TECHNI-TOOL®
Tomorrow's Products ▶ Shipped Today

Frontline Network Troubleshooting

Practical procedures and methods using the LinkRunner™ Network Multimeter

Our thanks to Fluke Networks for allowing us to reprint the following article.

When a network client goes down, a technician has the simplest job description around: make it work. In today's business climate, making it work quickly and economically is their top priority. In any network environment, giving technicians proper training, the right tools, and a solid methodology for using them makes network troubleshooting faster, saving technician time and getting network users back to productivity faster.

A troubleshooting mindset and methodology

The key to successful troubleshooting is for the technician to know how the network functions under normal conditions. This enables the technician to quickly recognize abnormal operation. Any other approach is little better than a shot in the dark.

Unfortunately, many networking products are not delivered with adequate performance specifications, theory of operation, or condensed technical data to aid in troubleshooting. The successful technician will thoroughly study whatever data is available, as well as develop in-depth insight into the function of all components and how to operate them. Finally, he or she will remember that conditions appearing to be serious defects are often the result of improper usage, configuration, or operator error.

The foundation of this insight is best gained from practical experience. The true troubleshooting master learns in the trenches, through trial and error, comparing notes with others, and discovering tried-and-true methods that are not taught in school. The following information can help shorten your learning curve and give you proven advice on how to isolate and solve network problems.

Successful troubleshooters quickly master the following basic concept: a few minutes spent evaluating symptoms can eliminate hours of time lost chasing the wrong problem. All information and reported symptoms must be evaluated in relation to each other, as well as how they relate to the overall operation of the network; only then can the technician gain a true understanding of what they indicate. Once you have collected data about the symptoms, you will need to conduct tests to validate or eliminate what you think the problem could be.

Five steps to successful troubleshooting

The successful technician follows a consistent methodology when approaching and solving any network related problem:

1. Document your network

Having access to up-to-date documentation such as physical and logical maps, performance baselines and audits, device inventories, configurations, address-to-host tables, etc., will dramatically reduce the amount of troubleshooting time spent in "discovery mode" where you are simply trying to figure out where the PC connects in the larger scheme of the overall network.

2. Collect all available information, and analyze the symptoms of the failure.

Ask yourself if you understand the symptoms. Can the user demonstrate the problem or can you recreate it? Determine if something was altered at the workstation or on the network just before the problem started.

Fluke Networks' LinkRunner is an ultra-portable, affordable tool for aiding technicians during the basic, early stages of network troubleshooting.

Designed as a simple tool for deployment throughout the organization, LinkRunner helps frontline technicians find basic connection problems and eliminate physical-layer issues before escalating the trouble-ticket to a more senior technician.

The LinkRunner is perfect for doing rapid on-site testing of basic network connection health as a means to solving many network problems, and as a necessary prelude to more complex problem resolution processes. When deployed by a technician at the problem site, the LinkRunner confirms a number of critical network operational parameters and provides information that can form the basis for solving problems with causes at higher levels in the ISO 7-layer stack.

3. Localize and isolate the problem

Reduce the scope of the problem. Is the problem related to a segment of the network or is it isolated to a single client? Within a single client, the problem can be further isolated to the network, the physical cabling or the workstation. As you will see, the process of collecting information and isolating the problem are often concurrent activities.

4. Correct the problem and verify problem resolution

Once isolated, identifying and repairing the specific fault should be simple. For network hardware, it is most expedient to replace a part, such as replacing a bad patch cable, changing hub/switch port or client network interface card (NIC). This step is complete when the user tests for the problem (what they were doing when it first occurred) to ensure that it has been repaired.

5. Document what you did

Go back to Step 1. Having a record of problems and their resolution (such as is offered in many trouble-ticketing applications) builds an internal knowledge base that can be referred to when similar problems occur in the future. This information speeds future troubleshooting sessions.

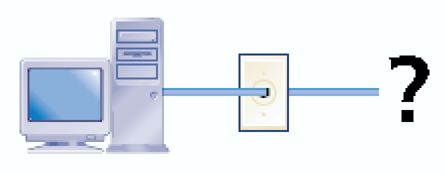

Is this trip really necessary?

Even with continuous improvement being made in operating system software reliability, “reboot your PC” is still the mantra of help desk technicians. Sadly, a cold-start reboot resolves so many otherwise inexplicable problems that it really is an unavoidable step. A good side effect of this step is that the problem is solved without the technician having to leave their desk.

Aside from asking the user to reboot the machine, further information gathering can take place over the phone with the help of the user before heading out to the workstation.

Most users can open a command prompt and report back to the technician the result of an IPCONFIG command. This step tells the technician whether the PC has an appropriate address for the subnet to which it is physically connected.

- If the PC is configured for Dynamic Host Control Protocol (DHCP), but returns a Windows default IP address (169.254.x.x), then the client is not communicating with the DHCP server.
- Portable computers will receive addressing suitable for the network they are connected to, but sometimes the DHCP lease for another network subnet is kept after the computer is moved. Have the user explicitly make a new request with these two commands from the command prompt:

```
C:\>ipconfig /release
```

```
C:\>ipconfig /renew
```

  Have the user attempt to use the network following receipt of a fresh IP address. If the IPCONFIG command reports that the DHCP operation cannot be performed, then the user is probably using a static IP configuration. Validate the reported IP address according to your network documentation.

- If the user reported a valid IP address, try pinging that address from your desk. If the user’s PC responds, then have the user attempt some other network activity, such as opening a web page or pinging the local router to verify basic connectivity. If these tests do not solve the problem, then a visit to the user’s computer is warranted.

Verifying the problem onsite with the client

Upon arrival at the suspect workstation, the information gathering process begins in earnest. Question the user about any action or activity that may have affected network performance. This is sometimes of limited value as the user is either unaware that many common actions related to the workstation or to the workspace can affect network performance, or they know full well that something they did was inappropriate and they are not about to admit it. Giving the user the benefit of the doubt, be sure to ask about any recent changes, even moving office furniture or installing a new screen saver.

Repeat the tests you asked the user to perform over the phone. A successful ping to a network server or off-net device immediately confirms that the workstation has Layer 3 connectivity to the network, all lower-layer tests are instantly deemed “not needed” and the tech can focus his or her attention elsewhere.

If Layer 3 connectivity cannot be validated, then you must start at Layer 1.

If the problem is a dropped or intermittent connection, a continuous ping sends an unending stream of echo request packets to the target device. Response time for each successful ping or non-response will be shown.

```
C:\> ping -t x.x.x.x
```

Lengthy response or dropped pings can be further investigated by running a trace route (TRACERT or PATHPING) to the target device. Trace route can tell you where along the network path the delay or dropped packet is occurring, and troubleshooting at Layer 1 begins there.

```
C:\> tracert x.x.x.x
```

or

```
C:\> pathping x.x.x.x
```

Is extended troubleshooting required?

If the problem has not yet been identified or isolated while verifying the user's problem report, a more detailed and lengthy investigation may be necessary. However, in much the same manner as a few quick tests were performed before leaving the help desk, it may still be possible to conclude this situation quickly.

Now that the user's report indicating an inability to log into the network has been verified, the first question for any technician involves whether the issue relates to the network or the user's PC. The next step is to determine whether the cable connecting the client to the network is in place, functioning properly, allowing the client access to the network. Ensuring these functions solves many problems and lays the proper foundation for solving more complex issues.

Solving network problems in a timely, cost-effective manner requires that frontline technicians have at their disposal a tool that will quickly verify the status of critical network functionality – a tool like the Fluke Networks LinkRunner Network Multimeter.

Frontline troubleshooting "Instant Tests":

1. Test for link

2. Check overall segment activity

3. Use DHCP as a diagnostic tool

4. Ping local and remote destinations

Test for link

Many network technicians have come to believe that the presence of a link light on a network interface card indicates the presence of link pulse. While reliable on some equipment, many link LEDs are controlled by software in the host system and are "turned on" when higher layer network activity is sensed. Some NICs also feature activity lights which indicate the presence of traffic; these tend to be more reliable than the link LED as a verification that network is present. However, neither of these LEDs offers any indication of speed or duplex setting, requiring additional investigation to verify.

Link may be determined by a process called auto-negotiation, where the two link partners exchange information about speed and duplex capabilities. Upon exchanging the information, the two link partners compare capabilities and then initiate communications at the highest common speed and duplex match. If one of the two partners is misconfigured or has flawed drivers, this process may fail to achieve a common setting, and communications may proceed intermittently or fail altogether.

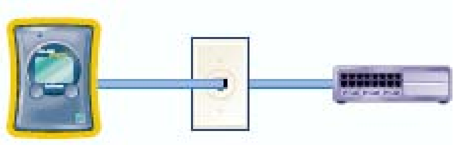

LinkRunner tests for the existence of link - a signal successfully sent and received over a single wire segment.

Upon connection to a network drop, LinkRunner first seeks to establish link with whatever "link partner" is found at the other end, whether that is a network connection (hub or switch) or a PC's network interface card. LinkRunner follows the process of auto-negotiation as specified in the IEEE 802.3 standard, and is a hardware-controlled link indicator, not software-controlled. After a successful auto-negotiation process, LinkRunner's LED will light bright green and the resulting speed and duplex setting is shown in the upper left corner of the LinkRunner screen.

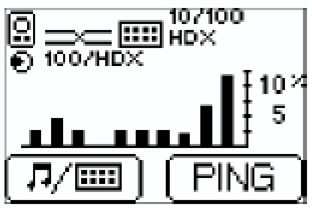

LinkRunner screen showing successful link, speed and duplex settings, and utilization.

Check overall segment activity

If network traffic is seen on the wire, LinkRunner indicates the level of that traffic in the main screen's utilization strip chart. Be reminded, however, that if connected to a single switch port (not shared media) the only traffic seen may be broadcast frames, which can be very intermittent on low traffic networks.

If you are testing a shared Ethernet environment, one which still uses hubs instead of switches, it is very likely that your network is operating in half duplex. Half duplex Ethernet is limited by the number of stations seeking to transmit at the same time, and by the size of the frame they are seeking to transmit. If too many stations attempt to transmit simultaneously, the Ethernet performance may suffer dramatic drops in performance due to collisions. If instead you are testing a network where each station is connected to a separate switch port, then the danger of excess collisions is negligible.

While the existence of collisions is a normal part of half duplex Ethernet operation, when the number of collisions begins to rise due to increasing traffic, the traffic volume will begin to rise at an increasing level because of the re-transmissions required. The result is a network that displays a performance curve that suddenly "falls off a cliff" as the number of frames sent, collisions, and re-transmitted packets spirals upward and a rapidly-increasing rate. As performance decreases, users will begin noticing delays and calling in trouble tickets.

In most networks, the level of traffic on the Ethernet is insignificant, and the problem is to be found elsewhere. By providing basic information on network segment usage statistics, the LinkRunner can provide essential clues to the source of network performance that is significantly below user expectation.

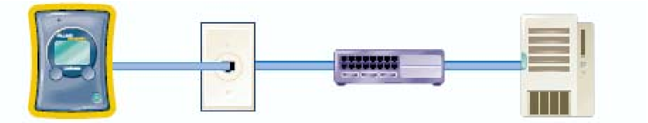

The success of DCHP address assignment confirms the ability of the client to talk on the network and successfully obtain an IP address - layer 1 through 3 verification in one step!

DHCP as diagnostic tool

If link can be established and utilization is reasonable, the user may then press the button corresponding to the ping test. LinkRunner will then attempt to obtain an IP address from the network's DHCP server. DHCP is normally a broadcast-based technology. As such, it requires either a separate DHCP server for each subnet (expensive and difficult to manage) or DHCP relay proxies or agents which are used to forward requests and replies between clients and servers when they are not on the same physical subnet. These directed broadcast helper applications on routers are the common tool used by large organizations that prefer to have DHCP servers run from a central location. The failure of either a client's or LinkRunner's automatic DHCP configuration could point to a problem with the DHCP relay system.

DHCP is available from most networks now, though LinkRunner can accept a manual or static IP configuration if necessary. The process of obtaining a DHCP address demonstrates the viability of the local cable, the local hub or switch port, and the network infrastructure all the way back to the DHCP server. In one simple operation most of the nearby network infrastructure has been validated up through Layer 3.

Ping local and remote destinations

Ping has been one of the most used networking troubleshooting tools in networking history. Included with every popular Internet-capable operating system, ping is one of the first steps in the network troubleshooting process used by most network technicians. Why is this simple utility so useful?

Ping is, in operation, much like SONAR used for oceanographic purposes. The ping utility sends a signal (typically an ICMP "echo request" frame) that "bounces" off the destination device (which generates an "echo reply") to tell the sender whether the destination system is there, and how long the signal took to reach the destination and return.

Once configured with an address from the DHCP server, LinkRunner will immediately initiate its ping test to the DNS (Domain Name Service) server and default router, both addresses being supplied by the DHCP configuration process. LinkRunner may also be preconfigured with up to four additional destination IP addresses, one of which may be selected for inclusion in the automatic ping test. Successfully pinging key network services such as web applications, user authentication, and so on, indicates that the services themselves should be available from the client location.

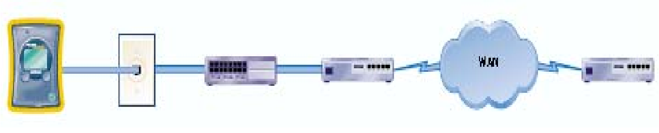

Ping may be performed against local systems and systems located across the Internet.

The simple success of a ping indicates that end-to-end Layer 3 connectivity exists between the two devices. The total roundtrip travel time for the request is easily compared to known values to provide a helpful diagnostic for more detailed analysis, if deeper analysis is required. However, ICMP requests are low-priority traffic and may be discarded if one of the routers along the way or the destination device is busy. This is why it is useful to send a series of pings to give the destination multiple opportunities to respond.

Servers outside the enterprise network may also be used as the target for pinging to verify WAN interconnectivity from the client station and local site to a remote site. If servers within the firewall respond to ping, but those outside the firewall do not, then network technicians might look at routers or other network boundary infrastructure for sources of the problem. If some servers respond while others do not, then network technicians can explore why particular network segments are unavailable. If pings are successful to both external and internal servers involving applications and servers but the client is not, in fact, receiving those services, it indicates the problem lies at a level beyond the physical transport. A successful ping implies that other traffic should be reaching the destination server, and any continued inability to access services is probably related to the server or the user's login account.

What's next?

If the problem has not been identified or resolved with the initial "instant tests," there are two clear directions for further action.

- If the tests demonstrate an inability to achieve an Ethernet link, then it's time to begin looking seriously at the network cable.
- If all instant tests are successful, showing link and reasonable network segment traffic levels, acquiring an address via DHCP, and

successfully pinging critical network servers, then the problem can be referred to higher technical levels for resolution at different network layers, to a user account administrator, or to a personal computer technician for exploration of workstation configuration issues.

Cable test

The first cable to check is the patch cable connecting the workstation or device to the network wall jack. For this, simply plug one end of the cable into the LinkRunner's network jack and the other end into the wiremap jack for a quick cable continuity and correctness test. If the patch cable is good, then it can be plugged back into the wall or floor jack and used as a component in further testing.

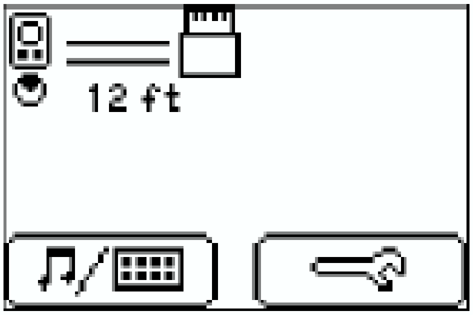

LinkRunner showing a good patch cable test ...

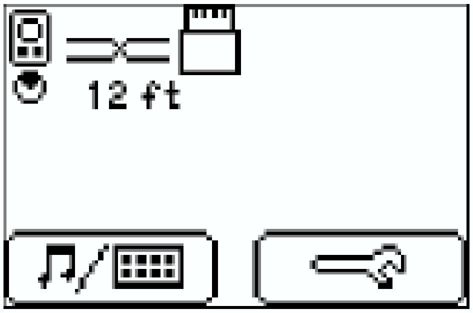

a crossover patch cable ...

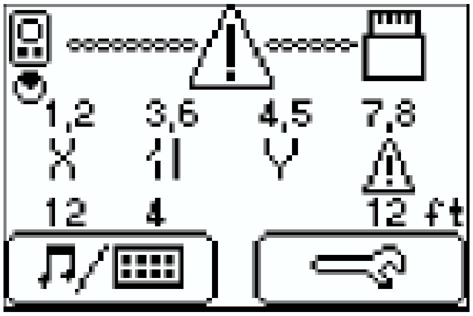

and a very bad patch cable.

The next step in finding the source of a cable problem is to trace the cable into the wiring closet and the local switch. In a crowded closet, picking the right cable out of a tangle coming out of a conduit can be difficult and time-consuming, but LinkRunner has two functions that can make finding the cable faster and easier. First, the LinkRunner can inject a tone onto a cable for audible tracing with a tone probe. With a standard tone probe, cables may be systematically inspected until the warbling tone identifying the proper cable is heard. This technique is useful when it's not certain that the cable is connected to the switch, or if no documentation is available indicating which wiring closet or switch the cable is supposed to go to.

Once the far end of the user's cable is located, the LinkRunner's Wiremap Adapter or one of the optional LinkRunner Cable ID Accessories may be used to terminate the link, and the entire horizontal cable run may be tested for continuity and proper pairing.

At the same time Link Runner is transmitting tone, it is also attempting to flash the switch port link light every three seconds. Discovering which port the cable is attached to is greatly facilitated by this feature. Once the port is located, try switching the cable to another unused port. Marginal or failed ports often still show link, so switching to a different port may solve the problem.

If the hub or switch port tests good, then the workstation is suspected of being the source of the problem. The workstation NIC can be verified by connecting LinkRunner directly. Just as with a hub or switch, LinkRunner will indicate the presence of link and the speed and duplex settings offered by the NIC. If link is present, try rebooting the PC or using a command line utility such as ping in order to initiate traffic for LinkRunner to monitor. If LinkRunner does not report any traffic, even though the PC claims to be transmitting, then verify the bindings and other configuration parameters on the PC. If LinkRunner reports both link and the PC's traffic, then proceed to diagnosing the PC's networking configurations.

Upper layer diagnostics

If the workstation establishes link on the network, the next step is to ensure that the addressing of the workstation is appropriate for the attached subnet. Confirm that the workstation is employing the proper protocol stacks, and that they are properly configured. Finally, the technician must verify that all necessary program components and libraries are in place. This is usually all verified by deleting the protocol or NIC from the workstation configuration and reinstalling it. If all of these components are in place and properly configured, and the workstation still does not show proper network and application connectivity, it is time to escalate the problem beyond the field technician level.

Proper tool, proper job

A low-cost tool with a simple interface, the LinkRunner is the right tool for broad deployment, placed in the hands (or on the belt) of department-level technicians. Just the simple act of always having a “known good” network device can serve to eliminate the doubt inherent to the ubiquitous laptop test.

As any network technician will attest, assuming that the client device is “known good” can lead to far more problems than it solves, involving not only technical issues, but political and organizational issues when support of the network and desktop clients comes from two separate groups. The issue is trickier technically (though simpler politically) when the device at the site of the trouble ticket isn’t a desktop computer, but is, instead, a network infrastructure component.

In this case, isolating the problem to (or eliminating it from) the link itself can rapidly point to the location of the problem when the normal device diagnostic screens aren’t available. When issues can be resolved at the department or group level, rather than escalating every problem to senior centralized IT technicians, problem solving is more cost-effective and efficient. More advanced tools can be deployed less widely, to the technicians and engineers trained in more advanced problem solving. The result of which is a problem resolution plan that ensures appropriate use of more expensive human and technical resources.

CLICK HERE to view Fluke Network's LinkRunner™ Pro

FLUKE networks®

1547 N. Trooper Road • P. O. Box 1117 • Worcester, PA 19490-1117 USA
Corporate Phone: 610-825-4990 • Sales: 800-832-4866 or 610-941-2400
Fax: 800-854-8665 or 610-828-5623 • Web: www.techni-tool.com