
Subject: [PATCH 0/3] Sysctl shadow management
Posted by [Pavel Emelianov](#) on Tue, 20 Nov 2007 11:38:36 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi guys!

You all know, that with multiple namespaces we have to take special care about sysctls. E.g. IPC sysctl handlers are equipped with kludges to alter the sysctl parameters of appropriate namespace. The same thing should be done for UTS namespace (but it is not - we have a BUG in mainstream) and (!) for network namespaces.

Unlike all the other namespaces, network will have to not just address different variables via same sysctl names, but to have different tables with different sysctl names. E.g. `/proc/sys/net/conf` have entries for devices, which differ across namespaces.

Eric currently have some work done in that directions, I like the approach in general very much, but it looks rather raw (Eric, take this in good part). You know, `ifdefs` in the middle of the code, explicit references to net namespace and so on and so forth.

So here's the RFC for a bit better sysctls shadow management.

I will provide 3 patches:

1. the sysctl shadows themselves;
2. using shadows in UTS namespace;
3. using shadows in IPC namespace;

If someone want I can send

4. example on how to create a `/proc/sys/net/conf/-like` structure with different names.

Using them in net namespace is already checked (I created sysctl entries with different names), but I don't have any patches against any David's tree yet. If we're OK with this set I will start talking to Andrew and David about who to send these patches to and making shadows for net-related sysctl variables.

Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

Containers mailing list
Containers@lists.linux-foundation.org

