

Analytics-Driven Networking: when the computer becomes the network

Inder Monga

Director, Energy Sciences Network

Division Director, Berkeley Lab

June 12th, 2020

Where does my perspective come from?

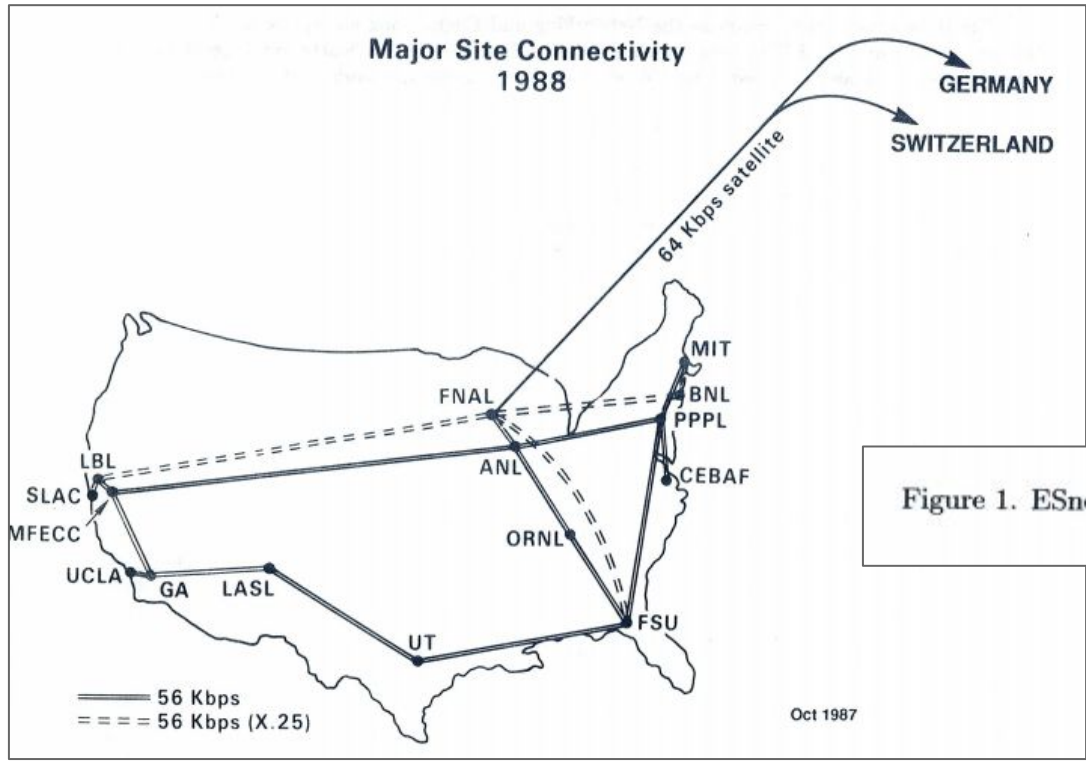
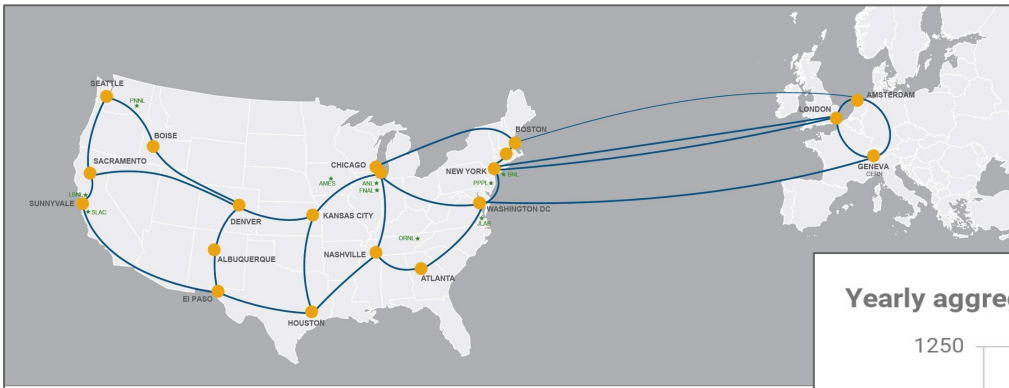
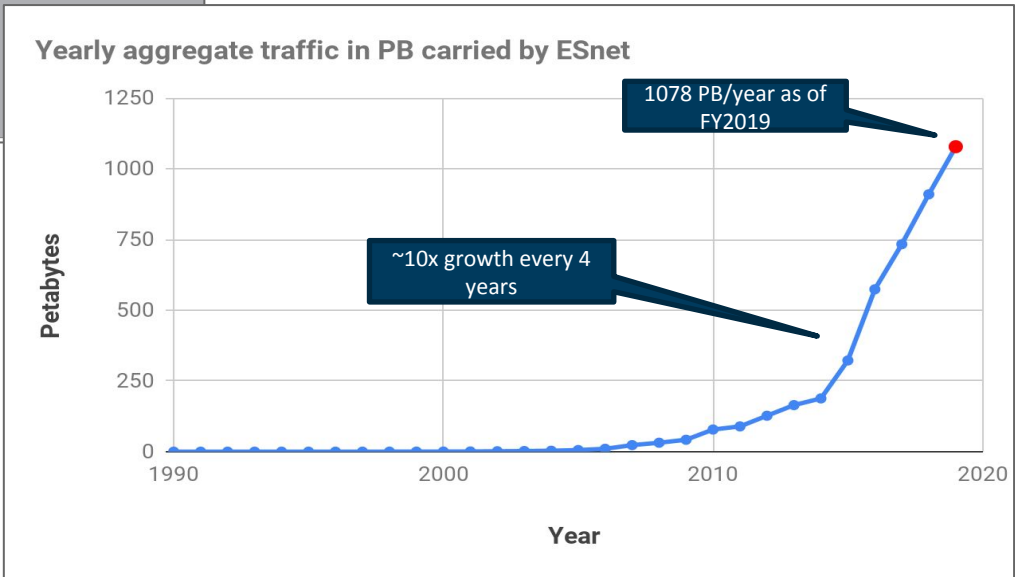


Figure 1. ESnet—major sites to be connected in 1988.

ESnet: Department of Energy's high-performance science network (multi-terabit) facility



Vision: Scientific progress will be **completely unconstrained** by the physical location of instruments, people, computational resources, or data.




Serves the Larger Research Complex at DOE




Provides connectivity to all of the DOE labs, experiment sites, and user facilities

Light sources data generation increasing by 10x




Year	Facility				
	ALS	APS	LCLS/LCLS-II	NSLS-II	SSRL
2021	3 PB	7 PB	30 PB	42 PB	15 PB
2028	31 PB	243 PB	300 PB	85 PB	15 PB

Terabits/sec streaming data expected



LCLS SLAC NATIONAL ACCELERATOR LABORATORY




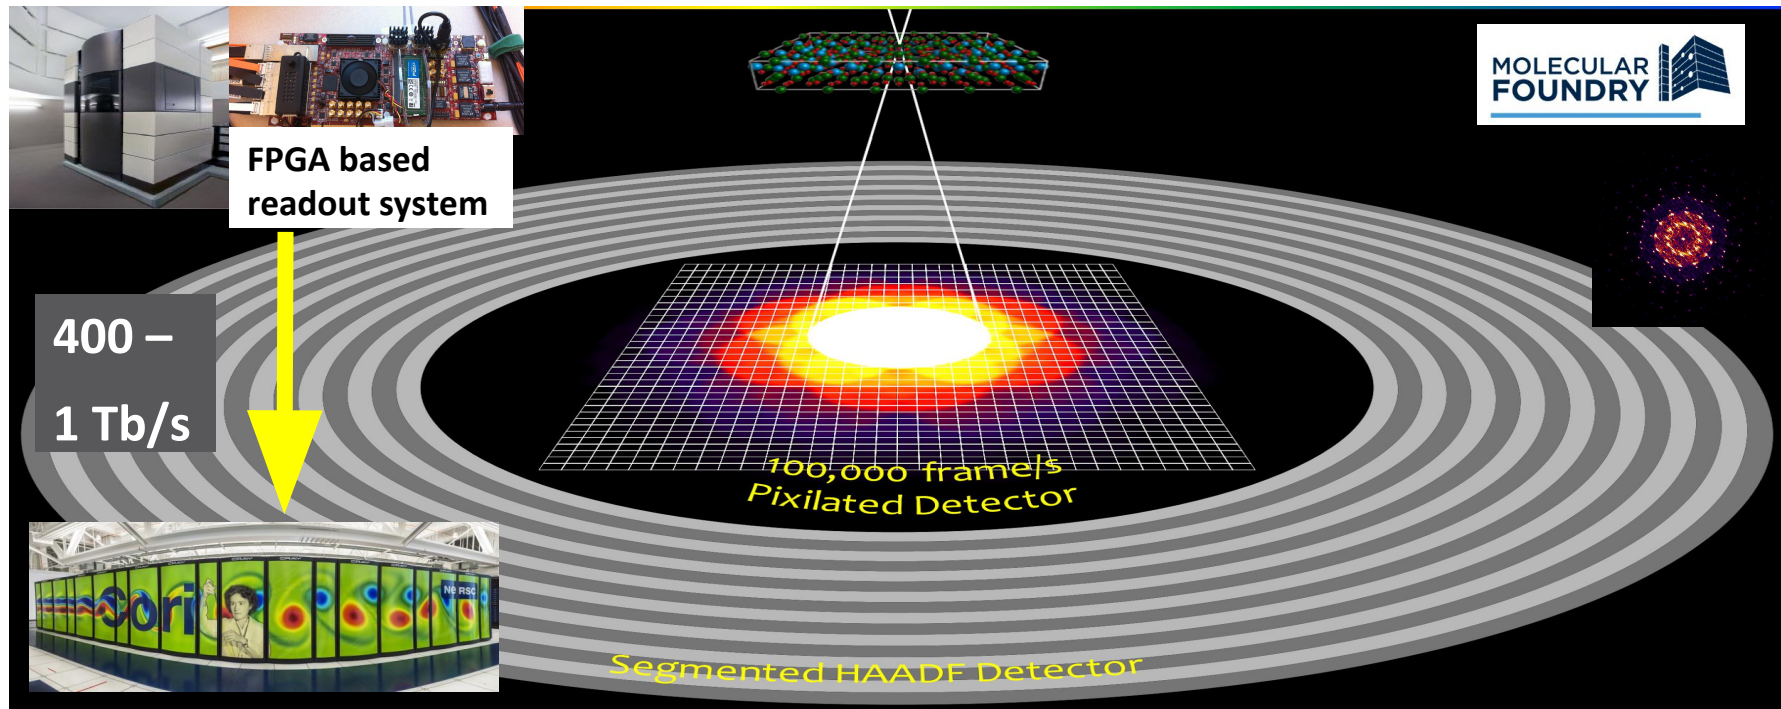
LU34 experiment: Taking Snapshots of O-O Bond Formation in Photosynthetic Water-Splitting Using Simultaneous X-ray Emission Spectroscopy and Crystallography – Y. Vital (LCLS PI)

Diffraction pattern from LU34

Certain experiments expected to stream data to HPCs at ~ 1 Tb/s



NCEM 4D-STEM: Near real-time processing necessitates on-demand coupling with supercomputer



On-demand use of cloud provides unprecedented scalability when needed

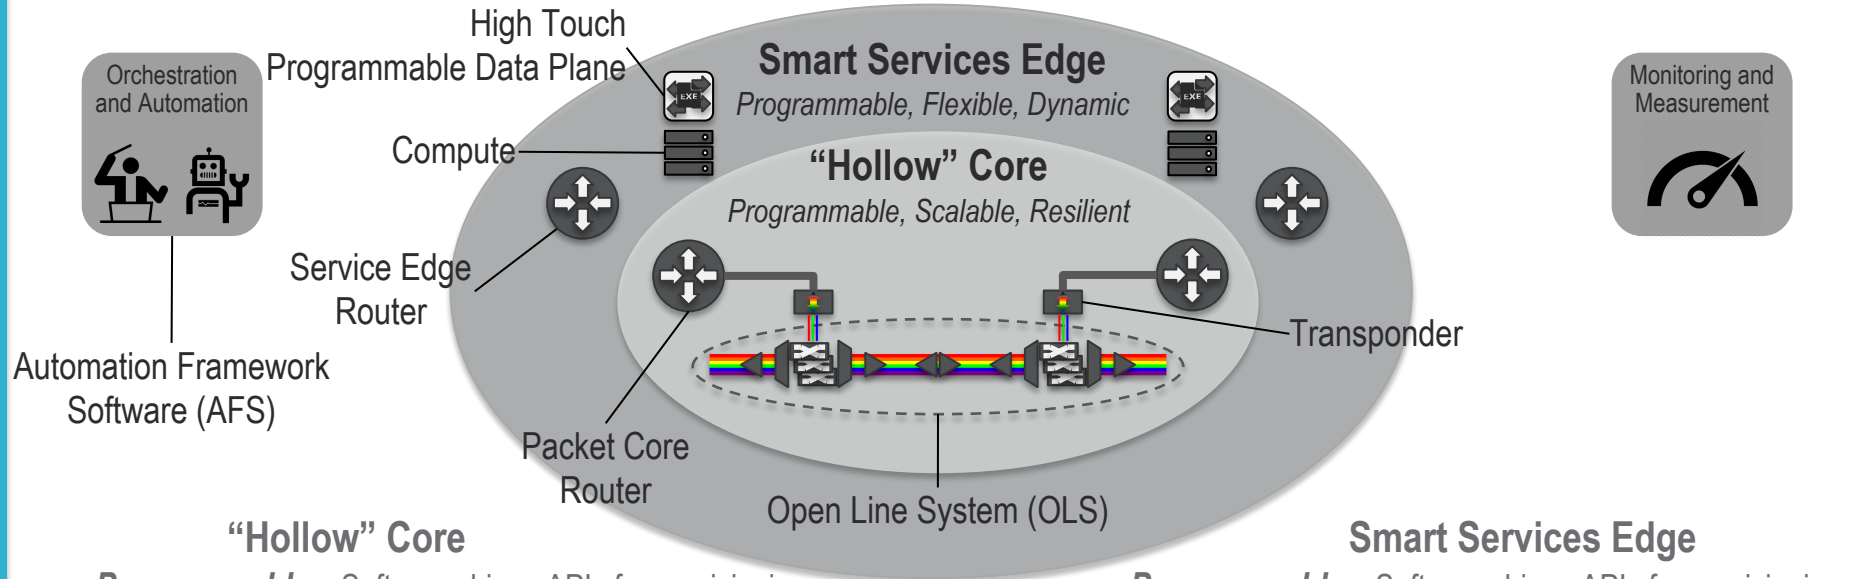


100Gbps peering with Google cloud to enable staging the data for this experiment



https://research.cs.wisc.edu/htcondor/HTCondorWeek2017/presentations/WedTimm_GCE.pdf

ESnet6: New greenfield network build in progress




"Hollow" Core

- **Programmable** – Software driven APIs for provisioning and monitoring.
- **Scalable** – Increased capacity and flexibility.
- **Resilient** – Protection and restoration using next generation Traffic Engineering (TE) protocols (e.g. Segment Routing (SR)).

Smart Services Edge

- **Programmable** – Software driven APIs for provisioning and telemetry.
- **Flexible** - Data plane programmable in conjunction with compute resources to prototype new services
- **Dynamic** – Dynamic instantiation of services.

Context behind my talk

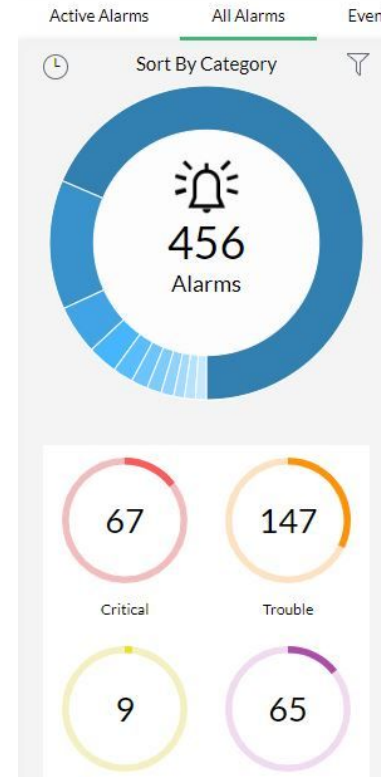


Most of us think of the doctor only when something breaks

We treat the network as the patient

- Manage through alarms
- Continuous monitoring of statistics
 - Graphs very similar to fitbit like health monitors
 - Long term planning or visual anomalies
- Network offers only crude statistical sampling
 - Per-flow measurement and performance management on all flows is considered near impossible

None of us know what's going on in the network just like we don't know what's happening in each of our organs




Network Error


Something went wrong on our end. Please try again later.

OK

Quantified Self: self-knowledge through data



Heart rate to diagnose onset of sickness



DonM99
Jogger
3 ✓ 0 🍊 1

07-31-2016 13:30

I've noticed the same thing a few illnesses ago. That was my worst one in a long time and it made my RHR go up 20+ so it would have been hard not to notice. It always goes up around 5-8 points before onset of symptoms. Obviously I'm sick now 🙄







Blood sugar correlated with steps



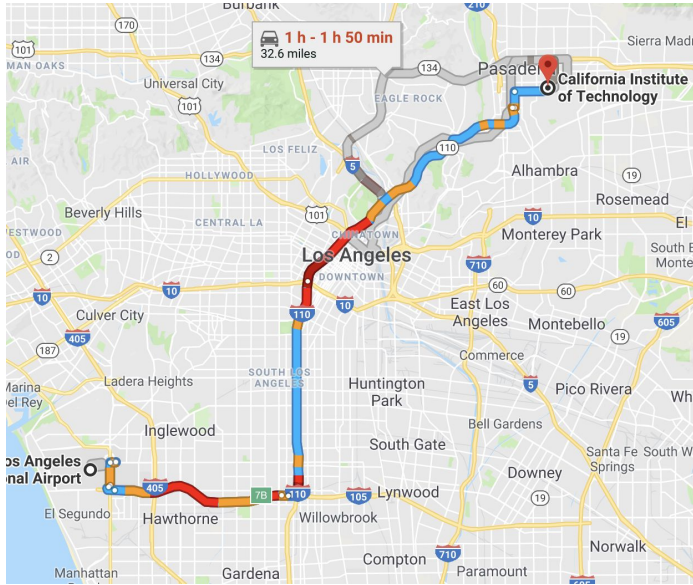
Twelve days of blood sugar data overlaid with steps (the shaded areas). Image: Eric Jain

Q1: What kind of telemetry do we need to scale/support the wide variety of apps?

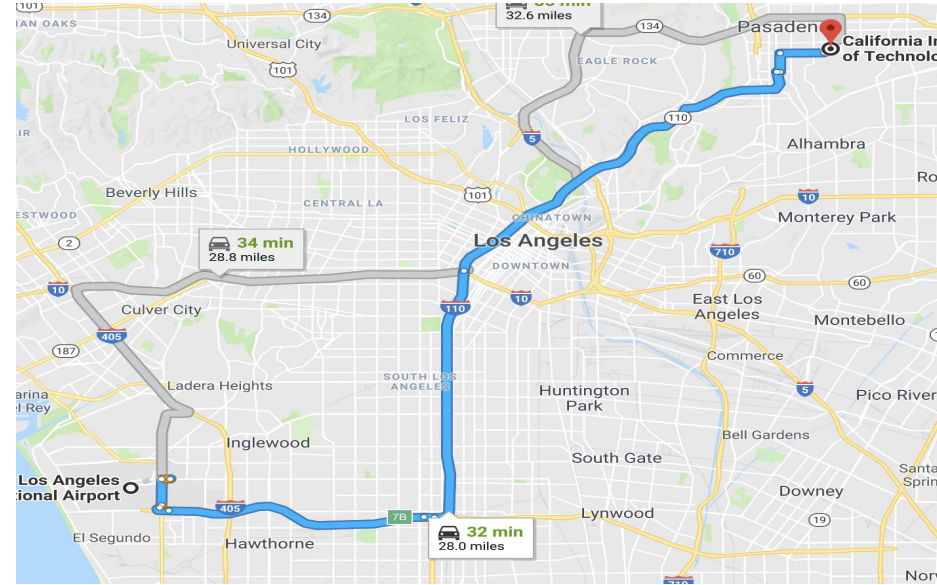


Some images from Intel website

Tools provide predictability in our everyday lives



LAX– Caltech, 6 pm:
1 hr – 1hr 50 min



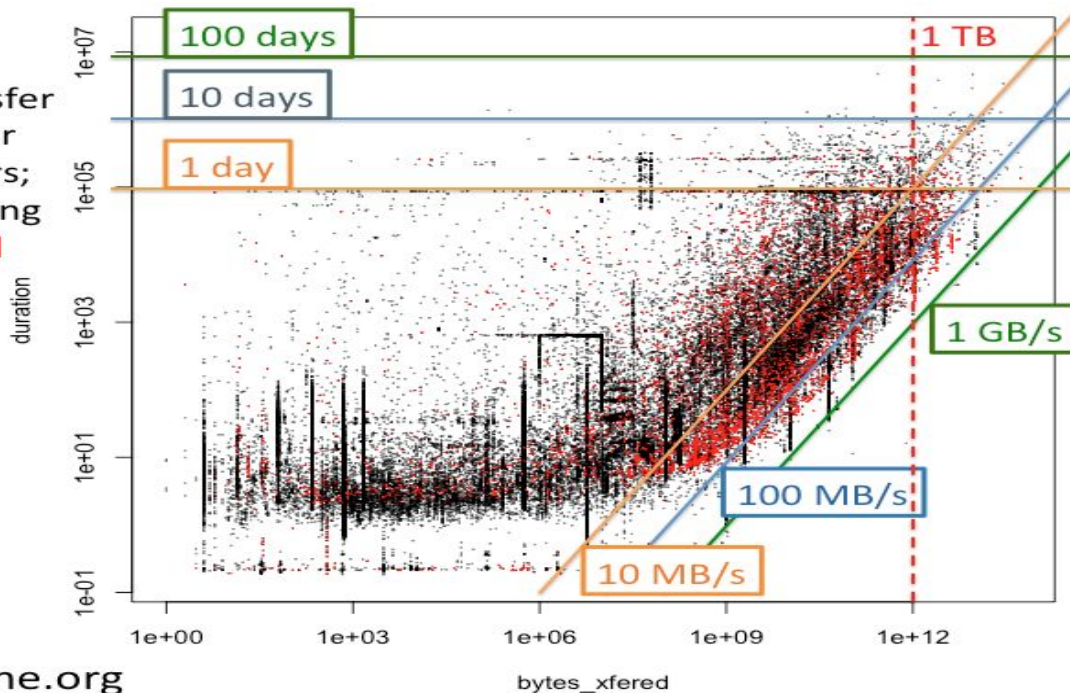
LAX– Caltech, 11 pm:
32 min

Q2: Why can't networks provide better predictability envelope for data transfer apps?

- Transfers over a shared network are not predictable
- Best-effort delivery can also mean worst-effort delivery



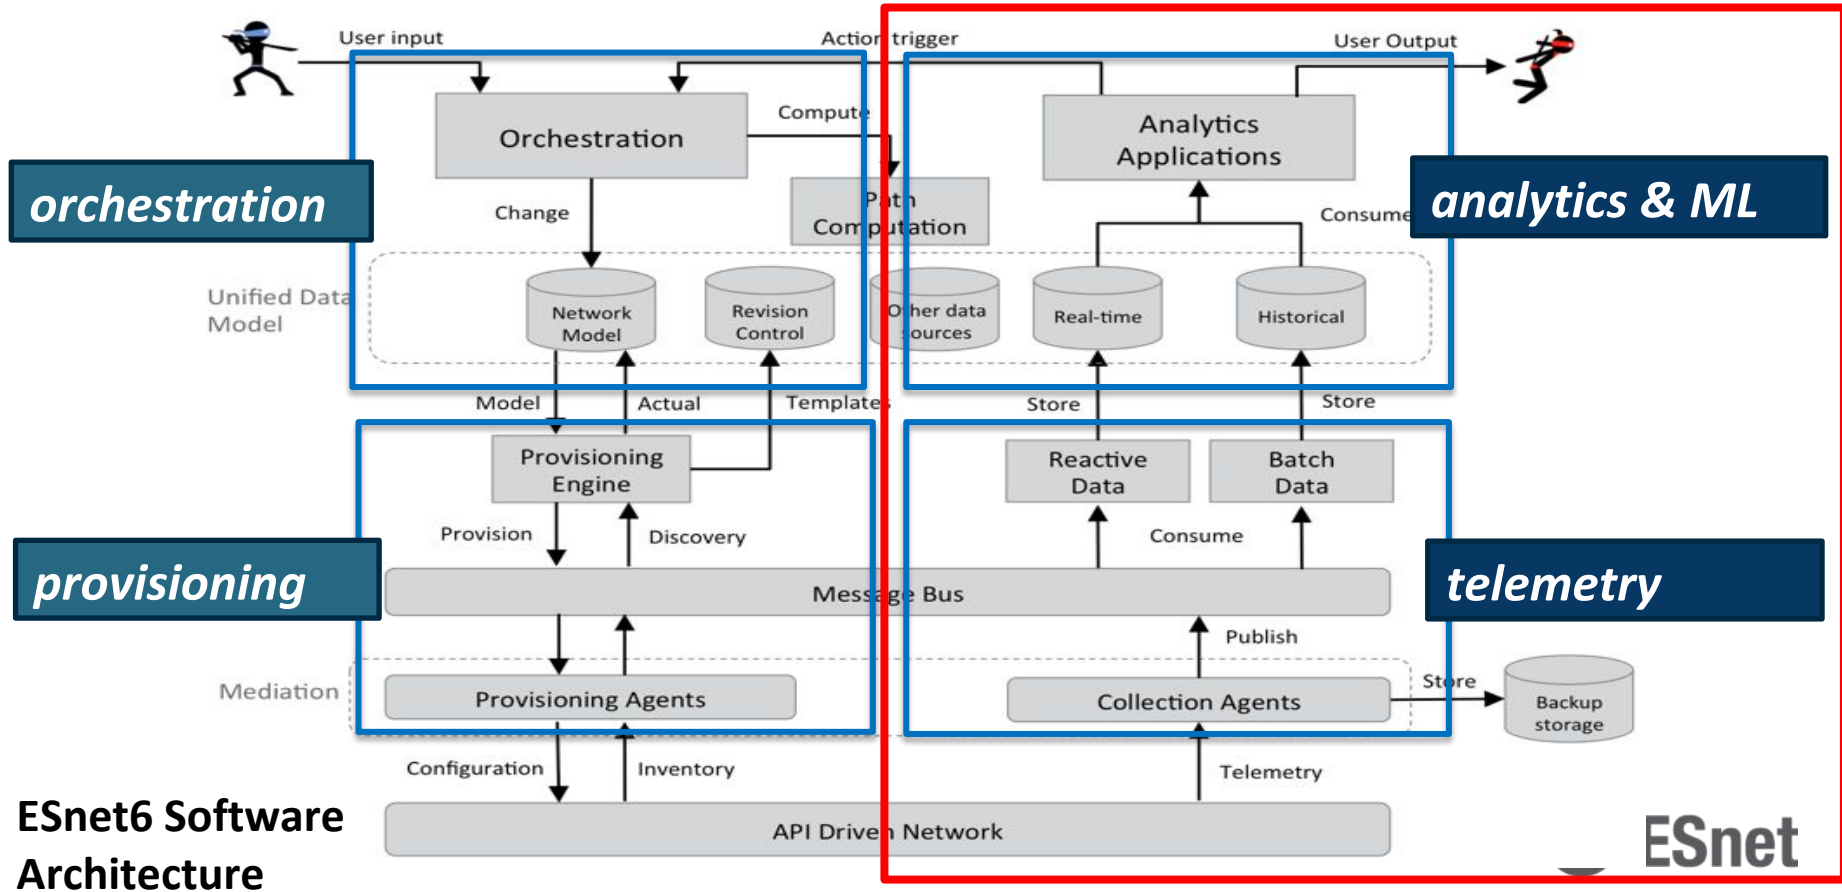
Globus Transfer requests over last two years; those involving NERSC in red



Vision: We need an **autopilot** for networks



Vision: We need an autopilot for networks



Telemetry - what do we get from the network?

- Network monitoring today:
 - SNMP counters: per-interface - **aggregate, polled**
 - Flow-based: Netflow/IPFIX - **approximate, sampled, delayed**

Current solutions are unable to detect e.g., micro-bursts, retransmissions and fine-grained performance dynamics

Packet-scope

- Software-defined hardware and software
- Programmatically deployable and customizable
- Accurate, precision timing
- Non-delayed, non-sampled (per-packet)

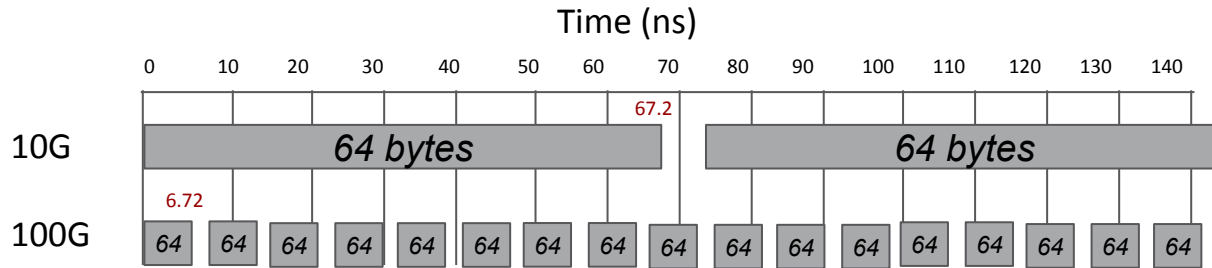
Technology enablers

- Software-Defined Networking
- Programmable network hardware with accurate timestamps (P4)
- High-speed packet processing libraries (XDP, DPDK...)




The Need for High Precision Timing

- Increasing speed of the network links
10G->40G->100G->400G->800G->...



At 100 Gbps, there can be as little as 6.7 nanoseconds between packets that need to be analyzed.

- Goal: going from **microsecond** precision to **nanosecond** precision



“The only reason for time is so that everything doesn't happen at once.”

How do we acquire per-flow telemetry for nx100Gbps flows?

- There are a variety of programmable network devices available today
- Requirements based on ESnet6 network design:
 - 100Gbit/s port speed and roadmap for higher speeds
 - Timing and performance guarantees
 - Easy programming (P4 style)
 - Established vendor
- We have done extensive evaluation of:
 - Netronome SmartNICs (**prototyping platform**)
 - Xilinx FPGAs (**current production platform of choice**)



Barefoot Tofino
(P4)



Xilinx Kintex UltraScale
(FPGA)+

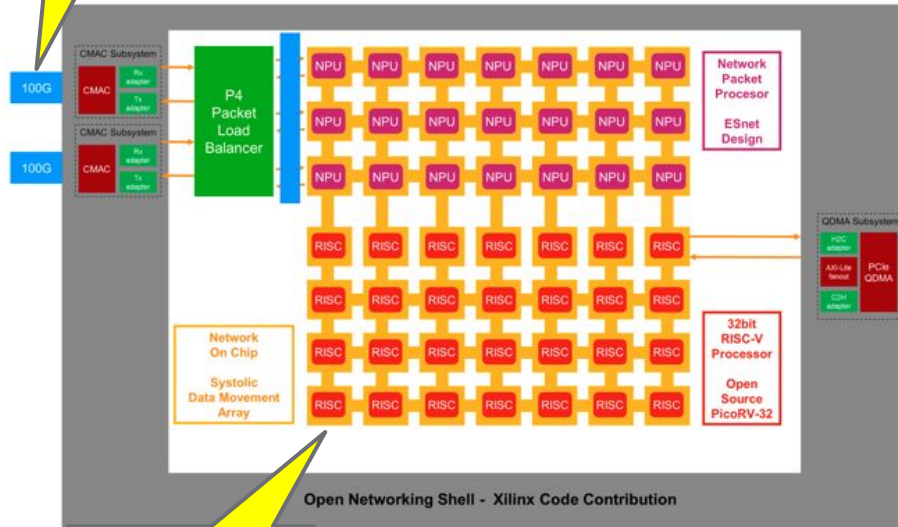


BareExchip NP-4
(NPU)



Specialized FPGA NIC to acquire per-packet telemetry


1.2 Tbps I/O



5 TFLOPS/card

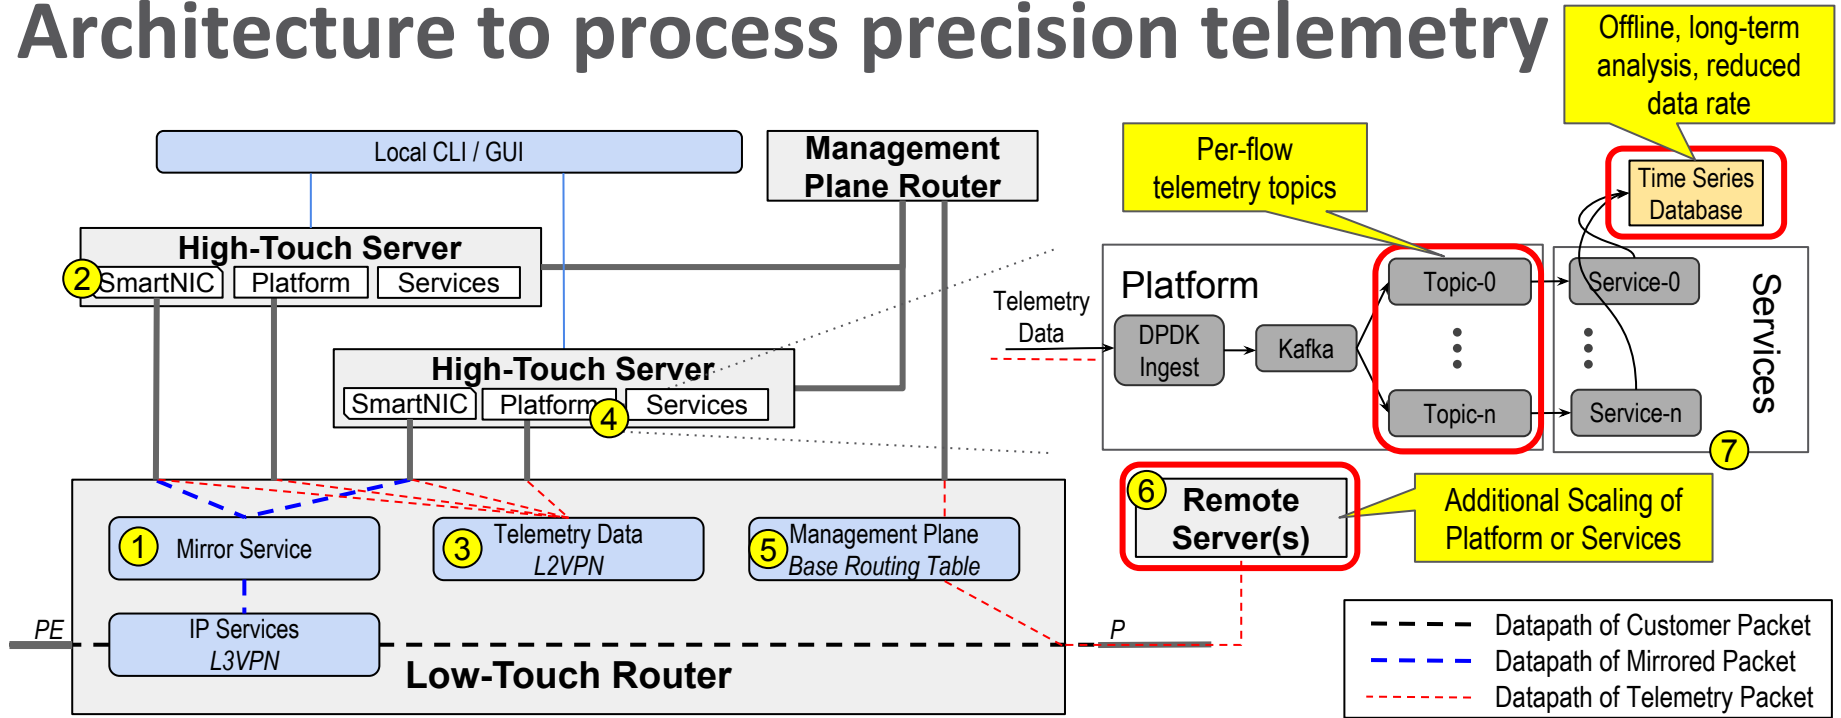
Berkeley eXtensible Processor Array

- Joint development with ESnet / CAG / Xilinx
- Open Source Hardware and SW design
- Nx100G real time processing of DSP / Packet Data



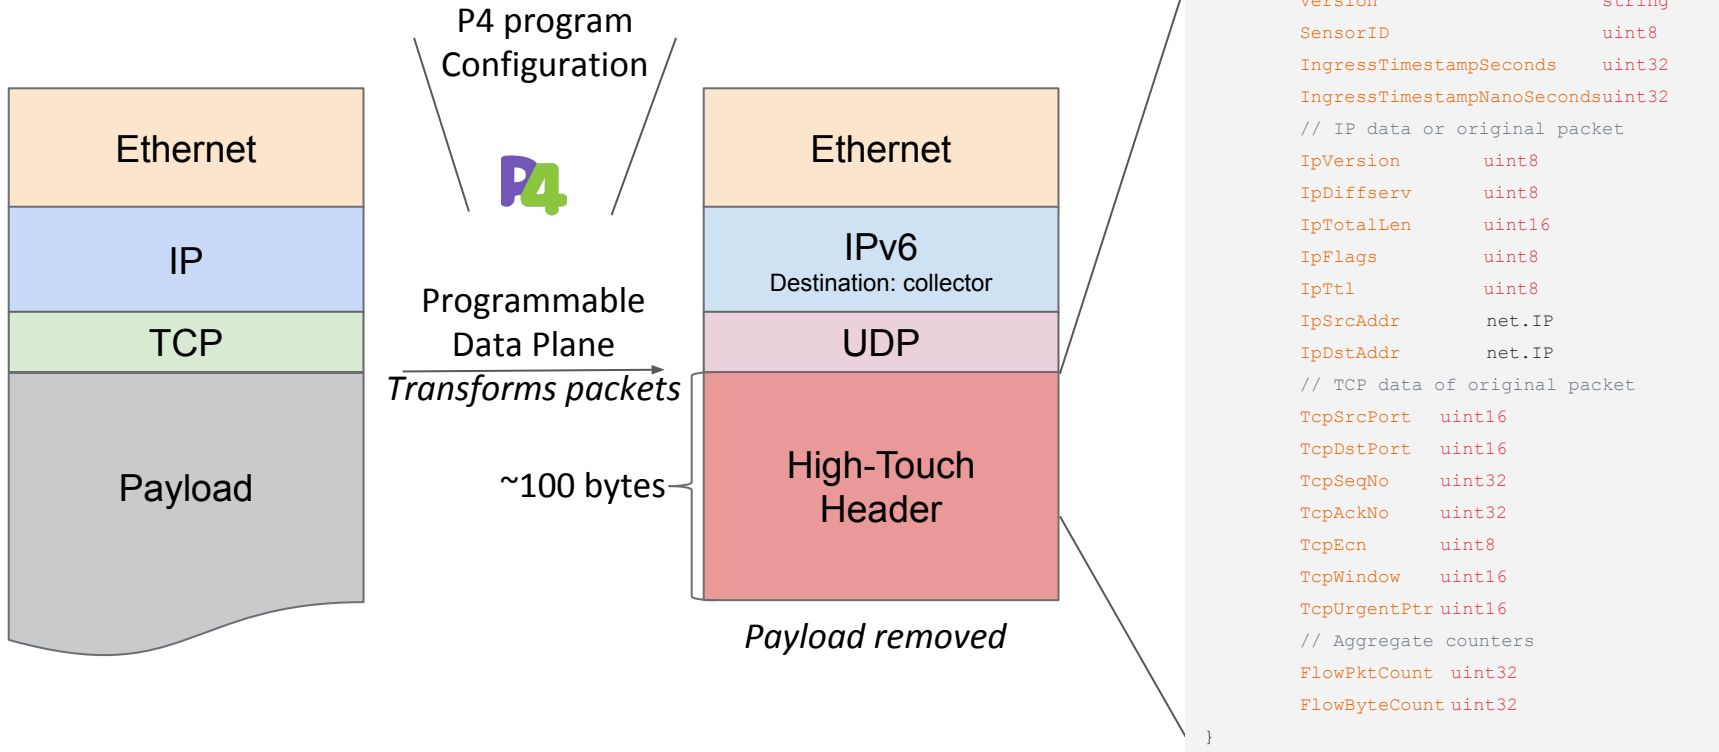
9.6 Tbps I/O
40 TFLOPS
80K DSP Cores
8K RISC-V Cores

Architecture to process precision telemetry



1. Mirror Service - Allows selective flows in the dataplane to be duplicated and sent to the SmartNIC for processing.
2. SmartNIC - Appends meta-data, timestamps and repackages packet for transmission to Platform code.
3. Telemetry Data L2VPN - Provides option to connect SmartNIC and Platform and bypass PCIe bus if needed.
4. Platform - Reads telemetry packets from the network and distributes information to High Touch Services.
5. Management Plane Base Routing Table - Provides connectivity to Remote Servers.
6. Remote Server - Hosts Platform components or Services (but not a SmartNIC). Telemetry data can be directed to Remote Servers.
7. Service - Reads data from the Platform and performs real-time analysis as well as inserts selected telemetry data into database.

Telemetry packets

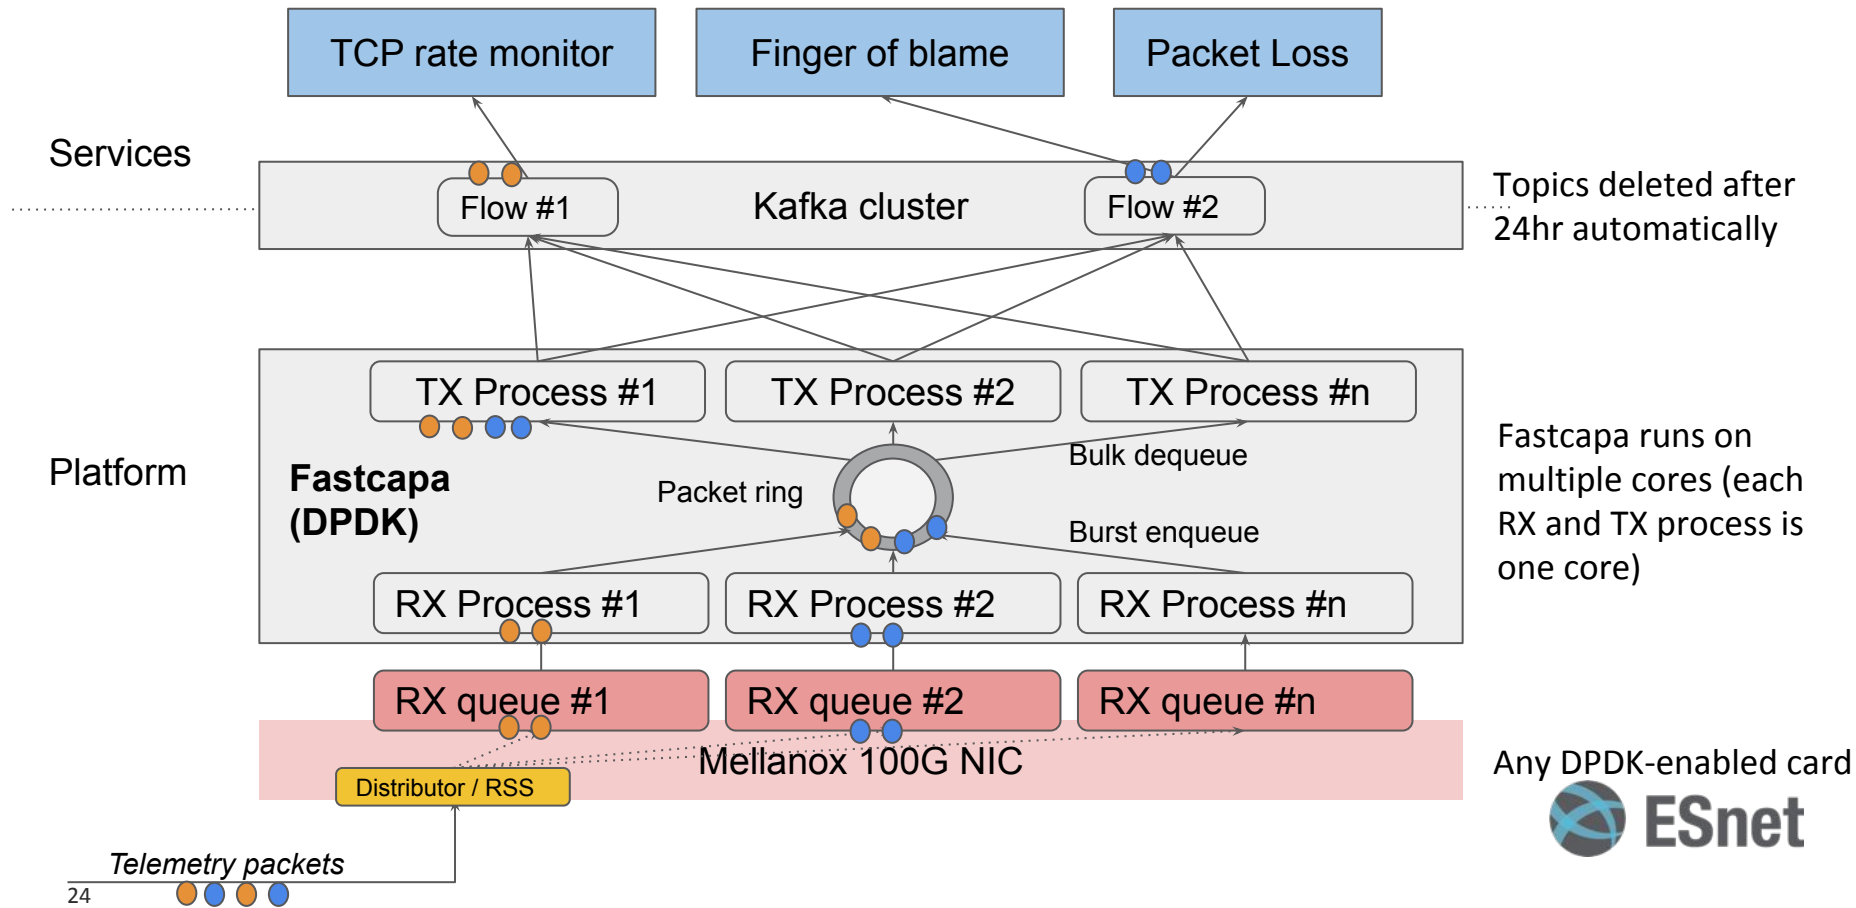


Copy of original packet
of a TCP flow

HighTouch
Telemetry Packet

HighTouch Telemetry
Packet Format v1

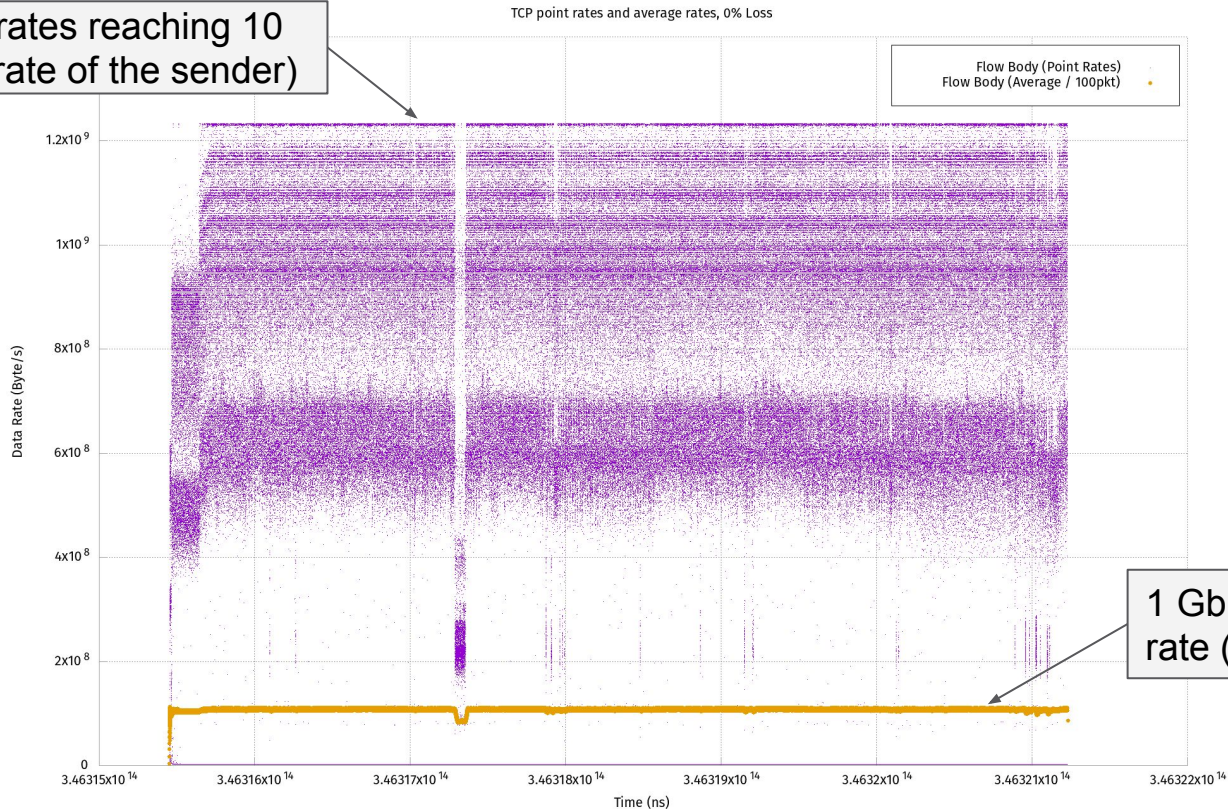
Telemetry as a service: supporting multiple apps



Example: Seeing the invisible bursts


1 Gbps iPerf Flow - 600,000 Packets

Per-packet rates reaching 10 Gbit/s (line rate of the sender)



Expected behavior? - 1% packet drop

Per-packet rates
maximum: 1 Gb/s



5 Mb/s average flow rate
(100 pkt window)


Note: only 23 packets were dropped all together, taking bandwidth down to 5 Mb/s from 1 Gb/s.



New ways to visualize goodput

The difference between the lines represents the performance of the data transfer, convergence is good, divergence is bad.


TCP average throughput and goodput with retransmissions shown, 1% Loss



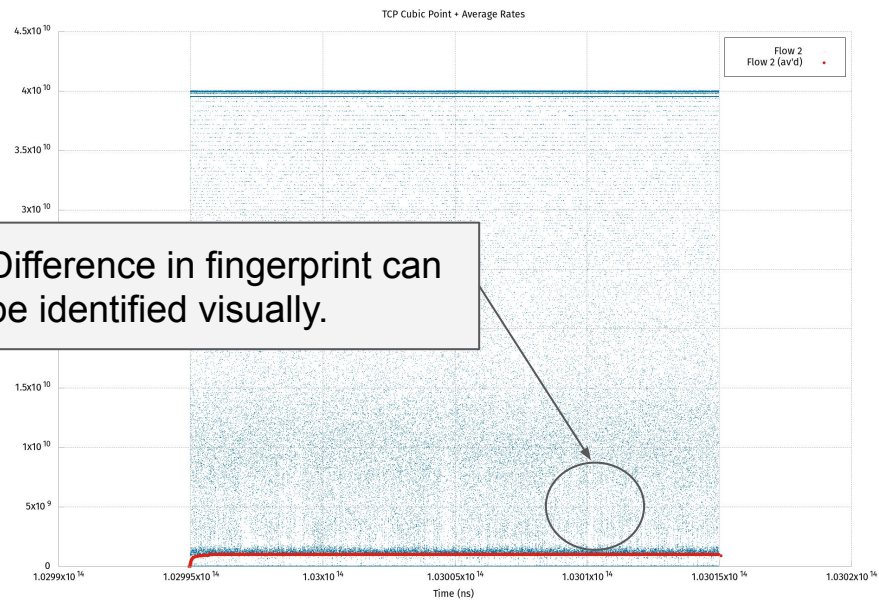
Flow throughput drops when a retransmission happens

Inferring congestion control algorithm behavior

BBR vs Cubic - Point Rates



TCP BBR




TCP Cubic


2 millions of data points shown (around 600.000 points a second generated)



BBR vs Cubic - Inter-Arrival Time Histogram



TCP BBR (delay-based)

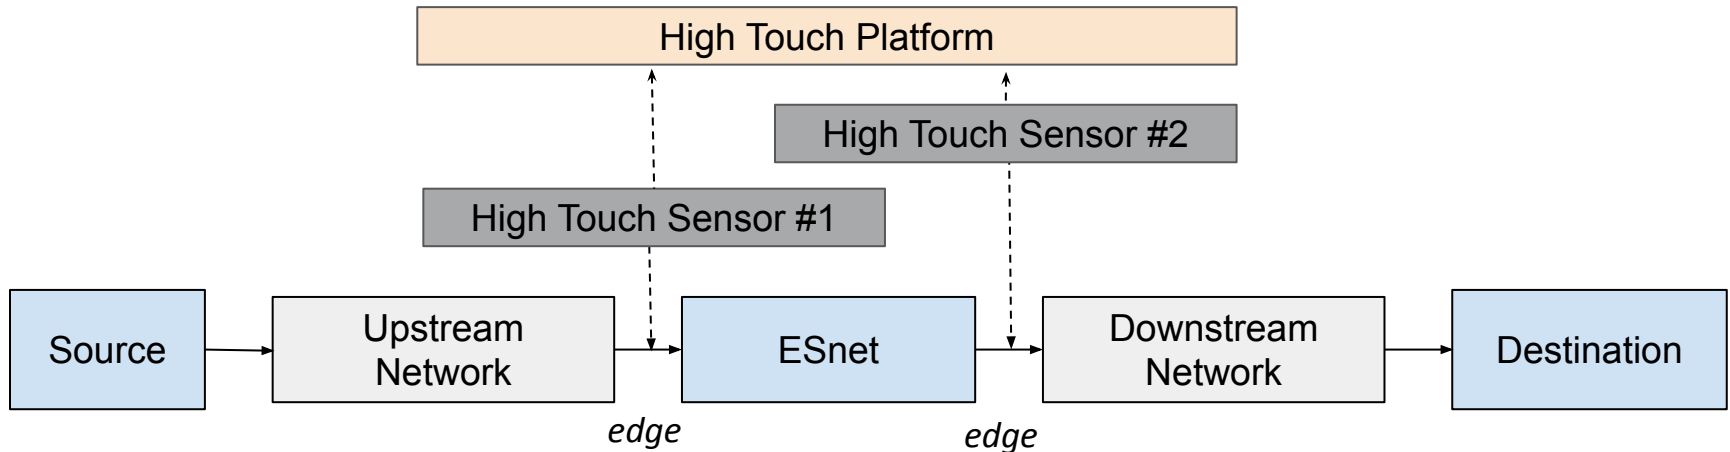


TCP Cubic (loss-based)

Difference in behaviour can be identified visually

BBR: inter-packet timing is more widespread than other congestion control algorithms.

Potential application: Finger of Blame



Upstream Network is experiencing packet loss:

- Both Sensor #1 and Sensor #2 see missing SEQ numbers (non-continuous stream has been observed)

ESnet is experiencing packet loss:

- No losses seen at Sensor #1, missing seq numbers at Sensor #2

Downstream Network is dropping:

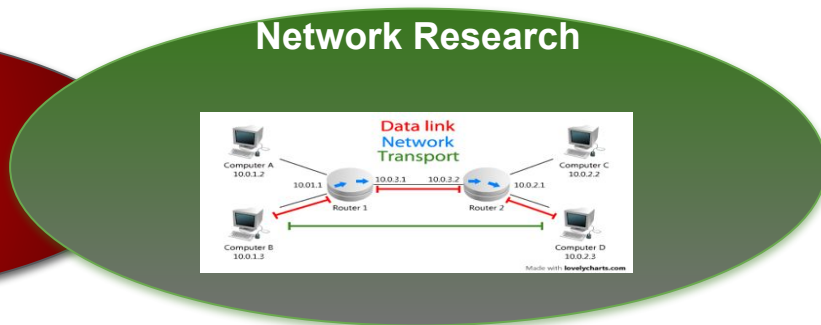
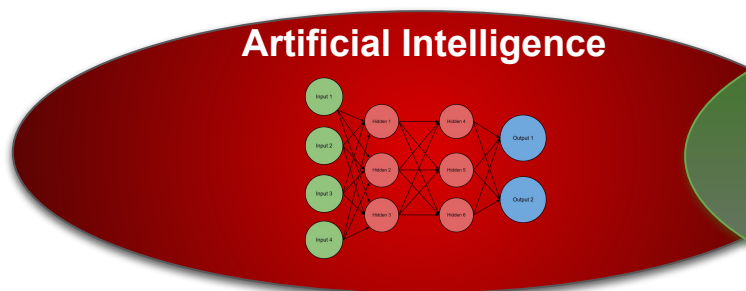
- Both Sensor #1 and Sensor #2 see repeated SEQ numbers





Creation of packet-scope (packet telescope!)


Analytics and ML: Analyzing the telemetry to create the autopilot



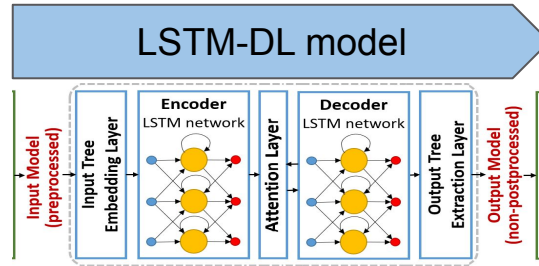

Two research themes explored by applying Deep Learning techniques

- Predicting congestion before scheduling large transfers
- Deploying Self-learning controllers

Case 1: Congestion-free transfers



LIVE Less busy than usual




Planning your next transfer?



- Real-time Data
 - perfSONAR (Loss, Throughput)
 - Traffic: SNMP data
 - Flow behavior: Netflow log



Predicting future x hours ($x=24$)



- Calculating “just-in-time paths” based on current traffic patterns
 - Look at congestion and utilization
 - Schedule maintenance
 - SW/HW upgrades
 - Opportunity for raising alarms
 - If ‘real’ is too high than ‘predicted’

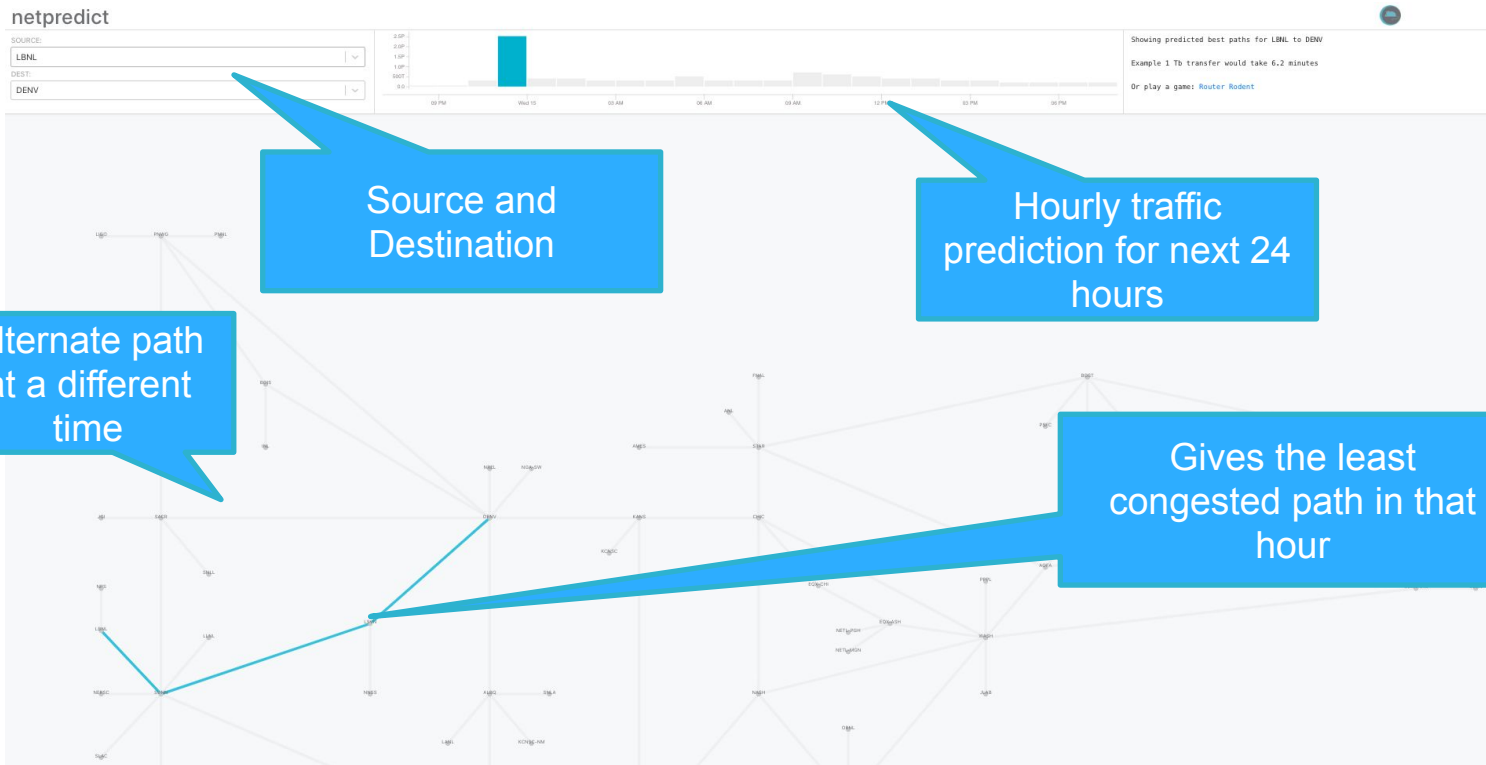
Tool Produced: NetPredict

(Linked with ESnet portal)

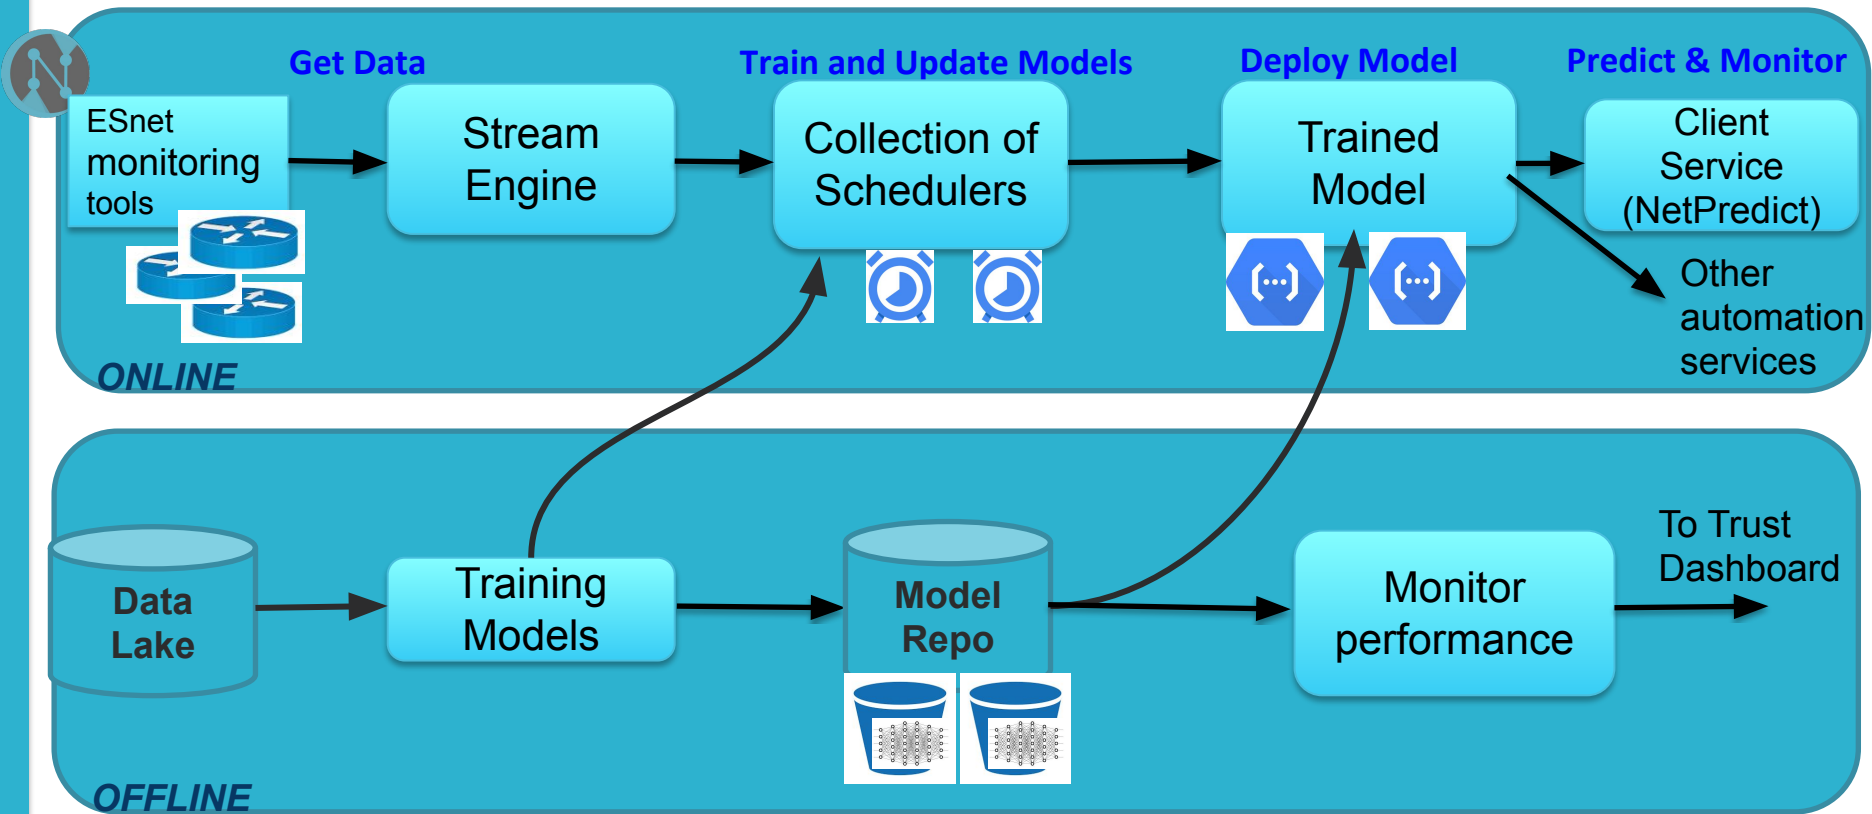
- Deployed on Google Cloud Platform
 - Different models can run at the same time to compute least congested paths
 - Estimates transfer completion time
- Trust dashboard
 - Real-time ML performance
 - Build engineer's confidence in predictions



NetPredict: Google Map for ESnet



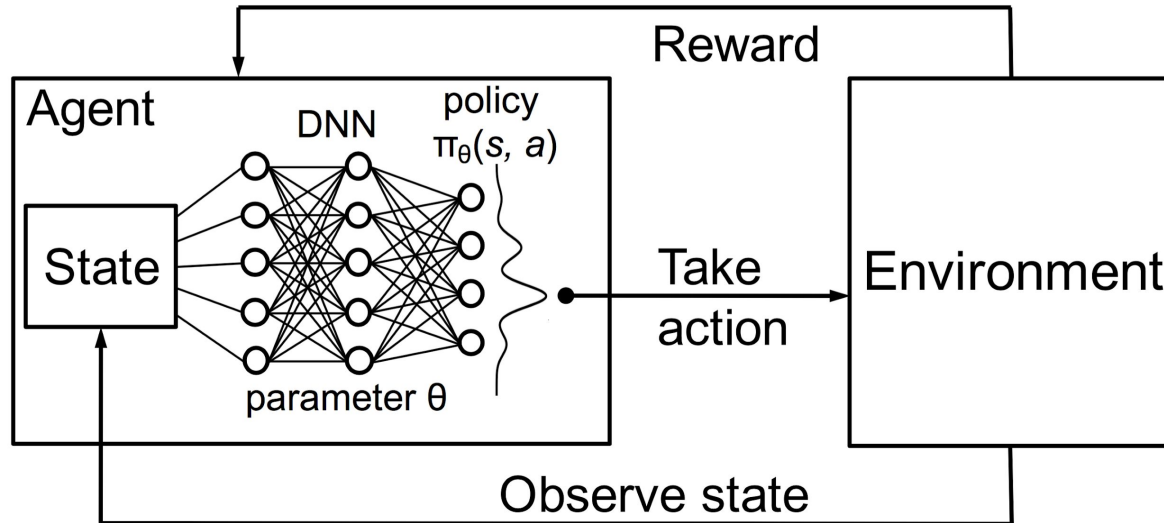
NetPredict System Architecture



Case 2: Self-Learning Network Controller

Using Episodic Q-Learning approach: $Q_t(s, a) = Q_{t-1}(s, a) + \alpha(R(s, a) + \gamma \max_{a'} Q(s', a') - Q_{t-1}(s, a))$

Developing **Model-Free approaches** to allow controller to learn “quickly” and optimal performance



Testing in Emulators and Cloud Testbed

The screenshot displays a network emulator interface with a central network topology diagram and several terminal windows. The topology diagram shows nodes labeled 'c0', 'pngw4', 'elpa3', 'sacr1', 'bois6', 'al bq7', 'sunn5', 'lsvn8', 'denv2', 'kans', and 'lbnl'. Solid blue lines represent physical connections, while dashed red lines represent logical connections. The terminal windows show network traffic logs and command outputs.

```
controller@controller-Virt... packet in 5 76:b9:5d:2a:9b:73 8a:99:cb:fd:e8:f9 4
controller@controller-Virt... packet in 5 8a:99:cb:fd:e8:f9 76:b9:5d:2a:9b:73 1
controller@controller-Virt... packet in 2 8a:99:cb:fd:e8:f9 76:b9:5d:2a:9b:73 5
controller@controller-Virt... packet in 5 8a:99:cb:fd:e8:f9 76:b9:5d:2a:9b:73 1
controller@controller-Virt... packet in 1 8a:99:cb:fd:e8:f9 76:b9:5d:2a:9b:73 3
controller@controller-Virt... packet in 3 8a:99:cb:fd:e8:f9 76:b9:5d:2a:9b:73 2
controller@controller-Virt... packet in 2 8a:99:cb:fd:e8:f9 76:b9:5d:2a:9b:73 5
controller@controller-Virt... packet in 7 8a:99:cb:fd:e8:f9 76:b9:5d:2a:9b:73 1
controller@controller-Virt... packet in 6 8a:99:cb:fd:e8:f9 76:b9:5d:2a:9b:73 1
controller@controller-Virt... packet in 2 76:b9:5d:2a:9b:73 8a:99:cb:fd:e8:f9 1
controller@controller-Virt... packet in 4 8a:99:cb:fd:e8:f9 76:b9:5d:2a:9b:73 2
controller@controller-Virt... packet in 5 76:b9:5d:2a:9b:73 8a:99:cb:fd:e8:f9 4
controller@controller-Virt... packet in 5 8a:99:cb:fd:e8:f9 76:b9:5d:2a:9b:73 1
controller@controller-Virt... packet in 2 8a:99:cb:fd:e8:f9 76:b9:5d:2a:9b:73 5
```

```
root@controller-VirtualBox: ~/Desktop/MiniNEM/Examples/ROUTING#
Host: lbnl!
64 bytes from 10.0.0.1: icmp_seq=77 ttl=64 time=0.067 ms
64 bytes from 10.0.0.1: icmp_seq=78 ttl=64 time=0.061 ms
64 bytes from 10.0.0.1: icmp_seq=79 ttl=64 time=0.105 ms
64 bytes from 10.0.0.1: icmp_seq=80 ttl=64 time=0.062 ms
64 bytes from 10.0.0.1: icmp_seq=81 ttl=64 time=0.092 ms
64 bytes from 10.0.0.1: icmp_seq=82 ttl=64 time=0.113 ms
64 bytes from 10.0.0.1: icmp_seq=83 ttl=64 time=0.060 ms
64 bytes from 10.0.0.1: icmp_seq=84 ttl=64 time=0.064 ms
64 bytes from 10.0.0.1: icmp_seq=85 ttl=64 time=0.060 ms
64 bytes from 10.0.0.1: icmp_seq=86 ttl=64 time=0.068 ms
64 bytes from 10.0.0.1: icmp_seq=87 ttl=64 time=0.057 ms
64 bytes from 10.0.0.1: icmp_seq=88 ttl=64 time=0.068 ms
64 bytes from 10.0.0.1: icmp_seq=89 ttl=64 time=0.062 ms
64 bytes from 10.0.0.1: icmp_seq=90 ttl=64 time=0.061 ms
64 bytes from 10.0.0.1: icmp_seq=91 ttl=64 time=0.065 ms
64 bytes from 10.0.0.1: icmp_seq=92 ttl=64 time=0.065 ms
64 bytes from 10.0.0.1: icmp_seq=93 ttl=64 time=0.060 ms
64 bytes from 10.0.0.1: icmp_seq=94 ttl=64 time=0.208 ms
64 bytes from 10.0.0.1: icmp_seq=95 ttl=64 time=0.098 ms
```

```
controller@controller-VirtualBox: ~/Music/deeproue-gym/tests
File Edit View Search Terminal Help
Taking predicted action 1
Taking predicted action 0
Taking predicted action 1
Taking predicted action 0
Taking predicted action 2
Taking predicted action 0
Taking predicted action 1
Taking predicted action 2
Taking predicted action 0
Taking random action 2
Taking predicted action 2
Taking predicted action 2
Taking predicted action 2
Taking predicted action 0
Taking predicted action 2
Taking predicted action 2
Taking predicted action 2
Taking predicted action 2
Taking predicted action 2
```

How do you now test these new ideas at scale?





NSF'S 10 BIG IDEAS



National Science Foundation
WHERE DISCOVERIES BEGIN



Genesis

Big Idea **Mid-scale Research Infrastructure**




- Many important potential experiments and facilities fall between the \$100K to \$4M¹ Major Research Instrumentation (MRI) program and the > \$70M Major Research Equipment and Facilities Construction (MREFC) account.
- This gap results in missed opportunities that may leave essential science undone.
- NSF needs a new agile process for funding experimental research capabilities in the mid-scale range.

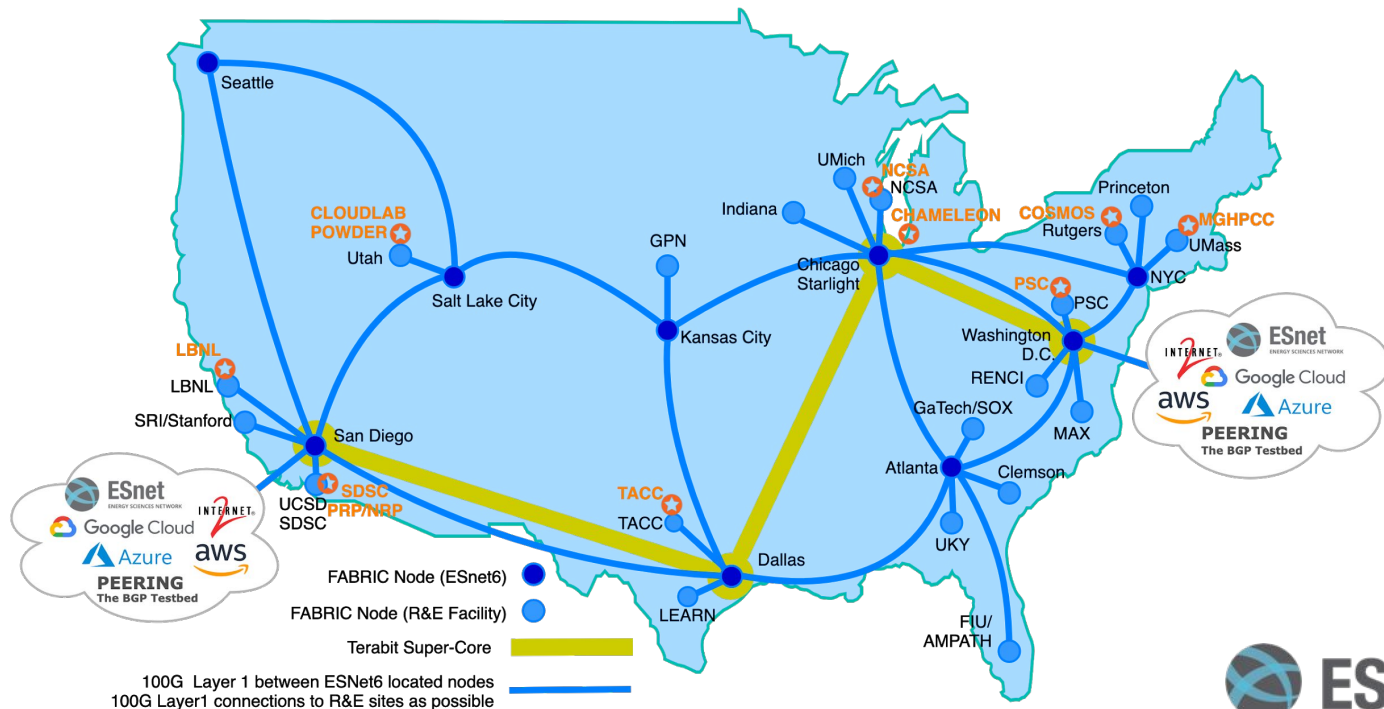
FABRIC is Adaptive Programmable Research Infrastructure for Computer Science and Science Applications - NSF MSRI project



This work is funded by NSF grant CNS-1935966



FABRIC



100G Layer 1 between ESNet6 located nodes
100G Layer1 connections to R&E sites as possible


For more details: whatisfabric.net



Why FABRIC?

- The mantra of the last 20 years – ‘Internet is showing its age.’
 - Applications designed around discrete points in the solution space
 - Inability to program the core of the network
- What changed?
 - Cheap compute/storage that can be put *directly in* the network
 - Multiple established methods of programmability (OpenFlow, P4, eBPF, DPDK, BGP flowspec)
 - Advances in Machine Learning/AI
 - Emergence of 5G, IoT, various flavors of cloud technologies
- Opportunity for the community to push the boundaries of distributed, stateful, ‘everywhere’ programmable infrastructure
 - More control *or* dataplane state, or some combination? Multiple architectures (co)exist in this space.
 - Network as a big-data instrument? Autonomous network control?
 - New protocols and applications that program the network?
 - Security as an integral component

FABRIC Node Concept

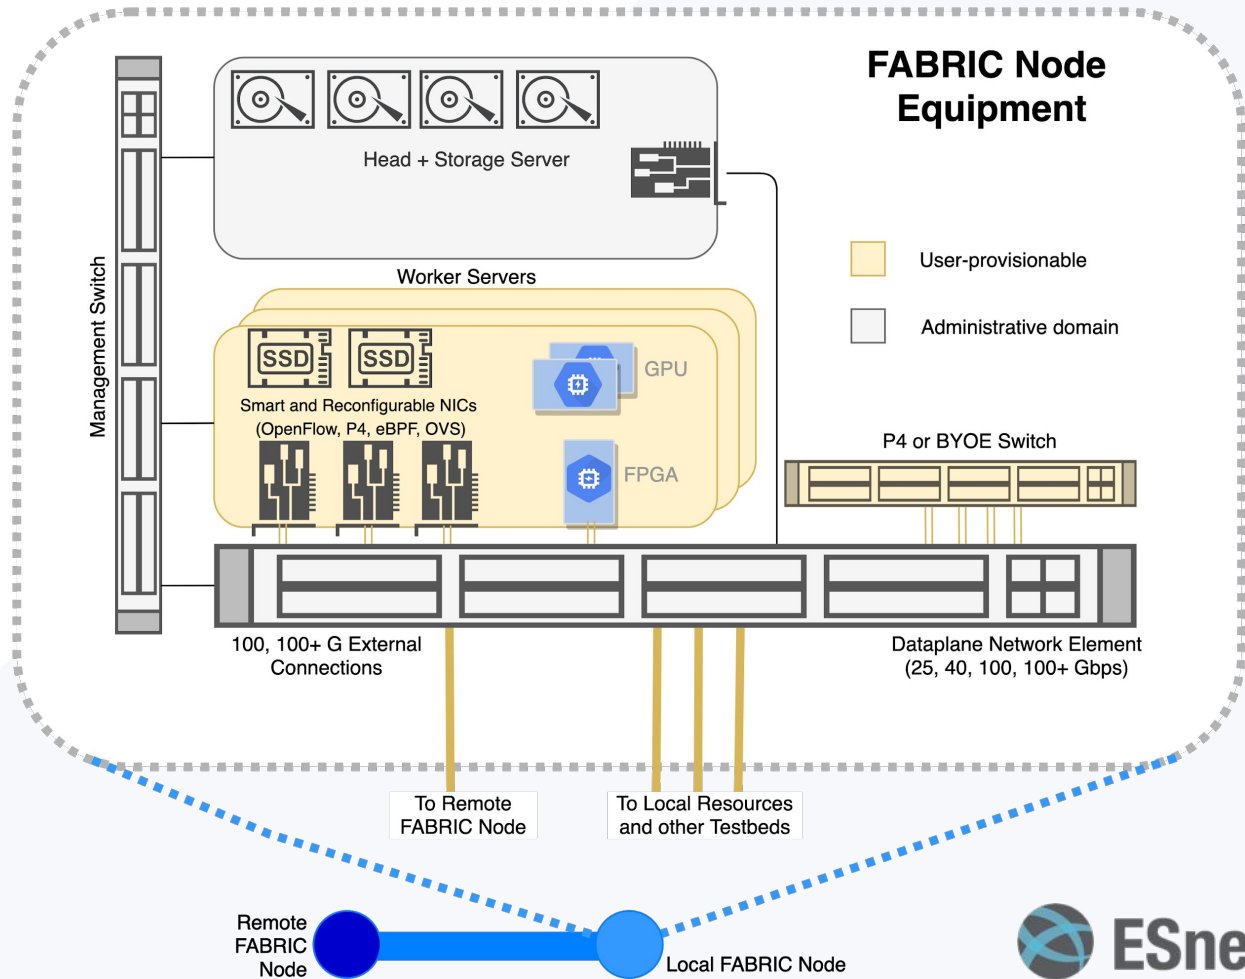


- At scale experimentation
 - 13 core nodes
 - multiple high-capacity *dedicated* optical links.
 - 16 *initial edge* nodes (also known as 'hanks') located on campuses, in lab datacenters to
 - provide base load, serve as gateways for facilities to connect to FABRIC

FABRIC Node ('hank') Design: Network + Compute

- We refer to it also as a 'disaggregated router'
- Network cards with high speed interfaces (25G, 40G, 100G. 200G+ in future)
 - Programmable interface cards (hardware OVS offload + DPDK)
 - Reconfigurable interface cards (FPGA and P4/network processors)
- High-performance servers equipped with
 - GPUs
 - FPGA compute accelerators
 - NVMe drives
 - Storage: User-provisionable short term & shared high volume. Not meant to be persistent.
- All ports interconnected by a 100G+ switch programmable through testbed control software
 - Acts as a 'patch panel' connecting various ports in the node together
- Users can fully interact with network, compute, storage
- Nodes are "sliceable" for experimenters to use simultaneously

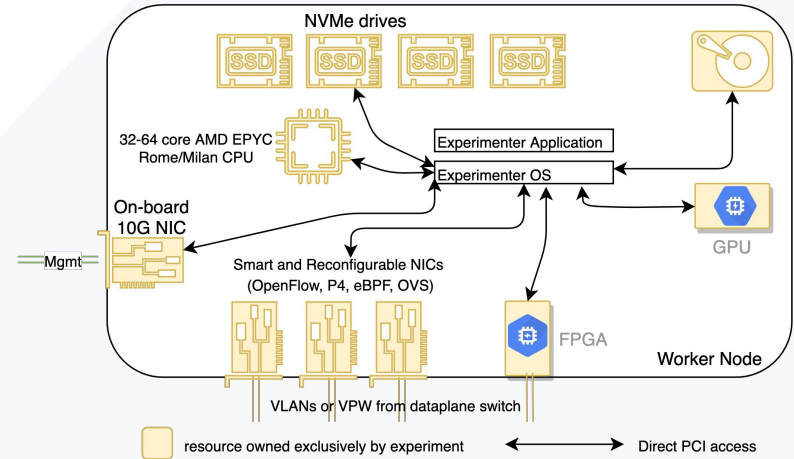
Conceptual FABRIC Node 'Hank' Overview



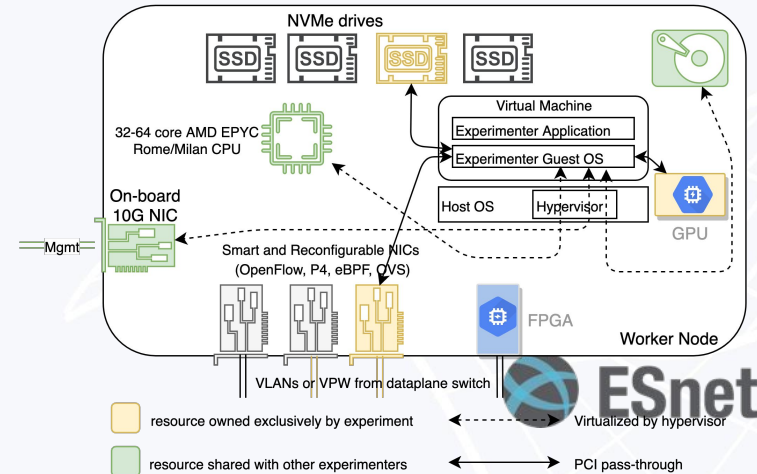
Node Level Programmability Abstractions

- Substantial compute and storage
- Main capabilities are various PCI cards in individual servers
 - NICs, GPUs, FPGAs
- Additional switches and BYOE hardware
- Depending on experimenter request can be provided as part of a bare-metal server or via PCI pass-through for VMs and containers

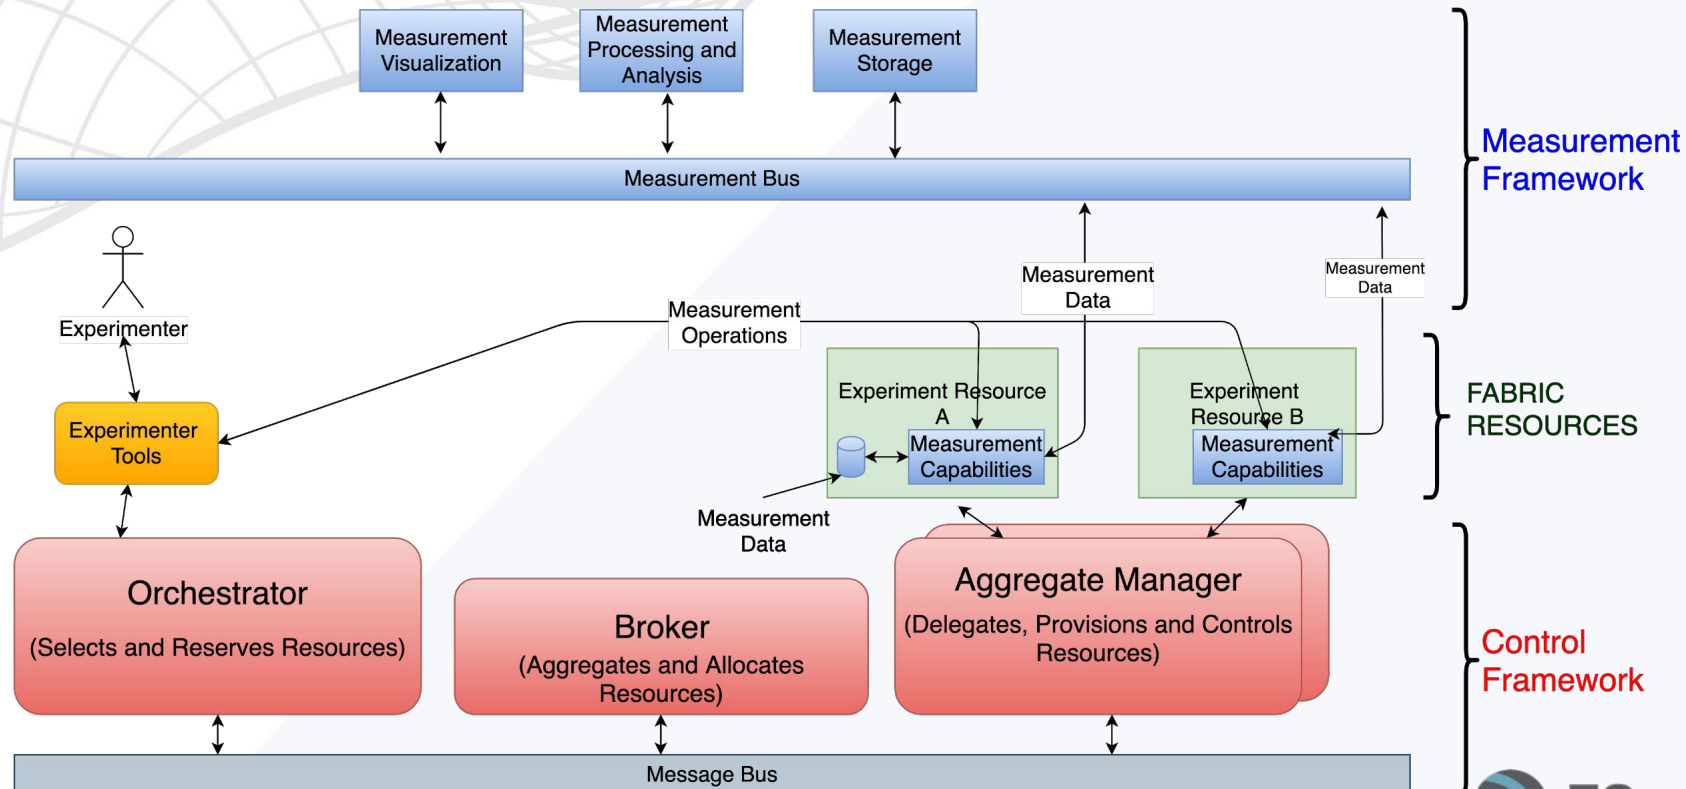
FABRIC experiment using a bare-metal server



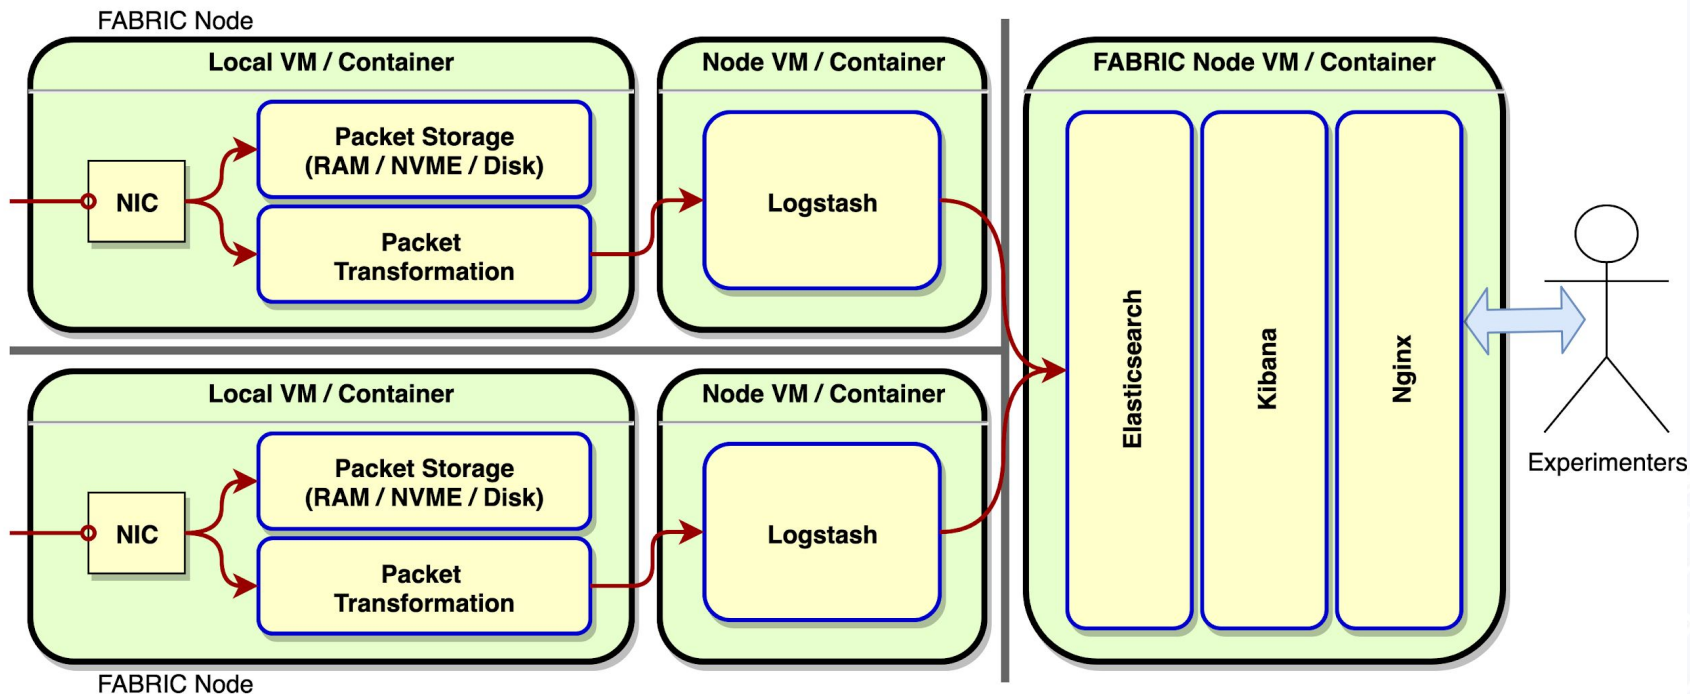
FABRIC experiment using a virtual machine



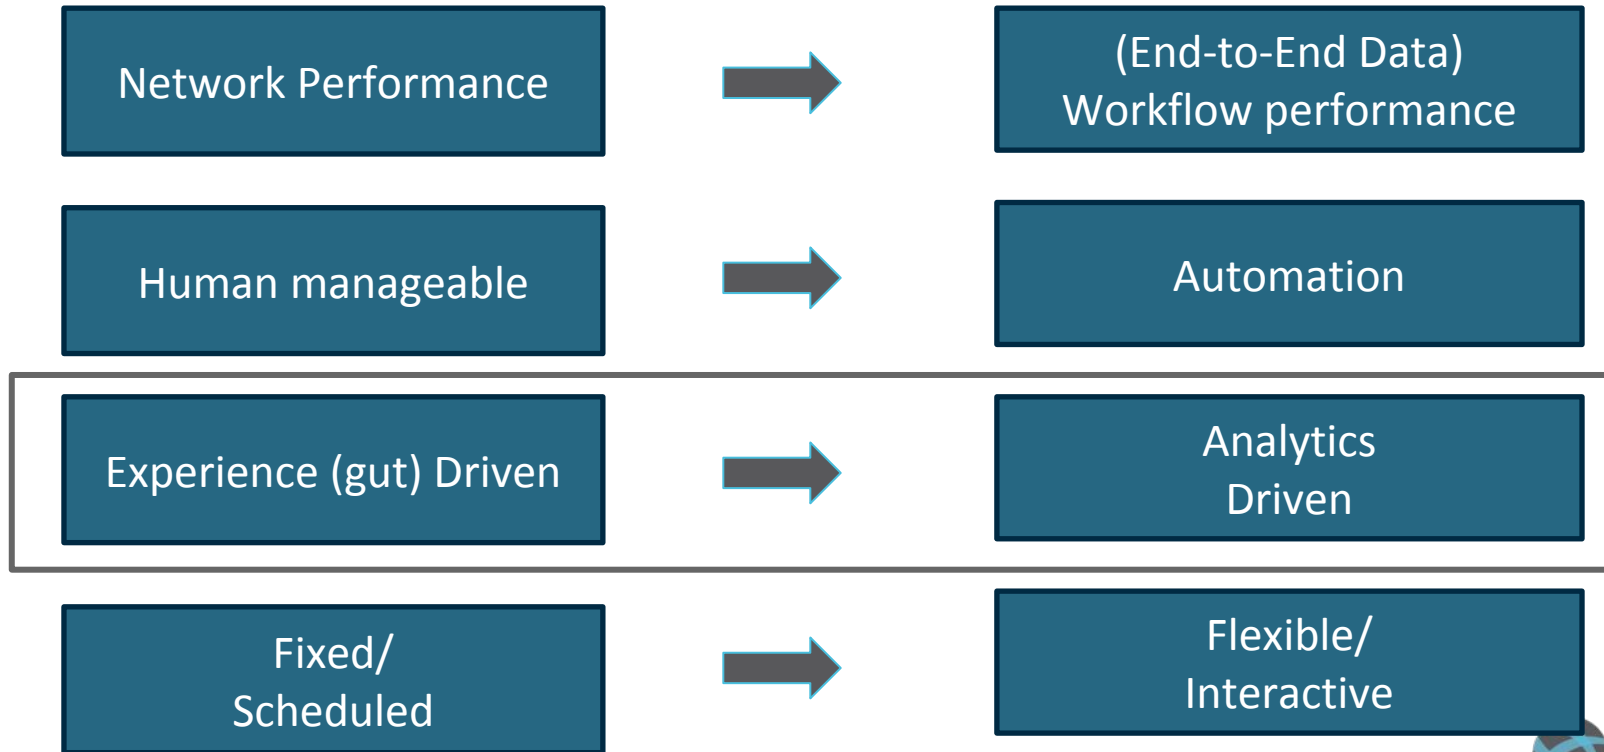
Control and Measurement Framework



FABRIC Measurement Data Processing/Analysis




Need for scale is changing how we manage, build, operate networks



Summary



Thanks to ESnet team and partners



Chin Guok
Yatish Kumar
Richard Cziva
Bruce Mah

ESnet6 - High Touch



Ana Butko
Farzad Fatollahi-Fard
Chin Guok
Yatish Kumar
Gordon Brebner
John Shalf



BX Processor




Mariam Kiran
Bashir Mohammed
Peter Murphy
Nandini Krishnamurthy

AI/ML for networks

FABRIC Leadership Team

Ilya Baldin (RENCI)




Anita Nikolich (IIT)



Inder Monga
(ESnet)



Jim Griffioen (UKY)



KC Wang (Clemson)



Dale Carder (ESnet)



Tom Lehman
(Virnao)



Paul Ruth (RENCI)



Zongming Fei (UKY)

