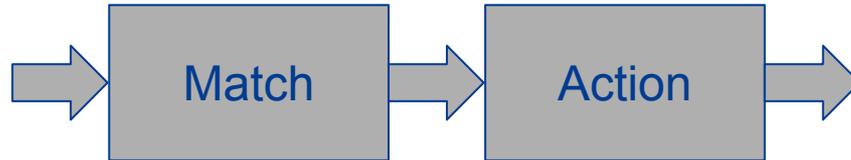
A blurred person in a dark suit is walking from left to right in the foreground. In the background, there is a modern building with a glass facade and a staircase. The scene is brightly lit, suggesting an outdoor or well-lit indoor environment.

NETRONOME

TC Flower Offload

Simon Horman
Netdev 2.2, Seoul

- Programmable datapaths are new and exciting
- Match-action tables are more static but widely used
- Packets traverse one or more table of match-action rules
 - Classifier seeds flow key
 - Flow key is optionally masked
 - Flow key is used to lookup rules of table
 - If a match is found the actions of the rule are executed
 - Otherwise some fall-through occurs



- TC Flower classifier allows matching packets against pre-defined flow key fields:
 - Packet headers: f.e. IPv6 source address
 - Tunnel metadata: f.e. Tunnel Key ID
 - Metadata: Input port
- TC actions allow packet to be modified, forwarded, dropped, etc...
 - pedit: modify packet data
 - mirred: output packet
 - vlan: push, pop or modify VLAN
 - ...

- TC Flower makes use of the flow disector to perform classification
- Resulting flow key is masked
 - Only one mask allowed per priority
- Used to lookup flows
- If a match is found then TC actions of rule are executed
- If not packet processing proceeds to next classifier
 - Standard TC behaviour

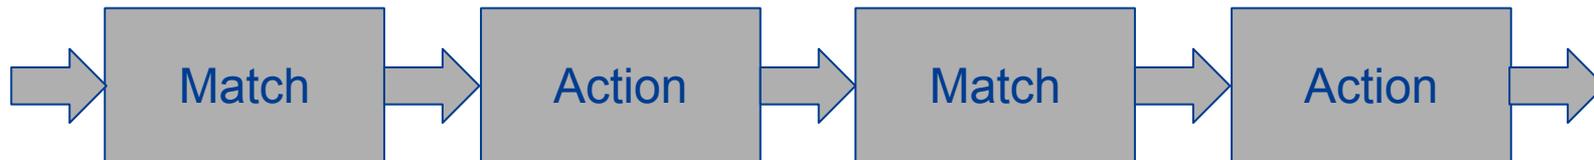
- Provide greater throughput
- Increase CPU core efficiency and scalability

- mqprio offload introduced in v2.6.39
- Introduced `ndo_setup_tc`
- u32 Classifier offload introduced in v4.6
- As of v4.15-rc4 BPF, flower and matchall classifier offload also supported
- Netronome NFP driver implements BPF and flower classifier offload

- skip_hw and skip_sw flags
 - Allow users to influence placement of flows by kernel
 - Default is to add to hardware and try to add to software
- in_hw and not_in_hw flags
 - Allow kernel to report presence of flow in hardware

- IPv6 label and neighbour discovery
- Maskable match of MPLS LSE fields
- GENEVE options

- Aim would be to allow enhanced rules to be written
 - By taking into account Conntrack state
- Scheme implemented by Open vSwitch datapath is:
 - Conntrack action passes packet to through conntrack
 - Packet is then classified for a second time; conntrack state may form part of flow key



A blurred person in a dark suit is walking past a modern glass building. The building features a prominent staircase with dark steps and a glass railing. The scene is brightly lit, suggesting an outdoor or well-lit indoor environment. The overall aesthetic is professional and modern.

NETRONOME

Thank You