


# TCP Flow Control Quiz

# Question 1


- Consider the network from Client to Server. Link data rates and one-way delays are same in each direction. Ignore other delays.



- What is the RTT?

# Question 1


- Consider the network from Client to Server. Link data rates and one-way delays are same in each direction. Ignore other delays.



- What is the RTT?


**20 ms**

## Question 2



- If UDP is used to send data from client to server, what is the maximum throughput that can be achieved? (ignoring header overheads)


## Question 2



- If UDP is used to send data from client to server, what is the maximum throughput that can be achieved? (ignoring header overheads)


**5 Mb/s**

# Question 3



- What is the Bandwidth Delay Product of the path?


# Question 3



- What is the Bandwidth Delay Product of the path?

**100,000 bits or 12,500 Bytes**


# Question 4



- How big should the server buffer be to achieve maximum TCP throughput?




# Question 4



- How big should the server buffer be to achieve maximum TCP throughput?


**$\geq 12,500$  Bytes**

# Question 5



- What is the maximum TCP throughput if the server buffer is 6000 Bytes?


# Question 5



- What is the maximum TCP throughput if the server buffer is 6,000 Bytes?

**2.4 Mb/s**

# Question 6




- In Linux, when an application opens a TCP connection, it can request memory to be allocated by the OS for the receiver. The OS allocates *double* of what is requested – this is to ensure there is enough space for received data *and* data structures used by the OS for TCP connections. For example:

```
iperf -s -w 20000
```

iperf requests 20,000 bits. How much memory is allocated?

# Question 6




- In Linux, when an application opens a TCP connection, it can request memory to be allocated by the OS for the receiver. The OS allocates *double* of what is requested – this is to ensure there is enough space for received data *and* data structures used by the OS for TCP connections. For example:

```
iperf -s -w 20000
```

iperf requests 20,000 bits. How much memory is allocated?

**40,000 bits**


# Question 7



- Of the memory allocated by the OS, assume 75% is for data received and 25% is for data structures used by the OS. What value should iperf server use for receiver window to achieve optimal TCP throughput?

`iperf -s -w ???`

# Question 7




- Of the memory allocated by the OS, assume 75% is for data received and 25% is for data structures used by the OS. What value should iperf server use for receiver window to achieve optimal TCP throughput?

`iperf -s -w ???`

**66,666 bits**

# Question 8




- Of the memory allocated by the OS, assume half is for data received and half is for data structures used by the OS. What value should iperf server use for receiver window to achieve optimal TCP throughput?

`iperf -s -w ???`



# Question 8



- Of the memory allocated by the OS, assume half is for data received and half is for data structures used by the OS. What value should iperf server use for receiver window to achieve optimal TCP throughput?

`iperf -s -w ???`

**100,000 bits**

Let's try some real experiments ...

<http://sandilands.info/sgordon/impact-of-bandwidth-delay-product-on-tcp-throughput>