

A pull-based e-mail architecture

Edson Kageyama
Graduate Program in
Computer Science
PUCPR
Curitiba PR, Brazil
edson@ppgia.pucpr.br

Carlos Maziero
Graduate Program in
Computer Science
PUCPR
Curitiba PR, Brazil
maziero@ppgia.pucpr.br

Altair Santin
Graduate Program in
Computer Science
PUCPR
Curitiba PR, Brazil
santin@ppgia.pucpr.br

ABSTRACT

Conventional e-mail systems are prone to problems that impact their scalability and dependability. E-mail systems operate following a “push-based” approach: the sender side server pushes the e-mails it wants to send to the corresponding receivers’ servers. This approach can impose processing and storage overhead on the receiver side. This paper presents an e-mail architecture in which messages are sent directly from senders to receivers using a “pull-based” approach. The sender stores locally all e-mails it intends to send, and notify their receivers using a global, distributed notification service. Receivers can then retrieve such notifications and decide if they want to receive the corresponding e-mails. If so, e-mails can be retrieved directly from their senders. This proposal is inspired from file sharing peer-to-peer systems, in which users locate and retrieve the contents they are looking for. A prototype was built to show the feasibility of the proposal.

Categories and Subject Descriptors

C.2.4 [Computer Systems Organization]: Computer-Communication Networks—*Distributed Systems, Distributed Applications*

General Terms

Reliability, Security

Keywords

E-mail, Peer-to-peer, Distributed systems

1. INTRODUCTION

Internet e-mail has become an extremely pervasive communication tool, due to its low cost, asynchronous operation and ability to easily transport different types of digital content. However, protocols for sending and receiving e-mails were designed to be simple, mainly because their initial use

was restricted to the academic community and the e-mail servers’ processing capacity was limited. With the widespread use of e-mail, some users started using this system to propagate malicious content, as *virus*, *spams*, and *scams*. Moreover, the increase in the use of e-mail systems led to the necessity of technologies capable to support this service in a more secure, reliable, and scalable manner.

Traditional e-mail systems operate following a “push-based” approach: the sender’s e-mail server pushes the e-mails she wants to send to the corresponding receivers’ servers. The receiver server should then accept the e-mail and deliver it to its destination mailbox. The receiver server can implement mechanisms to avoid receiving spam, like e-mail content analysis [19] and server white lists [7]. However, most techniques impose processing and storage overhead on the receiver side.

This article presents a distributed e-mail architecture in which e-mail is transported according to a “pull-based” approach, directly from senders to receivers. The sender stores locally all e-mail it intends to send, and notifies their receivers using a notification service built using a distributed hash table [17]. Receivers can then retrieve such notifications and decide if they want to receive the corresponding e-mails. If so, e-mails can be retrieved directly from their senders.

This article is structured as follows: section 2 reviews the conventional e-mail structure and the main threats to e-mail services; section 3 reviews some related work on *peer-to-peer* based e-mail systems; section 4 presents the architecture developed in this project; section 5 presents some prototype implementation details; finally, section 6 concludes the paper, discussing open problems and outlining perspectives.

2. THE E-MAIL SERVICE

The Internet e-mail architecture, as defined in [10], contains agents that act as senders and receivers of mail messages. The *Mail User Agent* (MUA) is the program used by the end user for reading and writing e-mails. The *Mail Transport Agent* (MTA) is responsible for receiving e-mails from the MUAs, forwarding them to their destinations (other MTAs), receiving e-mail from other MTAs, and keeping mailboxes for local users.

MUAs and MTAs use the *Simple Mail Transfer Protocol* (SMTP) [10] to forward e-mail to their destination servers. MUAs can use protocols like POP [14] and IMAP [3] for accessing their MTAs, to retrieve the e-mails stored on their mailboxes. E-mails are formatted according to [6, 16].

E-mail systems were initially designed to be simple, as

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAC’08 March 16-20, 2008, Fortaleza, Ceará, Brazil

Copyright 2008 ACM 978-1-59593-753-7/08/0003 ...\$5.00.

their usage was restricted to the academic community. This simple and straightforward structure was responsible for the e-mail popularity. On the Internet, e-mail systems gained importance and visibility, showing limitations and fragilities not initially considered, like privacy (e-mails are transferred in clear text format), sender authentication (SMTP does not provide effective mechanisms to authenticate senders, allowing e-mail forgery), and efficiency (e-mails can carry large attachments, but are sent and stored uncompressed).

Due to such fragilities, problems like spam, scam, and virus propagation compromise the security, performance, robustness, and usability of the current e-mail systems. Several techniques can be used to improve current e-mail systems, like *white/black lists* [7, 8], *content filtering* [19], and *sender authentication* mechanisms [1, 11, 21, 23].

3. PEER-TO-PEER AND E-MAIL

Peer-to-peer (*p2p*) systems present some advantages over traditional client/server systems, like more flexibility and scalability, and smaller costs. Primary *p2p* use was in anonymous file sharing over the Internet, but other application domains are appearing, like audio/video streaming, instant messaging, and distributed storage.

In the *p2p* approach, participants (*peers*) acts as clients and servers for the service being provided. Peers interact among each other in order to announce resources, to locate resources, and to retrieve/use them. Peer-to-peer systems can use a central server to maintain resource location and authentication services; others are decentralized and define “super-peers” in order to carry out resource location algorithms. Some are unstructured topology, while others provide a predefined topology to ease resource location [2].

A *Distributed Hash Table*, or DHT [17], is a special kind of decentralized *p2p* service that stores [*key, value*] pairs. This service is generally provided by a large amount of peers, allowing to store and retrieve information reliably and efficiently. Some DHT systems associate a password *p* and a validity time *t* to each entry: *p* protects the entry against unauthorized removals, and *t* allows the DHT itself to purge old entries. DHTs are used to build more complex services, like resource lookup, distributed file systems, and naming systems. Examples of DHT implementations include *Chord* [20] and *Pastry* [18].

Some studies were done about using the *p2p* approach to improve e-mail systems. In [9], the authors use a DHT for storing user certificates, e-mails and even mailboxes. For sending an e-mail, an *User Agent* (UA) retrieves the receiver certificate from the DHT, encrypts the e-mail body using the receiver’s public key and stores it back on the DHT. It also stores the encrypted message headers in the DHT, using the receiver’s e-mail address as the key. On the other side, the receiver UA uses its e-mail address to lookup for new message headers in the DHT; after decrypting the header, the UA uses the message ID as a key to retrieve the corresponding message body.

The *ePOST* system[13] implements an e-mail system on top of a *p2p* storage infrastructure called *POST* [12], which is built on top of the *Pastry* DHT [18]. In *ePOST*, each participant contributes to the DHT service and to the distributed e-mail storage. E-mails and metadata are encrypted and redundantly stored in the DHT. The local agent acts as a SMTP/IMAP server for the user’s MUA, allowing her to use any conventional e-mail client. Each e-mail

component (header, main body, attachments) is stored separately; thus, a large file sent as attachment to several destinations is stored only once in the DHT. E-mail delivery is done through a notification service called *Scribe*: the sender posts a notification to each receiver, containing the message header and references to the message body and attachments.

The work presented in [22] uses an hybrid *p2p* model, composed by communities and a central server, which provides authentication and peer location services. Each community comprises a set of nodes and a super-node. The super-node can act as a conventional MTA: all messages in a community are sent to its super-node, which can deliver them to local destinations (nodes in the same community) or forward them to other super-nodes, for remote delivery. In another operation mode, the super-nodes can route only the sending requests among peers, letting the e-mail transfers occur directly among them. If a super-node fails, its community elects another super-node; a recovery mechanism is provided to allow the new super-peer to rebuild the state of the previous one. Finally, *broker nodes* can be used to send messages to conventional e-mail servers.

There are other approaches for using *p2p* to improve e-mail systems. In [4], the authors implement a *p2p* system among e-mail servers, to share information about spam messages each one receives. This collective knowledge can improve the spam filtering done by each server.

To our knowledge, all systems proposed rely on *pushing* e-mails to a local server (super-node or local SMTP server), to a remote storage (DHT or remote SMTP server), or directly to the destination nodes. Consequently, the storage overhead does not impact the sender itself, but others. Also, spam/virus filtering should be done by the receivers, *after* receiving the e-mail. Finally, current DHT implementations have small limits on the size of the information that can be stored under each key, preventing the transfer of large messages. The next section proposes an architecture in which e-mails are stored in the sender side until they are *pulled* out by their destination nodes, if they accept to receive them.

4. A PULL-BASED P2P E-MAIL SYSTEM

This article presents a *p2p* e-mail architecture in which messages are transported according to a “pull-based” approach, directly from senders to receivers. The system is composed by a set of peers, each one acting on behalf of an user. Concerning the message flow, the system follows an unstructured approach, in which all peers have the same features. The e-mails are exchanged directly between the peers, there are no central authentication or storage servers. A DHT service is used to forward notifications from senders to receivers, and to store some control information.

A peer is a daemon that acts as an SMTP/IMAP server for its local client. Message format respects the RFC2822 specifications [16], allowing users to use standard client software (MUA) for e-mail handling. Each peer *p2pMTA* has four storage areas: a *spool area*, to store messages sent by the local MUA; an *inbox area*, in which it deposes messages to be read by the local MUA; a publicly accessible *outbox area*, holding messages to be pulled by receivers using HTTP; and a *key cache area*, to store public PGP keys of known peers. PGP keys [23] are used for encrypting messages and notifications, to guarantee their integrity and privacy. Peers organize themselves in groups, to improve message availability (see section 4.5).

Figure 1 shows an overview of this architecture; next sections detail how messages are transferred, and the mechanisms used for ensuring message privacy, integrity, and availability.

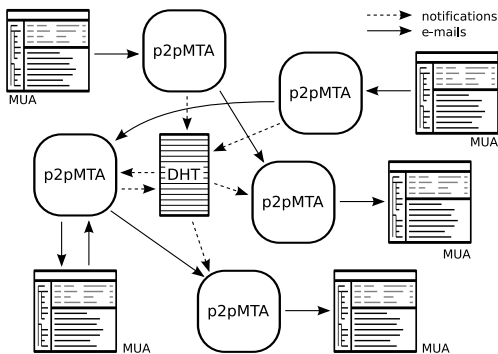


Figure 1: Architecture overview

4.1 Assumptions

We assume that the DHT provides three basic services:

- $put(x, v, p, t)$: register an entry defined by a key x , a value v , a password p , and a validity time t ; the DHT supports key collisions: two or more values can be stored under the same key x ;
- $get(x)$: retrieves all entries $e_i = [x_i, v_i]$ in which $x_i = x$;
- $del(x, v, p)$: deletes an entry defined by $[x, v]$, protected by a password p .

It should be noticed that password p only protects the entry against del operations, but does not forbid get operations. We consider that the DHT infrastructure purges all entries whose validity time was expired, according to [17].

Peer e-mail addresses have the format $name@group$, where $name$ define the peer identity and $group$ is the group of peers to which it belongs.

We assume that each peer has a pair of public/private PGP keys; its public key k is publicly available to the other peers, while its corresponding private key k' is securely kept by the peer. All hash operations are done using the SHA1 algorithm [5].

4.2 Starting an e-mail peer

When it starts, each peer i posts a *peer descriptor* d_i in the DHT, using its e-mail address as key. The descriptor d_i value contains the IP/port of the peer's outbox. The password p is chosen by the peer, and the validity time t is arbitrarily set as one day. This descriptor should be kept in the DHT by the peer, which renews or updates it when needed. If the peer leaves the network without removing its descriptor, the validity time t will ensure that it will be eventually purged from the DHT.

4.3 Sending and receiving an e-mail

The steps needed to transfer an e-mail m_i from a sender peer S to a receiver peer R are explained in the following; they are also presented in figure 2:

1. m_i is sent by the client MUA to its local peer S using SMTP; the message is stored in the spool area of S ;
2. S retrieves k_r from its key cache, or from R , using the URL $http://ip:port/pubkey$ (ip and $port$ are obtained from d_r). If no key is found, an error is raised;
3. m_i is encrypted using k_r : $m'_i = k_r\{m_i\}$, and a SHA1 hash h_i is generated from m'_i : $h_i = SHA_1(m'_i)$;
4. m'_i is then stored in S 's *outbox*, using h_i as its file name;
5. S posts a send notification ns_i to R in the DHT, using $ns:peer@group$ as key (where $peer@group$ is R 's address). The ns_i value contains some m_i headers ($sender@group$, $subject$, $date$, $size$, $name/type/size$ of attachments), h_i , and the DHT entry password p_i (a random value); ns_i value is encrypted using k_r . The validity time t is arbitrarily set to one week. ns_i is also copied locally in S .
6. Periodically, R searches for ns entries in the DHT; each ns_i is decrypted using its private key k'_r , to retrieve the message headers, the hash h_i , and the password p_i .
7. For each ns_i , R creates a message n_i in the local inbox area, using the m_i headers and an empty body;
8. The IMAP client at R retrieves all inbox messages and presents them to the user;
9. If the user decides to open n_i , R retrieves m'_i using the URL $http://ip:port/outbox/h_i$ (ip and $port$ are obtained from d_s); m'_i is then decrypted using k'_r to obtain m_i . Otherwise, if the user deletes an e-mail without reading it, the corresponding inbox entry is deleted.
10. After m'_i being retrieved or n_i deleted, ns_i is deleted from the DHT, and a receive notification nr_i is posted, using $nr:peer@group$ as key (where $peer@group$ is S 's address). nr_i value is the hash h_i , and the same p_i from ns_i is used to protect the DHT entry.
11. Periodically, S searches for nr entries for its pending messages; for each nr_i , it removes the DHT entry and the corresponding files from its outbox.

In this approach the storage overhead remains at the sender side. A spammer trying to send thousands of messages will consume its own storage space until the receivers decide to get them, to delete them, or to ignore them. Another positive aspect of this approach: a message already sent by the user (MUA) can be canceled, if the receiver did not yet receive it.

4.4 Multiple receivers

The steps shown in the last section present the system behavior in a simple case, in which a peer send a message to a single receiver. Now, a multiple-receiver message transfer scenario is depicted. This scenario is frequent in mailing lists, for instance. The changes in the previous scenario when sending a message m_i to several receivers R_x are:




Figure 2: Sending/receiving an e-mail

- m_i should be encrypted with the public key k_x of each receiver R_x . PGP does this easily: it cyphers m_i using a random session key ks ($m'_i = ks\{m_i\}$), then cyphers ks using k_x of each R_x , to get cyphered keys kc ($kc_x = k_x\{ks\}\forall R_x$); all kc_x are then appended to m'_i ;
- Next, S should post a ns_x for each receiver R_x ;
- Only when all ns_x are replied with nr_x or are expired, m_i files stored in S can be removed.

It is easy to see that the message is stored only once, even if the message has several receivers.

4.5 Peer groups

In a conventional e-mail system, the mail server acts as a temporary storage for e-mails, if their destination servers are offline. Our architecture proposal has no e-mail servers or intermediate temporary storage areas. This could cause problems, because sender and receiver should be both on-line to allow an e-mail to be transferred. To circumvent this restriction, peers are organized in *peer groups*.

Peers in a group replicate the contents of their outboxes, in order to provide higher availability of the e-mails awaiting to be retrieved by their receivers. A peer group is formed by peers that have some trust in each other. Each group has a name $gname$ and a descriptor gd_i registered in the DHT under a $group:gname$ key. The descriptor value contains a list of the group members' addresses. For each group, a special peer called *group master* is responsible for registering and maintaining its group descriptor. Currently, the group master is manually defined, as well as group members; more sophisticated techniques for building and maintaining groups are being investigated.

The outbox replication procedure is fairly simple: a given peer in a group retrieves lists of messages contained in each other peers' outboxes, from their URLs `http://ip:port/outbox/msglist`. Then it compares the list contents with the e-mails present in its local outbox; e-mails no more in the list should be deleted, because they were deleted by their owner; e-mails not in the outbox should be retrieved, to be replicated.

During a message transfer, if the sender peer S is offline, the receiver peer R can retrieve the sender group descriptor,

in order to know other peers where the message m'_i may be replicated. It will then try to retrieve the message from one of them (step 9 in section 4.3).

4.6 Garbage Collection

There are a fair amount of information being stored locally (at each peer) and in the DHT. This information should be deleted when no more needed:

- All DHT entries have validity times associated to them; when the entry age surpasses its validity time, the DHT infrastructure removes it, transparently [17].
- E-mails stored in the sender's outbox are deleted when (a) the receiver retrieved it and posted a corresponding receive notification; or (b) when the validity time of the corresponding entry in the DHT finishes.
- E-mails replicated among the group members are deleted during the replication procedure, if their owners deleted them from their respective outboxes.

5. PROTOTYPE

A prototype of this architecture was built, in order to show its feasibility. The prototype was built using the following COTS components: the *OpenDHT* infrastructure [17], which is a deployment of the *Bamboo* DHT implementation over the *PlanetLab* distributed platform [15]; the *GnuPG* package, an open source implementation of the *OpenPGP* standard; *Lighttpd*, a lightweight open source HTTP server; the *JES - Java E-mail* SMTP server; the *Binc* IMAP server; and the Mozilla Thunderbird e-mail client.

The prototype is fully functional, providing all services defined in the previous sections. Now, experiments are being carried out to estimate (a) the costs to transfer a message from a sender to a receiver, (b) the local storage costs in each peer, (c) the traffic generated by the replication mechanism. We intend to compare these costs with the corresponding costs in a conventional SMTP infrastructure. First results show that a slight reduction of the traffic and storage was obtained, mainly because *GnuPG* compresses the messages before encrypting them. So, all the messages transferred are compressed and encrypted. Also, the system scalability and its dependency on the DHT scalability should be evaluated.

6. CONCLUSION

In this paper we presented an alternative architecture for e-mail distribution over the Internet. In this architecture, the e-mails are stored in the sender side; senders post notifications in a global notification service provided by a distributed hash table (DHT). Receivers periodically query the DHT for new notifications; they can then retrieve the corresponding e-mails directly from the senders, or just ignore them. Our architecture got its inspiration from file sharing peer-to-peer systems, in which users locate the resources of their interest and retrieve only the files they choose. Also, messages are replicated among peers in order to improve their availability.

The main contribution of this architecture is to leave the storage overhead on the sender side; in this context, a spammer would fill its own outbox (and the outboxes of its group members) instead of receivers' ones. Also, if a receiver decides to not receive an e-mail, it will not be transferred from the sender, saving its bandwidth.

The architecture was implemented in a prototype using COTS components. Now we are carrying out experiments in order to compare it to the conventional approach, and to estimate its scalability and robustness.

There are some issues to be solved, and improvements that could be done in this proposal. First, peers should be reachable through HTTP (or other transport protocol), because receivers will retrieve their e-mails using it; this can be a problem for peers behind a firewall or NAT router. Also, the integration of this system to conventional e-mail systems should be investigated; a possible solution for this issue is to use *relay* peers that would create a bridge between both environments. In our prototype, the peer group is manually defined and static; other methods for defining and evolving peer groups should be investigated. Finally, as each e-mail is replicated among the sender group members, the receiver could retrieve parts of it (fragments) from each group member, to improve retrieval speed; this approach is already used in p2p file sharing services.

7. REFERENCES

- [1] E. Allman, J. Callas, M. Delany, M. Libbey, J. Fenton, and M. Thomas. DomainKeys identified mail (DKIM) signatures. IETF RFC-4871, 2007.
- [2] S. Androutsellis-Theotokis and D. Spinellis. A survey of peer-to-peer content distribution technologies. *ACM Computing Surveys*, 36(4), 2004.
- [3] M. Crispin. Internet message access protocol version 4 rev1. IETF RFC-3501, 2003.
- [4] E. Damiani, S. Vimercati, S. Paraboschi, and P. Samarati. P2P-based collaborative spam detection and filtering. In *IEEE Peer-to-Peer Computing*, 2004.
- [5] D. Eastlake and P. Jones. US secure hash algorithm 1 (SHA1). IETF RFC-3174, 2001.
- [6] N. Freed and N. Borenstein. Multipurpose internet mail extensions (MIME) part one: Format of internet message bodies. IETF RFC-2045, 1996.
- [7] R. J. Hall. How to avoid unwanted email. *Communications of the ACM*, 41(3), 1998.
- [8] J. Jung and E. Sit. An Empirical Study of Spam Traffic and the Use of DNS Black Lists. In *ACM Internet Measurement Conference*, 2004.
- [9] J. Kangasharju, K. Ross, and D. Turner. Secure and resilient peer-to-peer e-mail design and implementation. In *IEEE P2P Computing*, 2003.
- [10] J. Klensin. Simple mail transfer protocol. IETF RFC-2821, 2001.
- [11] J. Lyon and M. Wong. Sender ID: Authenticating e-mail. IETF RFC-4406, 2006.
- [12] A. Mislove and A. Post. POST: A secure, resilient, cooperative messaging system. In *USENIX HotOS*, 2003.
- [13] A. Mislove, A. Post, A. Haeberlen, and P. Druschel. Experiences in building and operating ePOST, a reliable peer-to-peer application. In *ACM EuroSys*, 2006.
- [14] J. Myers and M. Rose. Post office protocol - version 3. IETF RFC-1939, 2001.
- [15] L. Peterson, A. Bavier, M. Fluczynski, and S. Muir. Experiences implementing PlanetLab. In *USENIX OSDI*, 2006.
- [16] P. Resnick. Internet message format. IETF RFC-2822, 2001.
- [17] S. Rhea, B. Godfrey, B. Karp, J. Kubiatowicz, S. Ratnasamy, S. Shenker, I. Stoica, and H. Yu. OpenDHT: a public DHT service and its uses. In *ACM SIGCOMM*, 2005.
- [18] A. Rowstron and P. Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *ACM/IFIP/USENIX Middleware*, 2001.
- [19] M. Sahami, S. Dumais, D. Heckerman, and E. Horvitz. A bayesian approach to filtering junk E-mail. In *AAAI'98 Workshop on Learning for Text Categorization*, 1998.
- [20] I. Stoica, R. Morris, D. Karger, M. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *ACM SIGCOMM*, 2001.
- [21] M. Wong and W. Schlitt. Sender policy framework (SPF) for authorizing use of domains in e-mail, version 1. IETF RFC-4408, 2006.
- [22] Y. Zhao, S. Zhou, and A. Zhou. E-mail services on hybrid p2p networks. In *Grid and Cooperative Computing Conference*, 2004.
- [23] P. Zimmermann. *The Official PGP User's Guide*. The MIT Press, 1995.