

ICMP

This lecture describes the Internet Control Message Protocol (ICMP), a protocol used by routers to report problems routing packets over an IP network.

After this lecture you should be able to: state the contents and probable cause of each of the following ICMP messages: destination unreachable, TTL exceeded, echo request/response, and redirect; and list the sequence of packets transmitted and received by the ping and traceroute commands including the packets' source and destination addresses and their TTL field values.

Introduction

IP networks differ from many previous networks in that much of the network's management is implemented at the endpoints rather than within the network. Many functions, such as control of congestion (too much traffic for a link) and handling of link failures are performed by hosts and routers at the source and destination rather than by the core of the network. This makes IP networks scalable and reliable.

A protocol, the Internet Control Message Protocol (ICMP), defined in RFC 792, is used to implement many of the functions required for this distributed control model to work.

ICMP messages are not normally seen by users and are not used by most network applications. ICMP packets are generated and consumed by the software (typically operating system processes and drivers) that handle the IP packets before they are passed to applications.

The main purpose of ICMP is to handle error messages sent back to the source IP address as a result of routing failures such as routing table errors, excessive delays or routing loops, or queue overflows.

ICMP also includes an echo request function that can be used for diagnostic purposes. This function is used by the ping and traceroute utilities to map the current state of routing tables in routers between two points in the network.

Although ICMP packets have IP headers, they are not transport protocols like UDP or TCP, rather they implement a control protocol that affects the operation of the IP routing software. Like other control protocols like ARP, DHCP and SNMP, its functionality necessarily crosses protocol layers.


Many of the ICMP message types are rarely used. For example, quench messages that were originally meant to handle congestion have been deprecated.

In this lecture we will study only the most commonly encountered ICMP messages.

ICMP Packet Format

ICMP packets have IP headers so that they can be routed over an IP network. The protocol field is set to 1 (ICMP) rather than, for example, 17 for UDP or 6 for TCP.

The IP payload, the ICMP packet, begins with a one-byte type field followed by a one-byte code field and a 16-bit checksum. The contents of the rest of the ICMP packet depend on the type of ICMP message. For example, here is the payload of an ICMP Destination Unreachable message:



Many of the ICMP message responses include a copy of the offending IP header and the first 64 bits (8 bytes) of its payload which probably contain (at least part of) the higher-level protocol header. This data can be used for diagnostic purposes.

Common Message Types

The following sections describe the most common ICMP message types and a brief explanation of when they are generated.

Destination Unreachable

This message is generated by a router that cannot forward a packet. The code field gives more information

about the reason. The most common reason is code 0 which is “network unreachable.” This means there is no entry for that packet’s destination network in the host’s routing table. You will often see this message generated when a gateway’s link to the Internet is down.

The code “host unreachable” means that a host is not reachable (e.g. due to no response to an ARP request). Other less common codes include reassembly failed (fragmentation required but the no-fragment bit is set), or inability to use specific IP protocol values or UDP/TCP port values, possibly as a result of firewall restrictions.

Exercise 1: Write the contents of an ICMP Destination Unreachable packet with code 0 (Network Unreachable) including the IP header. Specify the values of any fields whose value are not defined by this question.

TTL Exceeded

An IP packet’s TTL field is decremented each time it is forwarded to another host. This ICMP message type indicates that the TTL field has been decremented to zero. Since the TTL field is usually initialized to a large value (e.g. 64) this message usually indicates a routing loop.


Another likely reason is the use of the traceroute utility described below.

Echo Request/Response

Echo request and echo response ICMP messages are used for diagnostic purposes, they are not error messages. These messages, typically generated and received by the ping utility, are the main tool for testing that IP frames are being routed between two hosts.

Unfortunately, these packets, like many other ICMP messages, can also be used by attackers to probe the layout of a network and so are often blocked from crossing firewalls.

The ICMP echo request/reply packets have a slightly different format:



The code is 8 for an echo request and 0 for an echo reply. The identifier and sequence number are

echoed back to the source so it can distinguish different echo requests which might become interleaved due to different packets having different routes or different transmission/queuing delays. The data is arbitrary data.

Redirect

This message might be generated by a router that detects that packets are being mis-routed to the wrong gateway on a local network. It tries to direct a host to use a different gateway. This sometimes happens when the default gateway changes and a host fails to update its routing table. It is an indication that these hosts should update their default routes (e.g. using DHCP).

Ping and Traceroute Utilities

The ping utility generates an ICMP echo request packet and displays received responses. The delay until the response is received indicates the round-trip time to the destination. The fraction of responses received indicates the packet loss rate along the path to/from the destination.

The traceroute utility generates packets (the protocol does not matter as long as the frame is passed on to the destination; ICMP echo request or UDP packets are the most common) with a TTL value that is set to a low-enough value to generate TTL Exceeded ICMP responses.

By incrementing the TTL starting from 1, the ICMP TTL Exceeded response will be received from routers at successively distant hops from the source host. Each packet is sent (typically) three times and the time until the response is received is displayed along with the source of the ICMP message.

Exercise 2: The host H1 in the network shown in the diagram below issues the command ‘traceroute H3’. Draw a table showing the ICMP messages transmitted and received by H1. Include five columns for source address, destination address, ICMP type, ICMP code and TTL value. Assume a single ping is used to discover each hop.

