

---

Subject: Is sharing same files between VEs possible with simfs?

Posted by [piavlo](#) on Sun, 20 Jan 2008 14:40:47 GMT

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

---

Hi, i'd like to initially clone some VE private area with rsync using the --link-dest= option for hard linking so that to VE with same unmodified file will actually use same inode and save space.

The question is will it work correctly with simfs? Meaning that if one VE modifies some shared file data inode this inode will be cloned with modifications, and the original inode link count will be reduced. Or will the original node be modified thus affecting all VEs sharing the file?

Thanks

Alex

---

---

Subject: Re: Is sharing same files between VEs possible with simfs?

Posted by [rickb](#) on Sun, 20 Jan 2008 19:49:32 GMT

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

---

Hi, the openvz fs does not have this ability. If this is critical for you, you should consider virtuoizzo and its vzfs filesystem, which accomplishes exactly what you want.

---

---

Subject: Re: Is sharing same files between VEs possible with simfs?

Posted by [piavlo](#) on Mon, 21 Jan 2008 07:23:41 GMT

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

---

rickb wrote on Sun, 20 January 2008 21:49Hi, the openvz fs does not have this ability. If this is critical for you, you should consider virtuoizzo and its vzfs filesystem, which accomplishes exactly what you want.

Well it's certainly not critical for me, just nice to have, especially considering it should not be hard to implement in container aware vfs code.

---

---

Subject: Re: Is sharing same files between VEs possible with simfs?

Posted by [szocske](#) on Wed, 23 Jan 2008 09:14:07 GMT

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

---

piavlo wrote on Mon, 21 January 2008 02:23rickb wrote on Sun, 20 January 2008 21:49Hi, the openvz fs does not have this ability. If this is critical for you, you should consider virtuoizzo and its vzfs filesystem, which accomplishes exactly what you want.

Well it's certainly not critical for me, just nice to have, especially considering it should not be hard to implement in container aware vfs code.

Overlay filesystems should be virtualisation-agnostic, but probably more hassle and less efficient than the virtio- specific integrated solution.

VServer has vhashify, but lacks many openVZ features we all like

---