

CONTAINER NETWORKING SOLUTIONS

Anjali Singhai Jain, Nupur Jain, Intel

Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Agenda

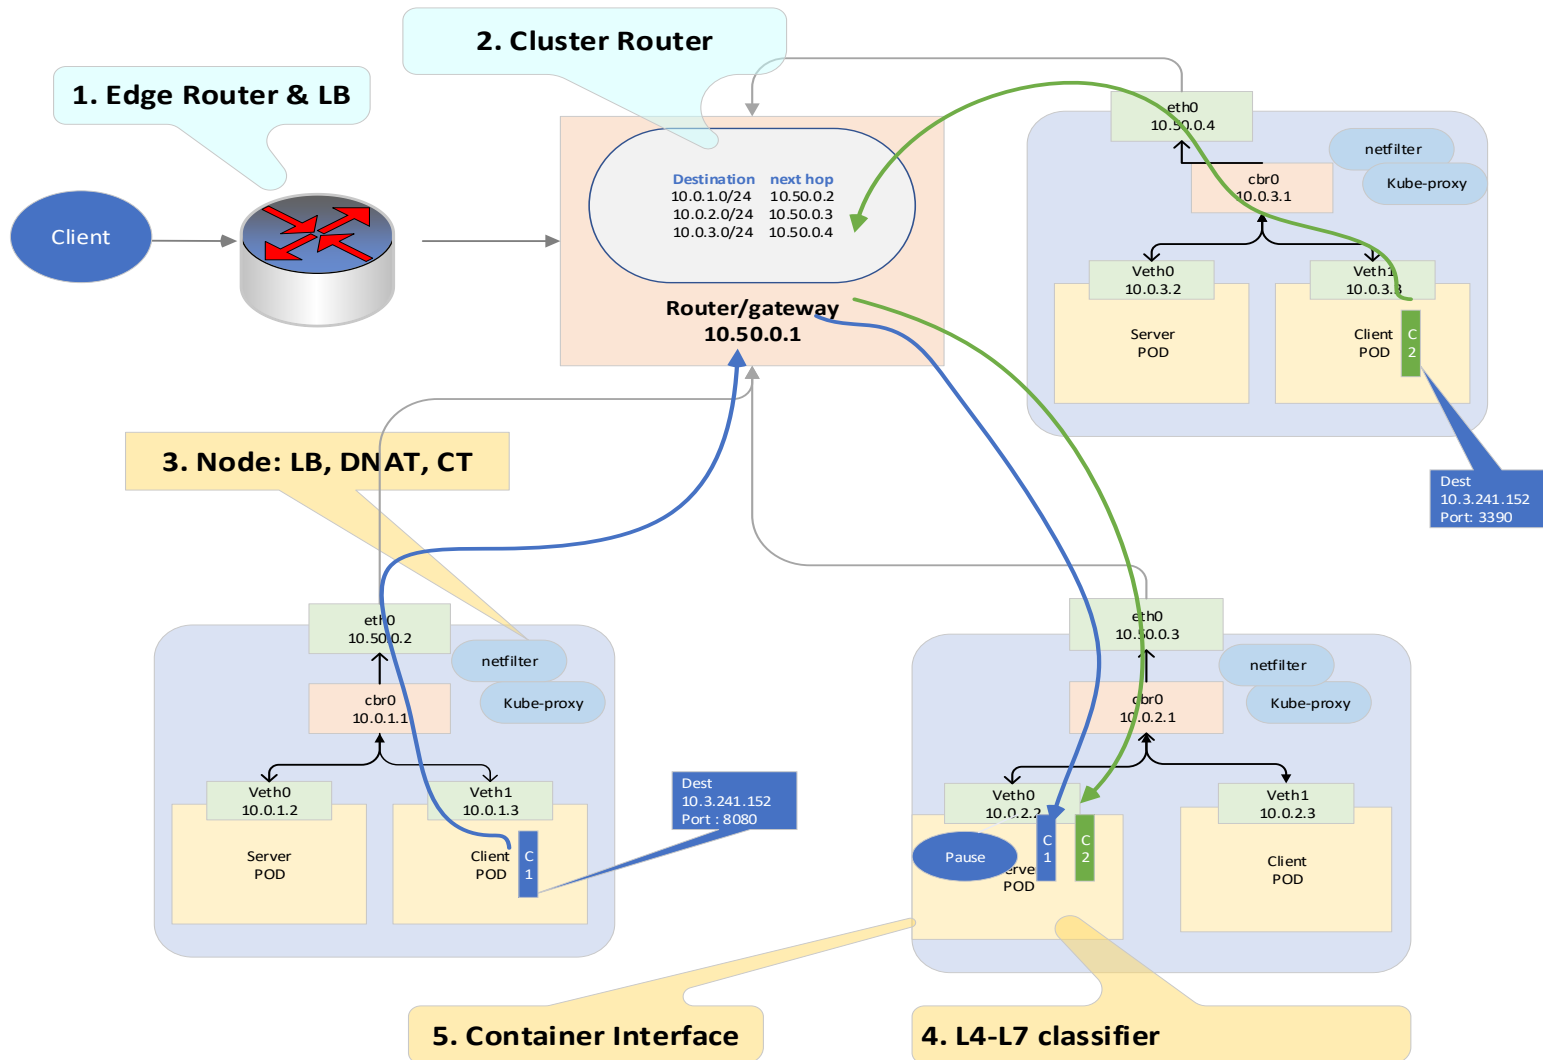
Introduction

- Container Network Functions & Interfaces
- Limitations
- Container Interface Classifier
- Community Solutions
- Our Proposed Solution
- Consolidation

Functionality Offload

- P4 Sample -Kubeproxy
- P4 Connection tracking
- P4 L4-L7 Classifier

Container Networking



Current State of Deployment



Need for Scale
CSP Container
Scale Needed
~10K





Device performance
Throughput,
Programmability,
Resource Scaling
has increased
significantly



Effective utilization
of the system
resources (cores,
memory, network)
to improve
performance and
packing of
containers



Ideal
Container
Network
Solution



5. Existing interfaces for containers

Shared (Pod Interfaces)

- Stacked netdevs on a PCIe PF netdev
- Assigned to container namespaces
- Examples : MACVLAN, IPVLAN, bridge

Dedicated (Pod Interfaces)

- SR-IOV VFs
- Too heavy
 - Separate PCIe config space
 - HW based packet replication for broadcast, multicast – higher PCI BW utilization

Why Hardware Accelerate ?

End to end Maximize throughput

- Avoid the SW long path which limits how much a Server Pod can handle.

End to end Native Scale out

- By Reducing Latency and Jitter introduced by kernel to user context switch in present AF_XDP memory model.
- Dedicated resources takes away the need for OS to schedule on a shared resource. OS overhead for managing resources is gone.
 - CPU scheduling, memory management etc.

Security & Isolation

- Queue level isolation.

Assignable Container interfaces using X-IOV (S-IOV) & User Interrupts

Hardware Assisted Virtualization

Highly scalable and high- performance sharing of I/O devices across isolated domains

Assignable device Interfaces =>User Container Interfaces

Platform Scalability using PASID

Support Virtual Device Composition

*RID and *PASID identifies the address space associated with the request

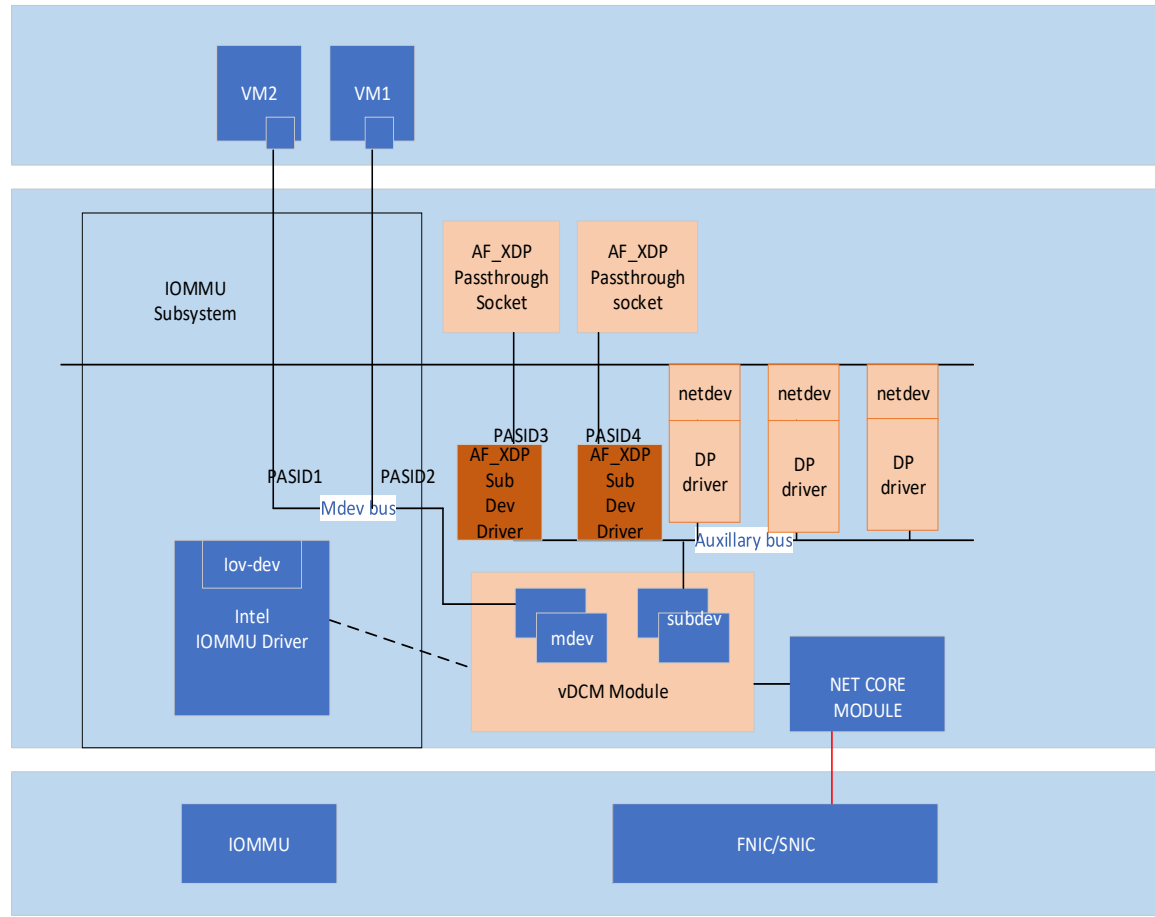
*ADI Memory Mapped regions

*RID = Requester ID

*PASID = Process Address Space Identifier


*ADI Assignable Device interfaces

SIOV Spec : Reference Number: 337679-001, Revision: 1.0



4. Container Interface Classifier: Solution

- L4-L7 Classifier and forwarder in the HW
- This extends the HW Offloaded vSwitch Classification End Point to the Container Interface.



Option1: HW Offload the classifier & forwarder

- AF_XDP raw_socket bound to a HW vPort/QP through Side band filters.
- Provide inline filters to be added in HW as part of TX packet.
 - ATR style in ADQ

Option2: HW Offload the Classifier

- Provide a meta data classification hint to kernel/user with a packet.
 - Flow mark or a 32bit hash value based on L4-L7 fields.

3. Node: Load Balancer, DNAT, CT

Existing Solutions

Kube-proxy

- Kernel Netfilters - Not performant
 - Iptables $O(n)$ chains proportional to size of cluster, in-place rule modifications not possible.
 - IPVS $O(1)$ - hash ipset but do not work well with other services requiring iptables for filtering
- Kernel with eBF/XDP - Accelerated

Connection Tracking

- Robust to syn floods but limited by max size
- No flexibility, fixed hash algorithm and field selection for hash

Kernel Overall not very flexible, latencies due to irq processing, context switching, slow API configuration interfaces

Community's Approach - eBPF/XDP

Benefits – Performant than the kernel

Designed as an alternative to DPDK.

Flexibility, code injected into the kernel

Ability to reload programs on-the-fly

Network Functions – Network Policy, Encryption, Load Balancer, Firewall, Monitoring etc

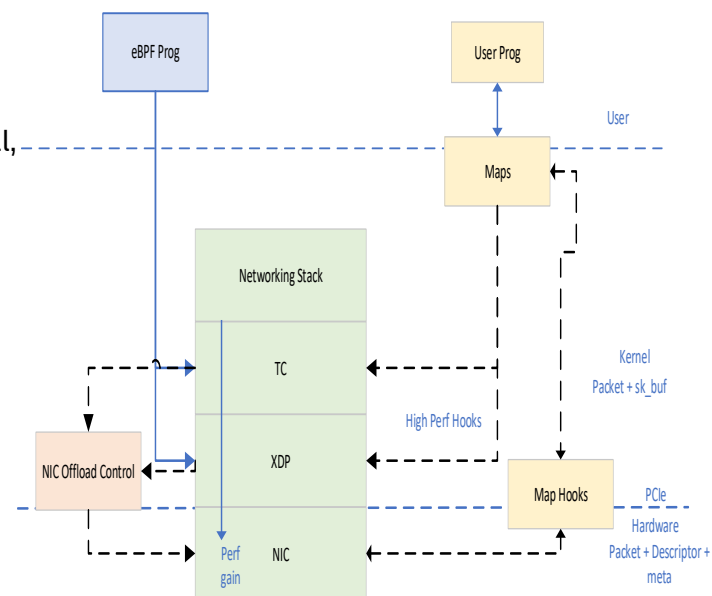
Plugins – Cilium, Calico, Katran (facebook), etc

Limitations –

- Not HW offloadable to ASICs.
 - Note: Netronome uses NPU general Purpose cores
- More cores required to scale connections
- Kernel/user space context switching
- General purpose CPUs and Memory architecture is not ideal for Deep table Lookup

Solution –

- Purpose built ASICs, purpose -built cores and dedicated Context aware caches may be the way to go



eBPF Implementation – Cilium

- Cilium does connect-time load balancing by hooking into the kernel XDP/TC hook on the receive.
- When a program tries to connect to a Kubernetes service, Cilium intercepts the connection attempt, load balances with DNAT's to directly connect to the backend pod's IP instead.

Throughput in queries/s – test run with fortio and nginx. 2 clients, 2 NodePort IPs, 2 backend Pods per NodePort

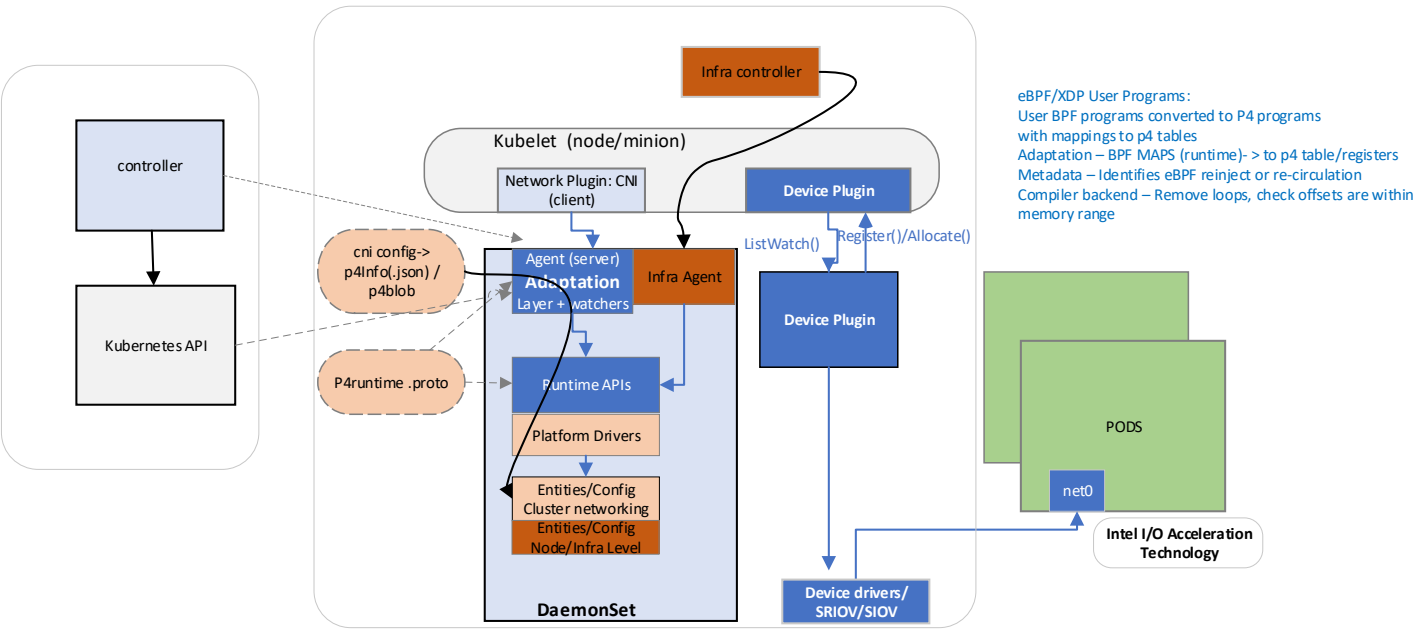
No XDP, 2 CLV interfaces			
Fortio + Nginx	Client 1	Client 2	Total
100B	47578.6	50411.5	97990.1
1500B	46174.4	49990.4	96164.8

XDP, 1 CLV interface			
Fortio + Nginx	Client 1	Client 2	Total
100B	64038.7	53970	118008.7
1500B	49623.5	48369.7	97993.2

No-XDP and XDP performed similar with 100,000+ connections. CPU consumption in both reaches to >8 cores with more sessions

We suspect, XDP path based on this experiment, does not take advantage of hardware XDP_REDIRECT queue designed to send packets to another interface, hence no significant performance gains

Our Approach – The Whole Datapath



*Programmable MAT tables.
Contract between control plane and data plane for runtime control

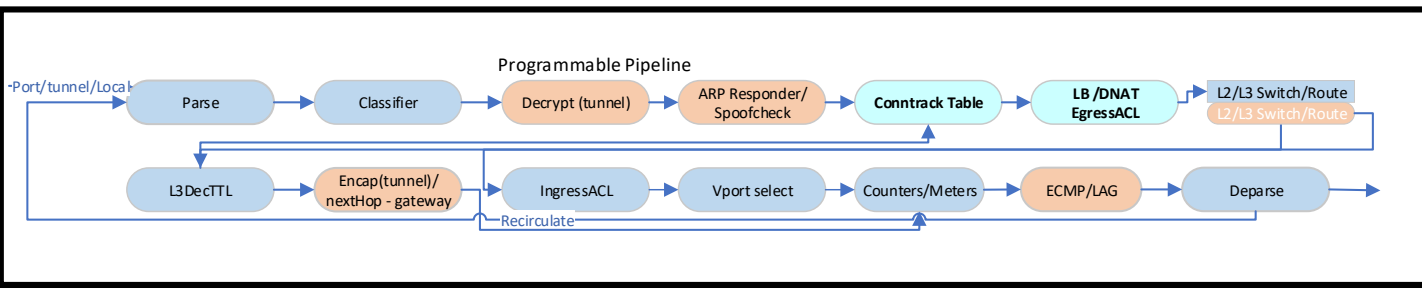
*Device capability defined by architecture. Eg. wildcard match -TCAM, Exact match-SRAM. Compiler responsible for mapping.

*Parallel lookups, conditional actions & atomicity

*Features are defined in the software. Faster introduction, verification, test and deployment.

*Programmer defined 1) parse graph; headers and orders. 2) Packet modifications

*Counters, meters, stateful registers, hash functions, ALU, TTLs, PRE



Node LB Data Plane P4

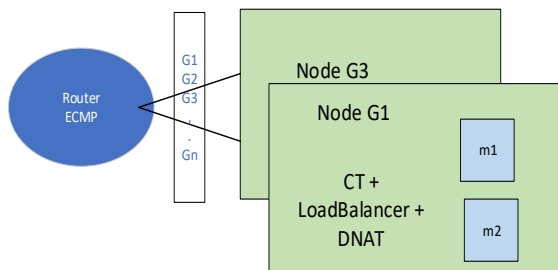


Table Entry	Key (ClusterIP, port)	Member group
t1	10.3.241.152	g1
t2	10.3.241.159	g2
t3	10.3.241.170	g3

Group Ref	Members
g1	m1, m2, m3
g2	m1, m2
g3	m1, m2

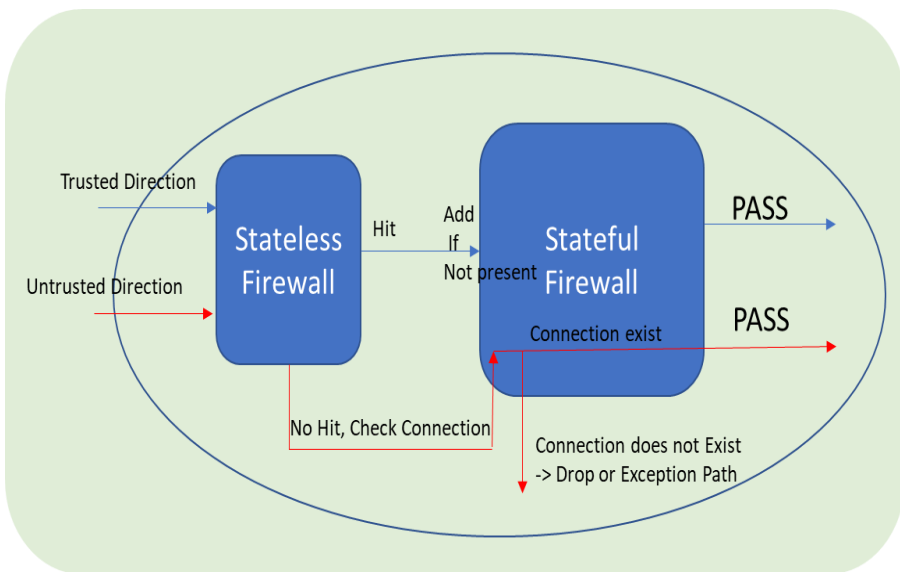
Member ref	Action Spec
m1	10.0.1.2, 8080, 02:42:88:7e:17:61
m2	10.0.2.2, 8080, f6:27:14:ce:4a:7d

```

1 // Proxy LB on cluster service IP to an endpoint on a POD
2 extern ActionSelector{
3     /// Construct a selection table for a given ActionProfile.
4     ActionSelector(ActionProfile action_profile,
5                     Hash<_> hash,
6                     SelectorMode_t mode,
7                     bit<32> max_group_size,
8                     bit<32> num_groups);
9
10    ActionSelector(bit<32> size, Hash<_> hash, SelectorMode_t mode);
11 }
12
13
14 control simple_lb(inout headers hdr,
15                  inout metadata meta,
16                  switch_uint32_t lb_table_size,
17                  inout standard_metadata_t standard_metadata) {
18
19     //Chose an extern hash or PNA hash
20     Hash<switch_uint32_t>(HashAlgorithm_t.CRC32) selector_hash;
21
22     ActionSelector(
23         1024, selector_hash, SelectorMode_t.FAIR) pod_selector;
24
25     // Pick an entry and apply DNAT
26     action set_nhopp(bit<48> pod_dmac, bit<32> pod_ipv4, bit<9> port) {
27         hdr.ethernet.dstAddr = pod_dmac;
28         hdr.ipv4.dstAddr = pod_ipv4;
29         hdr.tcp.dstPort = port;
30         standard_metadata.egress_spec = port;
31     }
32
33     // hash to use 5-tuple
34     // can be tcp, udp, sctp
35     table lb {
36         key = {
37             hdr.ipv4.dstAddr : exact;
38             hdr.tcp.dstPort : exact;
39             hdr.ipv4.srcAddr : selector;
40             hdr.ipv4.dstAddr : selector;
41             hdr.ipv4.tcp.srcPort : selector;
42             hdr.ipv4.tcp.dstPort : selector;
43             hdr.ipv4.protocol : selector
44         }
45
46         actions = {
47             NoAction;

```

CT P4 Data Plane

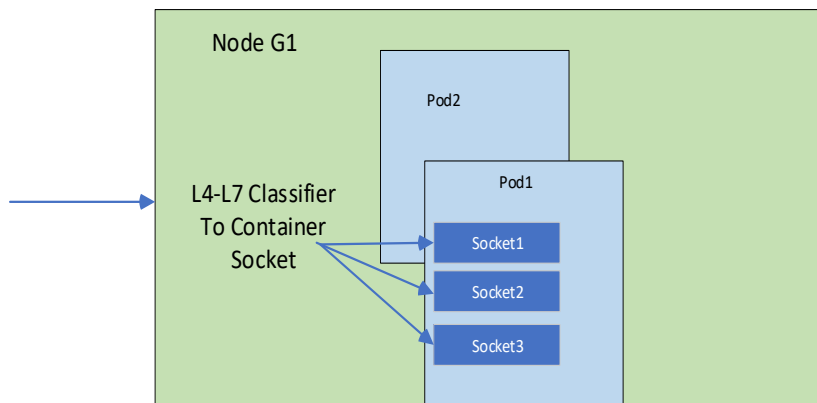


```

262 action ct_aging_table_hit (
263     @per_entry_state in    FlowId_t flow_id,
264     @per_entry_state inout ExpireTimeSelection_t expire_time)
265 {
266     // my_flow_id = flow_id;
267     if (modify_expire_soon_on_hit) {
268         expire_time = new_expire_time_selection; // timers per state
269         modify_entry_to_expire_soon();
270     } else if (update_expire_time) {
271         expire_time = new_expire_time_selection;
272         restart_expire_timer(); // sweep_count = 0
273     } else {
274         restart_expire_timer(); // sweep_count = 0
275     }
276 }
277
278 action ct_aging_table_miss() {
279     FlowId_t my_flow_id;
280     bool add_succeeded;
281     if (add_on_miss) {
282         my_flow_id = allocate_flow_id();
283         add_succeeded =
284             add_entry("ct_aging_table_hit", // name of action
285                     (ct_aging_table_hit_params_t)
286                     {flow_id = my_flow_id,
287                     expire_time = new_expire_time_selection});
288         // add_entry() initializes the new entry as if
289         // restart_expire_timer() had been called on it,
290     }
291 }
292
293 // CT 5 Tuple table
294 table ct_aging_table {
295     key = {
296         // P4 developer gets to select the key fields they want,
297         // e.g. This is a 5 tuple plus a field like zone id , used to guarantee
298         // that entries in different IP private address domains are
299         // unique.
300     }
301     actions = {
302         // @tableonly and @defaultonly are standard annotations
303         // defined in the P4_16 language specification.
304         @tableonly    ct_aging_table_hit;
305         @defaultonly ct_aging_table_miss;
306     }
307     // This is a straw man syntax. Probably want more than 'true' so
308     // that this code might contain other compile-time options.
309     add_on_miss = true;
310     idle_timeout = true;
311     const default_action = ct_aging_table_miss;
312 }
313
314 apply

```

L4-L7 P4 Classifier



```

1  control L4_L7_classifier(inout headers hdr,
2     inout metadata meta,
3     switch_uint32_t lb_table_size,
4     inout standard_metadata_t standard_metadata) {
5
6     // Forward to a Container Interface
7     action set_queue(bit<16> queue_id) {
8         standard_metadata.egress_spec = queue_id;
9     }
10
11    // Set a flow ID meta data
12    action set_flow_id(bit<32> mark) {
13        standard_metadata.flow_id = mark;
14    }
15
16    // exact match using n-tuple
17    // can be tcp, udp, sctp
18    table lb {
19        key = {
20            hdr.ipv4.dstAddr : exact;
21            hdr.udp.dstPort : exact;
22            hdr.udp.srcPort : exact;
23            hdr.ipv4.srcAddr : exact;
24            hdr.udp.quic.cid : exact
25        }
26
27        actions = {
28            NoAction;
29            set_queue;
30            flow_count.count();
31            set_flow_id;
32        }
33
34        const default_action = Drop;
35        counters = flow_counter;
36        size = lb_table_size; // can be configured
37        implementation = pod_selector;
38    }
39 }
40
41

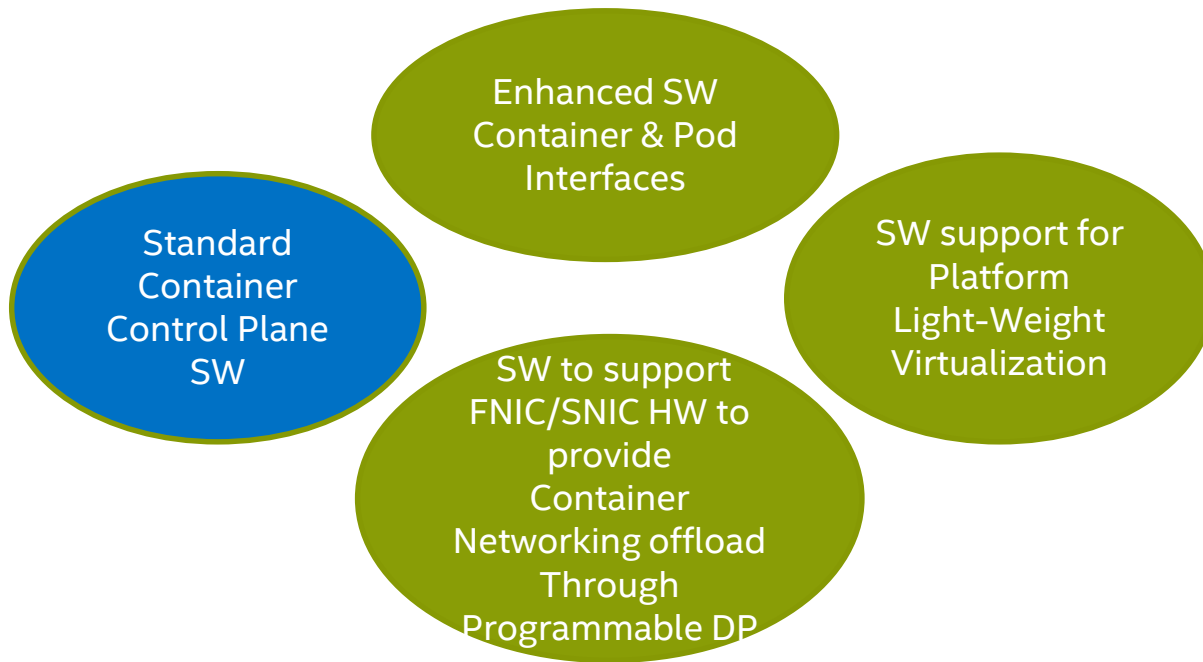
```


Opens

- Not every eBPF program can be HW offloaded as is. We are looking at all use cases.
- We would like to get community support in converting some well defined XDP implementations to p4 programs.
- P4 extensions or externs is an option for complete packet transformations like Crypto, checksum, packet replication etc.

Conclusion:

Components for Native Scale out of Container Networking



Components	
Standard (No Change)	
Need SW Enhancements to benefit from HW	

Contacts

P4 Code will be on github soon...

Please email for more info...

Contacts: Anjali Singhai Jain anjali.singhai@intel.com

Nupur Jain nupur.jain@intel.com

Intel Container Networking Team:

Anjali Singhai Jain, Nupur Jain, Amritha Nambiar, Pawel Szymanski, Shaopeng He, Phani Burra,	Dan Daly, Yadong Li, Sridhar Samudrala, Kiran Patil, Liang Cunming , Edwin Verplanke
---	---

eBPF Implementation Characterization – Cilium

Stack trace (NO XDP)

vmlinux	615.317s
nginx	55.704s
[Unknown]	24.101s
[Outside any known module]	24.101s
cls_bpf_classify	21.866s
↖ __tcf_classify ← tcf_classify_in	1.173s
↖ __tcf_classify ← tcf_classify ←	1.042s
[Unknown]	0.020s
↖ func@0x7a5bd4 ← func@0x79	0s
func@0x7a0570	0s
func@0x3527ee1	0s
ice.ko	23.259s
ip_tables.ko	20.728s
libc-2.13.so	11.371s
amplxe-perf	8.445s
cls_bpf.ko	8.339s
containerd-shim	8.209s

Stack trace (XDP)

[Unknown]	28.626s	52,704,500,000
[Outside any known module]	28.626s	52,704,500,000
cls_bpf_classify	17.400s	29,923,000,000
↖ bpf_prog_run_xdp ← i	8.194s	21,148,500,000
↖ __tcf_classify ← tcf_cl	1.463s	609,500,000
↖ __tcf_classify ← tcf_cl	1.273s	621,000,000
↖ ice_napi_poll ← napi_	0.236s	322,000,000
[Unknown]	0.035s	23,000,000
func@0x79fe60	0.015s	11,500,000
↖ func@0x7a19d5 ← fur	0.005s	0
func@0x3527ee1	0.005s	34,500,000
↖ func@0x7a0570 ← fur	0s	11,500,000
ice.ko	25.519s	26,553,500,000
ip_tables.ko	17.896s	32,798,000,000
libc-2.13.so	11.096s	22,482,500,000
containerd-shim	7.943s	12,788,000,000
cls_bpf.ko	7.367s	8,533,000,000

- XDP vs No XDP, CPU utilization is quite similar
- Benefits of XDP is being able to bypass the kernel stack in case of redirect.
- Redirect to external port requires dedicate HW Redirect TX queue.
- XDP Benefits can be derived from dedicated HW resources.

eBPF Implementation Characterization – Cilium - cont

Cilium chains in iptables

```
# Generated by iptables-save v1.6.1 on Tue Dec 1 13:22:26 2020
*raw
:PREROUTING ACCEPT [339745057:40541787768]
:OUTPUT ACCEPT [1966765:7841806195]
:CILIUM_OUTPUT_raw - [0:0]
:CILIUM_PRE_raw - [0:0]
-A PREROUTING -m comment --comment "cilium-feeder: CILIUM_PRE_raw" -j CILIUM_PRE_raw
-A OUTPUT -m comment --comment "cilium-feeder: CILIUM_OUTPUT_raw" -j CILIUM_OUTPUT_raw
-A CILIUM_OUTPUT_raw -o lxc+ -m mark --mark 0xa00/0xffffffff -m comment --comment "cilium: NOTRACK for proxy return traff
-A CILIUM_OUTPUT_raw -o cilium_host -m mark --mark 0xa00/0xffffffff -m comment --comment "cilium: NOTRACK for proxy return
-A CILIUM_PRE_raw -m mark --mark 0x200/0xf00 -m comment --comment "cilium: NOTRACK for proxy traffic" -j NOTRACK
COMMIT
# Completed on Tue Dec 1 13:22:26 2020
# Generated by iptables-save v1.6.1 on Tue Dec 1 13:22:26 2020
*mangle
:PREROUTING ACCEPT [339745170:40541800199]
:INPUT ACCEPT [1689410:7607891639]
:FORWARD ACCEPT [338055760:32933908560]
:OUTPUT ACCEPT [1966765:7841806195]
:POSTROUTING ACCEPT [340022530:4077515026]
:CILIUM_POST_mangle - [0:0]
:CILIUM_PRE_mangle - [0:0]
-A PREROUTING -m comment --comment "cilium-feeder: CILIUM_PRE_mangle" -j CILIUM_PRE_mangle
-A POSTROUTING -m comment --comment "cilium-feeder: CILIUM_POST_mangle" -j CILIUM_POST_mangle
-A CILIUM_PRE_mangle -m socket --transparent -m comment --comment "cilium: any-xpod redirect proxied traffic to host pro
-A CILIUM_PRE_mangle -p tcp -m mark --mark 0xd7ae0200 -m comment --comment "cilium: TPROXY to host cilium-dns-egress prox
-A CILIUM_PRE_mangle -p udp -m mark --mark 0xd7ae0200 -m comment --comment "cilium: TPROXY to host cilium-dns-egress prox
COMMIT
# Completed on Tue Dec 1 13:22:26 2020
# Generated by iptables-save v1.6.1 on Tue Dec 1 13:22:26 2020
*filter
:INPUT ACCEPT [1689410:7607891639]
:FORWARD DROP [0:0]
:OUTPUT ACCEPT [1966765:7841806195]
:CILIUM_FORWARD - [0:0]
:CILIUM_INPUT - [0:0]
:CILIUM_OUTPUT - [0:0]
-A INPUT -m comment --comment "cilium-feeder: CILIUM_INPUT" -j CILIUM_INPUT
-A FORWARD -m comment --comment "cilium-feeder: CILIUM_FORWARD" -j CILIUM_FORWARD
-A OUTPUT -m comment --comment "cilium-feeder: CILIUM_OUTPUT" -j CILIUM_OUTPUT
-A CILIUM_FORWARD -o cilium_host -m comment --comment "cilium: any->cluster on cilium host forward accept" -j ACCEPT
-A CILIUM_FORWARD -i cilium_host -m comment --comment "cilium: cluster->any on cilium host forward accept (nodeport)" -j
-A CILIUM_FORWARD -i lxc+ -m comment --comment "cilium: cluster->any on lxc+ forward accept" -j ACCEPT
-A CILIUM_FORWARD -i cilium_net -m comment --comment "cilium: cluster->any on cilium_net forward accept (nodeport)" -j AC
-A CILIUM_INPUT -m mark --mark 0x200/0xf00 -m comment --comment "cilium: ACCEPT for proxy traffic" -j ACCEPT
-A CILIUM_OUTPUT -m mark --mark 0xa00/0xffffffff -m comment --comment "cilium: ACCEPT for proxy return traffic" -j ACCEPT
-A CILIUM_OUTPUT -m mark ! --mark 0xe00/0xf00 -m mark ! --mark 0xd00/0xf00 -m mark ! --mark 0xa00/0xe00 -m comment --comm
COMMIT
# Completed on Tue Dec 1 13:22:26 2020
# Generated by iptables-save v1.6.1 on Tue Dec 1 13:22:26 2020
*nat
:PREROUTING ACCEPT [0:0]
:INPUT ACCEPT [0:0]
```

Cilium's own conntrack table

```

TCP IN 12.91.212.200:37386 -> 10.0.0.241:80 expires=17272887 RxFPackets=18796 RBytes=2057656 RxFlagsSeen=0x1b LastRtReport=17251794 RxFPackets=18915 TBytes=4544990 TxFlagsSeen=0
TCP OUT 13.91.212.200:33224 -> 10.0.0.55:80 expires=17275420 RxFPackets=83205 RBytes=19555155 RxFlagsSeen=0x1a LastRtReport=17254320 TxFPackets=80142 TBytes=9897097 TxFlagsSeen=0
TCP OUT 12.91.212.202:30007 -> 12.91.212.200:36342 service expires=17272760 RxFPackets=0 RBytes=6 RxFlagsSeen=0x00 LastRtReport=0 TxFPackets=0 TBytes=0 TxFlagsSeen=0x1f LastTIn
TCP OUT 12.91.212.202:38900 -> 12.91.212.200:30008 expires=17274846 RxFPackets=2 RBytes=120 RxFlagsSeen=0x14 LastRtReport=17253753 TxFPackets=1 TBytes=74 TxFlagsSeen=0x02 LastTIn
TCP OUT 12.91.212.200:43362 -> 10.0.0.246:80 expires=17275288 RxFPackets=141178 RBytes=33482100 RxFlagsSeen=0x1a LastRtReport=17254194 TxFPackets=137700 TBytes=17005280 TxFlagsT
TCP IN 13.91.212.200:40604 -> 10.0.0.55:80 expires=17275556 RxFPackets=90047 RBytes=9859771 RxFlagsSeen=0x1b LastRtReport=17250463 TxFPackets=91583 TBytes=21846155 TxFlagsSeen=0
TCP IN 13.91.212.200:50588 -> 10.0.0.55:80 expires=17273171 RxFPackets=52349 RBytes=5757932 RxFlagsSeen=0x1a LastRtReport=17252078 TxFPackets=53089 TBytes=4943739 TxFlagsSeen=0
TCP OUT 12.91.212.200:43004 -> 10.0.0.241:80 expires=17275006 RxFPackets=1816 RBytes=227571 RxFlagsSeen=0x1a LastRtReport=17253912 TxFPackets=933 TBytes=113493 TxFlagsSeen=0x1e
TCP IN 12.91.212.200:43512 -> 10.0.0.246:80 expires=17275288 RxFPackets=140974 RBytes=17409857 RxFlagsSeen=0x1e LastRtReport=17254195 TxFPackets=145276 TBytes=3427857 TxFlagsS
TCP OUT 13.91.212.202:40840 -> 13.91.212.200:30008 expires=17274814 RxFPackets=2 RBytes=120 RxFlagsSeen=0x14 LastRtReport=17253721 TxFPackets=1 TBytes=74 TxFlagsSeen=0x02 LastTIn
TCP IN 12.91.212.200:44096 -> 10.0.0.241:80 expires=17275421 RxFPackets=102519 RBytes=12669607 RxFlagsSeen=0x1e LastRtReport=17254328 TxFPackets=104458 TBytes=2488841 TxFlagsS
TCP IN 12.91.212.200:43468 -> 10.0.0.246:80 expires=17275288 RxFPackets=102334 RBytes=16095587 RxFlagsSeen=0x1e LastRtReport=17254195 TxFPackets=133864 TBytes=31706643 TxFlagsS
TCP OUT 12.91.212.202:30007 -> 12.91.212.200:42952 service expires=17275006 RxFPackets=0 RBytes=6 RxFlagsSeen=0x00 LastRtReport=0 TxFPackets=0 TBytes=0 TxFlagsSeen=0x1e LastTIn
TCP OUT 13.91.212.202:30008 -> 13.91.212.200:54258 service expires=17272887 RxFPackets=0 RBytes=9 RxFlagsSeen=0x00 LastRtReport=0 TxFPackets=0 TBytes=0 TxFlagsSeen=0x1b LastTIn
TCP OUT 12.91.212.202:54410 -> 12.91.212.200:30007 expires=17274731 RxFPackets=2 RBytes=120 RxFlagsSeen=0x14 LastRtReport=17253638 TxFPackets=1 TBytes=74 TxFlagsSeen=0x02 LastTIn
TCP OUT 12.91.212.202:54426 -> 12.91.212.200:30007 expires=17274731 RxFPackets=2 RBytes=120 RxFlagsSeen=0x14 LastRtReport=17253638 TxFPackets=1 TBytes=74 TxFlagsSeen=0x02 LastTIn
TCP OUT 12.91.212.202:30007 -> 12.91.212.200:35150 service expires=17274401 RxFPackets=0 RBytes=0 RxFlagsSeen=0x00 LastRtReport=0 TxFPackets=0 TBytes=0 TxFlagsSeen=0x1f LastTIn
TCP IN 12.91.212.200:42930 -> 10.0.0.55:80 expires=17275006 RxFPackets=1174 RBytes=144442 RxFlagsSeen=0x1e LastRtReport=17253913 TxFPackets=1292 TBytes=289662 TxFlagsSeen=0x1a
TCP IN 13.91.212.200:33946 -> 10.0.0.246:80 expires=17275768 RxFPackets=24996 RBytes=3086447 RxFlagsSeen=0x1a LastRtReport=17254672 TxFPackets=25670 TBytes=6078020 TxFlagsSeen=0
TCP OUT 12.91.212.202:30007 -> 12.91.212.200:30580 expires=17273496 RxFPackets=0 RBytes=6 RxFlagsSeen=0x00 LastRtReport=0 TxFPackets=0 TBytes=0 TxFlagsSeen=0x1b LastTIn
TCP IN 13.91.212.200:34256 -> 10.0.0.55:80 expires=17275769 RxFPackets=21454 RBytes=2649240 RxFlagsSeen=0x1a LastRtReport=17254675 TxFPackets=22597 TBytes=5255183 TxFlagsSeen=0
TCP OUT 13.91.212.200:33150 -> 10.0.0.55:80 expires=17275420 RxFPackets=88664 RBytes=19005294 RxFlagsSeen=0x1a LastRtReport=17254316 TxFPackets=77972 TBytes=9629102 TxFlagsSeen=0
TCP OUT 12.91.212.202:30007 -> 12.91.212.200:44004 service expires=17275421 RxFPackets=0 RBytes=9 RxFlagsSeen=0x00 LastRtReport=0 TxFPackets=0 TBytes=0 TxFlagsSeen=0x1e LastTIn
TCP OUT 13.91.212.202:30008 -> 13.91.212.200:55654 service expires=17273400 RxFPackets=0 RBytes=9 RxFlagsSeen=0x00 LastRtReport=0 TxFPackets=0 TBytes=0 TxFlagsSeen=0x1e LastTIn
TCP OUT 12.91.212.202:55190 -> 12.91.212.200:30007 expires=17274765 RxFPackets=2 RBytes=120 RxFlagsSeen=0x14 LastRtReport=17253672 TxFPackets=1 TBytes=74 TxFlagsSeen=0x02 LastTIn
TCP OUT 12.91.212.202:39914 -> 12.91.212.200:30008 expires=17274929 RxFPackets=2 RBytes=120 RxFlagsSeen=0x14 LastRtReport=17253836 TxFPackets=1 TBytes=74 TxFlagsSeen=0x02 LastTIn
TCP OUT 12.91.212.202:38932 -> 12.91.212.200:30008 expires=17274846 RxFPackets=2 RBytes=120 RxFlagsSeen=0x14 LastRtReport=17253753 TxFPackets=1 TBytes=74 TxFlagsSeen=0x02 LastTIn
TCP OUT 12.91.212.202:38888 -> 12.91.212.200:30008 expires=17274846 RxFPackets=2 RBytes=120 RxFlagsSeen=0x14 LastRtReport=17253753 TxFPackets=1 TBytes=74 TxFlagsSeen=0x02 LastTIn
TCP IN 12.91.212.202:30007 -> 12.91.212.200:44186 service expires=17275421 RxFPackets=0 RBytes=0 RxFlagsSeen=0x00 LastRtReport=0 TxFPackets=0 TBytes=0 TxFlagsSeen=0x1e LastTIn
TCP IN 12.91.212.200:44194 -> 10.0.0.246:80 expires=17275421 RxFPackets=115638 RBytes=14280861 RxFlagsSeen=0x1e LastRtReport=17254328 TxFPackets=118303 TBytes=28101171 TxFlagsS
TCP OUT 12.91.212.202:30007 -> 12.91.212.200:42722 service expires=17274983 RxFPackets=0 RBytes=8 RxFlagsSeen=0x00 LastRtReport=0 TxFPackets=0 TBytes=0 TxFlagsSeen=0x1e LastTIn

```

Q&A



