Subject: Re: [RFD] L2 Network namespace infrastructure
Posted by ebiederm on Sat, 23 Jun 2007 17:26:28 GMT
View Forum Message <> Reply to Message

Ben Greear <greearb@candelatech.com> writes:

> Patrick McHardy wrote:
>> Eric W. Biederman wrote:
>>
>>> -- The basic design
>>>
>>> There will be a network namespace structure that holds the global
>>> variables for a network namespace, making those global variables
>>> per network namespace.
>>>
>>> One of those per network namespace global variables will be the
>>> loopback device. Which means the network namespace a packet resides
>>> in can be found simply by examining the network device or the socket
>>> the packet is traversing.
>>>
>>> Either a pointer to this global structure will be passed into
>>> the functions that need to reference per network namespace variables
>>> or a structure that is already passed in (such as the network device)
>>> will be modified to contain a pointer to the network namespace
>>> structure.
>>>
>>
>>
>> I believe OpenVZ stores the current namespace somewhere global,
>> which avoids passing the namespace around. Couldn't you do this
>> as well?
>>
> Will we be able to have a single application be in multiple name-spaces?

A single application certainly. But then an application can be composed
of multiple processes which can be composed of multiple threads.

In my current patches a single task_struct belongs to a single network
namespace. That namespace is used when creating sockets. The sockets
themselves have a namespace tag and that is used when transmitting
packets, or otherwise operating on the socket.

So if you pass a socket from one process to another you can have
sockets that belong to different network namespaces in a single task.

Eric

_____
Containers mailing list