

Listserv Distribute Protocol

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard. Distribution of this memo is unlimited.

Abstract

This memo specifies a subset of the distribution protocol used by the BITNET LISTSERV to deliver mail messages to large amounts of recipients. This protocol, known as DISTRIBUTE, optimizes the distribution by sending a single copy of the message over heavily loaded links, insofar as topological information is available to guide such decisions, and reduces the average turnaround time for large mailing lists to 5-15 minutes on the average. This memo describes a simple interface allowing non-BITNET mailing list exploders (or other bulk-delivery scripts) to take advantage of this service by letting the BITNET distribution network take care of the delivery.

Introduction

Running a mailing list of 1,000 subscribers or more with plain "sendmail" while keeping turnaround time to a reasonable level is no easy task. Due mostly to its limited bandwidth in the mid-80's, BITNET has developed an efficient bulk delivery protocol for its mailing lists. Originally introduced in 1986, this protocol was refined little by little and now carries 2-6 million mail messages a day. In fact, this distribution mechanism implements a general-purpose delivery service which can be used by any user of BITNET or the Internet. Thus, a simple solution to the "sendmail" turnaround problem is to wrap the message and recipient list in a DISTRIBUTE envelope and pass it to a BITNET server for delivery. This may not be the best possible solution, but it has the advantage of being easy to implement.

In this document we will use the term "production" to refer to the normal operation of the mailing list (or bulk delivery application) you want to pipe through the DISTRIBUTE service. That is, the "production" options are those you should specify once everything is tested and you are confident that the setup is working to your

satisfaction. In contrast, "test" and "debug" options can be used to experiment with the protocol but should not be used for normal operation because of the additional bandwidth and CPU time required to generate the various informational reports.

Finally, it should be noted that the DISTRIBUTE protocol was developed to address a number of issues, some of them relevant only to BITNET, and has evolved since 1986 while keeping a compatible syntax. For the sake of brevity, this RFC describes only a small subset of the available options and syntax. This is why the syntax may appear unnecessarily complicated or even illogical.

1. Selecting an entry point into the DISTRIBUTE backbone

The first thing you have to do is to find a suitable site to submit your distributions to. For testing, and for testing ONLY, you can use:

LISTSERV@SEARN.SUNET.SE

For production use, however, you should select a DISTRIBUTE site in your topological vicinity: it would make no sense to pass your distributions from California to a server in Sweden if most of your recipients are in the US. If your organization is connected to BITNET and your BITNET system is part of the DISTRIBUTE backbone, this ought to be your best bet. Otherwise you will want to contact someone knowledgeable about BITNET (or the author of this RFC if you have no BITNET users). Make sure to run through the following checklist before sending any production traffic to the site in question:

- a. Do you have good connectivity to the host in question? Does the host, in general, have decent BITNET connectivity? There are still a few sites that insist on using 9.6k leased lines for BITNET in spite of having T1 IP access. You will want to avoid them.
- b. Send mail to the server with "show version" in the message body (not in the subject field, which is ignored). Is the server running version 1.7f or higher? If so, it should not have given you the following warning,

>>> This server is configured to use PUNCH format for mail <<<

which means that messages with lines longer than 80 characters cannot be handled properly. If the software version is less than 1.7f, the warning will not be present; instead, check the first (bottom) "Received:" field. If it does not say "LMail", do not use this server as it probably cannot handle messages with long lines.

Finally, make sure that the "Master nodes file" is not older than 2 months: there are a handful of sites which never update their tables due to staffing problems. They cannot be prevented from running LISTSERV, but you will certainly want to avoid them.

- c. How big is your workload? If you are planning to use the service for more than 10,000 daily recipients, you should get permission from the LISTSERV administrator, both as a matter of courtesy and to hear about any restrictions or regularly scheduled downtime they might have. For instance, some universities might not allow large distributions during prime time, or they may have several DISTRIBUTE machines and will want to make sure you use the "right" one. Send mail to "owner-listserv" at the host in question and give an estimate of the amount of daily messages and recipients you would like to submit. If your message bounces back with "No such local user" or the like, it means the server did not pass the above test (b) and you don't want to use it anyway.

An index of sites/hosts which have the required configuration, good connectivity, keep their tables up to date and have generally agreed to provide this service to anyone in their topological area will be published separately in the future.

2. Physical delivery of the DISTRIBUTE request

The distribution request is delivered via SMTP to the e-mail address obtained in step 1 (for instance, LISTSERV@SEARN.SUNET.SE). In fact, as long as you can somehow get mail to the server's host, you can use the service; SMTP is just the most convenient way of doing so.

2.1. Contents of MAIL FROM: field

You should set the MAIL FROM: field to the address of the person who maintains your mailing list or, generally speaking, to the address of a human being who can take action in case the message fails to reach the DISTRIBUTE server's host. This is a very rare occurrence.

2.2. Contents of RCPT TO: field

The RCPT TO: field points to the server's address (for instance, LISTSERV@SEARN.SUNET.SE).

2.3. Contents of the RFC822 header

After the DATA instruction, you must supply a valid RFC822 header with a "From:" field pointing to the mailbox that should receive notification of delivery problems, bounced mail, and so on. This can be the same as the MAIL FROM: field, an address of the type "owner-

xxxx@yourhost", etc. DO NOT PUT THE LIST SUBMISSION ADDRESS THERE, or you will get mailing loops.

For testing, the "From:" field should point to your own mailbox, so that you get the responses from the server.

As long as RFC822 syntax is respected, the only field that matters is the "From:" field (or "Sender:", "Resent-From:", etc.). In practice this means you can just pipe the distribution request into "mail listserv@whatever" and let your mail program build all the headers.

3. Format of the DISTRIBUTE request

The body of the message delivered to LISTSERV defines the recipients of the distribution and the text (header + body) of the RFC822 message you want to have delivered. The request starts with a "job card", followed by a DISTRIBUTE command, a list of recipients, and finally the message header and body.

3.1. Syntax of the JOB card

The purpose of the JOB card is to make sure that any spurious text inserted by mail gateways or the like is flushed and not erroneously interpreted as a command. It can optionally be used to associate a "job name" with the request, in case you want to use tools to assist you in processing the notifications you get from the DISTRIBUTE servers when running in test mode. The syntax is as follows:

```
//jobname JOB ECHO=NO
```

"jobname" can be anything as long as it does not contain blanks, and can be omitted. LISTSERV generally ignores case when parsing commands, so you can use "job" or "Job" if you prefer. The ECHO=NO keyword is required for production use, to suppress the "resource usage summary" you would otherwise get upon completion of your delivery. You may want to omit it when testing.

3.2. Syntax of the DISTRIBUTE command

Below the JOB card, you must supply the following line:

```
DISTRIBUTE MAIL
```

For production mode, do not specify anything else on that line. When testing, you should add ACK=MAIL in order to get an acknowledgement confirming the delivery. There are two other useful options: DEBUG=YES, which instructs the server to produce a report showing how the various recipients will be routed, but without actually

delivering the message; and TRACE=YES, which does the same but does deliver the message. Before making a "live" test with your actual recipients list, you should tack the DEBUG=YES option once to make sure you got all the parameters and syntax right, and get a rough idea of the efficiency of the distribution (see the section on performance).

3.3. Giving the list of recipients

The list of recipients follows the DISTRIBUTE line and is specified as follows:

```
//To DD *
user1@host1 BSMTP
user2@host2 BSMTP
/*
```

The two lines starting with a "/" have to be copied as-is. Each of the lines in between contains the address of one of the recipients, followed by a blank and by the word "BSMTP", which indicates that you do not want the header rewritten. There are four restrictions:

- a. The address must be a plain "local-part@hostname" - no name string, no angle bracket, no source route, etc. Bear in mind that the DISTRIBUTE server is not in the same domain as you: all the addresses should be fully qualified.
- b. If the local-part is quoted, it must be quoted from the first word on. Technically, RFC822 allows: Joe."Now@Home".Smith@xyz.edu, but for performance reasons this form is not supported. Just quote the first word to tell LISTSERV to run the address through the full parser: you would write "Joe"."Now@Home".Smith@xyz.edu instead.
- c. The local-part of the address may not start with an (unquoted) asterisk. You can bypass this restriction by quoting the local part and using a %-hack through the server's host:
"***JACK***%jack-ws.xyz.edu"@server-host.
- d. Blanks are not allowed anywhere in the address.

You can use the pseudo-domain ".BITNET" for BITNET recipients: it is always supported within DISTRIBUTE requests.

3.4. Specifying the message text

After the last recipient and the closing "/*", add the following line,

```
//Data DD *,EOF
```

followed by the RFC822 message (header + body) that you want delivered. The EOF option indicates that the message header and body will extend until the end of the message you are sending to the DISTRIBUTE server. If you are worried about extraneous data being appended by a gateway, remove the EOF option, add a closing "/" line after the end of the message, followed by a "// EOJ" card to flush any remaining text. This, however, will fail if the message itself contains a "/" line; you would have to insert a space before any such line.

4. Examples

Here is an (intentionally short) example to clarify the syntax:

```
----- cut here -----
//Test JOB
Distribute mail Ack=mail Debug=yes
//To DD *
joe@ws-4.xyz.edu BSMTP
jack@abc.com BSMTP
jim@tamvml.bitnet BSMTP
jill@alpha.cc.buffalo.edu BSMTP
james@library.rice.edu BSMTP
/*
//Data DD *,EOF
Date:      Tue, 19 Jan 1993 10:57:29 -0500
From:      Robert H. Smith <RHS@eta.abc.com>
Subject:   Re: Problem with V5.41
To:        somelist@some.host.edu
```

I agree with Jack, V5.41 is not a stable release. I had to fall back to V5.40 within 5 minutes of installation...

Bob Smith

```
----- cut here -----
```

Note: some of the hostnames are genuine, but the usernames are all fictitious.

You would get the following reply:

```
-----
Job "Test" started on 20 Feb 1993 01:09:40
```

```
> Distribute mail ack=mail debug=yes
Debug trace information:
```

```

ABC.COM                goes to SEARN    (213) - single recipient
ALPHA.CC.BUFFALO.EDU  goes to UBVM    (027) - single recipient
LIBRARY.RICE.EDU      goes to RICEVM1 (022) - single recipient
TAMVM1                 goes to TAIVM1   (247) - single recipient
WS-4.XYZ.EDU          goes to SEARN    (213) - single recipient

```

Path information:

```

TAIVM1  : UGA      RICEVM1  TAIVM1
UBVM    : UGA      UBVM
RICEVM1 : UGA      RICEVM1

```

```

(Debug) Mail forwarded to LISTSERV@UGA      for 3 recipients.
(Debug) Mail posted via BSMTTP to jack@ABC.COM.
(Debug) Mail posted via BSMTTP to joe@WS-4.XYZ.EDU.

```

Job "Test" ended on 20 Feb 1993 01:09:40

Summary of resource utilization

```

-----
CPU time:          0.086 sec           Device I/O:        6
Overhead CPU:     0.045 sec           Paging I/O:        5
CPU model:        9221                DASD model:       3380
-----

```

To actually perform the distribution and get an acknowledgement, you would change the first two lines as follows:

```

----- cut here -----
//Test JOB Echo=NO
Distribute mail Ack=mail
-----

```

And you would get the following reply:

```

-----
Mail forwarded to LISTSERV@UGA      for 3 recipients.
Mail posted via BSMTTP to jack@ABC.COM.
Mail posted via BSMTTP to joe@WS-4.XYZ.EDU.
-----

```

Finally, by removing the "Ack=mail" keyword you would perform a "silent" distribution without any acknowledgement, suitable for production mode.

5. Performance

The efficiency of the distribution depends mostly on the quality and accuracy of the topological information available to the DISTRIBUTE server (and, in some extreme cases, on system load). For BITNET recipients, the typical turnaround time for reasonably well connected systems is 5-15 minutes. Internet recipients fall in two categories: those which can be routed to a machine within or close to the recipient's organization (average turnaround time 5-20 minutes), and those for which no topological information is available at all. In that case the delivery can take much longer, but usually remains faster than with a vanilla sendmail setup. At the time being, topological information is available for most top-level domains outside the US and for many sub-domains of EDU and GOV.

You can measure the efficiency of the distribution using the DEBUG=YES option as explained above. Recipients which get forwarded to another server usually get delivered within 5-20 minutes (except to poorly connected sites or countries, for which not much can be done). Recipients which are handled locally are passed to a local SMTP agent whose efficiency depends very much on the amount of "burst" queries the local name server can handle in quick succession.

A number of projects are currently underway to investigate the feasibility of improving the quality of the topological information available to the DISTRIBUTE servers for the Internet.

Security Considerations

Security issues are not discussed in this memo.

Author's Address

Eric Thomas
Swedish University Network
Dr.Kristinas vaeg 37B
100 44 Stockholm, Sweden

E-mail: ERIC@SEARN.SUNET.SE