INTERNET-DRAFT Mark Day
Expires: September 13, 1998 Lotus Development Corporation

''HTTP Envy'' and Presence Information Protocols

draft-day-envy-00.txt

1. Status of this Memo

2. Abstract

There are a variety of proposals [Calsyn, Mohr] for building a presence information protocol as a variant or version of HTTP. This document summarizes why I believe that is not a good idea.

3. Introduction

HTTP is a remarkably ubiquitous protocol. It is not hard to find people who believe approximately one of the following:

"HTTP is everywhere. It could not be so pervasive without being good. Therefore it is good, and I should imitate it with my protocol."

"HTTP is everywhere. Therefore, if I base my protocol on HTTP, my protocol will also be everywhere."

In the rest of this document, I consider the reality of building a presence information protocol, and compare these realities to the HTTP fantasies. Other documents [Dusseault, Day] provide perspective on what a presence information protocol must be.

4. Crossing Firewalls Made Easy

In most organizations HTTP traffic crosses firewalls fairly easily. It is easy to claim (for instance) that a new super-duper protocol will use port 80 and thereby get all of HTTP's firewall-crossing virtues. There are typically two holes in this story.

The first hole is that these protocols typically don't get all the details right so that the protocol data tunnels through HTTP, with intermediates unaware that another protocol is being carried. Instead, there's a hand wave or two in the general direction of modifying or upgrading the deployed proxies and firewalls (a much harder problem than is usually acknowledged).

The second hole in the story is that firewall managers don't like HTTP tunnelling. It makes it difficult to block or control non-HTTP traffic without also slowing down the legitimate (real) HTTP traffic.

In the specific context of a presence information protocol, the HTTP/firewall fantasy has an additional hole: there is no "reverse path" in HTTP from the server to the client, and no particular reason to expect all of the proxies and firewalls involved in a particular client-to-server HTTP request to allow a distinct server-to-client request, even if both client and server agree on the value of such a request.

5. Reuse of well-understood technology

Even if a presence information protocol can't actually be HTTP (and thereby cross tall firewalls in a single bound), perhaps it should be like HTTP so that people can understand it readily and reuse their existing code for clients, proxies, and servers.

But this is revealed to be fantasy when we start looking at detailed proposals. The trouble is that HTTP itself is a reasonably well-understood protocol only when confined to GET requests. As soon as POST and PUT enter the mix, confusion usually follows. The IETF WebDAV working group is building an interoperable and workable semantics for creating and updating information on the Web: that is, WebDAV is an effort (as yet unfinished!) to fix the fact that PUT and POST are broken.

When a proposed protocol based on HTTP introduces new methods or headers, those methods must be related to the existing HTTP methods, headers, and ways of using them. It is not obvious that the use of HTTP saves any effort for reader, writer, or implementor when these issues are taken into account.

There is a relatively powerful and reusable piece of technology used in HTTP that is relevant to presence information and instant messaging: MIME. However, a presence information protocol can pick up MIME without dragging HTTP along with it.

7. Almost the Right Protocol

We might think that HTTP is enough like a presence information protocol, except for the "minor" addition of notifications and instant messaging, that we might as well leverage the existing facilities for retrieving state about people. This is not a bad theory, but properly applied it starts with LDAP instead of HTTP. That is, if we believe that the goal is to find information about people and resources, including contact and rendezvous information, the more natural starting point is a directory service. In this view, the role of a presence information protocol is to serve as an extension to the existing naming & locating services of a directory service, either by extending the directory protocol itself or by operating alongside it.

8. Let's Use URLs

Finally, we might think that URLs represent a particularly good way to represent people. A URL is not necessarily worse than other representations for a low-level location mechanism, hidden from users. But URLs are lousy for representing names for two reasons.

First, URLs are restricted to a weak subset of Latin-1. It is unconscionable that a system intended for global use should build in needless restrictions on accurately writing personal names.

Second, the definition of URLs recognizes the reality and value of aliases, but makes those facilities available primarily for the host name. Simply by using URLs as defined, We can recognize that two URLs are "the same" if the host parts are aliases for the same machine. But without defining additional mechanism, we cannot recognize that two URLs are "the same" if the remainders are aliases for the same person! And if we have to define that piece of machinery, it doesn't seem that we are getting much value from the use of URLs.

9. Conclusion

At least some of the arguments for HTTP as a basis for a presence information protocol -- ease of crossing firewalls, reuse of existing technology, use of similar protocol, and the applicability of URLs -- seem unconvincing. Perhaps there are other, better arguments for the use of HTTP. In the absence of such arguments, HTTP seems like a rather poor choice technically. Its popularity as a basis for presence information protocols seems more driven by fantasies of HTTP-like ubiquity than by rational thought.

10. References

[Calsyn] Martin Calsyn and Lisa Dusseault. "RVP: A Presence Notification Protocol." Internet-Draft draft-calsyn-rvp-01.txt.

[Day] Mark Day. "Requirements for Presence and Instant Messaging." Internet-Draft draft-day-rpim-00.txt.

[Dusseault] Lisa Dusseault. "Presence Information Protocol Requirements." Internet-Draft draft-dusseault-pipr-00.txt

[Mohr] Gordon Mohr. "Widely Hosted Object Data Protocol (WhoDP)". Internet-Draft draft-mohr-whodp-00.txt.

10. Author's Address

Mark Day
Lotus Development Corporation
55 Cambridge Parkway
Cambridge, MA 02142
USA

Mark_Day@lotus.com