

Solutions

Each answer shows the steps, then the **result**.

2.1

- `https://` is the **protocol**, which identifies the rules used to communicate with the server
- `www.exercises.lol` is the **domain name**, which identifies the **server** holding the resource
- `/library/physics.pdf` is the **path**, which identifies **which resource** on that server

2.2

- DNS holds the mapping from **domain name to IP address**, distributed across many name servers
- so that a browser given a name can obtain the address the data must actually be sent to

2.3

- **IPv4: 32 bits**, written as four denary numbers in the range 0-255, separated by dots
- **IPv6: 128 bits**, written as eight groups of four hexadecimal digits, separated by colons

2.4

- IPv4 has only about 4.3 billion addresses (2^{32}), and the number of devices connected to the internet has exhausted them
- IPv6's 128-bit addresses give an address space large enough that running out is not a practical concern

2.5

- a private address is not reachable directly from the internet, because it is unique only inside its own network
- so the device cannot be contacted – or attacked – directly from outside; an attacker must get past the router first

2.6

- an address that is issued to the device when it connects to the network, and which may be a different address the next time it connects

3.1 C

- each part of an IPv4 address must be in the range 0-255, so **A** is invalid
- **B** is an IPv6 address and **D** is a domain name, not an address at all

3.2

(a)

- the browser takes the **domain name** out of the URL and sends it to a **DNS name server**
- the name server looks up the domain name and finds the **IP address** it maps to, asking other name servers if it does not hold the record itself
- the address is returned to the browser, which then sends its request for the resource to that address

(b)

- a cached result answers the next request for the same name immediately, so the page loads faster and no query has to leave the machine
- and it removes an enormous amount of repeated traffic from the name servers, which could not otherwise handle every lookup for every page

(c)

- it passes the query on to another name server, working up the hierarchy – root, then top-level domain, then the **authoritative** server for that domain
- the address found there is returned back down the chain to the browser, and cached along the way

3.3

(a)

- a **public** IP address is unique across the whole internet and can be reached from it
- a **private** IP address is unique only within one network, and a router translates between the private addresses inside and the public address outside
- the 40 devices have **private** addresses; the router holds the single public one

(b)

- none of the 40 devices can be contacted directly from the internet, since their addresses have no meaning outside the company's own network
- so an attacker on the internet cannot reach a device without first getting through the router, which greatly reduces the exposed surface

(c)

- a **static** address is fixed and does not change
- a **dynamic** address is assigned when the device connects, typically by a DHCP server, and may differ next time
- a static address is necessary for anything that must be found at a known address – a web, mail or file **server**, because DNS records point at a fixed address

(d)

- subnetting divides one network into smaller **subnets**, using part of the host portion of the address as a subnet number
- one reason: traffic meant for a device in the same subnet need not be carried by the others, so congestion falls; or each subnet (for example the office and the guest wi-fi)

can be secured and managed separately