OAK RIDGE National Laboratory

From Path-Aware to Application-Aware Source Routing using Traffic Classes

U.S. DEPARTMENT OF ENERGY

ORNL IS MANAGED BY UT-BATTELLE LLC
FOR THE US DEPARTMENT OF ENERGY

Authors

Shashwitha Puttaswamy

Research Scientist at George Washington University

Mariam Kiran

AI Scientist at Oak Ridge Group, Lead to ORNL's Quantum Networking and Communications Group.

Background and Motivation : AI in Traffic Engineering

AI/ML techniques are increasingly used to learn network patterns and find optimal traffic engineering solutions.

Relied on TCP/UDP port numbers or deep packet inspection. This is often ineffective or impossible with modern encrypted traffic.

Uses attributes like flow size, duration, and packet info to identify behavior (e.g., 'elephant' vs. 'mice' flows).

ML-Driven Classification: Techniques like K-means, SVMs, and neural networks are now used to classify flows and infer application behavior even without inspecting the content

Key Takeaway:

TE is complex, involving decisions on capacity planning, flow allocation, and traffic splitting to minimize congestion.

This work contributes to In-band Network Telemetry (INT) and P4 programmability for self-engineering networks

Introducing edge-based traffic classification and a dynamic queuing mechanism to optimize network utilization and allocation.

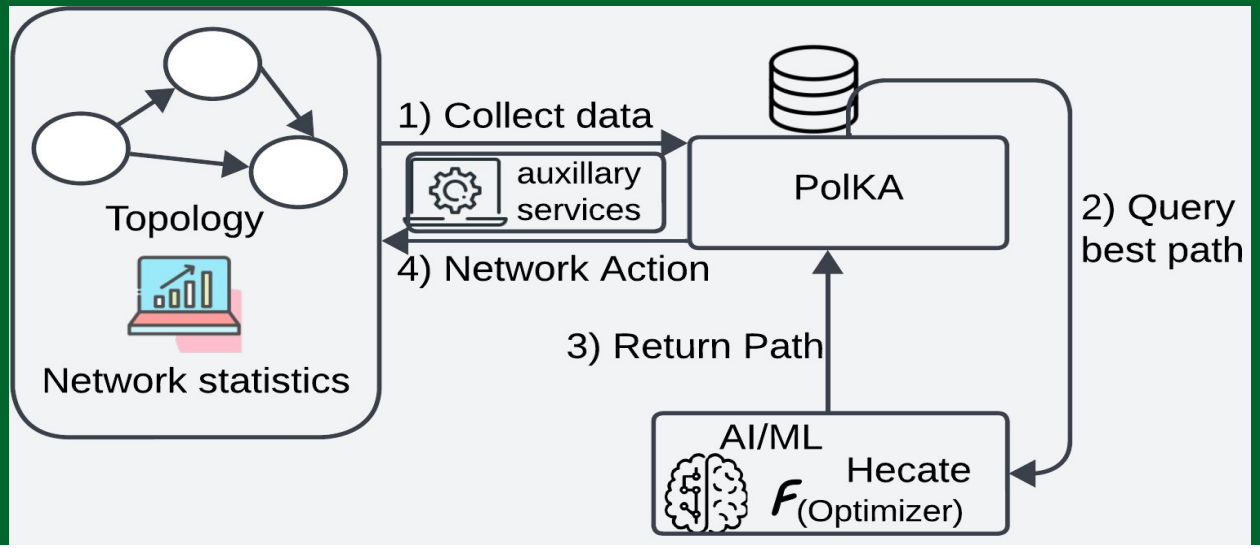

PolKA-Hecate integration framework

- Last year we presented Hecate-Polka service:
  - PolKA SR routing service
  - Hecate AI-Network Driven service
- Here we are adding finer granular decision making using traffic classes

Hecate's Seven Traffic Classes

Class 0: Low Latency

- Very low inter-arrival time (<0.005s).
- Small/medium packet size (e.g., between 60 and 200 bytes) and a sustained duration (e.g., greater than 5 seconds).
- These flows have minimal delay and jitter.

For VoIP, Online Gaming, and Video Conferencing.

Class 1: Priority Flows

- Destined for critical IP addresses.
- Moderate byte size.(e.g., less than 1000 KB), and a duration typically less than 60 seconds.

For essential applications of business, the control-plane traffic.

Class 2: High Throughput

- Large total bytes (>50,000 KB). High avg.
- Packet size. (e.g., greater than 1000 bytes, indicating full-sized data packets), and a significant duration (e.g., greater than 30 seconds).

For large file transfers,Software updates.

Class 3: Long-Living

- Long duration (>60s)and moderate total bytes (e.g., between 100 KB and 50,000 KB).

For SSH sessions, RDP, long-term data backups.

Class 4: Short-Living

- Very short duration (‹5s) small total bytes (e.g., less than 10 KB),
- Low number of packets (e.g., less than 10 Packets).

For DNS queries, NTP, IOT.

Class 5: Intent-Based

- Small average packet size (e.g., less than 100 bytes), relatively low IAT (e.g., less than 0.1 seconds,indicating some regularity), and a sustained duration (e.g., greater than 10 seconds).

Class 6:Default Flows

- Any traffic that does not match the criteria of the other six classes.

Specific DTN server IPs. Custom QOS needs. For specialized sensor data, store-and-forward traffic

DSCP Value

Class ID	Name	DSCP Value (Binary)	Priority Metric	Typical Applications
0	Low Latency	0b101110 (EF)	Latency	VoIP, Online Gaming, Video Conferencing
1	Priority	0b101000 (CS5)	Latency	Critical Business Applications, Control Plane Traffic
2	High Throughput	0b011010 (AF31)	Bandwidth	Large File Transfers, Software Updates, Video Streaming (
3	Long-Living Flows	0b010010 (AF21)	Connections	SSH Sessions, RDP, Long-term Data Backups
4	Short-Living Flows	0b001010 (AF11)	Latency	DNS Queries, NTP, Small HTTP Requests, IoT Heartbeats
5	Intent-Based Flows	0b001100 (Custom)	Custom (Latency/Conn)	DTN Data Transfers, Specialized Sensor Data
6	Default Flows	0b000000 (BE)	Any	Best-Effort Traffic, General Web Browsing

Hecate Architecture Design

Control & Data Plane Separation

The Hecate architecture is separated into distinct planes to enable intelligent, adaptive networking.

- **Data Plane:** P4-programmable edge routers are responsible for high-speed packet forwarding,
  - Packet parsing and header manipulation,
  - Match-action Processing
  - In-band Network Telemetry (INT) data generation
  - Egress queuing and scheduling (WRR).
  - Traffic Sampling / Mirroring

Applying dynamic Weighted Round Robin (WRR). WRR queues (Q_0 to Q_6) at each egress port.

- **Control Plane:** The "brain" of the architecture. It facilitates intelligent decision-making, orchestration, and dynamic configuration of the data plane in real-time.
  - Edge node classifier (ENC)
  - Core Network State Predictor (CNSP)
  - Dynamic Resource Orchestrator (DRO)

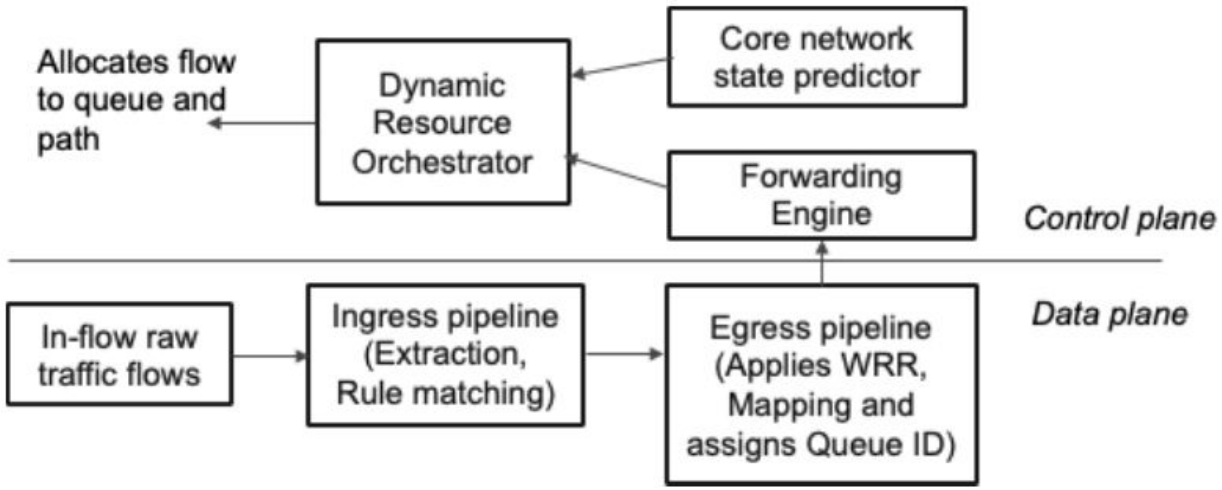

Control Plane

Edge Node Classifier:ENC

Uses ML (Decision Trees) to classify flows based on behavioral attributes (IAT, duration) and assigns a class ID (0-6).

Core Network State Predictor:CNSP

Collects and processes granular telemetry (INT) data from all P4 devices.

Predicts future traffic loads and identifies potential congestion.

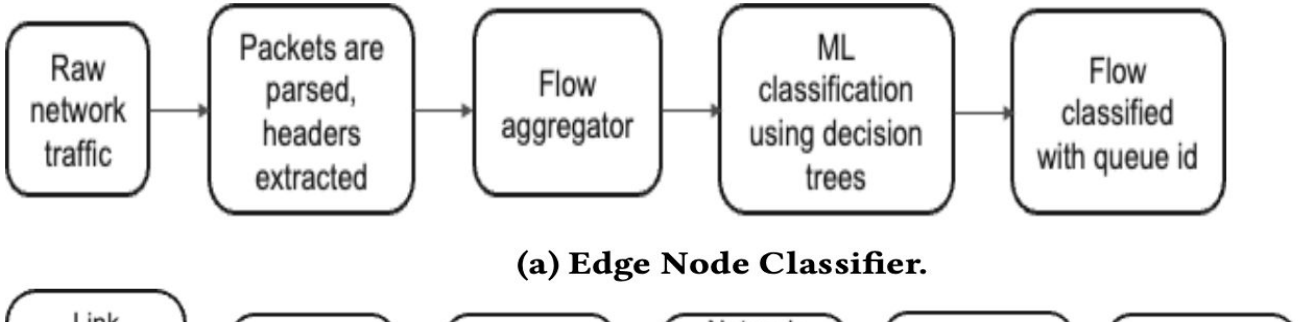

(a) Edge Node Classifier.

(b) Core Network Predictor.

Dynamic Resource Orchestrator:DRO

Central intelligence.Gathers data from ENC & CNSP, calculates queue weights, selects optimal paths (using Segment Routing), and pushes rules to the data plane.

- Evaluates the flow and calculates its weight.
- (DSCP) value to the flow, which signals its quality of service treatment throughout the network.
- Path selection algorithm (e.g., Least Connection, Weighted Least Connection, Least Response Time, Resource-Based) based on the flow's classified traffic class and the current network conditions.
- The policy for all network rules and security constraints.
- segment routing (SR) to determine the optimal explicit path for the flow to traverse, taking into account the selected load balancing objective and network constraints.

Fair Queuing: Dynamic Weighted Round Robin (WRR)

Dynamic & Fair Allocation.

This method ensures that all traffic class queues (QO-Q6) are emptied quickly and fairly, preventing starvation.

- Not Static: Weights are dynamically calculated and adjusted by the DRO based on the real-time network conditions and traffic volume.
- Work-Conserving: The scheduler intelligently bypasses any empty queue, ensuring no bandwidth is wasted on idle queues and maximizing network utilization.

Experiment & Simulation Setup

Traffic Generation: A Poisson distribution was used to model the arrival process (Inter-Arrival Time) for new flows, providing a realistic representation of network dynamics.

Dynamic Path Selection (by DRO):

Low Latency: The DRO prioritizes the Least Response Time to pick the fastest SR path.

High Throughput : The DRO uses a resource-based approach focusing on available bandwidth to pick an SR path with high capacity and low utilization.

Long-Living: The DRO uses weighted least connection to balance the number of active sessions across paths, considering their capacities.

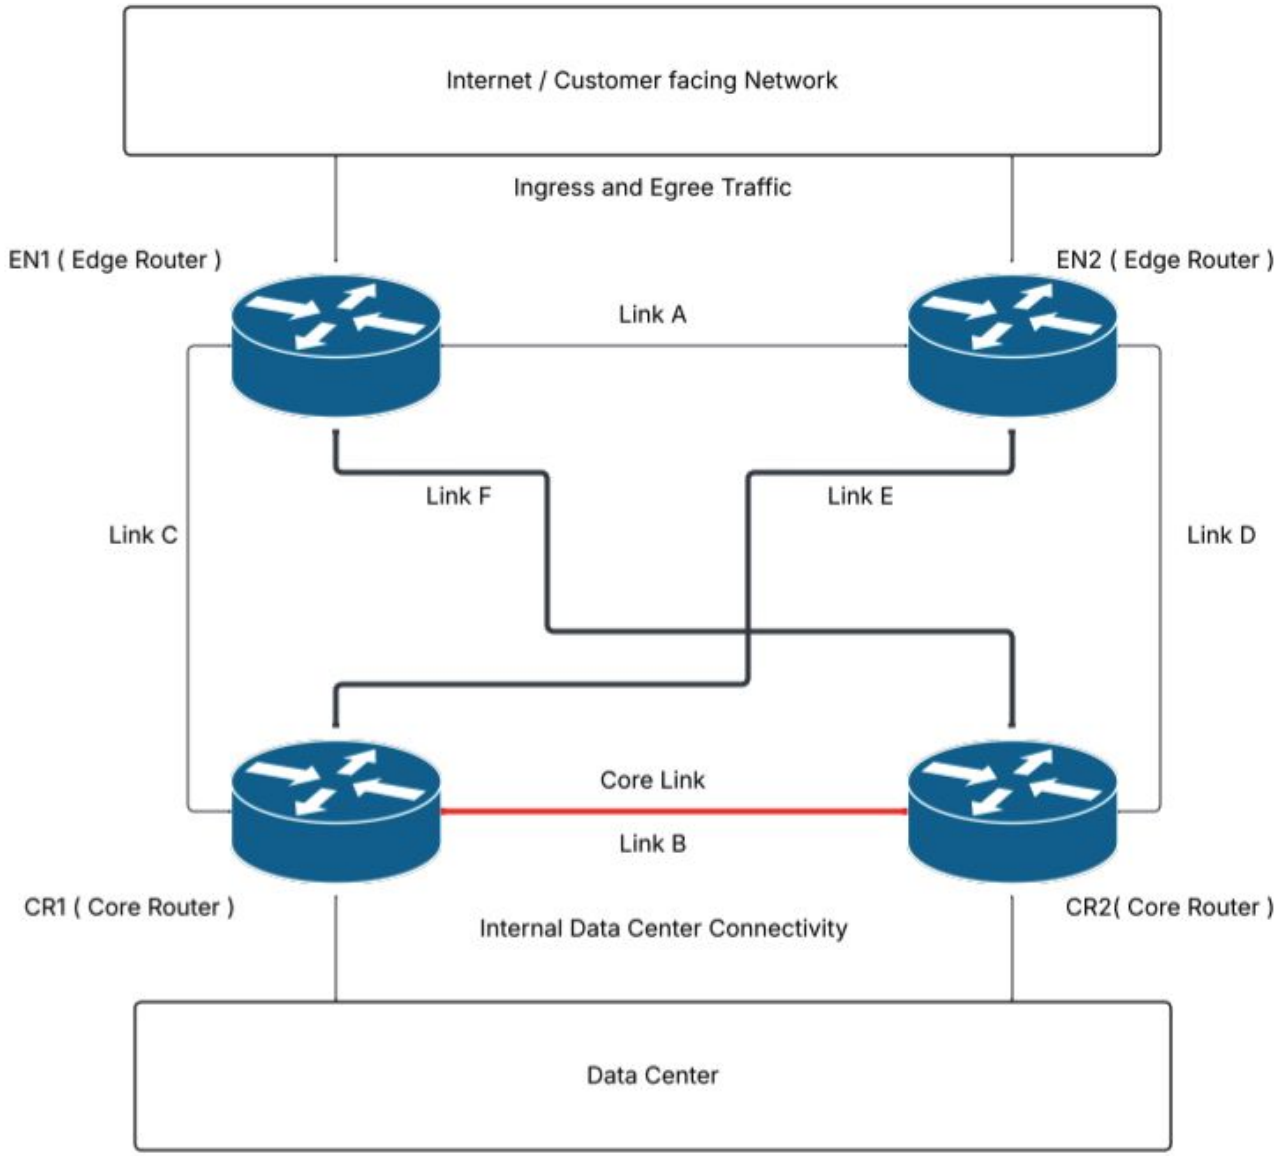

Results: Average Total Bandwidth Utilization

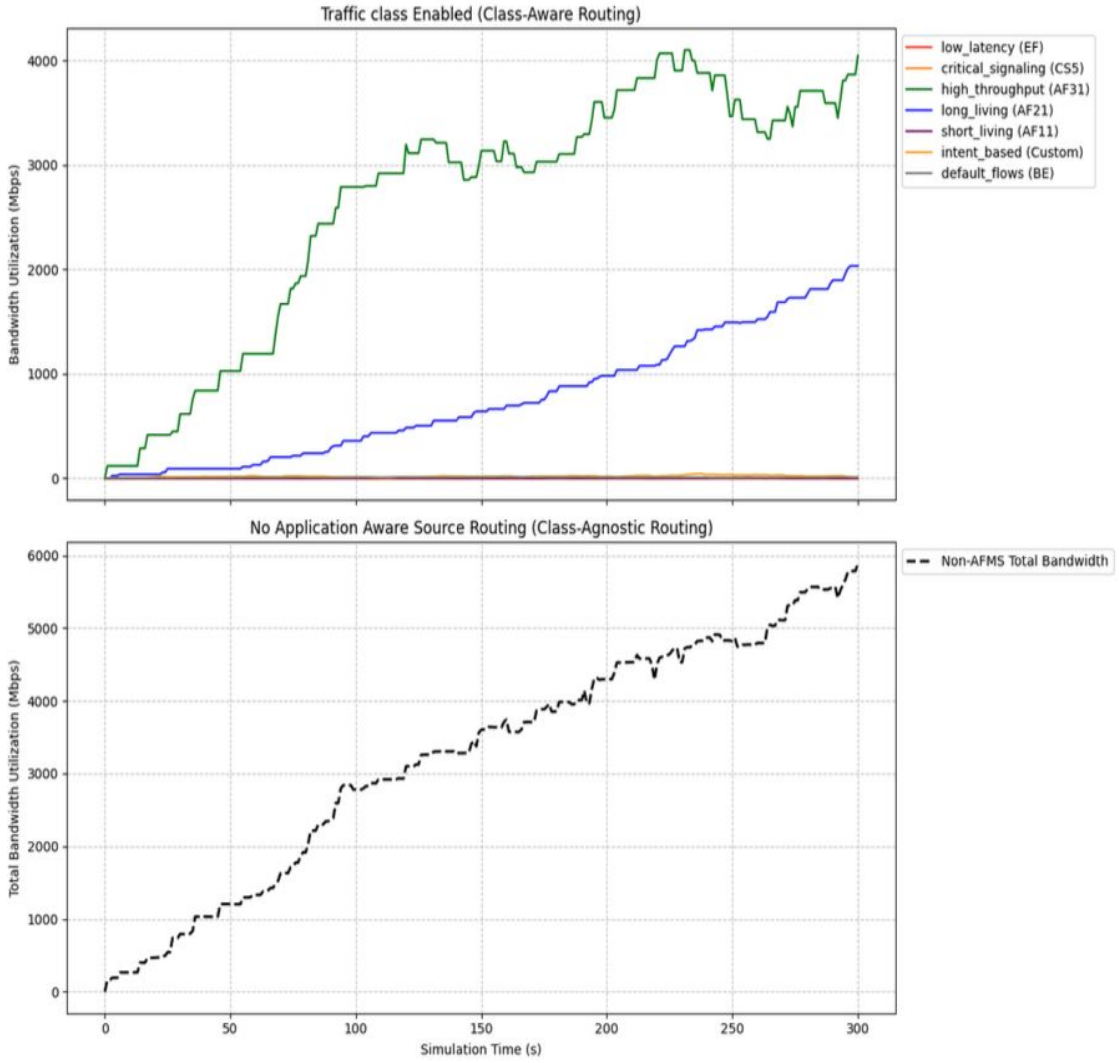

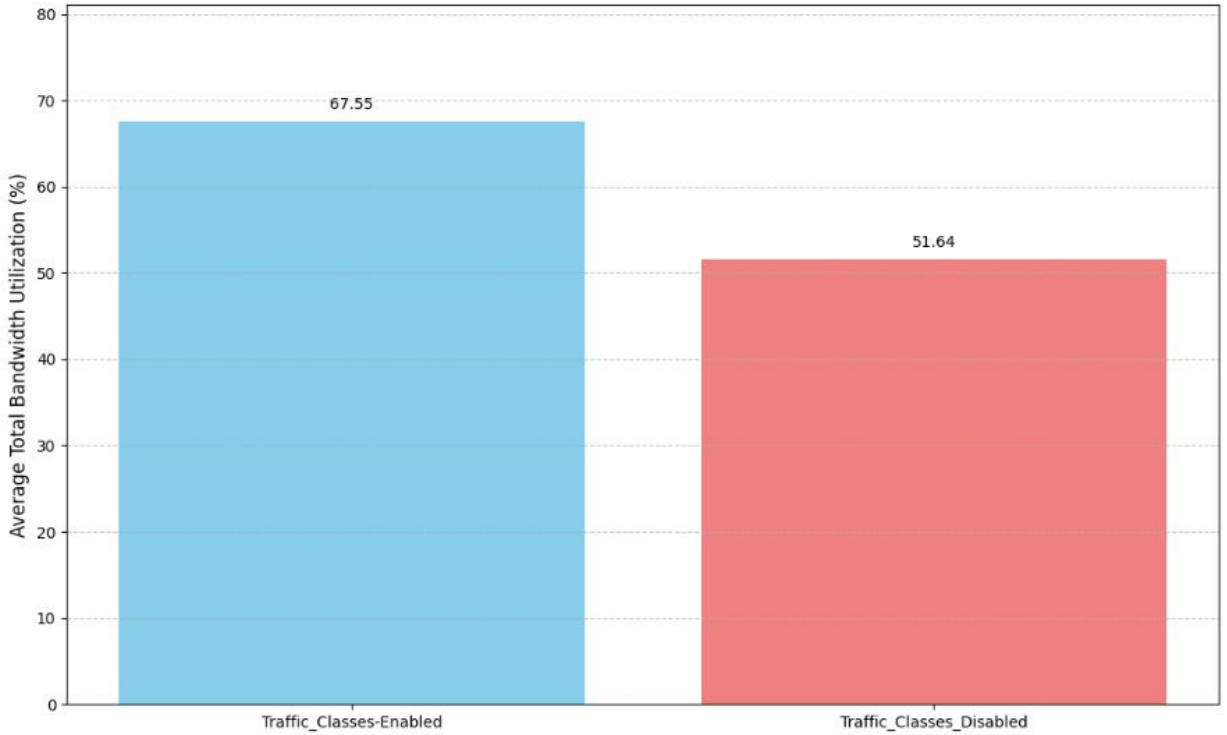

Enabling class-aware routing significantly improved average total bandwidth utilization from 51.84% to 67.59%.

Questions?
Thank you for attending.

Shashwitha Puttaswamy | sputtaswamy@email.gwu.edu
Mariam Kiran | kiranm@ornl.gov