BRIDGE

Bridges can be used to connect two or more LAN segments of the same type (e.g. Ethernet to Ethernet, or Token-Ring to Token-Ring). Like repeaters, bridges can extend the length of a network, but unlike repeaters they can also extend the capacity of a network, since each port on a bridge has its own MAC address. When bridges are powered on in an Ethernet network, they start to learn the network's topology by analysing the source addresses of incoming frames from all attached network segments (a process called backward learning).

Over a period of time, they build up a routing table . Unless the source and the destination are on different network segments, there is no need for the bridge to transfer an incoming frame to another network segment. If the source and the destination are on different segments, the bridge needs to be able to determine which segment the destination device belongs to.

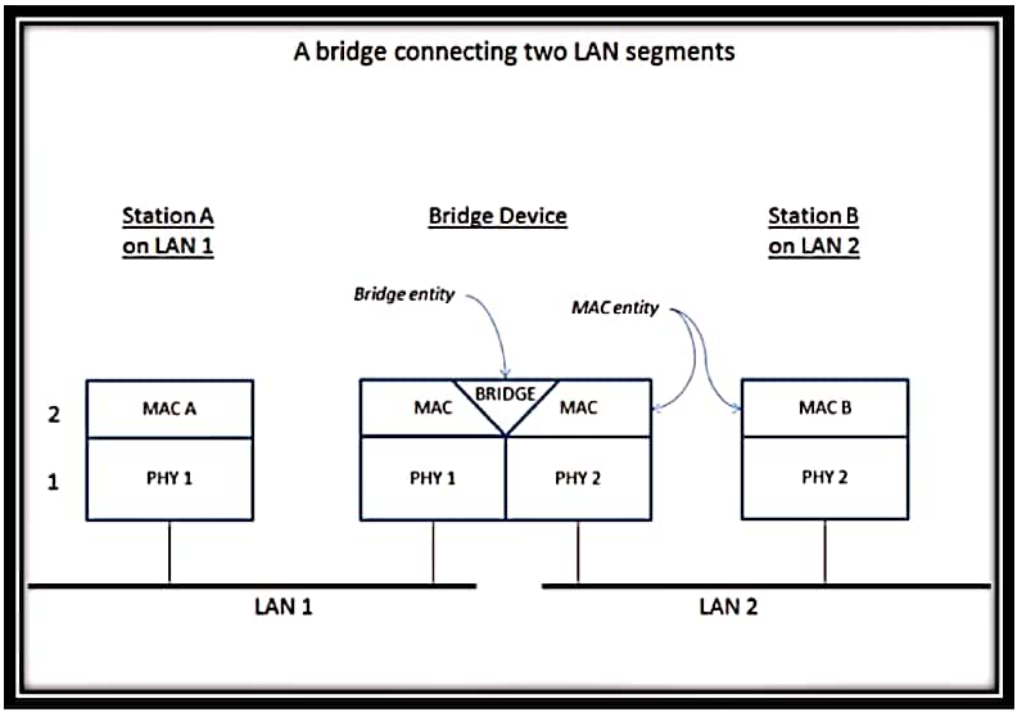

Fig: A high-level overview of network bridging, using the ISO/OSI layers and terminology

Types

There are four types of network bridging technologies: simple bridging, multiport bridging, learning or transparent bridging, and source route bridging.

Transparent bridging

A transparent bridge uses a forwarding database to send frames across network segments. The forwarding database is initially empty and entries in the database are built as the bridge receives frames. If an address entry is not found in the forwarding database, the frame is flooded to all other ports of the bridge, flooding the frame to all segments except the one from which it was received. By means of these flooded frames, the destination network will respond and a forwarding database entry will be created.

In the context of a two-port bridge, the forwarding database can be thought of as a filtering database. A bridge reads a frame's destination address and decides to either forward or filter. If the bridge determines that the destination node is on another segment on the network, it forwards (retransmits) the frame to that segment. If the destination address belongs to the same segment as the source address, the bridge filters (discards) the frame. As nodes transmit data through the bridge, the bridge establishes a filtering database of known MAC addresses and their locations on the network. The bridge uses its filtering database to determine whether a packet should be forwarded or filtered.

Simple bridging

The simple bridge connects two network segments. It typically operates transparently and decides on a packet-by-packet basis whether or not to forward from one network to the other. A store and forward technique is typically used such that, during forwarding, packet integrity is verified on the source network and CSMA/CD delays are accommodated on the destination network. In contrast to simple repeaters that were used to just extend the maximum reach of a segment, a bridge also reduces collisions by splitting the collision domain. Additionally, it lowers overall traffic by only forwarding those packets that are required to cross the bridge.

Multiport bridging

A multiport bridge connects multiple networks and operates transparently to decide on a packet-by-packet basis whether and where to forward. Like the simple bridge, a multiport bridge typically uses store and forward operation. The multiport bridge function is the basis for a network switch.

Source-Route Bridging (Token Ring)

The source-route bridging (SRB) algorithm was developed by IBM for Token Ring networks, and gets its name from the fact that routing information is placed in all inter-segment frames by the sending device. Bridges forward frames according to the routing information carried within the frame. A simple source-route bridging network is illustrated below.

Frame format

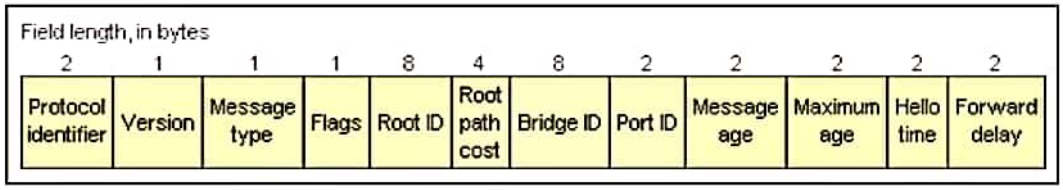

Fig: Bridge configuration message format

The fields of the bridge configuration message are described below.

Protocol identifier - contains the value zero.

Version - contains the value zero.

Message type - contains the value zero

Flags - only the first two bits are used. The topology-change bit , if set, signals a topology change, and the topology-change acknowledgment bit , if set, acknowledges receipt of a configuration message with the topology-change bit set.

Root ID - identifies the root bridge using its 2-byte priority followed by its 6-byte ID.

Root path cost - the root path cost from the bridge sending the configuration message to the root bridge.

Bridge ID - identifies the bridge sending the message using its 2-byte priority followed by its 6-byte ID.

Port ID - identifies the port from which the configuration message was sent.

Message age - indicates when the configuration message should be deleted.

Maximum age - contains the value zero.

Hello time - indicates the time period between root bridge configuration messages.

Forward delay - the time bridges should wait before transitioning to a new state after a topology change.

HUB

Hubs are used in Ethernet networks. A signal received at any port on the hub is retransmitted on all other ports. Network segments that employ hubs are often described as having a star topology, in which the hub forms the wiring centre of the star.

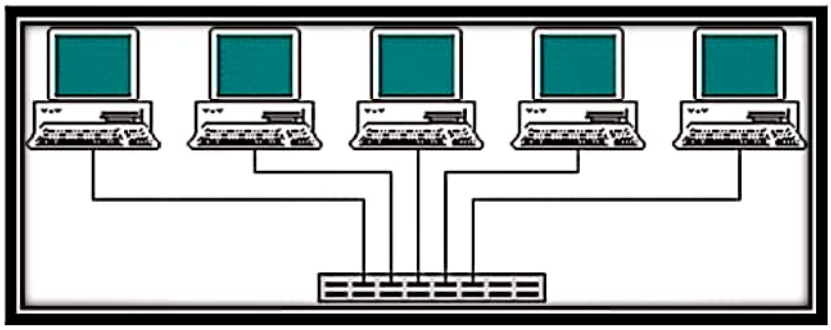

Fig: A hub in a star network configuration

Using a hub provides a degree of fault tolerance, because each network device has its own connection to the hub, and if a connection fails, only a single device is affected. Expanding the network is also easier, because many additional devices can be added to the network using a single hub, which is itself often connected to a network backbone. Hubs can be either active or passive. An active hub has its own power supply, and regenerates incoming frames before retransmitting them. Because signals are regenerated, each output port can connect a channel of up to 100 metres (the maximum allowed for twisted pair cables). Passive hubs simply relay the signal without regenerating it. Managed hubs allow administrators to enable or disable individual ports remotely, while intelligent hubs can autonomously close ports down if the occurrence of errors in transmitted packets exceeds a certain threshold.

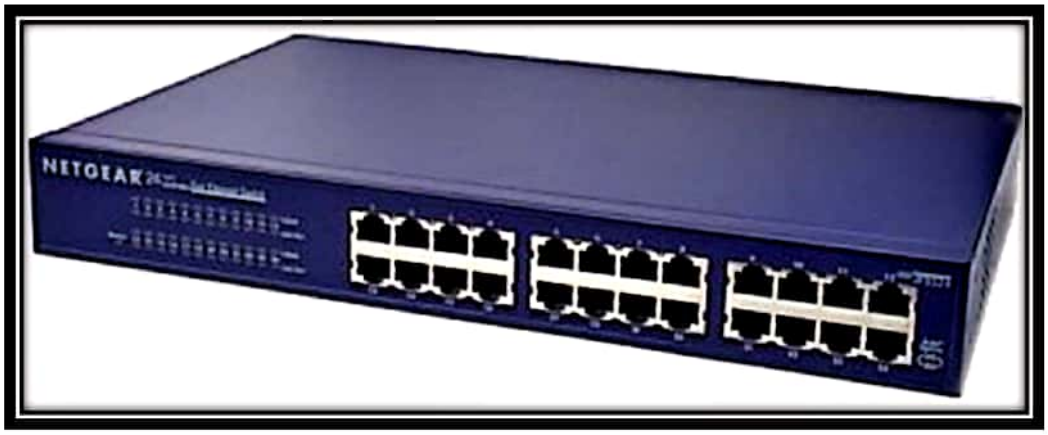

Fig: A 24-port hub

What Hubs Do

Hubs and switches serve as a central connection for all of your network equipment and handles a data type known as frames. Frames carry your data. When a frame is received, it is amplified and then transmitted on to the port of the destination PC.

In a hub, a frame is passed along or "broadcast" to every one of its ports. It doesn't matter that the frame is only destined for one port. The hub has no way of distinguishing which port a frame should be sent to. Passing it along to every port ensures that it will reach its intended destination. This places a lot of traffic on the network and can lead to poor network response times.

SWITCH

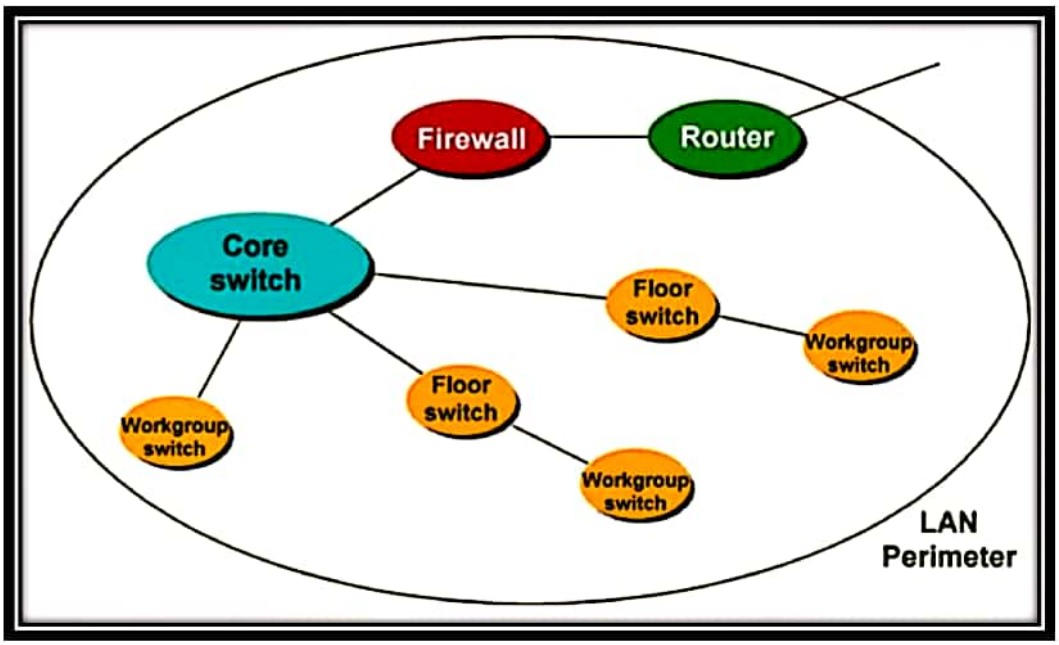

Fig: A core switch connects the high-level devices on the network

The switch is a relatively new network device that has replaced both hubs and bridges in LANs. A switch uses an internal address table to route incoming data frames via the port associated with their destination MAC address. Switches can be used to connect together a number of end-user devices such as workstations, or to interconnect multiple network segments.

A switch that interconnects end-user devices is often called a *workgroup switch*. Switches provide dedicated full-duplex links for every possible pairing of ports, effectively giving each attached device its own network segment This significantly reduces the number of intra-segment and inter-segment collisions.

The Hub and Switch Have Similar Roles

Each serves as a central connection for all of your network equipment and handles a data type known as frames. Frames carry your data. When a frame is received, it is amplified and then transmitted on to the port of the destination PC. The big difference between these two devices is in the method in which frames are being delivered.

In a hub, a frame is passed along or "broadcast" to every one of its ports. It doesn't matter that the frame is only destined for one port. The hub has no way of distinguishing which port a frame should be sent to. Passing it along to every port ensures that it will reach its intended destination. This places a lot of traffic on the network and can lead to poor network response times.

Additionally, a 10/100Mbps hub must share its bandwidth with each and every one of its ports. So when only one PC is broadcasting, it will have access to the maximum available bandwidth. If, however, multiple PCs are broadcasting, then that bandwidth will need to be divided among all of those systems, which will degrade performance.

A switch, however, keeps a record of the MAC addresses of all the devices connected to it. With this information, a switch can identify which system is sitting on which port. So when a frame is received, it knows exactly which port to send it to, without significantly increasing network response times.

And, unlike a hub, a 10/100Mbps switch will allocate a full 10/100Mbps to each of its ports. So regardless of the number of PCs transmitting, users will always have access to the maximum amount of bandwidth. It's for these reasons a switch is considered to be a much better choice then a hub.

ROUTER

A network environment that consists of several interconnected networks employing different network protocols and architectures requires a sophisticated device to manage the flow of traffic between these diverse networks. Such a device, sometimes referred to as an intermediate system, but more commonly called a router, must be able to determine how to get incoming packets (or datagrams) to the destination network by the most efficient route.

Routers gather information about the networks to which they are connected, and can share this information with routers on other networks. The information gathered is stored in the router's internal routing table, and includes both the routing information itself and the current status of various network links. Routers exchange this routing information using special routing protocols.

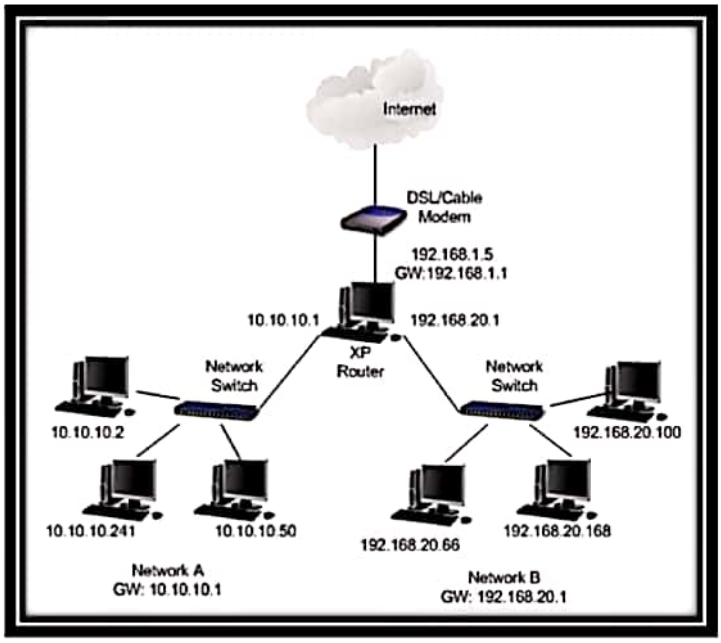

Computers, and other end-user devices attached to networks that form part of an internetwork, are often called hosts or end-systems. A network host does not know how to forward a datagram to a host on another network, and so it will forward the datagram to its local router (or default gateway). A datagram may traverse a number of networks, and hence a number of routers, as it travels from an end-system on the source network to an end-system on the destination network. At each intermediate router, a decision is made as to what is the optimum next hop.

The process undertaken by the router in transferring the incoming datagram to one of its output ports in this way is called switching, and routers are at the heart of packet-switching networks. Unlike bridges and switches, routers do not concern themselves with MAC addresses, and instead examine the IP address contained within a datagram to determine the address of the destination network.

VIRTUAL LAN

In computer networking, a single layer-2 network may be partitioned to create multiple distinct broadcast domains, which are mutually isolated so that packets can only pass between them via one or more routers; such a domain is referred to as a virtual local area network, virtual LAN or VLAN.

This is usually achieved on switch or router devices. Simpler devices only support partitioning on a port level (if at all), so sharing VLANs across devices requires running dedicated cabling for each VLAN. More sophisticated devices can mark packets through tagging, so that a single interconnect (trunk) may be used to transport data for multiple VLANs.

Grouping hosts with a common set of requirements regardless of their physical location by VLAN can greatly simplify network design. A VLAN has the same attributes as a physical local area network (LAN), but it allows for end stations to be grouped together more easily even if they are not on the same network switch. VLAN membership can be configured through software instead of physically relocating

devices or connections. Most enterprise-level networks today use the concept of virtual LANs. Without VLANs, a switch considers all interfaces on the switch to be in the same broadcast domain.

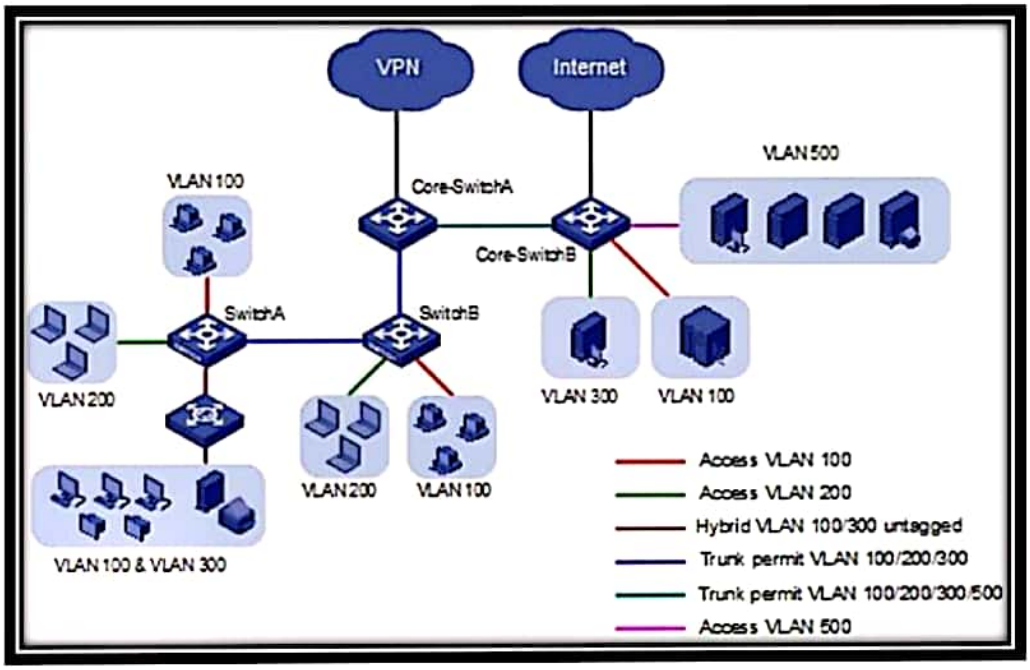

Protocols and design

IEEE 802.1Q

The protocol most commonly used today to configure VLANs is IEEE 802.1Q. The IEEE committee defined this method of multiplexing VLANs in an effort to provide multivendor VLAN support. Prior to the introduction of the 802.1Q standard, several proprietary protocols existed, such as Cisco's ISL (Inter-Switch Link) and 3Com's VLT (Virtual LAN Trunk). Cisco also implemented VLANs over FDDI by carrying VLAN information in an IEEE 802.10 frame header, contrary to the purpose of the IEEE 802.10 standard.

Both ISL and IEEE 802.1Q tagging perform "explicit tagging" - the frame itself is tagged with VLAN information. ISL uses an external tagging process that does not modify the existing Ethernet frame, while 802.1Q uses a frame-internal field for tagging, and therefore does modify the Ethernet frame. This internal tagging is what allows IEEE 802.1Q to work on both access and trunk links: frames are standard Ethernet, and so can be handled by commodity hardware.

Under IEEE 802.1Q, the maximum number of VLANs on a given Ethernet network is 4,094 (the 4,096 provided for by the 12-bit VID field minus reserved values 0x000 and 0xFFF). This does not impose the same limit on the number of IP subnets in such a network, since a single VLAN can contain multiple IP subnets. The VLAN limit is expanded to 16 million with Shortest Path Bridging.

Protocol-based VLAN's

In a switch that supports protocol-based VLANs, traffic is handled on the basis of its protocol. Essentially, this segregates or forwards traffic from a port depending on the particular protocol of that traffic; traffic of any other protocol is not forwarded on the port.

For example, it is possible to connect the following to a given switch:

A host generating ARP traffic to port 10

A network with IPX traffic to port 20

A router forwarding IP traffic to port 30

If a protocol-based VLAN is created that supports IP and contains all three ports, this prevents IPX traffic from being forwarded to ports 10 and 30, and ARP traffic from being forwarded to ports 20 and 30, while still allowing IP traffic to be forwarded on all three ports.