

Network Working Group
Internet-Draft
Obsoletes <[draft-wessels-icp-v2-01.txt](#)>
Expires: November 22, 1997

D. Wessels
National Laboratory for Applied
Network Research/UCSD
22 April 1997

Internet Cache Protocol (ICP), version 2
<[draft-wessels-icp-v2-02.txt](#)>

Status of this Memo

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

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is 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 ``[l-id-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 draft document describes the Internet Cache Protocol (ICP) as currently implemented in a couple of World-Wide Web proxy cache packages. A companion document (RFCXXXX, <[draft-wessels-icp-v2-appl-00.txt](#)>) describes the application of ICP to Web caches. ICP was initially developed by Peter Danzig, et. al. at the University of Southern California. It evolved as an important part of hierarchical caching on the Harvest research project.

1. Introduction

Internet-Draft

22 April 1997

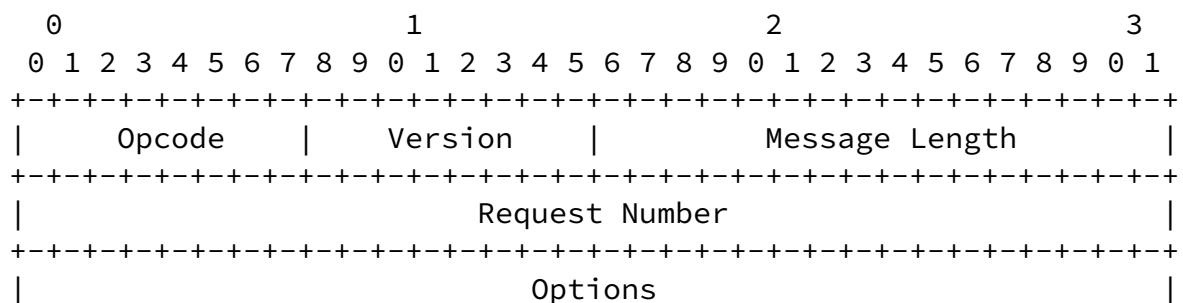
ICP is a message format used for communicating between Web caches. Although Web caches use HTTP[1] for the transfer of object data, caches benefit from a simpler, lighter communication protocol. ICP is primarily used in a cache mesh to locate specific Web objects in neighbor caches. One cache will send an ICP query to its neighbors. The neighbors will send back ICP replies indicating a ``HIT'' or a ``MISS.''

In current practice, ICP is implemented on top of UDP, but there is no requirement that it be limited to UDP. We feel that ICP over UDP offers features important to Web caching applications. An ICP query/reply exchange needs to occur quickly, typically within a second or two. A cache cannot wait longer than that before beginning to retrieve an object. Failure to receive a reply message most likely means the network path is either congested or broken. In either case we would not want to select that neighbor. As an indication of immediate network conditions between neighbor caches, ICP over a lightweight protocol such as UDP is better than one with the overhead of TCP.

In addition to its use as an object location protocol, ICP messages can be used for cache selection. Failure to receive a reply from a cache may indicate a network or system failure. The ICP reply may include information that could assist selection of the most appropriate source from which to retrieve an object.

ICP Message Format

The ICP message format consists of a 20-octet fixed header plus a variable sized payload:



return RTT measurements. The ICP_FLAG_SRC_RTT option is further described below.

Sender Host Address

The IPv4 address of the host sending the ICP message. This field should probably not be trusted over what is provided by `getpeername()`, `accept()`, and `recvfrom()`. There is some ambiguity over the original purpose of this field. In practice it is not used.

Payload

The contents of the Payload field vary depending on the Opcode, but most often it contains a null-terminated URL string.

[2.](#) ICP Opcodes

The following table shows currently defined ICP opcodes:

Value	Name
-----	-----
0	ICP_OP_INVALID
1	ICP_OP_QUERY
2	ICP_OP_HIT
3	ICP_OP_MISS
4	ICP_OP_ERR
5-9	UNUSED
10	ICP_OP_SECHO
11	ICP_OP_DECHO
12-20	UNUSED
21	ICP_OP_MISS_NOFETCH
22	ICP_OP_DENIED
23	ICP_OP_HIT_OBJ

ICP_OP_INVALID

A place holder to detect zero-filled or malformed messages. A cache must never intentionally send an ICP_OP_INVALID message. ICP_OP_ERR should be used instead.

ICP_OP_QUERY

A query message. NOTE this opcode has a different payload format than most of the others. First is the requester's IPv4 address, followed by a URL. The Requester Host Address is not that of the

NOTE: one problem with this approach is that while a system's echo port may be functioning perfectly, the cache software may not be running at all.

One of the following six ICP opcodes are sent in response to an ICP_OP_QUERY message. Unless otherwise noted, the payload must be the null-terminated URL string. Both the URL string and the Request Number field must be exactly the same as from the ICP_OP_QUERY message.

ICP_OP_HIT

An ICP_OP_HIT response indicates that the requested URL exists in this cache and that the requester is allowed to retrieve it.

ICP_OP_MISS

An ICP_OP_MISS response indicates that the requested URL does not exist in this cache. The querying cache may still choose to fetch the URL from the replying cache.

ICP_OP_ERR


An ICP_OP_ERR response indicates some kind of error in parsing or handling the query message (e.g. invalid URL).

ICP_OP_MISS_NOFETCH

An ICP_OP_MISS_NOFETCH response indicates that this cache is up, but is in a state where it does not want to handle cache misses. An example of such a state is during a startup phase where a cache might be rebuilding its object store. A cache in such a mode may wish to return ICP_OP_HIT for cache hits, but not ICP_OP_MISS for misses. ICP_OP_MISS_NOFETCH essentially means ``I am up and running, but please don't fetch this URL from me now.''

Note, ICP_OP_MISS_NOFETCH has a different meaning than ICP_OP_MISS. The ICP_OP_MISS reply is an invitation to fetch the URL from the replying cache (if their relationship allows it), but ICP_OP_MISS_NOFETCH is a request to NOT fetch the URL from the replying cache.

ICP_OP_DENIED



The receiving application must check to make sure it actually receives Object Size octets of data. If it does not, then it should treat the ICP_OP_HIT_OBJ reply as though it were a normal ICP_OP_HIT.

NOTE: the Object Size field does not necessarily begin on a 32-bit boundary as shown in the diagram above. It begins immediately following the NULL byte of the URL string.

UNRECOGNIZED OPCODES

ICP messages with unrecognized or unused opcodes should be ignored, i.e. no reply generated. The application may choose to note the anomalous behaviour in a log file.

3. ICP Option Flags

0x80000000 ICP_FLAG_HIT_OBJ

This flag is set in an ICP_OP_QUERY message indicating that it is okay to respond with an ICP_OP_HIT_OBJ message if the object data will fit in the reply.

0x40000000 ICP_FLAG_SRC_RTT

This flag is set in an ICP_OP_QUERY message indicating that the requester would like the ICP reply to include the responder's measured RTT to the origin server.

Upon receipt of an ICP_OP_QUERY with ICP_FLAG_SRC_RTT bit set, a cache should check an internal database of RTT measurements. If available, the RTT value MUST be expressed as a 16-bit integer, in units of milliseconds. If unavailable, the responder may either set the RTT value to zero, or clear the ICP_FLAG_SRC_RTT bit in the ICP reply. The ICP reply MUST not be delayed while waiting for the RTT measurement to occur.

This flag is set in an ICP reply message (ICP_OP_HIT, ICP_OP_MISS, ICP_OP_MISS_NOFETCH, or ICP_OP_HIT_OBJ) to indicate that the low 16-bits of the Option Data field contain the measured RTT to the host given in the requested URL. If ICP_FLAG_SRC_RTT is clear in the query then it MUST also be clear in the reply. If ICP_FLAG_SRC_RTT is set in the query, then it may or may not be set in the reply.

[4.](#) Security Considerations

The security issues relating to ICP are discussed in the companion document, RFCXXXX (<[draft-wessels-icp-v2-appl-00.txt](#)>).

[5.](#) References

[1] Fielding, R., et. al, "Hypertext Transfer Protocol -- HTTP/1.1", [RFC 2068](#), UC Irvine, January 1997.

[2] Berners-Lee, T., Masinter, L., and M. McCahill, "Uniform Resource Locators (URL)", [RFC 1738](#), CERN, Xerox PARC, University of Minnesota, December 1994.

[3] Bowman M., Danzig P., Hardy D., Manber U., Schwartz M., and Wessels D., "The Harvest Information Discovery and Access System", Internet Research Task Force - Resource Discovery, <http://harvest.transarc.com/>.

[4] Wessels D., Claffy K., "ICP and the Squid Web Cache", National

Laboratory for Applied Network Research, <http://www.nlanr.net/~wessels/Papers/icp-squid.ps.gz>

6. Acknowledgments

Paul A Vixie <paul@vix.com>

7. Author's Address:

Duane Wessels
National Laboratory for Applied Network Research
wessels@nlanr.net

K Claffy
National Laboratory for Applied Network Research
kc@nlanr.net

