

Troubleshooting Switched Environments

Ten years ago, the network was relatively simple. There were hubs, bridges and routers. Each was a discrete box, readily identifiable from the others. Troubleshooting was also simple. If you were attached to a hub, then the rules for troubleshooting a collision domain applied. At the point where the collision domain attached to a bridge, all errors stopped. Troubleshooting using a protocol analyzer was the best available option, and it was very effective once the user knew the basics of the network and the protocols in use.

Then switches appeared on the scene.

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? The troubleshooting difficulty originates from a basic 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 half duplex 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 falling price of switching technology, most new networks have a single station per port. So, in the case of half duplex connections, 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 need a routing service in the switch typically requiring considerable CPU resources in forwarding traffic. The network situation where it is appropriate to require routing on every single request and reply is very difficult to imagine, and this configuration should be avoided. 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.

Five techniques for troubleshooting a switch

There are five fundamental approaches used to gain visibility into a switch. Each of these techniques 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 controlled primarily by the availability of resources (which tools are available and/or pre-installed), and by the potential service interruption that will be created by using 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 nearly impossible to see all of the traffic flowing through a switch. 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, 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.




Figure 1 – A very basic switch scenario.

For simplicity, the troubleshooting model will be a server attached to a switch, as shown in **Figure 1**. Some descriptions will assume the user(s) in question are attached to the same switch, other descriptions will assume the user(s) in question are accessing the server through the uplink to either another switch or to a router. The troubleshooting scenario will start with a simple report that communications with the server are “slow.” This report tells the support staff almost nothing.

Method 1: Access the switch console via TELNET or the serial port

Senior network support staff or others with the password to the switch may choose to review the configuration of the switch during the troubleshooting process. The switch configuration is available by logging in through a TELNET session, or by attaching to the serial port of the switch and logging in. (**Figure 2**)



Figure 2 – Using the RS-232 console port.

OptiView™ Network Analysis Solutions have powerful troubleshooting capabilities to put you back in control of your switched network.

Plug into the network with an OptiView Integrated (portable) or Workgroup Analyzer and you'll have instant vision into your switches. Trace SwitchRoute shows you the exact path two devices use to communicate through your switched network. No more hidden problems.

You'll also see:

- Switch port utilization, errors and configuration
- Hosts connected by port
- VLANs on the switch with port membership
- Trunk ports and trunking protocol
- Multiport stats show you traffic levels across switch ports at a glance.

OptiView Link Analyzers provide real-time analysis and line rate packet capture for full duplex, switched Gigabit Ethernet links working with in-line taps. OptiView Protocol Expert software uses extensive seven-layer decodes to make it easy to identify and solve tough problems on switched segments. OptiView Console's SwitchTap port mirroring feature allows you to easily and safely configure mirror ports. It will even guide you to the best port to use with a Fluke Networks tool attached. You can then import the mirrored data into the OptiView Console database, trend it, analyze it, map it, and report on it.

Visit www.flukenetworks.com/optiview and check out the OptiView Network Analysis Solutions.

The switch configuration is available from either of these console access options, though the configuration information will not reveal misbehavior on the part of the switch. Whether the operating system of the switch has bugs, or whether the configuration is incomplete, is not likely to be evident from the configuration listing. This configuration data will instead be useful in guiding troubleshooting efforts to see if the switch is operating as expected. To validate the configuration, it will be necessary to utilize one or more of the other switch troubleshooting methods.

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. Considerable experience and theoretical knowledge is required to obtain benefit from several of these features.

Method 2: Connect to a spare (unused) port

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




Figure 3 – Monitor from any open port.

Connecting to a spare 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.

Unfortunately, the 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's designed to prevent unnecessary traffic from reaching ports where it does not belong. The protocol analyzer has not requested any traffic, and usually has not even transmitted a single frame.




Figure 4 – Switches forward traffic between the source and destination port. Very little traffic goes to other ports. The monitoring tool will see a few frames per minute instead of the thousands per second that may be passing between the stations and the server

The traffic forwarded to the monitored port will consist almost entirely of broadcast 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, and not often from truly unknown destinations. Many unwary technicians have seen the traffic distribution (nearly 100 percent broadcasts) and not noticed 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.

Since this view of the network is all but useless, it is necessary for the monitoring tool to solicit traffic. Soliciting traffic or

interrogating the broadcast domain is useful for network discovery and for finding other classes of problem, but will not aid a slow user connection problem significantly.

A better option, available from most switches, allows traffic from a selected port or ports to be copied to the monitoring port. (See **Figure 5**) This technique is usually referred to as port aliasing, port mirroring, or port spanning.

The ability to copy or mirror traffic is provided by most switch vendors to allow a monitoring tool to be connected to a configured monitor port on the switch. Older switches had a specific port that could be configured as this special monitoring port, but most newer switches can be configured to use any port as the output port for monitoring.

The implementation of this technique varies between vendors, but there are several common mirroring options. 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 almost all errors are filtered by the switch, and do not appear on the monitor port. For troubleshooting purposes, port mirroring can sometimes be quite ineffective because a whole class of problems is concealed by the switch in this way.




Figure 5 – Logical effect of configuring a mirror port.

In addition, the actual configuration must be performed either from the console (the RS-232 port on the switch) or a Telnet session. This often involves bringing a PC or terminal along with the monitoring tool so that the switch can be reconfigured as required for troubleshooting.

The mirror port is often a listen-only port, though a number of vendors permit configuring the port to be bidirectional. 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 mirrored port could be any other port on the switch, including uplink ports. The mirror could also be several or all other ports on the switch. The more ports included in the mirror, the less likely the output will include all of the traffic, and the capacity of the output port is rapidly exceeded.

Output capacity on the monitoring port is an important problem. The output port has a TX and RX path. It was already noted that the TX path from the monitoring device back to the switch may 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.

Assume you are seeking to monitor the traffic associated with a server connected to the switch at 100 Mbps in full duplex. At full duplex, the server's port – the TX path – is able to support 100 Mbps of traffic. At

the same time, the RX path is able to support 100 Mbps of traffic for an aggregate throughput 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. Any traffic on the switch port that exceeds 50 percent of the capacity of that link (200 Mbps) will be dropped.

If multiple ports are mirrored to the monitoring port, then this problem is potentially correspondingly worse. Since most switches operate at far below their capacity, the problem may not be noticed right away. Most user connections experience low single-digit utilization averages. Infrequently, there is a short but large burst of traffic.




Figure 6 – Output capacity is limited on a mirror port.

The situation may be mitigated by connecting the monitoring device to a higher speed port, one which has the native capacity to accept all of the mirrored output. If the mirror output port in **Figure 6** were a Gigabit port instead of a 100 Mbps port, then the aggregate 200 Mbps traffic load would be easily accommodated.

Method 3: Insert a hub into the link

Using a shared media hub involves a strategic placement of the monitoring tool. In many networks, most traffic will be received or transmitted by a shared resource such as a file server. Adding a shared-media hub

between the switch port and the file server allows an analyzer to be connected to the same collision domain as the file server, as shown in **Figure 7**. This technique enables the analyzer to see all the traffic to and from the file server, which assists the network support staff in diagnosing a wide range of problems, including user login failures, poor performance, and dropped connections.



Figure 7 – Using a hub to monitor a switched port.

This approach is impractical in most situations, particularly where there are multiple servers to be monitored. Where do you locate the hub? On all shared resources? 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, shared resources may be connected via a technology or connection speed that your monitoring tool may not support.

Using a shared media hub is still an effective way to monitor all traffic and errors present on a link. This is nearly the only way to actually see and analyze MAC Layer errors in a switched environment. Using SNMP to learn about them is useful. But, for good error analysis, there is nothing like seeing them with the diagnostic tool directly.

There are two major drawbacks to this method. The server link cannot be a full duplex connection or the resulting duplex mismatch will introduce more errors than

you are likely to reveal. Also, to diagnose problems with this method, a shared media hub is necessary. Many newer hubs are actually bridging devices masquerading as hubs. If you insert a hub that is not a shared media device, then you will not see the traffic you are looking for. It would be the equivalent of inserting another switch in the link, and would not benefit you at all. Dual speed hubs – such as 10/100 – may offer a collision domain at each speed, with a bridge between speeds. In this situation, you can use the dual speed hub as long as you are certain the monitored link is operating at the same speed as the diagnostic tool. Other hubs offer full bridging between all ports, and are therefore actually misleadingly labeled low-cost switches. They will not work for this method.

Method 4: Using a tap or splitter

This is somewhat similar to adding a shared media hub, except the tapped link may only be used for receiving and does not allow the monitoring tool to transmit.

The terms *tap* and *splitter* are potentially interchangeable, 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 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 monitoring connection. This loss of power clearly implies that if a link is already suffering from cable faults or excessive distance, the splitter could easily cause the link to fail by taking too much light from the primary transmit path. A splitter can easily cause a 3 dB loss in power over a




Figure 8 – Use of a tap or splitter

fiber link. Some transmitters are more robust than others, 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. Fiber optic splitters do not require power, as the split is accomplished by splicing fiber optic strands together very precisely. Since the splice is created in alignment with the inbound signal, it is vital to ensure the splitter input and output cables are connected correctly.

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 cables, 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 facing cable problems. Copper taps require power, as the signal is recovered and retransmitted to the monitor port. If properly designed, a copper tap will not drop or disrupt the tapped link if power is lost to the tap.

Tapping the line is an excellent way to see what is passing through a link. Once installed, the tap is invisible to the attached devices and may be utilized at any time without further disruption. Unfortunately, the link must be broken to insert the tap. Furthermore, the nature of a tap or splitter is that it will offer the data passing in each direction separately. That is, the transmit

path will be offered on one connection and the receive path on another.

To simultaneously monitor a request and response passing through the tapped link, it is necessary to have a monitoring tool with two input ports. Tools with dual inputs typically have the ability to separate each direction, or to integrate both data streams for analysis. The alternative is to examine traffic in one direction at a time, which is more difficult to analyze. There is no operational difference when monitoring half duplex or full duplex, and a tap is equally effective with either. You have the choice of monitoring a single direction separately with a typical single-input monitoring tool, or monitoring both directions simultaneously with a dual-input monitoring tool.




Figure 9 – Functional diagram of tap operation

Method 5: Query the switch using SNMP

Almost the only effective method of troubleshooting a switched network is to ask the switch itself how the network is behaving. This is 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.

Virtually all but the least expensive switches ship with SNMP management capability. The primary differences are in how granular the answers are. Some less expensive switches have SNMP agents that only offer information about the switch as a whole; other, more expensive switches offer very detailed information about each individual port.

SNMP is probably the most common and least intrusive method of monitoring a switched network. 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.




Figure 10 – Using SNMP to monitor a switch.

Because switches do not routinely forward errors, using SNMP is perhaps the best method of locating ports experiencing them. The switch may not forward the error, but it is certainly aware of the presence of errors.

There are a variety of MIBs available from most switches that support SNMP. Each MIB supported brings the console a slightly different or more detailed view of network conditions around the switch. In addition to private MIBs, which typically have customized support for each switch and level of operating code, the standard 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 good too.

RFC 1213 – MIB II

RFC 1643 – Ethernet-Like Interface MIB

RFC 2819 – RMON Ethernet

RFC 2021 – RMON 2

RFC 2613 – SMON

Many RFCs are 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) 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 usually enabled with a very common password when the switch is sold. SNMP passwords are called community strings, and are both case and punctuation sensitive. Community strings are transmitted in clear text, too, which in itself creates a security risk. While not widely deployed at this time, SNMP V3 offers encrypted communication to address this exposure. The most common default community string is *public*. It is shocking how many SNMP agents are accessible from the Internet with the *public* community string – even today.

At a minimum, the default community string should always be changed immediately. 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 providing a path to those SNMP agents may impose a variety of limits 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. In some cases it is necessary to upgrade the operating system on the switch before it is capable of supporting a desired MIB. Problems with this method include when an SNMP agent's implementation of a specific MIB is not accurate, and responses to queries are simply wrong. It doesn't happen that often, but programming errors occasionally result in inaccurate responses.

There are many reasons why your switch does not respond to a specific SNMP query. Once any access problems are resolved, SNMP offers a very useful tool for monitoring and trending.

Conclusion

A commonly used troubleshooting method is to wait for user complaints. 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. The problem with this method is that it is entirely reactive instead of proactive.

Ideally, the approach should be proactive. 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. Implementing tactics such as monitoring and trending switch port statistics and using tools that allow you see inside switches will take you from a troubleshooting mode to a trouble prevention mode.

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.

©2004 Fluke Corporation. All rights reserved.
Printed in U.S.A. 4/2004 2142429 A-ENG-N Rev A