

Troubleshooting LANs with Wireshark Packet Capture and Expert Analysis

Introduction

This application note is one in a series of local area network (LAN) troubleshooting papers from JDSU Communications Test and Measurement. Troubleshooting LAN issues covers a wide array of network problems and diagnostic scenarios that may include:

- evaluating network utilization over the course of a business day by link, virtual LAN (VLAN), or subnet
- detecting excessive broadcast or multi-case traffic
- finding “bandwidth hogs”
- understanding what protocols are present on the network (and determining whether they should be)
- identifying the “top-talkers” on the link—the IP devices that are consuming the most capacity
- experiencing application performance issues (slow web server response time or intermittent unavailability of an e-mail server).

Before network troubleshooting can begin, one must have a clear understanding of network test access. Testing tools used for network analysis and troubleshooting scenarios must be able to monitor the network traffic being tested. The most common means for monitoring a network is using the built-in port mirroring capabilities of a network device, such as the switch/router or to install a special “tap” device between the devices being analyzed, such as those between an application server and database server. Figures 1 and 2 show each test access mode for analyzing traffic between two servers.

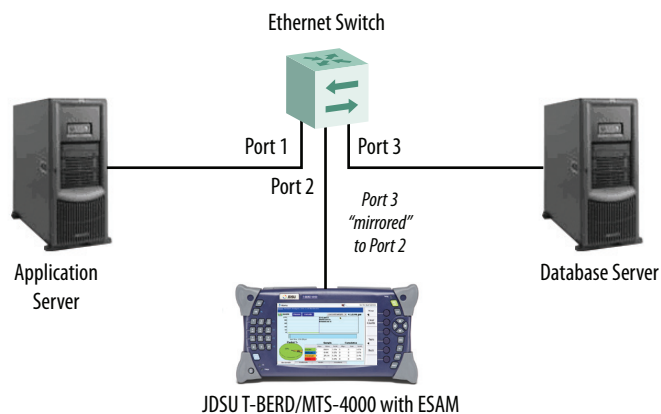


Figure 1: Test access via port mirroring

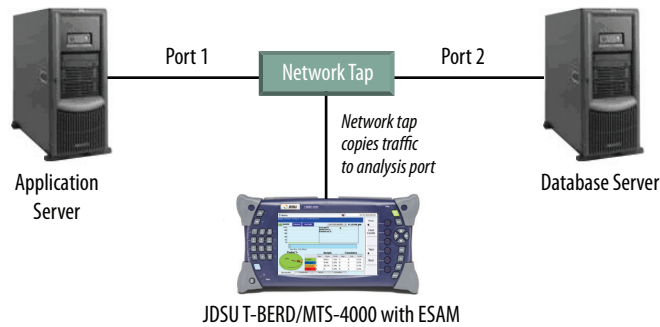


Figure 2: Test access with a network tap

In the port mirroring test access mode shown in Figure 1, the T-BERD®/MTS-4000 Enterprise Services Application Module (ESAM) is connected directly to a spare switch port (10M/100M/1000M) that is configured to copy all traffic to and from Port 3 (database traffic) to Port 2 (the test access port). Because a mirror port can copy traffic from both directions out to a single port, it will drop frames if full-duplex link utilization exceeds 50 percent. Aggregating taps perform similarly as they funnel both directions of traffic out of a single port. Likewise, if the full-duplex link utilization exceeds 50 percent it also drops frames. Some aggregating taps have internal buffers that allow them to compensate for bursts above 50 percent; however, this can result in incorrect timestamps when performing captures.

As a general rule, port mirroring is the preferred approach as there are generally spare switch ports and no interruption to production traffic.

The following table provides a summary of each test access mode:

Item	Port Mirroring	Aggregating Network Tap
Disruptive to network operation	No. A port mirror command does not interrupt normal production traffic.	Yes. Must install network taps during off-hours or as part of the production installation on critical network links.
Handle full line rate traffic	Handles up to 50-percent traffic utilization before dropping packets. Port mirroring may not be able to “keep up” on busy network links.	Handles traffic up to 50-percent utilization before dropping packets. Aggregating taps with buffers can compensate for bursts above 50 percent. A good network tap will not drop any production traffic, but may drop duplex traffic (on heavily loaded links) because the duplex traffic is combined into a single test access port.
Pass Layer 1 and Layer 2 Errors	No.	Depends on the tap. Some pass errors, some do not.
Require network device administrative privileges	Yes. Console access to the network device is required to enable port mirroring.	No.
Cost	Usually free, because most switches have a spare port.	Reputable 1000Base-T taps can cost \$1000+.

The following sections summarize a practical approach to conducting network analysis using the JDSU ESAM for the T-BERD/MTS-4000.

Network Analysis Workflow

There is no single method for analyzing network issues, and yet there are some best practices that experts use in their day-to-day troubleshooting activities. Figure 3 shows the JDSU-proposed network analysis workflow.

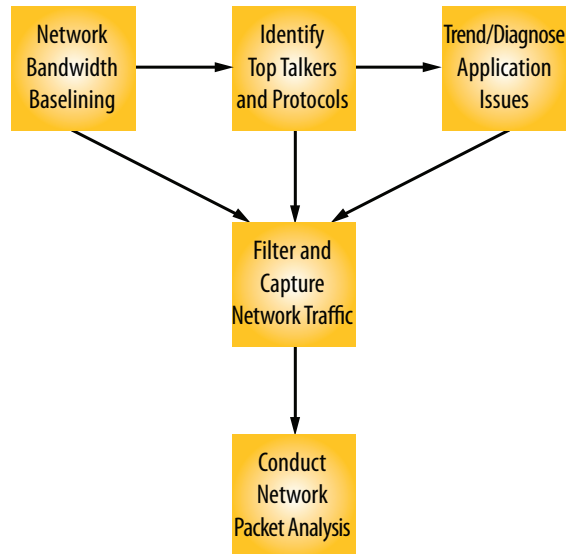


Figure 3: Best practices Network Analysis workflow

A second application note in this series, Troubleshooting LANs with Network Statistics Analysis, covers the details of the first three network troubleshooting steps (Network Bandwidth Baseline, Identifying Top Talkers and Protocols, and Trend/Diagnose Application Issues).

This application note covers these two steps:

1. **Filter and Capture Network Traffic:** At any of the steps listed above, it is important to conduct context-sensitive packet captures. For example, the ability to identify an Internet Protocol (IP) Top Talker and simply select it as a capture filter. The JDSU ESAM provides a simple capture filter user interface for the more common filter scenarios and also provides advanced deep packet inspection (DPI) filters that can search within the payload of packets.
2. **Network Packet Analysis:** The JDSU ESAM conducts analysis of the capture files directly using the popular open source software Wireshark. The ability of the test tool to perform expert analysis and diagnose common network problems within the packet file is also essential.

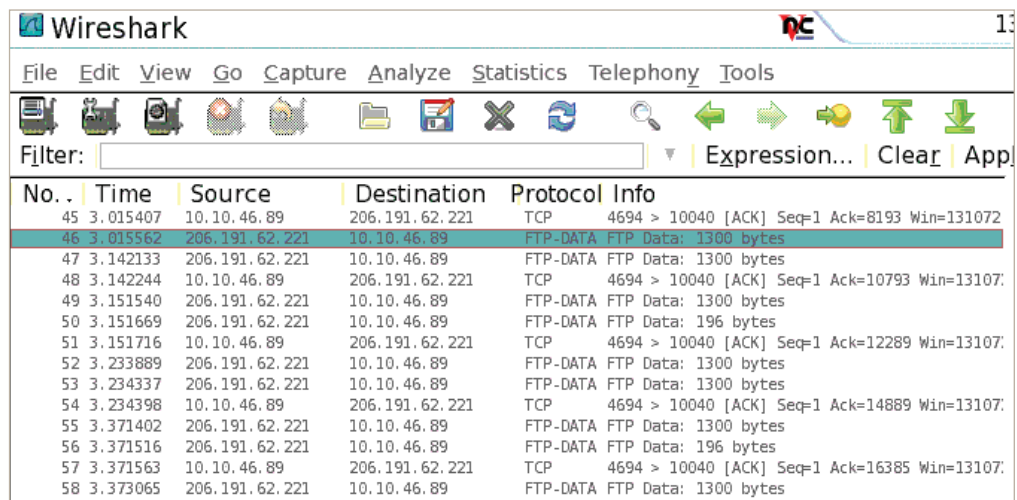
The following sections describe the details of packet filtering, capture, and analysis both with the Wireshark tool and the JDSU packet capture expert test feature J-Mentor.

Filtering and Packet Capture

The packet capture device must be able to capture packets at full Gigabit Ethernet (GigE) line speed. Wireshark-based captures use the network interface card (NIC) of the personal computer (PC) to perform the packet captures. The average-performing PC will drop packets even at line rates of 100 Mbps. Only very powerful (and expensive) workstations or servers with high performance NICs can capture at Gigabit Ethernet (GigE) wirespeed.

While PCs with Wireshark are sufficient for casual network “sniffing,” having a test tool that can capture all frames (regardless of size) at GigE speed is essential for performing accurate problem analysis and diagnosis.

The ESAM can capture up to 1 GB of packets at full line rates up to 1 GigE and store them natively in industry-standard packet capture (pcap) format. Wireshark is used to display and decode the packet captures directly on the ESAM user interface. Figure 4 shows an example of Wireshark running on the ESAM user interface.



No.	Time	Source	Destination	Protocol	Info
45	3.015407	10.10.46.89	206.191.62.221	TCP	4694 > 10040 [ACK] Seq=1 Ack=8193 Win=131072
46	3.015562	206.191.62.221	10.10.46.89	FTP-DATA	FTP Data: 1300 bytes
47	3.142133	206.191.62.221	10.10.46.89	FTP-DATA	FTP Data: 1300 bytes
48	3.142244	10.10.46.89	206.191.62.221	TCP	4694 > 10040 [ACK] Seq=1 Ack=10793 Win=13107
49	3.151540	206.191.62.221	10.10.46.89	FTP-DATA	FTP Data: 1300 bytes
50	3.151669	206.191.62.221	10.10.46.89	FTP-DATA	FTP Data: 196 bytes
51	3.151716	10.10.46.89	206.191.62.221	TCP	4694 > 10040 [ACK] Seq=1 Ack=12289 Win=13107
52	3.233889	206.191.62.221	10.10.46.89	FTP-DATA	FTP Data: 1300 bytes
53	3.234337	206.191.62.221	10.10.46.89	FTP-DATA	FTP Data: 1300 bytes
54	3.234398	10.10.46.89	206.191.62.221	TCP	4694 > 10040 [ACK] Seq=1 Ack=14889 Win=13107
55	3.371402	206.191.62.221	10.10.46.89	FTP-DATA	FTP Data: 1300 bytes
56	3.371516	206.191.62.221	10.10.46.89	FTP-DATA	FTP Data: 196 bytes
57	3.371563	10.10.46.89	206.191.62.221	TCP	4694 > 10040 [ACK] Seq=1 Ack=16385 Win=13107
58	3.373065	206.191.62.221	10.10.46.89	FTP-DATA	FTP Data: 1300 bytes

Figure 4: Wireshark running directly on the JDSU ESAM User Interface

Often, traffic volume is very high and requires filtering before packet capture, commonly referred to as a pre-capture filter. The most common form of pre-capture filter is:

IP Address (or Address Pair) AND Protocol (for example, hypertext transfer protocol [HTTP] and simple mail transfer protocol [SMTP]).

An IP Address Pair is commonly referred to as an IP Conversation. Figure 5 is an example of a Small-Medium Business (SMB) enterprise with 100 users accessing a corporate database.

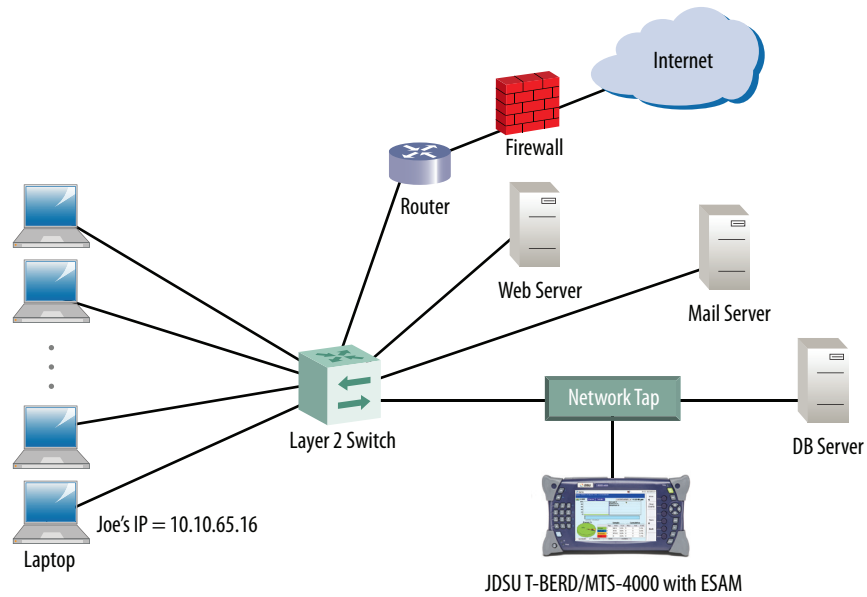


Figure 5: SMB with Corporate Database Server

Although simplified, Figure 5 illustrates 100 local users accessing the database, mail, and web servers within the corporate data center. Additionally, remote branch offices and telecommuters can access these resources via virtual private network (VPN) access and the Internet.

As a case study, let us introduce a hypothetical scenario where Joe experiences poor performance when accessing the database server. Because the database server link is critical, the network team at Joe's company has a dedicated network tap placed in line with the database server. All traffic to and from the database server is copied to the packet analyzer that, in this case, is the JDSU ESAM.

Without pre-capture filters, the JDSU ESAM captures all of the network traffic from all users, which is not a practical solution. Figure 6 illustrates the configuration of a pre-capture filter using the ESAM.




Figure 6: Configure Pre-capture Filter for IP = 10.10.65.16

Note that in addition to specifying the 10.10.65.16 IP address, both directions of communication will be captured.

Often a single server will host multiple applications. Looking back at Figure 5, imagine that all three functions (database, web, and e-mail servers) now reside on one physical server. Setting the filter to Joe's IP address would capture the traffic between Joe and all three applications.

The solution to this problem is to specify an IP address and protocol filter, as illustrated in Figure 7. Note that “TNS” is the transport layer protocol over which Oracle structured query language (SQL) database transactions are carried.

<input type="checkbox"/> All Hosts	<input checked="" type="checkbox"/> Single Host	<input type="checkbox"/> Conversation
IP Address	10.10.65.16	
Capture File Name	JDSU_2010_03_11_12_54_41_100M	
Filter By Protocol	TNS	

Figure 7: Configure Pre-capture Filter for IP = 10.10.65.16 and TNS

For more complex situations, IP and Protocol-based filters may be inadequate for capturing the desired traffic. Frequently, a database server is connected to a front-end web server or application server and the user traffic to the database server cannot be detected by IP address. The database server sees one IP address (for example, for the application server) and all of the user traffic is contained within a pool of transmission control protocol (TCP) connections. This can be thought of as “trunk” communication between the application server and the database server (all user traffic from application server IP to the database server IP).

In this scenario, DPI techniques are the appropriate means to troubleshoot this problem. DPI examines payload within the packet and the JDSU ESAM is equipped with advanced DPI technology.

For this example, the ESAM must first be configured with a custom DPI filter. Assume that Joe’s Oracle user ID is `joe_knap_2301`. This ID can usually be found in the payload of the packets sent to the database server.

The complex DPI based filter will search within the packet payload to detect and capture only the packets with Joe’s ID. Figure 8 illustrates the ease of configuring this type of complex DPI-based filter on the ESAM.

<input checked="" type="checkbox"/> TCP	<input type="checkbox"/> UDP	<input type="checkbox"/> LLC	<input type="checkbox"/> IPV 4	<input type="checkbox"/> IPV 6
Item	Value			
Filter Name				
Filter By IP Address	None			
Filter By MAC	None			
Filter By Port	None			
Filter by VLAN	None			
Payload 1 Setting	Search Payload			
Payload 1 String	joe_knap_2301			
Payload 2 Setting	Ignore			

Figure 8: Configuring a DPI Filter for “joe_knap_2301”

Note that no IP or port settings are configured for this address as they are not relevant; the ESAM will conduct DPI to detect Joe’s database communications and only capture those packets. After launching Wireshark directly on the ESAM user interface, the user can troubleshoot the database issue.

Expert Analysis

Wireshark packet captures provide a wealth of information, but it is very difficult for the average user to diagnose network problems. It requires significant network expertise is required to comb through the packet capture file and detect issues such as TCP retransmissions, Internet Control Message Protocol (ICMP) events, and others.

Expert packet analysis is essential to analyzing and diagnosing problems in poorly performing networks and applications. The JDSU ESAM provides J-Mentor, which is a packet capture expert that can analyze packet capture files and diagnose common network problems.

In this example, a 2 MB file download took over 60 seconds and the cause of this poor throughput must be diagnosed. Figure 9 shows screenshots of this capture file after opening in Wireshark and JDSU J-Mentor.

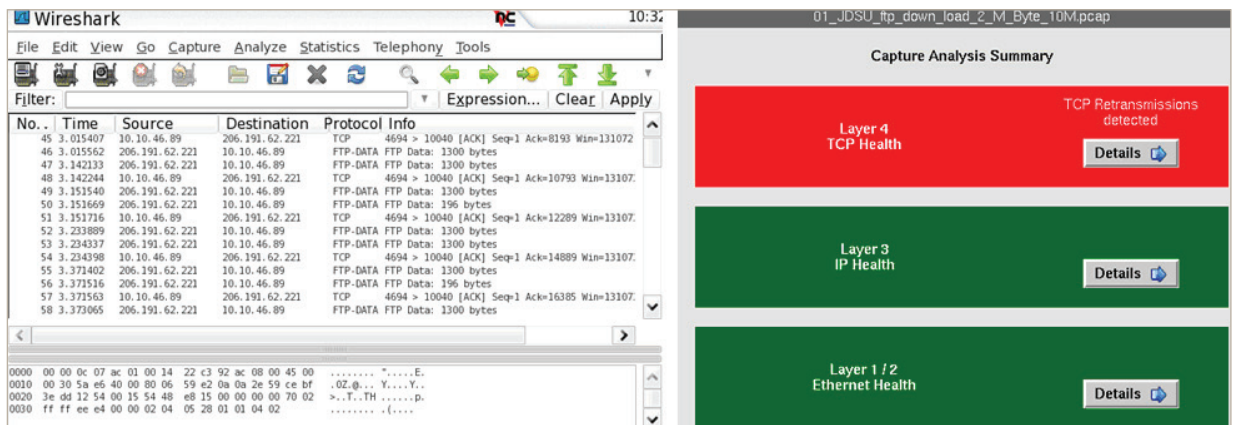


Figure 9: Wireshark Decodes versus J-Mentor Diagnosis of FTP Download

As the Wireshark screen on the left illustrates, users require expertise to navigate through the packet decodes and advanced analysis menu options. However, the J-Mentor screen on the right illustrates the ease with which network technicians can quickly isolate the problem to the problem network layer.

In this example, the network issue occurred at Layer 4 (TCP) and by clicking on the Details button, Figure 10 is presented.




Figure 10: Drilling into Layer 4 (TCP) Issues

The results show that a total of 24 retransmissions occurred. Next the user clicks on the TCP Retransmissions tab to further isolate the problem as Figure 11 illustrates.

TCP Retransmissions		
Source IP Address	Destination IP Address	Retransmissions
206.191.62.221	10.10.46.89	18
172.17.8.66	207.46.249.61	4
207.46.249.61	172.17.8.66	2

Figure 11: Isolating TCP Retransmissions to Source IP

As the figure shows, the FTP client (206.191.62.221) caused 18 of the retransmissions. This provides simple diagnostic information to the network troubleshooter and points to the problem source for further isolation and troubleshooting (determine if the host IP has half-duplex port issues, bad cabling, or a related problem.)

Half-duplex port issues also remain a cause of considerable headaches. Most networks list port settings: J-Mentor automatically detects these messages and provides a list of source MAC addresses that list half-duplex settings during the packet capture time interval, see Figure 12.

Half Duplex Ports			
Time (secs)	Source MAC Address	Platform	Port
162334362.284	c2:01:73:fe:00:00	Cisco 3725	FastEthernet0/0

Information
Cisco Discovery Protocol (CDP) messages were detected on this network and the table lists those MAC addresses and ports which advertised Half Duplex settings.

Recommendation
Locate the device with the source MAC address(es) and port(s) listed in the table and ensure that duplex settings are set to "full" and not "auto". It is not uncommon for a host to be set as "auto" and network device to be set as "auto", and the link incorrectly negotiates to half-duplex.

Figure 12: Isolating Source MAC using the Half-Duplex Setting

It is common to search for “bandwidth hogs” as the source of potential issues in poorly performing networks. Therefore, as Figure 13 shows, J-Mentor provides a listing of the Top Talkers detected within the packet capture file along with the number of bytes and frames for each talker.

IP Conversations							
Source IP Address	Destination IP Address	Frames S <- D	Bytes S <- D	Frames S -> D	Bytes S -> D	Total Frames	Total Bytes
172.17.8.66	161.58.73.170	71	7249	104	19848	175	27097
207.46.248.61	172.17.8.66	28	4814	17	8167	45	13101
172.17.8.66	4.2.2.1	4	639	4	311	8	950
208.38.50.34	172.17.8.66	4	555	2	369	6	924

Figure 13: Simple Display of IP Top Talkers

Conclusion

Troubleshooting LAN problems effectively requires a packet capture tool capable of capturing packets at full GigE line speed. The average-performing PC will drop packets even at line rates of 100 Mbps. While PCs with Wireshark are sufficient for casual network sniffing, having a test tool that can capture all frames (regardless of size) at GigE speed is essential for performing accurate problem analysis and diagnosis.

The JDSU ESAM for the T-BERD/MTS-4000 can capture up to 1 GB of network packets and store them natively in industry-standard pcap format. The ESAM provides a simple capture filter user interface for the more common filter scenarios and also provides advanced DPI filters that can search within the payload of packets. Wireshark runs natively on the ESAM display and J-Mentor provides expert packet capture diagnostics to provide best-practice troubleshooting for the less experienced network troubleshooter.

The ESAM provides a workflow-based interface that “walks” the user through the best practices approach toward solving a multitude of network problems. Figure 14 is the JDSU T-BERD/MTS-4000 platform with ESAM interface (and an optional fiber scope), and Figure 15 shows the workflow-based user interface.



Figure 14: JDSU T-BERD/MTS-4000 platform with the ESAM



Figure 15: Workflow-based user interface of the ESAM

The JDSU ESAM for the T-BERD/MTS-4000 provides comprehensive LAN testing capabilities with these features:

- Layer 1-7 protocol capture and expert analysis
- network connectivity
- network discovery
- a full range of physical media tests
- a workflow-based user interface
- a modular platform with many options:
 - Voice over IP (VoIP) phone emulation
 - optical power meter/visual fault locator
 - fiber inspection probe with automated pass/fail
 - Wireless fidelity (WiFi) testing
 - OTDR modules

Through its workflow-based intuitive user interface, the ESAM provides physical media tests including speed-certification of electrical Ethernet cabling, network connectivity tests, discovery, wirespeed deep-packet statistics, and wirespeed protocol capture and expert analysis using unique, in-depth JDSU J-Mentor capabilities. In addition, the ESAM is part of the modular JDSU T-BERD/MTS-4000 platform allowing additional options that include VoIP emulation, WiFi testing, IP video testing, optical power meters (OPMs), visual fault locators (VFLs), digital fiber inspection probes, and optical time domain reflectometers (OTDRs). Test connectivity can be obtained either electrically via a 10/100/1000 RJ45 Ethernet jack or via a small form-factor pluggable (SFP) for optical Ethernet.

Test & Measurement Regional Sales

NORTH AMERICA TEL: 1 866 228 3762 FAX: +1 301 353 9216	LATIN AMERICA TEL: +1 954 688 5660 FAX: +1 954 345 4668	ASIA PACIFIC TEL: +852 2892 0990 FAX: +852 2892 0770	EMEA TEL: +49 7121 86 2222 FAX: +49 7121 86 1222	WEBSITE: www.jdsu.com/esam
---	--	---	---	--