CLOUDFLARE

RFC3849 and RFC5737

Documentation! Documentation! Documentation!

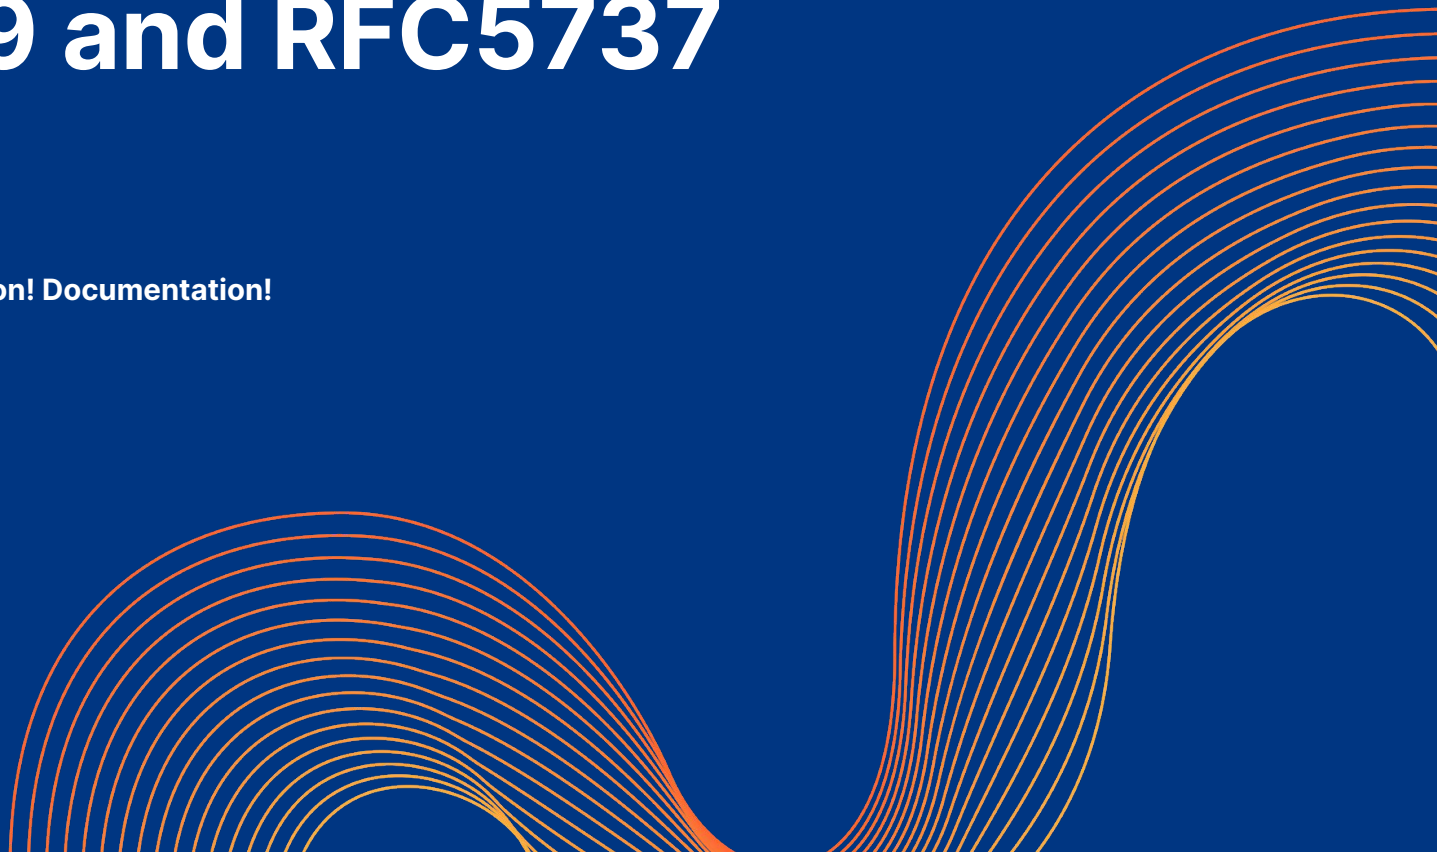

CLOUDFLARE

Documentation matters

CLOUDFLARE

Why?

CLOUDFLARE

Flipperty Flopperty

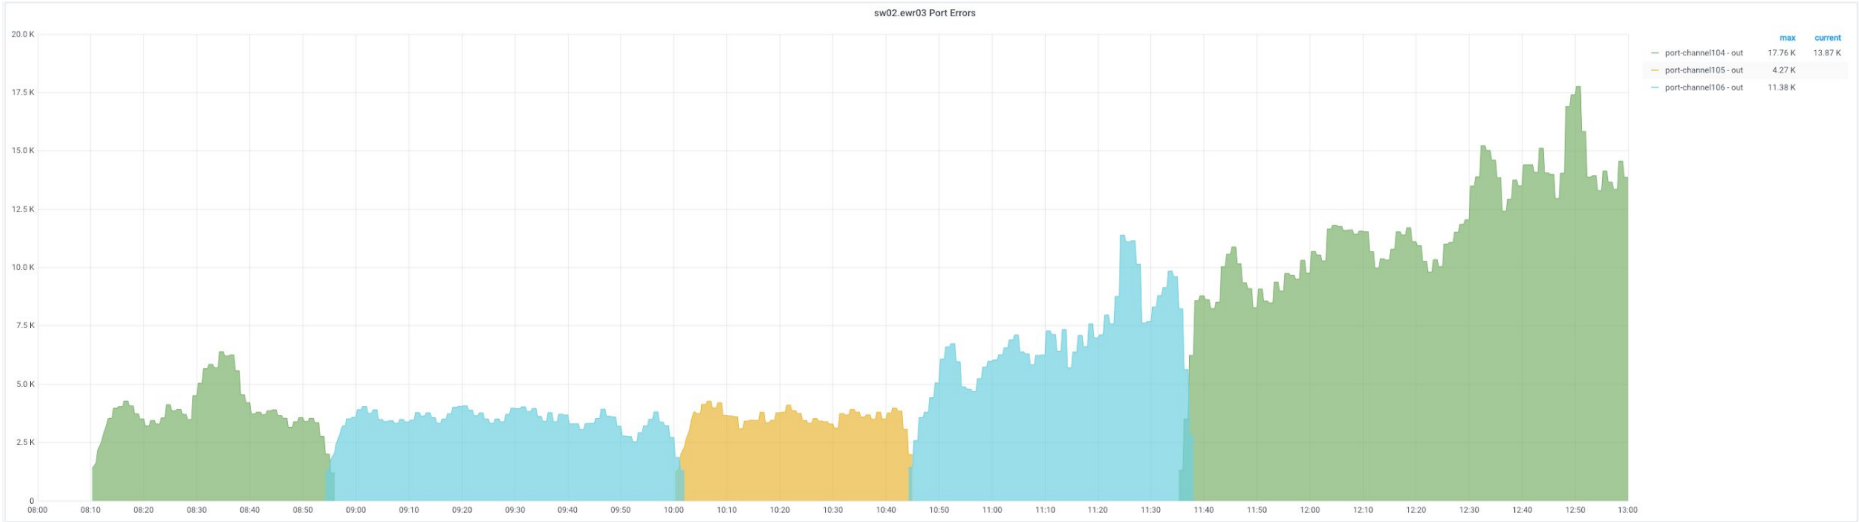

CLOUDFLARE

Nothing is straightforward

[redacted] added a comment - 18/Sep/19 10:15 AM - edited

so I think what's happening is that the Nexus receive a broken frame, because of cut-through the frames is forwarded to the dest [redacted] with a bad CRC on purpose (stomping on the Nexus)

```
[redacted]:~$ /sbin/ifconfig ext0 | grep err
RX errors 1746054969 dropped 195799 overruns 0 frame 1746054969
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
[redacted]:~$ /sbin/ifconfig ext0 | grep err
RX errors 1746064939 dropped 195799 overruns 0 frame 1746064939
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
[redacted]:~$ /sbin/ifconfig ext0 | grep err
RX errors 1746076243 dropped 195799 overruns 0 frame 1746076243
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

BUT for some reason the Nexus won't tell us where the broken frames are coming from (no CRC counter increasing)

so imho the issue here is a display issue more than a switch behavior

I'm pretty sure the problematic traffic comes from the same ToR since we would see CRC on the core01.ewr03 as well.

CLOUDFLARE

TCPDump!

CLOUDFLARE

TCPDump! (contd.)

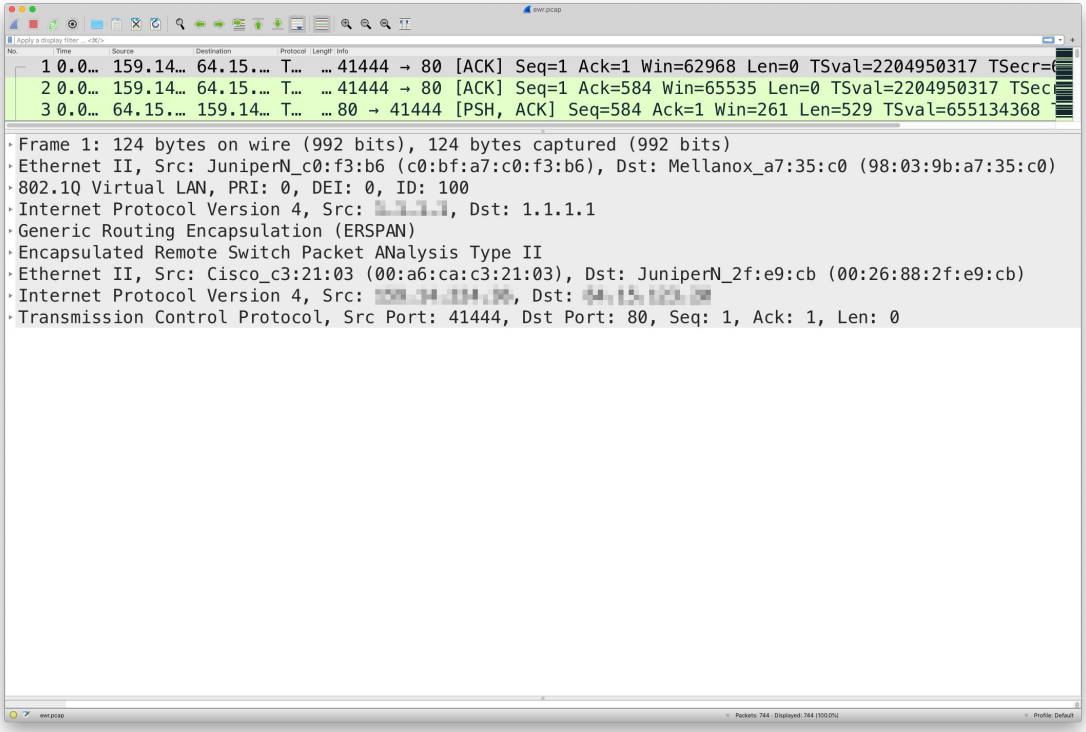

CLOUDFLARE

Combined problems == hard problems

[redacted] admin

> **Tom Strickx**
> Yup, but odd that the QFXs didn't spot it

nexus has a long feature of reading deeper into packets than it should. had similar issues in [redacted] with the N7K

CLOUDFLARE

Documentation?

CLOUDFLARE

And the penny drops

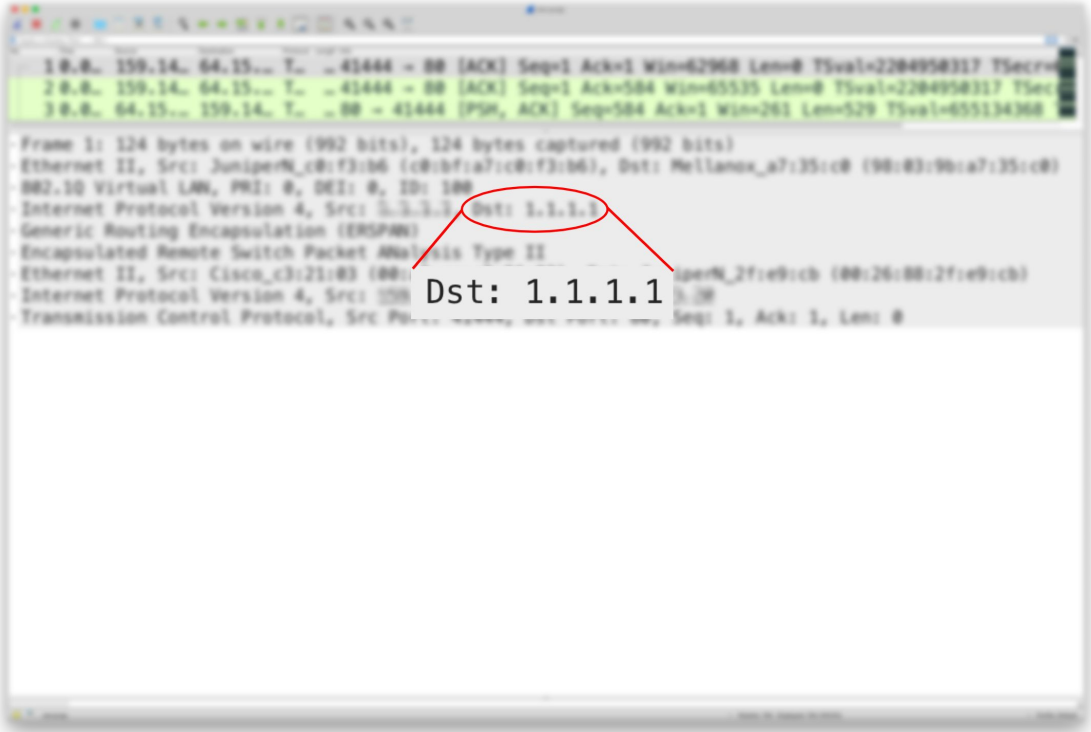

CLOUDFLARE

Documentation Matters!

https://www.cisco.com/c/en/us/products/collateral/switches/nexus-5000-series-switches/white-paper-c11-733110.pdf

You next need to define the source for the ERSPAN session: an interface, VLAN, or VSAN from which traffic will be mirrored. By default, traffic is monitored in both directions if the source is a port or PortChannel. To monitor traffic in only one direction, you must specify the transmit (tx) or receive (rx) parameter. To define a port and a PortChannel as the source of the ERSPAN session, use these commands:

```
switch(config-erspan-src)# source interface ethernet 1/1 [tx|rx|both]
switch(config-erspan-src)# source interface port-channel 101 [tx|rx|both]
```

To define a VLAN and a VSAN as the source of the ERSPAN session, use these commands:

```
switch(config-erspan-src)# source vlan 1
switch(config-erspan-src)# source vsan 1
```

The ERSPAN source session needs a defined destination IP address to which to forward traffic. The ERSPAN session can have only one destination IP address. Use this command to configure the destination IP address:

```
switch(config-erspan-src)# destination ip 1.1.1.1
```

Define the ERSPAN ID for the ERSPAN flow with this command:

```
switch(config-erspan-src)# erspan-id 1
```

To define a specific VRF to use instead of the global VRF, use either of these commands:

```
switch(config-erspan-src)# vrf default
```

or

```
switch(config-erspan-src)# vrf vrf-name
```

You can also define optional parameters: access control list (ACL), TTL, maximum transmission unit (MTU) for truncated ERSPAN, and differentiated services code point (DSCP) values.

In the source session, you can configure ACLs to filter packets in the ERSPAN session:

```
switch(config-erspan-src)# filter access-group erspan_acl_filter
```

To limit the number of jumps from the source to the destination and stop traffic from going to an unvented device, you can define the IP TTL. To configure the IP TTL for packets in ERSPAN traffic, use this command:

```
switch(config-erspan-src)# ip ttl 1
```

To prioritize ERSPAN traffic over some less important production traffic, you can change the DSCP value. By default, ERSPAN traffic is assigned a DSCP value of 0. To define a different DSCP value for packets in ERSPAN traffic, use this command:

```
switch(config-erspan-src)# ip dscp 1
```

If ERSPAN traffic is oversubscribing a network link, you can reduce the load by defining an MTU value. The MTU value for truncated ERSPAN packets can be between 64 and 1518 bytes. Use this command:

```
switch(config-erspan-src)# mtu 64
```

CLOUDFLARE

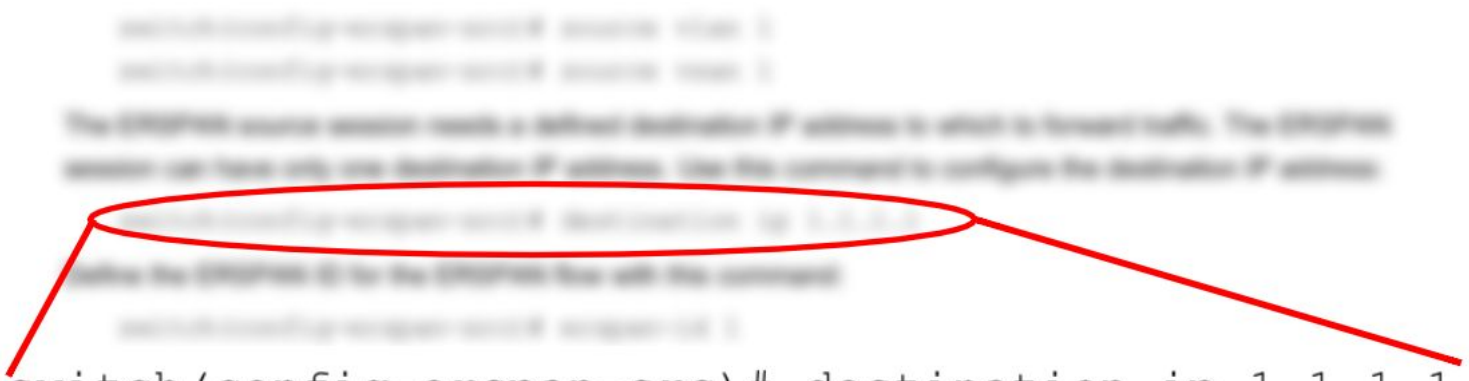

CLOUDFLARE

Welp

https://blog.cloudflare.com/fixing-reachability-to-1-1-1-1-globally/

Announcing 1.1.1.1: the fastest, privacy-first consumer DNS service

01/04/2018

Matthew Prince

CLOUDFLARE

Awkward

https://twitter.com/Cloudflare/status/1316062417416327169

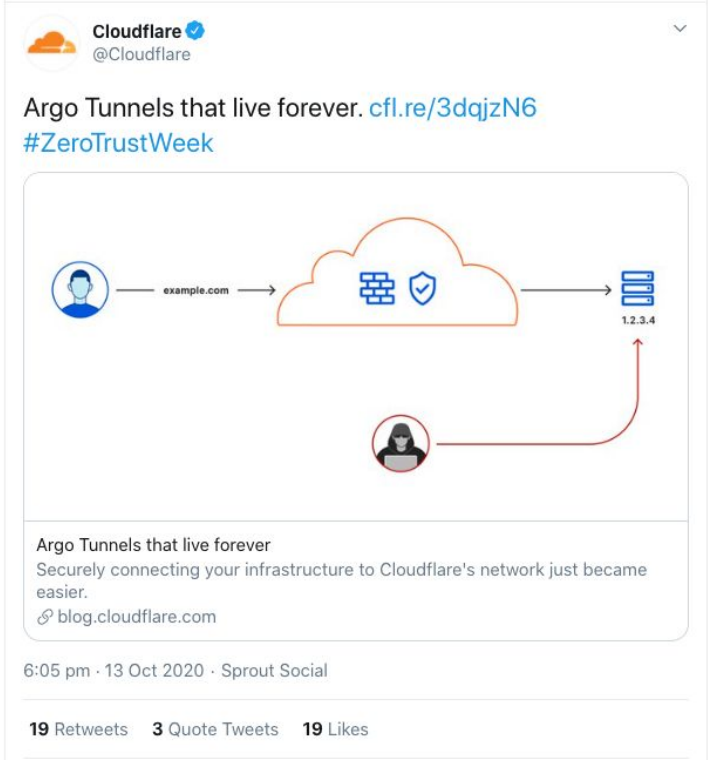

CLOUDFLARE

Awkward

https://twitter.com/Cloudflare/status/1316062417416327169

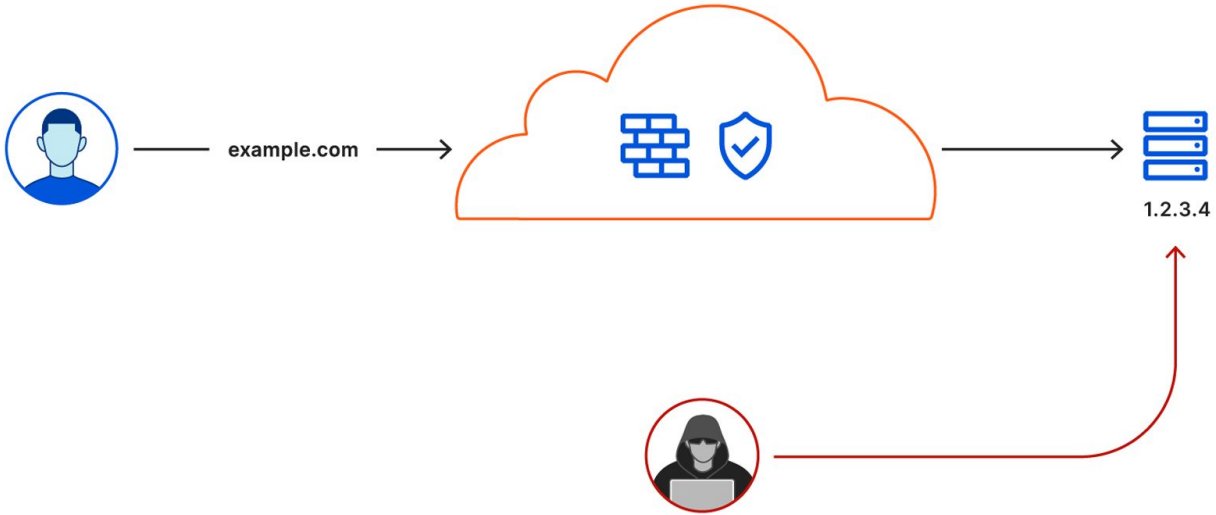

CLOUDFLARE

Awkward

https://twitter.com/Cloudflare/status/1316062417416327169

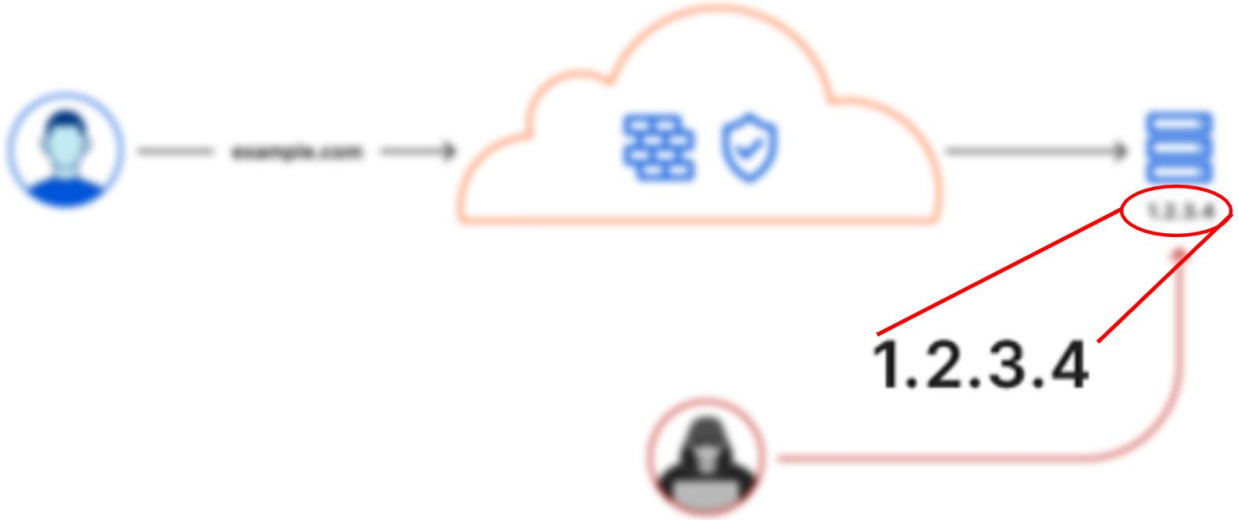

CLOUDFLARE

RFC3849 and RFC5737

July 2004

IPv6 Address Prefix Reserved for Documentation

January 2010

IPv4 Address Blocks Reserved for Documentation

CLOUDFLARE

For those in the back

RFC3849

and

RFC5737

CLOUDFLARE

Summary - RFC3849

2001:DB8::/32

CLOUDFLARE

Summary - RFC5737

192.0.2.0/24

198.51.100.0/24

203.0.113.0/24

CLOUDFLARE

Summary

CLOUDFLARE

Thank you!

@tstrickx

tstrickx@cloudflare.com