

Hi,

I'm trying to run a J2EE application server cluster in OpenVZ. Each J2EE application server instance in the cluster relies on multicast to find the other cluster members, so I need to get multicast working with OpenVZ. To start with, I'm working with a two node cluster, consisting of two VEs, each hosting a single J2EE server node. I've got multicast working when the VEs are both on the same HN, and so the J2EE cluster works fine in this situation. However, if the VEs are on separate HNs, multicast does not work between the VEs, and so the cluster members do not find each other. Note that all works fine if the two J2EE application server instances are on the HNs themselves (i.e. not running in VEs) - i.e. multicast is working fine between the two HNs. So, I'd be grateful for some help in trying to get multicast working when the VEs are on different HNs.

To clarify, current working configuration, with both VEs on the same HN:

```
|| HN1 [VE1] <--- multicast via bridge (vzbr0) ---> [VE2] ||
```

And desired configuration, with VEs on separate HNs:

```
|| HN1 [VE1] || <-- -multicast ---> || HN2 [VE2] ||
```

Environment details:

Host Kernel = ovzkernel-smp-2.6.9-023stab044.11.x86_64 (RHEL4 based kernel for x86_64).
Guest Operating System = centos-4-x86_64-minimal .

Following the thread at <http://forum.openvz.org/index.php?t=msg&goto=21090> , the VEs are configured to use veth, with the veths connected using a bridge (vzbr0). The routing tables are setup so that the (unicast) IPs for the VEs are routed via the bridge. This all works fine for unicast in all scenarios - i.e. between the VEs, between the hosting HN and VEs, and between other nodes and the VEs. It also works fine for multicast between the VEs if they are on the same HN. However, it does not work for multicast from the HNs to the VEs they are hosting, or from other nodes (including VEs on other nodes) to the VEs.

Steps carried out so far:

In attempting to get multicast working between VEs on separate machines (HNs), so far, I've carried out the following:

1. Added a route for the Class D multicast address to the bridge (vzbr0). This enabled multicast from the hosting HN to the VE it was hosting (verified this using ping). However, other nodes, including the VE on the other HN still could not send / receive multicast traffic from the VE.
2. Attempted to enable multicast forwarding for the bridge by executing "echo 1 > /proc/sys/net/ipv4/conf/vzbr0/mc_forwarding" on the HN . However, this was denied ("Operation

not permitted"). Perhaps the OpenVZ RHEL4 kernel does not allow this by default ? However, unsure if this is required to get multicast to work.

Questions:

So, no luck with the above. Does anyone have any suggestions as to how I could get multicast between VEs on different machines to work ?

In particular, is enabling the mc_forwarding setting required ? If so, is a RHEL4 kernel available with this enabled ?

Or is there something else I'm missing ?

Many thanks for any help on this.