

MBONE Deployment Working Group  
INTERNET-DRAFT  
Category: Informational  
<draft-ietf-mboned-mdh-05.txt>  
20 November 2000

Dave Thaler  
Microsoft  
Bernard Aboba  
Microsoft

## Multicast Debugging Handbook

### 1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

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."

The list of current Internet-Drafts can be accessed at  
<http://www.ietf.org/ietf/lid-abstracts.txt>.

To view the list Internet-Draft Shadow Directories, see  
<http://www.ietf.org/shadow.html>.

### 2. Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

### 3. Abstract

This document serves as a handbook for the debugging of multicast connectivity problems. In addition to reviewing commonly encountered problems, the draft summarizes publicly distributable multicast diagnostic tools, and provides examples of their use, along with pointers to source and binary distributions.

### 4. Introduction

In order to deploy multicast on a large scale, it is necessary for Network Operations Centers, support personnel and customers to be able to diagnose problems. Over the years a number of tools have been developed that can assist in the diagnostic process. This document serves as a handbook for the debugging of multicast connectivity problems. In addition to reviewing commonly encountered problems, the draft summarizes publicly distributable multicast diagnostic tools, and

provides examples of their use, along with pointers to source and binary distributions.

## 5. Usage scenarios

Multicast diagnostic tools are typically employed in one of the following settings:

Customer service or support	SDR mtrace RTPmon Dr. Watson
Network or system administrator	SDR mrinfo netstat mconfig mstat mtrace RTPmon tcpdump Dr. Watson Duppkts mrouted.dump, mrouted.cache
Network Operations Center	SDR mrtree map-mbone MVIEW Multicast heartbeat mwatch and mcollect asn asname

### 5.1. Customer service and support

ISPs offering multicast services are likely to encounter the following classes of customer questions:

Session announcement problems

Reception problems  
Multicast router problems

Below we discuss how each of these types of problems may be diagnosed.

#### 5.1.1. Session announcement problems

Session announcement questions are those which relate to the user's session announcement software. Sample complaints include:

No conferences were visible in the session announcement tool  
Conference X was not visible in the session announcement tool  
I can see conferences when dialed into POPA, but not POPB  
I can receive conferences listed in SDR, but sometimes when I join conferences via a Web site, I cannot receive them.

Session announcement questions are typically investigated via the following procedure:

1. **If only a specific session announcement is missing, check the session announcement from somewhere where it is being received, and find the group(s) and ports that the session utilizes, as well as the source IP addresses. If the problem is with all session announcements, find the information on any current session announcement which should be seen by the user.**
2. **Find out the user's IP address, if known, and the POP dialed into or router connected to. One way to determine the user's router given their IP address is to mtrace or traceroute to their address.**
3. **Do an mtrace between the session announcement's origin and the receiver on the sap.mcast.net group. If the mtrace succeeds, note any hops showing loss.**
4. **If the mtrace never gets past the receiver itself, check the NAsEs or routers with mstat -l to see if the relevant group has been joined. If not, the problem is probably with the receiver's host. Ask the user to check with Dr. Watson or a sniffer to see if the router is sending IGMP membership queries, and if the host is responding with membership reports and if so, what versions are being used.**
5. **If the sap.mcast.net group has been joined, but the mtrace failed, the problem is likely a multicast routing problem (see section 4.1.3).**
6. **If the mtrace succeeded, and one hop shows 100% loss, compare the output with the TTL of the session announcement. Users may download session descriptions from Web sites that they may not be in the position to receive, due to site or regional scoping. The loss may also be the**

result of a scoped boundary separating the originator and the receiver, which will also be indicated as such by mtrace.

**7. Otherwise, if the mtrace succeeded, look for hops showing non-negligible loss.** These typically denote points of congestion (see section 4.3.1). Note that if rate-limiting is installed on these congested links, session announcements are often lost since SAP traffic is given lower priority.

**8. If all else fails, request a network sniff from the user, and check whether it shows traffic to sap.mcast.net, and if so, from what sources, and what is being announced.**

#### 5.1.2. Reception problems

Reception questions are those where the user has successfully received the session announcement, but was unable to receive one or more media streams for the session joined. Sample complaints include:

- I joined conference X, but nothing happened
- I joined conference X, got video but no audio
- I joined conference X, and got intermittent audio
- I can't see source X, but source X can see me (or vice versa)

Reception questions are typically investigated via the following procedure:

**1. Check the session announcement, find the group(s) and ports that the session utilizes, as well as the source IP addresses.**

**2. Find out the user's IP address, if known, and the POP dialed into or router connected to.** One way to determine the user's router given their IP address is to mtrace or traceroute to their address.

**3. Check if the user is sending RTCP reports with RTPmon, and if so, what the loss rate is.**

**4. Do an mtrace between the source and the receiver on the relevant group.** If the mtrace succeeds, note any hops showing loss.

**5. If the mtrace never gets past the receiver itself, check the NASes or routers with mstat -l to see if the relevant group has been joined.** If not, the problem is probably with the receiver's host. Check with Dr. Watson to see if the router is sending IGMP membership queries, and if the host is responding with membership reports and if so, what versions are being used.

**6. If the relevant group has been joined, but the mtrace failed, the**

problem is likely a multicast routing problem (see section 4.1.3).

**7. If the mtrace succeeded, and one hop shows 100% loss, compare the output with the TTL of the session announcement.** The user may not be in a position to receive data from the source, due to site or regional scoping. The loss may also be the result of a scoped boundary separating the source and the receiver, which will also be indicated as such by mtrace.

**8. Otherwise, if the mtrace succeeded, look for hops showing non-negligible loss.** These typically denote points of congestion (see section 4.3.1).

**9. If all else fails, request a network sniff from the user, and check whether it shows traffic to the relevant group, and if so, from what sources.**

Other reception complaints include:

When I join my first conference, it works great. But then when I quit that and join another one, it doesn't work anymore.

Why is my modem light is still flashing even after I've quit SDR and VIC?

Such problems are often IGMP-related problems observed by a user connecting to the network using a host which is running a TCP/IP stack implementing IGMP v1. Such users will experience long leave latencies, with resulting poor reception and/or performance of other applications. Such problems can be distinguished from ordinary reception problems in that they typically do not occur for the first session joined, only for subsequent sessions. The solution consists of upgrading the user to an IGMP v2-capable stack. IGMP is described in [2].

IGMP-related questions are typically investigated by the following procedure:

**1. Obtain the vendor and version of the user's TCP/IP stack. Determine whether this stack is IGMP v2-enabled.**

**2. Ask the user to run Dr. Watson or a network sniffer and to indicate whether IGMP queries are being seen, whether responses are being sent, and if so, what version.**

### **5.1.3. Multicast router problems**

Multicast router questions are those which relate to the setup of a multicast router. Sample complaints include:

I can get video and audio when online with ISDN, but not with a modem, or vice versa.

I can't bring up a DVMRP tunnel to my site. Why not?

My router works great. Why can't I get multicast?

Why can't I source multicast?

Multicast routing questions are typically investigated via the following procedure:

- 1. Ask the user what the router vendor is, and what software version they have running.** Attempt to verify this information using `mrinto` or `mstat`.
- 2. Check whether this vendor and version supports multicast routing, and whether an upgrade to a later version is recommended.**
- 3. Ask for a copy of the router configuration file.**
- 4. Check whether the user has NAT enabled; this is incompatible with most multicast routing protocols, and so should be switched off.**
- 5. Find out the user's IP address(es), or if not known, the POP dialed into or router connected to.**
- 6. Check the loss rate and connectivity by doing an `mtrace` from various sources to the user's IP address.**
- 7. Check the user's router with `mstat -l` to see if it has joined any multicast groups, and check upstream routers to see if they are subscribed to any groups.**
- 8. When all else fails, request a network sniff and examine it to determine what multicast routing protocols are being run, if any.**

## **5.2. Network Operations Center**

A Network Operations Center (NOC) will typically receive a complaint after it has been investigated by customer support and determined to be a network-related issue. Although it is desirable for customer support to have performed the diagnostic tests described above, if this has not been done, NOC personnel will need to perform the tests themselves to isolate the cause of the problem. If the proper systems have been installed, in most cases, the NOC will already have been alerted to the problem prior to receiving referrals from customer support. The following diagnostic procedures are recommended:

**1. Regularly generate summaries based on RTCP receiver and sender**

reports, using RTP monitoring tools. Sample reports may include average loss rates experienced during sessions, or users whose loss rates exceed a particular threshold.

**2. Determine the source of the problems by doing mtraces between the sources and the receivers.**

**3. On a network monitoring station, keep track of the functioning of multicast-enabled hardware, either by doing periodic mtraces, or by using a heartbeat monitor to receive SNMP traps from equipment losing the heartbeat.**

**4. In order to keep track of group topologies, use mapping tools such as map-mbone, MVIEW, or mrtree, along with autonomous system mapping tools such as asn and asname.**

**5.3. Network or system administrator**

The NOC will escalate network engineering problems to network engineers and system administrators if their intervention is required. In order to understand the origin of the problem and repair it, it is necessary to diagnose the operations of individual systems and routers using router and system diagnostics such as netstat, mrimfo, mstat, mconfig, RTPmon, and mtrace, as well as network analysis tools such as tcpdump or Dr. Watson.

In smaller installations the network engineer or system administrator often doubles as customer support and network operations guru, in which case problems may be escalated without any triage (our condolences).

Typical classes of problems encountered by network engineers and system administrators include:

- Congestion and rate-limiting problems
- Multicast routing problems

**5.3.1. Congestion and rate limiting problems**

Congestion and rate limiting problems result in high packet loss with subsequent loss of session announcements and decrease in quality of audio and video. These problems may be investigated via the following procedure:

**1. Use RTPmon to check for receivers experiencing packet loss.**

**2. Do an mtrace from the source to the receiver on the relevant group in order to locate the problematic hops.**

3. **Do an mtrace in the opposite direction to help distinguish whether the problem is with the router or the link at that hop.**
4. **If the reverse mtrace shows similar loss at an hop adjacent to the lossy hop in the forward mtrace, odds are that the intermediate router is overloaded. Use mrimfo to check the fanout on the router. Overloaded routers are often the result of having too many tunnels.**
5. **If the reverse mtrace shows no problems near that hop, indicating that loss is one-way, then check the router on the upstream end of the link with mstat -nv to see if it has a rate-limit set on the link, and if the link traffic is near that limit.**
6. **If the reverse mtrace shows loss over the same link, the problem is likely to be link congestion. Use mstat -nv to see how much bandwidth is being used by multicast traffic. (If mstat fails, running an mtrace with the -T option may help to confirm link congestion, although the statistics can be misleading.)**
7. **If a congested link is suspected, but mstat failed, another indicator can be obtained by doing an mtrace from the session announcer to the destination on other groups joined by the receiver, such as the SAP group, and comparing loss statistics.**
8. **Check for unicast packet loss over the link using ping. Multicast (but not unicast) packet loss on a link with a rate limit is an indication that the link's multicast rate limit should be raised or eliminated entirely. Packet loss on a link without rate limiting is an indication of congestion. On such links it may be desirable to add a rate limit. Since DVMRP prunes are currently not retransmitted by most routers, prunes may be lost if no rate limit exists, which may further worsen the congestion problem.**
9. **Use mstat -gR to see whether a single group is using an inordinate amount of the link bandwidth. If so, use mstat to see whether a single source to that group is using an inordinate amount of the link bandwidth. If so, attempt to contact the source (contact information may be available in the session announcement).**

### 5.3.2. Multicast routing problems

Multicast routing problems include:

- Duplicate packets
- Injection of bogus routes (typically into DVMRP)
- Redistribution of unicast routes (via BGP or an IGP) into DVMRP
- Non-pruning routers

Duplicate packets are a symptom of routing loops. This problem may be investigated via the following procedure:

1. Use a program such as Duppkts to detect duplicate packets.
2. Use a network monitor or RTPmon to find the sources and receivers on the group.
3. Do an mtrace from the source(s) to the receivers in order to find the loop. Duplicates will also show up in mtrace output as hops with negative loss.

Bogus route injection problems may be investigated via the following procedure:

1. Dump the DVMRP routing table. The routing table may be examined remotely via mstat using the -r options, or locally (for mrouterd) by sending the USR1 signal to mrouterd, generating the /var/tmp/mrouterd.dump file.
2. Check the table for bogus routes (known as "martians"). Bogus routes include addresses reserved for use by private internets, as described in [9]. These routes include 10/8, 172.16/12, or 192.168/16, or more specific routes within these regions. Injecting a default route into the DVMRP backbone is also considered to be a bogus route.
3. Locate the origin of the bogus routes by doing an mtrace to an IP address in the bogus range.

Symptoms of unicast route redistribution are injection of a large number of unicast routes (25K+) into DVMRP. The problem may be investigated via the following procedure:

1. Examine the routing table. The DVMRP routing table may be examined remotely via mstat -r, or locally (for mrouterd) by sending the USR1 signal to mrouterd, generating the /var/tmp/mrouterd.dump file. For protocol independent multicast routing protocols (such as Sparse-Mode PIM), examine the unicast routing table.
2. Check if a single site is the predominant route injector. This site is likely to be the problem. One way to check this is to mtrace to addresses in a number of "suspect" prefixes.
3. If your router supports it, set a route limit on the DVMRP tunnel interface. A limit of 7000 routes is currently recommended. You may also wish to set "route-hog notification" at 5000 routes.

Non-pruning DVMRP routers are those which maintain groups in the

multicast routing table although there are no downstream subscribers. The problem can be solved via the following procedure:

- 1. Check the router version number using mstat or mrinfo.** Non-pruning routers include mrouterd versions prior to 3, Cisco Systems IOS prior to version 11.0(3), and Bay Networks implementations prior to 9.0.
- 2. Confirm lack of pruning as follows.** First, dump the multicast forwarding table. This can be done remotely with mstat -N, or locally (for mrouterd) by sending theUSR2 signal to mrouterd, generating the /var/tmp/mrouterd.cache file.
- 3. Check the forwarding table to see if an interface is in the outgoing interface list for every entry in the multicast forwarding table.** If so, it is likely that a non-pruner lies downstream in that direction.
- 4. You can confirm the existence of a non-pruner by creating a temporary, unadvertised, session and sending (preferably with a low data rate) data to that group.** After a few moments, dump the forwarding table again. If any interfaces appear in the outgoing interface list of the entry for your test stream, then non-pruners lie in those directions.
- 5. If a non-pruner exists downstream, use mrtree to follow the path of the data downstream to the non-pruning router(s).**
- 6. If your router supports it, enable the reject non-pruners option.** If not, and the unpruned interface is a tunnel, consider disabling the tunnel.

## 6. Appendix - Multicast Diagnostic Tools

### 6.1. Types of tools

Multicast diagnostic tools typically fall into the following categories:

RTP monitoring tools	RTPmon Msessmon RTPquality RTPdump RTPcast/RTPlisten Duppkts
Multicast router diagnostics	mrinfo netstat mconfig mstat mrouted.dump, mrouted.cache
Multicast traceroute	mtrace
MBONE mapping tools	mrtree map-mbone asn asname mcollect & mwatch
NOC tools	MVIEW Multicast heartbeat
Network analysis tools	tcpdump Dr. Watson

RTP monitoring tools are used to monitor the transmission quality and

popularity of individual sessions. Multicast router diagnostics are used to obtain information on the configuration and state of multicast routers. MBONE mapping tools are used to map out the topology for a particular group. These tools can show the topology at the level of individual systems, or at the level of autonomous system connections. Multicast traceroute tools are used to trace the path between a source and destination. Network Operations Center tools are used to monitor the state of network devices within an autonomous system. Network analysis tools (such as tcpdump and Dr. Watson) are used to analyze traffic on the network.

## 6.2. Facilities utilized

Multicast diagnostic tools typically rely on one or more of the following facilities:

RTCP source and receiver reports	RTPmon Msessmon RTPquality RTPdump RTPcast/RTPlisten Duppkts
SNMP MIBs	multicast heartbeat mconfig mstat MVIEW mrtree
IGMP trace facility	mtrace
IGMP ASK_NEIGHBORS message (DVMRP)	mrinfo mrtree map-mbone
Routing arbiter and WHOIS Databases	asn asname
Internal structures	tcpdump netstat mrouted.dump, mrouted.cache

Tools using RTCP reports analyze RTCP source and receiver reports, providing information on packet loss, inter-arrival jitter, bandwidth

availability, or listenership. These tools therefore will only work with RTP implementations which support RTCP reporting. Tools using SNMP MIBs perform queries for variables in the IGMP, multicast routing, DVMRP, and PIM MIBs. As a result, these tools depend on implementation of the relevant SNMP MIBs in the network devices that are being monitored. Tools based IGMP ASK\_NEIGHBORS messages use these messages to map portions of the MBONE, and thus will only work with routers implementing DVMRP. Tools based on IGMP tracing perform a multicast traceroute. These tools are typically only useful in cases where multicast routers along the path have a route back to the source.

Diagnostic tools may use more than one of these facilities. For example, mstat makes use of SNMP MIBs, as well as the IGMP ASK\_NEIGHBORS facility.

### 6.3. RTP monitoring tools

This class of tools provides information required to monitor the performance of RTP-based applications.

#### 6.3.1. RTPmon

##### Authors

David Bacher, Andrew Swan, and Lawrence A. Rowe  
{drbacher,aswan,rowe}@cs.berkeley.edu  
Computer Science Division - EECS  
University of California  
Berkeley, CA 94720-1776

##### Description

RTPmon allows network administrators or support personnel to monitor listenership as well as session quality experienced by subscribers. The tool also facilitates tracing the cause of problems resulting in quality degradation. To accomplish this task, RTPmon summarizes and analyzes information provided by RTCP source and receiver reports.

Receivers are displayed for a given sender in the form of a spreadsheet, with cells being filled in with metrics such as packet loss rate or jitter. Clicking on a cell displays a stripchart of statistics on packet loss rate, smoothed packet loss rate and jitter. From the stripchart it is possible to launch an mtrace between the sender and the receiver, a convenient way of diagnosing network problems along the multicast distribution path. Clicking on a receiver or sender displays summary information.

For groups with large memberships, the display may be limited to

members surpassing a given threshold in packet loss rate or jitter. Using RTPmon it is possible to sort receivers for a given sender according to maximum or average loss.

Further information is available in the RTPmon man page.

#### Example

For examples and further information, see the rtpmon man page, or:  
<http://bmr.c.berkeley.edu/~drbacher/projects/mm96-demo/>

#### Facilities used

RTCP source and receiver reports  
IGMP multicast trace (if installed)

#### Availability

RTPmon is available for UNIX and may be obtained from:  
<ftp://mm-ftp.cs.berkeley.edu/pub/rtpmon/>

Bug reports and suggestions should be sent to:  
[rtpmon@bmr.c.berkeley.edu](mailto:rtpmon@bmr.c.berkeley.edu).

### **6.3.2. RTPcast/RTPlisten, RTPquality, Duppkts, RTPdump, RTPtools, Mssesmon, Mpoll**

#### Author

Mpoll: Andrew Patrick <[andrew@calvin.dgbt.doc.ca](mailto:andrew@calvin.dgbt.doc.ca)>

#### Description

RTPcast listens to RTCP receiver reports and forwards data to another multicast group; RTPlisten then listens to that group. RTPdump listens for, and dumps RTP and RTCP packets. Duppkts listens on a multicast group and port, and reports the number of packets received and lost, as well as the number of duplicates. RTPquality listens to RTCP receiver reports and writes data on packet loss, as well as late and non-sequenced packets. RTPtools allows recording and playback of RTP sessions. Mssesmon provides a routemap of participants in RTP conferences as well as stripcharts of statistics on RTP packet loss and jitter. Mpoll is a survey collection tool that can be used to collect quality ratings during multicast sessions.

#### Example

Information on these tools is available from:

<http://sauce.mmlab.uninett.no/mice-nsc/tools.html>

#### Facilities used

RTCP source and receiver reports

#### Availability

Binaries for RTPcast/RTPlisten are available from:  
<ftp://sauce.uio.no/mice-nsc/util/rtp>

Source code for RTPquality is available from:  
<ftp://sauce.uio.no/mice-nsc/util/rtp/rtpqual.c>

Source code for RTPdump is available at:  
<ftp://sauce.uio.no/mice-nsc/util/rtpdump-1.0.tar.gz>

Source code for RTPtools is available at:  
<ftp://sauce.uio.no/mice-nsc/util/rtptools/rtptools-1.9.tar.gz>

Source and binaries for Msessmon is available at:  
<ftp://sauce.uio.no/mice-nsc/util/m sessmon/>

Source and binaries for Mpoll is available at:  
<ftp://sauce.uio.no/mice-nsc/util/mpoll/>

### 6.4. Multicast router diagnostics

This class of tools facilitates monitoring and management of multicast routers.

#### 6.4.1. `mrouted.dump`, `mrouted.cache`

##### Author

Bill Fenner, [fenner@research.att.com](mailto:fenner@research.att.com)

##### Description

Sending the USR1 signal to `mrouted` dumps the internal routing table to `/var/tmp/mrouted.dump`; sending the USR2 signal dumps the forwarding cache to `/var/tmp/mrouted.cache`.

Further information on `mrouted` and the `mrouted.dump` and `mrouted.cache` file formats is available in the `mrouted` man page.

##### Example

INTERNET-DRAFT

Multicast Debugging Handbook

20 November 2000

```
% cat mrouted.dump
vifs_with_neighbors = 2
```

## Virtual Interface Table

Vif	Name	Local-Address		M	Thr	Rate	Flags
0	ed0	128.31.107.1	subnet: 128.31.107/24	1	1	0	querier
			peers: 128.31.107.249 (3.8) (0xe)				
			groups: 239.109.100.200				
			224.0.0.2				
			224.0.0.4				
			pkts in : 4075				
			pkts out: 0				
1	ed0	128.31.107.1	tunnel: 204.67.107.11	1	32	500	
			peers: 204.67.107.11 (11.2) (0x1a)				
			pkts in : 0				
			pkts out: 2359				

## Multicast Routing Table (3801 entries)

Origin-Subnet	From-Gateway	Metric	Tmr	In-Vif	Out-Vifs
207.10.165.51/32	128.31.107.249	10	20	0	1
207.10.165.50/32	128.31.107.249	10	20	0	1
206.172.195.32/32	128.31.107.249	9	20	0	1
172/8	128.31.107.249	10	20	0	1
...					

```
% cat mrouted.cache
```

## Multicast Routing Cache Table (198 entries)

Origin	Mcast-group	CTmr	Age	Ptmr	IVif	Forwvifs
131.107.2.139/32	224.0.12.0	58s	7m	-	-1	
>131.107.2.139						
143.107.103.0/27	224.0.1.1	3m	2m	3m	0P	
>143.107.103.5						
128.232/16	224.0.1.1	4m	7m	4m	0P	
>128.232.2.209						
157.161/16	224.0.1.1	67s	6m	-	0	1
>157.161.114.2						
206.152.163/24	224.0.1.15	74s	7m	-	0	1
>206.152.163.21						
4.0.0.34/32	224.0.1.32	56s	4m	25s	0P	1p
>4.0.0.34						
137.39.2.254/32	224.0.1.32	3m	5m	-	0	1
>137.39.2.254						
137.39.43.32/30	224.0.1.32	38s	5m	-	0	1
>137.39.43.34						
...						

INTERNET-DRAFT

Multicast Debugging Handbook

20 November 2000

## Facilities used

Internal facilities (forwarding cache and routing table)

## Availability

The SNMP-capable mrouted distribution is available at:  
ftp://ftp.merit.edu/net-research/mbone/mirrors/mrouted/

### 6.4.2. mrinfo

#### Author

Van Jacobson, van@ee.lbl.gov

#### Description

mrinfo displays information about a multicast router; to do this, it uses the IGMP ASK\_NEIGHBORS message to discover the router's physical and virtual interfaces. Routers are also queried for their version number, and if this query is successful, for their metrics, thresholds, and flags. Results are printed in an indented list format similar to that for map-mbone.

#### Example

```
% mrinfo 192.80.214.199
192.80.214.199 (collegepk-mbonel.bbnplanet.net) [version 11.2,prune,mtrace,snmp]:
 128.167.252.196 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
 192.80.214.199 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
 192.41.177.196 -> 0.0.0.0 (local) [1/0/pim/querier/down/leaf]
 128.167.252.196 -> 128.167.254.165 (devo.sura.net) [1/32/tunnel/querier/down/leaf]
 128.167.252.196 -> 131.119.0.197 (paloalto-mbonel.bbnplanet.net)
    [1/64/tunnel/pim/querier]
 128.167.252.196 -> 199.94.207.2 (cambridgel-mbonel.bbnplanet.net)
    [1/32/tunnel/pim/querier]
 128.167.252.196 -> 137.39.43.34 (MBONE1.UU.NET) [1/32/tunnel/querier]
 128.167.252.196 -> 192.41.177.199 (wtm-ms2.bbnplanet.net) [1/16/tunnel/querier]
 128.167.252.196 -> 128.244.93.3 (sage.jhuapl.edu) [1/32/tunnel/querier]
 128.167.252.196 -> 192.221.34.22 (cdrn.bbnplanet.net) [1/32/tunnel/querier]
 128.167.252.196 -> 128.167.1.197 (cpk-ms1.ser.bbnplanet.com) [1/16/tunnel/querier]
 128.167.252.196 -> 134.205.93.150 (dilbert.sam.pentagon.mil) [1/32/tunnel/querier]
 128.167.252.196 -> 192.221.48.234 (atlanta3-mbonel.bbnplanet.net)
    [1/64/tunnel/pim/querier]
 128.167.252.196 -> 204.167.201.38 (dallas2-mbonel.bbnplanet.net)
    [1/64/tunnel/pim/querier]
 128.167.252.196 -> 205.130.85.3 (philipii.nap.edu) [1/32/tunnel/querier/down/leaf]
 128.167.252.196 -> 128.175.13.36 (pfet.nss.udel.edu) [1/32/tunnel/querier/down/leaf]
```

INTERNET-DRAFT

Multicast Debugging Handbook

20 November 2000

```

128.167.252.196 -> 192.41.177.197 (wtn-ms1.bbnplanet.net) [1/32/tunnel/querier]
128.167.252.196 -> 204.148.62.28 (mbone-e.ans.net) [1/32/tunnel/querier]
128.167.252.196 -> 205.128.246.2 (usnrctc.bbnplanet.net) [1/32/tunnel/pim/querier]

```

#### Facilities used

IGMP ASK\_NEIGHBORS message (DVMRP)

#### Availability

mrinfo is available for UNIX and is included in the SNMP-capable mouted distribution, available at:  
<ftp://ftp.merit.edu/net-research/mbone/mirrors/mouted/>

mrinfo is also available in the MVIEW distribution, available at:  
<ftp://ftp.merit.edu/net-research/mbone/mview/>

### 6.4.3. netstat

#### Author

Unknown

#### Description

On multicast-enabled systems, netstat is typically extended so as to provide information on virtual interfaces and the multicast forwarding cache (-g option), as well as multicast routing statistics (-gs option), and igmp behavior (-s option).

#### Example

```
%netstat -g
```

##### Virtual Interface Table

Vif	Thresh	Rate	Local-Address	Remote-Address	Pkts-In	Pkts-Out
0	1	0	128.15.2.120		16323	385
1	32	512	128.15.2.120	202.34.126.2	2	0

##### Multicast Forwarding Cache

Origin	Group	Packets	In-Vif	Out-Vifs:Ttls
128.15.2.120	224.2.195.166	281	0	
128.15.1.110	239.100.101.223	1660	0	
128.15.1.135	238.27.27.1	1660	0	
128.15.1.110	239.111.111.235	1660	0	
...				

```
%netstat -gs
```

## multicast forwarding:

```
182880 multicast forwarding cache lookups
8237 multicast forwarding cache misses
6736 upcalls to mrouterd
193 upcall queue overflows
5567 upcalls dropped due to full socket buffer
177 cache cleanups
7234 datagrams with no route for origin
    0 datagrams arrived with bad tunneling
    0 datagrams could not be tunneled
954 datagrams arrived on wrong interface
    0 datagrams selectively dropped
    0 datagrams dropped due to queue overflow
    0 datagrams dropped for being too large
```

```
%netstat -s
```

```
ip:
```

```
3807182 total packets received
0 bad header checksums
```

```
...
```

```
icmp:
```

```
40 calls to icmp_error
0 errors not generated 'cuz old message was icmp
```

```
...
```

```
igmp:
```

```
18504 messages received
0 messages received with too few bytes
48 messages received with bad checksum
2478 membership queries received
0 membership queries received with invalid field(s)
194 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
8510 membership reports sent
```

```
tcp:
```

```
10705 packets sent
    5536 data packets (1532081 bytes)
```

```
...
```

```
udp:
```

```
3104045 datagrams received
0 with incomplete header
```

```
...
```

## Facilities used

```
Netstat accesses system internal data structures in order to carry
out its function.
```

## Availability

netstat is included with a variety of operating systems, including UNIX, OS/2, and Windows. For further information, please consult the netstat man page or documentation.

### 6.4.4. mstat

#### Author

Dave Thaler, dthaler@microsoft.com

#### Description

mstat is a general purpose tool for obtaining router configuration and status information. In order to perform its task, mstat utilizes SNMP MIBs (such as the IGMP, multicast routing, PIM, and DVMRP MIBs), as well as ASK\_NEIGHBORS IGMP messages. mstat displays the contents of various MBONE-related data structures in various formats, depending on the options selected. Options include:

- G Show the PIM group table
- I Show the PIM interface table.
- K Show the cached IP multicast route table; works for all SNMP-capable routers.
- N Show the IP Multicast Next Hop Table.
- P Show the PIM neighbor table.
- a Show the alternate subnet table.
- b Show the scoped boundary table.
- d Show the DVMRP neighbor table.
- g Show the Group Summary table.
- i Show the DVMRP interface table; similar to an mrimf report.
- l Show the IGMP local group table.
- r Show the DVMRP routing table; similar to a portion of the mrouterd.dump file.
- t Show the DVMRP routing next hop table; similar to another portion of the mrouterd.dump file.
- v Show statistics corresponding to the DVMRP interface table.

#### Examples

```
% mstat
IP Multicast Route Table for bigco.com
Mcast-group      Origin-Subnet      InIf  UpTime  Tmr    Pkts      Bytes  RpF  Proto
NTP.MCAST.NET    0.0.0.0/32         0     245341  179    0         0      0    pim
NTP.MCAST.NET    128.232.0.49/32   7     206403  418    3056     293376  17   dvmp
NTP.MCAST.NET    128.232.2.209/32  7     206403  417    3027     290592  19   dvmp
```

INTERNET-DRAFT

Multicast Debugging Handbook

20 November 2000

NTP.MCAST.NET	143.107.103.5/32	7	592	218	3	228	3	dvmp
NTP.MCAST.NET	157.161.114.2/32	7	27703	517	411	31236	11	dvmp
IETF-2-VIDEO.MC	0.0.0.0/32	0	245349	175	0	0	0	pim
IETF-2-VIDEO.MC	206.152.163.21/32	7	242567	244	46887	4149336	3388	dvmp
MTRACE.MCAST.NE	0.0.0.0/32	0	1690	177	0	0	0	pim
MTRACE.MCAST.NE	194.104.0.25/32	7	405	483	2	792	0	dvmp
MTRACE.MCAST.NE	206.54.224.150/32	7	456	569	4	1072	4	dvmp
CISCO-RP-DISCOV	0.0.0.0/32	0	245534	0	0	0	0	pim
224.0.14.1	203.15.123.99/32	4	17	161	0	0	0	dvmp
224.0.92.3	171.68.201.39/32	4	174	4	0	0	0	dvmp
224.2.0.1	13.2.116.11/32	4	150	26	0	0	0	dvmp
224.2.0.1	128.32.38.218/32	4	147	30	0	0	0	dvmp
224.2.2.1	205.226.8.183/32	4	146	30	0	0	0	dvmp
224.2.20.165	13.2.116.11/32	4	55	119	0	0	0	dvmp
224.2.100.100	13.2.116.11/32	4	87	91	0	0	0	dvmp
SAP.MCAST.NET	164.67.63.7/32	4	114	64	1	855	0	dvmp
SAP.MCAST.NET	193.61.212.130/32	4	153	23	1	868	0	dvmp
SAP.MCAST.NET	199.94.220.184/32	4	26	152	1	416	0	dvmp
SAP.MCAST.NET	206.154.213.242/32	4	156	19	1	360	0	dvmp

...

Examples of the many other options are provided in the mstat man pages.

#### Facilities used

PIM, DVMP, IGMP, and multicast routing MIBs  
 IGMP ASK\_NEIGHBORS message (DVMP)

#### Availability

mstat is included in the SNMP-capable mouted distribution,  
 available at:  
<ftp://ftp.merit.edu/net-research/mbone/mirrors/mouted/>

mstat is also available in the MVIEW distribution, available at:  
<ftp://ftp.merit.edu/net-research/mbone/mview/>

#### 6.4.5. mconfig

##### Author

Dave Thaler, [dthaler@microsoft.com](mailto:dthaler@microsoft.com)

##### Description

mconfig allows the user to display and (if the community string is known) to modify the configuration of a multicast router implementing the DVMP MIB.

## Example

For more information on mconfig, please see the man page.

## Facilities used

DVMRP MIB

## Availability

mconfig is available for UNIX and is included in the SNMP-capable mrouted distribution, available at: <ftp://ftp.merit.edu/net-research/mbone/mirrors/mrouted/>

## 6.5. Multicast traceroute

### 6.5.1. mtrace

#### Author

Bill Fenner, [fenner@research.att.com](mailto:fenner@research.att.com)

#### Description

mtrace provides a facility by which to trace the path between a sender and a receiver of a particular group. This is particularly useful when used alongside a facility such as RTPmon, which allows you to identify problem source-receiver pairs.

Note that the utility of mtrace is often limited by the multicast topology. Where multicast and unicast topologies are not aligned (as is the case in many multicast-enabled networks) mtrace may not function.

For information on the details of the protocol, see reference [8].

#### Example

```
% mtrace 131.243.73.36 128.15.1.250 224.2.195.166
Mtrace from 131.243.73.36 to 128.15.1.250 via group 224.2.195.166
Querying full reverse path... * switching to hop-by-hop:
 0 bigman.bigco.com (128.15.1.250)
-1 * * littleman.bigco.com (128.15.1.249) DVMRP thresh^ 1
-2 * * * seamr1-gw.nwnet.net (192.35.180.201) DVMRP thresh^ 32
-3 * * seamr2-gw.nwnet.net (192.220.238.130) DVMRP thresh^ 0
-4 * * mcast.cac.washington.edu (140.142.116.1) DVMRP thresh^ 32
-5 * * * * dec3800-1-fddi-0.Sacramento.mci.net (204.70.164.29) didn't respond
```

INTERNET-DRAFT

Multicast Debugging Handbook

20 November 2000

-6 \* \* \*  
-7 \* \*

Resuming...

```
-5 dec3800-1-fddi-0.Sacramento.mci.net (204.70.164.29) DVMRP thresh^ 64
-6 dec3800-2-fddi-0.SanFrancisco.mci.net (204.70.158.61) DVMRP thresh^ 1
-7 mbone.nsi.nasa.gov (192.203.230.241) DVMRP thresh^ 64
-8 * * llnl-mr2.es.net (134.55.12.229) DVMRP thresh^ 64
-9 * * lbl-mr1.es.net (134.55.12.101) DVMRP thresh^ 8
-10 * * mr1.lbl.gov (131.243.64.184) DVMRP thresh^ 32
-11 * * ir40gw.lbl.gov (131.243.64.1) DVMRP thresh^ 0
-12 * * irals.lbl.gov (131.243.128.6) PIM thresh^ 0
-13 bl7-36.als.lbl.gov (131.243.73.36)
Round trip time 74 ms; total ttl of 72 required.
```

Waiting to accumulate statistics... Results after 10 seconds:

Source	Response Dest	Overall	Packet Statistics For Traffic From		
131.243.73.36	128.15.1.250	Packet	131.243.73.36 To 224.2.195.166		
v	___/ rtt 77 ms	Rate	Lost/Sent = Pct Rate		
131.243.73.1					
131.243.128.6	irals.lbl.gov				
v	^ ttl 1	6 pps	0/60	= 0%	6 pps
131.243.128.40					
131.243.64.1	ir40gw.lbl.gov				
v	^ ttl 2	13 pps	0/60	= 0%	6 pps
131.243.64.184	mr1.lbl.gov				
v	^ ttl 35	9 pps	0/60	= 0%	6 pps
198.128.16.13					
134.55.12.101	lbl-mr1.es.net				
v	^ ttl 35	0 pps	0/60	= 0%	0 pps
134.55.12.229	llnl-mr2.es.net				
v	^ ttl 69	0 pps	1/60	= 2%	0 pps
192.203.230.241	mbone.nsi.nasa.gov				
v	^ ttl 70	0 pps	0/59	= 0%	0 pps
204.70.158.61	dec3800-2-fddi-0.SanFrancisco.mci.net				
v	^ ttl 70	0 pps	0/59	= 0%	0 pps
204.70.164.29	dec3800-1-fddi-0.Sacramento.mci.net				
v	^ ttl 72	0 pps	0/59	= 0%	0 pps
140.142.116.1	mcast.cac.washington.edu				
v	^ ttl 72	0 pps	0/59	= 0%	0 pps
192.220.249.66					
192.220.238.130	seamr2-gw.nwnet.net				
v	^ ttl 72	0 pps	0/59	= 0%	0 pps
192.220.238.129					
192.35.180.201	seamr1-gw.nwnet.net				
v	^ ttl 72	0 pps	0/59	= 0%	0 pps
128.15.1.249	littleman.bigco.com				
v	___ ttl 72	0 pps	?/59		0 pps

INTERNET-DRAFT

Multicast Debugging Handbook

20 November 2000

```
128.15.1.250  128.15.1.250
Receiver      Query Source
```

#### Facilities used

IGMP multicast trace facility

#### Availability

mtrace is now distributed independently of mrouterd.

Source code is available from:

```
ftp://ftp.parc.xerox.com/pub/net-research/ipmulti/mtrace5.1.tar.Z
```

#### Binaries:

```
ftp://ftp.parc.xerox.com/pub/net-research/ipmulti/mtrace5.1-sparc-sunos41x.tar.Z
```

```
ftp://ftp.parc.xerox.com/pub/net-research/ipmulti/mtrace5.1-sparc-solaris2.tar.Z
```

```
ftp://ftp.parc.xerox.com/pub/net-research/ipmulti/mtrace5.1-alpha-osf1.tar.Z
```

```
ftp://ftp.parc.xerox.com/pub/net-research/ipmulti/mtrace5.1-sgi-irix.tar.Z
```

## 6.6. MBONE mapping tools

### 6.6.1. mrtree

#### Author

Dave Thaler, dthaler@microsoft.com

Andy Adams, ala@merit.edu

#### Description

mrtree uses a combination of IGMP and SNMP queries to discover the actual and potential multicast (sub)trees for a given source and group, rooted at a given router. An actual tree, discovered using the multicast routing MIB, consists of routers which are currently forwarding multicast traffic to a group from a given source. A potential tree, discovered using the DVMRP MIB, is one which would exist if every host were a member of the group.

#### Example

```
% mrtree mbone.merit.edu 224.2.143.24 204.62.246.73
Actual distribution tree rooted at mbone.merit.edu for group 224.2.143.24
and source 204.62.246.73...
0 mbone.merit.edu (198.108.2.20) [ver 3.8,prune,genid,mtrace],
  247390 pkts
1 cujo.merit.edu (198.108.60.97) [ver 3.6,prune,genid,mtrace], 333448
  6 pkts (1347%)
```

```
2 subnet: 198.108.60/24
2 shockwave.merit.edu (198.108.60.69) [ver 3.8,prune,genid,mtrace],
  1239130 pkts (500%)
1 tibia.cic.net (192.217.65.100) [ver 3.8,prune,genid,mtrace]
  ... (No response from tibia.cic.net)
2 fibula.cic.net (192.217.65.101) [ver 3.8,prune,genid,mtrace] ?
2 dcl2.gw.uiuc.edu (192.17.2.8) [ver 1.0] ?
2 goober.mci.net (204.70.104.45) [ver 3.6,prune,genid,mtrace] ?
  ... (goober.mci.net did not respond to DVMRP 'NEIGHBORS' msg)
1 a-wing.jvnc.net (130.94.40.6) [ver 3.3]
  ... (a-wing.jvnc.net does not support SNMP)
2 liberty-eth0/0.jvnc.net (130.94.40.1) [ver 10.2] ?
2 noc.hpc.org (192.187.8.2) [ver 3.8,prune,genid,mtrace] ?
2 liberty.jvnc.net (130.94.40.201) [ver 10.2] ?
2 dstest.ds.internic.net (198.49.45.4) [ver 3.8,prune,genid,mtrace] ?
2 cybercast.cc.nus.sg (137.132.9.70) [ver 3.6,prune,genid,mtrace] ?
  ... (cybercast.cc.nus.sg did not respond to DVMRP 'NEIGHBORS' msg)
```

#### Facilities used

DVMRP and multicast routing MIBs  
IGMP ASK\_NEIGHBORS message (DVMRP)

#### Availability

mrtree is available for UNIX and is included in the  
SNMP-capable mrouter distribution, available at:  
<ftp://ftp.merit.edu/net-research/mbone/mirrors/mrouter/>

mrtree is also available in the MVIEW distribution, available at:  
<ftp://ftp.merit.edu/net-research/mbone/mview/>

#### 6.6.2. map-mbone

##### Author

Pavel Curtis, [pavel@parc.xerox.com](mailto:pavel@parc.xerox.com)

##### Description

map-mbone is useful for discovering the topology within a DVMRP routing domain; to do this, it uses the IGMP ASK\_NEIGHBORS message to discover the neighbors of the starting router. If the -f (flooding) option is enabled (this is the default if no starting router is specified), then once these neighbors are discovered, they too are queried. This continues until the leaf routers are reached. This option should be used with care since it can result in excessive load on multicast routers.

If a starting router is specified but the -f option is not used, then the search terminates after the first hop routers are discovered, the output of map-mbone is very similar to that for mrinfo. Routers discovered by map-mbone are queried for their version numbers, and if this query is successful, for their metrics, thresholds, and flags, and the results are presented in an indented list format.

#### Example

```
% map-mbone 192.80.214.199
192.41.177.196: alias for 128.167.252.196

128.167.252.196 (collegepk-mbone1.bbnplanet.net): <v11.2>
  192.41.177.196: 192.41.177.196 [1/0/querier/down]
  192.80.214.199: 192.80.214.199 (collegepk-mbone1.bbnplanet.net) [1/0/querier]
  128.167.252.196: 205.128.246.2 (usnrctc.bbnplanet.net) [1/32/tunnel/querier]
                    204.148.62.28 (mbone-e.ans.net) [1/32/tunnel/querier]
                    192.41.177.197 (wtn-ms1.bbnplanet.net) [1/32/tunnel/querier]
                    128.175.13.36 (pfet.nss.udel.edu) [1/32/tunnel/querier/down]
                    205.130.85.3 (philipii.nap.edu) [1/32/tunnel/querier/down]
                    204.167.201.38 (dallas2-mbone1.bbnplanet.net) [1/64/tunnel/qu
                    192.221.48.234 (atlanta3-mbone1.bbnplanet.net) [1/64/tunnel/c
                    134.205.93.150 (dilbert.sam.pentagon.mil) [1/32/tunnel/querie
                    128.167.1.197 (cpk-ms1.ser.bbnplanet.com) [1/16/tunnel/querie
                    192.221.34.22 (cdrn.bbnplanet.net) [1/32/tunnel/querier]
                    128.244.93.3 (sage.jhuapl.edu) [1/32/tunnel/querier]
                    192.41.177.199 (wtn-ms2.bbnplanet.net) [1/16/tunnel/querier]
                    137.39.43.34 (MBONE1.UU.NET) [1/32/tunnel/querier]
                    199.94.207.2 (cambridge1-mbone1.bbnplanet.net) [1/32/tunnel/c
                    131.119.0.197 (paloalto-mbone1.bbnplanet.net) [1/64/tunnel/qu
                    128.167.254.165 (devo.sura.net) [1/32/tunnel/querier/down]
                    128.167.252.196 (collegepk-mbone1.bbnplanet.net) [1/0/querier

192.80.214.199 (collegepk-mbone1.bbnplanet.net): alias for 128.167.252.196
```

#### Facilities used

IGMP ASK\_NEIGHBORS message (DVMRP)

#### Availability

map-mbone is available for UNIX, and the software and manual pages are included in the SNMP-capable mrouted distribution, available at:  
<ftp://ftp.merit.edu/net-research/mbone/mirrors/mrouted/>

INTERNET-DRAFT

Multicast Debugging Handbook

20 November 2000

### 6.6.3. asn

#### Author

Dave Thaler, dthaler@microsoft.com

#### Description

asn gives the AS number of a given IP address by querying the routing arbiter database.

#### Example

```
% asn 141.213.10.41
AS237
```

#### Facilities used

Routing arbiter database

#### Availability

asn is included in the MVIEW distribution, available at:  
<ftp://ftp.merit.edu/net-research/mbone/mview/>

### 6.6.4. asname

#### Author

Dave Thaler, dthaler@microsoft.com

#### Description

asname gets the name of an AS, given the AS number by querying the WHOIS database.

#### Example

```
% asname 237
NSFNETTEST14-AS
```

#### Facilities used

WHOIS database

#### Availability

asname is included in the MVIEW distribution, available at:

<ftp://ftp.merit.edu/net-research/mbone/mview/>

## 6.7. Network Operations Center tools

These tools are suitable for use in a Network Operations Center.

### 6.7.1. MVIEW

#### Authors

Dave Thaler, [dthaler@microsoft.com](mailto:dthaler@microsoft.com)  
Andy Adams, [ala@merit.edu](mailto:ala@merit.edu)

#### Description

MVIEW uses utilities such as `mstat`, `mtrace`, `mrtree`, `asn` and `asname` in order to produce a graphical depiction of the multicast network topology and the actual and potential multicast trees for a given group and source.

#### Example

Further information on MVIEW as well as examples are available from:  
<http://www.merit.edu/net-research/mbone/mviewdoc/Welcome.html>

#### Facilities used

PIM, DVMRP, IGMP, and multicast routing MIBs (`mstat`)  
IGMP ASK\_NEIGHBORS message (`mrinfo`)  
Routing arbiter database (`asn`)  
WHOIS database (`asname`)

#### Availability

MVIEW is available for UNIX, and can be obtained from:  
<ftp://ftp.merit.edu/net-research/mbone/mview/>

Documentation is available as:  
<ftp://ftp.merit.edu/net-research/mbone/mviewdoc/>

### 6.7.2. Multicast heartbeat

#### Author

Many and various

#### Description

Devices implementing the multicast heartbeat listen on a designated group. If traffic is not observed on the group for a specified amount of time, an SNMP trap is generated. This allows multicast monitoring to be easily integrated into existing SNMP consoles. In situations where a shared-tree multicast routing protocol is used (such as sparse-mode PIM or CBT), it is recommended that the heartbeat generator be located close to the RP or core nodes, so as that loss of the heartbeat will correlate closely with loss of connectivity to the RP or core. Suitable heartbeat mechanisms include SNTTP, which uses the group 224.0.1.1 (ntp.mcast.net) and UDP port 123; and SAP, which uses the group 224.2.127.254 (sap.mcast.net) and UDP port 9875.

#### Example

For further information on SNTTP, consult [1].

#### Facilities used

- SNTTP (for time-based heartbeats)
- SAP (for session announcement heartbeats)
- SNMP traps (for alerts)

#### Availability

### 6.8. Network analysis tools

#### 6.8.1. Dr. Watson, the Network Detective's Assistant (DWTNDA)

##### Author

Karl Auerbach, karl@cavebear.com

##### Description

DWTNDA is a general purpose troubleshooting tool with some IP multicast tools (in addition to a fair number of non-multicast tools). For example it can watch IGMP "join" activity on a LAN and put up a real-time display in tabular format. It can generate some test packets, like IGMPv2 Leaves or Group Membership Requests. It can generate and respond to multicast pings (icmp, udp, or snmp based.) It will eventually acquire more sophisticated multicast facilities.

##### Example

See <http://www.cavebear.com/dwtnda/> for examples.

## Facilities used

This is a troubleshooting tool, so it will typically respond to packets that, strictly speaking, ought to go unanswered.

## Availability

DWTNDA runs on MS-DOS and Windows 95/98 and is free. Source is not provided. See <http://www.cavebear.com/dwtnda/> for various documents and download information.

### 6.8.2. Mtap

#### Author

Luis Fernando da Silva Barra, [barra@ax.apc.org](mailto:barra@ax.apc.org)  
Michael Stanton, [michael@omega.lncc.br](mailto:michael@omega.lncc.br)

#### Description

MTap is a tool for observing IP multicast packet traffic crossing a subnet, normally an Ethernet.

Each packet sent to an IP multicast group address (class D) is captured, and information is extracted concerning its origin, its size, and so forth. This information is summarized, permitting the determination of the current network load resulting from multicast traffic. Apart from global summaries, traffic information is summarized by group and by source, permitting the determination of the contribution of each group and each individual sender to global traffic. The data recorded are as follows: number of multicast packets and total multicast bytes passing through the network, load level, and date and time of the last packet received.

As well as processing packets sent to a multicast address, MTap also records separately multicast packets encapsulated in point-to-point packets. Thus we can also deal with traffic in DVMRP tunnels between multicast routers, and tunnel traffic data are recorded in the same way as for a group.

As well as recording the data, MTap also permits that individual packet data be exhibited in dump format at capture time, both for multicast packets and for tunneled packets.

In order to evaluate the impact which a group imposes on a subnetwork, MTap can enter or leave a multicast group, using the IGMP protocol. Thus traffic can be observed for a group which has no other members on the subnetwork.

In addition to passively observing and recording multicast traffic, MTap has a notification mechanism, which sets off an alarm whenever user-specified load levels are exceeded, either globally, by group or by tunnel. Notifications are also logged in a dedicated window.

#### Example

Further information on Mtap will be available from:  
<http://www.inf.puc-rio.br/~michael/GERENTE/tools>

#### Facilities used

Berkeley Packet Filter (BPF)

#### Availability

MTap uses a window-based user interface, developed using Tcl/Tk, and captures packets through the Berkeley Packet Filter (BPF). It can thus be ported to different platforms.

Mtap, which is still under development, has been ported to Linux and Solaris; minor problems related to packet capture have still to be resolved for the Solaris version. When it is released, it will be available from:

<http://www.inf.puc-rio.br/~michael/GERENTE/tools>

## 7. References

- [1] Mills, D., "Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI", RFC 2030, October 1996.
- [2] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.
- [3] McCloghrie, K., Farinacci, D., Thaler, D., "Internet Group Management Protocol MIB", Internet draft (work in progress), draft-ietf-idmr-igmp-mib-10.txt, June 1999.
- [4] Handley, M., Jacobson, V., "SDP: Session Description Protocol (Version 1)", RFC 2327, April 1998.
- [5] McCloghrie, K., Farinacci, D., Thaler D., "IP Multicast Routing MIB", Internet draft (work in progress), draft-ietf-idmr-multicast-routmib-10.txt, July 1999.
- [6] McCloghrie, K., Farinacci, D., Thaler, D., "Protocol Independent Multicast MIB", Internet draft (work in progress), draft-ietf-idmr-

pim-mib-07.txt, July 1999.

- [7] Thaler, D., "Distance Vector Multicasting Routing Protocol MIB", Internet draft (work in progress), draft-thaler-dvmpm-mib-09.txt, May 1998.
- [8] Fenner, W., Casner, S., "A "traceroute" facility for IP Multicast", Internet draft (work in progress), draft-ietf-idmr-traceroute-ipm-05.txt, June 1999.
- [9] Rekhter, Y. et al., "Address Allocation for Private Internets", RFC 1918, February, 1996.

## 8. Security Considerations

SNMP-based monitoring tools require that the manager be provided access to the relevant MIBs. In order to limit security risks, such access will typically be provided on a selective basis. For example, the authentication and access control facilities in SNMP v3 can be used to limit access to authorized users.

MBONE-mapping tools such as map-mbone should be used with care since in flooding mode they can result in excessive load on multicast routers.

Through use of RTP monitoring tools, it may be possible to obtain sensitive information on user viewing habits. In order to protect against this, encryption technologies such as IPSEC can be used to provide confidentiality.

## 9. Acknowledgments

Thanks to Karl Auerbach for the description of the Dr. Watson tool, and to Michael Stanton for the description of the Mtap tool.

## 10. Authors' Addresses

Dave Thaler  
Microsoft Corporation  
One Microsoft Way  
Redmond, WA 98052

Phone: 425-703-8835  
EMail: dthaler@microsoft.com

Bernard Aboba  
Microsoft Corporation  
One Microsoft Way  
Redmond, WA 98052

INTERNET-DRAFT

Multicast Debugging Handbook

20 November 2000

Phone: 425-936-6605

EMail: bernarda@microsoft.com

## 11. Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved. This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English. The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns. This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF

## 12. Expiration Date

This memo is filed as <draft-ietf-mboned-mdh-05.txt>, and expires July 1, 2001.