



Path Tracing (PT)


Clarence Filsfils

Ahmed Abdelsalam

Pablo Camarillo

The exact path from A to M is not known

- ECMP is a key in IP networks.
- 7 possible “valid” ECMP path
 - ABFM, ABGM, ACFM, ACGM, ACHM, ADGM, ADHM ...
- The path may be invalid
 - Routing or FIB corruption @ B



Path Tracing

- Deterministic Per-Packet Tracing
- Implemented at line rate in the base HW pipeline
 - No punting to CPU, no offload to co-processors
- Ultra-MTU-efficient: only 3 byte per hop!
 - 12-bit Interface, 8-bit Timestamp, 4-bit Load
- IPv6 with native SRv6 support
 - MPLS design also available
- Seamless Deployment
 - Interwork with legacy node

Design intuition

Dataplane Encapsulation

- Minimize NPU parsing
- Minimize # of Read/Write
- Minimize depth of Read/Write
- Maximize Read/Write at fixed positions
- Avoid Header Insert/Resize
- Minimize MTU

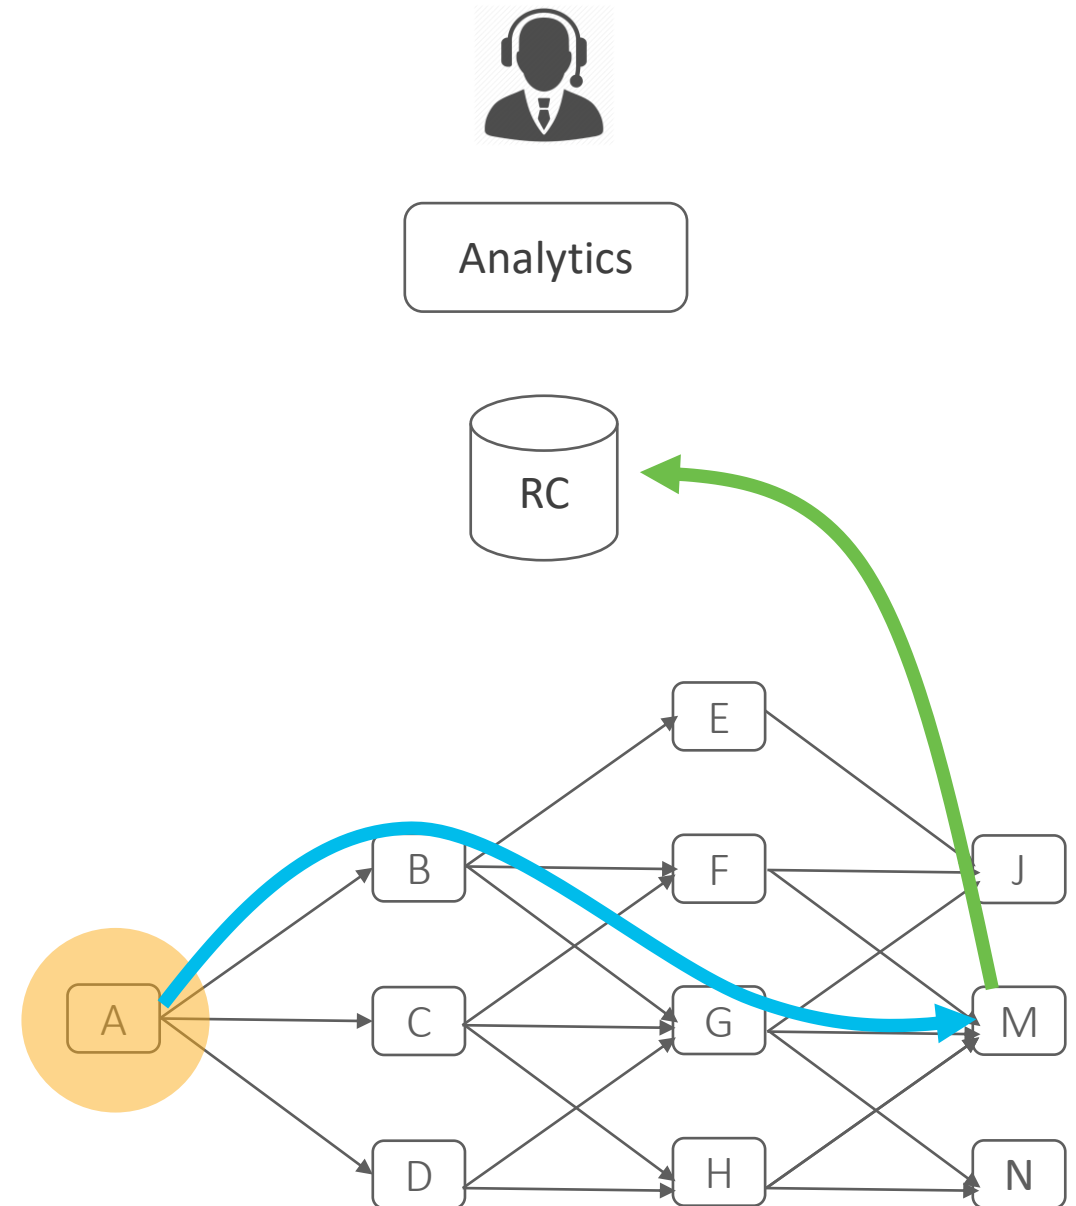
Per-Hop Collected Data

- Highly compressed
- Midpoint Compressed Data (MCD)
 - Only 3 Bytes per Hop!
 - 12-bit Interface ID, 4-bit Interface Load, 8-bit Timestamp

Roles and Data Model

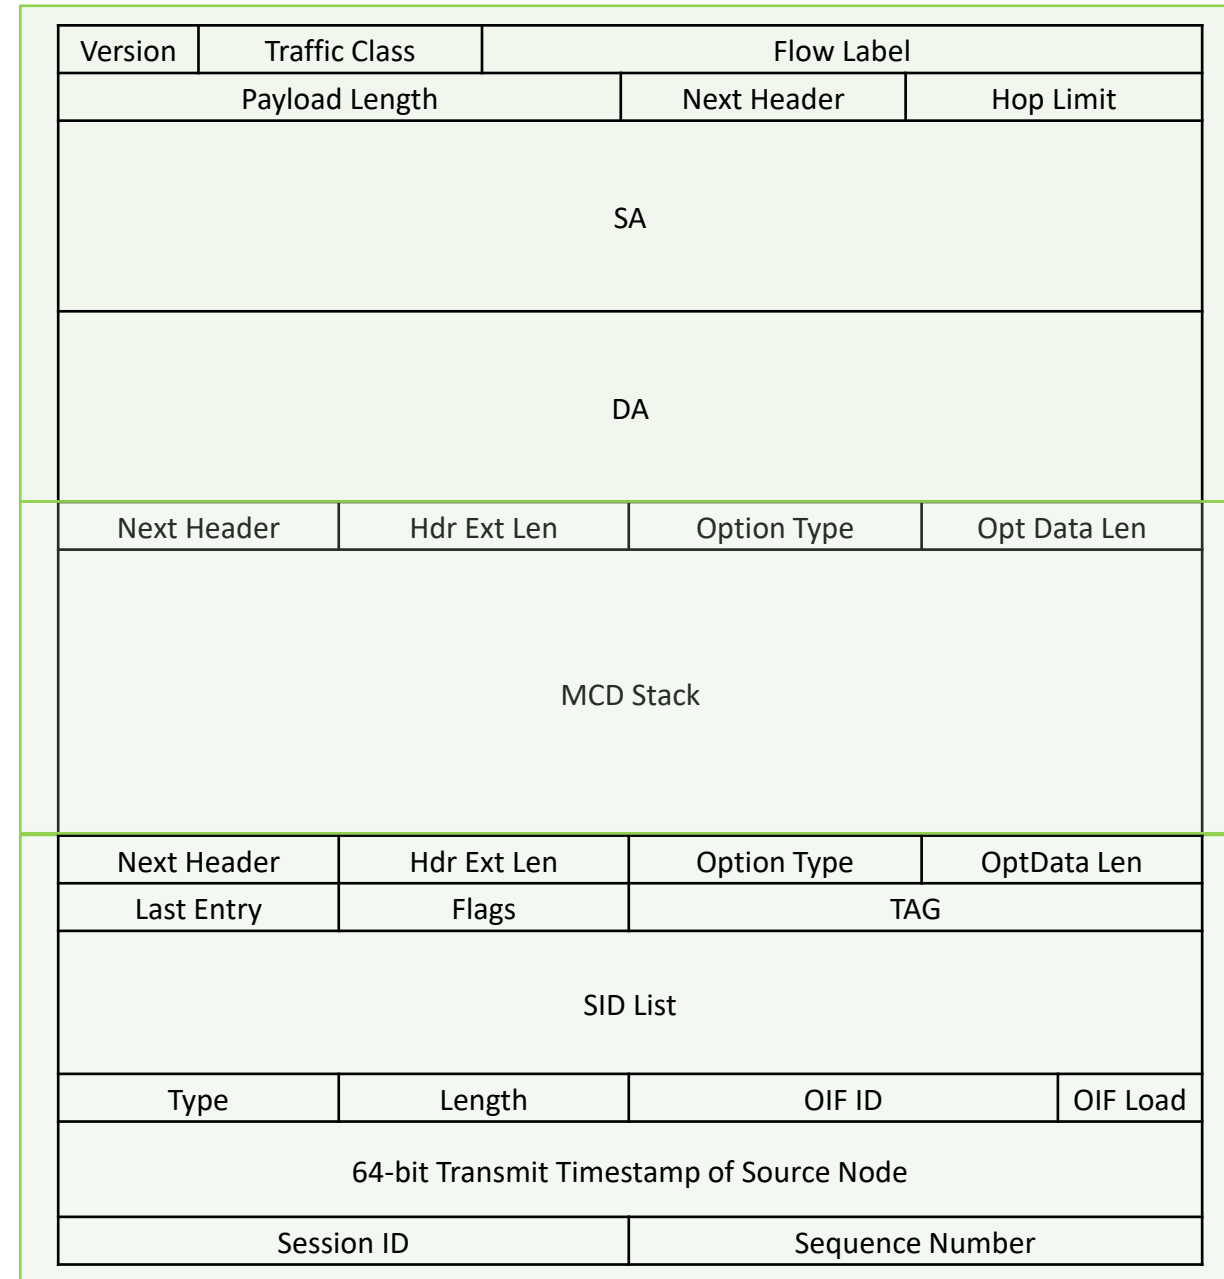
Source

- Originates the probe
- Collects transmission data



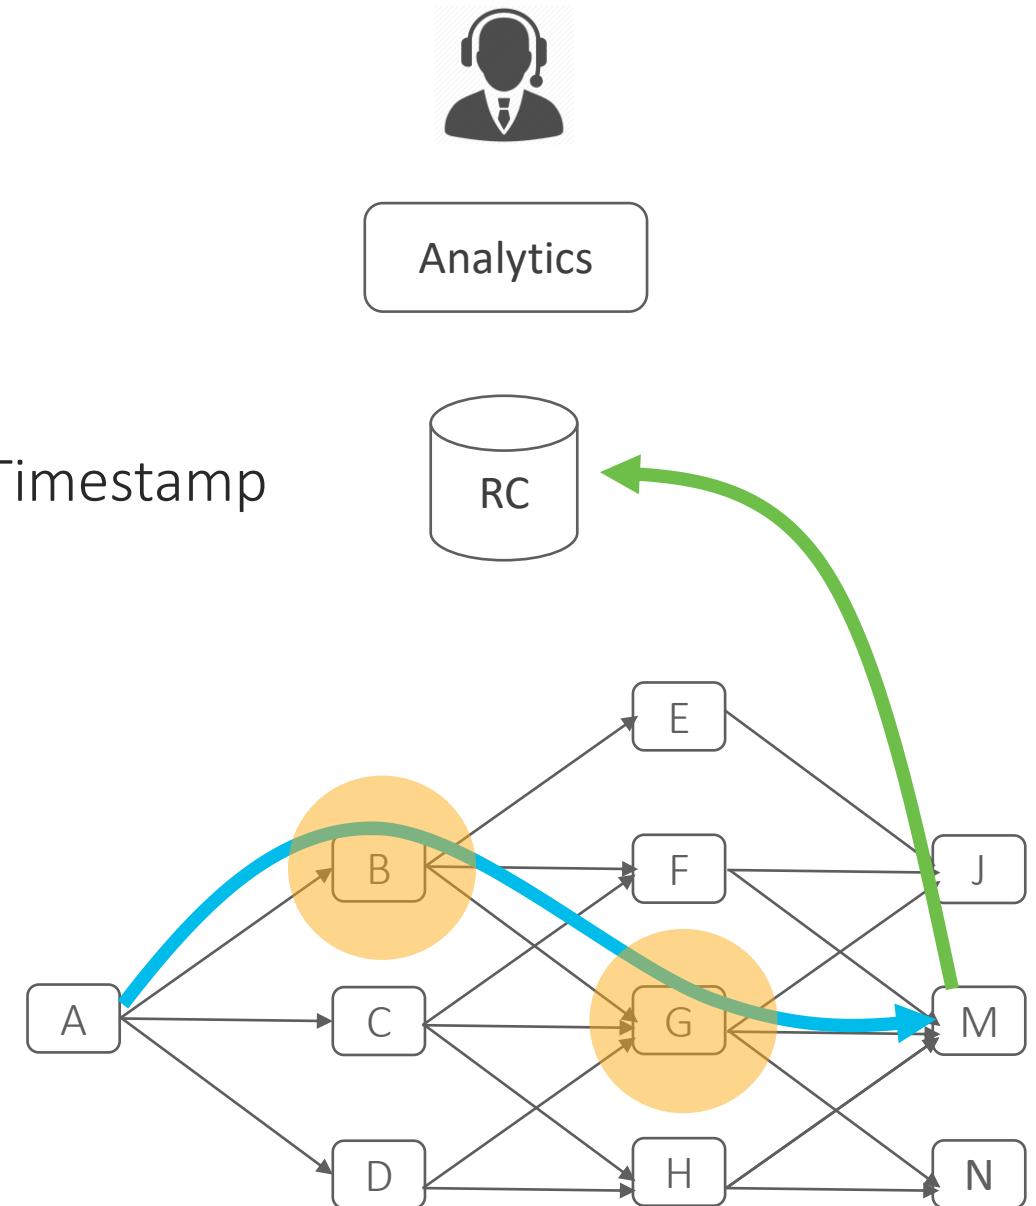
Packet Format

- IPv6 Header
 - SA, DA, DSCP, FL, ...
- HbH header
 - PT Option Type
 - MCD Stack
- SRH
 - SID List
 - SRH PT-TLV (TS, OIF ID, OIF Load)



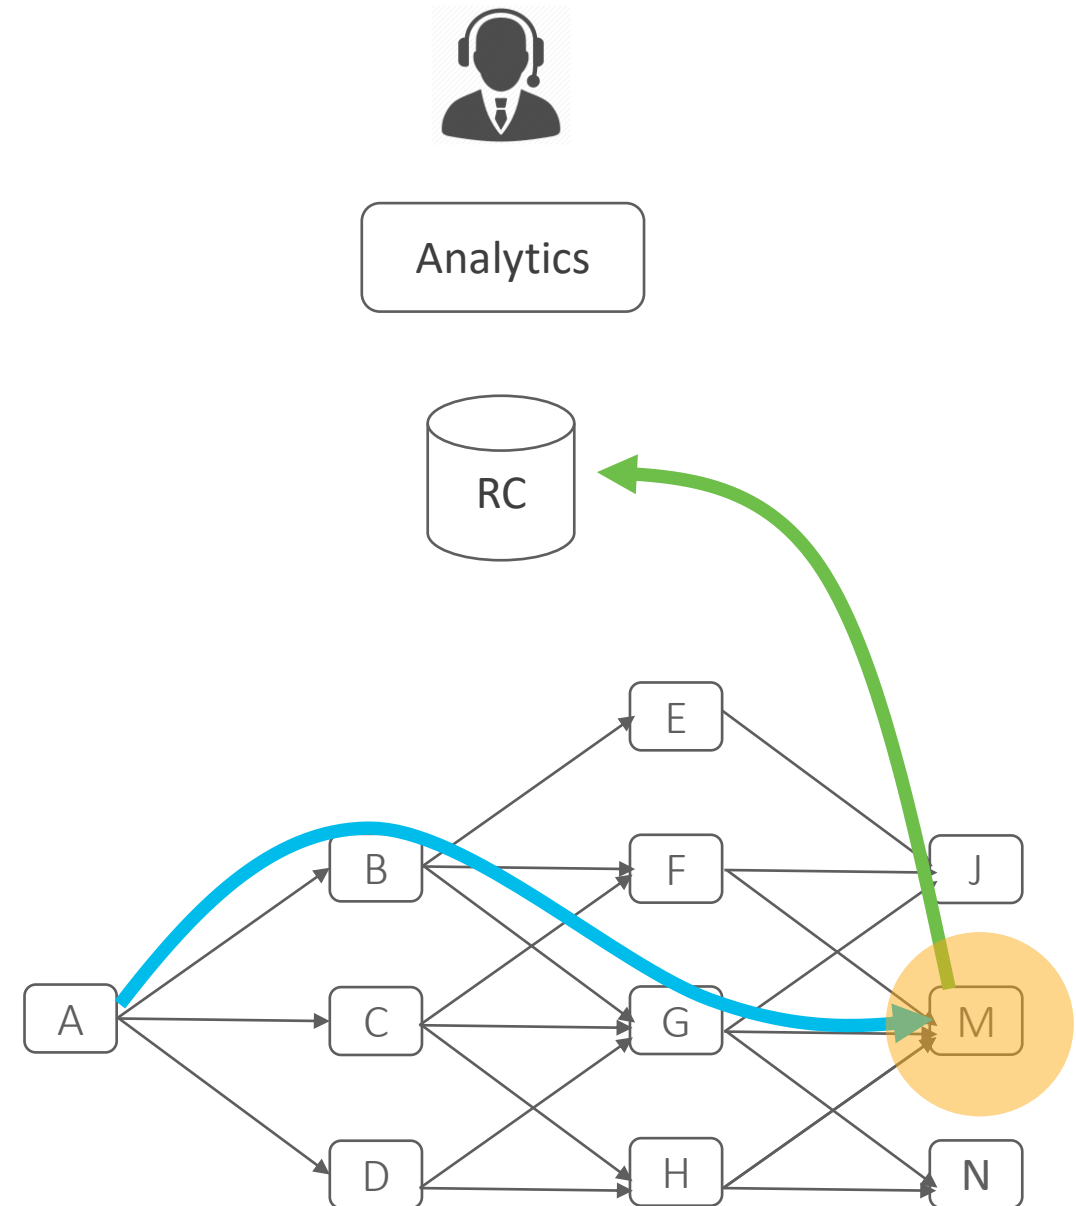
Midpoint

- Collects Midpoint Compressed Data (MCD)
- MCD
 - Only 3 byte per hop!
 - 12-bit Interface ID, 4-bit Interface Load, 8-bit Timestamp
- Shift & Stamp behavior
 - linerate



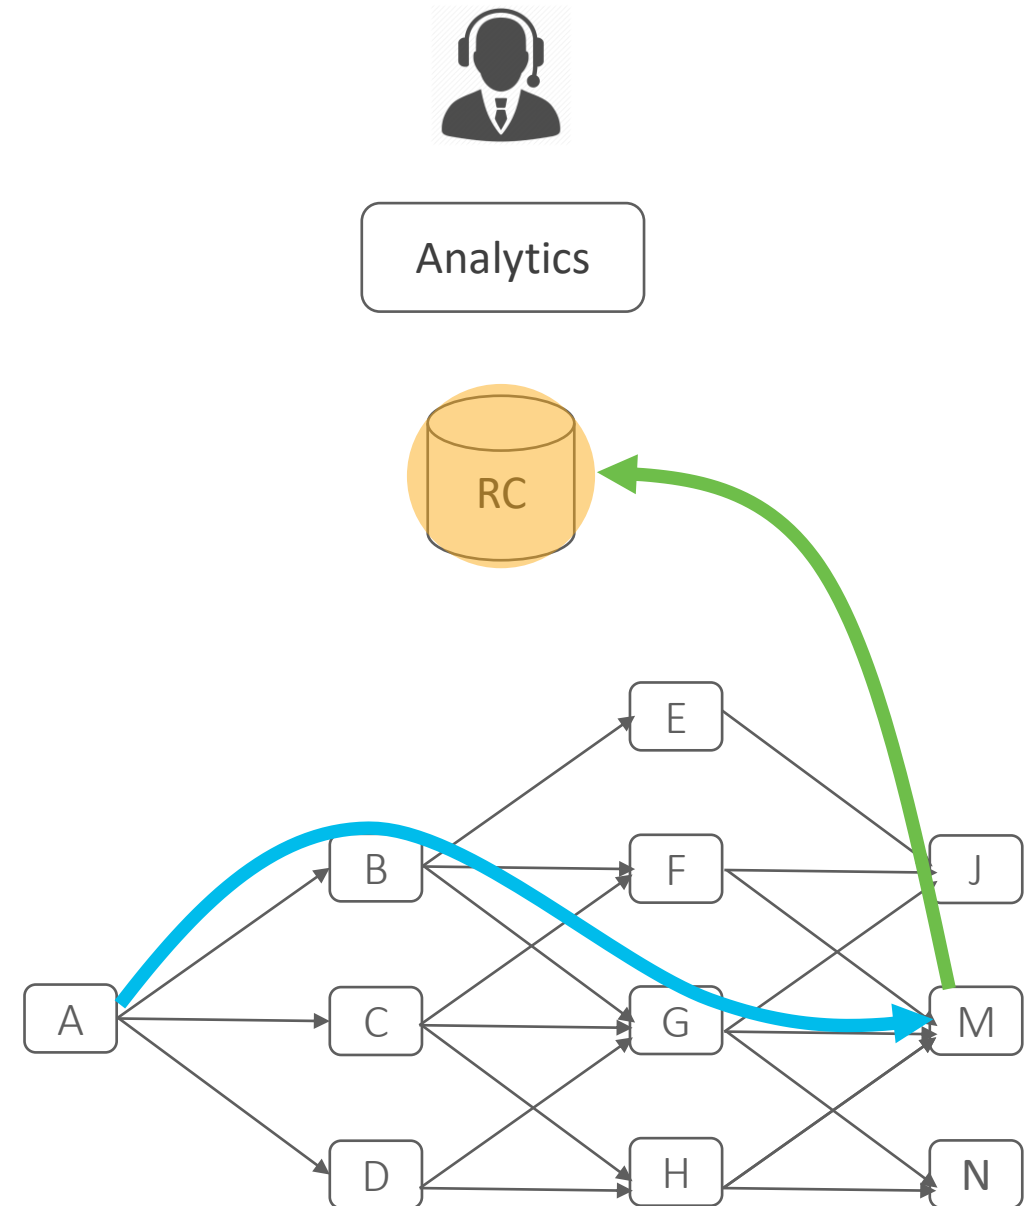
Sink

- Timestamp, Encapsulate and Forward (TEF)
 - Linerate
- Records reception data



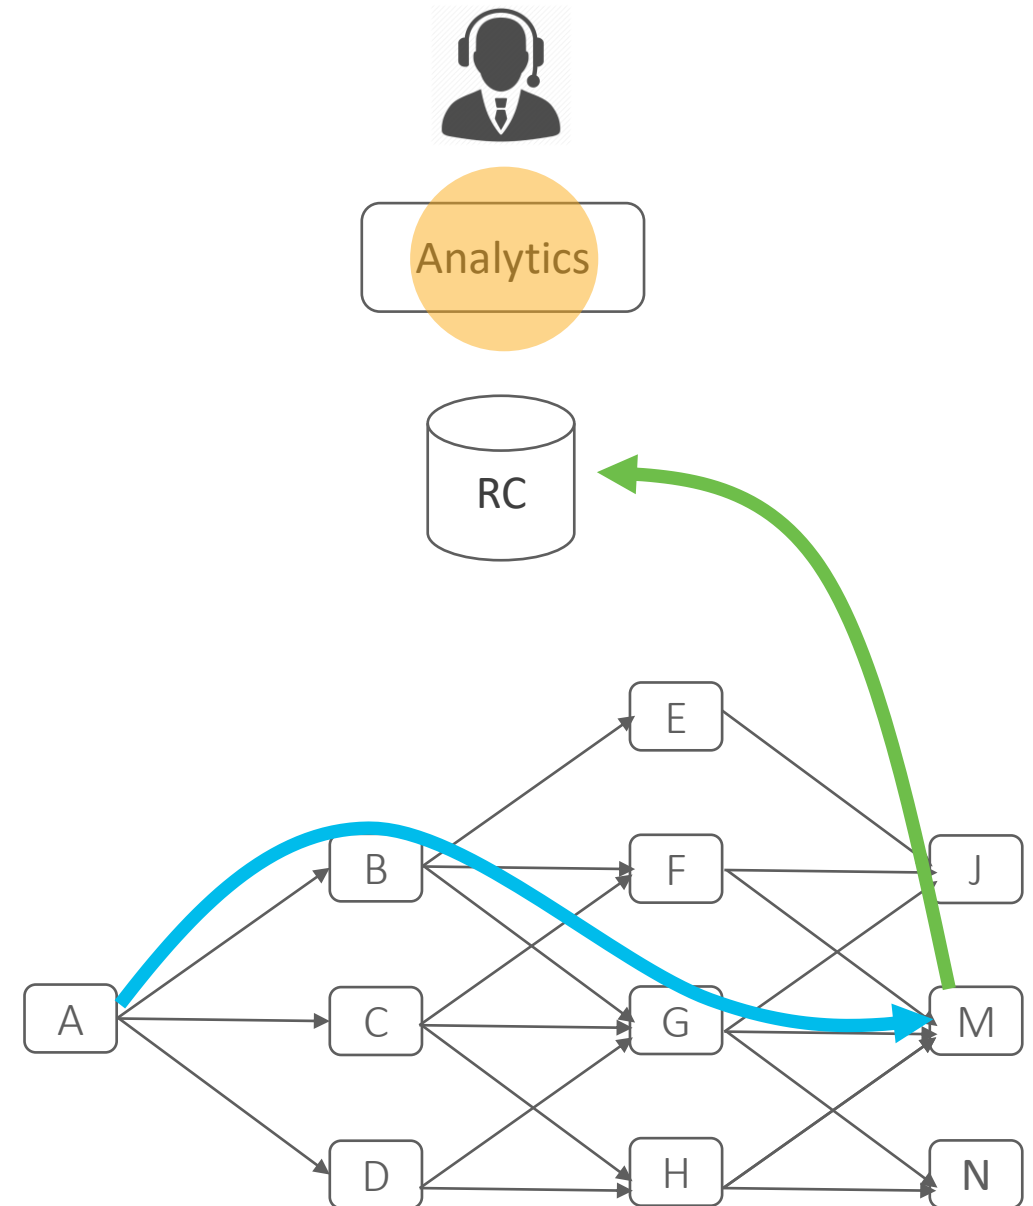
Regional Collector (RC)

- Ingest probes data in Time-Series Data store



Analytics

- Visualization
- Troubleshooting
- Alerts



Product, Deployment & Ecosystem

Product, Deployment & Ecosystem

- Cisco Shipping in CY22
 - C8000, NCS5500, ASR9000
- Strong Operator Interest
- Rich Eco-System
- Rich Open-Source



- Path Tracing is being standardized at IETF
 - <https://datatracker.ietf.org/doc/html/draft-filsfils-spring-path-tracing>
- More details and Demo: <https://www.segment-routing.net/path-tracing>

Conclusion

- Path Tracing
 - Operators can deterministically detect ECMP paths
 - Implemented at line rate in the base HW pipeline
 - Ultra-MTU-Efficiency
- Product, Deployment & Ecosystem
 - Cisco Shipping in CY22 (C8000, NCS55xx, ASR9000)
 - Strong Operator Interest
 - Rich Eco-System
 - Rich Open-Source
 - PT is being standardized at IETF

