

# Corosync i pacemaker w jednym stali domku

- czyli jak pozwolić adminowi spać w nocy II



## Agenda

- Zabawę 2 czas zacząć
- Co to jest corosync ?
- Co to jest pacemaker ?
- Zastosowania
- Podsumowanie

Maciej Miłaszewski

natan@iq.pl



# Z punktu klienta:



- Ma działać
- Wysoka dostępność
- Wysoka ważna „ważność serwisu”
- Optymalizacja kosztów
- Optymalizacja kosztów II
- Najlepiej żeby to było „cloud” i „virtual”



# Zabawe czas zacząć



Do HA:

- Heartbeat
- Corosync
- Ucarp
- Openais




„ .. nom, som takie sytuacje ”



# Na „bogato” i „dobrze”

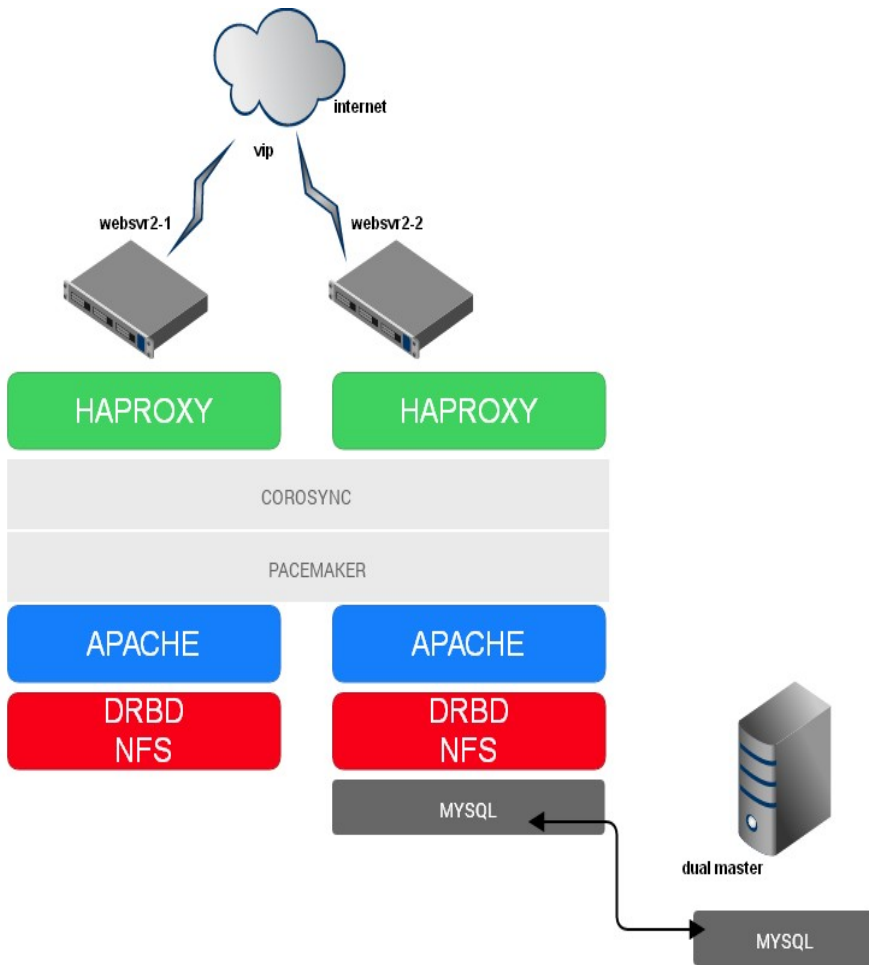
- Kasa
- Kasa
- HA
- Bo tak „cza”
- klient:

„Panowie ja mam tylko maly serwis”



# Taniej HA ?


„Albowiem 1i ar wy czempions!”



- Mniej maszyn
- Mniej problemów
- Wygoda
- Corosync
- Pacemaker
- Klient zadowolony



# Corosync



- Pochodna openais
- Dojrzałość
- Open source
- Prostota
- Wygoda
- Skalowalność

<http://corosync.github.io/corosync/>

<http://clusterlabs.org/>





# Przykładowa konfiguracja:

```
[root@vm1 ~]# cat /etc/corosync/corosync.conf
# Please read the corosync.conf.5 manual page compatibility: whitetank
```

```
totem {
  version: 2
  token 3000
  token_retransmits_before_loss_const: 10
  join: 60
  consensus: 3600
  secauth: off
  threads: 0
  rrp_mode: active

  interface {
    ringnumber: 0
    member {
      memberaddr: 10.211.0.201
    }

    member {
      memberaddr: 10.211.0.2
    }

    bindnetaddr: 10.211.0.0
    mcastport: 5405
    ttl: 1
  }
  transport: udpu
  interface {
    ringnumber: 1
    member {
      memberaddr: 10.222.0.201
    }

    member {
      memberaddr: 10.222.0.2
    }

    bindnetaddr: 10.222.0.0
    mcastport: 5405
    ttl: 1
  }
  transport: udpu
}
```

„Miałem raczej na myśli układ w którym Ty odwalasz czarną robotę, a Ja się na to patrzę i siem niecierpliwie „



„Miałem raczej na myśli układ w którym  
Ty odwalasz czarną robotę, a Ja się na  
to patrzę i siem niecierpliwie „



```
logging {  
    fileline: off  
    to_stderr: yes  
    to_logfile: yes  
    to_syslog: yes  
    logfile: /var/log/cluster/corosync.log  
    debug: off  
    timestamp: on  
    logger_subsys {  
        subsys: AMF  
        debug: off  
    }  
}  
  
amf {  
    mode: disabled  
}  
  
service {  
    # Load the Pacemaker Cluster Resource Manager  
    ver: 0  
    name: pacemaker  
}  
  
aisexec {  
    user: root  
    group: root  
}
```





# Pacemaker ?



- CRM – cluster resource manager
- Funkcjonalność
- Skalowaność
- Stabilność
- Prostota działania
- ..... ;)

<http://linux-ha.org/doc/man-pages/ap-ra-man-pages.html>

<http://clusterlabs.org/>



„Miałem raczej na myśli układ w którym  
Ty odwalasz czarną robotę, a Ja się na  
to patrzę i siem niecierpliwie „



- Funkcjonalność + Skalowaność + Stabilność + Prostota działania



„Miałem raczej na myśli układ w którym  
Ty odwalasz czarną robotę, a Ja się na  
to patrzę i siem niecierpliwie „



KOSZTEM ?



„Nie.. a nawet tak!  
Hańba Ci! ”



# UPIERDLIWOSCI KONFIGUROWANIA



# Przykładowa konfiguracja:

```
[root@vm1 ~]# crm_mon -1
```

```
=====
```

```
Last updated: Fri Mar 6 01:56:44 2015
```

```
Last change: Wed May 7 17:10:34 2014 via cibadmin on vm1
```

```
Stack: openais
```

```
Current DC: storage1 - partition with quorum
```

```
Version: 1.1.7-6.el6-
```

```
148fccfd5985c5590cc601123c6c16e966b85d14
```

```
2 Nodes configured, 2 expected votes
```

```
6 Resources configured.
```

```
=====
```

```
Online: [ storage1 vm1 ]
```

```
Resource Group: HA
```

```
  ipcloud (ocf::heartbeat:IPaddr2): Started storage1
```

```
  nfs (lsb:nfs): Started storage1
```

```
  fs_home (ocf::heartbeat:Filesystem): Started storage1
```

```
Master/Slave Set: ms_drbd [drbd]
```

```
  Masters: [ storage1 ]
```

```
  Slaves: [ vm1 ]
```

```
fs_nfs_home (ocf::heartbeat:Filesystem): Started vm1
```

„Będzie piękny, genialny, ale co najważniejsze skromny.”



```

root@vm1:~
node storage1
node vm1
primitive drbd ocf:linbit:drbd \
    params drbd_resource="nfs" \
    op start interval="0" timeout="240" \
    op stop interval="0" timeout="100" \
    op monitor interval="59s" role="Master" timeout="30s" \
    op monitor interval="60s" role="Slave" timeout="30s"
primitive fs_home ocf:heartbeat:Filesystem \
    params device="/dev/drbd0" directory="/home/www" fstype="ext4" \
    op start interval="0" timeout="60" \
    op stop interval="0" timeout="120" \
    op monitor interval="30s" \
    meta target-role="Started"
primitive fs_nfs_home ocf:heartbeat:Filesystem \
    params device="vnfs:/home/www" directory="/home/www" fstype="nfs" options="rw,noatime,tcp,soft,rsize=32768,wsiz=32768,intr,nfsvers=3,
lock" \
    op start interval="0" timeout="60s" \
    op stop interval="0" timeout="120s" \
    meta target-role="Started"
primitive ipcloud ocf:heartbeat:IPaddr2 \
    params ip="10.211.0.222" \
    op start interval="0" timeout="60s" \
    op stop interval="0" timeout="60s" \
    op monitor interval="20s" timeout="40s" \
    meta target-role="Started"
primitive nfs_lsb:nfs \
    op monitor interval="30s" \
    meta target-role="Started"
group HA ipcloud nfs fs_home
ms ms_drbd drbd \
    meta master-node-max="1" clone-max="2" clone-node-max="1" globally-unique="false" notify="true" target-role="Master"
location cli-prefer-fs_home fs_home \
    rule $id="cli-prefer-rule-fs_home" inf: #uname eq storage1
location drbd-fence-by-handler-www-ms_drbd ms_drbd \
    rule $id="drbd-fence-by-handler-www-rule-ms_drbd" $role="Master" -inf: #uname ne storage1
colocation fs_deps inf: HA ms_drbd:Master
colocation nfs_client_not_on_server -inf: fs_nfs_home HA
order fs_home_after_ms_drbd inf: ms_drbd:promote fs_home:start
order nfs_client_after_server inf: HA fs_nfs_home
property $id="cib-bootstrap-options" \
    dc-version="1.1.7-6.el6-148fccfd5985c5590cc601123c6c16e966b85d14" \
    cluster-infrastructure="openais" \
    expected-quorum-votes="2" \
    no-quorum-policy="ignore" \
    stonith-enabled="false" \
    last-lrm-refresh="1359026259" \
    maintenance-mode="false"
~

```



„Będzie piękny, genialny, ale co najważniejsze skromny.”



```
[root@vm1 ~]# crm configure show
```

```
node storage1
node vm1
primitive drbd ocf:linbit:drbd \
  params drbd_resource="nfs" \
  op start interval="0" timeout="240" \
  op stop interval="0" timeout="100" \
  op monitor interval="59s" role="Master" timeout="30s" \
  op monitor interval="60s" role="Slave" timeout="30s"
primitive fs_home ocf:heartbeat:Filesystem \
  params device="/dev/drbd0" directory="/home/www" fstype="ext4" \
  op start interval="0" timeout="60" \
  op stop interval="0" timeout="120" \
  op monitor interval="30s" \
  meta target-role="Started"
primitive fs_nfs_home ocf:heartbeat:Filesystem \
  params device="vnfs:/home/www" directory="/home/www" fstype="nfs" options="rw,noatime,tcp,soft,rsize=32768,wsiz=32768,intr,nfsvers=3,lock" \
  op start interval="0" timeout="60s" \
  op stop interval="0" timeout="120s" \
  meta target-role="Started"
primitive ipcloud ocf:heartbeat:IPAddr2 \
  params ip="10.211.0.222" \
  op start interval="0" timeout="60s" \
  op stop interval="0" timeout="60s" \
  op monitor interval="20s" timeout="40s" \
  meta target-role="Started"
primitive nfs lsb:nfs \
  op monitor interval="30s" \
  meta target-role="Started"
group HA ipcloud nfs fs_home
ms ms_drbd drbd \
  meta master-node-max="1" clone-max="2" clone-node-max="1" globally-unique="false" notify="true" target-role="Master"
location cli-prefer-fs_home fs_home \
  rule $id="cli-prefer-rule-fs_home" inf: #uname eq storage1
location drbd-fence-by-handler-www-ms_drbd ms_drbd \
  rule $id="drbd-fence-by-handler-www-rule-ms_drbd" $role="Master" -inf: #uname ne storage1
colocation fs_deps inf: HA ms_drbd:Master
colocation nfs_client_not_on_server -inf: fs_nfs_home HA
order fs_home_eter_ms_drbd inf: ms_drbd:promote fs_home:start
order nfs_client_after_server inf: HA fs_nfs_home
property $id="cib-bootstrap-options" \
  dc-version="1.1.7-6.el6-148fccfd5985c5590cc601123c6c16e966b85d14" \
  cluster-infrastructure="openais" \
  expected-quorum-votes="2" \
  no-quorum-policy="ignore" \
  stonith-enabled="false" \
  last-lrm-refresh="1359026259" \
  maintenance-mode="false"
```





„Będzie piękny, genialny, ale co najważniejsze skromny.”



## Table of Contents

ocf\_heartbeat\_anything — Manages an arbitrary service  
ocf\_heartbeat\_AoEtarget — Manages ATA-over-Ethernet (AoE) target exports  
ocf\_heartbeat\_apache — Manages an Apache web server instance  
ocf\_heartbeat\_AudibleAlarm — Emits audible beeps at a configurable interval  
ocf\_heartbeat\_ClusterMon — Runs crm\_mon in the background, recording the cluster status to an HTML file  
ocf\_heartbeat\_contrackd — This resource agent manages contrackd  
ocf\_heartbeat\_CTDB — CTDB Resource Agent  
ocf\_heartbeat\_db2 — Manages an IBM DB2 Universal Database instance  
ocf\_heartbeat\_Delay — Waits for a defined timespan  
ocf\_heartbeat\_drbd — Manages a DRBD resource (deprecated)  
ocf\_heartbeat\_Dummy — Example stateless resource agent  
ocf\_heartbeat\_eDir88 — Manages a Novell eDirectory directory server  
ocf\_heartbeat\_exportfs — Manages NFS exports  
ocf\_heartbeat\_Filesystem — Manages filesystem mounts  
ocf\_heartbeat\_fio — fio IO load generator  
ocf\_heartbeat\_ICP — Manages an ICP Vortex clustered host drive  
ocf\_heartbeat\_ids — Manages an Informix Dynamic Server (IDS) instance  
ocf\_heartbeat\_IPaddr — Manages virtual IPv4 addresses (portable version)  
ocf\_heartbeat\_IPaddr2 — Manages virtual IPv4 addresses (Linux specific version)  
ocf\_heartbeat\_IPsrcaddr — Manages the preferred source address for outgoing IP packets  
ocf\_heartbeat\_IPv6addr — Manages IPv6 aliases  
ocf\_heartbeat\_iscsi — Manages a local iSCSI initiator and its connections to iSCSI targets  
ocf\_heartbeat\_iSCSILogicalUnit — Manages iSCSI Logical Units (LUs)  
ocf\_heartbeat\_iSCSITarget — iSCSI target export agent  
ocf\_heartbeat\_jboss — Manages a JBoss application server instance  
ocf\_heartbeat\_LinuxSCSI — Enables and disables SCSI devices through the kernel SCSI hot-plug subsystem (deprecated)  
ocf\_heartbeat\_LVM — Controls the availability of an LVM Volume Group  
ocf\_heartbeat\_MailTo — Notifies recipients by email in the event of resource takeover  
ocf\_heartbeat\_ManageRAID — Manages RAID devices  
ocf\_heartbeat\_ManageVE — Manages an OpenVZ Virtual Environment (VE)  
ocf\_heartbeat\_mysql-proxy — Manages a MySQL Proxy daemon  
ocf\_heartbeat\_mysql — Manages a MySQL database instance  
ocf\_heartbeat\_nfsserver — Manages an NFS server  
ocf\_heartbeat\_nginx — Manages an Nginx web/proxy server instance  
ocf\_heartbeat\_oracle — Manages an Oracle Database instance  
ocf\_heartbeat\_oralsnr — Manages an Oracle TNS listener  
ocf\_heartbeat\_pgsq — Manages a PostgreSQL database instance  
ocf\_heartbeat\_postfix — Manages a highly available Postfix mail server instance  
ocf\_heartbeat\_proftpd — OCF Resource Agent compliant FTP script.  
ocf\_heartbeat\_Pure-FTPd — Manages a Pure-FTPd FTP server instance  
ocf\_heartbeat\_Raid1 — Manages a software RAID1 device on share  
ocf\_heartbeat\_Route — Manages network routes  
ocf\_heartbeat\_rsyncd — Manages an rsync daemon  
ains (DomUs)  
ocf\_heartbeat\_Xinetd — Manages an Xinetd service



„Cięcie! Po reżysersku to znaczy morda. „

KONIEC

