Practice Exam Questions for Nokia Quality of Service (exam number: 4A0-107)

The following questions will test your knowledge and help prepare you for the Nokia Quality of Service exam. Compare your responses with the Answer Key at the end of the document.

1. Which of the following is NOT a characteristic of policing?
   a. It smooths out traffic bursts.
   b. It drops the packets when the average peak rate is exceeded.
   c. It may additionally color the packets as they arrive.
   d. It buffers the packets at its output before sending them out.

2. Which of the following statements about the four-step QoS process (classification, policing/shaping, scheduling, and marking/remarking) executed for every packet entering a Nokia 7750 SR is TRUE?
   a. The process happens only before the packet goes through the switching fabric.
   b. The process happens only before the packet goes onto the egress port.
   c. Marking/remarking the packets is for future classification.
   d. A packet cannot be classified and policed on the same router.

3. Which of the following policies does NOT exist by default on a network port?
   a. A network QoS policy for every network interface configured on the port
   b. A slope policy for the port on ingress
   c. A network queue policy for the corresponding FFPC on ingress
   d. A network queue policy for the port on egress

4. Which of the following is a forwarding class type that is intended for network control traffic on a Nokia 7750 SR?
   a. High-1
   b. Assured
   c. Best Effort
   d. Expedited

5. Which of the following fields is in the MPLS header?
   a. 802.1p bits
   b. Precedence
   c. DSCP
   d. EXP

6. Which of the following statements about behavior-aggregate and multi-field classification on a Nokia 7750 SR is TRUE?
   a. Multi-field classification can use both MAC criteria and IP criteria in the same SAP-ingress QoS policy.
   b. Behavior aggregate classification can use both lsp-exp and DSCP in the same SAP-ingress QoS policy.
   c. Multi-field classification cannot use both IP criteria and IPv6 criteria in the same SAP-ingress QoS policy.
   d. Behavior aggregate classification cannot use both dot1p and lsp-exp in the same SAP-ingress QoS policy.

7. Which of the following criteria is an IP multi-field classification option?
   a. prec
   b. lsp-exp
   c. src-ip
   d. dot1p

8. The SAP-ingress policy shown in the exhibit is applied properly to a Nokia 7750 SR. A traffic stream is received from an IP address of 192.168.2.200 with DSCP and dot1p set to EF and 4, respectively. Which forwarding class and priority level is used for this traffic stream?

```
[pr:/configure qos sap-ingress "111"]
A:student@R01(PE1)# info
  ip-criteria {
    entry 10 {
      match {
        dscp ef
      }
      action {
        type accept
        fc "h1"
      }
    }
    entry 20 {
      match {
        src-ip {
          address 192.168.2.200
        }
      }
      action {
        type accept
        fc "af"
        priority low
      }
    }
  }
  dot1p 4 {
    fc "l1"
    priority high
  }
```

   a. FC H1, priority high
   b. FC H1, priority low
   c. FC AF, priority high
   d. FC AF, priority low

9. Which of the following statements about network QoS policies is FALSE?
   a. Classification rules must specify both an FC and a profile for the packets.
   b. A network policy is applied at a physical port.
   c. Classification can be performed with behavior-aggregate and/or multi-field statements.
   d. Classification using lsp-exp bits has the highest precedence.

10. Which of the following egress points does NOT support egress reclassification?
   a. Spoke-SDPs bound to an IES
   b. Spoke-SDPs bound to a VPRN interface
   c. SAPs
   d. Network interfaces

11. Which ports on the same FFPC share a buffer pool on a Nokia 7750 SR?
   a. Access ingress ports
   b. Access egress ports
   c. Network ingress ports
   d. Network egress ports

12. Which of the following statements about configuring queues to use shared or reserved buffer space is FALSE?
   a. When queues use shared buffer space, they need to compete with other queues for memory.
   b. Queues that use reserved buffer space will have memory available for their reserved portion.
   c. When a queue is not using its reserved portion, other queues can use it.
   d. Using reserved buffer space prevents a queue from being starved of buffer resources by other queues.

13. A profile-mode queue is configured with MBS = 100 KB, CBS = 50 KB, drop-tail low = 40% and the slope-policy is disabled. If the current queue depth is 70 KB, what will happen to an out-of-profile packet when it is transmitted to the queue?
   a. It will be stored in the reserved buffer space and remain out-of-profile.
   b. It will be stored in the reserved buffer space and become in-profile.
   c. It will contend for shared buffer space with packets from other queues.
   d. It will be discarded because the high-priority-only threshold has been reached.

14. In which of the following queues does a Nokia 7750 SR allocate a separate hardware queue for each destination Fast Forward Path Complex (FFPC)?
   a. An egress unicast queue defined in a network queue policy
   b. An ingress unicast queue defined in a SAP-ingress policy
   c. An ingress multipoint queue defined in a SAP-ingress policy
   d. An ingress multipoint queue defined in a network queue policy

15. Which of the following statements about CFHP policing hierarchy is FALSE?
    a. The root arbiter can be the parent of a tier 2 arbiter.
    b. A tier 1 arbiter can be the parent of a policer.
    c. A child policer can have a priority level from 0 to 10.
    d. Weight is used by a parent arbiter to allocate bandwidth among children of equal priority level.

16. Consider the exhibit. Both SAP-ingress 111 and policer-control-policy "CFHP" have been applied to a service SAP ingress. All policers are being offered traffic at their PIR. What is the maximum bandwidth that can be allocated to policer 2?

```
[pr:/configure qos]
A:student@R01(PE1)# info
  sap-ingress "111" {
    dscp ef {
      fc "ef"
    }
    dot1p 3 {
      fc "l2"
    }
    dot1p 5 {
      fc "l1"
      priority low
    }
    policer 1 {
      rate {
        pir 20000
        cir 20000
      }
      scheduler-parent {
        scheduler-name "T1"
        level 2
        weight 50
      }
    }
    policer 2 {
      rate {
        pir 30000
        cir 10000
      }
      scheduler-parent {
        scheduler-name "T1"
        level 2
        weight 50
      }
    }
```

```
    policer 3 {
        rate {
          pir 40000
          cir 10000
        }
        scheduler-parent {
          scheduler-name "root"
          level 6
        }
      }
      fc "af" {
        policer 3
      }
      fc "l1" {
        policer 1
      }
      fc "l2" {
        policer 2
      }
    }
  policer-control-policy "CFHP" {
      root {
        max-rate 100000
      }
      tier 1 {
        arbiter "T1" {
          rate 30000
          arbiter-parent {
            arbiter-name "root"
            level 4
          }
        }
      }
```

    a. 10 000 Kbps
    b. 15 000 Kbps
    c. 20 000 Kbps
    d. 30 000 Kbps

17. Which of the following is NOT an advantage of deploying queue-groups on a network port?
    a. It enables the redirecting of traffic for VPN services to a separate set of queues.
    b. It allows the rate parameters of queues in the queue-group to be configured using absolute values.
    c. It improves scalability by allowing the use of queues instead of the default set of policers used on network egress.
    d. It allows the buffer size of queues in the queue-group to be configured using either percentage or kilobytes.

18. A queue-group template can contain both queues and policers when applied under which conditions?
    a. When applied at the port level for access ingress
    b. When applied at the FP (Forwarding Pass) level for access ingress
    c. When applied at the FP (Forwarding Pass) level for network ingress
    d. When applied at the port level for network egress

19. Consider the exhibit. According to the SAP-ingress policy configured and applied, which queue will forwarding class AF's traffic use?

```
[pr:/configure qos queue-group-templates ingress queue-group "Q-Group1"]
A:student@R01(PE1)# info
   policer 1 {
   }
   policer 2 {
   }
[pr:/configure qos sap-ingress "10"]
A:student@R01(PE1)# info
   default-fc "af"
   policer 1 {
   }
   fc "af" {
      policer 2
      fp-redirect-group-policer
   }
[[pr:/configure]
A:student@R01(PE1)# card 1 fp 1 ingress access queue-group "Q-Group1" instance-id 1
[pr:/configure]
A:student@R01(PE1)# service epipe 1 sap 1/1/4 ingress qos sap-ingress policy-name 10 fp-redirect-group group-name Q-Group1 instance 1
```

    a. Policer 2 on the queue-group Q-Group1 instance 1.
    b. Policer 1 on the queue-group Q-Group1 instance 1.
    c. The default policer 1.
    d. AF traffic is dropped.

20. Which of the following statements about scheduling on the Nokia 7750 SR is FALSE?
    a. When using priority-mode queues, traffic marked as low-priority on ingress and scheduled in the above-CIR loop will be considered as in-profile at egress.
    b. When using priority-mode queues, traffic marked as high-priority on ingress and scheduled in the above-CIR loop will be considered as out-of-profile at egress.
    c. When using profile-mode queues, traffic without a pre-specified profile on ingress and scheduled in the above-CIR loop will be considered as out-of-profile at egress.
    d. When using profile-mode queues, traffic marked as in-profile on ingress and scheduled in the above-CIR loop will still be considered as in-profile at egress.

21. Which of the following is NOT an adaptation rule supported on the Nokia 7750 SR for calculating the operational CIR/PIR of a queue?
    a. Closest
    b. Average
    c. Min
    d. Max

22. Which of the following statements about the burst-limit of a queue is FALSE?

a. The burst-limit sets a limit on the maximum burst of packets transmitted during each scheduler visitation time.

b. Proper configuration of the burst-limit can prevent packet drops on the downstream nodes resulting from traffic bursts.

c. The lower the configured burst- limit value, the lower the burst allowed to enter the queue.

d. If the burst-limit is not configured, a default burst-limit value will be used by the queue.

23. Which of the following statements about the default scheduler on the Nokia 7750 SR is FALSE?

a. The default scheduler spends an equal amount of time servicing each queue that has traffic to be sent out.

b. Queue-type can be configured as best-effort, expedite, or auto-expedite.

c. A queue will be serviced as expedited if its type is configured as expedited, regardless of the forwarding classes mapped to it.

d. A queue configured as auto-expedited will be an expedited queue as long as no Best Effort FC is mapped to it.

24. Consider the exhibit. A hierarchical scheduler policy with the shown parameters is configured at a service ingress. If the offered rate of each queue is 10Mbps, what is the operational PIR and CIR for queue 3?

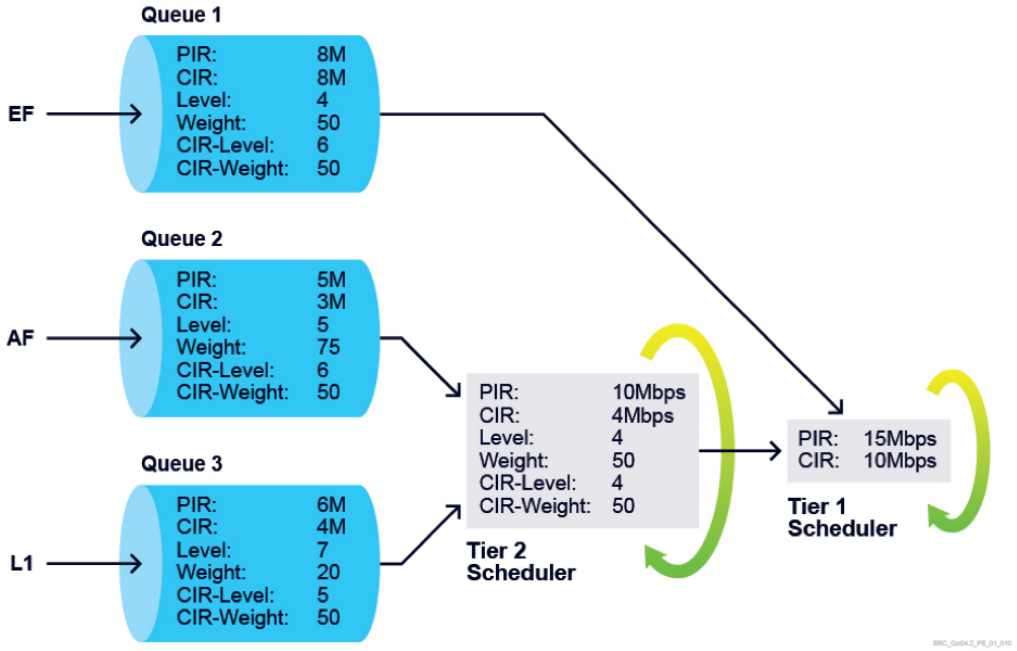

a. PIR = 5 Mbps, CIR = 0 Mbps

b. PIR = 4 Mbps, CIR = 4 Mbps

c. PIR = 6 Mbps, CIR = 4 Mbps

d. PIR = 6 Mbps, CIR = 1 Mbps

25. Which of the following statements describes the queue's weight parameter used in hierarchical scheduling?
    a. The weight specifies the priority of the queue for bandwidth allocation up to its CIR.
    b. The weight specifies the priority of the queue for bandwidth allocation above-CIR, up to its PIR.
    c. The weight is used to determine how this queue and other queues or schedulers at the same CIR-level share the available CIR bandwidth at that level.
    d. The weight is used to determine how this queue and other queues or schedulers at the same level share the available above-CIR bandwidth at that level.

26. Which of the following statements about configuring a hierarchical scheduler policy is FALSE?
    a. Several schedulers can be defined within a policy.
    b. A queue's parent scheduler must be a Tier-3 scheduler.
    c. The parent uses its children's CIR/PIR levels and weights to assign bandwidth.
    d. The fair share of within-CIR bandwidth for a queue (F-CIR) is calculated by that queue's parent.

27. While deciding how to share the available bandwidth among its children, when does a port scheduler consider the parameter "level" of a child?
    a. During the allocation of within-CIR bandwidth
    b. During the allocation of above-CIR bandwidth
    c. During the allocation of both within-CIR and above-CIR bandwidth
    d. The port scheduler does not use the level of a child for making bandwidth allocation decisions.

28. Consider the exhibit. Packets arrive at the VPLS SAP on PE1 with DSCP marking of af11 and Dot1p of 2. According to the QoS policies applied at routers PE1, P, and PE2, what is the DSCP value of packets egressing router PE2? (All the SDPs are MPLS-encapsulated, and all interfaces are using their default trust states).

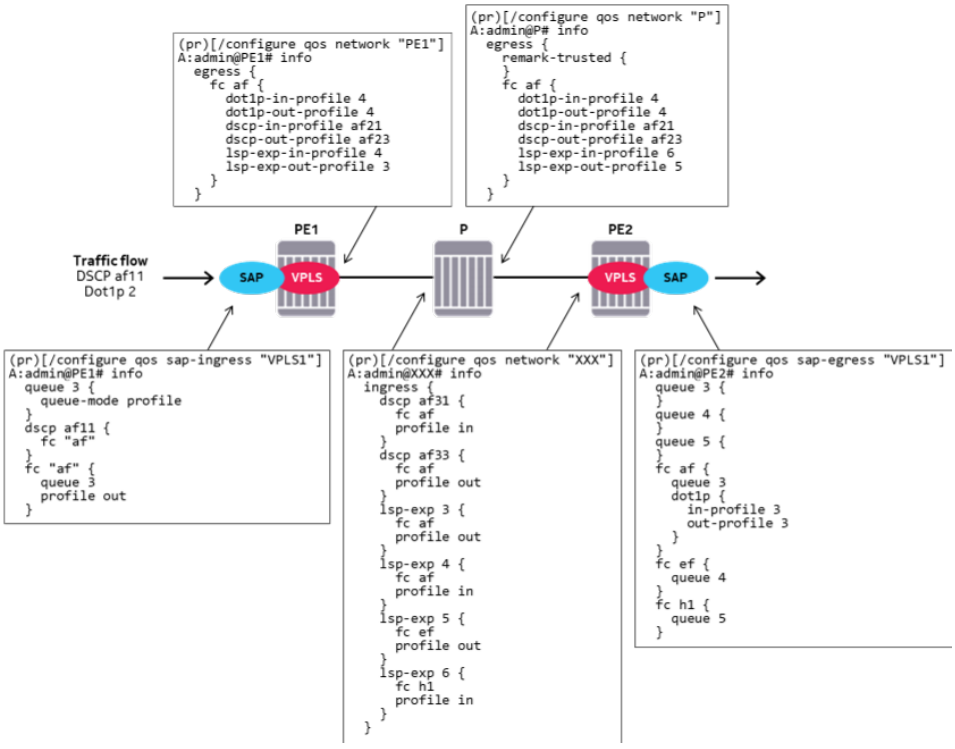

a. af11
b. af21
c. af23
d. af31

29. Which of the following types of delay is used to quantify the amount of time it takes a packet to egress one router and ingress the next-hop router?
a. Queuing
b. Processing
c. Serialization
d. Propagation

30. Which of the following statements defines the peak information rate for queues?
a. It is the rate at which a traffic stream is offered to a queue.
b. It is the rate at which packets are forwarded out of a queue.
c. It is the minimum rate at which packets may be forwarded out of a queue.
d. It is the maximum rate at which packets may be forwarded out of a queue.

31. Which of the following statements about the operation of a SAP-ingress QoS policy is FALSE?
    a. A SAP-ingress policy can use sub-classes to color traffic in or out of profile.
    b. An entry in the IP-criteria may identify the forwarding class and priority of the matching traffic.
    c. If no default forwarding class is configured, the forwarding class BE is the default.
    d. If using sub-classes to color the incoming traffic, the queue associated with the forwarding class must be configured as a priority queue.

32. Which of the following statements about the operation of token buckets is FALSE?
    a. A policer will drop non-compliant packets.
    b. The drain rate of a token bucket will vary with the ingress rate of the traffic.
    c. The maximum depth defines the size of the token bucket.
    d. A shaper will make non-compliant packets to continue waiting in the buffer.

33. A default scheduler will make two visitations to a queue that is configured with a PIR and a CIR. Which of the following statements is FALSE when referring to the scheduler's operation on the queue?
    a. If the traffic egressing the queue can put tokens in the CIR bucket, it is colored in-profile.
    b. The CIR parameter determines the drain rate used for the CIR bucket.
    c. When the CIR bucket gets full, the scheduler will stop its first visit to the queue.
    d. The burst limit of the queue defines the size of the CIR bucket.

34. Which of the following statements about the use of FIR with standalone queues is TRUE?
    a. The FIR bucket size is configured in kilobytes.
    b. The scheduler will visit best-effort queues within FIR before visiting expedited queues within CIR.
    c. The FIR rate cannot be configured to exceed the CIR rate.
    d. The FIR rate parameter can only be configured on SAP-ingress and SAP-egress queues.

35. Which of the following Class Fair Hierarchical Policing (CFHP) configuration descriptions is TRUE?
    a. The rates defined in the CFHP control policy can be either in kilobits per second (kbps) or port percentage.
    b. Both queues and policers in a SAP-ingress policy can be parented to arbiters in the same policer control policy.
    c. The tier 1 arbiter is also known as the "root" arbiter.
    d. Both the CFHP policer control policy and the SAP-ingress or SAP-egress policy need to be applied to the relevant SAP.

36. Which of the following statements about HQoS calculation of fair share of within-CIR bandwidth (F-CIR) is FALSE?
    a. The scheduler will calculate the fair share of bandwidth for each child of a given priority.
    b. The F-CIR calculation is based upon how much bandwidth is currently available for that specific priority level.
    c. The scheduler needs to consider the child's administrative CIR, CIR-level and CIR-weight.
    d. If sufficient bandwidth is available for the children, the F-CIR can be greater than the administrative CIR.

37. Which of the following descriptions about the configuration of an egress port scheduler is TRUE?
    a. If the egress port scheduler is to be applied to a network port, the rates need to be configured as a percentage.
    b. The egress port scheduler supports up to three levels of hierarchy.
    c. The egress port scheduler and the SAP-egress policy must be applied to the applicable SAP of the service.
    d. Both queues and policers can be parented back to the egress port scheduler.

38. Which of the following statements about utilizing an egress port scheduler and HQoS together is FALSE?
    a. The rates of the egress port scheduler can be configured in kilobits per second (kbps) or percentages.
    b. The SAP-egress queues and policers can be parented back to schedulers in the HQoS hierarchy, or directly to the egress port scheduler.
    c. Only the root scheduler in the HQoS hierarchy needs to be parented back to the egress port scheduler.
    d. The egress port scheduler, the HQoS scheduler and the SAP-egress policy must be applied to the applicable SAP of the service.

39. Which of the following statements about marking and remarking on the Nokia 7750 SR is FALSE?
    a. When a packet is encapsulated on the network egress of the first service provider router, the service provider adds QoS markings to the additional headers.
    b. Normally remarking of the data is not required if all of the routers are in the same QoS domain.
    c. The network interfaces can be configured as trusted or untrusted.
    d. Remarking on the network interface egress can only be done once the “remark-trusted” configuration is enabled.

40. When using the default network QoS policy on the Nokia 7750 SR, which two forwarding classes on the egress will be marked with the same MPLS EXP QoS information, merging the traffic into one forwarding class on the ingress of the next hop router?
    a. H1, EF
    b. EF, H2
    c. L1, AF
    d. AF, L2

Answer Key

1. A	11. C	21. B	31. D
2. C	12. C	22. C	32. B
3. B	13. D	23. A	33. D
4. A	14. B	24. B	34. B
5. D	15. C	25. D	35. D
6. B	16. B	26. B	36. D
7. C	17. C	27. B	37. D
8. A	18. D	28. A	38. D
9. B	19. A	29. D	39. D
10. D	20. A	30. D	40. C

About Nokia

At Nokia, we create technology that helps the world act together.

As a B2B technology innovation leader, we are pioneering networks that sense, think and act by leveraging our work across mobile, fixed and cloud networks. In addition, we create value with intellectual property and long-term research, led by the award-winning Nokia Bell Labs.

Service providers, enterprises and partners worldwide trust Nokia to deliver secure, reliable and sustainable networks today – and work with us to create the digital services and applications of the future.

Nokia OYJ
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Document code: (June) CID202017