

Is flow control necessary?



In this white paper we will explain the concept of the Ethernet switch buffer and flow control.

Hypercable®

81 rue des Carrières
Parc la Ronze
69440 - Taluyers - France
Tel: + 33 (0) 478 487 475
Voip: +33 (0) 962 305 810
Fax: + 33 (0) 478 453 182
Mail: info@hypercable.fr
Web: www.hypercable.fr
Code CEE: FR 19 384 007 894

About buffer memory in Ethernet terminology:

Ethernet switches rely on packet switching. The switch establishes a connection between two segments and keeps the connection long enough to send the current packet. Incoming packets are saved in the temporary memory area – a buffer.

If the destination segment is congested, the switch will hold on to the packet as it waits for the bandwidth to become available on the crowded segment. Buffers that are full present a problem. Therefore, the analysis of buffer sizes and strategies for handling overflows is of interest for the technically inclined network designer. Crowded segments cause many problems, so their impact on switch consideration is not important for most users, since networks should be designed to eliminate crowded, congested segments.

Traffic in multi-tier Ethernet networks, especially running (bursty) TCP/IP protocol, is not uniform. There are instantaneous traffic spikes that eventually fill up the FODU Wi700 buffer, which size is 128KB. At full 1Gbps traffic, the buffer gets filled in a roughly short period of time. The flow control allows mutually binding packet buffers of FODU Wi700 and interconnected user equipment; therefore, traffic bursts will not be dropped by FODU radios as user equipment is responsible for traffic equalization. For example, if the user switch has a 1MB packet buffer and obeys incoming pause packets from FODU Wi700, virtually FODU Wi700 will share the bigger buffer and benefit from it.

From theory to tests:

It is possible to generate traffic that does not exceed 10Mbps and freely passes a particular shaper but overrun FODU Wi700 128Kbyte buffer.

To show the above mentioned, we have conducted several tests with the following configurations:

1) FODU Wi700 link capacity is 100Mbps, both links are interconnected with the Ethernet tester over wired cables at 1Gbps speed, and the Ethernet tester does not obey PAUSE packets from FODU Wi700 – the flow control is off.

2) FODU Wi700 link capacity is 100Mbps, links are interconnected with Layer 2 switches which are connected to the Ethernet tester over wired cables at 1Gbps speed:

a) flow control is disabled on L2 switches;

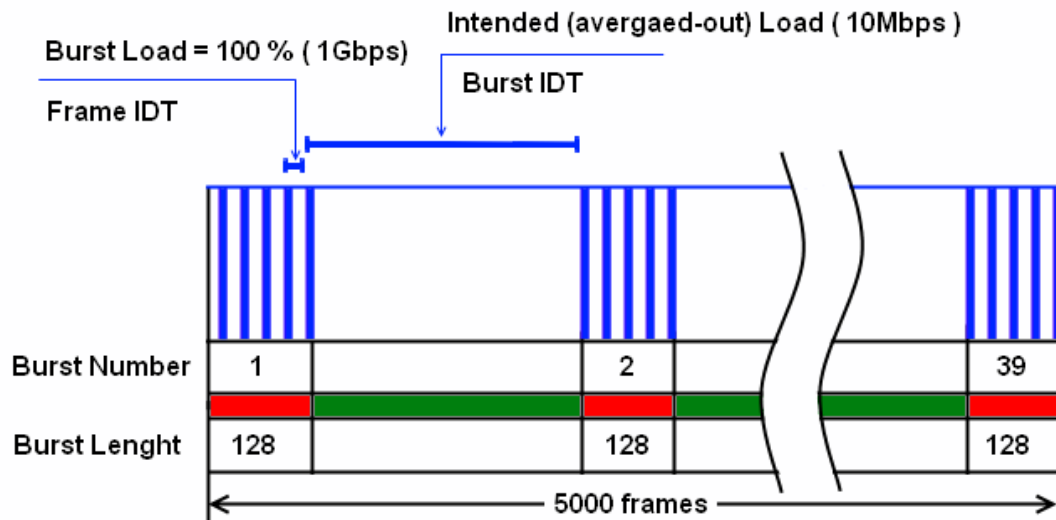
b) flow control is enabled (AUTO) on FODU Wi700 and enabled on Layer 2 switches. This means that Layer 2 switches will listen for FODU Wi700 pause packets and will transmit PAUSE packets to FODU Wi700.

FODU Wi700 Ethernet switch has a 128KB buffer size. Layer 2 switch has a 1MB buffer size. The Ethernet tester will generate traffic with 1000 bytes packet size in one direction which is 10Mbps, but with short bursts (128 frames) load at full 1Gbps speed.

Below is the example of a traffic stream from the Ethernet tester.

Is flow control necessary?


Below is the example of a traffic stream from the Ethernet tester:



Frame IDT: 12304 ns
Burst IDT: 153890390 ns


*IDT - inter-departure time

Test configuration No.1



Ethernet tester's Tx port is connected to FODU WI 700 LAN port 2 (in WEB GUI 'LAN4')

Is flow control necessary?

Test configuration No.2

Ethernet tester's Tx port is connected to L2 switch 1 to port 1 and FODU Wi700 LAN port 2 (in WEB GUI 'LAN4') is connected to L2 switch port 2

* pictures do not show the actual testing hardware.

Is flow control necessary?

Test configuration No.1

The flow control is in AUTO state on FODU Wi700 . That means if the flow control is disabled on L2 switch or in our case on the Ethernet tester, the FODU Wi700 switch will not send PAUSE packets to it.

The screenshot from Wi 700 WEB GUI below shows that the device which is connected to LAN4 is not obeying PAUSE packets from FODU Wi700 and does not transmit PAUSE packets to FODU Wi700

Ethernet configuration

Ethernet status and configuration

	LAN 1	LAN 2	LAN 3	LAN 4	WAN
Port state	Ok	Ok	Ok	Ok	Ok
Link	Off	Off	Off	1000 Mbps	1000 Mbps
Duplex (actual)	Full	Full	Full	Full	Full
Rx flow	Off	Off	Off	Off	On
Tx flow	Off	Off	Off	Off	On
Rx state	On	On	On	On	On
Tx state	On	On	On	On	On
Speed (set)	auto	auto	auto	auto	
Ethernet flowcontrol	<input type="radio"/> On <input type="radio"/> Off	<input type="radio"/> On <input type="radio"/> Off	<input type="radio"/> On <input type="radio"/> Off	<input type="radio"/> On <input type="radio"/> Off	<input type="radio"/> On <input checked="" type="radio"/> Off <input checked="" type="radio"/> Auto

Rollback on Execute configuration

Write to config file

System returned: Ok

Test No.	Burst packets	Tx packets	Rx packets	Discards on WAN	Discards on L2 switch	Packet loss
1	0	5000	5000	0	0	0%
2	64	5000	5000	0	0	0%
3	128	5000	2972	2028	0	41%
4	256	5000	1520	3480	0	70%

Transmitted and received packets

The meaning of the table:

Test No.:

- 1) no bursty traffic
 - 2) Half size of the FODU Wi700 buffer
 - 3) Full size of the FODU Wi700 buffer
 - 4) Double size of the FODU Wi700 buffer
- Tx packets – the number of packets sent from the Ethernet tester
 - Burst packets – the number of bursty packets that are included in Tx packets.
 - Rx packets – the number of packets received on the Ethernet tester
 - Discards on WAN – the number of packets discarded on FODU Wi700 1 WAN port
 - Discards on L2 switch – the number of packets discarded on L2 switch 1
 - Packet loss – number in % of packet loss from Ethernet tester Tx port to Rx port

Summary of test configuration No.1:

Here you can see that the number of discarded packets increases with the number of bursty packets. This will cause packet loss to **all customer data** which goes through the radio link.

Is flow control necessary?

Test configuration No.2 scenario A

The flow control is in AUTO state on FODU Wi700 . L2 switches are not configured to obey Wi700 PAUSE packets.

Test No.	Burst packets	Tx packets	Rx packets	Discards on WAN	Discards on L2 switch	Packet loss
1	0	5000	5000	0	0	0%
2	128	5000	2972	2028	0	41%
3	256	5000	1520	3480	0	70%
4	1152	5000	730	4270	0	85%

Transmitted and received packets

Test No.:

- 1) no bursty traffic
- 2) Full size of the FODU Wi700 buffer
- 3) Double size of the FODU Wi700 buffer
- 4) Full size of the FODU Wi700 buffer + L2 switch buffer

Summary of test configuration No.2 scenario A:

You can see that the switch with disabled flow control does not solve the problem with bursty traffic.

Test configuration No.2 scenario B

The flow control is in AUTO state on FODU Wi700 . L2 switches are configured to obey FODU Wi 700 PAUSE packets. The screenshot from Lumina WEB GUI below shows that FODU Wi 700 sends and receives PAUSE packets to/from the device which is connected to 's LAN4.

Ethernet configuration

Ethernet status and configuration

	LAN 1	LAN 2	LAN 3	LAN 4	WAN
Port state	Ok	Ok	Ok	Ok	Ok
Link	Off	Off	Off	1000 Mbps	1000 Mbps
Duplex (actual)	Full	Full	Full	Full	Full
Rx flow	Off	Off	Off	On	On
Tx flow	Off	Off	Off	On	On
Rx state	On	On	On	On	On
Tx state	On	On	On	On	On
Speed (set)	auto	auto	auto	auto	
Ethernet flowcontrol	<input type="radio"/> On <input type="radio"/> Off	<input type="radio"/> On <input type="radio"/> Off	<input type="radio"/> On <input type="radio"/> Off	<input type="radio"/> On <input type="radio"/> Off	<input type="radio"/> On <input type="radio"/> Off <input checked="" type="radio"/> Auto

Rollback on Execute configuration
Write to config file

System returned: Ok

Is flow control necessary?

Test No.	Burst packets	Tx packets	Rx packets	Discards on WAN	Discards on L2 switch	Packet loss
1	0	5000	5000	0	0	0%
2	128	5000	5000	0	0	0%
3	256	5000	5000	0	0	0%
4	1152	5000	2968	0	2032	41%

Transmitted and received packets

Test No.:

- 1) no bursty traffic
- 2) Full size of the FODU Wi700 buffer
- 3) Double size of the FODU Wi700 buffer
- 4) Full size of the FODU Wi700 buffer + L2 switch buffer

Summary of test configuration No.2 scenario B:

In this test we see that FODU Wi700 switch does not discard bursty packets that overflow the FODU Wi700 buffer size. Bursty traffic is discarded on the ingress port in L2 switch and that means that you can exclude the FODU Wi700 link from packet loss in case if the radio side is working well.

Note:

When we connect 100Mbps switches and run bandwidth tests, everything works smoothly because the traffic does not exceed the maximum link capacity and is already limited on the external switch capacity; therefore, the CFIP FODU Wi700 buffer, buffer is not filled up.

Hardware used in the tests:

Agilent N2X hardware tester
 FODU WI700 with electrical LAN ports
 HP ProCurve 3400 series switches
 CISCO Catalyst 2950 series switches

Conclusion

We recommend for all Radio Modem links:

- to eliminate crowded segments
- to use the flow control on our equipment and on the equipment it is connected to
- or use devices with traffic shaping support
- or/and if possible increase the buffer size on connected switches.

Is flow control necessary?