

# Disaster recovery with the HP StorageWorks Enterprise File Services Clustered Gateway Linux Version and HP StorageWorks XP Continuous Access white paper



Executive summary.....	2
Clustered Gateway XP Continuous Access configuration procedure .....	3
Preparing the cluster for XP Continuous Access failovers .....	3
Clustered Gateway XP Continuous Access failover procedure .....	8
For more information.....	10

## Executive summary


Developing a disaster tolerant file sharing infrastructure is an important goal in today's business environment. As file sharing has always been a business-critical activity, providing users with immediate and consistent access to shared files is a significant goal. To ensure that this shared data will be available, even in the case of a disaster, the HP StorageWorks XP Business Copy and HP StorageWorks XP Continuous Access families of high-availability data and disaster recovery tools enable real-time data mirroring between XP disk arrays and protect you from catastrophic failures. To meet this objective, HP provides a proof point and best practice for file sharing disaster recovery solution.

This white paper describes the process to configure an HP StorageWorks Enterprise File Services Clustered Gateway when XP Continuous Access is utilized in the solution. This paper also describes the process to fail over from a Primary XP to a Secondary XP in this same solution environment.

This white paper provides a:

- Clustered Gateway XP Continuous Access solution overview
- Clustered Gateway XP Continuous Access configuration procedure
- Clustered Gateway XP Continuous Access failover procedure

**Figure 1.** Simple SAN diagram



# Clustered Gateway XP Continuous Access configuration procedure

It is important to remember that XP LUN UIDs on the Primary are different than those found on the Secondary. Therefore, it is extremely important to develop and save the corresponding P-VOL to S-VOL relationships for manual failover purposes.

Before you begin, be sure to consult the STREAMS documentation to verify that you have the latest supported drivers.

This white paper describes the method of disaster recovery using HP StorageWorks XP RAID Manager. It should be noted that the Web Console can also be used for this purpose, but is not described in this white paper.

Installation and configuration of XP RAID Manager is beyond the scope of this white paper.

## Preparing the cluster for XP Continuous Access failovers

- Verify that XP RAID Manager is installed, configured, and started on the Clustered Gateway. For ease of use, install XP RAID Manager on all Clustered Gateway nodes. Installation on all nodes provides high availability for XP RAID Manager.

XP RAID Manager configuration files include:

```
/etc/horcm.conf  
/etc/services
```

To start XP RAID Manager:

```
horcmstart.sh 0  
horcmstart.sh 1
```

To set the XP RAID Manager instance:

```
export HORCMINST=0
```

- Again, due to the fact that the UIDs are different on the Primary XP and Secondary XP, it is necessary to create a LUN UID P-VOL to S-VOL relationship. Now determine the LUN UIDs on the Primary XP that have been imported into the Clustered Gateway. To obtain the imported Primary LUN UIDs, use “sandiskinfo –ailq > your\_filename\_sandiskinfo\_primary.” The following shows an example output from this command.

```
Disk: /dev/psd/psd1          (Membership Disk)  
  Uid: 6-HP:OPEN-V:R500:00032334:01:BB SAN info: 192.168.1.252:27  
  Vendor:          HP OPEN-V Capacity: 1024.68M  
  Local Device Paths: /dev/sda  
    partition 01: size 1023.80M type Linux (83) (PSMP/Active )  
Disk: /dev/psd/psd2          (Membership Disk)  
  Uid: 6-HP:OPEN-V:R500:00032334:01:BC SAN info: 192.168.1.252:27  
  Vendor:          HP OPEN-V Capacity: 1024.68M  
  Local Device Paths: /dev/sdb  
    partition 01: size 1023.80M type Linux (83) (PSMP/Active )  
Disk: /dev/psd/psd3          (Membership Disk)  
  Uid: 6-HP:OPEN-V:R500:00032334:01:BD SAN info: 192.168.1.252:27
```

```

Vendor:          HP OPEN-V Capacity: 1024.68M
Local Device Paths: /dev/sdc
    partition 01: size 1023.80M type Linux (83) (PSMP/Active )
Disk: /dev/psd/psd4
Uid: 6-HP:OPEN-V:R500:00032334:01:9C SAN info: 192.168.1.252:27
Vendor:          HP OPEN-V Capacity: 136368.75M
Local Device Paths: /dev/sde
Disk /dev/psd/psd4 doesn't contain a valid partition table
    unpartitioned: size 136368.75M type (none) (PSFS Filesystem)
Disk: /dev/psd/psd5
Uid: 6-HP:OPEN-V:R500:00032334:01:8D SAN info: 192.168.1.252:27
Vendor:          Capacity: 136368.75M
Local Device Paths: /dev/sdd
    unpartitioned: size 136368.75M type (none) (SUBDEV/psv1)
Disk: /dev/psd/psd6
Uid: 6-HP:OPEN-V:R500:00032334:01:AB SAN info: 192.168.1.252:27
Vendor:          Capacity: 136368.75M
Local Device Paths: /dev/sdf
    unpartitioned: size 136368.75M type (none) (SUBDEV/psv1)

```

- Having determined the Primary XP LUN UIDs, now determine the Secondary LUN UIDs, which have not been imported into the Clustered Gateway. To obtain the unimported Secondary read-only LUN UIDs, use "sandiskinfo -aulq > your\_filename\_sandiskinfo\_secondary." This information will be used later to determine and verify the corresponding LUN UID P-VOL to S-VOL relationship. The following shows an example output from this command. Notice that the LUN UIDs are different than those in the preceding example:

```

Disk Uid: 6-HP:OPEN-V-CM:R500:00032334:01:C1
SAN info:      192.168.1.252:27
Vendor:        HP OPEN-V-CM Capacity: 100.31M
Local Device Paths: /dev/sdg
Disk /dev/sdg doesn't contain a valid partition table
    unpartitioned: size 100.31M type (none)
Disk Uid: 6-HP:OPEN-V-CM:R500:00032334:01:C2
SAN info:      192.168.1.252:27
Vendor:        HP OPEN-V-CM Capacity: 100.31M
Local Device Paths: /dev/sdh
Disk /dev/sdh doesn't contain a valid partition table
    unpartitioned: size 100.31M type (none)
Disk Uid: 6-HP:OPEN-V:R500:00032334:01:BE
SAN info:      192.168.1.252:26
Vendor:        HP OPEN-V Capacity: 1024.68M
Local Device Paths: /dev/sdi

```

```

    partition 01: size 1023.80M type Linux (83) (PSMP/Active )
Disk Uid: 6-HP:OPEN-V:R500:00032334:01:BF
SAN info:      192.168.1.252:26
Vendor:        HP OPEN-V Capacity: 1024.68M
Local Device Paths: /dev/sdj
    partition 01: size 1023.80M type Linux (83) (PSMP/Active )
Disk Uid: 6-HP:OPEN-V:R500:00032334:01:C0
SAN info:      192.168.1.252:26
Vendor:        HP OPEN-V Capacity: 1024.68M
Local Device Paths: /dev/sdk
    partition 01: size 1023.80M type Linux (83) (PSMP/Active )
Disk Uid: 6-HP:OPEN-V:R500:00032334:01:00
SAN info:      192.168.1.252:26
Vendor:        HP OPEN-V Capacity: 136368.75M
Local Device Paths: /dev/sdl
Disk /dev/sdl doesn't contain a valid partition table
unpartitioned: size 136368.75M type (none)
Disk Uid: 6-HP:OPEN-V:R500:00032334:01:01
SAN info:      192.168.1.252:26
Vendor:        HP OPEN-V Capacity: 136368.75M
Local Device Paths: /dev/sdm
Disk /dev/sdm doesn't contain a valid partition table
unpartitioned: size 136368.75M type (none) (PSFS Filesystem)
Disk Uid: 6-HP:OPEN-V:R500:00032334:01:02
SAN info:      192.168.1.252:26
Vendor:        HP OPEN-V Capacity: 136368.75M
Local Device Paths: /dev/sdn
Disk /dev/sdn doesn't contain a valid partition table
unpartitioned: size 136368.75M type (none)

```

- Now that the Primary XP LUN UIDs and the Secondary XP LUN UIDs have been imported, determine the P-VOL to S-VOL relationship between the corresponding UIDs. The tool used to determine this P-VOL to S-VOL relationship is the XP RAID Manager `pairedisplay` command. Using XP RAID Manager "`pairedisplay -g <assigned CA group name> -ICA0 -fx > your_filename_pairedisplay,`" view the CA configuration and determine the P-VOL to S-VOL assigned for the XP Continuous Access LUNs

---

**Note**

Use CA0, like in the preceding example, if HORCMINST is currently 0. Otherwise, use the instance number that you are currently running).

---

```

Group PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,Fence,Seq#,P-LDEV# M
CGWXPCA disk1(L) (CL7-R , 0, 0)32334 1bb.P-VOL PAIR NEVER ,32334 1be -
CGWXPCA disk1(R) (CL8-R , 0, 3)32334 1be.S-VOL PAIR NEVER ,---- 1bb -
CGWXPCA disk2(L) (CL7-R , 0, 1)32334 1bc.P-VOL PAIR NEVER ,32334 1bf -
CGWXPCA disk2(R) (CL8-R , 0, 4)32334 1bf.S-VOL PAIR NEVER ,---- 1bc -
CGWXPCA disk3(L) (CL7-R , 0, 2)32334 1bd.P-VOL PAIR NEVER ,32334 1c0 -
CGWXPCA disk3(R) (CL8-R , 0, 5)32334 1c0.S-VOL PAIR NEVER ,---- 1bd -
CGWXPCA disk4(L) (CL7-R , 0, 3)32334 18d.P-VOL PAIR NEVER ,32334 100 -
CGWXPCA disk4(R) (CL8-R , 0, 0)32334 100.S-VOL PAIR NEVER ,---- 18d -
CGWXPCA disk5(L) (CL7-R , 0, 4)32334 19c.P-VOL PAIR NEVER ,32334 101 -
CGWXPCA disk5(R) (CL8-R , 0, 1)32334 101.S-VOL PAIR NEVER ,---- 19c -
CGWXPCA disk6(L) (CL7-R , 0, 5)32334 1ab.P-VOL PAIR NEVER ,32334 102 -
CGWXPCA disk6(R) (CL8-R , 0, 2)32334 102.S-VOL PAIR NEVER ,---- 1ab -

```

- It is now important to determine how each of the imported XP LUN UUIDs is being used. To determine their use, ask questions like: Are they being used for membership partitions? Which psd devices are they assigned to? Are they being used for psv devices? To determine this utilization on one of the Clustered Gateway nodes, save the contents of the membership partition by executing `"/opt/hpcfs/lib/mpdump > your_file_name_mpdump_primary."` The following shows an example output.

```
Current Product MP Version: 2
```

```
Membership Partition Version: 2
```

```
Membership Partitions:
```

```
6-HP:OPEN-V:R500:00032334:01:BB/1 (ONLINE)
```

```
6-HP:OPEN-V:R500:00032334:01:BC/1 (ONLINE)
```

```
6-HP:OPEN-V:R500:00032334:01:BD/1 (ONLINE)
```

```
Membership Partition Device Database (Version 1):
```

```
UID:6-HP:OPEN-V:R500:00032334:01:BB Label:psd1 (state=0x1/mask=00000000)
```

```
UID:6-HP:OPEN-V:R500:00032334:01:BC Label:psd2 (state=0x1/mask=00000000)
```

```
UID:6-HP:OPEN-V:R500:00032334:01:BD Label:psd3 (state=0x1/mask=00000000)
```

```
UID:6-HP:OPEN-V:R500:00032334:01:9C Label:psd4 (state=0x1/mask=00000000)
```

```
UID:6-HP:OPEN-V:R500:00032334:01:8D Label:psd5 (state=0x1/mask=00000000)
```

```
UID:6-HP:OPEN-V:R500:00032334:01:AB Label:psd6 (state=0x1/mask=00000000)
```

```
Membership Partition Volume Database (Version 2):
```

```
VOL:psv1 (stripesize=1024K)
```

```
Set 0: SUBDEV: 6-HP:OPEN-V:R500:00032334:01:8D/0 size=139640832K
```

```
SUBDEV: 6-HP:OPEN-V:R500:00032334:01:AB/0 size=139640832K
```

```
Membership Partition Host Registry (Version 3):
```

```
Host ID: 10.1.1.90 state=unfenced fencetype=0
```

```
Fence ID:21:00:00:e0:8b:08:0f:ea::192.168.1.252 state=0
```

```
Fence ID:21:01:00:e0:8b:28:0f:ea::192.168.1.252 state=0
```

- Using information gathered before this step, determine the actual LUN UID P-VOL to S-VOL relationship. For example:
  - Using the `your_filename_pairdisplay`, note that there is a P-VOL to S-VOL relationship between 1bb and 1be, where 1bb is the Primary and 1be is the Secondary. Note the last two characters of the P-VOL to S-VOL relationship. In this case, bb is related to be:
 

```
CGWXPCA disk1(L) (CL7-R , 0, 0)32334 1bb.P-VOL PAIR NEVER ,32334 1be -
CGWXPCA disk1(R) (CL8-R , 0, 3)32334 1be.S-VOL PAIR NEVER ,--- 1bb -
```
  - Using the `your_filename_sandiskinfo_primary`, you have a LUN UID:
 

```
6-HP:OPEN-V:R500:00032334:01:BB
```
  - Using the `your_filename_sandiskinfo_secondary`, you have a LUN UID:
 

```
6-HP:OPEN-V:R500:00032334:01:BE
```
  - Now note that:
 

```
6-HP:OPEN-V:R500:00032334:01:BB relates to 6-HP:OPEN-V:R500:00032334:01:BE
```
  - Do this for each of the LUN UIDs found in the `file_name_secondary` file.
- Now build an mpdump file for the Secondary XP. To do this on the same node, copy the Primary mpdump file so that you can edit this with the Secondary information, “`cp your_file_name_mpdump_primary your_file_name_mpdump_secondary.`”
- Now create the Secondary mpdump file using the P-VOL to S-VOL relationship process described in the preceding step. Edit the `your_file_name_mpdump_secondary` file and modify each LUN UID so that is equivalent to the Secondary LUN UID, which in turn corresponds to the Primary LUN UID. Save this file after all LUN UIDs have been modified. In this example, the Secondary file will be as follows:

```
Current Product MP Version: 2
Membership Partition Version: 2
Membership Partitions:
6-HP:OPEN-V:R500:00032334:01:BE/1 (ONLINE)
6-HP:OPEN-V:R500:00032334:01:BF/1 (ONLINE)
6-HP:OPEN-V:R500:00032334:01:C0/1 (ONLINE)
Membership Partition Device Database (Version 1):
UID:6-HP:OPEN-V:R500:00032334:01:BE Label:psd1 (state=0x1/mask=00000000)
UID:6-HP:OPEN-V:R500:00032334:01:BF Label:psd2 (state=0x1/mask=00000000)
UID:6-HP:OPEN-V:R500:00032334:01:C0 Label:psd3 (state=0x1/mask=00000000)
UID:6-HP:OPEN-V:R500:00032334:01:01 Label:psd4 (state=0x1/mask=00000000)
UID:6-HP:OPEN-V:R500:00032334:01:00 Label:psd5 (state=0x1/mask=00000000)
UID:6-HP:OPEN-V:R500:00032334:01:02 Label:psd6 (state=0x1/mask=00000000)
Membership Partition Volume Database (Version 2):
VOL:psv1 (stripesize=1024K)
Set 0: SUBDEV: 6-HP:OPEN-V:R500:00032334:01:00/0 size=139640832K
SUBDEV: 6-HP:OPEN-V:R500:00032334:01:02/0 size=139640832K
Membership Partition Host Registry (Version 3):
Host ID: 10.1.1.90 state=unfenced fencetype=0
Fence ID:21:00:00:e0:8b:08:0f:ea::192.168.1.252 state=0
Fence ID:21:01:00:e0:8b:28:0f:ea::192.168.1.252 state=0
```

## Clustered Gateway XP Continuous Access failover procedure

Understandably, the Clustered Gateway service will fail if the nodes lose access to the Primary XP for any reason because storage is no longer available. This failure could include power loss, physical link interruptions, and switch failures, to name a few.

- After a disaster it is important to inform the XP Continuous Access Secondary array that it is now the Primary. This step changes the Secondary XP LUNs from Read-only to Read/Write so that the Clustered Gateway can utilize the LUNs for storage. Fail over to the backup LUNs executing the following command (using CA1 if horcm1.conf is your backup):

```
horctakeover -g <assigned CA group name> -ICA1
```

As a precautionary step, it is advised to check whether the takeover was actually successful. To confirm the takeover was successful, use `pairdisplay -g <assigned CA group name> -ICA1 -fx`. The LUNs that were previously P-VOLs should now be S-VOLs. The following shows an example output.

```
Group PairVol(L/R) (Port#,TID, LU),Seq#,LDEV#.P/S,Status,Fence,Seq#,P-LDEV# M
CGWXPCA disk1(L) (CL7-R , 0, 0)32334 1bb.S-VOL PAIR NEVER ,---- 1be -
CGWXPCA disk1(R) (CL8-R , 0, 3)32334 1be.P-VOL PAIR NEVER ,32334 1bb -
CGWXPCA disk2(L) (CL7-R , 0, 1)32334 1bc.S-VOL PAIR NEVER ,---- 1bf -
CGWXPCA disk2(R) (CL8-R , 0, 4)32334 1bf.P-VOL PAIR NEVER ,32334 1bc -
CGWXPCA disk3(L) (CL7-R , 0, 2)32334 1bd.S-VOL PAIR NEVER ,---- 1c0 -
CGWXPCA disk3(R) (CL8-R , 0, 5)32334 1c0.P-VOL PAIR NEVER ,32334 1bd -
CGWXPCA disk4(L) (CL7-R , 0, 3)32334 18d.S-VOL PAIR NEVER ,---- 100 -
CGWXPCA disk4(R) (CL8-R , 0, 0)32334 100.P-VOL PAIR NEVER ,32334 18d -
CGWXPCA disk5(L) (CL7-R , 0, 4)32334 19c.S-VOL PAIR NEVER ,---- 101 -
CGWXPCA disk5(R) (CL8-R , 0, 1)32334 101.P-VOL PAIR NEVER ,32334 19c -
CGWXPCA disk6(L) (CL7-R , 0, 5)32334 1ab.S-VOL PAIR NEVER ,---- 102 -
CGWXPCA disk6(R) (CL8-R , 0, 2)32334 102.P-VOL PAIR NEVER ,32334 1ab
```

- Because this is a disaster scenario, reboot all of the cluster nodes without starting the Clustered Gateway service. To do this, execute the following commands on each node:

```
chkconfig pmxs off
reboot the node
```

- After the node has rebooted, log on as the root since you must reconfigure the Clustered Gateway to read the new LUNs. This is necessary because XP Continuous Access failover utilizes different LUN UIDs on the Primary and Secondary storage arrays. First, load the HP CGW modules with the `pmxs load` command on all nodes. Next, execute the `mpimport` command on all nodes so that the device numbering remains static, the membership information will be rebuilt, and the configuration file, as defined in the configuration step, will be read. To do this, execute the following commands:

```
/etc/init.d/pmxs load
/opt/hpcfs/lib/mpimport -s -M -f your_file_name_mpdump_secondary
```

- After the `mpimport` command has been executed, the Clustered Gateway service can be started. To do this, execute the following command:

```
/etc/init.d/pmxs start
```

- At this point, be sure that the Clustered Gateway service starts at boot time. To do this, execute the following command:  
`chkconfig pmxs on`
- When the Primary LUNs are ready to be imported back into the cluster, first shut down the services on all nodes again:  
`/etc/init.d/pmxs stop`
- A reverse resynchronization must be done to update the P-VOLs with any new changes from the S-VOLs. To do this, execute the following command:  
`pairresync -g <assigned device group name> -restore`
- To import the P-VOLs into the cluster, replacing the S-VOLs, execute the following command:  
`/opt/hpcfs/lib/mpimport -s -M -f your_file_name_mpdump_primary`
- After the mpimport command has been executed, the Clustered Gateway service can be started. To do this, execute the following command on each node:  
`/etc/init.d/pmxs start`

## For more information

[www.hp.com/go/storage](http://www.hp.com/go/storage)

© 2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

4AA1-4359ENW, August 2007

