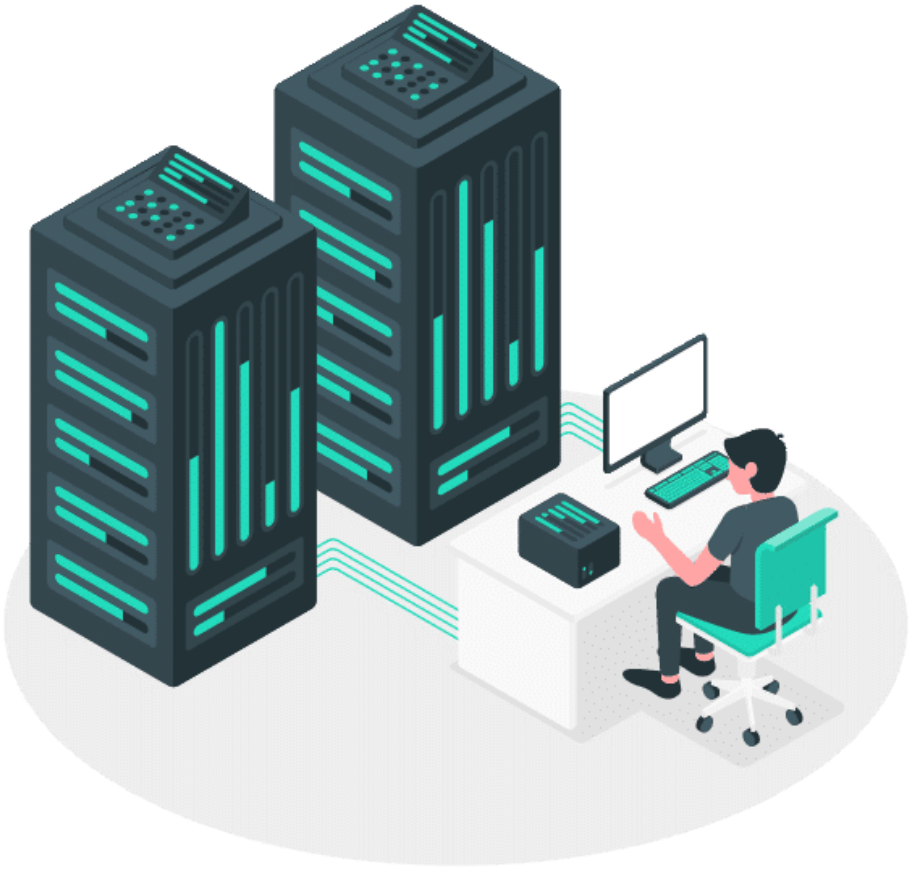

Haute Disponibilité

04.10.2024

—

BODNAR Georges

SOMMAIRE :

SOMMAIRE :..... 1
Installation de nginx.....1
Installation de phpmyadmin..... 2
Configuration de mariadb.....3
Installation de Wordpress.....4
Paramétrage du fichier wp-config.php..... 5
Paramétrage de PhpmyAdmin..... 8
Outils Haute Disponibilité(corosync et pacemaker) 9
Cloner la machine..... 10
Paramétrage et fonctionnement de corosync.....12
Attribution de ip virtuel..... 14
Réplication MySQL..... 16
Commande mysql maître..... 16
Commande mysql Esclave.....17
Commande Maître..... 17
Commande Esclave..... 17
En cas d'erreur..... 19
Cause probable :..... 19
Résolution :..... 19
1. Vérifier l'état actuel :..... 19

2. Identifier et ignorer l'erreur :..... 19
3. Si le problème persiste :.....20
4. Prévenir les erreurs à l'avenir :.....21
Utilisateur..... 21
Réplication Serveur de Fichiers DRBD.....22
Configuration de la ressource DRBD.....24
Configuration du système de fichiers.....24
Contraintes de colocation et ordre..... 24
Validation configuration DRBD.....24
Configuration du service FTP avec Pacemaker.....24
Tests suggérés..... 25
Gestion des erreurs et nettoyage.....25
Sauvegarde et restauration.....25
Gestion du Split Brain..... 25

Installation de nginx

Tout d'abord faut mettre a jour puis installer nginx :

apt upgrade

apt install nginx

```
Traitement des actions différées (« triggers ») pour man-db (2.11.2-2) ...
root@srv-bodnar:~# systemctl status nginx
● nginx.service - A high performance web server and a reverse proxy server
     Loaded: loaded (/lib/systemd/system/nginx.service; enabled; preset: enabled)
     Active: active (running) since Thu 2024-10-03 19:13:59 CEST; 57s ago
       Docs: man:nginx(8)
    Process: 1942 ExecStartPre=/usr/sbin/nginx -t -q -g daemon on; master_process on; (code=e>
    Process: 1943 ExecStart=/usr/sbin/nginx -g daemon on; master_process on; (code=exited, st>
   Main PID: 1966 (nginx)
      Tasks: 2 (limit: 2306)
     Memory: 1.7M
        CPU: 17ms
     CGroup: /system.slice/nginx.service
             ├─1966 "nginx: master process /usr/sbin/nginx -g daemon on; master_process on;"
             └─1967 "nginx: worker process"

oct. 03 19:13:59 srv-bodnar systemd[1]: Starting nginx.service - A high performance web serve>
oct. 03 19:13:59 srv-bodnar systemd[1]: Started nginx.service - A high performance web server>
```

Installation de phpmyadmin

apt install phpmyadmin

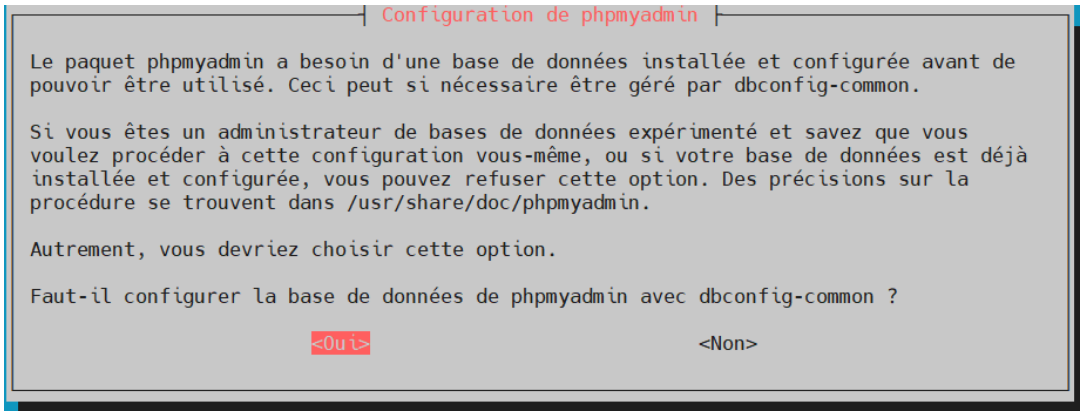


```
root@srv-bodnar:~# systemctl status mysql
● mariadb.service - MariaDB 10.11.6 database server
     Loaded: loaded (/lib/systemd/system/mariadb.service; enabled; preset: enabled)
     Active: active (running) since Thu 2024-10-03 19:06:26 CEST; 44min ago
       Docs: man:mariadbd(8)
             https://mariadb.com/kb/en/library/systemd/
   Main PID: 1557 (mariadbd)
     Status: "Taking your SQL requests now..."
      Tasks: 9 (limit: 2306)
     Memory: 213.1M
        CPU: 905ms
     CGroup: /system.slice/mariadb.service
             └─1557 /usr/sbin/mariadbd
```

Configuration de mariadb

D’abord faudra supprimer apache 2

rm -r apache2

Puis on rentre sur mysql

mysql -u root -p

Puis on tape copie colle ce texte

CREATE DATABASE wordpress; CREATE USER 'wordpress'@'localhost' IDENTIFIED BY 'sio'; GRANT ALL PRIVILEGES ON wordpress.* TO 'wordpress'@'localhost'; FLUSH PRIVILEGES;

EXIT;

Installation de Wordpress

Il faut d’abord télécharger le fichier de wordpress.

wget https://wordpress.org/latest.zip

Puis on le déplace dans :

mv latest.zip /var/www/html

On va pouvoir le désarchiver

apt install unzip

unzip latest.zip

On supprime le fichier désarchiver

rm latest.zip

On se déplace dans wordpress

cd wordpress

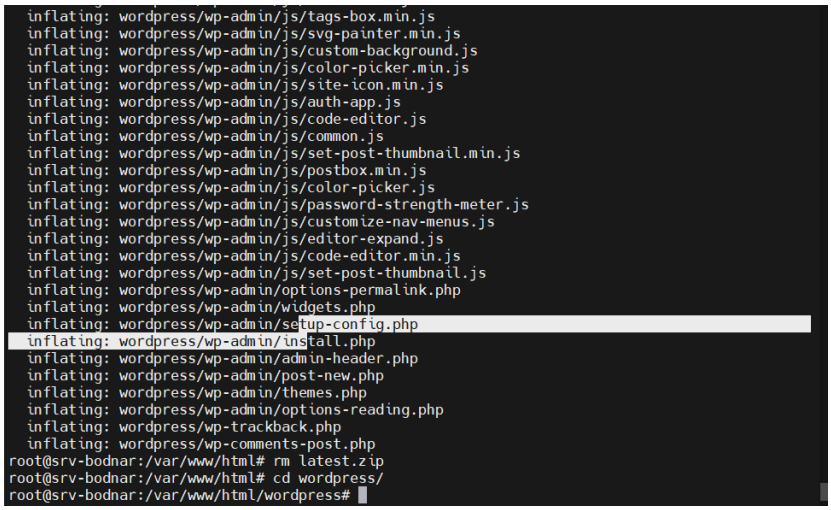

Après le transfert on va pouvoir déplacer mv wp-config.sample.php dans wp-config.php.

Si cela ne marche pas, on creer un fichier wp-config.php et après on fait le transfert.

```
 GNU nano 7.2                                  wp-config.php
<?php
/**
 * The base configuration for WordPress
 *
 * The wp-config.php creation script uses this file during the installation.
 * You don't have to use the website, you can copy this file to "wp-config.php"
 * and fill in the values.
 *
 * This file contains the following configurations:
 *
 * * Database settings
 * * Secret keys
 * * Database table prefix
 * * ABSPATH
 *
 * @link https://developer.wordpress.org/advanced-administration/wordpress/wp-config/
 *
 * @package WordPress
 */

// ** Database settings - You can get this info from your web host ** //
/** The name of the database for WordPress */
define( 'DB_NAME', 'database_name_here' );

/** Database username */
define( 'DB_USER', 'username_here' );
```

Paramétrage du fichier wp-config.php

```
// ** Database settings - You can get this info from your web host ** //
/** The name of the database for WordPress */
define( 'DB_NAME', 'wordpress' );

/** Database username */
define( 'DB_USER', 'wordpress' );

/** Database password */
define( 'DB_PASSWORD', 'toto' );

/** Database hostname */
define( 'DB_HOST', 'localhost' );

/** Database charset to use in creating database tables. */
define( 'DB_CHARSET', 'utf8' );

/** The database collate type. Don't change this if in doubt. */
define( 'DB_COLLATE', '' );

/**#@+
 * Authentication unique keys and salts.
 *
```

On recherche les lignes et on modifie les 3 defines.

On sauvegarde Ctrl+ O et ctrl+x.

```
        # SSL configuration
        #
        # listen 443 ssl default_server;
        # listen [::]:443 ssl default_server;
        #
        # Note: You should disable gzip for SSL traffic.
        # See: https://bugs.debian.org/773332
        #
        # Read up on ssl_ciphers to ensure a secure configuration.
        # See: https://bugs.debian.org/765782
        #
        # Self signed certs generated by the ssl-cert package
        # Don't use them in a production server!
        #
        # include snippets/snakeoil.conf;

        root /var/www/html/wordpress;

        # Add index.php to the list if you are using PHP
        index index.php;

        server_name _;

        location / {
                # First attempt to serve request as file, then
                # as directory, then fall back to displaying a 404.
                try_files $uri $uri/ =404;
        }
        #       location /phpmyadmin {
        #               root /usr/share;
        #               index index.php index.htm index.html;
        #               try_files $uri $uri/ /phpmyadmin/ index.php$query_string;
        #}
        # pass PHP scripts to FastCGI server
        #

        location ~ \.php$ {
                include snippets/fastcgi-php.conf;
#               try_files $uri $uri/ /phpmyadmin/ index.php$query_string;

        #       # With php-fpm (or other unix sockets):
                fastcgi_pass unix:/run/php/php8.2-fpm.sock;
#               fastcgi_param SCRIPT_FILENAME $document_root$fastcgi_script_name;
        #       # With php-cgi (or other tcp sockets):
        #       fastcgi_pass 127.0.0.1:9000;
        }

        # deny access to .htaccess files, if Apache's document root
        # concurs with nginx's one
        #
        #location ~ /\.ht {
        #       deny all;
        #}
}
```

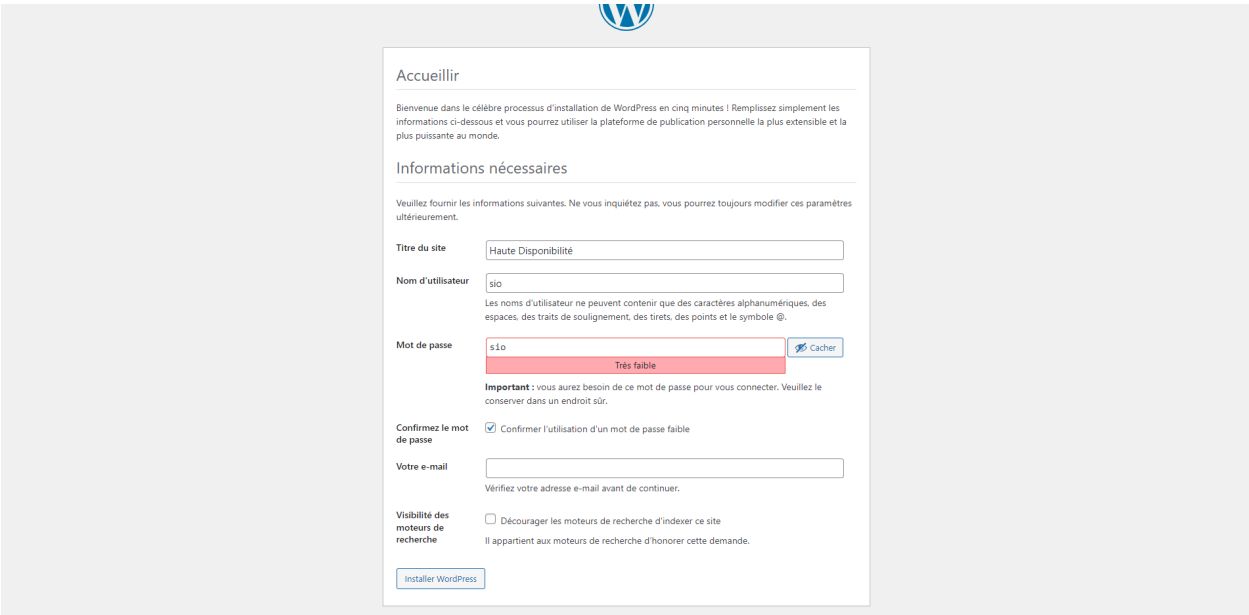

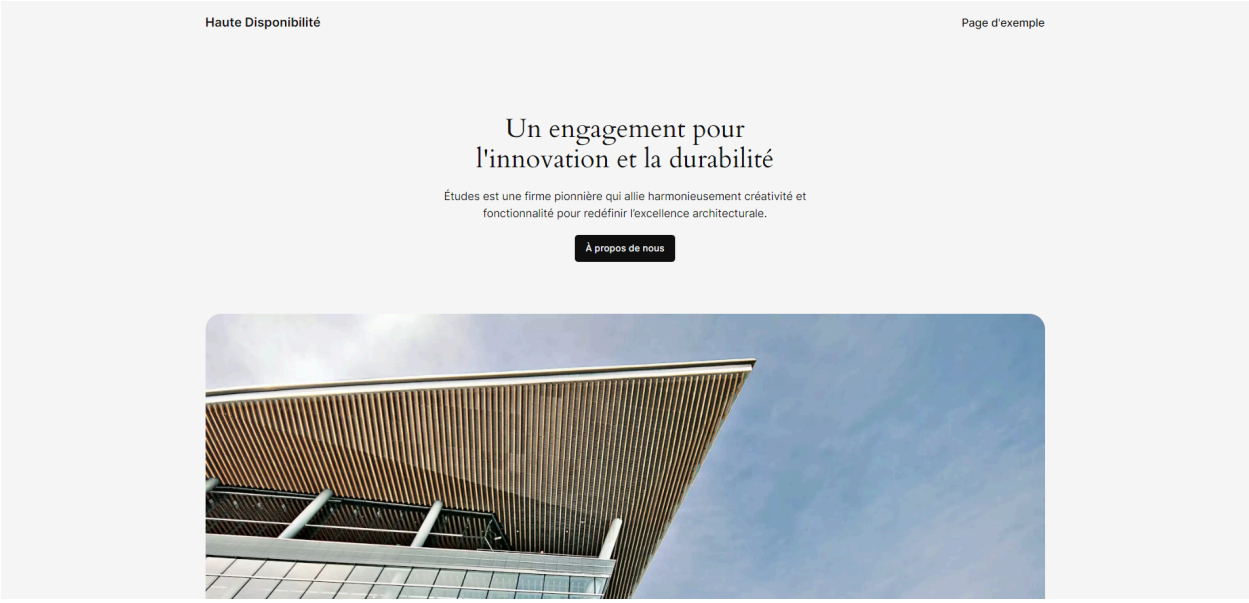
Haute Disponibilité
Page d'exemple
Un engagement pour
l'innovation et la durabilité
Études est une firme pionnière qui allie harmonieusement créativité et
fonctionnalité pour redéfinir l'excellence architecturale.
À propos de nous

Paramétrage de PhpmyAdmin

On va se diriger dans cette répertoire :

cd

```
        root /var/www/html;

        # Add index.php to the list if you are using PHP
        index index.php index.html index.htm index.nginx-debian.html;

        server_name;
        location / {
                # First attempt to serve request as file, then
                # as directory, then fall back to displaying a 404.
                try-files $uri $uri/ =404;
        }

        {
                # First attempt to serve request as file, then
          location /phpmyadmin {
                root /usr/share;
                index index.php index.html index.htm;
                try_files $uri $uri/ /phpmyadmin/index.php?$query_string;
                # pass PHP scripts to FastCGI server
                #
                location ~ \.php$ {
                try_files $uri $uri/ /phpmyadmin/index.php?$query_string;
                }
                **
        include snippets/fastcgi-php.conf;
        # With php-fpm (or other unix sockets):
        fastcgi_pass unix:/run/php/php8.2-fpm.sock;
        fastcgi_param SCRIPT_FILENAME $document_root$fastcgi_script_name;
        # With php-cgi (or other tcp sockets):
        #fastcgi_pass 127.0.0.1:9000;
        }

        # deny access to .htaccess files, if Apache's document root
        # concurs with nginx's one
        #
        #location ~ /\.ht {
        #       deny all;
        #}
}
```

Outils Haute Disponibilité(corosync et pacemaker)

apt install corosync pacemaker crmsh

corosync-keygen pour générer une clé

cp /etc/corosync/corosync.conf /etc/corosync/corosync.conf.save

Cette commande permet de sauvegarder le fichier de configuration.

On va se déplacer dans /etc/corosync/corosync.conf

On va pouvoir mettre nos ip.

```
# Please read the corosync.conf.5 manual page
totem {
        version: 2

        # Corosync itself works without a cluster name, but DLM needs one.
        # The cluster name is also written into the VG metadata of newly
        # created shared LVM volume groups, if lvmlockd uses DLM locking.
        cluster_name: bodnar

        # crypto_cipher and crypto_hash: Used for mutual node authentication.
        # If you choose to enable this, then do remember to create a shared
        # secret with "corosync-keygen".
        # enabling crypto_cipher, requires also enabling of crypto_hash.
        # crypto works only with knet transport
        crypto_cipher: none
        crypto_hash: none
}
logging {
        # Log the source file and line where messages are being
        # generated. When in doubt, leave off. Potentially useful for
        # debugging.
        fileline: off
        # Log to standard error. When in doubt, set to yes. Useful when
        # running in the foreground (when invoking "corosync -f")
        to_stderr: yes
        # Log to a log file. When set to "no", the "logfile" option
        # must not be set.
        to_logfile: yes
        logfile: /var/log/corosync/corosync.log
        # Log to the system log daemon. When in doubt, set to yes.
        to_syslog: yes
        # Log debug messages (very verbose). When in doubt, leave off.
        debug: off
        # Log messages with time stamps. When in doubt, set to hires (or on)
        #timestamp: hires
        logger_subsys {
                subsys: QUORUM
                debug: off
        }
}

quorum {
        # Enable and configure quorum subsystem (default: off)
        # see also corosync.conf.5 and votequorum.5
        provider: corosync_votequorum
}
```

```
quorum {
        # Enable and configure quorum subsystem (default: off)
        # see also corosync.conf.5 and votequorum.5
        provider: corosync_votequorum
}

nodelist {
        # Change/uncomment/add node sections to match cluster configuration

        node {
                # Hostname of the node
                name: bodnar-maitre
                # Cluster membership node identifier
                nodeid: 1
                # Address of first link
                ring0_addr: 172.31.4.21
                # When knet transport is used it's possible to define up to 8 links
                #ring1_addr: 192.168.1.1
        }
        node {
                # Hostname of the node
                name: bodnar-esclave
                # Cluster membership node identifier
                nodeid: 2
                # Address of first link
                ring0_addr: 172.31.4.22
                # When knet transport is used it's possible to define up to 8 links
                #ring1_addr: 192.168.1.1
        }

        # ...
}
```

```
oct. 11 11:28:17 srvBD-Maitre pacemaker-attrd[3033]:  notice: Pacemaker node attribute manager successfully started and accepting connections
oct. 11 11:28:17 srvBD-Maitre pacemaker-attrd[3033]:  notice: Setting #attrd-protocol[bodnar-maitre]: (unset) -> 4
oct. 11 11:28:17 srvBD-Maitre pacemaker-fenced[3031]:  notice: Pacemaker fencer successfully started and accepting connections
oct. 11 11:28:17 srvBD-Maitre pacemaker-attrd[3033]:  notice: Recorded local node as attribute writer (was unset)
oct. 11 11:28:17 srvBD-Maitre pacemaker-controld[3035]:  notice: Connecting to corosync cluster infrastructure
oct. 11 11:28:17 srvBD-Maitre pacemaker-controld[3035]:  warning: No quorum
oct. 11 11:28:17 srvBD-Maitre pacemaker-controld[3035]:  notice: Node bodnar-maitre state is now member
oct. 11 11:28:17 srvBD-Maitre pacemaker-controld[3035]:  notice: Pacemaker controller successfully started and accepting connections
oct. 11 11:28:17 srvBD-Maitre pacemaker-controld[3035]:  notice: State transition S_STARTING -> S_PENDING
oct. 11 11:28:18 srvBD-Maitre pacemaker-controld[3035]:  notice: Fencer successfully connected
root@srvBD-Maitre:~# nano /etc/corosync/corosync.conf
root@srvBD-Maitre:~# crm status
Status of pacemakerd: 'Pacemaker is running' (last updated 2024-10-11 11:29:41 +02:00)
Cluster Summary:
  * Stack: corosync
  * Current DC: bodnar-maitre (version 2.1.5-a3f44794f94) - partition WITHOUT quorum
  * Last updated: Fri Oct 11 11:29:41 2024
  * Last change:  Fri Oct 11 11:28:17 2024 by hacluster via crmd on bodnar-maitre
  * 2 nodes configured
  * 0 resource instances configured

Node List:
  * Node bodnar-esclave: UNCLEAN (offline)
  * Online: [ bodnar-maitre ]

Full List of Resources:
  * No resources

root@srvBD-Maitre:~#
```

Sauvegarde et on sort d'ici.

Cloner la machine

Après avoir fini on éteint la machine.

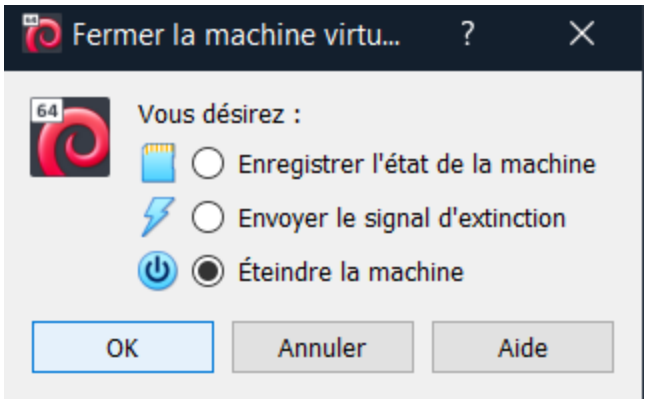

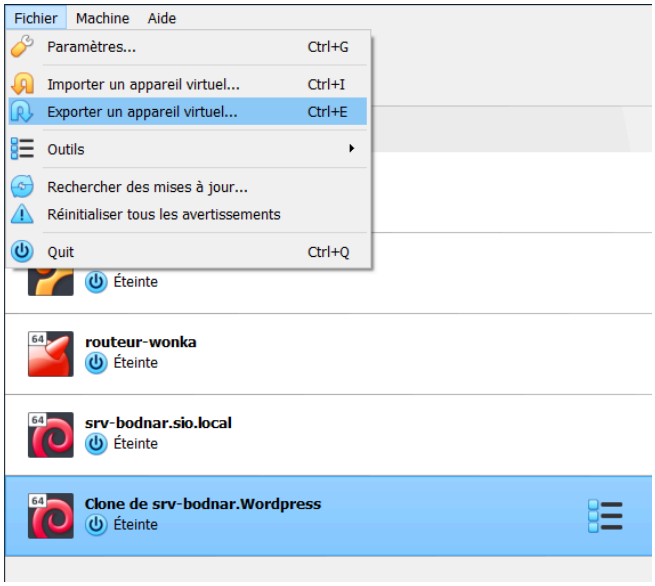
Fichier
Machine
Aide
Paramètres... Ctrl+G
Importer un appareil virtuel... Ctrl+I
Exporter un appareil virtuel... Ctrl+E
Outils
Rechercher des mises à jour...
Réinitialiser tous les avertissements
Quit Ctrl+Q
Éteinte
routeur-wonka
Éteinte
srv-bodnar.sio.local
Éteinte
Clone de srv-bodnar.Wordpress
Éteinte

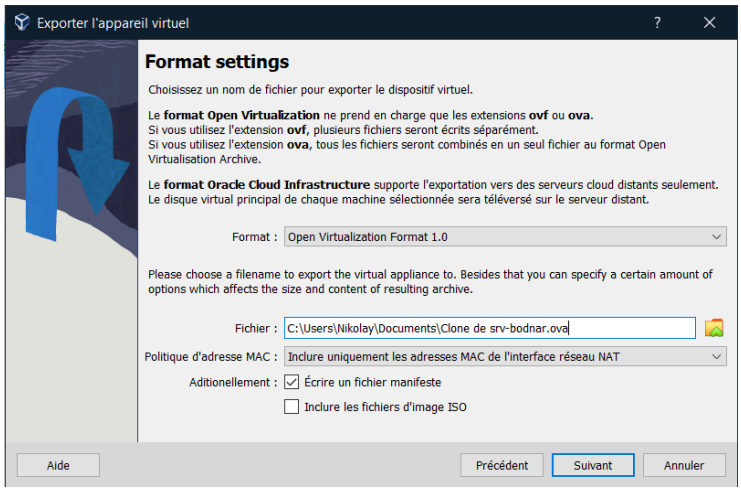
Exporter l'appareil virtuel
Format settings
Choisissez un nom de fichier pour exporter le dispositif virtuel.
Le format Open Virtualization ne prend en charge que les extensions ovf ou ova.
Si vous utilisez l'extension ovf, plusieurs fichiers seront écrits séparément.
Si vous utilisez l'extension ova, tous les fichiers seront combinés en un seul fichier au format Open Virtualisation Archive.
Le format Oracle Cloud Infrastructure supporte l'exportation vers des serveurs cloud distants seulement.
Le disque virtuel principal de chaque machine sélectionnée sera téléversé sur le serveur distant.
Format : Open Virtualization Format 1.0
Please choose a filename to export the virtual appliance to. Besides that you can specify a certain amount of options which affects the size and content of resulting archive.
Fichier : C:\Users\Nikolay\Documents\Clone de srv-bodnar.ova
Politique d'adresse MAC : Inclure uniquement les adresses MAC de l'interface réseau NAT
Additionellement : Écrire un fichier manifeste
Inclure les fichiers d'image ISO
Aide
Précédent
Suivant
Annuler

Paramétrage et fonctionnement de corosync

```
totem {
        version: 2

        # Corosync itself works without a cluster name, but DLM needs one.
        # The cluster name is also written into the VG metadata of newly
        # created shared LVM volume groups, if lvmlockd uses DLM locking.
        cluster_name: bodnar

        # crypto_cipher and crypto_hash: Used for mutual node authentication.
        # If you choose to enable this, then do remember to create a shared
        # secret with "corosync-keygen".
        # enabling crypto_cipher, requires also enabling of crypto_hash.
        # crypto works only with knet transport
        crypto_cipher: aes256
        crypto_hash: sha1
}
logging {
        # Log the source file and line where messages are being
        # generated. When in doubt, leave off. Potentially useful for
        # debugging.
        fileline: off
        # Log to standard error. When in doubt, set to yes. Useful when
        # running in the foreground (when invoking "corosync -f")
        to_stderr: yes
        # Log to a log file. When set to "no", the "logfile" option
        # must not be set.
        to_logfile: yes
        logfile: /var/log/corosync/corosync.log
        # Log to the system log daemon. When in doubt, set to yes.
        to_syslog: yes
        # Log debug messages (very verbose). When in doubt, leave off.
        debug: off
        # Log messages with time stamps. When in doubt, set to hires (or on)
        #timestamp: hires
        logger_subsys {
                subsys: QUORUM
                debug: off
        }
}

quorum {
        # Enable and configure quorum subsystem (default: off)
        # see also corosync.conf.5 and votequorum.5
        provider: corosync_votequorum
        expected_votes: 2
        two_nodes: 1
}
```

On doit modifier crypto_cipher et hash : [par défaut none]

Puis ajouter 2 lignes dans le quorum

expected_votes: 2

two_nodes: 1

Puis on redémarre avec systemctl restart corosync

```
root@srvBD-Maitre:~# crm status
Status of pacemakerd: 'Pacemaker is running' (last updated 2024-11-08 10:39:31 +01:00)
Cluster Summary:
  * Stack: corosync
  * Current DC: bodnar-maitre (version 2.1.5-a3f44794f94) - partition with quorum
  * Last updated: Fri Nov  8 10:39:32 2024
  * Last change:  Fri Nov  8 10:29:44 2024 by root via cibadmin on bodnar-maitre
  * 2 nodes configured
  * 1 resource instance configured

Node List:
  * Online: [ bodnar-esclave bodnar-maitre ]

Full List of Resources:
  * IPFailover  (ocf:heartbeat:IPaddr2):         Started bodnar-esclave
```

Cette partie on doit faire sur les 2 machines

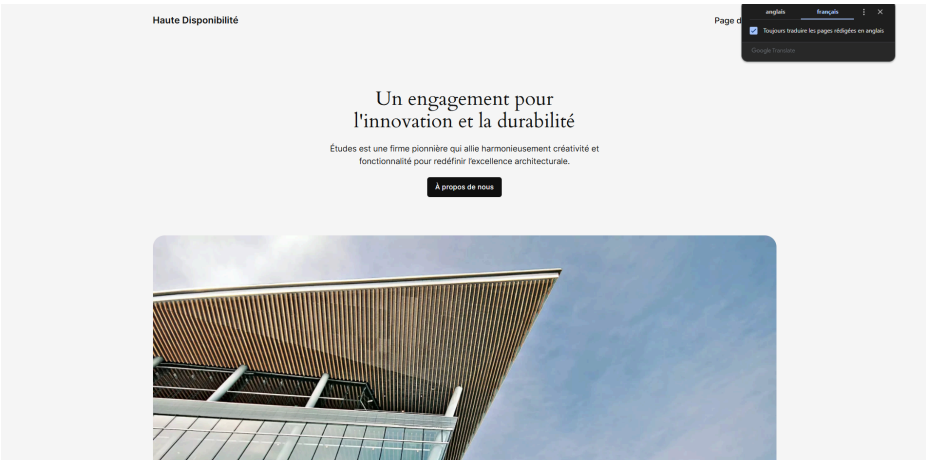

Attribution de ip virtuel

```
root@srvBD-Maitre:~# crm configure edit
root@srvBD-Maitre:~# export EDITOR=nano
root@srvBD-Maitre:~# crm configure edit
WARNING: (unpack_config)        warning: Blind faith: not fencing unseen nodes
root@srvBD-Maitre:~# systemctl restart corosync.service
root@srvBD-Maitre:~# crm status
Status of pacemakerd: 'Pacemaker is running' (last updated 2024-11-08 11:21:00 +01:00)
Cluster Summary:
  * Stack: corosync
  * Current DC: bodnar-esclave (version 2.1.5-a3f44794f94) - partition with quorum
  * Last updated: Fri Nov  8 11:21:01 2024
  * Last change:  Fri Nov  8 11:20:44 2024 by root via cibadmin on bodnar-maitre
  * 2 nodes configured
  * 1 resource instance configured

Node List:
  * Online: [ bodnar-esclave bodnar-maitre ]

Full List of Resources:
  * IPFailover  (ocf:heartbeat:IPaddr2):         Started bodnar-esclave

root@srvBD-Maitre:~# crm resource move IPFailover bodnar-maitre
INFO: Move constraint created for IPFailover to bodnar-maitre
INFO: Use `crm resource clear IPFailover` to remove this constraint
root@srvBD-Maitre:~# crm status
Status of pacemakerd: 'Pacemaker is running' (last updated 2024-11-08 11:21:25 +01:00)
Cluster Summary:
  * Stack: corosync
  * Current DC: bodnar-esclave (version 2.1.5-a3f44794f94) - partition with quorum
  * Last updated: Fri Nov  8 11:21:25 2024
  * Last change:  Fri Nov  8 11:21:21 2024 by root via crm_resource on bodnar-maitre
  * 2 nodes configured
  * 1 resource instance configured

Node List:
  * Online: [ bodnar-esclave bodnar-maitre ]

Full List of Resources:
  * IPFailover  (ocf:heartbeat:IPaddr2):         Started bodnar-maitre
```

```
  GNU nano 7.2
node 1: bodnar-maitre
node 2: bodnar-esclave
primitive IPFailover IPaddr2 \
        params ip=172.31.4.20 cidr_netmask=21 nic=enp0s3 iflabel=VIP \
        op monitor interval=30s timeout=20s
location cli-prefer-IPFailover IPFailover role=Started inf: bodnar-maitre
property cib-bootstrap-options: \
        have-watchdog=false \
        dc-version=2.1.5-a3f44794f94 \
        cluster-infrastructure=corosync \
        cluster-name=debian \
        stonith-enabled=false
```

Ici on devra ajouter le texte au dessus, ces commandes nous permettent de créer une IP virtuelle.

-crm resource move IPFailover abderma-maître | On met le maître en première position.

```
root@srvBD-Maitre:~# crm configure show
node 1: bodnar-maitre
node 2: bodnar-esclave
primitive IPFailover IPaddr2 \
        params ip=172.31.4.20 cidr_netmask=21 nic=enp0s3 iflabel=VIP \
        op monitor interval=30s timeout=20s
location cli-prefer-IPFailover IPFailover role=Started inf: bodnar-maitre
property cib-bootstrap-options: \
        have-watchdog=false \
        dc-version=2.1.5-a3f44794f94 \
        cluster-infrastructure=corosync \
        cluster-name=debian \
        stonith-enabled=false
root@srvBD-Maitre:~#
```

Réplication MySQL

Modifié le fichier sur maître : /etc/mysql/mariadb.conf.d/50-server.cnf

[mysqld]

#bind-address = 172.31.4.21

server_id = 101

log_bin = /var/log/mysql/mysql-bin.log

expire_logs_days = 10

max_binlog_size = 100M

binlog_do_db = wordpress

Modifié le fichier sur esclave : /etc/mysql/mariadb.conf.d/50-server.cnf

[mysqld]

bind-address = 172.31.4.22

server_id = 102

expire_logs_days = 10

max_binlog_size = 100M

master_retry_count = 20

replicate_do_db = wordpress

Enregistrer

Commande mysql maître

On rentre sur mysql : mysql -u root -p

On ajoute l'utilisateur : ADD USER exemple;

On donne les privilèges : GRANT ALL PRIVILEGES on *.* TO 'sio'@'%' IDENTIFIED BY 'sio';

flush privileges;

Puis on tape ces commande :

FLUSH TABLES WITH READ LOCK;

show master status;

```
MariaDB [(none)]> show master status;
+------------------+----------+--------------+------------------+
| File             | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+------------------+----------+--------------+------------------+
| mysql-bin.000011 |      342 | wordpress    |    |
+------------------+----------+--------------+------------------+
```

quit

Commande mysql Esclave

mysql -u sio -p

mysql> stop slave;

mysql> change master to

-> master_host='172.31.4.21',

-> master_user='sio',

-> master_password= ‘sio',

-> master_log_file='mysql-bin.000001',

-> master_log_pos=3921;

mysql> start slave;

Commande Maître

UNLOCK TABLES;

Commande Esclave

GRANT ALL PRIVILEGES ON *.* TO 'sio'@'%' IDENTIFIED BY 'sio';

mysql> show slave status \G;

```
MariaDB [(none)]> GRANT ALL PRIVILEGES ON *.* TO 'sio'@'%' IDENTIFIED BY 'sio';
ERROR 1045 (28000): Access denied for user 'sio'@'%' (using password: YES)
MariaDB [(none)]> show slave status \G
*************************** 1. row ***************************
               Slave_IO_State: Waiting for master to send event
                  Master_Host: 172.31.4.21
                  Master_User: sio
                  Master_Port: 3306
                Connect_Retry: 60
              Master_Log_File: mysql-bin.000011
          Read_Master_Log_Pos: 342
               Relay_Log_File: mysqld-relay-bin.000007
                Relay_Log_Pos: 555
        Relay_Master_Log_File: mysql-bin.000011
             Slave_IO_Running: Yes
            Slave_SQL_Running: Yes
          Replicate_Rewrite_DB:
              Replicate_Do_DB: wordpress
          Replicate_Ignore_DB:
           Replicate_Do_Table:
       Replicate_Ignore_Table:
      Replicate_Wild_Do_Table:
  Replicate_Wild_Ignore_Table:
                   Last_Errno: 0
                   Last_Error:
                 Skip_Counter: 0
          Exec_Master_Log_Pos: 342
              Relay_Log_Space: 865
              Until_Condition: None
               Until_Log_File:
                Until_Log_Pos: 0
           Master_SSL_Allowed: No
           Master_SSL_CA_File:
           Master_SSL_CA_Path:
              Master_SSL_Cert:
            Master_SSL_Cipher:
               Master_SSL_Key:
        Seconds_Behind_Master: 0
Master_SSL_Verify_Server_Cert: No
                Last_IO_Errno: 0
                Last_IO_Error:
               Last_SQL_Errno: 0
               Last_SQL_Error:
  Replicate_Ignore_Server_Ids:
             Master_Server_Id: 101
               Master_SSL_Crl:
           Master_SSL_Crlpath:
                   Using_Gtid: No
                  Gtid_IO_Pos:
      Replicate_Do_Domain_Ids:
  Replicate_Ignore_Domain_Ids:
                Parallel_Mode: optimistic
                    SQL_Delay: 0
          SQL_Remaining_Delay: NULL
      Slave_SQL_Running_State: Slave has read all relay log; waiting for more updates
             Slave_DDL_Groups: 0
Slave_Non_Transactional_Groups: 0
    Slave_Transactional_Groups: 0
1 row in set (0,001 sec)
```

En cas d'erreur

Le message d'erreur dans votre résultat indique un problème lié à une **clé primaire dupliquée** lors de l'exécution d'un événement de réplication sur l'esclave. Voici les détails et les étapes pour résoudre le problème.

Cause probable :

La réplication échoue parce que l'esclave tente d'insérer une entrée dans la table `wordpress.wp_options` avec une valeur de clé primaire `282` qui existe déjà. Cela se produit souvent lorsque :

1. L'esclave n'a pas été correctement synchronisé avec le maître au départ (des données incohérentes existent déjà).
2. Une opération conflictuelle a été exécutée directement sur l'esclave, entraînant une divergence avec le maître.

Résolution :

1. Vérifier l'état actuel :

Avant de corriger, assurez-vous que le problème ne vient que de cette erreur. Exécutez :

SHOW SLAVE STATUS\G;

- Confirmez que **`Slave_SQL_Running`** est sur `No` et que l'erreur mentionnée correspond bien à un conflit de clé primaire (erreur 1062).

---

2. Identifier et ignorer l'erreur :

Vous pouvez ignorer l'erreur spécifique pour avancer dans la réplication en exécutant les commandes suivantes sur l'esclave :

STOP SLAVE;

SET GLOBAL sql_slave_skip_counter = 1;

START SLAVE;

- **`SET GLOBAL sql_slave_skip_counter = 1;`** : Ignore l'événement problématique dans le journal binaire.

Ensuite, vérifiez l'état de la réplication :
SHOW SLAVE STATUS\G;

- Vérifiez que **`Slave_SQL_Running`** est passé à `Yes`.

---

3. Si le problème persiste :

Si l'erreur se reproduit ou si la réplication est bloquée sur plusieurs événements similaires :

- **Resynchronisez l'esclave avec le maître :**
    1. Sauvegardez les données actuelles de l'esclave si nécessaire.
    2. Exécutez une nouvelle copie des données depuis le maître vers l'esclave pour garantir une cohérence totale.
- **Exemple d'une resynchronisation manuelle :**

Arrêtez l'esclave :
STOP SLAVE;

1.
2. Sur le maître, créez une sauvegarde cohérente (avec `mysqldump` ou un outil similaire).
3. Importez la sauvegarde sur l'esclave.
4. Configurez à nouveau la réplication en spécifiant les bons paramètres (`CHANGE MASTER TO`).

```
MariaDB [(none)]> STOP SLAVE;
Query OK, 0 rows affected (0,008 sec)

MariaDB [(none)]> CHANGE MASTER TO
    -> master_host='172.31.4.21',
    -> master_user='sio',
    -> master_password='sio',
    -> master_log_file='mysql-bin.000012',
    -> master_log_pos=827978;
Query OK, 0 rows affected, 1 warning (0,025 sec)

MariaDB [(none)]> start slave;
Query OK, 0 rows affected (0,003 sec)

MariaDB [(none)]> SHOW SLAVE STATUS\G;
*************************** 1. row ***************************
               Slave_IO_State: Waiting for master to send event
                  Master_Host: 172.31.4.21
                  Master_User: sio
                  Master_Port: 3306
                Connect_Retry: 60
              Master_Log_File: mysql-bin.000012
          Read_Master_Log_Pos: 827978
               Relay_Log_File: mysqld-relay-bin.000002
                Relay_Log_Pos: 555
        Relay_Master_Log_File: mysql-bin.000012
             Slave_IO_Running: Yes
            Slave_SQL_Running: Yes
          Replicate_Rewrite_DB:
              Replicate_Do_DB: wordpress
          Replicate_Ignore_DB:
           Replicate_Do_Table:
       Replicate_Ignore_Table:
      Replicate_Wild_Do_Table:
```

---

4. Prévenir les erreurs à l'avenir :

- Évitez d'effectuer des opérations d'écriture directement sur l'esclave.
- Utilisez des options de réplication comme **GTID** ou **ROW** pour améliorer la fiabilité.
- Activez l'option `IGNORE_SERVER_IDS` si certains serveurs doivent être exclus des conflits.

Utilisateur

username:sio | mdp:sio

username:toto | mdp:toto

Réplication Serveur de Fichiers DRBD

Installer l'outil avec

apt install drbd-utils

Pour vérifier qu'il soit bien chargé : `lsmod | grep drbd`

```
root@srvBD-Maitre:~# modinfo drbd
filename:       /lib/modules/6.1.0-25-amd64/kernel/drivers/block/drbd/drbd.ko
alias:          block-major-147-*
license:        GPL
version:        8.4.11
description:    drbd - Distributed Replicated Block Device v8.4.11
author:         Philipp Reisner <phil@linbit.com>, Lars Ellenberg <lars@linbit.com>
srcversion:     19D914EA50F713FCCE48607
depends:        lru_cache,libcrc32c
retpoline:      Y
intree:         Y
name:           drbd
vermagic:       6.1.0-25-amd64 SMP preempt mod_unload modversions
sig_id:         PKCS#7
signer:         Debian Secure Boot CA
sig_key:        32:A0:28:7F:84:1A:03:6F:A3:93:C1:E0:65:C4:3A:E6:B2:42:26:43
sig_hashalgo:   sha256
signature:      84:48:23:8A:C8:C2:52:6B:94:E4:42:23:0D:89:56:CC:7B:EC:71:FB:
                E9:FB:19:BE:82:20:3A:72:26:72:97:26:BF:6F:E2:1E:A2:98:BC:35:
                9D:33:BB:98:CD:4D:20:D1:E0:D4:DF:EE:22:62:91:A7:1F:DA:E8:77:
                04:B1:A5:26:32:6F:C9:BC:62:30:F7:51:4C:C8:72:0A:B9:0C:96:E1:
                82:D7:B8:46:AC:31:F0:B3:ED:A5:92:EB:4E:FE:6D:D9:A9:5F:0D:70:
                A1:6E:51:6D:A9:A6:37:C4:94:1D:96:97:6F:F4:79:CA:79:F4:6E:CF:
                F1:64:0F:37:3F:6D:7C:48:56:98:70:21:50:AA:98:BE:6E:21:71:32:
                30:22:95:C0:04:C5:52:C4:D5:51:7C:B9:55:B3:03:B7:CA:2C:38:62:
                49:8B:79:C0:56:80:32:7F:79:B9:48:48:E8:97:C1:CA:30:D6:57:6F:
                0A:BA:87:E6:31:4E:57:19:8B:64:24:C1:37:2D:37:0C:62:BD:D3:1D:
                C6:45:17:AB:BB:98:05:51:2D:F4:E0:F1:BD:C1:6A:30:A8:7A:E8:C2:
                38:1D:9F:14:39:66:29:2D:9B:92:35:10:EA:E8:A7:47:73:15:C3:8E:
                10:24:EC:03:28:0A:C4:C5:8F:E3:F6:D4:3E:03:24:3A
parm:           allow_oos:DONT USE! (bool)
parm:           disable_sendpage:bool
parm:           proc_details:int
parm:           minor_count:Approximate number of drbd devices (1-255) (uint)
parm:           usermode_helper:string
root@srvBD-Maitre:~#
```

Pour savoir des informations sur le module.

Savoir sa version

```
root@srvBD-Maitre:~# cat /proc/drbd
cat: /proc/drbd: Aucun fichier ou dossier de ce type
root@srvBD-Maitre:~# systemctl status drbd
○ drbd.service - DRBD -- please disable. Unless you are NOT using a cluster manager.
     Loaded: loaded (/lib/systemd/system/drbd.service; disabled; preset: enabled)
     Active: inactive (dead)
root@srvBD-Maitre:~# version drbd.
-bash: version : commande introuvable
root@srvBD-Maitre:~# version drbd
-bash: version : commande introuvable
root@srvBD-Maitre:~# version drbd.service
-bash: version : commande introuvable
root@srvBD-Maitre:~# version nginx
-bash: version : commande introuvable
root@srvBD-Maitre:~# drbdadm --version
DRBDADM_BUILDTAG=GIT-hash:\ 409097fe02187f83790b88ac3e0d94f3c167adab\ reproducible\ build\,\ 2023-01-09\ 14:51:18
DRBDADM_API_VERSION=1
DRBD_KERNEL_VERSION_CODE=0x000000
DRBDADM_VERSION_CODE=0x091600
DRBDADM_VERSION=9.22.0
root@srvBD-Maitre:~#
```

```
losetup: /drbd-data/drbd.disk : échec de con
root@srvBD-Maitre:~# lsblk
NAME   MAJ:MIN RM  SIZE RO TYPE MOUNTPOINTS
sda      8:0    0   20G  0 disk
├─sda1   8:1    0   19G  0 part /
├─sda2   8:2    0    1K  0 part
└─sda5   8:5    0  975M  0 part [SWAP]
sdb      8:16   0    5G  0 disk
sr0     11:0    1 1024M  0 rom
root@srvBD-Maitre:~# fdisk /dev/sdb
```

Après avoir ajouter 1 disque il faut aller dans

nano /etc/drbd.d/resftp.res

```
  GNU nano 7.2                                                    /etc/drbd.d/resftp.res
resource resftp {
  protocol C;

  on srvBD-Maitre {
    device    /dev/drbd0;       # Périphérique DRBD
    disk      /dev/sdb;       # Partition physique
    address   172.31.4.21:7788;  # Adresse IP et port
    meta-disk internal;        # Métadonnées internes
  }

  on bodnar-esclave {
    device    /dev/drbd0;
    disk      /dev/sdb;
    address   172.31.4.22:7788;
    meta-disk internal;
  }
}
```

Puis taper drbdadm create-md resftp

```
root@srvBD-Maitre:~# lsblk
NAME   MAJ:MIN RM  SIZE RO TYPE MOUNTPOINTS
sda      8:0    0   20G  0 disk
├─sda1   8:1    0   19G  0 part /
├─sda2   8:2    0    1K  0 part
└─sda5   8:5    0  975M  0 part [SWAP]
sdb      8:16   0    5G  0 disk
sr0     11:0    1 1024M  0 rom
root@srvBD-Maitre:~# nano /etc/drbd.d/resftp.res
root@srvBD-Maitre:~# drbdadm adjust all
No valid meta data found
root@srvBD-Maitre:~# drbdadm adjust all
No valid meta data found
root@srvBD-Maitre:~# drbdadm create-md resftp
md_offset 5368705024
al_offset 5368672256
bm_offset 5368508416

Found some data

 ==> This might destroy existing data! <==

Do you want to proceed?
[need to type 'yes' to confirm] yes

initializing activity log
initializing bitmap (160 KB) to all zero
Writing meta data...
New drbd meta data block successfully created.
root@srvBD-Maitre:~# nano /etc/drbd.d/resftp.res
root@srvBD-Maitre:~#
```

https://marouanromualdo.wordpress.com/wp-content/uploads/2015/06/heartbeat.pdf

Configuration de la ressource DRBD

crm configure

primitive drbd0 ocf:linbit:drbd params drbd_resource=resftp op monitor role=Master interval=60s timeout=30s op monitor role=Slave interval=60s timeout=30s

ms ms-drbd0 drbd0 meta clone-max=2 clone-node-max=1 master-max=1 master-node-max=1 globally-unique=false notify=true

Configuration du système de fichiers

primitive fs0 ocf:heartbeat:Filesystem params fstype=ext4 directory=/home device=/dev/drbd0

Contraintes de colocation et ordre

colocation fs0-on-ms-drbd0 inf: fs0 ms-drbd0:Master

order ms-drbd0-before-fs0 mandatory: ms-drbd0:promote fs0:start

location ms-drbd0-master-on-intralabar ms-drbd0 rule role=master inf: #uname eq intralabar

Validation configuration DRBD

commit

crm_mon

Configuration du service FTP avec Pacemaker

crm configure primitive serviceFTP lsb:proftpd op monitor interval=30s

crm configure colocation ftp-on-fs0 inf: serviceFTP fs0

crm configure order fs0-before-ftp mandatory: fs0:start serviceFTP:start

crm configure colocation ftp-on-web inf: serviceFTP IPFailover serviceHTTP

commit

crm_mon

Tests suggérés

crm status

ifconfig

touch /home/user/Documents/essai1

crm node standby NOM_NŒUD_MAITRE

crm node online NOM_NŒUD_MAITRE

Gestion des erreurs et nettoyage

crm resource cleanup <id_resource>

crm_resource -P

Sauvegarde et restauration

cibadmin --query > sauvegarde_cluster.xml

cibadmin --replace --xml-file sauvegarde_cluster.xml

Gestion du Split Brain

umount /dev/drbd0

drbdadm disconnect resftp

drbdadm secondary resftp

drbdadm -- --discard-my-data connect resftp

drbdadm connect resftp

mount /dev/drbd0 /home