F i n i s a r

Application Note AN-2036

Frequently Asked Questions Regarding Finisar's 1000BASE-T SFPs (FCMJ-8521-3)

Finisar's 1000BASE-T GBIC transceivers (FCMJ-8521-3) are based on the SFP Multi Source Agreement (MSA). They are compatible with Gigabit Ethernet and 1000BASE-T standards as specified in IEEE Std. 802.3ab.

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 FCMJ-8521-3?**
   Yes. The FCMJ-8521-3 consumes a maximum of 1.20W under worst-case conditions (typical is < 1.0 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 FCMJ-8521-3 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 FCMJ-8521-3 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.

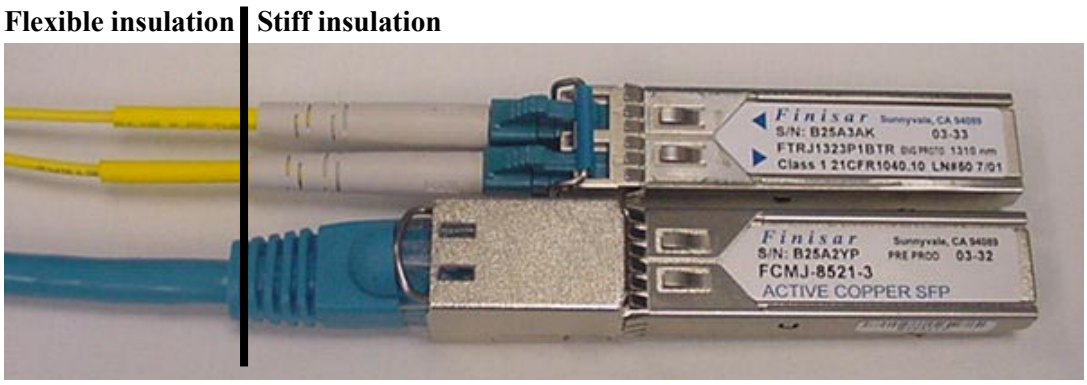

Picture 1: 1000BASE-T and optical SFP transceivers with cables inserted

3. What cable type is recommended for use with the FCMJ-8521-3?
The FCMJ-8521-3 was designed to operate using standard Cat 5 cable. 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. For longer cable runs (>20m), solid-core cable is recommended, where the lower resistance of this cable type is more appropriate.

4. Does the FCMJ-8521-3 support digital diagnostics as defined for SFPs?
Serial identification information at address A0h is supported per SFF-8472. However, the enhanced digital diagnostics functions are not supported. In addition, certain diagnostics functions unique to 1000BASE-T transceivers can be accessed at address 0xAC. See questions 5 and 18 for more information.

5. What is the PHY and how can it be accessed?
The Finisar FCMJ-8521-3 uses the Marvell 88E1111 Physical Layer IC (PHY) to convert between the SERDES 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.

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

7. How do you configure the module for 10/100/1000BASE-T operation?
The FCMJ-8521-3 can be used with a SGMII (without clocks) Rev. 1.5 interface. This interface supports 10, 100 and 1000 BASE-T modes of operation, as mentioned above. To enable SGMII mode, and advertise the proper speed, you will need to access the PHY via the 2 wire bi-directional serial interface. Please refer to Table 1 for details on enabling SGMII and advertising the proper speed/duplex. A simple power cycle will return the module to normal operation. *Note: you can only advertise one speed/duplex setting.*

PHY Address: 0xAC		
Register Address	**Write Data**	**Description**
0x1B	0x9084	Enable SGMII mode
0x09	0x0E00	Advertise 1000BASE-TX Full-Duplex
0x09	0x0D00	Advertise 1000BASE-TX Half-Duplex
0x04	0x0C81	Advertise 100BASE-TX Full-Duplex
0x04	0x0C41	Advertise 100BASE-TX Half-Duplex
0x04	0x0C21	Advertise 10BASE-TX Full-Duplex
0x04	0x0C11	Advertise 10BASE-TX Half-Duplex
0x00	0x8140	Software reset to allow changes to take effect

Table 1: Configuration settings for SGMII mode.

6. What are key features of the FCMJ-8521-3?

The FCMJ-8521-3 is designed to be a fiber-port-compliant device that functions properly in any slot designed for optical SFPs, with no hardware or software changes. This is accomplished by enabling hardware strapping options of the PHY during manufacture that allows 1000BASE-X auto negotiation to proceed normally (this is explained in more detail below). The FCMJ-8521-3 does not have link detection circuitry - the RX_LOS pin is internally grounded.

7. What is auto-negotiation?

Auto-Negotiation is the communication or handshake between two remote devices to first 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 an Institute of Electrical and Electronics Engineers (IEEE) shorthand identifier. 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 FCMJ-8521-3s are installed in the host systems, the 1000BASE-X auto-negotiation information is used to set the configuration options the FCMJ-8521-3 advertises during 1000BASE-T auto-negotiation (see Question 10 below).

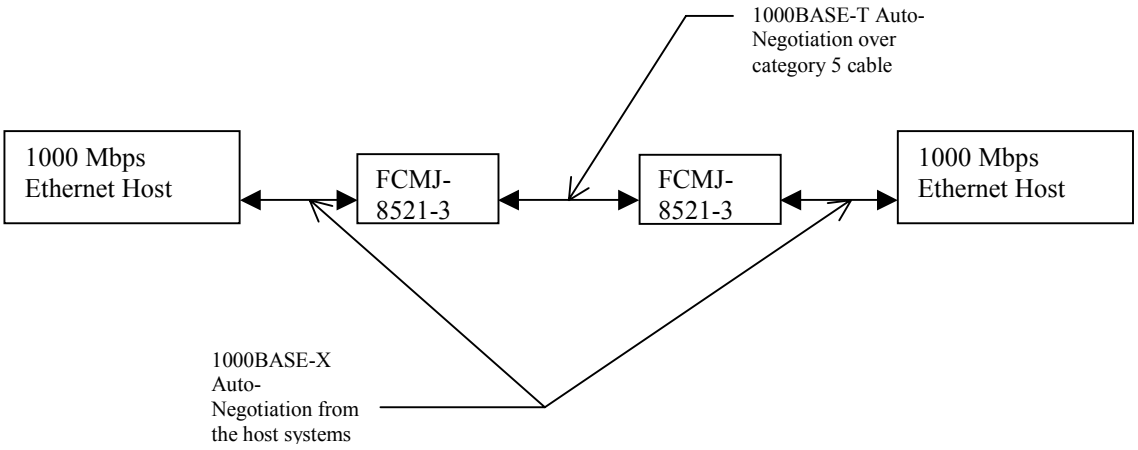

Figure 1: Link configuration of the FCMJ-8521-3

10. Does the FCMJ- 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 FCMJ-8521-3.

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

No, the 1000BASE-T auto-negotiation is fully supported by the FCMJ-8521-3. The FCMJ-8521-3 uses the fiber auto-negotiation (1000BASE-X AN) information it receives from the host to adjust the configuration options that it advertises during copper auto-negotiation.

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

No, to the host system, Finisar's FCMJ-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 FCMJ-8521-3.

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

The FCMJ-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 (4-8 seconds).

14. How are abilities resolved between the host and the FCMJ-8521-3 once copper auto-negotiation is complete?

Once the copper side completes auto-negotiation (1000BASE-T AN), the abilities received are compared with the abilities received from the host. If the abilities are the same, the fiber auto-negotiation (1000BASE-X AN) final handshake is performed and the link is established. If the abilities are not the same, the fiber auto-negotiation (1000BASE-X AN) is restarted.

15. Since RX_LOS is grounded on the FCMJ-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 FCMJ-8521-3 then the link is good. If the host is receiving auto-negotiation coded words, /C/ ordered sets, or junk 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.

16. Your data sheet indicates that the FCMJ-8521-3 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 the FCMJ-8521-3 is set to preferred master mode. If both link partners advertise the same preference, a pseudo random number generator will determine the master/slave choice.

17. The FCMJ-8521-3 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. Under a set of transmitter conditions specified by the manufacturer, the TX Fault pin is driven high. For optical SFPs, this functionality is often useful in indicating a possible eye-safety condition. For electrical SFPs, this functionality is typically of little value. The TX Fault pin is permanently grounded on the FCMJ-8521-3.

18. 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 **https://www.marvell.com/login/sign_up.do** to gain access.

19. 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.

20. 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. This allows the link partner to receive its own frames.

To enable line loopback, the Finisar FCMJ-8521-3 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.

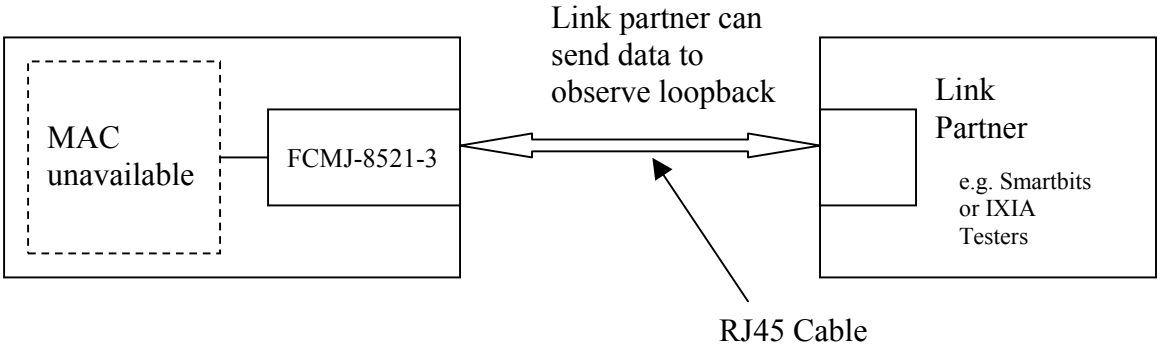

Figure 2: Line Loopback Setup

21. 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 SFPs PHY (data sent to the TX of the SFP is looped out 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, write to the following registers:

Reg 0x14 bit 3 = 0 Disable Fiber Auto-Negotiation
Reg 0x00 bit 15 = 1 Software reset. Needed to update register 20.3 setting
Reg 0x00 bit 14 = 1 Enable loopback.

After loopback test is completed, you can now re-enable fiber auto-negotiation by setting register 0x14 bit 3=1 and soft resetting the PHY.

22. Are there any noteworthy errata items pertaining to the Marvell PHY?
Yes, there are two issues 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.