
Subject: *SOLVED* chkpnt/restore time
Posted by [working_men](#) on Wed, 08 Aug 2007 13:07:04 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi!

I wanted to test how long it takes for chkpnt/restore to complete.
Trying to determine time vs ram

1. test:
tmpfs with 350MB urandom/zero file
-->chkpnt: 81sec
-->restore: 86sec
-->dumpsizes: 340MB

2. test:
C-program that allocates 660MB
-->chkpnt: 26sec
-->restore: 22sec
-->dumpsizes: 660MB

3. test:
installed mysql-server (140MB)
-->chkpnt: 0.5sec
-->restore: 4sec
-->dumpsizes: 30MB

My questions are:
How does openVZ generate the dump files?
or
Is there a difference between allocated/used memory?

It seems to me that it depends on which type of data is in RAM for the size of the dump and the time it takes.

Thanks alot for clearing this up for me!

thomas

Subject: Re: chkpnt/restore time
Posted by [Andrey Mirkin](#) on Thu, 09 Aug 2007 08:06:43 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello!

There are different types of memory pages which are used by processes.
We need to save whole pages which are allocated and used, but we need to save only

information about allocated pages.

We do not need to save pages which are read-only mapped from files (code sections and other data).

In your examples we have:

1. tmpfs

we have to save this file to dump, because it is stored only in memory and we can't map it from another place later during restore

That is why size of dump file is so huge

2. C-program

it is not clear what your program is doing. Just allocating memory or allocating and clearing it?

Can you please post here source code of this test program.

3. mysql

mysql is threaded application and much memory is used for pages with code. We do not save such pages as I wrote above, thus image size is not so huge.

Subject: Re: chkpnt/restore time

Posted by [working_men](#) on Thu, 09 Aug 2007 08:18:34 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi

Firstly thanks for the info.

I understand now why dump sizes vary.

Just for information the veryvery trival C-Code:

```
main()
{
int i;
for(i=0;i<200000000;i++)
{
    malloc(sizeof(int));
}
sleep(500000);
}
```

Thanks again

Thomas

Subject: Re: chkpnt/restore time

Posted by [Andrey Mirkin](#) on Fri, 10 Aug 2007 08:33:03 GMT

[View Forum Message](#) <> [Reply to Message](#)

Well, in this program you allocating 20000000 memory regions. For each such memory region kernel will allocate internal structure at least 16-byte length and we need to dump this structure during checkpointing.

If you will run simple program which doing malloc one time but for big memory region then you will see that image file will be small.
