECT BINLOG_GTID_POS('hcpg-2-bin.000001', 329);
+-----+
| BINLOG_GTID_POS('hcpg-2-bin.000001', 329) |
+-----+
| NULL                                     |
+-----+
```

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 on HCPG3 only

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

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

Copy the pb-cl-train1.202102101130-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.

If needed, open a MariaDB Command Prompt

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

```
C:\Temp>mysql -uroot -p < pb-cl-train1.202102101130-hcp3_4.sql
```

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

In the MariaDB Command Prompt, execute the following commands:

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

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

```
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| hcp3-3-bin.000001 |      329 |              | mysql              |
+-----+-----+-----+-----+
```

Note: The output or result from the show master status command will be used in command below ('hcp3-3-bin.000001', 329).

```
MariaDB [(none)]> SELECT BINLOG_GTID_POS('hcp3-3-bin.000001', 329);
```

```
+-----+-----+
| BINLOG_GTID_POS('hcp3-3-bin.000001', 329) |
+-----+-----+
| 3-1-14                                     |
+-----+-----+
```

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 on HCPG4 only

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

Copy the pb-cl-train1.202102101130-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.

If needed, open a MariaDB Command Prompt

```
C:\Windows\system32>cd \temp  
C:\Temp>mysql -uroot -p < pb-cl-train1.202102101130-hcpg3_4.sql  
Enter password: *****
```

In the MariaDB Command Prompt, execute the following commands:

```
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 command will be used in command below ('hcpg-4-bin.000001', 329).

```
MariaDB [(none)]> SELECT BINLOG_GTID_POS ('hcpg-4-bin.000001', 329);  
+-----+-----+  
| BINLOG_GTID_POS('hcpg-2-bin.000001', 329) |  
+-----+-----+  
| NULL                                     |  
+-----+-----+
```

Step 5 - Enable replication

HCPG1

Note: The position information is from the 'hcp2', 'hcg3' and 'hcg4'

SELECT BINLOG_GTID_POS output

If the result is empty or NULL, enter an empty string '' (2 single quotes)

If the result is a string of numbers separated by dashes '-',

example: '2-1-14' from hcp2 and '3-1-14' from hcg3, the hcg4 result was empty or NULL.

Use the entire string from both hcp2 and hcg3 enclosed by single quotes '2-1-14,3-1-14', do not enter a value from hcg4 since the result was empty or NULL.

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = ''; ← This is for a NULL value (2 single quotes).
```

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = '2-1-14,3-1-14'; ← This is value of 2-1-14 and 3-1-14.
```

Note: The IP address below is for hcp2

```
MariaDB [(none)]> CHANGE MASTER 'node2' TO master_host="10.1.1.12", master_port=3306,  
master_user="replication_user", master_password='Organ1cReplication',  
master_use_gtid=slave_pos, master_ssl=1;
```

Note: The IP address below is for hcp3

```
MariaDB [(none)]> CHANGE MASTER 'node3' TO master_host="10.1.1.13", master_port=3306,  
master_user="replication_user", master_password='Organ1cReplication',  
master_use_gtid=slave_pos, master_ssl=1;
```

Note: The IP address below is for hcp4

```
MariaDB [(none)]> CHANGE MASTER 'node4' TO master_host="10.1.1.14", master_port=3306,  
master_user="replication_user", master_password='Organ1cReplication',  
master_use_gtid=slave_pos, master_ssl=1;
```

HCPG2

Note: The position information is from the 'hcg1', 'hcg3' and 'hcg4'

SELECT BINLOG_GTID_POS output

If the result is empty or NULL, enter an empty string '' (2 single quotes)

If the result is a string of numbers separated by dashes '-',

example: '1-1-14' from hcg1 and '3-1-14' from hcg3, the hcg2 and hcg4 result was empty or NULL.

Use the entire string from both hcg1 and hcg3 enclosed by single quotes '1-1-14,3-1-14', do not enter a value from hcg4 since the result was empty or NULL.

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = ''; ← This is for a NULL value (2 single quotes).
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = '1-1-14,3-1-14'; ← This is value of 1-1-14 and
3-1-14.
```

Note: The IP address below is for hcp1

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11", master_port=3306,
master_user="replication_user", master_password='Organ1cReplication',
master_use_gtid=slave_pos, master_ssl=1;
```

Note: The IP address below is for hcp3

```
MariaDB [(none)]> CHANGE MASTER 'node3' TO master_host="10.1.1.13", master_port=3306,
master_user="replication_user", master_password='Organ1cReplication',
master_use_gtid=slave_pos, master_ssl=1;
```

Note: The IP address below is for hcp4

```
MariaDB [(none)]> CHANGE MASTER 'node4' TO master_host="10.1.1.14", master_port=3306,
master_user="replication_user", master_password='Organ1cReplication',
master_use_gtid=slave_pos, master_ssl=1;
```

HCPG3

Note: The position information is from the 'hcp1', 'hcp2' and 'hcp4'

SELECT BINLOG_GTID_POS output

If the result is empty or NULL, enter an empty string " (2 single quotes)

If the result is a string of numbers separated by dashes '-',

example: '1-1-14' from hcp1 and '2-1-14' from hcp2, the hcp4 result was empty or NULL.

Use the entire string from both hcp1 and hcp2 enclosed by single quotes '1-1-14,2-1-14',

do not enter a value from hcp4 since the result was empty or NULL.

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = ''; ← This is for a NULL value (2 single quotes).
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = '1-1-14,2-1-14'; ← This is value of 1-1-14 and
2-1-14.
```

Note: The IP address below is for hcp1

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11", master_port=3306,
master_user="replication_user", master_password='Organ1cReplication',
master_use_gtid=slave_pos, master_ssl=1;
```

Note: The IP address below is for hcp2

```
MariaDB [(none)]> CHANGE MASTER 'node2' TO master_host="10.1.1.12", master_port=3306,  
master_user="replication_user", master_password='Organ1cReplication',  
master_use_gtid=slave_pos, master_ssl=1;
```

Note: The IP address below is for hcp4

```
MariaDB [(none)]> CHANGE MASTER 'node4' TO master_host="10.1.1.14", master_port=3306,  
master_user="replication_user", master_password='Organ1cReplication',  
master_use_gtid=slave_pos, master_ssl=1;
```

HCPG4

Note: The position information is from the 'hcpg1', 'hcpg2' and 'hcpg3' SELECT BINLOG_GTID_POS output

If the result is empty or NULL, enter an empty string " (2 single quotes)

If the result is a string of numbers separated by dashes '-',

example: '1-1-14' from hcp1 and '3-1-14' from hcp3, the hcp2 result was empty or NULL.

Use the entire string from both hcp1 and hcp3 enclosed by single quotes '1-1-14,3-1-14', do not enter a value from hcp2 since the result was empty or NULL.

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = ""; ← This is for a NULL value (2 single quotes).
```

```
MariaDB [(none)]> SET GLOBAL gtid_slave_pos = '1-1-14,3-1-14'; ← This is value of 1-1-14 and  
3-1-14.
```

Note: The IP address below is for hcp1

```
MariaDB [(none)]> CHANGE MASTER 'node1' TO master_host="10.1.1.11", master_port=3306,  
master_user="replication_user", master_password='Organ1cReplication',  
master_use_gtid=slave_pos, master_ssl=1;
```

Note: The IP address below is for hcp2

```
MariaDB [(none)]> CHANGE MASTER 'node2' TO master_host="10.1.1.12", master_port=3306,  
master_user="replication_user", master_password='Organ1cReplication',  
master_use_gtid=slave_pos, master_ssl=1;
```

Note: The IP address below is for hcp3

```
MariaDB [(none)]> CHANGE MASTER 'node3' TO master_host="10.1.1.13", master_port=3306,  
master_user="replication_user", master_password='Organ1cReplication',  
master_use_gtid=slave_pos, 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
```

Step 7 – Start the Windows Services Wildfly and SAM VFS on all nodes and then start the shares in the HCP Gateway UI on **HCPG1**.

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 fallback to primary Gateway, please contact Hitachi Vantara Support.

Hitachi Vantara



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
HitachiVantara.com | community.HitachiVantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
HitachiVantara.com/contact