Digital Sovereignty in Practice: Leveraging In-Band Telemetry and ML for a Responsible Internet

Anestis Dalgkitsis, Jose Zerna Torres, Angelos Dimoglis, Luca Cetino, Marios Avgeris, Chrysa Papagianni, Paola Grosso

Informatics Institute, University of Amsterdam, Amsterdam, The Netherlands
{a.dalgkitsis, j.e.zernatorres, a.dimoglis, l.cetino, m.avgeris, c.papagianni, p.grosso} @uva.nl

1. Innovation, Contributions & High-Level Goals

As digital technologies continue to integrate into governance, commerce, and communication, it is essential for societies to maintain their autonomy and ensure that critical systems remain resilient against external manipulation or surveillance. A responsible Internet should empower not only providers of critical services but also individuals to access and choose the equipment handling their data. Users should also have enough control over the destiny of their data by specifying several requirements, such as trusted networking equipment and preferred geographical location. Moreover, it must enable users to verify whether operators act in good faith and trace incidents or attacks back to their root causes.

Emerging technologies such as programmable networks, In-band Network Telemetry (INT), Machine Learning (ML), and Intent Based Networking (IBN) play a pivotal role in realizing these goals. The advent of Programmable Data Planes (PDPs) has enabled INT, which offers significant advantages, including flexible programmability and real-time, detailed network visibility. This is achieved by embedding network state information directly into data packet headers, allowing the data plane to independently manage network measurements without intervention from the control plane.

Driven by the socio-economic outlook for the future of the Internet, we have designed, developed, and demonstrated a Responsible Internet proof-of-concept (PoC). Our work contributes to the field in three key ways, by:

- Enabling users to intuitively specify and monitor the path of their data on the Internet through an operator platform, and verify the trustworthiness of the path by using INT.
- Implementing a Reinforcement Learning (RL) approach in programmable devices that autonomously learns and selects optimal, secure routes based on INT-collected metrics and user preferences, facilitating dynamic path optimization directly in the data plane without control plane intervention.
- Conducting realistic experiments on the FABRIC network infrastructure, in order to validate the feasibility and practicality of this PoC.

2. Demonstrator Description & Operation

Fig. 1 provides an overview of the network topology and the components that constitute the demonstrator.

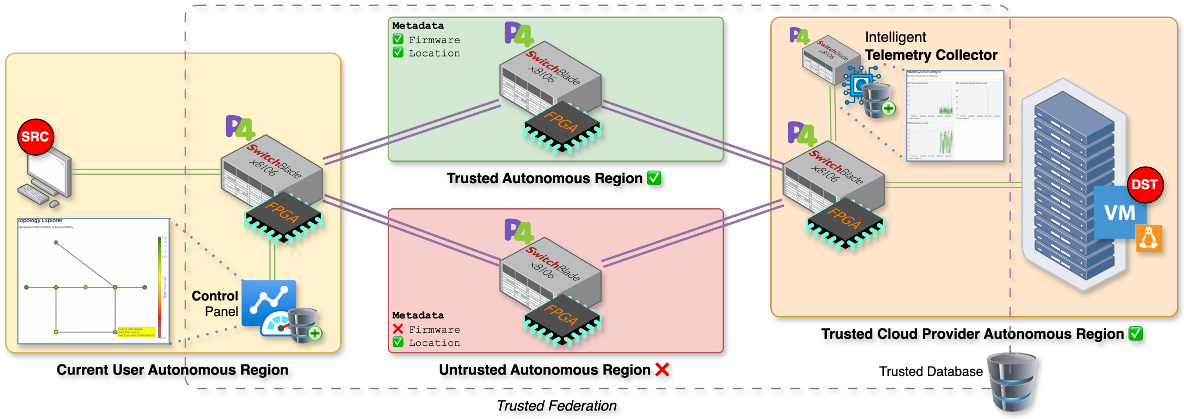
Metadata
✅ Firmware
✅ Location
P4
SwitchBlade x8106
FPGA
Trusted Autonomous Region ✅
Intelligent Telemetry Collector
SRC
Topology Explorer
Control Panel
Current User Autonomous Region
Metadata
❌ Firmware
✅ Location
Untrusted Autonomous Region ❌
DST
VM
Trusted Cloud Provider Autonomous Region ✅
Trusted Database
Trusted Federation

Figure 1: Experimental Demonstrator Topology.

2.1 Pervasive Telemetry & Trusted Flow Monitoring

Pervasive telemetry allows network elements to report metrics with a per-packet granularity. We utilize the INT approach on a fully P4-programmable data-plane FPGA hardware switches. The metadata of each networking device are appended to the traffic, conveying crucial service metrics to the Collection Point. The collected telemetry information is displayed through a portal provided by the service provider, referred to as the *Control Panel*. It allows the network to be transparent to the users by providing critical information about the traversed infrastructure, such as device firmware version and geo-location, used protocols. This information helps to evaluate the selected data paths.

2.2 User Data Flow Control

This action takes place through the Control Planes of the trusted operators via a platform called *Trusted Federation*. Aligned with the principles of IBN, the *Control Panel* allows the end user to interact safely with the PDP infrastructure and instruct in a high level the P4-programmable switches to modify the direction of their specific data flow to fit their preferences.

2.3 Intelligent Telemetry Collector

The system implements an RL framework, deployed directly in programmable network devices to realize intelligent path selection (In-Network RL). Operating entirely in the data plane, the RL agent processes telemetry data collected through INT as input states and incorporates the user-defined security intents into its decision-making process. A reward function evaluates the security compliance of potential paths based on these preferences, allowing the system to continuously update path selection actions through the agent's learning mechanism (i.e., intent activation and assurance).

3. PoC Demonstration Scenarios

In our demonstrations, we assume a critical client-server application scenario, where TCP traffic is sent from one end to the other. As the user data packets traverse through the network, the P4 programmable switches are appending telemetry information into the packets with metadata, leveraging traditional INT for collecting path tracing and trust information. At the last hop, the telemetry data are extracted and forwarded to the critical-application server. These data are processed on the server and then shared among the trusted Federation service providers.

Scenario I: User Intent-driven Traffic Control. The end-users connect to the *Control Panel* of the local service operator using their personal profiles to customize preferences, including trust settings for specific devices and geographical areas. The live topology map, generated using INT data collected by the PDP, clearly indicates whether the current path meets the end-user trust requirements for transmitting critical information to the application server. In the Control Panel, the users can specify trust criteria, prompting the network to reconfigure all P4-programmable switches to meet these requirements for the user data flows. Once confirmed, the end-users initiate application traffic and can monitor the path integrity in real-time through the *Control Panel*.

Scenario II: Autonomous Learning-driven In-network Control. The end-users initiate connections, the system collects INT metrics from the P4 programmable-switches along the data paths. A learning agent on the *Intelligent Telemetry Collector* processes the telemetry in real-time to compute optimal routes that meet the user trust preferences. As network conditions and security parameters change, an RL-based algorithm dynamically adjusts routing decisions, continuously improving path selection. Users can monitor these AI-driven decisions via the *Control Panel*.

4. Evaluation & Prototype PoC

We evaluate our demo by collecting different types of metadata and performing real-time path-control actions to enable Internet Transparency and Controllability. Specifically, we:

- embed the P4-programmable switch IDs and metadata between the ARs into the packet headers using INT (embedding metadata on each hop) based on the type of metric. The path is traced using unique identifiers for every switch along the path at low level, without any probes or trace-routing software. This approach ensures that the user is aware of the path trust that is being used at any given time.
- let the user set trust preferences and control their data path in the network in real-time, via securely interacting with the P4 programmable network switches.
- Validate the effectiveness of our SLA-based routing approach by measuring the convergence time to optimal paths and analyzing how quickly the system adapts to changes in security requirements or network conditions.