



#sf21vus

TCP Selective Acknowledgement (SACK)

...correctly interpreting this field can help you
be a better performance troubleshooter...

Services CTO
Global Customer Success
Riverbed Technologies
john.pittle@riverbed.com
[@end2endviz](https://twitter.com/end2endviz)

www.linkedin.com/in/john-pittle



#sf21vus

Abstract

- RFC 2018 introduced an optional ACK mechanism called "Selective Acknowledgement" (SACK)
- Understanding how to interpret SACK can help you quickly determine effect of Out of Sequence packets on overall application performance



#sf21vus

My Agenda

- Call out relevant RFCs
- Review TCP ACK Basics
- SACK Introduction, Decode Details, Examples
- Wireshark columns, filters, and colorization
- Adventures from the Field - Visualization Replay
- Wrap-Up

John Who?



#sf21vus

- Practicing Performance Engineering since 1980
- Protocol Analysis since 1991
- Professional Services with OPNET / Riverbed since 2005
- Love the mystery of a complicated performance issue
- Shaved off beard in 2003...





#sf21vus



#sf21vus • Online • September 12-17

Related RFCs

- RFC 793 – TCP (Original RFC – 1981)
- RFC 2018 – TCP Selective ACK Options (1996)
- RFC 2883 – D-SACK - An Extension to SACK ... (2000)

Review: TCP ACK Behavior



#sf21vus

- As long as packets arrive in the expected order, receiver will ACK every other packet (Default Behavior)
- If a packet arrives out of order, the receiver will immediately issue an ACK with a value equal to the SEQ that was expected

Review: TCP ACK Behavior



- Receiver will continue to ACK every packet until the expected packet is received
- If sender receives 4 ACKs with the same ACK number (aka Triple Duplicate ACK) he will retransmit the missing segment
 - Assumes TCP Fast Retransmit & Recovery (FRR) is available and enabled

! Heads Up !



#sf21vus

Important term on the next slide

Cumulative ACK (RFC 793)



No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	172.20.1.1	10.200.50.1	TCP	78	43650 → 8085 [SYN] Seq=0 Win=32850 Len=0 MSS=1436 TSval=525253167 TSecr=0 WS=8 SACK_PERM=1
2	0.000016	10.200.50.1	172.20.1.1	TCP	74	8085 → 43650 [SYN, ACK] Seq=0 Ack=1 Win=28960 Len=0 MSS=1460 SACK_PERM=1 TSval=1376522072 TSecr=525253167 WS=4
3	0.007680	172.20.1.1	10.200.50.1	TCP	66	43650 → 8085 [ACK] Seq=1 Ack=1 Win=263440 Len=0 TSval=525253168 TSecr=1376522072
4	0.008749	172.20.1.1	10.200.50.1	TCP	145	43650 → 8085 [PSH, ACK] Seq=1 Ack=1 Win=263440 Len=79 TSval=525253168 TSecr=1376522072
5	0.008758	10.200.50.1	172.20.1.1	TCP	66	8085 → 43650 [ACK] Seq=1 Ack=80 Win=28960 Len=0 TSval=1376522073 TSecr=525253168
6	0.030798	10.200.50.1	172.20.1.1	TCP	105	8085 → 43650 [PSH, ACK] Seq=1 Ack=80 Win=28960 Len=39 TSval=1376522075 TSecr=525253168
7	0.038758	172.20.1.1	10.200.50.1	TCP	66	43650 → 8085 [ACK] Seq=80 Ack=40 Win=263440 Len=0 TSval=525253171 TSecr=1376522075
8	0.043602	172.20.1.1	10.200.50.1	TCP	290	43650 → 8085 [PSH, ACK] Seq=80 Ack=40 Win=263440 Len=224 TSval=525253172 TSecr=1376522075
9	0.056458	10.200.50.1	172.20.1.1	TCP	2914	8085 → 43650 [ACK] Seq=40 Ack=304 Win=28960 Len=2848 TSval=1376522077 TSecr=525253172
10	0.056607	10.200.50.1	172.20.1.1	TCP	70	8085 → 43650 [PSH, ACK] Seq=2888 Ack=304 Win=28960 Len=4 TSval=1376522077 TSecr=525253172
11	0.064400	172.20.1.1	10.200.50.1	TCP	66	43650 → 8085 [ACK] Seq=304 Ack=40 Win=263440 Len=0 TSval=525253174 TSecr=1376522077

▶ Frame 7: 66 bytes on wire (528 bits), 66 bytes captured (528 bits)
▶ Ethernet II, Src: 40:01:d7:63:7c:50 (40:01:d7:63:7c:50), Dst: IntelCor_bb:d6:1c (00:1e:67:bb:d6:1c)
▶ Internet Protocol Version 4, Src: 172.20.1.1, Dst: 10.200.50.1
▶ Transmission Control Protocol, Src Port: 43650, Dst Port: 8085, Seq: 80, Ack: 40, Len: 0

The ACK in the TCP header is called the “**Cumulative ACK**”. The value reflects stream bytes received in order up to the point when the ACK packet was transmitted.

Receiver’s TCP declares that all bytes in the stream up to ACK-1 have been received. The next byte of TCP stream expected by the receiver should start with a SEQ equal to this ACK.

Selective Acknowledgement



#sf21vus

- RFC 2018 proposed an enhancement to the TCP ACK mechanism
- Selectively acknowledge segments that have arrived out of order
 - The sender won't have to retransmit those segments if he knows they've been received
 - But, this can't be accomplished with Cumulative ACK field alone, so a new field is needed

Selective ACK – A TCP Enhancement



#sf21vus

- New addition to the TCP Options field of the TCP header
- Up to four (4) contiguous out of order segments/segment ranges can be defined using SACK
 - Only three (3) if the TCP Timestamp option is also being used



#sf21vus

Enabling SACK

- On by default in modern TCP Stacks
- SACK is negotiated at connection start-up
- Decode the TCP Options in SYN and SYN+ACK and you'll see "SACK Permitted"
 - Meaning ... "I will process the SACK field if you send it to me"
- Each side can independently chose



#sf21vus

Intended Benefits


- Inform the sender with better info about the success of TCP segment delivery
- Goal is to minimize the amount of unnecessary retransmissions
- Will not necessarily change Congestion Control algorithms
- Any retransmission may still have a negative effect on the Congestion Window and related timers



Use during packet analysis



#sf21vus

- Interrogating the SACK fields will help you quantify the extent of Out of Sequence packets
- Use “Bytes in Flight” as a guiding metric
-  If in-flight data stays high no need to look any further
- If in-flight data constantly dips or hits zero; or you frequently see TCP slow-start, you may find the root cause is severe out of sequence packets

! Heads Up !



#sf21vus

Three new terms on the next slide

Wireshark is SACK Aware




- Wireshark can decode the SACK fields in the TCP Options section of the TCP layer
- “SACK Count” and “Left Edge / Right Edge” values can be displayed as columns in the decode summary section

Time for an Illustration...



#sf21vus



#sf21vus • Online • September 12-17

SACK Illustration #1



- Sender transmits a burst of 5 packets as follows:
 - Pkt 1 SEQ=11 Len=10
 - Pkt 2 SEQ=21 Len=10
 - Pkt 3 SEQ=31 Len=10
 - Pkt 4 SEQ=41 Len=10
 - Pkt 5 SEQ=51 Len=10

SACK Illustration #1




- Due to a network issue, the packets are received in the following order:
 - Pkt 1
 - Pkt 2
 - Pkt 4 ←
 - Pkt 5
 - Pkt 3

SACK Visualization #1




1st two packets arrive



SACK Visualization #1



1st two packets arrive



Left Edge = 41

Packet 4 arrives




Right Edge = 51



#sf21vus

SACK Visualization #1

1st two packets arrive



Left Edge = 41

Packet 4 arrives



Right Edge = 51

Packet 5 arrives



Right Edge = 61

Left Edge = 41

SACK Visualization #1




#sf21vus


1st two packets arrive



Packet 4 arrives



Packet 5 arrives



Packet 3 arrives



All data received up through and including byte 60, receiver is ready for TCP stream byte 61

Time for some decodes...



#sf21vus

SACK Wireshark Columns



#sf21vus

No.	Time	Source	Destination	Length	TCP SACK Count	TCP SACK Left Edge	TCP SACK Right Edge	Bytes in flight	Info
31	0.711313	Server	Client	1496				5704	8085 → 43650 [PSH, ACK] Seq=13341 Ack=2171 Win=32044 Len=1430 TSval=1376522
32	0.711325	Server	Client	2914				8552	8085 → 43650 [ACK] Seq=14771 Ack=2171 Win=32044 Len=2848 TSval=1376522143 T
33	0.711339	Server	Client	2914				11400	8085 → 43650 [ACK] Seq=17619 Ack=2171 Win=32044 Len=2848 TSval=1376522143 T
34	0.718880	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=10493 Win=263440 Len=0 TSval=525253239 TSec
35	0.718888	Server	Client	74				9982	8085 → 43650 [PSH, ACK] Seq=20467 Ack=2171 Win=32044 Len=8 TSval=1376522144
36	0.719162	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=13341 Win=260592 Len=0 TSval=525253239 TSec
37	0.719441	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 TSval=525253239 TSec
38	0.719446	Client	Server	78	1	17619	19043		[TCP Dup ACK 37#1] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 T
39	0.721344	Server	Client	1492				7130	8085 → 43650 [PSH, ACK] Seq=20475 Ack=2171 Win=32044 Len=1426 TSval=1376522
40	0.721452	Server	Client	1492				8556	8085 → 43650 [PSH, ACK] Seq=21901 Ack=2171 Win=32044 Len=1426 TSval=1376522
41	0.722115	Server	Client	1492				9982	8085 → 43650 [PSH, ACK] Seq=23327 Ack=2171 Win=32044 Len=1426 TSval=1376522
42	0.726388	Client	Server	86	2	20467,17619	20475,19043		[TCP Dup ACK 37#2] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 T
43	0.726408	Server	Client	1490				9982	[TCP Fast Retransmission] 8085 → 43650 [ACK] Seq=14771 Ack=2171 Win=32044 L

Urgent pointer: 0

Options: (24 bytes), No-Operation (NOP), No-Operation (NOP), Timestamps, No-Operation (NOP), No-Operation (NOP), SACK

▷ No-Operation (NOP)

▷ No-Operation (NOP)

▷ Timestamps: TSval 525253239, TSecr 1376522143

▷ No-Operation (NOP)

▷ No-Operation (NOP)

▷ SACK: 17619-19043

Kind: SACK (5)

Length: 10

left edge = 17619 (relative)

right edge = 19043 (relative)

[TCP SACK Count: 1]

All bytes through 14770r have been received and I'm ready for 14771r. But wait there's more....

SACK Wireshark Columns



#sf21vus

No.	Time	Source	Destination	Length	TCP SACK Count	TCP SACK Left Edge	TCP SACK Right Edge	Bytes in flight	Info
31	0.711313	Server	Client	1496				5704	8085 → 43650 [PSH, ACK] Seq=13341 Ack=2171 Win=32044 Len=1430 TSval=1376522
32	0.711325	Server	Client	2914				8552	8085 → 43650 [ACK] Seq=14771 Ack=2171 Win=32044 Len=2848 TSval=1376522143 T
33	0.711339	Server	Client	2914				11400	8085 → 43650 [ACK] Seq=17619 Ack=2171 Win=32044 Len=2848 TSval=1376522143 T
34	0.718880	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=10493 Win=263440 Len=0 TSval=525253239 TSec
35	0.718888	Server	Client	74				9982	8085 → 43650 [PSH, ACK] Seq=20467 Ack=2171 Win=32044 Len=8 TSval=1376522144
36	0.719162	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=13341 Win=260592 Len=0 TSval=525253239 TSec
37	0.719441	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 TSval=525253239 TSec
38	0.719446	Client	Server	78	1	17619	19043		[TCP Dup ACK 37#1] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 T
39	0.721344	Server	Client	1492				7130	8085 → 43650 [PSH, ACK] Seq=20475 Ack=2171 Win=32044 Len=1426 TSval=1376522
40	0.721452	Server	Client	1492				8556	8085 → 43650 [PSH, ACK] Seq=21901 Ack=2171 Win=32044 Len=1426 TSval=1376522
41	0.722115	Server	Client	1492				9982	8085 → 43650 [PSH, ACK] Seq=23327 Ack=2171 Win=32044 Len=1426 TSval=1376522
42	0.726388	Client	Server	86	2	20467,17619	20475,19043		[TCP Dup ACK 37#2] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 T
43	0.726408	Server	Client	1490				9982	[TCP Fast Retransmission] 8085 → 43650 [ACK] Seq=14771 Ack=2171 Win=32044 L

Urgent pointer: 0

Options: (24 bytes), No-Operation (NOP), No-Operation (NOP), Timestamps, No-Operation (NOP), No-Operation (NOP), SACK

▷ No-Operation (NOP)

▷ No-Operation (NOP)

▷ Timestamps: TSval 525253239, TSecr 1376522143

▷ No-Operation (NOP)

▷ No-Operation (NOP)

◀ SACK: 17619-19043

Kind: SACK (5)

Length: 10

left edge = 17619 (relative)

right edge = 19043 (relative)

[TCP SACK Count: 1]

All bytes through 14770r have been received and I'm ready for 14771r. But wait there's more....

I've also received one or more segment(s) out of order: 17619-19042r.

#sf21vus

SACK Wireshark Columns



#sf21vus

No.	Time	Source	Destination	Length	TCP SACK Count	TCP SACK Left Edge	TCP SACK Right Edge	Bytes in flight	Info
31	0.711313	Server	Client	1496				5704	8085 → 43650 [PSH, ACK] Seq=13341 Ack=2171 Win=32044 Len=1430 TSval=1376522
32	0.711325	Server	Client	2914				8552	8085 → 43650 [ACK] Seq=14771 Ack=2171 Win=32044 Len=2848 TSval=1376522143 T
33	0.711339	Server	Client	2914				11400	8085 → 43650 [ACK] Seq=17619 Ack=2171 Win=32044 Len=2848 TSval=1376522143 T
34	0.718880	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=10493 Win=263440 Len=0 TSval=525253239 TSec
35	0.718888	Server	Client	74				9982	8085 → 43650 [PSH, ACK] Seq=20467 Ack=2171 Win=32044 Len=8 TSval=1376522144
36	0.719162	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=13341 Win=260592 Len=0 TSval=525253239 TSec
37	0.719441	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 TSval=525253239 TSec
38	0.719446	Client	Server	78	1	17619	19043		[TCP Dup ACK 37#1] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 T
39	0.721344	Server	Client	1492				7130	8085 → 43650 [PSH, ACK] Seq=20475 Ack=2171 Win=32044 Len=1426 TSval=1376522
40	0.721452	Server	Client	1492				8556	8085 → 43650 [PSH, ACK] Seq=21901 Ack=2171 Win=32044 Len=1426 TSval=1376522
41	0.722115	Server	Client	1492				9982	8085 → 43650 [PSH, ACK] Seq=23327 Ack=2171 Win=32044 Len=1426 TSval=1376522
42	0.726388	Client	Server	86	2	20467,17619	20475,19043		[TCP Dup ACK 37#2] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 T
43	0.726408	Server	Client	1490				9982	[TCP Fast Retransmission] 8085 → 43650 [ACK] Seq=14771 Ack=2171 Win=32044 L

Urgent pointer: 0

Options: (24 bytes), No-Operation (NOP), No-Operation (NOP), Timestamps, No-Operation (NOP), No-Operation (NOP), SACK

▷ No-Operation (NOP)

▷ No-Operation (NOP)

▷ Timestamps: TSval 525253239, TSecr 1376522143

▷ No-Operation (NOP)

▷ No-Operation (NOP)

▷ SACK: 17619-19043

Kind: SACK (5)

Length: 10

left edge = 17619 (relative)

right edge = 19043 (relative)

[TCP SACK Count: 1]

All bytes through 14770r have been received and I'm ready for 14771r. But wait there's more....

I've also received one or more segment(s) out of order: 17619-19042r. What's missing?

#sf21vus

SACK Wireshark Columns



#sf21vus

No.	Time	Source	Destination	Length	TCP SACK Count	TCP SACK Left Edge	TCP SACK Right Edge	Bytes in flight	Info
31	0.711313	Server	Client	1496				5704	8085 → 43650 [PSH, ACK] Seq=13341 Ack=2171 Win=32044 Len=1430 TSval=1376522
32	0.711325	Server	Client	2914				8552	8085 → 43650 [ACK] Seq=14771 Ack=2171 Win=32044 Len=2848 TSval=1376522143 T
33	0.711339	Server	Client	2914				11400	8085 → 43650 [ACK] Seq=17619 Ack=2171 Win=32044 Len=2848 TSval=1376522143 T
34	0.718880	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=10493 Win=263440 Len=0 TSval=525253239 TSec
35	0.718888	Server	Client	74				9982	8085 → 43650 [PSH, ACK] Seq=20467 Ack=2171 Win=32044 Len=8 TSval=1376522144
36	0.719162	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=13341 Win=260592 Len=0 TSval=525253239 TSec
37	0.719441	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 TSval=525253239 TSec
38	0.719446	Client	Server	78	1	17619	19043		[TCP Dup ACK 37#1] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 T
39	0.721344	Server	Client	1492				7130	8085 → 43650 [PSH, ACK] Seq=20475 Ack=2171 Win=32044 Len=1426 TSval=1376522
40	0.721452	Server	Client	1492				8556	8085 → 43650 [PSH, ACK] Seq=21901 Ack=2171 Win=32044 Len=1426 TSval=1376522
41	0.722115	Server	Client	1492				9982	8085 → 43650 [PSH, ACK] Seq=23327 Ack=2171 Win=32044 Len=1426 TSval=1376522
42	0.726388	Client	Server	86	2	20467,17619	20475,19043		[TCP Dup ACK 37#2] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 T
43	0.726408	Server	Client	1490				9982	[TCP Fast Retransmission] 8085 → 43650 [ACK] Seq=14771 Ack=2171 Win=32044 L

Urgent pointer: 0

Options: (24 bytes), No-Operation (NOP), No-Operation (NOP), Timestamps, No-Operation (NOP), No-Operation (NOP), SACK

▷ No-Operation (NOP)

▷ No-Operation (NOP)

▷ Timestamps: TSval 525253239, TSecr 1376522143

▷ No-Operation (NOP)

▷ No-Operation (NOP)

◀ SACK: 17619-19043

Kind: SACK (5)

Length: 10

left edge = 17619 (relative)

right edge = 19043 (relative)

[TCP SACK Count: 1]

All bytes through 14770r have been received and I'm ready for 14771r. But wait there's more....

I've also received one or more segment(s) out of order: 17619-19042r. What's missing? 14771-17618r

#sf21vus



#sf21vus

Notice ACK Packet Len

No.	Time	Source	Destination	Length	TCP SACK Count	TCP SACK Left Edge	TCP SACK Right Edge	Bytes in flight	Info
31	0.711313	Server	Client	1496				5704	8085 → 43650 [PSH, ACK] Seq=13341 Ack=2171 Win=32044 Len=1430 TSval=1376522
32	0.711325	Server	Client	2914				8552	8085 → 43650 [ACK] Seq=14771 Ack=2171 Win=32044 Len=2848 TSval=1376522143 T
33	0.711339	Server	Client	2914				11400	8085 → 43650 [ACK] Seq=17619 Ack=2171 Win=32044 Len=2848 TSval=1376522143 T
34	0.718880	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=10493 Win=263440 Len=0 TSval=525253239 TSec
35	0.718888	Server	Client	74				9982	8085 → 43650 [PSH, ACK] Seq=20467 Ack=2171 Win=32044 Len=8 TSval=1376522144
36	0.719162	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=13341 Win=260592 Len=0 TSval=525253239 TSec
37	0.719441	Client	Server	66					43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 TSval=525253239 TSec
38	0.719446	Client	Server	78	1	17619	19043		[TCP Dup ACK 37#1] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 T
39	0.721344	Server	Client	1492				7130	8085 → 43650 [PSH, ACK] Seq=20475 Ack=2171 Win=32044 Len=1426 TSval=1376522
40	0.721452	Server	Client	1492				8556	8085 → 43650 [PSH, ACK] Seq=21901 Ack=2171 Win=32044 Len=1426 TSval=1376522
41	0.722115	Server	Client	1492				9982	8085 → 43650 [PSH, ACK] Seq=23327 Ack=2171 Win=32044 Len=1426 TSval=1376522
42	0.726388	Client	Server	86	2	20467,17619	20475,19043		[TCP Dup ACK 37#2] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 T
43	0.726408	Server	Client	1490				9982	[TCP Fast Retransmission] 8085 → 43650 [ACK] Seq=14771 Ack=2171 Win=32044 L

!! Heads Up !!



Prepare to observe awesome
Wireshark feature

How did we get those columns?

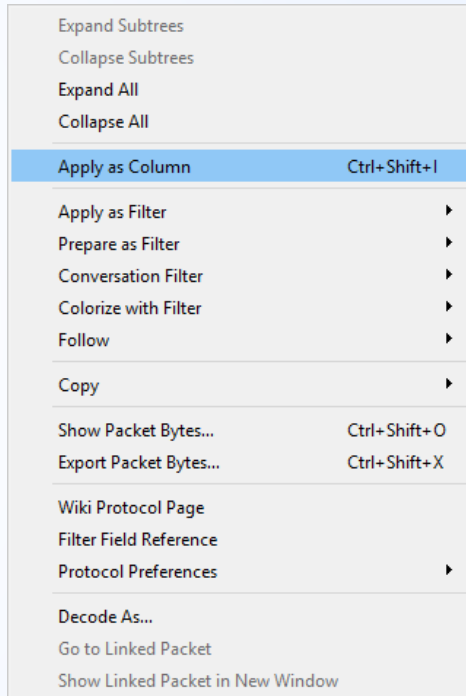


#sf21vus

Navigate to the TCP Options decodes
Highlight the field

```
▼ TCP Option - SACK 31873-36145
  Kind: SACK (5)
  Length: 10
  left edge = 31873 (relative)
  right edge = 36145 (relative)
  [TCP SACK Count: 1]
▼ [SEQ/ACK analysis]
  [rTT: 0.007680000 seconds]
```

Right
mouse
click



Receiver Side TCP Mechanics



#sf21vus



Receiver's ACK responses



#sf21vus

Pkt 1 SEQ=11 Len=10

- Pkt 1 Arrives, receiver starts delayed ACK timer, waits for a 2nd packet

Receiver's ACK responses



Pkt 2 SEQ=21 Len=10

- Pkt 2 Arrives, receiver cancels delayed ACK timer, sends:
 - ACK=31

Receiver's ACK responses



#sf21vus

Pkt 4 SEQ=41 Len=10

- Pkt 4 Arrives..., but wait!, it's out of order...
- receiver issues immediate ACK because...

Receiver's ACK responses



#sf21vus

Pkt 4 SEQ=41 Len=10

- ...the packet we just received is out of order
- Here's the ACK and SACK info...
 - ACK=31 SACK=41-51

Receiver's ACK responses



#sf21vus

Pkt 5 SEQ=51 Len=10

- Pkt 5 Arrives but it's also out of order, receiver issues another immediate ACK because packet is not the expected segment (i.e. it's OoO)
 - ACK=31 SACK=41-61
- *** Note: at this point receiver TCP stack is holding up to 2 packets in the receive buffer and can not hand off to the App ***

Receiver's ACK responses



- Pkt 3 Arrives, receiver issues:
 - ACK=61

Receiver's ACK responses

- Pkt 3 Arrives, receiver issues:
 - ACK=61
- Order is restored to the Force...



Another Example, Slightly More Complicated



#sf21vus


- Sender transmits a burst of 6 packets as follows:
 - Pkt 1 SEQ=11 Len=10
 - Pkt 2 SEQ=21 Len=10
 - Pkt 3 SEQ=31 Len=10
 - Pkt 4 SEQ=41 Len=10
 - Pkt 5 SEQ=51 Len=10
 - Pkt 6 SEQ=61 Len=10

Another Example, Slightly More Complicated



#sf21vus

- Due to a network problem, the packets are received in the following order:
 - Pkt 1
 - Pkt 3
 - Pkt 6
 - Pkt 4
 - Pkt 5
 - Pkt 2



SACK Visualization #2



#sf21vus

Packet 1 arrives

11 - 20




Start Delayed ACK Timer...

SACK Visualization #2



#sf21vus

Packet 1 arrives



Start Delayed ACK Timer...no ACK

Packet 3 arrives



ACK=21 SACK=31-41

SACK Visualization #2



#sf21vus

Packet 1 arrives

11 - 20



Start Delayed ACK Timer...no ACK

Packet 3 arrives

11 - 20

31-40



ACK=21 SACK=31-41

Packet 6 arrives

11 - 20

31-40

61-70




ACK=21 SACK=31-41, 61-71

SACK Visualization #2



#sf21vus



SACK Visualization #2



#sf21vus



Pop Quiz



Packet 5 arrives




ACK=21 SACK 31-71

How many packets
are in sender's
retransmit queue?




#sf21vus

Pop Quiz



How many packets
are in receiver's
queue?

SACK Visualization #2



We Made It!



#sf21vus



#sf21vus • Online • September 12-17

...more decodes...



#sf21vus

Example with two SACK blocks



#sf21vus

No.	Time	Source	Destination	Length	TCP SACK Count	TCP SACK Left Edge	TCP SACK Right Edge	Bytes in flight	Info
40	0.721452	Server	Client	1492				8556	8085 → 43650 [PSH, ACK] Seq=21901 Ack=2171 Win=32044 Len=1426 TSval=137652214
41	0.722115	Server	Client	1492				9982	8085 → 43650 [PSH, ACK] Seq=23327 Ack=2171 Win=32044 Len=1426 TSval=137652214
42	0.726388	Client	Server	86	2	20467,17619	20475,19043		[TCP Dup ACK 37#2] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 TSv
43	0.726408	Server	Client	1490				9982	[TCP Fast Retransmission] 8085 → 43650 [ACK] Seq=14771 Ack=2171 Win=32044 Len=
44	0.728899	Client	Server	86	2	20467,17619	21899,19043		[TCP Dup ACK 37#3] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 TSv
45	0.728910	Client	Server	86	2	20467,17619	21901,19043		[TCP Dup ACK 37#4] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 TSv
46	0.728915	Server	Client	1490				9982	[TCP Out-Of-Order] 8085 → 43650 [ACK] Seq=16195 Ack=2171 Win=32044 Len=1424
47	0.728920	Server	Client	1490				9982	[TCP Out-Of-Order] 8085 → 43650 [ACK] Seq=19043 Ack=2171 Win=32044 Len=1424
48	0.729045	Client	Server	86	2	20467,17619	23325,19043		[TCP Dup ACK 37#5] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 TSv
49	0.729054	Server	Client	1490				11406	8085 → 43650 [ACK] Seq=24753 Ack=2171 Win=32044 Len=1424 TSval=1376522145 TS
50	0.729058	Client	Server	86	2	20467,17619	23327,19043		[TCP Dup ACK 37#6] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 TSv
51	0.729063	Server	Client	1490				12830	8085 → 43650 [ACK] Seq=26177 Ack=2171 Win=32044 Len=1424 TSval=1376522145 TS
52	0.729486	Client	Server	86	2	20467,17619	24751,19043		[TCP Dup ACK 37#7] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 TSv

▷ Timestamps: TSval 525253240, TSecr 1376522143

▷ No-Operation (NOP)

▷ No-Operation (NOP)

▲ SACK: 20467-20475 17619-19043

Kind: SACK (5)

Length: 18

left edge = 20467 (relative)

right edge = 20475 (relative)

left edge = 17619 (relative)

right edge = 19043 (relative)

[TCP SACK Count: 2]



#sf21vus

Example with two blocks OOS

```
▲ SACK: 20467-20475 17619-19043  
  Kind: SACK (5)  
  Length: 18  
  left edge = 20467 (relative)  
  right edge = 20475 (relative)  
  left edge = 17619 (relative)  
  right edge = 19043 (relative)  
  [TCP SACK Count: 2]
```

A Different Wireshark Profile



#sf21vus

SACK-Sample-Conn43650 - Copy_anon.pcapng

File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help

Apply a display filter ... <Ctrl-/>

No.	Time	Source	Destination	Protocol	Length	SACK Count	Cumm ACK	SACK LE 1	SACK RE 1	SACK LE 2	SACK RE 2	Info
37	0.719441	172.20.1.1	10.200.50.1	TCP	66		14771					43650 → 8085 [A
38	0.719446	172.20.1.1	10.200.50.1	TCP	78	1	14771	17619	19043			[TCP Dup ACK 37
39	0.721344	10.200.50.1	172.20.1.1	TCP	1492		2171					8085 → 43650 [P
40	0.721452	10.200.50.1	172.20.1.1	TCP	1492		2171					8085 → 43650 [P
41	0.722115	10.200.50.1	172.20.1.1	TCP	1492		2171					8085 → 43650 [P
42	0.726388	172.20.1.1	10.200.50.1	TCP	86	2	14771	20467	20475	17619	19043	[TCP Dup ACK 37
43	0.726408	10.200.50.1	172.20.1.1	TCP	1490		2171					[TCP Fast Retra
44	0.728899	172.20.1.1	10.200.50.1	TCP	86	2	14771	20467	21899	17619	19043	[TCP Dup ACK 37
45	0.728910	172.20.1.1	10.200.50.1	TCP	86	2	14771	20467	21901	17619	19043	[TCP Dup ACK 37
46	0.728915	10.200.50.1	172.20.1.1	TCP	1490		2171					[TCP Out-Of-Ord
47	0.728920	10.200.50.1	172.20.1.1	TCP	1490		2171					[TCP Out-Of-Ord

! Heads Up !



Another cool Wireshark feature up next...



#sf21vus

Edit "Occurrence #"

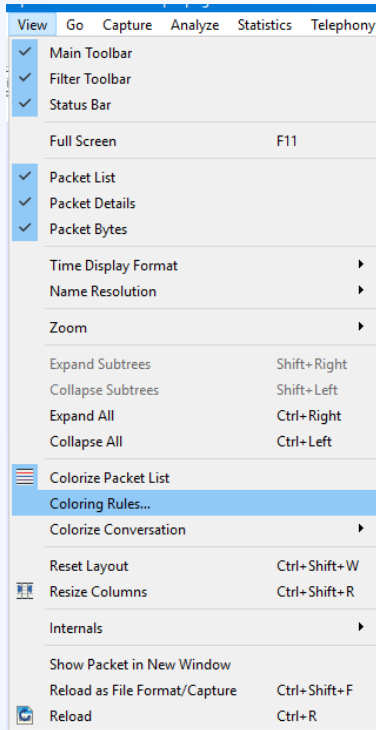
- Right mouse click on column header
- Select "Edit Column"

Apply a display filter ... <Ctrl-/>										
Title: SACK LE 2		Type: Custom		Fields: tcp.options.sack_le		Occurrence: 2				
No.	Time	Source	Destination	Protocol	Length	SACK Count	Cumm ACK	SACK LE 1	SACK RE 1	SA
37	0.719441	172.20.1.1	10.200.50.1	TCP	66		14771			
38	0.719446	172.20.1.1	10.200.50.1	TCP	78	1	14771	17619	19043	
39	0.721344	10.200.50.1	172.20.1.1	TCP	1492		2171			
40	0.721452	10.200.50.1	172.20.1.1	TCP	1492		2171			
41	0.722115	10.200.50.1	172.20.1.1	TCP	1492		2171			
42	0.726388	172.20.1.1	10.200.50.1	TCP	86	2	14771	20467	20475	

Coloring and Filtering Views



#sf21vus



Wireshark · Coloring Rules SACK	
Name	Filter
<input checked="" type="checkbox"/> SACK 2	<code>tcp.options.sack.count==3</code>
<input checked="" type="checkbox"/> SACK 2	<code>tcp.options.sack.count==2</code>
<input checked="" type="checkbox"/> SACK 1	<code>tcp.options.sack.count==1</code>

View with Filter & Color



#sf21vus

SACK-Sample-Conn43650 - anon.pcapng

File Edit View Go Capture Analyze Statistic

tcp.options.sack

No.	Time	Delta	Source	Destination	Length	Seq	NextSeq	ACK	LE1	RE1	LE2	RE2	LE3	RE3	LE4	RE4	Bytes in flight	R_Win
38	0.719446	0.000000000	172.20.1.1	10.200.50.1	78	2171	2171	14771	17619	19043							263440	
42	0.726388	0.006942000	172.20.1.1	10.200.50.1	86	2171	2171	14771	20467	20475	17619	19043					263440	
44	0.728899	0.002511000	172.20.1.1	10.200.50.1	86	2171	2171	14771	20467	21899	17619	19043					263440	
45	0.728910	0.000011000	172.20.1.1	10.200.50.1	86	2171	2171	14771	20467	21901	17619	19043					263440	
48	0.729045	0.000135000	172.20.1.1	10.200.50.1	86	2171	2171	14771	20467	23325	17619	19043					263440	
50	0.729058	0.000013000	172.20.1.1	10.200.50.1	86	2171	2171	14771	20467	23327	17619	19043					263440	
52	0.729486	0.000428000	172.20.1.1	10.200.50.1	86	2171	2171	14771	20467	24751	17619	19043					263440	
54	0.729496	0.000010000	172.20.1.1	10.200.50.1	86	2171	2171	14771	20467	24753	17619	19043					263440	
56	0.734330	0.004834000	172.20.1.1	10.200.50.1	86	2171	2171	16195	20467	24753	17619	19043					263440	
58	0.736839	0.002509000	172.20.1.1	10.200.50.1	78	2171	2171	19043	20467	24753							263440	
66	0.744708	0.007869000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	33297							257744	
68	0.744874	0.000166000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	34721							257744	
70	0.744948	0.000074000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	36145							257744	
72	0.744961	0.000013000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	37569							257744	
74	0.744971	0.000010000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	38993							257744	
75	0.745111	0.000140000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	40417							257744	
76	0.745116	0.000005000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	41841							257744	
78	0.752471	0.007355000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	43265							257744	
80	0.752601	0.000130000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	44689							257744	



#sf21vus

What about that “Duplicate SACK?” in RFC 2883?

- ⦿ DUP ACK Not == DUP SACK (Nothing in common)
- ⦿ Used by receiver to inform the sender “just received a duplicate stream bytes segment”
- ⦿ Check out Christian Reusch’s SharkFest 21 Virtual Europe Session #12
 - https://www.youtube.com/watch?v=_19UdEtj6Ak

Adventures from the Field



#sf21vus



#sf21vus • Online • September 12-17

Adventures from the Field



- Next we're going to look at actual results from a troubleshooting engagement involving crazy high levels out of sequence packets
- Not just high levels of OOS, but **crazy** high...



#sf21vus

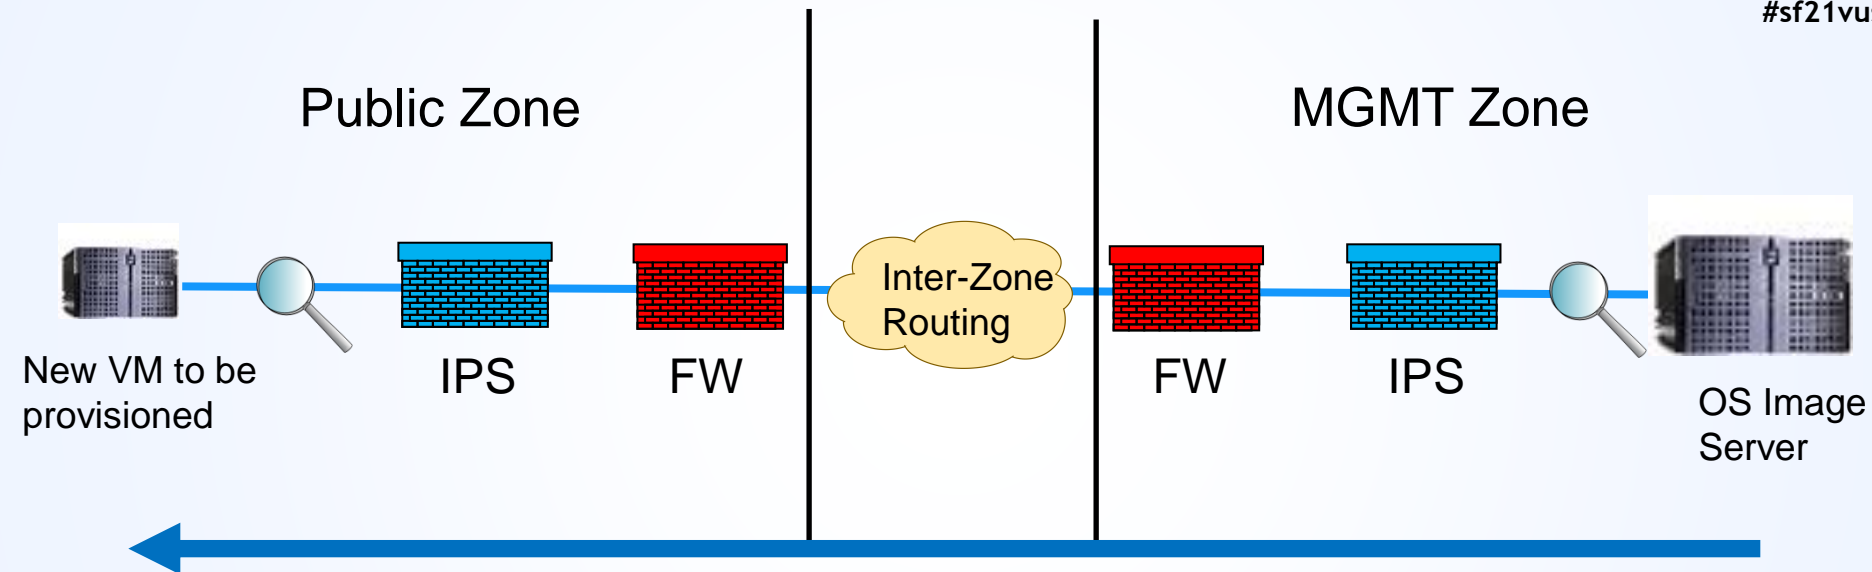
Scenario

- Pre-migration Performance Testing in Lab
- APP: Automation for virtual host provisioning
- Performance was so bad, customer was unable to “green-light” the migration



#sf21vus

Lab Configuration



Traffic: Client Downloading OS Configuration Build Details from Server
1 Gbps Links with no other active traffic



Wireshark Capture Sources



#sf21vus

Non-Technical Issues

- Finger pointing to the extreme...
- ...each vendor (3) is sure they are innocent and that it was the other vendor's issue
- Challenge: Help Customer figure out why throughput is so low, and help identify the vendor causing the problem

Advanced Analytics & Viz

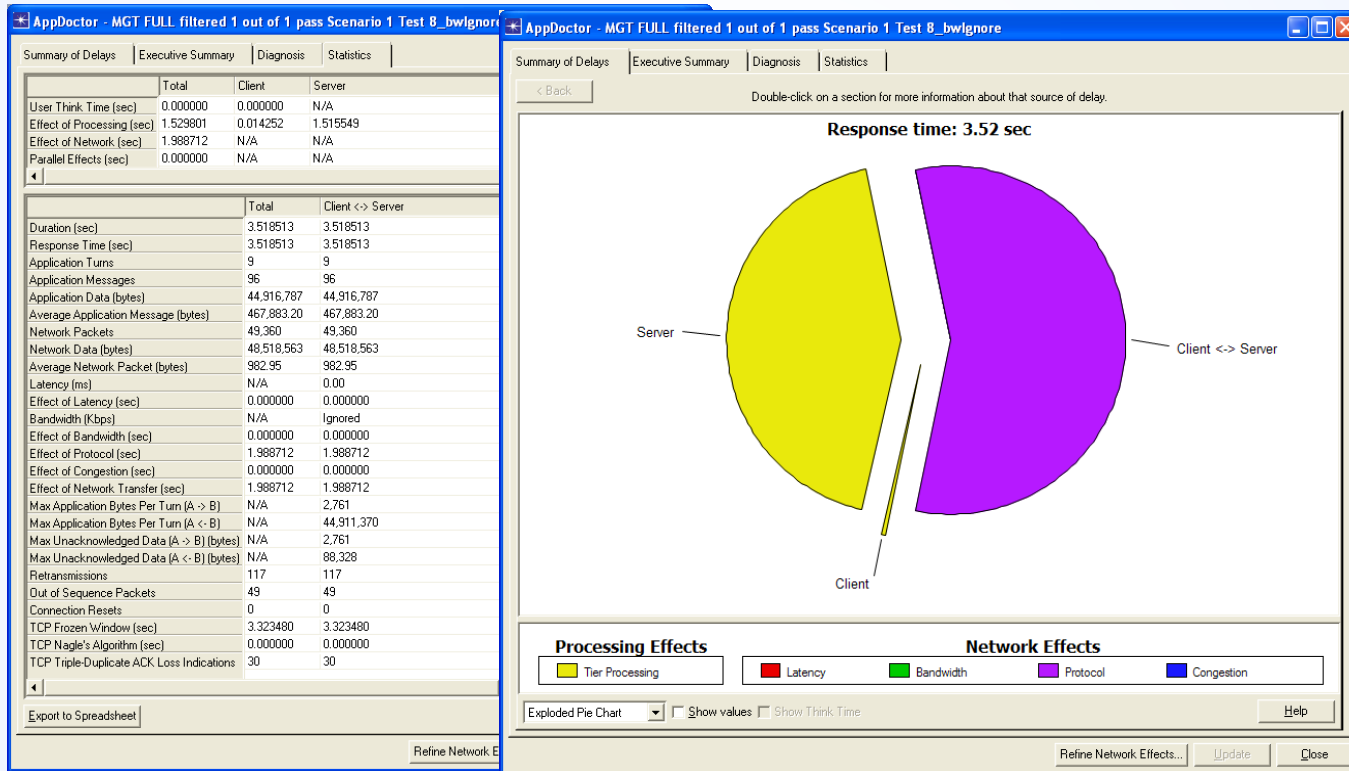


- Let's take a quick look at a visualization of the performance issue, before we deep dive into the packets
- ...with a little help from our Friends...

Summary of Delays



#sf21vus



Summary of Delays



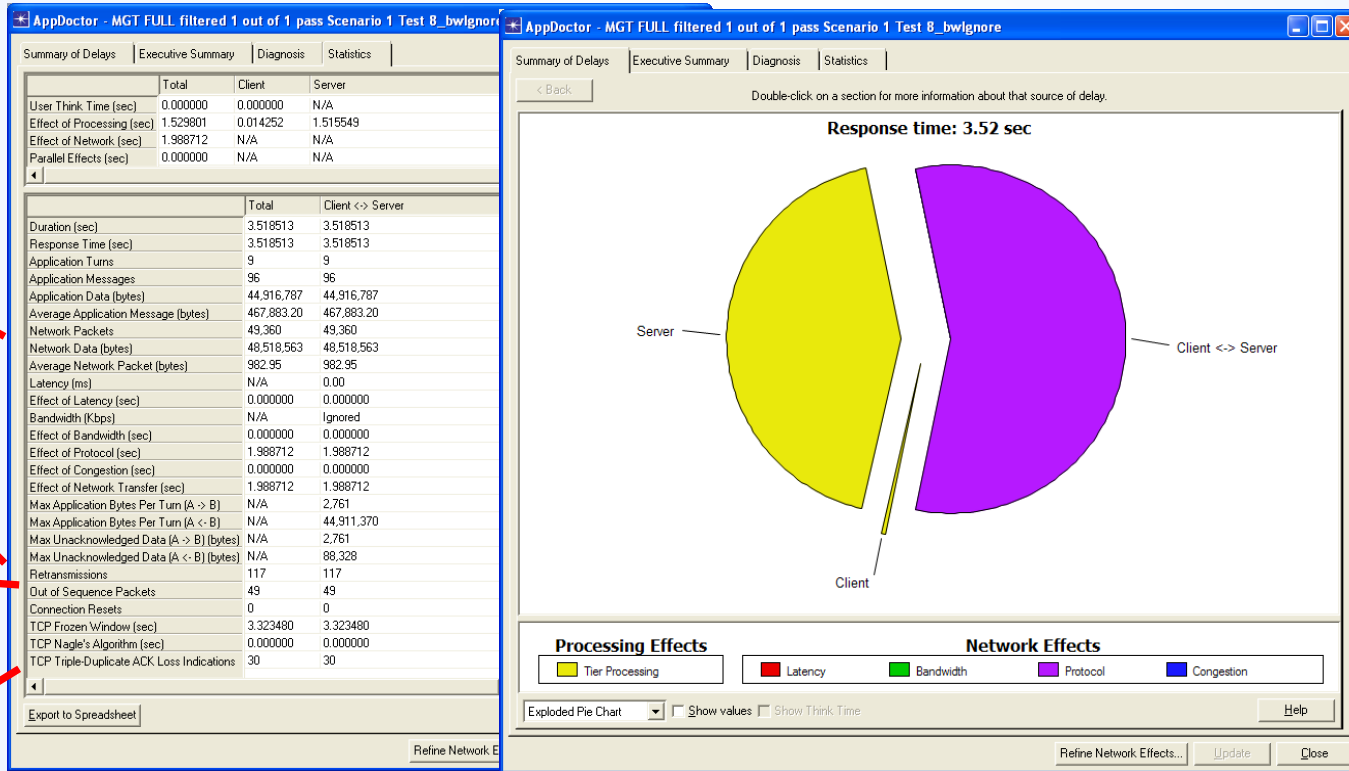
#sf21vus

48.5 MB Xfer

117 x Retrans

49 x OOS

20 x 3ACKs



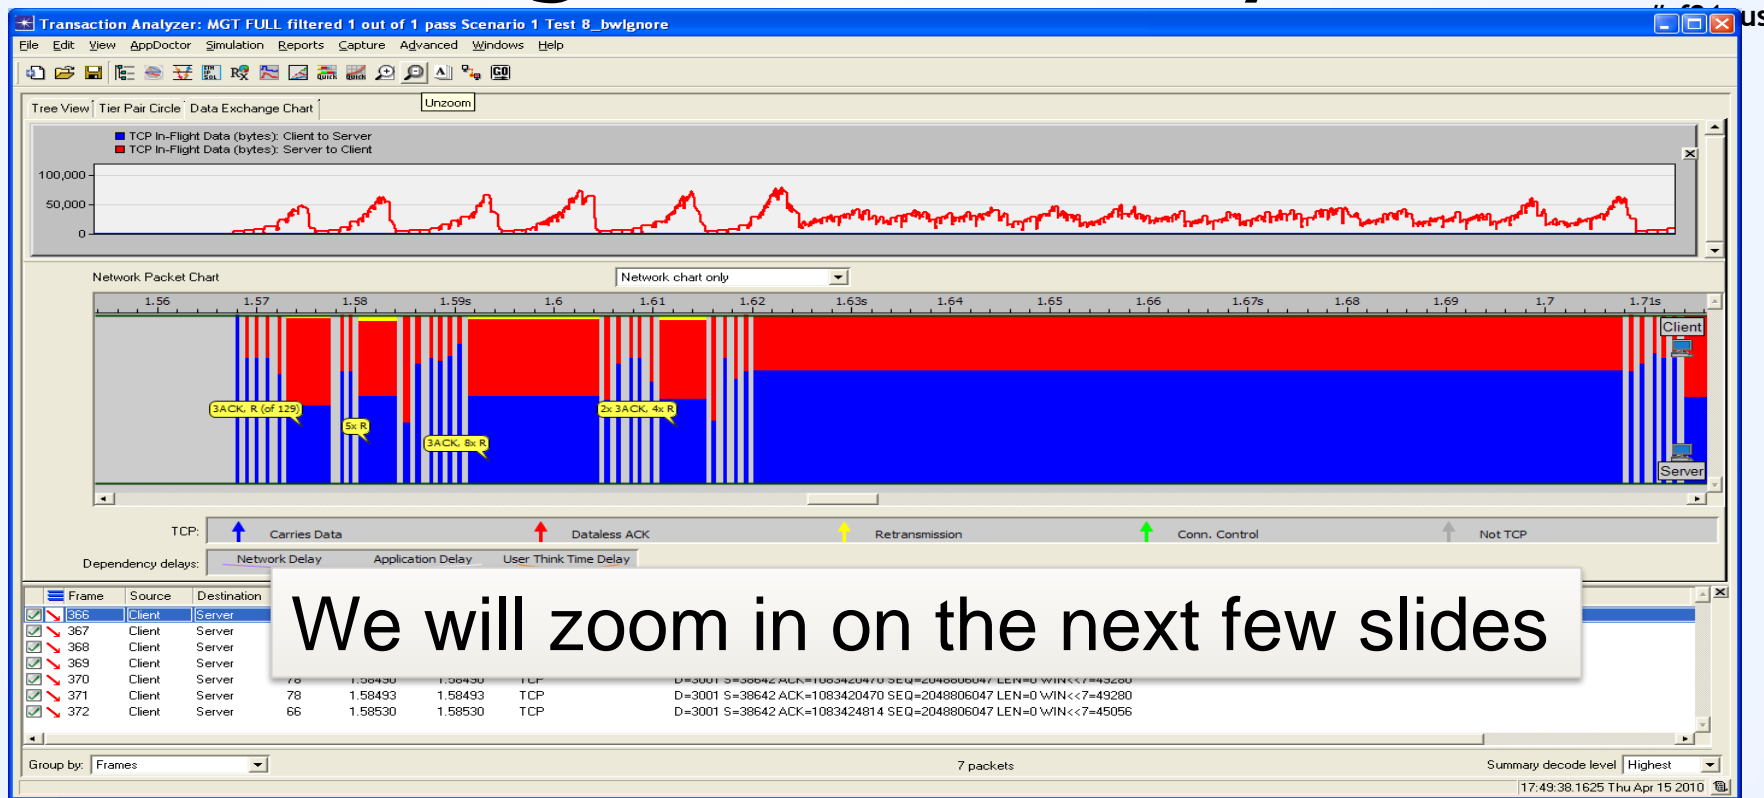
Review: Bytes in Flight



- What does this metric tell us?
- Answer: the number of stream bytes that are outstanding on the network before sender receives an ACK for the left most edge of the current burst
- Usually reflects the sender's Congestion Window




In-flight Data Analysis



Zoom #1 - In-flight Data




#sf21vus



Zoom #2 - In-flight Data




#sf21vus



Zoom #3 - In-flight Data



#sf21vus



Questions / Discussion



#sf21vus

OOS Visualization + SACK Analysis



#sf21vus


- The following section uses time lapse photography to step you through a 19 packet burst chosen at random
- The number, and **nature**, of out of sequence packets is crazy high and it's a nice example to illustrate how to interpret the SACK field

Impact on Performance



#sf21vus

- It's also a great example of how OOS can impact performance & how you can interpret "how bad is bad" with SACK



ACK Packets Corresponding to a Packet Burst



- These are the ACKs from the client

✓	↘	2721	Client	Server	78	1.68016	1.68016	TCP	D=3001 S=38642 ACK=1085576095 SEQ=2048806047 LEN=0 WIN<<7=44416
✓	↘	2722	Client	Server	86	1.68021	1.68021	TCP	D=3001 S=38642 ACK=1085576095 SEQ=2048806047 LEN=0 WIN<<7=44416
✓	↘	2723	Client	Server	86	1.68031	1.68031	TCP	D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0 WIN<<7=42880
✓	↘	2724	Client	Server	86	1.68034	1.68034	TCP	D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0 WIN<<7=42880
✓	↘	2725	Client	Server	86	1.68036	1.68036	TCP	D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0 WIN<<7=42880
✓	↘	2726	Client	Server	78	1.68040	1.68040	TCP	D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0 WIN<<7=38528
✓	↘	2727	Client	Server	86	1.68041	1.68041	TCP	D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0 WIN<<7=38528
✓	↘	2728	Client	Server	94	1.68044	1.68044	TCP	D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0 WIN<<7=38528
✓	↘	2729	Client	Server	94	1.68047	1.68047	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2730	Client	Server	94	1.68049	1.68049	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2731	Client	Server	94	1.68051	1.68051	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2732	Client	Server	94	1.68054	1.68054	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2733	Client	Server	94	1.68056	1.68056	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2734	Client	Server	86	1.68059	1.68059	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2735	Client	Server	78	1.68064	1.68064	TCP	D=3001 S=38642 ACK=1085592023 SEQ=2048806047 LEN=0 WIN<<7=28416
✓	↘	2736	Client	Server	78	1.68064	1.68064	TCP	D=3001 S=38642 ACK=1085592023 SEQ=2048806047 LEN=0 WIN<<7=28416
✓	↘	2737	Client	Server	78	1.68067	1.68067	TCP	D=3001 S=38642 ACK=1085593471 SEQ=2048806047 LEN=0 WIN<<7=27008
✓	↘	2738	Client	Server	78	1.68069	1.68069	TCP	D=3001 S=38642 ACK=1085593471 SEQ=2048806047 LEN=0 WIN<<7=27008
✓	↘	2739	Client	Server	66	1.68083	1.68083	TCP	D=3001 S=38642 ACK=1085605055 SEQ=2048806047 LEN=0 WIN<<7=17280

ACK Packets Corresponding to a Packet Burst



#sf21vus

- Each ACK corresponds to one (or more) of the 19 payload packets in our random sample
- We'll use these ACKs to determine the arrival order for the 19 packets



#sf21vus

Before we start.... a quick Pop Quiz:

1. Why are there so many ACKs, I thought receiver is supposed to ACK of every other packet?

✓	↘	2721	Client	Server	78	1.68016	1.68016	TCP	D=3001 S=38642 ACK=1085576095 SEQ=2048806047 LEN=0 WIN<<7=44416
✓	↘	2722	Client	Server	86	1.68021	1.68021	TCP	D=3001 S=38642 ACK=1085576095 SEQ=2048806047 LEN=0 WIN<<7=44416
✓	↘	2723	Client	Server	86	1.68031	1.68031	TCP	D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0 WIN<<7=42880
✓	↘	2724	Client	Server	86	1.68034	1.68034	TCP	D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0 WIN<<7=42880
✓	↘	2725	Client	Server	86	1.68036	1.68036	TCP	D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0 WIN<<7=42880
✓	↘	2726	Client	Server	78	1.68040	1.68040	TCP	D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0 WIN<<7=38528
✓	↘	2727	Client	Server	86	1.68041	1.68041	TCP	D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0 WIN<<7=38528
✓	↘	2728	Client	Server	94	1.68044	1.68044	TCP	D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0 WIN<<7=38528
✓	↘	2729	Client	Server	94	1.68047	1.68047	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2730	Client	Server	94	1.68049	1.68049	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2731	Client	Server	94	1.68051	1.68051	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2732	Client	Server	94	1.68054	1.68054	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2733	Client	Server	94	1.68056	1.68056	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2734	Client	Server	86	1.68059	1.68059	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2735	Client	Server	78	1.68064	1.68064	TCP	D=3001 S=38642 ACK=1085592023 SEQ=2048806047 LEN=0 WIN<<7=28416
✓	↘	2736	Client	Server	78	1.68064	1.68064	TCP	D=3001 S=38642 ACK=1085592023 SEQ=2048806047 LEN=0 WIN<<7=28416
✓	↘	2737	Client	Server	78	1.68067	1.68067	TCP	D=3001 S=38642 ACK=1085593471 SEQ=2048806047 LEN=0 WIN<<7=27008
✓	↘	2738	Client	Server	78	1.68069	1.68069	TCP	D=3001 S=38642 ACK=1085593471 SEQ=2048806047 LEN=0 WIN<<7=27008
✓	↘	2739	Client	Server	66	1.68083	1.68083	TCP	D=3001 S=38642 ACK=1085605055 SEQ=2048806047 LEN=0 WIN<<7=17280



#sf21vus

2. Why does the ACK packet size change between 66, 78, 86, and 94?

✓	↘	2721	Client	Server	78	1.68016	1.68016	TCP	D=3001 S=38642 ACK=1085576095 SEQ=2048806047 LEN=0 WIN<<7=44416
✓	↘	2722	Client	Server	86	1.68021	1.68021	TCP	D=3001 S=38642 ACK=1085576095 SEQ=2048806047 LEN=0 WIN<<7=44416
✓	↘	2723	Client	Server	86	1.68031	1.68031	TCP	D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0 WIN<<7=42880
✓	↘	2724	Client	Server	86	1.68034	1.68034	TCP	D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0 WIN<<7=42880
✓	↘	2725	Client	Server	86	1.68036	1.68036	TCP	D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0 WIN<<7=42880
✓	↘	2726	Client	Server	78	1.68040	1.68040	TCP	D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0 WIN<<7=38528
✓	↘	2727	Client	Server	86	1.68041	1.68041	TCP	D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0 WIN<<7=38528
✓	↘	2728	Client	Server	94	1.68044	1.68044	TCP	D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0 WIN<<7=38528
✓	↘	2729	Client	Server	94	1.68047	1.68047	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2730	Client	Server	94	1.68049	1.68049	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2731	Client	Server	94	1.68051	1.68051	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2732	Client	Server	94	1.68054	1.68054	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2733	Client	Server	94	1.68056	1.68056	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2734	Client	Server	86	1.68059	1.68059	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2735	Client	Server	78	1.68064	1.68064	TCP	D=3001 S=38642 ACK=1085592023 SEQ=2048806047 LEN=0 WIN<<7=28416
✓	↘	2736	Client	Server	78	1.68064	1.68064	TCP	D=3001 S=38642 ACK=1085592023 SEQ=2048806047 LEN=0 WIN<<7=28416
✓	↘	2737	Client	Server	78	1.68067	1.68067	TCP	D=3001 S=38642 ACK=1085593471 SEQ=2048806047 LEN=0 WIN<<7=27008
✓	↘	2738	Client	Server	78	1.68069	1.68069	TCP	D=3001 S=38642 ACK=1085593471 SEQ=2048806047 LEN=0 WIN<<7=27008
✓	↘	2739	Client	Server	66	1.68083	1.68083	TCP	D=3001 S=38642 ACK=1085605055 SEQ=2048806047 LEN=0 WIN<<7=17280



#sf21vus

3. Why is the receive window continuing to shrink?

✓	↘	2721	Client	Server	78	1.68016	1.68016	TCP	D=3001 S=38642 ACK=1085576095 SEQ=2048806047 LEN=0	WIN<<7=44416
✓	↘	2722	Client	Server	86	1.68021	1.68021	TCP	D=3001 S=38642 ACK=1085576095 SEQ=2048806047 LEN=0	WIN<<7=44416
✓	↘	2723	Client	Server	86	1.68031	1.68031	TCP	D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0	WIN<<7=42880
✓	↘	2724	Client	Server	86	1.68034	1.68034	TCP	D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0	WIN<<7=42880
✓	↘	2725	Client	Server	86	1.68036	1.68036	TCP	D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0	WIN<<7=42880
✓	↘	2726	Client	Server	78	1.68040	1.68040	TCP	D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0	WIN<<7=38528
✓	↘	2727	Client	Server	86	1.68041	1.68041	TCP	D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0	WIN<<7=38528
✓	↘	2728	Client	Server	94	1.68044	1.68044	TCP	D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0	WIN<<7=38528
✓	↘	2729	Client	Server	94	1.68047	1.68047	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0	WIN<<7=37120
✓	↘	2730	Client	Server	94	1.68049	1.68049	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0	WIN<<7=37120
✓	↘	2731	Client	Server	94	1.68051	1.68051	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0	WIN<<7=37120
✓	↘	2732	Client	Server	94	1.68054	1.68054	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0	WIN<<7=37120
✓	↘	2733	Client	Server	94	1.68056	1.68056	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0	WIN<<7=37120
✓	↘	2734	Client	Server	86	1.68059	1.68059	TCP	D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0	WIN<<7=37120
✓	↘	2735	Client	Server	78	1.68064	1.68064	TCP	D=3001 S=38642 ACK=1085592023 SEQ=2048806047 LEN=0	WIN<<7=28416
✓	↘	2736	Client	Server	78	1.68064	1.68064	TCP	D=3001 S=38642 ACK=1085592023 SEQ=2048806047 LEN=0	WIN<<7=28416
✓	↘	2737	Client	Server	78	1.68067	1.68067	TCP	D=3001 S=38642 ACK=1085593471 SEQ=2048806047 LEN=0	WIN<<7=27008
✓	↘	2738	Client	Server	78	1.68069	1.68069	TCP	D=3001 S=38642 ACK=1085593471 SEQ=2048806047 LEN=0	WIN<<7=27008
✓	↘	2739	Client	Server	66	1.68083	1.68083	TCP	D=3001 S=38642 ACK=1085605055 SEQ=2048806047 LEN=0	WIN<<7=17280



#sf21vus


4. Why do we see so many duplicate ACKs?

✓	↘	2721	Client	Server	78	1.68016	1.68016	TCP	D=3001 S=38642 ACK=1085576095 SEQ=2048806047 LEN=0 WIN<<7=44416
✓	↘	2722	Client	Server	86	1.68021	1.68021	TCP	D=3001 S=38642 ACK=1085576095 SEQ=2048806047 LEN=0 WIN<<7=44416
✓	↘	2723	Client	Server	86	1.68031	1.68031	TCP	D=3001 S=38642 ACK=1085577543 EQ=2048806047 LEN=0 WIN<<7=42880
✓	↘	2724	Client	Server	86	1.68034	1.68034	TCP	D=3001 S=38642 ACK=1085577543 EQ=2048806047 LEN=0 WIN<<7=42880
✓	↘	2725	Client	Server	86	1.68036	1.68036	TCP	D=3001 S=38642 ACK=1085577543 EQ=2048806047 LEN=0 WIN<<7=42880
✓	↘	2726	Client	Server	78	1.68040	1.68040	TCP	D=3001 S=38642 ACK=1085581887 EQ=2048806047 LEN=0 WIN<<7=38528
✓	↘	2727	Client	Server	86	1.68041	1.68041	TCP	D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0 WIN<<7=38528
✓	↘	2728	Client	Server	94	1.68044	1.68044	TCP	D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0 WIN<<7=38528
✓	↘	2729	Client	Server	94	1.68047	1.68047	TCP	D=3001 S=38642 ACK=1085583335 EQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2730	Client	Server	94	1.68049	1.68049	TCP	D=3001 S=38642 ACK=1085583335 EQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2731	Client	Server	94	1.68051	1.68051	TCP	D=3001 S=38642 ACK=1085583335 EQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2732	Client	Server	94	1.68054	1.68054	TCP	D=3001 S=38642 ACK=1085583335 EQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2733	Client	Server	94	1.68056	1.68056	TCP	D=3001 S=38642 ACK=1085583335 EQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2734	Client	Server	86	1.68059	1.68059	TCP	D=3001 S=38642 ACK=1085583335 EQ=2048806047 LEN=0 WIN<<7=37120
✓	↘	2735	Client	Server	78	1.68064	1.68064	TCP	D=3001 S=38642 ACK=1085592023 SEQ=2048806047 LEN=0 WIN<<7=28416
✓	↘	2736	Client	Server	78	1.68064	1.68064	TCP	D=3001 S=38642 ACK=1085592023 SEQ=2048806047 LEN=0 WIN<<7=28416
✓	↘	2737	Client	Server	78	1.68067	1.68067	TCP	D=3001 S=38642 ACK=1085593471 SEQ=2048806047 LEN=0 WIN<<7=27008
✓	↘	2738	Client	Server	78	1.68069	1.68069	TCP	D=3001 S=38642 ACK=1085593471 SEQ=2048806047 LEN=0 WIN<<7=27008
✓	↘	2739	Client	Server	66	1.68083	1.68083	TCP	D=3001 S=38642 ACK=1085605055 SEQ=2048806047 LEN=0 WIN<<7=17280

Questions / Comments



#sf21vus



#sf21vus

Pre-Departure Orientation



#sf21vus

- 19 Slide Journey
- The top portion of the slide shows you which packet in the burst has been received
- The bottom portion shows you the ACK and SACK values extracted from the corresponding ACK packets
- Each slide represents a new packet being received and the state of all previously received packets

Orientation



#sf21vus

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607



#sf21vus

Orientation

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,783	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,231	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,679	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,127	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,575	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,023	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,471	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

This is the frame order as seen in the Public capture (closest to Receiver)



#sf21vus

Orient...

This is the frame order as seen in the Mgmt capture – traffic in transit to Public (closest to Sender)

Pub Frame	Mgmt Frame	IP ID	SEQ #	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,448	1,085,603,607

This is the frame order as seen in the Public capture (closest to Receiver)

Orient...

This is the frame order as seen in the Mgmt capture – traffic in transit to Public (closest to Sender)

Pub Frame	Mgmt Frame	IP ID	SEQ #	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,448	1,085,586,231
		33,058	1,085,586,231	1,448	1,085,587,679
		33,060	1,085,587,679	1,448	1,085,589,127
		33,062	1,085,589,127	1,448	1,085,590,575
		33,064	1,085,590,575	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,448	1,085,603,607

Packet Just Received
Bracketed in Red

This is the frame order as
seen in the Public capture
(closest to Receiver)

Oriente

This is the frame order as seen in the Mgmt capture – traffic in transit to Public (closest to Sender)

Pub Frame	Mgmt Frame	IP ID	SEQ #	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,448	1,085,586,231
		33,058	1,085,586,231	1,448	1,085,587,679
		33,060	1,085,587,679	1,448	1,085,589,127
		33,062	1,085,589,127	1,448	1,085,590,575
		33,064	1,085,590,575	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,448	1,085,603,607

Packet Just Received
Bracketed in Red

This is the frame order as
seen in the Public capture
(closest to Receiver)

Frames previously received
bracketed in Green



#sf21vus

ACK Details for each packet received

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
			1,085,594,919	1,085,596,366	1,448	1,085,596,367
			1,085,596,367	1,085,597,814	1,448	1,085,597,815
			1,085,597,815	1,085,599,262	1,448	1,085,599,263
			1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

This is the TCP Header from ACK
Packet's Decode Summary

TCP D=3001 S=38642 ACK=1085576095 SEQ=2048806047 LEN=0 WIN<<7=44416

SACK: 1085578991-1085580439
left edge = 1085578991
right edge = 1085580439



#sf21vus

ACK Details for each packet received

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
			1,085,594,919	1,085,596,366	1,448	1,085,596,367
			1,085,596,367	1,085,597,814	1,448	1,085,597,815
			1,085,597,815	1,085,599,262	1,448	1,085,599,263
			1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

This is the TCP Header from ACK
Packet's Decode Summary

TCP D=3001 S=38642 ACK=1085576095 SEQ=2048806047 LEN=0 WIN<<7=44416

SACK: 1085578991-1085580439
left edge = 1085578991
right edge = 1085580439

This is the value of the SACK from
TCP Options Field



#sf21vus

Ready to start our Wild Ride....?

- ⦿ Fasten your seat belt..
- ⦿ Focus on the Cumulative ACK values and the SACK values as each packet is received..
- ⦿ Double check your understanding, ask if what you're seeing tracks with our earlier study of SACK
- ⦿ Prior to the start of this sequence, receiver had signaled that he's ready to receive the stream starting at byte:


1,085,576,095





#sf21vus

1st Packet Received



Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

After receipt of the above packet (in Red), the receiver issued the following ACK

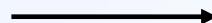
```
TCP D=3001 S=38642 ACK=1085576095 SEQ=2048806047 LEN=0 WIN<<7=44416
```

```
SACK: 1085578991-1085580439  
left edge = 1085578991  
right edge = 1085580439
```



#sf21vus

2nd Packet Received



Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

After receipt of the above packet (in Red), the receiver issued the following ACK

```
IP      Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64741
TCP    D=3001 S=38642 ACK=1085576095 SEQ=2048806047 LEN=0 WIN<<7=44416
```

```
SACK: 1085587679-1085589127 1085578991-1085580439
left edge = 1085587679
right edge = 1085589127
left edge = 1085578991
right edge = 1085580439
```



#sf21vus

3rd Packet Received

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64742
TCP D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0 WIN<<7=42880

SACK: 1085587679-1085589127 1085578991-1085580439
left edge = 1085587679
right edge = 1085589127
left edge = 1085578991
right edge = 1085580439

Notice the cumulative ACK has increased to a value of 1085577543

SACK Field has not changed



#sf21vus

4th Packet Received

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64743
 TCP D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0 WIN<<7=42880

SACK: 1085587679-1085590575 1085578991-1085580439

left edge = 1085587679
 right edge = 1085590575
 left edge = 1085578991
 right edge = 1085580439

← Right edge changed



#sf21vus

5th Packet Received

	Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
→	2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
	2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
	2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
	2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
	2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
	2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
	2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
	2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
	2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
	2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
	2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
	2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
	2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
	2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
	2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
	2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
	2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
	2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
	2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64744
TCP D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0 WIN<<7=42880

SACK: 1085578991-1085581887 1085587679-1085590575
left edge = 1085578991
right edge = 1085581887
left edge = 1085587679
right edge = 1085590575

Right edge changed



#sf21vus

5th Packet Received

	Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
→	2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
	2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
	2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
	2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
	2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
	2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
	2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
	2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
	2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
	2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
	2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
	2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
	2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
	2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
	2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
	2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
	2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
	2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
	2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64744
TCP D=3001 S=38642 ACK=1085577543 SEQ=2048806047 LEN=0 WIN<<7=42880

SACK: 1085578991-1085581887 1085587679-1085590575



left edge = 1085578991
right edge = 1085581887
left edge = 1085587679
right edge = 1085590575

*** and *** the order of the SACK blocks has changed



#sf21vus

Packet #6

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64745
TCP D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0 WIN<<7=38528

SACK: 1085587679-1085590575
left edge = 1085587679
right edge = 1085590575

Cumulative ACK is updated, and one SACK block removed



#sf21vus

Packet #7

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64746

TCP D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0 WIN<<7=38528

SACK: 1085596367-1085597815 1085587679-1085590575

left edge = 1085596367

right edge = 1085597815

left edge = 1085587679

right edge = 1085590575



#sf21vus

Packet #8

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64747
 TCP D=3001 S=38642 ACK=1085581887 SEQ=2048806047 LEN=0 WIN<<7=38528

SACK: 1085584783-1085586231 1085596367-1085597815 1085587679-1085590575

left edge = 1085584783
 right edge = 1085586231
 left edge = 1085596367
 right edge = 1085597815
 left edge = 1085587679
 right edge = 1085590575

SACK now indicates three groups of non-contiguous stream bytes have been received



#sf21vus

Packet #9

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64748
 TCP D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120

SACK: 1085584783-1085586231 1085596367-1085597815 1085587679-1085590575
 left edge = 1085584783
 right edge = 1085586231
 left edge = 1085596367
 right edge = 1085597815
 left edge = 1085587679
 right edge = 1085590575



#sf21vus

Packet #10

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64749
 TCP D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120

SACK: 1085596367-1085599263 1085584783-1085586231 1085587679-1085590575
 left edge = 1085596367
 right edge = 1085599263
 left edge = 1085584783
 right edge = 1085586231
 left edge = 1085587679
 right edge = 1085590575



#sf21vus

Packet #11

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64750
 TCP D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120

SACK: 1085587679-1085592023 1085596367-1085599263 1085584783-1085586231
 left edge = 1085587679
 right edge = 1085592023
 left edge = 1085596367
 right edge = 1085599263
 left edge = 1085584783
 right edge = 1085586231



#sf21vus

Packets #12 + 13

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64752
 TCP D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120

SACK: 1085596367-1085602159 1085587679-1085592023 1085584783-1085586231
 left edge = 1085596367
 right edge = 1085602159
 left edge = 1085587679
 right edge = 1085592023
 left edge = 1085584783
 right edge = 1085586231



#sf21vus

Packet #14

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64753
 TCP D=3001 S=38642 ACK=1085583335 SEQ=2048806047 LEN=0 WIN<<7=37120

SACK: 1085584783-1085592023 1085596367-1085602159
 left edge = 1085584783
 right edge = 1085592023
 left edge = 1085596367
 right edge = 1085602159

Notice that two of the dis-contiguous blocks are now contiguous; so we go from 3 blocks down to 2 blocks



#sf21vus

Packet #15

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64754
 TCP D=3001 S=38642 ACK=1085592023 SEQ=2048806047 LEN=0 WIN<<7=28416

SACK: 1085596367-1085602159
 left edge = 1085596367
 right edge = 1085602159

Notice the cumulative ACK has increased to a value of 1085592023 and we're down to just one dis-contiguous block



#sf21vus

Packet #16

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64755
TCP D=3001 S=38642 ACK=1085592023 SEQ=2048806047 LEN=0 WIN<<7=28416

SACK: 1085594919-1085602159
left edge = 1085594919
right edge = 1085602159

Left edge updated to reflect packet #16



#sf21vus

Packet #17

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64756
TCP D=3001 S=38642 ACK=1085593471 SEQ=2048806047 LEN=0 WIN<<7=27008

SACK: 1085594919-1085602159
left edge = 1085594919
right edge = 1085602159

Cumulative ACK is updated to reflect receipt of #17
No change to SACK fields



#sf21vus

Packet #18

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64757
TCP D=3001 S=38642 ACK=1085593471 SEQ=2048806047 LEN=0 WIN<<7=27008

SACK: 1085594919-1085603607
left edge = 1085594919
right edge = 1085603607

Right edge updated to reflect packet #18



#sf21vus

Packet #19

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

IP Internet Protocol, Src: 10.128.194.191 (10.128.194.191), Dst: 10.153.2.25 (10.153.2.25) ID=64758
TCP D=3001 S=38642 ACK=1085605055 SEQ=2048806047 LEN=0 WIN<<7=17280

New cumulative ACK reflects receipt of all 19 packets
(Plus, packet 20 not shown in the list....)



#sf21vus

After packet 19

Pub Frame	Mgmt Frame	IP ID	SEQ #	Last Byte	Len	Next Seq #
2,703	2,692	33,044	1,085,576,095	1,085,577,542	1,448	1,085,577,543
2,706	2,693	33,046	1,085,577,543	1,085,578,990	1,448	1,085,578,991
2,701	2,694	33,048	1,085,578,991	1,085,580,438	1,448	1,085,580,439
2,705	2,695	33,050	1,085,580,439	1,085,581,886	1,448	1,085,581,887
2,709	2,696	33,052	1,085,581,887	1,085,583,334	1,448	1,085,583,335
2,715	2,697	33,054	1,085,583,335	1,085,584,782	1,448	1,085,584,783
2,708	2,698	33,056	1,085,584,783	1,085,586,230	1,448	1,085,586,231
2,714	2,699	33,058	1,085,586,231	1,085,587,678	1,448	1,085,587,679
2,702	2,700	33,060	1,085,587,679	1,085,589,126	1,448	1,085,589,127
2,704	2,701	33,062	1,085,589,127	1,085,590,574	1,448	1,085,590,575
2,711	2,702	33,064	1,085,590,575	1,085,592,022	1,448	1,085,592,023
2,717	2,703	33,066	1,085,592,023	1,085,593,470	1,448	1,085,593,471
2,719	2,704	33,068	1,085,593,471	1,085,594,918	1,448	1,085,594,919
2,716	2,705	33,070	1,085,594,919	1,085,596,366	1,448	1,085,596,367
2,707	2,706	33,072	1,085,596,367	1,085,597,814	1,448	1,085,597,815
2,710	2,707	33,074	1,085,597,815	1,085,599,262	1,448	1,085,599,263
2,712	2,708	33,076	1,085,599,263	1,085,600,710	1,448	1,085,600,711
2,713	2,709	33,078	1,085,600,711	1,085,602,158	1,448	1,085,602,159
2,718	2,710	33,080	1,085,602,159	1,085,603,606	1,448	1,085,603,607

It's been a long, strange journey, but all data has finally been received



#sf21vus

Buffer Mgmt Issues?

- FIFO?
- LIFO?
- UFO?
- IFFY?

Questions / Comments



#sf21vus



End of Session





#sf21vus

- This is where the online presentation stopped
- A few more reminders and notes follow and can be viewed after the conference
- Slide deck will be uploaded to the Sharkfest Retrospective micro-site

Appendix



#sf21vus



#sf21vus

Reminders

- You can quickly determine presence of SACK in Wireshark using a display filter!
- “tcp.options.sack.count”
- You can easily add SACK related columns to Wireshark GUI



#sf21vus

Closing Remarks



- Focus on “Bytes in Flight” Data
 - If you see the congestion window constantly closing or reduced by half, then you need to figure out why
 - Interpreting SACK might help complete the picture
- It's easy to get lost drilling in to SACK fields...
- It's prudent to interpret some of them and make a high level assessment as to the extent of OOS packets
 - To interpret them you have to understand the RFC and expected behavior



#sf21vus

Closing Remarks

- If you have a lot of SACKs with 3 or 4 gaps declared, then OOS is “high” / “pervasive”
- If you have a few SACKs with only 1 gap each, then OOS may be less of a contributing factor to slow performance

Outcome from this Study



- Client was very pleased that we could help them understand the full extent of the OOS problem
 - Showed that packets are not “just a little out of sequence” but significantly out of sequence
 - Definitely impacted sender’s ability to maintain a large congestion window
 - Client re-evaluated plans to deploy more IPS devices

Outcome from this Study



- Client shared results with their IPS vendor which triggered a major investigation into stream and buffer management in the IPS

Outcome from this Study




- A few months later we tested a new model of IPS in Client's lab
- Some improvement but still a problem even at low throughput levels

Cause and Effect



#sf21vus





#sf21vus

Review: Effect of the OOS on the sender

- Potential Throughput Killer: Will likely trigger TCP congestion window reduction if he has to retransmit
- Dependent on the OS and patch level of the sender...and possibly the NIC driver (maybe)
- The RFC for SACK has a lot of “should”s and “may”s.
 - The implementer is allowed a lot flexibility in how they handle the SACK information provided by the receiver

Effect of the SACK field on the sender



#sf21vus

- Consider: should the sender retransmit just one missing segment, or if he can see from the SACK that lot's of different packets are missing should he retransmit all of them
- Also, sender has to maintain all packets in the retransmit queue until they've been ACK'd, possible stress on memory

Impact on Receiver



- He has to buffer all out of sequence packets
- Can not deliver any discontinuous stream bytes to the app until all missing packets are received
- Will generate more ACKs – one for each OOS packet received
- What happens if there are lots of gaps?
 - Remember SACK can only record up to 4 gaps (3 if timestamp option is also being used)

Ready to Wrap?



- Let's look at a few reminders...

SACK Wireshark Columns



#sf21vus

No.	Time	Source	Destination	Length	TCP SACK Count	TCP SACK Left Edge	TCP SACK Right Edge	Bytes in flight	Info
31	0.711313	Server	Client	1496				5704	8085 → 43650 [PSH, ACK] Seq=13341 Ack=2171 Win=32044 Len=1430 TSval=1376522
32	0.711325	Server	Client	2914				8552	8085 → 43650 [ACK] Seq=14771 Ack=2171 Win=32044 Len=2848 TSval=1376522143 T
33	0.711339	Server	Client	2914				11400	8085 → 43650 [ACK] Seq=17619 Ack=2171 Win=32044 Len=2848 TSval=1376522143 T
34	0.718880	Client	Server	66				43650	→ 8085 [ACK] Seq=2171 Ack=10493 Win=263440 Len=0 TSval=525253239 TSec
35	0.718888	Server	Client	74				9982	8085 → 43650 [PSH, ACK] Seq=20467 Ack=2171 Win=32044 Len=8 TSval=1376522144
36	0.719162	Client	Server	66				43650	→ 8085 [ACK] Seq=2171 Ack=13341 Win=260592 Len=0 TSval=525253239 TSec
37	0.719441	Client	Server	66				43650	→ 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 TSval=525253239 TSec
38	0.719446	Client	Server	78	1	17619	19043		[TCP Dup ACK 37#1] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 T
39	0.721344	Server	Client	1492				7130	8085 → 43650 [PSH, ACK] Seq=20475 Ack=2171 Win=32044 Len=1426 TSval=1376522
40	0.721452	Server	Client	1492				8556	8085 → 43650 [PSH, ACK] Seq=21901 Ack=2171 Win=32044 Len=1426 TSval=1376522
41	0.722115	Server	Client	1492				9982	8085 → 43650 [PSH, ACK] Seq=23327 Ack=2171 Win=32044 Len=1426 TSval=1376522
42	0.726388	Client	Server	86	2	20467,17619	20475,19043		[TCP Dup ACK 37#2] 43650 → 8085 [ACK] Seq=2171 Ack=14771 Win=263440 Len=0 T
43	0.726408	Server	Client	1490				9982	[TCP Fast Retransmission] 8085 → 43650 [ACK] Seq=14771 Ack=2171 Win=32044 L

Urgent pointer: 0

Options: (24 bytes), No-Operation (NOP), No-Operation (NOP), Timestamps, No-Operation (NOP), No-Operation (NOP), SACK

▷ No-Operation (NOP)

▷ No-Operation (NOP)

▷ Timestamps: TSval 525253239, TSecr 1376522143

▷ No-Operation (NOP)

▷ No-Operation (NOP)

▷ SACK: 17619-19043

Kind: SACK (5)

Length: 10

left edge = 17619 (relative)

right edge = 19043 (relative)

[TCP SACK Count: 1]

All bytes through 14770r have been received and I'm ready for 14771r. But wait there's more....

Power of Profiles



#sf21vus

SACK-Sample-Conn43650 - Copy_anon.pcapng

File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help

Apply a display filter ... <Ctrl-/>

No.	Time	Source	Destination	Protocol	Length	SACK Count	Cumm ACK	SACK LE 1	SACK RE 1	SACK LE 2	SACK RE 2	Info
37	0.719441	172.20.1.1	10.200.50.1	TCP	66		14771					43650 → 8085 [A
38	0.719446	172.20.1.1	10.200.50.1	TCP	78	1	14771	17619	19043			[TCP Dup ACK 37
39	0.721344	10.200.50.1	172.20.1.1	TCP	1492		2171					8085 → 43650 [P
40	0.721452	10.200.50.1	172.20.1.1	TCP	1492		2171					8085 → 43650 [P
41	0.722115	10.200.50.1	172.20.1.1	TCP	1492		2171					8085 → 43650 [P
42	0.726388	172.20.1.1	10.200.50.1	TCP	86	2	14771	20467	20475	17619	19043	[TCP Dup ACK 37
43	0.726408	10.200.50.1	172.20.1.1	TCP	1490		2171					[TCP Fast Retra
44	0.728899	172.20.1.1	10.200.50.1	TCP	86	2	14771	20467	21899	17619	19043	[TCP Dup ACK 37
45	0.728910	172.20.1.1	10.200.50.1	TCP	86	2	14771	20467	21901	17619	19043	[TCP Dup ACK 37
46	0.728915	10.200.50.1	172.20.1.1	TCP	1490		2171					[TCP Out-Of-Ord
47	0.728920	10.200.50.1	172.20.1.1	TCP	1490		2171					[TCP Out-Of-Ord

Power of Colorization



#sf21vus

SACK-Sample-Conn43650 - anon.pcapng

File Edit View Go Capture Analyze Statistics

tcp.options.sack

No.	Time	Delta	Source	Destination	Length	Seq	NextSeq	ACK	LE1	RE1	LE2	RE2	LE3	RE3	LE4	RE4	Bytes in flight	R_Win
38	0.719446	0.000000000	172.20.1.1	10.200.50.1	78	2171	2171	14771	17619	19043							263440	
42	0.726388	0.006942000	172.20.1.1	10.200.50.1	86	2171	2171	14771	20467	20475	17619	19043					263440	
44	0.728899	0.002511000	172.20.1.1	10.200.50.1	86	2171	2171	14771	20467	21899	17619	19043					263440	
45	0.728910	0.000011000	172.20.1.1	10.200.50.1	86	2171	2171	14771	20467	21901	17619	19043					263440	
48	0.729045	0.000135000	172.20.1.1	10.200.50.1	86	2171	2171	14771	20467	23325	17619	19043					263440	
50	0.729058	0.000013000	172.20.1.1	10.200.50.1	86	2171	2171	14771	20467	23327	17619	19043					263440	
52	0.729486	0.000428000	172.20.1.1	10.200.50.1	86	2171	2171	14771	20467	24751	17619	19043					263440	
54	0.729496	0.000010000	172.20.1.1	10.200.50.1	86	2171	2171	14771	20467	24753	17619	19043					263440	
56	0.734330	0.004834000	172.20.1.1	10.200.50.1	86	2171	2171	16195	20467	24753	17619	19043					263440	
58	0.736839	0.002509000	172.20.1.1	10.200.50.1	78	2171	2171	19043	20467	24753							263440	
66	0.744708	0.007869000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	33297							257744	
68	0.744874	0.000166000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	34721							257744	
70	0.744948	0.000074000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	36145							257744	
72	0.744961	0.000013000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	37569							257744	
74	0.744971	0.000010000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	38993							257744	
75	0.745111	0.000140000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	40417							257744	
76	0.745116	0.000005000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	41841							257744	
78	0.752471	0.007355000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	43265							257744	
80	0.752601	0.000130000	172.20.1.1	10.200.50.1	78	2171	2171	30449	31873	44689							257744	



#sf21vus

Firewall Effects

- Some firewalls will randomize the starting TCP SYN sequence number when new connections are created
- The receiver only knows the randomized version of the sequence numbers generated by the Firewall



#sf21vus

Firewall Effects

- When receiver creates ACKs with SACK values, the SACK sequence numbers will not match the sequence numbers in the TCP header known by the sender
 - Firewall will always restore the original sequence numbers in the TCP header only
 - ...but no guarantees for translating the SACK field
 - This generally makes the SACK field unusable for the sender



#sf21vus

Example of Firewall SEQ Randomization

- SACK sequence numbers bare no resemblance to the SEQ or ACK in the TCP header

```
103574
  ETH      Ethernet II, Src: Cisco_9b:58:00 (00:1a:30:9b:58:00), Dst: SunMicro_9d:78:ee (00:14:4f:9d:78:ee)
  IP       Internet Protocol, Src: 10.144.21.19 (10.144.21.19), Dst: 10.10.81.21 (10.10.81.21) ID=21991
  TCP     D=49242 S=1526 ACK=1472937932 SEQ=1369621428 LEN=0 WIN=151
          Source port: pdap-np (1526)
          Destination port: 49242 (49242)
          Sequence number: 1369621428
          Acknowledgement number: 1472937932
          Header length: 32 bytes
          Flags: 0x10 (ACK)
          0... .... = Congestion Window Reduced (CWR): Not set
          .0.. .... = ECN-Echo: Not set
          ..0. .... = Urgent: Not set
          ...1 .... = Acknowledgment: Set
          .... 0... = Push: Not set
          .... .0.. = Reset: Not set
          .... ..0. = Syn: Not set
          .... ...0 = Fin: Not set
          Window size: 151
          Checksum: 0xca1b [correct]
          [Good Checksum: True]
          [Bad Checksum: False]
          Options: (12 bytes)
          NOP
          NOP
          SACK: 3215007281-3215008661
          left edge = 3215007281
          right edge = 3215008661
```

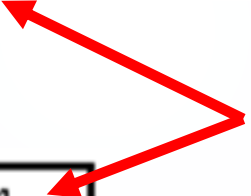
Example of Firewall SEQ Randomization



#sf21vus

- Zoom in from previous slide

```
➔ 103574
.... ETH      Ethernet II, Src: Cisco_9b:58:00 (00:1a:30:9b:58:00), Dst: SunMicro_9
.... IP       Internet Protocol, Src: 10.144.21.19 (10.144.21.19), Dst: 10.10.81.21
.... TCP      D=49242 S=1526 ACK=1472937932 SEQ=1369621428 LEN=0 WIN=151
```




```
SACK: 3215007281-3215008661
      left edge = 3215007281
      right edge = 3215008661
```


SACK is not a Promise



#sf21vus



Effect of out of sequence arrivals on the receiver



- Receiver is allowed to “reneg” if he runs out of buffer space

8. Data Receiver Reneging

Note that the data receiver is permitted to discard data in its queue that has not been acknowledged to the data sender, even if the data has already been reported in a SACK option. Such discarding of SACKed packets is discouraged, but may be used if the receiver runs out of buffer space.

Effect on Sender's NIC



#sf21vus

- What if TSO is enabled?
- What if TCP Chimney is enabled (Windows)?
- Who is managing the retransmit queue...the TCP Stack on the OS or the NIC?
- I pose these questions because they might be important...
- The specific NIC brand, driver version, and firmware version may impact answers to the above..

Effect on Sender's ESX Host NIC



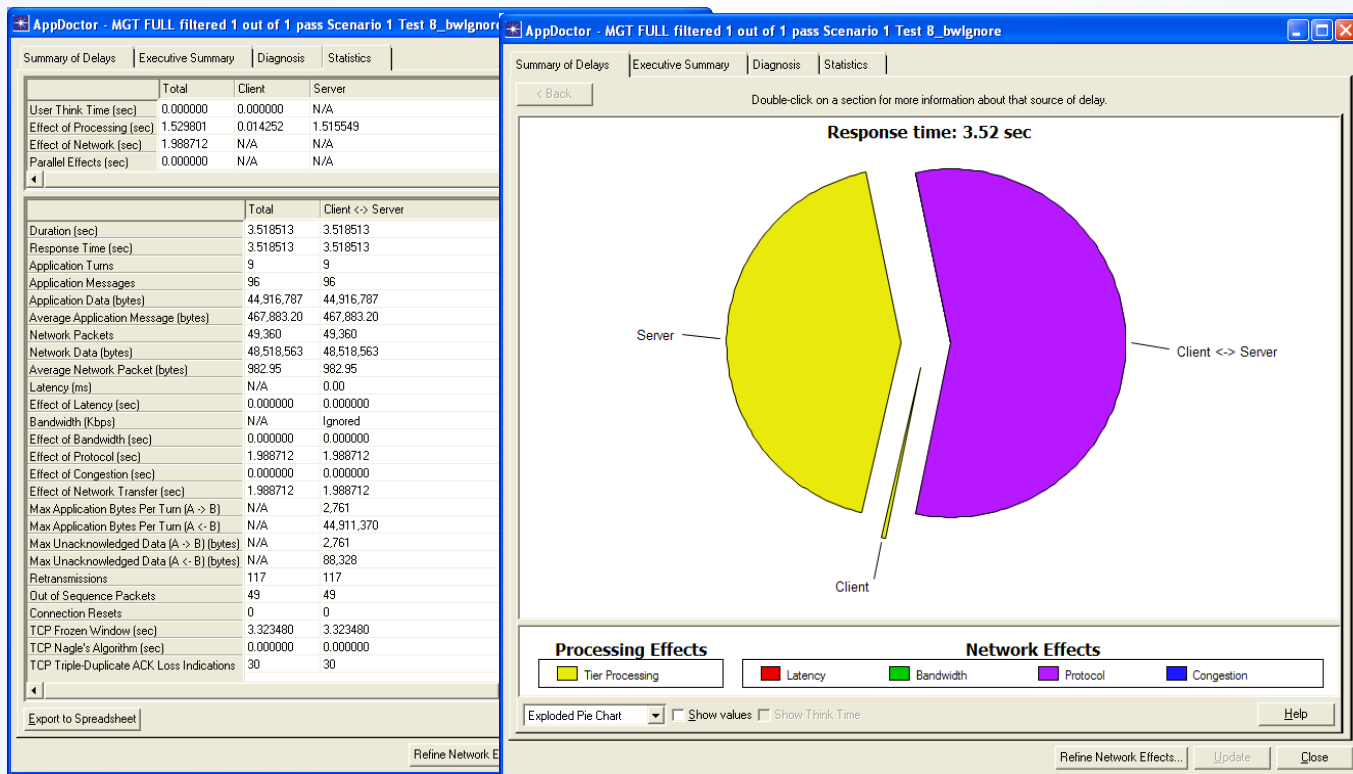
#sf21vus

- What if TSO is enabled on the Physical NIC?
- Who is managing the retransmit queue...the TCP Stack on ESX, NIC, OS or the vNIC?
- The specific NIC brand, driver version, and firmware version may impact answers to the above..

Supplemental Analytics Help



#sf21vus





End of Session

