
Subject: [PATCH][INET]: Fix accidentally broken inet(6)_hash_connect's port offset calculations.

Posted by [Pavel Emelianov](#) on Mon, 04 Feb 2008 08:10:37 GMT

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

The port offset calculations depend on the protocol family, but, as Adrian noticed, I broke this logic with the commit

5ee31fc1ecdcbc234c8c56dcacef87c8e09909d8

[INET]: Consolidate inet(6)_hash_connect.

Return this logic back, by passing the port offset directly into the consolidated function.

Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

Noticed-by: Adrian Bunk <bunk@kernel.org>

diff --git a/include/net/inet_hashtables.h b/include/net/inet_hashtables.h

index 48ac620..97dc35a 100644

--- a/include/net/inet_hashtables.h

+++ b/include/net/inet_hashtables.h

@@ -389,7 +389,7 @@ static inline struct sock *inet_lookup(struct net *net,
}

extern int __inet_hash_connect(struct inet_timewait_death_row *death_row,
- struct sock *sk,
+ struct sock *sk, u32 port_offset,
int (*check_established)(struct inet_timewait_death_row *,
struct sock *, __u16, struct inet_timewait_sock **),
void (*hash)(struct sock *sk));

diff --git a/net/ipv4/inet_hashtables.c b/net/ipv4/inet_hashtables.c

index 90f422c..9cac6c0 100644

--- a/net/ipv4/inet_hashtables.c

+++ b/net/ipv4/inet_hashtables.c

@@ -398,7 +398,7 @@ out:

EXPORT_SYMBOL_GPL(inet_unhash);

int __inet_hash_connect(struct inet_timewait_death_row *death_row,
- struct sock *sk,
+ struct sock *sk, u32 port_offset,
int (*check_established)(struct inet_timewait_death_row *,
struct sock *, __u16, struct inet_timewait_sock **),
void (*hash)(struct sock *sk))

@@ -413,7 +413,7 @@ int __inet_hash_connect(struct inet_timewait_death_row *death_row,
if (!snum) {
int i, remaining, low, high, port;

```

static u32 hint;
- u32 offset = hint + inet_sk_port_offset(sk);
+ u32 offset = hint + port_offset;
  struct hlist_node *node;
  struct inet_timewait_sock *tw = NULL;

@@ -502,7 +502,7 @@ EXPORT_SYMBOL_GPL(__inet_hash_connect);
int inet_hash_connect(struct inet_timewait_death_row *death_row,
    struct sock *sk)
{
- return __inet_hash_connect(death_row, sk,
+ return __inet_hash_connect(death_row, sk, inet_sk_port_offset(sk),
    __inet_check_established, __inet_hash_nolisten);
}

diff --git a/net/ipv6/inet6_hashtables.c b/net/ipv6/inet6_hashtables.c
index 43f3993..99fd25f 100644
--- a/net/ipv6/inet6_hashtables.c
+++ b/net/ipv6/inet6_hashtables.c
@@ -236,7 +236,7 @@ static inline u32 inet6_sk_port_offset(const struct sock *sk)
int inet6_hash_connect(struct inet_timewait_death_row *death_row,
    struct sock *sk)
{
- return __inet_hash_connect(death_row, sk,
+ return __inet_hash_connect(death_row, sk, inet6_sk_port_offset(sk),
    __inet6_check_established, __inet6_hash);
}

```

Subject: Re: [PATCH][INET]: Fix accidentally broken inet(6)_hash_connect's port offset calculations.

Posted by [davem](#) on Tue, 05 Feb 2008 11:14:44 GMT

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

From: Pavel Emelyanov <xemul@openvz.org>

Date: Mon, 04 Feb 2008 11:10:37 +0300

```

> The port offset calculations depend on the protocol family, but,
> as Adrian noticed, I broke this logic with the commit
>
> 5ee31fc1ecdcbc234c8c56dcacef87c8e09909d8
> [INET]: Consolidate inet(6)_hash_connect.
>
> Return this logic back, by passing the port offset directly into
> the consolidated function.
>
> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>
> Noticed-by: Adrian Bunk <bunk@kernel.org>

```

Applied, thanks.
