Topic 9 Fundamentals of Communication and Networking

Symbol

The term symbol refers to a pattern of bits represented by a signal. For example, a four bit symbol could be 0011.

Baud Rate

The baud rate is the median number of signal changes per second. 1 symbol change per second is 1 Baud, written as 1Bd.

Bit Rate

The bit rate is the number of bits which can be transmitted each second measured in bps (bits per second). The bitrate is the baud rate multiplied by the number of bits per signal.

Bandwidth

Bandwidth is the range of frequencies which can be transmitted and is measured in Hertz (Hz). Higher bandwidth gives a higher bit rate.

Latency

Latency is the difference between an action being initiated and its effect being noticed and is measured in milliseconds (ms). It can be best thought of as a delay, and usually increases with the distance data needs to travel.

Protocol

A protocol is a set of rules regarding communication between devices. This allows devices made by different manufacturers to easily communicate with one another.

Parallel and Serial Data Transmission

Serial Data Transmission

Serial data transmission sends data one bit at a time using a communication channel such as a copper cable or wireless signal. This method is often used for relative medium to long distance transmissions such as from a mouse to the computer.

Serial transmission does not suffer from crosstalk or skew, making it more reliable over longer distances. It is also cheaper than parallel data transmission since it uses a single line.

Parallel Data Transmission

This method uses many parallel communication lines to send many bits at the same time and is usually used between components within the computer. The more lines available, the more data that can be sent at the same time.

Each line within a parallel communication system will have very slightly different electrical properties, meaning that signals will take a very slightly different amount of time to travel across them. As a result, not all data sent at the same time will arrive at the same time, this problem is known as skew and increases as with the distance between components. In extreme cases this can lead to data overlapping or corruption.

Where parallel communication lines are close to each other crosstalk can occur, where data from one line leaks into or interferes with another, causing data corruption.

Parallel communication methods are more expensive because they use multiple lines, and because of this they are usually used over relatively short distances between internal computer components.

Synchronous and Asynchronous Data Transmission

Synchronous Data Transmission

Synchronous transmission uses a clock signal shared between the sender and receiver to control when data is sent. It is used within processor busses in the fetch-execute cycle. The clock signals are received in the same order in which they were sent, making them well suited to use in real time systems.

Asynchronous Data Transmission

This method removes the need for a clock, and instead uses start and stop bits to indicate transmission duration. The start bit can be either 0 or 1 and the stop bit must be the opposite of the start bit. The sender and receiver have to use the same baud rate but only need to synchronise their clocks for the duration of the transmission.

1	0	1	0	0	1	0	0	1	0	0
Stop Bit	Transmitted Data									Start Bit

Network Topology (Physical)

The physical network topology refers to how the devices within the network are physically laid out and connected.

Star Network Topology

In this topology, each client has its own direct connection to a central hub. The hub receives packets for all clients and delivers them to the correct place. A server is connected in the same way as any other client.

Advantages:

- Packets are sent directly to their recipients and are not seen by other clients.
- It is easy to add or remove clients.
- Each cable has one device, meaning there are no collisions.
- The failure of one cable does not affect the rest of the network.

Disadvantages:

- Expensive to install due to the amount of cable required.
- If the central hub fails, the entire network fails.

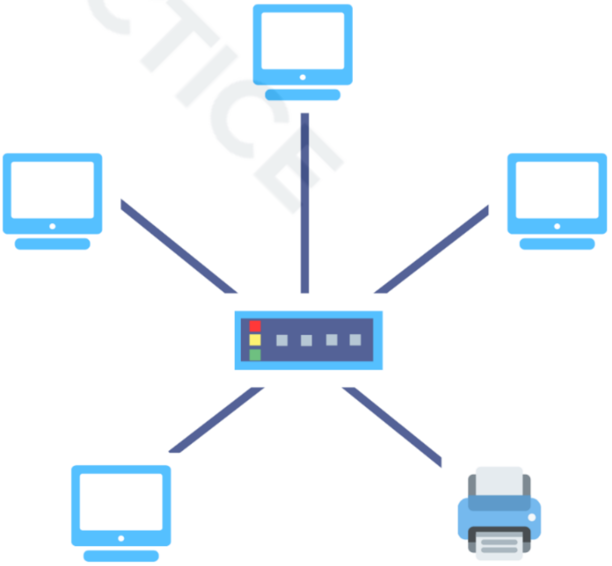

Bus Topology

A central bus connects clients along a single cable, known as a backbone. A terminator is placed at each end of the backbone. There is on central hub and servers can be connected anywhere on the backbone.

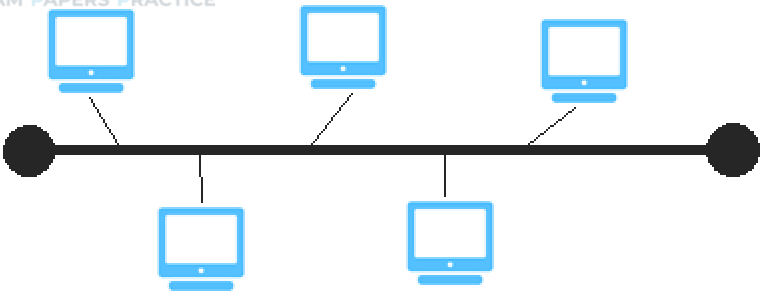

Advantages:

- There is no central hub, making the setup cheaper to install and less likely to fail.
- Less cabling is required making it cheaper to install.

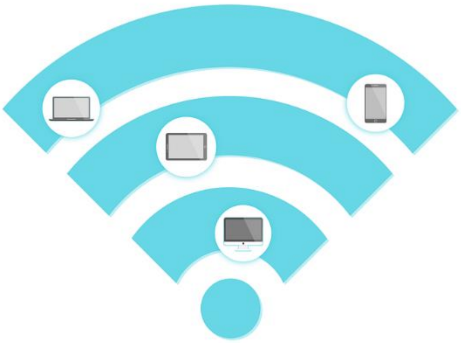

Disadvantages:

- Packets sent on the shared backbone can be seen by all clients.
- The backbone is used by all clients, introducing a risk of collisions.
- If the backbone fails, the whole network cannot be used

Logical Network Topology

The logical topology refers to the flow of data packets through the network. A logical bus network delivers packets to all clients on the network. A logical star network delivers packets to only the intended recipient.

Mixed Topologies

Physical and logical topologies can be mixed, for example, it is possible to setup a network in a physical star topology but to behave as a logical bus topology. This could be achieved by using a physical star layout when cabling but operating a bus protocol.

Hosts & Clients

A host is a device on a network which provides services to other systems. This is often a server and could provide file or print services. A client is a device which accesses services from another host on the network. Systems on a network can be both clients and hosts at the same time.

Client-Server Networking

Client-server networks use a central server to provide services to clients on the network. The server is connected to the network in the same way as the client but is usually a more powerful computer.

Clients request services from the server, which responds with the required services. The server could provide file storage, user accounts, print queue or email services. Sometimes a single server provides all services, whilst other times multiple servers are used.

This type of network is commonly used by schools and businesses to allow central management of clients. This improves security but requires additional expertise to set up and manage the network.

Peer to Peer Networking

This setup does not require a central server and instead services are provided by one or more of the clients themselves, with every client having equal status. One client might provide file services, whilst another manages storage for example.

The main disadvantage here is that a client must always be switched on and working in order to provide services, and if a client fails or is switched off the network is not fully operational. On the other hand, it is more cost effective as it does not require a powerful central server to run the network. It is also easier to setup and manage than client-server networks.

This arrangement is used by large file and multimedia sharing networks to provide high performance services without the need for a powerful server.

Wireless Networks

A wireless network allows communication without requiring a physical cable. They require a wireless access point and a wireless network adapter in the client. The wireless access point acts as a bridge between the wired network and the wireless.

WiFi

WiFi is the most commonly used method of providing wireless networks and refers to a wireless network based on international standards.

Traffic travelling over a wireless network should be encrypted using WPA or WPA2 to improve security. WPA stands for WiFi Protected Access and requires clients to enter a key before connecting to the network.

Wireless network security can also be improved by disabling SSID broadcast. The SSID (Service Set IDentifier) is the name used to identify the wireless network. Disabling SSID broadcast prevents wireless clients from seeing the network is there, allowing only clients who know of its existence and the SSID to connect.

MAC Address Filtering uses the MAC (Media Access Control) address built into every network card to control access to the network. A MAC whitelist allows only approved devices to connect to the network and prevents all other devices from connecting. On the other hand, a blacklist allows all devices to connect apart from a set of blocked addresses.

Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA)

This protocol is commonly used in wireless networks to avoid data collisions if multiple devices transmit at the same time.

When a device wants to transmit data, it listens to the communication channel to see if it is idle. If it is, the data is transmitted. If the channel is already in use, the device waits a random length of time and tries again. An exponential back off algorithm is used to increase the time period the device waits with each retry.
CSMA/CA suffers from a problem known as hidden nodes. This is where a device cannot see parts of the network where communication might be taking place. To get around this, another protocol called request to send/clear to send (RTS/CTS) is used. This adds an additional step into CSMA which requires the client to send a request to send to the server once it has detected the transmission media is clear.
If the server is idle, it responds with a clear to send message and the client sends its data. If the client does not receive a clear to send message it assumes the server is busy and waits before trying again.

The Internet Structure

The Internet is a global network built from many interconnected networks in different countries. The majority of Internet connectivity uses wired connections, with cables passing under oceans to connect different continents.

An ISP (Internet Service Provider) is a company or other organisation which provides access to the Internet to its customers. National internet service providers are larger ISPs who operate vast networks within a particular country, they provide access to regional and local ISPs who in turn provide access to homes and businesses.

Packet Switching

Packets are container which store data being transmitted over a network, and are labelled with both a sender and recipient address, with data being split into packets before being sent. Networks which operate in this way are called packet switched networks. Different packets may take different routes through the network before being reassembled by the recipient.

Sender: 11.55.6.2
Recipient: 22.11.5.6
Good,
TTL: 7
1 of 2

Sender: 11.55.6.2
Recipient: 22.11.5.6
Morning.
TTL: 7
2 of 2

Sender address	The address which sent the packet and to which any replies should be sent.
Recipient address	The destination address for the packet, allowing it to be correctly routed through the network.
Data	The data itself
TTL (Time to Live)	The number of routers a packet can pass through before it will be deleted.
Sequence number	The number of this particular packet along with the total number of packets. This allows the recipient to reassemble the packets in the correct order.

Routers and Gateways

A key requirement of The Internet is to connect different networks, allowing data to be sent between them. Routers and Gateways are network devices which allow two different networks to communicate.

Routers allow two networks which use the same protocol to communicate with one another. They hold tables containing information about the fastest route between different parts of the networks, and use these tables to route packets via the fastest possible route. These routing tables are updated frequently to ensure they always contain accurate information.

Gateways allow two networks to communicate even if they use different protocols. When the gateway receives a packet it removes the sender address, recipient address and other details, before adding these back in a way which conforms to the protocol of the destination network.

URLs (Uniform Resource Locators)

URLS (short for Uniform Resource Locators) are addresses used to access files via The Internet. A URL contains several different pieces of information as shown in the example below.

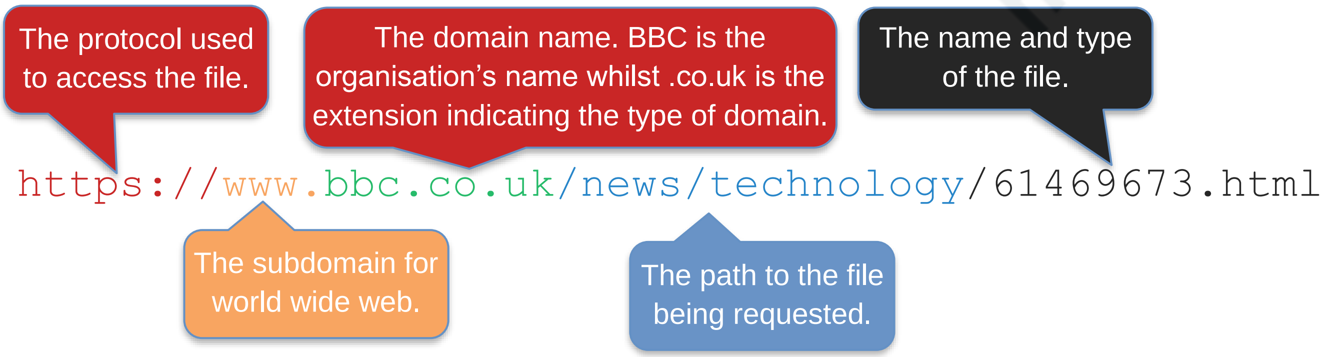
The protocol used to access the file.
The domain name. BBC is the organisation's name whilst .co.uk is the extension indicating the type of domain.
The name and type of the file.
https://www.bbc.co.uk/news/technology/61469673.html
The subdomain for world wide web.
The path to the file being requested.

Domain Names

A domain name identifies a resource accessible via The Internet using human readable words, letters and numbers. As shown in the example above, the domain name is made of both the organisation's name and a type such as .co.uk for UK Companies or .ac.uk for UK academic institutions.

Fully Qualified Domain Names (FQDNs)

FQDNs, short for Fully Qualified Domain Names, must include an exact resource and can be interpreted in only one way. The FQDL must include the full server host name.

The Domain Name Server System (DNS)

Every computer on a network is assigned and IP Address, these are often long and difficult for humans to remember and so domain names are used instead. A domain name maps an IP address to a human friendly, easy to remember name.

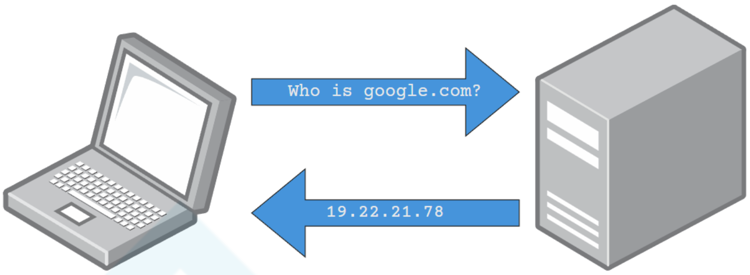

Domain	IP Address
bbc.co.uk	212.22.1.2
google.com	19.22.21.78
yahoo.com	98.215.22.21

A Domain Name Server is used to translate domain names to their corresponding IP address using a stored database of domain names and the IP addresses they map to. This allows a friendly domain name such as google.com to be entered into a web browser then mapped to the appropriate IP address, such as 19.21.22.78. If the domain name server does not have a record for the domain name in its own database, it will pass the request onto another domain name server.

Internet Registries

Internet registry organisations manage the allocation of IP addresses available for use on The Internet. Their role is becoming increasingly important as the number of available Internet IPv4 addresses is depleting.

Firewalls

A firewall is a network device which site between two networks, it controls which packets are allowed to pass between the networks and which are blocked, also known as dropped.

Packet Filtering

Packet filtering is a technique used by firewalls to decide which packets to allow and which to blocked based on the source and/or destination address of the packets. This allows packet sent from or to either specific devices, or certain areas of the network to be allowed or blocked (sometimes called deopped).

Stateful Inspection

This technique looks at the content of a packet, rather than just the sender and recipient address when deciding whether to allow the packet. This allows the firewall to provide more accurate and detailed filtering than by looking at just the sender and recipient address. Some firewalls will also look at active connections in order to block packets not related to current network activity.

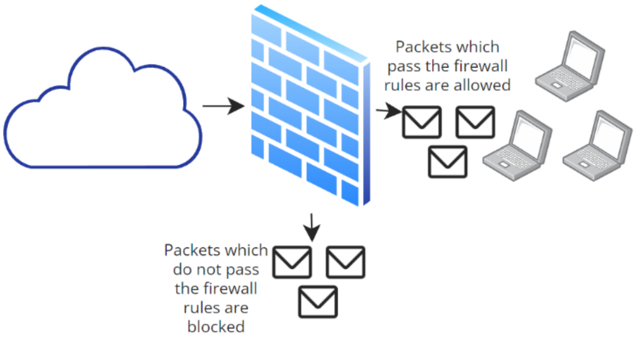
Packets which pass the firewall rules are allowed
Packets which do not pass the firewall rules are blocked

Proxy Servers

A proxy server is located between client devices and the firewall and hides the IP addresses of individual clients to improve privacy. The proxy server can also be configured to keep a copy of frequently accessed web pages, known as a cache. When a client attempts to access a website held in the cache, the proxy server will return its own copy of the page, reducing traffic and giving a faster response.

Proxy servers can be set up to keep a log of websites which are visited by clients, and to prevent access to some websites all together. This feature is commonly used by schools and employers to track web content being visited by their staff and pupils and to block access to websites which might be harmful.

The anonymity provided by proxy servers can be taken advantage of by someone wanting to hide their information and location whilst online. This can sometimes be to avoid an oppressive government, but is also used to carry out phishing or cyber attacks.

Encryption

Encryption is a process of scrambling data, preventing it being read or understood. Encryption uses keys to encrypt and decrypt data, and as with a physical lock, data which is intercepted cannot be read without the key. Encryption is particularly important when data is shared over a network or The Internet where other people may be able to intercept it.

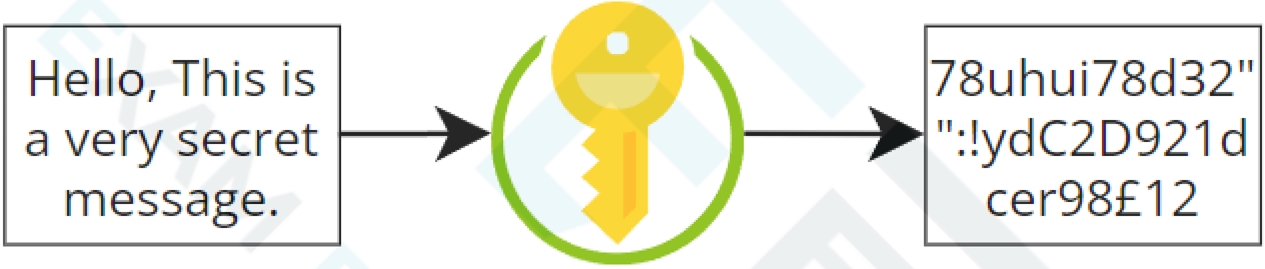
Hello, This is a very secret message.
78uhui78d32"":!ydC2D921dcer98£12

Symmetric Encryption

Symmetric encryption uses the same key to both encrypt and decrypt the data, meaning that before information can be sent, the sender and receiver must participate in a key exchange process so that both have a copy of the key.

This key exchange process is the main flaw with symmetric encryption, it is difficult to find a way to securely share the required key, and if the key is intercepted it can be used by anyone to decrypt and access the data.

Asymmetric Encryption

Asymmetric encryption uses two pairs of mathematically related keys, giving a total of four keys. One key in each pair is known as the private key and must be stored securely. The other is known as the public key and can be freely shared, often via The Internet. Messages encrypted with a public key can only be decrypted with the matching private key and not with any other key, including the public key.

To send a message, the sender encrypts the message using the recipients public key then sends the encrypted message via The Internet. The recipient will then use their private key to decrypt the message and read the contents. If the recipient needs to reply, they can do so using the sender's public key.

Digital Signatures and Certificates

Digital signatures form a key part of the asymmetric encryption process by allowing the recipient to verify both the identity of the sender and that message has not been tampered with during transmission. These steps show how a digital signature is created, checked, and applied to a message.

1) Hashing or a checksum algorithm are used to generate a digest of the message.
2) The sender uses their private key to encrypt the digest.
3) The digest is added to the end of the message.
4) The recipient's public key is used to encrypt the message, with the digest still at the end.

When the recipient receives the message, they follow these steps to decrypt it and check the identity.

1) The recipient's private key is used to decrypt the message.
2) The sender's public key is used to decrypt the digest.
3) The recipient runs the same checksum or hashing algorithm as the sender and checks to ensure that result matches the digest they received.

Because only the sender has access to their private key, the recipient can be sure that the digest really was created by them, and if the digest generated by the recipient matches that sent with the message, they can be sure that the message has not been tampered with.

Digital certificates also form a key part of the asymmetric encryption process by verifying who owns a particular key used to encrypt data. Digital certificates are issued by certificate authorities to individuals or organisations and verify that a key is genuinely issued to that organisation or person.

Malware - Viruses, Worms and Trojans

Malware malicious software which is designed to damage computer systems or the data held on them. When computers are able to communicate via The Internet or a network, it is much easier for malware to spread between them and so it is more important than ever to take precautions against malware.

Worms

A worm is a type of malware which is able to create copies of itself (known as self-replication). This can occur automatically by having the worm seek out other computers via a network, or by having users run a malicious file.

Trojans

A trojan is a type of malware which disguises itself as a normal, safe file in an attempt to trick users into opening it. Trojans disguise themselves as many different things, from a game or program a user may download from the Internet to document attached to an email they receive.

Viruses

A virus is a type of malware which is stored within another file, usually an executable file.

Malware Prevention

The most effective means to prevent malware is to make sure that all computers have antivirus software installed, that it is correctly configured and up to date. Antivirus programs scan files on computers for signs of malware and automatically remove any infected files. Most modern operating systems have this software built into them by default.

Organisations should also train their staff about the risks associated with malware, how to spot malware, and how to report a suspicious file or malware infection. Doing so will greatly reduce the risk that an employee will inadvertently install malware on their system.

Even with these and other precautions in place however it is impossible to completely avoid malware.

TCP / IP

The transmission control protocol / internet protocol, or TCP/IP for short, is used on local networks and The Internet to allow devices to communicate with each other. The protocol is made up of four different layers which each focus on a specific part of communication.

Application Layer	This layer interacts with application software such as an email client or web browser. It receives the data and chooses the most appropriate protocol to transmit it.
Transport Layer	Creates an end to end connection through the network between the sender and recipient known as a virtual path. Once the path has been created, data is split up into packets, each with a sequence number to identify its position to make sure packets are assembled in the correct number. The packet also contains the port number, which identifies the protocol being used.
Network Layer	Adds the source and destination IP addresses onto each packet. Routers operate at this layer and use the IP addresses to route packets through the network.
Link Layer	Controls the physical connections between network devices. This layer adds the MAC addresses to the packets. If a packet passes through multiple routers, the MAC addresses will change each time.

When sending data, the layers are processed from the top down, whereas to receive data they are processed from the bottom up.

Socket Addresses

A socket address is formed by joining an IP Address with a specific port number. This identifies not only a device, but also a particular port on that device and therefore the protocol being used.

The example below shows the socket address for port 443 on a device with the IP address 112.1.221.34. Notice how the colon is used to join the two pieces of information together.

```
112.1.221.34:443
```

Well Known Ports

These are commonly used ports associated with a particular networking protocol.

FTP (File Transfer Protocol) Ports 20 & 21

A protocol used for sending files between devices on a network, FTP clients receive files from FTP servers. FTP servers setup in non-anonymous mode require a username and password to access files, those in anonymous mode allow access to clients without needing a username or password.

SSH (Secure Shell) Port 22

A protocol used for remote management and control of devices via a network. SSH requires a username and password to access a device and all traffic is encrypted. SSH client software is needed to establish this connection and send commands to the remote device.

HTTP (Hypertext Transfer Protocol) Port 80

A protocol used to transfer web pages. Web servers store web pages in text form, whilst clients use application software to request and receive these files and display them as web pages.

HTTPS (Hypertext Transfer Protocol Secure) Port 443

Performs the same protocol as HTTP, but data is encrypted during transmission to keep data secure and prevent modification of data during transmission.

POP3 (Post Office Protocol version 3) Ports 110 & 995

A protocol used to receive emails from an email servers.

SMTP (Simple Mail Transfer Protocol) Ports 25, 587 & 465

A protocol used to send emails.

IP Address Structure

Networks can be divided up into smaller sections, known as subnets to reduce network traffic and make the network more organised and easier to manage.

To allow this, IP address are made up of a network identifier and a host identifier. As the names suggest, the network identifier identifies the subnet, whilst the host identifier identifies the individual host. Devices must still have a unique IP address, but devices on the same subnet will have the same network identifier.

A subnet mask is used with the IP address to work out which part of the address is the network identifier and which is the host identifier. To find the network ID, apply a binary AND to the IP Address and subnet mask as shown in the example below.

IP Address

`192.168.0.22`

Subnet Mask

`255.255.255.0`

IP Address Converted to Binary

`11000000.10101000.00000001.00010110`

Subnet Mask Converted to Binary

`11111111.11111111.11111111.00000000`

Binary AND Applied to IP Address and Subnet Mask to Find the Network ID

`11000000.10101000.00000000.00000000`

Network ID Converted from Binary

`192.168.0.0`

When designing subnets, it is important to consider the number of devices which need to use the subnet and how many subnets are needed. Assigning more bits to the network ID will allow a greater number of subnets to be created, but each subnet will support fewer devices. Assigning more its to the host ID will allow fewer subnets to be created, but each will support a larger number of devices.

Standards for IP Addresses

Two different versions of IP Addresses are in common use today, and it is important to know about both for your exam. Version 4 (IPv4) is the older standard whilst version 6 (IPv6) is the newer.

IPv4

The IPv4 standard used dotted quad numbers which take the format shown below, you have probably seen this type of IP address before. The address is made up of four 8 bit parts separated with a dot, meaning each part can be between 0 and 255 and one IPv4 address takes up 32 bits.

IPv4 allows for just over 4 billion unique addresses, however, with the rapid growth of The Internet more and more devices required internet routable IP addresses. As a result, the number of IPv4 addresses available for use on The Internet is rapidly running out and so a new standard IPv6 was developed to allow for a much greater number of addresses.

`172.0.0.1`

IPv6

IPv6 addresses are made up of eight blocks of four hexadecimal (a-f and 0-9) characters, with a colon between each. This arrangement takes up 128 bits, but allows for 10^{37} addresses, significantly more than IPv4.

```
2071:0cb8:85a3:8a5f:0100:0000:0370:7264
```

Public and Private IP Addresses

An IP address can be either public (known as routable) or private (known as non-routable). Public IP addresses must be globally unique to allow them to be used on the Internet, however, there are not enough IPv4 addresses to allow every device connected to a network to have its own IP Address. Global authorities are responsible for managing and assigning these public IP addresses to ensure that there are no duplications.

Instead, devices within a home or organisation are assigned private IP addresses. These must still be unique within the network, but od not need to be globally unique. Each home or organisation which needs internet access is assigned a single public IP address, although larger businesses may more than one. The NAT (Network Address Translation) process, which is explained later in this topic, allows these devices to share a single public IP address.

DHCP (Dynamic Host Configuration Protocol)

The DHCP protocol is used to assign IP addresses to devices on a network. It uses a pool of IP addresses, and allocates one to each device as it joins the network. These addresses are allocated for a fixed period of time, and once the time elapses or the device leaves the network, the address is returned to the pool for another device to use.

This avoids the manual work required to configure each device with an IP address and manage and record which addresses are available and which are in use. It also avoids IP addresses being wasted when a device leaves the network unnoticed.

NAT (Network Address Translation)

Network Address Translation (NAT) allows several devices on the same network to share one routable (public) IP address. Each device still needs its own non-routable (private) address. A router with two IP addresses, one of routable and one non-routable is used to perform the translation.

Devices send any packets which should be sent via The Internet to the router. When the router receives these packets, it adds a record of them to its database before replacing the sender IP address with its own routable IP address and sending the packet on.

When the router receives a response from The Internet, it checks its database to find the correct non-routable IP address before replacing the destination IP Address and sending the packet on.

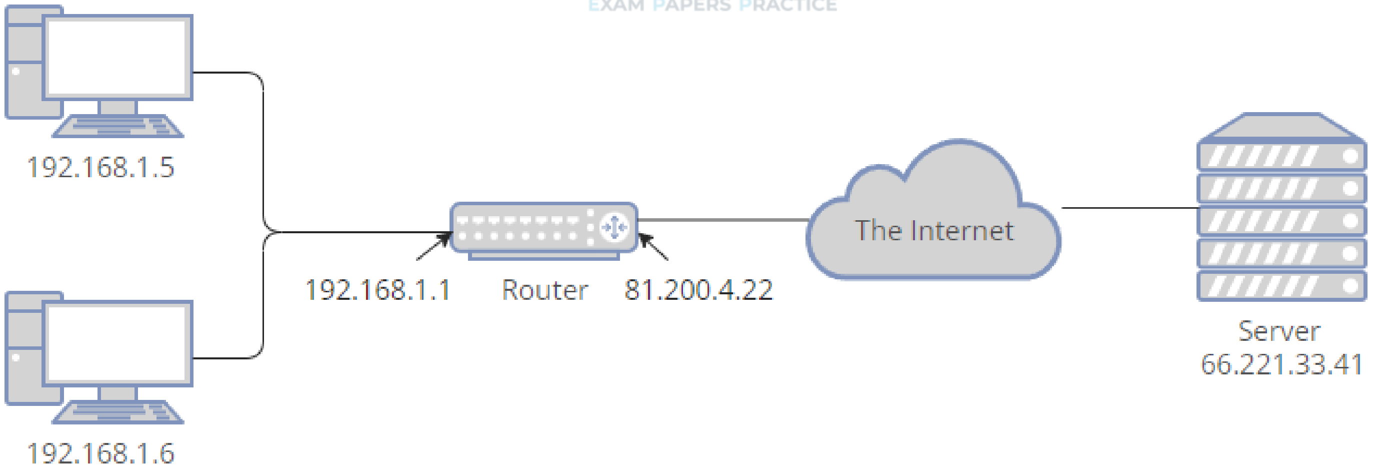

The picture above shows a small network with two computers using a router to share a connection to the Internet. Below we can see what happens when the computer at the top left sends a packet to a server via The Internet.

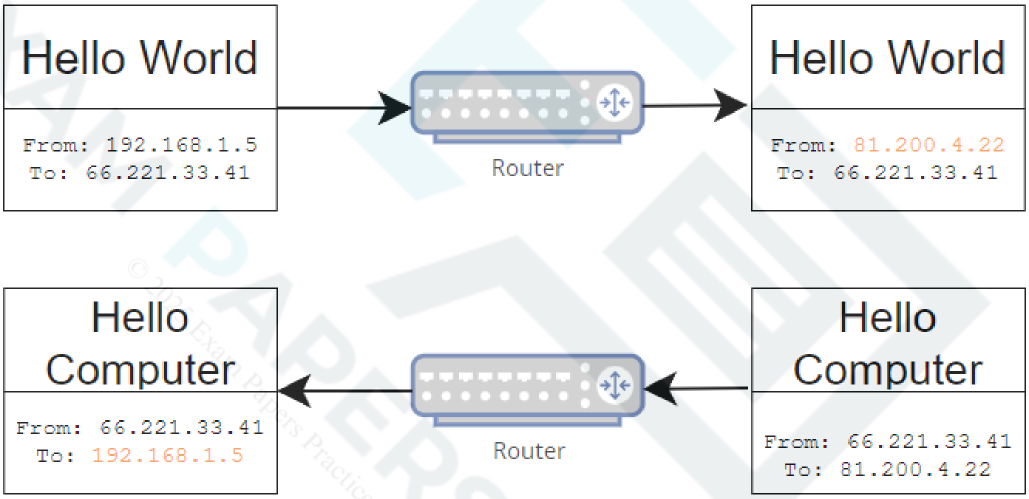

Port Forwarding

Port forwarding allows a remove device to connect to a computer or service in a private network via The Internet. Computers on private networks use non-routable addresses, which can't be used on the Internet; the port forwarding process works similar to NAT (described above) by using a router to provide access to these devices. The router is configured to forward any traffic which arrives on a particular port to a specific device on the network. This is often used to forward ports 80 and 443 to allow a web server to be accessed.

The Client Server Model

The client server model uses one or more powerful computers, known as servers, to store data and perform tasks centrally. Computers used by end users, known as clients, send request messages to the server, which replies with a response message. The response message might contain information, an error, or a confirmation that a task has been completed depending on the nature of the request.

Larger networks may have several specialist servers performing different tasks such as email servers, database servers, web servers and file servers, whilst a smaller network might have one server performing multiple tasks.

APIs

An API (Application Programming Interface) is a set of protocols which define how applications can communicate with one another. They set out how the communication should take place and what languages should be used, allowing one application to access data or perform tasks in another.

The WebSocket Protocol

The WebSocket protocol is an API which operates at the application layer of the TCP/IP stack, it provides a constant stream of information between two devices and is usually used between a web browser and a web server. The connection is full-duplex, meaning data can be sent by both devices at the same time. It provides a faster connection by reducing the size of packet headers and as such is often used for online gaming and video streaming.

CRUD & REST

CRUD

CRUD is an acronym which refers to the four main operations relating to data, Create, Retrieve, Update and Delete. When developing any application which works with data, it is important to make sure the application has the facility to implement all four operations.

Each CRUD statement has a corresponding SQL Statement to perform the action:

CRUD	SQL
Create	INSERT
Retrieve	SELECT
Update	UPDATE
Delete	DELETE

REST

Rest stands for Representational State Transfer and is an style of architecture for developing web services which use Hypertext Transfer Protocol to carry out the four CRUD operations. Web applications which adhere to this standard are known as RESTful. The REST protocol has four methods, each of which links to a CRUD and SQL operation.

CRUD	SQL	REST
Create	INSERT	POST
Retrieve	SELECT	GET
Update	UPDATE	PUT
Delete	DELETE	UPDATE

RESTFul applications allow a client and server to be programmed by different people independently and still interact successfully without needing to know details of how the other is programmed. Clients send requests to the server to perform one of the four actions. Once the request is complete, the server sends a response.

JSON

JSON stands for JavaScript Object Notation and is a text-based standard used to exchange data between systems. It is built around an ordered list of values containing name/value pairs. JSON is compact and easy for humans to read and create, it is also fast for computers to process. The example below shows data in the JSON format.

```
{"employees":[
    {"name":"Shyam", "email":"shyamjaiswal@gmail.com"},
    {"name":"Bob", "email":"bob32@gmail.com"},
    {"name":"Jai", "email":"jai87@gmail.com"}
]}
```

XML

XML Stands for Extensible Markup Language and is another standard for exchanging data. Whilst still text-based, XML is longer than JSON and every tag has to have a corresponding end tag. It is however seen as more flexible than JSON. The example below shows data held in the JSON format.

```
<bookstore>
  <book category="CHILDREN">
    <title lang="en">Harry Potter</title>
    <author>J K. Rowling</author>
    <year>2005</year>
    <price>29.99</price>
  </book>
</bookstore>
```

Thick and Thin Client Computing

Thin Client Networks

Thin client networks use powerful servers to carry out the majority of tasks and processing, whilst clients have only enough processing and storage needed to allow them to connect to the server.

This means that the clients themselves are cheap and it is easy to add new clients to the network. It also allows for greater control and easier management of security and updates since this all takes place on the server.

On the other hand, the servers themselves are expensive and require considerable expertise to setup and manage.

Thick Client Networks

Thick client networks use clients which have their own processing power and storage, allowing them to carry out some or all of the required tasks and processing. Servers can still be used, but they are often less powerful and only carry out specialist or centralised tasks, some thick client networks do not have a server at all.

This setup requires more powerful and therefore more expensive clients, making it more expensive to add new clients to the network. It also removes the extensive cost and expertise needed to setup a central server. Thick client networks also require smaller traffic volumes than thin clients since more of the processing is done locally. This reduces the likeliness of data collisions.

However, this also makes the network harder to maintain since the clients themselves also need to be kept up to date.

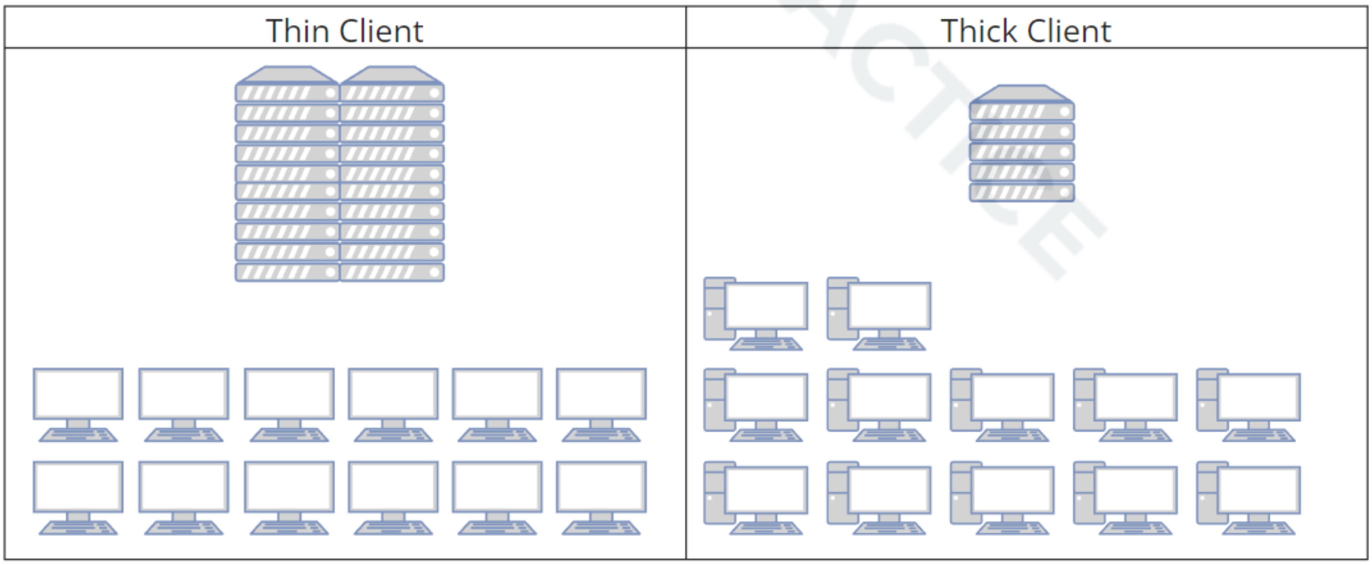
Thin Client
Thick Client