

Guide for Internet Standards Writers

Status of this Memo

This document specifies an Internet Best Current Practices for the Internet Community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

Abstract

This document is a guide for Internet standard writers. It defines those characteristics that make standards coherent, unambiguous, and easy to interpret. In addition, it singles out usage believed to have led to unclear specifications, resulting in non-interoperable interpretations in the past. These guidelines are to be used with RFC 2223, "Instructions to RFC Authors".

Table of Contents

1	Introduction	2
2	General Guidelines	2
2.1	Discussion of Security	3
2.2	Protocol Description	4
2.3	Target Audience	5
2.4	Level of Detail	5
2.5	Change Logs	6
2.6	Protocol Versions	6
2.7	Decision History	6
2.8	Response to Out of Specification Behavior	6
2.9	The Liberal/Conservative Rule	7
2.10	Handling of Protocol Options	8
2.11	Indicating Requirement Levels	9
2.12	Notational Conventions	9
2.13	IANA Considerations	10
2.14	Network Management Considerations	10
2.15	Scalability Considerations	10
2.16	Network Stability	11
2.17	Internationalization	11

2.18	Glossary	11
3	Specific Guidelines	12
3.1	Packet Diagrams	12
3.2	Summary Tables	13
3.3	State Machine Descriptions	13
4	Document Checklist	15
5	Security Considerations	16
6	References	16
7	Acknowledgments	18
8	Editor's Address	18
9	Appendix	19
10	Full Copyright Statement	20

1 Introduction

This document is a guide for Internet standard writers. It offers guidelines on how to write a standards-track document with clarity, precision, and completeness. These guidelines are based on both prior successful and unsuccessful IETF specification experiences. These guidelines are to be used with RFC 2223, "Instructions to RFC Authors", or its update. Note that some guidelines may not apply in certain situations.

The goal is to increase the possibility that multiple implementations of a protocol will interoperate. Writing specifications to these guidelines will not guarantee interoperability. However, a recognized barrier to the creation of interoperable protocol implementations is unclear specifications.

Many will benefit from having well-written protocol specifications. Implementers will have a better chance to conform to the protocol specification. Protocol testers can use the specification to derive unambiguous testable statements. Purchasers and users of the protocol will have a better understanding of its capabilities.

For further information on the process for standardizing protocols and procedures please refer to BCP 9/RFC 2026, "The Internet Standards Process -- Revision 3". In addition, some considerations for protocol design are given in RFC 1958, "Architectural Principles of the Internet".

2 General Guidelines

It is important that multiple readers and implementers of a standard have the same understanding of a document. To this end, information should be orderly and detailed. The following are general guidelines intended to help in the production of such a document. The IESG may require that all or some of the following sections appear in a

standards track document.

2.1 Discussion of Security

If the Internet is to achieve its full potential in commercial, governmental, and personal affairs, it must assure users that their information transfers are free from tampering or compromise. Well-written security sections in standards-track documents can help promote the confidence level required. Above all, new protocols and practices must not worsen overall Internet security.

A significant threat to the Internet comes from those individuals who are motivated and capable of exploiting circumstances, events, or vulnerabilities of the system to cause harm. In addition, deliberate or inadvertent user behavior may expose the system to attack or exploitation. The harm could range from disrupting or denying network service, to damaging user systems. Additionally, information disclosure could provide the means to attack another system, or reveal patterns of behavior that could be used to harm an individual, organization, or network. This is a particular concern with standards that define a portion of the Management Information Base (MIB).

Standards authors must accept that the protocol they specify will be subject to attack. They are responsible for determining what attacks are possible, and for detailing the nature of the attacks in the document. Otherwise, they must convincingly argue that attack is not realistic in a specific environment, and restrict the use of the protocol to that environment.

After the document has exhaustively identified the security risks the protocol is exposed to, the authors must formulate and detail a defense against those attacks. They must discuss the applicable countermeasures employed, or the risk the user is accepting by using the protocol. The countermeasures may be provided by a protocol mechanism or by reliance on external mechanisms. Authors should be knowledgeable of existing security mechanisms, and reuse them if practical. When a cryptographic algorithm is used, the protocol should be written to permit its substitution with another algorithm in the future. Finally, the authors should discuss implementation hints or guidelines, e.g., how to deal with untrustworthy data or peer systems.

Security measures will have an impact within the environment that they are used. Perhaps users will now be constrained on what they can do in the Internet, or will experience degradation in the speed of service. The effects the security measures have on the protocol's use and performance should be discussed.

The discussion of security can be concentrated in the Security Considerations section of the document, or throughout the document where it is relevant to particular parts of the specification. An advantage of the second approach is that it ensures security is an integral part of the protocol's development, rather than something that is a follow-on or secondary effort. If security is discussed throughout the document, the Security Considerations section must summarize and refer to the appropriate specification sections. This will insure that the protocol's security measures are emphasized to implementer and user both.

Within the Security Considerations section, a discussion of the path not taken may be appropriate. There may be several security mechanisms that were not selected for a variety of reasons: cost or difficulty of implementation, or ineffectiveness for a given network environment. By listing the mechanisms they did not use and the reasons, editors can demonstrate that the protocol's WG gave security the necessary thought. In addition, this gives the protocol's users the information they need to consider whether one of the non-selected mechanisms would be better suited to their particular requirements.

A document giving further guidance on security topics is in development. Authors should obtain a copy of the completed RFC to help them prepare the security portion of the standard.

Finally, it is no longer acceptable that Security Considerations sections consist solely of statements to the effect that security was not considered in preparing the standard.

Some examples of Security Considerations sections are found in STD 33/RFC 1350, STD 51/RFC 1662, and STD 53/RFC 1939. RFC 2316, "Report of the IAB Security Architecture Workshop", provides additional information in this topic area.

2.2 Protocol Description

Standards track documents must include a description of the protocol. This description must address the protocol's purpose, intended functions, services it provides, and, the arena, circumstances, or any special considerations of the protocol's use.

The authors of a protocol specification will have a great deal of knowledge as to the reason for the protocol. However, the reader is more likely to have general networking knowledge and experience, rather than expertise in a particular protocol. An explanation of it's purpose and use will give the reader a reference point for

understanding the protocol, and where it fits in the Internet. The STD 54/RFC 2328 was recommended to the STDGUIDE working group as providing a good example of this in its "Protocol Overview" section.

The protocol's general description must also provide information on the relationship between the different parties to the protocol. This can be done by showing typical packet sequences.

This also applies to the algorithms used by a protocol. A detailed description of the algorithms or citation of readily available references that give such a description is necessary.

2.3 Target Audience

RFCs have been written with many different purposes, ranging from the technical to the administrative. Those written as standards should clearly identify the intended audience, for example, designers, implementers, testers, help desk personnel, educators, end users, or others. If there are multiple audiences being addressed in the document, the section for each audience needs to be identified. The goal is to help the reader discover and focus on what they have turned to the document for, and avoid what they may find confusing, diverting, or extraneous.

2.4 Level of Detail

The author should consider what level of descriptive detail best conveys the protocol's intent. Concise text has several advantages. It makes the document easier to read. Such text reduces the chance for conflict between different portions of the specification. The reader can readily identify the required protocol mechanisms in the standard. In addition, it makes it easier to identify the requirements for protocol implementation. A disadvantage of concise descriptions is that a reader may not fully comprehend the reasoning behind the protocol, and thus make assumptions that will lead to implementation errors.

Longer descriptions may be necessary to explain purpose, background, rationale, implementation experience, or to provide tutorial information. This helps the reader understand the protocol. Yet, several dangers exist with lengthy text. Finding the protocol requirements in the text is difficult or confusing. The same mechanism may have multiple descriptions, which leads to misinterpretation or conflict. Finally, it is more difficult to comprehend, a consideration as English is not the native language of the many worldwide readers of IETF standards.

One approach is to divide the standard into sections: one describing the protocol concisely, while another section consists of explanatory text. The STD 3/RFC 1122/RFC 1123 and STD 54/RFC 2328 provides examples of this method.

2.5 Change Logs

As a document moves along the standards track, from Proposed to Draft or Draft to Full, or cycles in level, it will undergo changes due to better understanding of the protocol or implementation experience. To help implementers track the changes being made a log showing what has changed from the previous version of the specification is required (see Appendix).

2.6 Protocol Versions

Often the standard is specifying a new version of an existing protocol. In such a case, the authors should detail the differences between the previous version and the new version. This should include the rationale for the changes, for example, implementation experience, changes in technology, responding to user demand, etc.

2.7 Decision History

In standards development, reaching consensus requires making difficult choices. These choices are made through working group discussions or from implementation experience. By including the basis for a contentious decision, the author can prevent future revisiting of these disagreements when the original parties have moved on. In addition, the knowledge of the "why" is as useful to an implementer as the description of "how". For example, the alternative not taken may have been simpler to implement, so including the reasons behind the choice may prevent future implementers from taking nonstandard shortcuts.

2.8 Response to Out of Specification Behavior

A detail description of the actions taken in case of behavior that is deviant from or exceeds the specification is useful. This is an area where implementers often differ in opinion as to the appropriate response. By specifying a common response, the standard author can reduce the risk that different implementations will come in to conflict.

The standard should describe responses to behavior explicitly forbidden or out of the boundaries defined by the specification. Two possible approaches to such cases are discarding, or invoking error-handling mechanisms. If discarding is chosen, detailing the

disposition may be necessary. For instance, treat dropped frames as if they were never received, or reset an existing connection or adjacency state.

The specification should describe actions taken when a critical resource or a performance-scaling limit is exceeded. This is necessary for cases where a risk of network degradation or operational failure exists. In such cases, a consistent behavior between implementations is necessary.

2.9 The Liberal/Conservative Rule

A rule, first stated in STD 5/RFC 791, recognized as having benefits in implementation robustness and interoperability is:

"Be liberal in what you accept, and
conservative in what you send".

Or establish restrictions on what a protocol transmits, but be able to deal with every conceivable error received. Caution is urged in applying this approach in standards track protocols. It has in the past lead to conflicts between vendors when interoperability fails. The sender accuses the receiver of failing to be liberal enough, and the receiver accuses the sender of not being conservative enough. Therefore, the author is obligated to provide extensive detail on send and receive behavior.

To avoid any confusion between the two, recommend that standard authors specify send and receive behavior separately. The description of reception will require the most detailing. For implementations are expected to continue operating regardless of error received. Therefore, the actions taken to achieve that result, need to be laid out in the protocol specification. Standard authors should concern themselves on achieving a level of cooperation that limits network disruption, not just how to survive on the network. The appearance of undefined information or conditions must not cause a network or host failure. This requires specification on how to attempt acceptance of most of the packets. Two approaches are available, either using as much of the packet's content as possible, or invoking error procedures. The author should specify a dividing line on when to take which approach.

A case for consideration is that of a routing protocol, where acceptance of flawed information can cause network failure. For protocols such as this, the specification should identify packets that could have different interpretations and mandate that they be rejected completely or the nature of the attempt to recover some information from them. For example, routing updates that contain

more data than the tuple count shows. The protocol authors should consider whether some trailing data can be accepted as additional routes, or to reject the entire packet as suspect because it is non-conformant.

2.10 Handling of Protocol Options

Specifications with many optional features increase the complexity of the implementation and the chance of non-interoperable implementations. The danger is that different implementations may specify some combination of options that are unable to interoperate with each other.

As the document moves along the standard track, implementation experience shall determine the need for each option. Implementation shall show whether the option should be a mandatory part of the protocol or remain an option. If an option is not implemented as the document advances, it must be removed from the protocol before it reaches draft standard status.

Therefore, options shall only be present in a protocol to address a real requirement. For example, options can support future extensibility of the protocol, a particular market, e.g., the financial industry, or a specific network environment, e.g., a network constrained by limited bandwidth. They shall not be included as a means to "buy-off" a minority opinion. Omission of the optional item shall have no interoperability consequences for the implementation that does so.

One possible approach is to document protocol options in a separate specification. This keeps the main protocol specification clean and makes it clear that the options are not required to implement the protocol. Regardless of whether they appear within the specification or in a separate document, the text shall discuss the full implications of either using the option or not, and the case for choosing either course. As part of this, the author needs to consider and describe how the options are used alongside other protocols. The text must also specify the default conditions of all options. For security checking options the default condition is on or enabled.

There are occasions when mutually exclusive options appear within the protocol. That is, the implementation of an optional feature precludes the implementation of the other optional feature. For clarity, the author needs to state when to implement one or the other, what the effect of choosing one over the other is, and what

problems the implementer or user may face. The choice of one or the other options shall have no interoperability consequences between multiple implementations.

2.11 Indicating Requirement Levels

The BCP 14/RFC 2119, "Key words for use in RFCs to Indicate Requirement Level", defines several words that are necessary for writing a standards track document. Editors of standards track documents must not deviate from the definitions provided as they are intended to identify interoperability requirements or limit potentially harmful behavior. The capitalization of these words is the accepted norm, and can help in identifying an unintentional or unreasonable requirement. These words have been used in several RFCs the first instances being STD 3/RFC 1122/RFC 1123.

2.12 Notational Conventions

Formal syntax notations can be used to define complicated protocol concepts or data types, and to specify values of these data types. This permits the protocol to be written without concern on how the implementation is constructed, or how the data type is represented during transfer. The specification is simplified because it can be presented as "axioms" that will be proven by implementation.

The formal specification of the syntax used should be referenced in the text of the standard. Any extensions, subsets, alterations, or exceptions to that formal syntax should be defined within the standard.

The STD 11/RFC 822 provides an example of this. In RFC 822 (Section 2 and Appendix D) the Backus-Naur Form (BNF) meta-language was extended to make its representation smaller and easier to understand. Another example is STD 16/RFC 1155 (Section 3.2) where a subset of the Abstract Syntax Notation One (ASN.1) is defined.

The author of a standards track protocol needs to consider several things before they use a formal syntax notation. Is the formal specification language being used parseable by an existing machine? If no parser exists, is there enough information provided in the specification to permit the building of a parser? If not, it is likely the reader will not have enough information to decide what the notation means. In addition, the author should remember machine parseable syntax is often unreadable by humans, and can make the specification excessive in length. Therefore, syntax notations cannot take the place of a clearly written protocol description.

2.13 IANA Considerations

The common use of the Internet standard track protocols by the Internet community requires that unique values be assigned to parameter fields. An IETF WG does not have the authority to assign these values for the protocol it developed. The Internet Assigned Numbers Authority (IANA) is the central authority for the assignment of unique parameter values for Internet protocols. The authors of a developing protocol need to coordinate with the IANA the rules and procedures to manage the number space. This coordination needs to be completed prior to submitting the Internet Draft to the standards track.

A document is in preparation that discusses issues related to identifier assignment policy and guidelines on specific text to task IANA with its administration. Standard authors should obtain a copy of it when it is finalized as an RFC.

For further information on parameter assignment and current assignments, authors can reference STD 2, RFC 1700, "Assigned Numbers" (<http://www.iana.org>).

2.14 Network Management Considerations

When relevant, each standard needs to discuss how to manage the protocol being specified. This management process should be compatible with the current IETF Standard management protocol. In addition, a MIB must be defined within the standard or in a companion document. The MIB must be compatible with current Structure of Management Information (SMI) and parseable using a tool such as SMICng. Where management or a MIB is not necessary this section of the standard should explain the reason it is not relevant to the protocol.

2.15 Scalability Considerations

The standard should establish the limitations on the scale of use, e.g., tens of millions of sessions, gigabits per second, etc., and establish limits on the resources used, e.g., round trip time, computing resources, etc. This is important because it establishes the ability of the network to accommodate the number of users and the complexity of their relations. The STD 53/RFC 1939 has an example of such a section. If this is not applicable to the protocol, an explanation of why not should be included.

2.16 Network Stability

A standard should discuss the relationship between network topology and convergence behavior. As part of this, any topology that would be troublesome for the protocol should be identified. Additionally, the specification should address any possible destabilizing events, and means by which the protocol resists or recovers from them. The purpose is to insure that the network will stabilize, in a timely fashion, after a change, and that a combination of errors or events will not plunge the network into chaos. The STD 34/RFC 1058, as an example, has sections which discuss how that protocol handles the affects of changing topology.

The obvious case this would apply to is a routing protocol. However, an application protocol could also have dynamic behavior that would affect the network. For example, a messaging protocol could suddenly dump a large number of messages onto the network. Therefore, editors of an application protocol will have to consider possible impacts to network stability and convergence behavior.

2.17 Internationalization

At one time the Internet had a geographic boundary and was English only. The Internet now extends internationally. Therefore, data is interchanged in a variety of languages and character sets. In order to meet the requirements of an international Internet, a standard must conform to the policies stated in BCP 18/RFC 2277, "IETF Policy on Character Sets and Languages".

2.18 Glossary

Every standards track RFC should have a glossary, as words can have many meanings. By defining any new words introduced, the author can avoid confusing or misleading the implementers. The definition should appear on the word's first appearance within the text of the protocol specification, and in a separate glossary section.

It is likely that definition of the protocol will rely on many words frequently used in IETF documents. All authors must be knowledgeable of the common accepted definitions of these frequently used words. FYI 18/RFC 1983, "Internet Users' Glossary", provides definitions that are specific to the Internet. Any deviation from these definitions by authors is strongly discouraged. If circumstances require deviation, an author should state that he is altering the commonly accepted definition, and provide rationale as to the necessity of doing so. The altered definition must be included in the Glossary section.


If the author uses the word as commonly defined, she does not have to include the definition in the glossary. As a minimum, FYI 18/RFC 1983 should be referenced as a source.

3 Specific Guidelines


The following are guidelines on how to present specific technical information in standards.

3.1 Packet Diagrams

Most link, network, and transport layer protocols have packet descriptions. Packet diagrams included in the standard are very helpful to the reader. The preferred form for packet diagrams is a sequence of long words in network byte order, with each word horizontal on the page and bit numbering at the top:



In cases where a packet is strongly byte-aligned rather than word-aligned (e.g., when byte-boundary variable-length fields are used), display packet diagrams in a byte-wide format. The author can use different height boxes for short and long words, and broken boxes for variable-length fields:



3.2 Summary Tables

The specifications of some protocols are particularly lengthy, sometimes covering a hundred pages or more. In such cases, the inclusion of a summary table can reduce the risk of conformance failure by an implementation through oversight. A summary table itemizes what in a protocol is mandatory, optional, or prohibited. Summary tables do not guarantee conformance, but serve to assist an implementer in checking that they have addressed all protocol features.

The summary table will consist of, as a minimum, four (4) columns: Protocol Feature, Section Reference, Status, and References/Footnotes. The author may add columns if they further explain or clarify the protocol.

In the Protocol Feature column, list the protocol's characteristics, for example, a command word. We recommend grouping series of related transactions under descriptive headers, for example, RECEPTION.

Section reference directs the implementer to the section, paragraph, or page that describes the protocol feature in detail.

Status indicates whether the feature is mandatory, optional, or prohibited. The author can use either a separate column for each possibility, or a single column with appropriate codes. These codes need to be defined at the start of the summary table to avoid confusion. Possible status codes:

M	- must or mandatory
MN	- must not
O	- optional
S	- should
SN	- should not
X	- prohibited

In the References/Footnotes column authors can point to other RFCs that are necessary to consider in implementing this protocol feature, or any footnotes necessary to explain the implementation further.

The STD 3/RFC 1122/RFC 1123 provides examples of summary tables.

3.3 State Machine Descriptions

A convenient method of presenting a protocol's behavior is as a state-machine model. That is, a protocol can be described by a series of states resulting from a command, operation, or transaction. State-machine models define the variables and constants that

establish a state, the events that cause state transitions and the actions that result from those transitions. Through these models, an understanding of the protocol's dynamic operation as sequence of state transitions that occur for any given event is possible. State transitions can be detailed by diagrams, tables, or time lines.

Note that state-machine models are never to take the place of detailed text description of the specification. They are adjuncts to the text. The protocol specification shall always take precedence in the case of a conflict.

When using a state transition diagram, show each possible protocol state as a box connected by state transition arcs. The author should label each arc with the event that causes the transition, and, in parentheses, any actions taken during the transition. The STD 5/RFC 1112 provides an example of such a diagram. As ASCII text is the preferred storage format for RFCs, only simple diagrams are possible. Tables can summarize more complex or extensive state transitions.

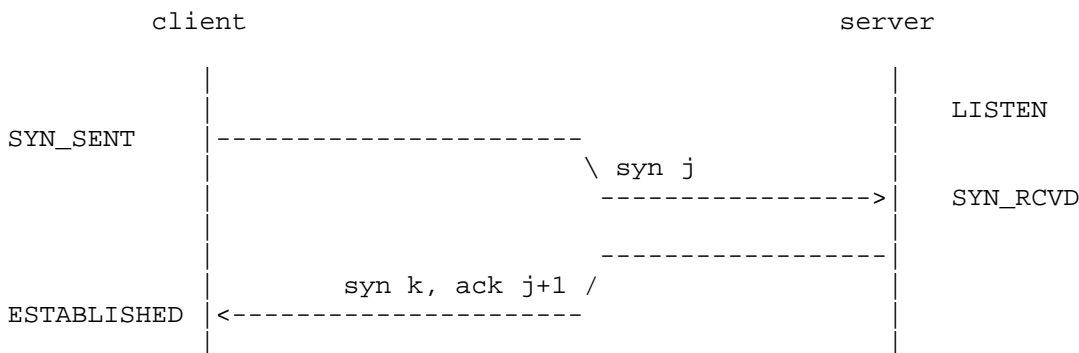
In a state transition table, the different events are listed vertically and the different states are listed horizontally. The form, action/new state, represents state transitions and actions. Commas separate multiple actions, and succeeding lines are used as required. The authors should present multiple actions in the order they must be executed, if relevant. Letters that follow the state indicate an explanatory footnote. The dash ('-') indicates an illegal transition. The STD 51/RFC 1661 provides an example of such a state transition table. The initial columns and rows of that table follow as an example:

Events	State					
	0 Initial	1 Starting	2 Closed	3 Stopped	4 Closing	5 Stopping
Up	2	irc,scr/6	-	-	-	-
Down	-	-	0	tls/1	0	1
Open	tls/1	1	irc,scr/6	3r	5r	5r
Close	0	tlf/0	2	2	4	4
TO+	-	-	-	-	str/4	str/5
TO-	-	-	-	-	tlf/2	tlf/3

The STD 18/RFC 904 also presents state transitions in table format. However, it lists transitions in the form n/a, where n is the next state and a represents the action. The method in RFC 1661 is preferred as new state logically follows action. In addition, RFC 904's Appendix C models transitions as the Cartesian product of two state machines. This is a more complex representation that may be

difficult to comprehend for those readers that are unfamiliar with the format. We recommend that authors present tables as defined in the previous paragraph.

A final method of representing state changes is by a time line. The two sides of the time line represent the machines involved in the exchange. The author lists the states the machines enter as time progresses (downward) along the outside of time line. Within the time line, show the actions that cause the state transitions. An example:



4 Document Checklist

The following is a checklist based on the above guidelines that can be applied to a document:

- o Does it identify the security risks? Are countermeasures for each potential attack provided? Are the effects of the security measures on the operating environment detailed?
- o Does it explain the purpose of the protocol or procedure? Are the intended functions and services addressed? Does it describe how it relates to existing protocols?
- o Does it consider scaling and stability issues?
- o Have procedures for assigning numbers been coordinated with IANA?
- o Does it discuss how to manage the protocol being specified? Is a MIB defined?
- o Is a target audience defined?
- o Does it reference or explain the algorithms used in the protocol?
- o Does it give packet diagrams in recommended form, if applicable?
- o Is there a change log?
- o Does it describe differences from previous versions, if applicable?
- o Does it separate explanatory portions of the document from requirements?
- o Does it give examples of protocol operation?

- o Does it specify behavior in the face of incorrect operation by other implementations?
- o Does it delineate which packets should be accepted for processing and which should be ignored?
- o If multiple descriptions of a requirement are given, does it identify one as binding?
- o How many optional features does it specify? Does it separate them into option classes?
- o Have all combinations of options or option classes been examined for incompatibility?
- o Does it explain the rationale and use of options?
- o Have all mandatory and optional requirements be identified and documented by the accepted key words that define Internet requirement levels?
- o Does it conform to the current internationalization policies of the IETF?
- o Are the recommended meanings for common Internet terms used?
- o If not, are new or altered definitions for terms given in a glossary?

5 Security Considerations

This document does not define a protocol or procedure that could be subject to an attack. It establishes guidelines for the information that should be included in RFCs that are to be submitted to the standards track. In the area of security, IETF standards authors are called on to define clearly the threats faced by the protocol and the way the protocol does or does not provide security assurances to the user.

6 References

- [RFC 791] Postel, J., "Internet Protocol (IP)", STD 5, RFC 791 September 1981.
- [RFC 904] Mills, D., "Exterior Gateway Protocol formal specification", RFC 904, April 1984.
- [RFC 1058] Hedrick, C., "Routing Information Protocol", STD 34, RFC 1058, June 1988.
- [RFC 1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.
- [RFC 1122] Braden, R., "Requirements for Internet Hosts -- Communication Layers", STD 3, RFC 1122, October 1989.

- [RFC 1123] Braden, R., "Requirements for Internet hosts -- Application and Support", STD 3, RFC 1123, October 1989.
- [RFC 1311] Postel, J., "Introduction to the STD Notes", RFC 1311, March 1992.
- [RFC 1350] Sollins, K., "The TFTP Protocol (Revision 2)", STD 33, RFC 1350, July 1992.
- [RFC 1661] Simpson, W., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.
- [RFC 1662] Simpson, W., "PPP in HDLC-like Framing", STD 51, RFC 1662, July 1994.
- [RFC 1700] Reynolds, J., and J. Postel, "Assigned Numbers", STD 2, RFC 1700, October 1994. (<http://www.iana.org>)
- [RFC 1939] Meyers, J., and M. Rose, "Post Office Protocol - Version 3", STD 53, RFC 1939, May 1996.
- [RFC 1958] Carpenter, B., "Architectural Principles of the Internet", RFC 1958, June 1996.
- [RFC 1983] Malkin, G., "Internet Users' Glossary", FYI 18, RFC 1983, August 1996.
- [RFC 2026] Bradner, S., "The Internet Standards Process -- Revision 3", RFC 2026, October 1996.
- [RFC 2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Level", BCP 14, RFC 2119, March 1997.
- [RFC 2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [RFC 2223] Postel, J. and J. Reynolds, "Instructions to RFC Authors", RFC 2223, October 1997.
- [RFC 2277] Alvestrand, H., "IETF Policy on Character Sets and Language", RFC 2277, January 1998.
- [RFC 2316] Bellovin, S., "Report of the IAB Security Architecture Workshop", RFC 2316, April 1998.

7 Acknowledgments

Peter Desnoyers and Art Mellor began the work on this document.
Others that contributed were:

Bernard Aboba
Harald T. Alvestrand
Fred Baker
Scott Bradner
Brian Carpenter
Robert Elz
Dirk Fieldhouse
Dale Francisco
Gary Malkin
Neal McBurnett
Thomas Narten
Craig Partridge
Vern Paxson
Mike O'Dell
Henning Schulzrinne
Kurt Starsinic
James Watt

8 Editor's Address

Gregor D. Scott
Director, Defense Information Systems Agency
ATTN: JIEO-JEBBC
Ft. Monmouth, NJ 07703-5613
USA

Phone: (732) 427-6856
Fax: (732) 532-0853
EMail: scottg@ftm.disa.mil

9 Appendix

CHANGES FROM DRAFT -06

The following changes were made following IESG review:

References to RFC 1543 were changed to RFC 2223 that obsoleted it.

In section 2.1, "export control" was dropped as a valid reason for not selecting a security mechanism. In addition, ambiguous or conflicting sentences were removed.

In section 2.1 reference made to RFC 2315 as an additional source of information.

Section 2.5 was changed to highlight the Change Log's purpose as assistance to implementers.

The IANA Considerations section (2.13) was rewritten to highlight that the IANA guidelines document is work in progress but should be used when it becomes available.

Section 3.4 Character Sets was deleted and replaced by section 2.17 Internationalization.

Spelling and grammar corrections were made.

CHANGES FROM DRAFT -05

A sentence pointing to a pending document that further addresses IANA considerations was added to section 2.13. The current draft of that document is draft-iesg-iana-considerations-02.txt. A clause stating that the IANA established the assignment policies was removed since it appeared to conflict with the intent of the referenced ID. Placeholders for the BCP and RFC number have been added to the text and reference section.

A new section (2.5) requiring change logs as documents progress along the standards track was added.

References to RFC 2044 were changed to RFC 2279 that obsoleted it.

Spelling and grammar corrections were made.

CHANGES FROM DRAFT -04

A paragraph pointing to a pending document that further addresses security was updated.

10 Full Copyright Statement

Copyright (C) The Internet Society (1998). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.