



OCP
SUMMIT

March 20-21
2018
San Jose, CA

OPEN. FOR BUSINESS.



An Evolution of Network Telemetry

Tal Mizrahi
Marvell

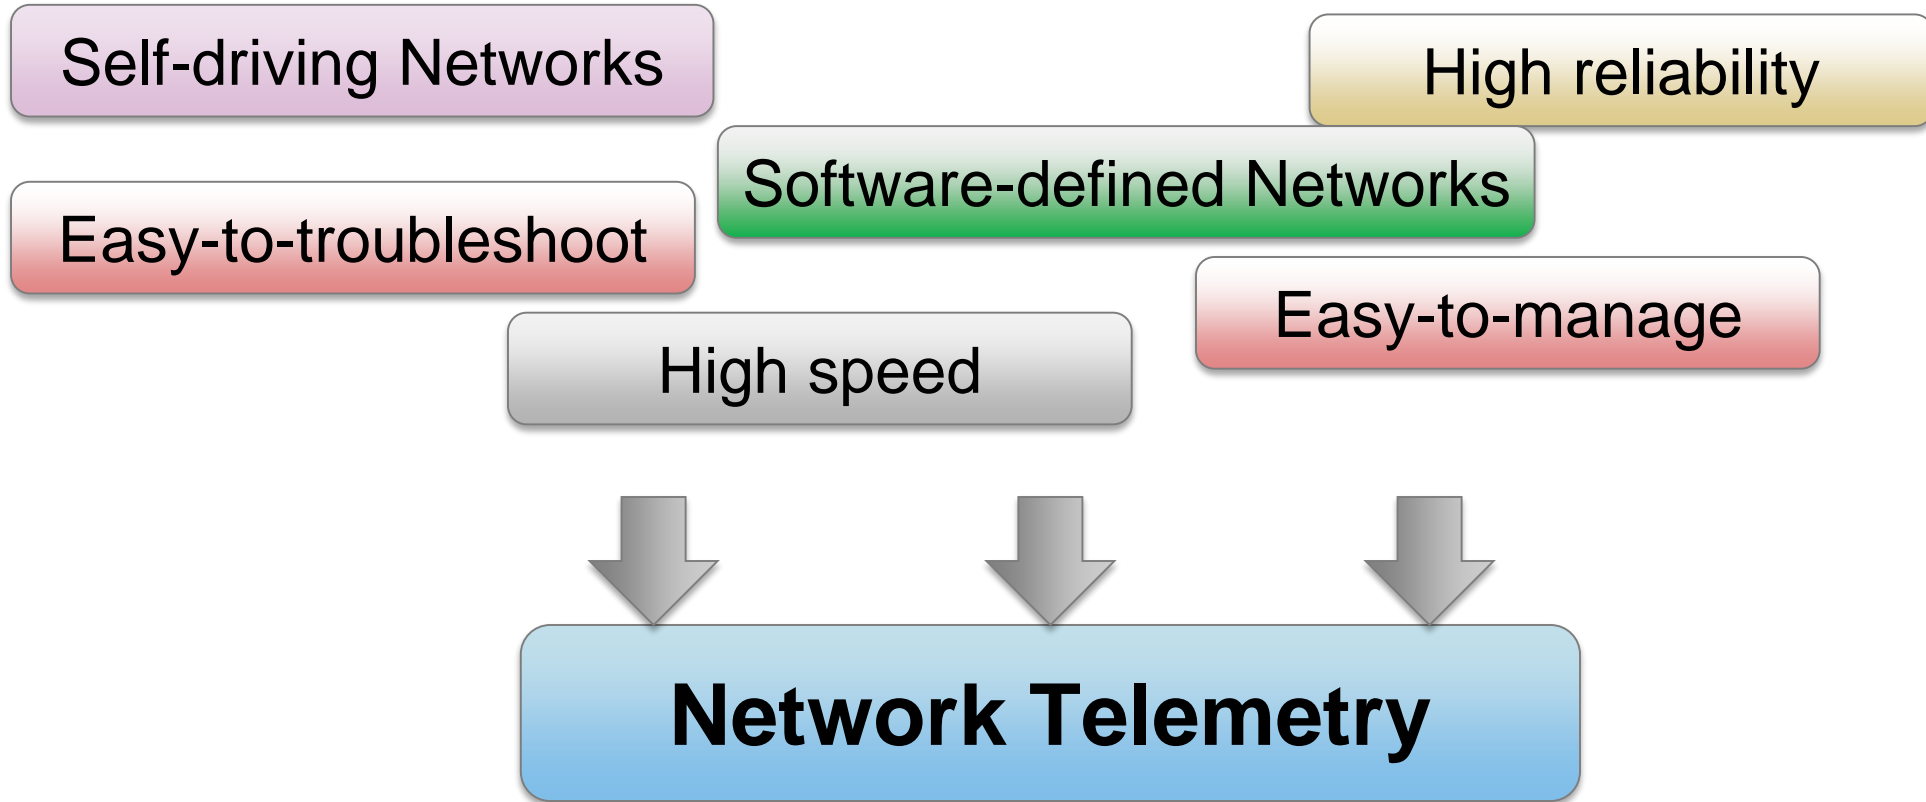
Telemetry BoF



OPEN. FOR BUSINESS.



Network Telemetry



Are we reinventing the wheel?



Ping / Traceroute


Reinventing the wheel?



```
C:\Windows\system32\cmd.exe

C:\>ping www.marvell.com -n 10

Pinging extranet.marvell.com [10.68.68.50] with 32 bytes of data:
Reply from 10.68.68.50: bytes=32 time=211ms TTL=57
```



Old-School Passive Monitoring



Counters

Per port

Queue State

Per flow

Latency

Per queue

⋮

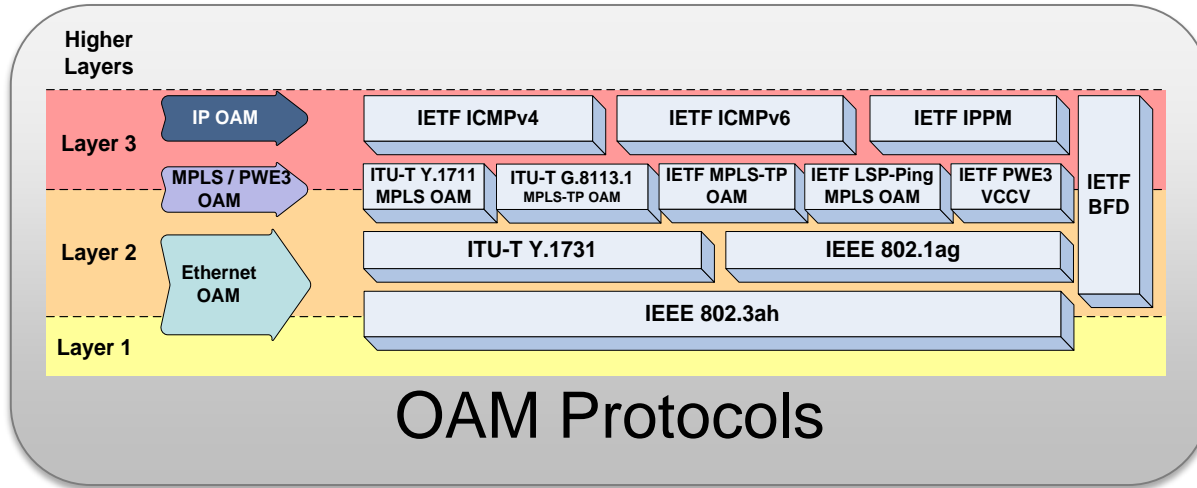
⋮

Reinventing the wheel?




Carrier Network OAM

Reinventing the wheel?

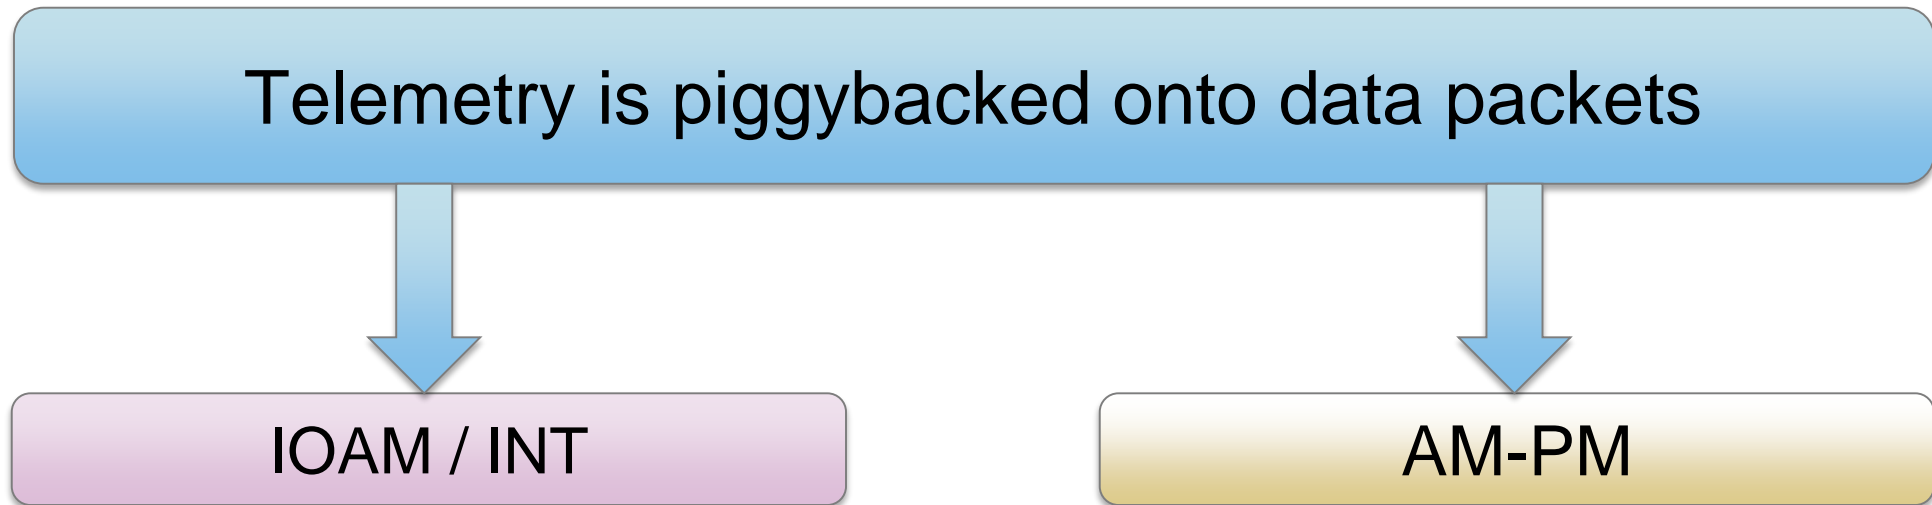


Active measurement / monitoring:



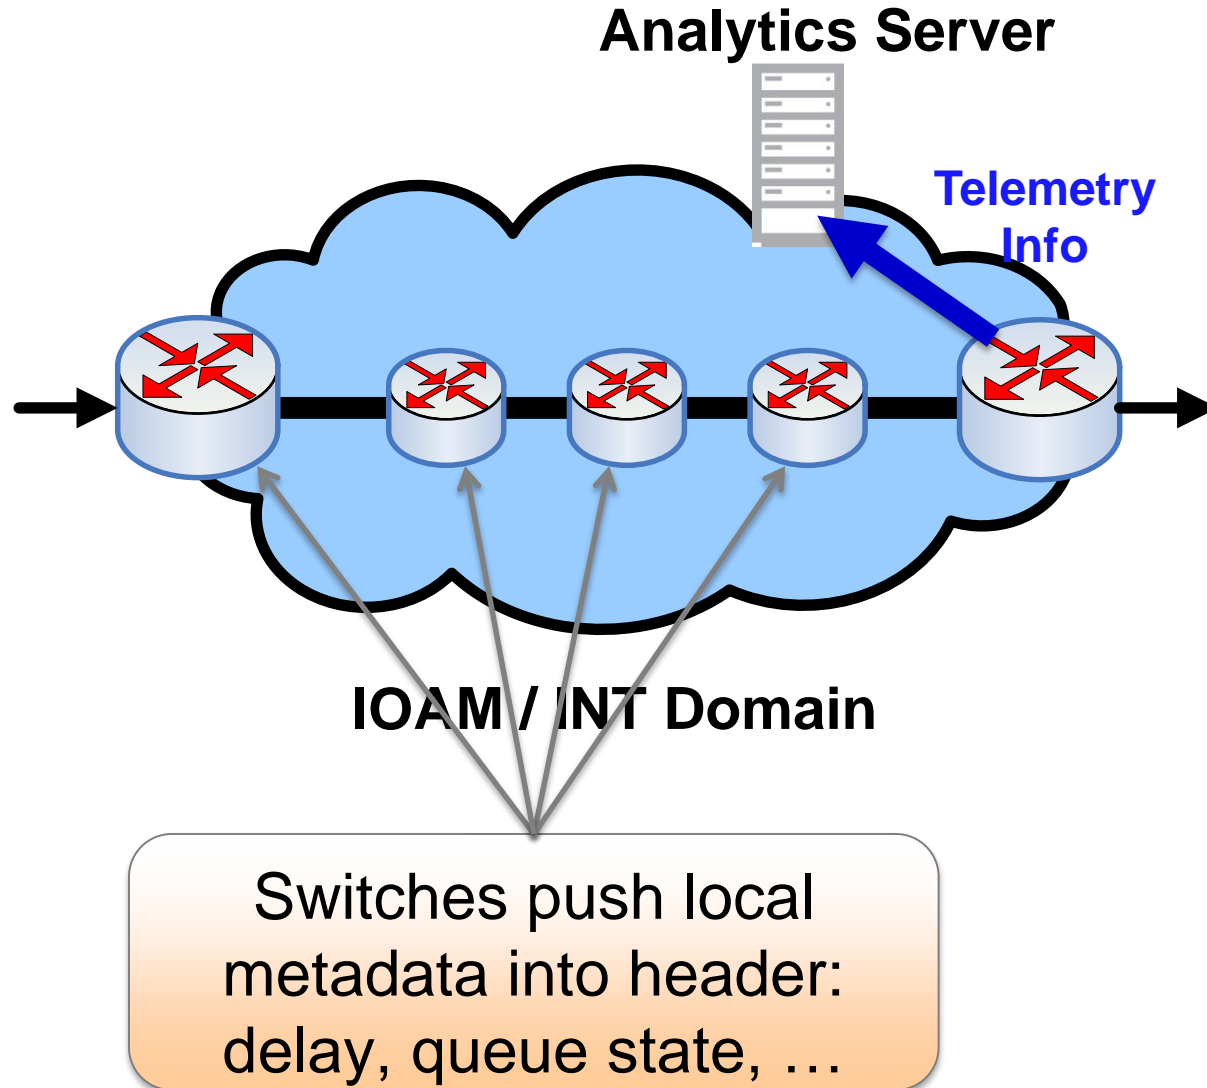
Piggybacked Measurement

Reinventing the wheel?



Piggybacked Metadata – IOAM / INT

Reinventing the wheel?




IOAM (IETF)
INT (P4)
+SAI TAM

AM-PM: Alternate Marking – Performance Measurement

(RFC 8321)

Reinventing the wheel? 



Packets Sent: 10,000
Packets Received: 9,500
Packets Lost: 500



Network Telemetry: An Evolution

Reinventing the wheel?



NO !



**Ping
Traceroute**



**Passive
Monitoring**

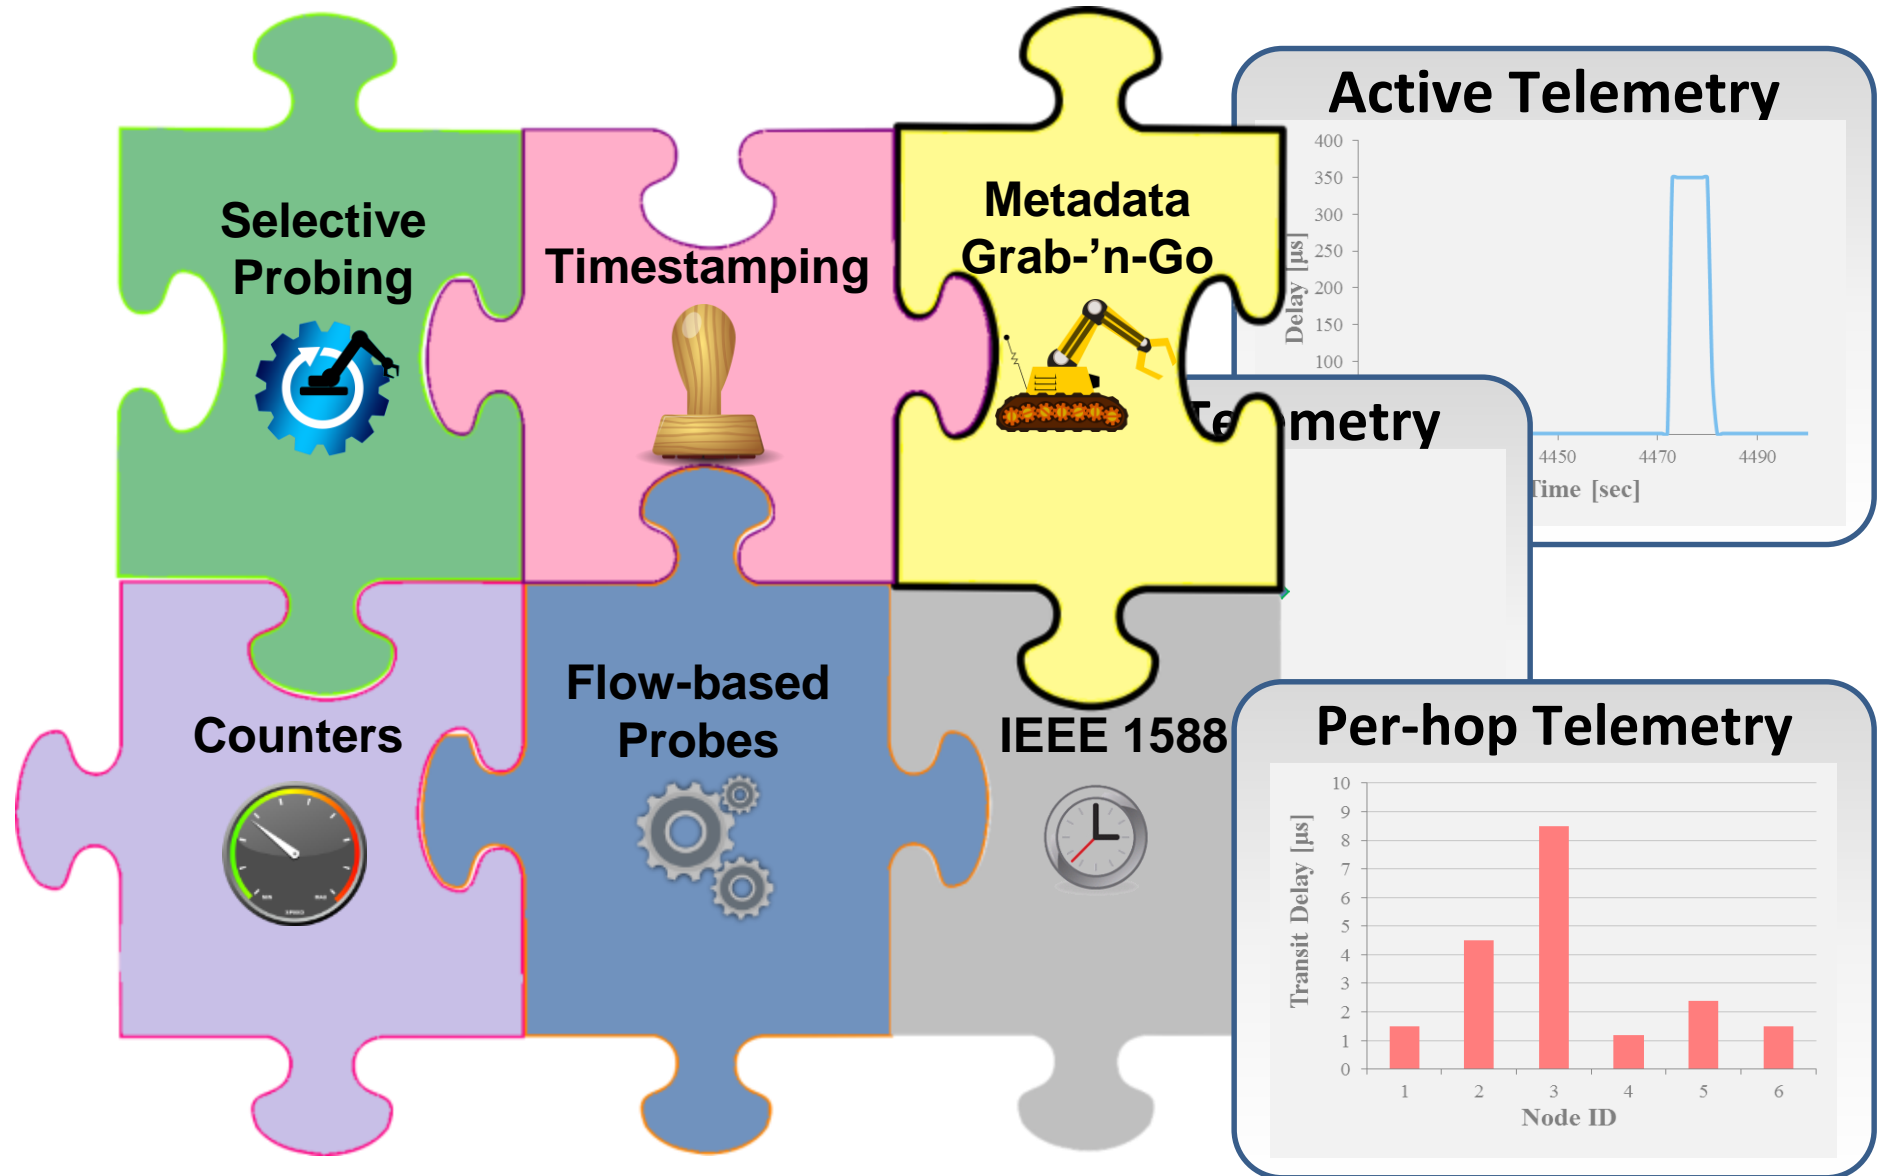


**Carrier
OAM**

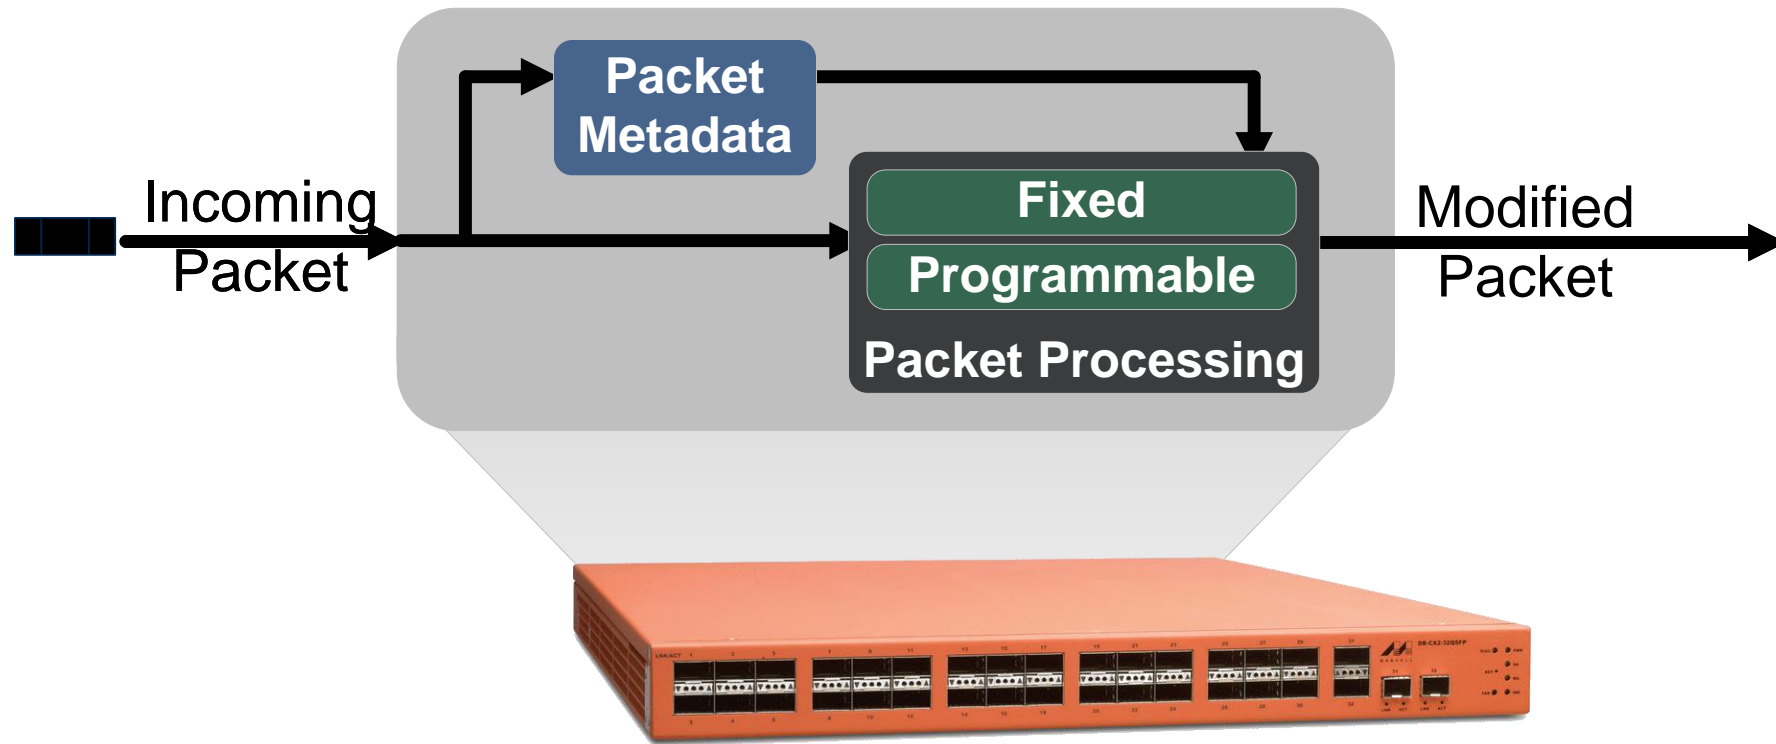


**IOAM / INT
AM-PM**

Marvell's Network Telemetry Solutions




Metadata: Grab-'n-Go



Marvell Prestera


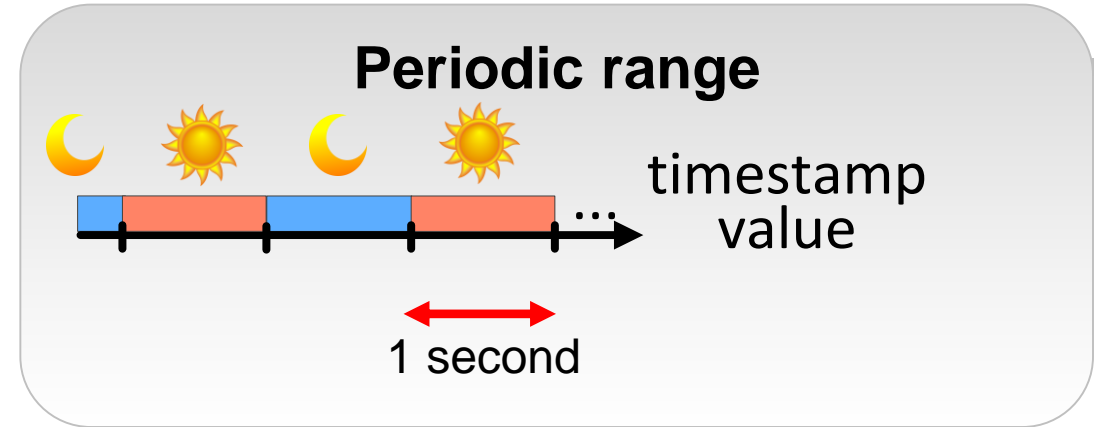
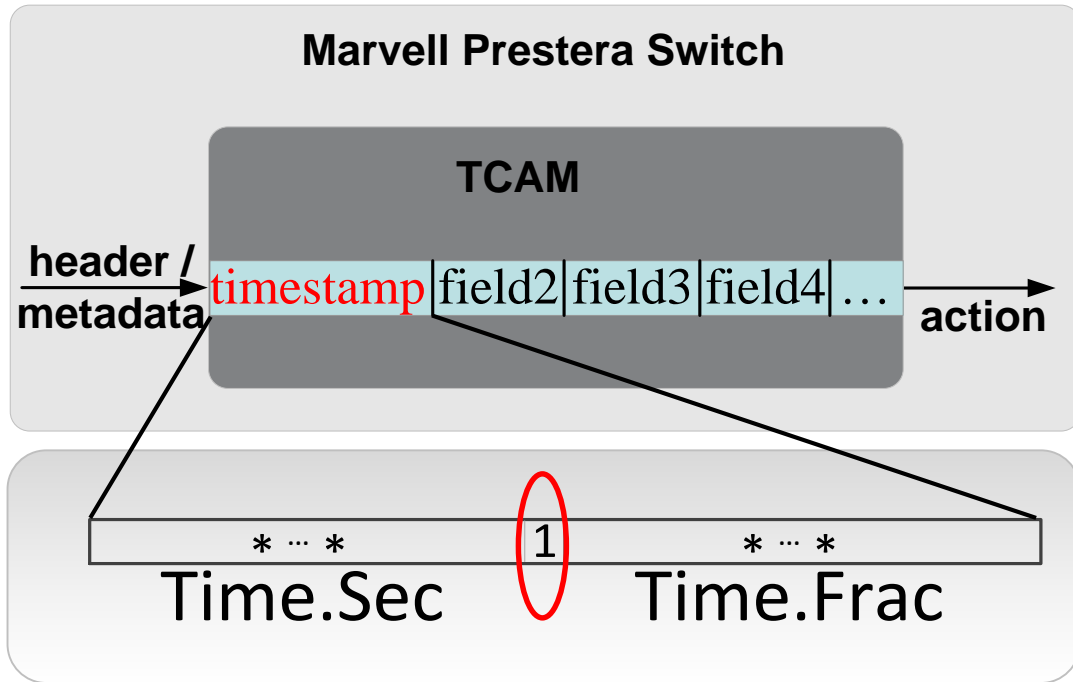
Timestamping Everything



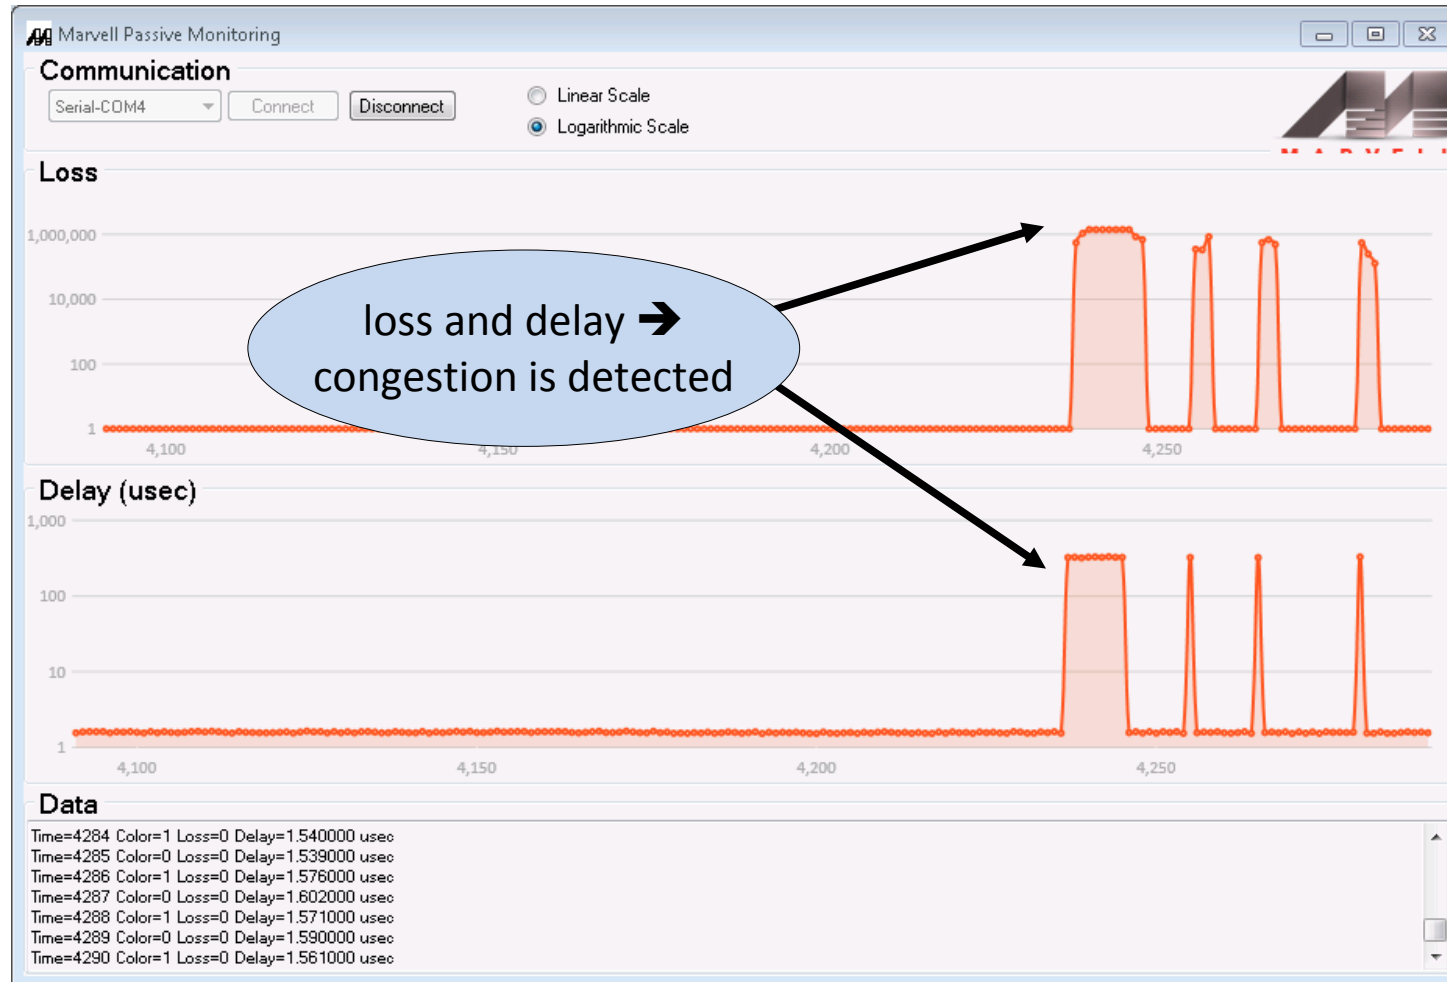
Marvell Prestera

- ✓ Periodic probing
- ✓ AM-PM
- ✓ TimeFlip


AM-PM using TimeFlip



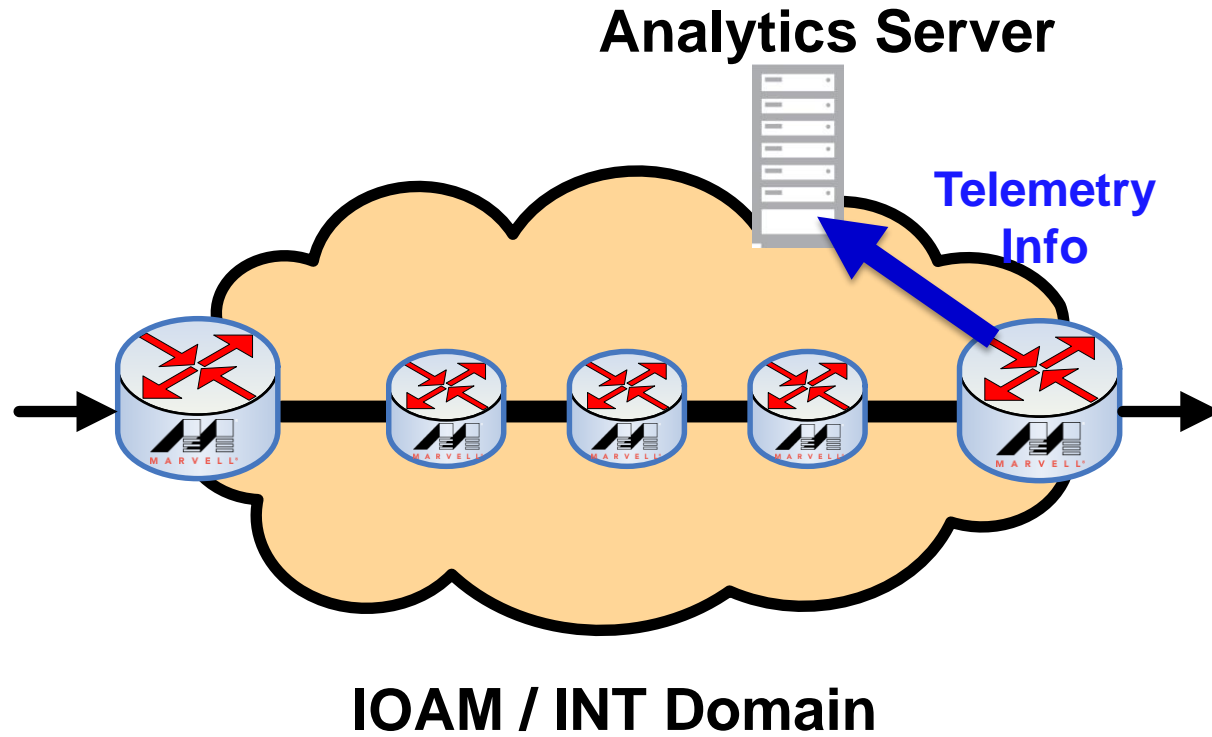
Per-Flow Congestion Detection using AM-PM




Selective Exporting



IOAM/INT: Periodic Exporting

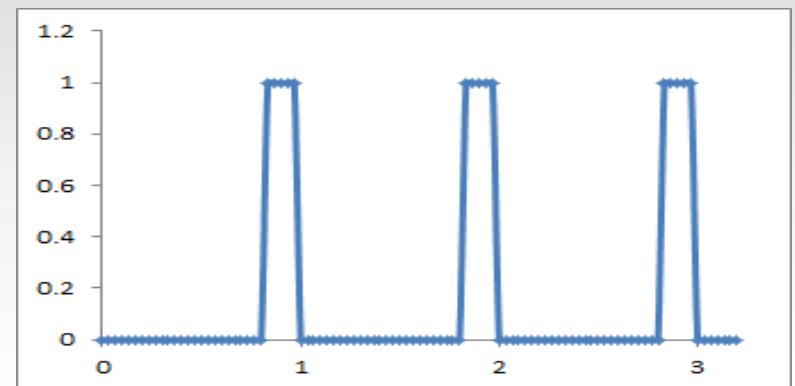


Export one packet per second




Per flow / per port

Export last ms of every second




IOAM/INT: Adaptive Exporting



Event-driven:

- ✓ On drop
- ✓ On congestion
- ✓ On high rate

Combining IOAM/INT with AM-PM





Challenge: export telemetry info without the expensive overhead of IOAM/INT.

AM-PM as a trigger for exporting telemetry info.

Per-hop Telemetry



The Network Telemetry Toolset



Marvell Prestera



OCP SUMMIT

References

- [1] Mizrahi, T., Vovnoboy, V., Nisim, M., G. Navon, and A. Soffer, “Network Telemetry Solutions for Data Center and Enterprise Networks”, Marvell white paper, 2018.
- [2] Brockners, F., Bhandari, S., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Mozes, D., Lapukhov, P., Chang, R. and D. Bernier "Data Fields for In-situ OAM", [draft-ietf-ippm-ioam-data-00](#), work in progress, 2017.
- [3] C. Kim et al., “[In-band network telemetry \(INT\)](#)”, P4 consortium, 2015.
- [4] Fioccola, G., Capello, A., Cociglio, M., Castaldelli, L., Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi, “Alternate Marking method for passive and hybrid performance monitoring”, [RFC 8321](#), 2018.
- [5] Mizrahi, T., Rottenstreich, O. and Y. Moses, “TimeFlip: Scheduling Network Updates with Timestamp-based TCAM Ranges”, IEEE INFOCOM, 2015.