

## Alternatives for Enhancing RTCP Scalability

1. Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress.''

To learn the current status of any Internet-Draft, please check the ``[id-abstracts.txt](#)'' listing contained in the Internet-Drafts Shadow Directories on [ds.internic.net](#) (US East Coast), [nic.nordu.net](#) (Europe), [ftp.isi.edu](#) (US West Coast), or [munnari.oz.au](#) (Pacific Rim).

The distribution of this memo is unlimited. It is filed as <[draft-aboba-rtpscale-02.txt](#)>, and expires August 1, 1997. Please send comments to the authors.

2. Abstract

This document discusses current issues with RTCP scalability as well as describing the advantages and disadvantages of possible solutions.

3. Requirements

In evaluating a solution to the RTCP scaling problem, it is important to have an understanding of the requirements that a proposed solution must meet. This document proposes four requirements:

- Decrease in convergence time
- Decrease in forwarding table entries
- Improvement in receiver reporting information
- Ability to adjust RTCP bandwidth fraction

3.1. Decrease in convergence time

Currently RTCP relies on multicasting of receiver reports to establish

a receiver-population estimate. This shared-state is then used by receivers to establish the receiver reporting interval. The RTCP reporting algorithm described in [1], provides for an initial RTCP reporting interval randomized on [1.25,3.75]. While the use of a

INTERNET-DRAFT

28 January 1997

minimum waiting time for the initial report allows a receiver to listen for other reports prior to sending, the reporting algorithm described in [1] does not mandate this. As a result, reporting hosts not implementing reconsideration (in which the reporting interval is readjusted prior to sending receiver reports) will send their reports within the initially calculated interval, resulting in a potential flood of initial reports.

### [3.2.](#) Decrease in forwarding table entries

As defined in [1], RTCP messages are sent on the same group as the original RTP transmission, but on adjacent ports. This is problematic for multicast routing protocols such as sparse mode PIM, which determine whether a given group will be hosted on the shared tree or source tree based on the sending rate. If an RTP group is placed on the source tree, then if RTCP messages are sent to the same group, then each host sending RTCP receiver reports will result in a forwarding table entry.

However, it should be noted that even if a shared-tree routing protocol (such as CBT) is used, a problem will still exist when such a domain is connected to a source-tree domain using a routing protocol such as MOSPF or DVMRP. This is because today there is no mechanism for a multicast area border router to summarize RTCP activity within a multicast routing domain. Therefore if a shared tree domain is connected to another domain (such as a DVMRP domain) then packets from each RTCP source within the shared tree domain will need to be forwarded into the source tree domain.

This problem is very common given that DVMRP functions today as the MBONE's multicast routing protocol. DVMRP generates forwarding cache entries on demand for each (source network, destination group) pair, and supports source-specific prune messages. The typical forwarding cache expiration time is 200 seconds. Since for a large conference the RTCP reporting interval will typically exceed the forwarding cache expiration time shortly after conference initiation, only a portion of

all (source network, destination group) pairs would be cached by a DVMRP implementation at a given time. Nevertheless, the router will expend considerable resources in adding and flushing forwarding cache entries, as well as in responding to source-specific prune requests. While future versions of DVMRP may prove more scalable, it is unlikely that these versions will be widely deployed within the next 18 months. As a result, summarization messages are desirable in multicast routing protocols, as well as a decrease in the number senders on the RTCP group.

### [3.3.](#) Improvement in receiver reporting information

The RTCP receiver report mechanism provides important information on listenership, reception quality, and bandwidth availability. This information is important useful both for engineering and business purposes. Engineers use reception quality information to diagnose

problems with the network. Sending systems can use reception quality and bandwidth availability information to modify transmission parameters such as compression ratio and sending rate. Business people use information on listenership to track the size of the audience, and reception quality information to measure the quality of the user experience.

Since in large conferences the algorithm described in [\[1\]](#) leads to infrequent receiver reporting, it can be argued that it does not meet the requirements for timely receiver reporting. Therefore any scaling "improvement" should not hamper the ability to collect this information, and probably should be expected to result in improved reporting.

### [3.4.](#) Ability to adjust RTCP bandwidth fraction

In low or asymmetric bandwidth situations, the fraction of session bandwidth reserved for RTCP may need to be decreased below the five percent specified in [\[1\]](#). For example, in Cable Internet it is common for downstream bandwidth to be a factor of twenty greater than upstream bandwidth; allowing five percent of downstream bandwidth to be used for receiver reporting upstream would result in congestion of the upstream channel. In such situations it is necessary to allow the bandwidth fraction to be adjusted.

#### [4.](#) Scaling alternatives

Several alternatives have been proposed for improving the scalability of RTCP. These include:

- Turning off RTCP receiver reports
- Improved convergence algorithms
- Use of separate multicast groups for RTP and RTCP
- Use of separate multicast groups for receiver and sender reports
- Unicasting of receiver reports
- Selective reporting
- Use of RTCP BYE messages
- Use of TTL scoping and summary messages
- Use of RTP translators

These alternatives are discussed in turn.

##### [4.1.](#) Turning off RTCP receiver reports

In some applications (such as satellite transmission) there may be no upstream channel for transmission of RTCP receiver reports. As a result, it may be desirable to allow receiver reporting to be turned off. Since there is no receiver reporting, there would be no way to estimate the receiver population.

Influence on RTCP receivers could be exercised via the SDP announcement mechanism. For example, Mark Handley has proposed that the

session profile specified in SDP via the "m=" line be used to define the desired RTCP behavior. This would allow for turning off of RTCP receiver reports, or other modifications to reporting behavior.

While this proposal would eliminate problems relating to RTCP bandwidth usage in dialup and asymmetric systems, it would deprive interested parties of information on listenership and reception quality. This will prevent senders from making adjustments to their transmission parameters, and would also prevent RTP monitors from providing listenership estimates needed to justify advertising rates.

However, it is not clear that these arguments are persuasive in all cases. RTCP receiver reports serve multiple purposes. One of these is to provide information on listenership; another is to provide information on reception quality and bandwidth availability. In some Cable Internet systems, it is possible to gain information on listenership via IGMP join and leave messages, since if the upstream and downstream channels are separated receivers do not necessarily hear each other's join or leave messages. Furthermore, it can be argued that the state of multicast diagnostics will eventually reach the point where RTCP receiver reports will no longer be needed for diagnostic purposes. Where receiver-based rate adjustment is used, there may not be a need for the sender to adjust their transmission parameters.

#### [4.2.](#) Improved convergence algorithms

Just as non-linear equation solvers can benefit from improved algorithms, it may be possible to improve RTP receiver population estimates via improvements in the population estimation algorithm. Improvements include reconsideration, as well as an initial population estimate, and the incorporation of additional information into the calculation.

The idea behind reconsideration, first proposed by Van Jacobson, is to improve the convergence time by calculating a new estimate of group size and reporting interval prior to sending the RTCP report. The report is only sent if the updated estimate is less than or equal to the original estimate; otherwise the report is not sent and the reporting interval is adjusted to the new estimate. Use of reconsideration greatly reduces the number of initial reports, since a host joining a conference already in progress will be brought to equilibrium prior to sending its first report. As a result, implementation of reconsideration should be made mandatory in future versions of the RTP specification.

In certain circumstances, it may be desirable to increase the dampening effect of reconsideration by increasing the initial minimum reporting interval. This will decrease the number of initial reports from sites with long delays, i.e. sites with a satellite downlink and POTS or ISDN uplink, as well as to decrease the effect of sites reporting early merely due to statistical considerations.

Currently the receiver population estimate has an initial value of 1, but it is possible for an initial population estimate to be supplied in the session announcement. Successive population estimates could then be derived via an averaging of the initial estimate and the receiver's own estimate, so that the effect of the initial estimate would die out over time. However, as in nonlinear equation solving, use of an initial estimate is only helpful if it accurate and the initial estimate is likely to prove less important than the algorithm which drives the system to steady state. While an improved initial population estimate would result in improved convergence times, were the initial estimate to be way off, convergence might be hindered.

The rate of convergence may be improved by allowing the receiver to use information on the rate at which receiver reports are being sent. For example, due to bandwidth limitations, receivers on higher bandwidth connections have greater knowledge of the overall receiver population. Thus if a receiver were to note a receiver report from a system with high bandwidth availability, that report could be weighed more heavily in determining the overall population estimate.

It is also possible to incorporate the overall receiver reporting rate into the reporting interval calculation. For example, if my current population estimate results in a receiver reporting interval of 3 minutes, and I am receiving 200 receiver reports/minute, it may be desirable to incorporate the rate of receiver reports into my calculations, increasing the reporting interval accordingly.

However, the utility of taking rate information into account is limited in the case of dialup clients, since they will only be able to receive a modest number of receiver reports/minute, and thus the rate at which receiver reports are flowing in may not be reflective of the overall receiver population.

While improved convergence algorithms can result in marked improvements in convergence times, and do result in more accurate population estimates, they do not result in a lower number of forwarding table entries or senders to the RTCP receiver report group. They also do not influence the steady state bandwidth usage for RTCP reporting.

#### [4.3.](#) Use of separate multicast groups for RTP and RTCP

Use of separate multicast groups for RTP and RTCP results in improved scalability for multicasting routing protocols such as sparse mode PIM, since the RTCP group can be placed on the shared tree due to its low sending rate. While this would markedly decrease the number of forwarding cache entries, the router will nevertheless expend considerable resources in adding and flushing forwarding cache entries, and setting up the individual hosts on the shared tree. For this reason,

we do not believe that moving RTCP to an adjacent group will be sufficient for very large conferences. Note also that using an adjacent group for RTCP does not improve the scalability for DVMRP, the MBONE multicast routing protocol, since this protocol has no notion of a shared tree.

INTERNET-DRAFT

28 January 1997

Use of separate multicast groups for RTP and RTCP does not affect convergence times or affect reporting accuracy. It also does not influence the steady state bandwidth usage for RTCP reporting.

#### [4.4.](#) Use of separate multicast groups for receiver and sender reports

Currently RTCP sender and receiver reports are sent to the same multicast group, and both RTP senders and receivers join this group. Were sender and receiver reports to be multicast on different groups, RTP receivers could be allowed to only join the sender report group, thus allowing them to avoid listening to receiver report traffic. RTP senders would still join both groups in order to receive feedback on listenership and transmission quality, and would need to provide information on the estimated receiver population within the sender reports.

While this proposal would lessen the problem for receivers, it would not improve things for senders, and might even make them worse. Since receivers would not longer be able to implement reconsideration as effectively, it would be likely that the senders would be deluged with initial receiver reports. Use of separate groups for receiver and sender reports would not result in a lower number of senders on the RTCP receiver report group, although it would decrease the number of forwarding table entries in protocol such as sparse mode PIM. This proposal would eliminate the bandwidth used by receivers for incoming RTCP reports, but would not address the problem with upstream bandwidth usage in asymmetric networks.

#### [4.5.](#) Unicasting of receiver reports

Instead of multicasting receiver reports, it is possible to unicast them back to the senders. This would allow RTP listeners to avoid receiver report traffic, while RTP senders would still receive feed-

back on listenership and transmission quality. In order to control the receiver report transmission rate, senders would need to provide information on the estimated receiver population within sender reports.

While this proposal would lessen the problem for receivers, as with the previous proposal, it might make things worse for senders, since receivers would no longer be able to reconsider as effectively. However, it would eliminate multicasting of RTCP receiver reports, which will be of benefit to overtaxed routers. This proposal would eliminate the bandwidth used by receivers for incoming RTCP reports, but would not address the problem with upstream bandwidth usage in asymmetric networks.

#### [4.6.](#) Selective reporting

In selective reporting, RTCP receiver reports are only sent by selected hosts. This method results in a reduction in the number of RTCP senders, with attendant reduction in the forwarding tables. It also is likely to result in lower convergence times. Assuming that the statistical sampling was adequate, the accuracy and timeliness of reporting need not be adversely affected. However, this method does not affect the steady state bandwidth allocated to RTCP.

The selection can occur via a polling or random selection mechanism, or it can occur by self-selection. Generally, random selection is preferred since self-selection brings with it the possibility of reporting implosions. For example, if receiver reports were only generated when packet loss exceeded a given threshold, then congestion and attendant packet loss would result in a large number of reports, making the situation worse.

Polling or random selection methods, while they show considerable promise, need to address security concerns. For example, if it is possible to get receivers to respond via a polling message, then that message should be authenticated to prevent leveraged denial of service attacks.



#### 4.7. Use of RTCP BYE messages

Given that receiver reporting intervals will tend to be very long for large conferences, it can be argued that absent an emergency, it makes sense to provide information on listenership, reception quality and bandwidth availability within the RTCP Bye message. Thus on leaving the conference, the receiver would send a message providing information on the time period in which they joined the conference, the overall reception quality and other information. Conventional receiver reports would then either not be sent at all, or would be sent with a TTL of 1. While such an approach would lessen the congestion problem at the beginning of a session, it would increase the size of the RTCP BYE message, resulting in a spike in RTCP bandwidth consumption at the end of a session. Given that no receiver population estimate would be available, it appears that this approach could actually worsen convergence times, unless it were combined with some kind of summarization mechanism. It would also not reduce the number of RTCP senders, or reduce the steady state RTCP bandwidth fraction.

#### 4.8. Use of TTL scoping and summary messages

In this approach, RTCP receiver reports would be sent with a TTL of 1, and a designated summarizer would be elected to provide summary reports with a larger TTL. This approach has the advantage of increasing the leverage of those RTCP receiver reports that are sent, and is likely to be particularly effective for conferences in which membership is densely distributed. However, in sparsely distributed conferences, the effect of summarization will be small unless multiple

levels of summarization are used. The designated summarizer would not necessarily be a router; it could also be another receiver, although this brings up the possibility of how a new summarizer could be elected if the current summarizer leaves the conference.

Since in this scheme receiver reports are not forwarded, non-summarizing RTP receivers should use the population estimate gleaned from local scope reports to calculate their reporting interval. Summarizers and RTP senders will then use global estimates gleaned from global scope summary reports to calculate their reporting interval. A recom-

mended bandwidth allocation for reporting is 25 percent for sender reports, 25 percent for summary reports, and 50 percent for receiver reports.

Since this proposal decreases the scope of RTCP announcements, it would substantially reduce the number of RTCP senders visible to the multicast backbone, and would decrease convergence times, since those messages that were sent would include more information on the receiver population. This proposal would also address the problems of interconnecting multicast routing domains, since the multicast area border router would be able to summarize RTCP behavior within its domain, rather than passing along all RTCP reports. However, it would also require substantial modifications in RTP client behavior.

#### 4.8.1. Summary reports

Via the use of summary reports, privacy can be provided while simultaneously providing senders with improved listenership and receiver quality reporting. This is possible because in the existing implementation much of the information gained from receiver reports is redundant. For example, if congestion results in packets being dropped on a particular link, then all receivers downstream from that link will report the problem. Redundant reporting obscures the information of greatest interest, which is information on global listenership and reception quality. Via introduction of summary reports, it is possible to provide more accurate and timely reporting on listenership and reception quality, as well as to address issues involved in connecting multicast routing domains.

Information useful in summary reports would include:

- Total number of systems being summarized
- Packets received and lost
- Histogram of reception quality (fraction lost)
- Histogram of receiver bandwidth
- Information on users registering

The total systems summarized number is used in order to provide for faster convergence times. Information on packets received and lost will typically be used by network engineers looking to diagnose problems with the MBONE. The histograms of reception quality and receiver bandwidth are propagated in order to allow senders to deduce information relating to the global user experience. In order to allow for

location of individuals participating in a conference, the summarizer may wish to relay information on the users in the conference. Alternatively, it may register users in a directory service via use of the LDAP extensions defined in [4].

#### 4.9. Use of RTP translators

Through the use of RTP translators, it is possible to divide RTP receivers into areas in much the same way as is accomplished by OSPF. Through the use of summary receiver reports, information on listenership and receiver report quality can be propagated between areas while reducing RTCP bandwidth usage and the RTCP sending population on the MBONE.

In order to insulate receivers within an area from external receiver reports, the RTP translator must filter external receiver reports, while allowing external summary reports and sender reports to pass through. Similarly, the RTP translator will listen to receiver reports generated from within its area, but will not pass them on to external areas. Instead, it would issue to external areas two types of summary reports. The first will be based on the packets it receives and will be identical to a conventional receiver report, except for the use of the summary report type; the second will be a true summary report based on area receiver reports. It is useful to allow receiver reports from RTP translators to pass through so as to allow diagnosis of internal distribution problems. The RTP translator will allow internal sender reports to pass through unmodified.

The introduction of RTP translators has several advantages:

- Improved convergence of the receiver population estimate
- Decreased RTCP bandwidth
- Improved privacy
- Listenership and reception quality information available to senders

While most of the above advantages are also available in the receiver summary approach, one advantage of the translator approach is that it provides for greater control by the network administrator. For example, since summary reports would be sent by RTP translators rather than by client summarizers, it would be possible for administrators to control the propagation of user information on a site-by-site basis, rather than on a per-session basis. For example, rather than having this sent in the summary report, it could be stored as a temporary attribute in the area directory server. This may be perceived as a substantial advantage by corporations looking to control access to directory information. For those cases where it is desirable to allow

external access to area registration information, the IP address of the area directory server may be advertised in the summary report. This allows senders with appropriate privileges to retrieve conference registration data from the area directory server via LDAP.

The RTP translator approach also has several major disadvantages. These include:

Aboba

[Page 9]

---

INTERNET-DRAFT

28 January 1997

Requires modifications to routers

Increases loading on routers that now must function as RTP translator

## [5.](#) Acknowledgements

Thanks to Steve Casner of Precept, Mark Handley of ISI, Bob Hinden of Ipsilon and Thomas Pfenning, Peter Ford and Stefan Solomon of Microsoft for useful discussions of this problem space.

## [6.](#) References

[1] H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson. "RTP: A Transport Protocol for Real-Time Applications." [RFC 1889](#), GMD Fokus, January 1996.

[2] H. Schulzrinne. "RTP Profile for Audio and Video Conferences with Minimal Control." [RFC 1890](#), GMD Fokus, January 1996.

[3] M. F. Speer, S. McCanne. "RTP Usage with Layered Multimedia Streams." [draft-speer-avt-layered-video-01.txt](#), Sun Microsystems, LBNL, June 1996.

[4] Y. Yaacovi, M. Wahl, K. Settle, T. Genovese. "Lightweight Directory Access Protocol: Extensions for Dynamic Directory Services." [draft-ietf-asid-ldapv3ext-02.txt](#), Microsoft, Critical Angle, October, 1996.

## [7.](#) Authors' Addresses

Bernard Aboba

Microsoft Corporation  
One Microsoft Way  
Redmond, WA 98052

Phone: 206-936-6605  
EMail: [bernarda@microsoft.com](mailto:bernarda@microsoft.com)