
Fraunhofer-Institut für Kommunikation, Informationsverarbeitung und Ergonomie FKIE

Lets talk about the IP multicast data plane

July 20th 2021

Henning Rogge



Lets talk about the IP multicast data plane

How to implement IP multicast routing in Linux?

What has been done?

- Multicast forwarding in user space

What has been done?

- Multicast forwarding in user space
 - Raw sockets

What has been done?

- Multicast forwarding in user space
 - Raw sockets

Pro: it works...

Con: it doesn't play nice

(Policy Routing, Firewalls)

What has been done?

- Multicast forwarding in user space
 - Raw sockets
 - Tun/Tap interface

What has been done?

- Multicast forwarding in user space
 - Raw sockets
 - Tun/Tap interface

Pro: easier to integrate

Con: getting traffic flowing is complicated

(IP multicast routes? Bridges?)

What has been done?

- Multicast forwarding in kernel space

What has been done?

- Multicast forwarding in kernel space
 - IP multicast routing

Pro: Fast, integrated, plays nice!

Con: Doesn't work for MANET

What has been done?

- Multicast forwarding in kernel space
 - IP multicast routing

Pro: Fast, integrated, plays nice!

Con: Doesn't work for MANET

Or does it?

(what new tools do we have?)

What's the trouble?

- Retransmission on same interface
- Selectively forwarding multicast (MPR?)
- Local traffic
- Suppressing duplicates

What's the trouble?

- Retransmission on same interface

Seems to be working on modern Linux kernels!

- Selectively forwarding multicast (MPR?)
- Local traffic
- Suppressing duplicates

What's the trouble?

- Retransmission on same interface
- Selectively forwarding multicast (MPR?)

Maybe some “dynamic” firewalling
(needs Ethernet source)?

- Local traffic
- Suppressing duplicates

What's the trouble?

- Retransmission on same interface
- Selectively forwarding multicast (MPR?)
- Local traffic

Could be made easier with virtual Ethernet devices

- Suppressing duplicates

What's the trouble?

- Retransmission on same interface
- Selectively forwarding multicast (MPR?)
- Local traffic
- Suppressing duplicates

We need custom code in the kernel... ***sigh***


What's the trouble?

- Retransmission on same interface
- Selectively forwarding multicast (MPR?)
- Local traffic
- Suppressing duplicates

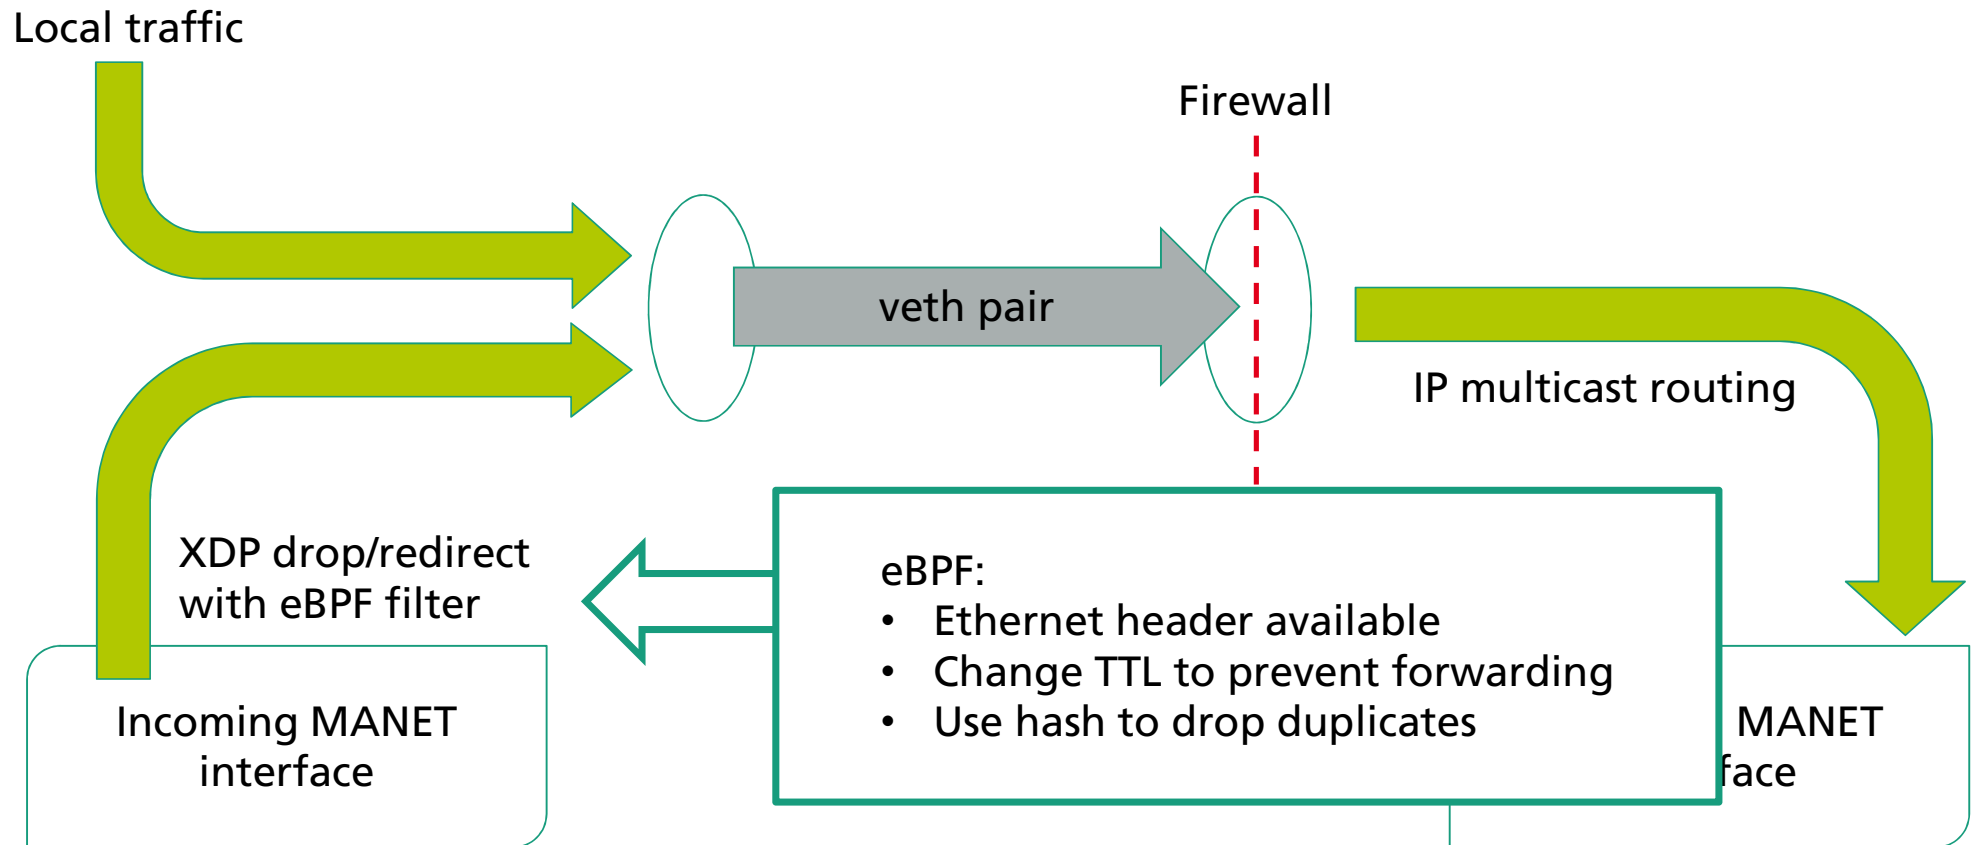
We need custom code in the kernel...

Maybe just use XDP and eBPF?

Idea for multicast data plane



Idea for multicast data plane



Questions?