SINE NOMINE
ASSOCIATES

Wireshark and AFS

Mark Vitale
Sine Nomine Associates
2021 AFS Technologies Workshop

Executive summary

SINE NOMINE ASSOCIATES

- Wireshark works well for debugging AFS and Rx
- Incremental improvements have been made over the past year
- More incremental improvements are in the works

What is Wireshark?

SINE NOMINE ASSOCIATES

"**Wireshark** is the world's foremost and widely-used network protocol analyzer. It lets you see what's happening on your network at a microscopic level…"

https://www.wireshark.org

Why Wireshark ?

SINE NOMINE ASSOCIATES

- AFS is a distributed network filesystem
  - although local diagnostics like BPF, fssync-debug, rxstats, etc. are great, sometimes there's not substitute for analyzing the communication between hosts
- Wireshark is best-of-breed for this task
  - although tcpdump also understands AFS and Rx, Wireshark understands them better.

Wireshark functions

SINE NOMINE ASSOCIATES

- live packet capture from almost any interface
- read or write almost any packet capture format
- deep packet analysis for almost any protocol
- rich capture and display filtering

Wireshark components

SINE NOMINE ASSOCIATES

- Wireshark graphical interface (GUI)
- tshark command line interface (CLI)
- dissectors – “hundreds of protocols”
  - built-in dissectors
    - in-tree; slower to develop; faster to load
    - Wireshark source epan/dissectors/*
      - packet-rx.c — the Rx dissector
      - packet-afs.c — the AFS dissector
  - plugin dissectors
    - in-tree or out-of-tree; faster to develop; slower to load
    - out-of-tree may be closed source for proprietary protocols
    - Wireshark source plugins/epan/*

Relevant dissectors for AFS

SINE NOMINE ASSOCIATES

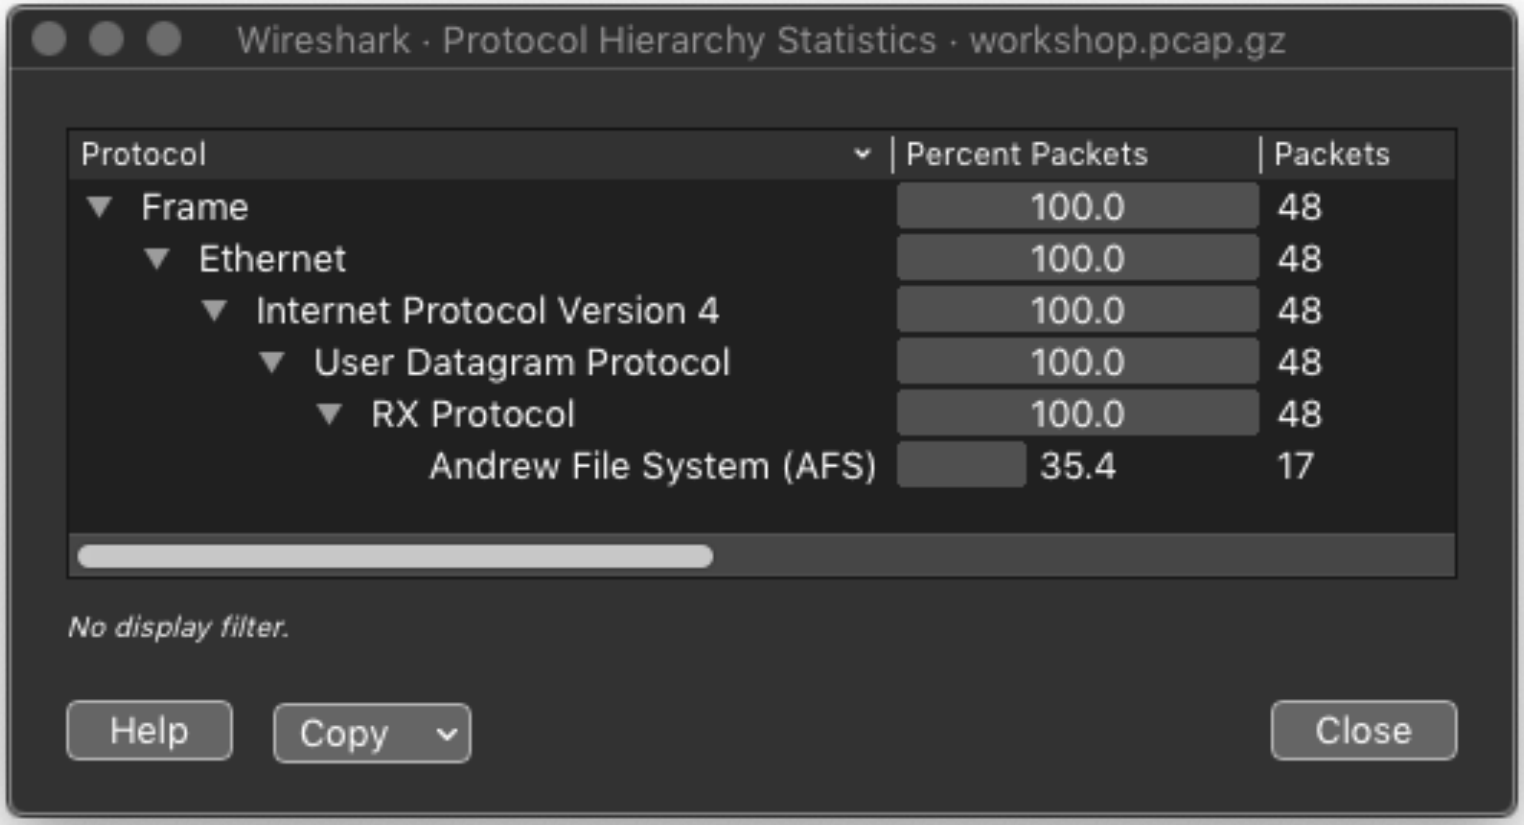

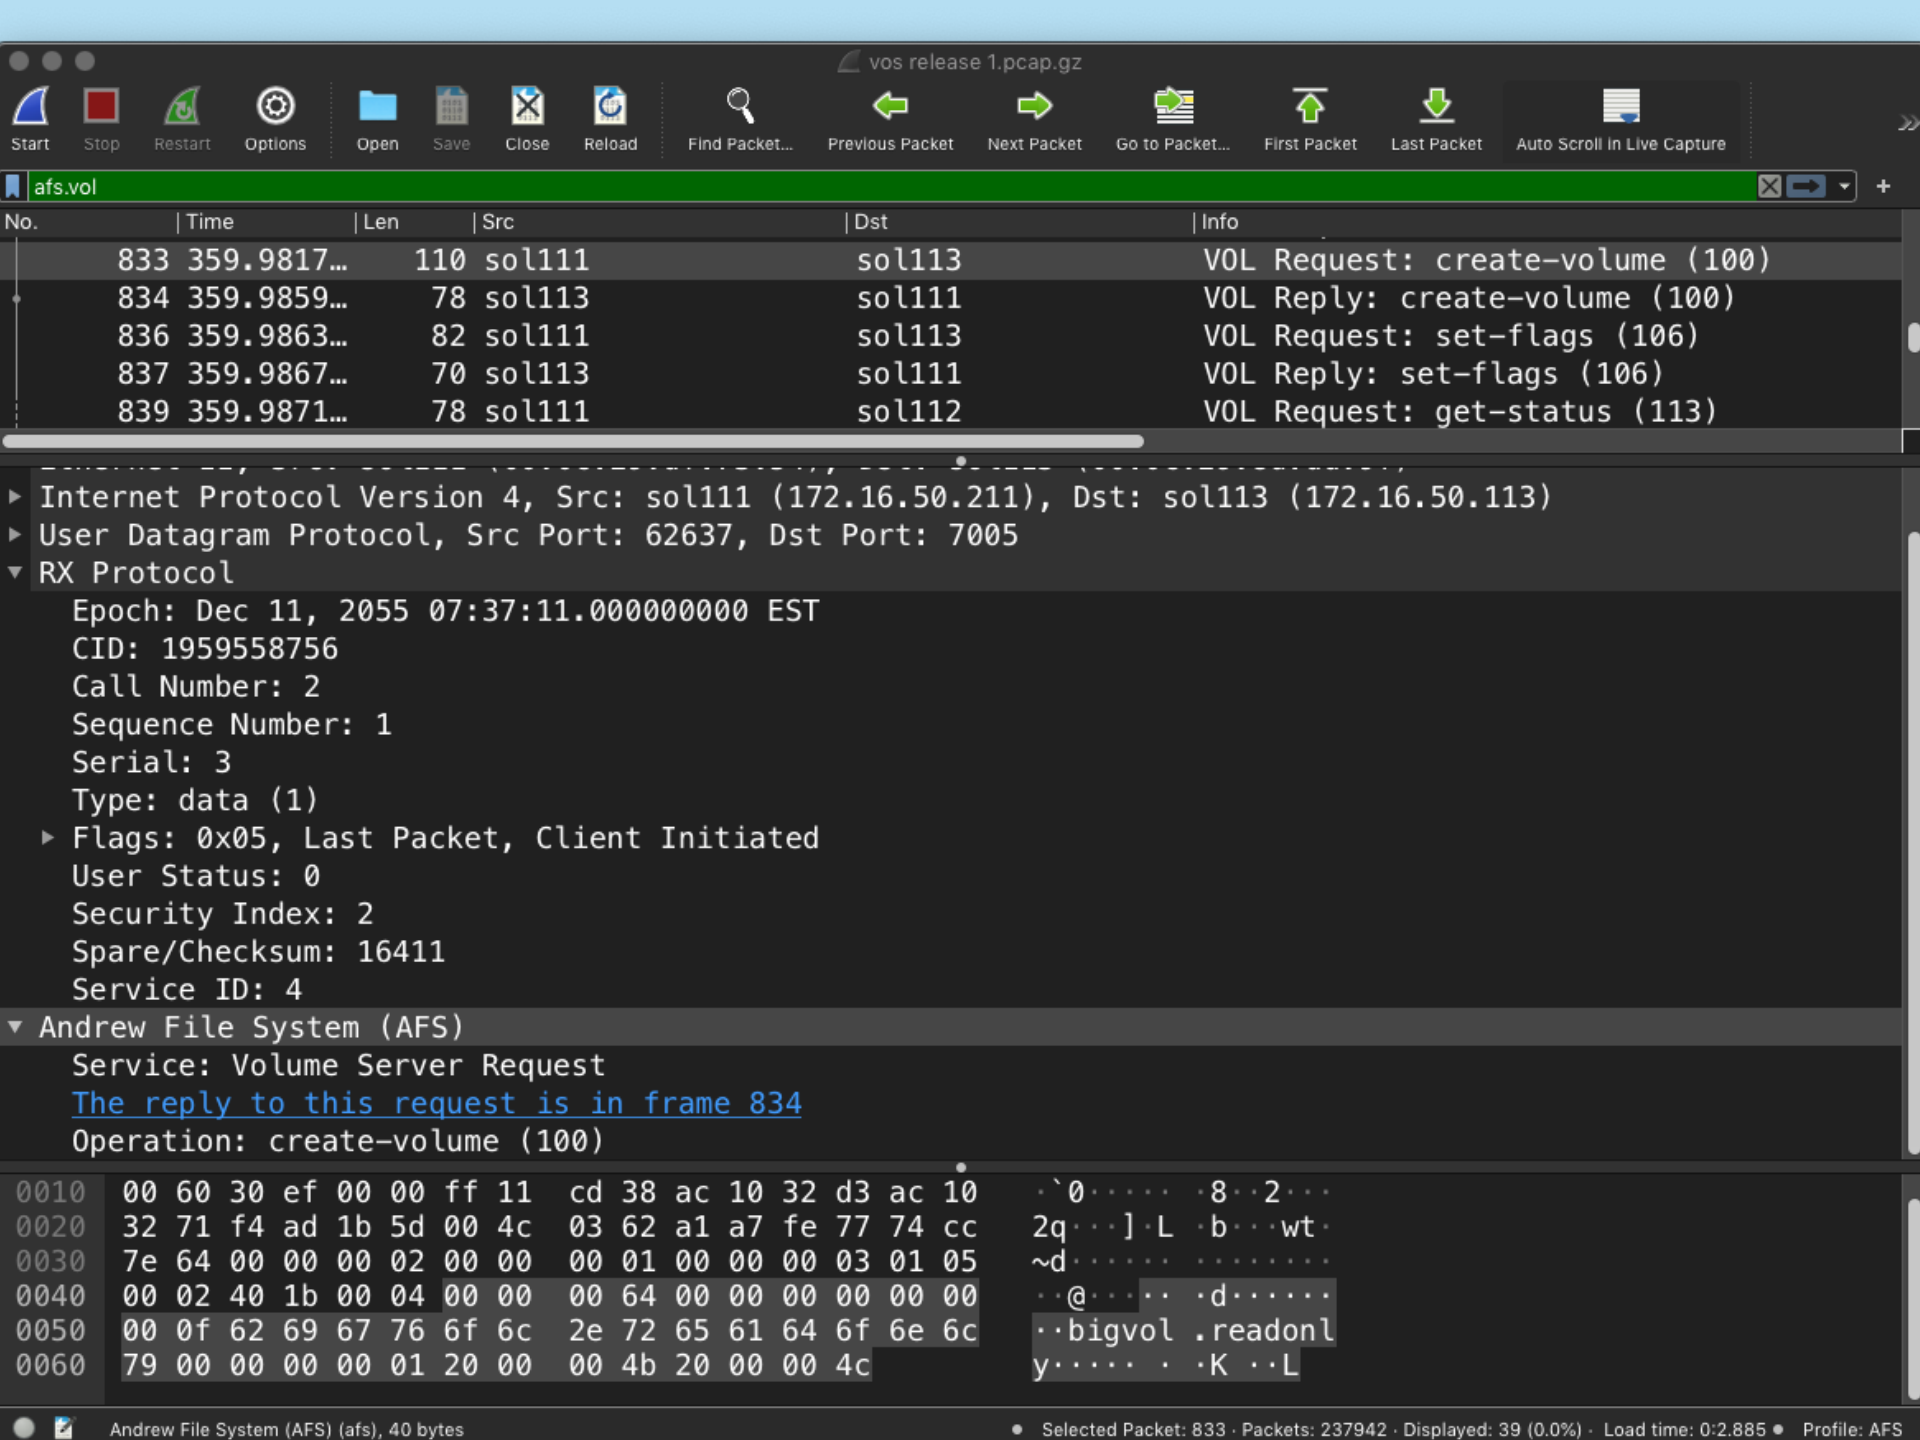
vos release 1.pcap.gz
Start
Stop
Restart
Options
Open
Save
Close
Reload
Find Packet...
Previous Packet
Next Packet
Go to Packet...
First Packet
Last Packet
Auto Scroll in Live Capture
afs.vol
No. | Time | Len | Src | Dst | Info
833 359.9817… 110 sol111 sol113 VOL Request: create-volume (100)
834 359.9859… 78 sol113 sol111 VOL Reply: create-volume (100)
836 359.9863… 82 sol111 sol113 VOL Request: set-flags (106)
837 359.9867… 70 sol113 sol111 VOL Reply: set-flags (106)
839 359.9871… 78 sol111 sol112 VOL Request: get-status (113)
Internet Protocol Version 4, Src: sol111 (172.16.50.211), Dst: sol113 (172.16.50.113)
User Datagram Protocol, Src Port: 62637, Dst Port: 7005
RX Protocol
Epoch: Dec 11, 2055 07:37:11.000000000 EST
CID: 1959558756
Call Number: 2
Sequence Number: 1
Serial: 3
Type: data (1)
Flags: 0x05, Last Packet, Client Initiated
User Status: 0
Security Index: 2
Spare/Checksum: 16411
Service ID: 4
Andrew File System (AFS)
Service: Volume Server Request
The reply to this request is in frame 834
Operation: create-volume (100)
0010 00 60 30 ef 00 00 ff 11 cd 38 ac 10 32 d3 ac 10 ·`0····· ·8··2···
0020 32 71 f4 ad 1b 5d 00 4c 03 62 a1 a7 fe 77 74 cc 2q···]·L ·b···wt·
0030 7e 64 00 00 00 02 00 00 00 01 00 00 00 03 01 05 ~d····· ·····
0040 00 02 40 1b 00 04 00 00 00 64 00 00 00 00 00 00 ··@····· ·d·····
0050 00 0f 62 69 67 76 6f 6c 2e 72 65 61 64 6f 6e 6c ··bigvol .readonl
0060 79 00 00 00 00 01 20 00 00 4b 20 00 00 4c y····· · ·K ··L
Andrew File System (AFS) (afs), 40 bytes
Selected Packet: 833 · Packets: 237942 · Displayed: 39 (0.0%) · Load time: 0:2.885
Profile: AFS

Rx dissector packet type support

SINE NOMINE ASSOCIATES

rx_packet_type	"tree" support	"info" support
1 data	pass to AFS dissector	pass to AFS dissector
2 ack	yes	yes
3 busy	yes	-
4 abort	yes	partial
5 ackall	yes	yes
6 challenge	yes	yes
7 response	yes	yes
8 debug	partial	-
9 params	partial	-
13 version	yes	yes*

Rx dissector preferences

SINE NOMINE ASSOCIATES

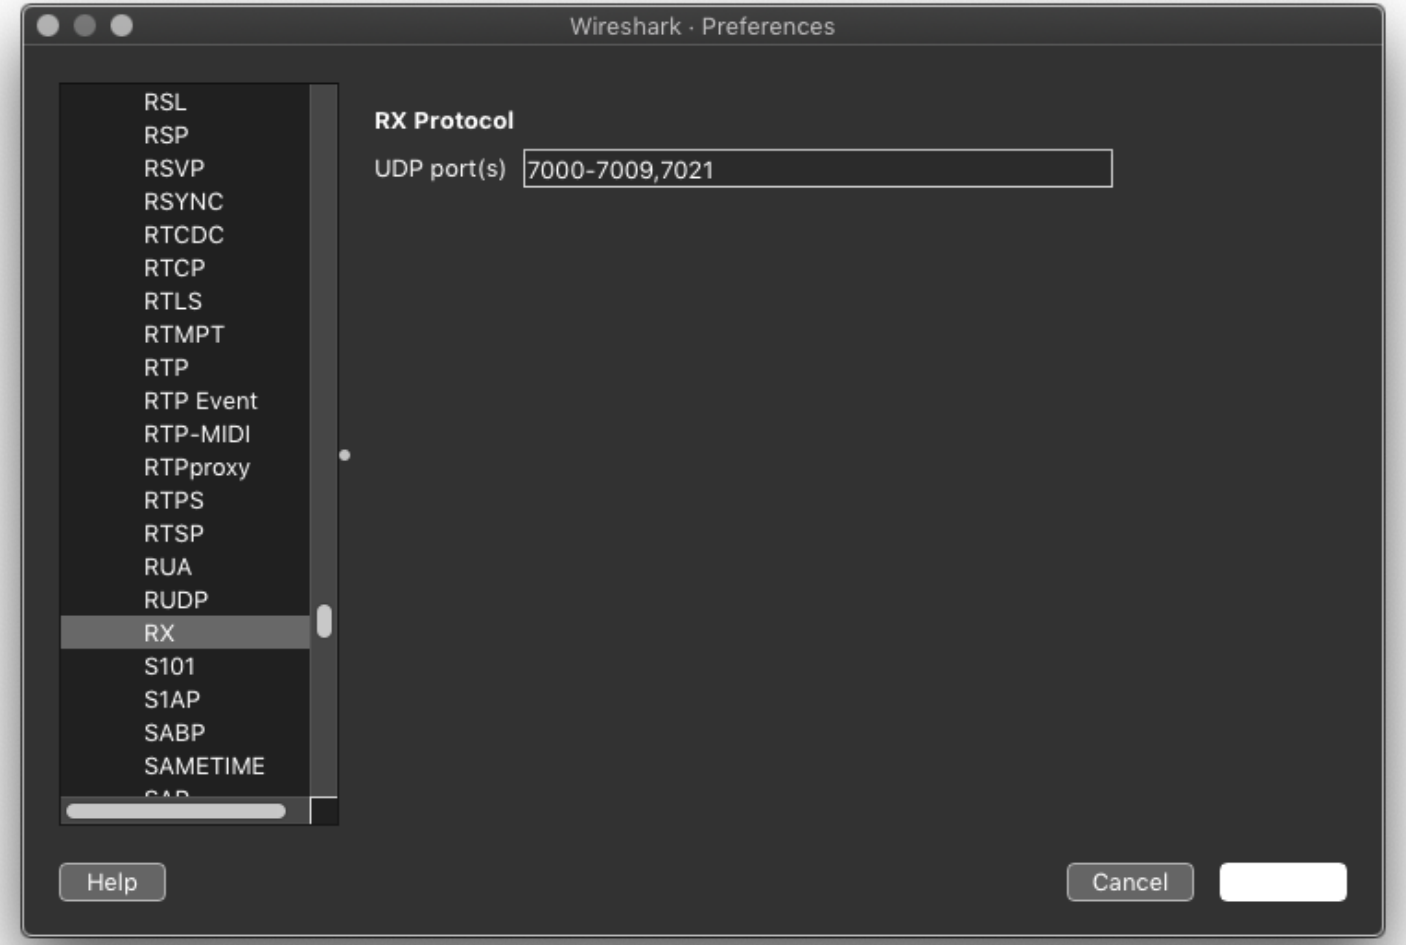

Rx dissector quirk

SINE NOMINE ASSOCIATES

- rx_header fields “out of order”

source code:

```
198     u_short serviceId;          /* service this packet is directed _to_ */
199     /* This spare is now used for packet header checkksum.  see
200      * rxi_ReceiveDataPacket and packet cksum macros above for details. */
201     u_short spare;
```

Wireshark dissector:

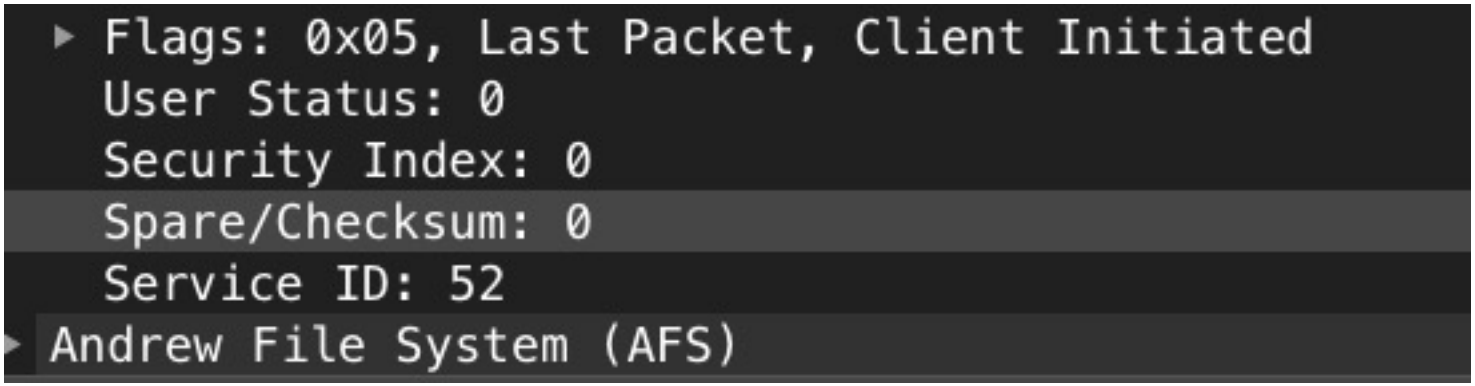

AFS dissector support

SINE NOMINE ASSOCIATES

- has mostly complete support for all RPCs, for all but one service
  - However, many RPC requests and replies are merely identified but their payloads are not dissected
- RXSTATS_* (service 409) is unknown

AFS dissector preferences

SINE NOMINE ASSOCIATES

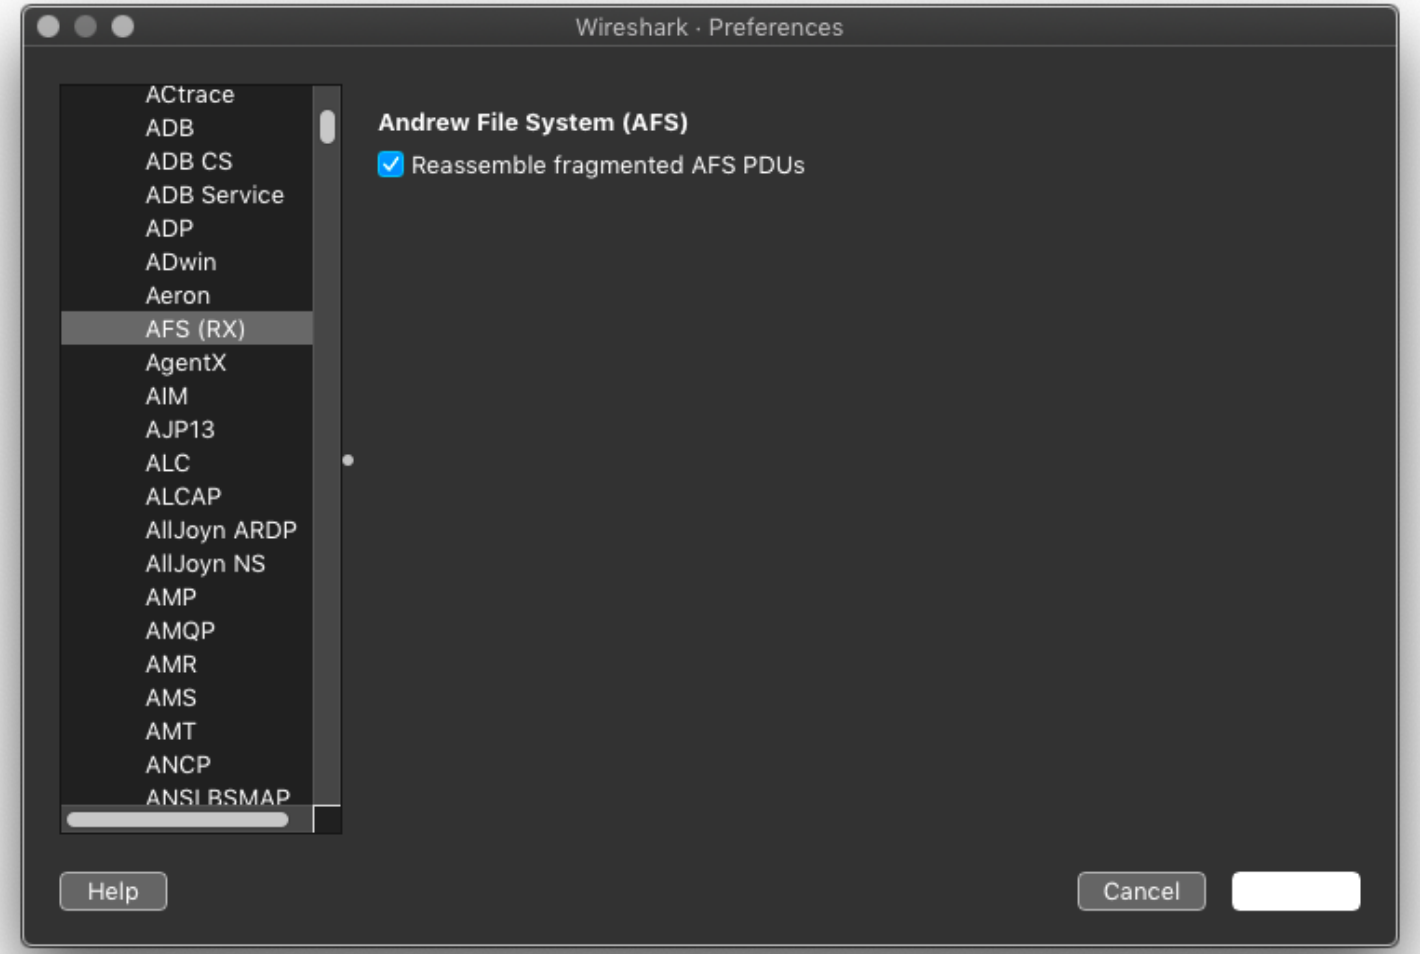

Packet reassembly

SINE NOMINE ASSOCIATES

- When enabled, this automatically consolidates multi-packet objects to allow detailed dissection
- pros:
  - avoid spurious “malformed packet” errors on some multi-packet RPCs (e.g. RXAFS_InlineBulkStatus)
- cons:
  - doesn’t scale well for large pcaps (or large RPCs, e.g. AFSVolRestore)

AFS dissector quirks

SINE NOMINE ASSOCIATES

- RXAFS_TellMeAboutYourself (65538) aka TMAY is reported as “get-capabilities”
- The dissector’s heuristic for identifying the Rx service may sometimes be confused by ephemeral ports, causing “Unknown” for the service or RPC name.

Protocol preference tips

SINE NOMINE ASSOCIATES

- the hard way:
  - Wireshark -> Preferences -> <prefs dialog box>
    - scroll through the giant list of protocols to find yours
- the expert shortcut:
    - in the packet list pane, select a packet
    - right-click -> Protocol Preferences -> <short protocol list>
    - select your protocol -> <context menu>
      - Open Andrew File System (AFS) preferences
      - [ ] Reassemble fragmented AFS PDUs
      - Disable AFS(RX)

Coloring rules

SINE NOMINE ASSOCIATES

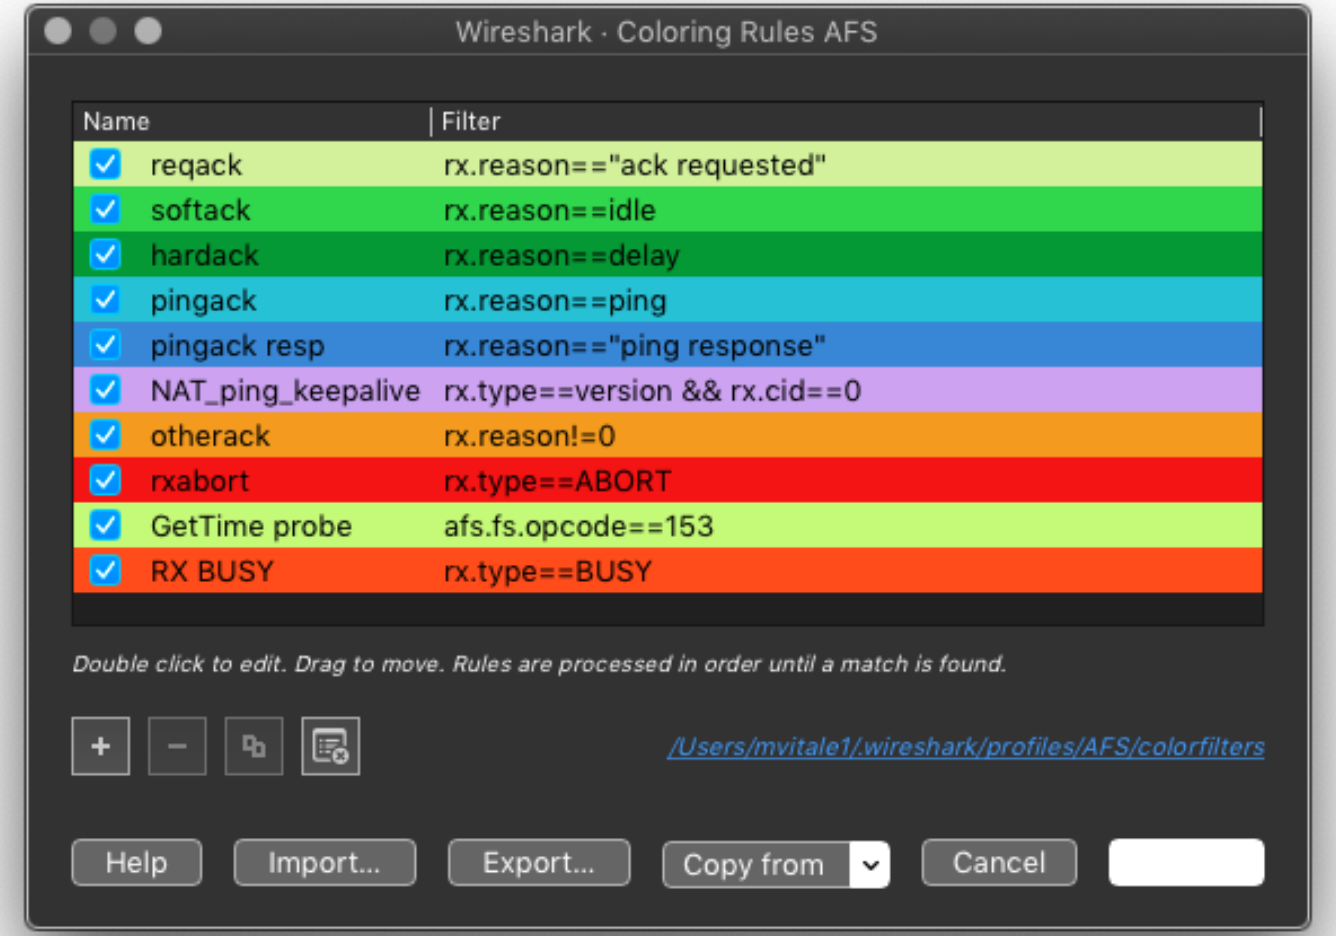

Performance tips

SINE NOMINE ASSOCIATES

- disable coloring rules
- disable reassembly
- disable unneeded name resolutions
- break up large packet traces (over 100MB) into smaller files
  - editcap or SplitCap

Recent improvements

commit	GA	comments
rx: display rx-ack reason string	3.4.0	adds ack reason to “info”
rx: decode version packets	3.4.0	distinguish –version from NAT PING
afs: add some “new” RPCs	3.4.0	RXAFS_CallBackRxConnAddr RXAFS_GetStatistics64 RXAFS_Link PR_ListSuperGroups AFSVolDumpV2 AFSVolPartitionInfo64
afs: make defragment / reassembly configurable	3.4.0	add AFS preference to enable/disable reassembly
afs: fix backup & butc RPC confusion	3.4.0	BUTC_* is now on correct port BACKUP_* support added
afs: correctly calculate padding for strings	2.6.15 3.0.10 3.2.3	bug fix for spurious “Malformed packet” for xdr strings (e.g. RXAFS_CreateFile filename)

Note: 2.6.16, 3.0.10, 3.2.2 GA Apr 09 2020; 3.4.0 GA Oct 29 2020

To-do list

SINE NOMINE ASSOCIATES

- fix VERSION packet bug
- ABORT packet improvements (including special case for VOTE_Beacon replies)
- support for DEBUG and BUSY packets
- distinguish ping ACK from MTU ping
- stop decoding rx epoch as an actual date

demo

Questions and discussion