STAC Update for Fast Data

Peter Lankford
Founder and Executive Director, STAC

peter.lankford@STACresearch.com

STAC® domains in capital markets

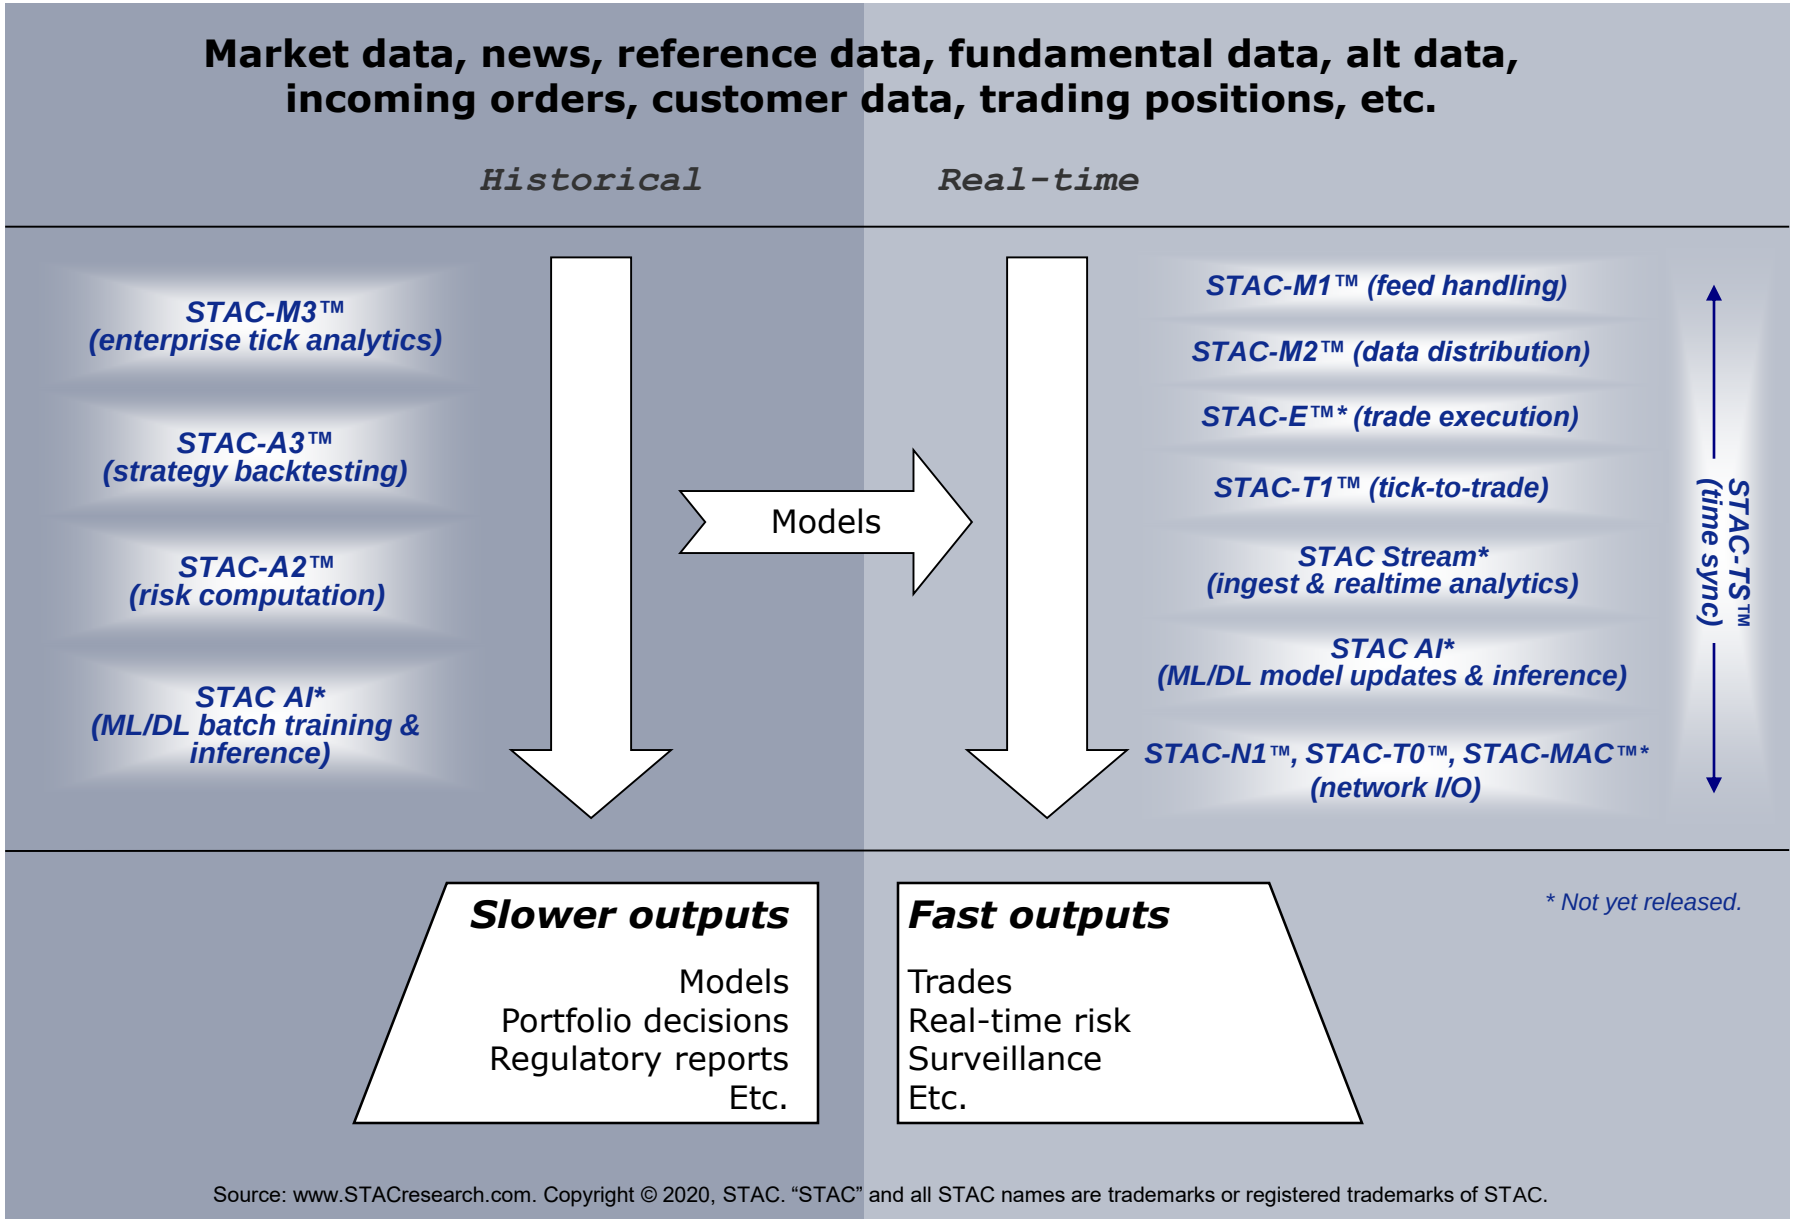

STAC® domains in capital markets

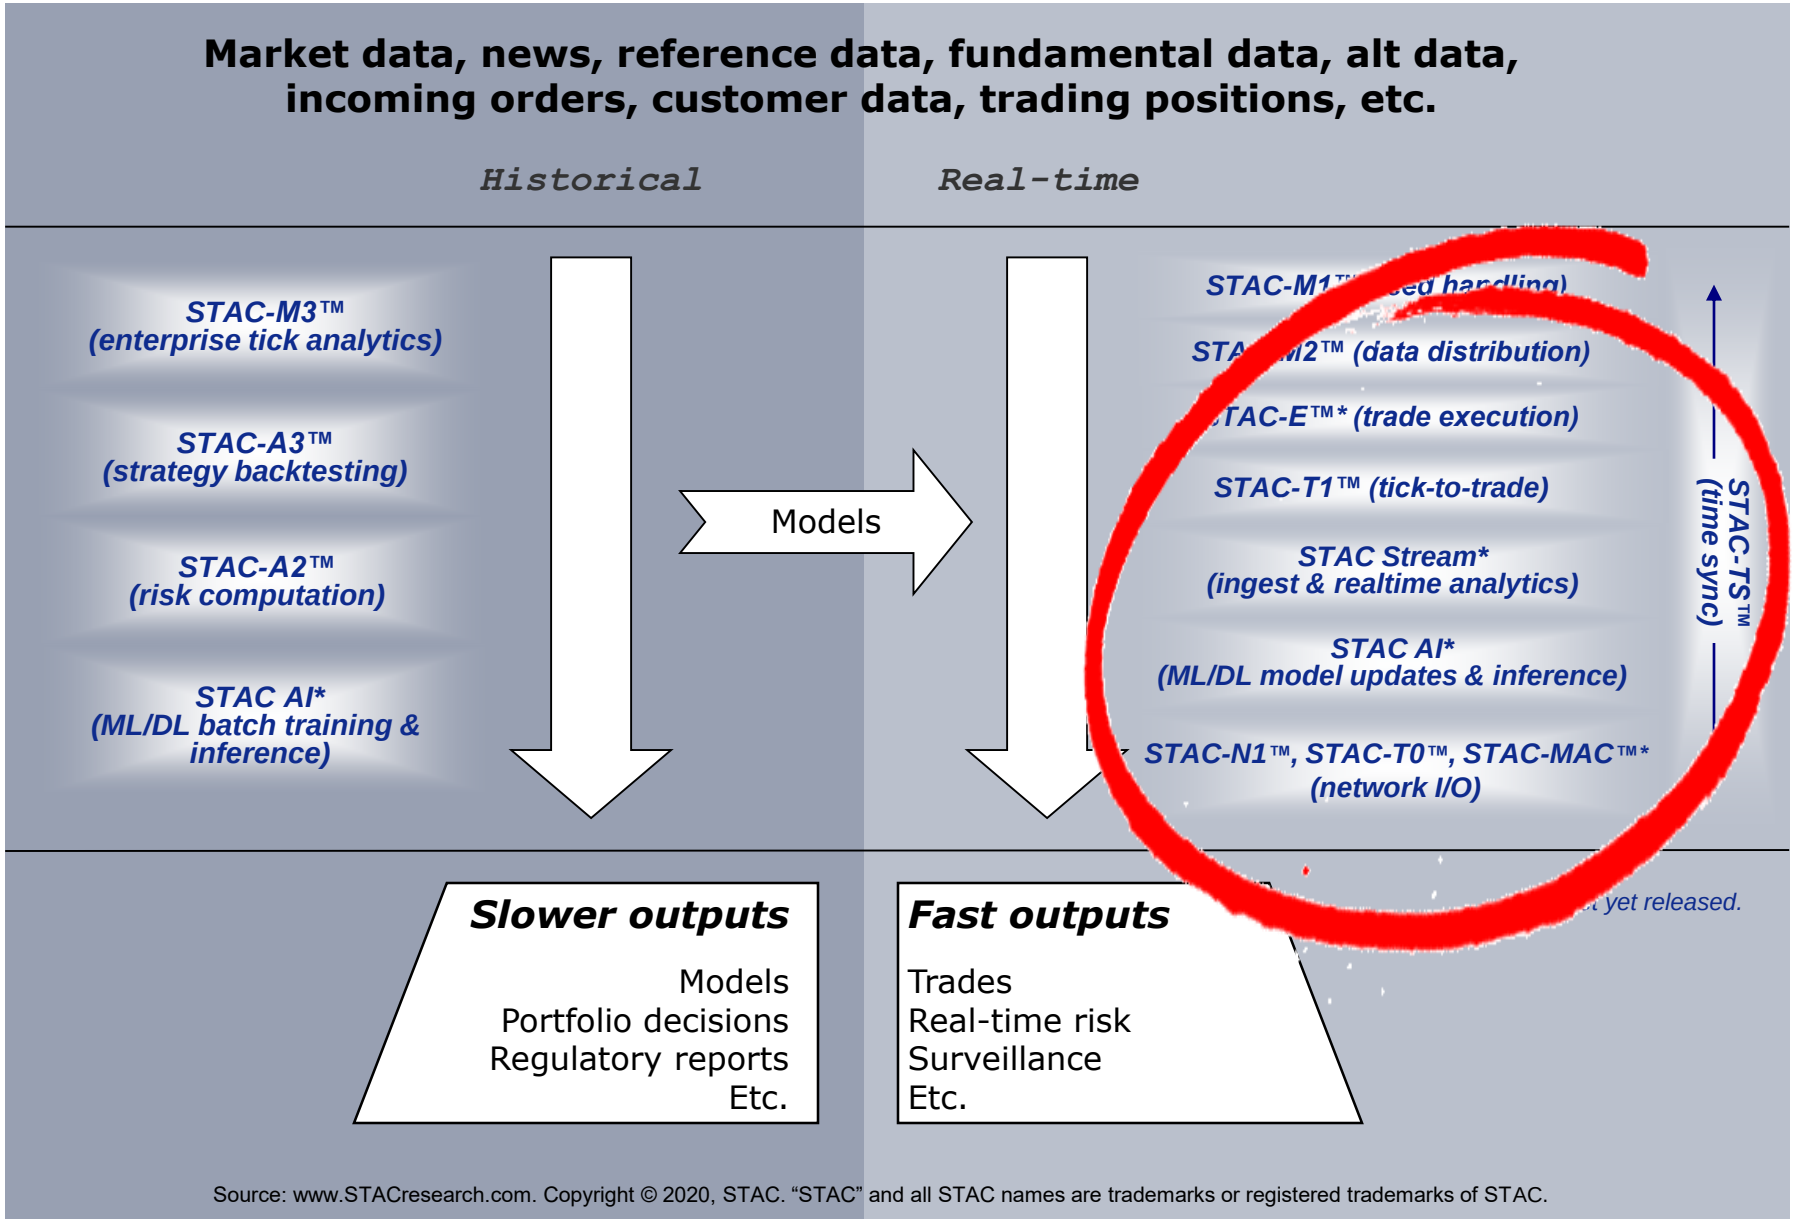

STAC-T0

- Tick-to-trade pattern
  - UDP in (507B or 68B frames x 3 rates)
  - TCP out (122B frames)
- Isolates network I/O latency
  - Latency that cannot be squeezed out of business logic
  - Does not co-mingle I/O latency with market-specific logic
- Extremely high accuracy
- Works with any trading platform (sw/hw)
- Key metric: Actionable Latency
  - STAC-T0.ACTIONABLE.LAT: $t_{FO} - t_{IND}$

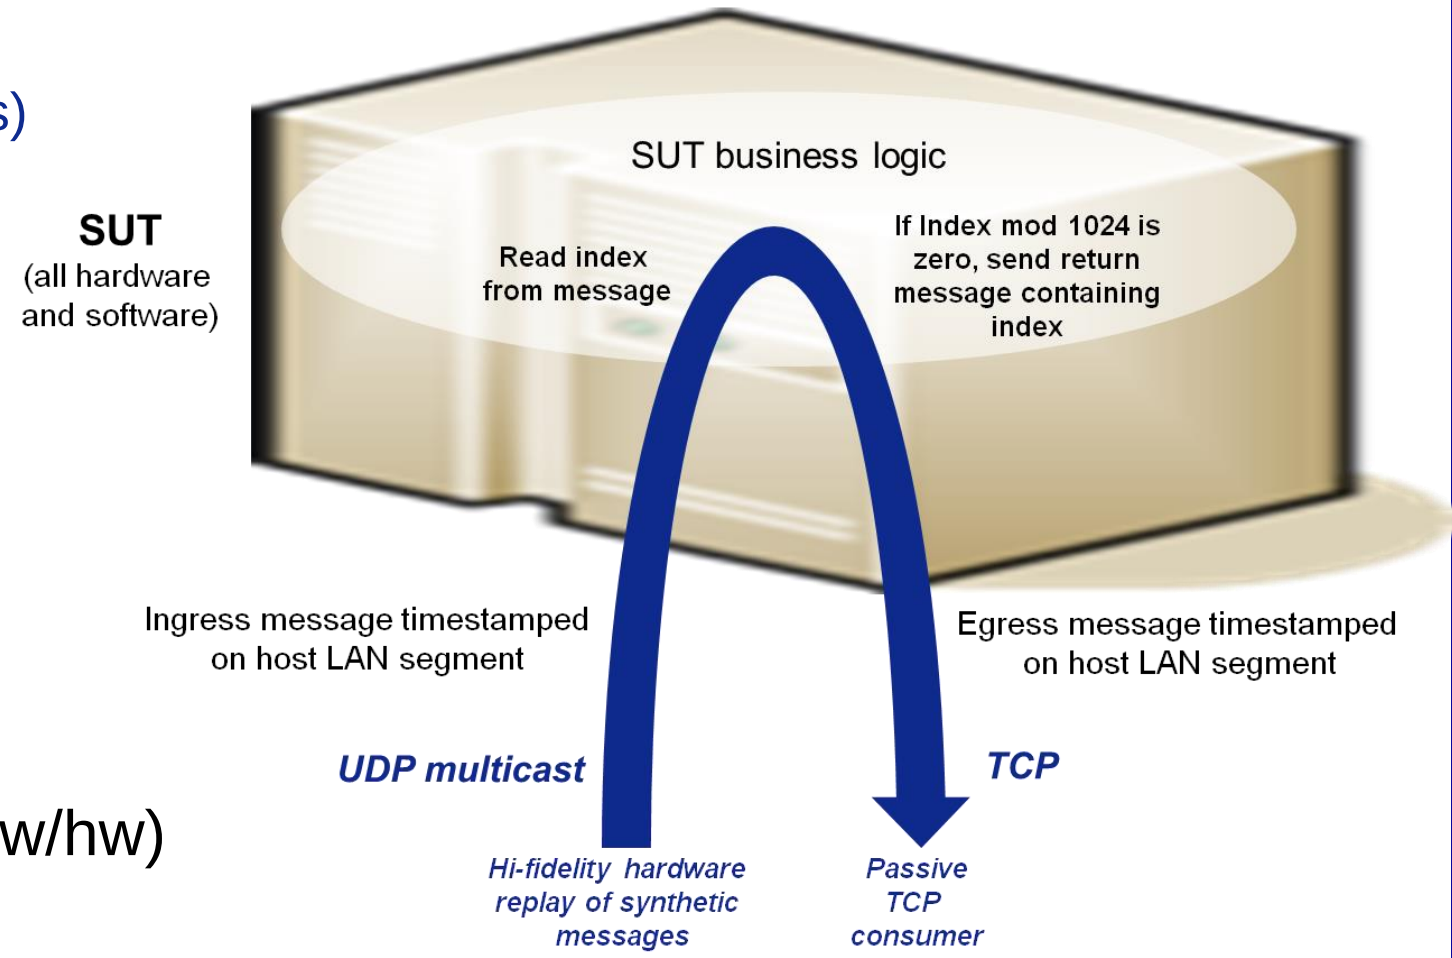

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Latest SUT: New solution from LDA and Xilinx

- LDA 644 MHz Framework, including:
  - 644 MHz/16-bit MAC/PCS
  - Lightspeed v2 TCP IP Core
- LDA SBM09P-3 FPGA board
- Xilinx® Virtex™ UltraScale+ VU9P-3 FPGA
- Xilinx Onload®
- Xilinx SFN8522 PLUS Ethernet adapter
- Lenovo server ThinkSystem SR665 with 2 x AMD EPYC™ 7742 processors
- Finisar FTL410QE1c QSFPs

www.STACresearch.com/XLX200514

Test setup

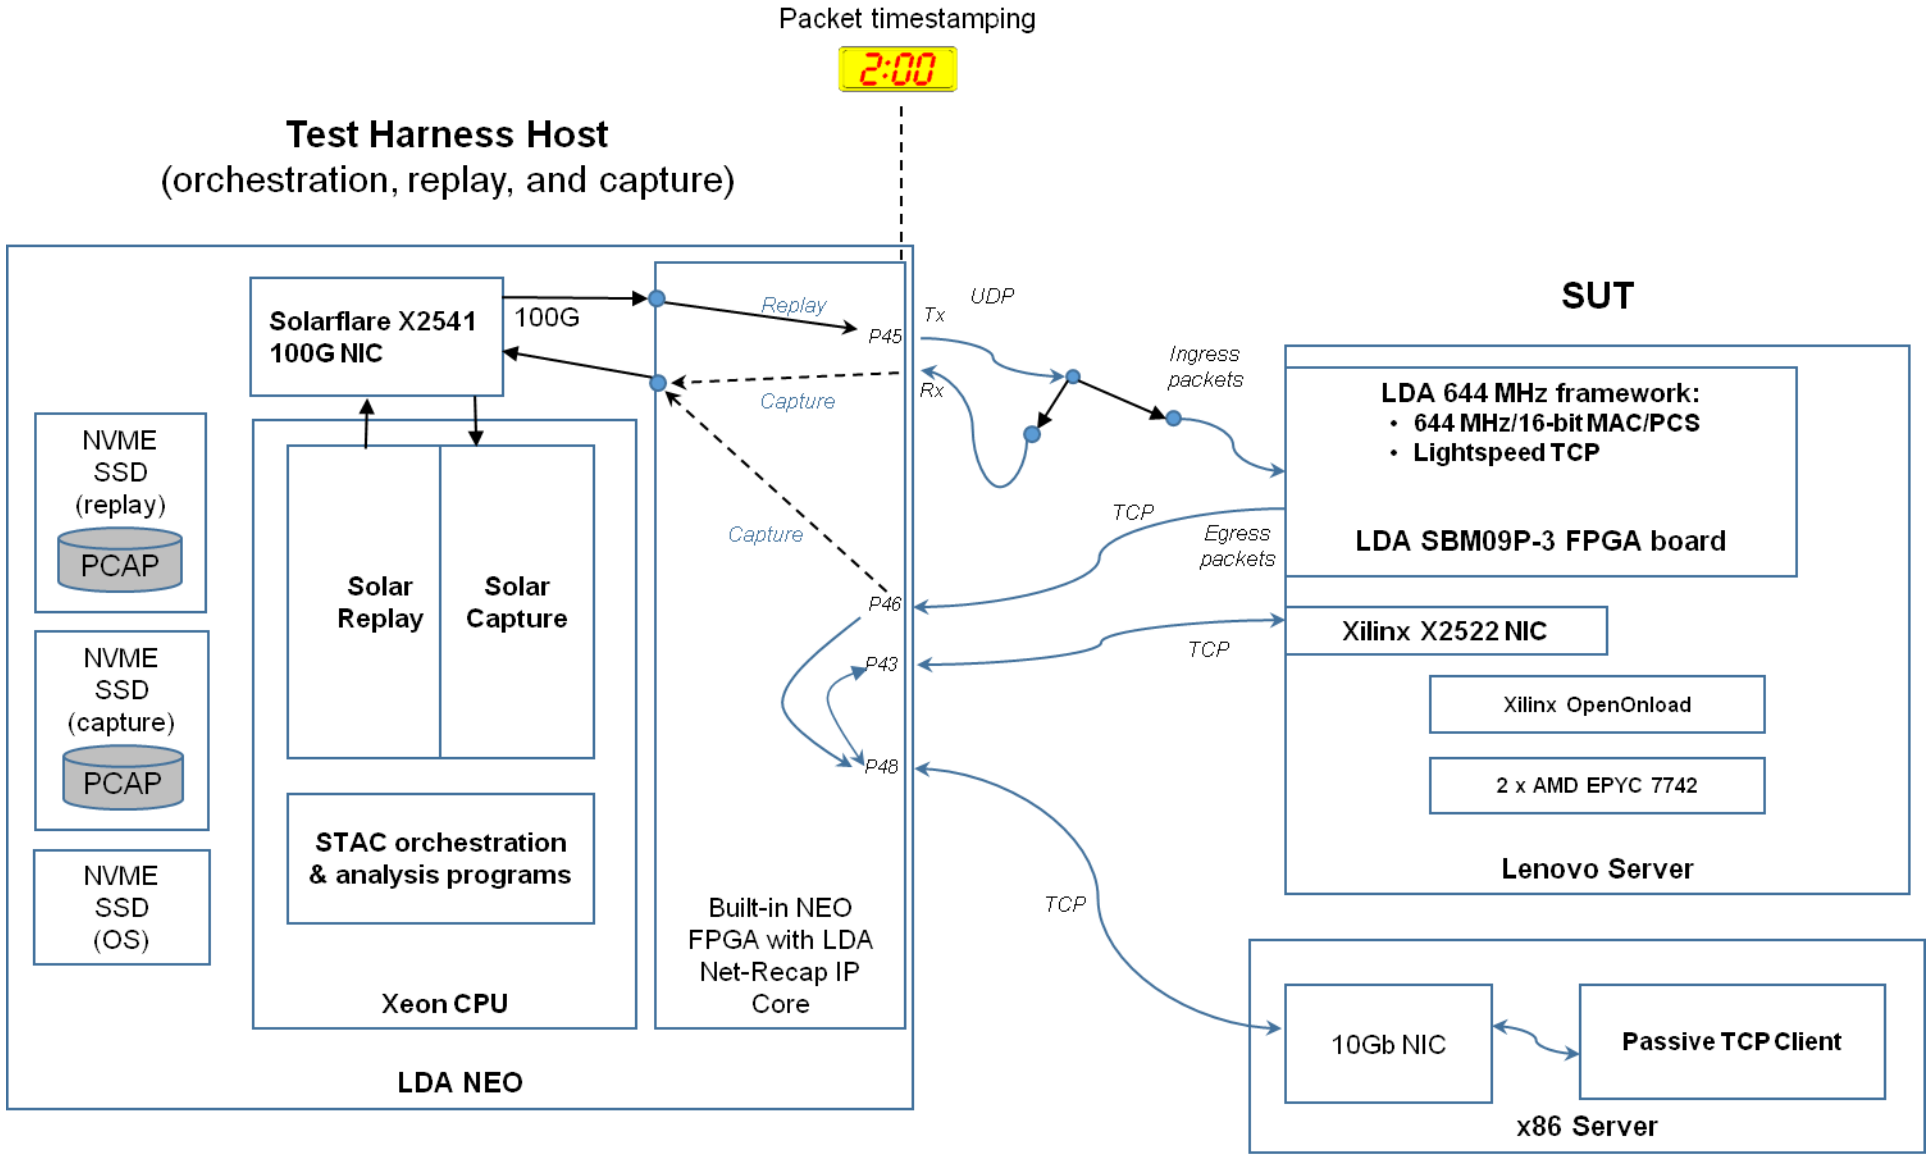

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Results highlights

- New records for both message sizes

Xilinx/LDA highlighted:

- For 68-byte frames at all ingress rates, actionable latency had a minimum of 24.2 nanoseconds (STAC-T0.β1.*.B.ACTIONABLE.MIN)
- For 68-byte frames, actionable latency had a mean of 27.9 to 29.2 nanoseconds, depending on ingress rate (STAC-T0.β1.*.B.ACTIONABLE.MEAN)

Arista under STAC-TS network timestamping & capture

- Arista DCS-7130-48L switch
- Arista MetaWatch 1.0.0 firmware
- Arista MOS 0.26.3 operating system
- Arista SFP-10G-SR fibre SPFs
- TimeTech Pulse Distribution Unit 10535
- Jaycar WC7802 RG316 Coax PPS cable - 3m

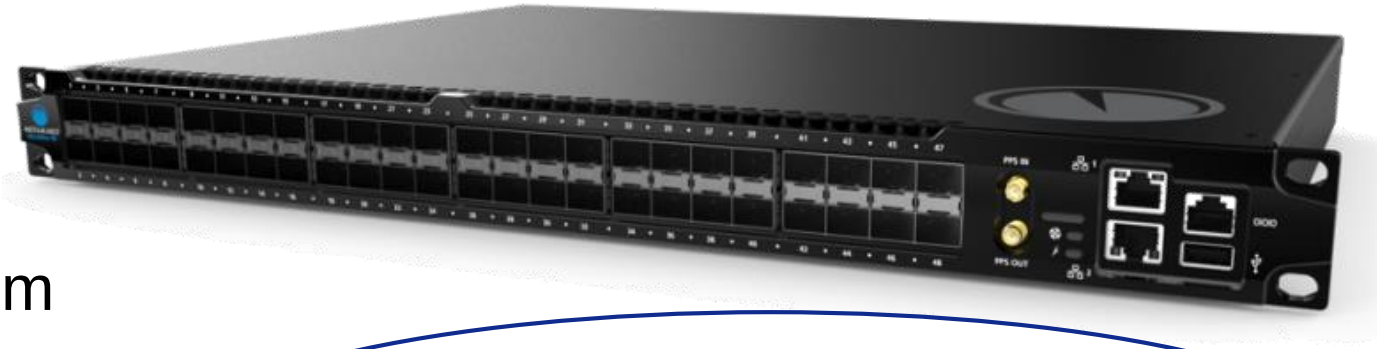

www.STACresearch.com/ARI200518

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-TS.PSE1 – Port sync error within a single device

- 47 ingress ports feeding 1 egress port
- Worst case between any two ports: 102 +/- 274 picoseconds

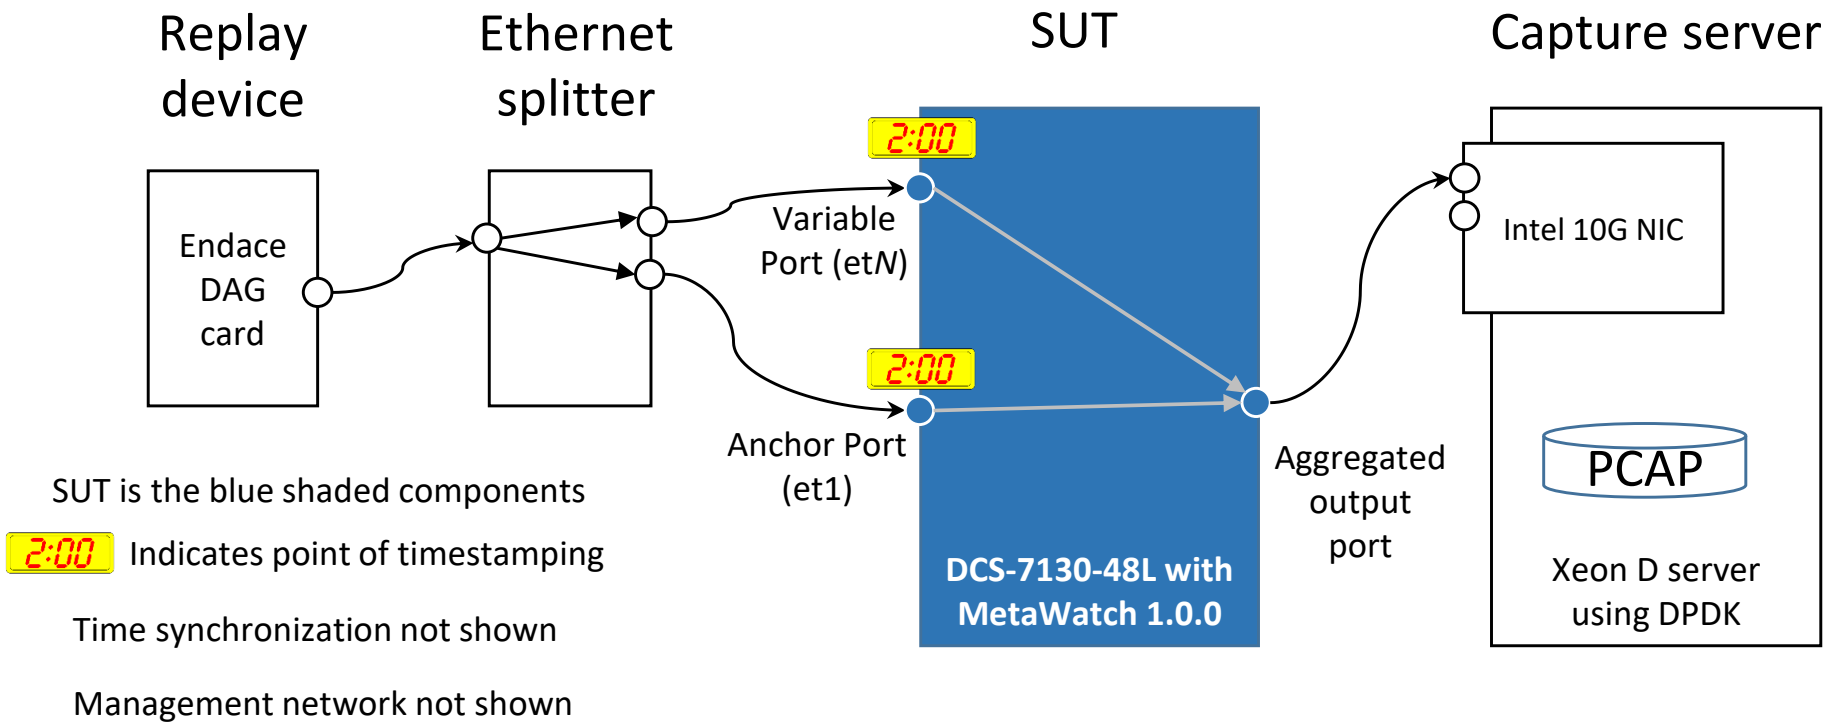

STAC® SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-TS.PSE2 – Port sync error across two devices

- 47 timestamping ports on each device
- Worst case sync among them: -118 +/- 681 picoseconds

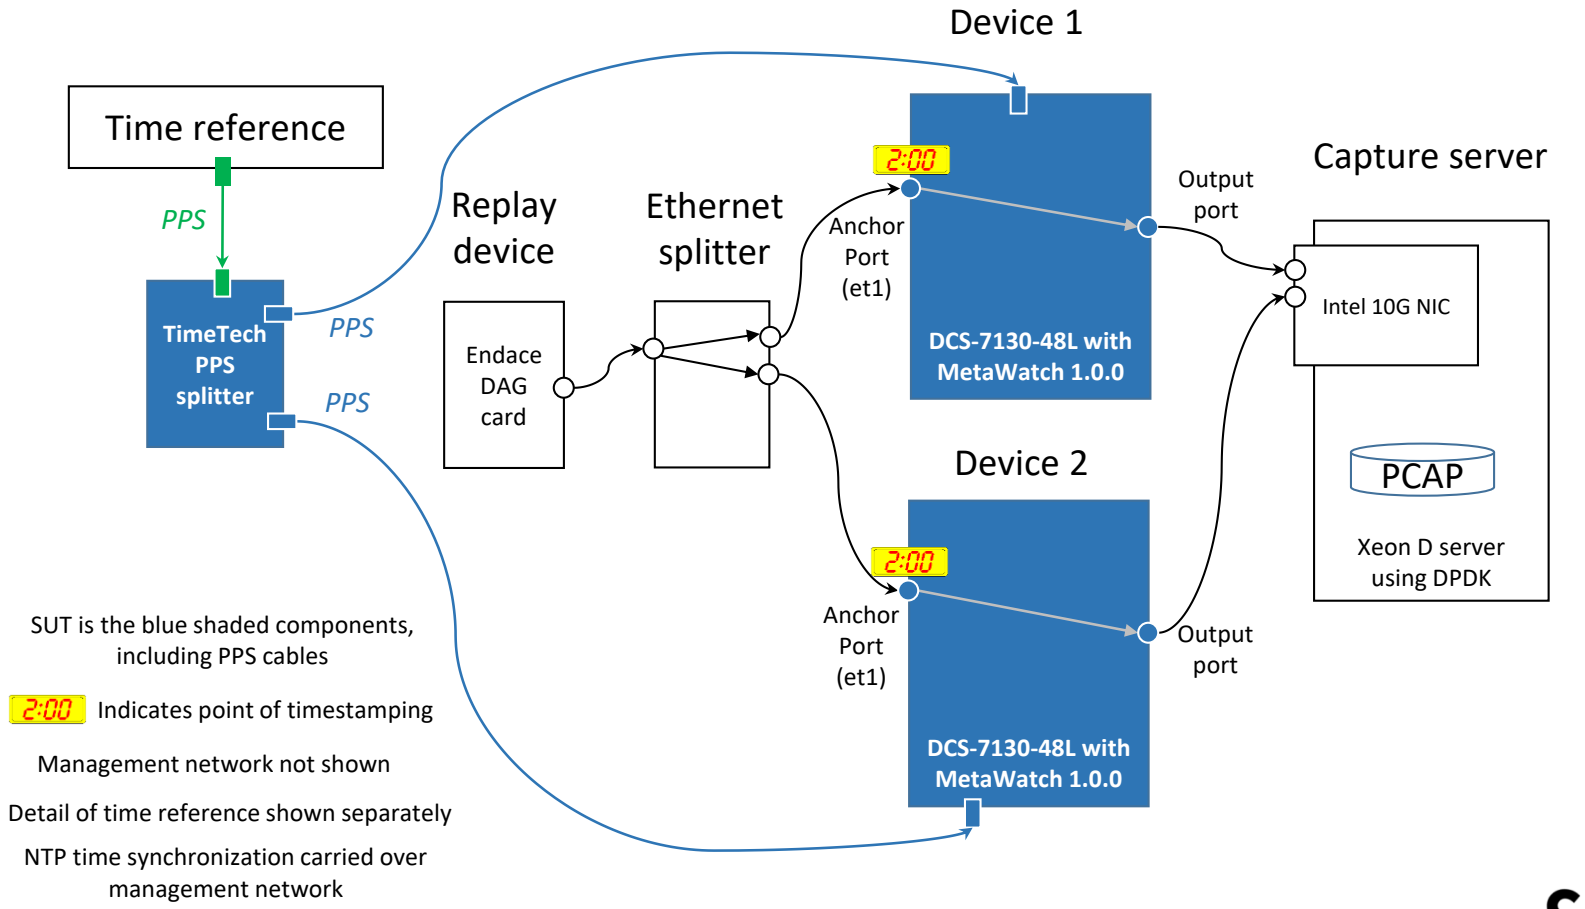

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-TS.CAP2 – Capture capacity by packet size and duration

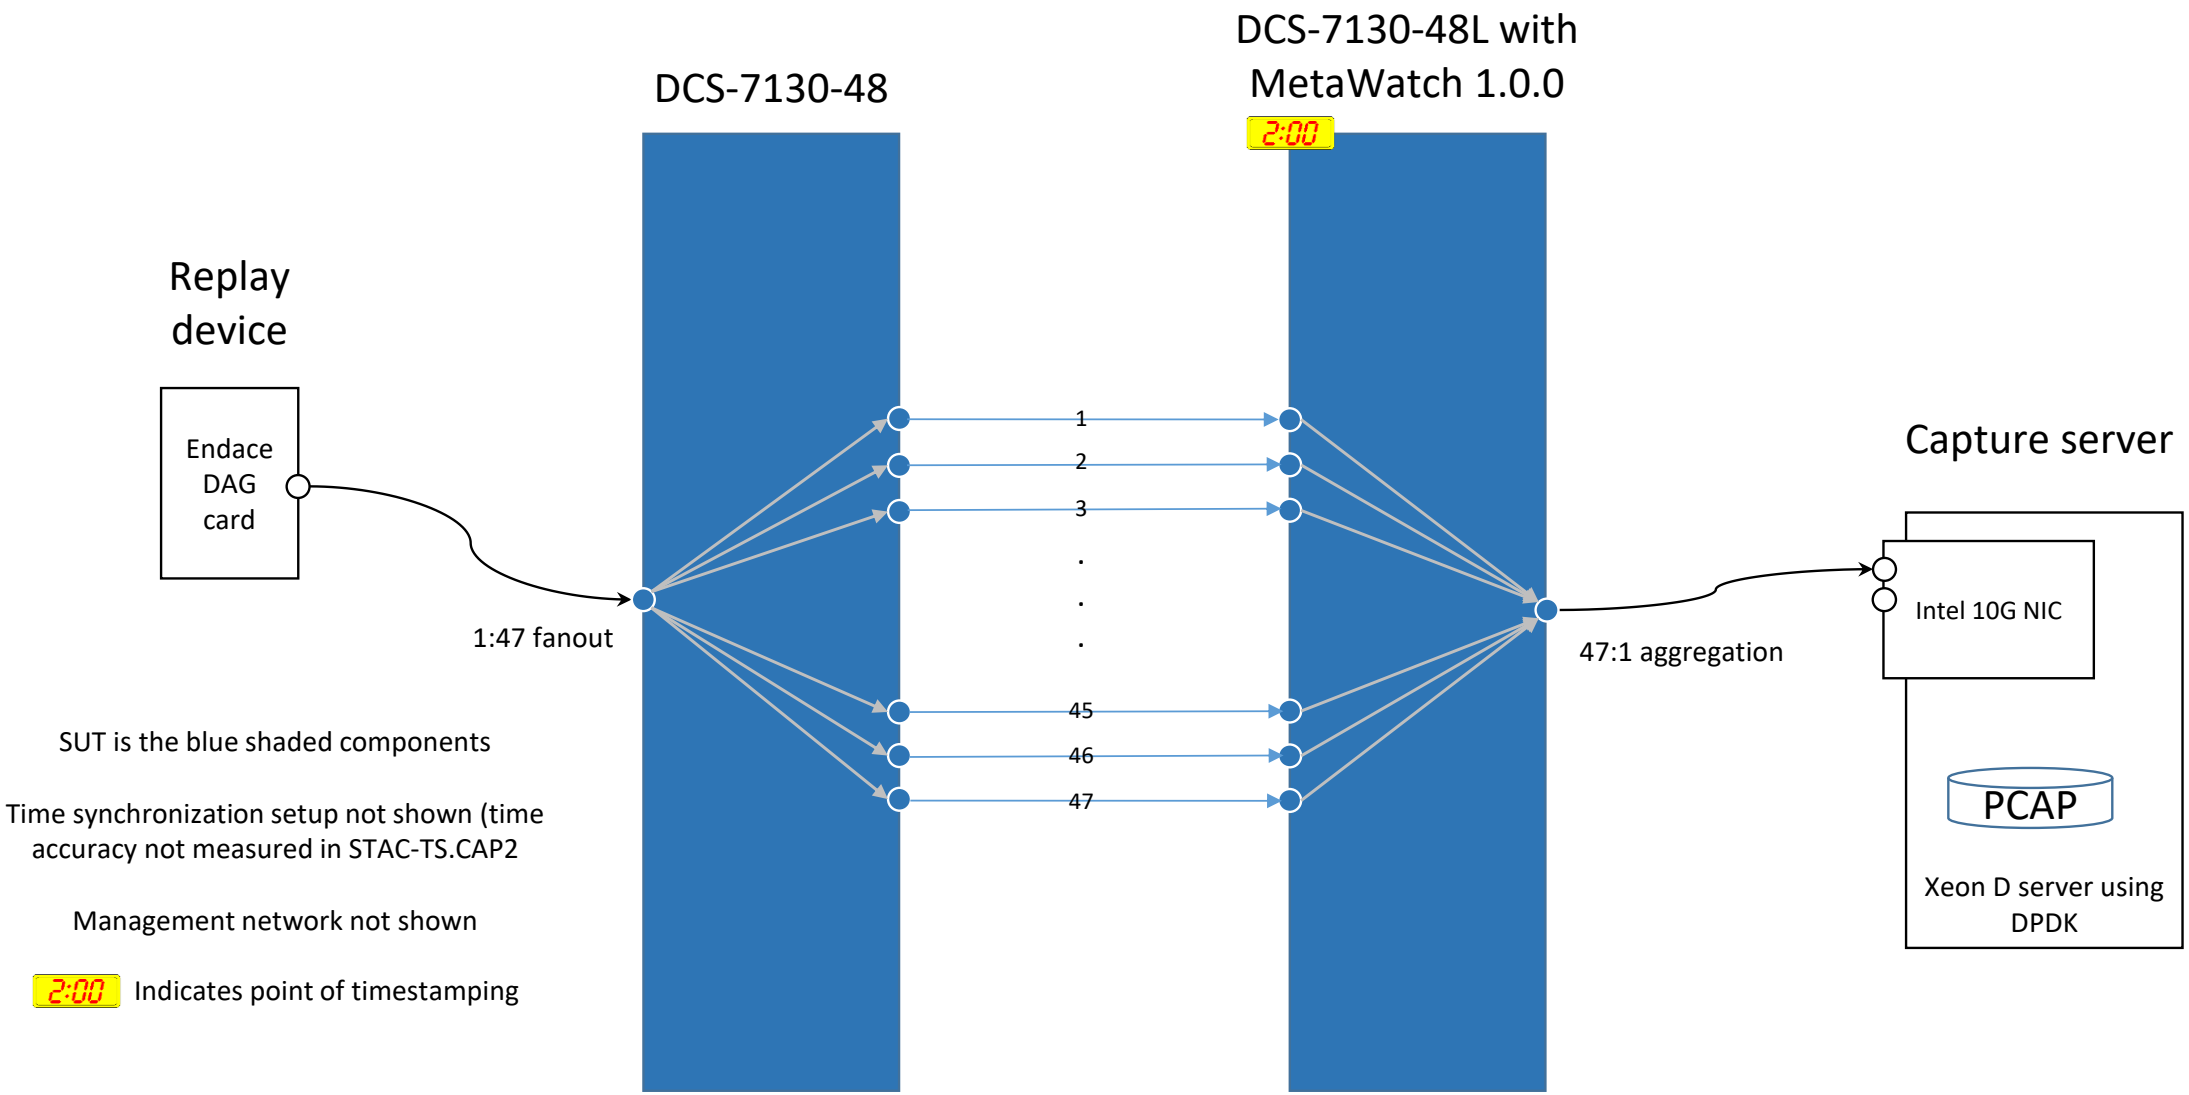

STAC® SECURITIES TECHNOLOGY ANALYSIS CENTER

Results

Max aggregate ingress bandwidth under increasing packet sizes and durations

	Duration of packet flow							
Bytes/ packet	100 microseconds	1 millisecond	10 milliseconds	100 milliseconds	500 milliseconds	1 second	100 seconds	
67	470.00	470.00	470.00	470.00	470.00	282.00	9.40	
131	470.00	470.00	470.00	470.00	470.00	277.30	9.40	←
257	470.00	470.00	470.00	470.00	470.00	274.95	11.75	
521	470.00	470.00	470.00	470.00	470.00	274.95	11.75	← Gbps
1031	470.00	470.00	470.00	470.00	470.00	274.95	11.75	
1283	470.00	470.00	470.00	470.00	470.00	274.95	11.75	←
1511	470.00	470.00	470.00	470.00	470.00	274.95	11.75	

STAC-TS.NTE1 – Absolute accuracy (network timestamp error)

- Test in steady state and holdover
- Three stages:
  - Holdover
  - Recovery
  - Steady state
- Holdover is 24 hours
- Vendor specifies the stabilization period for recovery (we report it)

Holdover results

- 24-hr drift: 1.4 µsec (STAC-TS.NTE1.HLDVR)

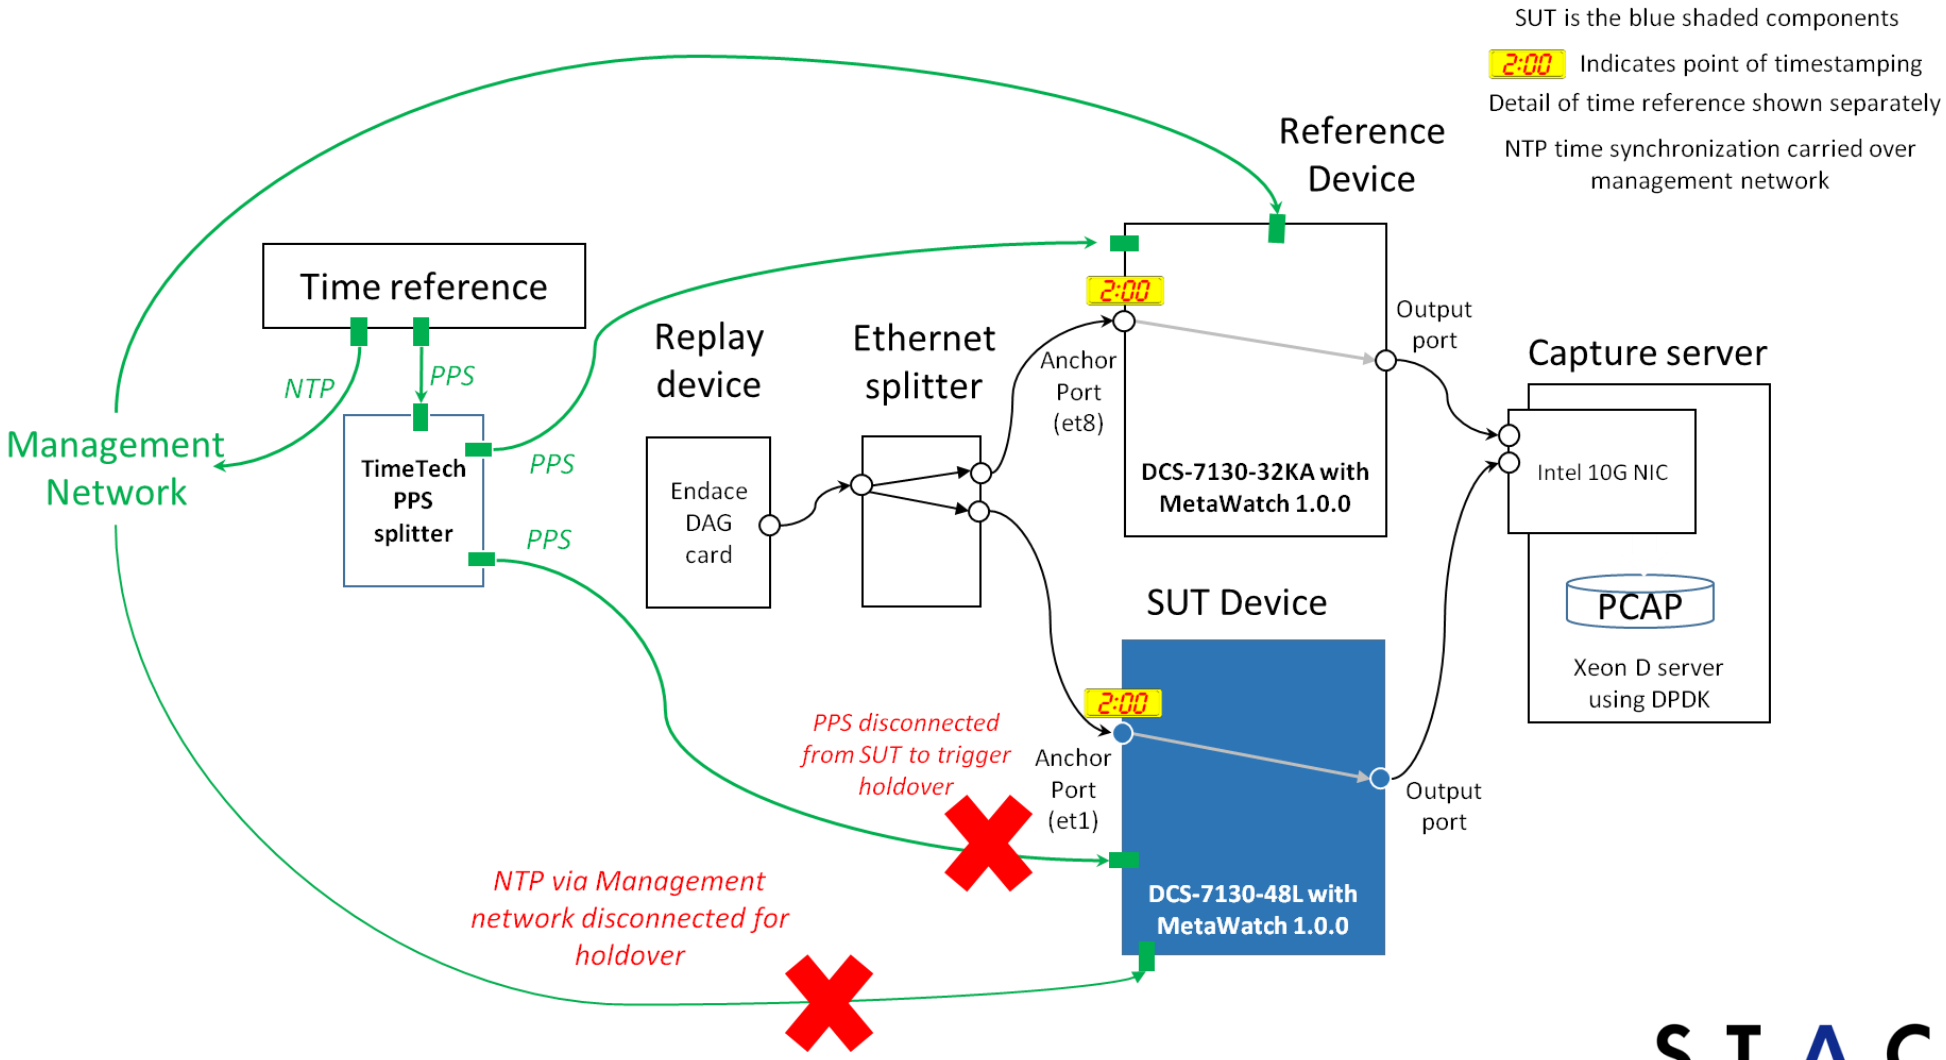

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Steady state results

- Steady state: 0 +/- 5 ns (STAC-TS.NTE1)
  - Worst-case absolute accuracy across all ports
  - Combines absolute accuracy of anchor port (measured against Reference Device) with port-sync errors
  - Reference Device is the switch config that we tested against a high-frequency scope a few years ago (STAC-TS.NTE2)
  - Interesting case, because Arista claims the Reference Device is probably less accurate than the SUT (one would expect, given that the SUT is much more recent)
  - That's consistent with the test results but not proven by them
  - STAC has to take the conservative stance
  - See the report for discussion

“STAC MAC”

- Reminder: Benchmarks for the latency of vendor-provided MAC and PHY for a given FPGA chip
- Sign up to be involved: www.STACresearch.com/nio
- Spec status: Agreed
  - Some interesting discussions. See STACesearch.com/resources/Network-IO-sig/meeting-note
- Tool status:
  - Leverage much of STAC-T0 Test Harness
  - Need to develop a few more things
- Expect ready in July/Aug

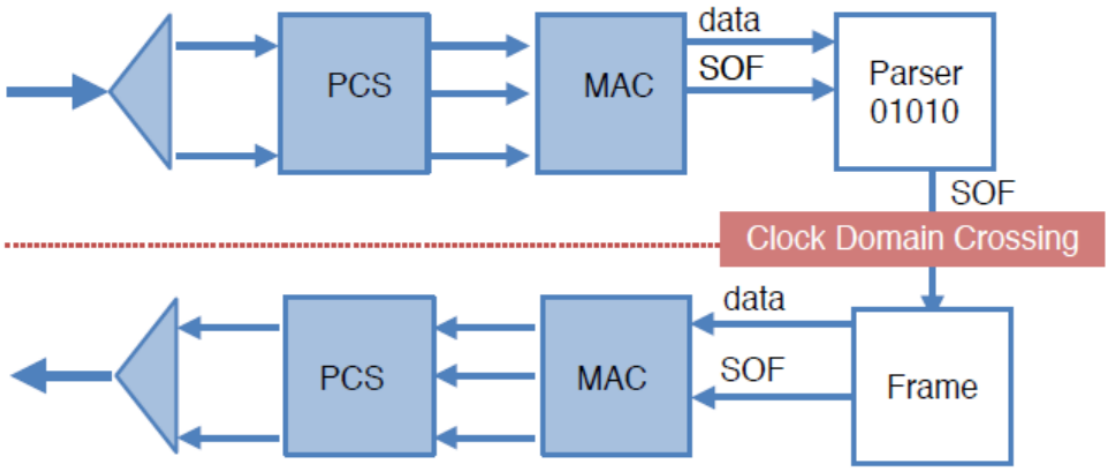

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

Stay tuned for the Innovation Roundup

- A few 5-minute presentations
- Vendors chose their own applause (!)
- Don’t forget to use your response card
- Update and “Send” it as often as you want

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER