

SharkFest `16 Europe

Windows Filesharing De-Mystified: SMB with a Eureka! Effect

October 19th 2016

Eddi Blenkers
packethunter



NetBIOS in the "good old days"





Ethernet: The new Yellow Cable

- Ethernet: Developed in the 1970s
- Became IEEE 802.3 in 1980
 - Unprecedented bandwidth of 10 MBit per second
 - Standardized frame format: DstMAC, SrcMAC, Frame Type
- Offered nearly instant transfer of data in a building
 - Still bound by the law of physics
 - Bits travel at 2/3 of the speed of light



NetBIOS: Remote File Access

- **Back in the days:**
 - BIOS = Basic Input / Output System
 - Among other things, facilitates access to hard disks
 - Operates on a block or sector level: Write 512 bytes to sector 63
- **The new thing: NetBIOS / SMB**
 - Later renamed to CIFS = Common Internet File System
 - Facilitates access to remote disks
 - Operates on a file level: Read / Write / Append for files
 - SMB runs on NetBIOS



*SMB was
designed as
Layer 2 protocol*





NetBIOS Frame Format

```
▸ Frame 9: 61 bytes on wire (488 bits), 61 bytes captured (488 bits) on interface 0
▸ IEEE 802.3 Ethernet
  ▸ Destination: Vmware_31:0d:01 (00:0c:29:31:0d:01)
  ▸ Source: Vmware_8e:87:a6 (00:0c:29:8e:87:a6)
  Length: 47
▸ Logical-Link Control
  ▸ DSAP: NetBIOS (0xf0)
  ▸ SSAP: NetBIOS (0xf0)
  ▸ Control field: U, func=UI (0x03)
▸ NetBIOS
  Length: 44 bytes
  Delimiter: EFFF (NetBIOS)
  Command: Add Name Response (0x0d)
  Status: Add name not in process (0)
  Name type: Unique name (0)
  Transmit Correlator: 0x0003
  ▸ Name to be added: WFW_HOST_1<03> (Messenger service/Main name)
    WFW_HOST_1
    0x03 (Messenger service/Main name)
  ▸ Name to be added: WFW_HOST_1<03> (Messenger service/Main name)
    WFW_HOST_1
    0x03 (Messenger service/Main name)
```



Selected NetBIOS Messages

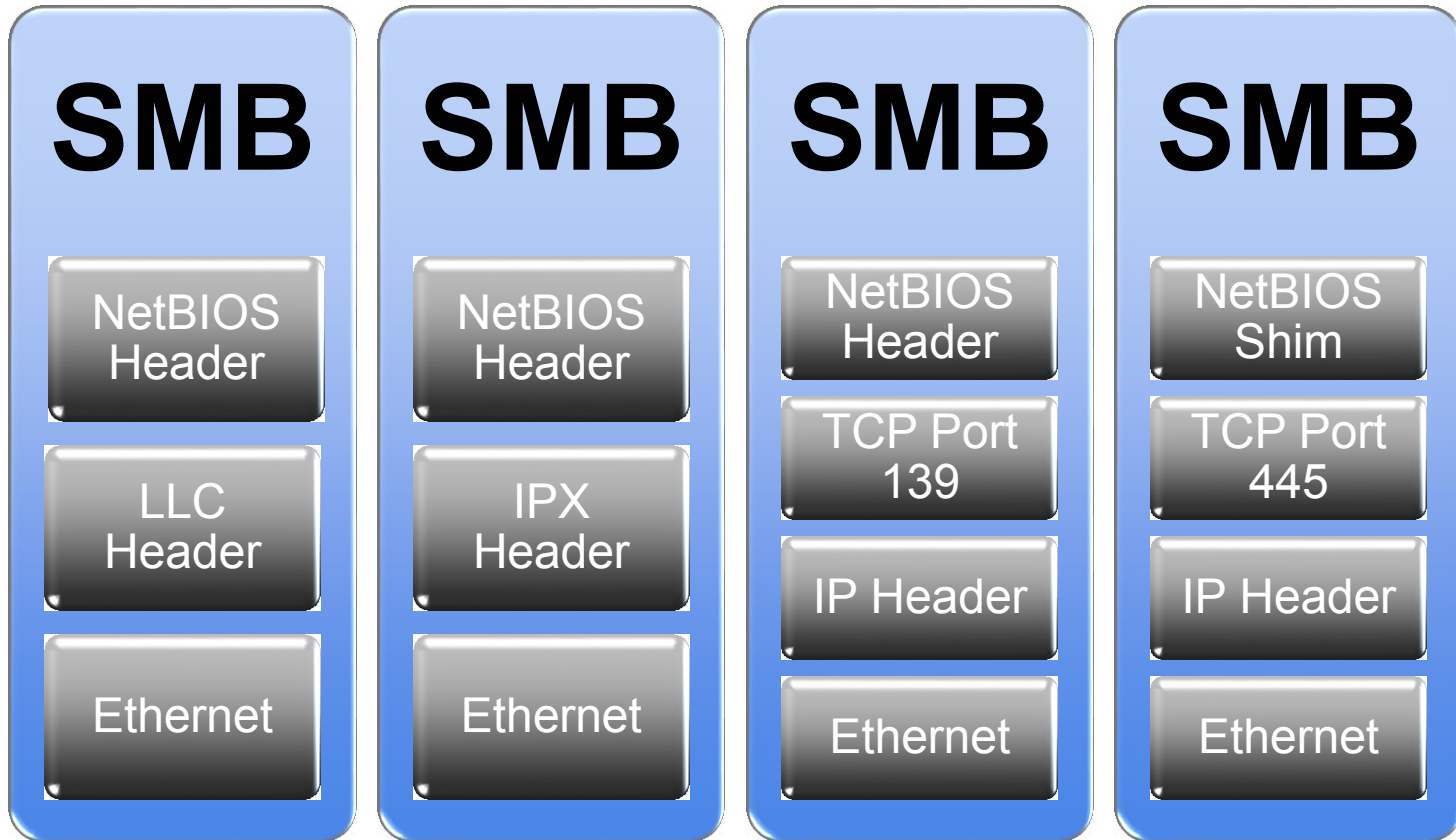
- **Name Registration and Resolution**
 - Translate hostname to MAC address
 - Ensure hostname is unique
 - Also handles group membership
- **Datagram Service, similar to UDP**
- **Session Service, similar to TCP**
 - Session setup is similar to TCP, but 2-way instead of 3-way
 - Keep alive messages for sessions
 - Used by SMB for file sharing



Evolution into an upper layer protocol



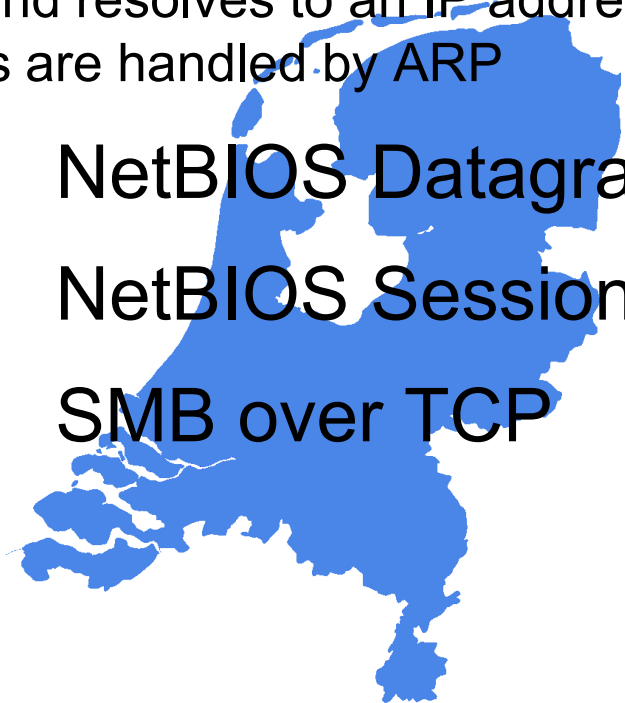
NetBIOS has moved to upper Layers





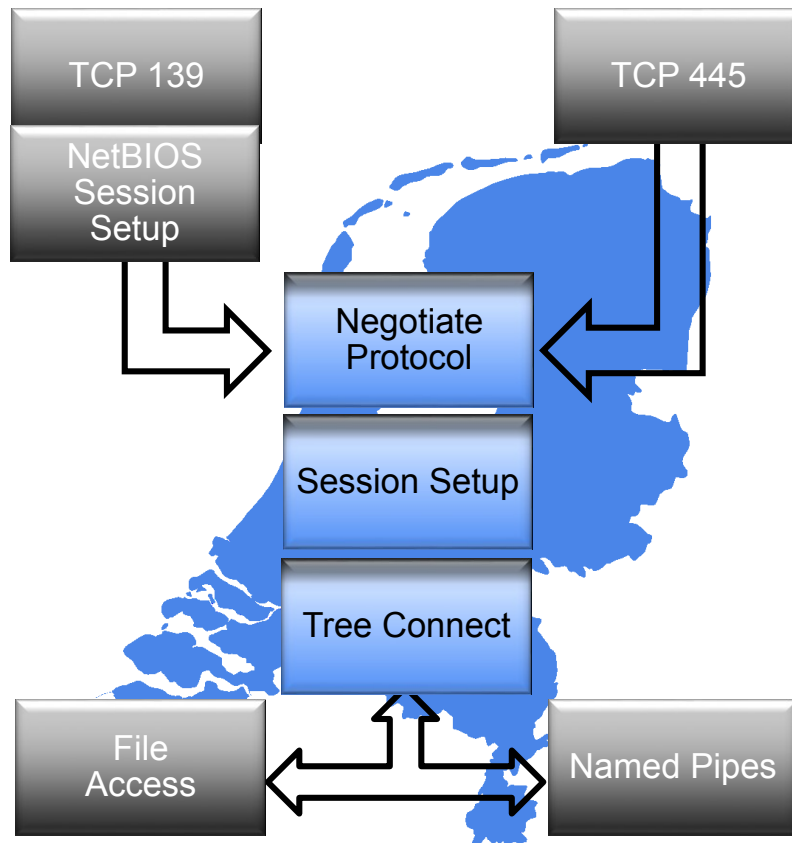
NetBIOS Services in IP networks

- **UDP 137** **NetBIOS Name Service**
 - Now registers and resolves to an IP address
 - MAC addresses are handled by ARP
- **UDP 138** **NetBIOS Datagram Service**
- **TCP 139** **NetBIOS Session Service**
- **TCP 445** **SMB over TCP**





SMB or NetBIOS: Remote File Access





Critical Part: Session Setup

- NetBIOS Session Setup: Only on Port 139

No.	Time	Source	Destination	Protocol	Info
...	5.906	172.16.0.100	172.16.0.124	TCP	1567→139 [SYN] Seq=0 Win=65535 Len=0 MSS=1460 TSval=0 TSecr=0 SACK_PERM=1
...	5.907	172.16.0.124	172.16.0.100	TCP	139→1567 [SYN, ACK] Seq=0 Ack=1 Win=65535 Len=0 MSS=1460 TSval=0 TSecr=0 SACK_PERM=1
...	5.907	172.16.0.100	172.16.0.124	NBSS	Session request, to *SMBSErv<00><00><00><00><00><00><00><00> from BALMUNG<00>
...	5.907	172.16.0.124	172.16.0.100	NBSS	Negative session response, Called name not present
...	5.907	172.16.0.100	172.16.0.124	TCP	1567→139 [ACK] Seq=73 Ack=7 Win=65530 Len=0 TSval=459157 TSecr=275285
...	5.908	172.16.0.100	172.16.0.124	TCP	1567→139 [FIN, ACK] Seq=73 Ack=7 Win=65530 Len=0 TSval=459157 TSecr=275285
...	5.908	172.16.0.124	172.16.0.100	TCP	139→1567 [ACK] Seq=7 Ack=74 Win=65463 Len=0 TSval=275285 TSecr=459157

▸ Frame 17: 138 bytes on wire (1104 bits), 138 bytes captured (1104 bits)

▸ Ethernet II, Src: AsustekC_00:ed:d1 (00:11:d8:00:ed:d1), Dst: Dell_cc:43:e3 (00:13:72:cc:43:e3)

▸ Internet Protocol Version 4, Src: 172.16.0.100, Dst: 172.16.0.124

▸ Transmission Control Protocol, Src Port: 1567, Dst Port: 139, Seq: 1, Ack: 1, Len: 62

▸ NetBIOS Session Service

- Message Type: Session request (0x81)
- Flags: 0x00
- Length: 68
- Called name: *SMBSErv<00><00><00><00><00><00><00><00> (Workstation/Redirector)
- Calling name: BALMUNG<00> (Workstation/Redirector)



Negotiate SMB Dialect

```
▸ Transmission Control Protocol, Src Port: 49675, Dst
▸ NetBIOS Session Service
▸ SMB (Server Message Block Protocol)
  ▸ SMB Header
  ▸ Negotiate Protocol Request (0x72)
    Word Count (WCT): 0
    Byte Count (BCC): 120
    ▸ Requested Dialects
      ▸ Dialect: PC NETWORK PROGRAM 1.0
      ▸ Dialect: LANMAN1.0
      ▸ Dialect: Windows for Workgroups 3.1a
      ▸ Dialect: LM1.2X002
      ▸ Dialect: LANMAN2.1
      ▸ Dialect: NT LM 0.12
      ▸ Dialect: SMB 2.002
      ▸ Dialect: SMB 2.???
```

```
▸ SMB2 (Server Message Block Protocol version 2)
  ▸ SMB2 Header
  ▸ Negotiate Protocol Request (0x00)
    ▸ StructureSize: 0x0024
      Dialect count: 5
    ▸ Security mode: 0x01, Signing enabled
      Reserved: 0000
    ▸ Capabilities: 0x0000007f, DFS, LEASING, LARGE MTU,
      Client Guid: 0eca552d-9378-11e6-aaf4-000c2903dfad
      NegotiateContextOffset: 0x0070
      NegotiateContextCount: 2
      Reserved: 0000
      Dialect: 0x0202
      Dialect: 0x0210
      Dialect: 0x0300
      Dialect: 0x0302
      Dialect: 0x0311
```



SMB Protocol Versions

- **Avoid SMB v1**
 - Designed for 16 bit systems
 - Wfw 3.1, LAN Manager, NT LM 0.12 and other dialects
 - Implementation details can turn a standard task into a night mare
 - Still widely distributed, especially in low cost NAS
- **Prefer SMB v2 or v3**
 - Version 2 was introduced with Windows Vista
 - Version 3 was introduced with Windows 8
 - Wireshark offers [display filters](#) and [smb2](#)



Advantages in SMB v2 and later

- Designed for 10 GBit to the desktop
- Supports larger files
- Pipelining
- Lot's of useful features, like branch cache, encryption and more





Authenticating the User

- Happens with the Session Setup Command
 - Domains need 1 turn, workgroups need 2 turns
 - More turns, if authentication fails
- Client offers multiple authentication methods
 - Usually NTLMv2 or Kerberos
 - NTLMv2 is challenge/response
- Safely ignore the return code "More processing required"

No.	Time	Protocol	Info
858	*REF*	SMB2	Session Setup Request, NTLMSSP_NEGOTIATE
+ 859	0.000	SMB2	Session Setup Response, Error: STATUS_MORE_PROCESSING_REQUIRED, NTLMSSP_CHALLENGE
860	0.000	SMB2	Session Setup Request, NTLMSSP_AUTH, User: DESKTOP-2AEFH76\Willi Wireshark
861	0.001	SMB2	Session Setup Response

SMB2 Header	
Session Setup Request (0x01)	
StructureSize: 0x0019	
Flags: 0	
Security mode: 0x01, Signing enabled	
Capabilities: 0x00000001, DFS	
Channel: None (0x00000000)	
Previous Session Id: 0x0000000000000000	
Security Blob: 604806062b0601050502a03e303ca0e300c060a2b060104...	
Offset: 0x00000058	
Length: 74	
GSS-API Generic Security Service Application Program Interface	
OID: 1.3.6.1.5.5.2 (SPNEGO - Simple Protected Negotiation)	
Simple Protected Negotiation	



Accessing Files





Get Ready to Access Files

- Mapping a share is done with "Tree Connect"
- The explorer will retrieve information about the target directory and stored files
- A "Create" Request is send when the file is opened
- Subsequent "Read" or "Write" Requests transfer data from and to the file



The full process to read a file

Many TCP Turns
hurt WAN users

Source	Destination	Protocol	Info
10.12.0.2	10.5.5.2	SMB	Negotiate Protocol Request
10.5.5.2	10.12.0.2	SMB	Negotiate Protocol Response
10.12.0.2	10.5.5.2	SMB	Session Setup AndX Request, User: \jwurst
10.5.5.2	10.12.0.2	SMB	Session Setup AndX Response
10.12.0.2	10.5.5.2	SMB	Tree Connect AndX Request, Path: \\10.5.5.2\infos
10.5.5.2	10.12.0.2	SMB	Tree Connect AndX Response
10.12.0.2	10.5.5.2	SMB	Trans2 Request, QUERY_FS_INFO, Query FS Device Info
10.5.5.2	10.12.0.2	SMB	Trans2 Response, QUERY_FS_INFO
10.12.0.2	10.5.5.2	SMB	Trans2 Request, QUERY_FS_INFO, Query FS Attribute Info
10.5.5.2	10.12.0.2	SMB	Trans2 Response, QUERY_FS_INFO
10.12.0.2	10.5.5.2	SMB	Trans2 Request, QUERY_PATH_INFO, Query File All Info, Path:
10.5.5.2	10.12.0.2	SMB	Trans2 Response, QUERY_PATH_INFO
10.12.0.2	10.5.5.2	SMB	Trans2 Request, QUERY_PATH_INFO, Query File Internal Info, Path:
10.5.5.2	10.12.0.2	SMB	Trans2 Response, QUERY_PATH_INFO
10.12.0.2	10.5.5.2	SMB	Trans2 Request, FIND_FIRST2, Pattern: *
10.5.5.2	10.12.0.2	SMB	Trans2 Response, FIND_FIRST2, Files: . .. documents papers
10.12.0.2	10.5.5.2	SMB	Trans2 Request, QUERY_PATH_INFO, Query File All Info, Path: \documents\Famous-Secrets.txt
10.5.5.2	10.12.0.2	SMB	Trans2 Response, QUERY_PATH_INFO
10.12.0.2	10.5.5.2	SMB	Trans2 Request, QUERY_PATH_INFO, Query File Internal Info, Path: \documents\Famous-Secrets.txt
10.5.5.2	10.12.0.2	SMB	Trans2 Response, QUERY_PATH_INFO
10.12.0.2	10.5.5.2	SMB	NT Create AndX Request, FID: 0x4001, Path: \documents\Famous-Secrets.txt
10.5.5.2	10.12.0.2	SMB	NT Create AndX Response, FID: 0x4001
10.12.0.2	10.5.5.2	SMB	Read AndX Request, FID: 0x4001, 4096 bytes at offset 0
10.5.5.2	10.12.0.2	SMB	Read AndX Response, FID: 0x4001, 521 bytes
10.12.0.2	10.5.5.2	SMB	Close Request, FID: 0x4001



Multiple IDs make a File unique

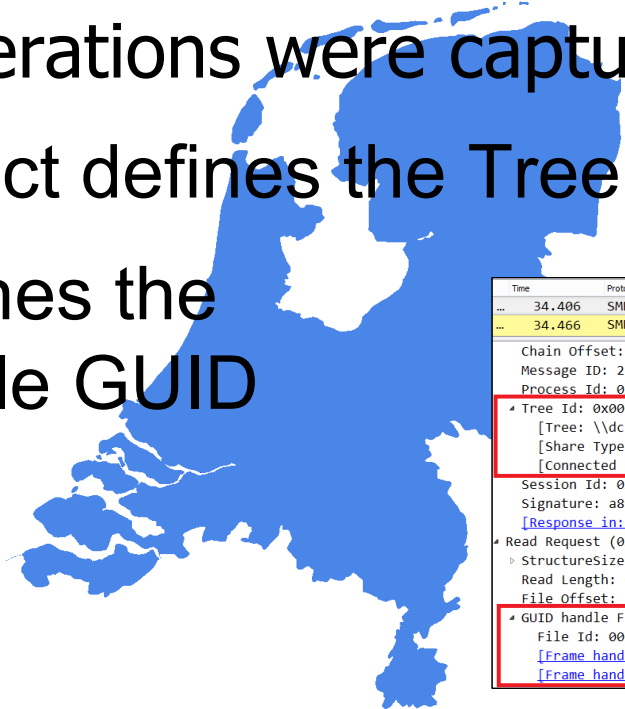
- The server assigns IDs for further reference
 - Process ID references a logon session
 - Tree ID references a share name
 - GUID references a file name
- The client uses only the IDs in requests
- Wireshark associates the IDs with names





Locating IDs in SMB Messages

- Wireshark will show the file names, if the relevant operations were captured
- Tree Connect defines the Tree ID
- Create defines the File ID or File GUID



```
Time      Protocol  Info
...      34.406   SMB2     Read Request Len:65536 Off:131072 File: Files\Pilz.JPG
...      34.466   SMB2     Read Response

Chain Offset: 0x00000000
Message ID: 251
Process Id: 0x000feff
  * Tree Id: 0x00000009 \\dc-hq-2008-1\Data
    [Tree: \\dc-hq-2008-1\Data]
    [Share Type: Physical disk (0x01)]
    [Connected in Frame: 107]
  Session Id: 0x0000400680001d
  Signature: a89b8190d26e2b7b6185748ed116305c
  [Response in: 7585]
  * Read Request (0x08)
    ▸ StructureSize: 0x0031
    Read Length: 65536
    File Offset: 131072
  * GUID handle File: Files\Pilz.JPG
    File Id: 00000059-0100-0000-5100-0000ffffffff
    [Frame handle opened: 7354]
    [Frame handle closed: 14700]
```



Analyzing SMB Traffic





Most useful Wireshark Commands

- Use the display filter **smb** or **smb2**
- Examine the protocol statistics
 - Locate frequent events
 - Locate long transaction times

- Locate smb errors:

smb.error_code > 0 or

smb.nt_status > 0 or

smb2.nt_status > 0

No.	Time	Protocol	Info
1...	54.041	SMB	NT Create AndX Response, FID: 0x0000, Error: STATUS_ACCESS_DENIED
110	54.091	TCP	51393→445 [ACK] Seq=1847 Ack=946 Win=64512 Len=0
111	54.571	ARP	Who has 192.168.199.134? Tell 192.168.199.1

SMB Command: NT Create AndX (0xa2)	
NT Status: STATUS_ACCESS_DENIED (0xc0000022)	
Flags: 0x98, Request/Response, Canonicalized Pathnames, Case Sensitivity	
Flags2: 0xc007, Unicode Strings, Error Code Type, Extended Security Negotiation, Sec	
Process ID High: 0	
Signature: 0000000000000000	
Reserved: 0000	
Tree ID: 2051 (\\192.168.199.134\STUFF)	
Process ID: 568	
User ID: 2050	
Multiplex ID: 192	

0000	00 50 56 c0 00 01 00 0c	29 1f cc 13 08 00 45 00	.PV.....).....E.
0010	00 4f 06 85 40 00 80 06	e4 4a c0 a8 c7 86 c0 a8	.0.@...J.....
0020	c7 01 01 bd c8 c1 62 13	00 d2 d5 8d e8 9f 50 18b.....P.
0030	f9 e6 36 03 00 00 00 00	00 23 ff 53 4d 42 a2 22	..6.....#.SMB.
0040	00 00 c0 98 07 c8 00 00	00 00 00 00 00 00 00 00
0050	00 00 03 08 38 02 02 08	c0 00 00 00 00 00	...8.....



Protocol Statistics

- Separate statistics for SMB and SMB2

Wireshark · SMB Service Response Time Statistics · MP1 (client 4-to-1)

Index	Procedure	Calls	Min SRT (s)	Max SRT (s)	Avg SRT (s)	Sum SRT (s)
SMB Commands						
4	Close	110	0.000001	0.001020	0.000528	0.058068
116	Logoff AndX	10	0.000003	0.000610	0.000337	0.003368
114	Negotiate Protocol	20	0.000003	0.000942	0.000565	0.011298
162	NT Create AndX	112	0.000001	0.001020	0.000538	0.060219
46	Read AndX	1	0.000539	0.000539	0.000539	0.000539
115	Session Setup AndX	51	0.000002	5.003432	0.981250	50.043750
117	Tree Connect AndX	40	0.000001	0.001027	0.000465	0.018599
113	Tree Disconnect	30	0.000001	0.001020	0.000469	0.014064
11	Write	111	0.000000	0.001221	0.000520	0.057667
47	Write AndX	27570	0.000000	0.004214	0.000587	16.173566
Transaction2 Sub-Commands						
1	FIND_FIRST2	667	0.000002	0.003572	0.001074	0.716515
2	FIND_NEXT2	112	0.001018	0.002613	0.001490	0.166862
16	GET_DFS_REFERRAL	75	0.000000	0.001013	0.000432	0.032432
7	QUERY_FILE_INFO	112	0.000000	0.001847	0.000551	0.061688
3	QUERY_FS_INFO	126	0.000001	0.001212	0.000387	0.048718
5	QUERY_PATH_INFO	379	0.000000	0.001968	0.000471	0.178440
NT Transaction Sub-Commands						

Display filter: Enter a display filter ...



Identify time consuming Transactions

- Look out for individual slow transactions
 - `smb.time > 0.2`
 - Your definition of "slow" depends on your infrastructure
- Look out for repeated, maybe useless requests
 - The previous screenshot showed 27.000 write requests
 - At 0.5 milliseconds each, we still spend 16 seconds waiting

```
Info
Write AndX Request, FID: 0x024d, 2 bytes at offset 5253
Write AndX Response, FID: 0x024d, 2 bytes
Write AndX Request, FID: 0x024d, 47 bytes at offset 5255
Write AndX Response, FID: 0x024d, 47 bytes
Write AndX Request, FID: 0x024d, 2 bytes at offset 5302
Write AndX Response, FID: 0x024d, 2 bytes
Write AndX Request, FID: 0x024d, 47 bytes at offset 5304
Write AndX Response, FID: 0x024d, 47 bytes
```



Safely ignore Notify Transactions

- The explorer will ask to be notified when a directory or file changes
- The response time for the notification depends on user and application behavior


Time	Protocol	Info
REF	SMB	NT Trans Request, NT NOTIFY, FID: 0x4000
8.862	SMB	NT Trans Response, FID: 0x4000, NT NOTIFY

Max Parameter Count: 32
Max Data Count: 0
Parameter Count: 0
Parameter Offset: 84
Data Count: 0
Data Offset: 0
Setup Count: 4
Function: NT NOTIFY (4)
NT NOTIFY Setup
▸ Completion Filter: 0x00000017, File Name Change, Directory Name Change, Attribute Change
▸ FID: 0x4000 ()



Avoid loading executables from shares

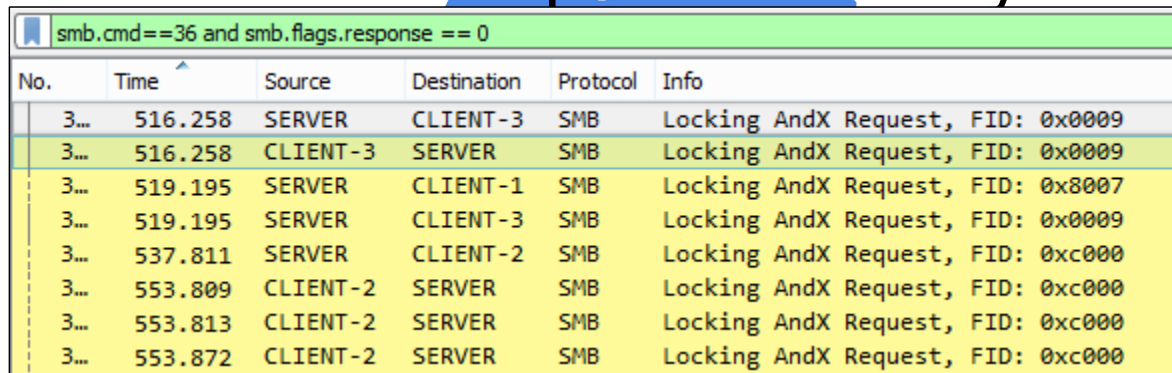
- Sometimes applications are stored on a file share
- The Windows prefetcher can load a program, even if the user did not click it
- Executables are treated like small page files
- Looong wait for RAS links





Watch out for Opportunistic Locking

- Files can be shared for reading and writing
- Clients can reserve a section of the file (lock)
- The server has to manage conflicting locks
- Watch out for lock requests send by the server



The image shows a Wireshark packet capture window with the filter `smb.cmd==36 and smb.flags.response == 0`. The table below represents the data shown in the packet list pane.

No.	Time	Source	Destination	Protocol	Info
3...	516.258	SERVER	CLIENT-3	SMB	Locking AndX Request, FID: 0x0009
3...	516.258	CLIENT-3	SERVER	SMB	Locking AndX Request, FID: 0x0009
3...	519.195	SERVER	CLIENT-1	SMB	Locking AndX Request, FID: 0x8007
3...	519.195	SERVER	CLIENT-3	SMB	Locking AndX Request, FID: 0x0009
3...	537.811	SERVER	CLIENT-2	SMB	Locking AndX Request, FID: 0xc000
3...	553.809	CLIENT-2	SERVER	SMB	Locking AndX Request, FID: 0xc000
3...	553.813	CLIENT-2	SERVER	SMB	Locking AndX Request, FID: 0xc000
3...	553.872	CLIENT-2	SERVER	SMB	Locking AndX Request, FID: 0xc000

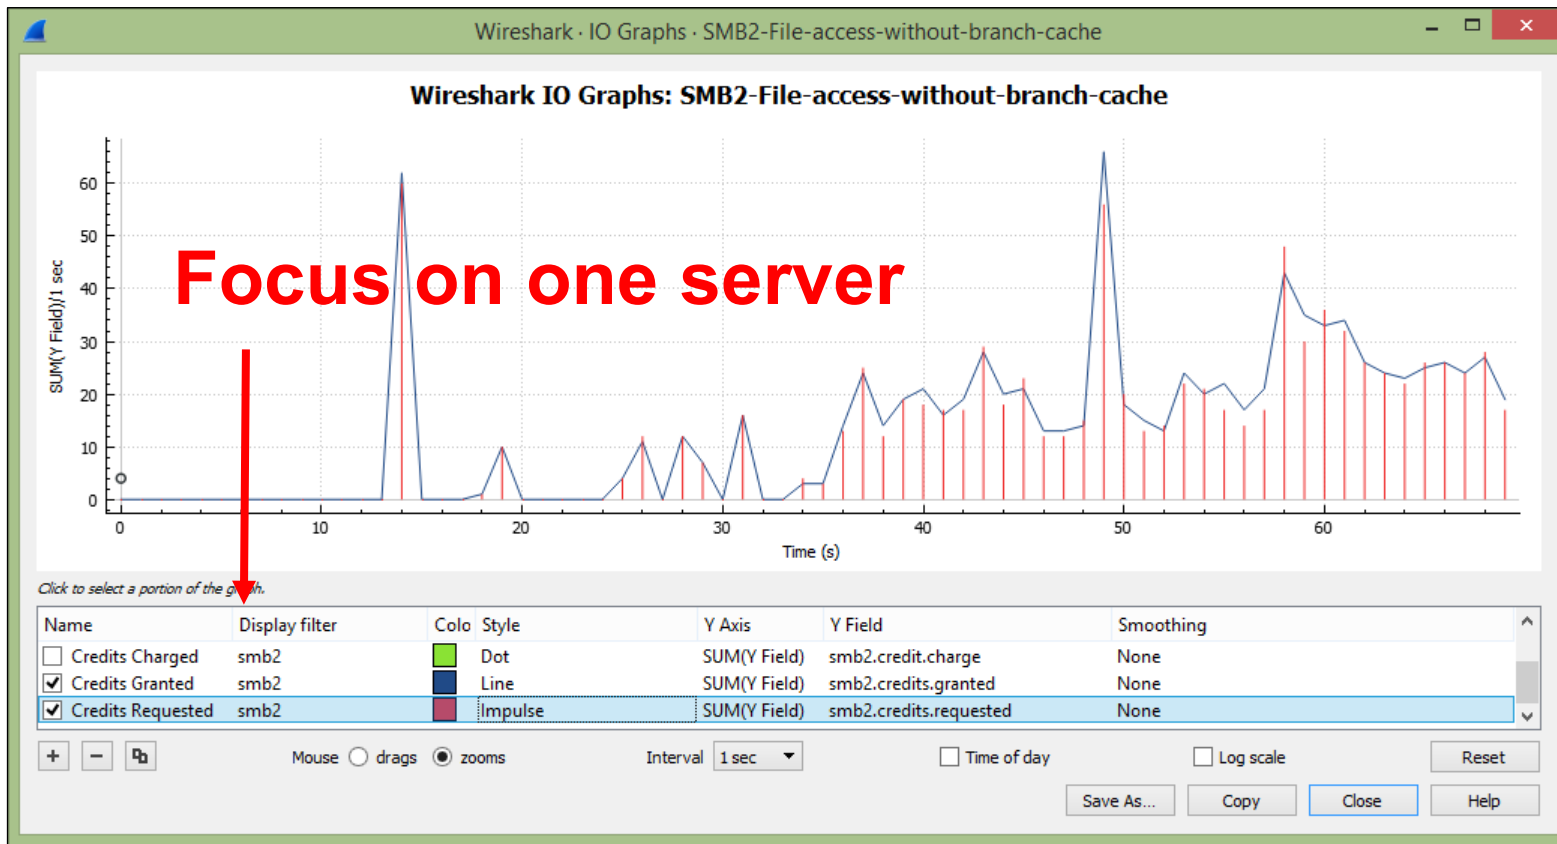


SMB2 Clients spend I/O Credits

- Since SMB2 clients have to spend credits for I/O Operations
 - 1 credit for each full 64 kByte of Data requested plus 1 credit for a fraction of 64 kByte
 - Cost depends on the number of requested bytes, not on the number of bytes read / written
- The server will replenish the credits
 - The client will always have at least 1 credit



Charting SMB2 Credits





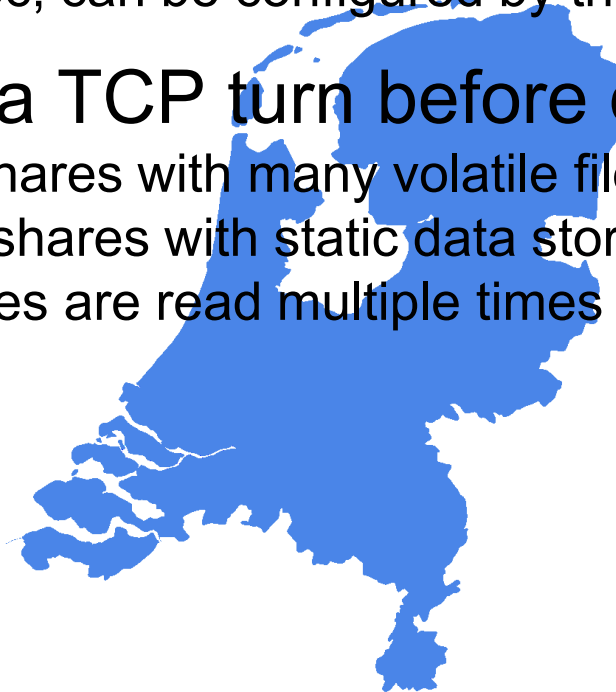
Windows Branch Cache can reduce WAN load

- Data from a network share is either cached on a server or distributed to all clients in the branch
- Data is split into chunks of 64 kByte
- Clients first retrieve a hash of the data and then ask in the LAN if someone has that chunk
- Reasonable secure implementation: cached chunks are encrypted



Considerations for the Branch Cache

- Caching is triggered by latency
 - Default: 80 msec, can be configured by the domain administrator
- Adds an extra TCP turn before data is read
 - Not useful for shares with many volatile files
 - Very useful for shares with static data stored in large files
 - Only useful if files are read multiple times





Do it yourself!

- Join the community on ask.wireshark.org
- Check out this question: <https://goo.gl/RW02jc>



The screenshot shows a web browser window with the URL <https://ask.wireshark.org/questions/55972/slow-writes-even-slower-reads-spanning-wan-to-netapp>. The page features the Wireshark logo, navigation tabs for Questions, Tags, Users, Badges, and Unanswered, and a search bar. The question title is "slow writes, even slower reads spanning WAN to Netapp". The user "packethunter" has asked the question. The content of the question is:

Win7 workstation -> LAN -> ASA -> Cisco ASR -> DMVPN -> ASR -> Palo Alto -> Nexus -> NetApp

0

We are experiencing the symptoms described in the title. This is not new, it predates me, and it happens at multiple spoke sites in our DMVPN. Each vendor just seems to point the finger at the other with no real data reinforcing their point. Cisco has cleared any real issues at the hardware level.



Questions?



Famous Last Words

- Tracefiles used in this presentation have been uploaded to the Wireshark Wiki:

<https://wiki.wireshark.org/SampleCaptures>

- More information on SMB and SMB2 on the Wireshark Wiki:

<https://wiki.wireshark.org/SMB>

<https://wiki.wireshark.org/SMB2>