

## Re: [PATCH 0/3] netfilter : 3 patches to boost ip\_tables performance

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-10/1843.html>

---

**From:** Patrick McHardy (*kaber\_at\_trash.net*)

**Date:** 10/07/05

Date: Fri, 07 Oct 2005 19:08:36 +0200

To: Andi Kleen <ak@suse.de>

Andi Kleen wrote:

> On Fri, Oct 07, 2005 at 04:38:02AM +0200, Harald Welte wrote:

>

>> On Wed, Oct 05, 2005 at 06:53:31PM +0200, Andi Kleen wrote:

>>

>>> Well you most likely wrecked local performance then when it's enabled.

There are lots of other hooks and conntrack/NAT already have a quite large negative influence on performance. Do you have numbers that show that enabling this actually causes more than a slight decrease in performance? Besides, most distributors enable all these options anyway, so it only makes a difference for a small group of users.

>> so you would favour a system that incorrectly deals with ICMP errors but

>> has higher performance?

>

> I would favour a system where development doesn't lose sight of performance.

I don't think we do.

> Perhaps there would be other ways to fix this problem without impacting

> performance unduly? Can you describe it in detail?

When an ICMP error is send by the firewall itself, the inner packet needs to be restored to its original state. That means both DNAT and SNAT which might have been applied need to be reversed. DNAT is reversed at places where we usually do SNAT (POST\_ROUTING), SNAT is reversed where usually DNAT is done (PRE\_ROUTING/LOCAL\_OUT). Since locally generated packets never go through PRE\_ROUTING, it is done in LOCAL\_OUT, which required enabling NAT in LOCAL\_OUT unconditionally. It might be possible to move this to some different hook, I didn't investigate it.

—

Linux-Kernel: Re: [PATCH 0/3] netfilter : 3 patches to boost ip\_tables performance

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to [majordomo@vger.kernel.org](mailto:majordomo@vger.kernel.org)  
More majordomo info at <http://vger.kernel.org/majordomo-info.html>  
Please read the FAQ at <http://www.tux.org/lkml/>