Finisar

Application Note AN-2036

Frequently Asked Questions Regarding Finisar's 1000BASE-T SFPs (FCxx-8520/1-3)

Finisar's 1000BASE-T SFP transceivers (FCMJ-8520/1-3 and FCLF-8520/1-3) are based on the SFP Multi Source Agreement (MSA). They are compatible with Gigabit Ethernet and 1000BASE-T standards as specified in IEEE Std. 802.3:2002.

The FCMJ-8520/1-3 versions are RoHS 5/6 whilst the FCLF-8520/1-3 versions are fully RoHS compliant.

This Application Note covers the most commonly asked questions about Finisar's 1000BASE-T SFPs.

1. **Are there special electrical supply considerations when using the FCxx-8520/1-3?**
   Yes. The FCxx-8520/1-3 consumes a maximum of 1.20W under worst-case conditions (typical is ~ 1.05 W). The maximum current draw is 375 mA, compared to 300 mA for a standard SFP. The electrical supply of the host should be examined to ensure that it can handle these increased demands.

2. **Will the extra length of the FCxx-8520/1-3 cause mechanical interference problems in my system?**
   While there is a chance of interference problems, the probability is small. If you compare the length of an FCxx-8520/1-3 to a fiber SFP, with their respective cables inserted, there is little difference in length. This is due to the strain relief features on LC connectors that protrude much farther than the insulation on a Cat 5 cable. An optical and 1000BASE-T SFP with cables inserted is shown in Picture 1 below.

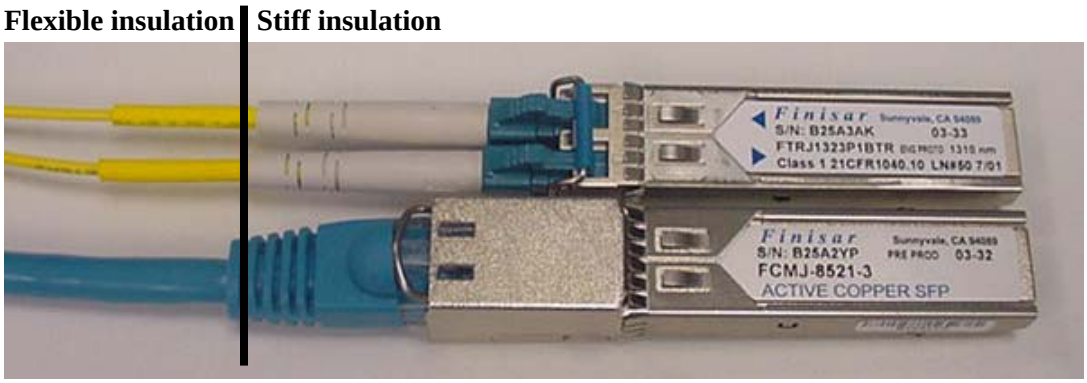

Picture 1: 1000BASE-T and optical SFP transceivers with cables inserted

3. **What cable type is recommended for use with the FCxx-8520/1-3?**
   The FCxx-8520/1-3 was designed to operate using standard Cat 5 cable that has been configured per IEEE802.3:2002. You can safely use Cat 5e or Cat 6 cables, as these are improved versions of Cat 5. For short cable runs, with intricate routing, stranded cable is recommended, as it is more flexible.

4. **What is the PHY and how can it be accessed?**
   The Finisar FCxx-8520/1-3 uses the Marvell 88E1111 Rev. B0 Physical Layer IC (PHY) to convert between the serial interface and 1000BASE-T interfaces. This chip has a number of useful features available on internal registers that can be accessed via the 2-wire bi-directional serial interface at address 0xAC. Each register is 2 bytes wide and details for accessing the register can be found at http://www.marvell.com .

5. **What is SGMII mode?**
   SGMII is a mode of communication between the MAC and PHY to allow for 10/100/1000BASE-T operation. In 100BASE-TX mode, the MAC still transmits to the PHY at 1.25 Gb/sec, but each byte is repeated 10 times. The PHY then converts this repeated data to 100BASE-TX format. The process is the same in 10BASE-T mode but each byte is repeated 100 times.

6. **How do you configure the module for 10/100/1000BASE-T operation including RxLOS functionality at all three data rates?**
   Only the FCxx-8520-1 has RxLOS functionality, for details on how to enable tri-rate functionality please refer to the table below. The FCxx-8520/1-3 can be used with a SGMII rev. 1.5 interface (without clocks).This interface supports 10, 100 and 1000 BASE-T modes of operation, as mentioned above. The table below shows how to enable SGMII and advertise all speeds and full/half-duplex using register writes to the PHY over the 2-wire serial interface (see Question 4). A simple power cycle will return the module to default operation.

PHY Address: 0xAC		
Register Address	**Write Data**	**Description**
0x1B	0x9084	Enable SGMII mode
0x09	0x0F00	Advertise 1000BASE-T Full/Half-Duplex
0x00	0x8140	Apply Software Reset
0x04	0x0DE1	Advertise 100/10BASE-T Full/Half-Duplex
0x00	0x9140	Apply Software Reset
For RxLOS pin of FCxx-8520-3 to function at 10/100/1000BASE-T		
0x18	0x4108	Enable RxLOS pin at 10/100/1000BASE-T

7. **What is auto-negotiation?**
   Auto-Negotiation is the communication or handshake between two remote devices to determine if the two devices can transfer data to one another and, if so, the specifics of data transfer such as data rate, flow control, and duplex traffic.

8. What does 1000BASE-T or 1000BASE-X mean?

This designation is defined in IEEE802.3:2002. The "1000" in the designation refers to the transmission speed of 1000 Mbps. The "BASE" refers to BASE band signaling, indicating that only Ethernet signals are carried on the medium. The "T" represents twisted-pair copper cable (for example Cat 5), and the "X" represents fiber optic cable.

9. What is the difference between 1000BASE-T and 1000BASE-X auto-negotiation?

1000BASE-T auto-negotiation is conducted over the Cat 5 cable between the two 1000BASE-T devices (see Figure 1 below). 1000BASE-X auto-negotiation is typically conducted between two host systems over fiber. In the case where FCxx-8521-3's are installed in the host systems, the 1000BASE-X auto-negotiation information is used to set the configuration options the FCxx-8521-3 advertises during 1000BASE-T auto-negotiation (see Question 11 below).

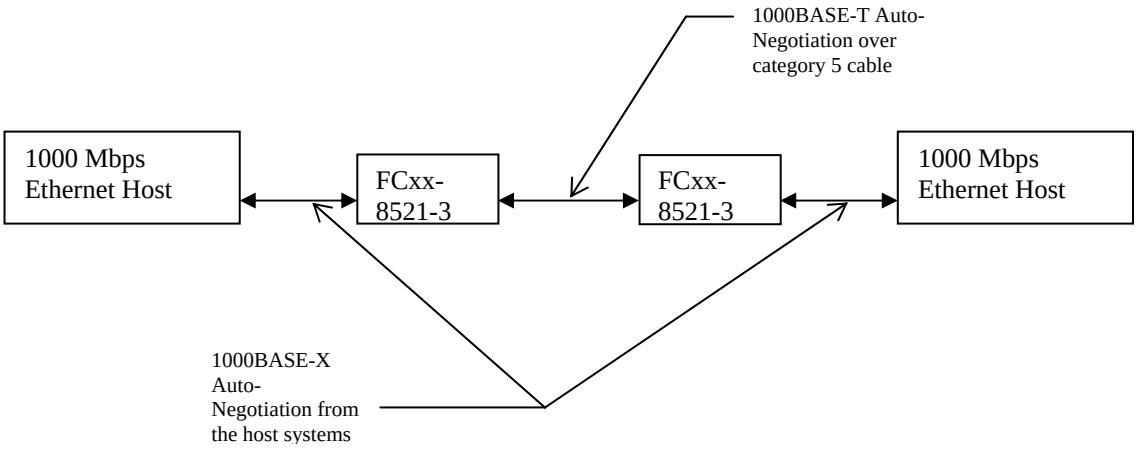

Figure 1: Link configuration of the FCxx-8521-3

10. What is the difference between the FCxx-8521-3 and the FCxx-8520-3?

The FCxx-8521-3 is designed to be a 1000BASE-X compliant device that functions properly in any slot designed for optical SFP's, with no hardware or software changes to the host system. The PHY is configured to perform 1000BASE-X auto-negotiation with the host system (see Question 11). The FCxx-8521-3 does not have link detection circuitry—the RxLOS pin is internally grounded.

The FCxx-8520-3 is similar to the FCxx-8521-3, but the 1000BASE-X auto-negotiation is disabled. For this device to operate, auto-negotiation should also be disabled in the host system. In the FCxx-8520-3, the RxLOS pin functions as a link indicator—RxLOS is asserted when the 1000BASE-T link is lost. Typically, this device is used in systems where the host cannot determine the link status without using RxLOS as a link indicator.

11. Does the FCxx- 8521-3 support 1000BASE-X auto-negotiation from the host system?

The 1000BASE-X auto-negotiation from the host systems is supported and should be enabled when using the FCxx-8521-3. This is not the case for the FCxx-8520-3, in which auto-negotiation must be disabled in the host system or a link cannot be established.

12. How do I disable 1000BASE-X auto-negotiation on the FCxx-8521-3?

PHY Address: 0xAC		
Register Address	**Write Data**	**Description**
0x16	0x0001	Select Fiber Register Bank
0x00	0x8140	Disable Auto-negotiation
0x16	0x0000	Return to Cu Register Bank

13. How do I enable 1000BASE-X auto-negotiation on the FCxx-8520-3?

If the system using an FCxx-8520-3 is compatible with 1000BASE-X auto-negotiation, this feature can be enabled as shown in the below table. Please note that depending on how the RxLOS signal is used by the system, enabling this feature might cause an interface problem with the host system.

PHY Address: 0xAC		
Register Address	**Write Data**	**Description**
0x16	0x0001	Select Fiber Register Bank
0x00	0x9140	Enable Auto-negotiation
0x16	0x0000	Return to Cu Register Bank

14. Does the host system need to support 1000BASE-T auto-negotiation?

No, the 1000BASE-T auto-negotiation is fully supported by either module (FCxx-8520/1-3). The FCxx-8521-3 uses the 1000BASE-X auto-negotiation information it receives from the host to adjust the configuration options that it advertises during copper auto-negotiation.

15. When using the FCxx-8521-3, does the host know that it is driving a copper transceiver?

No, to the host system, Finisar's FCxx-8521-3 will appear to be a fiber transceiver. From the host's perspective, acknowledgements received during the 1000BASE-X auto-negotiation process are coming from the remote link partner even though they are coming from the FCxx-8521-3.

16. When using the FCxx-8521-3, what keeps the 1000BASE-X auto-negotiation from finishing before the 1000BASE-T auto-negotiation?

The FCxx-8521-3 will hold back acknowledgement to the host until the copper auto-negotiation (1000BASE-T AN) is resolved, and then it will send the 1000BASE-X auto-negotiation advertisements and acknowledgement information. This way, the host can complete 1000BASE-X auto-negotiation

without knowing copper auto-negotiation was involved. This process will slow down 1000BASE-X auto-negotiation, because it will be delayed until the 1000BASE-T auto-negotiation is complete (<3 seconds).

17. How are abilities resolved between the host and the FCxx-8521-3 during auto-negotiation?

The module PHY is able to incorporate the host system's abilities in the cable-side auto-negotiation. More specifically, the module PHY first receives the abilities from the host system during 1000BASE-X auto-negotiation, but the 1000BASE-X auto-negotiation is not yet allowed to complete and no abilities are passed to the host system. The PHY then incorporates the host system's abilities into the 1000BASE-T auto-negotiation. Finally, the result of the 1000BASE-T auto-negotiation is passed on to the host system and the 1000BASE-X auto-negotiation is allowed to complete. The resolution of the 1000BASE-T auto-negotiation will therefore be compatible with the resolution of the 1000BASE-X auto-negotiation.

18. Since RxLOS is grounded on the FCxx-8521-3, what indication is there to determine whether the link is functioning or not?

If the host is receiving idles, /I/ ordered sets, from the FCxx-8521-3 then the link is good. If the host is receiving auto-negotiation coded words, /C/ ordered sets, or random data then the link is down. You can also access the PHY at register 0x11, bit 10 for real time link indication. 1=Link up, 0=Link down.

19. Your data sheet indicates that the FCxx-8520/1-3 is set to "Preferred Master" mode. What does this mean?

In 1000 BASE-T, one of the link partners becomes the master, and the other becomes the slave. The master uses its local clock to transmit data on the Cat 5 cable, while the slave uses the clock recovered from the data received from its link partner. During 1000BASE-T auto-negotiation, the link partners agree to 10/100/1000 Mb/s operation, and determine who is the master and who is the slave. Forcing master or slave mode can cause conflicts, so FCxx-8520/1-3's are set to preferred master mode. If both link partners advertise the same preference, a pseudo random number generator will determine the master/slave choice.

20. The FCxx-8520/1-3 data sheet indicates that TX_FAULT is not supported. What does this mean?

Pin 2 on the SFP connector is specified as the TX_FAULT output. In optical SFP's, the TX_FAULT pin is often used to indicate a possible eye-safety condition, and the pin is driven high under a set of transmitter conditions specified by the manufacturer. For electrical SFP's, this functionality is of no value. The TX_FAULT pin is permanently grounded on the FCxx-8520/1-3.

21. Can you provide greater detail on PHY registers?

The Marvell datasheet for the 88E1111 is confidential, and you must register at the Marvell extranet at http://www.marvell.com to gain access.

22. Is external loopback possible for this module, and do you have any other recommendations to aid in testing and debugging?

Although the PHY does have a test mode for external loopback, the regular operation of the PHY must be severely modified to work with an external loopback cable. By design all 1000BASE-T PHYs cancel the effects of their own transmissions from the received signals, also called Near End Cross Talk (NEXT) canceling. If line A is connected to line B, the crosstalk between the lines becomes 100%, and the PHY will automatically cancel the entire signal. Therefore, to enable external loopback, the NEXT canceling must be disabled, which can then make debugging, and testing ineffective.

Finisar recommends using line loopback and internal loopback for testing and debugging purposes in 1000BASE-T mode. This mode still requires that the PHY is reconfigured, but not in a way that could disguise problems in the system.

23. How does Line Loopback mode work?

Line loopback allows a link partner to send frames into the PHY to test transmit and receive data paths. Frames sent from a link partner into the PHY, before reaching the MAC interface pins are looped back and sent out on the line side. The link partner thus receives its own frames.

To enable line loopback, the Finisar FCxx-8520/1-3 must first establish copper link with another link partner. If auto-negotiation is enabled, both link partners should advertise the same speed and full duplex. If auto-negotiation is disabled, both link partners need to be forced to the same speed and full duplex. Once link is established, enable the line loopback mode by writing to register 0x14 bit 14.

0x14 bit 14 = 1 (Enable line loopback)
0x14 bit 14 = 0 (Disable line loopback)

Once the line loopback is enabled, the link partner can send data into the PHY.

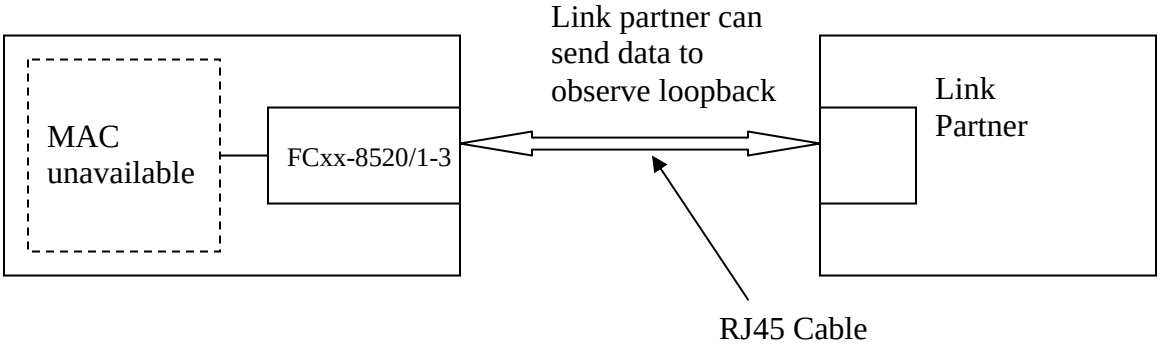

Figure 2: Line Loopback Setup

24. How do you configure the module for internal loopback testing?

Internal loopback testing is a quick way to check the integrity of the MAC to SFP connections. Data sent from the host system to the SFP is looped back internally in the SFP's PHY (data sent in to the TX pin of the SFP are looped out onto the RX pins). This requires register writes to the PHY to put it in loopback mode. For the internal loopback operation, fiber auto-negotiation should be disabled, both on the PHY and on the MAC. This is because the SFP mode of operation requires a handshake between the fiber auto-negotiation on the MAC side, and the copper auto-negotiation on the RJ45 side. Since the loopback mode automatically disables the receive functionality on the copper side, it will be impossible to carry out this handshake that is required by the SFP mode of operation.

To enable loopback, first disable 1000BASE-X auto-negotiation per Question 12. Then, set register 0x00 bit 14 = 1.

After the loopback test is completed, internal loopback can be disabled by setting register 0x00 bit 14 = 0 and 1000BASE-X auto-negotiation can be re-enabled per Question 13.

25. Are there any known problems with the Marvell PHY??

At the time of writing there are two problems that are associated with the current rev. B0 silicon of the PHY.

a. In SGMII mode, if the link partner is a 10Mb repeater hub, the Marvell PHY may potentially send frames to the MAC with an alignment error. This issue does not occur with a NIC or switch in 10Mb mode. It will only occur with link partners that are 10Mb repeater hubs.
b. Gigabit template testing of point A/B and C/D symmetry specified by the IEEE test mode may be slightly outside the limits specified in IEEE 802.3ab. Minor symmetry violations have no effect on performance, other than a possible effect in cable length performance. Finisar's 1000Base-T transceivers are tested with cables >100m in length, to make sure the parts operate error free at the max cable length.

26. What is the initialization sequence to optimize the pulse shape and improve BER?

PHY Address: 0xAC		
Register Address	**Write Data**	**Description**
0x1d	0x0006	Designate internal register
0x1e	0x4200	Write to internal register
0x1d	0x000A	Designate internal register
0x1e	0x0001	Write to internal register
0x00	0x9140	Software reset

27. Can the SGMII auto-negotiation be disabled?

Once the module is in SGMII mode, the SGMII auto-negotiation can be disabled using the register writes shown in the table below.

PHY Address: 0xAC		
Register Address	**Write Data**	**Description**
0x16	0x0001	Select register bank for host system side
0x00	0x8140	Turn off auto-negotiation and reset
0x16	0x0000	Return to register bank for cable side

28. When two modules are connected across the serial interface, are there any additional register writes required?

In this configuration, the data flows from the RD+/- lines of one module to the TD+/- lines of the other module, and vice versa.

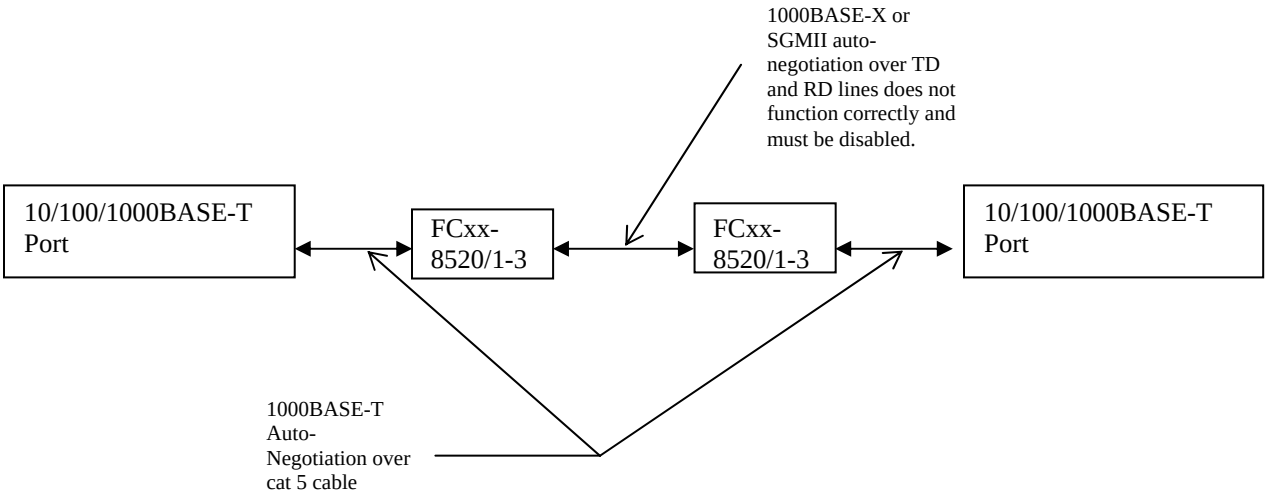

Figure 3: Connecting two modules across the serial interface

When using the FCxx-8521-3 in 1000BASE-T mode, the 1000BASE-X auto-negotiation must be disabled because of the intertwining of the 1000BASE-T and the 1000BASE-X auto-negotiations (see Question 12). Because the resolution of the 1000BASE-T auto-negotiations might result in the modules' having incompatibilities, the abilities of each module should be forced to disable flow control and only enable full-duplex. And, when the 1000BASE-T link is disconnected from one of the modules, the other 1000BASE-T has no way of knowing. The system must handle this situation by monitoring the PHY registers (see Question 18).

When using the FCxx-8520-3 in 1000BASE-T mode in this configuration, each module will link to its partner over the cat 5 cable because the 1000BASE-X auto-negotiation is already disabled (see Question 10). Data will pass between the modules without additional register writes. Again, when the 1000BASE-T link is disconnected from one of the modules, the other link has no way of knowing, and

the system must handle this situation by either monitoring the RxLOS pin or the PHY registers (see Question 18).

When using the FCxx-8520/1-3 in SGMII mode, the SGMII auto-negotiation must be disabled (see Question 27). Again, the abilities of each module should be forced to disable flow control and only enable full-duplex. Again, when the link across the cat 5 cable has been disconnected from one of the modules, the other link has no way of knowing. Finally, modules can be configured for operation at 10/100/100BASE-T (see Question 6)

For more details about monitoring the PHY registers, setting the abilities and speeds of each module, please see the Marvell documentation (see Question 21).

29. Can the PHY registers be accessed when the module TX_DISABLE is asserted?

No. Asserting TX_DISABLE puts the module into its hardware reset state, and the PHY cannot be accessed until the TX_DISABLE has been negated and the reset has completed.