Attachment 1

Demarcation of headers and content or payload in the Internet, and so-called application headers
David P. Reed
August 20, 2009

As I have noted in other comments to the FCC, CRTC, and the US House of Representatives, the key architectural feature of the Internet and its protocols is the separation of function between the IP transport network and the applications and services that use that network. I've likened this to an envelope or package, carried by the postal service or a network of package carrier services (such as Fedex or UPS). Since that analogy is strong, but may not be clear to a non-technical layman, let me outline the precise correspondence here.

I will use terminology used in my copy of one of the best references to all of the Internet protocols and their implementation, and I include two figures from that book:

> Douglas E. Comer, **Internetworking with TCP/IP, Volume 1: Principles, Protocols and Architecture, Third Edition**, Prentice-Hall, 1995.

This particular edition covers IP version 4. It does not fully cover the IP version 6 implementation of the IP protocol. The principles and architecture of IP v6 regarding header and content demarcation are similar, except for details of field sizes.

An Internet Protocol Datagram, or IP datagram, consists of a Datagram Header and a Datagram Data Area. See Figure 7.2 from page 92 of Comer. As noted in the caption of that figure, IP does not specify the format of the Datagram Data Area, because it is used to transport arbitrary data of the user or

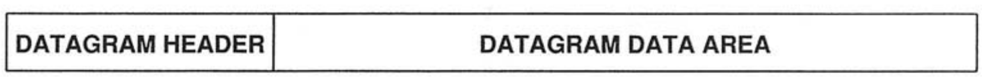

Figure 7.2 General form of an IP datagram, the TCP/IP analogy to a network frame. IP specifies the header format including the source and destination IP addresses. IP does not specify the format of the data area; it can be used to transport arbitrary data.

application's choosing between the source host computer and the destination host computer.
The Datagram Header, however, has a format that is precisely defined by the IP protocol standard. This standardization allows IP routers (also called gateways) to parse and understand the elements that are needed to route the datagram from source to destination, and to provide information about the delivery process. That is: the Datagram Header is precisely analogous to the outside of the envelope used by the postal service, or the shipping labels and markings on a package that are used by a package delivery carrier, whereas the Datagram Data Area is precisely analogous to the “inside of the envelope” or the contents of a package. The meaning and use of the Datagram Data Area, just like the contents of an envelope or package are a) unnecessary for the carrier to do its delivery job, and b) unmodified by the carrier in its delivery process. *All* of the information needed is provided in the Datagram Header.

The format of the Datagram Header is best illustrated by a second figure from Comer (figure 7.3 on page 92):

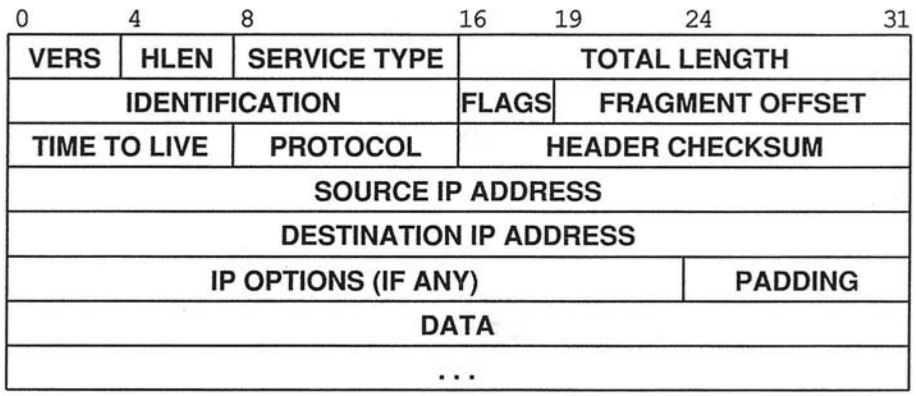

Figure 7.3 Format of an Internet datagram, the basic unit of transfer in a TCP/IP internet.

In this diagram, the Datagram Header consists of the fields preceding the field labeled DATA, which represents the Datagram Data Area noted above. Comer spends several chapters defining the various fields of the header and explaining how they are used in routing each Internet Datagram, etc., so for brevity I will not repeat that here.

However, to be very precise about how the routers in the network delimit the boundaries of the header and the payload, or envelope vs. content, I call attention to two fields of the header: HLEN and TOTAL LENGTH. HLEN is a 4-bit field that occupies the first octet (8-bit field) of the IP datagram. The HLEN field codes the length of the Datagram Header as a binary number, which is to be multiplied by 4 to calculate the size of the Datagram Header in octets, or equivalently, multipled by 32 to calculate the size of the Datagram Header in bits. TOTAL LENGTH is a 16-bit field that occupies the 3rd and 4th octets of the Datagram Header. The TOTAL LENGTH field encodes the size of the Internet datagram as a binary number, which directly represents the size of the Internet Datagram in octets, or can be multiplied by 8 to calculate the size of the Internet Datagram in bits. The size of the Datagram Data Area can be calculated by subtracting the length of the Datagram Header from the length of the Internet Datagram as a whole.

In other words, using fields contained in the first 32 bits of every Internet Datagram, a router calculates the size of the Datagram Header, the size of the Internet Datagram itself, which precisely defines the location and size of the Datagram Data Area.

So-called Application Headers

The content of Internet Datagrams (the Datagram Data Area above) is not intended to be used or disturbed by the transport networks within the Internet, but by the applications and services on host computers attached to the Internet transport that send and receive the Internet Datagrams. The concept of an “application header” is not part of the Internet standard, but is often confused with the headers of the datagram by informal use of the word “header” broadly. Since the term “application headers” sometimes appear to confuse discussions of Internet standards, I briefly discuss what so-called application headers are, and illustrate them here.

Just as we have standardized certain forms of correspondence frequently used in business and private communications (business letters, invoices, remittances) and certain kinds of package structures (internals of wine shipping boxes with separators and padding), Internet application and service programmers and designers have standardized a variety of useful structures to be used within the Datagram Data Area. Some, but not all, of these structures are so widely adopted that they are "built in" to the operating systems on most host computers – for example, DNS, TCP, UDP, FTP, ...

The structure of the Datagram Data Area contents for such applications are the subject of standards that enable communications among diverse endpoints, just as the semi-formal standard for the format of a "standard business letter" taught in schools serves to make business communications more effective and efficient.

Many of these standards include structures within the Datagram Data Area that contain headers with very rigid formats and content areas which can be used for anything. The best analogy for this is the concept used in "standard business letters" of an *enclosure.* A standard business letter contains some structured information such as date, address, salutation, signature, typist, ... while allowing for very diverse structures on other pages, usually noted to be enclosures. For purposes of handling multi-page contents of envelopes once the envelopes have been removed, the primary sheet will note the enclosures, both in number and type. Similarly, one may expect to find various standard items in certain types of packages, such as a "bill of lading", while the remainder can be almost anything at all.

So-called application headers are a loose term often used in network protocol contexts to describe the more strictly formatted portions of certain kinds of data used in particular application or service protocols, which are carried within the Datagram Data Area. They provide useful information or instructions provided by the sender of an application or service message to the ultimate receiver.

What is important about application headers arises from the fact that they are agreements between endpoints. There is no need for the network transporting the Internet datagrams containing these application headers to know what they are, what they mean, or to use any information from them to perform the function of delivering the messages properly to their destination host computer.

By analogy, the post office does not need to read or to understand the contents of the "standard business letter" in order to perform its function. The post office and an individual letter carrier can indeed guess, because of common knowledge, that the contents of a letter destined for a person at a business is *very likely* to contain a page formatted like a standard business letter, but it need not look at the structure of the letter to deliver it, and must not look at it if proper separation of function is to be preserved. As just one example of the problem with this, consider business letters between two Chinese-speaking correspondents in an English-speaking country. There is no need for the post office to understand a "standard Chinese Business Letter" between peers in order to deliver a letter with a proper envelope. There may not even be a direct mapping between the structure of a Chinese business person's correspondence and that of a standard English business person. Nor, and this is important, does the post office need to understand that the contents of a letter involve Chinese persons corresponding with each other, even if the names on the outside of the envelope *appear to be* transliterated Chinese characters.