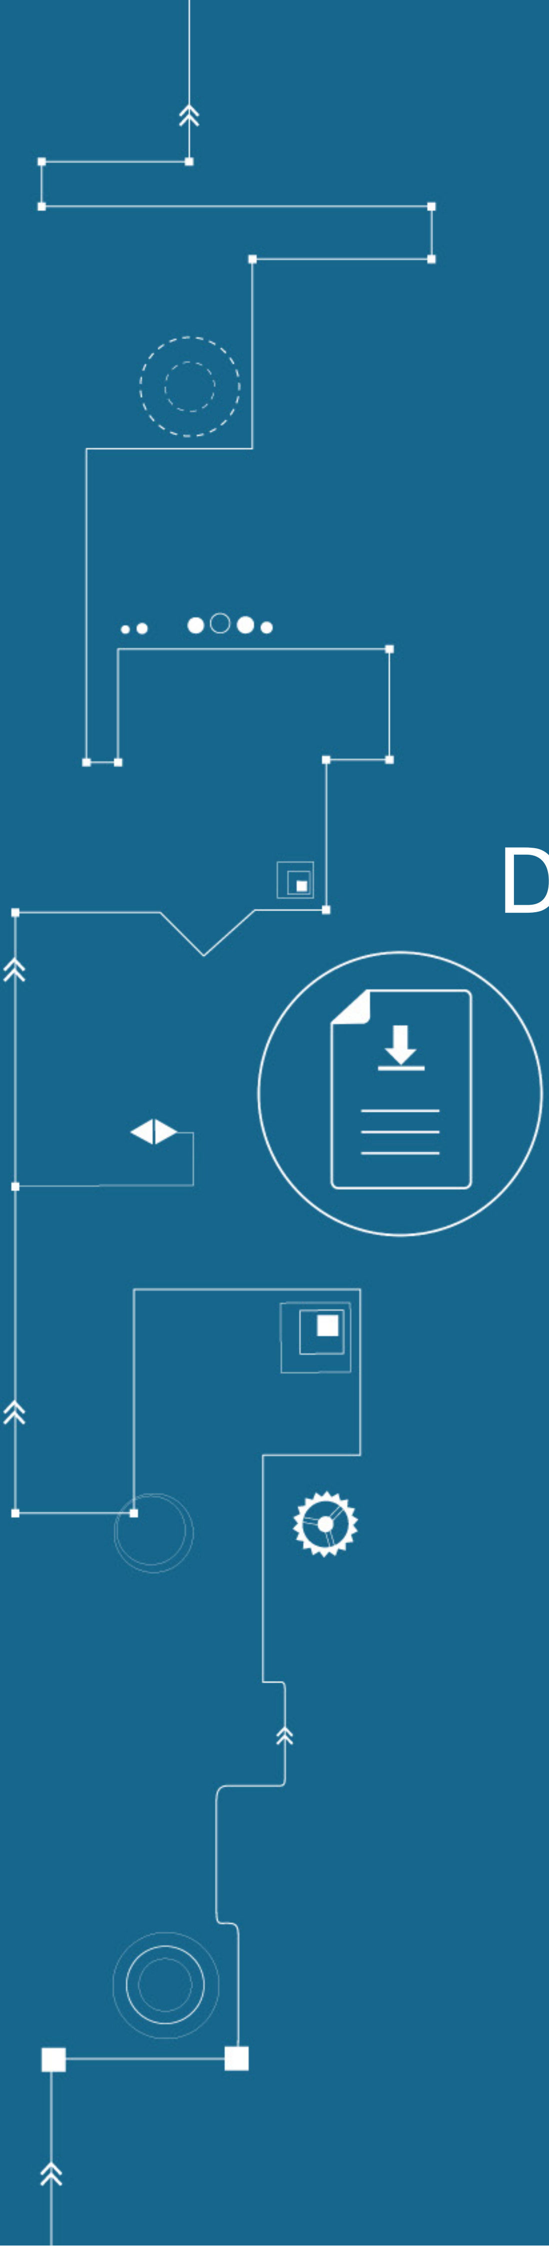

CUBRO
NETWORK VISIBILITY

EFFECT OF DUPLICATE PACKETS ON NETWORK

WHITE PAPER

JULY 2017

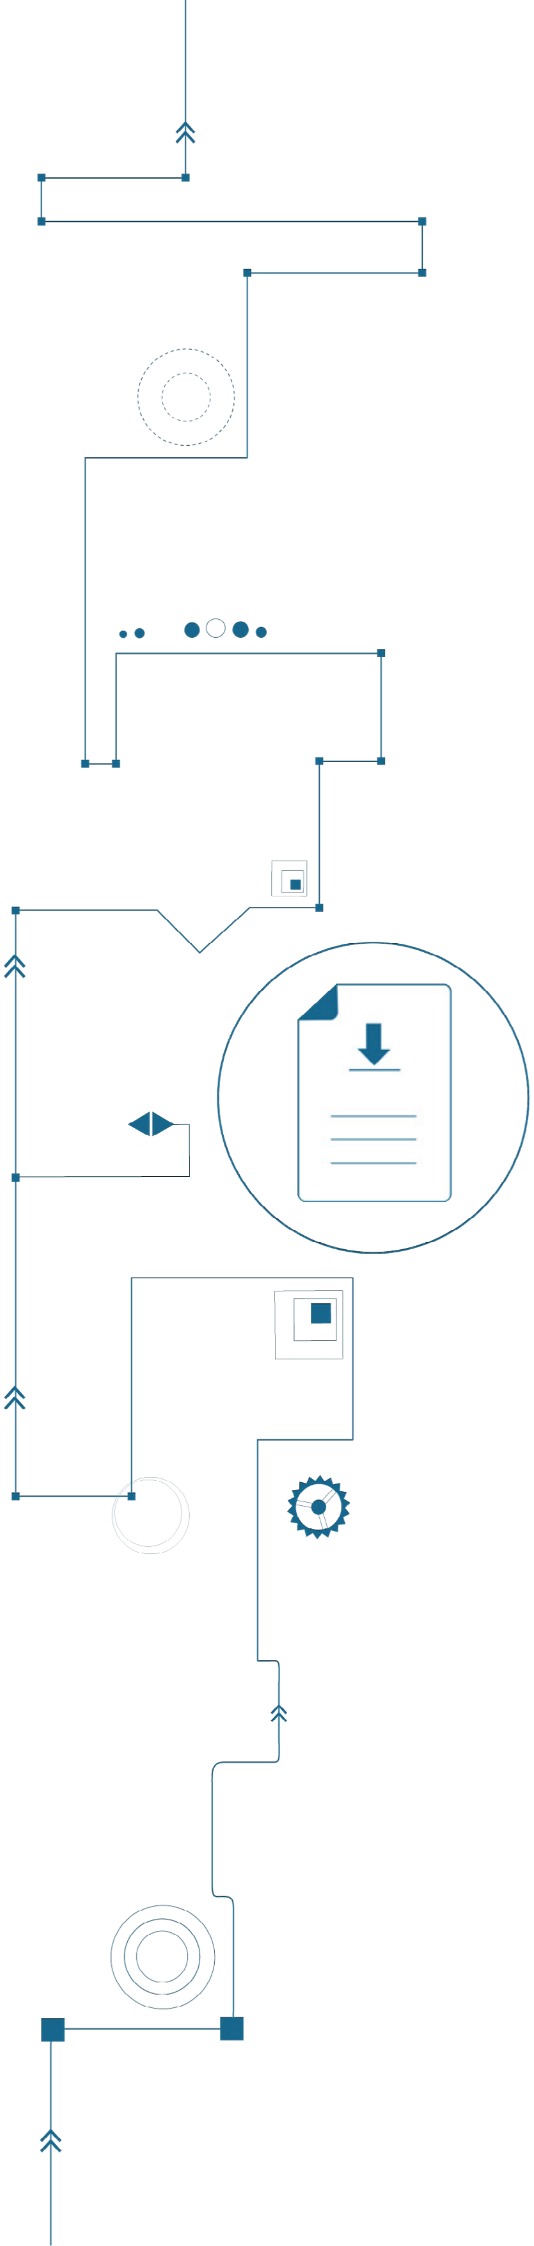

TABLE OF CONTENTS

Introduction

Packet duplication is a constant problem in all large-scale data center networks. These duplicate packets cause several problems if the network monitoring switch is not able to eliminate the copies.

Understanding Duplicate Packets

Duplicate packets lower the statistical accuracy of analysis, increase network link saturation, and can interfere with tools.

A duplicate packet is any packet that is identical to another packet. It is part of any network environment and is unavoidable. Networks that are tapped in multiple locations using SPAN or mirror ports send redundant copies of packets that can be recorded from many monitoring points across a network.

Network monitoring switches will receive multiple copies of the same packet as it traverses the network. Duplicate packets result in data being counted multiple times, leading to skewed trending statistics, such as application performance and utilization metrics. More time is also required to process and analyze the data.

Duplicate packets lower the statistical accuracy of analysis, increase network link saturation, and can interfere with tools. While some network packet brokers include advanced functions to enhance and automate network monitoring activities, the presence of these duplicate packets must be addressed in order to function at an optimal level.

Example of duplicate packet

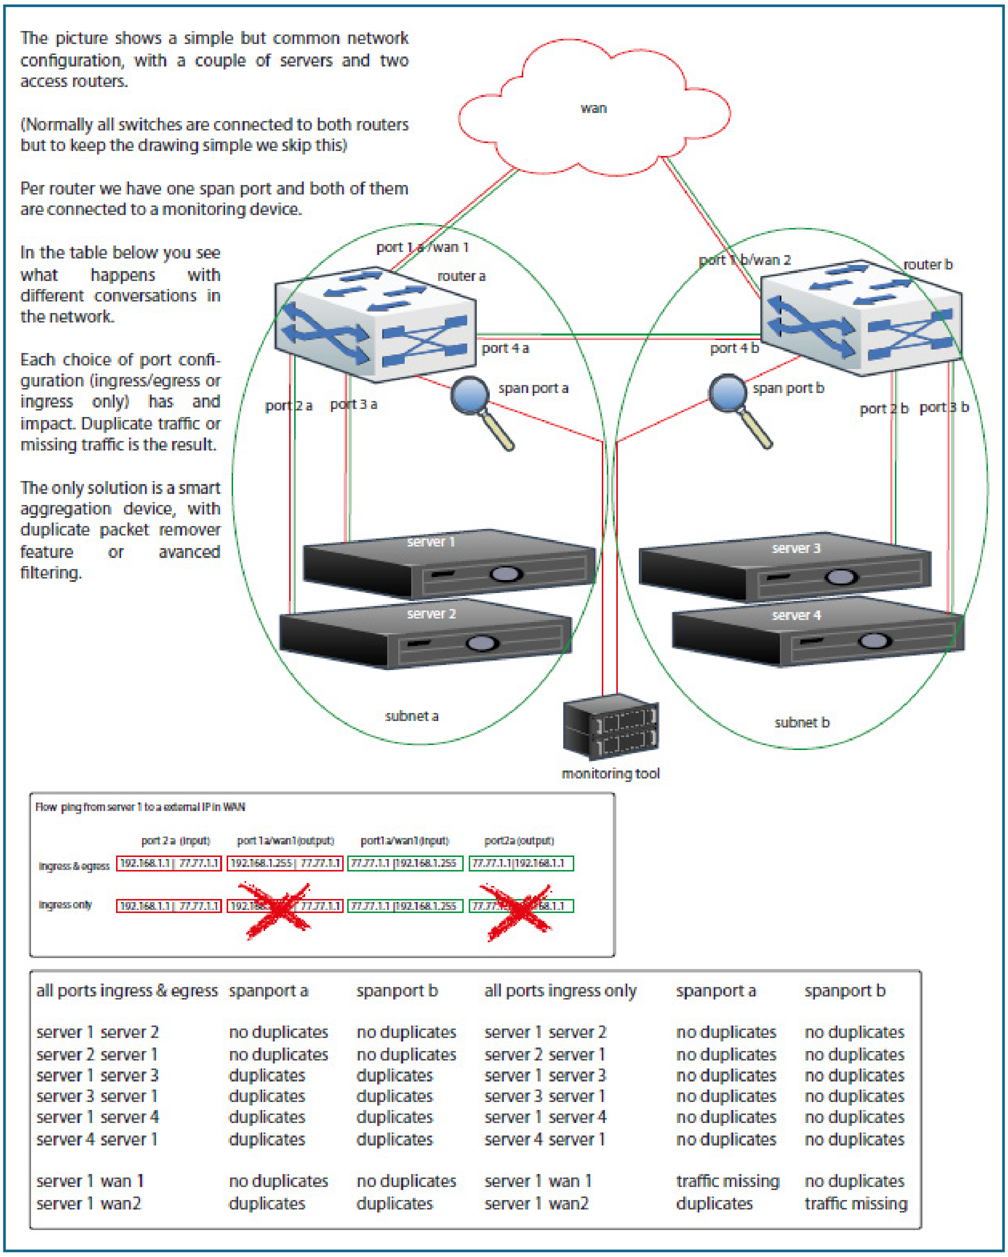

The picture shows a simple but common network configuration, with a couple of servers and two access routers.

(Normally all switches are connected to both routers but to keep the drawing simple we skip this)

Per router we have one span port and both of them are connected to a monitoring device.

In the table below you see what happens with different conversations in the network.

Each choice of port configuration (ingress/egress or ingress only) has and impact. Duplicate traffic or missing traffic is the result.

The only solution is a smart aggregation device, with duplicate packet remover feature or avanced filtering.

Flow ping from server 1 to a external IP in WAN

	port 2 a (input)	port 1a/wan1 (output)	port1a/wan1 (input)	port2a (output)
ingress & egress	192.168.1.1 | 77.77.1.1	192.168.1.255 | 77.77.1.1	77.77.1.1 |192.168.1.255	77.77.1.1|192.168.1.1
ingress only	192.168.1.1 | 77.77.1.1	192.168 [illegible] 77.77.1.1	77.77.1.1 |192.168.1.255	77.77.1 [illegible] 68.1.1

all ports ingress & egress	spanport a	spanport b	all ports ingress only	spanport a	spanport b
server 1 server 2	no duplicates	no duplicates	server 1 server 2	no duplicates	no duplicates
server 2 server 1	no duplicates	no duplicates	server 2 server 1	no duplicates	no duplicates
server 1 server 3	duplicates	duplicates	server 1 server 3	no duplicates	no duplicates
server 3 server 1	duplicates	duplicates	server 3 server 1	no duplicates	no duplicates
server 1 server 4	duplicates	duplicates	server 1 server 4	no duplicates	no duplicates
server 4 server 1	duplicates	duplicates	server 4 server 1	no duplicates	no duplicates
server 1 wan 1	no duplicates	no duplicates	server 1 wan 1	traffic missing	no duplicates
server 1 wan2	duplicates	duplicates	server 1 wan2	duplicates	traffic missing

Figure 1: The picture shows a simple but common network configuration, with a couple of servers and two access routers.

As an engineer, you want to have access to duplicate packets when necessary for troubleshooting, but you do not want those duplicate packets to skew network traffic summary statistics. Reducing duplicate packets as much as possible helps ensure your network is more efficient. It also allows your tools to be more accurate. Duplicate packets reduce statistical accuracy, which leads to higher perceived levels of traffic or network connections.

Deduplicating Packets

The solution to redundant packets in the network monitoring system is to eliminate copies before they arrive at the monitoring tools. It is important to note that identical packets are not exactly the same bit-for-bit. The packet header is inspected, and all fields must be identical for it to be a duplicate. However, there are some situations where the header has been modified slightly during the packet‘s journey. These situations require some fine-tuning of the deduplication settings to ignore those fields that were modified before the duplicate packet is received.

Packet deduplication removes duplicate packets and helps you avoid those situations. Deduplication in the context of packet broker networks (Tap Aggregation) is the ability to detect duplicates of a packet, allowing only the first packet and dropping other iterations of the same packet. If you experience duplicate packets, consider your analytical needs and network topology when deciding whether deduplication should be used. You most often encounter them when packets are traversing multiple routers and those routers are copying their traffic to the SPAN/mirror port.

Removing Unwanted Duplicate Packets

Removing duplicate traffic improves monitoring tool efficiency, accuracy, and recording space requirements. This enables monitoring tools to provide greater visibility while lowering overall costs. Removing duplicates from a saved packet capture can be more accurate than deduplication with the capture card. Observer has several more options than the capture card for ignoring packet header fields. These are header fields you choose to not examine (ignore) when determining if a packet is a duplicate. When all packet header fields are used as criteria (none are ignored) the capture card-based deduplication and Observer deduplication produce nearly the same results.

In some cases you may want to retain the duplicate packets. For example, when packets are being looped or when multiple VLANs are used with your hardware, you may want to keep the packets. Retaining a copy of duplicate packets and their traversal through both VLANs may be necessary when verifying whether the traffic was routed properly.
If you are attempting to find the source of duplicate packets in real time, do not deduplicate packets.

Conclusion

Since all the traffic that flows through is not useful, it is advisable to use a network packet broker which can remove redundant duplicate data. To save time, and processing power, duplicate packets, and other redundant data can be removed before reaching monitoring and security tools. However, it is also crucial that the network packet broker should not drop relevant original data. Cubro offers advanced NPBs with zero-loss advanced packet processing at full line rate. These packet brokers carefully sift out the duplicate data and transmit the original data to the monitoring tools