

IP & Network Forensics

The Technology Firm
Tony Fortunato



What to do?

- Some one calls and reports something is slow.
- You are moving a network from one location to another.
- Migrating users to a newer network.
- You have been asked to perform an application profile.
 - ✓ Proactive
 - ✓ Reactive
- If I capture a bunch of packets, what do I look for?
- How do I know if something is good or not if do not have a baseline?



Tony's Top Ten Customer Issues. (sort of)



1. Client Misconfiguration – Protocol Bindings
2. Client Misconfiguration – Ethernet Auto Negotiation
3. Client Misconfiguration – Cabling
4. Client Misconfiguration – Unnecessary Services

5. Server Misconfiguration – Protocol Bindings
6. Server Misconfiguration – Ethernet Auto Negotiation
7. Server Misconfiguration – Cabling

8. Network Device Misconfiguration – Full Duplex/Half Duplex
9. Network Device Misconfiguration – IP/Broadcast Issues
10. Network Device Misconfiguration – Spanning Tree

11. Application Issues
12. Latency and Throughput Issues
- 13. No Documentation of any kind**

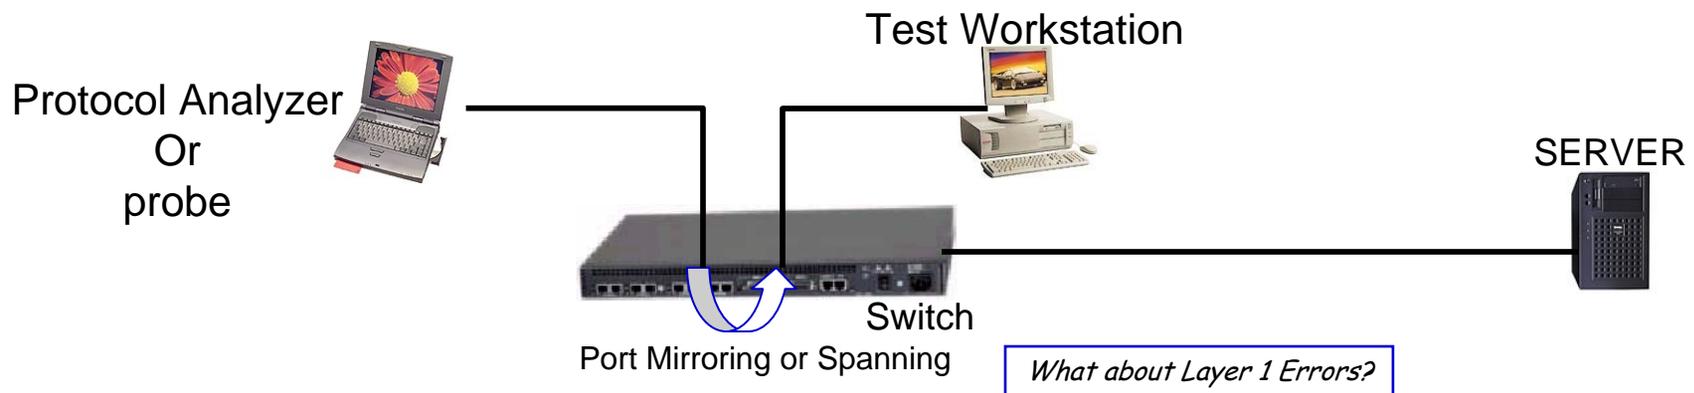
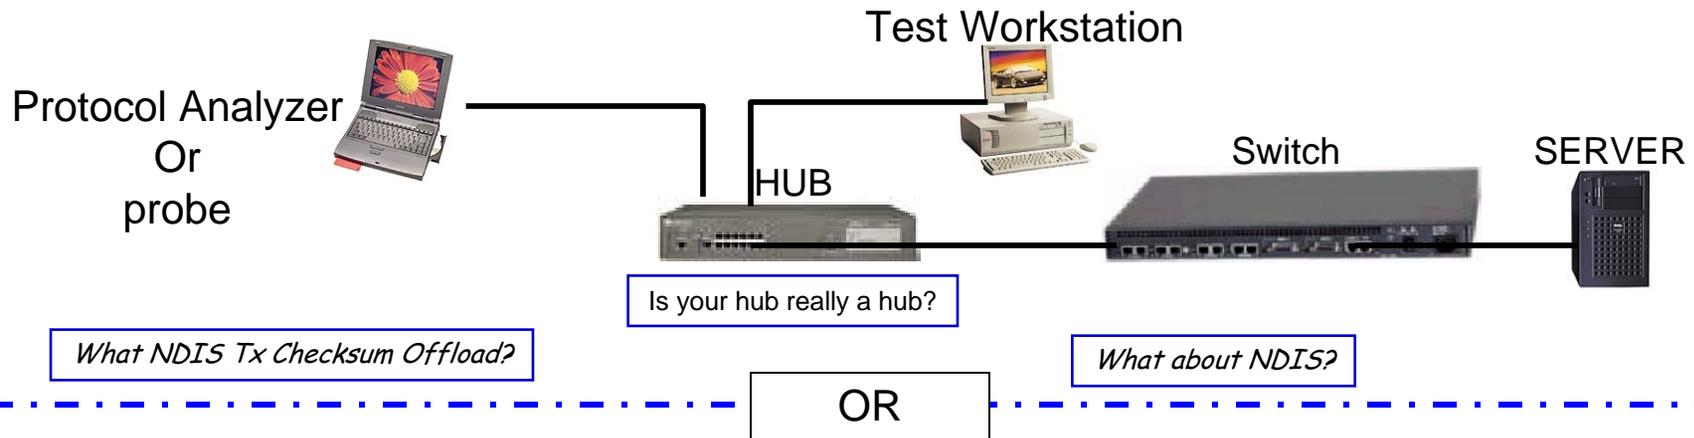
Login/Boot Up Process

- If your clients are connected to a hub, connect to any port at the same speed.
- If your clients are connected to switches, SPAN or MIRROR the client port to your analyzer.
- Use an SNMP poller or protocol analyzer to review the client PC while booting up.
- Ensure that the client isn't typing anything or logging in after the power up.
- Watch either the packet counters or hard drive access to document the end of the boot up process.

Capturing The Login/Boot Up Traffic



- Set up your analyzer to capture all traffic to and from the test workstation.
- The hub configuration is only applicable for Half-Duplex configurations.
- Full Duplex configurations require a full-duplex tap or mirror port.



Login/Boot Up Report - Template



Date		Location		Analyst	
PC OS		PC Vendor		Net Client	
Total Time		Slot/Port		Speed Duplex	
Total Bytes		Total Frames		Protocol Required	
DHCP Server		DNS Server		Protocol Found	
NDS Agent		DHCP Relay Agent Address		WINS Server	
List Of Devices That Responded		Known Bootup Processes			
Notes					
Suggested Changes					

Other Login/Boot Up Issues

- Keep login scripts under control.
 - ✓ Who is responsible for the login script?
 - ✓ How are changes to the login script managed or tested?
- Is the majority of your traffic local?
- Look for references to install servers or install technician PC's.
- Ensure that *Automatic Updates* aren't installing patches from the Internet.
- Use local time servers.
- Use local antivirus servers for signature updates.
- Note if more than one DHCP server respond to client requests.
- If asset inventory software exist, determine how much traffic it generates.
- If you are using roaming profiles, determine which servers are involved in this process and how much traffic it *costs* for your clients to have roaming profiles.

Create a Script



Use worksheet that outlines a script of tasks to be performed. Provide columns for starting and ending frame numbers.

<i>Description</i>	<i>Start Frame</i>	<i>End Frame</i>
Launch Application via Icon [app.exe].		
Enter Login Name [Joe Smith]		
Enter Login Password [don't document]		
Select Account Query from Main Screen		
Query Account 1234		

Record Frame Numbers



- Before each task is performed, note the frame number on the analyzer.
- After the task is complete note the frame number.
- Repeat this process for all items on your script.
- Make sure the person running the application waits for your signal before they move on to the next task.
- After the testing is complete, go through the capture file with the worksheet.
- If the application you are monitoring is generating a constant stream of data, set up an icon on your desktop to ping your default router. Then simply click the ping icon when you complete a task as a bookmark.

<i>Description</i>	<i>Start Frame</i>	<i>End Frame</i>
Launch Application via Icon [app.exe].	1	111
Enter Login Name [Joe Smith]	112	121
Enter Login Password [don't document]	122	144
Select Account Query from Main Screen	145	222
Query Account 1234	223	332

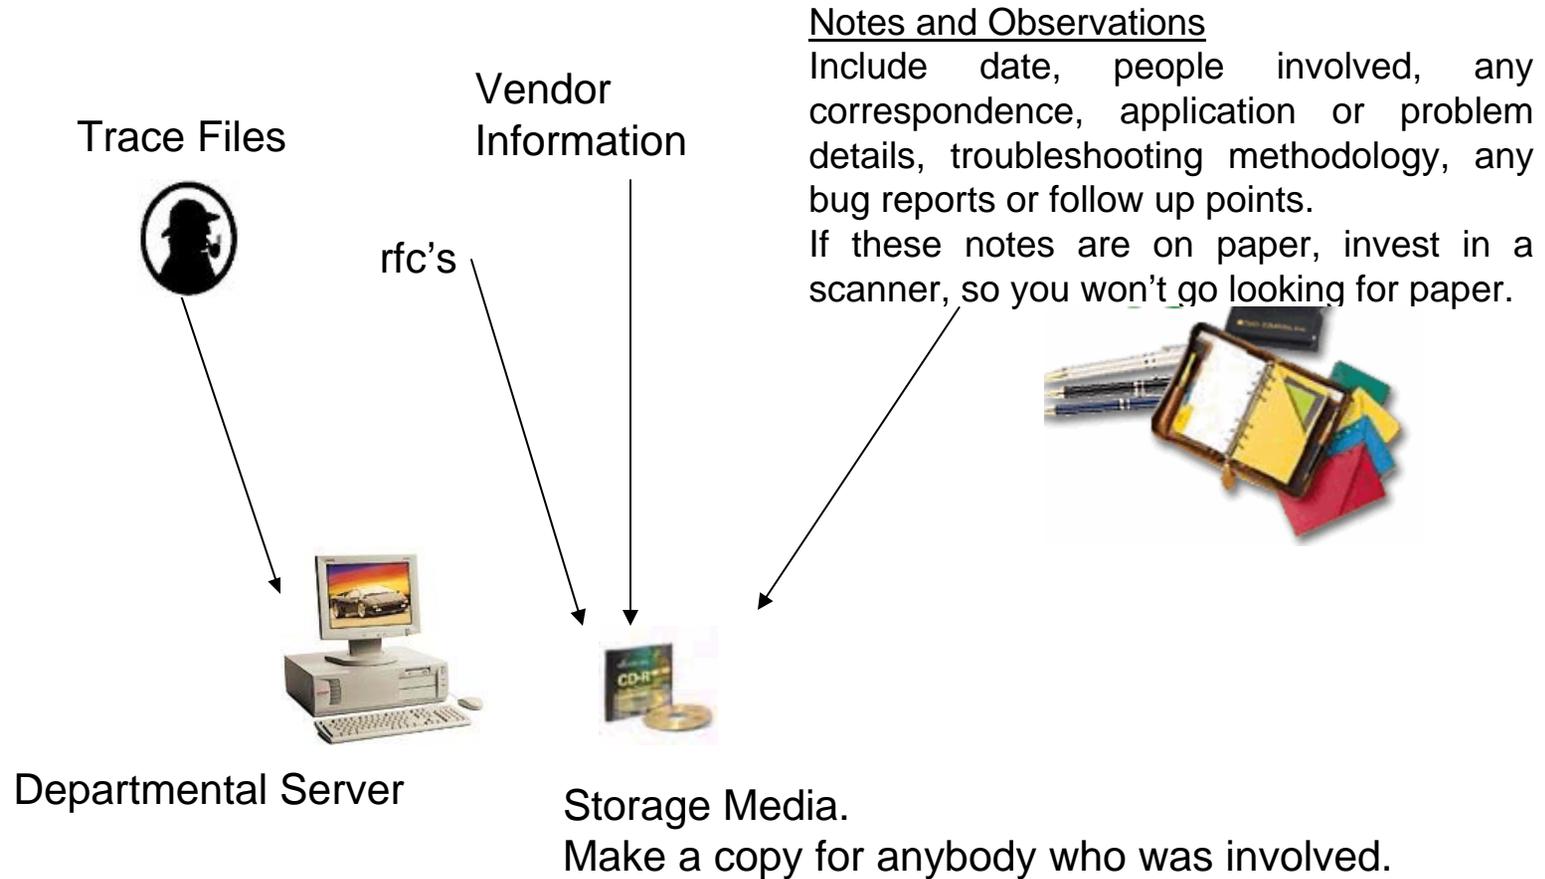
Reporting Your Results



- Now that you have gathered and calculated the traffic per task, prepare a chart with your findings.
- The only other variable to include is how many clients will be utilizing this application.
- More importantly, you need to predict how many estimated simultaneous clients will be accessing this application.

<i>Description</i>	<i>Start Frame</i>	<i>End Frame</i>	<i>Bytes</i>
Launch Application via Icon [app.exe].	1	111	120,000
Enter Login Name [Joe Smith]	112	121	10,000
Enter Login Password [don't document]	122	144	10,120
Select Account Query from Main Screen	145	222	60,000
Query Account 1234	223	332	132,022

Network Documentation Methods



What Client Misconfigurations?



The screenshot shows two overlapping dialog boxes. The top one is the 'Network' dialog, 'Bindings' tab, showing a tree view of network services. The bottom one is the 'Network Settings' dialog for a 'Xircom Card', showing various configuration options like Memory Address, I/O Port, Interrupt, and checkboxes for 'Early Transmit', 'Link Integrity', etc.

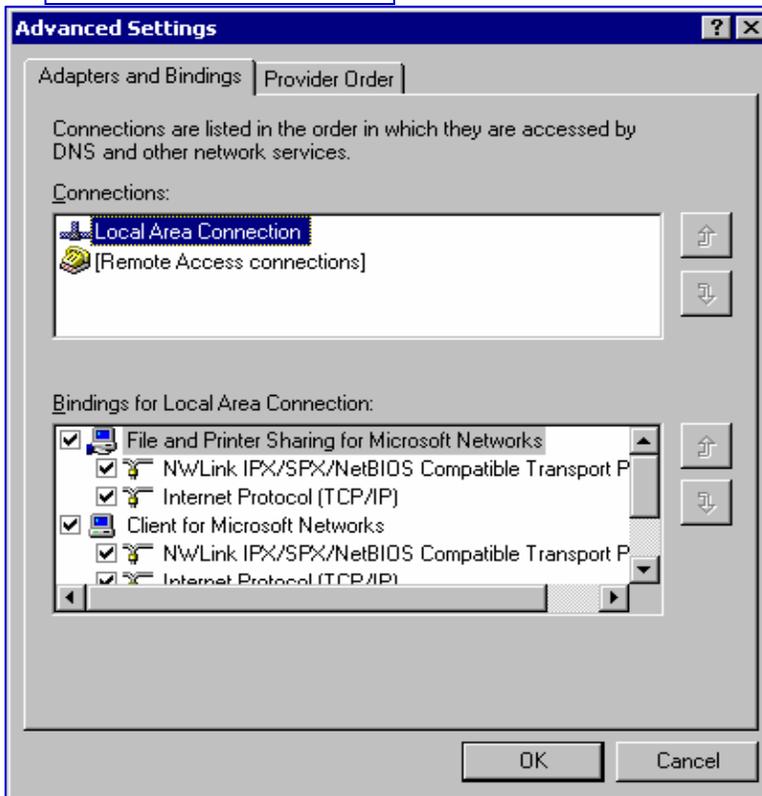
The screenshot shows the 'Advanced Settings' dialog box, 'Adapters and Bindings' tab. It displays a list of connections and a list of bindings for the selected 'Local Area Connection', including 'File and Printer Sharing for Microsoft Networks', 'Internet Protocol (TCP/IP)', and 'Client for Microsoft Networks'.

The screenshot shows the 'NwLink IPX/SPX Properties' dialog box, 'General' tab. It displays configuration options for the network adapter, including 'Adapter', 'Frame Type', and 'Network Number'.

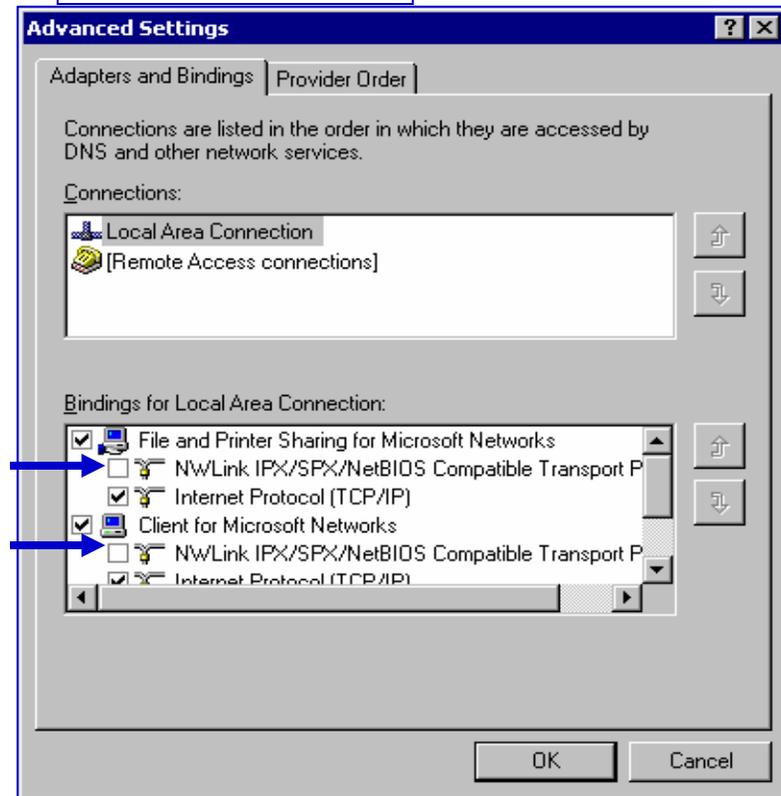
Windows 2000/XP Bindings Cleanup



Default Bindings



'Clean' Bindings



The IPX boxes should be unchecked since Microsoft clients use IP. If IPX is used for Netware/Novell Clients, bind it only to the Netware Client.

Broadcast Storms... The Truth!



T

- Many technologists believe that by installing more bandwidth and collapsing network architectures', the threat of a broadcast storm disappears.
- When several segments are collapsed into one large one, the chance of a broadcast storm increases.
- In summary, when you collapse many separate segments into one large one, the workstations/servers will have to process more broadcast packets. So how do you identify and minimize the source of your broadcast packets?
- You should plug a protocol aware tool into any switch port configured for a client VLAN and observe the broadcast protocols. NO spanning or mirroring is necessary.

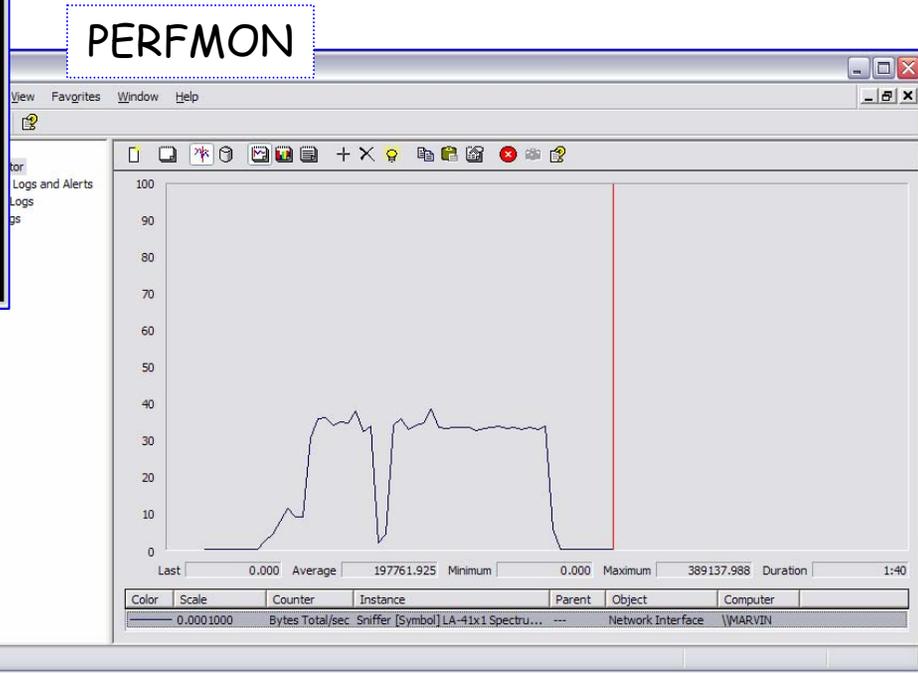
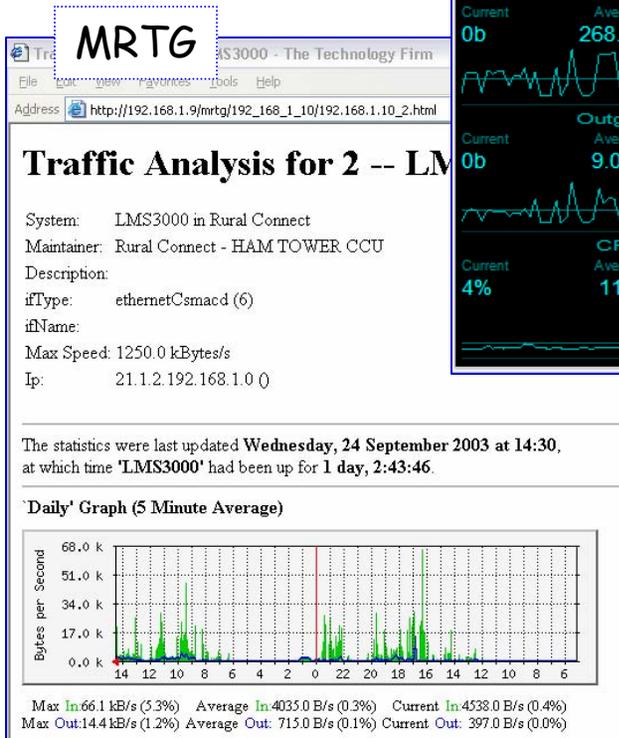
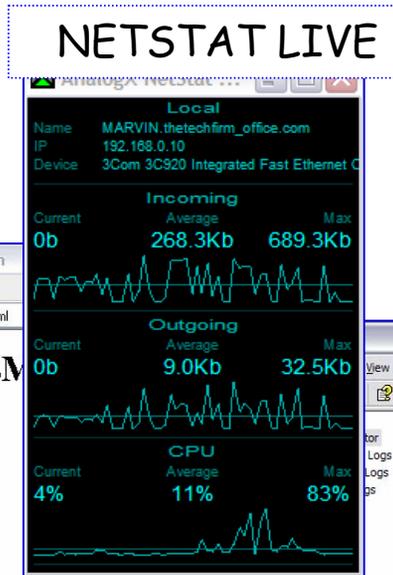
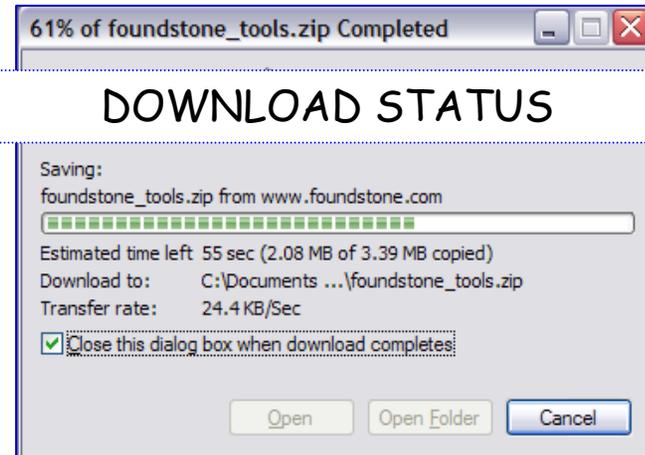
Why bother with a little broadcast packet?

- With the consolidation of multiple segments and the behavior of automatic settings, the number of broadcasts have been increasing over the past few years.
- Many applications rely on broadcasts to locate servers.
- Every broadcast packet requires an interrupt on the listening station for processing. Regardless if the client acts on the packet or not.

Getting A Line On Baselineing



- Different Tools, different results.
- Calibrate your tools.
- Understand what they are reporting.



Common Baselines



- Baselining is typically an afterthought and most people don't baseline for the following reasons:
 - ✓ When should I start?
 - ✓ What should a baseline look like?
 - ✓ What is a *good baseline*?
 - ✓ What if someone actually reads it and wants more information?
 - ✓ Why bother? Everything changes so fast around here.
 - ✓ I don't know how to baseline.
- Baselines should be:
 - ✓ Clearly defined. For example Bootup baseline, Login baseline, Application Baseline or Upgrade Baseline.
- As long as your goal is clear and the methodology is documented (and consistent), your baseline is correct.
- If you have performed a baseline correctly, you will typically find problems to fix along the way.
- Good tools to make baselining easier; (*free*)
 - ✓ AnalogX Netstat Live (www.analogx.com)
 - ✓ Microsoft Performance Monitor
 - ✓ Camstudio Screen Recorder (<http://www.rendersoftware.com>)
 - ✓ Gadwin Print Screen (<http://www.gadwin.com>)
 - ✓ Auto IT Scriptor (<http://www.hiddensoft.com/AutoIt>)
 - ✓ Netpeeker (www.netpeeker.com)
 - ✓ DRTCP (www.broadbandreports.com)

Fluke Optiview Console



The screenshot displays the Optiview Console Viewer interface for a network named MARVIN. The window title is "Optiview Console Viewer [MARVIN (129.196.203.000) - Archived 11:33 AM 1/19]". The interface includes a menu bar (File, View, Agent, Problem, Device, Tools, Reports, Network Map, Help) and a toolbar with icons for Agent, Print, Refresh, Options, Sort, Trace SR, Tools, Reports, Net Map, and RMON. Below the toolbar are tabs for Overview, Detail (selected), Trending, and Key Devices. A left-hand navigation tree shows a hierarchy of network components, with "Devices" expanded. The main pane displays a table of device details.

Name	NetBIOS Na...	IP Address	MAC Address
CHOW35024T03		172.020.160.003	00E07BDEAE60
LXKF38C9F		172.020.160.110	000400CF31F9
CHOW35024T05		172.020.160.005	00802D0FCAE1
LXKF677B4		172.020.160.106	0004006FEE2D
172.020.160.001		172.020.160.00...	000997BEE2C8
CHOW35024T06		172.020.160.006	00802DDD1AB2
CHOW35024T04		172.020.160.004	0060FDAB6ABD
327CHOW-A002	327CHOW-A002	172.020.160.202	0008743187E6
327CHOW-A003	327CHOW-A003	172.020.160.203	000874318450
327CHOW-A004	327CHOW-A004	172.020.160.212	0008743187D5
327CHOW-A005	327CHOW-A005	172.020.160.206	0008743187AE
327CHOW-A006	327CHOW-A006	172.020.160.208	0008743186CB
327CHOW-A007	327CHOW-A007	172.020.160.204	000874318A76
327CHOW-A008	327CHOW-A008	172.020.160.205	0008743187C9
327CHOW-A009	327CHOW-A009	172.020.160.207	0008743187BC
327CHOW-A010	327CHOW-A010	172.020.160.226	00C04F4D6866
327CHOW-A011	327CHOW-A011	172.020.160.217	00C04F4D6598
327CHOW-A012	327CHOW-A012	172.020.160.216	00C04F4D68E6
327CHOW-A013	327CHOW-A013	172.020.160.209	00C04F4D82AD
327CHOW-A015	327CHOW-A015	172.020.160.210	00C04F4D690F

NetBIOS Reporting



- Ensure that only one protocol is bound to the Microsoft client and Novell Client.

(200.200.240.000)

NetBIOS Inventory						
<u>Domain/Name</u>	<u>MAC Address</u>	<u>IP Address</u>	<u>OS Type</u>	<u>Protocols</u>		
				<u>IP</u>	<u>IPX</u>	<u>NetBEUI</u>
ADMIN						
CENSORNET327 <i>NetBIOS Services: NT Server</i>	000874-3185a8	172.020.160.051	NT Server 4.5	+		
BASGROUP						
BAS327 <i>NetBIOS Services: Master Browser</i>	3Com-36b09e	172.020.160.090	Windows 95/98	+		
LEARNING						
327CHOW-A002	000874-3187e6	172.020.160.202	Windows XP	+		

Wrong IP Mask



Fluke Optiview Console Report

IP Inventory			
<u>Subnet/Name</u>	<u>MAC Address</u>	<u>IP Address</u>	<u>Subnet Mask</u>
Subnet: 172.016.050.000			255.255.255.000
172.016.050.022	D-Link-f5297e	172.016.050.022	
172.016.050.026	TEKLOG-003951	172.016.050.026	
172.016.050.027	TEKLOG-003958	172.016.050.027	
172.016.050.028	TEKLOG-003e0e	172.016.050.028	
172.016.050.065	SYMBOL-483f06	172.016.050.065	
172.016.050.067	SYMBOL-436acd	172.016.050.067	
172.016.050.068	SYMBOL-43b43b	172.016.050.068	
172.016.050.069	SYMBOL-44014f	172.016.050.069	
172.016.050.070	SYMBOL-43b120	172.016.050.070	
172.016.050.071	SYMBOL-43b81b	172.016.050.071	255.255.000.000
 172.016.050.226	IBM-dede12	172.016.050.226	
<i>IP Services:</i>		<i>POP3, HTTP</i>	

Unnecessary Services



Fluke Optiview Console Report

IP Inventory							
Subject/Name	MAC Address	IP Address	Subnet Mask	SNMP	IP Services		
					HTTP	E-mail	Print
Subnet: 172.016.050.000			255.255.255.000				
172.016.050.022	D-Link-f5297e	172.016.050.022			•		
172.016.050.026	TEKLOG-003951	172.016.050.026					
172.016.050.027	TEKLOG-003958	172.016.050.027					
172.016.050.028	TEKLOG-003e0e	172.016.050.028					
172.016.050.065	SYMBOL-483f06	172.016.050.065					
172.016.050.067	SYMBOL-436acd	172.016.050.067					
172.016.050.068	SYMBOL-43b43b	172.016.050.068					
172.016.050.069	SYMBOL-44014f	172.016.050.069					
172.016.050.070	SYMBOL-43b120	172.016.050.070					
172.016.050.071	SYMBOL-43b81b	172.016.050.071	255.255.000.000				
172.016.050.226	IBM-dede12	172.016.050.226		•	•	•	
<i>IP Services: POP3, HTTP</i>							
172.016.050.230	HP-6e1733	172.016.050.230		•	•		•
172.016.050.235	HP-2b6cd1	172.016.050.235		•	•		•
172.016.050.236	HP-2b8cf8	172.016.050.236		•	•		•
172.016.050.237	HP-23f610	172.016.050.237		•	•		•
172.016.050.238	HP-2a2b1f	172.016.050.238		•	•		•
172.016.050.241	INTEL-60dc79	172.016.050.241		•			•
172.016.050.242	INTEL-60dd45	172.016.050.242		•			•
172.016.050.245	HP-29884e	172.016.050.245		•	•		•

Duplex Issues



- The 200 Mbits/sec means that this interface is 100 MB, but the fact that there are more than one device, indicates a possible hub or switch.

Fluke Optiview Console Report

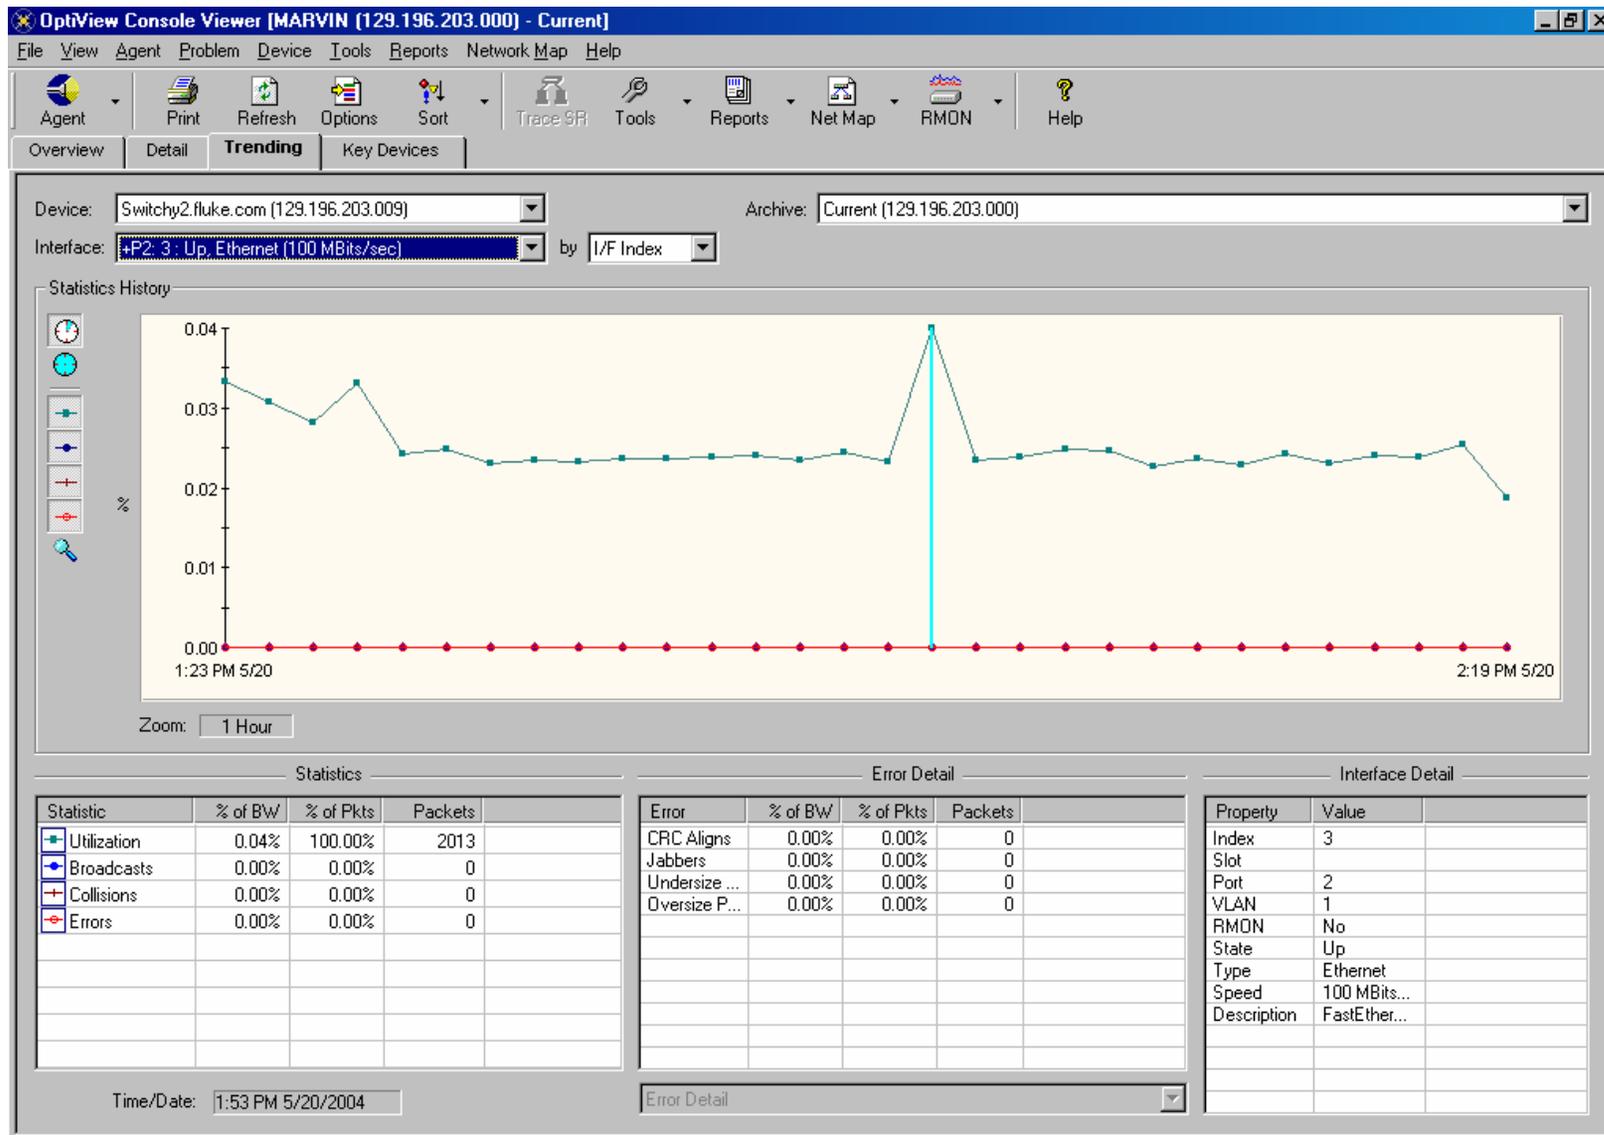
VLAN 1: Slot 1 : Port 12		(BayStack 450-24T - 12)		
Interface:	Up, Ethernet, RMON	Directly Connected Devices		
Speed:	200 Mbits/sec	<u>Name</u>	<u>MAC Address</u>	<u>IP Address</u>
MTU:	1514	OV129750	Dell-50e94d	172.016.050.192
		OV4184	INTEL-566d25	172.016.050.172

Who's on first?



Port 3		(BayStack 350-24T - 3)		
Interface:	Up, Ethernet, RMON	Directly Connected Devices		
Speed:	200 MBits/sec	<u>Name</u>	<u>MAC Address</u>	<u>IP Address</u>
MTU:	1514	327CHOW-A026	000874-3187ec	172.020.160.233
Utilization:	0.00%			
Port 4		(BayStack 350-24T - 4)		
Interface:	Up, Ethernet, RMON	Directly Connected Devices		
Speed:	200 MBits/sec	<u>Name</u>	<u>MAC Address</u>	<u>IP Address</u>
MTU:	1514	327CHOW-A027	000874-31857e	172.020.160.232
Utilization:	0.00%			
Port 5		(BayStack 350-24T - 5)		
Interface:	Up, Ethernet, RMON	Directly Connected Devices		
Speed:	200 MBits/sec	<u>Name</u>	<u>MAC Address</u>	<u>IP Address</u>
MTU:	1514	327CHOW-A005	000874-3187ae	172.020.160.206
Utilization:	0.00%			

Watching the Trends



Dissecting IP



Packet Layout (Ethernet II/Ethertype Format)

	DESTINATION MAC					SOURCE MAC						Type 800		A	B	
0000	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0010	C		D		E		F	G	H		I		J			
0020	J															
0030	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0040	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F

A - IP Version: **4 bits** (version 4) & IP Header Length: **8 bits**.

B - Type of Service: **8 bits**. Routers ignore these values by default

C – Total Length: **2 Bytes**. Indicates the total length of IP header and IP Payload.

D – Identification: **2 Bytes**. Used to identify a specific packet sent between two stations

E – Flags: **3 bits long**. One flag is used for fragmentation, and the other whether or not more fragments are to follow.

Fragment Offset: 13 bits. Used to indicate the offset of where this fragment begins.

F – Time to Live: **1 Byte**. How many links this datagram can travel before an IP router discards it.

G – Protocol: **1 Byte**. Indicates the Upper Layer protocol. (UDP or TCP)

H – Header Checksum: **2 Bytes**. The sending host performs a bit level integrity check on the IP header only.

I – Source Address: **4 Bytes**. Contains the IP address of the source.

J – Destination Address: **4 Bytes**. Contains the IP address of the source.

Microsoft Command – netstat -s



- Good way to a status on all your protocols.
- The following example shows a sample of what is returned when you type netstat -s.

```
C:\netstat -s

IP Statistics
Packets Received           = 2138
Received Header Errors     = 5
Received Address Errors    = 10
Datagrams Forwarded       = 0
Unknown Protocols Received = 0
Received Packets Discarded = 0
Received Packets Delivered = 2133
Output Requests           = 2098
Routing Discards          = 0
Discarded Output Packets   = 0
Output Packet No Route    = 0
Reassembly Required       = 0
Reassembly Successful     = 0
Reassembly Failures       = 0
Datagrams Successfully Fragmented = 0
Datagrams Failing Fragmentation = 0
Fragments Created        = 0

ICMP Statistics
                Received  Sent
Messages        55      14
Errors           0        0
Destination Unreachable  46      5
Time Exceeded   0        0
Parameter Problems 0        0
Source Quenchs  0        0
Redirects       0        0
Echos          0        0
Echo Replies   0        0
Timestamps     0        0
Timestamp Replies 0        0
Address Masks  0        0
Address Mask Replies 0        0

TCP Statistics
Active Opens           = 108
Passive Opens         = 1
Failed Connection Attempts = 0
Reset Connections     = 55
Current Connections   = 0
Segments Received     = 1278
Segments Sent         = 1367
Segments Retransmitted = 24

UDP Statistics
Datagrams Received    = 789
No Ports              = 66
Receive Errors        = 0
Datagrams Sent        = 694
```

ICMP Packet Info



	Layer 2 Information										IP					
0000	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
	IP Information															
0010	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
	IP		A		B		C		D		E		F			
0020	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
											F					
0030	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
											F					
0040	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F

Legend

- A = Type
 - ✓ The TYPE field identifies the ICMP message.
- B = Code
 - ✓ The CODE field provides further information about the associated TYPE field.
- C = Checksum
 - ✓ The CHECKSUM provides a method for determining the integrity of the message.
- D = Identifier
 - ✓ This field is used to correlate ICMP commands and responses.
- E = Sequence Number
 - ✓ This value is used to number commands/responses.
- F = Optional Data

Why should I 'Pathping'? (Windows 2000/XP)

- Provides information about network latency and network loss at intermediate hops between a source and destination.
- Transmits multiple ICMP Echo Request messages to each router between a source and destination over a period of time and then computes results based on the packets returned from each router.
- Because pathping displays the degree of packet loss at any given router or link, you get an idea which routers or subnets might be having network problems.
- Pathping performs the equivalent of the `tracert` command by identifying which routers are on the path.
- It then sends pings periodically to all of the routers over a specified time period and computes statistics based on the number returned from each.
- If you type `pathping` without parameters, you display the helpscreen.

Pathping Results



	Source to Here	This Node/Link	
Hop	RTT	Lost/Sent = Pct	Lost/Sent = Pct Address
0			MARVIN.Traflagar Rd N [10.44.10.145]
		0/ 100 = 0%	
1	16ms	0/ 100 = 0%	0/ 100 = 0% 10.44.10.10
		0/ 100 = 0%	
2	53ms	0/ 100 = 0%	0/ 100 = 0% 192.168.1.1
		20/ 100 = 20%	20% packet loss
8	114ms	20/ 100 = 20%	0/ 100 = 0% p4-7-3-0.r01.cncg100.us.bb.verio.net [129.250.9.189]
		4/ 100 = 4%	
9	102ms	24/ 100 = 24%	0/ 100 = 0% p16-7-0-0.r01.chcgil01.us.bb.verio.net [129.250.5.71]
		2/ 100 = 2%	
10	139ms	28/ 100 = 28%	2/ 100 = 2% p16-1-0-1.r20.asbnva01.us.bb.verio.net [129.250.5.103]
		0/ 100 = 0%	
11	108m	29/ 100 = 29%	3/ 100 = 3% p16-7-0-0.r02.asbnva01.us.bb.verio.net [129.250.2.83]
		0/ 100 = 0%	
12	115ms	26/ 100 = 26%	0/ 100 = 0% ge-1-1.a00.asbnva01.us.ra.verio.net [129.250.26.97]
		0/ 100 = 0%	
13	131ms	26/ 100 = 26%	0/ 100 = 0% ge-3-2.a00.asbnva01.us.ce.verio.net [168.143.105.58]
		1/ 100 = 1%	
14	---	100/ 100 = 100%	73/ 100 = 73% 216.239.47.102
		0/ 100 = 0%	
15	100ms	27/ 100 = 27%	0/ 100 = 0% 216.239.51.99

Trace complete.

IP Troubleshooting Commands



- Ping
 - ✓ May not help with throughput issues.
 - ✓ Used to verify if a IP Host is 'up'. In other words a reachability test.
 - ✓ UDP, TCP and port numbers are not tested.
 - ✓ Can be used to test Maximum MTU with the -f -l size option.
 - ✓ Some applications will put specific data as a signature
(NetTool; Jamie & Ted's Adventure)

- Pathping
- tracert or trace
 - ✓ Used to document a path from and to specified hosts.
 - ✓ Is just a ping with the Time To Live value starting from 1, incremented by one.

- Netstat -a
 - ✓ Displays Ethernet, IP, TCP and UDP statistics.

- Route print
 - ✓ Displays your local routing table.

UDP Overview

No negotiation or session establishment.

No sequencing or acknowledgements.

Identifies both source and destination port numbers.

Provides checksum of the entire packet, if implemented.

No Buffering or Flow Control.

No Segmentation for large blocks of data. Therefore applications must send data in small enough 'chunks' so that they are not larger than the IP MTU.

UDP HEADER



Source Port:

2 Byte field to identify the source application layer protocol.

Destination Port:

2 Byte field to identify the destination application layer protocol.

Length:

2 Byte field. Indicates the length of the UDP header and message. Value ranges from 8 to 65,515 bytes.

Checksum:

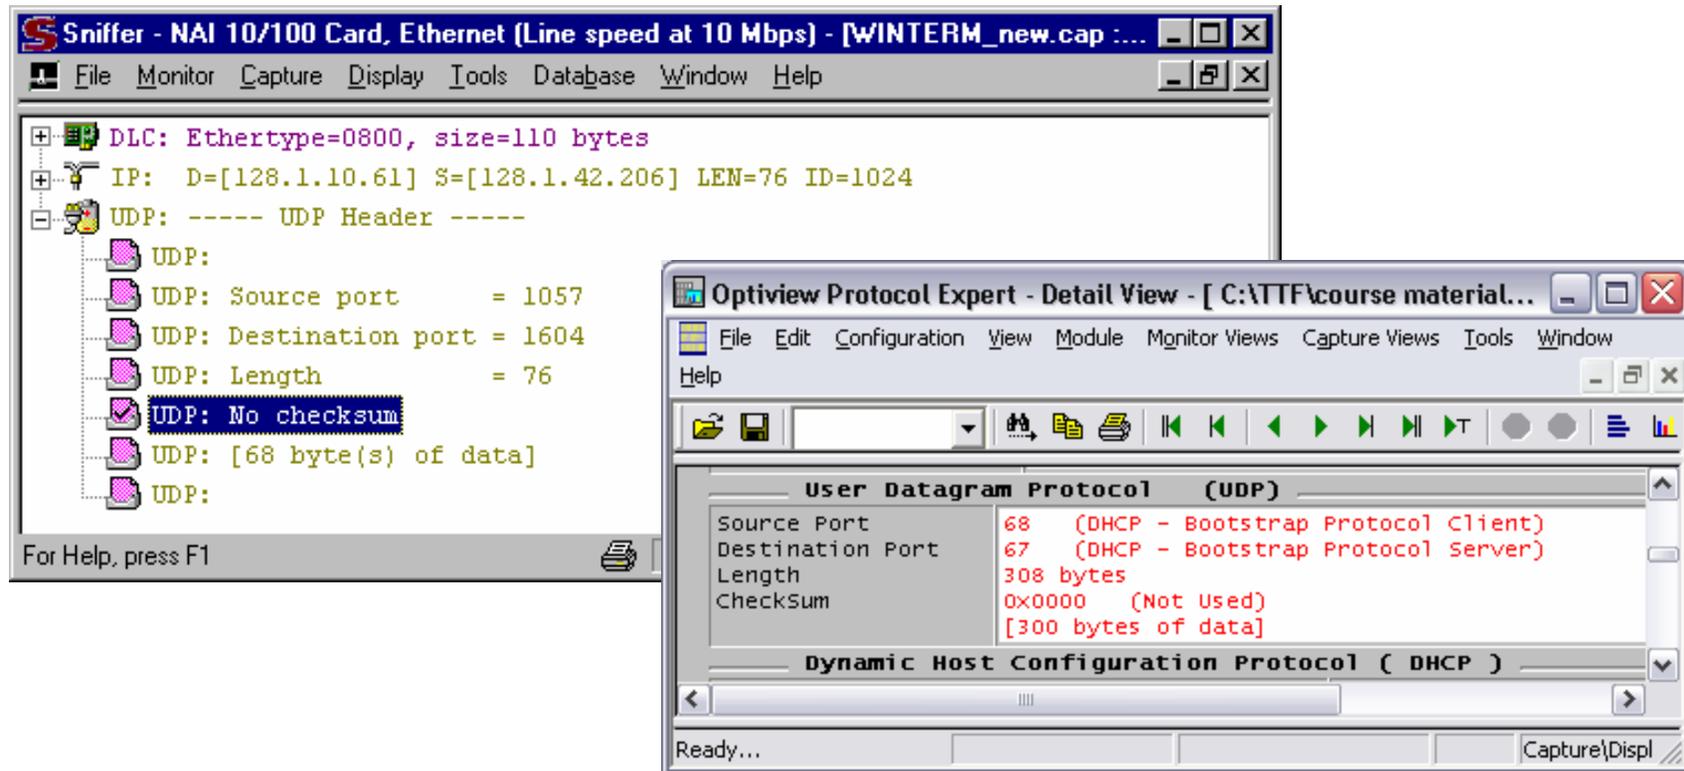
2 Byte field to provide a bit level integrity check for UDP. This value and optional and explained in more detail in another slide.

Windows 2000 always calculates a UDP checksum value.

UDP Checksums



There is a common implementation error in UDP checksums. Unlike the TCP checksum, the UDP checksum is optional; the value zero is transmitted in the checksum field of a UDP header to indicate the absence of a checksum. If the transmitter really calculates a UDP checksum of zero, it must transmit the checksum as all 1's (65535).



UDP Flooding



By default Routers do not forward IP broadcasts.

With the introduction of DHCP, multicast and broadcast applications analysts are modifying their routers to forward broadcasts.

With the addition of high density switches more clients may end up on the same broadcast domain [or VLAN].

Microsoft and NetBIOS applications utilize UDP port 137, 138 & 139.

Increasing the number of broadcasts on your segments may have catastrophic effects.

- ✓ Every broadcast received by a node requires an interrupt.
- ✓ Too many interrupts per second will hamper performance.
- ✓ Routers may experience buffer overflow symptoms.
- ✓ When this happens, bandwidth utilization is light and CPU load is largely unaffected.

Try to avoid generic UDP flooding and implement specific UDP port forwarding when possible.

TCP Header



Source Port:

2 Bytes to identify the source application layer protocol.

Destination Port:

2 Bytes to identify the destination application layer protocol.

Sequence Number:

4 Bytes. Indicates the outgoing bytes stream sequence number. When no data is to be sent the sequence number will be set to the next octet.

Acknowledgement Number:

4 Bytes. Provides a positive acknowledgement of all octets in the incoming byte stream.

Data Offset:

4 bits. Indicates where the TCP segment data begins.

Reserved:

6 bits. For future use.

Flags:

6 bits. Indicates one of six different flags.

Window:

2 Bytes. Indicates the number of Bytes of available space in the receive buffer of the sender.

Checksum:

2 Bytes. 2 Byte field to provide a bit level integrity check.

Urgent Pointer:

2 Bytes. Indicates the location of urgent data in the segment.

Options:

Indicates additional TCP Options.

TCP Three-way Handshake



The screenshot shows a Wireshark capture of a TCP three-way handshake. The table below represents the data shown in the packet list pane:

No.	Delta	Destination	Source	Protocol	Info
1	0.000000	192.168.1.216	192.168.1.210	TCP	1761 > ftp [SYN] Seq=575330 Ack=0 win=8192 Len=0
2	0.000335	192.168.1.210	192.168.1.216	TCP	ftp > 1761 [SYN, ACK] Seq=310955 Ack=575331 win=8760 Len=0
3	0.000243	192.168.1.216	192.168.1.210	TCP	1761 > ftp [ACK] Seq=575331 Ack=310956 win=8760 Len=0

Initiating Application Port Number (Port 21) FTP

Sequence Number + 1

The screenshot shows a NetworkMiner capture of a TCP three-way handshake. The table below represents the data shown in the packet list pane:

No.	Dest Address	Source Address	Summary
1	[207.219.36.66]	[142.86.42.187]	TCP: D=21 S=1027 SYN SEQ=156232 LEN=0 WIN=8192
2	[142.86.42.187]	[207.219.36.66]	TCP: D=1027 S=21 SYN ACK=156233 SEQ=2209979179 LEN=0 WIN=17520
3	[207.219.36.66]	[142.86.42.187]	TCP: D=21 S=1027 ACK=2209979180 WIN=8760

The delta value between frames 1 and 2 can be used as a TCP transport connect baseline value.

Other important information gathered from this handshake:

- Window Size
- SACK
- Maximum Segment Size
- Window Scale Option value

TCP MSS Issue



- Ensure that your servers are using the optimal TCP Maximum Segment Size. The Maximum for 10/100 Ethernet is 1460.

```
Transmission Control Protocol (TCP)
Source Port          524 (NCP over IP)
Destination Port    1029
Sequence Number     1891468473
Acknowledgement Number 619956437
Header Length       0x60
                   0110 .... 24 bytes - Header Length
                   .... 0000 Not Used
Flags               0x12
                   00.. .... Not Used
                   ..0. .... No URG
                   ...1 .... Acknowledgement
                   .... 0... No PSH
                   .... .0.. No RST
                   .... ..1. Synchronize
                   .... ...0 No FIN
Window Size         6144
Checksum            0x0524 (Correct)
Urgent Pointer      0
Options Present
Kind                2 (Maximum Segment Size)
Length              4 bytes
Maximum Segment Size 536
```

TCP Header NAI and Ethereal



Sniffer Portable - Local, Ethernet (Line speed at 100 Mbps) - ...

File Monitor Capture Display Tools Database Window Help

TCP: ----- TCP header -----

- TCP:
- TCP: Source port = 3547
- TCP: Destination port = 21 (FTP-ctrl)
- TCP: Initial sequence number = 1628246031
- TCP: Next expected Seq number = 1628246032
- TCP: Data offset = 32 bytes
- TCP: Reserved Bits: Reserved for Future Use (Not set)
- TCP: Flags = 02
- TCP: ..0. = (No urgent pointer)
- TCP: ...0 = (No acknowledgment)
- TCP: 0... = (No push)
- TCP:0.. = (No reset)
- TCP:1. = SYN
- TCP:0 = (No FIN)
- TCP: Window = 29856
- TCP: Checksum = 6603 (correct)
- TCP: Urgent pointer = 0
- TCP:
- TCP: Options follow
- TCP: Maximum segment size = 1460
- TCP: No-Operation
- TCP: Window scale Option
- TCP: Window scale factor = 3
- TCP: No-Operation
- TCP: No-Operation
- TCP: SACK-Permitted Option

Expert Decode Matrix Host Table Protocol Dist. Statistics

For Help, press F1

ftp active.enc - Ethereal

File Edit View Capture Analyze Help

No.	Abs Time	Size	Destination
5	18:40:01.799996	66	207.46.133.140

Frame 5 (66 bytes on wire, 66 bytes captured)

- Ethernet II, Src: 00:08:74:e1:28:c2, Dst: 00:08:00:02:00:04
- Internet Protocol, Src Addr: 192.168.0.2 (192.168.0.2), Dst Addr: 207.46.133.140
- Transmission Control Protocol, Src Port: 3547, Dst Port: 21
- Source port: 3547 (3547)
- Destination port: 21 (21)
- Sequence number: 1628246031
- Header length: 32 bytes
- Flags: 0x0002 (SYN)
- 0... = Congestion Window Reduced (CWR): Not set
- .0.. = ECN-Echo: Not set
- ..0. = Urgent: Not set
- ...0 = Acknowledgment: Not set
- 0... = Push: Not set
-0.. = Reset: Not set
-1. = Syn: Set
-0 = Fin: Not set
- Window size: 29856
- Checksum: 0x6603 (correct)
- Options: (12 bytes)
- Maximum segment size: 1460 bytes
- NOP
- Window scale: 3 (multiply by 8)
- NOP
- NOP
- SACK permitted

0000 00 50 ba 10 93 fb 00 08 74 e1 28 c2 0 ▶

0010 00 34 18 cb 40 00 80 06 cc 93 c0 a8 0 ▶

0020 85 8c 0d db 00 15 61 0d 10 0f 00 00 0 ▶

0030 74 a0 66 03 00 00 02 04 05 b4 01 03 0 ▶

0040 04 02 ▶

Filter:

TCP Header Fluke and Wildpackets



Fluke Optiview Protocol Expert - Detail View - [C:\Documen...]

File Edit Configuration View Histogram Module Monitor Views
Capture Views Tools Window Help

Transmission Control Protocol (TCP)

Source Port	3547
Destination Port	21 (File Transfer [Control])
Sequence Number	1628246031
Acknowledgement Number	0
Header Length	0x80
Flags	0x02
Window Size	29856
Checksum	0x6603 (Correct)
Urgent Pointer	0
Options Present	2 (Maximum Segment Size)
Kind	4 bytes
Maximum Segment Size	1460
Kind	1 (No Operation)
Kind	3 (Window Scale Factor)
Length	3 bytes
Shift Count	3
Kind	1 (No Operation)
Kind	1 (No Operation)
Kind	4 (TCP Selective Ack Permitted)
Length	2 bytes

Data/FCS

Data/Padding	[0 bytes]
Frame Check Sequence	0x8A9653FB (Correct)

Ready... Capture\

EtherPeek NX - [ftp active.enc]

File Edit View Capture Send Monitor Tools Window Help

Packet: 5 [X]

TCP - Transport Control Protocol

Source Port:	3547
Destination Port:	21 ftp
Sequence Number:	1628246031
Ack Number:	0
TCP Offset:	8 (32 bytes)
Reserved:	00000000
TCP Flags:	00000010 S.
Window:	29856
TCP Checksum:	0x6603
Urgent Pointer:	0

TCP Options:

Option Type:	2 Maximum Segment Size
Length:	4
MSS:	1460
Option Type:	1 No Operation
Option Type:	3 Window Scale Factor
Length:	3
Shift Count:	3
Option Type:	1 No Operation
Option Type:	1 No Operation
Option Type:	4 Sack Supported
Length:	2

FCS - Frame Check Sequence

Packets | Nodes | Protocols | Summary | Graphs | Log | Expert | Peer Map |

Done Local Area Connection

TCP FIN



The FIN (Finish) packet is initiated by the Client.

ACK = Sequence Number + 1

No.	Dest Address	Source Address	Summary
10	[192.168.1.102]	[192.168.1.101]	TCP: D=389 S=1029 FIN ACK=186469 SEQ=3918485085 LEN=0 WIN=5840
11	[192.168.1.101]	[192.168.1.102]	TCP: D=1029 S=389 ACK=3918485086 WIN=8459
12	[192.168.1.101]	[192.168.1.102]	TCP: D=1029 S=389 FIN ACK=3918485086 SEQ=186469 LEN=0 WIN=8459
13	[192.168.1.102]	[192.168.1.101]	TCP: D=389 S=1029 ACK=186470 WIN=5840

The FIN (Finish) packet is then sent by the server.

ACK = Sequence Number + 1

No.	Dest Address	Source Address	Summary
1	[192.168.1.101]	[192.168.1.103]	TCP: D=1024 S=389 FIN ACK=2891180479 SEQ=188168 LEN=0 WIN=8250
2	[192.168.1.103]	[192.168.1.101]	TCP: D=389 S=1024 ACK=188169 WIN=5840

Monitoring TCP Connections With netstat



```
C:\WINDOWS\Desktop>netstat -an
```

Active Connections

Proto	Local Address	Foreign Address	State
TCP	0.0.0.0:1033	0.0.0.0:0	LISTENING
TCP	0.0.0.0:1037	0.0.0.0:0	LISTENING
TCP	10.10.10.169:1025	10.10.10.150:139	TIME_WAIT
TCP	10.10.10.169:137	0.0.0.0:0	LISTENING
TCP	10.10.10.169:138	0.0.0.0:0	LISTENING
TCP	10.10.10.169:139	0.0.0.0:0	LISTENING
TCP	127.0.0.1:1034	0.0.0.0:0	LISTENING
TCP	216.254.151.2:1033	66.185.95.101:110	ESTABLISHED
TCP	216.254.151.2:1037	216.136.175.45:80	ESTABLISHED
TCP	216.254.151.2:137	0.0.0.0:0	LISTENING
TCP	216.254.151.2:138	0.0.0.0:0	LISTENING
TCP	216.254.151.2:139	0.0.0.0:0	LISTENING
UDP	10.10.10.169:137	*.*	
UDP	10.10.10.169:138	*.*	
UDP	127.0.0.1:1034	*.*	
UDP	216.254.151.2:137	*.*	
UDP	216.254.151.2:138	*.*	

Many programs exist that will display this same information in a GUI. (ie Netpeeker)

Full association consists of two endpoints may also includes the protocol name.

Half association only has one endpoint that also includes the protocol name.

```
C:\WINDOWS\Desktop>
```

TCP Port Numbers

- Port numbers are divided into three ranges: the Well Known Ports, the Registered Ports, and the Dynamic and/or Private Ports.
- The Well Known Ports are those from 0 through 1023.
- The Registered Ports are those from 1024 through 49151
- The Dynamic and/or Private Ports are those from 49152 through 65535
<http://www.isi.edu/in-notes/iana/assignments/port-numbers>
- Ports are used in the TCP [RFC793] to name the ends of logical connections which carry long term conversations. For the purpose of providing services to unknown callers, a service contact port is defined. This list specifies the port used by the server process as its contact port. The contact port is sometimes called the "well-known port".
- One simple security technique is to modify a default application TCP port number.
- Ensure that you have all your port numbers documented.
- Ensure that you are able to test various TCP/UDP port numbers with your troubleshooting tool of choice.
- Another benefit from many tools is the ability to add custom TCP/UDP port numbers that correspond to legacy applications.
- To discover all specific types of servers; [ie http servers] simply use software that can check for port 80 across your network.

TCP Receive Window Timer



- In the receiving window, a Delayed Acknowledgment Timer is set for those packets that arrive out of order. Remember, by default an acknowledgment is sent for every two sequenced packets, starting from the left-hand side of the window. If packets arrive out of order (if, for instance, 1 and 3 arrive but 2 is missing), an acknowledgment for two sequenced packets is not possible. When packets arrive out of order, a Delayed Acknowledgment Timer is set on the first packet in the pair. In the parenthetical example, a Timer is set on packet number 1.
- The Delayed Acknowledgment Timer is hard-coded for 200 milliseconds, or 1 /5 the Retransmit Timer. If packet 2 does not show up before the Delayed Acknowledgment Timer expires, an acknowledgment for packet 1, and only packet 1, is sent. No other acknowledgments are sent, including those for packets 3 through 8 that might have appeared. Until packet 2 arrives, the other packets are considered interesting, but useless. As data is acknowledged and passed to the Application layer, the receive window slides to the right, enabling more data to be received. Again though, if a packet doesn't show up, the window is not enabled to slide past it.

Configuring Delayed Acknowledgments



- A problem may occur with message block (SMB) write operations to a Windows 2000-based domain controller and may experience a delay of up to 200 milliseconds between file copies.
- If you review a trace of the problem, you notice that the delay occurs after the client sends the server an SMB Notify Change command with the FID entry that matches the FID entry of the target folder.
- Windows Explorer posts a Notify Change request on the network share, which asks to be notified if something changes in the folder that appears in the right pane of Windows Explorer. If a domain controller receives the Notify Change request, it does not respond to it immediately; it does not send packets for up to 200 milliseconds. At that point, a simple Transmission Control Protocol (TCP) acknowledgement (ACK) packet is sent and the file operation resumes as usual.
- This behavior is a result of the interaction between two core networking components of Windows 2000, TCP delayed ACKs, and thread prioritization on domain controllers.
- The Windows 2000-based domain controller, you can edit the TcpDelAckTicks registry value to adjust the TCP delayed ACK timer. If you change the TCP delayed ACK timer to a lower value, the server sends an ACK packet more frequently but at shorter intervals.
 1. Start Registry Editor (Regedt32.exe).
 2. Locate and click the following key in the registry, where Adapter GUID is the globally unique identifier (GUID) for the network adapter that connects to the clients:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces\Adapter GUID
 3. On the Edit menu, click Add Value, and then add the following registry value:
Value name: TcpDelAckTicks
Data type: REG_DWORD
Value data: You can set this value to a range from 0 to 6. The default setting is 2 (200 milliseconds).
 4. Quit Registry Editor.
 5. Restart Windows for this change to take effect.

TCP Window Information



- The Expert reports “Window Frozen” statistics.
 - ✓ Look to see if one particular device is freezing when the window is small. Why is it unable to use a reasonable size window?

- In the “Silly Window Syndrome,” the receiver keeps advertising a small window and sender keeps filling it with small packets.
 - ✓ Can you configure it to use a larger size?

- “Zero Window” symptom alerts you to stations that have closed their window.
 - ✓ Don’t worry if the window closes briefly at the beginning of a connection, then opens and maintains a reasonable size.
 - ✓ Do worry if a host frequently closes the window for long periods of time.
 - ✓ Do you see the window gradually growing smaller before it closes?

- A TCP MAY keep its offered receive window closed indefinitely. As long as the receiving TCP continues to send acknowledgments in response to the probe segments, the sending TCP MUST allow the connection to stay open.

TCP Window Scaling Option



- The computer will send a packet offering the Window Scale option, with a scaling factor of up to 5. If the target computer responds, accepting the Window Scale option in the SYN-ACK, then it is understood that any TCP window advertised by this computer needs to be left-shifted 5 bits from this point onward (the SYN itself is not scaled).
- The Large Windows option defines an implicit scale factor, which is used to multiply the window size value found in a TCP header to obtain the true window size. The TCP/IP stack supports a maximum window size of 1 GB. This Large Window option is negotiated when the TCP connection is established.
- The TCP Large Windows size is useful on fast networks (such as Gigabit Ethernet) with large round-trip times. To understand how this works, think of a water hose. To achieve maximum water flow, the hose should be full. As the hose increases in diameter and length, the volume of water necessary to keep it full increases. In networks, diameter equates to bandwidth, length is measured as round-trip time, and the volume of water is analogous to the TCP window size. On fast networks with large round-trip times, the TCP window size must be increased to achieve maximum TCP bandwidth.
- TCP performance depends not upon the transfer rate itself, but rather upon the product of the transfer rate and the round-trip delay. This "bandwidth delay product" measures the amount of data that would fill the pipe. It is the buffer space required at the sender and the receiver to obtain maximum throughput on the TCP connection over the path—in other words, the amount of unacknowledged data that TCP must handle in order to keep the pipeline full. So on fast networks with large round-trip times, having a large TCP Window helps by allowing for a greater amount of unacknowledged data.
- Windows 2000 uses window scaling automatically if the `TcpWindowSize` is set to a value greater than 64 KB, and the `Tcp1323Opts` registry parameter is set appropriately.

TCP Zero Window Example

- When you get into a Zero Window situation, it is quite normal to see the transmitting station send 1 byte packets.
- These packets are interpreted within 'retransmissions' by most analyzers.

No.	Dest Address	Source Address	Summary	Delta Time	R
28	[192.168.1.210]	[192.168.1.216]	HTTP: R Port=1151 Graphics Data	0.003.044	
29	[192.168.1.210]	[192.168.1.216]	HTTP: R Port=1151 Graphics Data	0.001.229	
30	[192.168.1.210]	[192.168.1.216]	HTTP: R Port=1151 Graphics Data	0.001.259	
31	[192.168.1.216]	[192.168.1.210]	TCP: D=80 S=1151 ACK=334277 WIN=2920	0.000.066	
32	[192.168.1.216]	[192.168.1.210]	TCP: D=80 S=1151 ACK=337197 WIN=0	0.000.586	
33	[192.168.1.210]	[192.168.1.216]	TCP: Retransmitted in frame 35; 21 Bytes of da	0.996.467	
34	[192.168.1.216]	[192.168.1.210]	TCP: D=80 S=1151 ACK=337197 WIN=0	0.000.413	
35	[192.168.1.210]	[192.168.1.216]	TCP: Retransmitted in frame 38; 21 Bytes of da	2.002.551	
36	[192.168.1.216]	[192.168.1.210]	TCP: D=80 S=1151 ACK=337197 WIN=0	0.000.397	
37	[192.168.1.216]	[192.168.1.210]	TCP: D=80 S=1151 ACK=337197 WIN=8760	2.574.534	
38	[192.168.1.210]	[192.168.1.216]	HTTP: R Port=1151 Graphics Data	0.001.578	
39	[192.168.1.210]	[192.168.1.216]	HTTP: R Port=1151 Graphics Data	0.001.229	
40	[192.168.1.210]	[192.168.1.216]	HTTP: R Port=1151 Graphics Data	0.001.236	

TCP Forced ACK



By default, Microsoft Windows TCP/IP will acknowledge every second packet.

No.	Source	Destination	Protocol	Info
777	VAIO	CR584842-A	NBSS	NBSS Continuation Message
778	VAIO	CR584842-A	NBSS	NBSS Continuation Message
779	CR584842-A	VAIO	TCP	1190 > nbssession [ACK] Seq=4655837 Ack=486327956
780	CR584842-A	VAIO	SMB	SMBreadBraw Request
781	VAIO	CR584842-A	NBSS	NBSS Continuation Message
782	VAIO	CR584842-A	NBSS	NBSS Continuation Message
783	CR584842-A	VAIO	TCP	1190 > nbssession [ACK] Seq=4655892 Ack=486330468
784	CR584842-A	VAIO	SMB	SMBreadBraw Request
785	VAIO	CR584842-A	NBSS	NBSS Continuation Message
786	VAIO	CR584842-A	NBSS	NBSS Continuation Message
787	CR584842-A	VAIO	TCP	1190 > nbssession [ACK] Seq=4655947 Ack=486333388
788	VAIO	CR584842-A	NBSS	NBSS Continuation Message
789	VAIO	CR584842-A	NBSS	NBSS Continuation Message
790	CR584842-A	VAIO	TCP	1190 > nbssession [ACK] Seq=4655947 Ack=486336308
791	VAIO	CR584842-A	NBSS	NBSS Continuation Message
792	VAIO	CR584842-A	NBSS	NBSS Continuation Message
793	CR584842-A	VAIO	TCP	1190 > nbssession [ACK] Seq=4655947 Ack=486338664

Output from Ethereal

Ensure you have the proper :

- Internal understanding of who is responsible for various technology components.
- Policies in place to support you.
- Tool for the proper job.
- Training to understand the data that the tool reports back to you.
- Design goals from your client.
- Plan to verify changes resulted the way you had hoped.
- Methodology to use your existing tools in your environment properly.

Start 'Documenting and Baselineing' today, tomorrow is always too late.