
Subject: RE: Network virtualization/isolation

Posted by [Leonid Grossman](#) on Sat, 25 Nov 2006 16:35:58 GMT

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

> -----Original Message-----

> From: netdev-owner@vger.kernel.org

> [mailto:netdev-owner@vger.kernel.org] On Behalf Of Eric W. Biederman

> Then the question is how do we reduce the overhead when we

> don't have enough physical network interfaces to go around.

> My feeling is that we could push the work to the network

> adapters and allow single physical network adapters to

> support multiple network interfaces, each with a different

> link-layer address. At which point the overhead is nearly

> nothing and newer network adapters may start implementing

> enough filtering in hardware to do all of the work for us.

Correct, to a degree.

There will be always a limit on the number of physical "channels" that a NIC

can support, while keeping these channels fully independent and protected at the hw level.

So, you will probably still need to implement the sw path, with the assumption that some containers (that care about performance) will get a separate

NIC interface and avoid the overhead, and other containers will have to use the sw path.

There are some multi-channel NICs shipping today so it would be possible to see the overhead between the two options (I suspect it will be quite noticeable), but for a general idea about what work could be pushed down to network adapters in the near future you can look at the pcisig.com I/O Virtualization Workgroup.

Once the single root I/O Virtualization spec is completed, it is likely to be supported by several NIC vendors to provide multiple network interfaces on a single NIC that you are looking for.

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>
