

strongSwan - Issue #939

UDP Encapsulation for IPv6 traffic

27.04.2015 16:04 - huanghanzhao huang

Status: Feedback	
Priority: Normal	
Assignee: Noel Kuntze	
Category: kernel	
Affected version:	Resolution:
Description	
<p>I tried to setup an IPv6 connection with following scenario</p> <p>10.0.0.1[Server] fec0::2 ---[nat-t]-- fec0::1[Client] when the Client's IP pass through the nat-t, the nat-t will change the Client's IP to other IP.</p> <p>And the configuration as followed -</p> <pre>[Client side] conn home left=fec0::1 keyexchange=ikev2 authby=secret right=fec0::2 rightsubnet=0.0.0.0/0 auto=add</pre> <pre>[Server side] conn psk left=fec0::2 leftsubnet=0.0.0.0/0 keyexchange=ikev2 authby=secret right=%any auto=add</pre> <p>However, after IKE SA established, server side cannot setup SAD successfully. It shows up following error log</p> <pre>received netlink error: Invalid argument (22) Unable to add SAD entry with SPI c2dc9aa0 received netlink error: Invalid argument (22) Unable to add SAD entry with SPI ced1801e Unable to install inbound and outbound IPsec SA (SAD) in kernel</pre>	
Related issues:	
Related to Feature #892: Android client and ipv6 gateway	Feedback 15.03.2015
Related to Issue #2231: Cannot create IPv4-via-IPv6 tunnel	Closed
Related to Issue #2354: Approval from apple regarding an application using st...	Closed

History

#1 - 28.04.2015 04:08 - huanghanzhao huang

PS:
My Linux kernel version is Linux 2.6.32.27.
And the strongswan version is 4.6.4.

#2 - 29.04.2015 10:02 - huanghanzhao huang

I try to set the SAs into Linux, just as the following:
[root@localhost ~]# ip xfrm state add src fd14:828:ba69:2::2 dst fd14:828:ba69:1:21c:f0ff:fefa:f3c0 proto esp spi 0x00000301 mode tunnel auth md5 0x96358c90783bbfa3d7b196ceabe0536b enc des3_ede 0xf6ddb555acfd9d77b03ea3843f2653255afe8eb5573965df encap espinudp 4500 4500

```
fd13:828:ba69:2::2
RTNETLINK answers: Invalid argument
[root@localhost ~]# setkey -D
No SAD entries.
[root@localhost ~]#
```

Can you tell me why?

#3 - 29.04.2015 10:05 - Martin Willi

Hi,

AFAIK the Linux kernel does not support UDP encapsulation for IPv6 SAs.

Regards
Martin

#4 - 29.04.2015 14:14 - nanjian5 Lee

```
/* enable UDP decapsulation for NAT-T sockets */
if (port &this->natt &&
!hydra->kernel_interface->enable_udp_decap(hydra->kernel_interface,
skt, family, this->natt)) {
DBG1(DBG_NET, "enabling UDP decapsulation for %s on port %d failed",
family AF_INET ? "IPv4" : "IPv6", this->natt);
}
```

when I run IPv6 and strongswan 5.1.1 with linux kernel 3.12.8, I can not see the error message above. Does it means that kernel supports UDP encapsulation for IPv6 SAs?

#5 - 29.04.2015 14:21 - nanjian5 Lee

```
/*
 * This is an encapsulation socket so pass the skb to
 * the socket's udp_encap_rcv() hook. Otherwise, just
 * fall through and pass this up the UDP socket.
 * up->encap_rcv() returns the following value:
 * =0 if skb was successfully passed to the encap
 *   handler or was discarded by it.
 * >0 if skb should be passed on to UDP.
 * <0 if skb should be resubmitted as proto -N
 */

/* if we're overly short, let UDP handle it */
encap_rcv = ACCESS_ONCE(up->encap_rcv);
if (skb->len > sizeof(struct udphdr) && encap_rcv != NULL) {
    int ret;

    ret = encap_rcv(sk, skb);
    if (ret <= 0) {
        UDP_INC_STATS_BH(sock_net(sk),
            UDP_MIB_INDATAGRAMS,
            is_udplite);
        return -ret;
    }
}
```

I dig into the source code of linux kernel 3.12.8 and find the above code in net\ipv6\udp.c

Dose it mean that kernel support UDP encapsulation? Then it should support IPv6 NAT_T SA.

#6 - 29.04.2015 16:24 - Martin Willi

when I run IPv6 and strongswan 5.1.1 with linux kernel 3.12.8, I can not see the error message above. Does it means that kernel supports UDP encapsulation for IPv6 SAs?

Most likely.

I dig into the source code of linux kernel 3.12.8 and find the above code in net\ipv6\udp.c

The mentioned code has been added in Linux 3.5, so possible that this is supported now. I've never tested it, though.

Regards
Martin

#7 - 30.04.2015 05:27 - nanjian5 Lee

received netlink error: Invalid argument (22)
Unable to add SAD entry with SPI c2dc9aa0
received netlink error: Invalid argument (22)
Unable to add SAD entry with SPI ced1801e
Unable to install inbound and outbound IPsec SA (SAD) in kernel =====
But I still get the same error message when running IPv6 and NAT-T with strongswan 5.1.1 and kernel 3.12.8.

#8 - 30.04.2015 05:29 - nanjian5 Lee

the same result when running the following command in Linux kernel 3.12.8 and 3.13.11:

```
[root@localhost ~]# ip xfrm state add src fd14:828:ba69:2::2 dst fd14:828:ba69:1:21c:f0ff:fefa:f3c0 proto esp spi 0x00000301 mode tunnel auth md5
0x96358c90783bbfa3d7b196ceabe0536b enc des3_ede 0xf6ddb555acfd9d77b03ea3843f2653255afe8eb5573965df encap espinudp 4500 4500
fd13:828:ba69:2::2
RTNETLINK answers: Invalid argument
```

#9 - 30.04.2015 09:19 - nanjian5 Lee

FYI, I find the related mail thread:
<https://www.mail-archive.com/users%40lists.strongswan.org/msg05280.html>

#10 - 09.03.2016 16:21 - Tobias Brunner

- Related to Feature #892: Android client and IPv6 gateway added

#11 - 25.01.2017 15:08 - Tobias Brunner

- Related to Issue #2231: Cannot create IPv4-via-IPv6 tunnel added

#12 - 27.05.2017 17:22 - Noel Kuntze

- Category set to kernel
- Status changed from New to Feedback

Still no support in Kernel 4.9.29.

#13 - 27.05.2017 17:22 - Noel Kuntze

- Affected version deleted (5.3.0)

#14 - 06.06.2017 21:27 - Noel Kuntze

- Related to Issue #2354: Approval from Apple regarding an application using StrongSwan VPN added

#15 - 06.05.2019 23:08 - Noel Kuntze

- Subject changed from IPSEC nat-t for IPv6 to UDP Encapsulation for IPv6 traffic
- Description updated
- Assignee set to Noel Kuntze

Development in progress.

#16 - 27.03.2020 17:43 - Rajesh Kay

Noel Kuntze wrote:

Development in progress.

Hi Noel,

Is this issue resolved or is there a work-around for this issue?

#17 - 27.03.2020 23:14 - Noel Kuntze

Sabrina Dubroca was working on it last in the beginning of this year.

#18 - 28.03.2020 06:35 - Rajesh Kay

Noel Kuntze wrote:

Sabrina Dubroca was working on it last in the beginning of this year.

Thank you Noel. I am facing this issue on a setup running 3.10 kernel and Strongswan 5.8.2. However, I am able to add XFRM State and Policy using the "ip" utility. Want to understand if this is an issue with Netlink or if the data passed to the Netlink for the XFRM_MSG_NEWSA?

Can you or Sabrina provide some pointers on your analysis so far?

#19 - 14.05.2020 12:06 - Tobias Brunner

The upcoming 5.8 kernel will be the first to support UDP encapsulation for IPv6 as [the patch for it](#) by Sabrina Dubroca has recently been accepted to the ipsec-next/linux-next tree.

#20 - 16.05.2020 13:12 - Phil Zhang

It took me days to figure out UDP Encap is incompatible with IPv6 outer protocol. Currently running FreeBSD 12.1

It should have worked as is but I forced the encapsulation because of the firewall on my dumb cable modem

#21 - 05.06.2020 19:49 - Carl-Daniel Hailfinger

And the patch is now merged in mainline Linux:

<https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=0146dca70b877b73c5fd9c67912b8a0ca8a7bac7>