
Subject: [ANNOUNCE] OpenVZ releases checkpointing/live migration of processes
Posted by [Kirill Korotaev](#) on Thu, 20 Apr 2006 17:14:23 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

OpenVZ team is proud to announce the release of the new checkpointing/restore feature. This feature allows to save (checkpoint) and restore the whole state of a Virtual Environment (VE, container) and do a live migration of a VE to another physical box while preserving process states and TCP/IP connections.

During live migration the in-kernel state of processes and their resources (including memory, registers, IPC, pids, open files, sockets, etc.) is saved and then restored on another machine. Since all network connections are preserved with all the in-progress requests, user doesn't experience interruption of service.

The feature is available on i686 and x86_64 architectures. Migration of 32bit VEs between i686 and x86_64 architectures is also supported. Current implementation works fine with complex applications like Oracle, Java, X apps.

Latest 2.6.16 OpenVZ kernel and tool packages with live migration support are available here:
<http://openvz.org/download/beta/kernel/>
<http://openvz.org/download/utls/>

GIT repository for all OpenVZ sources is available at
<http://git.openvz.org/>

Usage examples

~~~~~

New 'vzmigrate' utility is used for VE migration. Also, new commands for 'vzctl' allowing to dump and restore VE were introduced: 'chkpnt' and 'restore'.

To save current VE state with all processes:  
# vzctl chkpnt <VEID>

To restore VE after checkpointing:  
# vzctl restore <VEID>

To perform online migration of VE #101 to another machine:  
# vzmigrate --online destination.node.com 101  
without '--online' option vzmigrate does offline VE migration with VE start/stop.

With best regards,  
OpenVZ team.

---

---

Subject: Re: [ANNOUNCE] OpenVZ releases checkpointing/live migration of processes

Posted by [Jun OKAJIMA](#) on Sat, 22 Apr 2006 05:48:03 GMT

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

---

>  
>OpenVZ team is proud to announce the release of the new  
>checkpointing/restore feature. This feature allows to save (checkpoint)  
>and restore the whole state of a Virtual Environment (VE, container) and  
>do a live migration of a VE to another physical box while preserving  
>process states and TCP/IP connections.  
>  
>During live migration the in-kernel state of processes and their  
>resources (including memory, registers, IPC, pids, open files, sockets,  
>etc.) is saved and then restored on another machine. Since all network  
>connections are preserved with all the in-progress requests, user doesn't  
>experience interruption of service.  
>  
>The feature is available on i686 and x86\_64 architectures. Migration of  
>32bit VEs between i686 and x86\_64 architectures is also supported.  
>Current implementation works fine with complex applications like Oracle,  
>Java, X apps.  
>  
>Latest 2.6.16 OpenVZ kernel and tool packages with live migration  
>support are available here:  
><http://openvz.org/download/beta/kernel/>  
><http://openvz.org/download/utils/>  
>  
>GIT repository for all OpenVZ sources is available at  
><http://git.openvz.org/>  
>  
>Usage examples  
>~~~~~  
>  
>New 'vzmigrate' utility is used for VE migration. Also, new commands for  
>'vzctl' allowing to dump and restore VE were introduced: 'chkpnt' and  
>'restore'.  
>  
>To save current VE state with all processes:  
># vzctl chkpnt <VEID>  
>  
>To restore VE after checkpointing:  
># vzctl restore <VEID>

>  
>To perform online migration of VE #101 to another machine:  
># vzmigrate --online destination.node.com 101  
>without '--online' option vzmigrate does offline VE migration with VE  
>start/stop.  
>  
>

One Question ---.

Xen needs iSCSI or alike for a FS when you want to do a live migration,  
because it needs very same inode mapping.

How about this issue on OpenVZ?

For example, can I migrate a VE to a server which has same file tree on  
its HDD, but has different FS type and inode number?

--- Okajima, Jun. Tokyo, Japan.

---

Subject: Re: [ANNOUNCE] OpenVZ releases checkpointing/live migration  
of processes

Posted by [kir](#) on Sat, 22 Apr 2006 19:53:42 GMT

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

---

Jun OKAJIMA wrote:

>>OpenVZ team is proud to announce the release of the new  
>>checkpointing/restore feature. This feature allows to save (checkpoint)  
>>and restore the whole state of a Virtual Environment (VE, container) and  
>>do a live migration of a VE to another physical box while preserving  
>>process states and TCP/IP connections.

>>

>>During live migration the in-kernel state of processes and their  
>>resources (including memory, registers, IPC, pids, open files, sockets,  
>>etc.) is saved and then restored on another machine. Since all network  
>>connections are preserved with all the in-progress requests, user doesn't  
>>experience interruption of service.

>>

>>[....skipped....]

>>

>One Question ---.

>

>Xen needs iSCSI or alike for a FS when you want to do a live migration,  
>because it needs very same inode mapping.

>How about this issue on OpenVZ?

>

>

We do not have any special hardware requirements for live migration.

I.e. your laptop will work fine.

>For example, can I migrate a VE to a server which has same file tree on  
>its HDD, but has different FS type and inode number?

>

Yes you can; we just use rsync prior to freezing a VE to copy VE's files  
to the destination node, and another rsync after VE freeze.

---

---

Subject: Re: [ANNOUNCE] OpenVZ releases checkpointing/live migration  
of processes

Posted by [Jun OKAJIMA](#) on Sun, 23 Apr 2006 06:26:37 GMT

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

---

>

>>For example, can I migrate a VE to a server which has same file tree on  
>>its HDD, but has different FS type and inode number?

>>

>Yes you can; we just use rsync prior to freezing a VE to copy VE's files  
>to the destination node, and another rsync after VE freeze.

>

>

Good!.

What I am planning is, using OpenVZ on clients, not on server.

And, the purpose is fast booting by suspending technology.

For example, you put KDE/OOo to a VE, and then make a freezed image  
of it. Then, next time, you "boot" (actually it is a resuming) the image.

I suppose you can "boot"(=resume) your desktop very rapidly with this way.

How about this idea?

--- Okajima, Jun. Tokyo, Japan.

---

---

Subject: Re: [ANNOUNCE] OpenVZ releases checkpointing/live migration  
of processes

Posted by [dev](#) on Mon, 24 Apr 2006 05:27:49 GMT

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

---

> Good!.

>

> What I am planning is, using OpenVZ on clients, not on server.

> And, the purpose is fast booting by suspending technology.

> For example, you put KDE/OOo to a VE, and then make a freezed image

> of it. Then, next time, you "boot" (actually it is a resuming) the image.  
>  
> I suppose you can "boot"(=resume) your desktop very rapidly with this way.  
>  
> How about this idea?  
It should work fine. At least we tested xvnc inside VE and it was  
working great. Do you use xvnc inside VE?

Thanks,  
Kirill

---