ESnet Production SDN Solution Yatish Kumar, Chin Guok, Richard Cziva





This is our hardware solution

2 x 40G Ports Capable of full line rate telemetry of 100G router port

10 ns Timestamps Fully programmable header parsing: - TCP extraction / MPLS segment routing headers / ACLs

Fully programmable telemetry generator: **ESnet6 Custom Telemetry Packet:**



HT version	esnet6
HT sensorID.	1
HT original to seconds:	10/13
HT original to net	216308605
ID vorsion:	Z 10090090
	4
IF IFL.	0
IP UIIISERV.	U 70
IP total length:	79
IP Identification:	
IP flags:	000
IP frag offset:	0
IP ttl:	64
IP protocol:	17
IP checksum:	50760
IP source address:	192.169.50
IP destination address:	192.169.1.1
TCP source port:	0
TCP destination port:	8000
TCP sequence number:	1000
TCP ACK number:	0
TCP ECN flags:	000
TCP control flags:	000010
TCP window size:	8192
TCP urgent pointer:	0
Total s recorded:	21
Total bytes recorded:	1659







We place these cards EVERYWHERE.



\$. 600 **MSRP** NICs 100 \$ 60,000

That is a LOT of SDN for \$60k





We place these cards EVERYWHERE.



\$. **MSRP** 600 NICs 400 \$240,000

Let's do it again **!!!** 4 NICs for all.

We can keep going...







This is production not a testbed



We can build real operational use cases without fear of scale.

The hardware is in place day 1.

Software will not be built as a "Proof of Concept". But rather as an operational quality 24/7 service.







What can we do with it?



Fine Print: whilst we could do EVERY flow, that would be daft. Every interesting flow is a what we really plan to do.

Stateful tracking of TCP flow rate (packet_len)/(Delta Timestamps) Octets / ns which gives us Gbits / s

For EVERY flow on ESnet in real time..







Here's what we see when we zoom in



4! microseconds of a TCP buildup of a 10 Gbit/s flow ~12 Gbit/s



Dare I say it ?

ML can be applied here.





lperf







Not Just 40G



- (4 x 200G QSFP-DD) 800G (NOW)
- (4 x 400G QSFP-DD) 1600G (2020)



