

Network Working Group
Request for Comments: 1207
FYI: 7

G. Malkin
FTP Software, Inc.
A. Marine
SRI
J. Reynolds
ISI
February 1991

FYI on Questions and Answers
Answers to Commonly asked "Experienced Internet User" Questions

Status of this Memo

This FYI RFC is one of two FYI's called, "Questions and Answers" (Q/A), produced by the User Services Working Group of the Internet Engineering Task Force (IETF). The goal is to document the most commonly asked questions and answers in the Internet.

This memo provides information for the Internet community. It does not specify any standard. Distribution of this memo is unlimited.

Table of Contents

1. Introduction.....	1
2. Acknowledgements.....	3
3. Questions about the Internet.....	3
4. Questions About Other Networks and Internets.....	3
5. Questions About Internet Documentation.....	4
6. Questions About the Domain Name System (DNS).....	4
7. Questions About Network Management.....	7
8. Questions about Serial Line Internet Protocol (SLIP) and Point-to-Point Protocol (PPP) Implementations.....	9
9. Questions About Routing.....	11
10. Other Protocol and Standards Implementation Questions.....	11
11. Suggested Reading.....	12
12. References.....	13
13. Security Considerations.....	14
14. Authors' Addresses.....	15

1. Introduction

During the last few months, several people have monitored various major mailing lists and have extracted questions that are important or commonly asked. This FYI RFC is one of two in a series of FYI's which present the questions and their answers. The first FYI, FYI 4, presented questions new Internet users commonly ask and their answers.

The goal of this FYI is to codify the Internet lore so that network operations staff, especially for networks just joining the Internet, will have an accurate and up to date set of references from which to work. Also, redundancies are moved away from the electronic mailing lists so that the lists' subscribers do not have to read the same queries and answers over and over again.

Although the questions and their responses are taken from various mailing lists, they are presented here loosely grouped by related topic for ease of reading. First the question is presented, then the answer (or answers) as it appeared on the mailing list.

Sometimes the answers are abridged for better use of space. If a question was not answered on the mailing list, the editors provide an answer. These answers are not distinguished from the answers found on the lists. Sometimes, in order to be as complete as possible, the editors provide additional information that was not present in the original answer. If so, that information falls under the heading "Additional Information".

The answers are as correct as the reviewers can make them. However, much of this information changes with time. As the FYI is updated, temporal errors will be corrected.

Many of the questions are in first person, and the answers were directed to the originator of the question. These phrasings have not been changed except where necessary for clarity. References to the correspondents' names have been removed.

The Q/A mailing lists are maintained by Gary Malkin at FTP.COM. They are used by a subgroup of the User Services Working Group to discuss the Q/A FYIs. They include:

quail@ftp.com	This is a discussion mailing list. Its primary use is for pre-release review of the Q/A FYIs.
quail-request@ftp.com	This is how you join the quail mailing list.
quail-box@ftp.com	This is where the questions and answers will be forwarded-and-stored. It is not necessary to be on the quail mailing list to forward to the quail-box.

2. Acknowledgments

The following people deserve thanks for their help and contributions to this FYI Q/A: Jim Conklin (EDUCOM), John C. Klensin (MIT), Professor Kynikos (Special Consultant), Jon Postel (ISI), Marshall Rose (PSI, Inc.), David Sitman (Tel Aviv University), Patricia Smith (Merit), Gene Spafford (Purdue), and James Van Bokkelen (FTP Software, Inc.).

3. Questions about the Internet

3.1. How do I get statistics regarding the traffic on NSFNET?

Merit/NSFNET Information Services maintains a variety of statistical data at 'nis.nsf.net' (35.1.1.48) in the 'stats' directory. Information includes packet counts by NSS and byte counts for type of use (ftp, smtp, telnet, etc.). Filenames are of the form 'NSFyy-mm.type'.

Files are available for anonymous ftp; use 'guest' as the password.

The data in these files represent only traffic which traverses the highest level of the NSFNET, not traffic within a campus or regional network. Send questions/comments to nsfnet-info@merit.edu.

4. Questions About Other Networks and Internets

4.1. We have a user who would like to access a machine on "EARN/BITNET". I can't find anything on this in the domain name tables. Please, what is this, and how do I connect to it?

There are several machines on the Internet that act as gateways between the Internet and BITNET. Two examples are UICVM.UIC.EDU and CUNYVM.CUNY.EDU. You can address a mail message to user%nodename.bitnet@uicvm.uic.edu where the message will be passed from the Internet to BITNET.

Additional Information:

These same gateways, known as INTERBIT on the BITNET/EARN side, transfer mail from computers on that network which support SMTP mail headers, onto the Internet. (Many BITNET/EARN computers still do not support SMTP, which is not a part of the IBM protocol used, and it is not possible to send mail from those computers across the gateways into the Internet, in general.)

BITNET and EARN are the two largest of several cooperating networks which use the IBM RSCS/NJE protocol suite, but are not limited to IBM systems. These independently administered, interconnected networks function as a single, worldwide network directly connecting more than 3,300 computers in about 1,400, mostly higher-education, organizations worldwide. This worldwide network supports electronic mail, including mailing lists, sender-initiated file transfer, and short "interactive" messages.

BITNET, frequently used (outside of Europe) to refer to the whole worldwide network, technically refers to that portion in the United States, plus sites in other countries which are connected through the United States and do not have their own separately administered cooperating networks. More than 550 organizations in the U.S. participate in BITNET.

EARN is the European Academic Research Network. EARN links more than 500 institutions in Europe and several surrounding countries.

BITNET and CSNET merged organizationally on October 1, 1990, to form CREN, the Corporation for Research and Educational Networking. The two networks remain separate at the operational level level, however. (EARN and the other Cooperating Networks were not involved in this merger.)

5. Questions About Internet Documentation

5.1. Where do I get information regarding ordering documents related to GOSIP?

The complete information as issued by NIST is available online on the NIC.DDN.MIL host as PROTOCOLS:GOSIP-ORDER-INFO.TXT. The file contains pointers to contact people, ordering addresses, prices, and, in some cases, online pathnames, for various GOSIP related documents. In addition, the information as of August 1990 was published as an appendix to RFC 1169, "Explaining the Role of GOSIP" [1].

6. Questions About Domain Name System (DNS)

6.1. Is there a DNS Query server?

Actually, what you are looking for is the service that host 128.218.1.109 provides on port 5555 - you simply connect to that host at that port, type in a fully qualified domain name and it responds with an internet address and closes the connection. I

used it when I had a host that still only had /etc/hosts and it did just what I needed - which was basically a manual nslookup.

However, the vast majority of users will find it simpler to just use a DNS query tool and ask the DNS directly. This doesn't require much sophistication, and does allow the user to see how short names are expanded at the user's site rather than at 128.218.1.109 (wherever that is). For example, suppose a user wants to find out the address of a fully-qualified domain name "X.MISKATONIC.EDU", and also see what host and address are used when "Z" is typed as a host name.

Assuming the user is on a UNIX host and has a copy of the dig program, type:

```
    dig x.miskatonic.edu
and
    dig z
```

and the answers will appear. You are now on your way to becoming a DNS expert. There are other UNIX alternatives, e.g., nslookup, and similar programs for non-UNIX systems. Your local DNS guru certainly has one or more of these tools, and although they are often kept from the public, they are really quite easy to use for simple cases.

- 6.2. We have been having a frequent BIND failure on both our VAX and Solbourne that is traced to TCP domain queries from an IBM NSMAIN nameserver running in cache mode (UDP queries do not cause this problem, though it is usually a UDP resolution that is active upon the crash -- this resolution is an innocent victim).

I have discovered that something is trashing the hash areas (sometimes even as it is being recursively used in a resolution). Also, occasionally the socket/file descriptor for the TCP connection is changed to invalid entries causing a reply write fail (though this is not necessarily fatal, and the rest of the structure is not apparently altered).

Has any one else had frequent BIND failures (especially major domain sites that have heavy TCP domain loads)?

In both the case of BIND and the IBM implementation, often called FAL, there are multiple versions, with older versions being truly bad. Upgrade to recent version before exploring further.

BIND has always had a problem with polluting its own database.

These problems have been related to TCP connections, NS RRs with small TTLs, and several other causes. Experience suggests that the style of bug fixing has often been that of reducing the problem by 90% rather than eliminating it.

IBM's support for the DNS (outside of UNIX systems) is interesting in its techniques, encouraging in its improvement, but still somewhat depressing when compared to most other DNS software. IBM also uses terminology that varies somewhat from the usual DNS usage and preserves some archaic syntax, e.g., "...".

The combination of an old BIND and an old IBM server is just plain unpleasant.

- 6.3. Is the model used by the domain name system for host names that the owner of a name gets to choose its case?

The model used by the DNS is that you get to control at a specific point in the name space, and are hence free to select case as you choose, until points where you in turn give away control. As a practical matter, there are several implementations that don't do the right thing. IBM implementations often map everything into a single case.

- 6.4. According to RFC 1034 [2], section 4.2.1, one should not have to code glue RR's for name server's names unless they are below the cut. When I don't put glue RR's in, and do a query for NS records, the "additional" field is left blank. As far as I can tell, all other zones I query for NS records have this filled with the IP addresses of the NS hosts. Is this required or should I not be concerned that the additional field is empty?

The protocol says that an empty additional field is not a problem when the name server's name is not "below" the cut.

In practice, putting in the glue where it is not required can cause problems if the servers named in the glue are used for several zones. This is broken behavior in BIND. Not putting in glue can cause other problems in BIND, usually when the server name is difficult to resolve. So, the bottom line is to put glue in only when required, and don't use aliases or anything else tricky when it comes to identifying name servers.

7. Questions About Network Management Implementations

- 7.1. In reading the SNMP RFCs [3,4,5,6] I find mention of authentication of PDUs. Are there any standards for authentication mechanisms?

There is a working group of the IETF that is working on this problem. They are close to a solution, but nothing has yet reached RFC publication yet. Expect something solid and implementable by October of 1991.

- 7.2. Can vendors make their enterprise-specific variables available to users through a standard distribution mechanism?

Yes. But before someone submits a MIB, they should check it out themselves.

On uu.psi.com in pilot/snmp-wg/, there are two files

mosy-sparc-4.0.3.c

mosy-sun3-3.5

The first will run on a Sun-Sparc, the second will run on a Sun-3. After retrieving one of these files in BINARY mode via anonymous FTP, the submitter can run their MIB through it, e.g.,

```
% mosy mymib.my
```

Once your MIB passes, send it to:

mib-checker@isi.edu

If everything is OK, the mib-checker will arrange to have it installed in the /share/ftp/mib directory on venera.isi.edu.

Note: This processing does not offer an official endorsement. The documents submitted must not be marked proprietary, confidential, or the like.

- 7.3. I have a question regarding those pesky octet strings again. I use the variable-type field of the Response pdu to determine how the result should be displayed to the user. For example, I convert NetworkAddresses to their dotted decimal format ("132.243.50.4"). I convert Object Identifiers into strings ("1.3.6.1.2....").

I would LIKE to just print Octet Strings as strings. But,

this causes a problem in such cases as `atPhysAddress` in which the Octet string contains the 6 byte address instead of a printable ASCII string. In this case, I would want to display the 6 bytes instead of just trying to print the string.

MY QUESTION IS: Does anyone have a suggestion as to how I can determine whether I can just print the string or whether I should display the octet bytes. * Remember: I want to support enterprise specific variables too.

In general, there is no way that you can tell what is inside an OCTET STRING without knowing something about the object that the OCTET STRING comes from. In MIB-II [6], some objects are marked as `DisplayString` which has the syntax of OCTET STRING but is restricted to characters from the NVT ASCII character set (see the TELNET Specification, RFC 854 [7], for further information). These objects are:

- `sysDescr`
- `sysContact`
- `sysName`
- `sysLocation`
- `ifDescr`

If you want to be able to arbitrarily decide how to display the strings, without knowing anything about the object, then you can scan the octets, looking for any octet which is not printable ASCII. If you find at least one, you can print the entire string, octet by octet, in `"%02x:"` notation. If all of the octets are printable ASCII, then you can just `printf` the string.

- 7.4. If archived MIBs must be 1155-compatible [3], it would be nice if those who submit them check them first. Where are these MIB tools available for public FTP? Ideally, a simple syntax checker (that didn't actually generate code) would be nice.

In the ISODE 6.0 release there is a tool called MOSY which recognizes the 1155 syntax and produces a flat ASCII file. If you can run it through MOSY without problems then you are OK.

- 7.5. Suppose I want to create a private MIB object for causing some action to happen, say, do a reset. Should the syntax or this object specify a value such as:

Syntax:

```
INTEGER {  
    perform reset (1),  
}
```

even though there is only a single value? Or, is it ok to just allow a Set on this object with any value to perform the desired action? If the later, how is this specified?

For our SNMP manageable gizmos and doohickies with similar "action" type MIB variables, I've defined two values

Syntax:

```
INTEGER {  
    reset(1)  
    not-reset(2)  
}
```

And defined behavior so that the only valid value that the variable may be set to is "reset" (which is returned in the get response PDU) and at all other times a get/getnext will respond with "not-reset".

8. Questions about Serial Line Internet Protocol (SLIP) and Point-to-Point Protocol (PPP) Implementations

- 8.1. I seem to recall hearing that SLIP [8] will only run on synchronous serial lines. Is this true? ... is there something about SLIP which precludes it's being implemented over async lines?

Other way around: SLIP is designed for async lines and is not a good fit on sync lines. PPP [9, 10] works on either, and is what you should be implementing if you're implementing something.

- 8.2. Since we are very interested in standards in this area, could someone tell me were I can find more information on PPP?

Also, can this protocol be used in other fields than for the Internet (i.e., telecontrol, telemetry) where we see a profusion of proprietary incompatible and hard to maintain Point-to-Point Protocols?

PPP was designed to be useful for many protocols besides just IP. Whether it would be useful for your particular application should probably be discussed with the IETF's Point-to-Point Protocol Working Group discussion list. For general discussion: ietf-ppp@ucdavis.edu. To subscribe: ietf-ppp-request@ucdavis.edu

The PPP specification is available as RFC 1171 [9], and a PPP options specification is available as RFC 1172 [10].

In UnixWorld of April 1990 (Vol. VII, No. 4, Pg. 85), Howard Baldwin writes:

"Point-to-Point Protocol (PPP) has just been submitted to the CCITT from the Internet Engineering Task Force. It specifies a standard for encapsulating Internet Protocol data and other network layer (level three on ISO's OSI Model) protocol information over point-to-point links; it also provides ways to test and configure lines and the upper level protocols on the OSI Model. The only requirement is a provision of a duplex circuit either dedicated or switched, that can operate in either an asynchronous or synchronous mode, transparent to the data-linklayer frame.

"According to Michael Ballard, director of network systems for Telebit, PPP is a direct improvement upon Serial Line Internet Protocol (SLIP), which had neither error correction nor a way to exchange network address."

- 8.3. Does anyone know if there is a way to run a SLIP program on a IBM computer running SCO Xenix/Unix, with a multi-port serial board?

SCO TCP/IP for Xenix supports SLIP. It works. However, be warned: SCO SLIP works **only** with SCO serial drivers, so it will **not** work with intelligent boards that come with their own drivers. If you want lots of SLIP ports, you'll need lots of dumb ports, perhaps with a multi-dumb-port board.

Here's the setup -- SunOS 3.5, with the 4.3BSD TCP, IP & SLIP distributions installed. Slip is running between the "ttya" ports of two Sun 3/60's. "ping", "rlogin", etc., works fine, but a NFS mount results in "server not responding: RPC Timed Out".

SunOS 3.5 turns the UDP checksum off, which is legal and works okay over interfaces such as ethernet which has link-level checksumming. On the other hand, SLIP doesn't perform checksums thus running NFS over SLIP requires you to turn the UDP checksum on. Otherwise, you'll experience erratic behavior such as the one described above.

Save the older kernel and try,

```
% adb -k -w /vmunix /dev/kmem udpcksum?w 1
```

to patch up the kernel.

9. Questions About Routing

- 9.1. Some postings mentioned "maximum entropy routing". Could someone please provide a pointer to on-line or off-line references to this topic?

Try NYU CSD Technical Report 371: "Some Comments on Highly Dynamic Network Routing," by Herbert J. Bernstein, May 1988.

10. Other Protocol and Standards Implementation Questions

- 10.1. Does anyone recognize ethernet type "80F3"? I don't see it in RFC 1010, but I am seeing it on our net.

Ethernet type 0x80F3 is used by AppleTalk for address resolution. You must have Macs on your network which are directly connected to Ethernet. These packets are used by the Mac (generally at startup) to determine a valid AppleTalk node number.

Additional Information:

RFC 1010 is obsolete. Please consult RFC 1060 [11], the current "Assigned Numbers" (issued March 1990), which does list "80F3":

Ethernet		Exp. Ethernet		Description	References
-----		-----		-----	-----
decimal	Hex	decimal	octal		
33011	80F3	-	-	AppleTalk AARP (Kinetics)[XEROX]	

- 10.2. Does anyone know the significance of a high value for "Bad proto" in the output from netstat on Unix machines using ethernet? We're seeing values in the tens of thousands out of a few hundred thousand packets sent/received in all. Some "Bad proto" values are negative, too. (Off the scale?) Any help would be appreciated.

This probably indicates that you are getting tens of thousands of broadcast packets from some host or hosts on your network. You might want to buy or rent a LAN monitor, or install one of the public-domain packages to see what private protocol is guilty. "FYI on a Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices" (RFC

1147, FYI 2), [12] contains pointers to tools that may help you zero in on the problem.

10.3. Which RFC would explain the proper way to configure broadcast addresses when using subnets?

Consult RFC 1122, "Requirements for Internet Hosts -- Communication Layer" [13].

10.4. Can anyone tell me what .TAR files exactly are? Is it like ZIP or LZH for the IBM PC's? IF so, how do I go about getting a compressor/decompressor for .TAR files and what computer does this run on?

TAR stands for "Tape ARchive". It is a Unix utility which takes files, and directories of files, and creates a single large file. Originally intended to back up directory trees onto tape (hence the name), TAR is also used to combine files for easier electronic file transfer.

11. Suggested Reading

For further information about the Internet and its protocols in general, you may choose to obtain copies of the following works:

Bowers, K., T. LaQuey, J. Reynolds, K. Roubicek, M. Stahl, and A. Yuan, "Where to Start - A Bibliography of General Internetworking Information", RFC 1175, FYI 3, CNRI, U Texas, ISI, BBN, SRI, Mitre, August 1990.

Braden, R., Editor, "Requirements for Internet Hosts -- Communication Layer", RFC 1122, Internet Engineering Task Force, October 1989.

Braden, R., Editor, "Requirements for Internet Hosts -- Application and Support", RFC 1123, Internet Engineering Task Force, October 1989.

Comer, D., "Internetworking with TCP/IP: Principles, Protocols, and Architecture", Prentice Hall, New Jersey, 1989.

Frey, D. and R. Adams, "!!%@:: A Directory of Electronic Mail Addressing and Networks", O'Reilly and Associates, Newton, MA, August 1989.

Krol, E., "The Hitchhikers Guide to the Internet", RFC 1118, University of Illinois Urbana, September 1989.

LaQuey, T., Editor, "Users' Directory of Computer Networks", Digital Press, Bedford, MA, 1990.

Malkin, G., and A. Marine, "FYI on Questions and Answers - Answers to Commonly asked "New Internet User" Questions", RFC 1206, FYI 4, FTP Software, Inc., SRI, February 1991.

Postel, J., Editor, "IAB Official Protocol Standards", RFC 1140, Internet Activities Board, May 1990.

Quarterman, J., "Matrix: Computer Networks and Conferencing Systems Worldwide", Digital Press, Bedford, MA, 1989.

Reynolds, J., and J. Postel, "Assigned Numbers", RFC 1060, USC/Information Sciences Institute, March 1990.

Socolofsky, T., and C. Kale, "A TCP/IP Tutorial", RFC 1180, Spider Systems Limited, January 1991.

Stevens, W., "UNIX Network Programming", ISBN 0-13-949876-1, Prentice Hall, Englewood Cliffs, NJ, 1990.

Stine, R., Editor, "FYI on a Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices" RFC 1147, FYI 2, Sparta, Inc., April 1990.

12. References

- [1] Cerf, V., and K. Mills, "Explaining the Role of GOSIP", RFC 1169, IAB, NIST, August 1990.
- [2] Mockapetris, P., "Domain Names - Concepts and Facilities", RFC 1034, USC/Information Sciences Institute, November 1987.
- [3] Rose, M., and K. McCloghrie, "Structure and Identification of Management Information for TCP/IP-based Internets", RFC 1155, Performance Systems International, Hughes LAN Systems, May 1990.
- [4] McCloghrie, K., and M. Rose, "Management Information Base for Network Management of TCP/IP-based internets", RFC 1156, Hughes LAN Systems, Performance Systems International, May 1990.
- [5] Case, J., M. Fedor, M. Schoffstall, and J. Davin, "A Simple Network Management Protocol (SNMP)", RFC 1157, SNMP Research, Performance Systems International, Performance Systems International, MIT Laboratory for Computer Science, May 1990.
- [6] Rose, M., Editor, "Management Information Base for Network

Management of TCP/IP-based internets: MIB-II", RFC 1158, Performance Systems International, May 1990.

- [7] Postel, J., and J. Reynolds, "TELNET Protocol Specification", RFC 854, USC/Information Sciences Institute, May 1983.
- [8] Romkey, J., "A Nonstandard for Transmission of IP Datagrams over Serial Lines: SLIP", RFC 1055, June 1988.
- [9] Perkins, D., "The Point-to-Point Protocol: A Proposal for Multi-Protocol Transmission of Datagrams Over Point-to-Point Links", RFC 1171, CMU, July 1990.
- [10] Perkins, D., and R. Hobby, "The Point-to-Point Protocol (PPP) Initial Configuration Options", CMU, UC Davis, July 1990.
- [11] Reynolds, J., and J. Postel, "Assigned Numbers", RFC 1060, USC/Information Sciences Institute, March 1990.
- [12] Stine, R., Editor, "FYI on a Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices" RFC 1147, FYI 2, Sparta, Inc., April 1990.
- [13] Braden, R., Editor, "Requirements for Internet Hosts -- Communication Layer", RFC 1122, Internet Engineering Task Force, October 1989.

13. Security Considerations

Security issues are not discussed in this memo.

14. Authors' Addresses

Gary Scott Malkin
FTP Software, Inc.
26 Princess Street
Wakefield, MA 01880

Phone: (617) 246-0900
EMail: gmalkin@ftp.com

April N. Marine
SRI International
Network Information Systems Center
333 Ravenswood Avenue, EJ294
Menlo Park, CA 94025

Phone: (415) 859-5318
EMail: APRIL@nic.ddn.mil

Joyce K. Reynolds
USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695

Phone: (213) 822-1511
EMail: jkrey@isi.edu