International Conference on Computing, Networking and Communications (ICNC 2026)

# RAN–IDS: Adaptive Continual Learning for Intrusion Detection in Open RAN

*Nadia Niknami*

*Villanova University*

*Jie Wu*

*Temple University*

# Outline

Background & Motivation

Key Idea

Idea Overview

Evaluation & Conclusion

# Background

## Open-RAN

- Disaggregated RU / DU / CU architecture
- Near-RT and non-RT RIC control
- Multi-vendor, software-defined deployment

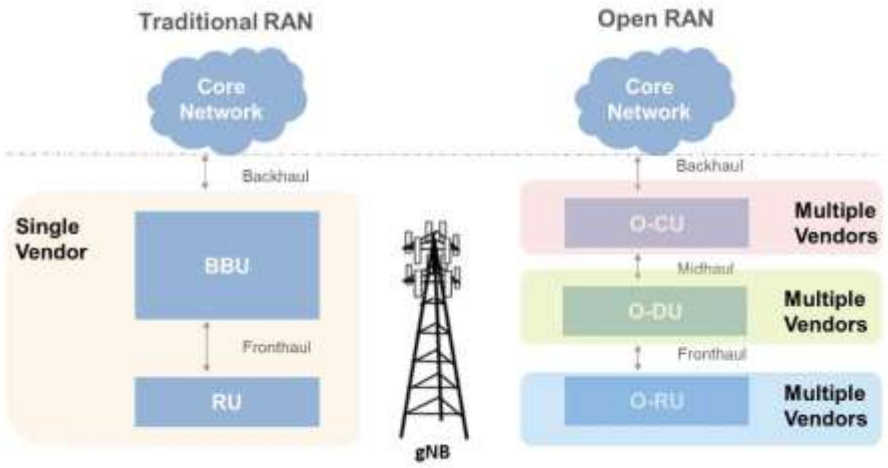


## Intrusion Detection System (IDS)

- Monitor network traffic for malicious behavior
- Typically trained offline on historical data
- Designed under relatively stable traffic distributions

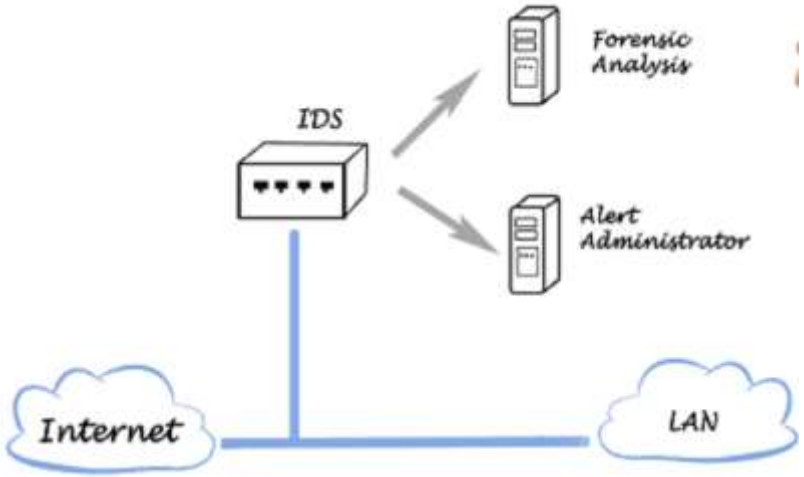

# Motivation

**Why O-RAN Breaks IDS?**

- Frequent RIC policy updates
- Dynamic slice scaling and reconfiguration
- Vendor-specific functional variations
- Continuous concept drift in control and data planes

**Why Existing IDS Cannot Cope?**

- Designed for stationary distributions
- Centralized retraining is too slow
- Naïve online updates → catastrophic forgetting

How can we design an IDS that adapts continuously without losing past knowledge for O-RAN?

# Key Idea: RAN–IDS

**A Continual Learning Framework Designed for O-RAN**

- Detect zero-day and evolving attacks
- Adapt continuously without catastrophic forgetting
- Slice-aware and interface-aware detection
- Hierarchical teacher–student architecture
- Operates under O-RAN latency constraints

**Key Idea:**
**Transforms IDS from static classification to architecture-aware adaptive learning**

# IDS Placement in O-RAN

- Near-RT RIC: control-plane monitoring
  - DU–CU midhaul: data-plane visibility
  - Backhaul: traffic toward core network

- Different interfaces exhibit distinct traffic dynamics and drift patterns

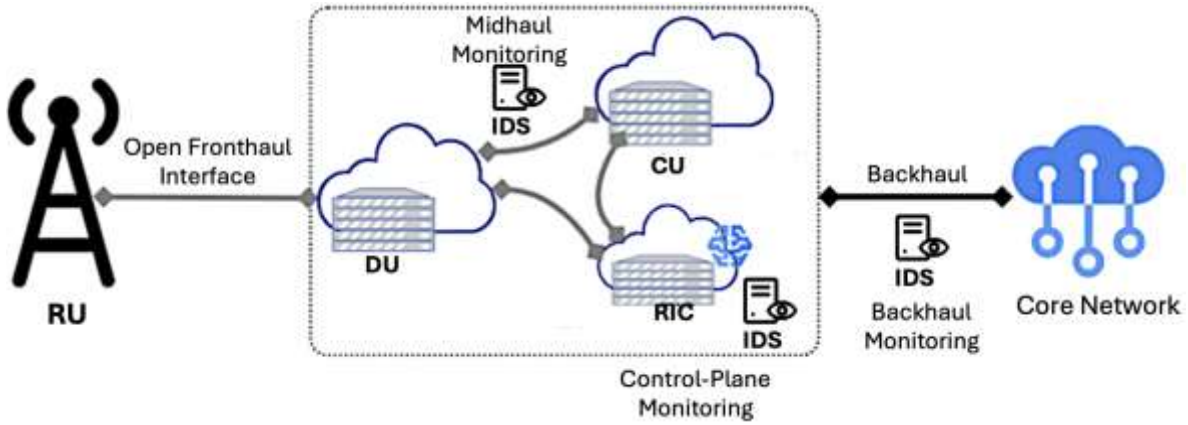

# Idea Overview

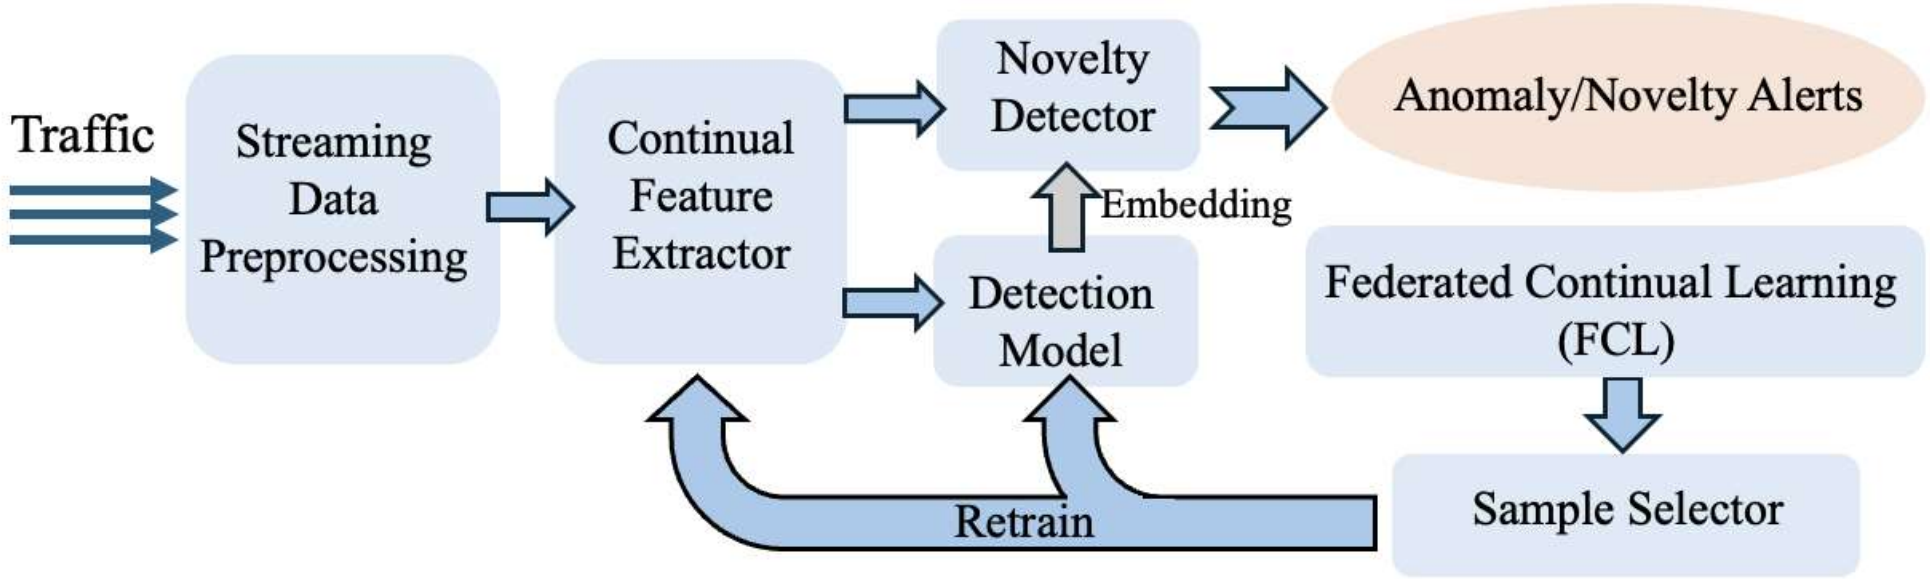


- Designed to align continuous traffic drift with controlled model adaptation.
- Combines streaming detection, replay-based CL, and hierarchical knowledge transfer.

# Idea Overview (Streaming & Novelty)

- Streaming preprocessing (no train/test split)
  - Continual feature extractor
  - Detection model and novelty detector

- Novelty Detector
  - Uses embeddings and confidence scores
  - Detects unseen and zero-day attacks
  - Slice-specific drift awareness

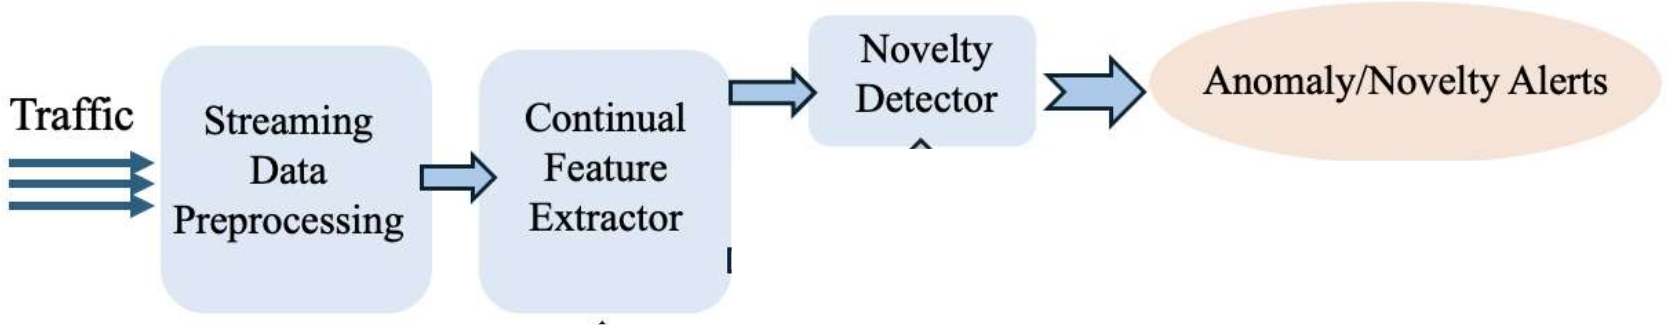


Operates fully online to match dynamic slice reconfiguration

# Idea Overview (Detection & FCL)

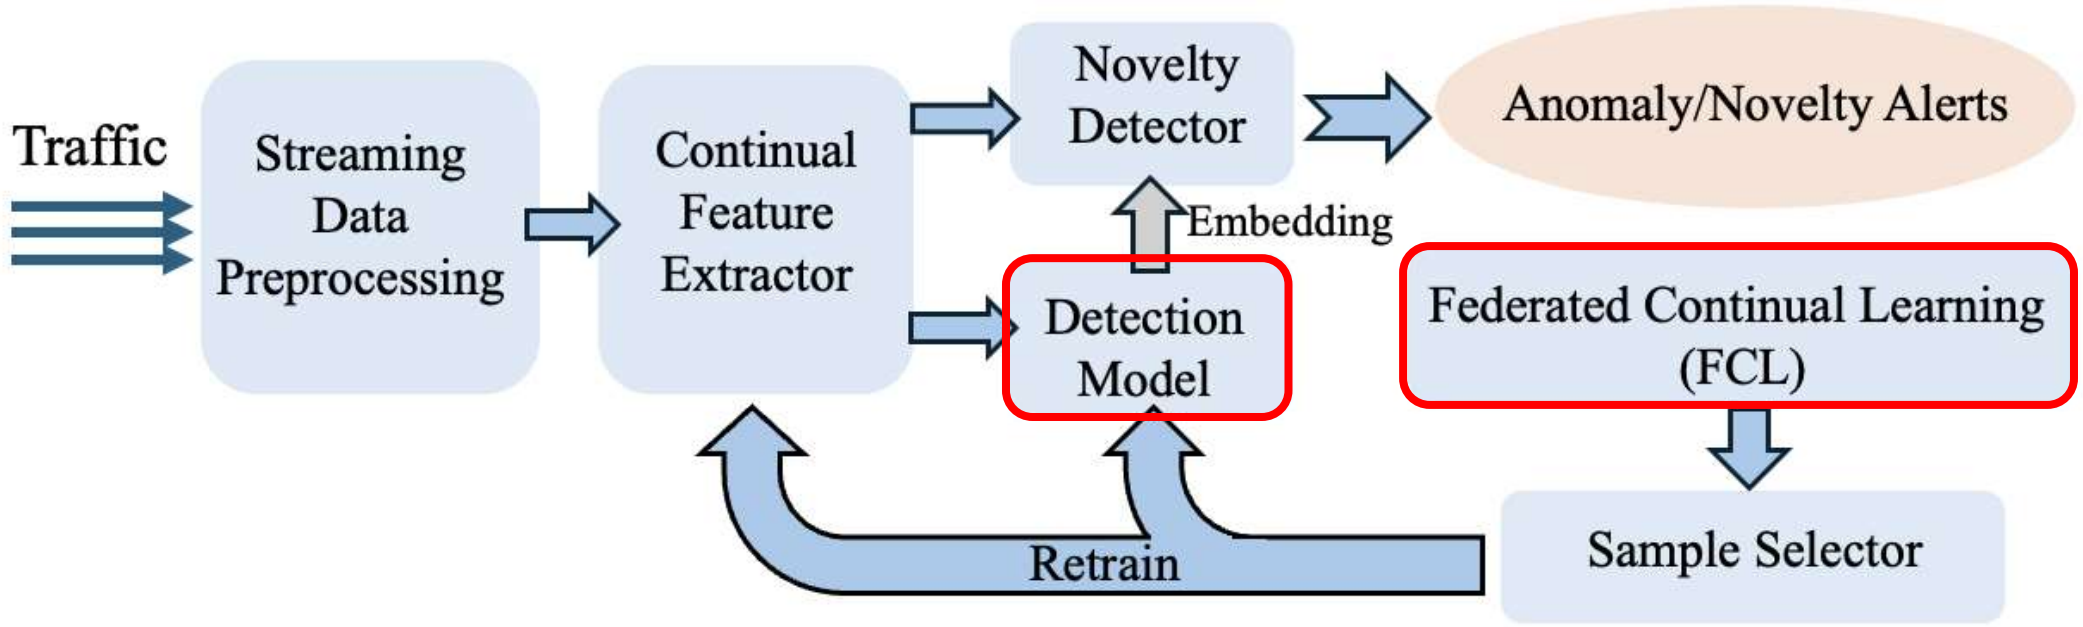


- FCL enables distributed adaptation without sharing raw traffic across vendors.
- Aggregates drift information across heterogeneous O-RAN nodes

# Continual Learning Loop

- Federated continual learning (no raw traffic sharing)
  - Replay buffer with sample selection
  - Mitigates catastrophic forgetting
- Replay and sample selection ensure stability under non-stationary traffic.
- Enables separation of global stability and local responsiveness.
- Essential under slice scaling and RIC policy drift.

# Teacher–Student

- The teacher–student hierarchy mirrors the O-RAN functional split (non-RT RIC → near-RT RIC / DU-CU)

- Teacher captures long-term global behavior
  - Students ensure low-latency inference
  - Separates learning capacity from deployment constraints

# Experimental Setup

Dataset:

- 5G-NIDD (realistic 5G traffic)
  - 9 sequential attack tasks
  - Streaming continual learning setting
  - Metric: F1 score

Baselines for Comparison

- ER, AGEM, LwF, EWC
  - AOC-IDS
  - Static non-adaptive model

# Evaluation: Overall Performance

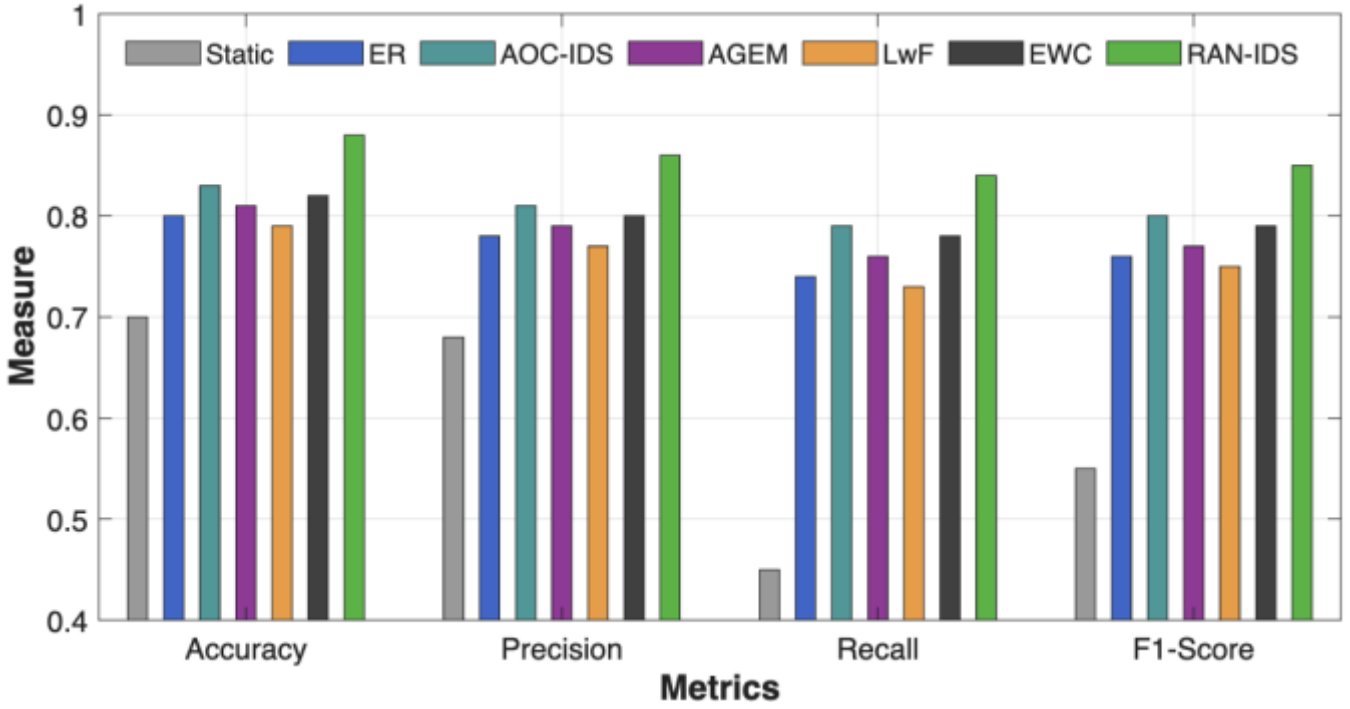


- RAN–IDS achieves highest F1, precision, recall
  - Static and naive CL degrade under drift
  - Hierarchical CL shows strongest robustness

# Evaluation: Performance and Ablation Study

TABLE I
ZERO–DAY DETECTION PERFORMANCE ON 5G–NIDD

| Method | F1 (%) | Recall (%) | Precision (%) |
|---|---|---|---|
| PCA | 58 | 50 | 68 |
| OC–SVM | 57 | 46 | 74 |
| Isolation Forest | 63 | 62 | 63 |
| Autoencoder (AE) | 67 | 64 | 71 |
| RAN–IDS (Ours) | **80** | **77** | **83** |

- Compared with PCA, OC-SVM, IF, AE
  - RAN–IDS achieves 80% F1
  - Significant improvement over baselines
- Removing CL causes major degradation
  - Replay is critical for memory retention
  - Sample selection improves forward transfer

TABLE II
ABLATION STUDY: IMPACT OF CONTINUAL LEARNING MECHANISMS

| Method | AVG (%) | FwdTrans (%) | BwdTrans (%) |
|---|---|---|---|
| RAN–IDS (Full) | **82.4** | **73.1** | **+1.8** |
| w/o CL | 68.7 | 58.4 | –12.6 |
| Naïve Fine-Tuning | 72.1 | 60.3 | –18.9 |

TABLE III
ABLATION STUDY: IMPACT OF REPLAY AND SAMPLE SELECTION

| Method | AVG (%) | FwdTrans (%) | BwdTrans (%) |
|---|---|---|---|
| RAN–IDS (Full) | **82.4** | **73.1** | **+1.8** |
| w/o Sample Selector | 78.9 | 67.8 | +0.4 |
| w/o Replay | 74.2 | 63.5 | –9.7 |

- Forward Transfer reflects zero-day generalization under evolving traffic.

# Evaluation: Task-wise Stability

- Baselines show gradual forgetting
  - RAN–IDS maintains stable performance
  - Effective adaptation across all tasks

- Demonstrates resistance to catastrophic forgetting.

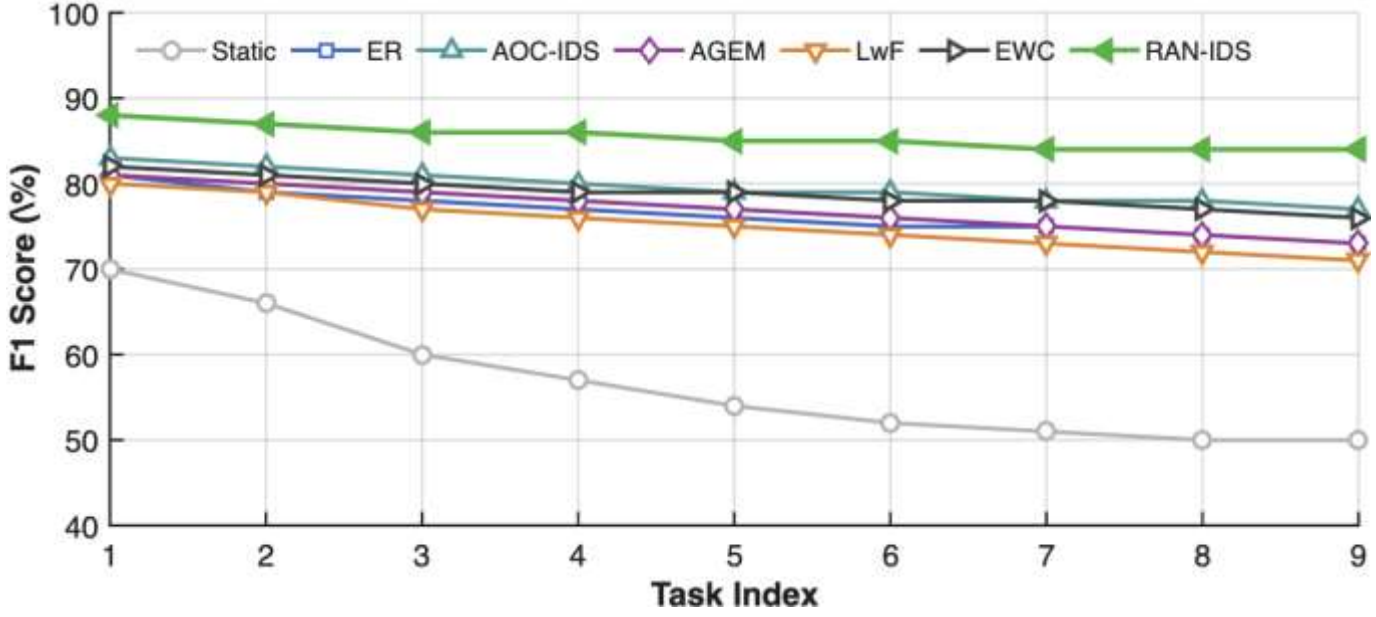

# Conclusion

- First architecture-aware continual IDS for O-RAN
- Mitigates catastrophic forgetting under interface drift
- Achieves superior zero-day detection with federated adaptation

Future Work:

- Integration with RIC closed-loop mitigation

# Thank you!

For questions this please email: nadia.niknami@villanova.edu