
Subject: [PATCH] vzctl: Add support for xt_tcpudp (aka ipt_tcp/ipt_udp) module introduced by x_tables

Posted by [lakostis](#) on Thu, 09 Nov 2006 07:48:35 GMT

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

In recent 2.6 kernels (2.6.16+) netfilter code is changed by new x_tables (Abstraction layer for {ip,ip6,arp}_tables). This patch updates iptables wrapper code in vzctl for this changes (adding xt_tcpudp support).

Signed-off-by: Konstantin A. Lepikhov <lakostis@altlinux.org>

```
etc/vz.conf      |  2 ++
include/linux/vzcalluser.h |  7 +++++-
man/vzctl.8      |  2 ++
src/lib/iptables.c |  1 +
5 files changed, 8 insertions(+), 4 deletions(-)
```

```
diff --git a/etc/vz.conf b/etc/vz.conf
index f1e9a42..5f23b90 100644
--- a/etc/vz.conf
+++ b/etc/vz.conf
@@ -29,5 +29,5 @@ DEF_OSTEMPLATE="fedora-core-4"
## Load vzwdog module
VZWDOG="no"

-IPTABLES="ipt_REJECT ipt_tos ipt_limit ipt_multiport iptable_filter iptable_mangle ipt_TCPMSS
ipt_tcpmss ipt_ttl ipt_length"
+IPTABLES="xt_tcpudp ipt_REJECT ipt_tos ipt_limit ipt_multiport iptable_filter iptable_mangle
ipt_TCPMSS ipt_tcpmss ipt_ttl ipt_length"
```

```
diff --git a/include/linux/vzcalluser.h b/include/linux/vzcalluser.h
index 360f246..31100b1 100644
--- a/include/linux/vzcalluser.h
+++ b/include/linux/vzcalluser.h
@@ -97,6 +97,7 @@ #define VE_IP_NAT_FTP_MOD (1U<<21)
#define VE_IP_NAT_IRC_MOD (1U<<22)
#define VE_IP_TARGET_REDIRECT_MOD (1U<<23)
#define VE_IP_MATCH_MAC_MOD (1U<<24)
+#define VE_IP_MATCH_TCPUDP_MOD (1U<<25)

/* these masks represent modules with their dependences */
#define VE_IP_IPTABLES (VE_IP_IPTABLES_MOD)
@@ -143,13 +144,15 @@ #define VE_IP_NAT_FTP (VE_IP_NAT_FTP_MO
#define VE_IP_NAT_IRC (VE_IP_NAT_IRC_MOD \
| VE_IP_NAT | VE_IP_CONNTRACK_IRC)
#define VE_IP_TARGET_REDIRECT (VE_IP_TARGET_REDIRECT_MOD \
- | VE_IP_NAT)
+ | VE_IP_NAT)
```

```

#define VE_IP_MATCH_MAC (VE_IP_MATCH_MAC_MOD \
+ | VE_IP_IPTABLES)
+#define VE_IP_MATCH_TCPUDP (VE_IP_MATCH_TCPUDP_MOD \
| VE_IP_IPTABLES)

/* safe iptables mask to be used by default */
#define VE_IP_DEFAULT \
- (VE_IP_IPTABLES | \
+ (VE_IP_IPTABLES | VE_IP_MATCH_TCPUDP | \
VE_IP_FILTER | VE_IP_MANGLE | \
VE_IP_MATCH_LIMIT | VE_IP_MATCH_MULTIPOINT | \
VE_IP_MATCH_TOS | VE_IP_TARGET_REJECT | \
diff --git a/man/vzctl.8 b/man/vzctl.8
index eec3155..0fb6eaa 100644
--- a/man/vzctl.8
+++ b/man/vzctl.8
@@ -311,7 +311,7 @@ You can use the following values for \fI
\flipt_length\fR, \flipt_conntrack\fR, \flipt_conntrack_ftp\fR,
\flipt_conntrack_irc\fR, \flipt_conntrack\fR, \flipt_state\fR,
\flipt_helper\fR, \fliptable_nat\fR, \flipt_nat_ftp\fR, \flipt_nat_irc\fR,
-\flipt_REDIRECT\fR \flixt_mac\fR.
+\flipt_REDIRECT\fR, \flixt_mac\fR, \flixt_tcpudp\fR.
.TP
\fBNetwork devices control parameters\fR
.IP "\fB--netdev_add\fR \fIname\fR"
diff --git a/src/lib/iptables.c b/src/lib/iptables.c
index 8d07264..93bbbc5 100644
--- a/src/lib/iptables.c
+++ b/src/lib/iptables.c
@@ -52,6 +52,7 @@ #endif
#endif VE_IP_MATCH_MAC
 {"xt_mac", VE_IP_MATCH_MAC_MOD, VE_IP_MATCH_MAC},
#endif
+ {"xt_tcpudp", VE_IP_MATCH_TCPUDP_MOD, VE_IP_MATCH_TCPUDP},
#endif /* VZCTL_ENV_CREATE_DATA */
{NULL, 0}
};

--
```

1.4.2.GIT

--
WBR et al.

Subject: Re: [PATCH] vzctl: Add support for xt_tcpudp (aka ipt_tcp/ipt_udp) module

introduced by x_tables

Posted by [Mishin Dmitry](#) on Thu, 09 Nov 2006 09:18:37 GMT

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

On Thursday 09 November 2006 10:48, Konstantin A. Lepikhov wrote:

> In recent 2.6 kernels (2.6.16+) netfilter code is changed by new x_tables
> (Abstraction layer for {ip,ip6,arp}_tables). This patch updates iptables
> wrapper code in vzctl for this changes (adding xt_tcpudp support).
xt_tcpudp previously were inside ip_tables module. So, for compatibility,
current kernels handles this module by VE_IP_IPTABLES mask.

```
>
> Signed-off-by: Konstantin A. Lepikhov <lakostis@altlinux.org>
> ---
> etc/vz.conf      |  2 ++
> include/linux/vzcalluser.h |  7 +++++--
> man/vzctl.8      |  2 ++
> src/lib/iptables.c |  1 +
> 5 files changed, 8 insertions(+), 4 deletions(-)
>
> diff --git a/etc/vz.conf b/etc/vz.conf
> index f1e9a42..5f23b90 100644
> --- a/etc/vz.conf
> +++ b/etc/vz.conf
> @@ -29,5 +29,5 @@ DEF_OSTEMPLATE="fedora-core-4"
> ## Load vzwdog module
> VZWDOG="no"
>
> -IPTABLES="ipt_REJECT ipt_tos ipt_limit ipt_multiport iptable_filter
> iptable_mangle ipt_TCPMSS ipt_tcpmss ipt_ttl ipt_length"
> +IPTABLES="xt_tcpudp ipt_REJECT ipt_tos ipt_limit ipt_multiport
> iptable_filter iptable_mangle ipt_TCPMSS ipt_tcpmss ipt_ttl ipt_length"
>
> diff --git a/include/linux/vzcalluser.h b/include/linux/vzcalluser.h
> index 360f246..31100b1 100644
> --- a/include/linux/vzcalluser.h
> +++ b/include/linux/vzcalluser.h
> @@ -97,6 +97,7 @@ #define VE_IP_NAT_FTP_MOD (1U<<21)
> #define VE_IP_NAT_IRC_MOD (1U<<22)
> #define VE_IP_TARGET_REDIRECT_MOD (1U<<23)
> #define VE_IP_MATCH_MAC_MOD (1U<<24)
> +#define VE_IP_MATCH_TCPUDP_MOD (1U<<25)
>
> /* these masks represent modules with their dependences */
> #define VE_IP_IPTABLES (VE_IP_IPTABLES_MOD)
> @@ -143,13 +144,15 @@ #define VE_IP_NAT_FTP (VE_IP_NAT_FTP_MO
> #define VE_IP_NAT_IRC (VE_IP_NAT_IRC_MOD \
>
> | VE_IP_NAT | VE_IP_CONNTRACK_IRC)
```

```

>
> #define VE_IP_TARGET_REDIRECT (VE_IP_TARGET_REDIRECT_MOD \
> - | VE_IP_NAT)
> + | VE_IP_NAT)
> #define VE_IP_MATCH_MAC (VE_IP_MATCH_MAC_MOD \
> + | VE_IP_IPTABLES)
> +#define VE_IP_MATCH_TCPUDP (VE_IP_MATCH_TCPUDP_MOD \
>
>     | VE_IP_IPTABLES)
>
> /* safe iptables mask to be used by default */
> #define VE_IP_DEFAULT \
> - (VE_IP_IPTABLES | \
> + (VE_IP_IPTABLES | VE_IP_MATCH_TCPUDP | \
> VE_IP_FILTER | VE_IP_MANGLE | \
> VE_IP_MATCH_LIMIT | VE_IP_MATCH_MULTIPOINT | \
> VE_IP_MATCH_TOS | VE_IP_TARGET_REJECT | \
> diff --git a/man/vzctl.8 b/man/vzctl.8
> index eec3155..0fb6eaa 100644
> --- a/man/vzctl.8
> +++ b/man/vzctl.8
> @@ -311,7 +311,7 @@ You can use the following values for \fI
> \fIlipt_length\fR, \fIflip_conntrack\fR, \fIflip_conntrack_ftp\fR,
> \fIflip_conntrack_irc\fR, \fIlipt_conntrack\fR, \fIlipt_state\fR,
> \fIlipt_helper\fR, \fIliptable_nat\fR, \fIflip_nat_ftp\fR, \fIflip_nat_irc\fR,
> -\fIlipt_REDIRECT\fR \fIxt_mac\fR.
> +\fIlipt_REDIRECT\fR, \fIxt_mac\fR, \fIxt_tcpudp\fR.
> .TP
> \fBNetwork devices control parameters\fR
> .IP "\fB--netdev _add\fR \fIname\fR"
> diff --git a/src/lib/iptables.c b/src/lib/iptables.c
> index 8d07264..93bbbc5 100644
> --- a/src/lib/iptables.c
> +++ b/src/lib/iptables.c
> @@ -52,6 +52,7 @@ #endif
> #ifdef VE_IP_MATCH_MAC
> {"xt_mac", VE_IP_MATCH_MAC_MOD, VE_IP_MATCH_MAC},
> #endif
> + {"xt_tcpudp", VE_IP_MATCH_TCPUDP_MOD, VE_IP_MATCH_TCPUDP},
> #endif /* VZCTL_ENV_CREATE_DATA */
> {NULL, 0}
> };
>
> --
> 1.4.2.GIT

```

--
Thanks,

Dmitry.

Subject: Re: [PATCH] vzctl: Add support for xt_tcpudp (aka ipt_tcp/ipt_udp) module introduced by x_tables

Posted by [lakostis](#) on Thu, 09 Nov 2006 12:56:56 GMT

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

> On Thursday 09 November 2006 10:48, Konstantin A. Lepikhov wrote:
>> In recent 2.6 kernels (2.6.16+) netfilter code is changed by new
>> x_tables
>> (Abstraction layer for {ip,ip6,arp}_tables). This patch updates iptables
>> wrapper code in vzctl for this changes (adding xt_tcpudp support).
> xt_tcpudp previously were inside ip_tables module. So, for compatibility,
> current kernels handles this module by VE_IP_IPTABLES mask.
But this module isn't been autoloaded. Let's try service vz start, and see
that xt_tcpudp is absent. So looks like it ovz kernel patch bug.

--
WBR et al.

Subject: Re: [PATCH] vzctl: Add support for xt_tcpudp (aka ipt_tcp/ipt_udp) module introduced by x_tables

Posted by [Mishin Dmitry](#) on Thu, 09 Nov 2006 14:05:23 GMT

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

On Thursday 09 November 2006 15:56, Konstantin A. Lepikhov wrote:

>
>> On Thursday 09 November 2006 10:48, Konstantin A. Lepikhov wrote:
>>> In recent 2.6 kernels (2.6.16+) netfilter code is changed by new
>>> x_tables
>>> (Abstraction layer for {ip,ip6,arp}_tables). This patch updates iptables
>>> wrapper code in vzctl for this changes (adding xt_tcpudp support).
>>
>> xt_tcpudp previously were inside ip_tables module. So, for compatibility,
>> current kernels handles this module by VE_IP_IPTABLES mask.
>
> But this module isn't been autoloaded. Let's try service vz start, and see
> that xt_tcpudp is absent. So looks like it ovz kernel patch bug.
It is not kernel bug. And it is not a bug at all. It is issue due to the one
vzctl and multiple kernels with and without this module. And I don't want to
add module dependencies in order to get xt_tcpudp loaded with x_tables.
When we'll switch stable branch to >=2.6.16 or add vzctl devel branch, then we
fix it as we done for all other modules. Right now, it is unnecessary work.

Workaround is simple, but if you prepare patches for vzctl AND kernel with compatibility issues in mind - welcome!

--

Thanks,
Dmitry.
