

Web Hosting z wykorzystaniem Solaris Zones

Dokument jawny

Gdańsk, 30 maja 2007 r.

możliwe rozwiązania

- nowy serwer – fizycznie nowa maszyna
 - najdroższe
 - najtrudniejsze w utrzymaniu
- VirtualHost
 - Najprostsze (?)
- Wirtualny Serwer
 - Vmware, Xen, OpenVZ / Virtuozzo, Linux-VServer, FreeBSD Jails, FreeVPS, zVM, Microsoft Virtual Server,
 - Solaris Zones

Co nam dają Solaris Zones + ZFS

- Bezpieczeństwo
- Aktualizacja oprogramowania
- Backup
- Administrowanie
 - sieć
 - zasoby dyskowe
 - CPU (*FSS* lub *Processor Sets*)
 - pamięć (*projects* lub *Memory Sets*)
- Współdzielenie zasobów

Założenia konfiguracyjne

- instalacja systemu Solaris 10
- konfiguracja zony globalnej
- konfiguracja i instalacja zony wzorcowej
- powielenie zony wzorcowej

Model rozwiązania

- **Oprogramowanie:**
 - **apache / PHP / MySQL z repozytorium CSW**
 - **binaria tylko do odczytu */opt/csw***
- **skonfigurowany**
 - **serwer WWW**
 - **baza danych dla użytkownika**
- **dostęp do zasobów przez FTP / SSH,**
- **Zasoby dyskowe:**
 - **rezerwacja - co najmniej 5 GB**
 - **quota - co najwyżej 10 GB**

Zona globalna - I

■ /etc/system

* Foil certain classes of bug exploits

```
set noexec_user_stack = 1
```

* Log attempted exploits

```
set noexec_user_stack_log = 1
```

■ IPMP

```
# cat /etc/hostname.bge0
```

```
wwhost.srv netmask + broadcast + group HAwww up
```

```
addif wwhost-bge0-probe.srv netmask + broadcast +  
-failover deprecated up
```

```
# cat /etc/hostname.bge1
```

```
wwhost-bge0-probe.srv netmask + broadcast +
```

```
-failover deprecated group HAwww up
```

■ pooladm / poolcfg

```
# pooladm -e
# pooladm -s
# pooladm
system default
    pool pool_default
        int      pool.sys_id 0
        boolean pool.active true
        [...]
        pset     pset_default
pset pset_default
    int      pset.sys_id -1
    [...]
    cpu
        int      cpu.sys_id 1
        [...]
        string   cpu.status on-line
    cpu
        [...]
```

Zona globalna - III

■ pooladm / poolcfg

```
# poolcfg -c 'modify pool pool_default
              (string pool.scheduler="FSS")
# poolcfg -c 'create pset vip-pset
              (uint pset.min=1; uint pset.max=1) '
# poolcfg -c 'create pool vip-pool'
# poolcfg -c 'associate pool vip-pool
              (pset vip-pset) '
# pooladm -c
```

żeby nie bootować serwera

```
# priocntl -s -c FSS -i class TS
# priocntl -s -c FSS -i pid 1
```

■ # dispadmin -d FSS

Tworzenie i konfiguracja wzorcowej zony

- Dopisanie do serwerów nazw rekordów IN A / IN PTR
- Przygotowanie sobie listy usług (SMF)

```
# svccfg extract > site.xml
```

- instalacja oprogramowania

```
# pkg-get -i mysql5
```

```
# pkg-get -i apache2*
```

```
# pkg-get -i php4*
```

```
# pkg-get -i php5*
```

Tworzenie zony (ZFS)

```
zname=template
```

```
ZoneQuota="10G"; ZoneRes="5G"
```

```
prefix=/zones/$zname
```

```
pprefix=zones/$zname
```

```
[[ -f /usr/sbin/zfs ]] && \  
    /usr/sbin/zfs create $pprefix
```

```
/usr/sbin/zfs set quota=$ZoneQuota $pprefix;
```

```
/usr/sbin/zfs set reservation=$ZoneRes $pprefix
```

Tworzenie zony (zasoby)

```
create -b  
set zonepath=$prefix  
set autoboot=true
```

```
add inherit-pkg-dir  
set dir=/opt
```

```
set pool=vip-pool
```

■ rctl

```
add rctl  
set name=zone.cpu-shares  
add value (priv=privileged,limit=3,action=none)
```

Instalacja zony

```
zonectfg -z $zname -f $prefix/$zname.config
zoneadm -z $zname verify
zoneadm -z $zname install
```

```
Preparing to install zone <template>.
Creating list of files to copy from the global zone.
Copying <427> files to the zone.
Initializing zone product registry.
Determining zone package initialization order.
Preparing to initialize <413> packages on the zone.
Initializing package <68> of <413>: percent complete: 16%
[...]
Initialized <413> packages on zone.
Zone <template> is initialized.
The file </zones/template/root/var/sadm/system/logs/install_log>
contains a log of the zone installation.
```

```
real      2:37.868
user      15.685
sys       15.906
```

Tworzenie zony III

■ `$prefix/root/etc/sysidcfg`

```
[...]
```

```
{hostname=template.srv
```

```
ip_address=172.16.17.18
```

```
netmask=255.240.0.0
```

```
protocol_ipv6=no
```

```
default_route=172.16.254.254}
```

```
timeserver=localhost
```

```
root_password=NP
```

```
# touch $prefix/root/etc/.NFS4inst_state.domain
```

```
# cp site.xml $prefix/root/var/svc/profile/site.xml
```

```
# cp csw*.xml $prefix/root/var/svc/manifest/site/
```

```
# zoneadm -z $zname ready
```

Tworzenie zony (dostosowanie)

```
# echo "apache2:80::::" >> $prefix/root/etc/project
# echo "mysql:3306::mysql:" >>
    $prefix/root/etc/project

# mkdir -p $prefix/root/etc/apache2
# cp -Rp@ $TemplateDir/etc/apache2/ $prefix/root/etc/

# mkdir -p $prefix/root/etc/mysql
# mkdir -p $prefix/root/home/mysql/

# cp -Rp@ $TemplateDir/home/mysql/ $prefix/root/home/
```

Instalacja zony - bootowanie

```
# zoneadm -z template boot
# ptime zlogin -C -e '#.' template
[Connected to zone 'test2' console]
SunOS Release 5.10 Version Generic_125101-01 64-bit
Copyright 1983-2007 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
Hostname: template
Loading smf(5) service descriptions: 86/86
Creating new rsa public/private host key pair
Creating new dsa public/private host key pair
rebooting system due to change(s) in /etc/default/init
[NOTICE: Zone rebooting]
SunOS Release 5.10 Version Generic_125101-01 64-bit
[...]
Hostname: template.srv
template.srv console login: #.
[Connection to zone 'test2' console closed]

real      4:04.973
```

SMF - apache

■ cswapache.xml

```
<method_context project='apache2'>
  <method_credential user='websrvd' />
  <method_environment>
    <envvar name='HOMEDIR' value='/home/apache2/' />
    <envvar name='APACHE_HOME' value='/opt/csw/apache2' />
    <envvar name='CONF_FILE' value='/etc/apache2/httpd.conf' />
    <envvar name='SERVER_CONFIG_FILE'
            value='/etc/apache2/httpd.conf' />
  </method_environment>
</method_context>
<exec_method name='start' type='method'
             exec='/opt/csw/lib/svc/method/svc-cswapache2 start'
             timeout_seconds='60'>
  <method_context>
    <method_credential user='websrvd' group='websrvd'
                      privileges='basic,!proc_session,!proc_info,
                                !file_link_any,net_privaddr' />
  </method_context>
</exec_method>
```

<http://www.sun.com/software/solaris/howtoguides/s10securityhowto.pdf>

SMF - mysql

■ cswmysql.xml

```
<method_context project='mysql'>
  <method_credential user='mysqld' group='mysqld' />
  <method_environment>
    <envvar name='HOMEDIR' value='/home/mysql' />
  </method_environment>
</method_context>
```

How to Clone a Zone ?

```
# zoneadm -z template halt
# zonecfg -z template export -f
                        /zones/template.cfg
# abra-ka-dabra.sh < /zones/template.cfg
                        > /zones/my-zone.cfg

# zfs create zones/my-zone
# chmod 700 /zones/my-zone
# zfs set quota=10G zones/test3
# zfs list
zones/my-zone      24.5K 10.0G  24.5K  /zones/my-zone
# zonecfg -z my-zone -f /zones/my-zone.cfg
```

How to Clone a Zone ?

```
# ptime zoneadm -z my-zone clone template
Cloning zonepath /zones/template...
```

```
real      1:35.424
```

```
#zfs list
```

```
zones/my-zone      373M  9.64G   373M   /zones/my-zone
```

```
# zoneadm -z my-zone list -v
```

ID	NAME	STATUS	PATH
-	my-zone	installed	/zones/my-zone

■ **`$prefix/root/etc/sysidcfg`**

■ ~~`# touch $prefix/root/etc/.NFS4inst_state.domain`~~

■ ~~`cp site.xml$prefix/root/var/svc/profile/site.xml`~~

How to Clone a Zone ?

```
# ptime zoneadm -z my-zone ready
real          9.679

# ptime zoneadm -z my-zone boot
real          0.057

# ptime zlogin -C -e '#.' my-zone
[Connected to zone 'test3' console]
SunOS Release 5.10 Version Generic_125101-01 64-bit
Copyright 1983-2007 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
Hostname: myzone.srv
Creating new rsa public/private host key pair
Creating new dsa public/private host key pair
my-zone.srv console login: #.
[Connection to zone 'test3' console closed]
real          14.888
```

Zarządzanie zainstalowaną zoną

zfs list

NAME	USED	AVAIL	REFER	MOUNTPOINT
zones	37.6G	29.8G	191M	/zones
zones/my-zone	374M	9.64G	374M	/zones/my-zone

```
# prctl -n zone.cpu-shares -r -v 25 -i zone my-zone
```

Jak usunąć zonę

```
#!/bin/bash
    if [ $# -ne 1 ]; then
        echo "usage: $0 <zone-name>"; echo "zones on this system"
        zoneadm list -vc | grep -v "0 global"
        exit 1
    fi
```

Zone=\$1

```
    echo "You are going to completely remove zone $Zone"
    read -p "Are you sure ? [y/n] : " OK
    [ "$OK" == "y" ] || exit 2
```

zoneadm -z \$Zone list -v

zfs list zones/\$Zone

```
    read -p "Are you sure that you are sure ??? [y/n] : " OK
    [[ "$OK" == "y" ]] || exit 3
    echo "Halting zone $Zone"
```

zoneadm -z \$Zone halt

```
    echo "uninstalling zone $Zone"
```

zoneadm -z \$Zone uninstall -F

```
    echo "deleting zone $Zone"
```

zonecfg -z \$Zone delete -F

```
    echo "destroying filesystem zones/$Zone"
```

zfs destroy zones/\$Zone

Plany na przyszłość

■ Solaris 7/07 (update 4)

```
zonecfg:testing1> add capped-memory  
zonecfg:testing1:capped-memory> set physical=512m  
zonecfg:testing1:capped-memory> set swap=512m  
zonecfg:testing1:capped-memory> end
```

■ Nevada – build 62

```
zonecfg:testing1> add capped-cpu  
zonecfg:testing1:capped-cpu> set ncpus=1.0  
zonecfg:testing1:capped-cpu> end
```

■ Tworzenie zony ze snapshotu

```
# zfs snapshot zones/template@snap1  
# zoneadm -z zone2 clone -s  
                  zones/template@snap1 my-zone
```

Linki

- **Dostawcy, oferujący Web Hosting bazujący na Solaris Zones**
- <http://www.joyent.com/>
- <http://www.gridzones.com/>
- <http://entic.net/servers>

Bibliografia

- <http://www.sun.com/bigadmin/content/zones/>
- <http://www.opensolaris.org/os/community/zones/>
- <http://www.solarisinternals.com/wiki/index.php/Zones>
- http://www.solarisinternals.com/wiki/index.php/Zones_Resource_Controls
- <http://docs.sun.com/app/docs/doc/817-1592>

- <http://www.sun.com/software/solaris/howtoguides/s10securityhowto.pdf>
- <http://www.sun.com/software/solaris/howtoguides/containersLowRes.jsp>
- http://www.sun.com/software/whitepapers/solaris10/grid_containers.pdf
- <http://www.sun.com/blueprints/0505/819-2679.pdf>
- <http://www.sun.com/blueprints/1205/819-5149.pdf>
- <http://www.sun.com/blueprints/1005/819-3734.pdf>
- <http://www.sun.com/blueprints/0605/819-2888.pdf>

- http://en.wikipedia.org/wiki/Solaris_Containers
- <http://www.cuddletech.com/blog/pivot/entry.php?id=803>

Pytania ?

Michał Łuczak <warf@wp-sa.pl>