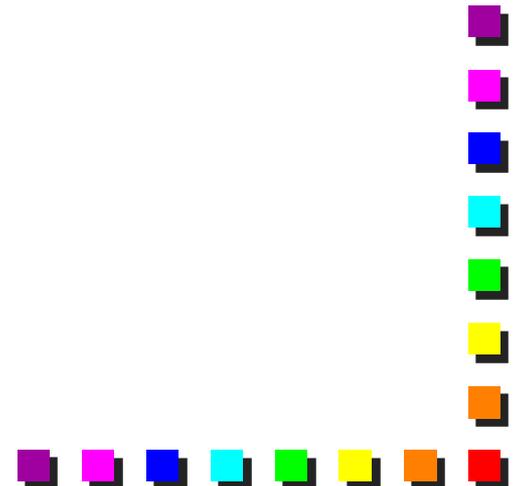


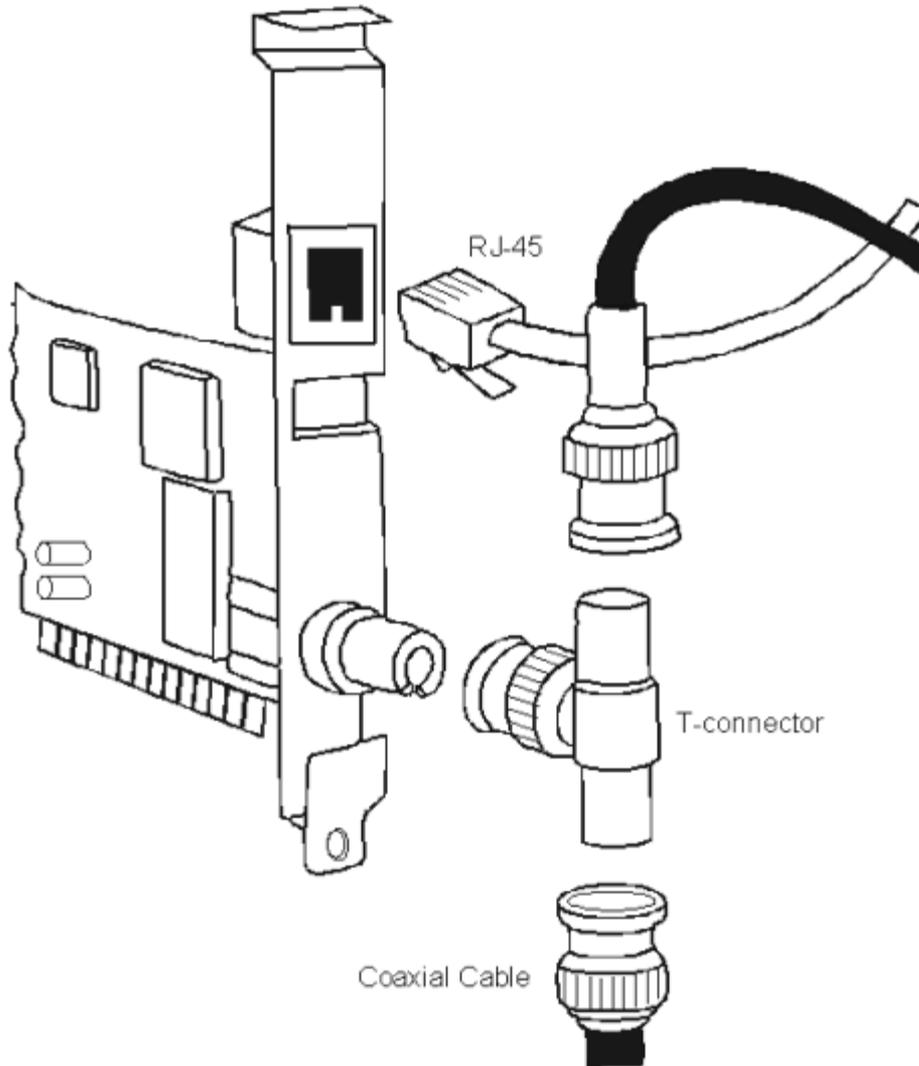
NICs and Cabling

Fulvio Riso

Politecnico di Torino



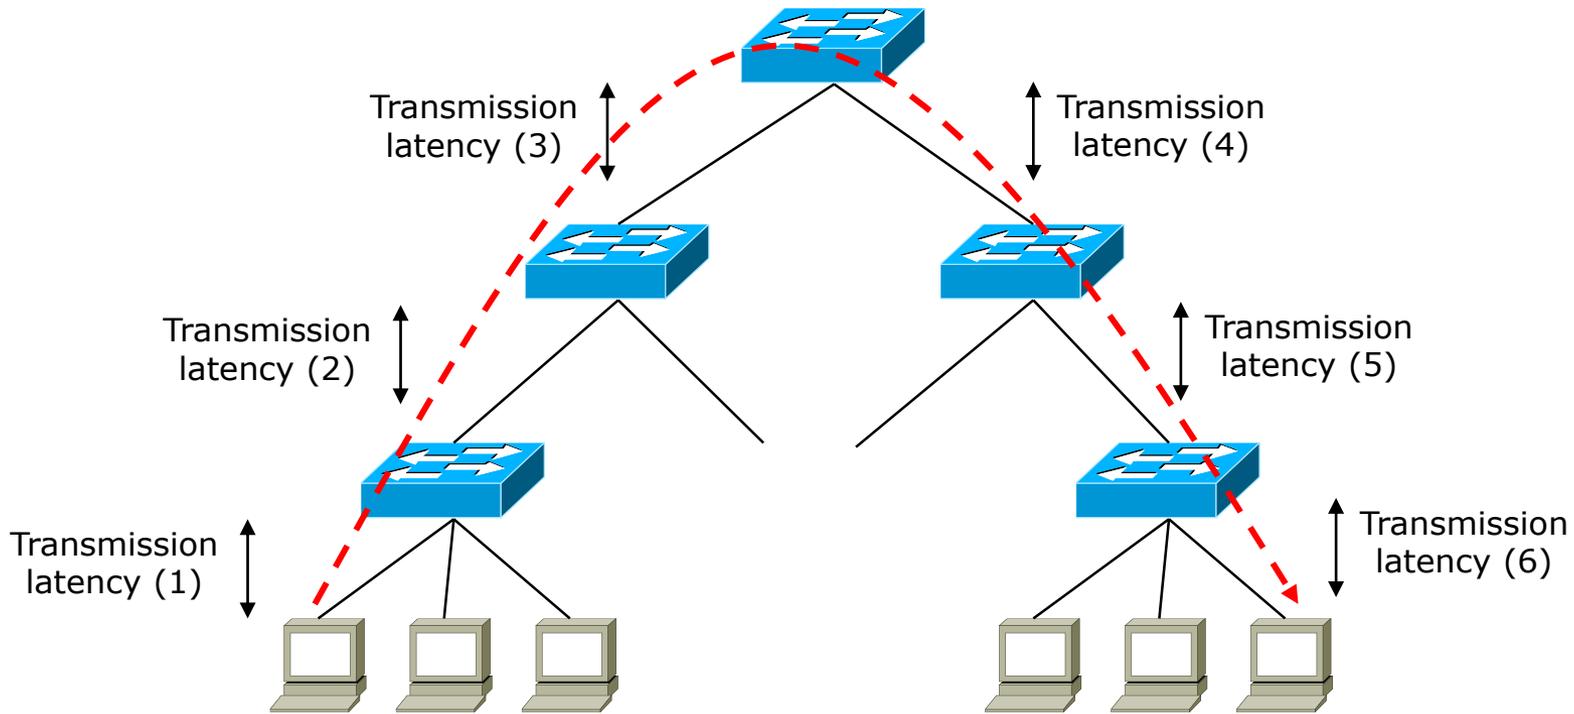
Original Ethernet: 10BaseT, 10Base2



10 Gigabit Ethernet

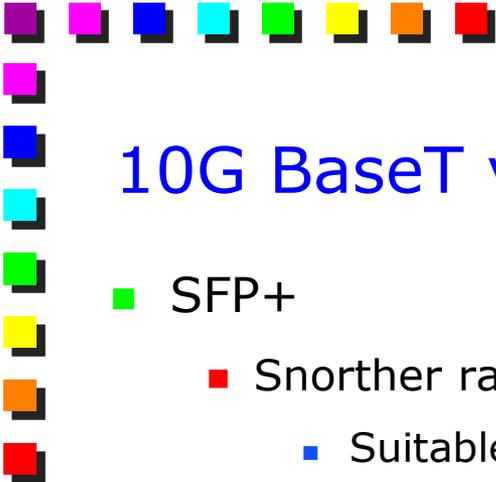
	Twinax	10GBase T	Multimode Fiber	Single mode fiber
Connector and cable	SFP+ direct attach copper	RJ-45 copper	Fiber optic	Fiber optic
Cable type	SFP+ direct attached twin axial cabling up to 10 m	Category-6 / 6A	MMF 62.5/50 μ m	SMF
Max Length	10m	100m with Cat-6A, 55m with Cat-6	300m	10Km
Latency	About 300ns	About 2.5us (~3000bytes)	negligible	negligible
Power (Intel, 2013)	6.2 W (dual port) X520-DA2	19W (dual port) X520-T2	6.8W (dual port) X520-SR2	

Why does NIC latency matter?



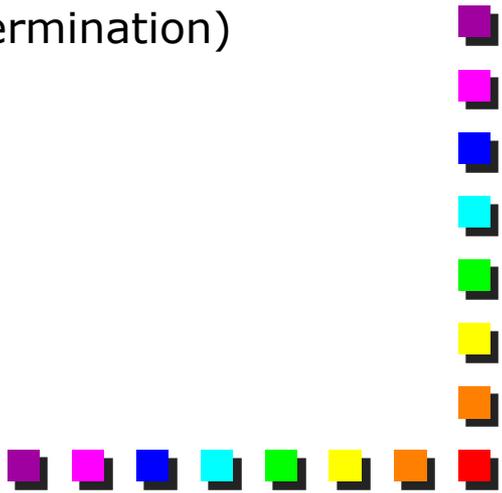
Total latency for 10GBaseT: $2.5\mu\text{s} \times 6 = 15\mu\text{s}$ (150K bit times, i.e., ~ 12.5 maximum size Ethernet frames)

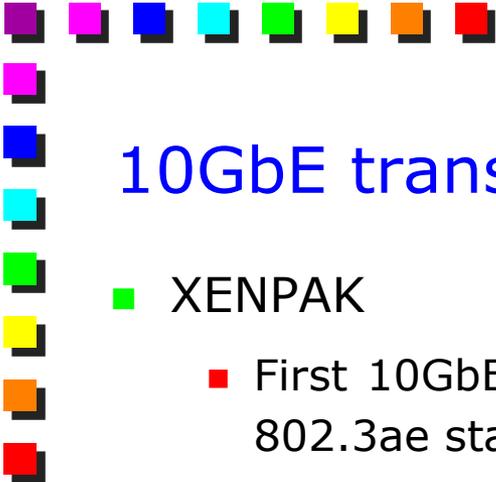
However, 10GBaseT covers only the access side while backbones use fibers, hence the total latency is usually far less.



10G BaseT vs SFP+ (twinax)

■ SFP+

- Snorther ranges
 - Suitable to Top of the Rack switches
 - Better latency (x10)
 - Better power consumption
 - Compatible with fiber optics cabling (just need to replace the SFP+ module)
 - Currently (2013) less expensive
 - Need to use pre-connected cabling (no in-field termination)
 - Not compatiblw ith lower speed (only 10Gbps)
- 

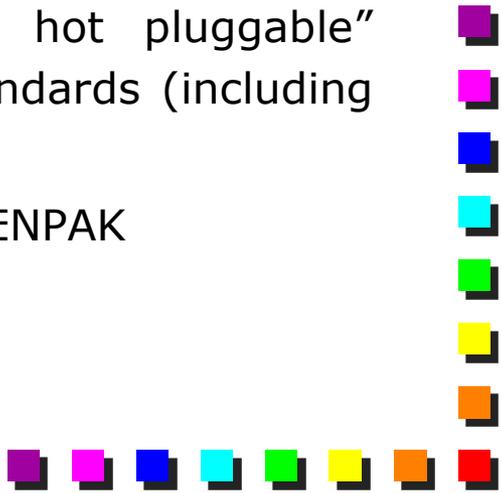


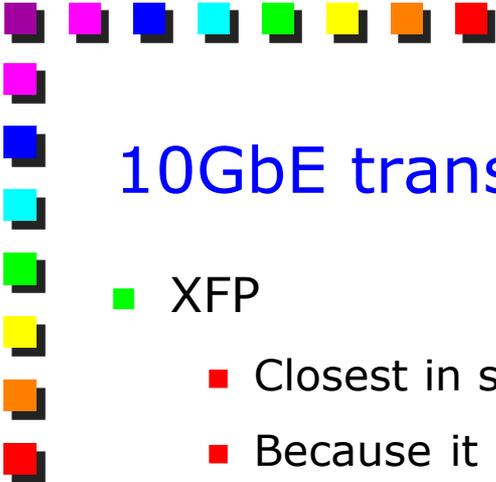
10GbE transceivers (1)

■ XENPAK

- First 10GbE pluggable transceivers on the market to support the 802.3ae standard transmission optics.
- Large, bulky and used mainly in LAN switches.
- “Hot pluggable” and support the new 802.3ak copper standard with vendors now producing transceivers to connect CX4 cables

■ X2

- Smaller brother of the XENPAK pluggable transceivers
 - About 2/3 the size of the XENPAK, same hot pluggable” specifications and supporting all the 10GbE standards (including copper)
 - Allows for more port density on switches than XENPAK
- 

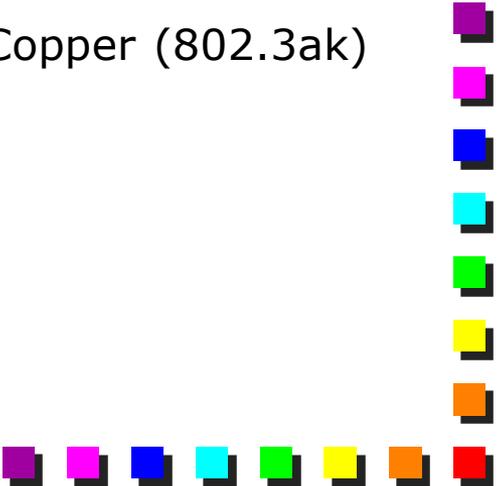


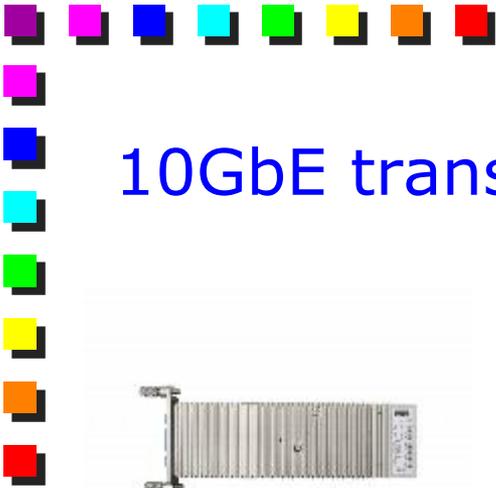
10GbE transceivers (2)

■ XFP

- Closest in size to the SFP pluggable transceiver
- Because it relies on a high-speed interface (10.3125Gbps), high-priced serializer/deserializer (SERDES) are required inside the switch to support it
- Over time, the cost of such SERDES will decline, but today they add an unacceptable cost to the base system. Still, the positive aspect of the XFP form factor is it will allow switch vendors to increase port density in a smaller area for cost savings
- Drawback is its inability to support the current Copper (802.3ak) or the 10GBASE-LX4 standards

■ SFP+

- Very compact
 - Allows 1G or 10G to reside in the same footprint
- 



10GbE transceivers (3)



XENPAK

SC Duplex connector



X2

SC Duplex connector



XPF

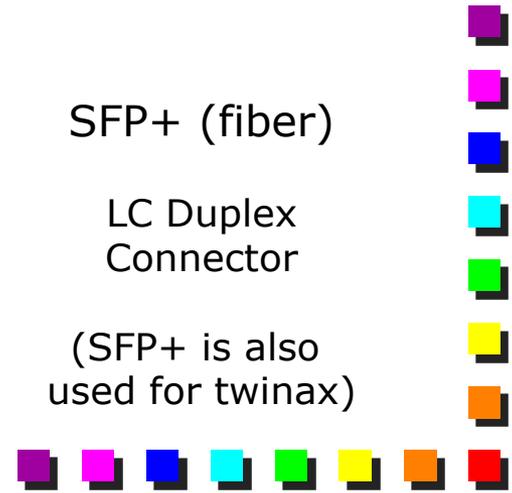
LC Duplex connector

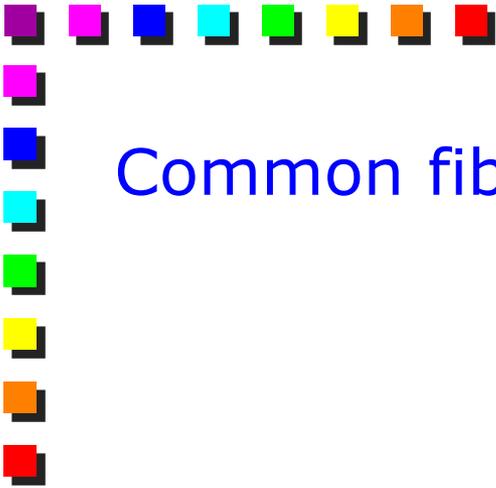


SFP+ (fiber)

LC Duplex Connector

(SFP+ is also used for twinax)





Common fiber optics connectors



LC (duplex) connector



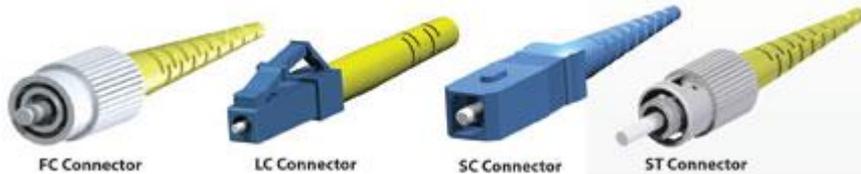
ST Connector
[known as "Straight Tip"; bayonet]



LC (duplex) connector
[known as "Little connector"; high density; very common for SFP and SFP+]



SC (duplex) connector
[known as "standard connector"; used in passive optical networks and for GBICs]

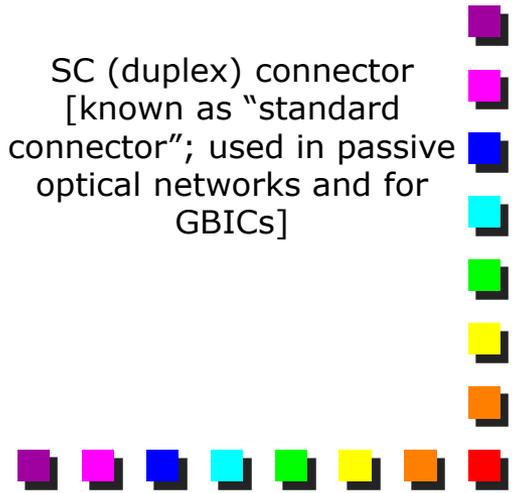


FC Connector

LC Connector

SC Connector

ST Connector

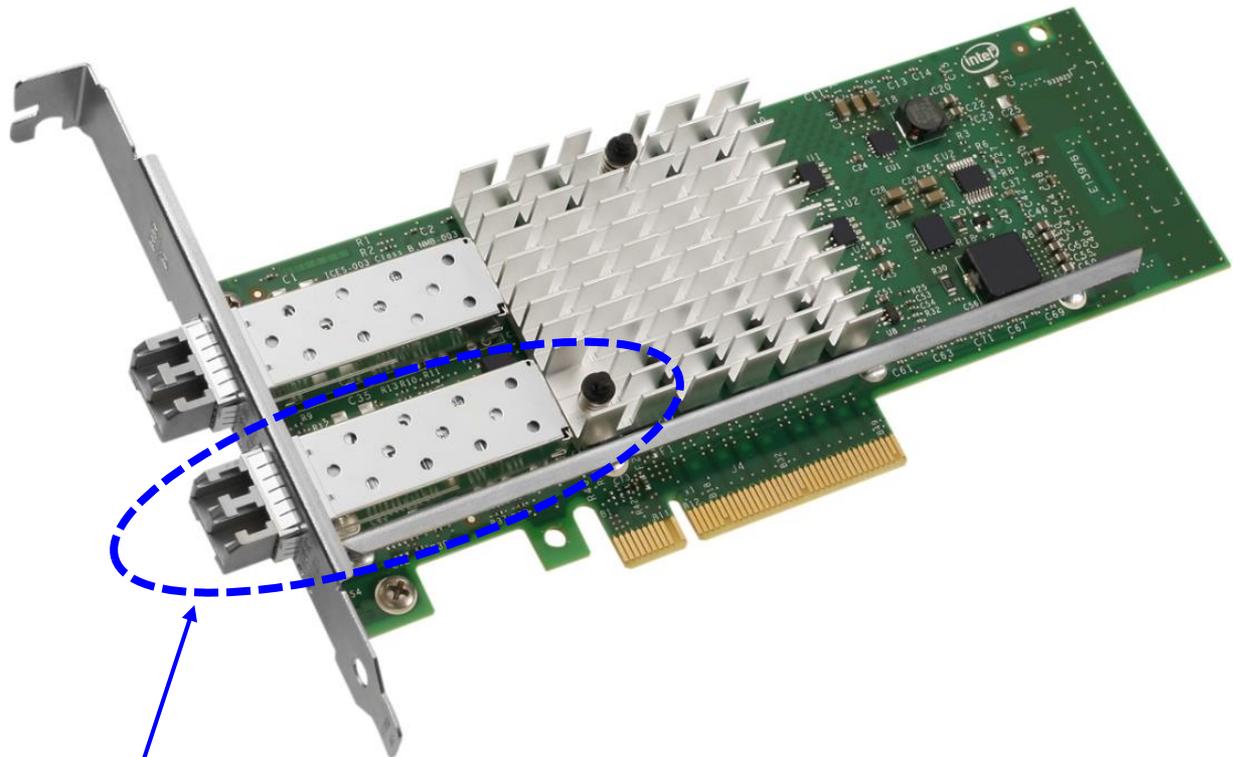


10 Gigabit Ethernet NICs: examples from Intel

	X520-DA2	X520-T2	X520-SR2	X520- LR1
Ethernet controller	Intel® 82599ES	Intel® 82599ES	Intel® 82599ES	Intel® 82599ES
Connector and cable	SFP+ direct attach copper	RJ-45 copper	Fiber optic	Fiber optic
Cable type	SFP+ direct attached twin axial cabling up to 10 m	Category-6 up to 55 m; Category-6A up to 100 m	MMF 62.5/50 µm up to 300 m	SMF up to 10 km
Slot type, maximum bus speed & slot width	PCI Express 2.0 5.0 GT/s lane x 8 lane	PCI Express 2.0 5.0 GT/s lane x 8 lane	PCI Express 2.0 5.0 GT/s lane x 8 lane	PCI Express 2.0 5.0 GT/s lane x 8 lane
Ports	Dual port	Dual port	Dual port	Single port
Supported slot height/s	Low profile and full height			
Intelligent offloads	Yes	Yes	Yes	Yes
Intel® Virtualization Technology (Intel® VT) for connectivity	On-chip QoS and traffic management, flexible port partitioning, Virtual Machine Device queues (VMDq), PCI-SIG* SR-IOV-capable	On-chip QoS and traffic management, flexible port partitioning, Virtual Machine Device queues (VMDq), PCI-SIG* SR-IOV-capable	On-chip QoS and traffic management, flexible port partitioning, Virtual Machine Device queues (VMDq), PCI-SIG* SR-IOV-capable	On-chip QoS and traffic management, flexible port partitioning, Virtual Machine Device queues (VMDq), PCI-SIG* SR-IOV-capable
Storage over Ethernet	iSCSI, FCoE, NFS	iSCSI, FCoE, NFS	iSCSI, FCoE, NFS	iSCSI, FCoE, NFS
iWARP/RDMA	NA	NA	NA	NA
Estimated street price (2013)	400€ + 50€ each cable	550€	900€	1100€ (single port)



Intel Converged Network Adapter X520-SR2 (MMF)

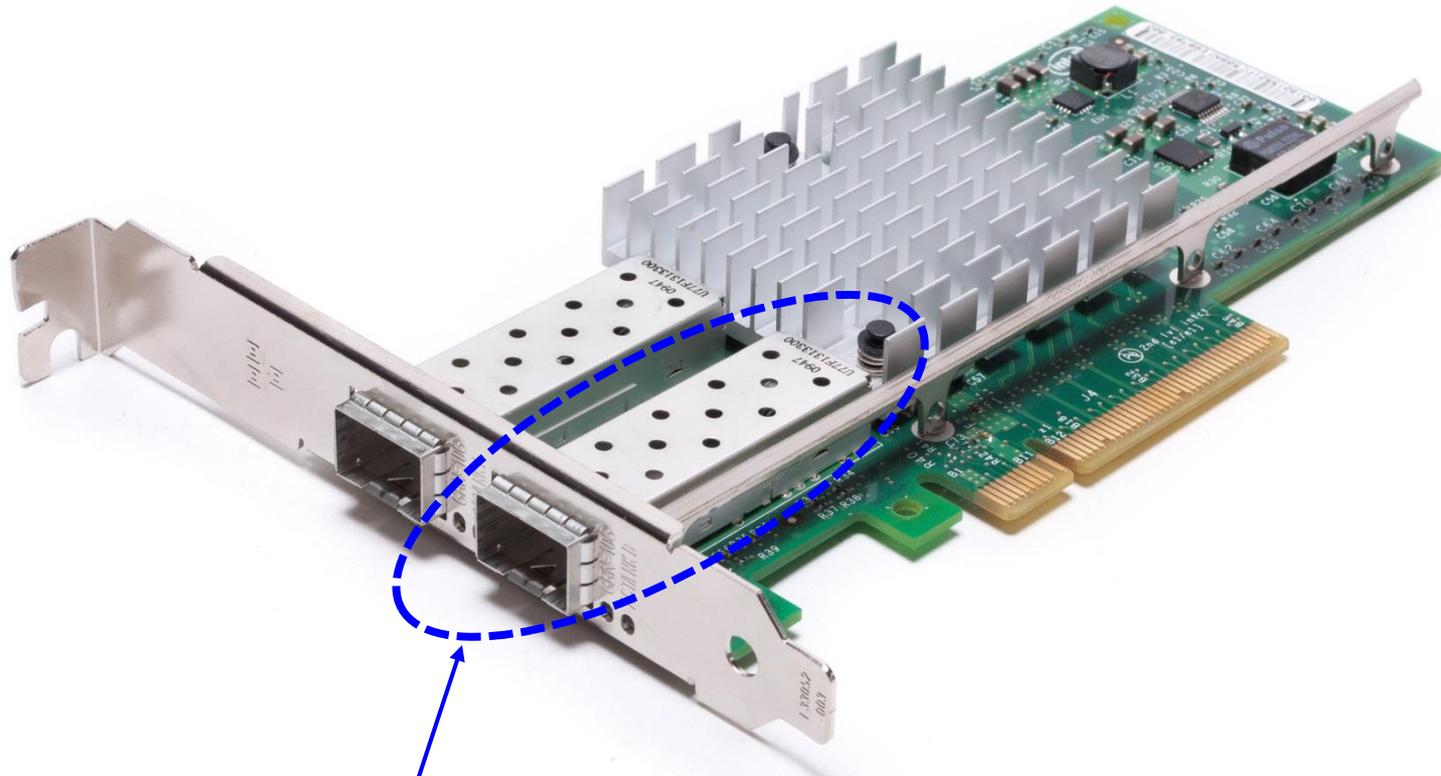


SFP+ fiber optic module (already plugged in the adapter)



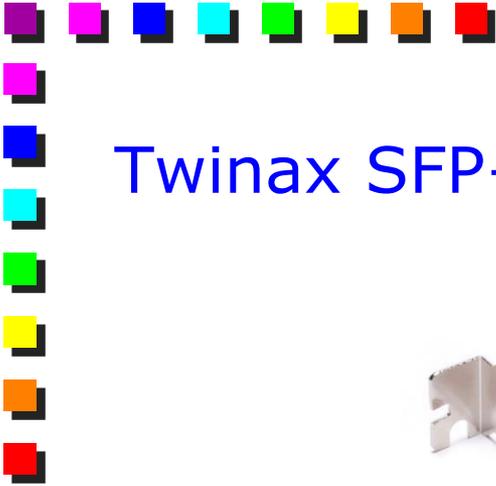


Intel Converged Network Adapter X520-DA2 (twinax)



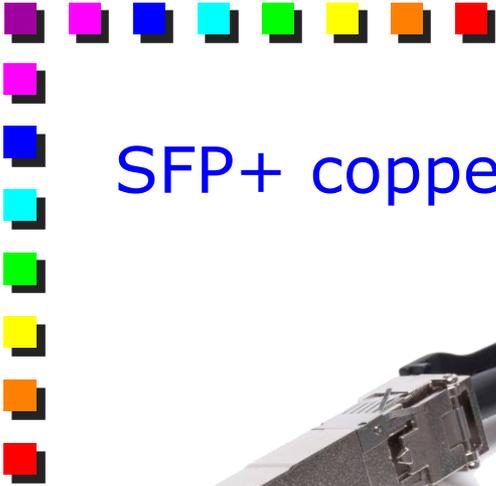
SFP+ fiber optic slot
(needs to plug-in a SFP+ module)





Twinax SFP+ module





SFP+ copper vs optic



SFP+ optic

SFP+ copper (twinax) cable

