
Subject: initcwnd adjustment using RHEL5 kernels?
Posted by [aquarapid](#) on Wed, 15 Jun 2011 07:46:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

When using a vanilla RHEL5 kernel, you can adjust the TCP initial congestion window (to 10, instead of the default 3, assuming a normal mss) by setting the initcwnd option on the default route. Something like:

```
ip route change default via x.x.x.x dev eth0 initcwnd 10
```

I can confirm that this works using the vanilla CentOS el5 kernel 2.6.18-238.9.1.el5.

I have been trying to make it work using the OpenVZ RHEL5 kernels (specifically using el5.028stab091.1, but I've also tried el5.028stab089.1). Eventually I want to use it inside a container, but for now I'm just testing using CT0 (one step at a time...).

However, I'm not having any luck at all. Setting the initcwnd via "ip route" succeeds, but I am still seeing the normal TCP slow start behavior when fetching files via HTTP from the machine. I've checked all the obvious problems, e.g.:

- * the default TCP send buffer size (set via net.ipv4.tcp_wmem) is indeed ≥ 10 times mss
- * make sure to flush the route cache between tests.
- * Restart the (HTTP) server application between tests for good measure.

Can anyone confirm that they have raised initcwnd in VE0 (or even better, inside a container) and actually seen it work?

J

Subject: Re: initcwnd adjustment using RHEL5 kernels?
Posted by [kir](#) on Wed, 15 Jun 2011 08:03:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

Please file a bug report to bugzilla.openvz.org if you want developers to look at the problem.

Subject: Re: initcwnd adjustment using RHEL5 kernels?
Posted by [aquarapid](#) on Wed, 15 Jun 2011 10:38:59 GMT
[View Forum Message](#) <> [Reply to Message](#)

Nevermind, I found the problem, which is quite interesting.

I was setting the sysctl net.ipv4.tcp_no_metrics_save=1 in my openvz sysctl.conf, which seems to kill the initcwnd adjustment and cause fallback to defaults.

As soon as I tracked this down and reset it back to the default value of 0, the initcwnd adjustment

works as expected.

J
