

QUICForge

Client-side Request Forgery in QUIC

Yuri Gbur and Florian Tschorsch

Yuri Gbur

- MSc in Computer Science at Technische Universität (TU) Berlin
- Security Consultant and Researcher at SEC Consult
- Berlin, Germany

yuri.gbur@posteo.com

y.gbur@sec-consult.com

 @yukonsec

Florian Tschorsch

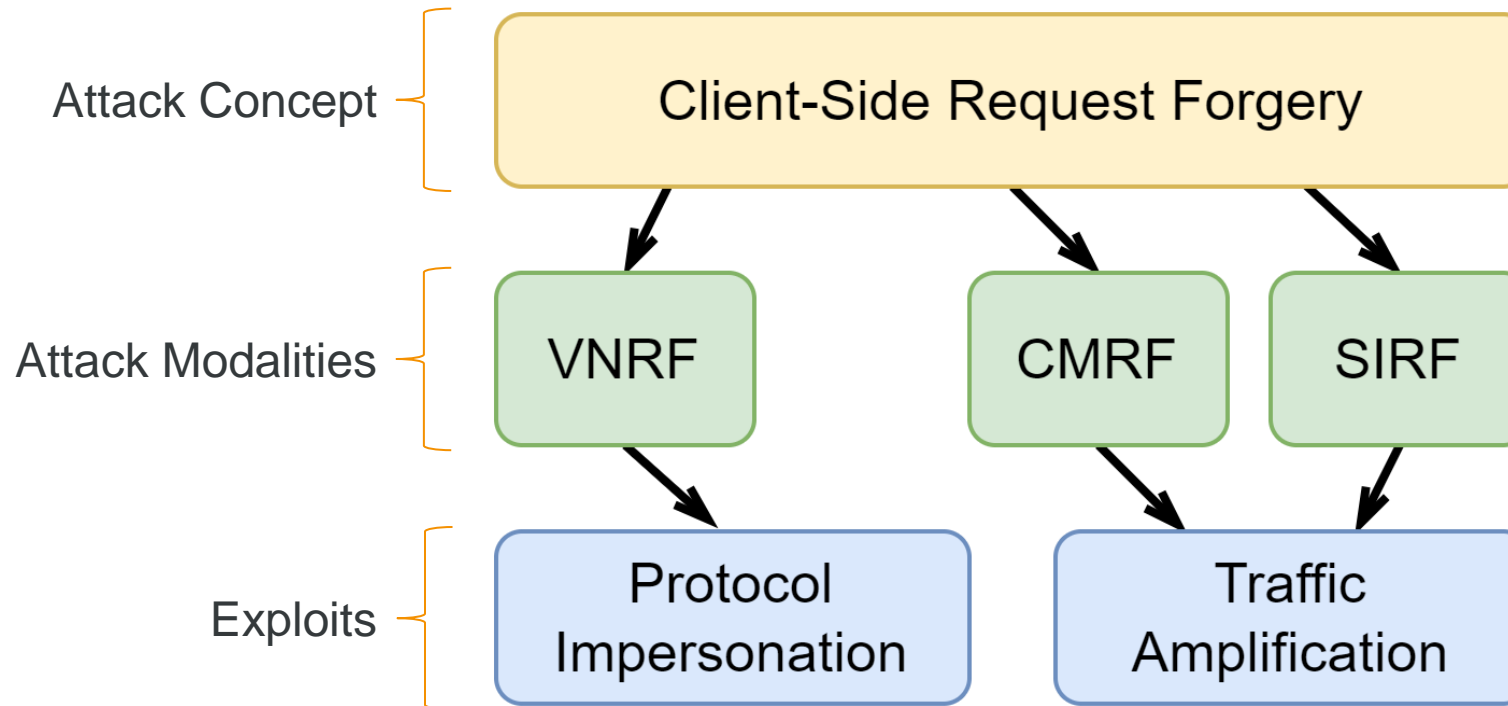
- Professor for Computer Science at Technische Universität (TU) Berlin
- Chair for Distributed Security Infrastructure (DSI)
- Berlin, Germany

florian.tschorsch@tu-berlin.de

 @flotschorsch

Research Goals

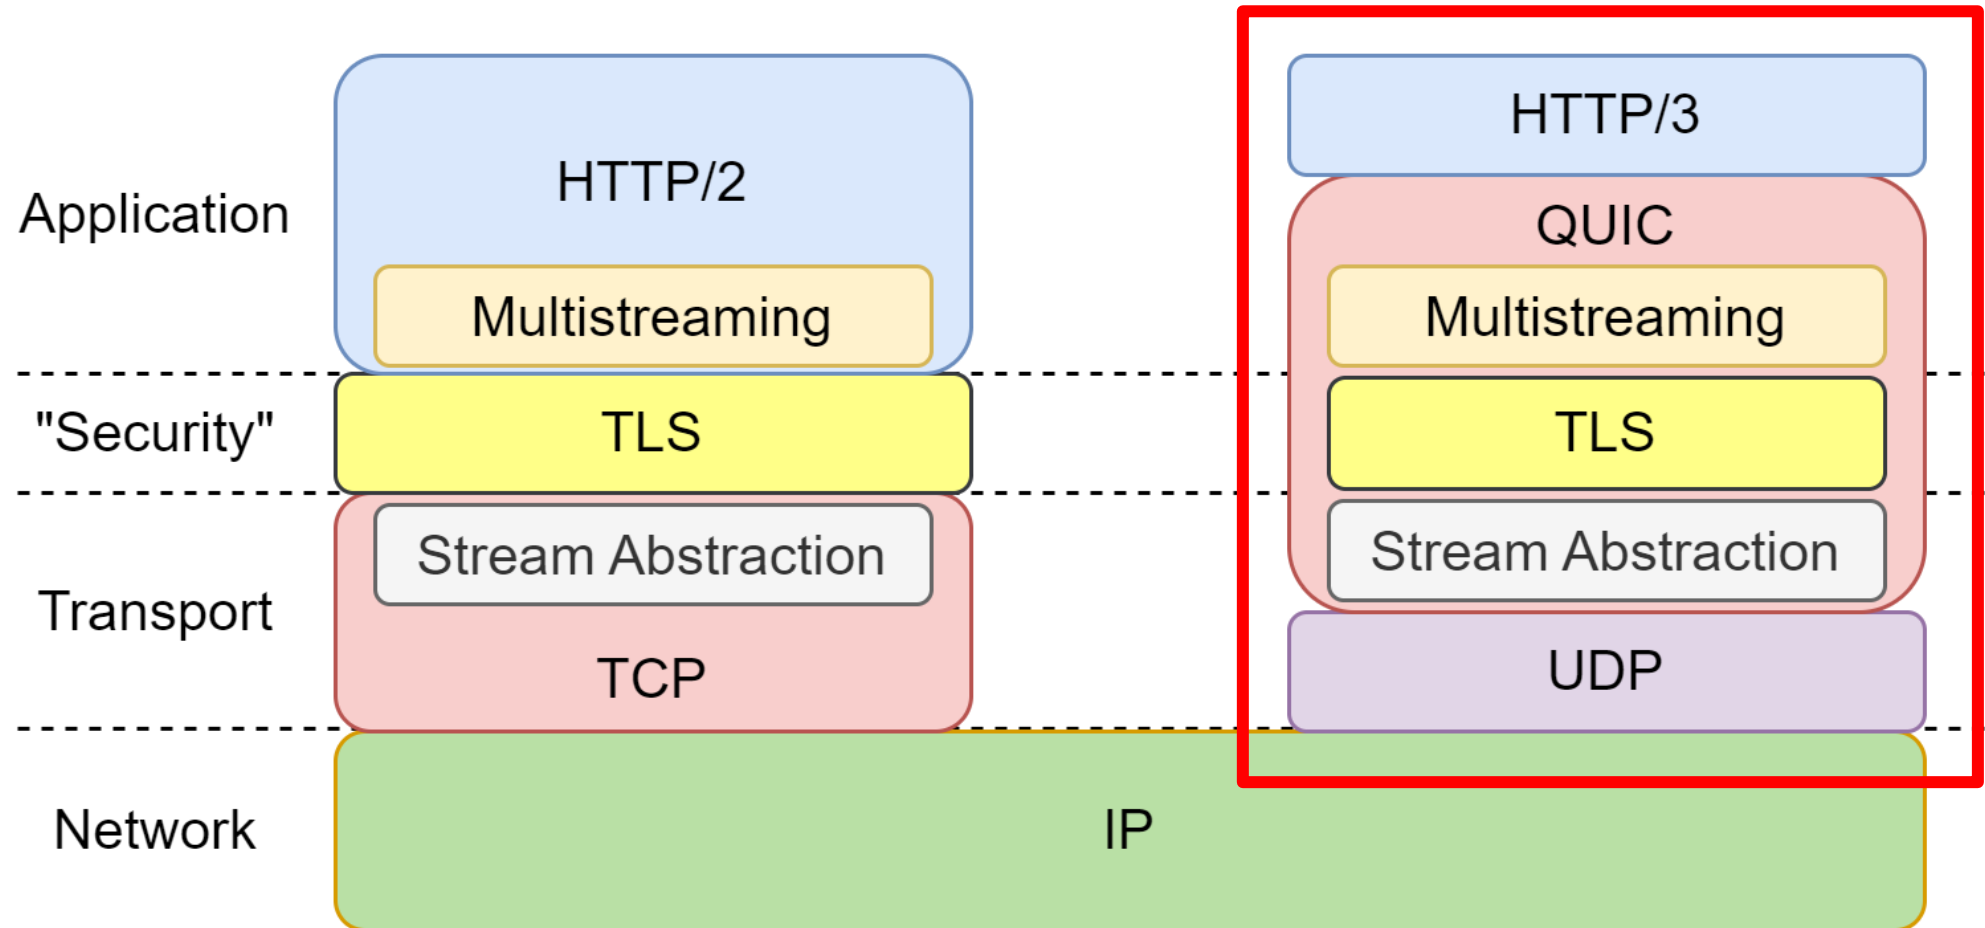
Attack Modalities VS Exploits



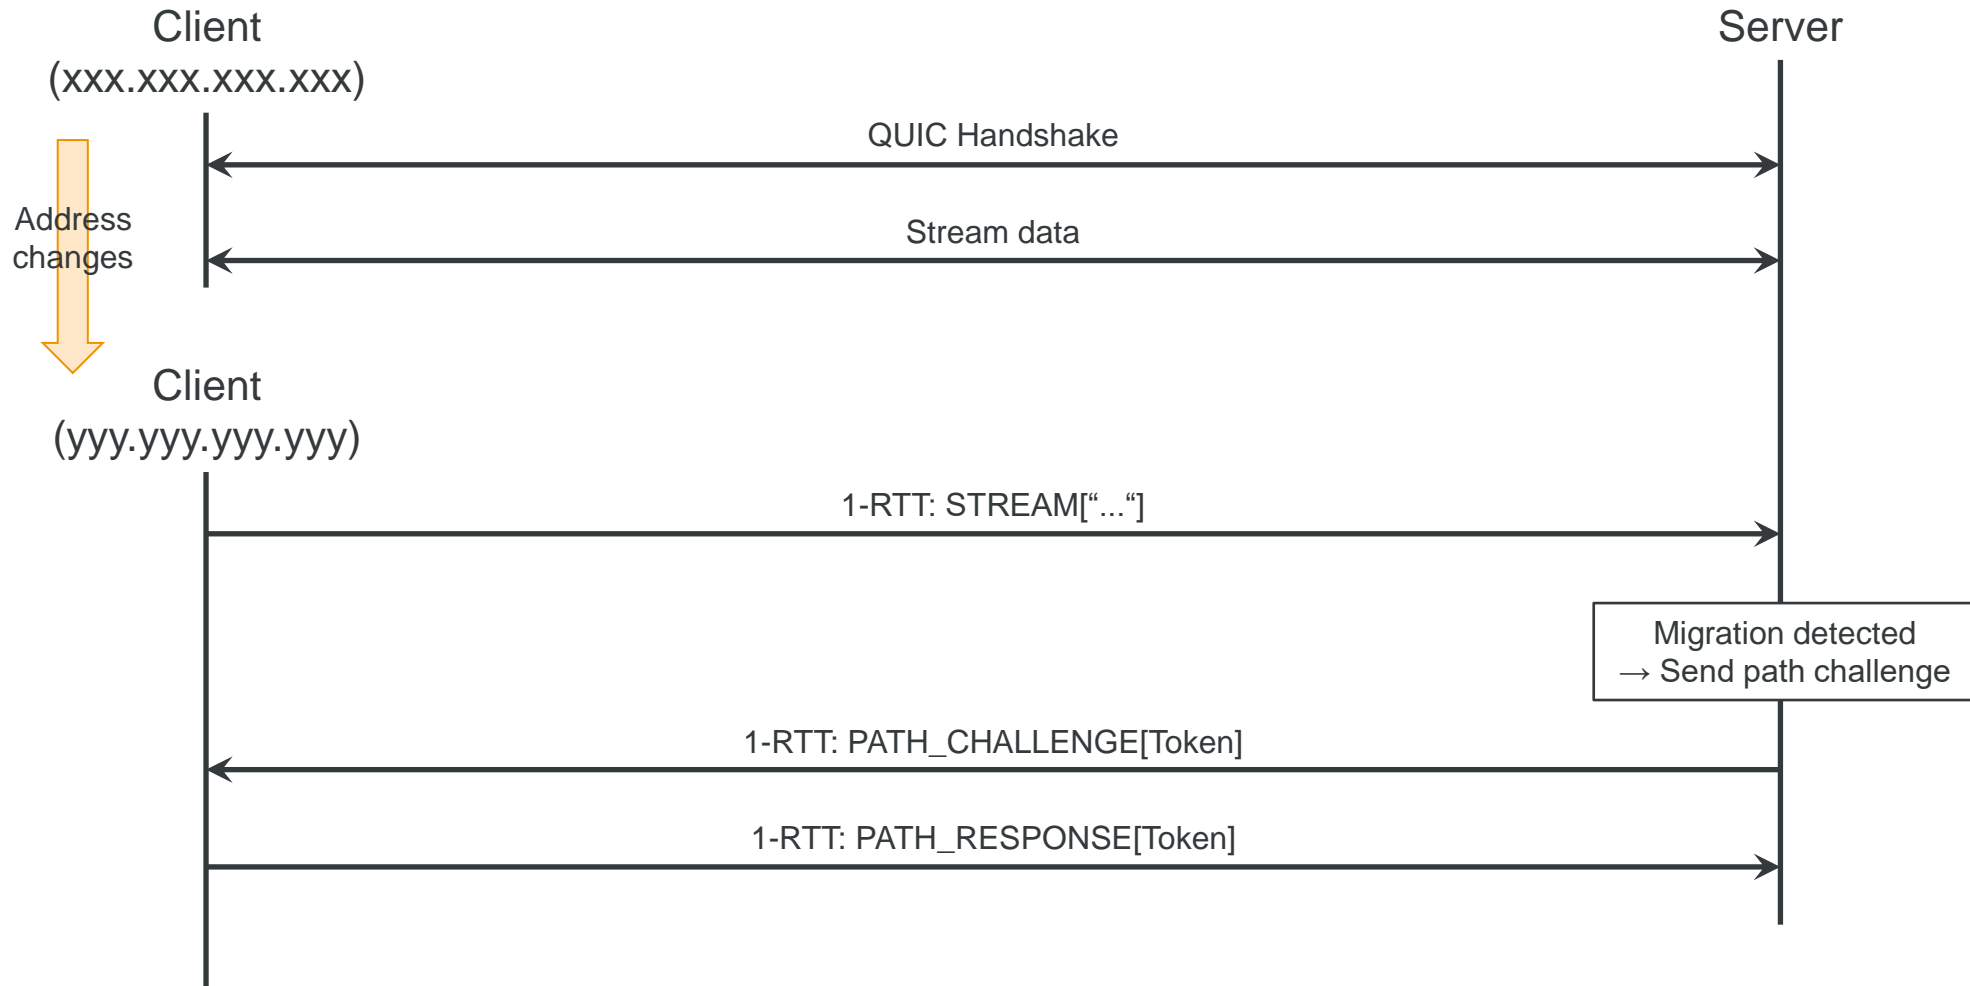
- Version Negotiation Request Forgery (VNRF)
- Connection Migration Request Forgery (CMRF)
- Server Initial Request Forgery (SIRF)

QUIC(K) Background

HTTP/2 VS HTTP/3




Connection Migration



Request Forgery

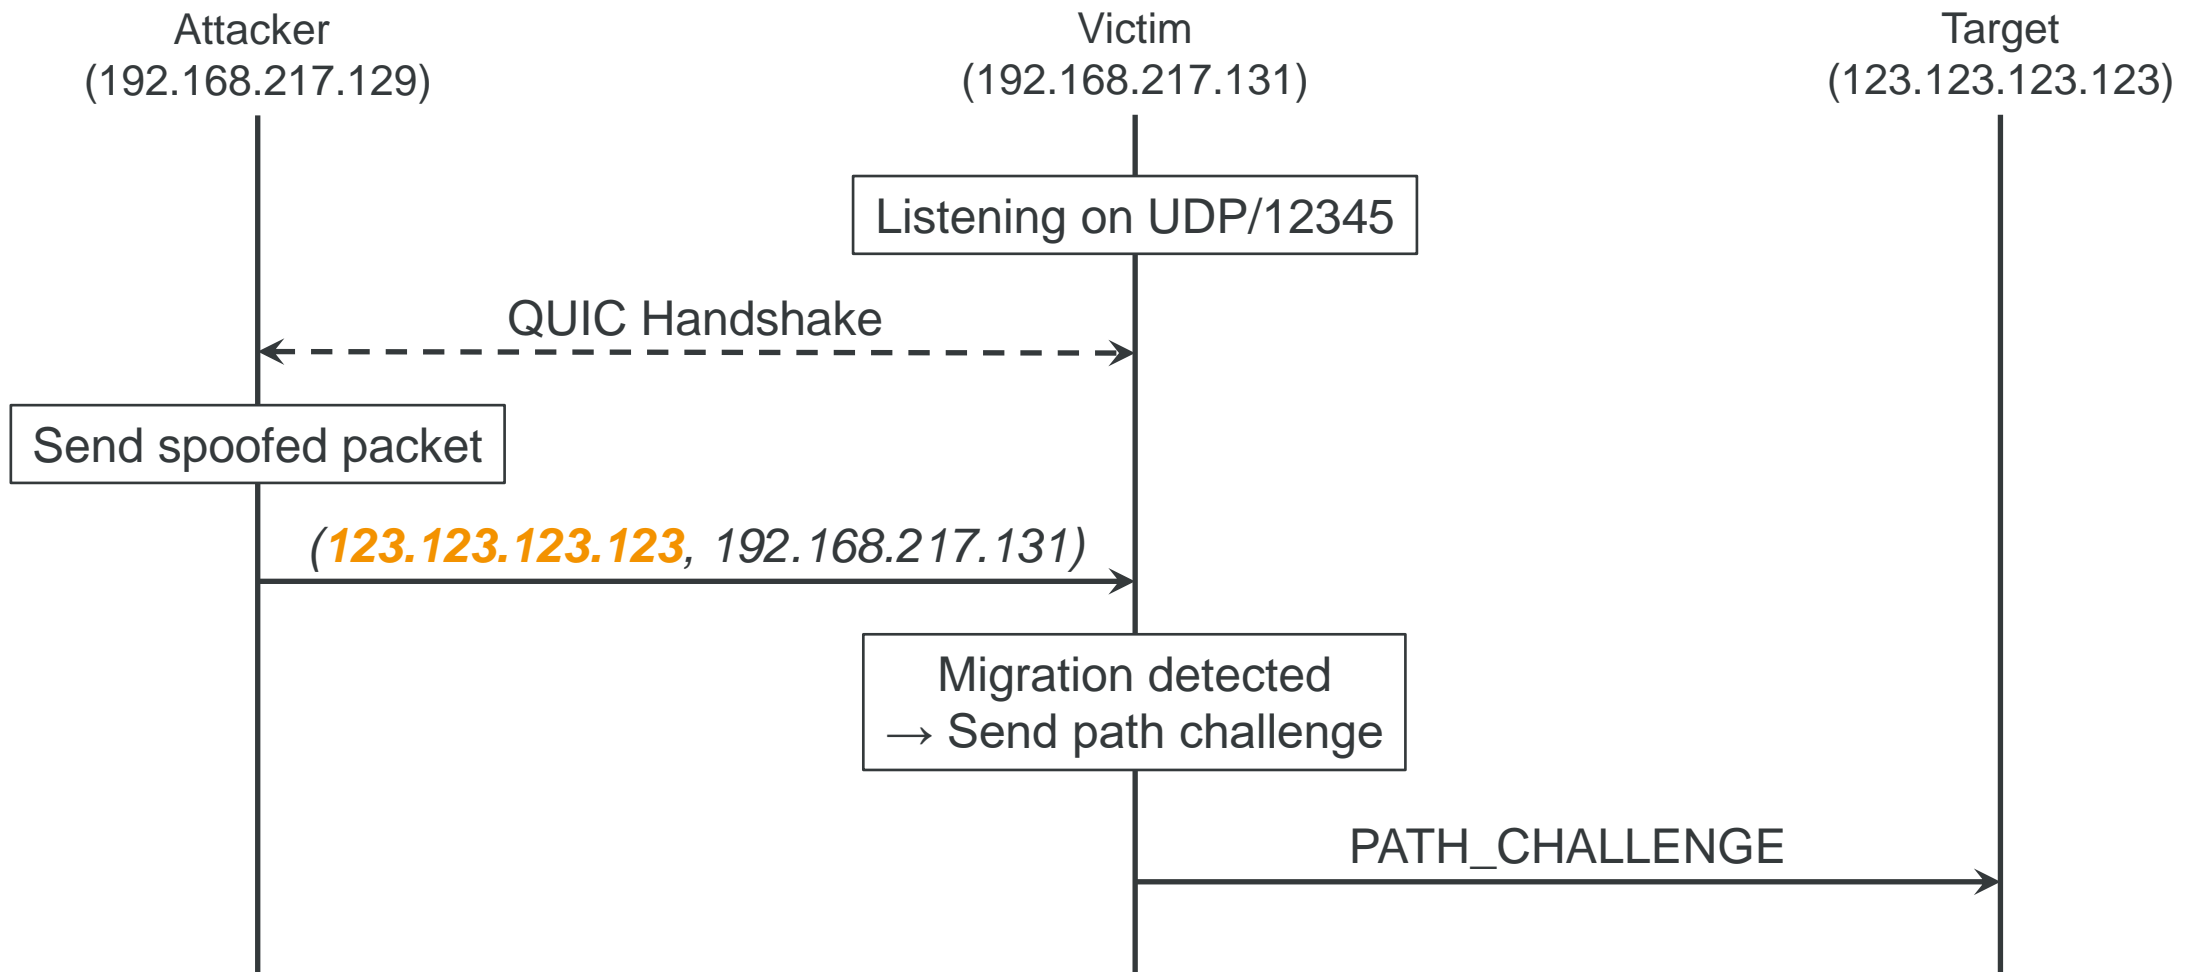
Client-side Request Forgery



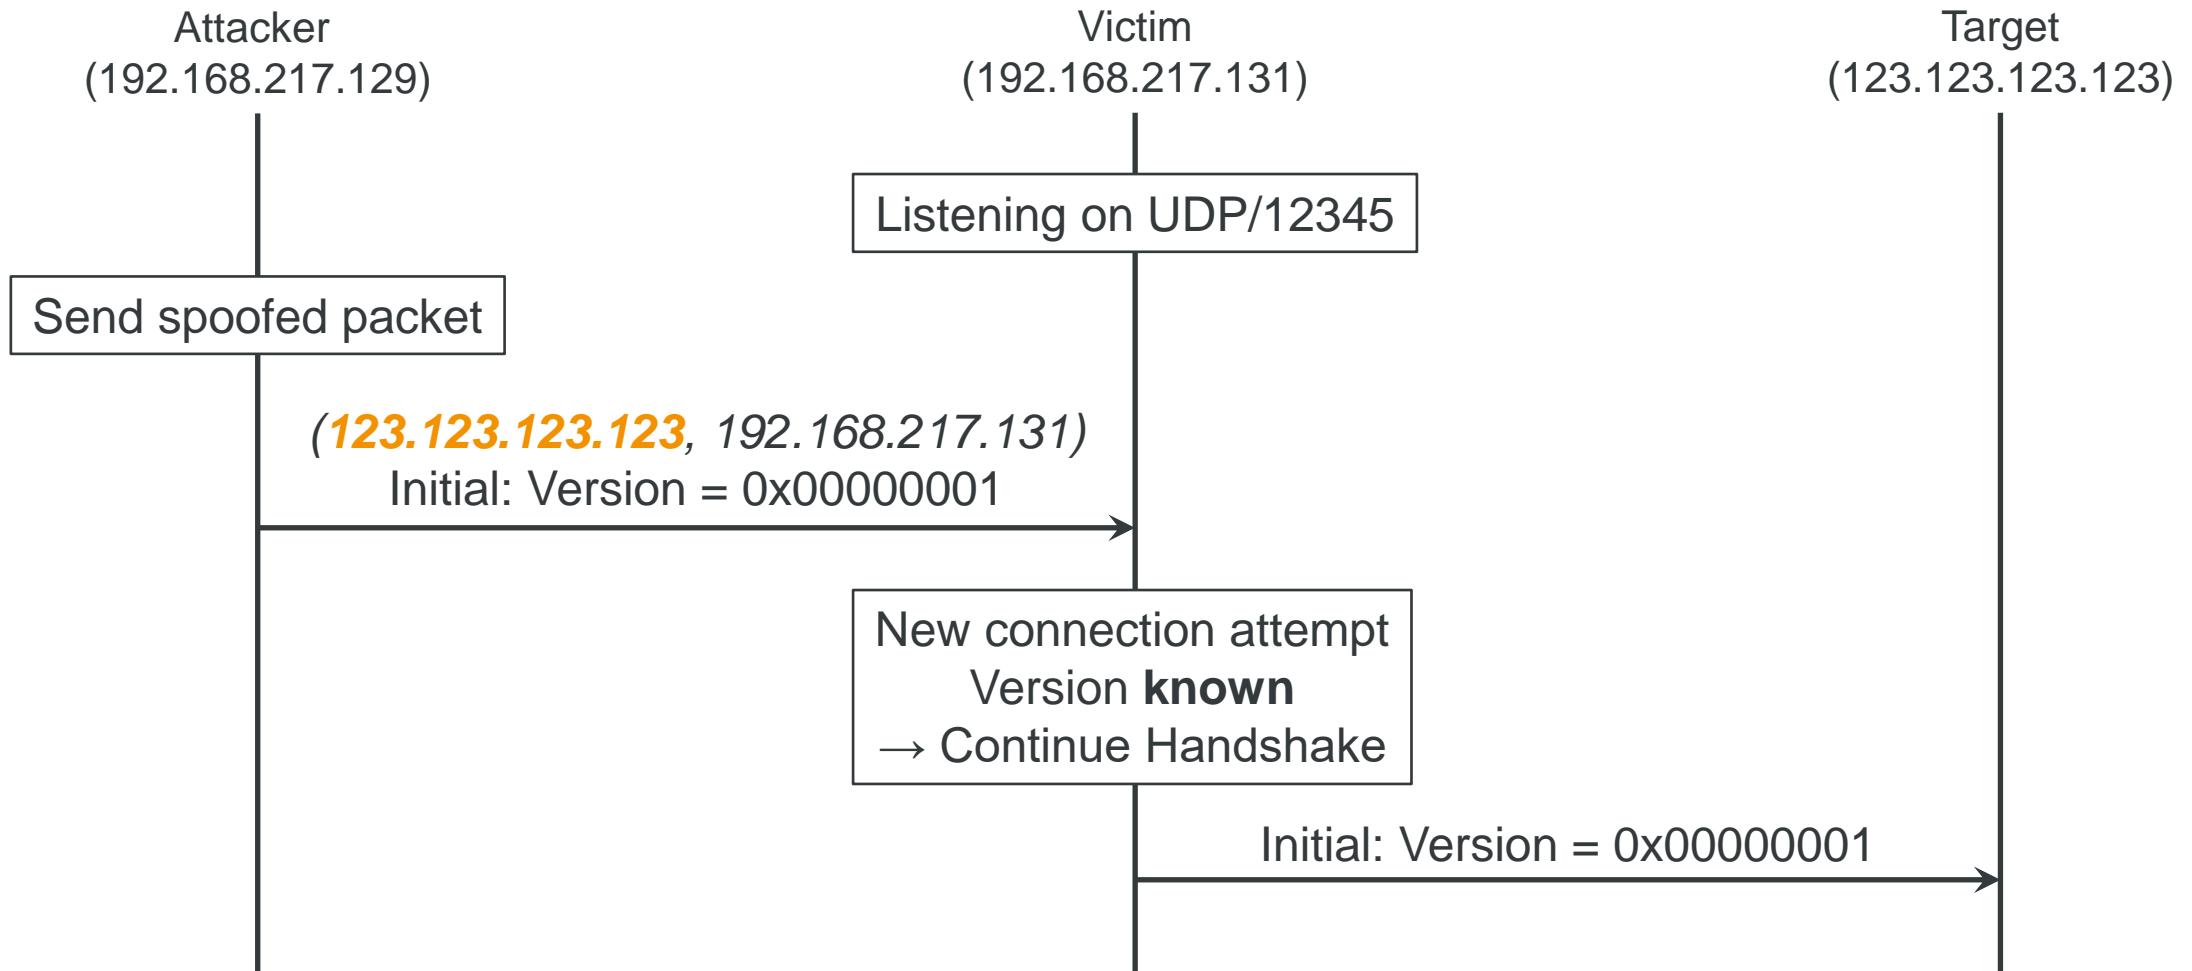
— Bypassing Network Restrictions

➔ Utilizing Victim Resources


Connection Migration Request Forgery (CMRF)



Server Initial Request Forgery (SIRF)

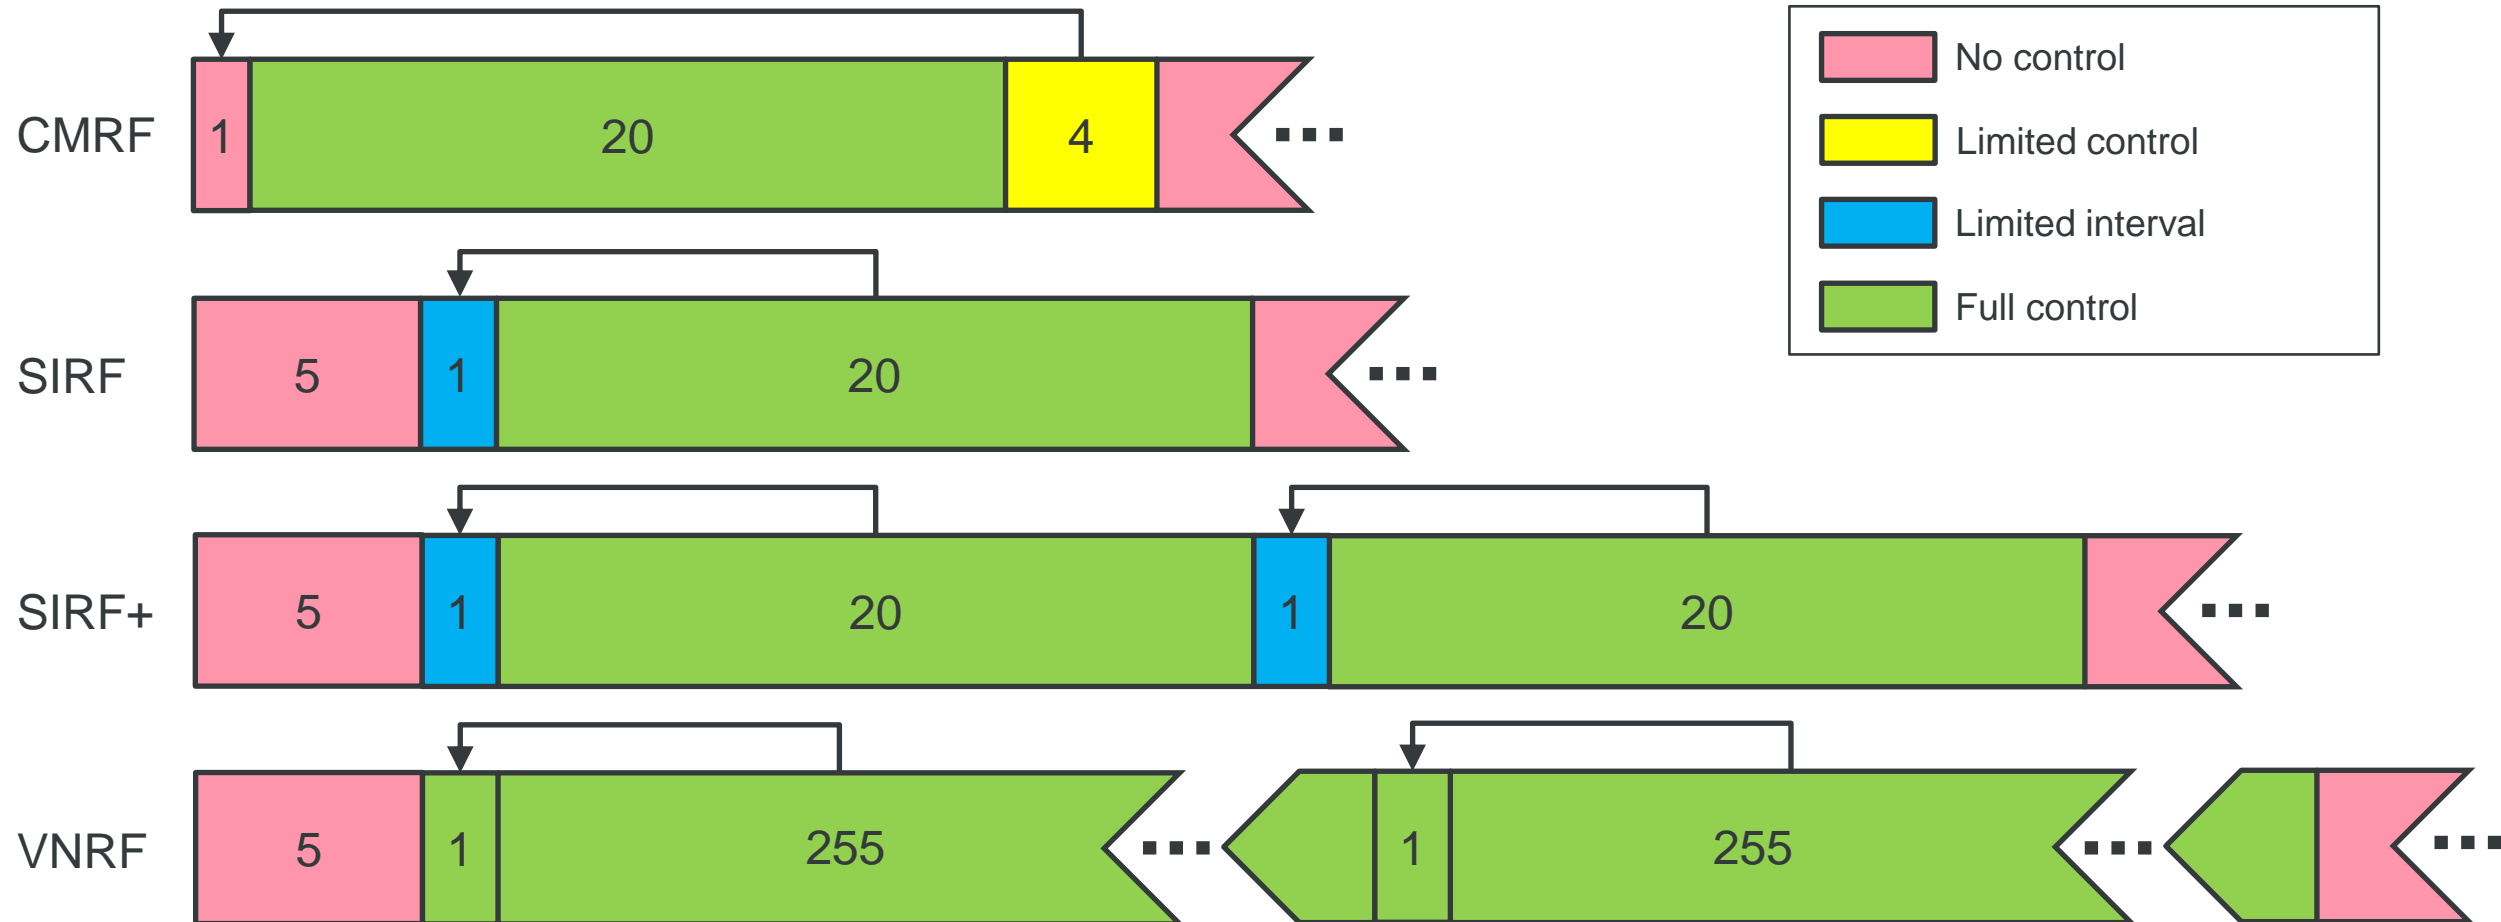


Version Negotiation Request Forgery (VNRFF)

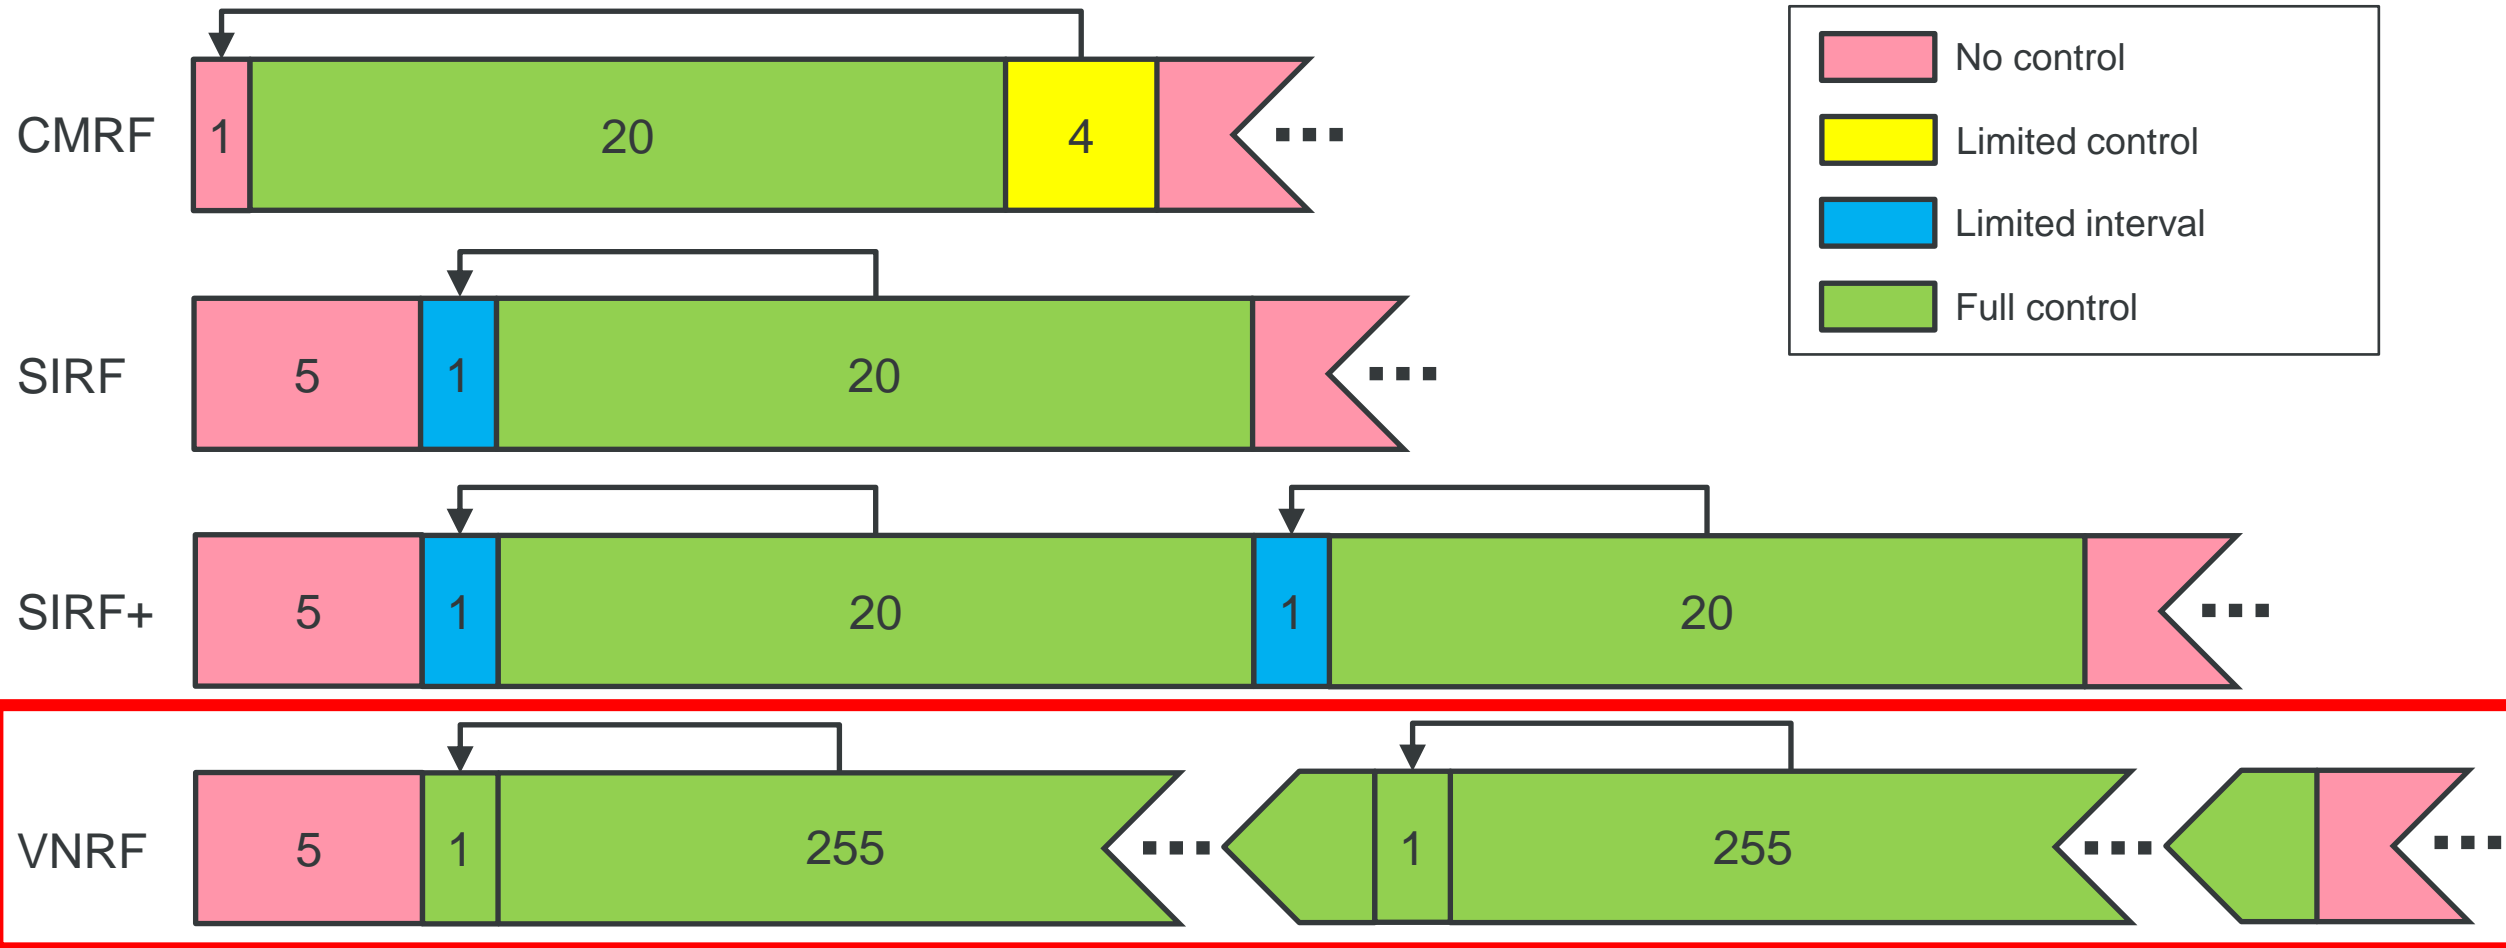


Protocol Impersonation

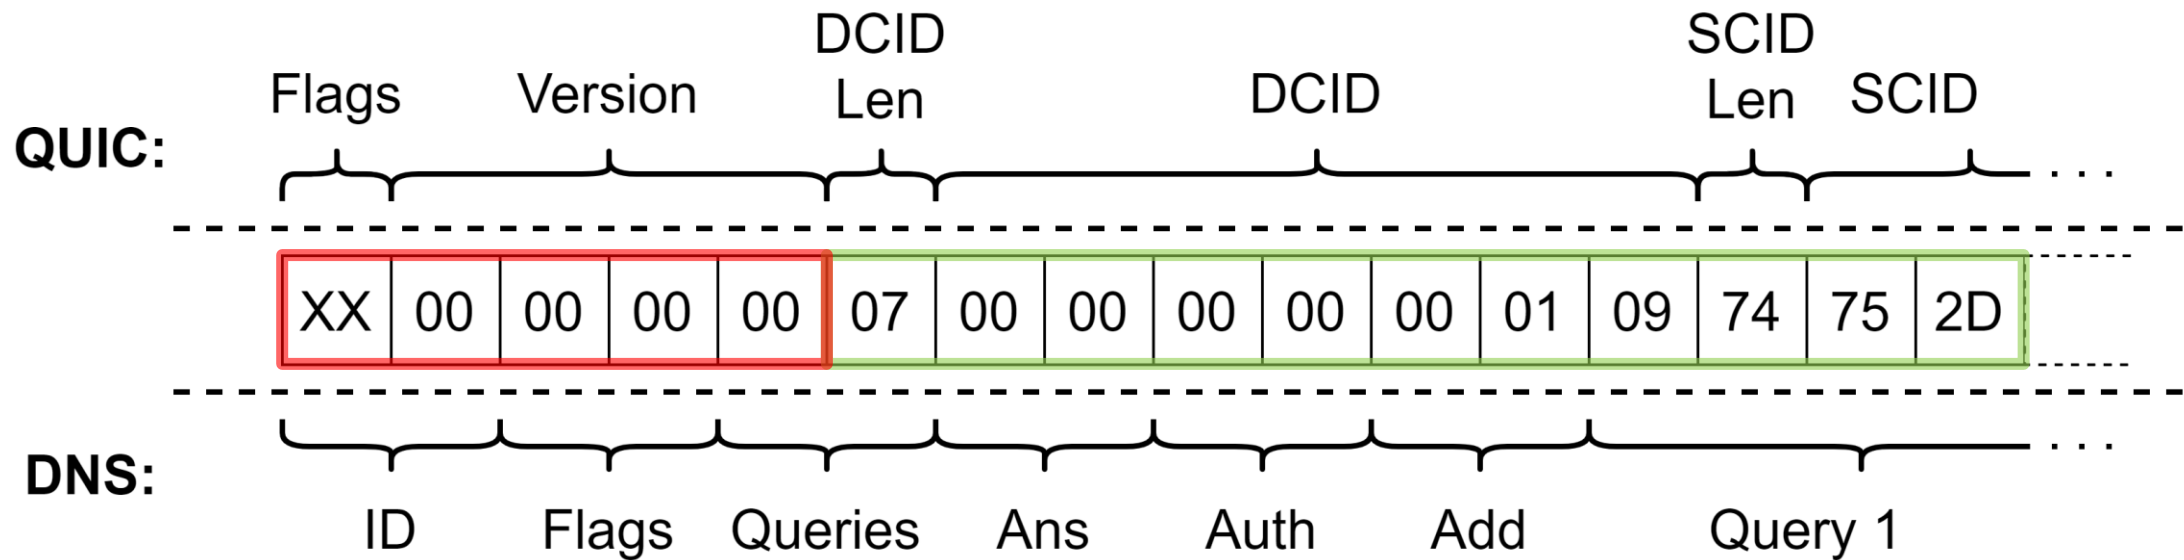
Controllable Bytes for Protocol Impersonation



Controllable Bytes for Protocol Impersonation



Impersonating DNS Requests with VNRF



Impersonating DNS Requests with VNRF (cont'd)

No.	Time	Source	Destination	Protocol	Length	Info
13	3.538438...	8.8.8.8	192.168.217.1...	QUIC	13...	Initial, SCID=00000000000109
14	3.538771...	192.168.217.1...	8.8.8.8	QUIC	200	Version Negotiation, DCID=00000000000109
15	3.558935...	8.8.8.8	192.168.217.1...	QUIC	152	53 → 12345 Len=110[Malformed Packet]

Frame 14: 200 bytes on wire (1600 bits), 200 bytes captured (1600 bits) on interface ens33, id 0
Ethernet II, Src: VMware_5e:6a:92 (00:0c:29:5e:6a:92), Dst: VMware_f6:95:1c (00:50:56:f6:95:1c)
Internet Protocol Version 4, Src: 192.168.217.131, Dst: 8.8.8.8
User Datagram Protocol, Src Port: 12345, Dst Port: 53
QUIC IETF
 QUIC Connection information
 [Packet Length: 158]
 1... .. = Header Form: Long Header (1)
 .100 1001 = Unused: 0x49
 Version: Version Negotiation (0x00000000)
 Destination Connection ID Length: 7
 Destination Connection ID: 00000000000109
 Source Connection ID Length: 116
 Source Connection ID: 752d6265726c696e0264650000100010000010001000001000100000100010000010001...
 Supported Version: v2-draft-01 (0x709a50c4)
 Supported Version: 1 (0x00000001)
 Supported Version: draft-32 (0xff000020)
 Supported Version: draft-31 (0xff00001f)
 Supported Version: draft-30 (0xff00001e)
 Supported Version: draft-29 (0xff00001d)
 Supported Version: Unknown (0x4a0ababa) (GREASE)

Frame (frame), 200 bytes Packets: 31 · Displayed: 3 (9.7%) Profile: Default

No.	Time	Source	Destination	Protocol	Length	Info
13	3.538438...	8.8.8.8	192.168.217.1...	DNS	13...	DNS Stateful operations (DSO) 0xc813[Malformed Packet]
14	3.538771...	192.168.217.1...	8.8.8.8	DNS	200	Standard query 0xc900 A tu-berlin.de A <Root> A <Root> A <Root> A
15	3.558935...	8.8.8.8	192.168.217.1...	DNS	152	Standard query response 0xc900 A tu-berlin.de A 10.150.7.69 A 172

Frame 14: 200 bytes on wire (1600 bits), 200 bytes captured (1600 bits) on interface ens33, id 0
Ethernet II, Src: VMware_5e:6a:92 (00:0c:29:5e:6a:92), Dst: VMware_f6:95:1c (00:50:56:f6:95:1c)
Internet Protocol Version 4, Src: 192.168.217.131, Dst: 8.8.8.8
User Datagram Protocol, Src Port: 12345, Dst Port: 53
Domain Name System (query)
 Transaction ID: 0xc900
 Flags: 0x0000 Standard query
 Questions: 7
 Answer RRs: 0
 Authority RRs: 0
 Additional RRs: 1
 Queries
 tu-berlin.de: type A, class IN
 <Root>: type A, class IN
 <Root>: type A, class IN
 <Root>: type A, class IN
 <Root>: type A, class IN
 <Root>: type A, class IN
 <Root>: type A, class IN
 Additional records
 <Root>: type Unused, class Unknown
 [Response In: 15]

Frame (frame), 200 bytes Packets: 31 · Displayed: 3 (9.7%) · Dropped: 0 (0.0%) Profile: Default

No.	Time	Source	Destination	Protocol	Length	Info
15	3.558935...	8.8.8.8	192.168.217.1...	QUIC	152	53 → 12345 Len=110[Malformed Packet]

Frame 15: 152 bytes on wire (1216 bits), 152 bytes captured (1216 bits) on interface ens33, id 0
Ethernet II, Src: VMware_f6:95:1c (00:50:56:f6:95:1c), Dst: VMware_5e:6a:92 (00:0c:29:5e:6a:92)
Internet Protocol Version 4, Src: 8.8.8.8, Dst: 192.168.217.131
User Datagram Protocol, Src Port: 53, Dst Port: 12345
QUIC IETF
 QUIC Connection information
 [Malformed Packet: QUIC]
 [Expert Info (Error/Malformed): Malformed Packet (Exception occurred)]
 [Malformed Packet (Exception occurred)]
 [Severity level: Error]
 [Group: Malformed]

Frame (frame), 200 bytes Packets: 31 · Displayed: 3 (9.7%) Profile: Default

No.	Time	Source	Destination	Protocol	Length	Info
15	3.558935...	8.8.8.8	192.168.217.1...	DNS	152	Standard query response 0xc900 A tu-berlin.de A 10.150.7.69 A 172

Frame 15: 152 bytes on wire (1216 bits), 152 bytes captured (1216 bits) on interface ens33, id 0
Ethernet II, Src: VMware_f6:95:1c (00:50:56:f6:95:1c), Dst: VMware_5e:6a:92 (00:0c:29:5e:6a:92)
Internet Protocol Version 4, Src: 8.8.8.8, Dst: 192.168.217.131
User Datagram Protocol, Src Port: 53, Dst Port: 12345
Domain Name System (response)
 Transaction ID: 0xc900
 Flags: 0x8080 Standard query response, No error
 Questions: 1
 Answer RRs: 5
 Authority RRs: 0
 Additional RRs: 0
 Queries
 Answers
 tu-berlin.de: type A, class IN, addr 10.150.7.69
 tu-berlin.de: type A, class IN, addr 172.31.25.70
 tu-berlin.de: type A, class IN, addr 10.150.7.68
 tu-berlin.de: type A, class IN, addr 10.150.7.67
 tu-berlin.de: type A, class IN, addr 10.150.7.70
 [Request In: 14]
 [Time: 0.020163079 seconds]

Impersonating DNS Requests with VNRF (cont'd)

No.	Time	Source	Destination	Protocol	Length	Info
13	3.538438...	8.8.8.8	192.168.217.1...	QUIC	13...	Initial, SCID=000000000000109
14	3.538771...	192.168.217.1...	8.8.8.8	QUIC	200	Version Negotiation, DCID=000000000000109
15	3.558935...	8.8.8.8	192.168.217.1...	QUIC	152	53 → 12345 Len=110[Malformed Packet]

▶ Frame 14: 200 bytes on wire (1600 bits), 200 bytes captured (1600 bits) on interface ens33, id 0

- ▶ Ethernet II, Src: VMware_5e:6a:92 (00:0c:29:5e:6a:92), Dst: VMware_f6:95:1c (00:50:56:f6:95:1c)
- ▶ Internet Protocol Version 4, Src: 192.168.217.131, Dst: 8.8.8.8
- ▶ User Datagram Protocol, Src Port: 12345, Dst Port: 53
- ▶ QUIC IETF
 - ▶ QUIC Connection information
 - [Packet Length: 158]
 - 1... = Header Form: Long Header (1)
 - .100 1001 = Unused: 0x49
 - Version: Version Negotiation (0x00000000)
 - Destination Connection ID Length: 7
 - Destination Connection ID: 000000000000109
 - Source Connection ID Length: 116
 - Source Connection ID: 752d6265726c696e026465000001000100000100010000010001000001000100000100010000010001...
 - Supported Version: v2-draft-01 (0x709a50c4)
 - Supported Version: 1 (0x00000001)
 - Supported Version: draft-32 (0xff000020)
 - Supported Version: draft-31 (0xff00001f)
 - Supported Version: draft-30 (0xff00001e)
 - Supported Version: draft-29 (0xff00001d)
 - Supported Version: Unknown (0x4a0ababa) (GREASE)

Frame (frame), 200 bytes Packets: 31 · Displayed: 3 (9.7%) Profile: Default
[Time: 0.020163079 seconds]

Impersonating DNS Requests with VNRF (cont'd)

No.	Time	Source	Destination	Protocol	Length	Info
13	3.538438...	8.8.8.8	192.168.217.1...	DNS	13...	DNS Stateful operations (DS0) 0xc813[Malformed Packet]
14	3.538771...	192.168.217.1...	8.8.8.8	DNS	200	Standard query 0xc900 A tu-berlin.de A <Root> A <Root> A <Root> A
15	3.558935...	8.8.8.8	192.168.217.1...	DNS	152	Standard query response 0xc900 A tu-berlin.de A 10.150.7.69 A 172.

Frame 14: 200 bytes on wire (1600 bits), 200 bytes captured (1600 bits) on interface ens33, id 0

- Ethernet II, Src: VMware_5e:6a:92 (00:0c:29:5e:6a:92), Dst: VMware_f6:95:1c (00:50:56:f6:95:1c)
- Internet Protocol Version 4, Src: 192.168.217.131, Dst: 8.8.8.8
- User Datagram Protocol, Src Port: 12345, Dst Port: 53
- Domain Name System (query)
 - Transaction ID: 0xc900
 - Flags: 0x0000 Standard query
 - Questions: 7
 - Answer RRs: 0
 - Authority RRs: 0
 - Additional RRs: 1
 - Queries
 - tu-berlin.de: type A, class IN
 - <Root>: type A, class IN
 - <Root>: type A, class IN
 - <Root>: type A, class IN
 - <Root>: type A, class IN
 - <Root>: type A, class IN
 - <Root>: type A, class IN
 - Additional records
 - <Root>: type Unused, class Unknown

[Response In: 15]

Frame (frame), 200 bytes

Packets: 31 · Displayed: 3 (9.7%) · Dropped: 0 (0.0%) Profile: Default

[Time: 0.020163079 seconds]

Mitigation

CID Reflection

- A server always chooses a fresh SCID, also for version negotiation

Hashing

- A „seed“ for a CID still chosen by the client
- The server uses a hash of the seed as DCID
- An attacker would need to calculate the inverse to create a meaningful payload


Masking

- QUIC headers get an additional field containing a masking value
- The masking value is randomly generated by the server
- The entire remaining header is XORed
- Client maintains control over DCID but payloads will appear „random“

Traffic Amplification

Path amplification VS Bandwidth Amplification

$$PAF = \frac{\text{\# Bytes from victim to target}}{\text{\# Bytes from attacker to victim with spoofed address}}$$



$$BAF = \frac{\text{\# Bytes from victim to target}}{\text{\# Bytes from attacker to victim}}$$

Amplification Pitfalls

“[...] not send more than three times the amount of data received on any unvalidated path.”

Minimum path requirements

- „QUIC must not be used if the network path cannot support 1200 bytes datagrams“
- Ensured through padding of initial packets and path challenges
- Small packets on new paths are an issue
- ***Server should send two separate path validations***

Unbalanced handshake sizes

- Server initial packets are larger than client initial packets
- ***Server initial packets should never be larger than 3*1200 bytes***

Amplification Pitfalls

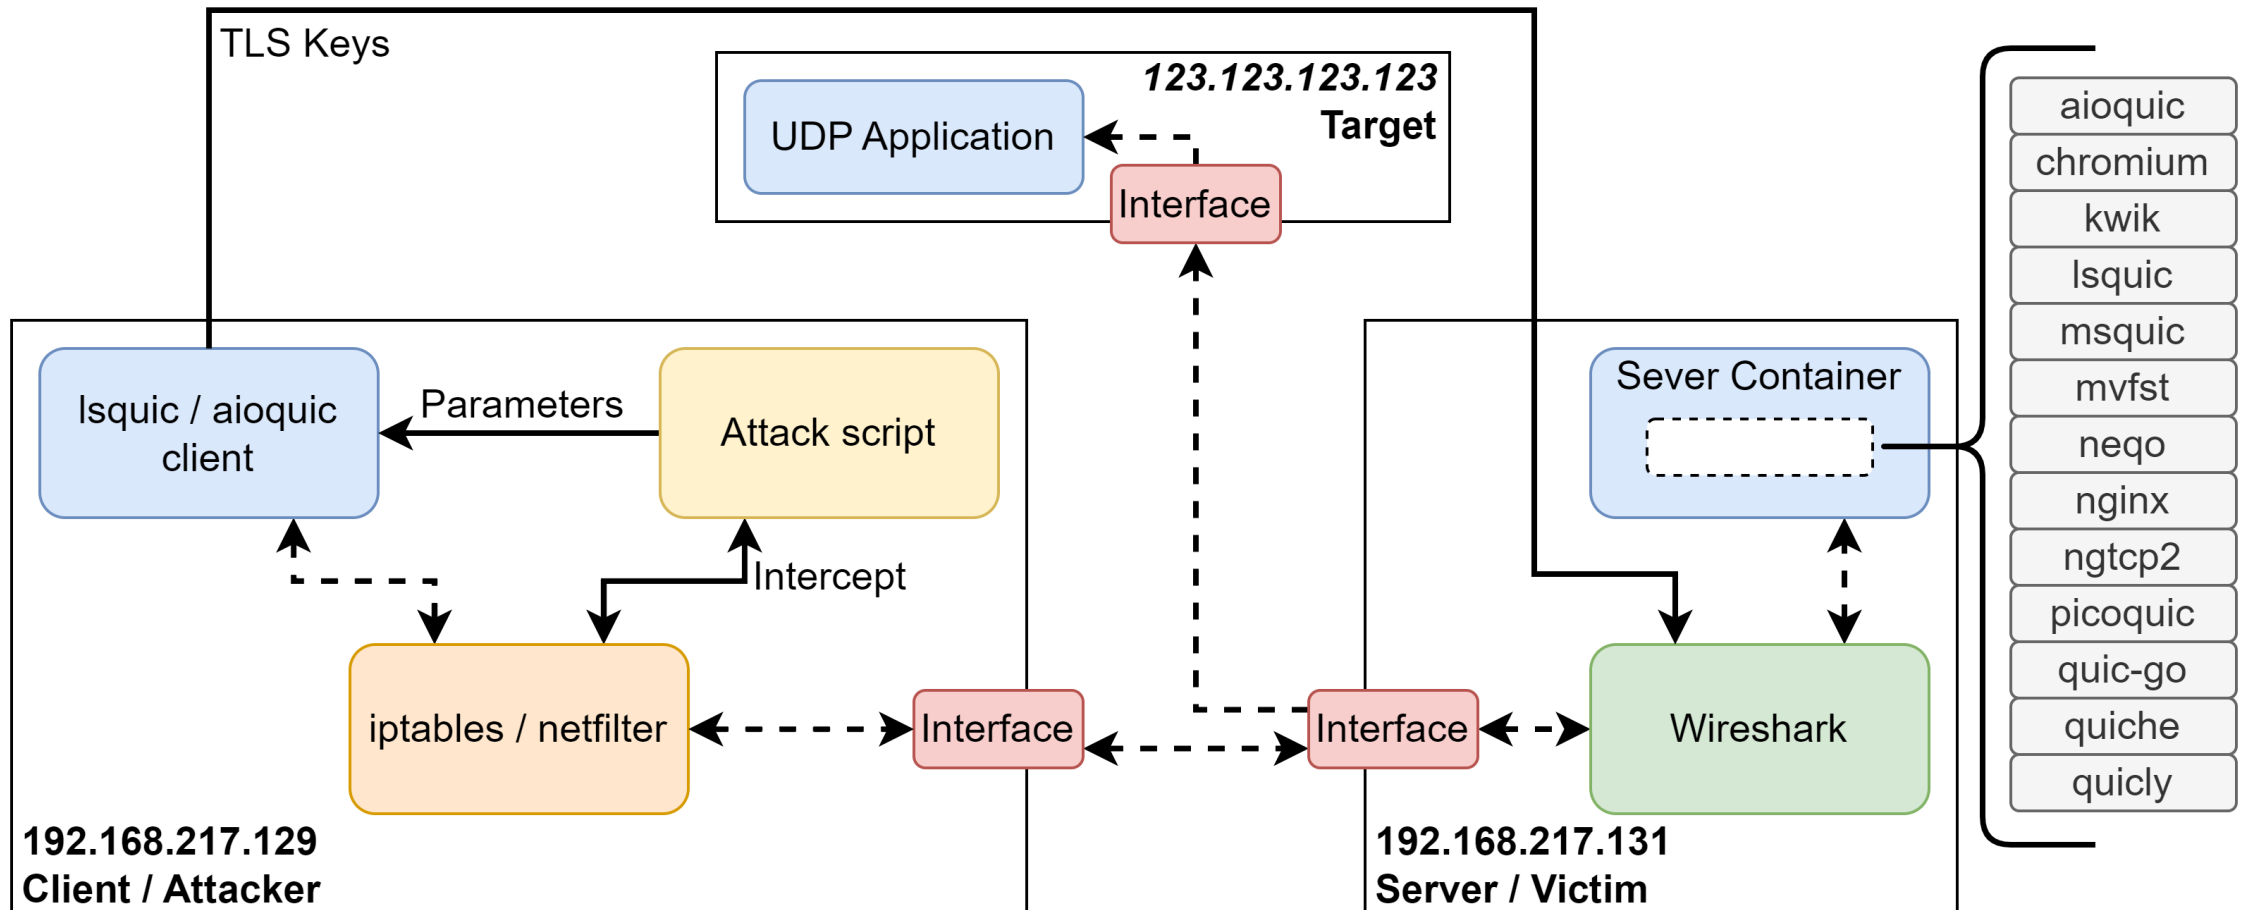
Reliability

- No „typical“ reliability in path challenges, server can send multiple challenges
 - Initial bursts
 - Re-send with timeout
- ***Multiple challenges definitely surpass amplification limits***
- Normal packets are re-sent if the acknowledgment is not received
- ***Server should not re-send server initial packets***
 - Retries for the initial messages have to be handled by the client.

RTFM

Evaluation


Evaluation Setup



Testing Implementations

	CMRF				SIRF			VNRF	
	Vuln.	Pad.	New CID	PAF > 3	Vuln.	PAF > 3	Ref. CID	Vuln.	CID>20
aioquic	✓	✗	✗	✗	✓	✓	✗	✓	✗
chromium	✓	✗	✗	✓	✓	✓	✗	✓	✓
kwik	✗	-	-	-	✓	✗	✗	✓	✓
lsquic	✓	✓	✓	✓	✓	✗	✗	✓	✗
msquic	✓	✗	✓	✗	✓	✓	✗	✓	✓
mvfst	✓	✓	✗	✓	✓	✓	✗	✓	✗
neqo	✓	✗	✓	✓	✓	✓	✗	✓	✓
nginx	✓	✗	✓	✓	✓	✗	✗	✓	✗
ngtcp2	✓	✗	✗	✗	✓	✗	✗	✓	✓
picoquic	✓	✗	✓	✓	✓	✗	✗	✓	✓
quic-go	✗	-	-	-	(✓)	✗	✗	✓	✓
quiche	✗	-	-	-	(✓)	✗	✗	✓	✓
quicly	✗	-	-	-	✓	✓	✗	✓	✓
Total	9	2	5	6	11(13)	6	0	13	9

Amplification



Conclusion

Conclusion

- **All 13 implementations were „vulnerable“ to at least one technique**
- Amplification protection is covered in the specification
- **Yet, significant PAF values uncovered**
 - 374.44 for CMRF
 - 22.1 for SIRF
- Vendors were notified and implemented
- **Protocol impersonation is possible with VNRF**
- Currently no built-in protection mechanism
 - changes to the specification should be considered

Thanks!

Blogpost with additional technical details:

<https://r.sec-consult.com/quic>



Thanks for listening!