

Hi Daniel,

On Fri, 2006-12-15 at 11:56 +0100, Daniel Lezcano wrote:

- > To be able to take a snapshot of the network container, we must ensure
- > it is freezed during the checkpoint, because we must ensure the
- > consistency in the host and with the peers network stack.
- >
- > We began the checkpoint/restart discussion with this point: how do we do
- > container's network freeze ?
- >
- > * The first step is to drop the traffic
- > - shall it be done with the sk_filter fields of the socket ?
- > - or with the netfilter NF_DROP/NF_STOLEN ?

I've tested with a very simple sk_filter and found that sk_filter cannot drop outgoing packets whereas it can drop incoming packets. Therefore, if the implementation is carefully done so as not to send outgoing ones when checkpointing, it can freeze container's network. But if there may be some accidental packet sending, maybe it fails to freeze the network. Or, we had better implement sk_filter for outgoing packets.

The following function is my simple sk_filter set in user level.
I have tested for both TCP and UDP socket.

Thanks,

Masahiko.

```
void filter_socket ( int fd )
{
    struct sock_fprog prg;
    struct sock_filter skf[1] = { 0, };

    prg.filter = skf;

    skf[0].code = BPF_RET|BPF_A;
    prg.len = 1;

    if ( setsockopt ( fd, SOL_SOCKET, SO_ATTACH_FILTER,
                    (void *) &prg, sizeof(prg)) < 0 )
        perror( "setsockopt SO_ATTACH_FILTER" );
}
```

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
