

Procedure: Cluster monitoring

Edge Airport France

Table of Contents

FAQ : Cluster monitoring

Scope

Description

Prerequisites

Connection Schema

Functioning of the cluster

Server's supervision web page

Checking the synchronization of cluster data

Validation of the correct functioning of the cluster (Corosync)

Description of the configuration file

 The first block indicates the state of the cluster

 The second block tells you which is the primary node, and where are the services

Verifying the correct functioning of the cluster

Commands for verifying the correct functioning of the cluster

VERIFICATION OF CLUSTER MANAGEMENT TOOLS

Edge Airport France

FAQ : Cluster monitoring

Scope

Testing and Commissioning Procedure of Cluster

Description

A server cluster is composed of 2 rigorously identical servers configured in normal / backup high availability. The first server in normal mode is called “primary”, the backup server is called “secondary”.

Prerequisites

At a minimum, each server uses 3 network adapters configured as follows:

- ETH1 = Main Network Interface = IP_Server
- ETH2 = bridged network interface for virtual machines = IP_Br0
- ETH3 = “Private” server synchronization network interface, direct link between the cluster nodes.

On HP servers, the HP_ILO management interface for monitoring the machine can be set to benefit from the information of the server's physical state (see ILO monitoring documentation).

The 2 servers are connected to each other by a link allowing to have the servers in 2 different and distant technical premises to ensure the physical integrity of the equipment and the non-propagation of a physical damage on one of the two rooms.

Connection Schema



Functioning of the cluster

The Linux services used for the Cluster are:

- Drbd = data replication between disk spaces
- Corosync = Configuration and scheduling of Cluster services
- Peacemaker = Monitoring cluster services

The services configured and monitored by the Cluster are:

- Apache = Web server
- MySQL = Database
- Samba = File Sharing
- Libvirt = KVM Virtualization Engine
- Libvirtguest = Virtualization Management Tools
- IP Cluster / Route Cluster = Active Network Node

All Linux services are controlled by Corosync, do not use the standard services of Linux daemons, do not use "services" or "systemctl" commands or automatic scripts like "samba". Any activation of the services by this type of command cancels the system monitoring by peacemaker and corosync.

Server's supervision web page

Goto « http://ip_server/web/system/ezmonitor »



Checking the synchronization of cluster data

In a terminal or by ssh access on one of the cluster nodes use the drbd-overview command

```
[root@bzvairsvr bzvadmin]# drbd-overview
0:server/0 Connected Primary/Secondary UpToDate/UpToDate /EdgeServer ext4 886G 769G 73G 92%
```

Here the 2 primary and secondary servers are perfectly synchronized at the data level since the status UpToDate is effective on both servers.

Primary / Secondary Uptodate / Uptodate shows the synchronization status of the 2 nodes of the cluster.

In case the DRBD service is not started correctly (Cluster out of service), it is possible to restart the server data synchronization service via the following command:

```
# service drbdserv -full-restart
```

Validation of the correct functioning of the cluster (Corosync)

To know the state of the services managed by the cluster via a terminal or by access ssh, use the command `crm status`

```
[root@bzvairsvr bzvadmin]# crm status
```

The command returns the configuration and cluster status

```
[root@dzacupsvr ~]# crm status
=====
Last updated: Sun Sep 23 08:21:21 2018
Last change: Tue Aug 28 09:42:27 2018 via crm_attribute on dzacupsvr2
Stack: corosync
Current DC: dzacupsvr2 (34212362) - partition with quorum
Version: 1.1.7-2.mga1-ee0730e13d124c3d58f00016c3376a1de5323cff
2 Nodes configured, unknown expected votes
11 Resources configured.
=====

Online: [ dzacupsvr dzacupsvr2 ]

named (lsb:named): Started dzacupsvr
Resource Group: services
  samba (lsb:smb): Started dzacupsvr
  apache (ocf::heartbeat:apache): Started dzacupsvr
  mysql (ocf::heartbeat:mysql): Started dzacupsvr
  libvirt (lsb:libvirt): Started dzacupsvr
  libvirt-guests (lsb:libvirt-guests): Started dzacupsvr
Master/Slave Set: drbdservClone [drbdserv]
  Masters: [ dzacupsvr ]
  Slaves: [ dzacupsvr2 ]
fsserv (ocf::heartbeat:Filesystem): Started dzacupsvr
Resource Group: iphd
  clusterip (ocf::heartbeat:IPaddr2): Started dzacupsvr
  clusterroute (ocf::heartbeat:Route): Started dzacupsvr
```

Description of the configuration file

The first block indicates the state of the cluster

Last updated: Sun Sep 23 08:21:21 2018

Last change: Tue Aug 28 09:42:27 2018 via crm_attribute on dzacupsvr2

Stack: corosync

Current DC: dzacupsvr2 (34212362) - partition with quorum

Version: 1.1.7-2.mga1-ee0730e13d124c3d58f00016c3376a1de5323cff

2 Nodes configured, unknown expected votes

11 Resources configured.

The second block tells you which is the primary node, and where are the services

Online: [dzacupsvr dzacupsvr2]

Resource Group: services

samba (lsb:smb): Started dzacupsvr

apache (ocf::heartbeat:apache): Started dzacupsvr

mysql (ocf::heartbeat:mysql): Started dzacupsvr

libvirtd (lsb:libvirtd): Started dzacupsvr

libvirt-guests (lsb:libvirt-guests): Started dzacupsvr

Master/Slave Set: drbdservClone [drbdserv]

Masters: [dzacupsvr]

Slaves: [dzacupsvr2]

fsserv (ocf::heartbeat:Filesystem): Started dzacupsvr

Resource Group: iphd

clusterip (ocf::heartbeat:IPAddr2): Started dzacupsvr

clusterroute (ocf::heartbeat:Route): Started dzacupsvr

⇒ the 2 servers are “online”, and each service is operational on the primary.

Verifying the correct functioning of the cluster

See the cluster configuration, use the following command:

```
# crm configure show
```

Example of a configuration file of the Abidjan cluster:

```
node 168430081: abjairsvr
node 168430082: abjairsvr2 \
  attributes standby=off
primitive apache apache \
  params configfile="/etc/httpd/conf/httpd.conf" \
  op start interval=0 timeout=120s \
  op stop interval=0 timeout=120s
primitive clusterip IPAddr2 \
  params ip=192.168.100.1 cidr_netmask=24 nic=en01
primitive clusterroute Route \
  params destination="0.0.0.0/0" gateway=192.168.100.254 \
  meta target-role=Started
primitive drbdserv ocf:linbit:drbd \
  params drbd_resource=server \
  op monitor interval=30s role=Slave \
  op monitor interval=29s role=Master \
  op start interval=0 timeout=240s \
  op stop interval=0 timeout=100s
primitive fsserv Filesystem \
  params device="/dev/drbd/by-res/server" directory="/EdgeServer" fstype=ext4 \
  op start interval=0 timeout=60s \
  op stop interval=0 timeout=60s \
  meta target-role=Started
primitive libvirt-guests systemd:libvirt-guests \
  meta target-role=Started
primitive libvirtd systemd:libvirtd \
  meta target-role=Started
primitive mysql systemd:mysql
primitive samba systemd:smb \
  meta target-role=Started
group iphd clusterip clusterroute
group services libvirtd libvirt-guests apache mysql samba
ms drbdservClone drbdserv \
  meta master-max=1 master-node-max=1 clone-max=2 clone-node-max=1 notify=true target-role=Started
colocation fs_on_drbd inf: fsserv drbdservClone:Master
order fsserv-after-drbdserv inf: drbdservClone:promote fsserv:start
order services-after-iphd inf: iphd services
order services_after_fsserv inf: fsserv services
colocation services_on_fsserv inf: services fsserv
colocation services_on_iphd inf: services iphd
property cib-bootstrap-options: \
  dc-version=1.1.15-1.mga5-e174ec8 \
  cluster-infrastructure=corosync \
  no-quorum-policy=ignore \
  stonith-enabled=false \
  have-watchdog=false \
  resource-stickiness=100 \
  last-lrm-refresh=1533836715
nsc_defaults rsc-options: \
  resource-stickiness=600
```

Commands for verifying the correct functioning of the cluster

for example « abjairsvr2 »

| DESIRED Action | SYSTEM Command |
|---|--|
| Checking the cluster status | service corosync status |
| See cluster nodes | crm node |
| See the cluster configuration | crm configure show |
| Edit cluster configuration | crm configure edit |
| Put a cluster node in standby time to change a configuration | crm node standby abjairsvr2 |
| Put back in service a node of the cluster (here secondary of abidjan) | crm node online abjairsvr2 |
| Change a cluster configuration parameter | crm configure rsc_defaults resource-stickiness=100 |
| View the status of a cluster service | crm resource libvirt-guests status |
| Purge a cluster service that does not start | crm resource cleanup libvirt-guests |
| Check whether or not a split brain exists (service that has migrated to a non-operational node) | grep "split-brain" /var/log/syslog |
| Move a service from one node to another (in the case of a split brain) | crm resource move libvirt-guests abjairsvr2 |
| Reattach a service to the cluster | crm resource manage libvirt-guests |
| Check that the configuration files are identical between the nodes of a server | crm cluster diff /etc/samba/smb.conf |

VERIFICATION OF CLUSTER MANAGEMENT TOOLS

| DESIRED Action | SYSTEM Commands | |
|---|--|------|
| See cluster nodes | systemctl status pacemaker | |
| See the cluster configuration | systemd-analyze verify pacemaker.service | |
| Edit cluster configuration | systemctl pacemaker.service reload | |
| Put a cluster node in standby time to change a configuration | systemd-delta pacemaker.service | |
| Put back in service a node of the cluster (here secondary of abidjan) | journalctl -u pacemaker | more |

From:
<https://oldwiki.embross-airport-services.com/> - **Documentation Embross (ex Edge Airport)**

Permanent link:
<https://oldwiki.embross-airport-services.com/doku.php?id=en:faq:materiel:cluster&rev=1538038253>

Last update: **27/09/2018 10:50**

Edge Airport France

Airport Manager Solutions

Phone: +33 553 801 366

Service commercial : contact@edge-airport.com

Support technique : support@edge-airport.com

Edge Airport France SAS au capital de 150 000 €

RCS Bergerac 529 125 346 Les Lèches TVA : FR53529125346 / EORI : FR52912534600039

Tel : +33(0)553 801 366 contact@edge-airport.com www.edge-airport.com