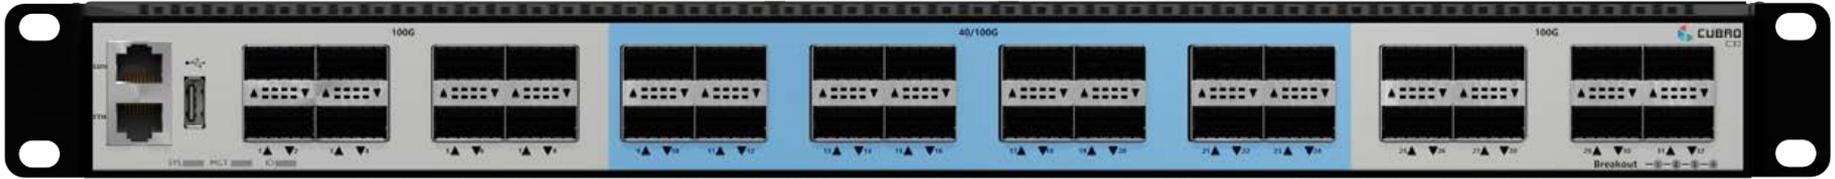


# Aggregator C32 Network Packet Broker

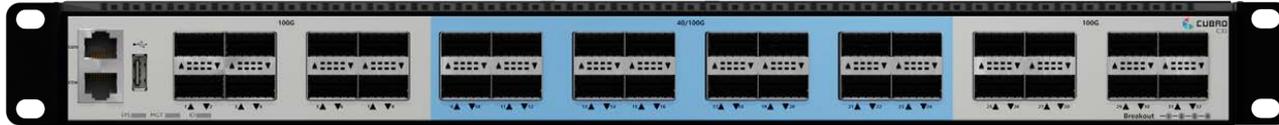
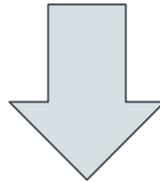
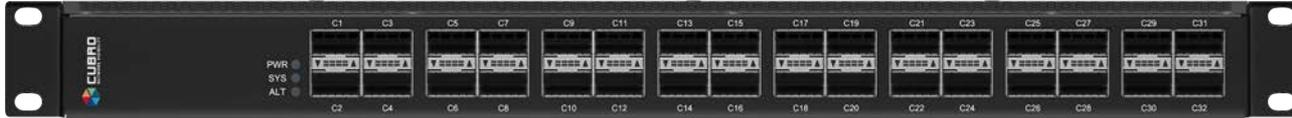
## APPLICATION NOTE

# Aggregator C32



- 16 x QSFP28 100G
- 16 x QSFP28 for 40G/100G – each of these ports can be used in 4 x 10G or 4 x 25G split mode
- When all ports are in split mode it supports 64 x 10G/25G + 16 x QSFP28 100G
- Each port can be used simultaneously as input and output and is totally independent from other ports
- Non-blocking architecture
- All ports are open – no software licence to enable ports

# AGG32100 replaced by C32



# Filtering capabilities



The Aggregator C32 supports up to 4000 parallel running IPv4/IPv6 filters. These filters can be used to redirect a selected part of the incoming traffic to a low bandwidth monitoring tool.

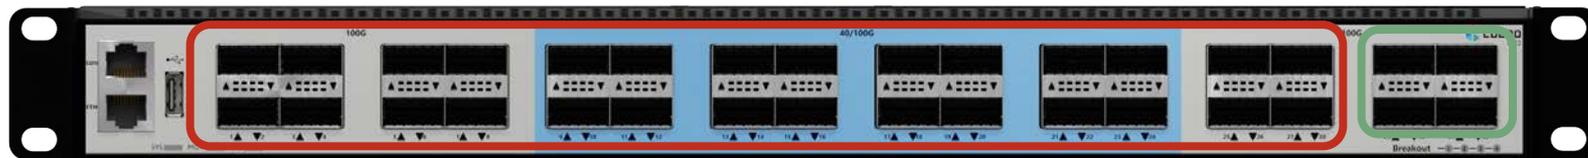
Filtering parameters include

Layer 2	Layer 3	Layer 4
MAC Src / Dst	IPv4 Src / Dst	Port Src / Dst
VLAN tag (QinQ)	IPv6 Src / Dst	TCP/UDP/SCTP/ ICMP
Ethertype	Protocol	

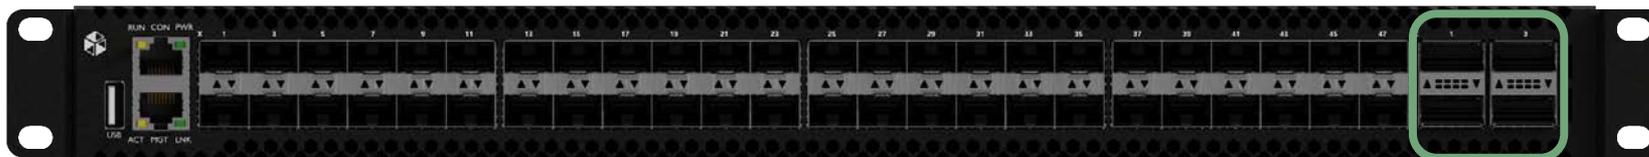
# Extending 100G interfaces for “advanced” NPB functions



The Cubro C32 can play a supporting role to extend the capabilities of more advanced Network Packet Brokers such as Omnia120, or EXA32100A. The simple and transparent design of the C32 makes it flexible when there is a requirement for additional 100G interfaces. The combination of C32 and the Omnia series can offer a powerful solution for layer 7 filtering. This means it is possible to filter on applications, keywords, or any wanted Regex.



C32 as 28: 4 aggregator

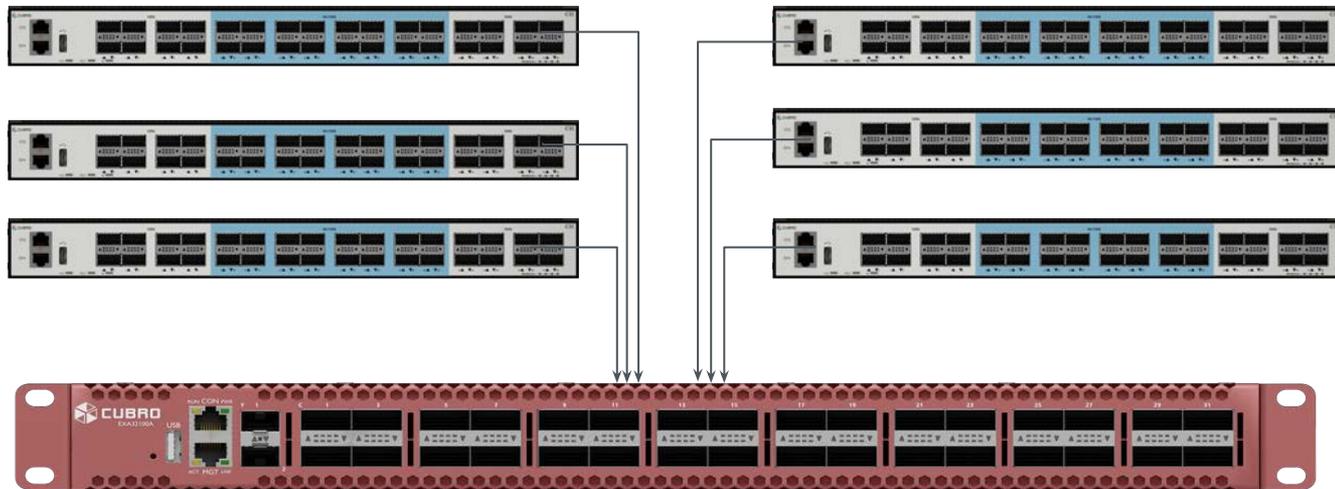


Omnia as advanced filtering / packet processing engine

Economic way to extend 100G port coverage



# Economical solution for high 100G port counts



Pre-aggregation to reduce input ports to EXA32100A/EXA64100

**Central** Packet Processing (filtering, load-balancing)

Central Filtering on EXA32100A allows easy management & operation

# General comparison to the Cubro EXA32100A



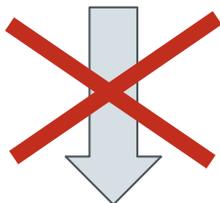
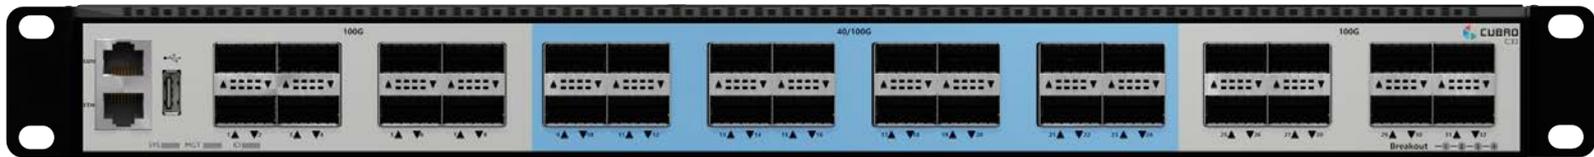
Cubro EXA32100A is the next level 100G Network Packet Broker which is designed to provide visibility in many network architectures and under several considerations. The C32 plays a supporting role to extend the quantity of 100G/40G interfaces, providing an attractive price per 100G interface.

The below comparison should help to decide when a EXA32100A or an Aggregator C32 is required:

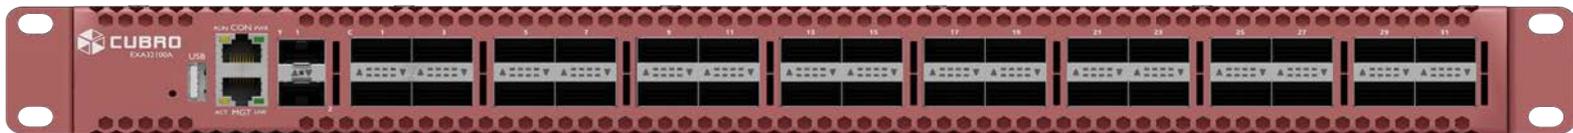
	Features	Aggregator C32	EXA32100A
<b>Hardware</b>	<b>Front Panel Ports</b>	16 x 40G/100G or 64 x 25G/10G ports 16 x 100G only	32 x 40G/100G; 2 x 25G or 130 x 25G/10G ports
	<b>Supported speed settings</b>	100G/40G/25G/10G	100G/50G/40G/25G/10G
	<b>Packet Buffer</b>	36 MB	22M
	<b>Number of fans</b>	4	5
	<b>Loop back port function</b>	NO	YES
	<b>Electrical footprint</b>	Input Power: 100-240V AC or 36-72 V DC Maximum Power Consumption: 224W	Input Power: 100-240 V AC or 36-72 V DC Maximum Power Consumption: 540W
<b>Traffic handling</b>	<b>Tunnel Removal</b>	NO	YES - GRE   VXLAN   MPLS   GTP   ERSPAN   many more
	<b>Packet Slicing</b>	NO	YES
	<b>Aggregation</b>	1:1; 1:n; n:1; n:n	1:1; 1:n; n:1; n:n
<b>Filtering</b>	<b>3 Stage Filtering</b>	NO	YES
	<b>Up to OSI Layer 4</b>	YES	YES
	<b>Filtering in OSI Layer 5</b>	NO	YES
	<b>Payload Filtering</b>	NO	YES
	<b>Filtering Fragmented traffic</b>	NO	YES
<b>Load-balancing</b>	<b>Flow-based Load-balancing</b>	YES	YES
	<b>Session-aware filtering (in virtual networks)</b>	NO	YES



# C32 → EXA32100A



No firmware upgrade from C32 to EXA32100A.  
**Totally different hardware**

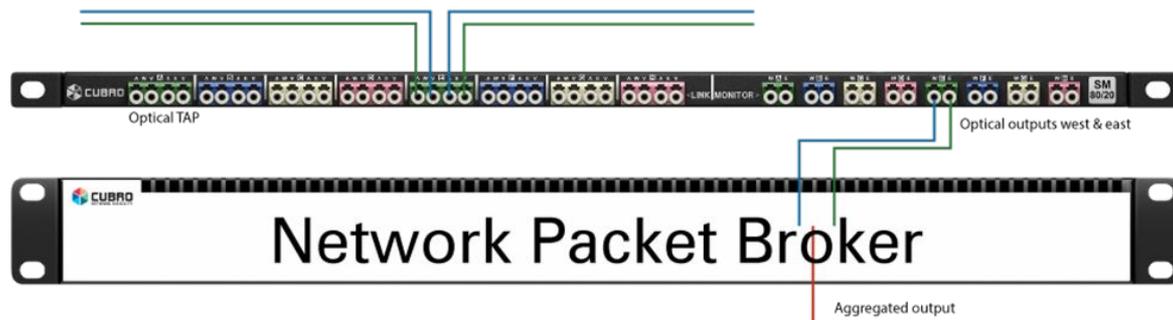


# Aggregation & Simple Filtering



The C32 can receive traffic from a single or multiple 40/100 Gbps link (s) via the monitoring ports of an inline tapping device. The incoming traffic can be further aggregated to a single or multiple outputs to connect analyzers and monitoring tools as required.

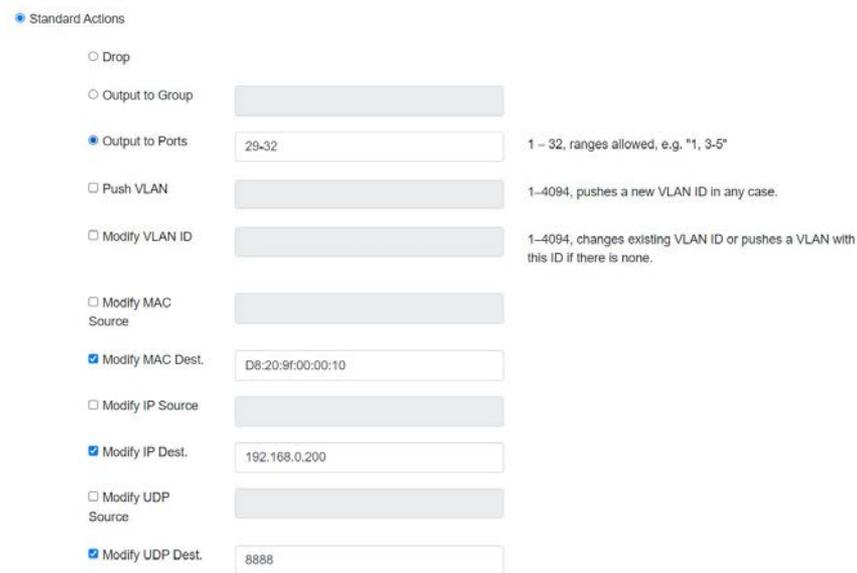
In the below example, the C32 aggregates up- and downstream traffic of a 100 Gbit link to a single output port for more economical usage of connected traffic probes/analytics systems.



By utilizing the various filtering capabilities of the C32 the user can further reduce traffic volume that needs to be processed, thus enabling quicker and more accurate analysis and troubleshooting. Moreover, incoming traffic can be VLAN tagged per physical port to allow easy identification at which physical port a packet original arrived.

The front-end programmable switch supports a wide range of different output actions so that filtered traffic can be provided to the appropriate tools.

- Forward Action
  - single port
  - multiple (parallel) ports
  - load-balanced ports
    - single load-balancing group
    - multiple load-balancing group
- Drop Action
  - delete filtered traffic
- Modify Egress traffic
  - Supports to modify header parameters like
    - MAC Src/Dest, IP Src/Dest, UDP/TCP Src/Dest.



The screenshot shows a configuration interface for 'Standard Actions'. It features a list of actions with radio buttons and input fields. The 'Output to Ports' action is selected and has the value '29-32' entered in its input field. Other actions like 'Drop', 'Output to Group', 'Push VLAN', 'Modify VLAN ID', 'Modify MAC Source', 'Modify IP Source', 'Modify UDP Source', and 'Modify UDP Dest.' are also visible but not selected. The 'Modify MAC Dest.' and 'Modify IP Dest.' actions are selected and have values 'D8:20:9f:00:00:10' and '192.168.0.200' respectively. The 'Modify UDP Dest.' action is selected and has the value '8888'. To the right of the input fields, there are descriptive notes for some actions, such as '1 - 32, ranges allowed, e.g. "1, 3-5"' for 'Output to Ports' and '1-4094, pushes a new VLAN ID in any case.' for 'Push VLAN'.

Distribute traffic as required

# CRC transparency



C32 platform not only allows 100% transparency to L2 protocols but also to CRC errors.

- Per default, the device will drop incoming CRC packets
- Via simple configuration option, the ingress and egress interfaces transparently receive and forward incoming CRC errors.
- This option allows the monitoring appliance to provide statistics about CRC errors of the live network

eth-0-9 (40G/100G QSFP+)

No

100 GBit

Force TX Up (Unidirectional Mode)

FEC

Only has an effect for 25Gbit/100Gbit.

Checksum Check

Checksum Recalculation

Activated

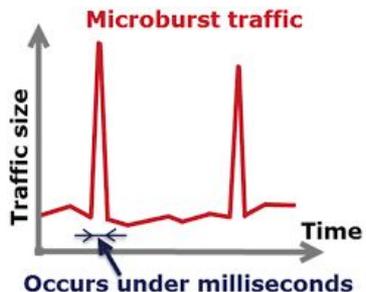
# Microburst detection & protection



## What are microbursts?

In Ethernet/IP networking, micro-bursting is a behavior seen on fast packet-switched networks, where rapid bursts of data packets are sent in quick succession, leading to periods of full line-rate transmission that can overflow packet buffers of the network stack, both in network endpoints and routers and switches inside the network. It can be mitigated by the network scheduler. In particular, micro-bursting is often caused by the use of the TCP protocol on such a network.

In an aggregation application where the aggregated output combines the traffic of several inputs it can happen that the egress port gets overloaded because of bursty input traffic – see following picture.



@Cubro

## How to overcome this problem?

If the input traffic can not be smoothed or shaped, the Aggregator needs to have a buffer that holds the data until there is again free bandwidth available to send the data. Basically two different concepts are available for buffering. The first one is to use a dedicated buffer per port while the second concept uses a centralized buffer that is available for any port that requires it. The C32 supports both concepts.

In the centralized buffer mode the Cubro C32 supports **36 MB of buffer**.

## Visibility via Port statistics

In the case of packet drop due to oversubscription, the port statistics of the dedicated ports will show a counter for the number of dropped packets. This is a useful indicator for the user, it signals that either the filtering needs to be extended or it is required to add more output ports for the forwarded traffic.

THANK YOU

