



PE310G4DBIR

Quad port Fiber 10 Gigabit Ethernet PCI Express Content Director Server Adapter Intel® FM10420 Based

Product Description

Silicom's 10 Gigabit Ethernet PCI Express content aware director Bypass server adapters is designed for servers and high-end appliances. The Silicom content aware director server adapter is designed with an on board smart routing architecture that enables packets to be redirected or dropped based on defined rules.



The Silicom's 10 Gigabit Ethernet content aware packet director reduces host system process since only packets that are defined to be targeted to the host systems are routed to the host; other packets can be routed to the other port or can be dropped by the content aware hardware routing architecture. The Silicom's 10 Gigabit Ethernet content aware packet director is targeted to network applications that needs to process, monitor or bypass packets based on defined rules.


The adapter supports three main modes of operation: Content Aware Bypass, Content Aware TAP and content Aware filtering NIC.

Content Aware Bypass

Silicom's 10 Gigabit Ethernet content aware director provides intelligent packet redirection capability where rules specify which packets are directed to the host system and which packets are directed to the other port (Bypass).

Content Aware TAP

Silicom's 10 Gigabit Ethernet content aware director provides intelligent packet redirection capability where all packets are directed to the other port (Bypassed) and rules specify which packets are copied to the host system (TAP).



Content Aware Filtering NIC

Silicom's 10 Gigabit Ethernet content aware provides intelligent packet redirection capability where rules specify which packets are directed to the host or dropped.

The Silicom 10 Gigabit Ethernet PCI Express content director server adapter is based on Intel FM10420 Ethernet controller and a L3 switch router. The Silicom's 10 Gigabit Ethernet PCI Express adapter is based on standard L2 driver and with the content director engine reduces CPU host system processing.

The Silicom 10 Gigabit Ethernet PCI Express content aware server adapter offers simple integration into any PCI Express X8 to 10Gigabit Network.

Key Features

Content Aware Director:

- Provides intelligent packet redirection capability where rules specify which packets are directed to the host system and which packets are directed to the other port (Bypass).
- Provides intelligent packet redirection capability where all packets are directed to the other port (Bypassed) and rules that specify which packets are copied to the host system (TAP).
- Provides intelligent packet filtering / drop capability where rules specify which packets are directed to the host or dropped.
- Provides redirection rules that can be defined using source IP/ destination IP / Source Port / Destination Port / VLAN tuples.
- Redirection and packet filtering / drop are performed by the hardware itself in wire speed and do not require any software and CPU host system power processing.
- Intelligent redirect mechanism is controllable via software.
- Intelligent routing mechanism is controllable via software.

Bypass / Disconnect:

- Bypass / Disconnect Ethernet ports on Power Fail, System Hangs or Software Application Hangs.
- Software programmable Bypass, Disconnect or Normal Mode.
- On Board Watch Dog Timer (WDT) Controller.
- Software programmable time out interval.
- Software Programmable WDT Enable / Disable counter.
- Software programmable Bypass Capability Enable / Disable.
- Software Programmable Disconnect Capability Enable / Disable.
- Software Programmable mode (Bypass, Normal or Disconnect mode) at Power up.
- Software Programmable mode (Bypass, Normal mode) at Power off.
- Independent Bypass operation in every two ports.
- Emulates standard NIC

-SRD: Fiber 1/10 Gigabit Ethernet 1000Base-SX / 10GBASE-SR:

- 10 Gigabit Fiber Ethernet port supports 10GBASE-SR (850nm LAN PHY)
- 1Gigabit Fiber Ethernet port supports 1000BASE-SX (850nm LAN PHY)
- 1/10Gigabit 850nm Small form Factor Pluggable (SFP+)

-LRD: Fiber 1/10 Gigabit Ethernet 1000Base-LX / 10GBASE-LR:

- 10 Gigabit Fiber Ethernet port supports 10GBASE-LR (1310nm LAN PHY)
- 1Gigabit Fiber Ethernet port supports 1000BASE-LX (1310nm LAN PHY)
- 1/10Gigabit 1310nm Small form Factor Pluggable (SFP+)

Common Key features:**Host Interface:**

- PCI Express X8 lane
- Support PCI Express Base Specification Revision 3.0, 8GT/s, 5GT/s or 2.5GT/s
- SR-IOV enabled to expose to 4 Virtual Functions (VF)

Intel FM10420 Features:

- Single-element 4MB shared memory
- L2/L3/L4/OpenFlow forwarding & ACLs
- Stateless load balancing to CPUs
- Datacenter Bridging (lossless Ethernet)
- 32K 40-bit TCAM entries
- 16K MAC & NextHop tables
- Up to 300Gbps High-bandwidth CPU interface
- 2x 50Gbps 8-lane PCIe interfaces
- 4x 25Gbps 4-lane PCIe interfaces
- Up to 8 25G/10G/2.5G/1G ports
- Up to 2 40G (4 x 10G)
- Up to 2 100G (4 x 25G)
- 300ns network latency (100GbE)
- 1000nS host-network latency

LAN Features:

- 256 queues per PCIe x8 interface
- SR-IOV (64 VFs per PCIe x8 interface)
- IP/TCP/UDP checksum
- Receive side scaling (RSS)
- TCP segmentation offload (TSO/LSO)
- LEDs indicator for link/Activity

Technical Specifications

Bypass Specifications:	
WDT Interval (Software Programmable):	3,276,800 mSec (3,276.8 Sec): Maximum 100 mSec (0.1 Sec) : Minimum WDT Interval = (2^wdt_interval_parameter)*(0.1) sec. wdt_interval_parameter: { Valid Range: 0-15}
Fiber 10 Gigabit Ethernet Technical Specifications – (10GBASE-SR) Adapters:	
IEEE Standard / Network topology:	Fiber Gigabit Ethernet, 1000Base-SX (850nM)
Data Transfer Rate:	10.3125GBd
Cables and Operating distance: Up to:	Multimode fiber: 62.5um, 160MHz/Km 13m* 62.5um, (OM1)200MHz/Km 16.5m * * Defined as half as the distance as specified in the optical transceiver
Optical Output Power:	Normal Mode (Bypass Off): Typical: TBD dBm (TX –Switch Normal – Fiber – LC/LC) Minimum: -7.3 dBm
Optical Receive Sensitivity:	Normal Mode (Bypass Off) Typical: TBD dBm Maximum: -11.1 dBm
Insertion Loss:	Bypass Mode: Insertion loss (Optical Power attenuation between TX to RX) (LC- fiber- switch- LC) Typical: TBD dB (From RX to TX) Maximum 1.9 dB
Fiber 10 Gigabit Ethernet Technical Specifications – (10GBASE-LR) Adapters:	
IEEE Standard / Network topology:	Fiber 10Gigabit Ethernet, 10GBASE-LR (1310nM LAN PHY)
Data Transfer Rate:	10.3125GBd
Cables and Operating distance: Up to:	Single-Mode: 5Km at 9um * * Defined as half as the distance as specified in the optical transceiver
Optical Output Power:	Normal Mode (Bypass Off): Typical: TBD dBm (TX –Switch Normal – Fiber – LC/LC) Minimum: -7.3 dBm
Optical Receive Sensitivity:	Normal Mode (Bypass Off) Typical: TBD dBm Maximum: -11.1 dBm
Insertion Loss:	Bypass Mode: Insertion loss (Optical Power attenuation between TX to RX) (LC- fiber- switch- LC) Typical: TBD dB (From RX to TX) Maximum 1.6 dB

– SRD: Fiber 1000BASE-SX / 10GBASE-SR Technical Specifications:

Optical Output Power (1G):	Minimum: -10.9 dBm
Optical Receive Sensitivity (1G):	Maximum: -15.6 dBm
Insertion Loss (1G)	Maximum: +1.9 dBm
Output Transmit Power (10G):	Minimum: -7.3 dBm
Optical Receive Sensitivity (10G):	Maximum: -11 dBm
Insertion Loss (10G):	Maximum: 1.6 dBm

– LRD: Fiber 1000BASE-LX / 10GBASE-LR Technical Specifications:

Optical Output Power (1G):	Minimum: -10.8 dBm m
Optical Receive Sensitivity (1G):	Maximum: -19 dBm
Insertion Loss (1G):	Maximum: +0.5dBm
Output Transmit Power (10G):	Minimum: -5.2dBm
Optical Receive Sensitivity (10G):	Maximum: -12.6 dBm
Insertion Loss (10G):	Maximum: +0.5 dBm

Operating Systems Support

Operating system support:	Linux
----------------------------------	-------

General Technical Specifications

Interface Standard:	PCI-Express Base Specification Revision 3.0(8 GTs)
Board Size:	Standard height long add-in card 193.04mm X 111.15mm (7.6"X 4.376")
PCI Express Card Type:	X8 Lane
PCI Express Voltage:	+3.3V ±9%, +12V ± 8%
External Voltage from external PW jack:	+12V ± 8%
PCI Connector:	Gold Finger: X8 Lane
Controller:	Intel FM10420

Holder:	Metal Bracket
Weight:	390g (13.76 oz)
Power Consumption:	12.48W: Full Traffic 10G all ports 11.52W: No Traffic, Link all ports 11.40W: Bypass Mode 11.40W: Disconnect Mode
Operating Humidity:	0%–90%, non-condensing
Operating Temperature:	0°C – 40°C (32°F – 104°F), Air flow requirement 200FLM
Storage:	-40°C–65°C (-40°F–149°F)
EMC Certifications:	FCC 47CFR Part 15:2013, Subpart B Class B Conducted emissions Radiated emissions EN 55022: 2010, Class B Conducted disturbance at mains terminals Conducted disturbance at telecommunication port Radiated disturbance EN 61000-3-2: 2006+A1(09)+A2(09) Harmonic current emissions EN 61000-3-3: 2008 Voltage fluctuations and flicker EN 55024: 2010 Immunity to electrostatic discharge (ESD) Radiated immunity to radio frequency electromagnetic field Conducted immunity to electrical fast transients / bursts (EFT/ B) Conducted immunity to voltage surges Conducted immunity to disturbances induced by radio frequency field Conducted immunity to voltage dips and short interruptions
MTBF*:	66 (Years) * The prediction was performed for 40°C Ambient temperature, GB Environmental condition. The reliability prediction was performed in accordance with Telcordia SR-332.
LEDs	
LEDs:	(1) Tri- color LEDs per port : Turns on link 10G (Green), Turns on link 1G (Yellow), Blinks on activity. (1) Bi- color LED per segment (2 ports): Off on Normal Mode Turns on Bypass Mode (Green) Turns on Disconnect Mode (Yellow)
LEDs location:	LEDs are located on the PCB, visible via holes in the metal bracket. Each 4 Yellow/Blue 1G/10G link/Act LEDs (1 LED per port) are located under their own LC connector port. (Each LC contains 2 ports).The yellow/ green for disconnect and bypass LED is located between the 2 ports in each segment- in the middle under each LC.
Connectors:	(4) LC

Functional Description

Director – Content Aware Bypass

Silicom's 10 Gigabit Ethernet content aware director Provides intelligent packet redirection capability where rules specify which packets are directed to the host system and which packets are directed to the other port (Bypass).




Figure 1: Content Aware Bypass Functional Block Diagram

Figure 1 illustrates functional block diagram of content aware Bypass:

Packets received in port A and meet rule are directed to port B, other packets are directed to port C (Bypass).

Packets received in port C and meet rule are directed to port D, other packets are directed to port A (Bypass).

Director – Content Aware TAP

Silicom's 10 Gigabit Ethernet content aware director Provides intelligent packet redirection capability where all packets are directed to the other port (Bypassed) and rules specify which packets are copied to the host system (TAP).




Figure 2: Content Aware TAP Functional Block Diagram

Figure 2 illustrates functional block diagram of content aware TAP:

Packets received in port A and meet rule are directed to ports B and C (TAP), other packets are directed to port C (Bypass).

Packets received in port C and meet rule are directed to ports D and A (TAP), other packets are directed to port A (Bypass).

Director – Content Filtering NIC

Silicom's 10 Gigabit Ethernet content aware provides intelligent packet redirection capability where rules specify which packets are directed to the host or dropped.

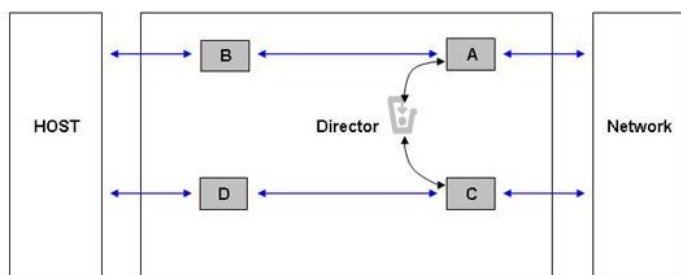


Figure 3: Content Aware Filtering NIC Functional Block Diagram

Figure 3 illustrates functional block diagram of content aware TAP:

Packets received in port A and meet rule, direct to port B. Packets received in port A and do not meet rule are dropped.

Packets received in port C and meet rule, direct to port D. Packets received in port C and do not meet rule are Dropt.

Load Balancing

Silicom’s 10 Gigabit Ethernet content aware director provides a load balancing of the traffic coming from the 4 x 10G external ports (0,2,4 and 6). The traffic is balanced, based on a defined hash configuration (5 tuple or other), to the 4x 10G internal interfaces (1,3,5 and 7)that are going to the host.

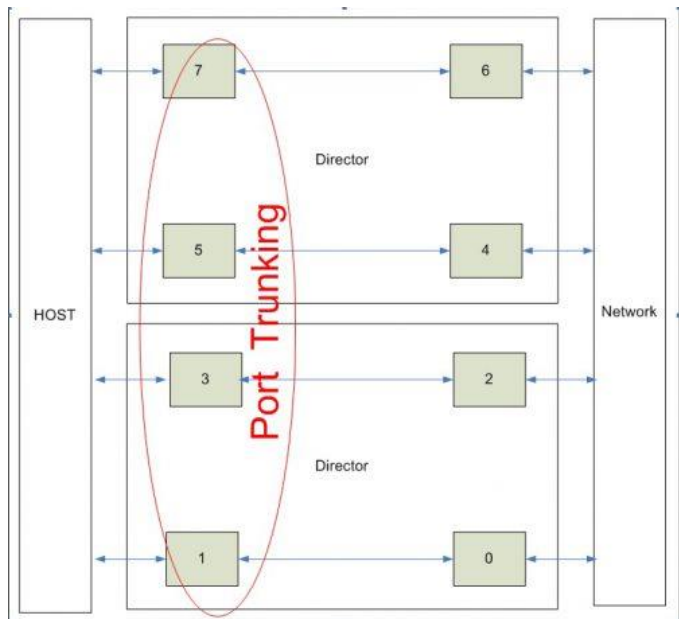


Figure 4: Load Balancing Functional Block Diagram

Figure 4 illustrates functional block diagram of Load Balancing:

Packets received in the 4 x 10G external ports (0,2,4 and 6) to be balanced based on a defined hash configuration (5 tuple or other), to the 4x 10G internal interfaces (1,3,5and 7) that are going to the host. If there is external port that is heavily loaded the traffic will go into the 4 internal interfaces with balanced load. An ISL tag is added to all incoming packets, it enables the host to know the source port, packets that are sent back from the host are stride from the ISL tag.

Tagged In-Line rule aware mode

Silicom’s 10 Gigabit Ethernet content aware director Provides intelligent packet redirection capability where rules specify which packets are directed to the host system and which packets are directed to the other ports (Bypass) but at the same time it will get these bypassed trficc into the host with a ISL tag marking that these packets are bypassed, per the rules that the host will issue to the Silicom’s 10 Gigabit Ethernet content aware director.




Figure 5: Tagged In-Line rule aware mode Block Diagram

Figure 5 illustrates functional block diagram of Tagged In-Line rule aware mode:

Port Group (0, 1), (2, 3), (4, 5), (6, 7) configured as VLAN groups.

Packets received in the 4 x 10G external ports (0,2,4 and 6) and meet and are directed to the other ports of the Vlan group with ISL tag that notify that match found. The original packet is sent to the “Output Port of Switch” in the rule matched policy

Packets received in the 4 x 10G external ports (0,2,4 and 6) and do not meet the rule are directed to the other ports of the Vlan group tagged with ISL tag that notify that no-match found.

Director: Rules Classification and capabilities

Director Capabilities

- The Redirector supports the following capabilities:
- Maximum total number of rules is 16K
- Each of the 16K rules can be defined to any port the on board multi-layer switch
- Each rule refers to incoming packet
- Rules are executed per order. First rule that matches will be executed
- Rules can be added and removed on the fly
- Each rule can include one or more classification fields. A rule match will be when all fields defined are match
- Each field can have a bit masking to check part of the classification field
- Per port statistics can be read, like packets count, errors, VLAN, and more
- Rules and action are done in wire speed at any packet size

Rules classification fields

Rules classification is done based on the first 128 bytes of the packets. The following list provides rules classification fields

- MAC address, source & destination
- IPv4 – source & destination IP
- IPv6 – source & destination IP
- L4 Port – source & destination port

- Ethernet Protocol – ethertype
- IP Protocol num
- VLAN ID tagging
- User defined fields
- DSCP – match the different services code point – the six most significant bits of the Type of Service octet (IPv4) or Traffic Class octet (IPv6).*
- IPv6 Flow Label*
- IP length*
- ISL Frame Type*
- ISL USER*
- Source & destination port range*
- VLAN priority*
- VLAN tag type*
- TCP flags*
- TOS – match Type of Service octet (IPv4) or Traffic Class octet (IPv6)*
- TTL field in a IPv4 header or Hop Limit in a IPv6 header*

*Future SW supports

Execution per rule

The following Executions per rule are supported:

- Drop – when a rule matches the packet will be dropped
- Redirect – when a rule matches redirect the packet to the defined destination port
- Mirror – when a rule matches copy packet also to a defined destination port

Director Advanced Features:

- Port trunking between the different Intel 10G ports to the Fulcrum 10G ports connected to it for load balancing between the different Intel ports
- Port trunking between the different External Fulcrum port to the external switch connected to it for load balancing between the different external ports
- Session balancing with L3/L4 hashing or other mechanism
- ISL (Inter Switch Link) Tagging per port can be added to the packets per configuration
- ISL Tagging can be removed and can be forward to specific port per the ISL index
- Quality of Service support with the following features:
 - Priority levels: 16 internal “switch” priorities, 8 or 16 VLAN priorities (optional use of CFI bit as an extra VLAN priority bit)
 - Arbitrary mapping of ingress VLAN priority to an internal VLAN priority

- Arbitrary mapping of an internal VLAN priority to egress VLAN priority
 - Arbitrary mapping of internal VLAN priority to switch priority
 - Arbitrary mapping of DSCP to switch priority, configurable priority source selection
 - Scheduler: 8 traffic classes, arbitrary mapping of switch priorities to traffic class, deficit weighted round-robin or strict priority
 - Notification: Two congestion notifications can be supported
 - Virtual output queue congestion notification (VCN) and Intel proprietary backward congestion notification (FCN)
 - Open Flow support (consistent with OpenFlow protocol standard)
 - sFlow support
 - User defined Packet transmission with two optional modes: 1. Simple mode – transmit on specific port. 2. Switched mode – where switch determines destination port/ports, or with specific information such as
 - whether or not egress processing rules should be applied
 - Storm Control Management – Switch can support a variety storm controller. Each storm controller can be programmable to define rat, condition (like unicast ICMP frames whose TTL is at most 1), frame type (can be
 - OR'ed), ingress & egress port ports. Actions: do nothing, drops frames to port (according to filter)
- *Future SW supports

Bypass / Disconnect

Bypass/Disconnect Mode of operation

Silicom's Bypass adapter supports the following mode states: Normal/inline, Bypass/Fail-To-Wire and Disconnect modes.

Normal/Inline mode

In Normal mode, the ports are independent interfaces (see Figure 6: Normal mode, one Bypass pair is illustrated).

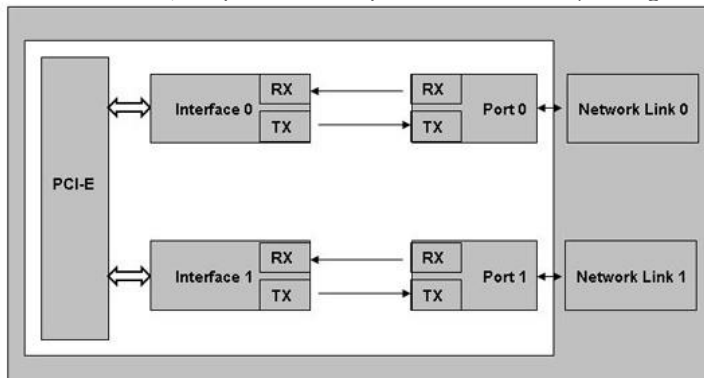
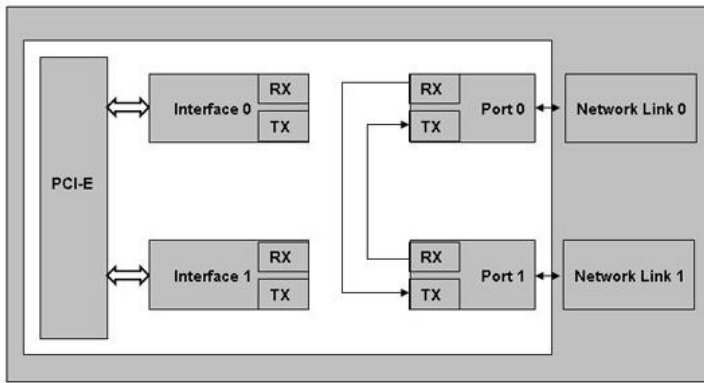


Figure 6: Normal Mode Functional Block Diagram

Bypass/Fail-To-Wire mode

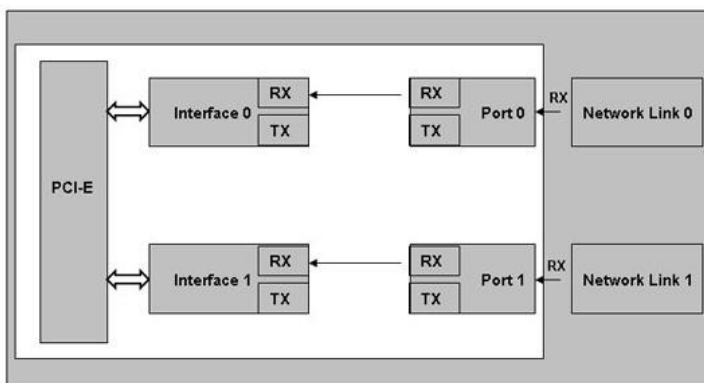
In Bypass mode, the connections of the Ethernet network ports are disconnected from the interfaces and switched over to the other port to create a crossed connection loop-back between the Ethernet ports. The connections of the interfaces are left unconnected. (See Figure 7: one Bypass pair illustrated)



Bypass Mode Functional Block Diagram

Disconnect mode

In Disconnect mode, the transmit connections of the interfaces are disconnected from the ports. The switch / router connected to the adapter does not detect link partner (See Figure 8):



Disconnect Mode Functional Block Diagram

Bypass/Disconnect Features

Silicom's Bypass server adapter supports software programmable to select Normal, Bypass or Disconnect modes. Silicom's Bypass adapters supports Disable Bypass, Disable Disconnected capabilities; hence, if those adapters receive Disable Bypass capability / Disable Disconnect commands, the adapter does not Bypass / does not Disconnect its Ethernet ports, The Disable Bypass Capabilities are reserved also after power off. This feature enables to emulate a standard NIC.

Key Features:

- All modes of operation are fully software configured with a well-defined and easy to use API
- Disable bypass/Disconnect for a standard NIC emulation and operation, retained even with power cycle
- Range of Watchdog timer timing, configurable for detecting appliance failure
- Watchdog timer reset function by application or self-reset by the product bypass driver
- Easy to read product status and product features capabilities
- Fast switching between mode with less than 10mS of transit time
- Compatible with all Silicom bypass products, one driver set and command set for all products

Order Information

P/N	Description	Notes
PE310G4DBIR-SRD	Quad port Fiber (SRD) 10 Gigabit Ethernet PCI Express Content Director Server Adapter	RoHS Compliant, X8 Gen 3, based on Intel FM10420, Standard height long add-in card 193.04mm X 111.15mm (7.6"X 4.376")
PE310G4DBIR-LRD	Quad port Fiber (LRD) 10 Gigabit Ethernet PCI Express Content Director Server Adapter	RoHS Compliant, X8 Gen 3, based on Intel FM10420, Standard height long add-in card 193.04mm X 111.15mm (7.6"X 4.376")

1V7