

Lecture 14 – DNS Quiz ANS

1. What is the primary purpose of the Domain Name System (DNS)?

- A. To route packets between autonomous systems
- B. To translate human-readable domain names into IP addresses
- C. To encrypt application-layer traffic
- D. To allocate IP address blocks to ISPs

ANS:

2. Why are IP addresses considered useful for machines but not for humans?

- A. They are always changing over time
- B. They are meaningful phrases that are easy to remember
- C. They help routing scale but are hard for humans to remember
- D. They contain information about the geographic location of the user

ANS:

3. Which of the following is a key scalability requirement for DNS mentioned in the lecture?

- A. DNS must encrypt all traffic by default
- B. DNS must handle many hosts, lookups, and updates efficiently
- C. DNS must operate only within local area networks
- D. DNS must store routing tables for all domains

ANS:

4. Which property is most closely associated with DNS being "highly available"?

- A. Having a single authoritative root name server
- B. Eliminating caching in resolvers
- C. Avoiding any single point of failure
- D. Using only TCP for DNS queries

ANS:

5. DNS is described as "lightweight and fast" primarily because:

- A. It is implemented using a link-layer protocol
- B. It uses UDP by default and typically completes in a single packet exchange
- C. It always uses encrypted transport
- D. It maintains long-lived TCP connections for every client

ANS:

6. What is a DNS name server?

- A. A host that only caches HTTP responses
- B. A server responsible for answering DNS requests for some set of domains
- C. Any router participating in BGP

D. A server that assigns IP addresses using DHCP

ANS:

7. In the DNS hierarchy, which servers are contacted first when resolving a name like cs.hofstra.edu?

- A. The .edu name servers
- B. The hofstra.edu name servers
- C. The root name servers
- D. The local web server for cs.hofstra.edu

ANS:

8. When a name server does not know the answer to a DNS query, what can it do according to the lecture?

- A. Drop the query without responding
- B. Always return a NXDOMAIN error
- C. Direct the client to another name server that has delegated authority
- D. Broadcast the query to all other name servers

ANS:

9. What is the role of caching in DNS, as illustrated by the cs.hofstra.edu example?

- A. To permanently store all DNS responses on the client
- B. To avoid repeatedly querying name servers for the same name within a time-to-live period
- C. To eliminate the need for recursive resolvers
- D. To ensure that only the root server is contacted for each query

ANS:

10. In practice, what is the typical relationship between a stub resolver and a recursive resolver?

- A. The stub resolver performs all iterative queries itself
- B. The recursive resolver only caches results from a single user
- C. The stub resolver sends a query to a recursive resolver, which performs the full lookup
- D. The recursive resolver runs on every end host

ANS:

11. Why are recursive resolvers often operated by ISPs or application providers like Google or Cloudflare?

- A. They must be physically located at every end host
- B. They benefit from aggregating many users' queries, improving cache hit rates and reducing server load
- C. They are required to maintain BGP routing tables
- D. They must provide DHCP services to clients

ANS:

12. Which transport protocol does DNS primarily use and why?

- A. TCP, because it provides reliable, ordered delivery for all lookups
- B. UDP, because it is lightweight and avoids connection setup, making lookups faster
- C. ICMP, because it is reserved for diagnostic messages
- D. SCTP, because it supports multi-homing

ANS:

13. Under what circumstances might DNS use TCP instead of UDP?

- A. When the client explicitly requests encryption for queries
- B. When messages are too large to fit in a single UDP packet or during large zone transfers
- C. When resolving local hostnames only
- D. When the resolver is behind a NAT

ANS:

14. In the DNS hierarchy, what is a "zone"?

- A. A set of IP addresses assigned to a single router
- B. The portion of the domain namespace under the control of a specific administrative authority
- C. A physical location where DNS servers must be deployed
- D. A backup copy of all DNS records worldwide

ANS:

15. Which organization is responsible for controlling the DNS root?

- A. IETF (Internet Engineering Task Force)
- B. IEEE (Institute of Electrical and Electronics Engineers)
- C. ICANN (Internet Corporation for Assigned Names and Numbers)
- D. ITU (International Telecommunication Union)

ANS:

16. What are top-level domains (TLDs)?

- A. Domains that are reachable only within a local network
- B. The zones directly below the root, such as .com, .org, or country codes like .uk
- C. Subdomains created by individual users
- D. Domains reserved exclusively for government use

ANS:

17. How is high availability typically achieved for a DNS zone?

- A. By hosting a single powerful primary server
- B. By having multiple authoritative name servers in a primary/secondary model
- C. By disabling caching across all resolvers

D. By forcing all queries to go through the root server

ANS:

18. What is anycast and how is it used for root DNS servers?

- A. A protocol for encrypting all DNS traffic between clients and servers
- B. A routing technique where many servers share the same IP address and the network routes clients to a nearby instance
- C. A method of compressing DNS messages to reduce overhead
- D. A way to tunnel DNS over HTTP for better performance

ANS:

19. How can DNS be used for load balancing?

- A. By having clients randomly generate destination IP addresses
- B. By returning multiple IP addresses for the same domain and allowing clients to choose among them
- C. By forcing all traffic through a single central proxy server
- D. By using DHCP to assign multiple IPs to the same host

ANS:

20. Which of the following is a challenge with geographical load balancing using DNS?

- A. DNS cannot return more than one IP address per domain
- B. It is impossible to map client IPs to any geographic information
- C. The name server often does not know the true network or geographic distance or performance between client and server
- D. Clients always ignore DNS responses from authoritative servers

ANS: