Offloading Encryption to QUIC Enabled NICs

Andy Gospodarek – Software Architect, Broadcom Eric Davis – System Architect, Broadcom Eric Spada – Distinguished Engineer, Broadcom

Overview of QUIC

Transport protocol that runs over UDP

Utilizes TLS1.3 for handshake/key exchange

Transmit and Receive Connection IDs are used to identify unique connections



Encryption is self-contained within each packet

Protoco	ol
✓ Fra	ame
~	Ethernet Address Resolution Protocol V Internet Protocol Version 4
	 Transmission Control Protocol
	Malformed Packet Transport Layer Security Ver Datagram Protocol Data Domain Name System
	QUIC IETF Simple Service Discovery Protec
	 Internet Protocol Version 6 User Datagram Protocol Multicast Domain Name System

^	Percent	Packets	Packets	Percent Bytes
		100.0	36214	100.0
		100.0	36214	1.6
	l	4.5	1631	0.2
		95.5	34569	2.0
			4	0.0
		14.6	5289	8.1
		0.0	1	0.0
		0.1	20	0.0
		6.8	2447	7.6
		80.8	29276	0.7
	1	4.1	1500	2.6
		0.4	160	0.0
			14	0.0
		76.2	27600	84.8
			2	0.0
		0.0	14	0.0
		0.0	14	0.0
		0.0	14	0.0



Hardware offload good; software only bad

Checksum Offload, LRO/GRO, GSO, Tunnel Offload, nTuple filtering, HW Timestamping, kTLS, TC, ...

What could QUIC hardware offload look like?

Leverage experience from kTLS Offload to support QUIC Offload

quic kernel module with userspace/netlink and driver interface





Tracks flow, connection ID, direction and QUIC session/crypto information for offload



#define QUIC_MAX_CONN_ID_LEN 20 #define QUIC_MAX_PKT_NUM_LEN 8

struct quic_flow_ctx {
 u8 conn_id_len;
 u8 conn_id[QUIC_MAX_CONN_ID_LEN];
 u8 pkt_num[QUIC_MAX_PKT_NUM_LEN];
 u8 key_phase;
 u8 cipher;
 u16 xid;

union {

struct cipher_aes_gcm_128 aes_gcm_128; struct cipher_aes_ccm_128 aes_ccm_128; struct cipher_chacha20_poly1305 chacha20_poly1305; cipher_info:

} cipher_info;

};

Receives offload requests via netlink



ip crypto show mode quic 4ed45802: rx 10.0.0.1:53227 10.0.0.100:443 cid e45ea092c337f4eb cipher aes_gcm_128 dev ens7f0np0







Adds offloaded connection info to drivers and hardware



enum quic_offload_dir { QUIC_OFFLOAD_DIR_RX, QUIC_OFFLOAD_DIR_TX };

int (*quic_dev_add)(struct net_device *netdev, struct sock *sk,

struct quic_flow_ctx *flow_ctx enum quic_offload_dir direction);



Returns XID to caller if flow can be offloaded

QUIC userspace components





Application/Library does initial connection setup/handshake

If application wants to offload; send request to offload via netlink

Application stores XID returned in netlink ext_ack cookie

Transmit datapath

XID passed into cmsghdr if data is to be encrypted by hardware

XID added to sk_buff extension in socket layer

Network driver formats TX buffer descriptor as needed based on sk_buff extension data

Hardware sends encrypted packet to peer

Receive datapath

Frames will be decrypted by hardware and validation reported in buffer descriptor

Driver will populate sk_buff extension with decrypt and authentication status from buffer descriptor

Socket will set cmsghdr and pass back to application, indicating XID, encryption, validation status

Sounds simple, right?

Conclusions and Next Steps

QUIC is well suited for device offload

Modifications to kernel and network drivers are clear

QUIC device offload will be part of bnxt_en driver updates

QUIC Libraries may take more work to increase adoption and usage

Plumbers Conference

Richmond, Virginia | November 13-15, 2023

