

Applying EtherPeek to Switched and Gigabit Ethernet Network Management

Executive Summary

This white paper characterizes the switched infrastructure in the context of network protocol analysis and shows how WildPackets' EtherPeek can be used to effectively analyze and troubleshoot problems. These issues are also significant with Gigabit analysis. This paper further describes how EtherPeek can provide proactive network management in switched topologies and how effective analysis can be performed when network traffic is segregated and isolated.

March 2003



Contents

Shared versus Switched Networks	1
EtherPeek in the Switched Network Environment.....	2
Developing a Methodology for Switched Network Analysis	6
Gigabit Ethernet Analysis	7
Aggregate Bandwidth Limits for Gig Mirroring	7
Capture Rate is Dependent on Packet Size	7
Capture Limitations are not a Significant Impediment	8
Span multiple 100 Mbit/sec nodes	8
Span a single Gig node	8
Span multiple Gig nodes	9
Span a combination of Gig and 100 Mbit/sec nodes	9
Conclusions and Summary	9

Applying EtherPeek to switched and Gigabit Ethernet network management

This white paper illustrates the switched infrastructure in the context of network protocol analysis and explains how WildPackets' EtherPeek™ can be used to analyze and troubleshoot problems. This paper will further describe how EtherPeek can be used as a key tool in providing proactive network management in switched topologies (in both classic 10/100 Mbit/sec networks as well as in contemporary Gigabit Ethernet networks) and how effective analysis can be performed when network traffic is segregated and isolated. (Note: "EtherPeek" in this paper will refer to both EtherPeek standard and EtherPeek NX™, WildPackets' expert analysis tool.)

WildPackets, Inc. developed the EtherPeek Ethernet LAN protocol analyzer in 1990 as a tool for capturing, decoding, and analyzing traffic in real-time for all devices in a LAN. The structure of the LAN itself, however, has changed dramatically in the ensuing years. Armed with information about how to apply EtherPeek to this new structure, the network engineer can integrate the protocol analysis process into the overall spectrum of tools and utilities available for network troubleshooting and analysis. This array includes switch statistics, available through vendor-specific switch management software, and the industry-standard SNMP/RMON statistical and packet gathering tools available today.

Shared versus Switched Networks

In the 1980's, the networking marketplace saw the beginning of the migration away from shared, coaxial Ethernet to the realm of twisted pair, hub-based Ethernet networks. Both coaxial connectivity and early twisted-pair hubs created an infrastructure in which every station's transmissions were visible to every other station within what is referred to as a "collision domain." Bridges could connect a series of collision domains into a larger infrastructure called a "broadcast domain." A broadcast domain consists of all the interconnected stations that receive each other's broadcast and multicast packets and is an area bounded by routers. A router, therefore, forms the end of a broadcast domain and serves as the gateway into a different broadcast domain.

In a network where all stations in a collision domain "see" all traffic from all the other stations in the collision domain, we say that the network medium (the cable system) is "shared." That is, everyone must compete with everyone else to take turns accessing the medium for transmitting packets.

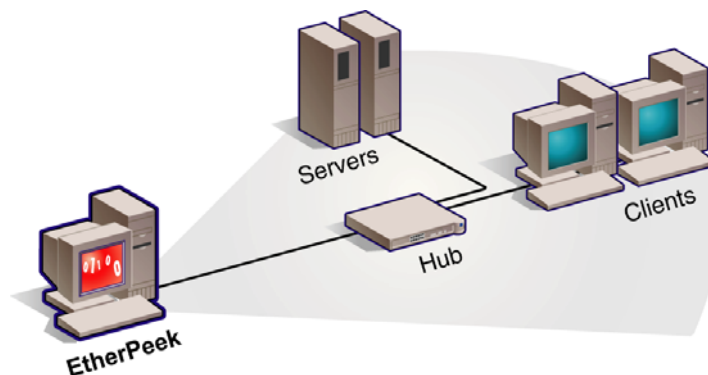


Figure 1 A shared-media network with EtherPeek.

Protocol analysis in a shared-media environment is straightforward, as depicted in Figure 1, above. By simply attaching an analyzer at any point in a collision domain, one can acquire 100% of the transmitted packets from all stations within that domain.

Over the decade of the 1990's, the networking marketplace saw dramatic increases in desktop computing power. As application programs grew in complexity and sophistication, the need to send large quantities of data as quickly as possible grew proportionally. The shared-media environment forced all of these communicators to compete with each other for the use of the media. This proved to be an inadequate solution. To facilitate the demands of these increasingly complex networks, the industry experienced an evolution from shared media to switched network infrastructures. Today star-wired LANs using switches as the central connecting points are pervasive, creating large meshed network topologies.

While switched networks provide part of the solution for efficient use of the network media and infrastructure, they bring with them some inherent restrictions and limitations to the protocol analysis engineer.

By their nature, switches do not forward all packets to all stations. Of course, broadcast and multicast packets continue to be forwarded out to all ports of a switch and, therefore, reach all the stations in the broadcast domain. This is identical to the shared-media model. Directed frames, however, are forwarded in a much more intelligent manner. A "directed frame" is one with a specific Ethernet address as the destination target address. It is intended for only one recipient. The switch evaluates the Ethernet destination address on all incoming packets and forwards them only through the single port to which the intended target machine is attached.

As a result of this behavior, the network benefits from a reduction in contention for network bandwidth and a corresponding reduction in Ethernet collisions and the resulting retransmissions. This can easily be seen if one considers a simple topology in which a single switch has two file servers and sixty workstations attached to it. At the same time that Workstation #1 is sending a packet to File Server #1, it is possible for Workstation #2 to send a packet to File Server #2. Neither workstation is required to wait for the other, as would have been the case in the older shared-media networking model.

EtherPeek in the Switched Network Environment

Now consider what happens when a network engineer attaches EtherPeek to some other port on the switch with the sixty workstations and two file servers just described, as shown in Figure 2, below. Since the packet from Workstation #1 is addressed to File Server #1, the switch only forwards that packet to the port to which File Server #1 is attached. The packet is not forwarded to the port to which EtherPeek is attached. EtherPeek does not see this packet, or any other packets that are directed to specific Ethernet destinations. This is the inherent nature of a switch and is normal, correct, and performance-enhancing behavior. If a workstation were to transmit a broadcast or multicast packet, then the switch would forward it out to all of its ports. EtherPeek would be able to capture broadcast and multicast packets since all of these packets would be sent by the switch to the EtherPeek port and to all other ports.

To overcome this inherent behavior, and to allow protocol analyzers to be attached effectively to switches, the switch manufacturers have implemented a mechanism by which the switch administrator can select a port to be dedicated to the analysis process. This is done through the switch management software and implies that the particular model of switch supports this type of configuration. In addition, even if a switch does not support the selection of a particular port for use with an analyzer, there are ways to work around the restriction and still effectively connect EtherPeek for protocol analysis.




Figure 2 EtherPeek attached to a switched network.

Configuring a Protocol Analysis Port

The most common term used to refer to the special port configured for use by a protocol analyzer is a “Mirror Port.” In the Cisco environment, it is called a “Span Port.” The table, below, shows a list of the terminology used by various vendors to refer to the configuration of a special protocol analyzer port in their switch equipment.

Imagine that a File Server is attached to Port #12 on a particular switch. For the sake of discussion, assume that EtherPeek is attached to Port #7 on the same switch. Through the switch management software a command is issued that “mirrors” Port #12 to Port #7. All traffic going in or out of Port #12 (to/from the File Server) is copied and sent out Port #7 (to EtherPeek). EtherPeek is now able to capture all traffic to and from the File Server. We say, “the File Server port is being mirrored.”

There are a variety of port mirroring options available from various vendors and, although the essence of the mechanism is the same (mirrored traffic is sent to the analyzer port), the options and configuration parameters vary from vendor to vendor. It is important to consult the switch documentation to understand the capabilities and configurations required to activate port mirroring.

Vendor	Terminology
Xtreme Switches	Port Mirroring
CISCO	Port Spanning
3 COM	Roaming Analysis Port
XYLAN	Mobile Port
Nortel Networks	Copy Streaming
NetScout	Roving Port
Foundry Works	Port Mirroring

Mirroring Multiple Ports Simultaneously

It is possible, with some vendors’ equipment, to mirror more than one port at a time. For example, several users may be complaining that they are having problems printing. EtherPeek

could capture the traffic to and from all of the users if the switch mirrored each of the users' ports back to the EtherPeek port.

A Virtual LAN (VLAN) environment is another case in which specialized port mirroring may come into play. Imagine the switched network that was previously described. This network has two file servers and sixty workstations attached to a single switch. When implementing a VLAN, one may decide that 30 of the workstations and one of the file servers will be Network #1 and the other 30 workstations and the other file server will be Network #2. Hence, the administrator of the switch has created two separate broadcast domains and, therefore, two separate "virtual" networks from the devices attached to a single switch.

It is also possible, with some vendors' equipment, to mirror all traffic within a VLAN and have a copy of each packet transmitted in the VLAN sent to the analyzer port.

The Problem with Aggregate Bandwidth

In any situation where an analyzer port is going to receive mirrored traffic from more than one mirror port, there is a possibility that the overall aggregate of packets from all of the mirrored ports may exceed the bandwidth capacity of the analyzer's mirror port itself. The other possible problem is that the analyzer might become overloaded by Gigabit line speed.

EtherPeek is attached using a 100 Mbps Ethernet connection to Port #7 on a switch. The switch is configured so that Port #12, #13, and #14 are mirrored onto Port #7. If the devices attached to the three-mirrored ports are each transmitting 45 Mbps (45% utilization on their individual 100 Mbps connections to the switch), then there will be 135 Mbps of aggregate traffic that the switch will try to send out Port #7 to EtherPeek. This will result in packets being dropped by the switch. The switch has no way to send 135 Mbps over a 100 Mbps Ethernet connection to EtherPeek.

If EtherPeek is going to be used with multiple port mirroring, then there must be consideration, and awareness, of the fact that packets exceeding the bandwidth capacity of the mirror port to which EtherPeek is attached will be dropped.

Analysis of Non-Managed Switches

Some switches do not have switch management software running in them. Therefore, there is no way to mirror ports on these non-managed devices. In this case, it is necessary to insert a simple, repeating hub between the switch and the device being analyzed and then attach EtherPeek to the hub.

In the picture below, you see the way that EtherPeek would be attached to a non-managed switch if the clients were complaining that they had problems accessing the Server. The cable from a client to the Switch is unplugged and a simple repeating hub (a \$40.00 4-port hub, for example) is inserted between the client and the Switch. EtherPeek is now attached to the Hub. Essentially a "Y"-cable has been created. EtherPeek can now capture all traffic between the client and the Switch.

This same logic could be applied if EtherPeek were being used to capture all traffic to and from the Server. The hub would be inserted between the Switch and the Server.




Figure 3 EtherPeek with a non-managed switch.

Remote Capture Tools

When it is necessary to analyze switched traffic on remote segments of a distributed network, remote capture tools are available which can acquire remotely captured trace files and send them back to a full-featured protocol analyzer like EtherPeek. WildPackets' PacketGrabber™ and RMONGrabber™ are remote capture tools that can be connected beyond a switch on a distributed network and send remotely captured packets back to be analyzed by EtherPeek.

WildPackets' PacketGrabber is the simplest solution. PacketGrabber can be used to troubleshoot enterprise networks by gathering trace files from workstations on remote segments and automatically send them by email or ftp for analysis in EtherPeek, as shown in Figure 5, below. WildPackets' PacketGrabber can be run from any machine in the network to capture a trace file, which can then be analyzed by EtherPeek.




Figure 4 PacketGrabber can automatically forward captured packets to EtherPeek.

WildPackets' RMONGrabber collects packets from an RMON probe directly into an EtherPeek capture window. You can connect to an RMON probe in a remote location and instantly view the packets in a local EtherPeek Capture window, identifying problem situations without requiring expensive and inconvenient off-site visits. When installed with EtherPeek NX, RMONGrabber can build Peer Maps and perform Expert Analysis on remote networks. With RMONGrabber, you can also capture from multiple probes at the same time

using multiple Capture windows. RMONGrabber allows filters and triggers to be set in order to restrict the traffic that is included in the captured file.




Figure 5 RMONGrabber collects network data from an RMON probe.

Attaching a remote data capture utility to a switched network is exactly like attaching EtherPeek. The requirements and constraints are identical. A switch port must be mirrored or a hub must be inserted to allow the remote capture machine to acquire directed frames.

WildPackets also has developed a trace file conversion utility called ProConvert™ that will convert trace files from most other vendors' protocol analyzers into EtherPeek-readable format.

Developing a Methodology for Switched Network Analysis

The guiding principle for analyzing a switched network is that “you cannot see all of the traffic at the same time.” This is a reflection of the fundamental distinction between the shared-media networks of the 1980's and today's switched infrastructures.

The methodology for analysis in a switched network environment focuses on two fundamental types of problems:

- You have a suspect: A particular user, server, router, or other device is suspected of having or causing a problem.
- You do not have a suspect: There is a need to assess the overall characteristics, performance, and statistical measurements related to the network as a whole.

When a suspect station is being analyzed, the methodology is relatively simple. It is only necessary to mirror (or insert a hub) to capture all of the traffic to and from the suspect station. Analysis then proceeds in a normal manner.

When no suspect is identified, it becomes critical that the statistics and other management information provided directly by the switch be examined. Managed switches will report a broad spectrum of statistical information concerning network performance and protocol behavior. Since the switch itself can “see” 100% of the traffic passing through it, the statistics and other information provided by the switch become the key piece of an overall network assessment.

Of course, when measuring the baseline performance in a network, it makes sense to expand on the basic switch statistics by capturing packets (by mirroring or with a hub) from each of the file servers and routers on the network. In general, all users will be communicating with the servers or through the routers. Hence, while there may be several hundred users, there will

probably be a much smaller number of servers or routers. The job of individually measuring traffic characteristics and protocol behavior from each server and router may seem onerous, but that is the consequence of accepting the benefits of switched network engineering as opposed to coaxial Ethernet or simple hubs.

Gigabit Ethernet Analysis

In the 1990's, the early adopters of Gigabit Ethernet were faced with analysis challenges that have become increasingly less significant as widespread deployment of Gigabit switches occurred. Early adopters were required to use expensive, inflexible hardware “pods” with limited feature sets to insert into a Gig link—this was due to the fact that Gigabit switching was in its early stages of evolution. Today all high-end Gig switches provide port mirroring and spanning capabilities that work exactly like the 10/100 Mbit/sec functions that preceded them. A network manager is no longer required to purchase specialized external hardware devices to make a connection to a Gigabit Ethernet link. The switch takes care of that for you!

EtherPeek NX supports Gigabit Ethernet NICs in exactly the same way it supports 10/100 Mbit/sec NICs. Of course, if you're going to use a Gig NIC in with a notebook computer, you'll need a PCI bus extender (since nobody manufactures a Gig PCMCIA card at the present time). In any case, the Gig card shows up in the Select Adapter dialog box and you simply capture from it. Simply attach the Gig card (using either a fiber or copper connection) to your switch and configure a Gig port as the mirror port.

One thing to keep in mind (and this is also true for a 100 Mbit/sec mirror port as well) is that the mirror port is uni-directional. That is, although the Ethernet is operating in full-duplex mode, there is nothing to “full-duplex” from the analyzer to the mirror port. Traffic only flows in one direction from the switch to the analyzer. Consequently, when one talks about a Gigabit Ethernet having a 2 Gig bandwidth capacity (1 Gig in each direction of the full-duplex link) this is not the case for a mirror port. You get a 1 Gig “pipe” from the switch to the analyzer.

Aggregate Bandwidth Limits for Gig Mirroring

There are four situations to consider when mirroring or spanning using a Gig port:

- Span multiple 100 Mbit/sec nodes
- Span a single Gig node
- Span multiple Gig nodes
- Span a combination of Gig and 100 Mbit/sec nodes

In each case, there are two limiting factors to consider. First, the Gigabit Ethernet NIC, PCI bus architecture, and analyzer-host-computer-CPU speed will combine to allow something less than full-rate (1 Gigabit/sec, continuous packet streaming). The second limiting factor is, of course, the fact that the switch itself can push no more than 1 Gigabit/sec of packet traffic out onto the mirror port's transmitter. So, you can't be sent more than 1 Gigabit/sec from the switch, and you can't actually move a full 1 Gigabit/sec of continuous data into the analyzer's buffer. However, these limitations are not normally a significant impediment to the analysis process.

Capture Rate is Dependent on Packet Size

A Gigabit Ethernet NIC, in conjunction with the associated device driver and operating system interface, form a limitation as to the rate at which an analyzer can acquire packets. In general, the number of packets, rather than their size, is the most significant factor. That is, as packet size gets smaller, the overall capture rate (in bytes per second) gets smaller. With maximum-size Ethernet packets (1518-bytes) the capture rate approaches more closely to 100% (1 Gbit/second). With minimum-size packets (64-bytes) the capture rate could fall

below 5% (that is, below 50 Mbits/second). If the maximum capture rate is exceeded, the analyzer's Gigabit NIC will simply drop the excess packets.

As luck would have it, applications that tend to produce smaller packets (terminal access applications such as Telnet) also tend to produce fewer packets (a user typing on a terminal screen is slow compared to the Gigabit data transfer rate). Applications that tend to produce larger packets (saving files, database, etc.) are the ones that often introduce the greatest load on the network. The consequence is that a typical network usually won't produce an aggregate of traffic that cannot be successfully captured with an analyzer.

Capture Limitations are not a Significant Impediment

Let's consider the four types of analysis situations that may be encountered and see what a real-world troubleshooting scenario would be. We'll assume that the packet sizes in use allow for a 60% (600 Mbits/second) capture rate. This is a scenario that is consistent with real-world consulting engagements that have been encountered at WildPackets.

Span multiple 100 Mbit/sec nodes

This would be the situation where several "power-users" were complaining about performance or intermittent network problems. You would span a group of ports and capture from all of the users. On the server side, you may have a database that resides on multiple servers. Troubleshooting SQL problems might require that all of the servers are spanned together. Perhaps you want to assess an entire VLAN in which case all of the nodes in the VLAN would be captured.

Assuming a 600 Mbit/second capture rate, this would reasonably allow full-rate capture from four file servers. Since a file server is a good candidate to take advantage of the 200 Mbit/sec bi-directional full-duplex 100 Mbit/second link, and assuming a normal, reasonable file I/O behavior, it's probably that four heavily used servers would not overload a single Gig mirror port. In more common practice one might not even encounter packet loss with as many as 7 or 8 servers (assuming that they were doing typical file-save/read operations).

A user's machine seldom takes advantage of the capabilities of a full-duplex link. A client machine must wait for a reply to each request that it makes, and, as a result, this leaves the client less able to send and receive at the same time (in full-duplex). Consequently, one could assume that 10 client machines, running 100 Mbit/second Ethernet, could be captured through a single Gigabit link.

In practice, it's possible to capture VLAN traffic from over 100 machines without packet loss assuming, of course, that the users are "typical" (from the standpoint that they perform much of their work locally with occasional email checking, web browsing, and file reading/writing).

In conclusion, you can surely capture 7 or 8 servers, 10 client machines, and possibly an entire VLAN using a Gigabit mirror port.

Span a single Gig node

In theory, a single Gig node (a large server, for example) could experience up to 200 Gbits/sec aggregate throughput (sending and receiving on a full-duplex link). Of course, the mirror port limitation of 1 Gbit/sec, as well as the analyzer limitations, form the upper boundary for capture capability.

A Gigabit-attached client (as with its 100 Mbit/second counterpart) normally doesn't take advantage of the capabilities of the full-duplex Gig link. In fact, the data rate for a client is often below the 600 Mbits/second capture limit suggested for the Gig card in the analyzer.

One might assume that 3 or 4 Gig-attached clients could be successfully captured through a single port span. Likewise, 1 or 2 Gig-attached servers could typically be captured using a single Gig span port.

The WildPackets consulting team has successfully used Gig span ports to do VLAN spanning for multiple Gig-attached clients and servers without noticeable packet loss.

Span multiple Gig nodes

In at least one case, WildPackets' on-site consulting engineers established a VLAN span for over 100 Gig-attached client machines. They were able to capture 512-Mbyte buffers full of data suitable for analysis without packet loss, thus causing confusion or inability to analyze the trace.

Span a combination of Gig and 100 Mbit/sec nodes

Bear in mind that the mirror/span port on the switch (to which the analyzer is Gig-attached) will impose a 1 Gbit/second maximum limit. From this the Gig NIC in the analyzer will be able to capture at a rate determined by the intermixture of small and large packets. Remember that the 100 Mbit/second packets will be arriving at the analyzer's port at the Gig speed (since the mirror port is a Gig port).

In practice, you may expect somewhere around a 600 Mbit/second capture rate, but this will possibly translate to many more 100 Mbit/second stations than Gig stations.

Conclusions and Summary

Through the use of port mirroring in both Gigabit and 10/100 Mbit/sec Ethernet networks (or the introduction of a mini-hub in a 10/100 Mbit/sec network), a network engineer can effectively use EtherPeek to analyze traffic in a switched Ethernet environment. EtherPeek captures the traffic to and from a suspect station through the correct application of port mirroring or placement of a mini-hub. Switch statistics provide a basis for the assessment of a switched network as a whole.

This document has focused on the switched network aspects of applying EtherPeek to the protocol analysis task in both classic 10/100 Mbit/sec networks as well as in contemporary Gigabit Ethernet networks. WildPackets' EtherPeek NX extends the decoding and reporting capabilities of EtherPeek with built-in expert analysis—by using advanced algorithms to locate problems in trace files that may be buried or may be outside the experience of the analyst.

It is no longer necessary to purchase expensive, inflexible hardware “pods” with limited feature sets to analyze a Gigabit Ethernet implementation. WildPackets provides comprehensive solutions to analyze and troubleshoot today's complex, sophisticated switched network infrastructures. When the features, capabilities, and price point of WildPackets' solutions are explored, it is reasonable to have every networking professional armed with the entire suite of EtherPeek analysis tools.

WildPackets Professional Services

WildPackets offers a full spectrum of unique professional support services, available on-site, online or through remote dial-in service.

WildPackets Academy

WildPackets Academy provides the most effective and comprehensive network and protocol analysis training available, meeting the professional development and training requirements of corporate, educational, government, and private network managers. Our instructional methodology and course design centers around practical applications of protocol analysis techniques for Ethernet and 802.11 wireless LANs.

In addition to classroom-taught Network Analysis Courses, WildPackets Academy also offers:

- Web-Delivered Training
- On-site and Custom Courseware Delivery
- The (T.E.N.) Technology, Engineering, and Networking Video Workshop Series
- On-site and Remote Consulting Services
- Instruction and testing for the Network Analysis Expert (NAX™) Certification

For more information about consulting and educational services, including complete course catalog, pricing and scheduling, please visit www.wildpackets.com/services. NAX examination and certification details are available at www.nax2000.com.

Live Online Quick Start Program

WildPackets now offers one-hour online Quick Start Programs on using EtherPeek NX/ EtherPeek and AiroPeek NX/AiroPeek, led by a WildPackets Academy Instructor. Please visit www.wildpackets.com for complete details and scheduling information.

About WildPackets, Inc.

WildPackets, a privately-held corporation, was founded in 1990 with a mission to create software-based tools to simplify the complex tasks associated with maintaining, troubleshooting, and optimizing evolving computer networks. WildPackets' patented, core "Peek" technology is the development base for EtherPeek™, TokenPeek™, AiroPeek™, and the NX™ family of expert packet analyzers. All are recognized as the analysis tools of choice for small, medium, and large enterprise customers, allowing IT Professionals to easily maximize network productivity. Information on WildPackets, WildPackets Academy, Professional Services, products, and partners is available at www.wildpackets.com.

WildPackets, Inc.
1340 Treat Blvd., Suite 500
Walnut Creek, CA 94597
925-937-3200
www.wildpackets.com

