LAN Fault Timing Analyzer (FTA) Software
User Information

Introduction

The LAN Fault Timing Analyzer (FTA) software is a technical tool that measures fault recovery times in self-healing Ethernet networks while in operation. Either ring or mesh structures may be measured. FTA determines the time required to re-establish the pre-existing Ethernet traffic flow after a fault occurs, and provides a display of the interrupt time in milliseconds (ms). Since test sessions will cover multiple fault experiences in order to characterize a self-healing LAN's behavior under stress, FTA also provides basic statistical data of minimum, maximum and average fault times.

FTA runs on any Windows®-based Personal Computer (PC). It needs a target node on the network to communicate with during the test. For best results, the target node should be a device that does not announce its IP address during the test. Typical targets are industrial PLCs and managed network devices such as the node in a managed switch or hub. PCs are not good targets because they frequently send out network IP announcements to accommodate moves and changes, and these announcements interfere with the fault timing tests.

The FTA software consists of a single executable file, compatible with any Windows PC running 32-bit Winsock network. Windows 95 or later versions should work properly. The computer needs TCP/IP networking software (normally already installed for the PC's NIC card) and it should be configured with an IP Address for use as a node on a LAN.

FTA is Freeware. To run FTA, users may copy the executable file (Windows-based file named *FTA.exe*) into a convenient folder and then double click on the file name to launch the FTA program on a dedicated PC. FTA software is available via the Internet, by going to the URL: www.garrettcom.com/redundancy.htm.

Theory of Operation

The FTA program sends ICMP ping packets in rapid succession across the network to the target device, and records the ICMP ping replies. A ping request with a normal response indicates the

network is working properly. A ping request with no response indicates a problem due to a fault or traffic interruption, and this condition starts the fault timer clock. Ping requests without responses are counted by FTA to indicate the duration of a fault, using a 1 millisecond (ms.) counter.

Upon experiencing a fault, a LAN with self-healing capabilities will recover using an alternate path for the traffic. FTA will then see responses to its ping requests again, and stops the FTA fault timer clock. The no-response interval is the fault recovery time.

Ping Target Selection

Any network node with an IP address can work as a ping target for FTA so long as it can respond to pings within a millisecond, a very slow rate by network speed standards. The target should never (or very rarely) announce itself after initialization. An industrial PLC, or the agent note in a spare managed switch, is a good choice for industrial LANs.

A TCP/IP target or host that does not send out unsolicited packets will always be the best choice for a ping target. PCs can cause erroneously short FTA fault recovery times because they send out packets frequently (usually every 12 to 15 seconds) to announce their presence on the LAN, a procedure that helps with administering and adapting to moves and changes in a network.

Ethernet switches in a LAN hold the addresses of attached nodes in their address table. To recover from a fault, the switches must re-direct some or all LAN traffic. To do this, the switches must re-learn nodes at new locations that are reachable only through the back-up path. A target that broadcasts an announcement, such as a PC, forces the intervening switches to re-learn that particular IP address. If it is the ping target, the next ping from FTA will get through. Since the target will respond, the FTA software will erroneously declare the fault recovery process complete, although other nodes, which did not announce themselves, may still be unreachable. At this point, path recovery is complete, but the fault condition may not be resolved. All switch buffers have to either be reset, or age out and re-learn all node addresses before fault recovery is complete.

FTA User Interface.

The FTA user interface is straightforward and easy to use. To initialize FTA, enter the IP address of the ping target device in the dialog box (Figure 1) Once started, the program will automatically terminate after 2 million ICMP ping requests (typically 6 hours of testing).

As each fault recovery or restoration test is performed, the timing data for that test will be displayed. When a series of tests are performed, statistics are maintained that provide minimum, average, and maximum fault recovery data over the life of the set of test sequences. At the completion of a test series, the user should record the test results data on an FTA Test Report Form, and / or hit “Stop” to use the PRINTSCREEN function to keep a record. To reset and start another test series, click on the “Reset” button. The target IP address will be left in place, and all test data fields will be reset and ready for new tests to begin.

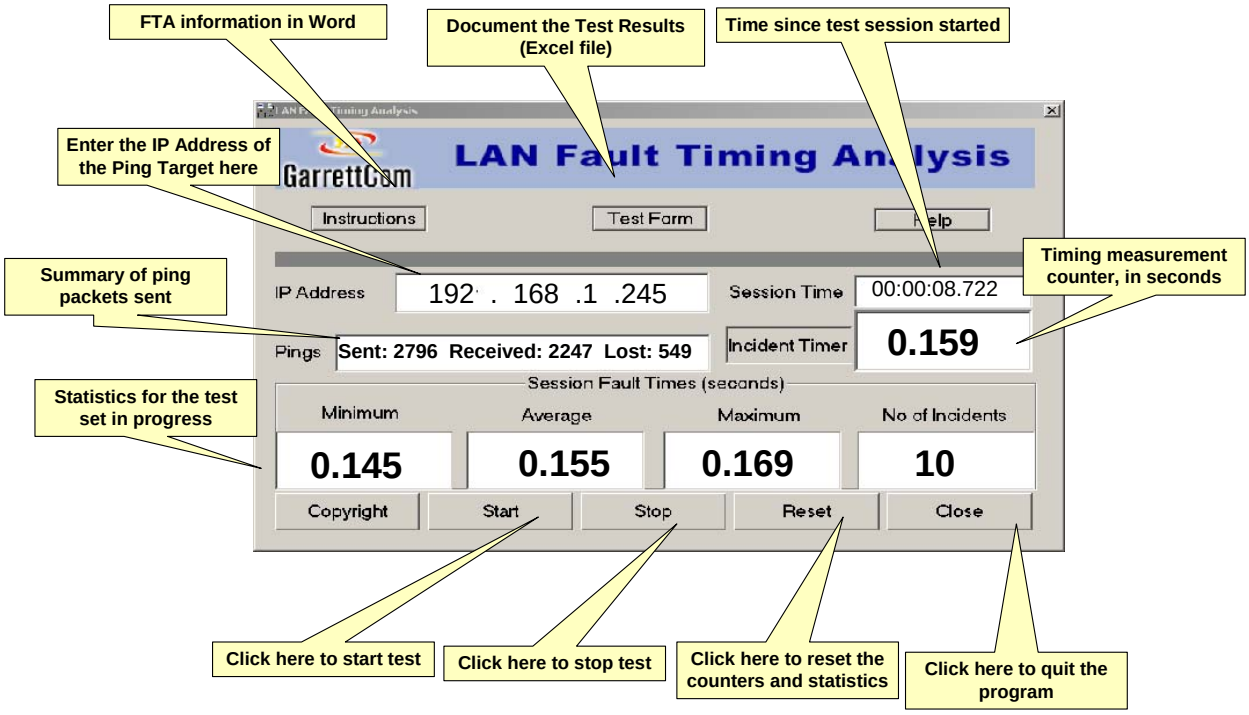

Figure 1. Screen shot of FTA in use.

The FTA program should be run on a dedicated PC. The rate at which packets are sent may be gated by the speed of the PC and its operating system. It is recommended that all unnecessary windows be closed during FTA tests, especially processor-intensive ones such as media players that could affect performance or create conflicts running the time-sensitive FTA software.

Setting up the Test LAN

FTA can measure self-healing LAN recovery times in both ring and mesh topologies. In setting up the test, it is only necessary to place the FTA computer issuing the timing pings at one point, and the ping target at another point such that the ping traffic between them traverses a normal path in which the fault condition to be measured will occur. The LAN must have a normal traffic path, and also a redundant traffic path available for use when there is a fault. An example of such a LAN set-up is in Figure 2, below.

During the test, the user will create the fault to be measured (e.g., by disconnecting a cable or pulling the power plug on a hub or switch). It is usually quite simple to connect the FTA computer and the ping target into appropriate connection points. Typically the FTA computer will be placed at a location simulating a Control Room or data collection station, and the ping target will be placed at a location simulating a field device such as an IED or PLC.

While complex mesh structures are interesting, ring structures are more common in industrial LANs. Rings are simpler in operation and also easier to understand for illustrating FTA testing principles. The diagram below, typical of simple ring topologies, shows a ring consisting of a Control Room ring-manager switch (labeled as Magnum 6K25) and four field switches (labeled as Magnum 6K16Vs) connected by fiber cable in a ring. The FTA computer is connected into a port of the Control Room switch, and the ping target (the IP node in a managed switch in this example) is connected into a port of a field switch. With all of the equipment turned on and operating normally, there is a normal traffic path from the FTA computer through the Control Room switch and through the left part of the ring to the ping target. When this traffic path is severed, fault recovery action will enable use of the traffic path from the Control Room switch to the ping target through the right side of the ring.

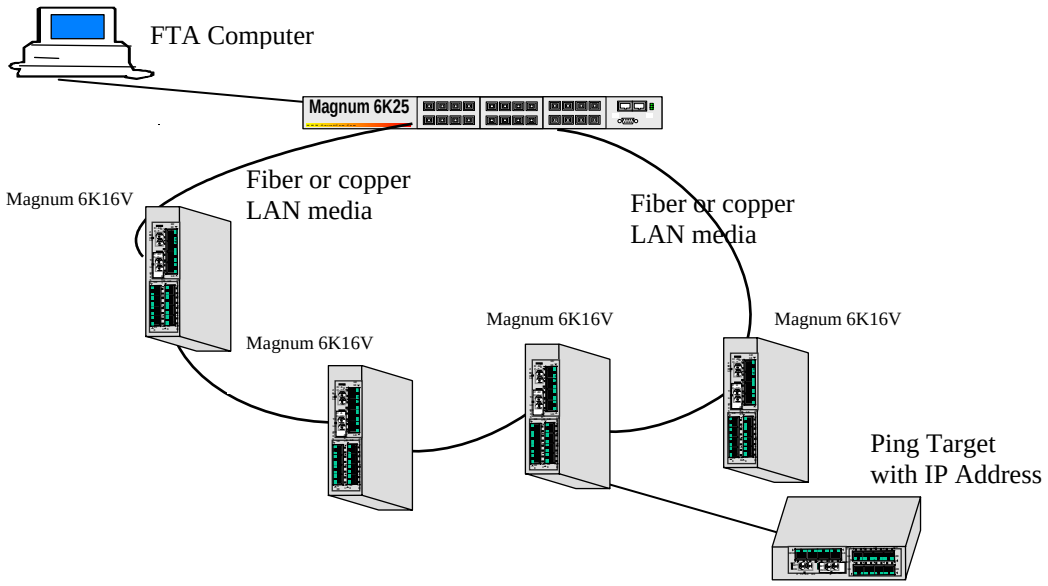

Figure 2. A typical self-healing ring set-up for testing fault recovery times using FTA.

FTA Operation

Once set-up is completed, the IP address of the target device must be entered into FTA. When the FTA program is started, FTA sends out ICMP ping request packets to the target at timed intervals and verifies that ICMP ping replies are received in a timely manner. When the operator creates a purposeful fault condition at the point of interest, FTA measures the time duration of the fault. The fault-recovered LAN operation will to carry some or all of the LAN traffic through a back-up path.

After a fault recovery, a typical test sequence is to cure the fault and restore normal LAN operation, with traffic passing through the original path as it was before the fault was introduced. FTA measures the traffic interrupt time of the restoration process as a timed event. Fault restoration is typically a mirror-image of a fault recovery, and fault restoration times are typically similar to fault recovery times. (In real life, of course, a fault is an unexpected event while a fault restoration is planned and managed).

Check-list for FTA Test

- Ping the target IP addresses as a test to verify that it is connected and working.
- Enter the target IP address in the FTA dialog box.
- Wait about 10 seconds for the FTA software and the LAN test set-up to stabilize, and confirm that the sending and receiving of ping packets is occurring smoothly.
- When the set-up has stabilized, it is possible to create a fault and begin to take data.
- While testing, maintain a minimum of a 10 second interval between a break (a fault created by the cable plug being taken out) and the restore (the cable plug inserted again) in order for the LAN elements to stabilize properly.
- If testing a ring, only the forwarding side of the ring should be broken to create a fault. If using STP or RSTP or S-Ring™ technology from GarrettCom, by STP rule, the ports lower in number will be forwarding and the higher ports will be blocking. (e.g., if the ring manager switch ports #1 and #3 are connected to the ring, then port #1 is forwarding and port #3 will be blocking for normal operation).
- If testing a mesh, make sure traffic is passing through the segment where the fault is created.

Testing Tips

Experience shows that fault recovery timings in self-healing LANs are highly variable. Several measurements of the same set-up, taken during a testing session, will yield different data. The number of variables in a network in operation is great, and the measured results of fault recoveries in the same set-up will vary accordingly. A single test reading, no matter how carefully done, is almost meaningless. For this reason, at least ten tests in a set should be taken to measure one condition, with the result simplistically characterized as the average time. Multiple set-ups are required in order to quantify fault recovery characteristics of a LAN with even moderate complexity.

When measuring sub-second or two-to-three-second recovery timings, expect the range of variation in timings to be 50% to 100%. When measuring tens of seconds (as with standard STP), expect the variation to be several seconds. Be prepared to observe a few spurious readings and suspiciously high or low test numbers during a test session, which may have to be discarded.

Measuring faults using pings on heavily loaded networks can be tricky. When the traffic is heaviest, some packets will inevitably be dropped. The way normal LANs prioritize traffic, ping packets are among those dropped first in a crunch. Also, for ping packets there is no re-try mechanism, so once a ping response is dropped it is gone forever. If you are using FTA on a LAN with maximum traffic, your measurement may be affected. Fortunately, one or two missed pings will appear as very short fault times (such as 5 milliseconds) and will be easily spotted as spurious readings in an FTA test session.

Finally, there is an important distinction between path recovery and fault recovery. Path recovery is fairly straightforward. It is defined as the LAN operating state where a new node can come on to the LAN and find a working path for traffic that enables it to use the LAN to communicate with another new node. Fault recovery, on the other hand, is defined as the LAN operating state where all existing nodes that previously communicated using the LAN elements can communicate again. The difference is that self-announcing ping targets can create a path for themselves through intervening switch-address-retention buffers, even though the rest of the devices have not and are prevented from communicating by the same address buffers. FTA measures fault recovery, and cannot distinguish between path recovery or fault recovery. For this reason, it is desirable to have complete documentation of a test to aid later analysis of the data. In case of doubt about a device being self-announcing and at what frequency, read the applicable device manual or use a sniffer to detect any announcement packets sent out.

Conclusions

Self-healing networks can be rings or meshes or combinations, and different vendors offer and/or support different topologies. Actual timing data on self-healing performance of a particular LAN configuration of interest is typically not available. Users want to know what their experience will be. FTA measures end-to-end time response of actual self-healing LANs in operation, with the set-up and test conditions and the device connection choices all controlled by the FTA test person.

With FTA, the performance of self-healing LANs can easily be measured during test sessions of an hour or two, and the resulting data recorded and catalogued for reference. FTA test data improves general understanding of the self-healing LAN phenomena and available topology choices, and leads to greater confidence in deploying such LANs in mission-critical applications.

For More Information

1. **IEEE 802.1 standards documents, 802.1D (STP) and 802.1W (RSTP)**
   http://standards.ieee.org/getieee802/802.1.html

2. **Ethernet Redundancy with Standards, a white paper**
   http://garrettcom.com/techsupport/papers/redundancy.pdf

3. **Link-Loss-Learn for Switches in Rings, a Technical Brief**
   http://www.garrettcom.com/techsupport/papers/lll_techbr.pdf

4. **S-Ring Adds Speed to Spanning Tree, a Technical Brief**
   http://www.garrettcom.com/techsupport/papers/sring_techbr.pdf

5. **Magnum MNS-6K User Documentation, Section 13 (STP) and 14 (RSTP)**
   http://www.garrettcom.com/techsupport/software/userguides/6k25_softugrel3.pdf

6. **Redundancy in Industrial Ethernet LANs, a web site navigation page**
   http://www.garrettcom.com/redundancy.htm

Fault Timing Analysis Test Report

GarrettCom™
Ethernet at Its Best℠

Test name: _____

Test By: _____ at _____ Date: _____

Redundancy Software used: _____ Tester's Contact Info, ph: _____

Release # _____ email: _____

MNS-6K or other net mgt software used: _____ Rev # : _____

Media types are: (check all that apply) Fiber___ RJ-45___ Speed: 10__ 100___ 1G___

Describe the ring or mesh members (mfg, model, etc): _____

Number of Ethernet units in the ring / mesh: Hubs: _____ Switches: _____ Other: _____

Describe nodes attached to the LAN, type, quantity, special features: _____

Distances of media: _____

Topology, brief description: _____

Ring, mesh, or combination: _____

Other products in the LAN, h/w and s/w: _____

FTA Software running on: (describe the PC) _____

Operating System on the FTA computer: _____

Ping Target is: _____

Notes: (Tester's comments, other ring activity noted during the test, where was fault introduced, etc.)

Does GarrettCom have permission to share this information with 3rd parties? Yes ____ or No ____

Results: Fault Times (in seconds, with 3 decimal places)

Minimum	Average	Maximum	No. of incidents

Please email your results to fta@garrettcom.com, or fax this completed form to 510-438-9072 Attn: FTA

www . GarrettCom . com