Subject: [PATCH] bridge: Reset IPCB on forward non-local packets in br_handle_frame_finish()
Posted by Vasily Averin on Wed, 02 Nov 2011 19:08:57 GMT
View Forum Message <> Reply to Message

if dst is not local br_handle_frame_finish() does not clone original skb and
forgets to reset IPCB before return to IP stack. it can lead to stack corruption
in icmp_send()

Signed-off-by: Vasily Averin <vvs@sw.ru>
---
 net/bridge/br_input.c | 5 +++--
 1 files changed, 3 insertions(+), 2 deletions(-)

diff --git a/net/bridge/br_input.c b/net/bridge/br_input.c
index f06ee39..6be8d00 100644
--- a/net/bridge/br_input.c
+++ b/net/bridge/br_input.c
@@ -93,10 +93,11 @@ int br_handle_frame_finish(struct sk_buff *skb)
    skb2 = skb;

   br->dev->stats.multicast++;
- } else if ((dst = __br_fdb_get(br, dest)) && dst->is_local) {
+ } else if ((dst = __br_fdb_get(br, dest)) != NULL) {
   skb2 = skb;
   /* Do not forward the packet since it's local. */
- skb = NULL;
+ if (dst->is_local) {
+   skb = NULL;
  }

  if (skb) {
-- 1.7.5.4

Subject: Re: [PATCH] bridge: Reset IPCB on forward non-local packets in br_handle_frame_finish()
Posted by Vasily Averin on Wed, 02 Nov 2011 19:11:38 GMT
View Forum Message <> Reply to Message

On 11/02/2011 11:08 PM, Vasily Averin wrote:
> if dst is not local br_handle_frame_finish() does not clone original skb and
> forgets to reset IPCB before return to IP stack. it can lead to stack corruption
> in icmp_send()

example of stack corruption
http://bugzilla.openvz.org/show_bug.cgi?id=2016

> Signed-off-by: Vasily Averin <vvs@sw.ru>
> ---
> net/bridge/br_input.c | 5 +++--
> 1 files changed, 3 insertions(+), 2 deletions(-)
>
> diff --git a/net/bridge/br_input.c b/net/bridge/br_input.c
> index f06ee39..6be8d00 100644
> --- a/net/bridge/br_input.c
> +++ b/net/bridge/br_input.c
> @@ -93,10 +93,11 @@ int br_handle_frame_finish(struct sk_buff *skb)
>   skb2 = skb;
>
>   br->dev->stats.multicast++;
> - } else if ((dst = __br_fdb_get(br, dest)) && dst->is_local) {
> + } else if ((dst = __br_fdb_get(br, dest)) != NULL) {
>   skb2 = skb;
>   /* Do not forward the packet since it's local. */
> - skb = NULL;
> + if (dst->is_local) {
> +   skb = NULL;
>   }
>
>   if (skb) {
> -- 1.7.5.4

Subject: Re: [PATCH] bridge: Reset IPCB on forward non-local packets in br_handle_frame_finish()
Posted by Stephen Hemminger on Wed, 02 Nov 2011 19:31:06 GMT
View Forum Message <> Reply to Message

On Wed, 02 Nov 2011 23:08:57 +0400
Vasily Averin <vvs@parallels.com> wrote:

> if dst is not local br_handle_frame_finish() does not clone original skb and
> forgets to reset IPCB before return to IP stack. it can lead to stack corruption
> in icmp_send()
>
> Signed-off-by: Vasily Averin <vvs@sw.ru>
> ---
> net/bridge/br_input.c | 5 +++--
> 1 files changed, 3 insertions(+), 2 deletions(-)
>
> diff --git a/net/bridge/br_input.c b/net/bridge/br_input.c
> index f06ee39..6be8d00 100644
> --- a/net/bridge/br_input.c
> +++ b/net/bridge/br_input.c
> @@ -93,10 +93,11 @@ int br_handle_frame_finish(struct sk_buff *skb)

>   skb2 = skb;
>
>   br->dev->stats.multicast++;
> - } else if ((dst = __br_fdb_get(br, dest)) && dst->is_local) {
> + } else if ((dst = __br_fdb_get(br, dest)) != NULL) {
>   skb2 = skb;
>   /* Do not forward the packet since it's local. */
> - skb = NULL;
> + if (dst->is_local) {
> +   skb = NULL;
>   }
>
>   if (skb) {

What kernel version are you using? There were several previous fixes
in br_netfilter to deal with this type of issue over the last year.

Subject: Re: [PATCH] bridge: Reset IPCB on forward non-local packets in br_handle_frame_finish()
Posted by Vasily Averin on Wed, 02 Nov 2011 20:03:53 GMT
View Forum Message <> Reply to Message

On 11/02/2011 11:31 PM, Stephen Hemminger wrote:
> On Wed, 02 Nov 2011 23:08:57 +0400
> Vasily Averin <vvs@parallels.com> wrote:
>
>> if dst is not local br_handle_frame_finish() does not clone original skb and
>> forgets to reset IPCB before return to IP stack. it can lead to stack corruption
>> in icmp_send()

> What kernel version are you using? There were several previous fixes
> in br_netfilter to deal with this type of issue over the last year.

Originally it was noticed on RHEL6-based kernel

You are right, in mainline this issue was fixed in br_nf_forward_ip() long time ago.

thank you,
 Vasily Averin

Subject: Re: [PATCH] bridge: Reset IPCB on forward non-local packets in br_handle_frame_finish()
Posted by davem on Wed, 02 Nov 2011 20:09:05 GMT
View Forum Message <> Reply to Message

From: Vasily Averin <vvs@parallels.com>
Date: Wed, 02 Nov 2011 23:08:57 +0400

> if dst is not local br_handle_frame_finish() does not clone original skb and
> forgets to reset IPCB before return to IP stack. it can lead to stack corruption
> in icmp_send()
>
> Signed-off-by: Vasily Averin <vvs@sw.ru>

Nothing is worse than posting a patch that doesn't even compile.

And I really mean _nothing_.
