Ethernet FAQs

What is Ethernet?

Ethernet is a type of network cabling and signalling specifications (OSI Model layers 1 [physical] and 2 [data link]) originally developed by Xerox in the late 1970. In 1980, Digital Equipment Corp. (DEC), Intel and Xerox (the origin of the term DIX, as in DEC/Intel/Xerox) began joint promotion of this baseband, CSMA/CD computer communications network over coaxial cabling, and published the "Blue Book Standard" for Ethernet Version 1. This standard was later enhanced, and in 1985 Ethernet II was released. The IEEE's (Institute of Electrical and Electronics Engineers') Project 802 then (after considerable debate) used Ethernet Version 2 as the basis for the 802.3 CSMA/CD network standard. The IEEE 802.3 standard is generally interchangeable with Ethernet II, with the greatest difference being the construction of the network packet header. A complete description of all Ethernet specifications is far outside the scope of this document. If this area interests you, you are encouraged to obtain copies of the IEEE 802.3 documents, and perhaps the ISO 8802-3 documents as well.

What is an 802.3 network?

That's IEEE-ish for Ethernet, but with a few small differences. The physical layer specifications are identical (though DIX Ethernet never specified standards for UTP and Fiber-Optic media) and the MAC sublayer are somewhat different.

What is CSMA/CD?

CSMA/CD is the media access control mechanism used by Ethernet and 802.3 networks; in other words, it determines how a packet of data is placed on the wire. CSMA/CD stands for "Carrier Sense Multiple Access, with Collision Detection". Before an Ethernet device puts a packet "on the wire", it listens to find if another device is already transmitting. Once the device finds the wire is clear, it starts sending the packet while also listening to hear if another device started sending at the same time (which is called a collision). Refer to the Q&A on collisions for more info about this phenomena.

What is a baseband network?

A baseband network is one that provides a single channel for communciations accross the physical medium (e.g., cable), so only one device can transmit at a time. Devices on a baseband network, such as Ethernet, are permitted to use all the available bandwidth for transmission, and the signals they transmit do not need to be multiplexed onto a carrier frequency. An analogy is a single phone line such as you usually have to your house: Only one person can talk at a time--if more than one person wants to talk everyone has to take turns.

What is a broadband network?

Simplisticly, it is the opposite of a baseband network. With broadband, the physical cabling is virtually divided into several different channels, each with its own unique carrier frequency, using a technique called "frequency division modulation". These different frequencies are multiplexed onto the network cabling in such a way to allow multiple simultaneous "conversations" to take place. The effect is similar to having several virtual networks traversing a single piece of wire. Network devices "tuned" to one frequency can't hear the "signal" on other frequencies, and visa-versa. Cable-TV is an example of a broadband network: multiple conversations (channels) are transmitted simultaneously over a single cable; you pick which one you want to listen to by selecting one of the frequencies being broadcast.

What is the OSI Model?

The Open Systems Interconnect (OSI) reference model is the ISO (International Standards Organization) structure for the "ideal" network architecture. This Model outlines seven areas, or layers, for the network. These layers are (from highest to lowest): 7.) Applications: Where the user applications software lies. Such issues as file access and transfer, virtual terminal emulation, interprocess communication and the like are handled here. 6.) Presentation: Differences in data representation are dealt with at this level. For example, UNIX-style line endings (CR only) might be converted to MS-DOS style (CRLF), or EBCIDIC to ASCII character sets. 5.) Session: Communications between applications across a net- work is controlled at the session layer. Testing for out-of-sequence packets and handling two-way communication are handled here. 4.) Transport: Makes sure the lower three layers are doing their job correctly, and provides a transparent, logical data stream between the end user and the network service s/he is using. This is the lower layer that provides local user services. 3.) Network: This layer makes certain that a packet sent from one device to another actually gets there in a reasonable period of time. Routing and flow control are performed here. This is the lowest layer of the OSI model that can remain ignorant of the physical network. 2.) Data Link: This layer deals with getting data packets on and off the wire, error detection and correction and retransmission. This layer is generally broken into two sub-layers: The LLC (Logical Link Control) on the upper half, which does the error checking, and the MAC (Medium Access Control) on the lower half, which deals with getting the data on and off the wire. 1.) Physical: The nuts and bolts layer. Here is where the cable, connector and signaling specifications are defined. There is also the undocumented but widely recognized ninth network layer: 9.) Bozone (a.k.a., loose nut behind the wheel): The user sitting at and using (or abusing, as the case may be) the networked device. All the error detection/correction algorithms in the world cannot protect your network from the problems initiated at the Bozone layer.

What does an ethernet packet look like?

The ethernet packet preamble is normally generated by the chipset. Software is responsible for the destination address, source address, type, and data. The chips normally will append the frame check sequence. | Preamble - | 62 bits | A series of alternating 1's and 0's used by the ethernet receiver to acquire bit synchronization. This is generated by the chip. | Start Of Frame Delimiter - | 2 bits | Two consecutive 1 bits used to acquire byte alignment. This is generated by the chip. Destination Ethernet Address - | 6 bytes | Address of the intended receiver. The broadcast address is all 1's. + Source Ethernet Address - | 6 bytes | The unique ethernet address of the sending | | station. Length or Type field - | 2 bytes | For IEEE 802.3 this is the number of bytes of data. For Ethernet I&II this is the type of packet. Types codes are > 1500 to allow both to coexist. The type code for IP packets is 0x800 46

bytes | Data - | to | Short packets must be padded to 46 bytes. | 1500 bytes | Frame Check Sequence - | 4 bytes | The FCS is a 32 bit CRC calculated using the AUTODIN II polynomial. This field is normally generated by the chip.

What is the difference between an Ethernet frame and a IEEE802.3 frame? Why is there a difference?
Ethernet was invented at Xerox Palo Alto Research Center and later became an international standard. IEEE handled making it a standard; and their specifications are slightly different from the original Xerox ones. Hence, two different types. 802.3 uses the 802.2 LLC to distinguish among multiple clients, and has a "LENGTH" field where Ethernet has a 2-byte "TYPE" field to distinguish among multiple client protocols. TCP/IP and DECnet (and others) use Ethernet_II framing, which is that which Xerox/PARC originated.

What is a SNAP header ??
Sub-Network Access Protocol, an extention to the original 802.2 data link level format. (SNAP is described in IEEE 802-1990) The 802.2 data link format replaced the Ethernet Protocol Type concept with two 8 bit fields; Source SAP, and Destination SAP. Unfortunately that causes problems with migration of protocols, and the lack of SAP space that is available. So one SAP as allocated for this scheme which greatly expands the available protocol space. When using the SNAP SAP the first 5 bytes of data are used as a protocol ID. The first 3 bytes should be a value allocated to you as a vendor id, the same as you get for Source address values. The is called the OUI (Organizationally Unique ID) The second 2 bytes is a protocol type. Note that this is 802.2 and applies across all 802 LAN media types. For translation bridging, there is a convention, if you set the OUI to zero, you are representing a mapped Ethernet frame. So that a bridge will translate such a frame back into Ethernet format, and not into an 802.3 frame format. 802.2 SNAP frame: MAC | DSAP | SSAP | UI | OUI | Type | data | | Header| 0xAA | 0xAA | 0x03 | 3bytes|2bytes| | This will appear the same on all 802 compliant LAN media. On 802.3, there will be a Length field between the SA and the DSAP but not on 802.5 or FDDI.

What is a MAC address?
It is the unique hexadecimal serial number assigned to each Ethernet network device to identify it on the network. With Ethernet devices (as with most other network types), this address is permanently set at the time of manufacturer, though it can usually be changed through software (though this is generally a Very Bad Thing to do).

Why must the MAC address to be unique?
Each card has a unique MAC address, so that it will be able to exclusively grab packets off the wire meant for it. If MAC addresses are not unique, there is no way to distinguish between two stations. Devices on the network watch network traffic and look for their own MAC address in each packet to determine whether they should decode it or not. Special circumstances exist for broadcasting to every device.

Is there a special numbering scheme for MAC addresses?
The MAC addresses are exactly 6 bytes in length, and are usually written in hexadecimal as 12:34:56:78:90:AB (the colons may be omitted, but generally make the address more readable). Each manufacturer of Ethernet devices applies for a certain range of MAC addresses they can use. The first three bytes of the address determine the manufacturer. RFC-1700 (available via FTP) lists some of the manufacturer-assigned MAC addresses. A more up-to-date listing of vendor MAC address assignments is available on ftp.lcs.mit.edu in pub/map/Ethernet-codes.

What is a preamble?
A seven octet field of alternating one and zero binary bits sent prior to each frame to allow the PLS circuitry to reach its steady state synchronization with received frame timing. (802.3 standard, page 24,42).

What is a Start Frame Delimiter (SFD)?
A binary sequence of '10101011' immediately following the preamble and indicating the beginning of a frame. (802.3 standard, page 24).

What does CRC mean?
Cyclical Redundancy Check - A method of detecting errors in a message by performing a mathematical calculation on the bits in the message and then sending the results of the calculation along with the message. The receiving work-station performs the same calculation on the message data as it receives it and then checks the results against those transmitted at the end of the message. If the results don't match, the receiving end asks the sending end to send again.

What is a broadcast address?
The unique address that identifies a packet as appropriate to all receiveing stations. In 802.3 any address in which the second byte is an odd number. (1,3,...F).

What exactly do 10Base5, 10BaseT, 10Base2, 10Broad36, etc mean?
These are the IEEE names for the different physical types of Ethernet. The "10" stands for signalling speed: 10MHz. "Base" means Baseband, "broad" means broadband. Initially, the last section as intended to indicate the maximum length of an unrepeated cable segment in hundreds of meters. This convention was modified with the introduction of 10BaseT, where the T means twisted pair, and 10BaseF where the F means fiber (see the following Q&A for specifics). This actually comes from the IEEE committee number for that media. In actual practice: 10Base2 Is 10MHz Ethernet running over thin, 50 Ohm baseband coaxial cable. 10Base2 is also commonly referred to as thin-Ethernet or Cheapernet. 10Base5 Is 10MHz Ethernet running over standard (thick) 50 Ohm baseband coaxial cabling. 10BaseF Is 10MHz Ethernet running over fiber-optic cabling. 10BaseT Is 10MHz Ethernet running over unshielded, twisted- pair cabling. 10Broad36 Is 10MHz Ethernet running through a broadband cable.

What does FOIRL mean?
Fiber Optic Inter Repeater Link. A "IEEE 802 standard" worked out between many vendors some time ago for carrying Ethernet signals across long distances via fiber optic cable. It has since been adapted to other applications besides connecting segments via repeaters (you can get FOIRL cards for PCs). It has been superseded by the larger 10BaseF standard.

What is coax cable?
Coaxial cable (coax) is a metallic electrical cable used for RF (radio frequency) and certain data communications transmission. The cable is constructed with a single solid or stranded center conductor that is surrounded by the dielectric layer, an insulating material of constant thickness and high resistance. A conducting layer of aluminum foil, metallic braid or a combination of the two encompass the dielectric and act as both a shield against interference (to or from the center conductor) and as the return ground for the cable. Finally, an overall insulating layer forms the outer jacket of the cable. Coaxial cable is generally superior in high-frequency applications such as networking. However, for shorter distances (up to 100 meters), UTP or STP cable is generally just as reliable when using differential modulation techniques (such as with 10BaseT). There are three types of RG-58 cable, as far as I can tell. There are probably other subtle differences, but for 10BASE2, impedance and velocity of propagation are the important ones. The table below summarizes: Cable Impedance Velocity ---------- ---------- ---------- RG-58A/U 50 ohms .66 or .78 RG-58C/U 50 ohms .66 RG-58/U 53.5 ohms .66 or .695

What does UTP, STP cabling mean?
Twisted pair cables. UTP is for UNshielded, twisted pair, while STP is for SHIELDED, twisted pair. UTP is what's typically installed by phone companies (though this is often not of high enough quality for high- speed network use) and is what 10BaseT Ethernet runs over. UTP is graded according to its data carrying ability (e.g., Level 3, Level 4, Level 5). 10BaseT Ethernet requires at least Level 3 cable. Many sites now install only Level-5 UTP, even though level 4 is more than sufficient for 10BaseT, because of the greater likelihood that emerging high-speed standards will require cable with better bandwidth capabilities. STP is typically used for Token-Ring networks, where it is commonly referred to IBM Type 1 (or 2, 3, 6, 8, etc); however there are several manufacturers of Ethernet equipment and interfaces that support Ethernet over STP. Nevertheless, Ethernet over STP is not officially defined in any standards. While there is a good level of interoperability with Ethernet over STP, (Lattisnet, developed by Synoptics, is the recognized de facto standard in this area), one should consider the long-term availability and cost of this non-standard scheme before planning new networks around it.

Are there any restrictions on how Ethernet is cabled?
Yes, there are many, and they vary according to the media used. First of all, there are distance limitations: 10Base2 limited to 185 meters (607 ft) per unrepeated cable segment. 10Base5 limited to 500 meters (1,640 ft) per unrepeated cable segment. 10BaseF depends on the signaling technology and medium used but can go up to 2KM. 10BaseT generally accepted to have a maximum run of 100-150M, but is really based on signal loss in Db's (11.5db maximum loss source to destination). 10Broad36 limited to 3,600 meters (almost 2.25 miles). Then there are limitations on the number of repeaters and cable segments allowed between any two stations on the network. There are two different ways of looking at the same rules: 1. The Ethernet way: A remote repeater pair (with an intermediate point-to-point link) is counted as a single repeater (IEEE calls it two repeaters). You cannot put any stations on the point to point link (by definition!), and there can be two repeaters in the path between any pair of stations. This seems simpler to me than the IEEE terminology, and is equivalent. 2. The IEEE way: There may be no more than five (5) repeated segments, nor more than four (4) repeaters between any two Ethernet stations; and of the five cable segments, only three (3) may be populated. This is referred to as the "5-4-3" rule (5 segments, 4 repeaters, 3 populated segments). It can really get messy when you start cascading through 10BaseT hubs, which are repeaters unto themselves. Just try to remember, that any possible path between two network devices on an unbridged/unrouted network cannot pass through more than 4 repeaters or hubs, nor more than 3 populated cable segments. Finally, 10Base2 is limited to a maximum of 30 network devices per unrepeated network segment with a minimum distance of 0.5m (1.5ft) between T-connectors. 10Base5 is limited to a maximum of 100 network devices per unrepeated segment, with a minimum distance of 2.5m (8.2ft) between taps/T's (usually indicated by a marker stamped on the cable itself every 2.5m). 10BaseT and 10BaseF are star-wired, so there is no minimum distance requirement between devices, since devices cannot be connected serially. You can install up to the Ethernet maximum of 1024 stations per network with both 10BaseT and 10BaseF.

Can I mix 10Base2 and 10Base5 cabling on a single segment?
It is not "legal", but the network police will not read you your rights and drag you away. Ideally, you should use a repeater (or bridge, router, etc...) between the different cabling types. However, in reality, it will work fine, as long as none of the other network parameters (lengths, numbers of stations, repeaters, etc) are near the limit of the specification.

What about wireless Ethernets? Are there any?
Yes, and no. Many vendors offer equipment for Ethernet across a variety of unbounded, or wireless, connections using lasers, microwaves, and spread-spectrum radio transmissions. However, none of these methods are organized by any standards body, so it is unlikely to find equipment from any two different manufacturers that work together.

When should I choose 10BaseT, when 10Base2 (or others)?
The specific environment and application must be considered when selecting your media type. However, there are some general rules-of-thumb that you can consider: Avoid using copper between buildings. The electrical disturbances caused by lightning, as well as naturally occurring differences in ground potential over distance, can very quickly and easily cause considerable damage to equipment and people. The use of fiber-optic cabling between buildings eliminates network cabling as a safety risk. There are also various wireless media available for inter-building links, such as laser, spread-

spectrum RF and microwave. However, wireless media is much more expensive and less reliable than fiber-optic, and should only be considered when it is impossible to get right-of-way for fiber-optic cable. 10Base2 (thin Ethernet or Cheapernet) is the least expensive way to cable an Ethernet network. However, the price difference between 10Base2 and 10BaseT (Ethernet over UTP) is rapidly diminishing. Still, for small, budget-conscious installations, 10Base2 is the most economical topology. The disadvantages of 10Base2 is that any break in the cable or poor connection will bring the entire network down, and you need repeaters if you have more than 30 devices connected to the network or the cable length exceeds 185 meters (607 feet). 10Base5 is generally used as a low-cost alternative to fiber-optic media for use as a backbone segment within a single building. It's extended length (500m or 1640ft), higher attached device count (100) and better noise resistance make 10Base5 well suited for use as a network trunk for one or more floors in a building. However, the high cost of connecting each device (in addition to the interface, you also need an external transceiver, or MAU, and an AUI cable) makes 10Base5 too expensive for most LAN installations, and like 10Base2, a single break or bad connection in the cable can bring the entire network down. 10BaseT is the most flexible topology for LANs, and is generally the best choice for most network installations. 10BaseT hubs, or multi-hub concentrators, are typically installed in a central location to the user community, and inexpensive UTP cabling is run to each network device (which may be 100m, or 330ft, from the hub). The signalling technology is very reliable, even in somewhat noisy environments, and 10BaseT hubs will usually detect many network error conditions and automatically shut-down the offending port(s) without affecting the rest of the network (unless, of course, the offending port was your server, shared printer, or router to the rest of the world). While the hardware is more expensive than 10Base2, the cabling is cheaper and requires less skill to install, making 10BaseT installation costs only slightly higher than 10Base2. The flexibility and reliability more than offset the marginally higher price. 10BaseF, and its predecessor, FOIRL, are the only recommended topologies for inter-building links. However, they need not be limited to this role. 10BaseF can also be run to the desktop, though the cost is prohibitively high in all but the most specialized environments (generally, extremely noisy manufacturing facilities, or very security-conscious installations). More commonly, FOIRL (and now, 10BaseF) is used inside buildings to form backbone networks and to connect wiring closets together.

Is it safe to run Unshield Twisted Pair next to power cable?
According to EIA/TIA-569, the standard wiring practices for running data cabling and companion to the above referenced EIA/TIA-568, you should not run data cable parallel to power cables. However, in reality, this should not be a problem with networks such as 10BaseT. 10BaseT uses differential signalling to pick the data signals off the wire. Since any interference from nearby power lines will usually affect all pairs equally, anything that is not canceled-out by the twists in the UTP should be ignored by the receiving network interface.

Can I connect the 10BaseT interface of two devices directly together, without using a hub?
Yes, but not more than 2 devices, and you also need a special jumper cable between the two 10BaseT ports: RJ45 pin RJ45 pin ======== ======== 1 <--[TX+]--------[RX+]--> 3 2 <--[TX-]--------[RX-]--> 6 3 <--[RX+]--------[TX+]--> 1 6 <--[RX-]--------[TX-]--> 2

Does my Ethernet coax have to be grounded? How?
Yes and no. The 10Base2 spec says the coax MAY be grounded at one and only one point, while the 10Base5 spec says the coax SHALL be grounded at one and only one point. Grounding your coax is generally a good idea; it allows static electricity to bleed off and, supposedly, makes for a safer installation. Further, many local electrical codes will require your network cabling to be grounded at some point. However, I have personally seen many Ethernet networks work with absolutely NO ground on the segment, and even a few unreliable segments become reliable when the one and only ground was removed. I'm not saying you should not ground your networks -- you should absolutely install cabling according to your electrical codes. On the other hand, if you do ground your cable, make sure you do so only at one point. Multiple grounds on an Ethernet segment will not only cause network errors, but also risk damage to equipment and injury to people. If you have a repeater on one end of the segment, this will usually automatically ground that end of the segment (you may want to check the repeater documentation and configuration to assure this is the case -- most repeaters can be set-up to NOT ground). If you don't have a repeater, you can get terminating resistors with ground straps attached.

What is a "segment"?
A piece of network wire bounded by bridges, routers, repeaters or terminators.

What is a "subnet"?
Another overloaded term. It can mean, depending on the usage, a segment, a set of machines grouped together by a specific protocol feature (note that these machines do not have to be on the same segment, but they could be) or a big nylon thing used to capture enemy subs.

What does "AUI" mean ?
Attachment Unit Interface, an IEEE term for the connection between a controller and the transceiver.

What is a transceiver?
A transceiver allows a station to transmit and receive to/from the common medium. In addition, Ethernet transceivers detect collisions on the medium and provide electrical isolation between stations. 10Base2 and 10Base5 transceivers attach directly to the common bus media, though the former usually use an internal transceiver built-onto the controller circuitry with a "T" connector to access the cable, while the latter use a separate, external transceiver and an AUI (or transceiver) cable to connect to the controller. 10BaseF, 10BaseT and FOIRL also usually use internal transceivers. Having said that, there also also external transceivers for 10Base2, 10BaseF, 10BaseT and FOIRL that can connect externally to the controller's AUI port, either directly or via an AUI cable.

What exactly does a repeater?

A repeater acts on a purely electrical level to connect to segments. All it does is amplify and reshape (and, depending on the type, possibly retime) the analog waveform to extend network segment distances. It does not know anything about addresses or forwarding, thus it cannot be used to reduce traffic as a bridge can in the example above.

What is a "hub"?
A hub is a common wiring point for star-topology networks, and is a common synonym for concentrator 10BaseT Ethernet and 10BaseF Ethernet and many proprietary network topologies use hubs to connect multiple cable runs in a star-wired network topology into a single network. hub. Hubs have multiple ports to attach the different cable runs. Some hubs (such as 10BaseT) include electronics to regenerate and retime the signal between each hub port.

What exactly does a bridge?
A bridge will connect to distinct segments (usually referring to a physical length of wire) and transmit traffic between them. This allows you to extend the maximum size of the network while still not breaking the maximum wire length, attached device count, or number of repeaters for a network segment.

What does a "learning bridge"?
A learning bridge monitors MAC (OSI layer 2) addresses on both sides of its connection and attempts to learn which addresses are on which side. It can then decide when it receives a packet whether it should cross the bridge or stay local (some packets may not need to cross the bridge because the source and destination addresses are both on one side). If the bridge receives a packet that it doesn't know the addresses of, it will forward it by default.

What is a remote bridge?
A bridge as described above that has an Ethernet interface on one side and a serial interface on the other. It would connect to a similar device on the other side of the serial line. Most commonly used in WAN links where it is impossible or impractical to install network cables. A high-speed modem (or T1 DSU/CSU's, X.25 PAD's, etc) and intervening telephone lines or public data network would be used to connect the two remote bridges together.

Is there a maximum number of bridges allowed on a network?
Per IEEE 802.1 (d), the maximum number of concatenated brides in a bridged LAN is 7. This number is rather arbitrary, however, and is based on simulations of application performance with expected bridge delays. In addition, the number assumes that all bridges are LOCAL (no remote WAN connections), and that the default Hold Time of 1 second is in place (this is the time after which a bridge will discard a frame it is holding). This prevents extra-late frame delivery. (i.e, a frame should never be delivered more than ~7 seconds after is it sent). I personally (Rich Seifert) find this to be much too long an allowance. My "rule of thumb" for bridged LANs is to limit the number of hops to 4, with not more than one of these being a WAN linked remote bridge.

What exactly does a router do?
Routers work much like bridges, but they pay attention to the upper network layer protocols (OSI layer 3) rather than physical layer (OSI layer 1) protocols. A router will decide whether to forward a packet by looking at the protocol level addresses (for instance, TCP/IP addresses) rather than the MAC address. Because routers work at layer 3 of the OSI stack, it is possible for them to transfer packets between different media types (i.e., leased lines, Ethernet, token ring, X.25, Frame Relay and FDDI). Many routers can also function as bridges.

So should I use a router or a bridge?
There is no absolute answer to this. Your network layout, type and amount of hosts and traffic, and other issues (both technical and non-technical) must be considered. Routing would always be preferable to bridging except that routers are slower and usually more expensive (due to the amount of processing required to look inside the physical packet and determine which interface that packet needs to get sent out), and that many applications use non-routable protocols (i.e., NetBIOS, DEC LAT, etc.). Rules of thumb: Bridges are usually good choices for small networks with few, if any, slow redundant links between destinations. Further, bridges may be your _only_ choice for certain protocols, unless you have the means to encapsulate (tunnel) the unroutable protocol inside a routable protocol. Routers are usually much better choices for larger networks, particularly where you want to have a relatively clean WAN backbone. Routers are better at protecting against protocol errors (such as broadcast storms) and bandwidth utilization. Since routers look deeper inside the data packet, they can also make forwarding decisions based on the upper-layer protocols. Occasionally, a combination of the two devices are the best way to go. Bridges can be used to segment small networks that are geographicly close to each other, between each other and the router to the rest of the WAN.

Are there problems mixing Bridging & routing?
Only if you plan on having bridged links in parallel with routed links. You need to be very careful about running bridges providing links in parallel to a router. Bridges may forward broadcast requests which will confuse the router there are lots of protocols you may not think of filtering (e.g. ARP, Apple ARP over 802.3 etc. etc.). Also, DECnet routers have the same MAC address on all ports. This will probably cause the bridge to think it is seeing an Ethernet loop.

What is a driver?
Typically the software that allows an Ethernet card in a computer to decode packets and send them to the operating system and encode data from the operating system for transmission by the Ethernet card through the network. By handling the nitty-gritty hardware interface chores, it provides a device-independent interface to the upper layer protocols, thereby making them more universal and [allegedly] easier to develop and use.

What is SQE? What is it for?
SQE is the IEEE term for a collision. (Signal Quality Error)

What is an SQE Test?
SQE Test (a.k.a. heartbeat) is a means of detecting a transceiver's inability to detect collisions. Without SQE Test, it is not possible to determine if your collision detector is operating properly. SQE Test is implemented by generating a test signal on the collision pair from the transceiver (or its equivalent) following every transmission on the network. It does not generate any signal on the common medium. The problem with SQE Test is that it is not part of the Ethernet Version 1.0 specification. Therefore, Version 1.0 equipment may not function with transceiver that generates the SQE Test signal. Additionally, IEEE 802.3 specifications state that IEEE 802.3 compliant repeaters must not be attached to transceivers that generate heartbeat. (This has to do with a jam signal that prevents redundant collisions from occurring on the network). Therefore, you must usually turn-off SQE Test (heartbeat) between the transceiver and an 802.3 repeater.

What means "IPG"?
The InterPacket Gap (more properly referred to as the InterFrame Gap, or IFG) is an enforced quiet time of 9.6 us between transmitted Ethernet frames.

What is a runt?
A packet that is below the minimum size for a given protocol. With Ethernet, a runt is a frame shorter than the minimum legal length of 60 bytes (at Data Link).

What causes a runt?
Runt packets are most likely the result of a collision, a faulty device on the network, or software gone awry.

What is a jabber?
A blanket term for a device that is behaving improperly in terms of electrical signalling on a network. In Ethernet this is Very Bad, because Ethernet uses electrical signal levels to determine whether the network is available for transmission. A jabbering device can cause the entire network to halt because all other devices think it is busy.

What causes a jabber?
Typically a bad network interface card in a machine on the network. In bizarre circumstances outside interference might cause it. These are very hard problems to trace with layman tools.

What is a collision?
A condition where two devices detect that the network is idle and end up trying to send packets at exactly the same time. (within 1 round-trip delay) Since only one device can transmit at a time, both devices must back off and attempt to retransmit again. The retransmission algorithm requires each device to wait a random amount of time, so the two are very likely to retry at different times, and thus the second one will sense that the network is busy and wait until the packet is finished. If the two devices retry at the same time (or almost the same time) they will collide again, and the process repeats until either the packet finally makes it onto the network without collisions, or 16 consecutive collision occur and the packet is aborted.

What causes a collision?
See above. Ethernet is a CSMA/CD (Carrier Sense Multiple Access/ Collision Detect) system. It is possible to not sense carrier from a previous device and attempt to transmit anyway, or to have two devices attempt to transmit at the same time; in either case a collision results. Ethernet is particularly susceptible to performance loss from such problems when people ignore the "rules" for wiring Ethernet.

How many collisions are too many?
This depends on your application and protocol. In many cases, collision rates of 50% will not cause a large decrease in perceived throughput. If your network is slowing down and you notice the percentage of collisions is on the high side, you may want try segmenting your network with either a bridge or router to see if performance improves.

How do I reduce the number of collisions?
Disconnect devices from the network. Seriously, you need to cut- down on the number of devices on the network segment to affect the collision rate. This is usually accomplished by splitting the segment into two pieces and putting a bridge or router in between them.

What is a late collision?
A late collision occurs when two devices transmit at the same time, but due to cabling errors (most commonly, excessive network segment length or repeaters between devices) neither detects a collision. The reason this happens is because the time to propagate the signal from one end of the network to another is longer than the time to put the entire packet on the network, so the two devices that cause the late collision never see that the other's sending until after it puts the entire packet on the network. Late collisions are detected by the transmitter after the first "slot time" of 64 byte times. They are only detected during transmissions of packets longer than 64 bytes. It's detection is exactly the same as for a normal collision; it just happens "too late." Typical causes of late collisions are segment cable lengths in excess of the maximum permitted for the cable type, faulty connectors or improper cabling, excessive numbers of repeaters between network devices, and defective Ethernet transceivers or controllers. Another bad thing about late collisions is that they occur for small packets also, but cannot be detected by the transmitter. A network suffering a measurable rate of late collisions (on

large packets) is also suffering lost small packets. The higher protocols do not cope well with such losses. Well, they cope, but at much reduced speed. A 1% packet loss is enough to reduce the speed of NFS by 90% with the default retransmission timers. That's a 10X amplification of the problem. Finally, Ethernet controllers do not retransmit packets lost to late collisions.

What is a jam?
When a workstation receives a collision, and it is transmitting, it puts out a jam so all other stations will see the collision also. When a repeater detects a collision on one port, it puts out a jam on all other ports, causing a collision to occur on those lines that are transmitting, and causing any non-transmitting stations to wait to transmit.

What is a broadcast storm?
An overloaded term that describes an overloaded protocol. :-). Basically it describes a condition where devices on the network are generating traffic that by its nature causes the generation of even more traffic. The inevitable result is a huge degradation of performance or complete loss of the network as the devices continue to generate more and more traffic. This can be related to the physical transmission or to very high level protocols.

How do I recognize a broadcast storm?
That depends on what level it is occurring. Basically you have to be aware of the potential for it beforehand and be looking for it, because in a true broadcast storm you will probably be unable to access the network. This can change dramatically for a higher level protocol. NFS contention can result in a dramatic DROP in Ethernet traffic, yet no one will have access to resources.

How can I prevent a broadcast storm?
Avoid protocols that are prone to it. Route when it is practical.

What is an Alignment Error ?
A received frame that does not contain an integer number of octets and contains a frame check sequence validation error. A frame in which the number of bits received is not an integer multiple of 8 and has a FCS (Frame Check Sequence) error. (802.3 standard, page 41)

What is *high* traffic on an Ethernet?
High traffic is when things start slowing down to the point they are no longer acceptable. There is not set percentage point, in other words. Xerox used to use a formula based on packet size over time, or something, but the issue has been significantly muddied by the plethora of protocols available and how they react to wire usage. I usually start paying attention over 40-50%, *or when things slow down*.

How can I test an Ethernet?
This depends on what level you want to test. The most basic test (a.k.a., "the fire test") is to connect a pair of devices to the network and see if they can communicate with each other. If you want to test the electrical integrity of the wire (i.e., will it carry a signal properly), a TDR or cable scanner that incorporates TDR and other functions, would be the most comprehensive tool (though a great deal cab be determined with a simple ohmmeter). If you need to test the performance or troubleshoot protocol transmission problems, you will need special and usually very expensive software, usually coupled with custom hardware, to capture, optionally filter, and analyze the network packets.