NEOX NETWORKS

WATCH YOUR NETWORK!

www.neox-networks.com

WHITEPAPER

NETWORK TAP VS SPAN/MIRROR PORT

PROFESSIONAL EXTRACTION OF NETWORK DATA
FOR MONITORING, ANALYSIS & SECURITY TOOLS

SecurITy made in Germany

TeleTrusT Quality Seal
www.teletrust.de/itsmig

1. Introduction

Monitoring, analysis and out-of-band security tools all have one thing in common: they need a reliable data source from which to obtain network data, and they depend on this source to function. But what is the best way to feed network data to these tools?

Many believe that they can simply configure a SPAN port, also called a mirror port, on existing switches to route out the network data. This port will then output a copy of the network data passing through the switch, depending on the configuration.

Others prefer to use Network TAPs (TAP stands for Test Access Point or Test Access Port), i.e. special devices that are looped into a network line and output a copy of the network data sent over this line from the productive network.

We will take a closer look at which of these two options should be chosen and the reasons for this on the following pages.

2. How a SPAN/Mirror port works

To know the advantages and disadvantages of a SPAN port, one must first understand how it works.

The concept itself is very simple. After the user has defined a free port on a switch as a SPAN port, incoming data packets are duplicated accordingly by the switch's operating system and output as a duplicate on that SPAN port.

At the same time, the SPAN port loses its function as a switch port, as all incoming packets are discarded by the switch.

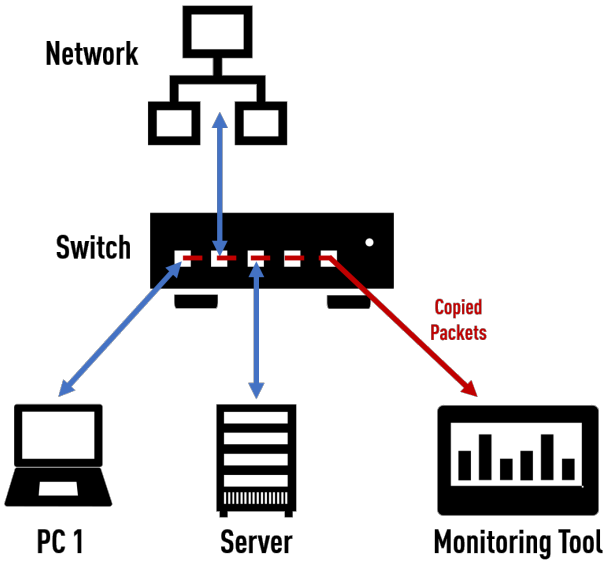

Figure 1: Functionality of a SPAN/Mirror port

3. How a Network TAP works

Unlike a SPAN port, which has to create the duplicates in software, a Network TAP works on OSI Layer 1, i.e. signal level, by means of a dedicated FPGA.

But beware! Really every TAP? Not at all! Please note that in the following we refer to a Network TAP which, like our PacketRaven models, relies on a corresponding FPGA chipset.
The term TAP is neither protected nor standardised. Some manufacturers continue to use regular switches for their Network TAPs and offer the TAP functionality by means of a fixed SPAN session. Of course, the following advantages that a TAP can bring do NOT apply to these and similar approaches!

A TAP is typically placed between 2 end devices and the connection of these two devices is looped through the network ports A and B. The TAP sits transparently in the network. Thus, the TAP sits transparently in the line to be monitored and, insofar as it is an FPGA-based TAP, can duplicate all packets exchanged by the end devices at the signal level.

This is also an important difference to the SPAN port, which, due to its mode of operation, can only duplicate valid Ethernet packets during processing, but a TAP picks up the data at the lowest OSI layer and will always and in principle make all exchanged information available to the connected system. As long as it is an Ethernet packet, it will be duplicated, no matter what size it is and how many tunnels or encapsulations are present. Even defective packets, which a switch discards without exception, are duplicated without loss while maintaining data integrity.

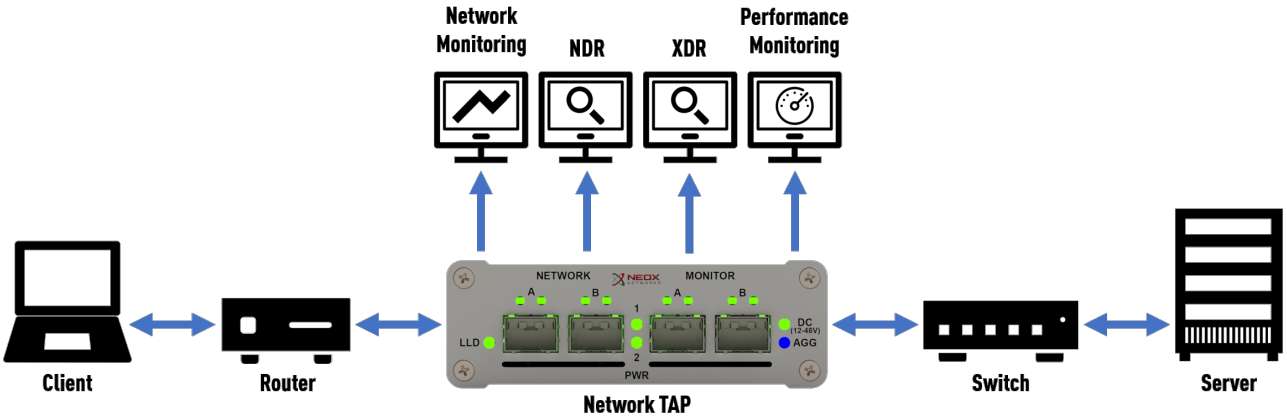

Figure 2: How a Network TAP works

4. Configuration effort and risk of misconfiguration

As already mentioned above, a SPAN port must be selected and defined manually. Of course, this first requires a free port that is to output the duplicated data from the system. Depending on the scenario, this alone can already lead to the first problems, as free ports on switches are sometimes scarce, depending on the environment.

The user must also be clear about exactly which connections he wants to monitor, as a switch can only duplicate packets on the input side; data leaving the switch can no longer be reached/duplicated by the software.

A large part of the risk of a SPAN port, however, lies in the danger of incorrect configuration. Since a SPAN port does not process incoming data and tries to discard it completely, a misconfiguration can lead to the administrator locking himself out of the system if he accidentally defines the only port with the MGMT IP as a SPAN port.

An important asset in packet analysis is the order of the packets themselves! If, for example, the data arrives in your monitoring in the wrong order, it can no longer make any evaluations. Since the switch has to „touch" each packet, figuratively speaking, there is always room for incorrect handling of the packets, which is justified by the best-effort principle, among other things, which we will discuss again below.

A wrong sequence is therefore not uncommon in heavily loaded networks. But not only the order itself can be affected, also the so-called timing of the data, ensured by the Inter Frame Gap (IFG), can change, which makes the accuracy of any analysis appear questionable.

However, even if none of the above side effects occur, time-accurate measurements are difficult to make with a SPAN port, simply because of the added latency (*Figure 3.*) to the duplicated data. While an FPGA-based Network TAP may add latency in the single-digit nanosecond range to the duplicated data, such a statement cannot be made for switches, as we are many factors higher here and can also be subject to enormous fluctuations.

Another risk of misconfiguration is the danger of duplicate packets. Here, too, it can happen very quickly that incorrectly made configurations on the SPAN session and the resulting duplicate data sets can cause an overbooking of your SPAN port or even of your connected solution itself.

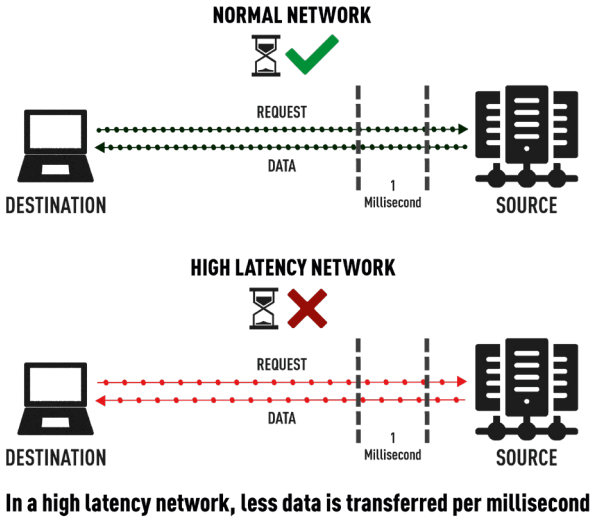

In a high latency network, less data is transferred per millisecond

Figure 3: Difference between normal and high latency. Each point in the network traffic illustrated above represents an Ethernet frame!

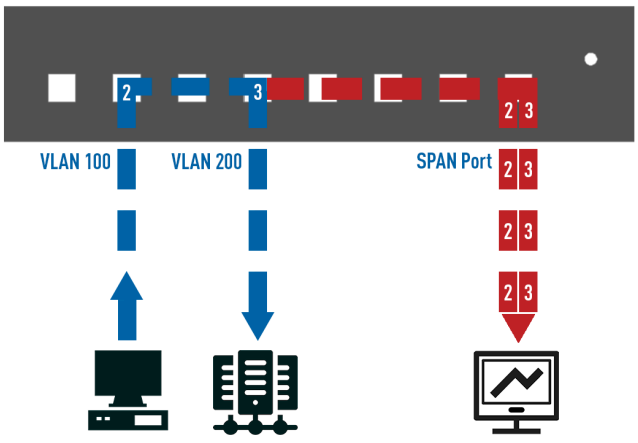

Figure 4: Problem - duplicate packages

Here is an example (*Figure 4.*) based on a VLAN-based SPAN session: If you want to have the data of VLAN 100 and VLAN 200 output on a SPAN port, the switch will output the data it receives with VLAN 100 on the SPAN port and then output the same data, but this time stamped with VLAN 200, also on the SPAN port. The result is then duplicate packets in the direction of your monitoring system.

The use of the parameter „Both" when setting up a SPAN session can also provoke such behaviour. „Both" refers to the send and receive direction of the transmitted data; in combination with a VLAN, the switch will accordingly duplicate all data that enters the VLAN and those that leaves it again; duplicate packets are also the result here, which will be visible on your monitoring system.

Up to this point, we have only talked about possible errors in the software-side setup of the switch. However, the physical component should not be disregarded. If a SPAN session has been set up for certain ports, but a colleague/employee, either unknowingly because of the SPAN session or for other reasons, simply connects the devices to be monitored to another port, you will lose the entire monitoring as long as the SPAN session has not been adjusted accordingly. Patching devices to other ports can quickly lead to the loss of important network transparency.

5. Visibility in the network

A Network TAP is, as already mentioned, always and basically 100% invisible in the network due to its FPGA-based construction! Since it operates on OSI layer 1, it also has no MAC or IP address and can neither be hacked nor compromised.

To be fair, however, it should be noted that a SPAN port is also invisible in the network, but not immune to attack. The switch itself offers great attack surfaces for intruders, who have ways and means to put a SPAN port out of operation, this is impossible with a Network TAP.

So while a TAP is always invisible and unattackable, a SPAN port is also invisible per se, but not unattackable.

6. Risk of compromise

A TAP offers no attack surface for any kind of compromise, it can neither be recognised nor hacked, and its construction does not even allow this due to the FPGA.

A SPAN port itself does not offer any attack surface, but the operating system of the switch, which is also responsible for duplicating and providing the data for the SPAN port, does!

Checking and installing regular updates on the switch are absolutely mandatory here, depending on the environment!

A Network TAP neither offers the mentioned attack surface nor does it need any updates.

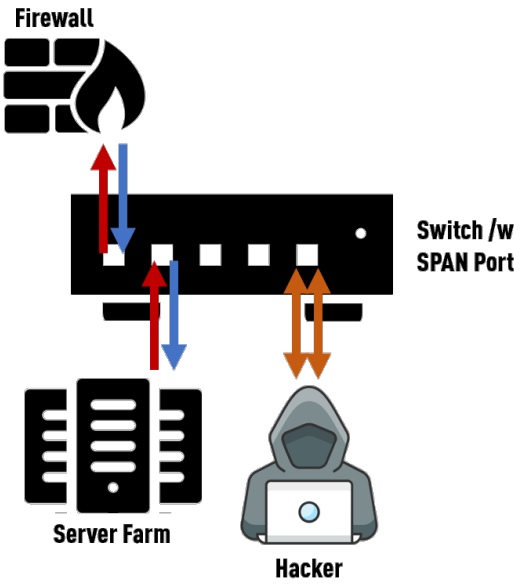

Figure 5: SPAN port compromise

7. Potential oversubscription

The danger of an oversubscription of the SPAN port is imminent. A SPAN port has several limitations in terms of throughput, one of which is of course the maximum theoretical bandwidth of the port, which is theoretically 1000 Mbps for a 1000Base-T link negotiation.

However, this value is difficult to achieve because the „source" of the data is a software component of the system, which has to do the duplication in addition to its regular task of switching.

But how exactly can a SPAN port be oversubscribed? If you look again at how exactly a SPAN port is set up, you will see that it is very easy to add too many source ports to the SPAN session and then exceed the maximum bandwidth of the SPAN port.

Here is an example: 2 servers and 4 clients are connected to a switch. If you want to read and evaluate all data exchanged between server and clients, all 6 ports are added to the SPAN session. One might now think that it is sufficient to monitor the 2 ports with the connected server.

One of the disadvantages of a SPAN session is that only incoming packets are duplicated, not outgoing packets.

If you want to do a full-duplex analysis of the traffic between clients and servers, the clients must also be included in the SPAN session and thus also all data traffic of the clients that is not intended for the servers, since a switch cannot make a distinction here.

Thus the sum of the duplicated packets can very quickly exceed the theoretical bandwidth of a SPAN port, not to mention the actual bandwidth of that port, which is limited by the switch's software.

The result is that not all data packets are forwarded to your connected solution.

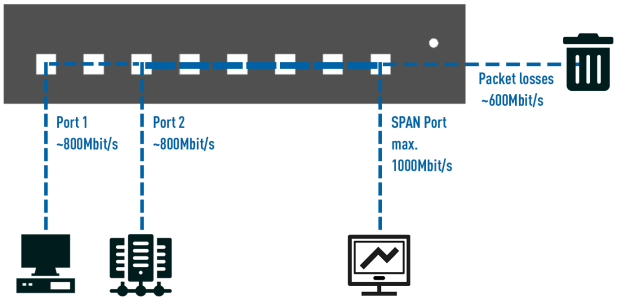

Figure 6: SPAN port oversubscription

A Network TAP can also help here! A TAP always works with the maximum theoretical bandwidth and can never be oversubscribed, even microbursts, no matter what size, are processed regularly and output as duplicates, the only limitation is the network itself!

8. Best-effort principle and high network load

This brings us directly to the next problem which, depending on the point of view, is also deliberate! The primary task of a switch is switching, for which the device was developed, produced and finally also purchased by the customer.
The possibility of setting up a SPAN session is only an additional extra offered by the manufacturer. This is accompanied by the so-called best-effort principle, which ensures that in the event of a high network load, regular operation and thus switching always has priority over other possible tasks of the switch, such as the generation of Netflow data or a SPAN session.

This characteristic, which all switches have in common, could be exploited by a potential attacker by artificially creating a high load in the network and thus forcing the switch to no longer be able to duplicate packets for the SPAN session, as it is too busy with its basic work. The best-effort principle thus allows the switch to give the SPAN session a low priority.

A Network TAP, on the other hand, has none of these restrictions. It is neither subject to a best-effort principle nor does it allow itself to be upset by a high network load; the FPGA ensures that all signals are always duplicated at all times. It is important to clarify once again that a TAP works on the signal level and not on the MAC or even IP level, it therefore has no concept for best-effort or the like, all signals are processed at full line rate without bottlenecks.

9. Completeness of the transmitted data

Here, too, we can directly follow on from the above point and talk about another important topic, namely the completeness of the data. In summary, it can be said that this is not always the case, precisely because of the best-effort principle already mentioned, under which the SPAN port operates.

Unfortunately, there are no counters or memories that buffer or at least count the unduplicated packets. This means that a SPAN port is not a reliable data source for your monitoring, let alone your security installation.

It should therefore be possible to see that a Network TAP, precisely because of the lack of any disadvantages of a SPAN port, guarantees the completeness as well as the integrity of the data at all times!

10. Integrity of the transmitted data

Especially with regard to security and perhaps even possible cyber insurance, data integrity is of incontrovertible importance!

Legal admissibility are the magic words here. Unfortunately, a SPAN port can never deliver data of this quality, because any VLANs that the switch stamps on outgoing data packets, for example, cannot be duplicated.

This means that important information is lost and the software-based duplication of the data does the rest in terms of possible manipulation of the data.

It can therefore never really be guaranteed that the duplicated data is an exact copy of the actual network traffic.

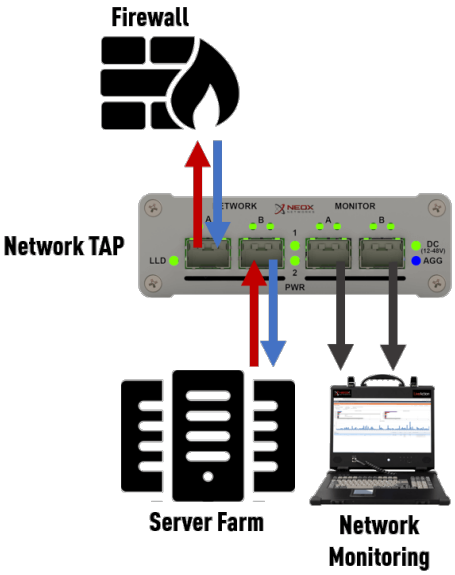

Figure 7: Portable TAP for Network Monitoring

11. Uni directional data diode design on the productive network

Here again, the Network TAP is ahead of the game: not only does a switch possibly cause higher latencies in the network due to an activated SPAN session, but also the blocking of incoming data on the SPAN port is a software-side function.

An open security hole, an unapplied patch or the like could allow an attacker to gain access to the switch and thus to your network via a compromised monitoring system.

All of this is impossible with a TAP, as a number of security features are used here, some of which are of a purely physical nature. One of these is the data diode function (*Figure 9.*), which all our electrified TAPs have in common.
Data sent back to the TAP by the monitoring system is either not passed through to the FPGA at all due to galvanic isolation.

In the case of our hybrid TAPs, the return of data via the monitoring ports to the network ports fails due to the optical-electrical-optical conversion, which again prevents a feedback into the network by means of physically unwired components.

With pure Fiber TAPs we can also implement this uni directional data diode design by means of special splitters (*Figure 8.*).

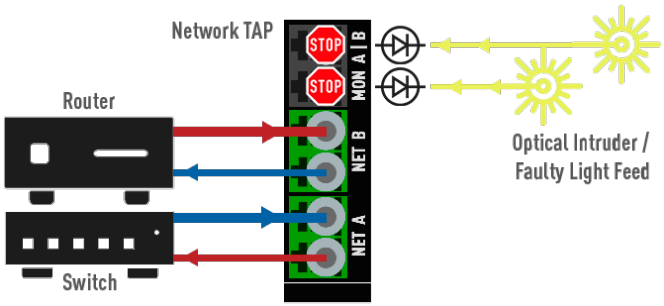

Figure 8: Modular TAP with optical filter and isolator

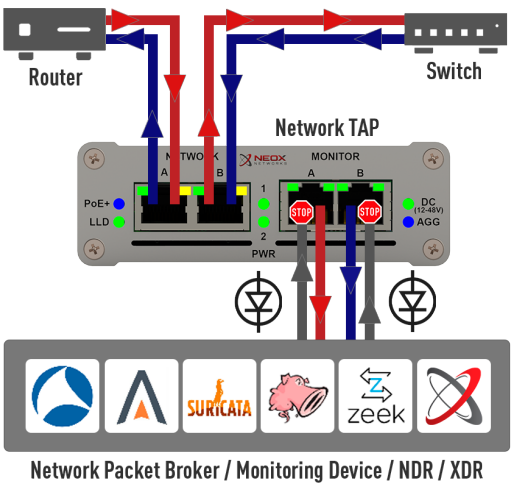

Figure 9: Portable TAP with Data Diode Function

12. Virtualisation

In virtual environments, so-called virtual SPAN/Mirror ports are used analogous to the physical network world. For this purpose we offer virtual TAPs (vTAPs). However, the advantages of these TAPs differ from their physical relatives and will soon be explained in more detail in another white paper.

13. Conclusion

It remains to be said: Tapping network traffic via a Network TAP has many advantages over routing it out via a SPAN port. Network TAPs are the only reliable data source that guarantee you 100% transparency, visibility and security in the network!

We will be happy to advise you on our various PacketRaven Network TAPs and their features, as well as on your individual requirements. You can find our contact details below!

Figure 10: Modular Fiber Network TAPs and chassis

Figure 11: Portable and rack-mountable Network TAPs