

IT Open Day 2022

Packet Capturing for Fun and Profit
Dr. Georg Sauthoff



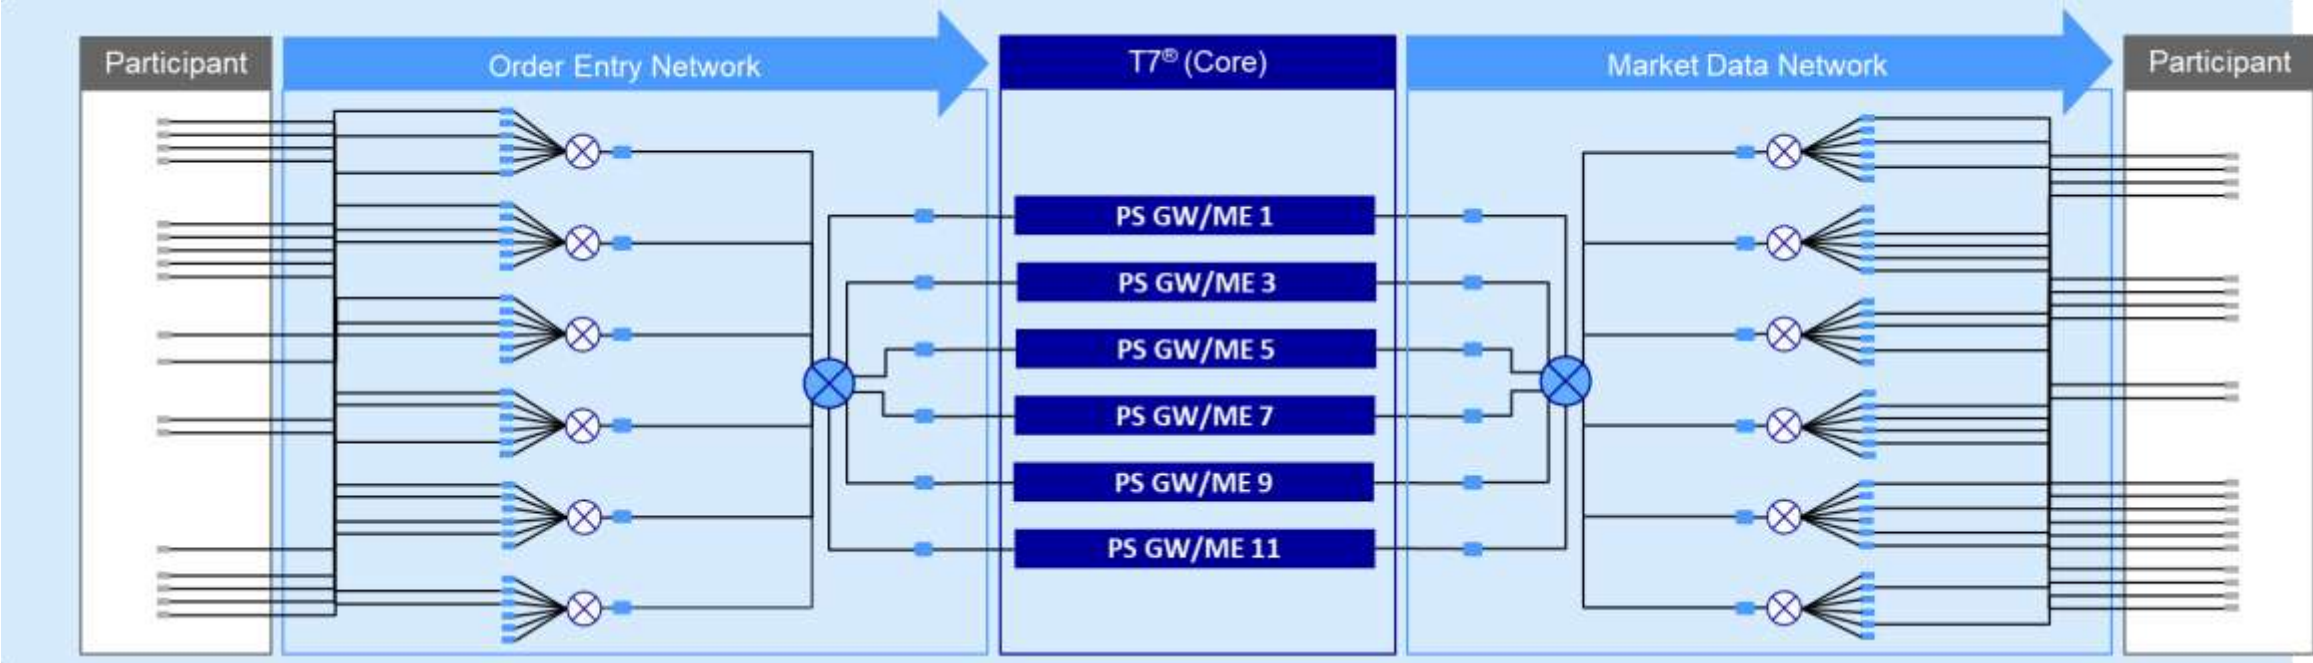
Outline

- Motivation
- Status Quo
- Non-Solutions
- T7 Solution

Motivation

- High Precision Timestamps (HPT) Services
- Troubleshooting
- Increasing volumes to capture
- Demand for extra TAPs to improve our monitoring

Motivation



Colo 2.0 Topology

- ⊗ Cisco 3548X Distribution Layer Switch
- ⊗ Cisco 3548X Access Layer Switch
- Optical TAP (capture and time measurement)

Status Quo

- Two-Socket Intel servers with a bunch of Solarflare NICs each
- Running SolarCapture Pro
- TAPs are aggregated and timestamped by Metamako devices
- Issues
 - Unclear future of SolarCapture
 - High CPU usage
 - Easy to drop packets even after extensive tuning

Non-Solutions

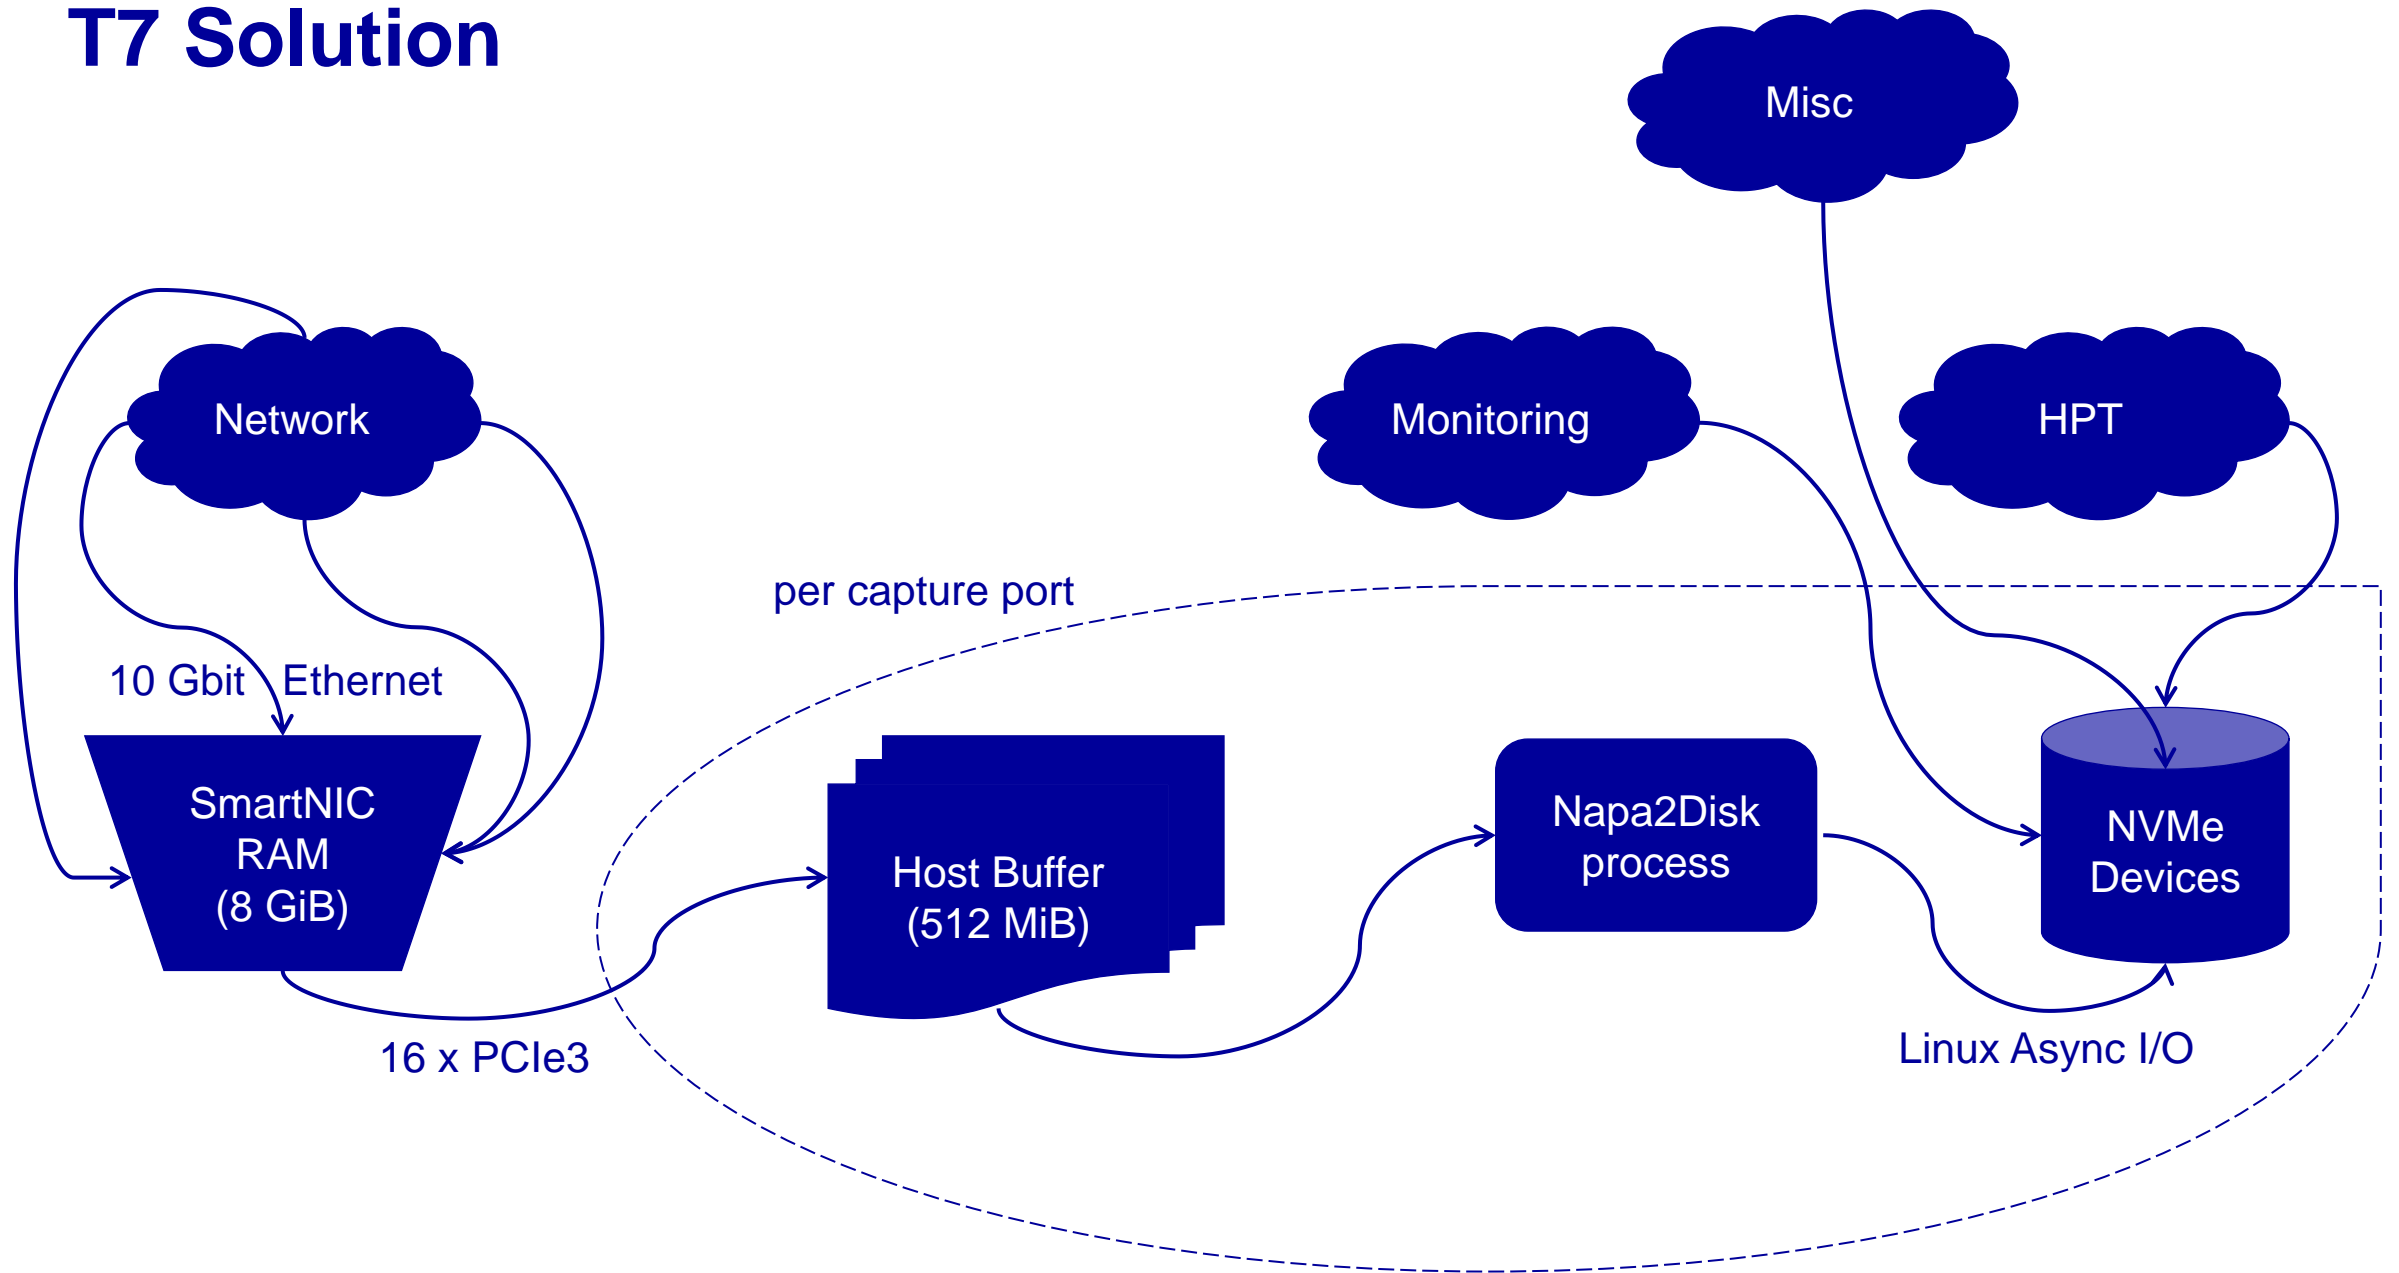
- SolarCapture Pro
- FMADIO capture appliance (as of 2021)

T7 Solution

- Requirements:
 - Capture 4 x 10 Gbit ports per host at line-rate/minimal packet size, in wire order
- Single Socket server with plenty of NVMe slots
- Linux Asynchronous Direct I/O
- Capture NIC with bonus features:
 - Big on-card buffers
 - Bulk transfers to host memory

(Sidenote: we continue to use Metamako devices for aggregation and timestamping)

T7 Solution



T7 Solution - SmartNIC

- Napatech NT100A01
- 4 ports
- FPGA on board
- 8 GiB on-card buffer
- Freely configurable host buffers
- Packets can be transferred in bulk, including PCAP headers
- No off-the-shelf capture application included,
but useful C API is available (-> Napa2Disk)
- Good documentation



NT100 Clock Skew



NT Clock Skew

- Stratum 3 TCXO
- In a real world datacenter environment
- When synchronizing against a 1PPS signal from a GPSDO

Napa2Disk CPU Usage

- One core per capture port
- ~ 60 % utilization at saturation
- (10 Gbit, minimal sized packets)



Summary

- T7 gets new capture servers
- New capture solution:
 - captures at line rate (minimal sized packets, in wire-order)
 - saves hardware resources
 - includes some safety margins
 - e.g. for unexpected scheduling delays



Thank you!

Georg Sauthoff

Xetra/Eurex Infrastructure Design & Analytics

Deutsche Börse AG

Mergenthalerallee 61

65760 Eschborn

E-mail monitoring@deutsche-boerse.com



Deutsche Börse AG opens up international capital markets for its customers. Its product and service portfolio covers the entire process chain – from pre-IPO services and the admission of securities, through securities and derivatives trading through the settlement of transactions and the provision of market information to the development and operation of electronic trading, clearing and settlement systems. With its process-oriented business model, Deutsche Börse increases the efficiency of capital markets. Committed employees are the key factor for innovation and further growth: without them, Deutsche Börse Group would not have developed into one of the most modern exchange organisations in the world. More than 5,000 employees work for the Group – a dynamic, motivated and international team.

© Deutsche Börse Group 2022

This publication is for informational purposes only. None of the information in this publication constitutes investment advice and does not constitute an offer to sell or a solicitation of an offer to purchase any contract, share or other financial instrument. This publication is not intended for solicitation purposes but only for use as general information. All descriptions, examples and calculations contained in this publication are for illustrative purposes only.

Deutsche Börse AG, Frankfurter Wertpapierbörse (FWB®, the Frankfurt Stock Exchange), Eurex Frankfurt AG, Eurex Deutschland and Eurex Clearing AG do not represent that the information in this publication is comprehensive, complete or accurate and exclude liability for any consequence resulting from acting upon the contents of this or another webpublication, in so far as no wilful violation of obligations took place or, as the case may be, no injury to life, health or body arises or claims resulting from the Product Liability Act are affected.

Securities traded on the Frankfurt Stock Exchange and Eurex derivatives (other than EURO STOXX 50® Index Futures contracts, EURO STOXX® Select Dividend 30 Index Futures contracts, STOXX® Europe 50 Index Futures contracts, STOXX® Europe 600 Index Futures contracts, STOXX® Europe Large/Mid/Small 200 Index Futures contracts, EURO STOXX® Banks Sector Futures contracts, STOXX® Europe 600 Banks/Industrial Goods & Services/Insurance/Media/Personal & Household Goods/Travel & Leisure/Utilities Futures contracts, Dow Jones Global Titans 50 IndexSM Futures contracts, DAX® Futures contracts, MDAX® Futures contracts, TecDAX® Futures contracts, SMIM® Futures contracts, SLI Swiss Leader Index® Futures contracts, RDXxt® USD - RDX Extended Index Futures contracts, Eurex inflation/commodity/weather/property and interest rate derivatives) are currently not available for offer, sale or trading in the United States nor may they be offered, sold or traded

by persons to whom US tax laws apply.

The fund shares listed in XTF Exchange Traded Funds® are admitted for trading on the Frankfurt Stock Exchange. Users of this information service who legally reside outside Germany are herewith advised that sale of the fund shares listed in XTF Exchange Traded Funds may not be permitted in their country of residence. The user makes use of the information at their own risk.

Legal validity of this disclaimer

In the event that individual parts of or formulations contained in this text are not, or are no longer, legally valid (either in whole or in part), the content and validity of the remaining parts of the document are not affected.

Trademarks

The following names and designations are registered trademarks of Deutsche Börse AG or an affiliate of Deutsche Börse Group:

1585®; A7®; Buxl®; C7®; CDAX®; CEF®; CEF alpha®; CEF ultra®; CFF®; Classic All Share®; Clearstream®; CX®; D7®; DAX®; DAXglobal®; DAXplus®; DB1 Ventures®; DBIX Deutsche Börse India Index®, Deutsche Börse®; Deutsche Börse Capital Markets Partner®; Deutsche Börse Commodities®; Deutsche Börse Venture Network®; Deutsches Eigenkapitalforum®; DivDAX®; eb.rexx®; eb.rexX Jumbo Pfandbriefe®; ERS®; eTriParty®; Eurex®; Eurex Bonds®; Eurex Clearing Prisma®; Eurex Improve®; Eurex Repo®; Euro GC®; ExServes®; EXTf®; F7®; FDAX®; FWB®; GC Pooling®; GCPI®; GEX®; Global Emission Markets Access – GEMA®; HDAX®; iNAV®; L-DAX®; L-MDAX®; L-SDAX®; L-TecDAX®; M7®; MDAX®; N7®; ODAX®;

ÖkoDAX®;PROPRIS®; REX®; RX REIT Index®; Scale®; SCHATZ-FUTURE®; SDAX®; ShortDAX®; StatistiX®; T7®; TecDAX®; Technology All Share®; TRICE®; USD GC Pooling®; VDAX®; VDAX-NEW®; Vestima®; Xscreen®, Xemac®, Xentric®, Xetra®; Xetra-Gold®; Xpect®; Xpider®; XTF®; XTF Exchange Traded Funds®; We make markets work®

The names and trademarks listed above do not represent a complete list and, as well as all other trademarks and protected rights mentioned in this publication, are subject unreservedly to the applicable trademark law in each case and are not permitted to be used without the express permission of the registered owner. The simple fact that this publication mentions them does not imply that trademarks are not protected by the rights of third parties.

The STOXX® indices, the data included therein and the trademarks used in the index names are the intellectual property of STOXX Ltd., Zug, Switzerland and/or its licensors. Eurex' derivatives based on the STOXX indices are in no way sponsored, endorsed, sold or promoted by STOXX and its licensors and neither STOXX nor its licensors shall have any liability with respect thereto.

STOXX iSTUDIO® is a registered trademark of STOXX Ltd., Zug, Switzerland.

TRADEGATE® is a registered trademark of Tradegate AG Wertpapierhandelsbank.

EEX® is a registered trademark of European Energy Exchange AG. Flexible is better.® is a registered trademark of Axioma, Inc.