Programmable & Trustworthy Data Paths for Sovereign Network Services

Anestis Dalgkitsis, José E. Zerna Torres, Angelos Dimoglis, Marios Avgeris, Chrysa Papagianni, Paola Grosso
University of Amsterdam, LAB42, Science Park 900, 1098 XH Amsterdam, The Netherlands
Email: {a.dalgkitsis, j.e.zernatorres, a.dimoglis, m.avgeris, c.papagianni, p.grosso}@uva.nl

***Abstract*—The integration of digital technologies across governance, commerce, and communication demands a more autonomous and resilient Internet, that protects against manipulation and surveillance. A Responsible Internet architecture must empower end users with fine-grained control over data handling, enabling a trust-aware, adaptive management of network behavior at the data plane. To this end, emerging technologies such as Programmable Data Planes, In-band Network Telemetry, In-Network Reinforcement Learning, and Intent-Based Networking are critical to achieving these goals. In this work, we present a Responsible Internet Proof-of-Concept that: Supports user-intent-driven path control and telemetry-based observability, integrates reinforcement learning agents within programmable forwarding devices for autonomous and secure data path selection, and demonstrates its feasibility and performance through deployment and evaluation on the international FABRIC testbed.**

***Index Terms*—Network Sovereignty, In-Network Reinforcement Learning, In-Band Network Telemetry, Data-plane Programmability, Trusted Communication, Intent-Based Networking.**

I. Overview

The heightened geopolitical tensions have propelled the urgency for governments and major corporations from all around the world to maintain firm control over their data infrastructures. Future network must be able to provide transparency and control, enforcing fine-grained customized user preferences, moving into a new era of explainability and automation, away from the black-box implementation of today.

Technology enablers such as Programmable Data Planes (PDP) and In-Band Network Telemetry (INT) provide data flow monitoring by embedding metadata on top of user traffic to allow for real-time path visibility [1], and enable the enforcement of network configuration changes directly within the data plane [2]. Reinforcement Learning (RL) at the level of the data plane can react instantly, enforcing automation rules set by the user to protect their data flow from external manipulation and surveillance of malicious actors. Intent-Based Networking (IBN) can take user preferences one step further, translating abstract ideas expressed with natural language into actionable steps that reconfigure the network by leveraging Large Language Models (LLM). This paradigm shift not only enhances user experience but also accelerates service deployment and adaptation, aligning perfectly with the goals of a highly automated, resilient network infrastructure.

We demonstrate a Proof-of-Concept (PoC) that embraces PDP, INT, IBN and In-Network Reinforcement Learning (IN-RL) to realize dynamic, trust-aware network control from a simple natural language description. Our demo showcases how real-time monitoring and in-network decision-making at the data plane can enforce user preferences at an instant to increase trustworthiness across the user data flows. Specifically, we show how traffic can be dynamically steered based on user-defined trust policies expressed with natural language.

979-8-3315-9777-1/25/$31.00 © 2025 IEEE

II. Innovation

This demo showcases a future network services PoC that enables user-driven, trust-aware, and adaptive control of the data plane. By combining INT with IN-RL on P4 programmable network elements [3], the system dynamically steers traffic based on real-time trust scores and user-defined policies. Specifically, our demo illustrates how deep programmability and AI-native infrastructure can enforce digital sovereignty and security guarantees directly in the network fabric.

III. PoC Implementation

Fig. 1 provides an overview of the network topology and the loosely coupled set of components that constitute the proposed framework.

We assume a critical client-server application scenario where the client initiates a TCP data flow to the server through the programmable data-plane. As the data packets that contain critical information traverse through the network, the switches are embedding telemetry data into the packet headers, leveraging INT. Once the traffic flow packets reach the last node in the path, the telemetry data are extracted and forwarded to the intelligent telemetry collector, while traffic is forwarded as expected to the end-user.

We use a P4-programmable switch as telemetry collector to allow for ultra-fast identification of changes in the trust level of the data flow path. It retrieves and processes telemetry data for each traffic flow, such as the vendor, location and firmware for each switch in the flow path to calculate the trust level of the path through the path tracing info. The trust properties of the active path are illustrated in real-time via a portal application provided by the local Internet Service Provider as retrieved by the telemetry server. These include a real-time abstract color-coded representation of the nodes, depending on the trust level of the switches in the active path.

IV. Main Demonstration Components

The main pillars of our demo will be summarized in the following paragraphs below.

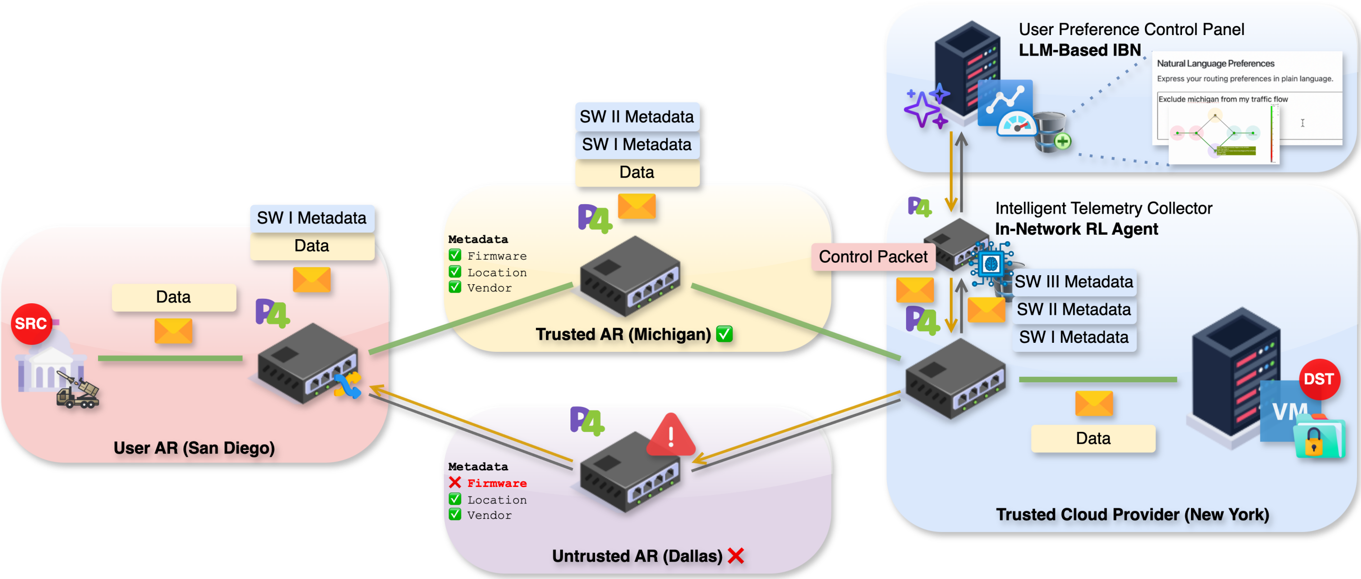

Fig. 1. Experimental demonstration topology overview. The Autonomous Regions (ARs) in parentheses reflect the actual location of the nodes and switches.

A. Pervasive Telemetry & Real-time Monitoring

Network elements report metrics at a per-packet granularity using an INT-based approach implemented on a fully P4-programmable data plane. Metadata are appended to the traffic, conveying critical service metrics to a collection point [2], located one hop before the network egress. The telemetry data are then forwarded from the collector switch to a telemetry server using Layer 2 embedded packets. Then the data are extracted and stored in a database for further data analysis and processing. The metadata include the switch ID, vendor, physical location, and active firmware version.

B. Node Trust Definition

The trustworthiness of a path is calculated as the minimum trust of the nodes participating in the path. Each node is assigned a trust level based on several device-specific metrics, such as geo-location, active firmware, vendor, operator, interactions with other nodes, and similar criteria [4], [5]. For demonstration purposes, we sort the regions and vendors arbitrarily and define the trust of a node via a weighted sum as follows:

$$trust = w_v \cdot vendor + w_l \cdot location + w_f \cdot firmware, \quad (1)$$

where w_v, w_l and w_f are the vendor, location and firmware weights specified by an overseeing authority, such as a company or the government.

C. In-Network Reinforcement Learning Agent

The metadata are extracted one hop before exiting the network and forwarded to the Intelligent Telemetry Collector, a P4-programmable switch that performs IN-RL at wire speed to guarantee that the user preferences are enforced. The Intelligent Telemetry Collector is responsible for the collection, aggregation, filtering, and correlation of the data and the IN-RL actuation for path selection.

The IN-RL agent runs in the collector switch as a Stochastic Learning Automaton (SLA) [6] to perform in-network path selection. Upon each packet arrival, the trust levels of all the nodes along the traversed path are aggregated using the minimum function to represent the weakest link. The reward is calculated by applying a square function to the path trust metric, accentuating the influence of higher trust levels in the SLA update process. This reward then drives the update of path selection probabilities, reinforcing the likelihood of choosing more trusted paths over time.

D. Web Portal

The telemetry information is displayed via a graphical user interface that provides:

- The **Real-time Topology map**, shown in Fig. 2, which is an abstract representation of the topology graph that shows in real-time the active path of the user data flows, and all elements in the network via aggregated information.
- The **Preferences** section, that allows the end-user to provide their preferences and control the path of their data. This form allows users to exclude specific regions or vendors of networking equipment from their data flow for business reasons or adjust the tolerance in trust degradation.
- The **AI-Agent** section allows data flow customization similar to the preferences section, but in contrast it uses receives high level abstract instructions in natural language. Inspired by modern LLM-based AI chat agents [7], it allows users without the required technical background to express the desired state and preferences

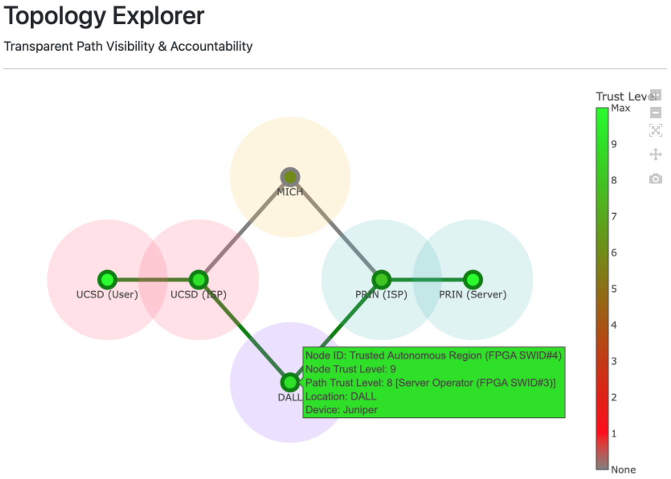

Topology Explorer

Transparent Path Visibility & Accountability

Node ID: Trusted Autonomous Region (FPGA SWID#4)
Node Trust Level: 9
Path Trust Level: 8 [Server Operator (FPGA SWID#3)]
Location: DALL
Device: Juniper

Fig. 2. Real-time topology explorer. The nodes are color-coded trust level according to the color bar on the right, whereas the background color differentiates the ARs. The green path indicates the currently active path of the user's data flow.

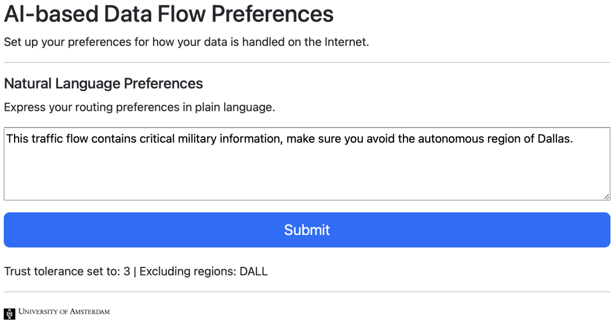

AI-based Data Flow Preferences

Set up your preferences for how your data is handled on the Internet.

Natural Language Preferences

Express your routing preferences in plain language.

This traffic flow contains critical military information, make sure you avoid the autonomous region of Dallas.

Trust tolerance set to: 3 | Excluding regions: DALL

Fig. 3. LLM-based IBN data flow preferences input from. The users can use natural language to express the preferences, without any domain-specific knowledge.

for their data flow using natural language, which the system can then interpret and apply.

E. LLM-Based IBN User Preferences Input

The Web Portal includes a tab where the users can express their preferences via natural languages in an easy and streamlined way. That allows users to set data-flow preferences without the need to learn any technical skills, domain-specific knowledge or perform actions that may break the production system. In this demo, we leverage LLM to build an IBN experiment that interacts with the network and invite viewers to participate. Local LLMs with Ollama [7] were used to develop a prototype that relies on the few-shot prompt technique to translate high level abstract user intents into network configurations. For the live demonstration we intent to use the OpenAI GPT-4o model.

F. Packet Header Structure & Collected Metadata

The structure of the packet Telemetry Headers is depicted in Fig. 4. Telemetry data such as the Path Tracing and Trust related information such as the brand, location and current firmware, is embedded to the packet as TCP Options

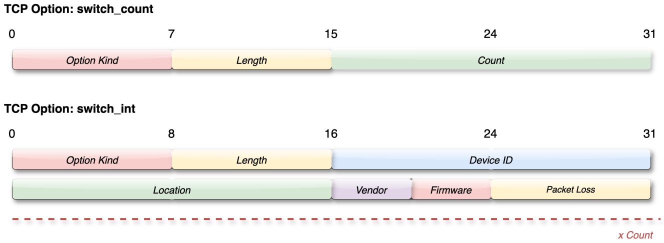

TCP Option: switch_count

0 7 15 24 31

Option Kind | Length | Count

TCP Option: switch_int

0 8 16 24 31

Option Kind | Length | Device ID

Location | Vendor | Firmware | Packet Loss

x Count

Fig. 4. Telemetry packet header structure overview.

The header structure of the INT-triggered control message, consists of the target Switch ID, the Port ID, along with flow identification metadata, for redirecting the flow to the desired path. We are collecting different types of performance, security and routing metrics to showcase our demo in action. The path is traced using unique identifiers for every switch along the path. This approach ensures that the user is aware of the path that is being used at any given time, as explained in [1].

G. The FABRIC Testbed & Dedicated Optical Links

The FABRIC testbed is an international research infrastructure designed for large-scale scientific experimentation [8]. It provides high-speed equipment to support bandwidth-intensive applications and real-time data transfers. The FABRIC TeraCore ring serves as the backbone, utilizing spectrum from the ESnet6 fiber footprint, an advanced high-speed network that connects US Department of Energy laboratories, third-party facilities, and international research organizations.

In this demo, we utilize a virtualized FABRIC network slice with shared connections, as depicted in Fig. 5. The slice is built using BMV2 switches and programmed with the P4-16 language [3]. It will be deployed and demonstrated live, allowing users to interact with it in real time. The TeraCore ring supporting the slice leverages optical transport equipment from Ciena.

V. Live Demo Scenario

The live demonstration consists of several steps, each showcasing a different component of the framework:

1) Traffic Flow Initiation: The demo begins by establishing a real-time TCP data flow between the user and the server via a remote terminal connection.

2) Manual Preferences Input: First, acting as the user, we navigate to the Preferences tab. Through this interface, the user can prompt the system to reconfigure the network elements based on their personal trust requirements. We define three levels of trust: low, medium, and high, considering security features, regulatory compliance, and digital sovereignty [4], [5]. Upon switching to the Live Topology tab Fig. 2, we observe the RL-agent dynamically identifying and selecting the most trusted available path on behalf of the user. In this case, the selected route passes through the Dallas, reflecting the user's trust preference.

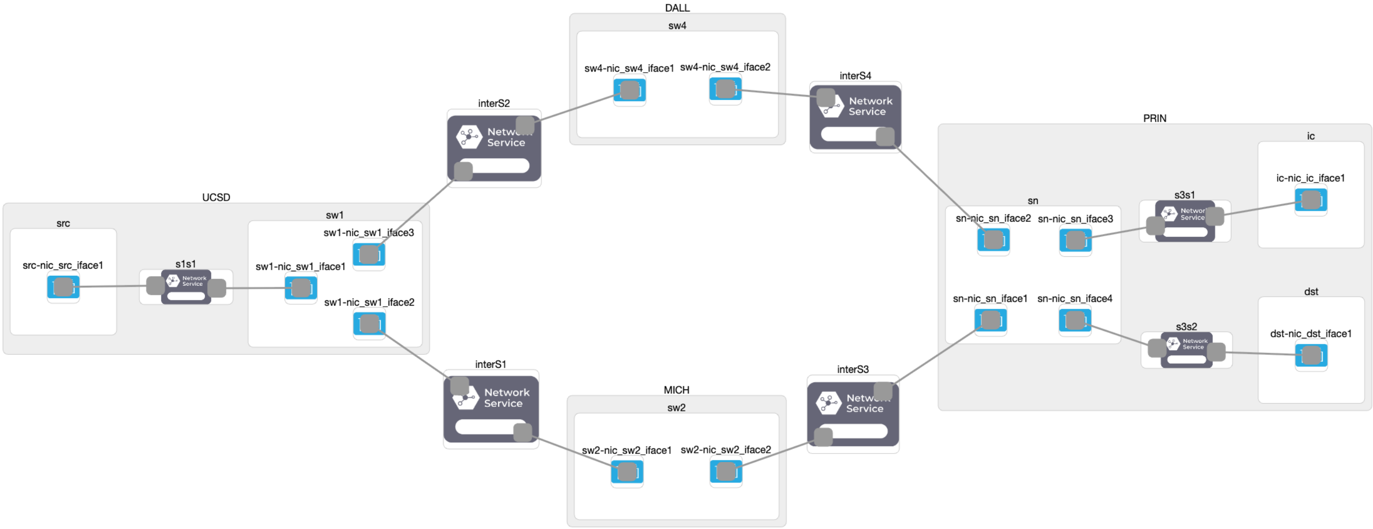

Fig. 5. Experimental demonstration FABRIC slice topology.

3) Custom Policy Enforcement with IN-RL: Next, we simulate a trust degradation event at the Dallas node by remotely pushing a firmware update that introduces a known security vulnerability. As a result, the calculated trust level of the node is reduced to low, making it incompatible with the specified trust preferences. The RL agent immediately detects the change through INT and shifts the traffic through the next most trusted AR Michigan. This transition is visualized in real time via the web app and the real-time network topology. We then simulate a new firmware patch on the Dallas node, which restores its original trust level. However, the RL agent maintains the current path to prevent any unnecessary fluctuations, demonstrating its stability-aware decision logic.

4) AI-enabled IBN Agent: In the next step, the demo attendees are invited to interact with the AI-Agent via the IBN portal, as shown in Fig. 3. They can describe their data flow preferences using natural language, for example requesting to exclude the Michigan AR for personal or business reasons. The LLM-based AI-Agent interprets this input and translates it into actionable network configurations which are applied immediately, resulting in a new path selection that aligns with the user's updated preferences.

5) Data & Strategy Explainability: The demonstration concludes by exploring a report generated by the RL-Agent. This report outlines the agent's training process and explains the path selection strategy in relation to the fluctuations in node trust and user-defined preferences. It provides transparency into the agent's decisions, showcasing how the framework adapts to real-time conditions.

VI. Conclusions

This PoC demonstrates a new direction for network services that emphasizes user-driven IBN control, trust awareness, and adaptive network behavior. With the combination of several programmable networking technologies, RL and LLM agents within the network infrastructure, we are demonstrating real-time monitoring and enforcement of policies on behalf of the user, extracted from an abstract natural language description. Our demo highlights the potential for future networks to become more transparent, resilient, and aligned with user-defined objectives, advancing the broader vision of a secure and autonomous Responsible Internet.

Acknowledgment

This work was supported by CIENA, the European research project DESIRE6G (101096466), the Dutch National Growth Fund project FNS and the Research Council NWO project CATRIN (NWA.1215.18.003). It was made possible by FABRIC and specifically FABRIC Across Borders.

References

[1] K. Papadopoulos, P. Papadimitriou, and C. Papagianni, "Deterministic and probabilistic p4-enabled lightweight in-band network telemetry," *IEEE Transactions on Network and Service Management*, 2023.

[2] A. Sgambelluri, A. Pacini, L. Valcarenghi, A. Dimoglis, C. Papagianni, F. Cugini, and F. Paolucci, "Pervasive Telemetry Solutions for 6G Systems and end-to-end Accelerated Services Monitoring," 2023.

[3] T. P. A. W. Group. In-band network telemetry dataplane specification v2.1. [Online]. Available: https://p4.org/p4-spec/docs/INT_v2_1.pdf

[4] N. Torkzaban, C. Papagianni, and J. S. Baras, "Trust-aware service chain embedding," in *2019 Sixth International Conference on Software Defined Systems (SDS)*. IEEE, 2019, pp. 242–247.

[5] S. K. Dhurandher and V. Mehra, "Multi-path and message trust-based secure routing in ad hoc networks," in *2009 International Conference on Advances in Computing, Control, and Telecommunication Technologies*. IEEE, 2009, pp. 189–194.

[6] C. Ünsal, J. S. Bay, J. A. Ball, W. T. Baumann, P. Kachroo, and H. F. Vanlandingham, "Intelligent navigation of autonomous vehicles in an automated highway system: Learning methods and interacting vehicles approach," 7 1998. [Online]. Available: http://hdl.handle.net/10919/30595

[7] F. Marcondes, A. Gala, and R. Magalhães, "Using ollama," in *Approach with Ollama*. Springer, 2025. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-031-76631-2_3

[8] I. Baldin, A. Nikolich, J. Griffioen, I. Monga, P. Ruth, X. Yufeng, J. Chase, and R. Ricci, "Fabric: A national-scale programmable experimental network infrastructure," *IEEE Internet Computing*, vol. 24, no. 4, pp. 38–47, 2020. [Online]. Available: https://ieeexplore.ieee.org/document/8972790