




Path Tracing

Clarence Filsfils, Ahmed Abdelsalam, Pablo Camarillo, *Sonia Ben Ayed*



How did the packet arrive from A to F?

- 3 possible “**valid**” ECMP paths
 - Any drop?
 - End-to-End Latency homogeneity?
- An **invalid path** is possible
 - Routing or FIB corruptions
- 40-year-old unsolved IP problem



Stamping Trajectory in PT Header

- Each transit router records in PT header:
 - Outgoing interface ID
 - Timestamp (with 60μs accuracy in WAN; 200ns accuracy in DC)
 - Load
- Highly compressed for low MTU overhead
 - Only 3 bytes per hop!
- Implemented at line rate: **Reports true packet experience**
- Native interworking with legacy nodes
 - Seamless deployment
- Hardware/XR feature (shipping in 7.8.1) with analytics app



Mature Eco-System

- PT Midpoint Shipping – IOS XR 7.8.1
 - Cisco 8000 (Silicon One Q200; native SDK)
 - NCS5700 (DNX2 – J2; native SDK)
 - ASR9000 (LS)
- Rich Eco-system
 - Cisco, Broadcom, Marvell, +others
 - SAI/SONiC in progress
 - Linux, FD.io VPP, P4, Wireshark, TCPDUMP
- Ongoing standardization
 - Path Tracing in SRv6 networks (ietf.org)



MARVELL





The bridge to possible