

Guide for Internet Standards Writers
<[draft-ietf-stdguide-ops-06.txt](#)>

Status of this Memo

This document is an Internet Draft. Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a "working draft" or "work in progress."

To learn the current status of any Internet-Draft, please check the "l1d-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ds.internic.net (US East Coast), nic.nordu.net (Europe), ftp.isi.edu (US West Coast), or munnari.oz.au (Pacific Rim).

Distribution of this document is unlimited.

This Internet Draft expires on 15 September 1998.

Abstract

This document is a guide for Internet standard writers. It defines those characteristics that make standards coherent, unambiguous, and easy to interpret. Also, 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 1543](#), "Instructions to RFC Authors."

This version of the document is a draft.

CHANGES FROM DRAFT -04

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

References to [RFC 1583](#) were changed to [RFC 2178](#) which obsoleted it.

Editorial changes and corrections were also 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. Place holders 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](#) which obsoleted it.

Spelling and grammar corrections were made.

Table of Contents

1	Introduction	3
2	General Guidelines	3
2.1	Discussion of Security	3
2.2	Protocol Description	5
2.3	Target Audience	6
2.4	Level of Detail	6
2.5	Change Logs	7
2.6	Protocol Versions	7
2.7	Decision History	7
2.8	Response to Out of Specification Behavior	7
2.9	The Liberal/Conservative Rule	8
2.10	Handling of Protocol Options	9
2.11	Indicating Requirement Levels	10
2.12	Notational Conventions	10
2.13	IANA Considerations	11
2.14	Network Management Considerations	11
2.15	Scalability Considerations	11
2.16	Network Stability	12
2.17	Glossary	12
3	Specific Guidelines	13
3.1	Packet Diagrams	13
3.2	Summary Tables	14
3.3	State Machine Descriptions	14
3.4	Character Sets	16
4	Document Checklist	16
5	Security Considerations	17
6	References	18

7	Acknowledgments	19
8	Editor's Address	19

[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 1543](#), "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. Implementors 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." Also, 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 implementors 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. For an implementor will find it easier to provide the security measures specified. While users will understand the security measures, and so have a higher level of trust in the Internet. Above all, new protocols and practices must not worsen overall Internet security.

A significant threat to the Internet are those individuals who are motivated and capable of exploiting circumstances, events, or vulnerabilities of the system to cause harm. Also, 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.

This discussion of the threat model and other assumptions should appear early in the standard. Doing so will establish a basis for the further discussion of security throughout the document.

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.

Additionally, the effects the security measures have on the protocol's use and performance should be discussed. Security measures will have an impact on the environment they are used in. Perhaps users will now be locked out of portions of the Internet

previously open to them, or users will experience a degradation in the speed of service. The user may decide to accept a greater risk in exchange for improved access or service. But the user must be able to make an informed decision. They need to understand the risks they are facing and the costs of reducing their risk.

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 make reference to the appropriate specification sections. This will insure that the protocol's security measures are emphasized to implementor 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; ineffectiveness for a given network environment; or export control. By listing the mechanisms they did not use and the reasons, editors can demonstrate that the protocol's WG gave security the necessary thought. Also, 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.

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 Draft Standard [RFC 2178](#) was recommended to the STDGUIDE working guide as providing a good example of this in it's "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, implementors, testers, help desk personnel, educators, end users, or others. If there are multiple audiences being addressed in the document, what sections are 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. Also, 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 Draft Standard [RFC 2178](#) 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 it in tracking the progress being made, the IESG requires each document to have a log showing what has changed from the previous version of the specification.

[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. Also, the knowledge of the "why" is as useful to an implementor 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 implementors from taking nonstandard shortcuts.

[2.8](#) Response to Out of Specification Behavior

Detail descriptions of the actions taken in case of behavior that is deviant from or exceeds the specification is useful. This is an area

where implementors 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 critical resource or performance scaling limits are 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 will be expected to accept any packet from the network without failure or malfunction. 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 differing interpretations and mandate that they be either 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. Doing so would make it clear that the options are not integral to the implementation of the protocol, and would keep the main protocol specification clean. 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 intended to be 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 may be occasions when mutually exclusive options appear within a 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 implementor 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. Also, 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 is working on. The Internet Assigned Numbers Authority (IANA) is the central coordinator for the assignment of unique parameter values for Internet protocols. The authors of a developing protocol that use a link, socket, port, protocol, etc., need to coordinate with the IANA the rules and procedures to be used to register constants and tags. This coordination needs to be completed prior to submitting the internet draft to the standards track. For further information on parameter assignment and current assignments, authors can reference STD 2/RFC 1700, "Assigned Numbers." The BCP /RFC , "Guidelines for Writing an IANA Considerations Section in RFCs," gives further guidance on determining identifier assignment policy and providing IANA with instructions to administer that policy.

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. Also 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 which would be troublesome for the protocol should be identified. Additionally, the specification should address any possible destabilizing events, and how 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 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 implementer. 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:

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Version| Prio. |                               Flow Label           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

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:

```

      0 1 2 3 4 5 6 7
+---+---+---+---+---+
|      Length N      |
+---+---+---+---+---+
|                      |
+   Address          +
      ...
+   (N bytes)       +
|                      |
+---+---+---+---+---+
|                      |
+ 2-byte field      +
|                      |
+---+---+---+---+---+

```


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 implementor 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 describe the feature, for example, a command word. We recommend grouping series of related transactions under descriptive headers, for example, RECEPTION.

Section reference directs the implementor 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 either use 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.

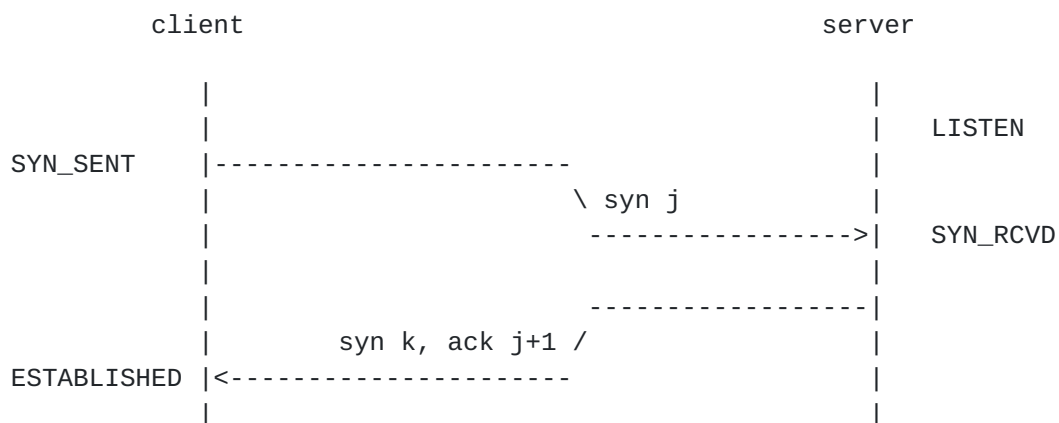
In a state transition table, read events vertically and states horizontally. The form, action/new state, represents state transitions and actions. Commas separate multiple actions, and succeeding lines are used as required. The authors should present multiple actions in the order they must be executed, 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:

	State					
	0	1	2	3	4	5
Events	Initial	Starting	Closed	Stopped	Closing	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
T0+	-	-	-	-	str/4	str/5
T0-	-	-	-	-	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. Also, [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:



3.4 Character Sets

At one time the Internet had a geographic boundary and was English only. Since the Internet now extends internationally, application protocols must assume that the contents of any text string may be in a language other than English. Therefore, new or updated protocols which transmit text must use ISO 10646 as the default Coded Character Set, and RFC 2279, "UTF-8, a transformation format of Unicode and ISO 10646" as the default Character Encoding Scheme. An exception is the use of US-ASCII for a protocol's controlling commands and replies. Protocols that have a backwards compatibility requirement should use the default of the existing protocol. This is in keeping with the recommendations of [RFC 2130](#), "The Report of the IAB Character Set Workshop held 29 February - 1 March 1996."

4 Document Checklist

The following is a checklist based on these 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 use the recommended Internet meanings for any terms use to specify the protocol?
- o 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](#) "Internet Protocol (IP)," J. Postel, September 1981.
- [RFC 904](#) "Exterior Gateway Protocol formal specification," D. Mills, April 1984
- [RFC 1058](#) "Routing Information Protocol," C. Hedrick, June 1988
- [RFC 1112](#) "Host extensions for IP multicasting," S. Deering, August 1989
- [RFC 1122](#) "Requirements for Internet Hosts -- Communication Layers," R. Braden, October 1989
- [RFC 1123](#) "Requirements for Internet hosts -- Application and Support," R. Braden, October 1989
- [RFC 1311](#) "Introduction to the STD Notes," J. Postel, March 1992
- [RFC 1350](#) "The TFTP Protocol (Revision 2)," K. Sollins, July 1992
- [RFC 1661](#) "The Point-to-Point Protocol (PPP)," W. Simpson, July 1994
- [RFC 1662](#) "PPP in HDLC-like Framing," W. Simpson, July 1994
- [RFC 1700](#) "Assigned Numbers," J. Reynolds, J. Postel, October 1994
- [RFC 1939](#) "Post Office Protocol - Version 3," J. Meyers, M. Rose, May 1996
- [RFC 1958](#) "Architectural Principles of the Internet," B. Carpenter, June 1996
- [RFC 1983](#) "Internet Users' Glossary," G. Malkin, August 1996
- [RFC 2026](#) "The Internet Standards Process -- Revision 3," S. Bradner, October 1996
- [RFC 2119](#) "Key words for use in RFCs to Indicate Requirement Level," S. Bradner, March 1997
- [RFC 2130](#) "The Report of the IAB Character Set Workshop held 29 February - 1 March 1996," C. Weider, C. Preston, K. Simonsen, H. Alvestrand, R. Atkinson, M. Crispin, P. Svanberg, April 1997

[RFC 2178](#) "OSPF Version 2," J. Moy, July 1997

[RFC 2279](#) "UTF-8, a transformation format of ISO 10646," F. Yergeau,
January 1998

RFC "Guidelines for Writing an IANA Considerations Section in
RFCs,"

[7](#) Acknowledgments

Peter Desnoyers and Art Mellor began the work on this document. The area directors that oversaw the STDGUIDE WG's efforts were Scott Bradner, Mike O'Dell, and John Curran. Others that contributed to this document were:

Bernard Aboba
Harald T. Alvestrand
Fred Baker
Brian Carpenter
Robert Elz
Dirk Fieldhouse
Dale Francisco
Gary Malkin
Neal McBurnett
Craig Partridge
Henning Schulzrinne
Kurt Starsinic
James Watt

[8](#) Editor's Address

Gregor D. Scott
Director, Defense Information Systems Agency
ATTN: JIE0-JEBBC
Ft. Monmouth, NJ 07703-5613
USA

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

This Internet Draft expires on 15 September 1998.

