



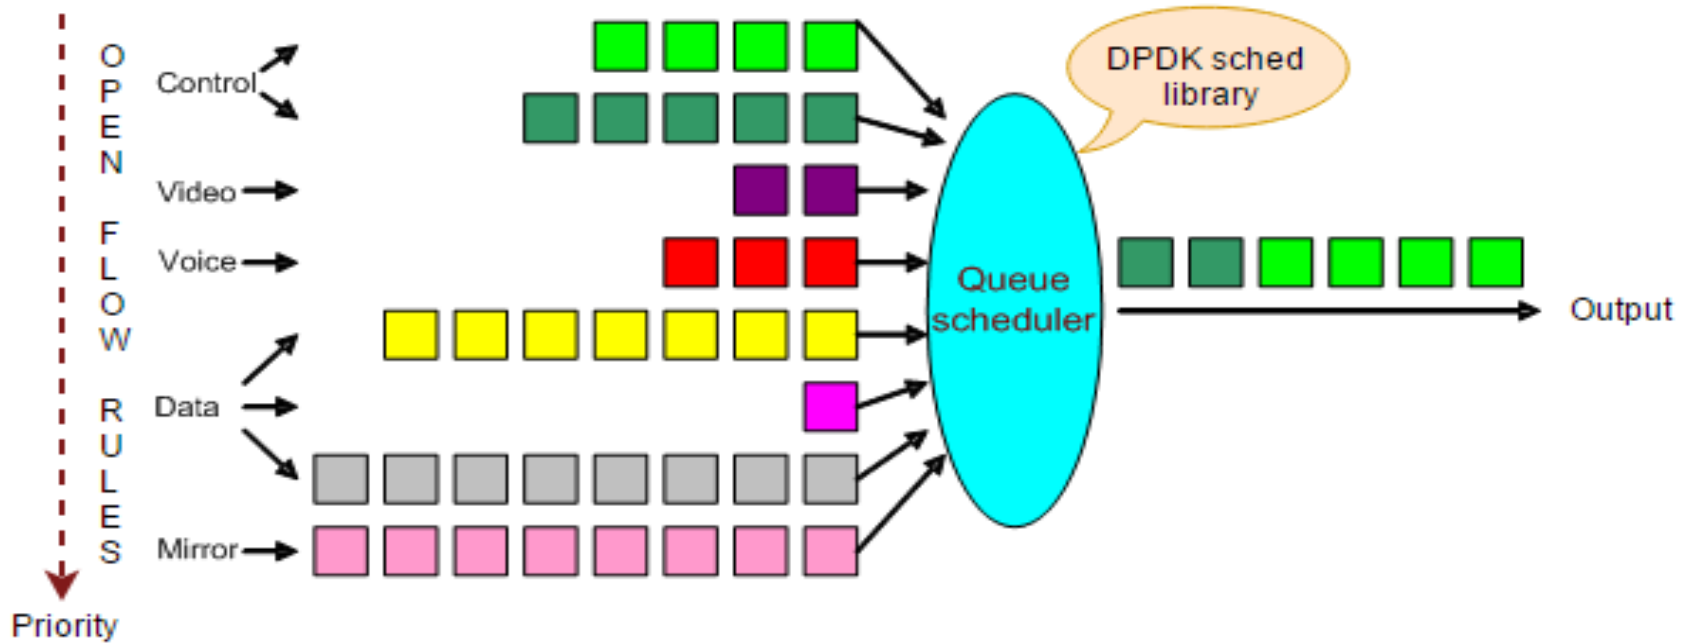
'Flow based QoS' for SDN- NFV platform using OVS and DPDK on Intel x86 Architecture

Karuppusamy M

Solution Overview

- The project objective is to implement 'Flow based QoS' feature for SDN-NFV platform using OVS and DPDK on Intel x86 architecture
- Enable support for packet classification and shaping on egress queue of OVS-DPDK, which is currently not supported
- Enable support for flow classification and shaping on egress even when there is no QoS marking in ingress stream
- **Approach:**
 - Classify traffic types
 - Assign higher priority queues for processing real time traffic
- Assign lower priority queues for processing best effort traffic with necessary rate limiting algorithms to shape the traffic
- **Benefits:**
 - Reduce packet loss, latency and jitter
 - Ensure deterministic performance of real time applications
 - Eliminate the need for dedicated custom hardware for QoS support

QoS with OVS-DPDK



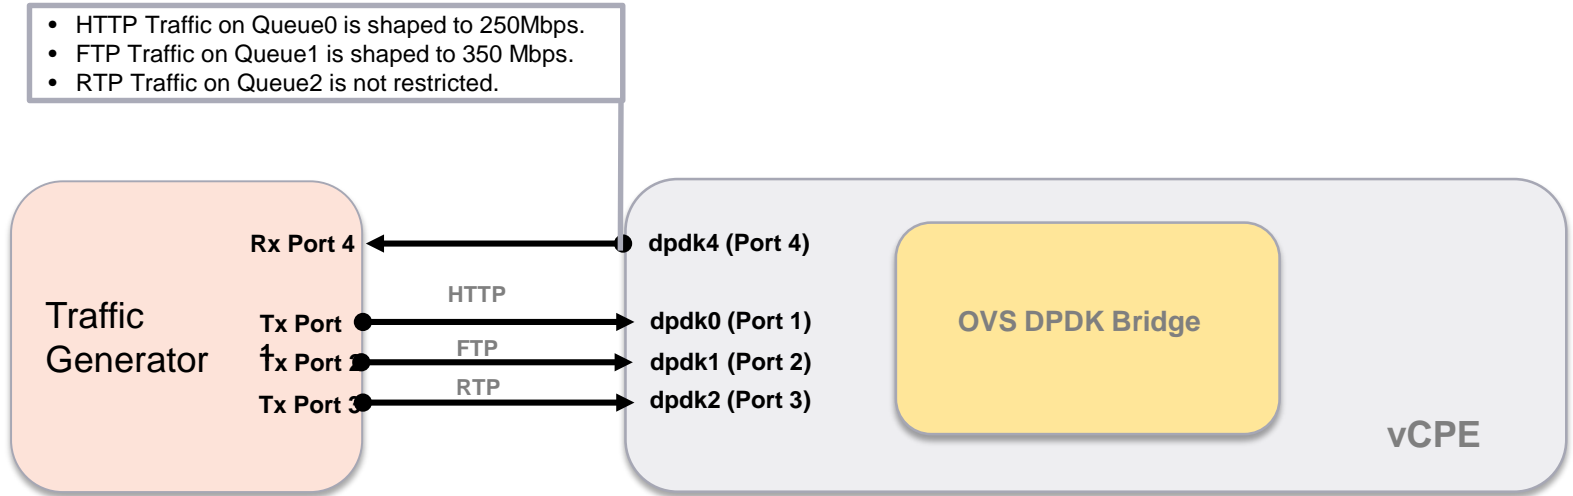
Intel Technologies leveraged

INTEL TECHNOLOGY	BENEFITS
Intel Xeon-E	Xeon-E based platform variant can fit in the mid / large branch office segment and central DC as the private/public cloud
Intel Xeon-D	Xeon-D based platform variant can fit in the mid segment deployment in the branch office
ATOM C2000	ATOM based platform fits for small office or home office environment
I350/i210 based NIC	Used as a physical interface to realize the DPDK accelerated network interface

Partner Technologies Leveraged

PARTNER TECHNOLOGY	BENEFIT TO POC
OpenvSwitch and DPDK	Enabling the QoS feature would accelerate the adoption of Intel x86 based virtualization platform in realizing networking use cases

QoS test topology



Platform configuration

```
[root@lab-atom ~]# lscpu
Architecture:          x86_64
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Little Endian
CPU(s):                8
On-line CPU(s) list:  0-7
Thread(s) per core:    1
Core(s) per socket:    8
Socket(s):              1
NUMA node(s):          1
Vendor ID:              GenuineIntel
CPU family:             6
Model:                 77
Model name:             Intel(R) Atom(TM) CPU C2758 @ 2.40GHz
Stepping:               8
CPU MHz:                2400.000
BogoMIPS:              4799.52
Virtualization:        VT-x
L1d cache:             24K
L1i cache:             32K
L2 cache:              1024K
NUMA node0 CPU(s):    0-7
[root@lab-atom ~]# lspci | grep -i ether
00:14.0 Ethernet controller: Intel Corporation Ethernet Connection I354 (rev 03)
00:14.1 Ethernet controller: Intel Corporation Ethernet Connection I354 (rev 03)
00:14.2 Ethernet controller: Intel Corporation Ethernet Connection I354 (rev 03)
00:14.3 Ethernet controller: Intel Corporation Ethernet Connection I354 (rev 03)
04:00.0 Ethernet controller: Intel Corporation I350 Gigabit Network Connection (rev 01)
04:00.1 Ethernet controller: Intel Corporation I350 Gigabit Network Connection (rev 01)
04:00.2 Ethernet controller: Intel Corporation I350 Gigabit Network Connection (rev 01)
04:00.3 Ethernet controller: Intel Corporation I350 Gigabit Network Connection (rev 01)
[root@lab-atom ~]#
[root@lab-atom ~]# free -g
              total        used          free      shared    buffers         cached
Mem:           15           10             4             0             0             0
-/+ buffers/cache:
              10           10             4
Swap:           7             0             7
[root@lab-atom ~]#
```

Details of the Intel x86 CPUs

Family: Atom
Manufacturer: Intel(R) Corporation
Version: Intel(R) Atom(TM) CPU C2758 @ 2.40GHz
Voltage: 1.6 V
External Clock: 100 MHz
Max Speed: 2600 MHz
Current Speed: 2400 MHz
Core Count: 8
Core Enabled: 8
Thread Count: 8
Characteristics: 64-bit capable

Family: Xeon
Manufacturer: Intel
Version: Intel(R) Xeon(R) CPU D-1548 @ 2.00GHz
Voltage: 0.0 V
External Clock: 100 MHz
Max Speed: 4000 MHz
Current Speed: 2000 MHz
Core Count: 8
Core Enabled: 8
Thread Count: 16

Characteristics: 64-bit capable, Multi-Core, Hardware Thread, Execute Protection, Enhanced Virtualization, Power/Performance Control

Details of the Intel x86 CPUs and NIC

Family: Xeon

Manufacturer: Intel

Version: Intel(R) Xeon(R) CPU E5-2603 v3 @ 1.60GHz

Voltage: 1.8 V

External Clock: 100 MHz

Max Speed: 4000 MHz

Current Speed: 1600 MHz

Core Count: 6

Core Enabled: 6

Thread Count: 6

Characteristics: 64-bit capable, Multi-Core, Execute Protection, Enhanced Virtualization, Power/Performance Control

Ethernet controller: Intel Corporation I350 Gigabit Network Connection (rev 01)

Ethernet controller: Intel Corporation I210 Gigabit Network Connection (rev 03)

Sample Configuration to enable QoS

- To enable tx multi queuing on dpdk0 interface

```
ovs-vsctl --no-wait set Interface dpdk0 options:n_txq=4
```

- To enable egress shaping of 100Mbps on single queue (q0) on dpdk0 interface

```
ovs-vsctl -- set port dpdk0 qos=@newqos -- --id=@newqos  
create qos type=none \ queues=0=@q0 -- --id=@q0 create  
Queue type=egress-shaper \ other-config:cir=12500000  
other-config:cbs=125000
```

Performance(w/o QoS) - Equal treatment for all the flows

Vhost_and_Phy_Testing_Config.tcc - Spirent TestCenter

File View Tools Actions Diagnostics Help

00:00:00 Technologies... Perspective Sequencer Reporter Wizards Summary... Manage Tags...

Test Configuration

Start Traffic Stop Traffic Manual Schedule

Tx State	Port Name	Tags	Scheduling Mode	Duration Mode	Duration	Burst Size	Inter Frame Gap	Inter Frame Gap Unit	Load Mode	Load Unit	Load	Random Min Load	Rar Load
●	Port //1/1	Click to ad...	Port Based	Continuous		1			Fixed	Percent (%)	10		
●	Port //1/2	Click to ad...	Port Based	Seconds	60	1			Fixed	Percent (%)	60		
●	Port //1/3	Click to ad...	Port Based	Seconds	60	1			Fixed	Percent (%)	60		
▶	Port //1/4	Click to ad...	Port Based	Seconds	60	1			Fixed	Percent (%)	40		
●	Port //1/5	Click to ad...	Port Based	Seconds	60	1			Fixed	Percent (%)	1		

Displaying Traffic Generators 1 - 5 | Total Traffic Generators: 5 | Selected 3 of 5

Vhost_and_Phy_Testing_Config:Results 1

Streams > Detailed Stream Results | Change Result View | 1 of 1 | Select Tx Ports: All Ports | Select Rx Ports: All Ports | Change Counter Mode: Latency-Jitter Moc

Resample

Name/ID	Tx L1 Rate (bps)	Rx L1 Rate (bps)	Tx Rate (fps)	Rx Rate (fps)	Rx Sig Count (Frames)	Rx Sig Rate (fps)	Short Term Avg Latency (us)	Avg Latency (us)	Min Latency (us)	Max Latency (us)	Avg Jitter (us)	Max Jitter
Q0/131072	600,000,153	333,330,709	140,978	78,321	3,475,395	78,321	26,227.47	26,184.22	47.28	26,268.72	2.34	54.2
Q1/196608	600,000,146	333,332,062	140,978	78,321	3,471,943	78,321	26,227.81	26,210.32	60.38	26,265.6	2.38	101.01
Q2/262144	400,000,098	333,333,415	93,985	78,321	3,474,735	78,321	26,226.18	26,159.11	9.58	26,244.8	3.4	70.22
Q3/327680	0	0	0	0	0	0	0	0	0	0	0	0

Performance(with QoS) – Prioritized treatment for RTP

Vhost_and_Phy_Testing_Config.tcc - Spirent TestCenter

File View Tools Actions Diagnostics Help

Chassis Apply 00:00:00 Technologies... Perspective Sequencer Reporter Wizards Summary... Manage Tags...

Test Configuration

Spirent TestCenter

- All Devices (Hosts, Routers, ...)
- All Profiles
- All Traffic Generators
- All Stream Blocks
- All Traffic Analyzers
- All Ports
 - Port //1/1
 - Port //1/2
 - Port //1/3
 - Port //1/4
 - Port //1/5
 - Settings

Start Traffic Stop Traffic Manual Schedule

Tx State	Port Name	Tags	Scheduling Mode	Duration Mode	Duration	Burst Size	Inter Frame Gap	Inter Frame Gap Unit	Load Mode	Load Unit	Load	Random Min Load	Rar Lo
	Port //1/1	Click to ad...	Port Based	Continuous		1			Fixed	Percent (%)	10		
	Port //1/2	Click to ad...	Port Based	Seconds	60	1			Fixed	Percent (%)	60		
	Port //1/3	Click to ad...	Port Based	Seconds	60	1			Fixed	Percent (%)	60		
	Port //1/4	Click to ad...	Port Based	Seconds	60	1			Fixed	Percent (%)	40		
	Port //1/5	Click to ad...	Port Based	Seconds	60	1			Fixed	Percent (%)	1		

Displaying Traffic Generators 1 - 5 | Total Traffic Generators: 5 | Selected 3 of 5

Vhost_and_Phy_Testing_Config:Results 1


Streams > Detailed Stream Results | Change Result View | 1 of 1 | Select Tx Ports: All Ports | Select Rx Ports: All Ports | Change Counter Mode: Latency-Jitter Moc

Resample

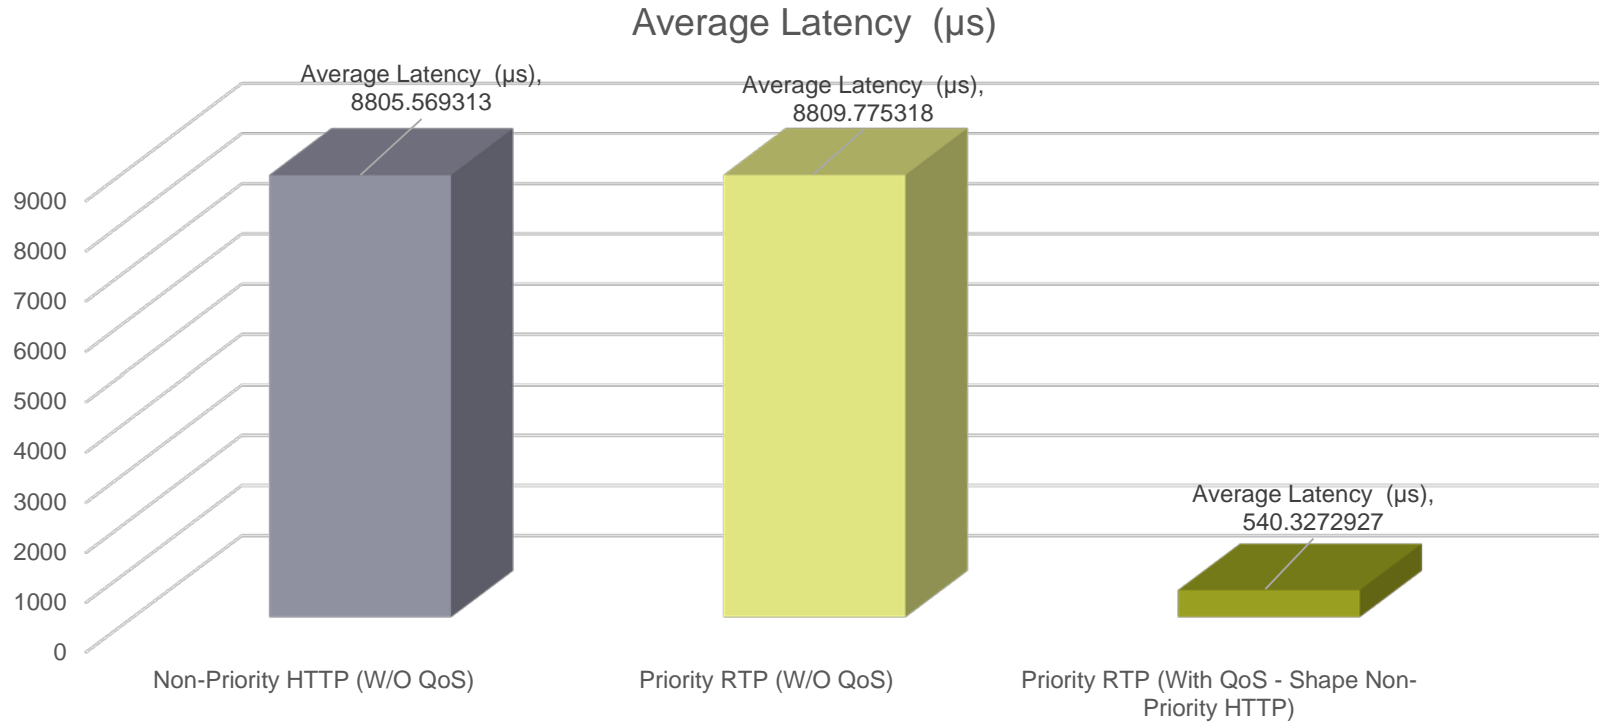
Basic Counters Errors Basic Sequencing Advanced Sequencing Histograms

Name/ID	Tx L1 Rate (bps)	Rx L1 Rate (bps)	Tx Rate (fps)	Rx Rate (fps)	Rx Sig Count (Frames)	Rx Sig Rate (fps)	Short Term Avg Latency (us)	Avg Latency (us)	Min Latency (us)	Max Latency (us)	Avg Jitter (us)	Max Jitter
Q0/131072	600,000,146	249,462,941	140,978	58,614	570,564	58,614	1,131.16	2,165.44	452.11	9,233,776.64	37.7	9,233,305
Q1/196608	600,000,146	333,547,180	140,978	78,371	760,073	78,371	1,348.47	2,121.07	793.29	9,255,301.12	25.02	9,254,502
Q2/262144	400,000,098	399,999,827	93,985	93,985	916,501	93,985	79.6	78.95	9.58	270.8	2.28	70.22
Q3/327680	0	0	0	0	0	0	0		0	0	0	0

Latency comparison with and without QoS (ATOM) on the Physical ports

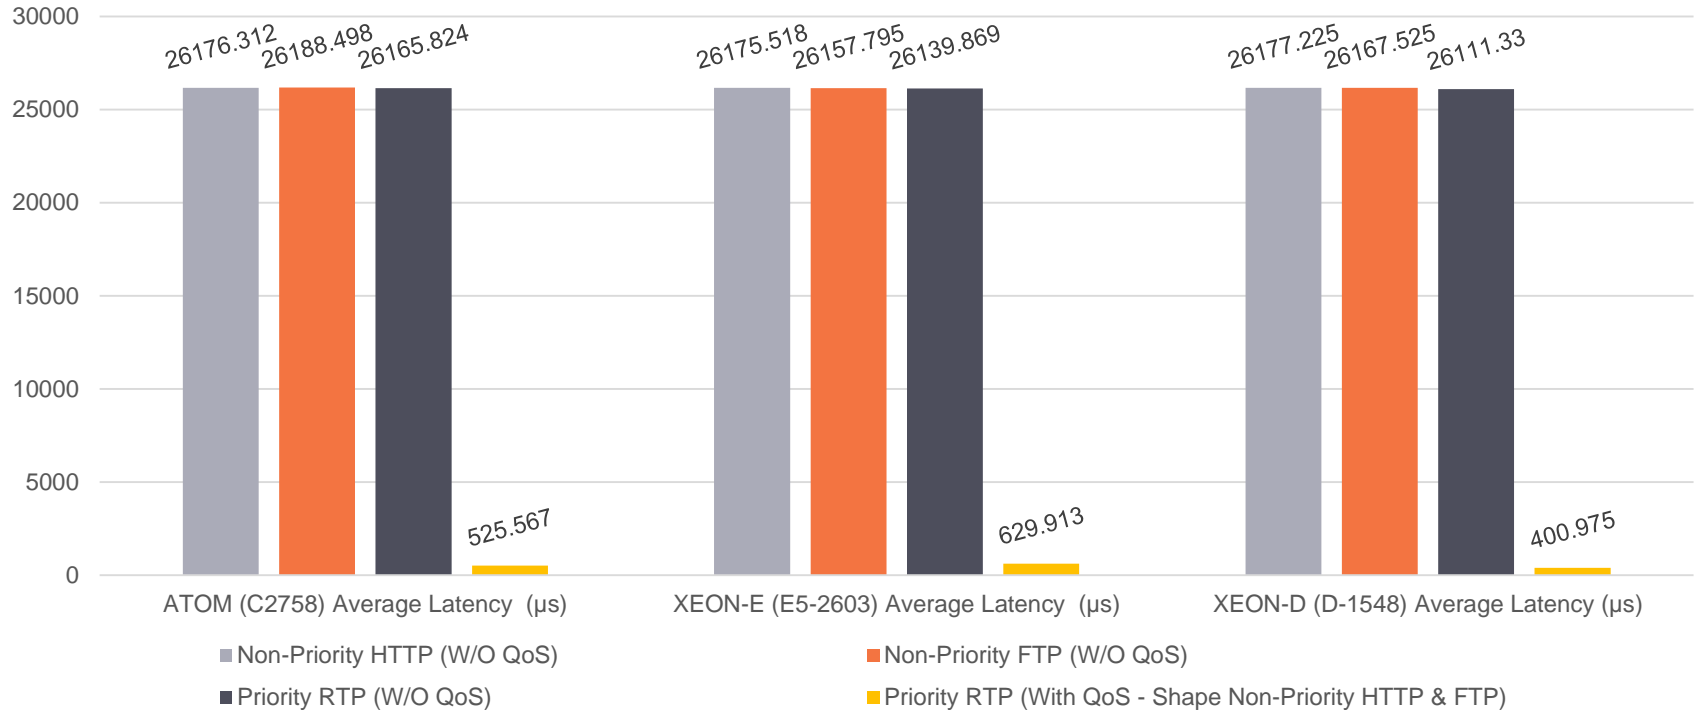


Latency comparison with and without QoS on DPDK vhost-user Interface (Xeon-E platform)



Latency across platforms

Average Latency Comparison on Physical interface Priority Queue (μ s)



Key Benefits / Use cases

Business Benefits

- To deploy various latency sensitive services which needs QoS handling in the OVS-DPDK environment
- Ability to create differentiated service chaining path for critical network functions

Use Case(s)

- Improving quality of video traffic
- Differentiated Service Function Chaining
- Enabling QoS profile support based on application priority

Market Potential

- Telecom Service Provider- For providing enterprise connectivity in a managed services model.
- Enterprise- To address their Branch networking requirement
- IOT- Industrial process automation and manufacturing customers who needs an IOT gateway as part of their IIOT offering.



Thank You