

freebsd-net: Re: kern/44355: After deletion of an IPv6 alias, the route to the whole subnet is removed too.

## Re: kern/44355: After deletion of an IPv6 alias, the route to the whole subnet is removed too.

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2004-08/0153.html>

---

*From:* B ( \$B?\_at\_L@C#:H (B)

*Date:* 08/29/04

Date: Sun, 29 Aug 2004 22:02:07 +0900  
To: Tilman Linneweh <arved@FreeBSD.org>

>>>> *On Fri, 27 Aug 2004 12:39:39 GMT,*  
>>>> *Tilman Linneweh <arved@FreeBSD.org> said:*

> *Synopsis: After deletion of an IPv6 alias, the route to the whole subnet is removed too.*  
> *Responsible-Changed-From-To: freebsd-bugs->freebsd-net*  
> *Responsible-Changed-By: arved*  
> *Responsible-Changed-When: Fri Aug 27 12:38:53 GMT 2004*  
> *Responsible-Changed-Why:*  
> *Old patches against IPv6, over to freebsd-net to decide if this PR is still*  
> *relevant*

> <http://www.freebsd.org/cgi/query-pr.cgi?pr=44355>

Hmm, this seems to be the same issue as that reported to the KAME project almost two years ago. The problem was then fixed there, but the fix does not seem to be merged to the FreeBSD repository.

The attached diff below is a similar fix for 5.2.1R. It cannot be applied to 4.x directly, but I guess it's easy to modify.

It would be nice if some committers could review the diff and (if appropriate) merge to the FreeBSD repository.

Thanks,

JINMEI, Tatuya  
Communication Platform Lab.  
Corporate R&D Center, Toshiba Corp.  
jinmei@isl.rdc.toshiba.co.jp

Index: in6.c

=====  
RCS file: /home/ncvs/src/sys/netinet6/in6.c,v  
retrieving revision 1.40  
diff -u -r1.40 in6.c

Re: kern/44355: After deletion of an IPv6 alias, the route to the whole subnet is removed too.

freebsd-net: Re: kern/44355: After deletion of an IPv6 alias, the route to the whole subnet is removed too.

```
--- in6.c 8 Nov 2003 23:36:32 -0000 1.40
+++ in6.c 29 Aug 2004 12:45:15 -0000
@@ -1,5 +1,5 @@
/* $FreeBSD: src/sys/netinet6/in6.c,v 1.40 2003/11/08 23:36:32 sam Exp $ */
-/* $KAME: in6.c,v 1.259 2002/01/21 11:37:50 keiichi Exp $ */
+/* $KAME: in6.c,v 1.334 2002/12/05 15:33:26 jinmei Exp $ */

/*
 * Copyright (C) 1995, 1996, 1997, and 1998 WIDE Project.
@@ -579,6 +579,14 @@
    */
    if ((error = in6_update_ifa(ifp, ifra, ia)) != 0)
        return (error);
+ if ((ia = in6ifa_ifpwithaddr(ifp, &ifra->ifra_addr.sin6_addr))
+ == NULL) {
+ /*
+ * this can happen when the user specify the 0 valid
+ * lifetime.
+ */
+ break;
+ }

    /*
    * then, make the prefix on-link on the interface.
@@ -617,6 +625,15 @@
    ((ifra->ifra_flags & IN6_IFF_AUTOCONF) != 0);
    pr0.ndpr_vltime = ifra->ifra_lifetime.ia6t_vltime;
    pr0.ndpr_pltime = ifra->ifra_lifetime.ia6t_pltime;
+ if ((error = in6_init_prefix_ltimes(&pr0)) != 0) {
+ /*
+ * Validation for lifetimes should have been done, so
+ * this should always succeed.
+ */
+ log(LOG_ERR, "in6_control: failed to initialize prefix"
+ " lifetimes\n");
+ return (error);
+ }

    /* add the prefix if not yet. */
    if ((pr = nd6_prefix_lookup(&pr0)) == NULL) {
@@ -632,48 +649,40 @@
        return (EINVAL); /* XXX panic here? */
    }
}
- if ((ia = in6ifa_ifpwithaddr(ifp, &ifra->ifra_addr.sin6_addr))
- == NULL) {
- /* XXX: this should not happen! */
- log(LOG_ERR, "in6_control: addition succeeded, but"
- " no ifaddr\n");
- } else {
- if ((ia->ia6_flags & IN6_IFF_AUTOCONF) != 0 &&
```

Re: kern/44355: After deletion of an IPv6 alias, the route to the whole subnet is removed too.

freebsd-net: Re: kern/44355: After deletion of an IPv6 alias, the route to the whole subnet is removed too.

```
- ia->ia6_ndpr == NULL) { /* new autoconfed addr */
- ia->ia6_ndpr = pr;
- pr->ndpr_refcnt++;
-
- /*
- * If this is the first autoconf address from
- * the prefix, create a temporary address
- * as well (when specified).
- */
- if (ip6_use_tempaddr &&
- pr->ndpr_refcnt == 1) {
- int e;
- if ((e = in6_tmpifadd(ia, 1)) != 0) {
- log(LOG_NOTICE, "in6_control: "
- "failed to create a "
- "temporary address, "
- "errno=%d\n", e);
- }
- }
- }
+
+ /* relate the address to the prefix */
+ if (ia->ia6_ndpr == NULL) {
+ ia->ia6_ndpr = pr;
+ pr->ndpr_refcnt++;
+
+ /*
+ * this might affect the status of autoconfigured
+ * addresses, that is, this address might make
+ * other addresses detached.
+ * If this is the first autoconf address from the
+ * prefix, create a temporary address as well
+ * (when required).
+ */
- pfxlist_onlink_check();
+ if ((ia->ia6_flags & IN6_IFF_AUTOCONF) &&
+ ip6_use_tempaddr && pr->ndpr_refcnt == 1) {
+ int e;
+ if ((e = in6_tmpifadd(ia, 1)) != 0) {
+ log(LOG_NOTICE, "in6_control: failed "
+ "to create a temporary address, "
+ "errno=%d\n", e);
+ }
+ }
+ }
+
+ /*
+ * this might affect the status of autoconfigured addresses,
+ * that is, this address might make other addresses detached.
+ */
+ pfxlist_onlink_check();
```

Re: kern/44355: After deletion of an IPv6 alias, the route to the whole subnet is removed too.

freebsd-net: Re: kern/44355: After deletion of an IPv6 alias, the route to the whole subnet is removed too.

```
+
    break;
}

case SIOCDIFADDR_IN6:
{
- int i = 0;
- struct nd_prefix pr0, *pr;
+ struct nd_prefix *pr;

    /*
     * If the address being deleted is the only one that owns
@@ -683,37 +692,12 @@
     * and the prefix management. We do this, however, to provide
     * as much backward compatibility as possible in terms of
     * the ioctl operation.
+ * Note that in6_purgeaddr() will decrement ndpr_refcnt.
    */
- bzero(&pr0, sizeof(pr0));
- pr0.ndpr_ifp = ifp;
- pr0.ndpr_plen = in6_mask2len(&ia->ia_prefixmask.sin6_addr,
- NULL);
- if (pr0.ndpr_plen == 128)
- goto purgeaddr;
- pr0.ndpr_prefix = ia->ia_addr;
- pr0.ndpr_mask = ia->ia_prefixmask.sin6_addr;
- for (i = 0; i < 4; i++) {
- pr0.ndpr_prefix.sin6_addr.s6_addr32[i] &=
- ia->ia_prefixmask.sin6_addr.s6_addr32[i];
- }
- /*
- * The logic of the following condition is a bit complicated.
- * We expire the prefix when
- * 1. the address obeys autoconfiguration and it is the
- * only owner of the associated prefix, or
- * 2. the address does not obey autoconf and there is no
- * other owner of the prefix.
- */
- if ((pr = nd6_prefix_lookup(&pr0)) != NULL &&
- (((ia->ia6_flags & IN6_IFF_AUTOCONF) != 0 &&
- pr->ndpr_refcnt == 1) ||
- ((ia->ia6_flags & IN6_IFF_AUTOCONF) == 0 &&
- pr->ndpr_refcnt == 0))) {
- pr->ndpr_expire = 1; /* XXX: just for expiration */
- }
-
- purgeaddr:
+ pr = ia->ia6_ndpr;
    in6_purgeaddr(&ia->ia_ifa);
+ if (pr && pr->ndpr_refcnt == 0)
+ prelist_remove(pr);
```

Re: kern/44355: After deletion of an IPv6 alias, the route to the whole subnet is removed too.

freebsd-net: Re: kern/44355: After deletion of an IPv6 alias, the route to the whole subnet is removed too.

```
        break;
    }

@@ -1177,24 +1161,26 @@
        in6_prefix_remove_ifid(iilen, oia);
    }

- /*
- * When an autoconfigured address is being removed, release the
- * reference to the base prefix. Also, since the release might
- * affect the status of other (detached) addresses, call
- * pfxlist_onlink_check().
+ /*
+ * Release the reference to the base prefix. There should be a
+ * positive reference.
    */
- if ((oia->ia6_flags & IN6_IFF_AUTOCONF) != 0) {
- if (oia->ia6_ndpr == NULL) {
- nd6log((LOG_NOTICE, "in6_unlink_ifa: autoconf'ed address "
- "%p has no prefix\n", oia));
- } else {
- oia->ia6_ndpr->ndpr_refcnt--;
- oia->ia6_flags &= ~IN6_IFF_AUTOCONF;
- oia->ia6_ndpr = NULL;
- }
+ if (oia->ia6_ndpr == NULL) {
+ nd6log((LOG_NOTICE,
+ "in6_unlink_ifa: autoconf'ed address "
+ "%p has no prefix\n", oia));
+ } else {
+ oia->ia6_ndpr->ndpr_refcnt--;
+ oia->ia6_ndpr = NULL;
+ }

+ /*
+ * Also, if the address being removed is autoconf'ed, call
+ * pfxlist_onlink_check() since the release might affect the status of
+ * other (detached) addresses.
+ */
+ if ((oia->ia6_flags & IN6_IFF_AUTOCONF))
+     pfxlist_onlink_check();
- }

    /*
    * release another refcnt for the link from in6_ifaddr.
```

---

freebsd-net@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-net>

To unsubscribe, send any mail to "freebsd-net-unsubscribe@freebsd.org"

Re: kern/44355: After deletion of an IPv6 alias, the route to the whole subnet is removed too.