Internet Engineering Task Force
INTERNET-DRAFT
draft-ietf-iptel-cpl-02.ps

IPTEL WG
Lennox/Schulzrinne
Columbia University

July 14, 2000
Expires: January, 2001

CPL: A Language for User Control of Internet Telephony Services

Status of this Memo

Copyright Notice

Abstract

The Call Processing Language (CPL) is a language that can be used to describe and control Internet telephony services. It is designed to be implementable on either network servers or user agent servers. It is meant to be simple, extensible, easily edited by graphical clients, and independent of operating system or signalling protocol. It is suitable for running on a server where users may not be allowed to execute arbitrary programs, as it has no variables, loops, or ability to run external programs.

This document is a product of the IP Telephony (IPTEL) working group of the Internet Engineering Task Force. Comments are solicited and should be addressed to the working group's mailing list at iptel@lists.research.bell-labs.com and/or the authors.

Contents

1 Introduction

The Call Processing Language (CPL) is a language that can be used to describe and control Internet telephony services. It is not tied to any particular signalling architecture or protocol; it is anticipated that it will be used with both SIP [1] and H.323 [2].

The CPL is powerful enough to describe a large number of services and features, but it is limited in power so that it can run safely in Internet telephony servers. The intention is to make it impossible for users to do anything more complex (and dangerous) than describing Internet telephony services. The language is not Turing-complete, and provides no way to write loops or recursion.

The CPL is also designed to be easily created and edited by graphical tools. It is based on XML [3], so parsing it is easy and many parsers for it are publicly available. The structure of the language maps closely to its behavior, so an editor can understand any valid script, even ones written by hand. The language is also designed so that a server can easily confirm scripts' validity at the time they are delivered to it, rather that discovering them while a call is being processed.

Implementations of the CPL are expected to take place both in Internet telephony servers and in advanced clients; both can usefully process and direct users' calls. This document primarily addresses the usage in servers. A mechanism will be needed to transport scripts between clients and servers; this document does not describe such a mechanism, but related documents will.

The framework and requirements for the CPL architecture are described in RFC 2824, "Call Processing Language Framework and Requirements." [4].

1.1 Conventions of this document

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in RFC 2119 [5] and indicate requirement levels for compliant CPL implementations.

In examples, non-XML strings such as `-action1-`, `-action2-`, and so forth, are sometimes used. These represent further parts of the script which are not relevant to the example in question.

> Some paragraphs are indented, like this; they give motivations of design choices, or questions for future discussion in the development of the CPL, and are not essential to the specification of the language.

2 Structure of CPL scripts

2.1 High-level structure

A CPL script consists of two types of information: *ancillary information* about the script, and *call processing actions.*

A call processing action is a structured tree that describes the decisions and actions a telephony signalling server performs on a call set-up event. There are two types of call processing actions: *top-level actions* are actions that are triggered by signalling events that arrive at the server. Two top-level action names are defined: incoming, the action performed when a call arrives whose destination is the owner of the script; and outgoing, the action performed when a call arrives whose originator is the owner of the script. *Sub-actions* are actions which can be called from other actions. The CPL forbids sub-actions from being called recursively: see section 9.

> Note: The names "action," "sub-action," and "top-level action" are probably not ideal. Suggestions for better names for these concepts are welcomed.

Ancillary information is information which is necessary for a server to correctly process a script, but which does not directly describe any actions. Currently, no ancillary information is defined, but the section is reserved for future extensions.

2.2 Abstract structure of a call processing action

Abstractly, a call processing action is described by a collection of nodes, which describe actions that can be performed or choices which can be made. A node may have several parameters, which specify the precise behavior of the node; they usually also have outputs, which depend on the result of the condition or action.

For a graphical representation of a CPL action, see figure 1. Nodes and outputs can be thought of informally as boxes and arrows; the CPL is designed so that actions can be conveniently edited graphically using this representation. Nodes are arranged in a tree, starting at a single root node; outputs of nodes are connected to additional nodes. When an action is run, the action or condition described by the top-level node is performed; based on the result of that node, the server follows one of the node's outputs, and that action or condition is performed; this process continues until a node with no specified outputs is reached. Because the graph is acyclic, this will occur after a bounded and predictable number of nodes are visited.

If an output to a node is not specified, it indicates that the CPL server should perform a node- or protocol-specific action. Some nodes have specific default actions associated with them; for others, the default action is implicit in the underlying signalling protocol, or can be configured by the administrator of the server. For further details on this, see section 11.

2.3 Location model

For flexibility, one piece of information necessary for the function of a CPL is not given as node parameters: the set of locations to which a call is to be directed. Instead, this set of locations is stored as an implicit global variable throughout the execution of a processing action (and its sub-actions). This allows locations to be retrieved from external sources, filtered, and so forth, without requiring general language support for such actions (which could harm the simplicity and tractability of understanding the language). The specific actions which add, retrieve, or filter location sets are given in section 6.

For the incoming top-level processing action, the location set is initialized to the empty set. For the outgoing action, it is initialized to the destination address of the call.

2.4 XML structure

Syntactically, CPL scripts are represented by XML documents. XML is thoroughly specified by [3], and implementors of this specification should be familiar with that document, but as a brief overview, XML consists of a hierarchical structure of tags; each tag can have a number of attributes. It is visually and

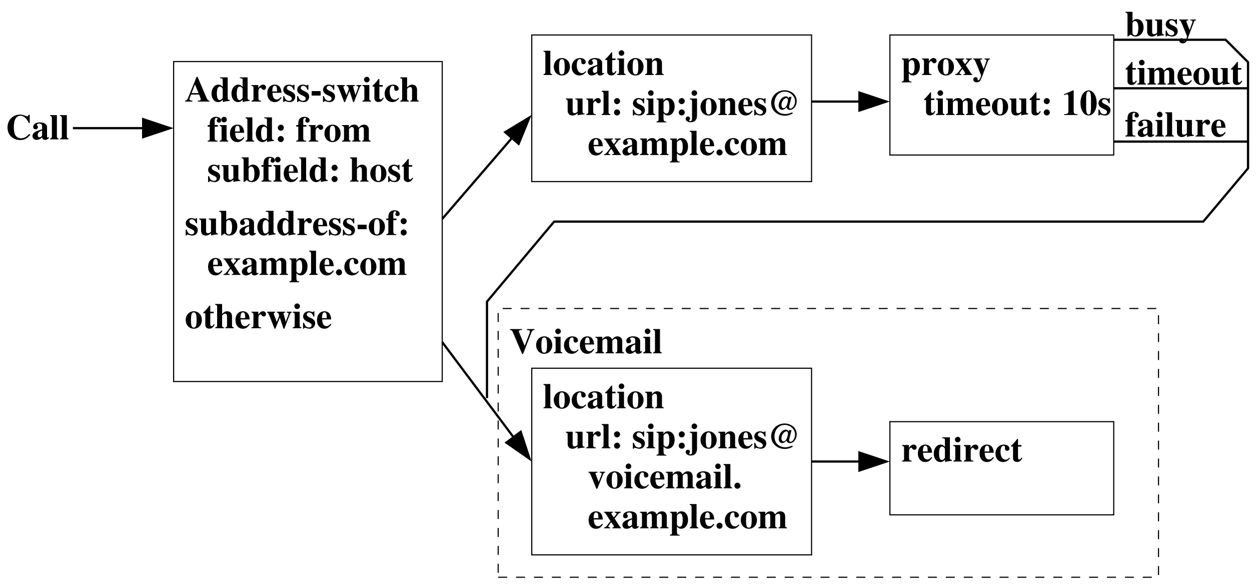

Figure 1: Sample CPL Action: Graphical Version

structurally very similar to HTML [6], as both languages are simplifications of the earlier and larger standard SGML [7].

See figure 2 for the XML document corresponding to the graphical representation of the CPL script in figure 1. Both nodes and outputs in the CPL are represented by XML tags; parameters are represented by XML tag attributes. Typically, node tags contain output tags, and vice-versa (with one exception; see section 2.3).

The connection between the output of a node and another node is represented by enclosing the tag representing the pointed-to node inside the tag for the outer node's output. Convergence (several outputs pointing to a single node) is represented by sub-actions, discussed further in section 9.

The higher-level structure of a CPL script is represented by tags corresponding to each piece of meta-information, sub-actions, and top-level actions, in order. This higher-level information is all enclosed in a special tag `cpl`, the outermost tag of the XML document.

A complete Document Type Declaration for the CPL is provided in Appendix A. The remainder of the main sections of this document describe the semantics of the CPL, while giving its syntax informally. For the formal syntax, please see the appendix.

3 Document information

This section gives meta-information about CPL scripts.

3.1 CPL Document Identifiers for XML

A CPL script list which appears as a top-level XML document is identified with the formal public identifier "-//IETF//DTD RFCxxxx CPL 1.0//EN". If this document is published as an RFC, "xxxx" will be replaced by the RFC number.

```
<?xml version="1.0" ?>
<!DOCTYPE cpl PUBLIC "-//IETF//DTD RFCxxxx CPL 1.0//EN" "cpl.dtd">

<cpl>
  <subaction id="voicemail">
    <location url="sip:jones@voicemail.example.com">
      <redirect />
    </location>
  </subaction>

  <incoming>
    <address-switch field="origin" subfield="host">
      <address subdomain-of="example.com">
        <location url="sip:jones@example.com">
          <proxy timeout="10">
            <busy> <sub ref="voicemail" /> </busy>
            <noanswer> <sub ref="voicemail" /> </noanswer>
            <failure> <sub ref="voicemail" /> </failure>
          </proxy>
        </location>
      </address>
      <otherwise>
        <sub ref="voicemail" />
      </otherwise>
    </address-switch>
  </incoming>
</cpl>
```

Figure 2: Sample CPL Script: XML Version

An CPL embedded as a fragment within another XML document is identified with the XML namespace identifier “http://www.ietf.org/internet-drafts/draft-ietf-iptel-cpl-02.txt”. If this document is published as an RFC, the namespace identifier will be “http://www.rfc-editor.org/rfc/rfcxxxx.txt”, where xxxx is the RFC number.

> Note that the URIs specifying XML namespaces are only globally unique names; they do not have to reference any particular actual object. The URI of a canonical source of this specification meets the requirement of being globally unique, and is also useful to document the format.

3.2 MIME Registration

As an XML type, CPL’s MIME registration conforms with “XML Media Types” [8] as well as RFC 2048 [9].

MIME media type name: application

MIME subtype name: cpl+xml

Mandatory parameters: none

Optional parameters: charset
As for `application/xml` in "XML Media Types."

Encoding considerations: As for `application/xml` in "XML Media Types."

Security considerations: See section 13, and section 10 of "XML Media Types."

Interoperability considerations: Different CPL servers may use incompatible address types. However, all potential interoperability issues should be resolvable at the time a script is uploaded; there should be no interoperability issues which cannot be detetected until runtime.

Published specification: This document.

Applications which use this media type: None publically released at this time, as far as the authors are aware.

Additional information: **Magic number:** None

File extension: `.cpl` or `.xml`

Macintosh file type code: "TEXT"

Person and e-mail address for further information:
Jonathan Lennox <lennox@cs.columbia.edu>
Henning Schulzrinne <hgs@cs.columbia.edu>

Intended usage: COMMON

Author/Change Controller: The IETF.

4 Script structure: overview

As mentioned, a CPL script consists of ancillary information, subactions, and top-level actions. The full syntax of the cpl node is given in figure 3.

Call processing actions, both top-level actions and sub-actions, consist of nodes and outputs. Nodes and outputs are both described by XML tags. There are four categories of CPL nodes: *switches*, which represent choices a CPL script can make; *location modifiers*, which add or remove locations from the location set; *signalling actions*, which cause signalling events in the underlying protocol; and *non-signalling actions,* which take an action but do not effect the underlying protocol.

5 Switches

Switches represent choices a CPL script can make, based on either attributes of the original call request or items independent of the call.

All switches are arranged as a list of conditions that can match a variable. Each condition corresponds to a node output; the output points to the next node to execute if the condition was true. The conditions are tried in the order they are presented in the script; the output corresponding to the first node to match is taken.

Tag:	cpl	
Parameters:	none	
Sub-tags:	ancillary	See section 10
	subaction	See section 9
	outgoing	Top-level actions to take on this user's outgoing calls
	incoming	Top-level actions to take on this user's incoming calls
Output:	outgoing	
Parameters:	none	
Output:	incoming	
Parameters:	none	

Figure 3: Syntax of the top-level cpl tag

There are two special switch outputs that apply to every switch type. The output not-present, which MAY occur anywhere in the list of outputs, is true if the variable the switch was to match was not present in the original call setup request. The output otherwise, which MUST be the last output specified if it is present, matches if no other condition matched.

If no condition matches and no otherwise output was present in the script, the default script action is taken. See section 11 for more information on this.

5.1 Address switches

Address switches allow a CPL script to make decisions based on one of the addresses present in the original call request. They are summarized in figure 4.

Node:	address-switch	
Outputs:	address	Specific addresses to match
Parameters:	field	origin, destination, or original-destination
	subfield	address-type, user, host, port, tel, display, password, or alias-type
Output:	address	
Parameters:	is	exact match
	contains	substring match (for display only)
	subdomain-of	sub-domain match (for host, tel only)

Figure 4: Syntax of the address-switch node

Address switches have two node parameters: field, and subfield. The mandatory field parameter allows the script to specify which address is to be considered for the switch: either the call's origin address (field "origin"), its current destination address (field "destination"), or its original destination (field "original-destination"), the destination the call had before any earlier forwarding was invoked. Servers MAY define additional field values.

The optional subfield specifies what part of the address is to be considered. The possible subfield values are: address-type, user, host, port, tel, and display. Additional subfield values MAY be defined: two additional ones, password and asn1 are defined specifically for SIP and H.323 respectively, in sections 5.1.1 and 5.1.2 below. If no subfield is specified, the "entire" address is matched; the precise meaning of this is defined for each underlying signalling protocol. Servers MAY define additional subfield values.

The subfields are defined as follows:

address-type This indicates the type of the underlying address; i.e., the URI scheme, if the address can be represented by a URI. The types specifically discussed by this document are sip, tel, and h323. The address type is not case-sensitive. It has a value for all defined address types.

user This subfield of the address indicates, for e-mail style addresses, the user part of the address. For telephone number style address, it includes the subscriber number. This subfield is case-sensitive; it may be not present.

host This subfield of the address indicates the Internet host name or IP address corresponding to the address, in host name, IPv4, or IPv6 format. For host names only, subdomain matching is supported with the subdomain-of match operator. It is not case sensitive, and may be not present.

port This subfield indicates the TCP or UDP port number of the address, numerically in decimal format. It is not case sensitive, as it MUST only contain decimal digits. It may be not present; however, for address types with default ports, an absent port matches the default port number.

tel This subfield indicates a telephone subscriber number, if the address contains such a number. It is not case sensitive (the telephone numbers may contain the symbols 'A' 'B' 'C' and 'D'), and might not be present. It may be matched using the subdomain-of match operator. Punctuation and separator characters in telephone numbers are discarded.

display This subfield indicates a "display name" or user-visible name corresponding to an address. It is a Unicode string, and is matched using the case-insensitive algorithm described in section 5.2. The contains operator may be applied to it. It may be not present.

For any completely unknown subfield, the server MAY reject the script at the time it is submitted with an indication of the problem; if a script with an unknown subfield is executed, the server MUST consider the not-present output to be the valid one.

The address output tag may take exactly one of three possible parameters, indicating the kind of matching allowed.

is An output with this match operator is followed if the subfield being matched in the address-switch exactly matches the argument of the operator. It may be used for any subfield, or for the entire address if no subfield was specified.

subdomain-of This match operator applies only for the subfields host and tel. In the former case, it matches if the hostname being matched is a subdomain of the domain given in the argument of the match operator; thus, `subdomain-of="example.com"` would match the hostnames "example.com", "research.example.com", and "zaphod.sales.internal.example.com". IP addresses may be given as arguments to this operator; however, they only match exactly. In the case of the tel subfield, the output matches if the telephone number being matched has a prefix that matches the argument of the match operator; `subdomain-of="1212555"` would match the telephone number "1 212 555 1212."

contains This match operator applies only for the subfield display. The output matches if the display name being matched contains the argument of the match as a substring.

5.1.1 Address switch mapping for SIP

For SIP, the origin address corresponds to the address in the From header; destination corresponds to the Request-URI; and original-destination corresponds to the To header.

The display subfield of an address is the display-name part of the address, if it is present. Because of SIP's syntax, the destination address field will never have a display subfield.

The address-type subfield of an address is the URI scheme of that address. Other address fields depend on that address-type.

For sip URLs, the user, host, and port subfields correspond to the "user," "host," and "port" elements of the URI syntax. The tel subfield is defined to be the "user" part of the URI if and only if the "user=phone" parameter is given to the URI. An additional subfield, password is defined to correspond to the "password" element of the SIP URI; however, use of this field is NOT RECOMMENDED for general security reasons.

For tel URLs, the tel and user subfields are the subscriber name; in the former case, visual separators are stripped. The host and port subfields are both not present.

For h323 URLs, the subfields are set as in section 5.1.2 below.

For other URI schemes, only the address-type subfield is defined by this specification; servers MAY set other pre-defined subfields, or MAY support additional subfields.

If no subfield is specified for addresses in SIP messages, the string matched is the URI part of the address. For "sip" URLs, all parameters are stripped; for other URLs, the URL is used verbatim.

5.1.2 Address switch mapping for H.323

For H.323, the origin address corresponds to the primary alias address in the sourceAddress field of the Setup-UUIE user-user information element, and to the Q.931 information element callingPartyNumber. If both fields are present, which one has priority is a matter of local server policy; the server SHOULD use the same resolution as it would use for routing decisions in this case. Similarly, the destination address corresponds to the primary alias address of the destinationAddress field, and to the Q.931 information element calledPartyNumber.

This discussion is based on H.323 version 4 [10], which is expected to be approved in November 2000.

The original-destination address corresponds to the redirectedNumber Q.931 information element, if it is present; otherwise it is the same as the destination address.

The mapping of H.323 addresses into subfields depends on the type of the alias address. An additional subfield type, alias-type, is defined for H.323 servers, corresponding to the type of the address. Possible values are dialedDigits, h323-ID, url-ID, transportID, email-ID, partyNumber, mobileUIM, and Q.931IE. If future versions of the H.323 specification define additional types of alias addresses, those names MAY also be used.

In versions of H.323 prior to version 4, dialedDigits was known as e164. The new name should be used.

The value of the address-type subfield for H.323 messages is "h323" unless the alias type is url-ID and the URL scheme is something other than h323; in this case the address-type is the URL scheme, as specified above for SIP.

If an alias address of type h323-ID is present anywhere among the sequence of aliases, the first such h323-ID alias address is used for the display subfield of the address. The values of all other subfields depend only on the first alias address in the sequence.

The following mappings are used for H.323 alias types:

dialedDigits, partyNumber, mobileUIM, and Q.931IE: the tel and user subfields are the string of digits, as is the "entire-address" form. The host and port subfields are not present.

url-ID with a "h323" URI: the user, host, and port subfields are set to the corresponding parts of the H.323 URL. The tel subfield is not present. The "entire-address" form corresponds to the entire URI.

url-ID with other URI schemes: the same mapping is used as for SIP, above.

email-ID: the user and host subfields are set to the corresponding parts of the e-mail address. The port and tel subfields are not present. The "entire-address" form corresponds to the entire e-mail address.

transportID: if the TransportAddress is of type "ipAddress," "ipSourceRoute," or "ip6Address," the host subfield is set to the "ip" element of the sequence, translated into the standard IPv4 or IPv6 textual representation, and the port subfield is set to the "port" element of the sequence represented in decimal. The tel and user fields are not present. The "entire-address" form is not defined. The representation and mapping of transport addresses is not defined for non-IP addresses.

5.2 String switches

String switches allow a CPL script to make decisions based on free-form Unicode strings present in a call request. They are summarized in figure 5.

Node:	string-switch	
Outputs:	string	Specific string to match
Parameters:	field	subject, organization, user-agent, language, or display
Output:	string	
Parameters:	is	exact match
	contains	substring match

Figure 5: Syntax of the string-switch node

String switches have one node parameter: field. The mandatory field parameter specifies which string is to be matched.

Currently five fields are defined. Three fields are currently applicable only to SIP, one is currently applicable only to H.323, and one is applicable to both.

The three fields which are applicable only to SIP are: subject, indicating the subject of the call; organization, indicating the originator's organization; and user-agent, indicating the program or device with which the call request was made. All these fields correspond to the contents of the SIP header fields with the same names.

The field applicable only to H.323 is display, which corresponds to the Q.931 information element of the same name.

> This is conventionally used for Caller-ID purposes, so arguably it should be mapped to the display subfield of an address-match with the field originator. However, since a) it is a message-level information element, not an address-level one, and b) the Q.931 specification [11] says only that "[t]he purpose of the Display information element is to supply display information that may be displayed by the user," it seems to be more appropriate to match it as a string instead.

The field appropriate both to SIP and H.323 is language. This field contains a list of RFC 1766 [12] language tags, separated by commas, corresponding to the SIP Accept-Language header and the H.323 language UUIE.

> Note that matching based on contains is likely to be much more useful than matching based on is, for this field.

Strings are matched as case-insensitive Unicode strings, in the following manner. First, strings are canonicalized to the "Compatibility Composition" (KC) form, as specified in Unicode Technical Report 15 [13]. Then, strings are compared using locale-insensitive caseless mapping, as specified in Unicode Technical Report 21 [14].

> Code to perform the first step, in Java and Perl, is available; see the links from Annex E of UTR 15 [13]. The case-insensitive string comparison in the Java standard class libraries already performs the second step; other Unicode-aware libraries should be similar.

The output tags of string matching are named string, and have a mandatory argument, one of is or contains, indicating whole-string match or substring match, respectively.

5.3 Time switches

Time switches allow a CPL script to make decisions based the time and/or date the script is being executed. They are summarized in figure 6.

Time switches are based closely on the specification of recurring intervals of time from the Internet Calendaring and Scheduling Core Object Specification (iCal COS), RFC 2445 [15].

> This allows CPLs to be generated automatically from calendar books. It also allows us to re-use the extensive existing work specifying time intervals.

The time-switch tag takes two optional parameters, tzid and tzurl, both of which are defined in RFC 2445 (sections 4.8.3.1 and 4.8.3.5 respectively). The TZID is the identifying label by which a time zone definition is referenced. If it begins with a forward slash (solidus), it references a to-be-defined global time zone registry; otherwise it is locally-defined at the server. The TZURL gives a network location from which an up-to-date VTIMEZONE definition for the timezone can be retrieved.

If a script is uploaded with a tzid and tzurl which the CPL server does not recognize or cannot resolve, it SHOULD diagnose and reject this at script upload time. If neither tzid nor tzurl are present, all non-UTC times within this time switch should be interpreted as being "floating" times, i.e. that they are specified in the local timezone of the CPL server.

Node:	time-switch	
Outputs:	time	Specific time to match
Parameters:	tzid	RFC 2445 Time Zone Identifier
	tzurl	RFC 2445 Time Zone URL
Output:	time	
Parameters:	dtstart	Start of interval (RFC 2445 DATE-TIME)
	dtend	End of interval (RFC 2445 DATE-TIME)
	duration	Length of interval (RFC 2445 DURATION)
	freq	Frequency of recurrence (one of "secondly", "minutely", "hourly", "daily", "weekly", "monthly", or "yearly")
	interval	How often the recurrence repeats
	until	Bound of recurrence (RFC 2445 DATE-TIME)
	count	Number of occurences of recurrence
	bysecond	List of seconds within a minute
	byminute	List of minutes within an hour
	byhour	List of hours of the day
	byday	List of days of the week
	bymonthday	List of days of the month
	byyearday	List of days of the year
	byweekno	List of weeks of the year
	bymonth	List of months of the year
	wkst	First day of week
	bysetpos	List of values within set of events specified

Figure 6: Syntax of the time-switch node

> Because of daylight-savings-time changes over the course of a year, it is necessary to specify time switches in a given timezone. UTC offsets are not sufficient, or a time-of-day routing rule which held between 9 am and 5 pm in the eastern United States would start holding between 8 am and 4 pm at the end of October.

Authors of CPL servers should be careful to handle correctly the intervals when local time is discontinuous, at the beginning or end of daylight-savings time. Note especially that some times may occur more than once when clocks are set back.

Time nodes specify a list of periods during which their output should be taken. They have two required parameters: dtstart, which specifies the beginning of the first period of the list, and exactly one of dtend or duration, which specify the ending time or the duration of the period, respectively. The dtstart and dtend parameters are formatted as iCal COS DATE-TIME values, as specified in section 4.3.5 of RFC 2445. Because time zones are specified in the top-level time-switch tag, only forms 1 or 2 (floating or UTC times) can be used. The duration parameter is given as an iCal COS DURATION parameter, as specified in section 4.3.6 of RFC 2445.

If no other parameters are specified, a time node indicates only a single period of time. More complicated sets periods intervals are constructed as recurrences. A recurrence is specified by including the freq parameter, which indicates the type of recurrence rule. No parameters other than dtstart, dtend, and duration SHOULD be specified unless freq is present.

The freq parameter takes one of the following values: secondly, to specify repeating periods based on an interval of a second or more; minutely, to specify repeating periods based on an interval of a minute or more; hourly, to specify repeating periods based on an interval of an hour or more; daily, to specify repeating periods based on an interval of a day or more; weekly, to specify repeating periods based on an interval of a week or more; monthly, to specify repeating periods based on an interval of a month or more; and yearly, to specify repeating periods based on an interval of a year or more. These values are not case-sensitive.

The interval parameter contains a positive integer representing how often the recurrence rule repeats. The default value is "1", meaning every second for a secondly rule, or every minute for a minutely rule, every hour for an hourly rule, every day for a daily rule, every week for a weekly rule, every month for a monthly rule and every year for a yearly rule.

The until parameter defines an iCal COS DATE or DATE-TIME value which bounds the recurrence rule in an inclusive manner. If the value specified by until is synchronized with the specified recurrence, this date or date-time becomes the last instance of the recurrence. If specified as a date-time value, then it MUST be specified in an UTC time format. If not present, and the count parameter is also not present, the recurrence is considered to repeat forever.

The count parameter defines the number of occurrences at which to range-bound the recurrence. The dtstart parameter counts as the first occurrence. The until and count parameters MUST NOT occur in the same time output.

The bysecond parameter specifies a comma-separated list of seconds within a minute. Valid values are 0 to 59. The byminute parameter specifies a comma-separated list of minutes within an hour. Valid values are 0 to 59. The byhour parameter specifies a comma-separated list of hours of the day. Valid values are 0 to 23.

The byday parameter specifies a comma-separated list of days of the week. MO indicates Monday; TU indicates Tuesday; WE indicates Wednesday; TH indicates Thursday; FR indicates Friday; SA indicates Saturday; SU indicates Sunday. These values are not case-sensitive.

Each byday value can also be preceded by a positive (+n) or negative (-n) integer. If present, this indicates the nth occurrence of the specific day within the monthly or yearly recurrence. For example, within a monthly rule, +1MO (or simply 1MO) represents the first Monday within the month, whereas -1MO represents the last Monday of the month. If an integer modifier is not present, it means all days of this type within the specified frequency. For example, within a monthly rule, MO represents all Mondays within the month.

The bymonthday parameter specifies a comma-separated list of days of the month. Valid values are 1 to 31 or -31 to -1. For example, -10 represents the tenth to the last day of the month.

The byyearday parameter specifies a comma-separated list of days of the year. Valid values are 1 to 366 or -366 to -1. For example, -1 represents the last day of the year (December 31st) and -306 represents the 306th to the last day of the year (March 1st).

The byweekno parameter specifies a comma-separated list of ordinals specifying weeks of the year. Valid values are 1 to 53 or -53 to -1. This corresponds to weeks according to week numbering as defined in [ISO 8601]. A week is defined as a seven day period, starting on the day of the week defined to be the week start (see wkst). Week number one of the calendar year is the first week which contains at least four (4) days in that calendar year. This parameter is only valid for yearly rules. For example, 3 represents the third week of the year.

Note: Assuming a Monday week start, week 53 can only occur when Thursday is January 1 or if it is a leap year and Wednesday is January 1.

The bymonth parameter specifies a comma-separated list of months of the year. Valid values are 1 to 12.

The wkst parameter specifies the day on which the workweek starts. Valid values are MO, TU, WE, TH, FR, SA and SU. This is significant when a weekly recurrence has an interval greater than 1, and a byday parameter is specified. This is also significant in a yearly recurrence when a byweekno parameter is specified. The default value is MO.

The bysetpos parameter specifies a comma-separated list of values which corresponds to the nth occurrence within the set of events specified by the rule. Valid values are 1 to 366 or -366 to -1. It MUST only be used in conjunction with another byxxx parameter. For example "the last work day of the month" could be represented as:

```
<time -timerange- freq="monthly" byday="MO,TU,WE,TH,FR" bysetpos="-1">
```

Each bysetpos value can include a positive (+n) or negative (-n) integer. If present, this indicates the nth occurrence of the specific occurrence within the set of events specified by the rule.

If byxxx parameter values are found which are beyond the available scope (ie, bymonthday="30" in February), they are simply ignored.

Byxxx parameters modify the recurrence in some manner. Byxxx rule parts for a period of time which is the same or greater than the frequency generally reduce or limit the number of occurrences of the recurrence generated. For example, freq="daily" bymonth="1" reduces the number of recurrence instances from all days (if the bymonth parameter is not present) to all days in January. Byxxx parameters for a period of time less than the frequency generally increase or expand the number of occurrences of the recurrence. For example, freq="yearly" bymonth="1,2" increases the number of days within the yearly recurrence set from 1 (if bymonth parameter is not present) to 2.

If multiple Byxxx parameters are specified, then after evaluating the specified freq and interval parameters, the Byxxx parameters are applied to the current set of evaluated occurrences in the following order: bymonth, byweekno, byyearday, bymonthday, byday, byhour, byminute, bysecond and bysetpos; then count and until are evaluated.

Here is an example of evaluating multiple Byxxx parameters.

```
<time dtstart="19970105T083000" duration="10M"
      freq="yearly" interval="2" bymonth="1" byday="SU" byhour="8,9"
      byminute="30">
```

First, the interval="2" would be applied to freq="YEARLY" to arrive at "every other year." Then, bymonth="1" would be applied to arrive at "every January, every other year." Then, byday="SU" would be applied to arrive at "every Sunday in January, every other year." Then, byhour="8,9" would be applied to arrive at "every Sunday in January at 8 AM and 9 AM, every other year." Then, byminute="30" would be applied to arrive at "every Sunday in January at 8:30 AM and 9:30 AM, every other year." Then the second is derived from dtstart to end up in "every Sunday in January from 8:30:00 AM to 8:40:00 AM, and from and 9:30:00 AM to 9:40:00 AM, every other year." Similarly, if the byminute, byhour, byday, bymonthday or bymonth parameter were missing, the appropriate minute, hour, day or month would have been retrieved from the dtstart parameter.

The iCal COS RDATE, EXRULE and EXDATE recurrence rules are not specifically mapped to components of the time-switch node. Equivalent functionality to the exception rules can be attained by using the ordering of switch rules to exclude times using earlier rules; equivalent functionality to the additional-date RDATE rules can be attained by using sub nodes (see section 9) to link multiple outputs to the same subsequent node.

The not-present output is never true for a time switch. However, it MAY be included, to allow switch processing to be more regular.

5.4 Priority switches

Priority switches allow a CPL script to make decisions based on the priority specified for the original call. They are summarized in figure 7.

Node:	priority-switch	
Outputs:	priority	Specific priority to match
Parameters:	none	
Output:	priority	
Parameters:	less	Match if priority is less than specified
	greater	Match if priority is greater than specified
	equal	Match if priority is equal to specified

Figure 7: Syntax of the priority-switch node

Priority switches take no parameters.

The priority tags take one of the three parameters greater, less, and equal. The values of these tags are the priorities specified in SIP [1]: in decreasing order, emergency, urgent, normal, and non-urgent. These values are matched in a case-insensitive manner. Outputs with the less parameter are taken if the priority of the message is less than the priority given in the argument; and so forth.

If no priority header is specified in a message, the priority is considered to be normal. If an unknown priority is given, the priority is considered to be equivalent to normal for the purposes of greater and less comparisons, but it is compared literally for equal comparisons.

Since every message has a priority, the not-present output is never true for a priority switch. However, it MAY be included, to allow switch processing to be more regular.

6 Location modifiers

The abstract location model of the CPL is described in section 2.3. The behavior of several of the signalling actions (defined in section 7) is dependent on the current location set specified. Location nodes add or remove locations from the location set.

There are three types of location nodes defined. *Explicit locations* add literally-specified locations to the current location set; *location lookups* obtain locations from some outside source; and *location filters* remove locations from the set, based on some specified criteria.

6.1 Explicit location

Explicit location nodes specify a location literally. Their syntax is described in figure 8.

Node:	location	
Outputs:	any node	
Parameters:	url	URL of address to add to location set
	clear	Whether to clear the location set before adding the new value

Figure 8: Syntax of the location node

Explicit location nodes have two node parameters. The mandatory url parameter's value is the URL of the address to add to the location set. Only one address may be specified per location node; multiple locations may be specified by cascading these nodes. The optional clear parameter specifies whether the location set should be cleared before adding the new location to it. Its value can be "yes" or "no", with "no" as the default.

Since the location is specified as a URL, all locations added in this manner are interpreted as url-ID addresses in H.323.

Basic location nodes have only one possible output, since there is no way that they can fail. (If a basic location node specifies a location which isn't supported by the underlying signalling protocol, the script server SHOULD detect this and report it to the user at the time the script is submitted.) Therefore, its XML representation does not have explicit output nodes; the `<location>` tag directly contains another node tag.

6.2 Location lookup

Locations can also be specified up through external means, through the use of location lookups. The syntax of these tags is given in figure 9.

Location lookup nodes have one mandatory parameter, and four optional parameters. The mandatory parameter is source, the source of the lookup. This can either be a URL, or a non-URL value. If the value of source is a URL, it indicates a location which returns the `application/url` media type. The server adds the locations returned by the URL to the location set.

Non-URL sources indicate a source not specified by a URL which the server can query for addresses to add to the location set. The only non-URL source currently defined is registration, which specifies all the locations currently registered with the server, using SIP REGISTER or H.323 RAS messages.

The lookup node also has four optional parameters. The timeout parameter which specifies the time, in seconds, the script is willing to wait for the lookup to be performed. If this is not specified, its default value is 30. The clear parameter specifies whether the location set should be cleared before the new locations are added.

The other two optional parameters affect the interworking of the CPL script with caller preferences and caller capabilities. These are defined in "SIP Caller Preferences and Callee Capabilities" [16]. By default, a CPL server SHOULD invoke caller preferences filtering when performing a lookup action; that is, it should honor any Accept-Location and Reject-Location headers of the original call request. The two parameters use and ignore allow the script to modify how the script applies caller preferences filtering. The use and ignore parameters both take as their arguments comma-separated lists of caller preferences parameters. If use is given, the server applies the caller preferences resolution algorithm only to those

Node:	lookup	
Outputs:	success	Action if lookup was successful
	notfound	Action if lookup found no addresses
	failure	Action if lookup failed
Parameters:	source	Source of the lookup
	timeout	Time to try before giving up on the lookup
	use	Caller preferences fields to use
	ignore	Caller preferences fields to ignore
	clear	Whether to clear the location set before adding the new values
Output:	success	
Parameters:	none	
Output:	notfound	
Parameters:	none	
Output:	failure	
Parameters:	none	

Figure 9: Syntax of the lookup node

preference parameters given in the use parameter, and ignores all others; if the ignore parameter is given, the server ignores the specified parameters, and uses all the others. Only one of use and ignore can be specified. The addr-spec part of the caller preferences is always applied, and the script cannot modify it.

Note: this is very SIP-specific. H.323 has no similar endpoint-capabilities and requested-capabilities mechanism.

Lookup has three outputs: success, notfound, and failure. Notfound is taken if the lookup process succeeded but did not find any locations; failure is taken if the lookup failed for some reason, including that specified timeout was exceeded. If a given output is not present, script execution terminates and the default action is taken.

Clients SHOULD specify the three outputs success, notfound, and failure in that order, so their script complies with the DTD given in Appendix A, but servers MAY accept them in any order.

6.3 Location filtering

A CPL script can also filter addresses out of the address set, through the use of a mechanism very similar to caller preferences: the remove-location node. The syntax of these nodes is defined in figure 10.

A remove-location node has the same effect on the location set as a Reject-Contact header in caller preferences [16]. The value of the location parameter is treated as though it were the addr-spec field of a Reject-Contact header; an absent header is equivalent to an addr-spec of "*" in that specification. If param and value are present, their values are comma-separated lists of caller preferences parameters and corresponding values, respectively, where the nth entry in the name list matches the nth entry in the value list. There MUST be the same number of parameters as values specified. These are treated, for location filtering

Node:	remove-location	
Outputs:	any node	
Parameters:	param	Caller preference parameter to apply
	value	Value of caller preference parameter
	location	Caller preference location to apply

Figure 10: Syntax of the remove-location node

purposes, as though they appeared in the params field of a Reject-Location header, as "; param=value" for each one.

Note: this is also very SIP-specific. H.323 has no similar endpoint-capabilities mechanism.

7 Signalling actions

Signalling action nodes cause signalling events in the underlying signalling protocol. Three signalling actions are defined: "proxy," "redirect," and "reject."

7.1 Proxy

Proxy causes the triggering call to be forwarded on to the currently specified set of locations. The syntax of the proxy node is given in figure 11.

After a proxy action has completed, the CPL server chooses the "best" response to the call attempt, as defined by the signalling protocol or the server's administrative configuration rules.

If the call attempt was successful, or if no output was specified which corresponded to the the best response, CPL execution terminates and the server returns to its default behavior (normally, to forward the best response upstream to the originator). Otherwise, one of the four outputs busy, noanswer, redirection, or failure is taken.

Note: future extension of the CPL to allow in-call or end-of-call actions will require success outputs to be added.

If no locations were present in the set, or if the only locations in the set were locations to which the server cannot proxy a call (for example, "http" URLs), the failure output is taken.

Proxy has three optional parameters. The timeout parameter specifies the time, in seconds, to wait for the call to be completed or rejected; after this time has elapsed, the call attempt is terminated and the noanswer branch is taken. If this parameter is not specified, the default value is 20 seconds if the proxy node has a no-answer output specified; otherwise the server SHOULD allow the call to ring for a reasonably long period of time (to the maximum extent that server policy allows).

Question: is having the default value dependent on script structure too ugly?

The second optional parameter is recurse, which can take two values, yes or no. This specifies whether the server should automatically attempt to place further call attempts to telephony addresses in redirection responses that were returned from the initial server. Note that if the value of recurse is yes, the redirection output to the script is never taken. In this case this output SHOULD NOT be present. The default value of this parameter is yes.

Node:	proxy	
Outputs:	busy	Action if call attempt returned “busy”
	noanswer	Action if call attempt was not answered before timeout
	redirection	Action if call attempt was redirected
	failure	Action if call attempt failed
Parameters:	timeout	Time to try before giving up on the call attempt
	recurse	Whether to recursively look up redirections
	ordering	What order to try the location set in.
Output:	busy	
Parameters:	none	
Output:	noanswer	
Parameters:	none	
Output:	redirection	
Parameters:	none	
Output:	failure	
Parameters:	none	

Figure 11: Syntax of the proxy node

The third optional parameter is ordering. This can have three possible values: parallel, sequential, and first-only. This parameter specifies in what order the locations of the location set should be tried. Parallel asks that they all be tried simultaneously; sequential asks that the one with the highest priority be tried first, the one with the next-highest priority second, and so forth, until one succeeds or the set is exhausted. First-only instructs the server to try only the highest-priority address in the set, and then follow one of the outputs. The priority of locations in a set is determined by server policy, though SIP servers SHOULD honor the q parameter of SIP registrations and the output of the caller preferences lookup algorithm. The default value of this parameter is parallel.

Once a proxy action completes, if control is passed on to other actions, all locations which have been used are cleared from the location set. That is, the location set is emptied if ordering was parallel or sequential; the highest-priority item in the set is removed from the set if ordering was first-only. In the case of a redirection output, the new addresses to which the call was redirected are then added to the location set.

7.2 Redirect

Redirect causes the server to direct the calling party to attempt to place its call to the currently specified set of locations. The syntax of this node is specified in figure 12.

Redirect immediately terminates execution of the CPL script, so this node has no outputs. It has one parameter, permanent, which specifies whether the result returned should indicate that this is a permanent redirection. The value of this parameter is either “yes” or “no” and its default value is “no.”

Node:	redirect	
Outputs:	none	
Parameters:	permanent	Whether the redirection should be considered permanent

Figure 12: Syntax of the redirect node

This corresponds to the SIP "moved permanently and "moved temporarily" (301 and 302) redirections. The contents of the location set are placed into the response's Contact header.

7.3 Reject

Reject nodes cause the server to reject the call attempt. Their syntax is given in figure 13.

Node:	reject	
Outputs:	none	
Parameters:	status	Status code to return
	reason	Reason phrase to return

Figure 13: Syntax of the reject node

This immediately terminates execution of the CPL script, so this node has no outputs.

This node has two arguments: status and reason. The status argument is required, and can take one of the values busy, notfound, reject, and error. Servers which implement SIP MAY also allow a numeric argument corresponding to a SIP status in the 4xx, 5xx, or 6xx range, but scripts SHOULD NOT use them if they wish to be portable.

The reason argument optionally allows the script to specify a reason for the rejection. CPL servers MAY ignore the reason, but ones that implement SIP SHOULD send them in the SIP reason phrase.

8 Other actions

In addition to the signalling actions, the CPL defines several actions which do not affect the telephony signalling protocol.

8.1 Mail

The mail node causes the server to notify a user of the status of the CPL script through electronic mail. Its syntax is given in figure 14.

The mail node takes one argument: a `mailto` URL giving the address, and any additional desired parameters, of the mail to be sent. The server sends the message containing the content to the given url; it SHOULD also include other status information about the original call request and the CPL script at the time of the notification.

Node:	mail	
Outputs:	any node	
Parameters:	url	Mailto url to which the mail should be sent

Figure 14: Syntax of the mail node

Mail nodes have only one output, since failure of e-mail delivery cannot reliably be known in real-time. Therefore, its XML representation does not have explicit output nodes: the `<mail>` tag directly contains another node tag.

Using a full `mailto` URL rather than just an e-mail address allows additional e-mail headers to be specified, such as `<mail url="mailto:jones@example.com?subject=lookup%20failed" />`.

8.2 Log

The Log node causes the server to log information about the call to non-volatile storage. Its syntax is specified in figure 15.

Node:	log	
Outputs:	any node	
Parameters:	name	Name of the log file to use
	comment	Comment to be placed in log file

Figure 15: Syntax of the log node

Log takes two arguments, both optional: name, which specifies the name of the log, and comment, which gives a comment about the information being logged. Servers SHOULD also include other information in the log, such as the time of the logged event, information that triggered the call to be logged, and so forth. Logs are specific to the owner of the script which logged the event. If the name parameter is not given, the event is logged to a standard, server-defined logfile for the script owner. This specification does not define how users may retrieve their logs from the server.

A correctly operating CPL server SHOULD NOT ever allow the log event to fail. As such, log nodes have only one output, and their XML representation does not have explicit output nodes. A CPL `<log>` tag directly contains another node tag.

9 Subactions

XML syntax defines a tree. To allow more general call flow diagrams, and to allow script re-use and modularity, we define subactions.

Two tags are defined for subactions: subaction definitions and subaction references. Their syntax is given in figure 16.

Subactions are defined through subaction tags. These tags are placed in the CPL after any ancillary information (see section 10) but before any top-level tags. They take one argument: id, a token indicating a script-chosen name for the subaction.

Tag:	subaction	
Subtags:	any node	
Parameters:	id	Name of this subaction
Pseudo-node:	sub	
Outputs:	none in XML tree	
Parameters:	ref	Name of subaction to execute

Figure 16: Syntax of subactions and sub pseudo-nodes

Subactions are called from sub tags. The sub tag is a "pseudo-node": it can be used anyplace in a CPL action that a true node could be used. It takes one parameter, ref, the name of the subaction to be called. The sub tag contains no outputs of its own; control instead passes to the subaction.

References to subactions MUST refer to subactions defined before the current action. A sub tag MUST NOT refer to the action which it appears in, or to any action defined later in the CPL script. Top-level actions cannot be called from sub tags, or through any other means. Script servers MUST verify at the time the script is submitted that no sub node refers to any sub-action which is not its proper predecessor.

> Allowing only back-references of subs forbids any sort of recursion. Recursion would introduce the possibility of non-terminating or non-decidable CPL scripts, a possibility our requirements specifically excluded.

Every sub MUST refer to a subaction ID defined within the same CPL script. No external links are permitted.

> If any subsequent version ever defines external linkages, it will use a different tag, perhaps XLink [17]. Ensuring termination in the presence of external links is a difficult problem.

10 Ancillary information

No ancillary information is currently defined for CPL scripts. If ancillary information, not part of any action, is found to be necessary for scripts in the future, it will be added to this section.

The (trivial) definition of the ancillary information section is given in figure 17.

> It may be useful to include timezone definitions inside CPL scripts directly, rather than referencing them externally with tzid and tzurl parameters. If it is, they will be included here.

Tag:	ancillary
Parameters:	none
Subtags:	none

Figure 17: Syntax of the ancillary tag

11 Default actions

When a CPL action reaches an unspecified output, the action it takes is dependent on the current state of script execution. This section gives the actions that should be taken in each case.

no location or signalling actions performed, location set empty: Look up the user's location through whatever mechanism the server would use if no CPL script were in effect. Proxy, redirect, or send a rejection message, using whatever policy the server would use in the absence of a CPL script.

no location or signalling actions performed, location set non-empty: (This can only happen for outgoing calls.) Proxy the call to the addresses in the location set.

location actions performed, no signalling actions: Proxy or redirect the call, whichever is the server's standard policy, to the addresses in the current location set. If the location set is empty, return not-found rejection.

noanswer output of proxy, no timeout given: (This is a special case.) If the noanswer output of a proxy node is unspecified, and no timeout parameter was given to the proxy node, the call should be allowed to ring for the maximum length of time allowed by the server (or the request, if the request specified a timeout).

proxy action previously taken: Return whatever the "best" response is of all accumulated responses to the call to this point, according to the rules of the underlying signalling protocol.

12 Examples

12.1 Example: Call Redirect Unconditional

The script in figure 18 is a simple script which redirects all calls to a single fixed location.

```
<?xml version="1.0" ?>
<!DOCTYPE cpl PUBLIC "-//IETF//DTD RFCxxxx CPL 1.0//EN" "cpl.dtd">

<cpl>
  <incoming>
    <location url="sip:smith@phone.example.com">
     <redirect />
    </location>
  </incoming>
</cpl>
```

Figure 18: Example Script: Call Redirect Unconditional

12.2 Example: Call Forward Busy/No Answer

The script in figure 19 illustrates some more complex behavior. We see an initial proxy attempt to one address, with further actions if that fails. We also see how several outputs take the same action, through the use of subactions.

12.3 Example: Call Screening

The script in figure 20 illustrates address switches and call rejection, in the form of a call screening script. Note also that because the address-switch lacks an otherwise clause, if the initial pattern did not match,

```
<?xml version="1.0" ?>
<!DOCTYPE cpl PUBLIC "-//IETF//DTD RFCxxxx CPL 1.0//EN" "cpl.dtd">

<cpl>
  <subaction id="voicemail">
    <location url="sip:jones@voicemail.example.com" >
      <proxy />
    </location>
  </subaction>

  <incoming>
    <location url="sip:jones@jonespc.example.com">
       <proxy timeout="8">
         <busy>
           <sub ref="voicemail" />
         </busy>
         <noanswer>
           <sub ref="voicemail" />
         </noanswer>
       </proxy>
    </location>
  </incoming>
</cpl>
```

Figure 19: Example Script: Call Forward Busy/No Answer

the script does not define any action. The server therefore proceeds with its default action, which would presumably be to contact the user.

```
<?xml version="1.0" ?>
<!DOCTYPE cpl PUBLIC "-//IETF//DTD RFCxxxx CPL 1.0//EN" "cpl.dtd">

<cpl>
  <incoming>
    <address-switch field="origin" subfield="user">
      <address is="anonymous">
         <reject status="reject"
            reason="I don't accept anonymous calls" />
      </address>
    </address-switch>
  </incoming>
</cpl>
```

Figure 20: Example Script: Call Screening

12.4 Example: Outgoing Call Screening

The script in figure 21 illustrates a script filtering outgoing calls, in the form of a script which prevent 1-900 (premium) calls from being placed.

```
<?xml version="1.0" ?>
<!DOCTYPE cpl PUBLIC "-//IETF//DTD RFCxxxx CPL 1.0//EN" "cpl.dtd">

<cpl>
  <outgoing>
    <address-switch field="original-destination" subfield="tel">
      <address subdomain-of="1900">
        <reject status="reject" reason="Not allowed to make 1-
900 calls." />
      </address>
    </address-switch>
  </outgoing>
</cpl>
```

Figure 21: Example Script: Outgoing Call Screening

12.5 Example: Time-of-day Routing

Figure 22 illustrates time-based conditions and timezones.

12.6 Example: Non-call Actions

Figure 23 illustrates non-call actions; in particular, alerting a user by electronic mail if the lookup server failed. The primary reason for the mail node is to allow this sort of out-of-band notification of error conditions, as the user might otherwise be unaware of any problem.

12.7 Example: A Complex Example

Finally, figure 24 is a complex example which shows the sort of sophisticated behavior which can be achieved by combining CPL nodes. In this case, the user attempts to have his calls reach his desk; if he does not answer within a small amount of time, calls from his boss are forwarded to his celphone, and all other calls are directed to voicemail.

```
<?xml version="1.0" ?>
<!DOCTYPE cpl PUBLIC "-//IETF//DTD RFCxxxx CPL 1.0//EN" "cpl.dtd">

<cpl>
  <incoming>
    <time-switch tzid="America/New-York"
        tzurl="http://zones.stds_r_us.net/tz/America/New-York">
      <time dtstart="20000703T090000" duration="8H"
            freq="weekly" byday="MO,TU,WE,TH,FR">
        <lookup source="registration">
          <success>
            <proxy />
          </success>
        </lookup>
      </time>
      <otherwise>
        <location url="sip:jones@voicemail.example.com">
          <proxy />
        </location>
      </otherwise>
    </time-switch>
  </incoming>
</cpl>
```

Figure 22: Example Script: Time-of-day Routing

```
<?xml version="1.0" ?>
<!DOCTYPE cpl PUBLIC "-//IETF//DTD RFCxxxx CPL 1.0//EN" "cpl.dtd">

<cpl>
  <incoming>
    <lookup source="http://www.example.com/cgi-bin/locate.cgi?user=jones"
            timeout="8">
      <success>
        <proxy />
      </success>
      <failure>
        <mail url="mailto:jones@example.com?subject=lookup%20failed" />
      </failure>
    </lookup>
  </incoming>
</cpl>
```

Figure 23: Example Script: Non-call Actions

```
<?xml version="1.0" ?>
<!DOCTYPE cpl PUBLIC "-//IETF//DTD RFCxxxx CPL 1.0//EN" "cpl.dtd">

<cpl>
  <subaction id="voicemail">
    <location url="sip:jones@voicemail.example.com">
      <redirect />
    </location>
  </subaction>

  <incoming>
    <location url="sip:jones@phone.example.com">
      <proxy timeout="8">
        <busy>
          <sub ref="voicemail" />
        </busy>
        <noanswer>
          <address-switch field="origin">
            <address contains="boss@example.com">
              <location url="tel:+19175551212">
                <proxy />
              </location>
            </address>
            <otherwise>
              <sub ref="voicemail" />
            </otherwise>
          </address-switch>
        </noanswer>
      </proxy>
    </location>
  </incoming>
</cpl>
```

Figure 24: Example Script: A Complex Example

13 Security considerations

The CPL is designed to allow services to be specified in a manner which prevents potentially hostile or mis-configured scripts from launching security attacks, including denial-of-service attacks. Because script runtime is strictly bounded by acyclicity, and because the number of possible script actions are strictly limited, scripts should not be able to inflict damage upon a CPL server.

Because scripts can direct users' telephone calls, the method by which scripts are transmitted from a client to a server MUST be strongly authenticated. Such a method is not specified in this document.

Script servers SHOULD allow server administrators to control the details of what CPL actions are permitted.

14 IANA considerations

This document registers the MIME type `application/cpl+xml`. See section 3.2.

15 Acknowledgments

This document was reviewed and commented upon by IETF IP Telephony Working Group. We specifically acknowledge the following people for their help:

The outgoing call screening script was written by Kenny Hom.

Paul E. Jones contributed greatly to the mappings of H.323 addresses.

The text of the time-switch section was taken (lightly modified) from RFC 2445 [15], by Frank Dawson and Derik Stenerson.

We drew a good deal of inspiration, notably the language's lack of Turing-completeness and the syntax of string matching, from the specification of Sieve [18], a language for user filtering of electronic mail messages.

Thomas F. La Porta and Jonathan Rosenberg had many useful discussions, contributions, and suggestions.

A The XML DTD for CPL

This section includes a full DTD describing the XML syntax of the CPL. Every script submitted to a CPL server SHOULD comply with this DTD. However, CPL servers MAY allow minor variations from it, particularly in the ordering of output branches of nodes. Note that compliance with this DTD is not a sufficient condition for correctness of a CPL script, as many of the conditions described above are not expressible in DTD syntax.

```
<?xml version="1.0" encoding="US-ASCII" ?>

<!--
    Draft DTD for CPL, corresponding to
    draft-ietf-iptel-cpl-01.
-->
```

```
<!-- Nodes. -->
<!-- Switch nodes -->
<!ENTITY % Switch 'address-switch|string-switch|time-switch|
                   priority-switch' >

<!-- Location nodes -->
<!ENTITY % Location 'location|lookup|remove-location' >

<!-- Signalling action nodes -->
<!ENTITY % SignallingAction 'proxy|redirect|reject' >

<!-- Other actions -->
<!ENTITY % OtherAction 'mail|log' >

<!-- Links to subactions -->
<!ENTITY % Sub 'sub' >

<!-- Nodes are one of the above four categories, or a subaction.
     This entity (macro) describes the contents of an output.
     Note that a node can be empty, implying default action. -->
<!ENTITY % Node     '(%Location;|%Switch;|%SignallingAction;|
                     %OtherAction;|%Sub;)?' >


<!-- Switches: choices a CPL script can make. -->

<!-- All switches can have an 'otherwise' output. -->
<!ELEMENT otherwise ( %Node; ) >

<!-- All switches can have a 'not-present' output. -->
<!ELEMENT not-present ( %Node; ) >

<!-- Address-switch makes choices based on addresses. -->
<!ELEMENT address-switch ( (address|not-present)+, otherwise? ) >
<!-- <not-present> must appear at most once -->
<!ATTLIST address-switch
   field         CDATA    #REQUIRED
   subfield      CDATA    #IMPLIED
>

<!ELEMENT address ( %Node; ) >

<!ATTLIST address
   is            CDATA    #IMPLIED
   contains      CDATA    #IMPLIED
```

```
   subdomain-of   CDATA    #IMPLIED
> <!-- Exactly one of these three attributes must appear -->


<!-- String-switch makes choices based on strings. -->

<!ELEMENT string-switch ( (string|not-present)+, otherwise? ) >
<!-- <not-present> must appear at most once -->
<!ATTLIST string-switch
   field         CDATA    #REQUIRED
>

<!ELEMENT string ( %Node; ) >
<!ATTLIST string
   is            CDATA    #IMPLIED
   contains      CDATA    #IMPLIED
>  <!-- Exactly one of these two attributes must appear -->

<!-- Time-switch makes choices based on the current time. -->

<!ELEMENT time-switch ( (time|not-present)+, otherwise? ) >
<!ATTLIST time-switch
   tzid          CDATA    #IMPLIED
   tzurl         CDATA    #IMPLIED
>

<!ELEMENT time ( %Node; ) >

<!-- Exactly one of the two attributes "dtend" and "duration"
     must occur. -->
<!-- The value of "freq" is
        (secondly|minutely|hourly|daily|weekly|monthly|yearly).  It is
        case-insensitive, so it is not given as a DTD switch. -->
<!-- None of the attributes following freq is meaningful unless freq
         appears. -->
<!-- At most one of "until" and "count" may appear. -->
<!-- The value of "wkst" is (MO|TU|WE|TH|FR|SA|SU).  It is
        case-insensitive, so it is not given as a DTD switch. -->
<!ATTLIST time
   dtstart       CDATA  #REQUIRED
   dtend         CDATA  #IMPLIED
   duration      CDATA  #IMPLIED
   freq          CDATA  #IMPLIED
   until         CDATA  #IMPLIED
   count         CDATA  #IMPLIED
```

```
   interval      CDATA  "1"
   bysecond      CDATA  #IMPLIED
   byminute      CDATA  #IMPLIED
   byhour        CDATA  #IMPLIED
   byday         CDATA  #IMPLIED
   bymonthday    CDATA  #IMPLIED
   byyearday     CDATA  #IMPLIED
   byweekno      CDATA  #IMPLIED
   bymonth       CDATA  #IMPLIED
   wkst          CDATA  "MO"
   bysetpos      CDATA  #IMPLIED
>


<!-- Priority-switch makes choices based on message priority. -->

<!ELEMENT priority-switch ( (priority|not-present)+, otherwise? ) >
<!-- <not-present> must appear at most once -->

<!ENTITY % PriorityVal '(emergency|urgent|normal|non-urgent)' >

<!ELEMENT priority ( %Node; ) >

<!-- Exactly one of these three attributes must appear -->
<!ATTLIST priority
   less          %PriorityVal;  #IMPLIED
   greater       %PriorityVal;  #IMPLIED
   equal         CDATA          #IMPLIED
>


<!-- Locations: ways to specify the location a subsequent action
     (proxy, redirect) will attempt to contact. -->

<!ENTITY % Clear  'clear (yes|no) "no"' >

<!ELEMENT location ( %Node; ) >
<!ATTLIST location
   url           CDATA    #REQUIRED
   %Clear;
>

<!ELEMENT lookup ( success,notfound?,failure? ) >
<!ATTLIST lookup
  source         CDATA     #REQUIRED
```

```
  timeout        CDATA     "30"
  use            CDATA     #IMPLIED
  ignore         CDATA     #IMPLIED
  %Clear;
>

<!ELEMENT success  ( %Node; ) >
<!ELEMENT notfound ( %Node; ) >
<!ELEMENT failure ( %Node; ) >

<!ELEMENT remove-location ( %Node; ) >
<!ATTLIST remove-location
   param         CDATA    #IMPLIED
   value         CDATA    #IMPLIED
   location      CDATA    #IMPLIED
>


<!-- Signalling Actions: call-signalling actions the script can
     take. -->

<!ELEMENT proxy ( busy?,noanswer?,redirection?,failure? ) >

<!-- The default value of timeout is "20" if the <noanswer> output
     exists. -->
<!ATTLIST proxy
   timeout       CDATA    #IMPLIED
   recurse       (yes|no) "yes"
   ordering      CDATA    "parallel"
>

<!ELEMENT busy ( %Node; ) >
<!ELEMENT noanswer ( %Node; ) >
<!ELEMENT redirection ( %Node; ) >
<!-- "failure" repeats from lookup, above. -->

<!ELEMENT redirect EMPTY >
<!ATTLIST redirection
   permanent     (yes|no) "no"
>


<!-- Statuses we can return -->

<!ELEMENT reject EMPTY >
```

```
<!-- The value of "status" is (busy|notfound|reject|error), or a SIP
     4xx-6xx status. -->
<!ATTLIST reject
   status        CDATA    #REQUIRED
   reason        CDATA    #IMPLIED
>

<!-- Non-signalling actions: actions that don't affect the call -->

<!ELEMENT mail ( %Node; ) >
<!ATTLIST mail
   url           CDATA    #REQUIRED
>

<!ELEMENT log ( %Node; ) >
<!ATTLIST log
   name          CDATA    #IMPLIED
   comment       CDATA    #IMPLIED
>


<!-- Calls to subactions. -->

<!ELEMENT sub EMPTY >
<!ATTLIST sub
   ref           IDREF    #REQUIRED
>


<!-- Ancillary data -->

<!ENTITY % Ancillary 'ancillary?' >

<!ELEMENT ancillary EMPTY >


<!-- Subactions -->

<!ENTITY % Subactions 'subaction*' >

<!ELEMENT subaction ( %Node; )>
<!ATTLIST subaction
   id            ID       #REQUIRED
>
```

```
<!-- Top-level actions -->

<!ENTITY % TopLevelActions 'outgoing?,incoming?' >

<!ELEMENT outgoing ( %Node; )>

<!ELEMENT incoming ( %Node; )>

<!-- The top-level element of the script. -->

<!ELEMENT cpl  ( %Ancillary;,%Subactions;,%TopLevelActions; ) >
```

B TODO

- Add many more examples, especially for the caller preferences parts of lookup and remove-location, and for time-switch. Other areas that are not currently addressed (or not addressed adequately) are address-switch subfields, string-switch, priority-switch, and log, and a number of parameters.
- Investigate if there is some way that H.323 endpoint characteristics can be usefully mapped to SIP caller preferences and callee capabilities.

C Changes from earlier versions

C.1 Changes from draft -01

The changebars in the Postscript and PDF versions of this document indicate significant changes from this version.

- Completely re-wrote changes to time switches: they are now based on iCal rather than on crontab.
- Timezone references are now defined within time switches rather than in the ancillary section. The ancillary section is now empty, but still defined for future use. To facilitate this, an explicit ancillary tag was added.
- Added XML document type identifiers (the public identifier and the namespace), and MIME registration information.
- Clarified that the not-present output can appear anywhere in a switch.
- Re-wrote H.323 address mappings. Added the alias-type subfield for H.323 addresses.
- Added the language and display string switch fields.
- Clarified why useless not-present outputs can appear in time and priority switches.

- Added the clear parameter to location and lookup nodes. (It had been in the DTD previously, but not in the text.)
- Weakened support for non-validating scripts from SHOULD to MAY, to allow the use of validating XML parsers.
- Added redirection output of proxy nodes.
- Clarified some aspects of how proxy nodes handle the location set.
- Added permanent parameter of redirect nodes.
- Add example script for outgoing call screening (from Kenny Hom)
- Updated example scripts to use the public identifier.
- Add omitted tag to example script for call forward busy/no answer
- Clarified in introduction that this document mainly deals with servers.
- Updated reference to RFC 2824 now that it has been published.
- Added explanatory text to the introduction to types of nodes.
- Numerous minor clarifications and wording changes.
- Fixed copy-and-paste errors, typos.

C.2 Changes from draft -00

- Added high-level structure; script doesn't just start at a first action.
- Added a section giving a high-level explanation of the location model.
- Added informal syntax specifications for each tag so people don't have to try to understand a DTD to figure out the syntax.
- Added subactions, replacing the old link tags. Links were far too reminiscent of gotos for everyone's taste.
- Added ancillary information section, and timezone support.
- Added not-present switch output.
- Added address switches.
- Made case-insensitive string matching locale-independent.
- Added priority switch.
- Deleted "Other switches" section. None seem to be needed.
- Unified url and source parameters of lookup.

- Added caller prefs to lookup.
- Added location filtering.
- Eliminated "clear" parameter of location setting. Instead, proxy "eats" locations it has used.
- Added recurse and ordering parameters to proxy.
- Added default value of timeout for proxy.
- Renamed response to reject.
- Changed notify to mail, and simplified it.
- Simplified log, eliminating its failure output.
- Added description of default actions at various times during script processing.
- Updated examples for these changes.
- Updated DTD to reflect new syntax.

D Authors' Addresses

Jonathan Lennox
Dept. of Computer Science
Columbia University
1214 Amsterdam Avenue, MC 0401
New York, NY 10027
USA
electronic mail: lennox@cs.columbia.edu

Henning Schulzrinne
Dept. of Computer Science
Columbia University
1214 Amsterdam Avenue, MC 0401
New York, NY 10027
USA
electronic mail: schulzrinne@cs.columbia.edu

References

[1] M. Handley, H. Schulzrinne, E. Schooler, and J. Rosenberg, "SIP: session initiation protocol," Request for Comments 2543, Internet Engineering Task Force, Mar. 1999.

[2] International Telecommunication Union, "Packet based multimedia communication systems," Recommendation H.323, Telecommunication Standardization Sector of ITU, Geneva, Switzerland, Feb. 1998.

[3] T. Bray, J. Paoli, and C. M. Sperberg-McQueen, "Extensible markup language (XML) 1.0," W3C Recommendation REC-xml-19980210, World Wide Web Consortium (W3C), Feb. 1998. Available at http://www.w3.org/TR/REC-xml.

[4] J. Lennox and H. Schulzrinne, "Call processing language framework and requirements," Request for Comments 2824, Internet Engineering Task Force, May 2000.

[5] S. Bradner, "Key words for use in RFCs to indicate requirement levels," Request for Comments 2119, Internet Engineering Task Force, Mar. 1997.

[6] D. Raggett, A. Le Hors, and I. Jacobs, "HTML 4.0 specification," W3C Recommendation REC-html40-19980424, World Wide Web Consortium (W3C), Apr. 1998. Available at http://www.w3.org/TR/REC-html40/.

[7] ISO (International Organization for Standardization), "Information processing — text and office systems — standard generalized markup language (SGML)," ISO Standard ISO 8879:1986(E), International Organization for Standardization, Geneva, Switzerland, Oct. 1986.

[8] M. Murata, S. S. Laurent, and D. Kohn, "XML media types," Internet Draft, Internet Engineering Task Force, June 2000. Work in progress.

[9] N. Freed, J. Klensin, and J. Postel, "Multipurpose internet mail extensions (MIME) part four: Registration procedures," Request for Comments 2048, Internet Engineering Task Force, Nov. 1996.

[10] International Telecommunication Union, "Packet based multimedia communication systems," Recommendation H.323 Draft v4, Telecommunication Standardization Sector of ITU, Geneva, Switzerland, July 2000. To be published November 2000.

[11] International Telecommunication Union, "Digital subscriber signalling system no. 1 (dss 1) - isdn user-network interface layer 3 specification for basic call control," Recommendation Q.931, Telecommunication Standardization Sector of ITU, Geneva, Switzerland, Mar. 1993.

[12] H. Alvestrand, "Tags for the identification of languages," Request for Comments 1766, Internet Engineering Task Force, Mar. 1995.

[13] M. Davis and M. Drst, "Unicode normalization forms," Unicode Technical Report 15, Unicode Consortium, Nov. 1999. Revision 18.0. Available at http://www.unicode.org/unicode/reports/tr15/.

[14] M. Davis, "Case mapping," Unicode Technical Report 21, Unicode Consortium, Nov. 1999. Revision 3.0. Available at http://www.unicode.org/unicode/reports/tr21/.

[15] F. Dawson and D. Stenerson, "Internet calendaring and scheduling core object specification (icalendar)," Request for Comments 2445, Internet Engineering Task Force, Nov. 1998.

[16] H. Schulzrinne and J. Rosenberg, "SIP caller preferences and callee capabilities," Internet Draft, Internet Engineering Task Force, Mar. 2000. Work in progress.

[17] S. DeRose, E. Maler, D. Orchard, and B. Trafford, "XML linking language (XLink)," Working Draft WD-xlink-20000221, World Wide Web Consortium (W3C), Feb. 2000. Available at http://www.w3.org/TR/xlink/.

[18] T. Showalter, “Sieve: A mail filtering language,” Internet Draft, Internet Engineering Task Force, May 2000. Work in progress.

Full Copyright Statement