
Subject: Re: [PATCH 1/2] virtualized ipt_REDIRECT

Posted by [dev](#) on Thu, 16 Mar 2006 13:07:02 GMT

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

Jason,

I fixed error path in do_ve_iptables().

Here is the patch which is committed in 2.6.8.

It will be also committed in coming 2.6.16.

Thanks,

Kirill

> Missed two defines when bringing the patch forward from 2.6.8.

> How embarrassing...

>

> Patch from Jason (jstubbs@work-at.co.jp):

> This patch virtualizes the ipt_REDIRECT iptables module.

>

> --

> Jason Stubbs

>

>

> -----

>

> diff -ur linux-2.6.15-openvz-025.014.orig/include/linux/nfcalls.h

linux-2.6.15-openvz-025.014/include/linux/nfcalls.h

> --- linux-2.6.15-openvz-025.014.orig/include/linux/nfcalls.h 2006-03-03 14:36:32.560909760
+0900

> +++ linux-2.6.15-openvz-025.014/include/linux/nfcalls.h 2006-03-03 15:20:12.223660488 +0900

> @@ -143,6 +143,7 @@

> DECL_KSYM_MODULE(iptable_nat);

> DECL_KSYM_MODULE(ip_nat_ftp);

> DECL_KSYM_MODULE(ip_nat_irc);

> +DECL_KSYM_MODULE(ipt_REDIRECT);

>

> struct sk_buff;

>

> @@ -170,6 +171,7 @@

> DECL_KSYM_CALL(int, init_iptable_nat, (void));

> DECL_KSYM_CALL(int, init_iptable_nat_ftp, (void));

> DECL_KSYM_CALL(int, init_iptable_nat_irc, (void));

> +DECL_KSYM_CALL(int, init_iptable_REDIRECT, (void));

> DECL_KSYM_CALL(void, fini_iptable_nat_irc, (void));

> DECL_KSYM_CALL(void, fini_iptable_nat_ftp, (void));

> DECL_KSYM_CALL(void, fini_iptable_nat, (void));

> @@ -194,6 +196,7 @@

> DECL_KSYM_CALL(void, fini_iptable_mangle, (void));

```

> DECL_KSYM_CALL(void, fini_iptables, (void));
> DECL_KSYM_CALL(void, fini_netfilter, (void));
> +DECL_KSYM_CALL(void, fini_iptable_REDIRECT, (void));
>
> DECL_KSYM_CALL(void, ipt_flush_table, (struct ipt_table *table));
> #endif /* CONFIG_VE_IPTABLES */
> diff -ur linux-2.6.15-openvz-025.014.orig/include/linux/ve_proto.h
linux-2.6.15-openvz-025.014/include/linux/ve_proto.h
> --- linux-2.6.15-openvz-025.014.orig/include/linux/ve_proto.h 2006-03-03 14:36:32.560909760
+0900
> +++ linux-2.6.15-openvz-025.014/include/linux/ve_proto.h 2006-03-03 14:38:42.914093064
+0900
> @@ -55,6 +55,7 @@
> extern int init_iptable_multiport(void);
> extern int init_iptable_tos(void);
> extern int init_iptable_REJECT(void);
> +extern int init_iptable_REDIRECT(void);
> extern void fini_netfilter(void);
> extern int fini_iptables(void);
> extern int fini_iptable_filter(void);
> @@ -62,6 +63,7 @@
> extern int fini_iptable_multiport(void);
> extern int fini_iptable_tos(void);
> extern int fini_iptable_REJECT(void);
> +extern int fini_iptable_REDIRECT(void);
> #endif
>
> #define VE_HOOK_INIT 0
> diff -ur linux-2.6.15-openvz-025.014.orig/include/linux/vzcalluser.h
linux-2.6.15-openvz-025.014/include/linux/vzcalluser.h
> --- linux-2.6.15-openvz-025.014.orig/include/linux/vzcalluser.h 2006-03-03 14:36:32.561909608
+0900
> +++ linux-2.6.15-openvz-025.014/include/linux/vzcalluser.h 2006-03-03 14:39:39.544483936
+0900
> @@ -80,6 +80,7 @@
> #define VE_IP_NAT_MOD (1U<<20)
> #define VE_IP_NAT_FTP_MOD (1U<<21)
> #define VE_IP_NAT_IRC_MOD (1U<<22)
> +#define VE_IP_TARGET_REDIRECT_MOD (1U<<23)
>
> /* these masks represent modules with their dependences */
> #define VE_IP_IPTABLES (VE_IP_IPTABLES_MOD)
> @@ -125,6 +126,8 @@
> | VE_IP_NAT | VE_IP_CONNTRACK_FTP)
> #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)

```

```

>
> /* safe iptables mask to be used by default */
> #define VE_IP_DEFAULT \
> diff -ur linux-2.6.15-openvz-025.014.orig/kernel/ve.c linux-2.6.15-openvz-025.014/kernel/ve.c
> --- linux-2.6.15-openvz-025.014.orig/kernel/ve.c 2006-03-03 14:36:33.253804424 +0900
> +++ linux-2.6.15-openvz-025.014/kernel/ve.c 2006-03-03 14:41:02.759833280 +0900
> @@ -75,6 +75,7 @@
> INIT_KSYM_MODULE(iptable_nat);
> INIT_KSYM_MODULE(ip_nat_ftp);
> INIT_KSYM_MODULE(ip_nat_irc);
> +INIT_KSYM_MODULE(ipt_REDIRECT);
>
> INIT_KSYM_CALL(int, init_netfilter, (void));
> INIT_KSYM_CALL(int, init_iptables, (void));
> @@ -100,6 +101,7 @@
> INIT_KSYM_CALL(int, init_iptable_nat, (void));
> INIT_KSYM_CALL(int, init_iptable_nat_ftp, (void));
> INIT_KSYM_CALL(int, init_iptable_nat_irc, (void));
> +INIT_KSYM_CALL(int, init_iptable_REDIRECT, (void));
> INIT_KSYM_CALL(void, fini_iptable_nat_irc, (void));
> INIT_KSYM_CALL(void, fini_iptable_nat_ftp, (void));
> INIT_KSYM_CALL(void, fini_iptable_nat, (void));
> @@ -124,6 +126,7 @@
> INIT_KSYM_CALL(void, fini_iptable_mangle, (void));
> INIT_KSYM_CALL(void, fini_iptables, (void));
> INIT_KSYM_CALL(void, fini_netfilter, (void));
> +INIT_KSYM_CALL(void, fini_iptable_REDIRECT, (void));
>
> INIT_KSYM_CALL(void, ipt_flush_table, (struct ipt_table *table));
> #endif
> diff -ur linux-2.6.15-openvz-025.014.orig/kernel/vecalls.c
linux-2.6.15-openvz-025.014/kernel/vecalls.c
> --- linux-2.6.15-openvz-025.014.orig/kernel/vecalls.c 2006-03-03 14:36:33.253804424 +0900
> +++ linux-2.6.15-openvz-025.014/kernel/vecalls.c 2006-03-03 14:42:35.671708528 +0900
> @@ -1592,6 +1592,13 @@
> if (err < 0)
> goto err_iptable_length;
> #endif
> +#if defined(CONFIG_IP_NF_TARGET_REDIRECT) || \
> + defined(CONFIG_IP_NF_TARGET_REDIRECT_MODULE)
> + err = KSYMIPTINIT(init_mask, ve, VE_IP_TARGET_REDIRECT,
> + ipt_REDIRECT, init_iptable_REDIRECT, ());
> + if (err < 0)
> + goto err_iptable_REDIRECT;
> +#endif
> return 0;
>
> /* ----- */

```

```

> @@ -1732,6 +1739,12 @@
>   ip_tables, fini_iptables, ());
>   err_iptables:
>   #endif
>   +#if defined(CONFIG_IP_NF_TARGET_REDIRECT) || \
>   +   defined(CONFIG_IP_NF_TARGET_REDIRECT_MODULE)
>   + KSYMIPTFINI(ve->_iptables_modules, VE_IP_TARGET_REDIRECT,
>   +   ipt_REDIRECT, fini_iphtable_REDIRECT, ());
>   +err_iphtable_REDIRECT:
>   +#endif
>   ve->_iptables_modules = 0;
>
>   return err;
> diff -ur linux-2.6.15-openvz-025.014.orig/net/ipv4/netfilter/ipt_REDIRECT.c
linux-2.6.15-openvz-025.014.orig/net/ipv4/netfilter/ipt_REDIRECT.c
> --- linux-2.6.15-openvz-025.014.orig/net/ipv4/netfilter/ipt_REDIRECT.c 2006-03-03
14:36:33.952698176 +0900
> +++ linux-2.6.15-openvz-025.014.orig/net/ipv4/netfilter/ipt_REDIRECT.c 2006-03-03
14:47:44.502759072 +0900
> @@ -17,6 +17,7 @@
>   #include <linux/inetdevice.h>
>   #include <net/protocol.h>
>   #include <net/checksum.h>
>   +#include <linux/nfcalls.h>
>   #include <linux/netfilter_ipv4.h>
>   #include <linux/netfilter_ipv4/ip_nat_rule.h>
>
> @@ -25,7 +26,7 @@
>   MODULE_DESCRIPTION("iptables REDIRECT target module");
>
>   #if 0
>   -#define DEBUGP printk
>   +#define DEBUGP ve_printk
>   #else
>   #define DEBUGP(format, args...)
>   #endif
> @@ -119,14 +120,36 @@
>   .me = THIS_MODULE,
>   };
>
>   +int init_iphtable_REDIRECT(void)
>   +{
>   +   return virt_ippt_register_target(&redirect_reg);
>   +}
>   +
>   +void fini_iphtable_REDIRECT(void)
>   +{
>   +   virt_ippt_unregister_target(&redirect_reg);

```

```

> +}
> +
> static int __init init(void)
> {
> - return ipt_register_target(&redirect_reg);
> + int err;
> +
> + err = init_iptable_REDIRECT();
> + if (err < 0)
> + return err;
> +
> + KSYMRESOLVE(init_iptable_REDIRECT);
> + KSYMRESOLVE(fini_iptable_REDIRECT);
> + KSYMMODRESOLVE(ipt_REDIRECT);
> + return 0;
> }
>
> static void __exit fini(void)
> {
> - ipt_unregister_target(&redirect_reg);
> + KSYMMODUNRESOLVE(ipt_REDIRECT);
> + KSYMUNRESOLVE(init_iptable_REDIRECT);
> + KSYMUNRESOLVE(fini_iptable_REDIRECT);
> + fini_iptable_REDIRECT();
> }
>
> module_init(init);
>
>
> -----
>
--- ./include/linux/nfcalls.h.iptredir 2006-03-03 14:43:05.000000000 +0300
+++ ./include/linux/nfcalls.h 2006-03-16 16:06:33.000000000 +0300
@@ -155,6 +155,7 @@ DECL_KSYM_MODULE(ipt_helper);
DECL_KSYM_MODULE(iptable_nat);
DECL_KSYM_MODULE(ip_nat_ftp);
DECL_KSYM_MODULE(ip_nat_irc);
+DECL_KSYM_MODULE(ipt_REDIRECT);
#endif

struct sk_buff;
@@ -182,6 +183,7 @@ DECL_KSYM_CALL(int, init_iptable_helper,
DECL_KSYM_CALL(int, init_iptable_nat, (void));
DECL_KSYM_CALL(int, init_iptable_nat_ftp, (void));
DECL_KSYM_CALL(int, init_iptable_nat_irc, (void));
+DECL_KSYM_CALL(int, init_iptable_REDIRECT, (void));
DECL_KSYM_CALL(void, fini_iptable_nat_irc, (void));
DECL_KSYM_CALL(void, fini_iptable_nat_ftp, (void));

```

```

DECL_KSYM_CALL(void, fini_iptable_nat, (void));
@@ -205,6 +207,7 @@ DECL_KSYM_CALL(void, fini_iptable_filter
DECL_KSYM_CALL(void, fini_iptable_mangle, (void));
DECL_KSYM_CALL(void, fini_iptables, (void));
DECL_KSYM_CALL(void, fini_netfilter, (void));
+DECL_KSYM_CALL(void, fini_iptable_REDIRECT, (void));

DECL_KSYM_CALL(void, ipt_flush_table, (struct ipt_table *table));
#endif /* CONFIG_VE_IPTABLES */
--- ./include/linux/vzcalluser.h.iptredir 2006-03-03 14:43:11.000000000 +0300
+++ ./include/linux/vzcalluser.h 2006-03-16 16:05:46.000000000 +0300
@@ -82,6 +82,7 @@ struct vzctl_ve_netdev {
#define VE_IP_NAT_MOD (1U<<20)
#define VE_IP_NAT_FTP_MOD (1U<<21)
#define VE_IP_NAT_IRC_MOD (1U<<22)
+#define VE_IP_TARGET_REDIRECT_MOD (1U<<23)

/* these masks represent modules with their dependences */
#define VE_IP_IPTABLES (VE_IP_IPTABLES_MOD)
@@ -127,6 +128,8 @@ struct vzctl_ve_netdev {
| VE_IP_NAT | VE_IP_CONNTRACK_FTP)
#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)

/* safe iptables mask to be used by default */
#define VE_IP_DEFAULT \
--- ./kernel/ve.c.iptredir 2006-03-03 14:43:05.000000000 +0300
+++ ./kernel/ve.c 2006-03-16 16:05:46.000000000 +0300
@@ -75,6 +75,7 @@ INIT_KSYM_MODULE(ipt_helper);
INIT_KSYM_MODULE(iptable_nat);
INIT_KSYM_MODULE(ip_nat_ftp);
INIT_KSYM_MODULE(ip_nat_irc);
+INIT_KSYM_MODULE(ipt_REDIRECT);

INIT_KSYM_CALL(int, init_netfilter, (void));
INIT_KSYM_CALL(int, init_iptables, (void));
@@ -99,6 +100,7 @@ INIT_KSYM_CALL(int, init_iptable_helper,
INIT_KSYM_CALL(int, init_iptable_nat, (void));
INIT_KSYM_CALL(int, init_iptable_nat_ftp, (void));
INIT_KSYM_CALL(int, init_iptable_nat_irc, (void));
+INIT_KSYM_CALL(int, init_iptable_REDIRECT, (void));
INIT_KSYM_CALL(void, fini_iptable_nat_irc, (void));
INIT_KSYM_CALL(void, fini_iptable_nat_ftp, (void));
INIT_KSYM_CALL(void, fini_iptable_nat, (void));
@@ -122,6 +124,7 @@ INIT_KSYM_CALL(void, fini_iptable_filter
INIT_KSYM_CALL(void, fini_iptable_mangle, (void));

```

```

INIT_KSYM_CALL(void, fini_iptables, (void));
INIT_KSYM_CALL(void, fini_netfilter, (void));
+INIT_KSYM_CALL(void, fini_iptable_REDIRECT, (void));

INIT_KSYM_CALL(void, ipt_flush_table, (struct ipt_table *table));
#endif
--- ./kernel/vecalls.c.iptredir 2006-03-03 16:16:06.000000000 +0300
+++ ./kernel/vecalls.c 2006-03-16 16:07:25.000000000 +0300
@@ -1577,11 +1577,24 @@ static int do_ve_iptables(struct ve_stru
    if (err < 0)
        goto err_iptable_length;
#endif
+#if defined(CONFIG_IP_NF_TARGET_REDIRECT) || \
+ defined(CONFIG_IP_NF_TARGET_REDIRECT_MODULE)
+ err = KSYMIPINIT(init_mask, ve, VE_IP_TARGET_REDIRECT,
+ ipt_REDIRECT, init_iptable_REDIRECT, ());
+ if (err < 0)
+ goto err_iptable_REDIRECT;
+#endif
    return 0;

/* ----- */

cleanup:
+#if defined(CONFIG_IP_NF_TARGET_REDIRECT) || \
+ defined(CONFIG_IP_NF_TARGET_REDIRECT_MODULE)
+ KSYMIPFINI(ve->_iptables_modules, VE_IP_TARGET_REDIRECT,
+ ipt_REDIRECT, fini_iptable_REDIRECT, ());
+err_iptable_REDIRECT:
+#endif
#if defined(CONFIG_IP_NF_MATCH_LENGTH) || \
    defined(CONFIG_IP_NF_MATCH_LENGTH_MODULE)
    KSYMIPFINI(ve->_iptables_modules, VE_IP_MATCH_LENGTH,
--- ./net/ipv4/netfilter/ipt_REDIRECT.c.iptredir 2004-08-14 14:55:10.000000000 +0400
+++ ./net/ipv4/netfilter/ipt_REDIRECT.c 2006-03-16 16:11:10.000000000 +0300
@@ -17,6 +17,7 @@
#include <linux/inetdevice.h>
#include <net/protocol.h>
#include <net/checksum.h>
+#include <linux/nfcalls.h>
#include <linux/netfilter_ipv4.h>
#include <linux/netfilter_ipv4/ip_nat_rule.h>

@@ -25,7 +26,7 @@ MODULE_AUTHOR("Netfilter Core Team <core
MODULE_DESCRIPTION("iptables REDIRECT target module");

#if 0
-#define DEBUGP printk

```

```

+#define DEBUGP ve_printk
#else
#define DEBUGP(format, args...)
#endif
@@ -115,14 +116,36 @@ static struct ipt_target redirect_reg =
.me = THIS_MODULE,
};

+int init_iptable_REDIRECT(void)
+{
+ return visible_ipt_register_target(&redirect_reg);
+}
+
+void fini_iptable_REDIRECT(void)
+{
+ visible_ipt_unregister_target(&redirect_reg);
+}
+
+static int __init init(void)
+{
- return ipt_register_target(&redirect_reg);
+ int err;
+
+ err = init_iptable_REDIRECT();
+ if (err < 0)
+ return err;
+
+ KSYMRESOLVE(init_iptable_REDIRECT);
+ KSYMRESOLVE(fini_iptable_REDIRECT);
+ KSYMMODRESOLVE(ipt_REDIRECT);
+ return 0;
+}

static void __exit fini(void)
{
- ipt_unregister_target(&redirect_reg);
+ KSYMMODUNRESOLVE(ipt_REDIRECT);
+ KSYMUNRESOLVE(init_iptable_REDIRECT);
+ KSYMUNRESOLVE(fini_iptable_REDIRECT);
+ fini_iptable_REDIRECT();
+}

module_init(init);

```
