AO88 (Rev. 12/06) Subpoena in a Civil Case

Issued by the
UNITED STATES DISTRICT COURT
Northern District of California

SMARTPHONE TECHNOLOGIES LLC,

Plaintiff,

v.

HUAWEI DEVICE USA INC. et al.,

Defendants.

Civil Action No. 6:12-cv-245
United States District Court for the Eastern District of Texas

SUBPOENA IN A CIVIL CASE

TO: The Internet Society
1775 Wiehle Ave
Suite 201
Reston, VA 20190

□ YOU ARE COMMANDED to appear in the United States District Court at the place, date, and time specified below to testify in the above case.		
PLACE OF TESTIMONY	COURTROOM DATE AND TIME	
☒ YOU ARE COMMANDED to appear at the place, date, and time specified below to testify at the taking of a deposition in the above case. Any organization that is subpoenaed for the taking of a deposition shall designate one or more officers, directors, or managing agents, or other persons who consent to testify on its behalf and may set forth, for each person designated, the matters on which the person will testify. Federal Rule of Civil Procedure 30(b)(6).		
PLACE OF DEPOSITION Jones Day, 51 Louisiana Avenue, N.W., Washington, D.C. 200001-2113	DATE AND TIME November 13, 2013, 9.00 a.m.	
☒ YOU ARE COMMANDED to produce and permit inspection and copying of the following documents or objects at the place, date, and time specified below (List documents or objects): Documents identified in Exhibit "A" attached to this subpoena.		
PLACE OF PRODUCTION Jones Day, 51 Louisiana Avenue, N.W., Washington, D.C. 200001-2113	DATE AND TIME November 6, 2013	
□ YOU ARE COMMANDED to permit inspection of the following premises at the date and time specified below.		
PREMISES	DATE AND TIME	
ISSUING OFFICER SIGNATURE AND TITLE Attorney for Huawei Device USA Inc.	DATE October 23, 2013	
ISSUING OFFICER'S NAME, ADDRESS AND PHONE NUMBER Keith B. Davis, kbdavis@jonesday.com,	Address: 2727 N. Harwood Street, Dallas Texas 75201 Telephone: (214) 969-4528	

(See Rule 45, Federal Rules of Civil Procedure, Parts (c), (d), and (e), on next page)

PROOF OF SERVICE

DATE PLACE

SERVED

SERVED ON (PRINT NAME) MANNER OF SERVICE

SERVED BY (PRINT NAME) TITLE

DECLARATION OF SERVER

I declare under penalty of perjury under the laws of the United States of America that the foregoing information contained in the Proof of Service is true and correct.

Executed on ______________________

SIGNATURE OF SERVER

ADDRESS OF SERVER

Rule 45, Federal Rules of Civil Procedure, Parts (c), (d), and (e), as amended on December 1, 2006:

(c) PROTECTION OF PERSONS SUBJECT TO SUBPOENAS.

(1) A party or an attorney responsible for the issuance and service of a subpoena shall take reasonable steps to avoid imposing undue burden or expense on a person subject to that subpoena. The court on behalf of which the subpoena was issued shall enforce this duty and impose upon the party or attorney in breach of this duty an appropriate sanction, which may include, but is not limited to, lost earnings and reasonable attorney's fee.

(2) (A) A person commanded to produce and permit inspection, copying, testing, or sampling of designated electronically stored information, books, papers, documents or tangible things, or inspection of premises need not appear in person at the place of production or inspection unless commanded to appear for deposition, hearing or trial.

(B) Subject to paragraph (d)(2) of this rule, a person commanded to produce and permit inspection, copying, testing, or sampling may, within 14 days after service of the subpoena or before the time specified for compliance if such time is less than 14 days after service, serve upon the party or attorney designated in the subpoena written objection to producing any or all of the designated materials or inspection of the premises—or to producing electronically stored information in the form or forms requested. If objection is made, the party serving the subpoena shall not be entitled to inspect, copy, test, or sample the materials or inspect the premises except pursuant to an order of the court by which the subpoena was issued. If objection has been made, the party serving the subpoena may, upon notice to the person commanded to produce, move at any time for an order to compel the production, inspection, copying, testing, or sampling. Such an order to compel shall protect any person who is not a party or an officer of a party from significant expense resulting from the inspection, copying, testing, or sampling commanded.

(3) (A) On timely motion, the court by which a subpoena was issued shall quash or modify the subpoena if it

(i) fails to allow reasonable time for compliance,

(ii) requires a person who is not a party or an officer of a party to travel to a place more than 100 miles from the place where that person resides, is employed or regularly transacts business in person, except that, subject to the provisions of clause (c) (3) (B) (iii) of this rule, such a person may in order to attend trial be commanded to travel from any such place within the state in which the trial is held;

(iii) requires disclosure of privileged or other protected matter and no exception or waiver applies; or

(iv) subjects a person to undue burden.

(B) If a subpoena

(i) requires disclosure of a trade secret or other confidential research, development, or commercial information, or

(ii) requires disclosure of an unretained expert's opinion or information not describing specific events or occurrences in dispute and resulting from the expert's study made not at the request of any party, or

(iii) requires a person who is not a party or an officer of a party to incur substantial expense to travel more than 100 miles to attend trial, the court may, to protect a person subject to or affected by the subpoena, quash or modify the subpoena or, if the party in whose behalf the subpoena is issued shows a substantial need for the testimony or material that cannot be otherwise met without undue hardship and assures that the person to whom the subpoena is addressed will be reasonably compensated, the court may order appearance or production only upon specified condition.

(d) DUTIES IN RESPONDING TO SUBPOENA

(1) (A) A person responding to a subpoena to produce documents shall produce them as they are kept in the usual course of business or shall organize and label them to correspond with the categories in the demand

(B) If a subpoena does not specify the form or forms for producing electronically stored information, a person responding to a subpoena must produce the information in a form or forms in which the person ordinarily maintains it or in a form or forms that are reasonably usable.

(C) A person responding to a subpoena need not produce the same electronically stored information in more than one form.

(D) A person responding to a subpoena need not provide discovery of electronically stored information from sources that the person identifies as not reasonably accessible because of undue burden or cost. On motion to compel discovery or to quash, the person from whom discovery is sought must show that the information sought is not reasonably accessible because of undue burden or cost. If that showing is made, the court may nonetheless order discovery from such sources if the requesting party shows good cause, considering the limitation of Rule 16(b)(2)(C). The court may specify conditions for the discovery.

(2) (A) When information subject to a subpoena is withheld on a claim that it is privileged or subject to protection as trial-preparation materials, the claim shall be made expressly and shall be supported by a description of the nature of the documents, communications, or things not produced that is sufficient to enable the demanding party to contest the claim.

(B) If information is produced in response to a subpoena that is subject to a claim of privilege or of protection as trial-preparation material, the person making the claim may notify any party that received the information of the claim and the basis for it. After being notified, any party must promptly return, sequester, or destroy the specified information and any copies it has and may not use or disclose the information until the claim is resolved. A receiving party may promptly present the information to the court under seal for a determination of the claim. If the receiving party disclosed the information before being notified, it must take reasonable steps to retrieve it. The person who produced the information must preserve the information until the claim is resolved.

(e) CONTEMPT. Failure of any person without adequate excuse to obey a subpoena served upon that person may be deemed a contempt of the court from which the subpoena issued. An adequate cause for failure to obey exists when a subpoena purports to require a nonparty to attend or produce at a place not within the limits provided by clause (ii) of subparagraph (c)(3)(A).

EXHIBIT A

The Internet Engineering Task Force is required to produce and permit inspection and copying of documents and things in its possession, custody or control that relate to the following categories of requests according to the following definitions and instructions.

DEFINITIONS & INSTRUCTIONS

1. This Subpoena seeks disclosure to the full extent of the Federal Rules of Civil Procedure and shall be interpreted as inclusive rather than exclusive.

2. These Requests are of a continuing nature and, as required by Federal Rule of Civil Procedure 26(e), You are required to timely provide supplemental documents and things if You obtain additional or different documents and things responsive to these Requests.

3. The following words and terms shall have the following meanings:

A. "Internet Standards Process RFC" shall mean and refer to the document RFC 2026, "The Internet Standards Process – Revision 3," October 1996, authored by Scott O. Bradner, attached as **Exhibit D**.

B. "Internet Standards-Related Publications" shall mean and refer to the documents described in Section 2 of the Internet Standards Process RFC, including Requests for Comments (RFCs), Internet-Drafts, publications of the IESG, IAB or the Internet community, and any other documents published by, or made available through, the IETF.

C. "You" and "Your," "IETF" and "Internet Engineering Task Force" means the Internet Engineering Task Force and the Internet Society, and all of its predecessors or successors (merged, acquired, or otherwise), parents, divisions, subsidiaries, and affiliates thereof, and all officers, agents, employees, in-house and outside counsel, and other persons acting on its behalf.

EXHIBIT B

The Internet Engineering Task Force is required to designate a representative to prepare for and testify on the following 30(b)(6) deposition topics according to the following definitions and instructions.

DEFINITIONS & INSTRUCTIONS

1. This 30(b)(6) notice seeks disclosure to the full extent of the Federal Rules of Civil Procedure and shall be interpreted as inclusive rather than exclusive.

2. The Definitions & Instructions in **Exhibit A** to this subpoena that requests the production of documents and things are incorporated herein by reference for the purposes of defining the 30(b)(6) deposition topics below.

DEPOSITION TOPICS

1. The authenticity of the RFC 2806 Publications.

2. The authenticity of any documents identified or produced in response to this subpoena.

3. Between 1997 and 2002, your awareness and knowledge of the RFC 2806 Publications.

4. Between 1997 and 2002, the availability of the RFC 2806 Publications, including the date on which the RFC 2806 Publications were first made published or available to members of the public.

5. Between 1997 and 2002, your practices regarding the availability of Internet Standards-Related Publications, including, for example, your practices regarding posting, distributing, displaying or disseminating Internet Standards-Related Publications.

6. Between 1997 and 2002, the process of creating an Internet Standard, including as described in the Internet Standards Process RFC.

REQUESTS

1. All Internet Standards-Related Publications relating to RFC 2806, including the documents identified below, any draft versions of such documents, and any other documents describing the use of URLs for telephony (collectively, "RFC 2806 Publications").

- "URLs for Telephone Calls," April 2000, authored by Antti Vaha-Sipila and/or Nokia Mobile Phones, attached as **Exhibit E**;
- "URLs for Telephony," February 23, 1998, authored by Antti Vaha-Sipila and/or Nokia Mobile Phones, attached as **Exhibit F**;
- "URLs for Telephony," August 26, 1997, authored by Antti Vaha-Sipila and/or Nokia Mobile Phones, attached as **Exhibit G**;
- "Conversational Multimedia URLs," December 16, 1997, authored by Pete Cordell, attached as **Exhibit H**; and
- "Uniform Resource Locators (URL)," December 1994, authored by Tim Berners-Lee, Larry Masinter, Mark McCahill, attached as **Exhibit I**.

2. Between 1997 and 2002, documents reflecting discussion(s) or comment(s) about the RFC 2806 Publications.

3. Between 1997 and 2002, documents concerning the availability of the RFC 2806 Publications, including showing the dates on which the RFC 2806 Publications were first made available, the persons who accessed the RFC 2806 Publications and the persons to whom the RFC 2806 Publications were sent.

4. Between 1997 and 2002, documents describing your practices regarding the availability of Internet Standards-Related Publications, including, for example, your practices regarding posting, distributing, displaying or disseminating Internet Standards-Related Publications.

5. Documents related to technical development of the RFC 2806 Publications.

```
                         *****************
                         *** TX REPORT ***
                         *****************


JOB NO.                  1545
ST. TIME                 10/24 17:02
PGS.                     5
SEND DOCUMENT NAME

TX IMCOMPLETE            -----
TRANSACTION OK           -----
ERROR                    kapfer@isoc.org                    Greg Kapfer
```

AO88 (Rev. 12/06) Subpoena in a Civil Case

Issued by the
UNITED STATES DISTRICT COURT
Northern District of California

SMARTPHONE TECHNOLOGIES LLC, **Plaintiff,** **v.** **HUAWEI DEVICE USA INC. et al.,** **Defendants.**	§ § § § § § § § § §	**Civil Action No. 6:12-cv-245** **United States District Court for the Eastern District of Texas** **SUBPOENA IN A CIVIL CASE**

TO: The Internet Society
1775 Wiehle Ave
Suite 201
Reston, VA 20190

□	YOU ARE COMMANDED to appear in the United States District Court at the place, date, and time specified below to testify in the above case.	
PLACE OF TESTIMONY		COURTROOM DATE AND TIME
☒	YOU ARE COMMANDED to appear at the place, date, and time specified below to testify at the taking of a deposition in the above case. Any organization that is subpoenaed for the taking of a deposition shall designate one or more officers, directors, or managing agents, or other persons who consent to testify on its behalf and may set forth, for each person designated, the matters on which the person will testify. Federal Rule of Civil Procedure 30(b)(6).	
PLACE OF DEPOSITION Jones Day, 51 Louisiana Avenue, N.W., Washington, D.C. 200001-2113		DATE AND TIME November 13, 2013, 9.00 a.m.
☒	YOU ARE COMMANDED to produce and permit inspection and copying of the following documents or objects at the place, date, and time specified below (List documents or objects): Documents identified in Exhibit "A" attached to this subpoena.	
PLACE OF PRODUCTION Jones Day, 51 Louisiana Avenue, N.W., Washington, D.C. 200001-2113		DATE AND TIME November 6, 2013
□	YOU ARE COMMANDED to permit inspection of the following premises at the date and time specified below.	
PREMISES		DATE AND TIME
ISSUING OFFICER SIGNATURE AND TITLE		

Exhibit D

Network Working Group
Request for Comments: 2026
BCP: 9
Obsoletes: 1602
Category: Best Current Practice

S. Bradner
Harvard University
October 1996

The Internet Standards Process -- Revision 3

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.

Abstract

This memo documents the process used by the Internet community for the standardization of protocols and procedures. It defines the stages in the standardization process, the requirements for moving a document between stages and the types of documents used during this process. It also addresses the intellectual property rights and copyright issues associated with the standards process.

Table of Contents

1. INTRODUCTION

This memo documents the process currently used by the Internet community for the standardization of protocols and procedures. The Internet Standards process is an activity of the Internet Society that is organized and managed on behalf of the Internet community by the Internet Architecture Board (IAB) and the Internet Engineering Steering Group (IESG).

1.1 Internet Standards

The Internet, a loosely-organized international collaboration of autonomous, interconnected networks, supports host-to-host communication through voluntary adherence to open protocols and procedures defined by Internet Standards. There are also many isolated interconnected networks, which are not connected to the global Internet but use the Internet Standards.

The Internet Standards Process described in this document is concerned with all protocols, procedures, and conventions that are used in or by the Internet, whether or not they are part of the TCP/IP protocol suite. In the case of protocols developed and/or standardized by non-Internet organizations, however, the Internet Standards Process normally applies to the application of the protocol or procedure in the Internet context, not to the specification of the protocol itself.

In general, an Internet Standard is a specification that is stable and well-understood, is technically competent, has multiple, independent, and interoperable implementations with substantial operational experience, enjoys significant public support, and is recognizably useful in some or all parts of the Internet.

1.2 The Internet Standards Process

In outline, the process of creating an Internet Standard is straightforward: a specification undergoes a period of development and several iterations of review by the Internet community and revision based upon experience, is adopted as a Standard by the appropriate body (see below), and is published. In practice, the process is more complicated, due to (1) the difficulty of creating specifications of high technical quality; (2) the need to consider the interests of all of the affected parties; (3) the importance of establishing widespread community consensus; and (4) the difficulty of evaluating the utility of a particular specification for the Internet community.

The goals of the Internet Standards Process are:
- technical excellence;
- prior implementation and testing;
- clear, concise, and easily understood documentation;
- openness and fairness; and
- timeliness.

The procedures described in this document are designed to be fair, open, and objective; to reflect existing (proven) practice; and to be flexible.

- These procedures are intended to provide a fair, open, and objective basis for developing, evaluating, and adopting Internet Standards. They provide ample opportunity for participation and comment by all interested parties. At each stage of the standardization process, a specification is repeatedly discussed and its merits debated in open meetings and/or public electronic mailing lists, and it is made available for review via world-wide on-line directories.

- These procedures are explicitly aimed at recognizing and adopting generally-accepted practices. Thus, a candidate specification must be implemented and tested for correct operation and interoperability by multiple independent parties and utilized in increasingly demanding environments, before it can be adopted as an Internet Standard.

- These procedures provide a great deal of flexibility to adapt to the wide variety of circumstances that occur in the standardization process. Experience has shown this flexibility to be vital in achieving the goals listed above.

The goal of technical competence, the requirement for prior implementation and testing, and the need to allow all interested parties to comment all require significant time and effort. On the other hand, today's rapid development of networking technology demands timely development of standards. The Internet Standards Process is intended to balance these conflicting goals. The process is believed to be as short and simple as possible without sacrificing technical excellence, thorough testing before adoption of a standard, or openness and fairness.

From its inception, the Internet has been, and is expected to remain, an evolving system whose participants regularly factor new requirements and technology into its design and implementation. Users of the Internet and providers of the equipment, software, and services that support it should anticipate and embrace this evolution as a major tenet of Internet philosophy.

The procedures described in this document are the result of a number of years of evolution, driven both by the needs of the growing and increasingly diverse Internet community, and by experience.

1.3 Organization of This Document

Section 2 describes the publications and archives of the Internet Standards Process. Section 3 describes the types of Internet standard specifications. Section 4 describes the Internet standards specifications track. Section 5 describes Best Current Practice RFCs. Section 6 describes the process and rules for Internet standardization. Section 7 specifies the way in which externally-sponsored specifications and practices, developed and controlled by other standards bodies or by others, are handled within the Internet Standards Process. Section 8 describes the requirements for notices and record keeping Section 9 defines a variance process to allow one-time exceptions to some of the requirements in this document Section 10 presents the rules that are required to protect intellectual property rights in the context of the development and use of Internet Standards. Section 11 includes acknowledgments of some of the people involved in creation of this document. Section 12 notes that security issues are not dealt with by this document. Section 13 contains a list of numbered references. Section 14 contains definitions of some of the terms used in this document. Section 15 lists the author's email and postal addresses. Appendix A contains a list of frequently-used acronyms.

2. INTERNET STANDARDS-RELATED PUBLICATIONS

2.1 Requests for Comments (RFCs)

Each distinct version of an Internet standards-related specification is published as part of the "Request for Comments" (RFC) document series. This archival series is the official publication channel for Internet standards documents and other publications of the IESG, IAB, and Internet community. RFCs can be obtained from a number of Internet hosts using anonymous FTP, gopher, World Wide Web, and other Internet document-retrieval systems.

The RFC series of documents on networking began in 1969 as part of the original ARPA wide-area networking (ARPANET) project (see Appendix A for glossary of acronyms). RFCs cover a wide range of topics in addition to Internet Standards, from early discussion of new research concepts to status memos about the Internet. RFC publication is the direct responsibility of the RFC Editor, under the general direction of the IAB.

The rules for formatting and submitting an RFC are defined in [5]. Every RFC is available in ASCII text. Some RFCs are also available in other formats. The other versions of an RFC may contain material (such as diagrams and figures) that is not present in the ASCII version, and it may be formatted differently.

```
*********************************************************
*                                                       *
*  A stricter requirement applies to standards-track    *
*  specifications:  the ASCII text version is the       *
*  definitive reference, and therefore it must be a     *
*  complete and accurate specification of the standard, *
*  including all necessary diagrams and illustrations.  *
*                                                       *
*********************************************************
```

The status of Internet protocol and service specifications is summarized periodically in an RFC entitled "Internet Official Protocol Standards" [1]. This RFC shows the level of maturity and other helpful information for each Internet protocol or service specification (see section 3).

Some RFCs document Internet Standards. These RFCs form the 'STD' subseries of the RFC series [4]. When a specification has been adopted as an Internet Standard, it is given the additional label "STDxxx", but it keeps its RFC number and its place in the RFC series. (see section 4.1.3)

Some RFCs standardize the results of community deliberations about statements of principle or conclusions about what is the best way to perform some operations or IETF process function. These RFCs form the specification has been adopted as a BCP, it is given the additional label "BCPxxx", but it keeps its RFC number and its place in the RFC series. (see section 5)

Not all specifications of protocols or services for the Internet should or will become Internet Standards or BCPs. Such non-standards track specifications are not subject to the rules for Internet standardization. Non-standards track specifications may be published directly as "Experimental" or "Informational" RFCs at the discretion of the RFC Editor in consultation with the IESG (see section 4.2).

```
********************************************************
*                                                      *
*   It is important to remember that not all RFCs      *
*   are standards track documents, and that not all    *
*   standards track documents reach the level of       *
*   Internet Standard. In the same way, not all RFCs   *
*   which describe current practices have been given   *
*   the review and approval to become BCPs. See        *
*   RFC-1796 [6] for further information.              *
*                                                      *
********************************************************
```

2.2 Internet-Drafts

During the development of a specification, draft versions of the document are made available for informal review and comment by placing them in the IETF's "Internet-Drafts" directory, which is replicated on a number of Internet hosts. This makes an evolving working document readily available to a wide audience, facilitating the process of review and revision.

An Internet-Draft that is published as an RFC, or that has remained unchanged in the Internet-Drafts directory for more than six months without being recommended by the IESG for publication as an RFC, is simply removed from the Internet-Drafts directory. At any time, an Internet-Draft may be replaced by a more recent version of the same specification, restarting the six-month timeout period.

An Internet-Draft is NOT a means of "publishing" a specification; specifications are published through the RFC mechanism described in the previous section. Internet-Drafts have no formal status, and are subject to change or removal at any time.

```
********************************************************
*                                                      *
*   Under no circumstances should an Internet-Draft    *
*   be referenced by any paper, report, or Request-    *
*   for-Proposal, nor should a vendor claim compliance *
*   with an Internet-Draft.                            *
*                                                      *
********************************************************
```

Note: It is acceptable to reference a standards-track specification that may reasonably be expected to be published as an RFC using the phrase "Work in Progress" without referencing an Internet-Draft. This may also be done in a standards track document itself as long as the specification in which the reference is made would stand as a complete and understandable document with or without the reference to the "Work in Progress".

3. INTERNET STANDARD SPECIFICATIONS

Specifications subject to the Internet Standards Process fall into one of two categories: Technical Specification (TS) and Applicability Statement (AS).

3.1 Technical Specification (TS)

A Technical Specification is any description of a protocol, service, procedure, convention, or format. It may completely describe all of the relevant aspects of its subject, or it may leave one or more parameters or options unspecified. A TS may be completely self-contained, or it may incorporate material from other specifications by reference to other documents (which might or might not be Internet Standards).

A TS shall include a statement of its scope and the general intent for its use (domain of applicability). Thus, a TS that is inherently specific to a particular context shall contain a statement to that effect. However, a TS does not specify requirements for its use within the Internet; these requirements, which depend on the particular context in which the TS is incorporated by different system configurations, are defined by an Applicability Statement.

3.2 Applicability Statement (AS)

An Applicability Statement specifies how, and under what circumstances, one or more TSs may be applied to support a particular Internet capability. An AS may specify uses for TSs that are not Internet Standards, as discussed in Section 7.

An AS identifies the relevant TSs and the specific way in which they are to be combined, and may also specify particular values or ranges of TS parameters or subfunctions of a TS protocol that must be implemented. An AS also specifies the circumstances in which the use of a particular TS is required, recommended, or elective (see section 3.3).

An AS may describe particular methods of using a TS in a restricted "domain of applicability", such as Internet routers, terminal servers, Internet systems that interface to Ethernets, or datagram-based database servers.

The broadest type of AS is a comprehensive conformance specification, commonly called a "requirements document", for a particular class of Internet systems, such as Internet routers or Internet hosts.

An AS may not have a higher maturity level in the standards track than any standards-track TS on which the AS relies (see section 4.1). For example, a TS at Draft Standard level may be referenced by an AS at the Proposed Standard or Draft Standard level, but not by an AS at the Standard level.

3.3 Requirement Levels

An AS shall apply one of the following "requirement levels" to each of the TSs to which it refers:

(a) Required: Implementation of the referenced TS, as specified by the AS, is required to achieve minimal conformance. For example, IP and ICMP must be implemented by all Internet systems using the TCP/IP Protocol Suite.

(b) Recommended: Implementation of the referenced TS is not required for minimal conformance, but experience and/or generally accepted technical wisdom suggest its desirability in the domain of applicability of the AS. Vendors are strongly encouraged to include the functions, features, and protocols of Recommended TSs in their products, and should omit them only if the omission is justified by some special circumstance. For example, the TELNET protocol should be implemented by all systems that would benefit from remote access.

(c) Elective: Implementation of the referenced TS is optional within the domain of applicability of the AS; that is, the AS creates no explicit necessity to apply the TS. However, a particular vendor may decide to implement it, or a particular user may decide that it is a necessity in a specific environment. For example, the DECNET MIB could be seen as valuable in an environment where the DECNET protocol is used.

As noted in section 4.1, there are TSs that are not in the standards track or that have been retired from the standards track, and are therefore not required, recommended, or elective. Two additional "requirement level" designations are available for these TSs:

(d) Limited Use: The TS is considered to be appropriate for use only in limited or unique circumstances. For example, the usage of a protocol with the "Experimental" designation should generally be limited to those actively involved with the experiment.

(e) Not Recommended: A TS that is considered to be inappropriate for general use is labeled "Not Recommended". This may be because of its limited functionality, specialized nature, or historic status.

Although TSs and ASs are conceptually separate, in practice a standards-track document may combine an AS and one or more related TSs. For example, Technical Specifications that are developed specifically and exclusively for some particular domain of applicability, e.g., for mail server hosts, often contain within a single specification all of the relevant AS and TS information. In such cases, no useful purpose would be served by deliberately distributing the information among several documents just to preserve the formal AS/TS distinction. However, a TS that is likely to apply to more than one domain of applicability should be developed in a modular fashion, to facilitate its incorporation by multiple ASs.

The "Official Protocol Standards" RFC (STD1) lists a general requirement level for each TS, using the nomenclature defined in this section. This RFC is updated periodically. In many cases, more detailed descriptions of the requirement levels of particular protocols and of individual features of the protocols will be found in appropriate ASs.

4. THE INTERNET STANDARDS TRACK

Specifications that are intended to become Internet Standards evolve through a set of maturity levels known as the "standards track". These maturity levels -- "Proposed Standard", "Draft Standard", and "Standard" -- are defined and discussed in section 4.1. The way in which specifications move along the standards track is described in section 6.

Even after a specification has been adopted as an Internet Standard, further evolution often occurs based on experience and the recognition of new requirements. The nomenclature and procedures of Internet standardization provide for the replacement of old Internet

Standards with new ones, and the assignment of descriptive labels to indicate the status of "retired" Internet Standards. A set of maturity levels is defined in section 4.2 to cover these and other specifications that are not considered to be on the standards track.

4.1 Standards Track Maturity Levels

Internet specifications go through stages of development, testing, and acceptance. Within the Internet Standards Process, these stages are formally labeled "maturity levels".

This section describes the maturity levels and the expected characteristics of specifications at each level.

4.1.1 Proposed Standard

The entry-level maturity for the standards track is "Proposed Standard". A specific action by the IESG is required to move a specification onto the standards track at the "Proposed Standard" level.

A Proposed Standard specification is generally stable, has resolved known design choices, is believed to be well-understood, has received significant community review, and appears to enjoy enough community interest to be considered valuable. However, further experience might result in a change or even retraction of the specification before it advances.

Usually, neither implementation nor operational experience is required for the designation of a specification as a Proposed Standard. However, such experience is highly desirable, and will usually represent a strong argument in favor of a Proposed Standard designation.

The IESG may require implementation and/or operational experience prior to granting Proposed Standard status to a specification that materially affects the core Internet protocols or that specifies behavior that may have significant operational impact on the Internet.

A Proposed Standard should have no known technical omissions with respect to the requirements placed upon it. However, the IESG may waive this requirement in order to allow a specification to advance to the Proposed Standard state when it is considered to be useful and necessary (and timely) even with known technical omissions.

Implementors should treat Proposed Standards as immature specifications. It is desirable to implement them in order to gain experience and to validate, test, and clarify the specification. However, since the content of Proposed Standards may be changed if problems are found or better solutions are identified, deploying implementations of such standards into a disruption-sensitive environment is not recommended.

4.1.2 Draft Standard

A specification from which at least two independent and interoperable implementations from different code bases have been developed, and for which sufficient successful operational experience has been obtained, may be elevated to the "Draft Standard" level. For the purposes of this section, "interoperable" means to be functionally equivalent or interchangeable components of the system or process in which they are used. If patented or otherwise controlled technology is required for implementation, the separate implementations must also have resulted from separate exercise of the licensing process. Elevation to Draft Standard is a major advance in status, indicating a strong belief that the specification is mature and will be useful.

The requirement for at least two independent and interoperable implementations applies to all of the options and features of the specification. In cases in which one or more options or features have not been demonstrated in at least two interoperable implementations, the specification may advance to the Draft Standard level only if those options or features are removed.

The Working Group chair is responsible for documenting the specific implementations which qualify the specification for Draft or Internet Standard status along with documentation about testing of the interoperation of these implementations. The documentation must include information about the support of each of the individual options and features. This documentation should be submitted to the Area Director with the protocol action request. (see Section 6)

A Draft Standard must be well-understood and known to be quite stable, both in its semantics and as a basis for developing an implementation. A Draft Standard may still require additional or more widespread field experience, since it is possible for implementations based on Draft Standard specifications to demonstrate unforeseen behavior when subjected to large-scale use in production environments.

A Draft Standard is normally considered to be a final specification, and changes are likely to be made only to solve specific problems encountered. In most circumstances, it is reasonable for vendors to deploy implementations of Draft Standards into a disruption sensitive environment.

4.1.3 Internet Standard

A specification for which significant implementation and successful operational experience has been obtained may be elevated to the Internet Standard level. An Internet Standard (which may simply be referred to as a Standard) is characterized by a high degree of technical maturity and by a generally held belief that the specified protocol or service provides significant benefit to the Internet community.

A specification that reaches the status of Standard is assigned a number in the STD series while retaining its RFC number.

4.2 Non-Standards Track Maturity Levels

Not every specification is on the standards track. A specification may not be intended to be an Internet Standard, or it may be intended for eventual standardization but not yet ready to enter the standards track. A specification may have been superseded by a more recent Internet Standard, or have otherwise fallen into disuse or disfavor.

Specifications that are not on the standards track are labeled with one of three "off-track" maturity levels: "Experimental", "Informational", or "Historic". The documents bearing these labels are not Internet Standards in any sense.

4.2.1 Experimental

The "Experimental" designation typically denotes a specification that is part of some research or development effort. Such a specification is published for the general information of the Internet technical community and as an archival record of the work, subject only to editorial considerations and to verification that there has been adequate coordination with the standards process (see below). An Experimental specification may be the output of an organized Internet research effort (e.g., a Research Group of the IRTF), an IETF Working Group, or it may be an individual contribution.

4.2.2 Informational

An "Informational" specification is published for the general information of the Internet community, and does not represent an Internet community consensus or recommendation. The Informational designation is intended to provide for the timely publication of a very broad range of responsible informational documents from many sources, subject only to editorial considerations and to verification that there has been adequate coordination with the standards process (see section 4.2.3).

Specifications that have been prepared outside of the Internet community and are not incorporated into the Internet Standards Process by any of the provisions of section 10 may be published as Informational RFCs, with the permission of the owner and the concurrence of the RFC Editor.

4.2.3 Procedures for Experimental and Informational RFCs

Unless they are the result of IETF Working Group action, documents intended to be published with Experimental or Informational status should be submitted directly to the RFC Editor. The RFC Editor will publish any such documents as Internet-Drafts which have not already been so published. In order to differentiate these Internet-Drafts they will be labeled or grouped in the I-D directory so they are easily recognizable. The RFC Editor will wait two weeks after this publication for comments before proceeding further. The RFC Editor is expected to exercise his or her judgment concerning the editorial suitability of a document for publication with Experimental or Informational status, and may refuse to publish a document which, in the expert opinion of the RFC Editor, is unrelated to Internet activity or falls below the technical and/or editorial standard for RFCs.

To ensure that the non-standards track Experimental and Informational designations are not misused to circumvent the Internet Standards Process, the IESG and the RFC Editor have agreed that the RFC Editor will refer to the IESG any document submitted for Experimental or Informational publication which, in the opinion of the RFC Editor, may be related to work being done, or expected to be done, within the IETF community. The IESG shall review such a referred document within a reasonable period of time, and recommend either that it be published as originally submitted or referred to the IETF as a contribution to the Internet Standards Process.

If (a) the IESG recommends that the document be brought within the IETF and progressed within the IETF context, but the author declines to do so, or (b) the IESG considers that the document proposes something that conflicts with, or is actually inimical to, an established IETF effort, the document may still be published as an Experimental or Informational RFC. In these cases, however, the IESG may insert appropriate "disclaimer" text into the RFC either in or immediately following the "Status of this Memo" section in order to make the circumstances of its publication clear to readers.

Documents proposed for Experimental and Informational RFCs by IETF Working Groups go through IESG review. The review is initiated using the process described in section 6.1.1.

4.2.4 Historic

A specification that has been superseded by a more recent specification or is for any other reason considered to be obsolete is assigned to the "Historic" level. (Purists have suggested that the word should be "Historical"; however, at this point the use of "Historic" is historical.)

Note: Standards track specifications normally must not depend on other standards track specifications which are at a lower maturity level or on non standards track specifications other than referenced specifications from other standards bodies. (See Section 7.)

5. BEST CURRENT PRACTICE (BCP) RFCs

The BCP subseries of the RFC series is designed to be a way to standardize practices and the results of community deliberations. A BCP document is subject to the same basic set of procedures as standards track documents and thus is a vehicle by which the IETF community can define and ratify the community's best current thinking on a statement of principle or on what is believed to be the best way to perform some operations or IETF process function.

Historically Internet standards have generally been concerned with the technical specifications for hardware and software required for computer communication across interconnected networks. However, since the Internet itself is composed of networks operated by a great variety of organizations, with diverse goals and rules, good user service requires that the operators and administrators of the Internet follow some common guidelines for policies and operations. While these guidelines are generally different in scope and style from protocol standards, their establishment needs a similar process for consensus building.

While it is recognized that entities such as the IAB and IESG are composed of individuals who may participate, as individuals, in the technical work of the IETF, it is also recognized that the entities themselves have an existence as leaders in the community. As leaders in the Internet technical community, these entities should have an outlet to propose ideas to stimulate work in a particular area, to raise the community's sensitivity to a certain issue, to make a statement of architectural principle, or to communicate their thoughts on other matters. The BCP subseries creates a smoothly structured way for these management entities to insert proposals into the consensus-building machinery of the IETF while gauging the community's view of that issue.

Finally, the BCP series may be used to document the operation of the IETF itself. For example, this document defines the IETF Standards Process and is published as a BCP.

5.1 BCP Review Process

Unlike standards-track documents, the mechanisms described in BCPs are not well suited to the phased roll-in nature of the three stage standards track and instead generally only make sense for full and immediate instantiation.

The BCP process is similar to that for proposed standards. The BCP is submitted to the IESG for review, (see section 6.1.1) and the existing review process applies, including a Last-Call on the IETF Announce mailing list. However, once the IESG has approved the document, the process ends and the document is published. The resulting document is viewed as having the technical approval of the IETF.

Specifically, a document to be considered for the status of BCP must undergo the procedures outlined in sections 6.1, and 6.4 of this document. The BCP process may be appealed according to the procedures in section 6.5.

Because BCPs are meant to express community consensus but are arrived at more quickly than standards, BCPs require particular care. Specifically, BCPs should not be viewed simply as stronger Informational RFCs, but rather should be viewed as documents suitable for a content different from Informational RFCs.

A specification, or group of specifications, that has, or have been approved as a BCP is assigned a number in the BCP series while retaining its RFC number(s).

6. THE INTERNET STANDARDS PROCESS

The mechanics of the Internet Standards Process involve decisions of the IESG concerning the elevation of a specification onto the standards track or the movement of a standards-track specification from one maturity level to another. Although a number of reasonably objective criteria (described below and in section 4) are available to guide the IESG in making a decision to move a specification onto, along, or off the standards track, there is no algorithmic guarantee of elevation to or progression along the standards track for any specification. The experienced collective judgment of the IESG concerning the technical quality of a specification proposed for elevation to or advancement in the standards track is an essential component of the decision-making process.

6.1 Standards Actions

A "standards action" -- entering a particular specification into, advancing it within, or removing it from, the standards track -- must be approved by the IESG.

6.1.1 Initiation of Action

A specification that is intended to enter or advance in the Internet standards track shall first be posted as an Internet-Draft (see section 2.2) unless it has not changed since publication as an RFC. It shall remain as an Internet-Draft for a period of time, not less than two weeks, that permits useful community review, after which a recommendation for action may be initiated.

A standards action is initiated by a recommendation by the IETF Working group responsible for a specification to its Area Director, copied to the IETF Secretariat or, in the case of a specification not associated with a Working Group, a recommendation by an individual to the IESG.

6.1.2 IESG Review and Approval

The IESG shall determine whether or not a specification submitted to it according to section 6.1.1 satisfies the applicable criteria for the recommended action (see sections 4.1 and 4.2), and shall in addition determine whether or not the technical quality and clarity of the specification is consistent with that expected for the maturity level to which the specification is recommended.

In order to obtain all of the information necessary to make these determinations, particularly when the specification is considered by the IESG to be extremely important in terms of its potential impact on the Internet or on the suite of Internet protocols, the IESG may, at its discretion, commission an independent technical review of the specification.

The IESG will send notice to the IETF of the pending IESG consideration of the document(s) to permit a final review by the general Internet community. This "Last-Call" notification shall be via electronic mail to the IETF Announce mailing list. Comments on a Last-Call shall be accepted from anyone, and should be sent as directed in the Last-Call announcement.

The Last-Call period shall be no shorter than two weeks except in those cases where the proposed standards action was not initiated by an IETF Working Group, in which case the Last-Call period shall be no shorter than four weeks. If the IESG believes that the community interest would be served by allowing more time for comment, it may decide on a longer Last-Call period or to explicitly lengthen a current Last-Call period.

The IESG is not bound by the action recommended when the specification was submitted. For example, the IESG may decide to consider the specification for publication in a different category than that requested. If the IESG determines this before the Last-Call is issued then the Last-Call should reflect the IESG's view. The IESG could also decide to change the publication category based on the response to a Last-Call. If this decision would result in a specification being published at a "higher" level than the original Last-Call was for, a new Last-Call should be issued indicating the IESG recommendation. In addition, the IESG may decide to recommend the formation of a new Working Group in the case of significant controversy in response to a Last-Call for specification not originating from an IETF Working Group.

In a timely fashion after the expiration of the Last-Call period, the IESG shall make its final determination of whether or not to approve the standards action, and shall notify the IETF of its decision via electronic mail to the IETF Announce mailing list.

6.1.3 Publication

If a standards action is approved, notification is sent to the RFC Editor and copied to the IETF with instructions to publish the specification as an RFC. The specification shall at that point be removed from the Internet-Drafts directory.

An official summary of standards actions completed and pending shall appear in each issue of the Internet Society's newsletter. This shall constitute the "publication of record" for Internet standards actions.

The RFC Editor shall publish periodically an "Internet Official Protocol Standards" RFC [1], summarizing the status of all Internet protocol and service specifications.

6.2 Advancing in the Standards Track

The procedure described in section 6.1 is followed for each action that attends the advancement of a specification along the standards track.

A specification shall remain at the Proposed Standard level for at least six (6) months.

A specification shall remain at the Draft Standard level for at least four (4) months, or until at least one IETF meeting has occurred, whichever comes later.

These minimum periods are intended to ensure adequate opportunity for community review without severely impacting timeliness. These intervals shall be measured from the date of publication of the corresponding RFC(s), or, if the action does not result in RFC publication, the date of the announcement of the IESG approval of the action.

A specification may be (indeed, is likely to be) revised as it advances through the standards track. At each stage, the IESG shall determine the scope and significance of the revision to the specification, and, if necessary and appropriate, modify the recommended action. Minor revisions are expected, but a significant revision may require that the specification accumulate more experience at its current maturity level before progressing. Finally, if the specification has been changed very significantly, the IESG may recommend that the revision be treated as a new document, re-entering the standards track at the beginning.

Change of status shall result in republication of the specification as an RFC, except in the rare case that there have been no changes at all in the specification since the last publication. Generally, desired changes will be "batched" for incorporation at the next level in the standards track. However, deferral of changes to the next standards action on the specification will not always be possible or desirable; for example, an important typographical error, or a technical error that does not represent a change in overall function of the specification, may need to be corrected immediately. In such cases, the IESG or RFC Editor may be asked to republish the RFC (with a new number) with corrections, and this will not reset the minimum time-at-level clock.

When a standards-track specification has not reached the Internet Standard level but has remained at the same maturity level for twenty-four (24) months, and every twelve (12) months thereafter until the status is changed, the IESG shall review the viability of the standardization effort responsible for that specification and the usefulness of the technology. Following each such review, the IESG shall approve termination or continuation of the development effort, at the same time the IESG shall decide to maintain the specification at the same maturity level or to move it to Historic status. This decision shall be communicated to the IETF by electronic mail to the IETF Announce mailing list to allow the Internet community an opportunity to comment. This provision is not intended to threaten a legitimate and active Working Group effort, but rather to provide an administrative mechanism for terminating a moribund effort.

6.3 Revising a Standard

A new version of an established Internet Standard must progress through the full Internet standardization process as if it were a completely new specification. Once the new version has reached the Standard level, it will usually replace the previous version, which will be moved to Historic status. However, in some cases both versions may remain as Internet Standards to honor the requirements of an installed base. In this situation, the relationship between the previous and the new versions must be explicitly stated in the text of the new version or in another appropriate document (e.g., an Applicability Statement; see section 3.2).

6.4 Retiring a Standard

As the technology changes and matures, it is possible for a new Standard specification to be so clearly superior technically that one or more existing standards track specifications for the same function should be retired. In this case, or when it is felt for some other reason that an existing standards track specification should be retired, the IESG shall approve a change of status of the old specification(s) to Historic. This recommendation shall be issued with the same Last-Call and notification procedures used for any other standards action. A request to retire an existing standard can originate from a Working Group, an Area Director or some other interested party.

6.5 Conflict Resolution and Appeals

Disputes are possible at various stages during the IETF process. As much as possible the process is designed so that compromises can be made, and genuine consensus achieved, however there are times when even the most reasonable and knowledgeable people are unable to agree. To achieve the goals of openness and fairness, such conflicts must be resolved by a process of open review and discussion. This section specifies the procedures that shall be followed to deal with Internet standards issues that cannot be resolved through the normal processes whereby IETF Working Groups and other Internet Standards Process participants ordinarily reach consensus.

6.5.1 Working Group Disputes

An individual (whether a participant in the relevant Working Group or not) may disagree with a Working Group recommendation based on his or her belief that either (a) his or her own views have not been adequately considered by the Working Group, or (b) the Working Group has made an incorrect technical choice which places the quality and/or integrity of the Working Group's product(s) in significant jeopardy. The first issue is a difficulty with Working Group process; the latter is an assertion of technical error. These two types of disagreement are quite different, but both are handled by the same process of review.

A person who disagrees with a Working Group recommendation shall always first discuss the matter with the Working Group's chair(s), who may involve other members of the Working Group (or the Working Group as a whole) in the discussion.

If the disagreement cannot be resolved in this way, any of the parties involved may bring it to the attention of the Area Director(s) for the area in which the Working Group is chartered. The Area Director(s) shall attempt to resolve the dispute.

If the disagreement cannot be resolved by the Area Director(s) any of the parties involved may then appeal to the IESG as a whole. The IESG shall then review the situation and attempt to resolve it in a manner of its own choosing.

If the disagreement is not resolved to the satisfaction of the parties at the IESG level, any of the parties involved may appeal the decision to the IAB. The IAB shall then review the situation and attempt to resolve it in a manner of its own choosing.

The IAB decision is final with respect to the question of whether or not the Internet standards procedures have been followed and with respect to all questions of technical merit.

6.5.2 Process Failures

This document sets forward procedures required to be followed to ensure openness and fairness of the Internet Standards Process, and the technical viability of the standards created. The IESG is the principal agent of the IETF for this purpose, and it is the IESG that is charged with ensuring that the required procedures have been followed, and that any necessary prerequisites to a standards action have been met.

If an individual should disagree with an action taken by the IESG in this process, that person should first discuss the issue with the ISEG Chair. If the IESG Chair is unable to satisfy the complainant then the IESG as a whole should re-examine the action taken, along with input from the complainant, and determine whether any further action is needed. The IESG shall issue a report on its review of the complaint to the IETF.

Should the complainant not be satisfied with the outcome of the IESG review, an appeal may be lodged to the IAB. The IAB shall then review the situation and attempt to resolve it in a manner of its own choosing and report to the IETF on the outcome of its review.

If circumstances warrant, the IAB may direct that an IESG decision be annulled, and the situation shall then be as it was before the IESG decision was taken. The IAB may also recommend an action to the IESG, or make such other recommendations as it deems fit. The IAB may not, however, pre-empt the role of the IESG by issuing a decision which only the IESG is empowered to make.

The IAB decision is final with respect to the question of whether or not the Internet standards procedures have been followed.

6.5.3 Questions of Applicable Procedure

Further recourse is available only in cases in which the procedures themselves (i.e., the procedures described in this document) are claimed to be inadequate or insufficient to the protection of the rights of all parties in a fair and open Internet Standards Process. Claims on this basis may be made to the Internet Society Board of Trustees. The President of the Internet Society shall acknowledge such an appeal within two weeks, and shall at the time of acknowledgment advise the petitioner of the expected duration of the Trustees' review of the appeal. The Trustees shall review the situation in a manner of its own choosing and report to the IETF on the outcome of its review.

The Trustees' decision upon completion of their review shall be final with respect to all aspects of the dispute.

6.5.4 Appeals Procedure

All appeals must include a detailed and specific description of the facts of the dispute.

All appeals must be initiated within two months of the public knowledge of the action or decision to be challenged.

At all stages of the appeals process, the individuals or bodies responsible for making the decisions have the discretion to define the specific procedures they will follow in the process of making their decision.

In all cases a decision concerning the disposition of the dispute, and the communication of that decision to the parties involved, must be accomplished within a reasonable period of time.

[NOTE: These procedures intentionally and explicitly do not establish a fixed maximum time period that shall be considered "reasonable" in all cases. The Internet Standards Process places a premium on consensus and efforts to achieve it, and deliberately foregoes deterministically swift execution of procedures in favor of a latitude within which more genuine technical agreements may be reached.]

7. EXTERNAL STANDARDS AND SPECIFICATIONS

Many standards groups other than the IETF create and publish standards documents for network protocols and services. When these external specifications play an important role in the Internet, it is desirable to reach common agreements on their usage -- i.e., to establish Internet Standards relating to these external specifications.

There are two categories of external specifications:

(1) Open Standards

Various national and international standards bodies, such as ANSI, ISO, IEEE, and ITU-T, develop a variety of protocol and service specifications that are similar to Technical Specifications defined here. National and international groups also publish "implementors' agreements" that are analogous to Applicability Statements, capturing a body of implementation-specific detail concerned with the practical application of their standards. All of these are considered to be "open external standards" for the purposes of the Internet Standards Process.

(2) Other Specifications

Other proprietary specifications that have come to be widely used in the Internet may be treated by the Internet community as if they were a "standards". Such a specification is not generally developed in an open fashion, is typically proprietary, and is controlled by the vendor, vendors, or organization that produced it.

7.1 Use of External Specifications

To avoid conflict between competing versions of a specification, the Internet community will not standardize a specification that is simply an "Internet version" of an existing external specification unless an explicit cooperative arrangement to do so has been made. However, there are several ways in which an external specification that is important for the operation and/or evolution of the Internet may be adopted for Internet use.

7.1.1 Incorporation of an Open Standard

An Internet Standard TS or AS may incorporate an open external standard by reference. For example, many Internet Standards incorporate by reference the ANSI standard character set "ASCII" [2]. Whenever possible, the referenced specification shall be available online.

7.1.2 Incorporation of Other Specifications

Other proprietary specifications may be incorporated by reference to a version of the specification as long as the proprietor meets the requirements of section 10. If the other proprietary specification is not widely and readily available, the IESG may request that it be published as an Informational RFC.

The IESG generally should not favor a particular proprietary specification over technically equivalent and competing specification(s) by making any incorporated vendor specification "required" or "recommended".

7.1.3 Assumption

An IETF Working Group may start from an external specification and develop it into an Internet specification. This is acceptable if (1) the specification is provided to the Working Group in compliance with the requirements of section 10, and (2) change control has been conveyed to IETF by the original developer of the specification for the specification or for specifications derived from the original specification.

8. NOTICES AND RECORD KEEPING

Each of the organizations involved in the development and approval of Internet Standards shall publicly announce, and shall maintain a publicly accessible record of, every activity in which it engages, to the extent that the activity represents the prosecution of any part of the Internet Standards Process. For purposes of this section, the organizations involved in the development and approval of Internet Standards includes the IETF, the IESG, the IAB, all IETF Working Groups, and the Internet Society Board of Trustees.

For IETF and Working Group meetings announcements shall be made by electronic mail to the IETF Announce mailing list and shall be made sufficiently far in advance of the activity to permit all interested parties to effectively participate. The announcement shall contain (or provide pointers to) all of the information that is necessary to support the participation of any interested individual. In the case of a meeting, for example, the announcement shall include an agenda that specifies the standards- related issues that will be discussed.

The formal record of an organization's standards-related activity shall include at least the following:

- o the charter of the organization (or a defining document equivalent to a charter);
- o complete and accurate minutes of meetings;
- o the archives of Working Group electronic mail mailing lists; and
- o all written contributions from participants that pertain to the organization's standards-related activity.

As a practical matter, the formal record of all Internet Standards Process activities is maintained by the IETF Secretariat, and is the responsibility of the IETF Secretariat except that each IETF Working Group is expected to maintain their own email list archive and must make a best effort to ensure that all traffic is captured and included in the archives. Also, the Working Group chair is responsible for providing the IETF Secretariat with complete and accurate minutes of all Working Group meetings. Internet-Drafts that have been removed (for any reason) from the Internet-Drafts directories shall be archived by the IETF Secretariat for the sole purpose of preserving an historical record of Internet standards activity and thus are not retrievable except in special circumstances.

9. VARYING THE PROCESS

This document, which sets out the rules and procedures by which Internet Standards and related documents are made is itself a product of the Internet Standards Process (as a BCP, as described in section 5). It replaces a previous version, and in time, is likely itself to be replaced.

While, when published, this document represents the community's view of the proper and correct process to follow, and requirements to be met, to allow for the best possible Internet Standards and BCPs, it cannot be assumed that this will always remain the case. From time to time there may be a desire to update it, by replacing it with a new version. Updating this document uses the same open procedures as are used for any other BCP.

In addition, there may be situations where following the procedures leads to a deadlock about a specific specification, or there may be situations where the procedures provide no guidance. In these cases it may be appropriate to invoke the variance procedure described below.

9.1 The Variance Procedure

Upon the recommendation of the responsible IETF Working Group (or, if no Working Group is constituted, upon the recommendation of an ad hoc committee), the IESG may enter a particular specification into, or advance it within, the standards track even though some of the requirements of this document have not or will not be met. The IESG may approve such a variance, however, only if it first determines that the likely benefits to the Internet community are likely to outweigh any costs to the Internet community that result from noncompliance with the requirements in this document. In exercising this discretion, the IESG shall at least consider (a) the technical merit of the specification, (b) the possibility of achieving the goals of the Internet Standards Process without granting a variance, (c) alternatives to the granting of a variance, (d) the collateral and precedential effects of granting a variance, and (e) the IESG's ability to craft a variance that is as narrow as possible. In determining whether to approve a variance, the IESG has discretion to limit the scope of the variance to particular parts of this document and to impose such additional restrictions or limitations as it determines appropriate to protect the interests of the Internet community.

The proposed variance must detail the problem perceived, explain the precise provision of this document which is causing the need for a variance, and the results of the IESG's considerations including consideration of points (a) through (d) in the previous paragraph. The proposed variance shall be issued as an Internet Draft. The IESG shall then issue an extended Last-Call, of no less than 4 weeks, to allow for community comment upon the proposal.

In a timely fashion after the expiration of the Last-Call period, the IESG shall make its final determination of whether or not to approve the proposed variance, and shall notify the IETF of its decision via electronic mail to the IETF Announce mailing list. If the variance is approved it shall be forwarded to the RFC Editor with a request that it be published as a BCP.

This variance procedure is for use when a one-time waving of some provision of this document is felt to be required. Permanent changes to this document shall be accomplished through the normal BCP process.

The appeals process in section 6.5 applies to this process.

9.2 Exclusions

No use of this procedure may lower any specified delays, nor exempt any proposal from the requirements of openness, fairness, or consensus, nor from the need to keep proper records of the meetings and mailing list discussions.

Specifically, the following sections of this document must not be subject of a variance: 5.1, 6.1, 6.1.1 (first paragraph), 6.1.2, 6.3 (first sentence), 6.5 and 9.

10. INTELLECTUAL PROPERTY RIGHTS

10.1. General Policy

In all matters of intellectual property rights and procedures, the intention is to benefit the Internet community and the public at large, while respecting the legitimate rights of others.

10.2 Confidentiality Obligations

No contribution that is subject to any requirement of confidentiality or any restriction on its dissemination may be considered in any part of the Internet Standards Process, and there must be no assumption of any confidentiality obligation with respect to any such contribution.

10.3. Rights and Permissions

In the course of standards work, the IETF receives contributions in various forms and from many persons. To best facilitate the dissemination of these contributions, it is necessary to understand any intellectual property rights (IPR) relating to the contributions.

10.3.1. All Contributions

By submission of a contribution, each person actually submitting the contribution is deemed to agree to the following terms and conditions on his own behalf, on behalf of the organization (if any) he represents and on behalf of the owners of any propriety rights in the contribution.. Where a submission identifies contributors in addition to the contributor(s) who provide the actual submission, the actual submitter(s) represent that each other named contributor was made aware of and agreed to accept the same terms and conditions on his own behalf, on behalf of any organization he may represent and any known owner of any proprietary rights in the contribution.

1. Some works (e.g. works of the U.S. Government) are not subject to copyright. However, to the extent that the submission is or may be subject to copyright, the contributor, the organization he represents (if any) and the owners of any proprietary rights in the contribution, grant an unlimited perpetual, non-exclusive, royalty-free, world-wide right and license to the ISOC and the IETF under any copyrights in the contribution. This license includes the right to copy, publish and distribute the contribution in any way, and to prepare derivative works that are based on or incorporate all or part of the contribution, the license to such derivative works to be of the same scope as the license of the original contribution.

2. The contributor acknowledges that the ISOC and IETF have no duty to publish or otherwise use or disseminate any contribution.

3. The contributor grants permission to reference the name(s) and address(es) of the contributor(s) and of the organization(s) he represents (if any).

4. The contributor represents that contribution properly acknowledge major contributors.

5. The contribuitor, the organization (if any) he represents and the owners of any proprietary rights in the contribution, agree that no information in the contribution is confidential and that the ISOC and its affiliated organizations may freely disclose any information in the contribution.

6. The contributor represents that he has disclosed the existence of any proprietary or intellectual property rights in the contribution that are reasonably and personally known to the contributor. The contributor does not represent that he personally knows of all potentially pertinent proprietary and intellectual property rights owned or claimed by the organization he represents (if any) or third parties.

7. The contributor represents that there are no limits to the contributor's ability to make the grants acknowledgments and agreements above that are reasonably and personally known to the contributor.

By ratifying this description of the IETF process the Internet Society warrants that it will not inhibit the traditional open and free access to IETF documents for which license and right have been assigned according to the procedures set forth in this section, including Internet-Drafts and RFCs. This warrant is perpetual and will not be revoked by the Internet Society or its successors or assigns.

10.3.2. Standards Track Documents

(A) Where any patents, patent applications, or other proprietary rights are known, or claimed, with respect to any specification on the standards track, and brought to the attention of the IESG, the IESG shall not advance the specification without including in the document a note indicating the existence of such rights, or claimed rights. Where implementations are required before advancement of a specification, only implementations that have, by statement of the implementors, taken adequate steps to comply with any such rights, or claimed rights, shall be considered for the purpose of showing the adequacy of the specification.

(B) The IESG disclaims any responsibility for identifying the existence of or for evaluating the applicability of any claimed copyrights, patents, patent applications, or other rights in the fulfilling of the its obligations under (A), and will take no position on the validity or scope of any such rights.

(C) Where the IESG knows of rights, or claimed rights under (A), the IETF Executive Director shall attempt to obtain from the claimant of such rights, a written assurance that upon approval by the IESG of the relevant Internet standards track specification(s), any party will be able to obtain the right to implement, use and distribute the technology or works when implementing, using or distributing technology based upon the specific specification(s) under openly specified, reasonable, non-discriminatory terms. The Working Group proposing the use of the technology with respect to which the proprietary rights are claimed may assist the IETF Executive Director in this effort. The results of this procedure shall not affect advancement of a specification along the standards track, except that the IESG may defer approval where a delay may facilitate the obtaining of such assurances. The results will, however, be recorded by the IETF Executive Director, and made available. The IESG may also direct that a summary of the results be included in any RFC published containing the specification.

10.3.3 Determination of Reasonable and Non-discriminatory Terms

The IESG will not make any explicit determination that the assurance of reasonable and non-discriminatory terms for the use of a technology has been fulfilled in practice. It will instead use the normal requirements for the advancement of Internet Standards to verify that the terms for use are reasonable. If the two unrelated implementations of the specification that are required to advance from Proposed Standard to Draft Standard have been produced by different organizations or individuals or if the "significant implementation and successful operational experience" required to advance from Draft Standard to Standard has been achieved the assumption is that the terms must be reasonable and to some degree, non-discriminatory. This assumption may be challenged during the Last-Call period.

10.4. Notices

(A) Standards track documents shall include the following notice:

"The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat."

(B) The IETF encourages all interested parties to bring to its attention, at the earliest possible time, the existence of any intellectual property rights pertaining to Internet Standards. For this purpose, each standards document shall include the following invitation:

"The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director."

(C) The following copyright notice and disclaimer shall be included in all ISOC standards-related documentation:

"Copyright (C) The Internet Society (date). 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 implmentation 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."

(D) Where the IESG is aware at the time of publication of proprietary rights claimed with respect to a standards track document, or the technology described or referenced therein, such document shall contain the following notice:

> "The IETF has been notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights."

11. ACKNOWLEDGMENTS

There have been a number of people involved with the development of the documents defining the IETF Standards Process over the years. The process was first described in RFC 1310 then revised in RFC 1602 before the current effort (which relies heavily on its predecessors). Specific acknowledgments must be extended to Lyman Chapin, Phill Gross and Christian Huitema as the editors of the previous versions, to Jon Postel and Dave Crocker for their inputs to those versions, to Andy Ireland, Geoff Stewart, Jim Lampert, and Dick Holleman for their reviews of the legal aspects of the procedures described herein, and to John Stewart, Robert Elz and Steve Coya for their extensive input on the final version.

In addition much of the credit for the refinement of the details of the IETF processes belongs to the many members of the various incarnations of the POISED Working Group.

12. SECURITY CONSIDERATIONS

Security issues are not discussed in this memo.

13. REFERENCES

[1] Postel, J., "Internet Official Protocol Standards", STD 1, USC/Information Sciences Institute, March 1996.

[2] ANSI, Coded Character Set -- 7-Bit American Standard Code for Information Interchange, ANSI X3.4-1986.

[3] Reynolds, J., and J. Postel, "Assigned Numbers", STD 2, USC/Information Sciences Institute, October 1994.

[4] Postel, J., "Introduction to the STD Notes", RFC 1311, USC/Information Sciences Institute, March 1992.

[5] Postel, J., "Instructions to RFC Authors", RFC 1543, USC/Information Sciences Institute, October 1993.

[6] Huitema, C., J. Postel, and S. Crocker "Not All RFCs are Standards", RFC 1796, April 1995.

14. DEFINITIONS OF TERMS

IETF Area - A management division within the IETF. An Area consists of Working Groups related to a general topic such as routing. An Area is managed by one or two Area Directors.

Area Director - The manager of an IETF Area. The Area Directors along with the IETF Chair comprise the Internet Engineering Steering Group (IESG).

File Transfer Protocol (FTP) - An Internet application used to transfer files in a TCP/IP network.

gopher - An Internet application used to interactively select and retrieve files in a TCP/IP network.

Internet Architecture Board (IAB) - An appointed group that assists in the management of the IETF standards process.

Internet Engineering Steering Group (IESG) - A group comprised of the IETF Area Directors and the IETF Chair. The IESG is responsible for the management, along with the IAB, of the IETF and is the standards approval board for the IETF.

interoperable - For the purposes of this document, "interoperable" means to be able to interoperate over a data communications path.

Last-Call - A public comment period used to gage the level of consensus about the reasonableness of a proposed standards action. (see section 6.1.2) online - Relating to information made available over the Internet. When referenced in this document material is said to be online when it is retrievable without restriction or undue fee using standard Internet applications such as anonymous FTP, gopher or the WWW.

Working Group - A group chartered by the IESG and IAB to work on a specific specification, set of specifications or topic.

15. AUTHOR'S ADDRESS

Scott O. Bradner
Harvard University
Holyoke Center, Room 813
1350 Mass. Ave.
Cambridge, MA 02138
USA

Phone: +1 617 495 3864
EMail: sob@harvard.edu

APPENDIX A: GLOSSARY OF ACRONYMS

ANSI:	American National Standards Institute
ARPA:	(U.S.) Advanced Research Projects Agency
AS:	Applicability Statement
FTP:	File Transfer Protocol
ASCII:	American Standard Code for Information Interchange
ITU-T:	Telecommunications Standardization sector of the International Telecommunication Union (ITU), a UN treaty organization; ITU-T was formerly called CCITT.
IAB:	Internet Architecture Board
IANA:	Internet Assigned Numbers Authority
IEEE:	Institute of Electrical and Electronics Engineers
ICMP:	Internet Control Message Protocol
IESG:	Internet Engineering Steering Group
IETF:	Internet Engineering Task Force
IP:	Internet Protocol
IRSG	Internet Research Steering Group
IRTF:	Internet Research Task Force
ISO:	International Organization for Standardization
ISOC:	Internet Society
MIB:	Management Information Base
OSI:	Open Systems Interconnection
RFC:	Request for Comments
TCP:	Transmission Control Protocol
TS:	Technical Specification
WWW:	World Wide Web

Exhibit E

Network Working Group
Request for Comments: 2806
Category: Standards Track

A. Vaha-Sipila
Nokia
April 2000

URLs for Telephone Calls

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Abstract

This document specifies URL (Uniform Resource Locator) schemes "tel", "fax" and "modem" for specifying the location of a terminal in the phone network and the connection types (modes of operation) that can be used to connect to that entity. This specification covers voice calls (normal phone calls, answering machines and voice messaging systems), facsimile (telefax) calls and data calls, both for POTS and digital/mobile subscribers.

Table of Contents

1. Introduction

1.1 New URL schemes

This specification defines three new URL schemes: "tel", "fax" and "modem". They are intended for describing a terminal that can be contacted using the telephone network. The description includes the subscriber (telephone) number of the terminal and the necessary parameters to be able to successfully connect to that terminal.

The "tel" scheme describes a connection to a terminal that handles normal voice telephone calls, a voice mailbox or another voice messaging system or a service that can be operated using DTMF tones.

The "fax" scheme describes a connection to a terminal that can handle telefaxes (facsimiles). The name (scheme specifier) for the URL is "fax" as recommended by [E.123].

The "modem" scheme describes a connection to a terminal that can handle incoming data calls. The term "modem" refers to a device that does digital-to-analog and analog-to-digital conversions; in addition to these, a "modem" scheme can describe a fully digital connection.

The notation for phone numbers is the same which is specified in [RFC2303] and [RFC2304]. However, the syntax definition is a bit different due to the fact that this document specifies URLs whereas [RFC2303] and [RFC2304] specify electronic mail addresses. For example, "/" (used in URLs to separate parts in a hierarchical URL [RFC2396]) has been replaced by ";". In addition, this URL scheme has been synchronized with [RFC2543].

When these URLs are used, the number of parameters should be kept to the minimum, unless this would make the context of use unclear. Having a short URL is especially important if the URL is intended to be shown to the end user, printed, or otherwise distributed so that it is visible.

1.2 Formal definitions

The ABNF (augmented Backus-Naur form) notation used in formal definitions follows [RFC2234]. This specification uses elements from the 'core' definitions (Appendix A of [RFC2234]). Some elements have been defined in previous RFCs. If this is the case, the RFC in question has been referenced in comments.

Note on non-unreserved characters [RFC2396] in URLs: the ABNF in this document specifies strings of raw, unescaped characters. If those characters are present in a URL, and are not unreserved [RFC2396], they MUST be escaped as explained in [RFC2396] prior to using the URL. In addition, when parsing a URL, it must be noted that some characters may have been escaped.

An example: ABNF notation "%x20" means a single octet with a hexadecimal value of "20" (in US-ASCII, a space character). This must be escaped in a URL, and it becomes "%20".

In addition, the ABNF in this document only uses lower case. The URLs are case-insensitive (except for the <future-extension> parameter, whose case-sensitivity is application-specific).

1.3 Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Compliant software MUST follow this specification.

2. URL schemes for telephone calls

2.1 Applicability

In this document, "local entity" means software and hardware that can detect and parse one or more of these URLs and possibly place a call to a remote entity, or otherwise utilize the contents of the URL.

These URL schemes are used to direct the local entity to place a call using the telephone network, or as a method to transfer or store a phone number plus other relevant data. The network in question may be

a landline or mobile phone network, or a combination of these. If the phone network differentiates between (for example) voice and data calls, or if the local entity has several different telecommunications equipment at its disposal, it is possible to specify which kind of call (voice/fax/data) is requested. The URL can also contain information about the capabilities of the remote entity, so that the connection can be established successfully.

The "tel", "fax" and "modem" URL schemes defined here do not use the hierarchical URL syntax; there are no applicable relative URL forms. The URLs are always case-insensitive, except for the <future-extension> parameter (see below), whose case-sensitivity is application specific. Characters in the URL MUST be escaped when needed as explained in [RFC2396].

2.2 "tel" URL scheme

The URL syntax is formally described as follows. For the basis of this syntax, see [RFC2303].

```
telephone-url         = telephone-scheme ":"
                        telephone-subscriber
telephone-scheme      = "tel"
telephone-subscriber  = global-phone-number / local-phone-number
global-phone-number   = "+" base-phone-number [isdn-subaddress]
                        [post-dial] *(area-specifier /
                        service-provider / future-extension)
base-phone-number     = 1*phonedigit
local-phone-number    = 1*(phonedigit / dtmf-digit /
                        pause-character) [isdn-subaddress]
                        [post-dial] area-specifier
                        *(area-specifier / service-provider /
                        future-extension)
isdn-subaddress       = ";isub=" 1*phonedigit
post-dial             = ";postd=" 1*(phonedigit /
                        dtmf-digit / pause-character)
area-specifier        = ";" phone-context-tag "=" phone-context-ident
phone-context-tag     = "phone-context"
phone-context-ident   = network-prefix / private-prefix
network-prefix        = global-network-prefix / local-network-prefix
global-network-prefix = "+" 1*phonedigit
local-network-prefix  = 1*(phonedigit / dtmf-digit / pause-character)
private-prefix        = (%x21-22 / %x24-27 / %x2C / %x2F / %x3A /
                        %x3C-40 / %x45-4F / %x51-56 / %x58-60 /
                        %x65-6F / %x71-76 / %x78-7E)
                        *(%x21-3A / %x3C-7E)
                        ; Characters in URLs must follow escaping rules
                        ; as explained in [RFC2396]
```

```
                        ; See sections 1.2 and 2.5.2
service-provider      = ";" provider-tag "=" provider-hostname
provider-tag          = "tsp"
provider-hostname     = domain ; <domain> is defined in [RFC1035]
                        ; See section 2.5.10
future-extension      = ";" 1*(token-char) ["=" ((1*(token-char)
                        ["?" 1*(token-char)]) / quoted-string )]
                        ; See section 2.5.11 and [RFC2543]
token-char            = (%x21 / %x23-27 / %x2A-2B / %x2D-2E / %x30-39
                        / %x41-5A / %x5E-7A / %x7C / %x7E)
                        ; Characters in URLs must follow escaping rules
                        ; as explained in [RFC2396]
                        ; See sections 1.2 and 2.5.11
quoted-string         = %x22 *( "\" CHAR / (%x20-21 / %x23-7E
                        / %x80-FF )) %x22
                        ; Characters in URLs must follow escaping rules
                        ; as explained in [RFC2396]
                        ; See sections 1.2 and 2.5.11
phonedigit            = DIGIT / visual-separator
visual-separator      = "-" / "." / "(" / ")"
pause-character       = one-second-pause / wait-for-dial-tone
one-second-pause      = "p"
wait-for-dial-tone    = "w"
dtmf-digit            = "*" / "#" / "A" / "B" / "C" / "D"
```

The URL starts with <telephone-scheme>, which tells the local entity that what follows is a URL that should be parsed as described in this document. After that, the URL contains the phone number of the remote entity. Phone numbers can also contain subaddresses, which are used to identify different remote entities under the same phone number. If a subaddress is present, it is appended to the phone number after ";isub=". Phone numbers can also contain a post-dial sequence. This is what is often used with voice mailboxes and other services that are controlled by dialing numbers from your phone keypad while the call is in progress. The <post-dial> sequence describes what and when the local entity should send to the phone line.

Phone numbers can be either "global" or "local". Global numbers are unambiguous everywhere. Local numbers are usable only within a certain area, which is called "context", see section 2.5.2.

Local numbers always have an <area-specifier>, which specifies the context in which the number is usable (the same number may have different interpretation in different network areas). The context can be indicated with three different prefixes. A <global-network-prefix> indicates that the number is valid within a numbering area whose global numbers start with <global-network-prefix>. Similarly, <local-network-prefix> means that the number is valid within a

numbering area whose numbers (or dial strings) start with it. A <private-prefix> is a name of a context. The local entity must have knowledge of this private context to be able to deduce whether it can use the number, see section 2.5.2. Additional information about the phone number's usage can be included by adding the name of the telephony services provider in <service-provider>, see section 2.5.10.

The <future-extension> mechanism makes it possible to add new parameters to this URL scheme. See section 2.5.11.

The <private-prefix>, <token-char> and <quoted-string> nonterminals may seem a bit complex at first, but they simply describe the set of octets that are legal in those nonterminals. Some octets may have to be escaped, see [RFC2396].

2.3 "fax" URL scheme

The URL syntax is formally described as follows (the definition reuses nonterminals from the above definition). For the basis of this syntax, see [RFC2303] and [RFC2304].

```
fax-url          = fax-scheme ":" fax-subscriber
fax-scheme       = "fax"
fax-subscriber   = fax-global-phone / fax-local-phone
fax-global-phone = "+" base-phone-number [isdn-subaddress]
                   [t33-subaddress] [post-dial]
                   *(area-specifier / service-provider /
                   future-extension)
fax-local-phone  = 1*(phonedigit / dtmf-digit /
                   pause-character) [isdn-subaddress]
                   [t33-subaddress] [post-dial]
                   area-specifier
                   *(area-specifier / service-provider /
                   future-extension)
t33-subaddress   = ";tsub=" 1*phonedigit
```

The fax: URL is very similar to the tel: URL. The main difference is that in addition to ISDN subaddresses, telefaxes also have an another type of subaddress, see section 2.5.8.

2.4 "modem" URL scheme

The URL syntax is formally described as follows (the definition reuses nonterminals from the above definitions). For the basis of this syntax, see [RFC2303].

```
modem-url          = modem-scheme ":" remote-host
modem-scheme       = "modem"
remote-host        = telephone-subscriber *(modem-params
                     / recommended-params)
modem-params       = ";type=" data-capabilities
recommended-params = ";rec=" data-capabilities
data-capabilities  = accepted-modem ["?" data-bits parity
                     stop-bits]
accepted-modem     = "V21" / "V22" / "V22b" /
                     "V23" / "V26t" / "V32" /
                     "V32b" / "V34" / "V90" /
                     "V110" / "V120" / "B103" /
                     "B212" / "X75" /
                     "vnd." vendor-name "." modem-type
data-bits          = "7" / "8"
parity             = "n" / "e" / "o" / "m" / "s"
stop-bits          = "1" / "2"
vendor-name        = 1*(ALPHA / DIGIT / "-" / "+")
modem-type         = 1*(ALPHA / DIGIT / "-" / "+")
```

The modem: URL scheme is also very similar to both the tel: and fax: schemes, but it adds the description of the capabilities of the remote entity. Minimum required compliance is listed in <modem-params> and recommended compliance is listed in <recommended-params>. For details, see section 2.5.9.

2.5 Parsing telephone, fax and modem URLs

2.5.1 Call type

The type of call is specified by the scheme specifier. "Tel" means that a voice call is opened. "Fax" indicates that the call should be a facsimile (telefax) call. "Modem" means that it should be a data call. Not all networks differentiate between the types of call; in this case, the scheme specifier indicates the telecommunications equipment type to use.

2.5.2 Phone numbers and their scope

<telephone-subscriber> and <fax-subscriber> indicate the phone number to be dialed. The phone number can be written in either international or local notation. All phone numbers SHOULD always be written in the international form if there is no good reason to use the local form.

Not all numbers are valid within all numbering areas. The <area-specifier> parameter, which is mandatory for local numbers, is used to indicate the locale within which this number is valid, or to qualify the phone number so that it may be used unambiguously. The

<area-specifier> can take three forms: <global-network-prefix>, <local-network-prefix> or <private-prefix>. These are used to describe the validity area of the phone number either in global numbering plan, local numbering plan, or in a private numbering plan, respectively.

If <area-specifier> is present, the local entity MUST NOT attempt to call out using the phone number if it cannot originate the call within the specified locale. If a <local-phone-number> is used, an <area-specifier> MUST be included as well.

There can be multiple instances of <area-specifier>. In this case, the number is valid in all of the given numbering areas.

The global prefix form is intended to act as the outermost context for a phone number, so it will start with a "+", followed by some part of an E.164 number. It also specifies the region in which the phone number is valid. For example, if <global-network-prefix> is "+358", the given number is valid only within Finland (country code 358) - even if it is a <global-phone-number>.

The local prefix form is intended to act as an intermediate context in those situations where the outermost context for a phone number is given by another means. One example of use is where the local entity is known to originate calls only within the North American Number Plan Area, so an "outermost" phone context can be assumed. The local context could, for example, be used to indicate the area code within which an associated phone number is situated. Thus "tel:456-7890;phone-context=213" would suffice to deliver a call to the telephone number "+1-213-456-7890". Note that the version including the <phone-context> implies further that the call can only be originated within the "area code 213" region.

The <private-prefix> form is intended for use in those situations where the context cannot be expressed with a start of a global phone number or a dialing string. The <private-prefix> is actually a name of a private context. The creator of the URL and the local entity have been configured to recognize this name, and as such they can interpret the number and know how they can utilize the number. For example, a private network numbering plan may be indicated by the name "X-COMPANY-NET", but the private dialling plan from the locales of the sender of the telephony URL and the local entity are different. The syntax of these tokens will be left for future specification. The ABNF above specifies the accepted characters that can be a part of <private-prefix>.

Unless the sender is absolutely sure that they share the same private network access digit string with the local entity, then they MUST NOT use a dialling plan number (a local phone number, or one qualified by a local context), as the result may be incorrect. Instead, they SHOULD use a global number, or if that is not possible, a private context as the last resort. If the local entity does not support dialling into the private network indicated by that context, then the request MUST be rejected. If it does, then it will use the access digit string appropriate for its locale.

Note that the use of <area-specifier> is orthogonal to use of the telephony service provider parameter (see 2.5.10); it qualifies the phone number, whilst the <service-provider> parameter indicates the carrier to be used for the call attempt.

For example, a large company may have private network interconnections between its sites, as well as connections to the Global Switched Telephone Network. A phone number may be given in "public network" form, but with a <service-provider> indicating that the call should be carried over the corporate network.

Conversely, it would be possible to represent a phone number in private network form, with a private context to indicate this, but indicate a public telephony service provider. This would request that the user agent convert the private network number plan address into a form that can be carried using the selected service provider.

Any telephone number MUST contain at least one <phonedigit> or <dtmf-digit>, that is, subscriber numbers consisting only of pause characters are not allowed.

International numbers MUST begin with the "+" character. Local numbers MUST NOT contain that character. International numbers MUST be written with the country (CC) and national (NSN) numbers as specified in [E.123] and [E.164]. International numbers have the property of being totally unambiguous everywhere in the world if the local entity is properly configured.

Local numbers MAY be used if the number only works from inside a certain geographical area or a network. Note that some numbers may work from several networks but not from the whole world - these SHOULD be written in international form, with a set of <area-specifier> tags and optional <service-provider> parameters. URLs containing local phone numbers should only appear in an environment where all local entities can get the call successfully set up by passing the number to the dialing entity "as is". An example could be a company intranet, where all local entities are located under a the same private telephone exchange. If local phone numbers are used, the document in which they are present SHOULD contain an indication of the context in which they are intended to be used, and an appropriate <area-specifier> SHOULD be present in the URL.

In some regions, it is popular to write phone numbers using alphabetic characters which correspond to certain numbers on the telephone keypad. Letters in <dtmf-digit> characters do not have anything to do with this, nor is this method supported by these URL schemes.

It should also be noted that implementations MUST NOT assume that telephone numbers have a maximum, minimum or fixed length, or that they would always begin with a certain number. Implementors are encouraged to familiarize themselves with the international standards.

2.5.3 Separators in phone numbers

All <visual-separator> characters MUST be ignored by the local entity when using the URL. These characters are present only to aid readability: they MUST NOT have any other meaning. Note that although [E.123] recommends the use of space (SP) characters as the separators in printed telephone numbers, spaces MUST NOT be used in phone numbers in URLs as the space character cannot be used in URLs without escaping it.

2.5.4 Converting the number to the local numbering scheme

After the telephone number has been extracted, it can be converted to the local dialing convention. (For example, the "+" character might be replaced by the international call prefix, or the international and trunk prefixes might be removed to place a local call.) Numbers that have been specified using <local-phone> or <fax-local-phone> MUST be used by the local entity "as is", without any conversions, unless the local entity decides to utilize the information in an optional <service-provider> parameter.

2.5.5 Sending post-dial sequence after call setup

The number may contain a <post-dial> sequence, which MUST be dialled using Dual Tone Multifrequency (DTMF) in-band signalling or pulse dialing after the call setup is complete. If the user agent does not support DTMF or pulse dialing after the call has been set up, <post-dial> MUST be ignored. In that case, the user SHOULD be notified.

2.5.6 Pauses in dialing and post-dial sequence

A local phone number or a post-dial sequence may contain <pause-character> characters which indicate a pause while dialing ("p"), or a wait for dial tone ("w").

Local entities MAY support this method of dialing, and the final interpretation of these characters is left to the local entity. It is RECOMMENDED that the length of each pause is about one second.

If it is not supported, local entities MUST ignore everything in the dial string after the first <pause-character> and the user SHOULD be notified. The user or the local entity MAY opt not to place a call if this feature is not supported and these characters are present in the URL.

Any <dtmf-digit> characters and all dial string characters after the first <pause-character> or <dtmf-digit> SHOULD be sent to line using DTMF (Dual Tone Multifrequency) in-band signaling, even if dialing is done using direct network signaling (a digital subscriber loop or a mobile phone). If the local infrastructure does not support DTMF codes, the local entity MAY opt to use pulse dialing. However, it should be noted that certain services which are controlled using DTMF tones cannot be controlled with pulse dialing. If pulse dialing is used, the user SHOULD be notified.

2.5.7 ISDN subaddresses

A phone number MAY also contain an <isdn-subaddress> which indicates an ISDN subaddress. The local entity SHOULD support ISDN subaddresses. These addresses are sent to the network by using a method available to the local entity (typically, ISDN subscribers send the address with the call setup signalling). If ISDN subaddressing is not supported by the caller, <isdn-subaddress> MUST be ignored and the user SHOULD be notified. The user or the local entity MAY opt not to place a call if this feature is not supported.

2.5.8 T.33 subaddresses

A fax number MAY also contain a <t33-subaddress>, which indicates the start of a T.33 subaddress [T.33]. Local entities SHOULD support this. Otherwise <t33-subaddress> MUST be ignored and the user SHOULD be notified. The user or the local entity MAY opt not to place a call if this feature is not supported.

2.5.9 Data call parameters

<modem-params> indicate the minimum compliance required from the local entity to be able to connect to the remote entity. The minimum compliance is defined as being equal to or a superset of the capabilities of the listed modem type. There can be several <modem-param> parameters, in which case compliance to any one of them will be accepted. <recommended-params> indicates the recommended compliance required from the local entity. This is typically the fastest and/or the most reliable modem type supported by the modem pool. The local entity can use this information to select the best number from a group of modem URLs. There can be several recommended modem types, which are equally desirable from the modem pool's point of view. <recommended-params> MAY NOT conflict with <modem-params>. If they do, the local entity MUST ignore the <recommended-params>.

The local entity MUST call out using compatible hardware, or request that the network provides such a service.

For example, if the local entity only has access to a V.22bis modem and the URL indicates that the minimum acceptable connection is V.32bis, the local entity MUST NOT try to connect to the remote host since V.22bis is a subset of V.32bis. However, if the URL lists V.32 as the minimum acceptable connection, the local entity can use V.32bis to create a connection since V.32bis is a superset of V.32.

This feature is present because modem pools often have separate numbers for slow modems and fast modems, or have different numbers for analog and ISDN connections, or may use proprietary modems that are incompatible with standards. It is somewhat analogous to the connection type specifier (typecode) in FTP URLs [RFC1738]: it provides the local entity with information that can not be deduced from the scheme specifier, but is helpful for successful operation.

This also means that the number of data and stop bits and parity MUST be set according to the information given in the URL, or to default values given in this document, if the information is not present.

The capability tokens are listed below. If capabilities suggest that it is impossible to create a connection, the connection MUST NOT be created.

If new modem types are standardized by ITU-T, this list can be extended with those capability tokens. Tokens are formed by taking the number of the standard and joining together the first letter (for example, "V"), number (for example, 22) and the first letter of the postfix (for example "bis" would become "b").

Proprietary modem types MUST be specified using the 'vendor naming tree', which takes the form "vnd.x.y", in which "x" is the name of the entity from which the specifications for the modem type can be acquired and "y" is the type or model of the modem. Vendor names MUST share the same name space with vendor names used in MIME types [RFC2048]. Submitting the modem types to ietf-types list for review is strongly recommended.

New capabilities MUST always be documented in an RFC, and they MUST refer to this document or a newer version of it. The documentation SHOULD also list the existing modem types with which the newly defined modem type is compatible with.

Capability	Explanation
V21	ITU-T V.21
V22	ITU-T V.22
V22b	ITU-T V.22bis
V23	ITU-T V.23
V26t	ITU-T V.26ter
V32	ITU-T V.32
V32b	ITU-T V.32bis
V34	ITU-T V.34
V90	ITU-T V.90
V110	ITU-T V.110
V120	ITU-T V.120
X75	ITU-T X.75
B103	Bell 103
B212	Bell 212
Data bits: "8" or "7"	The number of data bits. If not specified, defaults to "8".
Parity: "n", "e", "o", "m", "s"	Parity. None, even, odd, mark or space parity, respectively. If not specified, defaults to "n".
Stop bits: "1" or "2"	The number of stop bits. If not specified, defaults to "1".

2.5.10 Telephony service provider identification

It is possible to indicate the identity of the telephony service provider for the given phone number. <service-provider> MAY be used by the user-agent to place the call using this network, to enhance the user interface, for billing estimates or to otherwise optimize its functionality. It MAY also be ignored by the user-agent. <service-provider> consists of a fully qualified Internet domain name of the telephony service provider, for example ";tsp=terrifictelecom.com". The syntax of the domain name follows Internet domain name rules and is defined in [RFC1035].

2.5.11 Additional parameters

In addition to T.33 and ISDN subaddresses, modem types and area specifiers, future extensions to this URL scheme may add other additional parameters (<future-extension> in the BNF) to these URLs. These parameters are added to the URL after a semicolon (";"). Implementations MUST be prepared to handle additional and/or unknown parameters gracefully. Implementations MUST NOT use the URL if it contains unknown parameters, as they may be vital for the correct interpretation of the URL. Instead, the implementation SHOULD report an error.

For example, <future-extension> can be used to store application-specific additional data about the phone number, its intended use, or any conversions that have been applied to the number. Whenever a <future-extension> is used in an open environment, its syntax and usage MUST be properly documented in an RFC.

<future-extension> nonterminal a rephrased version of, and compatible with the <other-param> as defined in [RFC2543] (which actually borrows BNF from an earlier version of this specification).

2.6 Examples of Use

```
tel:+358-555-1234567
```

This URL points to a phone number in Finland capable of receiving voice calls. The hyphens are included to make the number more human-readable: country and area codes have been separated from the subscriber number.

```
fax:+358.555.1234567
```

The above URL describes a phone number which can receive fax calls. It uses dots instead of hyphens as separators, but they have no effect on the functionality.

```
modem:+3585551234567;type=v32b?7e1;type=v110
```

This phone number belongs to an entity which is able to receive data calls. The local entity may opt to use either a ITU-T V.32bis modem (or a faster one, which is compatible with V.32bis), using settings of 7 data bits, even parity and one stop bit, or an ISDN connection using ITU-T V.110 protocol.

```
tel:+358-555-1234567;postd=pp22
```

The above URL instructs the local entity to place a voice call to +358-555-1234567, then wait for an implementation-dependent time (for example, two seconds) and emit two DTMF dialing tones "2" on the line (for example, to choose a particular extension number, or to invoke a particular service).

```
tel:0w003585551234567;phone-context=+3585551234
```

This URL places a voice call to the given number. The number format is intended for local use: the first zero opens an outside line, the "w" character waits for a second dial tone, and the number already has the international access code appended to it ("00"). This kind of phone number MUST NOT be used in an environment where all users of this URL might not be able to successfully dial out by using this number directly. However, this might be appropriate for pages in a company intranet. The <area-specifier> which is present hints that the number is usable only in an environment where the local entity's phone number starts with the given string (perhaps singling out a company-wide block of telephone numbers).

```
tel:+1234567890;phone-context=+1234;vnd.company.option=foo
```

The URL describes a phone number which, even if it is written in its international form, is only usable within the numbering area where phone numbers start with +1234. There is also a proprietary extension "vnd.company.option", which has the value "foo". The meaning of this extension is application-specific. Note that the order of these parameters (phone-context and vnd.company.option) is irrelevant.

2.7 Rationale behind the syntax

2.7.1 Why distinguish between call types?

URLs locate resources, which in this case is some telecommunications equipment at a given phone number. However, it is not necessarily enough to know the subscriber number in order to successfully communicate with that equipment. Digital phone networks distinguish between voice, fax and data calls (and possibly other types of calls, not discussed in this specification). To be able to successfully connect to, say, a fax machine, the caller may have to specify that a fax call is being made. Otherwise the call might be routed to the voice number of the subscriber. In this sense, the call type is an integral part of the 'location' of the target resource.

The reason to have the call type in the scheme specifier is to make the URL simple to remember and use. Making it a parameter, much like the way modem parameters are handled now, will substantially reduce the human readability of this URL.

2.7.2 Why "tel" is "tel"?

There has been discussion on whether the scheme name "tel" is appropriate. To summarize, these are the points made against the other proposals.

callto	URL schemes locate a resource and do not specify an action to be taken.
telephone	Too long. Also, "tel" considered to be a more international form.
phone	Was countered on the basis that "tel" is more internationally acceptable.

2.7.3 Why to use E.164-style numbering?

E.164 refers to international telephone numbers, and the string of digits after the country code is usually a national matter. In any case, phone numbers are usually written as a simple string of numbers everywhere. Because of this, the syntax in this specification is intuitively clear to most people. This is the usual way to write phone numbers in business cards, advertisements, telephone books and so on.

It should be noted that phone numbers may have 'hierarchical' characteristics, so that one could build a 'forest' of phone numbers with country codes as roots, area codes as branches and subscriber numbers as leaves. However, this is not always the case. Not all areas have area codes; some areas may have different area codes depending on how one wants to route the call; some numbers must always be dialled "as is", without prepending area or country codes (notably emergency numbers); and area codes can and do change.

Usually, if something has a hierarchical structure, the URL syntax should reflect that fact. These URLs are an exception.

Also, when writing the phone number in the form described in this specification, the writer does not need to know which part of the number is the country code and which part is the area code. If a hierarchical URL would be used (with a "/" character separating the parts of the phone numbers), the writer of the URL would have to know which parts are which.

Finally, when phone numbers are written in the international form as specified here, they are unambiguous and can always be converted to the local dialing convention, given that the user agent has the knowledge of the local country and area codes.

2.7.4 Not everyone has the same equipment as you

There are several ways for the subscriber to dial a phone number:

- By pulse dialing. Typically old telephone exchanges. Usually this dialing method has only to be used to set up the call; after connecting to the remote entity, <post-dial> can be sent to the line using DTMF, because it will typically be processed by the remote entity, not the telephone network.

- By DTMF. These are the 'beeps' that you hear when you dial on most phones.

- By direct network signalling. ISDN subscribers and mobile phone users usually have this. There is no dial tone (or if there is, it is generated locally by the equipment), and the number of the called party is communicated to the telephone network using some network signalling method. After setting up the call, <post-dial> sequences are usually sent using DTMF codes.

2.7.5 Do not confuse numbers with how they are dialled

As an example, +123456789 will be dialled in many countries as 00123456789, where the leading "00" is a prefix for international calls. However, if a URL contains a local phone number 00123456789, the user-agent MUST NOT assume that this number is equal to a global phone number +123456789. If a user-agent received a telephony URL with a local number in it, it MUST make sure that it knows the context in which the local phone number is to be processed, or else the number MUST NOT be used. Equally, anyone sending a telephony URL MUST take into consideration that the recipient may have insufficient information about the phone number's context.

3. Comments on usage

These are examples of the recommended usage of this URL in HTML documents.

First of all, the number SHOULD be visible to the end user, if it is conceivable that the user might not have a local entity which is able to use these URLs.

```
  Telephone: <a href="tel:+3585551234567">+358-555-1234567</a>
```

Second, on a public HTML page, the telephone number in the URL SHOULD always be in the international form, even if the text of the link uses some local format.

Telephone: <a href="tel:+3585551234567">(0555) 1234567</a>

or even

For more info, call <a href="tel:+15554383785965">1-555-IETF-RULZ-OK</a>.

Moreover, if the number is a <local-phone-number>, and the scope of the number is not clear from the context in which the URL is displayed, a human-readable explanation SHOULD be included.

For customer service, dial <a href="tel:1234;phone-context=+358555">1234</a> (only from Terrific Telecom mobile phones).

4. References

[RFC1035] Mockapetris, P., "Domain Names - Implementation and Specification", STD 13, RFC 1035, November 1987.

[RFC1738] Berners-Lee, T., et al., "Uniform Resource Locators (URL)", RFC 1738, December 1994.

[RFC1866] Berners-Lee, T. and D. Connolly, "Hypertext Markup Language - 2.0", RFC 1866, November 1995.

[RFC2048] Freed, N., Klensin, J. and J. Postel, "Multipurpose Internet Mail Extensions (MIME) Part Four: Registration Procedures", RFC 2048, November 1996.

[RFC2119] Bradner, S., "Key Words for Use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2234] Crocker, D. and P. Overall, "Augmented BNF for Syntax Specifications: ABNF", RFC 2234, November 1997.

[RFC2303] Allocchio, C., "Minimal PSTN Address Format in Internet Mail", RFC 2303, March 1998.

[RFC2304] Allocchio, C., "Minimal FAX Address Format in Internet Mail", RFC 2304, March 1998.

[RFC2396] Berners-Lee, T., R. Fielding and L. Manister, "Uniform Resource Identifiers (URI): Generic Syntax", RFC 2396, August 1998.

[RFC2543] Handley, M., Schulzrinne, H., Schooler, E. and J. Rosenberg, "SIP: Session Initiation Protocol", RFC 2543, March 1999.

[E.123] ITU-T Recommendation E.123: Telephone Network and ISDN Operation, Numbering, Routing and Mobile Service: Notation for National and International Telephone Numbers. 1993.

[E.164] ITU-T Recommendation E.164/I.331 (05/97): The International Public Telecommunication Numbering Plan. 1997.

[T.33] ITU-T Recommendation T.33: Facsimile Routing Utilizing the Subaddress. 1996.

5. Security Considerations

It should be noted that the local entity SHOULD NOT call out without the knowledge of the user because of associated risks, which include

- call costs (including long calls, long distance calls, international calls and premium rate calls, or calls which do not terminate due to <post-dial> sequences that have been left out by the local entity)

- wrong numbers inserted on web pages by malicious users, or sent via e-mail, perhaps in direct advertising

- making the user's phone line unavailable (off-hook) for a malicious purpose

- opening a data call to a remote host, thus possibly opening a back door to the user's computer

- revealing the user's (possibly unlisted) phone number to the remote host in the caller identification data, and correlating the local entity's phone number with other information such as the e-mail or IP address

- using the same local number in different contexts, in which the number may have a different meaning

All of these risks MUST be taken into consideration when designing the local entity.

The local entity SHOULD have some mechanism that the user can use to filter out unwanted numbers. The local entity SHOULD NOT use rapid redialing of the number if it is busy to avoid the congestion of the (signaling) network. Also, the local entity SHOULD detect if the number is unavailable or if the call is terminated before the dialing string has been completely processed (for example, the call is terminated while waiting for user input) and not try to call again, unless instructed by the user.

6. Acknowledgements

Writing this specification would not have been possible without extensive support from many people.

Contributors include numerous people from IETF FAX, PINT, URI and URLREG mailing lists, as well as from World Wide Web Consortium and several companies, plus several individuals. Thanks to all people who offered criticism, corrections and feedback.

All phone numbers and company names used in the examples of this specification are fictional. Any similarities to real entities are coincidental.

7. Author's Address

Antti Vaha-Sipila
(quoted-printable: Antti V=E4h=E4-Sipil=E4)
Nokia Mobile Phones
P. O. Box 68
FIN-33721 Tampere
Finland

EMail: avs@iki.fi
antti.vaha-sipila@nokia.com

8. Full Copyright Statement

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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

Exhibit F

INTERNET-DRAFT
Expires 27-Aug-1998

A. Vaha-Sipila
Nokia
23-Feb-1998

URLs for Telephony
<draft-antti-telephony-url-04.txt>

Status of This Memo

Abstract

This document specifies URL (Uniform Resource Locator) schemes ''phone'', ''fax'' and ''modem'' for specifying the location of a terminal in the phone network and the connection types (modes of operation) that can be used to connect to that entity. This specification covers voice calls (normal phone calls, answering machines and voice messaging systems), facsimile (telefax) calls and data calls, both for POTS and digital/mobile subscribers.

Version History

| Changes to the previous versions are indicated by a bar in the
| left margin like in this section.

Contents

1. Introduction

1.1 New URL Schemes

URLs that designate phone or fax numbers that can be dialed have been brought forward in other Internet-Drafts. However, none of these has reached the RFC status. This document tries to remedy the situation. All interested parties are invited to submit comments on this Internet-Draft. Contact information can be found at the end of this document.

See also [CONV-URL] for more discussion on conversational URLs.

This specification defines three new URL schemes: "phone", "fax" and "modem". They are intended for describing a terminal that can be contacted using the telephone network. The description includes the subscriber (telephone) number of the terminal and the necessary parameters for successfully connecting to that terminal.

The "phone" scheme describes a connection to a terminal that handles normal voice telephone calls, a voice mailbox or another voice messaging system or a service that can be operated using DTMF codes.

The "fax" scheme describes a connection to a terminal that can handle telefaxes (facsimiles). The name (scheme specifier) for the URL is "fax" as recommended by [E.123].

The "modem" scheme describes a connection to a terminal that can handle incoming data calls. The term "modem" refers to a device that does digital-to-analog and analog-to-digital conversions; in addition to these, a "modem" scheme can describe a fully digital connection.

The notation for phone numbers is the same which is specified in [PSTN-ADDR] and [FAX-ADDR]. However, the syntax definition is a bit different due to the fact that this document specifies URLs whereas [PSTN-ADDR] and [FAX-ADDR] specify electronic mail addresses. For example, "/" (used in URLs to separate parts in a hierarchical URL [RFC1738]) has been replaced by ";".

1.2 Formal Definitions

Formal definitions follow [ABNF]. This specification uses elements from the 'core' definitions (Appendix A of [RFC2234]).

1.3 Requirements

Compliant software MUST follow this specification. Requirements are indicated by capitalized words as specified in [RFC2119].

2. URL Schemes for Telephone Calls

2.1 Applicability

In this document, "user agent" means software that can detect and parse one or more of these URLs and place a call to the remote terminal using hardware at its disposal.

These URL schemes are used to direct the user agent to place a call using the telephone network. The network in question may be a landline or mobile phone network. If the phone network differentiates between (for example) voice and data calls, or if the user agent has several different telecommunications equipment at its disposal, it is possible to specify which kind of call (voice/fax/data) is requested. The URL can also contain information about the capabilities of the remote entity, so that the connection can be established successfully.

None of the URL schemes do have a 'path' in them - they are always absolute. The URLs are always case-insensitive.

2.2 "phone" URL Scheme

The URL syntax is formally described as follows. For the basis of this syntax, see [PSTN-ADDR].

```
  telephone-url          = telephone-scheme ":"
                           telephone-subscriber
| telephone-scheme       = "phone"
  telephone-subscriber   = global-phone-number / local-phone-number
  global-phone-number    = "+" 1*phonedigit [isdn-subaddress]
                           [post-dial]
  local-phone-number     = 1*(phonedigit / dtmf-digit /
                           pause-character) [isdn-subaddress]
                           [post-dial]
  isdn-subaddress        = ";isub=" 1*phonedigit
  post-dial              = ";postd=" 1*(phonedigit / dtmf-digit
                           / pause-character)
  phonedigit             = DIGIT / visual-separator
  visual-separator       = "-" / "."
  pause-character        = one-second-pause / wait-for-dial-tone
  one-second-pause       = "p"
  wait-for-dial-tone     = "w"
  dtmf-digit             = "*" / "#" / "A" / "B" / "C" / "D"
```

2.3 "fax" URL Scheme

The URL syntax is formally described as follows (the definition reuses nonterminals from the definition above). For the basis of this syntax, see [PSTN-ADDR] and [FAX-ADDR].

```
fax-url                  = fax-scheme ":" fax-subscriber
fax-scheme               = "fax"
fax-subscriber           = fax-global-phone / fax-local-phone
fax-global-phone         = "+" 1*phonedigit [isdn-subaddress]
                           [t33-subaddress] [post-dial]
fax-local-phone          = 1*(phonedigit / dtmf-digit /
                           pause-character) [isdn-subaddress]
                           [t33-subaddress] [post-dial]
t33-subaddress           = ";tsub=" 1*phonedigit
```

2.4 "modem" URL Scheme

The URL syntax is formally described as follows. For the basis of this syntax, see [PSTN-ADDR].

```
modem-url                = modem-scheme ":" remote-host
modem-scheme             = "modem"
remote-host              = telephone-subscriber *modem-params
modem-params             = ";type=" data-capabilities
data-capabilities        = accepted-modem ["?" data-bits parity
                           stop-bits]
accepted-modem           = "V21" / "V22" / "V22b" /
                           "V23" / "V26" / "V32" /
                           "V32b" / "V34" / "V110" /
                           "V120" / "B103" / "B212" /
                           "X75"
data-bits                = "7" / "8"
parity                   = "n" / "e" / "o" / "m" / "s"
stop-bits                = "1" / "2"
```

2.5 Parsing telephone, fax and modem URLs

A. The type of call is specified by the scheme specifier. "phone" means that a voice call is opened. "Fax" indicates that the call should be a facsimile (telefax) call. "Modem" means that it should be a data call. Not all networks differentiate between the types of call; in this case, the scheme specifier indicates the telecommunications equipment type to use.

B. <telephone-subscriber> and <fax-subscriber> indicate the phone number to be dialed. The phone number can be written in either international or local notation. All phone numbers SHOULD always be written in the international form if there is no good reason to use the local form.

Any telephone number must contain at least one <DIGIT>, that is, subscriber numbers consisting only of non-numbers are not allowed.

International numbers MUST begin with the "+" character. Local numbers MUST NOT contain that character. International numbers MUST be written with the country (CC) and national (NSN) numbers as specified in [E.123] and [E.164]. Local numbers MAY be used if the number only works from inside a certain geographical area or a network. Note that some numbers may work from several networks but not from the whole world - these SHOULD be written in international form.

C. All <visual-separator> characters MUST be removed from the phone number by the user agent before using it do dial out. These cracaters are present only to aid readability: they MUST NOT have any other meaning. Note that although [E.123] recommends the use of space (SP) characters as the separators, spaces MUST NOT be used in these URLs.

D. After the telephone number has been extracted, it is converted to the format that the user agent can use to place the call. (For example, the "+" character might be replaced by the international call prefix, or the international and trunk prefixes might be removed to place a local call.) Numbers that have been specified using <local-phone> or <fax-local-phone> MUST be used by the user agent "as is", without any conversions.

E. The number may contain a <post-dial> sequence, which MUST be dialled using Dual Tone Multifrequency (DTMF) in-band signalling after the call setup is complete. If the user agent does not support DTMF, <post-dial> MUST be ignored. In that case, the user SHOULD be notified.

F. A local phone number or a post-dial sequence may contain <pause-character> characters which indicate a pause of 1 second while dialing ("p"), or a wait for dial tone ("w"). User agents SHOULD support this method of dialing. If it is not supported, user agents MUST ignore everything in the dial string after the first <pause-character> and the user SHOULD be notified. The user or the user agent MAY opt not to place a call if this feature is not supported.

Any <dtmf-digit> characters and all dial string characters after the first <pause-character> or <dtmf-digit> MUST be sent to line using DTMF (Dual Tone Multifrequency) in-band signaling, even if dialing is done using direct network signaling (a digital subscriber loop or a mobile phone).

G. A phone number MAY also contain an <isdn-subaddress> which indicates an ISDN subaddress. User agent SHOULD support ISDN subaddresses. These addresses are sent to the network by using a method available to the user agent (typically, ISDN subscribers send the address with the call setup signalling). If ISDN subaddressing is not supported by the caller, <isdn-subaddress> MUST be ignored and the user SHOULD be notified. The user or the user agent MAY opt not to place a call if this feature is not supported.

H. A fax number MAY also contain a <t33-subaddress>, which indicates the start of a T.33 subaddress [T.33]. User agents SHOULD support this. Otherwise <t33-subaddress> MUST be ignored and the user SHOULD be notified. The user or the user agent MAY opt not to place a call if this feature is not supported.

I. <modem-params> indicate the minimum compliance required from the user agent to be able to connect to the remote entity. The minimum compliance is defined as being equal to or a superset of the capabilities of the listed modem type.

The user agent MUST call out using compatible hardware, or request that the network provides such a service.

For example, if the user agent only has access to a V.22bis modem and the URL indicates that the minimum acceptable connection is V.32bis, the user agent MUST NOT try to connect to the remote host since V.22bis is a subset of V.32bis. However, if the URL lists V.32 as the minimum acceptable connection, the user agent can use V.32bis to create a connection since V.32bis is a superset of V.32.

This feature is present because modem pools often have separate numbers for slow modems and fast modems, or have different numbers for analog and ISDN connections, or may use proprietary modems that are incompatible with standards. It is somewhat analogous to the connection type specifier (typecode) in FTP URLs [RFC1738]: it provides the user agent with information that can not be deduced from the scheme specifier, but is helpful for successful operation.

This also means that the number of data and stop bits and parity MUST be set according to the information given in the URL or to default values, if the information is not present.

The capability tokens are listed below. If capabilities suggest that it is impossible to create a connection, the connection MUST NOT be created.

If new modem types are standardized by ITU-T, this list can be extended with those capability tokens. Tokens are formed by taking the name of the standard and joining together the first letter, number and the first letter of the postfix. New capabilities SHOULD then be documented in an RFC. New non-ITU-T capabilities (such as vendor-proprietary modem types) MUST be specified in a separate RFC.

Capability	Explanation
V21	ITU-T V.21
V22	ITU-T V.22
V22b	ITU-T V.22bis
V23	ITU-T V.23
V26t	ITU-T V.26ter
V32	ITU-T V.32
V32b	ITU-T V.32bis
V34	ITU-T V.34
V110	ITU-T V.110
V120	ITU-T V.120

X75	ITU-T X.75
B103	Bell 103
B212	Bell 212
Data bits: "8" or "7"	The number of data bits. If not specified, defaults to "8".
Parity: "n", "e", "o", "m", "s"	Parity. None, even, odd, mark or space parity, respectively. If not specified, defaults to "n".
Stop bits: "1" or "2"	The number of stop bits. If not specified, defaults to "1".

2.6 Examples of Use

```
phone:+358-55-1234567
```

This URL instructs the user agent to place a voice call to the specified number in Finland. The hyphens are included to make the number more human-readable: country and area codes have been separated from the subscriber number.

```
fax:+358.55.1234567
```

The above URL instructs the user agent to place a fax call to the specified number. It uses dots instead of hyphens as separators, but they have no effect on the functionality.

```
modem:+358551234567;type=v32b?7e1;type=v110
```

This URL instructs the user agent to place a data call to the specified number. The user agent may opt to use either a ITU-T V.32bis modem (or a faster one, which is compatible with V.32bis), using settings of 7 data bits, even parity and one stop bit, or an ISDN connection using ITU-T V.110 protocol.

```
phone:+358-55-1234567;postd=pp22
```

The above URL instructs the user agent to place a voice call to +358-55-1234567, then wait two seconds and emit two DMTF dialing tones "2" on the line (for example, to choose a particular extension number).

```
phone:0w003585512345 67
```

This URL places a voice call to the given number. The number format is intended for local use: the first zero opens an outside line, the "w" character waits for a second dial tone, and the number already has the international access code appended to it ("00"). This kind of phone number MUST NOT be used in an environment where all users of this URL might not be able to successfully dial out by using this number directly. However, this might be appropriate for pages in a company intranet.

3. References

[RFC2234] Augmented BNF for Syntax Specifications: ABNF. November 1997. D. Crocker et al. RFC 2234. <URL:ftp://ftp.ds.internic.net/rfc/rfc2234.txt>

[CONV-URL] Conversational Multimedia URLs. 1997. Pete Cordell. An Internet-Draft (work in progress). <URL:ftp://ftp.ds.internic.net/internet-drafts/draft-cordell-sg16-conv-url-00.txt>

[FAX-ADDR] Minimal FAX Address Format in Internet Mail. 1998. C. Allocchio. An Internet-Draft (work in progress). <URL:ftp://ftp.ds.internic.net/internet-drafts/draft-ietf-fax-minaddrfax-01.txt>

[PSTN-ADDR] Minimal PSTN Address Format in Internet Mail. 1998. C. Allocchio. An Internet-Draft (work in progress). <URL:ftp://ftp.ds.internic.net/internet-drafts/draft-ietf-fax-minaddrgen-02.txt>

[RFC1738] Uniform Resource Locators (URL). December 1994. T. Berners-Lee et al. RFC 1738. <URL:ftp://ftp.ds.internic.net/rfc/rfc1738.txt>

[RFC2119] Key Words for Use in RFCs to Indicate Requirement Levels. March 1997. S. Bradner. RFC 2119. <URL:ftp://ftp.ds.internic.net/rfc/rfc2119.txt>

[E.123] ITU-T Recommendation E.123: Telephone Network and ISDN Operation, Numbering, Routing and Mobile Service: Notation for National and International Telephone Numbers. 1993.

[E.164] ITU-T Recommendation E.164: Telephone Network and ISDN Operation, Numbering, Routing and Mobile Service: Numbering Plan for the ISDN Era. 1991.

[T.33] ITU-T Recommendation T.33: Facsimile Routing Utilizing the Subaddress. 1996.

4. Security Considerations

It should be noted that the user agent SHOULD NOT call out without the knowledge of the user because of associated risks, which include

- call costs (including long calls, long distance calls, international calls and prime rate calls)
- wrong numbers inserted on web pages by malicious users
- making the user's phone line unavailable (off-hook) for a malicious purpose
- opening a data call to a remote host, thus possibly opening a back door to the user's computer

The user agent SHOULD have some mechanism that the user can use to filter out unwanted numbers. The user agent SHOULD NOT use rapid redialing of the number if it is busy to avoid the congestion of the (signaling) network. Also, the user agent SHOULD detect if the number is unavailable or if the call is terminated before the dialing string has been completely processed (for example, the call is terminated while waiting for user input) and not try to call again, unless instructed by the user.

5. Authors' Addresses

Contact person for this specification:

Nokia Mobile Phones
Antti Vaha-Sipila
P. O. Box 68
FIN-33721 Tampere
Finland

Electronic mail: antti.vaha-sipila@nmp.nokia.com

Please include your name and electronic mail address in all communications. If you want to receive the newest version of this specification electronically, send mail to the address above.

This document expires on the 27th of August, 1998, or when a new version is released.

Exhibit G

INTERNET-DRAFT
Expires 28-Feb-1998

A. Vaha-Sipila
Nokia
26-Aug-1997

URLs for Telephony
<draft-antti-telephony-url-00.txt>

Status of This Memo

Abstract

This document specifies a URL (Uniform Resource Locator) scheme 'callto' for specifying a location for an entity in the phone network and the connection types (modes of operation) that can be used to connect to that entity. This specification covers voice calls (normal phone calls, answering machines and voice messaging systems), facsimile (telefax) calls and data calls, both for POTS and digital/mobile subscribers.

Contents

1. Introduction

1.1 New URL Schemes

This specification defines a new URL scheme, callto. It is intended for describing an entity that can be contacted using the telephone network. This URL scheme caters for voice calls (including DTMF based supplementary services), fax (facsimile) calls and data calls. URLs that designate phone or fax numbers that can be dialed have been brought forward in other (currently expired) Internet-Drafts. However, none of these has reached the RFC status. This document tries to remedy the situation. All interested parties are invited to submit comments on this Internet-Draft. Contact information can be found at the end of this document.

1.2 Formal Definitions

Rules are separated from definitions by an equal "=", literals are quoted with double quotes "", parentheses "(" and ")" are used to group elements, optional elements are enclosed in "[" and "]" brackets, and a set of elements where order is not significant is enclosed by "{" and "}" (note that this does not imply repetition; only that the elements in that group can appear in any order).

Elements may be preceded with n* to designate n repetitions of the following element; n defaults to 0. Single quotes '' are used to indicate elements that are not formally specified and are described in free text instead. Indentation indicates that the definition continues from the previous line.

1.3 Requirements

Compliant software MUST follow this specification. Requirements are indicated by capitalized words as specified in [RFC2119].

2. URL Schemes for Telephone Calls

2.1 Applicability

The callto URL scheme is used to direct the user agent to place a call using the telephone network. The network in question may be a landline or mobile phone network. If the phone network differentiates between (for example) voice and data calls, or if the user agent has several different telecommunications equipment at its disposal, it is possible to specify which kind of call (voice/fax/data) is requested. It is also possible to give information about the capabilities of the remote entity.

2.2 callto URL Scheme

The callto URL does not have a 'path' in it - it is always absolute. Everything is case-insensitive, expect for the "user-name" and "password", which may be case-sensitive in some environments.

Some characters are considered "reserved" and cannot be used in the free text fields (such as "user-name" and "password") without escaping them using URL encoding.

```
reserved                 = " " | ":" | "?" | ";" | "="
```

The URL syntax is formally described as follows:

```
callto-url               = scheme ":" scheme-specific-part
scheme                   = 'callto'
scheme-specific-part     = subscriber-id [type-specifier]
subscriber-id            = ["+"] phone-number
type-specifier           = [";type=" call-type]
call-type                = "voice" | "data" {[";params="
                           data-capabilities]
                           [";proto=" protocol-specifier]
                           [";user=" user-name] [";pass=" password]} |
                           "fax" [";params=" facsimile-capabilities]
phone-number             = 1*phonedigit [pause-character
                           *(phonedigit | dtmf-digit |
                           pause-character)]
phonedigit               = digit | "-"
pause-character          = "p" | "w"
digit                    = "0" | "1" | "2" | "3" | "4" | "5" |
                           "6" | "7" | "8" | "9"
dtmf-digit               = "*" | "#" | "A" | "B" | "C" | "D"
user-name                = 'a user name for authentication;
                           in URL-encoded notation'
password                 = 'a password for authentication;
                           in URL-encoded notation'
protocol-specifier       = "raw" | "ppp"
data-capabilities        = modem-type ["?" data-bits parity
                           stop-bits]
data-bits                = "7" | "8"
parity                   = "n" | "e" | "o" | "m" | "s"
stop-bits                = "1" | "2"
modem-type               = "V21" | "V22" | "V22b" |
                           "V23" | "V26" | "V32" |
                           "V32b" | "V34" | "V110" |
                           "V120" | "B103" | "B212" |
                           "X75"
facsimile-capabilities = "G1" | "G2" | "G3" | "G4"
```

"Subscriber-id" is the phone number to be dialed. This phone number SHOULD be written in international notation with country (CC) and area (NDC) codes, that is, both ITU-T CC and NSN elements SHOULD be present.

International numbers MUST begin with "+", which indicates that the number begins with a country code (CC). Hyphens are only to aid readability; they MUST NOT have any other meaning. Numbers that only work from inside a certain geographical area or from a certain network MUST NOT start with a "+". For the meaning of "p" and "w", see below.

"Type-specifier" specifies the call type (which kind of terminal equipment is used or what services are requested from the network). If not defined, it defaults to "voice". "Data-capabilities" is an optional field which indicates what kind of entity is answering at the other end. The "modem-type" field is present to indicate the line speed or protocol. The "modem-type" listed should be the one that offers the fastest connection. If several, incompatible modem types are connected to the same number, each of these is listed. Possible modem types are described below. Also, optional information on data bits, parity and stop bits may be provided for each "modem-type".

"Proto" is the definition of the protocol that is spoken over the data connection. If not specified, it defaults to "raw", which implies no higher-level protocol at either end (for example a terminal connection). "PPP" tells the user agent (client, caller) to start negotiating a PPP connection with the server.

"User-name" and "password" are the credentials used for authenticating the user. These direct the upper-level protocol or session that utilitises the newly created data connection. Use of this information is left to the user agent. When specified in a URL that directs the user agent to create a PPP connection this information SHOULD be used for PPP authentication. In a normal terminal ("raw") connection, this information MAY be utilised by automatic login scripts or a similar purpose.

"Facsimile-capabilities" inform the user agent about the remote fax capabilities. These are described below.

2.3 Parsing a callto URL

A. The "subscriber-id" is extracted. If it begins with a "+", it is an international number. This kind of a number is converted to the user agent's local format (for example, if the agent is a browser component that dials out, the "+" is replaced by the international call prefix, or if the country code matches the country code of user agent's home country, the "+" and the country code are replaced by a domestic call prefix).

A phone number may contain "p" or "w" characters which indicate a pause of 1 second while dialing, or a wait for user input, respectively. User agents SHOULD support this method of dialing. If it is not supported, user agents MUST ignore everything after the "p" or "w" characters. All digits after the first "p" or "w" character MUST be sent to line using DTMF (Dual Tone Multifrequency) in-band signaling, even if dialing is done using direct network signaling (a digital subscriber loop or a mobile phone).

B. If there is no "type-specifier" after the "subscriber-id", the user agent MUST open a normal voice call, and if the true nature of the call type can be defined later (for example, modem/fax autodetection), user agent MAY opt to change to the detected mode.

C. If there is a "type-specifier" after "subscriber-id", the user agent MUST call out using either a suitable terminal equipment (modem, telefax) or request that the network provides such a service. If call type is specified, calls MUST NOT be done in such a way that is incompatible with the recipient entity (such as calling a voice number with a modem, or calling a modem number with a telefax). Type specifiers are "voice" for normal voice call, or a call to a voice mailbox, or a voice messaging service; "fax" for a facsimile (telefax) call; "data" for a data call, with no processing of transferred data implied at either end by default; and data with a "ppp" protocol specifier for a data call where the caller (client) and the callee (server) should begin negotiating a PPP connection after the call has been established. What protocols are negotiated at this phase depends on the user agent's and server's capabilities and configuration. What happens after the call has been established is outside of the scope of this document.

D. If there are any parameters after "call-type", the user agent SHOULD parse the remote capabilities, such as the data bits/parity/ stop bits and remote user name and password and adjust itself appropriately. The capability tokens are listed below. If listed capabilities have been parsed and suggest that it is impossible to create a connection, the connection MUST NOT be created.

Capability	Explanation
V21	ITU-T V.21
V22	ITU-T V.22
V22b	ITU-T V.22bis
V23	ITU-T V.23
V26t	ITU-T V.26ter
V32	ITU-T V.32
V32b	ITU-T V.32bis
V34	ITU-T V.34
V110	ITU-T V.110
V120	ITU-T V.120
X75	ITU-T X.75
B103	Bell 103
B212	Bell 212
G1	ITU-T T.2 (G1) facsimile
G2	ITU-T T.3 (G2) facsimile
G3	ITU-T T.4 (G3) facsimile
G4	ITU-T T.6 (G4) facsimile
Data bits: "8" or "7"	The number of data bits. If not specified, defaults to "8".
Parity: "n", "e", "o", "m", "s"	Parity. None, even, odd, mark or space parity, respectively. If not specified, defaults to "n".
Stop bits: "1" or "2"	The number of stop bits. If not specified, defaults to "1".
Protocol specifier: "ppp" or "raw"	The higher level protocol to be used over the raw data stream. If not

specified, defaults to "raw" which implies that the data stream does not contain any higher level protocol (for example, a terminal connection). "ppp" implies that PPP handshake should be started over the data stream.

2.4 Examples of Use

```
callto:+358-55-1234567
```

This URL instructs the user agent to place a voice call to the specified number in Finland. The hyphens are included to make the number more human-readable: country and area codes have been separated from the subscriber number.

```
callto:+358-55-1234567;type=fax
```

The above URL instructs the user agent to place a fax call to the specified number.

```
callto:+358-55-1234567;type=data;params=v34?7e1;params=v110
```

This URL instructs the user agent to place a data call to the specified number. The user agent may opt to use either a ITU-T V.34 modem (or a slower one, which is compatible with V.34), using settings of 7 data bits, even parity and one stop bit, or an ISDN connection using ITU-T V.110 protocol.

```
callto:+358-55-1234567pp22
```

The above URL instructs the user agent to place a voice call to +358-55-1234567, then wait two seconds and emit two DMTF dialing tones "2" on the line (for example, to choose a particular extension number). (Most Hayes AT compatible modems accept commas "," as the pause characters.)

3. References

[RFC1738] Uniform Resource Locators (URL). December 1994. T. Berners-Lee et al. <URL:ftp://ftp.ds.internic.net/rfc/rfc1738.txt>

[RFC2119] Key words for use in RFCs to Indicate Requirement Levels. March 1997. S. Bradner. <URL:ftp://ftp.ds.internic.net/rfc/rfc2119.txt>

4. Security Considerations

It should be noted that the user agent SHOULD NOT call out without the knowledge of the user because of associated risks, which include

- call costs (including long calls, long distance calls, international calls and prime rate calls)

- wrong numbers inserted on web pages by malicious users
- making the user's phone line unavailable (off-hook) for a malicious purpose
- opening a data call to a remote host, thus possibly opening a back door to the user's computer

The user agent SHOULD have some mechanism that the user can use to filter out unwanted numbers. The user agent SHOULD NOT use rapid redialing of the number if it is busy to avoid the congestion of (signaling) network. Also, user agent SHOULD detect if the number is unavailable or if the call is terminated before the dialing string has been completely processed (for example, the call is terminated while waiting for user input) and not try to call again, unless instructed by the user.

5. Authors' Addresses

Contact person for this specification:

Nokia Mobile Phones
Antti Vaha-Sipila
P. O. Box 68
FIN-33721 Tampere
Finland

Electronic mail: antti.vaha-sipila@nmp.nokia.com

Please include your name and electronic mail address in all communications. If you want to receive the newest version of this specification electronically, send mail to the address above.

This document expires on the 28th of February, 1998, or when a new version is released.

Exhibit H

INTERNET-DRAFT
draft-cordell-sg16-conv-url-00.txt

Cordell
BT
Editor

Dec 16, 1997
Expires: 21 June 1998

Conversational Multimedia URLs

Status of this memo

{Editor's comments are in braces}

Abstract

The evolving technologies for real-time conversation over the Internet require URLs to provide user contact information. As there are many protocols (including some that are not Internet based) that can be used for inter-user conversation, this document describes a two stage transaction process for obtaining a URL that can be used to initiate conversation. The first stage involves retrieving a list of protocol specific URLs in a MIME encoded file. The MIME type enables an appropriate application to be launched which will analyse the presented URLs and select the most appropriate one. The second stage involves interpreting the protocol specific URL and initiating the conversation. The protocol specific URLs are encoded in a URL form so that they can be embedded directly into HTML pages. This allows the first stage to be omitted. The document describes the format of the MIME encoded list of URLs, and the format of a number of protocol specific URLs.

Contents

Revisions Since Last Version

* Removed call-id as a parameter

1. Introduction

Internet technology allows for real-time conversation to take place. It also provides a convenient method of obtaining user location information in the form of URLs. (Note: As used here, the term user can refer to a person, a machine, or any other entity a person or machine may care to have a conversation with.) These can describe Internet conversational protocols, and non-Internet based conversational mechanisms such as POTS. As there are a number of conversational protocols that can be used to contact a user, this document describes a two stage process for initiating conversation, with the first stage being optional. The first stage retrieves a list of protocol specific URLs in a MIME encoded file. This list is analysed and the most appropriate URL is selected. The second stage involves interpreting the protocol specific URL. The protocol specific URLs are in a form that can be directly embedded into HTML pages so that the first stage can be omitted.

The scheme presented here is designed to leverage as much as possible of existing infrastructure. As other technologies become common place (such as CMA [1] and vCard [2]) the mechanisms presented here may lapse.

The remainder of this document describes the format of the MIME file, and the format of a number of protocol specific URLs.

2. The MIME file

The first stage in the contact process is to obtain a list of possible contact mechanisms. To enable a single link to be placed in an HTML page, an indirection method is used wherein a link to a MIME encoded file is made. The MIME type of the file is:

```
APPLICATION/TALKTO
```

and the default extension is:

```
.tlk
```

The MIME file should be retrieved using HTTP. Files that contain time dependent protocol specific URLs should ensure that the files are marked as non-cacheable.

The MIME encoded file consists of ASCII text and lists a number of protocol specific URLs that can be used to contact a remote user. The section below describes a number of protocol specific URLs, but this should not be considered an exhaustive list.

Each protocol specific URL is presented on a separate line with no leading white space. The preferred line break convention is the one used for HTTP (CRLF), but applications must be tolerant to other line break conventions so that files can be readily edited on diverse hosts.

Each protocol specific URL may be followed by some white space and a comment. The comment should be in a form that can be presented to a user as part of a manual selection process. By default the comments are ignored. For example:

```
<protocol_specific_url>    ;My home number
<protocol_specific_url>    ;My bosses number
```

Lines which begin with white-space should be considered as comments and ignored.
The order of the URLs should be such that the most preferred URL is presented first, and the least preferred is presented last. When interpreting the file, if a URL is unsupported, or is not understood, it should be skipped. Endpoints are encouraged to take into account the preference order indicated by the file when selecting a URL, but this is not required. Parsing of the file may continue if a contact attempt fails.

Note that the file does not contain any other information such as the times when specific URLs are valid. This enables a simple file format that does not have to cope with arbitrary search sequences and the complications of time-zones. Therefore, strictly, the file is only valid at the time it is downloaded and the HTTP cache control attributes should be used to control its validity as required.

As with any file downloaded by HTTP, it can be a static file on a server or dynamically generated by an executable. The data for the latter may be uploaded by schemes such as the VoIP's CMA protocol[1].

Validation of who is allowed to obtain various types of location information can be done using WWW-Authentication and cookies. This document provides no additions to these HTTP mechanisms.

Example URLs for downloading the MIME file are:

```
        http://talkto.mycom.com/me.tlk
        http://talkto.mycom.com/cma.exe?me
```

{For consideration:
The above scheme is simple, but not extensible. It may be prudent to define a basic extension scheme to cope with any future problems. The follwing scheme is suggested for consideration.

If the line starts with a "+", then this line contains a parameter that is optional to interpret, i.e. parsing of the file can continue even if the parameter is not understood.
If the line starts with "*", then the line contains a parameter that must be understood. The rest of the file should only be interpreted if the parameter is understood, but earlier lines can be interpreted even if the paramter is not understood. This definition allows simple parser features and complex parser features to co-exist in the same file. e.g. a file might contain:

```
h323:pete@h323.bt.com
*time=17:00-8:30
h323:home@h323.bt.com
```

where time is a paramter to be defined in the future. Parsers that don't understand the time parameter could use the first URL, but not the second}.

3. Protocol Specific URLs

Protocol specific URLs describe contact information for a specific protocol. This section describes a number of these URLs, but this should not be considered an exhaustive list. Other suitable URLs include the IETF's SIP, VoIP's CMA, and Microsoft's CALLTO schemes. Although the main intention of these URLs is to describe conversational protocols, URLs such as CHAT and MAILTO may be appropriate as a last resort. Under certain circumstances RTSP URLs may also be useful.
This section starts with a description of some common elements. These are then used in the protocol specific URLs.

3.1. Common URL elements

This sub-section describes common elements from which the protocol specific URLs are constructed. A number of the elements use definitions from [3].

```
        network = packet-network | switched-network
        packet-network = "ip"  | "tls" | "udp" | "aal5"
                                 ; ip = IP connection without TLS
                                 ; tls = IP connection made over TLS
                                 ; udp = IP connection made over UDP -
                                 ;       this channel may be made reliable
                                 ;       using additional means
                                 ; aal5 = ATM AAL5 call
        switched-network = "pots" | "isdn" | "aal1"
                                 ; pots = GSTN or ISDN speech/audio call
                                 ; isdn = ISDN data call
                                 ; aal1 = ATM AAL1 call
        address = ip-address | phone-address

        ip-address =  hostport  ; hostport defined in [3]

        phone-address = global-phone-number *[ "&" global-phone-number ]
                                 ; global-phone-number defined in [4].
                                 ; H.323 endpoints do not support the
                                 ; wait for tone pause character
        param-list = param | param  param-list
        param = ";" h323-param
```

Telephone numbers in phone-address should always be presented in a full international form, including the "+" sign. It is the responsibility of endpoints and/or gatekeepers to convert these to location specific numbers.

3.2. H.323 URL

{Note: the format of this URL has been structured to have a basic form of h323:pete@h323.bt.com. This is because users are familiar with this format, and it is intuitive what it means. However, this does present problems when e-mail ids which include an @ are included in the URL. One solution is to include the e-mail @ in its escaped form, i.e. %40. Another option is to specify that parsers should be tolerant of duplicate @ signs. Yet another option is to use an alternative character to represent the @ in the basic URL form, i.e. h323:pete/bt.com. This appears less intuitive, and there may be many erroneous URLs generated as the number of /s at the beginning become very significant, such as in h323:/pete.bt.com which should resolve to an IP address only. }

There are two H.323 related URLs. The first form initiates a call directly based on the information in the URL. The second initiates a call based on information that is obtained by first issuing an H.323 LRQ.

For the first form, the scheme is:

```
        h323url = "h323" ":" [ "/" [ network ] "/" ] h323-address
                                         [h323-param-list]
```

and the second form is:

```
        lrqurl = "lrq" "://" ip-address [h323-param-list]
```

where:

```
        h323-address = user-part | address | user-part address
        user-part = user [ ":" type ] "@"
        user = 1*alphanum               ; alphanum defined in [3]
        type = "e164" | "h323id" | "email"
```

The 'network' part of the URL need only be present if the network is not of type IP (i.e. ip is the default network).

If an ip-address is used in the 'address' field, the 'user' and 'type' fields specify the information to be placed in the destinationInfo part of ARQ and destinationAddress part of SETUP. The 'type' field specifies the type of AliasAddress. If the user field starts with a digit, "*" or "#" the default type is "e164", otherwise it is "h323id".

If a 'phone-number' is used in the address field any 'user' and 'type' parts are placed in the remoteExtensionAddress part of SETUP, and the phone number is placed in the destinationInfo part of ARQ and destinationAddress part of SETUP. It is the responsibility of the receiving H.323 over ISDN gateway to transfer the remoteExtensionAddress to the destinationInfo part of ARQ and destinationAddress part of SETUP prior to making the onward call.

To place an aliasAddress containing an @ sign in the 'user' field, the escaped form of the @ sign must be used, i.e. %40.

If the 'address' field is of type ip-address this is placed in the destCallSignalAddress fields of both ARQ and SETUP.

The H.323 URL may have a number of parameters associated with it. If an endpoint does not know how to handle a parameter then it shall ignore the entire URL. At the time of writing the valid parameters are:

```
h323-param = cid-param | token-param | l2-param
cid-param = "cid" "=" UUID                ; UUID is specified in [5]
token-param = "token" "=" "0x" 1* hex
l2-param  = "l2" "=" ( "PPP" | "MPPP" | "SLIP" )
                                          ; Layer 2 format
```

The cid parameter encodes a UUID that should be placed in the conference ID field of the ARQ and SETUP messages. This field may appear a maximum of 1 time in the URL.

If a conference Identifier is specified, then the conferenceGoal should be "join" in the outgoing SETUP message, otherwise it should be "create".

The token field represents a hexadecimal representation of an octet sequnce. 0, 1 or more token parameters may be included in a URL.

The 'l2' parameter allows for different packetisation schemes to be used over switched network connections. If applicable, the default is PPP.

Note that an H.323 URL with a network type of ISDN indicates that H.323 is carried over the ISDN using a layer 2 protocol such as PPP (specified by the 'l2' parameter). It does not mean that the H.323 system should locate an H.320 gateway and use this to communicate over the ISDN. The **H.320 URL should be used to indicate this.**

Example H.323 URLs are as follows:

```
h323:pete@h323.bt.com
        AliasAddress = pete, AliasAddress type = h323id,
        destCallSignalAddress = h323.bt.com.

h323://pete@h323.bt.com
        Same as above

h323:646436@h323.bt.com
        AliasAddress = 646436, AliasAddress type = e164,
        destCallSignalAddress = h323.bt.com.

h323:pete@      -- This form requires a gatekeeper to determine a
        destCallSignalAddress
        AliasAddress = pete, AliasAddress type = h323id,
        destCallSignalAddress = GK supplied.
```

```
h323:pete.bt.com
        destCallSignalAddress = pete.bt.com


h323:/tls/pete%40bt.com:email@bt.com;token=0x5435;token=0xcdfe;cid=f81d4f
bf-7dec-11d0-a765-00a0c91e6bf6
        This call should be setup over a secure TLS channel.
        AliasAddress = pete@bt.com, AliasAddress type = email-ID,
        destCallSignalAddress = bt.com, Two tokens are supplied.
        A conference ID is also specified.

h323:/pots/+1-515-234-5645
        H.323 over PPP over GSTN.  destCallSignalAddress = an address of an
        H.323 over POTS gateway.  This may be gatekeeper provided.
        +1-515-234-5645 is placed in destinationInfo part of ARQ and
        destinationAddress part of SETUP.

lrq:pete@h323.bt.com
        Causes an LRQ to be performed first.
```

3.3. H.324 URL

The format of the H.324 URL is:

```
h324url = "h324" ":" [ "/" [ switched-network ] "/" ] phone-address
```

The default switched-network type is "pots". H.324i is denoted by having a switch-network type of "isdn".
An example URL is:

```
        h324:+1-515-234-5678
```

or:

```
        h324:/isdn/+1-515-234-5679&+1-515-234-5680
```

3.4. H.320 URL

The format of the H.320 URL is:

```
        h320url = "h320" ":" phone-address
```

The network type is always "isdn". An example is:

```
        h320:/isdn/+1-515-234-5679&+1-515-234-5680
```

3.5. POTS URL

The telephone number scheme for a basic voice call is defined in [4].

3.6. T.120 URL

The format of the T.120 URL is:

```
        t120url = "t120" ":" [ "/" [ network ] "/" ] address [t120-param-list]
```

The following parameters are valid: ???????

4. E-mail list

As many groups are interested in conversational URLs including SG16, VoIP, MMUSIC, PINT, TIPHON, URL-REG etc), a separate e-mail list has been set up. You can subscribe to the list by including the word "subscribe" in the message body text of an e-mail sent to the address:

h323-url-request@vocaltec.com

E-mail can be sent to the list at the following address:

h323-url@vocaltec.com

5. Security Considerations

Umm...

6. Acknowledgements

7. References

[1] "Service Interoperability Implementation Agreement," IMTC VoIP Forum

[2] vCard

[3] T. Berners-Lee, L. Masinter, M. McCahill, "Uniform Resource Locators (URL)," RFC1738, December 1994. {Editor's note: A new version of RFC1738 is being produced so this reference will have to be changed.}

[4] A. Vaha-Sipila, "URLs for Telephony," draft-antti-telephony-url-03.txt, 21 Nov 1997

[5] "Call Signalling Protocols and media Stream Packetization for Packet Based Multimedia Communications Systems," ITU-T Recommendation H.225 Version 2, January 1998

Appendix 1 - Complete ABNF

```
all-urls = h323url | lrqurl | h324url | h320url | t120url

h323url = "h323" ":" [ "/" [ network ] "/" ] h323-address [
h323-param-list ]
lrqurl = "lrq" "://" ip-address [ h323-param-list ]
h324url = "h324" ":" [ "/" [ switched-network ] "/" ] phone-address
h320url = "h320" ":" phone-address
t120url = "t120" ":" [ "/" [ network ] "/" ] address [ t120-param-list ]

network = packet-network | switched-network
packet-network = "ip"  | "tls" | "udp" | "aal5"
switched-network = "pots" | "isdn" | "aal1"

h323-address = user-part | address | user-part address
user-part = user [ ":" type ] "@"
user = 1*alphanum                ; alphanum defined in [3]
type = "e164" | "h323id" | "email"

address = ip-address | phone-address
```

```
ip-address =  hostport          ; hostport defined in [3]

phone-address = global-phone-number *[ "&" global-phone-number ]
                                ; global-phone-number defined in [4].
                                ; H.323 endpoints do not support the
                                ; wait for tone pause character

h323-param-list = ";" h323-param | ";" h323-param  h323-param-list

h323-param = cid-param | token-param | l2-param
cid-param = "cid" "=" UUID      ; UUID is specified in [5]
token-param = "token" "=" "0x" 1* hex
l2-param  = "l2" "=" ( "PPP" | "MPPP" | "SLIP" )
                                ; Layer 2 format

t120-param-list = ";" t120-param | ";" t120-param  t120-param-list
t120-param = {???????}
```

Author's Address

Pete Cordell
BT Labs
MLB 4/15
Martlesham Heath
Ipswich
IP5 3RE
UK
e-mail: pete.cordell@bt-sys.bt.co.uk

Exhibit I

Network Working Group
Request for Comments: 1738
Category: Standards Track

T. Berners-Lee
CERN
L. Masinter
Xerox Corporation
M. McCahill
University of Minnesota
Editors

December 1994

Uniform Resource Locators (URL)

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document specifies a Uniform Resource Locator (URL), the syntax and semantics of formalized information for location and access of resources via the Internet.

1. Introduction

This document describes the syntax and semantics for a compact string representation for a resource available via the Internet. These strings are called "Uniform Resource Locators" (URLs).

The specification is derived from concepts introduced by the World-Wide Web global information initiative, whose use of such objects dates from 1990 and is described in "Universal Resource Identifiers in WWW", RFC 1630. The specification of URLs is designed to meet the requirements laid out in "Functional Requirements for Internet Resource Locators" [12].

This document was written by the URI working group of the Internet Engineering Task Force. Comments may be addressed to the editors, or to the URI-WG <uri@bunyip.com>. Discussions of the group are archived at <URL:http://www.acl.lanl.gov/URI/archive/uri-archive.index.html>

2. General URL Syntax

Just as there are many different methods of access to resources, there are several schemes for describing the location of such resources.

The generic syntax for URLs provides a framework for new schemes to be established using protocols other than those defined in this document.

URLs are used to `locate' resources, by providing an abstract identification of the resource location. Having located a resource, a system may perform a variety of operations on the resource, as might be characterized by such words as `access', `update', `replace', `find attributes'. In general, only the `access' method needs to be specified for any URL scheme.

2.1. The main parts of URLs

A full BNF description of the URL syntax is given in Section 5.

In general, URLs are written as follows:

```
<scheme>:<scheme-specific-part>
```

A URL contains the name of the scheme being used (<scheme>) followed by a colon and then a string (the <scheme-specific-part>) whose interpretation depends on the scheme.

Scheme names consist of a sequence of characters. The lower case letters "a"--"z", digits, and the characters plus ("+"), period ("."), and hyphen ("-") are allowed. For resiliency, programs interpreting URLs should treat upper case letters as equivalent to lower case in scheme names (e.g., allow "HTTP" as well as "http").

2.2. URL Character Encoding Issues

URLs are sequences of characters, i.e., letters, digits, and special characters. A URLs may be represented in a variety of ways: e.g., ink on paper, or a sequence of octets in a coded character set. The interpretation of a URL depends only on the identity of the characters used.

In most URL schemes, the sequences of characters in different parts of a URL are used to represent sequences of octets used in Internet protocols. For example, in the ftp scheme, the host name, directory name and file names are such sequences of octets, represented by parts of the URL. Within those parts, an octet may be represented by

the chararacter which has that octet as its code within the US-ASCII [20] coded character set.

In addition, octets may be encoded by a character triplet consisting of the character "%" followed by the two hexadecimal digits (from "0123456789ABCDEF") which forming the hexadecimal value of the octet. (The characters "abcdef" may also be used in hexadecimal encodings.)

Octets must be encoded if they have no corresponding graphic character within the US-ASCII coded character set, if the use of the corresponding character is unsafe, or if the corresponding character is reserved for some other interpretation within the particular URL scheme.

No corresponding graphic US-ASCII:

URLs are written only with the graphic printable characters of the US-ASCII coded character set. The octets 80-FF hexadecimal are not used in US-ASCII, and the octets 00-1F and 7F hexadecimal represent control characters; these must be encoded.

Unsafe:

Characters can be unsafe for a number of reasons. The space character is unsafe because significant spaces may disappear and insignificant spaces may be introduced when URLs are transcribed or typeset or subjected to the treatment of word-processing programs. The characters "<" and ">" are unsafe because they are used as the delimiters around URLs in free text; the quote mark (""") is used to delimit URLs in some systems. The character "#" is unsafe and should always be encoded because it is used in World Wide Web and in other systems to delimit a URL from a fragment/anchor identifier that might follow it. The character "%" is unsafe because it is used for encodings of other characters. Other characters are unsafe because gateways and other transport agents are known to sometimes modify such characters. These characters are "{", "}", "|", "\", "^", "~", "[", "]", and "`".

All unsafe characters must always be encoded within a URL. For example, the character "#" must be encoded within URLs even in systems that do not normally deal with fragment or anchor identifiers, so that if the URL is copied into another system that does use them, it will not be necessary to change the URL encoding.

Reserved:

Many URL schemes reserve certain characters for a special meaning: their appearance in the scheme-specific part of the URL has a designated semantics. If the character corresponding to an octet is reserved in a scheme, the octet must be encoded. The characters ";", "/", "?", ":", "@", "=" and "&" are the characters which may be reserved for special meaning within a scheme. No other characters may be reserved within a scheme.

Usually a URL has the same interpretation when an octet is represented by a character and when it encoded. However, this is not true for reserved characters: encoding a character reserved for a particular scheme may change the semantics of a URL.

Thus, only alphanumerics, the special characters "$-_.+!*'(),", and reserved characters used for their reserved purposes may be used unencoded within a URL.

On the other hand, characters that are not required to be encoded (including alphanumerics) may be encoded within the scheme-specific part of a URL, as long as they are not being used for a reserved purpose.

2.3 Hierarchical schemes and relative links

In some cases, URLs are used to locate resources that contain pointers to other resources. In some cases, those pointers are represented as relative links where the expression of the location of the second resource is in terms of "in the same place as this one except with the following relative path". Relative links are not described in this document. However, the use of relative links depends on the original URL containing a hierarchical structure against which the relative link is based.

Some URL schemes (such as the ftp, http, and file schemes) contain names that can be considered hierarchical; the components of the hierarchy are separated by "/".

3. Specific Schemes

The mapping for some existing standard and experimental protocols is outlined in the BNF syntax definition. Notes on particular protocols follow. The schemes covered are:

ftp	File Transfer protocol
http	Hypertext Transfer Protocol
gopher	The Gopher protocol
mailto	Electronic mail address
news	USENET news

nntp	USENET news using NNTP access
telnet	Reference to interactive sessions
wais	Wide Area Information Servers
file	Host-specific file names
prospero	Prospero Directory Service

Other schemes may be specified by future specifications. Section 4 of this document describes how new schemes may be registered, and lists some scheme names that are under development.

3.1. Common Internet Scheme Syntax

While the syntax for the rest of the URL may vary depending on the particular scheme selected, URL schemes that involve the direct use of an IP-based protocol to a specified host on the Internet use a common syntax for the scheme-specific data:

```
//<user>:<password>@<host>:<port>/<url-path>
```

Some or all of the parts "<user>:<password>@", ":<password>", ":<port>", and "/<url-path>" may be excluded. The scheme specific data start with a double slash "//" to indicate that it complies with the common Internet scheme syntax. The different components obey the following rules:

user
An optional user name. Some schemes (e.g., ftp) allow the specification of a user name.

password
An optional password. If present, it follows the user name separated from it by a colon.

The user name (and password), if present, are followed by a commercial at-sign "@". Within the user and password field, any ":", "@", or "/" must be encoded.

Note that an empty user name or password is different than no user name or password; there is no way to specify a password without specifying a user name. E.g., <URL:ftp://@host.com/> has an empty user name and no password, <URL:ftp://host.com/> has no user name, while <URL:ftp://foo:@host.com/> has a user name of "foo" and an empty password.

host
The fully qualified domain name of a network host, or its IP address as a set of four decimal digit groups separated by ".". Fully qualified domain names take the form as described in Section 3.5 of RFC 1034 [13] and Section 2.1 of RFC 1123 [5]: a sequence of domain labels separated by ".", each domain label starting and ending with an alphanumerical character and possibly also containing "-" characters. The rightmost domain

label will never start with a digit, though, which syntactically distinguishes all domain names from the IP addresses.

port
The port number to connect to. Most schemes designate protocols that have a default port number. Another port number may optionally be supplied, in decimal, separated from the host by a colon. If the port is omitted, the colon is as well.

url-path
The rest of the locator consists of data specific to the scheme, and is known as the "url-path". It supplies the details of how the specified resource can be accessed. Note that the "/" between the host (or port) and the url-path is NOT part of the url-path.

The url-path syntax depends on the scheme being used, as does the manner in which it is interpreted.

3.2. FTP

The FTP URL scheme is used to designate files and directories on Internet hosts accessible using the FTP protocol (RFC959).

A FTP URL follow the syntax described in Section 3.1. If :<port> is omitted, the port defaults to 21.

3.2.1. FTP Name and Password

A user name and password may be supplied; they are used in the ftp "USER" and "PASS" commands after first making the connection to the FTP server. If no user name or password is supplied and one is requested by the FTP server, the conventions for "anonymous" FTP are to be used, as follows:

The user name "anonymous" is supplied.

The password is supplied as the Internet e-mail address of the end user accessing the resource.

If the URL supplies a user name but no password, and the remote server requests a password, the program interpreting the FTP URL should request one from the user.

3.2.2. FTP url-path

The url-path of a FTP URL has the following syntax:

```
<cwd1>/<cwd2>/.../<cwdN>/<name>;type=<typecode>
```

Where <cwd1> through <cwdN> and <name> are (possibly encoded) strings and <typecode> is one of the characters "a", "i", or "d". The part ";type=<typecode>" may be omitted. The <cwdx> and <name> parts may be empty. The whole url-path may be omitted, including the "/" delimiting it from the prefix containing user, password, host, and port.

The url-path is interpreted as a series of FTP commands as follows:

> Each of the <cwd> elements is to be supplied, sequentially, as the argument to a CWD (change working directory) command.
>
> If the typecode is "d", perform a NLST (name list) command with <name> as the argument, and interpret the results as a file directory listing.
>
> Otherwise, perform a TYPE command with <typecode> as the argument, and then access the file whose name is <name> (for example, using the RETR command.)

Within a name or CWD component, the characters "/" and ";" are reserved and must be encoded. The components are decoded prior to their use in the FTP protocol. In particular, if the appropriate FTP sequence to access a particular file requires supplying a string containing a "/" as an argument to a CWD or RETR command, it is

necessary to encode each "/".

For example, the URL <URL:ftp://myname@host.dom/%2Fetc/motd> is interpreted by FTP-ing to "host.dom", logging in as "myname" (prompting for a password if it is asked for), and then executing "CWD /etc" and then "RETR motd". This has a different meaning from <URL:ftp://myname@host.dom/etc/motd> which would "CWD etc" and then "RETR motd"; the initial "CWD" might be executed relative to the default directory for "myname". On the other hand, <URL:ftp://myname@host.dom//etc/motd>, would "CWD " with a null argument, then "CWD etc", and then "RETR motd".

FTP URLs may also be used for other operations; for example, it is possible to update a file on a remote file server, or infer information about it from the directory listings. The mechanism for doing so is not spelled out here.

3.2.3. FTP Typecode is Optional

The entire ;type=<typecode> part of a FTP URL is optional. If it is omitted, the client program interpreting the URL must guess the appropriate mode to use. In general, the data content type of a file can only be guessed from the name, e.g., from the suffix of the name;

the appropriate type code to be used for transfer of the file can then be deduced from the data content of the file.

3.2.4 Hierarchy

For some file systems, the "/" used to denote the hierarchical structure of the URL corresponds to the delimiter used to construct a file name hierarchy, and thus, the filename will look similar to the URL path. This does NOT mean that the URL is a Unix filename.

3.2.5. Optimization

Clients accessing resources via FTP may employ additional heuristics to optimize the interaction. For some FTP servers, for example, it may be reasonable to keep the control connection open while accessing multiple URLs from the same server. However, there is no common hierarchical model to the FTP protocol, so if a directory change command has been given, it is impossible in general to deduce what sequence should be given to navigate to another directory for a second retrieval, if the paths are different. The only reliable algorithm is to disconnect and reestablish the control connection.

3.3. HTTP

The HTTP URL scheme is used to designate Internet resources accessible using HTTP (HyperText Transfer Protocol).

The HTTP protocol is specified elsewhere. This specification only describes the syntax of HTTP URLs.

An HTTP URL takes the form:

```
http://<host>:<port>/<path>?<searchpart>
```

where <host> and <port> are as described in Section 3.1. If :<port> is omitted, the port defaults to 80. No user name or password is allowed. <path> is an HTTP selector, and <searchpart> is a query string. The <path> is optional, as is the <searchpart> and its preceding "?". If neither <path> nor <searchpart> is present, the "/" may also be omitted.

Within the <path> and <searchpart> components, "/", ";", "?" are reserved. The "/" character may be used within HTTP to designate a hierarchical structure.

3.4. GOPHER

The Gopher URL scheme is used to designate Internet resources accessible using the Gopher protocol.

The base Gopher protocol is described in RFC 1436 and supports items and collections of items (directories). The Gopher+ protocol is a set of upward compatible extensions to the base Gopher protocol and is described in [2]. Gopher+ supports associating arbitrary sets of attributes and alternate data representations with Gopher items. Gopher URLs accommodate both Gopher and Gopher+ items and item attributes.

3.4.1. Gopher URL syntax

A Gopher URL takes the form:

```
gopher://<host>:<port>/<gopher-path>
```

where <gopher-path> is one of

```
<gophertype><selector>
<gophertype><selector>%09<search>
<gophertype><selector>%09<search>%09<gopher+_string>
```

If :<port> is omitted, the port defaults to 70. <gophertype> is a single-character field to denote the Gopher type of the resource to which the URL refers. The entire <gopher-path> may also be empty, in which case the delimiting "/" is also optional and the <gophertype> defaults to "1".

<selector> is the Gopher selector string. In the Gopher protocol, Gopher selector strings are a sequence of octets which may contain any octets except 09 hexadecimal (US-ASCII HT or tab) 0A hexadecimal (US-ASCII character LF), and 0D (US-ASCII character CR).

Gopher clients specify which item to retrieve by sending the Gopher selector string to a Gopher server.

Within the <gopher-path>, no characters are reserved.

Note that some Gopher <selector> strings begin with a copy of the <gophertype> character, in which case that character will occur twice consecutively. The Gopher selector string may be an empty string; this is how Gopher clients refer to the top-level directory on a Gopher server.

3.4.2 Specifying URLs for Gopher Search Engines

If the URL refers to a search to be submitted to a Gopher search engine, the selector is followed by an encoded tab (%09) and the search string. To submit a search to a Gopher search engine, the Gopher client sends the <selector> string (after decoding), a tab, and the search string to the Gopher server.

3.4.3 URL syntax for Gopher+ items

URLs for Gopher+ items have a second encoded tab (%09) and a Gopher+ string. Note that in this case, the %09<search> string must be supplied, although the <search> element may be the empty string.

The <gopher+_string> is used to represent information required for retrieval of the Gopher+ item. Gopher+ items may have alternate views, arbitrary sets of attributes, and may have electronic forms associated with them.

To retrieve the data associated with a Gopher+ URL, a client will connect to the server and send the Gopher selector, followed by a tab and the search string (which may be empty), followed by a tab and the Gopher+ commands.

3.4.4 Default Gopher+ data representation

When a Gopher server returns a directory listing to a client, the Gopher+ items are tagged with either a "+" (denoting Gopher+ items) or a "?" (denoting Gopher+ items which have a +ASK form associated with them). A Gopher URL with a Gopher+ string consisting of only a "+" refers to the default view (data representation) of the item while a Gopher+ string containing only a "?" refer to an item with a Gopher electronic form associated with it.

3.4.5 Gopher+ items with electronic forms

Gopher+ items which have a +ASK associated with them (i.e. Gopher+ items tagged with a "?") require the client to fetch the item's +ASK attribute to get the form definition, and then ask the user to fill out the form and return the user's responses along with the selector string to retrieve the item. Gopher+ clients know how to do this but depend on the "?" tag in the Gopher+ item description to know when to handle this case. The "?" is used in the Gopher+ string to be consistent with Gopher+ protocol's use of this symbol.

3.4.6 Gopher+ item attribute collections

To refer to the Gopher+ attributes of an item, the Gopher URL's Gopher+ string consists of "!" or "$". "!" refers to the all of a Gopher+ item's attributes. "$" refers to all the item attributes for all items in a Gopher directory.

3.4.7 Referring to specific Gopher+ attributes

To refer to specific attributes, the URL's gopher+_string is "!<attribute_name>" or "$<attribute_name>". For example, to refer to the attribute containing the abstract of an item, the gopher+_string would be "!+ABSTRACT".

To refer to several attributes, the gopher+_string consists of the attribute names separated by coded spaces. For example, "!+ABSTRACT%20+SMELL" refers to the +ABSTRACT and +SMELL attributes of an item.

3.4.8 URL syntax for Gopher+ alternate views

Gopher+ allows for optional alternate data representations (alternate views) of items. To retrieve a Gopher+ alternate view, a Gopher+ client sends the appropriate view and language identifier (found in the item's +VIEW attribute). To refer to a specific Gopher+ alternate view, the URL's Gopher+ string would be in the form:

```
     +<view_name>%20<language_name>
```

For example, a Gopher+ string of "+application/postscript%20Es_ES" refers to the Spanish language postscript alternate view of a Gopher+ item.

3.4.9 URL syntax for Gopher+ electronic forms

The gopher+_string for a URL that refers to an item referenced by a Gopher+ electronic form (an ASK block) filled out with specific values is a coded version of what the client sends to the server. The gopher+_string is of the form:

```
+%091%0D%0A+-1%0D%0A<ask_item1_value>%0D%0A<ask_item2_value>%0D%0A.%0D%0A
```

To retrieve this item, the Gopher client sends:

```
    <a_gopher_selector><tab>+<tab>1<cr><lf>
    +-1<cr><lf>
    <ask_item1_value><cr><lf>
    <ask_item2_value><cr><lf>
    .<cr><lf>
```

to the Gopher server.

3.5. MAILTO

The mailto URL scheme is used to designate the Internet mailing address of an individual or service. No additional information other than an Internet mailing address is present or implied.

A mailto URL takes the form:

```
     mailto:<rfc822-addr-spec>
```

where <rfc822-addr-spec> is (the encoding of an) addr-spec, as specified in RFC 822 [6]. Within mailto URLs, there are no reserved characters.

Note that the percent sign ("%") is commonly used within RFC 822 addresses and must be encoded.

Unlike many URLs, the mailto scheme does not represent a data object to be accessed directly; there is no sense in which it designates an object. It has a different use than the message/external-body type in MIME.

3.6. NEWS

The news URL scheme is used to refer to either news groups or individual articles of USENET news, as specified in RFC 1036.

A news URL takes one of two forms:

```
news:<newsgroup-name>
news:<message-id>
```

A <newsgroup-name> is a period-delimited hierarchical name, such as "comp.infosystems.www.misc". A <message-id> corresponds to the Message-ID of section 2.1.5 of RFC 1036, without the enclosing "<" and ">"; it takes the form <unique>@<full_domain_name>. A message identifier may be distinguished from a news group name by the presence of the commercial at "@" character. No additional characters are reserved within the components of a news URL.

If <newsgroup-name> is "*" (as in <URL:news:*>), it is used to refer to "all available news groups".

The news URLs are unusual in that by themselves, they do not contain sufficient information to locate a single resource, but, rather, are location-independent.

3.7. NNTP

The nntp URL scheme is an alternative method of referencing news articles, useful for specifying news articles from NNTP servers (RFC 977).

A nntp URL take the form:

```
nntp://<host>:<port>/<newsgroup-name>/<article-number>
```

where <host> and <port> are as described in Section 3.1. If :<port> is omitted, the port defaults to 119.

The <newsgroup-name> is the name of the group, while the <article-number> is the numeric id of the article within that newsgroup.

Note that while nntp: URLs specify a unique location for the article resource, most NNTP servers currently on the Internet today are configured only to allow access from local clients, and thus nntp URLs do not designate globally accessible resources. Thus, the news: form of URL is preferred as a way of identifying news articles.

3.8. TELNET

The Telnet URL scheme is used to designate interactive services that may be accessed by the Telnet protocol.

A telnet URL takes the form:

```
telnet://<user>:<password>@<host>:<port>/
```

as specified in Section 3.1. The final "/" character may be omitted. If :<port> is omitted, the port defaults to 23. The :<password> can be omitted, as well as the whole <user>:<password> part.

This URL does not designate a data object, but rather an interactive service. Remote interactive services vary widely in the means by which they allow remote logins; in practice, the <user> and <password> supplied are advisory only: clients accessing a telnet URL merely advise the user of the suggested username and password.

3.9. WAIS

The WAIS URL scheme is used to designate WAIS databases, searches, or individual documents available from a WAIS database. WAIS is described in [7]. The WAIS protocol is described in RFC 1625 [17]; Although the WAIS protocol is based on Z39.50-1988, the WAIS URL scheme is not intended for use with arbitrary Z39.50 services.

A WAIS URL takes one of the following forms:

```
wais://<host>:<port>/<database>
wais://<host>:<port>/<database>?<search>
wais://<host>:<port>/<database>/<wtype>/<wpath>
```

where <host> and <port> are as described in Section 3.1. If :<port> is omitted, the port defaults to 210. The first form designates a WAIS database that is available for searching. The second form designates a particular search. <database> is the name of the WAIS database being queried.

The third form designates a particular document within a WAIS database to be retrieved. In this form <wtype> is the WAIS designation of the type of the object. Many WAIS implementations require that a client know the "type" of an object prior to retrieval, the type being returned along with the internal object identifier in the search response. The <wtype> is included in the URL in order to allow the client interpreting the URL adequate information to actually retrieve the document.

The <wpath> of a WAIS URL consists of the WAIS document-id, encoded as necessary using the method described in Section 2.2. The WAIS document-id should be treated opaquely; it may only be decomposed by the server that issued it.

3.10 FILES

The file URL scheme is used to designate files accessible on a particular host computer. This scheme, unlike most other URL schemes, does not designate a resource that is universally accessible over the Internet.

A file URL takes the form:

```
file://<host>/<path>
```

where <host> is the fully qualified domain name of the system on which the <path> is accessible, and <path> is a hierarchical directory path of the form <directory>/<directory>/.../<name>.

For example, a VMS file

```
DISK$USER:[MY.NOTES]NOTE123456.TXT
```

might become

```
<URL:file://vms.host.edu/disk$user/my/notes/note12345.txt>
```

As a special case, <host> can be the string "localhost" or the empty string; this is interpreted as `the machine from which the URL is being interpreted'.

The file URL scheme is unusual in that it does not specify an Internet protocol or access method for such files; as such, its utility in network protocols between hosts is limited.

3.11 PROSPERO

The Prospero URL scheme is used to designate resources that are accessed via the Prospero Directory Service. The Prospero protocol is described elsewhere [14].

A prospero URLs takes the form:

```
prospero://<host>:<port>/<hsoname>;<field>=<value>
```

where <host> and <port> are as described in Section 3.1. If :<port> is omitted, the port defaults to 1525. No username or password is allowed.

The <hsoname> is the host-specific object name in the Prospero protocol, suitably encoded. This name is opaque and interpreted by the Prospero server. The semicolon ";" is reserved and may not appear without quoting in the <hsoname>.

Prospero URLs are interpreted by contacting a Prospero directory server on the specified host and port to determine appropriate access methods for a resource, which might themselves be represented as different URLs. External Prospero links are represented as URLs of the underlying access method and are not represented as Prospero URLs.

Note that a slash "/" may appear in the <hsoname> without quoting and no significance may be assumed by the application. Though slashes may indicate hierarchical structure on the server, such structure is not guaranteed. Note that many <hsoname>s begin with a slash, in which case the host or port will be followed by a double slash: the slash from the URL syntax, followed by the initial slash from the <hsoname>. (E.g., <URL:prospero://host.dom//pros/name> designates a <hsoname> of "/pros/name".)

In addition, after the <hsoname>, optional fields and values associated with a Prospero link may be specified as part of the URL. When present, each field/value pair is separated from each other and from the rest of the URL by a ";" (semicolon). The name of the field and its value are separated by a "=" (equal sign). If present, these fields serve to identify the target of the URL. For example, the OBJECT-VERSION field can be specified to identify a specific version of an object.

4. REGISTRATION OF NEW SCHEMES

A new scheme may be introduced by defining a mapping onto a conforming URL syntax, using a new prefix. URLs for experimental schemes may be used by mutual agreement between parties. Scheme names starting with the characters "x-" are reserved for experimental purposes.

The Internet Assigned Numbers Authority (IANA) will maintain a registry of URL schemes. Any submission of a new URL scheme must include a definition of an algorithm for accessing of resources within that scheme and the syntax for representing such a scheme.

URL schemes must have demonstrable utility and operability. One way to provide such a demonstration is via a gateway which provides objects in the new scheme for clients using an existing protocol. If the new scheme does not locate resources that are data objects, the properties of names in the new space must be clearly defined.

New schemes should try to follow the same syntactic conventions of existing schemes, where appropriate. It is likewise recommended that, where a protocol allows for retrieval by URL, that the client software have provision for being configured to use specific gateway locators for indirect access through new naming schemes.

The following scheme have been proposed at various times, but this document does not define their syntax or use at this time. It is suggested that IANA reserve their scheme names for future definition:

afs	Andrew File System global file names.
mid	Message identifiers for electronic mail.
cid	Content identifiers for MIME body parts.
nfs	Network File System (NFS) file names.
tn3270	Interactive 3270 emulation sessions.
mailserver	Access to data available from mail servers.
z39.50	Access to ANSI Z39.50 services.

5. BNF for specific URL schemes

This is a BNF-like description of the Uniform Resource Locator syntax, using the conventions of RFC822, except that "|" is used to designate alternatives, and brackets [] are used around optional or repeated elements. Briefly, literals are quoted with "", optional elements are enclosed in [brackets], and elements may be preceded with <n>* to designate n or more repetitions of the following element; n defaults to 0.

```
; The generic form of a URL is:

genericurl     = scheme ":" schemepart

; Specific predefined schemes are defined here; new schemes
; may be registered with IANA

url            = httpurl | ftpurl | newsurl |
                 nntpurl | telneturl | gopherurl |
                 waisurl | mailtourl | fileurl |
                 prosperourl | otherurl

; new schemes follow the general syntax
otherurl       = genericurl

; the scheme is in lower case; interpreters should use case-ignore
scheme         = 1*[ lowalpha | digit | "+" | "-" | "." ]
```

```
schemepart     = *xchar | ip-schemepart
```

```
; URL schemeparts for ip based protocols:

ip-schemepart  = "//" login [ "/" urlpath ]

login          = [ user [ ":" password ] "@" ] hostport
hostport       = host [ ":" port ]
host           = hostname | hostnumber
hostname       = *[ domainlabel "." ] toplabel
domainlabel    = alphadigit | alphadigit *[ alphadigit | "-" ] alphadigit
toplabel       = alpha | alpha *[ alphadigit | "-" ] alphadigit
alphadigit     = alpha | digit
hostnumber     = digits "." digits "." digits "." digits
port           = digits
user           = *[ uchar | ";" | "?" | "&" | "=" ]
password       = *[ uchar | ";" | "?" | "&" | "=" ]
urlpath        = *xchar    ; depends on protocol see section 3.1

; The predefined schemes:

; FTP (see also RFC959)

ftpurl         = "ftp://" login [ "/" fpath [ ";type=" ftptype ]]
fpath          = fsegment *[ "/" fsegment ]
fsegment       = *[ uchar | "?" | ":" | "@" | "&" | "=" ]
ftptype        = "A" | "I" | "D" | "a" | "i" | "d"

; FILE

fileurl        = "file://" [ host | "localhost" ] "/" fpath

; HTTP

httpurl        = "http://" hostport [ "/" hpath [ "?" search ]]
hpath          = hsegment *[ "/" hsegment ]
hsegment       = *[ uchar | ";" | ":" | "@" | "&" | "=" ]
search         = *[ uchar | ";" | ":" | "@" | "&" | "=" ]

; GOPHER (see also RFC1436)

gopherurl      = "gopher://" hostport [ / [ gtype [ selector
                 [ "%09" search [ "%09" gopher+_string ] ] ] ] ]
gtype          = xchar
selector       = *xchar
gopher+_string = *xchar
```

```
; MAILTO (see also RFC822)

mailtourl      = "mailto:" encoded822addr
encoded822addr = 1*xchar               ; further defined in RFC822

; NEWS (see also RFC1036)
```

```
newsurl        = "news:" grouppart
grouppart      = "*" | group | article
group          = alpha *[ alpha | digit | "-" | "." | "+" | "_" ]
article        = 1*[ uchar | ";" | "/" | "?" | ":" | "&" | "=" ] "@" host

; NNTP (see also RFC977)

nntpurl        = "nntp://" hostport "/" group [ "/" digits ]

; TELNET

telneturl      = "telnet://" login [ "/" ]

; WAIS (see also RFC1625)

waisurl        = waisdatabase | waisindex | waisdoc
waisdatabase   = "wais://" hostport "/" database
waisindex      = "wais://" hostport "/" database "?" search
waisdoc        = "wais://" hostport "/" database "/" wtype "/" wpath
database       = *uchar
wtype          = *uchar
wpath          = *uchar

; PROSPERO

prosperourl    = "prospero://" hostport "/" ppath *[ fieldspec ]
ppath          = psegment *[ "/" psegment ]
psegment       = *[ uchar | "?" | ":" | "@" | "&" | "=" ]
fieldspec      = ";" fieldname "=" fieldvalue
fieldname      = *[ uchar | "?" | ":" | "@" | "&" ]
fieldvalue     = *[ uchar | "?" | ":" | "@" | "&" ]

; Miscellaneous definitions

lowalpha       = "a" | "b" | "c" | "d" | "e" | "f" | "g" | "h" |
                 "i" | "j" | "k" | "l" | "m" | "n" | "o" | "p" |
                 "q" | "r" | "s" | "t" | "u" | "v" | "w" | "x" |
                 "y" | "z"
hialpha        = "A" | "B" | "C" | "D" | "E" | "F" | "G" | "H" | "I" |
                 "J" | "K" | "L" | "M" | "N" | "O" | "P" | "Q" | "R" |
                 "S" | "T" | "U" | "V" | "W" | "X" | "Y" | "Z"
```

```
alpha          = lowalpha | hialpha
digit          = "0" | "1" | "2" | "3" | "4" | "5" | "6" | "7" |
                 "8" | "9"
safe           = "$" | "-" | "_" | "." | "+"
extra          = "!" | "*" | "'" | "(" | ")" | ","
national       = "{" | "}" | "|" | "\" | "^" | "~" | "[" | "]" | "`"
punctuation    = "<" | ">" | "#" | "%" | <">


reserved       = ";" | "/" | "?" | ":" | "@" | "&" | "="
hex            = digit | "A" | "B" | "C" | "D" | "E" | "F" |
                 "a" | "b" | "c" | "d" | "e" | "f"
escape         = "%" hex hex

unreserved     = alpha | digit | safe | extra
uchar          = unreserved | escape
xchar          = unreserved | reserved | escape
digits         = 1*digit
```

6. Security Considerations

The URL scheme does not in itself pose a security threat. Users should beware that there is no general guarantee that a URL which at one time points to a given object continues to do so, and does not even at some later time point to a different object due to the movement of objects on servers.

A URL-related security threat is that it is sometimes possible to construct a URL such that an attempt to perform a harmless idempotent operation such as the retrieval of the object will in fact cause a possibly damaging remote operation to occur. The unsafe URL is typically constructed by specifying a port number other than that reserved for the network protocol in question. The client unwittingly contacts a server which is in fact running a different protocol. The content of the URL contains instructions which when interpreted according to this other protocol cause an unexpected operation. An example has been the use of gopher URLs to cause a rude message to be sent via a SMTP server. Caution should be used when using any URL which specifies a port number other than the default for the protocol, especially when it is a number within the reserved space.

Care should be taken when URLs contain embedded encoded delimiters for a given protocol (for example, CR and LF characters for telnet protocols) that these are not unencoded before transmission. This would violate the protocol but could be used to simulate an extra operation or parameter, again causing an unexpected and possible harmful remote operation to be performed.

The use of URLs containing passwords that should be secret is clearly unwise.

7. Acknowledgements

This paper builds on the basic WWW design (RFC 1630) and much discussion of these issues by many people on the network. The discussion was particularly stimulated by articles by Clifford Lynch, Brewster Kahle [10] and Wengyik Yeong [18]. Contributions from John Curran, Clifford Neuman, Ed Vielmetti and later the IETF URL BOF and URI working group were incorporated.

Most recently, careful readings and comments by Dan Connolly, Ned Freed, Roy Fielding, Guido van Rossum, Michael Dolan, Bert Bos, John Kunze, Olle Jarnefors, Peter Svanberg and many others have helped

refine this RFC.

APPENDIX: Recommendations for URLs in Context

URIs, including URLs, are intended to be transmitted through protocols which provide a context for their interpretation.

In some cases, it will be necessary to distinguish URLs from other possible data structures in a syntactic structure. In this case, is recommended that URLs be preceeded with a prefix consisting of the characters "URL:". For example, this prefix may be used to distinguish URLs from other kinds of URIs.

In addition, there are many occasions when URLs are included in other kinds of text; examples include electronic mail, USENET news messages, or printed on paper. In such cases, it is convenient to have a separate syntactic wrapper that delimits the URL and separates it from the rest of the text, and in particular from punctuation marks that might be mistaken for part of the URL. For this purpose, is recommended that angle brackets ("<" and ">"), along with the prefix "URL:", be used to delimit the boundaries of the URL. This wrapper does not form part of the URL and should not be used in contexts in which delimiters are already specified.

In the case where a fragment/anchor identifier is associated with a URL (following a "#"), the identifier would be placed within the brackets as well.

In some cases, extra whitespace (spaces, linebreaks, tabs, etc.) may need to be added to break long URLs across lines. The whitespace should be ignored when extracting the URL.

No whitespace should be introduced after a hyphen ("-") character. Because some typesetters and printers may (erroneously) introduce a hyphen at the end of line when breaking a line, the interpreter of a URL containing a line break immediately after a hyphen should ignore all unencoded whitespace around the line break, and should be aware that the hyphen may or may not actually be part of the URL.

Examples:

```
Yes, Jim, I found it under <URL:ftp://info.cern.ch/pub/www/doc;
type=d> but you can probably pick it up from <URL:ftp://ds.in
ternic.net/rfc>.  Note the warning in <URL:http://ds.internic.
net/instructions/overview.html#WARNING>.
```

References

[1] Anklesaria, F., McCahill, M., Lindner, P., Johnson, D., Torrey, D., and B. Alberti, "The Internet Gopher Protocol (a distributed document search and retrieval protocol)", RFC 1436, University of Minnesota, March 1993. <URL:ftp://ds.internic.net/rfc/rfc1436.txt;type=a>

[2] Anklesaria, F., Lindner, P., McCahill, M., Torrey, D., Johnson, D., and B. Alberti, "Gopher+: Upward compatible enhancements to the Internet Gopher protocol", University of Minnesota, July 1993. <URL:ftp://boombox.micro.umn.edu/pub/gopher/gopher_protocol /Gopher+/Gopher+.txt>

[3] Berners-Lee, T., "Universal Resource Identifiers in WWW: A Unifying Syntax for the Expression of Names and Addresses of Objects on the Network as used in the World-Wide Web", RFC 1630, CERN, June 1994. <URL:ftp://ds.internic.net/rfc/rfc1630.txt>

[4] Berners-Lee, T., "Hypertext Transfer Protocol (HTTP)", CERN, November 1993.

<URL:ftp://info.cern.ch/pub/www/doc/http-spec.txt.Z>

[5] Braden, R., Editor, "Requirements for Internet Hosts -- Application and Support", STD 3, RFC 1123, IETF, October 1989. <URL:ftp://ds.internic.net/rfc/rfc1123.txt>

[6] Crocker, D. "Standard for the Format of ARPA Internet Text Messages", STD 11, RFC 822, UDEL, April 1982. <URL:ftp://ds.internic.net/rfc/rfc822.txt>

[7] Davis, F., Kahle, B., Morris, H., Salem, J., Shen, T., Wang, R., Sui, J., and M. Grinbaum, "WAIS Interface Protocol Prototype Functional Specification", (v1.5), Thinking Machines Corporation, April 1990. <URL:ftp://quake.think.com/pub/wais/doc/protspec.txt>

[8] Horton, M. and R. Adams, "Standard For Interchange of USENET Messages", RFC 1036, AT&T Bell Laboratories, Center for Seismic Studies, December 1987. <URL:ftp://ds.internic.net/rfc/rfc1036.txt>

[9] Huitema, C., "Naming: Strategies and Techniques", Computer Networks and ISDN Systems 23 (1991) 107-110.

[10] Kahle, B., "Document Identifiers, or International Standard Book Numbers for the Electronic Age", 1991. <URL:ftp://quake.think.com/pub/wais/doc/doc-ids.txt>

[11] Kantor, B. and P. Lapsley, "Network News Transfer Protocol: A Proposed Standard for the Stream-Based Transmission of News", RFC 977, UC San Diego & UC Berkeley, February 1986. <URL:ftp://ds.internic.net/rfc/rfc977.txt>

[12] Kunze, J., "Functional Requirements for Internet Resource Locators", Work in Progress, December 1994. <URL:ftp://ds.internic.net/internet-drafts /draft-ietf-uri-irl-fun-req-02.txt>

[13] Mockapetris, P., "Domain Names - Concepts and Facilities", STD 13, RFC 1034, USC/Information Sciences Institute, November 1987. <URL:ftp://ds.internic.net/rfc/rfc1034.txt>

[14] Neuman, B., and S. Augart, "The Prospero Protocol", USC/Information Sciences Institute, June 1993. <URL:ftp://prospero.isi.edu/pub/prospero/doc /prospero-protocol.PS.Z>

[15] Postel, J. and J. Reynolds, "File Transfer Protocol (FTP)", STD 9, RFC 959, USC/Information Sciences Institute, October 1985.

<URL:ftp://ds.internic.net/rfc/rfc959.txt>

[16] Sollins, K. and L. Masinter, "Functional Requirements for Uniform Resource Names", RFC 1737, MIT/LCS, Xerox Corporation, December 1994.
<URL:ftp://ds.internic.net/rfc/rfc1737.txt>

[17] St. Pierre, M, Fullton, J., Gamiel, K., Goldman, J., Kahle, B., Kunze, J., Morris, H., and F. Schiettecatte, "WAIS over Z39.50-1988", RFC 1625, WAIS, Inc., CNIDR, Thinking Machines Corp., UC Berkeley, FS Consulting, June 1994.
<URL:ftp://ds.internic.net/rfc/rfc1625.txt>

[18] Yeong, W. "Towards Networked Information Retrieval", Technical report 91-06-25-01, Performance Systems International, Inc. <URL:ftp://uu.psi.com/wp/nir.txt>, June 1991.

[19] Yeong, W., "Representing Public Archives in the Directory", Work in Progress, November 1991.

[20] "Coded Character Set -- 7-bit American Standard Code for Information Interchange", ANSI X3.4-1986.

Editors' Addresses

Tim Berners-Lee
World-Wide Web project
CERN,
1211 Geneva 23,
Switzerland

Phone: +41 (22)767 3755
Fax: +41 (22)767 7155
EMail: timbl@info.cern.ch

Larry Masinter
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94034

Phone: (415) 812-4365
Fax: (415) 812-4333
EMail: masinter@parc.xerox.com

Mark McCahill
Computer and Information Services,
University of Minnesota
Room 152 Shepherd Labs
100 Union Street SE

Minneapolis, MN 55455

Phone: (612) 625 1300
EMail: mpm@boombox.micro.umn.edu