SNIA DEVELOPER CONFERENCE

SDC 24

BY Developers FOR Developers

September 16-18, 2024
Santa Clara, CA

A New Era in Networking:

Insights from use of the Experimental Linux Kernel QUIC Driver with SMB3.1.1

Presented by: Steven French

Principal Software Engineer, Azure Storage, Microsoft

- This work represents the views of the author(s) and does not necessarily reflect the views of Microsoft Corporation
- Linux is a registered trademark of Linus Torvalds.
- Other company, product, and service names may be trademarks or service marks of others.

Who am I?

- Steve French smfrench@gmail.com
- Author and maintainer of Linux cifs vfs (for accessing Samba, Azure, Windows and various SMB3/CIFS based NAS appliances)
- Co-maintainer of the kernel server (ksmbd)
- Member of the Samba team (co-creator of the “net” utility)
- coauthor of SNIA CIFS Technical Reference, former SNIA CIFS Working Group chair
- Principal Software Engineer, Azure Storage: Microsoft

Outline

- Overview of QUIC
- Review of Xin Long's status update on QUIC in Linux kernel
- Discussion of cifs.ko changes needed
- Coming soon … what to look forward to
- Testing discussion

Why is QUIC important?
We could ask ChatGPT ...

ChatGPT

QUIC (Quick UDP Internet Connections) offers several advantages over TCP as the transport protocol for SMB3.1.1 (Server Message Block version 3.1.1), particularly in terms of performance, security, and reliability.

1. **Reduced Latency**: QUIC is designed to minimize latency by reducing the number of round trips required for establishing a connection. It achieves this by combining the initial handshake and encryption setup into a single round trip, resulting in faster connection establishment compared to TCP.
2. **Multiplexing**: QUIC supports multiplexing of multiple streams within a single connection. This enables concurrent transmission of multiple data streams over a single connection, reducing the overall latency and improving throughput, especially in scenarios where multiple files are being transferred simultaneously in SMB.
3. **Connection Migration**: QUIC allows for seamless connection migration between different network interfaces or IP addresses without disrupting ongoing data transmission. This is particularly advantageous for mobile devices or scenarios where network conditions change frequently, ensuring continuous connectivity and minimizing downtime.

Why QUIC (continued) ...

4. **Improved Security**: QUIC incorporates modern encryption protocols such as TLS 1.3 by default, providing robust encryption for data transmission and enhancing security compared to older versions of TLS used with TCP. Additionally, QUIC mitigates certain vulnerabilities associated with TCP, such as SYN flooding attacks, by utilizing UDP encapsulation.
5. **Adaptive Congestion Control**: QUIC includes built-in mechanisms for adaptive congestion control, which dynamically adjusts transmission parameters based on network conditions such as packet loss and latency. This helps optimize network utilization and ensures efficient use of available bandwidth, leading to improved performance and reliability.
6. **Faster Handshake**: QUIC's connection establishment process is optimized for efficiency, requiring fewer round trips compared to TCP's three-way handshake. This results in quicker initiation of data transfer sessions, reducing overall latency and improving user experience, especially for short-lived connections such as those typical in SMB file transfers.

By leveraging these advantages, using QUIC as the transport protocol for SMB3.1.1 can significantly enhance the performance, security, and reliability of file transfer operations, particularly in modern network environments where low latency, multiplexing, and adaptive congestion control are crucial requirements.

But it misses some of the obvious ...

- It solves the “port 445” problem (e.g. it can be hard form me to demo SMB3.1.1 when Spectrum blocks port 445). SMB3.1.1 mounts to anywhere … save your files more easily into the cloud ...
- It also could allow a mix of machine credentials (for QUIC encryption) and user credentials ... which may be helpful for some workloads
- If QUIC is also supported in unencrypted mode (as some QUIC drivers allow for testing) it will likely be faster than TCP (if encrypting it is still a good alternative to requiring SMB3.1.1 GCM encryption)
- Will there be cases where “double encrypting” (at QUIC layer and also for SMB3.1.1 layer) will be valuable?

What about alternatives for SMB3.1.1 to use to TCP?

- RDMA is great (smbdirect) for SMB3.1.1 but not available in all configurations
- SMB3.1.1 is fairly “network protocol agnostic” and has historically layered on various protocols not just TCP and RDMA
  - E.g. I did a demo with some HP developers years ago of cifs.ko and Samba on Linux using SCTP, and older servers supported NetBIOS
- What about newer protocols like Stanford’s Homa (see talk yesterday)?
  - Complementing TCP with Homa, Stanford’s Reliable, Rapid Request-Response Protocol | SDC 2024 (sniadeveloper.org)
- Additional experiments with different transport protocols with SMB3.1.1 would be welcome

But it still would be awesome due to ...

- No “head of line blocking”
- Better congestion control
- Faster session establishment
- Reduced retransmissions (and Forward Error Correction)
- Large percentage of internet traffic is already over QUIC (and growing)

So why don’t we have QUIC yet on Linux?

- Well … we do in userspace, multiple drivers
  - e.g. https://github.com/microsoft/msquic is well tested and works on Linux
- But unlike Windows, we don’t have a kernel driver
- And many of the open source QUIC drivers wouldn’t port well to Linux kernel (even if only porting minimal required function) and don’t follow kernel coding style
- Last year (June 24th) – Xin Long started a kernel project which is looking VERY promising

Recent Progress

Additional testing will be going on this week (e.g. at the SMB3.1.1 test event colocated with SNIA SDC)
Has been submitted to netdev for considerations for 6.12-rc. See:
[PATCH net-next 0/5] net: implement the QUIC protocol in linux kernel - Xin Long

Latest patch set getting feedback on netdev mailing list

Performance being investigated

May be addressable by optimizations like use of io_uring

Some initial theories about cause:

- Absence of Generic Segmentation Offload (GSO) for QUIC
- Additional data copy on the transmission (TX) path
- Extra encryption required for header protection in QUIC
- Longer header length for the stream data in QUIC

What about userspace code using it?

Samba team has done some initial experiments (so far successful) with quic and smbd userspace file server (thank you Metze!)
See the talk yesterday by Ralph Bohme (Samba team) which e.g. mentioned plans and funding to work on adding SMB3.1.1 quic support for Samba server

How do you find out if server supports QUIC?

An example of some ways to do it: It's Over 9000: Analyzing Early QUIC Deployments with the Standardization on the Horizon (zirngibl.github.io)
I also discussed with Xin the idea of writing some test tools to make it easier to discover supported targets

Is this just for Windows?

- Ned's slides and learn.microsoft.com explain more about Windows SMB3.1.1 over QUIC, but Visuality systems also supports SMB3.1.1 over QUIC
- QUIC will be a great addition to cifs.ko, ksmbd.ko and Samba server and tools

Windows QUIC SMB3.1.1 config

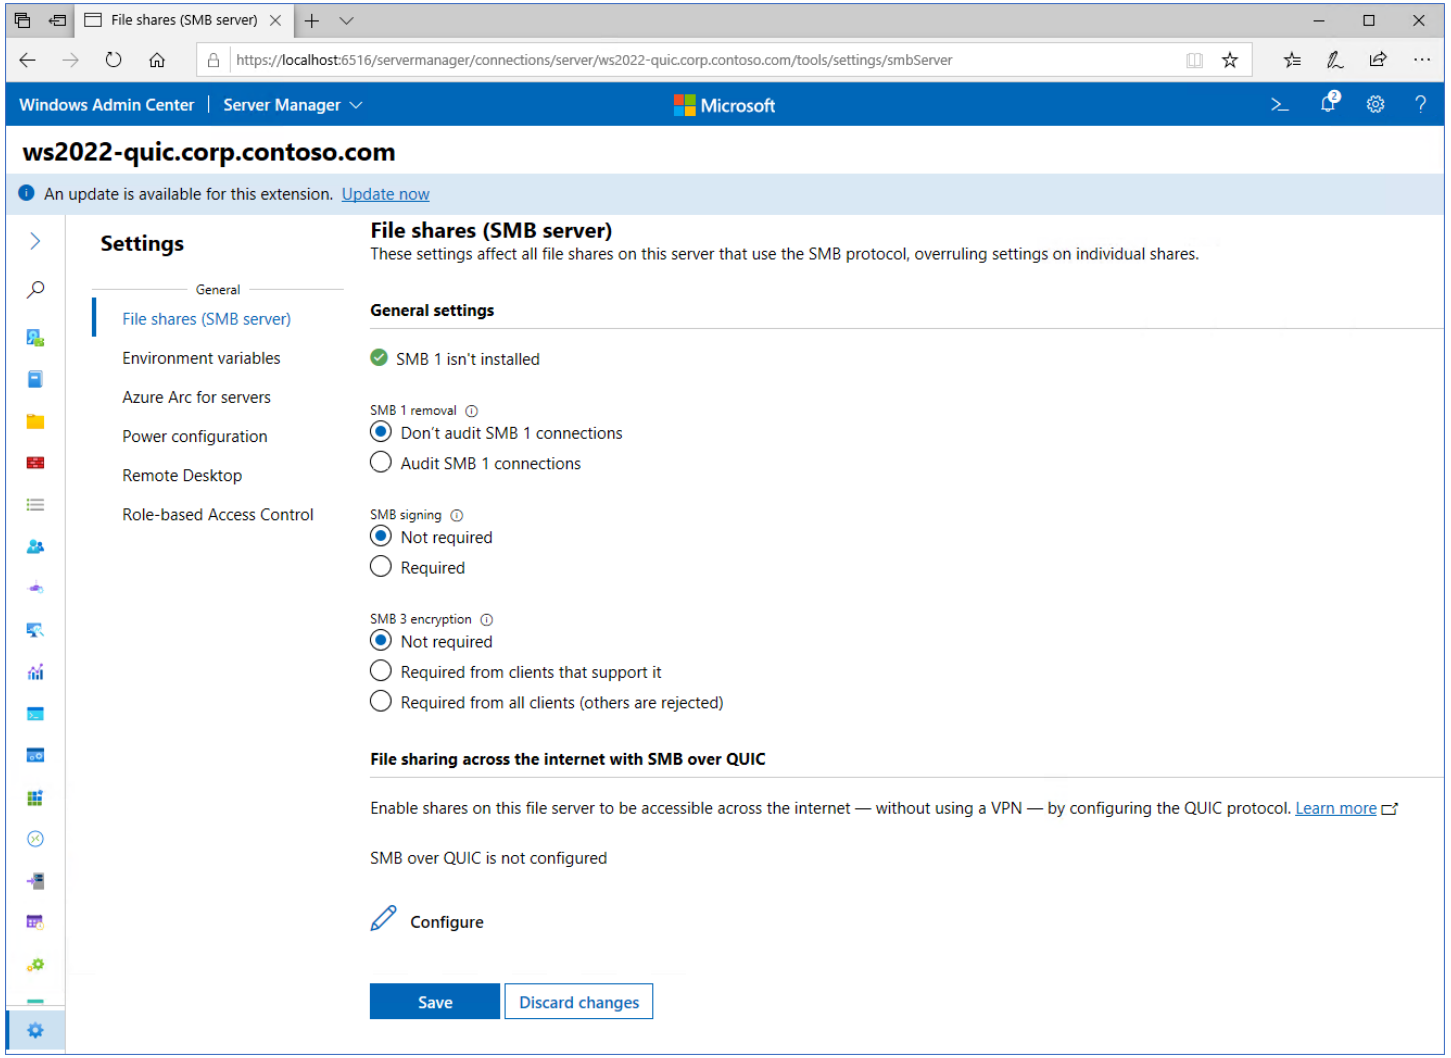

Example Windows mount from Ned's post

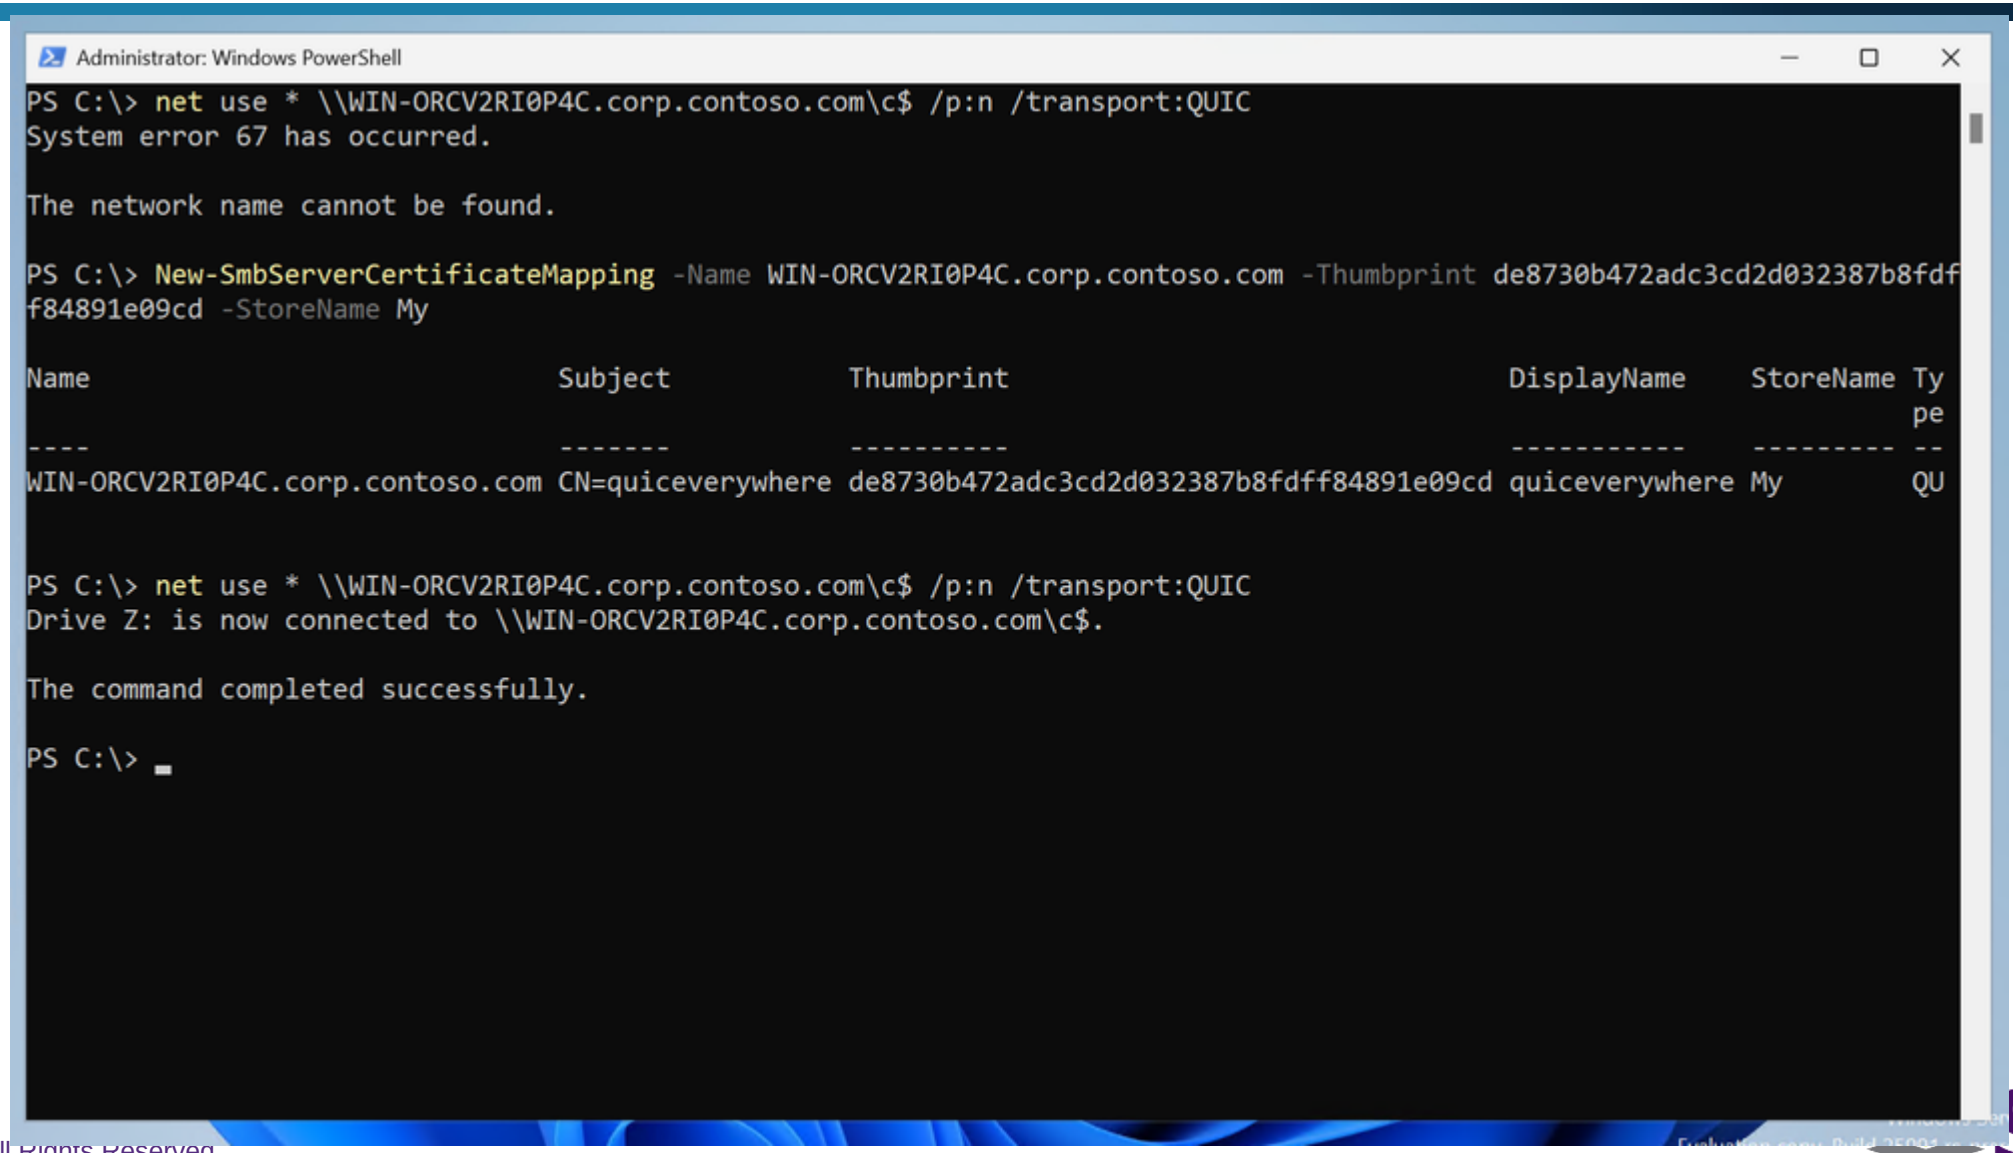

SDC 24

Diagram showing SMB3.1.1 over QUIC (for Windows)

- (Thanks to Obaid)

SMB/QUIC: Client

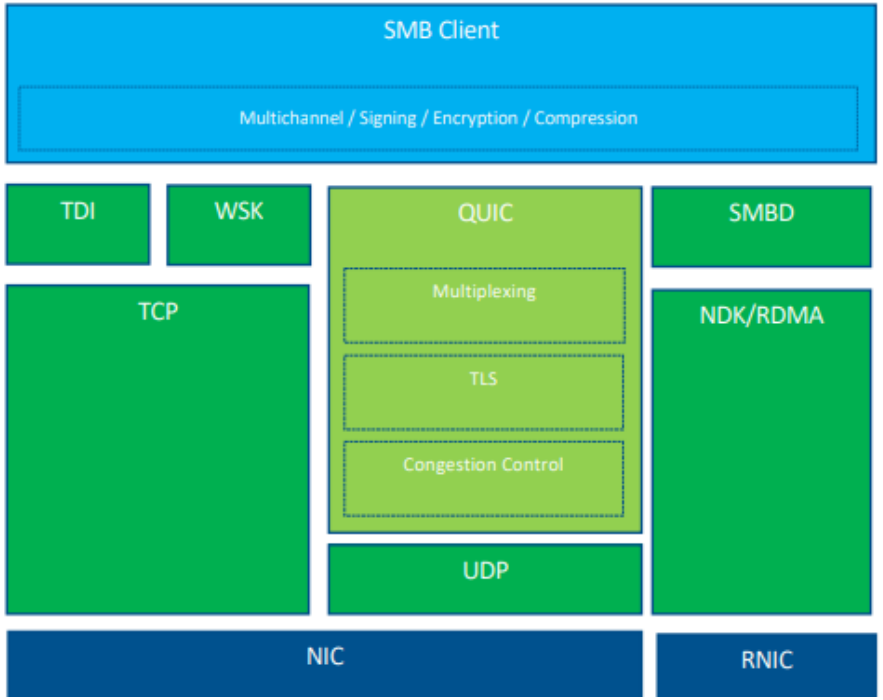

1. *Client opens \\ServerName\Share\foo.tst*
2. *Client resolves ServerName using DNS*
3. *Client attempts TCP and QUIC simultaneously**
4. *Client will start using whichever connects first*
5. *Client's multichannel will negotiate interfaces with server and will select most optimal protocols*
6. *Client sends SMB messages*

Instructions on setting up SMB3.1.1 QUIC in Windows

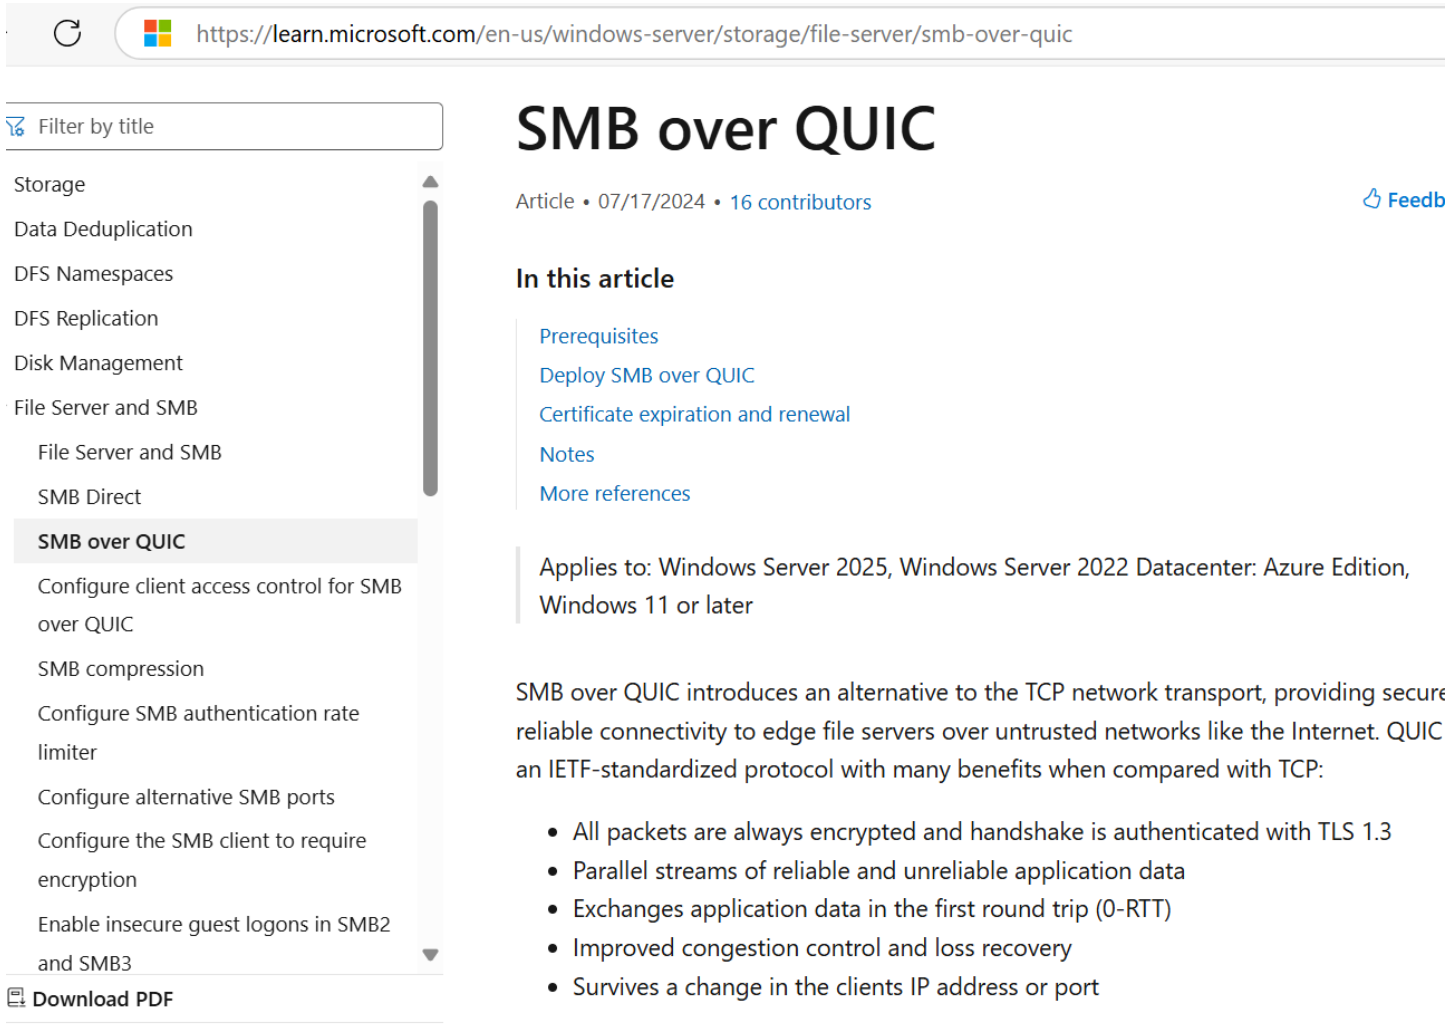

The New Kernel QUIC driver

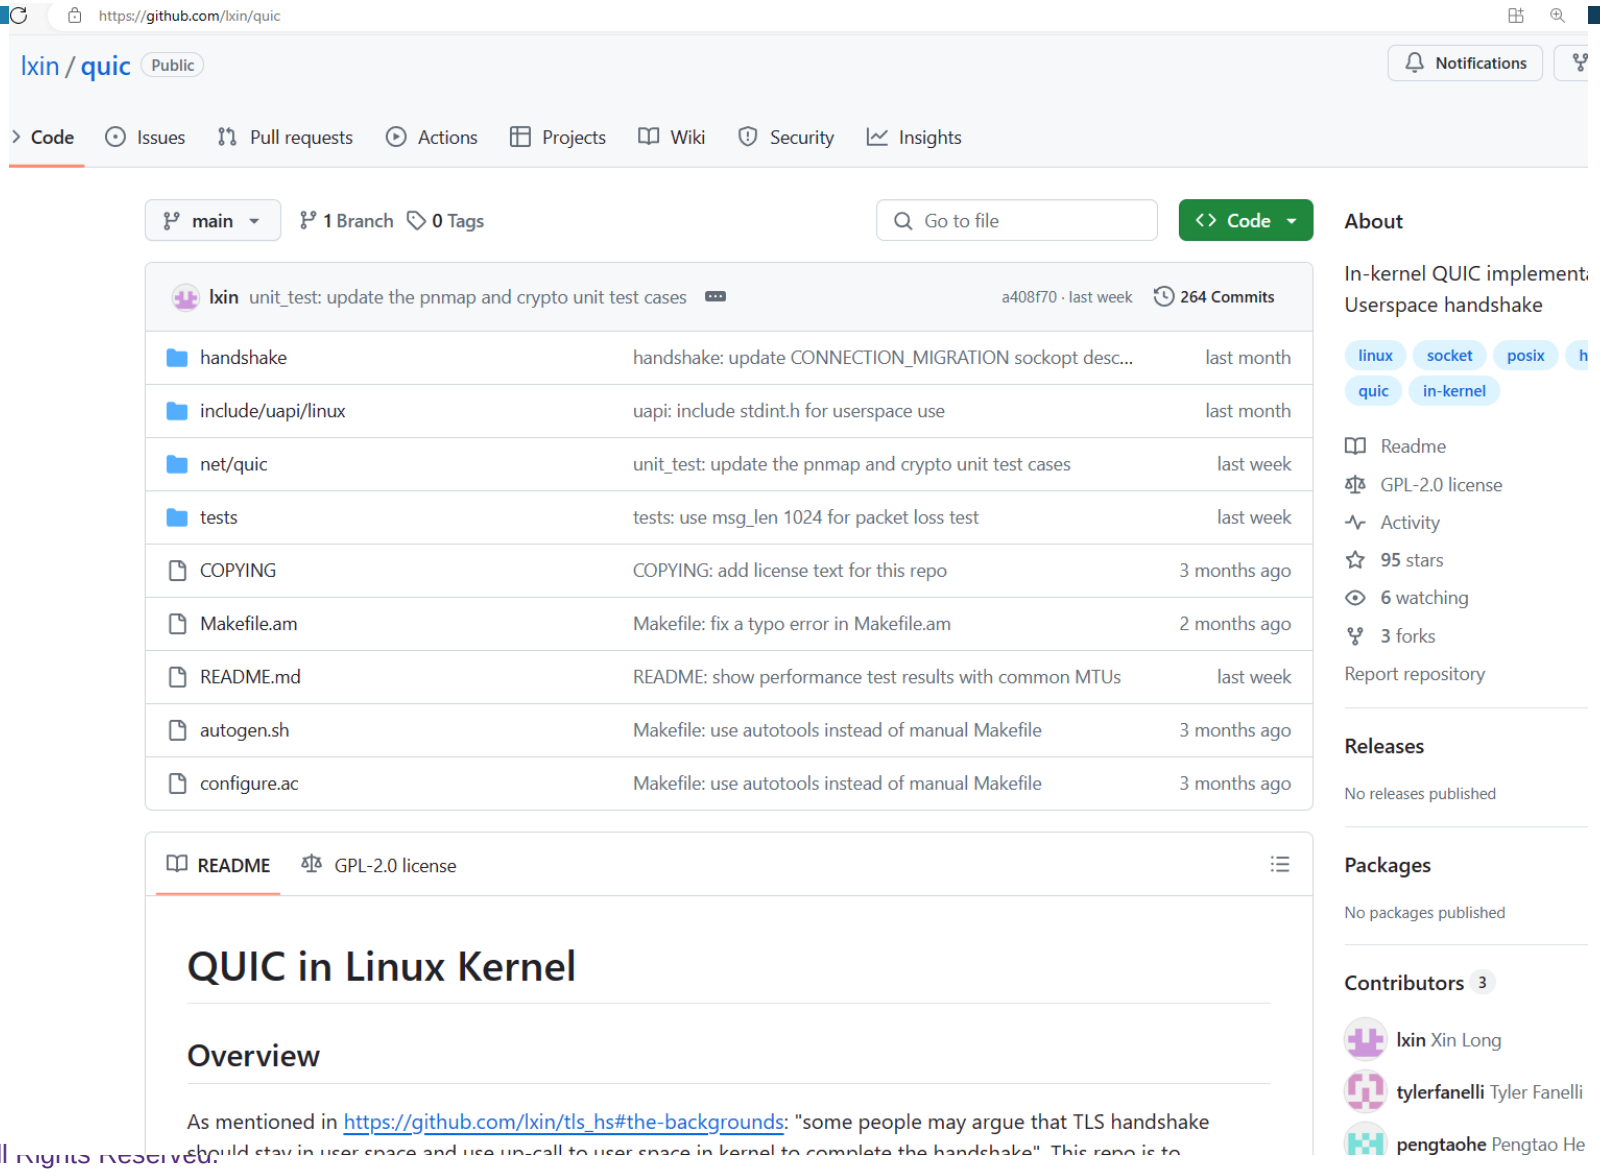

New Kernel QUIC driver

- What I am impressed with so far:
  - Kernel style and code looks good
  - Very easy to build the kernel driver and run their functional tests
  - Reasonable size (20.4KLOC about ½ of which is kernel code)
  - Checking this week's version: the size of the kernel pieces was 12.5KLOC in the new net/quic kernel directory. Reasonable
  - Good set of functional tests included (it passed all of them when I tried), and even last week a test was updated
  - Maintainers have been responsive to emails and questions
  - Unencrypted QUIC faster than TCP

Is QUIC faster than TCP (when encryption disabled)

with mtu = 1500, size 16384 and 65536, QUIC got better performance than TCP in tests Xin tried today

UNIT	size:1024		size:4096		size:16384		size:65536	
Gbits/sec	QUIC	TCP	QUIC	TCP	QUIC	TCP	QUIC	TCP
mtu:1500	2.17	2.49	3.59	8.36	6.09	15.1	6.92	16.2
no GSO		2.50		4.12		4.86		5.04
mtu:9000	2.47	2.54	7.66	7.97	14.7	20.3	19.1	31.3
no GSO		2.51		8.34		18.3		22.3

New Kernel QUIC driver

- Easy to build and install (see networking/quic.rst)

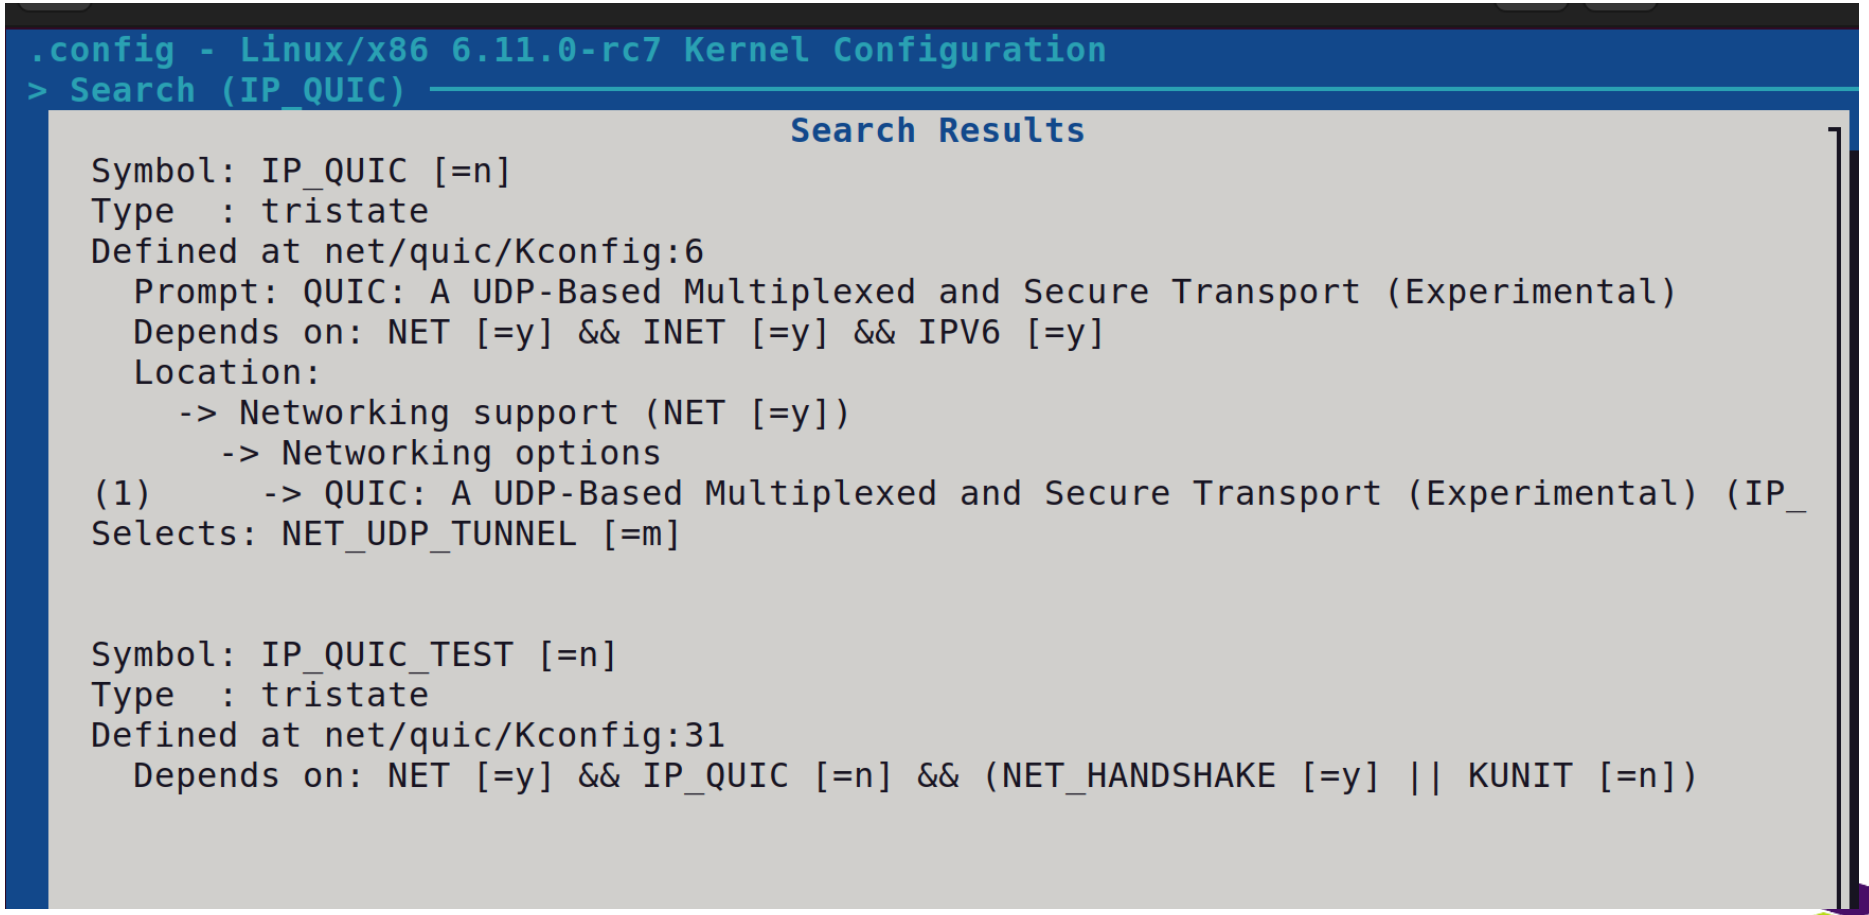

Kernel QUIC driver

- Easy to build and install (see quic/README.md)

```
root@smfrench-ThinkPad-P52:/home/smfrench/smb3-kernel# lsmod | grep quic
quic                   192512  0
ip6_udp_tunnel          16384  1 quic
udp_tunnel              32768  1 quic
root@smfrench-ThinkPad-P52:/home/smfrench/smb3-kernel# modinfo quic
filename:       /lib/modules/6.9.0-060900rc4-generic/extra/quic.ko
license:        GPL
description:    Support for the QUIC protocol (RFC9000)
author:         Xin Long <lucien.xin@gmail.com>
alias:          net-pf-10-proto-261
alias:          net-pf-2-proto-261
srcversion:     DF18010B3F937CB40FA73D2
depends:        udp_tunnel,ip6_udp_tunnel
retpoline:      Y
name:           quic
vermagic:       6.9.0-060900rc4-generic SMP preempt mod_unload modversions
```


Overview of the proposed kernel QUIC driver

Next 15 slides courtesy of Xin Long <lucien.xin@gmail.com>
See https://github.com/lxin/quic
Thanks to Xin and also for contributors to the code e.g.
Pengtao He <hepengtao@xiomi.com> and
Tyler Fanelli <tfanelli@redhat.com>

What we'll discuss today

- Background
- Implementation
- Usage
- Next Step

Background

- What is QUIC
- Why QUIC in Kernel is Needed
- In Kernel QUIC Evolution

What is QUIC

RFC9000 - A UDP-Based Multiplexed and Secure Transport

- UDP based
  - Connection Migration
- Multistreaming
- Secured by TLS 1.3
  - Rekeying/Session Resumption/0-RTT
- Transport Protocol
  - Flow Control/Congestion Control

Other RFCs:

- RFC8999 - Version Negotiation
- RFC9001 - Using TLS to Secure
- RFC9002- Congestion Control
- RFC9221 - Unreliable Datagram
- RFC9287 - Greasing the QUIC Bit
- RFC9368 - Version Negotiation 2
- RFC9369 - QUIC Version 2

Why QUIC in Kernel is Needed

- Kernel Consumers
  - SMB, NFS ...
- Socket APIs
  - listen/accept/connect/send/recv/close/get/setsockopt ...
- Avoiding Data Copies
  - zero-copy via sendfile()
- Offloading in NICs
  - Similar to TLS/IPsec offloading
- Less Interoperability Issues
  - Too many Userland QUIC implementations

In Kernel QUIC Evolution

In-Kernel QUIC != Crypto Offloading for Userland QUIC

- net: support QUIC crypto (a patchset posted on 2022)
- Offloading Encryption to QUIC Enabled NICs (LPC2023)
- TLS 1.3 Handshake in kernel/lib
- TLS Handshake netlink in kernel/net (upstream kernel)
- Long/Handshake Packets processed by libngtcp2

Implementation

- Idea
- Handshake Architecture
- User Data Architecture
- Socket Process

Idea

In-kernel QUIC with Userspace Handshake

- Userspace: Only raw TLS Handshake Messages creating and processing
- Kernel: Create IPPROTO_QUIC type socket running over UDP tunnels
  - No protocol number needed from IANA, like IPPROTO_MPTCP
  - Reasons why ULP layer is not used:
    - Connection Migration
    - Common Socket APIs
    - Transport Protocol
- Kernel consumers use handshake netlink to request a handshake

Handshake Architecture

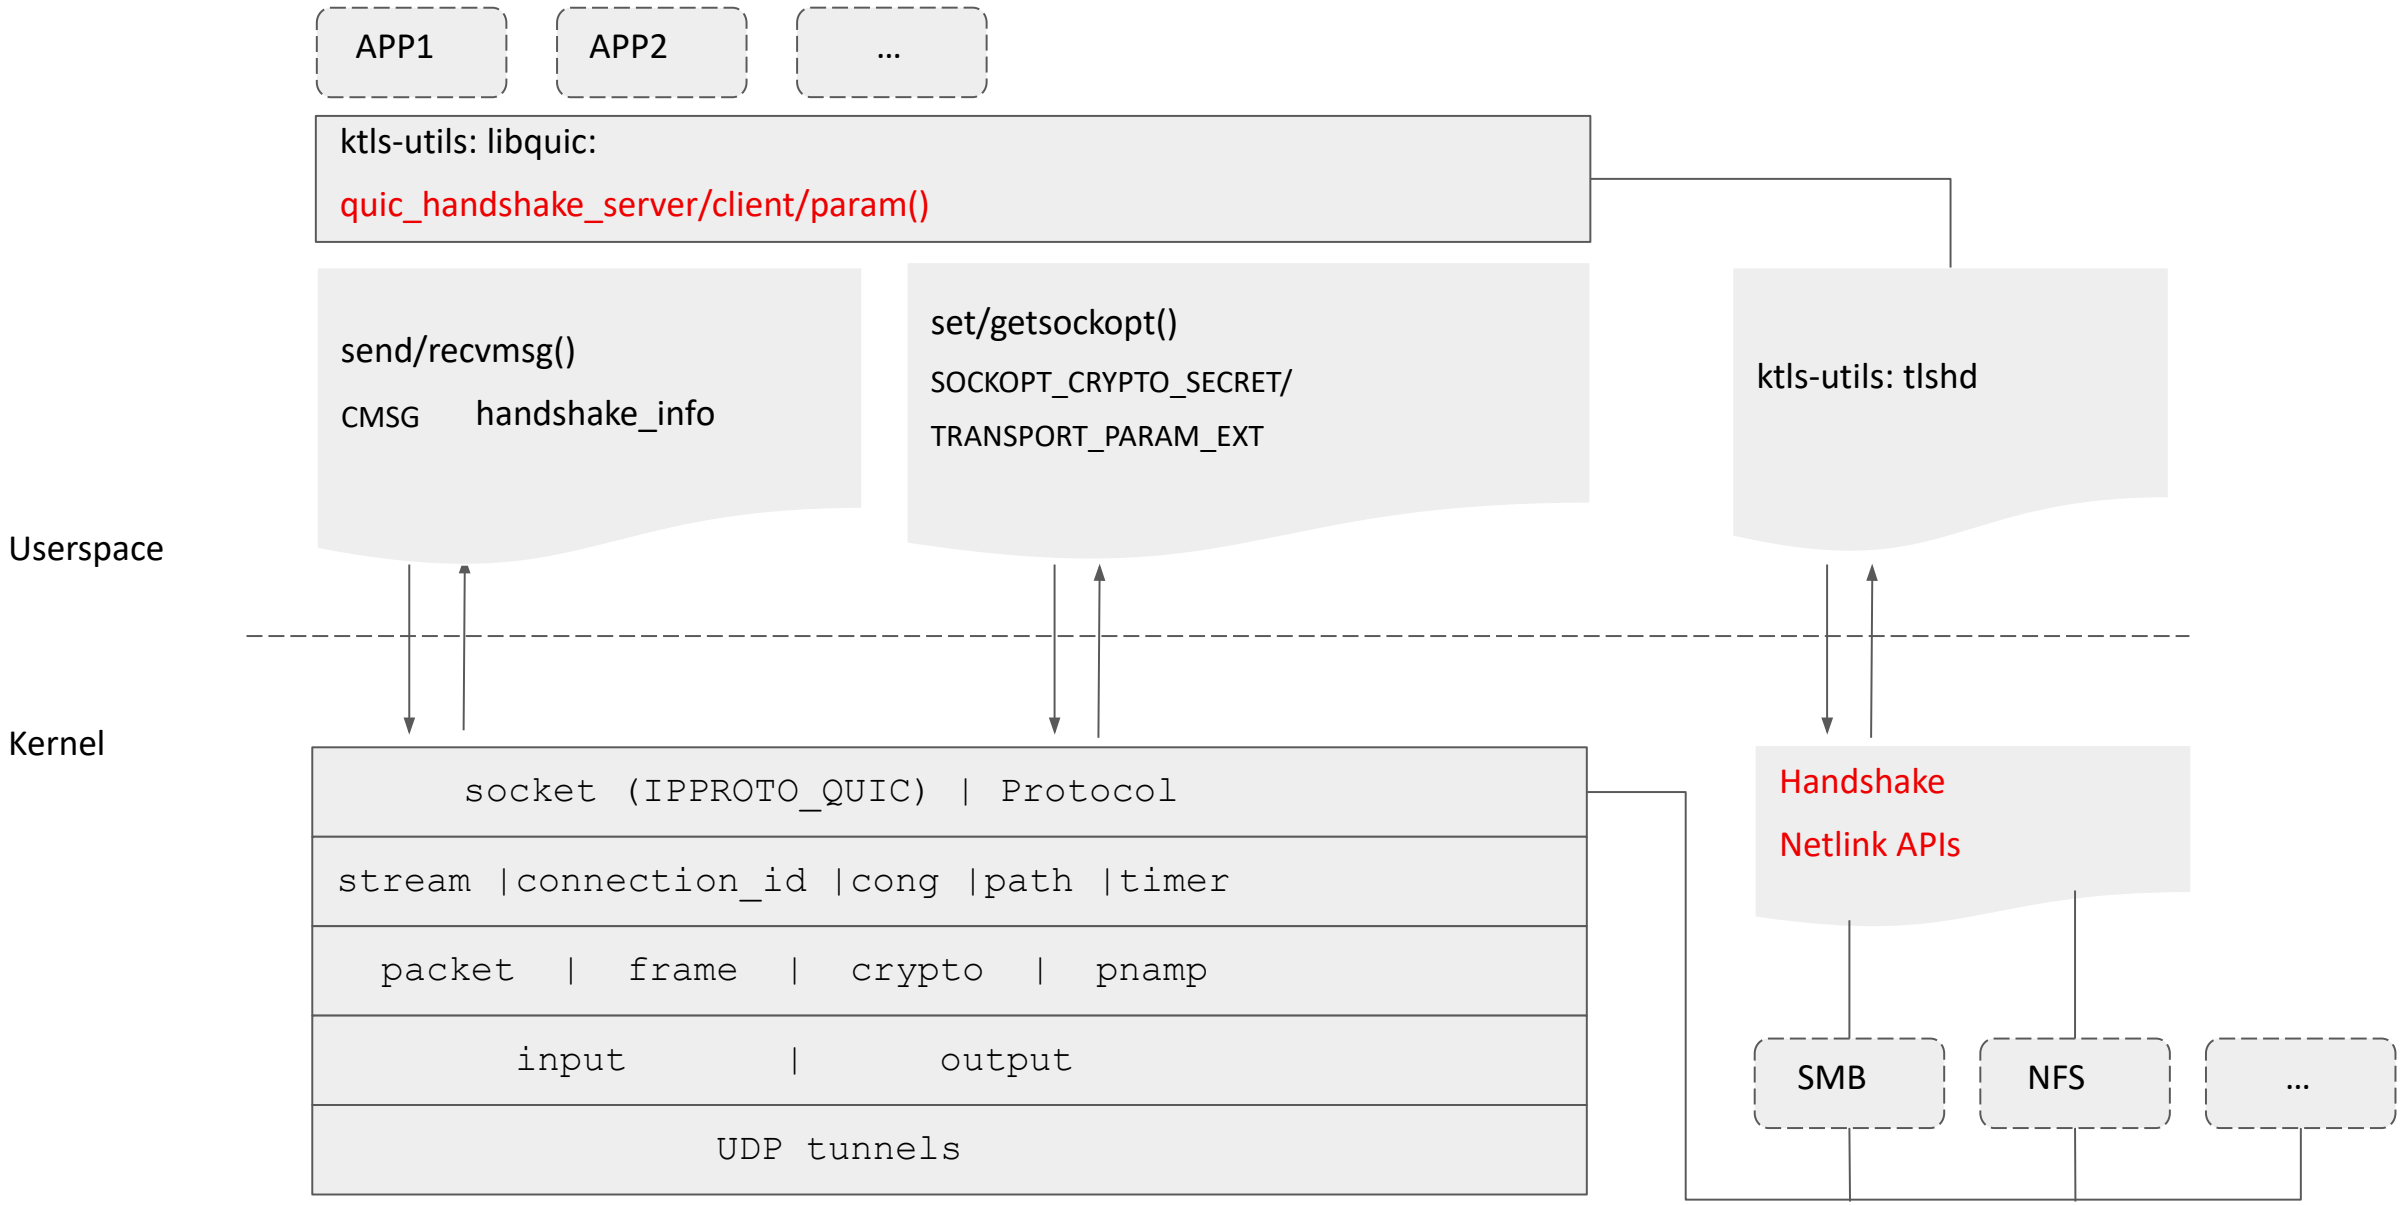

User Data Architecture

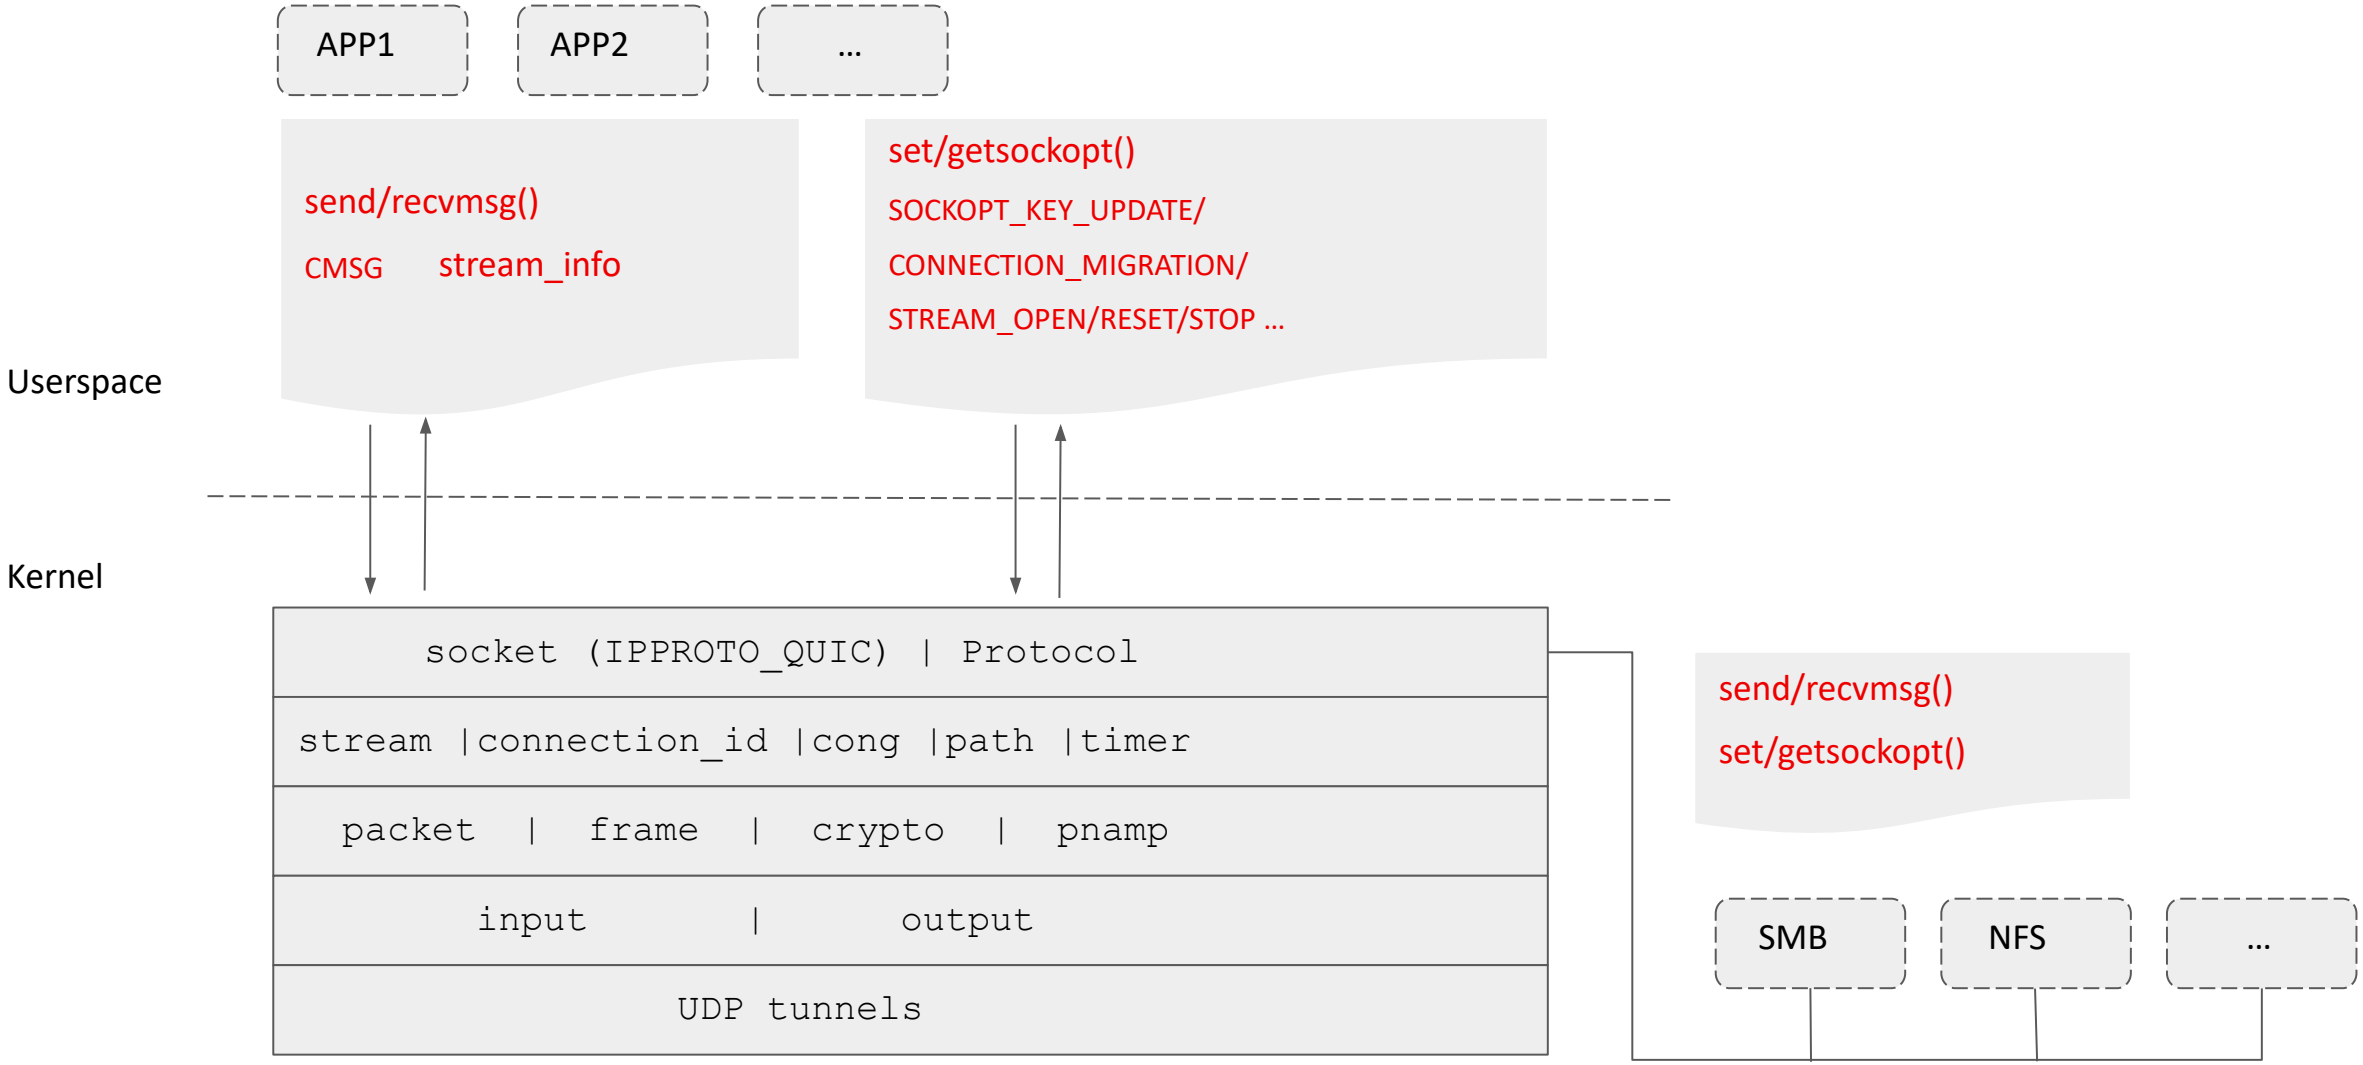

Socket Process

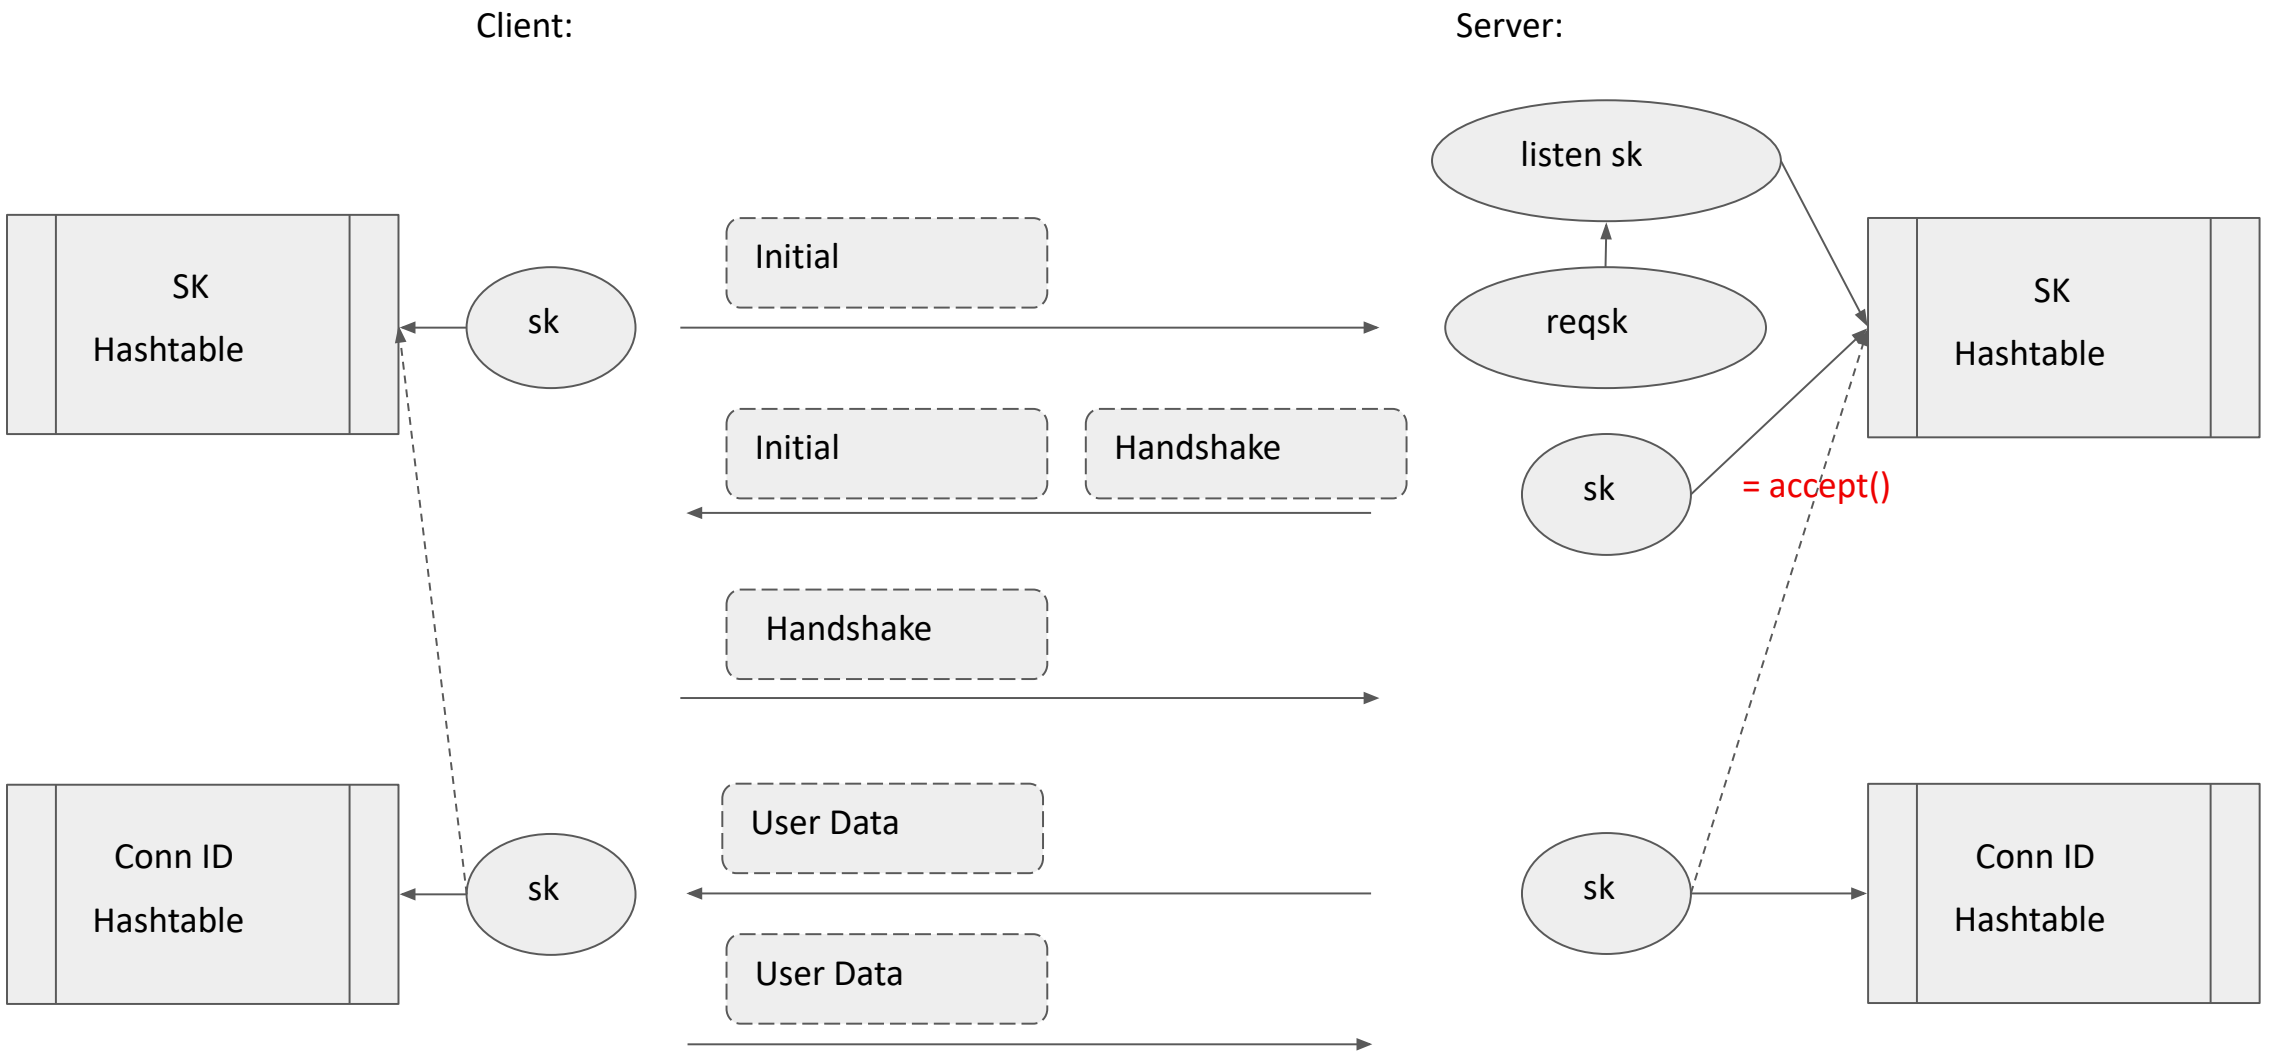

Usage

- From Userspace
- From Kernel
- Advanced APIs

From Userspace

- Client

```
sockfd = socket(IPPROTO_QUIC)
bind(sockfd)


connect(sockfd)
quic_client_handshake(sockfd)




sendmsg(sockfd)
close(sockfd)
```

- Server

```
listenfd = socket(IPPROTO_QUIC)
bind(listenfd)
listen(listenfd)




sockfd = accept(listenfd)
quic_server_handshake(sockfd, cert)

recvmsg(sockfd)
close(sockfd)
close(listenfd)
```

Sample: https://github.com/lxin/quic/blob/main/tests/sample_test.c

From Kernel

- Client

```
__sock_create(IPPROTO_QUIC, &sock)
kernel_bind(sock)


kernel_connect(sock)
tls_client_hello_x509(args:{sock})




kernel_sendmsg(sock)
sock_release(sock)
```

- Server

```
__sock_create(IPPROTO_QUIC, &sock)
kernel_bind(sock)
kernel_listen(sock)




kernel_accept(sock, &newsock)
tls_server_hello_x509(args:{newsock})

kernel_recvmsg(newsock)
sock_release(newsock)
sock_release(sock)
```

Sample: https://github.com/lxin/quic/blob/main/net/quic/sample_test.c

Advanced APIs

- Control Message
  - sendmsg() with cmsg stream info to set stream_id and flag
  - recvmsg() with cmsg stream info to get stream_id and flag
- Socket Options
  - SOCKOPT_KEY_UPDATE for rekeying
  - SOCKOPT_CONNECTION_MIGRATION for connection migration
  - SOCKOPT_STREAM_OPEN/RESET/STOP_SENDING for stream management
  - SOCKOPT_SESSION_TICKET for session resumption and 0-RTT Data
- Notification/Events
  - Connection Update
  - Stream Update

Man doc: https://github.com/lxin/quic/wiki/man

Examples: https://github.com/lxin/quic/tree/main/tests

Next Step

- HW crypto offloading
- Internet Draft For QUIC Sockets API Extensions

Summary from Xin Long

Kernel QUIC specialities:

- Enables use for Kernel Consumers
- Common easy to understand Socket APIs, like listen/accept/connect/send/recv/close.
- Avoid Data Copies from Users to Kernel via sendfile() syscall.
- Easy to cooperate with offloading in NICs, no configuration needed from userspace
- Less Interoperability Issues to implement QUIC in OS

WIP summary from Xin Long

What’s next for their driver (net/quic/quic.ko):

1. Improve the performance and complete the congestion control code (there's some testing data vs kTLS in github README)

2. Collaborate with Andy from Broadcom to implement the infrastructure for the Offloading inside HW NICs.

3. Work with another colleague to start an RFC doc standardizing the socket APIs for OS-level QUIC implementation.

What about the SMB3.1.1 pieces

- Changes on the SMB3.1.1 client side are small
- New socket type to support (see slide 26)
- ALPN (app layer protocol negotiation for TLS) Identification sequence used to identify the SMB2 protocol over QUIC is 0x73 0x6D 0x62 ("smb")
- See section 4.10 of MS-SMB2 and section 2.2.3.1.5 (SMB2_TRANSPORT_CAPABILITIES negotiate context which we will have to parse)

Next steps

- Continue following up with Dave and the network maintainers if we have updates on our SMB3.1.1 testing with it
- Continue building and testing Xin's git tree until it is in linux-next or mainline (when we can move to using that version)
- Add SMB3.1.1 code to fs/smb/client/transport.c to read/write over the new QUIC driver
- Add SMB3.1.1 code to process the transport capabilities SMB3.1.1 negotiate context to fs/smb/client/smb2pdu.c

More Next steps

- What about user space tooling?
  - updates to cifs-utils and/or Samba tools could help make this easier for users
  - New mount option for cifs.ko to force use of QUIC if available
- What about certificate setup and exchange?
  - Could Samba tooling help?
- Encourage code contributions, review and testing …
  - Would love some help on this exciting feature

Testing Improvements

Test … test … test ...

Additional tests are encouraged (quic or smb specific)

- See the tests subdirectory of the quic github tree
- The basic xfstests should be fine for testing SMB3.1.1 over QUIC, but we will need to come up with some reconnect tests (today mostly in the buildbot's "DFS" test group) to exercise reconnect with the new code, and will need to test various encryption options (QUIC only, SMB3.1.1 only and both QUIC and SMB3.1.1 encryption)

quic function tests

To run the functional tests

cd tests

make run

(runs 192 functional tests

then does various perf tests)

```
****** [Function Tests (PSK)] ******

HANDSHAKE DONE
STREAM TEST:
test1: PASS (not allowed send(MSG_SYN) to open a stream when last is not closed)
test2: PASS (use send(MSG_SYN) to open one stream)
test3: PASS (use send(MSG_SYN) to open next stream after last is closed)
test4: PASS (use send(MSG_SYN) to open next bidi stream after last is closed)
test5: PASS (not allowed to open a stream that is already closed with getsockopt(QUIC_SOCKOPT_STREAM_OPEN))
test6: PASS (use getsockopt(QUIC_SOCKOPT_STREAM_OPEN) to open a specific stream)
test7: PASS (use getsockopt(QUIC_SOCKOPT_STREAM_OPEN) to open next bidi stream)
test8: PASS (use getsockopt(QUIC_SOCKOPT_STREAM_OPEN) to open next uni stream)
test9: PASS (not allowed to open a stream that is already closed with sendmsg(QUIC_STREAM_FLAG_NEW))
test10: PASS (not allowed to open a stream twice with sendmsg(QUIC_STREAM_FLAG_NEW))
test11: PASS (sendmsg with a specific stream normally)
test12: PASS (not allowed to open a stream with sendmsg(sid == -1) if it the old one is not closed
test13: PASS (open next uni stream with sendmsg(sid == -1))
test14: PASS (open next bidi stream with sendmsg(sid == -1))
test15: PASS (open multiple stream and send on 1st one)
test16: PASS (open multiple stream and send on 2nd one)
test17: PASS (not allowed to send data on a closed stream)
test18: PASS (sendmsg with sid > max_streams_bidi in blocked mode)
test19: PASS (sendmsg with sid > max_streams_uni in blocked mode)
test20: PASS (return -EAGAIN in bidi non-blocked mode)
test21: PASS (sendmsg with sid > max_streams_bidi in non-blocked mode)
test22: PASS (return -EAGAIN in uni non-blocked mode)
test23: PASS (sendmsg with sid > max_streams_uni in non-blocked mode)
test24: PASS (getsockopt(QUIC_SOCKOPT_STREAM_OPEN) with sid > max_streams_bidi in blocked mode)
test25: PASS (getsockopt(QUIC_SOCKOPT_STREAM_OPEN) with sid > max_streams_uni in blocked mode)
test26: PASS (return -EAGAIN in bidi non-blocked mode)
test27: PASS (getsockopt(QUIC_SOCKOPT_STREAM_OPEN) with sid > max_streams_bidi in non-blocked mode)
test28: PASS (return -EAGAIN in uni non-blocked mode)
test29: PASS (sendmsg with sid > max_streams_uni in non-blocked mode)
test30: PASS (not allowed to reset a closed stream)
test31: PASS (not allowed to reset a stream that hasn't opened)
test32: PASS (reset a opened stream)
test33: PASS (not allowed to send data on a reset stream)
test34: PASS (not allowed to send data with FIN on a reset stream)
test35: PASS (not allowed to send data on a reset stream by peer stop_sending)
test36: PASS (not allowed to send data with FIN on a reset stream set by peer stop_sending)
CONNECTION TEST:
test1: PASS (retire source connection id 0)
test2: PASS (retire source connection id 1)
test3: PASS (not allow to retire a retired source connection id)
test4: PASS (not allow to retire all source connection id)
test5: PASS (retire multiple source connection id)
test6: PASS (retire max_count - 1 source connection id)
test7: PASS (retire dest connection id 0)
test8: PASS (retire dest connection id 1)
test9: PASS (not allow to retire a retired dest connection id)
test10: PASS (not allow to retire all dest connection id)
test11: PASS (retire multiple dest connection id)
test12: PASS (retire max_count - 1 dest connection id)
test13: PASS (connection migration is set)
test14: PASS (send message with new address)
```


Xfstest automation

Will also be important to add xfstest scenarios (e.g. local.config using a mount option that forces use of QUIC instead of TCP)

Example here is over TCP and runs 100s of fs tests

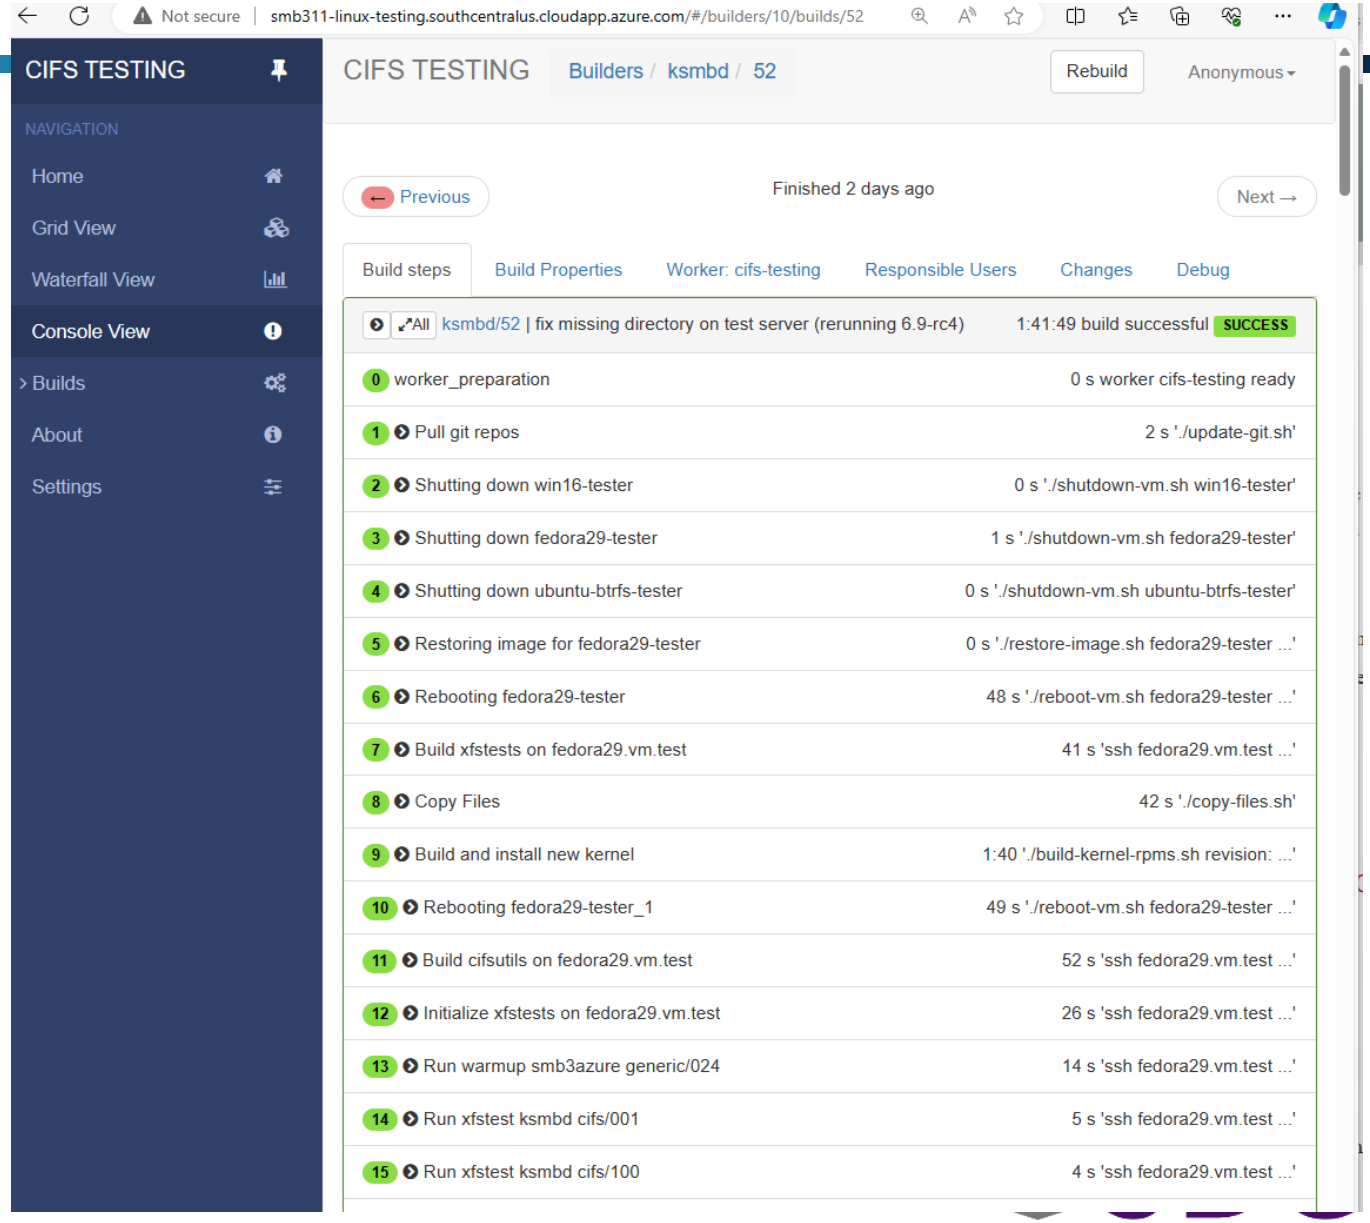

Thank you for your time

- Future is very bright!

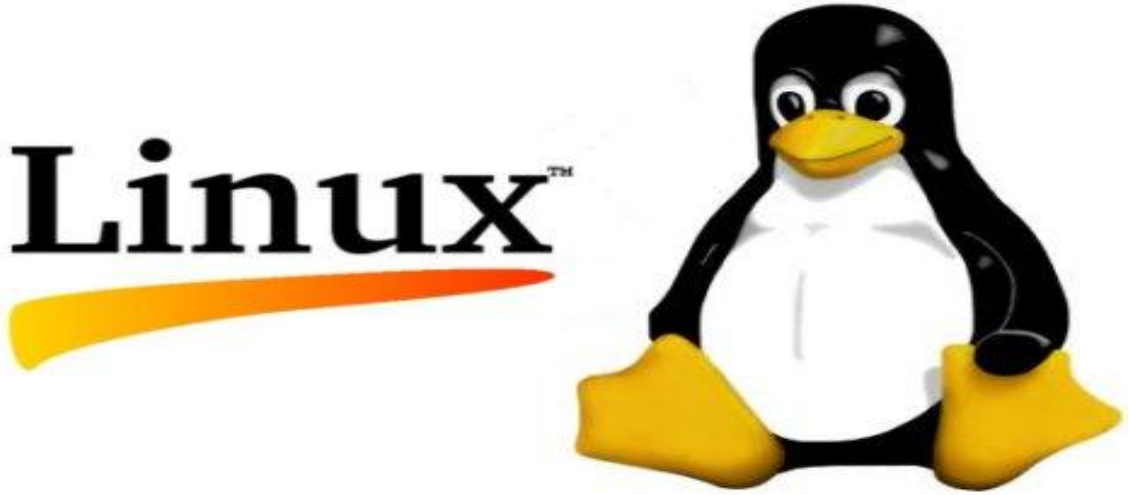

+ ***SMB3***

Additional Resources to Explore

https://github.com/lxin/quic
And an alternative open source implementation in userspace:
https://github.com/microsoft/msquic

https://techcommunity.microsoft.com/t5/storage-at-microsoft/smb-over-quic-is-ga-and-we-have-lots-of-news/ba-p/2928695 and there is also a newer post on SMB3.1.1 and QUIC by Ned

https://www.snia.org/sites/default/files/SDCEMEA/2021/snia-SMB-over-QUIC.pdf

https://learn.microsoft.com/en-us/windows-server/storage/file-server/smb-over-quic

Three research papers with interesting observations about areas where Linux QUIC drivers can improve perf:

https://publikationen.bibliothek.kit.edu/1000161904/152028985

https://www.net.in.tum.de/fileadmin/bibtex/publications/papers/jaeger2023quic.pdf
"QUIC is not Quick Enough" https://dl.acm.org/doi/10.1145/3589334.3645323

Please take a moment to rate this session.

Your feedback is important to us.

