

- Ethernet Media Access Controller (MAC)
 - Hardware peripheral for support of Link layer of Ethernet protocol
 - Embeds its own DMA for automatic dataflow control
 - Support 10, 100, and 1000 Mbps data transfer rates with MII/GMII/RMII and RGMII PHY interfaces

Application benefits

- Full IEEE 802.3 MAC standard compliance
- Enables the efficient development of applications based on TCP/IP model



The peripheral presented in these slides is a Media Access Controller, or MAC, for Ethernet protocol. It is fully compliant with the IEEE 802.3 standard.

With its own DMA and a data rate up to 1 gigabit, the peripheral provides high Ethernet performances.

The peripheral is involved in applications based on Internet networks. Such applications rely on the TCP/IP layer model as presented in the diagram.

The MAC is in charge of the Link layer of TCP/IP communication model.

Upper layers are managed by software. For example, Transport and Internet layers can be managed by the popular Light weight IP stack.

Finally, the physical layer, or PHY, is supported by external components and linked to an RJ45 connector.

The STM32MP1 Ethernet peripheral supports the following features:

- Operation modes and PHY support
 - 10M/100M/1G bit/s data rate
 - Full-duplex and half-duplex operations
 - (G)MII and (G)RMII interface to external PHY
- Processing control
 - Multi-layer packet filtering
 - Double VLAN processing
 - IEEE 1588-2008/PTPv2 support
 - Separate queues (Rx and Tx) for Audio-Video bridging support (AVB)
 - Supports network statistics with RMON/MIB counters (RFC2819/RFC2665)
- Offload processing
 - Preamble and start-of-frame data (SFD) insertion or deletion
 - Checksum checking of IPv4 header and TCP, UDP, or ICMP payload
 - Calculates and inserts IPv4 header and TCP, UDP, or ICMP payload checksums
 - ARP response
 - TCP segmentation
- Low-power mode
 - Remote wakeup packet and AMD Magic Packet™ detection



The key features of the STM32MP1 Ethernet MAC peripheral are presented in this slide.

The peripheral supports both full- and half-duplex modes of operation at either 10 or 100 Mbit/s. Auto-negotiation between the peripheral and the external PHY enables automatic configuration of the operation mode. The external PHY is supported through two interface types: The typical Media-Independent Interface, or MII, and the Reduced-MII interface that needs half the pins of the MII.

Among the advanced features supported by the peripheral, we can list:

- Multi-layer packet filtering
- Management of double VLAN tags
- Precision Timing Protocol support with high precision time-stamping of frames,
- Support of multiple queues for Audio-Video bridging

- Several network statistics registers available to monitor the connection quality.

In addition to the previous features, the peripheral brings several types of heavy processing offloading. It supports automatic management of preamble and start-of-frame tags, checksum checking for received frames and upper layers checksums computation and automatic TCP packets segmentation.

Two low-power modes are supported enabling power consumption saving. The Energy Efficient Ethernet mode for power control during transmission and a Sleep mode where the peripheral is put on hold, waiting to receive special wakeup packets to resume.

Ethernet datagram management overview

• Ethernet datagram offload processing

Preamble & SFD 1	Destination MAC (DA) 2	Source MAC (DA) 3	VLAN 4	Eth Type 5	Payload... 6	CRC 7
---------------------	---------------------------	----------------------	-----------	---------------	-----------------	----------

1. Preamble and Start-of-Frame tag automatic insertion (Rx) or deletion (Tx)
2. Destination MAC address filtering
3. Source MAC address filtering
4. Single or double VLAN tag insertion, replacement and deletion. VLAN-tag filtering.
5. Checks frame type and size (Rx) or Insert field (Tx)
6. Upper layer management
 - Payload Checksum computation and insertion (Tx) or checking (Rx)
 - Layer 4 (UDP/TCP) and Layer 3 (Internet) filtering capabilities
7. Datagram CRC computation (Tx) and checking (Rx)



This slide presents the offload processing managed by the peripheral on an Ethernet datagram.

You can see that most of the non-payload part of the datagram is efficiently managed in hardware.

The preamble and start frame delimiter (SFD) are basic synchronization patterns and are inserted or deleted automatically.

MAC address filtering is recommended to select only the frames that are relevant for your application. The MAC supports multiple filtering options for unicast or multicast address frames and perfect or hash filtering.

Single and double VLAN-tagged frames are supported. Double-tagging is used for routing complex traffic with two VLAN levels used simultaneously.

The Ethernet peripheral supports automatic insertion,

replacement and deletion of any sequence of outer and inner VLAN tags

The payload is composed of data from Transport or Internet layers. The peripheral can filter received frames depending on either port or IP addresses. The checksum is computed or checked for IPv4 headers and the TCP/UDP or the ICMP payload.

Finally, the CRC is computed for the whole datagram without taking into account the preamble and the start-of-frame tag.

Some of the features listed above are detailed in next slides.

Multi-Layer Packet Filtering

5

- The MAC supports following types of Rx packet filters

- MAC filter

- Unicast source (SA) address
- Unicast and multicast destination (DA) address
- Perfect or hash (6-bit index) filter

- VLAN tag-based filter

- Perfect or Hash (4-bit index) filter on C-VLAN and S-VLAN types

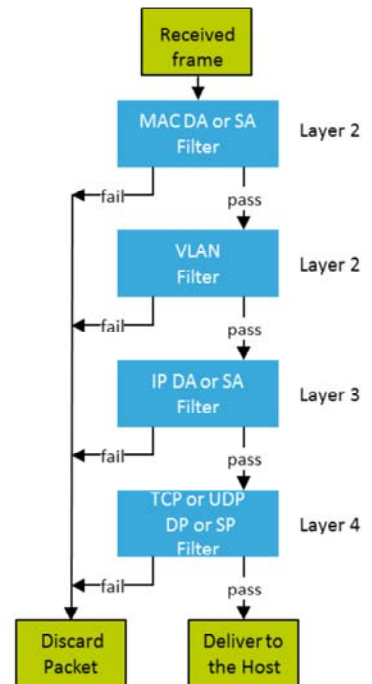
- Layer 3 filter

- IP source (SA) or destination (DA) address
- Perfect matching filter

- Layer 4 filtering


- Source (SP) or destination port (DP)
- Perfect matching filter

- Filters are cumulative from Layer 2 to Layer 4 filtering



The Ethernet MAC peripheral offers some filtering capabilities that can be applied to the received frames. The different filters are nested. Layer 2 is filtered first with the MAC addresses and VLAN tags; then the Internet layer's filter is applied based on IPv4 or IPv6 addresses; and finally, frames are filtered following the port number of UDP and TCP protocols. Any frame rejected by one of the active filters will be discarded and not delivered to the Host. For all layers, perfect or hash filtering is available.

- MAC embeds hardware support to following TCP/IP features
 - Integrity Checksum offload engine for IP header and TCP/UDP/ICMP payload
 - Transmit checksum calculation and insertion
 - Receive checksum calculation and comparison
 - IPv4 Address Recognition Protocol (ARP) offload
 - Automatic ARP request response with the device's MAC address is generated
 - TCP Segmentation Offload
 - Automatic split of large transmit TCP packets into multiple smaller packets
 - Maximum TCP packet size supported is 256 Kbytes
 - Maximum TCP packet size after segmentation is 16 Kbytes



We have seen in the previous slide that filtering capability on layer 3 and layer 4 was available. We will now see that other processing on these upper layers are offered by the peripheral.

Checksums of IPv4 header or checksums of TCP/UDP or ICMP data payload are computed by the hardware. These values are then either transmitted in the output packets or compared to the received ones to detect any error in the transmission.

Another of the Layer 3 offloading features is the automatic ARP protocol response by sending the device's MAC address to the requester without any software action. Automatic TCP packets segmentation is supported in hardware. The peripheral has the ability to split big TCP packets of up to 256 Kbytes into several smaller packets.


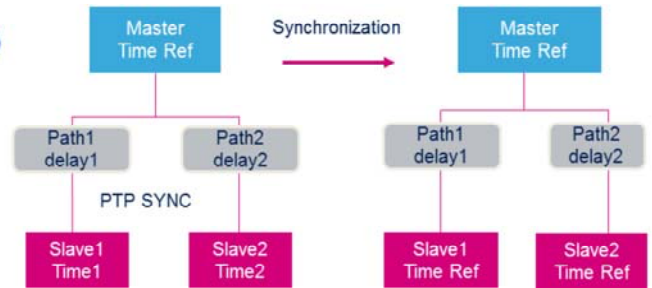
Precision Timing Protocol (PTP)

- PTP objectives

- Synchronize all nodes in a local network area (LAN) with very high accuracy ($< 1 \mu\text{s}$) using HW time stamping
- PTP protocol define synchronization messages between nodes and routers.

- STM32MP1 MAC features

- MAC is compliant with PTPv2 (IEEE 1588-2008) messages
- Accurate timing reference is based on hardware counter
 - 64 internal bits (32 bits for second and 32 bits for nanosecond counter)
 - Counter accuracy on HCLK is down to $\sim 4 \text{ ns}$ (@ 266 MHz)
 - Export timing reference through output pulse-per-second (PPS) signal and through timestamp snapshots triggered on external events



The precision timing protocol has been developed to support high-precision synchronization between several nodes of an Ethernet network. The targeted precision is approximately $1 \mu\text{s}$. This level of precision can only be achieved by hardware support for packet time-stamping.

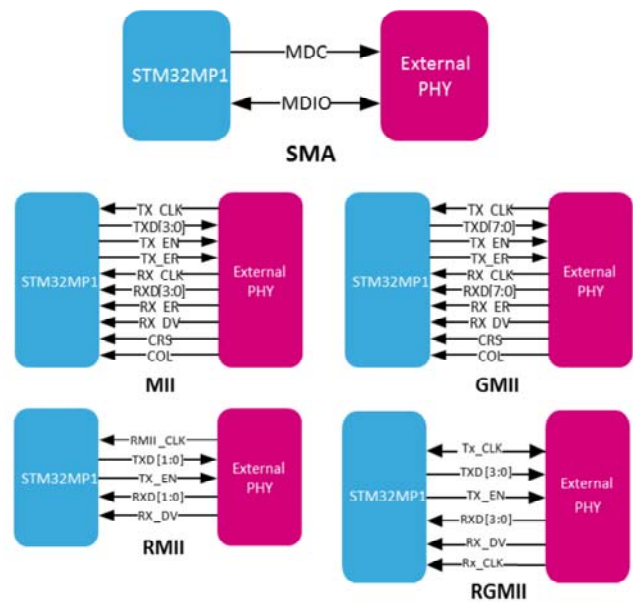
For this purpose, an accurate timing reference is maintained inside the peripheral in a 64-bit register. The reference can be adjusted by standardized synchronization messages between nodes of the network.

The internal reference timing is made available to the system through a pulse-per-second signal that can be output to an internal timer (TIM2 and TIM3) or a GPIO.

Timestamp snapshots are also available on up to 4 external events coming from timers or the CAN interface. These snapshots are stored in a FIFO accessible to the Host.

Media-Independent Interfaces (MII)

- PHY control interface
 - Station Management Agent (SMA)
 - Two-wire MDIO Interface to PHY
 - Enables control of PHY register
 - Compliant with IEEE 802.3 Clause 22
- PHY data interfaces
 - (Giga) Media Independent Interface: (G)MII
 - 16-signal interface for MII
 - 24-signal interface for GMII
 - Reduced (Giga) MII: R(G)MII
 - 7-signal interface for RMII
 - 12-signal interface for RGMII



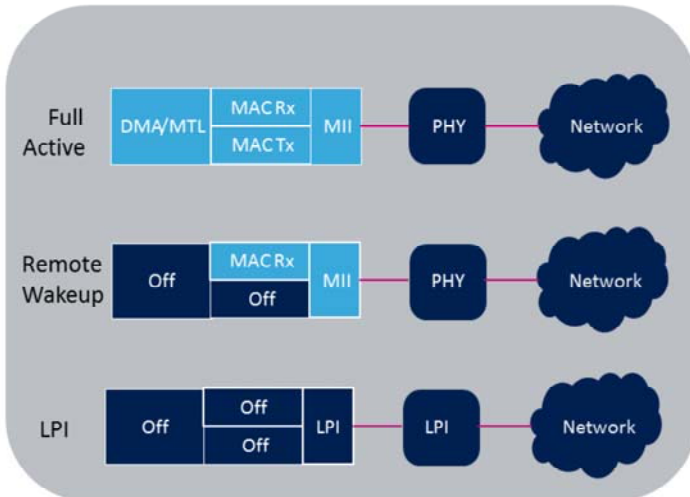
The external PHY is controlled by the peripheral through the Station Management Agent (SMA) that allows read and write access to PHY internal registers. This interface supports the MDIO protocol on a pair of wires. Read and Write operation codes are available.

Several Media Independent Interfaces are supported by the peripheral depending on the selected data rate.

For 10 and 100 Mbits data rate, the peripheral supports the classical Media-Independent Interface, or MII, that requires 16 signals between both devices, and the Reduced-MII that requires only 7 signals and hence allows IO saving.

For 1Gigabits data rate, the peripheral supports the Giga MII and its reduced pin count version of RGMII interface.

Remote wakeup frame detection



1. Remote Wakeup mode

- Application and Tx dock can be switched-off while PHY, MI interface and MAC Rx remain active
- Wake-up is controlled by network at link layer through AMD Magic or user-defined packets
- Wakeup frame detection is an event that can wake up the system from Stop mode

2. Energy-Efficient Ethernet

- EEE mode saves power by switching off the link between the MAC and PHY when no data are transmitted nor received. When no data is transferred, the MAC and PHY are in Low Power Idle (LPI) mode.
- The link wakes up quickly as soon as any data is transferred. There is no packet loss.

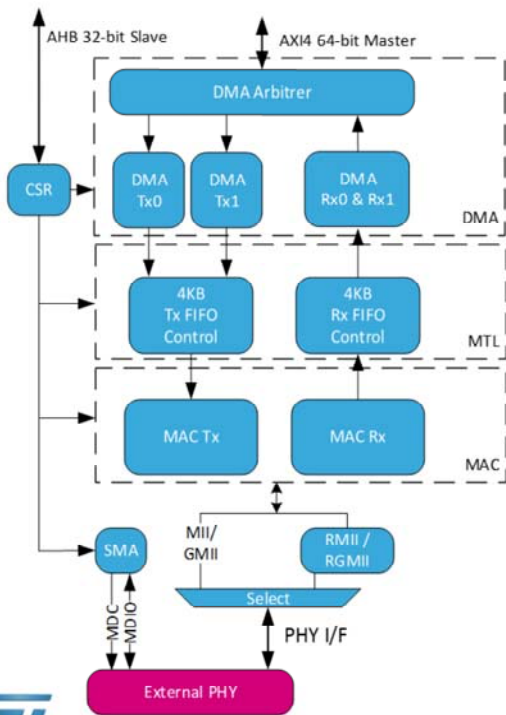
Two functional low-power modes are supported by the Ethernet peripheral.

The first one, the Remote Wakeup mode, puts the peripheral in a state where only the receive path is active, waiting to receive a special packet to wake up. This event can trigger a system Stop mode event.

All frames except the wakeup ones are dropped while the peripheral remains in this mode.

The second low-power mode, the Energy-Efficient Ethernet mode, works at finer granularity. The link between the MAC and the external PHY is maintained in the Low-power Idle mode (or LPI mode) while no data is transmitted nor received. The link returns to normal mode as soon as any data is transferred. There is no data loss in this mode.

Energy-Efficient Ethernet mode is available only in Full-duplex at 100 Mbit/s operation mode with the MII interface.



- The Ethernet peripheral embeds
 - Three master DMAs for direct memory interface
 - One DMA shared between two Rx queues
 - One DMA for each Tx queues (x2)
 - Internal FIFOs for Rx and Tx queues for dataflow management
 - 4KB for each FIFO shared between the queues
 - A media access controller (MAC) supporting 10/100/1000 Mbit/s data transfer rate in full- and half-duplex operations
 - A PHY interface block supporting Media Independent Interface (MII), Giga MII (GMII), Reduced MII (RMII) and Reduced GMII (RGMII) formats



This slide presents the peripheral block diagram.

The Ethernet peripheral embeds:

- Three direct memory interface for Receive and Transmit paths with an internal arbiter. In multiple queue operation, the Rx DMA is shared between the two queues while there is one Tx DMA per queue.
- Internal FIFOs for Rx and Tx queues for dataflow management
- A media access controller (MAC) supporting most functional features detailed in previous slides: Offload engines, Precision Timing Protocol, Power Management and MAC management counters for statistics gathering
- A Media Independent Interface with a dedicated block for the Reduced MII.

- Each functional block generates some interrupts but a single interrupt line is output and linked to the NVIC peripheral
- The PMT interrupt signal is also linked to the EXTI peripheral as a wakeup event

Functional block	Interrupts
DMA	Normal interrupts for correct transmission or reception Abnormal interrupts for any bus errors or buffer unavailability
MTL	Rx FFO overflow Tx FFO underflow
MAC	Timestamp interrupt: Set when specific events occur on PTP timing reference MMC interrupt: Set on statistics considerations LR interrupt: Set in Energy-Efficient Ethernet mode PMT interrupt: Set in Remote Wakeup mode PHY interrupt: Set by external PHY



An interrupt from the Ethernet peripheral can be generated as a result of various events. All these interrupt lines can be masked and converge to the same global interrupt signal linked to the NVIC peripheral.

An interrupt raised in Remote wakeup mode is redirected to the EXTI as a special event enabling a full system wakeup. Interrupts are reported from the three main functional blocks of the peripheral: the DMA, the MTL which manages the internal FIFOs for each receive and transmit path and the MAC itself, responsible for all functional part of the protocol. The DMA part generates normal interrupts when a packet is received or transmitted. It also raises interrupts for all bus error or buffer unavailability.

The MTL block generates interrupts when overflow is detected on the receive path or when an underflow is detected in the transmit path.

The MAC block generates interrupts linked to PTP protocol settings, MMC counters, Energy-Efficient Ethernet and

Remote Wakeup low-power modes. It also transfers the interrupt coming from the external PHY.

Mode	Description
Run	Active.
Sleep	Active. Peripheral interrupts cause the device to exit Sleep mode.
Stop	The Ethernet peripheral is able to detect frames while the system is in Stop mode, and wake it up via the EXTI controller.
Standby	Powered-down. The peripheral must be reinitialized after exiting Standby mode.

Here is an overview of the peripheral's status in each of the low-power modes.

Only PMT mode is available in Stop mode. In this mode, the peripheral waits for wakeup packets.

- The Ethernet peripheral is compliant with the following standards:
 - IEEE 802.3-2008 for Ethernet MAC, Gigabit Media Independent Interface (GMII), Media Independent Interface (MII)
 - IEEE 1588-2002 and IEEE 1588-2008 for Precision Timing Protocol (PTP)
 - IEEE 802.1Q-2005 for Virtual Bridged Local Area Networks (VLAN)
 - IEEE 802.1AS-2011 and 802.1-Qav-2009 for Audio Video (AV) traffic
 - IEEE 802.1-Qav specified credit-based shaper (CBS) algorithm for Transmit channels
 - IEEE 802.3az-2010 for Energy Efficient Ethernet (EEE)
 - RMI specification version 1.2 from RMI consortium
 - RGMII specification version 2.6 from HP/Marvell
- For more details, please refer to:
 - [UM1713](#): Developing applications on STM32Cube with LwIP TCP/IP stack (User manual)



The Ethernet peripheral is compliant with the following IEEE standards for support of:

- Ethernet MAC specification
- VLAN, PTP and AVB feature
- Energy Efficient Ethernet