

Hardware offload BOF

Shrijeet Mukherjee, Neil Horman

<https://etherpad.mozilla.org/2PlezMRjCF>

A fantasy agenda

Capabilities

a. Explicit list

Or Query serially, punt to higher level (explicit hierarchy)

Or Model each device uniquely with capability (no hierarchy)

b. Need to understand this for Switch Asic versus VEPA/EVB/SRIOV nic etc

Flow offload :

c. Manage as discrete devices or generic pipeline

How is interop measured aka how to avoid anarchy :)

d. Flow API scheme [John Fastabend]

d'. Model using P4 [Mihai]

f. TC scheme [Jiri Pirko]

f'. EZChip [Gilaad]

A fantasy agenda

Routing tables, FDB, MDB, ACL

g. Capacity indication, aka properties of tables [Roopa]

h. Fine grain capability (e.g. is it sufficient to ask if multicast is supported)

j. Table characteristics LPM versus Logical Hash based LPM's and practical implications

Device model : (Not mutually exclusive in anyway)

k. Maintaining operational consistency is KEY

Make switch look like NIC or vice versa e.g. is learning a basic capability ?

l. Model using OVS (inherently host based)

m. Model using rocker [Scott]

n. Switch Abstraction Interface [Sanjay]

n'. Intel [Uri]

n". Qualcomm [Olivari]

A fantasy agenda

Features :

l3 offloads [Hannes]

acl offloads [Pablo]

o. Load Balancing

p. Bonding ++ (MLAG and friends)

q. Stateful packet processing

Etherpad output

- **Hardware offload BoF**
- Netdev01 - Sun, Feb 15, 2015 1:00pm
-
- Major points:
- Preserve the Linux Networking Model
- Goal is to exchange ideas
-
- Capability Determination
- Patrick: Can the default route be removed? So that individual routes can be removed. Use route usage statistics to remove least used routes.
- David: Existing tools must continue to work.
- Signal an error if out of space, or
- Provide capacity initially and restrict usage to that limit
- Cares:
- Must be able to support the decisions we make
- Partick: A flag in the route add ??? [ed. help]
- Ben: Need a fairly good switch model to emulate the hardware devices.
- ????: The least recently used method proposed by Patrick may lead to periods of major route updates when network changes occur. With multiple routing suites each one needs to be given the capabilities (capacities) of the underlying switch
- Shrijeet: General agreement: minimal policy in the kernel, most in userspace
-
- Switch ASIC vs. SRIOV (NIC) Switching models
- Shrijeet: Do these two need to be different?
- Andy: Come to my talk tomorrow.
- Thomas: Bring all the DSA drivers onboard
- Gilaad: Need to define common features, plus need a way to include additional (unique) capabilities