Performance Evaluation of the GSI timing network

Enkhbold Ochirsuren

GSI, TOS

24.1.2024

Table of Contents

1. Overview

The objective of this study is to evaluate the performance of a timing network implemented on the White Rabbit switches (WRS). Beyond promising precise timing synchronization, the WR timing network ensures the reliable distribution of control data within a timing system. Therefore, it is essential to assess the timing network's performance with respect to standard qualification factors such as throughput and latency. These factors are evaluated in typical network traffic scenarios:

- Data Master (DM) broadcast: this scenario involves broadcasting control data from a central DM node to all recipients within a network
- Bunch-2-Bucket (B2B) unicast (many-to-one): in this case, control data is unicast from multiple B2B nodes to a single DM node
- service traffic (bi-directional, many-to-one, and one-to-many): bi-directional communication between multiple sources and a single destination (and vice-versa)
- mixed traffic (1 DM, 6 B2B, 6 service): it combines all traffics of the previous scenarios.

Bandwidth is a measure of the data volume that can pass through a network at any given time. It's not a measure of speed but rather a reflection of capacity, dependent on both throughput and latency.

Throughput represents the average amount of data that actually traverses a network within a certain time period.

Network latency is the amount of time that takes for data to move from its source to a destination across a network.

2. Test setup

The testbed consists of the following components:

- chassis: XenaBay
- software: Valkyrie2889 v1.41, ValkyrieManager r88.2
- configuration: https://github.com/GSI-CS-CO/network_testing/
  - GSI_Use_Case_test/performance_analysis_2023_11/Configuration/Valkyrie2889/*.v2889
  - GSI_Use_Case_test/performance_analysis_2023_11/Configuration/ValkyrieManager/xenabay_gsi_use_case_4_layers.vmcfg

The complete test setup is constructed using four 4 layers of WRS, each configured with its corresponding layer role (localmaster, distribution, access of v1.5.7). All WRSs, model WRS-3/18, are programmed with software release v6.1. Hardware versions of v3.3 and v3.4 are mixed within the test setup.

The XenaBay chassis serves as both the traffic generator and analyzer. Using the Valkyrie2889 software, advanced network tests can be performed according to the RFC 2889 specification. The provided throughput and forwarding rate tests are performed for individual traffic types (DM, B2B, and service), and the maximum throughput is measured.

The ValkyrieManager software is employed to generate a mix of different traffic types and measure their their frame loss statistics.

Note: in case of trouble reboot the XenaBay chassis and restart the Valkyrie software. Rebooting the chassis can be done by ValkyrieManager (Available Resources:reserve chassis -> Resource Properties: reboot chassis).

For DM broadcast and B2B unicast traffics, the test frames simulate the timing messages, ranging from 1 to 16 messages per frame. The layer 2 frame length of 90 bytes corresponds to 1 timing message per frame (preamble (7 octets), start frame delimiter (SFD, 1 octet) and interpacket gap (IPG, 12 octets) are not considered).

The service traffic represents any network traffic with all possible frame lengths, ranging from 64 to 1518 bytes.

3. Test results

Only the standard test results were recorded as files:

- valkyrie2889-Report-20231201-121634.pdf: DM broadcast, 1->17
- valkyrie2889-Report-20231201-110640.pdf: B2B unicast, 6->1, TX rate=1-100%
- valkyrie2889-Report-20231207-184526.pdf: B2B unicast, 17->1, TX rate=0,01-100%, overall rate=1%, iteration=3
- valkyrie2889-Report-20231208-164735.pdf: service traffic, 17<->1, TX rate=0,01-100%, overall rate=1%, iteration=3

~~valkyrie2889-Report-20231128-094655.pdf (B2B unicast, 16->1, TX rate=1-100%)~~
~~valkyrie2889-Report-20231128-115153.pdf (B2B unicast, 16->1, TX rate=0,01-1%)~~
~~valkyrie2889-Report-20231205-130023.pdf (B2B unicast, 16->1, TX rate=0,01-100%, overall rate=1%, iteration=2)~~
~~valkyrie2889-Report-20231128-152639.pdf (B2B unicast, 12->1, TX rate=0,01-1%)~~
~~valkyrie2889-Report-20231128-162710.pdf (B2B unicast, 8->1, TX rate=0,01-1%)~~
~~valkyrie2889-Report-20231128-173317.pdf (B2B unicast, 6->1, TX rate=0,01-1%)~~
~~valkyrie2889-Report-20231129-100803.pdf (B2B unicast, 7->1, TX rate=0,01-1%)~~
~~valkyrie2889-Report-20231129-111313.pdf (B2B unicast, 6->1, TX rate=1-100%)~~
~~valkyrie2889-Report-20231129-154848.pdf (B2B unicast, 6->1, TX rate=1-10%)~~

In a test case with multiple iterations (eg., B2B unicast, service traffic) only the average value of measurements is recorded (not value of each iteration). Hence, such intermediate values are given as low and high values in graph tables.

One-line commands to extract measurement values from a given report file:

- convert PDF to plain text: $ pdftotext -f 3 -l 6 -layout valkyrie2889-Report-20231201-121634.pdf
- extract ‘Latency’ values and sort: $ for str in Avg Min Max; do echo $str; grep "$str Latency" valkyrie2889-Report-20231201-121634.txt | tr -s ' ' | cut -d " " -f 5 | sort -g; done

3.1. DM Broadcast Forwarding Results, 4 layers

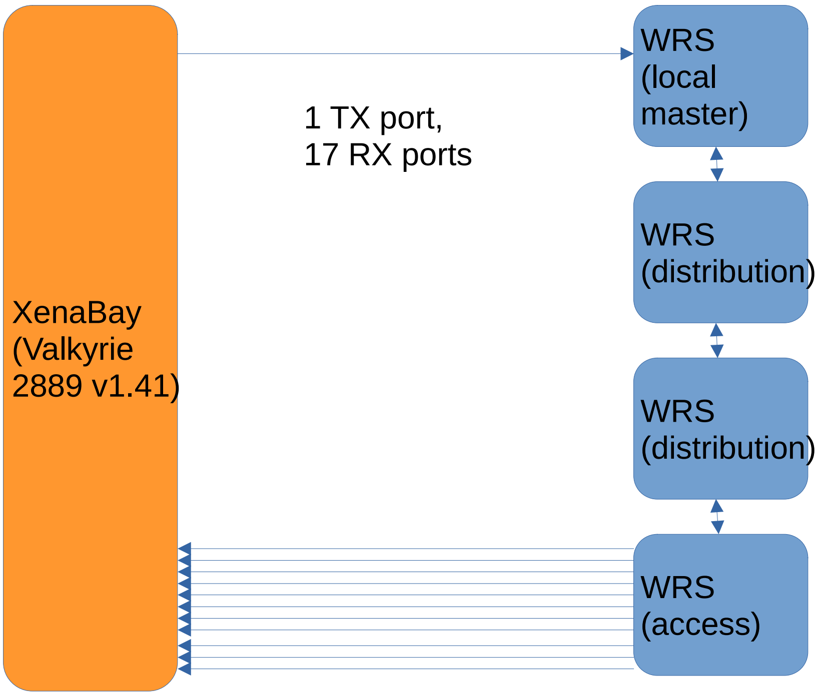

Throughput and latency were measured using the RFC 2889 broadcast forwarding test[1].

Throughput

Timing message broadcast (frame length = 90-750 bytes, rate = 0,1-100%) was generated from 1 source port and received at 17 destination ports. The measurements show that a single timing message could be broadcast up to 89% (corresponds to 889 Mb/s or 1011 Kframe/s) of the total data rate of the WR switch. Additionally, the throughput slightly increased to 98% (974 Mb/s or 158 Kframe/s) when multiple timing messages (eg., 16 messages) were sent per frame.

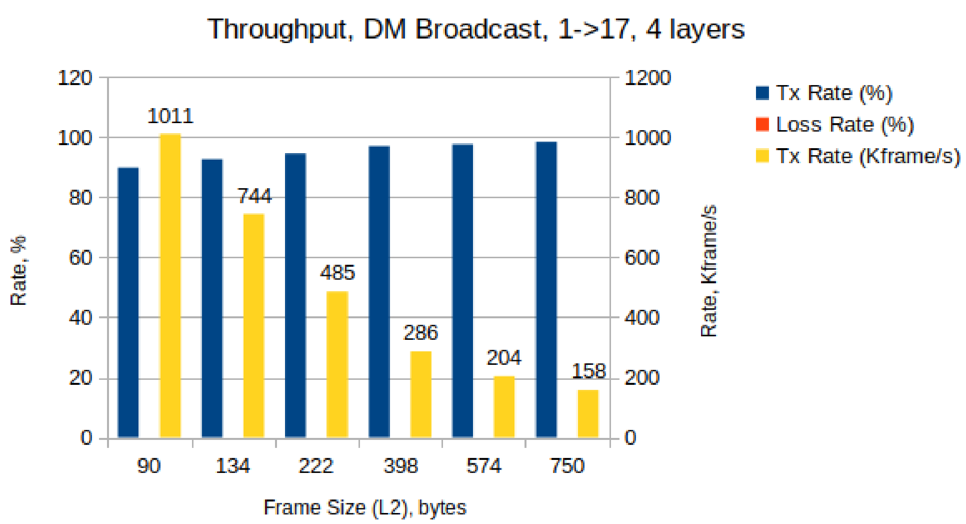

1 valkyrie2889-Report-20231201-121634.pdf

Timing messages	1	2	4	8	12	16
Frame size (L2), bytes	90	134	222	398	574	750
TX rate, %	89,85	92,59	94,54	97,07	97,66	98,46
TX rate, Kframe/s	1011	744	485	286	204	158
Loss rate, %	0	0	0	0	0	0
Loss, frames	0	0	0	0	0	0

Latency and Jitter

The average latency was reported to be around 9,62 to 12,66 us. Here, low and high values are boundary values from all ports. The difference of around 1 us was measured between broadcasting of a single timing message and (a few) multiple timing messages in a frame. The minimum and maximum latencies ranged from 9,39 to 10,84 us, and from 11,96 to 18,78 us, respectively.

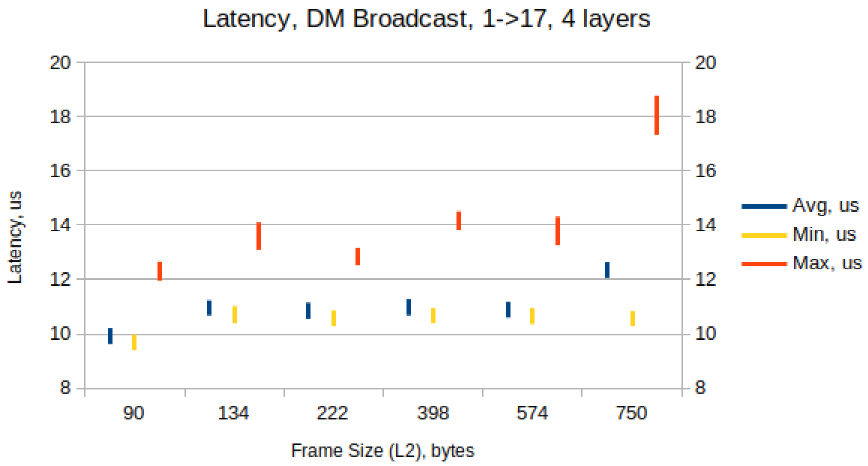

Timing messages	1	2	4	8	12	16
Frame size (L2), bytes	90	134	222	398	574	750
Avg (low), us	9,62	10,68	10,56	10,68	10,6	12,06
Avg (high), us	10,22	11,24	11,15	11,28	11,18	12,66
Min (low), us	9,39	10,39	10,28	10,39	10,36	10,28
Min (high), us	9,98	11,03	10,87	10,95	10,95	10,84
Max (low), us	11,96	13,11	12,54	13,83	13,27	17,33
Max (high), us	12,67	14,1	13,17	14,51	14,31	18,78

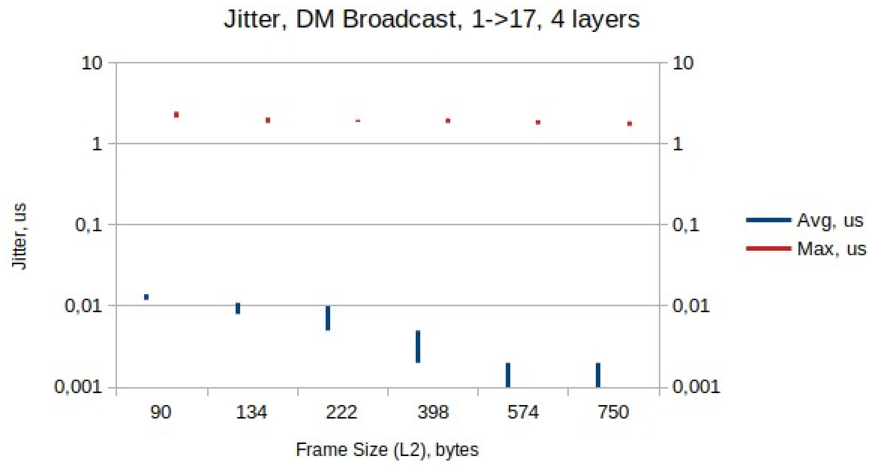

Timing messages	1	2	4	8	12	16
Frame size (L2), bytes	90	134	222	398	574	750
Avg (low), us	0,012	0,008	0,005	0,002	0,001	0,001
Avg (high), us	0,014	0,011	0,01	0,005	0,002	0,002
Min (low, high), us	0	0	0	0	0	0
Max (low), us	2,12	1,83	1,88	1,83	1,75	1,68
Max (high), us	2,5	2,12	1,99	2,07	1,97	1,9

1.2. B2B Unicast Forwarding Results, 4 layers

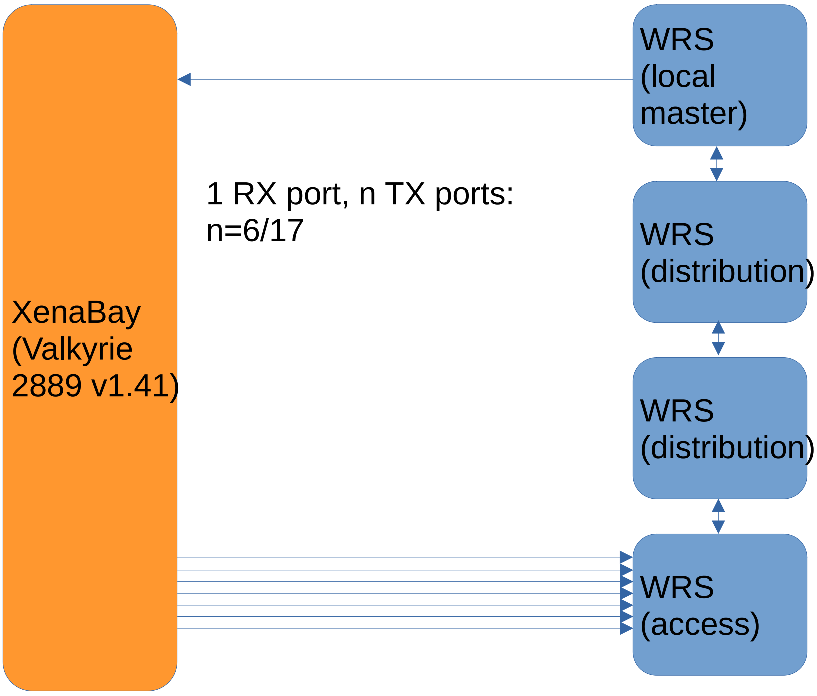

Throughput was measured using the RFC 2889 1:n partial mesh test. In general, unicast traffic from 17 west ports (B2B) was generated and forwarded to 1 east port (DM). Network throughput was measured depending on number of timing messages in a frame (frame length = 90-750 bytes), and number of west ports (n).

Throughput, 17 TX ports[2]

The first test series was done with all west ports (n = 17) and varying B2B unicast frame length (90-750 bytes). The maximum throughput varied for given frame lengths but never exceeded 0,95% (83,7 Kframe/s at 4 messages/frame). The minimum throughput was around 0,015% (2,04 Kframe/s at 2 messages/frame). Because report files do not keep the measurement results of each iteration, they are given in tables as low and high values.

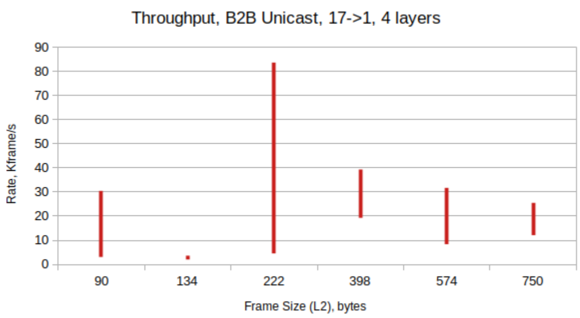

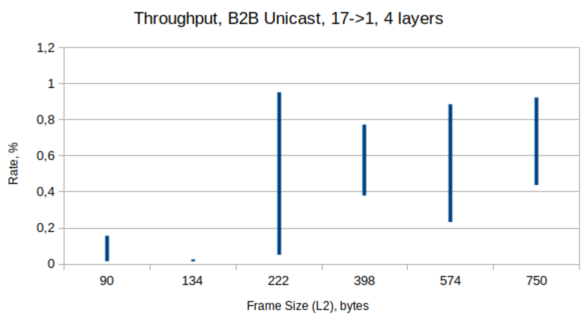

Timing messages	1	2	4	8	12	16
Frame size (L2), bytes	90	134	222	398	574	750
TX rate (low), %	0,016	0,0148	0,052	0,38	0,234	0,439

2 valkyrie2889-Report-20231207-184526.pdf

TX rate (high), %	0,158	0,026	0,953	0,774	0,887	0,924
TX rate (low), Kframe/s	3,1	2,04	4,55	19,3	8,36	12,11
TX rate (high), Kframe/s	30,46	3,59	83,7	39,35	31,72	25,51
TX rate (low), Mb/s	2,72	2,51	8,81	64,55	39,71	74,57
TX rate (high), Mb/s	26,81	4,43	162,04	131,6	150,75	157,13
Loss rate, %	0	0	0	0	0	0
Loss, frames	0	0	0	0	0	0

Throughput, 6 TX ports[3]

The second test series was dedicated to find out the trade-off between maximum number of west ports and maximum network throughput. Unicast traffic without loss was measured with 6 west ports and maximum throughput ranges 15-16% (150-160 Mb/s).

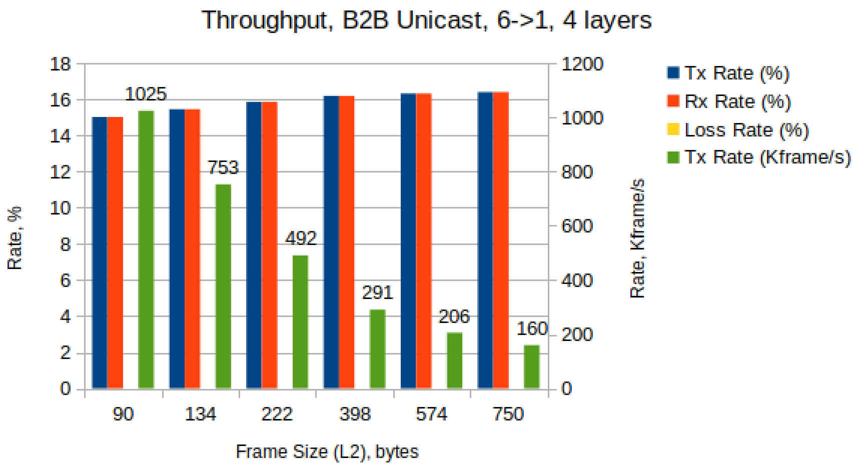

Timing messages	1	2	4	8	12	16
Frame size (L2), bytes	90	134	222	398	574	750
TX rate, %	15,03	15,46	15,87	16,2	16,33	16,41
TX rate, Kframe/s	1025	753	492	291	206	160
TX rate, Mb/s	901	927	952	972	980	984
Loss rate, %	0	0	0	0	0	0
Loss, frames	0	0	0	0	0	0

3 valkyrie2889-Report-20231201-110640.pdf

1.3. Service Traffic Forwarding Results, 4 layers

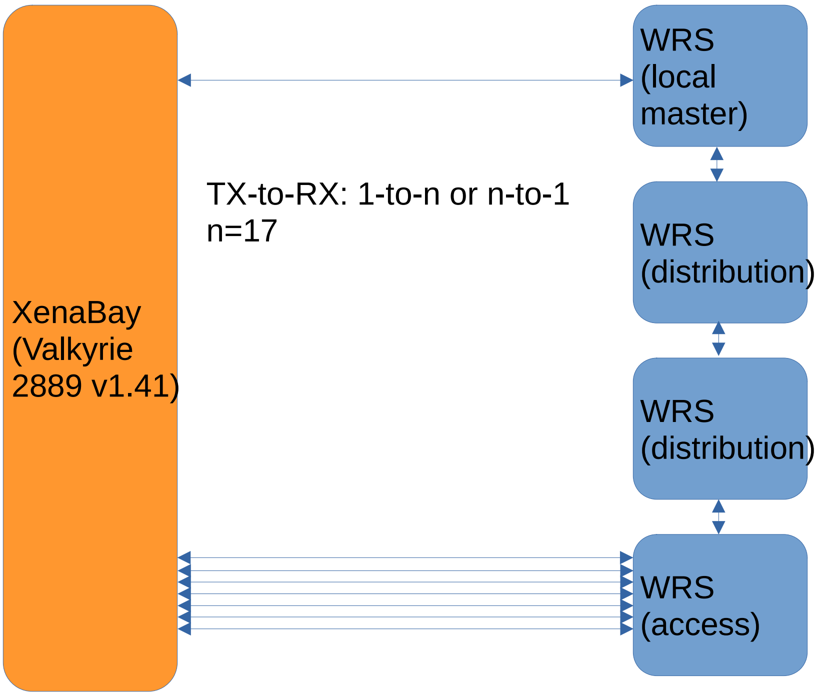

Throughput was measured by using the RFC 2889 1:N mesh test[4].

Throughput

Bi-directional traffic (64-1518 bytes) between 1 east port and 17 west ports was generated and analyzed. The maximum achieved throughput was around 0,00085% (142 Kb/s) of TX rate. Low and high values were taken from all iteration tests.

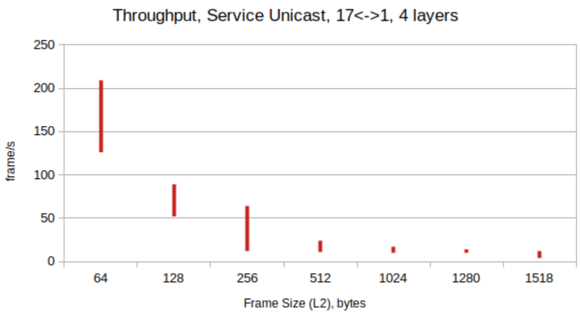

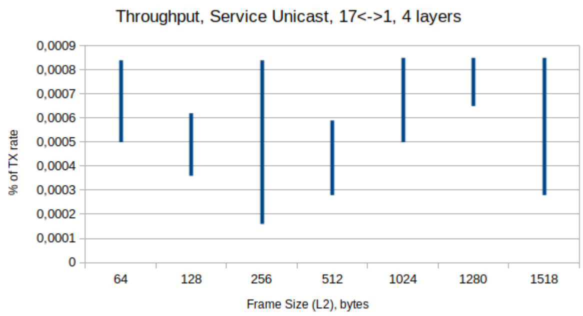

Frame size, bytes	64	128	256	512	1024	1280	1518
TX rate (low), %	0,0005	0,00036	0,00016	0,00028	0,0005	0,00065	0,00028
TX rate (high), %	0,00084	0,00062	0,00084	0,00059	0,00085	0,00085	0,00085
TX rate (low), Kframe/s	126	52	12	11	10	10	4
TX rate (high), Kframe/s	209	89	64	24	17	14	12
Loss rate, %	0	0	0	1	0	1	0
Loss, frames	0	0	0	0,12	0	0,19	0

4 valkyrie2889-Report-20231208-164735.pdf

1.4. Mixed Traffic, 4 layers

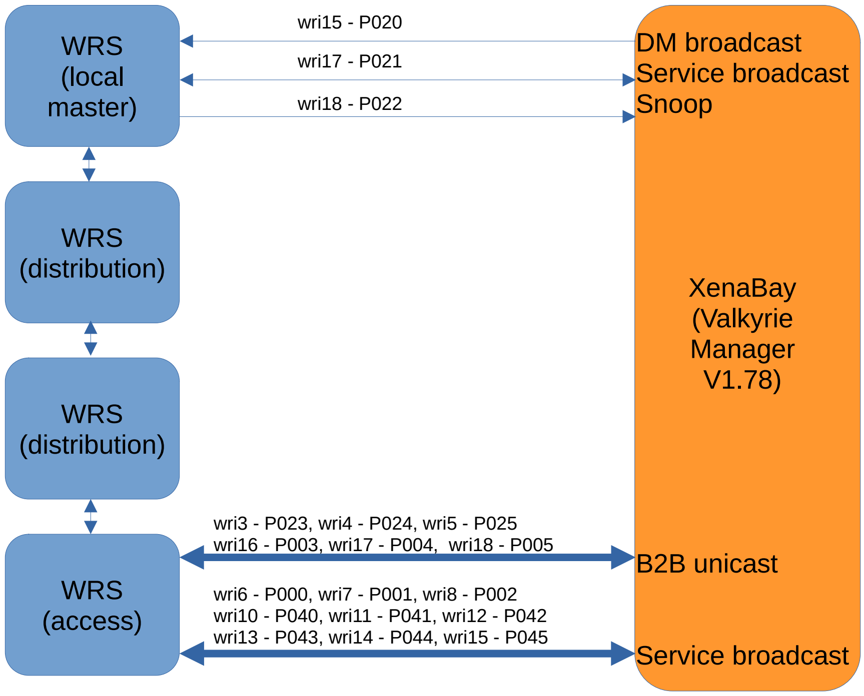

Traffic types of DM broadcast, B2B unicast, and service broadcast were generated with different data rate and frame length according to table below. The frame rate and length of the DM and B2B frames were used as iteration parameters for each measurement. Each run takes 5 minutes. The service broadcast traffic was generated with data rate of 142 Kb/s and frame length of 64-1518 bytes for all measurements.

Layer	WRS port	XenaBay port	Data rate, Mb/s	Frame size, bytes (Frame rate, Kframe/s)	Traffic
1	LM:wri15	P020	10, 20, 50, 100	90, 134 (113,64, 81,17)	DM
	LM:wri17	P021	0,14	64-1518 (21)	Service trunk
	LM:wri18	P022			(snooping)
4	A:wri3	P023	2, 4, 8, 16, 32	90, 134 (4,54, 3,25)	B2B
	A:wri4	P024	2, 4, 8, 16, 32	90, 134	B2B
	A:wri5	P025	2, 4, 8, 16, 32	90, 134	B2B
	A:wri6	P000	0,14	64-1518	Service
	A:wri7	P001	0,14	64-1518	Service
	A:wri8	P002	0,14	64-1518	Service
	A:wri10	P040	0,14	64-1518	Service
	A:wri11	P041	0,14	64-1518	Service
	A:wri12	P042	0,14	64-1518	Service
	A:wri13	P043	2, 4, 8, 16, 32	90, 134	B2B
	A:wri14	P044	2, 4, 8, 16, 32	90, 134	B2B
	A:wri15	P045	2, 4, 8, 16, 32	90, 134	B2B

	A:wri16	P003	0,14	64-1518	Service
	A:wri17	P004	0,14	64-1518	Service
	A:wri18	P005	0,14	64-1518	Service

DM Frame Loss

No frame loss in the DM broadcast was detected within probed data rates up to 100 Mb/s, and frame length of 90 and 134 bytes.

B2B Frame Loss

Frame loss in the B2B unicast was detected only with the frame length of 90 bytes:

B2B frame loss	B2B frame	B2B frame length, bytes
Detected	1 message/frame	90
Not detected	2, 4, 8 messages/frame	134, 222, 398

In addition, the frame loss occurred in higher data rates of the DM and B2B traffic.

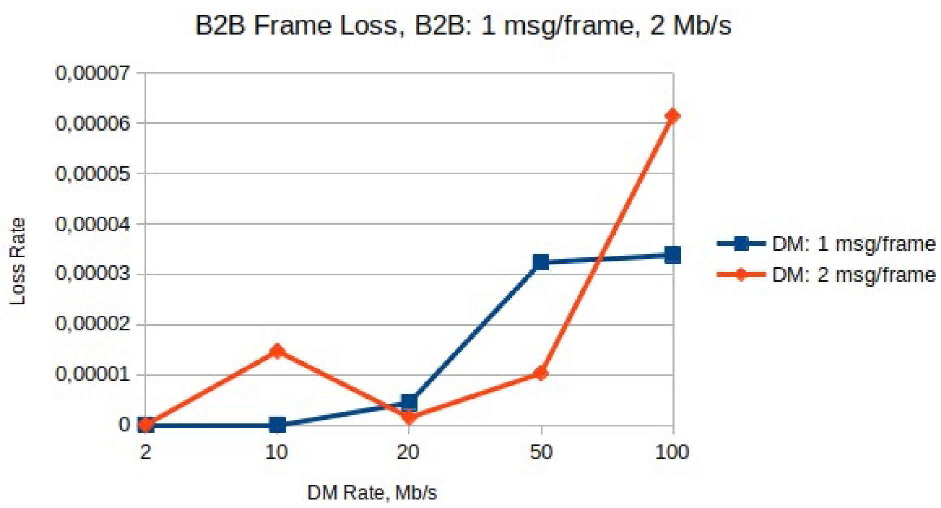

DM rate, Mb/s		10	20	50	100
DM rate, frame/s		11636	22727	56818	113636
DM msg/frame = 1 (frame = 90 bytes)	B2B frame loss, frames	0	3	22	23
	B2B frame loss rate	0	$4,4*10^{-6}$	$3,2*10^{-5}$	$3,3*10^{-5}$
DM msg/frame = 2 (frame = 134 bytes)	B2B frame loss, frames	10	1	7	42
	B2B frame loss rate	$1,5*10^{-5}$	$1,5*10^{-6}$	$1,0*10^{-5}$	$6,1*10^{-5}$

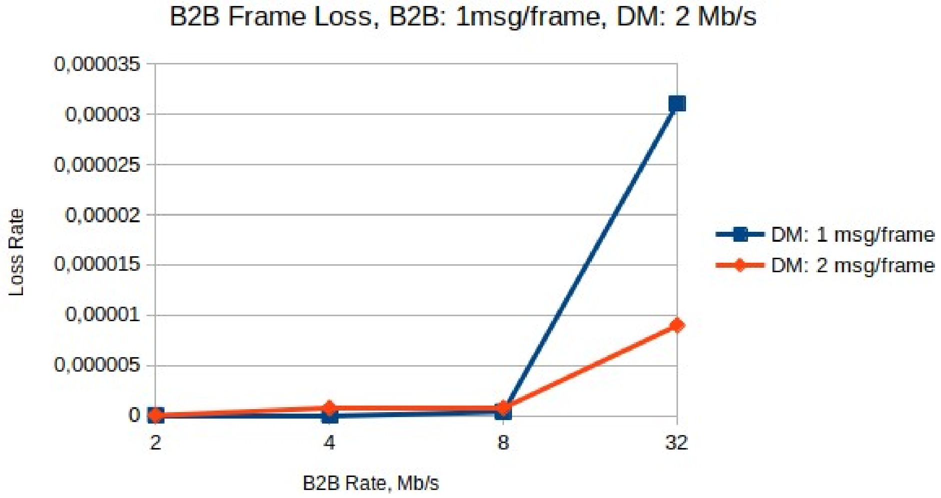

B2B rate, Mb/s		2	4	8	32
B2B rate, frame/s		2272	4545	9090	36363
DM msg/frame = 1 (frame = 90 bytes)	B2B frame loss, frames	0	0	1	339
	B2B frame loss rate	0	0	$3{,}66 \cdot 10^{-7}$	$3{,}1 \cdot 10^{-5}$
DM msg/frame = 2 (frame = 134 bytes)	B2B frame loss, frames	0	1	2	98
	B2B frame loss rate	0	$7{,}32 \cdot 10^{-7}$	$7{,}32 \cdot 10^{-7}$	$8{,}96 \cdot 10^{-6}$

Service Frame Loss

According the network configuration there are two uni-directional service traffic is enabled: upstream (layer 4 to layer 1) and downstream (layer 1 to layer 4). Frame loss occurred in the downstream traffic at rate of $1{,}85 \cdot 10^{-3}$.

Latency and Jitter

Latency were not evaluated, but one corner case measurement shows that it takes 9,64-10,87 us in average for all traffic types. The maximum latency of 22,55-23,36 us was measured for the DM broadcast and B2B unicast traffics.

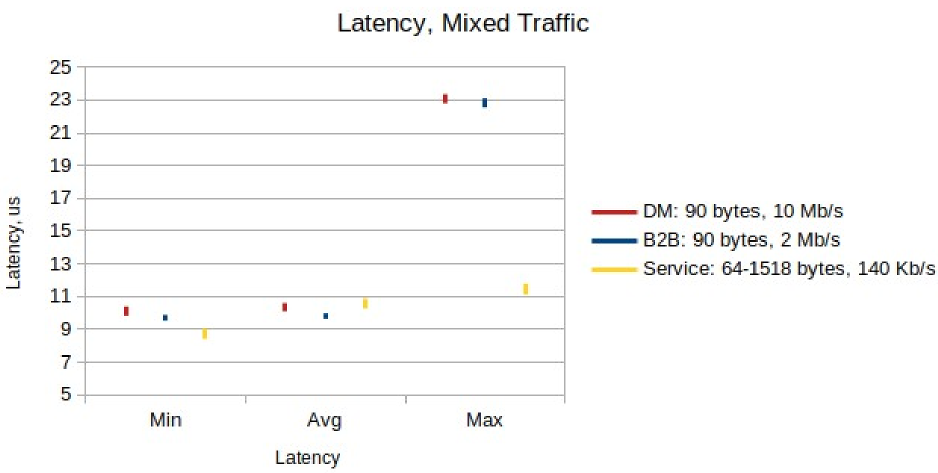

Traffic types	**Min latency, us**	**Avg latency, us**	**Max latency, us**
DM: 90 bytes, 10 Mb/s	9,83-10,4	10,1-10,61	22,8-23,36
B2B: 90 bytes, 2 Mb/s	9,53-9,88	9,64-9,98	22,55-23,11
Service: 64-1518 bytes, 140 Kb/S	8,41-9,05	10,28-10,87	11,1-11,77