

# rov\_enforcement\_dp.py — Data-Plane ROV Validation Module

*(Auxiliary validation of control-plane ROV enforcement inference)*

---

## 1. Purpose & Role in the Platform (Explicit Positioning)

rov\_enforcement\_dp.py is an auxiliary, non-core validation script.

It is **not** part of the continuous production pipeline and **does not contribute directly to the primary ROV enforcement score** used by the platform.

Instead, this module exists for **occasional, targeted validation** of the control-plane (CP) ROV enforcement results, in order to:

- empirically verify whether control-plane inferences reflect real packet-level routing behavior,
- detect possible control-plane overestimation,
- align the platform's methodology with peer-reviewed scientific research.

This design choice follows the methodology described here:

[https://www.researchgate.net/publication/324162583\\_Practical\\_Experience\\_Methodologies\\_for\\_Measuring\\_Route-Origin\\_Validation](https://www.researchgate.net/publication/324162583_Practical_Experience_Methodologies_for_Measuring_Route-Origin_Validation)

The authors explicitly conclude:

**“Control-plane experiments do not provide accurate information about ROV-enforcing ASes.”**

and

**“Although control-plane experiments are easier to launch, they are less accurate than our two new data-plane approaches.”**

For this reason, the platform intentionally implements **CP as the primary operational metric** and **DP strictly as a validation and audit mechanism**.

---

## 2. How Data-Plane Complements Control-Plane

### Control-Plane (What CP Measures)

CP observes **routing intent and propagation at the BGP control-plane level**:

- which ASNs propagate RPKI-Invalid announcements,
- which ASNs appear only on Valid paths,
- how often this happens across vantage points.

This is scalable and continuous, but **indirect**.

### Data-Plane (What DP Adds)

DP observes **actual packet delivery outcomes**, answering:

“Do packets routed via the ROA-valid path reach the destination while packets routed via the ROA-invalid path fail?”

This directly addresses the paper’s core criticism:

**“With uncontrolled control-plane experiments one cannot differentiate between ROV filtering and other causes such as traffic engineering.”**

DP therefore **grounds CP inference in real forwarding behavior**.

---

## 3. High-Level Concept (Validation-Oriented View)

The data-plane module compares **paired destinations**:

- one ROA-valid,
- one ROA-invalid,

and observes whether **real traffic** (ICMP / TCP) behaves differently.

If:

- valid traffic succeeds,
- invalid traffic fails,

then **real ROV filtering is likely occurring somewhere along the path**.

This does **not replace CP**, but **confirms or contradicts CP conclusions**.

---

## 4. Operational Flow (End-to-End)

1. Load explicitly defined DP targets (valid / invalid destination pairs).
  2. For each pair:
    - run traceroute to both destinations,
    - reconstruct AS-paths,
    - perform TCP connection attempts.
  3. Infer likely ROV behavior based on **differential reachability**.
  4. Aggregate raw evidence counters per ASN.
  5. Persist results in a dedicated DP table (`asn_rov_dp`).
  6. Use DP output **only for validation, not scoring**.
- 

## 5. Target Definition (Controlled Experiment Design)

The script requires a **JSON targets file** containing explicit probe pairs:

```
{  
  "valid_ip": "X.X.X.X",  
  "invalid_ip": "Y.Y.Y.Y",  
  "port": 80  
}
```

This mirrors the paper’s experimental setup using **paired prefixes with alternating ROAs**, ensuring:

- controlled conditions,
- reproducibility,
- isolation from live CP traffic.

As stated in the paper:

**“Two consecutive probes are executed: one compliant with ROA and one contradicting ROA.”**

---

## 6. Traceroute-Based Data-Plane Inference

### 6.1 Traceroute Execution

For each target pair, the script runs:

```
traceroute -n -w <timeout> -q 1 -m <max_ttl> <destination>
```

Results:

- hop IPs are collected,
- IPs are resolved to ASNs using Team Cymru WHOIS,
- an **observed AS-path** is reconstructed.

This implements the same “virtual AS-path” concept described in Section IV of the paper.

---

### 6.2 Traceroute Inference Logic (Exact Semantics)

The script compares:

- ASNs seen on the valid path,
- ASNs seen on the invalid path.

#### Inference rules:

- ASN appears **only on the valid path**
  - `tr_likely_rov += 1`
  - Positive evidence of ROV enforcement.
- ASN appears on invalid path (or both)
  - `tr_no_rov += 1`
  - No evidence of ROV filtering.

This exactly matches the paper's logic:

**“ASes that occur only on paths compliant to ROA are likely filtering invalid routes.”**

---

## 7. TCP Connectivity-Based Data-Plane Inference

### 7.1 TCP Probing

For each target pair, the script attempts:

```
socket.connect((ip, port))
```

to both:

- ROA-valid destination,
- ROA-invalid destination.

Multiple attempts can be configured to reduce noise.

---

### 7.2 TCP Inference Logic

- **Valid succeeds + Invalid fails**
  - `tcp_likely_rov += 1`
  - Strong packet-level evidence of filtering.
- **Both succeed or both fail**
  - `tcp_no_rov += 1`
  - Inconclusive, no inference.

This mirrors Section V of the paper:

**“The method is based on sending TCP connection initiation segments and observing whether replies are delivered.”**

Noise handling is explicit, as recommended:

**“Data-plane measurements result in significant noise and require filtering out randomness.”**

---

## 8. Evidence Aggregation Model (Why No Single Score)

Unlike CP, the DP module **intentionally does not compute a final enforcement score**.

Instead, it stores **raw evidence counters**, because:

- DP data is sparse,
- DP data is noisy,
- DP coverage is incomplete.

This directly follows the paper’s warning:

**“Complete reliance on data-plane measurements is not feasible due to noise and operational constraints.”**

---

## 9. Database Schema & Semantics

Table: `asn_rov_dp`

```
CREATE TABLE IF NOT EXISTS asn_rov_dp (  
    asn                INTEGER PRIMARY KEY,  
    tr_probed          INTEGER,  
    tr_likely_rov      INTEGER,  
    tr_no_rov          INTEGER,  
    tcp_probed         INTEGER,  
    tcp_likely_rov     INTEGER,  
    tcp_no_rov         INTEGER,  
    dp_last_updated    INTEGER  
);
```

## Field Semantics

- **asn**  
Autonomous System under evaluation.
- **tr\_probed**  
Number of traceroute-based opportunities involving this ASN.
- **tr\_likely\_rov**  
Traceroute events suggesting ROV enforcement.
- **tr\_no\_rov**  
Traceroute events showing no filtering evidence.
- **tcp\_probed**  
Number of TCP probe opportunities.
- **tcp\_likely\_rov**  
TCP events where valid succeeded and invalid failed.
- **tcp\_no\_rov**  
TCP events inconclusive or showing no filtering.
- **dp\_last\_updated**  
UNIX timestamp of last DP observation.

The table is **append-only in effect**, accumulating evidence across runs.

---

## 10. How DP Results Are Used (and How They Are Not)

### DP IS USED FOR:

- validating CP classifications,
- identifying CP false positives,
- supporting research-grade methodology claims,
- audit and due-diligence.

### DP IS NOT USED FOR:

- computing the main ROV enforcement score,
- feeding ML models directly,
- daily production scoring.

This separation is intentional and methodologically correct.

---

## 11. Why CP + DP Fully Aligns with the Scientific Study

The paper's central methodological conclusion is:

**“Complete reliance on control-plane measurements can be misleading.”**

But it also shows that:

- CP is scalable and operational,
- DP is realistic but noisy.

By implementing:

- **CP as the primary metric**, and



- DP as an auxiliary validation layer,

the platform **exactly follows the combined methodology advocated by the authors.**

This is not partial alignment — it is **full methodological alignment with a pragmatic, production-aware interpretation.**

---

## 12. Final Interpretation

Control-plane answers: “What does the Internet *say* it does?”

Data-plane answers: “What does the Internet *actually do*?”

Using both — with correct role separation — is the **only scientifically defensible approach.**