

Forum/Maillist Replication

From P4
< Forum

Replicating the maillist bi-directionally is not a feature that came with IPB. Our implementation tries to have as little an impact on the forum software as possible, which we hope will make the integration more maintainable. The two places we change existing code are in the post handling class and in the XMLRPC interface.

Linking emails to threads

In order to ensure that we post emails to the correct topics and that emails sent from the integration appear in the correct email threads, we store a map of post IDs against the message IDs. The maillist ID is also including so that we can support multiple maillist integrations at once.

If an incoming email has an "In-Reply-To" header, we find the refereed to post and post the email to the topic that the cross-referenced post belongs. This has a positive side effect that if a topic that started in a replicated forum is moved to a different forum, any associated emails from the mailling list will still be replicated into that topic.

Whenever a post is made in the forum, we check whether it's parent (the last post in the topic, or the post that it quotes in the same topic) was sent to the mailing list; if the post ID has an associated message ID. If it does, we add the post's body and parent's message ID to the queue of outgoing emails. If the post's parent was not sent to the mailing list, or if the post is the first in the topic, we check whether the topic is in a forum that should be replicated and if so, queue it without the parent post's ID. It will be sent with no "In-Reply-To" header. This again facilitates the movement of a replicated topic across forums without breaking the email thread's replication.

Replication process

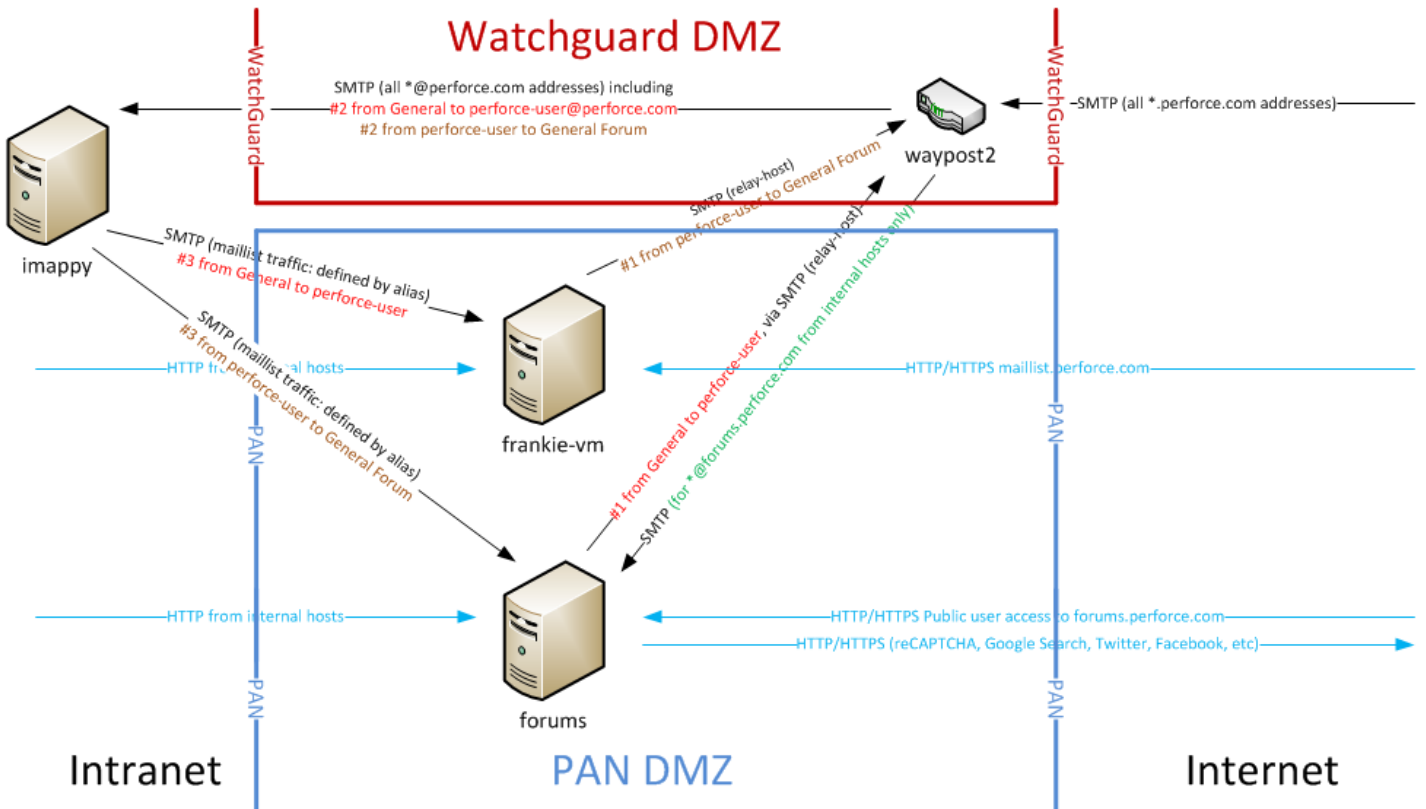
1. Outgoing emails

- In the forum code:
 1. On successful post, that post's content and some metadata is added to the maillist_mailqueue table.
- Every 5 minutes the IPB2Mailman.php script is executed by cron.
 1. Are these any emails in the maillist_mailqueue table? If so, iterate over them.
 1. Is the author of the post a member of one of the blocked groups? If so, skip this email.
 2. Construct the email's headers; if a parent post ID is set, resolve it's message ID and add the "In-Reply-To" header.
 3. Resolve the maillist's email address from the maillist table. (perforce-user@perforce.com for P4User)
 4. Format the email as UTF-8 plain text.
 5. Send email to maillist address.
 6. Email enters local postfix mailqueue.
 7. Email to delivered to waypost2 (the relay host) (see **red #1** in the diagram).
 8. Email to delivered to imappy (the recipient is a perforce.com address) (see **red #2** in the diagram).
 9. Recipient resolved to the maillist on frankie-vm by the aliases file (see **red #3** in the diagram).
 10. Email is delivered to frankie-vm where Mailman accepts the post to the mailing list and distributes it to the list members. (The sender address must match that of a member of the list: perforce-user-forum@forums.perforce.com).

2. Incoming emails

- A user on the forum machine is a member of the mailing list being replicated. All emails to the list are delivered to this user's Maildir folder structure where Dovecot makes them accessible via IMAP.
 1. Mailman sends the new post to all users on the list, including the user that we're using for replication.
 2. The emails are queued up in the local postfix mailqueue before being delivered to waypost2 (frankie-vm's replay host) (see **light brown #1** in the diagram).
 3. [Obsolete] Waypost2 is aware of forums.perforce.com and sends emails to this host directly to this host (never leaving the DMZ).
 4. Waypost2 routes the email to imappy (see **light brown #2** in the diagram).
 5. The new post's email is delivered to the forums machine (see **light brown #3** in the diagram) and stored in the replication user's Maildir to be made available via IMAP.
- Every 5 minutes the IPB2Mailman.php script is executed by cron
 1. Iterate over the maillist's configured in the maillist table.
 1. Log into the associated IMAP account.
 2. Iterate over any emails in the inbox.
 1. Does the email come from the correct mailing list? If not, skip this one.
 2. Does the email come from the mailling list integration's email account (did we send this)? If so, skip this one (we really don't want to recursively post the same message!).
 3. Format the email as BBCode/HTML.
 4. Do we know the user who sent the email? If so, we'll pretend to be them when we post.
 5. Is the email in reply to an existing post? If so, post the email as a reply to the topic. Otherwise, post a new topic, using the email subject as the topic title (stripping out the maillist tag).
 6. Did that fail? If so exit the loop now. We DO NOT want to continue posting emails if the replication has broken. This is very likely to cause messages to be posted out-of-order and/or become segmented over multiple new topics.
 7. Delete the email from the IMAP folder.

System layout



Maillist / Forums Integration Architecture Diagram 15 May 2013

Key:

- Mail Routing - Light Green - for *@forums.performe.com from internal hosts only - ? hops
- Mail Routing - Light Brown - from performe-user to General Forum - frankie-vm to forums - 3 hops
- Mail Routing - Red - from General Forum to performe-user - forums to frankie-vm - 3 hops
- Black - SMTP protocol
- Light Blue - HTTP/HTTPS protocols
- Red - WatchGuard DMZ
- Blue - PAN DMZ

Retrieved from "http://wiki.performe.com/?title=Forum/Maillist_Replication&oldid=178993"

- This page was last modified on 17 May 2013, at 16:00.
- This page has been accessed 4,795 times.