Internet Location Path System (ILPS)
draft-martins-ilps-00.txt

Status of this Document

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

   Internet-Drafts are draft documents valid for a maximum of six
months and may be updated, replaced, or obsoleted by other documents
at any time. It is inappropriate to use Internet-Drafts as reference
material or to cite them other than as ``work in progress.''

   To learn the current status of any Internet-Draft, please check
the ``1id-abstracts.txt'' listing contained in the Internet-Drafts
Shadow Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe),
munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or
ftp.isi.edu (US West Coast).

Abstract

   This document presents a new, uniform method of mapping IP
addresses, service identifiers and service-specific info. It does the
same work that is done by DNS and URLs. I realize this would be
extremally hard to implement, as there is already an immense working
base of DNS, but the idea provides so much benefits that I thought I
did not have the right not to post it (so maybe the URN folks use it
for something).

Introduction

   Currently a good number of Internet users are not experienced
computer users, and many of them aren't even willing to learn some
kinds of things. For this kind of people, the current model of DNS
and URLs is not a clear enought system.

   This document proposes a method that can join the three stages of
finding something in one, reasonable model. As this is designed to
mimic unix and msdos paths, I called it the "location path" system.

Table of Contents

This document is presented in five chapters and two appendices.

1.Description of the system

   The Location Path system presents the user with just a path - no
access scheme (http:// etc). Optionally, there are part-specifiers
(#) and parameters (?). A example location path could look like:

/org/ietf/ftp/1id-abstracts.txt

   It starts like the Domain Name System, just reversed: first the
top-level domain, then the more specific ones.

   This is also an oportunity to fix the problems people are
complaining of in the top-level domains. First, US domains could be
forced in /us/...; then, international bodies like the IETF or UNO
could be on /org/ietf or /org/uno and local US ones on /us/org/...

   Then the com domains can be split... etc... but this is outside
the scope of this document.

   A client application starts querying any top-level ILPS lookup
server about the /org path. The lookup server says it knows where
it is. The response is, in human-readable format, "yes - it's ILPS
on 123.456.789.0" (for example - for very dumb example).

   Then, as the reply says "ILPS", the client, using the same
connection, asks about /org/ietf. The server may know where this is,
or not. if it does, it will issue a similar response, probably giving
the IP address of the IETF's private lookup server. Then the client
will ask of /org/ietf/ftp, which the top-level server probably won't
know.

   On the first negative message, the client ends the connection and

connects to the last positive message - in this case the IETF lookup
server - using the specified protocol - in this case, ILPS lookup.

   Then it will ask for /ftp; the reply will be like "yes - it's ftp
on 132.151.19.0". As this response, while positive, specifies a
protocol that's not ILPS lookup, the client will close connection,
open a ftp connection to the specified IP address and consider all
remainder of the path to be the ftp path.

   Some other examples to make it clearer:

/us/serv/webcom/www/lalo/ox.html
   (means current http://www.webcom.com/lalo)

/fi/net/funet/irc#linux

/br/edu/ufsm/moo
   (telnet://moo.ufsm.br:7777 - note the port is reported by the
    lookup server, making things clearer and easier)

/us/serv/webcom/pop

etc...

2.The lookup protocol

   Location paths are resolved trough queries using a special
protocol. This protocol is designed to be very simple - indeed I
can't think of anything simplier.

   The connection is started by the client, trough connecting to a
well-known-port on the server. Of course the server's IP address has
to be specified.

   Then the server issues a welcome message, informing the client it
is ready for the lookup. Something like

Location Paths Lookup server at Somewhere ready.

   The only input that is accepted in this protocol is a parcial path
consisting of parts and slashes. Please note input doesn't need to be
preceded by a slash - you can either ask for "org" or "/org". End of
input is signaled by a linefeed character.

   Parts consist on sequences of letters, digits, underscores, and
dots (maybe other characters are acceptable; maybe even spaces. This
is up to the IETF). They're separated by slashes ("/").

   The suggested behaviour is to first ask only about the first part,
then the first two ones, etc, until the server issues a negative
reply. But this behaviour is not mandatory - if a programmer thinks
of something better, this protocol doesn't keep her from doing that.

Server responses are given in either three lines, if it is a
positive one, or one line, if it is a negative.

   The negative responses are just: "unknown". This can't be mistaken
for a positive reply, because in a positive response the first line
would consist only of numbers and dots.

   A positive response consists in three lines: the first one has the
IP address that corresponds to the queried path, in ascii format; the
second line has the protocol that should be used for the connection
(eg http or ftp or pop3). Finally, the third line informs the client
to which port number it is supposed to connect. If the default, well-
known port is to be used, the third line may optionally be blank.

   There's no command to end the session. The client may simply close
the IP connection.

   Example session: (The client input is marked with "%", and the
server's with "#", but please note there's no real "%" or "#" in the
protocol)

--------- Session example below ---------
#Welcome - ILPS lookup server at Somewhere Internet Provider ready
%org
#123.456.789.0
#ILPS-lookup
#
%org/ietf
#132.151.19.0
#ILPS-lookup
#
%org/ietf/ftp
#unknown
--------- The client closes connection and goes to 132.151.19.0 ----
#ILPS lookup server at the Internet Enginering Task Force ready
%org/ietf/ftp
#132.151.19.0
#ftp
#
--------- The client closes connection. Other example: ---------
#Welcome - ILPS lookup server at Universidade de Santa Maria ready
%br/edu/ufsm/moo
#132.1.30.2
#telnet
#7777
--------- End ---------

3.Proposed new top-level standard

   This proposal is not really part of the Location Path system, but
is included in this document for convenience.

   Companies, organisations and people in the .com domain are
complaining, recently, that there are not enought new domains
available. They propose either splitting or duplication of the .com
domain, with domains like "corp" or "bizz".

   This document proposes, first, that all domains belonging to the
United States be kept in the /us path. This is not for a personal
reason or even for the sake of organization. The reason for this is
that he top-level domains, outside /us, are kept for organisations
that are really international, like the IETF, the ONU or the world-
wide-corporative site of a multinacional corporation (and regional
sites would be in /us etc).

   The non-country top-level domains should be:

edu, gov, mil - for the same purpose these are used today;

net - almost the same as it is today, but organisations that are
      today on "org" domains and that have the Internet as its sole
      field of action like the IETF or w3c should be on "net".

serv (or ser) - this domain is to "com" like "net" is to "org"; it
      is for for-profit-companies that have the Internet as its sole
      OR MAIN field of action, like service providers, search engines,
      web-space providers etc.

info (or inf) - for companies and organisations on the field of
      public communications - radios, TVs, newspapers, magazines etc.

per - for individuals. Currently it is very rare to see domain names
      owned by individuals, but these already exist, and will become
      very common with technologies like cable modems. These shouldn't
      take up space in the com path. Of course there wouldn't be a
      top-level /per path; is there any person that is really
      international (the person herself, not trough an organisation)?

com - for any for-profit organisation that doesn't fit on the serv or
      info categories.

   This proposal was well thought of, so it should be considered.
Just the distinction of serv/inf/com is probably enought to solve
today's problems (as almost half com domains would go to serv).

   Also, it is recommended that all countries follow the same
standard.

4.Security considerations

    This document doesn't address any security issue.

Appendix A.References

    This document doesn't have any reference, as it addresses a
question that was never addressed before. Maybe the RFCs and STDs on
DNS and top-level domains should be considered references?

Appendix B.Author's Address

Fernando Mauro Martins - aka Lalo

postal: Rua Mario de Andrade, 912 - cep 14030-340
        Ribeirao Preto - SP - Brazil

phone: +55-11-263 7358 (home and operating base)

now how to REALLY contact me:

http://webcom.com/lalo | mailto:lalo@webcom.com