



CUBRO
NETWORK VISIBILITY

OMNIA & EXA48200 DEDUPLICATION


APPLICATION NOTE

13 MAR 2023

Author: Herbert Etlinger

TABLE OF CONTENTS

Deduplication Background.....	3
Deduplication on Cubro Omnia & EXA48200 family	4
Packet Size in Byte.....	5
Traffic Utilization in Gbit/s.....	5
Packets per second.....	5
Deduplication Operation.....	5
Deduplication Statistics	6
Summary	7



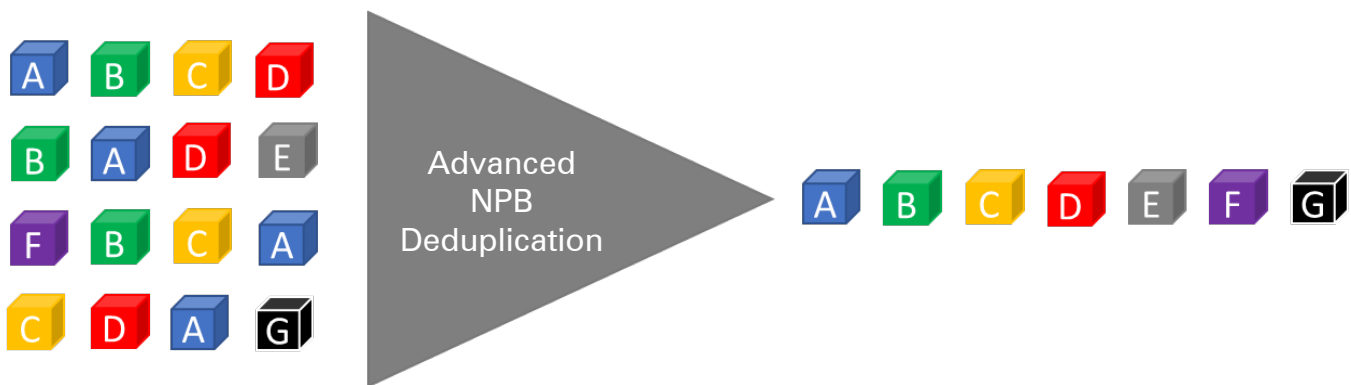
Deduplication Background

Tapping multiple points across a network can result in the same packet being copied multiple times, particularly with East/West traffic and traffic from SPAN ports. This can negatively affect the performance of various tools as they not only need to process the same traffic multiple times, but in many cases further expend processing resources performing deduplication themselves. In some cases, the duplicate packets could even cause false positives or skew reporting.

Performance resilience is the ability to ensure the performance of appliances in any data center environment. Eliminating duplicate traffic is critical for improving network performance. It is of key importance for performance monitoring, cybersecurity, or forensics appliances. Duplicates can cause a lot of issues. The obvious issue is that double the amount of data requires double the amount of processing power, memory, power, etc.

However, the main issue is false positives: errors that are not really errors or threats that are not really threats. One common way that duplicates affect analysis is by an increase in TCP out-of-order or retransmission warnings. Debugging these issues takes a lot of time, usually time that an overworked, understaffed network operations or security team does not have. In addition, any analysis performed on the basis of this information is probably not reliable, so this only exacerbates the issue.

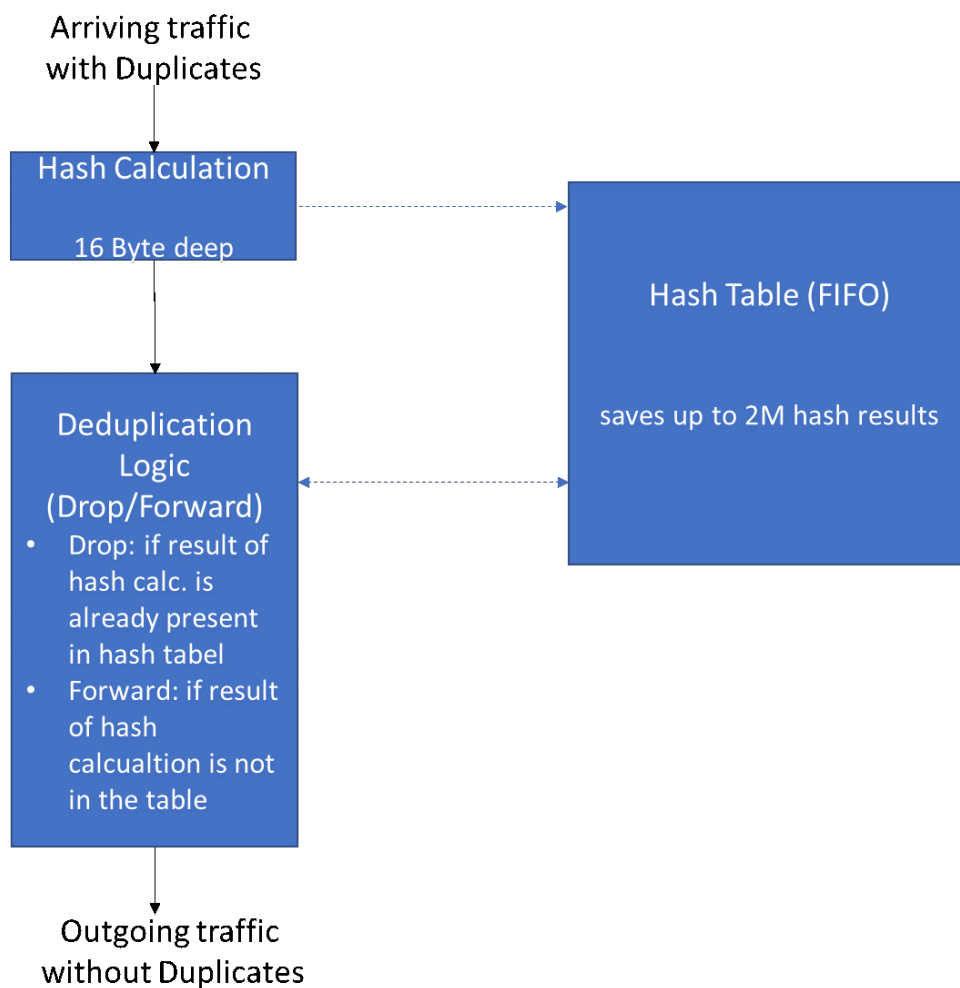
The Cubro Omnia and EXA48200 family of Advanced Network Packet Brokers offer a deduplication function to eliminate duplicated packets and thus to protect monitoring equipment from getting overloaded.



Deduplication on Cubro Omnia & EXA48200 family

Cubro provides a number of options that perform deduplication, according to user-selectable parameters, prior to forwarding traffic to monitoring tools. Offloading this resource intensive task frees a tool's processing resources allowing it to perform the tasks more efficiently they were designed to do.

When the deduplication function is switched on, Omnia/EXA48200 will calculate a checksum (=hash- key) over the full packet. Full packet means it uses every bit including payload for the hash-key calculation. The hash-key is a MD5 calculation, and the result is 16 byte long. Hash-keys are stored and every incoming packet (= one hash-key per packet) is checked against the memory. Once it is found that a hash-key is already in the memory the packet that is coming having the same hash-key is getting dropped.





Hash-keys are aged out after a certain time. This time is called detection window. Detection window of the Omnia/EXA48200 is 1000ms. Thus, when a duplicate packet arrives within 1000ms the Cubro Advanced NPB will find it and drop it. If the duplicate arrives later than the 1000ms the duplicate is not found.

Besides detection window, the size of the hash table is also important. Omnia & EXA48200 family support a hash-key storage of 4M values. When the hash table is getting full it is aging out (FIFO principle). The size of the hash table also decides the performance limits.

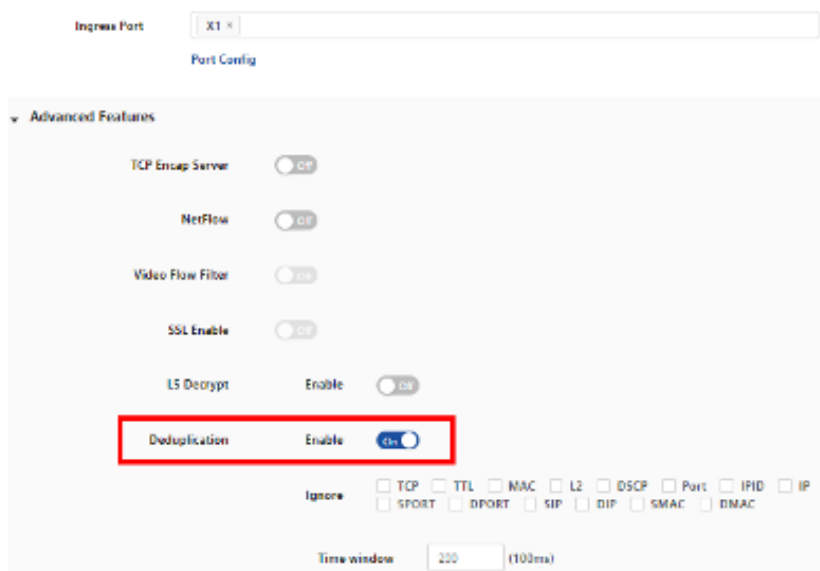
Packet Size in Byte	Traffic Utilization in Gbit/s	Packets per second
64	10	14 880 952
300	10	3 906 250
500	10	2 403 846
1000	10	1 225 490
1518	10	763 125

Small packets like 64 bytes make the hash table full much faster compared to the packet size of 1518 byte. Typical performance limit for deduplication is about 70Gbit/s for EXA48200 and 140Gbit/s for Omnia120 having in mind usual traffic patterns and packet size distributions.

Deduplication Operation

Setup of Deduplication is straight-forward and easy. In case the traffic arrives at 10G port 1 (X1) and should be sent out at 10G port 2 (X2) it requires following steps:

- a) Create a Forwarding Policy via drag and drop from X1 to X2
- b) Enable Deduplication on X1



Another important characteristic of the Omnia & EXA48200 is that it allows to blank-out certain fields in the packet.

By default Omnia & EXA48200 will compare full packets and will drop a packet if it looks 100% identical only.

Another important characteristic of the Omnia & EXA48200 family is that it allows to blank-out certain fields in the packet. By default, Omnia & EXA48200 will compare full packets and will drop a packet if it looks 100% identical only. In certain circumstances, it might happen that a packet is still a duplicate in terms of payload, but header might look different. This situation is very common when dealing with mirror ports when e.g. ingress mirroring has a single VLAN while egress mirroring has two VLAN tags.


To cope with this and still find the duplicates, the Omnia & EXA48200 allow to blank-out following fields from the hash calculation:

- TCP ... ignores Seq/Ack in the TCP header
- TTL ... Time to Live
- MAC ... ignores MAC Source and MAC Destination
- L2 ... ignores all Layer 2 including MAC and VLANs
- DSCP ... ignores Layer Quality of Service
- Port ... ignores input port
- IP ... ignore full IP header
- SPORT ... ignore L4 Source Port
- DPORT ... ignore L4 Destination Port
- SIP ... ignore Source IP
- DIP ... ignore Destination IP
- SMAC ... ignore Source MAC
- DMAC ... ignore Destination MAC

Ignore TCP TTL MAC L2 DSCP Port IPID IP
 SPORT DPORT SIP DIP SMAC DMAC

Deduplication Statistics

Statistics give you a good indication of how many duplicates have been found and eliminated.



Summary

The Cubro Omnia & EXA48200 family are perfect tools for deduplication to protect monitoring equipment from getting overloaded due to duplicate packets. These products are the latest generation of Advanced Network Packet Broker) based on a newly designed ASIC and ARM CPUs. The Omnia & EXA48200 products provide maximum flexibility because they combine the classic ASIC-based NPB with a high-performance CPU which enables them to run powerful applications on it.

Note: This document is applicable to Omnia Sessionmaster software package and EXA48200 Advanced NPB. For any further queries, please contact support@cubro.com.