



# ESnet

ENERGY SCIENCES NETWORK

# 1:1 Packet Sampling

**Bruce A. Mah, PhD**  
Energy Sciences Network  
Lawrence Berkeley National Laboratory

12th SIG-NGN Meeting  
26 October 2023



U.S. DEPARTMENT OF  
**ENERGY**  
Office of Science




# High-Touch Architecture and Design

FPGA (AMD/Xilinx  
Alveo U280)

High Touch Server

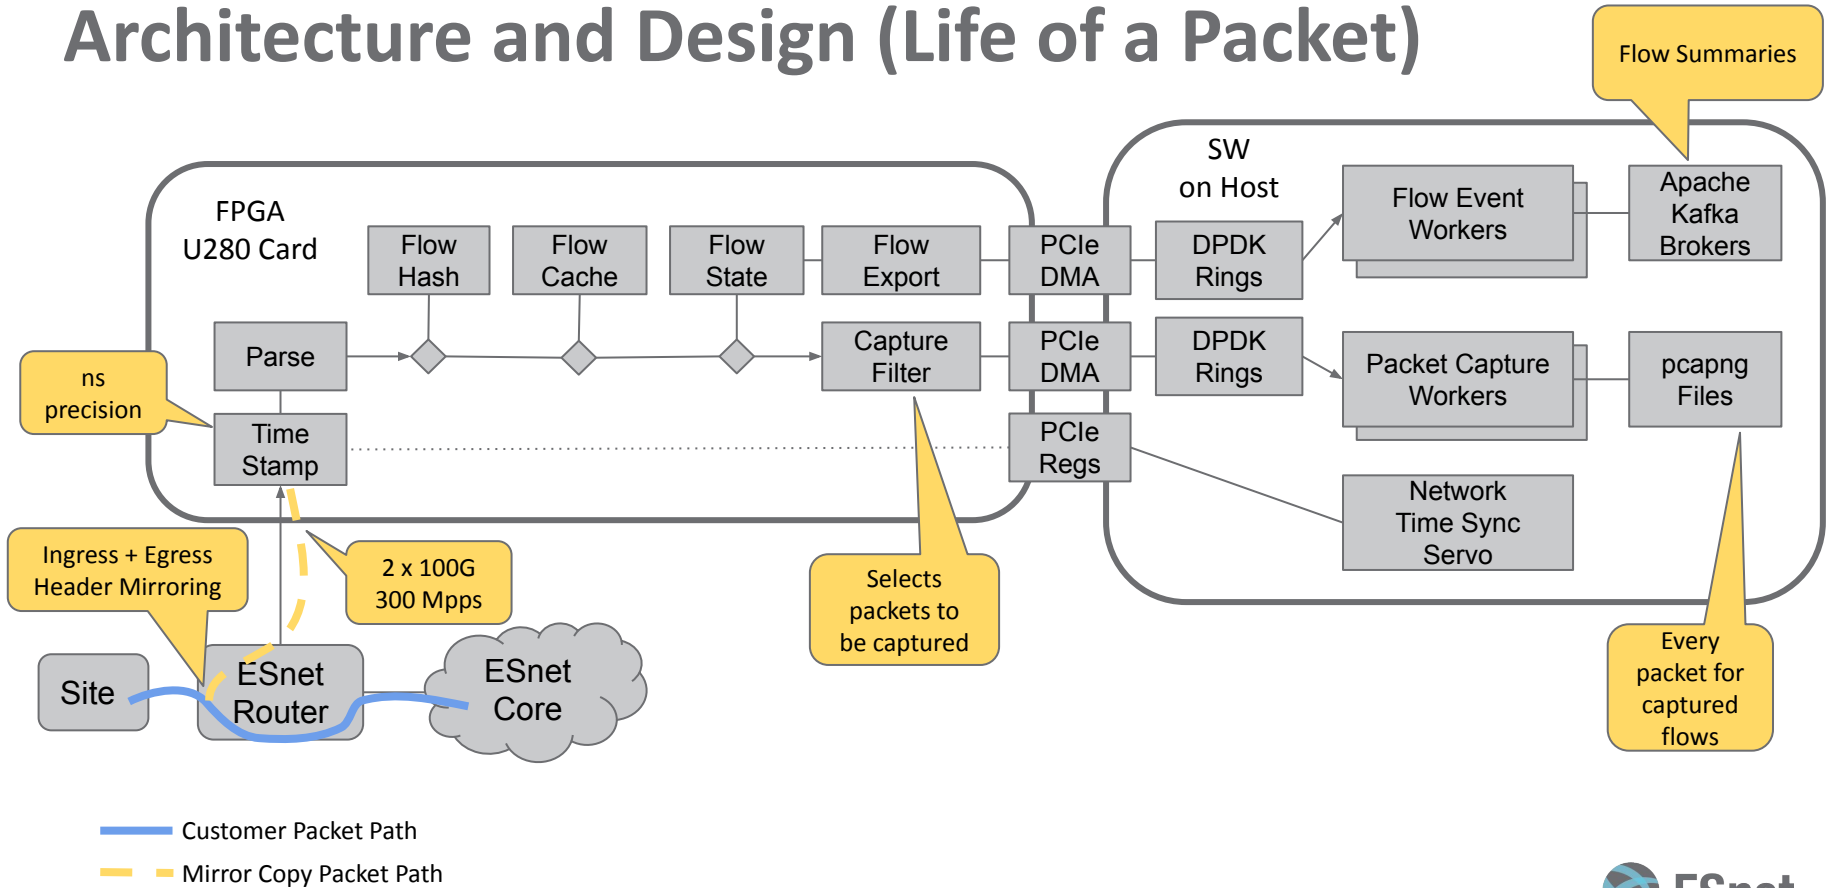
Router



A combination of programmable hardware and software, placed throughout the network, to provide new network services, such as high-precision network measurements.



# Architecture and Design (Life of a Packet)



# Flow Summary Record (JSON Representation)

```
{"reporting-interval":{"start-time":"2023-10-24T17:44:08,133416641+00:00","start-time-ns":1698169448133416641,"end-time":"2023-10-24T17:44:08,333491929+00:00","end-time-ns":1698169448333491929,"export-reason":"done"},"mirror-context":{"exporting-node":"lbnl59-ht1","router-name":"lbnl59-cr6","router-interface":"1\1\c17\3","direction":"in","vlan-id":47,"sap-name":"esnet_se-425","sap-type":"Layer3 Virtual Interface","sap-routing-instance":"Base","sap-bgp-policy-summary":"Site","sap-organization-name":"ESNET"},"flow-key":{"afi":"ipv4","ip-version":4,"ip-src":"X.X.X.X","ip-dst":"Y.Y.Y.Y","ip-proto-num":6,"ip-proto":"tcp","l4-src-port":43436,"l4-dst-port":80},"tcp-flags":{"syn-ack-only":false,"syn-only":true,"cwr":false,"ece":false,"urg":false,"ack":true,"psh":true,"rst":false,"syn":true,"fin":true},"tcp-seq":{"next-seq":2542394246,"loss-events":0,"retransmit-events":0,"retransmit-bytes":0},"packet-timestamps-ns":[1698169448133416641,1698169448199122842],"stats":{"packets":8,"bytes":503,"packet-size-histogram":{"bin-labels":["<64","64-127","128-255","256-511","512-1023","1024-2047","2048-4095","4096-65535"],"bin-lower-bounds":[0,64,128,256,512,1024,2048,4096],"bin-upper-bounds":[63,127,255,511,1023,2047,4095,65535],"bin-counts":[6,1,1,0,0,0,0,0]}}
```

# Flow Summary Records (JSON Representation)

```
{ "reporting-interval": { "start-time": "2023-10-24T17:44:08,133416641+00:00", "start-time-ns": 1698169448133416641, "end-time": "2023-10-24T17:44:08,333491929+00:00", "end-time-ns": 1698169448333491929, "export-reason": "done" }, "mirror-context": { "exporting-node": "lbn159-ht1", "router-name": "lbn159-cr6", "router-interface": "1\\1\\c17\\3", "direction": "in", "vlan-id": 47, "sap-name": "esnet_se-425", "sap-type": "Layer3 Virtual Interface", "sap-routing-instance": "Base", "sap-bgp-policy-summary": "Site", "sap-organization-name": "ESNET" }, "flow-key": { "afi": "ipv4", "ip-version": 4, "ip-src": "X.X.X.X", "ip-dst": "Y.Y.Y.Y", "ip-proto-num": 6, "ip-proto": "tcp", "l4-src-port": 43436, "l4-dst-port": 80 }, "tcp-flags": { "syn-ack-only": false, "syn-only": true, "cwr": false, "ece": false, "urg": false, "ack": true, "psh": true, "rst": false, "syn": true, "fin": true }, "tcp-seq": { "next-seq": 2542394246, "loss-events": 0, "retransmit-events": 0, "retransmit-bytes": 0 }, "packet-timestamps-ns": [1698169448133416641, 1698169448199122842], "stats": { "packets": 8, "bytes": 503, "packet-size-histogram": { "bin-labels": ["<64", "64-127", "128-255", "256-511", "512-1023", "1024-2047", "2048-4095", "4096-65535"], "bin-lower-bounds": [0, 64, 128, 256, 512, 1024, 2048, 4096], "bin-upper-bounds": [63, 127, 255, 511, 1023, 2047, 4095, 65535], "bin-counts": [6, 1, 1, 0, 0, 0, 0, 0] } } }
```

# Flow Summary Records (JSON Representation)

```
{"reporting-interval":{"start-time":"2023-10-24T17:44:08,133416641+00:00","start-time-ns":1698169448133416641,"end-time":"2023-10-24T17:44:08,333491929+00:00","end-time-ns":1698169448333491929,"export-reason":"done"},"mirror-context":{"exporting-node":"lbn159-ht1","router-name":"lbn159-cr6","router-interface":"1\1\c17\3","direction":"in","vlan-id":47,"sap-name":"esnet se-425","sap-type":"Layer3 Virtual Interface","sap-routing-instance":"Base","sap-bgp-policy-summary":"Site","sap-organization-name":"ESNET"},"flow-key":{"afi":"ipv4","ip-version":4,"ip-src":"X.X.X.X","ip-dst":"Y.Y.Y.Y","ip-proto-num":6,"ip-proto":"tcp","l4-src-port":43436,"l4-dst-port":80},"tcp-flags":{"syn-ack-only":false,"syn-only":true,"cwr":false,"ece":false,"urg":false,"ack":true,"psh":true,"rst":false,"syn":true,"fin":true},"tcp-seq":{"next-seq":2542394246,"loss-events":0,"retransmit-events":0,"retransmit-bytes":0},"packet-timestamps-ns":[1698169448133416641,1698169448199122842],"stats":{"packets":8,"bytes":503,"packet-size-histogram":{"bin-labels":["<64","64-127","128-255","256-511","512-1023","1024-2047","2048-4095","4096-65535"],"bin-lower-bounds":[0,64,128,256,512,1024,2048,4096],"bin-upper-bounds":[63,127,255,511,1023,2047,4095,65535],"bin-counts":[6,1,1,0,0,0,0,0]}}}
```

# Flow Summary Records (JSON Representation)

```
{"reporting-interval":{"start-time":"2023-10-24T17:44:08,133416641+00:00","start-time-ns":1698169448133416641,"end-time":"2023-10-24T17:44:08,333491929+00:00","end-time-ns":1698169448333491929,"export-reason":"done"},"mirror-context":{"exporting-node":"lbnl59-ht1","router-name":"lbnl59-cr6","router-interface":"1\1\c17\3","direction":"in","vlan-id":47,"sap-name":"esnet_se-425","sap-type":"Layer3 Virtual Interface","sap-routing-instance":"Base","sap-bgp-policy-summary":"Site","sap-organization-name":"ESNET"},"flow-key":{"afi":"ipv4","ip-version":4,"ip-src":"X.X.X.X","ip-dst":"Y.Y.Y.Y","ip-proto-num":6,"ip-proto":"tcp","l4-src-port":43436,"l4-dst-port":80},"tcp-flags":{"syn-ack-only":false,"syn-only":true,"cwr":false,"ece":false,"urg":false,"ack":true,"psh":true,"rst":false,"syn":true,"fin":true},"tcp-seq":{"next-seq":2542394246,"loss-events":0,"retransmit-events":0,"retransmit-bytes":0},"packet-timestamps-ns":[1698169448133416641,1698169448199122842],"stats":{"packets":8,"bytes":503,"packet-size-histogram":{"bin-labels":["<64","64-127","128-255","256-511","512-1023","1024-2047","2048-4095","4096-65535"],"bin-lower-bounds":[0,64,128,256,512,1024,2048,4096],"bin-upper-bounds":[63,127,255,511,1023,2047,4095,65535],"bin-counts":[6,1,1,0,0,0,0,0]}}
```



# Flow Summary Records (JSON Representation)

```
{"reporting-interval":{"start-time":"2023-10-24T17:44:08,133416641+00:00","start-time-ns":1698169448133416641,"end-time":"2023-10-24T17:44:08,333491929+00:00","end-time-ns":1698169448333491929,"export-reason":"done"},"mirror-context":{"exporting-node":"lbnl59-ht1","router-name":"lbnl59-cr6","router-interface":"1\1\c17\3","direction":"in","vlan-id":47,"sap-name":"esnet_se-425","sap-type":"Layer3 Virtual Interface","sap-routing-instance":"Base","sap-bgp-policy-summary":"Site","sap-organization-name":"ESNET"},"flow-key":{"afi":"ipv4","ip-version":4,"ip-src":"X.X.X.X","ip-dst":"Y.Y.Y.Y","ip-proto-num":6,"ip-proto":"tcp","l4-src-port":43436,"l4-dst-port":80},"tcp-flags":{"syn-ack-only":false,"syn-only":true,"cwr":false,"ece":false,"urg":false,"ack":true,"psh":true,"rst":false,"syn":true,"fin":true},"tcp-seq":{"next-seq":2542394246,"loss-events":0,"retransmit-events":0,"retransmit-bytes":0},"packet-timestamps-ns":[1698169448133416641,1698169448199122842],"stats":{"packets":8,"bytes":503,"packet-size-histogram":{"bin-labels":["<64","64-127","128-255","256-511","512-1023","1024-2047","2048-4095","4096-65535"],"bin-lower-bounds":[0,64,128,256,512,1024,2048,4096],"bin-upper-bounds":[63,127,255,511,1023,2047,4095,65535],"bin-counts":[6,1,1,0,0,0,0,0]}}
```

# Flow Summary Records (JSON Representation)

```
{"reporting-interval":{"start-time":"2023-10-24T17:44:08,133416641+00:00","start-time-ns":1698169448133416641,"end-time":"2023-10-24T17:44:08,333491929+00:00","end-time-ns":1698169448333491929,"export-reason":"done"},"mirror-context":{"exporting-node":"lbnl59-ht1","router-name":"lbnl59-cr6","router-interface":"1\1\c17\3","direction":"in","vlan-id":47,"sap-name":"esnet_se-425","sap-type":"Layer3 Virtual Interface","sap-routing-instance":"Base","sap-bgp-policy-summary":"Site","sap-organization-name":"ESNET"},"flow-key":{"afi":"ipv4","ip-version":4,"ip-src":"X.X.X.X","ip-dst":"Y.Y.Y.Y","ip-proto-num":6,"ip-proto":"tcp","l4-src-port":43436,"l4-dst-port":80},"tcp-flags":{"syn-ack-only":false,"syn-only":true,"cwr":false,"ece":false,"urg":false,"ack":true,"psh":true,"rst":false,"syn":true,"fin":true},"tcp-seq":{"next-seq":2542394246,"loss-events":0,"retransmit-events":0,"retransmit-bytes":0},"packet-timestamps-ns":[1698169448133416641,1698169448199122842],"stats":{"packets":8,"bytes":503,"packet-size-histogram":{"bin-labels":["<64","64-127","128-255","256-511","512-1023","1024-2047","2048-4095","4096-65535"],"bin-lower-bounds":[0,64,128,256,512,1024,2048,4096],"bin-upper-bounds":[63,127,255,511,1023,2047,4095,65535],"bin-counts":[6,1,1,0,0,0,0,0]}}
```

# Flow Summary Record (JSON Representation)

```
{"reporting-interval":{"start-time":"2023-10-24T17:44:08,133416641+00:00","start-time-ns":1698169448133416641,"end-time":"2023-10-24T17:44:08,333491929+00:00","end-time-ns":1698169448333491929,"export-reason":"done"},"mirror-context":{"exporting-node":"lbnl59-ht1","router-name":"lbnl59-cr6","router-interface":"1\1\c17\3","direction":"in","vlan-id":47,"sap-name":"esnet_se-425","sap-type":"Layer3 Virtual Interface","sap-routing-instance":"Base","sap-bgp-policy-summary":"Site","sap-organization-name":"ESNET"},"flow-key":{"afi":"ipv4","ip-version":4,"ip-src":"X.X.X.X","ip-dst":"Y.Y.Y.Y","ip-proto-num":6,"ip-proto":"tcp","l4-src-port":43436,"l4-dst-port":80},"tcp-flags":{"syn-ack-only":false,"syn-only":true,"cwr":false,"ece":false,"urg":false,"ack":true,"psh":true,"rst":false,"syn":true,"fin":true},"tcp-seq":{"next-seq":2542394246,"loss-events":0,"retransmit-events":0,"retransmit-bytes":0},"packet-timestamps-ns":[1698169448133416641,1698169448199122842],"stats":{"packets":8,"bytes":503,"packet-size-histogram":{"bin-labels":["<64","64-127","128-255","256-511","512-1023","1024-2047","2048-4095","4096-65535"],"bin-lower-bounds":[0,64,128,256,512,1024,2048,4096],"bin-upper-bounds":[63,127,255,511,1023,2047,4095,65535],"bin-counts":[6,1,1,0,0,0,0,0]}}}
```

# Packet Capture Service

“tcpdump for the network edge”

Desired capture specified by IP 5-tuple

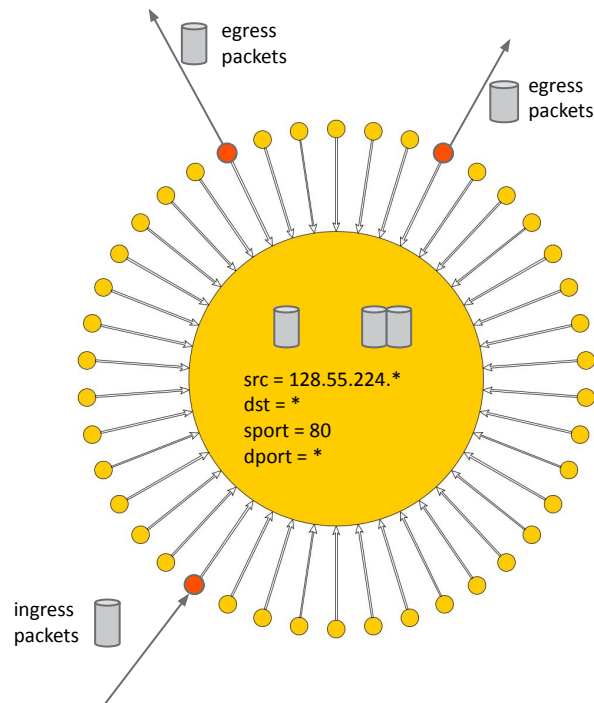
First 128 bytes of all (not a sampled subset) matching packets, saved to pcapng files

Single router (one at a time) today

Full-network application under development

- Simultaneous capture on all edge interfaces

- Merge matching packet captures to a single file



# Packet capture data

Like pcapNG file for multiple interface capture

Every “interface” is a router plus router port

```
Frame 1: 52 bytes on wire (416 bits), 52 bytes captured (416 bits) on interface lbn159-cr6
2/1/c2/1, id 6
  Interface id: 6 (lbn159-cr6 2/1/c2/1)
  Interface name: lbn159-cr6 2/1/c2/1
  Interface description:
  Encapsulation type: Raw IP (7)
```

Merging pcapNG files lets us see packets at ingress and egress

# Benefits of 1:1 Sampling

See every flow

- Background radiation traffic (often single packet flows)

- Normal (but short flows) are represented

More accurate view of flow sizes

- Exact packet counts, byte counts, timing

- Potentially finer-grain time behavior than traditional flow systems

# Benefits of 1:1 Sampling

Intra-flow visibility from flow summaries and individual packets

Insight into flow / TCP dynamics, packet loss

Multiple vantage points (see flow on ingress and egress)

# Questions, Suggestions, Ideas

Talk to us!

[hightouch@es.net](mailto:hightouch@es.net)