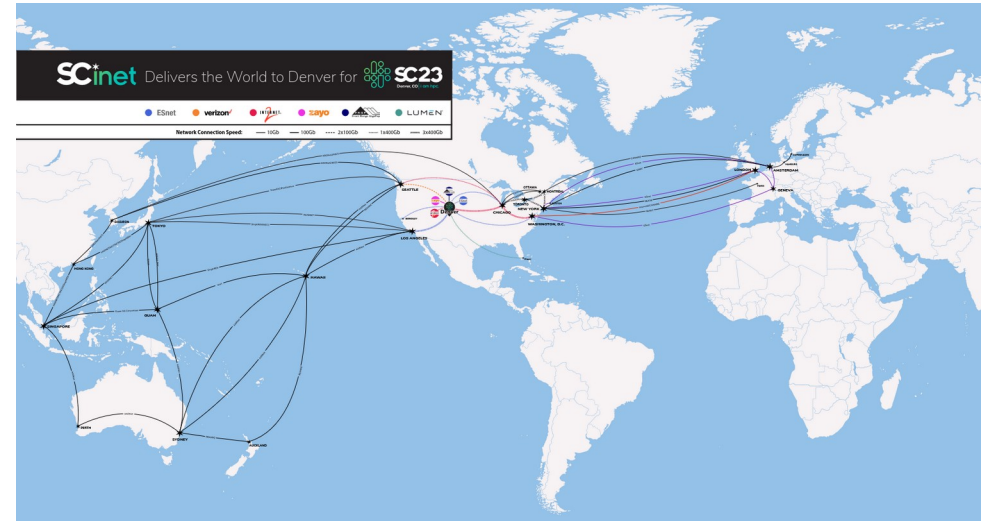


SC24 Scitags Demo

Tristan Sullivan, UVic
Beijing LHCONe Meeting
Oct. 12/24

SC24

- Held in Atlanta, Georgia, Nov. 17-22
- Scinet: Global network built specifically for SC
- Networking Research Exhibitions (NREs): Demonstrations of new network features, cutting-edge performance, etc.



SC23 Network Map

Scitags

- The Research Networking Technical Working Group was charged with improving visibility into R&E network usage
- Add a label at the packet level indicating the experiment to which the traffic belongs and the purpose of the traffic
- Implemented via IPv6 flow label
- [IETF RFC-Informational Draft](#) submitted June 2023: suggestion to look closer at using IPv6 extension headers instead

IPv6 Header

| Offset | Octet | 0 | | | | | | | | 1 | | | | | | | | 2 | | | | | | | | 3 | | | | | | | |
|--------|-------|-----------------------|---|---|---|----------------------|---|---|---|-------------------|---|----|----|----|----|----|----|--------------------|----|----|----|----|----|----|----|------------------|----|----|----|----|----|----|----|
| Octet | Bit | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 |
| 0 | 0 | <i>Version</i> | | | | <i>Traffic class</i> | | | | <i>Flow label</i> | | | | | | | | | | | | | | | | | | | | | | | |
| 4 | 32 | <i>Payload length</i> | | | | | | | | | | | | | | | | <i>Next header</i> | | | | | | | | <i>Hop limit</i> | | | | | | | |

- Overwriting flow label lightest-weight option
- Adding extension header requires changing next header field and inserting new header
- Next header values:
 - 0: Hop-By-Hop
 - 6: TCP
 - 60: Destination Options

How To Mark Packets?

What is eBPF?



Feature in the Linux kernel that allows user code to be injected into the kernel at runtime

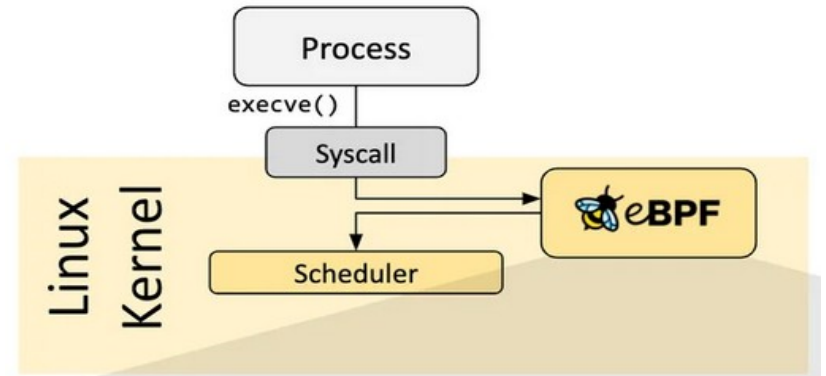
Can be a simpler replacement for kernel modules

How does it work?

First, the code is run through a verifier to ensure it is safe to execute

Then it is attached to some kernel hook

In our case, the code is run whenever a packet is sent



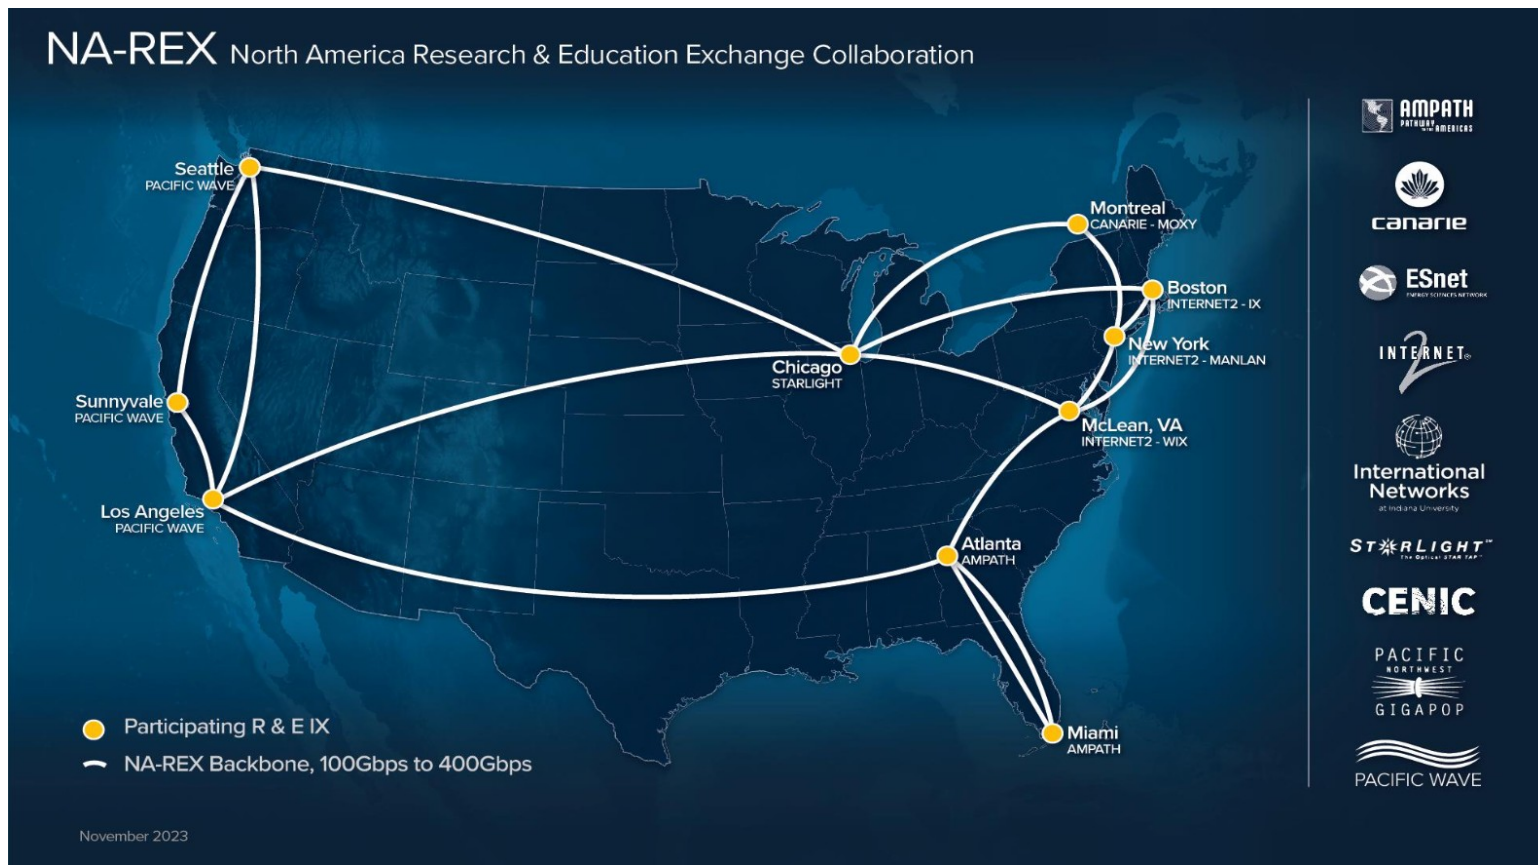
Network Path

UVic connected to CANARIE test link, which peers with Pacific Wave in Seattle before going to NA-RX

400G path end-to-end

UVic 400G DTN donated by Dell

At SC we'll use the Starlight booth and 400G DTN



From Joe Mambretti's presentation at LHCONE Catania Meeting

Demo Plan

- Try to achieve 400G from show floor to UVic
- Measure impact of flow label marking
- Measure impact of extension header marking