



PE3100G2DQIRL

Dual port Fiber 100 Gigabit Ethernet PCI Express Content Director Server Adapter Intel® based

Product Description

Silicom's 100/40 Gigabit Ethernet PCI Express content aware director server adapters are 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 100/40 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 100/40 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 100/40 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 100/40 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 100/40 Gigabit Ethernet content aware provides intelligent packet redirection capability where rules specify which packets are directed to the host or dropped.

The Silicom 100/40 Gigabit Ethernet PCI Express content director server adapter is based on Intel FM10420 Ethernet controller and a L3 switch router. The Silicom's 100 Gigabit Ethernet PCI Express adapter is based on standard L2 driver and with the content director engine reduces CPU host system processing.

The Silicom 100/40 Gigabit Ethernet PCI Express content aware server adapter offers simple integration into any PCI Express 2X8 to 100Gigabit 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.
- Support up to 2x100G / QSFP28 ports.
- 100G QSFP28 Ports support 100GBase-SR4 and 100GBase-LR4.
- 40G QSFP28 Ports support 40GBase-SR4 and 40GBase-LR4

Common Key features:

- Host Interface:
- PCI Express 2X8 lane
- Support PCI Express Base Specification Revision 3.0, 8GT/s, 5GT/s or 2.5GT/s

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
- 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

-QX4: QSFP28 100/40Gigabit Ethernet Technical Specifications Adapters:

QSFP28 (Quad Small Form-factor Pluggable) supports:	<p>40GBASE: XLPPPI interfaces according IEEE 802.3ba standard to support 40GBase-SR4, 40GBase-LR4 and 40GBase-CR4 interfaces.</p> <p>100GBASE: CAUI-4 interfaces according IEEE 802.3bm and IEEE 802.3bj standards , to support 100GBase-SR4, 100GBase-LR4 and 100GBase-CR4 interfaces.</p>
--	---

- ZS4: Fiber 100GBASE-SR4 Ethernet Technical Specifications:

IEEE Standard / Network topology:	Fiber Gigabit Ethernet, 100GBase-SR4 (850nm)
Data Transfer Rate:	103.125GBd
Cables and Operating distance: Up to:	Multimode fiber: 62.5um, (OM4) 100m

- ZL4: Fiber 100GBASE-LR4 Ethernet Technical Specifications:

IEEE Standard / Network topology:	Fiber 100Gigabit Ethernet, 100GBASE-LR4 (1310nm)
Data Transfer Rate:	103.1GBd
Cables and Operating distance: Up to:	Single-Mode: 10km

- QS41: Fiber 40GBASE-SR4 Ethernet Technical Specifications:

IEEE Standard / Network topology:	Fiber 40Gigabit Ethernet, 40GBASE-SR4 (840 to 860 nm LAN PHY). IEEE 802.3ba
Data Transfer Rate:	10.5 GBd per lane
Cables and Operating distance:	50um, (OM3) 1500 MHz*Km, 0.5 to 100 m 50um, (OM4) 3500 MHz*Km, 0.5 to 150 m
Output Transmit Power:	Maximum: 2.4 dBm per lane Minimum: -7.6 dBm per lane
Optical Receive Sensitivity:	Minimum: -5.4 dBm
Maximum Input Power:	Maximum: 2.4 dBm

- QS43: Fiber 40GBASE-SR4 Ethernet Technical Specifications:

IEEE Standard / Network topology:	Fiber 40Gigabit Ethernet, 40GBASE-SR4 (840 to 860 nm LAN PHY). IEEE 802.3ba
Data Transfer Rate:	10.5 GBd per lane

Cables and Operating distance:	50um, (OM3) 1500 MHz*Km, 0.5 to 300 m 50um, (OM4) 3500 MHz*Km, 0.5 to 400 m
Output Transmit Power:	Maximum: 0.5 dBm per lane Minimum: -7.5 dBm per lane
Optical Receive Sensitivity:	Minimum -7.5dBm
Maximum Input Power:	Maximum: 2.4 dBm
-QL4: Fiber 40GBASE-LR4 Ethernet Technical Specifications:	
IEEE Standard / Network topology:	Fiber 40Gigabit Ethernet, 40GBASE-LR4 (1264.5nm – 1277.5nm ; 1284.5nm – 1297.5nm ; 1304.5nm – 1317.5nm ; 1324.5nm – 1337.5nm LAN PHY). IEEE 802.3ba
Data Transfer Rate:	10.3125 GBd per lane
Cables and Operating distance:	SMF-28, 10Km
Output Transmit Power:	Maximum: 2.3 dBm per lane Minimum: -7.0dBm per lane
Optical Receive Sensitivity:	Maximum: -9.6 dBm
Maximum Input Power:	Maximum: 2.3 dBm
Operating Systems Support	
Operating system support:	Linux
General Technical Specifications	
Interface Standard:	PCI-Express Base Specification Revision 3.0(8 GTs)
Board Size:	Low profile short add-in card 167.64mm X 68.91mm (6.60"X 2.713")
PCI Express Card Type:	2x X8 Lane
PCI Express Voltage:	+12V ± 8%
PCI Connector:	Gold Finger: X16 Lane
Controller:	Intel FM10420
Holder:	Metal Bracket
Power Consumption –QX4:	15.48W
Power Consumption –ZS4:	19.68W
Power Consumption –ZL4:	19.32W
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
LEDs	
LEDs:	Two LEDs per port (1) Link 100/40/25/10Gbps LED: Turns on – link 100/25Gbps(Yellow) 40/10Gbps(Green) (1) Link/Act 100/40/25/10Gbps LED: Turns on – link, Blink – ACT 100/40/25/10Gbps(Green)
LEDs location:	LEDs are located on the PCB, visible by light guide in the metal bracket
Connectors:	(2) MPO

Functional Description

Silicom's 100 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

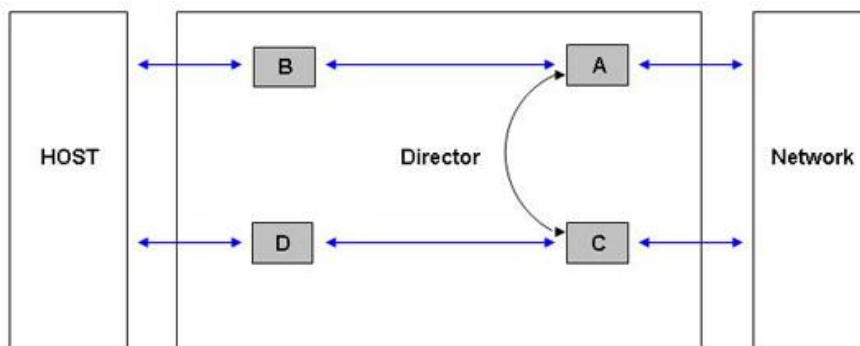


Figure 1 illustrates functional block diagram of content aware Bypass:

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 Awar 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).

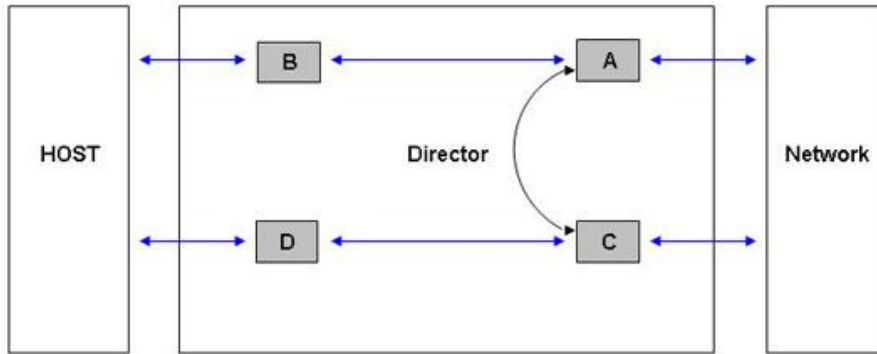


Figure 2: Content Aware TAP Functional Block Diagram

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.

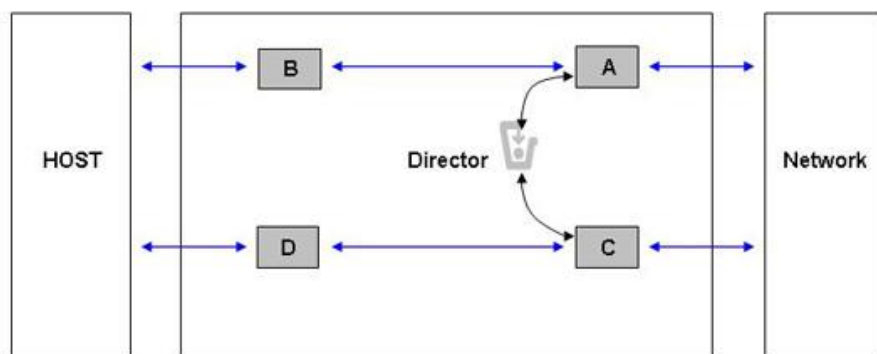


Figure 3: Content Aware Filtering NIC Functional Block Diagram

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.

Director – 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

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 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

Order Information

P/N	Description	Notes
PE3100G2DQIRL-QX4	Dual port Fiber 100 Gigabit Ethernet PCI Express Content Director Server Adapter	RoHS Compliant, 2x8 Gen 3, based on Intel FM10420
PE3100G2DQIRL-ZS4	Dual port Fiber (SR4) 100 Gigabit Ethernet PCI Express Content Director Server Adapter	RoHS Compliant, 2x8 Gen 3, based on Intel FM10420

PE3100G2DQIRL-ZL4	Dual port Fiber (LR4) 100 Gigabit Ethernet PCI Express Content Director Server Adapter	RoHS Compliant, 2x8 Gen 3, based on Intel FM10420 Only Support –LR4 with FEC enabled. As this is not part of IEEE standard, link partner should capable to support FEC enabled
PE3100G2DQIRL-QS41	Dual port Fiber (SR4) 40 Gigabit Ethernet PCI Express Content Director Server Adapter	RoHS Compliant, 2x8 Gen 3, based on Intel FM10420 on board support for Fiber SR4 100m on OM3 MMF
PE3100G2DQIRL-QS43	Dual port Fiber (LR4) 40 Gigabit Ethernet PCI Express Content Director Server Adapter	RoHS Compliant, 2x8 Gen 3, based on Intel FM10420 on board support for Fiber SR4 300m on OM3 MMF
PE3100G2DQIRL-QL4	Dual port Fiber (LR4) 40 Gigabit Ethernet PCI Express Content Director Server Adapter	RoHS Compliant, 2x8 Gen 3, based on Intel FM10420

1V4