In-band Network Telemetry using Data Plane Programming

Mauro Campanella (GARR), Tomas Martinek and Mario Kuka (CESNET), Joseph Hill (UvA), Matteo Gerola (FBK), Jakub Kabat (PSNC), Marinos Demolianis and Nikos Kostopulos (NTUA), Theodore Vasilopoulos(GRNET)

2nd Telemetry and Data Workshop
6 April 2022

www.geant.org

In-Band Network Telemetry (INT) short summary

INT alters structure of selected packets, in the fly, to collect and transport information in-band, in real time

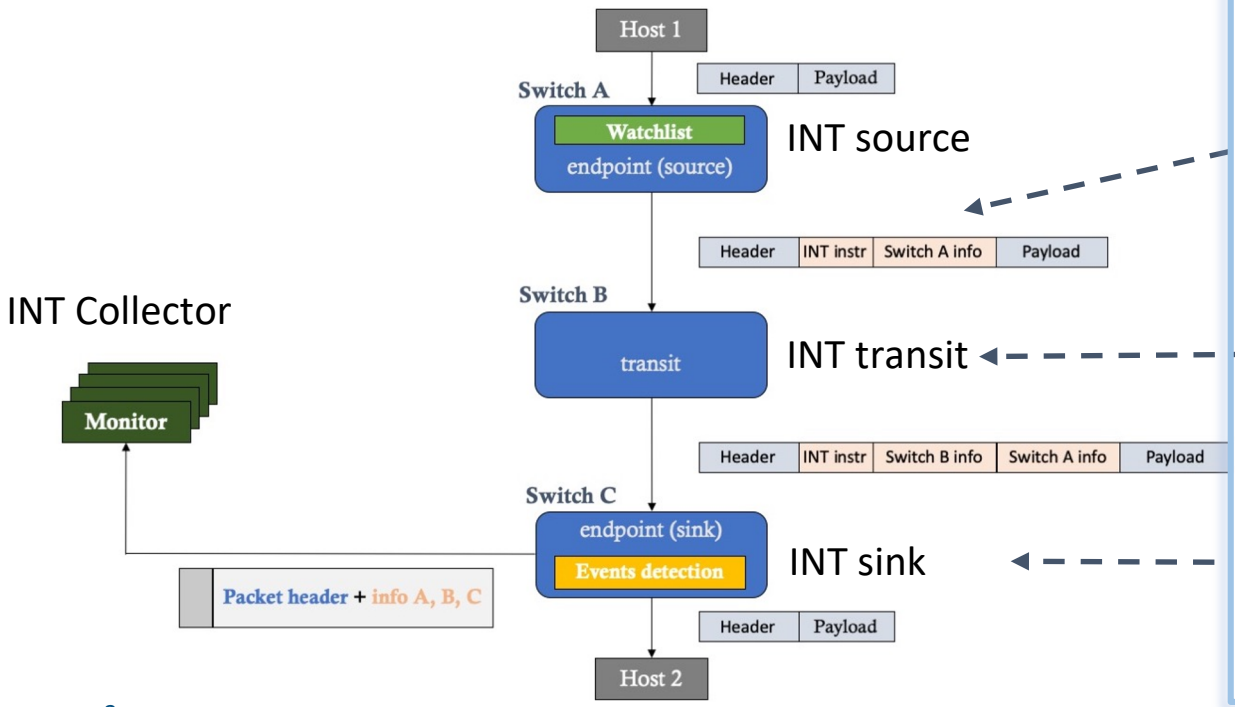

INT functions

INT source node adds a small **INT header** to **every chosen packet** containing e.g. Switch IDs, Interfaces IDs, Timestamps, Link and queue utilization

INT transit nodes add specific local

The last **INT sink** node extracts, may analyze and and sends informstion to a collector

INT as data plane programming application

- Alters the structure of every packet with "in-house" logic
- Applies "non switching" logic to packets
- Personalizes monitoring and fast-feed information to control plane
- Offers a tool that permits precision monitoring in faster, larger and more automated networks.

Investigation with tools started to be available in 2018:

Programming Protocol-Independent Packet Processors :
High level, C-like, coding language for controlling packet forwarding planes in networking devices

- New silicon, P4 enabled (Tofino, FPGA …), working at line speed enabling programming, Telemetry (push vs pull model),…

Our INT testbed over production NREN networks

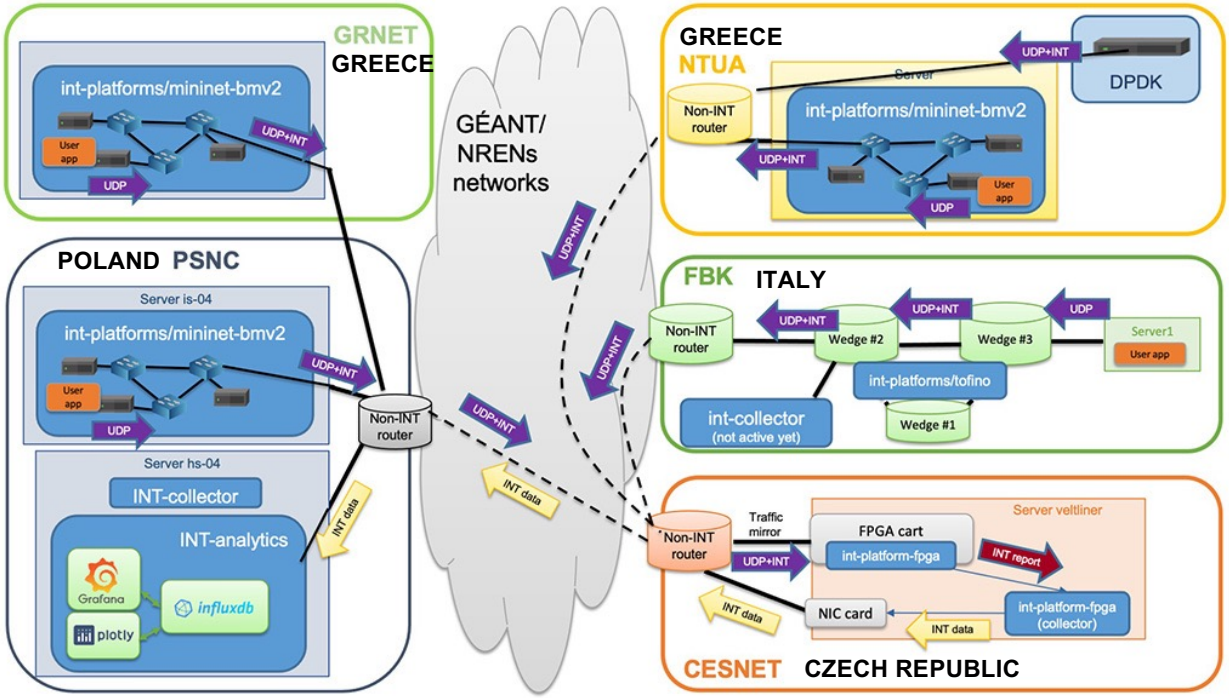

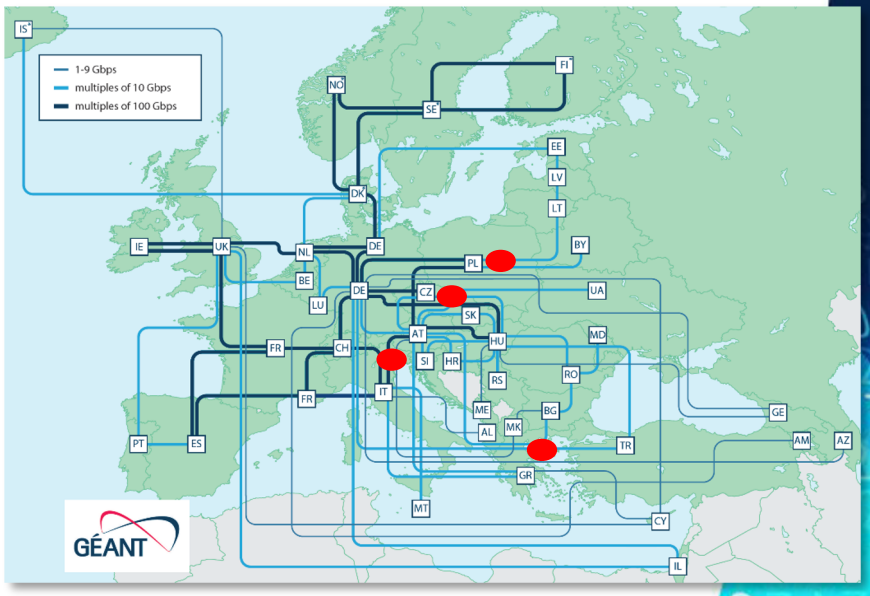

- Packet carries timestamps, sequence number in INT headers between Source (all) and Sink node (CESNET)
- UDP packets generated at constant rate ~1k to 300k pps

- 4 switch platforms
- UDP packets flow in NRENs networks
- Collected INT data in CESNET is sent to PSNC for collection and presentation.

INT test set-up (FBK to Cesnet)

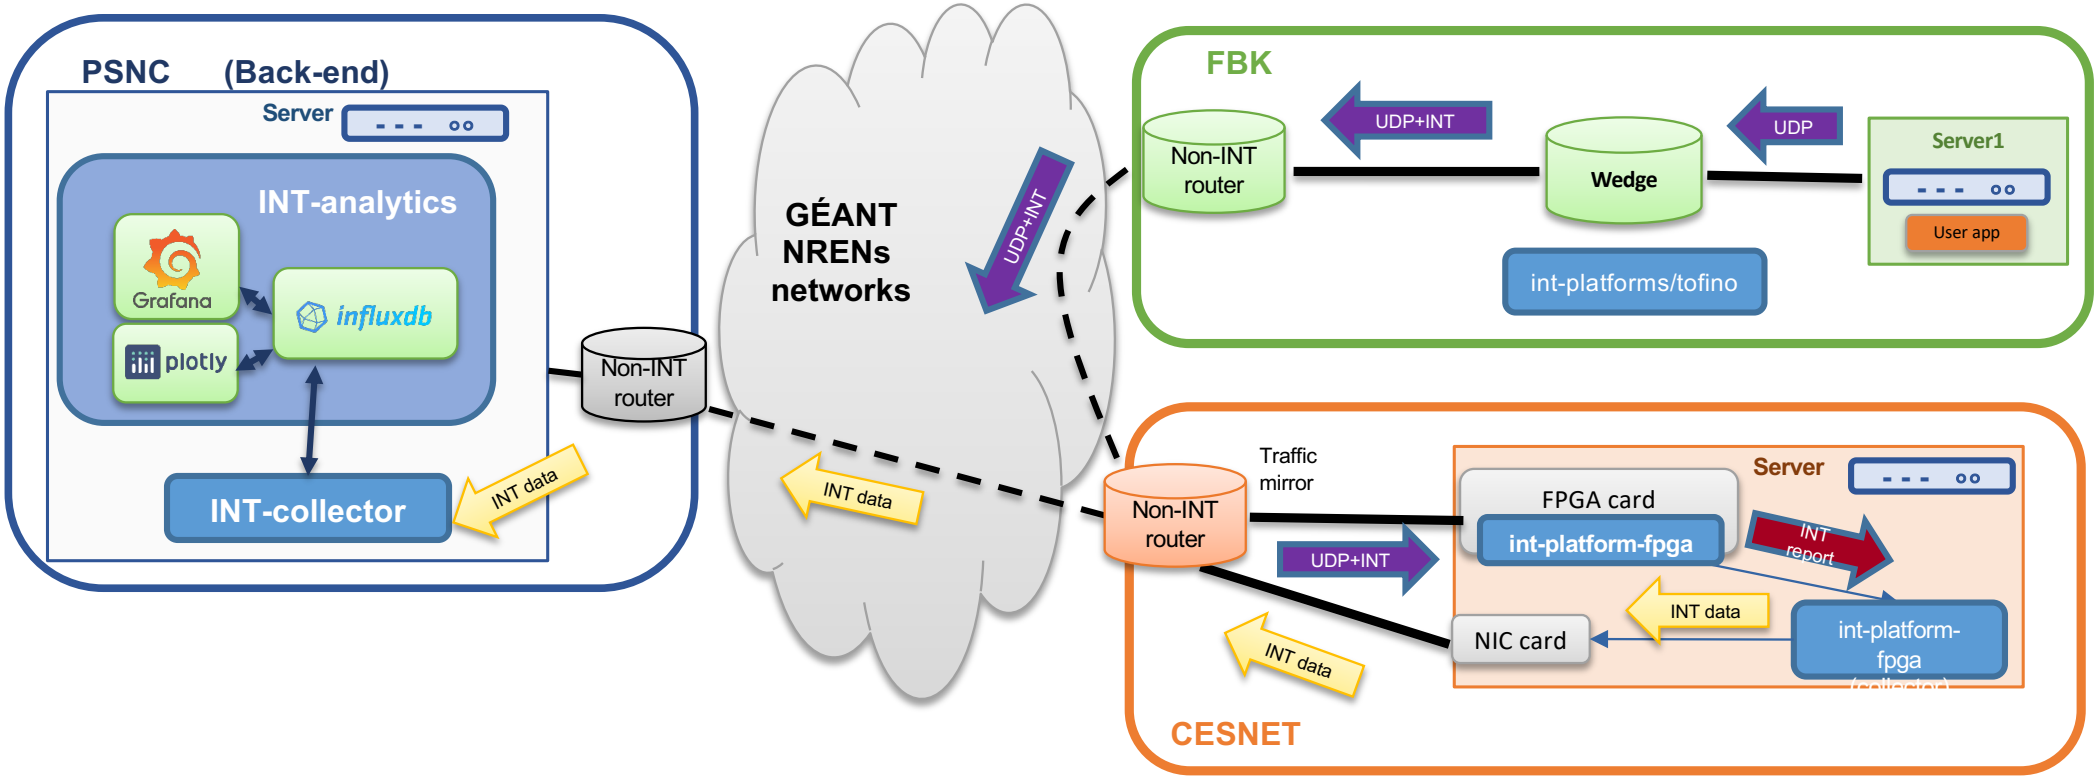

INT P4 code developed for these platforms (BMv2, Tofino on GitHuB)

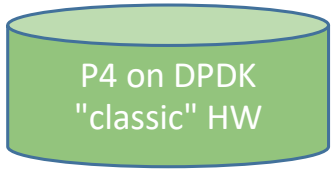

DPDK (Data Plane Development Kit) kernel software acceleration, available for 2 P4 to DPDK compilers (T4P4S and P4C-DPDK)

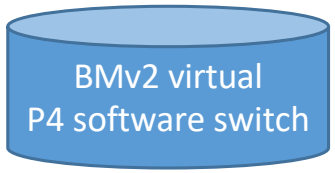

BMv2 Behavioral Model v2 – emulation of Tofino
Uses Mininet

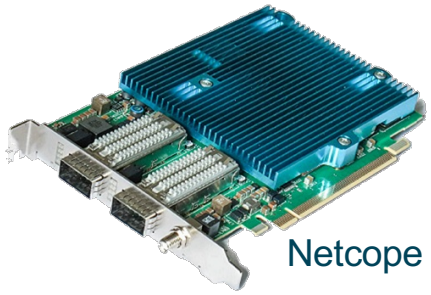

Netcope NFB-100G2

FPGA P4 compiler developed by Cesnet up to 2x100 Gbps

Edgecore Wedge100BF-32X

Tofino Asic engineered for dataplane programming by Barefoot (now Intel)

5 minutes of the INT tags in a single UDP flow PSNC to CESNET

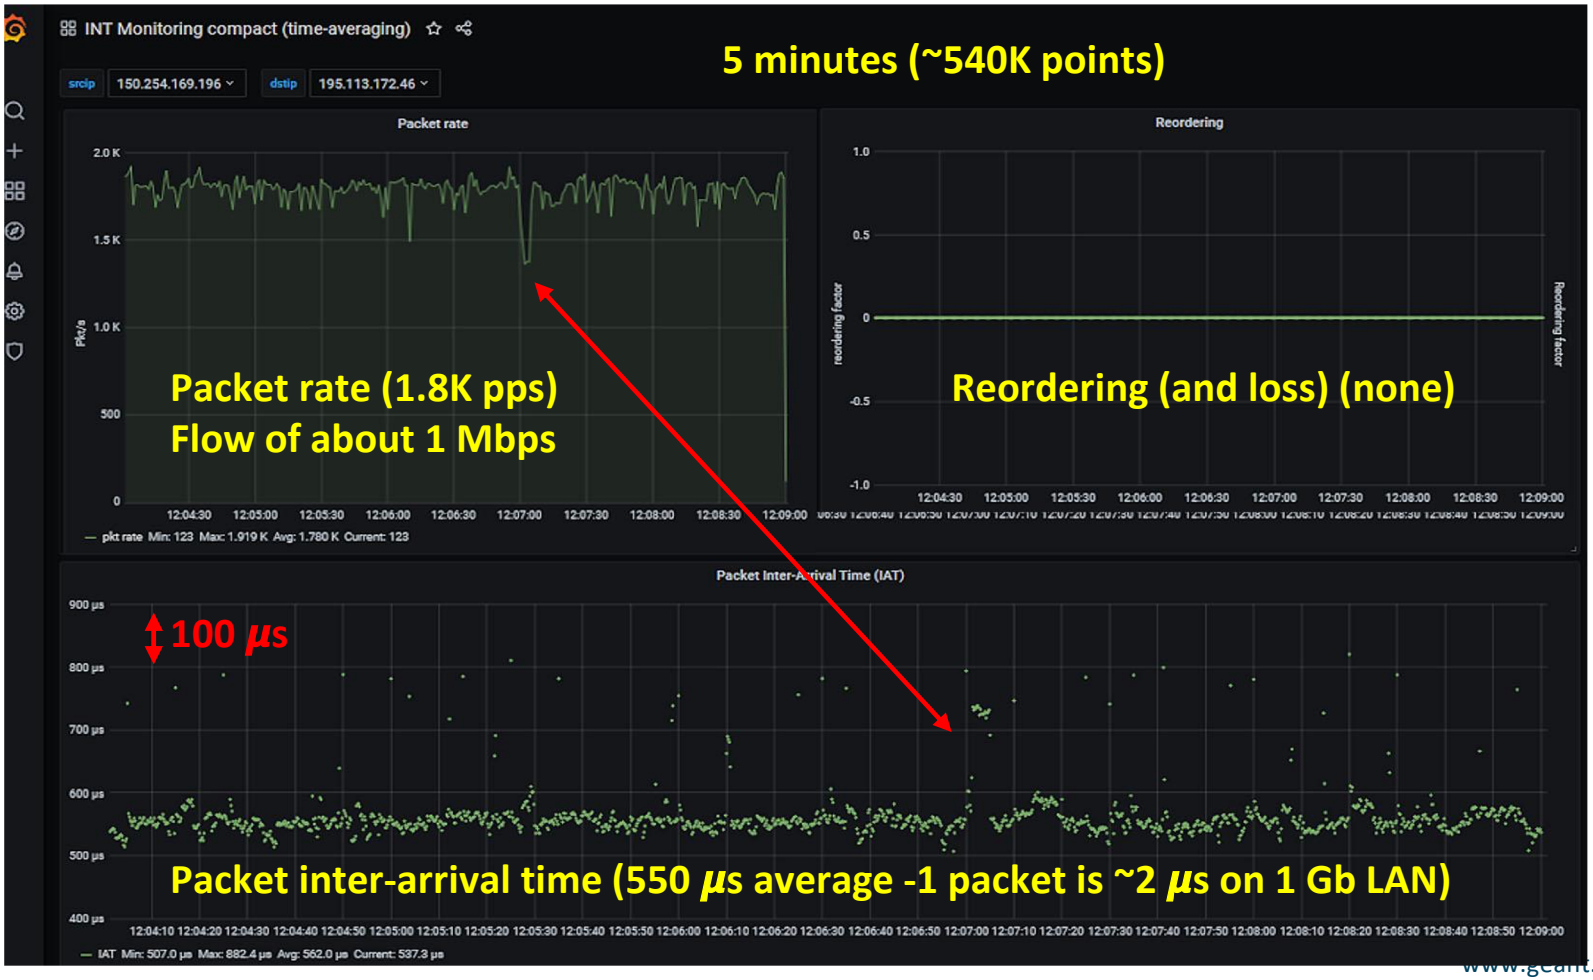

INT Platforms: Lessons Learned

- BMv2/mininet: good to set-up initial P4 program development (but performance/jitter ceiling)
- Tofino switches: feature-rich for P4
  - Clock timestamping, licensing issues
- FPGA card: fast, flexible HW
  - P4 compiler quality: vital;
  - HW insight is required
- INT-DPDK :
  - Promising performance up to few Gbps
  - Needs careful selection of Network Interface Cards, (timestamping, etc…)
  - P4 to DPDK compiler quality may vary

10 ms of inter-arrival packet time, 1 UDP flow of 260 K pps (~3.3 μs average) FBK to CESNET (1 packet is ~2 μs @ 1 Gb LAN)

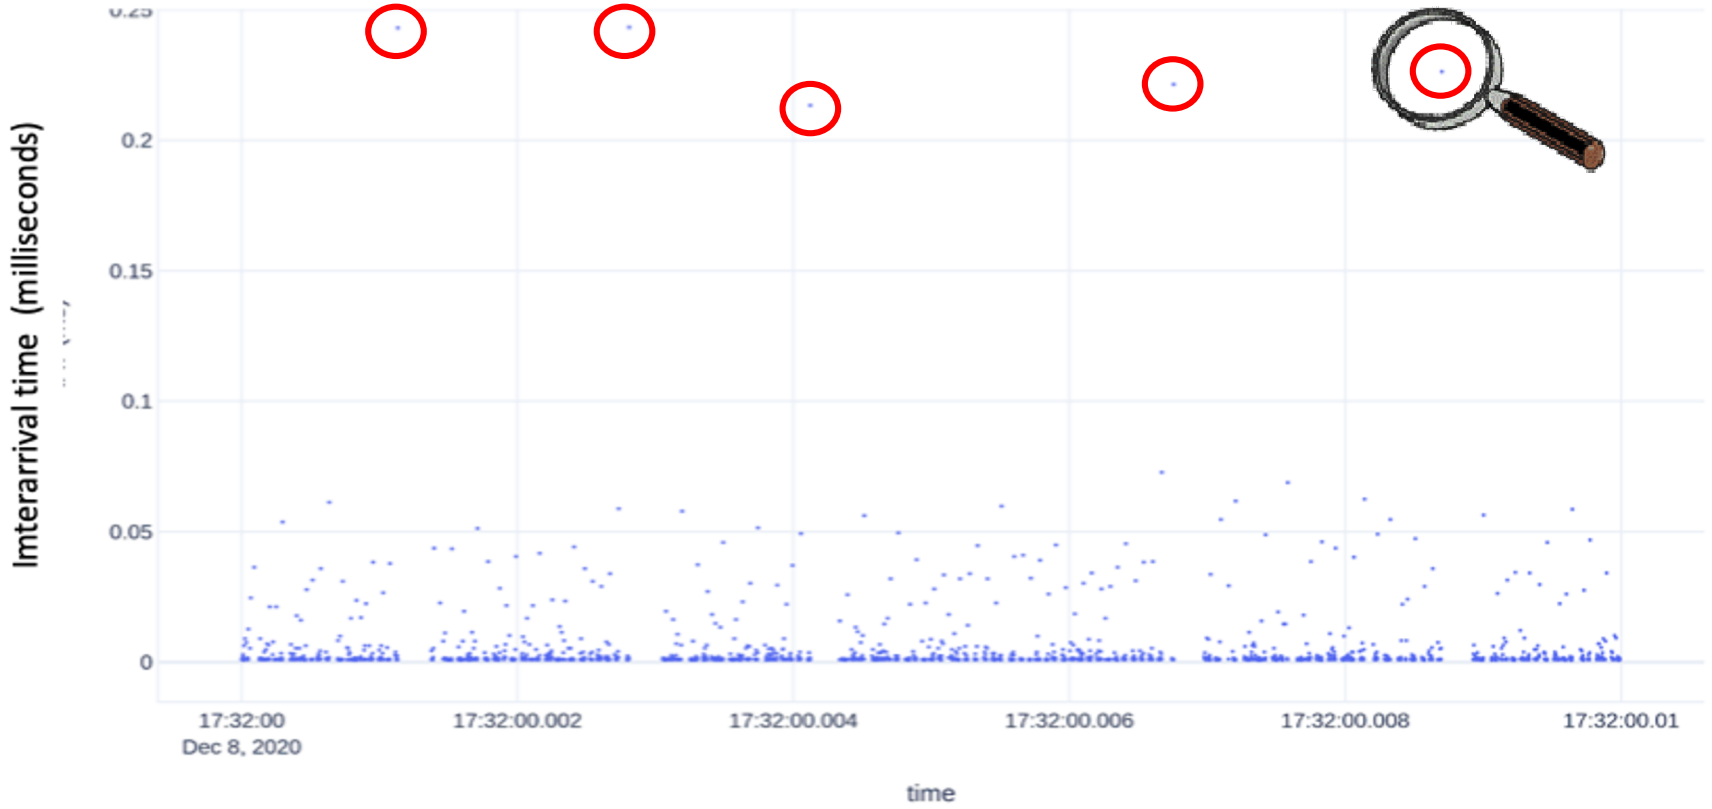

Current activity

- Bug fixes of the code for the four platforms
- Testing a new collector with rule based parsing to reduce data total size
- Systematic measurement / calibration of behaviour in a LAN environment

Then

- Systematic measurement in WAN of production environment

11

Results

- Data Plane Programming) (using P4) is not business-as-usual, requires specific expertise (fastly changing, mixing ICT, network and data)
- INT offers an agile tool for high frequency monitoring, supporting new control plane mechanisms and debugging in real time.
- The network behaviour at microsecond, or lower, scale looks complex, and transient effects are evident also at milliseconds range
- INT can operate at application level e2e between cooperating end-sites without imposing agreements on production backbones
- Time stamping is not a standard and precise time synchronization between cooperating node may be mandatory requirement for some use cases

12

Closing remarks

- Standard ICT HW works up to 1 Gbps links, with some optimization. Specific ASIs are required above 1Gbps
- High frequency monitoring needs to be understood to evaluate its usefulness, including when information is fed to with machine learning.
- INT/P4 use may/will generate and require handling of large amount of "raw" data, to be used for analytics and more.
- Having “big data” opens the possibility to use Machine Learning analisys tools

13

Acknowledgements

The effort reported has received essential contributions from many GÉANT participants and specifically:

Damu Ding (Italy), Federico Pederzolli (Italy), Pavel Benacek (Czech Republic),

Tim Chown (Jisc UK), Ivana Golub (PSNC Poland),

Xavier Jeannin (RENATER France)

Thanks !

More information on GÉANT Data Plane Programmability activity

- **Data Plane Programming** / **INT GEANT web page** https://wiki.geant.org/display/NETDEV/INT Includes all documents produced and a **pointer to GitHub INT P4 code**
- **Mailing list:** https://lists.geant.org/sympa/subscribe/int-discuss,
- **White Paper INT Tests in NREN networks** – DPP WP6 T1 white paper https://www.geant.org/Resources/Documents/GN4-3_White-Paper_In-Band-Network-Telemetry.pdf
- **DDoS Paper**: "In-Network Volumetric DDoS Victim Identification Using Programmable Commodity Switches", F. Pederzolli, M. Campanella and D. Siracusa, in IEEE Transactions on Network and Service Management, Vol. 18, Issue: 2, June 2021, page: 1191-1202, DOI: 10.1109/TNSM.2021.3073597 and at https://arxiv.org/abs/2104.06277
- The GÉANT **First Telemetry and Big Data Workshop** https://wiki.geant.org/display/PUB/Telemetry+and+Big+Data+Workshop

17

Non GEANT References

- **The Programmable Data Plane Reading List** : https://programmabledataplane.review/
- Oliver Michel, Roberto Bifulco, Gábor Rétvári, Stefan Schmid, "**The Programmable Data Plane: Abstractions, Architectures, Algorithms, and Applications"**, ACM Computing Surveys, Volume 54, Issue 4, May 2021, Article No.: 82, pp 1–36, https://doi.org/10.1145/3447868 10.36227/techrxiv.12894677.v1 https://www.univie.ac.at/ct/stefan/csur21.pdf
- "**A Survey on Data Plane Programming with P4: Fundamentals, Advances, and Applied Research"**, Frederik Hauser, Marco Häberle, Daniel Merling, Steffen Lindner, Vladimir Gurevich, Florian Zeiger, Reinhard Frank, and Michael Menth (50 pages).26 Jan 2021, to be published in" Communications Surveys & Tutorials (COMST) journal --https://arxiv.org/pdf/2101.10632.pdf
- Kaur, Sukhveer & Saluja, Krishan & Aggarwal, Naveen. (2021). "**A review on P4-Programmable data planes: Architecture, research efforts, and future directions"**. Computer Communications. 170. 10.1016/j.comcom.2021.01.027.
- **IOAM**: https://datatracker.ietf.org/wg/ioam/about/

Thank you

Any questions?

Mauro.Campanella@garr.it

www.geant.org

The research leading to these results has received funding from
the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 856726 (GN4-3).