Laura Chappell

Using ICMP to Troubleshoot TCP/IP Networks

Editor's Note: This article is based on Laura Chappell's upcoming book TCP/IP Analysis and Troubleshooting*, which will be available soon in both electronic and hard-bound format at http://www.podbooks.com.*

You don't need to feel like you are taking a shot in the dark when troubleshooting a TCP/IP network. The TCP/IP protocol suite includes Internet Control Message Protocol (ICMP), a message protocol that can help you identify network problems such as incorrect gateway settings, unavailable applications or processes, and fragmentation problems.

As a protocol analyst, I frequently use ICMP to troubleshoot network problems and examine network designs for faults. In fact, I often refer to ICMP as an "immediate gratification" protocol—one that can provide you with immediate help in troubleshooting network connectivity and design problems. If your job responsibilities include troubleshooting or testing TCP/IP networks, you should know ICMP inside and out.

This article explains the purpose and basic functionality of ICMP, the structure of ICMP messages, and the way to use ICMP to analyze a network. In addition, this article explains the development status of ICMPv6, the next release of ICMP.

UNDERSTANDING THE PURPOSE OF ICMP

In 1981, Jon Postel released Request For Comments (RFC) 792, which documents a host-to-host datagram protocol for reporting errors in datagram processing. (For more information about RFC 792, visit http://www.rfc-editor.org/rfc.html.) For example, RFC 792 stipulates that if a packet is sent to a gateway (or router) that does not know the route to a desired destination, that gateway should respond with an ICMP message that explains the problem.

ICMP messages provide feedback on communication problems such as the following:

- A client has been configured with the wrong IP address for its Domain Naming System (DNS) server. The destination device sends an ICMP message, indicating that this device does not support the DNS port (port 53).
- An application does not permit fragmentation of its communications, but fragmentation is required to communicate with the destination device. The router that would normally fragment the packet sends the source device an ICMP message, indicating that the packet could not be forwarded because the packet's "don't fragment" bit was set.
- A client sends all communications to a default router although another router offers the best route. The default router sends an ICMP message that includes the IP address of the router that provides the best route.
- A packet arrives at a router with a Time To Live (TTL) value of 1. All IP headers contain a Time to Live (TTL) value. Unlike the IPX hop count, which increments as the packet is forwarded through each router, the IP TTL value decrements as the IP packet is forwarded through each router. If an IP packet has a TTL value of 1, the router cannot decrement the TTL value by one and then forward the packet. Instead, the router discards the packet and sends an ICMP message, indicating that the packet's TTL expired in transit.

In addition, certain utilities use ICMP for testing and diagnostics. For example, to test communication to a destination device, the IP Packet Internet Grouper (PING) utility sends an ICMP echo to the device's IP address. If a communication route to the destination is available and the destination device

Illustration: Norman Felchle

is functioning properly, the destination device sends an ICMP echo reply packet.

To discover the routers in a route, many implementations of the Trace Route (TRACERT) utility send ICMP echo requests with varying TTL values. For example, the TRACERT utility will create and send an ICMP echo packet with a TTL value of 1. The router cannot set the TTL value to 0 and forward the packet. Instead, the router sends an ICMP message indicating that the destination device is unreachable because the TTL has been exceeded in transit.

The TRACERT utility now knows the IP address of the first router in the route. The TRACERT utility then increments the TTL to 2 and sends another ICMP echo request. The second router in the route responds and is added to the list of known routers in the route to the destination address. This process continues until the destination device receives the packet and sends an echo reply. (For more information about how to use the PING utility or the TRACERT utility for troubleshooting TCP/IP networks, see "Troubleshooting TCP/IP Networks: Building Your Toolkit," *NetWare Connection*, Jan. 1999, pp. 20–28. You can download this article from http://www.nwconnection.com/past.)

RECOGNIZING THE STRUCTURE OF ICMP MESSAGES

Before you can use ICMP packets to troubleshoot your company's network, you must understand the ICMP packet structure and know how network analyzers differentiate ICMP packets from other network traffic. ICMP packets are a little unusual. For example, ICMP packets do not rely on User Datagram Protocol (UDP) or Transmission Control Protocol (TCP). Instead, ICMP packets sit directly on the IP header. (See Figure 1 on p. 35.)

The protocol field in the IP header indicates that the ICMP header will follow the IP header. The assigned value for ICMP in this field is 1 (decimal). (Novell's LANalyzer for Windows automatically translates this value of 1 to ICMP.)

The Structure of Echo Request Packets

The structure of the ICMP packet depends on the type of information being exchanged and the function of the packet. For example, in Figure 1, the packet structure is simple. It includes the following fields:

- Type. This field identifies the packet as an echo request packet. (Echo request packets are type 8. For a list of types, see "Types of ICMP Messages" on p. 34.)
- Checksum. This field indicates the checksum value of the ICMP message, starting with the ICMP type field. (The checksum value does not include the value of the checksum field itself.)
- Code (optional). This field further defines the purpose of the ICMP message. This particular field is not used in ICMP echo messages. However, this field is used in the following messages: destination unreachable messages, redirect messages, alternate host address messages, time exceeded messages, and

Types of ICMP Messages

The following list defines the types of ICMP messages that can be sent on a network. This list is based on Internet Assigned Numbers Authority (IANA) documentation. To obtain the most current version of this list, visit http://www.iana.org or http://www.netanalysis.org.

Type	RFC Definition (if available)
0	Echo Reply (RFC 792) (used for PING reply)
1	Unassigned
2	Unassigned
3	Destination Unreachable (RFC 792)
4	Source Quench (RFC 792)
5	Redirect (RFC 792)
6	Alternate Host Address
7	Unassigned
8	Echo (RFC 792) (used for PING request)
9	Router Advertisement (RFC 1256)
10	Router Selection (RFC 1256)
11	Time Exceeded (RFC 792)
12	Parameter Problem (RFC 792)
13	Timestamp (RFC 792)
14	Timestamp Reply (RFC 792)
15	Information Request (RFC 792)
16	Information Reply (RFC 792)
17	Address Mask Request (RFC 950)
18	Address Mask Reply (RFC 950)
19	Reserved (for Security)
20–29	Reserved (for Robustness Experiment)
30	Traceroute (RFC 1393)
31	Datagram Conversion Error (RFC 1475)
32	Mobile Host Redirect
33	IPv6 Where-Are-You
34	IPv6 I-Am-Here
35	Mobile Registration Request
36	Mobile Registration Reply
37	Domain Name Request
38	Domain Name Reply
39	SKIP
40	Photuris
41–255	Reserved (JBP)

The initials JBP indicate that Jon B. Postel, one of the developers of the Internet protocol suite, is associated with these numbers. Postel helped shape the communications system of the Internet and millions of private networks until his death in October 1998. For example, Postel established the IANA, which coordinates the allocation and registration of Internet addresses and names. You can learn more about this Internet luminary at http://www.iana.org/postel/index.html.

parameter problem messages. (For a list of codes these ICMP messages use, visit the *NetWare Connection* web site at http://www.nwconnection.com.)

- Identifier. This field is used to match ICMP echo requests and replies.
- Sequence Number. This field is used to number ICMP echo requests and ICMP echo replies.

The Structure of Redirect Messages

Figure 2 shows the format of an ICMP redirect message (which is identified as a type 5 ICMP message). (See p. 38.) ICMP redirect messages are used to dynamically update a host's routing tables so that the host will use another router to reach a particular destination device.

A router sends an ICMP redirect message when that router knows that another router has a better route to the destination device. You will often see ICMP redirect messages on networks that use a default gateway setting. When preparing to transmit a packet, a host checks to see if it has a route entry for the desired destination device or network. If the host's routing table does not contain an entry for this information, the host sends the packet to a default gateway. The default gateway replies with the IP address of the router that offers a better route if one is known.

For example, in Figure 3, Katie is sending a packet to the FTP server on network 11.0.0.0. (See step 1 in Figure 3 on p. 40.) Since Katie's routing tables do not contain an entry for the FTP server, Katie sends the packet to the default gateway.

When the default gateway responds, it includes the original IP header, eight bytes of the datagram from Katie's packet, and the IP address of the router that Katie should use to reach the 11.0.0.0 network. (See step 2 in Figure 3.) You may wonder why the default gateway includes eight bytes of the datagram. The eight bytes contain the destination and source port fields. By examining these eight bytes, you can determine what process sent this packet.

The ICMP specifications stipulate that the redirecting gateway should forward the original datagram to the appropriate router. (See step 3 in Figure 3.)

CONNECTION-ORIENTED OR CONNECTIONLESS?

Is ICMP a connection-oriented or connectionless protocol? ICMP is connectionless because it does not require hosts to handshake before establishing a connection.

Connectionless protocols have advantages and disadvantages. The main advantage of connectionless protocols is less overhead: These protocols don't have the overhead of establishing a connection before sending a simple communication error message.

The main disadvantage of connectionless protocols is that delivery of ICMP messages is not guaranteed. If an ICMP message gets lost in transmission, the communication error must occur again, prompting another ICMP message to be transmitted.

ANALYZING YOUR COMPANY'S NETWORK WITH ICMP

Before you can use ICMP to troubleshoot your company's network, you must capture the ICMP traffic on that network. You can set up a network analyzer to capture all TCP/IP traffic and filter just the ICMP traffic (post-filtering), or you can set up a prefilter to capture just ICMP traffic (if the network analyzer you are using provides prefiltering capabilities). For example, I use post-filtering with Novell's LANalyzer for Windows and ManageWise, but I use prefiltering with Network Associates's Sniffer and Sniffer Pro. (See "Using Sniffer to Read ICMP Messages" on p. 38.)

After setting up the network analyzer to filter ICMP traffic, you should take a good look at the ICMP traffic that crosses the network. How many ICMP redirect messages do you see? It is typical to have some redirect messages (especially during start-up hours in the morning), but if one device is constantly being redirected before communicating with other devices on the network, you may need to assign that device a different default gateway.

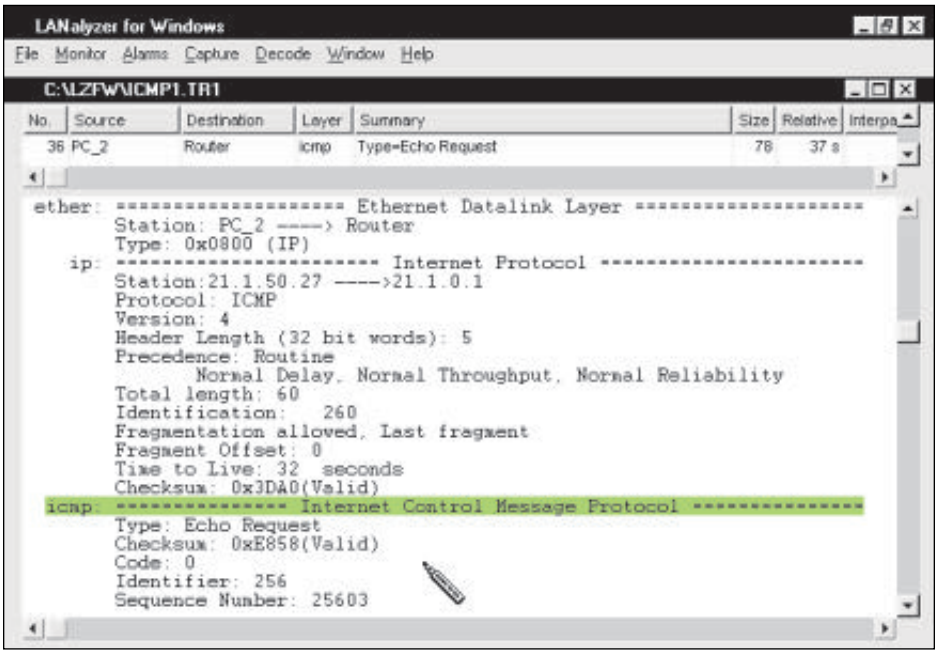

Figure 1. The ICMP header immediately follows the IP header.

Network and host unreachable messages may indicate a route or routing

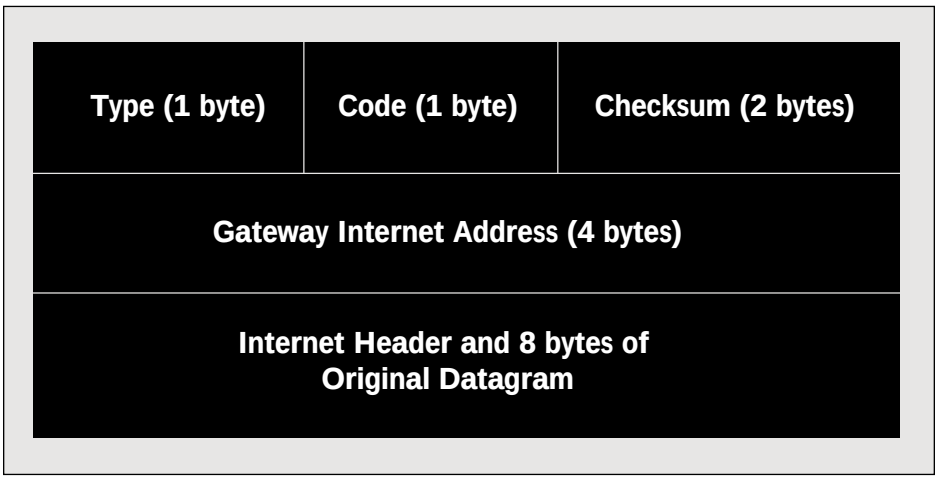

Figure 2. The ICMP redirect message format

failure. For example, if a router cannot forward a packet addressed to a certain device or network because that device or network is considered "down," the router will send a network unreachable or host unreachable message to the source device. This problem could be caused by a faulty IP stack on the destination device or by routing failures that have made a network unreachable.

Port unreachable messages, on the other hand, may indicate that a device is configured incorrectly. For example, if a device continually sends DNS queries to a specific IP address and receives port unreachable messages, the IP address for the DNS server may not be valid.

ICMP AND INTERNET SECURITY

Although ICMP messages are invaluable for troubleshooting TCP/IP networks, you should be aware that hackers find ICMP messages equally useful. For example, excessive port unreachable messages may be the first sign that a hacker is trying to discover what network services are running on a network. Port scanning utilities often use the simplistic approach of sending packets to a device and incrementing the destination port number by 1 in each packet. Port unreachable messages help determine which ports are not active, thereby identifying the ports or processes that are available on a system. Because hackers sometimes use port unreachable messages in this way, you should carefully examine these types of messages on your company's network.

You should also examine the echo request and echo reply messages being transmitted on the network. Hackers sometimes use echo requests to "discover" IP addresses of live devices on the network. If echo requests are being used in this way, the destination IP address is typically incremented by one in each message. For example, you will see an echo request sent to 10.0.0.1, an echo request sent to 10.0.0.2, an echo request sent to 10.0.0.3, and so on.

These types of requests may also be sent by a management product that is building a map of your company's network (and, therefore, has a legitimate reason for discovering devices). However, if an unknown or suspect device is performing this type of discovery, it can be the first sign that a hacker is attempting to get information about your company's network.

In addition, hackers use ICMP messages to cripple network devices. For example, if you find an excessive number of ICMP echo packets on a network, you may have cause for concern. An excessive number of ICMP echo packets may indicate a denial of service attack. A *denial of service attack* focuses on overloading or crippling a device to the point that it cannot provide services to other devices.

Because hackers can use ICMP messages to gain information about a network or to actually harm a network, many companies restrict devices from transmitting specific ICMP messages across their connection to the Internet. If your company's security policy doesn't cover ICMP messages, you may want to revise it to include such a restriction.

THE FUTURE OF ICMP

ICMPv6 will provide error and information reporting on IPv6 networks. The initial specification for ICMPv6 was developed in 1995 and documented in RFC 1885. It was updated in December

Using Sniffer to Read ICMP Messages

If you are using Network Associates's Sniffer, you will need to do some extra work when analyzing Internet Control Message Protocol (ICMP) packets. If an ICMP reply packet includes the original IP header, you must complete the following steps to analyze that packet. (These steps are highlighted in Figure 4 on p. 40.)

1. The ICMP message indicates that the port the client was trying to access is unreachable. The first step is to determine which port the client was trying to reach.
2. Highlight the destination IP address in the ICMP header. (As the main article explains, ICMP messages include part of the original packet that triggered the ICMP reply.) Be careful not to highlight the actual IP header that is placed after the Ethernet header.

   Sniffer automatically highlights the corresponding bytes in the hex window. (Novell's LANalyzer for Windows also automatically highlights these corresponding bytes.)
3. The port number fields are the next four bytes in the hex portion. Convert 0x0089 to a decimal number using the Windows calculator in scientific mode.

   0x0089 = 137 decimal = NetBIOS Name Service

   Now you know that 164.64.1.2 does not support NetBIOS name services. (To view a list of port numbers, visit http://www.netanalysis.org or http://www.iana.org.)

   If you see a high number of replies indicating that a port number is unreachable, you should consider reconfiguring the device that requests those services so the device sends its requests elsewhere.

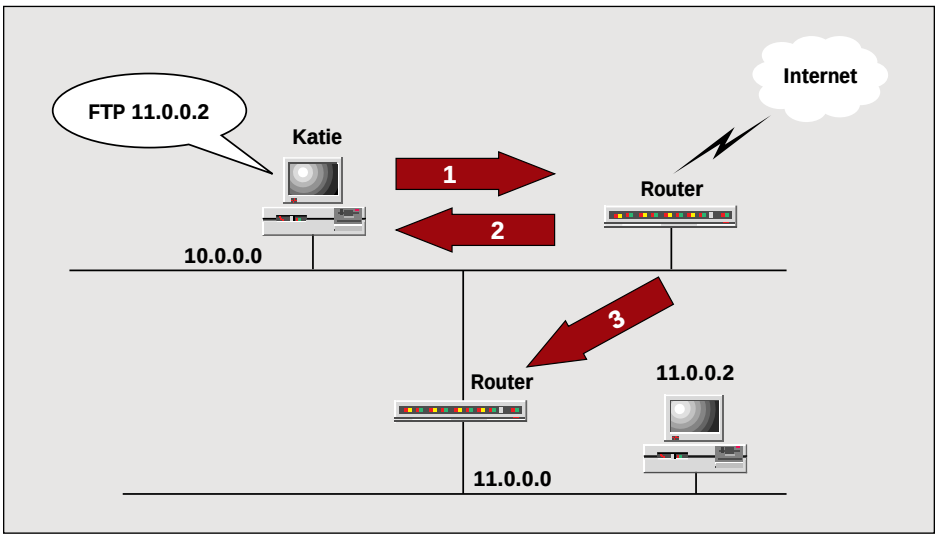

Figure 3. ICMP redirect messages are sent by routers.

1998 by RFC 2463. (For more information about RFC 2463, visit http://www.rfc-editor.org/rfc.html.)

In IPv6, value 58 in the Next Header field indicates that the packet is an ICMPv6 packet. ICMPv6 messages are grouped into two classes: error messages and informational messages. Error messages have message types from 0 to 127, and informational messages have message types from 128 to 255.

RFC 2463 defines the following message formats:

ICMPv6 Error Messages
1 Destination Unreachable
2 Packet Too Big
3 Time Exceeded
4 Parameter Problem

ICMPv6 Informational Messages
128 Echo Request
129 Echo Reply

The basic functionality of ICMPv6 is quite similar to ICMPv4. For example, RFC 2463 stipulates that ICMP messages should be transmitted under the following circumstances:

- A router should generate a destination unreachable message if a message cannot be delivered to its destination address. However, the router should not send this type of ICMP message if the packet is dropped due to congestion.
- A router should generate a packet too big message if a packet cannot be forwarded because it is larger than the Maximum Transmission Unit (MTU) of the next segment.
- A router should generate a time exceeded message if the packet's TTL has expired in transit.
- A device should generate a parameter problem message if the packet cannot be processed because the device detects a problem with a field in the IPv6 header or extension headers.
- A device should generate an echo request message and an echo reply message for troubleshooting and diagnostic purposes.

For more information about the types of messages being defined for ICMPv6, read RFC 2463. Other types of messages will be documented in separate drafts. (For more information about IPv6, visit http://www.ipv6.org, or see "IPv6: At the Starting Line," *NetWare Connection*, May 1999, pp. 6–17. You can download this article from http://www.nwconnection.com/past.)

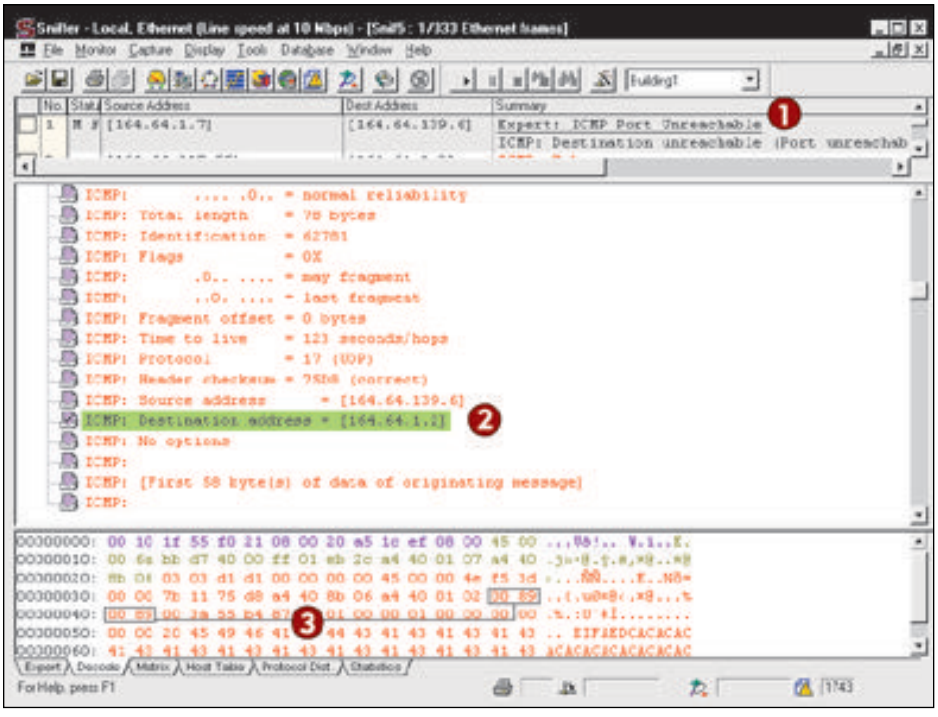

Figure 4. If you are using Network Associates's Sniffer or Sniffer Pro, you must perform three steps to decode ICMP messages.

CONCLUSION

I consider ICMP one of the key components of my online analysis and troubleshooting of TCP/IP networks. A quick look at the ICMP traffic patterns can save hours of troubleshooting and reconfiguration time. (For more information about TCP/IP analysis and troubleshooting, look for *TCP/IP Analysis and Troubleshooting*, a podbook that will soon be available at http://www.podbooks.com.)

A frequent contributor to NetWare Connection, *Laura Chappell is a senior protocol analyst for Network Analysis Institute. In addition to providing onsite and offsite network analysis, Ms. Chappell provides training classes on troubleshooting and optimization techniques. And she has actually seen Siegfried and Roy in Las Vegas. You can reach Ms. Chappell at lchappell@netanalysis.org.*