



Chipsmall Limited consists of a professional team with an average of over 10 year of expertise in the distribution of electronic components. Based in Hongkong, we have already established firm and mutual-benefit business relationships with customers from,Europe,America and south Asia,supplying obsolete and hard-to-find components to meet their specific needs.

With the principle of "Quality Parts,Customers Priority,Honest Operation,and Considerate Service",our business mainly focus on the distribution of electronic components. Line cards we deal with include Microchip,ALPS,ROHM,Xilinx,Pulse,ON,Everlight and Freescale. Main products comprise IC,Modules,Potentiometer,IC Socket,Relay,Connector.Our parts cover such applications as commercial,industrial, and automotives areas.

We are looking forward to setting up business relationship with you and hope to provide you with the best service and solution. Let us make a better world for our industry!



Contact us

Tel: +86-755-8981 8866 Fax: +86-755-8427 6832

Email & Skype: info@chipsmall.com Web: www.chipsmall.com

Address: A1208, Overseas Decoration Building, #122 Zhenhua RD., Futian, Shenzhen, China





fido2100 3-Port Industrial Ethernet DLR Switch with IEEE 1588

Data Sheet

Copyright © 2013 by Innovasic, Inc.

Published by Innovasic, Inc.
5635 Jefferson St. NE, Suite A, Albuquerque, New Mexico 87109 USA

fido® is a trademark of Innovasic, Inc.



Document #: IA211111101-04
UNCONTROLLED WHEN PRINTED OR COPIED

TABLE OF CONTENTS

1	Overview	7
1.1	Introduction to Device Level Ring Protocol.....	8
1.2	Introduction to IEEE-1588.....	9
2	Features	10
2.1	IEEE 802.3 MAC Details	11
2.1.1	IEEE 802.3 Full Duplex Flow Control	12
2.1.2	Half Duplex Flow Control and Broadcast/Multicast Storm Prevention ...	13
2.2	IEEE 1588 V2.....	13
2.2.1	IEEE 1588 Hardware Assist for Ordinary Clock.....	13
2.2.2	IEEE 1588 End to End Transparent Clock	16
2.3	Device Level Ring (DLR) Protocol Support	17
2.3.1	Network Loops.....	18
2.4	Cut Through Forwarding	18
2.5	Quality of Service	19
2.6	Unicast and Multicast Address Filtering	19
2.7	Statistics Counters.....	20
2.8	Buffer Management	20
2.9	Supported Network Topologies	21
3	LQFP Package.....	25
3.1	LQFP Package Pinout.....	25
3.2	LQFP Pin Listing	26
3.3	LQFP Package Physical Dimensions.....	30
4	BGA Package	31
4.1	BGA Package Pinout	31
4.2	BGA Pin Listing	32
4.3	BGA Package Physical Dimensions.....	36
5	Electrical Characteristics	37
5.1	Thermal Characteristics	38
5.2	PLL Supply	38
5.3	System Clock Input.....	38
5.4	RESET: Power-On-Reset, Hardware Reset and Software Reset.....	39
5.4.1	Power-On-Reset.....	39

5.4.2	Hardware Reset	39
5.4.3	Software Reset	39
5.5	Test Input	39
6	Host Bus Interface	40
6.1	READ, Host Bus Interface.....	40
6.2	WRITE, Host Bus Interface.....	41
7	External MII Interface	42
7.1	RECEIVE DATA, External MII Interface	42
7.2	TRANSMIT DATA, External MII Interface	43
8	Internal MII Interface	43
8.1	RECEIVE DATA, Internal MII Interface.....	44
8.2	TRANSMIT DATA, Internal MII Interface	45
9	Control and Status Registers	46
9.1	Register Map.....	46
9.2	Detailed Register Explanations.....	51
10	Revision History.....	84
11	For Additional Information	84

LIST OF FIGURES

Figure 1 fido2100 Top Level Block Diagram.....	7
Figure 2 IEEE 1588 Hardware Assist Block Diagram	14
Figure 3 Hierarchical Start Topology	21
Figure 4 Hybrid Daisy Chain and Hierarchical Star Topology	22
Figure 5 Ring Topology Media Redundancy.....	23
Figure 6 fido2100 LQFP Package Pinout	25
Figure 7 LQFP Package Physical Dimensions	30
Figure 8 fido2100 BGA Package Pinout.....	31
Figure 9 BGA Package Physical Dimensions	36
Figure 10 Host Interface, READ Timing.....	40
Figure 11 Host Interface, WRITE Timing	41
Figure 12 External MII Interface, RECEIVE DATA	42
Figure 13 External MII Interface, TRANSMIT DATA.....	43
Figure 14 Internal MII Interface, RECEIVE DATA	44
Figure 15 Internal MII Interface, TRANSMIT DATA.....	45

LIST OF TABLES

Table 1 128-Pin LQFP Pin List.....	26
Table 2 BGA Pin List.....	32
Table 3 Absolute Maximum Ratings	37
Table 4 ESD and Latch-Up Characteristics	37
Table 5 Recommended Operating Conditions	37
Table 6 DC Characteristics of I/O Cells	37
Table 7 Power Consumption (Nominal)	38
Table 8 Thermal Resistance Characteristics	38
Table 9 Host Interface, READ Timing	40
Table 10 Host Interface, WRITE Timing	41
Table 11 RECEIVE DATA, External MII Interface.....	42
Table 12 TRANSMIT DATA, External MII Interface	43
Table 13 RECEIVE DATA, External MII Interface.....	44
Table 14 TRANSMIT DATA, External MII Interface	45
Table 15 Register Map.....	46
Table 16 Document Revision History.....	84

1 Overview

The fido2100 is part of the Innovasic fido[®] family of real-time communication products. The fido2100 is a managed, 3-port Ethernet switch with support for IEEE-1588 and Device Level Ring (DLR) protocol. The fido2100 enhances the capabilities of Innovasic’s line of Industrial Ethernet products by providing an infrastructure element with low latency, time synchronization and redundancy for any embedded Industrial Ethernet solution.

The fido2100 is an IEEE 802.3 standard compliant, Layer 2 switch that is compatible with the IEEE 802.1D standard. It has three ports: two are external ports that function as physical ports of a product and one is an internal port that connects to the host CPU. Each of the three ports has an IEEE 802.3 compliant MAC and the ports are interconnected with each other through full wire speed, non-blocking switching logic. Figure 1 illustrates the top-level blocks of the fido2100 architecture.

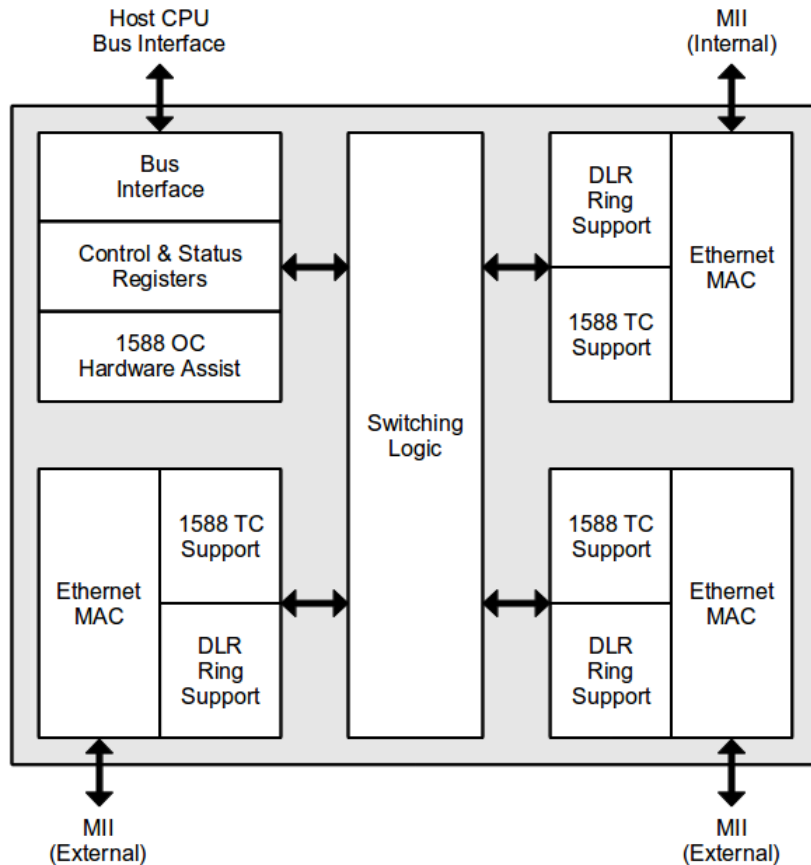


Figure 1 fido2100 Top Level Block Diagram

The pass-through of traffic is handled swiftly by using cut through forwarding to pass traffic to the other external port. This is accomplished without host CPU intervention, which keeps throughput time host interaction to a minimum and keeps pass through delay to a minimum.

The fido2100 also has the ability to operate in a DLR configuration, which provides media redundancy amongst its members. Host CPU source code is available to provide the ability for the device to be a DLR participant or a DLR supervisor using the DLR support functions. The fido2100 will default to linear topology operation, and when inserted into a ring, will automatically change to DLR operation as a DLR participant.

1.1 Introduction to Device Level Ring Protocol

There are a variety of Industrial applications in which Ethernet ring topologies are preferable to the star topologies common in enterprise networks. Ring networks provide inherent single-point fault tolerance and reduced connectivity costs. Device Level Ring (DLR) protocol provides a means for detecting, managing and recovering from faults in a ring-based network.

DLR supports three classes of devices:

- 1 Ring Supervisor – Ring Supervisors are required to send and process DLR beacon frames at the default beacon interval.
- 2 Ring Node, Beacon-based – These devices are required to process beacon frames.
- 3 Ring Node, Announce-based – These devices are not required to process DLR beacon frames, but must be capable of processing announce frames.

A DLR network consists of a Ring Supervisor and any number of Ring Nodes. Ring nodes incorporate fido2100 technology with at least two external ports. The Ring Supervisor is responsible for generating a “beacon” at regular intervals. These beacons traverse the ring in each direction. The ring supervisor must be capable of blocking DLR frames to avoid infinite propagation of beacons. Faults are detected when beacon traffic is interrupted. There are obviously a number of failure mechanisms and associated recovery strategies. For a detailed explanation of DLR refer to ODVA documentation - Volume II: EtherNet/IP Adaptation of CIP, chapter 9, section 9-5. For definition of the DLR EtherNet/IP object, refer to Volume II: EtherNet/IP Adaptation of CIP, chapter 5, section 5-5.

1.2 Introduction to IEEE-1588

The IEEE 1588 standard, also known as Precision Time Protocol (PTP), provides a means to synchronize participants in a network to a common time source. Each network participant has its own precision time source or “ordinary clock”. A PTP system consists of some number of ordinary clocks connected to a network. A grandmaster is elected based upon the quality of available time sources and all other participants synchronize directly to it.

PTP systems can be expanded through the use of “boundary clocks”. The boundary clock provides a means to bridge synchronization from one network segment to another. A synchronization master can then be elected for each network segment. The root timing reference remains the grandmaster.

PTP is a master-slave protocol. The master and network participants exchange synchronization messages. However, these messages will be delayed as they traverse the network due to the inherent latency of the network infrastructure. To compensate for this latency, PTP includes the concept of a transparent clock. The transparent clock automatically compensates for latency by modifying the timestamps in synchronization messages as they pass through a given device. The fido2100 provides hardware support for both transparent and ordinary clocks.

2 Features

- IEEE 802.3
 - 10/100Mbps, Half / Full Duplex
 - Full Duplex Flow Control
 - Half Duplex Back Pressure Flow Control
 - Broadcast / Multicast Storm Prevention

- IEEE 1588 V2
 - Hardware Assist for Ordinary Clock
 - End-to-end Transparent Clock

- DLR Beacons
 - Capability to process beacons from 100 ms down to 100 μ s
 - Provides auto-generation of beacon frames for supervisor functions

- Cut through forwarding
 - Cut through forwarding minimizes latency for High Performance Control
 - Applications such as CIP Motion

- IP Differentiated Services Code Point (DSCP) based Quality of Service (QoS)
 - 4 Prioritized Queues per Port

- Incoming Filtering on Unicast / Multicast Traffic
 - 256-entry Dynamic Unicast MAC Learning and Filtering for 2 External Ports
 - 128 buffers of 128 bytes for each port

- Statistics Counters for 2 External Ports

- Supported Network topologies:
 - Hierarchical Star with either one of the external ports
 - Daisy Chain / Hybrid Daisy Chain Hierarchical Star
 - DLR Media Redundancy with Hardware Support for a Single Fault Tolerant Ring Protocol

- Software for Switch & DLR Protocol Management

- Full Industrial Temperature Range -40 to +85C
- Packages:
 - 128 pin Plastic Low Profile QFT (LQFP), RoHS Compliant
 - 128 Ball Grid Array (10x10mm BGA), RoHS Compliant

2.1 IEEE 802.3 MAC Details

The fido2100 implements an IEEE 802.3 standard compliant MAC for each of the three ports. While beginning a transmission, the MAC will transmit 7 bytes of preamble pattern 0x55, followed by one byte of start of frame delimiter 0xD5, followed by the frame data and frame check sequence as per IEEE 802.3 specification.

While transmitting, the MAC will enforce an inter-frame gap period of 960 ns at 100Mbps speed and 9.6 ns at 10Mbps speed as per IEEE 802.3 specification. While receiving, the MAC can tolerate shrinkage of the inter-frame gap period down to 80 nanoseconds at 100Mbps speed and 800 nanoseconds at 10Mbps speed.

In full duplex mode the MAC will ignore carrier sense and collision detect signals from the PHY and is ready to transmit after satisfying the inter-frame gap period as per IEEE 802.3 specification. In half duplex mode, the MAC will wait for the current transmission to complete and the carrier sense signal to be de-asserted. Then it will monitor if the carrier sense signal gets asserted in the first 2/3 duration of the inter-frame gap period. If the carrier signal gets asserted within that period, the MAC will wait until the carrier sense signal gets de-asserted to restart inter-frame gap period again. If the carrier sense signal gets asserted in the last 1/3 duration of the inter-frame gap it will be ignored as per the IEEE 802.3 specification. The MAC is ready for transmission at the end of the inter-frame gap period. If a frame is not available for transmission at the end of the inter-frame gap, the carrier sense signal will be monitored for restarting inter-frame gap period.

If the frame transmission has commenced in half duplex mode and the collision signal gets asserted by the PHY, transmission of the current frame data will be truncated with a jam pattern of 4 bytes of 0xFF. At this point the collision back off and retransmission procedure will start using the truncated binary exponential back off algorithm as per IEEE 802.3 specification. At the end of enforcing a jamming pattern, the MAC will impose a delay before attempting to retransmit the frame. The delay is an integer multiple of a slot time, where one slot time is equal to 512 bits time. The number of slot times to delay before the n^{th} retransmission attempt is chosen as a uniformly distributed random integer r in the range:

$$0 \leq r < 2^k; \text{ where } k = \min(n, 10)$$

If sixteen transmission attempts fail with collision for the same frame, the frame will be dropped as per IEEE 802.3 specification. The mechanism used to generate the uniformly distributed random number is a free running 10 bit wide linear feedback shift register (LFSR). The LFSR will cycle through all states in 1024 transmit clocks of the port under consideration. Depending on the retransmission attempt number n , the most significant k bits of LFSR are used as delay integer r . Due to the free running nature of the LFSR, lock step random number generation in two different devices should not occur. However, a facility is provided whereby the firmware can load a random seed value into the LFSR anytime during operation.

The MAC will drop frames received with errors signaled by the PHY device and those with frame check sequence (CRC) errors. It will drop frames shorter than 64 bytes as per IEEE 802.3 specification. It will drop frames longer than maximum IEEE 802.3 tagged frame size of 1522 bytes. The MAC will also drop frames with alignment error, i.e. those having a non-integral number of bytes.

2.1.1 IEEE 802.3 Full Duplex Flow Control

The fido2100 implements IEEE 802.3 compliant full duplex flow control using PAUSE frames. In order to provide optimum performance, a special set of rules are followed for PAUSE frame generation and reception.

PAUSE frame reception is enabled on external ports operating in full duplex mode and transmission on the port through which the PAUSE frame was received will be suspended for specified pause duration in the PAUSE frame. PAUSE frame reception is always disabled on the internal port, since the host CPU will always have enough memory for its communications.

If speed and duplex modes of external ports are same and when buffer usage on either one of external ports reaches 95% of buffer capacity, a PAUSE frame will be generated from internal port to host CPU with maximum pause duration. If the buffer situation has not improved at the end of pause duration, additional PAUSE frames will be generated from internal port as needed. When buffer usage on both external ports reaches below 85% of buffer capacity, a PAUSE frame will be generated on internal port with zero pause duration to cancel pause. If speed and duplex modes of external ports are not same, PAUSE frames will not be generated.

For external ports, PAUSE frame generation on a port is enabled only for full duplex mode and when both external ports have same speed and duplex settings. When a first external port is operating in full duplex mode and buffer usage on internal port or second external port reaches 95% of buffer capacity a PAUSE frame will be generated on first external port with maximum pause duration. If the buffer situation has not improved at the end of pause duration, additional PAUSE frames will be generated from first external port as needed. When buffer usage on both internal and second external ports reaches below 85% of buffer capacity, a PAUSE frame will be generated on first external port with zero pause duration to cancel pause. If speed and duplex modes of external ports are not same, PAUSE frames will not be generated.

Obviously, for flow control on fido2100 to work properly, the integrated MAC on host CPU must support flow control as well. Most Ethernet microprocessors do support this capability and dual port product designer must ensure that the selected host CPU supports flow control.

2.1.2 Half Duplex Flow Control and Broadcast/Multicast Storm Prevention

The fido2100 implements back pressure flow control in half duplex mode. When both external ports are operating in same speed and half duplex mode, and buffer usage on the internal port or other external port reaches 95% of buffer capacity, back pressure flow control is activated on the first external port. This continues until the buffer usage on both internal and second external ports reaches below 85% of buffer capacity. When back pressure flow control is activated on a port, the port will continue to transmit frames while it has any. If there is no frame to transmit, the frame will keep the carrier sense signal active by sending the preamble pattern periodically, causing the neighboring node to back off from transmitting any frames. If speed and duplex modes of external ports are not same, back pressure will not be activated.

Excessive broadcast frames can overload the CPU of all devices on network, especially IO adapters and IO blocks, leading to poor performance. To avoid this, the fido2100 implements a broadcast storm prevention mechanism. When received broadcast data within a 100 millisecond period (1 second at 10Mbps) reaches about 1% of network bandwidth in a port operating at 100Mbps speed, additional broadcast frames received on that port within that period will be dropped. This process is repeated for every 100 millisecond period. In well-engineered networks, broadcast storms do not occur during normal operation and a 1% broadcast storm limit will never be reached.

Similarly every port is also monitored for received non-redundancy (non-DLR/BRP) multicast frames. When received non-redundancy multicast data within a 10 millisecond period (100 milliseconds at 10Mbps) reaches about 50% of network bandwidth in a port operating at 100Mbps speed, additional nonredundancy multicast frames received on that port within that period will be dropped. This process is repeated for every 10 millisecond period. In well-engineered networks, multicast storms do not occur during normal operation.

2.2 IEEE 1588 V2

2.2.1 IEEE 1588 Hardware Assist for Ordinary Clock

Figure 2 below shows a block diagram of the IEEE 1588 V2 hardware assist that can be used to implement an IEEE 1588 ordinary master or slave clock on an end device. The MII traffic of the internal port between the host CPU and the fido2100 is monitored by two independent 1588 frame detection logic blocks, one for the transmit channel and the other for the receive channel.

Whenever the timestamp point of a passing transmit/receive frame is reached, a snapshot of the system time counter is saved in a temporary register. When the appropriate 1588 frame type is also detected in the frame, the snapshot is saved from the temporary register to the transmit/receive snapshot register. The receive snapshot register contains a sixteen entry FIFO, while the transmit snapshot register is a

single entry register. With the transmit snapshot register, the snapshot will be locked to prevent overwriting by subsequent frames. The firmware must later clear the lock after it has read the snapshot. With the receive snapshot register, the firmware should read all snapshots while the FIFO is not empty. If the FIFO becomes full because the firmware is slow to read it, additional snapshots will not be saved until the FIFO has a free entry. To distinguish saved snapshots, the received frame sequence identifier and source port identity are also locked in the FIFO and may be read through appropriate registers along with the snapshot.

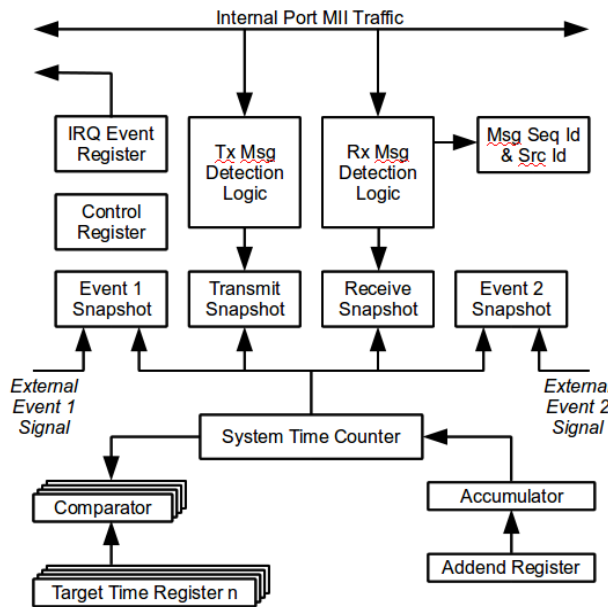


Figure 2 IEEE 1588 Hardware Assist Block Diagram

The supported frame mapping for frame detection logic is PTP over UDP over IPV4 over IEEE 802.3 tagged or untagged frames. The frames detected are based on the master mode bit setting in the time sync control register. In master mode, Sync frames are time stamped on transmit and Delay_Req frames are time stamped on receive. In slave mode, Delay_Req frames are time stamped on transmit and Sync frames are time stamped on receive. In addition, the sub-domain field of frames must match the subdomain configuration register for the snapshot to be saved.

The addend register, accumulator and system time counter comprise a frequency compensated clock under firmware control. The accumulator adds the addend register to itself on every incoming clock from the oscillator and the overflow signal of the accumulator causes the system time counter to increment. This provides a high precision tunable digital clock whose frequency of operation can be controlled by firmware through writing suitable values to the addend register.

While transmit and receive frame detection logic operate independently at 25MHz MII clock rates, the rest of the logic operates at 100MHz clock rate. The addend register and accumulator of frequency compensated clock are 32-bits wide, while the system time counter is 64-bits wide. The tunable precision for frequency compensation is better than 1 part per billion (1×10^{-9}). The procedure for realizing various nominal system time frequencies is described below.

- **FreqOscillator** is the clocking frequency of the time synchronization
- **FreqClock** is the nominal frequency at which the system time counter is to be incremented.
- **FreqDivisionRatio** = $\text{FreqOscillator} / \text{FreqClock}$ (This must always be > 1)
- **FreqCompensationValue** = frequency compensation value is the number held in the addend register, which is added to the accumulator once every $1/\text{FreqOscillator}$.

The equation for the FreqCompensationValue utilizes the precision of the accumulator and the FreqDivisionRatio. Since the accumulator is 32 bits, the following equation provides the value for the Addend register:

$$\text{FreqCompensationValue} = 2^{32} / \text{FreqDivisionRatio}$$

The hexadecimal representation of the FreqCompensationValue is the value that is written to the addend register. The following table gives examples of addend values based on a 100 MHz FreqOscillator.

Freq-oscillator	FreqClock	FreqDivisionRatio	FreqCompensationValue
100 MHz	83.33 MHz	1.2	0xD5555555
100 MHz	80 MHz	1.25	0xCCCCCCCD
100 MHz	66.66 MHz	1.5	0xAAAAAAB

The synchronization accuracy between the time master and the slave depend on cumulative accuracy of individual components used in the system. When a time slave is connected directly to the time master, a synchronization accuracy of less than 50 ns from master is easily possible. This accuracy is achievable with low cost 50PPM crystal clock oscillators. Turbulent air flow over the crystal oscillator should be avoided to prevent rapid temperature changes resulting in short term stability errors. Some PHY devices introduce random delays in transmit and receive paths, which can increase time synchronization errors significantly. It is recommended that the PHY devices be investigated for such behavior. Any asymmetry in transmit and receive paths should be accounted properly by firmware to avoid accumulation of errors. In a cascaded transparent clock configuration between the time master and the slave, the errors will increase linearly with the number of transparent clocks.

There are two 64-bit target time registers and comparators that can be used to generate an interrupt to the host CPU when the system time counter reaches the target time. This feature can be used by firmware to schedule/coordinate control loops/activities. While both target time registers can be used by firmware for any purpose, the second target time register has a special capability to generate a pulse per second signal. The pulse per second signal is required for internal IEEE 1588 compliance testing, but is not required on shipping products.

There are two additional snapshot registers for two external events. These snapshots can be used to synchronize system time with an external clock. For example, the event 1 can be used to synchronize internal IEEE 1588 system time as a slave with a global positioning system (GPS) signal as the master. Similarly, the event 2 can be used to synchronize internal IEEE 1588 system time as a master with an external clock in another module in the system as the slave. Both event snapshots have the same capability/behavior in implementation and they can be used interchangeably.

2.2.2 IEEE 1588 End to End Transparent Clock

The fido2100 implements an integrated IEEE 1588 V2 end to end transparent clock (E2E TC). This is a single step transparent clock in IEEE 1588 terminology. The supported frame mapping is PTP over UDP over IPV4 over IEEE 802.3 tagged or untagged frames. The E2E TC uses a non-synchronized free running timer as a time base. The free running timer runs at 100MHz, but the timestamp circuit uses both edges of the clock for sampling, and delay computation is suitably adjusted to make it effectively run at 200MHz, with 5 ns resolution.

In E2E TC mode, whenever an IEEE 1588 V2 Sync or Delay_Req frame is received, a receive timestamp is triggered on ingress at the timestamp point following the start of frame delimiter. The frame is parsed to extract the UDP checksum and correction field. When the received frame doesn't contain any errors and when the frame is transmitted through a port, a transmit timestamp is triggered on egress at the timestamp point. The residence time delay is then computed on the fly from the transmit and receive timestamps. The residence time delay is then added to the correction field in the frame with the UDP checksum and frame CRC adjusted on the fly.

Without transparent clocks, cascaded switches or boundary clocks will accumulate errors at exponential rate making time synchronization unstable. With transparent clocks the accumulation of errors is linear with the number of nodes. The implemented E2E TC has a theoretical worst case error of +/-15 ns per node assuming 50PPM source crystal oscillator stability.

2.3 Device Level Ring (DLR) Protocol Support

The fido2100 provides hardware support for a single-fault tolerant ring protocol. The DLR protocol support for non-supervisor mode and support for supervisor mode can be enabled by setting appropriate bits in the redundancy control register.

In non-supervisor mode, received non-erroneous DLR beacon frames from the active ring supervisor through either port will be automatically dropped after extracting state information and restarting the DLR beacon timers. Received non-erroneous DLR beacon frames may optionally be configured to be delivered to the host CPU by setting appropriate bits in the redundancy control register. Irrespective of this bit setting, beacon messages received from a different supervisor than active supervisor are always forwarded to host CPU. The fido2100 can be configured to interrupt host CPU when ring beacon frames from active supervisor are received through either port or are not received through either port within ring beacon timeout period. The fido2100 can also be configured to interrupt host CPU when a change of state is observed in ring beacon frames from active ring supervisor received through either port. Neighbor check and multicast sign on frames received from either external port will be forwarded only to host CPU and those from host CPU will be forwarded only through port number matching source port number field in frame. Source identifier and sequence identifier fields of received neighbor check frames are captured to identify received port later. Unicast sign on frames are treated as any other unicast frame.

In supervisor mode, either one of the external ports can be placed in blocked mode by setting appropriate port transmit/receive blocked bits in redundancy control register. All frame reception and transmission will be blocked by the port in blocked mode with the exception of some special frames. These special frames are described under port transmit/receive blocked bits in redundancy control register. In supervisor mode, ring beacon frames can be automatically generated and transmitted through unblocked ports to reduce the CPU load and their arrival through both ports are monitored. The fido2100 can be configured to interrupt the host CPU when DLR beacon frames are received through either port or are not received through either port within DLR beacon timeout period. The fido2100 can also be configured to interrupt the host CPU when a change of state is observed in DLR beacon frames are received from the active ring supervisor through either port. Received non-erroneous DLR beacon frames may optionally be configured to be delivered to host CPU by setting appropriate bits in redundancy control register. Irrespective of this bit setting, beacon messages received from a different supervisor than self is always forwarded to host CPU.

In supervisor mode, neighbor check and multicast sign on frames received from either external port will be forwarded only to the host CPU, and those from the host CPU will be forwarded only through the port number matching the source port number field in the frame. The source identifier and sequence identifier fields of received neighbor check frames and sign on frames are captured to identify the received port later. Link/neighbor status frames received from either external port will be forwarded

only to the host CPU. The source identifier and sequence identifier fields of received link/neighbor status frames are captured to identify the received port later.

2.3.1 Network Loops

The fido2100 has a built in safety feature to detect network loops. Undetected, network loops can cause severe network problems. Network loops may be present when the user inadvertently misconnects the network in such a way as to create a loop or when the user intentionally creates a DLR topology, but fails to configure a DLR supervisor. When a frame is received through one of the external ports with the source MAC address the same as the MAC address of the host CPU, the fido2100 will drop that frame and will set a port-specific bit in the switch event register indicating that such a frame was received. These bits can be read and cleared by firmware and can be used to flag the network fault situation to the user. Alternatively, the fido2100 can be configured to interrupt the host CPU when those bits are set.

It should be noted that in a properly configured DLR with a DLR supervisor, network loops may be present for short durations when the network is being reconfigured. Hence firmware should not flag the user of this fault condition when operating in proper DLR mode. The firmware can detect the proper DLR mode by the presence of DLR beacons on the network, whereas in non-DLR mode or an improperly configured DLR mode, DLR beacons would be absent on network.

2.4 Cut Through Forwarding

The fido2100 implements cut through forwarding with store and forward on contention at the transmitting port. COTS switches and standard switching integrated circuits typically support only the store and forward approach. With the store and forward approach a frame has to be completely received before being transmitted. Hence every frame will encounter a delay equal to the length of the frame and additional internal switching delay. This can become an issue for high performance applications such as CIP motion when a long chain of cascaded switches are involved, such as in daisy chain or DLR topology networks.

With cut through forwarding, when frame reception begins and when a sufficient number of bytes of a frame have been buffered, the frame is queued for transmission, subject to the following forwarding rules:

1. If the transmitting port is idle and ready for transmission, it will immediately start frame transmission while the reception is in progress.
2. If the transmitting port is not free immediately, it will start transmitting as soon as the port becomes ready for transmission, while the reception is in progress.
3. If the transmitting port doesn't become available before the entire frame is received, the received frame will be buffered completely for store and forward.

Because of cut through forwarding, non-IP frames from all transmit queues and IP frames from transmit queues 2 and 3 will face a delay of 2.8 μ s, other non-IEEE 1588 frames will face a delay of 4.9 μ s and IEEE 1588 time synchronization frames will face a delay of 7.3 μ s through the fido2100, when the transmitting port is free of contention. Cut through forwarding is disabled for a transmit port when the receiving port is operating at 10Mbps speed and the transmitting port is operating at 100Mbps speed. In this case, store and forward is used.

2.5 Quality of Service

In order to support high performance applications such as CIP motion, the fido2100 enforces quality of service based on IP Differentiated Services Code Point (DSCP). The fido2100 implements four prioritized transmit queues per port. Frames received with DSCP 59 are queued in the highest priority transmit queue (queue #1). Frames received with DSCP 55 are queued in the second highest priority transmit queue (queue #2). Frames received with DSCP 47 and 43 are queued in the third highest priority transmit queue (queue #3). Frames received with other DSCP values are queued in the lowest priority transmit queue (queue #4). In addition, DLR protocol frames are queued in the highest priority queue 1. When a port is ready to transmit the next frame, the highest priority frame is chosen from the current set of queued frames for transmission based on strict priority ordering. Within a given priority queue frames are transmitted in FIFO order.

2.6 Unicast and Multicast Address Filtering

The fido2100 implements 256 entry dynamic unicast MAC address learning and filtering mechanism. The unicast MAC address learning table is implemented as 4-way learning/look up table of 64 rows. The 48-bit source MAC address from a frame received on any external port is hashed into 6-bits by XOR operation of every sixth bit. The hash index identifies the row where one of four locations that is free is used for MAC address learning along with associated external port. When none of them is free, the 4th way location is always overwritten. Learning is performed on source MAC address only from frames with a unicast or multicast destination MAC address. Automatic aging of learned entries will happen in 4- 6 minutes, unless learning is refreshed within that period.

When a frame with unicast destination address is received through any port, 4-way simultaneous lookup is performed from entries at hash index location of destination MAC address. If the address is found, the frame is forwarded only through associated external port. If the address is not found, then it is forwarded through all other external ports. When a frame with unicast destination address is received through one of the external ports and the destination address matches host CPU MAC address, then that frame is forwarded only to internal port.

The fido2100 implements multicast filtering on the internal port for frames received through external ports and no multicast filtering is implemented for the external ports. The multicast filter consists of a 2048 bin hash table. The received multicast destination address is passed through a 32-bit Ethernet

CRC generator and the least significant 11 bits of the resulting CRC is used to index into the 2048 bin hash table. When the indexed bin in the hash table is set, the frame will be forwarded to the internal port and the other external port. If not set, the frame will be forwarded only to the other external port. By default, all 2048 bins of the multicast hash table are set to zero after power up or chip reset, and no multicast frames will be forwarded to the internal port. Since most devices such as I/O devices don't need to receive multicast frames, they don't need to change this default behavior. Note that the multicast filter table need not be set for receiving multicast ring frames.

2.7 Statistics Counters

The fido2100 implements statistics counters for the external ports. These counters can be used by firmware for the media and interface counters of the EtherNet/IP Ethernet Link object (Class Code: 0xF6). Most of these counters are 16 bits wide and to avoid roll over issues, firmware must read these registers at least once every 200 milliseconds or on demand, to calculate an increase in counter values. The firmware can then add the increase to firmware maintained 32 bit counters.

Counters are implemented for valid frames received with the group bit set in the destination address, valid frames received with a unicast destination address, total byte count of valid frames received, frames received with length greater than 1522 bytes, frames received with alignment error, frames received with a frame check sequence (CRC) error and frames received with other errors, including short runt frames less than 64 bytes in length.

Counters are also implemented for valid frames transmitted with the group bit set in the destination address, valid frames transmitted with a unicast destination address, total byte count of valid frames transmitted, frames transmitted after exactly one collision, frames transmitted after multiple collisions, frames dropped due to excessive collisions and frames truncated with error during transmission.

2.8 Buffer Management

The fido2100 implements an efficient packet buffering scheme using on board dual port SRAM memory. The buffers are 128 bytes in size to minimize memory wastage. Frames longer than 128 bytes are automatically fragmented to be stored in multiple buffers and reassembled during transmission.

Actual buffer usage depends on a number of factors such as network load, number of frames waiting to be transmitted at each node at the same time, the number of back to back frames arriving at the destination node through both external ports at the same time, etc. For some devices such as motion drives and I/O, buffer usage on any port may never exceed 10% of buffer capacity. For other devices such as network bridges, buffer usage on any port may be typically less than 50% of buffer capacity during normal operation. Irrespective of the type of end device, when buffer usage reaches 95% of buffer capacity on any port under heavily loaded network conditions, flow control will intervene to reduce buffer usage and will ensure frames are not being dropped.

2.9 Supported Network Topologies

Dual port products designed with this toolkit can be used in a wide variety of network configurations. This single hardware design will support various network topologies. Irrespective of the network configuration used, a dual port end device will always have a single IP address and a single MAC address. This simplifies both configuration requirements and firmware changes required.

Figure 3 shows hierarchical star topology with commercial off the shelf (COTS) switches. Either port of a dual port end device may be used to connect to COTS switch to realize hierarchical star network topology.

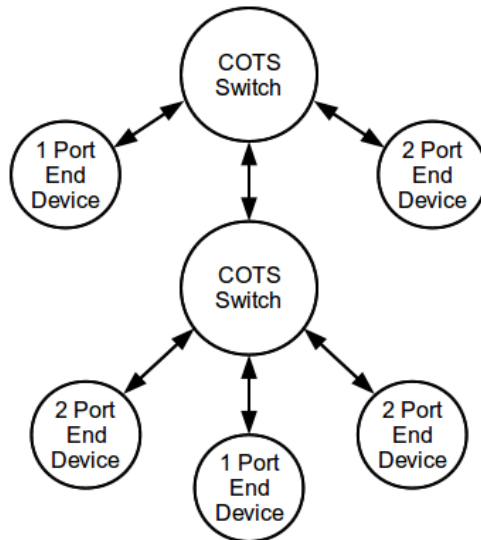


Figure 3 Hierarchical Start Topology

Figure 4 shows a hybrid daisy chain and hierarchical star topology. Single port end devices can be connected to a daisy chain at one of the ends of the chain, or through a standalone three port device that provides the daisy chain capability in the middle of chain. Figure 3 also shows a COTS switch in the middle of a daisy chain.




Figure 4 Hybrid Daisy Chain and Hierarchical Star Topology

It should be noted that a slower speed or half duplex link on main daisy chain will become a bottle neck, causing the entire network to operate at reduced capacity. Similarly, a single port device can be connected to either end of main daisy chain network, only if it operates at same/better speed and duplex mode as rest of the network. The recommended method for connecting slower speed/duplex devices is through a three port tap device to avoid performance issues.

Figure 5 shows a ring topology media redundancy that can tolerate single faults. One or more DLR supervisor(s) control the operation of ring. When multiple DLR supervisors are configured, one of them will be automatically chosen as active supervisor. Others will be in passive mode, until the active DLR supervisor fails or is removed from ring.

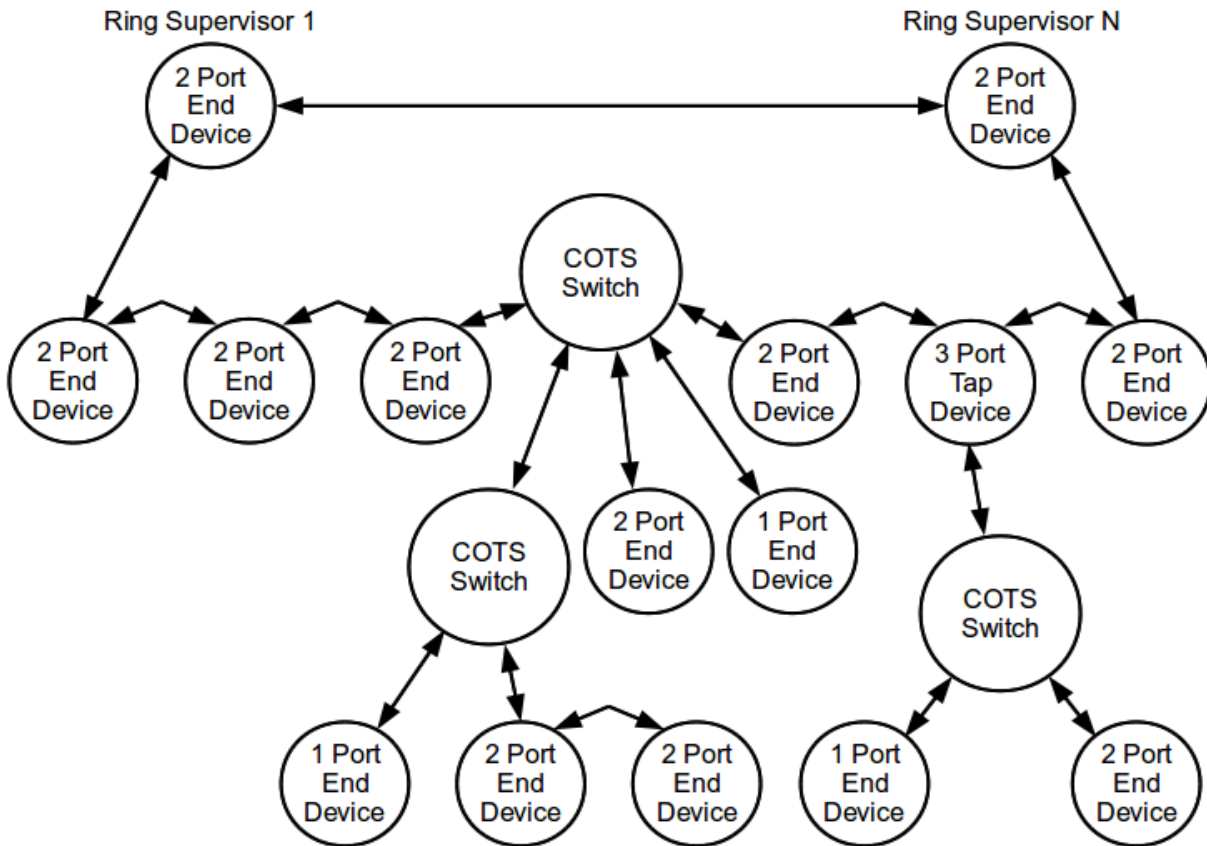


Figure 5 Ring Topology Media Redundancy

For IO applications, a COTS switch in the ring must support IEEE 802.3 tag based and IP DSCP based QoS with at least two prioritized queues, preferably four prioritized queues. The highest priority queue must be dedicated for ring protocol messages for predictable performance. For IEEE 1588 applications such as CIP motion, COTS switches directly on ring must support IEEE 1588 end-to-end transparent clock (or boundary clock) and IEEE 802.3 tag based and IP DSCP based QoS with four prioritized queues. The highest priority queue may be shared between IEEE 1588 and ring protocol messages. If not, COTS switches can be connected through three port tap device as shown in Figure 4.

The COTS switch must be configured so that ring protocol messages are not forwarded to devices connected to it. IGMP snooping and multicast filtering must be disabled on 2 ports of the COTS switch directly connected to ring to facilitate rapid network reconfiguration. Since unicast filtering cannot be disabled in a COTS switch, it may take couple of CIP connection RPI's for some data packets to be delivered correctly on some nodes after a network reconfiguration. Ring timing parameters such as beacon period and timeout must account for the presence of a COTS switch directly in ring. Typical network recovery time for a ring containing 50 nodes will be less than 1 millisecond for most types of faults on a loaded network with mostly EtherNet/IP frames. Further details of ring topology and protocol are described in, "Device Level Ring Protocol"; see CIP Networks Library, Volume 2, Chapter 9, Edition 1.7 or later (see www.odva.org for details).

It is possible to create complex ring topologies. Some examples are: a backbone ring connecting multiple star topologies; a backbone star connecting multiple ring topologies; a backbone ring connecting multiple ring topologies; and other combinations.

Most dual/three port devices operate as simple ring nodes and require no special configuration. By default they will power up in daisy chain mode and transparently switch to ring mode upon receiving beacon frames from ring supervisor.

3 LQFP Package

3.1 LQFP Package Pinout

The pinout for the fido2100 Industrial Ethernet Switch LQFP package is as shown in Figure 6. The corresponding pinout is provided in Table 1.




Figure 6 fido2100 LQFP Package Pinout