
Subject: [PATCH 0/12] L2 network namespace (v3)
Posted by [Mishin Dmitry](#) on Wed, 17 Jan 2007 15:51:14 GMT
[View Forum Message](#) <> [Reply to Message](#)

This is an update of L2 network namespaces patches. They are applicable to Cedric's 2.6.20-rc4-mm1-lxc2 tree.

Changes:

- updated to 2.6.20-rc4-mm1-lxc2
- current network context is per-CPU now
- fixed compilation without CONFIG_NET_NS

Changed current context definition should fix all mentioned by Cedric issues:

- the nsproxy backpointer is unnecessary now - thus removed;
- the push_net_ns() and pop_net_ns() use per-CPU variable now;
- there is no race on ->nsproxy between push_net_ns() and exit_task_namespaces() because they deals with different pointers.

=====
L2 network namespaces

The most straightforward concept of network virtualization is complete separation of namespaces, covering device list, routing tables, netfilter tables, socket hashes, and everything else.

On input path, each packet is tagged with namespace right from the place where it appears from a device, and is processed by each layer in the context of this namespace.

Non-root namespaces communicate with the outside world in two ways: by owning hardware devices, or receiving packets forwarded them by their parent namespace via pass-through device.

This complete separation of namespaces is very useful for at least two purposes:

- allowing users to create and manage by their own various tunnels and VPNs, and
- enabling easier and more straightforward live migration of groups of processes with their environment.

--
Thanks,
Dmitry.

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