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: B – DNS provides a mapping from human-readable domain names to IP addresses so hosts can communicate.

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: C – IP addresses are structured for scalable routing, but numeric strings are not meaningful or easily memorable for humans.

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: B – DNS must scale to many hosts, frequent lookups, and continuous updates across the Internet.

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: C – High availability requires DNS to avoid single points of failure so lookups continue despite failures.

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: B – DNS typically uses UDP, avoiding connection setup and allowing most queries and responses to fit in single packets.

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: B – A DNS name server is responsible for answering DNS queries for the domains under its authority.

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: C – DNS lookups begin at the root name servers, which then direct queries down the hierarchy.

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: C – DNS servers can return referrals pointing clients to other name servers that have authority over the requested name.

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: B – Caching stores DNS responses for a TTL so subsequent queries can be answered without repeating the full lookup.

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: C – The stub resolver on the client forwards queries to a recursive resolver, which contacts the necessary name servers and returns the final answer.

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: B – Operating recursive resolvers centrally allows providers to build large shared caches and reduce load on authoritative servers.

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: B – DNS typically uses UDP to minimize overhead and keep name lookups fast and lightweight.

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: B – Large responses and zone transfers can exceed UDP limits, so DNS falls back to TCP in these cases.

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: B – A zone represents the part of the DNS namespace managed by one administrative authority, which can delegate subzones to others.

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: C – ICANN manages the DNS root zone and coordinates top-level domains.

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: B – TLDs are the zones immediately below the root, including generic and country-code domains.

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: B – Zones commonly have multiple authoritative servers, with secondary servers replicating data from a primary for redundancy.

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: B – Anycast allows multiple physical servers to advertise the same IP address, so clients are routed to a nearby instance, improving availability and performance.

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: B – DNS servers can return several A records for one domain; clients choose one, distributing load across servers.

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: C – DNS-based geo-load balancing must infer proximity and performance from limited information, making accurate mapping difficult.