

Frontline LAN Troubleshooting Guide




Table of contents

Abstract	2
Introduction to troubleshooting	3
The best method	3
The process	4
Eight steps to successful troubleshooting	8
Troubleshooting the physical layer	16
Troubleshooting copper media	16
Copper cable tests	21
Troubleshooting fiber optic media	36
Fiber optic cable tests	41
Troubleshooting the network layer	47
Troubleshooting common network user complaints	47
Complaint: Can't connect	48
Complaint: Connection drops	53
Complaint: Network is slow	57
Troubleshooting switches	62
Typical switched network problems	62
Isolating the problem	64
Switch troubleshooting techniques	65
Method 1: Access the switch console	67
Method 2: Connect to an unused port	70
Method 3: Configure a mirror or span port	77
Method 4: Connect to a tagged or trunk port	84
Method 5: Insert a hub into the link	86
Method 6: Place the tester in series	90
Method 7: Place a Tap inline on a link	91
Method 8: Use SNMP-based management	98
Method 9: Use flow technology	106
Method 10: Set up a syslog server	109
Method 11: Use the server (host) resources	110
Method 12: Use a combination of the methods	112
Conclusion	113

Abstract

Local area networks are integral to the operation of many businesses today. Network engineers and network technicians have taken on the vital role of keeping these business-critical networks up and running. This guide provides these frontline network troubleshooters with practical advice on how to maintain LANs and solve common problems.

A local area network (LAN) is comprised of many elements: printers, monitors, PCs, IP phones, servers, storage hardware, networking equipment, security software, network applications, enterprise applications, office productivity applications, and more. In this guide, we will focus on layers 1 and 2 – the physical cable plant and switches. Network cabling and switches are the foundation of today's local area networks.

This guide begins with an introduction to LAN troubleshooting. We will review the troubleshooting process and the eight steps to successful troubleshooting. Next, will be information on troubleshooting physical layer problems. We will cover twisted-pair copper and fiber optic media. Advice on troubleshooting common network user complaints will follow. Common complaints include user connection issues and slow networks. The guide concludes with an in-depth discussion of troubleshooting switches. We will describe many switch troubleshooting methods. Frontline LAN troubleshooters who learn how to apply these methods will be able to solve network problems fast.

Introduction to troubleshooting

The best method

There is no “best” method for troubleshooting, just like there is no single tool that solves all of your networking problems. We will describe several different approaches to troubleshooting. Humor offers a way to illustrate the problem presented by troubleshooting.

The first example is the saying: “when the only tool you have is a hammer, everything begins to look like a nail.” This can be interpreted in many ways. One way to interpret this is that the user was able to bludgeon the network hard and long enough that the marginal or failed element was eliminated in some way, and the user concluded that bludgeoning is a suitable substitute for troubleshooting. Another way to interpret this is to describe how a person who has become very proficient at a particular network diagnostic product is able to apply that product to those situations for which it is technically unsuited. This is not because the product is capable of detecting certain classes of problems, but because the user can interpret the test results based on experience and knowledge and arrive at a conclusion that is close to correct.

The second example is the joke about how several blindfolded people were tasked with describing an elephant. All disagreed with the others because the only information they had was what each experienced directly. The person touching the trunk described how an elephant was like a large snake. The person touching the leg described how an elephant was like a tree. Descriptions of the tail and the flank of the elephant produced further contradictory descriptions. Each description was accompanied by the emphatic assurances of the person providing the description that their description was correct, because that is the only first-hand ex-

perience that person had of an elephant. To add confusion, they all agreed on what the skin felt like.




Figure 1: Personal experience often makes it difficult to accept another perspective or opinion.

If the person troubleshooting does not have a working knowledge of the technology, adequate information gathered from multiple points or information sources, or is lacking experience for a broader interpretation, then incorrect assumptions and conclusions are made. The accuracy and speed of troubleshooting depends on the knowledge, skill, and experience of each technician involved, and the tools at their disposal. It sometimes requires interpretation by an uninvolved third party who is able to provide an objective opinion.

The process

The key to successful troubleshooting is for the technician to know how the network functions under normal conditions. This enables the technician to recognize abnormal operation quickly.

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 and 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 or operator error.

The foundation of this insight is usually gained through formal training. However, the true troubleshooting master learns through trial and error, comparing notes with others, and discovering tried-and-true methods that are not often taught in school. Following a good formula or process for troubleshooting that includes careful documentation of your actions and your hypothesis for what might be causing each problem can help shorten your learning curve and at the same time shorten the time required to solve network problems.

Two extreme approaches to troubleshooting almost always result in disappointment, delay, or failure. On one extreme is the theorist, or *rocket scientist* approach. On the other is the practical or *caveman* approach. Since both of these approaches are extremes, the better approach is somewhere in the middle using elements of both.

The *rocket scientist* analyzes and re-analyzes the situation until the exact cause at the root of the problem has been identified and corrected with surgical precision. This sometimes requires taking a high-end protocol analyzer and collecting a huge sample (megabytes) of the network traffic while the problem is present and inspecting it in minute detail. While this process is reliable, few companies can afford to have their networks down for the hours – or days – it can take for this exhaustive analysis.

The *caveman's* first instinct is to start swapping cards, cables, hardware and software until miraculously the network begins operating again. This does not mean it is working properly, just that it is operating. Unfortunately, the troubleshooting section in some manuals actually recommends caveman-style procedures as a way to avoid providing information that is more technical. While this approach may achieve a change in symptoms faster, this

approach is not very reliable and the root cause of the problem may still be present. In fact, the parts used for swapping may include marginal or failed parts swapped out during prior troubleshooting episodes.

For the technician in search of a better way to troubleshoot, try the approach below. Once learned, the art of troubleshooting can be applied with very slight changes to almost any corrective situation. The process described below could be used to fix a lawn mower, a camera, or a software program.

Analyze the network as a whole rather than in a piecemeal fashion. One technician following a logical sequence will usually be more successful than a team of technicians, each with their own theories and methods for troubleshooting.

The logical technician asks the user questions, runs diagnostics, and thoroughly collects information. In a short time, he or she can analyze and evaluate the symptoms, zero-in on the root source of problems, make one adjustment or change one part, and cure the problem. The key is to simply isolate the smallest failing element and replace or reconfigure it. *Complete understanding of the cause of the failure is not required at this time.* The primary goal is to restore network operation rapidly. After the network is again running, further analysis may be undertaken – preferably in a lab environment.

There are many technicians with years of experience who have not yet mastered 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 then need to conduct tests to validate or eliminate what you think the problems could be. If adequate symptoms are known, perhaps the evaluation process is mental and does not involve the network or physical testing at all. Once you think you understand the problem, you must then verify it. At this stage, your efforts will be directed toward attempting to cause the problem to recur on demand.

Just as important: the logical technician always performs a checkout procedure on any repaired equipment or system, no matter how simple the repair. Far too often, the obvious problem is merely a symptom of another less-obvious problem, and until the source is eliminated, the situation will continue or reappear.

Once the problem has been solved, document and share the identifiable symptoms and the solution to that problem so that others do not have to reinvent what you have learned.

The last step is to provide feedback and training to the user. If they are informed of what action caused the problem or the nature of the problem and solution, then they either will avoid doing it again in the future or will be able to provide a more specific description for the next problem.

Eight steps to successful troubleshooting

1. Identify the exact issue or problem.
2. Recreate the problem if possible.
3. Localize and Isolate the cause.
4. Formulate a plan for solving the problem.
5. Implement the plan.
6. Test to verify that the problem has been resolved.
7. Document the problem and solution.
8. Provide feedback to the user.

Step 1. Identify the exact issue.

Defining the scope of the problem and deciding on the exact issue is important. Have the person who reported the problem explain how normal operation appears, and then demonstrate the perceived problem. If the reported issue is described as intermittent, instruct the user to contact you immediately if it ever happens again. It is very difficult to fix something that is clearly working just fine right now.

Do not discount what the user reports simply because it sounds implausible. The user does not have your knowledge of networking, and is probably describing the problem poorly. *Something* annoyed the user enough to contact you.

Note: *Has it ever worked? If the reported failure has never worked properly, then treat the situation as a new installation and not a troubleshooting event. The process and assumptions are completely different.*

Step 2. Recreate the problem.

Ask yourself if you understand the symptoms, and verify the reported problem yourself if possible. Problems are much easier to solve if they can be recreated on demand. Seeing the problem will allow you to observe error messages and various symptoms the user may not think important to relate, and may even provide the opportunity for you to collect network statistics during the event.

If the problem is intermittent, instruct the user what sort of symptoms are likely and provide a *written list* of what questions you are seeking answers to so the user can gather some of the information if you are unable to respond quickly enough to see it yourself. When possible, leave a diagnostic tool to gather information continuously. A protocol analyzer may be left gathering all traffic from the network and overwriting the buffer as it fills. Have the user halt its operation and/or store the current test results from other testers *immediately* upon rediscovering an intermittent problem.

Step 3. Localize and isolate the cause.

Once you have defined the problem, and recreated it if necessary, you should attempt to isolate that problem to a single device, connection, or software application. Reducing the scope of the problem in this way is where *divide-and-conquer* begins; the goal is to isolate the problem to the smallest element that could cause the problem. Test for and eliminate as many variables as possible. You may need to scan for a virus at this point.

Is there any normal function missing, or is there an abnormal response? Use the data gathered by your network monitoring tools to aid you in this process.

Determine whether *anything* was altered at that station or on the network just before the problem started. Often the user does not realize that changing something seemingly unrelated can cause problems on the network, such as rearranging the location of a portable heater or photocopier, or installing a new software application or adapter card. Do not discount the local environment when you are looking for change. Temperature changes (heat is often a problem), electrical use from adjacent spaces – including nearby businesses, time of day, and influences from electronic sources. Even the passage of an elevator, or use of a cordless phone, should be noted.

Can the problem be duplicated from another station, or using other software applications at the same station? Identify whether the problem is limited to one station, or one network resource such as a printer. Move one segment closer to the network resource and try again. If the problem goes away when you move closer to the network resource, test or replace the intervening infrastructure equipment.

If the problem affects an entire shared media segment, isolate the problem by reducing the variables to the fewest possible number. Try shortening the cable segment on a bus topology, or temporarily re-cabling a ring or star topology to create the smallest possible network for troubleshooting purposes. Try a different switch or hub. If the problem is on the same, shared media segment as the network resource, try turning off or disconnect all but two stations. Once those two are communicating, add more stations. If they are not communicating, check the physical layer possibilities such as the termination of the cable, the cable itself, or the specific ports used on the infrastructure equipment (hubs and switches).

If the problem can be isolated to a single station, try a different network adapter, a fresh copy of the network driver software (without using any of the network software or configuration files presently found on that station – delete them if necessary). Try accessing the network using a diagnostic tool from the existing network cable connection for that station. If the network connection seems intact, determine if only one application exhibits the problem. Try other applications from the same drive or file system. Compare configurations with a nearby but operational workstation. Try a fresh copy of the application software (again using none of the existing software or configuration files).

If only one user experiences the problem, check the network security and permissions for that user. Find out if any changes have been made to the network security that might affect this user. Has another user account been deleted that this user was made security equivalent to? Has this user been deleted from a security grouping within the network? Has an application been moved to a new location on the network? Have there been any changes to the system login script, or the user's login script? Compare this user's account with another user's account that is able to perform the desired task. Have the affected user log in and attempt the same task from a nearby station that is not experiencing the problem. Have the other user log in to the problem station and try the same task.

Step 4. Formulate a plan for solving the problem.

Once a single operation, application or connection is localized as the source of the problem, research and/or consider the possible solutions to the problem. Consider the possibility that some solutions to the problem at hand may introduce other problems.

Note 1: *To avoid unwanted repetition, and to make it possible to “back-out” any changes made should things get worse, be sure to carefully and completely document all actions taken during the problem resolution process. Copy all configuration files to a safe place before modifying them – especially on switches, routers, firewalls, and other key network infrastructure devices.*

Note 2: *It is advantageous to open a second terminal session into the switch or router where the commands required to reverse a configuration change are typed in and ready to execute prior to actually implementing the change in the first window. This is likely the fastest way to recover from changes that adversely affect your network.*

Step 5. Implement the plan.

Your actual solution to the problem may be replacing a network device, NIC, cable, or other physical component. If the problem is software, you may have to implement a software patch, reinstall the application or component or clean a virus infected file. If the problem is the user account, the user’s security settings or logon scripts may need to be adjusted.

For network hardware, it is most expedient to simply replace a part, and attempt to repair the part later. Another option is to change the connection to a spare port and cover or otherwise mark the suspect port. Remember than the goal is to restore full operation of the network as soon as possible.

Two avenues exist for solving software problems. The first option is to reinstall the problem software, eliminating possibly corrupted files and ensuring that all required files are present. This is an excellent way to ensure that the second option – reconfiguring

the software – works on the first try. Many applications allow for a software switch that tells the installation program to disregard any existing configuration files, which is a good way to avoid being misled by the error and duplicating it yet again. If this option is not evident, then it is often better to remove the application before reinstalling it.

If the problem is isolated to a single user account it is often faster to repeat the steps necessary to grant the user access to the problem application or operation as if the user had never been authorized before. By going through each of these steps in a logical order, you will probably locate the missing or incorrect element faster than by spot-checking. In some situations, it may be expedient to simply delete the whole account and start over.

Step 6. Test to verify that the problem has been resolved.

After you have implemented the solution, ensure that the entire problem has been resolved by having the user test for the problem again. Also, have the user quickly try several other normal operations with the equipment. It is not unheard of for a solution to one problem to cause other problems, and sometimes whatever was repaired turns out to be a symptom of another underlying problem.

Step 7. Document the problem and solution.

Documentation is useful for several reasons. First, documentation can be used for future reference to help you troubleshoot the same or similar problem. You can also use the documentation to prepare reports on common network problems for management and/or users, or to train new network users or members of the network support team.

Step 8. Provide feedback to the user.

There is often the temptation to fix the problem and leave. However, if a network user reported the problem they will appreciate knowing what happened. This will encourage them to report similar situations in the future, which will improve the performance of your network. Another reason for feedback is that if the user could have done something to correct or avoid the issue, it may reduce the number of future network problems.

A good working relationship between network support staff and the user community can significantly enhance your ability to keep the network running smoothly. Failure to take users seriously, or making unprofessional and condescending remarks can cause adversarial relations to develop, and can undermine your ability to do your job.

There is also a saying that 75% of fixing a problem is “fixing the user.” If the user does not agree that the problem has been taken to its conclusion (whether the problem has been corrected, or you have explained *to the user’s satisfaction* that a fix is impossible for the following technical, financial, or political reasons...), then you have not ended this support issue.

A place to start

As with everything else, do not assume that a short course and a book or two will make you into the networking equivalent of Sherlock Holmes. Take the time to learn one or two aspects of networking very well before seeking the next topic. Feel free to ask for help or guidance with everything else in the mean time. This approach will help you avoid making many of the common blunders.

The first suggested step in troubleshooting is to gather information. If you do not know what normal operation is like, and you do not know the technology used, and then it is difficult to gather information and symptoms about the current failure effectively.

Follow up on topics of interest from the first subject(s) of study, so that your knowledge expands from a central point. I hope that you will be working up the OSI Model as you progress. A significant number of senior networking specialists either have forgotten or never knew the basic operation of many elements of the network. Technology is changing very fast in this industry, and they have usually chosen to focus on the higher-layer aspects to the exclusion of developments in the lower layers. This causes them to make incorrect assumptions about some symptoms, and delays problem resolution accordingly. Since these people are often in positions where network architecture decisions are made a number of expensive upgrades have been purchased unnecessarily. Nobody knows it all, so ask for help when you are unsure. Consult multiple sources when the answer sounds too good to be true, or is questionable.

Similarly, each course or book offers insightful knowledge and experience in specific networking topics, but sometimes goes on to address topics that would be better left to other subject matter experts. One of the indications that you understand a topic or concept well enough is when you can identify the point where that happens, if it does.

Troubleshooting the physical layer

Troubleshooting copper media

General testing and installation issues

Most networks have converted from coax to Category 3, then to Category 5 or Category 5e, and soon to Category 6 or Category 6A UTP links. However, there is still a surprising amount of coax used for WAN and wireless, and in legacy network LAN segments where low bandwidth is still satisfactory. Fiber is approaching or at a price-point where it will rival Category 6A UTP in overall cost (material, installation, and network adapters). There are installation and maintenance issues that are different for all cable types. Below are several testing and troubleshooting issues for each cable type.

As mentioned, despite the conversion from legacy coax Ethernet to UTP implementations there is still a lot of coax used for WAN and wireless links. The coax type for thin Ethernet is 50 Ohm RG-58, while WAN links and 802.11 wireless antenna extension cables typically use 75 Ohm RG-59 and may experience virtually the same cable problems. Use of 93 Ohm RG-62 cable is no longer common in networking.

Types of UTP cable

The similarity between older standards and newer standards for a particular designation, such as Category 5, has created situations where a cable manufactured (and labeled) in compliance with an older version of a standard no longer meets the same designation in a new, similarly labeled standard. This is true for the TIA/EIA-568 standard as well as the ISO/IEC 11801 standard.

Category 5 or ISO/IEC Class D cables manufactured between 1995 and 1999 generally meet the requirements for TSB67. When TSB95

was published in 1999, most new cable was manufactured to meet those tighter test limits, but the actual cable labeling often did not change in a way that the average person would notice. The cables may have included a date too, either Category 5 (1999) or Category 5 (2000), for example. The same is true for Class D (1999) or Class D (2000). What changed was the marketing campaign from the manufacturers. The manufacturers began differentiating their cable from other manufacturer's products with fast sounding product names. When Addendum 5 to TIA/EIA-568-A was released, the same thing happened again. This time the cable label may have changed to Category 5e, Category 5 (2000), or Category 5 (2001). Plethoras of cables have names that were selected to suggest that they are perfect for Gigabit Ethernet and beyond. In fact, 1000BASE-T will run just fine on Category 5 cable that passes TSB95 test requirements. Category 5e is better performing cable than 1000BASE-T requires. Similarly, 10GBASE-T will run on cable that meets the test requirements for TIA/EIA TSB155 or ISO/IEC TR-24750. Category 6A and Class E_A links perform better than 10GBASE-T requires.

This labeling situation is important to the consumer in two very important ways. First, do not place too much faith in the cable labeling or cable product family names. Instead, rely upon the field cable tester results for performance to a selected cable grade. Sometimes a link does not pass the labeled performance level, other times the cable performs to a higher grade (when installed with excellent workmanship). Second, because the standards sometimes change without altering the name of the cable grade it is sometimes difficult to identify with which version of the standard the product actually complies. This did not have a particularly pronounced an effect on 100 MHz cables (Category 5

and Class D), but on the higher-speed cable systems the effect was alarming. The initial products from different manufacturers that were promoted as meeting the first drafts for Category 6 standards were fine when used as a complete end-to-end single-vendor system, but sometimes tested to a lower grade when mixed with Category 6 components from other vendors. You had to maintain the same vendor and product family throughout the entire link in order to achieve the expected performance. The early or pre-standard Category 6A and Class E_A links may well have similar results. Thus, if you are attempting to upgrade your cable plant to a higher level of performance it is necessary to retest each link in its final configuration to ensure that the link meets your expectations. Do not trust any labeling or marketing guarantee if you do not have a completely homogenous cable system installed at the same time from the same run of manufactured cable and connecting hardware. Even then, the installation workmanship may result in substandard performance. Test everything against the expected performance rating in its final *permanent link* and *channel link* configurations. The various link configurations of channel link, permanent link and the now obsolete basic link will be explained later. For now be aware that this largely translates to leaving the tested patch cables in place for the user, you cannot use one set of patch cables to test all of the links.

Naming of cables is an interesting subject too. The official designations are Category 5 or Category 5e. Attempts to promote Category 5E (upper case "E") as being better than Category 5e (lower case "e") are not supported by the standards, since this designation was created by marketing. The cable manufacturers also came out with cables labeled Category 6e in apparent anticipation that the standard would use that as the next designation, which it did not. The

TIA/EIA cable grade above Category 6 was chosen to be Category 6A (upper case “A”). There is no Category 7, other than cabling components for ISO/IEC 11801. Specifically, the connectors, punch down blocks, and so on are specified up to Category 7, but an assembled ISO link is specified as Class F. TIA/EIA has not published a Category 7 designation, and probably will not publish such.

After a cable analyzer (a frequency-based cable tester) Autotest function has failed a cable link, verify the following:

- Has the tester been appropriately configured for this Autotest?
- Has the correct link type been selected (Permanent Link or Channel)?
- Are you using the appropriate cable interface adapter for this test? Some third generation testers required interface adapters to match the installed cable for Permanent Link testing.
- Are you using the most current version of tester software? As noted above, standards change.
- Do the cable and connectors used in the installation match the performance settings of the Autotest selected?
- If the link is a TIA Category 6/ISO Class E, or Category 6A/ISO Class E_A installation, are all components matched appropriately? Some Category 6/ISO Class E and Category 6A/ISO Class E_A links installed before the standards body voted to approve that cable standard may not operate at the specified performance level when mixed with other vendor’s materials.
- Is the tester at ambient temperature and in calibration? Temperature will affect test results.
- Are the tester’s batteries adequately charged? Some test results become unreliable when the tester batteries fall below 20% of full charge on some cable testers.

- Have you carefully reviewed the installation quality of terminations and re-terminated where necessary?
- Are the cables too neatly “dressed?” If tie-wraps are too tight, or if high performance cable (Category 6/ISO Class E, and especially Category 6A/ISO Class E_A) is aligned perfectly in parallel for too great of a distance it may create problems which otherwise would not exist.

If test results pass or fail with a marginal (*) result, then examine the details to see if there is a point-source problem which could be corrected to improve the measured result in a retest. Run TDR or TDX tests and examine the graph for evidence of the fault location.

If the test failed without displaying any marginal test results (those marked with an asterisk *), and there were no wiremap failures, there is very little chance that *tuning* the cable will achieve a Pass test result for Category 6/Class E. If this is the case, use the advanced diagnostic features available from your cable analyzer to attempt to isolate the connection, cable, or patch cord as the source of failure. Start by running TDR or TDX tests and examine the graph for evidence of the fault location.

If the cable itself seems to be the source of the fault, or if you have a homogeneous Category 6/Class E system (all cable and connectors are part of a system from one vendor), then save complete test results and record the tester’s model, serial number, and software version. Contact the appropriate cable supplier, share your test results with them, and work to resolve the problem. Category 6/Class E products were sold for about two years before the standard was approved, and sometimes were not interoperable with other vendor’s Category 6/Class E products. Similar issues may exist for early Category 6A/Class E_A cable installations.

Copper cable tests

Wiremap

Wiremap failures are the easiest to locate as they involve opens, shorts, and pairing faults. Use wiremap test results and length measurements to isolate the location of termination, continuity, and pairing faults. Some split pair faults may require a distance-to-crosstalk test (such as TDX) which operates in a manner similar to a distance-to-fault test (length or TDR), and is described in the Advanced Cable Diagnostics section.

Most wiremap failures occur at cable terminations, either at the RJ45 (plug or jack), or at an intermediate cross connect or patch panel. Faults at the RJ45 can usually be seen by checking the wire colors carefully against T568A or T568B pinout colors, or by checking the RJ45 plug for wires that did not seat fully to the end of the connector when it was crimped. While checking for wires that were not fully seated, also try check to see if the correct type of RJ45 was used (stranded or solid wire pins) – though that is difficult once crimped (see Figure 2).

Using the wrong style of pin may cause intermittent connections after a period of time, though the cable usually works immediately after it is made.




Figure 2: Pin styles for crimping stranded and solid cable in an RJ45 plug.

Another source of RJ45 related problems is how well the connector was crimped. In the group of four bad crimps shown in Figure 3, the top left crimp pressed the end pins down adequately, but not the center pins. The top right crimp is exactly the opposite, the crimp tool pressed firmly in the center but both edges did not press adequately. The bottom two crimps show where pressure was applied firmly on one side of the crimp, but insufficient pressure was applied to pins on the other side and they were not adequately crimped. These four problems are usually associated with a low-cost crimp tool constructed with a plastic frame, where the plastic flexes as more pressure is applied. A multitude of other bad crimps is possible, including all pins being pressed evenly, but not far enough.




Figure 3: Examples of bad RJ45 crimping.

Partial crimps are likely when the tool does not ratchet down, and permits the RJ45 to be removed from the tool before it is fully crimped. Sometimes the crimp tool is damaged and one or more pins are not crimped at all. Sometimes the crimp tool is not rigid enough,

and it flexes to produce the problems shown in Figure 3. If a previous RJ45 plug was not crimped well then one of the wires in the jack may have been pushed flat, and may not extend out far enough to make contact with the pin in the RJ45 plug (see Figure 4).




Figure 4: (Left) RJ45 jack damaged by improperly crimped RJ45 plug pin. Outside pin on either side is permanently depressed. (Right) RJ45 plug which was not crimped properly.

The jack problem in Figure 4 can often be corrected by finding a thin pointed tool and carefully re-bending the pin so that its normal resting position places it in alignment with the other undamaged pins again. Take your time and do not over-bend the damaged pin in your attempt to fix the problem. Be aware that this attempt may void some product warranties. However, the risk in attempting this is small because the RJ45 jack is already damaged. If this problem is found in a classroom environment where students regularly make patch cables it may be appropriate to make a very short extension cable with a plug and jack, so that the extension cable is damaged by student cables instead of the equipment. The plug should be cut off and re-terminated.

Be careful to examine the plastic separating each pin in the RJ45 plug, as abuse or neglect may cause the plastic to bend over the pin and prevent the corresponding wire in the jack from making contact. This is a common problem with patch cables.



front view

top view

Figure 5: Two examples of damaged RJ45 plugs found while troubleshooting.

In Figure 5 the top RJ45 plug has easy to see damage to the plastic separation between pin connections. The two pins will not make contact, and the third might not. The bottom RJ45 plug has bent plastic also, but without comparing the separation distance it is hard to see that the right-most pin has too small of a gap for contact with the wire in the plug.

Also, examine the RJ45 jack to see if any of the wires have been bumped out of their track, and are shorting against an adjacent wire.




Figure 6: Pin out of place inside an RJ45 jack.

Length

The problem of having network cabling installed by someone untrained in the requirements still results in cable installations in excess of the maximum allowed 100 meters. The cable may simply be too long. If the cable is simply too long, look for coiled service loops that the installer may have left and remove one or more. Service coils in ceilings and walls were common (and useful) at the time of Category 5 cable, but coiled cable causes various crosstalk problems with Gigabit and 10 Gigabit Ethernet.

Also, check to see if the NVP setting for the tester is incorrect, which will result in inaccurate length measurements. NVP may be calculated by most cable analysis tools by simply measuring the physical length of a moderately long cable (at least 15 meters or 50 feet), and having the tester then calculate the length of the same cable. Adjust the tester's length calculation if necessary to obtain the NVP for that cable sample.

If one or more pairs of the cable are of substantially different lengths then check intermediate patch panels and interconnection points for loose wires and improper connections. Most such wiring faults are at these intermediate connecting points. Be aware that there will be minor length differences for each pair in almost all cables, as the twist rate varies on each pair.

The TIA/EIA-568-B standard directs that the length of the shortest pair determines the overall length of the cable. This means that a long cable could have one or more pairs that measure longer than the standard allows, and the test still passes.

If the cable is unexpectedly short then look is anywhere that facilities work or construction is underway or has recently been

performed. If you have a general idea where the cable path lies, then it should be relatively easy to estimate the location of the fault based on general length information. A common location for cut cables is the edge of new carpet, and doorways or other locations where the cable may have been pinched. Broken tabs on the RJ45 plug permit the plug to retract from good contact in the jack, and may cause opens over time.

The electronic length of a pair is affected by the dielectric insulation used on the wire. If one or two pairs in the cable have a different insulative material than the other pairs, the NVP – and therefore the length – will be markedly different (see Figure 7). Most high performance network cable has Teflon as an insulative material on each wire. However, for about a year in the mid-1990s there was a Teflon shortage after a fire in a key Teflon manufacturing plant. Until a new Teflon supply became available manufacturers experimented with using PVC as an insulator on the least-used wire pairs to reduce cost. The cable was generally available with either one or two pairs insulated in PVC and may be referred to as 3:1 or 2:2 cable. This mix of insulative material affects length, delay skew, and propagation delay. This type of cable is unlikely to perform adequately for Category 5e uses, and should be replaced.

LENGTH	
NO REMOTE DETECTED	
Pair	Length (ft)
1,2	320
3,6	303
4,5	305
7,8	304




Figure 7: Length measurement for a cable where three pairs use Teflon as the insulative material and one pair uses a PVC compound.

Insertion loss

Insertion Loss, more commonly known as Attenuation, is usually associated with cable length. The amount of signal lost grows proportionally to the length of the cable. Thus, the first place to check when trying to solve this problem is the overall length of the cable. Shortening the cable should help, if it can be done. Although the most logical cause, length is often not the source of the problem.

A far more common source of this problem is a very poor connection that often results from a loose cable, dirty or oxidized contacts, and so on. One bad patch cable can easily cause an entire link to fail. This type of problem increases Return Loss, which is why the Attenuation test was changed to Insertion Loss. Run TDR or TDX tests and examine the graph for evidence of the fault location.

Another source of this fault is the wrong Category cable used, such as Category 5e cable used for a link being tested to Category 6A limits. Again, run TDR or TDX tests and examine the graph for evidence of the fault location.

Near End Crosstalk (NEXT), ANEXT, and Power Sum

Excessive crosstalk, usually reported in the NEXT test results, originates in two places: inside the link (in-channel) and outside the link. Crosstalk originating inside the link is worst (has the greatest amplitude or leaks into the measured pair the most) nearest the transmission source where the transmitted signal is also the loudest. If the cable is left untwisted for more than the allowed 13mm (0.5 inch), the crosstalk will be correspondingly worse. For crosstalk, workmanship at each connection point should be examined.




Figure 8: Good workmanship example – pairs are untwisted only enough to terminate.

Re-terminate any connection with visibly untwisted wire. Try removing and re-punching the cable at intermediate cross connect locations if removing untwisted wire segments is not sufficient. The old Telco-style 66 blocks should not be used for network cabling, as they have very poor crosstalk and other test result performance. To satisfy Category 5e, 6, and 6A requirements the 110 style or other punchdown block should be marked for the level of performance it offers. Before spending time attempting to resolve crosstalk issues, refer to the advanced cable diagnostics section for additional tests to help pinpoint the source of crosstalk problems.

There is a special situation where a failing test result at a given frequency does not cause the overall test result to fail. The TIA and ISO standards include a so-called “4 dB rule.” If the insertion loss is less than 4 dB, then NEXT results pass regardless of the NEXT result (as long as ACR passes). A similar rule applies to measuring Return Loss where if the measured result is less than 3dB then the Return Loss test is informational only and is not used for Pass/Fail.

Noise

There are three general types of noise:

- Impulse noise that is more commonly referred to as voltage or current spikes induced on the cabling.
- Random (white) noise distributed over the frequency spectrum.
- Alien crosstalk (crosstalk from one cable to another adjacent cable).

Of the three, impulse noise is most likely to cause network disruptions. Most cable analyzers have impulse noise test capabilities. The 802.3 standard set the default threshold level for the detection of impulse noise at 264 mV in Clause 14.4.4. For higher-speed network applications such as 1000BASE-T, the threshold value for impulse noise detection is 40 mV in Clause 40.7.6. If there are very few pulses at this threshold level (less than 1 in 100 seconds), the cabling will be able to deliver very good support.

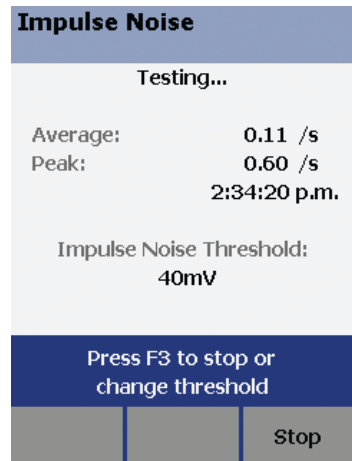


Figure 9: DTX-1800 Impulse Noise test.

Impulse and random noise sources include nearby electric cables and devices, usually with high current loads. These may include large electric motors, elevators, photocopiers, coffee makers, fans, heaters, welders, compressors, and so on. Another less obvious source is radiated emissions from transmitters, including: TV, radio, microwave, cell phone towers, hand-held radios, building security systems, avionics, and anything else that includes a transmitter

more powerful than a cell phone. Some cable analyzers will average this sort of noise out of the test results. The test also takes longer to run, as many additional measurements must be taken.

A small amount of noise “riding” on top of the network signaling does not materially affect the ability of the receivers in NIC cards and other active network devices to detect and interpret the network signals correctly. However, if the tester must average the noise out of the test results, it is likely that the network traffic will be disrupted by this noise.

Locate the noise source and move it or the cable, or convert that cable run to fiber. Finding the source can be problematic, as external noise sources are often intermittent. Use of a spectrum analyzer is often required to determine the frequency and magnitude of the noise. While searching for the source, be very aware of what is occurring in the area. The sudden absence of noise can be as helpful in locating the source as its continued presence. Discover what was just used or turned off.

Alien crosstalk is a special case of noise because it is induced by other cables in the same pathway. Anytime a UTP link is tested in a cabling bundling in which some links are active, the chances are very good that the tester will detect alien crosstalk – especially when the adjacent traffic is 100BASE-TX – and report the “external noise detected” message. Typically, alien crosstalk will not affect the reliability of network traffic operating at speeds below 10GBASE-T.

In general, detected noise will not impede or interfere with the reliable operation of the network if the following conditions are met:

- The cable analyzer completes an Autotest and the test results yields a “Pass.”
- The impulse noise test executed on the affected cabling links shows less than 0.01 average pulses per second when the detection threshold is set to 40 mV.
- If a link is tested in a bundle with active links, ensure that alien crosstalk will not interfere with the network operations: the link passes with a NEXT headroom of 3 dB or better over the required performance specification of the network application.

ACR-F or Equal Level Far End Crosstalk (ELFEXT)

Almost all far-end crosstalk results from the plug, the jack, or an inductive coupling in the mating of the two. Almost all near-end crosstalk results from capacitive coupling along the cable.

However, generally solving NEXT problems will eliminate most FEXT problems measured as ACR-F or ELFEXT. That leaves the electrical properties of the connections themselves.

First try replacing the RJ45 plug at the problem end of the link, and if that is not sufficient then try replacing the plug and jack with a mated pair from a cable system offered by a single vendor.

Return Loss

Return loss is a measure of all reflections that are caused by the impedance mismatches at all locations along the link. It indicates how well the cabling's characteristic impedance matches its rated impedance over a range of frequencies. The characteristic impedance of links tends to vary from higher values at low frequencies to lower values at the higher frequencies.

The termination resistance at both ends of the link must be equal to the characteristic impedance of the link to avoid reflections. A good match between characteristic impedance and termination resistance in the end equipment provides for a good transfer of power to and from the link and minimizes reflections. Return loss results vary significantly with frequency.

One small source of return loss is variations in the value of the characteristic impedance along the cable. This may be due to slight untwisting or separation of wires in the pairs, or due to variations in the metal of the wire and the uniformity of the insulation. The parameter Structural Return Loss (SRL) summarizes the impedance uniformity along the length of a cable, and is an indication of how consistent the manufacturing process for that cable was.

Another source of return loss is reflections from inside the installed link, mainly from connectors. Mismatches predominantly occur at locations where connectors are present. The main impact of return loss is not on loss of signal strength but rather the introduction of signal jitter. Signal reflections truly cause loss of signal strength but generally, this loss due to return loss does not create a significant problem.

Since return loss causes reflections, the TDR test is used to locate the discontinuities causing the problem. The more severe the return loss problem, the greater the amplitude of the problem on the TDR trace.




Figure 10: Sample high definition TDR test results from DTX-1800. This screen appears to have a very bad connection to the main tester, and a bad patch cable about 70 feet from the tester.

Propagation Delay

TIA/EIA-568-B permits up to 498 ns of propagation delay for the Permanent Link and up to 555 ns of propagation delay for the Channel Link, for all Categories. It is unlikely that this parameter could fail without other parameters failing as well. Failing propagation delay suggests inappropriate or bad cable in the link, or a cable which is simply too long.

Check the overall length of the cable. Inspect the cable closely to see if the correct type of cable was installed.

Delay Skew

TIA/EIA-568-B permits up to 44 ns of delay skew for the Permanent Link and up to 50 ns of delay skew for the Channel Link, for all Categories. Both of these numbers are quite generous. It is difficult to fail delay skew if good materials were used in the link. A delay skew failure is possible if wire pairs in a single cable have different insulative material on some pairs. See the discussion under the *Length* category above regarding different insulative materials. A failure is also possible if various lengths of twisted wire pairs were used as a patch cable or jumper at a connection point.

Varying the lengths of pairs at any point along the link probably indicates bad workmanship, as individual pairs should never be used for networking applications. This situation should cause other parameters to fail too.

Inspect the connection points in the link, and if the workmanship appears reasonable, you may have little choice but to replace the entire cable run. Test a sample of the new cable before installing it to be sure that your materials are not causing the problem.

Interpreting copper cable test results

Before troubleshooting a failing cable link, verify the tester configuration. This step is critical to obtaining accurate test results. At a minimum, verify that the correct test specification and link type has been selected. In addition, the test standards have evolved sufficiently that the requirements for a particular test may no longer be the same as what is loaded in the software of your tester. Check the tester manufacturer's web site for new tester software regularly, perhaps two or three times per year.

Unlike network failures, cable failures are approached in approximately the same manner whether the link is newly installed or if it has failed during operation. There are many instances where a poor quality link has been in service, but due to the operating environment and influences, it has stopped working. These influences include visible damage to the cable, as well as placing noise sources near the cable or moving the cable near a noise source. Another less obvious condition is that a new network physical layer implementation is now in use, such as an Auto-Negotiating network adapter that has linked at 1000 Mbps instead of the 100 Mbps that has been customary. This sort of condition could result from a new network adapter having been installed in the station, or from moving the connection to a different port on the hub or switch, or moving the connection to an entirely new hub or switch. Some ports will monitor the link for polarity faults (pair reversals) and crossover cables (transposed pairs), and correct for them internally. The newly connected port may not be doing that, and a preexisting cable fault is finally exposed. Table 1 suggests many common sources of failure, and the test

that will reveal them. This table is by no means the only source of these failures, nor does it portray the only test(s) that will reveal the listed failures.

Description	Open	Short	Reversed Pair	Crossed Pair	Split Pair	Length Problems	Delay Skew	Insertion Loss	NEXT	Return Loss	ACR-F	Alien Crosstalk
Cut, broken, or otherwise abused cable	•	•				•				•		
Damaged RJ45 plug or jack	•	•										
Mixed T56A and T568B color codes on same cable				•								
Different insulation material on some pairs							•					
Poor workmanship at cable junction or connector			•		•	•		•	•	•		
Improper wiring at cable junction or connector									•	•		
Improper, poor quality, or Telco rated RJ45 coupler									•	•		
Poor quality or lower-rated RJ45 plugs/jacks									•	•	•	•
Bad, or poor quality patch cord(s)									•	•		
Mixed use of 100 ohm and non-100 ohm cable										•		
Cable is too long or NVP is set incorrectly						•		•				
Untwisted or poorly twisted cable (includes too low of a cable rating, such as Cat 5 instead of Cat 6)									•	•	•	•
Cable ties too tightly fastened along cable									•	•		
External noise source near cable									•		•	•
Cable too closely aligned for a moderate to long distance – remove bindings and/or separate slightly												•

Table 1: Most likely cable test failures and causes

Troubleshooting fiber optic media

Tools

There is a narrow range of tools that may be used to troubleshoot fiber optic cabling. At the low end are effectively continuity testers. Intermediate level testing is performed to check that optical power levels are satisfactory across the link. Advanced diagnostics require an Optical Time Domain Reflectometer (OTDR), which is fairly expensive. If power levels are unsatisfactory, or if OTDR testing reveals a point-source problem, then cleaning and end-face inspection is appropriate.

Safety

Safety should be considered at all times when working with fiber optic cable. Wavelengths used in networking are outside of the visible light spectrum (the human eye begins to see violet light around 380nm, and stops seeing red light around 750nm). Many light sources used in networking are laser-based, and some are very powerful. You should never look straight into either a fiber optic cable end, or any fiber optic equipment jack. Place dust covers over unused equipment jacks, both to keep the connection clean and to prevent eye damage from the non-visible light being transmitted.




Figure 11: Visible light is below the wavelengths of light used in networking applications, at approximately 380nm to 750nm.

Safe viewing of a visible light source is best accomplished by pointing the end of the fiber at white paper, or holding the paper in front of the fiber connection point. Never look directly into any connection from which non-visible light may be emitting.

Continuity testing

One method for fiber continuity and pair polarity is tested using visible light. Popular sources for visible light include a standard white-light flashlight, as well as the many colors of very bright LED “keychain” lights now available. Special network flashlights are available which come with your choice of connector, including: SC, ST, and so on. The network flashlights typically offer focused bright red light, but from an incandescent source, not laser.



Figure 12: Flashlight manufactured for use with fiber optic cable.

Continuity may also be tested using visual fault locator (VFL) laser light sources, which operate in the visible light spectrum. VFL light sources are typically incandescent or laser based, but are most often Class II lasers operating at 650nm (red light).



Figure 13: VFL used to locate a break in a patch cable. Note that light will not penetrate all fiber jackets at a break.

While not always possible, due to the specific coatings used on a fiber optic cable, some cables will permit the location of a fault to be seen if there is a break or other severe fault in the cable.

Attenuation or loss testing

The terms *loss* and *attenuation* may be used interchangeable in relation to fiber optic cable, though loss can be attributed to a point-source fault. An Optical Loss Test Set (OLTS) is a special tester combining a light source and light power meter that tests for the total amount of light loss (attenuation) on a fiber link. The light source produces a continuous wave at specific wavelengths connected to one end of the fiber. A power meter with a photo detector is connected to the opposite end of the fiber link. The detector measures optical power at the same wavelengths produced by the light source. The light source may be LED or laser based, and is very similar to the type of light source used for networking applications. The measured result is then used to see if the required power budget has been met for the technology to be used on the link. Per TIA and ISO standards, which both define testing of installed fiber; an OLTS is a Tier 1 test device.

OTDR testing

An Optical Time Domain Reflectometer (OTDR) graphs the reflections and backscatter from a high-power light pulse sent into the test fiber in much the same manner as a TDR test graphs reflections on a test copper cable. When the pulse of light meets connections, breaks, cracks, splices, sharp bends or the end of the fiber, some amount of the light reflects back toward the OTDR where high-gain light detectors measure the strength of the reflection. In addition, a small amount of light is reflected back from the crystalline structure of the glass itself as backscatter, and is represented by the sloping trace along the length of the OTDR test result. The backscatter slope is used to measure attenuation. Close examination of the resulting graph reveals characteristic changes in the graph plot that may be interpreted as the

connections, breaks, cracks, splices, sharp bends, and so on mentioned above. As with a TDR, the delay between transmission of the light pulse and detection of any reflections may be interpreted as distance to the event. An OTDR trace is valuable because it makes it possible to certify that the workmanship and quality of the installation meets the design and warranty specifications, for current and future applications. With an OTDR, the performance of each splice and connector can be measured. Per TIA and ISO standards, an OTDR is a Tier 2 test device.

End-face inspection

Optical or video microscopes permit fiber end-face inspection, looking for dirt and contamination on fiber optic cable ends and on the end-equipment transmitters, or for problems with the end-face polish on fiber optic cables. Typical magnification is between 200x and 400x. One recent study indicated that more than 80% of all fiber problems related to contamination.

Types of fiber optic cable

The general assumption is that there is singlemode and multimode fiber, but it goes a bit deeper than that. Several examples are provided.

Some of the older multimode has been called “FDDI” fiber. This generation of fiber optic cable and glass quality called step-index fiber. The manufacturing process for this generation of older fiber optic cable left impurities, defects, and variations in the refractive index in the glass core. LED light sources were used with step-index fiber, and excited many modes on the fiber. Each “mode” represents a slightly longer path down the fiber where the light was traveling at greater angles away from straight down the center. Since distance is increased as the angle away from the center increases, the light

would arrive at the far end later in time than the light that traveled straight down the core. This caused a sharply transmitted pulse to arrive as a rounded bump. Firing pulses at high data rates resulted in the rounded bumps blurring together and becoming impossible for the receiver to distinguish one from the next. This is called modal dispersion.

The next generation of cable was called *graded index*, which uses a different composition of glass as you progress outward from the core, which causes the light rays to bend back toward the center. Instead of bouncing off the cladding, light in this type of fiber tends to flow more like a sinusoidal wave, often without quite touching the cladding. This type of fiber reduces the modal dispersion, permitting the transmitted signals to be recovered at greater distances than older step-index fiber.

Laser optimized multimode is manufactured with glass having a much more consistent refractive index. This permits the VCSEL lasers used with Gigabit Ethernet to excite fewer modes during transmission, which in turn results in less modal dispersion. The transmitted signal arrives more sharply defined at the far end, which permits higher signaling rates to be transmitted. The early laser optimized fiber that was first introduced in the mid 1990s is not capable of supporting 10 Gigabit. More recent formulations of glass and better manufacturing processes that result in a more tightly controlled refractive index began appearing as early as 1999, and are rated for 10 Gigabit. At the same time, since the glass core is wider, more modes may exist; the popularity of 62.5 μm multimode is waning in favor of the narrower 50 μm multimode fiber. Since there are fewer modes found in the 50 μm cable, the signal may be reliably recovered at greater distances and at higher data rates.

Singlemode fiber has had similar changes. The concept of singlemode fiber is that the core is so narrow that only one mode may exist at the wavelengths used - straight down the center. The basic construction is non-dispersion shifted fiber (NDSF) which worked very well at 1300/1310nm. This type of fiber did not work that well for 1550nm use. Instead, the cable was reformulated to move the optimum supported wavelength to 1550, and was called dispersion-shifted fiber (DSF).

When DWDM networking was introduced, it was discovered that the DSF fiber had some odd nonlinearities, and so non-zero dispersion shifted fiber (NZ-DSF) was created. Other more specialized fiber compositions and constructions are being developed now, such as polarization-maintaining (PM) fiber.

Research the parameters associated with any installed fiber before repurposing it to be used for a new technology.

Fiber optic cable tests

Testing fiber amounts to testing polarity, length, and attenuation. Short of laboratory grade equipment, there is as yet no convenient way to perform field-testing of many fiber properties.

Polarity may be verified using a visible light source such as a VFL or flashlight, and by testing both fibers in a pair simultaneously with an OLTS or OTDR.

Length may be obtained by examining the cable jacket markings, or from some OLTS. An OTDR will provide excellent length information.

Overall channel Attenuation may be measured with either an OLTS or an OTDR. OTDR test results can assist with loss budget calculations for a channel link by providing attenuation information about each detected event individually (see Figure 14).

EVENT TABLE				OFTM-5612	
Auto OTDR				06/22/2006 6:29:15 p.m.	
LOCATION (m)	dB@850nm	dB@1300nm	EVENT TYPE	STATUS	
0.00	N/A	N/A	OTDR PORT		
102.14	0.39	-0.22	GHOST SOURCE	PASS	
152.72	0.20	0.97	REFLECTION	FAIL	
164.11	1.17		LOSS	FAIL	
174.58	0.19	0.65	REFLECTION	PASS	
204.69	0.00	0.05	GHOST		
226.32	N/A	N/A	END		

↕ Scroll List, ↕ Select Field, Press EXIT to view SUMMARY

View Trace Sort Field View Details

Figure 14: OTDR screen showing event interpretation and loss.

Interpreting fiber test results

Polarity

Polarity does not actually fail, since the point of the test is to learn and mark or pair cables according to the polarity-pairing scheme employed within your network. Typically, polarity testing is accomplished as part of initiating the attenuation measurements. If the light source and light meter are not attached to the same fiber then they will not produce results.

Some networks pay little attention to polarity throughout the cable plant, and simply rely upon swapping the fiber connection at the equipment end if link is not established by how the fibers were first attached. When troubleshooting fiber problems an excellent first test is to swap the fibers attached to TX and RX at one end of the link. Whether the information learned results in correcting the pairing along the channel or simply accepting that the pair was swapped, a polarity problem is often solved very quickly.

Length

An OTDR will reveal the overall channel length, which may then be compared against the implementation specifications for the networking technology used. The OTDR may also reveal a link that is shorter than expected, and may be the result of a break in the cable.

If an OTDR is not available, then knowledge of the cable plant or access to the original installation certification documentation can be very beneficial. Utilizing the length markings on the cable jacket is another way to learn the length of each individual cable segment in the channel link. Again, the resulting channel length may then be compared against the implementation specifications for the networking technology used.

In either case, pay close attention to the modal bandwidth for the installed cable type. In many cases, this will have to be researched based on cable jacket markings, and then cross-referenced against distance limitations for the networking technology used when operating on cable with that modal bandwidth.

Attenuation test failure

Before troubleshooting, first ensure that:

- The number of adapters or splices is set correctly on the tester (for limits that use a calculated loss budget value).
- The correct fiber type is selected in the tester setup configuration.
- A valid power reference was set on the tester recently, under the same temperature conditions as you are now using it, and without disconnecting the patch cord attached to the light source afterward.

Use VFL to ensure you are on the correct fiber. A VFL will usually also isolate the location of broken or cracked fiber (see Figure 13).

Clean all fiber connections (plugs and jacks) in the problem path (including the output port on the end-equipment). Visually inspect the end-face of each cable. Look for cracks, scratches, or persistent contamination.

Quickly wiping the end-face with fiber grade cleaning alcohol may not be enough to dislodge some types of contamination (see Figure 15).




Figure 15: Fiber end-face contamination in the form of a mold or fungal growth.

Test individual patch cables with the OLTS. After setting the reference with a good patch cable, another patch cable should show close to zero loss. Any discrepancy should be investigated. If the problem is intermittent, then try flexing the patch cable while testing to see if changing the alignment opens or misaligns a crack or other damage. Do not exceed the bend radius while flexing.

If an OTDR is not available, then apply *divide and conquer* techniques – move closer to the other end and retest. Look for a large change in the loss results that does not correspond with how much of the link was removed by moving forward. Attenuation of the actual fiber in a LAN environment tends to be negligible, so base your judgment upon the standards limit of 0.75 dB loss for each connector pair.

There may be one or more dirty or damaged connections in the cabling. Clean all fiber end-faces and retest, or use the OTDR to locate bad connections.

A patch cord or fiber segment has the wrong core size. If the patch cords are the correct type, use the OTDR to look for mismatched fiber in the cabling.

The cabling has a bad fusion or mechanical splice or a sharp bend. Use the OTDR to locate these faults.

Inspect the cable path. Is the cable kinked or bent at an angle that exceeds the bend radius? Are cable fastenings causing microbends? See Figure 16.

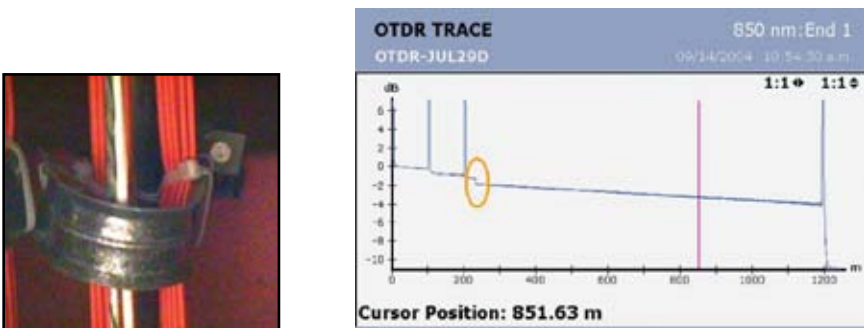


Figure 16: Microbend in fiber, and the corresponding signal loss as seen by an OTDR (circled).

Was a multimode coupler used on a singlemode cable run? The engineering tolerances for singlemode couplers are much more accurate, which prevents core misalignment and the corresponding loss of power. In addition, connectors and cables are only rated for a certain number of insertions. If the cable or connector has been used a great deal, then it could be becoming “sloppy” and not aligning the fiber correctly. An OTDR measurement would reveal the location of this and other related problems.

If the failure being resolved relates to a single device, it is sometimes due to the end-equipment transmitter. Either the connection inside the equipment is dirty, or the transmitter is not outputting adequate power. Try connecting the OLTS to the suspect port to obtain a power reading, and then compare that reading against other similar ports.

Troubleshooting the network layer

Troubleshooting common network user complaints

Users are a good barometer for the performance of your network. They are rarely reluctant to report perceived or actual substandard network performance. Unfortunately, they also rarely have the knowledge that would help you troubleshoot the problems they report. Furthermore, the problem descriptions provided by the user are sometimes imaginative, and may bear no relation to reality. Always remember that their lack of technical knowledge should not be interpreted as indicating that no problem exists. Something annoyed the user enough to contact you. Often the source of the problem is user error, misuse, or mismatch of expectations, and a few minutes of user training will make both of you happier.

Note: *Before you begin troubleshooting in earnest, verify that the desired server or service has operated successfully in the past from the problem location. The troubleshooting process is completely different if you are trying to solve a new installation problem instead of attempting to restore service to something that has been up and running.*

Below are three general categories of user complaint. Almost all user complaints fall into these three categories:

- Can't connect
- Dropped connection
- Poor performance

Some of the problems related to shared media (hubs), some to switched media, and some to both.

For each general problem category, a generalized troubleshooting tree is provided. Each step along the path is determined by the results of one or more tests. The discussion below does not include all possible variations, but instead forms an outline for how to troubleshoot. Conversely, the discussion provides a detailed series of steps. The detail is provided as an attempt to be clear about why that test is important, and what to look for. Do not take that to mean that every test should be performed. Use common sense to choose which steps to try, and which to bypass for now. Think of this list as a mental checklist. While troubleshooting you should be watching for all of these situations, and mentally checking them off. When one of the items on the list is not mentally checked off, then try the test.

Note: *If you change collision domains during the testing process, be sure to start with your mental list of collision domain tests again when you connect at the next location.*

Complaint: Can't connect

The following procedures assume that this server or service has been operating properly prior to this problem, and you have already:

- Cold-booted the station in question (a warm-boot does not reset all adapter cards). This will also apply any loaded but unapplied patches. In addition, some PnP devices seem to require two or three reboots to install fully.
- Verified that the station does not have any hardware failures.
- Verified that required network cables are present and properly connected.

- Verified that the network adapter is not disabled, and has valid addressing for the subnet (static or DHCP). Check also to see what the operating system NIC status reports for frames sent and received, if either is zero then investigate.
- Verified that nothing has been changed recently on the problem station, or on the server or that may have caused this problem, such as reconfiguring or adding software or hardware.

This problem is usually manifested when the user is unable to connect to a server or service. The user is often not able to discriminate between inability of the station to link to the network, and inability to connect to a particular server or service. Compared to the other problem categories, this is the easiest to isolate. Determine whether the problem is isolated to this station or a small group of stations (collision domain problem, including a single switch port) or if it affects many stations (broadcast domain or interconnected networks problem).

Before troubleshooting the hardware try connecting using your own login account, or have the affected user attempt the exact same network operation from another nearby station that is operating correctly. This is the fastest way to isolate user-account problems from network problems. If the first user is still unable to connect, then troubleshoot the first user's account. Observing the user during an alternate attempt may also reveal errors in the series of steps the user is accustomed taking in order to connecting to the network. A moment of training may prevent significant future frustration by the user.

Collision domain problems affect the local medium, and prevent reliable communications to the first Layer 2 or 3 infrastructure device – or the local server or service to which you are trying to connect. They typically result from:

- Bad cables
- Errors or excessive traffic on the local collision domain
- Blocked or mis-configured switch ports
- Failed or mis-configured station NIC
- Corrupted, unbound or mis-configured software drivers

Virtually all of these collision domain problems can be identified with an inline test while the user attempts to connect following a cold reboot. Rebooting a workstation is important as many operating system problems are difficult or impossible to recreate or isolate, and reloading the operating system clears these mysterious problems for a time.

Many users have both a wired and wireless NIC enabled. If the PC is trying to use a wireless NIC instead of a wired connection then the exact location or PC orientation may be preventing adequate connectivity. There are many blind spots in most networks, and some are quite small. Moving the PC even a few inches or rotating it slightly has been known to reestablish a wireless link. If people are congregated near the PC, it may be that they are blocking the signal.

Broadcast domain problems begin after a reliable MAC Layer link is established, and are typified by a failure to create a logical connection across a bridged environment. Included in this category are Network Layer issues that would prevent communications to servers and routers attached to this broadcast domain:

- Marginal or failed uplink port somewhere in the path, possibly the result of a bad cable.
- Broadcast storm or other excessive traffic within the broadcast domain (not necessarily traffic observed on the local port).
- ICMP errors present, or IP addressing incorrect for the local subnet, duplicate IP address.
- DNS and DHCP failures.
- Station or server improperly advertising routes.

Addressing and some other problems will be revealed by the same inline test that may be performed during the collision domain testing. Be sure to repeat collision domain testing if you change locations within the broadcast domain. If an address is obtained and/or correct, it may be necessary to either gather a protocol analyzer trace file for analysis or use network management software to interrogate infrastructure devices within the broadcast domain.

Interconnected network problems begin after a reliable link is established to the router offering a path out of the broadcast domain. The level of complexity usually increases and the level of access often decreases if the server or service resides beyond a WAN connection instead of residing on an adjacent LAN, but the process is similar:

- Unstable routing due to marginal or failed port somewhere beyond the broadcast domain, possibly the result of a bad cable.
- Trace Route failures, few if any Ping responses.
- Incorrect routing configurations, including DHCP request forwarding configuration for when the server is not on the local subnet and isolated VLANs.

- VPN problems, including MTU size.
- Firewall or other security blocking, such as login account or password problems.

Use of Ping and Trace Route will usually reveal the location where troubleshooting *cannot connect* problems should begin. For faster troubleshooting once a remote location is identified as being suspect, use network management to query the suspect infrastructure device and the infrastructure device immediately prior to it. One or the other should be showing errors of some type or excessive utilization. Establishing a reliable end-to-end connection at the Network Layer resolves most problems. Be sure to repeat collision domain and broadcast domain tests each time you move to a new location during the troubleshooting process.

If Ping responses are reliable but the link is still failing, try increasing the size of the Ping frame. This will reveal MTU size problems in the routed path. VPNs add overhead to the frame, and the user MTU must be correspondingly smaller. If end-to-end Network Layer delivery appears to be reliable, a protocol analyzer is almost your only remaining option. Capture and analyze the connection attempt. It may be necessary to repeat the capture from the server or service end of the link to ensure that requests are arriving or that responses are leaving.

If Ping and Trace Route are successful, try using Telnet to the required port. Successful Telnet connections establish link, but may not produce visible evidence of this. If the Telnet connection is refused then that service is not available, and a refused or failed connection is always obvious.

Complaint: Connection drops

Connections that drop may be caused by the same conditions that prevent a connection from being established in the first place. Consider also any of the situations described under the *cannot connect* heading.

The following procedures assume that this connection has been operating properly prior to this problem, and you have already:

- Cold-booted the station in question (a warm-boot does not reset all adapter cards). This will also apply any loaded but not applied patches. In addition, some PnP devices seem to require two or three reboots to install fully.
- Verified that the station does not have any hardware failures.
- Verified that required network cables are present and properly connected.
- Verified that the network adapter is not disabled, and has valid addressing for the subnet (static or DHCP). Check also to see what the operating system NIC status reports for frames sent and received, if either is zero then investigate.
- Verified that nothing has been recently changed on the problem station, or on the server or service that may have caused this problem, such as reconfiguring or adding new software or hardware.
- Eliminated potential station memory allocation problems and software conflicts on the station by loading only the minimum software required to operate a test application across the network. For this test disable any virus checking or security software, but turn it right back on after the test.

- Monitored the user's station for applications that are consuming microprocessor resources or hanging the system long enough to exceed connection timers, possibly a virus.

The reason for dropped connections is a logical or physical connectivity loss. This will be manifested by cable-related problems or by difficulties getting through a switch, bridge, router or WAN connection. Upper-layer protocols implement various timers that will terminate a station's logical connection if the timer expires without having heard from that station. Thus, if frames are being dropped across a switch, bridge, router or WAN connection, it is possible to lose your connection to the server or service while still operating perfectly on the local collision domain or broadcast domain.

Determine whether the problem is isolated to this station or a small group of stations (collision domain problem, including a single switch port) or if it affects many stations (broadcast domain or interconnected networks problem). Ask other users in the area if they have had similar problems. Ask also if the problem has been related to time of day, the type of query, or when some seemingly unrelated event or action in the vicinity takes place.

Collision domain problems affect the local medium, and disrupt communications to the first Layer 2 or 3 infrastructure device – or the local server or service to which you are trying to connect. They typically result from:

- Bad cables
- Marginal or intermittent station NIC, or port on hub or switch
- Errors or excessive traffic on the local collision domain
- Duplex mismatches
- Electrical noise and other environmental disruptions

Many collision domain problems related to dropped connections can be identified by disconnecting the user's station and attaching a tester in its place. Through the user's normal cable, exercise the network connection and attempt to reach the problem server or service. Restore the user's connection and leave an inline tester monitoring the link or a protocol analyzer gathering traffic and statistics. Instruct the user on what information to gather from the tester immediately after the connection fails again, and how to stop and save the captured traffic for later analysis.

Many users have both a wired and wireless NIC enabled. If the PC is trying to use a wireless NIC instead of a wired connection then the exact location or PC orientation may be preventing adequate connectivity. There are many blind spots in most networks, and some are quite small. Moving the PC even a few inches or rotating it slightly has been known to reestablish a wireless link. If people are congregated near the PC, it may be that they are blocking the signal.

Broadcast domain problems begin after a reliable MAC Layer link is verified, and are typified by a failure to maintain a logical connection across a bridged environment. Included in this category are Network Layer issues that would disrupt communications to servers and routers attached to this broadcast domain:

- Marginal or failed uplink port somewhere in the path, possibly the result of a bad cable.
- Spanning tree problems, possibly the result of a bad cable.
- Broadcast storm or other excessive traffic within the broadcast domain (not necessarily traffic observed on the local port).
- Duplex mismatches between ports somewhere in the path.
- Duplicate IP addresses.
- Station or server improperly advertising routes.

Continuously Ping the local router to check for lost frames on the broadcast domain. Use network management to interrogate the infrastructure devices in the path between the user's connection and the router, server or service, looking for errors or high utilization that appear to correspond to times when the connection was lost. Restore the user's station and use a protocol analyzer to monitor and/or capture traffic related to the problem server or service. If the problem is intermittent then leave the protocol analyzer gathering traffic. Instruct the user on how to stop the capture immediately should the connection again fail. This sort of problem is often intermittent, and it will be difficult to troubleshoot unless the user can summon you while it is happening or capture and examine a record of the communications leading up to the failure.

Interconnected network problems begin after a reliable link is verified to the router offering a path out of the broadcast domain. Reliable access to Internet servers or services is more problematic than reliable access to servers or services on adjacent LANs because of Internet Service Provider outages and denial of service attacks that are beyond the control and visibility of the local network support staff:

- Unstable routing due to marginal port or link somewhere beyond the broadcast domain, possibly the result of a bad cable.
- Excessive traffic across a low speed LAN or WAN link, possibly causing traffic to be discarded or buffer capacity to be exceeded.
- Trace Route and Ping response variability.
- Overloaded server or service.

Ping and Trace Route testing may reveal the location where troubleshooting Connections that drop problems should begin. Run

tests continuously in order to detect the probable location of fluctuating or intermittent problems. For faster troubleshooting once a remote location is identified as being suspect, use network management to query the suspect infrastructure device and the infrastructure device immediately prior to it. One or the other should be showing errors of some type or excessive utilization. Perform a throughput test that validates the path capacity between the user connection and the server or service; monitor the path with network management while the throughput test is performed to see if any errors appear. Establishing a reliable end-to-end connection at the Network Layer nearly always resolves dropped connection problems. Be sure to repeat collision domain and broadcast domain tests each time you move to a new location while troubleshooting using the divide-and-conquer process. If end-to-end Network Layer delivery appears to be reliable, a protocol analyzer is almost your only remaining option. Leave a protocol analyzer gathering traffic. Instruct the user on how to stop the capture immediately should the connection again fail. If the server or service is overloaded, or if there is something within the user's station that is consuming microprocessor resources or disrupting communications, the protocol analyzer trace file should identify which end of the connection to investigate first.

Complaint: Network is slow

Poor performance may be caused by the same conditions that prevent a connection from being established in the first place, or from the same conditions that cause connections to drop. Consider also any of the situations described under the *cannot connect* and *connections that drop* headings.

The following procedures assume that this connection has been operating properly prior to this problem, and you have already:

- Verified that nothing has been recently changed on the problem station, or on the server or service that may have caused this problem, such as reconfiguring or adding new software or hardware.
- Eliminated potential station memory allocation problems and software conflicts on the station by unloading all but the minimum software required to operate a test application across the network. For this test disable any virus checking or security software, but re-enable it immediately after the test.
- Tested the user's station for viruses and look for applications that are consuming disproportionate amounts of the microprocessor resources or hanging the system long enough to exceed connection timers.

The most common reasons for slow or poor performance include overloaded or underpowered servers, unsuitable switch or router configurations, traffic congestion on a low capacity link, and chronic frame loss. Tiered applications may suffer poor performance when any one of the servers in the tiered hierarchy suffers delays. Analyzing tiered applications can be tricky, as it is often difficult to map all of the dependencies.

Determine whether the problem is isolated to this station or a small group of stations (collision domain problem, including a single switch port) or if it affects many stations (broadcast domain or interconnected networks problem). Ask other users in the area if they have had similar problems, whether with the network or a particular application. Ask also if the problem has been related to time of day, the type of query, or when some seemingly unrelated event or action in the vicinity takes place.

Collision domain problems affect the local medium, and disrupt communications to the first Layer 2 or 3 infrastructure device – or the local server or service to which you are trying to connect. They typically result from:

- Bad cables
- Marginal or intermittent station NIC, or port on hub or switch
- Errors or excessive traffic on the local collision domain
- Duplex mismatches
- Electrical noise and other environmental disruptions

Many collision domain problems related to slow or poor performance can be identified by disconnecting the user's station and attaching a tester in its place. Through the user's normal cable, exercise the network connection and attempt to reach the problem server or service. Restore the user's connection and leave an inline tester monitoring the link or a protocol analyzer gathering traffic and statistics. Instruct the user on what information to gather from the tester immediately after the connection fails again, and how to stop and save the captured traffic for later analysis.

Business critical applications should rarely be delivered wirelessly, owing to the many possible impediments and generally lower available bandwidth. Check that NIC is being used to access the poorly performing server or service. If the circumstances dictate that wireless is appropriate, then use a spectrum analyzer to check for continuous or intermittent noise sources, or any other competition for the frequency band in use for this wireless link.

Broadcast domain problems begin after a reliable MAC Layer link is verified, and are typified by a failure to maintain a logical connection across a bridged environment. Included in this

category are Network Layer issues that would disrupt communications to servers and routers attached to this broadcast domain:

- Marginal or failed uplink port somewhere in the path, possibly the result of a bad cable.
- Spanning tree problems, possibly the result of a bad cable.
- Broadcast storm or other excessive traffic within the broadcast domain (not necessarily traffic observed on the local port).
- Duplex mismatches between ports somewhere in the path.
- Duplicate IP addresses.
- Station or server improperly advertising routes.

Continuously Ping the local router to check for lost frames on the broadcast domain. Use network management to interrogate the infrastructure devices in the path between the user's connection and the router, server or service, looking for errors or high utilization that appear to correspond to times when the connection was lost. Perform a throughput test to various points across the broadcast domain being careful to use the same uplinks as the desired traffic travels. Watch for inconsistent throughput results that will reveal duplex mismatches and other error related problems. Restore the user's station and use a protocol analyzer to monitor and/or capture traffic related to the problem server or service. Watch especially for ICMP errors and TCP retransmissions. If the poor performance problem is intermittent then leave the protocol analyzer gathering traffic. Instruct the user on how to stop the capture immediately should the connection again fail. This sort of problem is often intermittent, and it will be difficult to troubleshoot unless the user can summon you while it is happening or capture and examine a record of the communications leading up to the failure.

Interconnected network problems begin after a reliable link is verified to the router offering a path out of the broadcast domain. If performance is unsatisfactory all the time, then the problem is probably related to a poor configuration, inadequate capacity somewhere, or some other systemic problem. If performance varies and is not always unsatisfactory, then the problem is probably related to an error condition or is being affected by traffic from other sources.

- Unstable routing due to marginal port or link somewhere beyond the broadcast domain, possibly the result of a bad cable.
- Excessive traffic across a low speed LAN or WAN link, possibly causing traffic to be discarded or buffer capacity to be exceeded.
- Trace Route and Ping response variability.
- Overloaded server or service.

Ping and Trace Route testing may reveal the location where troubleshooting Slow or poor performance problems should begin. Run tests continuously in order to detect the probable location of fluctuating or intermittent problems. For faster troubleshooting once a remote location is identified as being suspect, use network management to query the suspect infrastructure device and the infrastructure device immediately prior to it. One or the other should be showing errors of some type or excessive utilization. Perform a throughput test that validates the path capacity between the user connection and the server or service; monitor the path with network management while the throughput test is performed to see if any errors appear. Establishing a reliable end-to-end connection at the Network Layer nearly always resolves dropped connection problems. Be sure to repeat collision domain and

broadcast domain tests each time you move to a new location while troubleshooting. If end-to-end Network Layer delivery appears to be reliable, a protocol analyzer is almost your only remaining option. Capture and analyze the connection attempt. If the server or service is overloaded, or if there is something within the user's station that is consuming microprocessor resources or disrupting communications, the protocol analyzer trace file should identify which end of the connection to investigate first.

Troubleshooting switches

Typical switched network problems

The problems found in a switched environment are generally the same as those experienced in a shared media environment. What happened, who did it, and how much? The primary difference is that answers need to relate back to a specific port.

Some of the issues that should be considered in a switched environment are:

- How busy is each port?
- How do you identify and track the source of errors?
- What is the source of a broadcast storm?
- Are bridge forwarding tables operating correctly?
- Which stations are attached to this port?
- Is the switch rate-limiting any protocols or ports?
- Is this port in a VLAN, and if so is it the same VLAN as the server or service?

How do you determine where to start looking for a reported problem in a switched network? Typically, it is not the switch that

causes the problem, but the inability to “see” inside it. This problem begins with the OSI Layer 2 bridging performed by a switch, and is exacerbated by enabling VLANs and other OSI Layer 3 and higher features and forwarding rules. Advanced switching features such as OSI Layer 4 and higher forwarding and load balancing require a strong knowledge of the switch configuration options to troubleshoot.

By installing a switch, you tend to create a collision domain on each port – that is simply the nature of a switch. If shared media hubs are attached to the port, then the collision domain may grow to the maximum size allowed for that Ethernet implementation. Due to the dropping price of switching technology, most new networks have a single station per port, so the collision domain is only a single cable link.

The entire switch tends to be part of a single broadcast domain, including any number of other switches connected in series or in parallel. If OSI Layer 3 features are enabled then multiple broadcast domains are created, equal to the number of VLANs. At the extreme, and if the switch features permit it, each port could be configured to be a separate broadcast domain. This configuration could reasonably be described as routed to the desktop. By creating a separate broadcast domain for each port, troubleshooting options are limited severely. A separate broadcast domain per port will also require a routing service in the switch to spend considerable CPU resources in forwarding traffic. The network situation where it is appropriate to require routing every single request and reply is very difficult to imagine, and this configuration should be avoided unless a very good reason could be found. Unfortunately, a less obvious form of this configuration is all

too common, and is found in networks where the servers are all located within one subnet or broadcast domain, and all users are in some number of other subnets or broadcast domains. Virtually all requests must still be routed. If maintenance activities must be limited to a single server room, then consider placing servers in separate VLANs. Then place the users that depend upon that server in the same VLAN. This configuration would allow the switch matrix to use OSI Layer 2 bridging for routine traffic, and only unusual or infrequent requests would be routed. If the server supports more than one user community, install additional network adapters in the server to maintain OSI Layer 2 connectivity to the users.

Isolating the problem

Almost the only effective method of troubleshooting a switched network is to ask the switch itself how the network is behaving. This is usually done with SNMP or by connecting to the console port of the switch. Obviously, directing queries through the console port is not desirable because you would have to physically touch every switch in the network. It is possible to minimize the impact of this alternative by setting up terminal servers that connect to the console ports. SNMP is a better choice most of the time because it allows you to make queries in-band from anywhere on the attached network, and it does that without any extra hardware. If you have implemented a network management system you may configure the switch to send an unsolicited response called an SNMP trap whenever utilization, errors, or some other parameter exceeds a specified threshold. Then use network management or a network monitoring tool to investigate what caused the threshold to be exceeded. There are some categories of problem which are

satisfied by asking the switch, but many which are not. Asking the switch has proactive as well as reactive network monitoring characteristics.

The alternative is to wait for user complaints. In most networks, this second option will be the more common solution. This method should not be discounted due to its simplicity – it is very effective. The user community has a very finely tuned subconscious sense of what the normal performance of the network is. Any perceived degradation of that sense of normal will result in a rapid complaint to the network support center. Once a user complains, you can start the troubleshooting process from his or her connection point. This method is entirely reactive.

Proactive efforts to prevent problems from affecting users include regularly interrogating each switch, and monitoring the quality of traffic on each switch port – just as any other segment would be monitored on a regular basis.

Once a problem has been reported or detected there are many ways to approach diagnosis, and each has positive and negative aspects.

Switch troubleshooting techniques

There are at least ten fundamental approaches used to gain visibility into a switch. Each technique offers a different view, and has both positive and negative aspects. Like many other situations related to networking, there is no single best answer. The most suitable solution will be determined primarily by the availability of resources (which tools are available and/or pre-installed), the skill level of the user, and by the potential service interruption that will from implementing that technique.

Even combined these techniques are not able to monitor the attached network as well as when hubs – instead of switches – were common. It is extremely difficult to see all of the traffic and errors that a switch is experiencing. Most troubleshooting assumes the traffic will pass between the station and an attached server or through the uplink. If two stations were passing information directly between themselves using peer-to-peer networking, the traffic would not pass through the uplink or to any other port on the switch. Unless you knew to look for it, it probably would not be detected. Errors tend to stop at the switch port, but could be forwarded depending on the nature of the error and the configuration of the switch. If forwarded, they are nearly always forwarded to only one other port.

For simplicity, the troubleshooting model will be a server attached to a switch as shown in Figure 17. Some descriptions will assume that the users in question are attached to the same switch;




Figure 17: A very basic switch scenario.

other descriptions will assume that the users in question are accessing the server through the uplink leading either to another switch or to a router. The troubleshooting scenario is based on a user complaint that communications with the server are “slow.” This problem description tells the support staff almost nothing. If instead of troubleshooting there is a security breach being investigated for

possible forensic and legal uses, there are additional considerations regarding the infallibility of the technique.

Note: *Information related to multiple techniques will be described in association with the method where it fits best. Much of the information in this discussion applies to methods other than the one it is described for, and this information may account for trivial to fundamental differences in your results.*

Method 1: Access the switch console

The switch configuration is available via multiple means, including:

- Logging in through a TELNET session
- Logging in through an SSH session
- Logging in through a web session
- Logging in through the serial port of the switch.

A variety of runtime troubleshooting aides are available from some switches, though the feature set for these troubleshooting aides is quite different depending on the vendor and switch model. Advanced operating system commands permit more detailed




Figure 18: Console Access.

examination of the forwarded traffic, but not in a user-friendly format. Considerable experience and theoretical knowledge is required to obtain benefit from several of these features.

Pros

This is a great way to troubleshoot, as evidenced by the popularity and prevalence of this technique. Considerable varieties of network problems now revolve around switch configurations and actions taken by the switch in accordance with those configurations.

Gaining access to the switch console is nearly always possible via some method. The increasingly ubiquitous presence of wireless service and data services from cell phones has made it possible to manage a network from almost anywhere on the planet. If the network management system is configured to send problem notifications to mobile devices, the problems reported may be investigated immediately.

If the problem is related to the configuration then console access will permit problem resolution.

Cons

Senior network support staff or others with the password to the switch appear to rely upon the configuration of the switch so heavily during the troubleshooting process that no other option is considered until this method has utterly failed to produce a resolution. Neglecting the other options may delay problem resolution and increase frustration. Not all problems will be revealed or solved via console access.

Routine console access commands reveal general utilization levels, but little or nothing about specific activity or root cause of a protocol failure. Furthermore, the information readily available via console access indicates what should be happening, but not always, what is happening and may not reveal misbehavior on the part of the switch. Whether the operating system of the switch has bugs, and whether the configuration is incomplete may not be evident from the configuration listing. In some cases, the configuration defaults are not revealed by dumping the configuration to the screen. Only the changes to the default configuration are shown, and it may be that a default setting is the source of network performance issue on your unique network.

Configuration data is useful in guiding troubleshooting efforts to see if the switch is operating as expected. However, configuration and performance validation requires one or more of the other switch troubleshooting methods.

For sensitive areas of the network, console access may not be permitted remotely, or may not be available from outside of a configured group of allowed addresses. Frequently the help desk and lower skilled support staff are not provided with the password, and thus are not allowed access to the console. Support staff with console access are often not involved in routine maintenance and troubleshooting activities. Consider how you would identify or solve many of your recent network performance issues if console access was completely denied to you.

Method 2: Connect to an unused port

The simplest approach to troubleshooting involves attaching a monitoring tool, such as a protocol analyzer, to any unused port on the switch.

Connecting to an unused switch port then allows the monitoring tool access to the attached broadcast domain without disrupting service anywhere. The attached tool has the same access to the broadcast domain as any other station.




Figure 19: Monitor from any open port.

It is not possible to determine if the unused port selected is in the same VLAN (broadcast domain) as the problem unless you have console access or network documentation that is atypically detailed and up-to-date. Also, keep in mind that even if your network documentation was current as of “now”, the problem you are looking for may be a configuration or cabling error that invalidates the documentation you are relying upon.

Passive monitoring

Pros

Passive monitoring requires no configuration or other effort. If you are patient, you will eventually see traffic from nearly all network devices attached to the broadcast domain. This is possible because of address table aging as well as basic bridge behavior.

Cons

Passive monitoring suggests that the monitoring device does not transmit. If no queries are sent, then two situations exist simultaneously: the switch never learns the MAC address of the monitoring device, and no queries means no responses. You should expect to see very little traffic under most normal networking conditions, and the traffic seen will mostly consist of broadcast service announcements or broadcast-based protocols such as ARP.

Connecting such a tool to the network would be an excellent method of obtaining a starting point for a security breach of the network, and the randomly received unsolicited traffic may itself present a security breach. Secure networks should investigate instances of switch ports having active link, but no associated MAC address (or an unknown MAC address).




Figure 20: Switches forward traffic between the source and destination port.

In Figure 20 very little traffic reaches the monitoring tool. The monitoring tool may see a few frames per second instead of the thousands per second that may be passing between the stations and the server.

Active monitoring

Pros

Active monitoring, or more correctly a mix of active and passive monitoring, is an excellent method of rapidly obtaining a list of stations on the broadcast domain as well as a good idea of what services may be available from those stations.

This method also helps to identify many common networking problems, including: duplicate IPs, mis-configured stations, and rogue DHCP servers or routers.

Cons

Actively soliciting traffic or interrogating the broadcast domain is useful for network discovery and for finding other classes of problem, but will not aid in resolving most slow connection problems.

Active discovery finds hosts but not reveal what they are doing.

Background for Method 2

Bridge behavior

A switch (which we are viewing as a multiport bridge) will only forward a very tiny amount of the traffic to the monitored port. This is appropriate behavior on the part of a bridging device, since it is designed to prevent unnecessary traffic from reaching ports where it does not belong. There are several bridging techniques. While possible with some switches to enable and configure other techniques, the default for Ethernet is transparent bridging.

In a transparent bridge, it is up to the bridge to discover where the destination station is located. This is accomplished by flooding unknown destination frames to all ports except where the frame originated. As soon as the destination station responds to a query then all bridges involved will know which port subsequent traffic should be sent to, and only that port. This is true for a single switch as well as each separate switch when attached in parallel or in a hierarchy of some sort.

The simplified method for this is for the switch to record the source MAC address for any traffic, and associate that address to a specific port. It does not matter if that port is servicing a single station, a hub, or an uplink to another switch. All traffic destined for that particular MAC address is forwarded to that port. Until the switch has observed a frame bearing a given MAC address in the source field it will not have a table entry to reference for traffic bearing that MAC address in the destination field, and all such destination traffic will be flooded.

Since there is no guarantee that the MAC address will remain associated with that port forever, there is an aging period for each entry in the bridging table. When an address is aged out of the table, the next frame sent to that discarded table entry MAC address will be flooded to all other ports again. The aging period value is affected by the switch vendor defaults, the switch configuration, static table entries, etc. Roaming wireless stations rely on aging and other techniques for rapid forwarding table reconvergence.

In review: traffic forwarded to the monitored port in Figure 20 will consist almost entirely of broadcast or multicast traffic, with a few frames resulting from unknown destinations appearing sporadically. These occasional frames are probably the result of aging of the

bridge forwarding table, or from broadcast-based protocols like ARP, and not often from truly unknown destinations. Many unwary technicians have seen the traffic distribution (nearly 100 percent broadcasts) but failed to notice the exceptionally low utilization level. This results in the incorrect diagnosis of a broadcast storm present, or that their network is experiencing unbelievably high broadcast rates as a part of normal operation.

Access control

An increasingly popular form of access control is the 802.1X protocol, which supports various authentication methods for gaining access to the host network. Each station connecting to the switch must authenticate before it is joined to the broadcast domain. There are many ways this can be implemented, including: fully locking the station out, placing the station in a “holding” state on an isolated broadcast domain until the authentication challenge is satisfied, and placing the station on an unsecured broadcast domain with Internet access but no local access unless it is authenticated. Implementation of 802.1X will result in passive monitoring challenges that include, but are far from limited to:

No traffic at all. No traffic is forwarded to the port by the switch until the attached station issues an authentication request, and then only authentication traffic will be observed until the station satisfies the authentication requirements.

Very little traffic. Only a limited subset of traffic may be observed until the station is authenticated. The specific types of traffic that may be seen falls into the “it depends” category, and will be different for each network, and possibly for different connection points within the network. A very limited list of the many possibilities includes:

- The switch may issue an authentication challenge, but nothing else.
- Spanning tree traffic may be seen.
- Inter-switch communications may be seen (such as LLDP).
- Flooded “unknown” destinations may be seen.
- Broadcast and multicast traffic may be seen.
- Traffic from within a quarantine broadcast domain may be seen.

However, despite any leakage the attaching station is not permitted to request any traffic from the protected network. The specific switch configuration may permit access to a quarantine network that permits visitors to access the Internet, but blocks access to network resources within the protected network.

Duplex and auto-negotiation

The default state for virtually all new Ethernet interfaces today is for Auto-Negotiation to be enabled. This is good. Many vendors describe it as bad, primarily because the design and support engineers do not have a solid understanding of how it works.

The likelihood of Auto-Negotiation actually failing to produce a viable link is very low. Some implementations do not operate correctly, but in almost all cases the vendor has learned about the problem and a new build of code is available to resolve the problem. If they are not aware of the problem, they will be very interested in working with you to resolve it. Review IEEE 802.3 clause 28 for specific details, however a very simplified description of Auto-Negotiation follows.

A negotiating station will send a handshake signal called an FLP (Fast Link Pulse). This is composed of a burst of normal link pulses common to 10BASE-T. The FLP defines the capabilities of the

negotiating station that sent them. If the link partner is also sending FLPs then they will compare the offered capabilities and will select the highest performance match in offered capabilities, then change to that link technology and begin communicating over it.

If the negotiating station does not detect FLPs from its link partner it will then attempt to detect the link partner's transmission speed. Speed detection is virtually always successful as long as both link partners support a common speed.

If the negotiating station did not detect FLPs from its link partner, it is required to choose half duplex for the connection. A negotiating station will not try to detect the link partner's duplex setting when FLPs are not received. This is the cause of most problems related to duplex problems. Far too many support engineers wrongly believe that the negotiating station will detect the duplex setting of the fixed-setting link partner.

Another misconception is that a duplex mismatch will cause a link failure. The link will experience errors if the duplex does not match, and the visible symptom will be slowness, but it will still pass traffic. If the duplex mismatch appears between two switches it is possible to review the reported interface errors and infer the duplex setting of each end of the link just from the nature of the reported errors.

From a monitoring and security perspective, the impact of a duplex mismatch is slight but still present. If the duplex mismatch is on the interface used by the passive monitoring tool then there probably will not be any perceived problem at all simply because it is not transmitting. If the duplex mismatch is on a link where both link partners are actively transmitting, then two factors come

into play. The link will experience lost frames due to errors appearing late in the transmission, and are therefore not retransmitted by the MAC Layer. These lost frames will require higher layers to recover from any problem that may result. If enough frames are detected as normal collided frames and the MAC Layer attempts to retransmit them, then the buffer will eventually back-up and start dropping frames intended for transmission on that interface. Lost frames could cause problems with problem analysis and forensic purposes, particularly as it is almost impossible to discover if frames were dropped, how many frames were dropped, or when.

The default state for some switches is to disable any port involved in too many collisions, regardless of whether or not it negotiated to half duplex. This results in normal 802.3 medium arbitration collisions on a shared media collision domain causing the switch to block or disable the affected port. As an example, Cisco calls this parameter “errdisable”.

Method 3: Configure a mirror or span port

Most managed switches permit configuration of one or more monitoring ports. These features allow traffic from a selected port or ports to be copied to the monitoring port. This technique is usually referred to as port aliasing, mirroring, or spanning.




Figure 21: Configure a mirror or span port for monitoring

The ability to copy or mirror traffic to a fixed or selected output port on the switch is provided by most switch vendors. Older switches had a specific port that could be configured as this special monitoring port, but most new switches may be configured to use any port as the output port for monitoring. Some permit multiple output ports while others support only one. A few models allow the monitored traffic to be forwarded from another switch, such as Cisco's RSPAN feature, though there are capacity and performance issues related to this feature.

The features offered by the switch for this type of traffic monitoring vary by vendor and switch model, but the basic functionality is usually the same. Traffic from a selected port is copied and sent to the monitoring port for analysis.




Figure 22: Logical effect of configuring a mirror or span.

The mirror port is often a listen-only (egress) port, though a number of vendors permit configuring the port to be bidirectional (ingress/egress). Configuring a mirror port on the switch permits the monitoring tool to see a copy of the actual traffic between our reportedly slow user connection and the server. The traffic source for the mirrored could be any other port on the switch, including an uplink port. The mirror source could also be several other ports on the switch, or even one or more VLANs.

Pros

Port mirroring is one of the most common and effective methods of troubleshooting a switched network. This technique permits the monitoring tool to observe traffic between two or more stations that passes through the switch. The most common usage is in association with a protocol analyzer.

Cons

The implementation of this technique varies between vendors, but there are several common choices. Note that in almost all cases the forwarding technique employed by the switch will also be used to filter data sent to the monitor port. This means errors are usually filtered by the switch, and do not appear on the monitor port. For troubleshooting purposes, port mirroring can sometimes be quite ineffective because entire classes of problems are concealed by the switch through some type of filtering.

Configuration of the mirror must be performed from a console session. This often involves bringing a PC along with the monitoring tool so that the switch can be reconfigured as required for troubleshooting. Lower level or third party technicians are often not in possession of the password, and there is a real risk that improperly configuring the mirror will result in a network disruption.

It is easily possible to oversubscribe the capacity of the output port. This leads to inexplicable missing traffic in the output data stream, and an underreporting of the monitored link utilization. As more ports are included in the mirror the chances increase that the capacity of the output port will be exceeded. An extreme instance of oversubscription is possible with switches that permit an entire VLAN to be mirrored. For data analysis and forensic purposes, it is quite important to know whether all traffic in the conversation was included for analysis, or that potential loss of some traffic has been accounted for. At the same time is important to understand the input capacity limits of the monitoring tool. For example, software protocol analyzer often cannot store anything close to line rate traffic.

Background for Method 3

Oversubscription

Output capacity on the monitor port is an important consideration. The output port has a TX and RX path. It was already noted that the TX path (ingress) from the monitoring device back to the switch might be blocked by the switch as part of the mirror configuration. Whether or not the TX path is blocked (whether the port is bidirectional or not), the RX path from the switch to the monitoring device is capacity limited. If you are mirroring a full duplex port of the same speed as the mirror output port, the switch may easily drop traffic without notifying you. In this regard, it does not matter whether the monitoring device is connected at half or full duplex; the inherent limit to the output path is the same.

Refer to Figure 23 where the traffic associated with a server connected to a switch at 100 Mbps in full duplex is mirrored to a monitor port. At full duplex both the TX path and the RX path are able to support 100 Mbps of traffic, for an aggregate throughput




Figure 23: Output capacity is limited on a mirror port.

potential of 200 Mbps. If you seek to mirror that traffic to another 100 Mbps port, you can only use the TX path from the switch to the monitoring tool. The amount of mirrored traffic is therefore limited to a maximum of 100 Mbps in this example. Despite a small amount of buffering provided by the switch to accommodate the inherent bursty nature of network traffic, traffic on the server's switch port which exceeds 50 percent of the combined capacity of the monitored full duplex port will probably be dropped.

If multiple ports are mirrored to the monitoring port, then the potential for this problem is correspondingly greater. Since most switches operate at far below their capacity, the problem may not be noticed right away.

The situation may be mitigated by connecting the monitoring device to a higher speed port, one that has the native output capacity to accept all of the mirrored traffic. If the mirror output port in Figure 23 were a Gigabit port instead of a 100 Mbps port, then the aggregate 200 Mbps traffic potential would be easily accommodated.

Switch forwarding techniques

Awareness of switch forwarding techniques and what will be forwarded once the decision is made seems to have fallen out of the pool of common knowledge. Assuming that you chose the correct port to span during a troubleshooting or forensic incident,

it is unlikely that you will see any MAC Layer errors that may be present on the spanned port. Most switches sold today default to, or are only capable of, the Store and Forward technique. However, many of the legacy switches deployed may offer or default to low-latency forwarding techniques. It is impossible to know which forwarding technique is employed without doing some research, and possibly some testing.

There are three common forwarding techniques employed, though other names have been used to describe them:

- Store and Forward (traditional OSI Layer 2 Bridge behavior)
- Cut-Through (forward after the destination MAC address is known)
- Modified Cut-Through (forward after 64 bytes have been received – this is equal to the 10/100 Ethernet half-duplex timing concept of slot time, and represents the cutoff for when a legal collision may be detected)

The two low-latency forwarding techniques lost popularity after only a few years. This is probably due to the comparatively slight performance improvement coupled with the increase in troubleshooting difficulty that is associated with low-latency forwarding. At least one additional combination technique existed, which has sometimes been called Error Sensing or Adaptive. In this last technique uses one of the low-latency techniques until the error level exceeds a fixed or configurable threshold, then it changes to Store and Forward. The trick is discovering whether the error went away, or the error is bad enough that the switch changed techniques. You may have paid extra for an intermittent error.

If one of the low-latency forwarding techniques is in use then any error observed may have come from the local collision domain, or could have come from anywhere within the attached broadcast domain – even several switches away.

If store and forward is used then you may safely assume that MAC Layer errors stop at the switch port (see Figure 24). If a low-latency forwarding technique is employed then a detected error could have originated from anywhere inside of the broadcast domain, not just “this side” of the switch port. This substantially alters troubleshooting considerations and assumptions.

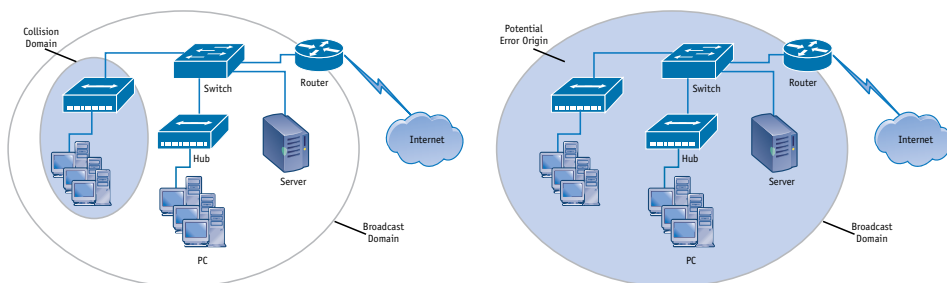


Figure 24: Potential error forwarding with Store and Forward (left) and low-latency (right).

At OSI Layer 2, a mirror output port will be subject to the same bridge forwarding technique as the rest of the switch. MAC Layer errors are almost never forwarded to a mirror output port.

The discussion so far has only considered OSI Layer 2 forwarding. The current marketplace offers a very broad and diverse feature set for switches, including: forwarding at OSI Layer 3, Layer 4, and Layers 5-7; load balancing; rate limiting; content-based forwarding

techniques; as well as proxy services and special buffering, filtering, and security options. As these higher-layer features are vendor and model dependent, all are outside of the scope of this basic description. Study the vendor documents to learn how the features operate in order to understand how to troubleshoot them. In many cases, troubleshooting these features will require observing traffic before and after the switch simultaneously. Begin by using a protocol analyzer to examine traffic and match the traffic to descriptions in the vendor documents. Once you can identify and explain correct behavior to someone else, you are ready to look for incorrect behavior.

Method 4: Connect to a tagged or trunk port

The tester may be connected to a VLAN trunk port, or a port with one or more VLANs associated to it. This is similar to a span or mirror port, so all of the pros and cons associated with that technique apply. Additionally, the tester must be capable of interpreting the VLAN tag or tags, and/or organizing them by broadcast domain in order to create a useful view of the network. If on a trunk port, the tester may be required to participate in the trunk management traffic, such as Cisco's VTP protocol.

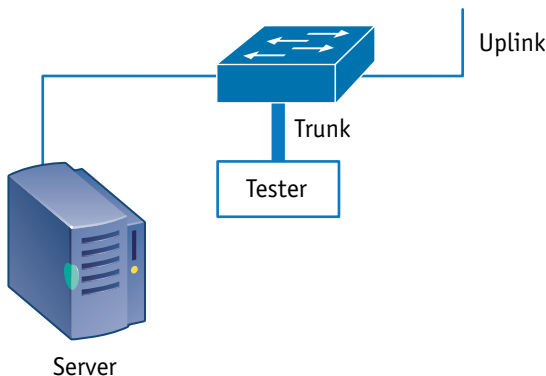


Figure 25: Connect to a VLAN tagged or VLAN trunk port.

Traffic observed on a tagged or trunk port will vary. Some possibilities include:

- The switch may present only VLAN tagged traffic to the port. Stations that are not participating in the VLAN will not be able to use the attached network resources.
- The switch may support both tagged and untagged traffic on the port, supporting both a VLAN trunk as well as local untagged stations. The native VLAN is often untagged on a trunk port.
- The switch may present traffic tagged for multiple VLANs to the port. In this configuration, there is rarely any expectation that an end-station will be attached, though servers may support multiple VLANs on a single connection.

Pros

Utilizing a VLAN trunk port would permit monitoring of a much larger portion of the attached network. Active discovery would benefit greatly by having broadcast domain access to multiple VLANs at once.

Cons

Few monitoring tools can take advantage of multiple VLANs simultaneously. Most are unable to do tagged discovery at all, relying instead on passive monitoring. Furthermore, the problem of oversubscription would be exacerbated by trying to monitor larger portions of the network. The switch should be applying normal forwarding rules to the trunk port, so only appropriate unicast destinations would be crossing the trunk, as well as unknown, aged, and broadcast traffic.

Method 5: Insert a hub into the link

This was perhaps the first switch troubleshooting method, and is still a common method of monitoring a problem related to a single switch port.

Using a shared media hub involves a strategic placement decision. The hub may be placed between switches, or on a client link. In many networks, most traffic of interest will be received or transmitted by a shared resource such as a file server, as shown in Figure 26.

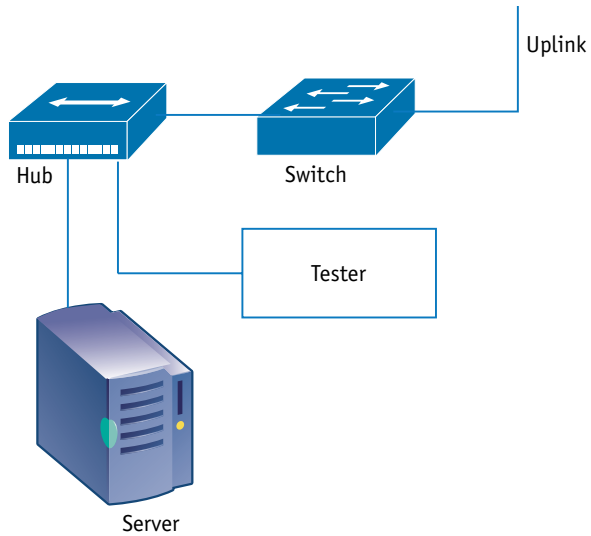


Figure 26: Inserting a hub inline between the attached station and the switch port.

For troubleshooting and forensic purposes, the monitoring tool may remain virtually invisible to the network as long as it does not transmit.

Note: Use of aggregating taps may replace this method simply because it has become almost impossible to buy a shared media hub.

Pros

Using SNMP to learn about traffic and errors is useful, but for good error analysis there is nothing like seeing them with the diagnostic tool directly. Placing a shared-media hub inline between a switch port and another interface allows a monitoring tool to be connected to the same collision domain. This technique enables the analyzer to see all traffic present. Access to all traffic greatly assists the network support staff in diagnosing a wide range of problems, including: user login failures, poor performance, and dropped connections.

This is one of the best troubleshooting situations for an issue related to a single station, since almost nothing is hidden from the monitoring tool.

Cons

This approach is impractical in most situations, particularly where there are multiple servers to be monitored. Where do you locate the hub – on all server links? If you choose to move a hub around as needed, are you prepared to interrupt the network long enough to install the hub? This delay is often long enough to cause dropped connections. Additionally, network resources such as servers may be connected via a technology or connection speed that your monitoring tool may not support.

Inserting a half duplex hub between two full duplex stations may significantly reduce the overall throughput of the link, introduce errors, and generally confuse the symptoms of the original problem or monitoring situation unless you are very aware of how your equipment works (switches, hubs, and testers).

Three common problems take place when a hub is inserted into the link:

- If the link was originally forced to full-duplex, then a new problem has been introduced into the link, as the hub will be operating in half duplex.
- If the “hub” used was actually a partial or full bridge and you have not gained anything.
- The link already had one or more hubs and the Ethernet architecture is now illegal, which causes late collisions.

Even where adding a hub does not introduce new problems, the results will depend on the monitoring tool used. Virtually all software protocol analyzers have some level of blindness to activity on the attached collision domain. Many basic software protocol analyzers will only see traffic that is of a legal size, because the NIC driver discards all traffic that is not a legal and fully formed frame with no errors. Some software protocol analyzers take advantage of features in certain NICs that permit the software protocol analyzer to see some level of errored traffic. To see all errors on the collision domain usually requires a hardware protocol analyzer or other tool with custom circuitry operating at the level of the Ethernet PHY (the part of the circuit that converts binary data to the correct signaling for the medium).

Background for Method 5

Hub architecture limits

- 10 Mbps Ethernet – up to four hubs may be used in series between any two distant PCs.
- 100 Mbps Ethernet – one or two hubs in series between any two distant PCs.

- Class I hubs may or may not be marked as such, but only one such hub may be used on a collision domain.
- Class II hubs are usually marked, and there may be one or two used on a collision domain, but not more.
- 1000 Mbps Ethernet – permits a single hub to be used on a collision domain, but they are simply unavailable for purchase. There are effectively no shared media Gigabit Ethernet hubs manufactured at this time.
- 10,000 Mbps Ethernet – 10 Gigabit Ethernet does not permit half-duplex operation at all, so there are no hubs for 10 Gigabit Ethernet.

Switch architecture limits

There are no hard limits to the number of bridges (switches operating at OSI Layer 2) which may be placed in series or parallel. Your architecture design will affect network performance though. The two most significant issues relate to broadcasts.

- If two bridges are operating in parallel, and both paths remain open, then the first broadcast frame seen by one will cause an instant broadcast storm. Bridges are required to forward all broadcasts to all ports except the port on which it was received. The next bridge does the same, and the original broadcast returns to the first port via the parallel path. Some switch configurations offer a “trust me” setting where you effectively promise the switch that you will not create a parallel path, and in return the switch will permit a new connection to access the network almost instantly (such as Cisco’s “portfast”).
- Even if Spanning Tree or some other mechanism is used to eliminate parallel paths, the greater the number of bridges participating in a single broadcast domain, the greater

the amount of appropriate broadcast traffic will be seen. Broadcasts represent a useful and necessary function in the network, so they cannot be eliminated. Placing too many stations on a broadcast domain will raise the level of background broadcast traffic to the point where it becomes a noticeable cause of poor performance. Each station on the broadcast domain is required to process each broadcast frame, which interrupts whatever it was supposed to be doing for you. The spanning tree *Max Age* timer will limit spanning tree enabled bridged networks to a maximum diameter.

Method 6: Place the tester in series

Inserting a monitoring tool in series in the link should avoid the problems related to inserting a hub in series, as the link could be operating in full duplex and not experience problems.




Figure 27: Inserting a monitoring tool in series.

Pros

Unless there was a parallel path to the station being tested (the server in Figure 27), such as wireless, there should be nothing the tester failed to see. Depending on the features of the monitoring tool, this technique could be very useful for any type of situation related to a single station or link. If the monitoring tool were inserted on a link going to the Internet, it could be used for verifying firewall effectiveness or for gathering forensic information for an investigation.

Even passively, the monitoring tool can track which link partner is the source of any MAC Layer errors, and can monitor requests made by the attached station to ensure that responses from the network returned. Most problems related to an inability to connect to a server or service should be easily identified.

Cons

The technique has significant potential problems.

If encryption is used on the link, then even being in series will not allow viewing of higher-Layer data.

Ethernet does not permit shared use of a single cable other than 10 Mbps coax. If the inline monitoring tool did not have an additional management port, either the monitoring tool would have to perform bridging in order to transmit, or it would have to be passive.

In-line monitoring tools are generally quite expensive, and are not widely deployed as a result.

Method 7: Place a Tap inline on a link

The terms tap and splitter are often used interchangeably, though splitter usually applies to fiber optic links.

On a fiber optic link, the splitter is rated by how much light is taken from the primary path and redirected to the monitoring path. Typical splitter ratings include ratios of 80:20, 70:30, or even 50:50. Using the first example, 80 percent of the light continues through the splitter to its original destination, and 20 percent of the light is redirected to the monitor output. It is important to match the tap to the cable type. For example, a multimode capable

tap cannot be used on a singlemode link. You should not use a multimode 50nm tap on a 62.5nm link. Most non-aggregating fiber taps are unpowered.

For copper cable links, depending on the complexity of the encoding process, a tap may be required to interpret the data much like the original destination receiver. Then the tap would retransmit the received data to the tap output. For this reason, it is important to match the tapped path to the capabilities of a copper tap. Many older copper taps are able to tap a 10/100 Ethernet link, but not a Gig Ethernet link. Some newer taps can only tap a Gig Ethernet link. Still other (newer) taps can tap all three speeds (10/100/1000). Cop-

per taps are usually powered, and will normally maintain the tapped link if power is lost. When power returns there may be a brief service interruption as the tap resets the link via relays.




Figure 28: Inserting a tap inline.

There are two significantly different types of tap available: the traditional or standard tap, and the new aggregating tap.

Standard Tap

Using standard taps means that the monitoring tool either sees the request or the response, but not both. In order to see both the request and the response the monitoring tool needs two analysis input ports, one for the TX path and one for the RX path.




Figure 29: A standard tap presents data from one direction on each output port.

Aggregating Tap

Aggregating taps provide both the request and response by design. Aggregating taps often also provide configurable ingress capability (the monitoring tool can send and receive through the tap output port). Aggregating taps may also offer more than one output port, so that two or more test tools can receive the same output traffic.

Pros

Tap use presents many advantages over configuring span or mirror ports, and avoids the problems related to inserting a hub in series. Due to their comparatively low cost, they may be placed inline on critical links and simply left unused until the need arises.




Figure 30: An aggregating tap combines data from both directions onto one or more output ports.

For a junior support person and for forensic purposes it is a good method because it is not necessary to have the password for the switch. The tap can be inserted in series quickly with a very short loss of link state on the affected link. Once installed, a tap also permits monitoring tools to be installed and removed at any time with no risk of network disruption, and are therefore good choices for both troubleshooting and forensic purposes.

Most taps are output (egress) only, and an attached monitoring tool will be invisible to the network. Many of the new aggregating taps may be configured to support input (ingress) as well as output, which thereby permit the monitoring tool to interrogate the network or respond to management queries by injecting traffic into the network through the monitored link. Both configurations can be viewed as having positive aspects, depending on the goal.

A standard tap provides a copy of all traffic on the link, including any errors that may be present. A standard tap does not experience oversubscription because it has a separate output for each direction (TX and RX). This permits the monitoring tool access on par with a hub in series without the drawback of forcing the link to half-duplex or creating a duplex mismatch. A standard tap is always output (egress) only.

Aggregating taps have a growing feature set, including the ability to filter specific traffic from the link and only send selected traffic to the monitor port. For high-speed links, this can be very useful, as it reduces the amount of traffic that must be examined by the monitoring tool or technician. Aggregating taps are new on the market, and are expected to undergo considerable feature evolution in the short term.

Most aggregating tap product lines offer single or multiple output ports, depending on the model. Additional output ports permit a variety of tools to view the same monitored data in parallel, usually for different purposes.

In a switched environment, an aggregating tap is overall the fastest, easiest method of gaining access to the data passing through a selected link. Aggregating taps are the new “hub” for troubleshooting.

Cons

One of the more significant issues related to using taps is the signal loss inherent with applying any copper or fiber inline tap. This loss of power clearly implies that if a link is already suffering from cable faults or excessive distance, the tap could easily cause the link to fail by taking too much signal from the primary transmit path. A splitter can easily cause a 3 dB loss in power on the tapped link. Some transmitters are more robust than others are, so even if installation of a splitter at one end causes the link to fail, it may still be possible to install the splitter at the other end of the link without causing it to fail.

Copper taps cause similar signal loss problems, as some of the signal is needed by the tap in order to read the passing traffic. For copper this is the equivalent of additional attenuation, and also may cause the tapped link to fail during the installation process if the link being tapped is very long or is already experiencing cable problems. Copper taps require power, as the signal is recovered and retransmitted to the monitor port. A high quality copper tap will not disconnect the tapped link if power is lost to the tap. Instead, there may be a brief interruption as relays reset the monitor state to

a pass-through mode. Alternatively, there may be a brief interruption as the relays reset the monitor state when power returns.

Direction counts, and problems and/or delays will be encountered if the fiber tap (and the occasional copper tap) is installed “backwards” resulting in a lack of output to the monitor port(s). The challenge is that a link requiring a tap is often highly utilized, and if the improper installation is not detected immediately, it could be weeks before the next maintenance window is available to reverse the connections.

The latest generation of aggregating tap uses a bridging function to combine the data streams from the RX and TX data paths. As a modified bridge link:

- Aggregating taps are subject to the oversubscription problem described in Method 3 (see Figure 23).
- Bursty traffic, which is common to networking, may or may not exceed the buffer inside the aggregating tap. The presence of the buffer may obscure some data loss, and would be a very significant concern for forensic analysis. Most reactive troubleshooting would not be affected greatly by data loss through oversubscription or by having a burst exceed the buffer, as the cause of a reactive troubleshooting problem is usually related to the majority of the traffic present.
- Filtering applied by the aggregating tap may be discarding the traffic of interest. A filter applied for a prior event may still be active, which would affect both troubleshooting and forensic activities. In addition, the act of filtering may cause the aggregating tap to drop additional frames due to CPU overhead in the tap.

MAC Layer errors are dropped because bridges do not forward errors. Everything else is flooded to the monitor port. This is perceived as a significant drawback by the industry since standard taps forward errors, and aggregating tap vendors are working hard to correcting this limitation.

Aggregating taps often have very limited bridging capabilities, and may not forward frames larger than a maximum sized VLAN tagged Ethernet frame (1522 octets). This means that 802.3as compliant frames which may be as large as 2000 octets may not be forwarded, and so-called jumbo frames may not be forwarded.

Typically, only aggregation taps that support “injection” (ingress traffic from the monitoring tool) have support for power over Ethernet (PoE), so inserting a standard tap will often turn off a PoE powered station. This is most noticeable when troubleshooting VoIP and wireless access points, which are likely to rely on PoE.

Since copper taps are usually active it is possible for a link to fail on the network, but the connection to the router remains up. That is, the link on one side of the tap is down, but the tap is holding the link to the router up. This will cause the router to believe that the failed link is still up, and it will forward traffic into the disconnected link. The problem will persist until someone arrives to troubleshoot the loss of connectivity, or forces the router interface into an administratively “down” state where a redundant backup link takes over. Fiber optic taps are passive, so this is not a problem.

Method 8: Use SNMP-based network management

SNMP was created in order to learn what was going on within distant areas of the network without having to monitor constantly or to be in each physical location. The SNMP management protocol permits long-term trend analysis as well as short-term detail analysis. SNMP is mostly based around a query/response model, which implies that a management station must be constantly querying the network in order to discover problems. To permit the network to inform the management station of a problem without waiting for the correct query to reveal the problem SNMP defines the ability to send an unsolicited response, called a trap. The trap function permits a managed device – the SNMP agent – to notify the management station that predefined conditions have been met or exceeded, and

attention is warranted. Receipt of some or all traps may cause the network management station to invoke user notification routines, such as e-mail and pager alerting.




Figure 31: Using SNMP to query network infrastructure devices remotely.

SNMP is probably the most common method of monitoring

a switched network today. The SNMP console does not have to be anywhere near the monitored device as long as there is a routed path to the target and security configurations permit the console to communicate with the agent in the switch.

Because switches do not routinely forward errors, using SNMP is perhaps the best method of locating ports experiencing errors. The switch may not forward an error, but it is certainly aware of the presence of errors. There are a variety of MIBs available from most switches that support SNMP. A MIB (management information base) is a dictionary of available queries accompanied by possible responses and the meaning of each response. Each MIB supported offers the management console a different or more detailed view of network conditions in or around the monitored device. In addition to private MIBs, which typically have customized support for each switch and level of operating code, the Standards or RFC-based MIBs can be used very effectively to monitor a switched network. In increasing order of detail, the following MIBs are useful for troubleshooting, though many others are useful too:

- RFC 1213 – MIB II
- RFC 1643 – Ethernet-Like Interface MIB
- RFC 2021 – RMON 2
- RFC 2819 – RMON Ethernet

RFCs may be updated or enhanced following their introduction, so always check the latest RFC index for updates. For example, RFC 1213 is updated or enhanced by at least five newer RFCs (2011, 2012, 2013, 2358 and 2665). In addition to the MIBs defined by these RFCs, which contain excellent information on utilization and errors, the bridge MIB (RFC 1493, 1525, and 2674) is very useful for troubleshooting. Security is a concern when using SNMP to monitor a network. If SNMP agents are unrestricted, then potentially anyone anywhere could be monitoring activity on your network or modifying your switch configurations. SNMP is often enabled by default with a very common default password when the switch or any other SNMP agent is sold. SNMP passwords are called

community strings, and are both case and punctuation sensitive. Community strings are transmitted in clear text, which in itself creates a security risk. SNMP version 3 offers authentication and encrypted communication to address this exposure.

At a minimum, the default community string should be changed immediately. SNMP agents may be configured to respond to different community strings for different levels of access, to queries from a specific subnet and no other, to queries from a specific IP address and no other, and many other configurations. The routers that provide a path to those SNMP agents may impose restrictions on SNMP. Firewalls may block SNMP entirely. If you are able to reach the agent using SNMP, the agent still has to support the MIB you are querying. Most vendors support the standard MIBs adequately; however, some vendors do not support standard MIBs as well. In some cases, it is necessary to upgrade the operating system on the switch before it is capable of supporting a desired public or private MIB.

There are many reasons why your switch may not respond to a specific SNMP query. Once the proper MIB is available any access problems are resolved, SNMP is a very useful tool for monitoring and trending.

Pros

Using network management systems to automate network monitoring is an excellent method for learning about changes in traffic patterns over time, as well as investigating network activity during troubleshooting or forensic events. It is possible, using SNMP, to obtain almost any information about your network, providing that the Agent supports the MIB.

For best effect, the network management system should be adjusted over time so that normal behavior is counted as such, and any abnormal behavior (excessive or missing traffic) is flagged for operator attention.

The specific resources available for detailed investigation will depend on capabilities built into the network infrastructure or deployed specifically for monitoring and diagnosing key network links or resources.

If the right resources are in place, a protocol analyzer capture file may be started automatically because of observed network behavior. Intrusion detection systems may be configured to monitor for symptoms that suggest that the network is under attack from external or internal agents.

Cons

With SNMP, you can get almost anything, including a packet capture file if the Agent supports it. However, most do not. Therefore, you are usually limited to knowing who used which protocol with to whom. If there is a protocol problem, or a problem with timing then that cannot be diagnosed unless a capture is available. Most devices reported as supporting RMON actually only support four of the nine groups. This limited subset is referenced by several names, including RMON Lite.

SNMP is a lower priority activity than forwarding traffic. If the Agent becomes busy, it may suspend gathering of SNMP statistics during some or all of the high-traffic event. This is the statistical equivalent of dropping frames. In many newer routers, the front-end ASICs are handling the routine traffic, and only unusual situations reach the CPU of the router. Using an SNMP console to query a

router may easily elevate the CPU utilization periodically to between 10% and 100%. This can be very alarming to the support staff. Multiple SNMP consoles querying a single device at the same moment have the ability to cause it to crash entirely, depending on how the device was engineered.

Despite being a traditional mainstay for discovering what normal is, or for locating abnormal behavior, many network management systems are so complex that they are never properly and completely configured, or require enough daily attention that the network support staff cannot adequately maintain them without people whose sole function is to operate the network management system. In many businesses, this quickly leads to a situation where the network management resource is abandoned, or only used for specific short-term situations. On an annual basis, the cost of keeping the network management resource current may approach or exceed the purchase cost.

Most networks use the common data paths for SNMP monitoring, and do not have alternate means of reaching distant network segments. If a key infrastructure path goes down, SNMP can only agree that the distant segment is unreachable, and cannot help troubleshoot unless an alternate path exists.

Background for Method 8

Intrusion detection examples using SNMP

Intrusion detection systems may be configured to monitor for symptoms that suggest that the network is under attack from external or internal agents, looking for such thing as:

- ipReasmFails (1.3.6.1.2.1.4.16): The number of failures detected by the IP reassembly algorithm. Monitored at hosts to determine possible attacks as well as network delivery problems.
- tcpAttemptFails (1.3.6.1.2.1.6.7): The number of times TCP connections have made a direct transition to the CLOSED state from either the SYN-SENT state or the SYN-RCVD state, plus the number of times TCP connections have made a direct transition to the LISTEN state from the SYN-RCVD state. This can be an indication of incoming attacks.
- udpNoPorts (1.3.6.1.2.1.7.2): The total number of received UDP datagrams for which there was no application at the destination port. This counter can be indicative of reconnaissance against your network.

Knowing which SNMP MIB you are using.

SNMP utilizes a great many MIBs – some of which are RFC based, and others are vendor and model specific. Failing to understand exactly what was asked may lead to misunderstandings. For example, in Figure 32 a server is queried using three different MIBs. In each case, the query is analogous to “how busy are you?”




Figure 32: Various queries related to current utilization.

The first query asks how busy the CPU is in that server. This is entirely dependent upon what applications are running on the server at that moment, and has little or nothing to do with network activity.

- RFC 2790 [Host Resources]: hrProcessorLoad
(1.3.6.1.2.1.25.3.3.1.2)

The average, over the last minute, of the percentage of time that this processor was not idle.

The second query asks how much traffic passed through the NIC in that server. This is entirely dependent upon traffic addressed to the server (including broadcasts), and may or may not be related to how busy the attached network is. For example, the attached network could be experiencing 35% traffic load, but the server is only accepting 7% of that traffic. The SNMP response would be 7%. Note that this query ignores the outbound traffic, which is another MIB OID.

- RFC 1213 [MIB II]: ifInOctets (1.3.6.1.2.1.2.2.1.10)

The total number of octets received on the interface, including framing characters.

The third query asks how busy network segment is where that NIC is attached. This is entirely focused on network activity, of which some of that traffic may or may not be destined for that server. Using the same example, the attached network could be experiencing 35% traffic load, but the server is only accepting 7% of that traffic. The SNMP response would be 35%. If the server is using a standard NDIS driver, it may only be seeing the “good” traffic, and may not be capable of seeing the errors even though the query is for a MIB

that would report the errors. A large number of RMON probes rely upon what amounts to a clone PC with a standard NIC operating with a standard NDIS driver. Many can see few if any errors.

- RFC 2819 [RMON]: etherStatsOctets (1.3.6.1.2.1.16.1.1.1.4)
The total number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets).

If the user is not aware of the exact meaning of each metric it is easy to be led to the wrong conclusion. This is particularly true considering the user interface for a network management system, which often uses red, yellow, and green to indicate status. Color does little to discriminate between different types of “busy”, you must already know the possible sources of the warning in order to interpret the warning properly.

MIB accuracy

Sometimes an SNMP agent’s implementation of a specific MIB is not accurate, and responses to queries are simply wrong. Less frequently, the SNMP Manager software does not correctly interpret a response, or there is a mismatch between the MIB versions in the Agent and Manager. It does not happen that often, but programming errors occasionally result in inaccurate responses. You might have the proper response, but be using different versions of the MIB on the monitored device and the management station. A new version of the MIB may change the meaning of the response. Example: you purchase a used network toaster. After taking it home, you download the latest MIB from the manufacturer’s website. When you put the first slice of bread in to be toasted and query the toaster with your SNMP management station, the toaster responds to your query for the cooking setting with a value of “3”.

The MIB in the old toaster had only three responses: 1="warmed slightly," 2="browned," 3="reduced to charcoal." Your new MIB, intended for the latest model of network toaster, has seven responses ranging from 1="warmed slightly to 7="reduced to charcoal." Due to the mismatch in MIBs, you believe that a setting of 3 means that your toast will be browned instead of burned.

Method 9: Use flow technology

Flow technologies are one of the answers to troubleshooting switched networks. It is very likely that flow technology use will join the other three core skills expected of someone involved in the IT industry (cable test, protocol analysis, and SNMP based network management). Using this technique the router becomes your inline diagnostic or management probe.

Flow technologies keep track of who has been talking with whom, using what protocol, how many bytes and packets were sent by each, and so on, and then a summary report of this is sent to the flow receiver. The amount of data gathered is reduced enormously over a protocol analyzer capture file. Only the summary reports cross the network to the flow receiver where statistics are compiled.




Figure 33: Switches may be configured to export flow summaries to a monitoring station for trending and analysis.

There are many different flow technologies in use today. Some choices include: NetFlow, IPFIX, J-Flow, cflowd, and sFlow.

Pros

Compared to SNMP, flow technologies offer some good trade-offs. The switch is not required to store reports about observed behavior for long periods (comparatively speaking). Flow reports typically time-out and are exported in 30 minutes or less, where SNMP studies may easily be keeping data for well over a day.

Hardware or software probes are not required. Flow data comes from the network infrastructure itself. In most cases, your existing infrastructure can already output flow summaries. Perhaps not the low-end switches near the user connections, but almost certainly, in the network core and on the WAN edges where you are often most interested. A few extra configuration lines and it is running.

The flow receiver could be anywhere in your enterprise network, though some WAN links may be expensive enough that flow receivers located in geographically strategic locations would be preferred. Sizing of the flow reports sent to the flow receiver may be conservatively estimated at between 3-5% of the monitored traffic rate. Depending on the nature of the traffic, it could be less than 1%.

Flow data is send constantly to the flow receiver. SNMP requires that the monitoring station poll the SNMP agent on a regular basis to obtain the data.

Cons

The most widely available flow technology at the time of this writing is NetFlow, created by Cisco. Alternatives are available if you do not have a Cisco infrastructure. RFC 3917 standardized a version of NetFlow as IPFIX, so this should become more available soon.

sFlow is a sampled flow technology defined in RFC 3176. It provides similar statistics on the amount of traffic, and who participated in the conversation. sFlow may be configured to sample every *n*th packet, or randomly. Since the traffic is always sampled, this technique is useful for growth planning, general trend analysis and go/no-go troubleshooting. Packet sampling makes it virtually impossible to report on sequences of packets from an individual transaction. Because of the sampling, this technique may not be suited to some security related activities. Unlike NetFlow and IPFIX, sFlow operates at OSI Layer 2 and will report statistics on non-IP traffic, though this advantage is slight, as IP has become by far the dominant protocol. Some platforms also support NetFlow or IPFIX sampling, though the same reduction in effectiveness applies.

Depending on when the flow ends, or the time-out period takes effect, the flow summary is usually not sent for anywhere between 1 and 30 minutes. It is not quite a real-time monitoring situation, though the constant nature of its reporting leads one to treat it that way.

Flow summaries are usually sent unencrypted, so it is possible to spoof them.

When compared with SNMP the flow summaries are likely to report slightly less traffic. For example, NetFlow summaries will report IP traffic but not other Layer 3 or lower layer traffic. Assuming that the correct MIB is queried SNMP will report all traffic on the interface.

Method 10: Set up a syslog server

Most infrastructure devices support sending syslog information to a syslog server for collection. Syslog is most often used for application and server management, and security auditing.

The level of detail reported is controlled by adjusting the setting for verbosity which is usually set to somewhere between catastrophic only to everything down to trivial (see RFC 3164 where messages types are defined). Messages sent include, but are

not limited to errors and events (login, login failure, process start and stop, etc.) and repetitive routine operations.




Figure 34: Most servers and other infrastructure devices may be configured to export syslog messages.

Pros

Syslog will report errors and events, and routine operations with no further administrative effort beyond a basic configuration in the switch that includes a destination IP address for the syslog server.

Since syslog is sent continuously (as needed) after it is configured, there is an on-going record of activity that is available for immediate or historical review and investigation.

Syslog is arguably one of the best tools for troubleshooting some types of authentication problems.

Cons

Syslog can generate vast amounts of useless data. Trying to sift through the logged messages looking for the source of a problem, or proactively searching for potential problems or security breaches can be particularly daunting. The mind-numbing quantity of unimportant messages which have to be examined and judged while looking for the information of interest has spawned the creation of syslog mining utilities and purchased applications which are able to catalog and group messages, and also provide flexible search features.

Syslog is prone to generating significant volumes of useless information in a short period if the logging level is too broad, or not reporting events that would be of interest if the logging level were too restricted.

Method 11: Use the server (host) resources

Virtually all computers and network adapters come with some sort of diagnostics. Vendor supplied diagnostics typically report most of the things that would affect day-to-day usage and operation.

Computer diagnostics

Each computer manufacturer has some level of hardware diagnostic application that either ships with the computer, or is downloadable

from the vendor web site. For the most part these diagnostics are related exclusively to the operation of the stand-alone hardware, but hardware failures may be directly related to network problems.

In addition, the NIC manufacturer often has downloadable diagnostic utilities to assist with NIC configuration and troubleshooting. Use these utilities to check for speed and duplex, as well as any reported errors. The NIC driver software installed on the computer does not make it easy to find this information, if it is available from the NIC driver at all.

Operating system diagnostics

Operating systems, such as Microsoft Windows and Unix or Linux, have various diagnostic capabilities.

Perhaps the two simplest Windows diagnostics are the *msconfig* utility and the NIC statistics window. Msconfig permits you to review the system configuration, including NIC driver information, and the NIC statistics show how much traffic the system has accepted from the network and how much it thinks it sent to the network. [These numbers usually relate directly to results from the MIB II query from Method 8. There could be much more traffic presented to the NIC than is being accepted, which would be revealed with RMON queries from Figure 33.]

It is difficult to choose a best example of a simple Unix or Linux diagnostic, as this is an Open Source operating system, and there are so many from which to choose. The diagnostics available from this type of operating system range from simple to highly capable and complex tools which have features very similar to the Third Party category below.

Third party diagnostics

Use third party diagnostics (such as a software protocol analyzer), on the station to isolate protocol related problems. There are simple protocol analysis products available from the Internet, as well as comprehensive purchased products that include many built-in report tools and libraries of common fault symptoms (usually called Experts).

Many other types of diagnostics are available, including SNMP based network management and specialized tools for a variety of network related diagnostics, investigation, and trending uses.

Method 12: Use a combination of the above methods

Some networking problems are satisfactorily addressed with a single troubleshooting method. Others require the combined results from two or more methods to properly quantify or isolate the situation.

One example of this might be to use a hardware protocol analyzer to monitor the input and output data paths from an infrastructure device – such as a switch – and then use another method to stimulate the link with some sort of traffic. The results of this test would determine whether the switch or other device had modified the traffic during transit, whether the traffic was being filtered by some security mechanism, by what priority the traffic was being forwarded (important for some types of application, such as VoIP), and what the switch latency is.

Conclusion

A commonly used troubleshooting method is to wait for user complaints. This method should not be discounted due to its simplicity and effectiveness. The user community has a very finely tuned subconscious sense of what the normal performance of the network is. Any perceived degradation of that sense of normal will result in a rapid complaint to the network support center. Once a user complains, you can start the troubleshooting process from his or her connection point. The problem with this method is that it is entirely reactive.

Ideally, the approach used by network support staff should be proactive. Proactive efforts to prevent problems from affecting users may include regularly interrogating each switch, and monitoring the quality of traffic on each switch port – just as any other segment would be monitored on a regular basis. Implementing tactics such as monitoring and trending switch port statistics and using tools that allow you to see inside switches will help to move from a reactive troubleshooting mode to a proactive trouble prevention mode. It is probably impossible to be completely proactive.

Regular training for all staff involved in network support or forensic activities is critical. There has been a disturbing trend over the last few years where theory, network design, and problem investigation and resolution for anything below the Network Layer has been overlooked or ignored because of the move from shared media to single user per switch port network designs. Since problems only affect a single user, it has been convenient to ascribe the problem to coincidence or to the hardware involved being outdated. Outdated in this case is often anything more than two or three years old. Many apparently inconsistent symptoms

or test results are ignored, where additional training would have allowed these symptoms and test results to suggest specific fault conditions or behavior. The lack of understanding of the underlying technology below the Network Layer has blinded a whole generation to these situations. The relatively recent transition of wireless technologies going from a bleeding-edge curiosity to a fully deployed and ubiquitous presence is forcing many to rediscover the impact of shared media and perhaps raises questions about what might be happening on the wired network. When faced with network support or forensic situations it is impossible to have too much training about how each element within the network behaves normally, so that abnormal behavior is recognized and unexpected symptoms can be explained and compensated for.

NETWORK SUPERVISION

Fluke Networks

P.O. Box 777, Everett, WA USA 98206-0777

Fluke Networks operates in more than 50 countries worldwide. To find your local office contact details, go to www.flukenetworks.com/contact.

©2008 Fluke Corporation. All rights reserved.
Printed in U.S.A. 4/2008 3331616 Rev A