Picosecond scale measurements for nanosecond scale trading

Dr Matthew Grosvenor

STAC Summit NYC, Spring 2017

Disclaimer:
Not STAC measurements

Trading in nanoseconds

120ns tick to trade

http://solarflare.com/solarflare-and-lda-technologies-harness-the-power-of-xilinx-fpgas

95ns switching

https://exablaze.com/exalink-fusion

49ns multiplexing

https://exablaze.com/exalink-fusion-mux

<5ns multicast

https://exablaze.com/exalink-fusion

https://www.metamako.com/products/metamux-48.html

How do you verify nanosecond scale operations?

Measurement resolution (precision)

	Precision
Device A	7.5ns
Device B	4.0ns
Device C	7.5ns
ExaNIC X10	6.2ns

Trading 4-50ns

vs

Measurement 4-7.5ns

Trading 4-50ns

vs

Measurement 4-7.5ns

50 ns ± 7.5ns

Trading 4-50ns

vs

Measurement 4-7.5ns

50 ns ± 7.5ns 30% ‘error margin’

Trading 4-50ns

vs

Measurement 4-7.5ns

50 ns ± 7.5ns 30% ‘error margin’

5 ns ± 4ns 80% ‘error margin’

42ns

42ns ± 8ns

42ns ± 8ns

~40% error margin

Measurement repeatability (accuracy)

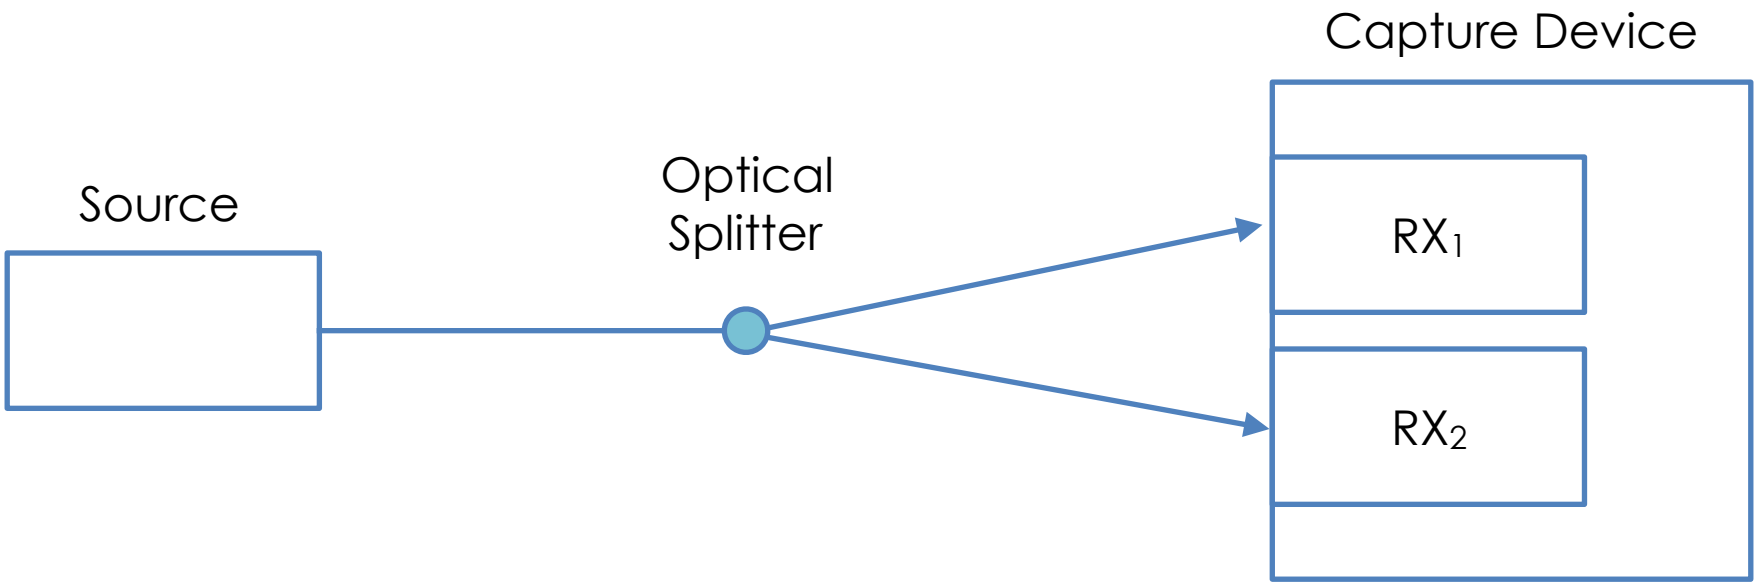
Capture Device
Source
Optical
Splitter
RX1
RX2

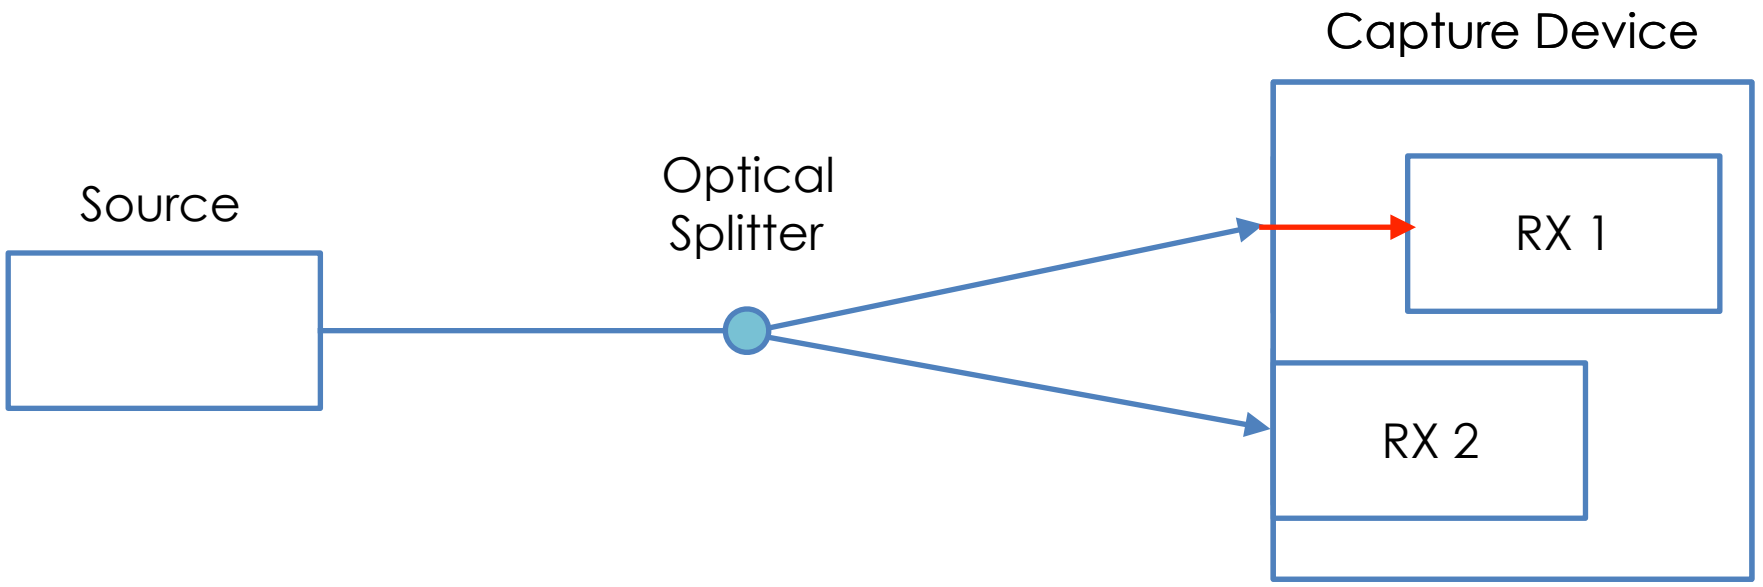

$$\text{error} = RX_2 - RX_1$$

Device A

resolution 7.5ns

10,000,000 samples

Error Distribution

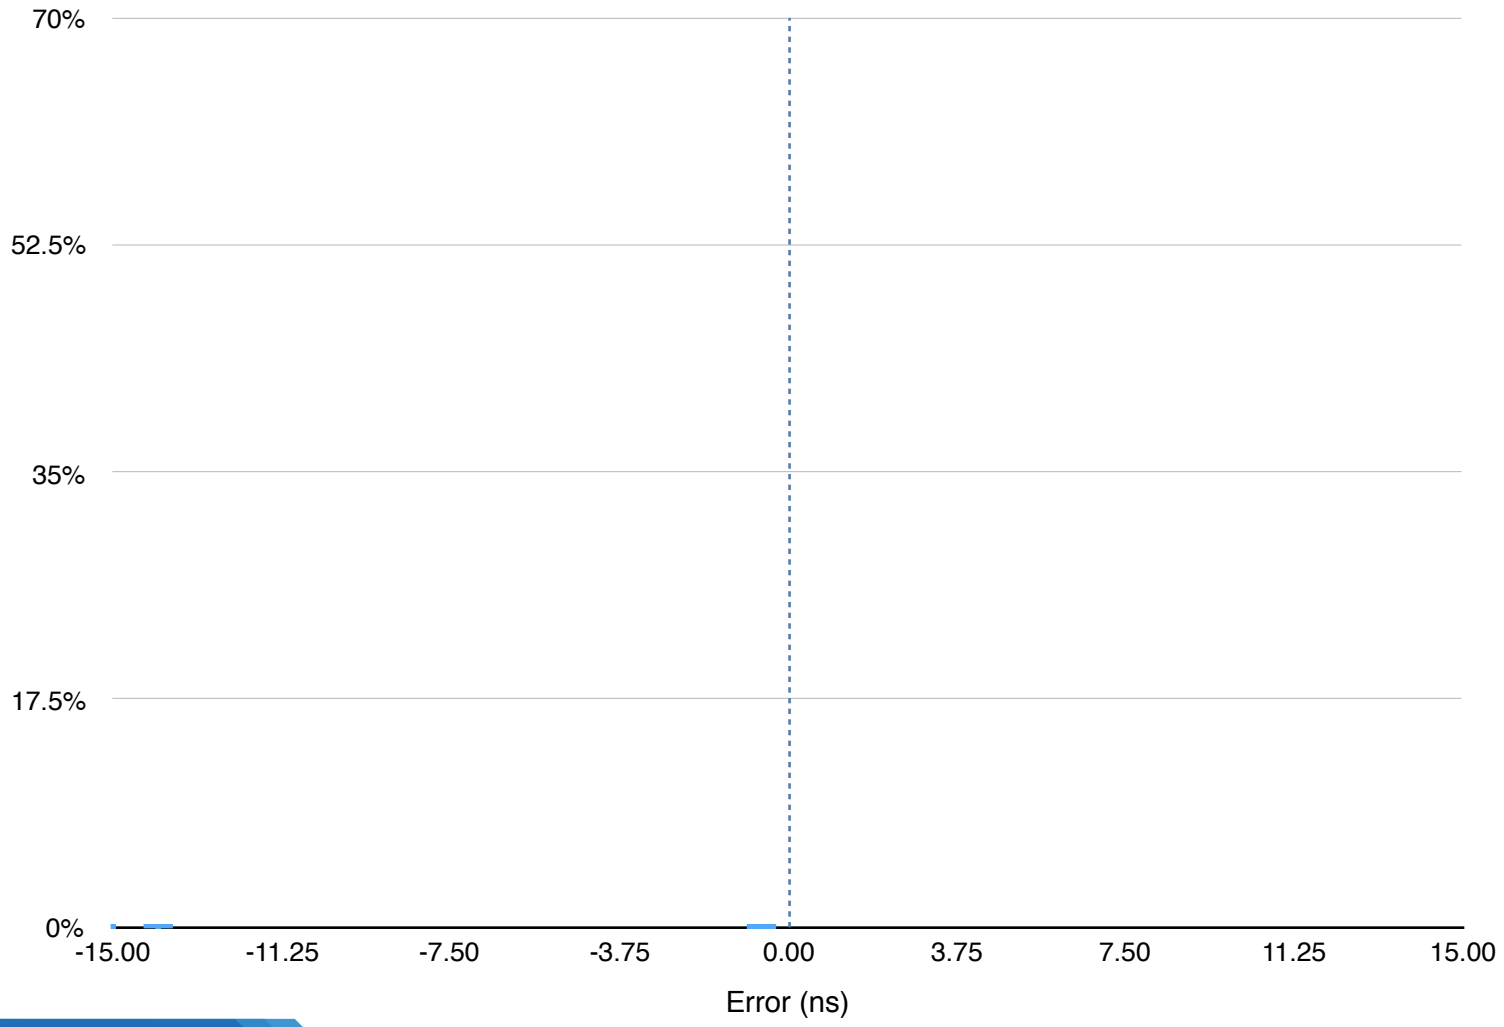

Error Distribution

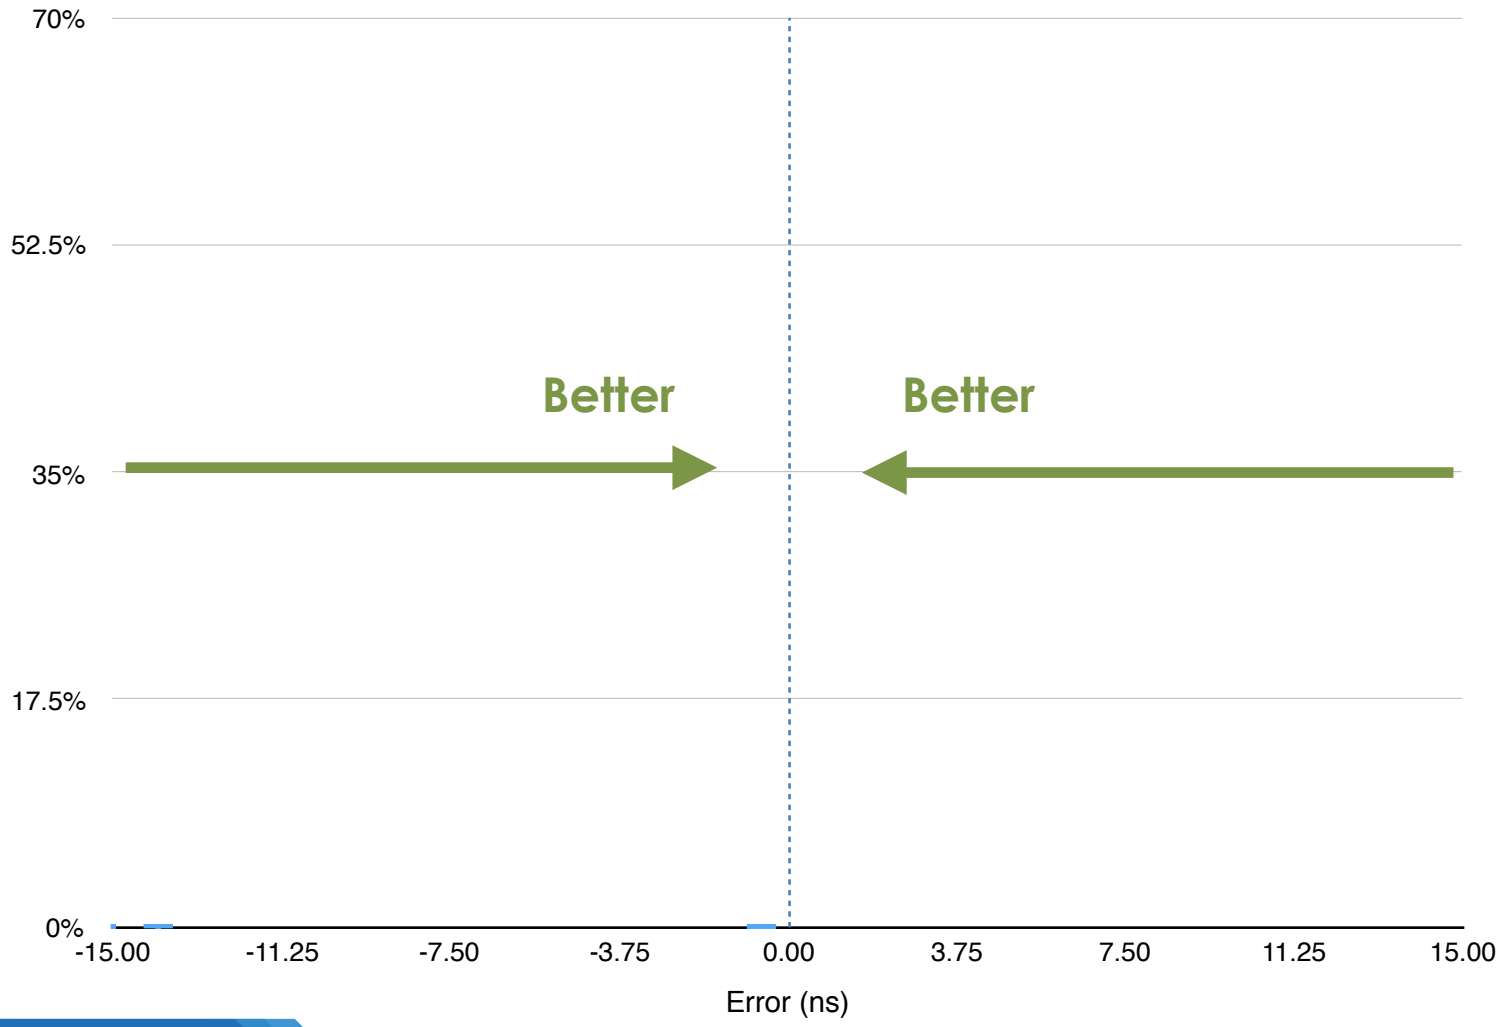

Error Distribution

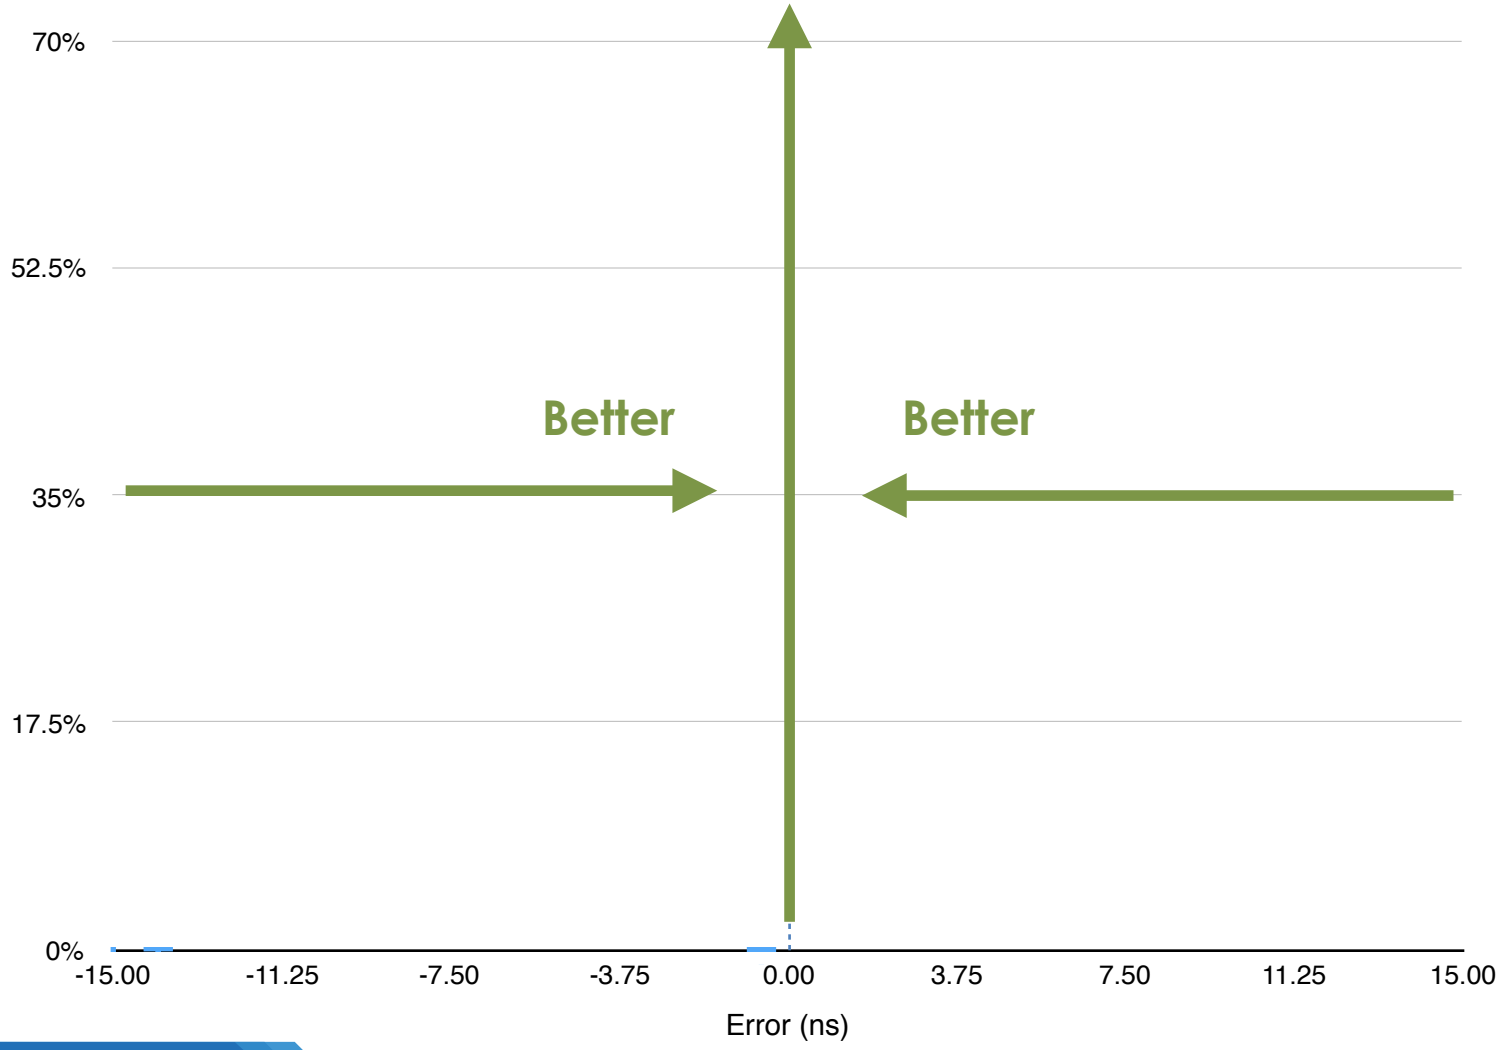

Device A

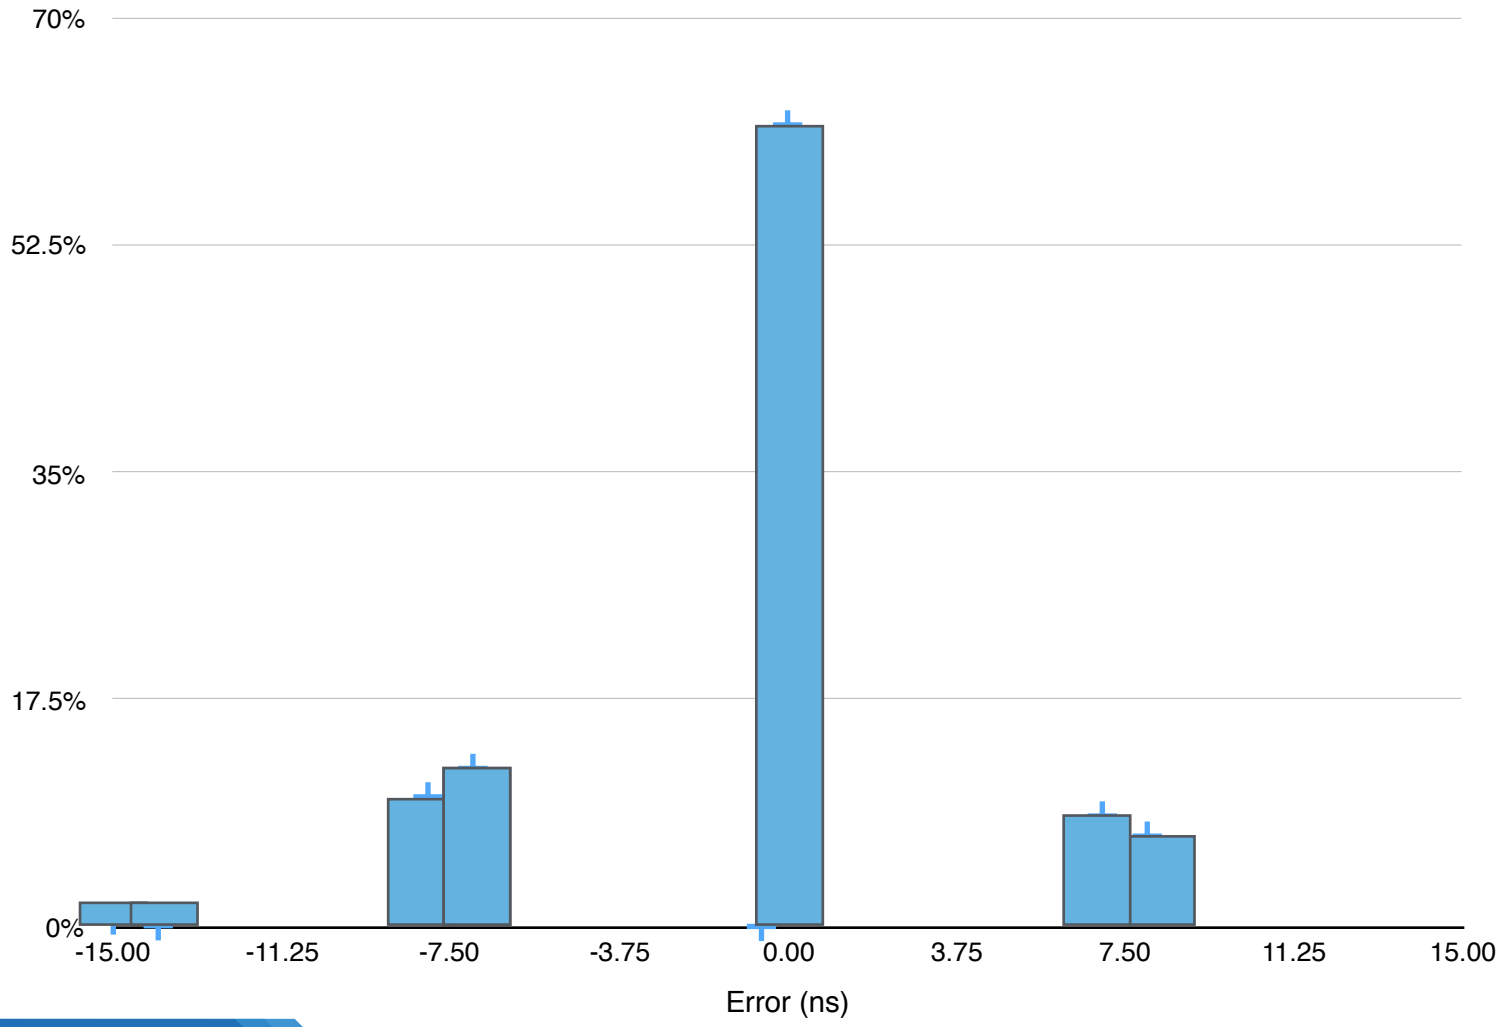

Device A

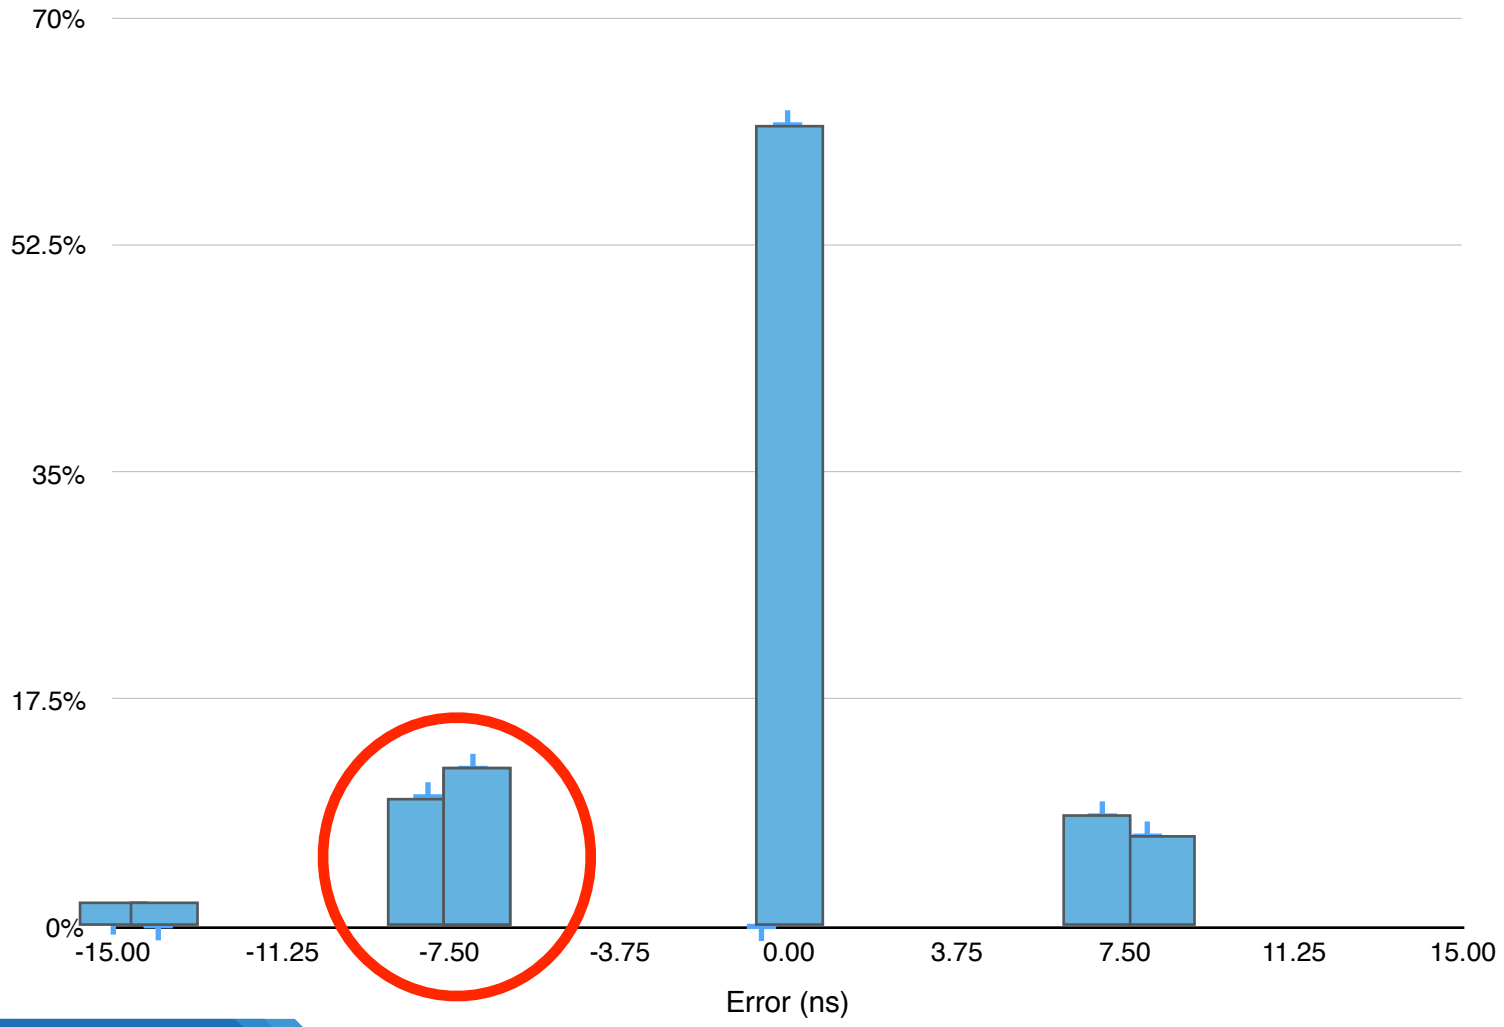

Gotcha #1

Measurement Precision ≠ Reporting Precision

7.5ns not evenly divisible by 1ns

Device A (cleaned)

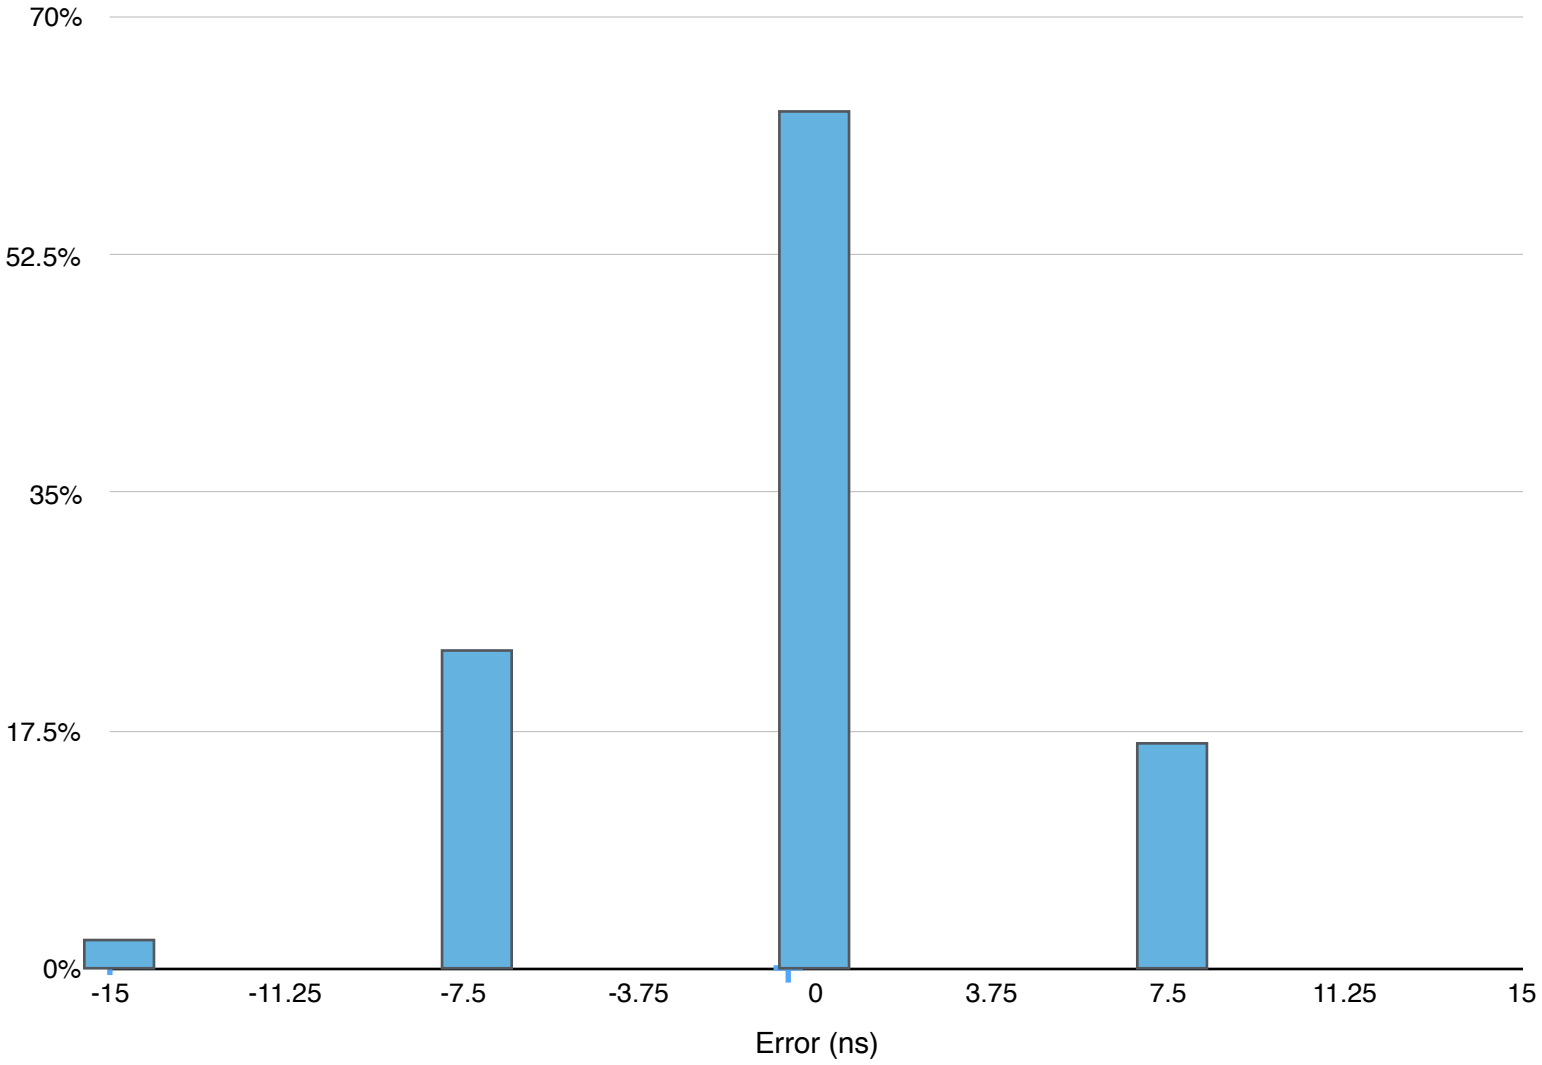

Device B

resolution 4ns

10,000,000 samples

Device B

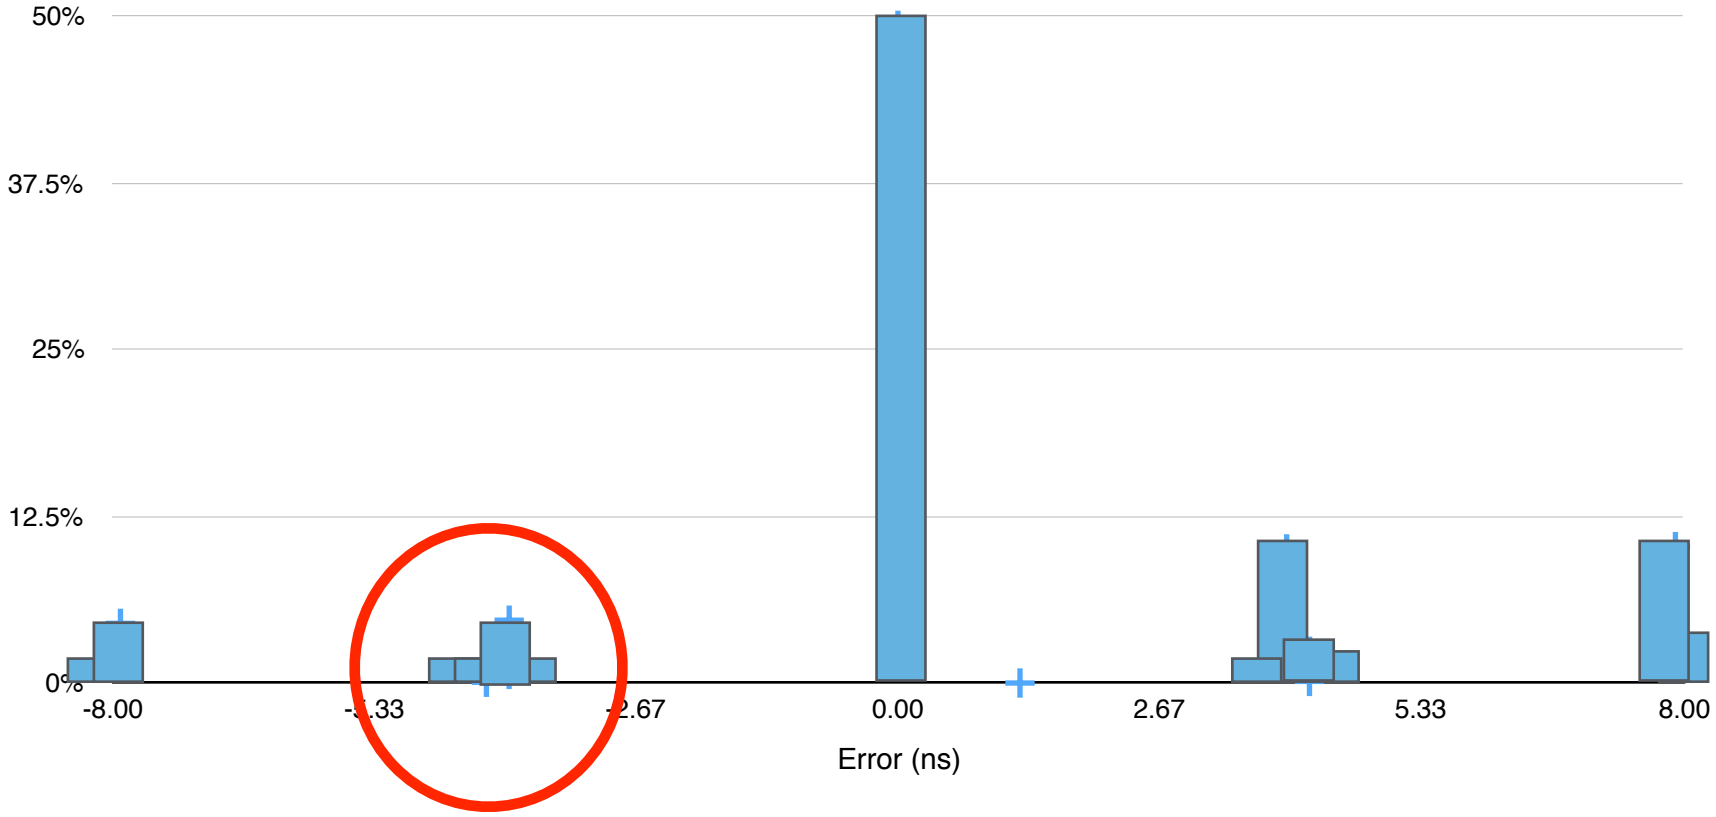

Gotcha #1 (again)

Measurement Precision ≠ Reporting Precision

4ns not representable as binary fraction

Tip #1

Round to nearest multiple of measurement clock period

Device B (cleaned)

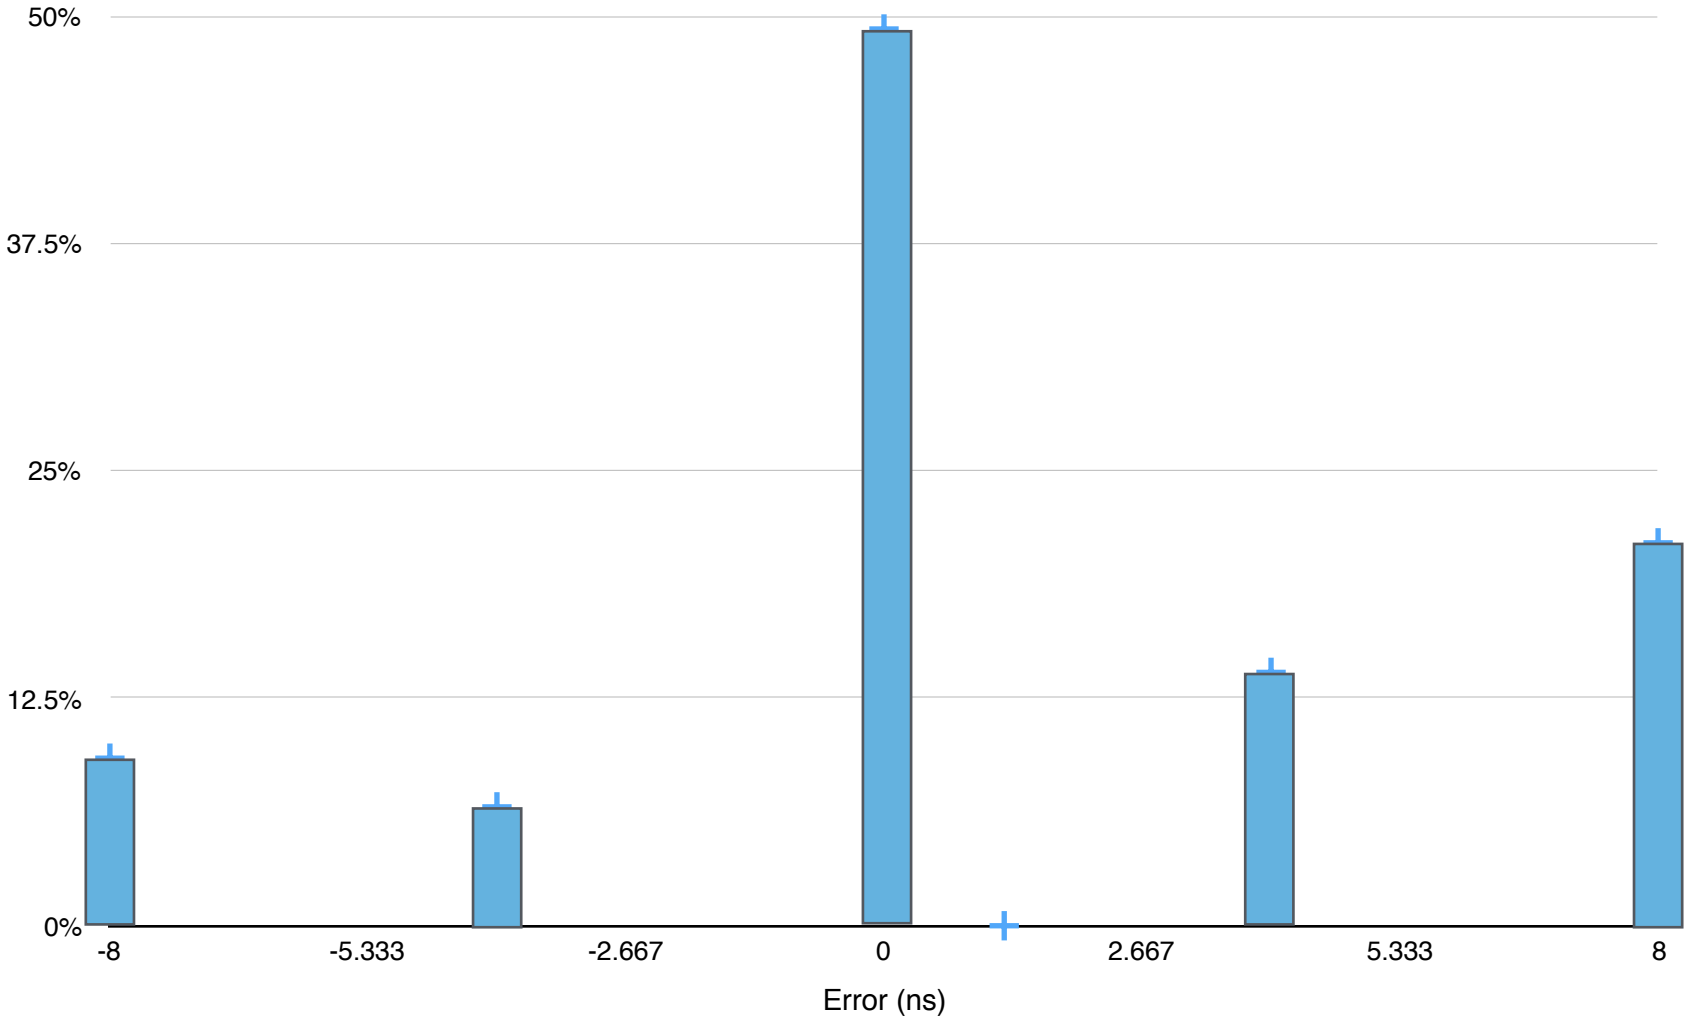

Device B (cleaned)

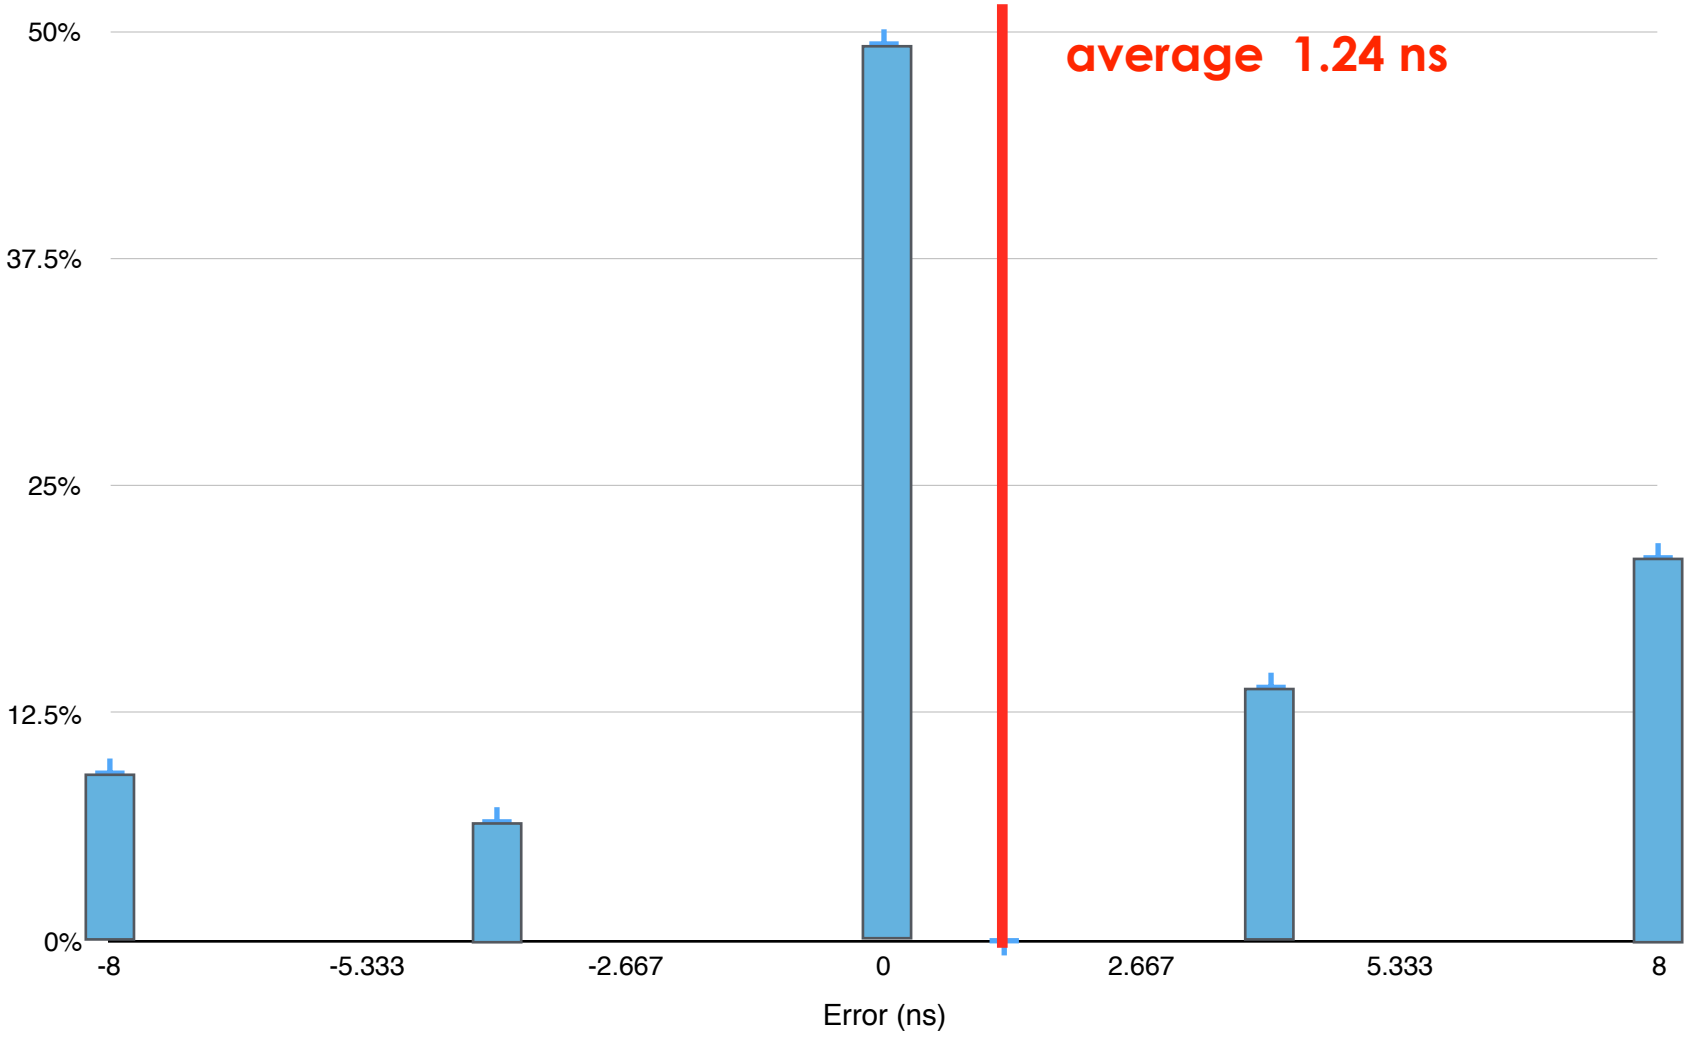

Device C

resolution 7.5ns

10,000,000 samples

Device C

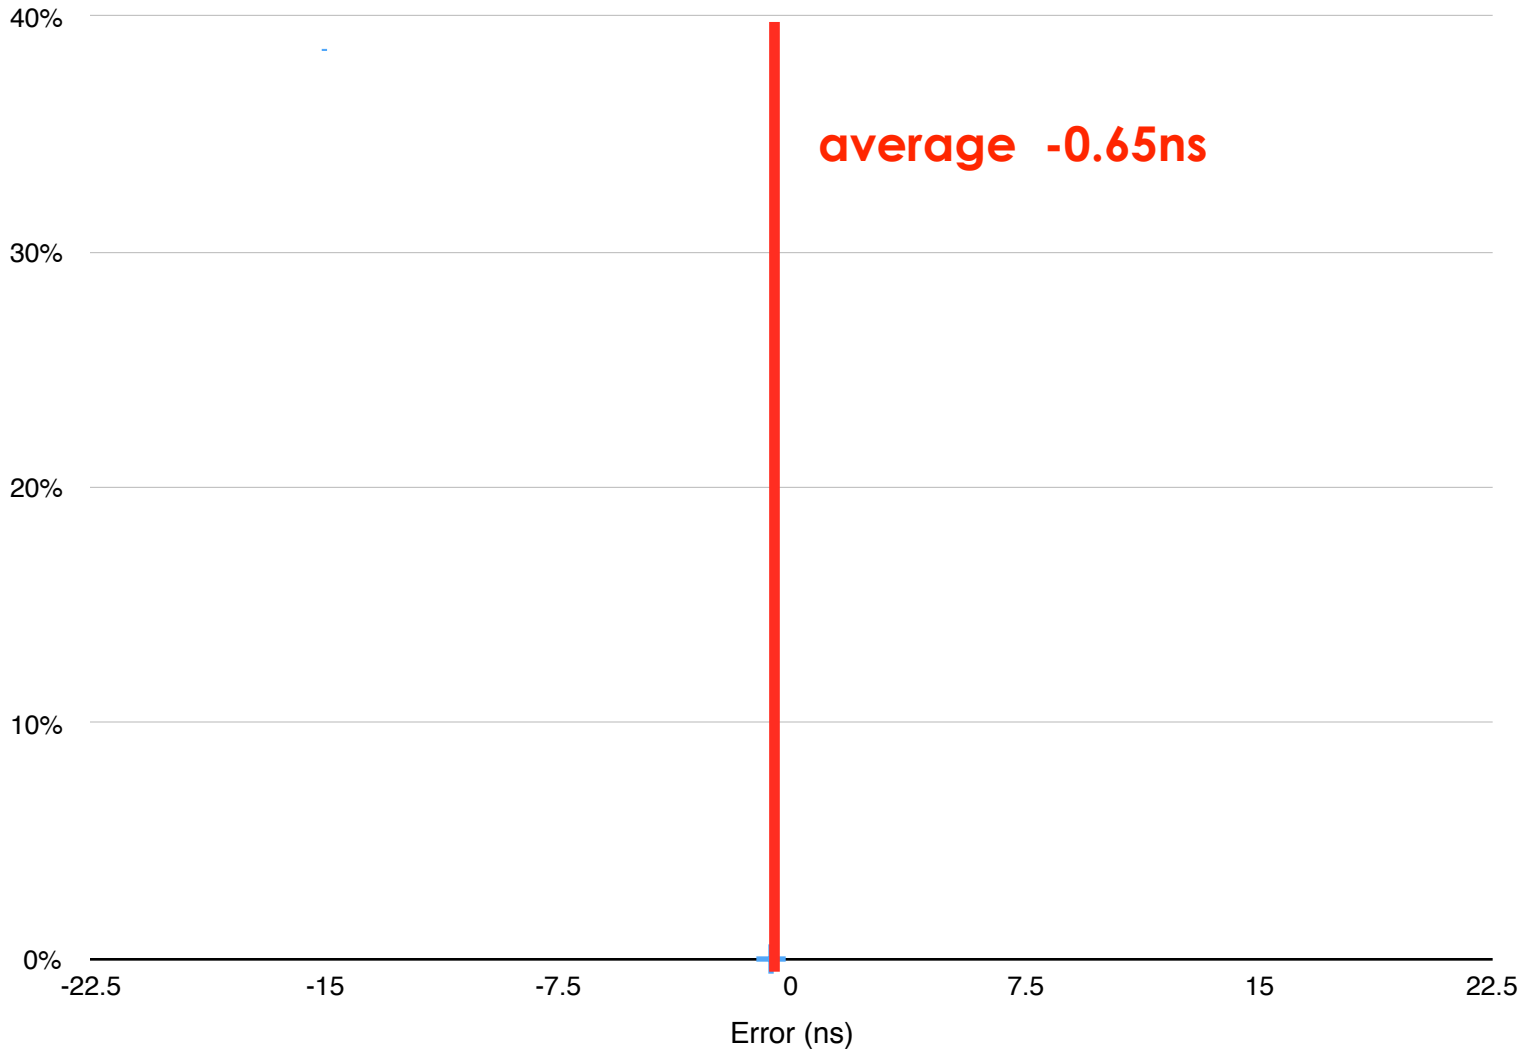

Device C (cleaned)

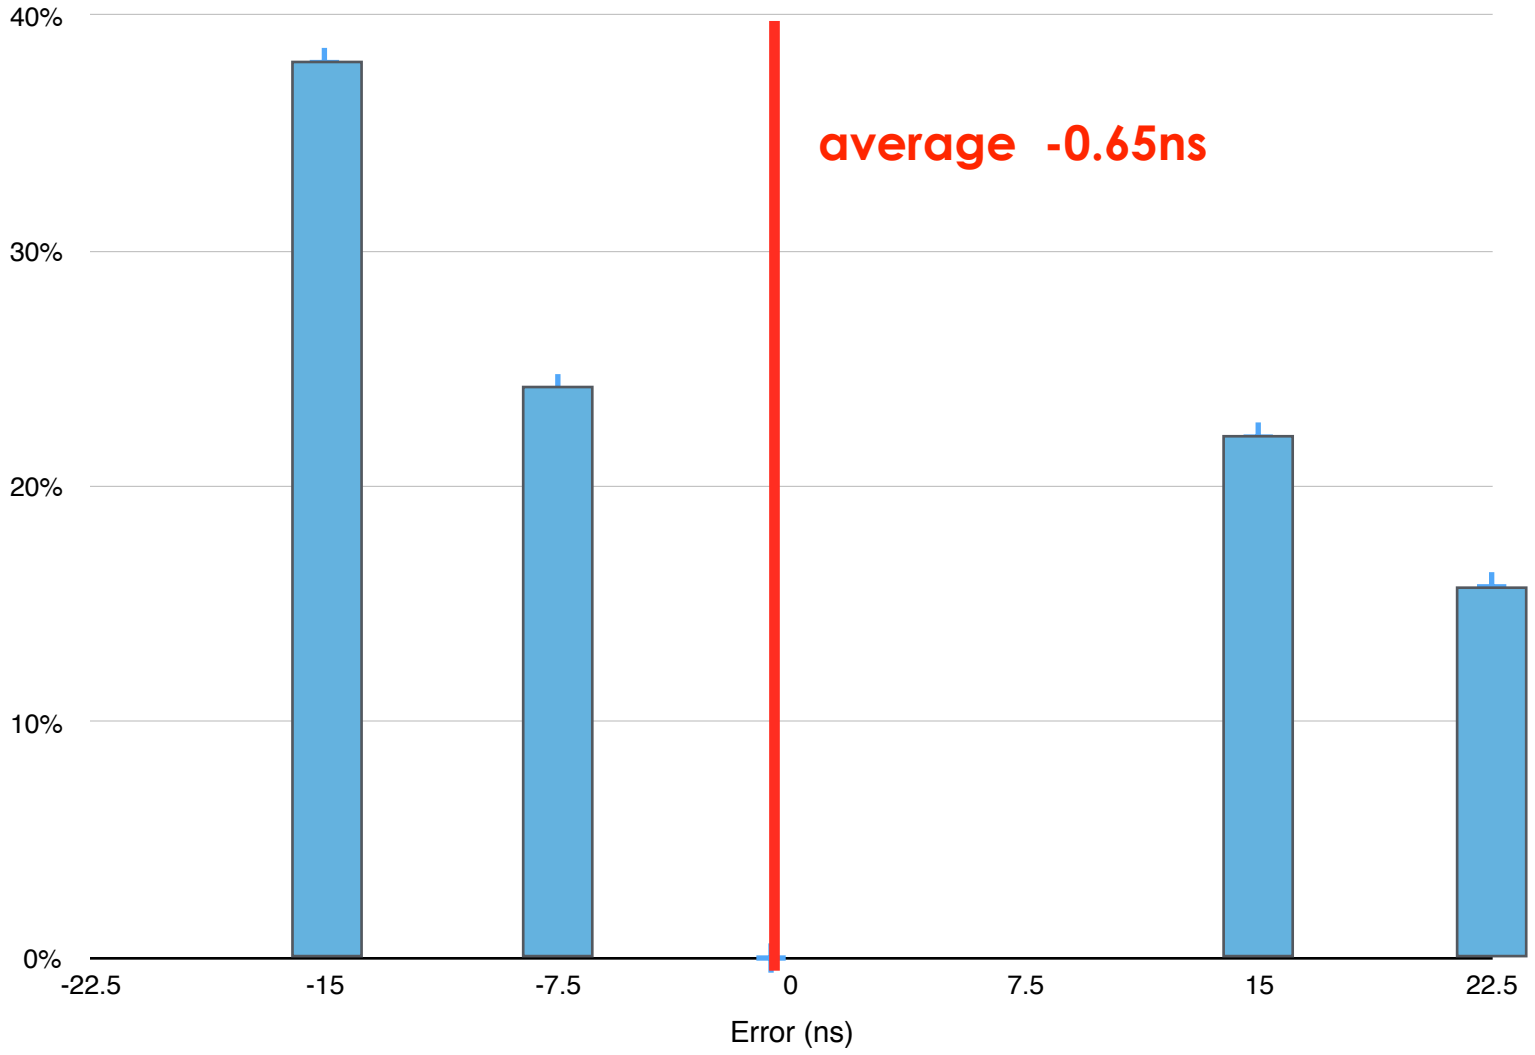

Device C (cleaned)

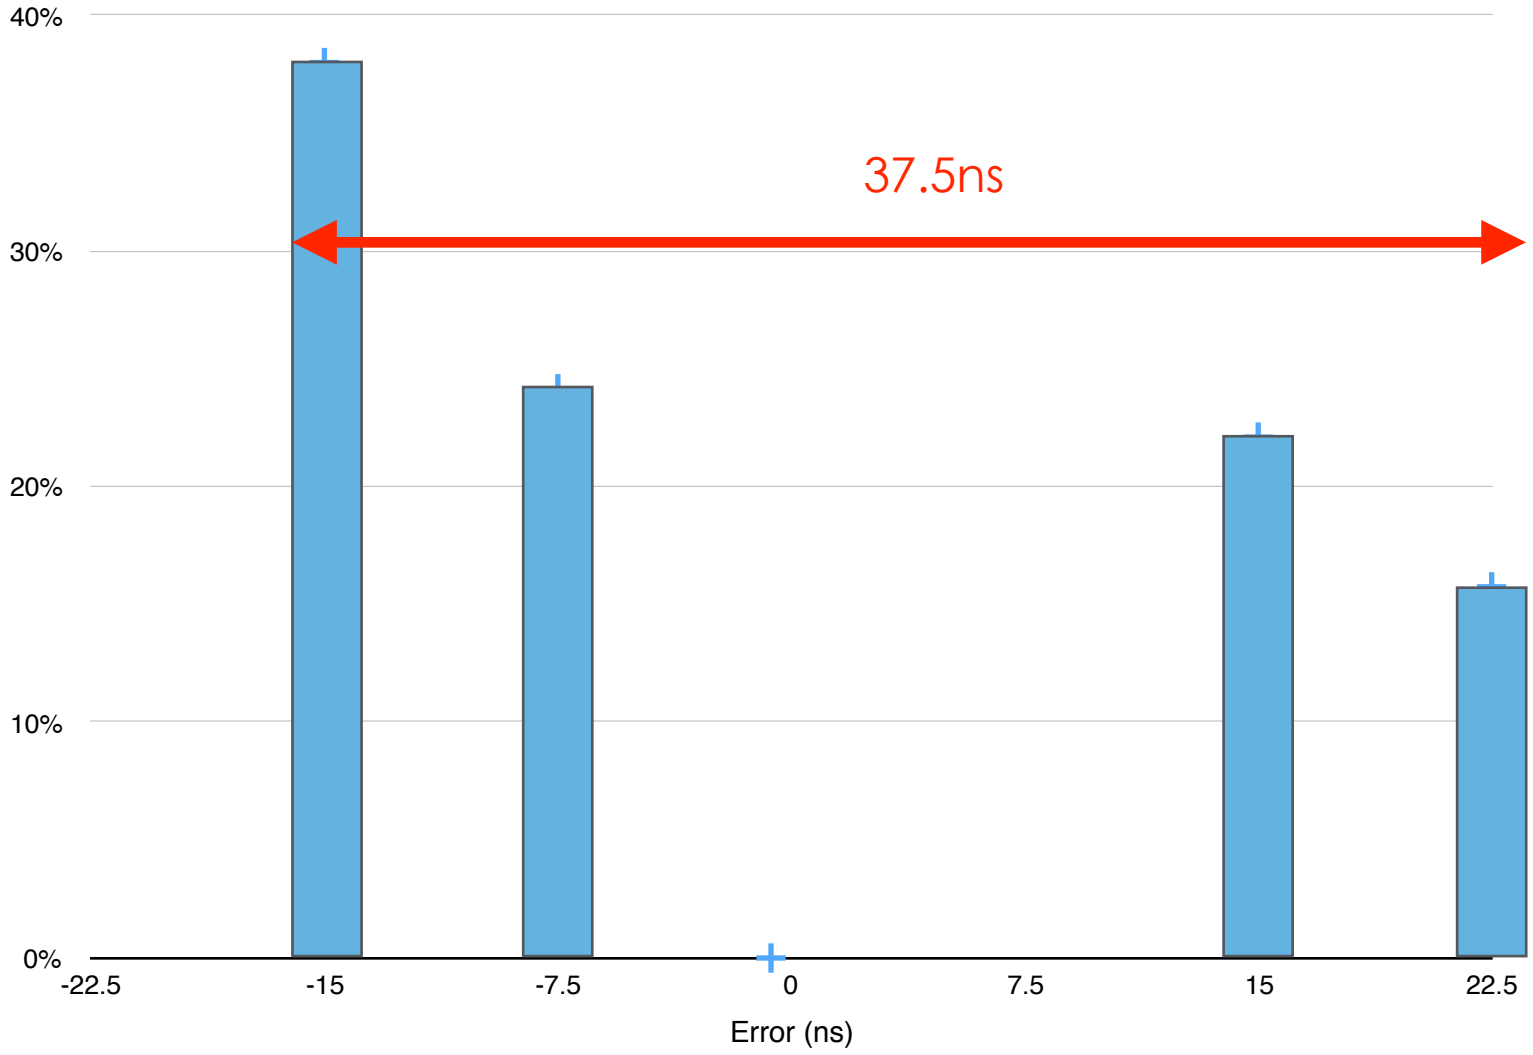

42ns ± 8ns

~40% error margin

42ns ± 18.75ns

42ns ± 18.75ns

~90% error margin!

Gotcha #2

Accuracy matters more than resolution

ExaNIC X10

resolution 6.2ns

10,000,000 samples

ExaNIC X10 (raw)

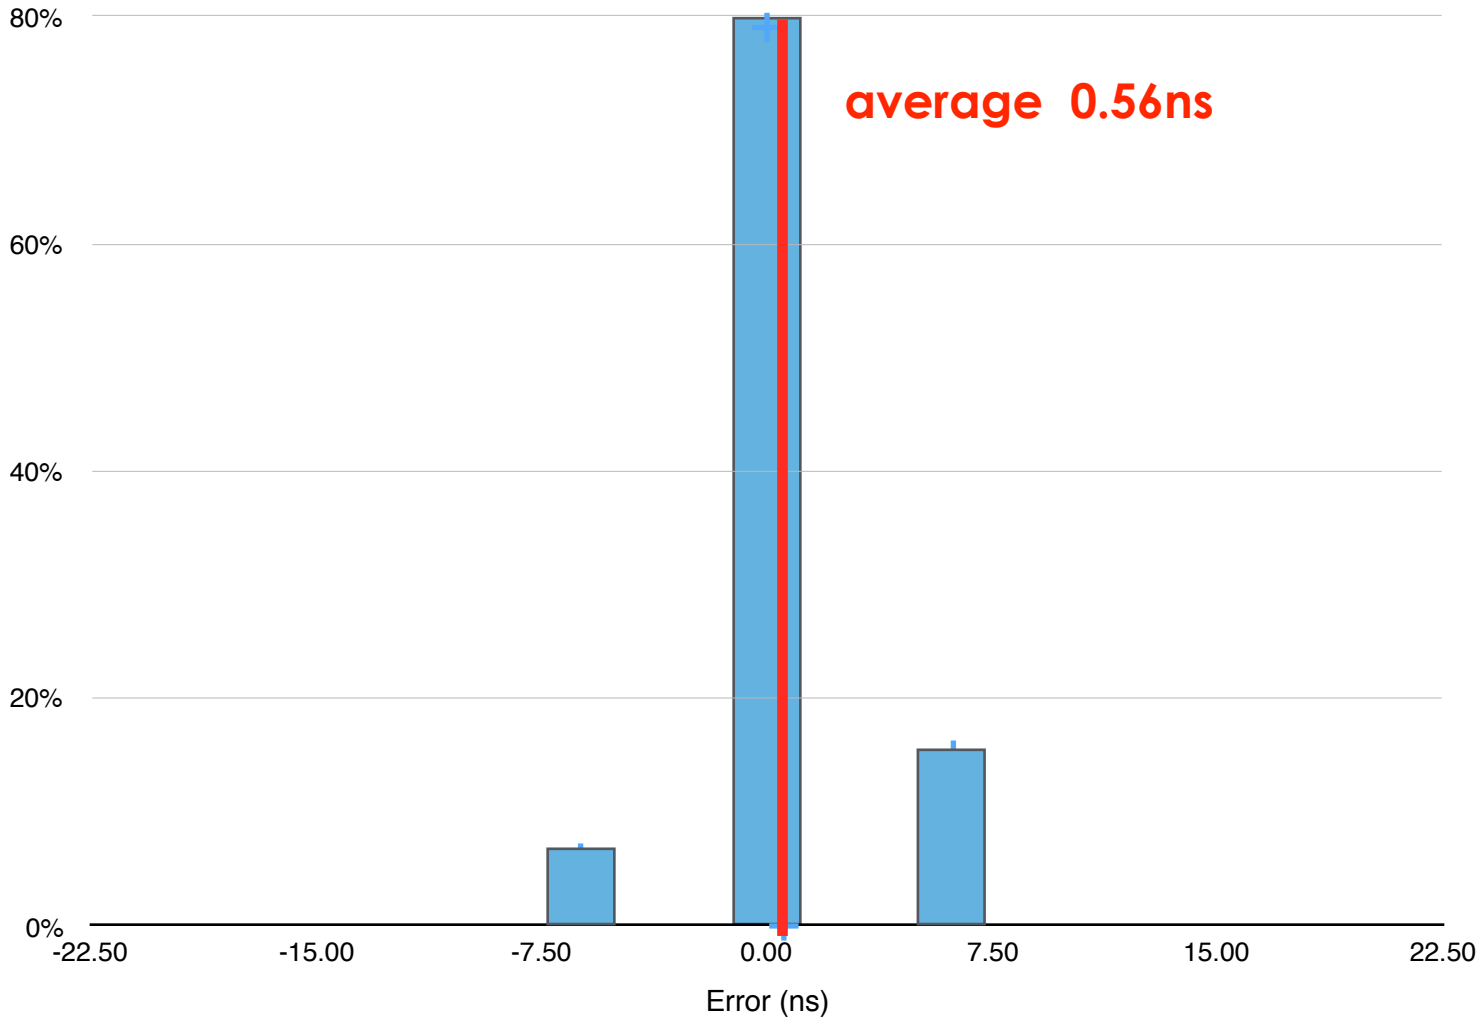

Remember tip #1?

Round to nearest multiple of clock period;

Tip #1B

Round to nearest multiple of clock period;

Or use a device/format that conveys the full precision

ExaNIC X10

resolution 6.2ns

accuracy ± 6.2ns

10,000,000 samples

Tip #2

Calibrate first!

Quantify Accuracy

Quantify Accuracy

Standard Deviation (σ) ?

Quantify Accuracy

Standard Deviation (σ) ?

$$\sigma = \sqrt{\frac{\sum_{i=1}^{n}(x_i - \bar{x})^2}{n-1}}$$

	Precision (ns)	Standard Dev (ns)
Device A	7.50	4.8
Device B	4.00	4.3
Device C	7.50	15.1
ExaNIC X10	6.20	2.7

ExaNIC HPT
High Precision Timing Capture

	Precision (ns)	Standard Dev (ns)
Device A	7.50	4.8
Device B	7.50	15.1
Device C	4.00	4.3
ExaNIC X10	6.20	2.7
ExaNIC HPT		

	Precision (ns)	Standard Dev (ns)
Device A	7.50	4.8
Device B	7.50	15.1
Device C	4.00	4.3
ExaNIC X10	6.20	2.7
ExaNIC HPT	0.25	

	Precision (ns)	Standard Dev (ns)
Device A	7.50	4.8
Device B	7.50	15.1
Device C	4.00	4.3
ExaNIC X10	6.20	2.7
ExaNIC HPT	0.25	

16x

	Precision (ns)	Standard Dev (ns)
Device A	7.50	4.8
Device B	7.50	15.1
Device C	4.00	4.3
ExaNIC X10	6.20	2.7
ExaNIC HPT	0.25	0.1

	Precision (ns)	Standard Dev (ns)
Device A	7.50	4.8
Device B	7.50	15.1
Device C	4.00	4.3
ExaNIC X10	6.20	2.7
ExaNIC HPT	0.25	0.1

27x

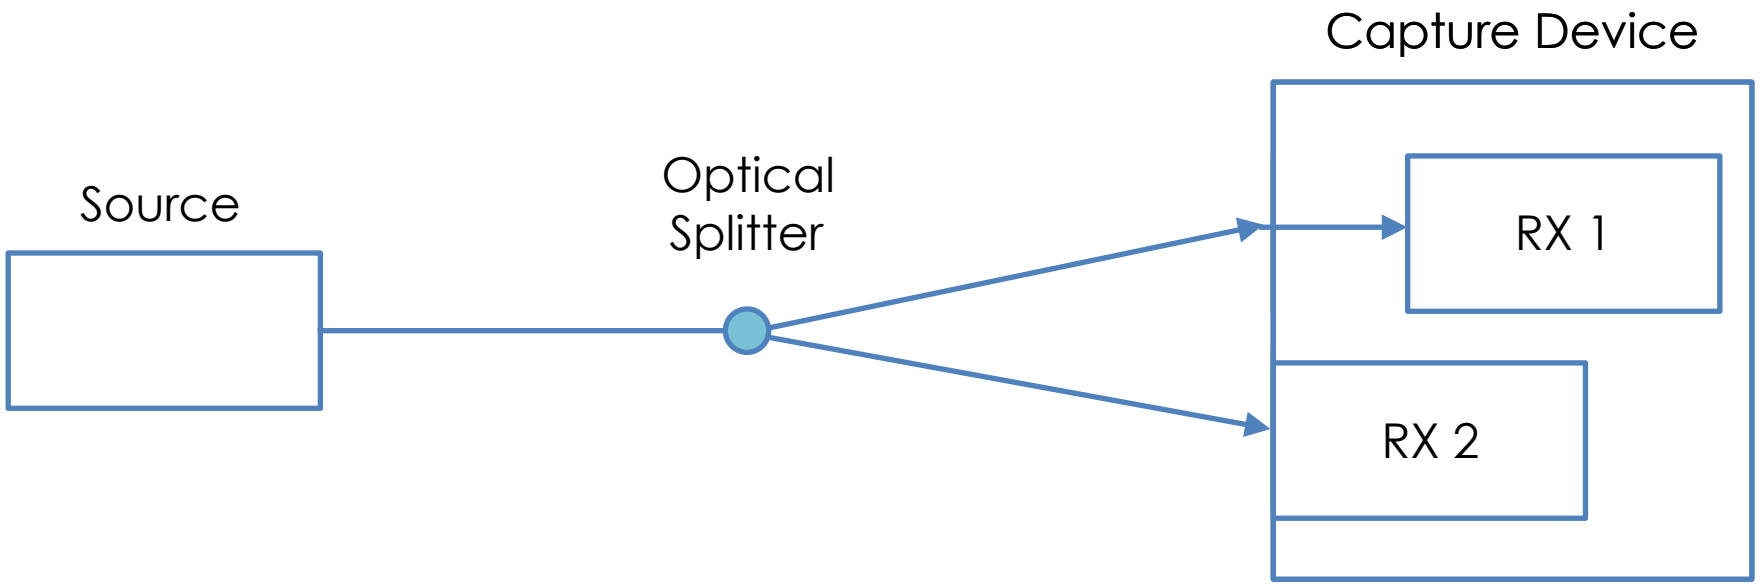

$$error = RX_2 - RX_1$$

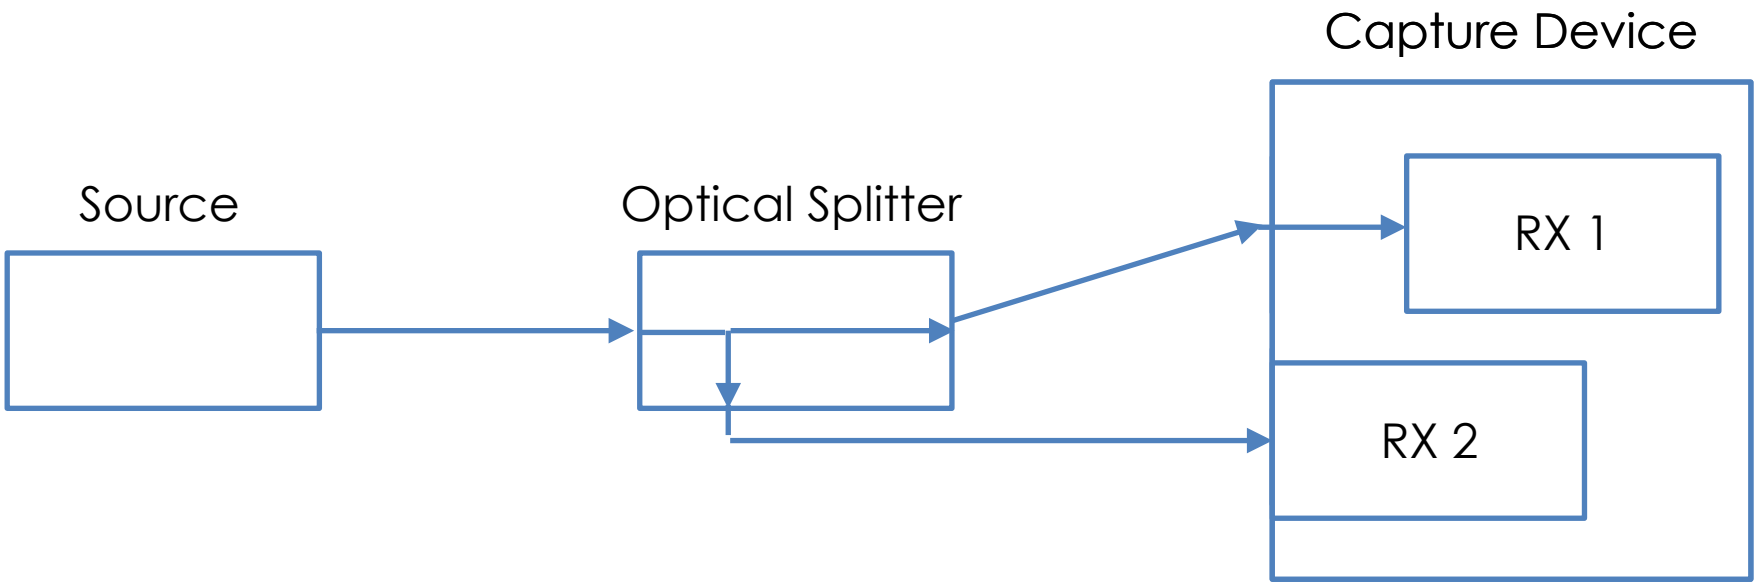

$$error = RX_2 - RX_1$$

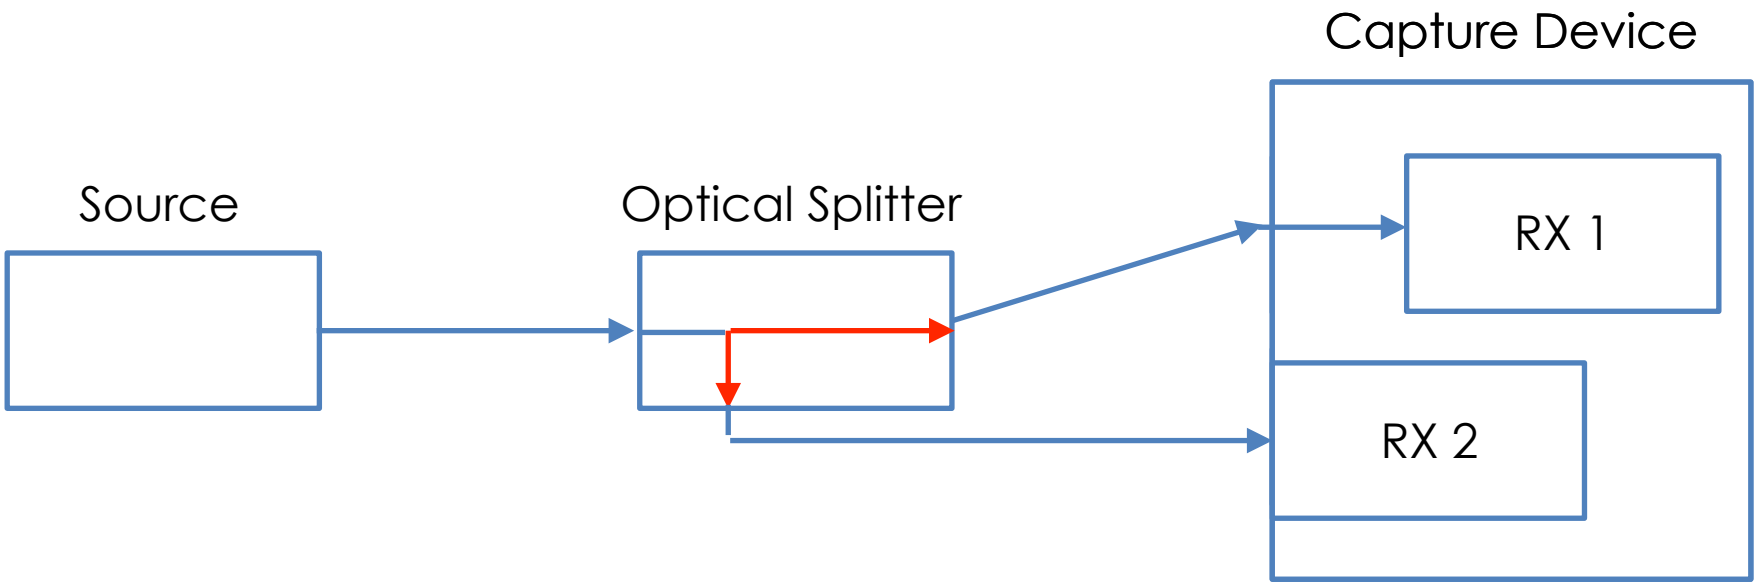

$$error = RX_2 - RX_1$$

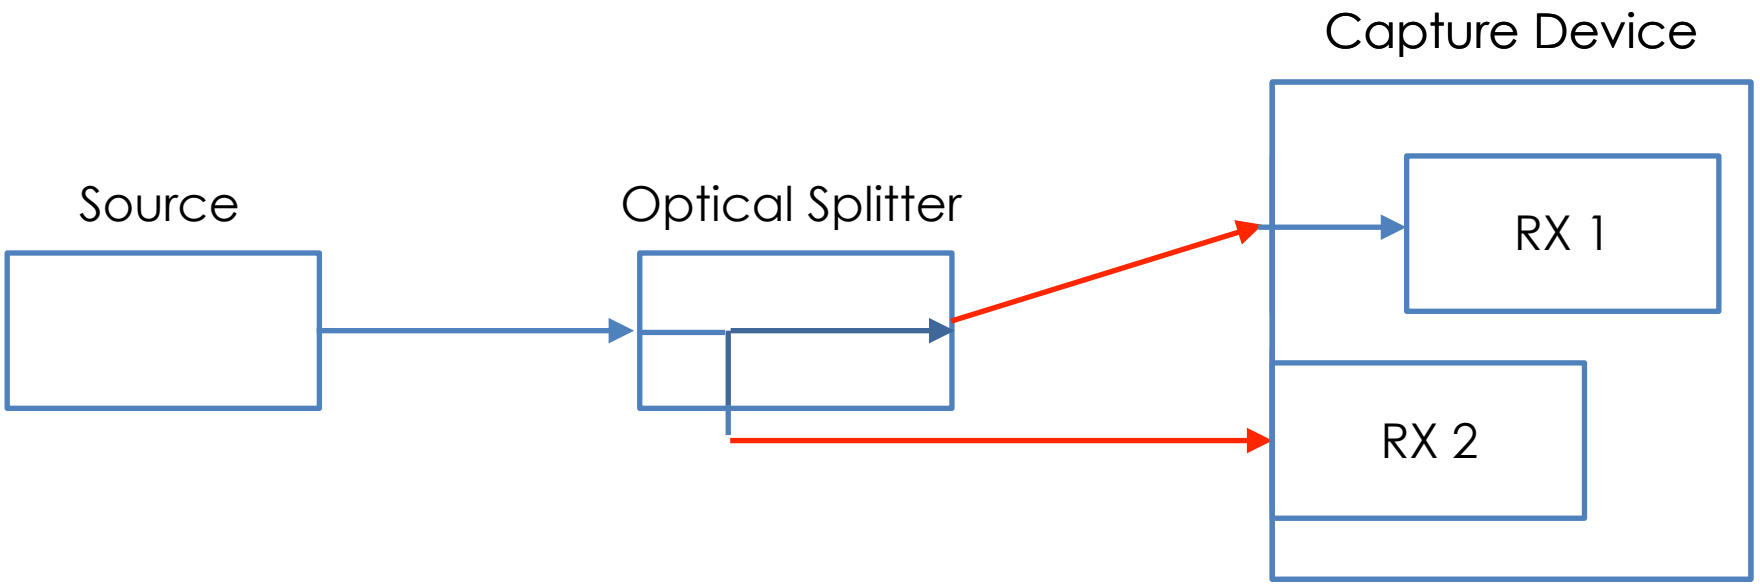

$$error = RX_2 - RX_1$$

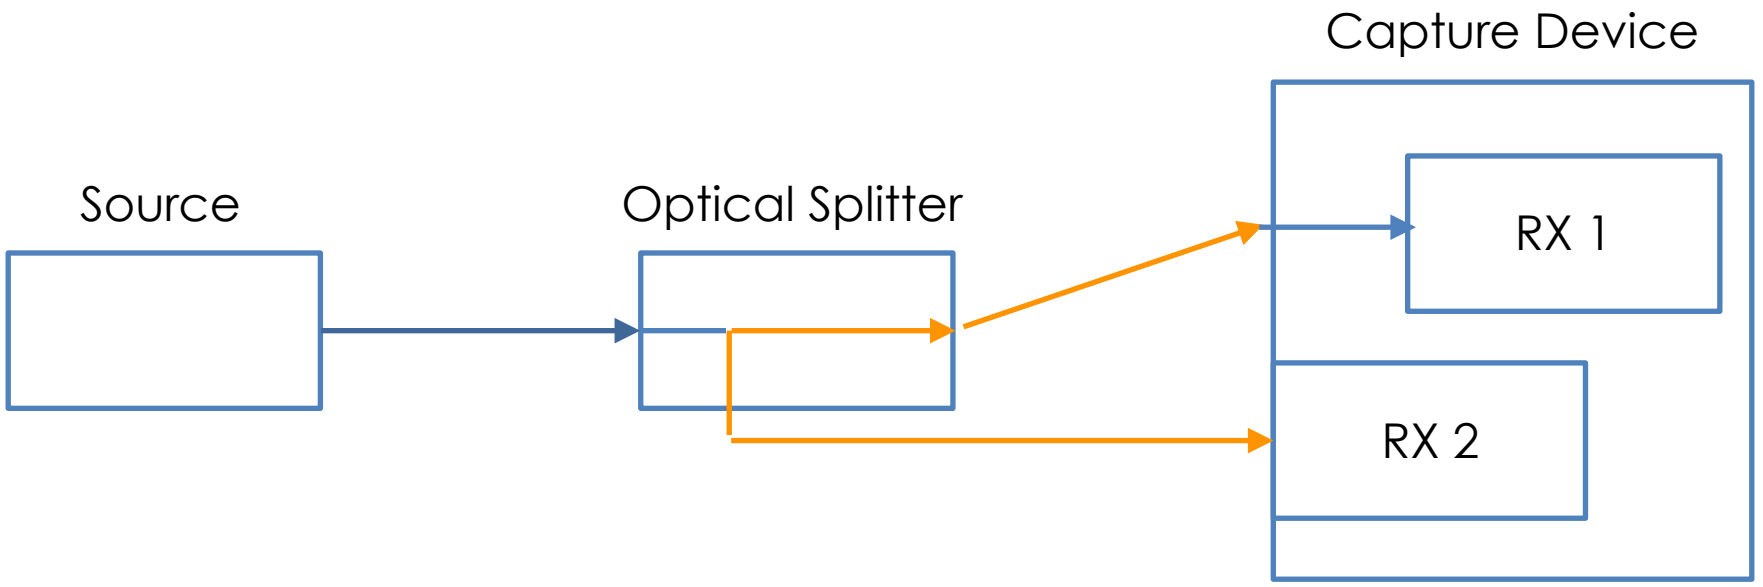

$$error = RX_2 - RX_1$$

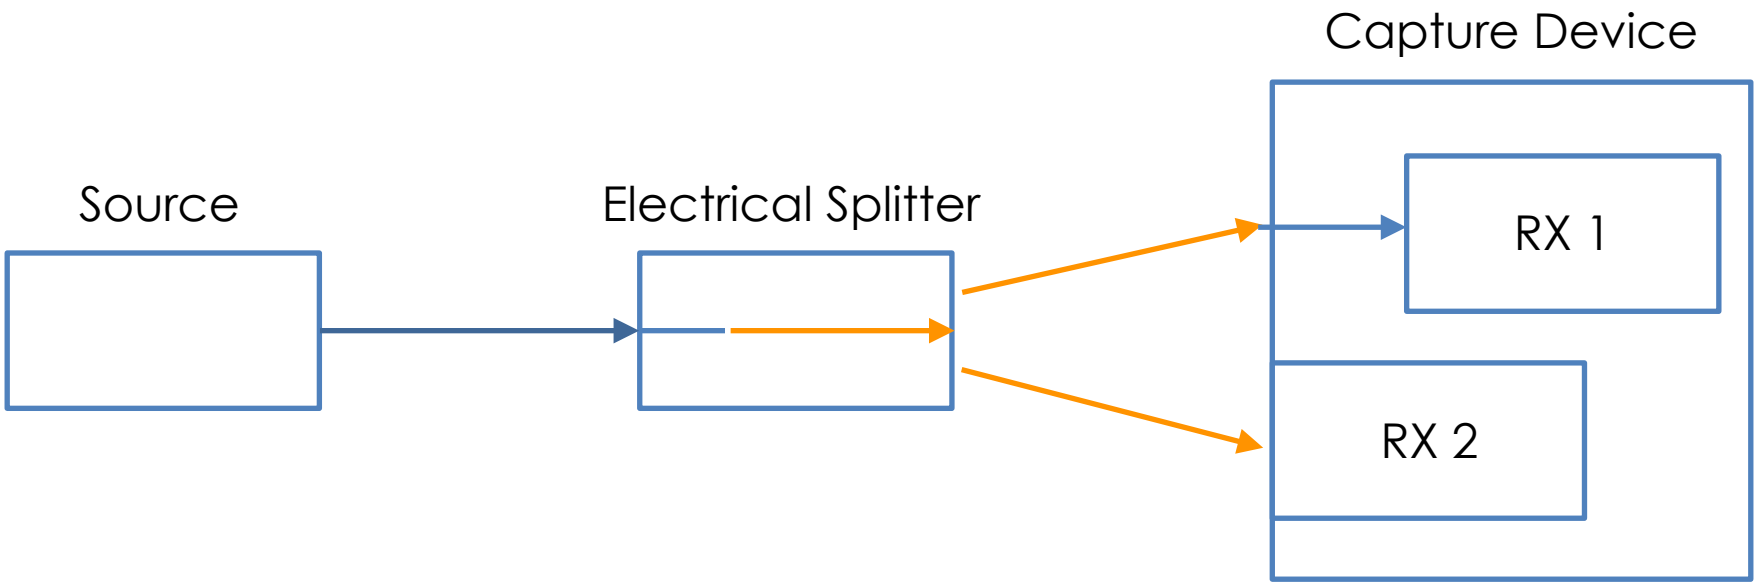

$$error = RX_2 - RX_1$$

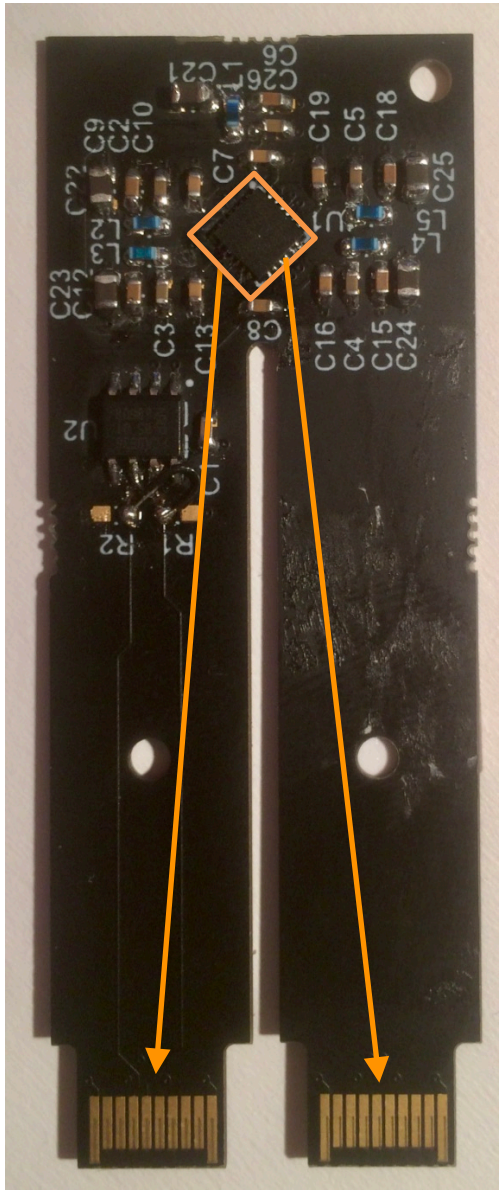
C21
C11
C23
C9
C2
C10
C19
C5
C18
C7
C25
C22
C23
C3
C13
C8
C16
C4
C15
C24
J2
R2
R1

ExaNIC HPT

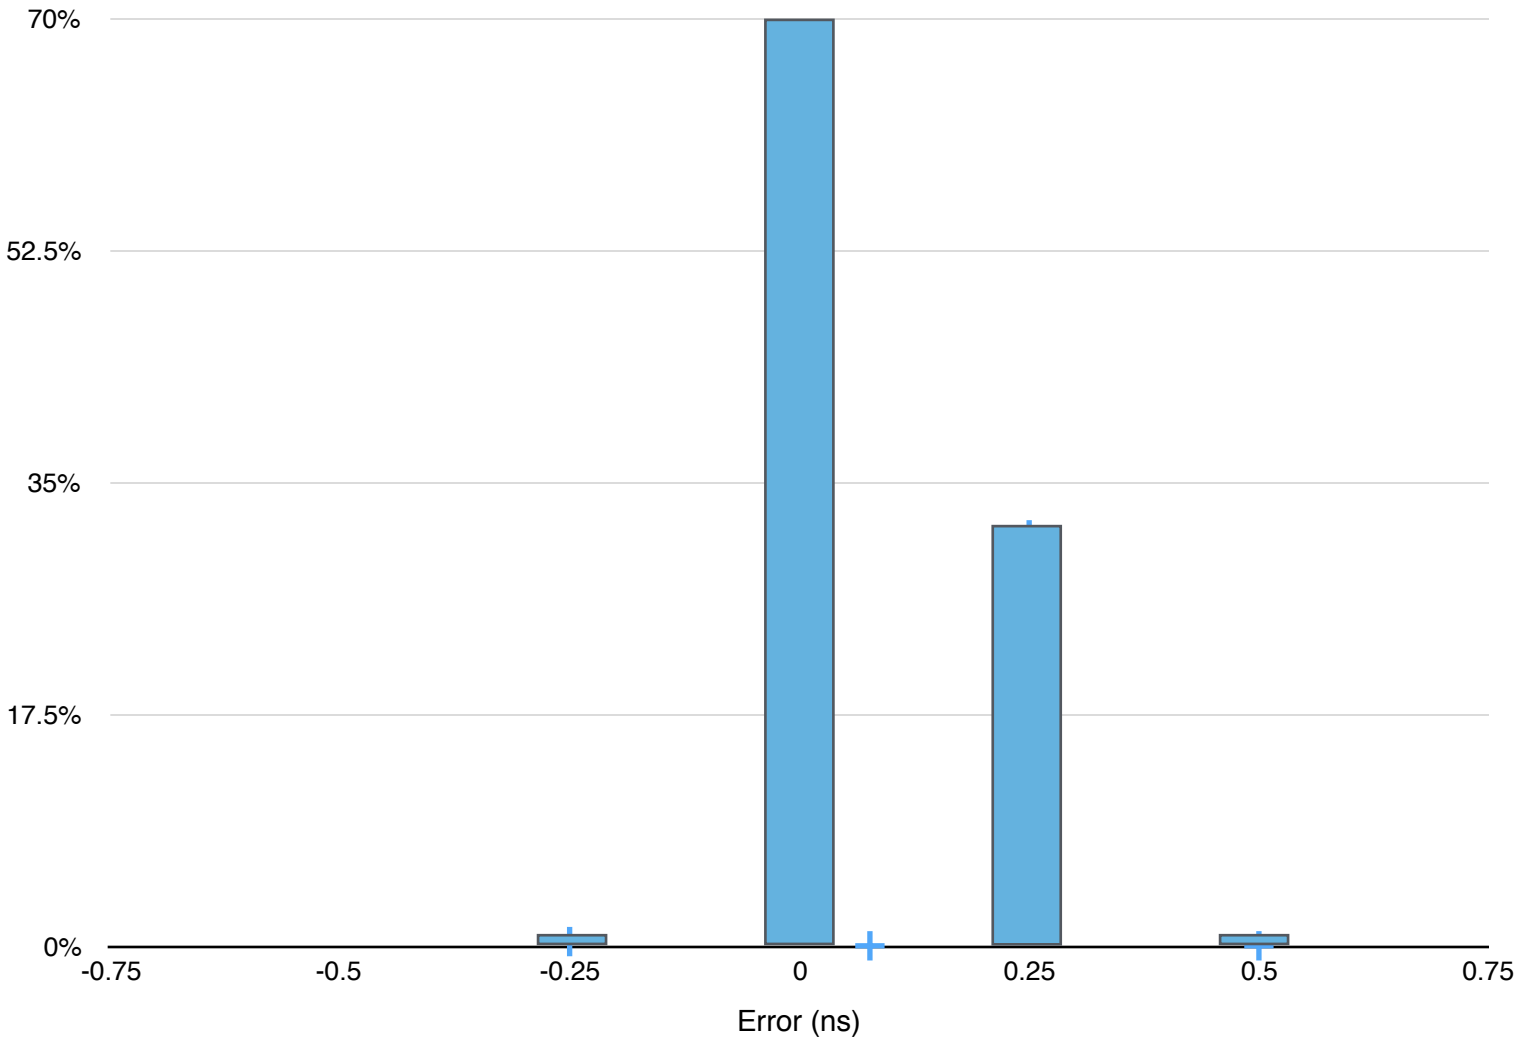

ExaNIC HPT

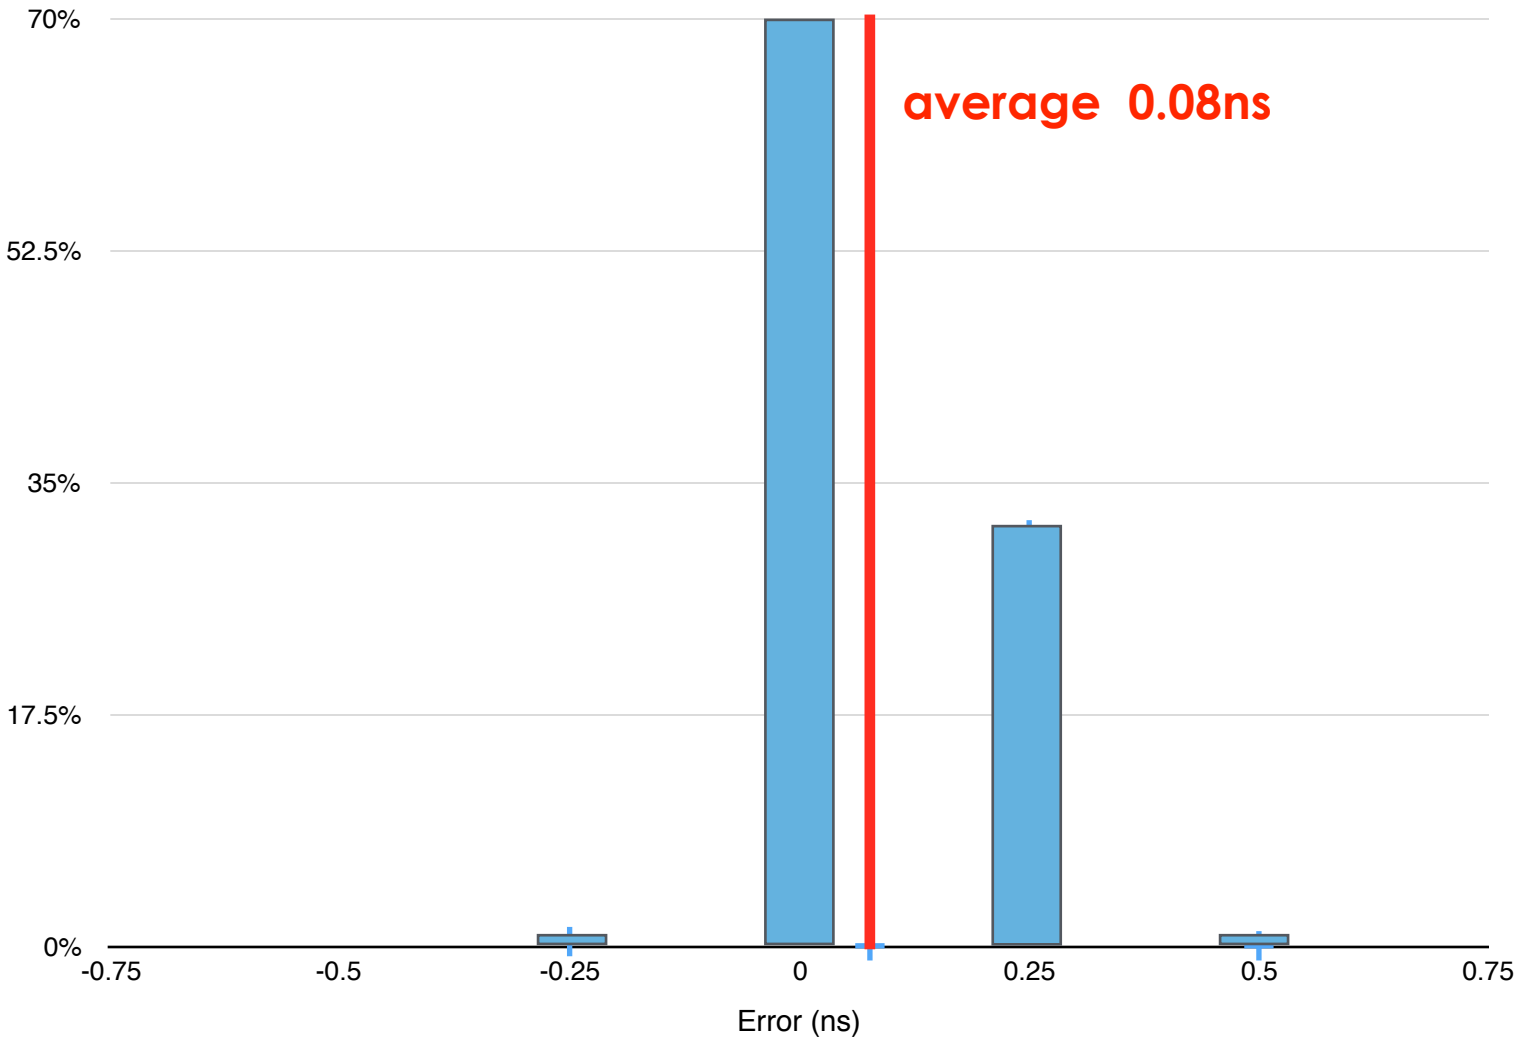

ExaNIC HPT

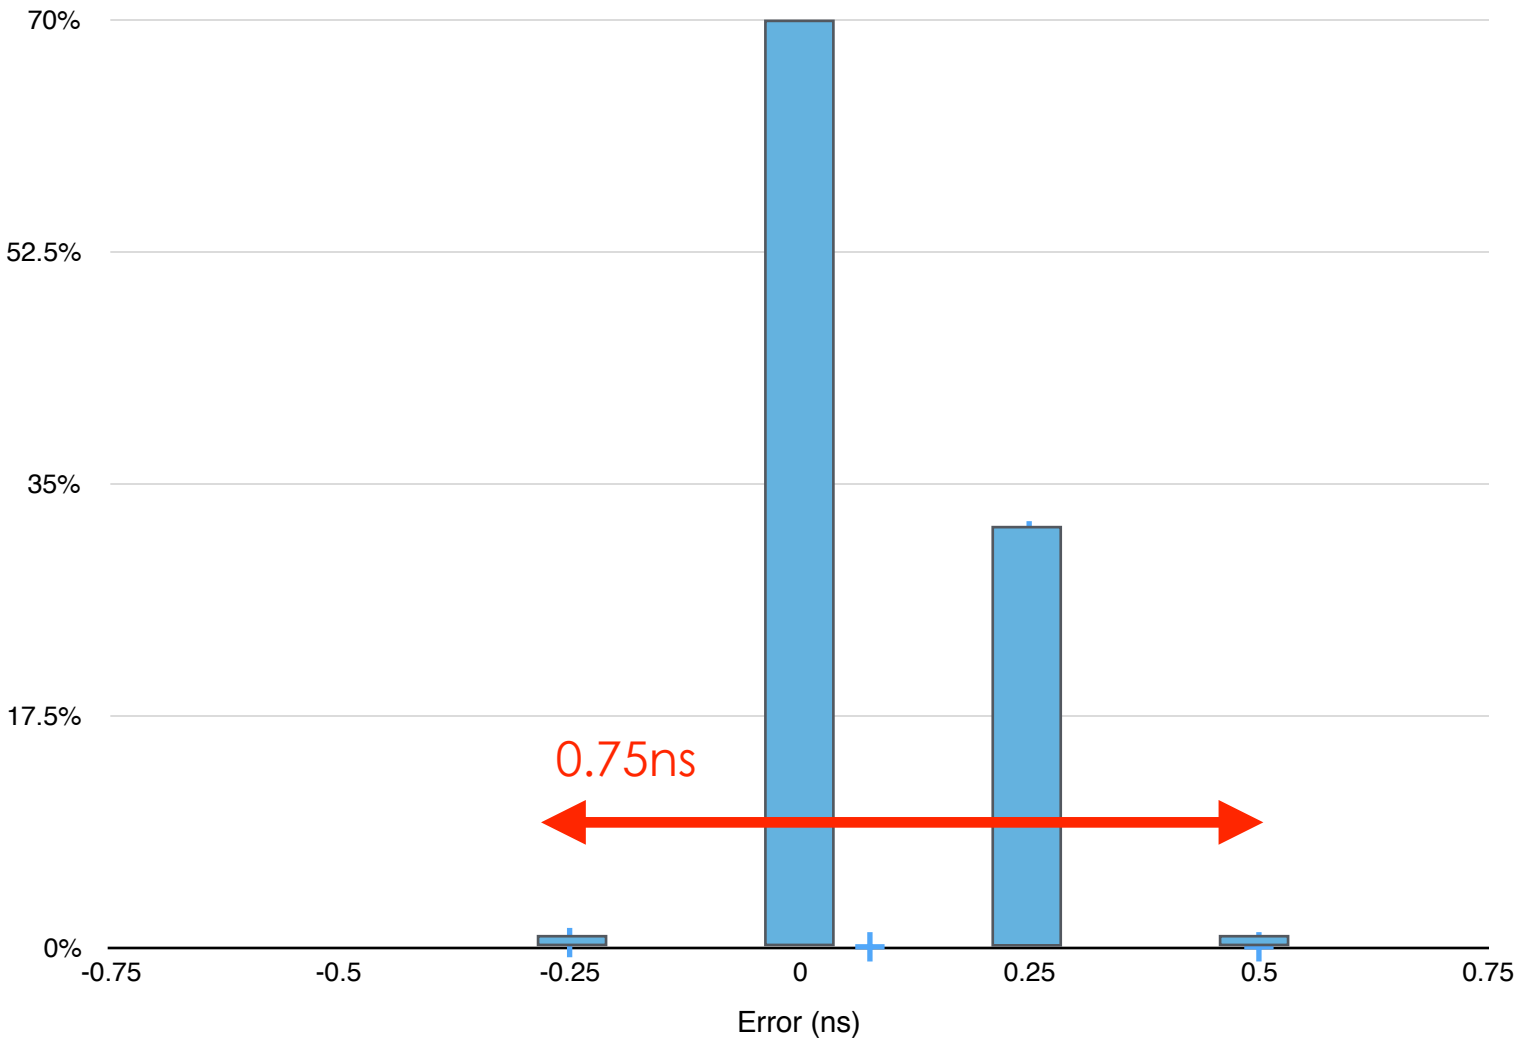

ExaNIC HPT vs X10

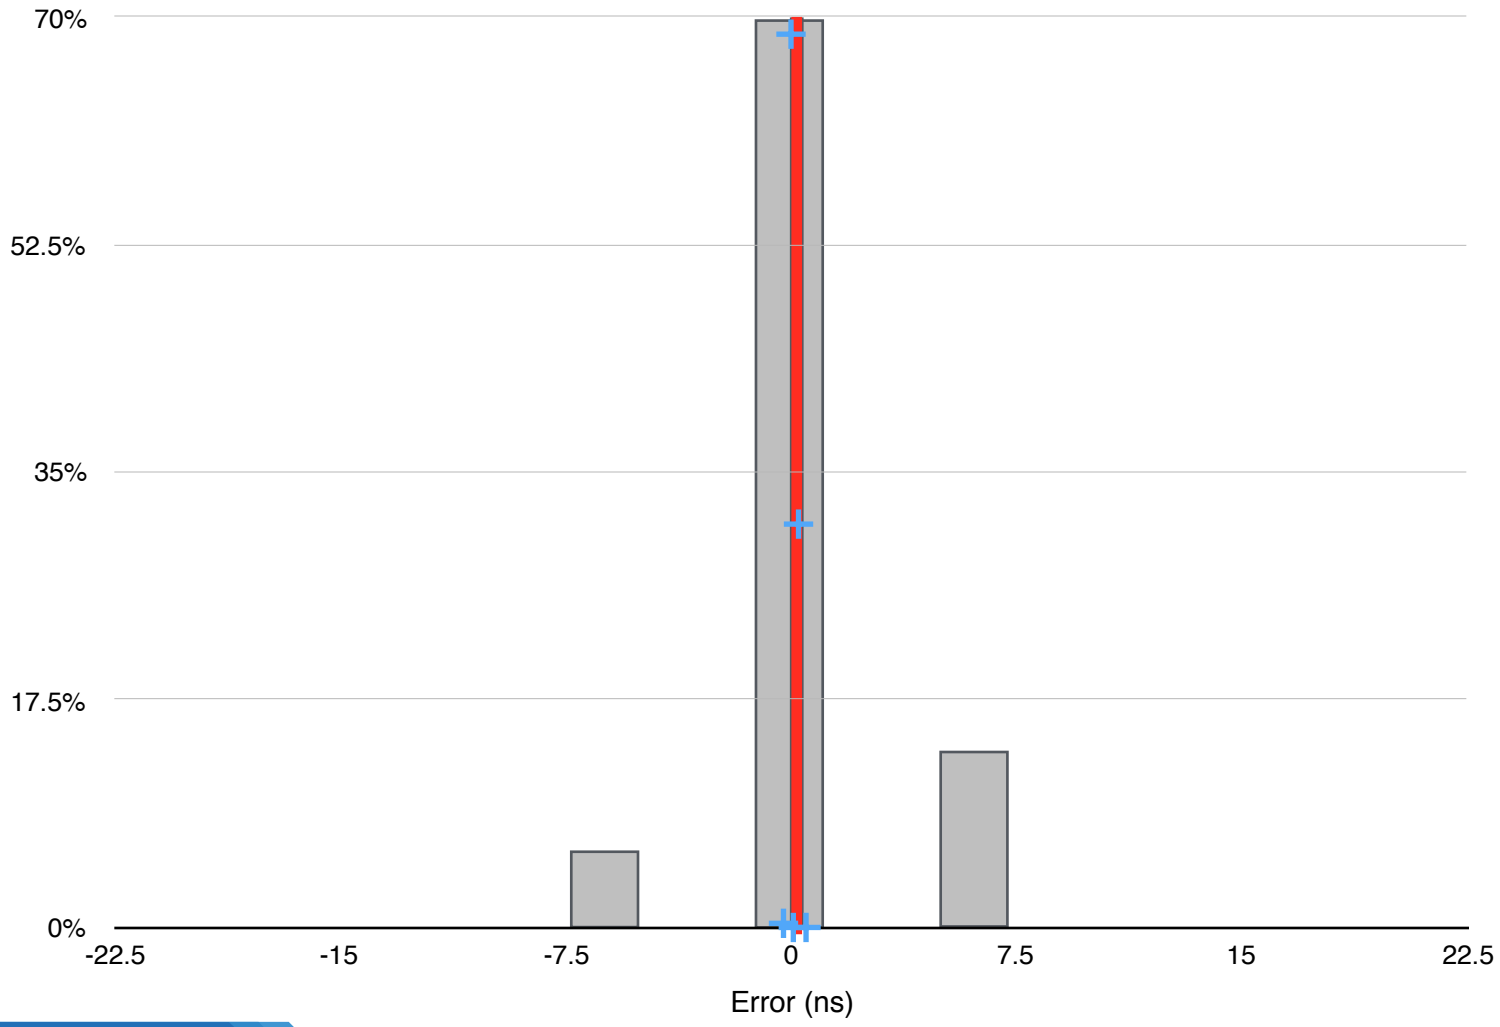

Picosecond measurement of something real?

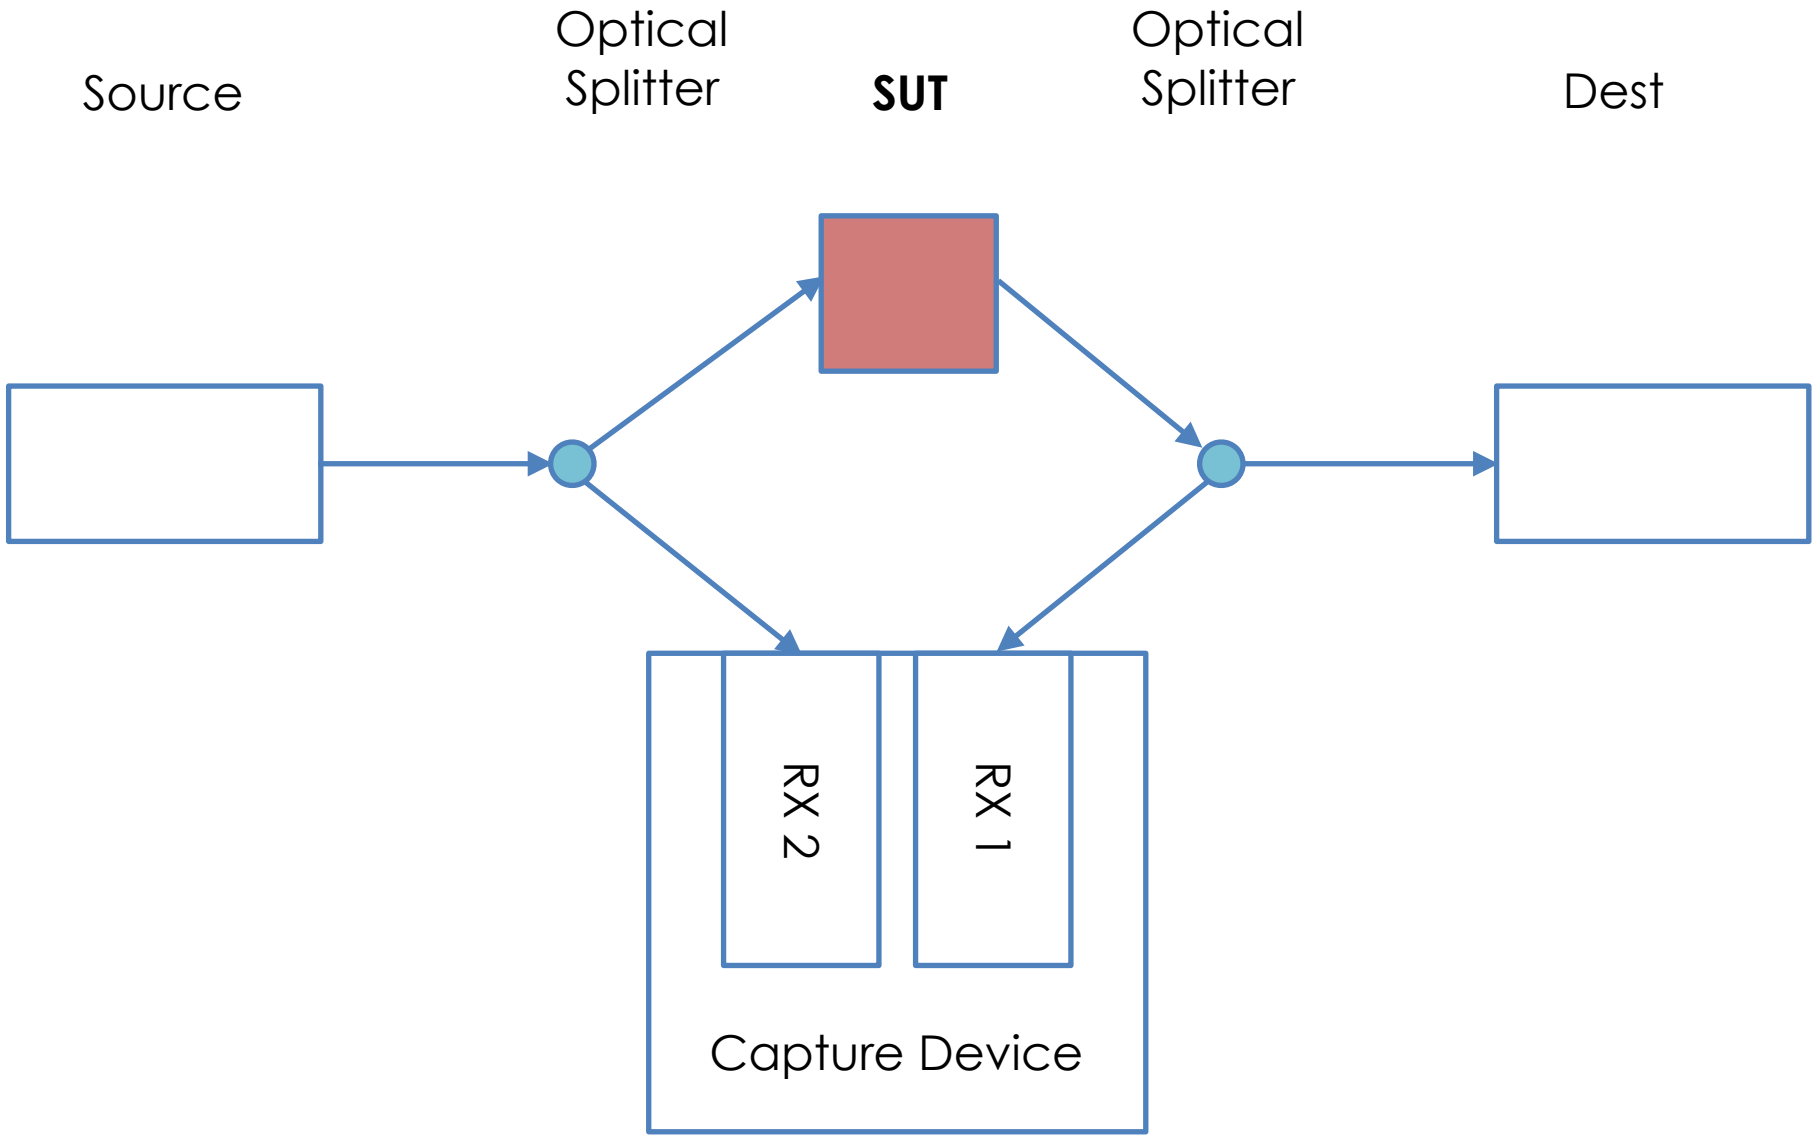
Source
Optical Splitter
SUT
Optical Splitter
Dest
RX 2
RX 1
Capture Device

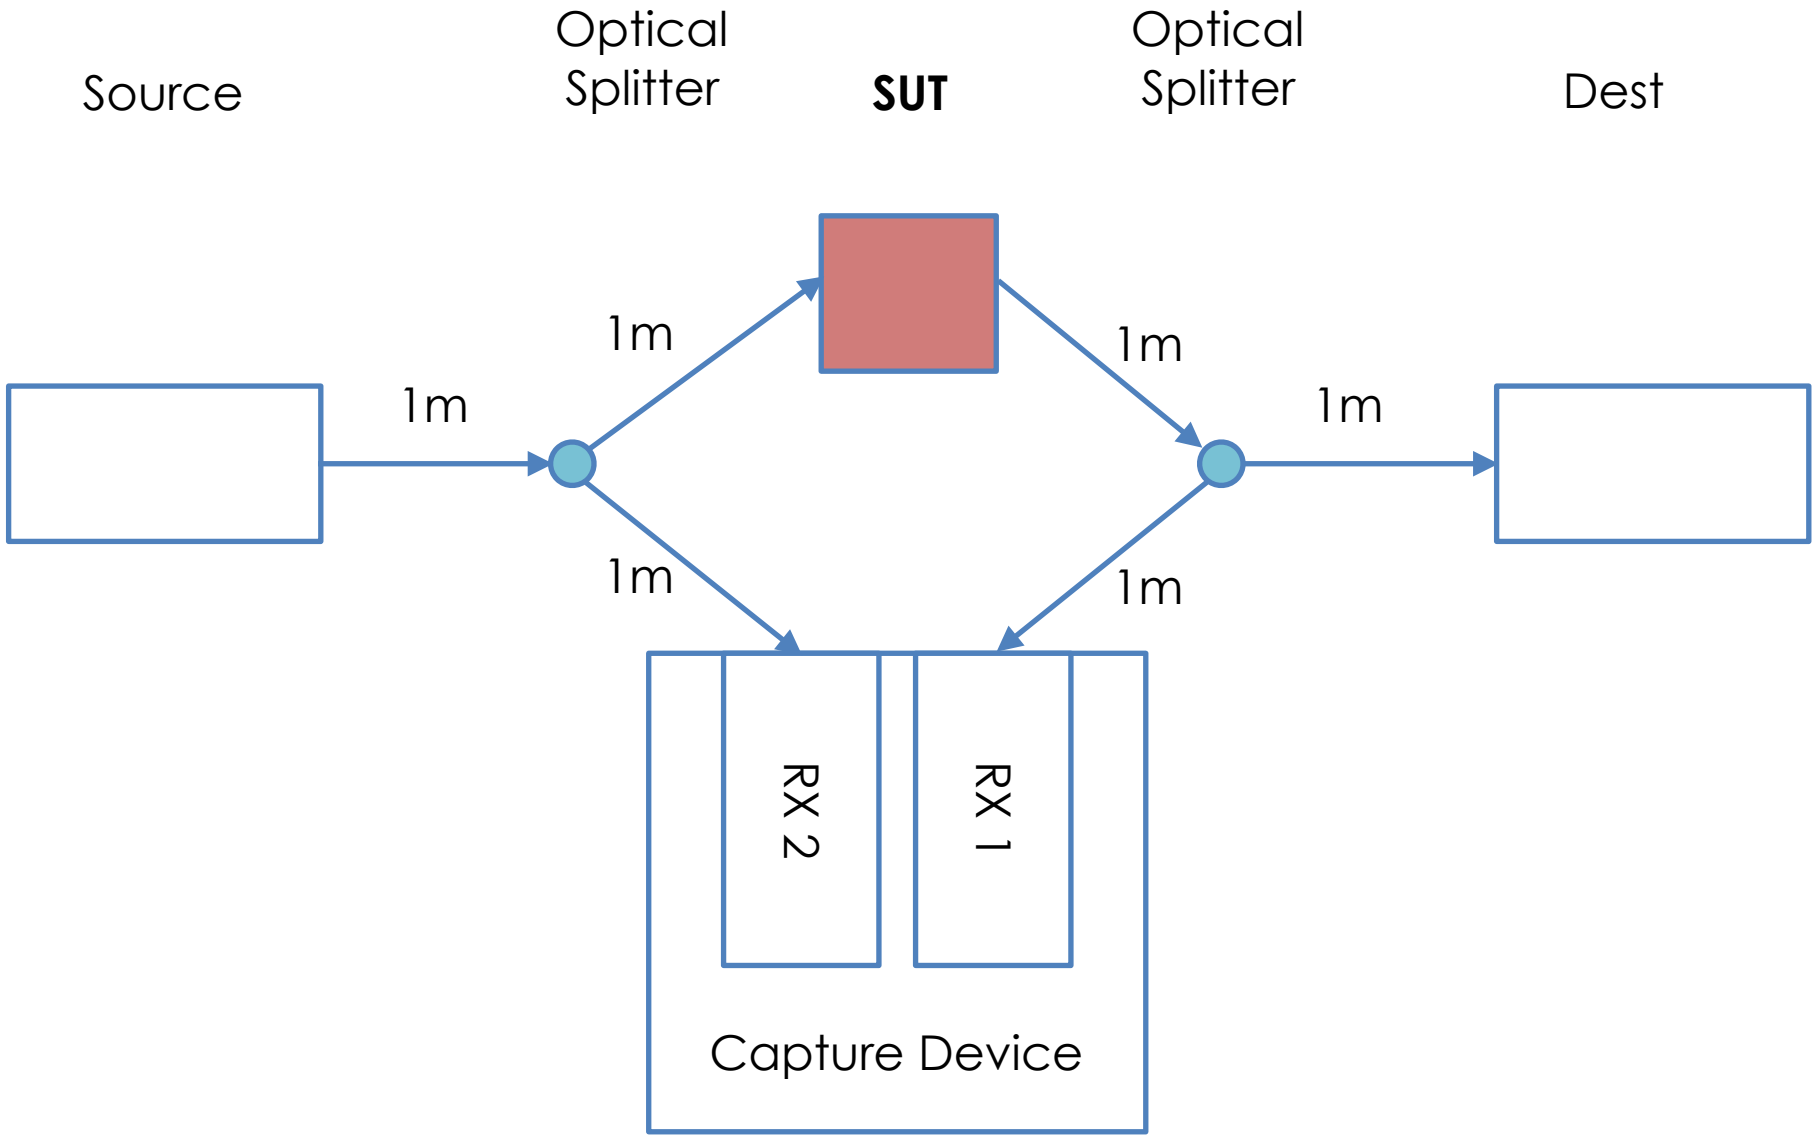
Source
Optical Splitter
SUT
Optical Splitter
Dest
1m
1m
1m
1m
1m
1m
RX 2
RX 1
Capture Device

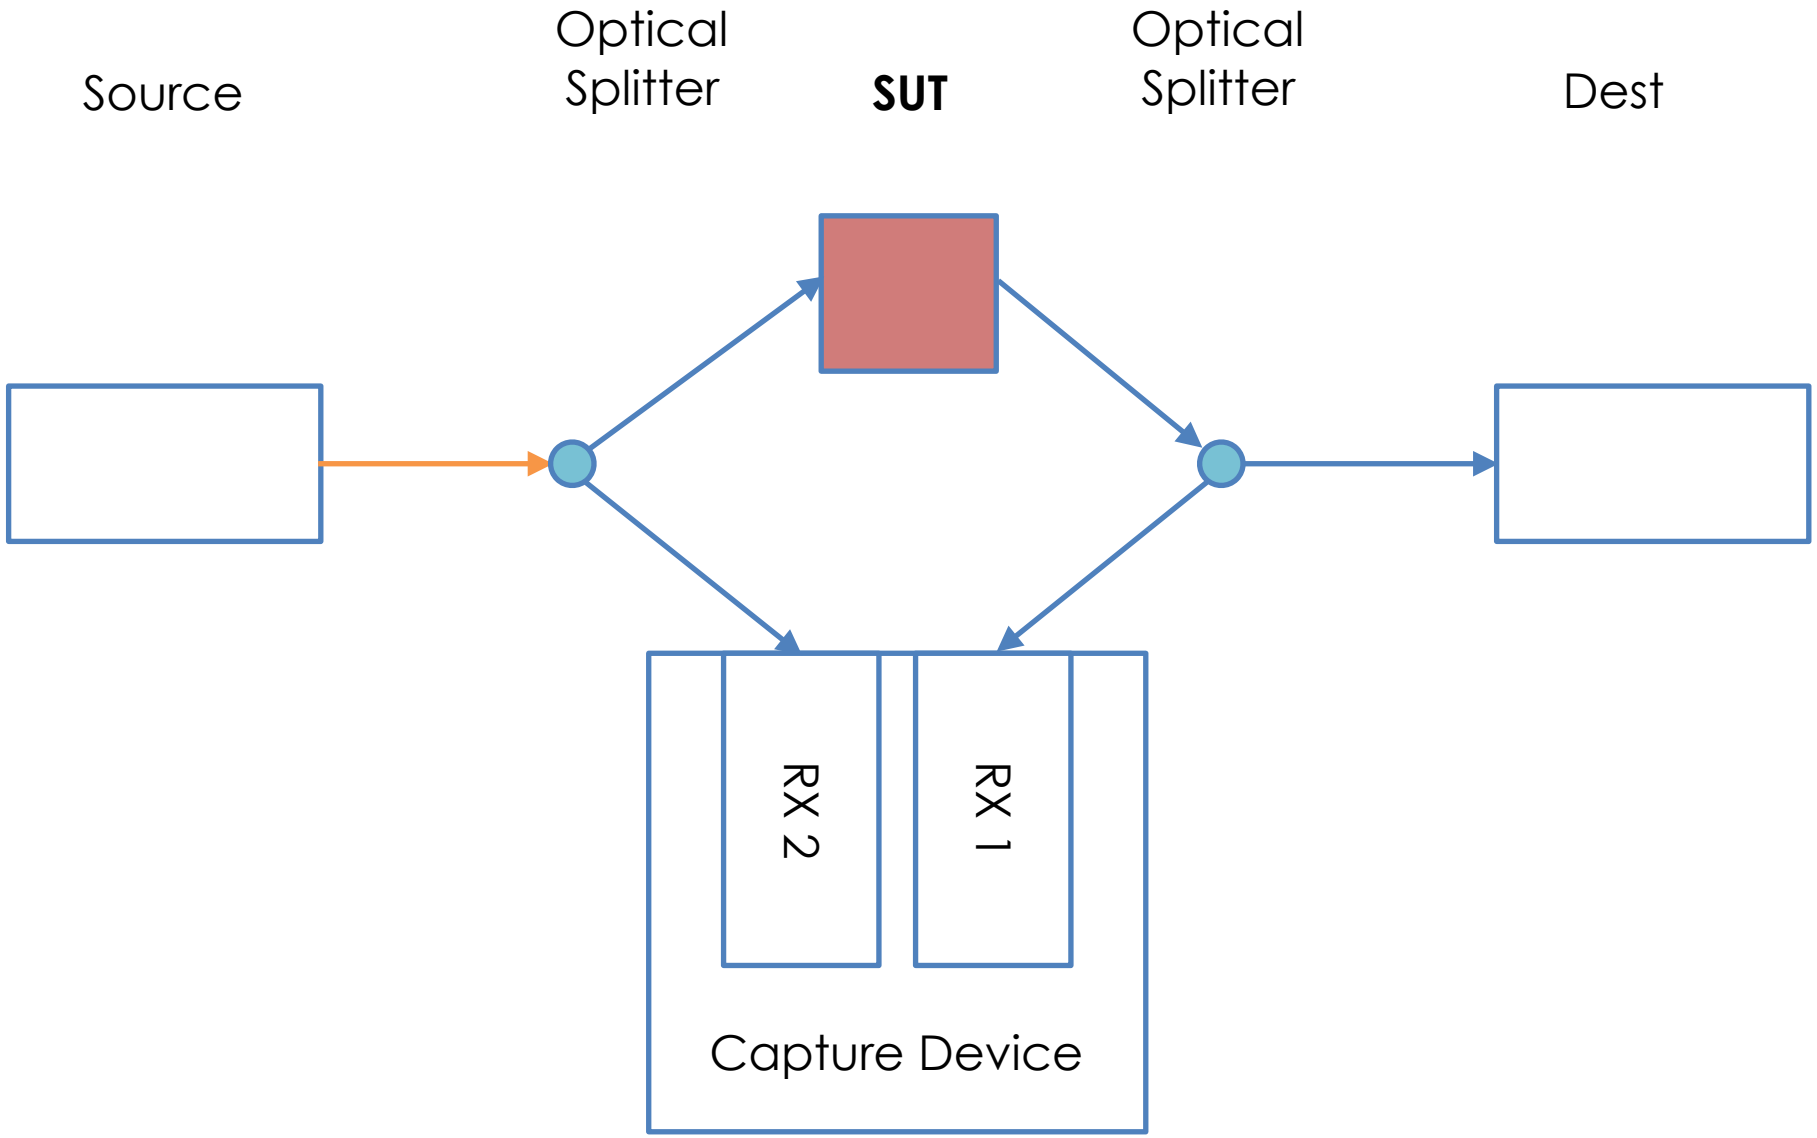
Source
Optical Splitter
SUT
Optical Splitter
Dest
RX 2
RX 1
Capture Device

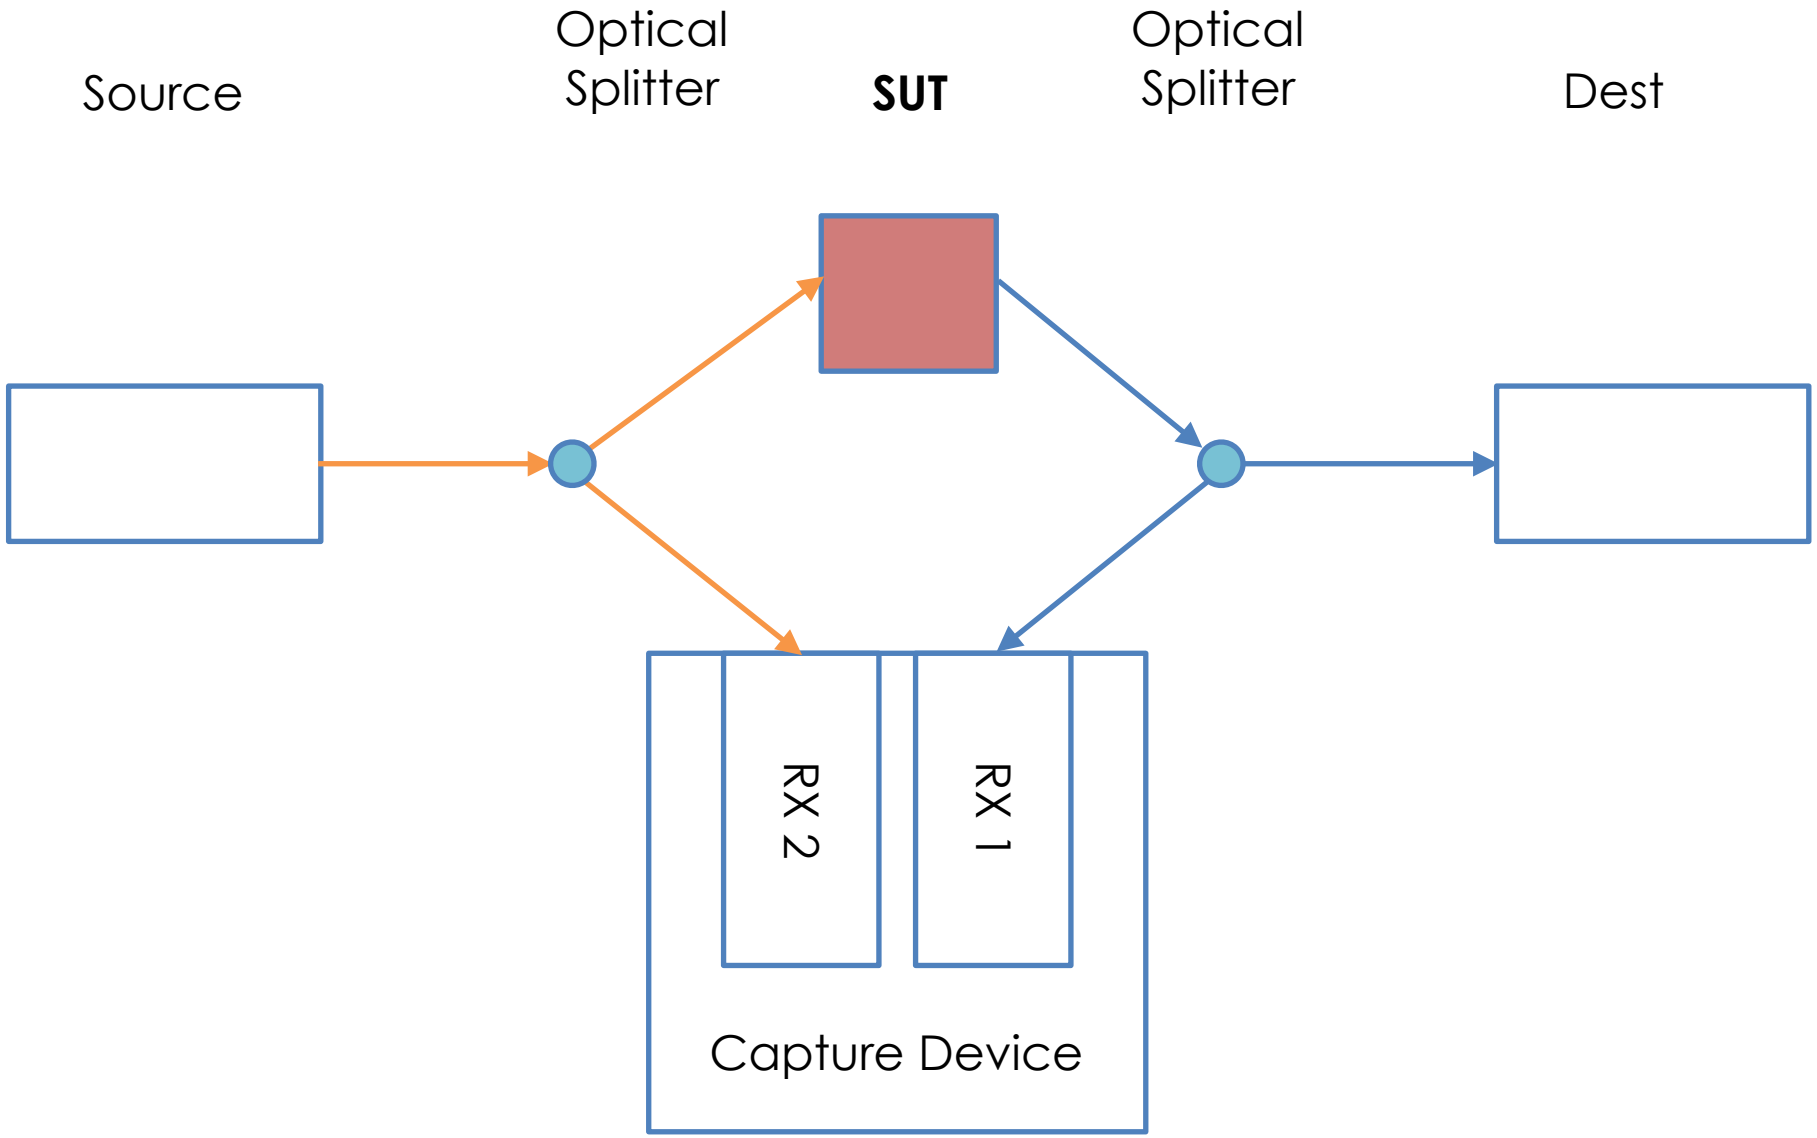

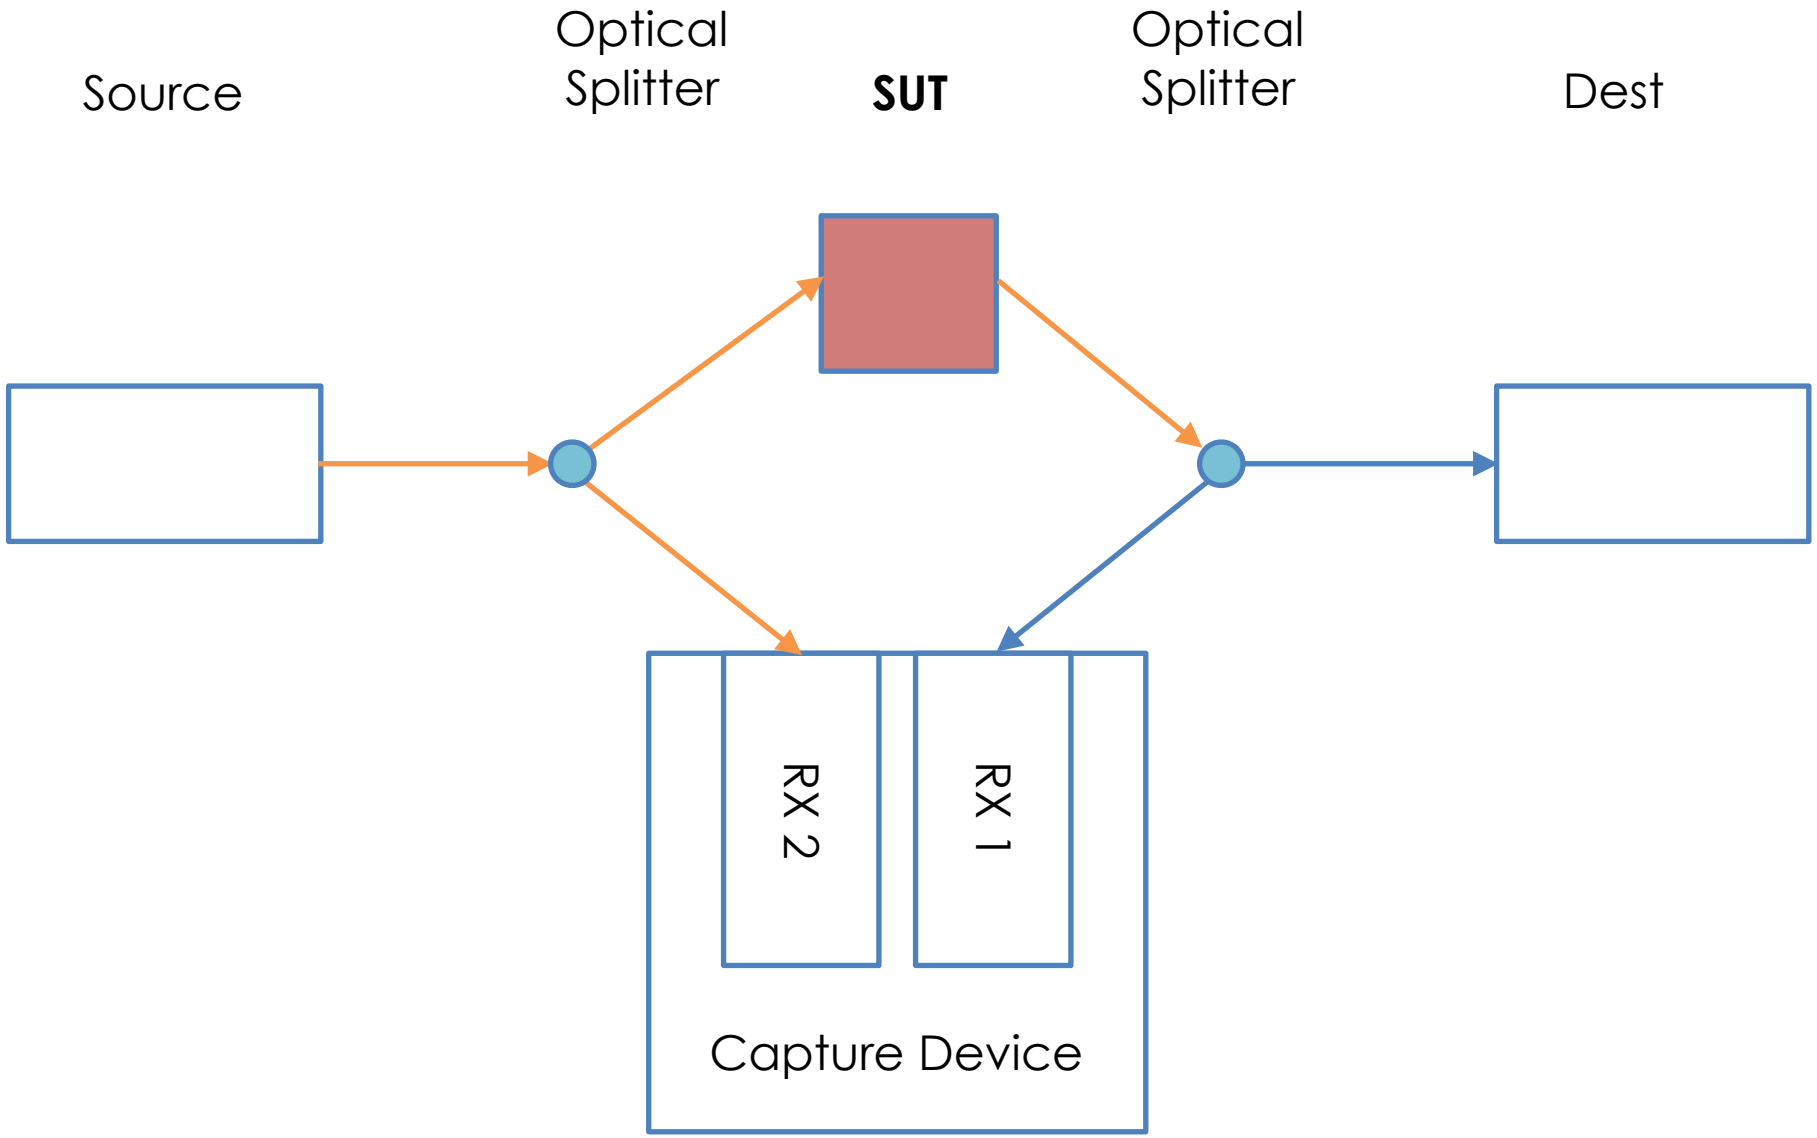
Source
Optical Splitter
SUT
Optical Splitter
Dest
RX 2
RX 1
Capture Device

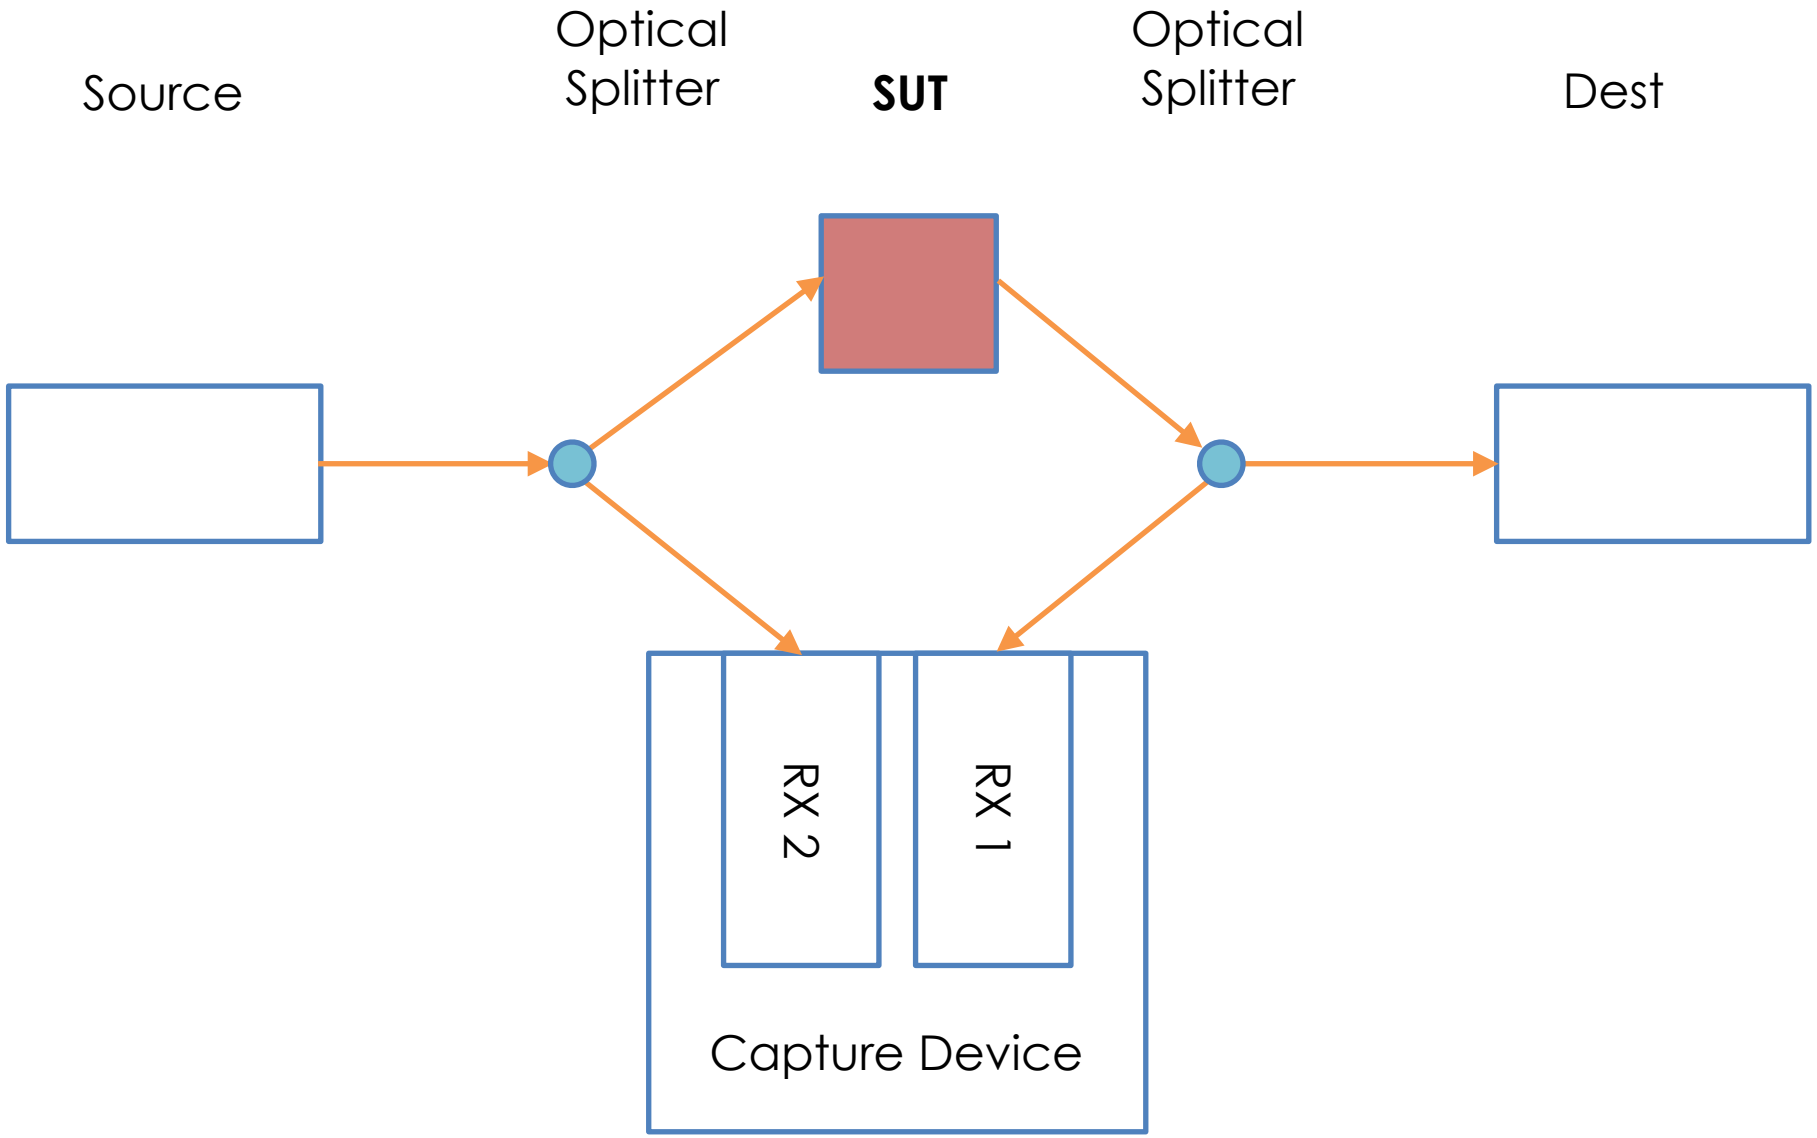
Source
Optical Splitter
SUT
Optical Splitter
Dest
RX 2
RX 1
Capture Device

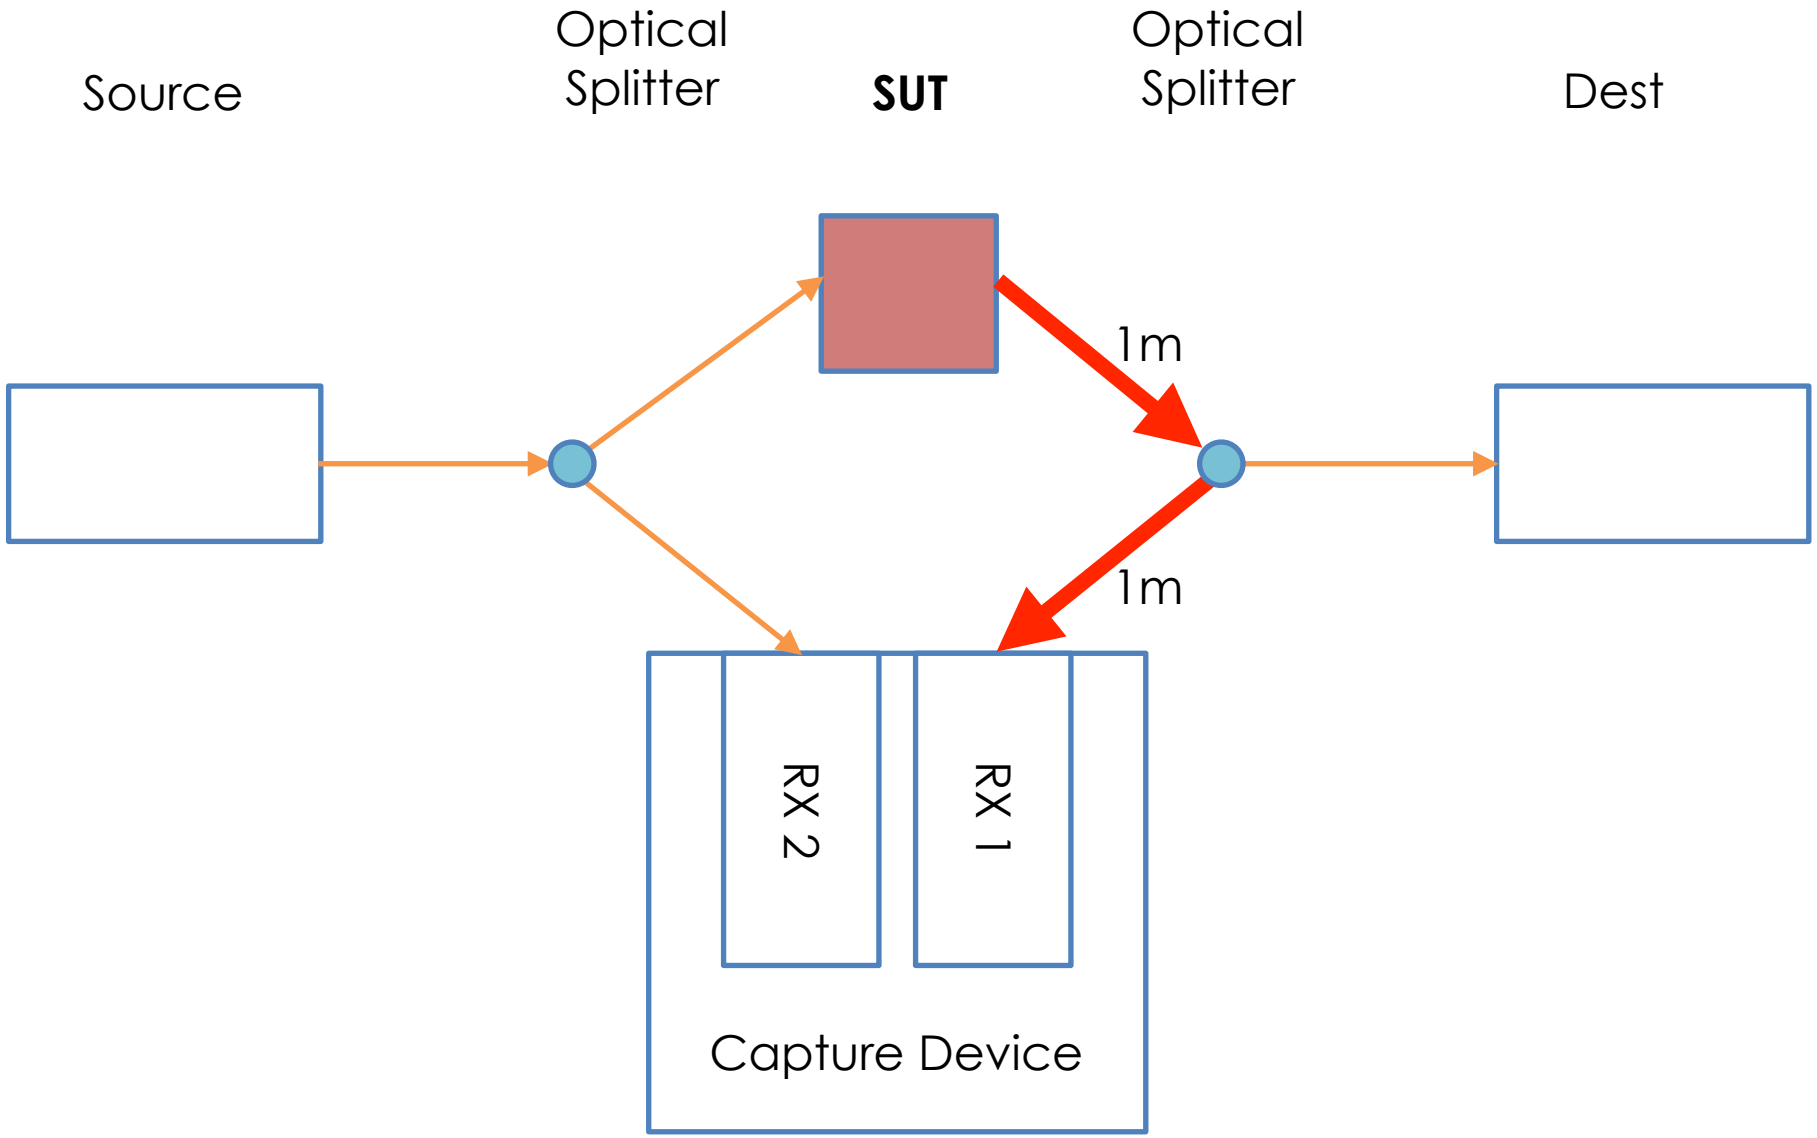
Source
Optical Splitter
SUT
Optical Splitter
Dest
1m
1m
RX 2
RX 1
Capture Device

How fast is the speed of light?

Exablaze Fibres

lengths 0.20m - 5.5m

Exablaze Fibres

lengths 0.20m - 5.5m

increments 0.25m - 0.5m

Exablaze Fibres

lengths 0.20m - 5.5m

increments 0.25m - 0.5m

total 18x

Speed of Light in Fibre

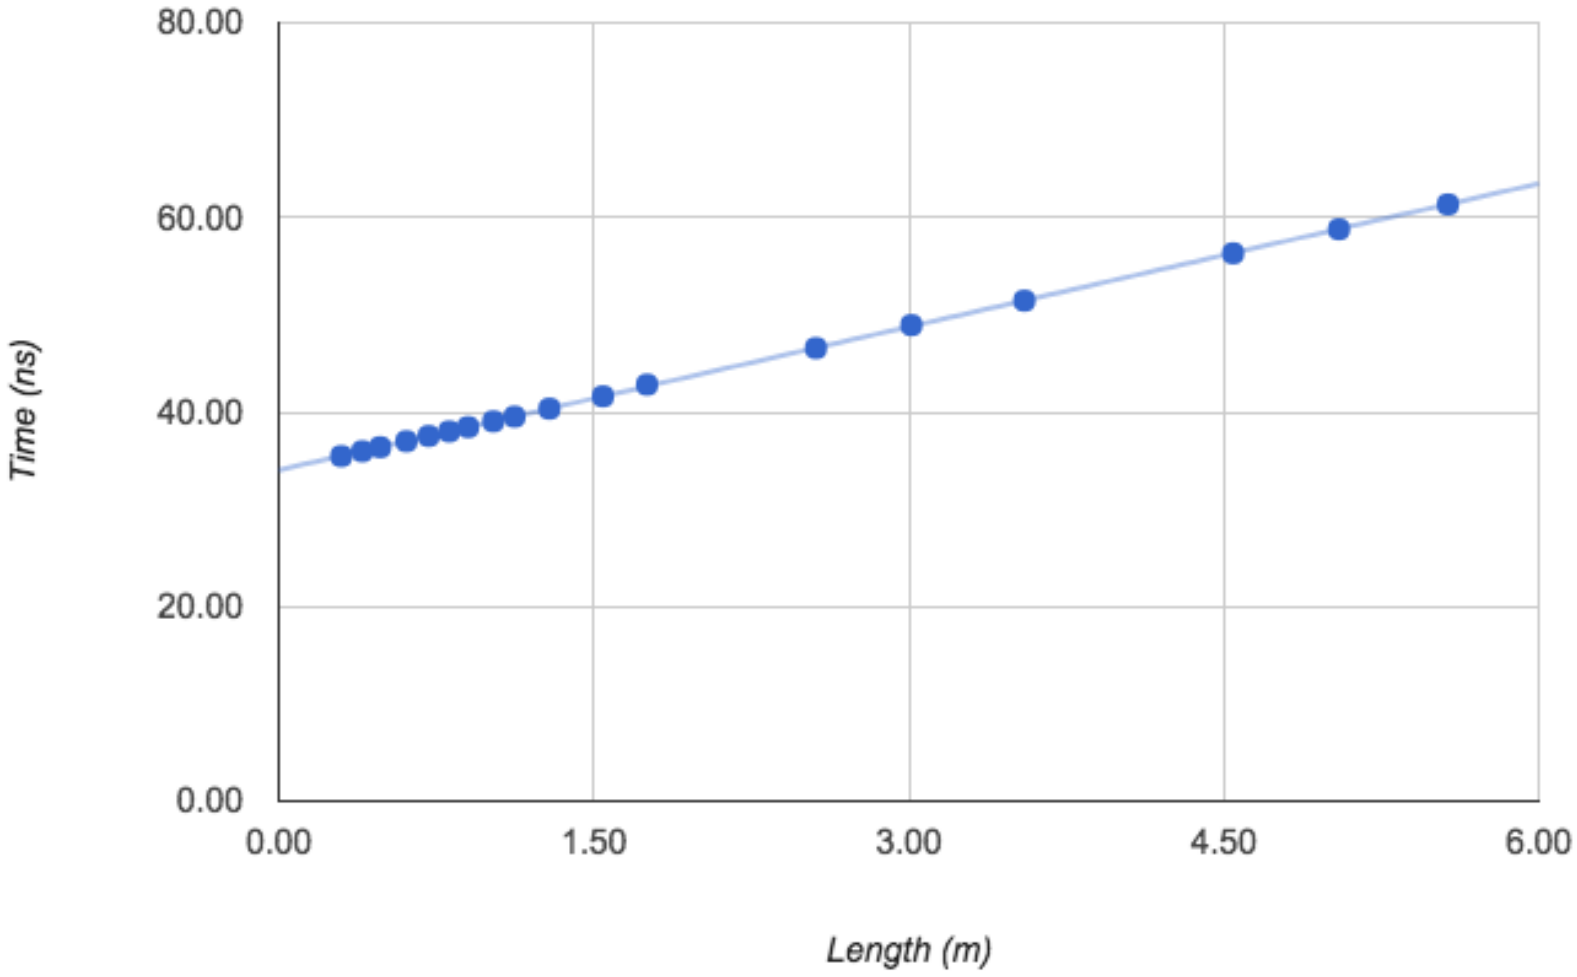

Speed of Light in Fibre

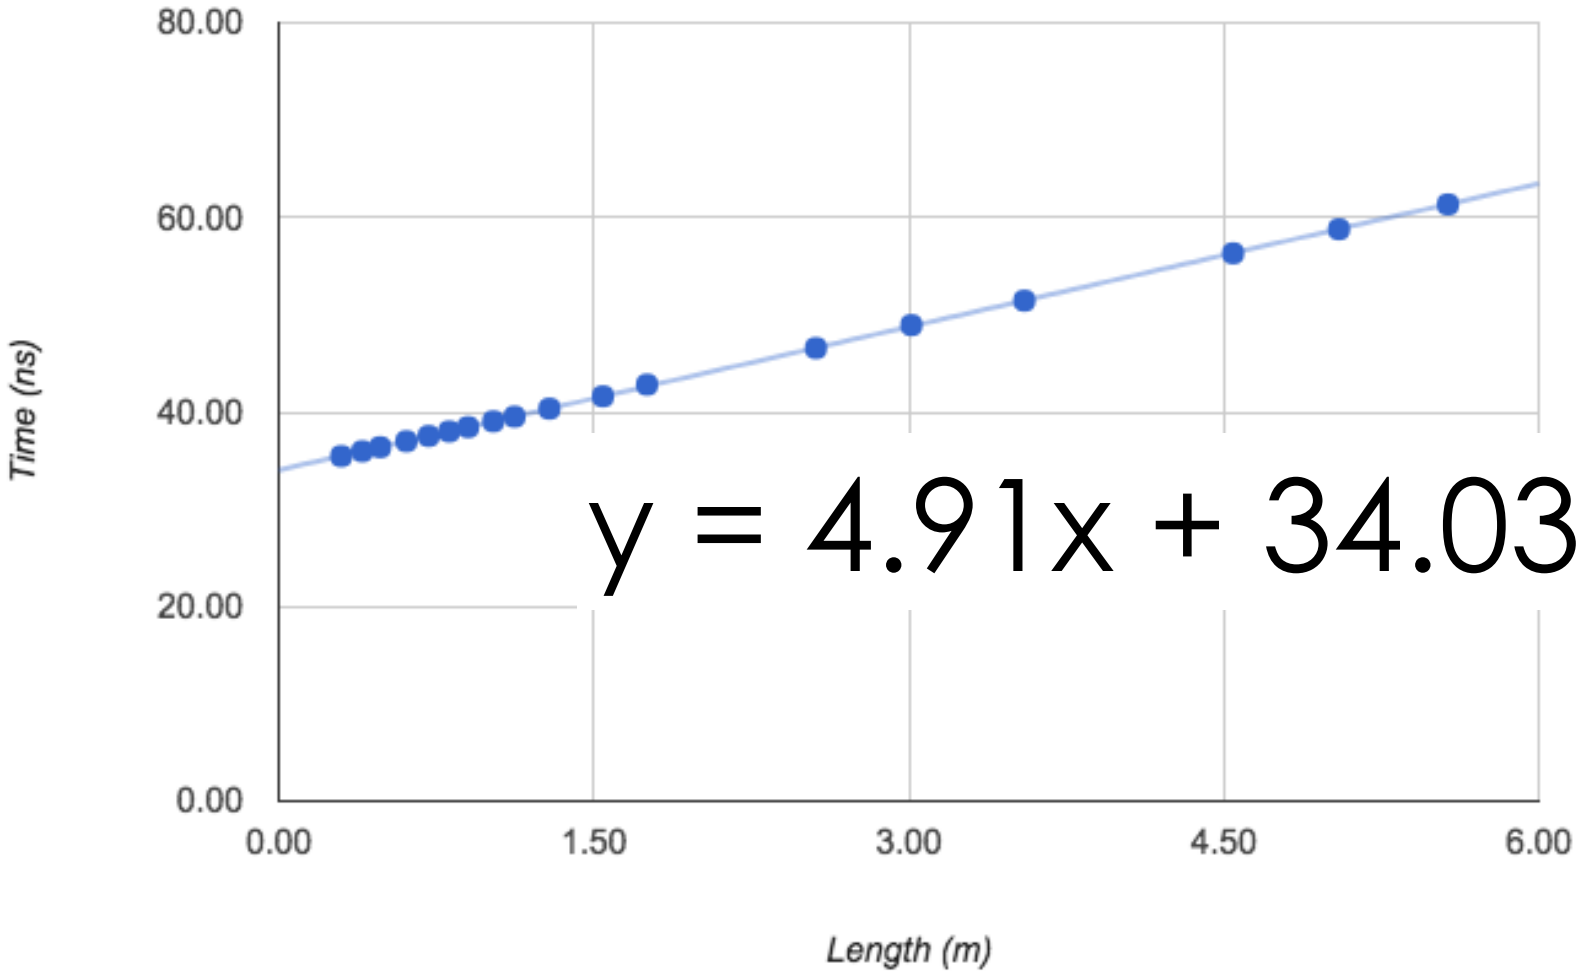

Speed of Light in Fibre

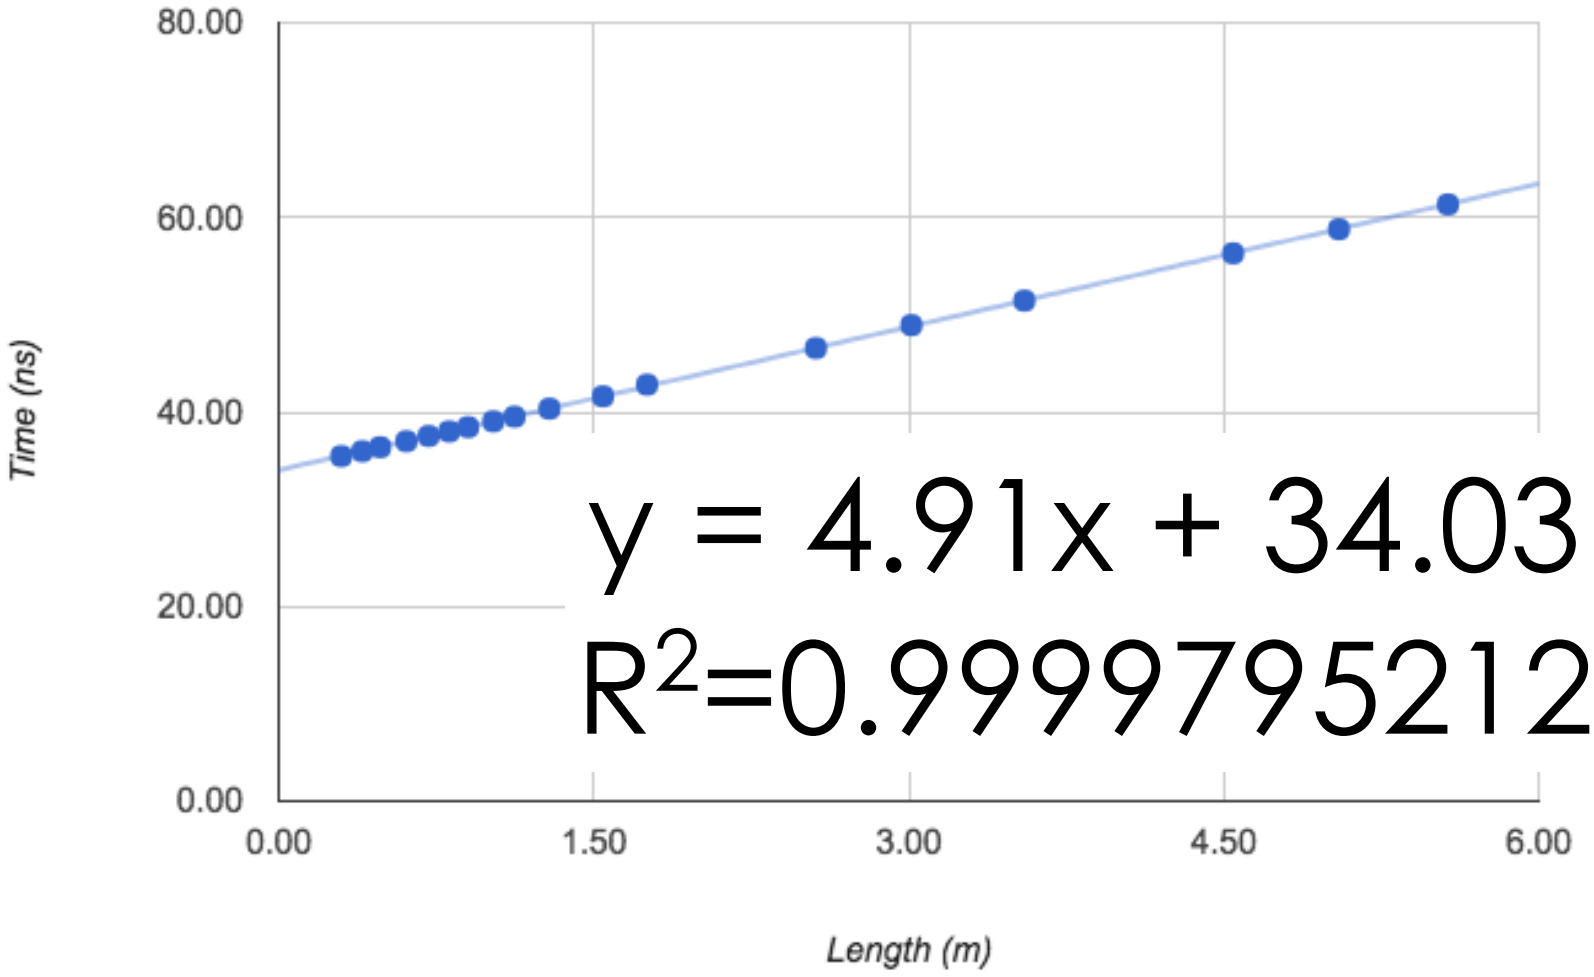

Speed of Light in Fibre

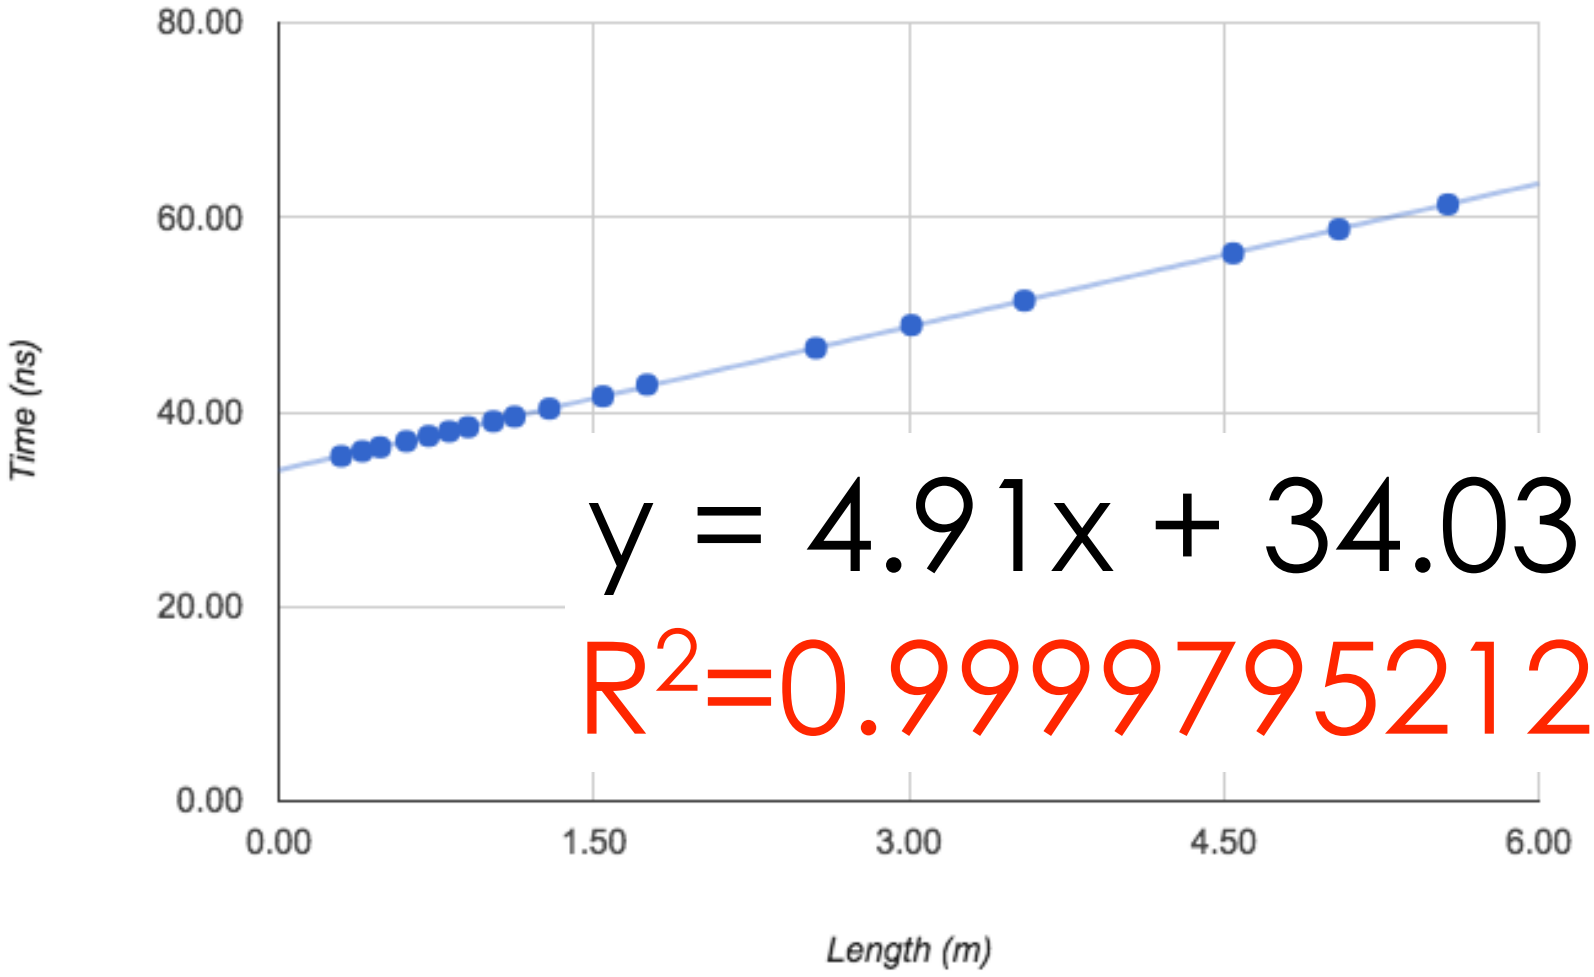

Speed of Light in Fibre

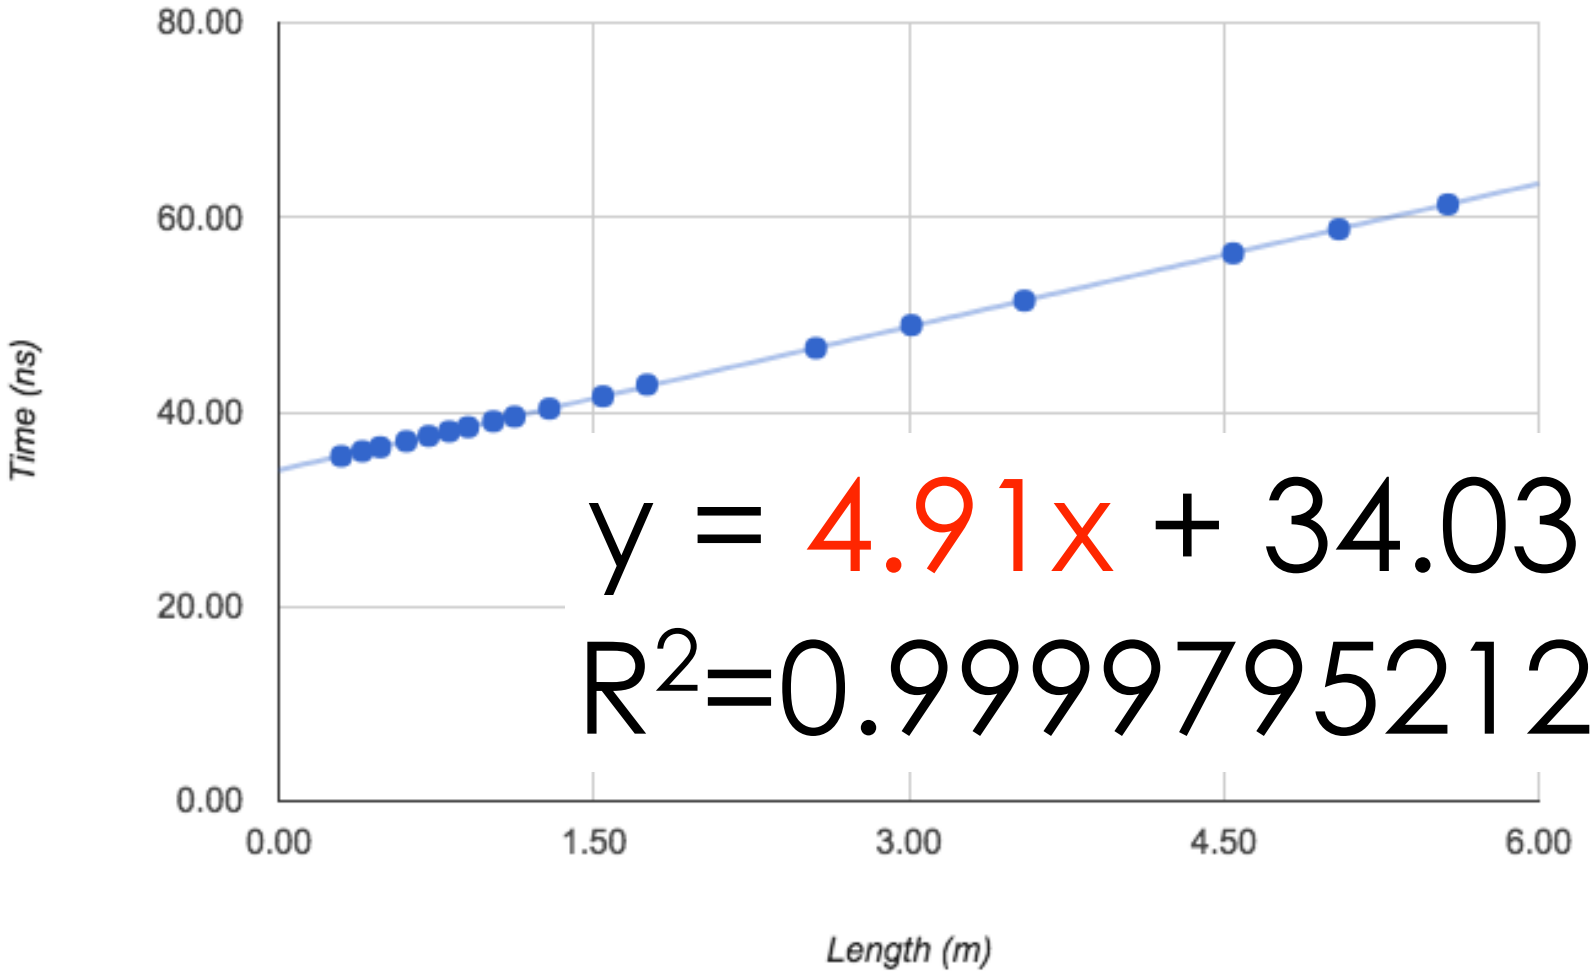

Tip #3

C in fibre 4.91ns/m

Speed of Light in Fibre

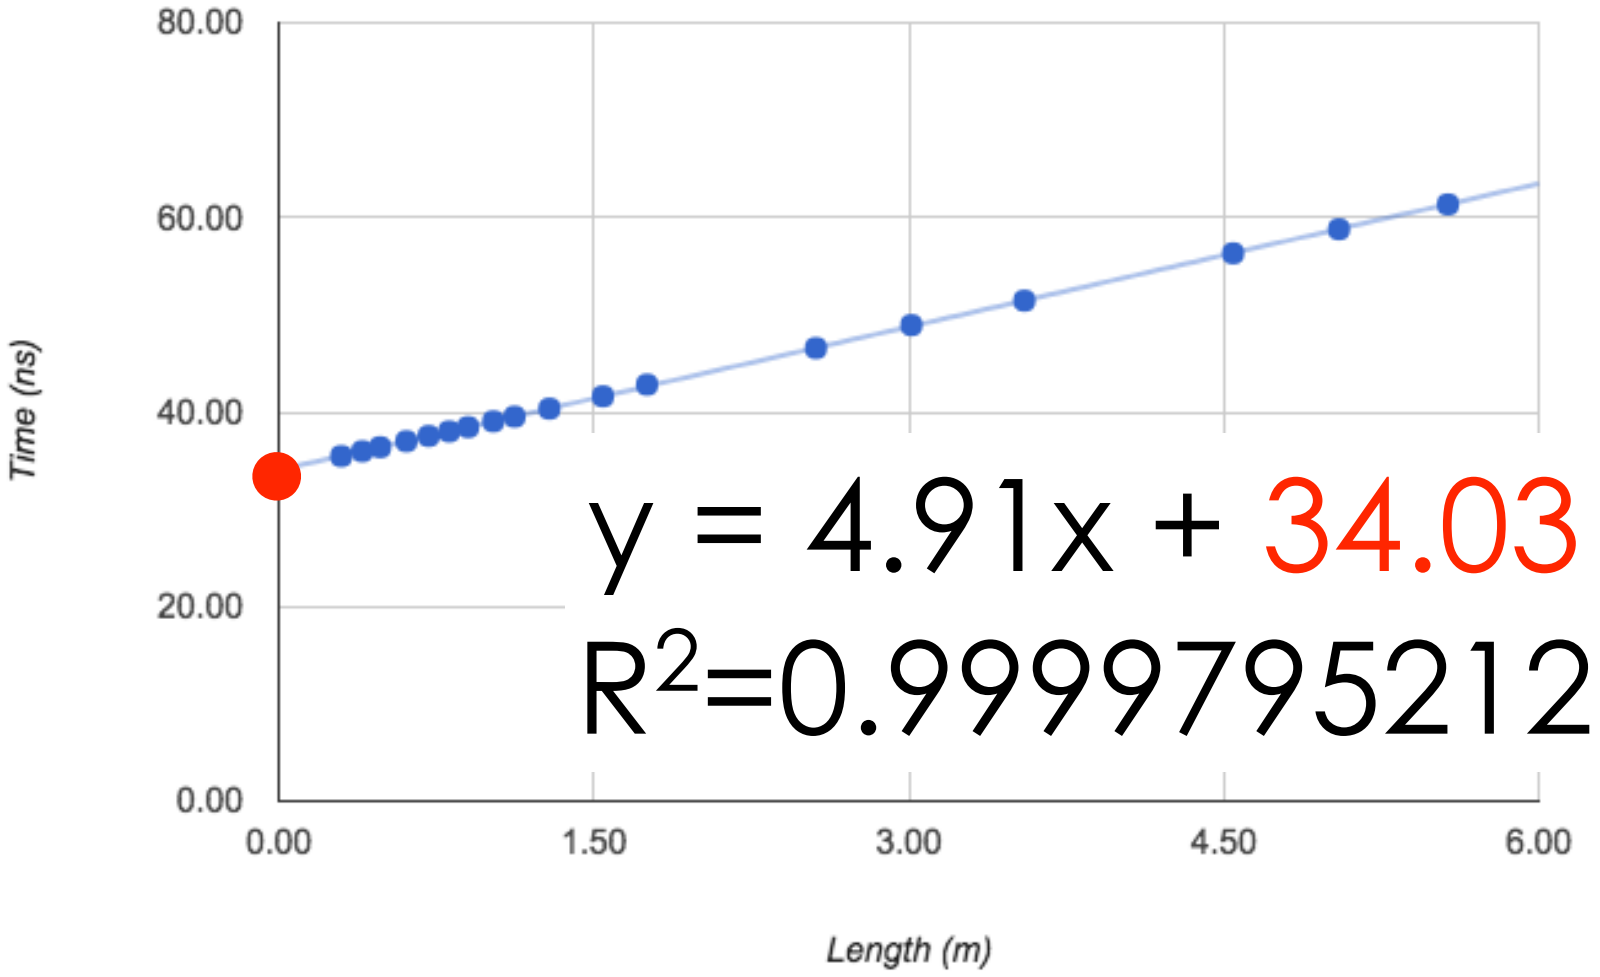

Copper?

Exablaze DA Cables

lengths 0.25m - 12m

Exablaze DA Cables

lengths 0.25m - 12m

increments 0.25m - 0.5m

Exablaze DA Cables

lengths 0.25m - 12m

increments 0.25m - 0.5m

gauges AWG24/30

Exablaze DA Cables

lengths 0.25m - 12m

increments 0.25m - 0.5m

gauges AWG24/30

total 25x

Speed of Light Copper (AWG30)

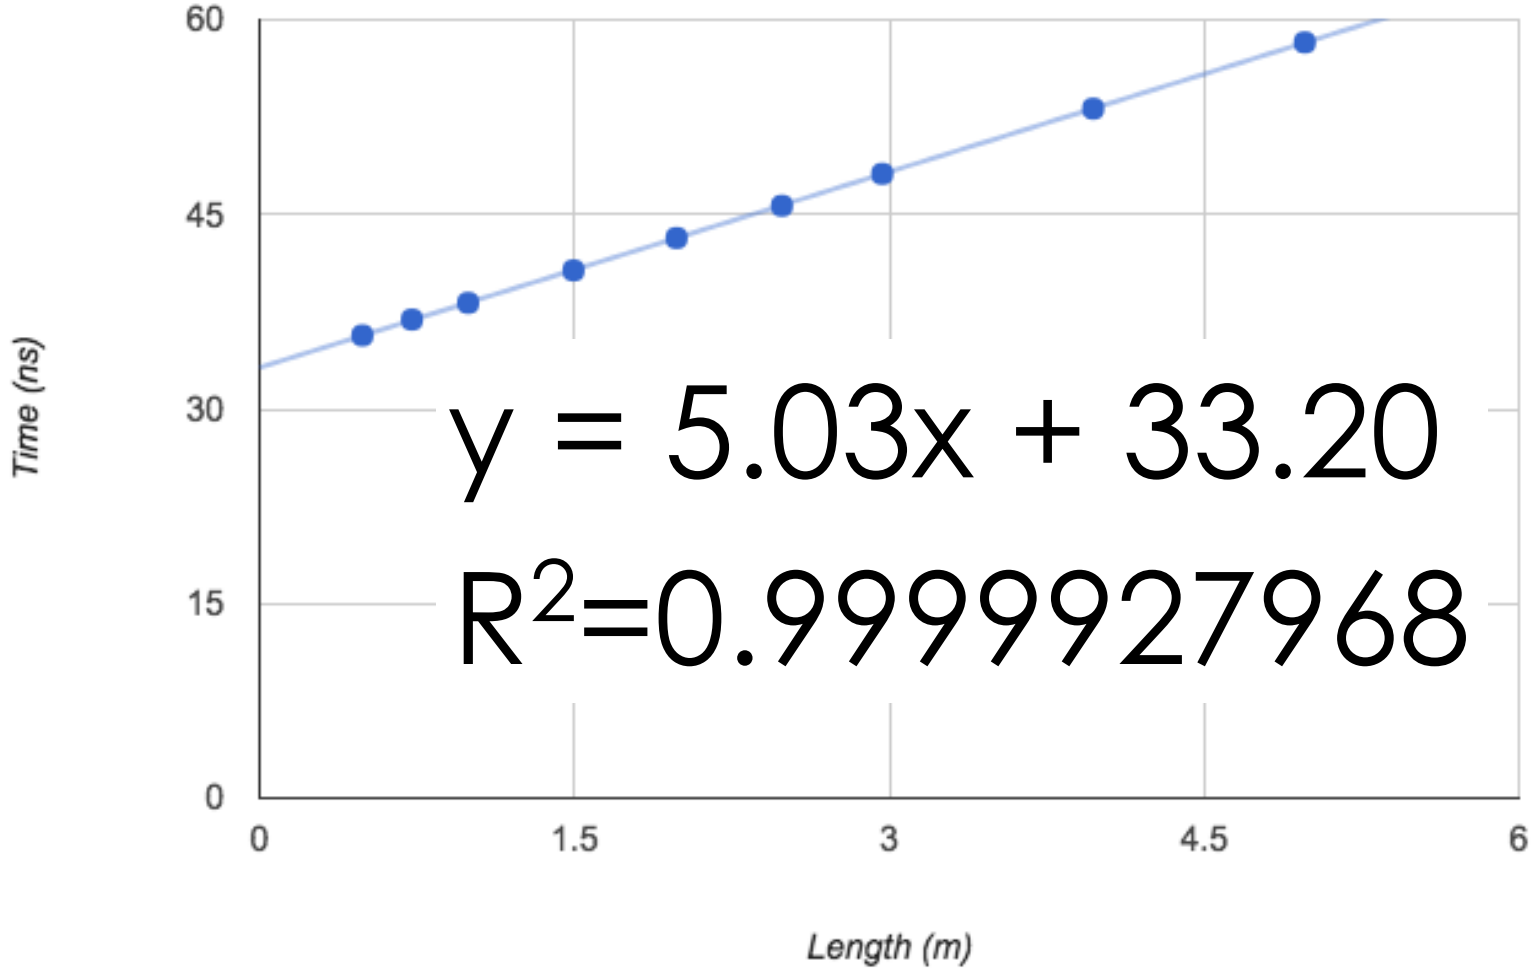

Fibre	4.91ns/m	34.03ns
AWG30	5.03ns/m	33.20ns

Fibre	4.91ns/m	34.03ns
AWG30	5.03ns/m	33.20ns

SFP cost = 830ps

Fibre	4.91ns/m	34.03ns
AWG30	5.03ns/m	33.20ns

Copper slower!

Gotcha #5

Copper isn't faster than fibre.

Speed of Light Copper (AWG24)

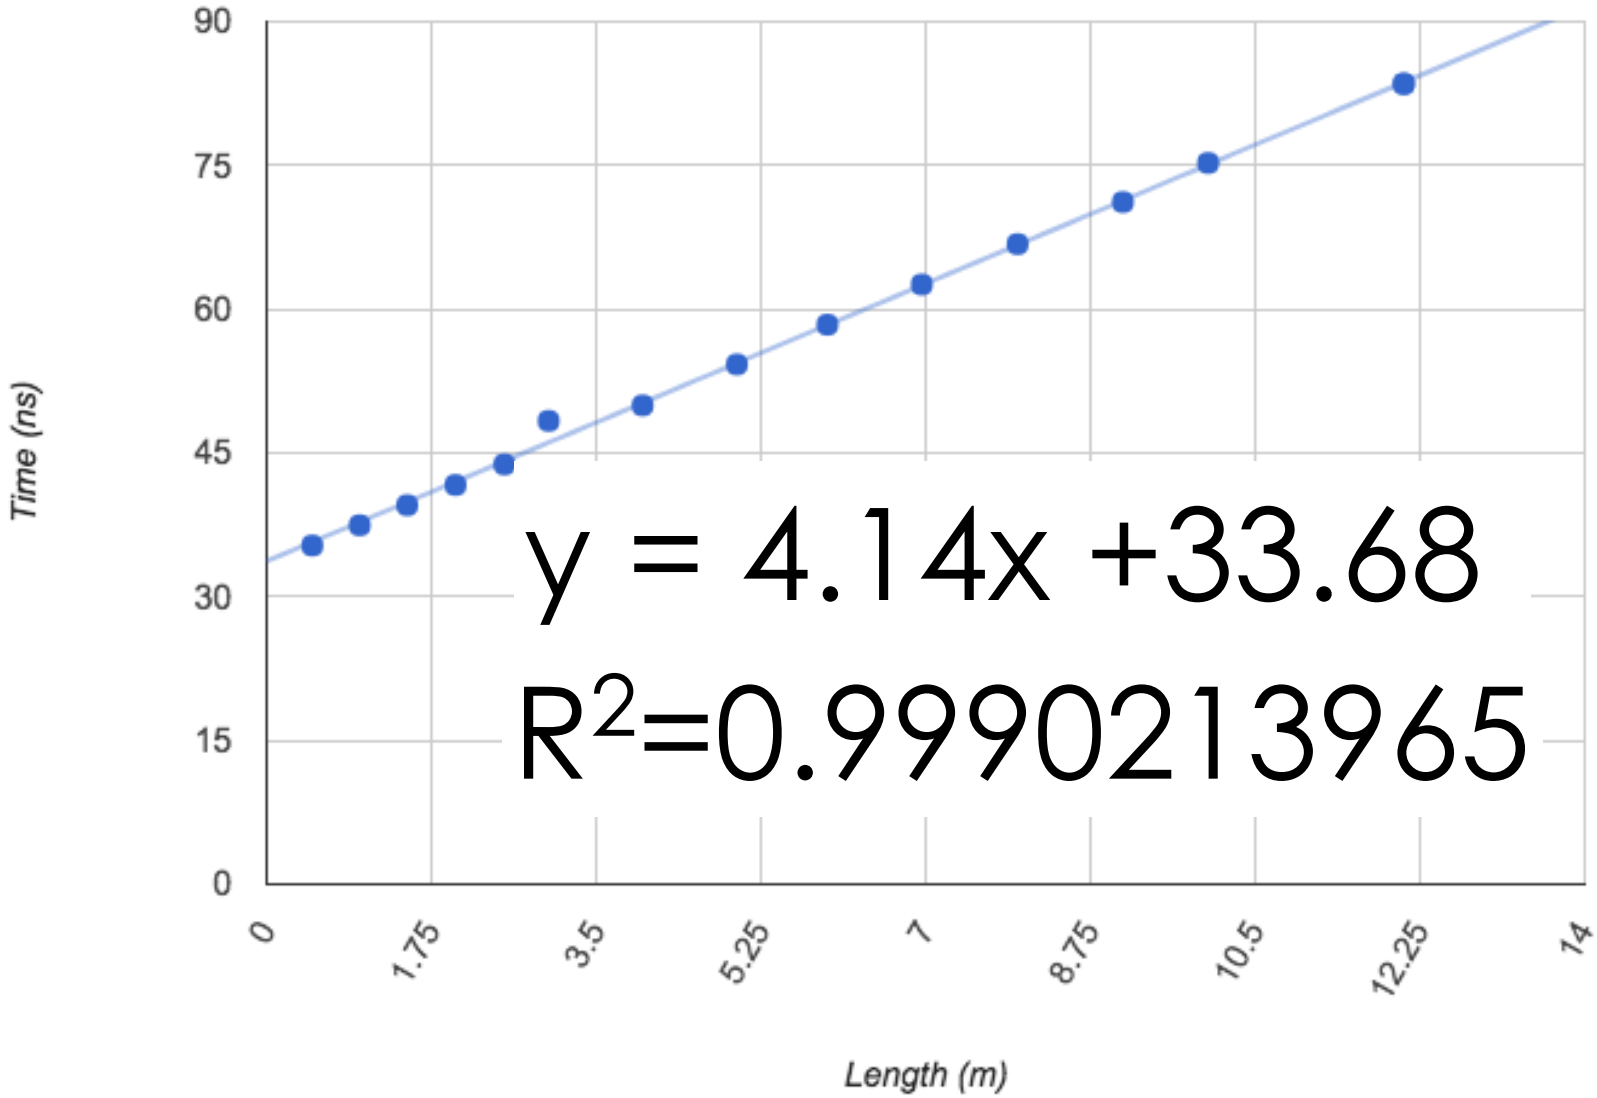

Speed of Light Copper (AWG24)

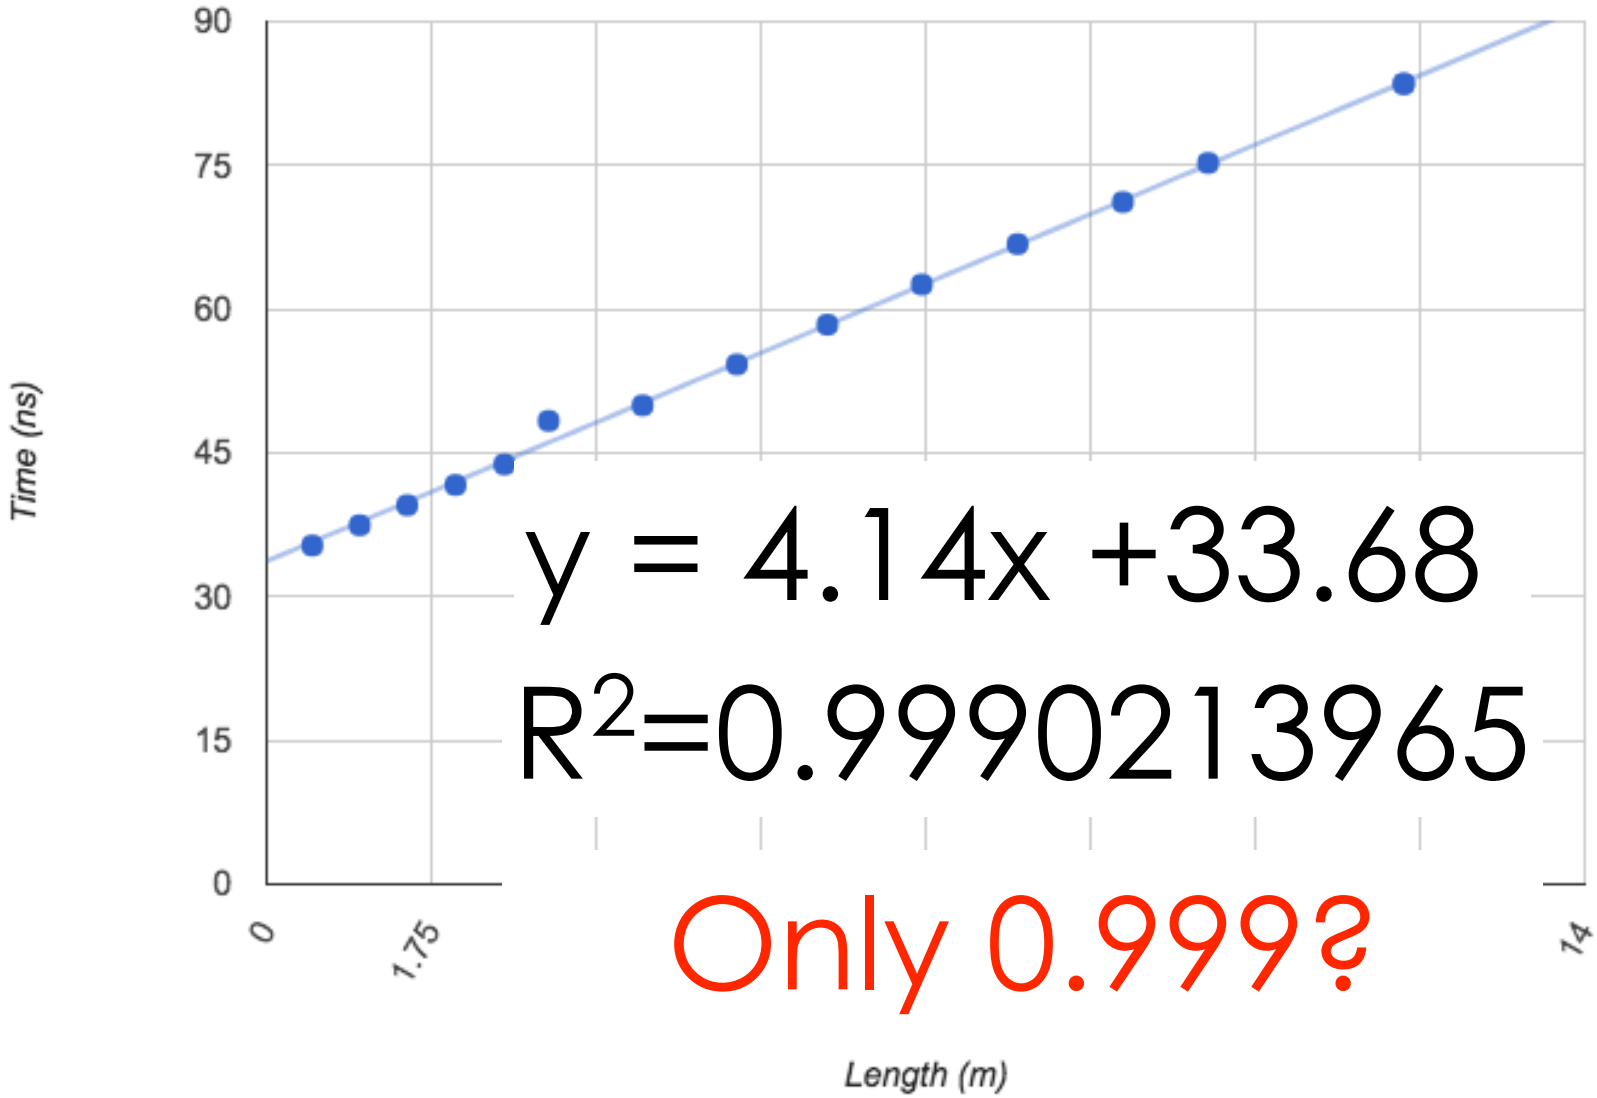

$y = 4.14x + 33.68$

$R^2 = 0.9990213965$

Only 0.999?

Speed of Light Copper (AWG24)

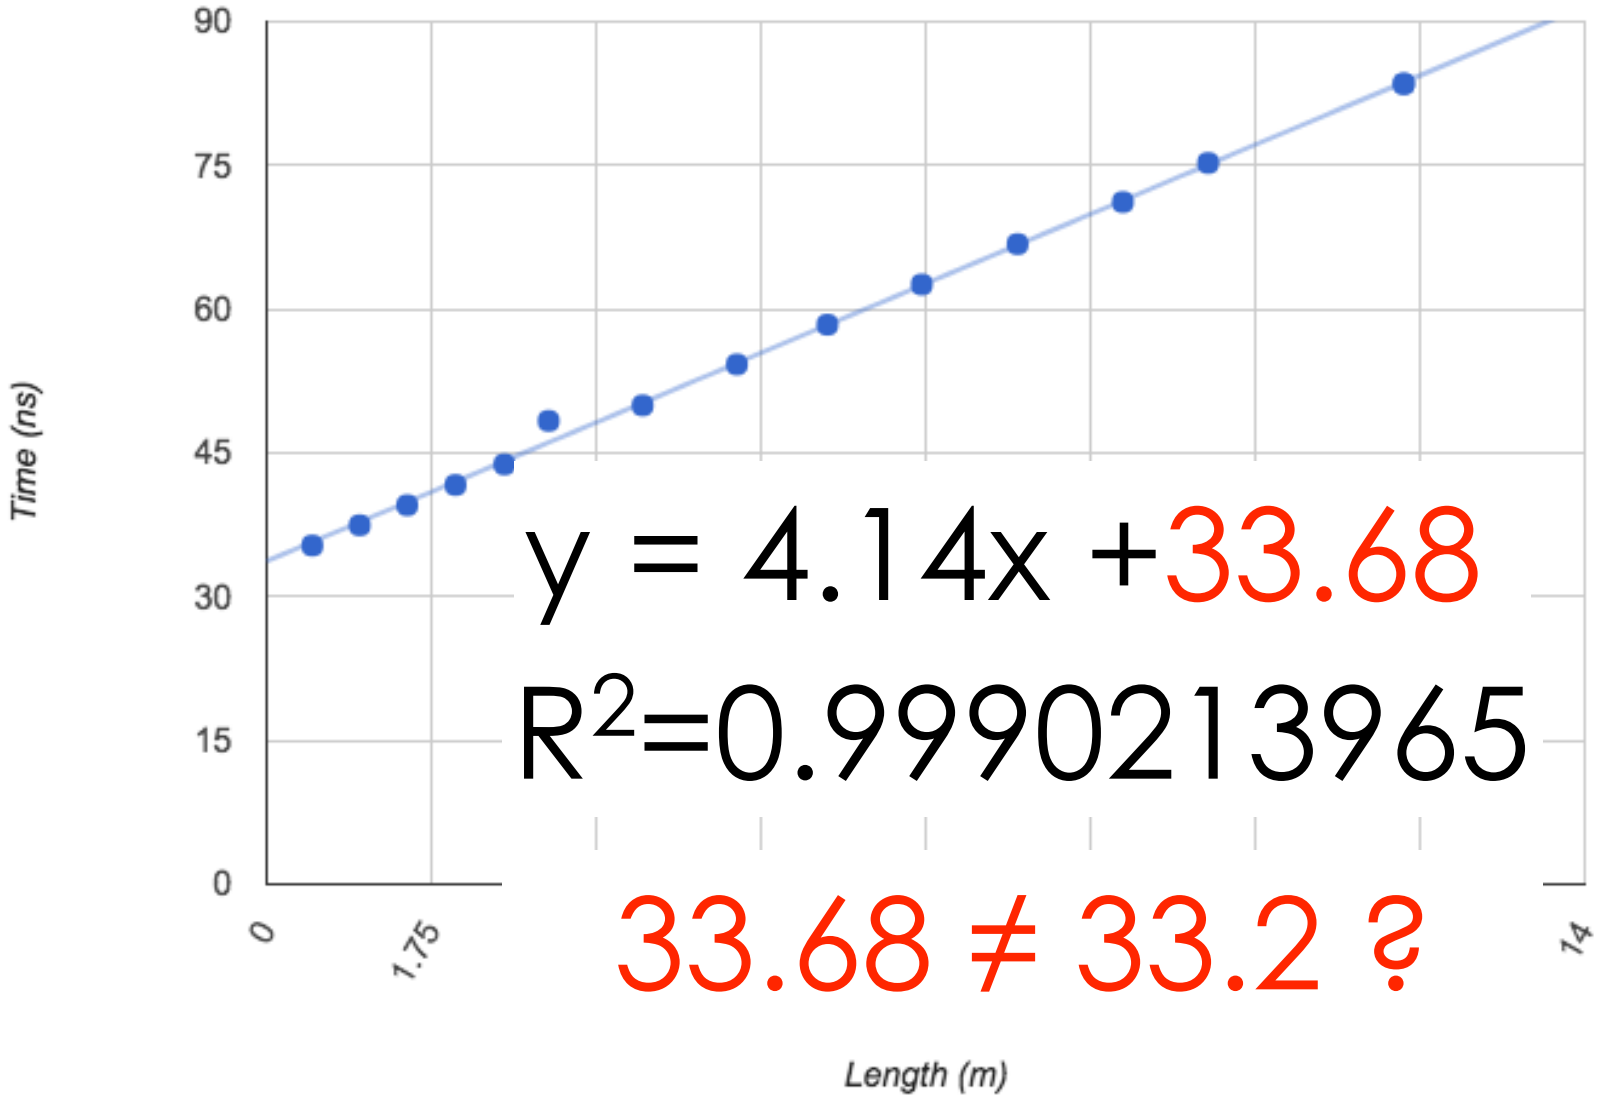

$y = 4.14x + 33.68$

$R^2 = 0.9990213965$

$33.68 \neq 33.2$?

Speed of Light Copper (AWG24)

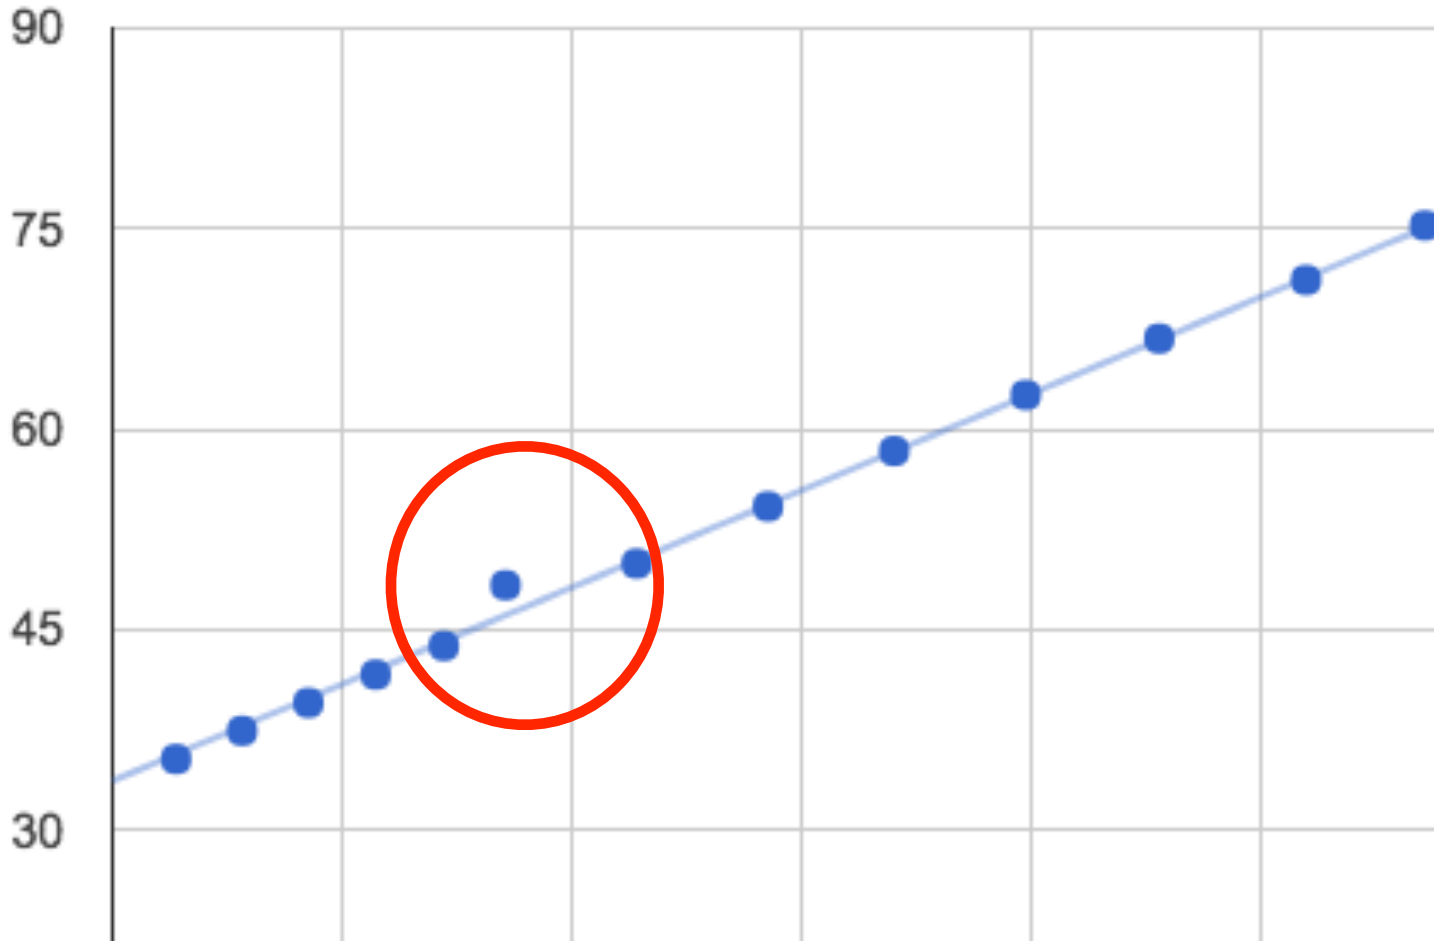

Speed of Light Copper (AWG24)

Length (m)	Time (ns)	ns/m
…	…	…
2.01	41.66	4.17
2.53	43.81	4.16
3.00	**48.35**	**5.03**
3.99	49.97	4.17
5.00	54.24	4.19
..	…	…

Active cable?

Active cable?

Active cable?

Length wrong?

Active cable? 

Length wrong? 

Active cable? 

Length wrong? 

Data bad?

Active cable? 

Length wrong? 

Data bad? 


BELDEN L-Skew 24A1B01-17 E357312-S


BELDEN L-Skew 24A1B01-17 E357312-S

Speed of Light Copper (AWG24)

Speed of Light Copper (AWG30)

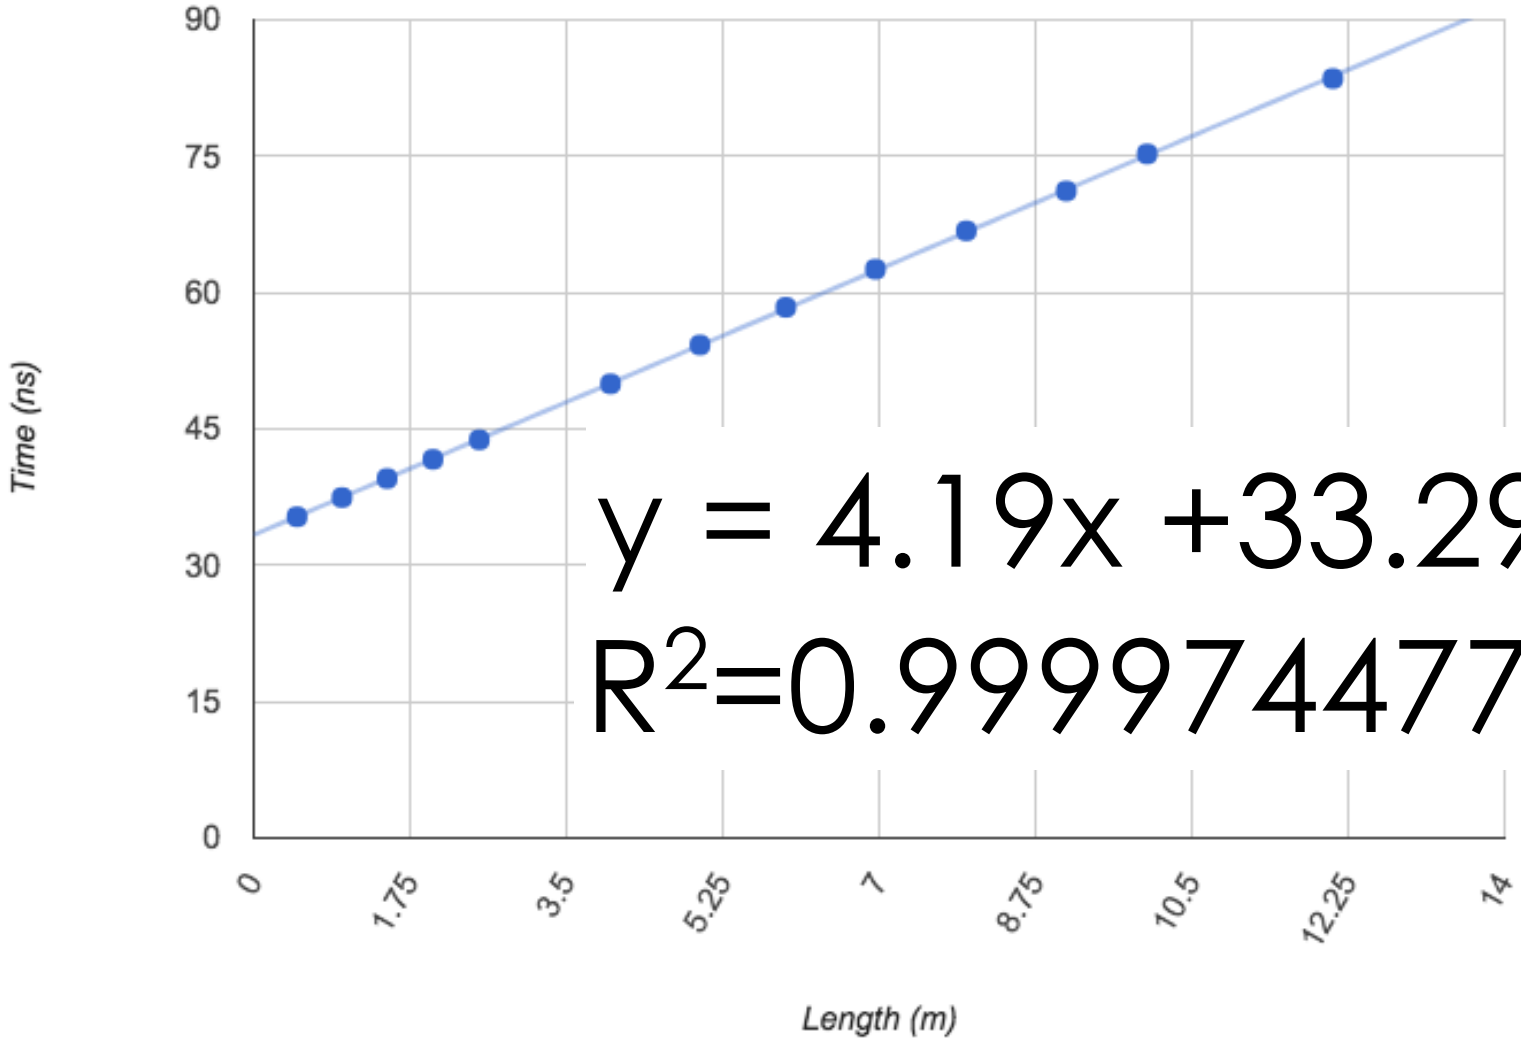

$y = 4.19x + 33.29$

$R^2 = 0.9999744775$

AWG24A	4.19ns/m
Fibre	4.91ns/m
AWG30	5.03ns/m
AWG24B	5.03ns/m

Remember Gotcha #3?

Copper isn't faster than fibre.

Gotcha #3B

Copper isn't always faster than fibre.

Remember Gotcha #3B?

Copper isn't always faster than fibre.

Tip #4?

Copper isn't always faster than fibre

Unless you get your copper from Exablaze…

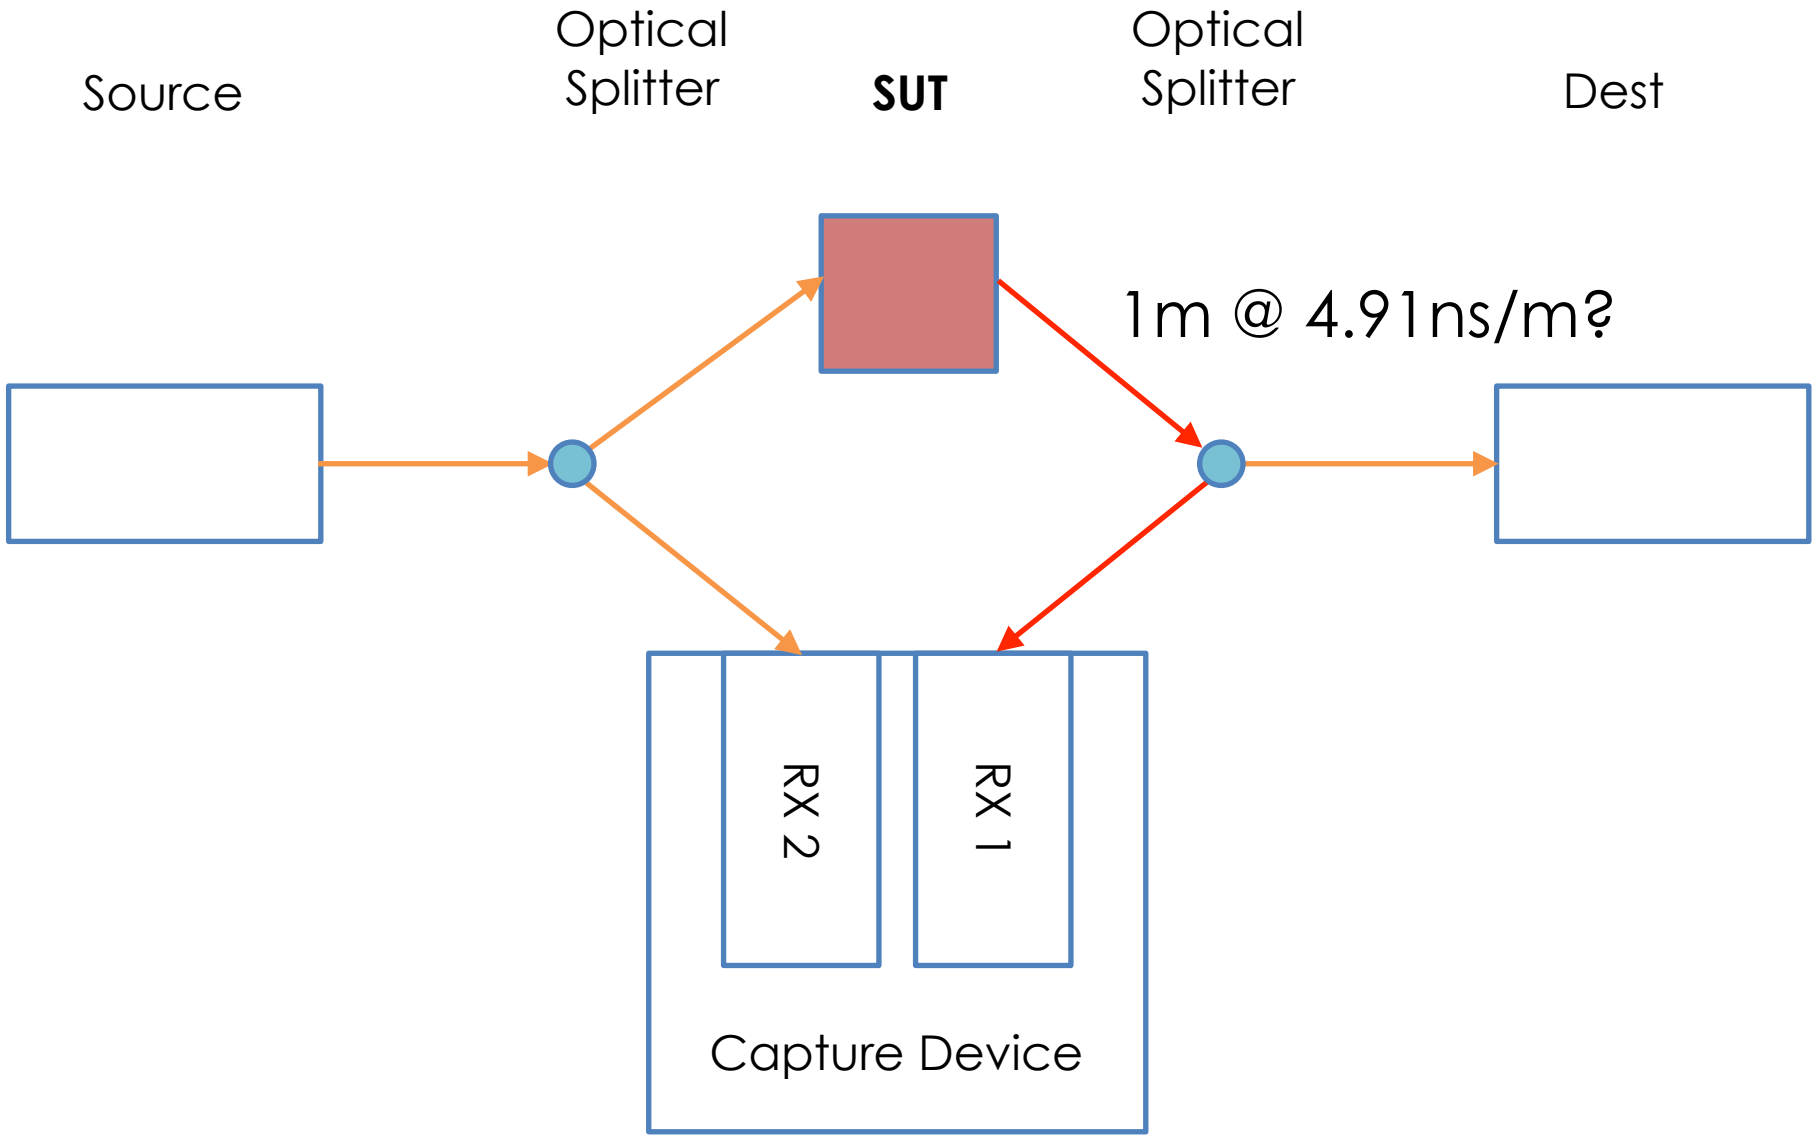
Source
Optical Splitter
SUT
Optical Splitter
Dest
1m @ 4.91ns/m?
RX 2
RX 1
Capture Device

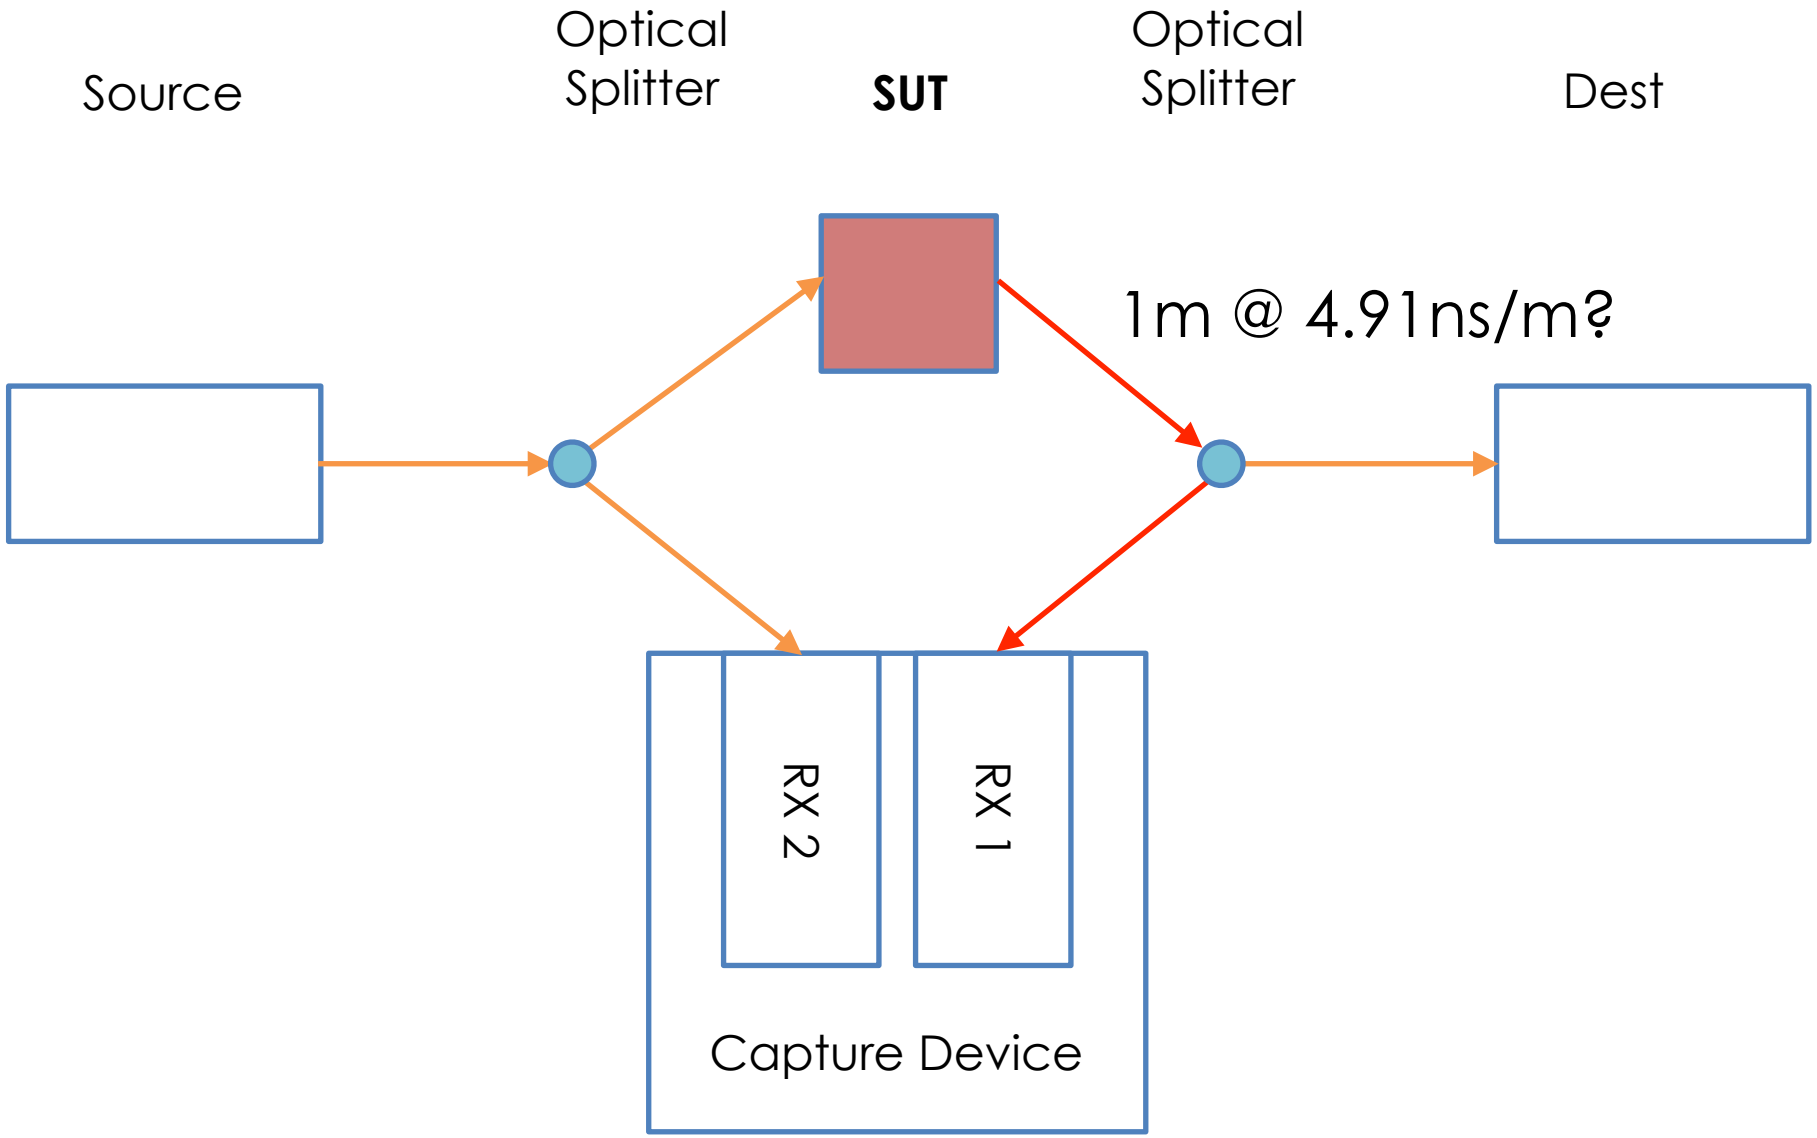
Source
Optical Splitter
SUT
Optical Splitter
Dest
1m @ 4.91ns/m?
RX 2
RX 1
Capture Device

1m @ 4.91ns/m = 4.91ns

1.0m @ 4.91ns/m = 4.91ns
1.1m @ 4.91ns/m = 5.40ns

1.0m @ 4.91ns/m = 4.91ns

1.1m @ 4.91ns/m = 5.40ns

500ps!

1.00m @ 4.91ns/m = 4.91ns
1.10m @ 4.91ns/m = 5.40ns
1.11m @ 4.91ns/m = 5.45ns

550ps!

Gotcha #4

1m isn't 1.00m

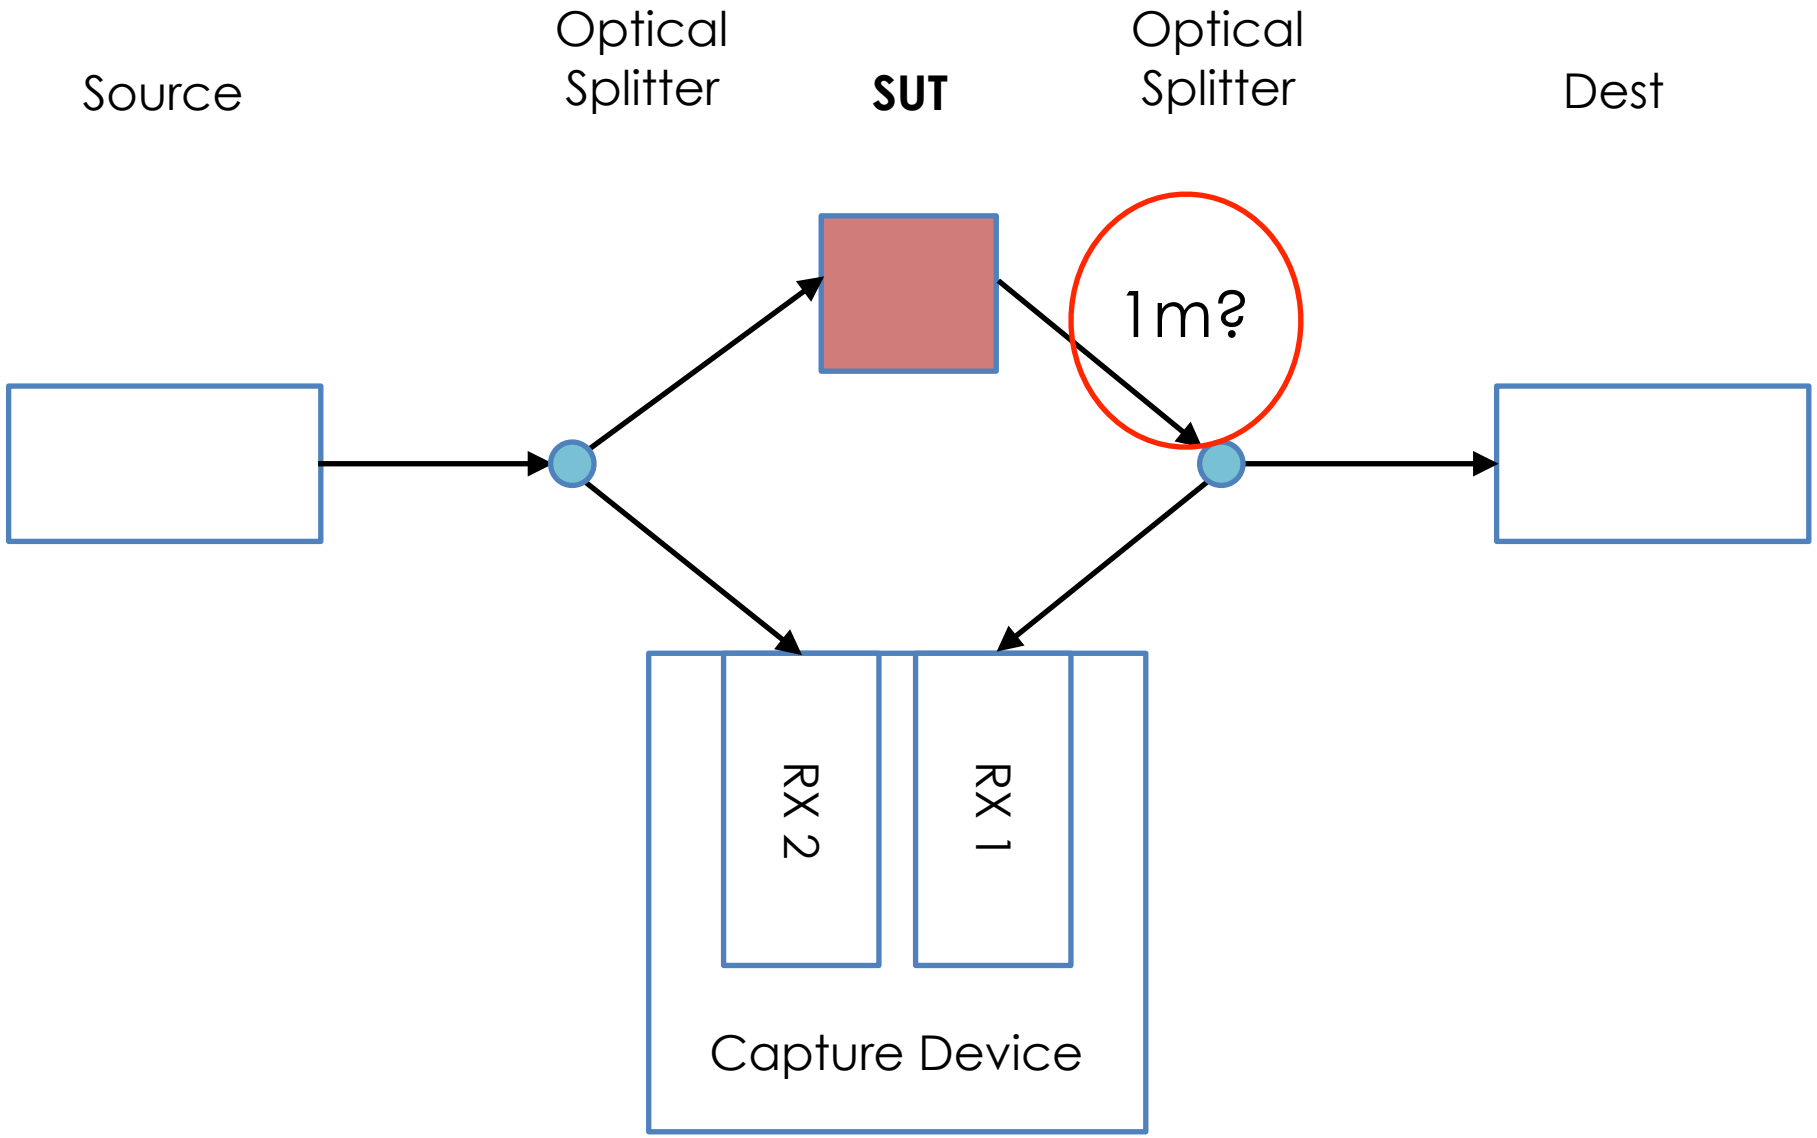
Source
Optical Splitter
SUT
Optical Splitter
Dest
1m?
RX 2
RX 1
Capture Device

Source **SUT** Dest

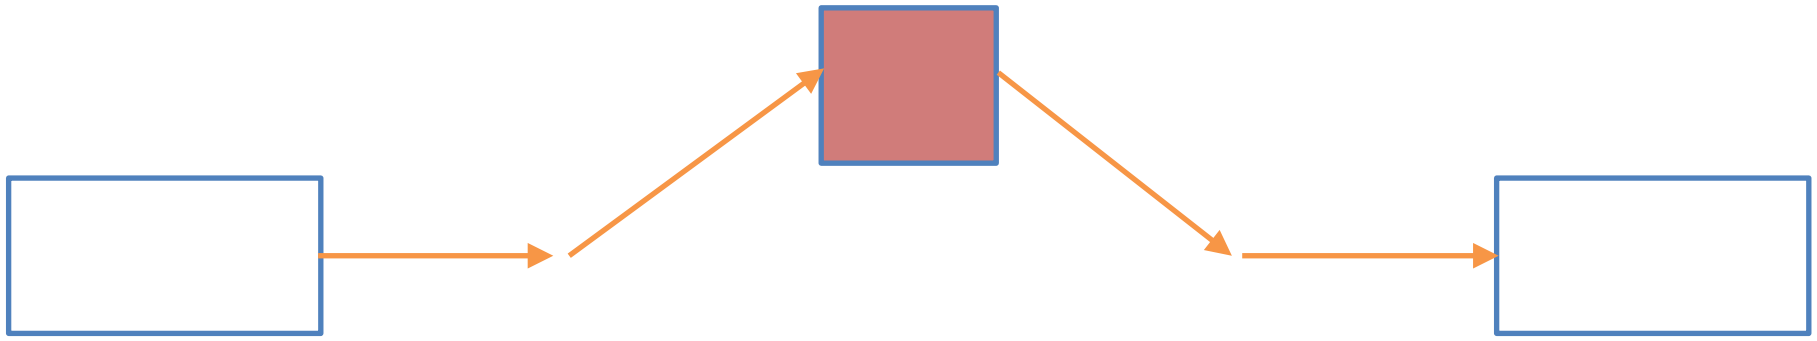

Source **SUT** Dest

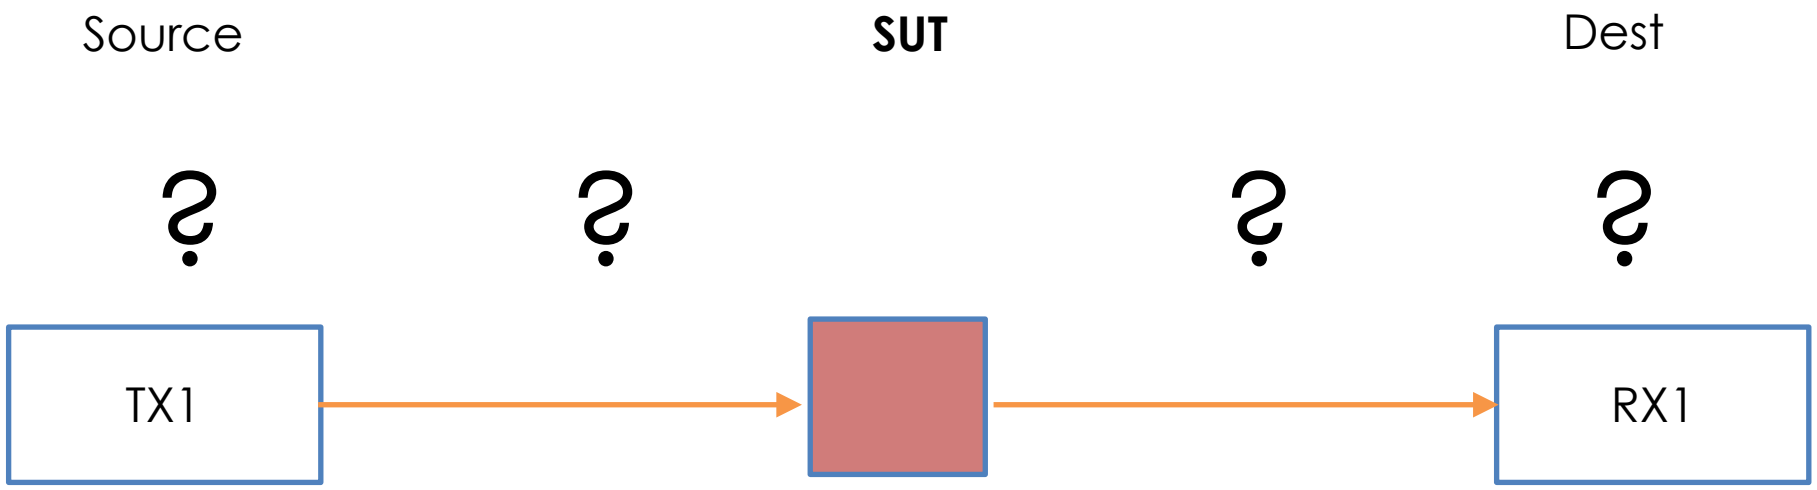
Source
SUT
Dest
?
?
?
?
TX1
RX1

Source **SUT** Dest

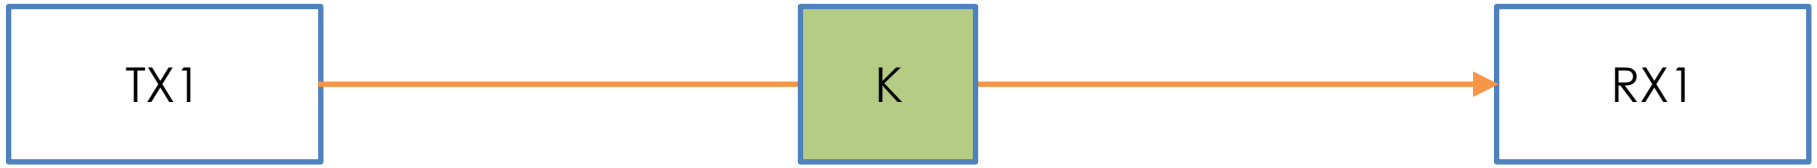

Source **SUT** Dest

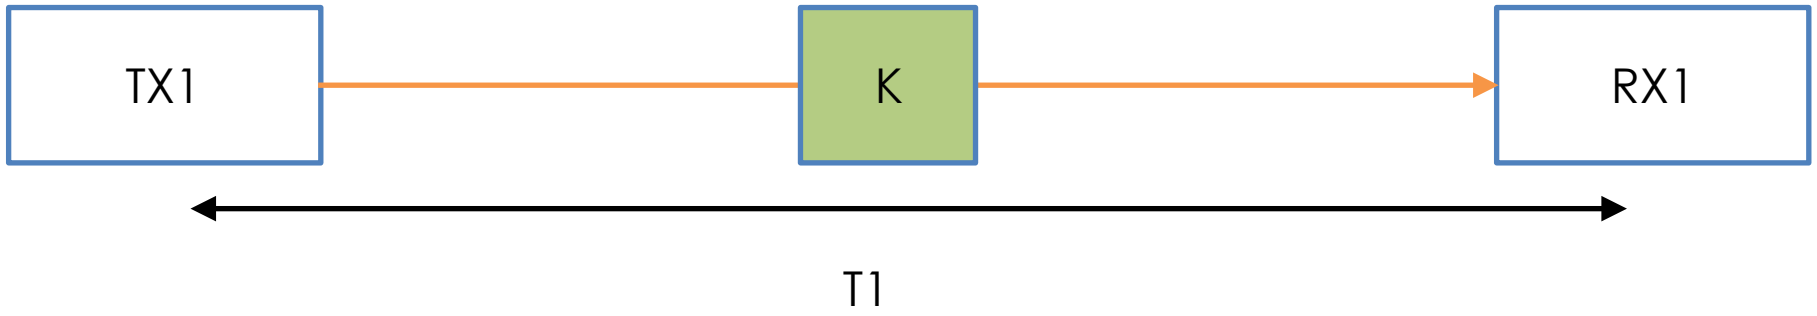

Source **SUT** Dest

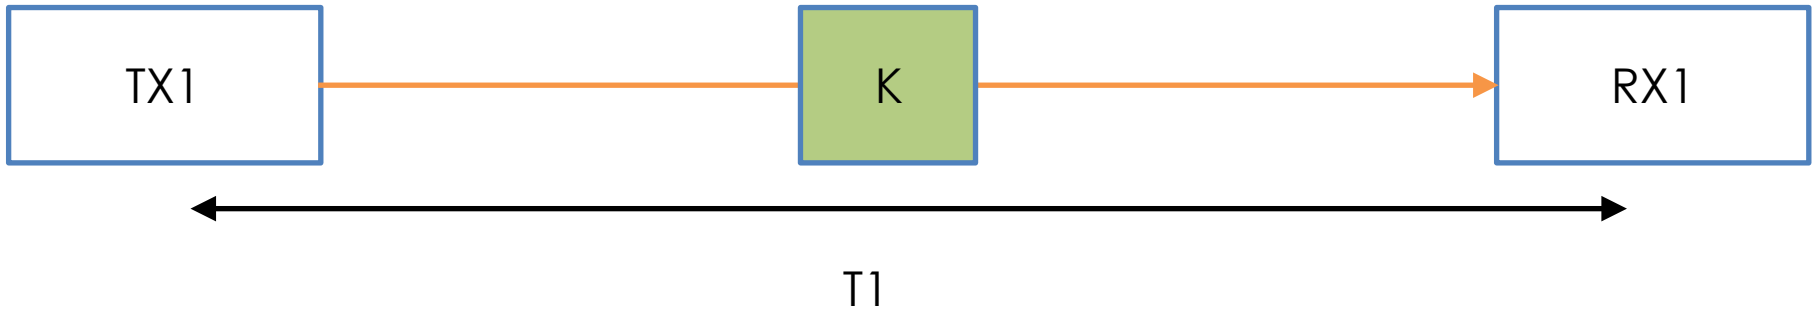

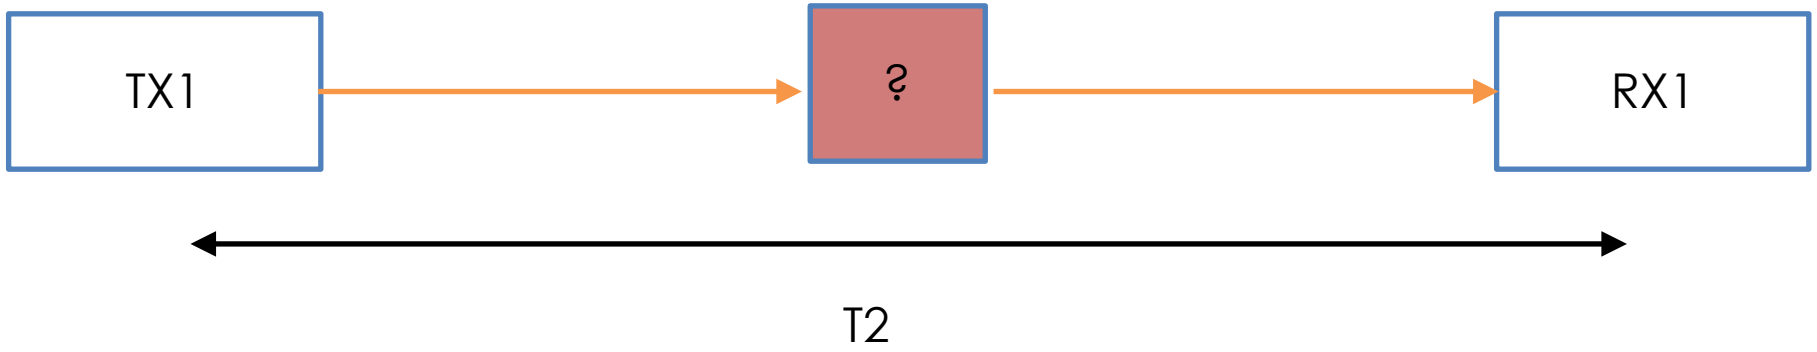

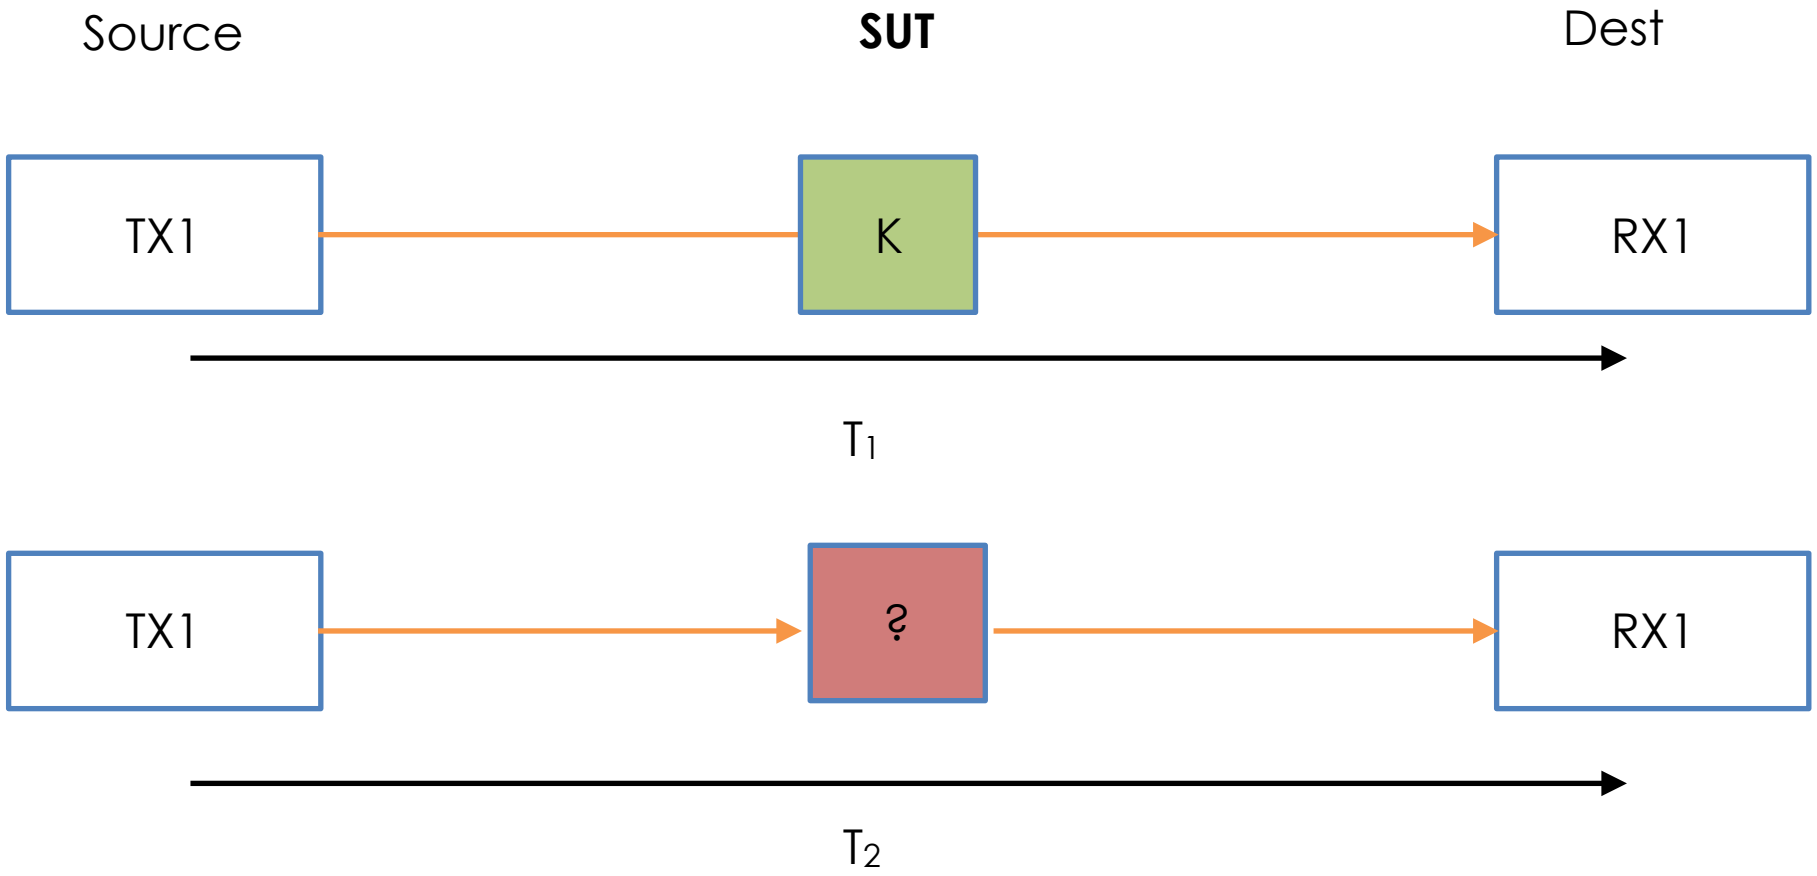

$$\text{Delay} = T_2 - T_1 - K$$

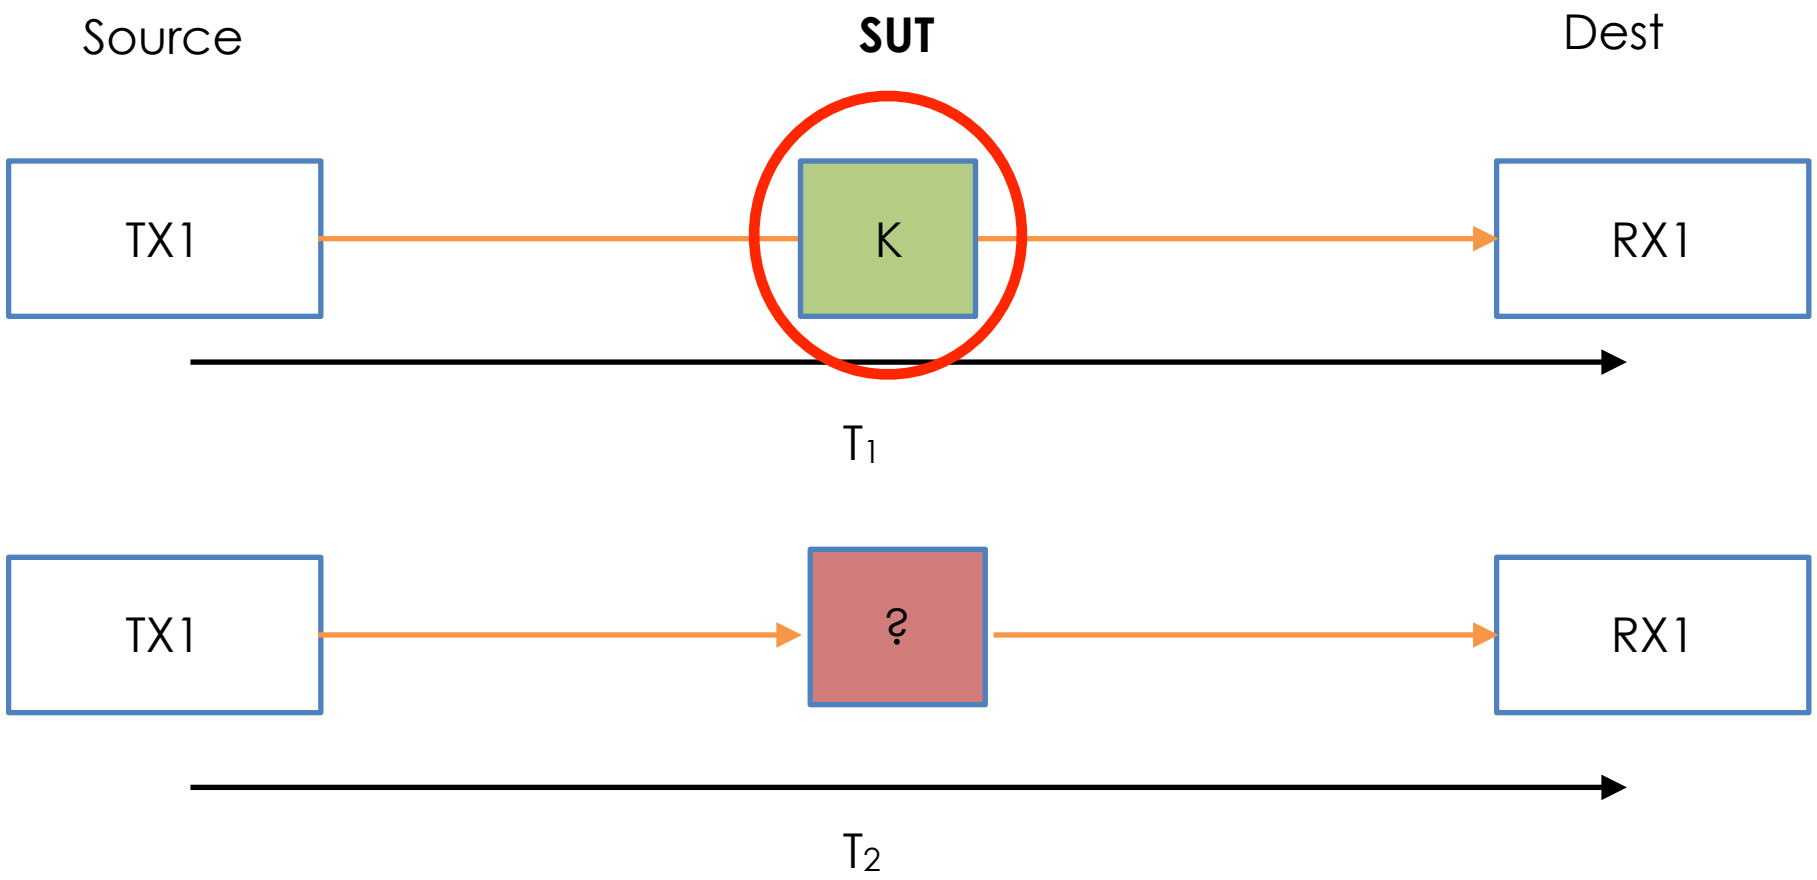

$$\text{Delay} = T_2 - T_1 - K$$

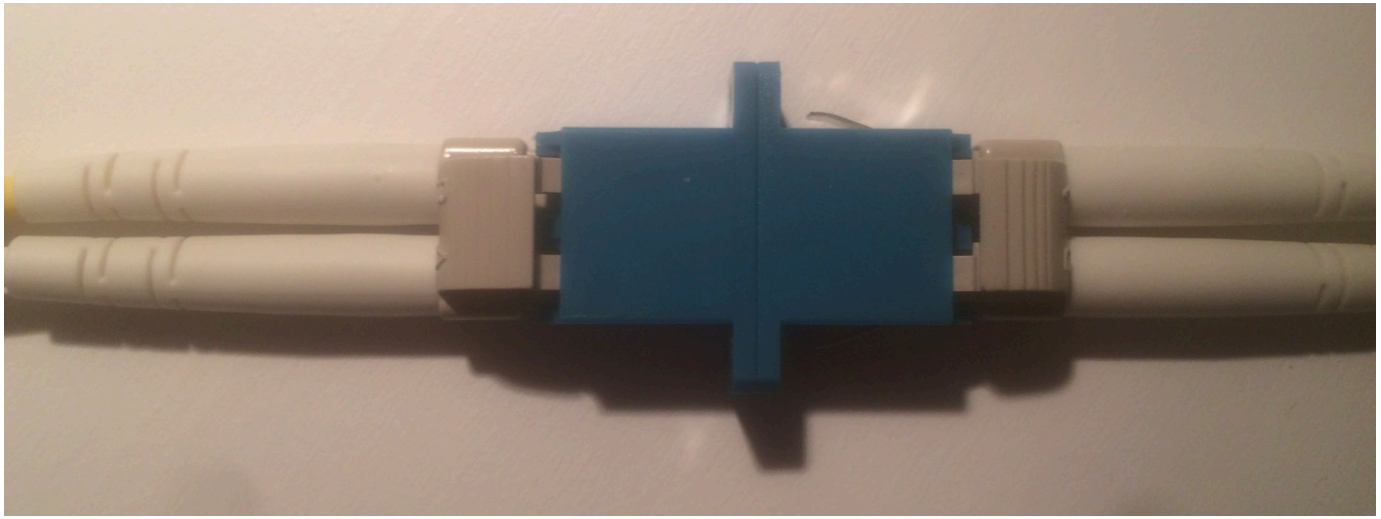

Source **SUT** Dest

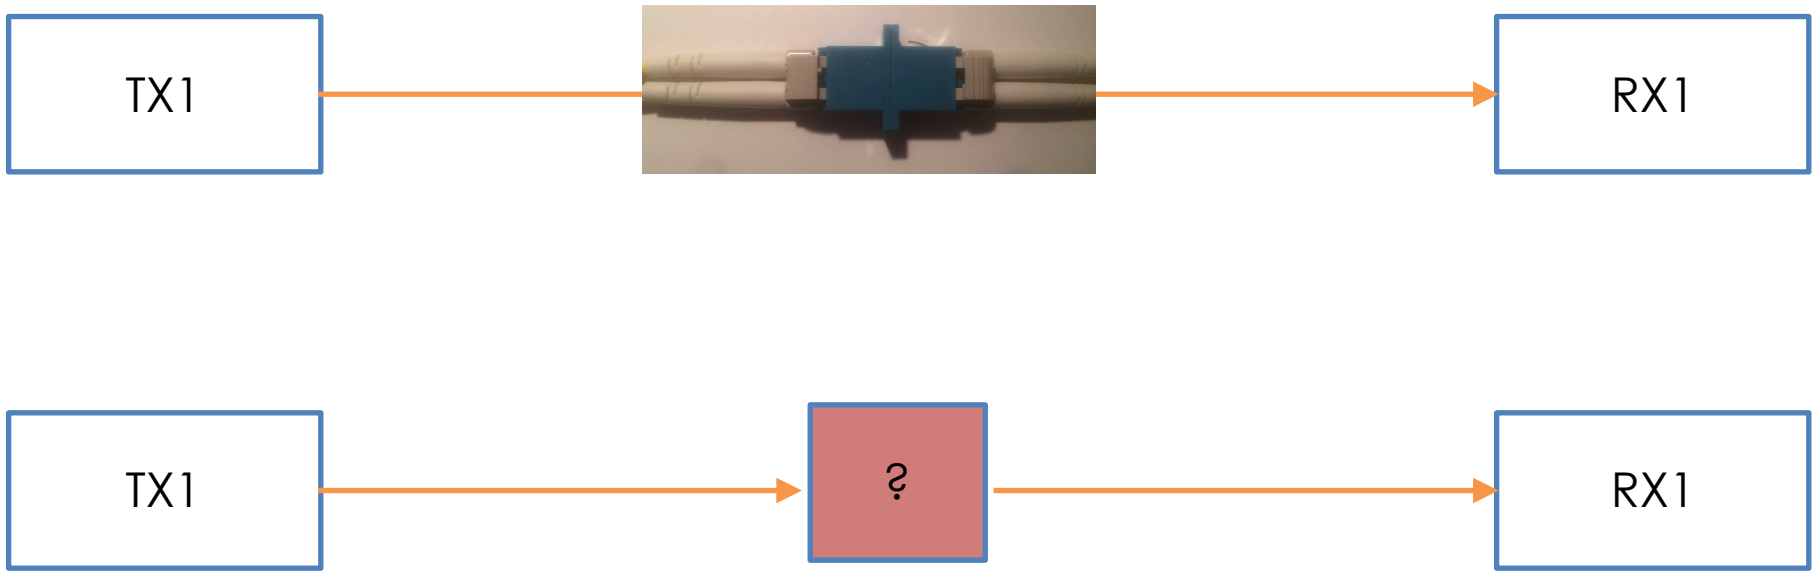


Finisar Sunnyvale, CA 94089
FTLX8571D3BCL
MADE IN MALAYSIA
850 nm
12-08
S/N: AM71E8J
Class 1 21CFR1040.10 LN#50 7/01

Source **SUT** Dest

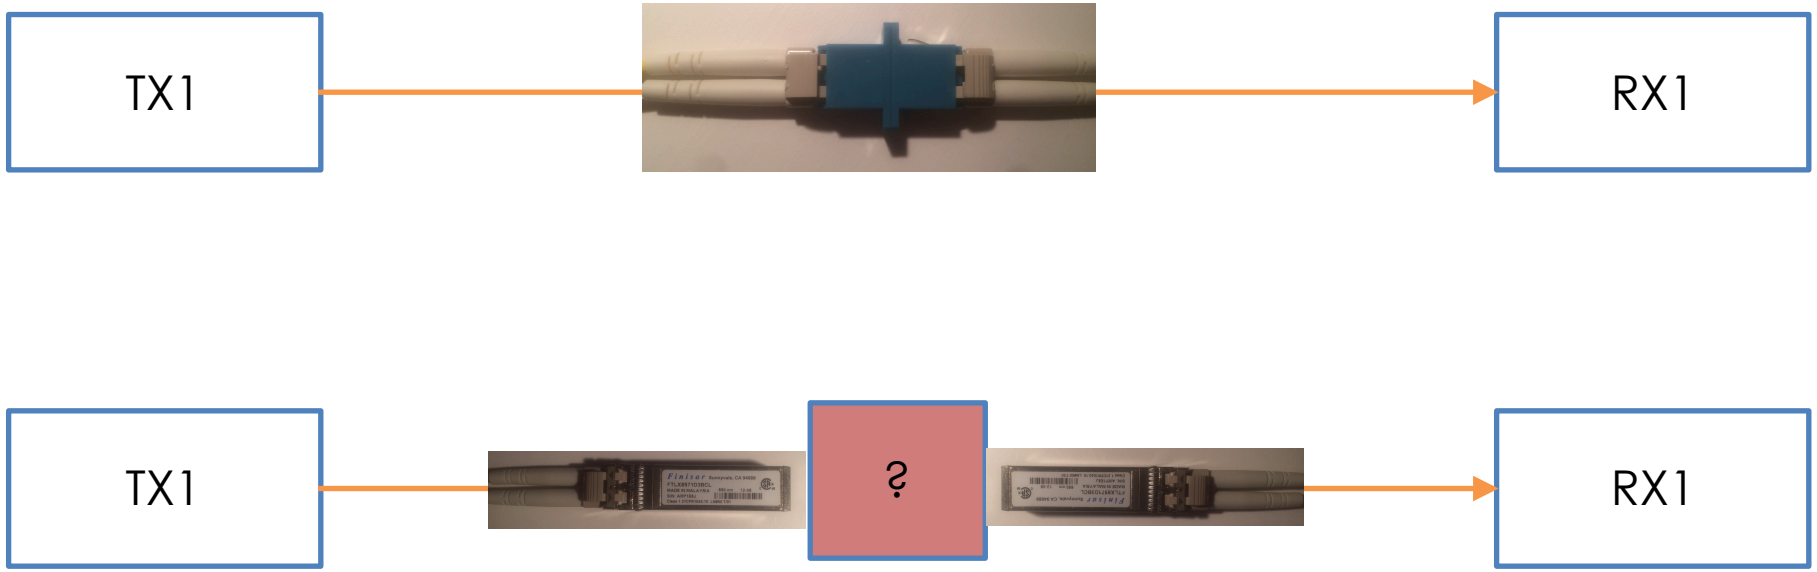

Source **SUT** Dest

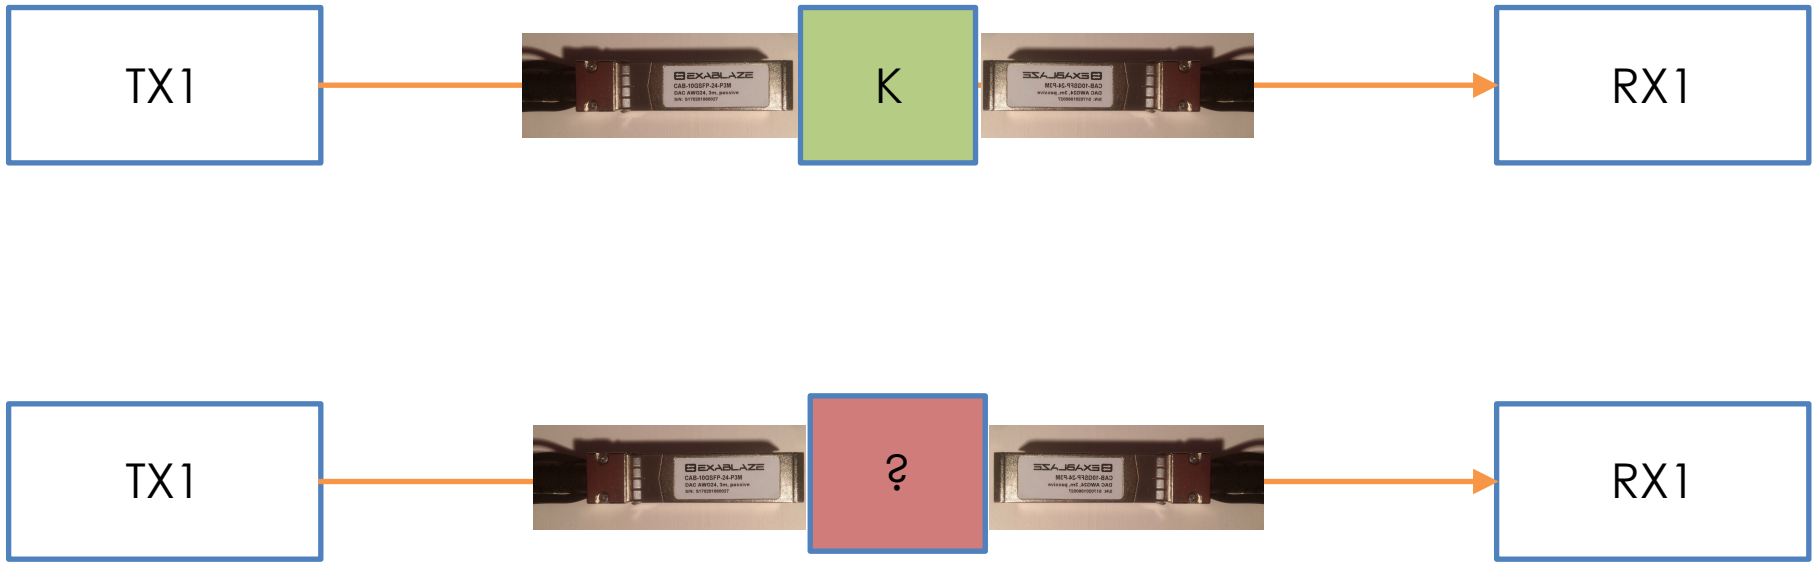

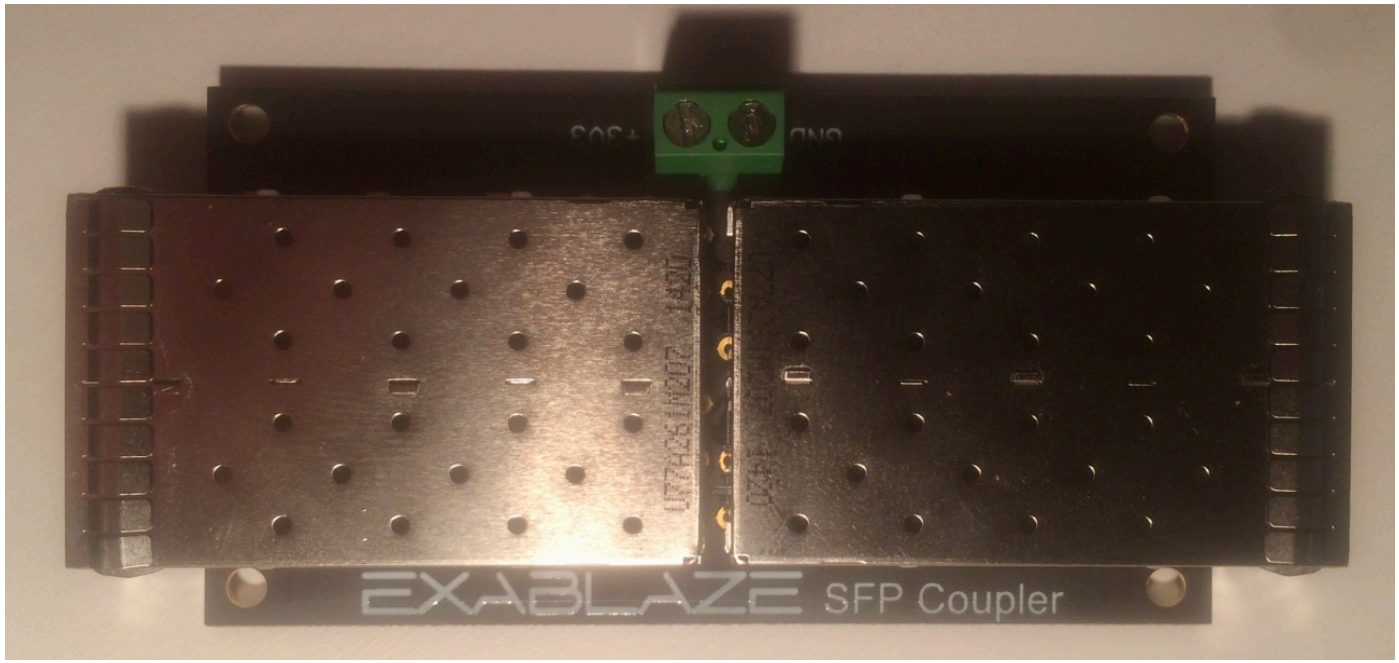
EXABLAZE SFP Coupler

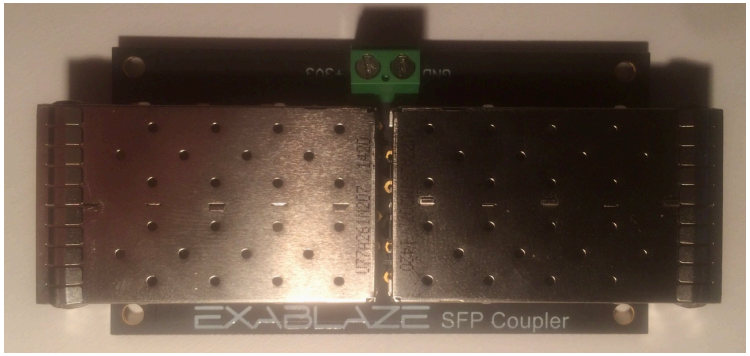
EXABLAZE SFP Coupler

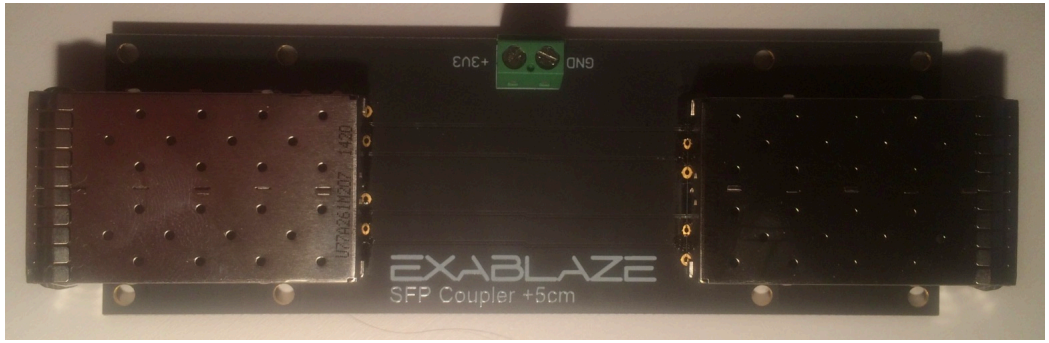
EXABLAZE
SFP Coupler +5cm

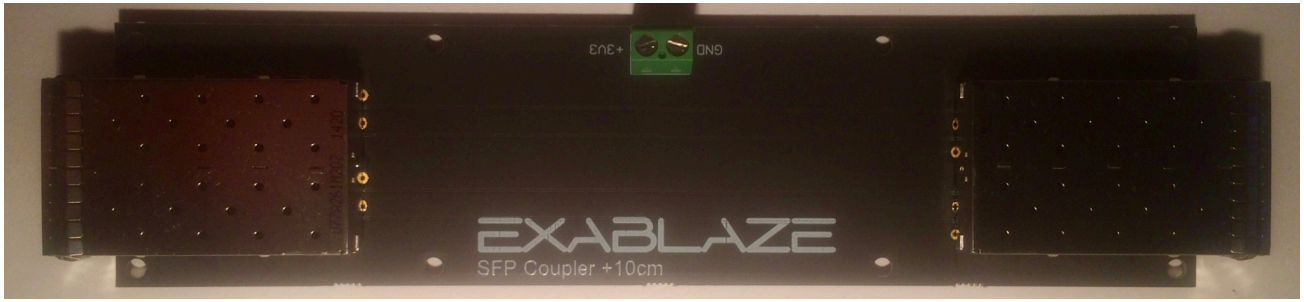
EXABLAZE
SFP Coupler +10cm

Speed of Light PCB

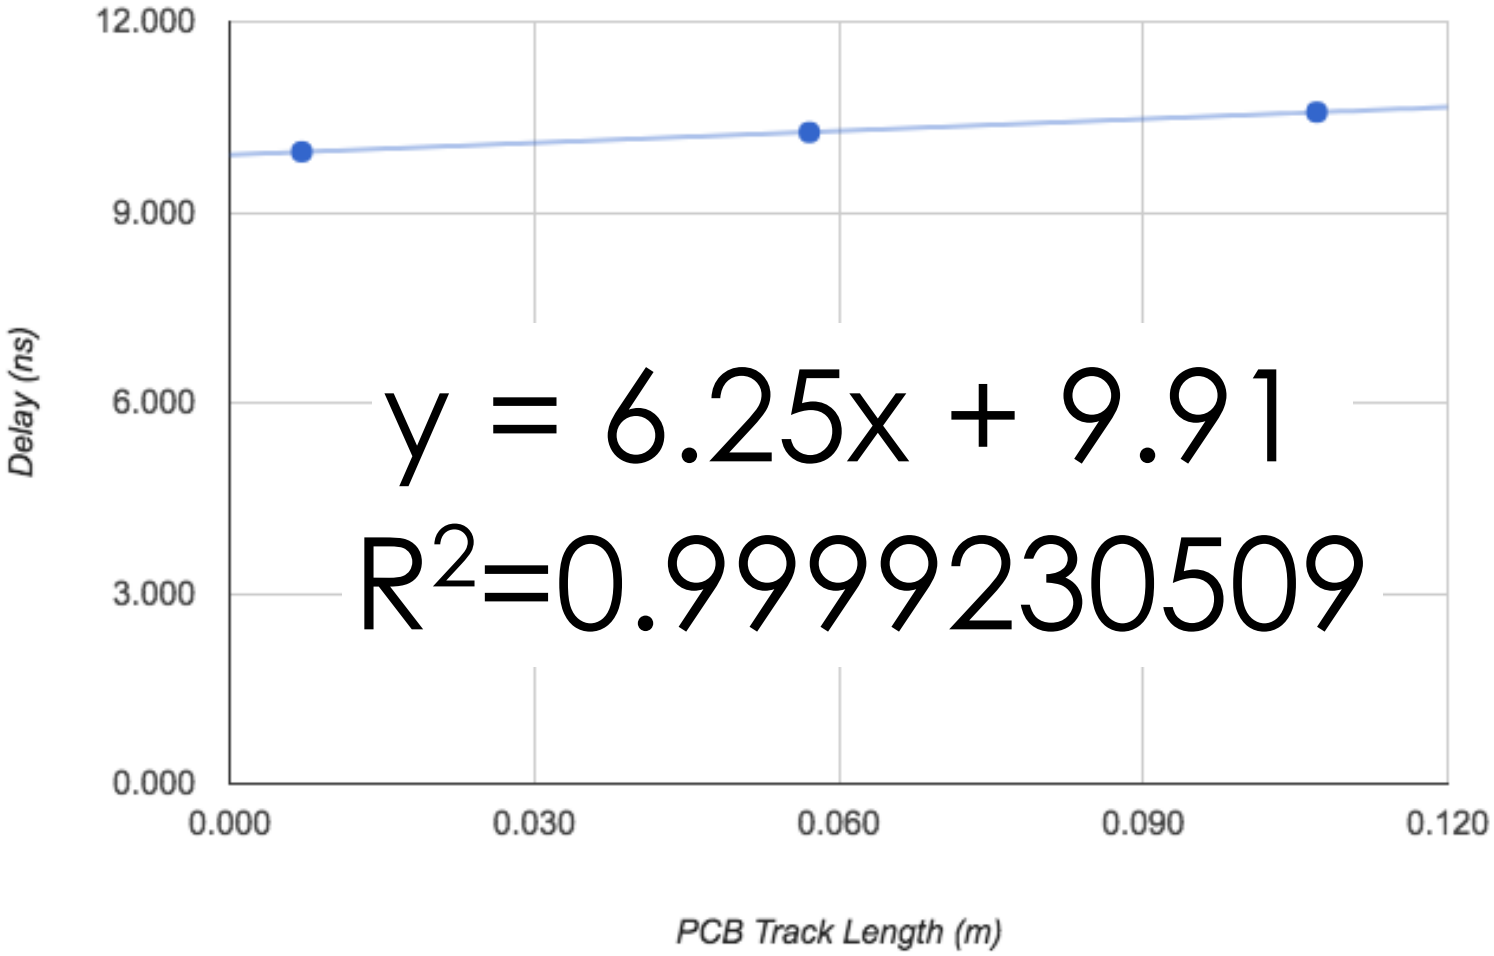

Speed of Light PCB

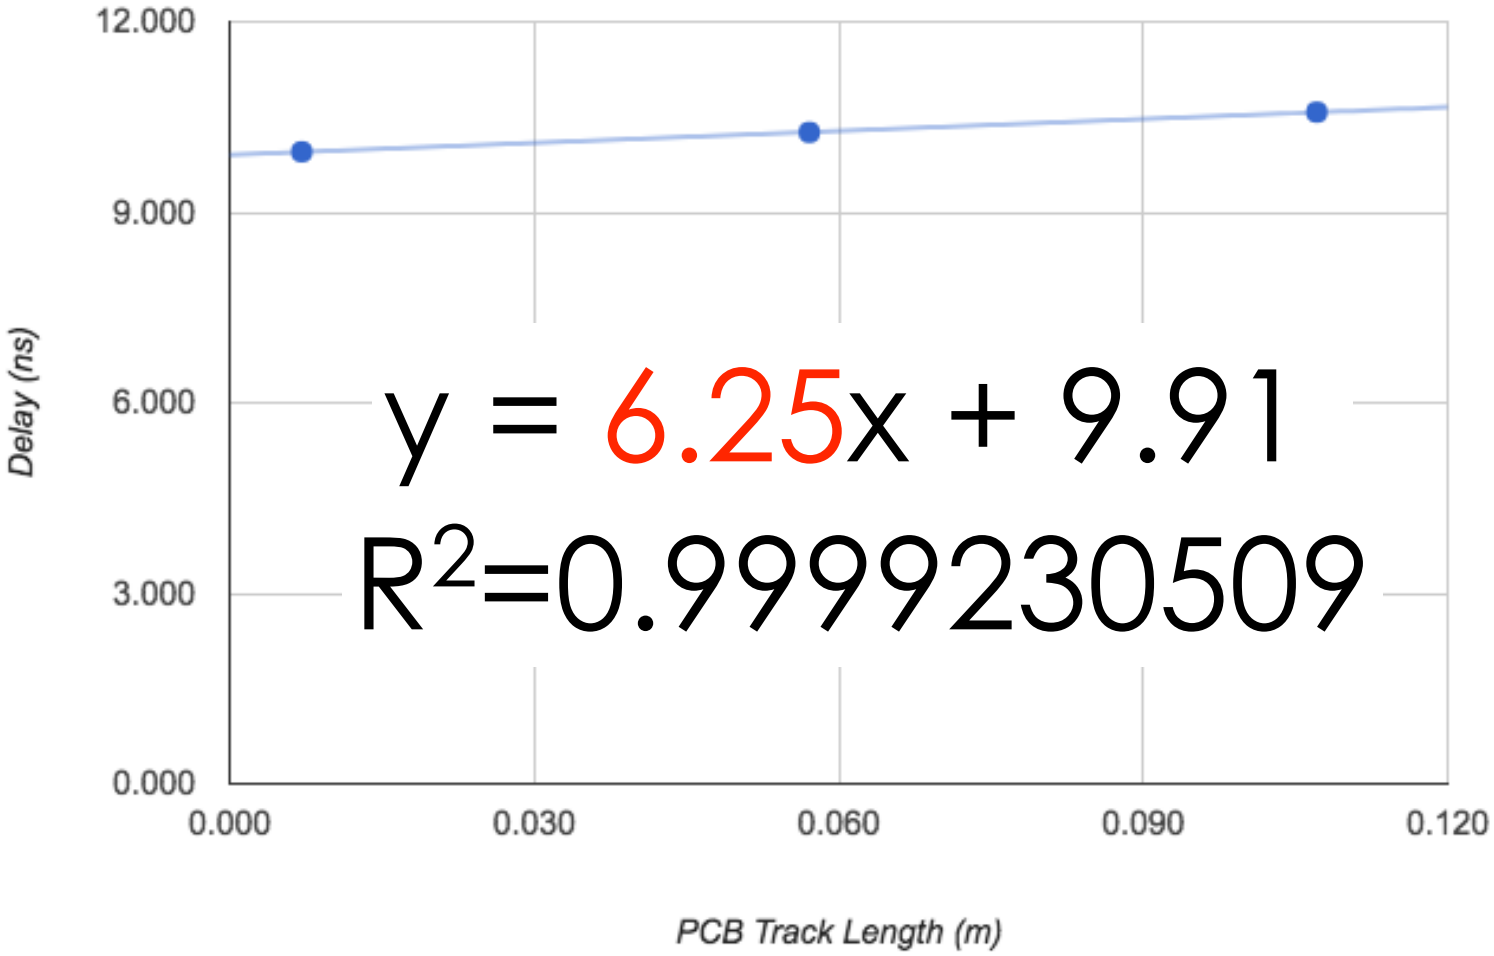

Speed of Light PCB

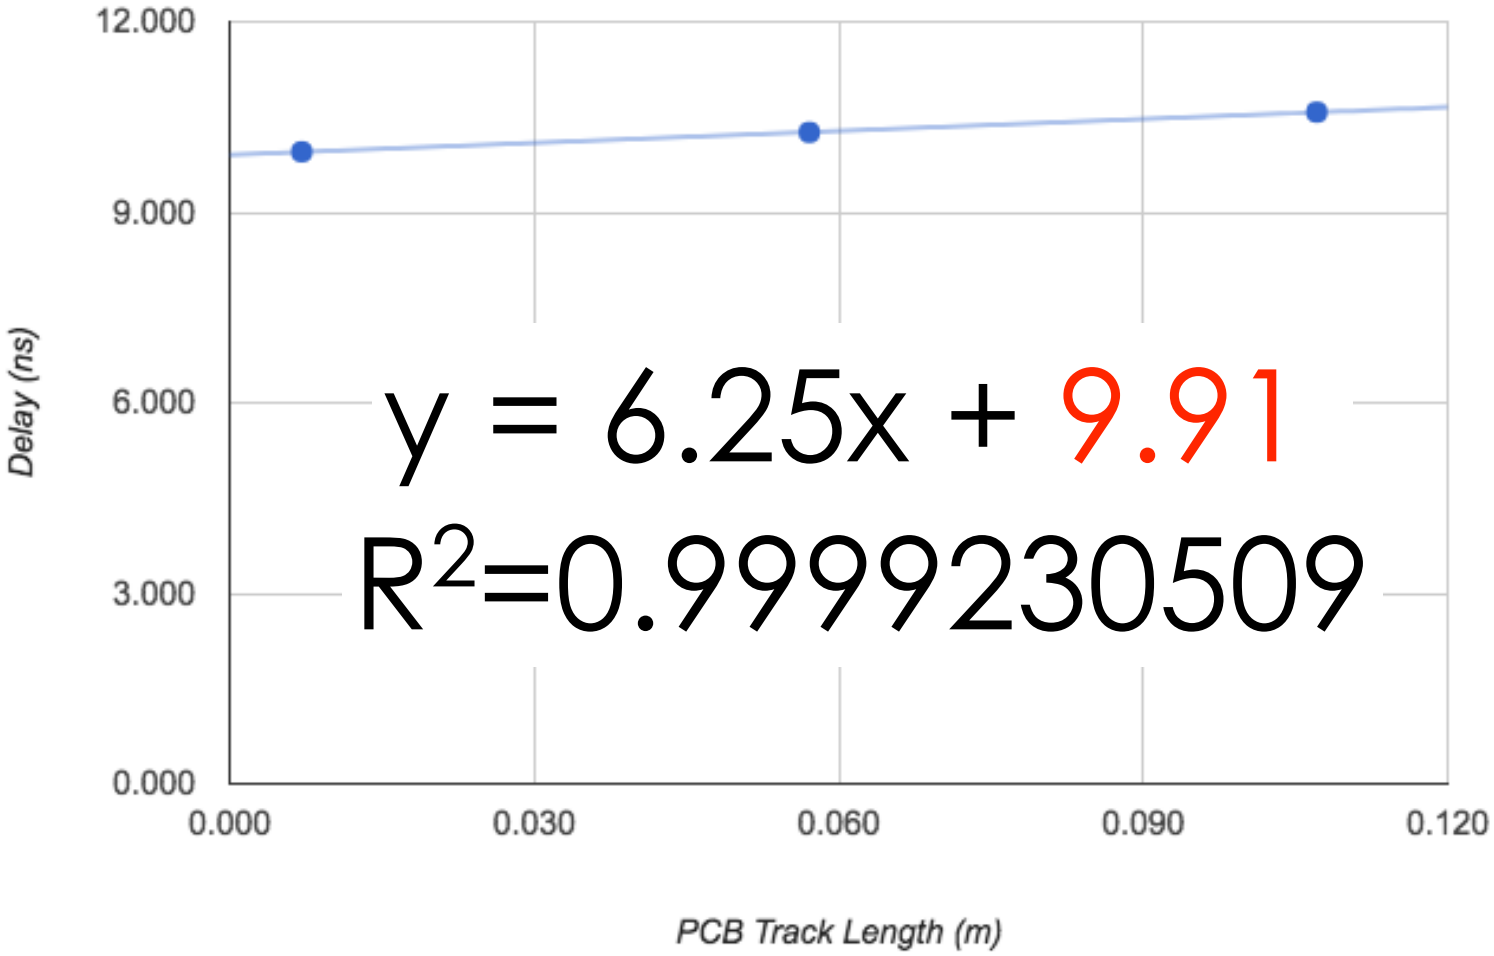

Source **SUT** Dest

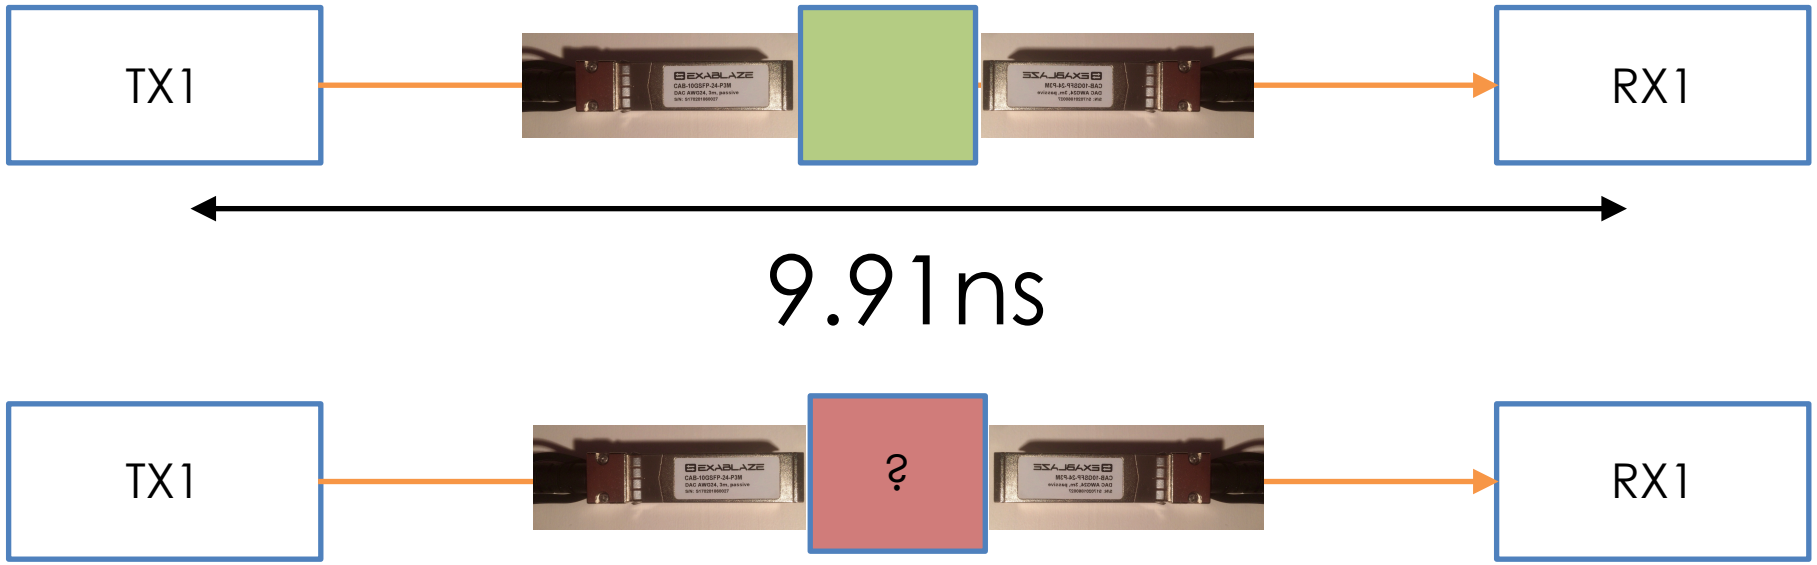

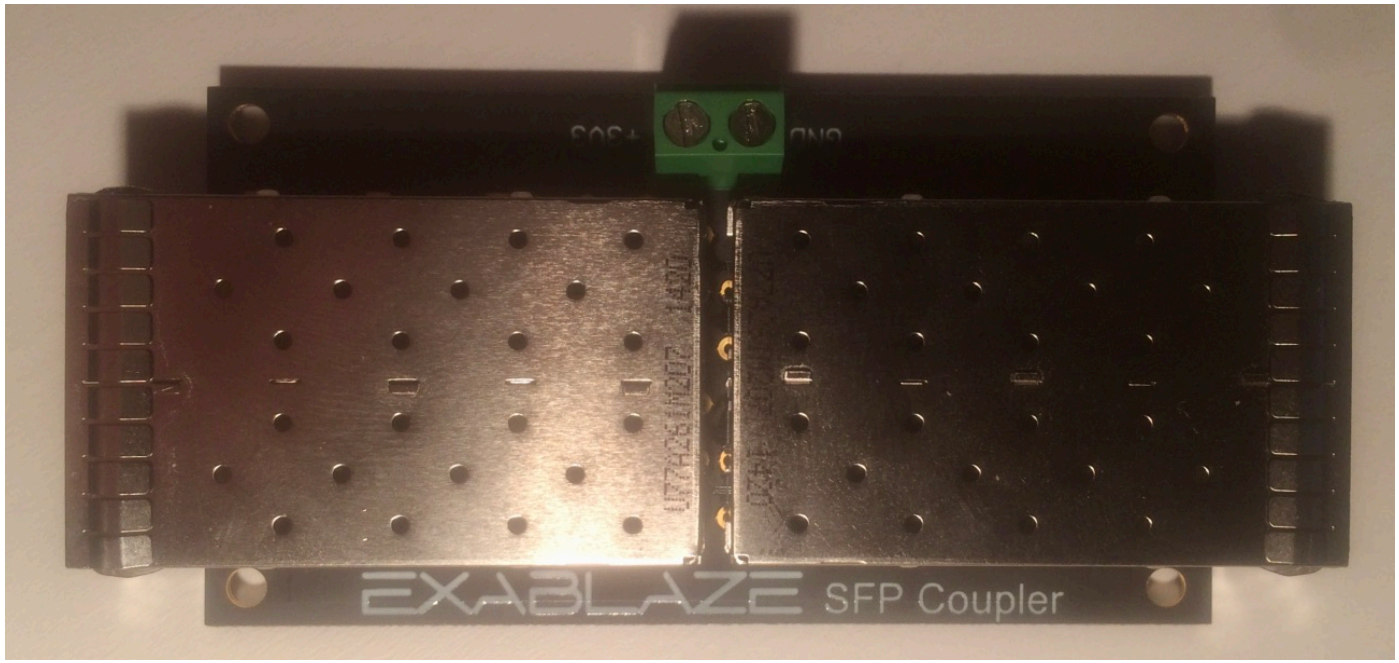
EXABLAZE SFP Coupler

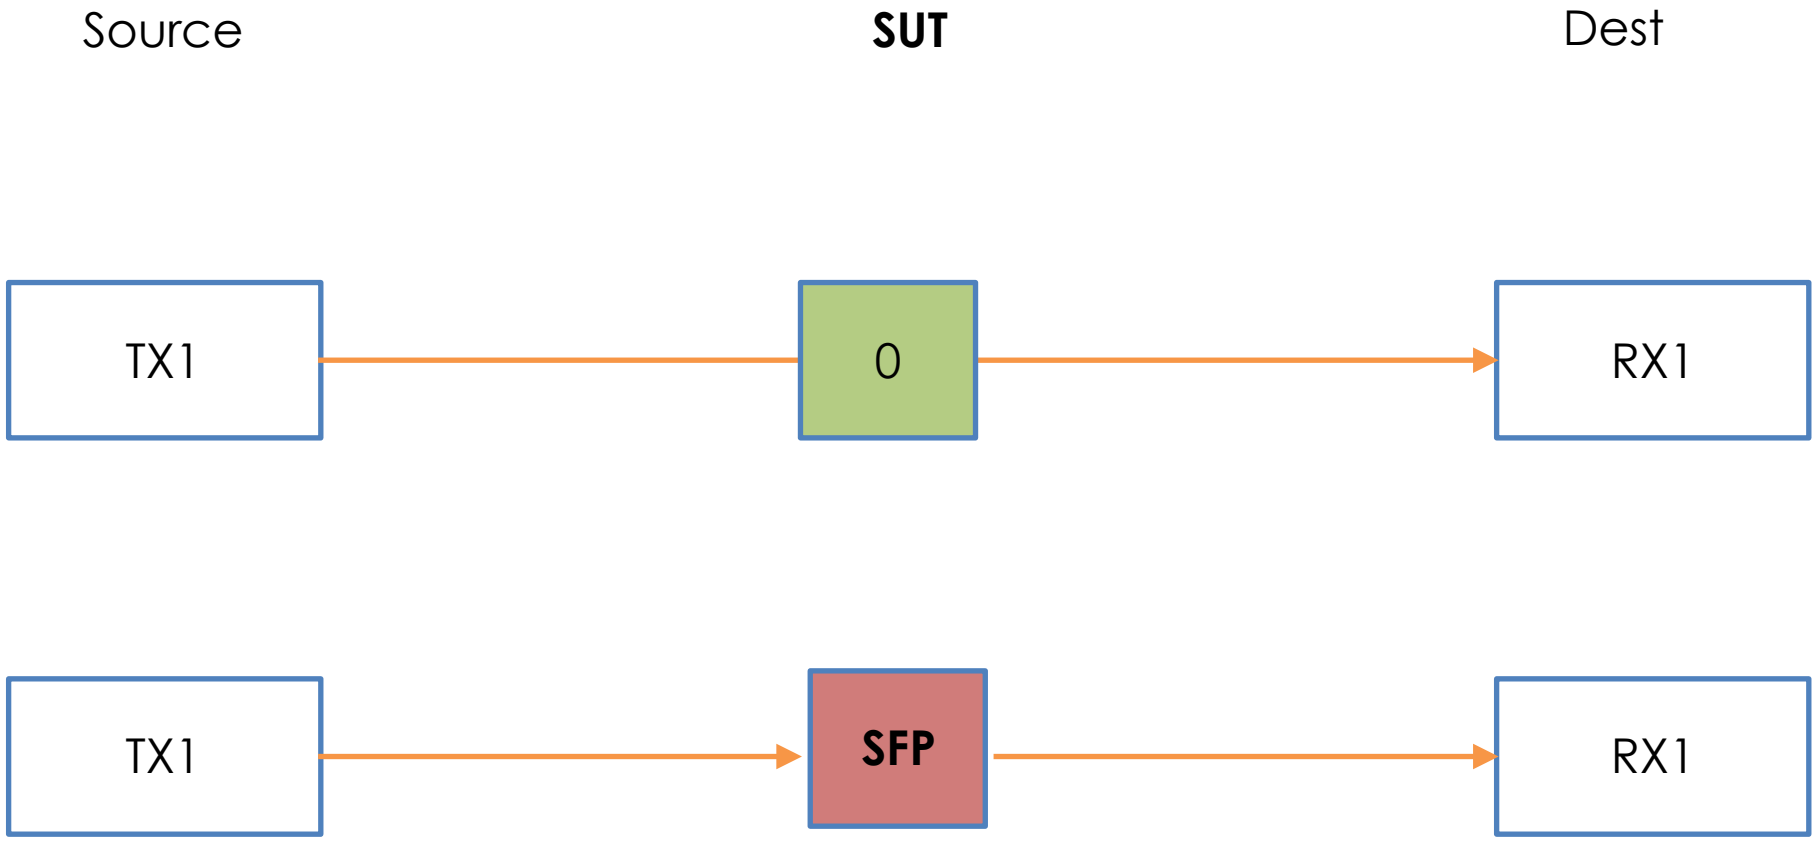
Source
SUT
Dest
TX1
0
RX1
TX1
SFP
RX1

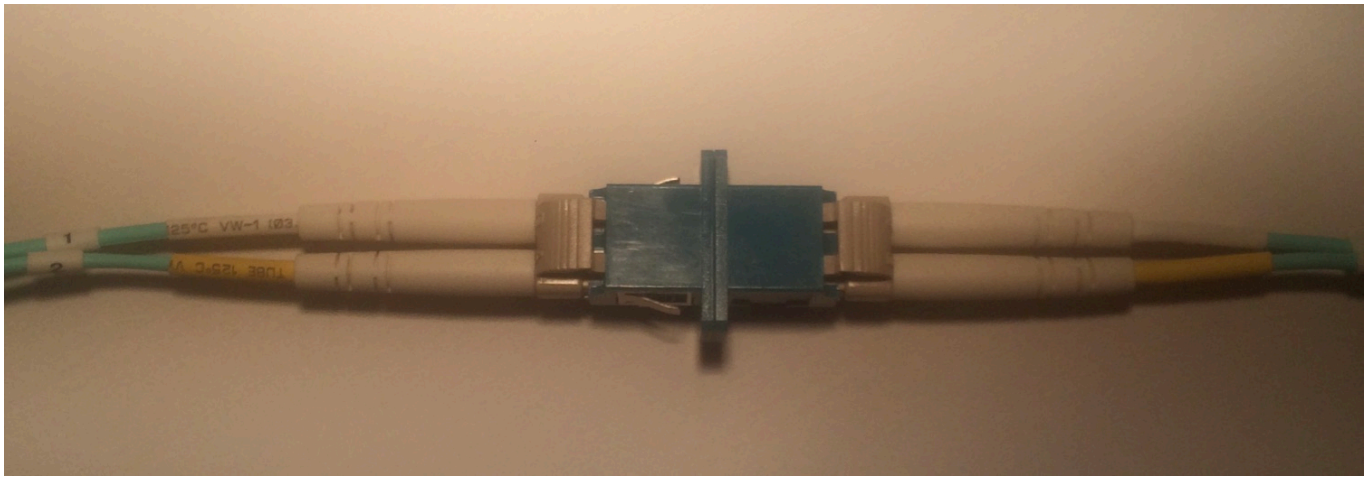

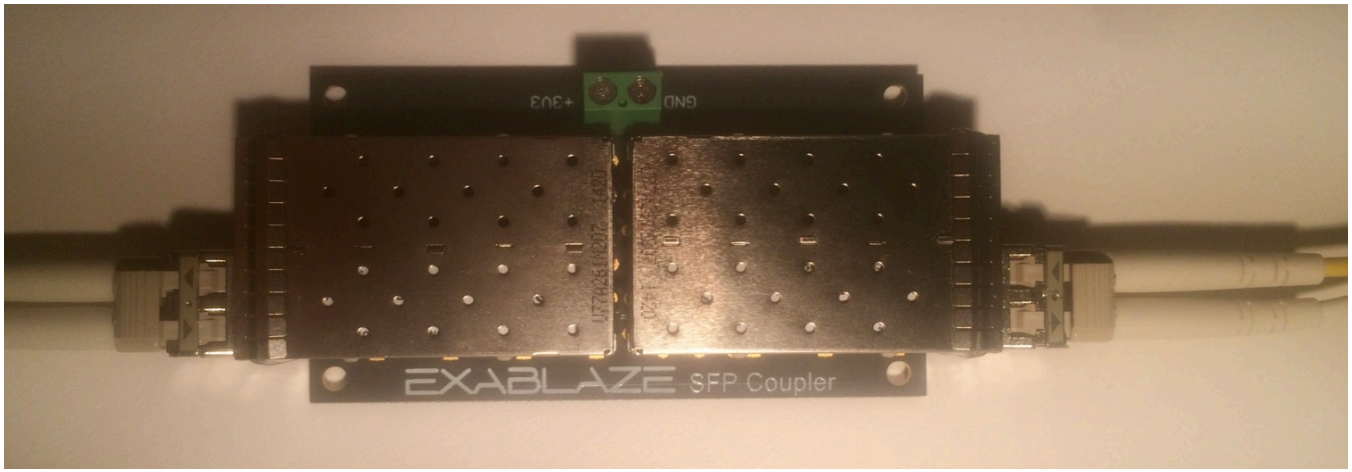
EXABLAZE SFP Coupler

Source **SUT** Dest

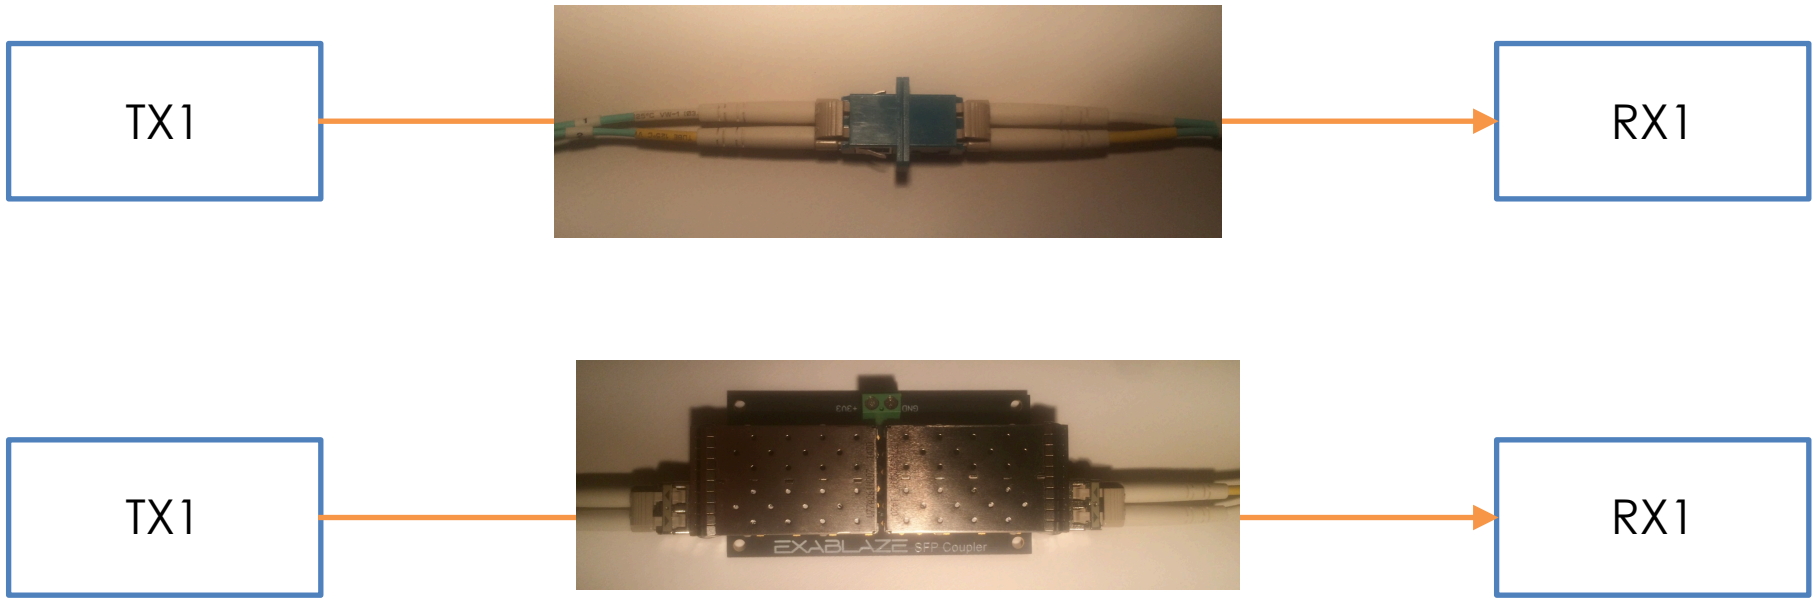

SFP Latency

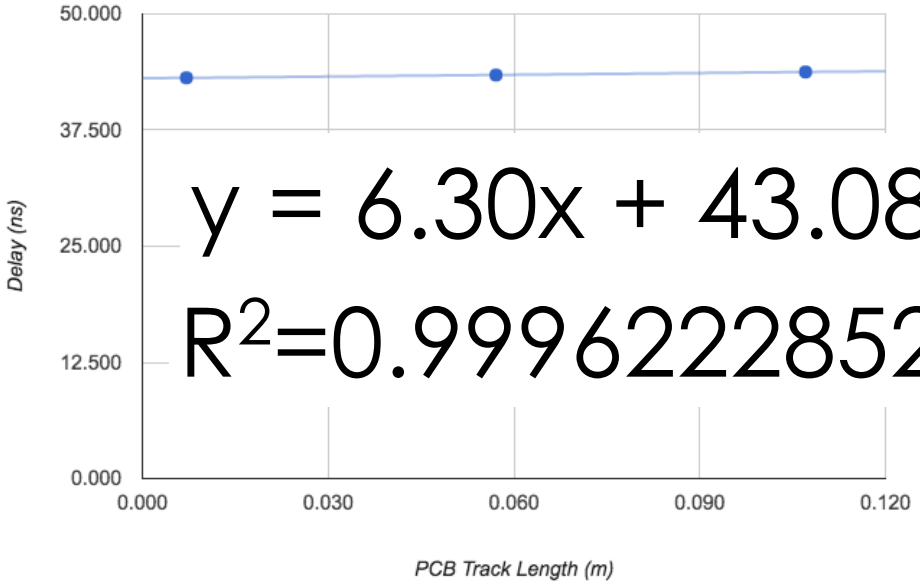

$y = 6.30x + 43.080$

$R^2=0.9996222852$

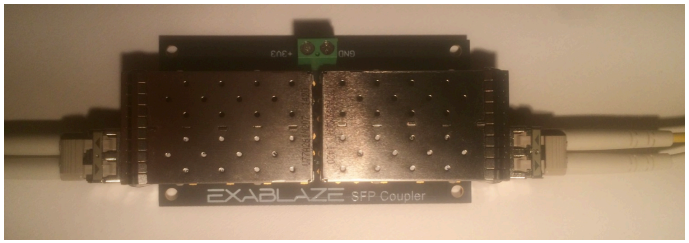

$y = 42.060$

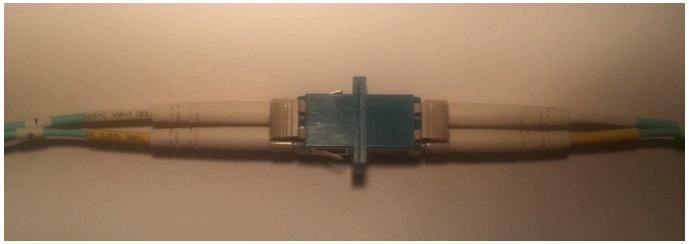

SFP Latency

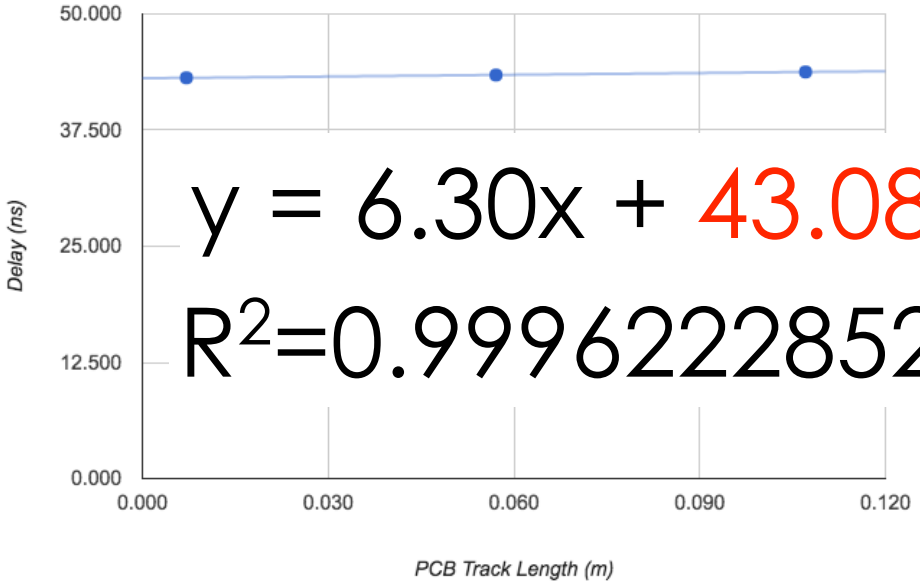

$y = 6.30x + 43.080$

$R^2=0.9996222852$

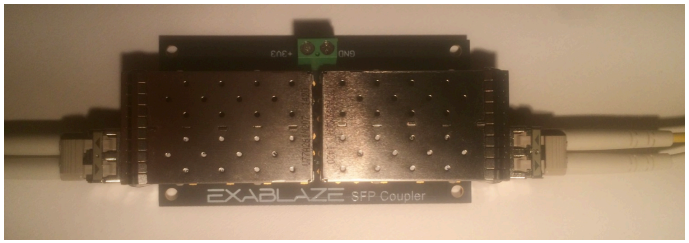

$y = 42.060$

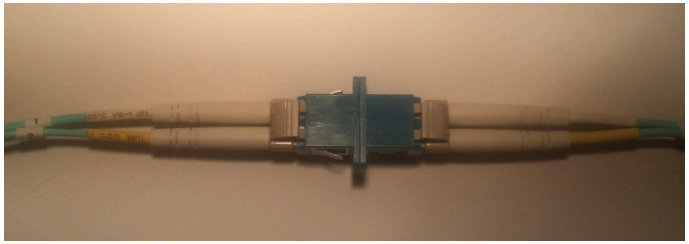

SFP latency = 43.080 - 42.060

SFP Latency

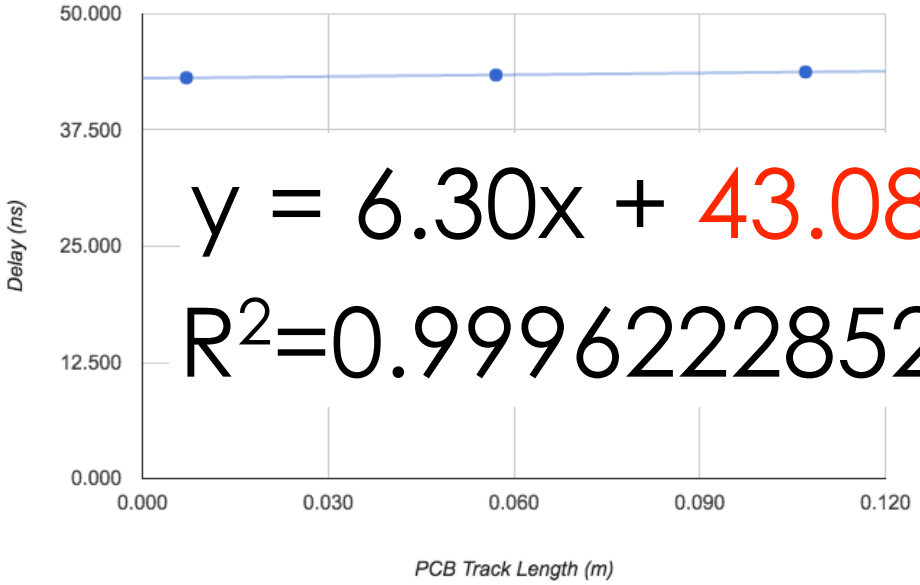

$y = 6.30x + 43.080$

$R^2=0.9996222852$

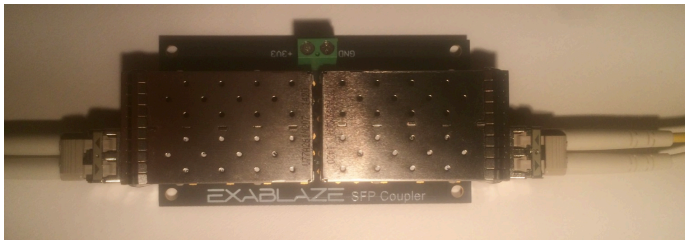

$y = 42.060$

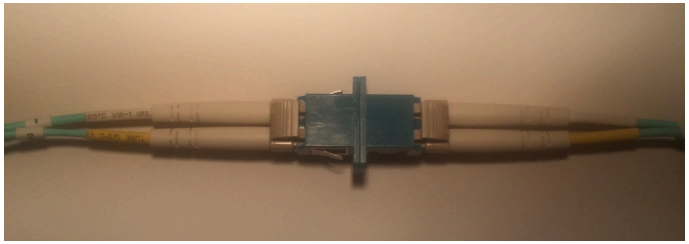

SFP latency = 43.080 - 42.060 = 1.02ns

SFP Latency

830ps vs 1020ps

SFP Latency

830ps vs 1020ps

1020 - 830 = 190ps

SFP Latency

830ps vs 1020ps

1020 - 830 = 190ps

HPT resolution 250ps

SFP Latency

830ps vs 1020ps

1020 - 830 = 190ps

HPT resolution 250ps

SFP ~ 925ps

Source **SUT** Dest

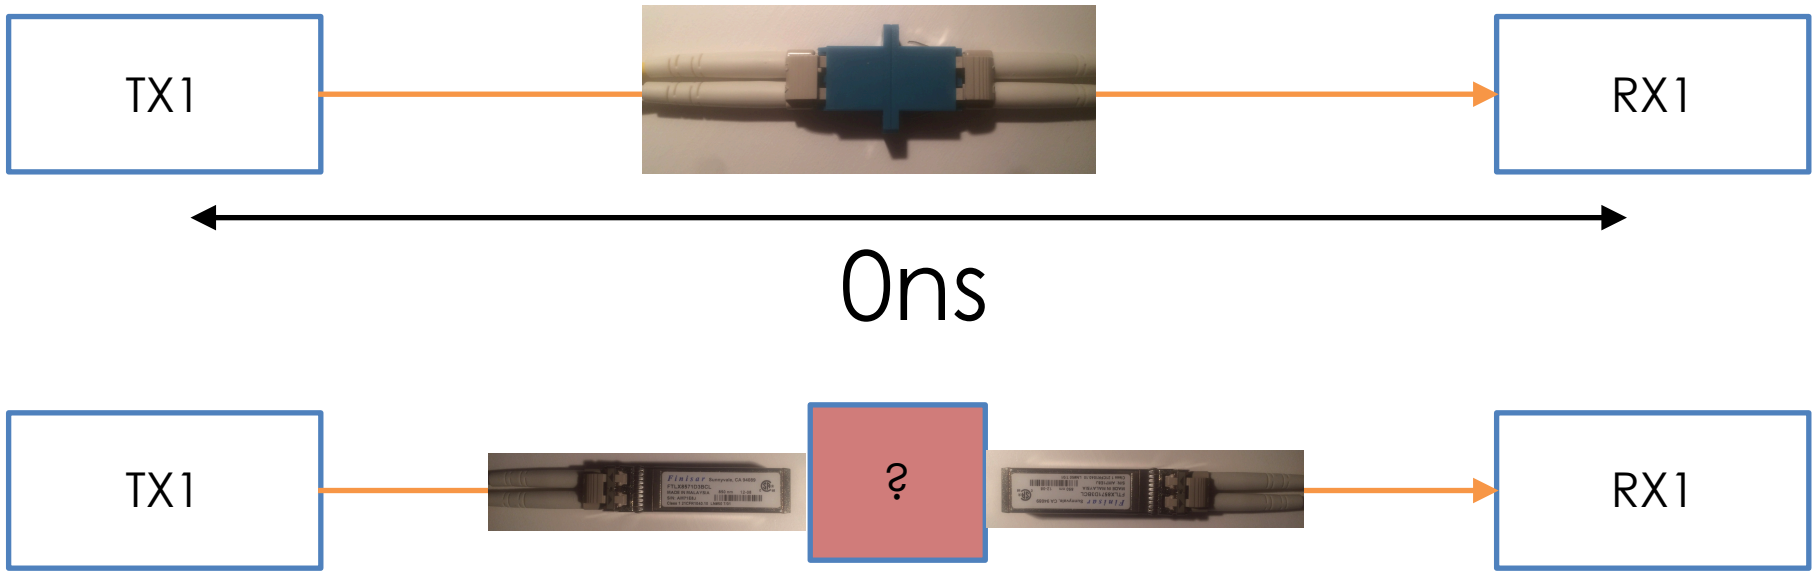

Source **SUT** Dest

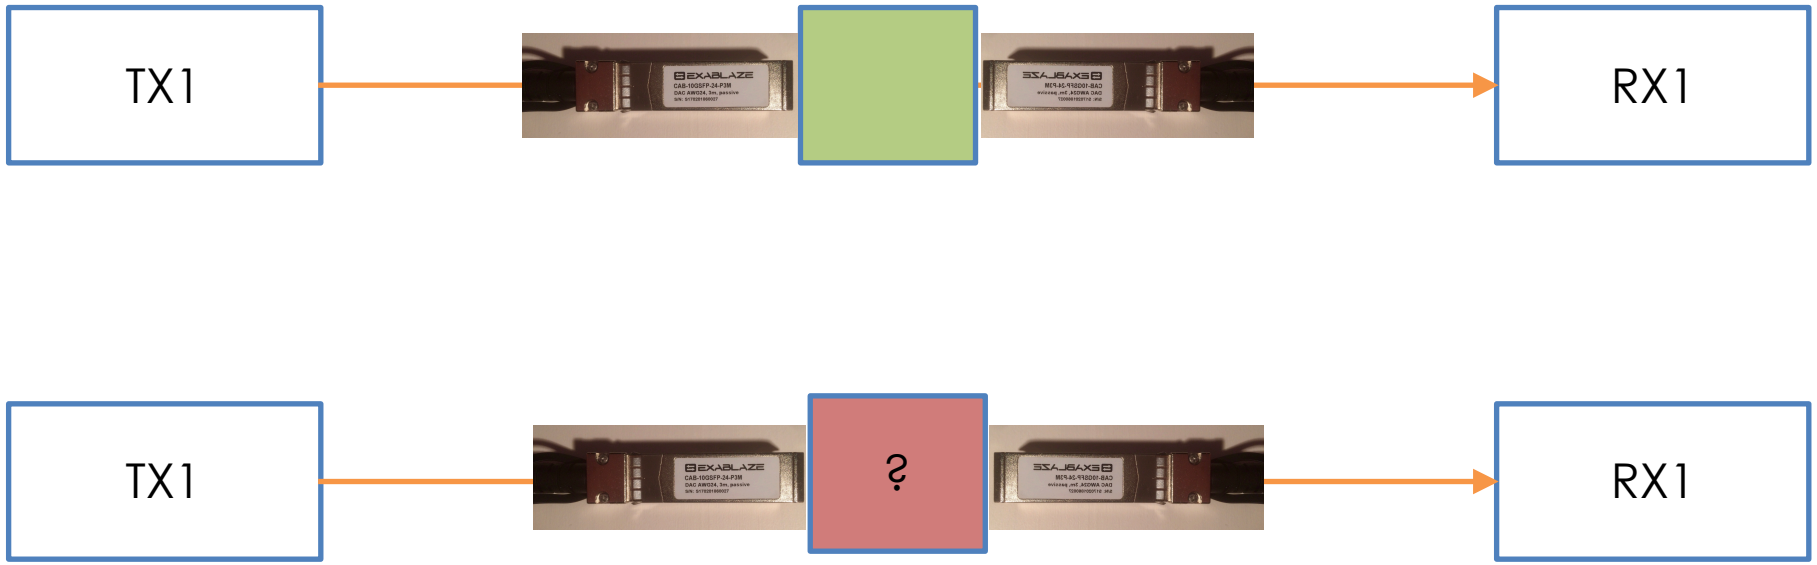

Source **SUT** Dest

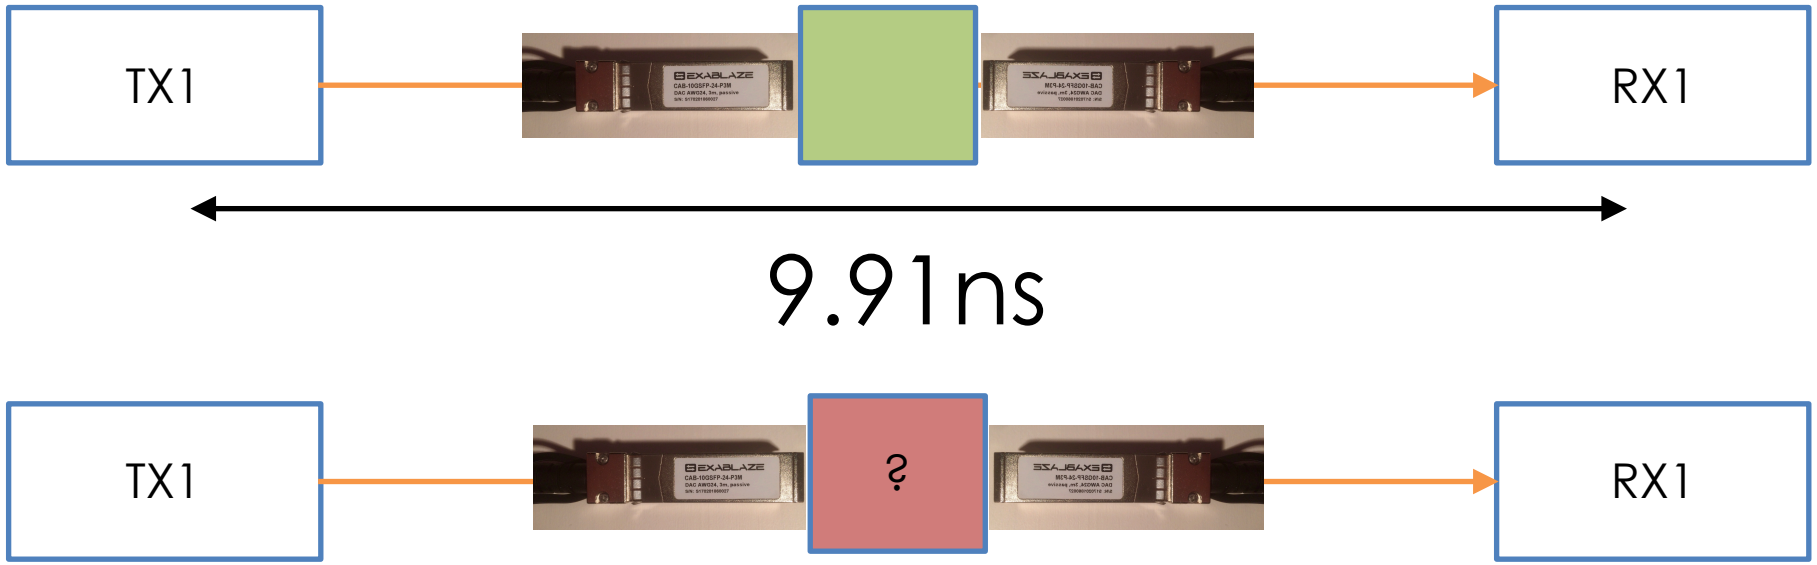

Network Multiplexor ?

Source

Exablaze FastMux

Dest

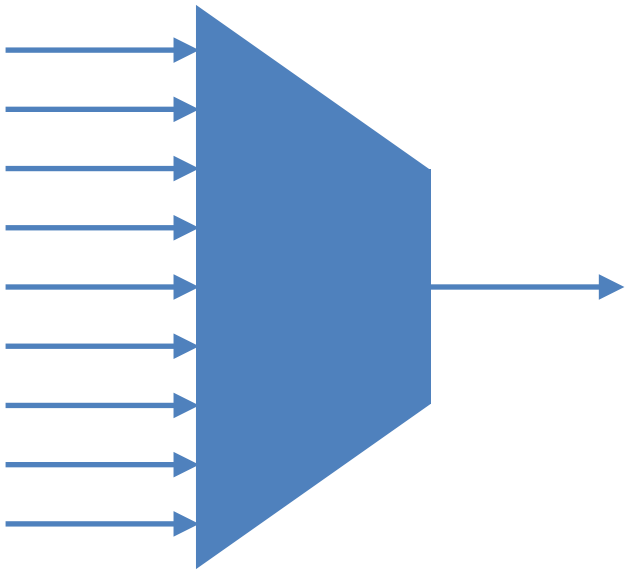

Source

Exablaze FastMux

Dest

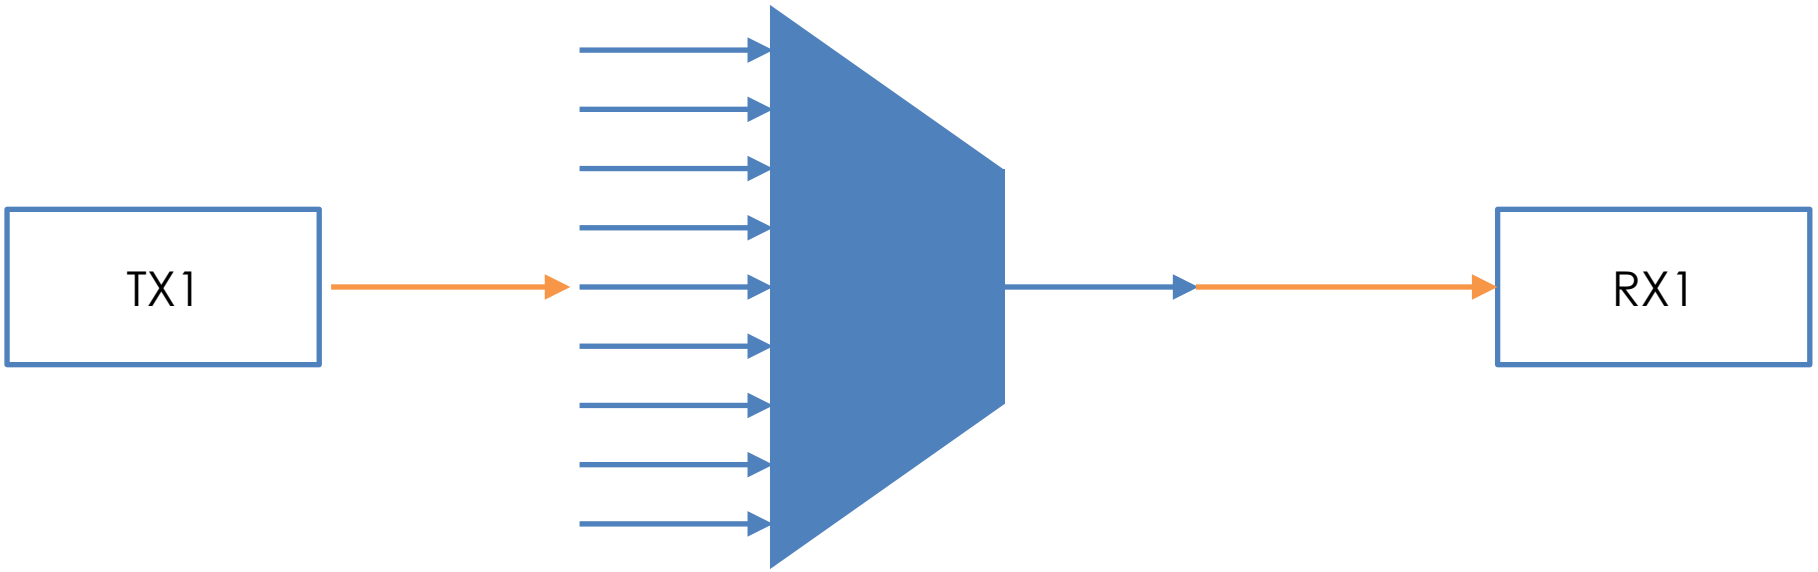

Source

Exablaze FastMux

Dest

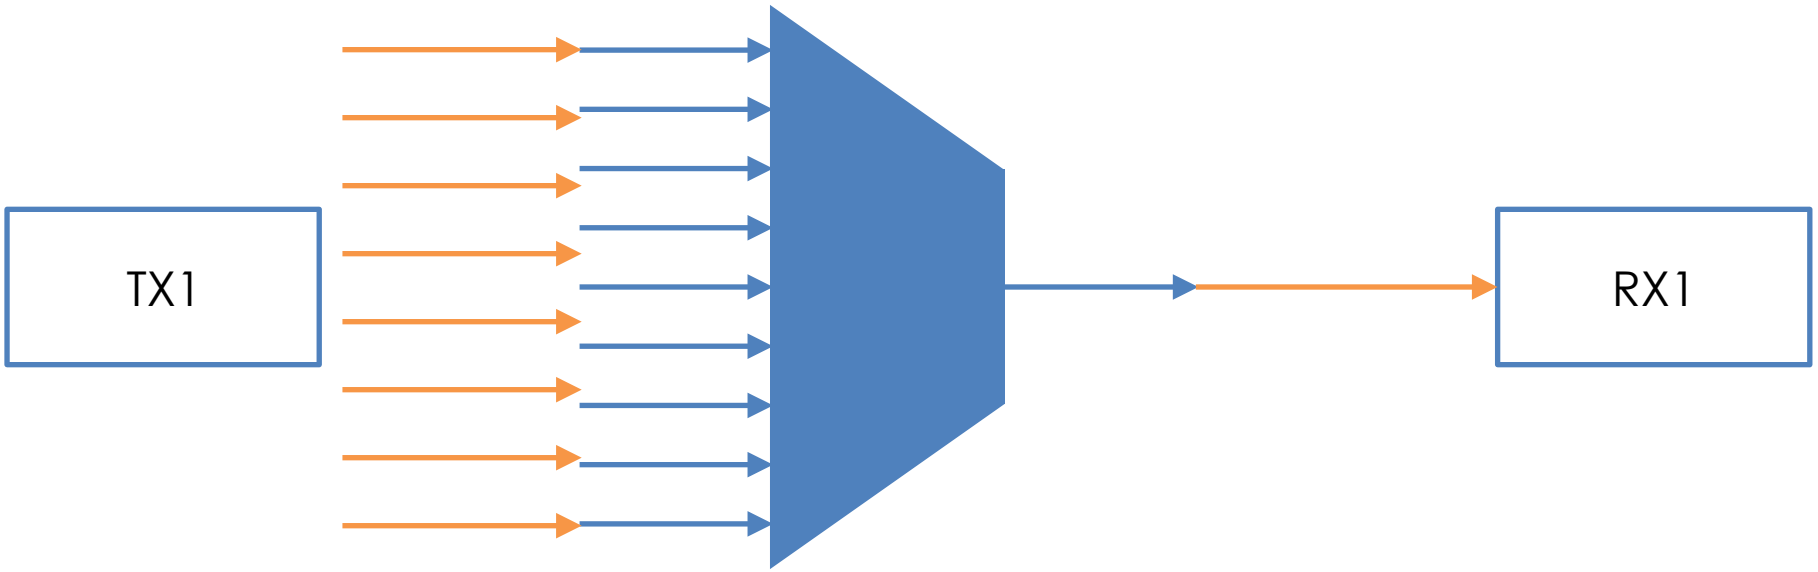

Performance Report – Exalink Fusion FastMux

This report details the performance of the ExaLINK Fusion **FastMux** firmware (ver 1.6.0). All measurements were taken using the **ExaNIC HPT**, High Precision Timing capture card (0.25ns resolution) from the start of a frame arriving, to the start of a frame departing the device. For this test, line card bay B was used and port B14 was configured as the upstream (egress) port. Ingress ports were configured in the following order: B4,2,1,12,6,5,16,15,10,9,8,13,7,3,11. Reported results are port-to-port at 10 Gb/s, delays due to SFP+ transceivers and signal propagation over fibers/cables are not included.

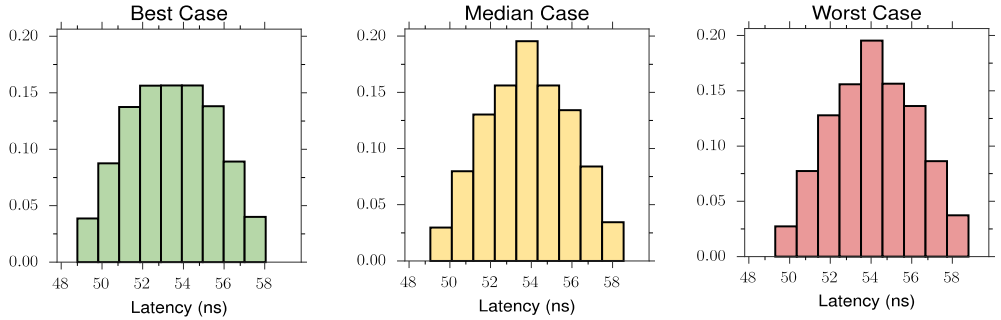

Figure 1: Latency distributions for (a) best case - B12, (b) median case - B15, and (c) worst case - B1

Figure 1 shows latency distributions for the best case (minimum of minimums), median case (median of medians) and worst case (maximum of maximums) configurations, found on ports B12, B15 and B1 respectively. The results show distributions strongly centered on a median of approximately 53ns and minimal variance between the cases.

Table 1 shows the performance of ports B1-B16 sorted in order of average latency. The absolute minimum latency of 48.79ns is achieved by ports B12, B11, B7 and B9. The median of medians latency of 53.79ns is found on the majority of ports. The median latency varies by 0.5ns between all ports. The latency for any given port varies by at most 9.50ns.

ExaLink Fusion FastMux 49ns

48.79ns

ExaLink Fusion FastMux 49ns

48.79ns (min)

53.79ns (med)

59.79ns (max)

Gotchas - Beware

Gotchas - Beware

#1 Device Precision ≠ Report Precision

Gotchas - Beware

#1 Device Precision ≠ Report Precision

#2 Accuracy matters more than resolution

Gotchas - Beware

#1 Device Precision ≠ Report Precision

#2 Accuracy matters more than resolution

#3B Copper isn't always faster than fibre.

Gotchas - Beware

#1 Device Precision ≠ Report Precision

#2 Accuracy matters more than resolution

#3B Copper isn't always faster than fibre.

#4 1m isn't 1.00m

Tips - Remember

Tips - Remember

#1 Round to nearest multiple of clock period
Or use full precise device/format

Tips - Remember

#1 Round to nearest multiple of clock period
Or use full precise device/format

#2 Quantify and calibrate accuracy first!

Tips - Remember

#1 Round to nearest multiple of clock period
Or use full precise device/format

#2 Quantify and calibrate accuracy first!

#3 C in fibre 4.91ns/m
C in copper 4.20 - 5.03ns/m
C in PCB 6.25ns/m

SFP latency 925ns

Tips - Remember

#1 Round to nearest multiple of clock period
Or use full precise device/format

#2 Quantify and calibrate accuracy first!

#3 C in fibre 4.91ns/m
C in copper 4.20 - 5.03ns/m
C in PCB 6.25ns/m

SFP latency 925ns

#4 Get your copper from Exablaze ;-)

Questions?


EXABLAZE