Newsgroups: comp.dcom.lans.ethernet
From: barr@tramp.Colorado.EDU (BARR DOUG)
Subject: Ethernet FAQ
Organization: University of Colorado, Boulder
Date: Tue, 5 Jan 1993 20:51:40 GMT

This has not been posted for a while, so I am taking the liberty of posting it:

Q: What is a runt?
A: 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 64 bytes (at Data Link).

Q: What causes a runt?
A: Runt packets can be caused accidentally or intentionally. If accidental, they are most likely the result of a faulty device on the network, or software gone awry. If intentional, they may be designed to be runts for a specific reason. SNMP (Simple Network Management Protocol) is often sent as runt packets so that many devices will simply ignore it.

Q: What is a jabber?
A: 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.

Q: What causes a jabber?
A: 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.

Q: What is a collision?
A: 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, etc.

Q: What causes a collision?
A: 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.

Q: What is a jam?
A: 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.

Q: What is a broadcast storm?
A: 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. There is a famous example of Banyan Vines bringing a huge network to its knees because of the addition of a single server, which brought the network to "critical mass" (this logic error has been corrected). NFS is famous for this type of failure.

Q: How do I recognize a broadcast storm?

A: 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.

Q: How can I prevent a broadcast storm?
A: Avoid protocols that are prone to it. Route when it is practical. Don't buy Ethernet. :-).

Q: What is *high* traffic on an Ethernet? 5%? 20%? 90%?
A: 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*. I've seen IPX segments that were slow with less than 20% usage.

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

Q: What means "heartbeat"? What is it for?
A: Heartbeat (a.k.a. SQE Test) is a means of detecting a transceiver's inability to detect collisions. The normal operation of an Ethernet will test the transceiver's power, transmitter and receiver; if any of these fail the station will not hear its own loopback. Without heartbeat, it is not possible to determine if your collision detector is operating properly. Heartbeat 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.

Note the older usage of this term to refer to the +-.7V carrier sense wave, although I haven't heard it used that way in a while (since SQE indicators became popular on transceivers).

Q: What means "CSMA/CD"?

A: Carrier Sense, Multiple Access, with Collision Detection, the MAC (Media Access Control) algorithm used by Ethernet to help avoid two devices on the same cable from transmitting at the same time, or at least recognize when this has happened so that the two devices can back-off and try again later.

Q: What means "IPG"?

A: 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.

Q: Does a NEMP (Nuclear Electro-Magnetic Pulse) affect an Ethernet?
A: The Russians have done the most research into the effects of NEMP, although the US and various European countries have also looked into it. I doubt that the results and theses from this work is available. Given my very limited understanding of the effect (as a layman), yes, I expect it would. Obviously, a fiber-optic network (since it is non-conducting) would have a greater chance for surviving NEMP. However, I suspect the EMF would not be significantly retarded by most system enclosures to prevent damage to the network interface (as well as the rest of the system internals) in spite of the lack of copper network cables acting as antennae.

Q: What means "promiscuous mode"?
A: A controller in promiscuous mode will receive all frames, regardless of destination address. Ethernet is promiscuous in that it allows any device on a segment to hear every packet on that segment if the card is so programmed. This is an obvious security issue. It used to be that there was no way around this besides encoding the packets themselves, but Synoptics recently released a secure Ethernet solution (blatant employee plug).

Q: How can I test an Ethernet?
A: You must be more specific. Do you wish to test the electrical integrity of the wire (ie, will it carry a signal properly) or do you wish to test the performance of it while running, etc? If the former, a TDR (see below) or cable scanner that incorporates and expands on the capabilities of a TDR would be the most comprehensive tool, though a great deal can be determined with a simple ohmmeter. The latter requires special and often very expensive software, usually combined with custom hardware, to capture, optionally filter, and analyze the network packets. The most basic test is to connect a pair of devices and see if they can communicate with each other, while monitoring any status indicators that the devices might provide.

Q: What is a "TDR"?

A: A Time-Domain Reflectometer is a tool used to detect cable faults. This device operates by sending a brief signal pulse down the cable and looking for its reflection to bounce back. By analyzing the reflected pulse, it is possible to make judgments about the quality of the cable segment. More advanced units can not only detect and identify the nature of the problem, but give a reasonably accurate indication of the problem's location (distance from the point of the test). There is also a device known as an OTDR, which is an Optical Time-Domain Reflectometer for fiber-optic cables.

Q: What means "BERT"?
A: Bit Error Rate Tester. This equipment is used to analyze the amount and types of errors that occur on a cable segment.

Q: What (free) tools are there to monitor/decode/etc an Ethernet?
A: There are many built into most Unix systems. Some cards for the PC come with utilities. There are several free ones available. Again, use archie.

Q: What is the difference between an Ethernet frame and a IEEE802.3 frame? Why are there two types? Why is there a difference?
A: 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, while NetWare defaults to 802.3.

Q: What is SNAP
A: Sub-Network Access Protocol

Q: Where can I find out which Protocols use which Ethernet type numbers?
A: Look at IETF RFC-1340 - Assigned Numbers RFC.

Q: What is UTP, STP?
A: Unshielded twisted pair, shielded twisted pair. UTP is what the phone companies typically use, though this is not always of high-enough quality for high-speed network use. STP is mostly from IBM. Either one can be used for Ethernet, but they have different electrical characteristics (impedance) and can't be mixed and matched freely. Some manufacturer's hubs and concentrator cards can be bought that will speak to either type of cable, so you CAN hook them together in a manner.

Q: What exactly means 10Base5, 10BaseT, 10Base2, 10Broad36, etc.
A: 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. 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:

10Base-2	Is 10MHz Ethernet running over thin, baseband coax. 10Base-2 is also commonly referred to as thin-Ethernet or Cheapernet.
10Base-5	Is 10MHz Ethernet running over standard (thick) base-band coax.
10Base-F	Is 10MHz Ethernet running over fiber-optic cabling.
10Base-T	Is 10MHz Ethernet running over unshielded, twisted-pair cabling.

Q: Are there any restrictions on how Ethernet is cabled?
A: Yes, there are many, and they vary according to the media used. First of all, there are distance limitations:

10Base-2	limited to 185 meters (607 ft) per unrepeated cable segment.
10Base-5	limited to 500 meters (1,640 ft) per unrepeated cable segment.
10Base-F	depends on the signaling technology and medium used but can go up to 2KM.
10Base-T	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).

Then there are limitations on the number of repeaters and cable segments allowed on a single network. There may be no more than five (5) repeated segments, nor more than four (4) repeaters on any Ethernet; 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 10Base-T 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, 10Base-2 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. 10Base-5 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).

I am not aware of any theoretical limit on the number of 10Base-T devices, and don't know the limitations for 10Base-F yet. (Can someone fill-in the blanks?)

Q: What is 10Base-F?
A: 10Base-F is an IEEE standard for 10mbps Ethernet over fiber-optic cabling. It defines the methodology and standard devices which, ideally, can permit one company's 10Base-F devices to interoperate with any others'.

Q: What means FOIRL?
A: 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 10Base-F standard.

Q: What about wireless LAN's? Are there any?
A: Yes. They typically use reflected or point-to-point infrared light, spread-spectrum RF or microwave RF transmission as as media. They are typically expensive, slow (relative to Ethernet) and are not yet a mature technology. There are special applications for light based (laser) repeaters.

Q: When should I choose 10BaseT, when 10Base2 (or others)?
A: 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 micro-wave. 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.

10Base-2 (thin Ethernet or Cheapernet) is the least expensive way to cable an Ethernet network. However, the price difference between 10Base-2 and 10Base-T (Ethernet over UTP) is rapidly diminishing. Still, for small, budget-conscious installations, 10Base-2 is the most economical topology. The disadvantages of 10Base-2 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).

10Base-5 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 10Base-5 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 10Base-5 too expensive for most LAN installations, and like 10Base-2, a single break or bad connection in the cable can bring the entire network down.

10Base-T is the most flexible topology for LANs, and is generally the best choice for most network installations. 10Base-T 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 10Base-T 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 10Base-2, the cabling is cheaper and requires less skill to install, making 10Base-T installation costs only slightly higher than 10Base-2. The flexibility and reliability more than offset the marginally higher price.

10Base-F, and its predecessor, FOIRL, are the only recommended topologies for inter-building links. However, they need not be limited to this role. 10Base-F 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, 10Base-F) is used inside buildings to form backbone networks and to connect wiring closets together.

Q: What are the advantages/disadvantages of a star like cabling?
A: Old style Ethernet bus wiring (ie, taking the cable from one machine to the next, and then to the next, etc) is prone to cable failure and quickly consumes allowed distances due to aesthetic wiring needs. If the wiring connection is broken at any point, the entire network (segment) fails - and the much greater number of connections increases the probability of a failure or break. On the other hand, it's pretty easy to do for a layman and may involve less actual wiring for small segments.

Star wiring eliminates the single point of failure of a common wire. A central hub has many connections that radiate out to hosts, if one of these hosts connections fails it usually doesn't affect the others. Obviously, however, the hub becomes a central point of failure itself, but studies show a quality hub is less likely to fail before a heavily used strand of coax.

There are a bunch of other reasons hubs are desirable, but this is the biggie.

Q: Is there an official "standard" punch down scheme for 10BaseT?
A: Get a copy of EIA-568, it covers all of that sort of stuff: horizontal, vertical, connectors, patch cords, cross-connects, etc.

Q: Is it safe to run Unshield Twisted Pair next to power cable (it is shielded)?

A: 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 10Base-T. 10Base-T 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.

Q: Why has the MAC address to be unique?
A: 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.

Q: Is there a special numbering scheme for MAC addresses?

A: 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-1340 (available via FTP) lists some of the manufacturer-assigned MAC addresses.

Q: What is a "segment"?
A: A piece of wire bounded by bridges, routers, or terminators. Some people consider wires on either side of a repeater separate segments, but they aren't really.

Q: What is a "subnet"?
A: 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 soviet subs.

Q: What is a fan-out? Is this device still used?
A: Fanout (a.k.a transceiver multiplexor, a.k.a. multiport transceiver, a.k.a. DELNI) allows multiple stations to connect to a single transceiver or transceiver-like device. They are still widely used.

Q: What means "AUI"?
A: Attachment Unit Interface, an IEEE term for the connection between a controller and the transceiver.

Q: What is a transceiver?
A: 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.

Q: What means "MAU"?
A: Medium Access Unit, an IEEE term for a transceiver. MAU is also commonly [mis]used to describe a Token-Ring Multi-Station Access Unit (MSAU). Refer to HUB for an explanation of MSAU.

Q: What exactly does a repeater?
A: 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.

Q: What is a "HUB"?
A: A hub is a common wiring point for star-topology networks, and is a common synonym for concentrator (though the latter generally has additional features or capabilities). Arcnet, 10Base-T Ethernet and 10Base-F Ethernet and many proprietary network topologies use hubs to connect multiple cable runs in a star-wired network topology into a single network. Token-Ring MSAUs (Multi-Station Access Units) can also be considered a type of hub, but don't let a token-ring bigot hear that. Hubs have multiple ports to attach the different cable runs. Some hubs (such as 10Base-T and active ArcNet) include electronics to regenerate and retime the signal between each hub port. Others (such as 10Base-F or passive Arcnet) simply act as signal splitters, similar to the multi-tap cable-TV splitters you might use on your home antenna coax (of course, 10Base-F uses mirrors to split the signals between cables). Token-Ring MSAUs use relays (mechanical or electronic) to reroute the network signals to each active device in series, while all other hubs redistribute received signals out all ports simultaneously, just as a 10Base-2 multi-port repeater would.

Q: What exactly does a bridge?
A: 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.

Q: What does a "learning bridge"?
A: 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.

Q: What is a remote bridge?
A: A bridge as described above that has an Ethernet (or token-ring) 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.

Q: What exactly does a router?

A: 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. Routing would always be preferable to bridging except for the fact 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).

Q: So should I use a router or a bridge?
A: 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. The following are the pros and cons of each:

Routing:
+ Can route between different media (although FDDI to Ethernet bridges are becoming common via the Translation Bridging standard).
+ There is isolation of Multicast & Broadcast packets at the MAC layer which helps to reduce broadcast storms.
+ Can run multiple active paths between sites in a mesh network to use links efficiently (bridging uses spanning tree to decide if a link is forwarding or in a back up state).
+ Takes part in higher level protocol so can provide more features (examples = logical zones in Appletalk, proxy ARP on IP).
+ Provide a clean cut off when connecting multiple management domains.
+ Only needs to know 'where next?' and so hides the detail of remote networks, whereas bridges must understand the whole topology of the net.

Bridging:
+ Much cheaper boxes.
+ Learning bridges virtually autoconfigure themselves.
+ Works with any protocol that conforms to the MAC level spec. some protocols such as DEC LAT & MOP can only be bridged.
+ Within a site uses IP address space more efficiently whilst providing some traffic segregation (address space is becoming a real scarce resource!).
+ Bridges are generally less complex devices, which usually translates to higher reliability.
+ Easy inter-vendor working via spanning tree standard (802.1d or DEC STP)

Q: Are there problems mixing Bridging & routing?
A: You should 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.

Q: What is a Kalpana EtherSwitch?
A: A device that works sort of like a bridge, but off a different principle. It's advantages are that it is extremely fast and can "bridge" more than one packet at a time (it is not limited to two interfaces as a traditional bridge is). Disadvantages are that it does not understand spanning tree and doesn't work well in many to one networks. You probably don't understand that, so ignore it.

Q: What is a driver?
A: 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. There are many other meanings to this word, but this is probably what you are looking for.

Q: What is NDIS, packet driver, ODI.?
A: NDIS is a Microsoft/3com puppy that allows "stacking" of multiple protocols for a single underlying driver. Essentially it allows a single Ethernet card in a PC (it's not limited to Ethernet) to speak many different network "languages", and usually at the same time.

A packet driver is another method of allowing multiple protocols to access the network interface at the same time. Developed and supported by FTP Software Inc, Clarkson University, BYU and, more recently, Crynwr Software, the packet driver spec (PDS) is used to provide a device independent interface to various TCP/IP applications, and often in combination with concurrent Novell access (IPX/SPX).

ODI is Novell and Apple's equivalent of NDIS. There are differences between the two specs, but not so much as to warrant description in this text.

The next logical question is "which one should I use?" There is no simple or obvious answer, except that you should use the one most commonly required by your software.

Q: Is there a troubleshooting guide for Ethernet?
A: Many. I suggest you check your local technical bookstore.
(Recommendations needed)

Q: What books are good about Ethernet LAN's?
A: There are many. The following are recommended by readers on this list:

"The Ethernet Management Guide - Keeping the Link" by Martin Nemzow. This book has good coverage of most of the average considerations of Ethernet, from what Manchester encoding is down to production segment traffic analysis.

Q: Where can I get IEEE803.x docs online?
A: Nowhere. IEEE documents must be ordered from the IEEE themselves. You can contact them at:

Institute of Electrical and Electronic Engineers
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
U.S.A.
(800) 678-IEEE

Q: Where can I get EIA/TIA docs online?
A: Nowhere? Must be ordered from:
Global Engineering
2805 McGaw Av
Irvine, CA 92714
phone 714-261-1455

Q: Where can I find the specifications of Ethernet equipment?
A: From the manufacturer of the product, probably.

Q: Where can I find IETF (Internet Engineering Task Force) documents?
A: These are available for anonymous FTP from a number of sites. One

known location is athos.rutgers.edu in /ietf. Drafts are also on athos in /internet-drafts.

--

RUCS User Services | Mark A. Medici, Systems Programmer III, User Services Division
Rutgers University Computing Services, New Brunswick, NJ 08903
[medici@gandalf.rutgers.edu] [908-932-2412]