
 6696 Fri Dec 19 04:18:42 2008

new/./arp.lm.txt
 Clearview IPMP manpages

1 System Administration Commands arp(1M)

5 NAME
 6 arp - address resolution display and control

8 SYNOPSIS
 9 arp hostname

12 arp -a [-n]

15 arp -d hostname

18 arp -f filename

21 arp -s hostname ether_address [temp] [pub] [trail]
 22 [permanent]

25 DESCRIPTION
 26 The arp program displays and modifies the Internet-to-MAC
 27 address translation tables used by the address resolution
 28 protocol (see arp(7P)).

31 With no flags, the program displays the current ARP entry
 32 for hostname. The host may be specified by name or by
 33 number, using Internet dot notation.

36 Options that modify the ARP translation tables (-d, -f, and
 37 -s) can be used only when the invoked command is granted the
 38 sys_net_config privilege. See privileges(5).

40 OPTIONS
 41 -a Display all of the current ARP entries. The definition
 42 for the flags in the table are:

44 d Unverified; this is a local IP address that is
 45 currently undergoing Duplicate Address Detection.
 46 ARP will not respond to requests for this address
 47 until Duplicate Address Detection completes.

50 o Old; this entry is aging away. If IP requests it
 51 again, a new ARP query will be generated. This
 52 state is used for detecting peer address changes.

55 y Delayed; periodic address defense and conflict
 56 detection was unable to send a packet due to

67 System Administration Commands arp(1M)

71 internal network use limits for non-traffic-
 72 related messages (100 packets per hour per inter-
 73 face). This occurs only on interfaces with very
 74 large numbers of aliases.

77 A Authority; this machine is authoritative for this
 78 IP address. ARP will not accept updates from
 79 other machines for this entry.

82 L Local; this is a local IP address configured on
 83 one of the machine's logical interfaces. ARP will
 84 defend this address if another node attempts to
 85 claim it.

88 M Mapping; only used for the multicast entry for
 89 224.0.0.0

92 P Publish; includes IP address for the machine and
 93 the addresses that have explicitly been added by
 94 the -s option. ARP will respond to ARP requests
 95 for this address.

98 S Static; entry cannot be changed by learned infor-
 99 mation. This indicates that the permanent flag
 100 was used when creating the entry.

103 U Unresolved; waiting for ARP response.

105 You can use the -n option with the -a option to dis-
 106 able the automatic numeric IP address-to-name transla-
 107 tion. Use arp -an or arp -na to display numeric IP
 108 addresses. The arp -a option is equivalent to:

110 # netstat -p -f inet

113 ...and -an and -na are equivalent to:

115 # netstat -pn -f inet

120 -d Delete an entry for the host called hostname.

121
 122 **Note that ARP entries for IPMP data and test addresses**
 123 **are managed by the kernel, and thus cannot be deleted.**

128 SunOS 5.11 Last change: 25 Jul 2006 2

135 System Administration Commands arp(1M)

139 -f Read the file named filename and set multiple entries
140 in the ARP tables. Entries in the file should be of
141 the form:

143 hostname MACaddress [temp] [pub] [trail] [permanent]

146 See the -s option for argument definitions.

149 -s Create an ARP entry for the host called hostname with
150 the MAC address MACaddress. For example, an Ethernet
151 address is given as six hexadecimal bytes separated by
152 colons. The entry will not be subject to deletion by
153 aging unless the word temp is specified in the com-
154 mand. If the word pub is specified, the entry will be
155 published, which means that this system will respond
156 to ARP requests for hostname even though the hostname
157 is not its own. The word permanent indicates that the
158 system will not accept MAC address changes for host-
159 name from the network.

161 Solaris does not implement trailer encapsulation, and
162 the word trail is accepted on entries for compatibil-
163 ity only.

165 arp -s can be used for a limited form of proxy ARP
166 when a host on one of the directly attached networks
167 is not physically present on a subnet. Another machine
168 can then be configured to respond to ARP requests
169 using arp -s. This is useful in certain SLIP confi-
170 gurations.

172 Non-temporary proxy ARP entries for an IPMP group are
173 automatically managed by the kernel. Specifically, if
174 the hardware address in an entry matches the hardware
175 address of an IP interface in an IPMP group, and the
176 IP address is not local to the system, this will be
177 regarded as an IPMP proxy ARP entry. This entry will
178 have its hardware address automatically adjusted in
179 order to keep the IP address reachable so long as the
180 IPMP group has not entirely failed.

182 ARP entries must be consistent across an IPMP group.
183 Therefore, ARP entries cannot be associated with
184 individual underlying IP interfaces in an IPMP group,
185 and must instead be associated with the corresponding
186 IPMP IP interface.

188 Note that ARP entries for IPMP data and test addresses
189 are managed by the kernel, and thus cannot be changed.

192 ATTRIBUTES

193 See attributes(5) for descriptions of the following attri-

194 butes:

198
199
200
201
202

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu

205 SEE ALSO
206 ifconfig(1M), netstat(1M), attributes(5), privileges(5),
207 arp(7P)

213 SunOS 5.11 Last change: 25 Jul 2006 3

 13679 Fri Dec 19 04:18:44 2008

new/./arp.7p.txt
 Clearview IPMP manpages

```

1 NAME
2   arp, ARP - Address Resolution Protocol

4 SYNOPSIS
5   #include <sys/fcntl.h>

7   #include <sys/socket.h>

9   #include <net/if_arp.h>

11  #include <netinet/in.h>

13  s = socket(AF_INET, SOCK_DGRAM, 0);

15  d = open ("/dev/arp", oflag);

17 DESCRIPTION
18  ARP is a protocol used to map dynamically between Internet
19  Protocol (IP) and Ethernet addresses. It is used by all Eth-
20  ernet datalink providers (network drivers) and can be used
21  by other datalink providers that support broadcast, includ-
22  ing FDDI and Token Ring. The only network layer supported
23  in this implementation is the Internet Protocol, although
24  ARP is not specific to that protocol.

26  ARP caches IP-to-link-layer address mappings. When an inter-
27  face requests a mapping for an address not in the cache, ARP
28  queues the message that requires the mapping and broadcasts
29  a message on the associated network requesting the address
30  mapping. If a response is provided, ARP caches the new map-
31  ping and transmits any pending message. ARP will queue a
32  maximum of four packets while awaiting a response to a map-
33  ping request. ARP keeps only the first four transmitted
34  packets.

36 APPLICATION PROGRAMMING INTERFACE
37  The STREAMS device /dev/arp is not a Transport Level Inter-
38  face (TLI) transport provider and may not be used with the
39  TLI interface.

41  To facilitate communications with systems that do not use
42  ARP, ioctl() requests are provided to enter and delete
43  entries in the IP-to-link address tables. Ioctls that
44  change the table contents require sys_net_config privilege.

46 SunOS 5.10          Last change: 2 Dec 2008          1

48 Protocols          arp(7P)

50  See privileges(5).

52  #include <sys/sockio.h>
53  #include <sys/socket.h>
54  #include <net/if.h>
55  #include <net/if_arp.h>
56  struct arpreq arpreq;
57  ioctl(s, SIOCSARP, (caddr_t)&arpreq);
58  ioctl(s, SIOCGARP, (caddr_t)&arpreq);
59  ioctl(s, SIOCDEARP, (caddr_t)&arpreq);

61  SIOCSARP, SIOCGARP and SIOCDEARP are BSD compatible ioctls.

```

```

62  These ioctls do not communicate the mac address length
63  between the user and the kernel (and thus only work for 6
64  byte wide Ethernet addresses). To manage the ARP cache for
65  media that has different sized mac addresses, use SIOCSXARP,
66  SIOCGXARP and SIOCDXARP ioctls.

68  #include <sys/sockio.h>
69  #include <sys/socket.h>
70  #include <net/if.h>
71  #include <net/if_dl.h>
72  #include <net/if_arp.h>
73  struct xarpreq xarpreq;
74  ioctl(s, SIOCSXARP, (caddr_t)&xarpreq);
75  ioctl(s, SIOCGXARP, (caddr_t)&xarpreq);
76  ioctl(s, SIOCDXARP, (caddr_t)&xarpreq);

78  Each ioctl() request takes the same structure as an argu-
79  ment. SIOCS[X]ARP sets an ARP entry, SIOCG[X]ARP gets an ARP
80  entry, and SIOC[X]ARP deletes an ARP entry. These ioctl()
81  requests may be applied to any Internet family socket
82  descriptors, or to a descriptor for the ARP device. Note
83  that SIOCS[X]ARP and SIOC[X]ARP require a privileged user,
84  while SIOCG[X]ARP does not.

86  The arpreq structure contains

88  /*
89  * ARP ioctl request
90  */
91  struct arpreq {
92      struct sockaddr arp_pa; /* protocol address */
93      struct sockaddr arp_ha; /* hardware address */
94      int arp_flags; /* flags */
95  };

97 SunOS 5.10          Last change: 2 Dec 2008          2

99 Protocols          arp(7P)

101  The xarpreq structure contains:

103  /*
104  * Extended ARP ioctl request
105  */
106  struct xarpreq {
107      struct sockaddr_storage xarp_pa; /* protocol address */
108      struct sockaddr_dl xarp_ha; /* hardware address */
109      int xarp_flags; /* arp_flags field values */
110  };
111  #define ATF_COM 0x2 /* completed entry (arp_ha valid) */
112  #define ATF_PERM 0x4 /* permanent (non-aging) entry */
113  #define ATF_PUBL 0x8 /* publish (respond for other host) */
114  #define ATF_USETRAILERS 0x10 /* send trailer pkcts to host */
115  #define ATF_AUTHORITY 0x20 /* hardware address is authoritative */

117  The address family for the [x]arp_pa sockaddr must be
118  AF_INET. The ATF_COM flag bits ([x]arp_flags) cannot be
119  altered. ATF_USETRAILERS is not implemented on Solaris and
120  altered. ATF_USETRAILER is not implemented on Solaris and
121  is retained for compatibility only. ATF_PERM makes the
122  entry permanent (disables aging) if the ioctl() request
123  succeeds. ATF_PUBL specifies that the system should respond
124  to ARP requests for the indicated protocol address coming
125  from other machines. This allows a host to act as an "ARP
126  server," which may be useful in convincing an ARP-only
127  machine to talk to a non-ARP machine. ATF_AUTHORITY indi-

```

127 cates that this machine owns the address. ARP does not
 128 update the entry based on received packets.

130 The address family for the arp_ha sockaddr must be
 131 AF_UNSPEC.

133 Before invoking any of the SIOC*XARP ioctls, user code must
 134 fill in the xarp_pa field with the protocol (IP) address
 135 information, similar to the BSD variant. The SIOC*XARP
 136 ioctls come in two (legal) varieties, depending on
 137 xarp_ha.sdl_nlen:
 138 1. if sdl_nlen = 0, it behaves as an extended BSD ioctl.
 139 The kernel uses the IP address to determine the network
 140 interface.
 141 2. if (sdl_nlen > 0) and (sdl_nlen < LIFNAMSIZ), the ker-
 142 nel uses the interface name in sdl_data[0] to determine
 143 the network interface; sdl_nlen represents the length of
 144 the string (excluding terminating null character).
 145 3. if (sdl_nlen >= LIFNAMSIZ), an error (EINVAL) is
 146 flagged from the ioctl.

148 SunOS 5.10 Last change: 2 Dec 2008 3

150 Protocols arp(7P)

152 Other than the above, the xarp_ha structure should be 0-
 153 filled except for SIOCSXARP, where the sdl_alen field must
 154 be set to the size of hardware address length and the
 155 hardware address itself must be placed in the
 156 LLADDR/sdl_data[] area. (EINVAL will be returned if user
 157 specified sdl_alen does not match the address length of the
 158 identified interface).

160 On return from the kernel on a SIOCGXARP ioctl, the kernel
 161 fills in the name of the interface (excluding terminating
 162 NULL) and its hardware address, one after another, in the
 163 sdl_data/LLADDR area; if the two are larger than can be held
 164 in the 244 byte sdl_data[] area, an ENOSPC error is
 165 returned. Assuming it fits, the kernel will also set
 166 sdl_alen with the length of hardware address, sdl_nlen with
 167 the length of name of the interface (excluding terminating
 168 NULL), sdl_type with an IFT_* value to indicate the type of
 169 the media, sdl_slen with 0, sdl_family with AF_LINK and
 170 sdl_index (which if not 0) with system given index for the
 171 interface. The information returned is very similar to that
 172 returned via routing sockets on an RTM_IFINFO message.

174 **The ARP ioctls have several additional restrictions and**
 175 **enhancements when used in conjunction with IPMP:**

177 * The ARP mappings for IPMP data and test addresses are
 178 managed by the kernel, and thus cannot be changed
 179 through any ARP ioctls, though they may be retrieved
 180 using *SIOCGARP* or *SIOCGXARP*.

182 * ARP mappings for a given IPMP group must be consistent
 183 across the group. Therefore, ARP mappings cannot be
 184 associated with individual underlying IP interfaces in
 185 an IPMP group, and must instead be associated with the
 186 corresponding IPMP IP interface.

188 * Proxy ARP mappings for an IPMP group are automatically
 189 managed by the kernel. Specifically, if the hardware
 190 address in a *SIOCSARP* or *SIOCSXARP* request matches
 191 the hardware address of an IP interface in an IPMP
 192 group, and the IP address is not local to the system,

193 the kernel will regard this as a IPMP Proxy ARP entry.
 194 This IPMP Proxy ARP entry will have its hardware
 195 address automatically adjusted in order to keep the IP
 196 address reachable (provided the IPMP group has not
 197 entirely failed).

199 ARP performs duplicate address detection for local
 200 addresses. When a logical interface is brought up (IFF_UP)
 201 or any time the hardware link goes up (IFF_RUNNING), ARP
 202 sends probes (ar\$spa == 0) for the assigned address. If a
 203 conflict is found, the interface is torn down. See
 204 ifconfig(1M) for more details.

206 ARP watches for hosts impersonating the local host, that is,
 207 any host that responds to an ARP request for the local
 208 host's address, and any address for which the local host is
 209 an authority. ARP defends local addresses and logs those
 210 with ATF_AUTHORITY set, and can tear down local addresses on
 211 an excess of conflicts.

213 ARP also handles UNARP messages received from other nodes.
 214 It does not generate these messages.

216 PACKET EVENTS
 217 The arp driver registers itself with the netinfo interface.
 218 To gain access to these events, a handle from
 219 net_protocol_lookup must be acquired by passing it the value
 220 NHF_ARP. Through this interface, two packet events are sup-
 221 ported:

223 SunOS 5.10 Last change: 2 Dec 2008 4

225 Protocols arp(7P)

227 Physical in - ARP packets received via a network inter face

229 Physical out - ARP packets to be sent out via a network
 230 interface

232 For ARP packets, the hook_pkt_event structure is filled out
 233 as follows:

235 hpe_ifp

237 Identifier indicating the inbound interface for pack-
 238 ets received with the "physical in" event.

240 hpe_ofp

242 Identifier indicating the outbound interface for pack-
 243 ets received with the "physical out" event.

245 hpe_hdr

247 Pointer to the start of the ARP header (not the eth-
 248 ernet header).

250 hpe_mp

252 Pointer to the start of the mblk_t chain containing the
 253 ARP packet.

255 hpe_mb

257 Pointer to the mblk_t with the ARP header in it.

```

259 NETWORK INTERFACE EVENTS
260 In addition to events describing packets as they move
261 through the system, it is also possible to receive notifica-
262 tion of events relating to network interfaces. These events
263 are all reported back through the same callback. The list
264 of events is as follows:

266 plumb

268 A new network interface has been instantiated.

270 SunOS 5.10 Last change: 2 Dec 2008 5
272 Protocols arp(7P)

274 unplumb

276 A network interface is no longer associated with ARP.

278 SEE ALSO
279 arp(1M), ifconfig(1M), privileges(5), if_tcp(7P), inet(7P),
280 netinfo(9F)

282 Plummer, Dave, An Ethernet Address Resolution Protocol or
283 Converting Network Protocol Addresses to 48 .bit Ethernet
284 Addresses for Transmission on Ethernet Hardware, RFC 826,
285 STD 0037, November 1982.

287 Malkin, Gary, ARP Extension - UNARP, RFC 1868, November,
288 1995

290 DIAGNOSTICS
291 Several messages can be written to the system logs (by the
292 IP module) when errors occur. In the following examples, the
293 hardware address strings include colon (:) separated ASCII
294 representations of the link layer addresses, whose lengths
295 depend on the underlying media (for example, 6 bytes for
296 Ethernet).

298 Node %x:%x ... %x:%x is using our IP address %d.%d.%d.%d on
299 %s.

301 Duplicate IP address warning. ARP has discovered another
302 host on a local network that responds to mapping
303 requests for the Internet address of this system, and
304 has defended the system against this node by re-
305 announcing the ARP entry.

307 %s has duplicate address %d.%d.%d.%d (in use by %x:%x ...
308 %x:%x); disabled.

310 Duplicate IP address detected while performing initial
311 probing. The newly-configured interface has been shut
312 down.

314 %s has duplicate address %d.%d.%d.%d (claimed by %x:%x ...
315 %x:%x); disabled.

317 Duplicate IP address detected on a running IP interface.
318 The conflict cannot be resolved, and the interface has
319 been disabled to protect the network.

321 SunOS 5.10 Last change: 2 Dec 2008 6
323 Protocols arp(7P)

```

```

325 Recovered address %d.%d.%d.%d on %s.

327 An interface with a previously-conflicting IP address
328 has been recovered automatically and reenabled. The con-
329 flict has been resolved.

331 Proxy ARP problem? Node '%x:%x ... %x:%x' is using
332 %d.%d.%d.%d on %s

334 This message appears if arp(1M) has been used to create
335 a published permanent (ATF_AUTHORITY) entry, and some
336 other host on the local network responds to mapping
337 requests for the published ARP entry.

```

```

*****
26006 Fri Dec 19 04:18:45 2008
new/./dhcpagent.txt
Clearview IPMP manpages
*****
 1 System Administration Commands          dhcpagent(1M)

5 NAME
6   dhcpagent - Dynamic Host Configuration Protocol (DHCP)
7   client daemon

 9 SYNOPSIS
10   dhcpagent [-a] [ -d n] [-f] [-v]

13 DESCRIPTION
14   dhcpagent implements the client half of the Dynamic Host
15   Configuration Protocol (DHCP) for machines running Solaris
16   software.

19   The dhcpagent daemon obtains configuration parameters for
20   the client (local) machine's network interfaces from a DHCP
21   server. These parameters may include a lease on an IP
22   address, which gives the client machine use of the address
23   for the period of the lease, which may be infinite. If the
24   client wishes to use the IP address for a period longer than
25   the lease, it must negotiate an extension using DHCP. For
26   this reason, dhcpagent must run as a daemon, terminating
27   only when the client machine powers down.

30   For IPv4, the dhcpagent daemon is controlled through
31   ifconfig(1M) in much the same way that the init(1M) daemon
32   is controlled by telinit(1M). dhcpagent can be invoked as a
33   user process, albeit one requiring root privileges, but this
34   is not necessary, as ifconfig(1M) will start it automati-
35   cally.

38   For IPv6, the dhcpagent daemon is invoked automatically by
39   in.ndpd(1M). It can also be controlled through ifconfig(1M),
40   if necessary.

43   When invoked, dhcpagent enters a passive state while it
44   awaits instructions from ifconfig(1M) or in.ndpd(1M). When
45   it receives a command to configure an interface, it brings
46   up the interface (if necessary) and starts DHCP. Once DHCP
47   is complete, dhcpagent can be queried for the values of the
48   various network parameters. In addition, if DHCP was used
49   to obtain a lease on an address for an interface, it con-
50   figures the address for use. When a lease is obtained, it
51   is automatically renewed as necessary. If the lease cannot
52   be renewed, dhcpagent will unconfigure the address, but the
53   interface will be left up and dhcpagent will attempt to
54   acquire a new address lease. dhcpagent monitors system
55   suspend/resume events and will validate any non-permanent
56   leases with the DHCP server upon resume. Similarly,

60 SunOS 5.11          Last change: 15 May 2008          1

```

```

67 System Administration Commands          dhcpagent(1M)

71   dhcpagent monitors link up/down events and will validate
72   any non-permanent leases with the DHCP server when the
73   downed link is brought back up.

76   For IPv4, if the configured interface is found to be
77   unplumbed, or to have a different IP address, subnet mask,
78   or broadcast address from those obtained from DHCP, the
79   interface is abandoned from DHCP control.
77   unplumbed, marked down, or to have a different IP address,
78   subnet mask, or broadcast address from those obtained from
79   DHCP, the interface is abandoned by DHCP control.

82   For IPv6, dhcpagent automatically plumbs and unplumbs
83   logical interfaces as necessary for the IPv6 addresses
84   supplied by the server. The IPv6 prefix length (netmask) is
85   not set by the DHCPv6 protocol, but is instead set by
86   *in.ndpd(1M)* using prefix information obtained by Router
87   Advertisements. If any of the logical interfaces created by
88   dhcpagent is unplumbed, or configured with a different IP
89   address, it will be abandoned from DHCP control. If the
90   link-local interface is unplumbed, then all addresses
91   configured by DHCP on that physical interface will be
92   removed.
93   For IPv6, dhcpagent automatically plumbs and unplumbs logi-
94   cal interfaces as necessary for the IPv6 addresses supplied
95   by the server. The IPv6 prefix length (netmask) is not set
96   by the DHCPv6 protocol, but is instead set by in.ndpd(1M)
97   using prefix information obtained by Router Advertisements.
98   If any of the logical interfaces created by dhcpagent is
99   unplumbed, marked down, or configured with a different IP
100  address, it will be abandoned by DHCP control. If the link-
101  local interface is unplumbed, then all addresses configured
102  by DHCP on that physical interface will be removed.

103  DHCP also acts as a mechanism to configure other information
104  needed by the client, for example, the domain name and
105  addresses of routers. Aside from the IP address, and for
106  IPv4 alone, the netmask, broadcast address, and default
107  router, the agent does not directly configure the worksta-
108  tion, but instead acts as a database which may be interro-
109  gated by other programs, and in particular by dhcpinfo(1).

112  On clients with a single interface, this is quite straight-
113  forward. Clients with multiple interfaces may present diffi-
114  culties, as it is possible that some information arriving on

```

115 different interfaces may need to be merged, or may be incon-
 116 sistent. Furthermore, the configuration of the interfaces is
 117 asynchronous, so requests may arrive while some or all of
 118 the interfaces are still unconfigured. To handle these
 119 cases, one interface may be designated as primary, which
 120 makes it the authoritative source for the values of DHCP
 121 parameters in the case where no specific interface is
 122 requested. See dhcpinfo(1) and ifconfig(1M) for details.

127 SunOS 5.11 Last change: 15 May 2008 2

134 System Administration Commands dhcpagent(1M)

138 For IPv4, the dhcpagent daemon can be configured to request
 139 a particular host name. See the REQUEST_HOSTNAME description
 140 in the FILES section. When first configuring a client to
 141 request a host name, you must perform the following steps as
 142 root to ensure that the full DHCP negotiation takes place:

```
144 # pkill dhcpagent
145 # rm /etc/dhcp/interface.dhc
146 # reboot
```

151 All DHCP packets sent by dhcpagent include a vendor class
 152 identifier (RFC 2132, option code 60; RFC 3315, option code
 153 16). This identifier is the same as the platform name
 154 returned by the uname -i command, except:

156 o Any commas in the platform name are changed to
 157 periods.

159 o If the name does not start with a stock symbol and
 160 a comma, it is automatically prefixed with SUNW.

162 Messages

163 The dhcpagent daemon writes information and error messages
 164 in five categories:

166 critical

168 Critical messages indicate severe conditions that
 169 prevent proper operation.

172 errors

174 Error messages are important, sometimes unrecoverable
 175 events due to resource exhaustion and other unexpected
 176 failure of system calls; ignoring errors may lead to
 177 degraded functionality.

180 warnings

182 Warnings indicate less severe problems, and in most
 183 cases, describe unusual or incorrect datagrams received
 184 from servers, or requests for service that cannot be
 185 provided.

188 informational

193 SunOS 5.11 Last change: 15 May 2008 3

200 System Administration Commands dhcpagent(1M)

204 Informational messages provide key pieces of information
 205 that can be useful to debugging a DHCP configuration at
 206 a site. Informational messages are generally controlled
 207 by the -v option. However, certain critical pieces of
 208 information, such as the IP address obtained, are always
 209 provided.

212 debug

214 Debugging messages, which may be generated at two dif-
 215 ferent levels of verbosity, are chiefly of benefit to
 216 persons having access to source code, but may be useful
 217 as well in debugging difficult DHCP configuration prob-
 218 lems. Debugging messages are only generated when using
 219 the -d option.

223 When dhcpagent is run without the -f option, all messages
 224 are sent to the system logger syslog(3C) at the appropriate
 225 matching priority and with a facility identifier LOG_DAEMON.
 226 When dhcpagent is run with the -f option, all messages are
 227 directed to standard error.

229 DHCP Events and User-Defined Actions

230 If an executable (binary or script) is placed at
 231 /etc/dhcp/eventhook, the dhcpagent daemon will automatically
 232 run that program when any of the following events occur:

234 BOUND and BOUND6

236 These events occur during interface configuration. The
 237 event program is invoked when dhcpagent receives the
 238 DHCPv4 ACK or DHCPv6 Reply message from the DHCP server
 239 for the lease request of an address, indicating success-
 240 ful initial configuration of the interface. (See also
 241 the INFORM and INFORM6 events, which occur when confi-
 242 guration parameters are obtained without address
 243 leases.)

246 EXTEND and EXTEND6

248 These events occur during lease extension. The event
 249 program is invoked just after dhcpagent receives the
 250 DHCPv4 ACK or DHCPv6 Reply from the DHCP server for the
 251 DHCPv4 REQUEST (renew) message or the DHCPv6 Renew or
 252 Rebind message.

254 Note that with DHCPv6, the server might choose to remove
 255 some addresses, add new address leases, and ignore

259 SunOS 5.11 Last change: 15 May 2008 4

266 System Administration Commands dhcpagent(1M)

270 (allow to expire) still other addresses in a given Reply
 271 message. The EXTEND6 event occurs when a Reply is
 272 received that leaves one or more address leases still
 273 valid, even if the Reply message does not extend the
 274 lease for any address. The event program is invoked just
 275 before any addresses are removed, but just after any new
 276 addresses are added. Those to be removed will be marked
 277 with the IFF_DEPRECATED flag.

280 EXPIRE and EXPIRE6

282 These events occur during lease expiration. For DHCPv4,
 283 the event program is invoked just before the leased
 284 address is removed from an interface and the interface
 285 is marked as down. For DHCPv6, the event program is
 286 invoked just before the last remaining leased addresses
 287 are removed from the interface.

290 DROP and DROP6

292 These events occur during the period when an interface
 293 is dropped. The event program is invoked just before the
 294 interface is removed from DHCP control. If the interface
 295 has been abandoned due the user unplumbing the inter-
 296 face, then this event will occur after the user's action
 297 has taken place. The interface might not be present.

300 INFORM and INFORM6

302 These events occur when an interface acquires new or
 303 updated configuration information from a DHCP server by
 304 means of the DHCPv4 INFORM or the DHCPv6 Information-
 305 Request message. These messages are sent using an
 306 ifconfig(1M) dhcp inform command or when the DHCPv6
 307 Router Advertisement 0 (letter 0) bit is set and the M
 308 bit is not set. Thus, these events occur when the DHCP
 309 client does not obtain an IP address lease from the
 310 server, and instead obtains only configuration parame-
 311 ters.

314 LOSS6

316 This event occurs during lease expiration when one or
 317 more valid leases still remain. The event program is
 318 invoked just before expired addresses are removed. Those
 319 being removed will be marked with the IFF_DEPRECATED
 320 flag.

325 SunOS 5.11 Last change: 15 May 2008 5

332 System Administration Commands dhcpagent(1M)

336 Note that this event is not associated with the receipt
 337 of the Reply message, which occurs only when one or more
 338 valid leases remain, and occurs only with DHCPv6. If all
 339 leases have expired, then the EXPIRE6 event occurs
 340 instead.

343 RELEASE and RELEASE6

345 This event occurs during the period when a leased
 346 address is released. The event program is invoked just
 347 before dhcpagent relinquishes the address on an inter-
 348 face and sends the DHCPv4 RELEASE or DHCPv6 Release
 349 packet to the DHCP server.

353 The system does not provide a default event program. The
 354 file /etc/dhcp/eventhook is expected to be owned by root and
 355 have a mode of 755.

358 The event program will be passed two arguments, the inter-
 359 face name and the event name, respectively. For DHCPv6, the
 360 interface name is the name of the physical interface.

363 The event program can use the dhcpinfo(1) utility to fetch
 364 additional information about the interface. While the event
 365 program is invoked on every event defined above, it can
 366 ignore those events in which it is not interested. The event
 367 program runs with the same privileges and environment as
 368 dhcpagent itself, except that stdin, stdout, and stderr are
 369 redirected to /dev/null. Note that this means that the event
 370 program runs with root privileges.

373 If an invocation of the event program does not exit after 55
 374 seconds, it is sent a SIGTERM signal. If does not exit
 375 within the next three seconds, it is terminated by a SIGKILL
 376 signal.

new/./dhcpageant.txt

7

379 See EXAMPLES for an example event program.

381 OPTIONS
382 The following options are supported:

384 -a

386 Adopt a configured IPv4 interface. This option is for
387 use with diskless DHCP clients. In the case of diskless

391 SunOS 5.11 Last change: 15 May 2008 6

398 System Administration Commands dhcpageant(1M)

402 DHCP, DHCP has already been performed on the network
403 interface providing the operating system image prior to
404 running dhcpageant. This option instructs the agent to
405 take over control of the interface. It is intended pri-
406 marily for use in boot scripts.

408 The effect of this option depends on whether the inter-
409 face is being adopted.

411 If the interface is being adopted, the following condi-
412 tions apply:

414 dhcpageant uses the client id specified in
415 /chosen:<client_id>, as published by the PROM or as
416 specified on a boot(1M) command line. If this value is
417 not present, the client id is undefined. The DHCP server
418 then determines what to use as a client id. It is an
419 error condition if the interface is an Infiniband inter-
420 face and the PROM value is not present.

422 If the interface is not being adopted:

424 dhcpageant uses the value stored in
425 /etc/default/dhcpageant. If this value is not present,
426 the client id is undefined. If the interface is Infini-
427 band and there is no value in /etc/default/dhcpageant, a
428 client id is generated as described by the draft docu-
429 ment on DHCP over Infiniband, available at:

431 <http://www.ietf.org>

435 -d n

437 Set debug level to n. Two levels of debugging are
438 currently available, 1 and 2; the latter is more ver-
439 bose.

442 -f

444 Run in the foreground instead of as a daemon process.

new/./dhcpageant.txt

8

445 When this option is used, messages are sent to standard
446 error instead of to syslog(3C).

449 -v

451 Provide verbose output useful for debugging site confi-
452 guration problems.

457 SunOS 5.11 Last change: 15 May 2008 7

464 System Administration Commands dhcpageant(1M)

468 EXAMPLES
469 Example 1 Example Event Program

472 The following script is stored in the file
473 /etc/dhcp/eventhook, owned by root with a mode of 755. It is
474 invoked upon the occurrence of the events listed in the
475 file.

```
478 #!/bin/sh
480 (
481 echo "Interface name: " $1
482 echo "Event: " $2
484 case $2 in
485 "BOUND")
486     echo "Address acquired from server "\
487         '/sbin/dhcpinfo -i $1 ServerID\'
488     ;;
489 "BOUND6")
490     echo "Addresses acquired from server " \
491         '/sbin/dhcpinfo -v6 -i $1 ServerID\'
492     ;;
493 "EXTEND")
494     echo "Lease extended for " \
495         '/sbin/dhcpinfo -i $1 LeaseTim\'" seconds"
496     ;;
497 "EXTEND6")
498     echo "New lease information obtained on $i"
499     ;;
500 "EXPIRE" | "DROP" | "RELEASE")
501     ;;
503 esac
504 ) >/var/run/dhcp_eventhook_output 2>&1
```

509 Note the redirection of stdout and stderr to a file.

```

512 FILES
513     /etc/dhcp/ifa.dhc
514     /etc/dhcp/ifa.dh6

516     Contains the configuration for interface. The mere
517     existence of this file does not imply that the confi-
518     guration is correct, since the lease might have expired.
519     On start-up, dhcpagent confirms the validity of the

523 SunOS 5.11             Last change: 15 May 2008             8

530 System Administration Commands             dhcpagent(1M)

534     address using REQUEST (for DHCPv4) or Confirm (DHCPv6).

537     /etc/dhcp/duid
538     /etc/dhcp/iaid

540     Contains persistent storage for DUID (DHCP Unique Iden-
541     tifier) and IAID (Identity Association Identifier)
542     values. The format of these files is undocumented, and
543     applications should not read from or write to them.

546     /etc/default/dhcpagent

548     Contains default values for tunable parameters. All
549     values may be qualified with the interface they apply to
550     by prepending the interface name and a period (".") to
551     the interface parameter name. The parameters include:
552     the interface parameter name.

554     To configure IPv6 parameters, place the string .v6
555     between the interface name (if any) and the parameter
556     name. For example, to set the global IPv6 parameter
557     request list, use .v6.PARAM_REQUEST_LIST. To set the
558     CLIENT_ID (DUID) on hme0, use hme0.v6.CLIENT_ID.

560     The parameters include:

562     RELEASE_ON_SIGTERM

564     Indicates that a RELEASE rather than a DROP should
565     be performed on managed interfaces when the agent
566     terminates. Release causes the client to discard the
567     lease, and the server to make the address available
568     again. Drop causes the client to record the lease in
569     /etc/dhcp/interface.dhc or /etc/dhcp/interface.dh6
570     for later use.

573     OFFER_WAIT

575     Indicates how long to wait between checking for
576     valid OFFERS after sending a DISCOVER. For DHCPv6,

```

```

577         sets the time to wait between checking for valid
578         Advertisements after sending a Solicit.

581         CLIENT_ID

583         Indicates the value that should be used to uniquely
584         identify the client to the server. This value can
585         take one of three basic forms:

589 SunOS 5.11             Last change: 15 May 2008             9

596 System Administration Commands             dhcpagent(1M)

600         decimal,data...
601         0xHHHHH...
602         "string..."

605         The first form is an RFC 3315 DUID. This is legal
606         for both IPv4 DHCP and DHCPv6. For IPv4, an RFC 4361
607         Client ID is constructed from this value. In this
608         first form, the format of data... depends on the
609         decimal value. The following formats are defined for
610         this first form:

612         1,hwtype,time,lla

614         Type 1, DUID-LLT. The hwtype value is an integer
615         in the range 0-65535, and indicates the type of
616         hardware. The time value is the number of
617         seconds since midnight, January 1st, 2000 UTC,
618         and can be omitted to use the current system
619         time. The lla value is either a colon-separated
620         MAC address or the name of a physical interface.
621         If the name of an interface is used, the hwtype
622         value can be omitted. For example: 1,,hme0

625         2,enterprise,hex...

627         Type 2, DUID-EN. The enterprise value is an
628         integer in the range 0-4294967295 and represents
629         the SMI Enterprise number for an organization.
630         The hex string is an even-length sequence of
631         hexadecimal digits.

634         3,hwtype,lla

636         Type 3, DUID-LL. This is the same as DUID-LLT
637         (type 1), except that a time stamp is not used.

640         *,hex

642         Any other type value (0 or 4-65535) can be used

```

643 with an even-length hexadecimal string.

645 The second and third forms of CLIENT_ID are legal
 646 for IPv4 only. These both represent raw Client ID
 647 (without RFC 4361), in hex, or NVT ASCII string for-
 648 mat. Thus, Sun and 0x53756E are equivalent.

655 SunOS 5.11 Last change: 15 May 2008 10

662 System Administration Commands dhcpageant(1M)

666 PARAM_REQUEST_LIST

668 Specifies a list of comma-separated integer values
 669 of options for which the client would like values.

672 REQUEST_HOSTNAME

674 Indicates the client requests the DHCP server to map
 675 the client's leased IPv4 address to the host name
 676 associated with the network interface that performs
 677 DHCP on the client. The host name must be specified
 678 in the /etc/hostname.interface file for the relevant
 679 interface on a line of the form

681 inet hostname

684 where hostname is the host name requested.

686 This option works with DHCPv4 only.

690 /etc/dhcp/eventhook

692 Location of a DHCP event program.

695 ATTRIBUTES

696 See attributes(5) for descriptions of the following attri-
 697 butes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsr
Interface Stability	Committed

710 SEE ALSO

711 dhcpageant(1), ifconfig(1M), init(1M), in.ndpd(1M),
 712 syslog(3C), attributes(5), dhcp(5)

715 Croft, B. and Gilmore, J., Bootstrap Protocol (BOOTP) RFC 951,
 716 Network Working Group, September 1985.

721 SunOS 5.11 Last change: 15 May 2008 11

728 System Administration Commands dhcpageant(1M)

732 Droms, R., Dynamic Host Configuration Protocol, RFC 2131,
 733 Network Working Group, March 1997.

736 Lemon, T. and B. Sommerfeld. RFC 4361, Node-specific Client
 737 Identifiers for Dynamic Host Configuration Protocol Version
 738 Four (DHCPv4). Nominum and Sun Microsystems. February 2006.

741 Droms, R. RFC 3315, Dynamic Host Configuration Protocol for
 742 IPv6 (DHCPv6). Cisco Systems. July 2003.

744 NOTES

745 DHCP can be performed on IPv4 logical interfaces just as
 746 with physical interfaces. When used on a logical interface,
 747 the daemon automatically constructs a Client ID value based
 748 on the DUID and IAID values, according to RFC 4361. The
 749 */etc/default/dhcpclient* *CLIENT_ID* value, if any,
 744 The dhcpageant daemon can be used on IPv4 logical interfaces,
 745 just as with physical interfaces. When used on a logical
 746 interface, the daemon automatically constructs a Client ID
 747 value based on the DUID and IAID values, according to RFC
 748 4361. The /etc/default/dhcpclient CLIENT_ID value, if any,
 750 overrides this automatic identifier.

753 As with physical IPv4 interfaces, the /etc/hostname.hme0:1
 754 and /etc/dhcp.hme0:1 files must also be created in order for
 755 hme0:1 to be automatically plumbed and configured at boot.
 756 In addition, unlike physical IPv4 interfaces, dhcpageant does
 757 not add or remove default routes associated with logical
 758 interfaces.

761 DHCP can be performed on IPMP IP interfaces to acquire and
 762 maintain IPMP data addresses. Because an IPMP IP interface
 763 has no hardware address, the daemon automatically constructs
 764 a Client ID using the same approach described above for IPv4
 765 logical interfaces. In addition, the lack of a hardware
 766 address means the daemon must set the "broadcast" flag in
 767 all *DISCOVER* and *REQUEST* messages on IPMP IP interfaces.
 768 Some DHCP servers may refuse such requests.

```
771 DHCP can be performed on IP interfaces that are part of an
772 IPMP group (to acquire and maintain test addresses). The
773 daemon will automatically set the *NOFAILOVER* and
774 *DEPRECATED* flags on each test address. Additionally, the
775 daemon will not add or remove default routes in this case.
776 Note that the actual DHCP packet exchange may be performed
777 over any active IP interface in the IPMP group. It is
778 strongly recommended that test addresses have infinite
779 leases. Otherwise, an extended network outage detectable
780 only by probes may cause test address leases to expire,
781 causing *in.mpathd(1M)* to revert to link-based failure
782 detection and trigger an erroneous repair.
```

```
785 With DHCPv6, the link-local interface must be configured
786 using /etc/hostname6.hme0 in order for DHCPv6 to run on hme0
787 at boot time. The logical interfaces for each address are
788 plumbed by dhcpageant automatically.
```

14328 Fri Dec 19 04:18:47 2008

new/./getsockopt.3socket.txt

Clearview IPMP manpages

1 Sockets Library Functions getsockopt(3SOCKET)

5 NAME

6 getsockopt, setsockopt - get and set options on sockets

8 SYNOPSIS

9 cc [flag ...] file ... -lsocket -lnsl [library ...]

10 #include <sys/types.h>

11 #include <sys/socket.h>

13 int getsockopt(int s, int level, int optname, void *optval,
14 int *optlen);

17 int setsockopt(int s, int level, int optname, const void *optval,
18 int optlen);

21 DESCRIPTION

22 The getsockopt() and setsockopt() functions manipulate
23 options associated with a socket. Options may exist at mul-
24 tiple protocol levels; they are always present at the upper-
25 most "socket" level.

28 The **level** argument specifies the protocol level at which
29 the option resides. To manipulate options at the socket
30 level, specify the **level** argument as **SOL_SOCKET**. To
31 manipulate options at the protocol level, supply the
32 appropriate protocol number for the protocol controlling the
33 option. For example, to indicate that an option will be
34 interpreted by TCP, set **level** to the protocol number of
35 TCP, as defined in the **<netinet/in.h>** header, or as
36 determined by using **getprotobyname(3SOCKET)**. Some socket
37 protocol families may also define additional levels, such as
38 **SOL_ROUTE**. Only socket-level options are described here.
28 When manipulating socket options, the level at which the
29 option resides and the name of the option must be specified.
30 To manipulate options at the "socket" level, level is speci-
31 fied as *SOL_SOCKET*. To manipulate options at any other
32 level, level is the protocol number of the protocol that
33 controls the option. For example, to indicate that an option
34 is to be interpreted by the TCP protocol, level is set to
35 the TCP protocol number. See *getprotobyname(3SOCKET)*.

41 The parameters *optval* and *optlen* are used to access option
42 values for setsockopt(). For getsockopt(), they identify a
43 buffer in which the value(s) for the requested option(s) are
44 to be returned. For getsockopt(), *optlen* is a value-result
45 parameter, initially containing the size of the buffer
46 pointed to by *optval*, and modified on return to indicate the
47 actual size of the value returned. Use a 0 *optval* if no
48 option value is to be supplied or returned.

51 The *optname* and any specified options are passed uninter-
52 preted to the appropriate protocol module for interpreta-
53 tion. The include file <sys/socket.h> contains definitions

54 for the socket-level options described below. Options at
55 other protocol levels vary in format and name.

58 Most socket-level options take an int for *optval*. For set-
59 sockopt(), the *optval* parameter should be non-zero to enable

63 SunOS 5.11

Last change: 7 Aug 2008

1

70 Sockets Library Functions

getsockopt(3SOCKET)

74 a boolean option, or zero if the option is to be disabled.
75 *SO_LINGER* uses a struct *linger* parameter that specifies the
76 desired state of the option and the linger interval. struct
77 *linger* is defined in <sys/socket.h>. struct *linger* contains
78 the following members:

80 *l_onoff* on = 1/off = 0

83 *l_linger* linger time, in seconds

87 The following options are recognized at the socket level.
88 Except as noted, each may be examined with getsockopt() and
89 set with setsockopt().

91 *SO_DEBUG* enable/disable recording of debugging
92 information

95 *SO_REUSEADDR* enable/disable local address reuse

98 *SO_KEEPAIVE* enable/disable keep connections alive

101 *SO_DONTROUTE* enable/disable routing bypass for outgo-
102 ing messages

105 *SO_LINGER* linger on close if data is present

108 *SO_BROADCAST* enable/disable permission to transmit
109 broadcast messages

112 *SO_OOBINLINE* enable/disable reception of out-of-band
113 data in band

116 *SO_SNDBUF* set buffer size for output

119 *SO_RCVBUF* set buffer size for input

122 SO_DGRAM_ERRIND application wants delayed error

129 SunOS 5.11 Last change: 7 Aug 2008 2

136 Sockets Library Functions getsockopt(3SOCKET)

140 SO_TIMESTAMP enable/disable reception of timestamp
141 with datagrams

144 SO_EXCLBIND enable/disable exclusive binding of the
145 socket

148 SO_TYPE get the type of the socket (get only)

151 SO_ERROR get and clear error on the socket (get
152 only)

155 SO_MAC_EXEMPT get or set mandatory access control on
156 the socket. This option is available only
157 when the system is configured with
158 Trusted Extensions.

161 SO_ALLZONES bypass zone boundaries (privileged).

164 SO_DOMAIN get the domain used in the socket (get
165 only)

168 SO_PROTOTYPE for socket in domains PF_INET and
169 PF_INET6, get the underlying protocol
170 number used in the socket. For socket in
171 domain PF_ROUTE, get the address family
172 used in the socket.

176 The SO_DEBUG option enables debugging in the underlying pro-
177 tocol modules. The SO_REUSEADDR option indicates that the
178 rules used in validating addresses supplied in a
179 bind(3SOCKET) call should allow reuse of local addresses.
180 The SO_KEEPAIVE option enables the periodic transmission of
181 messages on a connected socket. If the connected party fails
182 to respond to these messages, the connection is considered
183 broken and threads using the socket are notified using a
184 SIGPIPE signal. The SO_DONTROUTE option indicates that out-
185 going messages should bypass the standard routing facili-

186 ties. Instead, messages are directed to the appropriate net-
187 work interface according to the network portion of the des-
188 tination address.

195 SunOS 5.11 Last change: 7 Aug 2008 3

202 Sockets Library Functions getsockopt(3SOCKET)

206 The SO_LINGER option controls the action taken when unsent
207 messages are queued on a socket and a close(2) is performed.
208 If the socket promises reliable delivery of data and
209 SO_LINGER is set, the system will block the thread on the
210 close() attempt until it is able to transmit the data or
211 until it decides it is unable to deliver the information (a
212 timeout period, termed the linger interval, is specified in
213 the setsockopt() call when SO_LINGER is requested). If
214 SO_LINGER is disabled and a close() is issued, the system
215 will process the close() in a manner that allows the thread
216 to continue as quickly as possible.

219 The option SO_BROADCAST requests permission to send broad-
220 cast datagrams on the socket. With protocols that support
221 out-of-band data, the SO_OOBINLINE option requests that
222 out-of-band data be placed in the normal data input queue as
223 received; it will then be accessible with recv() or read()
224 calls without the MSG_OOB flag.

227 The SO_SNDBUF and SO_RCVBUF options adjust the normal buffer
228 sizes allocated for output and input buffers, respectively.
229 The buffer size may be increased for high-volume connections
230 or may be decreased to limit the possible backlog of incom-
231 ing data. The maximum buffer size for UDP is determined by
232 the value of the ndd variable udp_max_buf. The maximum
233 buffer size for TCP is determined the value of the ndd vari-
234 able tcp_max_buf. Use the ndd(lm) utility to determine the
235 current default values. See the Solaris Tunable Parameters
236 Reference Manual for information on setting the values of
237 udp_max_buf and tcp_max_buf. At present, lowering SO_RCVBUF
238 on a TCP connection after it has been established has no
239 effect.

242 By default, delayed errors (such as ICMP port unreachable
243 packets) are returned only for connected datagram sockets.
244 The SO_DGRAM_ERRIND option makes it possible to receive
245 errors for datagram sockets that are not connected. When
246 this option is set, certain delayed errors received after
247 completion of a sendto() or sendmsg() operation will cause a
248 subsequent sendto() or sendmsg() operation using the same
249 destination address (to parameter) to fail with the
250 appropriate error. See send(3SOCKET).

253 If the SO_TIMESTAMP option is enabled on a SO_DGRAM or a
 254 SO_RAW socket, the recvmsg(3XNET) call will return a times-
 255 tamp in the native data format, corresponding to when the
 256 datagram was received.

261 SunOS 5.11 Last change: 7 Aug 2008 4

268 Sockets Library Functions getsockopt(3SOCKET)

272 The SO_EXCLBIND option is used to enable or disable the
 273 exclusive binding of a socket. It overrides the use of the
 274 SO_REUSEADDR option to reuse an address on bind(3SOCKET).
 275 The actual semantics of the SO_EXCLBIND option depend on the
 276 underlying protocol. See tcp(7P) or udp(7P) for more infor-
 277 mation.

280 The SO_TYPE and SO_ERROR options are used only with get-
 281 sockopt(). The SO_TYPE option returns the type of the
 282 socket, for example, SOCK_STREAM. It is useful for servers
 283 that inherit sockets on startup. The SO_ERROR option
 284 returns any pending error on the socket and clears the error
 285 status. It may be used to check for asynchronous errors on
 286 connected datagram sockets or for other asynchronous errors.

289 The SO_MAC_EXEMPT option is used to toggle socket behavior
 290 with unlabeled peers. A socket that has this option enabled
 291 can communicate with an unlabeled peer if it is in the glo-
 292 bal zone or has a label that dominates the default label of
 293 the peer. Otherwise, the socket must have a label that is
 294 equal to the default label of the unlabeled peer. Calling
 295 setsockopt() with this option returns an EACCES error if the
 296 process lacks the NET_MAC_AWARE privilege or if the socket
 297 is bound. The SO_MAC_EXEMPT option is available only when
 298 the system is configured with Trusted Extensions.

301 The SO_ALLZONES option can be used to bypass zone boundaries
 302 between shared-IP zones. Normally, the system prevents a
 303 socket from being bound to an address that is not assigned
 304 to the current zone. It also prevents a socket that is bound
 305 to a wildcard address from receiving traffic for other
 306 zones. However, some daemons which run in the global zone
 307 might need to send and receive traffic using addresses that
 308 belong to other shared-IP zones. If set before a socket is
 309 bound, SO_ALLZONES causes the socket to ignore zone boun-
 310 daries between shared-IP zones and permits the socket to be
 311 bound to any address assigned to the shared-IP zones. If the
 312 socket is bound to a wildcard address, it receives traffic
 313 intended for all shared-IP zones and behaves as if an
 314 equivalent socket were bound in each active shared-IP zone.
 315 Applications that use the SO_ALLZONES option to initiate
 316 connections or send datagram traffic should specify the
 317 source address for outbound traffic by binding to a specific

318 address. There is no effect from setting this option in an
 319 exclusive-IP zone. Setting this option requires the
 320 sys_net_config privilege. See zones(5).

322 RETURN VALUES

327 SunOS 5.11 Last change: 7 Aug 2008 5

334 Sockets Library Functions getsockopt(3SOCKET)

338 If successful, getsockopt() and setsockopt() return 0. Oth-
 339 erwise, the functions return -1 and set errno to indicate
 340 the error.

342 ERRORS

343 The getsockopt() and setsockopt() calls succeed unless:

345 EBADF The argument s is not a valid file descrip-
 346 tor.

349 ENOMEM There was insufficient memory available for
 350 the operation to complete.

353 ENOPROTOPT The option is unknown at the level indi-
 354 cated.

357 ENOSR There were insufficient STREAMS resources
 358 available for the operation to complete.

361 ENOTSOCK The argument s is not a socket.

364 ENOBUFS SO_SNDBUF or SO_RCVBUF exceeds a system
 365 limit.

368 EINVAL Invalid length for IP_OPTIONS.

371 EHOSTUNREACH Invalid address for IP_MULTICAST_IF.

374 EINVAL Not a multicast address for
 375 IP_ADD_MEMBERSHIP and IP_DROP_MEMBERSHIP.

378 EADDRNOTAVAIL Bad interface address for IP_ADD_MEMBERSHIP
 379 and IP_DROP_MEMBERSHIP.

382 EADDRINUSE Address already joined for
 383 IP_ADD_MEMBERSHIP.

386 ENOENT Address not joined for IP_DROP_MEMBERSHIP.

393 SunOS 5.11 Last change: 7 Aug 2008 6

400 Sockets Library Functions getsockopt(3SOCKET)

404 EPERM No permissions.

407 EACCES Permission denied.

410 EINVAL The specified option is invalid at the
411 specified socket level, or the socket has
412 been shut down.

415 ATTRIBUTES
416 See attributes(5) for descriptions of the following attri-
417 butes:

421

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

422
423
424
425

428 SEE ALSO
429 ndd(1M), close(2), ioctl(2), read(2), bind(3SOCKET),
430 getprotobyname(3SOCKET), recv(3SOCKET), recvmsg(3XNET),
431 send(3SOCKET), socket(3SOCKET), socket.h(3HEAD), attri-
432 butes(5), zones(5), tcp(7P), udp(7P)

435 Solaris Tunable Parameters Reference Manual

459 SunOS 5.11 Last change: 7 Aug 2008 7

new./if_mpadm.txt

1

```
*****
4169 Fri Dec 19 04:18:48 2008
new./if_mpadm.txt
Clearview IPMP manpages
*****
1 System Administration Commands          if_mpadm(1M)

5 NAME
6   if_mpadm - administer interfaces in an IP multipathing group
6   if_mpadm - change operational status of interfaces within a
7   multipathing group

8 SYNOPSIS
9   if_mpadm -d | -r ifname
10  /usr/sbin/if_mpadm -d interface_name

13  /usr/sbin/if_mpadm -r interface_name

11 DESCRIPTION
12  The if_mpadm utility administers IP interfaces that are
13  part of an IP multipathing (IPMP) group. Currently,
14  administration is limited to offlining IP interfaces and
15  undoing previous offline operations.
16  Use the if_mpadm utility to change the operational status of
17  interfaces that are part of an IP multipathing group. If the
18  interface is operational, you can use if_mpadm -d to detach
19  or off-line the interface. If the interface is off-lined,
20  use if_mpadm -r to revert it to its original state.

17  When an IP interface is taken offline, all IP data traffic
18  that was flowing over the IP interface is moved to another
19  IP interface in the IPMP group. In addition, all *UP* IP
20  addresses hosted on the IP interface are brought down,
21  causing *in.mpathd(1M)* to stop probe-based failure
22  detection on the IP interface. As a result, an offline IP
23  interface will not be used for any inbound or outbound IP
24  traffic. Only IP interfaces that are in an IPMP group may
25  be brought offline. If the IP interface is the last
26  functioning interface in the IPMP group, the offline
27  operation will fail.

28
29  When an offline operation is undone, any IP addresses hosted
30  on that IP interface are brought *UP* and will be considered
31  by *in.mpathd* for probe-based failure detection. In
32  addition, provided the IP interface is otherwise active (see
33  *in.mpathd(1M)*), it will again be used to send and receive
34  IP data traffic for the IPMP group. Note that not all
35  offline operations can be undone. For instance, *in.mpathd*
36  may have offlined an IP interface because its hardware
37  address was not unique within its IPMP group. The
38  *ipmpstat* utility can be used to determine why an IP
39  interface is offline, identify which IP interfaces in a
40  group are being used for inbound and outbound IP traffic,
41  and more; see *ipmpstat(1M)*.
42  When a network interface is off-lined, all network access
43  fails over to a different interface in the IP multipathing
44  group. Any addresses that do not failover are brought down.
45  Network access includes unicast, broadcast, and multicast
46  for IPv4 and unicast and multicast for IPv6. Addresses
47  marked with IFF_NOFAILOVER do not failover. They are marked
48  down. After an interface is off-lined, the system will not
49  use the interface for any outbound or inbound traffic, and
```

new./if_mpadm.txt

2

```
32  the interface can be safely removed from the system without
33  any loss of network access.

36  The if_mpadm utility can be applied only to interfaces that
37  are part of an IP multipathing group.

43 OPTIONS
44  The *if_mpadm* utility supports the following options:
45  The if_mpadm utility supports the following options:

46  -d ifname          Offline the IP interface specified by
47                    *ifname*. If *ifname* is not in an
48                    IPMP group, or the offline would cause
49                    the IPMP group to lose network
50                    connectivity, the operation will fail.
42  -d interface_name Detach or off-line the interface speci-
43                    fied by interface_name.

52  -r ifname          Undo a previous offline of the IP
53                    interface specified by *ifname*. If
54                    *ifname* is not offline, the operation
55                    will fail.

46  -r interface_name Reattach or undo the previous detach
47                    or off-line operation on the interface
48                    specified by interface_name. Unless the
49                    -d option was used to detach or off-
50                    line the interface, this option will
51                    fail.

58 EXAMPLES
59  Example 1 Offlining an IP Interface
55  Example 1 Detaching an Interface

61  The following command offlines IP interface *under0*,
62  causing any IP packets that were being sent and received
63  through it to be handled by another IP interface in its
64  group.

66  example% if_mpadm -d under0

69  Example 2 Undoing a Previous Offline Operation
60 SunOS 5.11          Last change: 3 Sep 2002          1

72  Use the following command to undo the previous operation:

75  example% if_mpadm -r under0

67 System Administration Commands          if_mpadm(1M)

71  Use the following command to off-line or detach the inter-
72  face. All network access will failover from hme0 to other
73  interfaces in the same IP multipathing group. If no other
74  interfaces are in the same group, the operation will fail.

77  example% if_mpadm -d hme0
```

81 *Example 2 Reattaching an Off-line Interface*

84 *Use the following command to undo the previous operation.*
85 *Network access will fallback to hme0.*

88 *example% if_mpadm -r hme0*

79 ATTRIBUTES

80 See attributes(5) for descriptions of the following attri-
81 butes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu
Interface Stability	Unstable

94 SEE ALSO

95 ifconfig(1M), in.mpathd(1M), attributes(5)

97 DIAGNOSTICS

98 **cannot offline: no other functioning interfaces are in its**
99 **IPMP group.**

111 *off-line failed as there is no other functional interface*
112 *available in the multipathing group for failing over the*
113 *network access.*

101 Description:

103 **This message means that offlining the IP interface would**
104 **leave the IPMP group without network connectivity.**

106 **cannot offline: not a physical interface or not in an IPMP**
107 **group**

117 *This message means that other interfaces in the group*
118 *are failed over already or the multipathing configura-*
119 *tion was not suitable for completing a failover.*

126 SunOS 5.11 Last change: 3 Sep 2002 2

133 System Administration Commands if_mpadm(1M)

137 *off-line cannot be undone because multipathing configuration*

138 *is not consistent across all the interfaces in the group.*
109 Description:

111 **This means that the IP interface is not an underlying**
112 **interface in an IPMP group, and therefore is not**
113 **eligible to be offlined.**

142 *This message means that some interfaces in the IP mul-*
143 *tipathing group are not configured consistently with*
144 *other interfaces in the group, for example, one of the*
145 *interfaces in the group does not have an IFF_NOFAILOVER*
146 *address.*

192 SunOS 5.11 Last change: 3 Sep 2002 3

```

*****
5836 Fri Dec 19 04:18:50 2008
new./if_nameindex.txt
Clearview IPMP manpages
*****
1 Sockets Library Functions          if_nametoindex(3SOCKET)

5 NAME
6   if_nametoindex,          if_inde
7   if_freenameindex - routines to map Internet Protocol network
8   interface names and interface indexes

10 SYNOPSIS
11 cc [ flag... ] file... -lsocket [ library... ]
12 #include <net/if.h>

14 unsigned int if_nametoindex(const char *ifname);

17 char *if_inde
20 struct if_nameindex *if_nameindex(void)

23 void if_freenameindex(struct if_nameindex *ptr);

26 PARAMETERS
27   ifname      interface name

30   ifindex     interface index

33   ptr         pointer returned by if_nameindex()

36 DESCRIPTION
37 This API defines two functions that map between an Internet
38 Protocol network interface name and index, a third function
39 that returns all the interface names and indexes, and a
40 fourth function to return the dynamic memory allocated by
41 the previous function.

44 Network interfaces are normally known by names such as eri0,
45 s11, ppp2, and the like. The ifname argument must point to
46 a buffer of at least IF_NAMESIZE bytes into which the inter-
47 face name corresponding to the specified index is returned.
48 IF_NAMESIZE is defined in <net/if.h> and its value includes
49 a terminating null byte at the end of the interface name.

51 if_nametoindex()      The if_nametoindex() function returns
52 the interface index corresponding to
53 the interface name pointed to by the
54 ifname pointer. If the specified
55 interface name does not exist, the
56 return value is 0, and errno is set to

60 SunOS 5.11          Last change: 12 Dec 2003          1

```

```

67 Sockets Library Functions          if_nametoindex(3SOCKET)

71 ENXIO. If there was a system error,
72 such as running out of memory, the
73 return value is 0 and errno is set to
74 the proper value, for example, ENOMEM.

77   if_inde
78   if_inde
79   if_inde
80   if_inde
81   if_inde
82   if_inde
83   if_inde
84   if_inde
85   if_inde
86   if_inde
87   if_inde

89
90   if_nameindex()      The if_nameindex() function returns an
91   *if_nameindex()    The if_nameindex() function returns an
92 array of if_nameindex structures, one
93 structure per interface. The
94 if_nameindex structure holds the
95 information about a single interface
96 and is defined when the <net/if.h>
header is included:

98   struct if_nameindex
99   unsigned int if_index; /* 1, 2, ... */
100   char *if_name; /* "net0", ... */
101   char *if_name; /* null terminated name:
};

103 While any IPMP IP interfaces are
104 returned by *if_nameindex()*, the
105 underlying IP interfaces that comprise
106 each IPMP group are not returned.

108 The end of the array of structures is
109 indicated by a structure with an
110 if_index of 0 and an if_name of NULL.
111 The function returns a null pointer
112 upon an error and sets errno to the
113 appropriate value. The memory used for
114 this array of structures along with
115 the interface names pointed to by the
116 if_name members is obtained dynami-
117 cally. This memory is freed by the
118 if_freenameindex() function.

121   if_freenameindex() The if_freenameindex() function frees
122 the dynamic memory that was allocated
123 by if_nameindex(). The argument to
124 this function must be a pointer that
125 was returned by if_nameindex().

```

131 SunOS 5.11 Last change: 12 Dec 2003 2

138 Sockets Library Functions if_nametoindex(3SOCKET)

142 ATTRIBUTES
143 See attributes(5) for descriptions of the following attri-
144 butes:

148
149
150
151
152
153
154
155

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcs1 (32-bit) SUNWcs1x (64-bit)
MT Level	MT Safe

158 SEE ALSO
159 ifconfig(1M), if_nametoindex(3XNET), attributes(5), if(7P)

197 SunOS 5.11 Last change: 12 Dec 2003 3

```
*****
```

```
27477 Fri Dec 19 04:18:51 2008
```

```
new./if_tcp.txt
```

```
Clearview IPMP manpages
```

```
*****
```

```
1 Protocols if_tcp(7P)

3 NAME
4 if_tcp, if - general properties of Internet Protocol network
5 interfaces

7 DESCRIPTION
8 A network interface is a device for sending and receiving
9 packets on a network. It is usually a hardware device,
10 although it can be implemented in software. Network inter-
11 faces used by the Internet Protocol (IPv4 or IPv6) must be
12 STREAMS devices conforming to the Data Link Provider Inter-
13 face (DLPI). See dlpi(7P).

15 APPLICATION PROGRAMMING INTERFACE
16 An interface becomes available to IP when it is opened and
17 the IP module is pushed onto the stream with the I_PUSH
18 ioctl(2) command. (See streamio(7I)). The SIOCCLIFNAME
19 ioctl(2) is issued to specify the name of the interface and
20 to indicate whether it is IPv4 or IPv6. This may be ini-
21 tiated by the kernel at boot time or by a user program after
22 the system is running. Each interface must be assigned an IP
23 address with the SIOCCLIFADDR ioctl() before it can be used.
24 On interfaces where the network-to-link layer address map-
25 ping is static, only the network number is taken from the
26 ioctl() request; the remainder is found in a hardware
27 specific manner. On interfaces which provide dynamic
28 network-to-link layer address mapping facilities (for exam-
29 ple, 10Mb/s Ethernets using arp(7P)), the entire address
30 specified in the ioctl() is used. A routing table entry for
31 destinations on the network of the interface is installed
32 automatically when an interface's address is set.

34 Note IPMP IP interfaces cannot be created using the
35 procedure described above. Instead, ifconfig(1M) must be
36 used to create IPMP IP interfaces.

38 IOCTLS
39 The following ioctl() calls may be used to manipulate IP
40 network interfaces. Unless specified otherwise, the request
41 takes an lifreq structure as its parameter. This structure
42 has the form:

44 struct lifreq {
45 #define LIFNAMSIZ 32
46 char lifr_name[LIFNAMSIZ]; /* if name, e.g. "le1" */
47 union {
48 int lifru_addrflen; /* for subnet/token etc */
49 uint_t lifru_ppa; /* SIOCCLIFNAME */
50 } lifr_lifru;
51 union {
52 struct sockaddr_storage lifru_addr;
53 struct sockaddr_storage lifru_dstaddr;
54 struct sockaddr_storage lifru_broadaddr;
55 struct sockaddr_storage lifru_token; /* With lifr_addrflen */
56 struct sockaddr_storage lifru_subnet; /* With lifr_addrflen */
57 int lifru_index; /* interface index */
58 uint64_t lifru_flags; /* SIOC?LIFFLAGS */
59 }
```

```
62 Protocols if_tcp(7P)

64 int lifru_metric;
65 uint_t lifru_mtu;
66 int lifr_muxid[2]; /* mux id's for arp & ip */
67 struct lif_nd_req lifru_nd_req;
68 struct lif_ifinfo_req lifru_ifinfo_req;
69 zoneid_t lifru_zone; /* SIOC[GS]LIFZONE */
70 } lifr_lifru;

72 #define lifr_addrflen lifr_lifru.lifru_addrflen
73 #define lifr_ppa lifr_lifru.lifru_ppa /* Driver's ppa */
74 #define lifr_addr lifr_lifru.lifru_addr /* address */
75 #define lifr_dstaddr lifr_lifru.lifru_dstaddr
76 #define lifr_broadaddr lifr_lifru.lifru_broadaddr /* broadcast addr. */
77 #define lifr_token lifr_lifru.lifru_token /* address token */
78 #define lifr_subnet lifr_lifru.lifru_subnet /* subnet prefix */
79 #define lifr_index lifr_lifru.lifru_index /* interface index */
80 #define lifr_flags lifr_lifru.lifru_flags /* flags */
81 #define lifr_metric lifr_lifru.lifru_metric /* metric */
82 #define lifr_mtu lifr_lifru.lifru_mtu /* mtu */
83 #define lifr_ip_muxid lifr_lifru.lif_muxid[0]
84 #define lifr_arp_muxid lifr_lifru.lif_muxid[1]
85 #define lifr_nd lifr_lifru.lifru_nd_req /* SIOCCLIF*ND */
86 #define lifr_ifinfo lifr_lifru.lifru_ifinfo_req /* SIOC[GS]LIFLNKINFO */
87 #define lifr_zone lifr_lifru.lifru_zone /* SIOC[GS]LIFZONE */
88 };

90 SIOCCLIFADDR Set interface address.
92 SIOCGLIFADDR Get interface address.
94 SIOCCLIFDSTADDR Set point to point address for inter-
95 face.
97 SIOCGLIFDSTADDR Get point to point address for inter-
98 face.
100 SIOCCLIFFLAGS Set interface flags field. If the inter-
101 face is marked down, any processes
102 currently routing packets through the
103 interface are notified.

105 SIOCGLIFFLAGS Get interface flags.

107 SIOCGLIFCONF Get interface configuration list. This
108 request takes a *lifconf* structure (see
109 below) as a value-result parameter. The
110 *lifc_family* field can be set to
111 *AF_UNSPEC* to retrieve both *AF_INET*
112 and *AF_INET6* interfaces. The
113 *lifc_len* field should be set to the
114 size of the buffer pointed to by
115 *lifc_buf*.
116 request takes an lifconf structure (see
117 below) as a value-result parameter. The
118 *lifc_flags* field should usually be
119 set to zero, but callers that need low-
120 level knowledge of the underlying IP
121 interfaces that comprise an IPMP group
122 can set it to *LIFC_UNDER_IPMP* to
123 request that those interfaces be
124 included in the result.
125 Upon success, *lifc_len* will contain

106 SunOS 5.10 Last change: 14 Nov 2008 2
```

126 the length, in bytes, of the array of
 127 *lifreq* structures pointed to by
 128 *lifc_req*. For each *lifreq* structure,
 129 the *lifr_name* and *lifr_addr* fields
 130 will be valid.

108 Protocols if_tcp(7P)

110 below) as a value-result parameter. The
 111 lifc_family field can be set to
 112 AF_UNSPEC to retrieve both AF_INET and
 113 AF_INET6 interfaces. The lifc_flags
 114 field should be set to zero. The
 115 lifc_len field should be set to the size
 116 of the buffer pointed to by lifc_buf.
 117 Upon success, lifc_len will contain the
 118 length, in bytes, of the array of lifreq
 119 structures pointed to by lifc_req. For
 120 each lifreq structure, the lifr_name and
 121 lifr_addr fields will be valid.

133 SIOCGIFNUM Get number of interfaces. This request
 134 returns an integer which is the number
 135 of interface descriptions (struct
 136 lifreq) that will be returned by the
 137 SIOCGIFCONF ioctl; that is, it gives an
 138 indication of how large lifc_len has
 139 to be.
 128 indication of how large lifc_len has to
 129 be. This request takes an lifnum struc-
 130 ture (see below) as a value-result
 131 parameter. The lifn_family field should
 132 be set to AF_UNSPEC to count both
 133 AF_INET and AF_INET6 interfaces. The
 134 lifn_flags field should be initially set
 135 to zero.

141 This request takes a *struct lifnum*
 142 (see below) as a value-result parameter.
 143 The *lifn_family* field can be set to
 144 *AF_UNSPEC* to count both *AF_INET* and
 145 *AF_INET6* interfaces. The *lifn_flags*
 146 field should usually be set to zero, but
 147 callers that need low-level knowledge of
 148 the underlying IP interfaces that
 149 comprise an IPMP group can set it to
 150 *LIFC_UNDER_IPMP* to request that those
 151 interfaces be included in the count.

154 SIOCSLIFMTU Set the maximum transmission unit (MTU)
 155 size for interface. Place the request in
 156 the lifru_mtu field. The MTU can not
 157 exceed the physical MTU limitation
 158 (which is reported in the DLPI
 159 DL_INFO_ACK message).

161 SIOCGIFMTU Get the maximum transmission unit size
 162 for interface.

164 SIOCSLIFMETRIC Set the metric associated with the
 165 interface. The metric is used by rout-
 166 ing daemons such as in.routed(1M).

168 SIOCGIFMETRIC Get the metric associated with the
 169 interface.

171 SIOCGIFMUXID Get the ip and arp muxid associated with
 172 the interface.

174 SunOS 5.10 Last change: 14 Nov 2008 3

176 Protocols if_tcp(7P)

178 SIOCSLIFMUXID Set the ip and arp muxid associated with
 179 the interface.

181 SIOCGIFINDEX Get the interface index associated with
 182 the interface.

184 SIOCSLIFINDEX Set the interface index associated with
 185 the interface.

187 SIOCGIFZONE Get the zone associated with the inter-
 188 face.

190 SIOCSLIFZONE Set the zone associated with the inter-
 191 face. Only applies for zones that use
 192 the shared-IP instance.

194 SIOCLIFADDIF Add a new logical interface on a physi-
 195 cal interface using an unused logical
 196 interface number.
 179 unit number. If the physical interface
 180 is part of an IP multipathing group, the
 181 logical interface may be added to a dif-
 182 ferent physical interface in the same
 183 group. Upon return, the lifr_name field
 184 contains the name of the actual logical
 185 interface created.

198 SIOCLIFREMOVEIF Remove a logical interface by specifying
 199 its IP address or logical interface
 200 name.
 189 name. When the IP address is specified
 190 and the interface is part of an IP mul-
 191 tipathing group, the logical interface
 192 is removed from the physical interface
 193 in the group which holds the IP address.

202 SIOCSLIFTOKEN Set the address token used to form IPv6
 203 link-local addresses and for stateless
 204 address autoconfiguration.

206 SIOCGLIFTOKEN Get the address token used to form IPv6
 207 link-local addresses and for stateless
 208 address autoconfiguration.

210 SunOS 5.10 Last change: 14 Nov 2008 4

212 Protocols if_tcp(7P)

214 SIOCSLIFSUBNET Set the subnet prefix associated with
 215 the interface.

217 SIOCGLIFSUBNET Get the subnet prefix associated with
 218 the interface.

220 SIOCSLIFLNKINFO Set link specific parameters for the
 221 interface.

223 SIOCGLIFLNKINFO Get link specific parameters for the
 224 interface.

```

226 SIOCLIFDELND Delete a neighbor cache entry for IPv6.
228 SIOCLIFGETND Get a neighbor cache entry for IPv6.
230 SIOCLIFSETND Set a neighbor cache entry for IPv6.
232 SIOCCLIFUSESRC Set the interface from which to choose a
233 source address. The lifr_index field has
234 the interface index corresponding to the
235 interface whose address is to be used as
236 the source address for packets going out
237 on the interface whose name is provided
238 by lifr_name. If the lifr_index field is
239 set to zero, the previous setting is
240 cleared. See ifconfig(1M) for examples
241 of the usesrc option.
243 SIOCGLIFUSESRC Get the interface index of the interface
244 whose address is used as the source
245 address for packets going out on the
246 interface provided by lifr_name field.
247 The value is retrieved in the lifr_index
248 field. See ifconfig(1M) for examples of
249 the usesrc option.
251 SIOCGLIFSRCOF Get the interface configuration list for
252 interfaces that use an address hosted on
253 the interface provided by the
254 lifs_ifindex field in the lifsrcof
255 struct (see below), as a source address.
256 The application sets lifs_maxlen to the
258 SunOS 5.10 Last change: 14 Nov 2008 5
260 Protocols if_tcp(7P)
262 size (in bytes) of the buffer it has
263 allocated for the data. On return, the
264 kernel sets lifs_len to the actual size
265 required. Note, the application could
266 set lifs_maxlen to zero to query the
267 kernel of the required buffer size
268 instead of estimating a buffer size. The
269 application tests lifs_len <=
270 lifs_maxlen -- if that's true, the
271 buffer was big enough and the applica-
272 tion has an accurate list. If it is
273 false, it needs to allocate a bigger
274 buffer and try again, and lifs_len pro-
275 vides a hint of how big to make the next
276 trial. See ifconfig(1M) for examples of
277 the usesrc option.
279 SIOCTONLINK Test if the address is directly reach-
280 able, for example, that it can be
281 reached without going through a router.
282 This request takes an sioc_addrreq
283 structure (see below) as a value-result
284 parameter. The sa_addr field should be
285 set to the address to test. The sa_res
286 field will contain a non-zero value if
287 the address is onlink.
289 SIOCTMYADDR Test if the address is assigned to this
290 node. This request takes an sioc_addrreq

```

```

291 structure (see below) as a value-result
292 parameter. The sa_addr field should be
293 set to the address to test. The sa_res
294 field will contain a non-zero value if
295 the address is assigned to this node.
297 SIOCTMYSITE Test if the address is part of the same
298 site as this node. This request takes an
299 sioc_addrreq structure (see below) as a
300 value-result parameter. The sa_addr
301 field should be set to the address to
302 test. The sa_res field will contain a
303 non-zero value if the address is in the
304 same site.
306 The structure used by SIOCGLIFCONF has the form:
308 struct lifconf {
309     sa_family_t    lifc_family;
311 SunOS 5.10 Last change: 14 Nov 2008 6
313 Protocols if_tcp(7P)
315     int            lifc_flags;    /* request specific
316                                /* interfaces */
317     int            lifc_len;     /* size of assoc. buffer */
318     union {
319         caddr_t    lifcu_buf;
320         struct lifreq *lifcu_req;
321     } lifc_lifcu;
323 #define lifc_buf lifc_lifcu.lifcu_buf /* buffer address */
324 #define lifc_req lifc_lifcu.lifcu_req /* array of structs returned */
325 };
327 The structure used by SIOCGLIFNUM has the form:
329 struct lifnum {
330     sa_family_t    lifn_family;
331     int            lifn_flags; /* req. specf. interfaces */
332     int            lifn_count; /* Result */
333 };
335 The structure used by SIOCTONLINK, SIOCTMYADDR and SIOCTMYSITE has the fo
337 struct sioc_addrreq {
338     struct sockaddr_storage sa_addr; /* Address to test */
339     int                    sa_res; /* Result - 0/1 */
340 };
342 The structure used by SIOCGLIFSRCOF has the form:
344 struct lifsrcof {
345     uint_t    lifs_ifindex; /* addr on this interface */
346             /* used as the src addr */
347     size_t    lifs_maxlen; /* size of buffer: input */
348     size_t    lifs_len; /* size of buffer: output */
349     union {
350         caddr_t lifsu_buf;
351         struct lifreq *lifsu_req;
352     } lifs_lifsu;
353 #define lifs_buf lifs_lifsu.lifsu_buf /* buffer addr. */
354 #define lifs_req lifs_lifsu.lifsu_req /* array returned */
355 };

```

```

357 The following ioctl() calls are maintained for compatibility
358 but only apply to IPv4 network interfaces, since the data
359 structures are too small to hold an IPv6 address. Unless
360 specified otherwise, the request takes an ifreq structure as
361 its parameter. This structure has the form:

363 struct ifreq {
365 SunOS 5.10          Last change: 14 Nov 2008          7
367 Protocols          if_tcp(7P)

369 #define IFNAMSIZ 16
370 char ifr_name[IFNAMSIZ]; /* interface name - e.g. "hme0" */
371 union {
372     struct sockaddr ifru_addr;
373     struct sockaddr ifru_dstaddr;
374     struct sockaddr ifru_broadaddr;
375     short ifru_flags;
376     int ifru_metric;
377     int ifr_muxid[2]; /* mux id's for arp and ip */
378     int ifru_index; /* interface index */
379 } ifr_ifru;

381 #define ifr_addr ifr_ifru.ifru_addr /* address */
382 #define ifr_dstaddr ifr_ifru.ifru_dstaddr /* other end of p-to-p lin
383 #define ifr_broadaddr ifr_ifru.ifru_broadaddr /* broadcast address */
384 #define ifr_flags ifr_ifru.ifru_flags /* flags */
385 #define ifr_index ifr_ifru.ifru_index /* interface index */
386 #define ifr_metric ifr_ifru.ifru_metric /* metric */
387 };

389 SIOCSIFADDR Set interface address.
391 SIOCGIFADDR Get interface address.
393 SIOCSIFDSTADDR Set point to point address for interface.
395 SIOCGIFDSTADDR Get point to point address for interface.
397 SIOCSIFFLAGS Set interface flags field. If the inter-
398 face is marked down, any processes
399 currently routing packets through the
400 interface are notified.

402 SIOCGIFFLAGS Get interface flags.

404 SIOCGIFCONF Get interface configuration list. This
405 request takes an ifconf structure (see
406 below) as a value-result parameter. The
407 ifc_len field should be set to the size
408 of the buffer pointed to by ifc_buf. Upon
409 success, ifc_len will contain the length,
410 in bytes, of the array of ifreq struc-
411 tures pointed to by ifc_req. For each
412 ifreq structure, the ifr_name and
413 *ifr_addr* fields will be valid. While
414 any IPMP IP interfaces will be included
415 in the array, the underlying IP
416 interfaces that comprise those IPMP
417 groups will not be.
418 ifr_addr fields will be valid.

419 SunOS 5.10          Last change: 14 Nov 2008          8
421 Protocols          if_tcp(7P)

```

```

423 SIOCGIFNUM Get number of interfaces. This request
424 returns an integer which is the number of
425 interface descriptions (struct ifreq)
426 that will be returned by the SIOCGIFCONF
427 ioctl; that is, it gives an indication of
428 how large *ifc_len* has to be. While any
429 IPMP IP interfaces will be included in
430 the array, the underlying IP interfaces
431 that comprise those IPMP groups will not
432 be.
433 how large ifc_len has to be.

434 SIOCSIFMTU Set the maximum transmission unit (MTU)
435 size for interface. Place the request in
436 the ifr_metric field. The MTU has to be
437 smaller than physical MTU limitation
438 (which is reported in the DLPI
439 DL_INFO_ACK message).

441 SIOCGIFMTU Get the maximum transmission unit size
442 for interface. Upon success, the request
443 is placed in the ifr_metric field.

445 SIOCSIFMETRIC Set the metric associated with the inter-
446 face. The metric is used by routine dae-
447 mons such as in.routed(1M).

449 SIOCGIFMETRIC Get the metric associated with the inter-
450 face.

452 SIOCGIFMUXID Get the ip and arp muxid associated with
453 the interface.

455 SIOCSIFMUXID Set the ip and arp muxid associated with
456 the interface.

458 SIOCGIFINDEX Get the interface index associated with
459 the interface.

461 SIOCSIFINDEX Set the interface index associated with
462 the interface.

464 The ifconf structure has the form:

466 struct ifconf {
467     int ifc_len; /* size of assoc. buffer */
468     union {
470 SunOS 5.10          Last change: 14 Nov 2008          9
472 Protocols          if_tcp(7P)

474         caddr_t ifcu_buf;
475         struct ifreq *ifcu_req;
476     } ifc_ifcu;

478     #define ifc_buf ifc_ifcu.ifcu_buf /* buffer address */
479     #define ifc_req ifc_ifcu.ifcu_req /* array of structs returned */
480 };

482 IFF_ Flags
483 You can use the ifconfig(1M) command to display the IFF_
484 flags listed below (with the leading IFF_ prefix removed).
485 See the ifconfig(1M) manpage for a definition of each flag.

```

```

487 #define IFF_UP 0x000000001 /* Address is up */
471 #define IFF_UP 0x000000001 /* Interface is up */
488 #define IFF_BROADCAST 0x000000002 /* Broadcast address valid */
489 #define IFF_DEBUG 0x000000004 /* Turn on debugging */
490 #define IFF_LOOPBACK 0x000000008 /* Loopback net */

492 #define IFF_POINTOPOINT 0x000000010 /* Interface is p-to-p */
493 #define IFF_NOTRAILERS 0x000000020 /* Avoid use of trailers */
494 #define IFF_RUNNING 0x000000040 /* Resources allocated */
495 #define IFF_NOARP 0x000000080 /* No address res. protocol */

497 #define IFF_PROMISC 0x000000100 /* Receive all packets */
498 #define IFF_ALLMULTI 0x000000200 /* Receive all multicast pkts */
499 #define IFF_INTELLIGENT 0x000000400 /* Protocol code on board */
500 #define IFF_MULTICAST 0x000000800 /* Supports multicast */

502 #define IFF_MULTI_BCAST 0x000001000 /* Multicast using broadcast. add.
503 #define IFF_UNNUMBERED 0x000002000 /* Non-unique address */
504 #define IFF_DHCPRUNNING 0x000004000 /* DHCP controls interface */
505 #define IFF_PRIVATE 0x000008000 /* Do not advertise */

507 #define IFF_NOXMIT 0x000010000 /* Do not transmit pkts */
508 #define IFF_NOLOCAL 0x000020000 /* No address - just on-link subn
509 #define IFF_DEPRECATED 0x000040000 /* Address is deprecated */
493 #define IFF_DEPRECATED 0x000040000 /* Interface addr. deprecated */
510 #define IFF_ADDRCONF 0x000080000 /* Addr. from stateless addrconf

512 #define IFF_ROUTER 0x000100000 /* Router on interface */
513 #define IFF_NONUD 0x000200000 /* No NUD on interface */
514 #define IFF_ANYCAST 0x000400000 /* Anycast address */
515 #define IFF_NORTXCH 0x000800000 /* Don't xchange rout. info */

517 #define IFF_IPV4 0x000100000 /* IPv4 interface */
518 #define IFF_IPV6 0x000200000 /* IPv6 interface */
519 #define IFF_NOFAILOVER 0x000800000 /* in.mpathd test address */
520 #define IFF_FAILED 0x001000000 /* Interface has failed */
503 #define IFF_NOFAILOVER 0x000800000 /* No failover on NIC fail. */
504 #define IFF_FAILED 0x001000000 /* NIC has failed */

522 #define IFF_STANDBY 0x002000000 /* Interface is a hot-spare */
523 #define IFF_INACTIVE 0x004000000 /* Functioning but not used */
524 #define IFF_OFFLINE 0x008000000 /* Interface is offline */
506 #define IFF_STANDBY 0x002000000 /* Standby NIC-use on fail. */
507 #define IFF_INACTIVE 0x004000000 /* Stndby active or not? */
508 #define IFF_OFFLINE 0x008000000 /* NIC offlined */

510 SunOS 5.10 Last change: 14 Nov 2008 10
512 Protocols if_tcp(7P)

525 #define IFF_XRESOLV 0x010000000 /* IPv6 external resolver */

527 #define IFF_COS_ENABLED 0x020000000 /* If CoS marking is supported */
528 #define IFF_PREFERRED 0x040000000 /* Prefer as source address */
529 #define IFF_TEMPORARY 0x080000000 /* RFC3041 */
530 #define IFF_FIXEDMTU 0x100000000 /* MTU set with SIOCSLIFMTU */

532 #define IFF_VIRTUAL 0x200000000 /* Cannot send/receive pkts */
533 #define IFF_DUPLICATE 0x400000000 /* Local address in use */
534 #define IFF_IPMP 0x800000000 /* IPMP IP interface */

536 ERRORS
537 EPERM Calling process has insufficient privileges.

539 ENXIO The lifr_name member of the lifreq structure
540 contains an invalid value.

```

```

542 For SIOCSLIFSRCOF, the lifs_ifindex member of
543 the lifsrcof structure contains an invalid
544 value.

546 For SIOCSLIFUSESRC, this error is returned if
547 the lifr_index is set to an invalid value.

549 EBADADDR Wrong address family or malformed address.

551 EBUSY For SIOCSLIFFLAGS, this error is returned when
552 the order of bringing the physical interface
553 (for example, eri0) and a logical interface
554 associated with the same physical interface
555 (for example, eri0:1) up or down is violated.
556 The physical interface must be configured up
557 first and cannot be configured down until all
558 the corresponding logical interfaces have been
559 configured down.

561 EINVAL For SIOCSLIFMTU, this error is returned when
562 the requested MTU size is invalid. This error
563 indicates the MTU size is greater than the MTU
564 size supported by the DLPI provider or less
565 than 68 (for IPv4) or less than 1280 (for
566 IPv6).

568 For SIOCSLIFUSESRC, this error is returned if
569 either the lifr_index or lifr_name identify
570 interfaces that are already part of an existing
571 IPMP group.

573 SunOS 5.10 Last change: 14 Nov 2008 11
575 Protocols if_tcp(7P)

577 EEXIST For SIOCSLIFADDIF, this error is returned if the
578 lifr_name member in the lifreq structure
579 corresponds to an interface that already has
580 the PPA specified by lifr_ppa plumbed.

582 SEE ALSO
583 ifconfig(1M), in.routed(1M), ioctl(2), streamio(7I),
584 arp(7P), dlpi(7P), ip(7P), ip6(7P)

```

```
*****
71559 Fri Dec 19 04:18:53 2008
new./ifconfig.txt
Clearview IPMP manpages
*****
```

```
1 System Administration Commands          ifconfig(1M)

5 NAME
6   ifconfig - configure network interface parameters

8 SYNOPSIS
9   ifconfig interface [address_family] [address [/prefix_length]
10  [dest_address]] [addif address [/prefix_length]]
11  [removeif address [/prefix_length]] [arp | -arp]
12  [auth_algs authentication algorithm] [encr_algs encryption algorithm]
13  [encr_auth_algs authentication algorithm] [auto-revarp]
14  [broadcast address] [deprecated | -deprecated]
15  [preferred | -preferred] [destination dest_address]
16  [ether {address}] [failover | -failover] [group
17  [name | ""]] [index if_index] [ipmp] [metric n] [modlist]
20  [name | ""]] [index if_index] [metric n] [modlist]
18  [modinsert mod_name@pos] [modremove mod_name@pos]
19  [mtu n] [netmask mask] [plumb] [unplumb] [private
20  | -private] [nud | -nud] [set {address} [/netmask]]
21  [standby | -standby] [subnet subnet_address] [tdst
22  tunnel_dest_address] [token address/prefix_length]
23  [tsrc tunnel_src_address] [trailers | -trailers]
24  [up] [down] [usesrc {name | none}] [xmit | -xmit]
25  [encaplimit n | -encaplimit] [thoplimit n] [router
26  | -router] [zone zonename | -zone | -all-zones]

29   ifconfig [address_family] interface {auto-dhcp | dhcp} [primary]
30   [wait seconds] drop | extend | inform | ping
31   | release | start | status

34 DESCRIPTION
35   The command ifconfig is used to assign an address to a net-
36   work interface and to configure network interface param-
37   eters. The ifconfig command must be used at boot time to
38   define the network address of each interface present on a
39   machine; it may also be used at a later time to redefine an
40   interface's address or other operating parameters. If no
41   option is specified, ifconfig displays the current confi-
42   guration for a network interface. If an address family is
43   specified, ifconfig reports only the details specific to
44   that address family. Only privileged users may modify the
45   configuration of a network interface. Options appearing
46   within braces ({} ) indicate that one of the options must be
47   specified.

49 DHCP Configuration
50   The forms of ifconfig that use the auto-dhcp or dhcp argu-
51   ments are used to control the Dynamic Host Configuration
52   Protocol ("DHCP") configuration of the interface. In this
53   mode, ifconfig is used to control operation of
54   dhcpagent(1M), the DHCP client daemon. Once an interface is
55   placed under DHCP control by using the start operand, ifcon-
56   fig should not, in normal operation, be used to modify the
```

```
60 SunOS 5.11          Last change: 21 Jan 2007          1

67 System Administration Commands          ifconfig(1M)

71   address or characteristics of the interface. If the address
72   of an interface under DHCP is changed, dhcpagent will remove
73   the interface from its control.

75 OPTIONS
76   The following options are supported:

78   addif address

80   Create the next unused logical interface on the speci-
81   fied physical interface.
82   fied physical interface.
84   fied physical interface. If the physical interface is
85   part of a multipathing group, the logical interface can
86   be added to a different physical interface in the same
87   group.

83   all-zones

85   Make the interface available to every shared-IP zone on
86   the system. The appropriate zone to which to deliver
87   data is determined using the tnzonecfg database. This
88   option is available only if the system is configured
89   with the Solaris Trusted Extensions feature.

91   The tnzonecfg database is described in the tnzonecfg(4)
92   man page, which is part of the Solaris Trusted Exten-
93   sions Reference Manual.

96   anycast

98   Marks the logical interface as an anycast address by
99   setting the ANYCAST flag. See "INTERFACE FLAGS," below,
100  for more information on anycast.

103  -anycast

105  Marks the logical interface as not an anycast address by
106  clearing the ANYCAST flag.

109  arp

111  Enable the use of the Address Resolution Protocol
112  ("ARP") in mapping between network level addresses and
113  link level addresses (default). This is currently imple-
114  mented for mapping between IPv4 addresses and MAC
115  addresses.
```

122 SunOS 5.11 Last change: 21 Jan 2007 2

129 System Administration Commands ifconfig(1M)

133 -arp

135 Disable the use of the ARP on a physical interface.
136 **ARP cannot be disabled on an IPMP IP interface.**

139 auth_algs authentication algorithm

141 For a tunnel, enable IPsec AH with the authentication
142 algorithm specified. The algorithm can be either a
143 number or an algorithm name, including any to express no
144 preference in algorithm. All IPsec tunnel properties
145 must be specified on the same command line. To disable
146 tunnel security, specify an auth_alg of none.

148 It is now preferable to use the ipsecconf(1M) command
149 when configuring a tunnel's security properties. If
150 ipsecconf was used to set a tunnel's security proper-
151 ties, this keyword will not affect the tunnel.

154 auto-dhcp

156 Use DHCP to automatically acquire an address for this
157 interface. This option has a completely equivalent alias
158 called dhcp.

160 For IPv6, the interface specified must be the zeroth
161 logical interface (the physical interface name), which
162 has the link-local address.

164 primary

166 Defines the interface as the primary. The interface
167 is defined as the preferred one for the delivery of
168 client-wide configuration data. Only one interface
169 can be the primary at any given time. If another
170 interface is subsequently selected as the primary,
171 it replaces the previous one. Nominating an inter-
172 face as the primary one will not have much signifi-
173 cance once the client work station has booted, as
174 many applications will already have started and been
175 configured with data read from the previous primary
176 interface.

179 wait seconds

181 The ifconfig command will wait until the operation
182 either completes or for the interval specified,
183 whichever is the sooner. If no wait interval is
184 given, and the operation is one that cannot complete

185 immediately, ifconfig will wait 30 seconds for the

189 SunOS 5.11 Last change: 21 Jan 2007 3

196 System Administration Commands ifconfig(1M)

200 requested operation to complete. The symbolic value
201 forever may be used as well, with obvious meaning.

204 drop

206 Remove the specified interface from DHCP control
207 without notifying the DHCP server, and record the
208 current lease for later use. Additionally, for IPv4,
209 set the IP address to zero and mark the interface as
210 "down." For IPv6, unplumb all logical interfaces
211 plumbed by dhcpagent.

214 extend

216 Attempt to extend the lease on the interface's IP
217 address. This is not required, as the agent will
218 automatically extend the lease well before it
219 expires.

222 inform

224 Obtain network configuration parameters from DHCP
225 without obtaining a lease on IP addresses. This is
226 useful in situations where an IP address is obtained
227 through mechanisms other than DHCP.

230 ping

232 Check whether the interface given is under DHCP con-
233 trol, which means that the interface is managed by
234 the DHCP agent and is working properly. An exit
235 status of 0 means success.

238 release

240 Relinquish the IP addresses on the interface by
241 notifying the server and discard the current lease.
242 For IPv4, mark the interface as "down." For IPv6,
243 all logical interfaces plumbed by dhcpagent are
244 unplumbed.

247 start

249 Start DHCP on the interface.

255 SunOS 5.11 Last change: 21 Jan 2007 4

262 System Administration Commands ifconfig(1M)

266 status

268 Display the DHCP configuration status of the inter-
269 face.

273 auto-revarp

275 Use the Reverse Address Resolution Protocol (RARP) to
276 automatically acquire an address for this interface.
277 This will fail if the interface does not support RARP;
278 for example, IPoIB (IP over InfiniBand), and on IPv6
279 interfaces.

282 broadcast address

284 For IPv4 only. Specify the address to use to represent
285 broadcasts to the network. The default broadcast address
286 is the address with a host part of all 1's. A "+" (plus
287 sign) given for the broadcast value causes the broadcast
288 address to be reset to a default appropriate for the
289 (possibly new) address and netmask. The arguments of
290 ifconfig are interpreted left to right. Therefore

292 example% ifconfig -a netmask + broadcast +

295 and

297 example% ifconfig -a broadcast + netmask +

300 may result in different values being assigned for the
301 broadcast addresses of the interfaces.

304 deprecated

306 Marks the logical interface as deprecated. An address
307 associated with a deprecated interface will not be used
308 as source address for outbound packets unless either
309 there are no other addresses available on the interface
310 or the application has bound to this address explicitly.
311 The status display shows DEPRECATED as part of flags.
312 See for information on the flags supported by ifconfig.

315 -deprecated

317 Marks a logical interface as not deprecated. An address

321 SunOS 5.11 Last change: 21 Jan 2007 5

328 System Administration Commands ifconfig(1M)

332 associated with such an interface could be used as a
333 source address for outbound packets.

336 preferred

338 Marks the logical interface as preferred. This option is
339 only valid for IPv6 addresses. Addresses assigned to
340 preferred logical interfaces are preferred as source
341 addresses over all other addresses configured on the
342 system, unless the address is of an inappropriate scope
343 relative to the destination address. Preferred addresses
344 are used as source addresses regardless of which physi-
345 cal interface they are assigned to. For example, you can
346 configure a preferred source address on the loopback
347 interface and advertise reachability of this address by
348 using a routing protocol.

351 -preferred

353 Marks the logical interface as not preferred.

356 destination dest_address

358 Set the destination address for a point-to point inter-
359 face.

362 dhcp

364 This option is an alias for option auto-dhcp

367 down

369 Mark a logical interface as "down". (That is, turn off
370 the IFF_UP bit.) When a logical interface is marked
371 "down," the system does not attempt to use the address
372 assigned to that interface as a source address for out-
373 bound packets and will not recognize inbound packets
374 destined to that address as being addressed to this
375 host. Additionally, when all logical interfaces on a
376 given physical interface are "down," the physical inter-
377 face itself is disabled.

379 When a logical interface is down, all routes that
380 specify that interface as the output (using the -ifp
381 option in the route(1M) command or RTA_IFP in a
382 route(7P) socket) are removed from the forwarding table.

```

383          Routes marked with RTF_STATIC are returned to the table

387 SunOS 5.11          Last change: 21 Jan 2007          6

394 System Administration Commands          ifconfig(1M)

398          if the interface is brought back up, while routes not
399          marked with RTF_STATIC are simply deleted.

401          When all logical interfaces that could possibly be used
402          to reach a particular gateway address are brought down
403          (specified without the interface option as in the previ-
404          ous paragraph), the affected gateway routes are treated
405          as though they had the RTF_BLACKHOLE flag set. All
406          matching packets are discarded because the gateway is
407          unreachable.

410          encaps limit n

412          Set the tunnel encapsulation limit for the interface to
413          n. This option applies to IPv4-in-IPv6 and IPv6-in-IPv6
414          tunnels only. The tunnel encapsulation limit controls
415          how many more tunnels a packet may enter before it
416          leaves any tunnels, that is, the tunnel nesting level.

419          -encaps limit

421          Disable generation of the tunnel encapsulation limit.
422          This option applies only to IPv4-in-IPv6 and IPv6-in-
423          IPv6 tunnels.

426          encr_auth_algs authentication algorithm

428          For a tunnel, enable IPsec ESP with the authentication
429          algorithm specified. It can be either a number or an
430          algorithm name, including any or none, to indicate no
431          algorithm preference. If an ESP encryption algorithm is
432          specified but the authentication algorithm is not, the
433          default value for the ESP authentication algorithm will
434          be any.

436          It is now preferable to use the ipsecconf(1M) command
437          when configuring a tunnel's security properties. If
438          ipsecconf was used to set a tunnel's security proper-
439          ties, this keyword will not affect the tunnel.

442          encr_algs encryption algorithm

444          For a tunnel, enable IPsec ESP with the encryption algo-
445          rithm specified. It can be either a number or an algo-
446          rithm name. Note that all IPsec tunnel properties must
447          be specified on the same command line. To disable tunnel
448          security, specify the value of encr_alg as none. If an

```

```

449          ESP authentication algorithm is specified, but the

453 SunOS 5.11          Last change: 21 Jan 2007          7

460 System Administration Commands          ifconfig(1M)

464          encryption algorithm is not, the default value for the
465          ESP encryption will be null.

467          It is now preferable to use the ipsecconf(1M) command
468          when configuring a tunnel's security properties. If
469          ipsecconf was used to set a tunnel's security proper-
470          ties, this keyword will not affect the tunnel.

473          ether [ address ]

475          If no address is given and the user is root or has suf-
476          ficient privileges to open the underlying datalink, then
482          sufficient privileges to open the underlying device, then
477          display the current Ethernet address information.

479          Otherwise, if the user is root or has sufficient
480          privileges, set the Ethernet address of the interfaces
481          to address. The address is an Ethernet address
482          represented as x:x:x:x:x:x where x is a hexadecimal
483          number between 0 and FF. Similarly, for the IPoIB (IP
484          over InfiniBand) interfaces, the address will be 20
485          bytes of colon-separated hex numbers between 0 and FF.

487          Some, though not all, Ethernet interface cards have
488          their own addresses. To use cards that do not have their
489          own addresses, refer to section 3.2.3(4) of the IEEE
490          802.3 specification for a definition of the locally
491          administered address space. Note that all IP interfaces
492          in an IPMP group must have unique hardware addresses;
493          see *in.mpathd(1M)*.
497          administered address space. The use of multipathing
498          groups should be restricted to those cards with their
499          own addresses (see MULTIPATHING GROUPS).

496          -failover

498          Set *NOFAILOVER* on the logical interface. This makes
499          the associated address available for use by *in.mpathd*
500          to perform probe-based failure detection for the
501          associated physical IP interface. As a side effect,
502          *DEPRECATED* will also be set on the logical interface.
503          This operation is not permitted on an IPMP IP interface.
504          Mark the logical interface as a non-failover interface.
505          Addresses assigned to non-failover logical interfaces
506          will not failover when the interface fails. Status
507          display shows NOFAILOVER as part of flags.

506          failover

```

```

508 Clear *NOFAILOVER* on the logical interface. This is
509 the default. These logical interfaces are subject to
510 migration when brought up (see IP MULTIPATHING GROUPS).
511 Mark the logical interface as a failover interface. An
512 address assigned to such an interface will failover when
513 the interface fails. Status display does not show
514 NOFAILOVER as part of flags.
515

```

```

513 group [ name | "" ]
514 group [ name | "" ]

```

```

515 When applied to a physical interface, it places the
516 interface into the named group. If the group does not
517 exist, it will be created, along with one or more IPMP
518 IP interfaces (for IPv4, IPv6, or both). Any *UP*
519 addresses that are not also marked *NOFAILOVER* are
520 subject to migration to the IPMP IP interface (see IP
521 MULTIPATHING GROUPS). Specifying a group name of ""*
522 removes the physical IP interface from the group.
523 Insert the logical interface in the multipathing group
524 specified by name. To delete an interface from a group,

```

```

524 When applied to a physical IPMP IP interface, it renames
525 the IPMP group to have the new name. If the name
526 already exists, or a name of ""* is specified, it
527 fails. Renaming IPMP groups is discouraged. Instead,
528 the IPMP IP interface should be given a meaningful name
529 when it is created via the *ipmp* subcommand, which the
530 system will also use as the IPMP group name.

```

```

525 SunOS 5.11 Last change: 21 Jan 2007 8

```

```

532 System Administration Commands ifconfig(1M)

```

```

536 use a null string "". When invoked on the logical inter-
537 face with id zero, the status display shows the group
538 name.

```

```

533 index n

```

```

535 Change the interface index for the interface. The value
536 of n must be an interface index (if_index) that is not
537 used on another interface. if_index will be a non-zero
538 positive number that uniquely identifies the network
539 interface on the system.

```

```

542 ipmp

```

```

544 Create an IPMP IP interface with the specified name. An
545 interface must be separately created for use by IPv4 and
546 IPv6. The *address_family* parameter controls whether
547 the command applies to IPv4 or IPv6 (IPv4 if
548 unspecified). All IPMP IP interfaces have the *IPMP*
549 flag set.

```

```

551 metric n

```

```

553 Set the routing metric of the interface to n; if no
554 value is specified, the default is 0. The routing metric
555 is used by the routing protocol. Higher metrics have the
556 effect of making a route less favorable. Metrics are
557 counted as addition hops to the destination network or
558 host.

```

```

561 modinsert mod_name@pos

```

```

563 Insert a module with name mod_name to the stream of the
564 device at position pos. The position is relative to the
565 stream head. Position 0 means directly under stream
566 head.

```

```

568 Based upon the example in the modlist option, use the
569 following command to insert a module with name ipqos
570 under the ip module and above the firewall module:

```

```

572 example% ifconfig eri0 modinsert ipqos@2

```

```

575 A subsequent listing of all the modules in the stream of
576 the device follows:

```

```

578 example% ifconfig eri0 modlist
579 0 arp
580 1 ip
581 2 ipqos
582 3 firewall
583 4 eri

```

```

592 SunOS 5.11 Last change: 21 Jan 2007 9

```

```

599 System Administration Commands ifconfig(1M)

```

```

603 modlist

```

```

605 List all the modules in the stream of the device.

```

```

607 The following example lists all the modules in the
608 stream of the device:

```

```

610 example% ifconfig eri0 modlist
611 0 arp
612 1 ip
613 2 firewall
614 4 eri

```

```

619 modremove mod_name@pos

621 Remove a module with name mod_name from the stream of
622 the device at position pos. The position is relative to
623 the stream head.

625 Based upon the example in the modinsert option, use the
626 following command to remove the firewall module from the
627 stream after inserting the ipqos module:

629 example% ifconfig eri0 modremove firewall@3

632 A subsequent listing of all the modules in the stream of
633 the device follows:

635 example% ifconfig eri0 modlist
636 0 arp
637 1 ip
638 2 ipqos
639 3 eri

642 Note that the core IP stack modules, for example, ip and
643 tun modules, cannot be removed.

646 mtu n

648 Set the maximum transmission unit of the interface to n.
649 For many types of networks, the mtu has an upper limit,
650 for example, 1500 for Ethernet. This option sets the
651 FIXEDMTU flag on the affected interface.

```

658 SunOS 5.11 Last change: 21 Jan 2007 10

665 System Administration Commands ifconfig(1M)

```

669 netmask mask

671 For IPv4 only. Specify how much of the address to
672 reserve for subdividing networks into subnetworks. The
673 mask includes the network part of the local address and
674 the subnet part, which is taken from the host field of
675 the address. The mask contains 1's for the bit positions
676 in the 32-bit address which are to be used for the net-
677 work and subnet parts, and 0's for the host part. The
678 mask should contain at least the standard network por-
679 tion, and the subnet field should be contiguous with the
680 network portion. The mask can be specified in one of
681 four ways:

```

```

683 1. with a single hexadecimal number with a leading
684 0x,

686 2. with a dot-notation address,

688 3. with a "+" (plus sign) address, or

690 4. with a pseudo host name/pseudo network name
691 found in the network database networks(4).
692 If a "+" (plus sign) is given for the netmask value, the
693 mask is looked up in the netmasks(4) database. This
694 lookup finds the longest matching netmask in the data-
695 base by starting with the interface's IPv4 address as
696 the key and iteratively masking off more and more low
697 order bits of the address. This iterative lookup ensures
698 that the netmasks(4) database can be used to specify the
699 netmasks when variable length subnetmasks are used
700 within a network number.

702 If a pseudo host name/pseudo network name is supplied as
703 the netmask value, netmask data may be located in the
704 hosts or networks database. Names are looked up by first
705 using gethostbyname(3NSL). If not found there, the names
706 are looked up in getnetbyname(3SOCKET). These interfaces
707 may in turn use nsswitch.conf(4) to determine what data
708 store(s) to use to fetch the actual value.

710 For both inet and inet6, the same information conveyed
711 by mask can be specified as a prefix_length attached to
712 the address parameter.

715 nud

717 Enables the neighbor unreachability detection mechanism
718 on a point-to-point physical interface.

```

724 SunOS 5.11 Last change: 21 Jan 2007 11

731 System Administration Commands ifconfig(1M)

```

735 -nud

737 Disables the neighbor unreachability detection mechanism
738 on a point-to-point physical interface.

741 plumb

743 For a physical IP interface, open the datalink
744 associated with the physical interface name and set up
745 the plumbing needed for IP to use the datalink. When
746 used with a logical interface name, this command is used
747 to create a specific named logical interface on an

```

748 **existing physical IP interface.**
 742 *Open the device associated with the physical interface*
 743 *name and set up the streams needed for IP to use the*
 744 *device. When used with a logical interface name, this*
 745 *command is used to create a specific named logical*
 746 *interface. An interface must be separately plumbed for*
 747 *use by IPv4 and IPv6. The address_family parameter con-*
 748 *trols whether the ifconfig command applies to IPv4 or*
 749 *IPv6.*

750 **An interface must be separately plumbed for IPv4 and**
 751 **IPv6 according to the *address_family* parameter (IPv4**
 752 **if unspecified). Before an interface has been plumbed,**
 753 **it will not be shown by *ifconfig -a*.**
 751 *Before an interface has been plumbed, the interface will*
 752 *not show up in the output of the ifconfig -a command.*

755 **Note that IPMP IP interfaces are not tied to a specific**
 756 **datalink and are instead created with the *ipmp***
 757 **subcommand.**

759 private

761 Tells the in.routed routing daemon that a specified log-
 762 ical interface should not be advertised.

765 -private

767 Specify unadvertised interfaces.

770 removeif address

772 Remove the logical interface on the physical interface
 773 **specified that matches the address specified.**
 769 *specified that matches the address specified. When the*
 770 *interface is part of a multipathing group, the logical*
 771 *interface will be removed from the physical interface in*
 772 *the group that holds the address.*

775 router

777 Enable IP forwarding on the interface. When enabled, the
 778 **interface is marked *ROUTER*, and IP packets can be for-**
 779 **warded to and from the interface. Enabling *ROUTER* on**
 780 **any IP interface in an IPMP group applies the flag to**
 781 **all IP interfaces in that IPMP group.**
 778 *interface is marked ROUTER, and IP packets can be for-*
 779 *warded to and from the interface.*

784 -router

786 Disable IP forwarding on the interface. IP packets are
 787 **not forwarded to and from the interface. Disabling**
 788 ***ROUTER* on any IP interface in an IPMP group disables**
 789 **it on all IP interfaces in that IPMP group.**
 785 *not forwarded to and from the interface.*

800 System Administration Commands ifconfig(1M)

804 set

806 Set the address, prefix_length or both, for a logical
 807 interface.

810 standby

812 **Mark the physical IP interface as a *STANDBY* interface.**
 813 **If an interface is marked *STANDBY* and is part of an**
 814 **IPMP group, the interface will not be used for data**
 815 **traffic unless another interface in the IPMP group**
 816 **becomes unusable. When a *STANDBY* interface is**
 817 **functional but not being used for data traffic, it will**
 818 **also be marked *INACTIVE*. This operation is not**
 819 **permitted on an IPMP IP interface.**
 808 *Marks the physical interface as a standby interface. If*
 809 *the interface is marked STANDBY and is part of the mul-*
 810 *tipathing group, the interface will not be selected to*
 811 *send out packets unless some other interface in the*
 812 *group has failed and the network access has been failed*
 813 *over to this standby interface.*

815 *The status display shows "STANDBY, INACTIVE" indicating*
 816 *that that the interface is a standby and is also inac-*
 817 *tive. IPF_INACTIVE will be cleared when some other*
 818 *interface belonging to the same multipathing group fails*
 819 *over to this interface. Once a fallback happens, the*
 820 *status display will return to INACTIVE.*

822 -standby

824 **Clear *STANDBY* on this interface. This is the default.**
 825 **Turns off standby on this interface.**

827 subnet

829 Set the subnet address for an interface.

832 tdst tunnel_dest_address

834 Set the destination address of a tunnel. The address
 835 should not be the same as the dest_address of the tun-
 836 nel, because no packets leave the system over such a
 837 tunnel.

840 thoplimit n

842 Set the hop limit for a tunnel interface. The hop limit
 843 value is used as the TTL in the IPv4 header for the
 844 IPv6-in-IPv4 and IPv4-in-IPv4 tunnels. For IPv6-in-IPv6
 845 and IPv4-in-IPv6 tunnels, the hop limit value is used as
 846 the hop limit in the IPv6 header.

```

849     token address/prefix_length

854 SunOS 5.11             Last change: 21 Jan 2007             13

861 System Administration Commands             ifconfig(1M)

865     Set the IPv6 token of an interface to be used for
866     address autoconfiguration.

868     example% ifconfig eri0 inet6 token ::1/64

873     trailers

875     This flag previously caused a nonstandard encapsulation
876     of IPv4 packets on certain link levels. Drivers supplied
877     with this release no longer use this flag. It is pro-
878     vided for compatibility, but is ignored.

881     -trailers

883     Disable the use of a "trailer" link level encapsulation.

886     tsrc tunnel_src_address

888     Set the source address of a tunnel. This is the source
889     address on an outer encapsulating IP header. It must be
890     an address of another interface already configured using
891     ifconfig.

894     unplumb

896     For a physical or IPMP interface, remove all associated
897     logical IP interfaces and tear down any plumbing needed
898     for IP to use the interface. For an IPMP IP interface,
899     this command will fail if the group is not empty. For a
900     logical interface, the logical interface is removed.
897     Close the device associated with this physical interface
898     name and any streams that ifconfig set up for IP to use
899     the device. When used with a logical interface name, the
900     logical interface is removed from the system. After this
901     command is executed, the device name will no longer
902     appear in the output of ifconfig -a.

902     An interface must be separately unplumbed for IPv4 and
903     IPv6 according to the *address_family* parameter (IPv4
904     if unspecified). Upon success, the interface name will
905     no longer appear in the output of *ifconfig -a*.

```

```

908     up

910     Mark a logical interface *UP*. As a result, the IP
911     module will accept packets destined to the associated
912     address (unless the address is zero), along with any
913     associated multicast and broadcast IP addresses.
914     Similarly, the IP module will allow packets to be sent
915     with the associated address as a source address.
907     Mark a logical interface "up". This happens automati-
908     cally when assigning the first address to a logical
909     interface. The up option enables an interface after an
910     ifconfig down, which reinitializes the hardware.

918     usesrc [ name | none ]

920     Specify a physical interface to be used for source
921     address selection. If the keyword none is used, then any
922     previous selection is cleared.

926 SunOS 5.11             Last change: 21 Jan 2007             14

933 System Administration Commands             ifconfig(1M)

937     When an application does not choose a non-zero source
938     address using bind(3SOCKET), the system will select an
939     appropriate source address based on the outbound inter-
940     face and the address selection rules (see
941     ipaddrsel(1M)).

943     When usesrc is specified and the specified interface is
944     selected in the forwarding table for output, the system
945     looks first to the specified physical interface and its
946     associated logical interfaces when selecting a source
947     address. If no usable address is listed in the forward-
948     ing table, the ordinary selection rules apply. For exam-
949     ple, if you enter:

951     # ifconfig eri0 usesrc vni0

954     ..and vni0 has address 10.0.0.1 assigned to it, the
955     system will prefer 10.0.0.1 as the source address for
956     any packets originated by local connections that are
957     sent through eri0. Further examples are provided in the
958     EXAMPLES section.

960     While you can specify any physical interface (or even
961     loopback), be aware that you can also specify the vir-
962     tual IP interface (see vni(7D)). The virtual IP inter-
963     face is not associated with any physical hardware and is
964     thus immune to hardware failures. You can specify any
965     number of physical interfaces to use the source address
966     hosted on a single virtual interface. This simplifies
967     the configuration of routing-based multipathing. If one
968     of the physical interfaces were to fail, communication
969     would continue through one of the remaining, functioning

```

970 physical interfaces. This scenario assumes that the
971 reachability of the address hosted on the virtual inter-
972 face is advertised in some manner, for example, through
973 a routing protocol.

975 Because the ifconfig preferred option is applied to all
976 interfaces, it is coarser-grained than the usesrc
977 option. It will be overridden by usesrc and setsrc
978 (route subcommand), in that order.

980 The use of the usesrc option is mutually exclusive of
981 the IPMP *group* and *standby* subcommands. That is, if
982 an interface is already part of a IPMP group or
983 specified as a *STANDBY* interface, then it cannot be
984 specified with a usesrc option, and vice-versa.
976 the IP multipathing ifconfig options, group and standby.
977 That is, if an interface is already part of a IP mul-
978 tipathing group or specified as a standby interface,
979 then it cannot be specified with a usesrc option, and
980 vice-versa. For more details on IP multipathing, see
981 in.mpathd(1M) and the .

988 SunOS 5.11 Last change: 21 Jan 2007 15

995 System Administration Commands ifconfig(1M)

999 xmit

1001 Enable a logical interface to transmit packets. This is
1002 the default behavior when the logical interface is up.

1005 -xmit

1007 Disable transmission of packets on an interface. The
1008 interface will continue to receive packets.

1011 zone zonename

1013 Place the logical interface in zone zonename. The named
1014 zone must be active in the kernel in the ready or run-
1015 ning state. The interface is unplumbed when the zone is
1016 halted or rebooted. The zone must be configured to be an
1017 shared-IP zone. zonecfg(1M) is used to assign network
1018 interface names to exclusive-IP zones.

1021 -zone

1023 Place IP interface in the global zone. This is the
1024 default.

1027 OPERANDS

1028 The interface operand, as well as address parameters that
1029 affect it, are described below.

1031 interface

1033 A string of one of the following forms:

1035 o name physical-unit, for example, eri0 or cel

1037 o name physical-unit:logical-unit, for example,
1038 eri0:l

1040 o ip.tunN or ip6.tunN, for tunnels

1041 If the interface name starts with a dash (-), it is
1042 interpreted as a set of options which specify a set of
1043 interfaces. In such a case, -a must be part of the
1044 options and any of the additional options below can be
1045 added in any order. If one of these interface names is
1046 given, the commands following it are applied to all of
1047 the interfaces that match.

1049 -a

1054 SunOS 5.11 Last change: 21 Jan 2007 16

1061 System Administration Commands ifconfig(1M)

1065 Apply the command to all interfaces of the specified
1066 address family. If no address family is supplied,
1067 either on the command line or by means of
1068 /etc/default/inet_type, then all address families
1069 will be selected.

1072 -d

1074 Apply the commands to all "down" interfaces in the
1075 system.

1078 -D

1080 Apply the commands to all interfaces not under DHCP
1081 (Dynamic Host Configuration Protocol) control.

1084 -u

1086 Apply the commands to all "up" interfaces in the
1087 system.

1090 -Z

1092 Apply the commands to all interfaces in the user's
1093 zone.

1096 -4

1098 Apply the commands to all IPv4 interfaces.

1101 -6

1103 Apply the commands to all IPv6 interfaces.

1107 address_family

1109 The address family is specified by the address_family
1110 parameter. The ifconfig command currently supports the
1111 following families: inet and inet6. If no address family
1112 is specified, the default is inet.

1114 ifconfig honors the DEFAULT_IP setting in the
1115 /etc/default/inet_type file when it displays interface
1116 information. If DEFAULT_IP is set to IP_VERSION4, then

1120 SunOS 5.11 Last change: 21 Jan 2007 17

1127 System Administration Commands ifconfig(1M)

1131 ifconfig will omit information that relates to IPv6
1132 interfaces. However, when you explicitly specify an
1133 address family (inet or inet6) on the ifconfig command
1134 line, the command line overrides the DEFAULT_IP set-
1135 tings.

1138 address

1140 For the IPv4 family (inet), the address is either a host
1141 name present in the host name data base (see hosts(4))
1142 or in the Network Information Service (NIS) map hosts,
1143 or an IPv4 address expressed in the Internet standard
1144 "dot notation".

1146 For the IPv6 family (inet6), the address is either a
1147 host name present in the host name data base (see
1148 hosts(4)) or in the Network Information Service (NIS)
1149 map ipnode, or an IPv6 address expressed in the Internet
1150 standard colon-separated hexadecimal format represented
1151 as x:x:x:x:x:x:x:x where x is a hexadecimal number
1152 between 0 and FFFF.

1155 prefix_length

1157 For the IPv4 and IPv6 families (inet and inet6), the
1158 prefix_length is a number between 0 and the number of
1159 bits in the address. For inet, the number of bits in the

1160 address is 32; for inet6, the number of bits in the
1161 address is 128. The prefix_length denotes the number of
1162 leading set bits in the netmask.

1165 dest_address

1167 If the dest_address parameter is supplied in addition to
1168 the address parameter, it specifies the address of the
1169 correspondent on the other end of a point-to-point link.

1172 tunnel_dest_address

1174 An address that is or will be reachable through an
1175 interface other than the tunnel being configured. This
1176 tells the tunnel where to send the tunneled packets.
1177 This address must not be the same as the interface des-
1178 tination address being configured.

1181 tunnel_src_address

1186 SunOS 5.11 Last change: 21 Jan 2007 18

1193 System Administration Commands ifconfig(1M)

1197 An address that is attached to an already configured
1198 interface that has been configured "up" with ifconfig.

1201 INTERFACE FLAGS

1202 The ifconfig command supports the following interface flags.

1203 The term "address" in this context refers to a logical
1204 interface, for example, eri0:0, while "interface " refers to
1205 the physical interface, for example, eri0.

1207 ADDRCONF

1209 The address is from stateless addrconf. The stateless
1210 mechanism allows a host to generate its own address
1211 using a combination of information advertised by routers
1212 and locally available information. Routers advertise
1213 prefixes that identify the subnet associated with the
1214 link, while the host generates an "interface identifier"
1215 that uniquely identifies an interface in a subnet. In
1216 the absence of information from routers, a host can gen-
1217 erate link-local addresses. This flag is specific to
1218 IPv6.

1221 ANYCAST

1223 Indicates an anycast address. An anycast address identi-
1224 fies the nearest member of a group of systems that pro-
1225 vides a particular type of service. An anycast address

1226 is assigned to a group of systems. Packets are delivered
 1227 to the nearest group member identified by the anycast
 1228 address instead of being delivered to all members of the
 1229 group.

1232 BROADCAST

1234 This broadcast address is valid. This flag and POINTTO-
 1235 POINT are mutually exclusive

1238 CoS

1240 This interface supports some form of Class of Service
 1241 (CoS) marking. An example is the 802.1D user priority
 1242 marking supported on VLAN interfaces. For IPMP IP
 1243 interfaces, this will only be set if all interfaces in
 1244 the group have CoS set.
 1245 marking supported on VLAN interfaces.

1247 DEPRECATED

1249 This address is deprecated. This address will not be
 1250 used as a source address for outbound packets unless

1251 SunOS 5.11 Last change: 21 Jan 2007 19

1258 System Administration Commands ifconfig(1M)

1251 there are no other addresses on this interface or an
 1252 application has explicitly bound to this address. An
 1253 IPv6 deprecated address is part of the standard
 1254 mechanism for renumbering in IPv6 and will eventually be
 1255 deleted when not used. For both IPv4 and IPv6,
 1256 *DEPRECATED* is also set on all *NOFAILOVER* addresses,
 1257 though this may change in a future release.
 1258 IPv6 deprecated address will eventually be deleted when
 1259 not used, whereas an IPv4 deprecated address is often
 1260 used with IP network multipathing IPv4 test addresses,
 1261 which are determined by the setting of the NOFAILOVER
 1262 flag. Further, the DEPRECATED flag is part of the stan-
 1263 dard mechanism for renumbering in IPv6.

1259 DHCPRUNNING

1261 The logical interface is managed by *dhcpageant(1M)*.
 1272 DHCP

1274 DHCP is used to manage this address.

1264 DUPLICATE

1266 The logical interface has been disabled because the IP
 1267 address configured on the interface is a duplicate. Some

1268 other node on the network is using this address. If the
 1269 address was configured by DHCP or is temporary, the sys-
 1270 tem will choose another automatically, if possible. Other-
 1271 wise, the system will attempt to recover this address
 1272 periodically and the interface will recover when the
 1273 conflict has been removed from the network. Changing the
 1274 address or netmask, or setting the logical interface to
 1275 up will restart duplicate detection. Setting the inter-
 1276 face to down terminates recovery and removes the DUPLI-
 1277 CATE flag.

1280 FAILED

1282 The *in.mpathd* daemon has determined that the interface
 1283 has failed. *FAILED* interfaces will not be used to
 1284 send or receive IP data traffic. If this is set on a
 1285 physical IP interface in an IPMP group, IP data traffic
 1286 will continue to flow over other usable IP interfaces in
 1287 the IPMP group. If this is set on an IPMP IP interface,
 1288 the entire group has failed and no data traffic can be
 1289 sent or received over any interfaces in that group.
 1290 The interface has failed. New addresses cannot be
 1291 created on this interface. If this interface is part of
 1292 an IP network multipathing group, a failover will occur
 1293 to another interface in the group, if possible

1292 FIXEDMTU

1294 The MTU has been set using the -mtu option. This flag is
 1295 read-only. Interfaces that have this flag set have a
 1296 fixed MTU value that is unaffected by dynamic MTU
 1297 changes that can occur when drivers notify IP of link
 1298 MTU changes.

1301 INACTIVE

1303 The physical interface is functioning but is not used to
 1304 send or receive data traffic according to administrative
 1305 policy. This flag is initially set by the *standby*
 1306 subcommand and is subsequently controlled by
 1307 *in.mpathd*. It also set when *FAILBACK=no* mode is
 1308 enabled (see *in.mpathd(1M)*) to indicate that the IP
 1309 interface has repaired but is not being used.
 1310 Indicates that the interface is not currently being used
 1311 for regular traffic by the system. New addresses cannot

1312 IPMP

1314 Indicates that this is an IPMP IP interface.
 1317 SunOS 5.11 Last change: 21 Jan 2007 20

1324 System Administration Commands ifconfig(1M)

1328 be created on this interface. The flag is set automati-
 1329 cally on standby interfaces. It can also be set when the

1330 *system detects that a failed interface has been repaired*
 1331 *and FAILBACK=no is configured in /etc/default/mpathd.*
 1332 *The flag is cleared when the interface fails or when a*
 1333 *failover to that interface occurs.*

1317 LOOPBACK

1319 Indicates that this is the loopback interface.

1341 MIP

1343 *Indicates that mobile IP controls this interface.*

1322 MULTI_BCAST

1324 Indicates that the broadcast address is used for multi-
 1325 cast on this interface.

1328 MULTICAST

1330 The interface supports multicast. IP assumes that any
 1331 interface that supports hardware broadcast, or that is a
 1332 point-to-point link, will support multicast.

1335 NOARP

1337 There is no address resolution protocol (ARP) for this
 1338 interface that corresponds to all interfaces for a dev-
 1339 ice without a broadcast address. This flag is specific
 1340 to IPv4.

1343 NOFAILOVER

1345 **The address associated with this logical interface is**
 1346 **available to *in.mpathd* for probe-based failure**
 1347 **detection of the associated physical IP interface.**
 1369 *This address will not failover if the interface fails.*
 1370 *IP network multipathing test addresses must be marked*
 1371 *nofailover.*

1350 NOLOCAL

1352 The interface has no address , just an on-link subnet.

1359 SunOS 5.11 Last change: 21 Jan 2007 21

1366 System Administration Commands ifconfig(1M)

1370 NONUD

1372 NUD is disabled on this interface. NUD (neighbor
 1373 unreachability detection) is used by a node to track the
 1374 reachability state of its neighbors, to which the node
 1375 actively sends packets, and to perform any recovery if a
 1376 neighbor is detected to be unreachable. This flag is
 1377 specific to IPv6.

1380 NORTEXCH

1382 The interface does not exchange routing information. For
 1383 RIP-2, routing packets are not sent over this interface.
 1384 Additionally, messages that appear to come over this
 1385 interface receive no response. The subnet or address of
 1386 this interface is not included in advertisements over
 1387 other interfaces to other routers.

1390 NOXMIT

1392 Indicates that the address does not transmit packets.
 1393 RIP-2 also does not advertise this address.

1396 OFFLINE

1398 **The interface is offline and thus cannot send or receive**
 1399 **IP data traffic. This is only set on IP interfaces in**
 1400 **an IPMP group. See *if_mpadm(1M)* and *cfgadm(1M)*.**
 1422 *Indicates that the interface has been offlined. New*
 1423 *addresses cannot be created on this interface. Inter-*
 1424 *faces in an IP network multipathing group are offlined*
 1425 *prior to removal and replacement using dynamic reconfi-*
 1426 *guration.*

1403 POINTOPOINT

1405 Indicates that the address is a point-to-point link.
 1406 This flag and BROADCAST are mutually exclusive

1409 PREFERRED

1411 This address is a preferred IPv6 source address. This
 1412 address will be used as a source address for IPv6 com-
 1413 munication with all IPv6 destinations, unless another
 1414 address on the system is of more appropriate scope. The
 1415 DEPRECATED flag takes precedence over the PREFERRED
 1416 flag.

1423 SunOS 5.11 Last change: 21 Jan 2007 22

```

1430 System Administration Commands          ifconfig(1M)

1434 PRIVATE

1436     Indicates that this address is not advertised. For RIP-
1437     2, this interface is used to send advertisements. How-
1438     ever, neither the subnet nor this address are included
1439     in advertisements to other routers.

1442 ROUTER

1444     Indicates that IP packets can be forwarded to and from
1445     the interface.

1448 RUNNING

1450     Indicates that the required resources for an i nterface
1451     Indicates that the required resources for an interface
1452     are allocated. For some interfaces this also indicates
1453     that the link is up. For IPMP IP interfaces, *RUNNING*
1454     is set as long as one IP interface in the group is
1455     active.
1456     that the link is up.

1457 STANDBY

1459     Indicates that this physical interface will not be used
1460     for data traffic unless another interface in the IPMP
1461     group becomes unusable. The *INACTIVE* and *FAILED*
1462     flags indicate whether it is actively being used.
1463     Indicates that this is a standby interface to be used on
1464     failures. Only interfaces in an IP network multipathing
1465     group should be designated as standby interfaces. If
1466     this interface is part of a IP network multipathing
1467     group, the interface will not be selected to send out
1468     packets unless some other interface in the group fails
1469     over to it.

1465 TEMPORARY

1467     Indicates that this is a temporary IPv6 address as
1468     defined in RFC 3041.

1471 UNNUMBERED

1473     This flag is set when the local IP address on the link
1474     matches the local address of some other link in the sys-
1475     tem

1478 UP

1480     Indicates that the logical interface (and the associated
1481     physical interface) is up. The IP module will accept
1482     packets destined to UP addresses (unless the address is
1483     zero), along with any associated multicast and broadcast
1484     IP addresses. Similarly, the IP module will allow
1485     packets to be sent with an UP address as a source
1486     address.

```

```

1507     Indicates that the interface is up, that is, all the
1508     routing entries and the like for this interface have
1509     been set up.

1489 SunOS 5.11          Last change: 21 Jan 2007          23

1496 System Administration Commands          ifconfig(1M)

1500 VIRTUAL

1502     Indicates that the physical interface has no underlying
1503     hardware. It is not possible to transmit or receive
1504     packets through a virtual interface. These interfaces
1505     are useful for configuring local addresses that can be
1506     used on multiple interfaces. (See also the *usesrc*
1507     used on multiple interfaces. (See also the -usesrc
1508     option.)

1510 XRESOLV

1512     Indicates that the interface uses an IPv6 external
1513     resolver.

1516 LOGICAL INTERFACES
1517     Solaris TCP/IP allows multiple logical interfaces to be
1518     associated with a physical network interface. This allows a
1519     single machine to be assigned multiple IP addresses, even
1520     though it may have only one network interface. Physical net-
1521     work interfaces have names of the form driver-name
1522     physical-unit-number, while logical interfaces have names of
1523     the form driver-name physical-unit-number:logical-unit-
1524     number. A physical interface is configured into the system
1525     using the plumb command. For example:

1527     example% ifconfig eri0 plumb

1532     Once a physical interface has been "plumbed", logical inter-
1533     faces associated with the physical interface can be config-
1534     ured by separate -plumb or -addif options to the ifconfig
1535     command.

1537     example% ifconfig eri0:1 plumb

1542     allocates a specific logical interface associated with the
1543     physical interface eri0. The command

1545     example% ifconfig eri0 addif 192.168.200.1/24 up

```

1550 allocates the next available logical unit number on the eri0
1551 physical interface and assigns an address and prefix_length.

1555 SunOS 5.11 Last change: 21 Jan 2007 24

1562 System Administration Commands ifconfig(1M)

1566 A logical interface can be configured with parameters (
1567 address, prefix_length, and so on) different from the physi-
1568 cal interface with which it is associated. Logical inter-
1569 faces that are associated with the same physical interface
1570 can be given different parameters as well. Each logical
1571 interface must be associated with an existing and "up" phy-
1572 sical interface. So, for example, the logical interface
1573 eri0:1 can only be configured after the physical interface
1574 eri0 has been plumbed.

1577 To delete a logical interface, use the **unplumb** or
1578 **removeif** options. For example,
1603 To delete a logical interface, use the *-unplumb* or *-removeif*
1604 options. For example,

1580 example% ifconfig eri0:1 down unplumb

1582 will delete the logical interface **eri0:1**.

1584 IP MULTIPATHING GROUPS

1586 Physical interfaces that share the same IP broadcast domain
1587 *must* be collected into a single IP Multipathing (IPMP)
1588 group using the **group** subcommand. Each IPMP group has an
1589 associated IPMP IP interface, which can either be explicitly
1590 created (the preferred method) by using the **ipmp**
1591 subcommand or implicitly created by **ifconfig** in response
1592 to placing an IP interface into a new IPMP group.
1593 Implicitly-created IPMP interfaces will be named *ipmp_N*
1594 where *N* is the lowest integer that doesn't conflict with
1595 an existing IP interface name or IPMP group name.

1597 Each IPMP IP interface is created with a matching IPMP group
1598 name, though it can be changed using the **group** subcommand.
1599 Each IPMP IP interface hosts a set of highly-available IP
1600 addresses. These addresses will remain reachable so long as
1601 at least one interface in the group is active, where
1602 "active" is defined as having at least one UP address and
1603 having **INACTIVE**, **FAILED**, and **OFFLINE** clear. IP
1604 addresses hosted on the IPMP IP interface may either be
1605 configured statically or configured through DHCP via the
1606 **dhcp** subcommand.
1611 will delete the logical interface *eri0:1*.

1608 Interfaces assigned to the same IPMP group are treated as

1609 equivalent and monitored for failure by **in.mpathd**.
1610 Provided that active interfaces in the group remain, IP
1611 interface failures (and any subsequent repairs) are handled
1612 transparently to sockets-based applications. IPMP is also
1613 integrated with the Dynamic Reconfiguration framework (see
1614 **cfgadm(1M)**), which enables network adapters to be replaced
1615 transparently to sockets-based applications.

1613 MULTIPATHING GROUPS

1614 Physical interfaces that share the same IP broadcast domain
1615 can be collected into a multipathing group using the group
1616 keyword. Interfaces assigned to the same multipathing group
1617 are treated as equivalent and outgoing traffic is spread
1618 across the interfaces on a per-IP-destination basis. In
1619 addition, individual interfaces in a multipathing group are
1620 monitored for failures; the addresses associated with failed
1621 interfaces are automatically transferred to other function-
1622 ing interfaces within the group.

1617 The IP module automatically load-spreads all outbound
1618 traffic across all active interfaces in an IPMP group.
1619 Similarly, all **UP** addresses hosted on the IPMP IP
1620 interface and will be distributed across the active
1621 interfaces to promote inbound load-spreading. The
1622 **ipmpstat(1M)** utility allows many aspects of the IPMP
1623 subsystem to be observed, including the current binding of
1624 IP data addresses to IP interfaces.

1626 When an interface is placed into an IPMP group, any **UP**
1627 logical interfaces are "migrated" to the IPMP IP interface
1628 for use by the group, unless:
1625 For more details on IP multipathing, see *in.mpathd(1M)* and
1626 the *.* See *netstat(1M)* for per-IP-destination information.

1630 * The logical interface is marked **NOFAILOVER**
1631 * The logical interface hosts an IPv6 link-local address.
1632 * The logical interface hosts an IPv4 0.0.0.0 address.

1634 Likewise, once an interface is in a group, if changes are
1635 made to a logical interface such that it is **UP** and not
1636 exempted by one of the conditions above, it will also
1637 migrate to the associated IPMP IP interface. Logical
1638 interfaces never migrate back, even if the physical
1639 interface that contributed the address is removed from the
1640 group.

1642 Each interface placed into an IPMP group may be optionally
1643 configured with a "test" address that **in.mpathd** will use
1644 for probe-based failure detection; see **in.mpathd(1M)**.
1645 These addresses must be marked **NOFAILOVER** (using the
1646 **failover** subcommand) prior to being marked **UP**. Test
1647 addresses may also be acquired through DHCP via the **dhcp**
1648 subcommand.

1649 For more background on IPMP, please see the "IPMP
1650 Administrative Overview" and "IPMP Configuration Tasks"
1651 chapters of the administrator documentation.

1655 CONFIGURING IPV6 INTERFACES

1656 When an IPv6 physical interface is plumbed and configured
1657 "up" with *ifconfig*, it is automatically assigned an IPv6
1658 link-local address for which the last 64 bits are calculated
1659 from the MAC address of the interface.

1661 example% ifconfig eri0 inet6 plumb up

1666 The following example shows that the link-local address has
1667 a prefix of fe80::/10.

```
1669     example% ifconfig eri0 inet6
1670     ce0: flags=2000841<UP,RUNNING,MULTICAST,IPv6>
```

1674 SunOS 5.11 Last change: 21 Jan 2007 25

1681 System Administration Commands ifconfig(1M)

```
1685         mtu 1500 index 2
1686         inet6 fe80::a00:20ff:fe8e:f3ad/10
```

1691 Link-local addresses are only used for communication on the
1692 local subnet and are not visible to other subnets.

1695 If an advertising IPv6 router exists on the link advertising
1696 prefixes, then the newly plumbed IPv6 interface will auto-
1697 configure logical interface(s) depending on the prefix
1698 advertisements. For example, for the prefix advertisement
1699 2001:0db8:3c4d:0:55::/64, the autoconfigured interface will
1700 look like:

```
1702     eri0:2: flags=2080841<UP,RUNNING,MULTICAST,ADDRCONF,IPv6>
1703         mtu 1500 index 2
1704         inet6 2001:0db8:3c4d:55:a00:20ff:fe8e:f3ad/64
```

1709 Even if there are no prefix advertisements on the link, you
1710 can still assign global addresses manually, for example:

```
1712     example% ifconfig eri0 inet6 addif \
1713     2001:0db8:3c4d:55:a00:20ff:fe8e:f3ad/64 up
```

1718 To configure boot-time defaults for the interface eri0,
1719 place the following entry in the /etc/hostname6.eri0 file:

```
1721     addif 2001:0db8:3c4d:55:a00:20ff:fe8e:f3ad/64 up
```

1724 Configuring IPv6/IPv4 tunnels

1725 An IPv6 over IPv4 tunnel interface can send and receive IPv6
1726 packets encapsulated in an IPv4 packet. Create tunnels at
1727 both ends pointing to each other. IPv6 over IPv4 tunnels
1728 require the tunnel source and tunnel destination IPv4 and

```
1729 IPv6 addresses. Solaris 8 supports both automatic and con-
1730 figured tunnels. For automatic tunnels, an IPv4-compatible
1731 IPv6 address is used. The following demonstrates auto-tunnel
1732 configuration:
```

```
1734     example% ifconfig ip.atun0 inet6 plumb
1735     example% ifconfig ip.atun0 inet6 tsrc IPv4-address \
1736     ::IPv4 address/96 up
```

1740 SunOS 5.11 Last change: 21 Jan 2007 26

1747 System Administration Commands ifconfig(1M)

```
1751 where IPv4-address is the IPv4 address of the interface
1752 through which the tunnel traffic will flow, and IPv4-
1753 address, ::<IPv4-address>, is the corresponding IPv4-
1754 compatible IPv6 address.
```

1757 The following is an example of a configured tunnel:

```
1759     example% ifconfig ip.tun0 inet6 plumb tsrc my-ipv4-address \
1760     tdst peer-ipv4-address up
```

1765 This creates a configured tunnel between my-ipv4-address and
1766 peer-ipv4-address with corresponding link-local addresses.
1767 For tunnels with global or site-local addresses, the logical
1768 tunnel interfaces need to be configured in the following
1769 form:

```
1771     example% ifconfig ip.tun0 inet6 addif my-v6-address peer-v6-address up
```

1776 For example,

```
1778     example% ifconfig ip.tun0 inet6 plumb tsrc 109.146.85.57 \
1779     tdst 109.146.85.212 up
1780     example% ifconfig ip.tun0 inet6 addif 2::45 2::46 up
```

1785 To show all IPv6 interfaces that are up and configured:

```
1787     example% ifconfig -au6
1788     ip.tun0: flags=2200851<UP,POINTOPOINT,RUNNING,MULTICAST,NUD,IPv6>
1789         mtu 1480 index 3
1790         inet tunnel src 109.146.85.57 tunnel dst 109.146.85.212
1791         tunnel security settings --> use 'ipsecconf -ln -i ip.tun1'
1792         tunnel hop limit 60
1793         inet6 fe80::6d92:5539/10 --> fe80::6d92:55d4
1794     ip.tun0:1: flags=2200851<UP,POINTOPOINT,RUNNING,MULTICAST,NUD,IPv6>
```

```
1795          mtu 1480 index 3
1796          inet6 2::45/128 --> 2::46
```

1801 In the output above, note the line that begins with "tunnel
1802 security settings". The content of this line varies

1806 SunOS 5.11 Last change: 21 Jan 2007 27

1813 System Administration Commands ifconfig(1M)

1817 according to whether and how you have set your security set-
1818 tings. See "Display of Tunnel Security Settings," below.

1820 Configuring IPv4/IPv6 Tunnels

1821 An IPv4 over IPv6 tunnel interface can send and receive IPv4
1822 packets encapsulated in an IPv6 packet. Create tunnels at
1823 both ends pointing to each other. IPv4 over IPv6 tunnels
1824 require the tunnel source and tunnel destination IPv6 and
1825 IPv4 addresses. The following demonstrates auto-tunnel con-
1826 figuration:

```
1828 example% ifconfig ip6.tun0 inet plumb tsrc my-ipv6-address \  
1829         tdst peer-ipv6-address my-ipv4-address \  
1830         peer-ipv4-address up
```

1835 This creates a configured tunnel between my-ipv6-address and
1836 peer-ipv6-address with my-ipv4-address and peer-ipv4-address
1837 as the endpoints of the point-to-point interface, for exam-
1838 ple:

```
1840 example% ifconfig ip6.tun0 inet plumb tsrc fe80::1 tdst fe80::2 \  
1841         10.0.0.208 10.0.0.210 up
```

1846 To show all IPv4 interfaces that are up and configured:

```
1848 example% ifconfig -au4
1849 lo0: flags=1000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4> mtu 8232 index 1
1850     inet 127.0.0.1 netmask ff000000
1851 eri0: flags=1004843<UP,BROADCAST,RUNNING,MULTICAST,DHCP,IPv4> mtu 1500 \  
1852     index 2
1853     inet 172.17.128.208 netmask ffffffff broadcast 172.17.128.255
1854 ip6.tun0: flags=10008d1<UP,POINTOPOINT,RUNNING,NOARP,MULTICAST,IPv4> \  
1855     mtu 1460
1856     index 3
1857     inet6 tunnel src fe80::1 tunnel dst fe80::2
1858     tunnel security settings --> use 'ipsecconf -ln -i ip.tun1'
1859     tunnel hop limit 60 tunnel encapsulation limit 4
1860     inet 10.0.0.208 --> 10.0.0.210 netmask ff000000
```

1865 In the output above, note the line that begins with "tunnel
1866 security settings". The content of this line varies accord-
1867 ing to whether and how you have set your security settings.
1868 See "Display of Tunnel Security Settings," below.

1872 SunOS 5.11 Last change: 21 Jan 2007 28

1879 System Administration Commands ifconfig(1M)

1883 Display of Tunnel Security Settings

1884 The ifconfig output for tunneled interfaces indicates secu-
1885 rity settings, if present, for a tunnel. The content of the
1886 line showing your settings differs depending on how you have
1887 made your settings:

1889 o If you set your security policy using the ifconfig
1890 -auth_algs, -encr_algs, and -encr_auth_algs options
1891 and do not use ipsecconf(1M), ifconfig displays
1892 your settings for each of these options.

1894 o If you set your security policy using ipsecconf(1M)
1895 with the tunnel keyword (the preferred method),
1896 ifconfig displays:

```
1898         tunnel security settings --> use 'ipsecconf -ln -i ip.tun1'
```

1901 ...in effect, hiding your settings from those
1902 without privileges to view them.

1904 If you do not set security policy, using either
1905 ifconfig or ipsecconf, there is no tunnel security
1906 setting displayed.

1908 EXAMPLES

1909 Example 1 Using the ifconfig Command

1912 If your workstation is not attached to an Ethernet, the net-
1913 work interface, for example, eri0, should be marked "down"
1914 as follows:

```
1917 example% ifconfig eri0 down
```

1921 Example 2 Printing Addressing Information

1924 To print out the addressing information for each interface,
1925 use the following command:

1928 example% ifconfig -a

1932 Example 3 Resetting the Broadcast Address

1938 SunOS 5.11 Last change: 21 Jan 2007 29

1945 System Administration Commands ifconfig(1M)

1949 To reset each interface's broadcast address after the net-
1950 masks have been correctly set, use the next command:

1953 example% ifconfig -a broadcast +

1957 Example 4 Changing the Ethernet Address

1960 To change the Ethernet address for interface ce0, use the
1961 following command:

1964 example% ifconfig ce0 ether aa:1:2:3:4:5

1968 Example 5 Configuring an IP-in-IP Tunnel

1971 To configure an IP-in-IP tunnel, first plumb it with the
1972 following command:

1975 example% ifconfig ip.tun0 plumb

1980 Then configure it as a point-to-point interface, supplying
1981 the tunnel source and the tunnel destination:

1984 example% ifconfig ip.tun0 myaddr mydestaddr tsrc another_myaddr \
1985 tdst a_dest_addr up

1990 Use ipsecconf(1M), as described above, to configure tunnel
1991 security properties.

1994 Example 6 Configuring 6to4 Tunnels

1997 To configure 6to4 tunnels, use the following commands:

2000 example% ifconfig ip.6to4tun0 inet6 plumb

2004 SunOS 5.11 Last change: 21 Jan 2007 30

2011 System Administration Commands ifconfig(1M)

2015 example% ifconfig ip.6to4tun0 inet6 tsrc IPv4-address 6to4-address/64 up

2020 IPv4-address denotes the address of the encapsulating inter-
2021 face. 6to4-address denotes the address of the local IPv6
2022 address of form 2002:IPv4-address:SUBNET-ID:HOSTID.

2026 The long form should be used to resolve any potential con-
2027 flicts that might arise if the system administrator utilizes
2028 an addressing plan where the values for SUBNET-ID or HOSTID
2029 are reserved for something else.

2033 After the interface is plumbed, a 6to4 tunnel can be config-
2034 ured as follows:

2037 example% ifconfig ip.6to4tun0 inet6 tsrc IPv4-address up

2042 This short form sets the address. It uses the convention:

2045 2002:IPv4-address::1

2049 The SUBNET-ID is 0, and the HOSTID is 1.

2052 Example 7 Configuring IP Forwarding on an Interface

2055 To enable IP forwarding on a single interface, use the fol-
2056 lowing command:

2059 example% ifconfig eri0 router

2064 To disable IP forwarding on a single interface, use the fol-
2065 lowing command:

2070 SunOS 5.11 Last change: 21 Jan 2007 31

2077 System Administration Commands ifconfig(1M)

2081 example% ifconfig eri0 -router

2085 Example 8 Configuring Source Address Selection Using a Vir-
2086 tual Interface

2089 The following command configures source address selection
2090 such that every packet that is locally generated with no
2091 bound source address and going out on qfe2 prefers a source
2092 address hosted on vni0.

2095 example% ifconfig qfe2 usesrc vni0

2100 The ifconfig -a output for the qfe2 and vni0 interfaces
2101 displays as follows:

```
2104 qfe2: flags=1100843<UP,BROADCAST,RUNNING,MULTICAST,ROUTER,IPv4> mtu
2105 1500 index 4
2106 usesrc vni0
2107 inet 1.2.3.4 netmask fffffff0 broadcast 1.2.3.255
2108 ether 0:3:ba:17:4b:e1
2109 vni0: flags=20011100c1<UP,RUNNING,NOARP,NOXMIT,ROUTER,IPv4,VIRTUAL>
2110 mtu 0 index 5
2111 srcof qfe2
2112 inet 3.4.5.6 netmask ffffffff
```

2116 Observe, above, the usesrc and srcof keywords in the ifcon-
2117 fig output. These keywords also appear on the logical
2118 instances of the physical interface, even though this is a
2119 per-physical interface parameter. There is no srcof keyword
2120 in ifconfig for configuring interfaces. This information is
2121 determined automatically from the set of interfaces that
2122 have usesrc set on them.

2126 The following command, using the none keyword, undoes the
2127 **effect of the preceding *ifconfig* *usesrc* command.**
2100 *effect of the preceding ifconfig usersrc command.*

2130 example% ifconfig qfe2 usesrc none

2136 SunOS 5.11 Last change: 21 Jan 2007 32

2143 System Administration Commands ifconfig(1M)

2147 Following this command, ifconfig -a output displays as fol-
2148 lows:

```
2151 qfe2: flags=1100843<UP,BROADCAST,RUNNING,MULTICAST,ROUTER,IPv4> mtu
2152 1500 index 4
2153 inet 1.2.3.4 netmask fffffff0 broadcast 1.2.3.255
2154 ether 0:3:ba:17:4b:e1
2155 vni0: flags=20011100c1<UP,RUNNING,NOARP,NOXMIT,ROUTER,IPv4,VIRTUAL>
2156 mtu 0 index 5
2157 inet 3.4.5.6 netmask ffffffff
```

2161 Note the absence of the usesrc and srcof keywords in the
2162 output above.

2165 Example 9 Configuring Source Address Selection for an IPv6
2166 Address

2169 The following command configures source address selection
2170 for an IPv6 address, selecting a source address hosted on
2171 vni0.

2174 example% ifconfig qfe1 inet6 usesrc vni0

2179 Following this command, ifconfig -a output displays as fol-
2180 lows:

```
2183 qfe1: flags=2000841<UP,RUNNING,MULTICAST,IPv6> mtu 1500 index 3
2184 usesrc vni0
2185 inet6 fe80::203:baff:fe17:4be0/10
2186 ether 0:3:ba:17:4b:e0
2187 vni0: flags=2002210041<UP,RUNNING,NOXMIT,NONUD,IPv6,VIRTUAL> mtu 0
2188 index 5
2189 srcof qfe1
```

```

2190      inet6 fe80::203:baff:fe17:4444/128
2191      vni0:1: flags=2002210040<RUNNING,NOXMIT,IPV6,VIRTUAL> mtu 0
2192      index 5
2193      srcof qfel
2194      inet6 fec0::203:baff:fe17:4444/128
2195      vni0:2: flags=2002210040<RUNNING,NOXMIT,NOXMIT,IPV6,VIRTUAL> mtu 0
2196      index 5
2197      srcof qfel
2198      inet6 2000::203:baff:fe17:4444/128

2202 SunOS 5.11      Last change: 21 Jan 2007      33

2209 System Administration Commands      ifconfig(1M)

2213      Depending on the scope of the destination of the packet
2214      going out on qfel, the appropriately scoped source address
2215      is selected from vni0 and its aliases.

2218      Example 10 Using Source Address Selection with Shared-IP
2219      Zones

2222      The following is an example of how the usesrc feature can be
2223      used with the zones(5) facility in Solaris. The following
2224      commands are invoked in the global zone:

2227      example% ifconfig hme0 usesrc vni0
2228      example% ifconfig eri0 usesrc vni0
2229      example% ifconfig qfe0 usesrc vni0

2234      Following the preceding commands, the ifconfig -a output for
2235      the virtual interfaces would display as:

2238      vni0: flags=20011100c1<UP,RUNNING,NOARP,NOXMIT,ROUTER,IPV4,VIRTUAL>
2239      mtu 0 index 23
2240      srcof hme0 eri0 qfe0
2241      inet 10.0.0.1 netmask ffffffff
2242      vni0:1:
2243      flags=20011100c1<UP,RUNNING,NOARP,NOXMIT,ROUTER,IPV4,VIRTUAL> mtu 0
2244      index 23
2245      zone test1
2246      srcof hme0 eri0 qfe0
2247      inet 10.0.0.2 netmask ffffffff
2248      vni0:2:
2249      flags=20011100c1<UP,RUNNING,NOARP,NOXMIT,ROUTER,IPV4,VIRTUAL> mtu 0
2250      index 23
2251      zone test2
2252      srcof hme0 eri0 qfe0
2253      inet 10.0.0.3 netmask ffffffff
2254      vni0:3:
2255      flags=20011100c1<UP,RUNNING,NOARP,NOXMIT,ROUTER,IPV4,VIRTUAL> mtu 0

```

```

2256      index 23
2257      zone test3
2258      srcof hme0 eri0 qfe0
2259      inet 10.0.0.4 netmask ffffffff

2263      There is one virtual interface alias per zone (test1, test2,
2264      and test3). A source address from the virtual interface

2268 SunOS 5.11      Last change: 21 Jan 2007      34

2275 System Administration Commands      ifconfig(1M)

2279      alias in the same zone is selected. The virtual interface
2280      aliases were created using zonecfg(1M) as follows:

2283      example% zonecfg -z test1
2284      zonecfg:test1> add net
2285      zonecfg:test1:net> set physical=vni0
2286      zonecfg:test1:net> set address=10.0.0.2

2291      The test2 and test3 zone interfaces and addresses are
2292      created in the same way.

2295      Example 11 Turning Off DHCPv6

2298      The following example shows how to disable automatic use of
2299      DHCPv6 on all interfaces, and immediately shut down DHCPv6
2300      on the interface named hme0. See in.ndpd(1M) and
2301      ndpd.conf(4) for more information on the automatic DHCPv6
2302      configuration mechanism.

2305      example% echo ifdefault StatefulAddrConf false >> /etc/inet/ndpd.conf
2306      example% pkill -HUP -x in.ndpd
2307      example% ifconfig hme0 dhcp release

2311 FILES
2312      /etc/netmasks

2314      Netmask data.

2317      /etc/default/inet_type

2319      Default Internet protocol type.

```

2322 ATTRIBUTES
 2323 See attributes(5) for descriptions of the following attri-
 2324 butes:

2334 SunOS 5.11 Last change: 21 Jan 2007 35

2341 System Administration Commands ifconfig(1M)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu
Interface Stability for command-line options	Committed
Interface Stability for command output	Uncommitted

2357 SEE ALSO
 2358 dhcpinfo(1), dhcpagent(1M), in.mpathd(1M), in.ndpd(1M),
 2359 **in.routed(1M), ippstat(1M), ipsecconf(1M), netstat(1M),**
 2332 *in.routed(1M), ipsecconf(1M), ndd(1M), netstat(1M),*
 2360 zoneadm(1M), zonecfg(1M), ethers(3SOCKET),
 2361 gethostbyname(3NSL), getnetbyname(3SOCKET), hosts(4),
 2362 inet_type(4), ndpd.conf(4), netmasks(4), networks(4),
 2363 nsswitch.conf(4), attributes(5), privileges(5), zones(5),
 2364 arp(7P), ipsec(7P), ipsecconf(7P), tun(7M)

2367 DIAGNOSTICS
 2368 ifconfig sends messages that indicate if:

- 2370 o the specified interface does not exist
- 2372 o the requested address is unknown
- 2374 o the user is not privileged and tried to alter an
 2375 interface's configuration

2377 NOTES
 2378 Do not select the names broadcast, down, private, trailers,
 2379 up or other possible option names when you choose host
 2380 names. If you choose any one of these names as host names,
 2381 it can cause unusual problems that are extremely difficult
 2382 to diagnose.

2400 SunOS 5.11 Last change: 21 Jan 2007 36

```

*****
13177 Fri Dec 19 04:18:55 2008
new/./in.mpathd.lm.txt
Clearview IPMP manpages
*****
1 System Administration Commands          in.mpathd(1M)

5 NAME
6 in.mpathd - IP multipathing daemon
6 in.mpathd - daemon for network adapter (NIC) failure detec-
7 tion, recovery, automatic failover and fallback

8 SYNOPSIS
9 /usr/lib/inet/in.mpathd

12 DESCRIPTION
14 The in.mpathd daemon performs Network Interface Card (NIC)
15 failure and repair detection. In the event of a NIC failure,
16 it causes IP network access from the failed NIC to failover
17 to a standby NIC, if available, or to any another opera-
18 tional NIC that has been configured as part of the same net-
19 work multipathing group. Once the failed NIC is repaired,
20 all network access is restored to the repaired NIC.

14 The *in.mpathd* daemon performs failure and repair detection
15 for IP interfaces that have been placed into an IPMP group
16 (or optionally, for all IP interfaces on the system). It
17 also controls which IP interfaces in an IPMP group are
18 "active" (being used by the system to send or receive IP
19 data traffic) in a manner which is consistent with the
20 administrator's configured policy.

23 The *in.mpathd* daemon can detect IP interface failure and
24 repair through two methods: by monitoring the *IFF_RUNNING*
25 flag for each IP interface (link-based failure detection),
26 and by sending and receiving ICMP probes on each IP
27 interface (probe-based failure detection). Link-based
28 failure detection is instantaneous and is always enabled
29 (provided the network driver supports the feature);
23 The in.mpathd daemon can detect NIC failure and repair
24 through two methods: by monitoring the IFF_RUNNING flag for
25 each NIC (link-based failure detection), and by sending and
26 receiving ICMP echo requests and replies on each NIC
27 (probe-based failure detection). Link-based failure detec-
28 tion requires no explicit configuration and thus is always
29 enabled (provided the NIC driver supports the feature);
30 probe-based failure detection must be enabled through the
31 configuration of one or more test addresses (described
32 below), but tests the entire IP interface send and receive
33 path. The *ipmpstat(1M)* utility can be used to check which
34 failure detection methods are enabled.
32 below), but has the benefit of testing the entire NIC send
33 and receive path.

37 If only link-based failure detection is enabled, then the
38 health of the interface is determined solely from the state
39 of the IFF_RUNNING flag. Otherwise, the interface is con-
40 sidered failed if either of the two methods indicate a
41 failure, and repaired once both methods indicate the failure
42 has been corrected. Not all interfaces in a group need to be
43 configured with the same failure detection methods.

```

```

46 As mentioned above, to perform probe-based failure detection
47 *in.mpathd* requires a test address on each IP interface for
48 the purpose of sending and receiving probes. Each address
49 must be marked *NOFAILOVER* (see *ifconfig(1M)*) and
50 *in.mpathd* will be limited to probing targets on the same
51 subnet. Each address may be configured statically or
52 acquired via DHCP. To find targets, *in.mpathd* first
53 consults the routing table for routes on the same subnet,
54 and uses the specified next-hop. If no routes match, it
55 sends all-hosts ICMP probes and selects a subset of the
56 systems that respond. Thus, for probe-based failure
57 detection to operate, there must be at least one neighbor on
58 each subnet that responds to ICMP echo request probes. The
59 *ipmpstat(1M)* utility can be used to display both the
60 current probe target information and the status of sent
61 probes.
45 As mentioned above, in order to perform probe-based failure
46 detection in.mpathd needs a special test address on each NIC
47 for the purpose of sending and receiving probes on the NIC.
48 Use the ifconfig command -failover option to configure these
49 test addresses. See ifconfig(1M). The test address must
50 belong to a subnet that is known to the hosts and routers on
51 the link.
62

64 Both IPv4 and IPv6 are supported. If an IP interface is
65 plumbed for IPv4 and an IPv4 test address is configured then
66 *in.mpathd* will start sending ICMPv4 probes over that IP
67 interface. Similarly, if an IP interface is plumbed for
68 IPv6 and an IPv6 test address is configured then *in.mpathd*
69 will start sending ICMPv6 probes over that IP interface.
70 However, note that *in.mpathd* will ignore IPv6 test
71 addresses that are not link-local. If both IPv4 and IPv6
72 are plumbed, it is sufficient to configure only one of the
73 two, that is, either an IPv4 test address or an IPv6 test
74 address. If both IPv4 and IPv6 test addresses are
75 configured, *in.mpathd* probes using both ICMPv4 and ICMPv6.
54 The in.mpathd daemon can detect NIC failure and repair by
55 two methods, by sending and receiving ICMP echo requests and
56 replies on each NIC, and by monitoring the IFF_RUNNING flag

```

```

78 As mentioned above, *in.mpathd* also controls which IP
79 interfaces in an IPMP group are "active" (used by the system
80 to send and receive IP data traffic). Specifically,
81 *in.mpathd* tracks the administrative configuration of each
82 IPMP group and attempts to keep the number of active IP
83 interfaces in each group consistent with that configuration.
84 Therefore, if an active IP interface fails, *in.mpathd* will
85 activate an *INACTIVE* interface in the group, provided one
86 exists (it will prefer *INACTIVE* interfaces that are also
87 marked *STANDBY*). Likewise, if an IP interface repairs and
88 the resulting repair leaves the IPMP group with more active
89 interfaces than the administrative configuration specifies,
90 *in.mpathd* will deactivate one of the interfaces
91 (preferably one marked *STANDBY*), except when the
92 *FAILBACK* variable is used, as described below. Similar
93 adjustments will be made by *in.mpathd* when offlining IP
94 interfaces (for instance, in response to *if_mpadm(1M)*).

```

67 System Administration Commands *in.mpathd(1M)*

71 for each NIC. The link state on some models of NIC is indicated by the *IFF_RUNNING* flag, allowing for faster failure detection when the link goes down. The *in.mpathd* daemon considers a NIC to have failed if either of the above two methods indicates failure. A NIC is considered to be repaired only if both methods indicate the NIC is repaired.

79 The *in.mpathd* daemon sends the ICMP echo request probes to on-link routers. If no routers are available, it sends the probes to neighboring hosts. Thus, for network failure detection and repair, there must be at least one neighbor on each link that responds to ICMP echo request probes.

86 *in.mpathd* works on both IPv4 and IPv6. If IPv4 is plumbed on a NIC, an IPv4 test address is configured on the NIC, and the NIC is configured as part of a network multipathing group, then *in.mpathd* will start sending ICMP probes on the NIC using IPv4.

93 In the case of IPv6, the link-local address must be configured as the test address. The *in.mpathd* daemon will not accept a non-link-local address as a test address. If the NIC is part of a multipathing group, and the test address has been configured, then *in.mpathd* will probe the NIC for failures using IPv6.

101 Even if both the IPv4 and IPv6 protocol streams are plumbed, it is sufficient to configure only one of the two, that is, either an IPv4 test address or an IPv6 test address on a NIC. If only an IPv4 test address is configured, it probes using only ICMPv4. If only an IPv6 test address is configured, it probes using only ICMPv6. If both type test addresses are configured, it probes using both ICMPv4 and ICMPv6.

97 The *in.mpathd* daemon accesses three variable values in */etc/default/mpathd*: *FAILURE_DETECTION_TIME*, *FAILBACK* and *TRACK_INTERFACES_ONLY_WITH_GROUPS*.

102 The **FAILURE_DETECTION_TIME** variable specifies the probe-based failure detection time. The shorter the failure detection time, the more probe traffic. The default value of **FAILURE_DETECTION_TIME** is 10 seconds. This means that IP interface failure will be detected by **in.mpathd** within 10 seconds. The IP interface repair detection time is always twice the value of **FAILURE_DETECTION_TIME**. Note that failures and repairs detected by link-based failure detection are acted on immediately, though **in.mpathd** may ignore link state changes if it suspects that the link state is flapping due to defective hardware; see *DIAGNOSTICS*. The *FAILURE_DETECTION_TIME* variable specifies the NIC failure detection time for the ICMP echo request probe method of detecting NIC failure. The shorter the failure

119 detection time, the greater the volume of probe traffic. The default value of *FAILURE_DETECTION_TIME* is 10 seconds. This means that NIC failure will be detected by *in.mpathd* within 10 seconds. NIC failures detected by the *IFF_RUNNING* flag

115 By default, **in.mpathd** limits failure and repair detection to IP interfaces that are configured as part of a named IPMP group. Setting **TRACK_INTERFACES_ONLY_WITH_GROUPS** to **no** enables failure and repair detection on all IP interfaces, even if they are not part of a named IPMP group. IP interfaces that are tracked but not part of a named IPMP group are considered to be part of the "anonymous" IPMP group. In addition to having no name, this IPMP group is special in that its IP interfaces are not equivalent and thus cannot take over for one another in the event of an IP interface failure. That is, the anonymous IPMP group can only be used for failure and repair detection, and provides no high-availability or load-spreading.

126 SunOS 5.11 Last change: 8 Sep 2006 2

130 As described above, when **in.mpathd** detects that an IP interface has repaired, it activates it so that it will again be used to send and receive IP data traffic. However, if **FAILBACK** is set to **no**, then the IP interface will only be activated if no other active IP interfaces in the group remain. However, the interface may subsequently be activated if another IP interface in the group fails.

133 System Administration Commands *in.mpathd(1M)*

137 being cleared are acted on as soon as the *in.mpathd* daemon notices the change in the flag. The NIC repair detection time cannot be configured; however, it is defined as double the value of *FAILURE_DETECTION_TIME*.

143 By default, *in.mpathd* does failure detection only on NICs that are configured as part of a multipathing group. You can set *TRACK_INTERFACES_ONLY_WITH_GROUPS* to *no* to enable failure detection by *in.mpathd* on all NICs, even if they are not part of a multipathing group. However, *in.mpathd* cannot do failover from a failed NIC if it is not part of a multipathing group.

152 The *in.mpathd* daemon will restore network traffic back to the previously failed NIC, after it has detected a NIC repair. To disable this, set the value of *FAILBACK* to *no* in */etc/default/mpathd*.

138 FILES
139 */etc/default/mpathd* Contains default values used by the *in.mpathd* daemon.
140

143 ATTRIBUTES
144 See *attributes(5)* for descriptions of the following attributes:
145

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsr

156 SEE ALSO
 157 `ifconfig(1M), ipmpstat(1M), if_mpadm(1M), icmp(7P), icmp6(7P)`
 176 `ifconfig(1M), attributes(5), icmp(7P), icmp6(7P),`

160 DIAGNOSTICS

161 **IP interface `*interface_name*` has a hardware address which**
 162 **is not unique in group `*group_name*`; offlining**
 180 *Test address address is not unique; disabling probe based*
 181 *failure detection on interface_name*
 163 Description:

165 For probe-based failure detection, load-spreading, and
 166 other code IPMP features to work properly, each IP
 167 interface in an IPMP group must have a unique hardware
 168 address. If this requirement is not met, `*in.mpathd*`
 169 will automatically offline all but one of the IP
 170 interfaces with duplicate hardware addresses.

172 **IP interface `*interface_name*` now has a unique hardware**
 173 **address in group `*group_name*`; onlining**
 174 Description:

176 The previously-detected duplicate hardware address is
 177 now unique, and therefore `*in.mpathd*` has brought
 178 `*interface_name*` back online.

181 **Test address `*address*` is not unique in group; disabling**
 182 **probe-based failure detection on `*interface_name*`**
 183 Description:

186 For `in.mpathd` to perform probe-based failure detection,
 187 **each test address in the group must be unique.**
 186 *each test address in the group must be unique. Since the*
 187 *IPv6 test address is a link-local address derived from*
 188 *the MAC address, each IP interface in the group must*

190 SunOS 5.11 Last change: 8 Sep 2006 3

197 System Administration Commands `in.mpathd(1M)`

201 **No test address configured on interface `*interface_name*`;**
 202 **disabling probe-based failure detection on it**
 203 Description:
 203 *have a unique MAC address.*

206 For `*in.mpathd*` to perform probe-based failure detection
 207 on an IP interface, it must be configured with a test
 208 address: IPv4, IPv6, or both.

211 NIC interface_name of group group_name is not plumbed for
 212 IPv[4|6] and may affect failover capability
 213 Description:

216 All NICs in a multipathing group must be homogeneously
 217 plumbed. For example, if a NIC is plumbed for IPv4, then
 218 all NICs in the group must be plumbed for IPv4. The
 219 **STREAMS modules pushed on all NICs must also be identical.**
 215 *streams modules pushed on all NICs must be identical.*

219 No test address configured on interface interface_name disa-
 220 bling probe-based failure detection on it
 221 Description:

224 In order for `in.mpathd` to perform probe-based failure
 225 detection on a NIC, it must be configured with a test
 226 address: IPv4, IPv6, or both.

222 The link has come up on interface_name more than 2 times in
 223 **the last minute; disabling repair until it stabilizes.**
 231 *the last minute; disabling failback until it stabilizes.*
 224 Description:

227 To limit the impact of interfaces with intermittent
 228 hardware (such as a bad cable), `*in.mpathd*` will not
 229 consider an IP interface with a frequently changing link
 230 state as repaired until the link state stabilizes.
 235 In order to prevent interfaces with intermittent
 236 hardware, such as a bad cable, from causing repeated
 237 failovers and failbacks, `in.mpathd` does not failback to
 238 interfaces with frequently fluctuating link states.

234 **Invalid failure detection time of `*time*`, assuming default**
 235 **of 10000 ms**
 242 *Invalid failure detection time assuming default 10000*
 236 Description:

239 An invalid value was encountered for
 240 FAILURE_DETECTION_TIME in the `/etc/default/mpathd` file.

244 **Too small failure detection time of `*time*`, assuming minimum**
 245 **of 100 ms**
 251 *Too small failure detection time of time assuming minimum*
 252 *100*
 246 Description:

251 SunOS 5.11 Last change: 8 Sep 2006 4

258 System Administration Commands in.mpathd(1M)

262 The minimum value that can be specified for
263 FAILURE_DETECTION_TIME is currently 100 milliseconds.

267 Invalid value for FAILBACK value
268 Description:

271 Valid values for the boolean variable FAILBACK are yes
272 or no.

276 Invalid value for TRACK_INTERFACES_ONLY_WITH_GROUPS value
277 Description:

280 Valid values for the boolean variable
281 TRACK_INTERFACES_ONLY_WITH_GROUPS are yes or no.

285 Cannot meet requested failure detection time of time ms on
286 (inet[6] interface_name) new failure detection time for
287 group group_name is time ms
288 Description:

291 The round trip time for ICMP probes is higher than
292 necessary to maintain the current failure detection
293 time. The network is probably congested or the probe
294 targets are loaded. in.mpathd automatically increases
295 the failure detection time to whatever it can achieve
296 under these conditions.

300 Improved failure detection time time ms on (inet[6]
301 interface_name) for group group_name
302 Description:

305 The round trip time for ICMP probes has now decreased
306 and in.mpathd has lowered the failure detection time
307 correspondingly.

311 **IP interface failure detected on interface_name**
318 *NIC failure detected on interface_name*
312 Description:

317 SunOS 5.11 Last change: 8 Sep 2006 5

324 System Administration Commands in.mpathd(1M)

328 ***in.mpathd* has detected a failure on *interface_name*,
329 and has set the *IFF_FAILED* flag on *interface_name*,
330 ensuring that it will not be used for IP data traffic.
335 in.mpathd has detected NIC failure on interface_name,
336 and has set the IFF_FAILED flag on NIC interface_name.**

333 **IP interface repair detected on *interface_name***

340 *Successfully failed over from NIC interface_name1 to NIC
341 interface_name2*
334 Description:

337 ***in.mpathd* has detected a repair on *interface_name*,
338 and has cleared the *IFF_FAILED* flag. Depending on the
339 administrative configuration, the *interface_name* may
340 again be used for IP data traffic.
345 in.mpathd has caused the network traffic to failover
346 from NIC interface_name1 to NIC interface_name2, which
347 is part of the multipathing group.**

351 *NIC repair detected on interface_name*
352 Description:

355 *in.mpathd has detected that NIC interface_name is
356 repaired and operational. If the IFF_FAILED flag on the
357 NIC was previously set, it will be reset.*

361 *Successfully failed back to NIC interface_name*
362 Description:

365 *in.mpathd has restored network traffic back to NIC
366 interface_name, which is now repaired and operational.*

343 The link has gone down on interface_name
344 Description:

347 ***in.mpathd* has detected that the *IFF_RUNNING* flag for
348 *interface_name* has been cleared, indicating the link
349 has gone down.
374 in.mpathd has detected that the IFF_RUNNING flag for NIC
375 interface_name has been cleared, indicating the link has**

376 *gone down.*

353 The link has come up on interface_name
354 Description:

357 **in.mpathd* has detected that the *IFF_RUNNING* flag for*
358 **interface_name* has been set, indicating the link has*
384 *in.mpathd has detected that the IFF_RUNNING flag for NIC*
385 *interface_name has been set, indicating the link has*
359 *come up.*

363 SunOS 5.11 Last change: 8 Sep 2006 6

370 System Administration Commands in.mpathd(1M)

426 SunOS 5.11 Last change: 8 Sep 2006 7

```

*****
16248 Fri Dec 19 04:18:56 2008
new/./ipmpstat.lm.txt
*** NO COMMENTS ***
*****
1 System Administration Commands          ipmpstat(1M)
3 NAME
4     ipmpstat - display IPMP subsystem status
6 SYNOPSIS
8     ipmpstat [-n] [-o _field_ [-P]] -a|-g|-i|-p|-t
10 DESCRIPTION
12     The *ipmpstat* command concisely displays information about
13     the IPMP subsystem. It supports five different output
14     modes, each of which provides a different view of the IPMP
15     subsystem (address, group, interface, probe, and target),
16     described below. At most one output mode may be specified
17     per invocation, and the displayed information is guaranteed
18     to be self-consistent. It also provides a parsable output
19     format which may be used by scripts to examine the state of
20     the IPMP subsystem. Only basic privileges are needed to
21     invoke ipmpstat, with the exception of probe mode which
22     requires all privileges.
24 OPTIONS
26     -a
28         Display IPMP data address information ("address"
29         output mode).
31     -g
33         Display IPMP group information ("group" output mode).
35     -i
37         Display IP interface information ("interface" output
38         mode).
40     -n
42         Display IP addresses numerically, rather than
43         attempting to resolve them to hostnames. This option
44         may be used in any output mode.
47     -o _field_[,...]
49         Display only the specified output fields, in order.
50         The list of field names is case-insensitive and
51         comma-separated. The field names that are supported
52         depend on the selected output mode, described below.
53         The special field name "all" may be used to print all
54         fields for a given output mode.
56     -p
58         Display IPMP probe information ("probe" output mode).
60     -t

```

```

62         Display IPMP target information ("target" output mode).
64     -P
66         Display using a machine-parsable format, described
67         below. If this option is specified, an explicit list
68         of fields must be specified using the *-o* option.
71 OUTPUT MODES
73     Address Mode
75         Address mode displays the state of all IPMP data
76         addresses on the system. The following output fields are
77         supported:
79     ADDRESS      The hostname (or IP address) associated with
80                 the information. Note that because duplicate
81                 down addresses may exist, the address must be
82                 taken together with the GROUP to form a
83                 unique identity. For a given IPMP group, if
84                 duplicate addresses exist, at most one will
85                 be displayed, and an up address will always
86                 take precedence.
88     STATE        The state of the address. Either *up* if the
89                 address is IFF_UP (see ifconfig(1M)), or
90                 *down* if the address is not IFF_UP.
92     GROUP        The IPMP IP interface hosting the address.
94     INBOUND      The underlying IP interface that will receive
95                 packets for this address. This may change in
96                 response to external events such as IP
97                 interface failure. If this field is empty,
98                 then the system will not accept IP packets
99                 sent to this address (e.g., because the
100                address is down or because there are no
101                active IP interfaces left in the IPMP group).
103     OUTBOUND     The underlying IP interfaces that will send
104                 packets using this source address. This may
105                 change in response to external events such as
106                 IP interface failure. If this field is
107                 empty, then the system will not send packets
108                 with this address as a source (e.g., because
109                 the address is down or because there are no
110                 active IP interfaces left in the IPMP group).
111
112         If *-o* is not specified, all output fields are displayed.
113
114     Group Mode
115
116         Group mode displays the state of all IPMP groups on the
117         system. The following output fields are supported:
118     GROUP        The IPMP IP interface name associated with
119                 the information. For the anonymous group
120                 (see *in.mpathd(1M)*), this field will be
121                 empty.
122
124     GROUPNAME    The IPMP group name. For the anonymous
125                 group, this field will be empty.
126
127     STATE        The state of the group:

```

```

129      *ok*      All interfaces in the group are
130                usable.
131      *degraded* Some (but not all) interfaces in
132                the group are usable.
133      *failed*   No interfaces in the group are
134                usable.

136      FDT       The probe-based failure detection time. If
137                probe-based failure detection is disabled,
138                this field will be empty.

140      INTERFACES The list of underlying IP interfaces in the
141                group. The list is divided into three parts:

143                1. Active interfaces are listed first and not
144                   enclosed in any brackets or parenthesis.
145                   Active interfaces are those being used by
146                   the system to send or receive data traffic.

148                2. *INACTIVE* interfaces are listed next and
149                   enclosed in parenthesis. *INACTIVE*
150                   interfaces are those that are functioning,
151                   but not being used according to
152                   administrative policy.

154                3. Unusable interfaces are listed last and
155                   enclosed in brackets. Unusable interfaces
156                   are those that cannot be used at all in
157                   their present configuration (e.g.,
158                   *FAILED* or *OFFLINE*).

160      If *-o* is not specified, all output fields are displayed.

162      Interface Mode

164      Interface mode displays the state of all IP interfaces
165      that are tracked by *in.mpathd* on the system. The
166      following output fields are supported:

168      INTERFACE The IP interface name associated with the
169                information.

171      ACTIVE     Either *yes* or *no*, depending on if the IP
172                interface is being used by the system for
173                IP data traffic.

175      GROUP      The IPMP IP interface associated with the IP
176                interface. For IP interfaces in the
177                anonymous group (see *in.mpathd(1M)*), this
178                field will be empty.

180      FLAGS      Assorted information about the IP interface:

182                i Unusable due to being *INACTIVE*.

184                s Marked *STANDBY*.

186                m Nominated to send/receive IPv4 multicast
187                for its IPMP group.

189                b Nominated to receive IPv4 broadcast for
190                its IPMP group.

192                M Nominated to send/receive IPv6 multicast
193                for its IPMP group.

```

```

195                d Unusable due to being *down*.

197                h Unusable due to being brought *OFFLINE* by
198                *in.mpathd* because of a duplicate
199                hardware address

201      LINK      The state of link-based failure detection:

203                *up*      The link is up.

205                *down*    The link is down.

207                *unknown* The network driver does not
208                detect link state changes.

210      PROBE     The state of probe-based failure detection:

212                *ok*      Probes detect no problems.

214                *failed*   Probes detect failure.

216                *unknown* Probes cannot be sent since no
217                suitable probe targets are known.

219                *disabled* Probes have been disabled because
220                a unique IP test address has not
221                been configured.

223      STATE     The overall state of the interface:

225                *ok*      The interface is online and
226                functioning properly based on the
227                configured failure detection
228                methods.

230                *failed*   The interface is online but has a
231                link state of *down* or a probe
232                state of *failed*.

234                *offline* The interface is offline.

236                *unknown* The interface is online but may
237                or may not be functioning because
238                the configured failure detection
239                methods are in *unknown* states.

241      If *-o* is not specified, all output fields are displayed.

243      Probe Mode

245      Probe mode displays information about the probes being
246      sent by *in.mpathd*. Unlike other output modes, this
247      mode runs until explicitly terminated using *^C*. The
248      following output fields are supported:

249      TIME       The time the probe was sent, relative to when
250                *ipmpstat* was started. If the probe was
251                sent prior to starting *ipmpstat*, the time
252                will be negative.

253

255      PROBE     An identifier representing the probe. The
256                identifier will start at zero and will
257                monotonically increment for each probe sent
258                by *in.mpathd* over a given interface. To
259                enable more detailed analysis by packet

```

```

260 monitoring tools, this identifier matches the
261 *icmp_seq* field of the ICMP probe packet.

263 INTERFACE The IP interface the probe was sent on.

265 TARGET The hostname (or IP address) of the target
266 the probe was sent to.

268 NETRTT The network round-trip-time for the probe.
269 This is the time between when the IP module
270 sends the probe and when the IP module
271 receives the ack. If *in.mpathd* has
272 concluded that the probe has been lost, this
273 field will be empty.

275 RTT The total round-trip-time for the probe.
276 This is the time between when *in.mpathd*
277 starts executing the code to send the probe,
278 and when it completes processing the ack. If
279 *in.mpathd* has concluded that the probe has
280 been lost, this field will be empty. Spikes
281 in the total round-trip time that are not
282 present in the network round-trip time
283 indicate that the local system itself is
284 overloaded.

286 RTTAVG The average round-trip-time to *TARGET* over
287 *INTERFACE*. This aids identification of
288 slow targets. If there is insufficient data
289 to calculate the average, this field will be
290 empty.

292 RTTDEV The standard deviation for the
293 round-trip-time to *TARGET* over *INTERFACE*.
294 This aids identification of jittery targets.
295 If there is insufficient data to calculate
296 the standard deviation, this field will be
297 empty.

299 If *-o* is not specified, all fields except for *RTTAVG*
300 and *RTTDEV* are displayed.

302 Target Mode
303
304 Target mode displays IPMP probe target information. The
305 following output fields are supported:

307 INTERFACE The IP interface name associated with the
308 information.

310 MODE The probe target discovery mode:

312 *routes* Probe targets found via the
313 routing table.

315 *multicast* Probe targets found via multicast
316 ICMP probes.

318 *disabled* Probe-based failure detection is
319 disabled.

321 TESTADDR The hostname (or IP address) that will be
322 used for sending and receiving probes. If a
323 unique test address has not been configured,
324 this field will be empty. Note that if an IP
325 interface is configured with both IPv4 and

```

```

326 IPv6 test addresses, probe target information
327 will be displayed separately for each test
328 address.

330 TARGETS A space-separated list of probe target
331 hostnames (or IP addresses), in firing order.
332 If no probe targets could be found, this
333 field will be empty.

335 If *-o* is not specified, all output fields are displayed.

337 OUTPUT FORMAT

339 By default, ipmpstat uses a human-friendly tabular format
340 for its output modes, where each row contains one or more
341 fields of information about a given object, which is in turn
342 uniquely identified by one or more of those fields. In this
343 format, a header identifying the fields is displayed above
344 the table (and after each screenful of information), fields
345 are separated by whitespace, empty fields are represented by
346 **-, and other visual aids are used.

347
348 Machine-parsable format also uses a tabular format, but is
349 designed to be efficient to programmatically parse.
350 Specifically, machine-parsable format differs from
351 human-friendly format in the following ways:

353 * No headers are displayed.

355 * Fields with empty values yield no output, rather than
356 **-.

358 * Fields are separated by a single *:, rather than
359 variable amounts of whitespace.

361 * If multiple fields are requested, and a literal *: or
362 \* occur in a field's value, they are escaped by
363 prefixing them with \*.
364
365 EXAMPLES

367 Use the machine-parsable output format to create a *ksh*
368 function that outputs the failure detection time of a given
369 IPMP IP interface:

371 getfdt() {
372     ipmpstat -gP -o group,fdt | while IFS=: read group fdt; do
373         [[ "$group" = "$1" ]] && { echo "$fdt"; return; }
374     done
375 }

377 ATTRIBUTES

379 See attributes(5) for descriptions of the following
380 attributes:

382 /usr/sbin

384
385
386
387
388
389
390
391

```

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu
Interface Stability	Committed
Machine-Parsable Format	Committed

392
393
394

Human-Friendly Format

Not-an-Interface

397 **SEE ALSO**

399 `in.mpathd(1M), ifconfig(1M), if_mpadm(1M)`

20769 Fri Dec 19 04:18:57 2008

new/./route.lm.txt

Clearview IPMP manpages

```

1 System Administration Commands          route(1M)

5 NAME
6   route - manually manipulate the routing tables

8 SYNOPSIS
9   route [-fnvq] sub-command [ [modifiers] args]

12  route [-fnvq] [-p [-R root-dir]] add | delete [modifiers] destination gatew
13    [args]

16  route [-fnvq] change | get [modifiers] destination
17    [gateway [args]]

20  route [-fn] monitor [modifiers]

23  route [-fnvq] flush [modifiers]

26  route -p [-R root-dir] show

29 DESCRIPTION
30  route manually manipulates the network routing tables. These
31  tables are normally maintained by the system routing daemon,
32  such as in.routed(1M) and in.ripngd(1M).

35  route supports a limited number of general options, but a
36  rich command language. Users can specify an arbitrary
37  request that can be delivered by means of the programmatic
38  interface discussed in route(7P).

41  route uses a routing socket and the new message types
42  RTM_ADD, RTM_DELETE, RTM_GET, and RTM_CHANGE. While only
43  superusers can modify routing tables, the RTM_GET operation
44  is allowed for non-privileged users.

46 OPTIONS
47   -f          Flush the routing tables of all gateway
48              entries. If you use the -f option in conjunc-
49              tion with any of the route sub-commands,
50              route flushes the gateways before performing
51              the sub-command. Specify the table to flush
52              by placing the -inet or -inet6 modifier
53              immediately after the -f option. If unspeci-
54              fied, flushing IPv4 (-inet) routes is the
55              default.

```

```

67 System Administration Commands          route(1M)

71   -n          Prevent attempts to print host and network
72              names symbolically when reporting actions.
73              This option is useful when name servers are
74              unavailable.

77   -p          Make changes to the network route tables per-
78              sistent across system restarts. The operation
79              is applied to the network routing tables
80              first and, if successful, is then applied to
81              the list of saved routes used at system
82              startup. In determining whether an operation
83              was successful, a failure to add a route that
84              already exists or to delete a route that is
85              not in the routing table is ignored. Particu-
86              lar care should be taken when using host or
87              network names in persistent routes, as
88              network-based name resolution services are
89              not available at the time routes are added at
90              startup.

93   -q          Suppress all output.

96   -R root-dir Specify an alternate root directory where
97              route applies changes. This option is ignored
98              unless used in conjunction with the -p
99              option. When -R is specified, route changes
100             are applied only to the list of saved routes
101             to be used at startup, not to the network
102             routing tables. In addition, certain checks,
103             such as the existence of network interfaces
104             such as the existence of network interfaces
105             used with -ifp, are skipped. This can be use-
106             ful from within JumpStart scripts, where the
107             root directory of the system being modified
             is in a location other than /.

110  -v          Print additional details in verbose mode.

113  Subcommands
114    The following subcommands are supported:

116  add          Add a route.

119  change       Change aspects of a route (such as its gateway).

```

133 System Administration Commands route(1M)

137 delete Delete a specific route.

140 flush Remove all gateway entries from the routing
141 table.

144 get Look up and display the route for a destination.

147 monitor Continuously report any changes to the routing
148 information base, routing lookup misses, or
149 suspected network partitionings.

152 show Display the list of routes to be applied at sys-
153 tem startup. Can be used only in conjunction with
154 the -p option.

158 The add and delete sub-commands have the following syntax:

160 route [-fnvq] cmd destination gateway [metric/netmask]

165 where cmd is add or delete, destination is the destination
166 host or network, and gateway is the next-hop intermediary
167 through which packets should be routed. Modifiers described
168 in OPERANDS can be placed anywhere on the command line.

171 The get and change sub-commands have the following syntax:

173 route [-fnvq] cmd destination [gateway [metric/netmask]]

178 where cmd is get or change, destination is the destination
179 host or network, and gateway is the next-hop intermediary
180 through which packets should be routed. Modifiers described
181 in OPERANDS can be placed anywhere on the command line.

184 The monitor sub-command has the following syntax:

186 route monitor [-inet | -inet6]

199 System Administration Commands route(1M)

203 OPERANDS
204 route executes its sub-commands on routes to destinations by
205 way of gateways.

207 Destinations and Gateways
208 By default, destination and gateway addresses are inter-
209 preted as IPv4 addresses. All symbolic names are tried first
210 as a host name, using getipnodebyname(3SOCKET). If this
211 lookup fails in the AF_INET case, getnetbyname(3SOCKET)
212 interprets the name as that of a network.

215 Including an optional modifier on the command line before
216 the address changes how the route sub-command interprets it.

219 The following modifiers are supported:

221 -inet Force the address to be interpreted as an IPv4
222 address, that is, under the AF_INET address fam-
223 ily.

226 -inet6 Force the address to be interpreted as an IPv6
227 address, that is, under the AF_INET6 address fam-
228 ily.

232 For IPv4 addresses, routes to a particular host are by
233 default distinguished from those to a network by interpret-
234 ing the Internet address specified as the destination. If
235 the destination has a local address part (that is, the por-
236 tion not covered by the netmask) of 0, or if the destination
237 is resolved as the symbolic name of a network, then the
238 route is assumed to be to a network; otherwise, it is
239 presumed to be a route to a host.

242 You can force this selection by using one of the following
243 modifiers:

245 -host Force the destination to be interpreted as a host.

248 -net Force the destination to be interpreted as a net-
249 work.

253 For example:

265 System Administration Commands route(1M)

269	Destination	Destination Equivalent
270		
271	128.32	-host 128.0.0.32
272	128.32.130	-host 128.32.0.130
273	-net 128.32	128.32.0.0
274	-net 128.32.130	128.32.130.0

278 Two modifiers avoid confusion between addresses and keywords
 279 (for example., host used as a symbolic host name). You can
 280 distinguish a destination by preceding it with the `-dst`
 281 modifier. You can distinguish a gateway address by using the
 282 `-gateway` modifier. If the destination is directly reachable
 283 by way of an interface requiring no intermediary IP router
 284 to act as a gateway, this can be indicated by using the
 285 `-interface` or `-iface` modifier.

288 In the following example, the route does not refer to an
 289 external gateway (router), but rather to one of the
 290 machine's interfaces. Packets with IP destination addresses
 291 matching the destination and mask on such a route are sent
 292 out on the interface identified by the gateway address. For
 293 interfaces using the ARP protocol, this type of route is
 294 used to specify that all matching destinations are local to
 295 the physical link. That is, a host could be configured to
 296 ARP for all addresses, without regard to the configured
 297 interface netmask, by adding a default route using this com-
 298 mand. For example:

```
300 example# route add default hostname -interface
```

305 where gateway address hostname is the name or IP address
 306 associated with the network interface over which all match-
 307 ing packets should be sent. On a host with a single network
 308 interface, hostname is usually the same as the nodename
 309 returned by the `uname -n` command. See `uname(1)`.

312 For backward compatibility with older systems, directly
 313 reachable routes can also be specified by placing a 0 after
 314 the gateway address:

```
316 example# route add default hostname 0
```

331 System Administration Commands route(1M)

335 This value was once a route metric, but this metric is no
 336 longer used. If the value is specified as 0, then the desti-
 337 nation is directly reachable (equivalent to specifying
 338 `-interface`). If it is non-zero but cannot be interpreted as
 339 a subnet mask, then a gateway is used (default).

342 With the `AF_INET` address family or an IPv4 address, a
 343 separate subnet mask can be specified. This can be specified
 344 in one of the following ways:

- 346 o IP address following the gateway address . This is
- 347 typically specified in decimal dot notation as for
- 348 **inet_addr(3SOCKET) rather than in symbolic form.**
- 348 *inet_addr(3SOCKET) rather than in symbolic form.*

- 350 o IP address following the `-netmask` qualifier.

- 352 o Slash character and a decimal length appended to
- 353 the destination address.

356 If a subnet mask is not specified, the mask used is the sub-
 357 net mask of the output interface selected by the gateway
 358 address, if the classful network of the destination is the
 359 same as the classful network of the interface. Otherwise,
 360 the classful network mask for the destination address is
 361 used.

364 Each of the following examples creates an IPv4 route to the
 365 destination 192.0.2.32 subnet with a subnet mask of
 366 255.255.255.224:

```
368 example# route add 192.0.2.32/27 somegateway
369 example# route add 192.0.2.32 -netmask 255.255.255.224 somegateway
370 example# route add 192.0.2.32 somegateway 255.255.255.224
```

375 For IPv6, only the slash format is accepted. The following
 376 example creates an IPv6 route to the destination 3ffe:: with
 377 a netmask of 16 one-bits followed by 112 zero-bits.

```
379 example# route add -inet6 3ffe::/16 somegateway
```

384 In cases where the gateway does not uniquely identify the
 385 output interface (for example, when several interfaces have
 386 the same address), you can use the `-ifp ifname` modifier to

390 SunOS 5.11 Last change: 26 Jun 2006 6

397 System Administration Commands route(1M)

401 specify the interface by name. For example, **-ifp lo0**
 402 associates the route with the **lo0** interface. If the named
 403 interface is an underlying interface in an IPMP group, then
 404 requests to add a route will automatically be translated to
 405 the corresponding IPMP IP interface, and requests to delete
 406 or change a route on an underlying interface will fail.
 401 specify the interface by name. For example, *-ifp lo0* associ-
 402 ates the route with the *lo0* interface.

409 Routing Flags
 410 Routes have associated flags that influence operation of the
 411 protocols when sending to destinations matched by the
 412 routes. These flags can be set (and in some cases cleared,
 413 indicated by ~) by including the following modifiers on the
 414 command line:

Modifier	Flag	Description
-interface	~RTF_GATEWAY	Destination is directly reachable
-iface	~RTF_GATEWAY	Alias for interface modifier
-static	RTF_STATIC	Manually added route
-nostatic	~RTF_STATIC	Pretend route was added by kernel or routing daemon
-reject	RTF_REJECT	Emit an ICMP unreachable when matched
-blackhole	RTF_BLACKHOLE	Silently discard packets during updates
-proto1	RTF_PROTO1	Set protocol specific routing flag #1
-proto2	RTF_PROTO2	Set protocol specific routing flag #2
-private	RTF_PRIVATE	Do not advertise this route
-multirt	RTF_MULTIRT	Creates the specified redundant route
-setsrcc	RTF_SETSRCC	Assigns the default source address

435 The optional modifiers *-rtt*, *-rttvar*, *-sendpipe*, *-recvpipe*,
 436 *-mtu*, *-hopcount*, *-expire*, and *-ssthresh* provide initial
 437 values to quantities maintained in the routing entry by
 438 transport level protocols, such as TCP. These can be indivi-
 439 dually locked either by preceding each modifier to be locked
 440 by the *-lock* meta-modifier, or by specifying that all ensu-
 441 ing metrics can be locked by the *-lockrest* meta-modifier.

444 Some transport layer protocols can support only some of
 445 these metrics. The following optional modifiers are sup-
 446 ported:

448 *-expire* Lifetime for the entry. This optional modifier
 449 is not currently supported.

452 *-hopcount* Maximum hop count. This optional modifier is
 453 not currently supported.

456 *-mtu* Maximum MTU in bytes.

461 SunOS 5.11 Last change: 26 Jun 2006 7

468 System Administration Commands route(1M)

472 *-recvpipe* Receive pipe size in bytes.

475 *-rtt* Round trip time in microseconds.

478 *-rttvar* Round trip time variance in microseconds.

481 *-sendpipe* Send pipe size in bytes.

484 *-ssthresh* Send pipe size threshold in bytes.

487 *-secattr* Security attributes of the route. This modifier
 488 is available only if the system is configured
 489 with the Solaris Trusted Extensions feature.

491 The *-secattr* modifier has the following format:

493 min_sl=val,max_sl=val,doi=val,cipso

495 or:

497 sl=VAL,doi=VAL,cipso

499 In the first form, above, the *val* for *min_sl*
 500 and *max_sl* is a sensitivity label in either hex
 501 or string form. The *val* for *doi* is a non-
 502 negative integer. The route will apply only for
 503 packets with the same domain of interpretation
 504 as defined by the *doi* value and within the
 505 accreditation range defined by the *min_sl* and
 506 *max_sl* values. The *cipso* keyword is optional
 507 and set by default. Valid *min_sl*, *max_sl* and
 508 *doi* keyword/value pairs are mandatory. Note
 509 that if *val* contains a space, it must be pro-
 510 tected by double quotes.

512 The second form, above, is equivalent to speci-
 513 fying the first form with the same *VAL* for
 514 *min_sl* and *max_sl*. The second form should be
 515 used for the *get* command, because *get* uses only
 516 a single sensitivity label.

519 Compatibility

520 The modifiers host and net are taken to be equivalent to
 521 -host and -net. To specify a symbolic address that matches
 522 one of these names, use the dst or gateway keyword to dis-
 523 tinguish it. For example: -dst host

527 SunOS 5.11 Last change: 26 Jun 2006 8

534 System Administration Commands route(1M)

538 The following two flags are also accepted for compatibility
 539 with older systems, but have no effect.

543	Modifier	Flag
544		
545	-cloning	RTF_CLONING
546	-xresolve	RTF_XRESOLVE

550 The -ifa hostname modifier is also accepted, but has no
 551 effect.

553 FILES
 554 /etc/defaultrouter List of default routers
 557 /etc/hosts List of host names and net addresses
 560 /etc/networks List of network names and addresses

563 ATTRIBUTES
 564 See attributes(5) for descriptions of the following attri-
 565 butes:

569	ATTRIBUTE TYPE	ATTRIBUTE VALUE
570		
571	Availability	SUNWcsu
572		
573		

576 SEE ALSO
 577 uname(1), in.ripngd(1M), in.routed(1M), netstat(1M),
 578 routed(1M), ioctl(2), getipnodebyname(3SOCKET),
 579 getnetbyname(3SOCKET), inet_addr(3SOCKET), defaultrouter(4),
 580 hosts(4), networks(4), attributes(5), ARP(7P), ip(7P),
 581 route(7P), routing(7P)

583 DIAGNOSTICS
 584 add [host| network] destination:gateway flags

586 The specified route is being added to the tables. The
 587 values printed are from the routing table entry supplied
 588 in the ioctl(2) call. If the gateway address used was
 589 not the primary address of the gateway (the first one

593 SunOS 5.11 Last change: 26 Jun 2006 9

600 System Administration Commands route(1M)

604 returned by getipnodebyname(3SOCKET)) the gateway
 605 address is printed numerically as well as symbolically.

608 delete [host| network] destination:gateway flags
 609 change [host| network] destination:gateway flags

611 As add, but when deleting or changing an entry.

614 destination done

616 When the -f flag is specified, or the flush sub-command
 617 is used, each routing table entry deleted is indicated
 618 with a message of this form.

621 Network is unreachable

623 An attempt to add a route failed because the gateway
 624 listed was not on a directly-connected network. Give the
 625 next-hop gateway instead.

628 not in table

630 A delete operation was attempted for an entry that is
 631 not in the table.

634 entry exists

636 An add operation was attempted for a route that already
 637 exists in the kernel.

640 routing table overflow

642 An operation was attempted, but the system was unable to
 643 allocate memory to create the new entry.

646 insufficient privileges

648 An attempt to add, delete, change, or flush a route
 649 failed because the calling process does not have
 650 appropriate privileges.

new/./route.lm.txt

11

653 NOTES

654 Specifying that destinations are local (with the -inter-
655 facemodifier) assumes that the routers implement proxy ARP,

659 SunOS 5.11 Last change: 26 Jun 2006 10

666 System Administration Commands route(1M)

670 meaning that they respond to ARP queries for all reachable
671 destinations. Normally, using either router discovery or RIP
672 is more reliable and scalable than using proxy ARP. See
673 in.routed(1M) for information related to RIP.

676 Combining the all destinations are local route with subnet
677 or network routes can lead to unpredictable results. The
678 search order as it relates to the all destinations are local
679 route are undefined and can vary from release to release.

new/./route.lm.txt

12

725 SunOS 5.11 Last change: 26 Jun 2006 11

```

*****
14100 Fri Dec 19 04:18:58 2008
new/./route.7p.txt
Clearview IPMP manpages
*****

```

1 Protocols route(7P)

```

5 NAME
6 route - kernel packet forwarding database

8 SYNOPSIS
9 #include <sys/types.h>
10 #include <sys/socket.h>
11 #include <net/if.h>
12 #include <net/route.h>

14 int socket(PF_ROUTE, SOCK_RAW, int protocol);

17 DESCRIPTION
18 UNIX provides some packet routing facilities. The kernel
19 maintains a routing information database, which is used in
20 selecting the appropriate network interface when transmit-
21 ting packets.

24 A user process (or possibly multiple co-operating processes)
25 maintains this database by sending messages over a special
26 kind of socket. This supplants fixed size ioctl(2)'s speci-
27 fied in routing(7P). Routing table changes may only be car-
28 ried out by the superuser.

31 The operating system may spontaneously emit routing messages
32 in response to external events, such as receipt of a re-
33 direct, or failure to locate a suitable route for a request.
34 The message types are described in greater detail below.

37 Routing database entries come in two flavors: entries for a
38 specific host, or entries for all hosts on a generic subnet-
39 work (as specified by a bit mask and value under the mask).
40 The effect of wildcard or default route may be achieved by
41 using a mask of all zeros, and there may be hierarchical
42 routes.

45 When the system is booted and addresses are assigned to the
46 network interfaces, the internet protocol family installs a
47 routing table entry for each interface when it is ready for
48 traffic. Normally the protocol specifies the route through
49 each interface as a direct connection to the destination
50 host or network. If the route is direct, the transport
51 layer of a protocol family usually requests the packet be
52 sent to the same host specified in the packet. Otherwise,
53 the interface is requested to address the packet to the
54 gateway listed in the routing entry, that is, the packet is
55 forwarded.

```

67 Protocols route(7P)

```

71 When routing a packet, the kernel attempts to find the most
72 specific route matching the destination. If no entry is
73 found, the destination is declared to be unreachable, and a
74 routing-miss message is generated if there are any listeners
75 on the routing control socket (described below). If there
76 are two different mask and value-under-the-mask pairs that
77 match, the more specific is the one with more bits in the
78 mask. A route to a host is regarded as being supplied with a
79 mask of as many ones as there are bits in the destination.

82 A wildcard routing entry is specified with a zero destina-
83 tion address value, and a mask of all zeroes. Wildcard
84 routes are used when the system fails to find other routes
85 matching the destination. The combination of wildcard routes
86 and routing redirects can provide an economical mechanism
87 for routing traffic.

90 One opens the channel for passing routing control messages
91 by using the socket call shown in the section above. There
92 can be more than one routing socket open per system.

95 Messages are formed by a header followed by a small number
96 of sockaddr, whose length depend on the address family.
97 sockaddrs are interpreted by position. An example of a type
98 of message with three addresses might be a CIDR prefix
99 route: Destination, Netmask, and Gateway. The interpretation
100 of which addresses are present is given by a bit mask within
101 the header, and the sequence is least significant to most
102 significant bit within the vector.

105 Any messages sent to the kernel are returned, and copies are
106 sent to all interested listeners. The kernel provides the
107 process ID of the sender, and the sender may use an addi-
108 tional sequence field to distinguish between outstanding
109 messages. However, message replies may be lost when kernel
110 buffers are exhausted.

113 The protocol parameter specifies which messages an applica-
114 tion listening on the routing socket is interested in see-
115 ing, based on the the address family of the sockaddr
116 present. Currently, you can specify AF_INET and AF_INET6 to
117 filter the messages seen by the listener, or alternatively,
118 you can specify AF_UNSPEC to indicate that the listener is
119 interested in all routing messages.

```

133 Protocols route(7P)

137 The kernel may reject certain messages, and will indicate
138 this by filling in the rtm_errno field of the rt_msghdr
139 struct (see below). The following codes may be returned:

141 EEXIST If requested to duplicate an existing entry

144 ESRCH If requested to delete a non-existent entry

147 ENOBUFS If insufficient resources were available to
148 install a new route.

151 EPERM If the calling process does not have appropriate
152 privileges to alter the routing table.

156 In the current implementation, all routing processes run
157 locally, and the values for rtm_errno are available through
158 the normal errno mechanism, even if the routing reply mes-
159 sage is lost.

162 A process may avoid the expense of reading replies to its
163 own messages by issuing a setsockopt(3SOCKET) call indicat-
164 ing that the SO_USELOOPBACK option at the SOL_SOCKET level
165 is to be turned off. A process may ignore all messages from
166 the routing socket by doing a shutdown(3SOCKET) system call
167 for further input.

170 By default, underlying IP interfaces in an IPMP group are
171 not visible to routing sockets. As such, routing sockets
172 will not receive events related to underlying IP interfaces
173 in an IPMP group. For consistency, when an IP interface is
174 placed into an IPMP group, *RTM_DELADDR* messages will be
175 generated for each of the *IFF_UP* addresses that are not
176 migrated to the corresponding IPMP IP interface, and an
177 *RTM_IFINFO* message will be sent indicating that the
178 interface is now down. Similarly, when an underlying
179 interface is removed from an IPMP group, an *RTM_IFINFO*
180 message will be sent indicating that the interface is again
181 up, and *RTM_NEWADDR* messages will be generated for each
182 *IFF_UP* address found on the interface.

185 The *RT_AWARE* socket option at the *SOL_ROUTE* level allows
186 an application to indicate its awareness of certain
187 features, which controls routing socket behavior. The
188 supported values are:

190 *RTAW_DEFAULT* Default awareness.

192 *RTAW_UNDER_IPMP* IPMP underlying interface awareness.
193 When this is enabled, underlying IP

194 interfaces in an IPMP group will remain
195 visible to the routing socket and
196 events related to them will continue to
197 be generated.

200 An *RTM_ADD* request tied to an underlying IP interface in
201 an IPMP group will be translated to an *RTM_ADD* request for
202 its corresponding IPMP IP interface. All routing socket
203 requests other than *RTM_ADD* and *RTM_GET* will fail when
204 issued on an underlying IP interface in an IPMP group.

207 If a route is in use when it is deleted, the routing entry
208 is marked down and removed from the routing table, but the
209 resources associated with it are not reclaimed until all
210 references to it are released.

213 The RTM_IFINFO, RTM_NEWADDR, and RTM_ADD messages associated
214 with interface configuration (setting the IFF_UP bit) are
215 normally delayed until after Duplicate Address Detection
216 completes. Thus, applications that configure interfaces and
217 wish to wait until the interface is ready can wait until
218 RTM_IFINFO is returned and SIOCGLIFFLAGS shows that
219 IFF_DUPLICATE is not set.

221 Messages
222 User processes can obtain information about the routing
223 entry to a specific destination by using a RTM_GET message.

229 SunOS 5.11 Last change: 25 July 2006 3

236 Protocols route(7P)

240 Messages include:

242 #define RTM_ADD 0x1 /* Add Route */
243 #define RTM_DELETE 0x2 /* Delete Route */
244 #define RTM_CHANGE 0x3 /* Change Metrics, Flags, or Gateway */
245 #define RTM_GET 0x4 /* Report Information */
246 #define RTM_LOSING 0x5 /* Kernel Suspects Partitioning */
247 #define RTM_REDIRECT 0x6 /* Told to use different route */
248 #define RTM_MISS 0x7 /* Lookup failed on this address */
249 #define RTM_LOCK 0x8 /* fix specified metrics */
250 #define RTM_OLDADD 0x9 /* caused by SIOCADDRT */
251 #define RTM_OLDDEL 0xa /* caused by SIOCDELRT */
252 #define RTM_RESOLVE 0xb /* request to resolve dst to LL addr */
253 #define RTM_NEWADDR 0xc /* address being added to iface */
254 #define RTM_DELADDR 0xd /* address being removed from iface */
255 #define RTM_IFINFO 0xe /* iface going up/down etc. */

259 A message header consists of:

```

261 struct rt_msghdr {
262     ushort_t rtm_msglen; /* to skip over non-understood messages */
263     uchar_t rtm_version; /* future binary compatibility */
264     uchar_t rtm_type; /* message type */
265     ushort_t rtm_index; /* index for associated ifp */
266     pid_t rtm_pid; /* identify sender */
267     int rtm_addrs; /* bitmask identifying sockaddrs in msg */
268     int rtm_seq; /* for sender to identify action */
269     int rtm_errno; /* why failed */
270     int rtm_flags; /* flags, incl kern & message, e.g., DONE */
271     int rtm_use; /* from rtentry */
272     uint_t rtm_inits; /* which values we are initializing */

274 struct rt_metrics rtm_rmx; /* metrics themselves */
275     };

279 where

281 struct rt_metrics {
282     uint32_t rmx_locks; /* Kernel must leave these values alone */
283     uint32_t rmx_mtu; /* MTU for this path */
284     uint32_t rmx_hopcount; /* max hops expected */
285     uint32_t rmx_expire; /* lifetime for route, e.g., redirect */
286     uint32_t rmx_recvpipe; /* inbound delay-bandwidth product */
287     uint32_t rmx_sendpipe; /* outbound delay-bandwidth product */
288     uint32_t rmx_ssthresh; /* outbound gateway buffer limit */
289     uint32_t rmx_rtt; /* estimated round trip time */
290     uint32_t rmx_rttvar; /* estimated rtt variance */
291     uint32_t rmx_pksent; /* packets sent using this route */

```

295 SunOS 5.11 Last change: 25 July 2006

4

302 Protocols route(7P)

```

306     };

308     /* Flags include the values */

311 #define RTF_UP 0x1 /* route usable */
312 #define RTF_GATEWAY 0x2 /* destination is a gateway */
313 #define RTF_HOST 0x4 /* host entry (net otherwise) */
314 #define RTF_REJECT 0x8 /* host or net unreachable */
315 #define RTF_DYNAMIC 0x10 /* created dynamically (by redirect) */
316 #define RTF_MODIFIED 0x20 /* modified dynamically (by redirect) */
317 #define RTF_DONE 0x40 /* message confirmed */
318 #define RTF_MASK 0x80 /* subnet mask present */
319 #define RTF_CLONING 0x100 /* generate new routes on use */
320 #define RTF_XRESOLVE 0x200 /* external daemon resolves name */
321 #define RTF_LLINFO 0x400 /* generated by ARP */
322 #define RTF_STATIC 0x800 /* manually added */
323 #define RTF_BLACKHOLE 0x1000 /* just discard pkts (during updates) */
324 #define RTF_PRIVATE 0x2000 /* do not advertise this route */
325 #define RTF_PROTO2 0x4000 /* protocol specific routing flag #2 */

```

```

326 #define RTF_PROTO1 0x8000 /* protocol specific routing flag #1 */

328 /* Specifiers for metric values in rmx_locks and rtm_inits are */

330 #define RTV_MTU 0x1 /* init or lock _mtu */
331 #define RTV_HOPCOUNT 0x2 /* init or lock _hopcount */
332 #define RTV_EXPIRE 0x4 /* init or lock _expire */
333 #define RTV_RPIPE 0x8 /* init or lock _recvpipe */
334 #define RTV_SPIPE 0x10 /* init or lock _sendpipe */
335 #define RTV_SSTHRESH 0x20 /* init or lock _ssthresh */
336 #define RTV_RTT 0x40 /* init or lock _rtt */
337 #define RTV_RTTVAR 0x80 /* init or lock _rttvar */

339 /* Specifiers for which addresses are present in the messages are */

341 #define RTA_DST 0x1 /* destination sockaddr present */
342 #define RTA_GATEWAY 0x2 /* gateway sockaddr present */
343 #define RTA_NETMASK 0x4 /* netmask sockaddr present */
344 #define RTA_GENMASK 0x8 /* cloning mask sockaddr present */
345 #define RTA_IFP 0x10 /* interface name sockaddr present */
346 #define RTA_IFA 0x20 /* interface addr sockaddr present */
347 #define RTA_AUTHOR 0x40 /* sockaddr for author of redirect */
348 #define RTA_BRD 0x80 /* for NEWADDR, broadcast or p-p dest addr

351 SEE ALSO
352     ioctl(2), setsockopt(3SOCKET), shutdown(3SOCKET),
353     routing(7P)

355 NOTES
356     Some of the metrics may not be implemented and return zero.
357     The implemented metrics are set in rtm_inits.

```

361 SunOS 5.11

Last change: 25 July 2006

5

```

*****
12781 Fri Dec 19 04:18:59 2008
new/./socket.h.txt
Clearview IPMP manpages
*****
 1 Headers                                socket.h(3HEAD)

5 NAME
6   socket.h, socket - Internet Protocol family

8 SYNOPSIS
9   #include <sys/socket.h>

12 DESCRIPTION
13   The <sys/socket.h> header defines the unsigned integral type
14   sa_family_t through typedef.

17   The <sys/socket.h> header defines the sockaddr structure
18   that includes the following members:

20   sa_family_t  sa_family    /* address family */
21   char         sa_data[]    /* socket  address (variable-length
22                          data) */

26 libxnet Interfaces
27   The <sys/socket.h> header defines the msg_hdr structure for
28   libxnet interfaces that includes the following members:

30   void         *msg_name    /* optional address */
31   socklen_t    msg_namelen  /* size of address */
32   struct iovec *msg_iov     /* scatter/gather array */
33   int          msg_iovlen   /* members in msg_iov */
34   void         *msg_control  /* ancillary data, see below */
35   socklen_t    msg_controllen /* ancillary data buffer len */
36   int          msg_flags    /* flags on received message */

40   The <sys/socket.h> header defines the cmsghdr structure for
41   libxnet that includes the following members:

43   socklen_t    cmsg_len     /* data byte count, including hdr */
44   int          cmsg_level   /* originating protocol */
45   int          cmsg_type    /* protocol-specific type */

49   Ancillary data consists of a sequence of pairs, each con-
50   sisting of a cmsghdr structure followed by a data array. The
51   data array contains the ancillary data message, and the
52   cmsghdr structure contains descriptive information that
53   allows an application to correctly parse the data.

60 SunOS 5.11                Last change: 03 Aug 2006                1

```

```

67 Headers                                socket.h(3HEAD)

71   The values for cmsg_level will be legal values for the level
72   argument to the getsockopt() and setsockopt() functions. The
73   SCM_RIGHTS type is supported for level SOL_SOCKET.

76   Ancillary data is also possible at the socket level. The
77   <sys/socket.h> header defines the following macros for use
78   as the cmsg_type values when cmsg_level is SOL_SOCKET.

80   SCM_RIGHTS    Indicates that the data array contains the
81                  access rights to be sent or received.

84   SCM_UCRED    Indicates that the data array contains a
85                  ucred_t to be received. The ucred_t is the
86                  credential of the sending process at the time
87                  the message was sent. This is a Sun-specific,
88                  Evolving interface. See ucred_get(3C).

92   The IPv4 data formats generally use the same values for data
93   passed back in cmsghdr as for setsockopt() to enable the
94   feature. The IPv4 data formats are listed below with the
95   associated payload for each.
95   associated payload for each.

97   IPPROTO_IP
98   IP_RECVDSTADDR

100   ipaddr_t, IP address

103   IPPROTO_IP
104   IP_RECVOPTS

106   variable-length IP options, up to 40 bytes

109   IPPROTO_IP
110   IP_RECVIF

112   uint_t, ifIndex number

115   IPPROTO_IP
116   IP_RECVSLLA

118   struct sockaddr_dl, link layer address

121   IPPROTO_IP
122   IP_RECVTTL

126 SunOS 5.11                Last change: 03 Aug 2006                2

```

133 Headers socket.h(3HEAD)

137 uint8_t

140 SOL_SOCKET
141 SO_RECVUCRED

143 ucred_t - cmsghdr.cmsg_type is SCM_UCRED, not
144 SO_RECVUCRED

148 The IPv6 data formats use different values for enabling the
149 option and for passing the value back to the application.
150 **The IPv6 data formats are listed below with the associated**
151 *The IPv6 data formats are listed below with the associated*
151 payload for each.

153 IPPROTO_IPV6
154 IPV6_RECVPKTINFO

156 in_pktinfo, cmsg_type IPV6_PKTINFO

159 IPPROTO_IPV6
160 IPV6_RECVTCLASS

162 uint_t, cmsg_type IPV6_TCLASS

165 IPPROTO_IPV6
166 IPV6_RECVPATHMTU

168 ip6_mtuinfo, cmsg_type IPV6_PATHMTU

171 IPPROTO_IPV6
172 IPV6_RECVHOPLIMIT

174 uint_t, cmsg_type IPV6_HOPLIMIT

177 IPPROTO_IPV6
178 IPV6_RECVHOPOPTS

180 variable-length IPv6 options, cmsg_type IPV6_HOPOPTS

183 IPPROTO_IPV6
184 IPV6_RECVDSTOPTS

186 variable-length IPv6 options, cmsg_type IPV6_DSTOPTS

192 SunOS 5.11

Last change: 03 Aug 2006

3

199 Headers socket.h(3HEAD)

203 IPPROTO_IPV6
204 IPV6_RECVRTHDR

206 variable-length IPv6 options, cmsg_type IPV6_RTHDR

209 IPPROTO_IPV6
210 IPV6_RECVRTHDRDSTOPTS

212 variable-length IPv6 options, cmsg_type IPV6_DSTOPTS

216 The <sys/socket.h> header defines the following macros to
217 gain access to the data arrays in the ancillary data associ-
218 ated with a message header:

220 MSG_DATA(cmsg)

222 If the argument is a pointer to a cmsghdr structure,
223 this macro returns an unsigned character pointer to the
224 data array associated with the cmsghdr structure.

227 MSG_NXTHDR(mhdr, cmsg)

229 If the first argument is a pointer to a msghdr structure
230 and the second argument is a pointer to a cmsghdr struc-
231 ture in the ancillary data, pointed to by the
232 msg_control field of that msghdr structure, this macro
233 returns a pointer to the next cmsghdr structure, or a
234 null pointer if this structure is the last cmsghdr in
235 the ancillary data.

238 MSG_FIRSTHDR(mhdr)

240 If the argument is a pointer to a msghdr structure, this
241 macro returns a pointer to the first cmsghdr structure
242 in the ancillary data associated with this msghdr struc-
243 ture, or a null pointer if there is no ancillary data
244 associated with the msghdr structure.

247 MSG_SPACE(len)

249 Given the length of an ancillary data object,
250 MSG_SPACE() returns the space required by the object
251 and its cmsghdr structure, including any padding needed
252 to satisfy alignment requirements. This macro can be
253 used, for example, to allocate space dynamically for the
254 ancillary data. This macro should not be used to

258 SunOS 5.11 Last change: 03 Aug 2006 4

265 Headers socket.h(3HEAD)

269 initialize the `cmsg_len` member of a `cmsghdr` structure.
270 Use the `CMSG_LEN()` macro instead.

273 `CMSG_LEN(len)`

275 Given the length of an ancillary data object, `CMSG_LEN()`
276 returns the value to store in the `cmsg_len` member of the
277 `cmsghdr` structure, taking into account any padding
278 needed to satisfy alignment requirements.

282 The `<sys/socket.h>` header defines the `linger` structure that
283 includes the following members:

```
285 int l_onoff /* indicates whether linger option is enabled */
286 int l_linger /* linger time, in seconds */
```

290 The `<sys/socket.h>` header defines the following macros:

292 `SOCK_DGRAM` Datagram socket

295 `SOCK_STREAM` Byte-stream socket

298 `SOCK_SEQPACKET` Sequenced-packet socket

302 **The `<sys/socket.h>` header defines the following macros for
303 use as the `*level*` argument of `*setsockopt()*` and
304 `getsockopt()*`.**

302 *The `<sys/socket.h>` header defines the following macro for
303 use as the `level` argument of `setsockopt()` and `getsockopt()`.*

306 **`*SOCKET*` Options to be accessed at the socket level,
307 not the protocol level.**

305 `SOCKET` Options to be accessed at socket level, not
306 protocol level.

309 **`*SO_ROUTE*` Options to be accessed at the routing socket
310 level, not the protocol level.**

313 The `<sys/socket.h>` header defines the following macros for
314 use as the `option_name` argument in `getsockopt()` or `set-`
315 `sockopt()` calls:

317 `SO_DEBUG` Debugging information is being recorded.

320 `SO_ACCEPTCONN` Socket is accepting connections.

327 SunOS 5.11 Last change: 03 Aug 2006 5

334 Headers socket.h(3HEAD)

338 `SO_BROADCAST` Transmission of broadcast messages is sup-
339 ported.

342 `SO_REUSEADDR` Reuse of local addresses is supported.

345 `SO_KEEPAIVE` Connections are kept alive with periodic
346 messages.

349 `SO_LINGER` Socket lingers on close.

352 `SO_OOBINLINE` Out-of-band data is transmitted in line.

355 `SO_SNDBUF` Send buffer size.

358 `SO_RCVBUF` Receive buffer size.

361 `SO_ERROR` Socket error status.

364 `SO_TYPE` Socket type.

367 `SO_RECVUCRED` Request the reception of user credential
368 ancillary data. This is a Sun-specific,
369 Evolving interface. See `ucrd_get(3C)`.

372 `SO_MAC_EXEMPT` Mandatory Access Control (MAC) exemption
373 for unlabeled peers. This option is avail-
374 able only if the system is configured with
375 Trusted Extensions.

378 `SO_ALLZONES` Bypass zone boundaries (privileged).

382 The `<sys/socket.h>` header defines the following macros for
383 use as the valid values for the `msg_flags` field in the
384 `msghdr` structure, or the `flags` parameter in `recvfrom()`,
385 `recvmsg()`, `sendto()`, or `sendmsg()` calls:

```

387     MSG_CTRUNC     Control data truncated.

393 SunOS 5.11         Last change: 03 Aug 2006           6

400 Headers                                     socket.h(3HEAD)

404     MSG_EOR        Terminates a record (if supported by the pro-
405                    tocol).

408     MSG_OOB        Out-of-band data.

411     MSG_PEEK       Leave received data in queue.

414     MSG_TRUNC      Normal data truncated.

417     MSG_WAITALL    Wait for complete message.

421     The <sys/socket.h> header defines the following macros:

423     AF_UNIX         UNIX domain sockets

426     AF_INET        Internet domain sockets

430     The <sys/socket.h> header defines the following macros:

432     SHUT_RD         Disables further receive operations.

435     SHUT_WR         Disables further send operations.

438     SHUT_RDWR      Disables further send and receive operations.

441     libsocket Interfaces
442     The <sys/socket.h> header defines the msg_hdr structure for
443     libsocket interfaces that includes the following members:

445     void            *msg_name        /* optional address */
446     socklen_t        msg_namelen     /* size of address */
447     struct iovec     *msg_iov        /* scatter/gather array */
448     int              msg_iovlen     /* # elements in msg_iov */
449     caddr_t          msg_accrights    /* access rights sent/received */

```

```

453     The msg_name and msg_namelen parameters specify the destina-
454     tion address when the socket is unconnected The msg_name can
455     be specified as a NULL pointer if no names are desired or

459 SunOS 5.11         Last change: 03 Aug 2006           7

466 Headers                                     socket.h(3HEAD)

470     required. The msg_iov and msg_iovlen parameters describe
471     the scatter-gather locations, as described in read(2). The
472     msg_accrights parameter specifies the buffer in which access
473     rights sent along with the message are received. The
474     msg_accrightslen specifies the length of the buffer.

476     ATTRIBUTES
477     See attributes(5) for descriptions of the following attri-
478     butes:

482
483
484
485
486

```

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard

```

489     SEE ALSO
490     accept(3SOCKET), accept(3XNET), bind(3SOCKET), bind(3XNET),
491     connect(3SOCKET), connect(3XNET), getpeername(3SOCKET),
492     getpeername(3XNET), getpeerucred(3C), getsockname(3SOCKET),
493     getsockname(3XNET), getsockopt(3SOCKET), getsockopt(3XNET),
494     libsocket(3LIB), listen(3SOCKET), listen(3XNET),
495     recv(3SOCKET), recv(3XNET), recvfrom(3SOCKET),
496     recvfrom(3XNET), recvmsg(3SOCKET), recvmsg(3XNET),
497     send(3SOCKET), send(3XNET), sendmsg(3SOCKET),
498     sendmsg(3XNET), sendto(3SOCKET), sendto(3XNET),
499     setsockopt(3SOCKET), setsockopt(3XNET), shutdown(3SOCKET),
500     shutdown(3XNET), socket(3SOCKET), socket(3XNET),
501     socketpair(3SOCKET), socketpair(3XNET),
502     ucred_get(3C)attributes(5), standards(5)

```

`new/./socket.h.txt`

9

525 SunOS 5.11

Last change: 03 Aug 2006

8