

LAN9646 Silicon Errata and Data Sheet Clarification

This document describes known silicon errata for functional revisions A1 of the LAN9646. The silicon errata discussed in this document are for silicon revisions as listed in [Table 1](#). The silicon revision can be determined by the device's top marking. A summary of LAN9646 silicon errata is provided in [Table 2](#).

Some errata work arounds may require modifying register values. If the system design does not include a processor to manage the switch, a small 8-bit PIC or AVR microcontroller can be used to configure the switch via the I²C or SPI interface. These low-cost microcontrollers are available in packages as small 8-pins, with integrated oscillator and non-volatile program memory. The microcontroller does not need a MII or RMII connection to the switch. Alternatively, the switch can be configured by a remote computer via the in-band management feature of the switch. The default port for in-band management is port 7.

TABLE 1: AFFECTED SILICON REVISIONS

Part Numbers	Silicon Revision	Package Top Mark
LAN9646	A1	MXX

TABLE 2: SILICON ISSUE SUMMARY

Item Number	Silicon Issue Summary	Affected Silicon Revisions
1.	Register settings are needed to improve PHY receive performance	A1
2.	Transmit waveform amplitude can be improved (1000BASE-T, 100BASE-TX, 10BASE-Te)	A1
3.	Port 6 Default RGMII ingress timing does not comply with the RGMII specification	A1
4.	Energy Efficient Ethernet (EEE) feature select must be manually disabled	A1
5.	Toggling PHY Powerdown can cause errors or link failures in adjacent PHYs	A1
6.	Certain PHY registers must be written as pairs instead of singly	A1
7.	SGMII auto-negotiation does not set bit 0 in the auto-negotiation code word	A1
8.	SGMII port link details from the connected SGMII PHY are not passed properly to the port 7 GMAC	A1
9.	Register settings are required to meet data sheet supply current specifications	A1
10.	Automatic SPI Data Out Edge Select may cause issues	A1
11.	1000BASE-T Transmitter Distortion fails to meet IEEE compliance specification	A1
12.	No Pause frames are generated for ingress rate limiting with an EEE link	A1
13.	Link drop can occur when back pressure is enabled in 100BASE-TX half-duplex mode	A1
14.	Port may hang in half-duplex mode when No Excessive Collision Drop feature is enabled	A1
15.	SGMII registers are not initialized by hardware reset	A1
16.	When tail tag is enabled, frame length field check fails for 802.3 frames	A1
17.	Transmission halt with late collisions	A1
18.	Single-LED Mode Setting Requires Two Register Writes	A1
19.	IBA Must Be Disabled When Using I2C	A1
20.	Transmission halt with Half-Duplex and VLAN	A1

LAN9646

TABLE 2: SILICON ISSUE SUMMARY (CONTINUED)

Item Number	Silicon Issue Summary	Affected Silicon Revisions
21.	Frame Length Field Check feature does not work well when the actual frame length is less than 64 bytes	A1
22.	Flooded unicast frames not limited by ingress rate limiter when limiting mode set to "01"	A1

Silicon Errata Issues

IMPORTANT NOTE

Multiple errata workarounds in this document call for changing PHY registers for each PHY port. PHY registers 0x0 to 0x1F are in the address range 0xN100 to 0xN13F, while indirect (MMD) PHY registers are accessed via the PHY MMD Setup Register and the PHY MMD Data Register.

Before configuring the PHY MMD registers, it is necessary to set the PHY to 100 Mbps speed with auto-negotiation disabled by writing to register 0xN100-0xN101. After writing the MMD registers, and after all errata workarounds that involve PHY register settings, write register 0xN100-0xN101 again to enable and restart auto-negotiation.

[addr]	[data]	
0xN100-0xN101	0x2100	Disable auto-neg, force 100M full-duplex
[PHY register initialization goes here]		
0xN100-0xN101	0x1340	Enable auto-neg, restart auto-neg

Module 1: Register settings are needed to improve PHY receive performance

DESCRIPTION

The default receiver settings are not optimized. Receive errors may occur, especially at longer cable lengths.

END USER IMPLICATIONS

For best receiver performance, users should write the following PHY MMD registers. This is done individually for each port (1-5) using any of the management interfaces: MDC/MDIO, I²C, SPI, or in-band.

Work Around

Before writing the PHY MMD registers, it is necessary to set the PHY to 100 Mbps speed with auto-negotiation disabled by writing to register 0xN100-0xN101. After writing the MMD registers, and after all errata workarounds that involve PHY register settings, write register 0xN100-0xN101 again to enable and restart auto-negotiation. See details in the important note at the beginning of the Silicon Errata Issues section.

Write to the following MMD registers for each PHY port [1-5]:

[MMD]	[register]	[data]
0x01	0x6F	0xDD0B
0x01	0x8F	0x6032
0x01	0x9D	0x248C
0x01	0x75	0x0060
0x01	0xD3	0x7777
0x1C	0x06	0x3008
0x1C	0x08	0x2000 (Note 1)

Note 1: The value of this register may read back as either 0x2000 or 0x2001. Bit 0 is read-only, and is not a fixed value.

PLAN

This erratum will not be corrected in a future revision.

Module 2: Transmit waveform amplitude can be improved (1000BASE-T, 100BASE-TX, 10BASE-Te)

DESCRIPTION

The transmit waveform amplitude can be improved for 10BASE-Te, 100BASE-TX and 1000BASE-T.

END USER IMPLICATIONS

With the default settings, the waveform amplitude may be outside the specifications in some corner case conditions, and the transmitter may not be fully compliant with the IEEE standard. This may degrade performance under some conditions.

Work Around

Before writing the PHY MMD registers, it is necessary to set the PHY to 100 Mbps speed with auto-negotiation disabled by writing to register 0xN100-0xN101. After writing the MMD registers, and after all errata workarounds that involve PHY register settings, write register 0xN100-0xN101 again to enable and restart auto-negotiation. See details in the important note at the beginning of the Silicon Errata Issues section.

Write to the following MMD registers for each PHY port [1-5]:

[MMD]	[register]	[data]
0x1C	0x4	0x00D0

PLAN

This erratum will not be corrected in a future revision.

Module 3: Port 6 Default RGMII ingress timing does not comply with the RGMII specification

DESCRIPTION

The RGMII defining document specifies a typical data-to-clock setup time into the receiver (switch signals TXD6_[3:0], TX_ER6 and TX_EN6 to TX_CLK6) of 1.8ns. However, port 6 requires additional setup time in order to avoid ingress data errors on this interface. Refer to the data sheet for details.

This issue does not occur on port 7, and there are no timing issues with the MII and RMII modes.

END USER IMPLICATIONS

Careful analysis of the RGMII timing must be performed, considering the timing of both connected devices, and relative signal delay times on the PCB.

Work Around

Additional PCB trace delay may be needed on the TX_CLK6 clock signal into port 6 of the switch. Another option is to set the ingress delay bit [4] in register 0x6301. Refer to the data sheet for timing details.

PLAN

This erratum will not be corrected in a future revision

Module 4: Energy Efficient Ethernet (EEE) feature select must be manually disabled

DESCRIPTION

The EEE feature is enabled by default, but it is not fully operational. It must be manually disabled through register controls. If not disabled, the PHY ports can auto-negotiate to enable EEE, and this feature can cause link drops when linked to another device supporting EEE.

END USER IMPLICATIONS

If the link partner does not support EEE, then EEE will not be activated and there will be no problem with link drops. This is also true if auto-negotiation is disabled, since EEE is activated only if auto-negotiated between the two link partners.

If the link partner is not known, or if the link partner is EEE capable, then the EEE feature should be manually disabled to avoid link drop problems.

Work Around

Disable EEE by writing to the following registers. This is done individually for each PHY port (1-5) using any of the management interfaces: MDC/MDIO, I²C, SPI, or in-band.

EEE is disabled by clearing bits [2:1] of the PHY indirect register: MMD 7, address 3Ch.

MMD register:

[MMD]	[register]	[data]
7	0x3C	0x0000

After writing to the MMD register, it is necessary to write to register 0xN100-0xN101 to restart auto-negotiation.

[address]	[data]
0xN100-0xN101	0x1340

PLAN

This erratum will not be corrected in a future revision.

Module 5: Toggling PHY Powerdown can cause errors or link failures in adjacent PHYs

DESCRIPTION

The PHY power down is controlled by bit 11 in registers 0xN100-0xN101. It provides separate power down control for each PHY. When a PHY is brought out of power down by clearing this bit, the resulting power surge can disrupt an adjacent PHY, causing data errors or temporary link down on that port.

END USER IMPLICATIONS

Data errors or link down can occur in an active PHY when its neighbor is brought out of power down mode.

Work Around

Avoid dynamically changing the power down state of any PHYs if other PHYs may be linked and active. Only change the power up or power down state of a PHY when no other PHYs on the chip are linked and possibly passing traffic.

PLAN

This erratum will not be corrected in a future revision.

Module 6: Certain PHY registers must be written as pairs instead of singly

DESCRIPTION

When using SPI, I²C, or in-band register access, writes to certain PHY registers should be performed as 32-bit writes instead of 16-bit writes.

The PHY control and status registers are 16-bit registers. They occupy the byte address range 0xN100 to 0xN13F, where N is the port number. Registers from 0xN100 to 0xN11F function normally and can be written either 16- or 8-bits at a time.

An error in the register access logic causes all writes from 0xN120 to 0xN13F to be 32-bit writes. For example, a 16-bit write to register 0xN122-0xN123 also results in all zeros being written to register 0xN120-0x121. Also, a 16-bit write to register 0xN120-0xN121 causes all zeros to be written to 0xN122-0xN123.

END USER IMPLICATIONS

This issue is relevant only to write operations, not to reads. When writing only 16-bits to a register in this address range, all zeros will be written to the adjacent register. This may change the PHY settings and cause the PHY to malfunction.

Work Around

To avoid writing zero to an adjacent register, always write the registers in this address range in pairs as 32-bits:

0xN120 - 0xN123

0xN124 - 0xN127

0xN128 - 0xN12B

0xN12C - 0xN12F

0xN130 - 0xN133

0xN134 - 0xN137

0xN138 - 0xN13B

0xN13C - 0xN13F

Note that some of these registers are not defined in the data sheet. In order to avoid writing inappropriate data to any undefined register, the register should be read before writing (i.e. read-modify-write). Do not assume that undocumented registers should be all zeros.

PLAN

This erratum will not be corrected in a future revision.

Module 7: SGMII auto-negotiation does not set bit 0 in the auto-negotiation code word

DESCRIPTION

Port 7 SGMII auto-negotiation sends the incorrect code word “0x0000”, instead of “0x0001” as defined in the SGMII specification. If the connected SGMII device requires that bit 0 be set, then this can prevent successful auto-negotiation. Some SGMII devices require that bit 0 is set, while other devices don't care.

END USER IMPLICATIONS

Many SGMII devices ignore the value of the code word received during auto-negotiation, but some will not auto-negotiate properly when an incorrect value is received.

Work Around

If the connected SGMII device does not care about the value of the code word, then this workaround is not needed.

The workaround is to write to port 7 SGMII register 0x1F0004. This only needs to be done once after the chip is powered up or reset.

Using either the SPI, I²C, or in-band management interface, perform the following writes:

[addr]	[data]	
0x7200	0x001F0004	(32-bits)
0x7206	0x01A0	(16-bits)

PLAN

This erratum will not be corrected in a future revision.

Module 8: SGMII port link details from the connected SGMII PHY are not passed properly to the port 7 GMAC

DESCRIPTION

Link information from the connected SGMII PHY is not automatically forwarded from the SGMII interface to the internal PCS block. The GMAC therefore does not automatically know the speed (10/100/1000) and duplex of the link established by the PHY connected to the SGMII interface. If the PCS configuration and the PHY link conditions do not match, the port will not function properly.

END USER IMPLICATIONS

This applies to port 7 operated in the (default) SGMII MAC mode. It is not applicable for (non-default) SGMII PHY mode. The SGMII interface can operate at any of three speeds: 10, 100 or 1000 Mbps. The speed is determined by the link speed of the SGMII PHY attached to port 7 and is communicated by auto-negotiation to port 7. Though this information is received at port 7, it is not forwarded to the port 7 PCS.

The GMAC defaults to 1000 Mbps, so if the SGMII link is always 1000 Mbps, then there is no problem and a workaround is not needed.

If the SGMII speed is variable or less than 1000 Mbps, software must actively read the link details from one register and write it to another register whenever a link up condition occurs or when there is a change to the link speed.

Work Around

Software must actively read the link status (up or down), speed and duplex from one SGMII register and write it to another SGMII register whenever the link conditions change. Two methods are available for detecting a change in link speed:

1. Polling. Software reads SGMII indirect register 0x1f8002.
2. SGMII AN Complete interrupt.

The link speed is read from the SGMII indirect register 0x1F8002, bits [3:2]:

10 = 1000 Mbps
01 = 100 Mbps
00 = 10 Mbps

The speed must be written to SGMII indirect register 0x1F0000, bits [6, 13]:

10 = 1000 Mbps
01 = 100 Mbps
00 = 10 Mbps

Use either the SPI, I²C or in-band management interface to access the indirect SGMII registers. Because they are indirect, two writes are needed to write to an SGMII register. An SGMII read is accomplished by a write followed by a read:

[addr]	[data]	
Write 0x7200	0x001F8002	set up to read reg 1F8002
Read 0x7206	0xFFFF	read reg 1F8002 (16-bits)
Write 0x7200	0x001F0000	set up to write reg 1F0000
Write 0x7206	0x1140	write reg 1F0000

When writing indirect register 0x1F0000, use the following values for these link speeds:

0x1140 for 1000 Mbps full duplex
0x3100 for 100 Mbps full duplex
0x1100 for 10 Mbps full duplex

PLAN

This erratum will not be corrected in a future revision.

Module 9: Register settings are required to meet data sheet supply current specifications

DESCRIPTION

The power supply current specifications in the data sheet are based on the following register settings. Without these register changes, the power consumption will be somewhat higher during auto-negotiation, and in 10BASE-T and 100BASE-TX modes. There is no change to 1000BASE-T power consumption, assuming that EEE is disabled.

END USER IMPLICATIONS

If the following register settings are not made, the chip can dissipate more power than indicated in the data sheet for the AVDDH and AVDLL power rails when in modes other than 1000BASE-T.

Work Around

Write to the following MMD registers for each PHY port [1-5]:

[MMD]	[register]	[data]
0x1C	0x13	0x6EFF
0x1C	0x14	0xE6FF
0x1C	0x15	0x6EFF
0x1C	0x16	0xE6FF
0x1C	0x17	0x00FF
0x1C	0x18	0x43FF
0x1C	0x19	0xC3FF
0x1C	0x1A	0x6FFF
0x1C	0x1B	0x07FF
0x1C	0x1C	0x0FFF
0x1C	0x1D	0xE7FF
0x1C	0x1E	0xEFFF
0x1C	0x20	0xEEEE

PLAN

This erratum will not be corrected in a future revision.

Module 10: Automatic SPI Data Out Edge Select may cause issues

DESCRIPTION

Automatic SPI Data Out Edge Select is a feature that is normally enabled in register 0x0100. It detects the SPI clock frequency and selects either the rising clock edge or falling clock edge to clock out the SPI data based on that frequency. The behavior is not fully predictable when the SPI clock frequency is near 25MHz, which may cause the SPI interface to stop functioning. Also, it does not adapt to changes in the SPI clock frequency.

END USER IMPLICATIONS

The SPI interface may stop functioning if the inappropriate clock edge is selected, or if the clock rate changes. Generally there is no problem when operating at lower clock rates, such as below 15 MHz.

Work Around

When operating the SPI above 15 MHz, it is suggested to disable the automatic feature by clearing register 0x0100 bit 1, and at the same time setting bit 0 to the desired value to manually select the mode of operation.

PLAN

This erratum will not be corrected in a future revision.

Module 11: 1000BASE-T Transmitter Distortion fails to meet IEEE compliance specification**DESCRIPTION**

The device's 1000BASE-T Transmitter Distortion is approximately 40mV, versus the <10mV indicated in the IEEE specification.

END USER IMPLICATIONS

It is unlikely this specification failure will impact system performance. The following link to the Gigabit Transmit Distortion Testing document on the IEEE802.org website also questions the validity of this measurement:

http://www.ieee802.org/3/axay/public/may_07/sefidvash_1_0507.pdf

IEEE testing calls for <10mV peak transmitter distortion for at least 60% of the UI within the eye opening. However, this measurement might not be valid, as the transmit distortion test is sensitive to transmit jitter. Refer to the explanation below, taken from the aforementioned IEEE document.

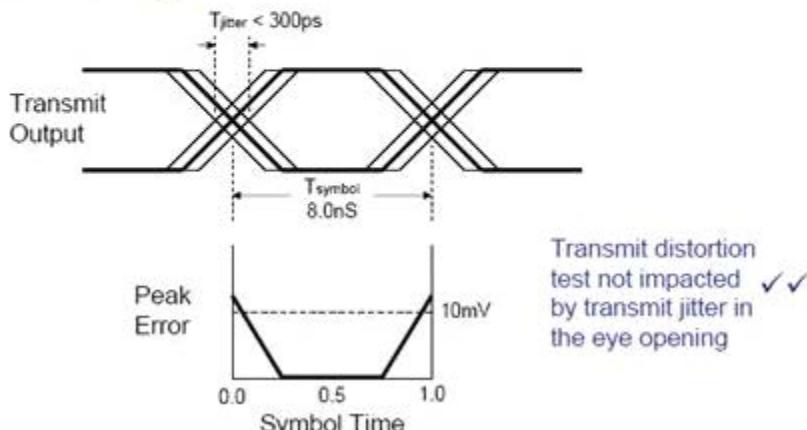
The Gigabit Transmit Distortion Testing document indicates:

- On page 6, a contradiction between Transmit Jitter and Transmit Distortion requirements:

FIGURE 1: IEEE DOCUMENT PAGE 6

Consistency Between Transmit Jitter and Transmit Distortion Specifications

- PHY operating with allowed transmit jitter will not be able to meet 10mV distortion spec *if the distortion spec must be met at the edges*



- On page 7:
 - The transmit distortion test is sensitive to transmit clock jitter during the rise/fall time.
 - It is recommended to change the requirement to use at least 30%, instead of at least 60%, of the UI within the eye opening for the <10mV peak transmitter distortion.

FIGURE 2: IEEE DOCUMENT PAGE 7

Defining the “Settled” Interval for Transmit Distortion Testing

- Currently defined transmit distortion test is sensitive to transmit clock jitter during the rise/fall time of the transmitter
 - Error voltage will be contaminated by jitter during transitions
 - Portion of error contributed by distortion cannot be determined during transitions
 - Appropriate place to apply test is after the rise/fall time where the waveform has settled to its final value
- Clause 40.6.1.2.3 specifies a 5ns rise/fall time (note #3)
 - 3ns of 8ns (37.5%) of UI will be free from effects of jitter
 - Recommend to use 30% for ease of testing
 - Ie: measure error voltage at 10 phases, require 3 of these measurements to be below 10mW

Extensive testing has been performed to ensure the device can inter-operate with different Gigabit PHY link partners.

Work Around

None.

PLAN

This erratum will not be corrected in a future revision.

Module 12: No Pause frames are generated for ingress rate limiting with an EEE link

DESCRIPTION

When an Energy Efficient Ethernet (EEE) link is established with another device, and ingress rate limiting is set up, the port may not generate Pause frames in response to ingress traffic exceeding the rate limit. It also assumes that the Ingress Rate Limit Flow Control Enable bit in register 0xN403 has been set.

END USER IMPLICATIONS

If Pause frames are not generated by the switch, then the link partner will not be able to regulate the rate at which it sends traffic, making ingress rate limiting ineffective. Note that when the flow control function is enabled for ingress rate limiting, the ingress port will not drop packets when the rate limit is exceeded - it relies only on flow control for limiting the ingress rate.

Work Around

The problem can be resolved by writing to the following three global registers. Note that these registers may not be documented in the data sheet.

Global registers:

[addr]	[data]
0x03C0	0x4090
0x03C2	0x0080
0x03C4	0x2000

PLAN

This erratum will not be corrected in a future revision.

Module 13: Link drop can occur when back pressure is enabled in 100BASE-TX half-duplex mode

DESCRIPTION

When back pressure is enabled for 100BASE-TX half duplex mode, CRS-based back pressure is the default mode. In this mode, if the switch forwards long packets and the link partner is set up to detect and respond to jabber, then the link partner may drop link. The link down condition is temporary.

END USER IMPLICATIONS

If all of the above conditions are met, then this problem is likely to occur, which will be disruptive, even though it is self-healing. If any of the above conditions are not present, then the problem will not occur. In general, half-duplex is not common. It is also very uncommon for NICs or switches to implement jabber-based link drop since they are normally full-duplex. This function is seen mostly in hubs, which are half duplex.

Work Around

The workaround is to change the back pressure mode from CRS-based to collision-based by clearing bit 5 in register 0x0331. This completely eliminates the link drop problem. This register can be written using the SPI, I²C, or in-band management interface, but not via the MIIM interface.

Global register:

[addr]	[data]
0x0331	0xD0

PLAN

This erratum will not be corrected in a future revision.

Module 14: Port may hang in half-duplex mode when No Excessive Collision Drop feature is enabled

DESCRIPTION

This issue is seen when two of these switch devices are connected together and are configured in the same way. The two devices use the same back-off algorithm and their back-off can become synchronized, causing lock-up.

END USER IMPLICATIONS

“No excessive collision drop” is a feature that may be used for half duplex to potentially improve collision performance. It is controlled in register 0x0300, and is disabled by default. If it is enabled and the link partner is a similar device, then lock up can occur on the link. It will persist until the link is broken (either physically or by register) and re-established.

Work Around

The problem is avoided by enabling the Alternate back-off algorithm when using the No excessive collision drop feature. This is done by setting bit 7 in register 0x0330. This register can be written using the SPI, I²C, or in-band management interface, but not via the MIIM interface.

PLAN

This erratum will not be corrected in a future revision.

Module 15: SGMII registers are not initialized by hardware reset

DESCRIPTION

The reset pin RST_N and the Soft Hardware Reset control bit do not reset the SGMII registers.

END USER IMPLICATIONS

When asserted, reset pin RST_N and Soft Hardware Reset bit (register 0x300, bit 1) will reset all device registers to their default state, with the exception of the SGMII registers. The SGMII registers are unaffected by RST_N or Soft Hardware Reset, and they retain their settings. An additional step is required to reset the SGMII registers, but only when the user feels that a complete device reset is required. Typically it is never necessary to take this action. The SGMII registers are automatically initialized at power up.

Work Around

If a reset of the SGMII interface is required, set the SGMII reset bit located in SGMII register 0x1F0000, bit 15.

Write the following global registers:

[addr]	[data]	
0x7200	0x001F0000	(32-bits)
0x7206	0x9140	(16-bits)

PLAN

This erratum will not be corrected in a future revision.

Module 16: When tail tag is enabled, frame length field check fails for 802.3 frames

DESCRIPTION

The comparison of the length field of the Ethernet frame with the actual length of the data field portion of the frame fails for the ingress packets with tail tag. This issue is not applicable to packets with the type field in the frame.

END USER IMPLICATIONS

The packets will be dropped when length check fails for the packet.

Work Around

Do not set the length check in bit 3 of the register 0x0330 (Global Switch MAC Control Register 0). The Microchip provided driver disables the length check.

PLAN

This erratum will not be corrected in a future revision.

Module 17: Transmission halt with late collisions

DESCRIPTION

Section 4 of the IEEE 802.3 Specification details Carrier Sense Multiple Access / Collision Detection (CSMA/CD) parameters when operating in half-duplex mode. The first 512 bit times are designated as the slotTime, which is the maximum amount of time allowed for a collision to occur. If a link partner is configured incorrectly, where the PHY is linking in half-duplex mode but the MAC is configured in full-duplex mode, there is a chance that the link partner will generate a collision after the first 512 bit times, violating the IEEE 802.3 specification. These late collisions, combined with other factors, can cause the switch port transmitter to lock up and stop sending packets. The receiver will still function.

END USER IMPLICATIONS

If this erratum occurs, the switch will stop transmitting data to the half-duplex port, making it seem the half-duplex link partner has stopped communicating to the network. The more traffic there is, the greater the risk of the violating link partner generating a late collision that will affect the port.

Work Around

Ideally, the link partner that is violating the specification would need to be updated so the MAC and PHY are correctly configured to the same duplex setting. If the link partner cannot be modified to conform to the IEEE 802.3 specification, the switch can be re-configured to full-duplex when late collisions are detected to avoid a lock up condition. Of note, each switch port functions independently. Therefore, any work around must be implemented separately for each port.

Method 1:

To avoid transmitter lock up, when a port is linked in half-duplex mode, the software should monitor the TxLateCollision MIB counter (MIB Index 0x16). If the number is ever non-zero, the software should force the link to function in full-duplex mode by disabling auto-negotiation and setting full-duplex and the appropriate speed in the PHY Basic Control Registers (addresses 0xN100 - 0xN101).

Method 2:

To detect transmitter lockup, see the work around section of [Module 20: Transmission halt with Half-Duplex and VLAN](#).

PLAN

This erratum will not be corrected in a future revision.

Module 18: Single-LED Mode Setting Requires Two Register Writes

DESCRIPTION

The PHY Port LEDx_0 pin does not go low in the presence of link activity when Single-LED Mode is selected in the MMD LED Mode Register.

END USER IMPLICATIONS

When Single-LED Mode is selected, an additional bit in a different register must also be written. When this is done, the LED pins will function correctly. When Tri-Color Dual-LED Mode (the default mode) is selected, do not set the additional register bit.

Work Around

To fully enable Single-LED Mode:

- Single-LED mode is selected by setting bit 4 in the MMD LED Mode Register (PHY MMD address 2, register 0). (Write 0x11 to this register.)
- Implement the workaround by writing 0xfa00 to register 0xN13C-0xN13D (PHY register 0x1E). This sets bit 9, which corrects Single-LED mode operation for the LEDx_0 pin.

Note: Due to a separate errata module about needing to perform 32-bit writes to certain PHY registers, this needs to be a 32-bit write of data 0xfa00_0300 to registers 0xN13C-0xN13F.

Each port has its own LED mode setting.

PLAN

This erratum will not be corrected in a future revision.

Module 19: IBA Must Be Disabled When Using I²C

DESCRIPTION

When using I²C for device management, it may not function correctly if In-Band Management (IBA) is enabled. In particular, problems can occur when accessing PHY registers.

END USER IMPLICATIONS

If IBA is enabled while using I²C to access PHY registers, reads and writes may not occur correctly, and the I²C may become non-functional. If this occurs, the chip may need to be reset. This problem can be avoided by following the work-around. It is possible to alternate between IBA and I²C, but IBA must be disabled whenever I²C is in use.

Work Around

There is a pin configuration strap option to enable / disable IBA. Do not strap it high (enabled) unless IBA will actually be used.

IBA can also be enabled / disabled via the In-Band Management (IBA) Control Register (address 0x0104 - 0x0107). This register bit indicates the status of the configuration strap, and is used to override it.

The best advice is to enable IBA only if it is being used. Otherwise leave it disabled. Make sure it is disabled if I²C is in use.

PLAN

This erratum will not be corrected in a future revision.

Module 20: Transmission halt with Half-Duplex and VLAN

DESCRIPTION

A port can stop transmitting if it is operating in half-duplex mode and 802.1Q VLAN is enabled. The half-duplex link can be either 10BASE-T or 100BASE-TX. Note that 1000BASE-T is always full-duplex.

Depending on the traffic at the time of lockup, the problem may only show itself as blocked transmission on that port, or it may cause all packet forwarding through the switch to be blocked. Depending on packet size and traffic rate in each direction, the time for the problem to occur is variable. It may take occur after a few seconds, or it may occur only after hours or days of operation.

The problem has not been observed with full-duplex connections, nor when 802.1Q VLAN is disabled (see Switch Lookup Engine Control 0 Register (address 0x0310)).

Transmission of flow control (Pause) packets is not affected by this issue.

END USER IMPLICATIONS

When a port stops transmitting due to this issue, it remains stopped until the switch is reset. It causes traffic to back up within the switch, which consumes resources both within the transmit queue for that port, and in the main switch packet memory. If transmit queue resources are exhausted first, then the only impact of the problem is that traffic does not egress the effected port, and packets are forwarded normally to other ports. However, if the switch packet memory resources are exhausted first, then the switch will no longer forward packets between any ports, and all subsequent packets received by the switch will be dropped.

forwarding of packets to other ports may or may not be affected.

There is no direct indication that the issue has occurred. It can only be detected indirectly.

Work Around

Avoidance: As described above, this problem can be avoided by not using 802.1Q VLAN on a half-duplex port.

Detection: There is no interrupt or status bit which will indicate that the problem has occurred. The problem can be detected by monitoring the following MIB Counters which are available for each port:

- RxDropPackets
- TxDropPackets

In normal operation, both counters should always be zero for each port. If the problem occurs and the transmit queue of the blocked port is exhausted before the main packet memory, then the TxDropPackets counter will increase for the blocked port as packets continue to be forwarded to that port but cannot be transmitted. The TxDropPackets counters will remain zero for all other ports. The RxDropPackets counters will remain zero for all ports.

If the problem occurs and the main packet memory is exhausted before the transmit queue of the blocked port, then the RxDropPackets counters will increment for all ports that are receiving traffic. The TxDropPackets counters will remain zero for all ports.

The TxUnicastPkts counter can also be monitored for all ports that may be subject to this issue. The counter will stop incrementing when the problem occurs. However, it will continue to increment if the port is generating flow control (Pause) packets.

The problem may also be detected by monitoring the following resource utilization registers:

- Packet Memory Available Block Count (PMAVBC) in bits [26:16] of global register 0x03AC-0x03Ax
- TX Queue Blocks Used Count (TXQBU) in bits [10:0] of port registers 0xNA10-0xNA13

When the switch is idle, PMAVBC = 0x7e8 and TXQBU = 0.

With normal traffic, these values do not change substantially. PMAVBC will normally be > 0x580, and each TXQBU will be < 0x200. When the problem occurs, these counters will be outside of these limits, and they will remain fixed until the problem is cleared. TXQBU remains within the normal range for ports not directly experiencing lockup.

For example, when the problem occurs and the transmit queue of the blocked port is exhausted before the main packet memory, the following values have occurred: PMAVBC = 0x5e4, TXQBU = 0x204 on the blocked port, and TXQBU = 0 on all other ports.

LAN9646

When the problem occurs and the main packet memory is exhausted before the transmit queue of the blocked port, the following values have occurred: PMAVBC = 0x0060, TXQBU = 0x0788 on the blocked port, and TXQBU = 0 on all other ports.

Clearing the problem: The issue can be cleared either by strobing the RESET_N pin, or by setting and then clearing the Global Software Reset bit in the Global Chip ID 3 Register (address 0x0003). Note that while the Global Software Reset bit does not clear many registers, it does clear the VLAN Table and the Port Default Tag 0 and 1 registers (addresses 0xN000 and 0xN001).

PLAN

This erratum will not be corrected in a future revision.

Module 21: Frame Length Field Check feature does not work well when the actual frame length is less than 64 bytes

DESCRIPTION

The Frame Length Field Check feature will discard any received frame if the actual length does not match the value in the frame length field in the header. Because this test is applied even when the length field value is less than the minimum legal packet length, some legal minimum size packets may get dropped.

END USER IMPLICATIONS

Some frames, such as Spanning Tree Protocol frames, start out very small - for example 0x26 or 0x27 bytes. This length value is inserted into the frame length field of the header when the frame is created. Before it is transmitted, the network driver or MAC will pad it up to the minimum IEEE frame size. However, the frame length field is typically not updated to match.

Because the Frame Length Field Check does not make an exception for the case of small values in the frame length field, it drops these packets, even though they are valid.

Work Around

Do not enable the Frame Length Field Check feature, found in the Switch MAC Control 0 Register.

PLAN

This erratum will not be corrected in a future revision.

Module 22: Flooded unicast frames not limited by ingress rate limiter when limiting mode set to "01"

DESCRIPTION

Flooded unicast frames are not limited by the ingress rate limiter when the limiting mode (bit [3:2] of register 0xN403) is set to "01".

END USER IMPLICATIONS

The ingress rate limiter is located before the MAC address lookup table in the ingress processing pipeline. Therefore, a unicast frame can't be judged as learned or unlearned (flooded) at the rate limiter. The ingress rate limiter performs correctly in mode "10" and "11", where broadcast and multicast frames are limited. The ingress rate limiter also performs correctly in mode "00", where all frames, including learned unicast, unlearned (flooded) unicast, broadcast, and multi-cast, are limited.

Work Around

No work around exists for this erratum.

PLAN

This erratum will not be corrected in a future revision.

APPENDIX A: DOCUMENT REVISION HISTORY

Revision Level & Date	Section/Figure/Entry	Correction
DS80001138A (07-09-24)	All	Initial release.

THE MICROCHIP WEB SITE

Microchip provides online support via our WWW site at www.microchip.com. This web site is used as a means to make files and information easily available to customers. Accessible by using your favorite Internet browser, the web site contains the following information:

- **Product Support** – Data sheets and errata, application notes and sample programs, design resources, user's guides and hardware support documents, latest software releases and archived software
- **General Technical Support** – Frequently Asked Questions (FAQ), technical support requests, online discussion groups, Microchip consultant program member listing
- **Business of Microchip** – Product selector and ordering guides, latest Microchip press releases, listing of seminars and events, listings of Microchip sales offices, distributors and factory representatives

CUSTOMER CHANGE NOTIFICATION SERVICE

Microchip's customer notification service helps keep customers current on Microchip products. Subscribers will receive e-mail notification whenever there are changes, updates, revisions or errata related to a specified product family or development tool of interest.

To register, access the Microchip web site at www.microchip.com. Under "Support", click on "Customer Change Notification" and follow the registration instructions.

CUSTOMER SUPPORT

Users of Microchip products can receive assistance through several channels:

- Distributor or Representative
- Local Sales Office
- Field Application Engineer (FAE)
- Technical Support

Customers should contact their distributor, representative or Field Application Engineer (FAE) for support. Local sales offices are also available to help customers. A listing of sales offices and locations is included in the back of this document.

Technical support is available through the web site at: <http://microchip.com/support>

Note the following details of the code protection feature on Microchip products:

- Microchip products meet the specifications contained in their particular Microchip Data Sheet.
- Microchip believes that its family of products is secure when used in the intended manner, within operating specifications, and under normal conditions.
- Microchip values and aggressively protects its intellectual property rights. Attempts to breach the code protection features of Microchip product is strictly prohibited and may violate the Digital Millennium Copyright Act.
- Neither Microchip nor any other semiconductor manufacturer can guarantee the security of its code. Code protection does not mean that we are guaranteeing the product is "unbreakable" Code protection is constantly evolving. Microchip is committed to continuously improving the code protection features of our products.

This publication and the information herein may be used only with Microchip products, including to design, test, and integrate Microchip products with your application. Use of this information in any other manner violates these terms. Information regarding device applications is provided only for your convenience and may be superseded by updates. It is your responsibility to ensure that your application meets with your specifications. Contact your local Microchip sales office for additional support or, obtain additional support at <https://www.microchip.com/en-us/support/design-help/client-support-services>.

THIS INFORMATION IS PROVIDED BY MICROCHIP "AS IS". MICROCHIP MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WHETHER EXPRESS OR IMPLIED, WRITTEN OR ORAL, STATUTORY OR OTHERWISE, RELATED TO THE INFORMATION INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE, OR WARRANTIES RELATED TO ITS CONDITION, QUALITY, OR PERFORMANCE.

IN NO EVENT WILL MICROCHIP BE LIABLE FOR ANY INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL, OR CONSEQUENTIAL LOSS, DAMAGE, COST, OR EXPENSE OF ANY KIND WHATSOEVER RELATED TO THE INFORMATION OR ITS USE, HOWEVER CAUSED, EVEN IF MICROCHIP HAS BEEN ADVISED OF THE POSSIBILITY OR THE DAMAGES ARE FORESEEABLE. TO THE FULLEST EXTENT ALLOWED BY LAW, MICROCHIP'S TOTAL LIABILITY ON ALL CLAIMS IN ANY WAY RELATED TO THE INFORMATION OR ITS USE WILL NOT EXCEED THE AMOUNT OF FEES, IF ANY, THAT YOU HAVE PAID DIRECTLY TO MICROCHIP FOR THE INFORMATION.

Use of Microchip devices in life support and/or safety applications is entirely at the buyer's risk, and the buyer agrees to defend, indemnify and hold harmless Microchip from any and all damages, claims, suits, or expenses resulting from such use. No licenses are conveyed, implicitly or otherwise, under any Microchip intellectual property rights unless otherwise stated.

Trademarks

The Microchip name and logo, the Microchip logo, Adaptec, AVR, AVR logo, AVR Freaks, BesTime, BitCloud, CryptoMemory, CryptoRF, dsPIC, flexPWR, HELDO, IGLOO, JukeBlox, KeeLoq, Kleer, LANCheck, LinkMD, maXStylus, maXTouch, MediaLB, megaAVR, Microsemi, Microsemi logo, MOST, MOST logo, MPLAB, OptoLyzers, PIC, picoPower, PICSTART, PIC32 logo, PolarFire, Prochip Designer, QTouch, SAM-BA, SenGenuity, SpyNIC, SST, SST Logo, SuperFlash, Symmetricom, SyncServer, Tachyon, TimeSource, tinyAVR, UNI/O, Vectron, and XMEGA are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

AgileSwitch, ClockWorks, The Embedded Control Solutions Company, EtherSynch, Flashtec, Hyper Speed Control, HyperLight Load, Libero, motorBench, mTouch, Powermite 3, Precision Edge, ProASIC, ProASIC Plus, ProASIC Plus logo, Quiet-Wire, SmartFusion, SyncWorld, TimeCesium, TimeHub, TimePictra, TimeProvider, and ZL are registered trademarks of Microchip Technology Incorporated in the U.S.A.

Adjacent Key Suppression, AKS, Analog-for-the-Digital Age, Any Capacitor, AnyIn, AnyOut, Augmented Switching, BlueSky, BodyCom, ClockStudio, CodeGuard, CryptoAuthentication, CryptoAutomotive, CryptoCompanion, CryptoController, dsPICDEM, dsPICDEM.net, Dynamic Average Matching, DAM, ECAN, Espresso T1S, EtherGREEN, EyeOpen, GridTime, IdealBridge, IGaT, In-Circuit Serial Programming, ICSP, INICnet, Intelligent Parallelizing, IntelliMOS, Inter-Chip Connectivity, JitterBlocker, Knob-on-Display, MarginLink, maxCrypto, maxView, memBrain, Mindi, MiWi, MPASM, MPF, MPLAB Certified logo, MPLIB, MPLINK, mSiC, MultiTRAK, NetDetach, Omnipresent Code Generation, PICDEM, PICDEM.net, PICkit, PICtail, Power MOS IV, Power MOS 7, PowerSmart, PureSilicon, QMatrix, REAL ICE, Ripple Blocker, RTAX, RTG4, SAM-ICE, Serial Quad I/O, simpleMAP, SimpliPHY, SmartBuffer, SmartHLS, SMART-I.S., storClad, SQI, SuperSwitcher, SuperSwitcher II, Switchtec, SynchroPHY, Total Endurance, Trusted Time, TSHARC, Turing, USBCheck, VariSense, VectorBlox, VeriPHY, ViewSpan, WiperLock, XpressConnect, and ZENA are trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

SQTP is a service mark of Microchip Technology Incorporated in the U.S.A.

The Adaptec logo, Frequency on Demand, Silicon Storage Technology, and Symmcom are registered trademarks of Microchip Technology Inc. in other countries.

GestIC is a registered trademark of Microchip Technology Germany II GmbH & Co. KG, a subsidiary of Microchip Technology Inc., in other countries.

All other trademarks mentioned herein are property of their respective companies.
© 2024, Microchip Technology Incorporated and its subsidiaries.

All Rights Reserved.

ISBN: 9781668348994

For information regarding Microchip's Quality Management Systems, please visit www.microchip.com/quality.



Worldwide Sales and Service

AMERICAS

Corporate Office
2355 West Chandler Blvd.
Chandler, AZ 85224-6199
Tel: 480-792-7200
Fax: 480-792-7277
Technical Support:
<http://www.microchip.com/support>
Web Address:
www.microchip.com

Atlanta

Duluth, GA
Tel: 678-957-9614
Fax: 678-957-1455

Austin, TX

Tel: 512-257-3370

Boston

Westborough, MA
Tel: 774-760-0087
Fax: 774-760-0088

Chicago

Itasca, IL
Tel: 630-285-0071
Fax: 630-285-0075

Dallas

Addison, TX
Tel: 972-818-7423
Fax: 972-818-2924

Detroit

Novi, MI
Tel: 248-848-4000

Houston, TX

Tel: 281-894-5983

Indianapolis

Noblesville, IN
Tel: 317-773-8323
Fax: 317-773-5453
Tel: 317-536-2380

Los Angeles

Mission Viejo, CA
Tel: 949-462-9523
Fax: 949-462-9608
Tel: 951-273-7800

Raleigh, NC

Tel: 919-844-7510

New York, NY

Tel: 631-435-6000

San Jose, CA

Tel: 408-735-9110
Tel: 408-436-4270

Canada - Toronto

Tel: 905-695-1980
Fax: 905-695-2078

ASIA/PACIFIC

Australia - Sydney
Tel: 61-2-9868-6733
China - Beijing
Tel: 86-10-8569-7000
China - Chengdu
Tel: 86-28-8665-5511
China - Chongqing
Tel: 86-23-8980-9588
China - Dongguan
Tel: 86-769-8702-9880
China - Guangzhou
Tel: 86-20-8755-8029
China - Hangzhou
Tel: 86-571-8792-8115
China - Hong Kong SAR
Tel: 852-2943-5100
China - Nanjing
Tel: 86-25-8473-2460
China - Qingdao
Tel: 86-532-8502-7355
China - Shanghai
Tel: 86-21-3326-8000
China - Shenyang
Tel: 86-24-2334-2829
China - Shenzhen
Tel: 86-755-8864-2200
China - Suzhou
Tel: 86-186-6233-1526
China - Wuhan
Tel: 86-27-5980-5300
China - Xian
Tel: 86-29-8833-7252
China - Xiamen
Tel: 86-592-2388138
China - Zhuhai
Tel: 86-756-3210040

ASIA/PACIFIC

India - Bangalore
Tel: 91-80-3090-4444
India - New Delhi
Tel: 91-11-4160-8631
India - Pune
Tel: 91-20-4121-0141
Japan - Osaka
Tel: 81-6-6152-7160
Japan - Tokyo
Tel: 81-3-6880- 3770
Korea - Daegu
Tel: 82-53-744-4301
Korea - Seoul
Tel: 82-2-554-7200
Malaysia - Kuala Lumpur
Tel: 60-3-7651-7906
Malaysia - Penang
Tel: 60-4-227-8870
Philippines - Manila
Tel: 63-2-634-9065
Singapore
Tel: 65-6334-8870
Taiwan - Hsin Chu
Tel: 886-3-577-8366
Taiwan - Kaohsiung
Tel: 886-7-213-7830
Taiwan - Taipei
Tel: 886-2-2508-8600
Thailand - Bangkok
Tel: 66-2-694-1351
Vietnam - Ho Chi Minh
Tel: 84-28-5448-2100

EUROPE

Austria - Wels
Tel: 43-7242-2244-39
Fax: 43-7242-2244-393
Denmark - Copenhagen
Tel: 45-4485-5910
Fax: 45-4485-2829
Finland - Espoo
Tel: 358-9-4520-820
France - Paris
Tel: 33-1-69-53-63-20
Fax: 33-1-69-30-90-79
Germany - Garching
Tel: 49-8931-9700
Germany - Haan
Tel: 49-2129-3766400
Germany - Heilbronn
Tel: 49-7131-72400
Germany - Karlsruhe
Tel: 49-721-625370
Germany - Munich
Tel: 49-89-627-144-0
Fax: 49-89-627-144-44
Germany - Rosenheim
Tel: 49-8031-354-560
Israel - Ra'anana
Tel: 972-9-744-7705
Italy - Milan
Tel: 39-0331-742611
Fax: 39-0331-466781
Italy - Padova
Tel: 39-049-7625286
Netherlands - Drunen
Tel: 31-416-690399
Fax: 31-416-690340
Norway - Trondheim
Tel: 47-7288-4388
Poland - Warsaw
Tel: 48-22-3325737
Romania - Bucharest
Tel: 40-21-407-87-50
Spain - Madrid
Tel: 34-91-708-08-90
Fax: 34-91-708-08-91
Sweden - Gothenberg
Tel: 46-31-704-60-40
Sweden - Stockholm
Tel: 46-8-5090-4654
UK - Wokingham
Tel: 44-118-921-5800
Fax: 44-118-921-5820