



# PhoeniX G2 ATSC 3.0 Traffic Shaping on External Devices

# CONTENTS

## Table of contents

1.	Ethernet traffic burst management .....	4
1.1.	Problem description.....	4
1.2.	Why not IEEE 802.3x Flow Control .....	5
1.3.	Traffic shaping mechanism .....	5
1.4.	QoS .....	6
2.	Traffic shaping and QoS configuration guidelines .....	7
3.1	Traffic shaping implementation guideline on a Cisco device .....	7
3.	Test description and results .....	8
3.1.	Test results with standard FW and without traffic shaper .....	8
3.2.	Test results with standard FW and with traffic shaper on Cisco CBS350.....	9
3.3.	Test results with FW with increased buffer size and without traffic shaping on CISCO switch .....	9
3.4.	Test results with FW with increased buffer size and with traffic shaper on Cisco CBS350 .....	10
4.	Conclusion .....	10

# DISCLAIMER

The following document describes the effects of the bursty ATSC 3.0 traffic on the Phoenix G2 point-to-point microwave radio link with limited Ethernet buffer performance using Cisco equipment.

The document contains a description of the problem and ways to reduce the effects of the problem by using Traffic Shaping and QoS; Shaper and QoS configuration implementation on Cisco CBS350 switch; and test results of the proposed problem prevention scenarios.

SAF Tehnika JSC has no liability for typing errors in this document or damages of any kind that result from the use of this document.

To get up-to-date information about features explained in this document, please contact SAF representative.

# 1. Ethernet traffic burst management

## 1.1. Problem description

A microwave link might be a bottleneck in the network. An Ethernet switch within a microwave radio (e.g., Phoenix G2) implies limited buffer memory. For instance, the Phoenix G2's buffer size for LAN1 port is 94kB, which provides a relatively low data transfer latency. Therefore, during the busiest traffic hours and possible modulation downshifts due to the ACM function, when link capacity is decreased, bursts of traffic may congest the Ethernet switch in the radio, causing packet loss.

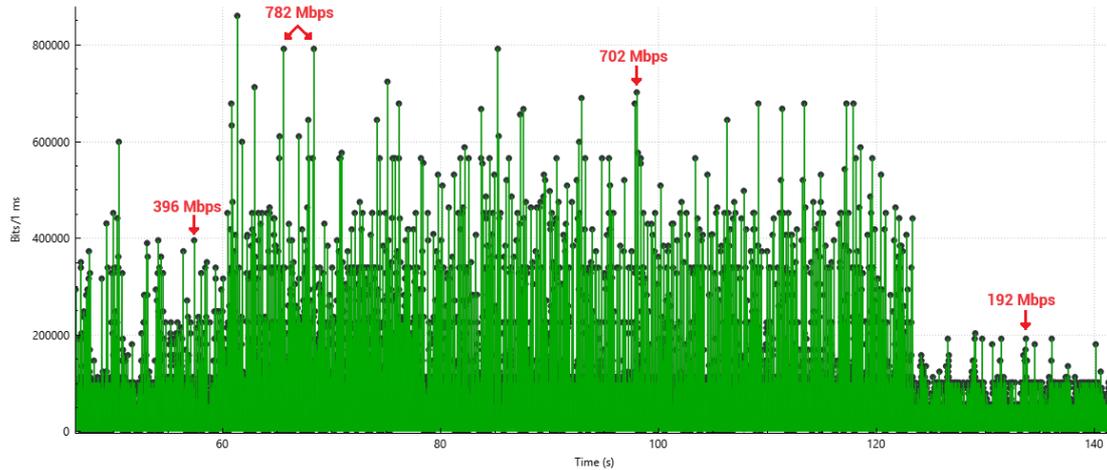


Figure 1 – 25 Mbps ATSC 3.0 traffic capture example

Bursts of almost 1Gbps (Figure 1) coming into the microwave radio can be transmitted through the link only at the speed that the link can perform, which generally is significantly less than 1Gbps. Since ingress frames are coming faster than they can be transmitted, excessive frames are stored in an Ethernet buffer until they are transmitted. In case of a large ingress burst, received frames congest the buffer, and the following incoming frames are discarded until the frames on WAN are successfully transmitted and some buffer space is freed. It is possible to observe discarded frames in the "In\_Discards" field on the LAN port on the "Counters" page in Phoenix G2's Web GUI.

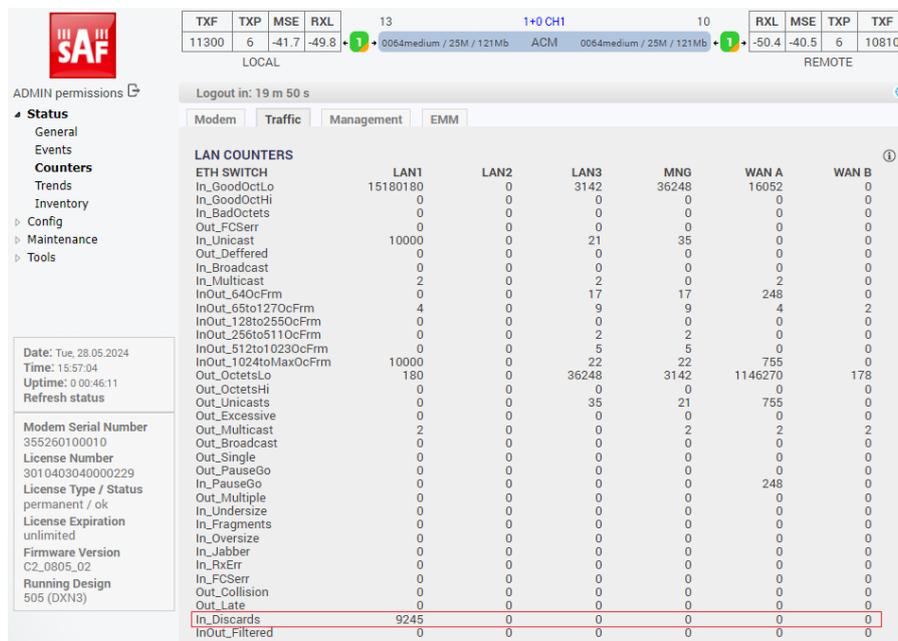
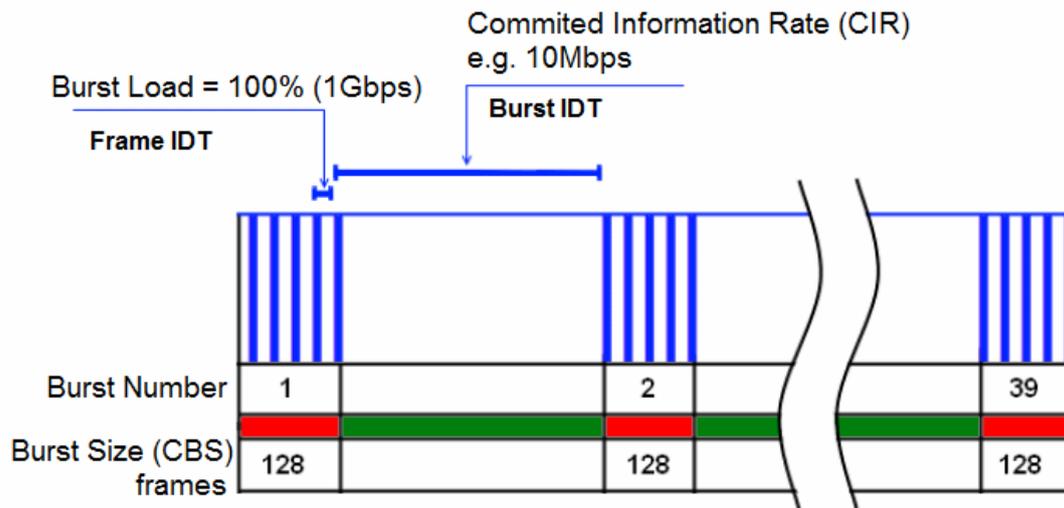


Figure 2 – Traffic counters page in Web GUI

To avoid discarding on a radio with a small Ethernet buffer, the ingress burst size (Committed Burst Size - CBS) should not exceed the buffer capacity, and the average traffic rate (Committed Information Rate - CIR) should not exceed the microwave link capacity. Since a

single Ethernet buffer is shared for a full duplex link, it is recommended not to exceed the burst size of more than half of the buffer size.



**For Phoenix G2 burst length MUST NOT exceed 47 kB**

For instance, bursts:  
 968 frames per 64 bytes  
 40 frames per 1518 bytes

\*IDT - inter-departure time

Figure 3 – Traffic burst description

## 1.2. Why not IEEE 802.3x Flow Control

IEEE 802.3x flow control allows the nodes to control traffic rates by pausing the operation of the link at the other end in case of congestion in the Ethernet buffer. A pause frame is being sent from one port that is experiencing congestion to another to stop sending traffic until the condition clears. After receiving a pause frame, the device stops sending any data packets, preventing their loss during the congestion period.

Problems arise when several flows arrive over two or more links whose total exceeds an output link's capacity, or when a flow enters a switch at a higher rate than the one it goes out. Ultimately, this will lead to the exhaust of any amount of buffering in the Ethernet switch. However, **blocking the sending port will also delay all flows over that link, even those that do not cause any congestion.**

## 1.3. Traffic shaping mechanism

Traffic shaping implementation in the switch/router before the microwave radio is another method used to avoid traffic congestion in the radio. Excess packets are stored and, if necessary, discarded in the equipment with a configured traffic shaper. Traffic with an acceptable average rate per second (CIR) may be transmitted in large bursts (Figure 3) at the line rate (e.g. 1Gbps). Such bursty traffic may cause congestion and packet loss in the radio even if the radio link provides capacity equal to or greater than the average rate of the incoming traffic. **To avoid packet loss in the radio due to oversized bursts, it is necessary to limit the burst length on the output of the external network device (e.g. router or switch).** The size of the burst is determined by the Token Bucket algorithm.

For example, a burst comes into the Phoenix G2 from a LAN 1Gbps interface. The radio link data rate is ~200Mbps FDX (Figure 4). In this example, transmitting the same frame through the WAN interface takes approximately five times longer than receiving from the LAN interface. Since frames are received faster than they are transmitted, a bottleneck occurs and the buffer in the Phoenix G2 gets filled in a relatively short time. Once the buffer is full, the following received frames will be discarded until there is enough buffer space available for at least one frame. When a frame is transmitted through the WAN, the buffer space for one frame is released, and one ingress frame can be placed inside the buffer instead of being discarded. Figure 4 shows the incoming burst, which exceeds the 94kB buffer capacity – frames marked with green are the ones that successfully pass through the radio link, while the red frames show discarded frames due to buffer overflow.

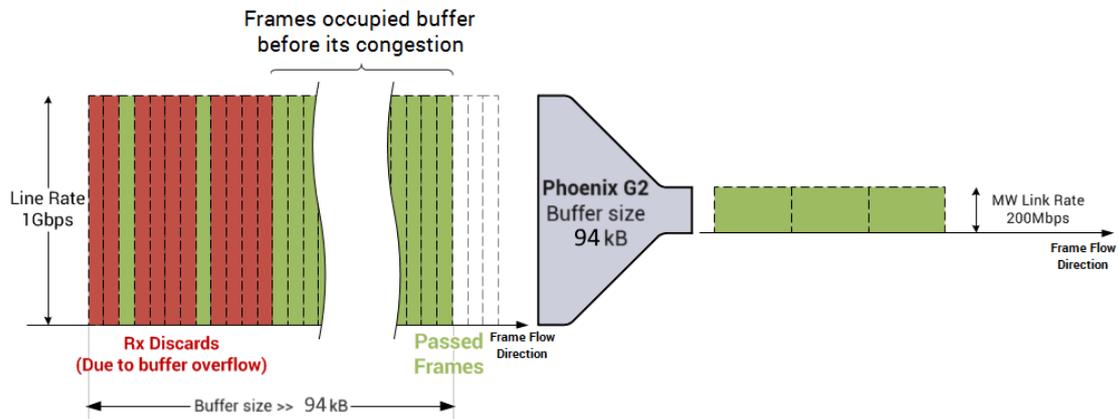


Figure 4 – Phoenix G2 effect on the oversized burst

Figure 5 shows the ingress bursts whose size is greater than the buffer capacity. In the Fig. 5(a) 190kB bursts come at a line rate of 1Gbps to the microwave radio unit with a buffer size of 94kB. Despite the average rate is 200Mbps, the part of the frames at the end of the burst are discarded (marked red) because the oversized burst quickly congests the buffer. When the burst ends, the buffer becomes freed up and it is congested again by the next incoming burst. In the example in Fig. 5(b), Traffic Shaping is enabled, and the token bucket (egress burst size) is configured to a value lower than half of the buffer size. As a result, traffic is injected into the network in smaller 50kB bursts, but with the same average data rate of 200Mbps. 50kB bursts do not congest the Ethernet buffer and there are no discards due to buffer congestion.

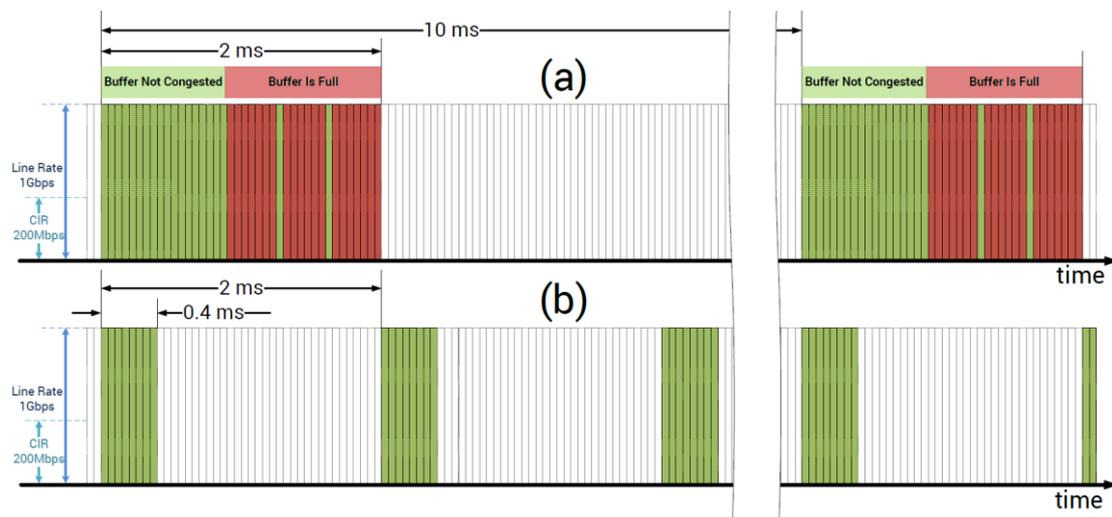


Figure 5 – Burst without shaping (a) and after shaping (b)

#### 1.4. QoS

To ensure good performance for real-time applications like broadcasting, streaming, and VoIP in a congested network, such packets must be prioritized for transmission over bulk traffic, and packet drop must be minimized. At the same time, traffic should be shaped for the value, which is below link capacity.

To prioritize certain traffic, classification must be made first. Traffic may be classified according to QoS and CoS marks like the DSCP field in an IP header and the 802.1p field in an Ethernet MAC header. Traffic may also be classified by means of VLAN tagging, source and destination addresses, or MPLS EXP field values.

According to the pre-formed classification, packets are placed into different queues. The queuing and scheduling mechanism ensures the appropriate transmission priority to the different classes of traffic. Packets from the queue with the highest priority are transmitted in the first instance and therefore have the lowest delay time and discarding rate over packets from other queues. Increased bandwidth requirements from the low-priority traffic will not affect the transmission quality of the high-priority real-time traffic. This will work in the opposite way – in case high-priority traffic requires more bandwidth, it will be guaranteed no matter what, and limited only by the physical channel capacity.

Consider an example with four incoming streams, each marked with its own DSCP tag (e.g., 1, 2, 3, 4). The classifier distributes incoming packets to the different priority queues. Frames from FIFO queue Q1 are transmitted with the highest priority, while frames from FIFO queue Q4 are transmitted with the lowest priority. In case Q1 requires more bandwidth, extra bandwidth is ensured at the expense of queues with lower priority. Therefore, if the total bandwidth requirement from four queues exceeds the overall transmission channel capacity, frames placed in the queue with the lowest priority will be dropped first.

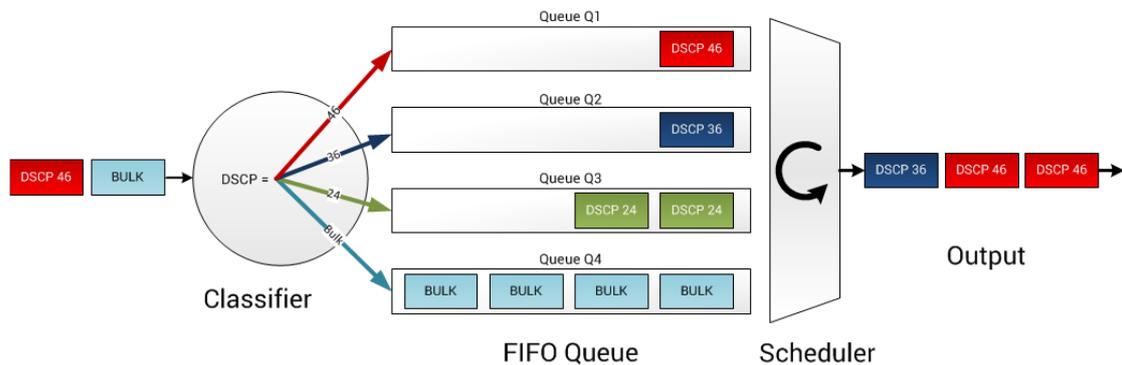


Figure 6 – QoS mechanism example

## 2. Traffic shaping and QoS configuration guidelines

All examples below are for informational purposes only and are not intended to be implemented on a network without further testing. Please refer to your switch/router manual for more information about the traffic shaper and QoS configuration.

If more than one traffic channel is used to transfer data on Phoenix G2, allocating the speed limits approximately equally between all traffic channels on Phoenix G2 is recommended for better link performance.

In the following examples, all egress traffic to the Phoenix-G2 is shaped to CIR=65MB Mbps and CBR=60kB. The lower priority traffic is assigned to queue 1, and the highest priority is assigned to queue 7.

### 3.1 Traffic shaping implementation guideline on a Cisco device

Assign traffic from port 2 to queue 7 and traffic from port 4 to queue 1, queue 1 has the lowest priority, queue 7 has the highest priority:

```
CBS350# configure terminal
CBS350(config)# interface GigabitEthernet2
CBS350(config-if)# qos cos 7
CBS350(config-if)# exit
CBS350(config)# interface GigabitEthernet4
CBS350(config-if)# qos cos 1
CBS350(config-if)# exit
```

Configure egress traffic shaping on GigabitEthernet3 to limit the CIR to 65000 kbit/s and CBS to 60000 bytes (no drops of the Phoenix G2 interface). Configure egress traffic shaping for queues on GigabitEthernet3 to limit the CIR and CBS of the first queue to 22000 kbit/s and 15000 bytes, and the seventh queue to 43000 kbit/s and 45000 bytes:

```
CBS350(config)# interface GigabitEthernet3
CBS350(config-if)# traffic-shape 65000 700000
CBS350(config-if)# traffic-shape queue 1 22000 15000
CBS350(config-if)# traffic-shape queue 7 43000 680000
CBS350(config-if)# exit
CBS350(config)# end
CBS350# write memory
```

### 3. Test description and results

PhoeniX G2 link capacity is set to 70 Mbps, with no shaping on the PhoeniX G2. Two bursty streams with a 1 Gbps continuous burst load are generated at the Ethernet tester output. The 1<sup>st</sup> test stream has an L2 bitrate of 40 Mbps, the burst size is changed to observe the queuing impact on packet drops. The 2<sup>nd</sup> stream has an L2 bitrate of 20 Mbps, and the burst length is 80. Frame size is 1518 bytes for both streams.

For the standard FW, tests were performed with a burst length equal to:

- 1) Half the size of the PhoeniX G2 buffer\*
- 2) Full size of the PhoeniX G2 buffer\*
- 3) Double the size of the PhoeniX G2 buffer\*
- 4) About 10x of the PhoeniX G2 buffer\*

\*Size of PhoeniX G2 buffer for the standard firmware.

The results for high-priority traffic only are represented. We disregard any lower-priority traffic loss in this test.

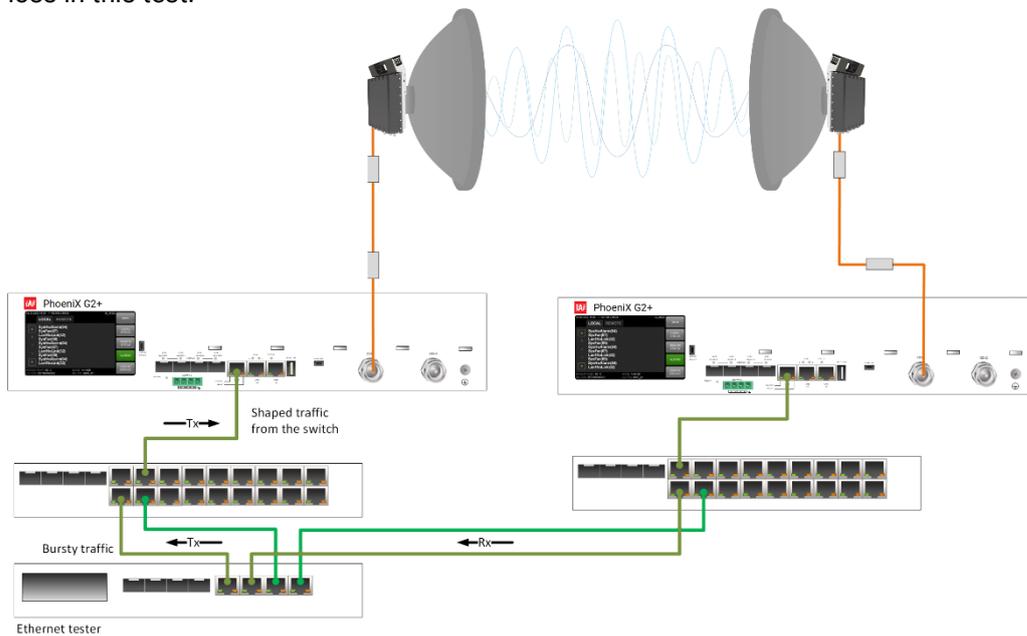


Figure 7 – Test setup

#### 3.1. Test results with standard FW and without traffic shaper

Test traffic is sent through the setup with the PhoeniX G2 link with standard FW, no processing on the CISCO switches.

Results of 10-minute tests without traffic shaper and standard FW for Q7 traffic:

Test No.	Burst Length, frames	Tx Packets	Rx Packets	Discards on PhoeniX G2	Packet Loss Rate
1	30	1961330	1961330	0	0.00%
2	60	1961330	1732506	228824	11.67%
3	120	1961330	964323	997007	50.83%
4	580	1961330	305019	1656311	84.45%

• Tx Packets – the number of packets sent from the Ethernet tester;

- Burst Length – the number of bursty frames included in Tx packets;
- Rx Packets – the number of packets received by the Ethernet tester;
- Discards on Phoenix G2 radio – the number of packets discarded on Phoenix G2;
- Packet Loss Rate – the percentage of packet loss from Ethernet tester Tx port to Rx port.

Summary of the test without traffic shaper:

Here you can see that the number of discarded packets increases with the number of bursty packets. This will result in the loss of all customer data transmitted over the radio link.

### 3.2. Test results with standard FW and with traffic shaper on Cisco CBS350

The traffic shaper on the CISCO switch is configured for a CIR of 65 Mbps and CBS of 60000 bytes on port 3. Egress traffic shaping for queues on the interface to Phoenix-G2 is limited to CIR and CBS for the first queue (less important) to 22000 kbit/s and 15000 bytes, and the seventh queue (more important) to 43000 kbit/s and 45000 bytes:

Results of 10-minute tests using traffic shaper for Q7 traffic:

Test No.	Burst Length, frames	Tx Packets	Rx Packets	Discards on Phoenix G2	Packet Loss Rate
1	30	1961330	1961330	0	0.00%
2	60	1961330	1961330	0	0.00%
3	120	1961330	1961330	0	0.00%
4	580	1961330	1473155	488175	24.89%

Summary of the test using traffic shaper on Cisco CBS350:

In this test, the Phoenix G2 switch does not discard bursty packets from the tester that exceed its buffer size. Bursty traffic is discarded only on egress from Cisco CBS350 when the Cisco CBS350 buffer is full. In this configuration, the Phoenix G2 link can be excluded from packet loss if the radio part is working normally.

### 3.3. Test results with FW with increased buffer size and without traffic shaping on CISCO switch

Test traffic is sent through the setup with the Phoenix G2 link with FW with increased buffer size, no processing on the switches.

Results of 10-minute tests with FW with increased buffer size for Q7 traffic:

Test No.	Burst Length, frames	Tx Packets	Rx Packets	Discards on Phoenix G2	Packet Loss Rate
1	30	1961330	1961330	0	0.00%
2	60	1961330	1961330	0	0.00%
3	120	1961330	1961330	0	0.00%
4	580	1961330	1382915	578415	29.49%

- Burst Length – the number of bursty frames included in Tx packets;
- Tx Packets – the number of packets sent from the Ethernet tester;
- Rx Packets – the number of packets received by the Ethernet tester;
- Discards on Phoenix G2 radio – the number of packets discarded on Phoenix G2;
- Packet Loss Rate – the percentage of packet loss from Ethernet tester Tx port to Rx port.

Summary of the test without traffic shaper and FW with increased buffer size:

Here, you can see that there is an improvement compared with the results with standard FW without shaping and even with the used switch with traffic shaping. The danger of losing customer data transmitted over the radio link is significantly lowered compared to the standard FW, and also by using the switch buffer.

### 3.4. Test results with FW with increased buffer size and with traffic shaper on Cisco CBS350

Test traffic is sent through the setup with the Phoenix G2 link with FW with increased buffer size, and shaping is used on the switch.

Results of 10-minute tests without traffic shaper and FW with increased buffer size for Q7 traffic:

Test No.	Burst Length, frames	Tx Packets	Rx Packets	Discards on Phoenix G2	Packet Loss Rate
1	30	1961330	1961330	0	0.00%
2	60	1961330	1961330	0	0.00%
3	120	1961330	1961330	0	0.00%
4	580	1961330	1653681	307649	15.69%

- Tx Packets – the number of packets sent from the Ethernet tester;
- Burst Length – the number of bursty frames included in Tx packets;
- Rx Packets – the number of packets received by the Ethernet tester;
- Discards on Phoenix G2 radio – the number of packets discarded on Phoenix G2;
- Discards on Huawei switch – the number of packets discarded on Huawei switch;
- Packet Loss Rate – the percentage of packet loss from Ethernet tester Tx port to Rx port.

Summary of the test without traffic shaper and FW with increased buffer size:

Here, you can see that there is an improvement compared with results with standard FW without shaping, and it is better than results with FW with increased buffer size without traffic shaping. The probability of losing customer data transmitted over the radio link is significantly lowered compared to the standard FW.

## 4. Conclusion

For very bursty traffic, such as ATSC 3.0, when using Phoenix G2 radios, it is recommended to use the following measures to avoid traffic loss:

- Phoenix G2 with the firmware with increased buffer size. The buffer size is up to 700kB (exact value depends on the packet size and the interface used).
- Phoenix G2 with the standard firmware and additional external network devices:
  - with traffic shaping support,
  - with sufficient buffer size,
  - with QoS features for better traffic prioritization.