Graph Neural Network-Based DDoS Protection for Data Center Infrastructure

Kartikeya Sharma
karsharma@equinix.com
Equinix

Craig Jacobik
cjacobik@equinix.com
Equinix

Agenda

- **The Growing Threat Landscape:** Why data centers are prime targets?
- **Limitations of Traditional Methods:** Where current approaches fall short
- **What are Graph Neural Networks?**
- **Our Solution:** Heterogeneous Graph U-Nets
- **System Architecture & Implementation:** Real-world deployment strategies
- **Experimental Results:** Performance across three datasets
- **Limitations, Future Directions & Q&A**

The Growing Threat Landscape

The Growing DDoS Threat to Data Centers

$340.2B

Global DC Market in 2024

3.47 Tbps

Recent Attack Size

5X

DDoS Attack Growth since 2016

Critical Challenge for Providers

- **Provider infrastructure is a high-value target**
- **Single point of failure affects ALL tenants**
- **Attacks threaten power, cooling, network backbone**
- **Cascading failures impact entire data center**

Limitations of Traditional Methods

Why Traditional Methods Fall Short

Signature-based

- Cannot detect novel attack variants
- Require constant signature updates

Threshold-based

- Excessive false positives
- Cannot adapt to traffic variations

Traditional ML

- Ignores relational context

Needs a structure aware and adaptive detection

What are Graph Neural Networks?

What are Graph Neural Networks?

Graph Neural Networks are a class of deep learning models that learn from connected data, helping us uncover hidden patterns in complex networks.

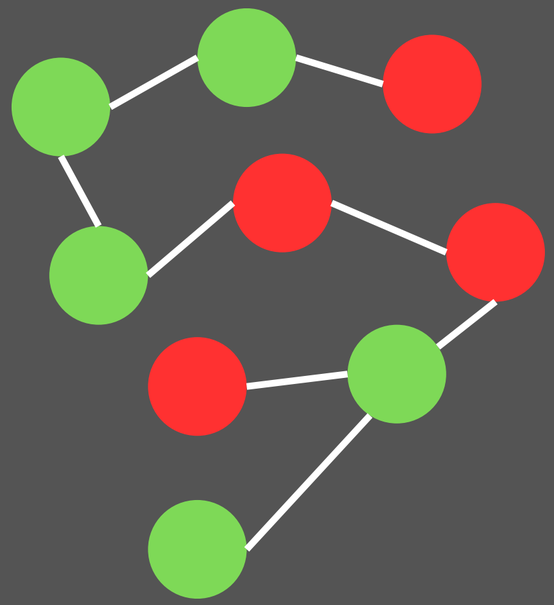

What are Graph Neural Networks?

Nodes (also known as vertices) represent entities or objects in a graph.

Edges represent the relationships or connections between nodes.

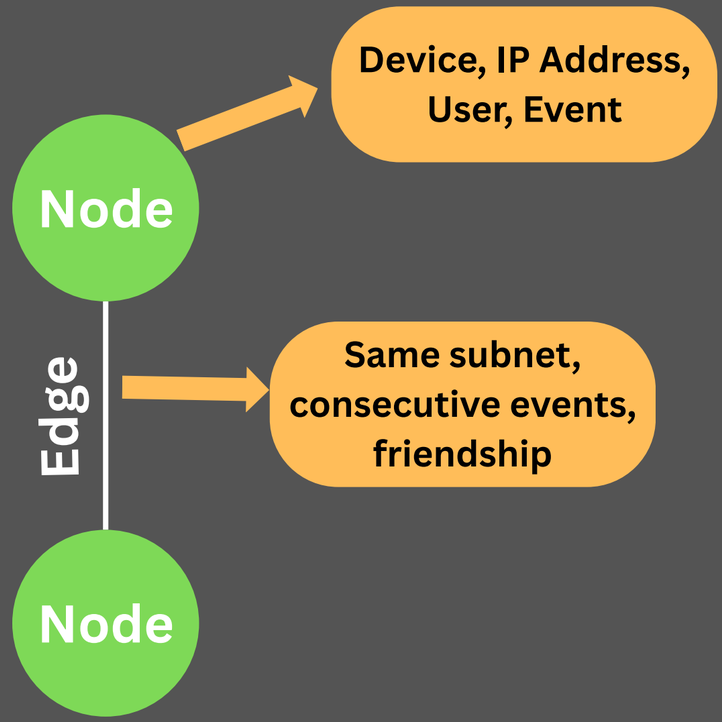

What are Graph Neural Networks?

GNNs learn rich node representations, called embeddings using Message Passing.

Message Passing Framework

1. **Aggregate:** Collect information from neighboring nodes
2. **Combine:** Merge neighbor info with node's own features
3. **Update:** Generate new node representation
4. **Repeat:** Multiple layers for broader context

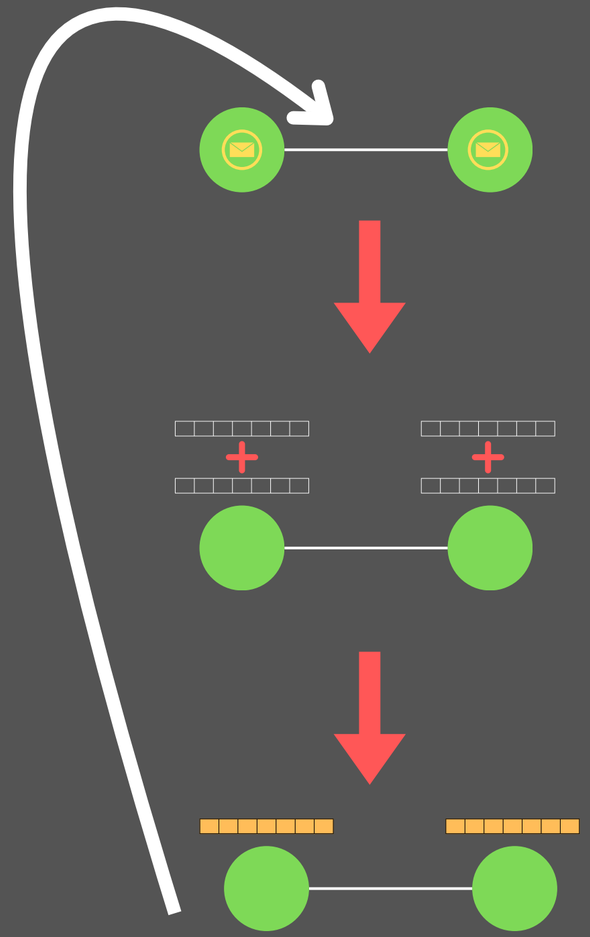

What are Graph Neural Networks?

GNNs learn rich node representations, called embeddings using Message Passing.

Message Passing Framework

1. **Aggregate:** Collect information from neighboring nodes
2. **Combine:** Merge neighbor info with node's own features
3. **Update:** Generate new node representation
4. **Repeat:** Multiple layers for broader context

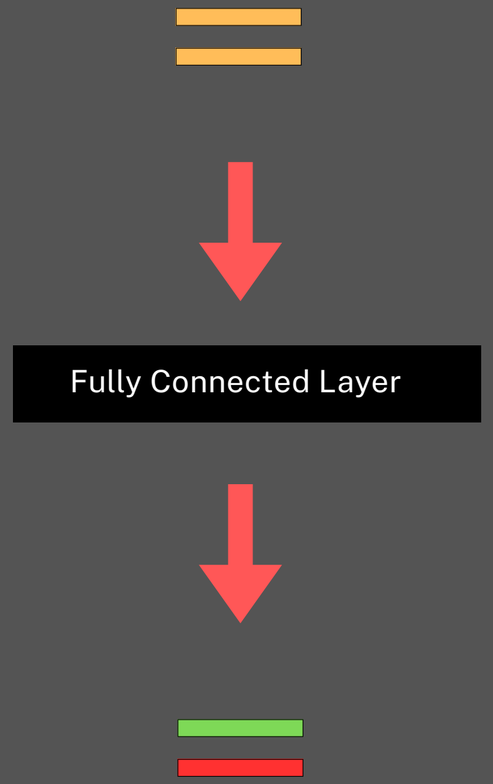

What are Graph Neural Networks?

GNNs have found applications in various domains, including:

- Social network analysis
- Molecular property prediction
- Knowledge graph completion
- Recommender systems

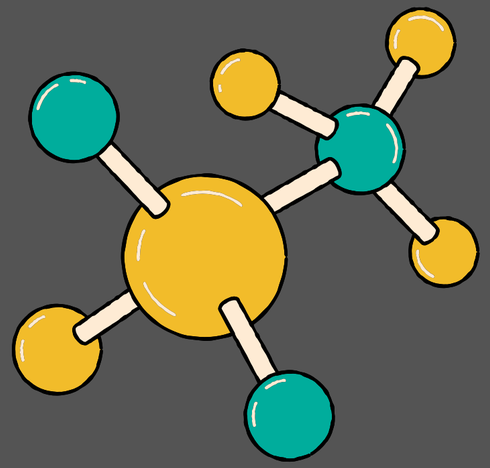

The GNN Intuition for the DDoS Detection

Network traffic will turn into a Graph.

Graph Neural Network will capture relational dependencies.

Leads to Context-aware, scalable and generalizable detection.

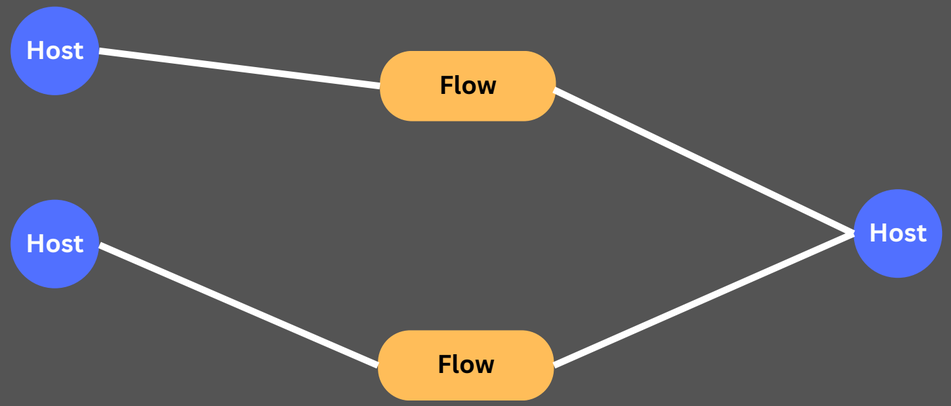

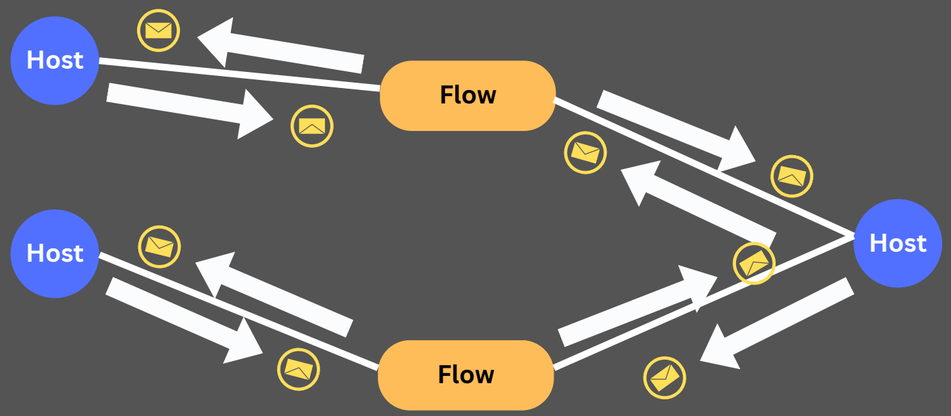

Our Solution

Network Traffic as a Graph

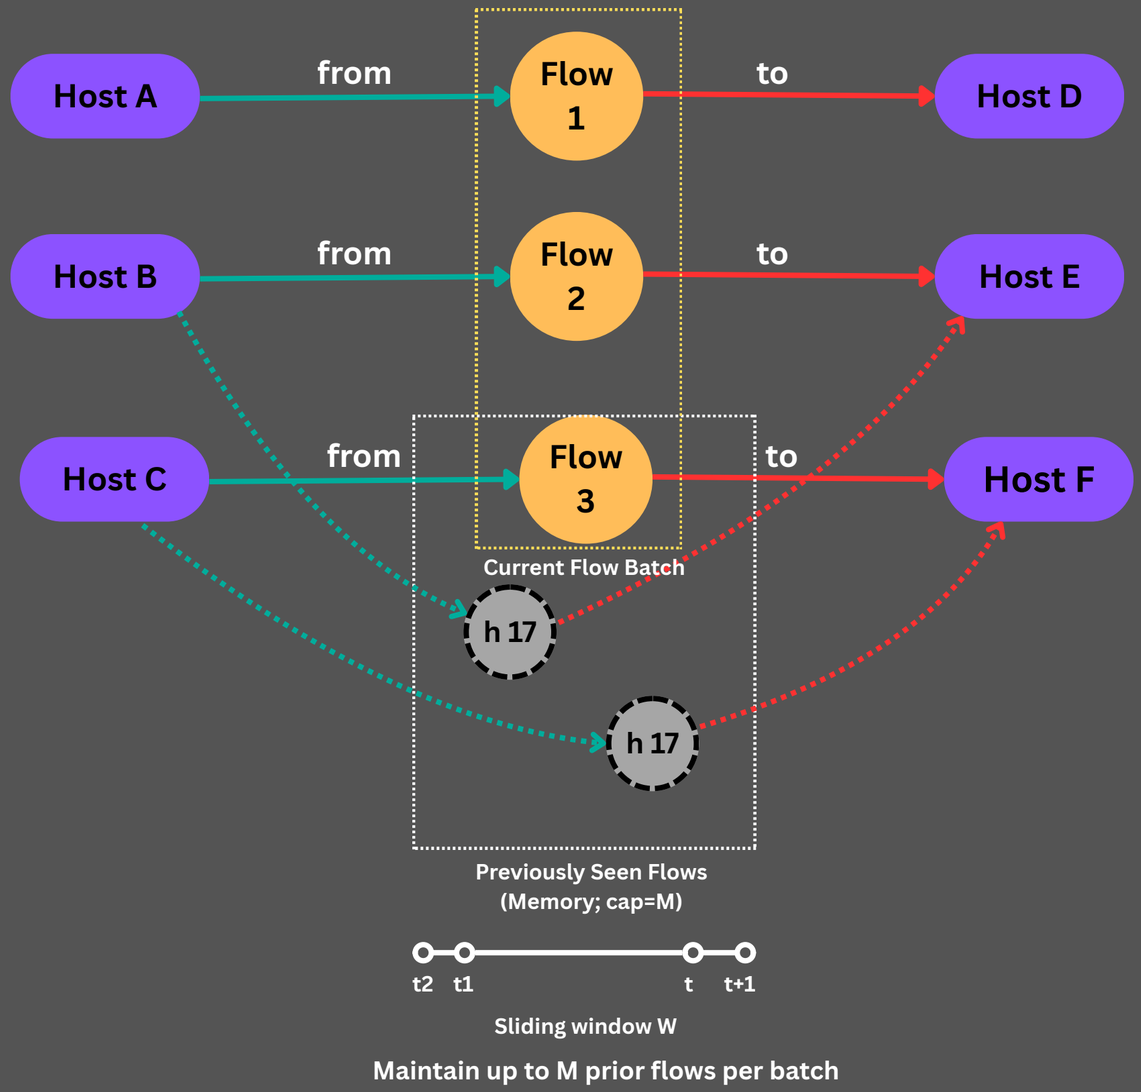

Network Traffic as a Graph

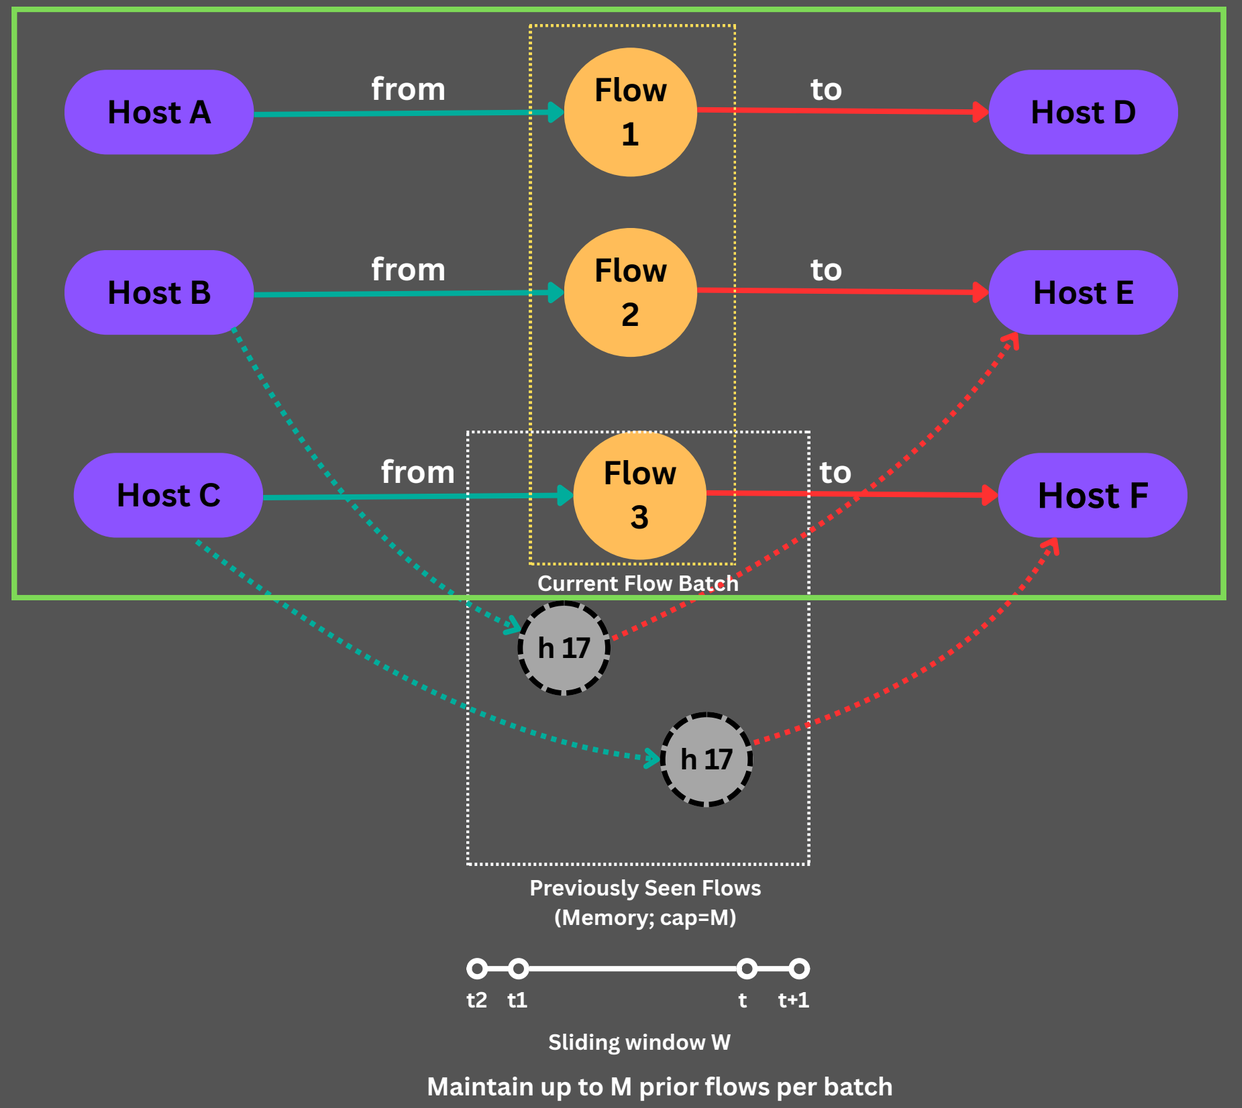

Network Traffic as a Graph

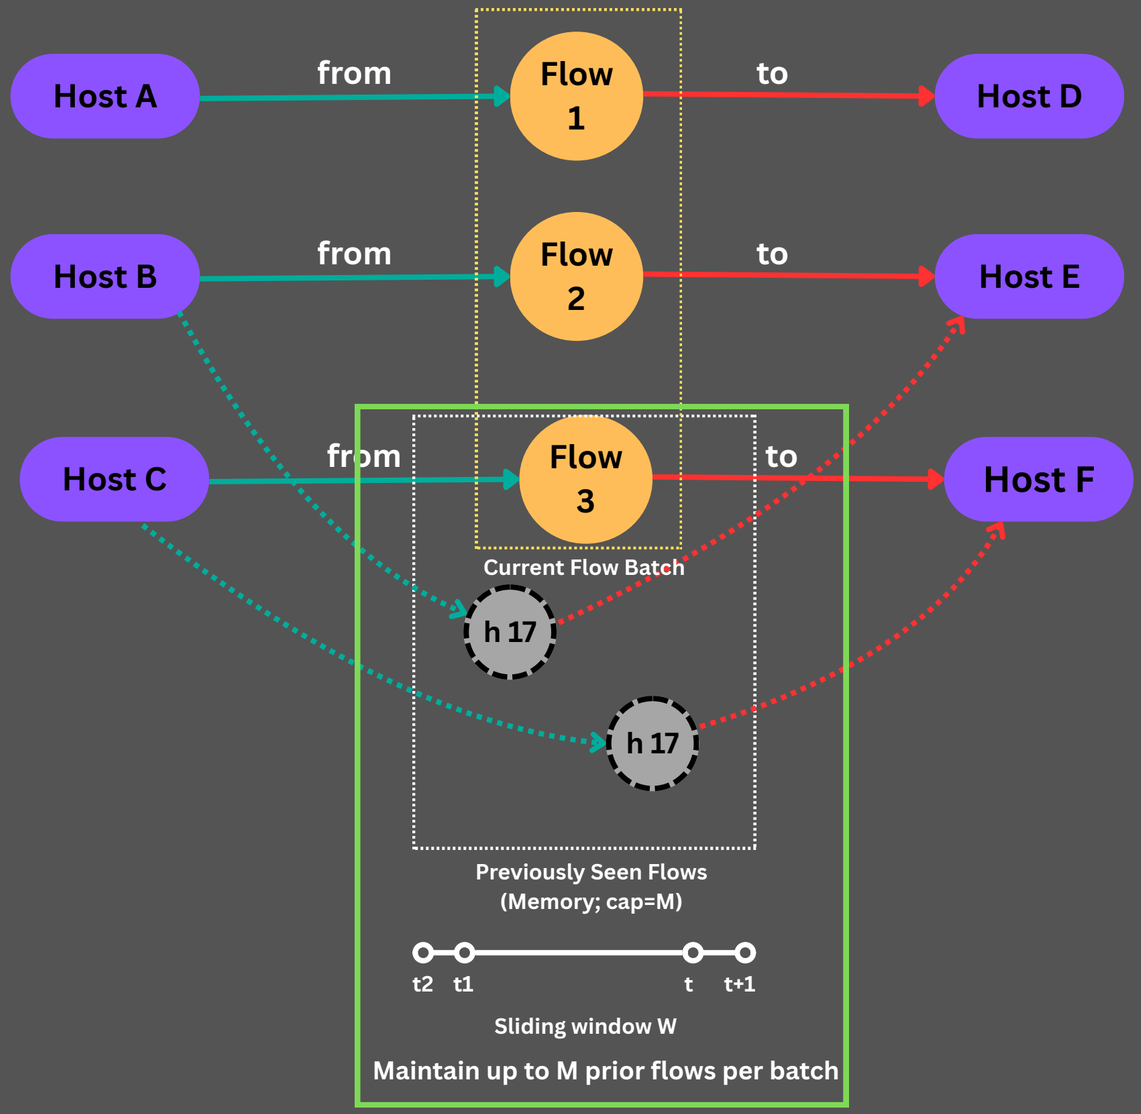

Heterogeneous Graph U-Nets Architecture

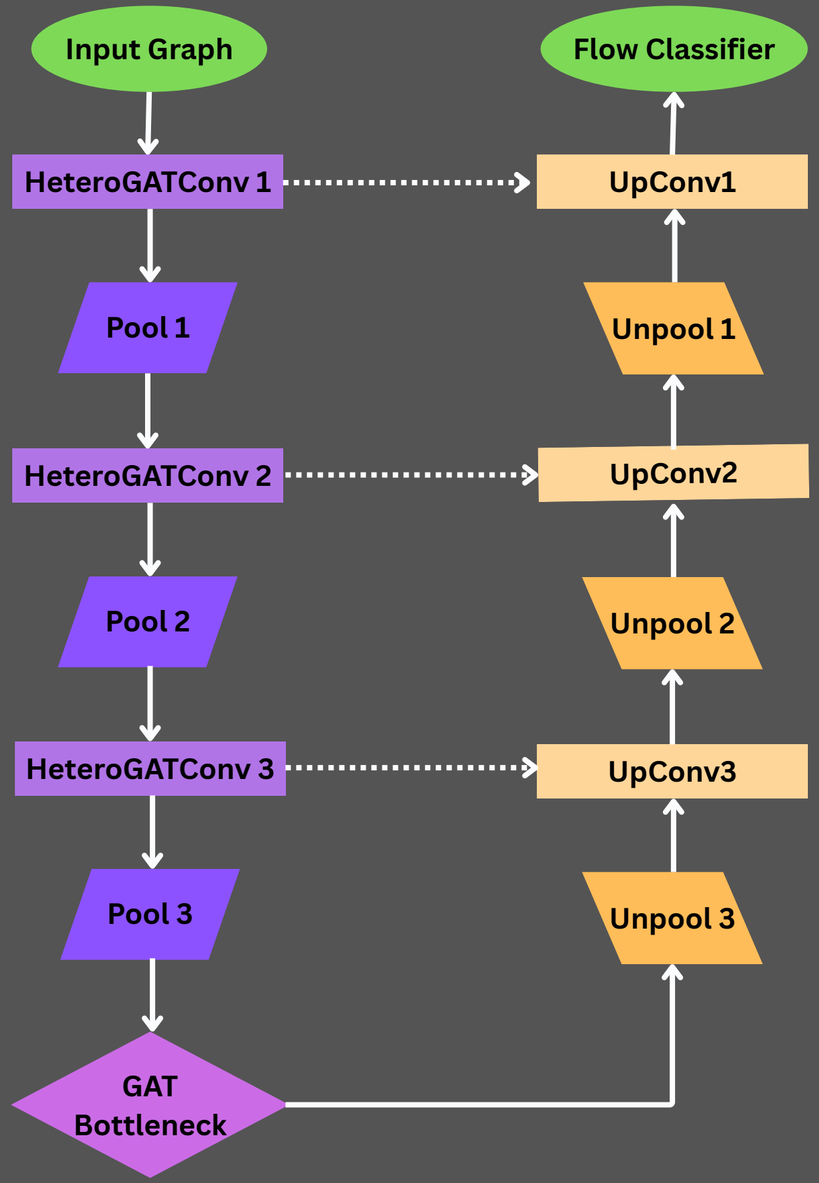

Heterogeneous Graph U-Nets Architecture

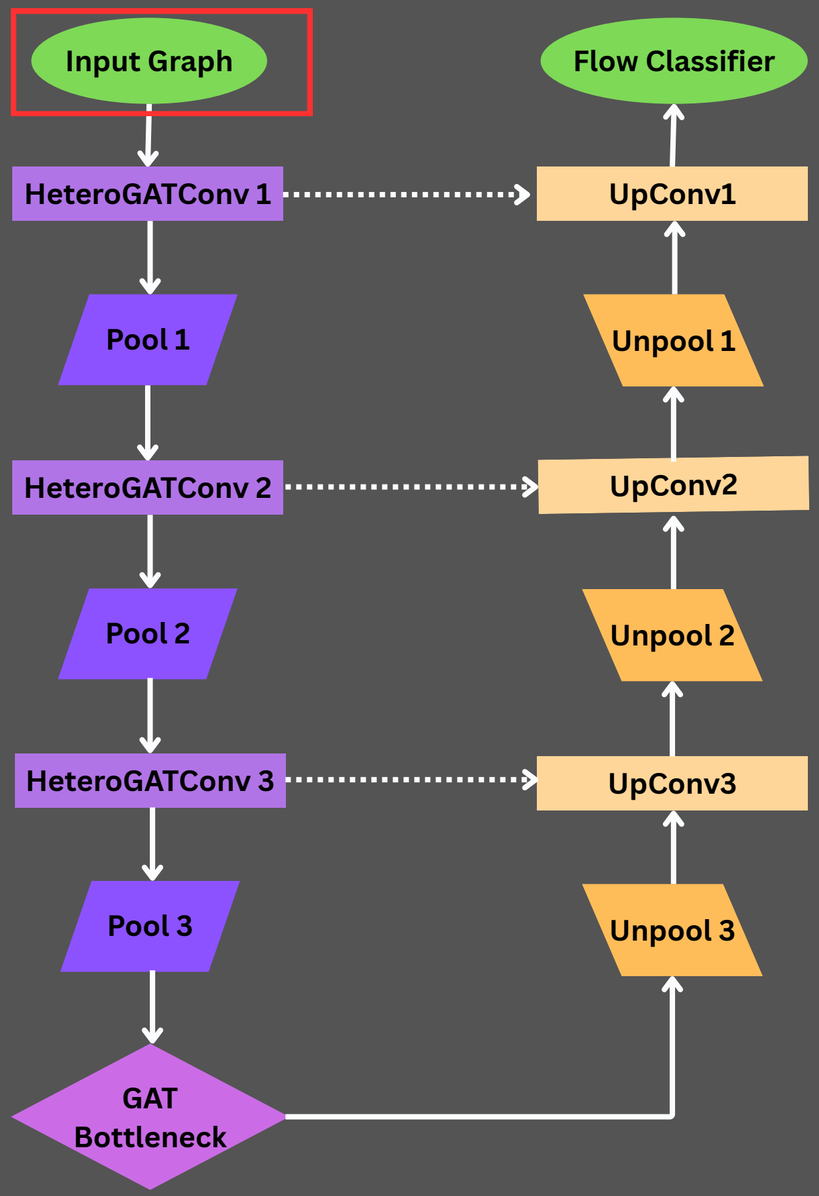

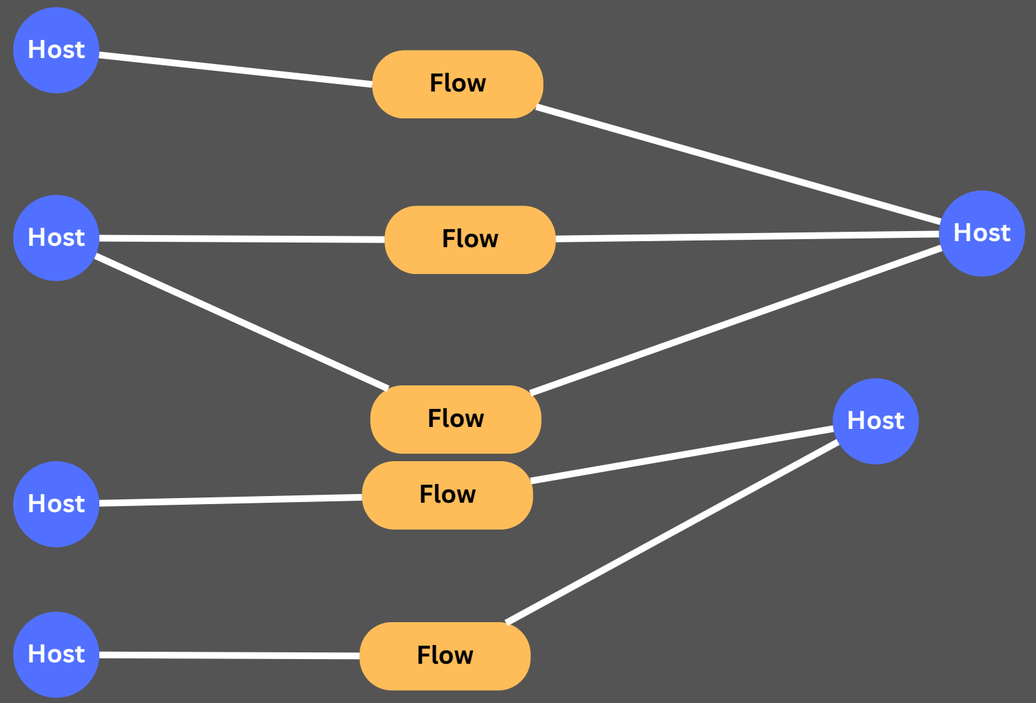

Heterogeneous Graph U-Nets Architecture

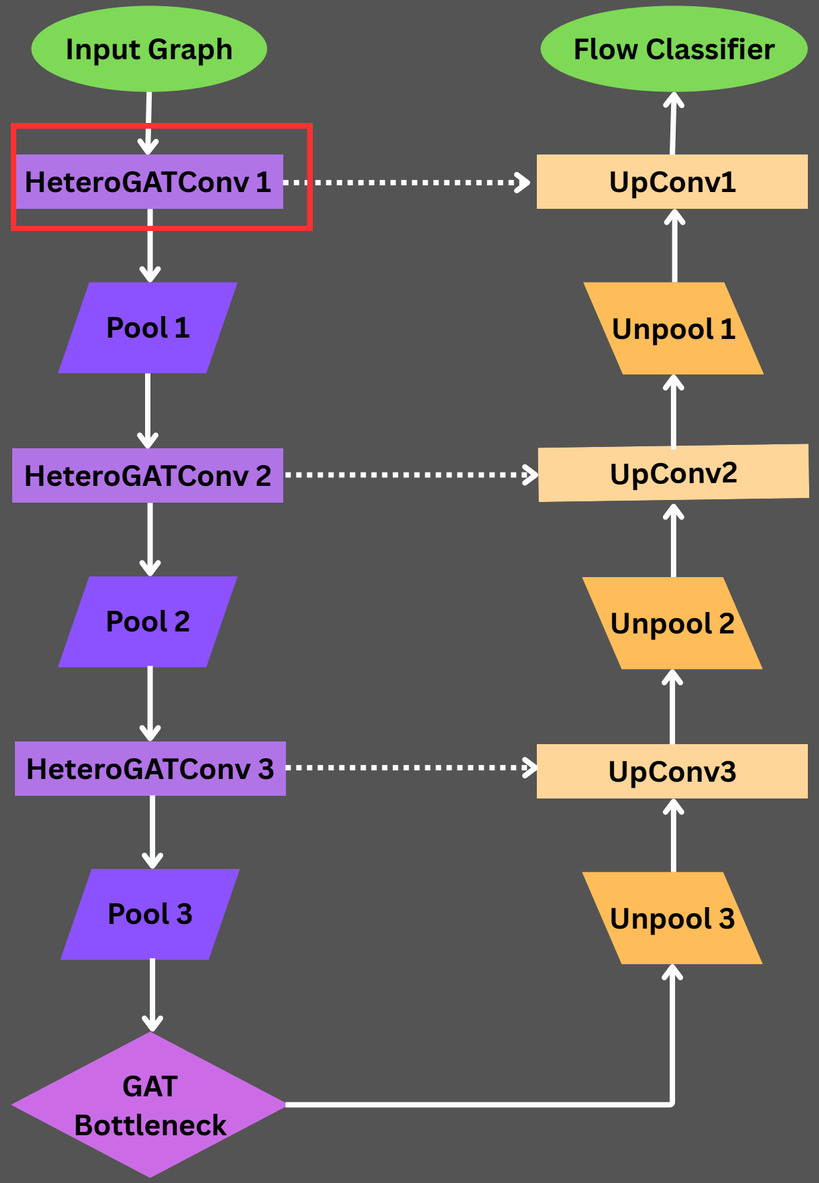

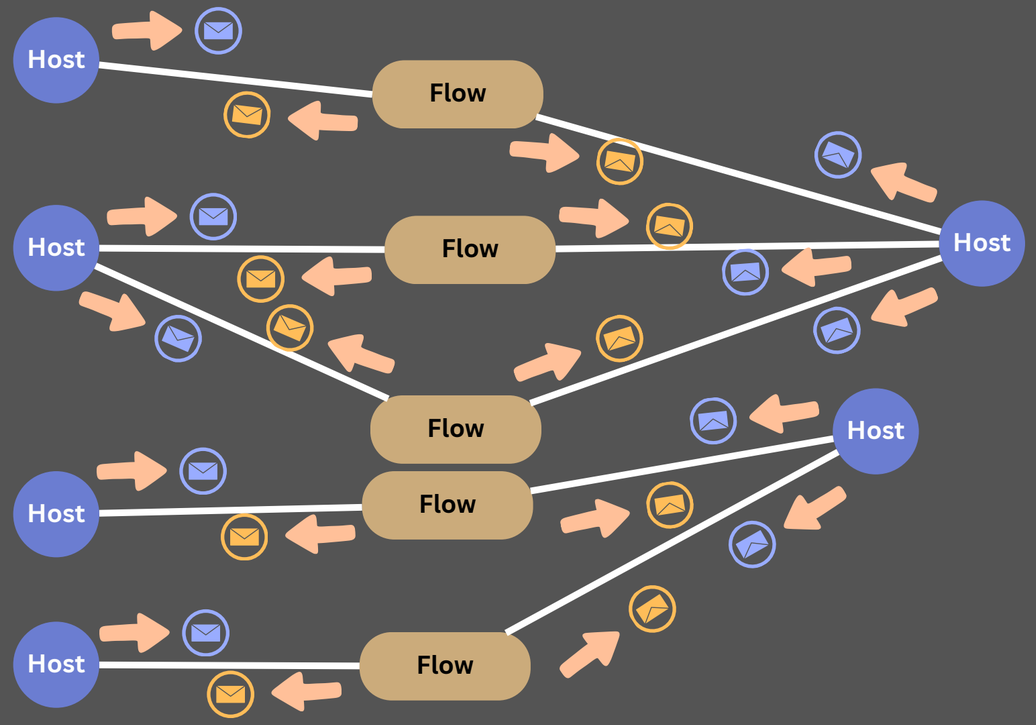

Heterogeneous Graph U-Nets Architecture

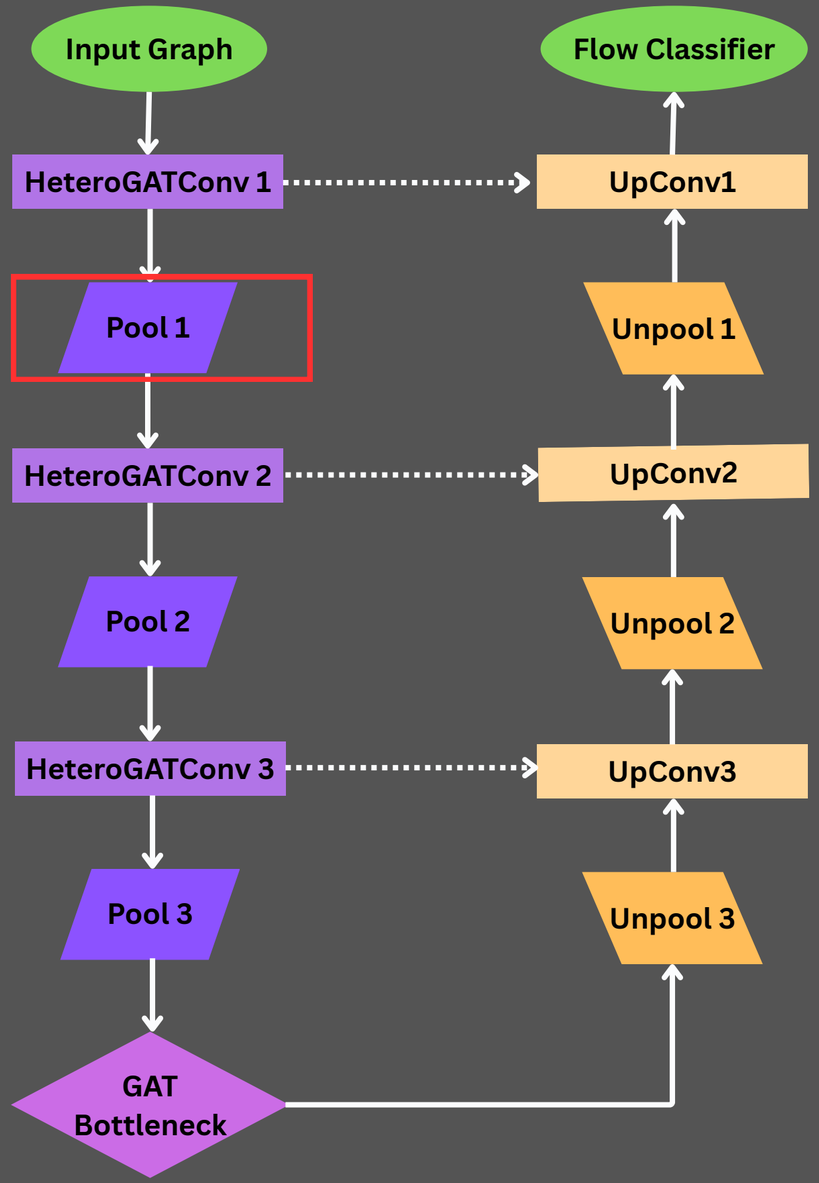

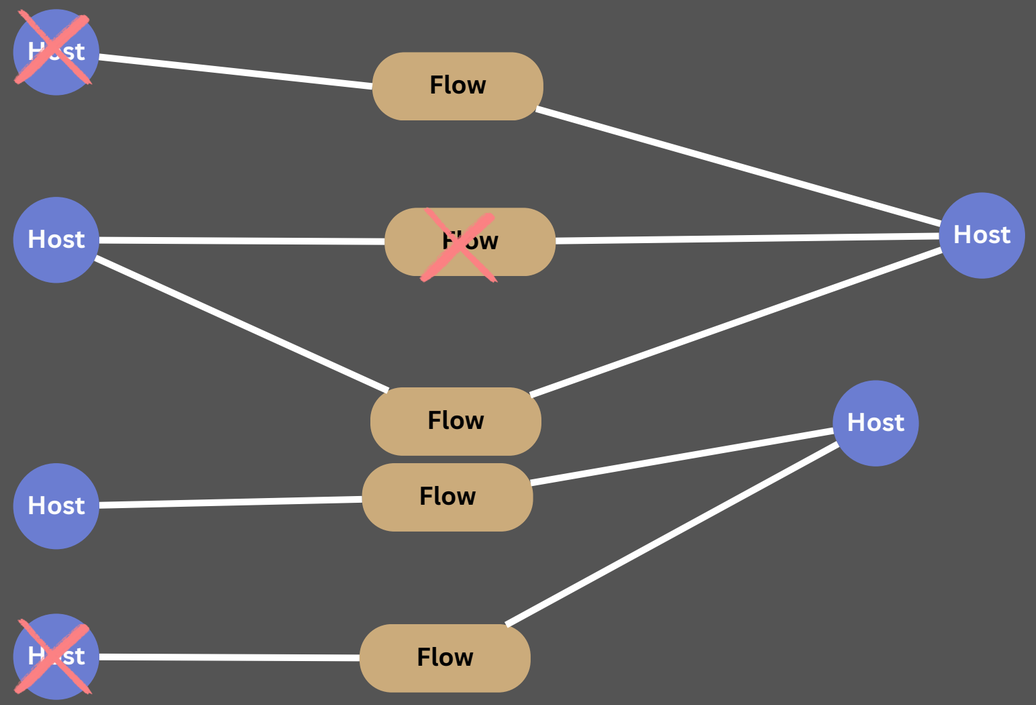

Heterogeneous Graph U-Nets Architecture

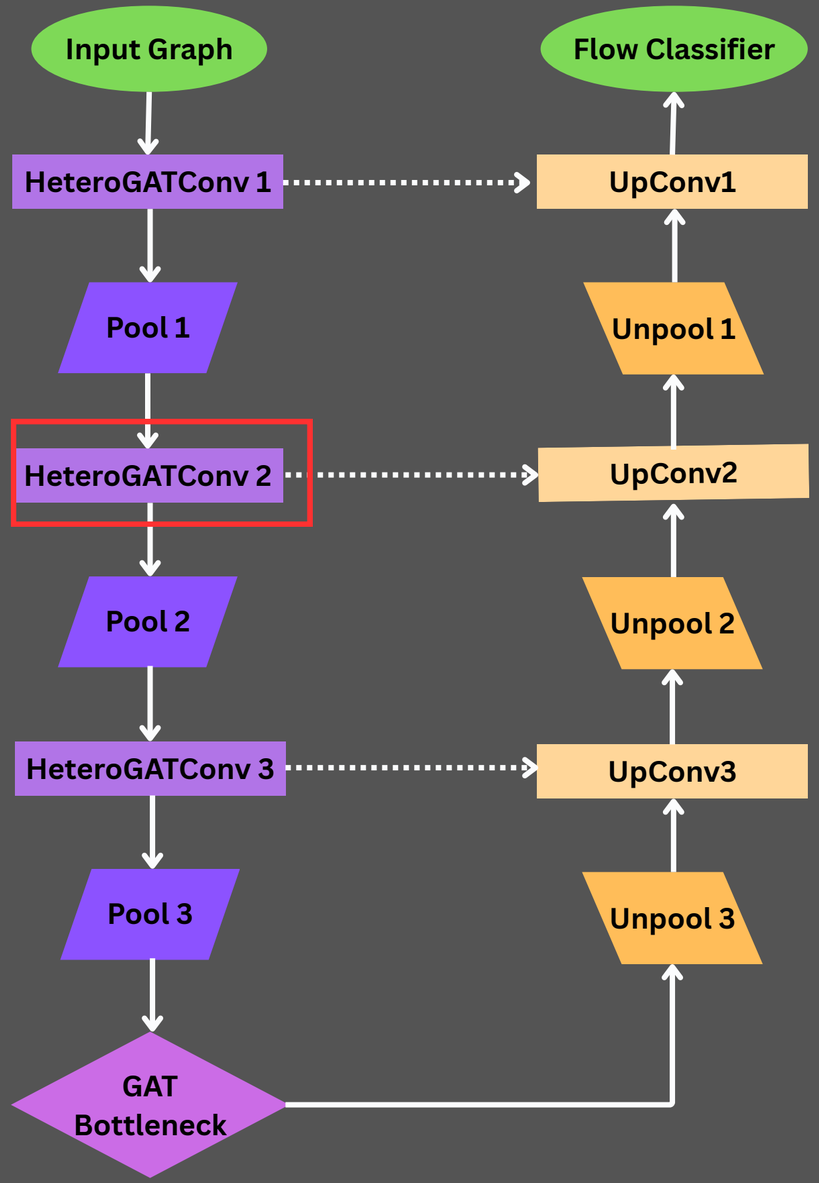

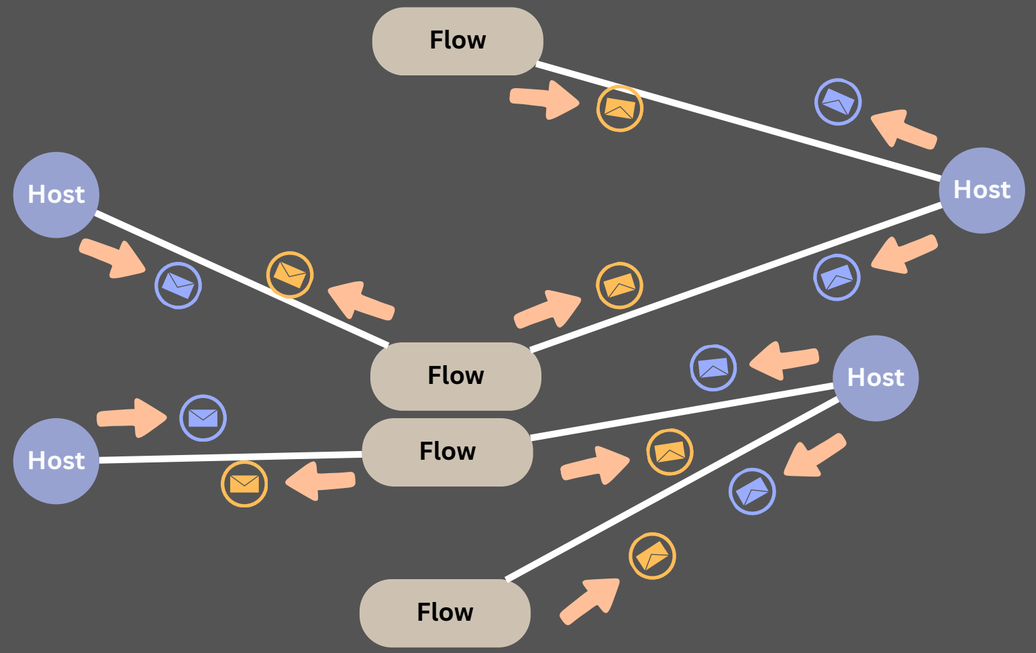

Heterogeneous Graph U-Nets Architecture

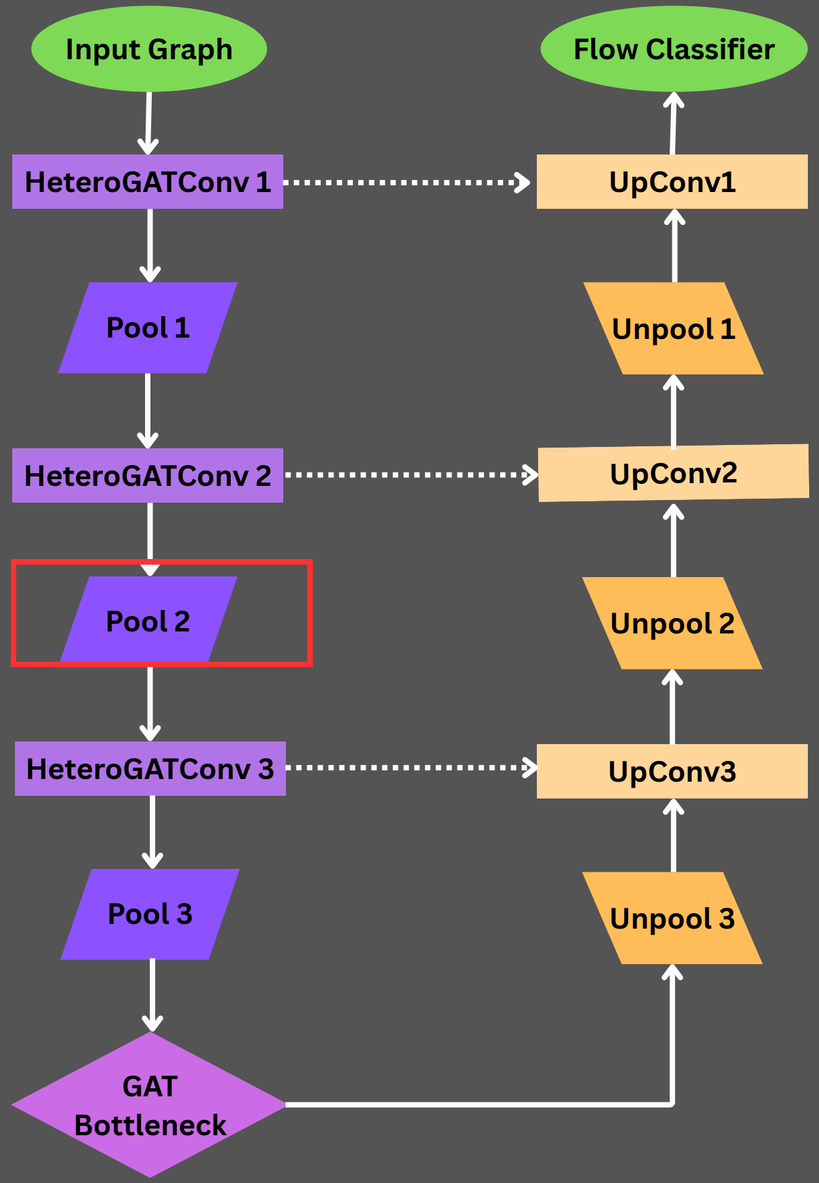

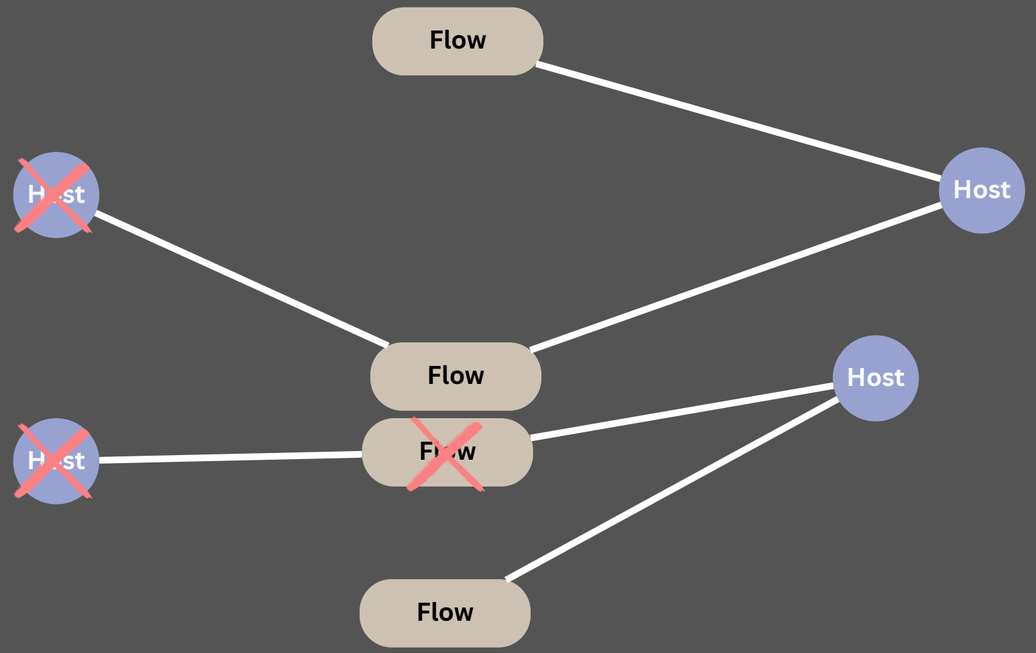

Heterogeneous Graph U-Nets Architecture

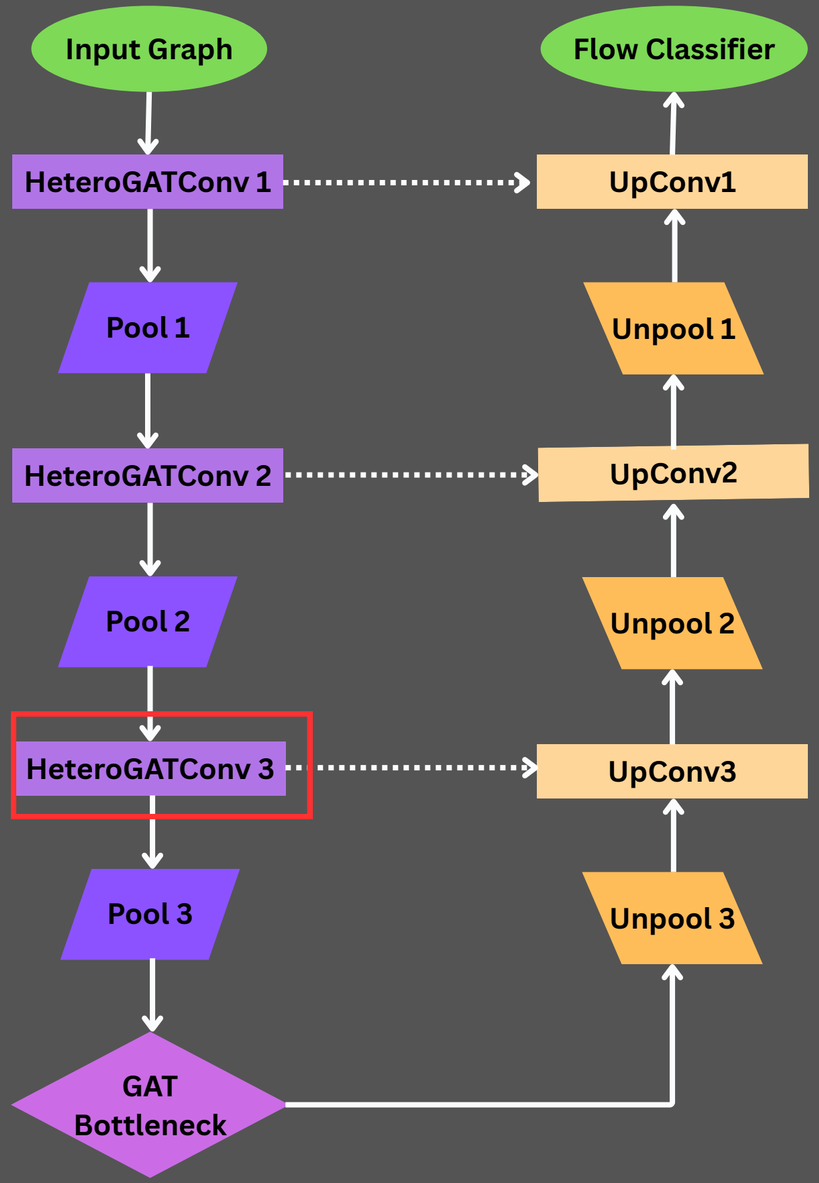

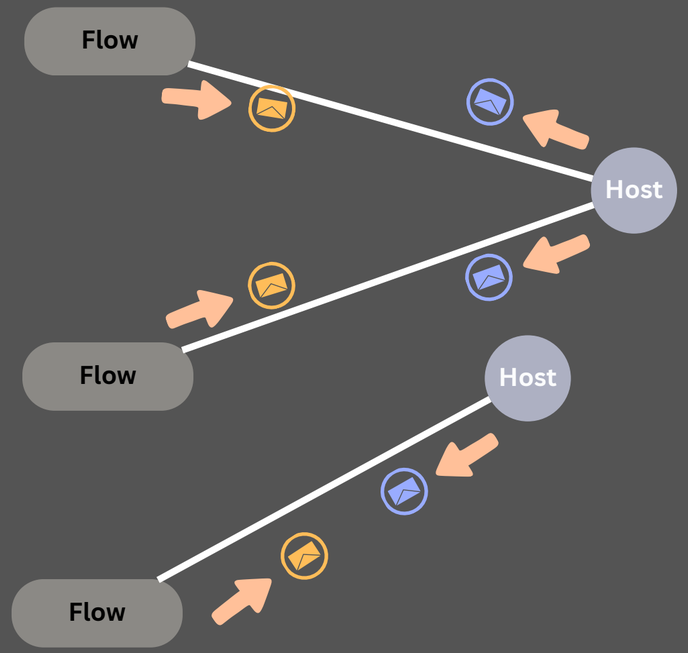

Heterogeneous Graph U-Nets Architecture

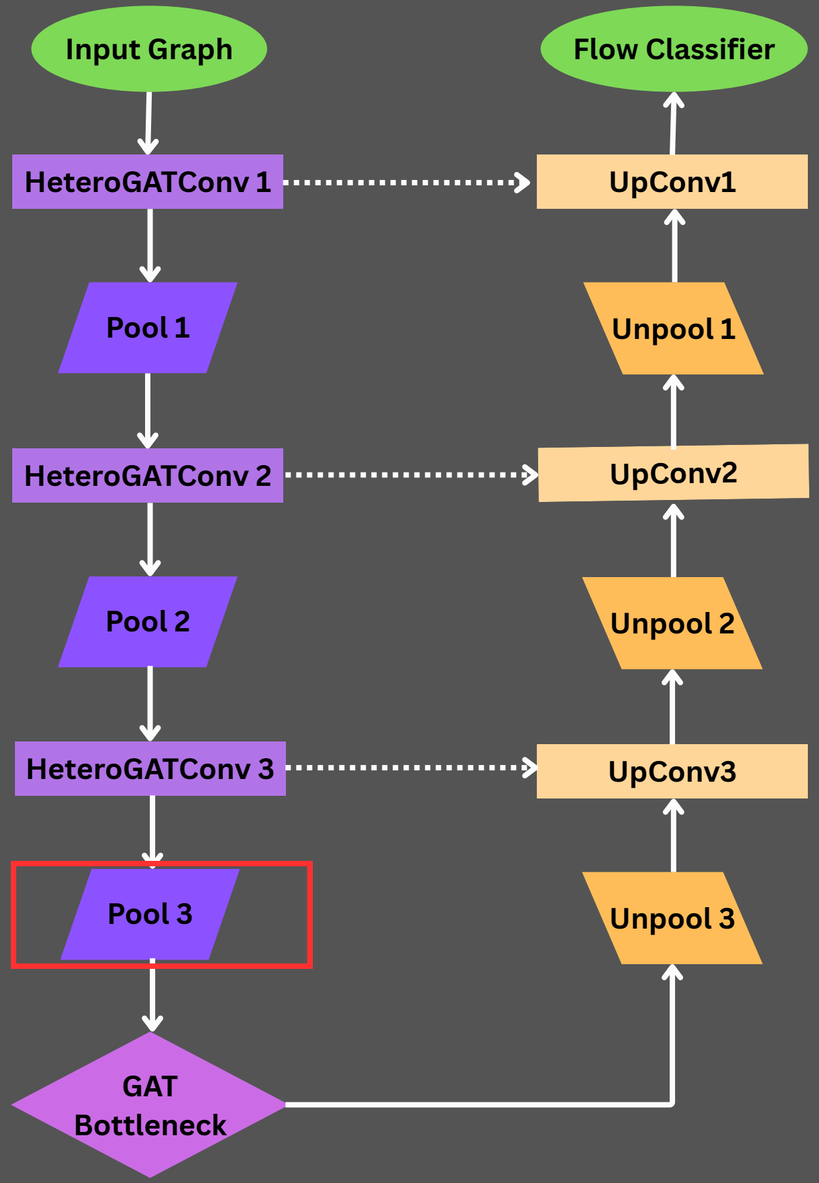

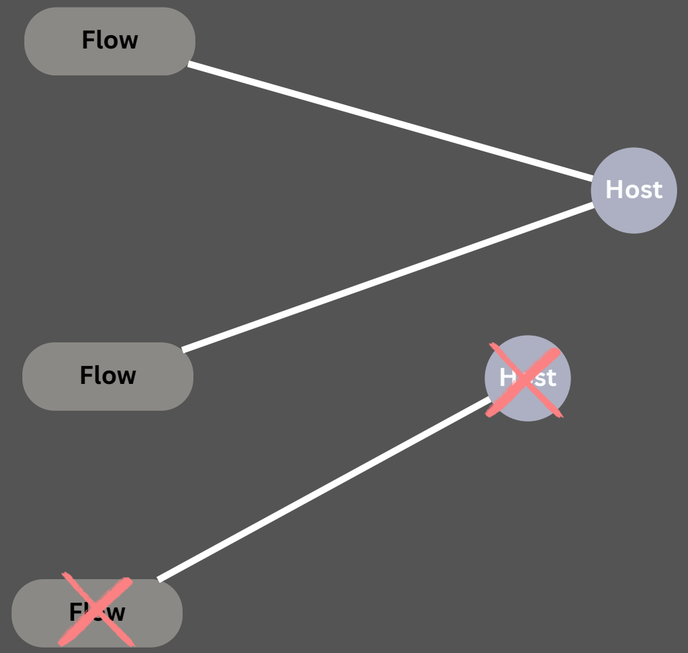

Heterogeneous Graph U-Nets Architecture

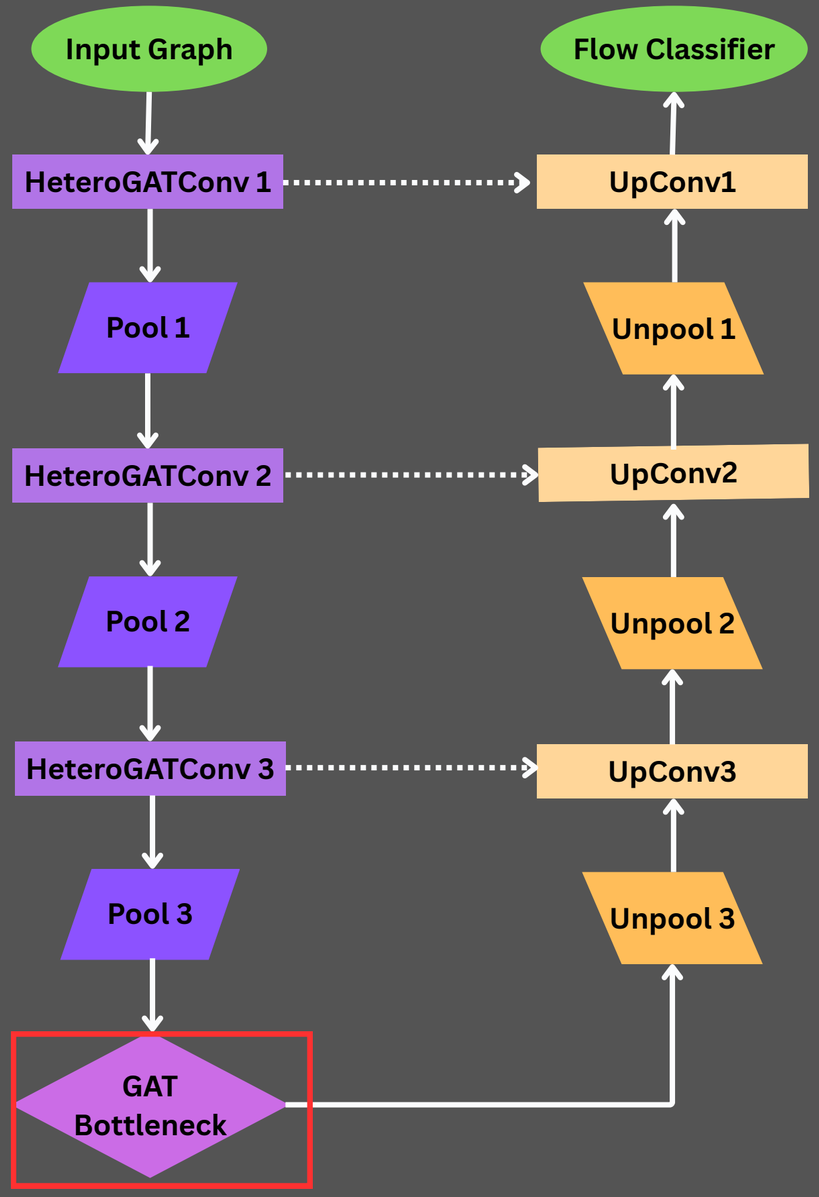

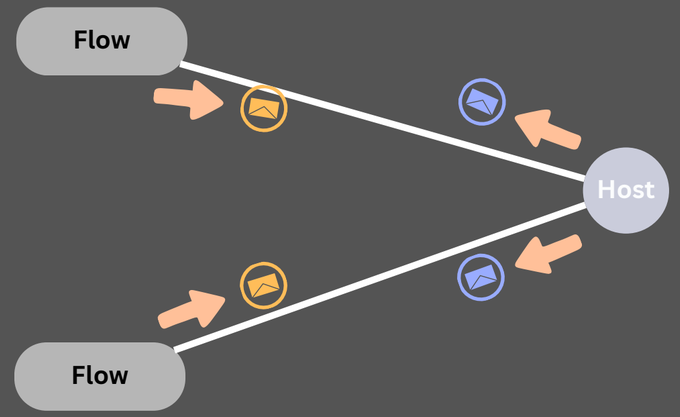

Heterogeneous Graph U-Nets Architecture

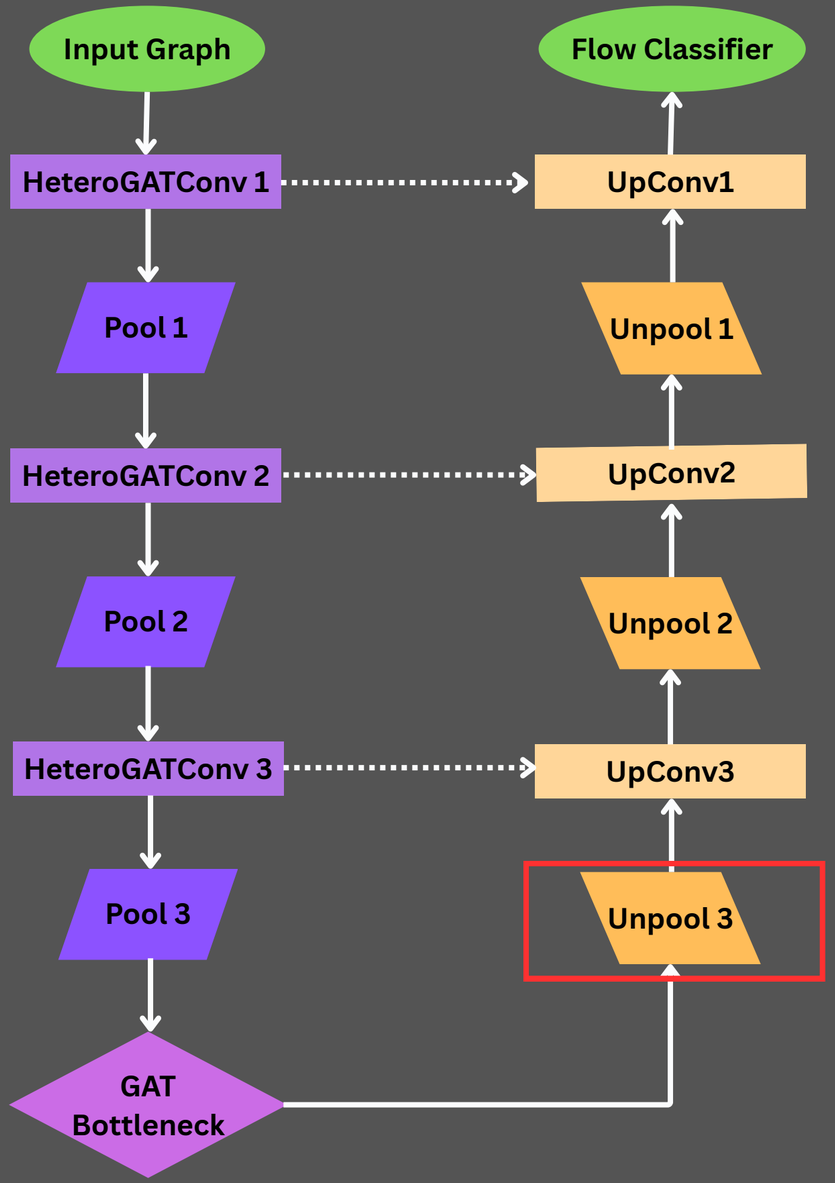

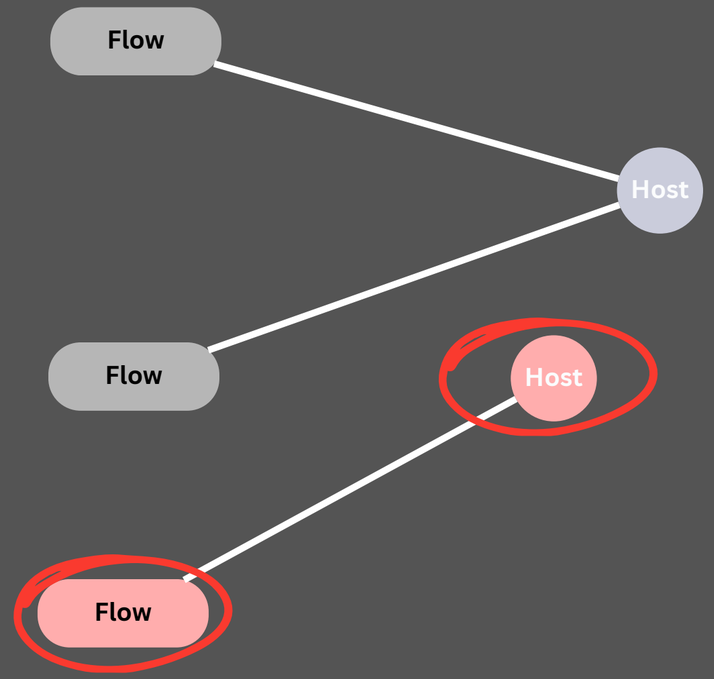

Heterogeneous Graph U-Nets Architecture

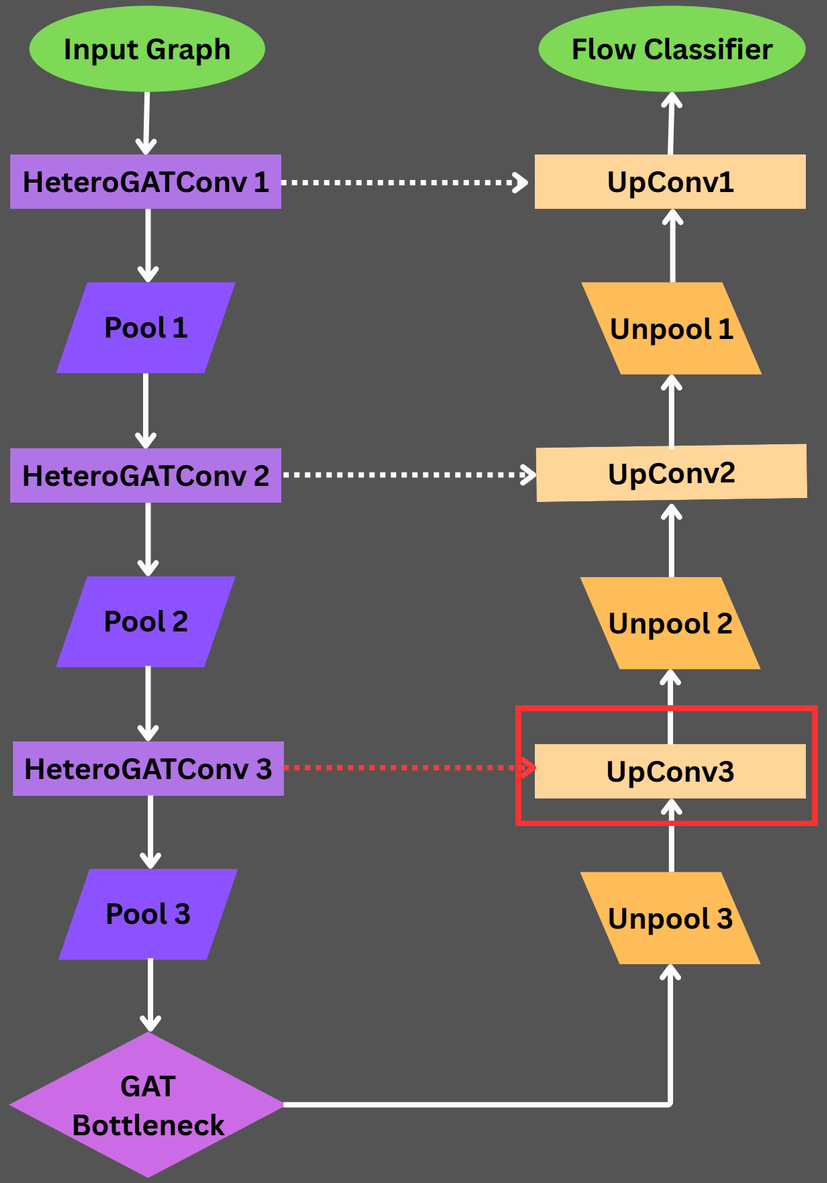

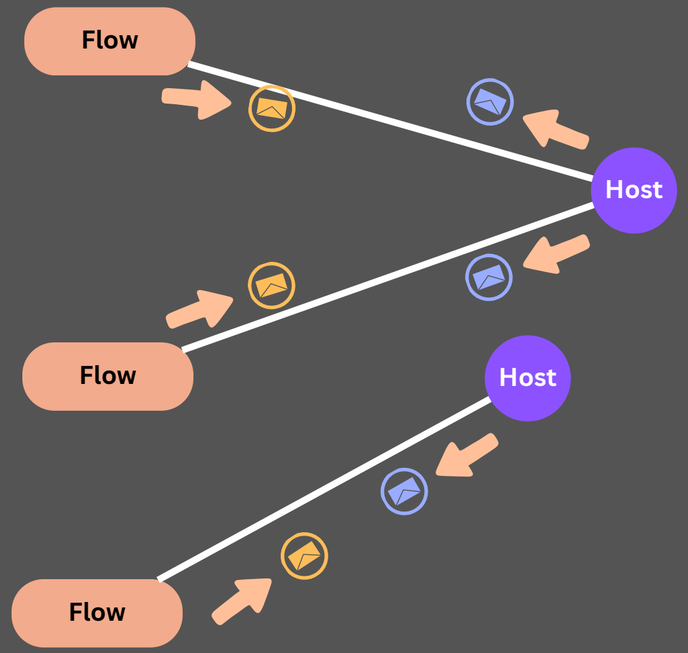

Heterogeneous Graph U-Nets Architecture

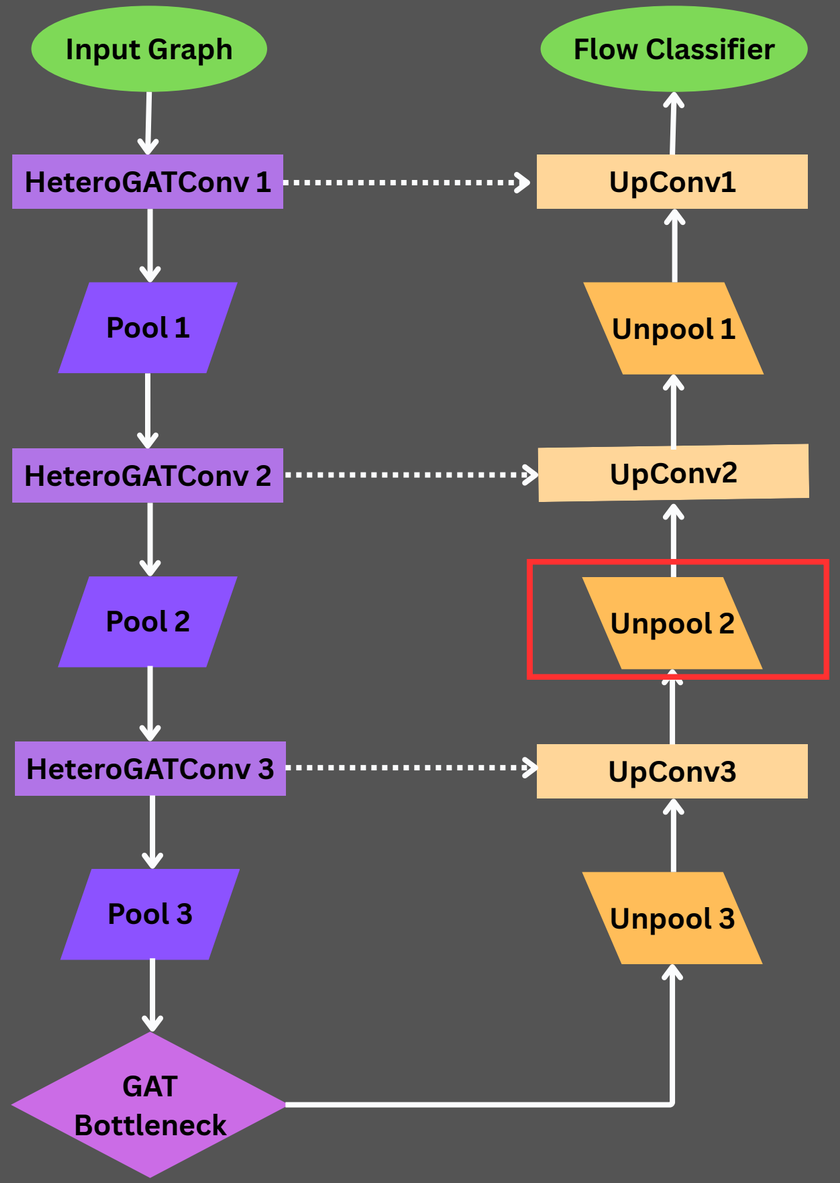

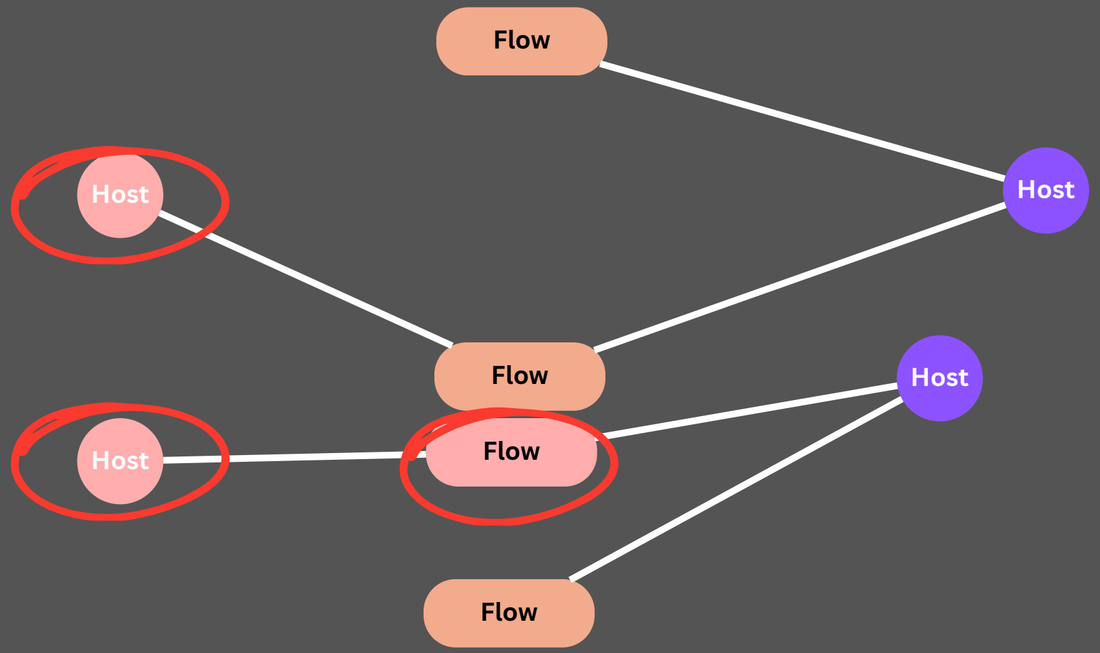

Heterogeneous Graph U-Nets Architecture

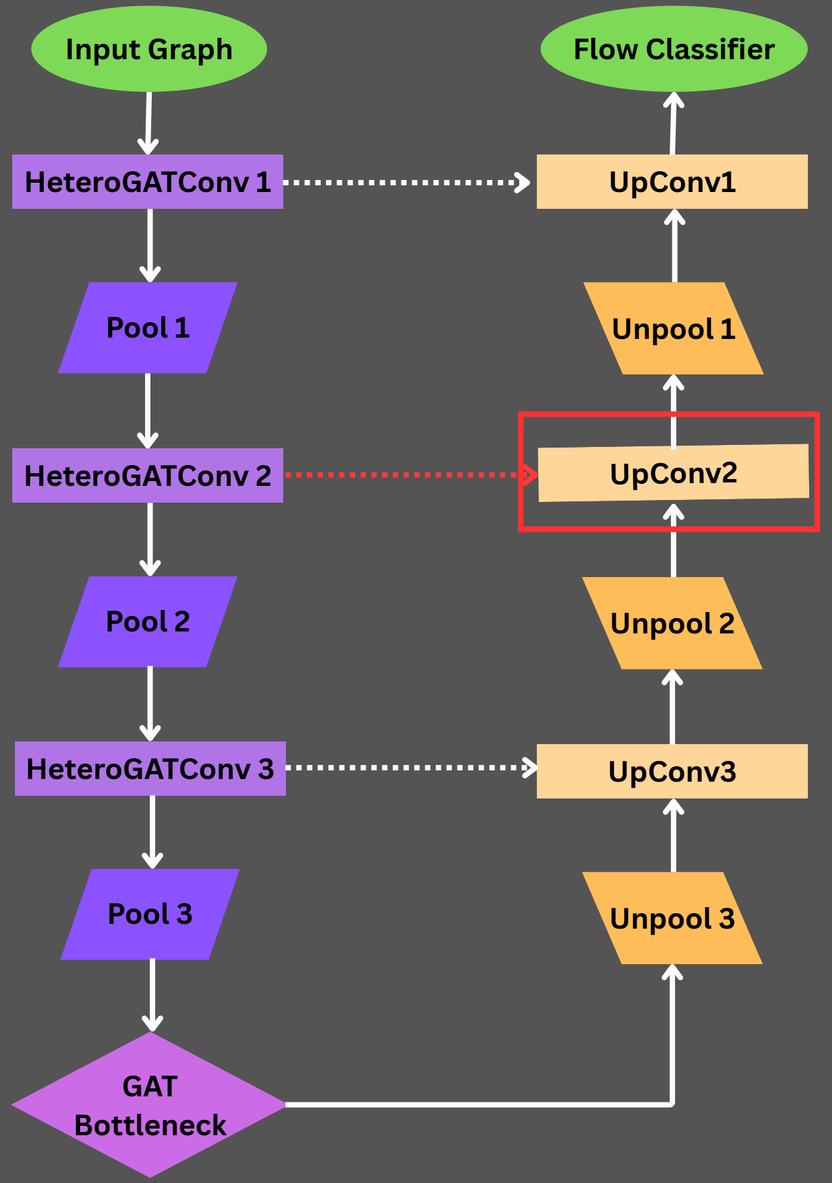

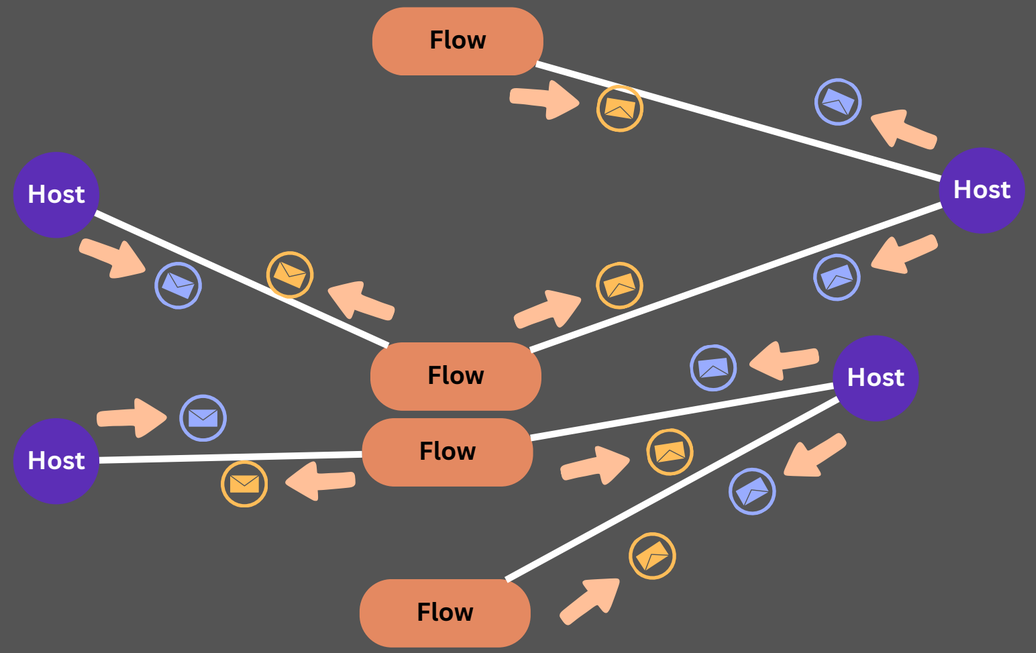

Heterogeneous Graph U-Nets Architecture

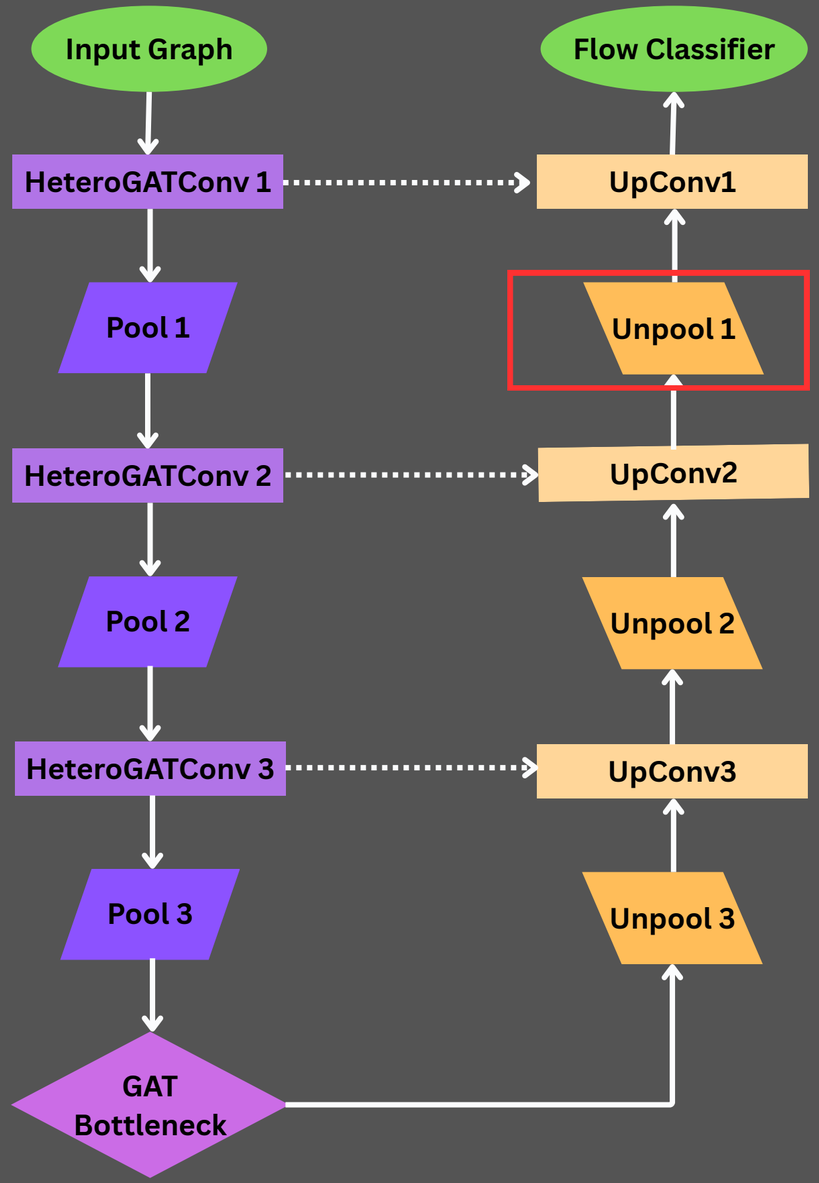

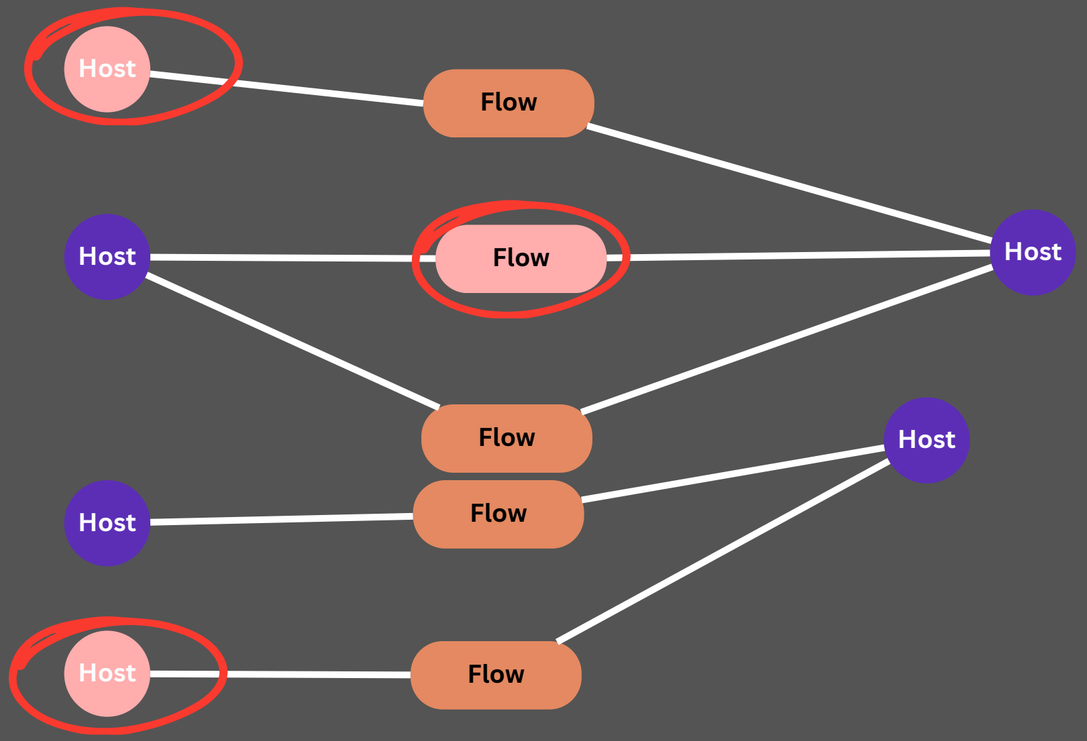

Heterogeneous Graph U-Nets Architecture

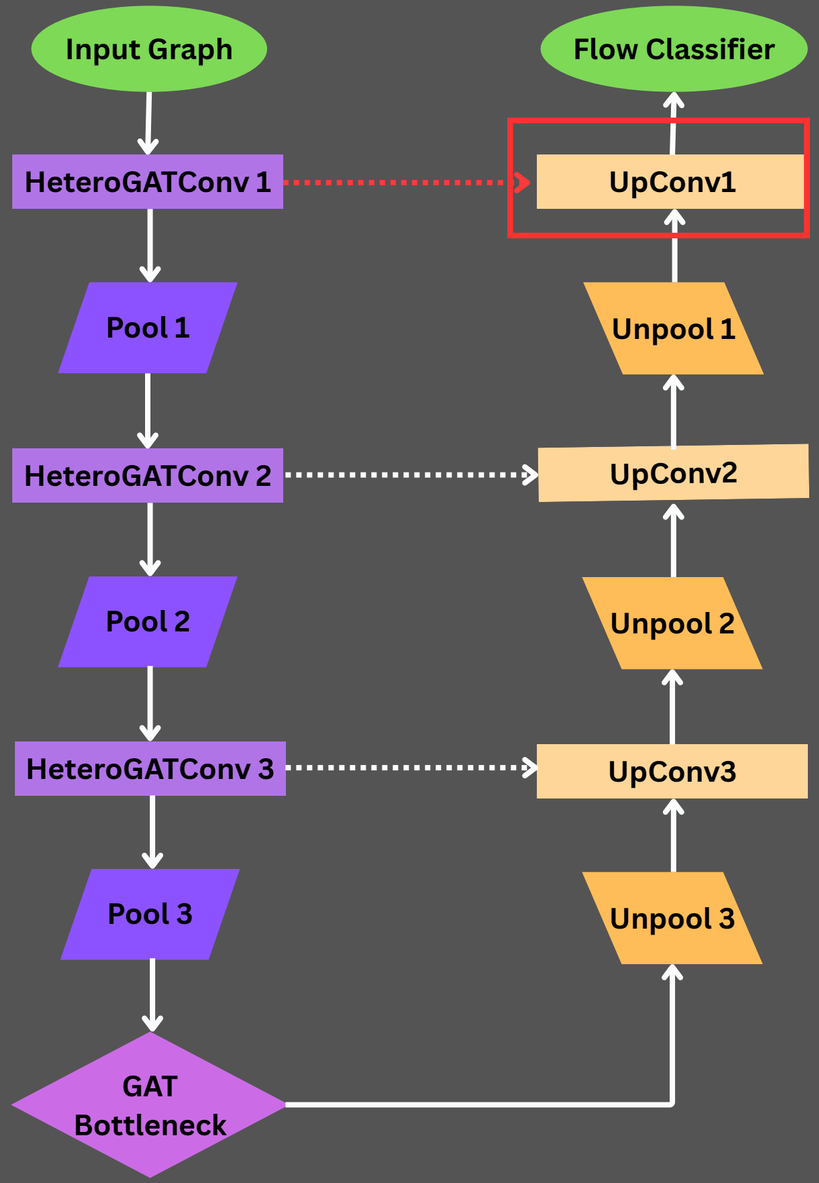

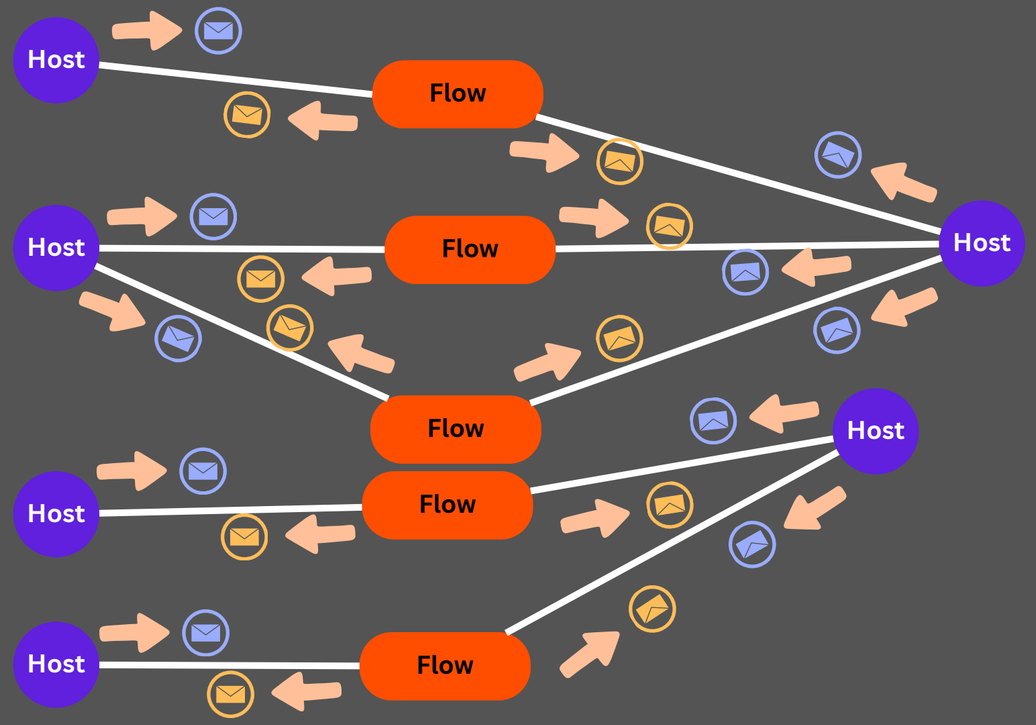

Heterogeneous Graph U-Nets Architecture

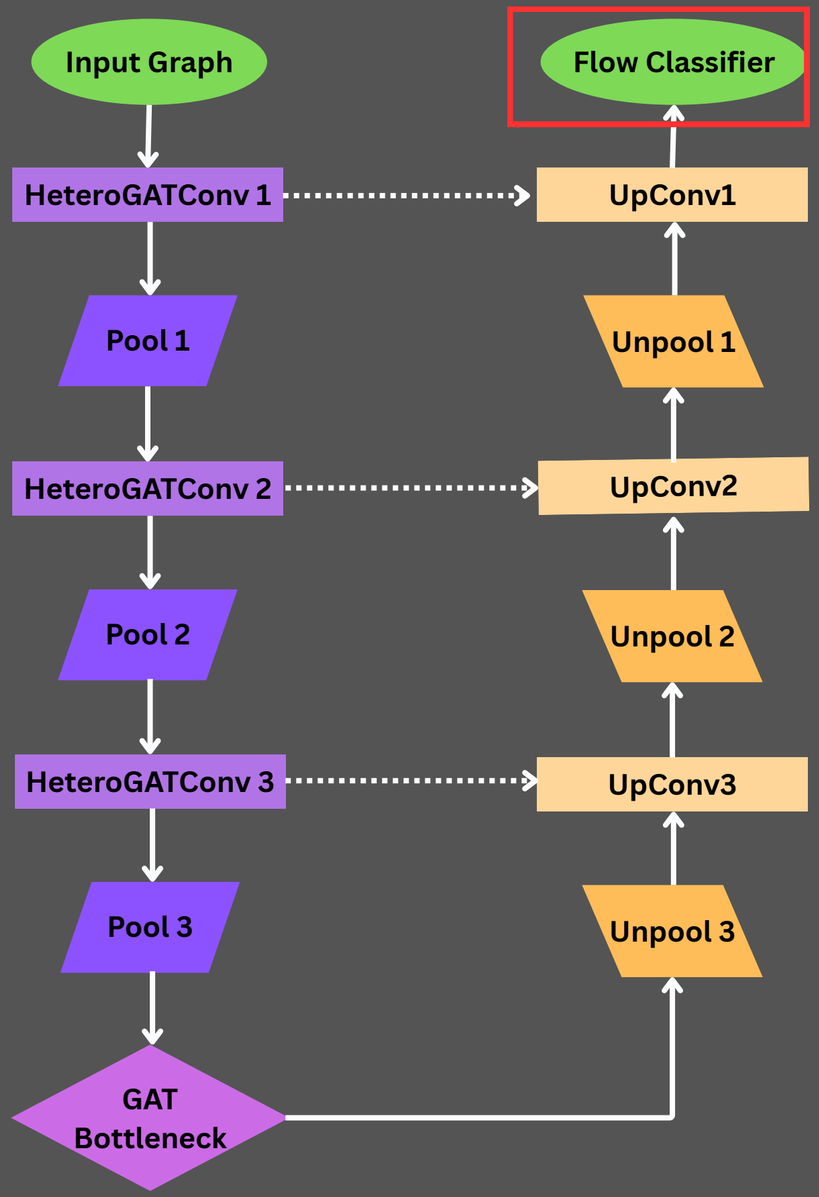

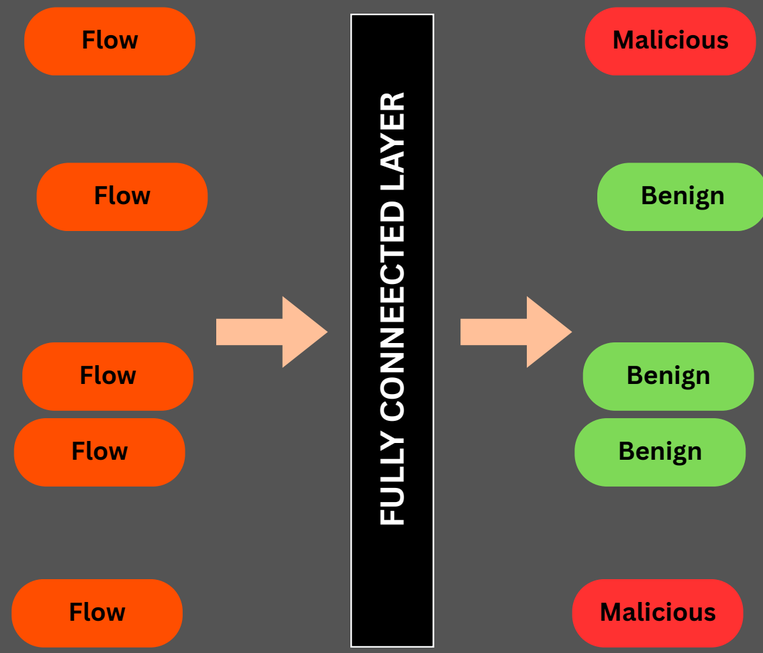

System Architecture & Implementation

Three-Stage Detection Pipeline

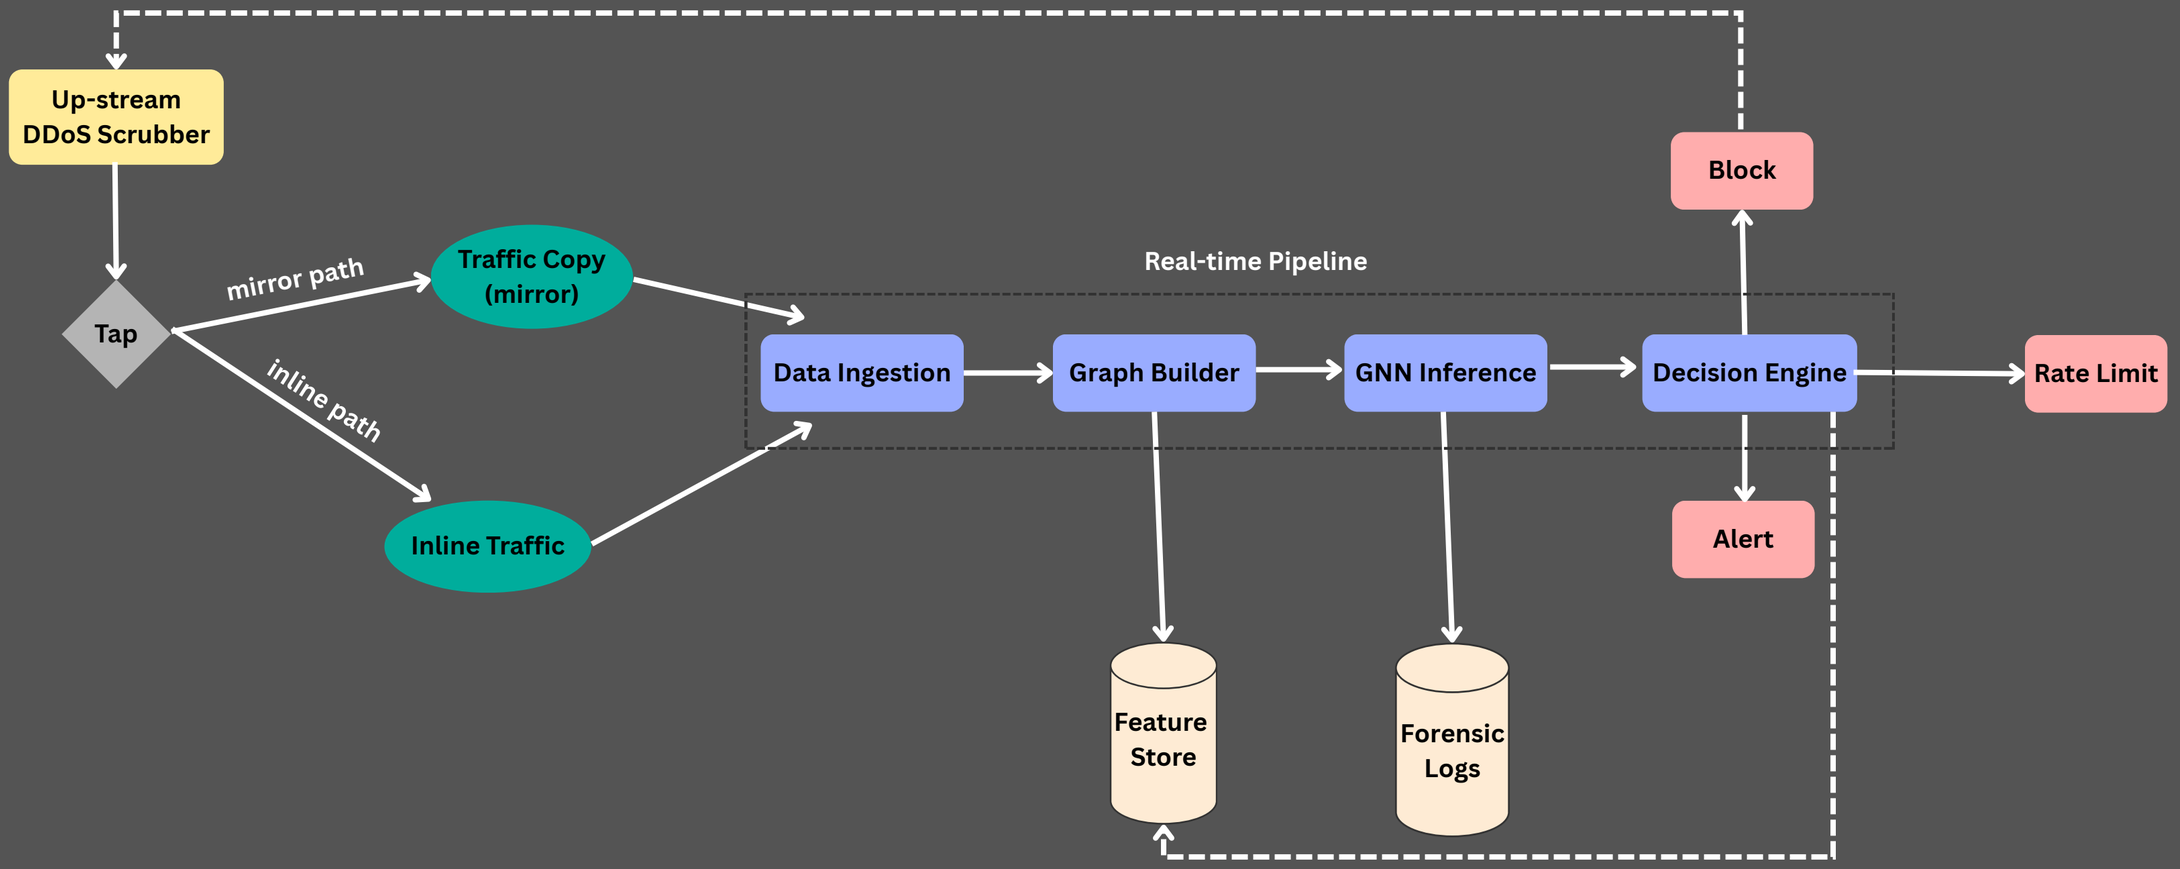

Three-Stage Detection Pipeline

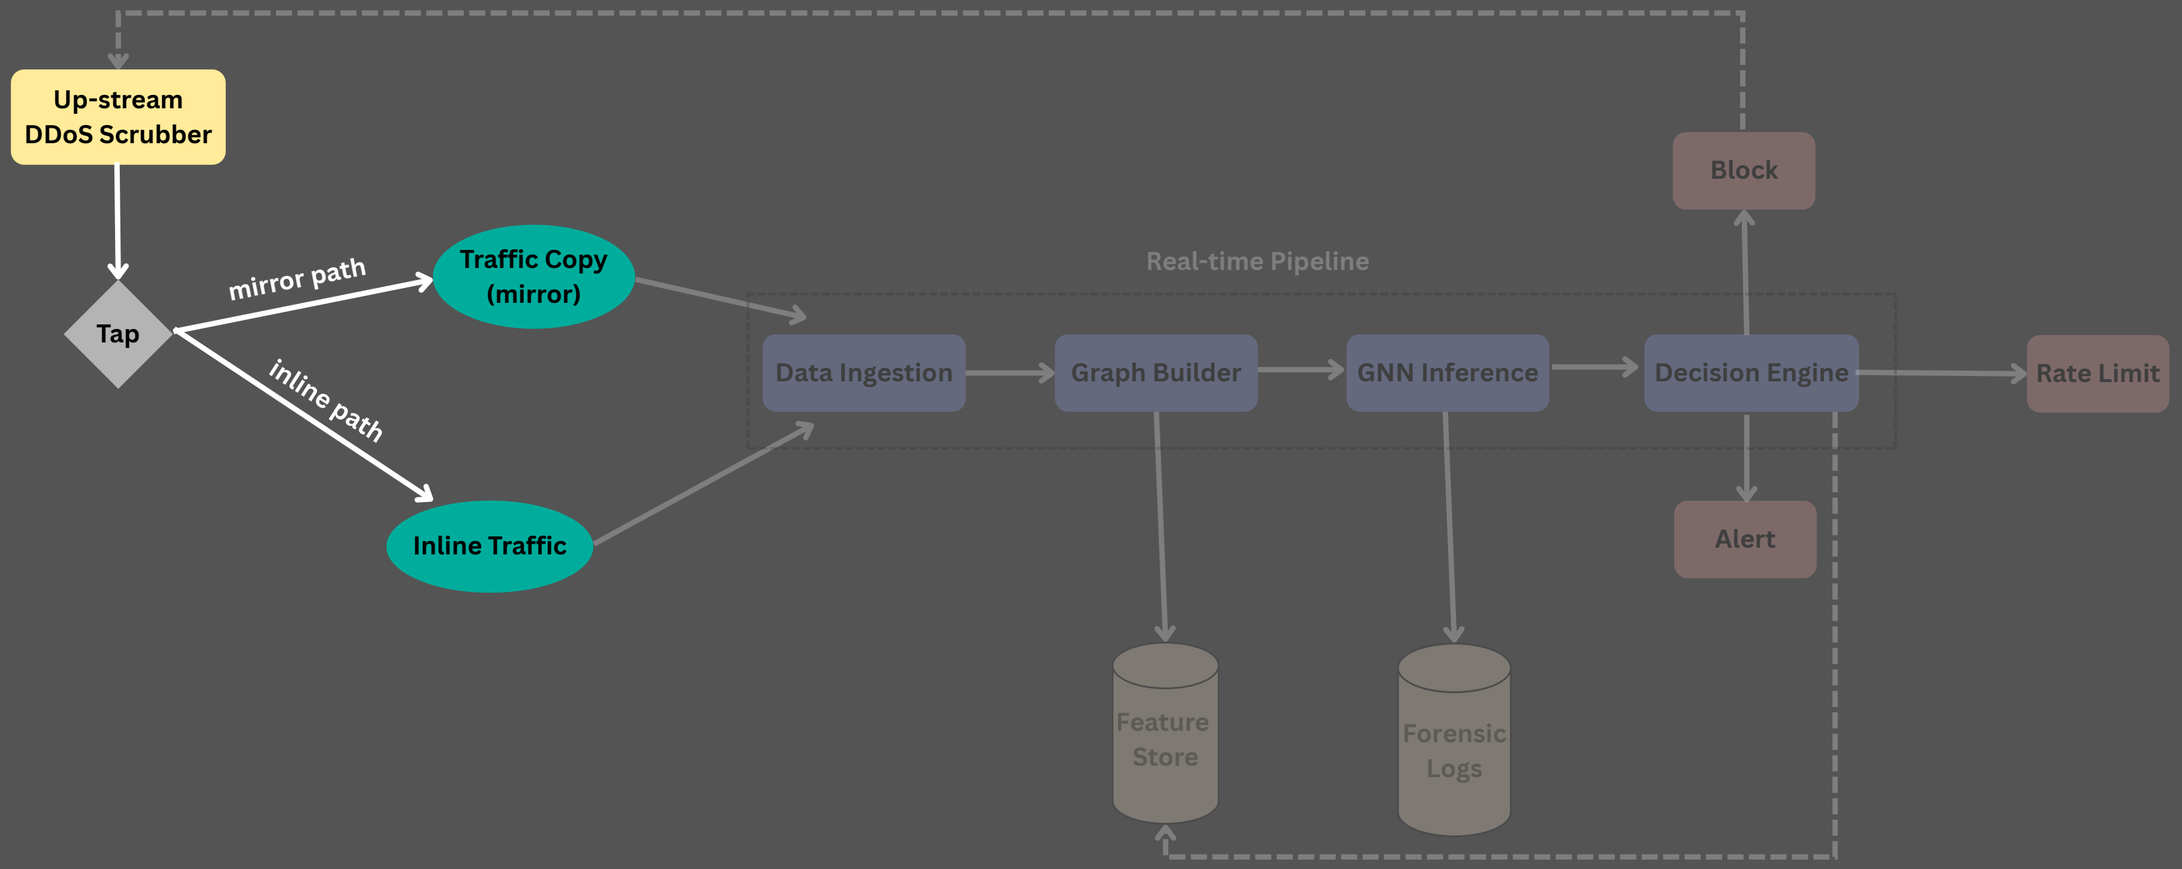

Three-Stage Detection Pipeline

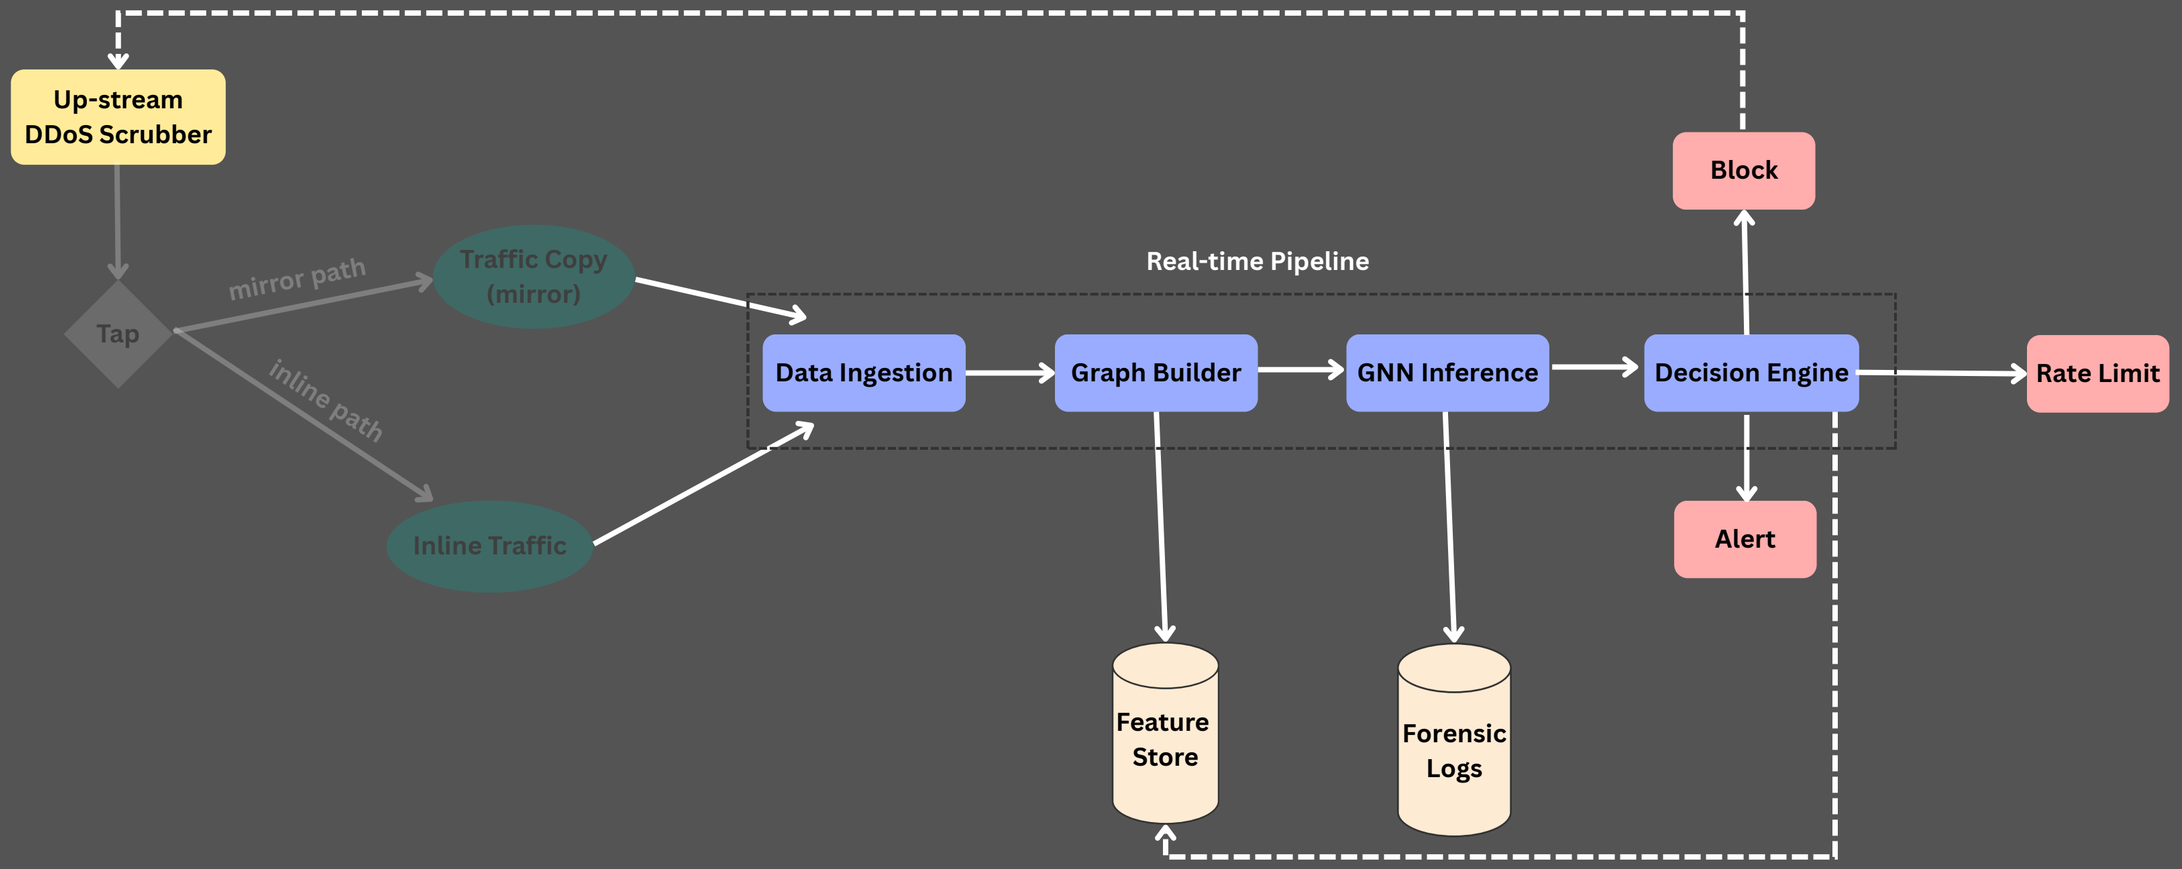

Integration in Data Centers

Deployment Modes

Inline Mode

- **Real-time enforcement**
- **Automatic blocking/rate-limiting**
- **Grey-zone analyst review**

Passive Mode

- **Zero added latency**
- **Auxiliary detection signal**
- **Complements existing systems**

Mitigation Actions

- **High confidence:** Automatic blocking
- **Medium confidence:** adaptive rate-limiting
- **Low confidence:** SOC Analyst review

Phased Deployment Strategy

PHASE 0

- **Alert-only mode**
- **Deploy on management network**
- **Build Trust**

PHASE 1

- **Automated rate-limiting**
- **Inter-DC links**

PHASE 2

- **Hard-blocking**
- **Tenant edges**
- **Recurring signatures back to upstream**

Experimental Results

Experiment Evaluation Setup

Three Diverse Datasets

Dataset	Total Flows	Attack Types	Environment
CIC IDS 2017	2.4M	LOIC based DDoS	Enterprise
CIC DDoS 2019	20M	11 attack families including DNS, LDAP, UDP-Lag, SYN	Mixed OS testbed
BCCC Cloud 2024	700K	17 TCP-SYN variants	AWS VPC

Baseline Comparisons

- **E-GraphSAGE: Edge-level classification approach using Graph Sage.**
- **GNN-RNIDS: Node-level classification using a custom GNN.**

Performance Results

CIC Datasets (2017 & 2019)

- All models achived near perfect performance (>99.9% across all metrics)
- Clean well-curated patterns are easily detected.

BCCC Cloud Dataset 2024

Model	Accuracy	Precision	Recall	F1 Score
Heterogeneous Graph U-Net	0.972 ± 0.001	0.985 ± 0.002	0.945 ± 0.010	0.960 ± 0.002
GNN RNIDS	0.957 ± 0.005	0.946 ± 0.016	0.982 ± 0.018	0.940 ± 0.007
E-GraphSage	0.947 ± 0.014	0.922 ± 0.029	0.994 ± 0.010	0.914 ± 0.021

Limitations, Future Directions & Q&A

Current Limitations & Future Directions

Limitations

- **Training Data:** Public datasets may not capture all DC-specific patterns
- **Scalability Boundaries:** >100K concurrent flows requires distributed implementation
- **Feature Engineering:** Still requires domain expertise for optimization
- **Adversarial Testing:** GNN-specific attacks not yet evaluated
- Need **inference-time profiling** & **ablation studies**

Future Research Directions

- **Online/continual learning** for evolving threats
- **Federated GNNs** across datacenters
- **Hardware acceleration (GPU/FPGA)**
- **Encrypted traffic adaptation** & **zero-day testing**
- **Ablation studies t**o optimize architecture
- **Adversarial robustness testing**

Key Takeaways

- **Heterogeneous Graph U-Nets** excels in DDoS detection.
- **Superior Performance:** 96% F1 score on realistic cloud traffic
- **Practical Deployment:** Minimal disruption, phased rollout
- **Multi-scale Detection:** Catches both local and distributed attacks
- **Production Ready:** Sub-millisecond latency, horizontal scaling
- **Future Proof:** Adaptable to emerging threats

Github Link: ***github.com/kartikeyas00/heterogeneous-graph-unets-ddos***

Questions?

Thank You!