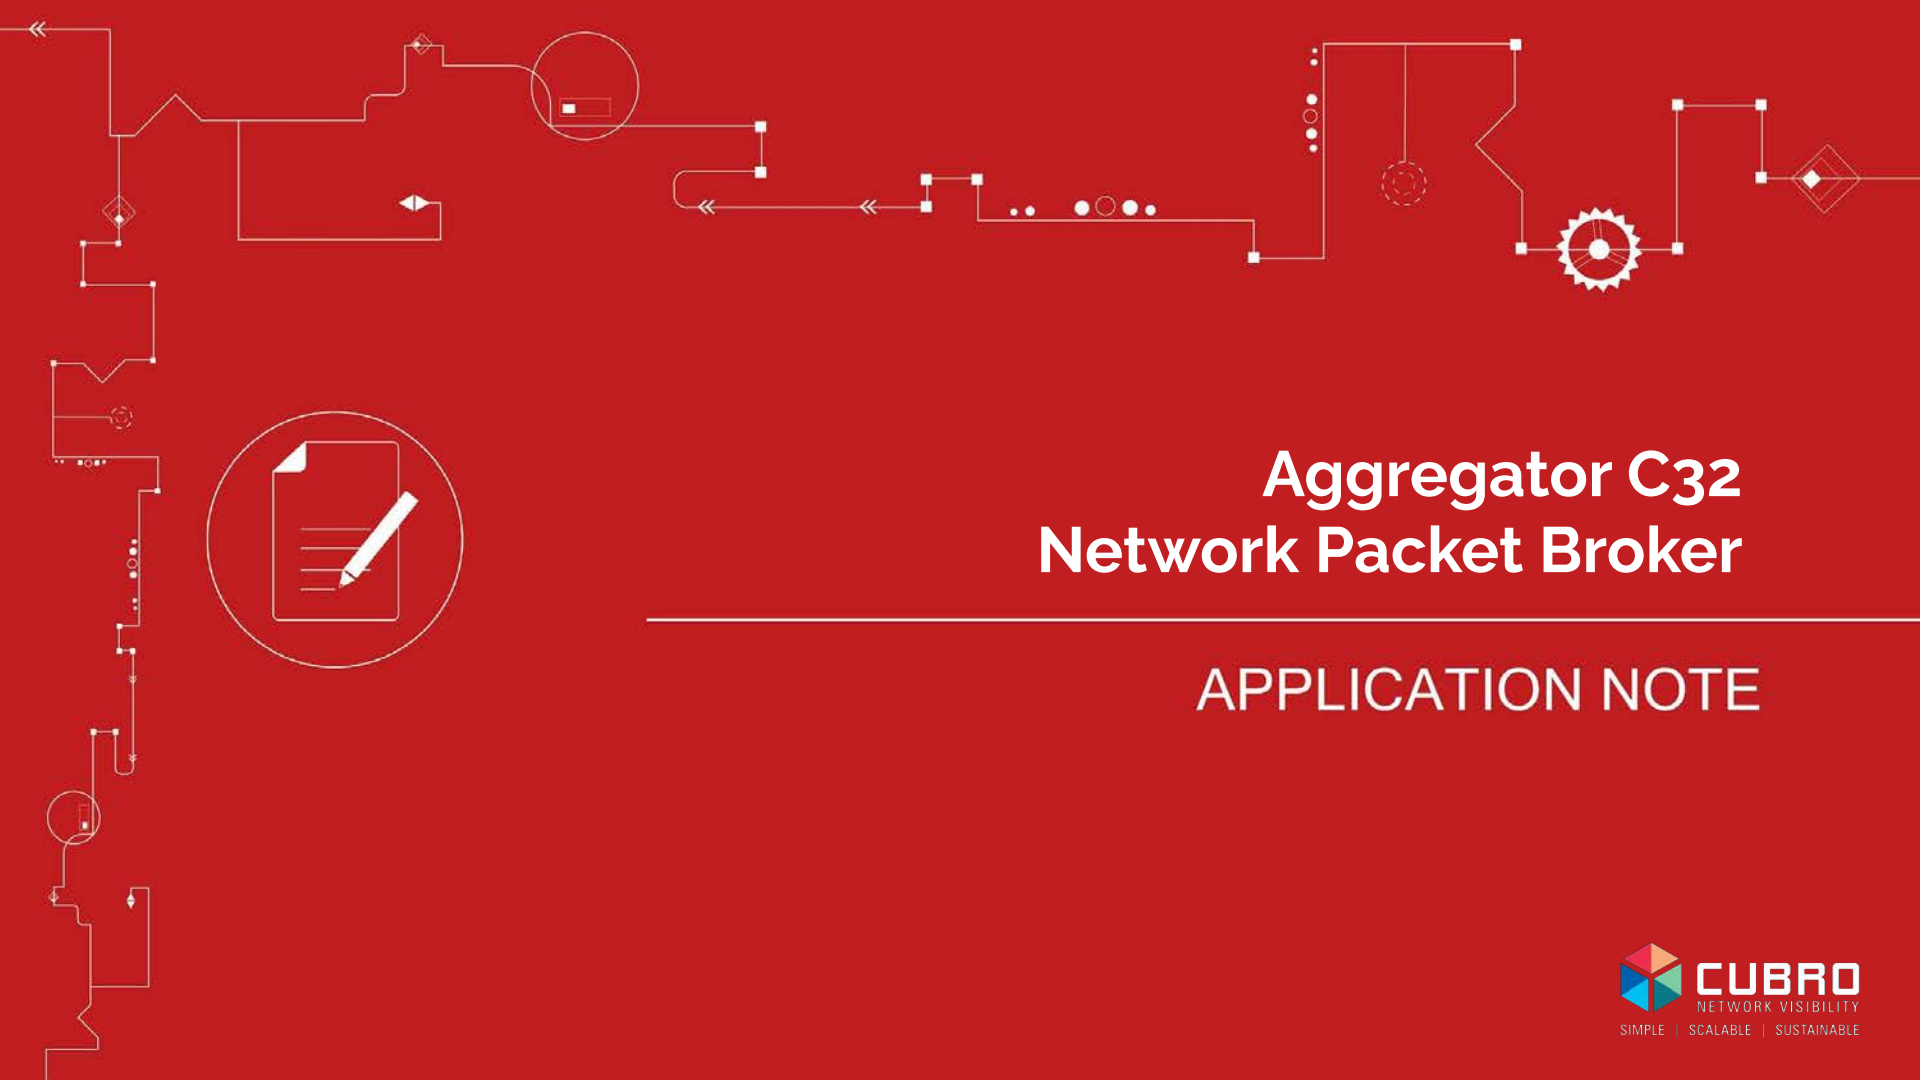

Aggregator C32
Network Packet Broker

APPLICATION NOTE

CUBRO
NETWORK VISIBILITY
SIMPLE | SCALABLE | SUSTAINABLE

Aggregator C32

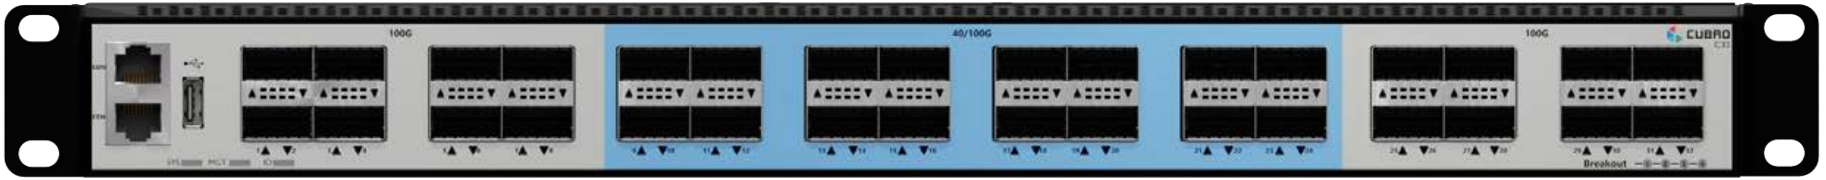

- 16 x 40G / 100G (QSFP/QSFP28)
- 16 x 40G / 100G (QSFP7QSFP28 with break-out to 4 x 10G / 25G
    - When all ports are in split mode it supports 64 x 10G / 25G
- Each port can be used simultaneously as input and output and is totally independent from other ports
- Non-blocking architecture
- All ports are open – no software license to enable ports

AGG32100 replaced by C32

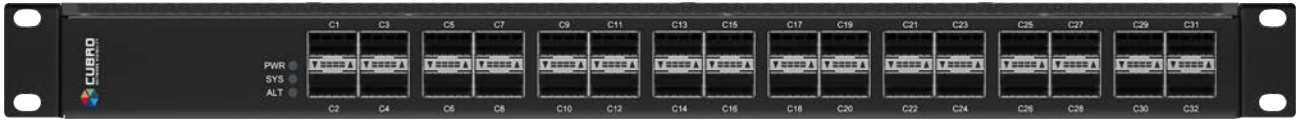

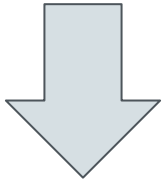

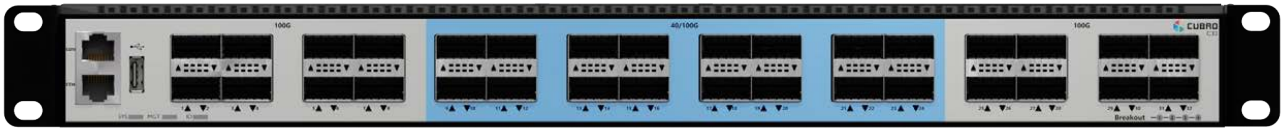

Filtering capabilities

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.

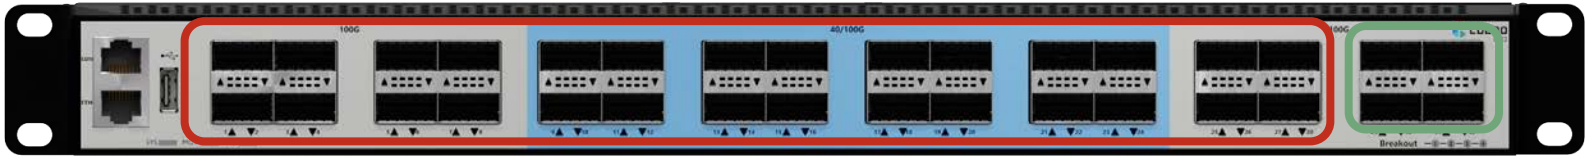

C32 as 28: 4 aggregator

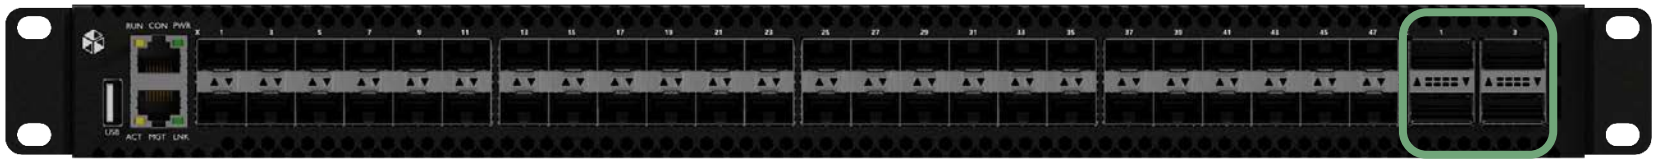

Omnia as advanced filtering / packet processing engine

Economic way to extend 100G port coverage

CUBRO NETWORK VISIBILITY
SIMPLE | SCALABLE | SUSTAINABLE

Economical solution for high 100G port counts

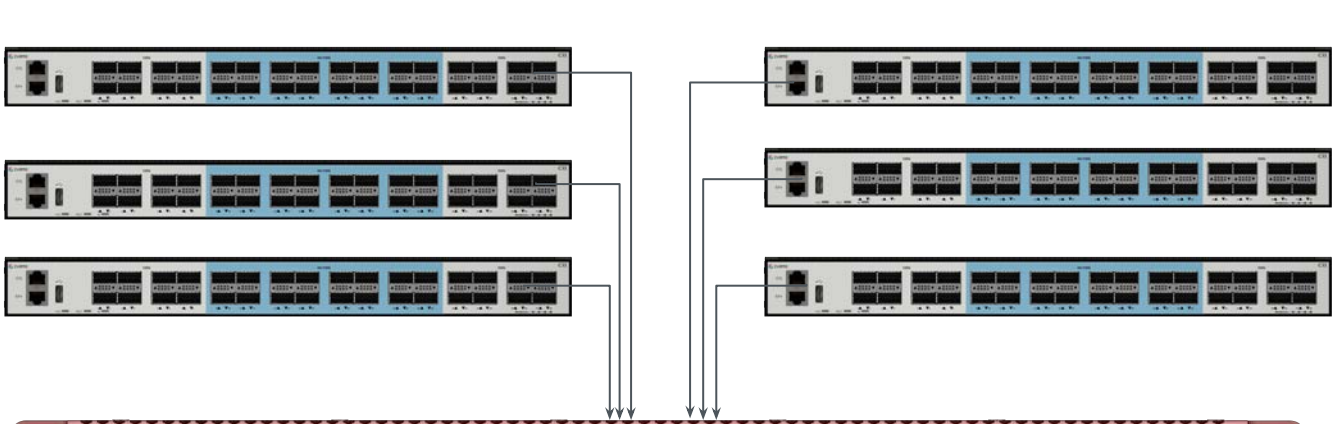

Pre-aggregation to reduce input ports to EXA32100A/EXA64100

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
Hardware	**Front Panel Ports**	16 x 40G/100G or 64 x 10G ports 16 x 100G only	32 x 40G/100G; 2 x 25G or 130 x 25G/10G ports
	Supported speed settings	100G/40G/10G	100G/50G/40G/25G/10G
	Packet Buffer	36 MB	22M
	Number of fans	4	5
	Loop back port function	NO	YES
	Electrical footprint	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
Traffic handling	**Tunnel Removal**	NO	YES - GRE | VXLAN | MPLS | GTP | ERSPAN | many more
	Packet Slicing	NO	YES
	Aggregation	1:1; 1:n; n:1; n:n	1:1; 1:n; n:1; n:n
Filtering	**3 Stage Filtering**	NO	YES
	Up to OSI Layer 4	YES	YES
	Filtering in OSI Layer 5	NO	YES
	Payload Filtering	NO	YES
	Filtering Fragmented traffic	NO	YES
Load-balancing	**Flow-based Load-balancing**	YES	YES
	Session-aware filtering (in virtual networks)	NO	YES

C32→ EXA32100A

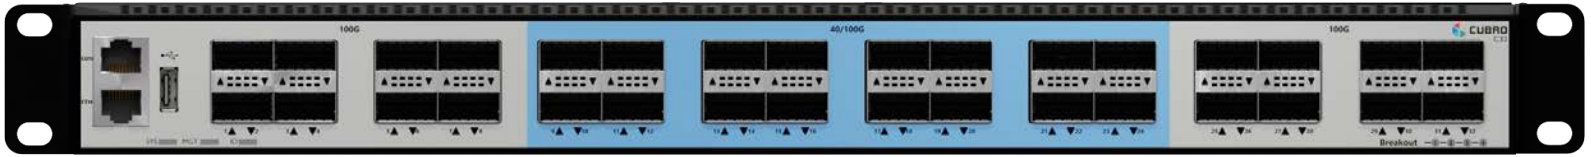

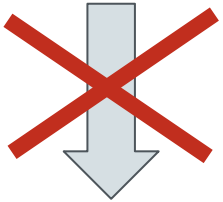

No firmware upgrade from C32 to EXA32100A.
Totally different hardware

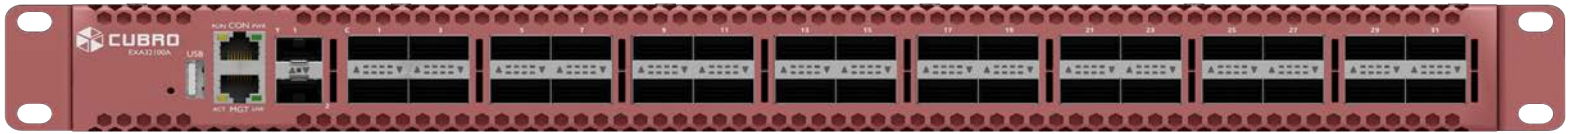

CUBRO NETWORK VISIBILITY
SIMPLE | SCALABLE | SUSTAINABLE

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.

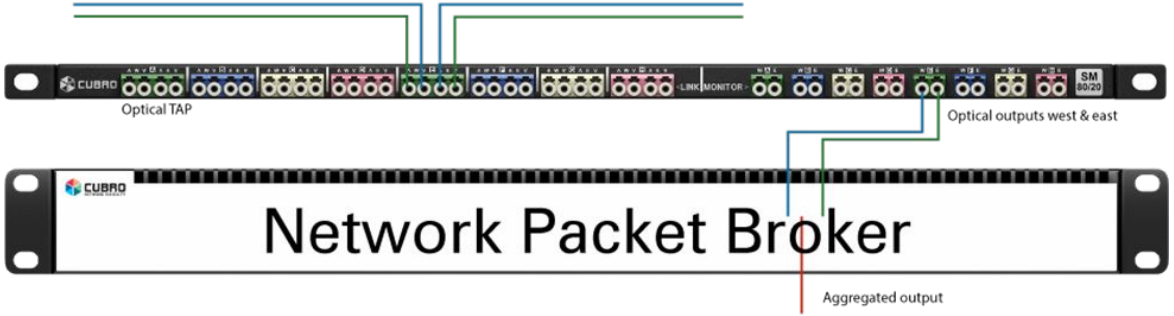

By utilizing the various filtering capabilities of the C32 the user can further reduce traffic volume that needs to be processed, thus enabling quicker and more accurate analysis and troubleshooting. Moreover, incoming traffic can be VLAN tagged per physical port to allow easy identification at which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

Output Actions

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.

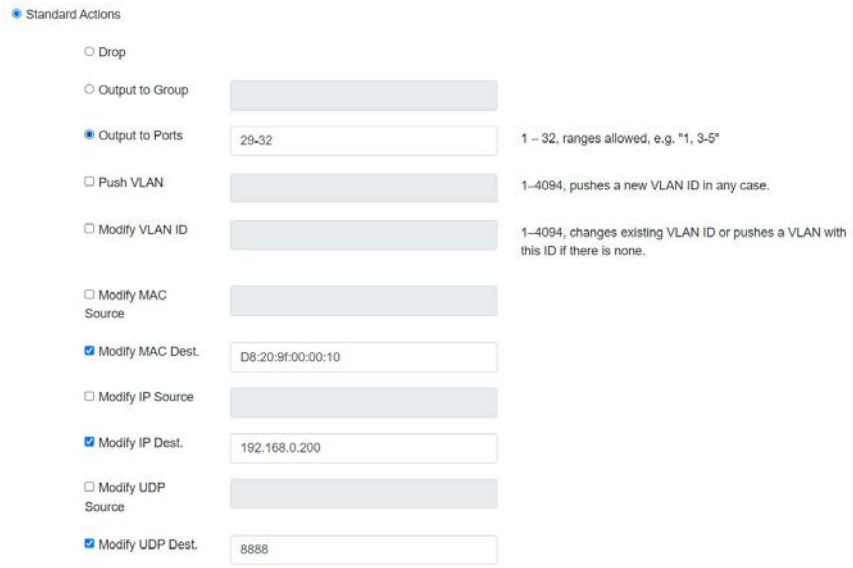

Distribute traffic as required

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

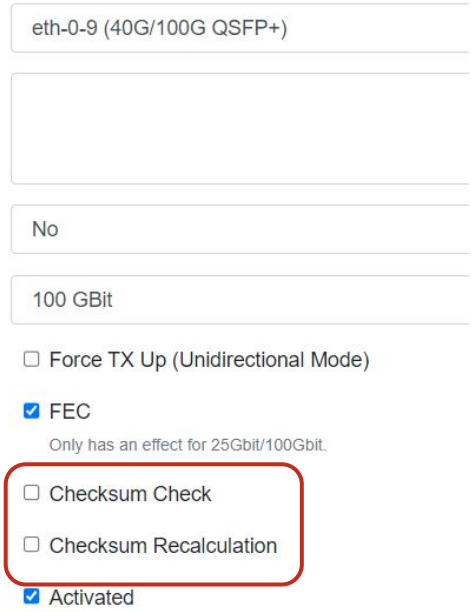

Microburst detection & protection

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.

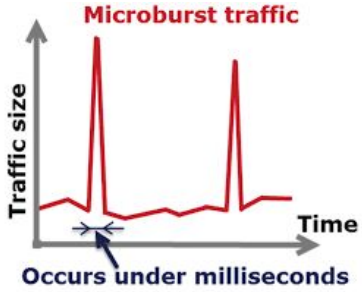

How to overcome this problem?

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