strongSwan - Feature #1230

connmark module, CONNMARK restore is too broad

08.12.2015 21:16 - Saso Slavicic

Status:	Closed	**Start date:**	08.12.2015
Priority:	Normal	**Due date:**	
Assignee:	Tobias Brunner	**Estimated time:**	0.00 hour
Category:	libcharon		
Target version:	5.4.0		
Resolution:	Fixed		

Description

When connmark module is active, CONNMARK restore is inserted into mangle table OUTPUT chain.

This rule is too vague, as it will restore all marks between two hosts. When two clients from the same NAT IP connect, two identical rules will be inserted (can also be seen on connmark doc page: https://wiki.strongswan.org/projects/strongswan/wiki/Connmark).

When Windows hosts connect, L2TP traffic always uses 1701 as both src and dst ports. Unfortunately conntrack can no longer properly track these udp streams, as both will have identical src/dst tuple. Only traffic to the client that last sent inbound packet is active as all outbound packets will have the same mark. The other client simply timeouts...

Now, I don't think this can be solved with iptables magic alone, so l2tp daemon has to have some control where to send packets (eg. set marks on the outgoing packets).
Unfortunately CONNMARK restore will overwrite any mark that has already been set on the packet.

I propose to add at least mark=0 match to this rule. Adding ipsec policy dir out would be nice to further limit what traffic strongswan modifies, but for some reason the first L2TP response packet does not have this set so the connection cannot be established if this match is also set.

A sample patch is attached (against 1212-ipt-alignment branch).
I have patched my xl2tpd to set correct tunnel marks on the socket (using SO_MARK) and this finally allows 2 Windows hosts behind same NAT to use L2TP.

Associated revisions

Revision c659d369 - 10.03.2016 17:26 - Tobias Brunner

connmark: Don't restore CONNMARK for packets that already have a mark set

This allows e.g. modified versions of xl2tpd to set the mark in
situations where two clients are using the same source port behind the
same NAT, which CONNMARK can't restore properly as only one conntrack entry
will exist with the mark set to that of the client that sent the last packet.

Fixes #1230.

History

#1 - 09.03.2016 12:22 - Tobias Brunner

- Tracker changed from Issue to Feature

- Category set to libcharon

- Status changed from New to Feedback

- Assignee set to Tobias Brunner

- Target version set to 5.4.0

- Resolution set to Fixed

> I propose to add at least mark=0 match to this rule.

Yes, makes sense. I applied a slightly modified version of the patch in the *1229-1230-connmark-fixes* branch.

I have patched my xl2tpd to set correct tunnel marks on the socket (using SO_MARK) and this finally allows 2 Windows hosts behind same NAT to use L2TP.

Perhaps you could add some information on this to the connmark page? Maybe even a patch? Thanks!

#2 - 10.03.2016 17:26 - Tobias Brunner

- Status changed from Feedback to Closed

Thanks a lot for the updates to connmark!

Files

connmark-mark_match.patch	1.29 KB	08.12.2015	Saso Slavicic