MSGTRK BOF
INTERNET-DRAFT
draft-jones-msgtrk-def-01.txt
Expires: September 1999

G. Jones [TomorrowSys@yahoo.com]
B. Ernst [bruce_ernst@lotus.ssw.com]
G. Vaudreuil [gregv@ons.octel.com]

Basic Definition of Message Tracking

Status of this Memo

Abstract

This document defines message tracking as a prelude to the creation of a message tracking model. Message tracking is a messaging management function; it provides the ability to find out, after the fact, the path that a particular message took through the messaging system, the current status of that message, and its characteristics.

Definition

Message tracking refers, in its simplest form, to determining the path an RFC822 message has taken, its current location, and its characteristics. Message tracking allows the originator of a message to issue a request about a previously sent message, the answer to which contains the delivery status, delivery time, delivered recipients, and other information about the message. This is different from the delivery status notification (DSN)

function in use today, because DSNs are requested at the time of submission and are generated automatically; alternatively a tracking request is generated independently of the previously submitted message's status and is done so on demand.

This capability is analogous to the service provided by carriers of conventional paper mail - the ability to quickly locate where a package is, and to determine whether or not the package has been delivered to its destination. An Internet-standard approach will allow development of message tracking applications that can operate in a multi-vendor messaging environment, and will encourage operation of the function across administrative boundaries.

One might ask: why should there be a standard for message tracking when Internet domains will be unwilling to open themselves up to outside tracking requests ? This is one reason why we design and implement Internet standards. So that there is a reliable, secure, agreed upon mechanism for message tracking that people will be willing to use.

Companies have implemented and are implementing message tracking today. Standardization of this technique will aid the Internet user community, make Internet messaging more profitable, and fulfill a key messaging management need.

Reasons for Message Tracking

Message tracking is useful for determining the whereabouts and status of "lost" messages, and for several other purposes:

o When there is a lack of trust in the messaging system, such as when an originator claims a message failed to be delivered, the point of failure may be isolated. This includes messages that were never delivered or messages that were delivered incorrectly. Message tracking thus adds to the overall reliability of the mail system;

o Per-message information can be used for accounting, billing, and performance purposes. Traffic can be itemized on a per-origin or per-destination basis by system or originator. This typically involves two steps - collection of message traffic data, followed by the gehe time they are submitted to an MTA up until the time a network of MTAs discharges the message onward to another entity (e.g. a proprietary mail server, IMAP server, and so on).

o Message tracking information adds security in that the origins of potential security threats can be more precisely determined. If a system were flooded with traffic, for instance, the origin of this traffic would become known. Message tracking information is suitable for routine security audits containing the details of messaging traffic over specific time intervals;

o End-to-end delivery time could be measured;

o Message tracking would aid in message loop detection, since unique message identifiers of looping messages, when these exist, would be recorded multiple times;

o Performance characteristics about the type of messaging traffic could be determined, such as when an inbound message causes the creation of multiple outbound messages, and the percentage of messages that were actually delivery reports or receipts. This is valuable for performance measurement, among other reasons;

o Standardized message tracking information acts as a bridge between dissimilar messaging systems and dissimilar messaging communities;

Tracking Messages on the Public Internet

One might ask: why bother to track messages if a majority of public Internet traffic is point to point; messages don't live long enough to be trackable, and are not an interesting event to track since you always know the next point ? Just because you know where a message went that doesn't mean you know what happened to it, how fast it got there, or what was in it. As the Internet is used more and more for commercial/official purposes a logging function is commonly embedded in the messaging system internally. Even if most of the message traffic is point to point, this point-to-point traffic is inter-domain, across firewalls, and thus it is even more important to have a reliable tracking mechanism that organizations can agree on. It is something that intra-domain messaging users want. Even if 95% of transactions are point to point, the 5% that is non point-to-point is still a huge amount of traffic, and this is exactly the traffic that users will want to track. Once messaging traffic enters an intranet or domain of any size it invariably encounters a more hierarchical routing structure.

Who is Allowed to Track Messages

Only the originators of messages are allowed to track their messages. Optionally, an originator may delegate this responsibility to a third party, but this is left for future study.

How Tracking is Done: Requests and Responses

The originator will issue a message tracking request using the Unique Message Identifier plus security information. The originator (of both the message and the query at this point) will receive optional response criteria such as the message disposition, delivered recipients, delivery time, and the names of MTAs that handled the message.

Security for Message Originators

One option for message security is that the originator calculates a hash A to be equal to the hash of the message ID + time stamp + a per-user secret. The user then calculates hash B to be the hash of A. The user includes B in the submitted message, and retains A. Later, when the user makes a message tracking request to the messaging system or tracking entity, it submits A in the racking request. The entity receiving the tracking request then uses A to calculate B, since it was already provided B, verifying that the requestor is authentic. Summarily

A = H(message ID + time stamp + secret)
B = H(A)

If the originator of a message were to delegate his or her tracking request to a third party by sending them A, this would be vulnerable to snooping over unencrypted sessions, but the user can decide on a message-by-message basis if this risk is acceptable.

Three Possible Architectures

There are ways of accomplishing message tracking without mandating the addition of large amounts of new infrastructure on the participants. Optionally, if more infrastructure is proven to be a good and necessary thing, it should be considered.

In all cases, messages are only tracked from the time they are submitted to an MTA up until the time a network of MTAs discharges the message onward to another entity (e.g. a proprietary mail server, IMAP server, and so on).

The three architectural alternatives offered by the start-up working group to date might be called "ask later", "ask now", and "ask someone else."

Under "ask later", a user requests tracking as a service when submitting a message, and then at a later time issues a separate tracking request to the mail system. The user receives a response to the request from the tracking entity. This has the advantage of being deployable within the existing SMTP infrastructure.

Under "ask now", a user requests tracking as a service while submitting a message, and receives a step-by-step report contemporaneously from each MTA that handles the message. This provides the user with a high level of service, but also causes extra overhead: an additional message generated for each hop the original message takes.

Under "ask someone else", the user issues a separate message tracking request to an entity other than the messaging system at a later time. The user receives a response to the request from this same third-party tracking entity. This has the advantage of allowing tracking to occur when the messaging process has failed but the platform is still working. It also off-loads the tracking function from the messaging system itself. It may,

however, require new infrastructure in order to support it.

One possibility would be to implement "ask now" and "ask later" as SMTP extensions. One could implement "ask someone else" as a UDP- or TCP-based protocol, among other options.

Acknowledgments

Thanks to all those who participated in the message tracking meetings. Many thanks to Ned Freed and Harald Alvesrand for the hashing material.