



PE325G2DSIR

Dual Port Fiber 25/10G Ethernet PCIe Director Server Adapter Intel® FM10420 Based

Product Description

Silicom's 25 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 25 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 25 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 25 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 25 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 25 Gigabit Ethernet content aware provides intelligent packet redirection capability where rules specify which packets are directed to the host or dropped.

The Silicom 25 Gigabit Ethernet PCI Express content director server adapter is based on Intel FM10420 Ethernet controller and a L2/L3 switch router.

The Silicom's 25 Gigabit Ethernet PCI Express adapter is based on standard L2/L3 driver and with the content director engine reduces CPU host system processing. The Silicom 25 Gigabit Ethernet PCI Express content aware server adapter offers simple integration into any PCI Express X8 to 25Gigabit 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 2x25G / SFP28 ports.
- 25G SFP28 Ports support 25GBase-SR and 25GBase-LR.

SFP28 25Gigabit Ethernet:

- 25Gigabit Ethernet Adapter with SFP cage support:
 - ZS: **Fiber 25 Gigabit Ethernet 25GBASE-SR:**
 - 25GBase-SR with 25G 850nM Small Form Factor Pluggable (SFP28)
 - ZL: **Fiber 25 Gigabit Ethernet 25GBASE-LR:**
 - 25GBase-LR with 25G 1310nM Small Form Factor Pluggable (SFP28)
- ZS: **Fiber 25 Gigabit Ethernet 25GBASE-SR:**
 - 25 Gigabit Fiber Ethernet port supports 25GBASE-SR (850nM LAN PHY)
- 25Gigabit 850nM Small form Factor Pluggable (SFP28)
- ZL: **Fiber 25 Gigabit Ethernet 25GBASE-LR:**
 - 25 Gigabit Fiber Ethernet port supports 25GBASE-LR (1310nM LAN PHY)
- 25Gigabit 1310nM Small form Factor Pluggable (SFP28)

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

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
- 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

SFP28 25Gigabit Ethernet Technical Specifications Adapters:

SFP28 (Small Form Factor Pluggable) supports:	SFI interfaces supports 25GBase-R PCS and 25 Gigabit PMA in order to connect with SFP28 to 25GBase-SR
25GBase-SR SFP28: IEEE Standard / Network topology:	Fiber 25Gigabit Ethernet, 25GBASE-SR (850nm LAN PHY)
25GBase-SR SFP28: Data Transfer Rate:	25.78125Gbps
25GBase-SR SFP28: Cables and Operating	35m on 50/125µm MMF 100m over M5F MMF (50/125 um OM4)

distance Up to:	70m over M5E MMF (50/125um OM3)
25GBase-LR SFP28: IEEE Standard / Network topology:	Fiber 25Gigabit Ethernet, 25GBASE-LR (1310nm LAN PHY)
25GBase-LR SFP28: Data Transfer Rate:	25.78125Gbps
25GBase-LR SFP28: Cables and Operating distance Up to:	10km on 9/125um
-ZS: Fiber 25GBASE-SR Technical Specifications:	
Output Transmit Power (25G):	Minimum: -9.1 dBm Maximum: 2.4 dBm
Optical Receive Sensitivity (25G):	Maximum: -5.6 dBm
-ZL: Fiber 25GBASE-LR Technical Specifications:	
Output Transmit Power (25G):	Minimum: -5.0 dBm Maximum: 3.0 dBm
Optical Receive Sensitivity (25G):	Maximum: -11.4 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:	X8 Lane
PCI Express Voltage:	+12V ± 8%
PCI Connector:	Gold Finger: X8 Lane

Controller:	Intel FM10420
Holder:	Metal Bracket
Operating Humidity:	0%–90%, non-condensing
Operating Temperature:	0°C – 40°C (32°F – 104°F), Air flow requirement 200LFM
Storage:	-40°C–65°C (-4°F–149°F)
EMC Certifications:	<p>FCC 47CFR Part 15:2013, Subpart B Class B</p> <p>Conducted emissions</p> <p>Radiated emissions</p> <p>EN 55022: 2010, Class B</p> <p>Conducted disturbance at mains terminals</p> <p>Conducted disturbance at telecommunication port</p> <p>Radiated disturbance</p> <p>EN 61000-3-2: 2006+A1(09)+A2(09)</p> <p>Harmonic current emissions</p> <p>EN 61000-3-3: 2008</p> <p>Voltage fluctuations and flicker</p> <p>EN 55024: 2010</p> <p>Immunity to electrostatic discharge (ESD)</p> <p>Radiated immunity to radio frequency electromagnetic field</p> <p>Conducted immunity to electrical fast transients / bursts (EFT/ B)</p> <p>Conducted immunity to voltage surges</p> <p>Conducted immunity to disturbances induced by radio frequency field</p> <p>Conducted immunity to voltage dips and short interruptions</p>
LEDs	
LEDs:	<p>(2) LEDs per port</p> <p>Link Led:</p> <p>Green on 25G link</p> <p>Blue on 10G Link</p> <p>Link/Act Led:</p> <p>Turns on link (Green),</p> <p>Blinks on activity (Green)</p>
LEDs location:	LEDs are located on the PCB, visible by light guide in the metal bracket
Connectors:	(2) SFP28 / zSFP+ Cages

Director Functional Description

Silicom's 25 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)

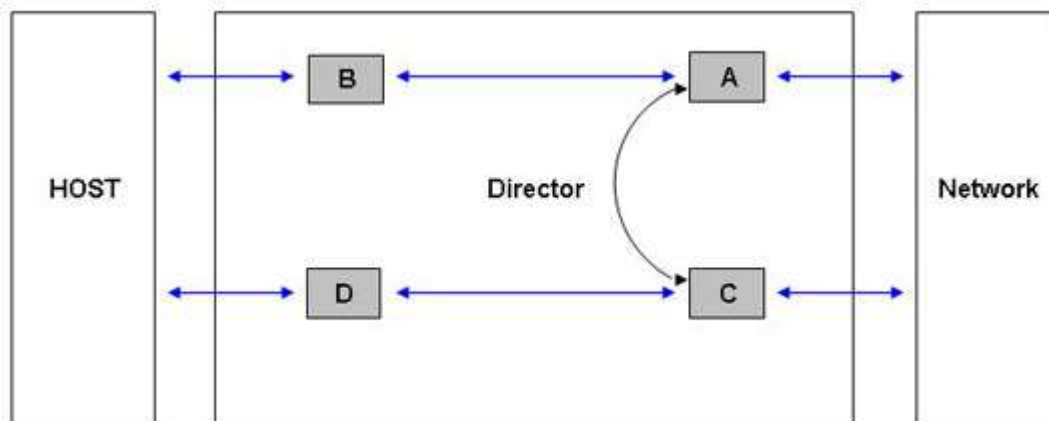


Figure 1: Content Aware Bypass Functional Block Diagram

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 25 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).

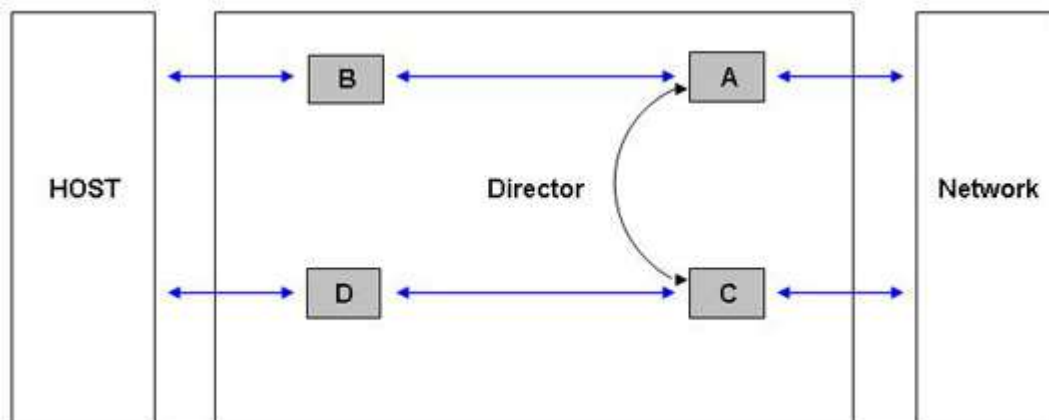


Figure 2: Content Aware TAP Functional Block Diagram

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

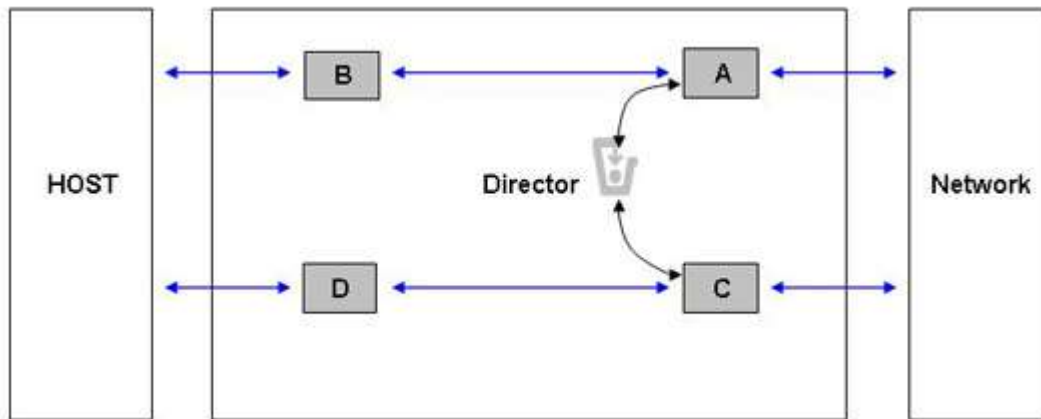


Figure 3: Content Aware Filtering NIC Functional Block Diagram

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: 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 10/25G ports to the Fulcrum 10/25G 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
PE325G2DSIR-XR	Dual Port Fiber 25/10G (DA) Ethernet PCIe Director Server Adapter	RoHS Compliant, x8 Gen 3, based on Intel FM10420
PE325G2DSIR-ZS	Dual Port Fiber 25G (SR) Ethernet PCIe Director Server Adapter	RoHS Compliant, x8 Gen 3, based on Intel FM10420
PE325G2DSIR-ZL	Dual Port Fiber 25G (LR) Ethernet PCIe Director Server Adapter	RoHS Compliant, x8 Gen 3, based on Intel FM10420
PE325G2DSIR-SR	Dual Port Fiber 10G (SR) Ethernet PCIe Director Server Adapter	RoHS Compliant, x8 Gen 3, based on Intel FM10420
PE325G2DSIR-LR	Dual Port Fiber 10G (LR) Ethernet PCIe Director Server Adapter	RoHS Compliant, x8 Gen 3, based on Intel FM10420