

WHITE PAPER

Duplicate Packets: Are They Good or Bad and Best Practices for Managing Them

Overview

Duplicate packets of monitoring data can come from several sources, including the use of SPAN ports and the geographic location of the data captures. For instance, a normally configured SPAN port (which is frequently used to connect monitoring tools to the network) can generate multiple copies of the same packet (see Figure 1). This is because SPAN ports on a network switch are usually configured to copy ingress and egress data from every switch data port. The SPAN port output to a network monitoring tool includes duplicate copies of every packet that goes into and then out of the switch which means multiple copies of packets are being sent to the SPAN port for every data port being used on that network switch. These copies are exact duplicates of the original packet. Even when optimally configured, a SPAN port may generate between one and four copies of a packet. The duplicate packets can represent as much as 50% of the network traffic being sent to a monitoring tool.

It also matters where you capture monitoring data. If you capture it at the ingress and then again in the core, you may have copied the same data twice, unless you selectively screened the data at the time of data capture. This double capture is in addition to whatever duplicates were made by the core switches themselves.



Even when optimally configured, a SPAN port may generate between one and four copies of a packet. The duplicate packets can represent as much as 50% of the network traffic being sent to a monitoring tool.

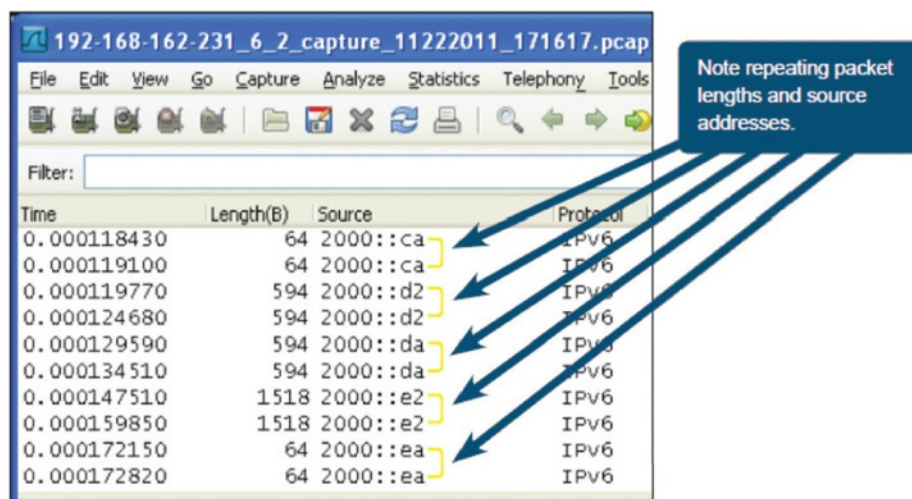


Figure 1. Examples of duplicate network traffic

All of these duplicate packets create two major problems. They prevent your tools from working at their peak efficiency and they keep you from reaching your ultimate network goals—improved network performance and better network management. IT professionals sometimes accept the duplicate packets produced by SPAN ports as a necessary byproduct of using the device, even though they clearly cause several problems that prevent effective network monitoring.

At the same time, duplicate packets can be indicators of some other issues happening on the network. So the existence of duplicate packets may not be all bad.

Duplicate Packets Can Help Isolate Visibility Architecture Design Flaws

As mentioned earlier, while duplicate packets are generally thought of as bad, they can have a positive benefit. This is because their existence can be an indicator of several network issues that should be addressed. Some examples include: malfunctioning hardware devices, a flaw in your visibility architecture that is resulting in the creation of too many copies of the same information, or a flaw or malfunction in your data filtering device (SPAN session programming, network packet broker, etc.) for network monitoring data.

In this case, a well-designed packet broker has the ability to turn de-duplication capabilities on and off. By turning de-duplication capabilities off periodically for troubleshooting purposes, an IT engineer has the ability to observe the amount of duplicate monitoring data on the network. This lets them know if they need to make any adjustments to the network or reprogram any of their data filters.

Duplicate Packets Diminish Monitoring Port Bandwidth

If a significant number of the packets sent to a monitoring tool are duplicates, then the effective bandwidth capability of the monitoring tool link is severely decreased. Let us say your monitoring tool has a 10 gigabit (10G) Ethernet interface. If one-half of the bandwidth is used to send duplicate packets, then the monitoring tool has an effective bandwidth of only 5 Gigabytes. Every duplicate packet transmitted reduces the bandwidth capability of the monitoring tools.

Compounding this problem is the fact that monitoring tools have struggled to keep pace with the increasing speed and capability of other network devices. Now add the additional burden of reduced effective bandwidth, because of a significant amount of traffic is duplicate packets. The monitoring tool simply cannot keep up with the amount of traffic arriving from the network.

Slower tools and duplicate packets reduce the effective bandwidth capability, which causes data jams that result in dropped packets and lost data. Monitoring tools are not as effective as they should be if they cannot keep up with network traffic.



Slower tools and duplicate packets reduce effective bandwidth capability, which causes data jams that result in dropped packets and lost data.

Duplicate Packets Reduce Packet Capture Storage Capability

A packet capture device records and stores the network traffic for a number of uses: regular performance analysis, forensic analysis after a network event or anomaly, security analysis, and compliance initiatives. If the network traffic sent to the capture device contains a large number of duplicate packets, then the amount of useful information these data recorders can hold is drastically reduced. Half of the data being stored may be duplicate packets that provide no real benefit for analysis or other uses. At an average packet capture cost of \$30,000 per terabyte (TB), duplicate packets waste large amounts of money.

By eliminating duplicate packets, you can store more useful data, which allows you to better analyze your data and improve your network security. Plus, with fewer duplicates taking up storage space, the captured network traffic can be stored longer, giving your IT department a larger data window for analysis and for tracking network events.

A few packet capture devices have limited filtering ability for eliminating duplicate packets. Frequently, however, implementing filtering capabilities can increase set-up time and complexity. Plus, an improper filtering set-up introduces the possibility of filtering out wanted packets or slowing the capture rate. Also, it does not address the issue of reduced bandwidth caused by duplicate packets being sent to the packet capture device in the first place.

Duplicate Packets Decrease Processing Power of Analyzer Tools

Most network analyzer tools do more than just capture and store data. They have to process the incoming data stream to analyze the data and prepare it for future analysis. Many tools have a difficult time keeping pace with regular network traffic. Just the regular 10G data stream sent to a monitoring tool can overrun the capacity of the tool, especially when a 10G monitoring device is attached to a 40G network.

This problem is exacerbated by the extra effort required for network analyzing tools to process duplicate packets. Original packets have to be queued while they wait for unnecessary duplicate packets to be processed. Once again, data bottlenecks result, which overflows the tool's buffer.

This data traffic back-up can force the monitoring tool port into premature overload. Packets may be dropped altogether, never even being received or processed by the monitoring tool. Even worse, network engineers may never be aware that there are lost packets or how much data was lost. It is difficult to perform effective network monitoring and analysis under these conditions.

Duplicate Packets Interfere with Network Analysis and Tactical Troubleshooting

Duplicate packets are equivalent to false data when it comes to analyzing network performance or troubleshooting. Network statistics can be significantly skewed when a large percentage of the traffic being analyzed is duplicate information. The problems created by duplicate packets can hinder proper analysis for performance improvement, impede the ability to pinpoint problem areas, or even worse, delay the ability to restore service if a network is down.

All IT departments want to exceed their MTTR objectives. Duplicate packets can create a major obstacle in meeting these uptime goals.

How to Eliminate Duplicate Packets

All of these issues caused by duplicate packets being sent to monitoring tools by SPAN ports lead to one overarching result—inefficient monitoring that reduces the return on your investment in monitoring tools.

Ineffective monitoring can cause problems that ripple through the data center. It leads to more man-hours spent addressing network problems and performance, more network downtime, more risks associated with losing IT capability, and less return on IT investments due to lower productivity and capability.



By eliminating duplicate packets, you can store more useful data, which allows you to better analyze your data and improve your network security.

When a SPAN port produces duplicate packets that interfere with a monitoring tool's ability to function properly, it creates a major obstacle to using monitoring tools to achieve higher network performance. Proactive IT departments are looking for ways to eliminate the negative impact of these duplicate packets.

Some monitoring tool vendors have introduced a de-duplication feature in an attempt to deal with the problem of duplicate packets interfering with their tool's functionality. The trouble with this approach is that most tools already have difficulty keeping up with network traffic. The additional burden of removing duplicate packets added to the tool's processing requirements only increases their inability to keep up. It is not an ideal solution. Plus, this approach does not address the issue of duplicate packets reducing the effective bandwidth of network traffic to the monitoring tool or the issue of wasting capture storage space. A better method is available that addresses all the issues associated with duplicate packets.

A Network Packet Broker: The Best Solution to Duplicate Packets

Network packet brokers (NPBs), which are capable of advanced functions, are the best solution to the problem of duplicate packets sent to a monitoring tool from a SPAN port. Besides resolving the issue of connectivity by allowing multiple monitoring tools of all types to connect to a single Tap or SPAN port, an NPB with advanced features is capable of removing duplicate packets at full line rate before forwarding traffic to the monitoring tools. This means network monitoring tools, from packet analyzers to packet capture devices, no longer have to deal with the problems associated with duplicate packets, like reduced bandwidth, overwhelmed tools, dropped packets, and wasted storage capacity.

The most important characteristic of de-duplication capabilities in an NPB is the size of the de-duplication window. This is the amount of time, after having seen the first copy of a packet, that a duplicate packet can be detected and removed. A large de-duplication window and the ability to configure the window size make the de-duplication feature extremely powerful.

Now you can use the tools as you intended, and get the performance and efficiency benefits from network monitoring you intended.

But the benefits of advanced network packet brokers do not end with packet de-duplication. They can also provide high-speed network optimization of data, voice, and video when you use the following features:

- Load balancing
- Centralized control of traffic



...a network packet broker with advanced features is capable of removing duplicate packets at full line rate before forwarding traffic to the monitoring tools.

- Data filtering
- Traffic prioritizing packet trimming
- MPLS stripping
- Automated IT response to monitoring alerts

With features like these, it is easy to see how integrating an NPB technology can increase the performance of your network monitoring tools.

Summary

A network packet broker allows IT teams to simultaneously connect a wide array of monitoring tools to the network. Best-in-class packet brokers do more than solve connection problems; they also aggregate and filter data according to the type of monitoring tool. Some NPBs include advanced functions to enhance and automate network monitoring activities, functions like extended burst protection and packet de-duplication at line rate. When duplicate packets get transmitted to monitoring tools, they hinder effective network monitoring. Eliminating duplicate packets is a critical step to achieve monitoring goals and improve network performance.

Duplicate packets commonly created by SPAN ports create problems for the network monitoring tools connected to the port. The duplicate packets reduce effective tool bandwidth, waste tool processing power, and consume tool storage capacity, reducing their effectiveness. Some solutions to the duplicate packet problem, like duplicate packet removal at the monitoring tool, have been less than ideal. Monitoring tools are already over-taxed in handling and processing network traffic, and they usually do not have the processing resources to handle an additional task that is processing intensive.

Eliminating duplicate packets with a network monitoring switch can reduce network monitoring tool processing load by 50% (see Figure 2). Performing de-duplication with an NPB means duplicate packets no longer reduce the effective bandwidth of traffic flow to the network monitoring tool. This method also eliminates data clutter. Monitoring tools can process and store twice as much useful data without the unneeded duplicates being transmitted to them.



Eliminating duplicate packets with a network monitoring switch can reduce network monitoring tool processing load by 50%.

Part Name	Total Tx Rate (Fps)	Total Tx Rate (Fps)	Controlled traffic without packet de-duplication*
Port 1/6/1	1,488,095	0	
Port 1/6/1	0	1,488,095	

Part Name	Total Tx Rate (Fps)	Total Tx Rate (Fps)	With Vision ONE
Port 1/6/1	1,488,095	0	
Port 1/6/1	0	744.04	

Figure 2. Example of improved network monitoring tool load

If you are investing time and money in network monitoring tools to improve network performance and efficiency, do not hamper their effectiveness with streams of duplicate packets. Advanced network packet broker technology can help you gain greater value from your network monitoring tools, making them more effective, improving their performance, and increasing their uptime.

Learn more at: www.keysight.com

For more information on Keysight Technologies’ products, applications or services, please contact your local Keysight office. The complete list is available at: www.keysight.com/find/contactus

