CERTLABZ CERTIFICATION BOARD

Certified IPv6 Deployment Specialist (CIDS)

Ten Sample Questions

These ten questions are written at the difficulty of the live exam. Six carry a figure, a data table, a configuration or a code listing; four are plain scenario questions, as on the live paper. Every one is a scenario a working network engineer would meet, and none of them appears in the live question bank.

The live exam is multistage adaptive: how you score on one stage selects the difficulty of the next, so no two candidates sit the same paper. It is a performance-based exam: each form can carry up to 5 to 12 performance-based questions alongside the multiple-choice questions, which means you build, sequence or complete something rather than only choosing from a list. Answers and full explanations follow each question here; the live exam does not disclose item-level answers.

Question 1

Domain 1: IPv6 Foundations and Addressing

CODE

Emma Wilson must write the address 2001:0db8:0000:0000:0000:0000:0000:0001 in its shortest valid form.

```
full : 2001:0db8:0000:0000:0000:0000:0000:0001
rule : drop leading zeros; compress one run of zero groups with ::
```

Listing 1.1 The address and the compression rules.

Which is the correct shortest form?

A. 2001:db8:0:1, dropping the middle zero groups and writing the remaining groups separated by single colons only.

B. 2001:db8::1, dropping the leading zeros and compressing the run of zero groups with a single double colon.

C. 2001::db8::1, using a double colon both before and after the db8 group to compress the surrounding zeros efficiently.

D. 2001.0db8::1, keeping the leading zeros on the first group and separating the site prefix with a dot for clarity.

Correct answer: B

2001:db8::1 drops the leading zeros in each group and compresses the single run of all-zero groups with one ::. A double colon may appear only once, so 2001::db8::1 is invalid, 2001:db8:0:1 omits groups incorrectly, and IPv6 never uses dots between hex groups.

Question 2 — Domain 1: IPv6 Foundations and Addressing — DATA

James Carter classifies four IPv6 addresses.

Address / range	Type
fe80::/10	(to identify)
2000::/3	(to identify)
fc00::/7	(to identify)
::1	(to identify)

Table 2.1 Addresses to classify.

Which classification is correct?

A. fe80::/10 global unicast, 2000::/3 loopback, fc00::/7 link-local, ::1 unique local, matching each range to a type.

B. fe80::/10 link-local, 2000::/3 global unicast, fc00::/7 unique local, ::1 loopback, matching each range to its type.

C. fe80::/10 loopback, 2000::/3 unique local, fc00::/7 global unicast, ::1 link-local, assigning each a distinct type.

D. fe80::/10 unique local, 2000::/3 link-local, fc00::/7 loopback, ::1 global unicast, giving each range a type.

Correct answer: B
fe80::/10 is link-local, 2000::/3 is global unicast, fc00::/7 is unique local (private), and ::1 is the loopback. The other options swap these types around.

Question 3 — Domain 2: Subnetting and Address Planning — DATA

Olivia Davis plans addressing for a site allocated a /48.

Question	Answer needed
Prefix per LAN	(to decide)
Subnets available in the /48	(to decide)
Planning priority	(to decide)

Table 3.1 Planning decisions.

Which set of decisions is correct?

A. Use /112 per LAN, giving 256 subnets in the /48, and prioritise conserving addresses because IPv6 space is scarce and limited.

B. Use /64 per LAN, giving 65,536 subnets in the /48, and prioritise a clean, aggregatable structure over conserving addresses.

C. Use /64 per LAN, giving 256 subnets in the /48, and prioritise squeezing subnets as tightly as possible to save address space.

D Use one /64 for the whole site, giving a single subnet, and prioritise using the smallest possible allocation from the provider.

Correct answer: B
The standard is /64 per LAN, and the 16 bits between /48 and /64 give 65,536 subnets. IPv6 planning prioritises clean, aggregatable structure, not conservation, because the address space is vast. Non-/64 LANs break SLAAC, the subnet count in a /48 is 65,536 not 256, and one /64 for a whole site wastes the hierarchy a /48 provides.

Question 4

Domain 3: Neighbor Discovery, SLAAC and DHCPv6

CONFIG

Liam Anderson finds that hosts on a segment have only link-local addresses.

```
host address : fe80::... only (no global address)
router       : interface up, IPv6 enabled
router config : no Router Advertisements being sent
DHCPv6       : not in use
```

Listing 4.1 Observed state on the segment.

Why do the hosts lack a global address, and what is the fix?

A DNS is down, so the hosts cannot resolve the router, and restoring DNS will let SLAAC assign the global addresses again.

B No Router Advertisements are being sent, so SLAAC cannot proceed; enable RAs (advertising the prefix) on the router.

C The switch is blocking IPv4, so the hosts fall back to link-local only, and re-enabling IPv4 will restore global IPv6 addressing.

D The hosts need static addresses, because SLAAC cannot assign a global address without a DHCPv6 server present on the segment.

Correct answer: B
SLAAC builds a global address from the prefix in a Router Advertisement. With no RAs being sent, hosts can only form link-local addresses, so the fix is to enable RAs advertising the prefix on the router. DNS, IPv4 and a DHCPv6 server are not required for SLAAC to assign a global address.

Question 5

Domain 4: Routing and Transition Mechanisms

CONFIG

Ava Thompson debugs IPv6 where small requests work but large transfers stall.

```
small packets : OK over IPv6
large transfers : stall / time out
firewall rule : drop all ICMPv6, including "packet too big"
```

Listing 5.1 Symptoms and a firewall rule.

What is happening, and what is the fix?

A. The route to the destination is missing, so large packets are dropped, and adding a default IPv6 route will restore the transfers.

B. Path MTU Discovery is broken because ICMPv6 "packet too big" is filtered, so permit that ICMPv6 type to fix the MTU black hole.

C. DNS is intermittently failing on large lookups, so the transfers stall, and pointing hosts at a different resolver resolves the issue.

D. IPv4 is interfering with the IPv6 transfers, so disabling IPv4 on the hosts will stop the large IPv6 transfers from stalling.

Correct answer: B
IPv6 routers do not fragment; the source relies on Path MTU Discovery, which uses ICMPv6 "packet too big" messages. Dropping them creates an MTU black hole where small packets pass but large ones vanish. Permitting that ICMPv6 type restores PMTUD. A missing route, DNS or IPv4 would not produce this size-dependent pattern.

Question 6

Domain 5: DNS, Services and Applications

DATA

A dual-stack website loads slowly for some users after IPv6 was enabled.

Fact	Detail
A record	Correct, IPv4 path works
AAAA record	Published
IPv6 path	Broken end to end
Symptom	Slow load, then it works

Table 6.1 Observations for the dual-stack site.

What is the cause, and what is the correct fix?

A. The A record is wrong, so clients fail over to IPv6 slowly, and correcting the IPv4 address will remove the delay for all users.

B. Clients try the broken IPv6 path first and stall before falling back to IPv4; fix the IPv6 path, or remove the AAAA until IPv6 works.

C. DNS caching is the cause, so lowering the record time-to-live will stop the slow loads without any change to the IPv6 path itself.

D. The site needs a second AAAA record for redundancy, and publishing an additional IPv6 address will resolve the slow initial load.

Correct answer: B
A valid AAAA record with a broken IPv6 path makes clients attempt IPv6 first and stall until they fall back to IPv4, which is the slow-then-works symptom. The fix is to repair the IPv6 path end to end, or withdraw the AAAA record until IPv6 works. The A record is fine, and caching or an extra AAAA does not address a broken path.

Question 7

Domain 2: Subnetting and Address Planning

A team assigns /112 prefixes to its LANs to "save address space". What will this break, and what is the correct approach?

A. Nothing breaks, because a /112 provides plenty of addresses per LAN and conserving IPv6 space is a sensible design priority.

B. SLAAC and standard host autoconfiguration break, because they assume a 64-bit interface identifier; use /64 per LAN as standard.

C. Routing stops entirely, because IPv6 routers cannot forward any prefix longer than /64, so all inter-LAN traffic will fail.

D. Only reverse DNS breaks, and the fix is to keep the /112 prefixes while delegating the reverse zone at a /64 boundary instead.

Correct answer: B
SLAAC and several IPv6 features assume a /64 with a 64-bit interface identifier, so a /112 LAN breaks stateless autoconfiguration. IPv6 address space is vast, so there is no need to conserve it this way; the correct approach is /64 per LAN. Routing does not stop at /64, and the problem is broader than reverse DNS.

Question 8 Domain 4: Routing and Transition Mechanisms

An organisation is moving from IPv4 to IPv6 and wants a low-risk transition. Why is dual stack usually preferred over relying on tunnels long term?

A. Tunnels are always faster than native IPv6, so the organisation should build its entire long-term network design around permanent tunnelling.

B. Native dual stack runs both protocols independently, without tunnelling overhead, MTU issues and extra failure points, so it is more reliable.

C. Dual stack means running IPv6 only, which removes IPv4 immediately and is therefore the lowest-risk way to complete the transition.

D. Tunnels require no maintenance and never fail, so long-term reliance on them is the recommended end state for an IPv6 transition.

Correct answer: B
Dual stack serves both protocols natively, avoiding the overhead, MTU problems and additional failure points that tunnels introduce, which makes it simpler and more reliable for a gradual transition. Tunnels are a temporary bridge, not a destination; dual stack does not remove IPv4 instantly, and tunnels do need maintenance and can fail.

Question 9 Domain 6: IPv6 Security and Deployment

An enterprise enables IPv6 on its hosts, but its firewalls and intrusion detection still inspect only IPv4. What is the danger, and what is required?

A. There is no danger, because the existing IPv4 firewall and intrusion detection rules automatically apply to IPv6 traffic as well.

B. IPv6 traffic passes uninspected, so attacks over IPv6 evade the controls; security parity across both protocols is required before exposure.

C The only danger is slower performance, because inspecting two protocols is slower, and the fix is to disable IPv6 inspection to keep speed.

D IPv6 cannot be attacked, so inspecting only IPv4 is sufficient, and the enterprise simply needs to document that IPv6 is exempt from review.

Correct answer: B
Enabling IPv6 without extending firewalls and intrusion detection to it leaves an uninspected parallel path that attacks can use to evade controls. The requirement is security parity: IPv6 must be firewalled, monitored and inspected to the same standard as IPv4. IPv4 rules do not automatically cover IPv6, and IPv6 is certainly attackable.

Question 10 Domain 3: Neighbor Discovery, SLAAC and DHCPv6

An administrator blocks all ICMPv6 at the host firewall to "harden" it, and IPv6 connectivity becomes unreliable. Why?

A ICMPv6 carries the actual user data in IPv6, so blocking it drops the payload and that is what makes connectivity unreliable.

B Blocking all ICMPv6 breaks Neighbor Discovery and Path MTU Discovery, which IPv6 depends on, so essential ICMPv6 types must be permitted.

C Blocking ICMPv6 disables DNS resolution over IPv6, and re-enabling only DNS traffic through the firewall will fully restore connectivity.

D ICMPv6 is unrelated to IPv6 operation, so the unreliability has another cause and the ICMPv6 block can safely be left exactly as it is.

Correct answer: B
IPv6 relies on ICMPv6 for core functions such as Neighbor Discovery (address resolution, router and prefix discovery) and Path MTU Discovery, so a blanket block breaks the protocol. Filtering must be selective, permitting the essential ICMPv6 types. ICMPv6 does not carry user data, and the cause is the block itself, not DNS or an unrelated issue.

Published by the CertLabz Certification Board, a division of CertLabz LLC. Companion to objectives document CIDS-001. https://certlabz.com/certifications/certified-ipv6-deployment-specialist.html