

# EJFAT

---

ESnet-Jefferson Lab FPGA Accelerated Transport

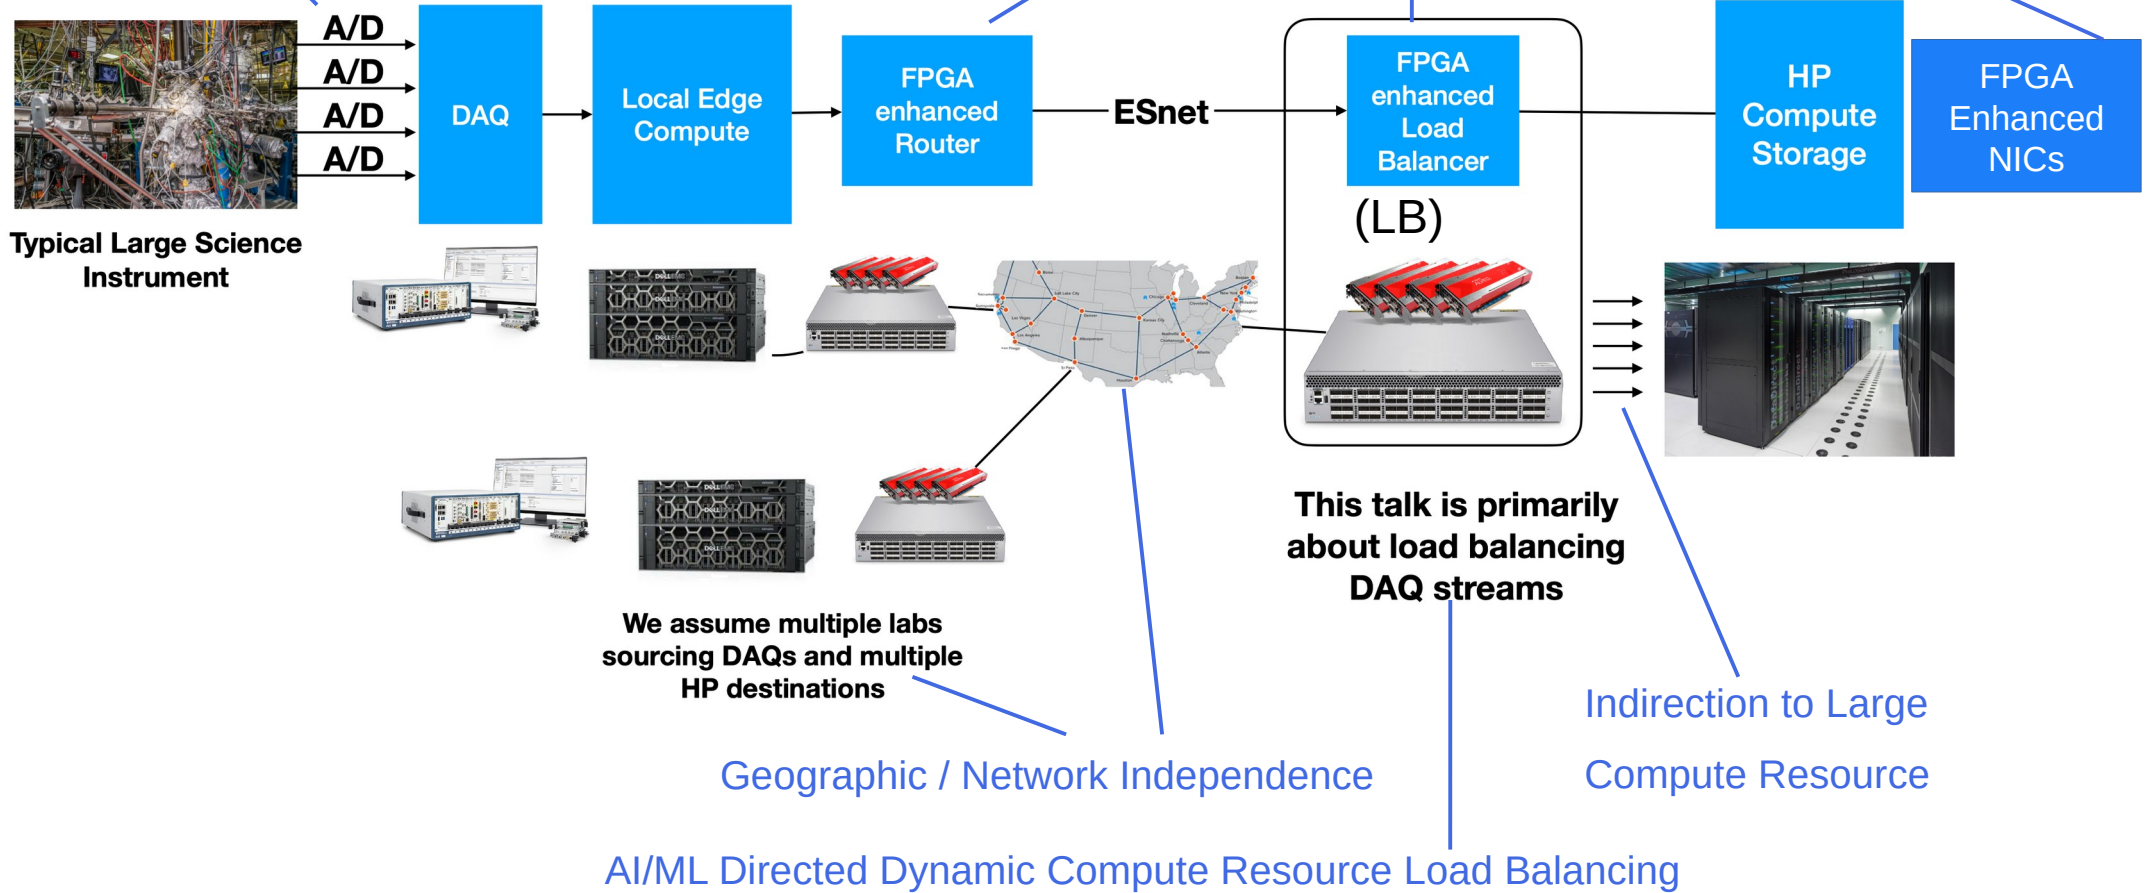
---

Michael Goodrich , Carl Timmer, Vardan Gyurjyan,  
David Lawrence , Graham Hayes ( JLAB )  
Yatish Kumar , Stacey Sheldon (ESnet)

# EJFAT = Edge to Core System Architecture: Workflows Steered by AI/ML

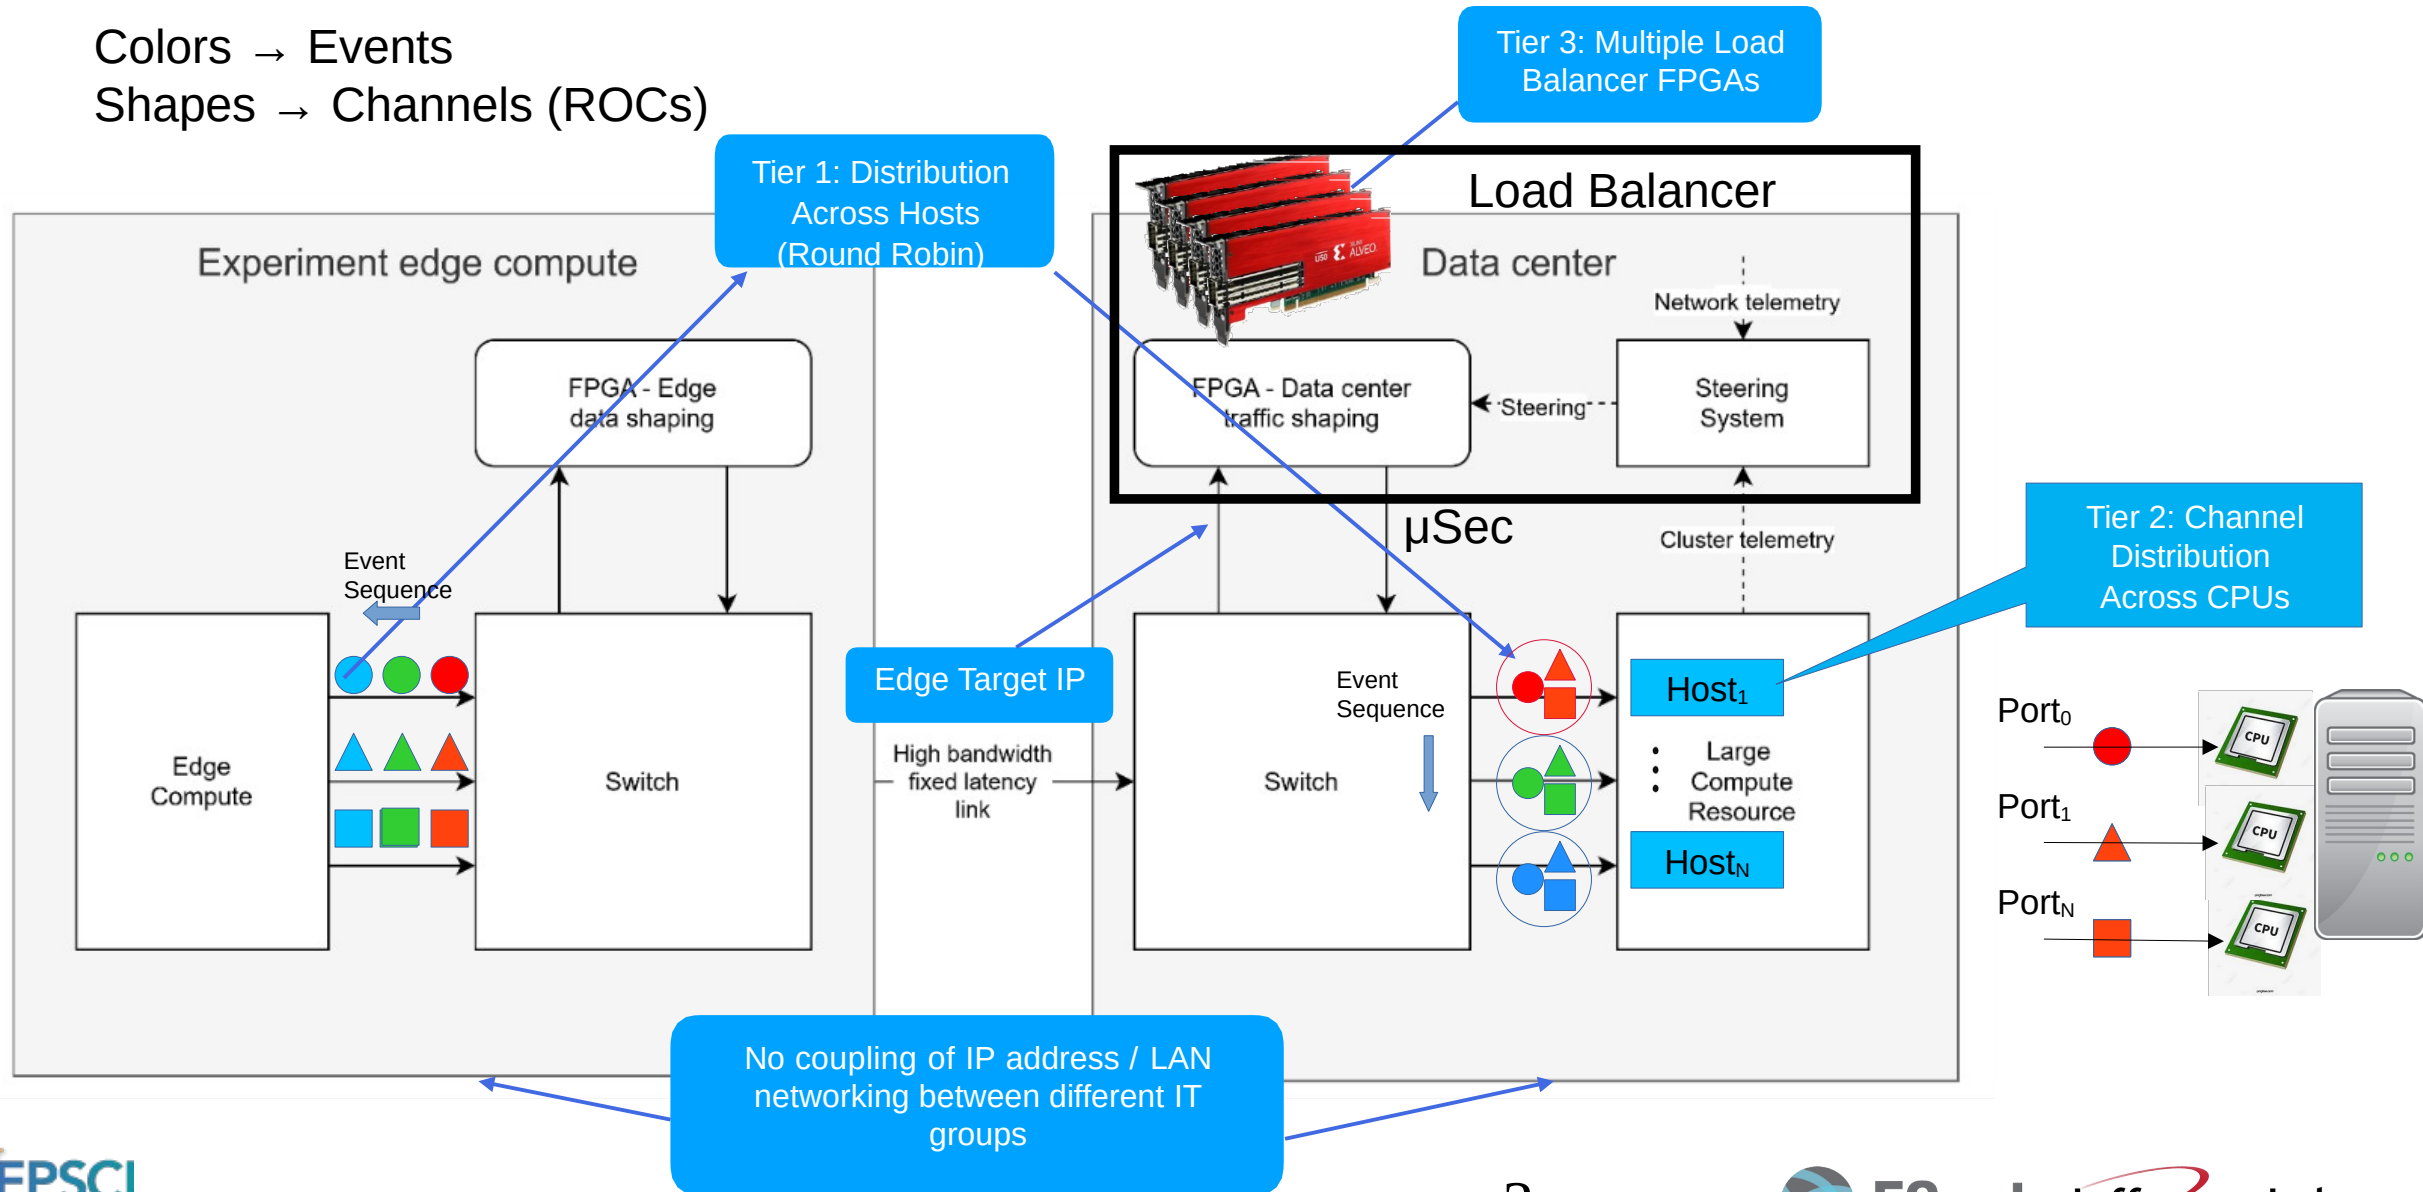
Loss-less UDP Streaming Triggered

Opportunistic FPGA Based Network Acceleration



# Channel Aggregation + Three Tier Horizontal Scaling

Colors → Events  
Shapes → Channels (ROCs)




# LB: Xilinx U280 FPGA (PCIe) + Host

- **Data Plane (DP):** FPGA FW = RTL + P4
  - Packet Filtering, ARP, Ping
  - P4: Data Base for UDP Hdr Rewrites
- **Control Plane (CP):** Host
  - DP DB Maintenance
  - Monitor Network / Core Telemetry
  - AI/ML Steerage / Feedback
    - Upstream: Experiment / DAQ
    - Downstream: Core Computing
    - Core Resource Provisioning




# Load Balancer:

# User Interface





Typical Large Science Instrument

Data Source UDP Packet Payload



Load Balancer meta-data

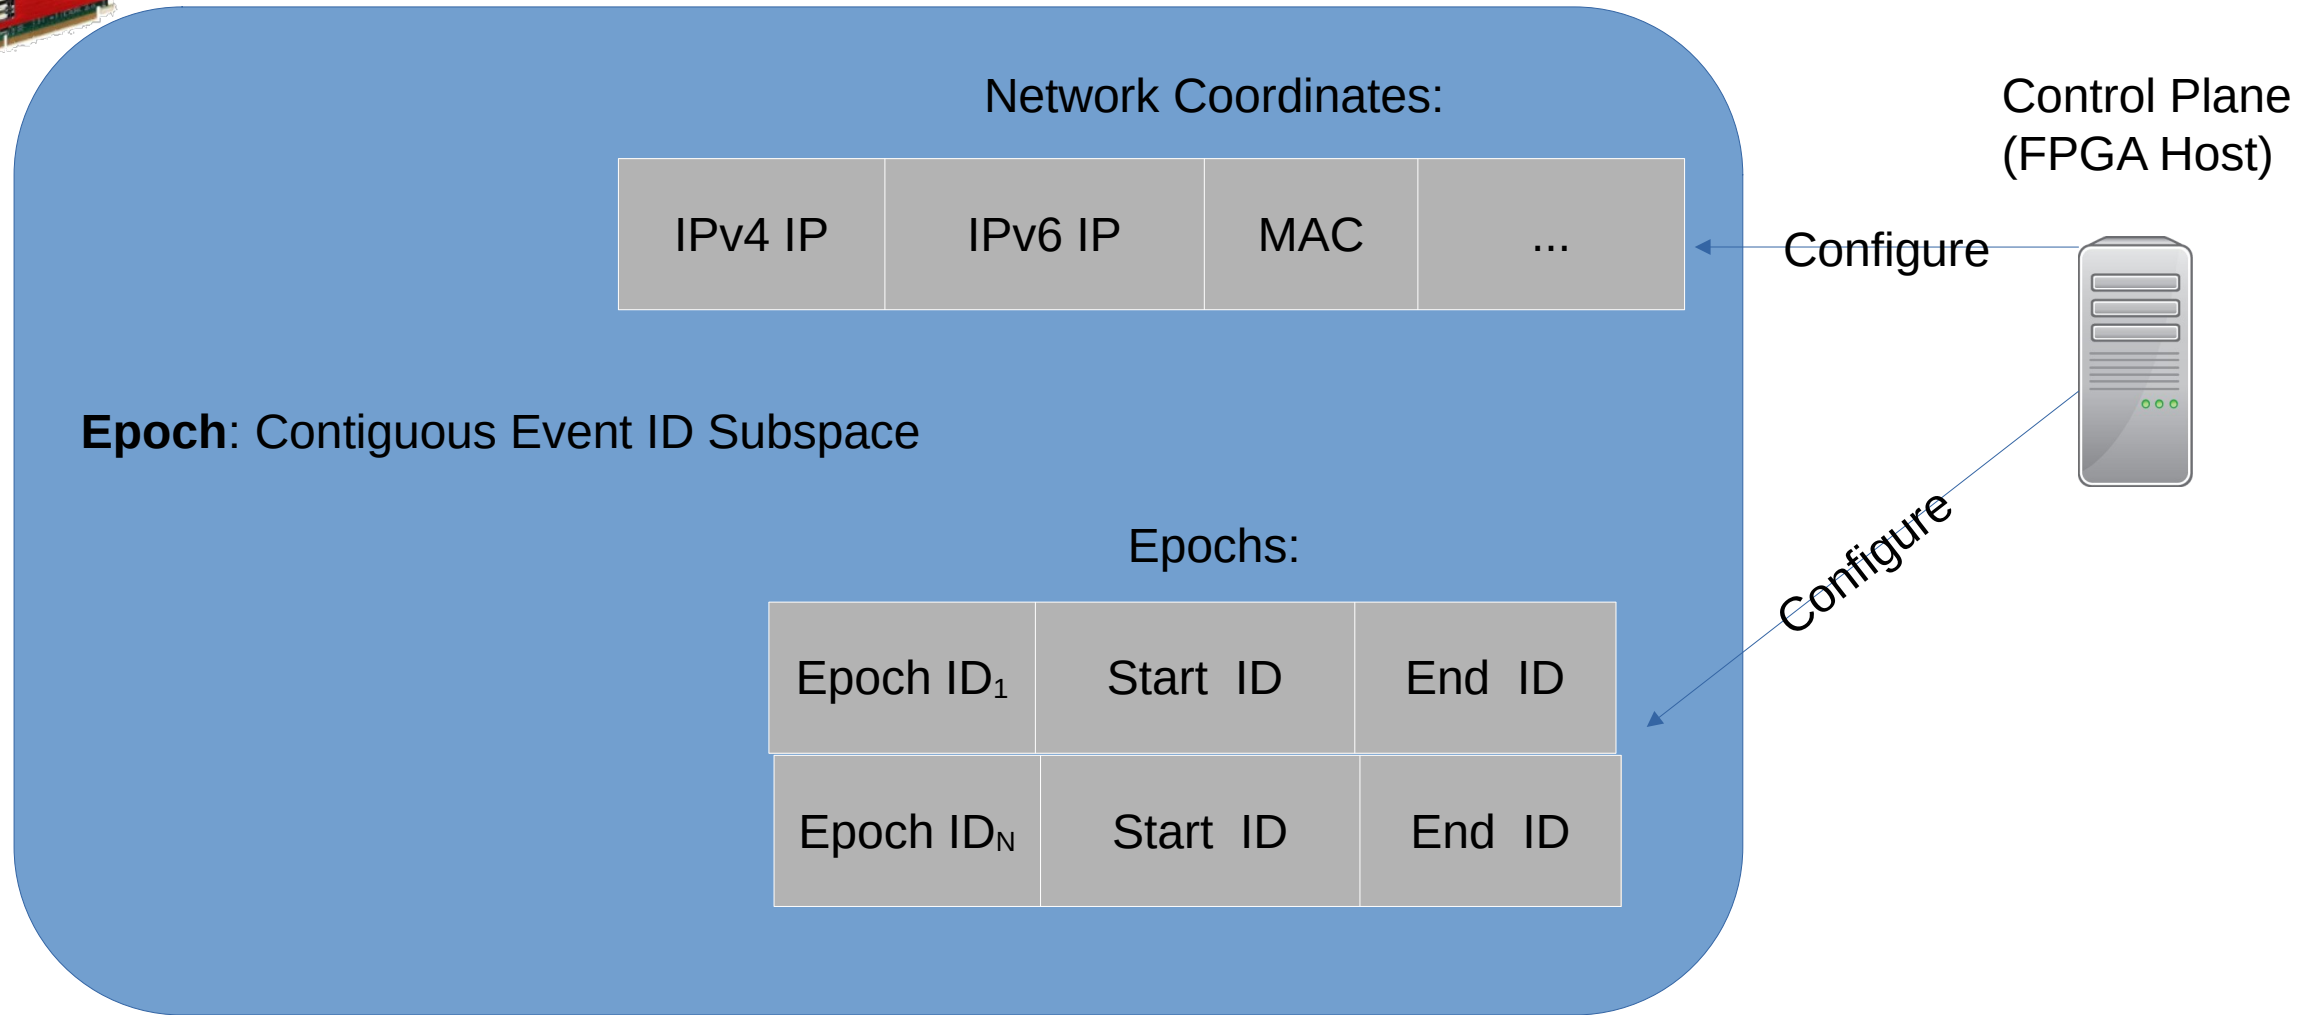


Reassembly meta-data

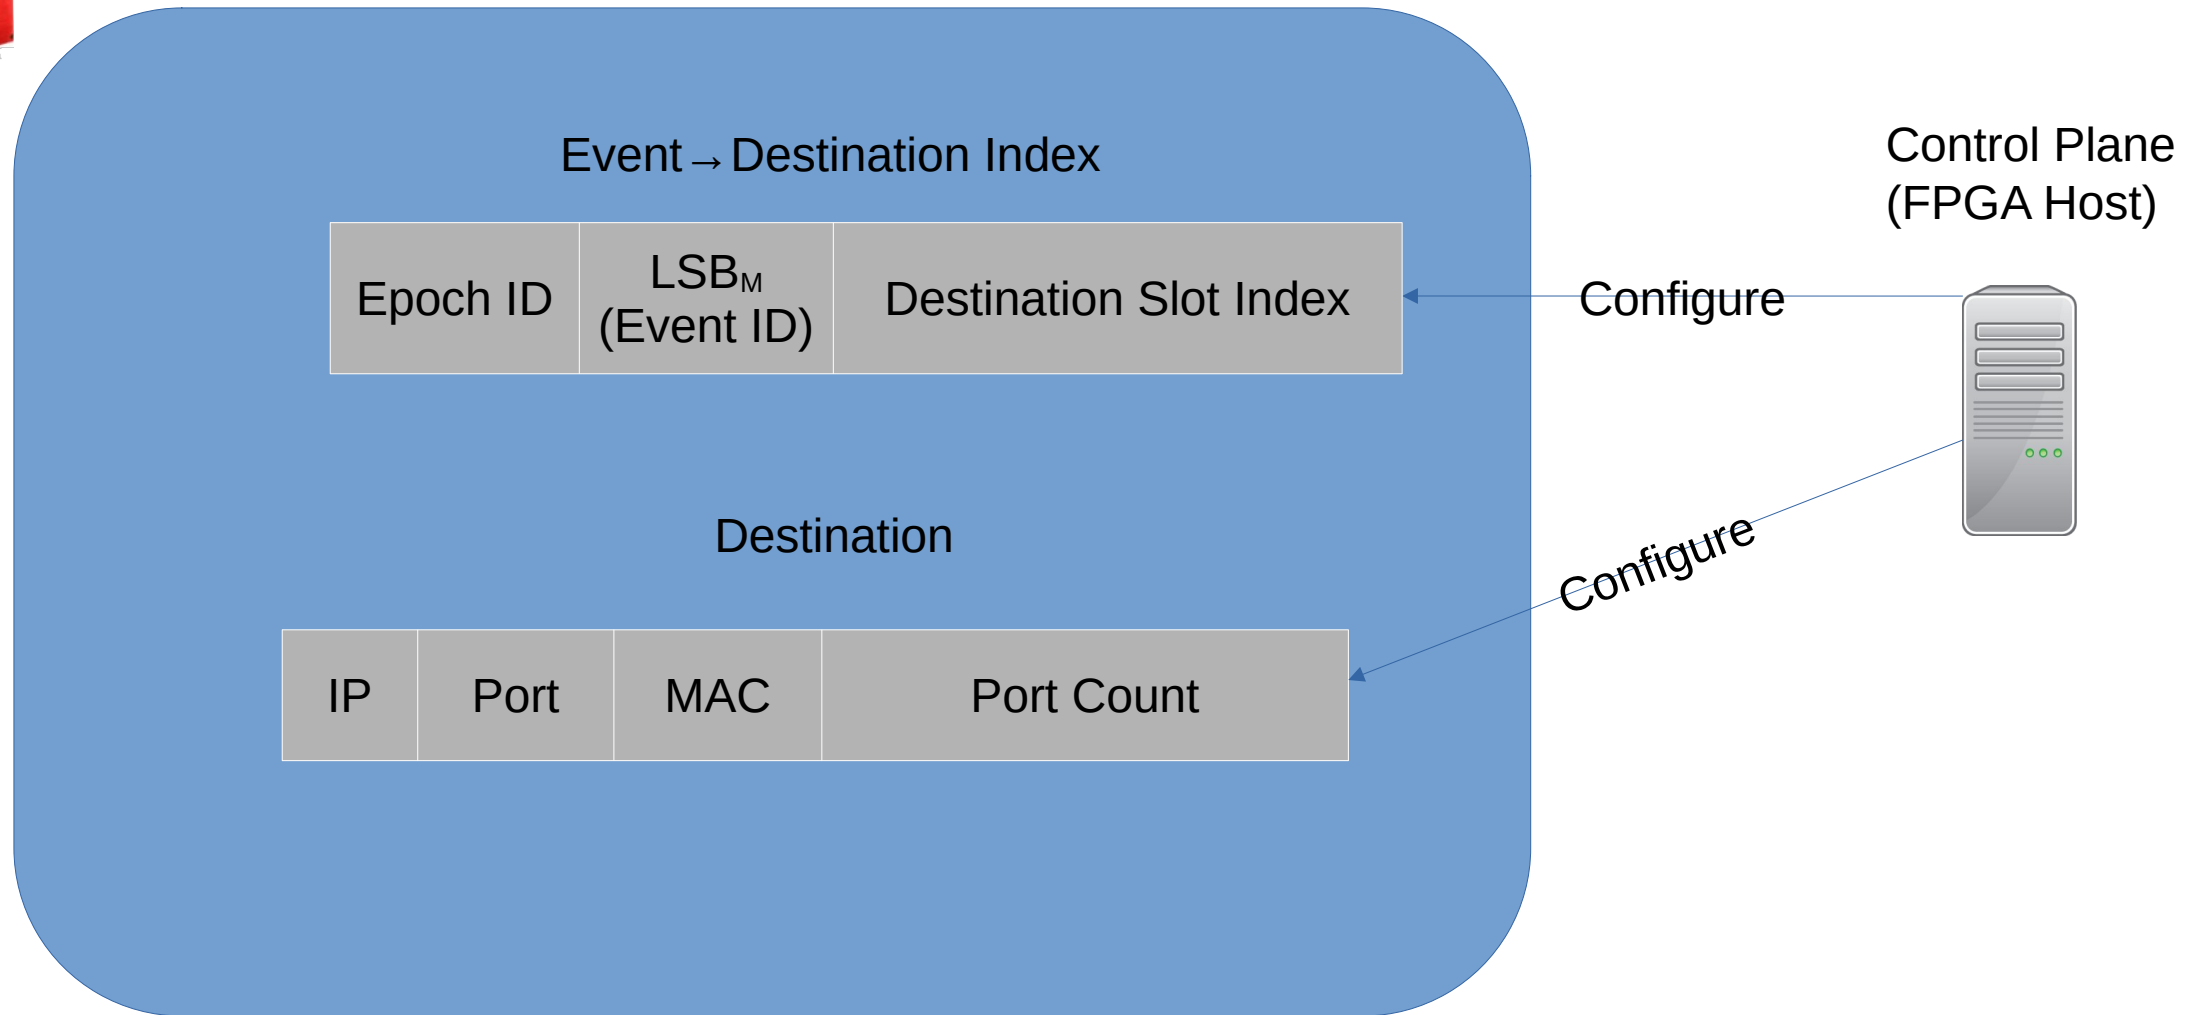
## Compute Host



# LB Data Plane (FPGA): Network / Calendar Setup




# LB Data Plane (FPGA): Destination Setup for Each Epoch



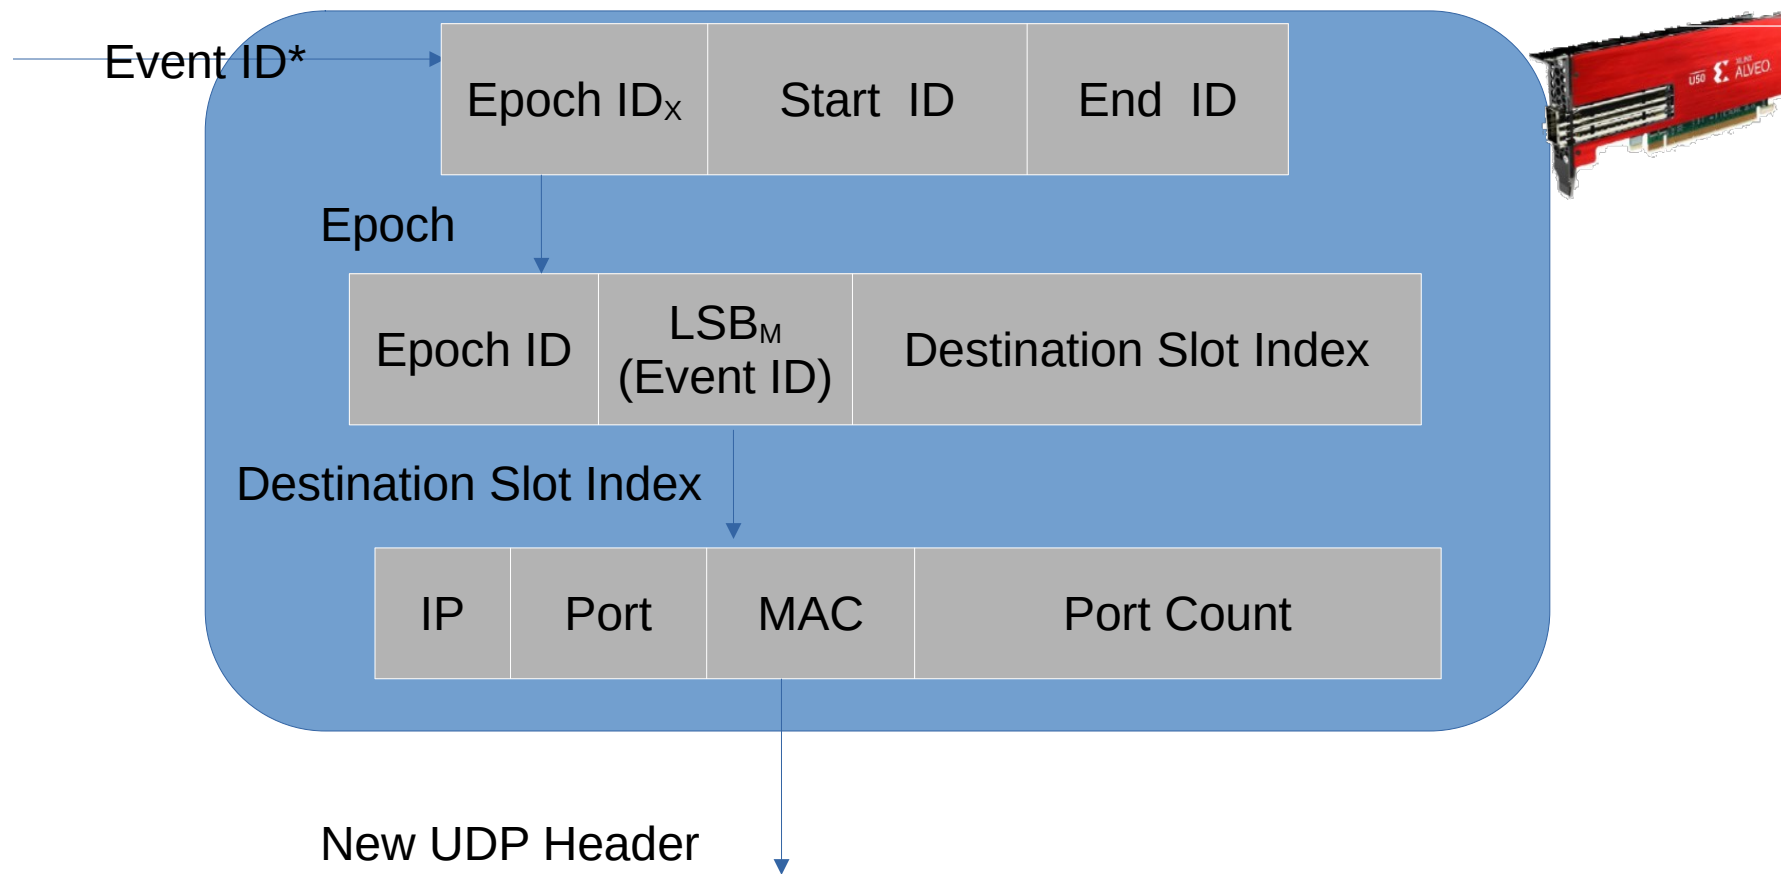
**Epoch:** Contiguous Event ID Subspace

# LB Data Plane (FPGA): Packet Processing



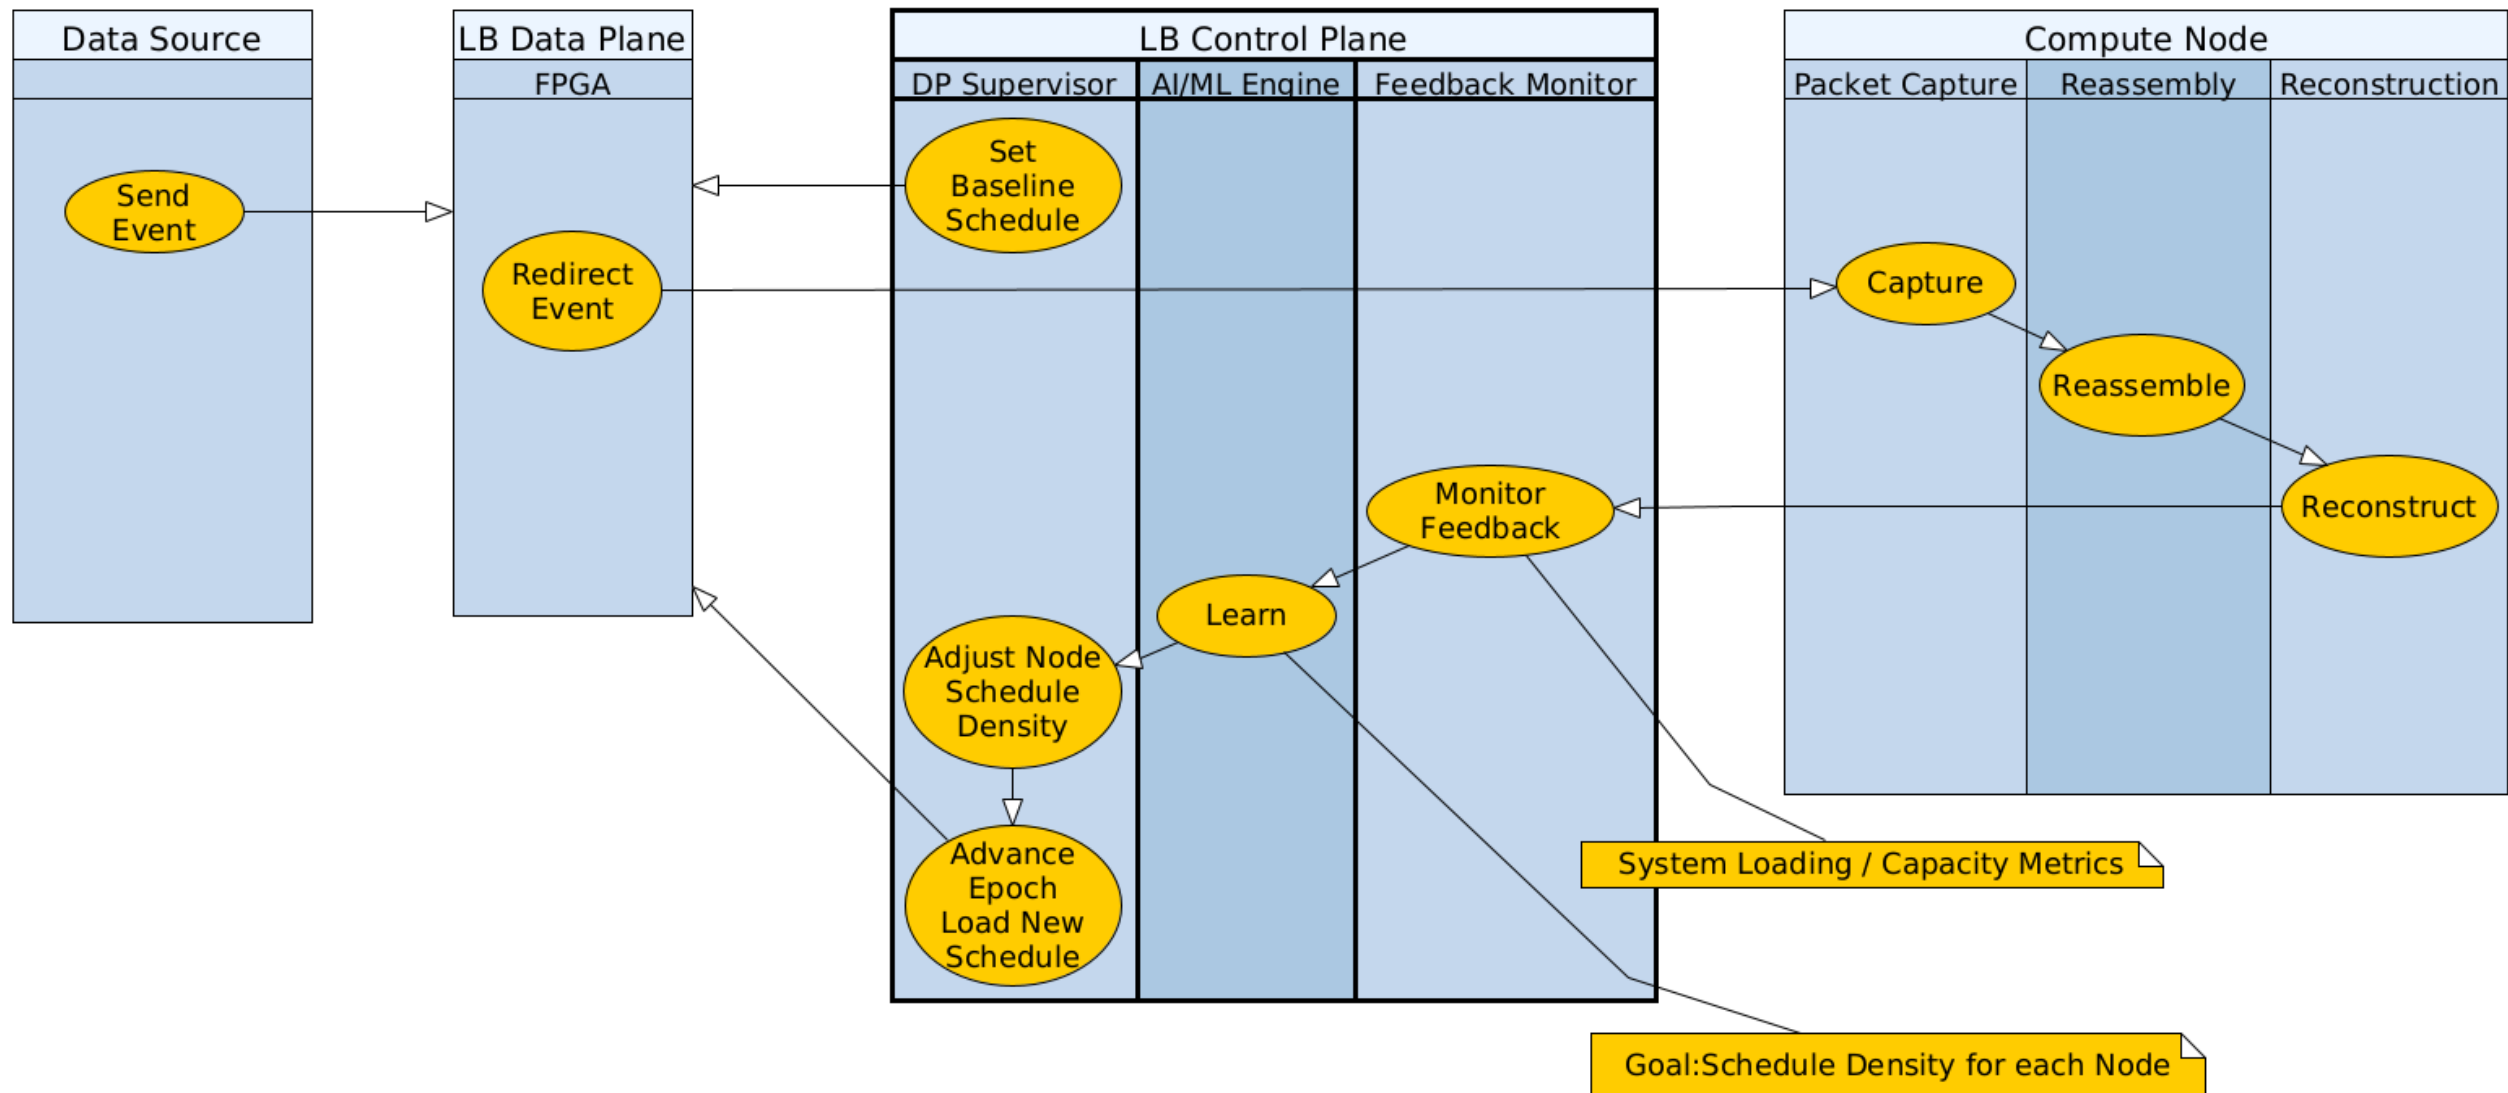
Typical Large Science Instrument

\*Event ID = Tick

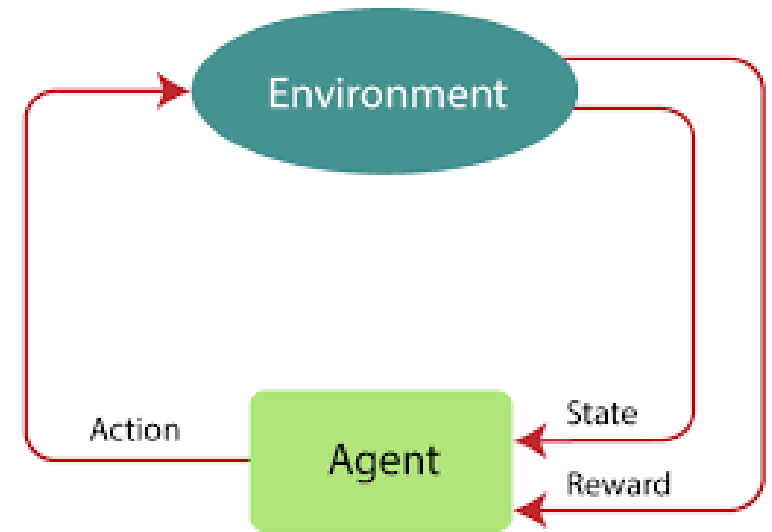
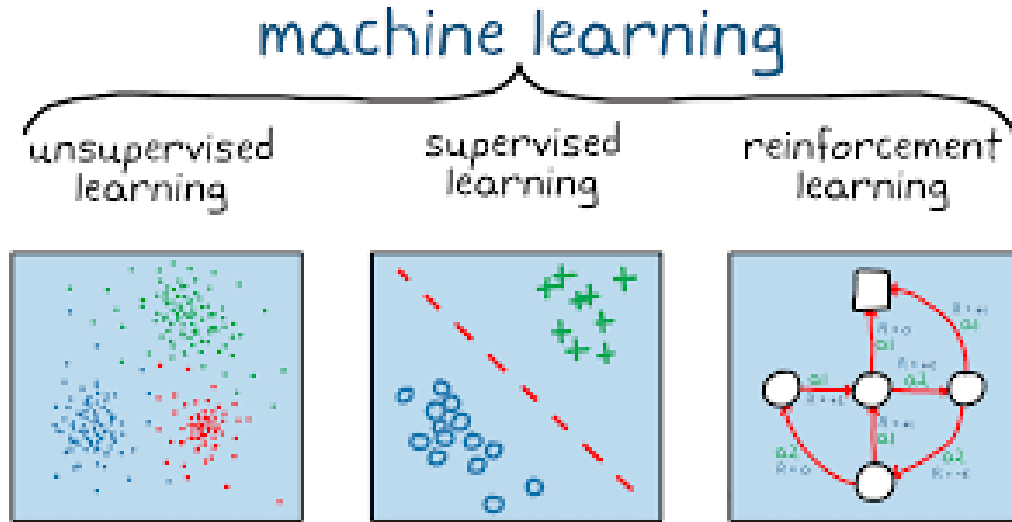




# LB Control Plane (FPGA Host):



# Reinforcement Learning



## Q Learning:

- Many Variants
- Exploration / Exploitation
- Exploration → Learning
- Exploitation → Control


$$\text{New } Q(s, a) = Q(s, a) + \alpha [R(s, a) + \gamma \max_{a'} Q'(s', a') - Q(s, a)]$$

Learning Rate:  $\alpha$       Discount Rate:  $\gamma$

Labels for the equation components:

- New Q value for the state and action:**  $\text{New } Q(s, a)$
- Current Q values:**  $Q(s, a)$
- Reward for taking an action in a state:**  $R(s, a)$
- Maximum expected future reward:**  $\max_{a'} Q'(s', a')$
- Current Q values:**  $Q(s, a)$

# Reinforcement Learning -> Farm Scheduling



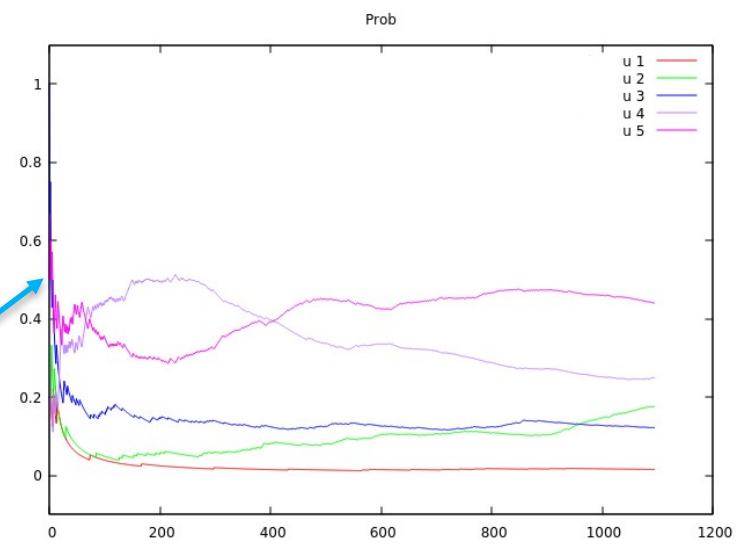

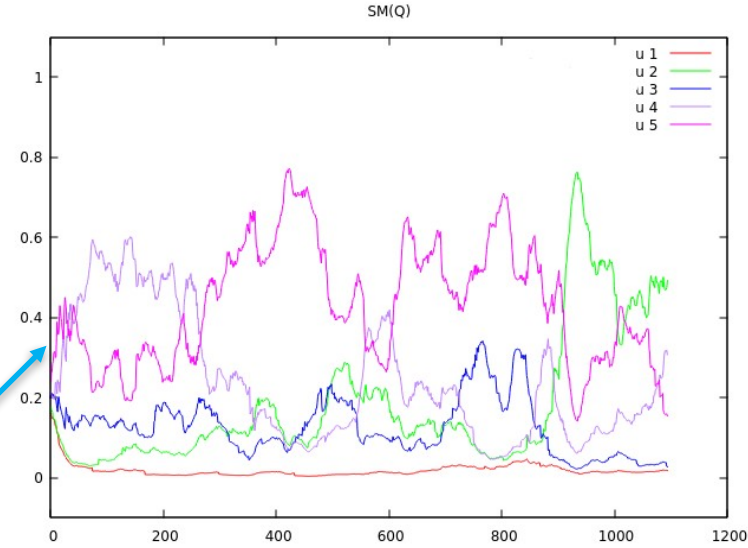
Rewards: Feedback from Nodes ->  
Relative Availability for New Work

Q Learning: SARSA

Action Selection(Q):


SoftMax Policy +  
Weighted Sampling

LaPlace's Rule of Succession(Action) ->  
Probability of Scheduling Each Node

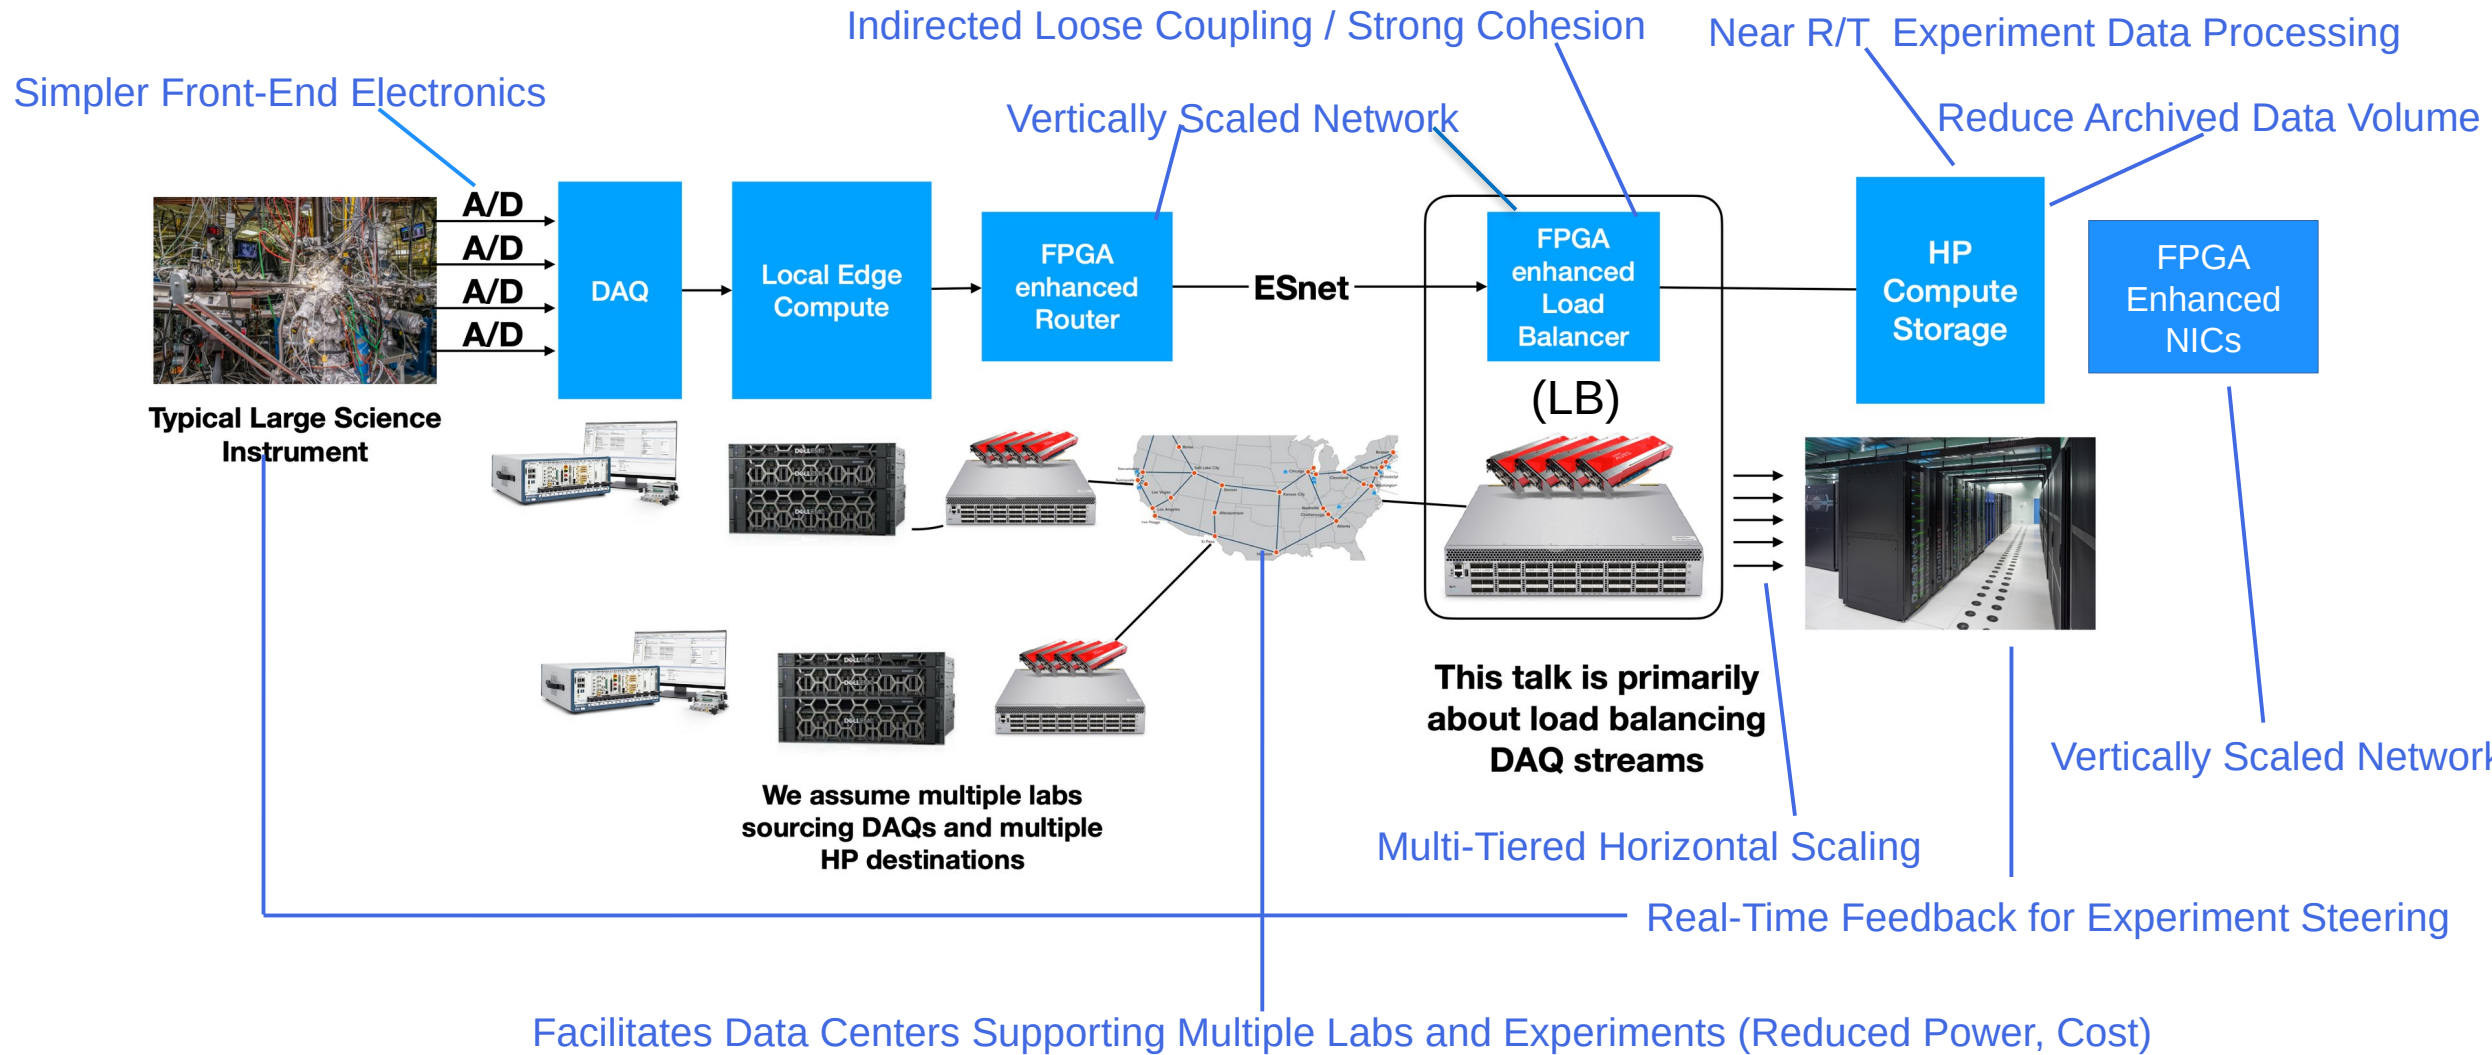


# Q Learning -> LB Scheduling

Epoch ID	LSB <sub>M</sub> (Event ID)	Destination Slot Index
0	LSB <sub>M</sub> (Event ID)	Destination Slot Index
0		
0		
...		
1		
1		
1		



# Benefits



# EJFAT - Summary

---

- Supports Triggered + Streaming Detector, Data Center Workflows
- Data Event Aggregation
- Data Channel to Port Distribution
- Real-Time UDP Packet Re-Direction with Fixed  $\mu$ Sec Latency
- Real-Time AI/ML Guided Destination Load Balancing
- Real-Time AI/ML Guided Cluster Resource Provisioning
- Decouples Edge and Cluster, Geographically and Network
- Running at 100Gbs – can support Up To 200Gbs

# Questions ?